

CONSUMER DEMAND UNDER PRICE UNCERTAINTY: EMPIRICAL EVIDENCE FROM THE MARKET FOR CIGARETTES

Mark Coppejans, Donna Gilleskie, Holger Sieg, and Koleman Strumpf*

Abstract—We develop a demand model for goods that are subject to habit formation. We show that consumption plans of forward-looking individuals depend on preferences, current period prices, and individual beliefs about the evolution of future prices. Moreover, an increase in price uncertainty reduces consumption along the optimal path. With smoking as our application, we test the predictions of our model using a unique data set of prices for cigarettes and the restricted-use version of the National Education Longitudinal Study. Our estimation results suggest that teenagers who live in metropolitan areas with a large amount of cigarette price volatility have, on average, significantly lower levels of cigarette consumption.

I. Introduction

THE goal of this paper is to analyze consumer demand in markets with large price uncertainty. Our analysis differs from most previous consumer demand studies which assume either that individuals face little uncertainty about future prices or that uncertainty about future prices has a negligible impact on demand. While this may be a plausible assumption in many applications, there are clearly cases in which expectations and uncertainty about future prices cannot be ignored. In these circumstances consumer demand will be affected by price uncertainty.¹ In this paper, we develop a demand model for goods that are subject to habit formation. We show that consumption plans of forward-looking individuals depend not only on preferences and current period prices, but also on individual beliefs about the evolution of future prices. Moreover, a mean preserving spread in the price distribution and, hence, an increase in price uncertainty reduces consumption along the optimal path.

One purpose of this paper is to quantify the effects of price uncertainty on consumer demand in volatile markets. Our application focuses on the demand for cigarettes. Our empirical analysis draws on a number of different data sets. We use a unique price data set collected by the Bureau of Labor Statistics (BLS) which provides price measures at the metropolitan-area level. Our analysis thus avoids aggregation bias inherent in data at the state or federal level. Our empirical findings suggest that there is much price variation among the set of metropolitan areas analyzed in this study. Furthermore, estimates based on aggregate time series often

underestimate the amount of price volatility experienced at the local level.

To estimate the effects of price uncertainty on consumption decisions, we need to characterize price expectations and construct measures of price volatility. First, we focus on historical volatility.² These measures are relevant if individuals form adaptive price expectations, that is, if individuals infer future prices by looking at price realizations in the preceding periods. Second, we compute more sophisticated measures of price expectations. These measures capture the idea that individuals must forecast future price realizations. Ideally forecasts should be based on the true data-generating process (Muth, 1961). We explore regime-switching models proposed by Hamilton (1989, 1990) to model the time series properties of prices. Our empirical findings suggest that for the majority of metropolitan areas studied in this paper, there are two distinctly different regimes of price changes. There are time periods which are fairly stable and exhibit only small changes in prices. These periods are followed by short periods which are much more volatile and exhibit large swings in prices. In these periods, predicted confidence intervals for future prices are quite large.

We then investigate whether the demand for cigarettes is affected by price volatility. We focus on the behavior of young individuals who may be most susceptible to large swings in prices because their disposable income is relatively low compared to adults. This part of the analysis is based on a restricted-use version of the National Education Longitudinal Study (NELS) which is collected by the National Center for Educational Statistics (NCES). We merge the NELS with BLS data on prices and price volatilities using geographic identifiers in the NELS. We then estimate demand models to quantify the impact of price volatility on the demand for cigarettes among teenagers. We find that individuals who live in metropolitan areas with a large amount of price volatility have, on average, significantly lower levels of cigarette consumption. Moreover, these individuals are less likely to start consuming cigarettes. Models based on forecasted price volatility fit the data slightly better than models based on historical volatility measures. However, formal non-nested hypothesis tests often fail to distinguish between the alternatives. The results also provide some evidence that young individuals are forward looking. If teenagers were myopic, price volatility measures would have little explanatory power for observed choices. Our findings suggest the opposite: teenagers re-

² This approach of measuring price volatility is in the spirit of the adaptive expectation hypothesis which is typically attributed to Nerlove (1958).

Received for publication September 9, 2004. Revision accepted for publication April 10, 2006.

* Barclays Global Investors, University of North Carolina at Chapel Hill and NBER, Carnegie Mellon University and NBER, and Kansas School of Business, respectively.

¹ The seminal paper on market equilibrium under price uncertainty is Muth (1961). Some influential recent studies are Wolak and Kolstad (1991) who study input demand under price uncertainty, Appelbaum and Ullah (1997) who analyze production decisions under price uncertainty, and Hall and Rust (2002) and Osborne (2004) who consider inventory decisions under price uncertainty. Finally, Erdem, Imai, and Keane (2003) estimate a model of brand choice under price uncertainty.

spond to increased price uncertainty by reducing their consumption.

There are three strands of the empirical literature on rational addiction that are closely related to this study. Most prior empirical studies of the rational addiction model follow Becker and Murphy (1988) and analyze first-order conditions that prices and quantities need to satisfy, given individuals' quadratic utility functions. Chaloupka (1991) and Becker, Grossman, and Murphy (1991, 1994) apply this methodology and find that tobacco consumption typically responds to lagged, current, and future prices as predicted by rational addiction theory.³ A second line of research develops alternative tests of forward-looking behavior focusing on behavioral responses to changes in tax policy (Gruber and Koszegi, 2001) or health shocks (Arcidiacono, Sieg, and Sloan, 2007). Finally, there are a number of empirical studies that primarily focus on smoking initiation of teenagers.⁴

The rest of the paper is organized as follows. In section II, we present a demand model with habit formation and characterize the relationship between consumption of addictive goods and price expectations. Section III focuses on measuring price uncertainty. This part of the analysis is based on monthly price data collected by the BLS. Section IV investigates the impact of price uncertainty on the consumption of cigarettes using a sample drawn from the NELS. Section V summarizes the main findings and offers some conclusions that can be drawn from our analysis.

II. Price Uncertainty, Expectations, and Consumer Demand

The starting point of our analysis is a consumer demand model that accounts for habit formation and uncertainty about future prices.⁵ Consider an individual who can consume two types of goods: a good which is subject to habit formation denoted by a_t and a composite private good denoted by c_t . The stock characterizing habit formation, S_t , evolves according to the law of motion $S_{t+1} = \delta S_t + a_t$, where δ is the rate of depreciation of the stock. Individuals rank alternatives according to a utility function $u(c_t, a_t, S_t)$ that satisfies standard regularity assumptions imposed in the habit formation literature.⁶ Individuals are forward looking with a planning horizon of T periods. Individuals maximize expected intertemporal utility:

$$E \left(\sum_{t=1}^T \beta^{t-1} u(c_t, a_t, S_t) \right), \quad (1)$$

³ Chaloupka and Warner (2000) provide an overview of the existing empirical literature on the rational addiction model.

⁴ Some recent examples include DeCicca, Kenkel, and Mathios (2002) and Gilleskie and Strumpf (2005).

⁵ Becker and Murphy (1988) develop the basic rational addiction model without uncertainty. Orphanides and Zervos (1995) consider uncertainty about addiction, but not price uncertainty.

⁶ These assumptions are smoothness, concavity, complementarity of a and S , and negativity.

where β is the discount factor.⁷ Thus if $\beta = 0$, individuals are myopic. If $\beta > 0$ individuals are forward looking. Individuals face a sequence of budget constraints given by

$$c_t + p_t a_t = y_t, \quad (2)$$

where p_t is the gross-of-tax price of a at time t and y_t denotes income at time t . We have conveniently normalized the price of the composite private good to be equal to 1.⁸ Prices for the addictive good evolve according to a stochastic law of motion. Individuals do not have perfect foresight. Instead, they have subjective beliefs characterizing the distribution of future prices. Price expectations are given by the transition density, $f(p_{t+1}|p_t)$.

