Upside Down: Exploring the Ecosystem of Dark Web Data Markets

Bogdan Covrig[1], Enrique Barrueco Mikelarena[1], Constanta Rosca[1], Catalina Goanta[2], Gerasimos Spanakis[1], and Apostolis Zarras[3]

[1] Maastricht University
[2] Utrecht University
[3] Delft University of Technology

Abstract Large-scale dark web marketplaces have been around for more than a decade. So far, academic research has mainly focused on drug and hacking-related offers. However, data markets remain understudied, especially given their volatile nature and distinct characteristics based on shifting iterations. In this paper, we perform a large-scale study on dark web data markets. We first characterize data markets by using an innovative theoretical legal taxonomy based on the Council of Europe's Cybercrime Convention and its implementation in Dutch law. The recent Covid-19 pandemic showed that cybercrime has become more prevalent with the increase of digitalization in society. In this context, important questions arise regarding how cybercrime harms are determined, measured, and prioritized. We propose a determination of harm based on criminal law qualifications and sanctions. We also address the empirical question of what the economic activity on data markets looks like nowadays by performing a comprehensive measurement of digital goods based on an original dataset scraped from twelve marketplaces consisting of approximately 28,000 offers from 642 vendors. The resulting analysis combines insights from the theoretical legal framework and the results of the measurement study. To our knowledge, this is the first study to combine these two elements systematically.

1 Introduction

The rise of cryptocurrencies led to a flourishing marketplace environment on the dark web [5]. Since the first pioneering marketplace (i.e., Silkroad) has started using a technological customer infrastructure involving Tor, Escrow payments, and Bitcoin, hundreds of other platforms have followed suit. Once infamous for the drug trade, dark web marketplaces have also gradually become ecosystems to monetize unlawfully obtained data, ranging from stolen data to data dumps or blatant scams. The value of data traded like this is estimated to be USD 1.5 trillion annually, and it includes personal, corporate, and financial data [26]. While law enforcement has been increasingly active in taking down markets and vendors offering for sale drugs, guns, and other traditionally illicit goods, data markets remain underexplored.

Given their volatile nature and scale, dark web marketplaces are inherently challenging to investigate [3]. The adversarial ecosystem of dark marketplaces forces these environments to adapt constantly. This means business models, infrastructure, and market features (e.g., reputation systems) often do not have time to mature, and that new market iterations mushroom at a breakneck pace, often with different characteristics. Each year of activity may shape a different picture in this space, so it is essential to constantly generate new datasets, as most studies in this field often rely on older data. However, using empirical methods for automated data collection makes it easier to compare activity across marketplaces and understand how market structures evolve.

Data markets are hard to define, as the variety of things available for sale makes it difficult to categorize what *data* may be. For instance, data can range from personal information, aggregated public information, and software to online services. Simultaneously, while certain practices may develop around the virtual spaces where data is offered for sale, a vast spectrum of environments is used by vendors or providers, including dark web marketplaces, Tor forums, or private conversations. So far, most investigations on data markets are either tech journalism reports addressing individual incidents or investigations by cybersecurity companies, which often lack transparent and reproducible methodologies [10]; it is often unclear how these studies calculate their estimations, as they typically use sources from the web and not analyses of original data obtained from *.onion* pages. In addition, scientific studies focus on classifying hacking products [15] or the commoditization of cybercrime with particular emphasis on *Business-to-Business* (B2B) and *Business-to-Consumer* (B2C) transactions [22].

In this paper, we contribute to the existing literature on dark web measurement in two significant ways. First, we answer the question of how cybercrime harms on dark web data markets can be classified according to international standards of criminal law, as well as their incorporation in the Dutch legal system. Second, we explore data market economic activity by performing a comprehensive measurement of digital goods sold on twelve marketplaces.

In summary, we make the following main contributions:

- We build on earlier research and critically reflect on the criteria used to define and categorize data offers available on dark web marketplaces, including those proposed on the marketplaces we investigate.
- We propose a novel legal theoretical framework for the criminal qualification of economic activities on dark web data markets, including their punishment, as the foundation of further comparative law research that can complement web measurement methodologies.
- We report on economic indicators relating to offers, commissions, and vendors to describe the economic activity of the data markets in our original large-scale dataset, and we compare these results to earlier findings of studies focused on the commoditization of cybercrime.
- We discuss potential new features of dark web marketplaces based on their current iteration (e.g., the platformization or standardization of dark web marketplaces).

2 Crimes on Data Markets: A Legal Framework

The interest in dark web marketplaces, in general, and in data-related transactions, in particular, has been on the rise. However, even when zooming into the latter, the vast majority of these studies have a particular economic focus [11,14]. Such markets rely on and amplify the intermediation of cybercrime, and given the complex web of applicable rules and resulting legal uncertainty [18], expert legal analyses have not been used so far to classify dark web marketplace activities for the purpose of measurement studies. In this section, we endeavor to define data markets and qualify data products and services as crimes under international and Dutch law.

2.1 Defining Data Markets

Web marketplaces are virtual spaces connecting vendors with buyers. On the dark web, transactional environments are formed (*i*) randomly or structurally around communication networks such as forums (e.g., threads designed explicitly for transactions), or (*ii*) structurally around platforms that systematize transactions both for vendors and buyers (e.g., e-commerce platforms).

