International Journal of Drug Policy 73 (2019) 293–300

Contents lists available at ScienceDirect

International Journal of Drug Policy

journal homepage: www.elsevier.com/locate/drugpo

Research Paper

Triangulating web & general population surveys: Do results match legal cannabis market sales?

Jonathan P. Caulkins[a,*], Steve Davenport[b], Anhvinh Doanvo[a], Kyle Furlong[a], Aatir Siddique[a], Michael Turner[a], Beau Kilmer[b]

[a] Carngie Mellon University, Heinz College, 5000 Forbes Ave., Pittsburgh, PA 15213, USA
[b] RAND Drug Policy Research Center, 1776 Main St., Santa Monica, CA 90401, USA

ARTICLE INFO

Keywords:
Cannabis
Web surveys
General population surveys
Markets
Legalization
Drug policy

ABSTRACT

Background: This paper combines complementary attributes of web and general population surveys to estimate cannabis consumption and spending in Washington State. It compares those estimates to legal sales recorded by the state's seed-to-sale tracking system, and thus exploits a rare opportunity to contrast two independent estimates for the same cannabis market. This sheds light on the question of whether nontrivial amounts of black market sales continue even after a state allows licensed production and sale.

Methods: Prevalence of past-month use is estimated from the 2015/16 U.S. National Survey on Drug Use and Health, adjusted for under-reporting. Estimates of consumption and spending per user broken down by age, gender, and frequency of use are developed from RAND's 2013 survey of cannabis users in Washington State. Supply side estimates come from the Washington State Liquor and Cannabis Board's seed-to-sale tracking system. They are expressed in terms of spending, equivalent-weight of flowers, and THC, with THC for edibles imputed using a machine learning technique called random forests.

Results: For the period July 1, 2016 to June 30, 2017, Washington's seed-to-sale data record sales from licensed cannabis stores of $1.17B and across all products an amount of THC that is equivalent to roughly 120–150 MT of flower. Survey responses suggest that amounts spent and quantities consumed are larger than that, perhaps on the order of $1.66B and over 200 MT, respectively.

Conclusion: A perfect match is not expected because of sales to tourists, residual black market activity, production for medical purposes, and diversion across state lines. Nonetheless, the results suggest that three years after state-licensed stores opened, there remained considerable consumption of cannabis supplied outside of the licensed system.

Introduction

Reuter and Greenfield (2001) observed that estimating the size of drug markets is difficult but rewarding because a sense of scale can be informative. There are multiple metrics of market size. Perhaps the most common is prevalence, meaning the number of people who consumed a drug within the last month or year, but prevalence has limitations because there is such heterogeneity in intensity of use (Burns, Caulkins, Everingham, & Kilmer, 2013).

Daily and near daily (DND) users consume much more, perhaps three times more, grams per day of use as do infrequent users (van Laar, Frijns, Trautmann, & Lombi, 2013; Caulkins & Kilmer, 2013, 2013b; also data below). Therefore, one DND user can buy and use at a rate per month that is on the order of 75 times greater than a once a month user. That would not be a problem with monitoring markets via prevalence if the numbers of light and heavy users went up and down in the same proportions. However, that is not always the case. Everingham, Peter Rydell, and Caulkins (1995) showed that the relative numbers of light and heavy cocaine users shifted drastically in the U.S. during the 1980s. Likewise, the proportion of current (past-month) U.S. cannabis users who report daily or near-daily use increased from 11% in 1992 to 36% in 2013 (Caulkins, Kilmer, & Kleiman, 2016). Hence, user counts can vary over time in ways that differ sharply from variation in market demand, spending, and consumption.

Measures that factor in intensity of use can be more informative, so in this paper, we estimated quantity consumed and amount spent in a particularly interesting market, the state-legal but federally-banned market for cannabis in Washington State (Pardo, 2014). We did this by

* Corresponding author.
E-mail address: caulkins@andrew.cmu.edu (J.P. Caulkins).

https://doi.org/10.1016/j.drugpo.2019.06.010

drawing upon the distinctive attributes of general population surveys (GPS) to measure prevalence and web-based surveys to measure intensity of use.

Because the market is state-legal, we were able to contrast those estimates with estimates produced from Washington's seed-to-sale tracking system. That contrast sheds light not only on total use, but also on how much appears to have been supplied outside of legal channels. That, in turn, has implications for how much enforcement is necessary to suppress supply operating outside of a legal system.

Thus, we produced market estimates using two strategies common in the literature. The first, sometimes referred to as "demand-side estimates", starts with estimates of numbers of users (e.g., Wilkins, Bhatta, & Casswell, 2002; Legleye et al., 2008; Kilmer et al., 2013; van Laar et al., 2013; Giommoni, 2014). The second starts with measures pertaining to supply, such as acreage under cultivation (e.g., Bouchard, 2008; Kalacska & Bouchard, 2011; van der Giessen, van Ooyen-Houben, & Moolenaar, 2016). Since the supply-side estimates were so strong, being rooted in Washington's regulatory tracking system, contrasting them shed some light on the effectiveness of efforts to combine GPS and web-based surveys.

There is now a growing and vibrant literature on Washington's experience with state-legalization. Much focuses on outcomes for adolescents (e.g., Mason, Hanson, Fleming, Ringle, & Haggerty, 2015; Fleming, Guttmannova, Cambron, Rhew, & Oesterle, 2016; Cerdá et al., 2017), policy implications for other jurisdictions (Cambron, Guttmannova, & Fleming, 2017; Carnevale, Kagan, Murphy, & Esrick, 2017), or other general descriptions of the market (e.g., Jensen & Roussell, 2016; Pacula, Jacobson, & Maksabedian, 2016; Subbaraman & Kerr, 2016; Roffman, 2016; Dilley, Hitchcock, McGroder, Greto, & Richardson, 2017). No prior paper has sought to compare the sizes of the legal and residual black markets in Washington State, although Smart, Caulkins, Kilmer, Davenport, and Midgette (2017) and Caulkins et al. (2018) did use the seed-to-sale data to look at other aspects of the market. Smart et al. (2017) described state-legal sales of the various types of cannabis, e.g., they noted the rapid expansion in sales of extract-based products, documented the high potency of flower, and described the sharp declines in the (tax-inclusive) retail price per gram of flower. Caulkins et al. (2018) extended this work to describe prices and potencies of various extract-based products, reported prices at the processor level, and noted that the ratio of retail to processor level prices was curiously stable at around 3:1.