Since we abstract from saving decisions, we can substitute the budget constraint into the utility function and define

$$w(y_t, p_t, a_t, S_t) = u(y_t - p_t a_t, a_t, S_t). \quad (3)$$

Under the assumptions made above, we can express the dynamic optimization problem faced by a forward-looking individual using the following recursive representation:

$$V_t(y_t, S_t, p_t) = \max_{a_t \in [0, y_t/p_t]} w(y_t, p_t, a_t, S_t) + \beta \int V_{t+1}(y_{t+1}, \delta S_t + a_t, p_{t+1}) f(p_{t+1}|p_t) dp_{t+1}, \quad (4)$$

where $V_t(\cdot)$ denotes the value function at time t . The state variables are y_t , S_t , and p_t . The decision variable is a_t .

We are primarily interested in characterizing the relationship between the consumption of the addictive good a_t and the beliefs that individuals hold about future prices. To get more precise results, we follow Orphanides and Zervos (1995) and assume that the preferences of individuals can be characterized by the following function:

$$u(c_t, a_t, S_t) = \ln(c_t) + \ln(a_t) + S_t^\psi (-\phi + \gamma a_t), \quad (5)$$

where ψ , γ , and ϕ are parameters of the model. Given this additional assumption, we can prove the following result:

Proposition 1. Under the assumptions made above, a mean preserving spread in the price distribution for each period will reduce smoking along the optimal path.

⁷ Alternatively, one could assume that individuals engage in hyperbolic discounting as suggested by Harris and Laibson (2001) and Gruber and Koszegi (2001) or make systematic mistakes as in Bernheim and Rangel (2004). Our main argument rests on the notion that individuals are forward looking.

⁸ For simplicity, we assume that there is no savings. This is a reasonable assumption for young individuals. Our analysis also largely abstracts from stockpiling, which is an interesting aspect of consumer behavior of frequently consumed goods as discussed, for example, by Hendel and Nevo (2002). However, teenagers are less likely to have such sophisticated behavioral patterns. First, it is harder for teenagers to store large amounts of cigarettes, especially if their parents do not want them to smoke. Second, they are less likely to make bulk purchases largely because of cash constraints or legal issues associated with purchase of cigarettes.

An increase in the variance of future prices implies an increase in the variance of future consumption. As a consequence individuals will substitute from the risky consumption good that is subject to price uncertainty to the consumption good without price uncertainty.⁹

Our model suggests three hypotheses that can be tested empirically. First, myopic and forward-looking individuals will reduce consumption of the addictive good in response to a current period price increase. Second, myopic individuals are not concerned about future prices. In particular, their behavior does not depend on future price expectations. Third, forward-looking individuals are concerned about future prices. An increase in the variance of future prices will decrease the current period consumption of the addictive good. In the remaining sections of this study, we provide an empirical investigation of these hypotheses.

III. Measuring Price Uncertainty

A. Data

The purpose of this paper is to analyze whether consumer decisions are affected by price uncertainty in markets that exhibit large fluctuations in prices. Our application focuses on the market for cigarettes. Studying the demand for cigarettes is interesting because prices have been fluctuating significantly over the past ten to fifteen years. To investigate the relationship between price volatility and consumption decisions, we first need to measure price volatility. Our price data come from the Price Indices for Tobacco and Smoking Products collected by the Bureau of Labor Statistics (BLS). This data set consists of monthly price series for a number of metropolitan areas in the United States covering the time period from 1986 to 2002.¹⁰ An advantage of the BLS data is that prices are measured on a disaggregate level. Our analysis thus avoids aggregation bias. Prices are also sampled on a monthly basis, which allows us to focus on price variation within shorter periods. Empirical analysis based on quarterly or yearly data is likely to underestimate the significant amount of price variation in the underlying price processes. The BLS data are also reliable, cover a large time period, and are available upon request from the BLS.¹¹

The BLS sample contains price indices for a large number of metropolitan areas in the United States. In this part of the analysis, we restrict attention to 27 metropolitan areas. The BLS reports an index that we converted into price per carton using the ACCRA (2003) data which includes quar-

terly prices, inclusive of all excise taxes, for a carton of cigarettes.¹²

Table 1 reports descriptive statistics for the 27 metro areas in our sample. It reports means, standard deviations, minimums, and maximums for price levels over the sixteen-year period from 1986:12 to 2002:11. The reported minimum typically occurred during the beginning of the time period; the maximum, toward the end. We, therefore, find that prices increased on average by approximately 100% during the observation period. To illustrate the basic properties of our price data, we provide plots for four metro areas in our sample. Figure 1 suggests that prices of tobacco products increased substantially throughout the time period in four metropolitan areas. However, there are also time periods in the sample in which prices decreased. A comparison of all metropolitan areas in the sample shows that there is significant heterogeneity in prices among geographic entities in the United States.

A large majority of the price series exhibit a strong upward trend. This suggests that the price process may not be stationary in levels. To investigate these issues more formally, we consider the following baseline model to test for stationarity:

$$p_{it} = a_i + b_i p_{it-1} + c_i t + e_{it}, \quad (6)$$

where p_{it} denotes price levels of city i at time t and e_{it} is a white-noise error term. Based on this model, we construct four different stationarity tests. Our first test statistic, denoted by T-I in table 1, gives the p -values for the null $c_i = 0$ (assuming $|b_i| < 1$). Hence small p -values are evidence in support of a nonstationary model because of the time trend. The second test, denoted by T-II, is the test statistic—not the p -value—for the null $b_i = 1$ and $c_i = 0$ (Hamilton, 1994). Approximate 5% and 1% critical values for this test statistic are 6.48 and 8.72. DF is the Dickey-Fuller test, which is the t -statistic for the null that $b_i = 0$ in the model

$$\Delta p_{it} = a_i + b_i p_{it-1} + c_i t + d_i \Delta p_{it-1} + e_{it}, \quad (7)$$

where $\Delta p_{it} = p_{it} - p_{it-1}$. The 95% confidence interval for this test statistic is $(-3.69, -0.62)$. Finally, we report the modified Dickey-Fuller (MDF) test statistic for the null that $b_i = 0$ and $c_i = 0$ in the above model.

Table 1 reports the results for the four stationarity tests. Our findings suggest that prices may not be stationary in levels. The p -values for the first test statistic are low, and the second test statistic is often above the critical levels at commonly used significance levels. The two versions of the DF test show similar results.

It is also interesting to ask the question whether a subset of the metropolitan areas are driven by common stochastic trends which may explain some of the short-term and

⁹ Proposition 1 follows from the fact that the value function is concave in prices. A proof is provided in the appendix.

¹⁰ Past research has primarily relied on either the Tobacco Institute's weighted-average price by state or data collected by ACCRA (2003).

¹¹ Alternatively one could rely on commercially available data from sources such as A. C. Nielsen or IRI. These data sets allow researchers to focus on even higher frequencies such as weekly observations. See, for example, the work by Erdem and Keane (1994) or Erdem et al. (2003).

¹² For each metro area, we normalize the BLS index so that February 1993 is unity and then multiply by the ACCRA price from the first quarter of 1993.