Data markets on the regular web are often categorized based on the object of the transaction (e.g., goods versus services) or based on the parties involved in the transaction (e.g., consumers or traders). This classification can affect the applicable legal regime, i.e., whether one can have ownership rights over data can distinguish goods from services. However, these criteria do not capture the legal harms arising on such markets. First, although the business models and industries formed around data are becoming complex [29], they are not often linked to legal qualifications under applicable legal regimes. Second, given the features of dark markets, such as pseudonymity and cybercriminality, it is often difficult to distinguish between consumers and traders. Other classifications focus on the type of data transacted and acknowledge personal, corporate, and financial data [26]. Still, as neither the law (e.g., personal data as defined by the GDPR is a highly interpretable concept) nor business practice led to harmonized criteria of cataloging data, such concepts remain vague.

The most popular approach to categorizing data-related transactions on the dark web focuses on the functional description of listings, which is the starting point of the categories employed by the marketplaces themselves. Dark web marketplaces as data markets can be thus said to have three essential characteristics:

- *From the perspective of the object of the transactions*, they are inherently illegal and, in consequence, operate in an adversarial setting for the purpose of which typical legal classifications of data goods and services are irrelevant.
- *From the perspective of the transacting parties*, pseudonymity makes it challenging to specify the nature of the actors in a transaction (B2B/B2C/C2C).
- *From the perspective of the nature of the listings*, while some items can be identified based on their functions (e.g., malware, tutorials, or guides), this clustering may sometimes result in overlaps (e.g., selling access to a cash-out bank account is very similar to selling access to a porn or a Netflix account).

2.2 Dark Web Marketplaces and Criminal Qualifications

Legal criminal qualifications of activities on dark web data marketplaces reflect a considerable research gap due to at least the following factors:

Lack of legal harmonization. Criminal rules governing cybercrime are inherently national. In other words, how cybercrime is regulated, interpreted, and enforced is left to the discretion of sovereign states [18]. Translating these standards into computational frameworks entails mapping national rules to create annotation taxonomies, which requires considerable interdisciplinary efforts involving computational methodologies and comparative law.

Regulatory debt. With fast-evolving business models leading to new iterations of dark web marketplaces, the law is often criticized for lagging behind, as new technologies may render existing legal frameworks obsolete [2]. For criminal law, this may result in a constant need to generate new interpretations for existing legal standards or draft regulation in a future-proof manner.

Prioritizing legal practice. As cybercrime often leads to harm of real people, the literature on the dark web aims to assist law enforcement authorities optimize compliance with criminal rules [2]. Thus, theoretical frameworks based on expert knowledge are essential in making legal standards computational.

As such, we propose an exploratory legal taxonomy that aims to classify listings on dark web data marketplaces. Such a taxonomy makes two main contributions: (*i*) it complements existing academic research on measurement, economics, and criminology by systematically mapping a legal regime applicable to cybercrime, and (*ii*) it offers an approach to assessing cybercrime harms by highlighting maximum penalties imposed at the national level in the Netherlands. This theoretical framework shows the potential of exploring cybercrime from a comparative law and computational perspective, as further research can explore additional jurisdictions so that comparisons between legal regimes can be made.

Cybercrimes can be qualified according to general and special rules. To provide insights into criminal qualifications systematically, we depart from the Council of Europe's Convention on Cybercrime. The Convention governs a wide range of cybercrimes, ranging from hacking to interfering with computer systems. It also criminalizes certain Internet activities, but it does not harmonize the level of the sanctions imposed on these crimes since that is left to the discretion of the ratifying countries. This is why it is essential to map further the implementation of the Convention in a national legal system. To this end, we chose to report on the Netherlands, a jurisdiction where law enforcement has been actively pursuing the reduction of cybercrime through international and European coordinated actions. Dutch courts have also been increasingly dealing with cases relating to phishing and hacking. A complete overview of the Dutch implementation of the Convention can be found in Table 1, where the sanctions applicable to the cybercrimes comprised therein can also be consulted. It must be noted that this overview reflects legal statutes and thus constitutes a theoretical and descriptive rendition of the legal regime applicable in the Netherlands as a result of the ratification of the Cybercrime Convention.

Table 1. Legal qualifications of cybercrime according to the Cybercrime Convention and Dutch criminal law

Convention	Dutch Criminal Code	Maximum Jail Sentence (NL)
Illegal Access	Hacking	Max. 4 years
Illegal Interception	Tapping over Telecoms	Max. 2 years
	Placing of Tapping	Max. 4 years
Data Manipulation	Intentional Interference	Max. 4 years
Systems Interference	Intentional Sabotage	Max. 15 years
Misuse of Devices	Placing of Tapping	Max. 15 years
Computer-Related Forgery	Forgery	Max. 6 years
	Skimming	Max. 6 years
Computer-Related Fraud	Fraud	Max. 4 years

2.3 The Challenges of Mapping Legal Regimes

Mapping the criminal legal regime applicable to dark web data markets has several limitations. First, as the Cybercrime Convention pre-dates dark web marketplaces operating at scale, crimes such as illegal access and interception, data manipulation, systems interference, misuse of devices as well as computer-related fraud and forgery, reflect a criminal landscape with less intermediation than the supply chains amplified by the dark web. Courts, though, may interpret that in providing the tools with which additional crimes are perpetrated, for instance, sellers of malware can be held accountable for the crimes of their buyers.

Second, criminal courts' application of any legal rules will entail a level of discretion that goes hand in hand with the evidence presented in a criminal indictment, based on procedural safeguards. Evidence is also used to prove a perpetrator's criminal intent, essential in sentencing. Two provisions relevant to the Convention's implementation on Dutch law have not been integrated (see Table 1) since they deal with negligent interference and negligent sabotage, and they go beyond the seller-buyer relationship this paper focuses on.