While the primary substantive contribution is estimating the size of the legal and extra-legal cannabis markets in Washington State in the period following state-legalization, the primary methodological innovation is combining data from GPS and web-based surveys to estimate drug market size.

GPS provide representative estimates, but often only of prevalence or days of use. Also, a relatively small number of high-frequency users account for a large share of all consumption (Caulkins, 2016; Kleiman, 2001), and GPS interview few heavy users because heavy use is not common. Fewer than one-in-five Washingtonians reported past-year cannabis use in the 2016 NSDUH, and only about one-in-five past-year users consumed daily or near daily, so only about one in 25 members of Washington's general population are high-frequency users. Web-surveys have various strengths and limitations (Fricker & Schonlau, 2002; Van Gelder, Bretveld, & Roeleveld, 2010). Notably for present purposes, a focused web survey can ask more detailed questions about quantities and types of products and attract a disproportionate number of heavy users. However, since respondents self-select into these surveys, they are not representative. Therefore, combining information from the two types of surveys may leverage their complementary strengths.[1]

[1] One referee noted that this process might instead combine the complementary *weaknesses* of the two data sources.

Data

General population survey data

The National Survey on Drug Use and Health (NSDUH) monitors drug use and related behavior through annual surveys. Respondents are selected through a national probability sample of the civilian, non-institutionalized population, 12 years of age and older, living in the 50 states or the District of Columbia. Youths, Blacks, and Hispanics are over-sampled to improve the precision of estimates for these subpopulations. Households (or more precisely "dwelling units") are randomly selected by SAMHSA to receive screening questions to participate in the NSDUH. Then, individuals are randomly selected to participate in the data-collection interview from households that submit responses to the screening questions which determine their eligibility. In 2015–2016, 3778 out of 4844 Washington households (77.9%) successfully submitted eligible screening responses. From these households, the weighted interview response rate, which is equivalent to the AAPOR standard definition *RR2,* was 68.2% (1878 out of 2668 individuals).

Data on age and gender of past month users were accessed using the Restricted-use Data Analysis System (R-DAS), which provided Washington State-specific estimates for 2015 and 2016 combined. R-DAS cannot be queried for a single state's distribution of frequencies of past-month use, so those relative frequencies were obtained at the national level using the Public-use Data Analysis System (P-DAS) broken down by age and gender.

The survey figures were adjusted for under-reporting as in Kilmer et al. (2013).

RAND's 2013 web survey

In 2013, the RAND Corporation conducted a web survey of cannabis consumption among a convenience sample of users in Washington named the Cannabis Consumption Survey (CCS) (Kilmer et al., 2013). The survey was widely publicized. Over the ten days it was available from June 24th to July 3rd, it attracted 3488 respondents, of whom 2783 were from Washington State. 1687 of them reported using in the past-month, including 1216 who answered enough of the detailed questions on numbers of use-days and past-month consumption to be included in analysis of grams consumed per day of use. The survey also asked other questions, including questions about gender, age, race, education, and employment status. Kilmer et al. (2013) Appendix A provides details about the CCS.

It is important to note that we did not use web survey data to estimate prevalence or any other market aggregate; we only used the web survey to estimate quantities consumed per day of use conditional on frequency of use and demographic variables. As with any opt-in survey, there is potential for self-selection bias, although quantities consumed per day conditional on frequency of use may be less vulnerable to those concerns than would be, say, the distribution of frequencies of use.

The survey collected data on the quantity of cannabis consumed and how it was obtained. To improve accuracy, respondents were asked multiple questions about consumption, including one with a visual prompt of a picture of some cannabis next to a ruler, coin, and credit card (the amount of cannabis displayed was randomly assigned). In addition, respondents categorized whether their consumption for each day in the past week was typical, light, or heavy compared to their normal daily consumption.

The survey was conducted before the first state-licensed adult use stores opened (July 2014), but after personal possession and use was legalized (December 2012). Furthermore, even before 2012, Washington State had a robust collection of (unregulated) bricks-and-mortar "medical dispensaries", and medical recommendations were easy to obtain. Prices and product variety have continued to evolve even after stores opened, so of course it would be preferable to have more recent data, but these are the latest of which we were aware that asked detailed questions about quantities consumed per day.

Washington State Liquor and Cannabis Board data

The Washington State Liquor and Cannabis Board (WSLCB) oversees the legal cannabis industry. WSLCB mandates that all transactions and conversions of cannabis products, "from seed to sale" be reported in a database (Miller, 2017; Smart et al., 2017). The variables of primary interest here were those that recorded the amount paid, weight, and potency.

We reported spending here inclusive of state excise taxes but not local or state sales taxes to maintain consistency with past analyses (e.g., Smart et al., 2017; Caulkins et al., 2018). Following Smart et al. (2017) and Washington regulations, total THC content was calculated as THC (decarboxylated) plus 0.877 times THC-A (carboxylated) to account for changes in mass during decarboxylation.

The data field 'inventorytpe' distinguishes ten types of retail products. The two most common are "usable marijuana" (Washington State law's term for dried flower) and "marijuana extract for inhalation" (extracts) which includes products such as wax, kief, shatter, oils and distillates for portable vaporizers (Caulkins et al., 2018).

The seed-to-sale database is a distributed system; businesses are responsible for inputting their own data. There is little to prevent a business from entering bad data, so we dropped some observations to remove anomalous values. In particular, when totaling amounts spent, we used all 38,767,747 observations for usable marijuana and 9,053,623 for extracts for inhalation. For calculations involving weight and potency we omitted observations that had:

1 NULL values for Total THC (2,634,459 for flower & 10,937 for extracts)
2 Potency above 100.0% (980 for flower & 1597 for extracts)
3 Weights larger than what is legally permissible for retail sale (29,104 for flower & 10,472 for extracts).

