

Curbing Habit Formation

The Effects of Tobacco Control Policies in a Dynamic Equilibrium *

Francisco Pareschi¹ and Gaston Lopez²

¹Northwestern University

May 1, 2024

LATEST VERSION

Abstract

We study the equilibrium effects of tobacco control policies. To accurately assess the impact of tobacco regulation on consumption, it is essential to anticipate firm responses. We highlight that consumers' dependence on cigarettes, which we refer to as *consumer inertia*, introduces dynamic incentives for firms. Thus, we develop a dynamic oligopoly model and estimate it using product-level data and a panel of smokers. Leveraging large tax fluctuations and a policy that forced approximately 40% of products out of the market, we show that smokers have significant addiction and brand loyalty. We then propose a tractable equilibrium notion to compute market outcomes and show that dynamic competition models under inertia capture firm behavior well. We use this framework to examine the counterfactual effect of caps on nicotine concentration and uniform packaging, which are expected to decrease smokers' addictiveness and brand loyalty, respectively. We highlight that firms' price and portfolio responses account for a substantial part of the policies' impact. Moreover, more simplistic models of static competition would have led to significantly different conclusions about expected companies' responses.

*First Version: October 2023. A previous version of this paper circulated as "Reduction Consumer Inertia in Tobacco Markets". We are grateful to David Besanko, Vivek Bhattacharya, Igal Hendel, Gastón Illanes, Rob Porter, Mar Reguant, and Bill Rogerson for many detailed comments and discussions. We also thank Felipe Berruti, Nick Buchholz, Sebastian Fleitas, Eilidh Geddes, Alexander MacKay, Matt O'Keefe, Anran Li, Jose Salas, Molly Schnell, David Stillerman, Patricia Triunfo, Maren Vairo, Benjamin Vatter, Jingyuan Wang, Tomas Wilner, Hans Zhu, and seminar participants at the IO Graduate Students Seminar at Northwestern for helpful comments. We thank Soledad Fernandez, Gonzalo Chiappara, and the Scanntech team for generously sharing the data.

1 Introduction

Tobacco kills 8 million people every year around the world and more than 480,000 in the US alone. Although governments have discouraged its consumption through taxation and regulation, the industry remains resilient. Authorities are now considering a new generation of innovative policies to combat the tobacco epidemic. In 2022, for instance, the Food and Drug Administration (FDA) proposed a plan to develop a product standard that would establish a maximum nicotine level to reduce the addictiveness of cigarettes [FDA, 2022]. In addition, several countries have started implementing uniform packaging to decrease the appeal of cigarettes [WHO, 2022].¹ Whereas this type of regulation aims to affect consumer behavior, understanding how tobacco companies will respond is crucial for anticipating the overall impact of these policies on consumption. Echoing this concern, a UK government review warned that uniform packaging could “reduce brand loyalty, causing smokers to switch to cheaper brands and encouraging price competition between manufacturers” [Chantler, 2014, pp. 5].

This paper studies the equilibrium effect of tobacco control policies. In particular, we study whether firms’ responses amplify or undo the direct impact of regulation on consumers. We stress that smokers’ dependence on cigarettes makes the problem of the firm inherently dynamic. Smokers have two well-known sources of habit formation. They become *addicted* to tobacco due to nicotine intake and develop persistent *brand loyalty* to the products they smoke.² Under these forms of habit formation, future consumption becomes a function of current purchases. Thus, firms have incentives to consider the long-term implications of their decisions. The economic literature has highlighted that such incentives modify how firms price and offer products. For instance, Klemperer [1987a] found that consumer inertia encourages firms to lower prices to attract a larger customer base and then raise them to profit from the locked-in consumers. Also, the sacrifice of profits required to lure customers to a new product can deter their introduction. Since Bain [1956]’s seminal work, economists have identified consumer inertia as a major barrier to entry. Therefore, a careful analysis of tobacco regulation must account for and evaluate firms’ dynamic incentives.

To empirically assess the equilibrium effects of regulation in tobacco markets, we develop and estimate a dynamic model of competition that accounts for consumers’ and firms’ responses. Consumers exhibit addiction and brand loyalty, while forward-looking firms choose prices and product portfolios. We apply this model to the Uruguayan cigarette industry and provide robust empirical

¹Australia was the first country to pass plain packaging legislation in 2012. Since then, France (2017), United Kingdom (2017), New Zealand (2018), Norway (2018), Ireland (2018), Hungary (2019), Thailand (2019), Uruguay (2019), Saudi Arabia (2020), Slovenia (2020), Turkey (2020), Belgium (2021), Canada (2022), Singapore (2020), Israel (2020), Netherlands (2021), and Denmark (2022) have enacted some form of plain packaging policy.

²Product loyalty is significantly higher in the tobacco industry than for other consumption goods. US consumers make 60% of cigarette purchases on their focal product [Dawes, 2014], while the average across products is approximately (29-33)% (Uncles et al., 1994; Bhattacharya, 1997).

evidence that a dynamic competition model is necessary to capture firm behavior realistically. Finally, we simulate the equilibrium effect of the FDA’s nicotine caps and uniform packaging, two policies expected to decrease the degree of habit formation in the industry. We find that curtailing consumers’ habit formation is an effective tool to curb tobacco consumption and that firms’ responses account for a substantial part of the overall effect. Our results suggest that these policies discourage firms from capturing new consumers by restricting their ability to retain them in the future, which tends to reinforce the direct effect. More simplistic models of static competition would have missed these dynamic considerations and led to different conclusions since demand also becomes substantially more elastic.

Connecting our model to the data involves not only identifying its primitives but also overcoming significant computational constraints. To identify its primitives, we combine product-level data and a panel of smokers with rich variation from the Uruguayan experience. We leverage two primary sources of variation to identify consumers’ addiction and brand loyalty. First, there are notable tax oscillations.³ Tax-driven price swings allow us to observe asymmetric responses to prices of consumers with and without dependence, which informs the extent of addiction. Second, a regulation forbade firms from offering multiple products under the same brand name and forced approximately 40% of products out of the market. The choices of customers who “lost” the product to which they were loyal identify the preferences of consumers who are not used to any particular product but still face addiction. To alleviate the computational constraints, we propose a tractable equilibrium concept that restricts firms’ informational requirements. While such restrictions make the equilibrium easy to compute, they preserve several key strategic incentives that arise in the Markov Perfect Equilibrium (MPE). The computational simplicity of the equilibrium enables the application of the method of simulated moments (MSM) to estimate costs.⁴

Estimates suggest a high degree of consumer inertia. Current smokers are willing to pay nearly two times the observed average price for any cigarette and more than three times the average price to repeat their product choice. In this regard, the mean own-price elasticity is around -0.9, which implies that firms price in the inelastic region of the demand curve. These estimates are low compared with other industries but consistent with the scarce literature that treats cigarettes as differentiated products [Ciliberto and Kuminoff, 2010, Liu et al., 2015, Tuchman, 2019]. The

³These fluctuations arise from (1) governmental priority shifts regarding tobacco control and (2) setting specific taxes at nominal values.

⁴Although our equilibrium notion renders MSM computationally feasible, we must overcome a few hurdles. First, we discuss how to address potential equilibrium multiplicity using the absorbing steady state of the game without product assortment. We then observe that prices depend on fixed costs through the next period’s portfolio probabilities. Thus, we cannot split the problem into two steps by recovering marginal costs from static first-order price conditions and then solving the entry/exit dynamic game to estimate fixed costs. However, participation choices and prices define distinct combinations of marginal cost and continuation probabilities that could rationalize them. Therefore, conduct still aids identification.

estimated production costs are small, implying that taxes represent more than 90% of firms' total marginal costs. Under these conditions, firms' dynamic and strategic incentives play a critical role in determining industry dynamics. For instance, we find that firms avoid lowering prices too much to prevent stealing consumers from rivals and sparking fiercer competition in subsequent periods. Our estimates suggest this form of business-stealing is comparable to the business-stealing effect within the firm portfolio.

Next, we provide evidence that our theoretical model of competition under inertia captures firm behavior realistically. The model explains two market features that would be hard to capture in a static competition model. First, it explains why firms set low markups despite highly inelastic demand. Indeed, the fact that firms are pricing in the inelastic region of the demand curve already suggests that a static model cannot accurately depict firm behavior. In our dynamic competition model, firms internalize consumers' contribution to their *long-term* profits. Therefore, firms are willing to sacrifice short-term profits to retain and attract new customers. As a result, they set lower markups than those implied by a static model for any elasticity level. Second, under the estimated levels of habit formation, the model generates significant price discounts when introducing a new product. We observe that predicted and observed introductory pricing strategies—which occasionally implied setting prices at cost—are similar and not caused by cost changes or consumer preferences.

Finally, we use our framework to simulate the counterfactual effect of nicotine caps and uniform packaging. We find that supply responses are relevant to understanding the impact of both policies. The common feature of these policies is that they lower consumers' habit formation, addictiveness, and brand loyalty respectively. Although lowering inertia increases demand elasticity substantially (up to three times in some policy scenarios), it also reduces the long-term value customers have for the firm. Thus, firms lower their investments in bringing new customers to the market, constraining price decreases and reducing consumption. We show this effect is equivalent to raising firms' costs. Nonetheless, declines in consumer inertia can facilitate the introduction of new products. If the policy shifts demand towards less attractive products without directly affecting the aggregate market, as with uniform packaging, disadvantaged products can generate more profits. Moreover, endogenous entry cost decreases because stealing consumers from established products becomes more accessible. As a result, product availability and consumption can increase.

Our results demonstrate that limiting firms' ability to retain customers in the future—which can take many forms, such as reducing addiction, inducing more competition, or adding regulatory uncertainty—is a valuable tool to limit consumption. Unlike taxation, it can reduce consumption without increasing the burden on consumers, which makes it especially attractive for regulators.

Related Literature We contribute to the policy literature on tobacco control, understanding industry dynamics under consumer inertia, and the methodological body of work on empirical dynamic oligopolies. From the policy point of view, while many studies have investigated the effect of multiple tobacco regulations, ours is one of the few studies accounting for firm responses and industry dynamics.⁵ Different from previous work, we compute tobacco firms' equilibrium price and portfolio strategies within an empirical dynamic oligopoly. In that sense, our work also relates to Barahona et al. [2020], which accounts for firms' responses to evaluate the equilibrium effect of food labeling policies, and Abi-Rafeh et al. [2023], which analyzes firms' dynamic incentives to study the impact of sin taxes and advertisement restrictions in the sugar-sweetened beverage industry. Our results support using non-price policies to reduce sin-goods consumption, as they can achieve their goals without shifting the burden to consumers, which can have negative distributional consequences, as recently illustrated by Conlon et al. [2022].

Our paper also advances the understanding of industry dynamics under consumer inertia. We build on the modern research on dynamic price competition in this context [Dubé et al., 2009, Arie and E. Grieco, 2014], and introduce entry and exit considerations following the framework laid out by Benkard [2004], Farrell and Katz [2005], Besanko et al. [2014, 2019] to study games of dynamic competition under learning-by-doing, and network externalities. Our results capture Bain [1956]'s intuition of brand loyalty as a barrier to entry and are in the same spirit as Fleitas [2017]'s findings.⁶ Similar to previous studies, we find a non-monotonic relationship between prices and brand loyalty.⁷ Our work suggests this pattern remains even after introducing entry and exit decisions. Finally, our estimates indicate that firms do not engage in entry-deterrence or exit-inducing behavior despite having rational incentives to do so, as initially noted by Klemperer [1987b] and recently incorporated into a legal theory of predation [Fumagalli and Motta, 2013].

We also relate to many empirical papers taking models of consumer inertia to the data. Our central contribution to this empirical literature is to provide direct evidence that firms are forward-looking and behave close to theory predictions. Our approach is reminiscent of Benkard [2004], which estimates all primitives of the model without ever solving the equilibrium and then compares equilibrium outcomes with data. Although we use the equilibrium to estimate firms' costs, we do so in a way that does not fully rationalize the data, letting us test the model's predictive power. Re-

⁵See [Levy et al., 2019] for an in-depth discussion about tobacco control from an economic and marketing perspective. A few exceptions are Ciliberto and Kuminoff [2010], evaluating the effect of the 1997 Master Settlement Agreement (MSA) on firms' ability to collude, and Qi [2013]'s study about industry dynamics following the 1971 cigarette advertising ban in the United States.

⁶These results differ from the few papers exploring the relationship between inertia and participation choices in simple theoretical frameworks [Farrell and Shapiro, 1988, Beggs and Klemperer, 1992, Gabszewicz et al., 1992], suggesting that higher inertia would facilitate entry due to increased industry profits.

⁷This relation appears to be a robust feature of competition under inertia and aligns with most of the literature, both when the seller cannot discriminate between consumers [Dubé et al., 2009, Arie and E. Grieco, 2014, Fabra and García, 2015] and when they do [Cabral, 2016].

garding our identification and estimation strategy, price oscillations have been previously used to identify switching costs [Pakes et al., 2021] and brand loyalty [Dubé et al., 2010], while transitory shocks to the choice set have also been used to identify switching cost and inattention [Handel, 2013]. We also relate to a smaller literature estimating costs in dynamic models of price competition. The recent empirical work on price competition under inertia generally takes firms' costs as given [Dubé et al., 2009, MacKay and Remer, 2021] or uses solution-free approaches to estimate them [Fleitas, 2017, Pavlidis and Ellickson, 2017].

Finally, we contribute to expanding the empirical tools available to analyze dynamic oligopolies. Recent work has stressed the limitations of MPE for empirical work. In response, Weintraub et al. [2008], Fershtman and Pakes [2012], Benkard et al. [2015] and Ifrach and Weintraub [2017], among others, have proposed relaxations of the MPE, which allow for tractable computation of equilibria. A common feature of these approaches is that they are applied to games that restrict firms' dynamic controls to affect their own states only. Hence, assuming rivals' states are fixed at some stationary level as is the case of Weintraub et al. [2008] or evolving according to a quasi-exogenous process as in Ifrach and Weintraub [2017], does not eliminate many of the relevant strategic incentives present in the MPE of these games. However, in several dynamic games of price competition, such as competition under inertia, network externalities, and learning-by-doing, firms' dynamic strategies (prices) affect all player's states' transitions (demand). This interaction, in turn, gives place to rich strategic incentives and industry dynamics—see, for instance, Farrell and Katz [2005], Besanko et al. [2010, 2014]. Our approach is innovative in that it preserves these strategic interactions while still being tractable. To achieve this tradeoff, we must limit firms' strategic incentives to be homogeneous across rivals or groups of rivals.

2 Industry Background

Next, we describe the Uruguayan tobacco market. We observe that consumer choices are persistent. Moreover, we argue that firm behavior is consistent with a model of competition under habit formation, a point we revisit in Section 5.3.

2.1 Data

We use two primary data sources. First, we use store scanner data, which provides information on the quantity and price paid for cigarettes sold between 2006 and 2019. The final sample includes around 100 stores scattered across 40 regions. Aggregate sales in our sample closely track the aggregate national sales according to the Uruguayan tax reports. The market structure is simple.

Three players participate in the Uruguayan cigarettes market.⁸ Monte Paz, a national firm, holds around 75% of the market, while the two multinationals, Philip Morris and British American Tobacco (BAT), account for 20% and 5%, respectively. There are between 20 and 30 products in the market. However, many of these products share prices, observable characteristics, and are introduced and retired simultaneously. In most of the analysis, we bundle similar products together and refer to them as products or segments. We distinguish between each firm flagship products, other regular products, the light category (low in tar), and other products with special characteristics (slim, longer, etc.). In total, we work with nine product segments.

We also leverage an individual panel built by the International Tobacco Control Policy Evaluation Project (ITC). The ITC contacted and interviewed individuals every other year from 2006 to 2014. Initially, the panel included only smokers, but if they decided to quit between interviews, they stayed in the sample. This panel contains information about individuals' smoking status in each wave, quantity smoked, brand, the price paid, age of initiation, time smoking the current brand, demographics, etc. The final sample includes around 1,300 individuals and almost 3,000 choice events. Additionally, we use other relevant information, such as the population survey (Encuesta Continua de Hogares), to obtain demographics and smoking prevalence at the regional level. Appendix A summarizes this information and presents the details of the product aggregation.

2.2 Addiction and Brand Loyalty

Table 1 shows the proportion of consumers who repeat choices across waves of the individual panel. First, our sample's overall quitting rate is between 15% and 25%, and the smoking rate for people not smoking in the previous wave is around 21%. Second, on average, 70% of smokers repeat their product choice in the next wave. These figures show that consumer choices are highly persistent. However, we do not know, *a priori*, whether it is due to persistent consumer preferences or structural state dependence [Heckman, 1981]. Indeed, this is the crucial identification challenge in our analysis, to which we come back in Section 4.

2.3 Investing and Harvesting

In 2004, Uruguay ratified the Framework Convention on Tobacco Control of the WHO, implementing a wide array of tax and non-price policies to regulate the tobacco industry. These policies included prohibiting all advertising of tobacco products, smoking in enclosed public places, and

⁸Observe that we are excluding three potentially relevant substitute products. First, we exclude roll-your-own tobacco products. Second, we are also omitting the smuggled cigarette products. Finally, note that our sample period coincides with the rise of e-cigarette sales in the US [Tuchman, 2019]. However, Uruguay forbade the commercialization of e-cigarettes. Although individuals can buy them abroad and bring them as personal objects, according to tobacco use surveys, less than 10% of individuals have ever used them between 2006 and 2012. Hence, we do not expect e-cigarettes to impact our results significantly.

Table 1: Switching matrix for smokers and non-smokers

Smoker Status	Year	Repeat Choice	Smoke
Non Smoker	2010		0.180
	2012		0.226
	2014		0.244
Smoker	2008	0.676	0.823
	2010	0.592	0.773
	2012	0.679	0.807
	2014	0.760	0.867

Note: Probabilities are computed between waves of the ITC survey, which usually span over two years.

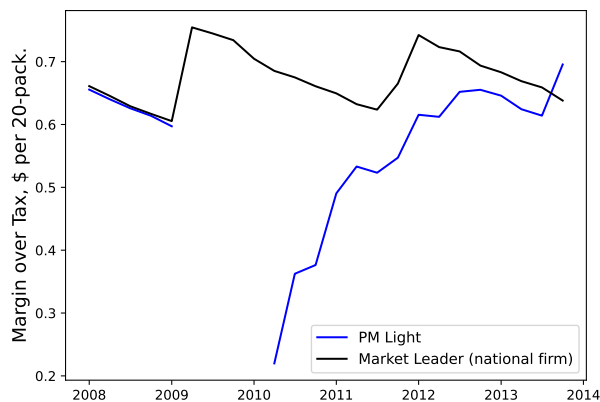
imposing sizeable pictorial health warnings on tobacco products’ packaging. In 2009, the government passed the “one-presentation-per-brand” regulation, which required producers to use a different brand name for each product (see DeAtley et al. [2018] for a detailed analysis of the policy compliance). Before 2009, all firms structured their product portfolios similarly. They sold several brands. Each brand had a “main” product (the bestseller) and, sometimes, secondary products, usually in the light segment or presenting special characteristics. Thus, all firms had to retire between 20% and 50% of their product portfolio from the market, mainly affecting the light product segment.

Philip Morris suffered the largest impact from the policy because a larger share of its products shared brand names, and it could not replace products immediately. For instance, Philip Morris had to discontinue sub-varieties of the Marlboro brand, which accounted for more than 25 % of total sales.⁹ Despite Philip Morris’ flagship brands retaining a fraction of the consumers whose products disappeared, its market share decreased in the months following the policy. Hence, one year after the policy began, the company reintroduced products in the light segment, setting strikingly low prices—with respect to the market average and Philip Morris’ prices before the policy was passed. Figure 1 shows per-cigarette unit margins relative to taxes. While unit margins were around 0.5-0.7 before the policy, they dropped to about 0.10-0.25 between January and March 2010. They returned to the original level around 2014.

We stress that firms’ response to the policy is consistent with investing in consumers in a model of competition under habit formation. If consumers’ choices are entirely due to persistent preferences, then firms have no incentives to invest or harvest consumers, and the model reverts to a static

⁹Informal talks with industry agents suggest that international property rights agreements delayed the introduction of new brands to the country. Philip Morris International sued the Uruguayan government because of this policy. They considered it violated international property rights agreements. Philip Morris claimed that the sudden prohibition to commercialize several trademark products under the Marlboro brand caused sizable pecuniary damage. See Figure A.5 for a representation of how it affected Philip Morris’ portfolio. The Uruguayan government finally won the case, and the norm remains.

Figure 1: Philip Morris' real unit margin.



Note: The unit margin is computed against the taxes. It does not include any other marginal costs the firm might have while producing cigarettes. 25% of the drop was due to nominal price drops, while the other 75% was due to tax increases not passed to prices.

competition model. For instance, it would require artificial, often implausible, cost changes to justify introducing new products by pricing them at cost, as in Figure 1. However, if consumers are state-dependent, firms can affect future demand by changing prices today, which appears consistent with firms' aggressive penetration pricing strategies.

3 A dynamic model of competition under inertia

This section presents a dynamic competition model for consumers with addiction and product loyalty. To simplify notation, we assume firms produce one product. In the general model, a fixed number of firms make portfolio and pricing decisions. Although the extension to multi-product firms is almost immediate, we stress any additional assumption when required. Section 3.2 analyzes firms' equilibrium behavior focusing on a pure strategy Markov perfect equilibria [Maskin and Tirole, 1988]. In Section 3.3, we present an alternative, empirically tractable equilibrium concept.

3.1 Model Setup

Firms & Time Horizon The industry evolves over discrete time in an infinite horizon. We denote each period by $t \in \mathbb{N}$. There are F firms. Firm f decides whether to offer its product (which we also denote by f) at period t and sets its price. Consumers' choice set at t is $\mathbb{J}_t \in \{0, 1\}^F$, with $\mathbb{J}_{f_t} = 1$ if product f is offered at t . We call the set of all possible choice sets \mathcal{J} , of dimension 2^F .

Demand Demand is based on the differentiated product discrete choice model but also incorporates dynamic elements of consumer choice. In particular, we allow for habit formation in smoking

(addiction) [Ciliberto and Kuminoff, 2010] and product loyalty [Dubé et al., 2010]. Consumer i in period t chooses a single product or the outside option—*not to smoke*. Consumers have endogenous, time-varying, individual preferences. Utility depends on the state $z \in \{1, \dots, N\}$, i.e., the product they patronize, or $z = 0$ if they were not smoking the previous period. Consumer i 's utility from consuming product j in market t , if she was in state z , is

$$\begin{aligned} u_{ijt}(z, \mu^D) &= \delta_t + \delta_j + \sum_r \sum_k (D_i^r X_j^k) \gamma^{kr} + \eta_0 1\{z \neq 0\} + \eta_1 \{z = j\} + \varepsilon_{ijt} \quad \text{if } j \neq 0 \\ u_{i0t}(z, \mu^D) &= \varepsilon_{i0t} \quad \text{otherwise} \end{aligned} \quad (1)$$

$\sum_r \sum_k (D_i^r X_j^k)$ represents the individual-specific static component of utility: X_j^k are observable product characteristics, and D_i^r denotes demographic variables. Consumer demographics define N consumer types. δ_t is the mean valuation for cigarettes in period t , δ_j is product j 's mean utility. Furthermore, consumers' utility is state-dependent. First, individuals get extra utility η_0 (if positive) from consuming any inside goods if they were previously affiliated with any product. We coin this term “addiction” since it makes consumers more likely to choose an inside good if they previously consumed any cigarette. Finally, η_1 indicates that individuals get higher utility (if $\eta_1 > 0$) from the good they are affiliated to than from any other, and we denote it “product loyalty”. μ^D summarizes all consumer preferences: $\mu^D = (\delta_t, \delta_j, \gamma^{kr}, \eta_0, \eta_1)$. ε_{ijt} is a type I extreme value error term.

