ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 86, No. 5 (September, 2018), 1651–1687

MARKET POWER AND THE LAFFER CURVE

EUGENIO J. MIRAVETE
Department of Economics, The University of Texas at Austin, Centre for Competition Policy/UEA, and CEPR

KATJA SEIM
Wharton School, University of Pennsylvania, CEPR, and NBER

JEFF THURK
Department of Economics, University of Notre Dame

Econometrica, Vol. 86, No. 5 (September, 2018), 1651–1687

MARKET POWER AND THE LAFFER CURVE

EUGENIO J. MIRAVETE
Department of Economics, The University of Texas at Austin, Centre for Competition Policy/UEA, and CEPR

KATJA SEIM
Wharton School, University of Pennsylvania, CEPR, and NBER

JEFF THURK
Department of Economics, University of Notre Dame

We study commodity taxation and characterize the Laffer curve, a trade-off between tax rates and revenue, in noncompetitive markets. Pricing in these markets leads to incomplete tax pass-through and agents re optimize their purchase and pricing decisions in response to any tax change. We use detailed data from Pennsylvania, a state that monopolizes retail sales of alcoholic beverages, to estimate a model of demand for horizontally differentiated products that ties consumers' demographic characteristics to heterogeneous preferences for spirits. We find that under the state's current tax policy, spirits are overpriced. Distillers respond to decreases in the tax rate by increasing wholesale prices, which limits the state's revenue gain to only 13% of the incremental tax revenue predicted under the common assumption of perfect competition. The strategic response of noncompetitive firms to changes in taxation therefore flattens the Laffer curve significantly.

KEYWORDS: Laffer curve, commodity taxation, market power, public monopoly pricing.

1. INTRODUCTION

THE EMPIRICAL PUBLIC ECONOMICS LITERATURE almost exclusively assumes perfect competition when considering how to set commodity taxes to raise tax revenue with minimal allocative distortions. In such settings, the amount of tax revenue raised is purely a function of demand and supply elasticities. At the same time, the industrial organization literature documents the existence of market power across a wide variety of settings. For example, in homogeneous product industries, such as cement, transportation costs result in highly localized competition among few firms. Even in industries with an abundance of competitors, such as beer, horizontal differentiation between products bestows market power in the part of the product characteristic space where the firm competes.

In this paper, we study how strategic pricing by firms impacts optimal taxation. Acknowledging the ability of firms to strategically respond to changes in tax policy alters the relationship between commodity tax rates and the amount of tax revenue generated—

Eugenio J. Miravete: miravete@eco.utexas.edu
Katja Seim: kseim@wharton.upenn.edu
Jeff Thurk: jthurk@nd.edu

We are grateful to the editor and four referees for their guidance and thorough reading of our manuscript. We also thank Thomas Krantz at the Pennsylvania Liquor Control Board as well as Jerome Janicki and Mike Ehtesham at the National Alcohol Beverage Control Association for access to the data. We are grateful for comments received at several seminar and conference presentations, and in particular to Jeff Campbell, Kenneth Hendricks, and Joel Waldfogel. We are solely responsible for any errors.

 https://doi.org/10.3982/ECTA12307

the Laffer curve—derived under perfect competition.[1] Why might optimal taxation differ once we recognize that firms frequently have pricing power? In an oligopoly, firms price at a level where demand is elastic. If elasticity of demand increases with price, the introduction of a tax entails larger quantity distortions than in the competitive case. Since oligopolistic firms face downward sloping residual demands, they have an incentive to revise their pricing decisions in response to a tax, introducing an additional factor beyond demand and supply elasticities in determining the amount of tax revenue generated. Such pricing responses, if they are significant, change features of the Laffer curve so any taxation policy that fails to account for these responses is likely to be ineffective. This may be particularly true once one allows for the heterogeneous cross-price effects associated with differentiated products.

We present a theoretical derivation of the Laffer curve that allows for the existence of market power among taxed firms. We then empirically characterize the determinants of the Laffer curve for commodity taxes; an important source of government revenue in the U.S. (17.4% of revenue) as well as in the average developed country (32.7% of revenue).[2] Earlier studies of the Laffer curve focus on competitive general equilibrium effects of labor and capital taxation. For example, Trabandt and Uhlig (2011) use a neoclassical growth model with heterogeneous household preferences to generate calibrated labor and capital Laffer curves.[3]

Our focus on commodity taxation differs from earlier work in two regards. First, we are able to provide empirical evidence for imperfectly competitive markets in an environment where the nature of strategic interaction between oligopolistic firms and their response to taxes is likely better understood than the role of frictions or concentration in either labor or capital markets. Second, we exploit detailed industry data in estimation, which allows us to generate robust empirical estimates of consumer demand and firm market power. In contrast, authors who employ sufficient statistics to study optimal taxation typically abstract from any market power effects (e.g., Auerbach (1985); Chetty (2009); and for contexts similar to ours Chetty, Looney, and Kroft (2009), and Evans, Ringel, and Stech (1999)).[4] The combination of estimated consumer preferences with a model of oligopolistic pricing allows us to empirically determine tax pass-through rates that reflect not only elasticities of demand and supply but also changes in markups as taxes change.[5] Foreshadowing our results, that allows both consumers and imperfectly competitive firms to respond to changes in tax policy has a significant impact on optimal taxation and the characterization of the Laffer curve.

[1] Atkinson and Stiglitz (2015) commented on the surprising persistence of the assumption of perfect competition among firms in the 2015 reprint of their classic textbook on Public Economics: "We went on to emphasize that the model underlying much of the Lectures—and much of public economics—was the Arrow-Debreu model of competitive general equilibrium. Looking back a third of a century later, we are struck that little seems to have changed in this respect."

[2] https://taxfoundation.org/sources-government-revenue-oecd-2016/.

[3] Early theoretical works on the Laffer curve include Ireland (1994), Novales and Ruiz (2002), and Schmitt-Grohé and Uribe (1997). Empirical efforts at assessing the optimality of labor taxation include Stuart (1981) and Fullerton (1982) in static contexts, and Pecorino (2011) who allowed for hours worked to entail human capital accumulation.

[4] An exception is Stolper (2016) who studied consumer pass-through of unit gasoline taxes in markets of differing levels of concentration. The current paper provides the foundation for studying the effect of imperfect competition on consumer pass-through and tax revenue in environments with ad valorem taxes.

[5] We provide an empirical extension to Weyl and Fabinger (2013) who analyzed the impact of imperfect competition on tax pass-through under several model environments but do not address asymmetric firms with horizontally-differentiated products.

To show that a Laffer curve relationship exists for an imperfectly competitive industry, we first use a simple model of monopoly taxation to explore the equilibrium interactions between firms and consumers under alternative levels of the tax rate. We show that raising commodity tax rates beyond a critical level indeed decreases government revenues under very general demand conditions. The shape and location of the Laffer curve depend not only on the tax rate and consumer demand elasticity, but also on the firm's response to taxation. Moreover, we show that the tax rate and the firm's price response are strategic substitutes under most empirically relevant consumer preference specifications. Therefore, the firm response limits any change in retail price induced by a change in the tax rate. Effective tax policy thus depends on anticipating the response of the firm to tax changes—failing to do so would result in the firm's price response unraveling, at least partially, the government's objective.

We evaluate these predictions empirically within the context of Pennsylvania's taxation of distilled spirits. The production, distribution, and sale of alcoholic beverages is the second largest beverage industry in the United States (behind soft drinks) and an important source of government tax revenues. Consequently, the regulation of alcoholic beverages has received considerable attention beyond our work. Seim and Waldfogel (2013), Aguirregabiria, Ershov, and Suzuki (2016), and Illanes and Moshary (2017) studied the effect of entry on prices, spatial competition, and product offerings in different states. Conlon and Rao (2015) and Miller and Weinberg (2017) evaluated how alcohol pricing regulations affect collusive behavior by producers.

Motivated by the widespread use of sales taxes to generate government revenue, our analysis focuses on one particular aspect of optimal commodity taxation: the effect of a common tax rate for different products on the overall level of tax revenues. Of related interest is Miravete, Seim, and Thurk (2017), where we study the disproportionate incidence on certain firms and consumers of current uniform taxes relative to optimal subsidy-free product-specific taxes.[6]

Based on detailed price and quantity data for 2002–2004 across all retail liquor stores in Pennsylvania, we estimate the response elasticities of both upstream distillers and consumers to the state's choice of tax rate. We exploit several notable features of the data. First, in Pennsylvania the state monopolizes both the wholesale and retail distribution of alcoholic beverages and applies a pricing rule that translates wholesale prices into retail prices via a single, uniform ad valorem tax. Thus, distillers effectively choose retail prices taking into account the tax. Second, as a product's retail price is by law common across the state at any point in time, differences in consumer preferences materialize in differences in product purchases. This allows us to let consumer tastes vary systematically across products and demographics leading to more flexible substitution patterns and better estimates of consumer demand and distiller market power. Third, the fact that we observe both wholesale and retail prices enables us to estimate consumer demand without placing any restrictions on distiller conduct and market power ex ante. Our estimation therefore admits the possibility that firms are price-takers and cannot react strategically to changes in policy.

We leverage these features to estimate product-level demand across a variety of consumer types for 312 products produced by 37 firms. The estimated demand model combined with data on wholesale prices and a model of oligopoly pricing reveals that upstream

[6]Similarly, Griffith, O'Connell, and Smith (2017) evaluate the use of product-specific corrective taxes to minimize external health costs from ethanol consumption though they ignore the strategic pricing response of retailers and firms to changes in taxation.

firms in the industry enjoy considerable market power, earning 35 cents in profit for every dollar of revenue. They can therefore react strategically to changes in the tax rate. We use the estimated equilibrium model to trace out the Laffer curve under alternative assumptions of firm conduct and behavior. This allows us to quantify the impact of firm price decisions on the characteristics of the entire Laffer curve rather than just local deviations from observed behavior.

To illustrate the implications of not accounting for the strategic responses of firms to changes in tax policy, we compare market equilibria when the state can either perfectly anticipate the distillers' response to changes in its taxation policy (a Stackelberg equilibrium) and, alternatively, when firms do not respond at all to policy changes, our so-called "naïve" equilibrium. To complete this counterfactual analysis, we also evaluate how the optimal response of upstream firms to tax policy varies with upstream firm market conduct, ranging from single-product pricing to full collusion, and how such a wholesale pricing response translates into changes in the shape and position of the Laffer curve.

We show the shape and location of the Laffer curve in our context depends not only on the interaction of the tax rate and consumers' downstream product demand responses but also on upstream market power and firms' strategic pre-tax price setting. The estimated Laffer curve for a naïve policymaker is much steeper than for the policymaker who correctly anticipates price responses, and this is particularly true as conduct among upstream firms becomes more collusive. Not accounting for price response of taxed upstream firms therefore leads to poor policy recommendations. For instance, a naïve regulator would have concluded that the state could increase tax revenues 7.75% (or $28.74 million) by reducing the ad valorem tax from 53.4% to 30.68%. This reduction in the ad valorem tax would have increased profits for all upstream firms but we show that they could do even better by increasing their wholesale prices by 3.79%, or 34 cents, on average. Thus, the estimated model indicates that upstream prices and the tax rate are strategic substitutes in oligopoly as predicted by our theoretical model.

While this change in upstream price may appear small, the fact that the distillers' price on the elastic region of demand leads to a large change in quantity demanded by consumers. Ultimately, the firm response enables distillers to convert 87% of the incremental tax revenue into firm profits, or equivalently the response limits revenue gains to only 12.97% of the forecasted incremental revenues ($3.73 vs. $28.74 million). It is important to note that this substantial undoing of the state's revenue objective requires no coordination among firms. Where distillers instead to collude in setting wholesale prices the equilibrium effects are worse: in this case, we predict that the tax revenue raised at the naïve tax rate would be 1.36% lower than the amount raised at the current tax rate.

Alternatively, a regulator attempting to maximize tax revenue and endowed with perfect-foresight would have predicted the upstream response and instead would only decrease the tax rate to 39.31%. While the state does manage to increase tax revenue 2.23%, profits among upstream distillers increase 30.80% as does their share of integrated industry profits (from 29.5% to 34.9%). The presence of market power among firms therefore flattens the Laffer curve as these taxed firms strategically respond to convert incremental tax revenue to profits. Moreover, this result extends to any industry in which firms have market power.

We conclude that naïve tax policy—or wrongly assuming perfectly competitive behavior—is largely ineffective at increasing tax revenues. This highlights the significance of the policymaker's ability to account for the responses of firms and consumers to policies. As such, it underscores the importance of recent efforts by, for example, the Congressional Budget Office to consider *Dynamic Scoring* of new proposed legislation by accounting for the response of firms, workers, and consumers to changes in government policy; a direct response to the *Lucas Critique*.

The remainder of paper proceeds as follows. In Section 2, we provide a simple model of taxation under imperfect competition to illustrate the key mechanisms underlying our results. We then describe the data, the Pennsylvania pricing rule, the upstream distillery market, and consumption patterns across demographic groups in Section 3. Section 4 contains an equilibrium discrete choice model of demand for horizontally differentiated spirits. The model incorporates the features of the current pricing regulations while allowing for (but not imposing) imperfect competition in the upstream distillery market. In Section 5, we discuss the estimation procedure and results that benefit from unique features of our data in identifying rich substitution patterns across products. We also show how we use the estimated model to infer upstream market power among distillers. In Section 6, we rely on the estimated demand model and alternative models of upstream conduct to show that this market power has significant implications for the shape and location of the Laffer curve. We conclude in Section 7. In the Supplemental Material (Miravete, Seim, and Thurk (2018)), we describe additional data sources, descriptive statistics, robustness of demand estimates, and other results.

2. A SIMPLE MODEL OF THE LAFFER CURVE UNDER MARKET POWER

We begin by presenting a simple model of monopoly excise taxation. Our goal is to illustrate how a tax authority's choice of tax rate affects the tax revenue it generates when allowing for optimal price responses by taxed firms to changes in policy. Consider the case where a monopoly firm supplies a single product at a constant marginal cost c. The monopolist chooses the pre-tax price, which we denote by p^w representing the wholesale price in our empirical context, for a given tax rate $\tau \geq 0$. The pre-tax price, together with the chosen tax rate, implies a tax-inclusive retail price, p^r, of

$$p^r = (1+\tau)p^w. \tag{1}$$

The monopolist then chooses p^w to maximize profits $\Pi(p^w) = (p^w - c)D(p^r)$, given consumer demand $D(p^r)$ at the tax-inclusive price, which requires

$$D(p^r) + (p^w - c)D'(p^r)(1+\tau) = 0, \tag{2}$$

or equivalently, in terms of the Lerner index,

$$\frac{p^w - c}{p^w} = \frac{-D(p^r)}{D'(p^r)(1+\tau)} \cdot \frac{1+\tau}{p^r} = \frac{-1}{\varepsilon(p^r)}. \tag{3}$$

This standard inverse-elasticity pricing rule relates the pre-tax (wholesale) markup of the monopolist to the inverse of the demand elasticity evaluated at the tax-inclusive (retail) price. The monopolist thus sets p^w so that at the tax-inclusive price, demand is elastic.