The total number of observations dropped (2,664,543 for flower and 23,006 for extracts) was slightly less than the sum of the numbers that would be dropped at each of these three steps because some observations violated more than one of the three restrictions. Overall, these filters reduced the number of observations by 6.9%–36,110,437 for usable marijuana and by 0.3%–9,030,627 for extracts for inhalation.

Quantity and potency were reported in straightforward ways for usable marijuana and extracts, but not for solid or liquid edibles, which account for about 9% of spending on cannabis products. Therefore, we used text-analytic algorithms to impute THC content of edibles. We then converted THC content into an "equivalent weight" of usable marijuana by dividing by an assumed rate of THC extracted per gram of usable marijuana. All other product categories together accounted for only 4% of spending, so their equivalent weight was proxied more simply based on spending.

Demand side estimates

Numbers of past-month users

SAMHDA's R-DAS tool for the 2015 and 2016 NSDUH survey stratifies Washington State's past-month cannabis users across age and gender as follows.

Disaggregating by frequency of use would be desirable because there can be a positive correlation between consumption per day and number of days of use in the past month (Zeisser et al., 2011; Caulkins & Kilmer, 2013, 2013b). Unfortunately the R-DAS tool does not allow this for an individual state, so we were forced to assume that the relative (but not absolute) frequencies of past-month use within each of the bins in Table 1 match the corresponding distributions for the same demographic groups as reported in the 2015 NSDUH survey for the entire country.

This simplification was supported by two pieces of circumstantial evidence but cannot be definitively supported with any data of which we were aware. First, Washington State's ratio of past-month to past-year users recorded in the 2015–2016 NSDUH state estimation (12.0% / 18.9% = 0.63) is the same as the ratio for the country (8.6% / 13.7% = 0.63).[2] Second, Lapham et al. (2017) obtained 22,095 responses to a brief cannabis use survey administered in primary care settings in Washington State between March 2015 and February 2016. Within that sample, 20.3% of past-year users reported using daily or almost daily. That matches the 21.4% of past-year users in the 2016 NSDUH who reported 21 or more days of use in the past-month (a common definition of daily or near daily use). Lapham et al.'s other categories (less than monthly, monthly, and weekly) do not map neatly into ranges of past-month days of use, but seem broadly consistent. For instance, 50% of Lapham et al.'s respondents use "less than monthly" and 51% of NSDUH's past-year users report using 0, 1, or 2 times in the past month.

Estimating intensity of use by subgroup

RAND's 2013 web survey was conducted after Washington State had legalized personal possession and use, but before stores opened, and at that time, most reported cannabis consumption took the traditional form of flowers. Average grams of flower consumed per day of use was calculated for each respondent to RAND's 2013 web survey similarly to Kilmer et al. (2013), drawing primarily on two questions asked for each of the past seven days:

1 *Did you use [cannabis]?* Response options: Yes, No, Don't Know
2 *Was the amount you used for that day typical for you?* Response options: Light Day, Typical Day, Heavy Day, Don't Know

Later questions asked respondents to specify a quantity in grams, ounces, or another unit of consumption for what constitutes a Light, Typical, or Heavy day of consumption. Over 70% answered in grams or ounces. Responses in other units, such as joints or bowls, were converted to grams using other questions. For instance, users who prefer reporting consumption in joints were asked (1) how many joints they smoke on a light, typical or heavy day and (2) how many grams they put in their usual joint using pictures of various size piles of marijuana alongside everyday items (a coin, ruler, and credit card.) to calibrate answers. The top 1% of respondents' answers were extremely large, so they were top-coded, making the maximum weight consumed per day of use 8.86 g.

We computed past-week consumption by multiplying past-week days of use by grams per day of use for each type of day (light, typical, or heavy), and adding. Dividing by the number of past-week days of use yielded the respondent's average consumption per day of use. Appendix A reports the results by age, gender, and frequency of use.

We improved these estimates by controlling for differences between web survey respondents and the general population. For example, Fig. 1 shows that greater levels of educational attainment tend to be associated with less being used per day of use. To the extent that web survey respondents tend to be more educated (e.g., because only people with internet access could participate), then controlling for education adjusts these estimates upward.

We used the following regression to predict grams per day of use among past-month users:

[2] https://www.samhsa.gov/samhsa-data-outcomes-quality/major-data-collections/state-reports-NSDUH-2016.

Table 1
Past-Month Prevalence. Based on the 2015–2016 NSDUH GPS Surveys.

		Point Estimate	95% Confidence Interval		General Population
			Lower Limit	Upper Limit	
Total	–	755,000	609,960	900,040	6,029,000
12-17 Years	Male	22,000	12,200	31,800	272,000
12-17 Years	Female	19,000	9,200	28,800	260,000
18-25 Years	Male	82,000	54,560	109,440	379,000
18-25 Years	Female	73,000	53,400	92,600	367,000
26-34 Years	All genders	191,000	124,360	257,640	923,000
35+ Years	All genders	369,000	243,560	494,440	3,828,000

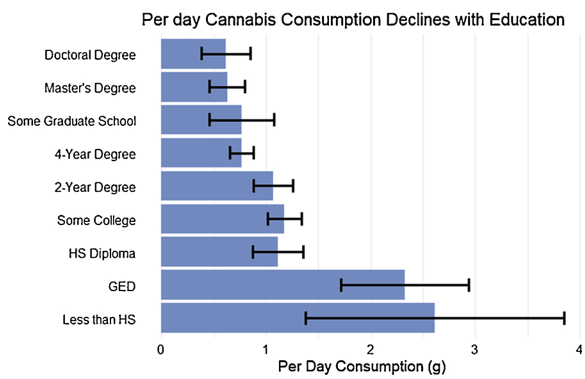

Fig. 1. Average quantity consumed per day of use broken down by level of education.

$$Grams\ per\ use\ day_i\ =\ \beta_1 FreqGenderAge_i + \beta_2 Edu_i\ +\ \beta_3 Emp_i + \beta_4 Race_i\ + \varepsilon_i$$

where *FreqGenderAge* stands for a set of forty dummy variables, one for each combination of the interaction between frequency of use, gender and age, and where Edu_i, Emp_i, and $Race_i$ represent control variables for level of education, employment status, and race, respectively. This model, whose estimated coefficients are given in Appendix A, replicated differences in the unadjusted data concerning consumption intensity between those without a high school diploma and those with post-secondary education and found these differences to be statistically significant. White and Asian respondents consumed on average 0.1-0.3 g per day less than other major racial groups; however, these differences were not statistically significant at the 0.05 level. The only positive racial coefficient ($\beta = +0.32$) from our model was for American-Indian/Alaskan Natives. The model also indicated that retirees consume approximately 0.3 g per day *more* than do full-time workers; this may reflect medical use of cannabis among retirees, as this coefficient ceased to be statistically significant after dropping medical users.