Under these assumptions, today's choices depend on the decisions of the previous period. Thus, demand at t is a function of product characteristics, prices, and lagged market shares $S_{t-1} \in [0, 1]^{F \times N}$ for every product-consumer type, taking into consideration the available choice set, \mathbb{J}_t . Note that the demand depends on the whole vector of market shares for each product by each type of consumer. This dependence is essential to define the states of the game and suggests it grows exponentially with the number of products and consumer types. Letting M be the market size, which we fix throughout time, we can write the demand for product f at time t as $D_{ft}(p_t, S_{t-1}, \mathbb{J}_t; \mu^D) = M \times S_{ft}(p_t, S_{t-1}, \mathbb{J}_t; \mu^D)$

Discussion - Consumer Preferences We simplify consumer preferences in three dimensions. First, we depart from classic rational addiction models [Becker and Murphy, 1988] in two ways: individuals are myopic and do not experience heterogeneous dependence on cigarettes due to the intensity of past consumption. Our assumption about myopic consumers is not uncommon in the literature—see, for instance, Tuchman [2019]—and it arises from the challenge to differentiate it

from forward-looking behavior empirically.¹⁰ Moreover, Arcidiacono et al. [2007] shows that the distinction between myopic and forward-looking behavior has almost identical implications for the estimation of addiction. Our counterfactuals focus on reducing addictiveness and brand loyalty. Thus, the myopic model approximates consumers’ behavior well, even if it does not capture the underlying mechanism. Our assumption about homogeneous addictiveness is empirically supported. Although we would have expected long-term consumers or heavy smokers to be less likely to quit, our data does not show significant heterogeneity in quitting rates (or product switching) across the age profile of smokers or the time smoking.¹¹

Second, we model demand as a discrete choice process despite the intensive margin having a relevant role in cigarette consumption. Although the intensive margin is relevant to accurately characterize the demand and policy evaluation, accounting for the intensive margin would significantly expand the model’s state space. In that case, firms would have to track the entire distribution of consumption. Third, we assume there is no unobserved persistent preference heterogeneity. Introducing unobserved persistent heterogeneity would also increase the size of the state substantially as it expands the number of consumer types to evaluate. We have chosen to follow this simplistic demand representation because it can still capture firm behavior well, which is the primary objective of our work. The fact that we can accurately represent firm behavior suggests that, in practice, companies might only track some naive approximations of consumer heterogeneity. For instance, not accounting for unobserved persistent preferences might imply that we incorrectly assign some choice persistence to state-dependent utility. However, if firms made similar “mistakes”, our model would still be a good approximation of the actual market dynamics.¹²

Variable Profits Prices, demand, and marginal production costs (c_{ft}) determine per-period profits $\pi_f(p_t, S_{t-1}, \mathbb{J}_t, c_{ft}, \mu^D) = (p_{ft} - c_{ft})D_{ft}(p_t, S_{t-1}, \mathbb{J}_t; \mu^D)$. We assume that all firms’ marginal costs c_{ft} are public information. We decompose products’ marginal costs into a time-invariant, product-specific component and a time-varying term common to all products. Thus, we can write marginal costs as $c_{ft} = c_t + c_f$. This assumption makes particular sense in the tobacco industry: the marginal cost of producing a cigarette is well-known, stable, and largely homogeneous across firms.

Fixed Costs Each period, firms decide whether to offer their products. We assume firms are established in the market and do not need to pay entry costs to provide new products. They only

¹⁰There is, however, some evidence in favor of rational addiction [Becker and Murphy, 1988, Chaloupka and Warner, 2000, Gruber and Köszegi, 2001, Arcidiacono et al., 2007]

¹¹Table A.5 presents linear regressions confirming these results. There is, however, some heterogeneity in the amount of cigarettes consumed in the past.

¹²This observation highlights the usefulness of considering firm behavior to separately identify unobserved persistent preference from structural state dependence. We leave the exploration of this idea for future research.


need to pay a fixed cost Θ_f^{FC} to keep their products in the market. This assumption is sensible since it is relatively easy for established firms to introduce and sell new products if they are profitable. Even outsiders to the tobacco industry can import international brands to distribute them nationally.

Fixed costs are private information, i.i.d realizations across products and time, from distribution F_{FC} . Fixed costs are the only source of firms' private information. That is, all firms know each other's marginal costs and the distribution of fixed costs but not the specific realizations of the latter. Moreover, let $\chi_{ft} \in \{0, 1\}$ denote product f 's participation choice, where a value of 1 indicates that the product will be offered at period $t + 1$. The assumption of private information of fixed costs is usual in the literature to ensure that an equilibrium exists. In our case, although it does not guarantee existence, it provides tractability.

Timing of Events The timing of the stage game is as follows

1. At the beginning of the period, all firms observe past market shares (current customer base), the product portfolio, production costs, and consumer preferences. Then, they set prices to compete in the product market.
2. Market shares realize.
3. Costs and consumer preferences update, and firms privately draw fixed cost shocks.
4. Firms make portfolio decisions and pay fixed costs accordingly. New products enter the market with zero market share.

Figure 2: Stage Timeline



Our timing assumption is similar to Besanko et al. [2014] and differs from several papers in the literature, such as Wollmann [2018] and Fan and Yang [2020]. While the latter papers intend to highlight the effect that product positioning has on subsequent pricing effects, we focus on the dynamic incentives that firms have to price. Among these incentives, we want to characterize the possibility of excluding rivals' products from the market, which our timing assumption allows us to do.¹³

¹³Our interest in the exclusionary effects of pricing is motivated by the fact that tobacco companies have been

Transition Dynamics Today's prices, past market shares, and the current choice set determine the next period's customer base. That is, $S_{ft} = S_f(p_t, S_{t-1}, \mathbb{J}_t; \mu^D)$ if $\mathbb{J}_{jt} = 1$ and 0 otherwise. Equivalently, participation choices fully determine the next period's industry structure. Because firms' fixed costs are private information, entry and exit decisions are random variables from the perspective of the rivals. Thus, firms only need to know rivals' participation probabilities when forming an expectation over future states. We call participation probabilities $\phi_f \in [0, 1]$. Finally, marginal costs and consumer preferences follow an exogenous transition, $dF(c_{t+1}|s, \mathbb{J}_t, c_t, p_t, \chi_t) = dF(c_{t+1}|c_t)$, $dF(\delta_{t+1}|s, \mathbb{J}_t, c_t, p_t, \chi_t) = dF(\delta_{t+1}|\delta_t)$. Note that δ_t is the only component of μ^D that varies throughout time.

State Space All payoff relevant variables are past market shares (by consumer type), industry structure, marginal costs, and consumer preferences. Although firms observe the private information shocks before making participation choices, we show how to integrate them to keep the state space equivalent to a game of complete information. Thus, let the commonly observed vector of state variables be \mathbb{X}_t . This is defined as $\mathbb{X}_t = (S_{t-1}, \mathbb{J}_t, c_t, \delta_t)$, where $S_{t-1} \in [0, 1]^{F \times N}$, $c_t \in \mathbb{R}$, $\delta_t \in \mathbb{R}$, and $\mathbb{J}_t \in \{0, 1\}^F$.¹⁴

Firms Objective and Choices Firms set prices and make portfolio decisions to maximize expected discounted profits.

$$V_f(S_{t-1}, \mathbb{J}_t, c_t, \delta_t) = \max_{p_f, \chi_f} \mathbb{E} \left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \{ \pi_f(p_\tau, S_{\tau-1}, \mathbb{J}_\tau, c_\tau, \delta_\tau) - \chi_{f\tau} \Theta_f \} | \mathbb{X}_t, \Theta_f \right] \quad (2)$$

where the expectation is taken over current firms' participation actions, future values of the actions, private shocks, and state variables.

Bellman Equation If firm behavior is given by a Markov strategy profile σ , then we can write firms' expected profits recursively. At the last stage, when private information costs are realized, the value of the firm under the current period choice set, market shares, costs, and preferences is

accused of predatory pricing practices. Indeed, the national firm, Monte Paz, sued Philip Morris for predatory pricing due to the pricing at costs we saw in Section 2. In that instance, prices were proved to be below costs, and the defendant did not argue that price drops were due to cost shocks nor that there was any significant change in the firms' cost structure. In the first instance, Philip Morris was found guilty of predatory pricing practices but acquitted in 2018 by a higher court. This lawsuit provides rare insights into Philip Morris' motivation behind its aggressive price strategy. Philip Morris alleged that "Philip Morris reduced its suggested prices to consumers and dropped wholesale prices" [to] ". . . revert the dramatic market share lost due to Monte Paz response with respect to OPPB. . .". Authors translation from Spanish version.

¹⁴Observe that $\mathbb{J}_{jt} = 0$ cannot be interpreted as if lagged shares were 0 since it internalizes that the product is not currently being offered. In contrast, $\mathbb{J}_{jt} = 1, S_{j,t-1} = 0$ indicates that product j does not have a loyal base, but it is available in the market. Nevertheless, it is true that if $\mathbb{J}_{jt} = 0$ then $S_{j,t-1} = 0$ for all j .

$$U_f(S_t, \mathbb{J}_t, c_{t+1}, \delta_{t+1}, \Theta_f | \sigma) = -\sigma_f^\chi(S_t, \mathbb{J}_t, c_{t+1}, \delta_{t+1}, \Theta_f) \Theta_f + \beta \int V_f(S_t, \mathbb{J}, c_{t+1}, \delta_{t+1} | \sigma) dP(\mathbb{J} | \sigma^\chi) \quad (3)$$

where \mathbb{J} represents a random variable whose elements are possible choice sets, and belong to \mathcal{J} .

Then, moving backward to the first stage, firms set prices taking into consideration the state $(S_{t-1}, \mathbb{J}_t, c_t, \delta_t)$ and continuation payoffs before participation choices are taken: integrating U_f over Θ_f . The value at this point can be written as

$$V_f(S_{t-1}, \mathbb{J}_t, c_t, \delta_t | \sigma) = \pi_f(\sigma^p, S_{t-1}, \mathbb{J}_t, c_t, \delta_t) + \int \left(\int U_f(S_t(\sigma^p), \mathbb{J}_t, c_{t+1}, \delta_{t+1}; \Theta_f | \sigma) \times dF(c_{t+1} | c_t) dF(\delta_{t+1} | \delta_t) \right) dF_{\Theta_f} \quad (4)$$

Note, also, that preferences and costs update between the beginning and end of the period.

3.2 Optimal choices

We solve the stage game backward to analyze firms' decisions, from the entry/exit phase to the price-setting stage.

3.2.1 Firms' Participation

Suppose there are only two firms. If a firm participates in the market, its continuation payoff - omitting shares, costs, and preferences- is $\beta \left(V_1((1,1))\sigma_2^\phi + V_1((1,0))(1 - \sigma_2^\phi) \right)$ while if it does not participate, it is $\beta \left(V_1((0,1))\sigma_2^\phi + V_1((0,0))(1 - \sigma_2^\phi) \right)$. Therefore, firm 1 participates in the market if and only if $\Theta_1 \leq \beta \left(\sigma_2^\phi (V_1(1,1) - V_1(0,1)) + (1 - \sigma_2^\phi) (V_1(1,0) - V_1(0,0)) \right) = \bar{\Theta}_1(\sigma_2^\phi)$. More generally, the threshold $\bar{\Theta}_1$ depend on the states of the game $S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}$. Thus, we can write optimal participation policies as the following cutoff rule

$$\sigma_f^\chi(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}, \Theta_f) = \begin{cases} 1 & \text{if } \Theta_f < \bar{\Theta}_f(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}, \sigma_{-f}^\phi) \\ 0 & \text{o/w} \end{cases}$$

where the threshold is the difference in expected continuation payoffs between participating in the market or not, taking into consideration that other products participate according to rule σ_{-f}^ϕ . Therefore, equilibrium participation policies solve the fixed-point problem

$$\sigma_f^\phi(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}) = F_\Theta(\bar{\Theta}(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}, \sigma_{-f}^\phi)) \quad \forall f \quad (5)$$

3.2.2 Firms' Pricing

Next, observe that, in the single-product case, representing participation choices by the probability of offering product f is without loss of information –in Appendix, B.1, we discuss the assumptions to obtain an equivalent representation for multi-product firms.¹⁵ Then, rewriting firm f 's problem taking participation probabilities as controls and integrating U_f over realizations of Θ_f we get $U_f = -E \left[\Theta_f \times 1\{\Theta_f \leq \bar{\Theta}(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}, \sigma_{-f}^\phi)\} \right] + \beta E[V(S_t, \mathbf{J}, c_{t+1}, \delta_{t+1}) | \sigma^\phi]$, where the second expectation is taken over all possible choice sets \mathbf{J} , according to participation probabilities σ^ϕ and exogenous distributions. Therefore, in the price-setting stage¹⁶, the Bellman equation of firm f is

$$\begin{aligned} V_f(S_{t-1}, \mathbb{J}_t, c_t, \delta_t | \sigma) = & \max_{p_f} \{ \pi_f(p_f, \sigma_{-f}^p, S_{t-1}, \mathbb{J}_t, c_t, \delta_t) - \\ & \int (E \left[\Theta_f \times 1\{\Theta_f \leq \bar{\Theta}(S_t(p_f, \sigma_{-f}^p), \mathbf{J}, c_{t+1}, \delta_{t+1}, \sigma_{-f}^\phi)\} \right] + \\ & \beta E[V(S_t(p_f, \sigma_{-f}^p), \mathbf{J}, c_{t+1}, \delta_{t+1}) | \sigma^\phi] dF(c_{t+1} | c_t)) dF(\delta_{t+1} | \delta_t) \} \end{aligned} \quad (6)$$

Taking derivatives of Equation 6 with respect to p_f , applying an envelope condition on firms' own entry/exit optimal choices, and noting that $\sum_{\mathbf{J} \in \mathcal{J}} \frac{\partial Pr(\mathbf{J} | \sigma^\phi)}{\partial \sigma_r^\phi} V_f(\mathbf{J}) = E[V_f | \mathbf{J}_r = 1] - E[V_f | \mathbf{J}_r = 0]$, we get FOC for dynamic prices:

$$\begin{aligned} p_{ft} = & \underbrace{\left(c_f - \frac{\beta}{M} \frac{\partial EV_f}{\partial S_{ft}} \right)}_{\text{Virtual Cost}} - \underbrace{\frac{S_{ft}}{\frac{\partial S_{ft}}{\partial p_{ft}}}}_{\text{Static Markup}} - \underbrace{\sum_{k: \mathbf{J}_{kt}=1, k \neq f} \frac{\frac{\partial S_{kt}}{\partial p_{ft}}}{\frac{\partial S_{ft}}{\partial p_{ft}}} \left(\frac{\beta}{M} \frac{\partial EV_f}{\partial S_k} \right)}_{\text{Dynamic Business Stealing}} \\ & + \underbrace{\frac{\beta}{M} \sum_{k: \mathbf{J}_{kt}=1} \left(\sum_{r=1, r \neq f}^F (E[V_f | \mathbf{J}_r = 1] - E[V_f | \mathbf{J}_r = 0]) \frac{\partial \sigma_r^\phi}{\partial S_{kt}} \right) \frac{\partial S_{kt}}{\partial p_{ft}}}_{\text{Entry Deterrence/Exit Inducing}} \frac{\partial S_{ft}}{\partial p_{ft}} \end{aligned} \quad (7)$$

¹⁵Letting $\mathbb{X}' = (S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1})$, observe that $\sigma_f^\chi(\mathbb{X}', \Theta_f) = 1\{\Theta_f \leq \bar{\Theta}_f(\mathbb{X}')\}$, hence $\sigma_f^\phi(\mathbb{X}') = \int 1\{x \leq \bar{\Theta}_f(\mathbb{X}')\} dF_\Theta(x) = F(\bar{\Theta}_f(\mathbb{X}'))$ and $\bar{\Theta}_f(\mathbb{X}') = F_\Theta^{-1}(\sigma_f^\phi(\mathbb{X}'))$.

¹⁶In this section we assume there is only one consumer type, which reduces the states of the game to the market shares by product. This is to obtain expressions that are analogous to the ones we would get in our equilibrium as defined in Section 3.3.

Investing and Harvesting The first two terms on the RHS are almost identical to firms’ optimal pricing without consumer inertia and would also appear in the case of a monopolist. They indicate that firms mark up some measure of their cost following a rule that depends on the elasticity of demand. However, the key difference is that when there is habit formation, customers become a valuable asset that firms invest in. An additional customer of product j changes firm f discounted expected long-term profits by $\psi_{fj} = \beta \frac{\partial EV_f}{\partial S_{jt}}$. This is the value of an additional customer to the firm. In our model, firms invest in bringing consumers to their products by decreasing prices. Therefore, the additional value of a customer to the long-term profits of the firm enters the price FOC as the negative of a cost. We follow Besanko et al. [2010] and call the difference between the marginal costs and this additional value *virtual costs*.¹⁷ We can interpret an increase in the incentives to invest in the customer base as a decrease in the virtual cost of serving them. Then, firms mark up virtual costs based on consumers’ elasticity. However, note that these markups are a function of the products’ locked-in customer base. The more customers a firm has, the more inelastic the residual demand is, and the higher the markup firms want to set. In other words, if a firm has more locked-in customers, it can extract more value from them; that is, it can “harvest” them more. Therefore, the first two terms of Equation 7 recast the well-known investing and harvesting incentives [Farrell and Klemperer, 2007] into a FOC that resembles the usual inverse elasticity pricing rule.

Strategic Incentives Third, firms internalize the business stealing effect on all products in the market, even if they do not jointly control them. Firms consider their prices’ impact over all other products in the market because they understand that stealing customers from rivals will trigger a competitive response in the future. In principle, this effect can create upward or downward pressure on prices. In the multi-product case, firms have two sources of business stealing effects within products in their portfolio. First, the usual static effect. Second, an additional term that considers the long-term losses that sealing customers from other products in its own portfolio generates. Finally, participation choices introduce a fourth term with no counterpart in the static case. Firms can affect rivals’ participation by changing their next-period loyal base. This new mechanism creates incentives to deter rivals’ participation by lowering their access to the market. These terms illustrate that firms have *rational* incentives to induce exit or prevent entry in markets with inertia. The relative size of these strategic incentives is an empirical matter.

3.3 Equilibrium

This section introduces a framework that facilitates using our dynamic model in empirical applications. First, note that the MPE is not a suitable equilibrium concept in our setting. The pay-

¹⁷Note that the virtual costs depend on the size of the firm’s locked-in customer base, but also on other firms’ customer base. A similar intuition is encountered in Hortaçsu et al. [2022], though they refer to virtual costs as opportunity costs.

off relevant variables are past market shares (by consumer type), industry structure, time-varying marginal costs, and mean cigarette valuation. Suppose, for instance, ten relevant products are in the market, and we can summarize the population into four relevant consumer types. Even if we work with a coarse approximation of the value function using ten grid points, the state space would be approximately $11^{(10 \times 4)}$ plus the size of the exogenous shocks. Neither researchers nor market participants can track this state space for computational reasons.

We aim to construct a tractable equilibrium concept that captures the key strategic interactions between firms. This objective is particularly challenging because, in the game we study, firms' dynamic controls affect the evolution of their own and rivals' states. This property contrasts with "capital accumulation" games, where firms' dynamic controls only affect their own state's evolution. The fact that firms' controls do not affect rival states limits some of the strategic considerations and facilitates the summary of fringe agents' states into a few aggregates or moments. Weintraub et al. [2008]' Oblivious Equilibrium (OE) and Ifrach and Weintraub [2017]'s Moment-based Markov Equilibrium (MME) leverage this intuition to construct a computationally feasible approximation of the MPE in capital accumulation games. In the OE case, each firm tracks its own state and assumes rivals are at their stationary level. To the extent that firms do not respond to rivals' positions, this equilibrium concept severely limits the range of strategic responses. On the other hand, in an MME, firms condition their strategies on some summary statistics of their rivals' states. As a result, such relaxation can account for many of the relevant strategic incentives present in the MPE of capital accumulation games.

However, in the MME, the aggregate moments evolve following a Markov process entirely independent of firms' own strategies. Thus, the straightforward application of this concept to our game would limit some of the strategic interactions we presume exist in the MPE of dynamic pricing games such as ours. For instance, while the MME would lead to strategic considerations in the investing-harvesting tradeoff, it would eliminate the dynamic business stealing and entry-deterrence/exit-inducing incentives highlighted in the previous section, which have been shown to impact industry dynamics significantly [Besanko et al., 2010, 2014]. In this section, we describe the natural extension of MME to allow firms to influence the evolution of the aggregate states through their pricing. We refer to this equilibrium concept as Moment-based Markov Equilibrium with Dynamic Spillovers (MME-S). In Appendix Online I.6, we show that our equilibrium concept approximates the underlying MPE well and improves over the canonical application of the OE and MME.¹⁸

¹⁸Although the MPE is computationally infeasible, we can compare the equilibrium outcomes that result in the long run, for whose computation we do not need to solve it, to those generated by alternative equilibrium definitions.

3.3.1 Moment-based Markov Equilibrium with Dynamic Spillovers

Under our approach, firms' best responses depend only on a subset of all the payoff-relevant states, I_f . This subset contains information on the aggregate shares of products in their own portfolio, any other relevant market shares (which might be dominant firms or close competitors), and an aggregate state representing the total sales of all remaining. Formally, firm f tracks the market shares of T^f products with $\#T^f \leq F$. S_{t-1}^f represents the vector of past market shares of all products that belong to T^f and $\bar{S}_{t-1}^f = \sum_{k \notin T^f} S_{k,t-1}$ the sum of all past shares of non-tracked products. We summarize this information in the firm specific vector $z_t^f = (S_{t-1}^f, \bar{S}_{t-1}^f)$. In addition, firms have information about the current choice set, the common component of costs, and mean valuations for cigarettes. Thus, $\xi_t = (\mathbb{J}_t, c_t, \delta_t)$ also belong to I_f . Moment-based strategies are functions from the space information set \mathcal{S}_f to the space of actions (prices and entry/exit decisions): $\tilde{\sigma}_f = \tilde{\sigma}_f(z_t, \xi_t) : \mathcal{S}_f \rightarrow [0, 1]^{\mathbb{J}_f} \times \mathbb{R}$.

Our main innovation is allowing firms to influence the aggregate state's transition kernels through their dynamic controls. The transition kernel of ξ_t is known and does not require any approximation. On the other hand, firms' f transition kernel of next period market shares, when firm f plays moment-based strategy $\tilde{\sigma}_f^p$ and rivals play $\tilde{\sigma}_{-f}^p$ is defined by

$$S^e(z_t^f, \xi_t; \tilde{\sigma}) = E [S(\{S_{ijt-1}\}, \xi_t; \tilde{\sigma}_f, \tilde{\sigma}_{-f}(I_{-f})) | I_f; \tilde{\sigma}] \quad (8)$$

Even though we are restricting firms' information set to circumvent the curse of dimensionality, firms still internalize the effect of their prices on rivals' next-period shares: they are determined by the derivatives of the *expectation*, $\frac{S_{-f}^e(z_t^f, \xi_t; \tilde{\sigma})}{\partial p_f}$. Note as well that firms take this expectation over both $\{S_{ijt-1}\}$ and I_{-f} , which are unknown from firm f 's perspective. To construct this expectation, we follow Fershtman and Pakes [2012] and Ifrach and Weintraub [2017]. Concretely, firms draw realizations of last period market shares (the customer base) from the long-run distribution of states over the recurrent class, which we can forward-simulate from any guess of firms' strategies $\tilde{\sigma}$. Naturally, firms can then condition on their information set, which makes some distributions more likely than others. For instance, a firm has perfect information about the shares they track. Moreover, if a firm had 99% of the market share in the last period, particular distributions of rivals' shares are possible. Then, for each realization of market shares, it is easy to construct a corresponding realization of rivals' information set I_{-f} since it implies summing over different groups of market shares. Finally, firms need to evaluate rivals' policies $\tilde{\sigma}_{-f}$ at the corresponding realizations of rivals' information set to construct a realization of rivals' prices $p_{-f}(I_{-f})$. The construction of expected static profits is analogous.