TABLE 1.—DESCRIPTIVE STATISTICS

City	Price Levels (per carton)							
	Mean	SD	Min	Max	T-I	DF	MDF	T-II
Anchorage	19.23	5.93	10.55	30.84	0.19	-1.21	0.75	1.10
Atlanta	17.77	6.97	10.07	34.15	0.06	-1.21	2.51	1.78
Baltimore	21.43	9.65	9.48	45.10	0.17	-0.25	2.66	1.92
Boston	23.76	9.35	11.93	51.54	0.05	-0.44	1.94	2.10
Chicago	22.74	8.37	11.16	42.63	0.07	-0.48	1.91	1.90
Cincinnati	18.07	7.38	8.99	35.81	0.03	-1.42	1.76	2.44
Cleveland	20.18	7.25	10.73	40.16	0.04	-0.43	2.03	2.19
Dallas	20.41	6.17	11.66	35.86	0.00	-1.76	1.72	4.76
Denver	16.83	4.75	9.97	27.72	0.01	-1.67	1.58	4.12
Detroit	22.21	8.37	10.18	42.38	0.00	-2.26	3.15	4.99
Honolulu	21.72	9.53	8.98	41.27	0.05	-1.47	1.90	2.02
Houston	19.32	4.19	11.46	28.00	0.02	-1.81	1.66	3.80
Kansas City	23.30	9.60	11.79	44.08	0.00	-1.79	2.48	6.01
Los Angeles metro	23.29	9.27	10.68	43.31	0.10	-1.37	1.05	1.46
Miami	21.10	6.69	13.33	38.00	0.11	-0.52	1.48	1.33
Milwaukee	19.60	6.28	10.53	34.71	0.02	-1.76	1.76	2.88
Minneapolis	26.04	12.02	10.30	53.05	0.00	-2.06	2.29	5.25
New York City metro	24.18	10.06	12.28	53.84	0.36	0.61	4.09	2.46
Philadelphia	18.82	6.70	9.83	37.87	0.02	-1.10	1.56	2.63
Pittsburgh	18.57	6.08	10.58	36.99	0.08	-0.47	1.55	1.66
Portland	22.03	9.10	10.89	43.42	0.03	-0.96	2.01	2.51
San Diego	21.15	7.45	9.73	36.63	0.01	-1.59	1.32	3.52
San Francisco	23.55	9.84	10.56	43.85	0.10	-1.56	1.47	1.39
Seattle	23.89	9.09	11.18	45.93	0.10	-1.15	1.93	2.03
St. Louis	17.28	5.54	10.39	31.33	0.07	-1.01	1.23	1.69
Tampa	20.07	6.46	11.11	34.35	0.03	-1.62	1.54	2.53
Washington, DC	19.64	8.24	9.69	40.76	0.03	-0.96	1.65	2.43

long-term interactions of the time series. To investigate these issues we perform a panel data unit root test suggested by Levin and Lin (1993). This test is based on the following model specification:

$$\Delta p_{it} = a_i + a_t + b_i p_{it-1} + \sum_{j=1}^{k_i} d_{ij} \Delta p_{it-j} + e_{it}. \tag{8}$$

The null hypothesis is the joint hypothesis that $b_i = 0$ for all i . Since the asymptotic distribution of the test statistic is hard to approximate, we follow Cecchetti, Nelson, and Sonora (2002) and use bootstrap techniques to compute the p -value associated with the test statistic. We implement the test focusing on a subsample of our data set which consists of the four cities used in figure 1. We find that the p -value for the LL test is 0.30, which is not strong evidence against our modeling approach.

We therefore difference the data and run the same tests on the differenced data. Our results suggest that first-differencing the data yields stationary time series. We conclude that it is reasonable to model prices in first differences as a stationary time series. We therefore estimate all formal pricing models reported in the next subsection in first differences. To illustrate the main properties of the data in first differences, we also plot the time series for four metro areas in figure 1. The plots suggest that there are time periods that are characterized by large volatilities in prices. The last few years in our sample are very good examples of these high-volatility time periods. At the same time, there appear to be periods with fairly low variation in prices. For

example, price changes are much smaller in the middle of our observation period.

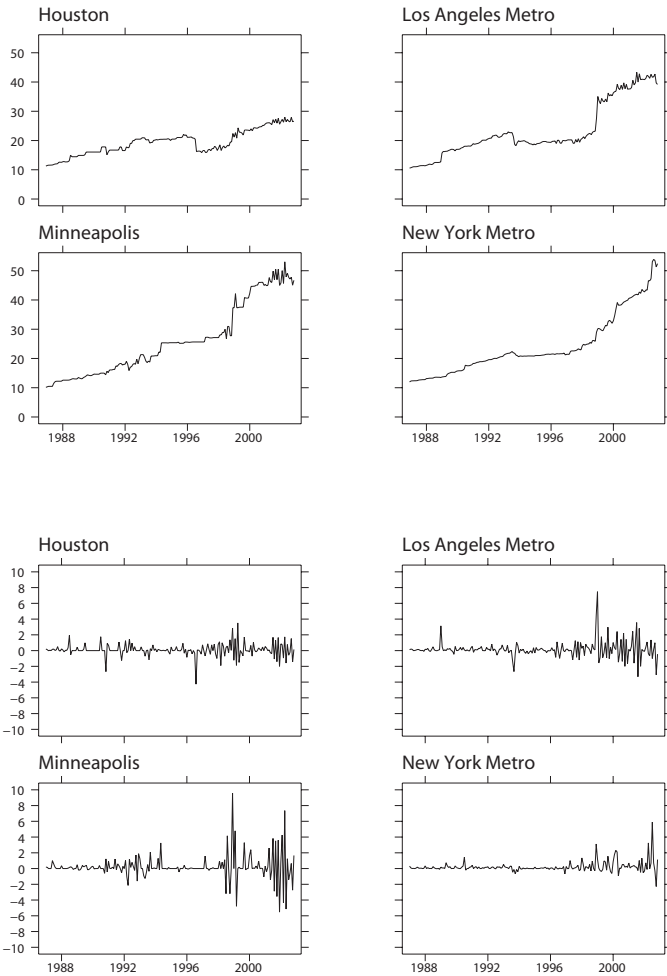
Some of the differences in estimates for the cities in our sample may be attributed to the sampling period. For example, it is possible that we would get different estimates and hence different predictions of the price volatilities if we began or ended the price series at different points of times. Lagged price changes lead to initial-condition problems in estimation and thus affect forecasting. However, our data set is relatively large, including monthly observations for sixteen years. We expect that the type of initial-condition problems discussed above are important in shorter data sets with yearly or quarterly observations.¹³

To investigate these issues more rigorously, we analyze whether large metropolitan areas lead smaller metropolitan areas in pricing behavior. In particular, we use the price series for New York, Chicago, and Los Angeles and investigate whether lagged values of these series have any explanatory power in the price process of smaller metro areas in their vicinities. We estimate the following model:

$$\Delta p_{it} = a_i + b_i \Delta p_{it-1} + c_i \Delta p_{NY,t-1} + d_i \Delta p_{Chi,t-1} + e_i \Delta p_{LA,t-1} + e_{it}. \tag{9}$$

¹³ We observe prices in the BLS data set for a much longer time period than we observe smoking choices in the NELs. Thus our analysis of smoking behavior reported in section IV of this paper does not require us to forecast price volatilities at the beginning or the end of the BLS sample period.

FIGURE 1.—PRICES (\$/CARTON) IN LEVELS (TOP PANEL) AND DIFFERENCES (BOTTOM PANEL)



We calculate an F-test for the null hypothesis that $c_i = d_i = e_i = 0$ for the 24 cities that are not New York, Chicago, or Los Angeles. We find that the null hypothesis is rejected for 15 out the 24 cities at 5%. But we can only reject the null at 1% three times. We then pick the major city that mostly influences a minor city (via significance of F -test) and estimate the following model:

$$\Delta p_{it} = a_i + b_i \Delta p_{it-1} + c_i \Delta p_{NY, Chi, or LA, t-1} + e_{it}. \quad (10)$$

We test the null hypothesis that $c_i = 0$. Our evidence suggests that New York seems to have some influence on other Northeastern cities such as Philadelphia, Boston, Pittsburgh, or Detroit. The results for all other cities are inconclusive. Sometimes we obtain counterintuitive negative point estimates. We thus conclude that there is only limited evidence suggesting the existence of spillover between larger and smaller cities.

Some of the price variation observed in our sample is a direct result of changes in state and federal tax policies that were implemented during the past decade. In 1983, the federal excise tax was doubled from \$0.08 per pack to \$0.16 per pack. This rate held for a decade, when the federal rate was increased another \$0.08 to \$0.24 per pack. In 1997, legislation passed increasing the federal tax on cigarettes to \$0.34 per pack in 2000 and \$0.39 per pack in 2002. Some of the cross-sectional variation of prices is due to differences in state tax policies. Table 2 summarizes the main features of policies during the last decade for the states in our sample. It highlights the cross-sectional and time series variation in tax rates. While all states had low rates in 1987, most doubled or tripled the tax in a series of changes over the next fifteen years. And while many tobacco-producing states maintained low rates, there were sharp increases in

TABLE 2.—ANNUAL CIGARETTE TAXES (IN CENTS PER PACKAGE)

State	1987 Tax	2002 Tax	1987–2002 Mean	1987–2002 Std Dev	# of Changes
Alaska	16	100	46.82	35.78	2
California	10	87	46.35	28.74	3
Colorado	15	20	19.71	1.21	1
Connecticut	26	111	46.82	18.77	5
Delaware	14	24	21.06	4.70	2
District of Columbia	13	65	46.06	23.15	4
Florida	21	33.9	30.81	4.98	2
Georgia	12	12	12.00	0	0
Hawaii	28	120	63.82	30.12	4
Illinois	20	98	41.53	20.44	4
Indiana	10.5	55.5	17.26	9.99	2
Maryland	13	100	37.12	24.77	4
Massachusetts	26	151	53.94	32.93	3
Michigan	21	125	53.94	31.42	3
Missouri	13	17	15.12	2.06	1
New Jersey	25	150	51.82	32.67	4
New York	21	150	57.47	35.28	5
Ohio	14	55	22.88	9.08	3
Oregon	27	128	47.18	27.45	5
Pennsylvania	18	100	30.41	18.99	4
Texas	20.5	41	35.94	8.22	2
Virginia	2.5	2.5	2.50	0	0
Wisconsin	25	77	43.59	16.77	5

other initially low tax states like California. The data, therefore, indicate that individuals in our sample of metro areas face a wide range of market conditions.