With this regression model, we adjusted estimates for average daily consumption for each group defined by frequency of last month's use (1–3 days, 4–9 days, 10–19 days, or 20+ days), gender, and age while holding level of education, employment status, and race constant at rates matching last month cannabis users in Washington according to NSDUH data. As Fig. 2 indicates, even after controlling for these demographic variables, high-frequency users continue to consume considerably more per day than do less frequent users. (Appendix A gives the specific values and associated confidence intervals.)

Daily and near-daily users reported consuming 1.3–2.0 grams per day of use, depending on age and gender, whereas low frequency male users consumed perhaps slightly more than 0.5 g per day, and low-frequency female users considerably less than that. Except for the oldest age category (age 55+), women reported consuming less per day than men do. For daily and near-daily users the gap was modest, with women consuming 70–80% as much per day as men, but the gap was larger among less frequent users.

Direct estimate of statewide consumption

We multiplied the number of past-month users in each group by the corresponding average grams consumed per month to generate an estimate of monthly consumption. We then multiplied this monthly consumption estimate by 12 for an annual estimate of 200 metric tons (MT).[3] This should be understood to be the quantity consumed if, as at the time of the 2013 RAND web survey, most consumption was in the form of cannabis flower.

Illicit activity is often prone to under-reporting. Kilmer et al. (2013) increased their estimate of the size of Washington's cannabis market by a factor of 1.22 (with a range between 1.02 and 1.43) to account for under-reporting. It is possible that underreporting has diminished since then, but 2013 was already after possession and use had been legalized, so we retained the same factor, in part to facilitate comparisons. We multiplied the estimate above by 1.22 for a revised point estimate of 244 MT.

There is considerable uncertainty around this point estimate. Each subpopulation size estimated by R-DAS has an associated confidence interval, as do the estimates for grams consumed per day of use and the range of possible adjustments for under-reporting. We ran a Monte Carlo simulation (using 100,000 trials) varying these parameters (c.f., van der Giessen et al., 2016). That parametric uncertainty generated a 90% interval estimate of 180–289 metric tons with a median of 231 metric tons.

The RAND survey obtained many responses from individuals who reported consuming primarily home-grown cannabis, and they typically consume significantly more per day than do others. We re-ran the analysis without those respondents, and the results were about 10–15% lower: 180 MT before adjusting for underreporting and 220 MT after. The corresponding 90% interval estimate from Monte Carlo simulation was 160–255 MT with a median of 205 metric tons.

The 2013 RAND CCS did not collect spending data at the same level of granularity, so we did not estimate the dollar value of the market in this way.

Raking to reweight web survey respondents to match GPS prevalence data

As an alternative to averaging web survey respondents' answers by population group, we used individuals' responses directly by employing a raking algorithm to find weights for those individuals that make the web survey's (weighted) group proportions mirror those in the GPS data (Izrael, Hoaglin, & Battaglia, 2004).[4] As Table 2 shows, web survey respondents were older, more likely to be male and much more likely to

[3] RAND survey data were collected in just one month, so if there is seasonality in consumption that might bias this calculation.

[4] Only a subset of the 3,488 web survey respondents reported consumption and spending. After recoding the top 1% of consumers to the 99th percentile, we had 723 usable responses for spending and 1,111 for consumption. Analysis restricted to those who did not grow most of what they consumed reduced these numbers to 964 for consumption and 692 for spending.

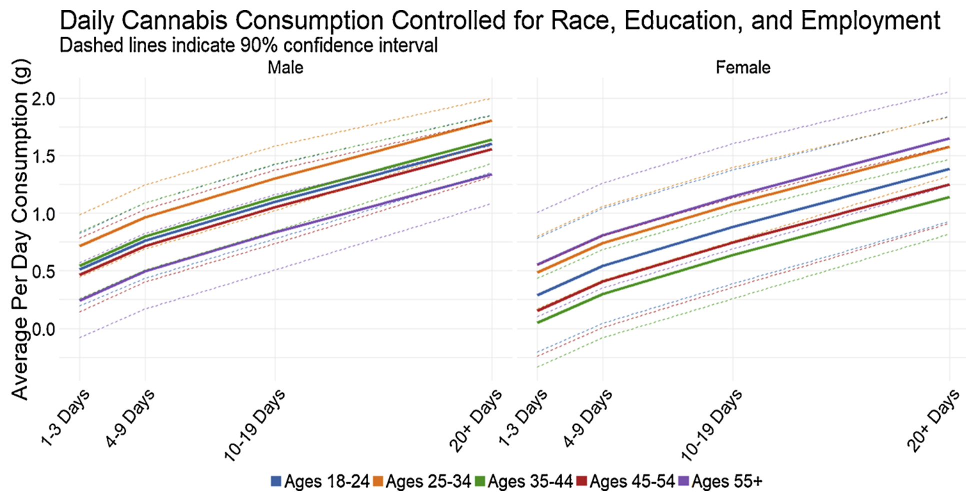

Fig. 2. Average quantities consumed per day of use after controlling for covariates.

Table 2
Subgroup Population Proportions in the Web and GPS Surveys.