Third, theoretical insights from Dutch law cannot be extrapolated to other legal systems with potentially different criminal public policies. While we do not tackle the problem of applicable law in this exploratory framework, it is worth noting that determining what criminal law applies to cyberspace is in itself a fascinating albeit highly complex question. Therefore, the location and nationality of sellers on the dark web may play a role in applying different or even diverging rules. This is why harmonization is necessary for this field.

Finally, other criminal rules are also applicable. Most prominently, these rules include criminalizing the making, distribution, or possession of child pornography, copyright-protected content, or other forms of illegal content. Interestingly, all the markets included in this paper have provisions excluding the sale of child pornography, and some prohibit even certain cybercrimes covered by the Convention. Dark web marketplaces can be seen as private legal orders which make their own rules regarding the conduct allowed on the platform. Most platforms draw up general terms and conditions to deal with rights and obligations for both vendors and buyers, ranging from contractual to moral standards. These terms also touch upon digital goods/services, such as government data.

3 Measuring Activity on Data Markets

3.1 Data Collection

To compile a list of marketplaces to crawl, we monitored the active marketplaces listed in the `#Markets` category on *Onion.Live*, a clearnet Tor Network directory created to monitor and study popular `.onion` hidden services. Having excluded marketplaces with niche specializations such as guns, drugs and/or cannabis, as well as local marketplaces (with the exception of Hydra, which is reportedly one of the largest and most resilient markets, having been in existence since 2015), we selected twelve omnibus marketplaces: *Asean*, *Big Blue*, *Darkfox*, *Dark Market*, *Deepsea*, *Empire*, *Hydra*, *Icarus*, *Neptune*, *Torrez*, *Versus*, and *White House*. After an initial exploration of these marketplaces, we targeted our collection of data to the offers listed under the *"Digital products"* and *"Fraud"* (where available) marketplace categories only, as these categories were most likely to contain offers of interest, essential for our study.

Since the available offers in the marketplaces continuously alter (i.e., new offers appear and the old ones get removed), we periodically crawled each marketplace to generate a more representative corpus of these offers. Our crawler was based on `Selenium`, a software-testing framework for web applications that can programmatically control a real web browser (Google Chrome connected to Tor in our experiments). This approach allowed us to retrieve the entire content of a rendered offer, which may not be possible if we used a simple command-line tool like `wget`. We scraped the content included in the categories mentioned above in a period of four months (June – September 2020). It must be noted that during this time, two of the markets (Empire and Icarus) were taken down, and thus we were able to crawl only a portion of these two markets. Another important fact is that all of the markets were available in English, except for Hydra, which was available in Russian.

Most of these marketplaces attempt to keep their activities away from prying eyes, especially those of automated bots designed to extract information of the marketplaces' activities. As such, they have deployed CAPTCHA mechanisms to protect themselves. To overcome this hurdle, we initiated the crawling process by logging in into the markets, manually solving any necessary CAPTCHAs, and storing the login and CAPTCHA cookies. The crawler then used these cookies to collect the data from the marketplaces without any barriers. Someone could claim that the process could become entirely automated, using machine learning techniques able to solve CAPTCHA challenges [27]. We did not try this approach because it requires long training periods, and the number of images needed to model each type of CAPTCHA made it unfeasible for us, given the number of markets we scraped, most of which used a different kind of CAPTCHA challenge. In addition, human intervention would still be needed when solving logical puzzles present in many markets as anti-DDoS measures. However, we found ways to minimize the amount of human intervention necessary by taking advantage of blind spots in these markets' bot detection algorithms or by exploiting bugs in the sites' implementation. These include:

1. Switch the onion circuit through which the crawler accesses the market.
2. Rotate trough different mirrors of the same market before or when getting blocked by the market.
3. Log out and log back in before reaching the threshold and flagged as a bot.
4. Go through the search results of a category and save the links that are then accessed randomly to avoid sequential scraping of the products.
5. Wait random times between visits.

For markets where we could not avoid being flagged as a bot, either we made the crawler notify us and wait until we would intervene if graphical puzzles were present, or in the case of markets that only required regular CAPTCHAs to be solved, an email was sent, and the solved CAPTCHA was read and submitted to the market to resume the scraping.

3.2 Data Preprocessing

Our crawler exfiltrated all the data available when visiting an offer from a marketplace. However, as the various marketplaces differ from each other, so do their data representation, which can produce misleading results. Therefore, the quality of data and their representation is considered the most critical step before running any analysis. As such, we have to bring the data to such a state that our algorithms can easily parse and interpret it.

Duplicates: The dataset duplicates were identified and removed. The reposts of the offers (e.g., exact title match but different price, date, or description) were kept for correlations and future analysis.

Prices: The offers' prices are displayed in different currencies depending on the market preferences. To have a more accurate view regarding the offers' prices, we normalized them by exchanging the displayed value to USD using the exchange rate recorded on 31 August 2020. From the price analysis, we dismissed the prices equal to zero (free offers) and higher than USD 1,000,000. After that, we discarded the outliers identified as prices lower/higher than two times the standard deviation below/above the average value for each market.

Vendors: Due to the pseudonymized nature of the scraped markets and the lack of vendor identification upon their registration, the vendors' cross-market identity cannot be fully recognized. We considered *unique vendors* by matching their exact username across markets. In addition, we anonymized the usernames of the vendors in the presented results. We attributed common first names to the vendors, keeping them consistent cross-table.

