An EU-focused analysis of drug supply on the AlphaBay marketplace*

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: October 23, 2017

*This document has been prepared by the author for the European Monitoring Centre for Drugs and Drug Addiction (EMCDDA) under contract CT.17.SAT.0063.1.0. 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.

The AlphaBay marketplace, which was in operation between December 2014 and July 2017, reportedly became the leading marketplace during that time. In this report, we present an analysis of sales on AlphaBay, with a 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:

- AlphaBay did become a very large marketplace, with daily sales overall exceeding 600,000 euros at its peak in early 2017. By itself, it grossed at least as much revenue over 2014–2017 as all other marketplaces combined between 2011–2015 [8].
- EU-based suppliers represent approximately a quarter of all drug sales; this is down from 46% for marketplaces previously studied [8] in the 2011–2015 interval.
- EU-originating drugs primarily came from Germany, the Netherlands, and the United Kingdom.
- Cannabis, cocaine and other stimulants altogether represented a majority of all EU-based drug sales.
- Supply of New Psychoactive Substances (NPS) remained very modest with revenues below EUR 2,500 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.
- Save for the decreasing share of European sales, most of the trends observed in this report confirm what we had previously found for other marketplaces in the 2011–2015 interval [8]. In other words, the ecosystem, as a whole, appears relatively stable over time.

1 Introduction

By using a combination of network-level anonymity technology [11] and pseudonymous [16] or anonymous [5, 19] payment systems, 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 [7].

This report is a sequel to, and builds up on, a previous report released in 2016 [8]. In that previous effort, we analyzed data collected by Soska and Christin [17], spanning several years (late 2011–early 2015) including the "early days" of online anonymous marketplaces; in the present report, we extend our analysis to cover the AlphaBay marketplace throughout most of its existence (December 2014–May 2017; AlphaBay was shut down in July 2017 by law enforcement). Both reports share some amount of prose – particularly when results are near-identical in both the datasets considered in our previous report [8] and here – but the data presented in the present report is completely new.

As in our previous work [8], we focus on analyzing drug supply reportedly originating from the European Union, and derive relationships between financial revenues and actual quantities (weights, volumes, units) of products being sold.

This report is organized as follows. We briefly discuss how online anonymous marketplaces operate and their history in Section 2. We move to discussing our methodology, which shares many of the same assumptions and limitations as the original study by Soska and Christin, along with certain particularities linked to AlphaBay, in Section 3. We turn to analyzing the collected data in Section 4 before drawing brief conclusions in Section 5.

2 Background

Narcotics have been traded on the Internet pretty much since the advent of the World Wide Web. The earlier trading platforms (e.g., the Hive, Pondman, RACResearch, OVDB, and others) were primarily discussion forums. While quite open, these forums catered to relatively limited numbers of users, and offered very weak anonymity guarantees, leading their patrons to take relatively high risks when ordering online.

The creation of the Silk Road website in 2011 marked a drastic change in the way drugs were traded online [7]. The combination of network-level anonymizing technology (Tor [11] or i2p [1]), crypto-currencies, such as Bitcoin [16], with better privacy protection than traditional online financial instruments, and media exposure [6, 14] ushered a new era in the online drug trade. We refer the reader to Christin [7], Martin [15], and Soska and Christin [17], for a thorough description of how modern online anonymous marketplaces (or "cryptomarkets") are designed. In a nutshell, cryptomarkets resemble e-commerce marketplaces such such as eBay or the Amazon marketplace; however, the level of anonymity they promise, sometimes a bit too confidently, facilitates the trade of illicit or illegal items.

2.1 A short history of cryptomarkets

Between February 2011 and November 2013, Silk Road [7] was pretty much in a monopoly position. While other cryptomarkets appeared – Black Market Reloaded, Sheep Marketplace, and Atlantis, notably – their revenues were comparatively extremely small [17]. However, in November 2013, the Silk Road operator was identified, arrested, and the site taken down. This resulted in major changes to the ecosystem: former also-rans (Black Market Reloaded and Sheep Marketplace) all of sudden saw an influx of Silk Road customers coming to patronize their markets; numerous short-lived markets (e.g., Black Flag) unsuccessfully tried to fill the hole left by Silk Road's demise. This chaotic situation stabilized about a month later, with the emergence of Silk Road 2.0, ran by former Silk Road staff, and using an interface visually strikingly similar to Silk Road's; a number of other marketplaces (Pandora, Hydra, Evolution, Agora, Cloud 9, ...) also appeared within a few months, resulting in a diverse ecosystem, that flourished throughout 2014. Indeed, the aggregate revenue of all markets quickly exceeded what Silk Road was earning in its heyday [17]. But again, law enforcement intervened in November 2014, through Operation Onymous [18], which led to the shutdown of Silk Road 2.0, as well as that of a number of lesser-known sites (e.g., Cloud 9). The immediate effect was that traffic primarily concentrated into the Agora and Evolution marketplaces from November 2014.

2.2 The AlphaBay marketplace

The AlphaBay marketplace was reportedly designed in mid-2014 [2]. It went online on December 26, 2014, shortly after Operation Onymous took place. Similar to the Evolution marketplace, AlphaBay was reportedly started by "carders," i.e., people who had been trading pilfered credit card numbers and other banking credentials. However, AlphaBay quickly started offering listings for narcotics as well. AlphaBay was originally a fairly small marketplace, overshadowed by Evolution and Agora. By mid-2015, it started to get considerably increased exposure as Evolution closed its doors, and reportedly became one of the leading markets later that year. By 2016, it was, supposedly, the undisputed leader in the cryptomarket space. Through our analysis in Section 4, we will be able to confirm these historical accounts.

3 Collection Methodology and Data

3.1 Data collection

For this report, we use data collected by our research group between March 18, 2015, and May 24, 2017 of the AlphaBay marketplace. We have, all in all, collected 27 (full) scrapes of the AlphaBay marketplace over that time interval.

Our data collection infrastructure is identical to the platform we used for other marketplaces, and which we described in previously published work [17]. In short, we devised 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 a companion paper [17].

The data solely collected from the AlphaBay marketplace represents more than 471 GB of storage. The

SQLite database containing the parsed information is approximately 20 GB. In comparison, the data from all marketplaces in 2011–2015 used in the previous report [8] corresponded to a 19 GB database, and used more than 3 TB of storage [8].

3.2 Data classification and processing

As was the case with other marketplaces, once all listings have been parsed and stored into a database, data 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. We use exactly the same methodology as in our previous report [8], which we summarize here.

Item categories As discussed in our 2015 paper [17], categories self-reported by sellers, e.g., "Stimulants/Cocaine" 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 our 2015 paper [17], but uses different categories of interest. In an effort to provide heads-to-heads comparison with other marketplaces, we consider the same following 22 categories as in the previous report [8]:

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 fentanyls, 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.

In the commission of our previous report, we evaluated this redesigned 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, when compared to our baseline. 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.

A key caveat, however, is that the classifier can only perform as well as its training set; that is, if the samples used for training are incorrectly labeled, the classifier will not be able to rectify the error. In the reminder of this report, and similar to our previous efforts [8, 17], the training set used is Evolution data, which has the advantage of presenting an overwhelmingly large number of correctly labeled items. Unfortunately, there are categories in which classification is inherently ambiguous, even for an expert manually labeling samples.

This is particularly true of NPS, which tend to be bundled with the type of chemicals they attempt to emulate—a notable exception are NPS Hallucinogens, which tend to be properly labeled. On the other hand, opioids are much more muddled. For instance, fentanyl is frequently not labeled as "Fentanyl" (which itself would map to NPS (Opioids)) but as the more general "Opioids" category. As a result, like in the previous reports, it is hard to assess through this automated classification system the actual proportion of NPS, as opposed to heroin and morphine, in opioids. Classification difficulties do not only affect NPS, but might also happen with other categories. For instance, Oxycontin could be classified as an opioid or as a prescription drug, and both are actually correct.

To improve the quality of the training, we have two options worth considering in future work:

- Manually label a large number of items and use them as training. While this approach should provide a very reliable training set, it does not scale to very large numbers, and it is unclear that it will actually achieve higher performance than our current training. Indeed, consider a 1,000-strong training sample. This would probably would take a few hours to label (using multiple experts to ensure agreement). Yet, a single error would have a (very roughly) equivalent impact to 100 errors in a 100,000-strong pre-labeled dataset.

- Inject artificial entries containing strings characteristic of entries that may be ambiguous. For instance, we could create a couple of synthetic listings for fentanyl, carfentanyl, and other NPS or interest, and feed them to our existing training set. This has the advantage of keeping the benefits of our large, mostly well-labeled, training set; on the other hand, it might also bias classification toward expecting certain strings.

We emphasize however, that, by and large, classification is correct. A manual evaluation of classification "errors" revealed that misclassified items were actually properly classified, but in a category more general than we would have wished.

