

An EU-focused analysis of drug supply on the online anonymous marketplace ecosystem*

Nicolas Christin

Carnegie Mellon University
School of Computer Science (ISR) and Engineering & Public Policy
RMCIC, Room 2202
4720 Forbes Ave
Pittsburgh, PA 15213, USA
nicolasc@cmu.edu

This version: January 9, 2017

*This document has been prepared by the author for the European Monitoring Centre for Drugs and Drug Addiction (EMCDDA) under contract CC.16.SAT.0050.1.0. While the author is professionally affiliated with Carnegie Mellon University, this work was performed as an independent consultant. This paper represents the position of the author, at the time of the writing, and not that of Carnegie Mellon University.

Executive summary

Online anonymous marketplaces are a relatively recent technological development that enables sellers and buyers to transact online with far stronger anonymity guarantees than on traditional electronic commerce platforms. This has led certain individuals to engage in transactions of illicit or illegal goods.

This reports presents an analysis of the online anonymous marketplace data collected by Soska and Christin [13] over late 2011–early 2015. In this report, we focus on drug supply coming from the European Union. Keeping in mind the limitations inherent to such data collection, we found that, for the period and the marketplaces considered:

- EU-based suppliers represented a significant share of all drug revenue—approximately 46% of all drug sales.
- EU-originating drugs primarily came from Germany, the Netherlands, and the United Kingdom.
- Cocaine and other stimulants altogether represented a majority of all EU-based drug sales.
- Supply of New Psychoactive Substances (NPS) was heavily concentrated in the United Kingdom, and remained very modest with revenues in the order of EUR 3,000 per day at market peak.
- Marketplace vendors primarily catered in the retail space, but there was evidence of larger (bulk-level) sales. Volume-based discounting tended to occur, albeit at relatively modest levels.
- Half of the vendors specialized in one type of drug; and half of the drug sellers tended to stick to a given weight echelon.

1 Introduction

By using a combination of network-level anonymity technology [8] and pseudonymous payment systems [12], online anonymous marketplaces are a relatively recent technological development that enables sellers and buyers to transact online with far stronger anonymity guarantees than on traditional electronic commerce platforms. Unfortunately, as a by-product of this anonymity, certain individuals have been using this technology to engage in transactions of illicit or illegal goods, as exemplified by most transactions on the well known Silk Road marketplace [6].

Of interest to us in this report, is the extent of narcotic trafficking on online anonymous marketplaces. Previous efforts [6,13] have aimed at characterizing the entire ecosystem of online anonymous marketplaces, and have shown that, thus far, narcotic trafficking is a large share of the entire economy supported by this technology.

In this report, we analyze data previously collected by Soska and Christin [13], spanning several years (late 2011–early 2015) including the “early days” of online anonymous marketplaces (Silk Road, the original “modern” online anonymous marketplace, opened its doors in February 2011). Different from Soska and Christin’s work, which attempted a general characterization of the entire ecosystem, here, we primarily focus on analyzing drug supply reportedly originating from the European Union. In addition, we perform additional analysis of the relationship between financial revenues and actual quantities (weights, volumes, units) of products being sold.

This report is organized as follows. We next provide background on online anonymous marketplaces in Section 2, discuss our methodology, which shares many of the same assumptions and limitations as the original study by Soska and Christin, in Section 3. We turn to analyzing the collected data in Section 4 before drawing brief conclusions in Section 5.

2 Background on Online Anonymous Marketplaces

The earliest modern online anonymous marketplaces, more commonly known as “darknet marketplaces,” or, more rarely as “cryptomarkets [11],” appeared in early 2011, and were a follow-up to older drug forums which benefited from technological advances to guarantee better anonymity to their visitors. We refer the reader to Christin [6], Martin [11], and Soska and Christin [13], among others, for a thorough description of how the various technologies at play interact with each other.

At a very high level, online anonymous marketplaces consist of websites similar to other electronic commerce businesses such as eBay or the Amazon marketplace. The key difference is that those websites usually run as Tor hidden services [8] or, more rarely, as i2p “eep sites” [3], which allows them to conceal the location of the server on which they run, and force their patrons (sellers and buyers) to use anonymizing technology, thereby (hopefully) evading prosecution and arrest.

From a user interface standpoint, Figure 1 shows, as an example, the AlphaBay marketplace. Information available on AlphaBay is reasonably representative of what can typically be found on an online anonymous marketplace; that is, most other online anonymous marketplaces contain similar information (or a subset of it).

The screenshot shows the AlphaBay Market interface. At the top, there is a navigation bar with links for HOME, SALES, MESSAGES (1), ORDERS, LISTINGS, BALANCE, FEEDBACK, FORUMS, API, and SUPPORT. A user is logged in as 'BTC' with an 'Autoshop' and 'Logout' option. The main content area displays search results for 'Stimulants / Cocaine'. On the left, there is a 'BROWSE CATEGORIES' sidebar with various drug categories and their respective counts. The search results list several items, including 'FERRAGAMO' batch arrived 8TH June is 10/10, 'FE 100%' items, and 'AUGUST PROMOTION' items. Each listing includes a buy price in USD and a quantity left.

(a) Overview

This screenshot shows a detailed view of a specific item listing. The item is titled '★★'FERRAGAMO' batch arrived 8TH June is 10/10 ★★ 1G ★ PURE QUALITY PERUVIAN COCAINE ★ UK TO UK ★ FAST DELIVERY ★★'. The listing includes a vendor profile for 'The Honest Cocaine Company', a 'LATEST NEWS FOR JUNE' section, and a table of features. The purchase price is listed as USD 79.99, and there are 'Buy Now' and 'Queue' buttons. The 'Description' tab is selected, showing the text: 'LATEST NEWS FOR JUNE Product quality has increased on latest batch that arrived on 8th June. All our batches have been around 75% but this batch is SUPREME and will last 1-3 months.'

Product class	Features	Origin country	Features
Physical package	United Kingdom	United Kingdom	United Kingdom
Quantity left	Never	Ships to	Ireland, United Kingdom
Ends in	Never	Payment	Escrow

(b) Item listing

Figure 1: Example of an online anonymous marketplace. This figure shows the AlphaBay online anonymous marketplace. The screen in 1(a) displays various drug listings; 1(b) shows a specific item listing.

Description	Bids	Feedback	Refund Policy
Listing Feedback			
Buyer	Date	Time	Comment
⊕ [blurred]	August 16, 2016	19:25	Excellent as always
⊕ [blurred]	August 16, 2016	14:56	Super fast delivery. Thanks as always
⊕ [blurred]	August 16, 2016	03:04	Excellent again. Next day as promised
⊕ [blurred]	August 15, 2016	23:07	
⊕ [blurred]	August 15, 2016	13:17	Fast delivery and nice product. Great vendor.
⊕ [blurred]	August 15, 2016	12:56	Had a slight mixup but these guys customer service is beyond epic! Never had such comms! Loyal customer!
⊕ [blurred]	August 15, 2016	12:55	Had a slight mixup but these guys customer service is beyond epic! Never had such comms! Loyal customer!
⊕ [blurred]	August 15, 2016	10:56	2DD, excellent as always. Highly recommend, thanks :)
⊕ [blurred]	August 15, 2016	09:03	Excellent, 2DD! Have not tried but smells good. Stealth 3/5. Definitely will use again.
⊕ [blurred]	August 14, 2016	20:26	NDD - very good quality
⊕ [blurred]	August 14, 2016	01:15	ok
⊕ [blurred]	August 11, 2016	13:37	ordered Thurs arrived Tues... bang on weight
⊕ [blurred]	August 11, 2016	09:45	
⊕ [blurred]	August 11, 2016	01:13	thanks

Figure 2: **Example of item feedback.** AlphaBay offers rich feedback information compared to some of its competition.