		Web Survey	GPS	Difference
Age	18-24 Years	11.9%	26.0%	−14.1%
	25-34 Years	31.1%	25.3%	5.8%
	35+ Years	57.0%	48.9%	8.1%
Gender	Male	77.1%	52.4%	24.7%
	Female	22.9%	47.6%	−24.8%
Frequency	1-3 Days	3.9%	28.1%	−24.2%
	4-9 Days	7.65%	17.8%	−10.2%
	10-19 Days	9.6%	12.9%	−3.2%
	20+ Days	78.9%	41.1%	37.7%

be daily or near-daily users than were typical past-month users in Washington according to GPS data. The algorithm iteratively adjusted a vector of weights until the web survey proportions matched those in the GPS data first on age, then gender, then also on frequency of use.

For past-month users, the *unweighted* monthly averages from the web survey were $214 for spending and 37 g for consumption. These unweighted values were too high because the web survey over-sampled heavy users. After raking, the average monthly spending and quantity consumed fell by roughly one-third to $152 and 21.8 g, respectively. The corresponding estimate of consumption per day of use fell from 1.55 g to 1.06 g.

We re-ran the analysis after removing people who reported growing a majority of the cannabis they used, which reduced the averages (after raking) to $150 spent per month and 18.8 g consumed per month, with the average daily quantity consumed falling to 0.92 g.

Dividing our estimate of $150 spent per month by 18.8 g consumed per month implies spending an average of about $8 per gram consumed. That is below what Smart et al. (2017) report for the average price of cannabis in state-licensed stores from October 2014 – September 2016, as may be expected because (1) prices have been falling, (2) we did not remove people who grow some but not the majority of their cannabis and (3) black-market products might sell for less than the taxed price in state-licensed stores.

Multiplying $150 spent per month and 18.8 g consumed per month by 12 months per year, 755,000 past-months users recorded by the GPS, and 1.22 to adjust for under-reporting produced statewide market estimates of $1.66 billion and 208 metric tons.[5] Retaining the own-growers did not change the spending estimate, but increased the consumption estimate to 241 MT.

[5] $150 x 12 x 755,000 x 1.22 = $1.66billion and 18.8g x 12 x 755,000 x 1.22 = 208MT.

Supply-side estimates

We produced supply-side estimates from Washington's seed-to-sale database for July 1, 2016 through June 30, 2017. This period covers one full calendar year, commences after unlicensed medical dispensaries were closed, and terminates before data collection in Washington State was contracted to a new vendor. It is slightly later than the period covered by the demand side estimates.

Our analysis focused on four classes of products that represent 93.98% of observations; usable marijuana (68.59%), marijuana extract for inhalation (16.02%), solid marijuana infused edible (7.53%), and liquid marijuana infused edibles (1.83%).

Producing spending totals is a straightforward summing of transactions' recorded price, leaving negative price observations in the calculation since those represent product returns.

Estimating product weight and THC content is generally straightforward for usable marijuana products and inhalation extracts as these observations have a field "usable weight" that appears to directly reflect the product's weight (exclusive of packaging), and most products were produced from a single batch of "parent" flower whose cannabinoid content was established via a laboratory test. Their THC and CBD levels were inherited from their parent flower lots.

For usable marijuana, 6.38% of observations had two or more parent plants. It was not clear how to combine the potencies of multiple parents, so we computed THC totals for the 93.62% of observations that had a single parent, and scaled those totals up proportionally to account for the items with multiple parents. This assumes that products with multiple parents are not systematically more or less potent than are those with a single parent. (This was not an important issue with other product classes. Only 0.00003% of all marijuana extracts for inhalation, 0.02% of solid edible, and 0.022% of liquid edible observations had multiple parents.)

Parallel calculations cannot be performed for edible products because the database does not list usable weights for them. Also, for a small but significant proportion of edible products, particularly liquid edibles, the database does not seem to provide reliable potency values. We identified quantities of primary chemicals (meaning THC for THC products and CBD for CBD products) for 45% of the edibles observations via text analysis of their product names. For example, if the product's name was "EH MILK MINIS 100mg" and its THC potency exceeded its CBD potency, we inferred that its THC content was 100 mg. Washington State law limits THC in edibles to 100 mg per package; the

Table 3a
Cross-Validated Mean Error per Observation of Primary Chemical by Product (Numbers in parenthesis are the standard error of the models' estimated mean error).

Model	THC content (mg) in THC-based solid edibles	THC content (mg) in THC-based liquid edibles	CBD content (mg) in CBD-based solid edibles	CBD content (mg) in CBD-based liquid edibles
Sample means-based estimation	39.9	38.6	69.9	69.9
	(0.129)	*(0.382)*	*(6.27)*	*(0.946)*
Price-based Linear regression	19.4	21.1	62.9	39.2
	(0.123)	*(0.188)*	*(6.40)*	*(0.437)*
Random forest exploiting product names	4.48	5.11	18.8	7.53
	(0.341)	*(0.277)*	*(3.23)*	*(0.573)*

Table 3b
Cross-Validated Mean Errors per Observation of Usable Weight by Product (Numbers in parenthesis represent the standard error of the models' estimated mean error).

Model	Usable Weight (mg) in THC-based solid edibles	Usable Weight (mg) in THC-based liquid edibles	Usable Weight (mg) in CBD-based solid edibles	Usable Weight (mg) in CBD-based liquid edibles
Sample means-based estimation	683.1	550.4	911.5	391.5
	(3.572)	*(15.32)*	*(30.24)*	*(35.07)*
Price-based linear regression	496.1	484.6	786.0	382.2
	(5.821)	*(6.217)*	*(26.13)*	*(14.90)*
Random Forest exploiting product names	76.59	74.09	167.4	67.6
	(4.910)	*(6.713)*	*(27.40)*	*(9.400)*

0.5% of observations whose names suggested contents exceeding 100 mg of HTC were dropped.