B. Expected Price Volatility

The theoretical model studied in section II suggests that consumption decisions can depend on beliefs that individuals hold about future prices. We construct two measures to characterize expected price volatility. The first measure is based on historical volatility. For that we consider the price variation in the preceding periods. This measure is consistent if individuals have adaptive expectations. However, individuals may be more rational and recognize that they need to forecast future price realizations. Extrapolating historical realizations may not be the best way to do that. To forecast prices, it is desirable to use a formal time series model of the price process. If individuals have correct expectations, beliefs will be based on objective price transitions; these can be estimated by an econometrician. To formulate measures characterizing expectations about future price uncertainty, we estimate regime-switching models pioneered by Hamilton (1989, 1990). We consider a first-order autoregressive regime switching model that can be written as

$$\Delta p_t = \mu_{s_t} + \rho_{s_t} \Delta p_{t-1} + \epsilon_{s_t}, \tag{11}$$

where s_t is the (unobserved) state of the time series process at time t . In a regime-switching model the parameters of the autoregressive process, μ_{s_t} and ρ_{s_t} , and the distribution of the error terms depend on the state of the process. This feature of the model allows us to capture the fact that prices are stable in some periods and highly volatile in other periods. We assume that ϵ_{s_t} is i.i.d. $N(0, \sigma_s^2)$. For notational simplicity, let us write the density of Δp_t conditional on $s_t = j$ and Δp_{t-1} as

$$f(\Delta p_t | s_t = j, \Delta p_{t-1}; \theta), \tag{12}$$

where θ is the parameter vector to be estimated.

The evolution of the state of the process is modeled as the outcome of an unobserved J -state Markov chain. For simplicity let us consider a two-regime model ($J = 2$). The Markov transition matrix for a two-regime model is given by

$$Q = \begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix}, \tag{13}$$

where $q_{ij} = \Pr\{s_t = i | s_{t-1} = j\}$. Denote the history of price changes up to time $t - 1$ as $\Delta \vec{p}_{t-1}$. The probability that the process is in state $s_t = j$ conditional on $\Delta \vec{p}_{t-1}$ is written as $\Pr\{s_t = j | \Delta \vec{p}_{t-1}; \theta\}$. Given that we observe a realization Δp_t , we draw inference about the state of the process by iterating the following two equations:

$$\begin{aligned} & \Pr\{s_t = i | \Delta \vec{p}_t; \theta\} \\ &= \frac{f(\Delta p_t | s_t = i, \Delta p_{t-1}; \theta) \Pr\{s_t = i | \Delta \vec{p}_{t-1}; \theta\}}{\sum_{j=1}^2 f(\Delta p_t | s_t = j, \Delta p_{t-1}; \theta) \Pr\{s_t = j | \Delta \vec{p}_{t-1}; \theta\}} \end{aligned} \tag{14}$$

and

$$\begin{bmatrix} \Pr\{s_{t+1} = 1 | \Delta \vec{p}_t; \theta\} \\ \Pr\{s_{t+1} = 2 | \Delta \vec{p}_t; \theta\} \end{bmatrix} = \begin{bmatrix} q_{11} & q_{21} \\ q_{12} & q_{22} \end{bmatrix} \begin{bmatrix} \Pr\{s_t = 1 | \Delta \vec{p}_t; \theta\} \\ \Pr\{s_t = 2 | \Delta \vec{p}_t; \theta\} \end{bmatrix}. \tag{15}$$

Equations (14) and (15) completely characterize the stochastic evolution of the state of the process. We have a sample of price changes observed over a sequence of T periods. The likelihood function for the data is given by the following equation:

$$L = \prod_{t=1}^T \sum_{j=1}^J \Pr\{s_t = j | \Delta \vec{p}_{t-1}; \theta\} f(\Delta p_t | s_t = j, \Delta p_{t-1}; \theta). \tag{16}$$

The likelihood function does not have a closed-form analytical solution, but needs to be computed using an Expectation-Maximization (EM) algorithm. In the EM algorithm, we start with an initial guess for the probabilities of each state and then iterate forward using equations (14) and (15) to compute the conditional probabilities characterizing each state at time t (Hamilton, 1994).

We estimate the regime-switching models for each of the 27 metro areas. Table 3 reports the point estimates of the parameters of the different regime-switching models. Estimated standard errors are reported in parentheses. Table 3 suggests that there is substantial heterogeneity among the metropolitan areas in our data set, as the parameter estimates differ considerably among the 27 metro areas. The estimates for the means, μ_j , are typically positive in both regimes. They are often significantly different from 0. This result reflects the earlier observation that prices were mostly increasing during the observation period. We also find that the point estimates for ρ_s are often negative in both regimes. This result suggests that there is some mean reversion in the data. A period of positive price changes is likely to be followed by a period with negative price changes.

Table 3 shows that two-state regime-switching models fit the data better than simple AR(1) specifications, at least for a large number of metro areas. The point estimates suggest that regime 1 is characterized by large changes in prices accompanied with large volatility. These are the periods of price wars or changes in tax or regulatory policies. Regime 2, in contrast, is fairly stable and shows only modest amounts of volatility and price changes.

We have also conducted a formal test to distinguish between one-state and two-state regime-switching models