The main page, in Figure 1(a), displays various categories of items available for sale, as represented on the left-hand menu, including – of interest to this report – “drugs,” that is, primarily illegal narcotics or psychotropic drugs, and prescription drugs, i.e., medicines. As evidenced in the figure, different items may be sold by different vendors. Usually, the marketplace only acts as a broker – similar to eBay – that ensures that transactions complete to the satisfaction of both buyers and sellers.

Specific item listings, as shown in Figure 1(b), contain numerous pieces of information: in the case of AlphaBay, an origin country (here, UK), potential shipping destinations (here, UK and Ireland), a vendor name (blurred), a description of the item, and, crucially, user feedback.

Figure 2 shows more specifically what this feedback contains. On AlphaBay, feedback consists of a timestamp, a short message, a rating (represented here by the green “plus” signs), and a string allowing, to some extent, to differentiate between buyers. This latter field is usually not present on the vast majority of marketplaces – AlphaBay remains an exception.

This feedback is particularly important for our analysis. As described in Christin [6], mandatory feedback (which is the case on the majority of marketplaces) is an excellent proxy for sales. We can indeed obtain an idea of the sales volume, simply by correlating the feedback timestamps with the item prices (and the quantity, when available.) For instance, the product whose feedback is presented in Figure 2 is sold for \$79.99 (see Figure 1(b)); we see six pieces of feedback were deposited on August 15, 2016. We can then infer that the vendor sold $\$79.99 \times 6 = \479.94 worth of the item on that specific day.

Table 1: **Markets crawled.** The table, taken from Soska and Christin [13], describes which markets were crawled, the time the measurements spanned, and the number of snapshots that were taken. * denote market sites seized by the police, † voluntary shutdowns, and ‡ (suspected) fraudulent closures (owners absconding with escrow money).

Marketplace	Measurement dates	# snapshots
Agora	12/28/13–06/12/15	161
Atlantis‡	02/07/13–09/21/13	52
Black Flag‡	10/19/13–10/28/13	9
Black Market Reloaded†	10/11/13–11/29/13	25
Tor Bazaar*	07/02/14–10/15/14	27
Cloud 9*	07/02/14–10/28/14	27
Deep Bay‡	10/19/13–11/29/13	24
Evolution‡	07/02/14–02/16/15	43
Flo Market‡	12/02/13–01/05/14	23
Hydra*	07/01/14–10/28/14	29
The Marketplace†	07/08/14–11/08/14	90
Pandora‡	12/01/13–10/28/14	140
Sheep Marketplace‡	10/19/13–11/29/13	25
Silk Road* ¹	11/22/11–07/24/12	133
	06/18/13–08/18/13	31
Silk Road 2.0*	11/24/13–10/26/14	195
Utopia*	02/06/14–02/10/14	10

3 Collection Methodology and Data

3.1 Data collection

For this report, we exclusively rely on the data collected by Soska and Christin [13]. Soska and Christin built a special-purpose web crawler, using heavily parallelized connections to gather considerable amounts of data in relatively short amounts of time. A full exposition of the technical details can be found in their paper [13]. Table 1 summarizes the data collected. Roughly speaking, the data spans late 2011–early 2015, and represents more than 3 TB of storage. The SQLite database containing the parsed information is approximately 19 GB.

¹The November 2011–July 2012 Silk Road data comes from a previously reported collection effort, with publicly available data [6].

3.2 Data classification and processing

While all listings have been previously parsed and stored into a database, data still needs to be further processed to be amenable to analysis. We in particular need to identify the type of product being sold, the quantities and volumes of product being sold, and the country of origin of the items.

Item categories As discussed in earlier work [13], categories self-reported by sellers, e.g., “Stimulants/Cocaine” in Figure 1, are often incorrect (e.g., we have seen weapons being categorized under “plants”). Instead we determine the type of product by performing automated text analysis of the item description. The process is analogous to that described in earlier work [13]. A key difference, however, is that the categories of interest for the present study differ from those in the earlier work [13]. As such we need to re-run the classifying process, and re-evaluate its accuracy. In this work, we consider the following 22 categories:

1. *Drug categories of primary interest:*

- Cannabis: All forms of cannabis products (resin, herbal, oil, seeds, ...)
- Cocaine: Cocaine products.
- Dissociatives: Ketamine, GHB, GBL.
- Hallucinogens: LSD and related, but excluding psychedelics.
- Stimulants: All stimulants other than cocaine, including (meth)amphetamine, MDMA, MDA...
- Opioids: Heroin, opium, analgesics (e.g., oxycodone)
- NPS (Cannabinoids): Synthetic cannabinoids including spice, K2, ...
- NPS (Dissociatives): Synthetic dissociatives such as methoxetamine (MXE), dextromethorphan (DXM).
- NPS (Hallucinogens): Synthetic hallucinogens including 25i-NBOMe, 4-ACO-DMT, 2C-B, ...
- NPS (Opioids): Synthetic opioids (including fentanyl, MT-45, ...)
- NPS (Synthetic Stimulants): Other New Psychoactive Substances not classified above, e.g., mephedrone, 4-fluoroamphetamine, ...

2. *Other drugs:*

- Benzodiazepines: Benzodiazepine, Valium, Rivotril, Xanax, “downers” that are used as an anti-anxiety muscle relaxant and can be sleep-inducing.
- Prescription: Prescription drugs.
- Psychedelics: Mushrooms and other psychedelics.
- Sildenafil: Viagra and related products.
- Steroids: Steroid products.

3. *Non-drugs:*

- Drug paraphernalia: Bongs, pipes, scales, ...
- Digital goods: All forms of digital goods (including forgeries, credit card numbers, e-books, etc...).
- Electronics: Electronic items and components
- Misc: Miscellaneous items not categorized in any other category.
- Tobacco: tobacco products, including e-cigarettes.
- Weapons: all sorts of firearms, weapons, etc.

Similar to previous work [13], we evaluated the classifier using 10-fold cross validation. The overall precision and recall were both (roughly) 0.97, meaning in plain English that the classifier gets things right about 97% of the time. We evaluated the classifier on data from the Agora marketplace when trained with samples from the Evolution marketplace and vice-versa to ensure that the classifier was not biased to only perform well on the distributions it was trained on. We show the resulting confusion matrix in Figure 3: classification performance is very strong for all categories. Like in previous work [13], “Misc” is occasionally confused with “Digital goods” and “Prescriptions” are occasionally confused with “Benzos,” which in fact is not necessarily surprising. NPS classification is usually pretty good, despite potential confusion with other categories (e.g., prescription drugs).