Furthermore, we redefine the value function. When firm f plays a moment-based strategy $\tilde{\sigma}_f'$ and

rivals follow strategies $\tilde{\sigma}_{-f}$, the moment-based value of firm f is,

$$\tilde{V}_f(I_f|\tilde{\sigma}'_f, \tilde{\sigma}_{-f}) = \pi^{e(f)}(I_f; \tilde{\sigma}_f, \tilde{\sigma}_{-f}) + \beta \int V_f(I'_f|\tilde{\sigma}'_f, \tilde{\sigma}_{-f}) pr(I'_f|I_f, \tilde{\sigma}'_f, \tilde{\sigma}_{-f}) \quad (9)$$

Then, we can define our equilibrium concept:

Definition 1. MME-S

The equilibrium consists of

1. *Price and participation policies $(\tilde{\sigma}^{*p}, \tilde{\sigma}^{*\phi}) : \mathcal{I} \rightarrow R^F \times [0, 1]^F$*
2. *Expected discounted value of current and future net cash flow conditional on own strategies $\tilde{\sigma}'$, rivals' strategies $\tilde{\sigma}$ at any information set I_f : $\{\tilde{V}_f(I_f|\tilde{\sigma}'_f, \tilde{\sigma}_{-f}) \text{ for } f \in \{1, \dots, F\}\}$*

such that

1. *Strategies $\tilde{\sigma}_f^*$ are optimal when rivals behave according to $\tilde{\sigma}_{-f}^*$ at every information set I_f for all f*
2. *Firms' beliefs are consistent with equilibrium play in the sense of Equation 8 and Equation 9.*

Discussion There are a few issues worth noting about our equilibrium definition. First, as in the OE and MME cases, ours is not an equilibrium in the proper sense but rather an approximation of the underlying MPE. The approximation S^e is not necessarily Markov, even if the transition of the underlying variables is. Therefore, firms would be better off conditioning on more information, particularly in longer horizons of the states. In this regard, our choice of the variables that compose firms' information sets is somewhat arbitrary. While Fershtman and Pakes [2012] offers an alternative solution, we are also interested in computing equilibrium policies outside the recurrent class. As highlighted by Aguirregabiria et al. [2021], machine learning and artificial intelligence are promising avenues to tackle this issue. In Online Appendix I.6, we use forward simulation to show that expected market shares and prices are close to realized outcomes.

Second, observe that summarizing rivals' market shares by summing them is equivalent to assuming that strategic incentives are homogenous across rivals. This assumption is admittedly strong. For instance, a firm might have a more significant incentive to foreclose a competitor closer to its own products in the product space than firms with far-off products. We can accommodate this concern by establishing a more complex information set for firms. However, this would increase the problem's dimensionality and make it harder to solve. The right tradeoff depends on each specific situation.

Existence & Multiplicity In Online Appendix II, we use simulations to show that an equilibrium exists for a wide range of parameters. However, we do not have a proof of existence. In particular, we cannot use the available results in the literature [Doraszelski and Satterthwaite, 2010, Escobar, 2013] to prove it. Nevertheless, the lack of proof of existence does not arise due to the game’s dynamic nature or firms’ participation decisions. Consumer inertia introduces heterogeneity into the demand, making it a special case of the mixed logit distribution with discrete types. Thus, we cannot ensure that the optimal price correspondence is convex-valued conditional on rivals’ strategies and taking continuation values fixed.¹⁹ Furthermore, there is no guarantee that the equilibrium, if it exists, is unique. The sources of multiplicity are hard to isolate as well. On the one hand, dynamic games of price competition with consumer inertia, network externalities, or learning-by-doing are likely to present multiple equilibria [Besanko et al., 2010, Reguant and Pareschi, 2021]. On the other hand, participation costs also introduce multiplicity, even in static games [Pesendorfer and Schmidt-Dengler, 2008].

Implementation The algorithm for computing the equilibrium is standard and based on approximated value function iteration using parametric interpolation methods (Chebyshev polynomials). We initialize the algorithm using the value function’s value at the steady state of the game without entry and exit at every possible choice set. The computation of this steady state is simple since we can circumvent the curse of dimensionality by imposing equilibrium restrictions at a steady state (see Online Appendix I.3). At each iteration step, firms observe their information sets and choose optimal policies by evaluating payoff-relevant states according to their beliefs. We update firms’ values using the new policies and iterate until convergence. Online Appendix I describes the algorithm in detail.

4 Estimation

Although we can derive some theoretical regularities from the model, the long-run equilibrium outcomes that result from decreasing habit formation depend on the model’s primitives. For instance, it is known that lowering consumer inertia can lead to lower or higher prices [Dubé et al., 2009], more or less entry [Farrell and Shapiro, 1988], and ultimately more or less consumption. There are three sets of primitives to recover from data: consumer preferences, marginal costs, and participation costs. We first describe the empirical demand specification and stress the key assumptions.

¹⁹This is a well-known problem in the literature, which has prevented the development of general existence results for games of imperfect competition with mixed logit demand, even in static settings. As noted by Caplin and Nalebuff [1991]: “Without any restrictions on market demand, it may be that two extreme strategies, either charging a high price to a select group of customers (for whom the product is well positioned) or charging a low price to a mass market, both dominate the strategy setting an intermediate price. This issue has been a major stumbling block in the study of existence”. See Aksoy-Pierson et al. [2013] for a notable exception.

We then move to the supply side. Finally, we summarize the estimation procedure.

4.1 Econometric Model

4.1.1 Demand

We construct individual consumption probabilities conditional on past choices and aggregate market shares in the following way.

Consumption Probabilities Let $\bar{u}_{ijtm}(z; \mu^D) = \delta_{jtm} + \sum_r \sum_k (D_i^r X_j^k) \gamma^{kr} + \eta_0 1\{z \neq 0\} + \eta_1 \{z = j\}$, and assume ε_{ijtm} has the type-I extreme value distribution i.i.d. across individuals, products, markets, and time. Then the probability of consuming product j , conditional on being affiliated to the product z last period is,

$$s_{ijtm}(z; \mu^D) = \frac{\exp(\bar{u}_{ijtm}(z; \mu^D))}{1 + \sum_k \exp(\bar{u}_{iktm}(z; \mu^D))} \quad (10)$$

where $\mu^D = \{\delta, \gamma, \eta\}$

Market Level Shares Then, market-level shares depend on state and demographic-specific choice probabilities $\{s_{ijtm}(z; \mu^D)\}$, and the joint distribution of demographics and affiliations at market m at period t . We express them as $S_{jtm}(\mu^D) = \int s_{ijtm}(z; \mu^D) dF_{tm}(D_i, z)$. If aggregate market share under affiliation z is $S_{jtm}(z; \mu^D) = \int s_{ijtm}(z; \mu^D) dF(D_i|z)$ ²⁰, then, taking into consideration that affiliation is a discrete random variable, aggregate market shares can be expressed as $S_{jtm}(w; \mu^D) = \sum_{z \in \{1, \dots, J\}} w_{tm}(z) S_{jtm}(z; \mu^D) + w_{tm}(0) S_{jtm}(0; \mu^D)$, where $w_{tm}(z) = S_{z,t-1,m}$.

To construct the outside option, we assume every store could sell cigarettes to 35% of the store's customers, which was the national smoking rate in 2001. To determine the number of stores' customers, we assume they are currently selling cigarettes to a proportion of customers that coincides with the current smoking prevalence within the market in which stores operate. At the aggregate level, this implies that the outside option oscillates between 30% and 45% over our sample period.

²⁰Although we observe the marginal distribution of states $dF_{tm}(z)$ —determined by past market shares—and the marginal distribution of demographics $dF_{tm}(D_i)$ from population surveys, we do not know the joint distribution of demographics and affiliations at the market level. To construct it, we leverage the individual level data and assume that the distribution of demographics conditional on previously patronized products is not store specific, i.e., $dF(D_i|z)$. Nevertheless, we let the demographics conditional on not smoking vary across stores. Then, we can express the joint distribution of demographics and affiliations using the market-specific distribution of states and the average distribution of demographics conditional on states.

4.1.2 Supply

There are two sources of statistical noise on the supply side: a common knowledge shock to the time-varying marginal costs and an unobserved, private information shock to fixed costs. For the first part, we assume the common marginal costs are taxes plus an unobservable part: $c_{kt} = \theta_k^{vc} + tax_t + \sigma_\varepsilon \varepsilon_t^c$, where $\{\theta_k^{vc}\}$ are parameters to recover from data, tax_t is observable, and ε_t is an unobserved marginal cost shock common to all products. It is distributed $N(0, 1)$ and σ_ε is a parameter we wish to estimate.

The second source of statistical noise comes from random fixed costs. We assume all firms face the same fixed-cost distribution. Its average value decreases with the number of products they sell $\mu^{FC}(N) = \theta_S e^{-\theta_R(N-1)}$, where θ_S regulates the scale of the mean value of the fixed cost distribution, while θ_R regulates the rate at which they decrease with the number of products. (θ_S, θ_R) are parameters to recover from data. Finally, we assume the distribution of participation costs is exponentially distributed, simplifying the computation of continuation values and their derivatives.²¹ The frequency of the model is annual and we parametrize $\beta = 0.95$ accordingly.

4.2 Identification & Estimation Procedure

We estimate the model in three stages. First, we recover the demand parameters. Then, we compute the exogenous variables' transitions directly from the data. Finally, we solve the model to recover the firms' primitives.

Step 1: Consumer Preferences

We combine individual-level data with aggregate market shares to estimate demand. We maximize the likelihood of observing individual i choosing product j at market m at time t , restricting product mean values to be consistent with the observed shares across markets. In particular, we follow Goolsbee and Petrin [2004] and maximize the individual likelihood constraining market shares to be consistent with the observed in the product-level data.²²

²¹Exponential distributions of fixed costs are prevalent in the dynamic entry/exit literature (see Pakes et al. [2007]) because of its memoryless property: expectations, conditional on participation, can be expressed in closed form, getting rid of complicated integrals. In our setting, using this assumption, we can express expected fixed costs conditional on offering the product as $E \left[\Theta_k \times 1\{\Theta_k \leq \bar{\Theta}(\mathbb{X}', \sigma_{-k}^\phi)\} \right] = \sigma_k^\phi \times \mu_k^{FC} - (1 - \sigma_k^\phi) \times \bar{\Theta}(\mathbb{X}', \sigma_{-k}^\phi)$

²²We make two adjustments to the standard procedure. Although consumers report choices over eight quarters, we assume the relevant time horizon for firms is one year. Hence, we simulate an intermediate choice for which we have no information, which allows us to estimate dynamic preferences at the annual level. Second, we do not observe the stores where consumers shop. Thus, we assign them to markets probabilistically. These probabilities depend on the chosen products over the periods, consumer and market demographics, and the size of stores.

$$\begin{aligned} \min_{\theta} \quad & \frac{1}{N} \sum_i \sum_{j_i} 1\{\mathfrak{I}_i = j\} \times \log(s_{ijtm}(\mathfrak{I}_{i,t-1}; \delta, \theta)) \\ \text{s.t.} \quad & S_{jtm}(S_{j,t-1,m}, \delta, \eta, \lambda, \gamma) = \hat{S}_{jtm} \end{aligned} \quad (11)$$

In the second stage, we decompose mean utility through linear regression²³,


$$\delta_{jtm} = \delta_{im} + \delta_{jm} - \alpha p_{jt} + \Delta \delta_{jtm} \quad (12)$$

Identification Although consumers' choices to smoke or to quit and which product to choose are highly persistent, it could be due to persistent consumer preferences or structural state dependence [Heckman, 1981]. This distinction has crucial implications for firm behavior. In the former case, firms cannot influence their future demand by changing their actions today. Hence, they would not face any dynamic incentive. Therefore, distinguishing between the two scenarios is the key identification challenge in our analysis. Next, we present the sources of variation in our data that help us address this challenge.

Figure 3 show that the tax policy has been inconsistent over the years, leading to swings in the tax level. These oscillations are a helpful tool to identify smokers' addiction. The ideal experiment is to observe identical individuals facing the same market conditions (prices, choice sets, etc), choosing with and without dependence on cigarettes. The price swings induced by these tax oscillations provide variation that resembles this ideal experiment. For instance, initial tax increases make smokers more likely to quit smoking –see quitting rates in Table 1. Then, the posterior return of taxes to a low level allows us to compare the choices of those who switched out in a similar context as before they quit. To sum up, we can identify the addiction parameter by tracking people over the years and comparing the asymmetry of switching behavior between high and low tax periods.

²³To make the equilibrium computation tractable, we assume firms disregard the store-specific variation in mean utilities. Thus, we extract the product-specific component of mean utility (δ_j), which does not vary throughout time nor stores, and a time-specific valuation for cigarettes (δ_t) that does not change across stores.²⁴

Figure 3: Real Taxes.



Then, we exploit substantial variation in the product portfolio and relative prices to identify product loyalty. We observe large shocks to the product portfolio due to the 2009 policy that forced firms to sell only one product under the same brand name. Consequently, several consumers had to make new active choices, which aids in identifying brand loyalty, as in Handel [2013]. Formally, it provides an additional moment for identifying brand loyalty: the choice probabilities of an individual who faces addiction yet is not locked into any product. Note, however, that we cannot identify individual-specific valuations for different products since the policy effectively eliminated products from the choice set.²⁵ Thus, our strategy uses this moment to determine consumers' persistent preferences for products *characteristics*. Then, any choice persistence that we cannot attribute to these persistent preferences is assigned to brand loyalty. We can also leverage relative price. Our data also contains large relative price swings due to the aggressive reintroduction of products –see Figure 1. The intuition on how transitory relative price changes identify brand loyalty is analogous to the price variation we used to identify addiction.

Finally, in the second stage, we face the usual endogeneity problem between prices and unobserved utility $\Delta\delta_{jtm}$. However, our institutional setting and the available data make this issue less relevant. First, firms set prices at the national level, and almost all stores abide by suggested prices. Thus, p_{jt} is unlikely to be correlated with time-product-market and product-market unobserved shocks. Second, we can control for time-market and product fix-effects. Lastly, we instrument for product-time unobserved shocks using taxes. During our sample period, there were changes in excise and value-added taxes, which created variation at the time-product level.

²⁵In that sense, our setting is different from Handel [2013], where individuals are forced to choose without changing the choice set. However, we can think of his setting as products going out of the market and being introduced anew.

Step 2: Exogenous states' transitions.

Then, we use our estimates of $\hat{\delta}_t$ and observed taxes to compute the two exogenous transitions: the common component of marginal costs and the mean valuation of inside products. We assume they follow independent AR(1) processes. Then, we recover the parameters of the tax process by fitting an AR(1) process to the data: $tax_t = \mu^{tax} + \rho^{tax}tax_{t-1} + \sigma_{\varepsilon^{tax}}\varepsilon_t^{tax}$; $\hat{\delta}_t = \mu^{\hat{\delta}} + \rho^{\hat{\delta}}\hat{\delta}_{t-1} + \sigma_{\varepsilon^{\hat{\delta}}}\varepsilon_t^{\hat{\delta}}$

Step 3: Firms Costs

After estimating demand primitives and exogenous state transitions, there are $J + 3$ remaining parameters: $\theta = \{\theta_k^{vc}, \sigma_{\varepsilon^c}, \theta_S, \theta_N\}$. We estimate them using the simulated method of moments (MSM). According to our model, discrete participation decisions are described by the policy $\tilde{\sigma}^\chi(z_t^f, \xi_t; \varepsilon_t^c, \Theta_{kt}; \theta)$ (recall that $z_t = (S_{t-1}^f, \bar{S}_{t-1}^f)$, and $\xi_t = (\mathbb{J}_t, c_t, \delta_t)$). Thus, at the true parameters θ_0 , $\chi_{kt} = \sigma_k^\chi(z_t^f, \xi_t; \varepsilon_t^c, \Theta_{kt}; \theta_0)$. Equivalently, prices are determined by the optimal policies and marginal cost shocks (they do not depend on the fixed cost's realization of the private shocks): $p_{kt} = \sigma_k^p(z_t^f, \xi_t; \varepsilon_t^c; \theta_0)$. Thus, given the observed data $\{\chi_{kt}, p_{kt}, \{z_t^f\}_f, \xi_t\}_{i=1}^{J \times T}$ ²⁶, an MSM estimator of θ_0 can be generated from the conditional expectations: $E[\chi_{kt} - E[\sigma_k^\chi(z_t^f, \xi_t; \varepsilon_t^c, \Theta_{kt}; \theta_0)|z_t, \xi_t]|z_t, \xi_t] = 0$, $E[p_{kt} - E[\sigma_k^p(z_t^f, \xi_t; \varepsilon_t^c; \theta_0)|z_t, \xi_t]|z_t, \xi_t] = 0$. Online Appendix I.7.2 shows how to use importance sampling to reduce the computational burden of the estimation procedure.

Moments & Identification Average prices in the data –given estimated levels of inertia and conduct²⁷– are informative about marginal costs. Additionally, we leveraged the correlation between participation choices and observed states (tax rates and customer bases throughout time), which is standard in the entry/exit literature. In fact, we could have used the score of the likelihood of participation decisions as moments in the data. Interestingly, the correlation between prices and the customer base also informs fixed costs. Suppose we observe a large drop in the loyal base without a significant price response. This points out that firms assess the product's probability of leaving the market to be high, which informs fixed costs. We are unaware of previous work that exploits firms' prices to identify fixed costs in dynamic settings –see Berry and Pakes [2000] for early work suggesting a similar approach. Similarly, larger unobserved cost shocks implied

²⁶The number of observations is not precisely $J \times T$ since prices are only observable if the products are currently in the market. We do not make it explicitly in the notation to avoid overloading it.

²⁷Although in our setting conduct does not directly identify marginal costs (conditional on prices and demand estimates), it still provides valuable information about the joint distribution of fixed and marginal costs. For instance, a firm that knows a product will likely exit the market will not invest in building a large customer base and increase prices accordingly. We argue that prices and participation choices provide two distinct sequences of marginal cost and continuation probabilities that rationalize them. In a way, we highlight that prices and participation FOC are not collinear. Hence, the static FOC inversion has an equivalent representation in the dynamic model. Figure I.4 presents a graphical illustration of this argument. There, we show the sequence of marginal and fixed costs that generate alternative BAT's flagship average prices and participation rates. Then, observed prices and participation select different cuts of the plane.

lower prices and lower pass-through, while higher marginal costs mean higher prices and lower pass-through. Hence, the level and correlation of prices with taxes help tease apart the product's marginal cost from the common unobserved shock.

Multiplicity We address potential equilibrium multiplicity in two stages. First, we use different methods to argue that the dynamic pricing game without entry and exit has a unique equilibrium for large regions of the parameter space. Then, we use the steady state of this game to provide natural initial values across different parameterizations, which acts as a way of choosing one of the equilibria in the game with entry and exit. We refer the reader to Online Appendix II for details.

Discussion We do not follow any classic approaches to estimating marginal and fixed costs in dynamic models. The standard approach involves inverting static FOC to recover marginal costs and then using those marginal costs to compute static profits, which allows us to recover fixed costs by solving the entry/exit game by maximum likelihood—see for instance [Igami, 2017, Igami and Uetake, 2020, Elliott, 2022]. This approach is not feasible because fixed costs influence prices through the choice set probabilities—see Equation 7. This issue is not specific to our model but is a general concern for many dynamic pricing games. Similarly, we cannot apply standard solution-free dynamic estimation methods such as Bajari et al. [2007] because we observe a single market (the national market), which is common in many situations. Although we observe substantial variation in observed states, computing policy functions from data alone would be challenging. Lastly, we cannot follow MacKay and Remer [2021] reduce form estimation of dynamic incentives since we do not observe the firms' marginal costs.

5 Structural Results & Model Validation

5.1 Consumer Preferences

Habit Formation Inertia is high. Indeed, smokers are willing to pay almost two times the average price for any cigarette. Moreover, they are willing to pay around three times the average cigarette price to repeat their product choice. Naturally, firms do not only target repeated customers, which allows these consumers to pay lower prices than their willingness to pay. There is also a modest amount of consumer heterogeneity. In particular, educated and young customers are less price-sensitive and value light products more. Table 2 presents a summary of consumer preference estimates. Overall, the demand model accurately captures switching patterns between products and in and out of smoking—see Table C.1.

Table 2: Demand Estimates

		Complete Secondary	Working Age
Real Price Per Cig	-0.931	0.046	0.421
s.e	(0.033)	(0.032)	(0.029)
Light		0.080	0.100
s.e		(0.147)	(0.154)
Premium		-0.194	0.259
s.e		(0.136)	(0.124)
Addiction	2.007		
s.e	(0.055)		
Brand Loyalty	3.437		
s.e	(0.045)		
N Individuals Observations	2850		
N Markets	12422		

Elasticity These estimates imply that the mean own-price elasticity in the market is low, around -0.9. The small elasticities are unlikely due to a failure in the identification strategy. Such low demand elasticity estimates are a robust result among studies treating cigarettes as differentiated products. Ciliberto and Kuminoff [2010], Liu et al. [2015] and more recently Tuchman [2019] obtain estimates in the range of $[-1.35, -0.64]$. Additionally, the implied aggregate market elasticity is slightly below 0.4, which aligns with a large body of work computing smoking elasticities. These figures suggest firms are pricing in the inelastic portion of the demand curve. This pricing behavior is consistent with our dynamic competition model since companies have incentives to constrain the markup they set to retain and capture new consumers from whom they can profit in the future. A static competition model could not capture the observed pricing patterns. We return to this point in Section 5.3.

Table 3: Demand Elasticity

	Own-Price	Agg. Market
Baseline Estimates	-0.853	-0.346
Prev. Literature	$[-1.35, -0.64]$	$[-0.5, -0.15]$

Note: Median elasticity in previous literature refers to Ciliberto and Kuminoff [2010], Liu et al. [2015] and Tuchman [2019]. Evans and Farrelly (1998) reported the aggregate elasticity range.

5.2 Firms' Costs & Pricing Incentives

Firms' Costs Overall, the estimated marginal costs of production (without considering taxes) are small and homogeneous, which aligns with accounting estimates. For all firms, taxes represent more than 90% of marginal costs. The mean value of the fixed-cost distribution is low and relatively constant as firms add new products. The expected fixed costs paid represent around 40% of the variable profits the firms make. Table 4 presents the full estimation results.