For the remaining edibles observations, we imputed cannabinoid content based on the relationship observed between cannabinoid content and predictor variables in the seed-to-sale database among the observations whose cannabinoid content could be inferred from the product name. This was done separately for solids and liquids and separately for THC- vs. CBD products, where an item was deemed to be a CBD product if its CBD potency exceeded its THC potency. Approximately 11.5% of edibles were CBD products by this definition, a much greater proportions than for usable marijuana (1%) or inhalation products (2%). THC content for CBD-dominant products was calculated assuming that the ratio of THC to CBD content was the same as the reported ratio of THC to CBD potency.

We attempted to do this imputation via regression using the predictor variables price, time, the interaction between the two, and personal income in the store's census tract as a proxy for labor and other business operating costs. We calculated the mean error per observation via five-fold cross validation to evaluate performance in calculating cannabinoid content on the observations that had been labeled based on their product name. The regression model only marginally outperformed sample-means-based estimation, which is to say assuming that every product's value matches the mean value in the training set. (See Tables 3a and 3b.)

What turned out to work much better was a machine-learning algorithm called random forests. Random forests generate many decision trees that recursively partition the data according to its features. Once trained, the model takes a 'plurality vote' among these decision trees to make its predictions. We used as features the frequencies of every different word in the product names field except numbers and measures of weight, along with price and local income. (Details are provided in Appendix C.) As Tables 3a and 3b show, the random forest models substantially reduced estimation errors.

Table 4 summarizes the results of these supply-side estimates, distinguishing direct measurement from the random forest estimation for edibles, as well as some extrapolations to estimate the total weight of usable marijuana (i.e., flowers) or its equivalent that was sold. In particular, the seed-to-sale data suggest that about 7.5 g of flower lot material is used to produce each gram of hydrocarbon wax or $C0_2$ hash oil. The average potency of these products is 70–75% (Caulkins et al., 2018). That might suggest that about ten grams of usable marijuana is processed for each gram of THC appearing in extracts. However, leaves, stems, and other plant material are also used in the extraction process, so not all of the THC in extracts comes from flower. Nonetheless, we estimated the weight of marijuana flower used to produce extracts by dividing total THC content for extracts and edibles by 10%, and term this the Usable Weight Equivalent" (UWE) of flowers. For the remaining minor product types (last column), we simply scaled the weight of usable flower by the ratio of spending on other products relative to spending on usable flower (i.e., 85.32 MT * ($47.1B / $763.9B) = 5.3 MT). Summing across product types produced final market size estimates of $1.17B for total sales and 147.1 MT for the weight of usable marijuana or its equivalent in derivatives or other plant material.

Dividing by 10% is conservative with respect to our overall conclusion that cannabis consumption exceeds sales in licensed stores. If one thought the average potency of all material used to create extracts was 13.3%, not the 20% typical of flowers sold as flowers (Smart et al., 2017), then the previous step would involve dividing by 15%, not 10%, and the total sales would be 128.2 MT, not 147 MT.

Another, simpler approach sidesteps the random forest model estimation and conversion of THC estimates back into flower. In particular, estimating UWE by scaling the amount purchased as usable marijuana (85.32 M T) by the ratio of total spending on cannabis products ($1.17B) to spending on usable marijuana ($763.9 M) produces an estimate of 130.3 MT. If we scale based on both flower and extract purchases, the result is 137.3 MT.

All these estimates are fairly close, and in deference to various uncertainties and assumptions in the analysis, might best be characterized as falling within a range from 120 to 150 MT.

Discussion

The supply-side estimates of sales from Washington State's licensed cannabis stores for July 1, 2016 to June 30, 2017 were $1.17B and 120–150 MT. The demand side estimates were significantly greater.

Multiplying general population survey (GPS) numbers of past-month users separated by age, sex, and frequency of use by the intensities of use for those groups observed in RAND's 2013 web-survey produced a market estimate of 220 MT after omitting most home-growing users. The raking method to re-weight web survey respondents produced market size estimates of $1.66 billion and 208 metric tons.

Table 4
Supply-side Estimates of Washington's State-Legal Cannabis Market, July 1st 2016–June 30th, 2017.

Quantity	Usable Marijuana	Extracts for Inhalation	Solid Edibles	Liquid Edibles	Other	Total
Spending ($M)	$763.9	$264.6	$76.7	$26.2	$47.1	$1,166.7
Product Weight (MT)[a]	85.32	7.58				
THC or Estimated THC (MT)[a]	18.97	5.35	**0.23**	**0.07**		
Price of THC ($ per g)	$40	$49	$340	$366		
THC Content / Product Weight	22.2%	70.6%				
Assumed THC yield per gram of usable MJ used to make extracts		10%	10%	10%		
Equivalent Usable Weight (MTE)		53.5	2.3	0.7	5.3	147.1

Light grey cells were computed directly from the data.
Dark grey cells are estimated via random forest models.
White cells are calculated within the table.
[a]Product weight and THC content for usable marijuana is computed directly for observations with a single parent flower and then scaled by dividing by the proportion of observations that had a single parent.

These figures are greater than Kilmer et al. (2013)'s estimate that 175 MT were consumed in 2013, the year before state-licensed stores opened.

Demand-side estimates of consumption by Washington residents should be larger than sales recorded in state-licensed stores if Washington residents consume cannabis that is purchased from other states' legal markets, purchased on the black market, produced by a medical cooperative, or home-grown.

The demand-side estimates include consumption by those under 21. Stores are not supposed to sell to those under 21 and compliance checks suggest that they rarely do.[6] That said, it is likely that some items purchased from stores by adults are consumed by those who are under 21.

Conversely, seed-to-sale tracking system estimates could tend to be larger because (1) Consumption has been increasing and the supply side estimate comes from a slightly later time period (July 1, 2016 – June 30th 2017 vs. an annual estimate based on the 2015 and 2016 NSDUH surveys), and (2) Some cannabis purchased from licensed stores may be consumed by non-resident tourists or be sent (illegally) beyond Washington's borders (Hansen, Miller, & Weber, 2017).

We do not think error in the supply side estimate can account for much of the gap because most of the market is usable marijuana and extract-based products, and estimation of their totals is straightforward given the seed-to-sale data. We are also skeptical that Washington State residents' state-legal purchases from other states would be large; the price differences would not seem to warrant traveling to Oregon or Colorado to purchase cannabis, particularly given that transporting those goods back across state lines is illegal.

