IV. Price and Purity Models

Methodological Changes

A number of methodological changes have been made for the price and purity econometric models. The most significant modification is the adoption of the Expected Purity Hypothesis (EPH), which develops an empirical model of price based on the assumption that it is the buyer's perception of purity at the time of the transaction, not actual purity of the drug, that determines the price he or she is willing to pay for the drug. Illicit drugs are what economists refer to as "experience goods"; purchasers often cannot readily assay the quality of the drug purchased until it is consumed, which generally occurs after a price is negotiated and the deal is completed. Hence, it is typically not the actual purity of the drug that governs the negotiated price at the time of the transaction, but rather the supposed or expected purity of the drug. For example, one might observe that most transactions of a particular drug at a particular time, place, and transaction size are 60 to 80 percent pure, but that a minority have very low or even zero purity, even though the price paid for these very-low-purity drugs is not noticeably lower. The view implicitly adopted by past statistical methods was that purchasers of these low-purity observations were knowingly paying much more, sometimes ten or more times as much, per pure gram than were most customers. The view implicit in the EPH models is that these customers were ripped off; they paid a price typical of 60 to 80 percent pure transactions because they thought or expected that they were buying drugs that were 60 to 80 percent pure. In the EPH model, these low-purity transactions are not discarded; they represent a real cost to customers. Therefore, they are incorporated into expectations of the pure quantity contained in purchases, on average, rather than assuming that they represent fully informed purchases.[14]

The adoption of the EPH has two important implications for the way the data get analyzed. First, observations with low purity, but not zero purity, are retained in the analysis, provided they meet other general criteria for inclusion. Second, price is estimated through a two-step procedure where expected purity rather than actual purity is included in the price regression model. Expected purity is the predicted value obtained from a first-stage regression where actual purity is estimated as a function of all other observable information available to the buyer and included in the data (amount, city, quarter, year). Because expected purity is far less volatile than actual purity, the EPH model generally produces smoother price series, even when relatively fewer data points are available (e.g., when estimating prices for a specific city). Failing to use the EPH model can either inflate or suppress the estimated price level somewhat, depending on the details of the distribution of purities observed and whether and how many low-purity observations are discarded. Thus, it is not appropriate to compare the level of prices produced by an EPH model with that produced using a non-EPH method.

---

[14] This framework for estimating price series from STRIDE has been widely adopted in the economics literature. See, for example, Caulkins, J.P. (1994), *Developing Price Series for Cocaine*, MR-317-DPRC, RAND, Santa Monica, CA; Saffer H., and F.J. Chaloupka (1995), "The Demand for Illicit Drugs," National Bureau of Economic Research Working Paper #5238, August; Grossman M., F.J. Chaloupka, and C.C. Brown (1998), "The Demand for Cocaine by Young Adults: A Rational Addiction Approach," *Journal of Health Economics*, Vol 17, No. 4, pp. 427–474; DeSimone, J. (2002), "Illegal Drug Use and Employment," *Journal of Labor Economics*, 20(4), pp. 952–977.

A related change is the use of inflation-adjusted price, not inflation-adjusted price per pure gram, as the dependent variable in the statistical regression models. Because quantity levels are defined in terms of grams unadjusted for purity, it no longer makes sense to estimate the models in terms of the inflation-adjusted price per pure gram. Further, including actual (or expected) purity in the denominator of the dependent variable (i.e., estimating price per pure gram) causes the coefficient on the actual (or expected) purity variable on the right-hand side of the equation to be biased and hence will generate biased predictions.

A third methodological improvement is that both the purity and price models are estimated using hierarchical modeling (HM), which offers at least three principal advances over the previous methodology employed. First, it directly accounts for the nested nature of the data being used and adjusts standard errors and variance-covariance matrices to account for the fact that specific clusters of observations (in this case, observations from the same city) are not entirely independent. The error terms across observations from within a city are allowed to be correlated to account for city-specific unmeasured components of price (or purity). Second, it adjusts the variance-covariance matrix to account for unequal variances in error terms across different cities, which could result because of different unobservables that exist across cities and different sample sizes across cities. Finally, HM is highly flexible and allows each city to have unique relationships between price and the other independent variables. The methodology employed in previous reports allowed the price levels to differ from city to city (through a city-specific intercept term), but the relationship between other variables, such as the amount of the transaction, and price was assumed to be constant across all cities. The interpretation of this restriction is that within a specific market level, quantity discounts across cities are all the same. This assumption is likely to be overly restrictive. With HM, the relationship between price and amount (or any other independent variable) can vary across cities and over time. The specific form of the HM model employed here is a random coefficients model.

