Tokyo, Denver, Helsinki, Lisbon or the Professor? A Framework for Understanding Cybercriminal Roles in Darknet Markets

***Abstract*—There is comparatively little information about the roles and the separation of these roles within financially-motivated cybercrime online. As Darknet Markets (DNMs) are online fora, roles can often be conflated with membership or user types within such fora, e.g., administrator, new user, etc. The insights presented in this paper are grounded in a Conversation Analysis of underground forum threads in combination with Social Network Analysis of the relationships between actors in these fora and an automated analysis of the thematic scope of their communications using NLP techniques. This results in a more nuanced understanding of roles, and the power relationships between roles, as they emerge through and are defined by linguistic interactions. Based on this mixed methods approach, we developed a dynamic typology of three key roles within DNMs that goes beyond a basic supply-demand logic: *entrepreneurs*, *influencers* and *gatekeepers*. A closer analysis of these roles can contribute to a better understanding of emerging trends in a forum and allow for the identification and prioritisation of high-risk targets.**

***Index Terms*—Darknet, Cybercriminal Roles, Social Network Analysis, Natural Language Processing, Conversation Analysis, Mixed Methods**

I. Introduction

With digital technologies becoming pervasive across society, law enforcement agencies are increasingly engaged in tackling ever more sophisticated cybercrime. In 2017, the UK National Cybercrime Unit (NCCU) published its findings from research on debriefing arrested cybercriminals [1]. The report identified that cybercrime is not a solitary and anti-social activity, but one wherein online social interactions play a critical role in the recruitment, training and professional advancement of criminals. As such, investigating these social interactions is important to understanding the dynamics leading to initial engagement in cyber crime, continued careers and (potentially) retirement.

While prior work has led to valuable insights in terms of socio-demographic characteristics, motivations and behavioural patterns of cyber offenders (e.g., [2]–[5]), there is comparatively little information available about the roles and separation of roles within cyber crime communities. For instance, do entry-level criminals specialise in one form of criminal activity, and later broaden their remit, or does the reverse occur, with early offenders having a broad but shallow skillset that is then specialised as they become "professionals"? Additionally, there is a noticeable bias towards self-reporting surveys, which rely on the willingness and ability of individuals to be forthcoming and to articulate their motivation(s) to participate in cybercrime. This privileges individuals who feel confident and safe enough to do so at the expense of those who could play a prominent role in Darknet fora but, for various reasons, might be reluctant to disclose such information. This leads to a limited view of the dynamics of different user roles in such fora.

In this paper, we combine a qualitative analysis with novel techniques in the area of Natural Language Processing (NLP) and Social Network Analysis (SNA), enabling a corpus-based approach that incorporates all users and their communications in Darknet fora. This approach allows for a shift of the research focus from self-reporting surveys to a more systematic approach leading to a multifaceted understanding of such roles and how individuals move between them. More specifically, the key contributions of this study are as follows:

- We present a dynamic typology, which goes beyond a basic supply-demand logic (cf. Section II). More specifically, we provide an in-depth and qualitatively interesting understanding of roles and power relations between roles as they emerge through and are defined by linguistic interactions.
- We describe a novel unsupervised learning methodology to automatically categorise offenders within this dynamic role typology, which allows for cybercriminal forums and marketplaces to be subdivided into usefully-delineated sub-communities, and for identifying key users playing prominent roles in these communities.

Additionally, a closer analysis by law enforcement investigators of such roles in any Darknet Market (DNM) can contribute to a better understanding of emerging trends in a forum and enable the identification and prioritisation of high-risk targets according to different mission briefs.

The remainder of this paper is structured as follows. In Section II, we provide an overview of the related work. We discuss the data used for this study in III and describe the unsupervised learning methods developed to support our qualitative analysis in Section IV. In Section V, we present our findings of our qualitative study. Section VI highlights our mixed methods approach to developing a new typology of roles in DNMs. Finally, in Section VII, we discuss the implications and limitations of our findings for improving the analysis of underground fora.

Identify applicable funding agency here.

II. Related Work

Within the area of deep web social network visualisation, a majority of prior research has focused on plotting data towards a surface level (e.g., mapping onion sites based on their connection to each other through URLs [6]). However, none of these works focused on analysing communications within DNM offender communities.