Table 4: Cost Estimates, Exogenous Process, and Price Decomposition

	Production Cost	Virtual Cost	Mkp	Static BS	Dynamic BS (Own)	Dynamic BS (Other)	Exit/Entry
MP Flagship	0.0 (0.17)	-0.379	2.819	0.089	0.163	0.128	-0.015
MP Regular	0.31 (0.05)	2.128	0.724	0.027	0.922	0.101	-0.009
MP Light	0.17 (0.07)	0.655	0.996	0.03	0.859	0.099	-0.01
MP Specials	0.0 (0.06)	0.527	0.964	0.028	0.86	0.098	-0.01
PM Flagship	0.11 (0.09)	0.674	1.531	0.035	0.04	0.103	-0.001
PM Light	0.17 (0.13)	1.15	0.931	0.047	0.147	0.088	-0.001
PM Regular	0.02 (0.04)	1.387	0.891	0.039	0.168	0.088	-0.001
BAT Standard	-0.01 (0.03)	1.154	0.959	0.016	0.004	0.026	-0.0
BAT Premium	0.18 (0.17)	1.803	0.876	0.012	0.022	0.025	-0.0
	Other Costs						
$\mu \theta_S$	0.42 (0.09)						
$\mu \theta_R$	40.44 (44.82)						
$\mu \sigma_\varepsilon^2$	0.35 (0.34)						

Note: The exogenous process estimates are: $tax_t = 0.243 + 0.860tax_{t-1} + \sigma_{\varepsilon^{tax}} \varepsilon_t^{tax}$; $\hat{\delta}_t = -0.143 + 0.007\hat{\delta}_{t-1} + \sigma_{\varepsilon^{\delta}} \varepsilon_t^{\delta}$. These estimates are obtained from linear regression using 44 quarters. Because the persistence of $\hat{\delta}_t$ is quite low, in our baseline estimates and counterfactuals, we assume firms believe $\hat{\delta}_t = \hat{\delta}_t$, fixed over time. The supply estimates are obtained using importance sampling. The coefficients reported the mean of the obtained distributions. The variance estimates are the following: $\sigma_{mc}^2 = 0.0351$, $\sigma_{\theta_S}^2 = 0.041$, $\sigma_{\theta_R}^2 = 110.35$, $\sigma_\varepsilon^2 = 0.026$. Standard errors are computed using standard formulas from MSM. We decompose prices using the multi-product version of the first-order conditions of the model presented in Section 3. We evaluate the different terms of the decomposition at the observed states of the industry and average them for every product. These are preliminary results.

Pricing Incentives Decomposition: The Role of Strategic Incentives Preference and cost primitives determine firms' pricing incentives. Table 4 shows that virtual costs differ even though production costs and taxes are homogeneous across products. The reason is that consumers from different products have significantly different long-term value for the firm. For instance, locking a consumer in Monte Paz's flagship product is considerably more valuable than making it loyal to a BAT product. Then, we can use our FOC decomposition to gain some understanding of the relevance of strategic incentives. First, we observe that for relatively more minor products within a firm's portfolio (for instance, Monte Paz' Specials and Regular or Philip Morris Regular), dynamic business stealing effects with respect to other products within their portfolio are substantially more important than the static ones. This means that firms have less incentive to expand smaller product lines once we consider the dynamic implications than in a static model.

Finally, our estimated primitives suggest that firms lack significant incentives to induce rivals' exit or deter entry. While our model introduces rational incentives to predate, the existing asymme-

try among firms indicates that ousting rivals from the market is either exceedingly challenging or not lucrative. On the contrary, firms appear more inclined towards softening competition to avoid potential reprisals that could spark fiercer competition in subsequent periods. Indeed, we observe that the dynamic business stealing effects with respect to products of other firms are almost as important as the static business stealing effects—and in some cases even more important, for instance, the case of BAT products. This implies that firms set prices higher than they would otherwise to avoid stealing consumers from rivals. This “cooperative” feature distinguishes competition under inertia from other dynamic pricing models, such as learning-by-doing [Cabral and Riordan, 1994, Besanko et al., 2014] or network externalities [Farrell and Katz, 2005], where firms typically exhibit a more aggressive stance to undercut rivals and secure a favorable competitive position in the future.²⁸

5.3 Model Validation - Firm Behavior under Habit Formation

Before using the structural estimates to evaluate the effect of nicotine caps and uniform packaging, we provide empirical that firms’ pricing strategies are consistent with forward-looking behavior under the estimated addiction and brand loyalty levels. We highlight two features that allow us to validate our assumptions. First, the model rationalizes relatively low unit margins despite exceptionally low elasticities (in this section, we interpret unit margins as the difference between prices and estimated marginal costs, *not* virtual costs). Second, our model accommodates the striking price drops used to introduce new products, observed in the data and discussed in Section 2.

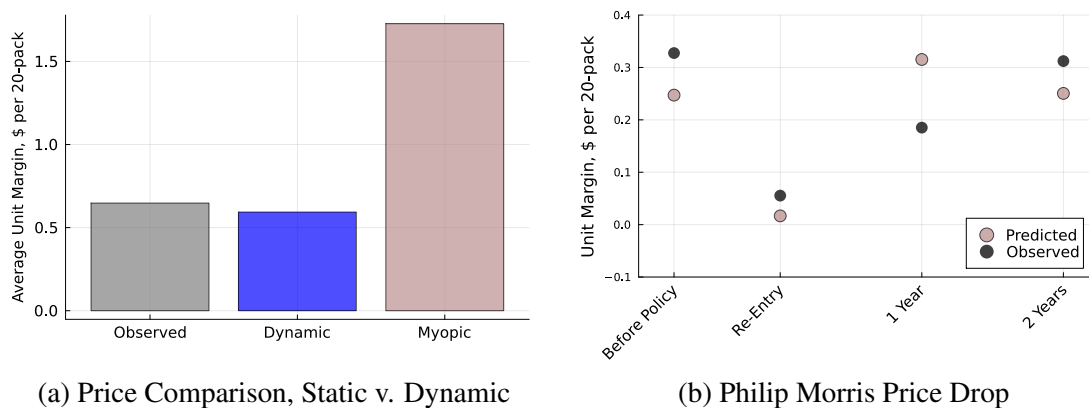
Low Markups Despite Inelastic Demand First, we show that observed elasticity is consistent with firms’ forward-looking behavior. We take the estimated consumer primitives but assume that all costs firms face are taxes. Then, we solve the model both for forward-looking and myopic firms. Moreover, we solve equilibrium for all possible choice sets, assuming products would remain in the choice set indeterminately in every case. These assumptions make our results independent of any supply estimate.

Then, we compare the unit margins in each case. The unit margins myopic firms would set are approximately three times larger than those observed and those set by forward-looking firms. This result is consistent with the idea that firms are forward-looking and that they lower prices to account for the long-term value of customers. In other words, consumers’ state dependence partially explains the observed low elasticity. For firms, low consumer responsiveness to prices is compensated with lower virtual costs, and prices stay relatively low despite strikingly low elasticity be-

²⁸Exploring this possibility is particularly relevant in our setting, considering that Philip Morris faced a lawsuit for predatory pricing due to its aggressive penetration pricing strategy. Our analysis confirms that the aggressive pricing strategies observed in the data are consistent with investing in capturing new customers and not anti-competitive behavior.

cause they internalize the long-term profits, not just the short-term gains. The fact that we assume all costs firms face are taxes makes this result even more striking since it implies that we would need significant negative marginal production costs to rationalize a static competition model.

Figure 4: Model Validation - Firm Behavior under Habit Formation



Note:

The figure shows the observed and predicted unit margins for the Philip Morris light segment at several periods. At the time of re-entry, we set the Philip Morris light segment's loyal base to 0 and evaluated prices according to the equilibrium policies. We average the observed and simulated values for the first two quarters of 2010 (re-entry), as tax changes make it hard to determine whether the firm was setting prices considering the new tax level. Then, the one, two, three, and four years after evaluating simulated and observed prices in the first quarter of every year.

Introductory Pricing: Investing in Consumers Next, we argue that firms' strategies to introduce new products are consistent with our conduct model. More generally, this means that firms' price responses to changes in the size of their loyal customer base are consistent with our competition model at the estimated consumer inertia. Figure 4b compares Philip Morris's aggressive introductory pricing strategy with the model's prediction. The model predicts a sharp margin drop consistent with Philip Morris's pricing around marginal cost. If persistent choices were due to persistent preferences, firms would not have incentives to modify prices as a response to changes in the customer base. Our model captures their response to the loyal customer base, which suggests that firms' beliefs about habit formation are not too different from our estimates.

Overall, at the estimated levels of inertia, our competition model captures optimal markups and price response to changes in the customer base well. Indeed, we have seen that assuming consumers do not derive disutility from prices, that persistent choices are due to persistent preferences, or that firms are myopic would have generated less sound patterns. This analysis confirms that our model of competition, embedded with our empirical estimates, is a good approximation of the actual market dynamics. Finally, Table 5 shows that under the estimated primitives, the long-run steady state of the economy fits shares, prices, concentration ratios, switching patterns, and elasticity very well.

Table 5: Comparison of Actual and Simulated Moments

Statistics	Observed	Policies	Long-Run Simulation
Smoking Rate	0.216	0.216	0.24
AveragePrice	2.274	2.5	2.24
N Products	6.657	6.293	5.987
Switching	0.82	–	0.733
Elasticity	-0.853	–	-1.284
HHI	5782.408	–	6204.18

Note: The simulation column reflects the average across states within the stationary long-run distribution. In this column, the tax process is fixed at \$1.85.

6 Reducing Habit Formation in Tobacco Markets: Nicotine Caps and Uniform Packaging Policies

In this section, we use our estimates to analyze the effect of two policies: capping nicotine content and standardizing tobacco product packaging. We focus on these policies for two reasons. First, while they are currently under discussion in the United States and other countries, there is little empirical evidence to evaluate their impact. Second, they are salient examples of policies aimed at reducing the degree of habit formation in tobacco consumption—either by lowering addiction or consumers’ loyalty to particular products. Nonetheless, our framework is more general and can be used to study various policies such as taxation, increasing the lump sum costs to offer new products or even competition policy. We evaluate the impact of the FDA’s nicotine caps by reducing the estimated “addiction” parameter and analyze uniform packaging by lowering both the mean cigarette valuation and consumer loyalty. We solve equilibrium policies for each counterfactual demand parametrization and simulate the industry 2,500 times. All the results we present in this section are the average over the long-term stationary distribution of states when costs are equal to \$1.8 per pack of 20 cigarettes.

6.1 Nicotine Caps

In 2018, the FDA issued an advanced notice of public rulemaking to establish a maximum nicotine level in cigarettes—and potentially other tobacco products—to minimally or non-addictive levels.²⁹ Nicotine is the main addictive chemical in cigarettes and is responsible for the habit-forming properties of tobacco. The FDA’s goal is to reduce the addiction potential of cigarettes and, conse-

²⁹This proposal is both technically and legally feasible under the 2009 Family Smoking Prevention and Tobacco Control Act (TCA) [Berman et al., 2018]

quently, help addicted smokers quit and limit the amount of new smokers becoming regular users. Results from clinical trials suggest that reducing nicotine to approximately 0.4 mg or less per gram of tobacco allows consumers to quit smoking at higher rates without compensatory effects. Our results isolate the equilibrium impact of reducing addiction from other potential changes in consumer preferences. This analysis complements previous results by providing a framework to understand how firms would respond when selling a less addictive product.

Holding firms’ strategies fixed, if nicotine caps eradicate tobacco dependence (i.e., we shift the addiction parameter to 0.0), the Uruguayan smoking rate would decrease 33.4% to slightly below 13.5% of the population. When firms adjust their strategies, they reinforce the effect, reducing the smoking rate up to 20% more, as shown in Table 6.³⁰ Although average prices decrease around 8.2% because of the policy, they are 4.8% higher than if firms did not adjust their strategies. Moreover, the number of products in the market decreases slightly. Overall, such price and portfolio re-optimization allow firms to increase their average profits by 9.0%

Table 6: Equilibrium outcomes under nicotine caps.

	Baseline	No Addiction/Baseline Strategies	No Addiction/Equilibrium
Smoking Population (%)	20.2	13.45	10.99
Average Price	2.57	2.25	2.36
Product Lines	6.12	6.08	5.77
Profits	355.38	134.36	146.46


Note: For each scenario, we take demand and supply primitives, solve equilibrium policies according to the definition in Section I.2, and simulate the industry for 150 periods 2,500. The outcomes reported in this table are the average over the visited states in the long-run station distribution (after burning the first 100 periods). In Figure C.3, we present the distribution of one key summary statistic of the states: aggregate consumption. In the “Baseline” scenario, the addiction parameter (η_0) is 2.007, and the brand loyalty (η_1) is 3.437. For the “No Addiction/Baseline Strategies”, we set parameters $\eta_0 = 0.0$, $\eta_1 = 2.437$, but the economy is forward-simulated using the equilibrium policies at the baseline addiction level. We fix costs at \$1.85 per 20-cigarette pack in all scenarios. We compute the average smoking rates assuming that the available market size represents 35.6 % of the population, which was the average smoking rate in Uruguay in 2000. Figure C.4 shows average consumption, prices and number of products for several intermediate values of the addiction parameter, with and without firm responses.

To illustrate the role of supply responses, we show how firms’ optimal pricing changes as we eliminate addiction and how that impacts long-run consumption. Figure 5 depicts these changes for the case of the national firm’s flagship product—the intuition holds for any product in the market. The solid black curve indicates the optimal pricing function when addiction is at the baseline level. It determines the average optimal price for each size of its own customer base. The optimal pricing function is a steep function of the customer base: Under current levels of addiction, firms are willing to offer a low price to attract new customers when not many individuals consume their product. As the customer base grows, the firm increases its price to exploit the –partially– locked-in customers.

³⁰These results are consistent with simulation exercise carried out in the US and New Zealand, which found reducing nicotine to non-addictive levels would lower cigarette smoking rates between 40% and 80% [Laugesen and Grace, 2015, Apelberg et al., 2018]. Although we can use these studies as a benchmark, they rely heavily on expert elicitation of crucial parameters.

The direct effect of eliminating addiction is to reduce *smokers'* valuation for cigarettes. So, if firms do not adjust their strategies, the demand for cigarettes drops, and the average long-run consumption decreases as well. If firms do not internalize the new addictiveness scenario, they would significantly lower prices to partially recover their lost customers—consumption drops despite prices decreasing from 2.8 to 2.45. However, once firms internalize that cigarettes are no longer addictive, this is not their optimal response. Firms recognize that they cannot retain customers and profit from them in the future as efficiently as before. Therefore, consumers' long-term value to the firm decreases. This, in turn, discourages firms from lowering prices to attract consumers to the market—a decrease in the long-term value of consumers is equivalent to increasing the virtual cost of serving them. Thus, the optimal pricing function shifts up even though the aggregate cigarette demand decreases. Therefore, we observe a price increase over the relevant customer base range, which further drops consumption and rationalizes why firms tend to reinforce the policy's direct effect.

Figure 5: Price Policy - Monte Paz, Flagship, Addiction



Note: The solid curves (black and blue) represent Monte Paz's optimal pricing function for its flagship product. It establishes the firm's average price for each size of its customer base, that is, the average across all possible prices optimally set depending on rivals' customer bases and the number of smokers. The black policy is calculated when the addiction parameter is $\eta_0 = 2.007$, and the brand loyalty at $\eta_1 = 3.437$. The blue line represents the case where $\eta_0 = 0.0$, $\eta_1 = 2.437$. We fix costs at \$1.85 per 20-cigarette pack in all scenarios. The dashed lines represent the average market share of the product in the long-run stationary distribution, obtained by simulating the economy 2,5000 times. Because these are the average shares, they also represent the average customer base of the product, that is, how many consumers, on average, were consuming the product in the previous period. The black line is the average market share of Monte Paz's flagship product under the baseline scenario. The pink vertical dashed line represents the average customer base when addiction is $\eta_0 = 0.0$ eliminated, but firms continue to play the strategies at the baseline addiction level (blue policies). The blue vertical line represents average market shares when addiction is $\eta_0 = 0.0$, and firms play their optimal strategies according to this scenario. The intersection between the solid black line and the vertical dashed line represents Monte Paz's flagship product's average prices and market shares across the stationary distribution of states in the long run. The intersection between the black policy and the vertical pink line represents prices and shares when the firms play the baseline strategies, but addiction is set at $\eta_0 = 0.0$. Finally, the intersection between the blue policy and the blue vertical dashed line is the prices and quantities the product sells under the equilibrium without addiction.

Robustness The FDA recognizes that nicotine caps could have other unintended results beyond making cigarettes non-addictive [FDA, 2018]. In particular, there is concern that consumers could increase their cigarette consumption to maintain their nicotine intake. Our model did not explic-

itly characterize the relationship between cigarette consumption, addiction, and nicotine content. Although we have access to nicotine and could model demand for nicotine explicitly, the range of variation in nicotine content is outside what the medical literature considers necessary to eliminate the addictive properties of tobacco. While recent clinical trials suggest that moderate reductions in nicotine content could lead to increased smoking intensity and more significant exposure to harmful chemicals, they observe that sufficiently large reductions encourage consumers to quit smoking without compensatory effects [Berman and Glasser, 2019]. Thus, according to these studies, extrapolating compensatory behaviors from a demand model estimated from variation of nicotine within the addictive range would be inappropriate to evaluate the effect of limiting nicotine to non or minimally addictive levels.

6.2 Uniform Packaging: Industry Lobby and Economic Considerations


Australia became the first country to implement uniform packaging for tobacco products in 2012. The policy forced firms to sell all their products in standardized packages with no branding or logos – Figure A.4 illustrates how packages should look. The policy aims to reduce the appeal of tobacco products, particularly to young people, and to enhance the effectiveness of health warnings. The tobacco industry has opposed this policy, arguing that it would reduce brand loyalty, causing smokers to switch to cheaper brands and encouraging price competition between manufacturers [Chantler, 2014]. Since 2017, many countries have adopted similar policies, including the UK, France, and New Zealand. However, the empirical evidence on the effectiveness of this policy is scarce and usually confounded with the effect of other measures implemented simultaneously. Our counterfactual analysis allows us to isolate the potential impact of uniform packaging on tobacco consumption. To consider policymakers and industry arguments, we evaluate the effects of this policy by reducing both the mean valuation for cigarettes and consumers' loyalty.

6.2.1 Price Responses

If firms adjust their prices but keep portfolio strategies fixed, uniform packaging would reduce consumption in most scenarios we analyze. Figure 6 shows the percentage change in long-term average consumption for several levels of brand loyalty and mean valuations (we assume that reducing brand loyalty does not increase quitting at observed prices but only increases switching between products). Contrary to industry claims, we show that even if the effect of the policy does not reduce consumers' appeal for cigarettes, the policy can have significantly favorable effects. If the uniform packaging eliminates brand loyalty, but cigarettes' valuation does not change, consumption would still decrease by approximately 6.8%. In the worst case scenario, when loyalty decreases around 30%, the average smoking rate in the country increases at most 0.68%. If the policy achieves a decline in the valuation of cigarettes of at least 15%, then the policy would reduce

consumption no matter how much brand loyalty declines.

Figure 6: Percentage change in long-term average consumption under uniform packaging.



Note: Note: For each counterfactual scenario, we take demand and supply primitives, solve equilibrium policies according to the definition in Section I.2, and simulate the industry for 150 periods 2,500. This figure compares and reports the percentage change in average consumption over the visited states in the long-run station distribution (after burning the first 100 periods). All comparisons are in relation to the baseline scenario where the addiction parameter is $\eta_0 = 2.007$, and the brand loyalty is $\eta_1 = 3.437$, which is represented in the top-right corner by a black dot. We fix costs at \$1.85 per 20-cigarette pack in all scenarios. Green colors indicate that consumption decreases from our average consumption in our baseline scenario, while red indicates that consumption increases.

Even though our results contradict the tobacco industry’s claims, their arguments are sensible. According to our findings, uniform packaging would make consumers significantly more price-sensitive and lead to more frequent product switching. Indeed, in the worst-case scenario from a policy point of view, i.e., when cigarette valuation is not affected, the average price elasticity increases up to 2.5 times, and the average probability that a consumer repeats its product choice over time decreases by a factor of 3.5. Table 7 shows these figures. In static models, such an increase in demand elasticity would lead to substantial price drops and consumption increases, which is at the core of tobacco companies’ opposition to uniform packaging. Indeed, we show that if firms were myopic, prices would decline approximately 15% –see Figure C.5f. However, we have provided robust evidence that firms also consider the long-term implications of their pricing decisions. Indeed, firms’ change in dynamic incentives as we lower consumer loyalty reconciles this paradoxical result of lower consumption under more elastic demand.

A decline in brand loyalty has two main effects on firms’ optimal pricing strategies, which we illustrate in Figure 7 taking the national firm flagship product as an example. First, as consumers become more price-sensitive, firms’ markups decrease and become less sensitive to the size of the customer base. This effect alone would substantially lower prices, particularly for firms with a large equilibrium customer base (clockwise rotation), as shown in Figure 7a. However, eliminating brand loyalty decreases the long-term value of an additional customer to the firm. This value decreases because consumers are significantly more likely to switch to other products in the future.

Table 7: Equilibrium outcomes under uniform packaging (no change in mean utility).

	No Loyalty			75% Loyalty	
	Baseline	Fixed Portfolios	Equilibrium	Fixed Portfolios	Equilibrium
Consumer Behavior					
Smoking Population (%)	20.37	17.11	18.29	20.46	20.95
Elasticity	-1.56	-3.55	-3.55	-2.47	-2.47
Switching	0.82	0.24	0.23	0.63	0.64
Price Decomposition					
Price	2.59	2.53	2.54	2.55	2.55
Virtual Costs	0.26	1.08	1.08	0.65	0.65
Markups	1.9	0.76	0.76	1.25	1.25
Profits	368.29	307.18	308.79	361.45	360.24
Portfolio Decisions					
Product Lines	6.3	6.49	7.54	6.44	7.0
HHI	6009.58	5839.82	5635.3	6108.36	6108.26
Endog. Entry Cost	0.09	0.05	0.05	0.08	0.08

Note: For each scenario, we take demand and supply primitives, solve equilibrium policies according to the definition in Section I.2, and simulate the industry for 150 periods 2,500. The outcomes reported in this table are the average over the visited states in the long-run station distribution (after burning the first 100 periods). In Figure C.3, we present the distribution of one key summary statistic of the states: aggregate consumption. In the “Baseline” scenario, the addiction parameter (η_0) is 2.007, and the brand loyalty (η_1) is 3.437. For the “No Loyalty” scenarios, we set parameter $\eta_1 = 0.0$, and for the “75% Loyalty” is $\eta_1 = 2.43$. In the “Fixed Portfolios” columns, the industry is simulated using firms’ equilibrium portfolio strategies computed under the “Baseline” parametrization. In the columns under “Equilibrium”, both price and portfolio equilibrium policies correspond to the parameters indicated in each column. In all cases, mean valuation is at its baseline level $\delta_i = -0.1468$. We fix costs at \$1.85 per 20-cigarette pack in all scenarios. We compute the average smoking rates assuming that the available market size represents 35.6 % of the population, which was the average smoking rate in Uruguay in 2000. Figure C.5 shows consumption, elasticity, switching behavior, prices, virtual costs, and markups for intermediate levels of brand loyalty.


As a result, virtual costs rise as brand loyalty falls. The increase in virtual costs, in turn, shifts up the policy, i.e., the prices firms charge for any level of their customer base. The joint effect of virtual costs and static markups is to rotate the price policy clockwise and shift it up. According to our estimates, the long-run average virtual cost increases by a factor of four when we eliminate brand loyalty— see Table 7. The offsetting effect of decreasing markups and raising virtual costs leads to a small, non-monotonic effect on prices and reconciles the possibility of large increases in demand elasticity with drops in consumption.³¹

6.2.2 Portfolio Adjustments

Next, we evaluate how portfolio adjustments influence uniform packaging’s overall impact. Figure 8 compares its effect on long-term consumption when firms keep their portfolio decisions fixed

³¹The non-monotonic pattern between brand loyalty, prices, and consumption results from the initial market concentration, where a single product captures more than 50% of the market. In this case, initial drops in brand loyalty make the market leader’s demand substantially more elastic without affecting the probability a customer returns as much, which leads to significant price drops. However, as brand loyalty decreases, the market leader loses its dominant position, and the investing incentives dominate. This pattern is consistent with several recent papers exploring the relationship between prices and switching costs [Dubé et al., 2009, Arie and E. Grieco, 2014, Fabra and García, 2015, Cabral, 2016].