Price/Purity Model Specifications

We estimate the price/purity model for each distribution level for each drug. As described above, the model is estimated using a two-step procedure that involves first estimating purity. The purity model is implemented without zero-potency observations. The previous report estimated the purity model as a logistic model because the dependent variable should be constrained to be between 0 and 1. However, using a simple linear specification of the random coefficients model, very few cases have an estimated purity that exceeds those bounds. Thus, a linear specification of the model is retained here. In addition, in contrast to the previous report, the amount of the transaction (measured as weight in grams) is included as an additional regressor.

The empirical specification of the random coefficient purity model can be written as

$$Purity_{ijk} = \alpha_{0k} + \alpha_{1k} time_{ij} + \alpha_{2k} AMT_{ijk} + \varepsilon_{ijk} \qquad (1)$$

$$\alpha_{0k} = \gamma_0 + u_{0k}$$

$$\alpha_{1k} = \gamma_1 + u_{1k}$$

$$\alpha_{2k} = \gamma_2 + u_{2k}$$

$$\varepsilon_{ijk} \sim N(0,\sigma^2)$$

$$\begin{pmatrix} u_{0k} \\ u_{1k} \\ u_{2k} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_{00} & \tau_{01} & \tau_{02} \\ \tau_{10} & \tau_{11} & \tau_{12} \\ \tau_{20} & \tau_{21} & \tau_{22} \end{pmatrix} \right)$$

where $time_{ij}$ is a vector of dummy variables representing a year-quarter (i.e., 90 quarters over 22.5 years) and AMT_{ijk} is the raw weight of the i^{th} observation in city *k* at time *j*. The coefficient α_{0k} represents the intercept for city k, α_{1k} is a vector for the time coefficient for city k, and α_{2k} is the amount coefficient for city k. The terms, γ_0, γ_1, and γ_2, respectively, are the overall mean estimates for the intercept, time, and amount effects. The random coefficients for the intercept, amount, and time are assumed to be independently and identically distributed as specified above. The interpretation of the slope coefficients in a random coefficient model is similar to that in other OLS models; α_{1k} is a vector of coefficients identifying time (year-quarter) effects for each city *k*. Test statistics used to evaluate the appropriateness of this functional form versus a simpler random-intercept model are provided in Appendix B.

Equation (1) is estimated as a general linear random effects model, and the estimates of expected purity generated from this model are then used in the second stage price model. For the few cases in which the predicted purity was below 0 percent or above 100 percent, we modified the prediction to equal either 0.5 or 99.5 percent, respectively.

Price is modeled using another general linear random effects model. This model was described in McCullagh and Nelder (1989) and used in ONDCP (2001).[15] The equation is the following:

$$E(\text{real price}_{ijk}|\gamma_{0k}, \gamma_{1k}, \gamma_{2k}) = \exp[[\gamma_{0k} + \gamma_{1k}\text{time}_j + \gamma_{2k}[\ln(\text{AMT}_{ijk})+\ln(\text{predicted purity}_{ijk})] \quad (2)$$

$$\gamma_{0k} = \beta_0 + \varepsilon_{0k}$$
$$\gamma_{1k} = \beta_1 + \varepsilon_{1k}$$
$$\gamma_{2k} = \beta_2 + \varepsilon_{2k}$$

$$\begin{pmatrix} \varepsilon_{0k} \\ \varepsilon_{1k} \\ \varepsilon_{2k} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_{00} & \tau_{01} & \tau_{02} \\ \tau_{10} & \tau_{11} & \tau_{12} \\ \tau_{20} & \tau_{21} & \tau_{22} \end{pmatrix}\right)$$

$$\text{Var}(\text{real price}_{ijk}|\gamma_{0k}, \gamma_{1k}, \gamma_{2k}) = \lambda^2 [E(\text{real price}_{ijk}|\gamma_{0k}, \gamma_{1k}, \gamma_{2k})]^2 \quad (3)$$

In Equation (2), the real price for observation *i* in period *j* in city *k* is estimated as a function of time effects, city effects, and the sum of the natural logarithm of amount and the natural logarithm of expected purity. Note that this last term is just a long-form way of specifying the price in terms of the natural logarithm of expected pure grams.[16] Average mean effects of these control variables on price are captured through the coefficients β_0, β_1, and β_2. The coefficients γ_{0k}, γ_{1k}, and γ_{2k} are the city-specific intercept, time, and ln(expected pure grams) coefficients, respectively, each of which comes from a common normal distribution centered at 0. We estimate Equation (2) as a hierarchical generalized linear model with a log link function, a gamma error structure, and a constant coefficient of variation. This model transforms the dependent variable to the desirable log form during estimation, which essentially puts the coefficient estimates in percentage terms. Equation (3) shows the conditional variance function of this model, where λ is the coefficient of variation of the real price.[17] Additional test statistics that provide information on the appropriateness of the functional form employed here versus a simpler random intercept model are provided in Appendix C.

Summing the natural logarithms of amount and predicted purity gives us identification in the model, which is necessary because all of the explanatory variables in the purity model of Equation (1) are in the price model of Equation (2). Furthermore, such summing provides a specification of price in which price is a function of the expected pure grams involved in the transaction.

