1. Introduction

Illicit drugs are ultimately consumer goods, and like other goods in modern societies, they are provided primarily through markets. Prices play a prominent role in understanding, analyzing, and intervening in markets of all kinds, illicit as well as licit. Obtaining national or even local price and purity information for illicit drugs is challenging, however, for a variety of reasons. Some challenges are largely unavoidable, for example, the need to rely on administrative datasets not designed for tracking prices. Other problems, such as the fact that drugs are not generally sold in standardized quantities or qualities, can largely be resolved by using appropriate statistical techniques. Because of these complexities, greater effort has traditionally been devoted to collecting and reporting data related to demand and quantities consumed (e.g., National Household Survey on Drug Abuse (NHSDA)-based estimates of numbers of users) than to data on prices. That is unfortunate, because (1) prices affect drug use and consumption; (2) many outcomes of interest relate to expenditure, which is the product of price and quantity consumed; and (3) price data are a potentially important tool for understanding the workings of drug markets and interventions intended to control those markets.

This report continues a series produced by the Office of National Drug Control Policy (ONDCP) that seeks to improve understanding of trends in prices for five major illicit drugs. It provides updated estimates of the price and purity of powder cocaine, crack cocaine, heroin, and d-methamphetamine and of the price of marijuana in the United States from 1981 through the second quarter of 2003, using data from the Drug Enforcement Agency's (DEA's) System to Retrieve Information from Drug Evidence (STRIDE) database.

Since the 1970s, the Intelligence Division of DEA has been recording in STRIDE information obtained from seizures, purchases, and other drug acquisition activities conducted by undercover agents and informants from federal and, in some locations, local law enforcement agencies. STRIDE is a forensic database, designed primarily to control the inventory of drug acquisitions in the laboratories and to provide scientific data regarding the quality and quantity of the substances collected, for judicial processes. The data included in the STRIDE database represent only those acquisitions that are sent to a DEA laboratory for analysis and thus exclude most of the very large number of purchases and seizures that are made by state and local agencies. Because the data are not collected for analytical purposes, they do not reflect a random sample of all drug transactions that occur within any geographic location.[1,2] Instead they represent a "convenience sample," or observations that are obtained in response to purposeful decisions made by law enforcement agencies investigating specific drug-related activities. The timing and location of encounters are not only unrepresentative they are erratic. The number of observations from a given location can vary dramatically from year to year. This has implications for how these data can be used, in terms of both which observations should be retained within a sample and the statistical methods that should be used to analyze them.[3, 4]

---

[1] Frank, R.S. (1987), "Drugs of Abuse: Data Collection Systems of DEA and Recent Trends," *Journal of Analytical Toxicology*, Vol. 11, pp. 237–241 (Nov./Dec.).
[2] Manski, Charles F., John V. Pepper, and Carol V. Petrie (eds.) (2001), *Informing America's Policy on Illegal Drugs: What We Don't Know Keeps Hurting Us*, National Academy Press, Washington, DC.
[3] Horowitz, Joel L. (2001), "Should the DEA's STRIDE Data Be Used for Economic Analysis of Markets for Illegal Drugs?" *JASA*, Vol. 96, No. 456, pp.1254–1271.
[4] Manski, Charles F., John V. Pepper, and Carol V. Petrie (eds.) (2001), *Informing America's Policy on Illegal Drugs: What We Don't Know Keeps Hurting Us*, National Academy Press, Washington, DC.

Even with these limitations, the STRIDE database is the best source of information on illicit drug prices and purity currently available. No other database provides as much objective information on the characteristics of specific drug acquisitions over time or for as many geographic areas in the United States. Furthermore, although the data represent a convenience sample rather than a probability sample, they may still convey valid and useful information regarding changes in price and purity if they are used properly. Indeed, there are parallel examples in the business world of price indices constructed from convenience samples, including the ACCRA Local Cost of Living Index. Such indices are constructed and examined despite their known limitations, in large measure because they are highly correlated with other data that are drawn from independent, probabilistic samples. In the case of illicit drugs, there is a growing literature demonstrating that price series generated from the STRIDE data are also significantly correlated with related series constructed from independent probability samples, such as trends in drug use and drug-related consequences.[5,6,7] Further, the STRIDE data may also be quite informative about trends in purity. STRIDE's limitations are most problematic for assessing absolute levels of prices and purities, rather than trends, but even in the case of absolute levels, STRIDE data can be informative.