Figure 7: Price Decomposition - Monte Paz, Flagship




(a) Effect on Static Markups

(b) Effect on Virtual Costs

Note: This figure represents the policies when marginal costs equal 1.8 US\$. The shaded region indicates all the possible values the policy might take for a given size of the product's customer base. The policies are constructed by solving the model at the average parameters of the importance sampling distribution under the indicated levels of inertia. Dashed vertical lines represent the average customer base for each level of inertia. illustrates the effect of eliminating brand loyalty from the baseline estimates (holding mean valuations fixed) for the national firm' flagship product

and when they optimally adjust them. We find that the portfolio adjustments have a significant influence on consumption. Specifically, the expected number of available products increases as brand loyalty drops. If uniform packaging eliminates loyalty without affecting cigarette valuation, firms will offer, on average, 7.5 product lines, which is significantly larger than the 6.5 product lines we observe in the baseline scenario. Thus, for intermediate levels of brand loyalty, which coincide with the most significant decrease in prices, consumption might increase up to 2.9%. Furthermore, in this case, the mean valuation for cigarettes needs to drop around 35% to ensure consumption decreases for any level of brand loyalty.

Figure 8: Percentage change in long-term average consumption under uniform packaging.



(a) Fixed Participation Policies

(b) Full Equilibrium

Note: For each counterfactual scenario, we take demand and supply primitives, solve equilibrium policies according to the definition in Section I.2, and simulate the industry for 150 periods 2,500. This figure compares and reports the percentage change in average consumption over the visited states in the long-run station distribution (after burning the first 100 periods). All comparisons are in relation to the baseline scenario where the addiction parameter is $\eta_0 = 2.007$, and the brand loyalty is $\eta_1 = 3.437$, which is represented in the top-right corner by a black dot. We fix costs at \$1.85 per 20-cigarette pack in all scenarios. Green colors indicate that consumption decreases from our average consumption in our baseline scenario, while red indicates that consumption increases.

Two factors establish the negative correlation between brand loyalty and the number of available products. First, demand becomes more symmetric across products. If we assume all products were offered in the market, eliminating brand loyalty would decrease the HHI by around 650 points –see Table 7. Thus, relatively disadvantaged products can generate more profits at the expense of the market leaders. To the extent that the current market structure is very asymmetric, this implies that multiple smaller products can reach a sustainable customer base while the market leaders suffer “tolerable” losses. New, small product lines are offered without sacrificing the presence of the established products. Second, the endogenous entry costs decline. Endogenous entry costs are the initial investment necessary to gain a sustainable customer base. We measure it by computing the difference between a product’s value when it has no loyal customer base versus its value with its steady-state level of customers. The gap between initial and steady state values goes from around 10% in our baseline scenario to 5% when there is no brand loyalty. In other words, firms are more willing to introduce new products when they can free-ride on their rivals’ efforts to attract customers.

Overall, our results show that uniform packaging can positively affect cigarette consumption but can also backfire. On the one hand, it would decrease the value of customers for firms, discouraging them from attracting new consumers to the market and lowering consumption. On the other hand, it might facilitate the introduction of new varieties as demand becomes more symmetric and endogenous entry costs decrease. These results are consistent with studies that analyze the effects of uniform packaging after its implementation in a handful of countries. Although the evidence is anecdotal and the analysis descriptive, they have found a small impact on prices and a more significant increase in the number of available varieties, without a large effect on aggregate consumption [Breton et al., 2018, Underwood et al., 2020, Moodie et al., 2022].

Robustness Uniform packaging could also reduce product differentiation beyond its effect on brand loyalty. If that were the case, the policy would induce more price competition without discouraging firms from attracting new consumers. Although it could still prevent them from investing in new consumers by decreasing industry profits, it would not do so directly through the probability of retaining customers. Hence, in this case, the investing incentives would not decrease as much, potentially leading to more price competition and a more significant increase in consumption. In Figure C.6, We show that if uniform packaging eliminates all vertical differentiation between products without affecting brand loyalty, consumption might increase up to 8%.

7 Conclusions

The theoretical literature posits that firms are strategically motivated to adapt their pricing and product strategies to foster consumers' habit formation and maximize long-term profitability. Our empirical findings strongly support these theories, demonstrating that firms indeed adjust their behavior to exploit consumer inertia. Moreover, we show that relatively simple adjustments to the traditional workhorse models in economics are sufficient to take these models to the data, which allows us to understand better the effect of policies in markets with habit formation. In particular, in the case of sin goods, we stress that policymakers should leverage these incentives to design policies that discourage consumption through companies' responses. This strategy has the advantage of decreasing their use without shifting the burden to consumers as much.

References

- Rossi Abi-Rafeh, Pierre Dubois, Rachel Griffith, and Martin O'Connell. The effects of sin taxes and advertising restrictions in a dynamic equilibrium. Technical report, Technical Report 1480, TSE Working Paper 1480, 2023.
- Daniel A Ackerberg. A new use of importance sampling to reduce computational burden in simulation estimation. *QME*, 7:343–376, 2009.
- Victor Aguirregabiria, Allan Collard-Wexler, and Stephen P Ryan. Dynamic games in empirical industrial organization. In *Handbook of Industrial Organization - Volume 4.*, pages 225–343. Elsevier, 2021.
- Margaret Aksoy-Pierson, Gad Allon, and Awi Federgruen. Price competition under mixed multinomial logit demand functions. *Management Science*, 59(8):1817–1835, 2013.
- Benjamin J Apelberg, Shari P Feirman, Esther Salazar, Catherine G Corey, Bridget K Ambrose, Antonio Paredes, Elise Richman, Stephen J Verzi, Eric D Vugrin, Nancy S Brodsky, et al. Potential public health effects of reducing nicotine levels in cigarettes in the united states. *New England Journal of Medicine*, 378(18):1725–1733, 2018.
- Peter Arcidiacono, Holger Sieg, and Frank Sloan. Living rationally under the volcano? an empirical analysis of heavy drinking and smoking. *International Economic Review*, 48(1):37–65, 2007.
- Guy Arie and Paul L E. Grieco. Who pays for switching costs? *Quantitative Marketing and Economics*, 12:379–419, 2014.

- Joe S Bain. Advantages of the large firm: production, distribution, and sales promotion. *Journal of marketing*, 20(4):336–346, 1956.
- Patrick Bajari, C Lanier Benkard, and Jonathan Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007.
- Nano Barahona, Cristóbal Otero, Sebastián Otero, and Joshua Kim. Equilibrium effects of food labeling policies. *Available at SSRN 3698473*, 2020.
- Gary S Becker and Kevin M Murphy. A theory of rational addiction. *Journal of political Economy*, 96(4):675–700, 1988.
- Alan Beggs and Paul Klemperer. Multi-period competition with switching costs. *Econometrica: Journal of the Econometric Society*, pages 651–666, 1992.
- C Lanier Benkard. A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies*, 71(3):581–611, 2004.
- C Lanier Benkard, Przemyslaw Jeziorski, and Gabriel Y Weintraub. Oblivious equilibrium for concentrated industries. *The RAND Journal of Economics*, 46(4):671–708, 2015.
- Micah L Berman and Allison M Glasser. Nicotine reduction in cigarettes: literature review and gap analysis. *Nicotine and Tobacco Research*, 21(Supplement_1):S133–S144, 2019.
- Micah L Berman, Patricia J Zettler, and David L Ashley. Anticipating industry arguments: The us food and drug administration’s authority to reduce nicotine levels in cigarettes. *Public Health Reports*, 133(4):502–506, 2018.
- Steve Berry and Ariel Pakes. Estimation from the optimality conditions for dynamic controls. *Manuscript, Yale University*, pages 841–890, 2000.
- David Besanko, Ulrich Doraszelski, Yaroslav Kryukov, and Mark Satterthwaite. Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica*, 78(2):453–508, 2010.
- David Besanko, Ulrich Doraszelski, and Yaroslav Kryukov. The economics of predation: What drives pricing when there is learning-by-doing? *American Economic Review*, 104(3):868–97, 2014.
- David Besanko, Ulrich Doraszelski, and Yaroslav Kryukov. How efficient is dynamic competition? the case of price as investment. *American Economic Review*, 109(9):3339–64, 2019.
- Magdalena Opazo Breton, John Britton, Yue Huang, and Ilze Bogdanovica. Cigarette brand diversity and price changes during the implementation of plain packaging in the united kingdom. *Addiction*, 113(10):1883–1894, 2018.

- Luis Cabral. Dynamic pricing in customer markets with switching costs. *Review of Economic Dynamics*, 20:43–62, 2016.
- Luis MB Cabral and Michael H Riordan. The learning curve, market dominance, and predatory pricing. *Econometrica: Journal of the Econometric Society*, pages 1115–1140, 1994.
- Andrew Caplin and Barry Nalebuff. Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica: Journal of the Econometric Society*, pages 25–59, 1991.
- Frank J Chaloupka and Kenneth E Warner. The economics of smoking. *Handbook of health economics*, 1:1539–1627, 2000.
- Cyril Chantler. Standardised packaging of tobacco. Technical report, Secretary of State for Health, 2014.
- Jiawei Chen. How do switching costs affect market concentration and prices in network industries? *The Journal of Industrial Economics*, 64(2):226–254, 2016.
- Federico Ciliberto and Nicolai V Kuminoff. Public policy and market competition: how the master settlement agreement changed the cigarette industry. *The BE Journal of Economic Analysis & Policy*, 10(1), 2010.
- Christopher Conlon and Jeff Gortmaker. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.
- Christopher Conlon, Nirupama Rao, and Yinan Wang. Who pays sin taxes? understanding the overlapping burdens of corrective taxes. *Review of Economics and Statistics*, pages 1–27, 2022.
- John Dawes. Cigarette brand loyalty and purchase patterns: an examination using us consumer panel data. *Journal of Business Research*, 67(9):1933–1943, 2014.
- Carl De Boor and Carl De Boor. *A practical guide to splines*, volume 27. springer-verlag New York, 1978.
- Teresa DeAtley, Eduardo Bianco, Kevin Welding, and Joanna E Cohen. Compliance with uruguay’s single presentation requirement. *Tobacco control*, 27(2):220–224, 2018.
- Ulrich Doraszelski and Mark Satterthwaite. Computable markov-perfect industry dynamics. *The RAND Journal of Economics*, 41(2):215–243, 2010.
- Michaela Draganska, Michael Mazzeo, and Katja Seim. Beyond plain vanilla: Modeling joint product assortment and pricing decisions. *QME*, 7(2):105–146, 2009.

- Jean-Pierre Dubé, Günter J Hitsch, and Peter E Rossi. Do switching costs make markets less competitive? *Journal of Marketing research*, 46(4):435–445, 2009.
- Jean-Pierre Dubé, Günter J Hitsch, and Peter E Rossi. State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445, 2010.
- Jonathan T Elliott. Investment, emissions, and reliability in electricity markets. *Working Paper*, 2022.
- Juan F Escobar. Equilibrium analysis of dynamic models of imperfect competition. *International Journal of Industrial Organization*, 31(1):92–101, 2013.
- Natalia Fabra and Alfredo García. Market structure and the competitive effects of switching costs. *Economics Letters*, 126:150–155, 2015.
- Ying Fan and Chenyu Yang. Competition, product proliferation, and welfare: A study of the us smartphone market. *American Economic Journal: Microeconomics*, 12(2):99–134, 2020.
- Joseph Farrell and Michael L Katz. Competition or predation? consumer coordination, strategic pricing and price floors in network markets. *The Journal of Industrial Economics*, 53(2):203–231, 2005.
- Joseph Farrell and Paul Klemperer. Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization*, 3:1967–2072, 2007.
- Joseph Farrell and Carl Shapiro. Dynamic competition with switching costs. *The RAND Journal of Economics*, pages 123–137, 1988.
- FDA. Tobacco product standard for nicotine level of combusted cigarettes. Technical report, FDA, 2018.
- FDA. Fda announces plans for proposed rule to reduce addictiveness of cigarettes and other combusted tobacco products, <https://www.fda.gov/news-events/press-announcements/fda-announces-plans-proposed-rule-reduce-addictiveness-cigarettes-and-other-combusted-tobacco>. Technical report, FDA, 2022.
- Chaim Fershtman and Ariel Pakes. Dynamic games with asymmetric information: A framework for empirical work. *The Quarterly Journal of Economics*, 127(4):1611–1661, 2012.
- Sebastian Fleitas. Dynamic competition and price regulation when consumers have inertia: Evidence from medicare part d. *Unpublished mimeo, University of Arizona*, 2, 2017.

- Chiara Fumagalli and Massimo Motta. A simple theory of predation. *The Journal of Law and Economics*, 56(3):595–631, 2013.
- Jean Gabszewicz, Lynne Pepall, and Jacques-Francois Thisse. Sequential entry with brand loyalty caused by consumer learning-by-using. *The Journal of Industrial Economics*, pages 397–416, 1992.
- Austan Goolsbee and Amil Petrin. The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica*, 72(2):351–381, 2004.
- Jonathan Gruber and Botond Köszegi. Is addiction “rational”? theory and evidence. *The Quarterly Journal of Economics*, 116(4):1261–1303, 2001.
- Benjamin R Handel. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82, 2013.
- Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the econometric society*, pages 1029–1054, 1982.
- James J Heckman. Statistical models for discrete panel data. *Structural analysis of discrete data with econometric applications*, 114:178, 1981.
- Ali Hortaçsu, Aniko Oery, and Kevin R Williams. Dynamic price competition: Theory and evidence from airline markets. Technical report, National Bureau of Economic Research, 2022.
- Bar Ifrach and Gabriel Y Weintraub. A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies*, 84(3):1106–1150, 2017.
- Mitsuru Igami. Estimating the innovator’s dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy*, 125(3):798–847, 2017.
- Mitsuru Igami and Kosuke Uetake. Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies*, 87(6):2672–2702, 2020.
- Kenneth L Judd. *Numerical methods in economics*. MIT press, 1998.
- Paul Klemperer. The competitiveness of markets with switching costs. *The RAND Journal of Economics*, pages 138–150, 1987a.
- Paul Klemperer. Entry deterrence in markets with consumer switching costs. *The Economic Journal*, 97:99–117, 1987b.

- Murray Laugesen and Randolph C Grace. Excise, electronic cigarettes and nicotine reduction to reduce smoking prevalence in new zealand by 2025. *The New Zealand Medical Journal (Online)*, 128(1420):72, 2015.
- David T Levy, Frank Chaloupka, Eric N Lindblom, David T Sweanor, Richard J O’connor, Ce Shang, and Ron Borland. The us cigarette industry: an economic and marketing perspective. *Tobacco regulatory science*, 5(2):156–168, 2019.
- Hong Liu, John A Rizzo, Qi Sun, and Fang Wu. How do smokers respond to cigarette taxes? evidence from china’s cigarette industry. *Health economics*, 24(10):1314–1330, 2015.
- Alexander MacKay and Marc Remer. Consumer inertia and market power. *Available at SSRN 3380390*, 2021.
- Eric Maskin and Jean Tirole. A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica: Journal of the Econometric Society*, pages 549–569, 1988.
- Crawford Moodie, Janet Hoek, David Hammond, Karine Gallopel-Morvan, Diego Sendoya, Laura Rosen, Burcu Mucan Özcan, and Yvette van der Eijk. Plain tobacco packaging: progress, challenges, learning and opportunities. *Tobacco Control*, 31(2):263–271, 2022.
- W Ross Morrow and Steven J Skerlos. Fixed-point approaches to computing bertrand-nash equilibrium prices under mixed-logit demand. *Operations research*, 59(2):328–345, 2011.
- Ariel Pakes, Michael Ostrovsky, and Steven Berry. Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *the RAND Journal of Economics*, 38(2):373–399, 2007.
- Ariel Pakes, Jack R Porter, Mark Shepard, and Sophie Calder-Wang. Unobserved heterogeneity, state dependence, and health plan choices. Technical report, National Bureau of Economic Research, 2021.
- Polykarpos Pavlidis and Paul B Ellickson. Implications of parent brand inertia for multiproduct pricing. *Quantitative Marketing and Economics*, 15:369–407, 2017.
- Martin Pesendorfer and Philipp Schmidt-Dengler. Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies*, 75(3):901–928, 2008.
- Shi Qi. The impact of advertising regulation on industry: The cigarette advertising ban of 1971. *The RAND Journal of Economics*, 44(2):215–248, 2013.
- Mar Reguant and Francisco Pareschi. Bounding outcomes in counterfactual analysis. Technical report, working paper, 2021.

- Oleksandr Shcherbakov. Measuring consumer switching costs in the television industry. *The RAND Journal of Economics*, 47(2):366–393, 2016.
- Anna E Tuchman. Advertising and demand for addictive goods: The effects of e-cigarette advertising. *Marketing science*, 38(6):994–1022, 2019.
- David Underwood, Sizhong Sun, and Riccardo AMHM Welters. The effectiveness of plain packaging in discouraging tobacco consumption in australia. *Nature Human Behaviour*, 4(12):1273–1284, 2020.
- Gabriel Y Weintraub, C Lanier Benkard, and Benjamin Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.
- WHO. Tobacco plain packaging: global status update. Technical report, World Health Organization, 2022.
- Thomas G Wollmann. Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406, 2018.

A Appendix: Market Description

A.1 Data Summary Statistics

Table A.1: Data Summary Statistics

	All	Final Sample		All	Final Sample
Markets	71	38	Individuals	2208	1305
Stores	3601	93	Observations	4978	2850
Observations	130880	12422			

A.2 Market Segmentation

Our supply-side model cannot accommodate many products within a firm portfolio. Next, we show that there are well-defined market segments over which consumers have similar preferences, firms price uniformly, and such that all products within it are introduced/retired at the same time. This allows us to reduce the number of products in the choice set without losing much information. Next, we describe what are the main cigarettes segments from the perspective of consumers in the Uruguayan cigarette market. Then we show that firms set almost identical prices for all products within a segment. In turn, we argue that following product or segment shares provide almost as much information. Finally, we show that the entry and exit of products coincide with the evolution of these market segments.

A.2.1 Vertical and Horizontal Differentiation

There are eight clearly differentiated product segments. First, we differentiate between four types of cigarettes: flagship products (or leader products), other products with normal levels of tar, light cigarettes (low in tar), and products with special characteristics such as slightly longer than normal, slimmer, no filter, etc. Flagship products are the market best-selling product of each firm and are generally associated with brands that have a long tradition in the Uruguayan market, such as Nevada, Coronado (Monte Paz), Fiesta and Marlboro (Philip Morris), Pall Mall (BAT). Light products might share brand names with the flagship products but are low in tar and generally less popular. Finally, other regular products are “common” cigarettes, with similar levels of tar as flagship products, but that do not belong to a leader brand.

Within each of these categories, there are two types of vertical “qualities”: standard and premium. The low-cost segment, which is common in other countries, did not develop in Uruguay. Table A.2 shows market shares across these dimensions and firms. The premium segment is small, and


only BAT and Philip Morris sell premium cigarettes. However, these segments are not completely isolated markets. There exists non-trivial substitution between premium-light categories.

Table A.2: Market Share by Segment and Firm.


	Leader			Other Regular			Light			Specials		
	MP	BAT	PM	MP	BAT	PM	MP	BAT	PM	MP	BAT	PM
Standard	0.559	0.050	0.133	0.024	-	0.011	0.073	-	0.022	0.002	-	-
Premium	-	0.018	0.054	-	-	-	-	0.009	0.035	-	-	-

A.2.2 Price Evolution by Segment

Firms price products within each one of these categories uniformly. In some cases, prices are identical across products within a segment. For instance, the flagship and light products within the standard segment (see Monte Paz’ Nevada Filtro and Nevada Blanco-California prices). Figure A.1a shows prices for premium products. We see that all products within a firm-premium category have similar prices (it is even hard to distinguish different lines). The evolution of prices in the standard quality category is slightly more complicated. In particular, there are two price levels in this segment. The most common price is the higher one. So there is not exactly one price for medium products. However, the bulk of the segment shares are determined by one price per firm.



(a) High Segment Prices.



(b) Medium Segment Prices.

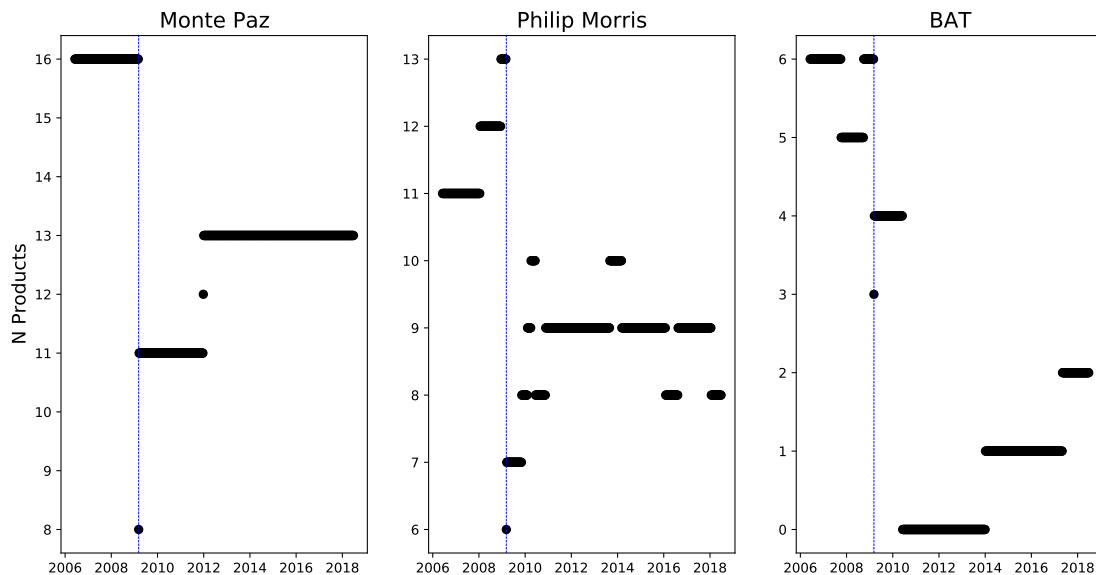
Overall, this evidence suggests that firms set prices uniformly across products within segments and sometimes across segments. Next, we show that if firms set uniform prices for all products within a segment and consumers value these products, then tracking individual-level products or segment aggregates should generate similar strategies and equilibrium outcomes.

A.2.3 Product Entry & Exit

We have omitted product entry and exit so far. However, it is evident that to reduce the number of products, we need all products within a segment to be introduced and retired at the same time. Next, we describe the entry and exit patterns of products within a segment.

Monte Paz had a presence in four segments before the one-presentation-per-brand policy: flagship, light, other regular, and special products. After the policy, all light and special products were eliminated. They reintroduced the light products right away and waited two years to re-introduce the special products. They reintroduced these segments by launching one fewer product in each category (so, in each category, one product was not “substituted”). Hence, while Monte Paz product portfolio did not reach the exact same size as before, its shares within each segment were similar (see Figure A.2). Philip Morris suffered a large shock as a consequence of the policy. They lost all their light and retired other regular products shortly after. While they re-introduced an almost identical portfolio of light products after 10 months, they never reintroduced the other regular products. BAT did not face a significant shock as a consequence of the policy since its light segment was small. Indeed, after the policy, it reintroduced one product to recompose its light segment. However, in 2010 it took all its products off the market and re-entered in 2014 with just its premium products (two products) in 2014.

Figure A.2: Evolution of firms’ product portfolio.



A.2.4 Considered Segments

Therefore, we use a four-level criterion to define each segment: similar consumer preferences, priced uniformly, being introduced or retired together within a short time, and similar marginal costs. The natural partition following these criteria would be to have eight product segments by firm. However, we simplify the problem in the following way. We only consider segments that had been observed sometime in our sample. This implies, for instance, that Monte Paz cannot have premium products or that Philip Morris and BAT do not have special products. Then, we assume that BAT segments only differentiate between the premium and standard dimensions since the light products within each vertical dimension represent a negligible share. On the other hand, for Philip Morris, we do not consider the premium-standard differentiation because they price them similarly, just setting a price gap between them that is constant throughout time. Moreover, Philip Morris introduced premium and non-premium products in the light segment simultaneously, following the one-presentation-per-brand. Thus, we define nine segments (4 Monte Paz, 3 Philip Morris, 2 BAT). Table A.3 shows summary statistics by segment.