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. We use exactly the same strategy as in our previous report [8] To summarize, we infer volume and quantities from the item listing titles, based on "regular expression matching." For instance, we 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 in our previous work, that with about seventeen regular expressions, we were able to correctly infer most of the item weights and quantities.

As a particular improvement over our our previous report, we noticed a couple of particularly egregious mistakes with stimulants. While rare, these mistakes could have a drastic impact in overall computations. For instance, a bulk sale of several thousand pills of MDMA was incorrectly labeled as a single pill, under-estimating volumes by several orders of magnitude. We manually labeled such items (a couple dozens of them in all, out of more than 100,000 of items) after discovering them.

As in our previous report, we evaluated the classification algorithms by picking 200 items at random, manually labeling them, and comparing the manual labels with those obtained programmatically. From our AlphaBay corpus, of those 200 items, 133 were drug-related, and were thus useful for our purposes (the others were discarded). Out of these 133 items, we could infer quantities and volumes for 99 items (i.e., 74% of the time). Manual inspection indicated that 119 items actually had a quantity and/or a volume specified. Finally, we were able to infer both the correct volume and quantity on 97 of the 99 items. In other words, we successfully extracted the volume and quantity more than 81% of the time (97/119) it was available. In slightly under 17% of the cases we completely failed to extract any information. In the last 2% of the cases, we extracted the correct volume, but underestimated the quantity of items. In this evaluation set, we never overestimated quantities or volumes, which means all of our estimates were conservative.

As in our previous effort, 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. We removed such packaged deals from consideration whenever possible – simply assigning them an unknown volume.

Origin countries Last, we also need to assert which countries products are shipping from.[1] 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,

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

so that buyers can properly estimate shipping times. On AlphaBay, determining the origin country task is slightly easier than in our previous effort [8], as the entries vendors can provide are far more constrained – e.g., there is no "free-form" answer possible. As a result, we can use the manually-labeled system we used in our previous study. When several countries are mentioned, we label the country as "Other."

Currency conversion For each item, AlphaBay provides a direct conversion into US dollars on its website. Rather than collecting item prices in Bitcoin, which can be prone to errors due to the rapid fluctuations of the currency, we use this dollar value as a baseline, and then convert it to euros using the USD-EUR exchange rate at the time of the sale.

3.3 Assumptions and limitations

As discussed in previous work [7, 8, 17], a study of this magnitude, on field data, relies on a number of assumptions, and suffers from a number of limitations, which we discuss next.

Lack of buyer information AlphaBay provides obfuscated buyer information, e.g., a buyer whose account name is "`nicolasc`" would be reported in feedback as "`n...c`," but unfortunately, that information is not readily amenable to analysis. In particular, we cannot determine buyer location. In addition, other marketplaces on which we had reported previously [8], and with which we want to compare results for AlphaBay, frequently feature even less information. As a result, we do not perform any buyer analysis in this study.

Incomplete data coverage As is the case in all studies of this type [17], coverage is necessarily 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. This can lead to pernicious errors: For instance, Dolliver [12] appears to have under-estimated sales on Silk Road 2.0 by several orders of magnitude [17].

Bulk items Following the same methodology as in our previous report [8], we filter out all items with a sales price greater than $10,000 from the analysis in the next sections. The rationale, presented in our previous work [17], is that 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). However, Aldridge and Décary-Hétu 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. Manual inspection of AlphaBay listings confirms the presence of such bulk listings, which leads us to carry out a separate analysis here.

We examine 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 $n = 1,488$ 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 $n = 165$.

(Note, however, that this is a lower bound, as we cannot claim perfect coverage.) 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 $n = 25$. We manually inspect these $n = 25$ items. All of them appear plausible:

- 6 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.
- 5 listings are for MDMA crystals sold in bulk: 1 kg (1 listing), 2.5 kg (2 listings), and 5 kg (1 listing)
- 5 listings are bulk cocaine sales: 125 g (1 listing), 250 g (3 listings) and 1 kg (1 listing)
- 3 listings are for bulk sales of ecstasy: 2 for 10,000 pills, and 1 for 50,000 pills. Assuming approximately 200 mg per pill (this information is not available in the listings), this would represent 2–10 kg of ecstasy.
- 3 listings are "combinations," bundling cocaine (50–130 g) with ecstasy (2,000–7,500 pills), and in two cases with ketamine as well (100–130 g).
- 2 listings are bulk ketamine sales for 2 and 3 kg, respectively.
- 1 listing is for 5 kg of amphetamine "speed" paste.

20 of these listings have only one sale; three have two sales; one has three sales; and one has eight sales. In total we evaluate the revenue from all of these sales to approximately 619,000 euros—as we will see later, this is quite negligible when compared to the overall EU-originated sales on AlphaBay (approximately 1% of all EU sales).

Automated classification Our scrapes have collected data on a total of 101,999 items in total, which means that we need to perform automated classification as any manual processing would not scale to such sizes. As shown above, our algorithms fortunately appear to have good accuracy, and generally err on the side of producing conservative estimates, rather than inflated ones.

4 Data analysis

We next turn to analyzing the data we collected. Different from our previous report [8], in which we relied on previously reported data [17], we are using data here that has not been described before. Thus, we start with an overview of AlphaBay as a whole, before delving into European Union specifics.

4.1 Evolution of sales on AlphaBay

We start by a stacked plot, in Figure 1, representing the overall evolution of sales on AlphaBay during (most of) its history. Each vertical dotted line corresponds to the time we collected one scrape – we have 27 such scrapes, from March 2015 through May 2017. Each (stacked) curve represents a specific (set of) countries.

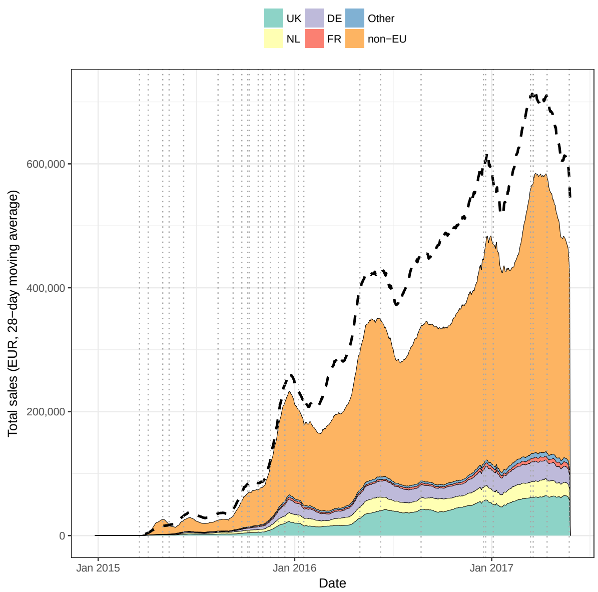

(a) Total sales over time (per-country breakdown)

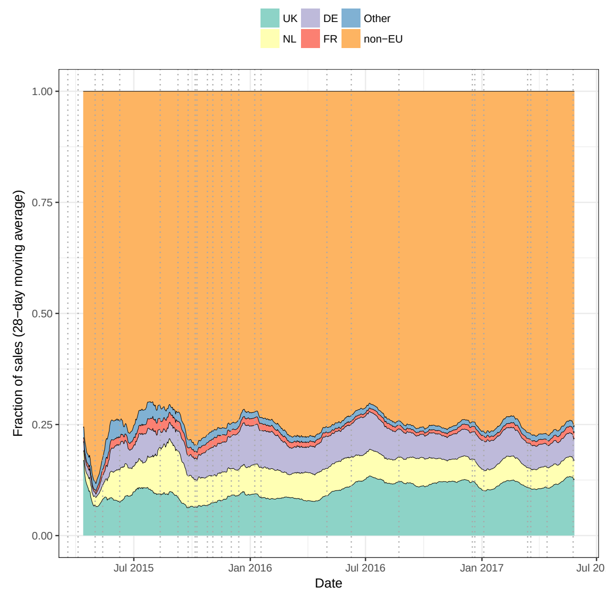

(b) Relative sales over time (per-category breakdown)

Figure 1: **Evolution of sales on AlphaBay over time (per-country breakdown).** The vertical dotted lines represent times at which we obtained a scrape of AlphaBay. The left-hand plot represents a breakdown per country. The dashed line is the total amount of sales. The white area represents sales for which we have a record, but for which we cannot recover the corresponding item listing, and thus cannot infer the country. The orange area are non-EU sales. The right-hand plot presents the same information, but on a relative scale, excluding items for which we do not have a corresponding listing.

The UK is in green, the Netherlands in yellow, Germany in purple, France in red, all other EU countries in blue, and non-EU countries in orange.