Methodological Changes

In an effort to be responsive to criticisms raised by the National Research Council regarding past price indices constructed from STRIDE data[8], we have made a number of methodological changes that enable us to use the data more carefully. First, in the development of price and purity indices, we have tried not to aggregate across different drug forms unless the drug forms are indistinguishable to the buyer at the time of the transaction. In addition, price and purity series are estimated only for drug forms for which sufficient data exist across time. For example, the current report presents price and purity series separately for cocaine hydrochloride (essentially powder cocaine) and cocaine base (predominantly crack cocaine). Observations pertaining to other forms of cocaine are dropped from the analysis because they (1) are physically distinguishable from the other two forms of cocaine and (2) are insufficient in number to permit estimates of their own price and purity series. A similar approach is taken with the other drugs. Thus it is important to clarify what is meant by specific drug names used in this report:

- *Powder cocaine* refers to cocaine hydrochloride.
- *Crack cocaine* refers to cocaine base.
- *Heroin* refers only to heroin base and heroin hydrochloride.
- *Methamphetamine* refers only to the d-forms of methamphetamine.

---

[5] Caulkins, Jonathan P. (1999), "Can Supply Factors Suppress Marijuana Use by Youth?" Federation of American Scientists' *Drug Policy Analysis Bulletin*, Issue No. 7, pp. 3–5; Caulkins, Jonathan P. (2001), "The Relationship Between Prices and Emergency Department Mentions for Cocaine and Heroin," *American Journal of Public Health*, Vol., 91, No. 9, pp. 1446–1448.
[6] Saffer and Chaloupka (1999), "The Demand for Illicit Drugs" *Economic Inquiry* 37(3): pp. 401–411.
[7] DeSimone J. and M. Farrelly. 2003. "Price and Enforcement Effects on Cocaine and Marijuana Demand" *Economic Inquiry* 41(1): 98-115; DeSimone, J. (2001), "The Effect of Cocaine Prices on Crime," *Economic Inquiry* 39(4), pp. 627–643.
[8] Manski, Charles F., John V. Pepper, and Carol V. Petrie (eds.) (2001), *Informing America's Policy on Illegal Drugs: What We Don't Know Keeps Hurting Us,* National Academy Press, Washington, DC.

- *Marijuana* refers to plant material (and not whole plants or seeds).

A second methodological change from previous reports is the use of a random coefficient regression model, which enables observations from one city to have a unique relationship between price and quantity that is different from the price/quantity relationships in other cities. The justification for this model, which was empirically tested and validated, is the possibility that drug markets behave differently in different locations. Therefore, instead of imposing the same relationship between price and quantity across all locations for which we have data, the random coefficient model groups observations by cities and then estimates the relationship between price and quantity by city, using all the available data. Predicted standardized prices (and purities) for each city are calculated for each quarter or year from this model. These predicted standardized prices (and purities) are then weighted to generate the national price (purity) indices reported here.

This report incorporates two further methodological changes in an attempt to capture and describe the considerable variability observed in the price and purity of illicit drugs. First, price and purity series for selected cities are presented, along with the national series. There can be clear and sometimes rather pronounced differences in price and purity across cities, just as there is geographic variation in prices for licit goods, such as houses. Examining only aggregate, national series that represent some composite or average of the city-specific series can obscure the extent of this spatial variation. City-specific series make it possible to evaluate how price and purity move in geographically smaller markets. They also help confirm whether apparent "national" trends in the aggregate series are really nationwide trends, and not merely trends in some regions, and even whether the apparent national trends might be spurious artifacts that emerge because of STRIDE's nonrandom sampling. Second, estimates of the variation in predicted prices (and purities) across cities are represented by the identification of the 25th and 75th percentile value of standardized predicted prices (purities) in addition to the average price (purity) index.[9] Often the gap between the 25th and 75th percentile values, known as the *interquartile range*, is wide, showing that at any given point and time, a wide range of prices and purities can be observed. Hence, successive samples—say, from different quarters—can yield somewhat different averages even if there is no change in the underlying distribution of prices. The 25th and 75th percentile observations enable the reader to judge whether trends in these series are truly meaningful without imposing any sort of distributional assumptions on the series on which they are based. Changes in price (or purity) that remain within this interquartile range are more likely to reflect only sampling changes over time, rather than real movements in price (or purity), in contrast to changes that extend beyond the interquartile range. The interquartile range should not be interpreted as confidence intervals, because a substantial amount of variation (50 percent) still exists outside of these bands.