---

[15] (1) McCullagh, P., and J.A. Nelder (1989), *Generalized Linear Models*, Second Edition, Chapman and Hall, London, Chapter 8; Office of National Drug Control Policy (2001), "The Price of Illicit Drugs: 1981 through the Second Quarter of 2000," Washington, DC; Littell, Ramon C. , George A. Milliken, Walter W. Stroup, and Russell D. Wolfinger (1996), *SAS System for Mixed Models*', SAS Institute, Inc., Cary, NC, Chapter 11.

[16] In other words, [(ln AMT) + ln (expected purity)] = ln (AMT*expected purity). When we multiply the weight by expected purity we are generating an estimate of the weight in expected pure grams.

[17] McCullagh P. and Nelder J.A. (1989). Generalized Linear Models, Second Edition. Chapman and Hall, London, p. 285.

As discussed previously, transactions that deviate significantly from the norm could have large effects on the coefficient estimates and the predicted prices. These deviations could occur, for example, from miscoded data or poor bargaining on the part of a DEA agent. While the gross-outlier restrictions eliminate many of these observations from the sample as a whole, some still survive because they are gross outliers for a time period or a specific quantity level (even if they are not gross outliers for the full period or all quantities combined). To reduce the potential influence of these outliers, extreme-residual analysis is performed following criteria employed in previous reports. Operationally, the model is estimated and the residuals from the model are kept, standardized, and plotted. Using the same criterion as that employed in previous reports, observations that fell beyond 3.09 were deleted (so the probability of deleting good data is set to 0.002). This process of reestimating the model and deleting residuals is continued until no extreme residuals exist in the sample. Most of the models lost between 3 and 5 percent of the sample, with the smallest loss being the third quantity level for d-methamphetamine (2.2 percent) and the largest loss being the first quantity level for marijuana (7.4 percent).

Because these models are highly parameterized and separate models are run for specific quantity levels that are determined after gross outliers have been deleted, it is difficult in some cases to achieve convergence. Problems with convergence could be due to the data not fitting well in our imposed functional form, or they could be caused by outliers in a city with few observations, which would make it difficult to identify a city-specific slope coefficient. To assist in convergence, a simpler second stage price model is estimated to identify the first round of extreme residuals for each quantity level. These simpler models allow for only a random intercept, so slope coefficients are forced to be constant across cities. With the reduced parameterization of the model, the models quickly converge, and it is possible to identify many extreme residuals. Once these first round extreme residuals are dropped, the full random coefficients model expressed in Equations (2) and (3) is estimated and convergence is obtained for all quantity levels. Future work should explore the feasibility of model alternatives to accommodate or downweight the influence of extreme residuals within a particular quantity level.

Marijuana Price Model Specification

A modified model is developed for marijuana because no information on marijuana purity is available and because the smaller number of observations makes it difficult to fit the highly parameterized model developed for the other drugs.

The marijuana model specification is as follows:

$$E(\text{real price}_{ijk}|\gamma_{0k},\gamma_{1k},\gamma_{2k}) = \exp[\gamma_{0k} + \beta_1 \text{quarter}_{ij} + \gamma_{1k}\text{year}_{ij} + \gamma_{2k}\ln(\text{AMT}_{ijk})] \quad (4)$$

$$\gamma_{0k} = \alpha_0 + \varepsilon_{0k}$$
$$\gamma_{1k} = \alpha_1 + \varepsilon_{1k}$$
$$\gamma_{2k} = \alpha_2 + \varepsilon_{2k}$$

$$\begin{pmatrix} \varepsilon_{0k} \\ \varepsilon_{1k} \\ \varepsilon_{2k} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_{00} & \tau_{01} & \tau_{02} \\ \tau_{10} & \tau_{11} & \tau_{12} \\ \tau_{20} & \tau_{21} & \tau_{22} \end{pmatrix}\right)$$

$$\text{Var}(\text{real price}_{ijk}|\gamma_{0k},\gamma_{1k},\gamma_{2k}) = \sigma^2\,[E(\text{real price}_{ijk}|\gamma_{0k},\gamma_{1k},\gamma_{2k})]^2 \quad (5)$$

There are two primary differences between this model and the model described Equation (2). First, because no information is available on the purity of marijuana, price is estimated in terms of actual amount and not expected pure weight (i.e., the log(AMT)). Second, due to the significantly smaller sample size, it is not possible to estimate individual city-specific interacting time trends for quarters and years as we did in the previous models. Instead, the model includes year trends (captured through a series of dichotomous indicators for each year) and quarter trends separately. Further, only the year effect varies randomly across cities; the influence of specific quarters is fixed across all years and cities. Finally, for the first quantity level of marijuana, the specification further restricts year to be constant across cities, so γ_{1k} is set equal to just α_1 in that model only. The strategy of deleting extreme residuals employed in the previous models is also used here.