Figure 1(a) represents the evolution of sales in absolute value, over time, represented in euros, as a 28-day moving average. The thick dashed line denotes the *total* amount of sales. The discrepancy – materialized by a white area – between the total amount of sales and the sum of all sales from countries that could be identified is due to partially incomplete coverage. Specifically, we use item feedback to count the number of sales (assuming each piece of feedback corresponds to one sale, which, by and large, has shown to be a reasonable approximation in the past [17]). On AlphaBay, this feedback is displayed in two locations: each *item page* features a list of all pieces of feedback recorded for that item; each *user page* features a list of all pieces of feedback recorded for that user. Both sets should be consistent – that is, the set of all pieces of feedback for all users, and the set of all pieces of feedback for all items should be the same. However, scraping is rife with practical difficulties: a given scrape may take a couple of days to complete, and in the meantime certain item listings may be removed; the scrape may abnormally terminate before having visited every single page of the website, due to the site going down, or the scraper being logged out accidentally or voluntarily. As a result, it may happen that feedback gleaned from a user page is "orphaned," that is, does not have a corresponding item listing we can refer to. This orphaned feedback is characterized by the white area in Figure. 1(a): we know a vendor sold something for a given value on that day, but we do not necessarily have all information for the listing—in particular, we cannot be sure where the item reportedly ships from, and we do not have the listing description, but merely the listing title.[2] We could potentially infer the shipping origin by looking into shipping information for other items from the same vendor, but certain vendors operate from multiple countries at once; instead we prefer to adopt the conservative approach of marking these listings as "unidentified." We can see that the white area widens in times we have fewer scrapes per unit of time (e.g., mid-2016), which is consistent with our previous work [17], in which we discovered that perfect coverage often requires several scrapes to fill in the missing pieces. The small glitch in early 2015, in which the volume of non-EU sales appears slightly higher than the total volume is due to rounding errors in computing moving averages.

Overall, the plot shows that AlphaBay gained tremendous popularity toward the end of 2015, and then continued its steady climb pretty much until its demise. The apparent decline, at the end of collection interval, is probably an artifact of imperfect coverage: while, for earlier scrapes, incomplete data could be recovered in a subsequent scrape, such is not the case for the last scrape; this phenomenon is related to "censored data" in statistics, and fully explained in our previous paper [17]. Comparing these results with those obtained for the entire marketplace ecosystem in 2011–2015 [17], we can see that AlphaBay *itself* managed to be roughly comparable in size with the cumulative totals for all marketplaces in previous years; in 2017, it was roughly twice as large as the original Silk Road in its heyday (summer 2013).

The right-hand side plot in Figure 1(b) shows the same information, but expressed as a fraction of total sales. We see that all EU trade is roughly a quarter of the total trade we could identify, and that this remained roughly constant over time.

[2] There are also pieces of feedback gleaned from item pages, that do not have a user page associated with them. However, this is less of an issue, since the item page contains all the information needed to classify a specific listing.

Validating our numbers We can immediately use these findings to validate our analysis, by comparing it with established "ground truth." Our numbers are conservative: in particular, the spikes we see in the curves of Figure 1(a) and Figure 2(a) frequently correspond to times at which we launch a scrape, which indicates that we are likely to miss non-negligible amounts of data in-between scrapes. As we further discuss below, we are also most certainly underestimating revenue from digital goods transactions.

However, our numbers are also roughly consistent with estimates provided by other sources. The criminal complaint for forfeiture against Mr. Alexandre Cazes [3], who was alleged to be the AlphaBay founder and operator, says that "AlphaBay was the largest dark web market and its annual sales were estimated to be in the hundreds of millions of dollars," and subsequently estimates

> "[...] between May 2015 and February 2017, Bitcoin addresses associated with AlphaBay conducted approximately 4,023,480 transactions, receiving approximately 839,087 Bitcoin and sending approximately 838,976 Bitcoin. This equals approximately US$450 million in deposits to AlphaBay."

The estimates coming from our scrapes yield US $222,932,839 (and 2,223,992 transactions). We believe the $450 million dollar from the complaint might be an overestimate, due to currency mixing that might result in double-counting. Furthermore, each sale on AlphaBay can represent up to *two* Bitcoin transactions – one from the buyer to put money in their account, one from the seller to "cash out" once the item has been sold.

On the other hand, our own estimates are on the conservative side. In particular, 1) we ignored bulk items as discussed above, 2) we assume that each piece of feedback corresponds exactly to one sale (when sometimes one piece of feedback might actually corresponds to multiple quantities being part of the same sale), 3) we underestimate digital goods sales, due to a quirk in the way certain purveyors or credit card numbers do their business,[3] and 4) we do not have perfect coverage, especially in mid-2016.

Consequently, our numbers should provide a reasonably good, conservative estimate of AlphaBay trade. In particular, our estimates are likely to be far more accurate in times we have been aggressively scraping, most notably in 2015 and in early 2017.

Categories over time We next provide, in Figure 2 an different graphical representation of the overall sales volume, broken down, this time, by product categories. This plot shows worldwide sales. As before, the white area represents the transactions for which we could not identify a category, due to an orphaned piece of feedback[4] Figure 2(b) shows that AlphaBay did primarily start as a digital goods business, and then, when the (then-leading) Evolution marketplace went down in March 2015, started picking up narcotics traffic; after the Agora marketplace shut down in August 2015, the uptick in traffic is even more significant. Ultimately, the proportions we observe are roughly similar to those we had seen in the overall ecosystem in 2011–2015 [17, Fig. 7]: cannabis, cocaine, and MDMA/ecstasy represent approximately 2/3 of all trade.

[3]Many stolen credit card number vendors list their items in generic form, with a price of zero, instead leaving the specifics in the shipping costs—presumably to obfuscate their stocks and possibly to reduce the commissions imposed by the marketplace operator. For instance, a listing would be for "credit card dumps," with a price of zero, but with shipping options for various types of cards at various prices. Because we cannot determine which cards are purchased, we simply list ignore sales volumes from these listings.

[4]We could have attempted to classify these transactions based on the listing title, but would have probably run into inconsistencies with the item-description-based classification we perform for all items.

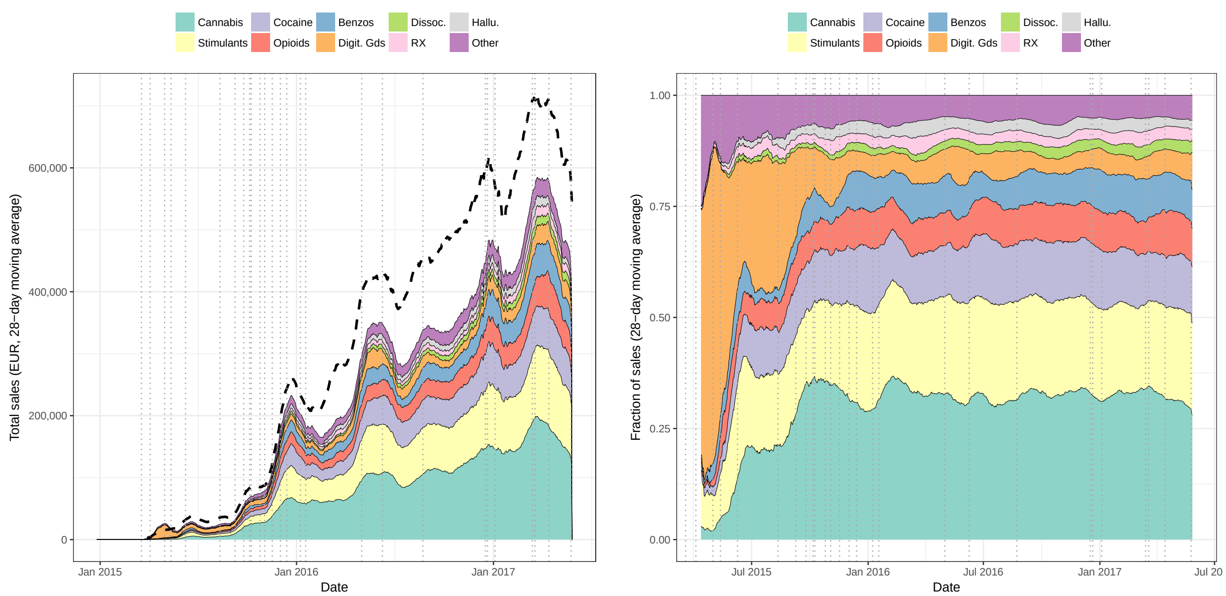

(a) Total sales over time (per-category breakdown)

(b) Relative sales over time (per-category breakdown)

Figure 2: **Evolution of sales on AlphaBay over time (per-category breakdown).** The left-hand plot represents a breakdown per category. The dashed line is the total amount of sales. The white area represents sales that we have a record of, but for which we cannot recover the corresponding item listing and thus cannot infer the category. The right-hand plot presents the same information, but on a relative scale (excluding items for which we do not have a corresponding listing). The vertical dotted lines represent times at which we obtained a scrape of AlphaBay.