Table A.3: Summary Statistics by Segment

Segment	Number of Products	Average Price	Shares	Time in Sample
Monte Paz Leader	2.0	2.3	0.595	1.0
Monte Paz Other Regular	4.0	2.15	0.022	1.0
Monte Paz Light	4.5	2.3	0.086	1.0
Monte Paz Special	3.5	2.21	0.001	0.69
Philip Morris Leader	4.0	2.24	0.233	1.0
Philip Morris Light	3.5	2.25	0.049	0.94
Philip Morris Other Regular	4.0	1.9	0.006	0.17
BAT Leader	1.0	1.92	0.049	0.29
BAT Premium	3.5	2.08	0.004	0.63

A.3 State Dependence Reduced Form Test

Here we apply Shcherbakov [2016]’s reduced form strategy to get a first sense of which force dominates. We use our aggregate data source and regress current shares on lagged shares of the same product, average shares of all other products in the choice set, and prices.

$$s_{ijt} = \alpha_{0pj} + \beta_0 s_{ij,t-1} + \beta_1 s_{i,(-j),t-1} + \gamma_t + \gamma_j + \varepsilon_{ijt}$$

Table A.4: Structural State Dependence versus Spurious Dependence

	(1) OLS	(2) IV
own lagged share	0.886*** (0.003)	0.676*** (0.034)
lagged mean shares others	0.010*** (0.003)	-0.348*** (0.069)
real price per cig	-0.011*** (0.001)	-0.015*** (0.001)
N	188927	188927
R ²	0.929	0.892
Prod, Time FE	Yes	Yes
IVs	None	Lagged Prices

Although the OLS estimate of β_0 determines the correlation between current and lagged choices, yesterday's decisions are influenced by factors that are permanent throughout time but not observed by the econometrician. These unobserved factors make lagged choices endogenous to the error term. We use exogenous shifters of lagged choices as instruments to isolate the effect of past decisions on current ones from persistent factors.³² Table A.4 shows the results, indicating that state dependence is a relevant force in our setting.

A.4 Other Tables

Table A.5: Extensive and Intensive Margin of Consumption

	EM1	EM2	EM3	EM4	EM5	IM1	IM2	IM3	IM4	IM5
Previous Smoker	0.641 (0.024)	0.638 (0.024)	0.599 (0.027)	- -	0.303 (0.061)	3.907 (1.358)	3.586 (1.324)	-7.018 (1.018)	- -	-38.296 (1.729)
Quantity Smoked	- -	- -	0.004 (0.001)	0.003 (0.001)	- -	- -	- -	0.670 (0.017)	0.619 (0.017)	- -
log(Addiction Stock)	- -	- -	- -	- -	0.061 (0.009)	- -	- -	- -	- -	7.291 (0.233)
R ²	0.252	0.256	0.306	0.315	0.312	0.004	0.058	0.486	0.486	0.377
N	2770.0	2770.0	2544.0	2544.0	2544.0	1944.0	1944.0	1873.0	1873.0	1873.0
Product FE	No	No	No	Yes	No	No	No	No	Yes	No
Demos FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes

Note:

³²Recall cigarette prices are uniform across stores. Hence, prices are exogenous to store-time specific shocks.

A.5 Other Figures

Figure A.3: Aggregate National Sales and Sample Aggregate Sales




Figure A.4: Plain Packages



Figure A.5: Philip Morris' portfolio and products affected by OPP



B Appendix: Additional Modeling Details

B.1 Dynamic model for multi-product firms

The only relevant difference between the dynamic model for multi-product and single-product firms is how they choose product portfolios, as extending pricing decisions to the multi-product case is straightforward.

In the single firm case, optimal participation decisions are characterized by a threshold rule. Thus, participation probabilities characterize the entire participation strategy without loss of information. On the other hand, if multi-product firms observe the realization of fixed costs for all products simultaneously, they must consider the probability of each possible bundle, which is not a threshold strategy rule. In principle, we could solve for each bundle's equilibrium probabilities as in Draganska et al. [2009]

Let \mathfrak{J}_f be one possible bundle offered by firm f at t . Then, taking into consideration that shocks to fixed costs are private information of the firms, the probability of offering this bundle –omitting all other states– is given by

$$Pr(\mathfrak{J}_f) = \int 1 \left\{ \mathfrak{J}_f = \operatorname{argmax}_J \left(\beta \int_{\mathfrak{J}_{-f}} V(J, \mathfrak{J}_{-f}) dF(\mathfrak{J}_{-f} | \sigma_{-f}^\phi) - \sum_{k:J_k=1} \Theta_k \right) \right\}$$

The indicator function $1 \left\{ \mathfrak{J}_f = \operatorname{argmax}_J \left(\beta \int_{\mathfrak{J}_{-f}} V(J, \mathfrak{J}_{-f}) dF(\mathfrak{J}_{-f} | \sigma_{-f}^\phi) - \sum_{k:J_k=1} \Theta_k \right) \right\}$ defines an area of integration $A^{\mathfrak{J}_f}$. This area is the region of realizations of $\Theta_f \in \mathbb{R}^K$ over which offering \mathfrak{J}_f is optimal. Moreover, we can define product-specific areas of integration for each bundle. That is, the area of integration under which product $k(f)$ is offered under state is $A_{k(f)} = \{\Theta_f \in \cup_{\{\mathfrak{J}_f: \mathfrak{J}_{fk}=1\}} A_k^{\mathfrak{J}_f}\}$. Hence, we can re-express $\chi_{k(f)}(\Theta_f) = 1\{\Theta_f \in A_{k(f)}\}$. Note that $\chi_{k(f)}(\Theta_f)$ depends on the entire vector of fixed cost shocks of firm f but not on shocks of other firms since these are private still information.

Then, the value of firm f before participation choices are made is

$$U_f(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}; \sigma) = -E \left[\sum \Theta_{k(f)} 1\{\Theta_f \in A_{k(f)}(S_t, \mathbf{J}_t, c_{t+1}, \delta_{t+1}, \sigma_{-f}^\phi)\} \right] + \beta E[V(S_t, \mathfrak{J}, c_{t+1}, \delta_{t+1}) | \sigma^\phi]$$

However, working with the derivatives of such expectations with respect to prices would be complex. In particular, it would make the solution to the pricing problem significantly more involved.

We simplify the problem by assuming fixed cost shocks are also private information within the firm. That is, the departments in charge of deciding on the participation of different products do not need to communicate with each other. Nevertheless, they optimize the company's operations, considering the profits of the whole firm. Hence, we still account for possible cannibalization effects.

C Robustness Checks and Additional Results

C.1 Model Fit and Robustness

Table C.1: Predicted and Observed Switching

	Outside	MP Flag.	MP Light	MP Other	PM Flag.	PM Prem	PM Light	PM Light Prem	PM Other	BAT	BAT Prem
Outside	0.78, 0.78	0.17, 0.09	0.01, 0.03	0.0, 0.01	0.01, 0.04	0.02, 0.03	0.01, 0.01	0.01, 0.01	nan, nan	nan, nan	0.0, 0.01
MP Flag	0.18, 0.13	0.74, 0.79	0.04, 0.02	0.0, 0.01	0.01, 0.02	0.02, 0.02	0.0, 0.01	0.0, 0.01	nan, 0.0	0.01, 0.03	nan, 0.01
MP Light	0.16, 0.26	0.28, 0.11	0.5, 0.5	0.04, 0.02	0.03, 0.05	0.03, 0.03	0.03, 0.01	0.01, 0.01	nan, 0.01	nan, 0.06	0.02, 0.01
MP Other	0.25, 0.35	0.11, 0.16	0.19, 0.05	0.45, 0.3	0.06, 0.07	0.06, 0.04	0.18, 0.02	nan, 0.01	nan, 0.01	0.23, 0.07	nan, 0.02
PM Flag.	0.18, 0.23	0.07, 0.1	0.03, 0.03	nan, 0.02	0.69, 0.57	0.04, 0.03	nan, 0.01	0.02, 0.01	nan, 0.0	nan, 0.05	nan, 0.01
PM Prem	0.2, 0.27	0.11, 0.12	0.09, 0.04	nan, 0.02	0.04, 0.05	0.61, 0.48	nan, 0.01	0.04, 0.01	nan, 0.01	nan, 0.05	0.02, 0.01
PM Light	0.31, 0.42	0.23, 0.17	nan, 0.05	nan, 0.02	0.23, 0.07	nan, 0.05	0.36, 0.26	0.09, 0.01	0.2, 0.01	0.2, 0.08	0.11, 0.02
PM Light Prem	0.26, 0.41	0.26, 0.19	nan, 0.06	0.07, 0.02	0.21, 0.07	0.17, 0.06	0.1, 0.02	0.45, 0.15	nan, 0.01	nan, 0.08	nan, 0.02
PM Other	0.38, 0.48	0.5, 0.19	nan, 0.06	0.25, 0.03	0.5, 0.09	nan, 0.05	nan, 0.02	nan, 0.01	nan, 0.09	nan, 0.11	nan, 0.01
BAT	0.1, 0.39	0.2, 0.14	0.05, 0.04	0.11, 0.02	0.2, 0.06	0.05, 0.05	0.16, 0.02	0.05, 0.01	nan, 0.01	0.57, 0.59	nan, 0.01
BAT Prem	0.75, 0.4	nan, 0.18	nan, 0.06	nan, 0.03	nan, 0.07	nan, 0.07	nan, 0.02	nan, 0.02	nan, 0.01	nan, 0.08	0.75, 0.19

Note: Switching rates are both calculated at eight quarter periods, which represents the frequency at which we observe individuals.

Figure C.1: Compare Observed and Simulated Participation Probabilities.

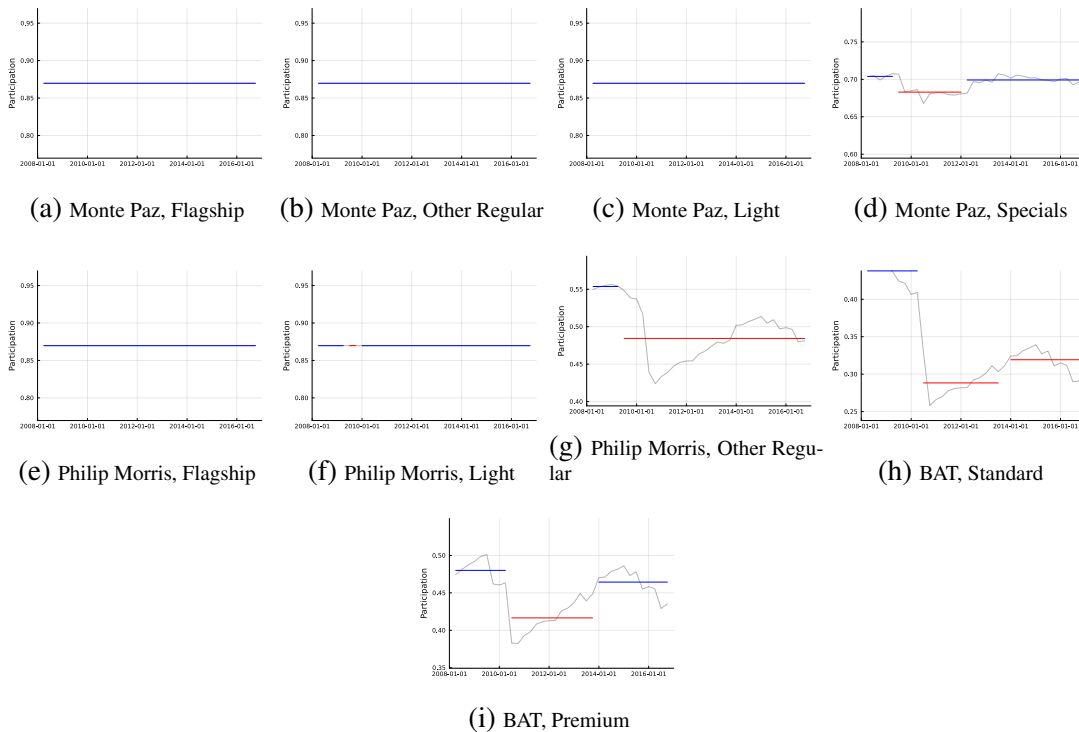
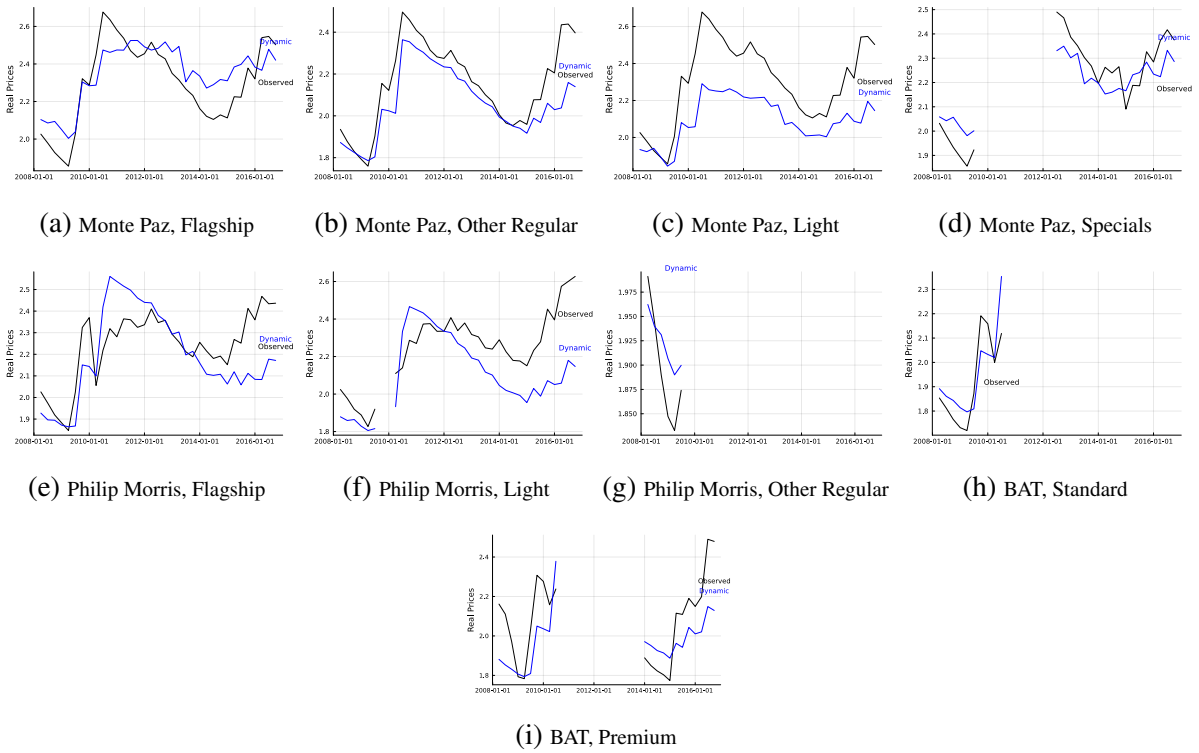


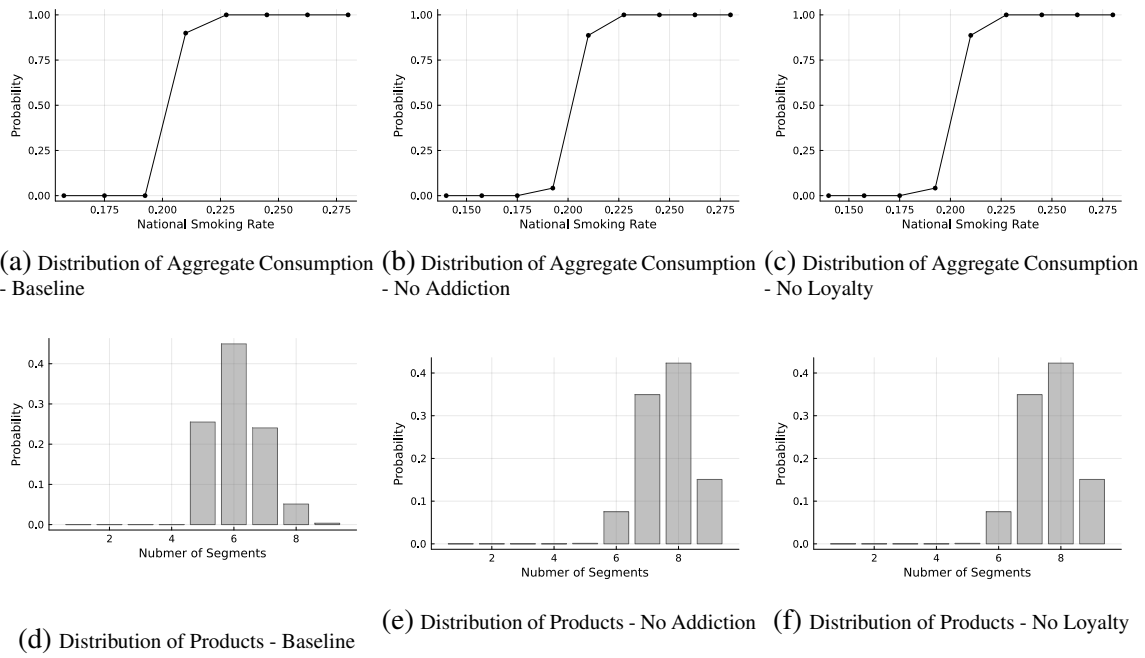
Figure C.2: Compare Observed and Simulated Prices.



C.2 Additional Results

C.2.1 Stationary Distribution of Consumption and Product Portfolios


Figure C.3: Stationary Distribution of Consumption and Product Portfolios



Note: The smoking rates are assuming the total market size is 35.6%, according to the smoking rate in 2000. Moreover, we assume there is no heterogeneity in quantities smoked across different consumers.

C.2.2 Addiction

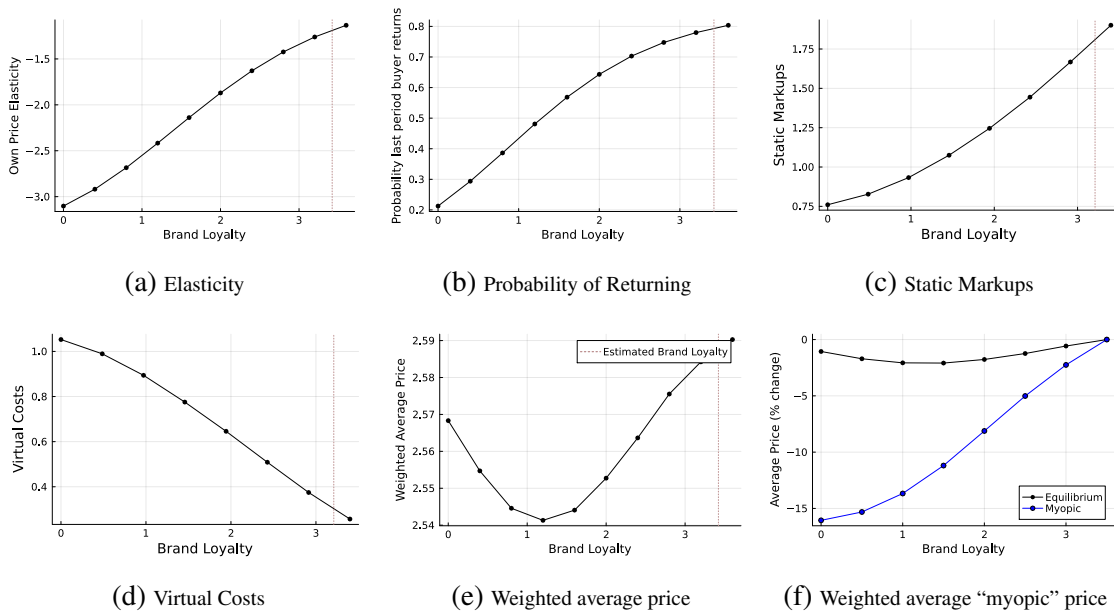
Figure C.4: Equilibrium Outcomes under Alternative Levels of Addiction



Note: For each level of inertia, we solve for equilibrium policies for prices and market participation. Then, we simulate the industry evolution over 50 periods, replicated 2,500 times from a null industry scenario. The depicted results are the averages across time and simulations. For the results where we fix firms' strategies, we solve the equilibrium at the baseline level of inertia and then simulate the economy modifying consumer preferences. The average smoking rates are computed assuming that the available market size represents 35.6 % of the population, which was the average smoking rate in Uruguay in 2000. The vertical dotted brown line references estimated levels of inertia.

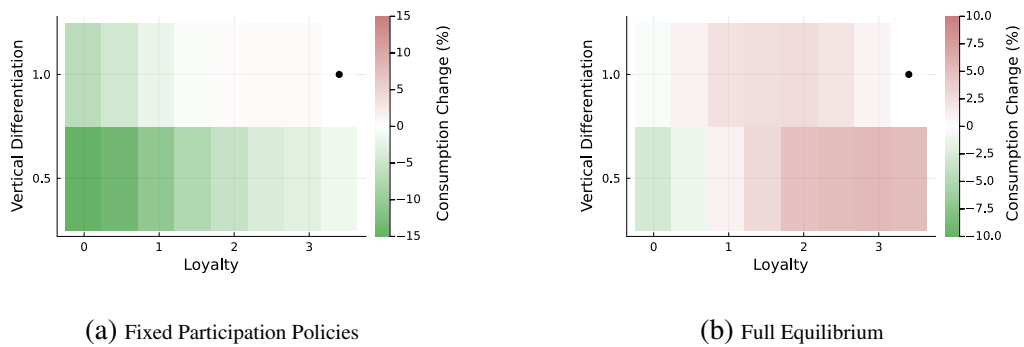
C.2.3 Brand Loyalty

Figure C.5: Summary Statistics for Intermediate Levels of Brand Loyalty



Note: This figure represents the policies when marginal costs equal 1.8 US\$. The shaded region indicates all the possible values the policy might take for a given size of the product's customer base. The policies are constructed by solving the model at the average parameters of the importance sampling distribution under the indicated levels of inertia. Dashed vertical lines represent the average customer base for each level of inertia.

Figure C.6: Vertical Differentiation



Online Appendix

Curbing Habit Formation: The Effects of Tobacco Control Policies in a Dynamic
Equilibrium

Francisco Pareschi and Gastón López

May 1, 2024

I Appendix: Computational Implementation

I.1 Equilibrium Computation Algorithm

We compute the equilibrium iterating in approximate best responses. We approximate the value function using Chebyshev Polynomials over Chebyshev nodes and solve for a fixed point in the space of the polynomials coefficients. In each step, we perform three tasks. First, we compute the best participation and price response given the trial value function approximation, rivals policies and firms beliefs. Then, we update firms beliefs according to new policies. Finally, we update values at the selected nodes and update the approximation, i.e., the polynomial coefficients.

I.1.1 Approximation

There are three type of states. The choice sets, which are discrete. The exogenous variables: preferences and costs. And the customer base (lagged market shares), which is continuous. Allowing for exogenous states (c_t, δ_t) is simple. We assume each of the exogenous states (c_t, δ_t) follow an independent AR(1) process, and discretize them using Tauchen's method. We call the resulting nodes $e \in \mathcal{E}$ and the transition matrix Π^e . On the other hand, the vector of own customer bases (S_{t-1}^f) and the tracked market moments (\bar{S}_{t-1}^f) aggregate is firm specific. We denote the whole vector z^f . To avoid solving the game for a large grid of possible values of z^f , we approximate the value function with a basis approximation such that for each possible choice set \mathfrak{J} and exogenous state e and firm f , we have

$$V_f(x, \mathfrak{J}, e; q) = \sum_{n=0}^N q_n^{f, \mathfrak{J}, e} T_n(x) \quad (13)$$

where each T_n is a Chebyshev basis of order n , and $\{q_n^{f, \mathfrak{J}, e}\}_{n \in \{1, \dots, N\}}$ are coefficients to be computed. Thus, we do not compute the fixed point of the value function but a fixed point of the coefficients $\{q_n^{f, \mathfrak{J}, e}\}_{n \in \{1, \dots, N\}, \mathfrak{J} \in \mathcal{J}}$, for each possible choice set.