Categories: The available offers were categorized, and 15,377 of them were also sub-categorized. Given that many of the categories and sub-categories on the markets were relatively generic and not always comparable for the selected markets, we categorized the offer titles ourselves. We explored topic modeling and strategies deployed in previous studies, such as human labeling of the complete dataset. Natural language processing techniques, such as topic modeling

Table 2. Keywords assigned to categories

Category	Keywords
General data	*account, database, plaintext, leads, accounts, streaming, hacked, voter, vpn, mobile, hacking, email, voters, cracked, crack, records, record, porn, clone, access, config, mba, checker, emails, dtabase, sentry, numbers, buffered*
Banking & tokens	*card, carding, balance, credit, money, bank, cvv, cashout, gift, egift, carded, cards*
E-learning	*method, tutorial, guide, hack, amazon, make, get*
PII	*psd, template, license, statement, passport, ssn, dob, fullz, utility*
Other	*fraud, snapshot, month, mac, pack, android, paypal, market, login, live, bitcoin, generator, btc, usa*

yielded poor results. Instead of a labor-intensive human labeling process that would be difficult to apply in other contexts, we opted for a simple heuristic. We categorized the offers by selecting the most prevalent terms (or unigrams) in the offer titles across the platforms. We grouped the most prevalent and relevant unigrams (e.g., cvv, passport, porn) after discarding irrelevant terms (e.g., premium, lifetime, or other descriptors in the titles), the removal of which would not affect the nature of the offer. To do so, we manually scanned through the complete titles for each of the 100 most prevalent and relevant unigrams.

3.3 Economic Activity on Data Markets: Strange Facts

We explored various metrics that shed light on the economic activity taking place on the scraped marketplaces by showcasing relevant descriptive statistics. Based on the generated keywords and the categories previously found on the markets, the following categories were produced: (*i*) *General Data*, (*ii*) *Banking & tokens*, (*iii*) *E-learning*, (*iv*) *Personal Identifiable Information (PII)*, and (*v*) *Other*. Each offer was placed under categories based on its title containing at least one keyword from Table 2. This resulted in a coverage of 85,58%, meaning that almost nine out of ten offers could be categorized based on the selected unigrams, and some titles were considered in more than one category.

We compared the categorization results of the unigram method to those 11,261 product titles that had sub-category information from the labels used on the scraped markets. To do so, we manually grouped all different sub-categories present in the data and derived from the different markets (in total, there are 45, some overlapping) into the five categories that we constructed when using the unigram method. This way, we could compare for which of these products there is an agreement in the categorization. The method we applied finds that there is a perfect match for only 46.70% of the titles. That percentage rises to 69.87% when considering products that the unigram assigns to two categories. By manually inspecting some of the titles where there was a disagreement, we observed that our method is superior in classifying many titles more accurately. Similarly, many sub-categories of markets that have broad titles (e.g., *Other*) are more accurately classified by our method (e.g., into *E-learning* or *PII*).

Table 3. Overview of markets. Some vendors participate in more than one market; this is why the total number of vendors differs from the expected one.

Market	No. of offers	No. of vendors	Sum of offers ($)
Asean	4,043	36	30,663.95
Big Blue	1,669	78	31,381.80
Darkfox	1,300	34	21,446.80
Dark Market	3,629	127	73,956.12
Deepsea	4,210	111	46,757.14
Empire	2,690	135	38,176.93
Hydra	204	204	11,434.43
Icarus	4,091	37	28,961.87
Neptune	2,352	23	16,187.55
Torrez	615	14	5,399.99
Versus	901	25	5,330.83
White House	2,842	88	145,977.37
Total	**28,546**	**642**	**455,674.78**

Table 4. Overview of categories

Category	No. of offers	No. of vendors	Sum of offers ($)	Max. jail time (years)
General data	15,219	277	248,603.16	**2 – 15**
Banking & tokens	5,673	269	125,888.12	**2 – 6**
E-learning	3,821	181	27,920.27	**4**
PII	2,689	150	42,996.53	**6**
Other	10,410	573	186,039.71	**n/a**

We subsequently looked at the sum of all offers on the twelve marketplaces. This represents the total offers (goods and services) advertised on the platforms (Table 3), which is one way of estimating the value of the total supply of data economy in the scraped categories (Table 4). However, prior research has shown the problems with such estimates [20]. These offers sometimes do no reflect real prices but are rather scam (i.e., meant to deceive buyers), spam, or may employ techniques such as *"holding price"* (i.e., raising the price so much that no one can afford to buy it), in an attempt to keep the offer listing open, while having the (temporary) intention of not selling, or marking the offer as not in stock.

An extreme example of holding price found in our dataset for the offers *"out of stock Wowcher accounts with balance Auto Delivery & Lifetime Warranty"*, as well as for *"not working Auto Delivery & Warranty"*, both listed at the skyrocketing value of USD 11,111,100,000.00. As the titles indicate, they are listed as *"out of stock"* and respectively *"not working"*. Particular attention can be paid to spam offers (Table 5). It seems that certain vendors list the same offer on the same market up to 25 times. Perhaps, this may be a marketing strategy aimed at making an offer more visible when browsing through listings to increase the number of sold items. However, what works for one vendor on one market (e.g., vendor *"Barb"* on *"Dark Market"*), might not apply to other vendors on other markets (e.g., vendor *"Phil"* on *"Neptune"*). Additional attention needs to be paid to 17 free offers in our dataset (USD 0.00): 10x *virtual camwhore* (offered by the same vendor), 1x *porn tutorial*, 2x *porn accounts* (offered by the same vendor), 1x *Spotify*, and 2x *bitcoin exchange accounts*.