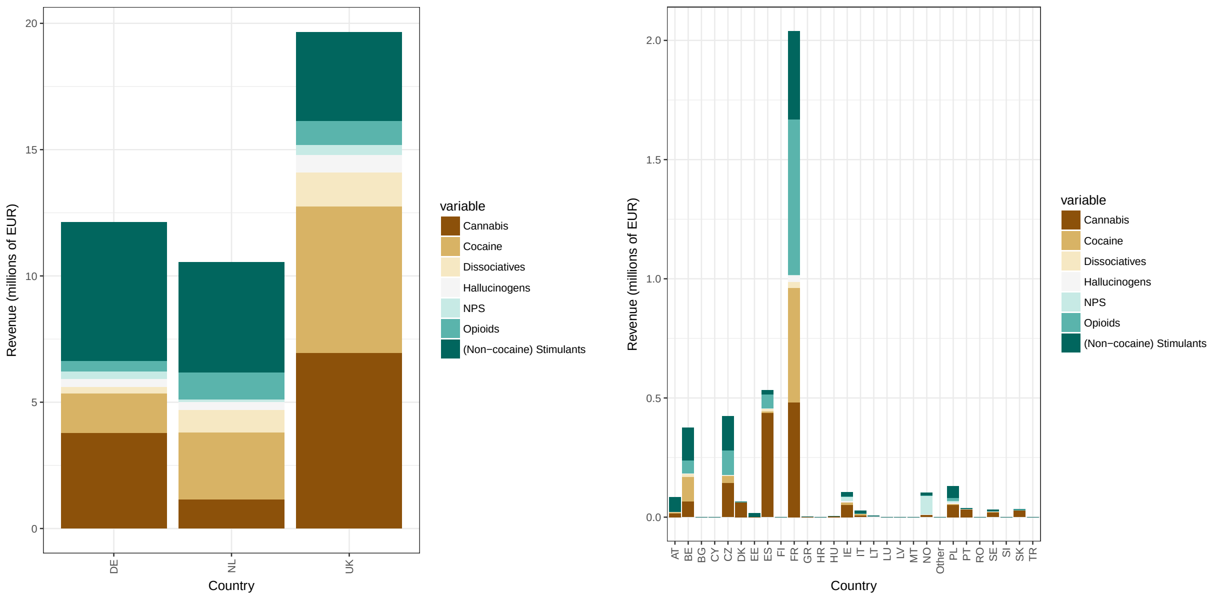

(a) Breakdown by revenue (major countries)

(b) Breakdown by revenue (other countries)

Figure 3: **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 (3(a)).

The share of benzodiazepines, and, to a lesser extent, of opioids, seems to have grown a bit compared to 2011–2015. On the other hand, prescription drugs went down (perhaps because some of them have been reclassified as opioids or benzodiazepines). The ecosystem appears relatively stable overall.

4.2 Sales from EU sellers

We next focus solely on the drugs originating from the European Union. For the seven categories of drugs of primary interest (see Section 3.2), Figures 3 and 4 presents a breakdown of sales originating from the European Union (plus Norway and Turkey) by country. Both plots are stacked plots. NPS are aggregated in a single category. Figure 3 represents the aggregate amount of transactions over our entire data collection interval (December 26, 2014–May 27, 2017).

Revenue analysis As we can see in Figure 3, as was the case for the marketplace ecosystem as a whole between 2011 and 2015 [8], the vast majority of sales originating from the EU comes from (the same) three countries: the United Kingdom, with approximately EUR 19.7 million total sales for the seven drug categories of interest – to be compared with the EUR 20.3 million that were sold between 2011–2015 on other marketplaces [8]; Germany, with EUR 12.1 million sales (down from 26.6 million for 2011–2015 [8]); and the Netherlands, with slightly more than EUR 10.6 million sales (compared to 17.9 million in the previous study [8]). Differences are more marked among "second-tier" countries: France (EUR 2.0M) is the only country with a gross revenue higher than EUR 1M; Spain (EUR 533K), the Czech Republic (EUR 424K), and Belgium (EUR 375K) are also noteworthy. By contrast, in our previous study, Belgium

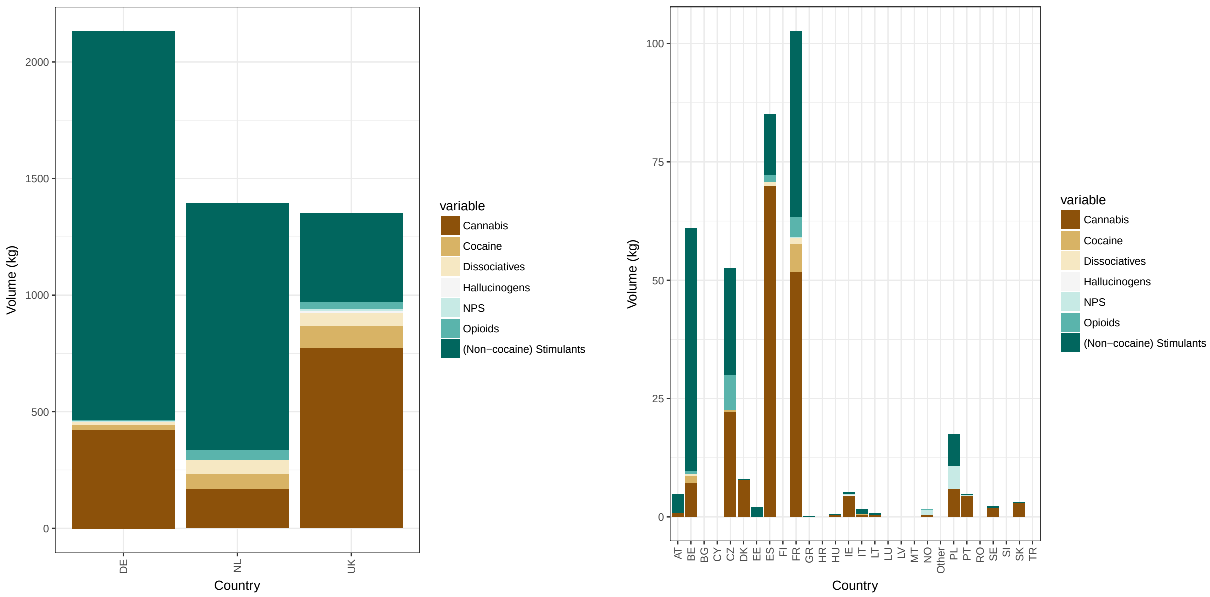

(a) Breakdown by volume (major countries)

(b) Breakdown by volume (other countries)

Figure 4: **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 (4(a)).

(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) were above EUR 1M. This may be due to the fact we are only looking at AlphaBay, as opposed to a collection of marketplaces. Whereas, in our previous study, some marketplaces (e.g., Utopia, Hydra) were seemingly more catering to Europeans, AlphaBay tends to be more international. It would be interesting, as future work, to look at revenues of European-based marketplaces such as Hansa or Valhalla, to see how they might change the picture.

We had discovered in our previous study that the top three countries were primarily selling stimulants other than cocaine, that is, MDMA, ecstasy and related products [8]. The situation appears to have changed a bit, with revenues that now appear more evenly distributed between cannabis, cocaine, and other stimulants; with a second tier (opioids, hallucinogens, and dissociatives) also roughly evenly distributed. The Netherlands appear to sell significantly less cannabis than other countries, and, proportionally, more cocaine and stimulants. France, our fourth country on the list, surprisingly to us seems to garner a fair amount of revenues in opioids. NPS overall remain quite small (in the order of EUR 100–300K for the leading countries), but again, it is quite possible a lot of NPS opiates are instead classified as opioids.

Volume analysis Figure 4 shows a breakdown by volume. Results are generally consistent with those of Figure 3: Germany (2,130 kg overall), the Netherlands (1,392 kg overall), and the United Kingdom (1,352 kg overall) dominate the ecosystem; these are the only countries where products shipped exceed, in aggregate, a metric ton. Interestingly, the UK generates more revenue with less volume. In particular, the volume of stimulants from the UK is much smaller than to those sold in Germany and the Netherlands, yet

Table 1: **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 sales[5]		**Total sales**
	Volume (g)	Revenue (EUR)	Revenue (EUR)
European Union (plus Norway, Turkey)	5,228,918	46,363,102	51,400,357
Rest of the world	16,341,359	116,631,207	152,751,443
Total	21,570,277	162,994,309	204,151,800

revenue is only slightly less. By manually inspecting our results, we discovered that while Germany and the Netherlands primarily cater MDMA and ecstasy pills, the UK tends to sell more stimulants such as Adderall and Ritalin, whose dosages are far smaller than MDMA and ecstasy, at a similar price.