Quantities and volumes For a number of the analyses of interest in this report, we also needed to extract quantities and volumes from each listing. This is not something that was done in the original work [13]. For a vast majority of the marketplaces we studied, quantities and weights or volumes are not explicitly specified in the listings and need to be inferred from context; this is complicated by the fact that certain users use imperial units, while others use metric units. We eventually settled on a very simple strategy, of inferring volume and quantities from the item listing titles, based on “regular expression matching.” For instance, we would scan item titles for number patterns followed by the characters “G” or “grams” to infer how many grams were sold in that specific listing. While, at first glance, this seems like an error-prone heuristic, we discovered that with about seventeen regular expressions, we were able to correctly infer most of the item weights and quantities.

We evaluated the classification algorithms by picking 200 items at random, manually labeling them, and comparing the manual labels with those obtained programmatically. Of those 200 items, 142 were drug-related, and were thus useful for our purposes (the others were discarded). Out of these 142 items, we could infer quantities and volumes for 107 items (i.e., 75% of the time). Manual inspection indicated that 128 items actually had a quantity and/or a volume specified. Finally, we were able to infer both the correct volume and quantity on 98 of the 107 items. In other words, we successfully extracted the volume and quantity more than 76% of the time it was available. In slightly over 16% of the cases we completely failed to extract any information. In slightly over 6% of the cases, we extracted the correct volume, but underestimated the quantity of items. In the remaining <1% of the cases we extracted the correct quantity, but underestimated the volume. In this evaluation set, we never overestimated quantities or volumes, which means all of our estimates were conservative.

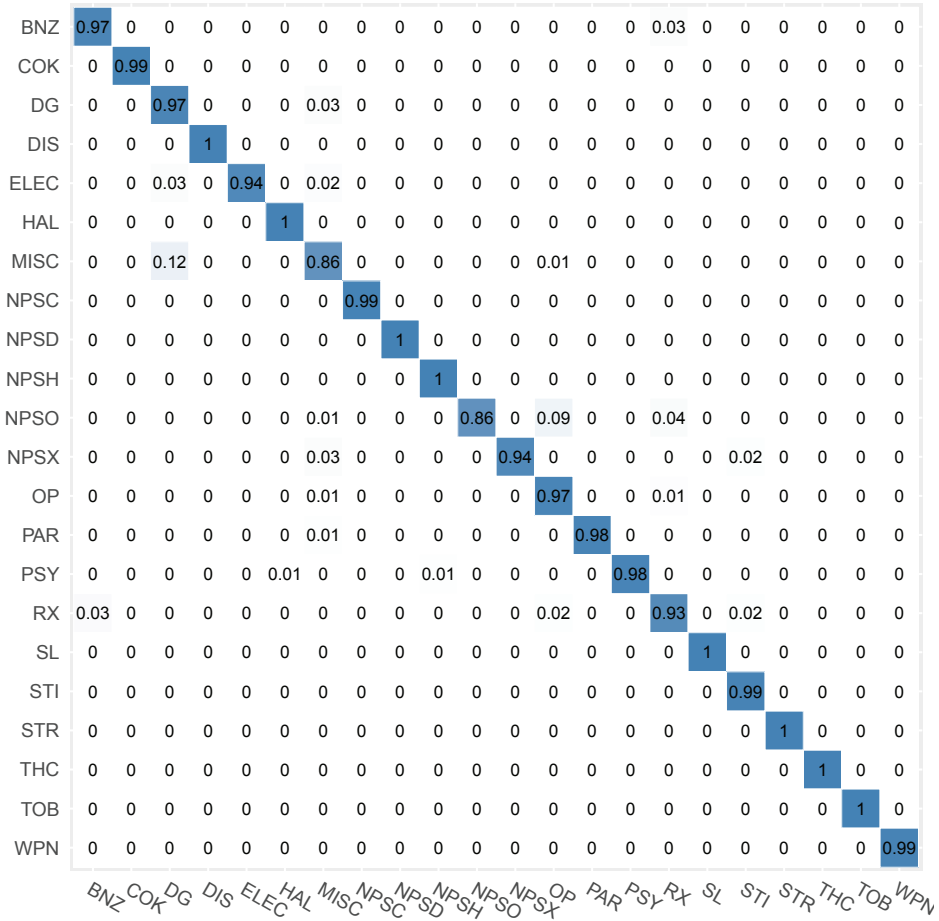


Figure 3: **Classifier confusion matrix.** BNZ: Benzos, COK: Cocaine, DG: Digital Goods, DIS: Dissociatives, ELEC: Electronics, HAL: Hallucinogens, MISC: Miscellaneous, NPSC: NPS (Cannabis), NPSD: NPS (Dissociatives), NPSH: NPS (Opioids), NPSO: NPS (Other), NPSX: NPS (Stimulants), OP: Opioids, PAR: Drug Paraphernalia, PSY: Psychedelics, RX: Prescription drugs, SL: Sildenafil, STI: Stimulants, STR: Steroids, THC: Cannabis, TOB: Tobacco, WPN: Weapons.

Extraction failed when, for instance, certain items consisted of “package deals” such as a (small) dose of MDMA coupled with a (small) dose of cannabis.

Origin countries Last, we also need to assert which countries products are shipping from.² Vendors typically indicate where they are shipping from in a specific field, and have no incentive to lie, as buyers can verify the postmark when they receive parcels. On the contrary, they have a strong incentive to be truthful, so that buyers can properly estimate shipping times. On a number of marketplaces, however, the field indicating where the product is shipping from allows for free-form entries. Instead of selecting a particular country, some vendors elect to be more vague, e.g., “a small central EU country,” or can type several locations, e.g., “UK, Belgium and Germany.” To resolve such ambiguities, we manually label every one of the 371 single different entries we come across in our database. Of those, we identify 149 denoting an EU origin (including the UK, and along with Turkey and Norway). To determine EU origin, we adopt a conservative approach: if several countries are mentioned, we require that all countries listed be in the EU. For instance, an item listing “USA, UK, Singapore, India” as its origin is considered as a non-EU item. Our manual mapping reduces the list of 149 origins to the 28 member states, Turkey, and Norway. When several countries are mentioned, we label the country as “EU – Other.”

Currency conversion Most sales on underground marketplaces are made in bitcoins, a currency that has shown very high volatility [5], which in turn is a frequent source of measurement errors (e.g., the infamous estimate that Silk Road had grossed about \$1.2 Bn when the reality was far more modest [10]). To address this issue, we convert bitcoins into their values in euros *at the time of each observed sale*. To perform the conversion, we use the average daily exchange rate on Bitstamp, as reported by bitcoincharts.com [1], on the day of the sale. When (some of the older) data includes only prices in US dollars (i.e., the original price in bitcoins was not recorded), we convert prices from US dollars to euros using the USD-EUR exchange rate at the time of the sale.