To characterize the monopolist's optimal price response to a change in tax policy, we make use of the retail price definition in equation (1) and totally differentiate the first-order condition in equation (2) with respect to p^w and τ to obtain

$$\frac{dp^w}{d\tau} = \frac{-1}{1+\tau} \cdot \frac{(2p^w - c)D'(p^r) + p^r(p^w - c)D''(p^r)}{2D'(p^r) + (p^w - c)(1+\tau)D''(p^r)}. \tag{4}$$

For convenience, we define $\eta(\tau)$ as the elasticity of the monopolist's optimal pre-tax price to a change in the tax rate τ. Using the inverse-elasticity rule (3), the firm's optimal response elasticity is

$$\eta(\tau) \equiv \frac{dp^w}{d\tau} \cdot \frac{\tau}{p^w} = \frac{-1}{1+\tau} \cdot \frac{\left(p^w - \dfrac{p^w}{\varepsilon(p^r)}\right)D'(p^r) - \dfrac{p^w \times p^r}{\varepsilon(p^r)} D''(p^r)}{2D'(p^r) - \dfrac{p^r}{\varepsilon(p^r)} D''(p^r)} \cdot \frac{\tau}{p^w} \quad \text{or}$$

$$\eta(\tau) = \frac{-\tau}{1+\tau} \cdot \frac{\left(1 - \dfrac{1}{\varepsilon(p^r)}\right) - \kappa(p^r)}{2 - \kappa(p^r)}, \tag{5}$$

where the key element that determines the sign of $\eta(\tau)$ is $\kappa(p^r)$, the curvature of demand given by

$$\kappa(p^r) = \frac{D''(p^r)D(p^r)}{\left[D'(p^r)\right]^2}. \tag{6}$$

At optimal prices when $\varepsilon(p^r) < -1$, the response elasticity $\eta(\tau)$ is negative for linear or concave demand since $D''(p^r) \leq 0$ ensures that $\kappa(p^r) \leq 0$. While linear demand is commonly used for algebraic convenience, theoretical models frequently rely on concave demand (Tirole (1989, Section 1.1)). The strategic substitutability of p^w and τ arises for other demand systems more broadly. For instance, equation (5) indicates that $\eta(\tau)$ is negative when $\kappa(p^r) \in [0, 1)$ and the monopolist prices on the elastic region of demand. The curvature condition of $\kappa(p^r) < 1$ describes the class of log-concave demand functions, including both concave and somewhat convex demand functions.[7] Even for demand systems with higher curvature, with $\kappa(p^r) \in [1, 2)$, it is possible for $\eta(\tau)$ to be negative, depending on the relative magnitudes of $(1 - \varepsilon(p^r)^{-1})$ and $\kappa(p^r)$.[8] Isoelastic demand which is common feature of consumer demand models in the macroeconomic literature (e.g., Dixit–Stiglitz CES preferences) is a limiting case where $\kappa(p^r) = (1 - \varepsilon(p^r)^{-1})$, and thus $\eta(\tau) = 0$. In this case, firms (by assumption) do not alter their pricing decisions in response to changes in tax policy. As log-concavity characterizes the majority of demand specifications commonly used in economic analysis, however, we focus on demand systems with $\kappa(p^r) < 1$ going forward. This includes the empirically relevant discrete choice models of demand based on Type I extreme value distributed errors that we rely on in our empirical analysis below (see Fabinger and Weyl (2016, Appendix 3)).

Having established the interaction between the monopolist's price and the tax rate, we now explore how the retail price paid by the consumer and tax revenue vary with the tax rate τ. To highlight the role of the firm response in these relationships, we contrast outcomes under a fixed and under an optimally chosen wholesale price p^w. Consider first the tax elasticity of the retail price, which we denote by $\psi(\tau)$, with $\psi(\tau) \equiv \frac{dp^r}{d\tau} \cdot \frac{\tau}{p^r}$. Relying

[7]If $\kappa(p^r) < 1$, it follows from the definition of curvature that $D''(p^r)D(p^r) - [D'(p^r)]^2 < 0$, which is the condition for demand to be log-concave.

[8]We restrict attention to demand systems with $\kappa < 2$ since $\kappa(p^r) \in [0, 2)$ ensures that the revenue function $R(p^r) = p^r D(p^r)$ is concave in p^r, or equivalently, that the marginal revenue function is decreasing, a common demand restriction in models of imperfect competition.

on the dependency of the retail price on the tax rate in equation (1) and the wholesale price response in equation (4), the tax elasticity of the retail price is given by

$$\psi(\tau) = \frac{\tau}{1+\tau} + \eta(\tau). \tag{7}$$

The firm's incentive to reduce price as the tax rate rises thus reduces the retail price responsiveness to tax rate changes relative to the case where wholesale price does not respond. This results in a more muted quantity response by the consumer, highlighting the firm's ability to affect tax revenue receipts through its price choice.

The government's tax revenue function is given by

$$T(\tau) = \left(p^r - p^w\right)D\left(p^r\right) = \tau p^w D\left((1+\tau)p^w\right), \tag{8}$$

and the effect of a change in the tax rate on tax revenue is

$$\begin{aligned}\frac{dT(\tau)}{d\tau} &= p^w D\left(p^r\right) + \tau p^w D'\left(p^r\right)p^w + \frac{dp^w}{d\tau}\left(\tau D\left(p^r\right) + \tau p^w D'\left(p^r\right)(1+\tau)\right) \\ &= p^w D\left(p^r\right)\left[\left(1 + \frac{\tau}{1+\tau}\varepsilon\left(p^r\right)\right) + \eta(\tau)\left(1+\varepsilon\left(p^r\right)\right)\right].\end{aligned} \tag{9}$$

Tax revenue is thus not monotonic in the tax rate; the sign of $dT(\tau)/d\tau$ depends on the relative magnitudes of the tax rate, the equilibrium demand elasticity $\varepsilon(p^r)$, and the equilibrium pre-tax price response elasticity $\eta(\tau)$. Certain tax rates fall into what Arthur Laffer called the "Prohibitive Range," when an increase in the tax rate leads to a reduction in tax revenue. $dT(\tau)/d\tau$ is negative when

$$\frac{dT(\tau)}{d\tau} < 0 \quad \Longleftrightarrow \quad 1 + \frac{\tau}{1+\tau}\varepsilon\left(p^r\right) + \eta(\tau)\left(1+\varepsilon\left(p^r\right)\right) < 0. \tag{10}$$

Equation (10) highlights that when—akin to perfect competition—the pre-tax price p^w is fixed and $\eta(\tau) = 0$, being in the prohibitive range of the Laffer curve requires a sufficiently elastic consumer demand at the tax-inclusive equilibrium price:

$$\varepsilon\left(p^r(\tau)\right) < \varepsilon^\circ(\tau) = -\frac{1+\tau}{\tau}. \tag{11}$$

For instance, for a tax rate of 50% (similar to the 53.4% tax charged in Pennsylvania) to be in the prohibitive region of the Laffer curve, the demand elasticity $\varepsilon(p^r)$ at the corresponding equilibrium tax-inclusive price needs to be below $\varepsilon^\circ(0.5) = -3$. For many more moderate taxation schemes, such as sales taxes, which in the U.S. reached only 9.45% across states in 2015,[9] demand for the affected products is unlikely to be sufficiently elastic for the observed tax rates to be near or beyond the peak of the Laffer curve; the critical value for the demand elasticity is -11 when $\tau = 10\%$. This elasticity-based mechanism for the Laffer curve in the absence of a strategic supply response resembles the incentive mechanism relied on in the macro literature to generate a Laffer curve: higher income taxes reduce workers' labor supply to eventually sufficiently reduce the labor tax base for income tax revenues to fall. See Trabandt and Uhlig (2011, Proposition 2).

[9] See https://taxfoundation.org/state-and-local-sales-tax-rates-2015/.

Consider now the more realistic case where the monopolist re-optimizes its pricing decision after a change of the tax rate. Substituting (5) for $\eta(\tau)$ in condition (10), the prohibitive range of the Laffer curve arises when

$$\frac{2-\kappa(p^r)+2\tau}{\tau}+\varepsilon(p^r)+\frac{1}{\varepsilon(p^r)}<0. \tag{12}$$

For this inequality to hold over the elastic range of demand where the monopolist prices, it suffices that

$$\varepsilon(p^r)<\varepsilon^\star(\tau,\kappa)=-\frac{2-\kappa(p^r)+\tau}{\tau}. \tag{13}$$

How does the implicit relationship between the tax rate and the elasticity of demand described by equation (13) compare to the case when $\eta(\tau)=0$? When demand is log-concave and $\kappa(p^r)\in[0,1)$, as in our empirical setting, we can show that for any given tax rate τ:

$$\varepsilon^\star(\tau,\kappa)=-\frac{2-\kappa(p^r)+\tau}{\tau}<-\frac{1+\tau}{\tau}=\varepsilon^\circ(\tau). \tag{14}$$

When allowing for a strategic response by the monopolist to chosen tax rates, demand at tax-inclusive retail prices thus needs to be *more elastic* than under fixed wholesale prices for a tax rate increase to push tax revenues down the slippery slope of the Laffer curve. For example, for the above tax rate of 50% to be in the prohibitive region, the demand elasticity at the resulting tax-inclusive prices now needs to be less than -5 or -4 for $\kappa(\tau)$ equal to 0 or 0.5, respectively. This compares to the above critical value for the elasticity of -3 when $\eta(\tau)=0$. For a given tax rate, the difference between the two elasticity cutoff values converges to zero as demand becomes more convex, reflecting that $\varepsilon^\star(\tau,\kappa)\to\varepsilon^\circ(\tau)$ as $\kappa(p^r)\to 1$; it is highest when demand is linear and $\kappa(p^r)=0$.

This result also implies that the tax revenue maximizing tax rate may be higher when accounting for the reduction in pre-tax price by the monopolist than when the pre-tax price does not adjust. Expressing equation (14) in terms of the optimal tax rate given the demand responsiveness at the tax-inclusive prices results in

$$\tau^\star(\varepsilon,\kappa)=-\frac{2-\kappa(p^r)}{1+\varepsilon(p^r)}>-\frac{1}{1+\varepsilon(p^r)}=\tilde{\tau}^\circ\big(\varepsilon(p^r)\big). \tag{15}$$

Note that $\tilde{\tau}^\circ(\varepsilon)$ is not necessarily equal to $\tau^\circ(\varepsilon)$, the optimal tax rate with fixed p^w, since the demand elasticities are evaluated at different tax-inclusive retail prices, $p^r=(1+\tau^\star)p^w$ when p^w responds and $p^r=(1+\tau^\circ)p^w$ with fixed p^w. Locally, with small changes in the tax-inclusive retail price, however, and supported by the empirical evidence we present below, the difference between $\tilde{\tau}^\circ(\varepsilon)$ and $\tau^\circ(\varepsilon)$ is sufficiently small that $\tau^\star(\varepsilon,\kappa)>\tau^\circ(\varepsilon(p^r))$.

Moving from the case when $\eta(\tau)=0$ to one where $\eta(\tau)<0$ may not only shift the Laffer curve to the right, but may also make it flatter. The firm response adds the last term in equation (9), $\eta(\tau)(1+\varepsilon(p^r))$, which is positive because of the strategic substitutability of the tax rate and the wholesale price, $\eta(\tau)<0$. Thus, the stronger the firm response, the flatter the Laffer curve becomes.

In summary, the theoretical model suggests that the downward sloping part of the Laffer curve arises if demand is sufficiently elastic relative to the tax rate and the tax rate elasticity of the monopolist's chosen price. Even in the absence of a price response, there is a role for empirical analysis in determining the elasticity of demand under alternative tax-inclusive prices to characterize the tax revenue function. Accounting for the firm's price response further requires estimates of firm market power via $\eta(\tau)$. Our analysis of Pennsylvania's spirit pricing therefore allows us to empirically assess the equilibrium responses of upstream firms and consumers.

We could extend this analysis to various homogeneous good oligopoly models along the lines of the framework for analyzing tax incidence put forth in Weyl and Fabinger (2013) though theoretical results do not exist for the differentiated products we study in our empirical application (Fabinger and Weyl (2016, Appendix E)). Here, the main challenge lies in the fact that a tax rate increase leads to substitution not only to the outside option, but also to other taxed products; the resulting overall change in tax revenue reflects varying changes in product sales due to heterogeneity in products' costs and characteristics. Further, firms and consumers respond differentially to a tax rate change based on variation in market power and preferences, respectively. We account for these effects empirically in characterizing the Laffer curve across all spirits products offered in Pennsylvania. As in the analysis here, we are particularly interested in comparing the tax revenue expected by a naïve regulator who mistakenly expects firm prices to remain fixed after a tax change, $\eta(\tau) = 0$, to that realized by an agency that correctly anticipates firm responses to its actions, $\eta(\tau) < 0$.

3. PLCB PRICING AND SALES DATA

We now describe our data and the institutional details that inform our theoretical modeling and econometric specification. Most pertinent are Pennsylvania's regulations governing retail prices of spirits products and the frequency and duration of temporary wholesale price adjustments. We also explore the nature of competition in the upstream distiller market that mitigates the effectiveness of such pricing regulations. Finally, we document the heterogeneity of consumer preferences for different types of spirits.

3.1. *Data: Quantities Sold, Prices, and Characteristics of Spirits*

We obtained store-level panel data from the Pennsylvania Liquor Control Board (*PLCB*) under the Pennsylvania Right-to-Know Law. The data contain daily information from 2002 to 2004 on quantities sold and gross receipts at the UPC and store level for all spirits and wines, together with each product's wholesale price according to well-defined pricing periods.[10]

As of January 2003, the *PLCB* operated a system of 593 state-run retail stores across the state.[11] We drop wholesale and outlet stores and combine sales of stores in the same zip code, that results in a total of 456 local markets. We link store locations with data on local population and demographic characteristics using the 2000 Census. The *PLCB* opened and closed several stores over the time period of our sample. We take these entry/exit decisions as exogenous shifts in the demographic composition of the potential pool of each store's consumers.[12]

[10]Pennsylvania allows for the controlled entry of private retailers in the the sale of beer.

[11]See Seim and Waldfogel (2013) for an account of the welfare losses induced by the limited store entry allowed in the Pennsylvania market.

[12]Appendix A of the Supplemental Material (Miravete, Seim, and Thurk (2018)) describes how we construct demographic variables based on assigning Census block groups to their closest store and how we deal with store

Each store carries a multitude of products. We focus on the spirits category as it represents the majority of *PLCB* sales at 60.8% of system revenue. Spirits further constitute a well-defined and mature product category with a small number of easily measurable product characteristics, including the type of spirit, the alcohol content, the possible addition of fruit or other flavors, and the product's country of origin.[13]

We restrict our sample to 375 ml, 750 ml, and 1.75 L bottles of spirits products, representing 80.9% of total spirit category sales by volume and 91.6% by revenue. Many of these products are available across stores but are rarely purchased. We therefore only include the highest selling products that together account for 80% of sales in each bottle size and spirit type category. Consequently, a popular 750 ml bottle of E&J Brandy (average retail price of $9.95) is in our final sample, while a rare 750 ml bottle Remy Martin Louis XIII Cognac (average retail price of $1,078) is not.[14] We also exclude tequila as a segment, as it accounts for less than 2% of sales. Together, these two restrictions allow us to drop a total of 1,240 products from our sample. Our final sample consists of 3,377,659 observations of market and time-period level purchases of 312 products that span brandy, cordials, gin, rum, vodka, and whiskey for three bottle sizes. The final sample represents 56.8% and 63.2% of total bottle sales and spirits revenue from off-premise consumption outside of a restaurant or bar, respectively.