Two studies have incorporated textual features found across DNMs in a Social Network Analysis. For example, using data collections performed by The University of Arizona's Artificial Intelligence Lab, Arnold et al. [6] developed Cyber-Threat Intelligence (CTI) tools to proactively monitor online hacker communities which leverage a social network analysis approach to identify cyber threats across major DNMs. Their findings showed that fraud, breached accounts and hacking tools were the most prominent cyber threats to companies and their customers. Similarly, Rios et al. [7] combined SNA with Text Mining techniques to detect overlapping extremist communities in Dark Web portals in order to identify potential homeland security threats. However, none of these studies included user-level analyses of social interactions on DNM communities.

Lane et al. [8] proposed an Event Analysis of Systemic Teamwork (EAST) approach to represent and analyse trading activities on a DNM and identify vulnerabilities for potential market disruption. Their analysis led to useful insights into the steps required for a user to start engaging in DNM activities and to buy and sell illegal products. However, information on how users build their reputation, move up the ladder, or resolve disputes within such communities remained limited, because they only had access to about 40 screenshots of pages accessible to users without interaction with other actors and their analysis was restricted to the perspectives of buyers and sellers.

To our knowledge, there is only one study that focused on identifying key users in terms of roles, influence levels, and their social relationships: Huang et al. [9] proposed a topic-based social network analysis approach with unsupervised clustering methods for identifying the key members and their associated roles in the Chinese cyber fraud underground economy. Based on the results of their Latent Dirichlet Allocation (LDA) analysis, they attributed user roles based on the keywords of topics detected in user communications. However, their role typology was based on prior work in the area of the underground economy of credit card fraud [10], and, hence, limited to "attack originators", "buyers", "droppers", "shoppers", "runners", and "other sellers". Such a typology, which is mainly based on forum structure rather than the role users play in a community, or even titles assigned to users during their registration, may not reflect the actual roles users adopt in the market.

In this paper, we combine novel SNA and NLP techniques with an extensive qualitative analysis, allowing for a bottom-up, dynamic approach of detecting different user roles in cybercriminal communities on the Darknet, and leading to a multifaceted understanding of such roles and how individuals move between them.

III. Data

For this analysis, we make use of the DNM Corpus: a large dataset collected between 2013 and 2015 [11]. In particular, we targeted a discussion forum within this collection, the Evolution forum, which acted as support area for an underground marketplace dealing in a number of different illicit goods.

The Evolution dataset contains 509,225 messages written by 21,946 different users in total, with on average 23.2 messages per user and 53.1 tokens per message. Each individual in the dataset contributed to on average 11.3 different threads. In Table I, we show the 25 most active users in terms of number of messages and number of different threads to which they contributed.

TABLE I
THE 25 MOST ACTIVE USERS IN THE EVOLUTION DATASET.

User	# Messages	# Threads
FRIM	3,818	1,560
wefinance	3,268	1,077
Yasuo	3,035	788
themostseekrit	2,884	2,012
penissmith	2,833	1,293
Kimble	2,755	2,154
LudoTilMortem	2,663	2,345
EludingHell	2,456	1,023
BlueHighSky	2,162	1,709
Scattermind	2,096	1,450
Trippyy	1,749	1,011
ScoobyJew	1,637	794
nswgreat	1,587	736
moka	1,579	753
elmachico777	1,538	631
leon-trotsky	1,473	587
scrufffe	1,459	469
evilsmile	1,436	506
Ozzyz	1,430	657
First	1,423	550
Cypher	1,411	795
Grandeur	1,398	367
mountainhigh9	1,362	426
andigatel	1,291	454
bitbybit	788	648

The raw data provided in [11] captures fora as scraped at several semi-regular intervals by the dataset authors. This leads to heavy redundancy within the data, as threads may be captured at multiple times. However, this redundancy is also useful, as it helps to guard against intermittent faults in the crawling process. Our approach to parsing the data takes a *latest-version-first* view – of all pages captured within the crawling process, we treat as canonical the most recent version, only parsing older pages where they were not captured in later scrapes. We note that capturing pages from older scrapes is an important step in handling this data, as many thousands of threads and user profile pages are not present at all in the most recent scrapes of each forum. Differences could be attributed to crawling failures in later scrapes, incomplete coverage as part of the crawling processes, or to administrator action in taking down or hiding discussion threads over time.

Parsing of the data proceeded in two stages within the scrape history of each community. First, user profile pages

were processed to build up a dataset of users and associated information from their profile pages (e.g., PGP public keys, membership status). Next, discussion thread pages were parsed in order to associate posts (including textual content and metadata such as posting time, subforum, etc.) with the user that authored them. Where quotations of other users could be identified within the text of a user's post, these quotations were separated from the authored text, to avoid contamination of the NLP analysis. It sometimes occurred that user profile pages were not captured in the scrapes due to sites protecting access to those pages, or where users were observed posting for whom no profile page had been seen (either due to people using guest accounts, or due to incomplete coverage of profile pages in the crawls). In these cases, new user entries were created on the fly during the second stage of parsing, using such metadata as was available about the author account from the post metadata.