At each iteration k , we solve for $\{q_n^{f, \mathfrak{J}, e}\}$. To solve for these coefficients, we must pick a set of nodes $\{z_i^f(\mathfrak{J}, e) \in \mathcal{S}^{f, \mathfrak{J}, e}\}$ for each firm f , choice set \mathfrak{J} and exogenous state e , where to solve the value function –and corresponding policies.³³ Choosing the nodes $\mathcal{S}^{f, \mathfrak{J}, e}$ appropriately is key to getting sound convergence properties of the approximation [Judd, 1998, De Boor and De Boor, 1978]. We pick Chebyshev nodes over each dimension and solve the problem at the intersection of every dimension.³⁴ Note that we still face a curse of dimensionality on the dimension of each

³³The i subscript indicate that it refers to particular nodes instead of the continuous vector z^f

³⁴There is an additional complexity in our setting: the state space is not represented by "rectangles" since shares live in the simplex. We explore different solutions to this issue. In our main specification we add more nodes into each

firm portfolio—not on the number of firms. We have solved the problem with up to five products per firm.

After evaluating the values at each node, $V_f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e)$, –see below– we find the optimal coefficients by OLS. That is we minimize the sum of the square distance of $\varepsilon_i^{f, \mathfrak{J}, e} = V_f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e) - V_f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; q)$, that is

$$\hat{q}_n^{f, \mathfrak{J}, e} = \arg \min_{q \in \mathbb{R}^N} \sum_i (\varepsilon_i^{f, \mathfrak{J}, e})^2$$

Finally, observe that taking derivatives of the value function (or any approximation) can be performed efficiently since it amounts to taking derivatives of a polynomial.

I.1.2 Updating Firms Beliefs

We enter each step k of the iteration holding $\hat{q}_{k-1}^{f, \mathfrak{J}, e}$, σ_{k-1}^p , σ_{k-1}^ϕ in memory. Here, σ^p , σ^ϕ refers also to prices and probabilities of participation for each product j at each node $z_i^f(\mathfrak{J}, e)$ for each choice set \mathfrak{J} and exogenous state e , evaluated by each firm f .³⁵

In the first step we solve for iteration k optimal prices σ_k^p , holding values and rivals policies constant. In order to solve these policies, we need to compute firm f expected profits and transitions for every possible vector of price p'_f , firm f might set.

Appendix I.2 describes how firms form beliefs about payoff relevant states from information sets I_f , i.e from (z_i^f, \mathfrak{J}, e) . Here, it suffices to say that for each realization of the information set I_f firms can compute rivals customer base $S_{-f, t-1}^f$, which also allow them for construct rivals information sets I_{-f}^f .³⁶ Then, firms can evaluate rivals' actual policies $\sigma_{k, -f}$ at I_{-f}^f to assess rivals expected actions. Naturally, I_{-f}^f is not an element of (z_i^f, \mathfrak{J}, e) . Hence, we need to interpolate the policy function of rivals $-f$ to evaluate it at I_{-f}^f . We do this using the same Chebyshev basis approximation as in the value function.

This information is enough to compute expected profits and transitions, because are both determined by the demand for each product j at each node $z_i^f(\mathfrak{J}, e)$. Integrating over the conditional distribution of types $w[n|k]$, conditional on lagged consumption, and firms assessment customer bases $S_{s, t-1}^f$ for each product s and rivals prices $\sigma_{-f}^p(I_{-f}^f)$ we get

horizontal dimension and drop the nodes that lie outside the simplex. In a robustness check, we use a transformation of the state space that allows us to project points from rectangles to the simplex and vice versa. We find the former to work better in our case.

³⁵Note that the product j defines the firm f

³⁶The suprascripts reflect the agent forming beliefs.

$$S_{jt}^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma) = \sum_{s=1}^F \sum_{n=1}^N S_{snjt}(p_{ft}, \mathfrak{J}, e; \sigma_{-f}(I_{-f}^f); \mu^D) \times S_{s,t-1}^f \times w[n|k] \quad (14)$$

I.1.3 Policy Step

Participation Policy Step Due to our timing assumption, participation decisions are made at the end of the period, once new market shares realize. Optimal prices at iteration k determine σ_k^p . Hence, firms' perceived vector of demand is $S_t^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma_k^p)$, following Equation 14. Then, participation best responses are determined by evaluating the right-hand side of the participation fixed-point problem, using $V_f(x, \mathfrak{J}, e; \hat{q}_k)$ evaluated at $x = S_{jt}^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma_k^p)$. At this point, we can either update firms' best responses or perform several iterations of the fixed point problem to get a better approximation of optimal participation probabilities at the current value function approximations.

Price Policy Step At the beginning of the period firms optimize over prices. Thus, to compute price's best responses, we must calculate values and participation derivatives at the demand generated by each trial price p_j at every possible node $z_i^f(\mathfrak{J}, e)$ for each choice set and exogenous states, holding rivals' strategies fixed: $S_t^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma_{k-1,-f}^p, p_j)$.

$S_t^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma_{k-1,-f}^p, p_j)$ is immediate to compute from our information hold in memory. Evaluating values and participation derivatives is also straightforward. To evaluate the derivatives of the value function, we can take derivatives of the approximated polynomials from the previous step, $V_f(x, \mathfrak{J}, e; \hat{q}_{k-1})$ and evaluate it at $S_t^f(z_i^f(\mathfrak{J}, e), \mathfrak{J}, e; \sigma_{k-1,-f}^p, p_j)$. In the case of participation derivatives, we can use the fact that the participation threshold is a function of the value function and the distribution of fixed costs. Indeed, under a distributional assumption (exponentially distributed fixed costs) we can compute the derivative of participation policies in closed form using the implicit function theorem —see Appendix I.5.

Hence, evaluating the right-hand side of Equation 7 is not computationally demanding. By doing so, we can compute the next guess of prices towards its fixed point. Finally, we might speed up computations by not solving the fixed point at every iteration k but only make decent progress toward it. Although we can reach the value function's fixed point faster by not forcing prices to be optimal on each step, we ensure that equilibrium prices are indeed optimal once we reach the fixed point of the value function operator.

I.1.4 Value function update

Finally, we update the values at each interpolating node z_i , choice set \mathfrak{J} and exogenous state e , $V_{k+1}(z_i, \mathfrak{J}, e)$, using $\sigma_{kf}^p(z_i^f, \mathfrak{J}, e)$, $\sigma_{kf}^\phi(S^f, \mathfrak{J}, e')$, and $V_{fk}(S^f, \mathfrak{J}', e', \hat{q}_k)$. While computing variable

profits at optimal prices and expected continuation payoffs (from previously computed values) is immediate, obtaining the expected value of fixed costs conditional on participating is slightly more difficult. However, under our assumption that fixed costs are distributed exponentially with mean values θ_{FC} we can compute the next period's values in closed form as³⁷

$$\begin{aligned}
V_{f,k+1}(z_i^f, \mathfrak{I}, e) &= M \times S_{jt}^f(z_i^f, \mathfrak{I}, e; \sigma_k^p) \times (\sigma_{kj(f)}^p(z_i^f, \mathfrak{I}, e) - c_{j(f)}(e)) - \\
&\sum_{e' \in \mathcal{E}} \left(\phi_{kj(f)}(S_{jt}^f(z_i^f, \mathfrak{I}, e; \sigma_k^p), \mathfrak{I}, e') \times \theta_{FC} - \left[1 - \phi_{kj(f)}(S_{jt}^f(z_i^f, \mathfrak{I}, e; \sigma_k^p), \mathfrak{I}, e') \right] \times \bar{\Theta}(S_{jt}^f(z_i^f, \mathfrak{I}, e; \sigma_k^p), \mathfrak{I}, e'; \phi_k) Pr(e'|e) \right) + \\
&\sum_{e' \in \mathcal{E}} \sum_{\mathfrak{I}' \in \mathcal{I}} V_{f,k}(S_{jt}^f(z_i^f, \mathfrak{I}, e; \sigma_k^p), \mathfrak{I}', e'; \hat{q}_k) Pr(\mathfrak{I}'|\phi_k) Pr(e'|e)
\end{aligned} \tag{15}$$

Then, we proceed to the next iteration. Let Ψ^p and Ψ^ϕ refer to the RHS of Equation 7 and Equation 5 respectively, and denote Equation 15 the Bellman Equation. Algorithm 1 describes the algorithm used to solve for the equilibrium.

Algorithm 1 Equilibrium Solver

- 1: Choose projection nodes $\mathcal{S} \rightarrow$ state space is $\mathcal{X} = \mathcal{S} \times \mathcal{I}$.
- 2: Set $V_f^0[x] \quad \forall x \in \mathcal{X}, f$

- 3: **while** $crit^V > tol^V$ **do**
- 4: **while** $crit^\phi > tol^\phi$ or $iter^\phi < \maxiter(k_\phi)$ **do**
- 5: $Pr(\mathfrak{I}'|\mathfrak{I}) \leftarrow \sigma_{k_\phi-1}^\phi$
- 6: $\sigma_{k_\phi}^\phi \leftarrow \Psi^\phi \left(V_f(S_{jt}^f(\sigma_{k_\phi-1}^p); \hat{q}_{k-1}), S_{jt}^f(\sigma_{k_\phi-1}^p), Pr(\mathfrak{I}'|\mathfrak{I}) \right)$
- 7: $crit^\phi = \|\sigma_{k_\phi+1(k_p)}^\phi - \sigma_{k_\phi(k_p)}^\phi\| / \|\sigma_{k_\phi(k_p)}^\phi\|$
- 8: **end while**

- 9: Set $\sigma_f^{p,0(k_V)}$
- 10: **while** $crit^p > tol^p$ or $iter^p < \maxiter(k_p)$ **do**
- 11: $\nabla_S V_f(S_{k_p-1}^f(\sigma_{k_p-1}^p); \hat{q}_{k-1}) \leftarrow$ derivative of Chebyshev polynomials
- 12: $\nabla_S \sigma_k^\phi$ at $S_{k_p-1}^f(\sigma_{k_p-1}^p) \leftarrow$ by IFT
- 13: $\sigma_{k_p}^p \leftarrow \Psi^p \left(S_{k_p-1}^f, \nabla_S V_f, \nabla_S \sigma_k^\phi, Pr(\mathfrak{I}'|\mathfrak{I}), Pr(e'|e) \right)$
- 14: $S_{k_p}^f \leftarrow \sigma_{k_p}^p$ according to Equation 14
- 15: $crit^p = \|\sigma_{k+1(k_V)}^p - \sigma_{k(k_V)}^p\| / \|\sigma_{k(k_V)}^p\|$
- 16: **end while**

- 17: Update $V_f^k \leftarrow$ according to Equation 15
- 18: Compute $\{\hat{q}_k\}$ by OLS.
- 19: $crit^V = \|V_k - V_{k-1}\| / \|V_{k-1}\|$
- 20: **end while**

³⁷We are omitting the explicit dependence of z_i^f on \mathfrak{I} and e to avoid loading the notation.

I.2 Beliefs Computation

Constructing expected customer base for non-tracked products Let $\Omega(\mathbb{J})$ be a random variable representing the shares of the MPE's recurrent class of the pricing game without entry and exit under choice set \mathbb{J} . Moreover, let $pr(\Omega(\mathbb{J})|I_f)$ be its distribution of market shares conditional on firm f information set (note that $\mathbb{J} \in I_f$).

Then, for a given realization ω_f from the distribution $pr(\Omega(\mathbb{J})|I_f)$ firms can construct the vector of customer bases as

$$S_{k,t-1}^{e(f)}(\omega_f, I_f) = \begin{cases} S_{k,t-1}^f & k \in T^f \\ \omega_k|I_f & \text{otherwise} \end{cases} \quad (16)$$

where the ω_f argument reflects that $S_{k,t-1}^{e(f)}(\cdot, I_f)$ is a random variable, whose distribution is determined by $pr(\Omega(\mathbb{J})|I_f)$. The key limitation of this approach is that the private information component of I_f , $(S_{t-1}^f, \bar{S}_{t-1}^f)$ might not be observed in the support of $\Omega(\mathbb{J})$ because it is a continuous variable. Hence, we perform simple linear interpolations to construct firms beliefs about the vector of customer bases, conditional on their information set I_f .

Then, firms can leverage this information to construct a distribution of rivals information sets $pr(I_{-f}|I_f)$ simply aggregating rivals tracked and non-tracked shares according to $S_{k,t-1}^{e(f)}$. Let ζ_f be a realization of $pr(I_{-f}|I_f)$, then firms evaluate rivals' actual strategies at their perceived information sets: $\sigma_{-f}(\zeta_f)$.

Constructing shares without observing consumer-type/past-choice shares. For a given vector of prices $p_t = (\sigma_f^p(I_f), \sigma_{-f}^p(\zeta_f))$, choice set \mathbb{J} , and consumer preferences μ^D , we can compute the probability that consumer-type n , previously consuming product k , chooses product j . However, constructing the market share of product j at time t requires knowledge on the joint distribution of customer bases and product types, which firms do not possess. We leverage our assumption on the conditional distribution of types conditional on previous consumption $w[n|k]$ to construct market shares.

$$S_{jt}^{e(f)}(\omega_f, \zeta_f, I_f; \sigma) = \sum_{k=1}^F \sum_{n=1}^N S_{knjt}(p_{jt}, \mathbb{J}; \sigma_{-f}(\zeta_f), \mu^D) S_{k,t-1}^{e(f)}(\omega_f, I_f) w[n|k]$$

Note that using this information firms can construct the distribution of next period tracked shares and their aggregates, which defines the Markov process on information sets: $pr(I'_f|I_f, \sigma)$.

Expected Payoff Moreover, firms can construct the expected payoff and value functions. For the per-period profits, we write payoffs as

$$\pi^{e(f)}(I_f; \sigma) = M \times \int \left\{ S^{e(f)}(\omega_f, \zeta_f, I_f; \sigma) \times (\sigma_f^p - c_f) - E_{\Theta} \left[\Theta_f \times 1\{\Theta_f \leq \bar{\Theta}(I_f, \sigma_{-f}^{\phi}(\zeta_f))\} \right] \right\} pr(\zeta_f | I_f) pr(\omega_f | I_f) \quad (17)$$

I.3 Absorbing Steady State in Dynamic Game without Entry and Exit

We have mentioned that we select the MPE by initializing the algorithm at the absorbing steady state. In this section, we describe how we compute the absorbing steady state.

Algorithm 2 Equilibrium at absorbing state for fixed choice sets — based on MacKay and Remer [2021].

- 1: Initialize price policy's derivative: $\nabla_s \sigma_0^p$.
- 2: **while** $crit > tol$ **do**
- 3: At each iteration k , initialize $\sigma_0^p(k)$

- 4: **while** $crit^p > tol^p$ **do**
- 5: $S^{ss}(p_{k_p(k)}) \leftarrow S = S(p_{k_p(k)}, S)$.
- 6: $\nabla_S V_f(S^{ss}(p_{k_p(k)})) \leftarrow \nabla_S V_f = \nabla_S \pi_f + \nabla_p \pi_f \nabla_S \sigma_k^p + \nabla_S V_f^T [\nabla_S S + \nabla_p S \nabla_S \sigma_k^p]$
- 7: $p^{k_p+1(k)} \leftarrow$

$$p_f^{k_p+1(k)} = \left(c_f - \frac{\beta}{M} \frac{\partial V_f}{\partial S_f} \right) - \frac{S_f}{\frac{\partial S_{f_t}}{\partial p_f}} - \sum_{r: \mathbf{J}_r=1, k \neq f} \frac{\frac{\partial S_k}{\partial p_f}}{\frac{\partial S_f}{\partial p_f}} \left(\frac{\beta}{M} \frac{\partial V_f}{\partial S_r} \right)$$
- 8: $crit^p = \|p^{k+p+1(k)} - p^{k_p(k)}\| / \|p^{k_p(k)}\|$
- 9: **end while**

- 10: $\frac{\partial \sigma_{k+1}^p}{\partial S} \leftarrow$ numerically differentiating p_k^* .
- 11: $crit = \|\nabla_S \sigma_{k+1}^p - \nabla_S \sigma_k^p\| / \|\nabla_S \sigma_k^p\|$
- 12: **end while**

The crucial aspect of the computation is that we can leverage restrictions about the absorbing steady-state. In particular, we solve for the steady state shares for any guess of prices from $S^{ss}(\mathbf{J}) = S(S^{ss}, \mathbf{J}, P)$. Additionally, we can solve for the value function derivatives (at any market share) given a guess of the price policy derivatives. Therefore, we solve for steady state prices from a guess of price policy derivatives, iterating on firms' FOC. Finally, we can update our price policies' derivatives guess, by numerically differentiating these prices. Algorithm 2 describes the algorithm for any of these choice sets.

Therefore, we can circumvent the curse of dimensionality that would otherwise arise from solving for the equilibrium at each point in the state space. Hence, we can quickly obtain the value at such a point. The following algorithm details our computations. We solve the steady state for every choice set \mathbf{J} , assuming it will remain constant in the future.

I.4 Multi-Product Price Inversion

Once we move into the multi-product case, we carry out price updates (line 14 in Algorithm 1 and line 7 in Algorithm 2) using Morrow and Skerlos [2011] as described in Conlon and Gortmaker [2020]. It requires minimal changes to adapt Morrow and Skerlos [2011] inversion to the dynamic setting used in this paper. Recall, that the derivative of market shares with respect to prices can be broken up into two parts

$$\frac{\partial S}{\partial p}(p) = \Lambda(p) - \Gamma(p)$$

where Λ is a diagonal matrix with diagonal elements

$$\Lambda_{jj} = \sum_w \sum_n (-\alpha_n) S_j(w, n) dF(D_n|w) dF(w)$$

and Γ is a dense matrix with elements

$$\Gamma_{jk} = \sum_w \sum_n \alpha_n S_j(w, n) S_k(w, n) \frac{\partial S_k(w)}{\partial p_j} dF(D_n|w) dF(w)$$

with $dF(D_n|w)$ and $dF(w)$ denoting the weights of consumer types conditional on consuming product w in the past and $dF(w)$ indicates the share of product w consumers.

Then, price FOC can be written as

$$p = c + \Lambda(p)^{-1} (\mathcal{O} \cdot \Gamma(p)) (p - c) - \Lambda^{-1}(p) \left\{ S + \frac{\beta}{M} \times \nabla_p S \times E[\nabla_S V] + \frac{\beta}{M} \nabla_p S \times \sum_{-f} (E[V|\mathbf{J}_{-f} = 1] - E[V|\mathbf{J}_{-f} = 0]) \nabla_S \phi_{-f} \right\}$$

I.5 Participation Problem under Exponentially Distributed Fixed Costs

As shown by Doraszelski and Satterthwaite [2010] the participation problem can be expressed in terms of participation thresholds, or participation probabilities. The latter representation usually turns out to be more useful. Thus, we can express the participation problem as

$$\max_{\phi_j} -E[\Theta_j \times 1\{\Theta_j \leq F^{-1}(\phi_j)\}] + \beta E[V_f(\mathbf{J})]$$

Now, assuming fixed costs are exponentially distributed, we have

$$E[\Theta_f \times 1\{\Theta_f \leq F^{-1}(\phi_f)\}] = \phi\mu - (1 - \phi)F^{-1}(\phi_f)$$

where μ is the unconditional mean of the fixed cost distribution. Moreover,

$$F^{-1}(\phi_f) = -\mu \times \log(1 - \phi_j)$$

FOC of the participation problem boil down to

$$F_j^\phi = \beta (E[V_f | \mathfrak{I}_r = 1] - E[V_f | \mathfrak{I}_r = 0]) + \mu \log(1 - \phi_j) = 0$$

Therefore, applying the implicit function theorem, it is easy to see that

$$\nabla \phi_s = -(\nabla F_\phi^\phi)^{-1} \nabla F_s^\phi \quad (18)$$

Furthermore, these Jacobian matrices are easy to characterize and derive in closed form. The diagonal of $\frac{\partial \phi}{\partial s}$ is determined by the effect of higher participation probabilities on participation costs, while off diagonal effects' are influenced by how rivals participation influence participation threshold. Thus,

$$\frac{\partial F_j^\phi}{\partial \phi_k} = \begin{cases} -\frac{\mu}{1-\phi_j} & \text{if } j = k \\ \beta \{ (E[V_f | \mathfrak{I}_j = 1, \mathfrak{I}_k = 1] - E[V_f | \mathfrak{I}_j = 0, \mathfrak{I}_k = 1]) - (E[V_f | \mathfrak{I}_j = 1, \mathfrak{I}_k = 0] - E[V_f | \mathfrak{I}_j = 0, \mathfrak{I}_k = 0]) \} & \text{if } j \neq k \end{cases}$$

On the other hand, ∇F_s^ϕ is simply the gradient of the participation threshold with respect to loyal bases.