Table I reports the number, prices, and characteristics of products in our sample, both in aggregate and by type of spirit, which highlight significant product differentiation. The average proof is 75.33; 37.40% of products are imported; and 16.3% of products contain flavor add-ins.[15] Vodkas and whiskeys have significantly larger market shares (31.88% and 24.41%, resp.) than rum (16.18%), cordials (13.38%), brandy (7.24%), or gin (6.91%). The differences in product variety within each category mirror the differences in market shares, with only approximately one-half as many brandy and gin varieties as vodkas while 28.9% of the products are whiskeys. Flavored products are primarily cordials and brandies and to a lesser extent vodkas and rums. We also see variation in domestic versus imported varieties across spirit types: 58.89% of whiskeys and 51.61% of cordials in our sample are imported, but imported products comprise less than half of the products of the other spirit types. We complement these product characteristics obtained from the *PLCB* with data on spirit product quality from Proof66.com, a company that aggregates spirits ratings into a single quality score for each rated product. The quality score is largely informative within, but not across, spirit types, and we therefore do not report it in Table I.

We denote spirits as expensive when their simple averaged price exceeds the mean price of other spirits of the same type and bottle size. Table I shows that expensive spirits are purchased nearly as often as cheaper varieties, but are less likely to be flavored or domestically produced and have higher proof. The 750 ml bottle is the most popular size

openings and closings. We also document that the large majority of spirits are sold at every store. This alleviates concerns about assortment differences between stores leading to potential competition for consumers between stores.

[13]In contrast, wines have hard-to-measure quality determinants and a large number of products with limited life cycles leading to tiny, highly volatile market shares. For example, within the popular 750 ml bottle category, the top-100 selling wines (out of 4,675) constitute only 45% of total 750 ml wine revenue.

[14]Similarly, a 375 ml bottle of Captain Morgan could be excluded from the sample if its sales rank among 375 ml bottles is too low while other bottle sizes of Captain Morgan are included. In practice, this did not occur.

[15]In 16th century England, if a pellet of gunpowder soaked in a spirit could still burn, the spirit was determined to be "proof," and thus taxed at a higher rate. Gunpowder soaked in rum will ignite only if the alcohol by volume exceeds 57.15%. To simplify, since 1848 in the U.S., a 100 proof corresponds to a spirit with 50% alcohol by volume content. See Jensen (2004).

TABLE I

PRODUCT CHARACTERISTICS BY SPIRIT TYPE[a]

	Products	Price	Share	% Flavored	% Imported	Proof
By Spirit Type:						
BRANDY	26	14.37	7.24	30.77	26.92	76.15
CORDIALS	62	14.02	13.38	32.26	51.61	55.82
GIN	28	15.03	6.91	3.57	28.57	83.42
RUM	40	13.61	16.18	10.00	17.50	74.03
VODKA	66	16.75	31.88	21.21	40.91	81.60
WHISKEY	90	16.67	24.41	0.00	58.89	80.98
By Price and Size:						
EXPENSIVE	150	20.30	46.89	12.00	64.67	77.82
CHEAP	162	10.92	53.11	17.90	22.84	72.46
375 ml	48	8.89	15.19	8.33	47.92	75.10
750 ml	170	14.44	50.20	21.76	44.71	72.95
1.75 L	94	20.56	34.61	6.38	37.23	78.77
ALL PRODUCTS	312	15.07	100.00	16.30	37.40	75.33

[a]"Price" is the simple average price between 2002 and 2004 and across products in each category. "Share" is based on number of bottles sold. "Cheap" ("Expensive") products are those products whose mean price is below (above) the mean price of other spirits in the same spirit type and bottle size.

of product in terms of unit sales and product variety, accounting for 50.20% of bottles sold and 54.5% of available spirits products, closely followed by the 1.75 L bottle with a share of 34.61% of bottles sold and 30.1% of available spirit products. The smallest bottles we consider, those in the 375 ml format, account for 15.2% of bottles sold and 15.4% of spirit varieties.

Finally, upstream firms produce brands in particular bottle sizes. For instance, our final sample is composed of 198 brands (e.g., Captain Morgan) but 88 of these brands are available only in the 750 ml bottle size while 1 and 31 brands come only in the 375 ml and 1.75 L size, respectively. The remaining 78 brands were offered in several bottle sizes (e.g., Diageo sold Captain Morgan in 375 ml, 750 ml, and 1.75 L sizes).[16]

3.2. *The Mechanics of the Pricing Regulation*

The *PLCB* acts as a monopolist in the retail distribution of wine and spirits; the Pennsylvania State Legislature exerts regulatory oversight over several aspects of the daily operations of the stores. Most notably, as per Pennsylvania Liquor Code (47 P.S. Section 1–101 *et seq*.) and Pennsylvania Code Title 40, the legislature imposes a uniform pricing formula with a constant wholesale price markup that the *PLCB* applies both across products and across stores. Prices of spirits are thus identical across the state at a point in time and follow a common pricing/taxation rule known to all consumers and upstream manufacturers.

The legislature has modified this rule only infrequently over the years. From 1937 until 1980, the retail price for all products reflected a 55% markup over wholesale cost for all gins and whiskeys and a 60% markup for other spirits. In 1980, the legislature reduced

[16]This pattern is reflected in the raw data as well where 958 of the potential 1,192 brands were offered only in one bottle size, usually the 750 ml format. Table B.I in Appendix B of the Supplemental Material (Miravete, Seim, and Thurk (2018)) provides additional descriptive statistics on the distribution of spirit prices by type and size of bottle.

the markup to 25% for all products, but introduced a per-unit handling fee, the *Logistics, Transportation, and Merchandise Factor* (*LTMF*), of initially \$0.81, rising to \$0.85 by 1982. The legislature instituted the current 30% markup in 1993 when it also modified the unit fee to vary by bottle size to better reflect transportation costs from the *PLCB*'s centralized warehouses to the retail stores. The *LTMF* unit fee for the 375 ml, 750 ml, and 1.75 L bottles in our sample amounts to \$1.05, \$1.20, and \$1.55, respectively. For the average product, the *LTMF* fee accounts for 26.7% of the final retail markup. In addition, consumers also have to pay an 18% sales tax, the "Johnstown Flood Tax," on all liquor purchases.[17] Accordingly, the retail price p^r of a given product with wholesale price, p^w, is calculated as[18]

$$p^r = \left[p^w \times 1.30 + LTMF\right] \times 1.18. \tag{16}$$

Of primary concern for this paper is the uniform markup, an ad valorem tax, applied to all products, amounting to $(1.30 \times 1.18 - 1)$, or 53.4%.

The *PLCB* has limited ability to depart from this uniform percent markup rule. It operates seven outlet stores close to the state borders in an effort to mitigate any *border bleed* of consumers who illegally import lower-priced products into Pennsylvania from neighboring states. While these stores offer wines and spirits at discounted prices, the *PLCB* remains within the uniform markup policy by primarily selling products in the outlet stores not found in regular stores, for example, multipacks or unusual bottle sizes for a particular product. Controlling for these stores has little qualitative or quantitative effect on our results. Related robustness checks are reported in Appendix C.2 of the Supplemental Material (Miravete, Seim, and Thurk (2018)).

The *PLCB* purchases bottles of spirits directly from upstream distillers at wholesale prices p^w. Because of the legislated pricing formula, retail price p^r is driven by the wholesale pricing decisions p^w of the *PLCB*'s suppliers and any change in the wholesale price results in a change to the retail price passed on to consumers.

Wholesale prices can change for only two reasons. First, for most products, distillers can temporarily modify the wholesale price at set intervals that we denote as "pricing periods." Such temporary price changes—generally price reductions—last four or five weeks and typically coincide with the month of year. The *PLCB* places some limitations on the frequency of temporary price changes: distillers can put products on sale up to four times a year, or once per quarter. A product can thus go on sale for one pricing period, but not for two in a row. Distillers also need to submit any proposed sale prices to the *PLCB* at least five months before the start of the promotion.

Second, upstream firms can permanently change the wholesale price of a product, thereby also changing the reference price for temporary price changes. Permanent price changes typically take effect at the beginning of the first of four-week long intervals that *PLCB* accounting rules employ as subdivisions of quarters. Similar to temporary price reductions, distillers need to request permanent price increases (but not permanent price decreases, which are effective immediately) with lead time, by the start of one quarter prior to the desired date of the price increase. The periodicity of permanent price changes is therefore slightly different from that of temporary price changes. Since temporary wholesale price changes account for 84.8% of price changes in our sample, we

[17]The original 10% tax was instituted in 1936 to provide \$41 million for the rebuilding of the flood-ravaged town of Johnstown. Despite reaching the funding goal after the initial six years, the tax was never repealed, but instead rose to 15% in 1963 and to 18% in 1968.

[18]An additional 6% sales tax is then applied to the posted price to generate the final price paid by the consumer.

TABLE II

PERCENT OF PRODUCTS PLACED ON SALE OVER THE YEAR[a]

	Spring	Summer	Fall	Winter	Holiday	Year	Times
By Spirit Type:							
BRANDY	30.77	50.00	34.62	26.92	34.62	59.26	2.37
CORDIALS	40.32	48.39	30.65	45.16	43.55	61.29	2.35
GIN	46.43	39.29	50.00	39.29	39.29	63.64	2.24
RUM	47.50	40.00	50.00	32.50	42.50	57.45	2.12
VODKA	50.00	60.61	57.58	39.39	50.00	76.81	2.24
WHISKEY	58.89	51.11	48.89	42.22	53.33	65.71	2.51
By Price and Size:							
EXPENSIVE	51.33	59.33	50.67	47.33	56.00	75.44	2.22
CHEAP	45.68	41.36	41.98	32.10	37.65	55.23	2.50
375 ml	14.58	18.75	20.83	8.33	6.25	30.91	1.39
750 ml	50.59	53.53	45.88	46.47	51.18	71.66	2.14
1.75 L	61.70	59.57	59.57	42.55	58.51	72.28	2.91
ALL PRODUCTS	48.40	50.00	46.15	39.42	46.47	65.31	2.34

[a]"Cheap" ("Expensive") products are those products whose mean price is below (above) the mean price of other spirits in the same spirit type and bottle size. We define the "Holiday" season as the two pricing periods that encompass Thanksgiving through the end of the year. Statistics reflect the percent of products with a temporary price reduction during the corresponding season except for "Times," which denotes the average number of times that spirits in each category are on sale during a year.

aggregate daily data on prices and quantity sold to the level of pricing periods, resulting in 34 periods from 2002 to 2004. Note that the delay between the request and effectiveness of either permanent or temporary price adjustments limits the ability of the distillers to respond to temporary demand shocks—an issue we revisit when discussing price endogeneity concerns in Section 5. We discuss the periodicity of the price series further in Appendix A of the Supplemental Material (Miravete, Seim, and Thurk (2018)).

Table II presents descriptive statistics for changes in temporary price. First, distillers temporarily change a product's price 2.3 times a year on average. While not all products experience a temporary price change, the majority do; 65.31% of spirits are on sale at least once in a given year. This is true across spirit types, with distillers changing the price of vodkas, expensive varieties, and all but the smallest bottles more frequently than the rest. There is also a seasonal pattern of price changes across spirit types as distillers are more likely to change a product's price during the summer and less likely during the winter. Over the holidays, defined as pricing periods that overlap with Thanksgiving through the end of the year, distillers place 46.47% of our spirit products, ranging from 34.62% of brandies to 53.33% of whiskeys, on sale at least once, but change the price of 375 ml bottles rarely. The combination of variation in monthly price changes, both temporary and permanent, and differences in the amount of the price changes is the primary source of price variation that we exploit in the estimation of our demand model.

3.3. *The Upstream Distillers*

While the distiller market saw several large mergers post-2004, there were no mergers or acquisitions of relevance in the distiller segment during the sample period and the market is relatively unconcentrated. Table III shows that the market leader, Diageo, accounted for 21.60% of revenues and 24.48% of bottle sales in the sample; more than double the size of its two largest competitors at the time: Bacardi and Beam. Meanwhile,

TABLE III

THE UPSTREAM MARKET[a]

Firm	Products	Share of Spirit Market		Top Selling Product	
		By Revenue	By Quantity	Name	Type
Diageo	63	21.60	24.48	Captain Morgan	Rum
Bacardi	22	8.92	9.79	Bacardi Light Dry	Rum
Beam	32	9.86	9.01	Windsor Canadian	Whiskey
Other Firms (34)	195	59.62	56.72	SKYY (Campari)	Vodka

[a]Upstream distillers sorted in descending order according to quantity (bottles) share.

34 smaller firms collectively hold significant market share: 59.62% share of revenue and 56.72% of bottle sales.

The largest upstream firms—Diageo, Bacardi, and Beam—operate product portfolios that extend into all spirit types and bottle sizes while 21 of the 34 smaller distillers operate product portfolios of less than five products and eight are single product firms. There is, however, substantial heterogeneity in product offerings even among the top three distillers. For example, Diageo has a relatively balanced portfolio where rums, vodkas, and whiskeys generate 19.6%, 31.8%, and 24.4% of its total revenue, respectively. In contrast, 70.2% of Bacardi's revenue comes from rums compared to just 4.1% for Beam.[19]

While the overall market appears competitive with a Herfindahl–Hirschman Index (HHI) based on bottle sales of only 930.3, distillers have more market power in some regions of the product space than others. For example, the HHI is approximately 3,000 for rums, while the brandy and gin segments are moderately concentrated with HHIs around 2,000. The cordial, vodka, and whiskey segments are the most competitive with low concentration measures (all less than 1,400). Horizontal differentiation of products within a spirit class would provide further market power. An accurate characterization of the response of distillers to changes in government tax policy therefore requires estimation of patterns in consumer preferences that motivate the observed extent of differentiation across and within these different product segments.

3.4. *Evidence of Preference Heterogeneity*

To conclude the description of the primary data features, we now document systematic variation in consumer preferences along different demographic profiles. Throughout the analyses, we rely on four primary demographic attributes of stores' market areas: income, educational attainment, and the prevalence of minority and young consumers. We use categorical data on income by minority status to fit generalized beta distributions. These allow us to draw random samples of income for our estimation from market-specific continuous distributions that vary by demographic trait, and to estimate the share of high-income households (incomes above $50,000). We similarly obtained information on educational attainment by minority status to derive the share of the minority and white

[19]Table B.II in Appendix B of the Supplemental Material (Miravete, Seim, and Thurk (2018)) reports the market shares by spirit type, bottle size, and price range. Bacardi acquired Grey Goose from Sidney Frank in August 2004. Since *PLCB* requires a five-month advance notice for any temporary price reduction, we assume that Sidney Frank manages the brand in our estimation and counterfactuals.