A number of other improvements have also been made. Notably, distribution levels, referred to more accurately in this report as *quantity levels*, are identified on the basis of amounts purchased, unadjusted for potency, rather than on pure quantities. This change facilitates interpretation, since it is sometimes more natural to think of quantity levels in terms of actual quantities transacted rather than of quantities adjusted for purity. More fundamentally, it ensures that wholesale transactions with very low purities (including rip-offs) are not inadvertently

[9] The construction of the interquartile range estimates is based on the model and only holds as far as the model holds. They do not represent actual data points.

lumped together with smaller retail transactions. STRIDE observations span a continuum of quantities; there are not just a few well-defined transaction sizes, as there are in some licit markets (for example, milk in supermarkets is sold predominantly in pints, quarts, half-gallons, and gallons). Hence, boundaries between quantity levels are not well-defined, and their selection is somewhat arbitrary. Thus, it might be just as reasonable to include among "retail"-level cocaine transactions either transactions of 0.2 to 4.0 grams or transactions of 0.1 to 2.0 grams. We sought to define quantity levels with roughly equal numbers of observations in each level but with round-number boundaries. For all drugs except powder cocaine, three levels were identified, each of which in most cases contains between 25 and 50 percent of the total number of observations for that drug. In the case of powder cocaine, enough data were available to identify a fourth quantity level. Given that these boundaries were largely data-driven, little meaning should be assigned to the labels applied to them.

Another improvement is the adoption of the expected purity hypothesis (EPH). Illicit drugs are what economists refer to as "experience goods"; purchasers often cannot readily assay the quality of the drug until it is consumed, which generally occurs after a price is negotiated and the deal is completed. Hence, the *actual* purity of the drug does not typically govern the negotiated price at the time of the transaction, but rather the supposed or *expected* purity of the drug. For example, it might be observed that most transactions of a particular drug at a particular time, place, and transaction size are 60 to 80 percent pure, but a minority have very low or even zero purity although the price paid for these very low-purity drugs is not noticeably lower. The view implicitly adopted by past statistical models was that purchasers of low-purity drugs were knowingly paying much more—sometimes ten or more times as much—per pure gram than were most customers because actual purity (and not expected purity) was included in the model. 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 80percent pure. These low-purity transactions are not discarded; they represent a real cost to customers. In the EPH, they are incorporated into expectations of the pure quantity contained in purchases, on average, rather than being assumed to represent fully informed purchases.

The adoption of the EPH has two important implications for the way the data get analyzed. First, observations involving low 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 that is reported in the database (e.g., amount, city, quarter, year).[10] 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 method and a non-EPH method.

[10] For detailed information regarding the specification of the price and purity model see the accompanying report, *Technical Report for the Price and Purity of Illicit Drugs Through 2003*.

Another, related change is that the price of the transaction, not the price per pure gram, is the dependent variable in the statistical regression models. This is an improvement for a technical, statistical reason. The old methods included amount and purity in the denominator of the dependent variable, which biases estimates of the coefficients on the amount and purity variables on the right-hand side of the equation and hence leads ultimately to biased price estimates. Predictions of the price of one expected pure gram can still be generated with this new approach simply by multiplying the coefficient estimates of the regression model by the value of the corresponding independent variable for each transaction, with purity set to 100 percent and amount set to 1 gram. However, because the current models are estimated for different quantity levels, amount is not actually set to 1 gram. Instead, the predicted price of an amount given by the midpoint for each range evaluated at 100 percent purity is calculated and then scaled up or down by a factor of proportionality to generate the equivalent price per expected pure gram.