Table 5. Top 10 spam offers

Offer	# Posts	Market	Vendor	# Sold	Price ($)
Credit Cards #1	25	Dark Market	Barb	189	20.00
Carding Software Setup	16	Dark Market	Barb	65	5.00
E-Gift Cards	14	Icarus	Connie	*n/a*	1.00
Debit Cards	13	Dark Market	Scott	16	20.00
Spotify Account	13	Neptune	Phil	0	1.92
RealityKings Account	11	Neptune	Phil	0	7.23
Credit Cards #2	10	Dark Market	Barb	44	15.00
TeamSkeet Account	10	Neptune	Phil	0	7.23
Torrent Accounts	10	Big Blue	Karen	*n/a*	20.52
Credit Cards #3	9	Dark Market	Holly	60	15.00

Table 6. Sum of Offers (with and without outliers)

Market	Sum of offers ($) with outliers	Sum of offers ($) without outliers	Difference
Asean	51,948.35	30,663.95	21,284.40
Big Blue	390,481.80	31,381.80	359,100.00
Darkfox	35,099.62	21,446.80	13,652.82
Dark Market	118,516.85	73,956.12	44,560.73
Deepsea	82,374.17	46,757.14	35,617.03
Empire	638,176.93	38,176.93	600,000.00
Hydra	18,511.24	11,434.43	7,076.81
Icarus	35,507.87	28,961.87	6,546.00
Neptune	31,184.52	16,187.55	14,996.97
Torrez	7,344.94	5,399.99	1,944.95
Versus	7,564.39	5,330.83	2,233.56
White House	486,177.37	145,977.37	340,200.00
Total	**1,902,888.05**	**455,674.78**	**1,447,213.27**

There is no way of identifying holding prices that are not outliers. Thus, we decided to remove listings that were two standard deviations above and below the mean when calculating the sum of offers. Since we report the sum of offers rather than means of price listings, there was no need to remove zero-price listings. The results suggest that removing the outliers does not only have a significant impact on the sum of offers, but it also substantially reduces the differences between the sum of offers of the markets Table 6.

We continued the analysis by inspecting the distribution of vendors across multiple markets (Figure 1). While the majority of vendors are present on one market (78.63%), vendors with the same name are present on anywhere between two and five markets (19,97%), and a few are present on up to six and nine markets (1,4%). Similarly, offers are also posted on multiple markets, albeit by different vendors. Table 7 and Table 8 provide more insights into which vendors and offers can be mostly found across the markets we inquired into.

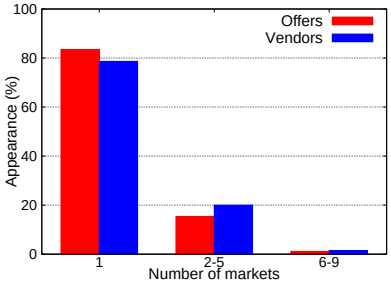

Figure 1. Offers and vendors that appear in multiple markets

Table 7. Top 10 vendors and offers by the number of sold units

Vendor	# Units	Offer	Market	Category	Price ($)	# Units
Eleven	13,838	Netflix Account #1	Empire	General data	3.35	4,712
Erica	9,211	Doordash USA Account	Empire	General data	2.99	2,354
Steve	5,523	Grubhub Account	Empire	General data	0.99	1,501
Billy	4,712	TryCaviar Account	Empire	General data	0.75	1,493
Jim	4,572	Get a free iPhone	Empire	E-learning	4.99	1,491
Bob	4,475	Make 2500$ a day on Bet365	Empire	E-learning	4.99	1,425
Robin	3,920	Tip Jar	Versus	Other	1.00	1,409
Sam	3,603	2 Brazzers Accounts	Empire	General data	1.95	1,338
Will	3,336	Netflix Account #2	Empire	General data	0.80	1,327
Murray	3,269	Deep Web Onion Links List	Empire	General data	1.30	1,280

Table 8. Top 10 vendors and offers by the number of markets

Vendor	# Markets	# Offers	Sum of offers ($)	Offer	# Markets	Category	Vendor	Price ($)
Will	9	3,120	30,312.86	Atlas Quantum Database	9	General data	Will	9.99
Jim	8	789	4,850.62	JobStreet Database	9	General data	Will	9.99
Robin	8	4,058	16,836.61	Money Bookers Database	9	Banking & tokens	Will	9.99
Mike	7	172	4,098.12	United Kingdom Mobile Numbers	9	General data	Will	9.99
Dustin	7	264	4,181.57	Australia Mobile Numbers	9	General data	Will	9.99
Lucas	6	81	1,440.39	Germany Mobile Numbers	9	General data	Will	9.99
Nancy	6	116	4,017.56	Italy Mobile Numbers	9	General data	Will	9.99
Jonathan	6	364	1,990.92	Oregon Voter Database	9	General data	Will	9.99
Karen	6	1,545	24,533.33	Canadian Business Database	8	General data	Will	4.50 - 65.00
Max	6	328	5,643.88	Canadian Residential Database	8	General data	Robin	9.00 - 65.95

We also explored the associations between the number of offers, the number of vendors, and the sum of offers on the various marketplaces. For this, we ran Pearson correlation tests on the elements mentioned above. The correlation between the number of offers and vendors is close to zero ($r = -.03, ns$). The correlations between the number of offers and the sum of offers ($r = .44, ns$) and between the number of vendors and the sum of offers ($r = .25, ns$) are moderate and weak, respectively, but both not statistically significant, which does not surprise because of the low sample size.