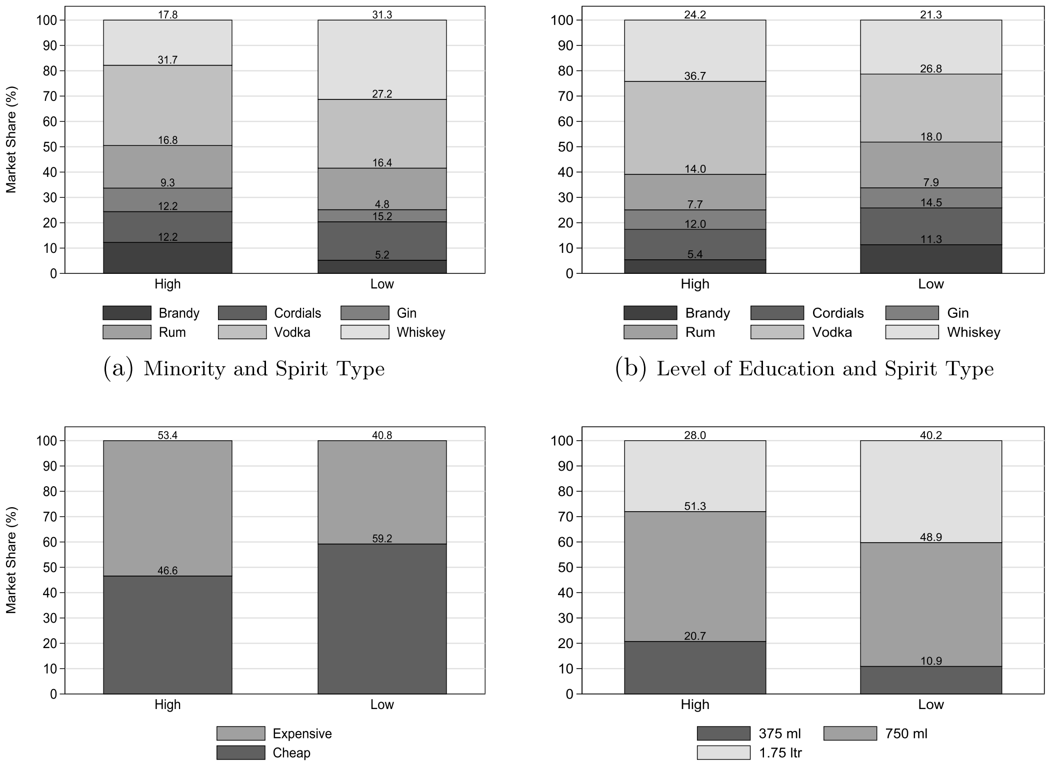

(a) Minority and Spirit Type

(b) Level of Education and Spirit Type

(c) Income and Price

(d) Age and Bottle Size

FIGURE 1.—Spirit consumption and demographics. Notes: We compare markets in the top ("High") and bottom ("Low") quintile of each demographic trait. "Minority" is the percent of the market population who identifies as "nonwhite." "Education" is defined as the percent of the population with some college experience. "Income" is the percent of households with high income, defined as greater than $ 50,000. "Age" is the percent of the population between 21 and 29 years of age.

population with at least some college education in each area. Lastly, we employ the unconditional share of each market's population between the ages of 21 and 29.

We show differences in preferences by assigning the store markets into quintiles based on each demographic trait—the share of high-income households, the share of nonwhite or minority households, the share of residents with some college education, and the share of residents in their twenties. Figure 1 compares the purchase patterns of the top and bottom quintiles. Markets with a greater share of minorities have substantially higher sales of vodka, gin, and brandy but lower sales of whiskey. In areas where the population has more residents with college experience, however, not only vodka but also whiskey is more popular, while rum and brandy have lower sales. In markets with a larger share of high income residents, we observe larger purchases of expensive spirits indicating that high income consumers are more willing to buy expensive spirits, presumably because these spirits tend to be of higher quality.[20] Finally, as the share of young residents increases, so do sales of smaller bottles.

[20]The Proof66.com data confirm that price and quality are positively correlated, particularly for cordials, gins, and whiskeys.

Our analysis exploits this wide variation in observable preferences across demographic groups in three ways. First, this preference heterogeneity allows us to capture rich substitution patterns to best reflect the purchase decisions of consumers, leading to more robust elasticity estimates both across products and for spirits as a whole. This, in turn, results in more accurate estimates of upstream firm market power, and thus, the distillers' ability to respond to changes in the downstream tax rate ($\eta(\tau)$ of the model in Section 2). Second, the large amount of heterogeneity among products and distillers combined with heterogenous consumer preferences suggests that the upstream response to changes in tax policy will vary by firm and product leading to heterogenous effects across consumers.

Finally, the wide variation in observed preferences provides us with an opportunity to explore the characteristics of the Laffer curve, such as location and shape, across consumer segments. We confirm that an inverted U-shaped curve characterizes the relationship between the tax rate and tax revenue for any particular demographic group. The statewide Laffer curve we present is thus not an inadvertent outcome of having aggregated the non-Lafferian responses of heterogeneous consumers to a change in tax policy.

4. EMPIRICAL MODEL

In this section, we describe a static model of oligopoly price competition with differentiated goods. We assume that each period upstream spirit manufacturers simultaneously choose wholesale prices (p^w) to maximize profits. The downstream firm, the *PLCB*, takes these prices as given and generates the final retail price by applying a single markup and a per-unit handling fee that varies by bottle size. Finally, consumers in each market choose the product that maximizes their utility. We solve the model backwards, first presenting downstream consumer demand and then progressing to the profit-maximization problem of the upstream spirit manufacturers.

4.1. *Downstream Market—A Discrete Choice Model of Demand for Spirits*

In modeling demand for spirits as a function of product characteristics and prices, we follow the large literature on discrete-choice demand system estimation using aggregate market share data (e.g., Berry, Levinsohn, and Pakes (1995) (*BLP*) and Nevo (2001)). In pricing period t, consumer i in market l obtains the indirect utility from consuming a bottle of spirit $j \in J_{lt}$ given by

$$\begin{gathered} u_{ijlt} = x_j \beta_i^* + \alpha_i^* p_{jt}^r + [h_t \quad q3_t]\gamma + \xi_{jlt} + \overline{\epsilon}_{ijlt}, \\ \text{where } i = 1, \dots, M_{lt};\ j = 1, \dots, J_{lt};\ l = 1, \dots, L;\ t = 1, \dots, T. \end{gathered} \tag{17}$$

The $n \times 1$ vector of observed product characteristics x_j is fixed over time, though the availability of different products changes over time due to product introductions/removals or store closings/openings. The holiday dummy variable h_t indicates whether period t coincides with the end-of-year holiday season from Thanksgiving to the New Year, while the summer dummy variable $q3_t$ captures periods which overlap with the months of July, August, and September. We denote the retail price of product j at time t by p_{jt}^r; it is constant across markets l at point t. We further allow utility to vary across products, markets, and time via the time and location-specific product valuations ξ_{jlt}, which are common knowledge to consumers, upstream firms, and the *PLCB* but unobserved by the econometrician.[21]

We characterize consumer i in market l by a d-vector of observed demographic attributes D_{il} including education, race, age, and income. To allow for individual heterogeneity in purchase behavior and alleviate the restrictive substitution patterns generated by the Independence of Irrelevant Alternatives (IIA) property of the multinomial logit model, we assume that the distribution of consumer preferences over product characteristics and prices follows a multivariate normal distribution:

$$\begin{pmatrix} \alpha_i^* \\ \beta_i^* \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_{il} + \Sigma \nu_{il}, \quad \nu_{il} \sim N(0, I_{n+1}), \tag{18}$$

where Π is a $(n+1) \times d$ matrix of coefficients that measures the effect of observable individual attributes on the consumer valuation of spirit characteristics, while Σ measures the covariance in unobserved preferences across characteristics. We restrict $\Sigma_{jk} = 0\ \forall k \neq j$, and estimate only the variance in unobserved preferences for characteristics.

We follow Grigolon and Verboven (2014) in assuming that the unobserved individual preferences for products, $\overline{\epsilon}$, are correlated across spirits of the same type. We observe $g = 6$ distinct spirit types in the data (brandy, cordials, gin, rum, vodka, and whiskey) and define group zero to be the outside good. In the resulting random coefficient nested logit model or *RCNL*, we can decompose the idiosyncratic valuation into

$$\overline{\epsilon}_{ijlt} = \zeta_{igt} + (1 - \rho)\epsilon_{ijlt}, \tag{19}$$

where the "nesting parameter" $\rho \in [0, 1]$, ϵ is distributed i.i.d. multivariate type I extreme value, and the distribution of ζ_{igt} is such that $\overline{\epsilon}_{ijlt}$ is also distributed extreme value.

Together, equations (18) and (19) encompass a range of demand specifications. When $\Sigma = 0$ and $\rho > 0$, the model collapses to the nested logit; as ρ approaches one, consumers view products within each spirit type as perfect substitutes. Alternatively, when $\rho = 0$ but $\Sigma > 0$ the model collapses to the random coefficients model of *BLP*. When both $\Sigma = 0$ and $\rho = 0$, the model returns the simple multinomial logit choice probabilities.

We assume that in each time period, a consumer selects either one bottle of the J_{lt} spirits available in her market l or opts to purchase the outside option denoted by $j = 0$. We define the potential market, M_{lt}, as all purchases of alcoholic beverages, including spirits, beer, and wine, for off-premise consumption.[22] According to Haughwout, Lavallee, and Castle (2015), the average drinking-age Pennsylvanian consumed 124.2, 120.5, and 121.0 liters of alcoholic beverages in 2002, 2003, and 2004, respectively, of which off-premise purchases accounted for 79.8%. Beer makes up 90% of total consumption by volume: the average drinking-age Pennsylvanian consumes the equivalent of nearly five 375 ml bottles of beer per week, but only approximately four 750 ml bottles of wine and spirits during the

[21]Our assumption that firms observe the full distribution of consumer preferences is a simplification to abstract from second-degree price discrimination within brand by bottle size. Accounting for such information asymmetries would require a demand model with less flexible substitution patterns than the ones that motivate the observed horizontal product differentiation.

[22]Hendel (1999) highlighted that the present static discrete choice approach has limitations when individuals purchase several products or multiple bottles of the same product at the same time. Hendel and Nevo (2006) further showed that static demand estimates overestimate own-price elasticities and underestimate cross-price elasticities when consumers make dynamic purchase decisions. We discuss potential issues and biases associated with stockpiling in Appendix C.5 of the Supplemental Material (Miravete, Seim, and Thurk (2018)) and provide evidence that suggests stockpiling is not an issue in our data.

year. The potential market M_{lt} for location l in period t is simply the prorated potential off-premise consumption per-capita based on the length of pricing period t scaled by the market's population over the age of 21, and the outside option represents beer and wine purchases expressed in 750 ml bottle equivalents.

Given these preliminaries, the set of individual-specific characteristics that lead to the optimal choice of spirit j is given by

$$A_{jt}\big(p^r_{\cdot t}, x_\cdot, \xi_{\cdot t}; \theta\big) = \big\{(D_{il}, \nu_{il}, \epsilon_{\cdot lt}) | u_{ijlt} \geq u_{iklt}\ \forall k = 0, 1, \dots, J_{lt}\big\}, \tag{20}$$

where we summarize all model parameters by θ. We decompose the deterministic portion of the consumer's indirect utility in equation (17) into a common part shared across consumers, δ_{jlt}, and an idiosyncratic component, μ_{ijlt}:

$$\begin{aligned} \delta_{jlt} &= x_j\beta + \alpha p^r_{jt} + [h_t \quad q3_t]\gamma + \xi_{jlt}, \\ \mu_{ijlt} &= \big[x_j \quad p^r_{jt}\big](\Pi D_{il} + \Sigma\nu_{il}). \end{aligned} \tag{21}$$

We take advantage of the additive specification of utility to integrate over the distribution of ϵ_{it} giving rise to A_{jt} analytically. The probability that consumer i purchases product j in market l in period t is then

$$s_{ijlt} = \frac{\exp\left(\dfrac{\delta_{jlt} + \mu_{ijlt}}{1-\rho}\right)}{\exp\left(\dfrac{I_{iglt}}{1-\rho}\right)} \times \frac{\exp(I_{iglt})}{\exp(I_{ilt})}, \tag{22}$$

where

$$\begin{aligned} I_{iglt} &= (1-\rho)\ln\sum_{m=1}^{J_g}\exp\left(\frac{\delta_{mlt} + \mu_{imlt}}{1-\rho}\right), \\ I_{ilt} &= \ln\left(1 + \sum_{g=1}^{G}\exp(I_{iglt})\right). \end{aligned} \tag{23}$$

Last, to derive product j's aggregate market share in each location, we integrate over the distributions of observable and unobservable consumer attributes D_{il} and ν_{il}, denoted by $P_D(D_i)$ and $P_\nu(\nu_i)$, respectively. Thus, the model predicts a market share for product j in market l at time t of

$$s_{jlt} = \int_{\nu_l}\int_{D_l} s_{ijlt}\, dP_D(D_i)\, dP_\nu(\nu_i), \tag{24}$$

which we evaluate by Monte Carlo simulation using Halton draws from the empirical distribution of ν. For each market l, we simulate the consumption choices of 200 randomly drawn heterogeneous consumers who vary in their demographics and income according to the empirical distributions from the 2000 Census for market l. We exploit the availability of data on the distribution of income and of educational attainment conditional on minority status in generating demographic attributes for the simulated set of consumers. Since the ambient population of stores changes with store openings and closings over the course of the sample, we allow the simulated set of agents to change accordingly. See

Appendix A of the Supplemental Material for further details (Miravete, Seim, and Thurk (2018)).