Comparison with non-EU sales Table 1 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. First, we notice that drug sales are an overwhelming majority (>90%) of all EU sales. This is less so the case for the rest of the world (≈76%), which makes sense: any item whose origin cannot be established with certainty will be marked as "unknown," i.e., outside of the EU. This includes, in particular, a large number of digital goods (e-books, credit card numbers, stolen credentials, ...) whose origin is listed as unknown, even a number of these sales probably originate from Europe. This finding is consistent with our previous report [8].

On the other hand, in our previous report, we noted that EU countries represented roughly 46% of all revenue, but only 34% of all volumes, for 2011–2015 [8]. The percentages, here, are markedly smaller: 28.4% of all revenue, and 24.4% of all volumes. There are many possible explanations, including:

- Following the arrests of prominent EU vendors (SuperTrips, Shiny Flakes...) EU traffic on anonymous marketplaces did overall decrease compared to the rest of the world; or
- A larger portion of vendors is not labeling their origin country properly; in fact we noticed that 26,316 items in our corpus (i.e., slightly above a quarter) list "worldwide" as their origin—only 12,170 of these items were classified as digital goods; or
- A number of prominent EU vendors moved to marketplaces seemingly catering more to Europeans (e.g., Hansa, Valhalla, francophone forums, ...) or to their own, single-vendor, shops.

[5] In the seven categories of interest.

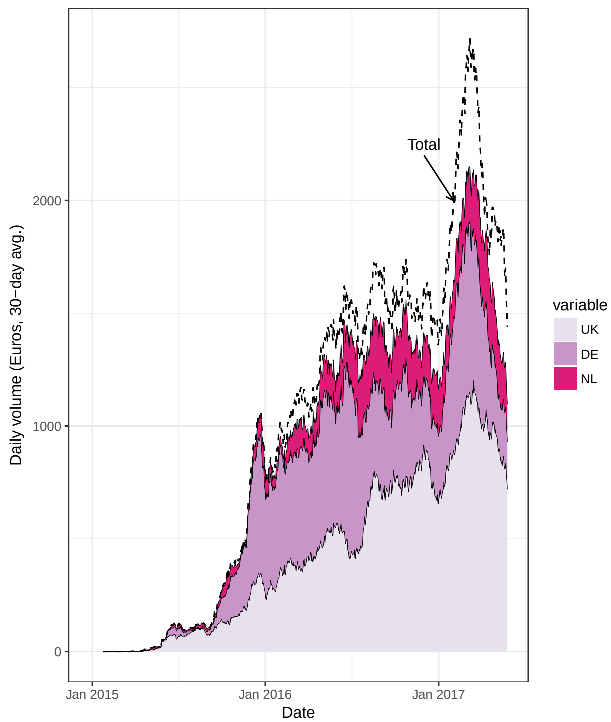

Figure 5: **Breakdown of NPS sales originating from the European Union (plus Norway and Turkey).**

4.3 New Psychoactive Substances

In an effort to compare data from AlphaBay with that from older marketplaces, we next focus on New Psychoactive Substances, or NPS. Our previous measurements [8] showed that NPS accounted for a very small fraction of all narcotics traffic on online anonymous marketplaces – in the order of EUR 3,000/day at most. With reports of "opioid epidemics," it stands to reason that the amount of NPS being sold online would also increase.

Figure 5 provides a finer-grained view of the sales of NPS on online anonymous marketplaces, over time, as a stacked plot. The high-level takeaway is that these sales remain modest, and in-line with what we had previously observed in the 2011–2015 interval [8]. NPS originate primarily from the UK and Germany, with the Netherlands a distant third, and all other EU countries accounting for negligible volumes. We emphasize again that many opioids (e.g., fentanyl variants) are actually labeled as opioids by our classifier – rather than NPS, which certainly causes this graph to be an underestimate. However, when it comes to "legal highs," manual inspection confirms that online anonymous marketplaces are certainly not a hub of activity. Most NPS sold on AlphaBay (and classified as such) appear to be hallucinogens.

4.4 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, and compare our findings with those of our previous report [8].

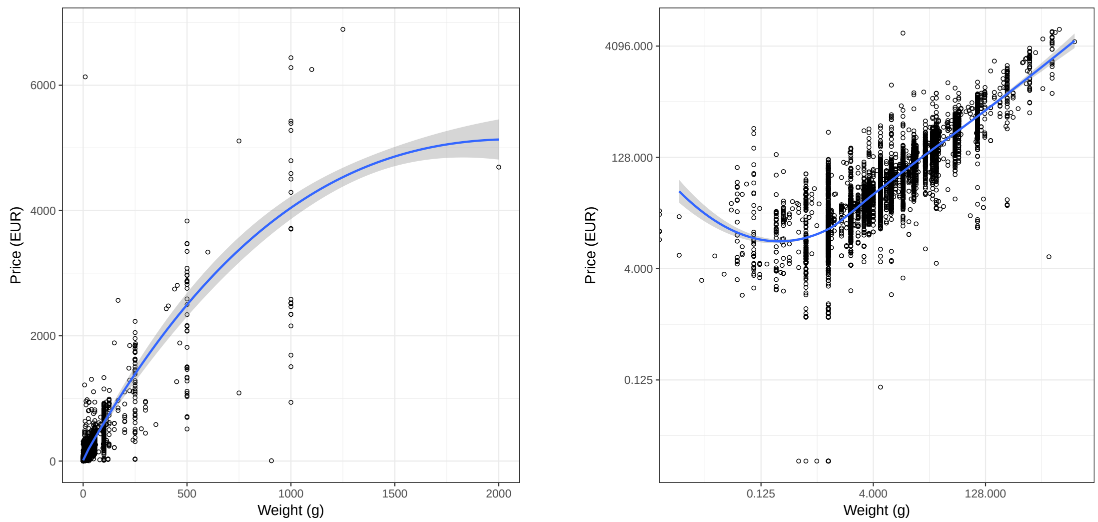

(a) Linear scale

(b) Logarithmic scale

Figure 6: **Cannabis prices as a function of weight.** The curve is a local polynomial regression fitting, the gray shade corresponds to the 95% confidence interval.

Cannabis Figure 6 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 6(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 6(b)). The blue curve corresponds to a non-parametric regression (using a local polynomial regression fitting, [10]), 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 results closely mirror those of our previous report [8, Fig. 7].

The most common units sold are 1g (1,239 items, mean price EUR 12, standard deviation EUR 20), 5g (976 items, mean price EUR 49, standard deviation EUR 28), and 10g (910 items, mean price EUR 86, standard deviation EUR 204). All of these values are in line with what we had previously observed.[6]

As in the previous report, the high standard deviations can be explained by the large dispersion due to a range of different products (oils, edibles, etc) being classified as cannabis; and as before, we note the presence of a few items with very small volumes (close to zero gram); some of these are sample offers, some (particularly those with high cost) are the few items for which our heuristics for extracting quantities have failed.

Cocaine and Other Stimulants Figure 7 shows similar scatter plots for cocaine products. Figure 7(a) uses a linear scale, and Figure 7(b), a logarithmic scale. We use again a local polynomial regression, which

[6]For comparison purposes, we had found 1,610 1-gram listings (μ = EUR 17, σ = EUR 16); 1,745 5-gram listings (μ = EUR 58, σ = EUR 39); and 1,165 10-gram listings (μ = EUR 99, σ = EUR 55) in our study of 2011–2015 data [8].

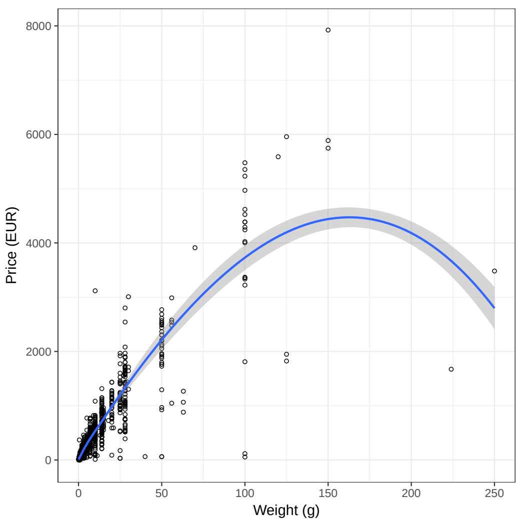

(a) Cocaine (Linear scale)