TABLE 3.—MARKOV-SWITCHING ESTIMATES

City	q_{11}	q_{22}	μ_1	μ_2	ρ_1	ρ_2	σ_1^2	σ_2^2
Anchorage	0.676 (0.052)	0.600 (0.057)	0.195 (0.095)	0.000 (0.006)	-0.285 (0.118)	0.000 (0.005)	0.914 (0.131)	0.000 (0.012)
Atlanta	0.657 (0.053)	0.549 (0.058)	0.246 (0.094)	0.000 (0.006)	-0.416 (0.112)	-0.000 (0.008)	0.882 (0.123)	0.000 (0.013)
Baltimore	0.782 (0.053)	0.491 (0.076)	0.313 (0.095)	0.000 (0.010)	-0.249 (0.099)	0.000 (0.013)	0.805 (0.107)	0.000 (0.021)
Boston	0.920 (0.051)	0.943 (0.028)	0.421 (0.232)	0.127 (0.036)	-0.401 (0.127)	-0.074 (0.085)	3.102 (0.559)	0.106 (0.018)
Chicago	0.965 (0.099)	0.975 (0.019)	0.409 (0.202)	0.090 (0.021)	-0.470 (0.216)	-0.094 (0.085)	1.249 (0.290)	0.036 (0.007)
Cincinnati	0.699 (0.055)	0.415 (0.068)	0.234 (0.097)	0.000 (0.010)	-0.520 (0.100)	-0.000 (0.012)	1.002 (0.136)	0.000 (0.024)
Cleveland	0.904 (0.057)	0.936 (0.028)	0.323 (0.155)	0.124 (0.021)	-0.456 (0.131)	-0.342 (0.066)	1.500 (0.276)	0.036 (0.007)
Dallas	0.765 (0.058)	0.511 (0.070)	0.208 (0.123)	0.002 (0.012)	-0.529 (0.095)	-0.019 (0.018)	1.683 (0.225)	0.000 (0.063)
Denver	0.758 (0.046)	0.508 (0.068)	0.181 (0.092)	-0.000 (0.008)	-0.511 (0.059)	-0.000 (0.009)	0.965 (0.130)	0.000 (0.024)
Detroit	0.868 (0.059)	0.916 (0.032)	0.328 (0.192)	0.112 (0.023)	-0.261 (0.131)	-0.155 (0.062)	2.185 (0.403)	0.044 (0.010)
Honolulu	0.609 (0.059)	0.633 (0.054)	0.371 (0.108)	0.000 (0.006)	-0.314 (0.128)	-0.000 (0.007)	0.997 (0.141)	0.000 (0.011)
Houston	0.791 (0.042)	0.528 (0.065)	0.144 (0.081)	-0.000 (0.007)	-0.467 (0.095)	-0.000 (0.009)	0.768 (0.093)	0.000 (0.011)
Kansas City	0.894 (0.084)	0.964 (0.018)	0.459 (0.394)	0.079 (0.016)	-0.708 (0.249)	-0.002 (0.041)	4.583 (1.163)	0.033 (0.007)
Los Angeles metro	0.885 (0.060)	0.933 (0.028)	0.315 (0.226)	0.079 (0.025)	-0.224 (0.130)	0.066 (0.068)	2.861 (0.541)	0.053 (0.012)
Miami	0.685 (0.067)	0.711 (0.056)	0.338 (0.117)	0.048 (0.017)	-0.214 (0.113)	-0.972 (0.018)	1.056 (0.166)	0.016 (0.020)
Milwaukee	0.962 (0.137)	0.981 (0.050)	0.185 (0.269)	0.109 (0.031)	-0.287 (0.365)	-0.153 (0.094)	1.403 (0.411)	0.070 (0.010)
Minneapolis	0.729 (0.047)	0.556 (0.065)	0.383 (0.152)	0.000 (0.011)	-0.683 (0.097)	-0.000 (0.006)	2.727 (0.357)	0.000 (0.036)
New York metro	0.848 (0.064)	0.923 (0.028)	0.430 (0.177)	0.080 (0.012)	0.043 (0.149)	0.089 (0.047)	1.312 (0.251)	0.014 (0.005)
Philadelphia	0.936 (0.054)	0.958 (0.025)	0.349 (0.178)	0.055 (0.019)	-0.367 (0.138)	-0.131 (0.072)	1.813 (0.333)	0.030 (0.006)
Pittsburgh	0.660 (0.059)	0.566 (0.060)	0.291 (0.120)	0.000 (0.007)	-0.403 (0.116)	-0.000 (0.009)	1.137 (0.169)	0.000 (0.023)
Portland	0.696 (0.052)	0.406 (0.067)	0.364 (0.110)	0.000 (0.013)	-0.624 (0.096)	-0.000 (0.011)	1.279 (0.166)	0.000 (0.026)
San Diego	0.818 (0.045)	0.710 (0.064)	0.305 (0.135)	0.041 (0.015)	-0.367 (0.091)	-0.982 (0.011)	1.903 (0.259)	0.006 (0.028)
San Francisco	0.909 (0.072)	0.954 (0.023)	0.287 (0.244)	0.115 (0.022)	0.046 (0.262)	-0.166 (0.076)	2.080 (0.457)	0.053 (0.008)
Seattle	0.710 (0.053)	0.534 (0.061)	0.319 (0.098)	-0.000 (0.010)	-0.083 (0.111)	-0.000 (0.010)	0.938 (0.126)	0.000 (0.017)
St. Louis	0.738 (0.058)	0.424 (0.073)	0.195 (0.080)	0.000 (0.010)	-0.542 (0.099)	-0.000 (0.015)	0.674 (0.089)	0.000 (0.023)
Tampa	0.706 (0.057)	0.413 (0.069)	0.214 (0.093)	0.000 (0.011)	-0.344 (0.104)	-0.000 (0.014)	1.031 (0.140)	0.000 (0.030)
Washington, DC	0.651 (0.078)	0.840 (0.036)	0.500 (0.267)	0.062 (0.013)	-0.523 (0.167)	-0.012 (0.019)	3.482 (0.691)	0.018 (0.018)

for a subset of the metro areas in the data set. Determining the number of states in a regime-switching model is, however, complicated because standard regularity assumptions imposed in likelihood ratio tests are not met (Hansen, 1992). One of the problems encountered here is that the null hypothesis involves a restriction on the boundary. For these types of tests there are no general asymptotic results available. We therefore rely on bootstrapping algorithms to construct p -values for these tests. Our findings indicate that for the majority of metropolitan areas in our samples we can

reject the null hypothesis that there is only one state in the regime-switching model.¹⁴

As a final robustness check, we have also considered GARCH(1,1) models (Bollerslev, 1986). We have estimated

¹⁴ We also estimated AR models with more than one lag and found that the AR(1) specification is sufficient to capture the main regularities in the data. We also performed a number of sensitivity tests to investigate whether adding an additional state to our model would change the main empirical results. All of these tests suggested that adding a third regime to the model does not improve the fit of the model.

a separate GARCH(1,1) model for each city. As with the switching model, we find that there tends to be mean reversion in the differenced prices. In general the GARCH(1,1) models do not perform as well as the regime-switching models. (Details are available in Coppejans et al., 2006.)

IV. Price Volatility and Demand

A. Data

Despite a great interest in the United States in understanding youth smoking behavior, few nationally representative data sets are available that chronicle the behavior of the same children over multiple periods of time. The National Education Longitudinal Study (NELS) of 1988 is one exception. NELS, a continuing study sponsored by the U.S. Department of Education's National Center for Education Statistics, began in 1988 with the specific purpose of collecting information on educational, vocational, and personal development of a nationally representative sample of eighth graders as they transition from middle school into high school, through high school, and into postsecondary institutions and the workforce. Approximately 24,500 eighth graders in more than 1,000 public and private schools in all fifty states participated in the first wave of the study. In addition to the student questionnaires, supplementary questionnaires were administered to the students' parents, teachers, and school principals and provide a wealth of information on the early social and academic environment of the students. Through special agreement with the U.S. Department of Education, we obtained access to restricted-use NELS data that include geographic information.

The first follow-up, administered in the spring of 1990, includes responses from approximately 17,500 of the students from the 1988 base year interview, while the second follow-up, administered in the spring of 1992, includes approximately 16,500 students from the original cohort. One of the many unique features of the NELS data is that youth who leave high school prior to graduation continue to be interviewed throughout the longitudinal study and are asked the same questions pertaining to smoking behavior. It is therefore possible to examine the smoking behavior of all youth, including those not represented in other national school-based surveys such as Monitoring the Future. The NELS data contain information on the student's background, upbringing, early family environment, early school environment, and other behaviors. It provides many variables that have been found to be significant risk factors for smoking such as school performance, religious affiliation, family structure and living arrangement, and parental education. Since parents are surveyed in the base year and second follow-up, it is possible to obtain time-varying information on family background and socioeconomic characteristics that the student would not be as informed about. In the first and second follow-up, school principals and

TABLE 4.—SMOKING DYNAMICS IN NELS

Behavior	Cigarette Use (conditional on smoking)			
	Smoke Any	Smoke 1–5 Cigs	Smoke 6–10 Cigs	Smoke 11+ Cigs
Full Sample	12.80	59.64	20.32	20.04
1988	4.03	60.93	16.56	22.52
1990	13.91	66.15	17.70	16.15
1992	20.56	54.99	22.83	22.18

Persistence: % who continue prior behavior	Unconditional	Prior Behavior	
		Prior Nonsmoker	Prior Smoker
1990	97.09	88.07	62.94
1992	86.27	88.13	74.66

Note: Sample size is 11,146 person-year observations. All numbers are percentages.

teachers continue to be surveyed, making it possible to control for important school environmental characteristics as well.