4.2. *An Oligopoly Model for Upstream Distillers*

Wholesale prices p^w are the outcome of an upstream market equilibrium given the *PLCB*'s pricing rule and consumer demand for spirits. We now present a flexible model of upstream behavior that places few restrictions on firm conduct while allowing for robust estimates of upstream market power. Each firm $f \in F$ produces a subset J_t^f of the $j = 1, \dots, J_t$ products and faces several competitors from the set F of distillers. In each period t, the upstream firms simultaneously choose the vector of wholesale prices $\{p_{jt}^w\}_{j \in J_t^f}$ to maximize period t profit

$$\max_{\{p_{jt}^w\}} \sum_{j \in J_t^f} \left(p_{jt}^w - c_{jt}\right) \times \underbrace{\sum_{l=1}^{L} M_{lt} s_{jlt}\left(p^r\left(p^w\right), x, \xi; \theta\right)}_{\substack{\text{statewide demand for} \\ \text{product } j \text{ in period } t}}, \tag{25}$$

where c_{jt} denotes the marginal cost of producing spirit j in period t. Given the static nature of the firms' pricing decisions, we omit the period t subscripts for the sake of clarity going forward.[23] Define $s_j(p^r, x, \xi; \theta) = \sum_{l=1}^{L} M_l s_{jl}(p^r, x, \xi; \theta)$ the state-wide demand for product j. Assuming a pure strategy Bertrand–Nash equilibrium in wholesale prices, upstream firm f chooses prices p_j^w $\forall j \in J^f$ to solve the set of first-order conditions

$$s_j\left(p^r\left(p^w\right), x, \xi; \theta\right) + \sum_{m \in J^f} \left(p_m^w - c_m\right) s_m\left(p^r\left(p^w\right), x, \xi; \theta\right) \times \frac{\partial s_m}{\partial p_j^w} = 0. \tag{26}$$

The term $\frac{\partial s_m}{\partial p_j^w}$ is the change in quantity sold for product m in response to a change in the retail price induced by the wholesale price change, of product j. Rearranging equation (26) yields

$$p^w = c - \underbrace{\left[O^w * \Delta^{w\prime}\right]^{-1} \times s\left(p^r\left(p^w\right), x, \xi; \theta\right)}_{\text{vector of wholesale markups}}, \tag{27}$$

where O_t^w denotes the ownership matrix for the upstream firms with element (j, m) equal to one if goods j and m are in J^f. The matrix Δ^w captures changes in demand due

[23]Table II documents that the average product goes on sale only 2.3 times per year; 76.6% of products go on sale three times or less in a year. The *PLCB*'s limits on temporary price reductions per year thus do not constrain upstream pricing for the majority of products, and we do not address any dynamic considerations to the timing of pricing decisions over the course of the year.

to changes in wholesale price:

$$
\begin{aligned}
\Delta^w = \Delta^d \Delta^p &= \begin{bmatrix} \frac{\partial s_1}{\partial p_1^r} & \cdots & \frac{\partial s_1}{\partial p_J^r} \\ \vdots & \ddots & \vdots \\ \frac{\partial s_J}{\partial p_1^r} & \cdots & \frac{\partial s_J}{\partial p_J^r} \end{bmatrix} \begin{bmatrix} \frac{dp_1^r}{dp_1^w} & \cdots & \frac{dp_1^r}{dp_J^w} \\ \vdots & \ddots & \vdots \\ \frac{dp_J^r}{dp_1^w} & \cdots & \frac{dp_J^r}{dp_J^w} \end{bmatrix} \\
&= \begin{bmatrix} \frac{\partial s_1}{\partial p_1^r} & \cdots & \frac{\partial s_1}{\partial p_J^r} \\ \vdots & \ddots & \vdots \\ \frac{\partial s_J}{\partial p_1^r} & \cdots & \frac{\partial s_J}{\partial p_J^r} \end{bmatrix} \begin{bmatrix} 1.534 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 1.534 \end{bmatrix}.
\end{aligned} \tag{28}
$$

Δ^d is the matrix of changes in quantity sold in period t due to changes in retail price with element (r, m) equal to $\frac{\partial s_r}{\partial p_m^r}$. An element (m, j) of Δ^p, the matrix of changes in retail price due to changes in wholesale price, equals $\frac{dp_m^r}{dp_j^w}$.

Villas-Boas (2007) shows that in vertical retail pricing markets, Δ^p can be a complicated object. Here, in contrast, the state's regulation of alcohol sales simplifies and constrains downstream price responses significantly. For example, under the current pricing rule, $\frac{dp_j^r}{dp_j^w}$ is simply 1.534, the uniform markup that translates a change in the wholesale price for product j to a change in the product's retail price, and for $j \neq m$, the retail price for product m does not respond to a change in the wholesale price for product j, resulting in a diagonal price response matrix Δ^p.

We rely on this model of firm behavior in the following ways. First, we use the firms' first-order conditions in equation (27) to back out product-level marginal costs given demand estimates and observed wholesale and retail prices. This enables us to measure wholesale markups, and thus the degree of market power in the upstream market. Second, we evaluate the optimal distiller response to changes in *PLCB* tax policy, which we denote as η above. Third, the upstream model allows us to consider alternative forms of upstream conduct, by varying the definition of the ownership matrix O^w. We exploit this flexibility to explore how different conduct assumptions impact the upstream response to taxes and the shape and location of the Laffer curve.

5. ESTIMATION, IDENTIFICATION, AND DEMAND MODEL RESULTS

Our estimation approach follows Nevo (2001) adapted to the institutional features of the Pennsylvania spirits market. We take advantage of the fact that distillers and the *PLCB* set a single price per product, which allows us to identify the contribution to demand of heterogeneity in tastes across the state separately from the contribution of price as a time varying, common, shifters of demand. We discuss the response of distillers to downstream consumer demand shocks as a potential source of price endogeneity that threatens identification of the price coefficient. We then present the model estimates and the implied demand elasticities and upstream response elasticities.

5.1. *Estimating the Random Coefficients and Demographic Interactions*

We begin with a description of the first of the three stages of our estimation procedure, in which we estimate the contributions of unobserved taste heterogeneity Σ and demographic interactions Π to the deviations from mean utility, μ_{ijlt}, as well as the nesting parameter ρ, controlling for location and product-time fixed effects. We employ a generalized method of moments (*GMM*) estimator to estimate these first stage parameters $\theta_A = \{\Sigma, \Pi, \rho\}$ by interacting a structural demand side error $\omega(\theta_A)$ with instruments Z. To define ω, we rewrite each product's mean utility δ_{jlt} from equation (21) as

$$\begin{aligned} \delta_{jlt} &= x_j\beta + \alpha p^r_{jt} + [h_t \quad q3_t]\gamma + \xi_{jlt} \\ &= \upsilon_l + \upsilon_{jt} + \Delta\xi_{jlt}. \end{aligned} \tag{29}$$

We thus decompose the mean utility into a location fixed effect, a product-time fixed effect, and deviations thereof, $\Delta\xi_{jlt}$. The location fixed effect υ_l captures systematic variation across markets in either the preference for spirits relative to beer and wine or in the local preference for alcoholic beverages in aggregate relative to the Pennsylvania average. In the absence of disaggregate data on purchases of all alcoholic beverages, we use the Pennsylvania average as the potential market size for all locations. To capture seasonality and other variation in tastes over time, we include product-time fixed effects υ_{jt}. These fixed effects also reflect the effect of product characteristics, price, and seasonal buying, $x_j\beta + \alpha p_{jt} + [h_t\ q3_t]\gamma$, on a product's mean utility.

Equation (29) highlights an advantage to our setting: since price does not vary across locations l, we are able to control for its mean contribution to utility via product-time fixed effects, which we then use in a second stage estimation to isolate α. The remaining structural error ω therefore represents deviations in unobserved product valuations within a store, $\Delta\xi_{jlt}$, from these mean product-time valuations after controlling for the average taste for spirits in market l.

We solve for the structural error $\omega(\theta_A) = \Delta\xi_{jlt}$ using the following algorithm. For a given guess at θ_A, we find the mean-utility levels $\delta_{jlt}(\upsilon_l, \upsilon_{jt}, S_{jlt}; \theta_A)$ that set the predicted market share of each product, s_{jlt} in equation (24), equal to the market share observed in the data, S_{jlt}.[24] Following Somaini and Wolak (2015), we use a within transformation of δ to remove the store and product-time fixed effects υ_l and υ_{jt}, leaving only ω. Define Z^+ as the within transformation of the instruments matrix Z; for example, for instrument k, $Z^{+,k}_{jlt} = Z^k_{jlt} - \overline{Z}^k_{jt} - \overline{Z}^k_l$. The *GMM* estimator exploits the fact that at the true value of the parameters $(\Sigma^\star, \Pi^\star, \rho^\star)$, the instruments Z^+ are orthogonal to the structural errors $\omega(\Sigma^\star, \Pi^\star, \rho^\star)$, that is, $E[Z^{+'}\omega(\theta^\star)] = 0$. The *GMM* estimates solve

$$\hat{\theta}_A = \underset{\theta_A}{\arg\min}\big\{\omega(\theta_A)'Z^+W^+Z^{+'}\omega(\theta_A)\big\}, \tag{30}$$

where W^+ is the weighting matrix, representing a consistent estimate of $E[Z^{+'}\omega\omega'Z^+]$.[25] Finding a global solution to such a highly nonlinear problem is difficult. As suggested by

[24]We rely on the contraction mapping outlined in Appendix I of *BLP* and modified for the RCNL model by Grigolon and Verboven (2014, Appendix A). In order to ensure convergence to consistent stable estimates, we follow the advice of Dubé, Fox, and Su (2012, Section 4.2) and set the norm for the mean value contraction equal to 1e-14.

[25]We first assume homoscedastic errors and use $W^+ = [Z^{+'}Z^+]^{-1}$ to derive initial parameter estimates. Given these estimates, we solve for the structural error ω and construct $E[Z^{+'}\omega\omega'Z^+]^{-1}$ as a consistent estimate for W^+.

Dubé, Fox, and Su (2012), we employ the Knitro Interior/Direct algorithm using several initial conditions to ensure robustness of our results.

5.1.1. *Identification of Random Coefficients and Demographic Interactions*

Identification of Σ and ρ comes from correlation between a product's market share and its characteristics relative to other more or less similar products; see Berry and Haile (2014). We construct two instruments similar to those used in Bresnahan, Stern, and Trajtenberg (1997). First, we employ the number of products in the market that share product j's characteristic. For example, to identify random taste variation for imported products, we count, for a given imported product sold in market l, the total number of competing imported products of the same bottle size in that market. Similarly, to identify the nesting parameter, ρ, we use the total number of competing products of the same spirit type in market l, separately for each spirit type. Second, we use each product's quality score and compute the average distance, measured in squared deviations, of product j to other products that share its characteristic. Thus, for the above imported product, this would be the average distance in quality scores from other imported products in location l. This instrument provides additional identifying power since it allows for differential effects of introducing a high-quality imported product—say a Scotch whiskey—into a market with other high-quality imported products versus into a market populated by largely low-quality imported products—say lower quality Canadian whiskeys.

We base identification of Π on correlation between a product's market share in a given store market and the demographics of the population served by each store. We thus interact the above instruments with the prevalence of a given demographic attribute in each market. For example, we identify the differential taste of young households for vodkas by interacting our earlier instruments for vodka with the share of young consumers in each market. To identify how the price response varies with income, we interact the count of competing products that share each focal product's price category (cheap vs. expensive) with the share of households in the market with income above \$50,000. These instruments proposed by Waldfogel (2003) are valid if there are no demand spillovers from consumers in other similar markets.

5.2. *Estimating Mean Utility Coefficients*

In the second and third stages of estimation, we recover the product-time fixed effects $v_{jt}(\hat{\theta}_A)$ from the first-stage estimates $\hat{\theta}_A$. We express v_{jt} as a function of price and seasonal indicators, controlling for product fixed effects $\overline{v}_j$,

$$v_{jt} = \alpha p_{jt} + [h_t \quad q3_t]\gamma + \overline{v}_j + \Delta v_{jt}. \tag{31}$$

Equation (31) highlights the potential for price endogeneity, to the extent that a product's price responds to common time-varying preference variation, such as seasonal variation in consumption. Since the *PLCB*'s pricing rule applies a fixed markup to wholesale price irrespective of local or seasonal demand responses, its pricing cannot respond to such unobserved demand shocks. But the predictable link between wholesale and retail prices opens the possibility to spirit prices being endogenous not because of the pricing practices of the *PLCB* but because of the pricing behavior of upstream distillers whose chosen wholesale prices in equation (27) reflect, through market shares, the unobserved common tastes for product characteristics of spirits, Δv_{jt}.

In principle, the fact that distillers need to request both temporary and permanent changes to their wholesale price a number of months before the new price takes effect mitigates such endogeneity concerns. Prices thus only respond to predictable variation in a product's demand over time. At the same time, none of our product characteristics vary across time, limiting our ability to flexibly represent such time varying preference heterogeneity at the level of the product. We therefore use instrumental variables techniques to estimate the parameters in equation (31).

We use two sets of price instruments. First, we rely on the contemporaneous average price of a given product from liquor control states outside of the Northeast and Mid-Atlantic regions.[26] Our identifying assumption is that cost shocks are national, since products are often produced in a single facility, but demand shocks are largely regional, linked to differences in climate and demographics, which correlate with consumption as we establish in Section 3.4. We thus focus on product prices in control states that are distant from Pennsylvania and have sufficiently different demographic profiles and climates to support the assumption that their demand shocks are uncorrelated with Pennsylvania's. For example, states such as Idaho, North Carolina, Oregon, and Wyoming have at least a 50% higher share of Hispanics than Pennsylvania's low 6%, while only North Carolina has a significant prevalence of African Americans (at 22% double Pennsylvania's 11% share), but also a very different climate. As a result, the appeal of product categories like whiskey whose consumption peaks during cold months varies significantly between Pennsylvania and the states used to construct the instrument. Given that most price changes in the data are temporary, rather than permanent, a concern with this instrument would be possible coordination of sales across the control states. We find little evidence of such coordination, however.[27]

We add to this instrument changes in the price of major inputs, sugar and corn, interacted with spirit type indicators to account for exogenous cost changes across spirit types: while a major input for cordials and rums is sugar, corn is an input to gins, vodkas, and whiskeys. Including input price interactions for barley, glass, oats, rice, rye, sorghum, and wheat did not change the estimates.

In the third and final estimation step, we use the estimated product fixed effects $\overline{\upsilon}_j$ from equation (31) and project them onto observable product characteristics x_j to isolate systematic variation in demand by spirit type, proof, and bottle size. Variation in prices over time identifies the price coefficient, exploiting the fact that distillers do not change the wholesale prices p^w for all products at the same time, which introduces variation in relative prices. We identify seasonality and mean preferences for time-invariant product characteristics such as proof and spirit type from systematic variation in market shares of spirits by time period and by characteristic.

[26]An "alcohol control state" is a state that has monopoly over the wholesale or retail sales of some or all categories of alcoholic beverages. As of 2016, this list includes: Alabama, Idaho, Iowa, Maine, Maryland (Montgomery, Somerset, and Wicomico counties), Michigan, Mississippi, Montana, New Hampshire, North Carolina, Ohio, Oregon, Pennsylvania, Utah, Vermont, Virginia, West Virginia, and Wyoming. Washington state privatized its alcohol retail distribution in 2012.

[27]We define a price reduction in our price instruments as "temporary" when price falls for one month. We then test whether these indicators are correlated with temporary price reductions in Pennsylvania using Kendall's τ, which we find to be low (on average 0.03 across state-pairs). This indicates that these temporary price changes do not appear to be correlated.

5.3. *Estimation Results*

Table IV presents the demand estimates of our preferred specification of the *RCNL* model using the three-step procedure outlined above. The parameters are precisely estimated. Estimated demand for spirits increases during the summer and the holiday season. On average, consumer valuations of brandy, cordials, and whiskey exceed gin (our reference category), while rums and vodkas are, on average, less valued. Consumers prefer 1.75 L to smaller bottles reflecting that when a brand is available in several bottle sizes, the 1.75 L is more popular, all else equal. Consumers also prefer imported, flavored, and high proof products.