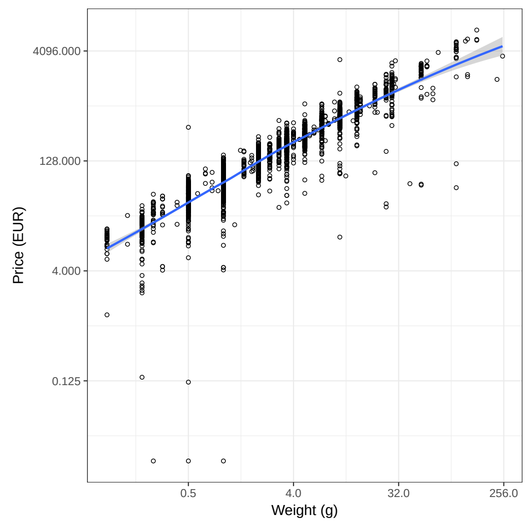

(b) Cocaine (Logarithmic scale)

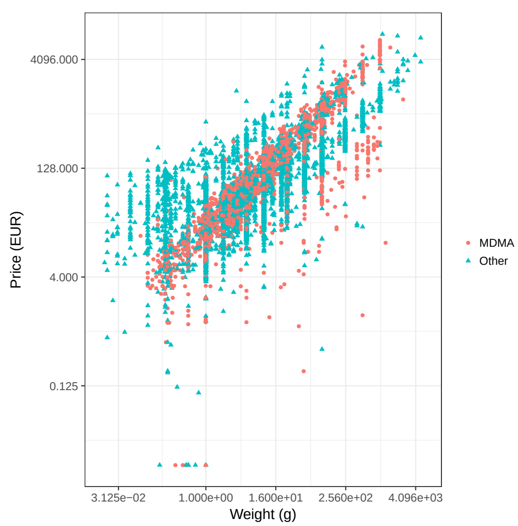

(c) Stimulants other than cocaine

Figure 7: **Cocaine and other stimulant prices as a function of weight.** The curve in 7(a) and 7(b) is a local polynomial regression fitting, the gray shade corresponds to the 95% confidence interval.

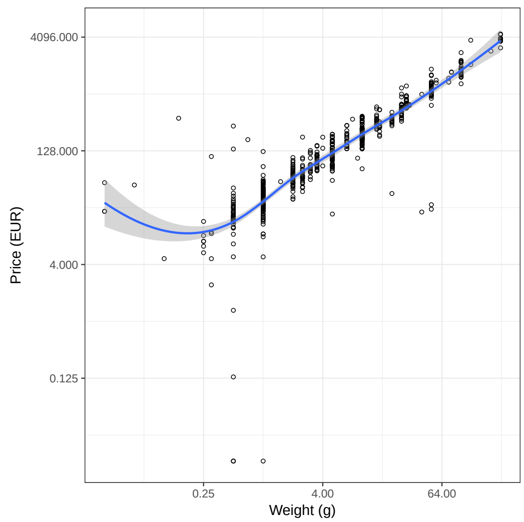

Figure 8: **Dissociative (Ketamine) prices as a function of weight.** The curve are local polynomial regression fittings, the gray shade corresponds to the 95% confidence interval. Note the different estimators depending on the type of dissociative sold. We removed the (separate) plots for GHB and other dissociatives (PCP), as there were too few datapoints.

shows that volume discounting effect is markedly more pronounced here than it was in the case of cannabis products; note however that the single outlier bringing down the curve on the right is not a misclassification *per se* – this item is a set of ground coca *leaves*, that need additional purifying and processing, and is thus considerably cheaper per unit than actual cocaine. This graph is very similar to what we observed in our previous report [8, Fig. 8].

For cocaine, the most common unit sold is 1g (626 items, mean price EUR 68, standard deviation EUR 22).

Stimulants (other than cocaine), in Figure 7(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. The weight is defined 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.

Different from our previous report, which showed extreme dispersion at low volumes for stimulants other than cocaine and MDMA, we also see a near linear relationship for these products. This is primarily due to additional data processing and manual labeling of a few egregiously wrong quantities (see Section 3.2).

Dissociatives We next turn to dissociatives. Figure 8 provides a similar scatter plot in log-log scale for ketamine. Consistent with our previous report [8, Fig. 9], 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.

The most common unit sold (for Ketamine) is 1g (158 items, mean price EUR 30, standard deviation

EUR 13).

Other drugs: Hallucinogens, NPS (Hallucinogens), and Opioids We next turn to the other three categories of drugs of interest, depicted in Figure 9, in which all scatter plots are using a log-log scale. For Hallucinogens and NPS, the results are strikingly similar to those in our previous report [8, Fig. 10]. Hallucinogens, i.e., primarily LSD, represented in Figure 9(a), show two clear clusters: there is a large price variance for low, single use, type of doses (250 μg and less). As in the previous report, this appears to 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 9(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 9(c), on the other hand show far less price dispersion than we had observed in our previous research. The regression, using again a local polynomial regression fitting suggests two modes: below one gram, where prices increase modestly (dispersion due to quality heterogeneity seems to dominate), and above one gram, where we observe modest volume discounting; for high volumes ($>$16 g), the scarcity of data makes the regression less precise.

4.5 Vendor diversification

Following the same methodology as in our companion report [8], 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.

As before, we will use the coefficient of diversity defined by Soska and Christin [17], whose definition we summarize here. 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 $\forall s_j \in S, \ \sum_{i=1}^{|\mathcal{C}|} \mathcal{C}_i(s_j) = 1$. 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:

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

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

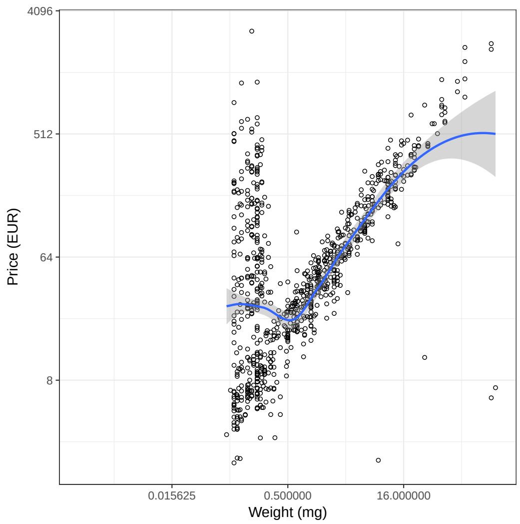

(a) Hallucinogens

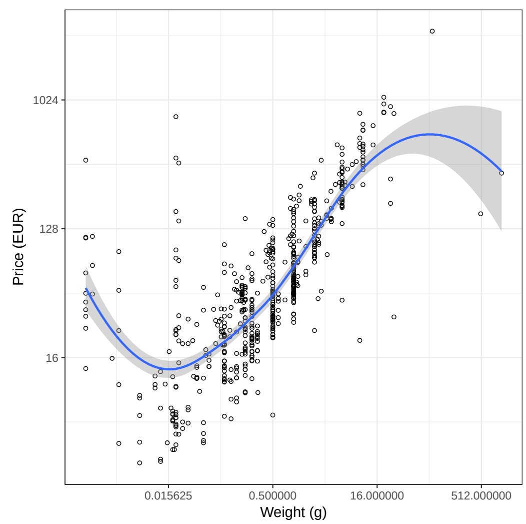

(b) NPS (Hallucinogens)

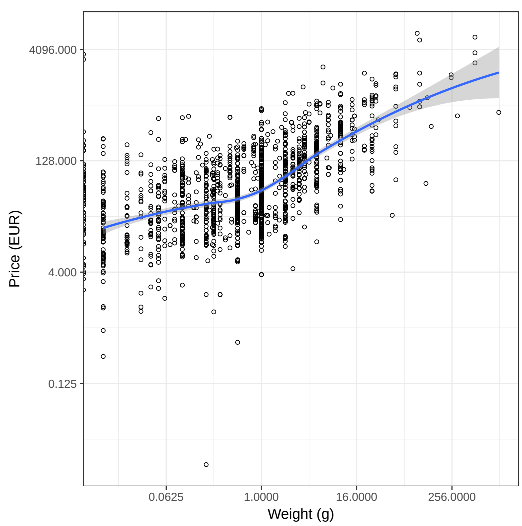

(c) Opioids

Figure 9: **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.