I.6 Equilibrium Behavior - Goodness of Fit

I.6.1 Beliefs Consistency

Figure I.1: MSE Expectation versus Observed Outcomes - Market Shares

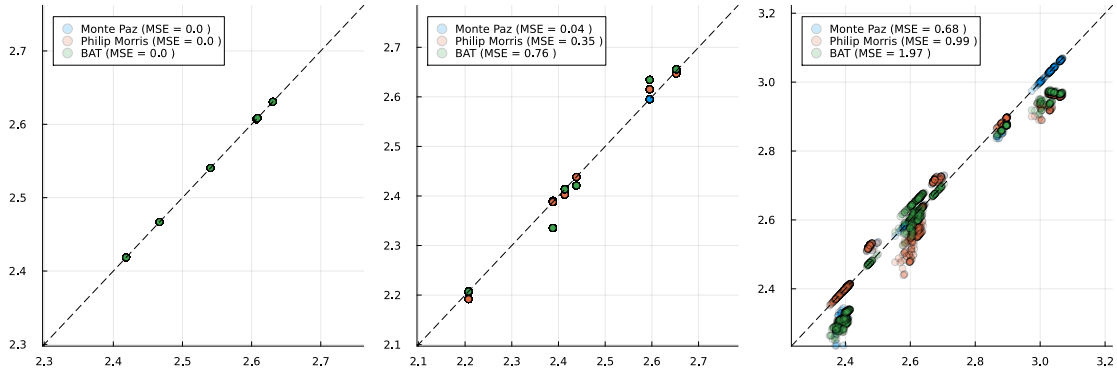


Figure I.2: MSE Expectation versus Observed Outcomes - Prices


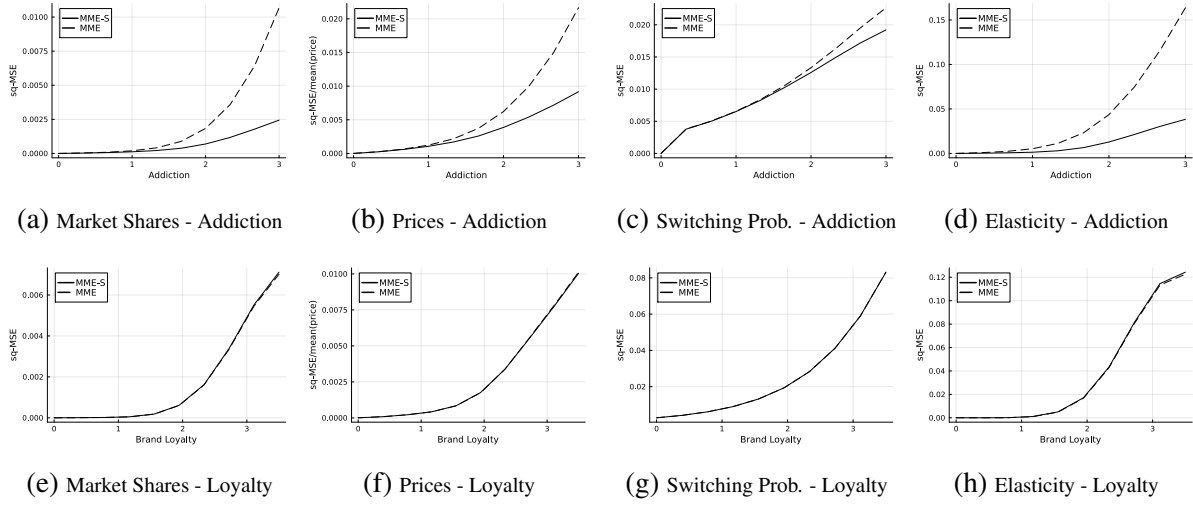


Table I.1: MSE comparison across models

	No Dynamics		No Entry		Baseline	
	Price MSE	Share MSE	Price MSE	Share MSE	Price MSE	Share MSE
Monte Paz	6.30e-09	2.72e-11	3.53e-02	1.07e-04	6.83e-01	9.52e-04
Philip Morris	5.31e-09	4.13e-11	3.54e-01	1.17e-03	9.94e-01	2.18e-03
British American Tobacco	1.86e-08	9.71e-11	7.57e-01	2.20e-03	1.97e+00	2.97e-03

I.6.2 Long-Run Outcomes under Alternative Equilibrium Definitions

Figure I.3: MSE of equilibrium outcomes over the stationary distribution - Addiction vs. Loyalty



I.7 Supply Estimation

I.7.1 Moment Selection and Empirical Objective Function

Let $q_1(\chi_{kt}, \mathbb{X}_{kt}, \theta) = \chi_{kt} - E[\sigma_k^\chi(\tilde{S}_{t-1}^f, \mathbb{J}_t, \delta_t, tax_t; \varepsilon_t, \Theta_{kt}; \theta_0) | \mathbb{X}_{kt}]$, $q_2(p_{kt}, \mathbb{X}_{kt}, \theta) = p_{kt} - E[\sigma^p(\tilde{S}_{t-1}, \mathbb{J}_t, \delta_t, tax_t; \varepsilon_t; \theta) | \mathbb{X}_{kt}]$, and $q(\chi_{kt}, p_{kt}, \mathbb{X}_{kt}; \theta)$ the vertical stack of q_1 and q_2 . Then we can form unconditional moments for chosen basis of \mathbb{X}_{kt} -or other suitable instruments, $h_1(\mathbb{X}_{kt})$ and $h_2(\mathbb{X}_{kt})$ as:

$$g(\chi_{kt}, p_{kt}, \mathbb{X}_{kt}; \theta) = E \begin{pmatrix} q_1(\chi_{kt}, \mathbb{X}_{kt}, \theta) \times h_1(\mathbb{X}_{kt}) \\ q_2(p_{kt}, \mathbb{X}_{kt}, \theta) \times h_2(\mathbb{X}_{kt}) \end{pmatrix} = 0 \quad \text{at } \theta = \theta_0$$

If the model is overidentified, then we can construct the objective function G such that θ_0 minimizes it,

$$G(\theta) = g(\mathbb{X}_{kt}; \theta)' W q(\mathbb{X}_{kt}; \theta)$$

where W is a positive definite weighting matrix.

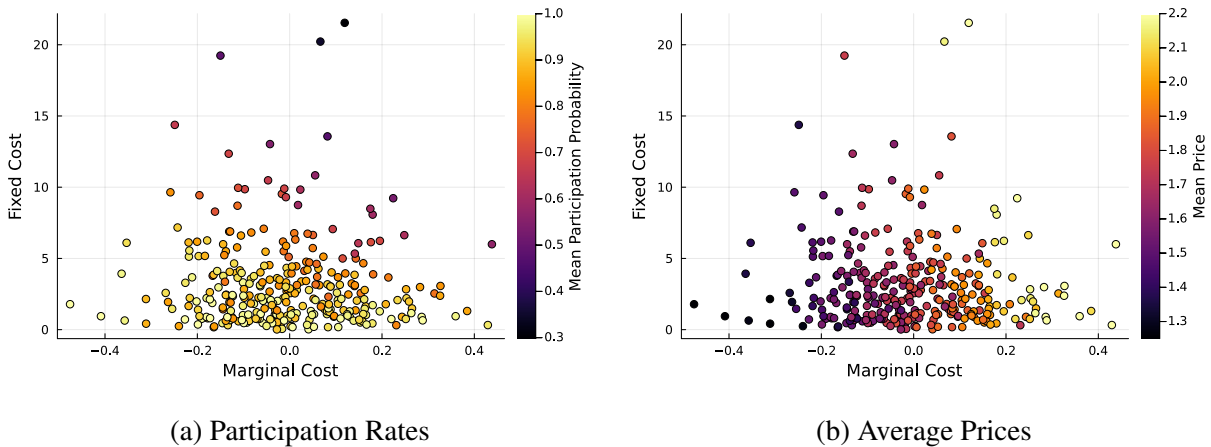
For each candidate parameter θ , we solve the equilibrium, compute the expected prices and participation functions, and construct the empirical version $G(\theta)$: $G_n(\theta)$. Our estimator is defined by

$$\hat{\theta} = \arg \min_{\theta} G_n(\theta)$$

Hansen [1982] establish the regularity conditions that ensure asymptotic normality of $\hat{\theta}$.

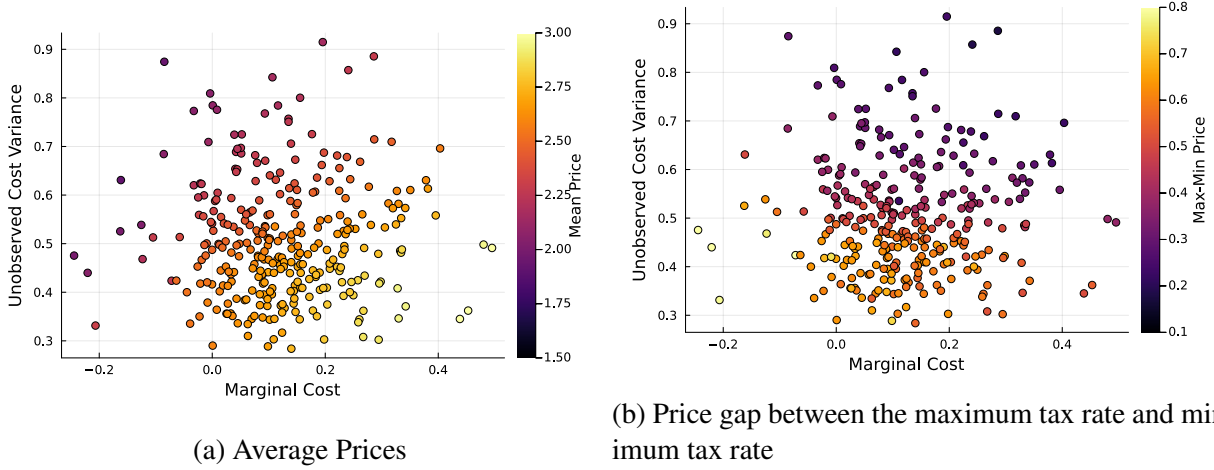
Identification — Plane cuts implied by data. In Section 4.2 we argue that conduct still help us identify marginal and fixed costs despite the fact that we cannot break up the problem and identify marginal costs from prices and conduct, and fixed costs from participation choices and conduct. Here, we provide a graphical representation that illustrate that prices and participation choices contain distinct information about the sequence of marginal and fixed costs that can rationalize it. In other words, the participation and price optimality equations do not appear to be colinear. Figure I.4 shows the sequence of marginal and fixed costs that generate alternative BAT’s flagship average prices and participation rates. Each dot represents a simulated parameter. Their color indicates the average price (right panel) they would set in our sample and the average participation probability (left panel). Then, observed prices and participation select different cuts of the plane

Figure I.4: Identification of Fixed and Marginal costs, an illustration with BAT prices and participation rates.



We then observe that cost pass-through together with conduct also help us identify the unobserved cost shocks. Figure I.5 illustrate how identification works. These plots show all possible mean prices (left panel) and price variation (the gap between the price at the max tax with respect to the min tax) that alternative pairs of marginal costs and unobserved cost variances could generate. Then, we use the observed mean prices and price variation to select the region of the marginal cost-unobserved cost variance plane that is compatible with the data.

Figure I.5: Identification of Unobserved Variance, an illustration with Monte Paz prices



Moment Selection Following the intuition in our previous arguments, we pick the basis (h_1, h_2) to match the average price and participation rates, their correlation with taxes, and their correlation with their customer bases. Then, we pick the corresponding basis h_1, h_2 . h_1 are almost identical to h_2 , and composed of dummy indicators for each product j , taxes, the products' own customer base, and other products' customer base. The first J moments are equivalent to match the average price and participation by product. The elements of h_2 are interacted by an indicator of whether these products are currently being offered in the market, because if they are not firms do not make price decisions. Hence, our empirical objective function is

$$Gn(\theta) = \begin{pmatrix} \frac{1}{T} \sum_t q_{11t} \\ \vdots \\ \frac{1}{T} \sum_t q_{1Jt} \\ \frac{1}{T} \sum_t q_{1jt} tax_t \\ \frac{1}{T} \sum_t q_{1jt} S_{jt-1} \\ \frac{1}{T} \sum_t q_{1jt} S_{jt-1}^{other} \\ \frac{1}{T_{active}} \sum_{t: \mathbb{J}_{1t}=1} q_{21t} \\ \vdots \\ \frac{1}{T_{active}} \sum_{t: \mathbb{J}_{tJ}=1} q_{2Jt} \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} tax_t \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} S_{jt-1} \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} S_{jt-1}^{other} \end{pmatrix}' W \begin{pmatrix} \frac{1}{T} \sum_t q_{11t} \\ \vdots \\ \frac{1}{T} \sum_t q_{1Jt} \\ \frac{1}{T} \sum_t q_{1jt} tax_t \\ \frac{1}{T} \sum_t q_{1jt} S_{jt-1} \\ \frac{1}{T} \sum_t q_{1jt} S_{jt-1}^{other} \\ \frac{1}{T_{active}} \sum_{t: \mathbb{J}_{1t}=1} q_{21t} \\ \vdots \\ \frac{1}{T_{active}} \sum_{t: \mathbb{J}_{tJ}=1} q_{2Jt} \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} tax_t \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} S_{jt-1} \\ \frac{1}{T_{active}} \sum_j \sum_{t: \mathbb{J}_{tJ}=1} q_{2jt} S_{jt-1}^{other} \end{pmatrix}$$

I.7.2 Importance Sampling

Finding the $\hat{\theta}$ that minimizes $G_n(\theta)$ can imply many evaluations of $g(\chi_{kt}, p_{kt}, \mathbb{X}_{kt}; \theta)$. Evaluating this function implies solving the equilibrium many times. We follow Akerberg [2009] and use importance sampling together with a change of variable, to avoid solving the game for each evaluation of the parameters. To this extent, we perturb the econometric model and add uncertainty in the production costs and the mean value of the fixed costs' distributions.

Hence, we re-write variable and fixed costs parameters as

$$\begin{aligned}\theta_k^{vc} &= \mu_k^{vc} + \sigma^{vc} \varepsilon_k^{vc} \\ \theta_S &= e^{\mu^S + \sigma^S \varepsilon^S} \\ \theta_R &= e^{\mu^R + \sigma^R \varepsilon^R} \\ \sigma_\varepsilon &= e^{\mu^{\sigma\varepsilon} + \sigma^{\sigma\varepsilon} \varepsilon^{\sigma\varepsilon}}\end{aligned}$$

where $(\varepsilon_k^{vc}, \varepsilon^S, \varepsilon^R, \varepsilon^{\sigma\varepsilon}) \sim N(0, I)$.

Under this reformulation, we re-write the objective function in terms of $\mu = (\{\mu_k^{vc}\}, \mu^S, \mu^R, \mu^{\sigma\varepsilon})$, $\sigma = (\sigma_k^{vc}, \sigma^S, \sigma^R, \sigma^{\sigma\varepsilon})$, and importance sampling noise ε^{IS} .

$$G_n(\mu, \sigma) = \frac{1}{T \times N} \sum_{t,k} E_{\varepsilon^{IS}} g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; \mu, \sigma, \varepsilon^{IS})' W E_{\varepsilon^{IS}} g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; \mu, \sigma, \varepsilon^{IS})$$

Although we now need to compute the expectation $E_{\varepsilon^{IS}} g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; \mu, \sigma, \varepsilon^{IS})$, we make a change of variable and use importance sampling to reduce the cost of doing so. The change of variable is such that

$$u_s^l = \mu^l + \sigma^l \varepsilon_s^l \quad l = \{vc, S, R, \sigma\varepsilon\}$$

Let u_s be the vector of all new variables. Furthermore, let $f(u_s | \mu, \sigma)$ be the density function of u obtained by the change of variables formula, and define u 's importance sampling density to be a normal distribution $N(\mu_0, \Sigma_0)$ with density function $g(u)$, which does not depend on parameters $\{\mu^l, \sigma^l\}$.

Then the simulated moment is

$$\tilde{E}_{\varepsilon^{IS}} g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; \mu, \sigma, \varepsilon^{IS}) = \frac{1}{S} \sum_s g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; u_s) \frac{f(u_s | \mu, \sigma)}{g(u_s)}$$

where u_s are draws from $g(\cdot)$. Finally, we can write the importance sampling method of MSM as

$$\hat{\mu}, \hat{\sigma} = \arg \min \tilde{G}_n(\mu, \sigma) = \frac{1}{T \times N \times S} \sum_{t,k,s} \left(g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; u_s) \frac{f(u_s | \mu, \sigma)}{g(u_s)} \right)' W \left(g(\mathbb{X}_{kt}, p_{kt}, \chi_{kt}; u_s) \frac{f(u_s | \mu, \sigma)}{g(u_s)} \right) \quad (19)$$


The critical aspect of Equation 19 is that u_s remains constant as μ, σ changes. Therefore, we only need to evaluate $q(p_{ft}, \mathbb{X}_t; u_s)$ S times and not for every guess of the parameters μ, σ . Instead, for each trial of the parameters μ, σ , we recompute the importance sampling weights $\frac{f(u_s | \mu, \sigma)}{g(u_s)}$, a much simpler problem.

II Appendix: Convergence and Multiplicity

II.1 Convergence & Existence

As we commented in Section 3, there are no guarantees that an equilibrium exists or is unique. However, our algorithm converges for wide regions of the parameter space, which we take as strong evidence that existence is not a significant issue—see Figure 6. We observe some non-convergence at high inertia values, but we believe these are more closely related to issues with the parametric approximations when the model becomes highly non-linear than evidence of non-existence.

Figure 6: Convergence Plots



II.2 Multiplicity in game without entry and exit

We then explore equilibrium multiplicity in the game without entry and exit, using several techniques. First, we use our simulation design. We draw 5000 simulated primitives with replacement and solve the game without entry and exit for each draw. We then regress equilibrium outcomes on primitives and find that they explain more than 99% of the equilibrium outcomes variation. Table 2 shows the results of the regressions on shares, and Table 3 on prices.

Table 2: Regression of equilibrium shares of the game without entry and exit, on primitives

	Share Leader (new draw)	Share Leader (original)	Share Follower (new draw)	Share Follower (original)
	(1)	(2)	(3)	(4)
(Intercept)	-0.189*** (0.002)	-0.189*** (0.002)	-0.225*** (0.002)	-0.225*** (0.002)
addiction	0.089*** (0.000)	0.089*** (0.000)	0.069*** (0.000)	0.069*** (0.000)
addiction sq	-0.002*** (0.000)	-0.002*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
brand loyalty	0.126*** (0.001)	0.126*** (0.001)	0.065*** (0.001)	0.065*** (0.001)
brand loyalty sq	-0.009*** (0.000)	-0.009*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
compensated utility	0.184*** (0.001)	0.184*** (0.001)	0.163*** (0.001)	0.163*** (0.001)
compensated utility sq	-0.011*** (0.000)	-0.011*** (0.000)	-0.009*** (0.000)	-0.009*** (0.000)
addiction x brand loyalty	-0.011*** (0.000)	-0.011*** (0.000)	-0.004*** (0.000)	-0.004*** (0.000)
brand loyalty x utility	-0.009*** (0.000)	-0.009*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)
Estimator	OLS	OLS	OLS	OLS
N	4,997	4,997	4,997	4,997
R^2	0.996	0.996	0.995	0.995

Note:

Table 3: Regression of equilibrium prices of the game without entry and exit, on primitives

	Price Leader (new draw)	Price Leader (original)	Price Follower (new draw)	Price Follower (original)
	(1)	(2)	(3)	(4)
(Intercept)	1.736*** (0.006)	1.736*** (0.006)	1.918*** (0.006)	1.918*** (0.006)
addiction	0.014*** (0.001)	0.014*** (0.001)	-0.018*** (0.001)	-0.018*** (0.001)
addiction sq	0.014*** (0.000)	0.014*** (0.000)	0.012*** (0.000)	0.012*** (0.000)
brand loyalty	-0.139*** (0.002)	-0.139*** (0.002)	-0.206*** (0.002)	-0.206*** (0.002)
brand loyalty sq	0.043*** (0.000)	0.043*** (0.000)	0.046*** (0.000)	0.046*** (0.000)
compensated utility	0.293*** (0.004)	0.293*** (0.004)	0.194*** (0.004)	0.194*** (0.004)
compensated utility sq	-0.022*** (0.001)	-0.022*** (0.001)	-0.017*** (0.001)	-0.017*** (0.001)
addiction x brand loyalty	0.045*** (0.000)	0.045*** (0.000)	0.042*** (0.000)	0.042*** (0.000)
brand loyalty x utility	0.015*** (0.001)	0.015*** (0.001)	0.018*** (0.001)	0.018*** (0.001)
Estimator	OLS	OLS	OLS	OLS
N	4,997	4,997	4,997	4,997
R^2	0.992	0.992	0.990	0.990

Note:

We complement this analysis with more formal methods to look for multiplicity robustly. First, we use the method suggested by Reguant and Pareschi [2021] to bound all feasible counterfactual

outcomes in a simplified version of the game.³⁸ The equilibrium bounds are narrow for large regions of the parameter space, indicating that the equilibrium is likely to be unique. Moreover, we employ homotopy to find all equilibria of the game as in Besanko et al. [2010]. Again, we find no evidence of multiplicity.³⁹

II.3 Equilibrium Selection

Then, we use the solution of the steady state of the game without entry and exit as the initial values to start the algorithm that solves the full game. Although this method does not guarantee that the selected equilibrium is the same for every parameter, we find evidence supporting it. Concretely, we regress equilibrium outcomes on primitives. Table 4, Table 5, and Table 6 show the results of the regressions. We find that the primitives almost perfectly explain the equilibrium outcomes after we follow such equilibrium selection.

³⁸The simplified version is equivalent to the model in Dubé et al. [2009] and Chen [2016]. The main difference to our baseline model is that firms face a single consumer whose affiliation evolves over time. Under that specification, the equilibrium of the game exists, and computations simplify.

³⁹pending

Table 4: Regression of Equilibrium Participation on Primitives

	Participation Leader		Participation Follower	
	(1)	(2)	(3)	(4)
(Intercept)	-2.332*** (0.068)	-3.459*** (0.876)	0.176* (0.079)	-10.562** (3.246)
addiction	0.750*** (0.022)	1.027*** (0.227)	-0.165*** (0.025)	2.236** (0.843)
addiction sq	-0.042*** (0.002)	-0.055** (0.017)	0.030*** (0.002)	-0.086 (0.064)
brand loyalty	0.772*** (0.017)	1.209*** (0.332)	-0.100*** (0.020)	4.812*** (1.229)
brand loyalty sq	-0.046*** (0.001)	-0.081* (0.033)	0.006*** (0.001)	-0.547*** (0.123)
utility	1.506*** (0.038)	1.936*** (0.282)	0.069 (0.045)	0.646 (1.045)
utility sq	-0.182*** (0.005)	-0.187*** (0.028)	0.014* (0.006)	0.108 (0.103)
addiction x brand loyalty	-0.089*** (0.003)	-0.149*** (0.040)	0.026*** (0.003)	-0.537*** (0.147)
addiction x utility	-0.166*** (0.006)	-0.218*** (0.039)	0.067*** (0.007)	0.032 (0.145)
brand loyalty x utility	-0.165*** (0.005)	-0.276*** (0.048)	0.031*** (0.006)	-0.161 (0.178)
Estimator	OLS	OLS	OLS	OLS
N	2,241	154	2,241	154
R^2	0.978	0.957	0.987	0.889

This is evidence that our procedure effectively selects the same equilibrium every time, even though multiple equilibria might exist. The only region where the explanatory power of primitives is relatively low is for prices at high inertia levels (when the average probability of repeating product choices is above 85%). Nevertheless, it is hard to disentangle the effect of multiplicity from deficient parametric approximations in highly non-linear regions of the parameter space. In any case, we flag that region as problematic.

Table 5: Regression of Equilibrium Prices on Primitives

	Price, No Entry/Exit, Leader		Price, MPE, Leader		Price, No Entry/Exit, Follower		Price, MPE, Follower	
	(1)	(2)	(3)	(4)	(5)	(6)		
(Intercept)	1.989*** (0.030)	1.648*** (0.052)	0.250 (2.043)	2.208*** (0.030)	2.310*** (0.040)	-6.818 (4.174)		
addiction	-0.071*** (0.009)	-0.126*** (0.016)	0.546 (0.530)	-0.114*** (0.010)	-0.070*** (0.013)	0.704 (1.084)		
addiction sq	0.021*** (0.001)	0.041*** (0.002)	-0.033 (0.040)	0.019*** (0.001)	0.014*** (0.001)	0.077 (0.082)		
brand loyalty	-0.201*** (0.008)	-0.211*** (0.013)	0.409 (0.774)	-0.273*** (0.008)	-0.085*** (0.010)	5.353*** (1.581)		
brand loyalty sq	0.048*** (0.001)	0.058*** (0.001)	-0.017 (0.078)	0.051*** (0.001)	0.014*** (0.001)	-0.665*** (0.159)		
utility	0.152*** (0.017)	0.389*** (0.029)	0.365 (0.658)	0.031 (0.017)	-0.035 (0.023)	-1.526 (1.344)		
utility sq	-0.002 (0.002)	-0.036*** (0.004)	-0.077 (0.065)	0.006* (0.002)	0.018*** (0.003)	0.502*** (0.132)		
addiction x brand loyalty	0.054*** (0.001)	0.089*** (0.002)	-0.039 (0.093)	0.052*** (0.001)	0.006*** (0.002)	-0.428* (0.189)		
addiction x utility	0.024*** (0.003)	0.024*** (0.005)	-0.004 (0.091)	0.027*** (0.003)	0.019*** (0.004)	0.502*** (0.187)		
brand loyalty x utility	0.031*** (0.002)	0.026*** (0.004)	0.062 (0.112)	0.036*** (0.002)	-0.004 (0.003)	-0.098 (0.229)		
Estimator	OLS	OLS	OLS	OLS	OLS	OLS		
<i>N</i>	2,241	2,087	154	2,241	2,087	154		
<i>R</i> ²	0.992	0.985	0.908	0.991	0.991	0.485		