These are only average valuations. We also allow for rich variation in valuations across demographic groups by interacting the young, minority, and college-educated indicators with spirit type, bottle size, and proof. The estimates of Π reveal significant heterogeneity in tastes for spirits. Demand becomes steeper as consumers become wealthier, consistent with the increased consumption of expensive spirits by "high income" consumers presented in Figure 1. Consumers of different demographic groups often do not favor the same spirit types: Although the average consumer values vodkas less than gins, individuals with some college education strongly favor them, as they do whiskeys. These consumers also dislike brandy, while minorities rank brandy over gin and then whiskey. Young consumers have a marked preference for rum and 750 ml bottles.

Despite the large number of demographic interactions included in Π, we still estimate statistically significant random coefficients (Σ) for proof and 750 ml bottles. We also obtain a significant estimate of the spirit type nesting parameter ρ, indicating that products within a spirit type are closer substitutes for each other than those of a different spirit category. To corroborate this point, we calculate the ratio of the average cross-price elasticity between each focal product and all products that share its spirit type and the average cross-price elasticity between the product and all remaining products. This ratio amounts to 17.1 for brandies—suggesting that the average cross-price elasticity within brandies is 17.1 times the cross-price elasticity of a brandy and other products; 6.5 for cordials; 18.1 for gins; 9.8 for rums; 7.4 for vodkas; and 5.8 for whiskeys. In Table D.I in Appendix D of the Supplemental Material (Miravete, Seim, and Thurk (2018)), we also document the estimated best substitute for a variety of popular products. The demand estimates thus generate sensible substitution patterns across products.

Turning to the own-price responsiveness of demand, the estimates in Table IV imply an average elasticity for a given product of −3.86. Rums, brandies, and cordials have less elastic demand than vodka and whiskey products. Similarly, demand for 375 ml bottles is less elastic than for 1.75 L bottles, with the medium-sized 750 ml bottles in-between. We also find that cheap spirit products have less elastic demand than more expensive products. See Table V. These elasticities are similar to those reported in Conlon and Rao (2015) and Aguirregabiria, Ershov, and Suzuki (2016) who both estimated alcohol demand using disaggregate data. They are also robust to our choice of price instruments. In Appendix C.1 of the Supplemental Material (Miravete, Seim, and Thurk (2018)), we document that the price coefficient is stable when using alternative instruments in the second stage of estimation and that accounting for price endogeneity due to upstream responses, as expected, entails a larger estimated price coefficient than a simple OLS specification.

The model estimates imply an average price elasticity of off-premise spirit demand of −2.8. That is, a 1% increase in the retail price of all spirits leads to a 2.8% decrease in the aggregate quantity of off-premise spirits demanded. Earlier literature, however, documents less elastic demand estimates for alcoholic beverages in general, and spirits in

TABLE IV

RCNL DEMAND ESTIMATES[a]

	Mean Utility	Random Coeff.	Demographic Interactions (Π)			
	(β)	(Σ)	Income	Young	Minority	College
PRICE	−0.3062 (0.0036)		0.1151 (0.0036)			
HOLIDAY	0.3153 (0.0057)					
SUMMER	0.0557 (0.0049)					
375 ml	−2.9554 (0.5608)					
750 ml	−7.5816 (0.4037)	0.5939 (0.3061)		22.7684 (3.2953)	0.4025 (0.0844)	4.9886 (0.2976)
BRANDY	0.3882 (0.6902)			0.8616 (0.2288)	1.3978 (0.0231)	−0.8738 (0.0518)
CORDIALS	0.2977 (0.7163)					
RUM	−4.7646 (0.8355)			11.5406 (2.9485)	−0.1628 (0.0146)	0.6795 (0.0426)
VODKA	−1.9611 (0.4835)			4.9747 (0.6656)	−0.3713 (0.0233)	4.2314 (0.2701)
WHISKEY	0.3875 (0.5123)			1.2203 (0.2059)	−0.9270 (0.0231)	0.9549 (0.0554)
FLAVORED	3.7007 (0.4848)			−4.9731 (0.7219)	−0.5111 (0.0374)	−3.2395 (0.1943)
IMPORTED	1.3598 (0.3519)	0.1912 (0.5134)				
PROOF	15.1897 (1.6844)	1.2575 (0.2505)		−26.0064 (4.2377)	1.6695 (0.0913)	−5.5765 (0.4402)
QUALITY	3.9347 (2.1101)					
CONSTANT	−15.3884 (1.8244)					
NEST (ρ)		0.1225 (0.0139)				

[a] Robust standard errors are reported in parentheses. Estimates for random coefficients Σ, demographic interactions Π, and nesting parameter ρ are based on *GMM* estimation using 3,377,659 observations in 12,957 markets given by store-period pairs and 200 simulated agents in each market. Mean utility estimates for price, holiday, and summer are based on the projection of estimated product-time fixed effects from the *GMM* estimation onto corresponding characteristics plus product fixed effects after controlling for price endogeneity. Remaining mean utility estimates based on the projection of the estimated product fixed effects onto the remaining observable product characteristics.

particular. Based on a review of the literature, Leung and Phelps (1993) concluded that the price elasticity of demand for distilled spirits is −1.5. Two effects likely drive the discrepancy between our results and this work. First, rather than estimating the elasticity of total alcohol consumption, we exclude the presumably less price sensitive consumption in bars and restaurants. Second, the earlier studies use state or national consumption data whereas we have detailed local data on consumption choices. In Appendix C.2 of the Supplemental Material (Miravete, Seim, and Thurk (2018)), we show that aggregation in our data drives the price coefficient, and consequently the estimated elasticity for spirits toward zero. We also show in Table C.IV that values of the price coefficient that generate elasticities consistent with Leung and Phelps (1993) imply demand curves which are inconsistent with upstream profit-maximization.

Lastly, the demand estimates translate into sensible variation in estimated elasticities across markets that differ in consumer demographics. In line with the consumption differences from Section 3.4, markets with high concentrations of wealthy and educated consumers exhibit product demands that are relatively less elastic, while markets with concentrations of young people tend to have more elastic demands for individual spirits. Demand for spirits as a category is less elastic in markets with large minority populations and markets with greater levels of income and education.

The analysis so far has focused on estimating a measure of the elasticity of demand for alcoholic beverages (ε) that reflects the systematic differences between products. Heterogeneity in consumer preferences for these differentiated products translates into variation in product-level elasticities that filters through to the aggregate elasticity for spirits, one of the two key determinants of optimal tax rates. The second is the strategic price response of distillers (η), the estimation of which we turn to next.

5.4. *The Upstream Marketplace*

To characterize the upstream response to changes in the downstream tax rate, we require estimates of distiller marginal cost. We use our demand estimates and the data on wholesale price and quantity sold to back out product-level marginal costs for each pricing period that imply that distillers maximize profit following the model of distiller conduct from Section 4.2 and the associated optimality conditions in equation (27). Since we observe both wholesale and retail prices, we do not impose any supply-side restrictions on prices during estimation, which alleviates the identification concerns of Villas-Boas and Hellerstein (2006). Here, we assume that firms set wholesale prices that internalize the effect of each product's price on the remaining products in their portfolio; the ownership matrix O^w simply reflects the firms' portfolios during the sample period.

The resulting cost estimates $\hat{c}$ are reasonable (see middle columns of Table V). The marginal cost of expensive products is on average 2.5 times that of inexpensive products; for a subset of brandies and whiskeys with detailed product information, we find that the marginal cost of products that are aged for five or more years is 1.3 and 1.4 times the cost of nonaged products for brandies and whiskeys, respectively. Lastly, imported products are 1.8 times more costly than nonimported products on average.

We use these cost estimates to assess market power in the distillery market. The last columns in Table V indicate that the average Lerner index is 34.66%: out of every dollar received from the *PLCB* distillers take home 35 cents in profits. Diageo, Bacardi, and Beam products generate average margins of 32.3%, 32.2%, and 38.4%, respectively.[28] Small bottles have higher margins (54% on average) than large bottles (26% on average)

[28]Financial disclosures for seven public distillers during the sample period report gross profit margins that are comparable to the estimated Lerner indeces, albeit while deriving from the companies' aggregate sales across states. For instance, the average gross profit margin among public distillers amounts to 37.80%; Diageo's gross profit margin (30.04%) is lower than Beam's (46.06%).

TABLE V

ESTIMATED PRODUCT ELASTICITIES, UPSTREAM MARGINAL COSTS, AND UPSTREAM MARKET POWER[a]

		$\varepsilon_{jt}(p^r)$		$\hat{c}_{jt}$		Lerner_{jt}	
	Price	Avg	SD	Avg	SD	Avg	SD
By Firm:							
Bacardi	17.66	−4.21	1.58	7.80	5.37	32.21	13.83
Beam	12.89	−3.39	1.03	4.89	2.91	38.35	12.19
Diageo	17.08	−4.15	1.44	7.38	4.79	32.34	11.42
Other Firms (34)	15.23	−3.81	1.41	6.40	4.43	35.07	15.95
By Price and Size:							
EXPENSIVE	20.37	−4.73	1.37	9.43	4.62	25.94	8.15
CHEAP	11.04	−3.04	0.84	3.79	1.98	42.92	14.75
375 ml	9.16	−2.54	0.83	2.71	2.07	53.54	20.32
750 ml	14.43	−3.76	1.23	5.99	3.57	34.35	11.28
1.75 L	21.16	−4.68	1.39	9.34	5.13	26.30	6.80
ALL PRODUCTS	15.63	−3.86	1.41	6.53	4.51	34.66	14.70

[a]"Price" is measured in dollars; "Average" refers to the simple average of the corresponding category while "SD" is the standard deviation. Appendix D of the Supplemental Material (Miravete, Seim, and Thurk (2018)) presents the full elasticity distributions by spirit type and bottle size. Estimated marginal costs ($\hat{c}_{jt}$) are based on product-level marginal costs for each pricing period using the firm first-order conditions in equation under the assumption that firms set prices for products in their observed product portfolios. "Lerner_{jt}" the Lerner index for product j in period t defined as $100 \times \frac{p^w_{jt} - \hat{c}_{jt}}{p^w_{jt}}$.

due to their less elastic demand. There is less heterogeneity across spirit types as the average whiskey generates a margin of 32.5% compared to 37.6% and 36.7% for the average brandy and rum, respectively.

These estimates of upstream market power suggest that distillers can do much to counteract changes in *PLCB* policy. The upstream response is more complex than in the single-product monopoly model of Section 2: as above, a modification in *PLCB* policy generates incentives for distillers to adjust wholesale price and, indirectly through the pricing rule, retail price, to offset the effect of a tax rate change on demand. With multiple products, there is always the possibility of substitution across products rather than just to the outside option, however, and firms coordinate pricing across their full portfolio of products.

To highlight the combined effect of these inputs into the firms' wholesale price response, Table VI first summarizes the responsiveness of retail prices to an increase in the tax rate, holding fixed wholesale prices at the levels observed in the data. The *PLCB* pricing formula implies that on average, a 1% increase in the tax rate beyond the level observed in the data translates into a 0.57% increase in retail prices. The observed variation in tax elasticities reflects differences in price levels across producers and product categories, as well as variation in the relative contribution to final retail prices of the per-unit handling fee, which we denote as $LTMF$ in Section 3.2 and hold constant throughout. It amounts to a larger share of the retail markup for cheap products than for expensive products, contributing to the lower tax elasticity of retail prices for cheap products.

In the right-most columns of Table VI, we describe the upstream price response to such a tax rate increase by calculating the percent change in wholesale price as the tax rate rises by 1%—or η in the simple monopoly model above. As in the simple model from Section 2, we observe the tax rate and upstream wholesale prices are strategic substitutes, that is, $\eta < 0$. Across products, distillers reduce wholesale price by an average of 0.20% when the tax rate rises by 1% above the observed level, but there is significant hetero-

TABLE VI

RETAIL AND WHOLESALE PRICE RESPONSES TO CHANGES IN TAX POLICY[a]

	Retail Price Response (ψ) [$\eta = 0$]		Wholesale Price Response [η]	
	Avg	SD	Avg	SD
By Firm:				
Bacardi	0.58	0.03	−0.18	0.09
Beam	0.57	0.02	−0.19	0.08
Diageo	0.59	0.02	−0.17	0.08
Other Firms (34)	0.56	0.05	−0.22	0.17
By Price and Size:				
EXPENSIVE	0.60	0.02	−0.13	0.05
CHEAP	0.55	0.04	−0.27	0.16
375 ml	0.52	0.05	−0.43	0.21
750 ml	0.58	0.03	−0.18	0.07
1.75 L	0.58	0.02	−0.14	0.05
ALL PRODUCTS	0.57	0.04	−0.20	0.14

[a]"Avg" and "SD" are the sales-weighted (bottles) average and standard deviation of the corresponding category. "Retail Price Response" is the percent change in retail price from a 1% increase in the ad valorem tax rate, holding fixed upstream prices. "Wholesale Price Response" is the percent change in wholesale price given a 1% increase in the ad valorem tax rate, assuming upstream conduct based on product ownership. Retail and wholesale price response calculated at the observed tax rate in the data of $\tau = 53.4\%$.

geneity in the response. The reduction in wholesale price is greatest among cheap and 375 ml products. This estimated strategic substitutability of wholesale prices and tax rates suggests that the naïve retail price responsiveness estimates in column 1 overstate the equilibrium increase in retail prices after a tax rate hike. In the following section, we formally assess the magnitude of the distiller response and its effect on mitigating the retail price pass-through of alternative tax rates.

6. LAFFER CURVES: POLICYMAKER FORESIGHT AND MARKET CONDUCT

We now use our estimates of spirit demand and upstream marginal costs to measure effect of distiller pricing responses on state tax revenue. The goal of this section is to characterize the Laffer curve that traces tax revenue as the *PLCB* varies its ad valorem tax, τ, while holding other aspects of the pricing regulation fixed. We emphasize τ as the central policy instrument since control and noncontrol states use ad valorem taxes in the regulation of alcohol. As taxation of goods is commonly done with ad valorem taxes (e.g., sales taxes) our results can therefore be extended to a broad set of industries. Our focus on the *PLCB*'s ability to generate tax revenue for the state's general fund is motivated by ongoing efforts in the state legislature at reforming the Pennsylvania wine and spirits retail markets. For example, the Pennsylvania Legislature's Act 39, which took effect in August 2016, granted the *PLCB* limited pricing flexibility in allowing it to price its best-selling items, defined as the top-selling 150 SKUs, "in a manner that maximizes the return on the sale of those items."[29]

[29]Omnibus Amendment to Pennsylvania's Liquor Code, Act of June 8, 2016, P.L. 273, No. 39. Full text available at http://www.legis.state.pa.us/cfdocs/legis/li/uconsCheck.cfm?yr=2016&sessInd=0&act=39. At the time of this writing, the *PLCB* has taken advantage of this pricing flexibility only a limited number of times.

A regulator and 37 firms producing 312 products may interact in multiple ways. The degree to which the *PLCB* understands and internalizes the optimal responses of distillers and consumers to changes in tax policy is unclear. Rather than taking a stand on the ability of the *PLCB* to anticipate agents' responses, we choose to use the model to evaluate the influence of distillers on the final outcome of the policy under different beliefs or information assumptions on the part of the regulator.