3.3 Assumptions and limitations

As discussed in the original work on the subject [13], a study of this magnitude, on field data, relies on a number of assumptions. In addition, it inherently suffers from a number of limitations. We next discuss these assumptions and limitations.

Lack of buyer information Most marketplaces do not give any information about product buyers. A handful of marketplaces provide obfuscated buyer information (see Figure 2 for an example on AlphaBay), but unfortunately, that information is not readily amenable to analysis. We are unable to determine, for instance, buyer location. One possibility, to determine where items are shipped, would be to look at the list of *possible* destinations, similar to the original study of the Silk Road marketplace [6]. However, this is a very poor proxy for actual sales. For instance, a given vendor may agree to ship to all of Europe, but their sales may be concentrated to one or two countries. In addition, a significant number of vendors agree to ship

²We discuss in Section 3.3 why it is extremely difficult to determine where products are actually being shipped *to*.

worldwide, making it all but impossible to determine where products actually ship. For these reasons, we elected to simply not consider possible shipping destinations.

Incomplete data coverage Even though we have coverage of a number of marketplaces, as described in Table 1, coverage is, as explained in the original study this work is based on [13], imperfect, given that it is very difficult to ensure that a “scrape” of an online anonymous marketplace is complete, particularly when that marketplace is large. In addition, the data solely covers the collection interval of earlier work [13]. We are not, for instance, analyzing any data from February 2015 on. In particular, we are not considering AlphaBay in our analysis, given that it still represented very modest volumes at the time we stopped collection. We also were not able to fully parse certain marketplaces. Furthermore, there may be certain (brief) temporal gaps in data collection of the marketplaces we were able to parse. A full account of the issues we faced is given in the paper fully describing the data collection infrastructure [13].

Bulk items We also filter out all items with a sales price greater than \$10,000 from the analysis in the next sections. As discussed in previous work [13], such items are rare, and tend to more often reflect a technique used by vendors discourage customers from buying a specific out-of-stock item, without removing the listing (and thus discarding the reputational data associated with it). We note that Aldridge and Décary-Héту found, on the Silk Road marketplace, $n = 52$ high-priced items that were apparently legitimate, and corresponding to bulk sales [4]. Such items are eliminated from our analysis, which could bias our work against bulk sales. We note, however, that 1) these items represent a very small fraction of all listings, and 2) sales are rare (given the amount of money at stake). As a result, we do not think the data was particularly biased by our filtering techniques. In fact, external validation of sales volumes – by comparing it to data provided by law enforcement in times of arrests – showed that our estimates were very close to the actual sales realized [13].

To put this assertion to the test, we examined all items with a sales price *consistently* greater than \$10,000. That is, the price needed to be in the top three quartiles of all prices, and should not have been greater than 100 times the minimum price for the item. We found $n = 2,131$ such items. We then filter out all of the items for which we did not have any records of any sale having taken place: this brings the number of items to consider to $n = 202$. Finally, we look only at the items shipping from the EU, Norway or Turkey: this further brings down the number of items to consider to $n = 93$. We manually inspect these $n = 93$ items, removing manually all those explicitly marked as “sold out” or with quantities that make it implausible they would sell at such prices (e.g., “1g of weed”). We are left with $n = 31$ plausible candidates:

- 11 of those are “custom orders” for very specific individuals, and for which no detail is available—those could be bulk purchases, or the price might have been artificially inflated to prevent others from ordering these listings.
- 11 listings correspond to 1 kg of MDMA crystals; 1 listing corresponds to 3 kg.
- 6 listings are for ecstasy pills—usually 160–200 mg per pill, sold in batches of 5000 (two batches are for 10,000 pills, and is more expensive).
- The last 3 items correspond to 1 kg of cocaine, heroin, weed, respectively




Figure 4: **Breakdown of sales revenue originating from the European Union (plus Norway and Turkey) by country.** For readability, the three major countries are represented on a different scale (4(a)).

21 of these listings have only one sale; 6 have 2 sales; 4 have 3 sales. Considering that each of these items sells for – usually – around USD 10,000, the overall impact on our measurements appears rather negligible.

Ultimately, we believe that the data we have give a fairly accurate depiction of what was happening in the early years of online anonymous marketplaces, at the time the ecosystem was growing quite significantly.

Automated classification Manual inspection of all data is impractical at the scale we are considering (128,618 items in total), and thus we need automated classifiers and analysis routines to do most of the processing we rely on. Sometimes, automated classification or data extraction fails. However, the numbers we earlier reported on the performance of our algorithms tend to substantiate these algorithms 1) have good accuracy and 2) tend to err on the side of producing conservative estimates, rather than inflated ones.

4 Data analysis

We next turn to analyzing the data we collected. We first look at sales volumes originating from the European Union (and Turkey and Norway), and compare them with sales originating outside of the European Union. We then perform an analysis of quantities being sold, before looking at vendor characteristics.

4.1 Sales from EU sellers

For the seven categories of drugs of primary interest (see Section 3.2), Figures 4 and 5 presents a breakdown of sales originating from the European Union (plus Norway and Turkey) by country. Both plots are stacked

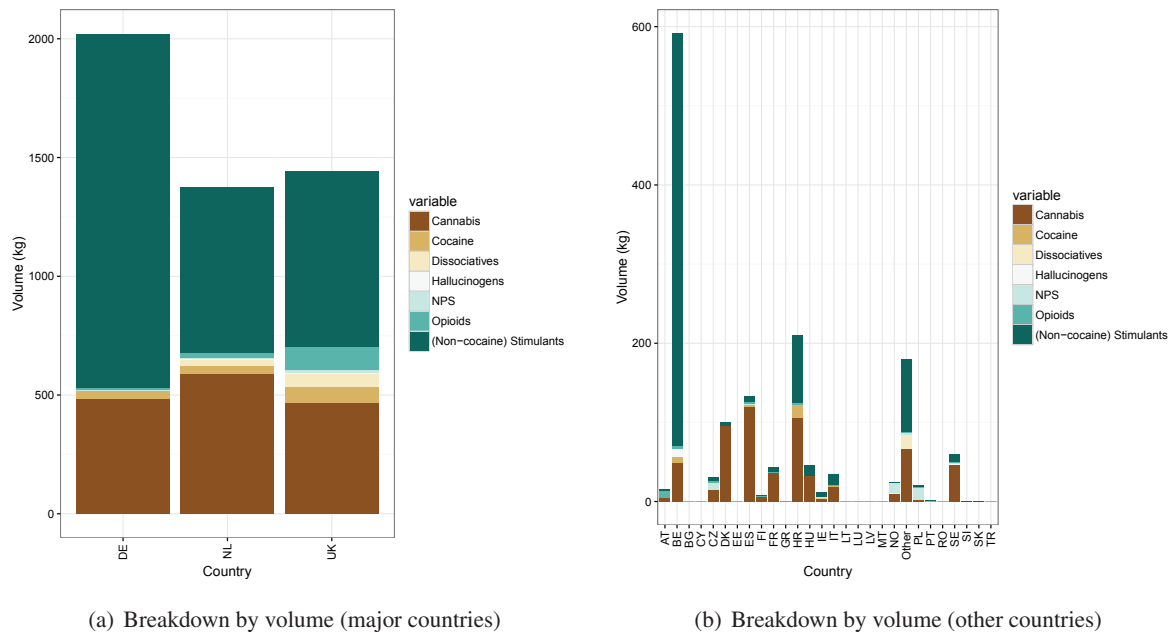


Figure 5: **Breakdown of sales volumes originating from the European Union (plus Norway and Turkey) by country.** For readability, the three major countries are represented on a different scale (5(a)).

plots. NPS are aggregated in a single category here. Figure 4 represents the aggregate amount of transactions over our entire data collection interval (November 22, 2011–February 16, 2015). We caution, however, against extrapolations of yearly revenues given 1) the rapid changes in the whole ecosystem in that period [13], as well as 2) our own data collection limitations (see Section 3.3).