We model the behavior of youths who are observed in each year (1988, 1990, and 1992) of the survey; we do not model attrition from the full sample. We keep only those youths who were on grade during the sample period or who were permanent dropouts (12,954 youths). We are forced to drop 2,237 youths for whom smoking behavior is unobserved. Because prices differ by state, another 270 are dropped if we cannot identify the state in which they live or go to school, 196 are dropped if they do not reside in the same state in all three waves, and 18 are deleted since important variables are missing. We finally omit individuals who do not live in one of the cities for which we have detailed price data. This leaves a sample consisting of 11,146 person-year observations.

Information on smoking behavior is collected in each wave of the survey. In each year, youths are asked, "How many cigarettes do you currently smoke in a day?" Responses are limited to the following categories: do not smoke, smoke less than one cigarette a day, smoke one to five cigarettes, smoke about a half pack (6–10), smoke more than half a pack but less than two packs (11–39), and smoke two packs or more (40+).

In general, adolescent smokers are older white youths with lower test scores and socioeconomic status than nonsmokers. They are more likely to have older siblings, to have siblings who dropped out of school, to have one parent absent from the home, and to report no religion.

The top panel of table 4 shows the rapid increase in smoking participation between 1988 and 1992. Among the 935 youth observed smoking at some point in the sample, only 16% began in 1988 while 45% started in 1990 and 39% started in 1992. The dramatic increase in smoking rates is not surprising given that smoking initiation typically occurs during the late teens. We also form indicators of the quantity smoked conditional on being a smoker. There are no clear trends in conditional use reported in table 4. Table 4 also shows that participation behavior is relatively persistent.

TABLE 5.—ESTIMATES OF SMOKING BEHAVIOR FROM NELS: HISTORICAL VOLATILITY MEASURES

	(Logit) Participation		(Cox proportional hazard) Hazard of Participation		(Multinomial Logit) Total Consumption: cigs/day					
					1-5	6-10	>10	1-5	6-10	>10
$Price_{it}$	-0.105 (0.021)	-0.053 (0.026)	-0.076 (0.025)	-0.048 (0.027)	-0.096 (0.026)	-0.145 (0.041)	-0.084 (0.042)	-0.052 (0.031)	-0.092 (0.052)	-0.000 (0.056)
$StdDev(Price_{it})$	-0.172 (0.096)	-0.053 (0.110)	-0.160 (0.105)	-0.043 (0.110)	-0.095 (0.116)	-0.107 (0.204)	-0.581 (0.231)	-0.034 (0.126)	0.138 (0.229)	-0.380 (0.265)
Year FE	Yes	Yes	No	No		Yes			Yes	
Individual Covariates	No	Yes	No	Yes		No			Yes	
N	11,146	11,146	10,437	10,437		11,146			11,146	
$\log L$	-3,981.28	-3,315.47	-7,595.68	-7,478.76		-5,331.13			-4,417.88	
Fitted Values										
$Y Data$	0.128	0.128	0.287	0.152	0.076	0.026	0.026	0.076	0.026	0.026
$Y Min(Std Dev(Price_{it}))$	0.147	0.132	0.324	0.158	0.081	0.028	0.043	0.078	0.022	0.034
$Y Max(Std Dev(Price_{it}))$	0.057	0.106	0.118	0.120	0.051	0.016	0.001	0.064	0.053	0.004

Note: Standard errors are in parentheses.

The bottom panel shows that over 85% of individuals continue with their most recent behavior in 1990 and 1992. This panel also shows that only an eighth of nonsmokers begin smoking in either 1990 or 1992. Alternatively, the percentage of smokers who continue smoking rises by 10 percentage points between 1990 and 1992.

B. Empirical Evidence

Our objective is to investigate whether cigarette demand is sensitive to price levels and price volatility. One hypothesis is that individuals are less likely to start smoking, and to consume fewer cigarettes if they already smoke, when prices are highly volatile. As we have seen in section II, theory predicts that greater price variation or higher prices make smoking less attractive for forward-looking and risk-averse individuals. To investigate these hypotheses, we separately estimate logit probabilities of cigarette smoking participation, Cox proportional hazard models of smoking initiation, and multinomial logit probabilities of total smoking consumption (smoking intensity is reported as a categorical variable in NELS). For each specification we are interested in how cigarette price levels and volatility influence smoking behavior.¹⁵ The equations we consider are

$$Y_{it} = \beta_0 + \beta_1 \times P_{it} + \beta_2 \times E_t[StdDev(P_{it+1})] + \gamma' X_{it} + e_{it}, \quad (17)$$

where Y_{it} is a measure of smoking behavior for individual i in period t , P_{it} is the cigarette price he faces, $E_t[StdDev(P_{it+1})]$ is the expected next period price volatility, and X_{it} are additional covariates. With larger Y_{it} indicating more smoking, one null hypothesis is that $\beta_1 < 0$: individuals reduce smoking if prices increase. The second

null hypothesis is that $\beta_2 < 0$: individuals smoke less if they face more price uncertainty.¹⁶

The dependent variables, Y_{it} , are a smoking indicator for participation logit models; first time smoking for Cox proportional hazards, and four smoking categories (with non-smoking the omitted category) for the total consumption multinomial logits. The individual covariates, X_{it} , are gender, race, age, previous smoking status, standardized test scores, religion, dropout indicator, sibling dropout indicator, family composition, family socioeconomic status, parents' education, income, and employment status, guardian's age, and school characteristics.¹⁷

Table 5 presents estimates of the parameters of demand models using historical volatility measures. We use the standard deviation of monthly cigarette prices over 24 months prior to the individual's survey date as the measure of price volatility. This presumes individuals have adaptive expectations about prices, and the two-year window is used since this is the typical period between interviews. Our estimates are broadly consistent with the two main hypotheses. The first two columns show that higher prices and price volatility reduce smoking participation. The only drawback is that the estimated coefficients of the price volatility measures have large standard errors once we allow for observed covariates.

The last three rows in these columns report the estimates that imply economically important effects. The fitted value Y is the proportion that smoke in the participation specifications, the relative hazard in the Cox hazards, and the proportion smoking in the listed category in total smoking. The fitted values reflect predictions using observed covariates (and the full set of parameter estimates) and then forcing the price standard deviation to the maximum or

¹⁶ We also reestimated the models using the expected future price instead of the current price and found no significant differences in the parameter estimates.

¹⁷ An alternative empirical framework that nests both discrete and continuous choice aspects is given, for example, by Gupta (1988) who uses a multinomial logit model of brand choice, and a cumulative logit model of purchase quantity.

¹⁵ An interesting extension of our analysis would look at smoking and drinking decisions jointly. Decker and Schwartz (2000) provide some evidence that higher alcohol prices decrease both alcohol consumption and smoking participation, suggesting a complementarity in consumption.

TABLE 6.—ESTIMATES OF SMOKING BEHAVIOR FROM NELS:88: FORECASTED VOLATILITY MEASURES

	(Logit) Participation				(Cox proportional hazard) Hazard of Participation						(Multinomial Logit) Total Consumption: cigs/day		
	1-5	6-10	>10	1-5	6-10	>10	1-5	6-10	>10	1-5	6-10	>10	
$Price_{it}$	-0.087 (0.020)	-0.029 (0.025)	-0.069 (0.023)	-0.031 (0.026)	-0.078 (0.025)	-0.115 (0.040)	-0.086 (0.041)	-0.032 (0.030)	-0.042 (0.051)	0.007 (0.054)			
$E_t[StdDev(Price_{it+1})]$	-0.201 (0.049)	-0.173 (0.056)	-0.162 (0.055)	-0.144 (0.056)	-0.164 (0.061)	-0.233 (0.100)	-0.278 (0.100)	-0.137 (0.065)	-0.239 (0.113)	-0.247 (0.115)			
Year FE	Yes	Yes	No	No	Yes	Yes	Yes	Yes	Yes	Yes			
Individual Covariates	No	Yes	No	Yes	No	No	No	No	Yes	Yes			
N	11,146	11,146	10,437	10,437	11,146	11,146	11,146	11,146	11,146	11,146			
$\log L$	-3,974.43	-3,310.82	-7,592.53	-7,475.62	-5,325.83	-5,325.83	-5,325.83	-5,325.83	-5,325.83	-5,325.83			
Fitted Values													
$Y Data$	0.128	0.128	0.210	0.101	0.076	0.026	0.026	0.076	0.026	0.026			
$Y Min(Std Dev(Price_{it}))$	0.158	0.148	0.257	0.121	0.090	0.033	0.035	0.086	0.031	0.031			
$Y Max(Std Dev(Price_{it}))$	0.095	0.104	0.157	0.078	0.060	0.018	0.016	0.065	0.020	0.020			

Note: Standard errors are in parentheses.

minimum in the data. Even after including a wide range of individual characteristics (column 2), a shift from the minimum to maximum price volatility in the data would reduce smoking participation by 1.7 percentage points. This is nearly a 10% reduction from the mean observed smoking rate.