6.1. *The Laffer Curve and Naïve Policymakers: Mechanical Effect*

The first alternative we consider is an application of the famous critique of Lucas (1976) in which the *PLCB* mistakenly believes that upstream firms lack the interest or ability to react to changes in policy: $\eta(\tau) = 0$. This is the so-called "mechanical effect" of taxation where we assume firms do not respond to changes in the tax policy, although we allow for reoptimization by consumers to changes in retail price. We call this the "Naïve" equilibrium.

We solve for this equilibrium by varying the *PLCB*'s tax rate from 10% to 90% using one percentage point increments. Throughout we incorporate into the markup the Johnstown alcohol tax of 18% so current policy corresponds to a tax rate of 53.4%. For each tax level and fixed vector of wholesale prices, we solve for the new vector of retail prices in equation (16) and rely on the estimated demand system to predict consumer spirit purchases and *PLCB* tax revenue.

Figure 2 plots the resulting absolute value of the demand elasticity of spirits as a category (ε in Section 2) and the *PLCB*'s tax revenue as a function of the ad valorem tax τ for fixed wholesale prices. Panel (a) shows that off-premise spirit demand becomes increasingly elastic as the tax rate (and retail price) increases. Spirit demand is elastic for even small values of the tax, indicating that the wholesale prices in the data, assumed constant here, generate retail prices on the elastic portion of the demand curve for a wide variety of tax rates.

In panel (b), we plot the associated tax revenue the *PLCB* would generate under the alternative rates. For small values of the tax rate, the Laffer curve is steep, reflecting lower demand elasticities for spirits as a category. As we move past the peak into the "prohibitive" range, the slope becomes flatter as the demand elasticity stabilizes. Under the

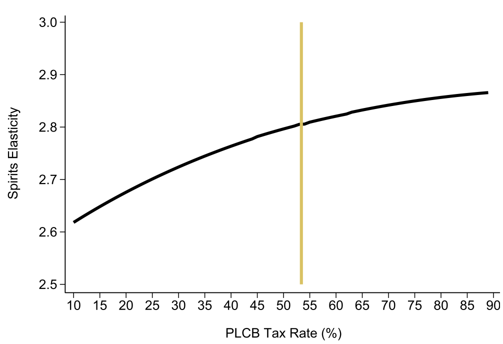

(a) Spirit Demand Elasticity

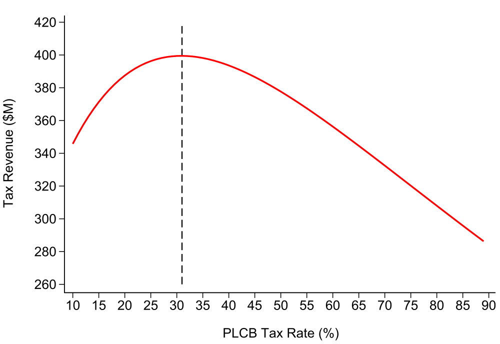

(b) Laffer Curve – Naïve

FIGURE 2.—Tax rate, spirit demand elasticity, and the Laffer curve. Notes: Absolute value of the spirit demand elasticity and the PLCB Laffer curve under constant wholesale prices. "PLCB Tax Rate" is in percent and includes the 18% Johnstown Flood tax. The solid vertical line corresponds to the current policy; the dashed line to the tax revenue maximizing policy.

current 53.4% tax rate, the agency forgoes significant tax revenue by overpricing spirits. The *PLCB* could maximize tax revenue with a tax rate of 30.7%, leading to an estimated increase in tax revenue of 7.8%, or $28.7 million, and a decrease in retail prices of 13%. In response, consumption would increase by 47.5%, or 28.6 million bottles, which in turn would lead to a 51% increase in upstream profits, or $79.7 million.

6.2. *The Laffer Curve and Naïve Policymakers: Behavioral Response*

Next, we allow distillers to exploit the market power we demonstrate in Section 5.4 by responding to the *PLCB*'s naïve policy. The difference between the intended goal analyzed in the Naïve equilibrium and the realized goal in what we call the "Response" equilibrium identifies the degree to which upstream firms can unravel the *PLCB*'s policy. This is commonly known as the "behavioral response" (Saez (2001)). The taxation literature typically considers only behavioral responses related to the general equilibrium effect of taxation on the entry and exit of firms in a perfectly competitive industry. As we consider firms with market power, the behavioral response we consider here encompasses the wholesale pricing response to taxes as a function of the nature of competition among upstream firms. The difference between Naïve and Response equilibria thus establishes the potential error in assuming perfect competition among firms.

In panel (a) of Figure 3, we present the upstream firm response elasticity η under different choices of the tax rate. We solve for the response elasticity numerically by comparing Bertrand–Nash equilibrium wholesale prices for each incremental 1% change in the *PLCB* ad valorem tax rate (τ). Three characteristics of the upstream response stand out. First, the upstream response to changes in the *PLCB* tax rate varies significantly across products, as illustrated through the middle 50% interquartile range for η (shaded region). This reflects differences in product elasticities driven by differences in product characteristics, firm portfolios, and consumer demand. Second, the upstream firm response becomes more muted at higher tax rates reflecting the curvature of log-concave demand increasing in price (Fabinger and Weyl (2016)). Third, at no point in the range of tax rates we consider do firms choose not to respond, consistent with the existence of market power and in stark contrast with the predictions of models of commodity taxation in competitive environments so common in public finance.

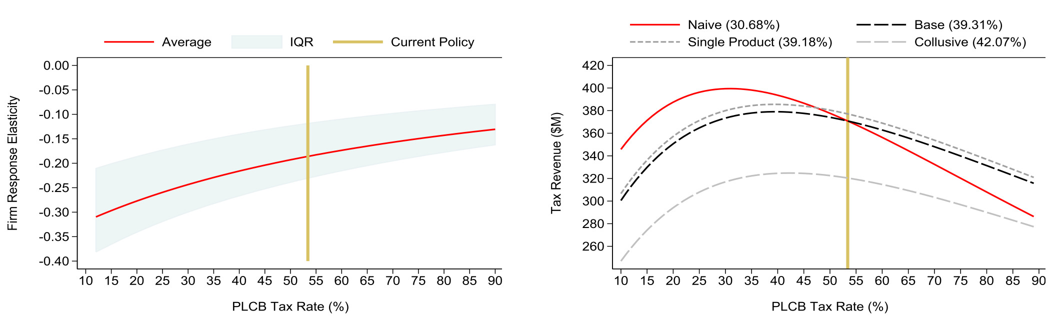

(a) Upstream Firm Response Elasticity (η)

(b) Naïve and Response Laffer Curves

FIGURE 3.—Distiller response and the Laffer curve. Notes: Panel (a) presents the average upstream response elasticity and their 50% interquartile range (IQR) when firms optimally set wholesale prices. Panel (b) presents "Response" Laffer curves accounting for the profit-maximizing response of distillers under alternative assumptions on upstream conduct. Revenue-maximizing tax rate for each Laffer curve indicated in parentheses. Vertical line corresponds to the current 53.4% policy.

TABLE VII

MAXIMIZING TAX REVENUE AND DISTILLER RESPONSES[a]

	Base		Single Product		Collusive	
	Naïve	Response	Naïve	Response	Naïve	Response
Markup (%)	30.68	30.68	30.90	30.90	29.15	29.15
Percent Change:						
– Bottles	47.52	34.59	46.34	34.22	56.47	34.98
– Distiller Price (p^w)	0.00	3.79	0.00	3.61	0.00	5.83
– Retail Price (p^r)	−13.36	−10.45	−13.22	−10.45	−14.35	−9.90
– Distiller Profit	51.33	56.22	50.22	55.35	60.62	62.51
– Tax Revenue (T)	7.75	1.01	7.49	1.15	9.75	−1.36
Elasticities:						
– Spirits Demand (ε)	−2.63	−2.73	−2.60	−2.70	−2.78	−2.93
– Wholesale Price Response (η)	0.00	−0.24	0.00	−0.23	0.00	−0.33
– Retail Price Response (ψ)	0.67	0.47	0.66	0.47	0.68	0.40
Consumer Pass-Through	100.00	69.04	100.00	70.06	100.00	57.18

[a] "Markup" is in percent and includes the 18% Johnstown Flood tax. "Percent Change" is the percent change in the corresponding statistic from its value under the tax rate in the data ($\tau = 53.4\%$) and wholesale prices implied by the assumed competitive conduct. Wholesale prices based on estimates of upstream marginal costs presented in Table V. "Retail Price Response" is the percent change in retail price from a 1% increase in the ad valorem tax rate (τ) including the upstream firm response. Retail and wholesale price response calculated at the markup in row 1. "Consumer Pass-Through" is the share of an incremental tax change borne by consumers. Where applicable, statistics are the sales-weighted (bottles) average.

In panel (b), we compare the Naïve and Response Laffer curves to assess the aggregate implications of ignoring firm market power. Response Laffer curves are flatter and shift down and to the right, crossing the Naïve Laffer curve at the observed equilibrium in the data. The change in location and shape reflects distillers maximizing profits by moving their wholesale prices in the opposite direction of any change in the *PLCB* tax rate. We again find that current tax policy overprices spirits though the revenue-maximizing tax rate increases from 30.7% when we hold wholesale prices fixed to 39.3% when we allow for the strategic firm response with the current product portfolios, that is, the "Base."

To summarize the consequences of naïve policy making, we show in Table VII the aggregate effect on prices, consumption, and tax revenue of the *PLCB* reducing its tax rate to 30.7%, which it naïvely believes to be the revenue maximizing tax level. In response, upstream firms increase wholesale prices by 3.8% on average. Appendix E in the Supplemental Material (Miravete, Seim, and Thurk (2018)) contains detailed results for individual firms. The upstream response, which increases distiller profit gains by an additional 4.89 percentage points relative to the naïve scenario, limits the *PLCB*'s gain in tax revenue to 1% of profit at current rates, or only 12.97% of the envisioned tax revenue of a naïve regulator.

We summarize the impact on consumers using two measures. The average response elasticity of the retail price with respect to the tax rate captures the extent to which the tax decrease translates into retail price decreases. It is only 0.47, instead of 0.67 under the naïve policy. Upstream firms raising wholesale price with the tax rate cut limits the retail price decrease to 10.5%, instead of 13.4%. Second, we report consumer pass-through rates which we construct as the average share of the incremental tax under the new percent markup that is reflected in the retail price. The *PLCB*'s mechanical pricing rule implies a 100% retail price pass-through when wholesale prices cannot adjust; in the naïve scenario, the price thus changes by the full amount of the incremental tax. Under the response equilibrium, we calculate the change in retail price when not only the markup but also wholesale prices change, as a share of the same incremental tax. With wholesale price adjustments, we find that only 69.0% of the tax decrease in moving to a tax rate of 30.7% feeds through to the retail price. Quantity consumed thus increases by a more limited amount than anticipated by the *PLCB*, and the optimal wholesale pricing response to changes in τ nearly fully undermines the achievement of the *PLCB*'s tax revenue goal.

To highlight the effect of the degree of imperfect competition on the location and shape of the Laffer curve, we also consider two departures from Bertrand–Nash pricing. In the first, a product manager in each firm chooses the price for her product without internalizing the effects of that price choice on the demand for other products in the firm's portfolio. We call this form of conduct "Single Product" to represent the most competitive behavior possible within a differentiated products Bertrand pricing oligopoly. In the second, we allow all firms to jointly set prices, and we call this form of conduct "Collusive." In each alternative conduct scenario, we solve for the upstream response using the same marginal cost estimates as in the base conduct case.[30]

We plot these Response Laffer curves alongside the Naïve and Base Response Laffer curves in Figure 3. The figure highlights several implications of departing from multi-product Bertrand–Nash pricing. Not surprisingly, the intensity of upstream competition affects the level of tax revenue the *PLCB* realizes. Across tax rates, revenues are lowest when upstream conduct is collusive. For the tax rate in the data of 53.4%, for example, the ability of upstream firms to jointly set wholesale prices reduces tax revenue by 13.6% relative to base conduct (Table E.I). In contrast, revenues are uniformly higher under single-product pricing than under multi-product pricing, reflecting the more competitive upstream market. The difference is not pronounced, however: for this same tax rate of 53.4%, the increase in tax revenue over the base conduct case is only 1.7%. This similarity between tax revenue and distiller profit under the two forms of conduct reflects the large distillers' broad product portfolios covering multiple spirit types that we describe in Section 3.3. This limits the role of business stealing across products in the firm's portfolio that pricing under the base form of conduct internalizes (see Section 4.2).

Figure 3 also demonstrates that the Response Laffer curves shift down and to the right and flatten under all three conduct scenarios.[31] We find similar patterns if instead of focusing on Pennsylvania in aggregate, we compare the Naïve and Response Laffer curves for subsets of store markets at the top and bottom of the distributions of income, educational attainment, age, and size of the minority population. The patterns here are thus not the result of pooling markets with different demographic makeups and, consequently, different distributions of demand elasticities (see Appendix E of the Supplemental Material (Miravete, Seim, and Thurk (2018))). We thus conclude that the response of distillers with imperfectly competitive market conduct generally erodes the *PLCB*'s naïve policy.

[30]We compute the Laffer curves under the alternative forms of firm conduct as follows. First, we use the estimates of marginal cost (Table V) and the firm first-order conditions (Equation (27)) under single-product and collusive ownership matrices O^w to solve for wholesale prices at the observed tax rate of 53.4%. This enables construction of the Naïve Laffer curve condition on firm conduct. We construct the corresponding Response and Stackelberg equilibria by varying the *PLCB* tax rate τ and resolving for the Bertrand–Nash equilibrium.

[31]For simplicity, we omit the naïve Laffer curves for collusive and single-product upstream pricing from Figure 3.

Table VII illustrates that the ability of upstream firms to erode the *PLCB*'s naïve policy is higher the less competitive the upstream market. We contrast the naïve and response equilibria for all three forms of upstream conduct. In line with additional market power reducing the *PLCB*'s ability to raise revenue broadly, we find that the revenue-maximizing tax rate is 29.2% when wholesale prices are fixed at collusive upstream prices, but 30.9% when wholesale prices are fixed at single-product upstream prices. The upstream response to a move to such naïvely optimal tax rates is similar to the above when distillers price as single-product firms: tax revenue increases only by 1.2%, instead of 7.5% under the naïve equilibrium, relative to tax revenue at the current tax rates and single-product wholesale prices. At the other extreme, a collusive response by upstream firms to moving to the naïvely optimal tax rate results in a 1.4% *reduction* in expected revenue relative to revenue under the current tax rate and collusive wholesale prices. Hence, not accounting for market power among regulated firms leads to a suboptimal policy recommendation with more significant unanticipated effects on tax collection for less competitive taxed industries.