Revenue analysis These limitations notwithstanding, a pattern clearly emerges from Figure 4. The vast majority of sales originating from the EU comes from three countries: Germany, with approximately EUR 26.6 million total sales for the seven drug categories of interest; the United Kingdom with slightly over EUR 20.3 million sales; and the Netherlands, with slightly more than EUR 17.9 million sales. There is then a precipitous drop – Belgium (EUR 4.7M), Croatia (EUR 2.3M), Sweden (EUR 1.3M), Spain (EUR 1.2M) and “Others,” i.e., those items purporting to ship from multiple possible locations (EUR 1.1M), are the only countries to be over EUR 1M in total sales.

We further see that, among the top four countries, Germany (EUR 14.5M), the Netherlands (EUR 8.8M) and Belgium (EUR 3.6M) are primarily selling stimulants other than cocaine, that is, MDMA, ecstasy and related products. In Germany and the Netherlands, cocaine and cannabis sales are also significant (EUR 5.6M of cannabis sales originate from Germany, EUR 3.7M originate from the Netherlands; EUR 4.2M of cocaine sales originate from Germany, while EUR 2.6M originate from the Netherlands). The United Kingdom, on the other hand, presents a more balanced revenue across all drugs: non-cocaine stimulants represent roughly EUR 5.6M of all drugs sales, cannabis accounts for EUR 4.9M, and cocaine, EUR 4.6M. UK sellers also appear to dominate the market for dissociatives (EUR 1.7M) or NPS (EUR 852K), which are far more modest

Table 2: **Comparison of drug vs. other sales in the European Union, and the rest of the world.** Volumetric breakdowns are not given for total sales, given that volumes make no sense for certain items, e.g., digital goods.

	Drug sales³		Total sales
	Volume (g)	Revenue (EUR)	Revenue (EUR)
European Union (plus Norway, Turkey)	5,523,695	79,012,948	86,126,606
Rest of the world	10,877,448	93,332,764	115,736,370
Total	16,401,144	172,345,712	201,862,977

in other countries.

Volume analysis Figure 5 shows a similar breakdown, but this time, by volume. The general trends observed with respect to financial revenue hold here as well: Germany (2,027 kg overall), the Netherlands (1,279 kg overall), and the United Kingdom (1,172 kg overall) dominate the ecosystem; these are the only countries where products shipped exceed, in aggregate, a metric ton. Due to the vastly different prices per unit, in this volumetric representation, cocaine, opioids, hallucinogens are far less represented than in the revenue representation; conversely, cannabis is far more significantly represented. We will discuss in Section 4.3 the various prices per quantity we observed.

Comparison with non-EU sales Table 2 compares sales originating from the European Union (plus Norway and Turkey) to those originating from other countries, both for the drugs in the seven categories of interest, and for all products. We first notice that, for both EU countries and the rest of the world, drug sales represent an overwhelming majority of the economic revenue of these marketplaces. This is even more the case in the European Union than in the rest of the world; this discrepancy can be explained by the fact that a lot of digital goods being sold (credit card numbers, ebooks, logins, etc) are actually shipped digitally and thus are not usually classified as having a specific origin by their sellers. These digital goods represent a non-negligible portion of the overall trade [13]. In terms of drug sales, we see that EU countries represent roughly 46% of all revenue, but only 34% of all volumes. This may be explained by the fact that cannabis sales are overall more prevalent than what we see in the EU [13]. Because cannabis is generally priced at a lower price per unit than other drugs, volumes observed in the rest of the world stand to be higher.

4.2 New Psychoactive Substances

We next turn our attention to New Psychoactive Substances, or NPS. The legal status of many of these substances, frequently referred to as “legal highs,” is murky, and subject to rapid changes. The high-level

³In the seven categories of interest.

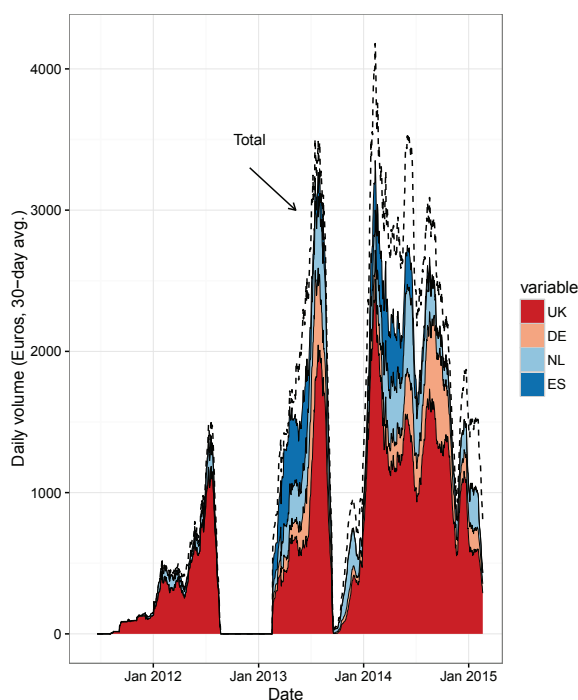


Figure 6: **Breakdown of NPS sales originating from the European Union (plus Norway and Turkey).**

takeaway from our measurements is that NPS represent a very small proportion of all trade on online anonymous marketplaces, and that the vast majority of NPS sold on these marketplaces apparently come from the United Kingdom.

Figure 6 provides a finer-grained view of the sales of NPS on online anonymous marketplaces, over time, as a stacked plot. Keep in mind that the data is subject to the same limitations as the data described in the original paper [13], which explains, for instance, the collection gaps in late 2012—most markets were not functional at that point, and we have scant data for those that were, like Black Market Reloaded or Sheep marketplace. For readability purposes, all datapoints, here, represent averages over a 30-day moving windows. The crux of the figure is that NPS volumes rarely represent more than EUR 3,000/day. Interestingly enough, most of the NPS being sold on online anonymous marketplaces in the time interval of our study were hallucinogens – synthetic cannabinoids, dissociatives, opioids and stimulants are almost negligible. Figure 6 confirms that the majority of NPS seem to originate from the UK; Germany, Netherlands, and Spain also contribute, albeit far less significantly. The top line correspond to the aggregate of all countries including those not represented individually on the plot.