Finally, we looked at how much markets make based on their commissions (Table 9). For this purpose, we investigated what commission rates platforms apply by manually checking the general terms, FAQ sections, and other descriptive materials markets make available to their users. Some markets seem to be more buyer-friendly by only charging lump percentage fees to vendors, others only charge buyers, and some markets charge both vendors and buyers. A few markets have set up progressive commission rates, so they charge vendors depending on, for instance, the transaction price bracket they can be found in. The commission information was not available for all markets. Wherever it was available, we were able to estimate the maximum commission turnover. This is based on the maximum commission calculated on reported sales, which could be accounted for. As it can be observed in Table 10, the commission turnover varies across markets (e.g., Neptune's turnover resulted in a mere USD 2.03), and it can highlight whether data transactions are profitable for marketplaces.

Additionally, most markets also have rules on vendor bonds, namely a one-time fee that users need to pay to obtain vendor status. Not all markets charge bonds, and some deem them refundable. If marketplaces apply bonds to all ven-

Table 9. Fees and commission rates

Market	Vendor fee (%)	Buyer fee (%)	Vendor bond ($)
Asean	3	*n/a*	400
Big Blue	1.5 - 3.5	0.5 - 2	250
Darkfox	4 - 5	*n/a*	150
Dark Market	*n/a*	5	750
Deepsea	2 - 4	*n/a*	150
Empire	*n/a*	*n/a*	*n/a*
Hydra	1.5 - 5	no fee	300
Icarus	*n/a*	*n/a*	*n/a*
Neptune	*n/a*	4	125
Torrez	4 - 5	no fee	250
Versus	*n/a*	*n/a*	*n/a*
White House	5	no fee	400

Table 10. Estimated marketplace turnover

Market	Maximum Bond ($)	Maximum Commission ($)	Estimated Turnover ($)
Asean	1,617,200.00	*n/a*	1,617,200.00
Big Blue	417,250.00	*n/a*	417,250.00
Darkfox	195,000.00	54.40	195,054.40
Dark Market	2,721,750.00	5,864.17	2,727,614.17
Deepsea	631,500.00	4,826.68	636,326.68
Hydra	61,200.00	*n/a*	61,200.00
Neptune	294,000.00	2.03	294,002.03
Torrez	153,750.00	*n/a*	153,750.00
White House	1,136,800.00	*n/a*	1,136,800.00
Total	**7,228,450.00**	**10,747.28**	**7,239,197.28**

dors equally and if the bonds are non-refundable, they can make a considerable turnover based on this business model. However, many markets indicate in their terms and conditions or FAQs that bonds may be waived for various reasons (e.g., if vendors can prove their legitimate activity on other marketplaces). It is, therefore, problematic to make more accurate estimates of the marketplace turnover. In Table 10, we estimate the turnover based on the maximum turnover based on bonds and commissions. Overall, for the markets where this data was available, we could estimate that across six marketplaces, the overall bond and commission turnover amounts to USD 7,239,197.28. This does not account for additional business models such as the monthly rent which Hydra charges its vendors, as we have incomplete information regarding whether this fee is charged to all vendors, or only active vendors. If applicable to all vendors, this fee (USD 100.00) would generate an additional monthly turnover of USD 20,400.00 for the 204 vendors identified in our snapshot. Moreover, Hydra also refers to a commission charged for disputes started, meaning that economic incentives may be linked to reputation costs even more directly, as disputes can end up costing money in commissions paid, and not only sales lost as a reputation cost.

4 Discussion

4.1 Main Findings

Dark web data markets prove difficult to operationalize empirically. The field lacks methodologies for measuring basic characteristics such as the classification of data products. We explored topic modeling and strategies deployed in previous studies, such as human labeling of the complete dataset. Topic modeling yielded poor results. Instead of a labor-intensive and dataset-dependent process of human labeling, we opted for a simple heuristic. This approach successfully allowed us to identify offers for three categories: *Banking & tokens*, *E-learning*, and *PII*. Yet, most offers were assigned to the *General data* and *Other* categories.

These categories were then used to understand the legal risks according to the Cybercrime Convention implementation in the Netherlands. *General data*, the category featuring most listings and vendors, may reflect crimes sanctionable with incarceration where the maximum punishment ranges from 2 to 15 years. The high ceiling of this range is driven by the crime of intentional sabotage, which certainly includes the use and distribution of worms, viruses, trojans, and ransomware. However, punishments may only go beyond 6 years if lives are in danger or lost due to intentional sabotage. This may be the case when ransomware is used to take hospitals or other essential service operators hostage. This category of harms arising from dark web data marketplaces may deserve more individual attention or categorization in further research. What is interesting is that after setting aside the most harmful malware listings as described by the crime of systems interference (Convention) or intentional sabotage (Dutch Criminal Code), the category of *General data*, which includes the sale of data dumps affecting millions of individuals, leads to less severe punishments (max. 4 years) for the category of *PII*, which entails forging individual documents based on sold templates (max. 6 years). Similarly, the crime of skimming, associated with the category of *Banking & tokens*, leads to a higher punishment (6 years) than the crime of hacking (4 years). In the light of the measurement we completed, as *General data* is the most popular listing category, it is worth asking whether the current hierarchy of punishments in Dutch law is fit to tackle the realities of the dark web. Further research should explore the role of knowledge relating to criminal punishments on the activity of dark web markets.

The economic activity on the selected markets was examined in terms of the offers listed, activity across the 12 markets, best-performing products, and marketplace turnover. It was found that markets differ significantly in terms of the number of offers, vendors, and the sum of offers, with no weak and moderate correlations among them, meaning that there are markets with relatively few vendors but with a large number of offers and vice versa, but also there are markets with a large number of vendors, a large number of offers, and a large sum of offers. Consequently, the question arises whether it is possible to speak about `the dark web data marketplace` or it is more accurate to see this marketplace as a collection of markets with significant differences among them.