In our judgment, the two most plausible explanations for why demand-side estimates are so much larger are: (1) The demand side estimates err by being substantially too high or (2) Even after state legalization, there remains considerable consumption of illegal cannabis in Washington State.

With respect to errors, these estimates are affected by a small number of respondents who report very high rates of use. There is no way to know whether that accurately reflects the skewed character of consumption or false reports. However, many of those reporting very high rates of consumption wrote in explanations that describe special circumstances, the most common being very heavy use for medical purposes that was supplied by home-growing – and we showed the effects of dropping those responses.

The 2013 RAND web survey was fielded from June 24th to July 3rd, so if consumption is higher than average during that season, that could inflate consumption estimates. Also, average potency may have been higher in 2016 than in 2013, and if users titrate in response to higher potency, the weight consumed per person may have been lower in 2016.

So the demand-side estimate may be wrong. However, it is also possible that all of the point estimates are correct, the spending gap reflects illicit-market purchases, and the additional gap in consumption reflects own-growing. One plausible story that would reconcile the point estimates is that Washingtonians: (1) Bought 130 MT at $9 per gram (average) from state-licensed stores for $1.17B, (2) Bought 70 MT at $7 per gram from the black market for $0.49B, and (3) Grew 30 MT, spending on growing supplies but not the cannabis itself. That would account for $1.66B in total spending on cannabis with only $1.17B going to state-licensed stores, and 230 M T of total consumption with only 130 M T bought at state-licensed stores. Thus, in this scenario, about 30% of spending and 40% of consumption would occur outside the state-licensed system.

None of the estimates are precise enough to view this story as being "right" in its particulars, but we consider the results here as suggestive that Washingtonians may have continued to consume nontrivial amounts of own-grown and black-market purchased cannabis even after state-licensed stores were in full operation. Admittedly, given the limitations of the data and the nature of the problem, this interpretation must be understood to be a judgment, not something that can be known with certainty.

Conclusions and further work

There are reasons why the specific results above may not generalize to other markets and so replication elsewhere may be fruitful. The analysis pertained to a market that was state-legal but prohibited by federal law, and that was still relatively young and with falling prices. Unlike Colorado, Washington State did not have a state-regulated medical marijuana system before state-legalization of production for non-medical use. Hence, even if Washington still had own-growing and black market sales circa 2016, that does not mean that own-growing and black market sales would be as prominent a decade or more after

[6] https://data.lcb.wa.gov/dataset/Marijuana-Compliance-Checks-Data-Lens-/u2ct-nbza.

national legalization, or that they will remain at those levels in Washington State if national legalization does not occur. Replacing black markets with legal sales may be an extended process - not something that happens rapidly upon passage of a law or the opening of the first store. This raises an interesting question of how aggressively law enforcement should chase out the residual black market during such transition times.

There are at least three areas for further work. One is improving the estimates from the seed-to-sale database of the uncommon product forms, such as marijuana mixes. The second is improving estimates of the intensities of use produced by web surveys, or other sources. The third is replicating this work for other markets, such as those in other US states that have state-legal markets and in Canadian provinces.

Perhaps the most basic conclusion of this work is that it is possible to combine web- and GPS surveys, leveraging their complementarities, to produce estimates that seem sensible, and triangulating market sizes simultaneously from both the demand and supply sides can provide an instructive contrast.

Declaration of Competing Interest

None to declare.

Acknowledgments

We gratefully acknowledge Claude Setodji for help with the raking algorithm; Krista Kinnard, Lauren Renaud, and Edson Severnini for generous help on multiple aspects of this work; and two anonymous referees for their many helpful comments. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

Bouchard, M. (2008). Towards a realistic method to estimate cannabis production in industrialized countries. *Contemporary Drug Problems, 35*(2-3), 291–320.

Burns, R. M., Caulkins, J. P., Everingham, S. S., & Kilmer, B. (2013). Statistics on cannabis users skew perceptions of cannabis use. *Frontiers in Psychiatry, 4*, 138.

Cambron, C., Guttmannova, K., & Fleming, C. B. (2017). State and national contexts in evaluating cannabis laws: A case study of Washington state. *Journal of Drug Issues, 47*(1), 74–90.

Carnevale, J. T., Kagan, R., Murphy, P. J., & Esrick, J. (2017). A practical framework for regulating for-profit recreational marijuana in US States: Lessons from Colorado and Washington. *The International Journal of Drug Policy, 42*, 71–85.

Caulkins, J. P. (2016). The real dangers of marijuana. *National Affairs,* (26), 21–34. Available at: http://www.nationalaffairs.com/publications/detail/the-real-dangers-of-marijuana.

Caulkins, J. P., & Kilmer, B.(with M. G.) (2013a). Estimating the size of the EU Cannabis Market. In F. Trautmann, B. Kilmer, & P. Turnbull (Eds.). *Further insights into aspects of the illicit EU drugs Market* (pp. 289–323). Luxembourg: Publications Office of the European Union Part I, Report 4.

Caulkins, J. P., & Kilmer, B.(with M. G.) (2013b). Estimating the size of the EU Cannabis Market. In B. Kilmer, & P. Turnbull (Eds.). *Further insights into aspects of the illicit EU drugs Market / franz trautmann* (pp. 289–323). Luxembourg: Publications Office of the European Union Part I, Report 4.

Caulkins, J. P., Bao, Y., Davenport, S., Fahli, I., Guo, Y., Kinnard, K., ... Kilmer, B. (2018). Big data on a big new market: Insights from Washington State's legal cannabis market. *The International Journal of Drug Policy, 57*, 86–94.

Caulkins, J. P., Kilmer, B., & Kleiman, M. A. R. (2016). *Marijuana legalization: What everyone needs to know* (2nd edition). Oxford University Press.

Cerdá, M., Wall, M., Feng, T., Keyes, K. M., Sarvet, A., Schulenberg, J., ... Hasin, D. S. (2017). Association of state recreational marijuana laws with adolescent marijuana use. *JAMA Pediatrics, 171*(2), 142–149.