IV. Methodology

This study was set up as follows: first, we applied Conversation Analysis (CA) on a sample of Evolution threads relating to scamming support and shipping (Section V. Next, in order to assess our findings in a more systematic way, we applied the Social Network Analysis approach (Section IV-A) and used our best performing topic modelling approach (Section IV-B) to the conversations of key users identified by this analysis. Finally, based on a combination of our qualitative and quantitative approaches, we developed a new typology of roles in DNMS (Section VI). We start this section by discussing the quantitative methods used to support our research.

A. Cybercriminal Network Analysis

Social Network Analysis is a popular method to investigate social structures through the use of Graph theory [12]. It transforms networked structures into nodes (e.g. individuals within a network) and edges (i.e. the ties or links between the nodes). In this study, we apply a weighted Social Network Analysis approach[1] to identify the most important nodes in the Evolution dataset (i.e., "community influencers") based on their contributions to different forum conversations (or threads). More specifically, we consider two users as linked if they have both contributed to the same thread and the total number of threads to which they both contributed as a measure to weigh the strength of the link between them. This resulted in a weighted graph with 21,393 nodes (or active users) and 4,485,425 edges.

A key aspect of our approach is community detection. Within Social Network theories, the idea is that a large network can be divided into smaller sub-structures [14]. Identifying such sub-communities in Darknet fora can be useful to gain insight into how a forum as a whole is likely to behave. For example, if memberships of these sub-communities overlap, it can be expected that trends occurring in one sub-community may spread rapidly across the entire forum. To automatically detect sub-communities and their members in the Evolution dataset, we applied Clauset-Newman-Moore greedy modularity maximization [15], which begins with each node in its own community and joins the pair of communities that most increases modularity until no such pair exists. This approach identified 56 sub-communities in the Evolution dataset. As can be seen in Figure 1, the Evolution dataset is comprised of 4 larger sub-communities, which contain 11,725; 6,645; 1,885; and 926 users. The other sub-communities did not include more than 51 nodes.

Fig. 1. Sub-communities detected in the Evolution dataset

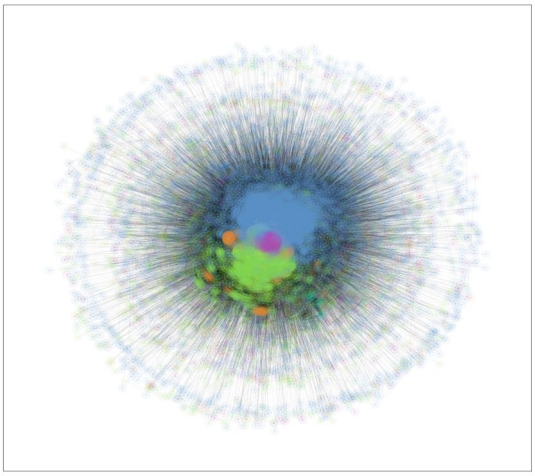

Next, within each of the 4 larger sub-communities, we applied the following centrality measures:

- Degree Centrality: this is a measure of the number of links a user has in the forum. The more connections a user has, the more influential he or she can be.
- Local Clustering Coefficient: this metric measures the extent to which each user is located within a tight "cluster" of neighbouring users in a forum. In other words, it provides information on the probability that a user(A)'s linked nodes (A-B, A-C) are also linked to each other (B-C).
- Eigenvector Centrality: this measures the importance of a user based on his/her connections to other important users in the forum.

The users that display high centrality scores potentially have a strategic position in the sub-community and, hence, can play a significant role in influencing the communication or information flow within a forum. Therefore, we extracted users who yielded above average scores for all three centrality measures (Av. Deg. Cent. = 0.16; Av. Loc. Cl. Coef. = 0.58; Av. Eig. Cent. = 0.01).

The results of the community detection approach and the centrality measures were aggregated to generate a list of potential users of interest, and were used to inform our

[1]We used the Python package Networkx [13]

qualitative analysis as discussed in Section V. Table II, shows the 40 most influencial users according to our Social Network Analysis.

TABLE II
THE 25 MOST IMPORTANT USERS IN THE EVOLUTION DATASET ACCORDING TO THE SNA.