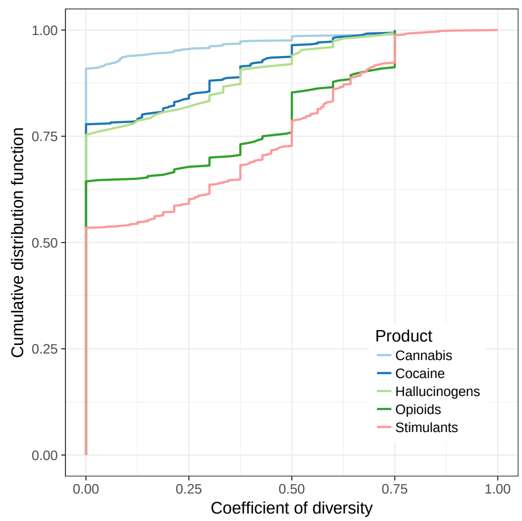

Figure 10: **Coefficient of diversity for vendors, by product, across volume tiers.**

Table 2: **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	≥1000g	≥1000 g	≥160 mg	≥1000 g	≥1000 g

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 2. We use a simple three-tier distinction between retail, middle-market, and bulk sales. We base our distinction between tiers on the EMCDDA's own classification [13]. 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 2, we plot the cumulative distribution function of the coefficient of diversity of their vendors in Figure 10. Results are nearly identical to what we had observed in the 2011–2015 study [8, Fig. 11]. That is, the figure shows that an overwhelming (≈90%) majority of cannabis vendors stay within one volume tier (coefficient of diversity of zero). The rest

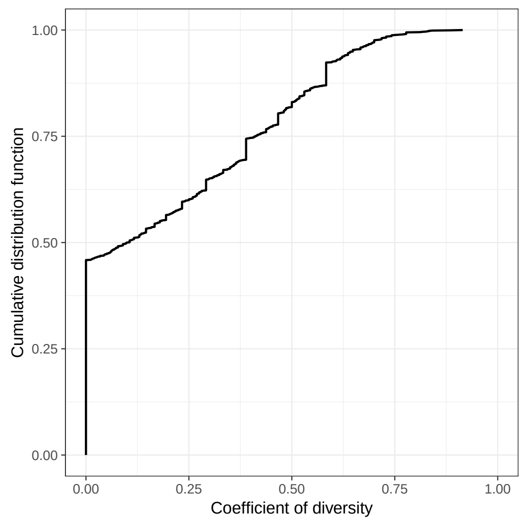

Figure 11: **Coefficient of diversity for vendors, across product categories.**

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 within 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. Likewise, hallucinogens present almost the same behavior as they did in our previous measurements.

Cocaine and opioid present more diversity than they did in the previous study. In particular, we observe that a number of sellers sell across tiers, some giving out free (or cheap) samples, and are also selling in far larger quantities.

As in our previous study [8], stimulants other than cocaine present the most diversity, with most vendors sticking to the retail tier, but some vendors selling across multiple tiers with a number of items in each tier. This can be explained by the ease of shipping fairly large quantities of pills stealthily.

We checked a few instances in our database that suggest that vendors selling in multiple echelons tend to be "superstores," 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. In other words, vendor behavior, overall, has not changed much compared to our previous measurements [8].

Diversification in terms of products offered In all, 1,956 AlphaBay vendors reportedly ship from the European Union. 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 11 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 (329 cases) and stimulants other than cocaine (225) vendors, which is not very surprising given that those are very frequently sold items. On the other hand, only 59 vendors purely specialize in opioids. The other half is far more diverse, and typically those vendors sell from a couple of categories. For instance, vendors selling dissociatives, hallucinogens, or NPS 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.

Out of the 1,956 EU vendors, 1,321 sell drugs in one of the seven categories of interest. Of those 1,321 vendors, 171 (or 12.94%) also sell other types of drugs (e.g., prescription drugs). 464 (35.12%) also sell non-drug products: we found this diversity surprising in our previous study [8], but the consistency in results shows it was not a fluke. Of those, the majority appears to be digital goods, and sold in very small quantities – as complementary income of sorts.

133 of those 464 vendors that sell non-drug products remain confined to one drug category (primarily cannabis, and stimulants other than cocaine). 93 of those 464 vendors have on the other hand high diversity coefficients (>0.5). Many of these vendors sell cannabis, stimulants (sometimes including cocaine), and opioids.

5 Conclusions

We have presented an analysis of the AlphaBay marketplace, through most of its 2.5-year history. For the sake of comparison, most of the analysis attempted to reproduce what we had done for the online anonymous marketplace ecosystem as a whole in 2011–2015 [8, 17]. We first found that AlphaBay did become a major player in the marketplace ecosystem, ultimately exceeding the aggregate sales of all other marketplaces we had monitored previously, and even exceeding – on its own – the amount of daily sales from the whole ecosystem. We emphasize that our data is not fully complete and likely to produce (slight) under-estimates of total revenues and volume; however it does not cause bias toward one specific item or category.

Focusing on products reportedly shipping from the European Union, we found a number of similarities between AlphaBay, and the marketplaces we had previously considered [8]: EU-originating drugs still primarily came from Germany, the Netherlands, and the United Kingdom. Many countries did not register any significant sales.

Likewise, 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.

Further, “recreational drugs,” that is cannabis, cocaine and other stimulants (amphetamines, MDMA, ecstasy,...), cocaine altogether continue to represent a majority of all EU-based sales. Opioids and dissociatives formed a next tier; hallucinogens and NPS were more modest. As we had noted in our previous study [8], NPS volumes remained comparatively very modest – with revenues below EUR 2,500 at market peak.

Similarly, the vendor ecosystem remains 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. These results were remarkably close to those of our previous study, hinting some overall stability.

In terms of differences between this report and our previous study, the total share of EU-based supplier seems to have decreased – roughly accounting for a quarter of all revenue and a quarter of all volumes, as opposed to approximately 45% in the previous study.

Acknowledgment

This work greatly benefited from extensive discussions with Kyle Soska and Jeremy Thomas.

References

[1] I2P: The internet invisible project. `http://www.geti2p.net`.

[2] United States of America vs. Alexandre Cazes, June 2017. United States District Court, Eastern District of California. Indictment 1:17CR-0144-LJO-SKO.

[3] United States of America vs. Alexandre Cazes – verified compaint for foreiture in rem, July 2017. United States District Court, Eastern District of California. Case 1:17-at-00557.

[4] Judith Aldridge and David Décary-Hétu. Hidden Wholesale: The drug diffusing capacity of online drug cryptomarkets. *International Journal of Drug Policy*, 35:7–15, 2016.

[5] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society, 2014.

[6] A. Chen. The underground website where you can buy any drug imaginable. `http://gawker.com/5805928/the-underground-website-where-you-can-buy-any-drug-imaginable`. Last accessed May 1, 2012.

[7] 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.

[8] Nicolas Christin. An EU-focused analysis of drug supply on the online anonymous marketplace ecosystem, 2016. EMCDDA report for contract CC.16.SAT.0050.1.0.

[9] Nicolas Christin. An EU-focused analysis of drug supply on the alphabay marketplace, 2017. EMCDDA report for contract CT.17.SAT.0063.1.0.

[10] William S Cleveland, Eric Grosse, and William M Shyu. Local regression models. *Statistical models in S*, 2:309–376, 1992.

[11] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.

[12] Diana Dolliver. Evaluating drug trafficking on the Tor network: Silk Road 2, the sequel. *International Journal of Drug Policy*, 2015.

[13] European Monitoring Centre for Drug and Drug Addition. EMCDDA drug prices indicator, protocol for reporting price data, 2016.

[14] A. Greenberg. An interview with a digital drug lord: The Silk Road's Dread Pirate Roberts (Q&A), August 2013. `http://www.forbes.com/sites/andygreenberg/2013/08/14/an-interview-with-a-digital-drug-lord-the-silk-roads-dread-pirate-roberts-qa/`.

[15] James Martin. *Drugs on the dark net: How cryptomarkets are transforming the global trade in illicit drugs*. Springer, 2014.

[16] S. Nakamoto. Bitcoin: a peer-to-peer electronic cash system, October 2008. Available from `http://bitcoin.org/bitcoin.pdf`.

[17] 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.

[18] U.S. Attorney's Office, Southern District of New York. Dozens of online "dark markets" seized pursuant to forfeiture complaint filed in Manhattan federal court in conjunction with the arrest of the operator of Silk Road 2.0, November 2014. `http://www.justice.gov/usao/nys/pressreleases/November14/DarkMarketTakedown.php`.

[19] Nicolas van Saberhagen. Cryptonote v2.0. `https://cryptonote.org/whitepaper.pdf`, 2013.

A Per-tier sales amounts (AlphaBay and other marketplaces)

This appendix provides five tables describing the various amounts of sales in each sales tier (retail, middle-market, and bulk) defined by EMCDDA, as introduced in the associated reports [8, Table 3], [9, Table 3]. Table 3 provides, for six families of drugs (Cannabis, Cocaine, Opioids, MDMA, Stimulants other than cocaine and MDMA, and Hallucinogens) the total weight of drugs sold in each tier, the number of corresponding transactions, and the associated monetary revenue. All items priced USD 10,000 or more are excluded from this computation per the discussion in our previous reports. Table 3 provides these measures for the AlphaBay marketplace. We are not able to identify volumes for every single transaction, and that sporadic errors might result from our automated treatment – all limitations described in the main report [9] apply. We noticed a few spurious errors in bulk sales of hallucinogens (due to vendors, for instance, erroneously labeling micrograms as grams), which we manually corrected. As sanity check, we manually inspected all bulk sales of hallucinogens, opioids and cocaine, and confirmed they are all correct.