Interestingly, the sum of offers of around USD 455,000 is overwhelmingly lower than figures that are floating around in market research (e.g., USD 1.5 trillion annually [26]). Even when looking at the most popular ten offers (Table 7), there is a massive difference between how many units of the first and tenth most popular offers were sold (4,712 vs. 1,280), based on reported sales of one of the largest marketplaces in our dataset (Empire). The fact that not all marketplaces list sales is an interesting finding in itself. On earlier marketplaces such as the Silk Road, sale reports reflected a vendor's reputation [5]. However, reputation building takes time and maintenance (e.g., dealing with fake sales, fake reviews), and more volatile markets may not have sufficient resources to develop these systems under the adversarial circumstances they need to operate.

Furthermore, we find several similar listings and vendor names across different platforms. This finding is important because, in its early days, the dark web had a handful of markets that invested many resources into devising original solutions to improve their resilience. Nowadays, with ±20% of vendors operating on up to five marketplaces, this is reminiscent of developments on the regular web due to the platformization of digital transactions by intermediaries. Platforms simplify economic activity and reduce the necessary literacy skills and transaction costs. While the Internet's first dark web marketplaces had to be built from scratch, current marketplaces may take the form of Platform-as-a-Service, which may explain the proliferation of both marketplaces as well as the business models they support. This proliferation can signify a variety of marketing techniques in getting more business out of data transactions.

The cross-posting of items likely inflates estimates of cybercrime revenues, at least with respect to dark web data markets. Of course, we only included twelve markets, and our data collection is a snapshot rather than a longitudinally collected dataset. Nevertheless, our findings do question how much value should be attached to popular estimates of how the value is associated with data on the dark web. More generally, one may wonder how accurate any predictions about dark web data markets are, considering the difficulty of defining which product categories should fall under the umbrella of *dark web data markets*.

4.2 Limitations of the Study

For two of the markets (Empire and Icarus), we only have partial data, as they were taken down during our scraping. We also acknowledge that these statistics do not reflect longitudinal data but are a snapshot of the economic activity at the time of scraping. Like any online market, dark web data markets may be subject to constant changes. However, longitudinal of continuous scraping is difficult considering the technical measures (e.g., CAPTCHA) taken by markets. Moreover, certain variables of interest are difficult or impossible to capture. Closed transactions, and consequently, the actual number of sales, revenue, and profit, are not reported and are impossible to retrieve. Furthermore, categorization, or natural language processing tasks, prove difficult due to spelling mistakes, jargon, and abbreviations. Addressing these limitations might reveal additional characteristics and patterns compared to the ones presented in this paper.

4.3 Future Research

Future research should further compare and develop the categorization of titles of dark web data markets as well as address how to collect data over an extended period so that results across different studies can be more directly comparable. While academic literature on dark web marketplaces is growing, studies remain disparate and use complementary yet uncoordinated approaches from the perspective of a vast array of disciplines to investigate what is happening on dark web marketplaces. Yet, the nature of these marketplaces and the business models behind some of the offers (e.g., data markets) are highly volatile, so coordination may lead to more clarity regarding dark web marketplaces as objects of study.

This is not to say that further angles cannot be added to this already vast body of knowledge. A comparative analysis of cybercrimes in different legal systems (e.g., all EU countries or EU countries and the United States) could support additional or alternative computational measures of the activities offered on the dark web market to design public policy on this matter. Research may also include exploring relationships between what is offered on the markets by whom and against which price on the one hand and the governance structure of the platforms, including the terms of service. More generally, research into the trust mechanisms developed and applied on dark web markets and comparisons with offline equivalents could contribute to understanding the role of contracts and trust.

Finally, the platformization of dark web marketplaces warrants further investigation. While exploring and contextualizing dark web data markets, we encountered services offered for the creation of markets. Similarly to creating a WordPress website, one can create a dark web marketplace. It is interesting to discover the business models behind such platformization on dark web data markets, and dark web marketplaces in general, and how the platformization interacts with the economic activity in those places.

5 Related Work

Exploring the specifics of dark web activity has attracted the interest of many computer science studies. Research has explored anonymity and privacy regarding the use of the dark web [1] and tried to answer the questions about which actors can be found therein [23]. Trust and engagement in dark web forums have also been studied widely. In particular, there have been studies exploring the popularity of listings either by looking at trust mechanisms [25] or predicting demands for drugs (on such markets) using Wikipedia views [16]. Similarly, there has been research that explores the connection of dark web markets with the global drug supply chain [4]. Researchers have also been trying to provide forensic frameworks for assisting with the investigation of dark web markets [6]. Previous works have also attempted to automatically identify drug traffickers based on writing and photography styles [28], or detect multiple identities of the same vendor over different dark markets, again using photography style [24].

Researchers have looked into the factors that contribute to criminal performance and which influence the advertised price for offers like dumps and account credentials [8], the signals of trust used by vendors to indicate trustworthiness within their advertisements for stolen data [9], the structure and organisation of underground forums [7], and the type of interventions that can be applied [10]. Researchers have also developed tools for the automated analysis of cybercrime markets operating on underground forums [17], for profiling their member users and identifying top vendors in these communities [13].

Researchers have also studied the cybercrime and stolen dark web data markets by comparing the distribution of victim nations in stolen data markets and examining variations between Open and Dark Web operations [19], evaluating the factors influencing pricing for stolen identities [21], and categorizing products offered on marketplaces specializing in malicious hacking products [15]. A few classifications and measurement studies draw insights from legal scholarship, but expert legal analyses are generally lacking from this work [12].