6.3. *The Laffer Curve and Policymakers With Perfect Foresight*

We now compare the Naïve and Response equilibria to one in which the regulator has perfect foresight and correctly anticipates the distiller response. Graphically, we observe this "Stackelberg" equilibrium as the tax rate that maximizes tax revenue of the response Laffer curve conditional on our assumption of upstream conduct. This strategic pricing game between upstream distillers and the *PLCB* as the downstream retailer resembles the classic double-marginalization problem considered in the vertical contracting literature (e.g., Villas-Boas (2007) and Mortimer (2008)). Since the *PLCB* has traditionally committed to a tax rate through legislative oversight, we assume—in contrast to the contracting literature—that the *PLCB* moves first in choosing the tax rate τ (equivalent to a downstream markup), before upstream firms respond by setting the wholesale price p^w. For each conduct assumption, we solve for the tax rate τ that maximizes *PLCB* tax revenue, the peak of the Laffer curves of Figure 3, given the firms' wholesale price responses. We summarize the resulting outcome in Table VIII.

The comparison of the Naïve, Response, and Stackelberg regimes reveals how the regulator has to alter its policy in order to accommodate the optimal distiller response and still achieve the objective of maximizing tax revenues. The Stackelberg equilibrium under the base assumption of upstream firms choosing prices for all products in their portfolio, entails the lowering of the tax rate from 53.4% to only 39.3%, instead of 30.7% under the naïve equilibrium. Current policy thus overprices spirits and is in the "prohibitive range" of the Laffer curve regardless of whether our regulator is naïve or has perfect foresight. The average retail price would be 6.48% (or $1.00) lower in the Stackelberg equilibrium, relative to 10.45% ($1.61) in the base Response equilibrium. While the *PLCB* is able to generate higher revenue under this lower tax rate, distillers are the clear winners as the 30.8% increase in distiller profits far outpaces the *PLCB*'s 2% growth in tax revenue. As a share of the sum of upstream profit and tax revenue, distillers would account for 34.9% under the Stackelberg equilibrium, compared to only 29.5% under current prices. The Stackelberg equilibria under alternative assumptions on upstream conduct depart in similar ways from the respective Naïve equilibria. Appendix E of the Supplemental Material (Miravete, Seim, and Thurk (2018)) contains detailed results, and Appendix C.4 of the Supplemental Material (Miravete, Seim, and Thurk (2018)) investigates the robustness of our overpricing finding to alternative demand specifications.

TABLE VIII

TAX REVENUE AND REGULATOR FORESIGHT[a]

	Base	Single Product	Collusive
Markup (%)	39.31	39.18	42.07
Percent Change:			
– Bottles	19.62	19.84	14.35
– Distiller Price (p^w)	2.21	2.15	2.45
– Retail Price (p^r)	−6.48	−6.60	−4.65
– Distiller Profit	30.80	31.06	24.08
–Tax Revenue (T)	2.23	2.27	1.40
Elasticities:			
– Spirits Demand (ε)	−2.76	−2.73	−2.96
– Wholesale Price Response (η)	−0.21	−0.21	−0.28
– Retail Price Response (ψ)	0.45	0.45	0.38
Consumer Pass-Through	70.28	71.25	60.07

[a]"Markup" is in percent and includes the 18% Johnstown Flood tax. "Percent Change" is the percent change in the corresponding statistic from its value under the tax rate in the data ($\tau = 53.4\%$) and wholesale prices implied by the assumed competitive conduct. Wholesale prices based on estimates of upstream marginal costs presented in Table V. "Retail Price Response" is the percent change in retail price from a 1% increase in the ad valorem tax rate (τ) including the upstream firm response. Retail and wholesale price response calculated at the markup in row 1. "Consumer Pass-Through" is the share of an incremental tax change borne by consumers.

Figure 4 investigates heterogeneity in the distillers' profit impact of a move to revenue-maximizing tax rates under the three policy regimes, contrasting the top three distillers with upstream distillers as a group. It suggests that as the *PLCB* lowers tax rates, Diageo and Bacardi benefit from their large product portfolios in extracting a greater share of industry profit across all scenarios. See Table E.I in Appendix E of the Supplemental Material (Miravete, Seim, and Thurk (2018)) for detail. We also illustrate the effect of moving to revenue-maximizing tax rates for consumers in terms of total expenditure on spirits products. Consistent with the current markup being too high across all policy regimes, we

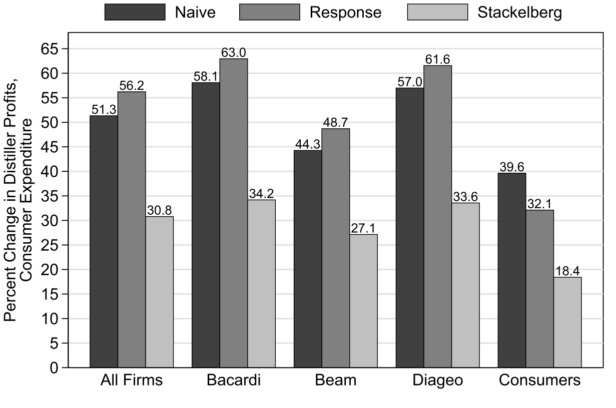

FIGURE 4.—Tax incidence across equilibria. Notes: Figure presents percent change in profits (distillers) and liquor expenditure (consumers) from estimated equilibrium under different policy regimes. Distiller pricing reflects observed product ownership.

find that consumer expenditure uniformly increases by between 18.4 and 39.6% across regimes. The largest increases arise under the Naïve and Response equilibria where the average retail price declines induce significant increases in spirit consumption that drive up spending.

7. CONCLUDING REMARKS

We study the relationship between commodity taxation and tax revenues in a noncompetitive industry. We show the existence of a Laffer curve with an optimal, tax revenue maximizing, rate that depends not only on the elasticity of demand but also on the tax authority's ability to anticipate the pricing response of taxed firms to commodity taxation. Accounting for both effects, we show that reducing the tax rate below typical levels of sales taxes would raise total tax revenue only for products with highly elastic demands. This indicates that current sales taxes are below optimal levels for the vast majority of products. For demand elasticities in the range of those generally estimated for differentiated consumer products, the optimal commodity tax rate is significantly higher, but indeed similar to commonly observed excise tax rates, such as those on alcoholic beverages.

Against this backdrop, we empirically analyze alcohol taxation in the Pennsylvania spirits market. We show the current 53.4% ad valorem tax is excessively high irrespective of any anticipation by the tax authority of a strategic price response by distillers to its choice of tax rate. Our estimated demand system for differentiated spirits, combined with a game-theoretic model of distiller pricing that accounts for the response of upstream firms to changes in tax policy, allows us to characterize Laffer curves in a wide variety of scenarios. We find that upstream market power and price responses mitigate the effect of any tax policy change on tax revenue, reducing the ability of the tax authority to drive revenue generation via tax policy changes.

We point to oligopolistic firms' strategic price responses as the most important behavioral effect of taxation rather than the entry and exit channel used in the existing public economics literature. Our results suggest that assuming perfect competition among firms in their pricing has the potential to generate poor policy recommendations: Regardless of the policymaker's objective—be it tied to tax revenues, overall consumption levels, or equilibrium prices—we show that ignoring firms' price responses neglects their ability to undo the realization of the policymaker's objective. Recent work by Fowlie, Reguant, and Ryan (2016) similarly points to the role of market power by firms and the associated allocative inefficiencies as a limiting factor for the realization of social benefits of corrective taxation. While theoretical work going back to Musgrave (1959), Bishop (1968), and Buchanan (1969) has long recognized the need to account for the nature of competition in optimal commodity and corrective taxation, our work highlights the empirical relevance of their conclusions, in particular since perfectly competitive markets are regrettably rare in practice.

REFERENCES

AGUIRREGABIRIA, V., D. ERSHOV, AND J. SUZUKI (2016): "Estimating the Effect of Deregulation in the Ontario Wine Retail Market," Report, University of Toronto. [1653,1674]

ATKINSON, A. B., AND J. S. STIGLITZ (2015): *Lectures on Public Economics*. Princeton, NJ: Princeton University Press. [1652]

AUERBACH, A. J. (1985): "The Theory of Excess Burden and Optimal Taxation," in *Handbook of Public Economics, Vol. 1*, ed. by A. J. Auerbach and M. Feldstein. New York, NY: North-Holland. [1652]

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890. [1666]

BERRY, S. T., AND P. A. HAILE (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797. [1672]

BISHOP, R. L. (1968): "The Effects of Specific and Ad Valorem Taxes," *The Quarterly Journal of Economics*, 82, 198–217. [1685]

BRESNAHAN, T. F., S. STERN, AND M. TRAJTENBERG (1997): "Market Segmentation and the Sources of Rents From Innovation: Personal Computers in the Late 1980," *RAND Journal of Economics*, 28, S17–S44. [1672]

BUCHANAN, J. M. (1969): "External Diseconomies, Corrective Taxes, and Market Structure," *American Economic Review*, 59, 174–177. [1685]

CHETTY, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488. [1652]

CHETTY, R., A. LOONEY, AND K. KROFT (2009): "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99, 1145–1177. [1652]

CONLON, C. T., AND N. RAO (2015): "The Price of Liquor Is Too Damn High: The Effects of Post and Hold Pricing," Report, Columbia University. [1653,1674]

DUBÉ, J.-P., J. T. FOX, AND C.-L. SU (2012): "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 80, 2231–2267. [1671, 1672]

EVANS, W. N., J. S. RINGEL, AND D. STECH (1999): "Tobacco Taxes and Public Policy to Discourage Smoking," *Tax Policy and the Economy*, 13, 1–55. [1652]

FABINGER, M., AND E. WEYL (2016): "The Average-Marginal Relationship and Tractable Equilibrium Forms," Report, University of Tokyo. [1656,1659,1680]

FOWLIE, M., M. REGUANT, AND S. P. RYAN (2016): "Market-Based Emissions Regulation and Industry Dynamics," *Journal of Political Economy*, 124, 249–302. [1685]

FULLERTON, D. (1982): "On the Possibility of an Inverse Relationship Between Tax Rates and Government Revenues," *Journal of Public Economics*, 19, 3–22. [1652]

GRIFFITH, R., M. O'CONNELL, AND K. SMITH (2017): "Design of Optimal Corrective Taxes in the Alcohol Market," Report, Institute for Fiscal Studies. [1653]

GRIGOLON, L., AND F. VERBOVEN (2014): "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation," *Review of Economics and Statistics*, 66, 916–935. [1667,1671]

HAUGHWOUT, S. P., R. A. LAVALLEE, AND I.-J. P. CASTLE (2015): "Apparent per capita Alcohol Consumption: National, State, and Regional Trends, 1977–2013," Surveillance Report 102, National Institute on Alcohol Abuse and Alcoholism. [1667]

HENDEL, I. (1999): "Estimating Mulitple-Discrete Choice Models: An Application to Computerization Returns," *Review of Economic Studies*, 66, 423–446. [1667]

HENDEL, I., AND A. NEVO (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74, 1637–1673. [1667]

ILLANES, G., AND S. MOSHARY (2017): "What Is the Marginal Effect of Entry? Evidence From a Natural Experiment in Liquor Licensure," Report, Northwestern University. [1653]

IRELAND, P. N. (1994): "Supply-Side Economics and Endogenous Growth," *Journal of Monetary Economics*, 33, 559–571. [1652]

JENSEN, W. B. (2004): "The Origin of Alcohol Proof," *Journal of Chemical Education*, 81, 1258. [1660]

LEUNG, S. F., AND C. E. PHELPS (1993): "My Kingdom for a Drink ... ? A Review of the Price Sensitivity of Demand for Alcoholic Beverages," in *Economics and the Prevention of Alcohol-Related Problems*, ed. by M. Hilton and G. Bloss. Bethesda, MD: National Institute on Alcohol Abuse and Alcoholism, 1–32. [1675, 1676]

LUCAS, R. E. (1976): "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labor Markets*, ed. by K. Brunner and A. Meltzer. New York, NY: North-Holland. [1679]

MILLER, N. H., AND M. C. WEINBERG (2017): "Understanding the Price Effects of the MillerCoors Joint Venture," *Econometrica*, 85 (6), 1763–1791. [1653]

MIRAVETE, E. J., K. SEIM, AND J. THURK (2017): "One Markup to Rule Them All: Taxation by Liquor Pricing Regulation," Report, University of Texas. [1653]

——— (2018): "Supplement to 'Market Power and the Laffer Curve'," *Econometrica Supplemental Material*, 86, https://doi.org/10.3982/ECTA12307. [1655,1659,1661-1664,1667,1669,1674,1676,1677,1681-1684]

MORTIMER, J. H. (2008): "Vertical Contracts in the Video Rental Industry," *The Review of Economic Studies*, 75 (1), 165–199. [1683]

MUSGRAVE, R. (1959): *The Theory of Public Finance: A Study in Public Economy*. New York: McGraw-Hill. [1685]

NEVO, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342. [1666,1670]

NOVALES, A., AND J. RUIZ (2002): "Dynamic Laffer Curves," *Journal of Economic Dynamics & Control*, 27, 181–206. [1652]

PECORINO, P. (2011): "Tax Rates and Tax Revenues in a Model of Growth Through Human Capital Accumulation," *Journal of Monetary Economics*, 36, 527–539. [1652]

SAEZ, E. (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 68, 205–229. [1680]

SCHMITT-GROHÉ, S., AND M. URIBE (1997): "Balanced-Budget Rules, Distortionary Taxes, and Aggregate Instability," *Journal of Political Economy*, 105, 976–1000. [1652]

SEIM, K., AND J. WALDFOGEL (2013): "Public Monopoly and Economic Efficiency: Evidence From the Pennsylvania Liquor Control Board's Entry Decisions," *American Economic Review*, 103, 831–862. [1653, 1659]

SOMAINI, P., AND F. WOLAK (2015): "An Algorithm to Estimate the Two-Way Fixed Effect Models," *Journal of Econometric Methods*, 5, 105–130. [1671]

STOLPER, S. (2016): "Who Bears the Burden of Energy Taxes? The Role of Local Pass-Through," Working paper, Harvard University. [1652]

STUART, C. E. (1981): "Swedish Tax Rates, Labor Supply, and Tax Revenues," *Journal of Political Economy*, 89, 1020–1038. [1652]

TIROLE, J. (1989): *Industrial Organization*. Cambridge, MA: The MIT Press. [1656]

TRABANDT, M., AND H. UHLIG (2011): "The Laffer Curve Revisited," *Journal of Monetary Economics*, 58, 305–327. [1652,1657]

VILLAS-BOAS, S. B. (2007): "Vertical Relationships Between Manufacturers and Retailers: Inference With Limited Data," *Review of Economic Studies*, 74, 625–652. [1670,1683]

VILLAS-BOAS, S. B., AND R. HELLERSTEIN (2006): "Identification of Supply Models of Retailers and Manufacturer Oligopoly Pricing," *Economics Letters*, 90, 132–150. [1676]

WALDFOGEL, J. (2003): "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *RAND Journal of Economics*, 34, 557–568. [1672]

WEYL, E., AND M. FABINGER (2013): "Pass-Through as an Economic Tool: Principles of Incidence Under Imperfect Competition," *Journal of Political Economy*, 121, 528–583. [1652,1659]

Co-editor Liran Einav handled this manuscript.

Manuscript received 17 February, 2014; final version accepted 2 May, 2018; available online 25 July, 2018.