Table 4 provides the same information as Table 3, but this time *including* items priced at more than USD 10,000 we manually inspected [9, Section 3.3]. The only differences are for bulk sales of cocaine, bulk sales of MDMA, and bulk sales of stimulants.

Table 5 shows the same information as Table 3, but this time for the marketplaces studied in our previous report [8] rather than AlphaBay. This table excludes items priced at USD 10,000 or more; we performed manual correction and validation of most bulk items.

Table 6 expands on the results described in Table 5, this time including items priced USD 10,000 or more.

Finally, Table 7 combines the results of Tables 4 and 6 by simply adding all quantities to give a holistic picture of the entire marketplace ecosystem. Again, all limitations and caveats discussed in our companion reports [8, 9] apply. In particular, it is likely that the numbers reported are *under*estimates of the actual volumes, given the conservativeness of our computations.

Table 3: **Breakdown of sales by market tier for AlphaBay.** We use the tiers provided by EMCDDA (see Table 3 in the associated report). This table only represents a subset of all transactions – those for which we can identify a weight and quantity. While we did our best to manually remove egregious outliers that were clearly incorrect, parsing errors may introduce some errors in the weight computations. This table does **not** include items which sold for more than USD 10,000.

Product	Tier	Weight (g)	Nr. Transactions	Revenue (EUR)
Cannabis	Retail	1,361,345	167,802	10,238,221
	Middle-market	149,961	617	575,700
	Bulk	31,350	30	111,383
Cocaine	Retail	135,758	71,345	8,284,447
	Middle-market	53,024	1908	1,830,040
	Bulk	0	0	0
Opioids	Retail	17,518	29088	1,208,555
	Middle-market	70,716	10,797	1,612,659
	Bulk	0	0	0
MDMA	Retail	160,349	46,882	1,799,377
	Middle-market	1,000,627	13,965	3,927,724
	Bulk	68,250	66	271,025
Other Stimulants	Retail	199,929	58,081	2,414,878
	Middle-market	1,335,925	17,037	3,022,693
	Bulk	486,050	413	591,465
Hallucinogens	Retail	15,818,797 ($\times 10^{-6}$)	22,226	856,072
	Middle-market	12,055,450 ($\times 10^{-6}$)	726	187,713
	Bulk	440,000 ($\times 10^{-6}$)	2	4,491

Table 4: **Breakdown of sales by market tier for AlphaBay.** We use the tiers provided by EMCDDA (see Table 3 in the associated report). This table only represents a subset of all transactions – those for which we can identify a weight and quantity. While we did our best to manually remove egregious outliers that were clearly incorrect, parsing errors may introduce some errors in the weight computations. This table **does** include items which sold for more than USD 10,000.

Product	Tier	Weight (g)	Nr. Transactions	Revenue (EUR)
	Retail	1,361,345	167,802	10,238,221
Cannabis	Middle-market	149,961	617	575,700
	Bulk	31,350	30	111,383
	Retail	135,758	71,345	8,284,447
Cocaine	Middle-market	54,149	1913	1,883,479
	Bulk	3,000	3	89,952
	Retail	17,518	29088	1,208,555
Opioids	Middle-market	70,716	10,797	1,612,659
	Bulk	0	0	0
	Retail	160,349	46,882	1,799,377
MDMA	Middle-market	1,000,627	13,965	3,927,724
	Bulk	84,000	72	361,012
	Retail	199,929	58,081	2,414,878
Other Stimulants	Middle-market	1,335,925	17,037	3,022,693
	Bulk	491,050	414	600,377
	Retail	15,818,797 ($\times 10^{-6}$)	22,226	856,072
Hallucinogens	Middle-market	12,055,450 ($\times 10^{-6}$)	726	187,713
	Bulk	440,000 ($\times 10^{-6}$)	2	4,491

Table 5: **Breakdown of sales by market tier for the other markets.** We use the tiers provided by EMCDDA (see Table 3 in the associated report). This table only represents a subset of all transactions – those for which we can identify a weight and quantity. While we did our best to manually remove egregious outliers that were clearly incorrect, parsing errors may introduce some errors in the weight computations. This table does **not** include items which sold for more than USD 10,000. * manual inspection did confirm the item was listed at a very small price despite the bulk quantity specified. It is not clear whether this was due to an (unlikely) parsing error, or if, on the other hand, the seller used this low price as a signaling mechanism to engage potential customers in paying out-of-band (more likely), e.g., to avoid marketplace commission fees.

Product	**Tier**	**Weight (g)**	**Nr. Transactions**	**Revenue (EUR)**
Cannabis	Retail	1,514,386	178,617	12,149,478
	Middle-market	418,985	1454	1,628,581
	Bulk	226,000	202	612,144
Cocaine	Retail	120,220	61,340	8,302,118
	Middle-market	39,882	1,230	1,449,678
	Bulk	2,000	2	0.01*
Opioids	Retail	18,299	36,227	2,007,932
	Middle-market	120,583	8,467	1,572,175
	Bulk	10,000	3	7,487
MDMA	Retail	221,627	72,873	5,222,583
	Middle-market	1,795,668	21,563	10,981,699
	Bulk	90,100	51	69,338
Other Stimulants	Retail	180,106	65,719	3,269,786
	Middle-market	946,671	12,465	3,768,916
	Bulk	447,900	392	595,407
Hallucinogens	Retail	23,418,830 ($\times 10^{-6}$)	24,616	1,767,312
	Middle-market	22,038,050 ($\times 10^{-6}$)	997	405,144
	Bulk	225,000 ($\times 10^{-6}$)	1	5,006

Table 6: **Breakdown of sales by market tier for the other markets.** We use the tiers provided by EMCDDA (see Table 3 in the associated report). This table only represents a subset of all transactions – those for which we can identify a weight and quantity. While we did our best to manually remove egregious outliers that were clearly incorrect, parsing errors may introduce some errors in the weight computations. This table **does** include items which sold for more than USD 10,000.

Product	Tier	Weight (g)	Nr. Transactions	Revenue (EUR)
Cannabis	Retail	1,514,386	178,617	12,149,478
	Middle-market	418,985	1454	1,628,581
	Bulk	227,000	203	619,080
Cocaine	Retail	120,220	61,340	8,302,118
	Middle-market	39,882	1,230	1,449,678
	Bulk	3,000	3	8,426
Opioids	Retail	18,299	36,227	2,007,932
	Middle-market	120,583	8,467	1,572,175
	Bulk	11,000	4	20,653
MDMA	Retail	221,627	72,873	5,222,583
	Middle-market	1,795,668	21,563	10,981,699
	Bulk	125,200	80	384,077
Other Stimulants	Retail	180,106	65,719	3,269,786
	Middle-market	946,671	12,465	3,768,916
	Bulk	447,900	392	595,407
Hallucinogens	Retail	23,418,830 ($\times 10^{-6}$)	24,616	1,767,312
	Middle-market	22,038,050 ($\times 10^{-6}$)	997	405,144
	Bulk	317,500 ($\times 10^{-6}$)	1	30,211

Table 7: **Breakdown of sales by market tier for all markets.** We use the tiers provided by EMCDDA (see Table 3 in the associated report). This table only represents a subset of all transactions – those for which we can identify a weight and quantity. While we did our best to manually remove egregious outliers that were clearly incorrect, parsing errors may introduce some errors in the weight computations. This table **does** include items which sold for more than USD 10,000.

Product	**Tier**	**Weight (g)**	**Nr. Transactions**	**Revenue (EUR)**
	Retail	2,875,732	346,419	22,387,700
Cannabis	Middle-market	568,946	2,071	2,204,282
	Bulk	258,350	233	730,464
	Retail	255,978	132,685	16,586,565
Cocaine	Middle-market	94,031	3,143 3,333,158	
	Bulk	6,000	6	98,378
	Retail	35,817	65,315	3,216,487
Opioids	Middle-market	191,299	19,264	3,184,834
	Bulk	11,000	4	20,653
	Retail	381,976	119,755	7,021,960
MDMA	Middle-market	2,796,294	35,528	14,909,423
	Bulk	209,200	152	745,090
	Retail	380,036	123,802	5,684,664
Other Stimulants	Middle-market	2,282,596	29,502	6,791,609
	Bulk	938,950	806	1,195,784
	Retail	39,237,627 ($\times 10^{-6}$)	46,842	2,623,384
Hallucinogens	Middle-market	34,093,500 ($\times 10^{-6}$)	1,723	592,858
	Bulk	757,500 ($\times 10^{-6}$)	4	34,702

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 100 academic research papers in computer networks and security. Major recent publications include:[7]

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 280 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 280 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 230 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. (Cited 80 times in other academic publications. The most comprehensive, to-date, study of online anonymous markets.)

[7] Different 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.