Variations in revenues follow the growth and decline of the overall online anonymous marketplace ecosystem. We caution against a hasty interpretation of the observed decreases in early 2015: this data represents the ecosystem immediately after Operation Onymous (when a number of marketplaces were taken down), and corresponds to the end of our measurement interval. This means that it is statistically less reliable than earlier data, for reasons fully detailed in previous work [13], having mostly to do with

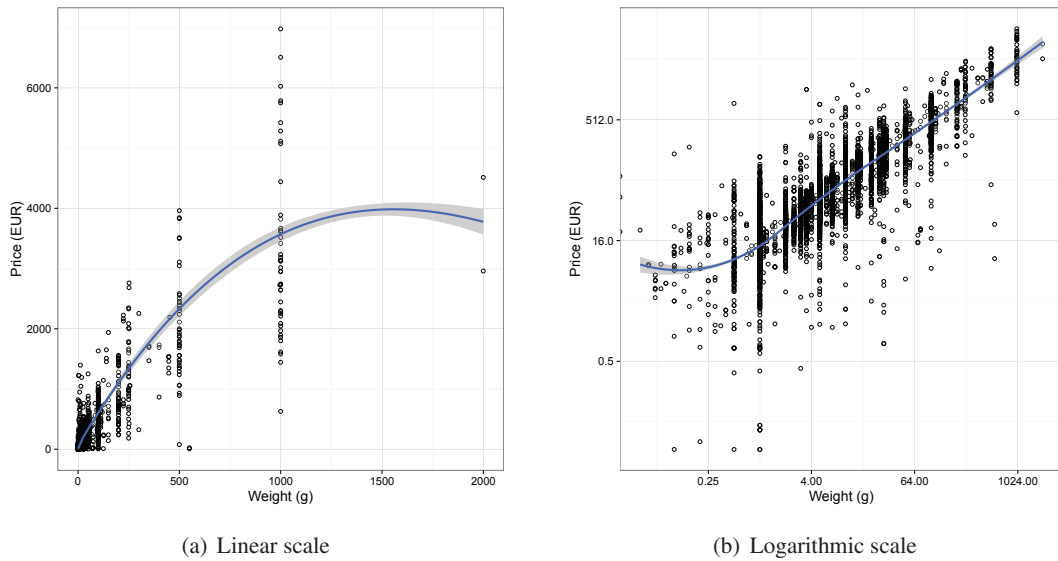


Figure 7: **Cannabis prices as a function of weight.** The curve is a local polynomial regression fitting, the gray shade corresponds to the 95% confidence interval.

incomplete coverage of every single scrape.


4.3 Transaction amounts broken down by drug and level

We next turn to a discussion of the transaction amounts broken down by drug and by level.


Cannabis Figure 7 is a scatter plot, in which each point represents the weight (x coordinate) and price (y coordinate) of each cannabis item in our dataset for which 1) we could infer the weight and 2) we observed at least one transaction. The left hand side (Figure 7(a)) shows this scatter plot on a linear scale. Because the vast majority of items correspond to small quantities, we find it useful to present the same data on a logarithmic-logarithmic scale (Figure 7(b)). The blue curve corresponds to a non-parametric regression (using a local polynomial regression fitting, [7]), with 95% confidence interval in the gray shade. There appears to be a modest volume discounting effect; however we caution that the number of observations at high volumes are considerably smaller than at low volumes, and thus regression fits are probably more questionable at high volumes.

The most common units sold are 5g (1,745 items, mean price EUR 58, standard deviation EUR 39), 1g (1,610 items, mean price EUR 17, standard deviation EUR 16), and 10g (1,165 items, mean price EUR 99, standard deviation EUR 55). The high standard deviations are explained by the fact that various products (oils, edibles, etc) are classified as cannabis, so that there is quite a large price dispersion.


We note the presence of a few items with very small volumes (close to zero gram); some of these are sample offers, some are the few items for which our heuristics for extracting quantities might have failed.



(a) Cocaine (Linear scale)



(b) Cocaine (Logarithmic scale)



(c) Stimulants other than cocaine

Figure 8: **Cocaine and other stimulant prices as a function of weight.** The curve in 8(a) and 8(b) is a local polynomial regression fitting, the gray shade corresponds to the 95% confidence interval.




Figure 9: **Dissociative (Ketamine) prices as a function of weight.** The curve are local polynomial regression fittings, the gray shade corresponds to the 95% confidence interval. Note the different estimators depending on the type of dissociative sold. We removed the (separate) plots for GHB and other dissociatives (PCP), as there were too few datapoints.


Cocaine and Other Stimulants We next present, in Figure 8, similar plots for cocaine products. the first plot (Figure 8(a)) represents a linear scale, and the second plot (Figure 8(b)) represents a logarithmic scale. We use again a local polynomial regression, which shows that volume discounting effect is markedly more pronounced here than it was in the case of cannabis products.

The most common unit sold is 1g (664 items, mean price EUR 84, standard deviation EUR 30).


Stimulants (other than cocaine), in Figure 8(c), exhibit a near linear relationship for products labeled as MDMA (in red). Other products include various types of drugs (anything labeled “ecstasy,” “speed,” “meth,” ... would end up in this category), of highly varying quality. We point out here that we defined the weight as the product of the unit weight by the number of units. For instance, somebody selling 1,000 pills of 200mg MDMA pills would be considered as selling 200 grams. The price dispersion observed is compounded by the fact that a number of sellers offer “lottery sales,” which consist of a single item being sold at a very discounted price to multiple buyers, only one of whom will actually “win” it.

Dissociatives We next turn to dissociatives. Figure 9 provides a similar scatter plot in log-log scale for ketamine. There is not much bulk discounting, apparently. Data for other products (GHB/GBL, and PCP and others) was too limited to provide meaningful regressions and is omitted from the plot (it would show disjoint clusters, but with only a few points).


The most common unit sold (for Ketamine) is 1g (143 items, mean price EUR 40, standard deviation EUR 18).



(a) Hallucinogens



(b) NPS (Hallucinogens)



(c) Opioids

Figure 10: **Prices as a function of weight for hallucinogens, NPS (hallucinogens), and opioids.** All plots are in log-log scale; the curves are local polynomial regression fitting, the gray shades correspond to the 95% confidence interval.

Other drugs: Hallucinogens, NPS (Hallucinogens), and Opioids We next turn to the other three categories of drugs of interest, depicted in Figure 10, in which all scatter plots are using a log-log scale. Hallucinogens, i.e., primarily LSD, represented in Figure 10(a), show two clear clusters: there is a large price variance for low, single use, type of doses ($250 \mu\text{g}$ and less). This may be partially due to measurement errors at such low levels, and partially due to vendors offering samples for (nearly) free, or as part of lotteries. As the volume grows, the price grows as well somewhat linearly.

In Figure 10(b), we show that NPS hallucinogens (NBOMe, DMT, ...) present the same behavior albeit with considerably larger weights (grams as opposed to milligrams) than LSD, which is unsurprising. The higher quantities (and to a certain extent, lower quantities as well) have high uncertainty due to limited data. Interestingly, these hallucinogens constitute the vast majority of the NPS sales in our database, as, apparently, during our measurement intervals, synthetic cannabinoids were not very well represented on underground marketplaces).