The hazard estimates in the third and fourth columns show that both price levels and volatility reduce smoking take-up. To gauge the importance of the volatility effect, the last three rows report relative hazards implied by the parameters. After controlling for individual characteristics (column 4), an increase in the price standard deviation from the minimum to maximum would reduce the relative hazard by 4.4 percentage points or 17% of the relative hazard at the mean. The remaining six columns show that higher prices and price volatility reduce total cigarette consumption. The multinomial logits suggest that price variation markedly reduces heavy smoking intensity—smoking more than half a pack per day—and shifts individuals into the omitted nonsmoking category. The last three rows again show these effects are large even when including individual covariates.¹⁸

Table 6 reports estimates for the same demand models studied in table 5. Here we use forecasted volatilities based on regime-switching models instead of historical volatility measures. In general, we find that the main results of this study are robust to different measurements of price volatility. Estimates of the main effects are significant even after we control for observed covariates and fixed effects. We view this finding as strong evidence supporting our main hypothesis that even young individuals are forward looking and respond to increased price uncertainty by reducing consumption as predicted by our theoretical model. As another final robustness check, we use the parameter estimates of the GARCH models described in section IIIB to form price volatility measures. Price volatility continues to

reduce smoking participation and intensity. (Details are available in Coppejans et al., 2006.)

Finally, we also perform a series of non-nested tests of model selection comparing the specifications that use historical and forecasted price volatility. We use McFadden (1974) pseudo R^2 as a measure of goodness of fit. Model selection tests include Akaike’s information criterion, the Bayesian information criterion, and Vuong’s (1989) likelihood ratio test. We find that the specifications using forecasted price volatilities outperform those using historical price volatilities. The BIC statistics provide strong support in favor of the forecasted model. The Vuong test gives more tentative support, since the statistics are not statistically significant. Still, the statistic is negative in all cases, which provides some weak support in favor of the forecasted model. (Details are available in Coppejans et al., 2006.)

V. Conclusions

In this paper we focus on the demand for goods that are subject to habit formation or addiction in the presence of price uncertainty. Our theoretical model predicts that forward-looking individuals form beliefs about the distribution of prices in the future. Moreover, individual consumption plans depend crucially on beliefs they hold about future prices. To test the main implications of this model, we have assembled a unique data set to analyze the market for cigarettes. Our empirical findings suggest that consumers face considerable uncertainty about future market conditions. Prices and market conditions also vary significantly among the set of metropolitan areas analyzed in this study. The variation in price uncertainty across space and time thus allows us to test whether individuals respond to price uncertainty as predicted by our theoretical model.

We have constructed two types of measures of expected price variability: one based on adaptive expectations using historical volatilities and another based on forecasted price volatility using regime-switching and GARCH models. We have estimated reduced-form models of cigarette consumption that are based on the restricted-use version of NELS

¹⁸ The results reported in table 5 may be subject to omitted variable problems. Any metro-level variables that we have excluded from the analysis or are not measured precisely and are correlated with price volatilities would bias the results.

that allows us to match individuals to metropolitan areas. The empirical evidence confirms the main predictions of our model. We find that teenagers who live in metropolitan areas with large amounts of price volatility have, on average, significantly lower levels of cigarette consumption than individuals in low-volatility areas. Models based on forecasted price volatility fit the data better than models based on historical volatility measures. We thus conclude that young individuals are forward looking and respond to changes in price uncertainty.

Understanding the role that uncertainty and risk aversion play in determining consumption decisions of addictive goods has important policy implications. Individuals often face significant uncertainty about future tax policies. This uncertainty about future taxes is likely to affect consumer choices. Moreover, federal and state governments sometimes try to change behavior by announcing policies that may be implemented in the future. Our findings suggest that these policy announcements may be effective if they permanently change the beliefs that individuals hold about future prices. If, on the other hand, an announced tax increase is perceived to be temporary or if individuals believe that it is not likely to be implemented, then it will have, at best, modest effects on individual consumption. Tax policies thus not only affect prices in the period that they are announced or enacted, but they also affect beliefs about future prices. Announced policy changes can have large immediate effects if they are perceived to be credible.¹⁹

Our analysis provides ample scope for future research. Our findings illustrate the need to control for price expectations and uncertainty in empirical demand analysis. These findings, thus, raise a number of questions regarding the common practice of ignoring uncertainty about future prices in demand analysis or assuming perfect foresight about prices. Our estimates of the pricing processes reflect the observed equilibria in the regional markets. From the perspective of consumers that is all that matters. It is, however, an interesting question to ask what supply-side models would yield price processes that are similar to the ones we observe in the data. Future research should help us understand how supply-side conditions interact with demand models of the type considered in this paper to generate equilibria that exhibit not only large price fluctuations across time but also the large degree of spatial price dispersion.

We would like to thank Daron Acemoglu, editor of the review, and two anonymous referees for helpful comments. We would also like to thank Tim Bollerslev, Martin Gaynor, Bruce Hansen, Carolyn Levine, Frank Sloan, and V. Kerry Smith for comments and suggestions. We also thank Derek Brown and Justin Trogden for research assistance. The data used in this paper were provided by the Bureau of Labor

Statistics and by the Department of Education. We would like to thank Bill Cook, Roger von Haefen, Pat Jackman, David Johnson, and Lois Orr at the BLS. Financial support was provided by grants from the NIH-NIAAA (R01 AA12162-01) and the National Institute of Child Health and Human Development (1R01 HD42256-01). The opinions expressed here do not necessarily reflect those of Barclays Global Investors.

REFERENCES

- ACCRA, *ACCRA Cost of Living Index*, Arlington, VA (2003).
- Appelbaum, E., and A. Ullah, "Estimation of Moments and Production Decisions under Uncertainty," *this REVIEW* 79:4 (1997), 631–637.
- Arcidiacono, P., H. Sieg, and F. Sloan, "Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking," *International Economic Review* 48:1 (2007), 37–65.
- Becker, G., M. Grossman, and K. Murphy, "Rational Addiction and the Effect of Price on Consumption," *American Economic Review* 81:2 (1991), 237–242.
- "An Empirical Analysis of Cigarette Addiction," *American Economic Review* 84:3 (1994), 396–418.
- Becker, G., and K. Murphy, "A Theory of Rational Addiction," *Journal of Political Economy* 96:4 (1988), 675–700.
- Bernheim, D., and A. Rangel, "Addiction and Cue-Triggered Decision Processes," *American Economic Review* 94:5 (2004), 1558–1590.
- Bollerslev, T., "Generalized Autogressive Conditional Heteroskedasticity," *Journal of Econometrics* 31 (1986), 307–327.
- Cecchetti, S., M. Nelson, and R. Sonora, "Price Index Convergence Among United States Cities," *International Economic Review* 43 (2002), 1081–1109.
- Chaloupka, F., "Rational Addictive Behavior and Cigarette Smoking," *Journal of Political Economy* 99:4 (1991), 722–742.
- Chaloupka, F., and K. Warner, "The Economics of Smoking" (pp. 1151–1227), in *Handbook of Health Economics* (North Holland, 2000).
- Coppejans, M., D. Gilleskie, H. Sieg, and K. Strumpf, "Consumer Demand under Price Uncertainty: Empirical Evidence from the Market of Cigarettes," NBER working paper (2006).
- DeCicca, P., D. Kenkel, and A. Mathios, "Putting Out the Fires: Will Higher Taxes Reduce the Onset of Youth Smoking?" *Journal of Political Economy* 110:1 (2002), 144–169.
- Decker, S., and A. Schwartz, "Cigarettes and Alcohol: Substitutes or Compliments?" NBER working paper 7535 (2000).
- Erdem, T., S. Imai, and M. Keane, "Brand and Quantity Choice Dynamics Under Price Uncertainty," *Quantitative Marketing and Economics* 1 (2003), 5–64.
- Erdem, T., and M. Keane, "Decision Making Under Uncertainty: Capturing Choice Dynamics in a Turbulent Consumer Market," *Marketing Science* 16 (1994), 1–21.
- Gilleskie, D., and K. Strumpf, "The Behavioral Dynamics of Youth Smoking," *Journal of Human Resources* 40:4 (2005), 822–866.
- Gruber, J., and B. Koszegi, "Is Addiction Rational: Theory and Evidence," *Quarterly Journal of Economics* 116:4 (2001), 1261–1304.
- Gupta, S., "Impact of Sales Promotions on When, What, and How Much to Buy," *Journal of Marketing Research* 25 (1988), 342–355.
- Hall, G., and J. Rust, "Econometric Models for Endogenously Sampled Time Series: The Case of Commodity Price Speculation in the Steel Market," working paper (2002).
- Hamilton, J., "A New Approach to the Economic Analysis of Nonstationary Time Series and Business Cycles," *Econometrica* 57 (1989), 357–384.
- "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics* 45 (1990), 39–70.
- *Time Series Analysis* (Princeton, NJ: Princeton University Press, 1994).
- Hansen, B., "The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP," *Journal of Applied Econometrics* 7 (1992), S61–S82.
- Harris, C., and D. Laibson, "Dynamic Choice of Hyperbolic Consumers," *Econometrica* 69:4 (2001), 935–958.
- Hendel, I., and A. Nevo, "Measuring the Implications of Sales and Consumer Stockpiling Behavior," working paper (2002).