Dilley, J. A., Hitchcock, L., McGroder, N., Greto, L. A., & Richardson, S. M. (2017). Community-level policy responses to state marijuana legalization in Washington State. *The International Journal of Drug Policy, 42*, 102–108.

Everingham, S., Peter Rydell, C., & Caulkins, J. P. (1995). Cocaine consumption in the US: Estimating past trends and future scenarios. *Socio-economic Planning Sciences, 29*(No. 4), 305–314.

Fleming, C. B., Guttmannova, K., Cambron, C., Rhew, I. C., & Oesterle, S. (2016). Examination of the divergence in trends for adolescent marijuana use and marijuana-specific risk factors in Washington State. *Journal of Adolescent Health, 59*(3), 269–275.

Fricker, R. D., & Schonlau, M. (2002). Advantages and disadvantages of Internet research surveys: Evidence from the literature. *Field Methods, 14*(4), 347–367.

Giommoni, L. (2014). The retail value of the illicit drug market in Italy: A consumption-based approach. *Global Crime, 15*(1-2), 27–50.

Hansen, B., Miller, K. S., & Weber, C. (2017). *How extensive is inter-state diversion of recreational marijuana?* .

Izrael, D., Hoaglin, D. C., & Battaglia, M. P. (2004). To rake or not to rake is not the question anymore with the enhanced raking macro. *Proceedings of the Twenty-Ninth Annual SAS Users Group International Conference*. Paper 207. Available at https://www.researchgate.net/profile/Michael_Battaglia/publication/237510628_To_Rake_or_Not_To_Rake_Is_Not_the_Question_Anymore_with_the_Enhanced_Raking_Macro/links/53ed0cf80cf2981ada11e468.pdf.

Jensen, E. L., & Roussell, A. (2016). Field observations of the developing legal recreational cannabis economy in Washington State. *The International Journal of Drug Policy, 33*, 96–101.

Kalacska, M., & Bouchard, M. (2011). Using police seizure data and hyperspectral imagery to estimate the size of an outdoor cannabis industry. *Police Practice and Research, 12*(5), 424–434.

Kilmer, B., Caulkins, J. P., Midgette, G., Dahlkemper, L., MacCoun, R. J., & Pacula, R. L. (2013). *Before the grand opening: Measuring washington state's marijuana market in the last year before legalized commercial sales.* Santa Monica, CA: RAND.

Kleiman, M. A. R. (2001). Controlling drug use and crime with testing, sanctions, and treatment. In P. B. Heymann, & W. N. Brownsberger (Eds.). *Drug addiction and drug policy: The struggle to control dependence* (pp. 168–192). Cambridge, MA: Harvard University Press.

Lapham, G. T., Lee, A. K., Caldeiro, R. M., McCarty, D., Browne, K. C., Walker, D. D., & Bradley, K. A. (2017). Frequency of cannabis use among primary care patients in Washington State. *Journal of the American Board of Family Medicine, 30*(6), 795–805.

Legleye, S., Legleye, S., Ben Lakhdar, C., Legleye, S., Ben Lakhdar, C., Spilka, S., ... Spilka, S. (2008). Two ways of estimating the euro vsalue of the illicit market for cannabis in France. *Drug and Alcohol Review, 27*(5), 466–472.

Mason, W. A., Hanson, K., Fleming, C. B., Ringle, J. L., & Haggerty, K. P. (2015). Washington State recreational marijuana legalization: Parent and adolescent perceptions, knowledge, and discussions in a sample of low-income families. *Substance Use & Misuse, 50*(5), 541–545.

Miller B., Making it legal: The tech implications of regulation recreational marijuana, *Government Technology* (March), 2017, Available at http://www.govtech.com/policy/Making-It-Legal-Tech-Implications-of-Regulating-Recreational-Marijuana.html.

Pacula, R. L., Jacobson, M., & Maksabedian, E. J. (2016). In the weeds: A baseline view of cannabis use among legalizing states and their neighbours. *Addiction, 111*(6), 973–980.

Pardo, B. (2014). Cannabis policy reforms in the Americas: A comparative analysis of Colorado, Washington, and Uruguay. *The International Journal of Drug Policy, 25*(4), 727–735.

Reuter, P., & Greenfield, V. (2001). Measuring global drug markets. *World economics, 2*(4), 159–173.

Roffman, R. (2016). Legalization of cannabis in Washington State: How is it going? *Addiction, 111*(7), 1139–1140.

Smart, R., Caulkins, J. P., Kilmer, B., Davenport, S., & Midgette, G. (2017). Variation in cannabis potency and prices in a newly-legal market: Evidence from 30 million cannabis sales in Washington State. *Addiction, 112*(12), 2167–2177.

Subbaraman, M. S., & Kerr, W. C. (2016). Marijuana policy opinions in Washington State since legalization: would voters vote the same way? *Contemporary Drug Problems, 43*(4), 369–380.

van der Giessen, M., van Ooyen-Houben, M. M., & Moolenaar, D. E. (2016). Estimating the production, consumption and export of cannabis: The Dutch case. *The International Journal of Drug Policy, 31*, 104–112.

Van Gelder, M. M. H. J., Bretveld, R. W., & Roeleveld, N. (2010). Web-based questionnaires: the future in epidemiology? *American Journal of Epidemiology, 172*(11), 1292–1298.

van Laar, M., Frijns, T., Trautmann, F., & Lombi, L. (2013). Sizing the cannabis market: A demand-side and user-specific approach in seven European countries. *Current Drug Abuse Reviews, 6*(2), 152–164.

Wilkins, C., Bhatta, K., & Casswell, S. (2002). A' demand side' estimate of the dollar value of the cannabis black market in New Zealand. *Drug and Alcohol Review, 21*(2), 145–151.

Zeisser, C., Thompson, K., Stockwell, T., Duff, C., Chow, C., Vallance, K., & Lucas, P. (2011). A 'standard joint'? The role of quantity in predicting cannabis-related problems. *Addiction Research & Theory, 20*(1), 82–92.