Opioids, presented in Figure 10(c), show considerable price dispersion. The regression, using again a local polynomial regression fitting, is not particularly conclusive but suggests modest volume discounting; for high volumes ($>16 \text{ g}$), the scarcity of data makes the regression very imprecise.

4.4 Vendor diversification

We next examine the range of products and volumes vendors offer. We start by looking at if, and how, vendors diversify in terms of volumes they offer. That is, we try to determine whether vendors who stay within their market echelon (e.g., always selling small quantities), or if, on the other hand, they diversify their offerings, and to which extent. We then discuss whether vendors who sell drugs *also* sell other types of goods.

Throughout the following discussion we will use a coefficient of diversity, as defined by Soska and Christin [13]. In short, we divide all items of interest in a set \mathcal{C} of groups. For instance, \mathcal{C} could denote a set of volume tiers, or a set of item categories. Let \mathcal{S} be the set of all sellers based in the EU (plus Norway and Turkey) across all marketplaces. We define $\mathcal{C}_i(s_j)$ as the normalized value of the i -th group for seller j such that $\forall s_j \in \mathcal{S}, \sum_{i=1}^{|\mathcal{C}|} \mathcal{C}_i(s_j) = 1$. For instance, if vendor j 's revenue comes from 50% of items in group 1, 25% in group 2, and 25% in group 3, then $\mathcal{C}_1(s_j) = 0.5, \mathcal{C}_2(s_j) = 0.25, \mathcal{C}_3(s_j) = 0.25$. We can then define the coefficient of diversity for seller s_j as:

$$c_d(s_j) = \left(1 - \max_i (\mathcal{C}_i(s_j))\right) \frac{|\mathcal{C}|}{|\mathcal{C}| - 1}.$$

Intuitively, the coefficient of diversity is measuring how invested a seller is into their most popular group, normalized so that $c_d \in [0, 1]$. A vendor s_j with a coefficient of diversity of zero sells only items from a specific group; a vendor s_j with a coefficient of diversity of one gets exactly the same amount of revenue from each group of items.

Diversification in terms of volumes offered We use the coefficient of diversity so defined to examine whether vendors who sell large quantities sell also small quantities. To do so, we define quantity tiers for each drug category in Table 3. We use a simple three-tier distinction between retail, middle-market, and bulk




Figure 11: Coefficient of diversity for vendors, by product, across volume tiers.

Table 3: Quantity tiers for each drug of interest.

	Cannabis	Cocaine	Hallucinogens	Opioids	Stimulants
Retail	<100 g	<10 g	<8 mg	<1 g	<10 g
Middle-market	100–999 g	10–999 g	8–159 mg	1–999 g	10–999 g
Bulk	≥ 1000 g	≥ 1000 g	≥ 160 mg	≥ 1000 g	≥ 1000 g

sales. We base our distinction between tiers on the EMCDDA’s own classification [9]. For hallucinogens, the tiers are usually expressed in number of doses (50 doses or less, 50–1000 doses, more than 1000 doses); since we are using grams as a base unit, we convert this to volumetric tiers using a baseline of 160 microgram doses (which is close to the arithmetic mean of what we observed.) We exclude from this discussion NPS, as they are too heterogeneous a set to provide meaningful comparisons, *and* the uncertainty on their legal status makes it hard to pick appropriate thresholds.

Here, we have $|\mathcal{C}| = 3$. Then, for each group of drugs listed in Table 3, we plot the cumulative distribution function of the coefficient of diversity of their vendors in Figure 11. The figure shows that an overwhelming ($\approx 90\%$) majority of cannabis vendors stay within one volume tier (coefficient of diversity of zero). The rest are more spread out, without any noticeable jumps; almost no vendor has a coefficient of diversity greater than 0.75. In other words, most vendors stay in a single volume tier, but a minority sell across two tiers; almost no one has meaningful sales across three tiers. It is quite rare for a vendor to sell both bulk-size quantities and small volumes at the same time, but some vendors selling larger quantities sometimes offer “testing samples” to their customers. Cocaine shows a similar picture. On the other hand, hallucinogens, opioids, and especially stimulants, present more diversity, with most vendors sticking to the retail tier, but some vendors selling across multiple tiers with a number of items in each tier.




Figure 12: **Coefficient of diversity for vendors, across product categories.**

Anecdotally checking our database suggests that vendors selling in multiple echelons tend to be “super-stores,” that also carry more than one type of drug. They also are more likely to have higher sales volumes. Conversely, vendors who stick to one echelon (typically, the lowest one) tend to specialize in one item, and to have relatively low sales volumes.

Diversification in terms of products offered We next look at the diversity across products being sold. We define \mathcal{C} here as the set of drugs of interest: Cannabis, Cocaine, Dissociatives, Hallucinogens, NPS, Opioids, and Stimulants. We plot, in Figure 12 the corresponding cumulative distribution function of the coefficient of diversity for all vendors. We see that approximately half of all vendors specialize in exactly one category – this is frequently the case for cannabis (566 cases) and stimulants other than cocaine (422) vendors, which is not very surprising given that those are very frequently sold items. The other half is far more diverse, and typically those vendors sell from a couple of categories. For instance, vendors selling dissociatives hallucinogens rarely only sell from these categories. A very small number of vendors have a coefficient of diversity close to one, denoting that they sell a little bit of everything (e.g., MDMA, Ketamine, DMT and weed; LSD, weed, MDMA, and mushrooms, ...). Those vendors usually focus purely on individual or small doses.

Finally, we examined all 3,305 vendors reportedly shipping from the European Union. Out of those, 2,062 sell drugs in one of the seven categories of interest. Of those 2,062 vendors, 346 (or 16.78%) also sell other types of drugs (e.g., prescription drugs). Surprisingly to us, 897 (43.50%) also sell non-drug products. Of those, the vast majority appears to be digital goods. 318 of those 897 vendors that sell non-drug products remain confined to one drug category (primarily cannabis, and stimulants other than cocaine). 214 of those 897 vendors have on the other hand high diversity coefficients (>0.5). Many of these vendors sell cannabis, stimulants (sometimes including cocaine), and opioids. Only a handful of vendors (five) selling non-drugs are bulk vendors – three sell large amounts of hallucinogens, one sells opioids, and one sells stimulants other

than cocaine.

Vendors selling under multiple aliases or marketplaces The entire discussion above assumes that every vendor account denotes a unique vendor. However, it is not infrequent that a vendor sells on more than one marketplace [13]. We thus try to identify, among the 3,305 vendor accounts earlier identified, which ones belong to the same person(s).

To do so, we use the exact same heuristics as in our earlier work [13]: vendors with the same name, or whose names only differ in character case (e.g., “Sally Seller” vs. “sally seller,” or “Sally seller”), but operating on different marketplaces are assumed to be identical. Likewise, we use results from the Grams search engine [2], which includes a seemingly manually curated database of vendor aliases, to assert different monikers belong to the same person.