¹⁹ Gruber and Koszegi (2001) provide some empirical evidence in favor of this hypothesis.

- Khawaja, A., "Health Insurance, Habits and Health Outcomes: A Dynamic Stochastic Model of Investments in Health," working paper (2006).
- Levin, A. and C. Lin, "Unit Root Tests in Panel Data: New Results," Board of Governors of the Federal Reserve System working paper (1993).
- McFadden, D., "The Measurement of Urban Travel Demand," *Journal of Public Economics* 3 (1974), 303–328.
- Muth, J., "Rational Expectations and the Theory of Price Movements," *Econometrica* 29 (1961), 315–335.
- Nerlove, M., "Adaptive Expectations and Cobweb Phenomena," *Quarterly Journal of Economics* 73 (1958), 227–240.
- Orphanides, A., and D. Zervos, "Rational Addiction with Learning and Regret," *Journal of Political Economy* 103:4 (1995), 739–758.
- Osborne, T., "Market News in Commodity Price Theory: Application to the Ethiopian Grain Market," *Review of Economic Studies* 71:1 (2004), 133–164.
- Vuong, Q., "Likelihood Ratio Test for Model-Selection and Non-Nested Hypotheses," *Econometrica* 57 (1989), 307–334.
- Wolak, F., and C. Kolstad, "A Model of Heterogeneous Input Demand Under Price Uncertainty," *American Economic Review* 81:3 (1991), 514–38.

APPENDIX

Proof of Proposition 1

To prove the result, we first show that a mean preserving spread in the price distribution for a single period will reduce smoking along the optimal path under the functional forms used in the main analysis.

Suppose prices in a single period $t = s$ are stochastic. Let p_s follow the density $f(p_s|\sigma)$ where σ characterizes the dispersion. The path of all other variables is known with certainty. We continue to assume that period utility follows equation (7), the budget constraint follows equation (4), and the addiction stock satisfies the law of motion in equation (1). An individual's optimization problem is

$$\max_{a,c} \int \sum_t \beta^{t-1} (\ln(c_t) + \ln(a_t) + S_t^\psi (\phi + \gamma a_t)) f(p_s|\sigma) dp_s$$

$$s.t. \quad S_{t+1} = \delta S_t + a_t$$

$$c_t = y_t - p_t a_t \quad (A1)$$

$$\{S_t, S_{t-1}, \dots\}, \{y_t \forall t\}, \{p_t \forall t \neq s\} \in \Omega_t,$$

where Ω_t is the information set at time t (we specify below when $\sigma \in \Omega_t$).

After substituting in the constraints and optimizing over the smoking choices, we have the first-order conditions,

$$\int \left(\frac{-p_t}{y_t - p_t a_t} + \frac{1}{a_t} + \gamma S_t^\psi + \beta W_t(a_t) \right) f(p_s|\sigma) dp_s = 0 \quad \forall t, \quad (A2)$$

where

$$\begin{aligned} W_t(a_t) &\equiv \psi S_{t+1}^{\psi-1} (\phi + \gamma a_{t+1}) + \sum_{r>t} \beta^{r-t-1} \\ &\times \left(\frac{-p_r}{y_r - p_r a_r} + \frac{1}{a_r} + \gamma S_r^\psi + \beta \psi S_{r+1}^{\psi-1} (\phi + \gamma a_{r+1}) \right) \frac{\partial a_r}{\partial a_t} \quad (A3) \\ &= \psi S_{t+1}^{\psi-1} (\phi + \gamma a_{t+1}). \end{aligned}$$

The second equality in equation (A3) follows from applying the $t + 1$ first-order condition, and $W_t(a_t)$ depends on a_t through its effect on S_{t+1} and a_{t+1} . Along the optimal path the second-order condition must be satisfied (the value function is concave).

$$\frac{\partial f \left(\frac{-p_t}{y_t - p_t a_t} + \frac{1}{a_t} + \gamma S_t^\psi + \beta W_t(a_t) \right) f(p_s|\sigma) dp_s}{\partial a_t} < 0. \quad (A4)$$

Finally, $c_t > 0 \rightarrow y_t - p_t a_t > 0$, which implies

$$\frac{\partial \frac{-p_t}{y_t - p_t a_t}}{\partial p_t} = \frac{-y_t}{(y_t - p_t a_t)^2} < 0 \quad (A5)$$

$$\frac{\partial^2 \frac{-p_t}{y_t - p_t a_t}}{\partial p_t^2} = \frac{-2a_t y_t}{(y_t - p_t a_t)^3} < 0 \Leftrightarrow y_t - p_t a_t > 0.$$

This means the first-order condition is concave in current prices, since p_t directly enters equation (A2) only through the first term.

Now we consider the effect of a mean preserving spread in the price distribution (an increase in σ) for the single period $t = s$. Suppose initially that this is unanticipated until $t = s$, so optimal $a_t \forall t < s$ are unaffected. The mean preserving spread reduces the left side of equation (A2) due to the concavity result in equation (A5). To maintain optimality, a_s (the only free variable at $t = s$) adjusts: a_s must decrease since this will increase (A2) via (A4). $a_t \forall t > s$ decline by an induction argument. The addiction stock $S_t \equiv \delta S_{t-1} + a_{t-1}$ falls, since the induction assumption states a_{t-1} declines and so S_{t-1} falls for $t > s + 1$ (S_s is unchanged). This means a_t declines due to the usual adjacent complementarity argument: the complementarity and negativity assumptions on preferences (Orphanides and Zervos, 1995) imply

$$\psi, \gamma \geq 0 \rightarrow \frac{\partial f \left(\frac{-p_t}{y_t - p_t a_t} + \frac{1}{a_t} + \gamma S_t^\psi + \beta W_t(a_t) \right) f(p_s|\sigma) dp_s}{\partial S_t} \geq 0. \quad (A6)$$

So when S_t decreases, equations (A2), (A4), and (A6) require that a_t (the only free variable in equation [A2]) also decreases. Now suppose that the mean preserving spread is anticipated in some period $p \leq s$. This will reduce a_p following the adjacent complementarity argument above, and the smoking level continues to decline during and after the change in σ . Combining all the results,

$$\frac{\partial a_t}{\partial \sigma} \leq 0 \quad \forall t \geq p, \quad (A7)$$

where p is the first period where the mean preserving spread is anticipated.

The fact that a mean preserving spread in the price distribution for each and every period will reduce smoking along the optimal path can then be shown by induction follows from repeated application of the argument above.