A final methodological change is the use of simpler weights when generating the national price and purity indices as a weighted average of the various city-specific series. Various weighting schemes can be used. Each has advantages and disadvantages, and the relative merits of each approach depend in part on the purpose for which the national price or purity series will be used. Past reports sought to weight city-specific prices by a proxy for the quantity of the drug consumed in each city, where that proxy was based on drug-specific emergency department episodes recorded by the Drug Abuse Warning Network (DAWN). That approach is not unreasonable. However, the current report adopts the simpler and more transparent approach of weighting city-specific estimates by the relative size of the city as indicated by its population. This approach provides a national series that might be interpreted as the national average price seen by potential users (who reside in cities with enough data to estimate city-specific prices). The previous method attempted to estimate the national average price paid by current users (who reside in cities with enough data to estimate city-specific prices and DAWN rates). Neither average is intrinsically of greater interest than the other. We prefer the former because it can incorporate price data from any city, not just cities for which DAWN estimates can be created, and because population estimates are reliably and universally understood.[11]

A number of minor technical adjustments have also been made; these adjustments are explained in detail in the accompanying technical report. The purpose of all of the revisions is to improve the scientific methodology employed so that more-accurate information can be obtained from the STRIDE data. Again, it is important to stress, that *given these revisions, it is not appropriate to compare the level of price and purity estimates in this report with those presented in earlier reports, and users of this report are strongly advised not to draw inferences from such comparisons. Comparisons of levels within this report, however—e.g., between different cities or years—are, of course, appropriate.*

General Comments About the Presentation of the Results

All prices in this report represent the standardized real price per one expected pure gram (or in the case of marijuana, one bulk gram), adjusted for inflation and expressed in 2002 dollars. As in previous reports in this series, results are presented in a series of graphs and tables. The

[11] Just to emphasize that there are many different reasonable definitions of a national price series, a third alternative would be constructing a series weighted by the number of teenagers in each city. That might be interpreted as the national average price seen by people vulnerable to initiating use of that drug (and who reside in cities with price data).

statistical models underlying these graphs and tables are described in the accompanying technical report. Annual figures and tables of the price per expected pure gram and purity of powder cocaine, crack cocaine, heroin, and methamphetamine observed in purchase transactions are presented, as are annual figures and tables of the price per bulk gram of marijuana. (STRIDE does not include information on potency of marijuana in its observations.) Tables showing the predicted quarterly prices are presented in the Appendix to this report. All tables and figures are based on drug acquisitions within the 50 states and the District of Columbia, and only purchases and purchase attempts are used to estimate price and purity trends in Section 2.

For each drug, the figures and tables present prices and purities at several quantity levels. It is clear from this and other research that prices vary dramatically across these levels; enforcement pressure creates substantial price markups at each stage of the distribution chain between source countries and consumers in the United States. Likewise, for some drugs in some times and places, it is not uncommon for purity to be diluted as the drugs move down the distribution chain. However, as mentioned earlier, definitions of quantity level boundaries and names are somewhat arbitrary. For example, DEA prefers to call the lowest market level the *street* level, reserving the term *retail level* for the next higher level of distribution (e.g., 2.0 to 10.0 grams, for cocaine). Academics prefer using the term *retail* for transactions between a seller and a user, not between two dealers. The data in STRIDE contain information on transactions that range substantially in size, from very small (e.g., 0.1 grams) to very large (multiple kilograms). Illicit drug transactions do not all occur in specific, round lot sizes.[12] Hence, dividing this continuum into discrete levels to represent specific markets is inevitably somewhat arbitrary. To facilitate exposition, we use the terms Q1, Q2, Q3, and Q4 in this report to refer to incrementally higher quantity levels based on the amount of drug involved in the transaction. Readers are advised not to read too much into the labels, as they merely indicate natural breaks in the data and are not intended to convey any scientific meaning.

Because ONDCP is also interested in knowing how purity varies as drugs move down the distribution chain, Section 3 describes the actual purity of powder cocaine, heroin, and d-methamphetamine in transactions reported in all the U.S. seizure and purchase acquisitions available in STRIDE. The inclusion of the seizure data dramatically increases the number of observations, particularly of large drug acquisitions, and, for some drugs, provides greater insights into purity differences across quantity levels.

[12] The accompanying technical report provides information plots of the quantities transacted for particular substances in Appendix A.