To the best of our knowledge, this paper presents the first systematic study at the intersection of expert legal knowledge and web measurement approaches. It explores the characteristics of twelve dark web marketplaces focusing on the data economy in particular, instead of dark web markets in their entirety (e.g., including drugs) or on a sub-category of cybercrimes (e.g., malware economy).

6 Conclusion

This paper set out to understand and describe the criminal and economic activity on dark web data markets by focusing on two research questions: how to use criminal law insights from international and Dutch law to sketch an exploratory legal framework applicable to dark web data markets and how to measure such markets using an original large-scale dataset of twelve scraped marketplaces.

References

1. A. S. Beshiri, A. Susuri, et al. Dark Web and Its Impact in Online Anonymity and Privacy: A Critical Analysis and Review. *Journal of Computer and Communications*, 7(3):30–43, 2019.
2. M. Chertoff. A Public Policy Perspective of the Dark Web. *Journal of Cyber Policy*, 2(1):26–38, 2017.
3. N. Christin. Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace. In *The Web Conference*, 2013.
4. M. Dittus, J. Wright, and M. Graham. Platform Criminalism: The 'Last-Mile' Geography of the Darknet Market Supply Chain. In *The Web Conference*, 2018.
5. C. Goanta. The Private Governance of Identity on the Silk Road. *Frontiers in Blockchain*, 3, 2020.
6. D. R. Hayes, F. Cappa, and J. Cardon. A Framework for More Effective Dark Web Marketplace Investigations. *Information*, 9(8), 2018.
7. T. J. Holt. Exploring the Social Organisation and Structure of Stolen Data Markets. *Global Crime*, 14(2-3):155–174, 2013.

8. T. J. Holt, Y.-T. Chua, and O. Smirnova. An Exploration of the Factors Affecting the Advertised Price for Stolen Data. In *eCrime Researchers Summit*, 2013.
9. T. J. Holt, O. Smirnova, and A. Hutchings. Examining Signals of Trust in Criminal Markets Online. *Journal of Cybersecurity*, 2(2):137–145, 2016.
10. A. Hutchings and T. J. Holt. The Online Stolen Data Market: Disruption and Intervention Approaches. *Global Crime*, 18(1):11–30, 2017.
11. T. S. Hyslip and T. J. Holt. Assessing the Capacity of DRDoS-For-Hire Services in Cybercrime Markets. *Deviant Behavior*, 40(12):1609–1625, 2019.
12. S. Kaur and S. Randhawa. Dark Web: A Web of Crimes. *Wireless Personal Communications*, 112(4):2131–2158, 2020.
13. W. Li and H. Chen. Identifying Top Sellers in Underground Economy Using Deep Learning-Based Sentiment Analysis. In *IEEE joint intelligence and security informatics conference*, 2014.
14. M. Macdonald and R. Frank. Shuffle Up and Deal: Use of a Capture–Recapture Method to Estimate the Size of Stolen Data Markets. *American Behavioral Scientist*, 61(11):1313–1340, 2017.
15. E. Marin, A. Diab, and P. Shakarian. Product Offerings in Malicious Hacker Markets. In *IEEE conference on intelligence and security informatics (ISI)*, 2016.
16. S. Miller, A. El-Bahrawy, M. Dittus, M. Graham, and J. Wright. Predicting Drug Demand With Wikipedia Views: Evidence From Darknet Markets. In *The Web Conference*, 2020.
17. R. S. Portnoff, S. Afroz, et al. Tools for Automated Analysis of Cybercriminal Markets. In *The Web Conference*, 2017.
18. M. Shillito. Untangling the Dark Web: An Emerging Technological Challenge for the Criminal Law. *Information & Communications Technology Law*, 28(2), 2019.
19. O. Smirnova and T. J. Holt. Examining the Geographic Distribution of Victim Nations in Stolen Data Markets. *American Behavioral Scientist*, 61(11), 2017.
20. K. Soska and N. Christin. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem. In *USENIX security symposium*, 2015.
21. C. M. Steel. Stolen Identity Valuation and Market Evolution on the Dark Web. *International Journal of Cyber Criminology*, 13(1):70–83, 2019.
22. R. Van Wegberg et al. Plug and Prey? Measuring the Commoditization of Cybercrime via Online Anonymous Markets. In *USENIX security symposium*, 2018.
23. M. Wang, X. Wang, J. Shi, Q. Tan, Y. Gao, M. Chen, and X. Jiang. Who Are in the Darknet? Measurement and Analysis of Darknet Person Attributes. In *International Conference on Data Science in Cyberspace (DSC)*, 2018.
24. X. Wang, P. Peng, C. Wang, and G. Wang. You Are Your Photographs: Detecting Multiple Identities of Vendors in the Darknet Marketplaces. In *Asia Conference on Computer and Communications Security*, 2018.
25. F. Wehinger. The Dark Net: Self-Regulation Dynamics of Illegal Online Markets for Identities and Related Services. In *European Intelligence and Security Informatics Conference*, 2011.
26. E. Wilson. Disrupting Dark Web Supply Chains to Protect Precious Data. *Computer Fraud & Security*, 2019(4):6–9, 2019.
27. A. Zarras, I. Gerostathopoulos, and D. M. Fernández. Can Today's Machine Learning Pass Image-Based Turing Tests? In *Information Security Conference*, 2019.
28. Y. Zhang, Y. Fan, et al. Your Style Your Identity: Leveraging Writing and Photography Styles for Drug Trafficker Identification in Darknet Markets Over Attributed Heterogeneous Information Network. In *The Web Conference*, 2019.
29. Z. Zheng, J. Zhu, and M. R. Lyu. Service-Generated Big Data and Big Data-as-a-Service: An Overview. In *IEEE International Congress on Big Data*, 2013.