Using this set of heuristics, the 3,305 vendor accounts in the European Union appear to map to 2,180 unique entities, 1,271 of which sell drugs. 226 (17.78%) vendors sell at least two different types of drugs, and 683 (53.74%) vendors sell drugs and other products (e.g., digital goods).

These heuristics are admittedly imperfect. If a vendor uses two radically different aliases, and Grams is not aware of them belonging to the same person, we would consider these aliases to be distinct vendors. Likewise, it is possible that somebody impersonates a vendor by reusing their moniker on a different marketplace; in fact, this may be a good strategy to profit from an establishing vendor’s reputation and gain customers. For this reason we have focused in the rest of our analysis on considering unique “aliases,” without attempting any reconciliation; but wanted to give a quantitative idea of the potential differences between monikers and actual vendors.

5 Conclusions

We have performed an analysis of the online anonymous marketplace data collected by Soska and Christin [13] over late 2011–early 2015. Complementary to Soska and Christin’s original analysis, we have primarily focused on drug supply coming from the European Union (extended to include Turkey and Norway). We found that, for the period and the marketplaces considered, EU-originating drugs primarily came from Germany, the Netherlands, and the United Kingdom. Other countries, such as Belgium, Spain, Sweden, or Croatia, were far behind. Many countries did not register any significant sales. Nevertheless, EU-based suppliers represented a significant share of all revenue—approximately 46% of all drug sales, when overall, EU-based vendors (drugs and non-drugs) represented 43% of all sales.

Slightly different from the overall ecosystem, in which cannabis-related sales appeared more prevalent [13], stimulants and cocaine altogether represented a majority of all EU-based sales. Hallucinogens, opioids were far less significant; and dissociatives and NPS were even more modest. A complementary study of NPS indicates that the supply was heavily concentrated in the United Kingdom. The same appeared to hold true for dissociatives. However, NPS volumes remained comparatively very modest – with revenues in the order of EUR 3,000 at market peak.

We saw that, while vendors on these marketplaces primarily cater in the retail space, with individual item weights and volumes frequently corresponding to personal amounts, there is evidence of much larger

(bulk-like) sales. Regression-based analyses show that volume-based discounting tend to occur, albeit at relatively modest levels.

Finally, the ecosystem is roughly split in half: half of the vendors are specializing in one type of drug, while the other half is far more diversified. Slightly less than half of the drug sellers tend to stick to a given weight echelon, while others present a more diverse set of offerings.

Acknowledgment

This work greatly benefited from extensive discussions with, and feedback from, Andrew Cunningham, Teodora Groshkova and Mahmood Sharif. Kyle Soska was instrumental in helping redesign the classification primitives.

References

- [1] Bitcoin charts. <http://bitcoincharts.com/markets/>.
- [2] Grams: Search the darknet. <http://grams7enufi7jmdl.onion>.
- [3] I2P: The internet invisible project. <http://www.geti2p.net>.
- [4] Judith Aldridge and David Décary-Héту. Hidden Wholesale: The drug diffusing capacity of online drug cryptomarkets. *International Journal of Drug Policy*, 35:7–15, 2016.
- [5] R. Boehme, N. Christin, B. Edelman, and T. Moore. Bitcoin: Economics, technology, and governance. *The Journal of Economic Perspectives*, 29(2):213–238, Spring 2015.
- [6] N. Christin. Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd World Wide Web Conference (WWW'13)*, pages 213–224, Rio de Janeiro, Brazil, May 2013.
- [7] William S Cleveland, Eric Grosse, and William M Shyu. Local regression models. *Statistical models in S*, 2:309–376, 1992.
- [8] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
- [9] European Monitoring Centre for Drug and Drug Addition. EMCDDA drug prices indicator, protocol for reporting price data, 2016.
- [10] James T. Graves, Alessandro Acquisti, and Nicolas Christin. Big data and bad data: On the sensitivity of security policy to imperfect information. *The University of Chicago Law Review*, 83(1):117–137, 2016.
- [11] James Martin. *Drugs on the dark net: How cryptomarkets are transforming the global trade in illicit drugs*. Springer, 2014.
- [12] S. Nakamoto. Bitcoin: a peer-to-peer electronic cash system, October 2008. Available from <http://bitcoin.org/bitcoin.pdf>.
- [13] K. Soska and N. Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security'15)*, pages 33–48, Washington, DC, August 2015.

Author biography

Nicolas Christin is an Associate Research Professor at Carnegie Mellon University, jointly appointed in the School of Computer Science (Institute for Software Research) and in Engineering & Public Policy. He is a core faculty in CyLab, the university-wide information security institute, and also has affiliations with the Information Networking Institute and the department of Electrical and Computer Engineering.

He holds a Diplôme d'Ingénieur from École Centrale Lille (1999), and M.S. (2000) and Ph.D. (2003) degrees in Computer Science from the University of Virginia.

After a post-doctoral fellowship in the School of Information at the University of California, Berkeley, he joined Carnegie Mellon in 2005. He worked for three years as resident faculty at CMU CyLab Japan, before returning to Carnegie Mellon's main campus in 2008. He subsequently served as Associate Director of the Information Networking Institute between 2008 and 2013, and as a research faculty in the Electrical and Computer Engineering department between 2013 and 2016.

His research interests are primarily in computer and information security; most of his work is at the boundary of systems and policy research, with a slant toward networking aspects. He has most recently focused on online crime, security economics, digital currencies, and psychological aspects of computer security. His research combines field measurements and mathematical modeling.

He has published over 90 academic research papers in computer networks and security. Major recent publications include:⁴

Content Availability, Pollution and Poisoning in Peer-to-Peer File Sharing Networks (with Andreas S. Weigend and John Chuang). In Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05), pages 68-77. Vancouver, BC, Canada. June 2005. (Cited over 250 times in other academic publications, and in an amicus brief to the U.S. Supreme Court (MGM vs. Grokster).)

Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In Proceedings of the 22nd International World Wide Web Conference (WWW'13), pages 213-224. Rio de Janeiro, Brazil. May 2013. (Cited over 210 times in other academic publications. Extensively referenced by U.S. and international press, including The Economist, Forbes, Le Nouvel Observateur, Marketplace Radio among many others.)

Of Passwords and People: Measuring the Effect of Password-Composition Policies (with Saranga Komanduri, Richard Shay, Patrick Gage Kelley, Michelle Mazurek, Lujo Bauer, Lorrie Cranor and Serge Egelman). In Proceedings of the 2011 ACM Conference on Human Factors in Computing Systems (CHI 2011), pages 2595-2604. Vancouver, BC, Canada. May 2011. (Cited over 185 times in other academic publications.)

Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem (with Kyle Soska). In Proceedings of the 24th USENIX Security Symposium (USENIX Security'15), pages 33-48. Washington, DC. August 2015. (Already cited 26 times in other academic publications, despite being very recent. The most comprehensive, to-date, study of online anonymous markets.)

⁴Different from other fields where publishing in journals is the norm, in Computer Science, the most prestigious publication venues are conference proceedings of highly selective conferences, where acceptance rates typically represent between 10 and 20% of all submitted papers.