User	Eig. Cent. (Av = 0.01)	Deg. Cent. (Av = 0.16)	Loc. Clust. Coef. (Av = 0.58)
Thinkpad	0.0197	0.1733	0.9495
Buddhacus	0.0199	0.1781	0.9162
anewusername	0.0200	0.1795	0.9078
powerhacks	0.0203	0.1820	0.8954
dave2014	0.0202	0.1814	0.8946
potus50	0.0208	0.1842	0.8896
anonbuyer1997	0.0203	0.1831	0.8849
volumatic	0.0203	0.1840	0.8816
PhyishNet	0.0204	0.1848	0.8717
Mr_Sickness	0.0207	0.1859	0.8697
Cajun	0.0210	0.1876	0.8687
DarkNet563	0.0210	0.1876	0.8661
Suckafree	0.0202	0.1848	0.8655
bigwang45feet	0.0204	0.1867	0.8610
FrankySauce	0.0205	0.1870	0.8604
acarhitme	0.0212	0.1907	0.8475
reelwananigga	0.0213	0.1926	0.8465
Jimmy911	0.0213	0.1929	0.8420
1292FOR	0.0211	0.1918	0.8403
headcrusher	0.0212	0.1943	0.8182
finalfantasy1337	0.0210	0.1949	0.8118
stalinpub	0.0216	0.1996	0.8063
lolapalooza	0.0213	0.1993	0.7986
kmekme	0.0212	0.1988	0.7928
whocares	0.0221	0.2016	0.7911

B. Analysing the Thematic Scope of DNM Communications

In most NLP approaches, the training data is prelabelled with the required information to perform a categorisation task (i.e., the "ground truth"). However, detecting the thematic scope of Evolution users' communications required an unsupervised learning approach, because no information on the presence of different topics was available in the dataset. Hence, we focused on developing a methodology that could reveal linguistic patterns from the unlabelled data and provide us with an understanding of the thematic clusters that could be inferred automatically from each user's collection of messages. For the purposes of this study, we took a text clustering approach, in which the similarity between different text samples can be measured by using one or more similarity functions.

With regard to textual data in which the objects can be of different granularities (e.g., documents, paragraphs, sentences or words), clustering methods have shown promising results for, e.g., browsing or organising documents and summarising large text corpora [16]. Standard practice for vector data is to use the K-means algorithm or Latent Dirichlet Allocation (LDA). The first technique divides a set of text samples into k disjoint clusters, each described by the centroid of the text samples in the cluster. The algorithm then attempts to select centroids that minimise the within-cluster sum-of-squares (or inertia) [17]. LDA, from its part, is a Bayesian probabilistic model, which also assumes a collection of k clusters. The latter algorithm has been applied successfully on social media communications, because it assumes that each document instance is a mixture of a small number of topics and that each word can be clustered into one of these topics [18].

Because the texts provided to the learner are unlabelled, no actual categorisation is performed and, hence, there is no evaluation of the accuracy of the output of the similarity algorithm. Therefore, we calculated the mean Silhouette Coefficient (SC), which is a measure to validate the consistency within the resulting clusters of data. More specifically, the Silhouette Coefficient is defined for each sample and is composed of two scores:

1) x: The mean distance between a sample and all other points in the same class.
2) y: The mean distance between a sample and all other points in the next nearest cluster.

$$SC = \frac{y - x}{max(x, y)}$$

Values for this score range between [-1, 1]. Values near 0 indicate overlapping topic clusters. Negative values tend to indicate that samples have been assigned to a wrong cluster [19].

To obtain a detailed model of the thematic scope represented in a user of interest's communications, and to avoid topic bias that is potentially present when analysing larger posts, we decided to split up each post into sentences and treat each sentence as a separate object. However, due to the nonstandard use of punctuation in the DNM dataset, none of the default sentence tokenisers included in NLP tools (e.g., NLTK, Scikit-Learn and Spacy) produced accurate results in our preliminary experiments. Hence, we trained a custom sentence tokeniser on the raw text messages in the Evolution dataset. Again, as we did not have any ground truth labels for this task, we trained an unsupervised learning model, which incorporated abbreviation words, collocations, and words that start sentences to detect sentence boundaries in the Evolution dataset. Next, we removed urls and punctuation – keeping emoticons intact –, tokenised all sentences, and extracted the following textual feature types:

- **content words**: Bag-of-Words with stopwords removed;
- **stemmed content words**: stemmed Bag-of-Words, and stopwords removed[2];
- **bigrams**: word bigrams, no stopwords removed;
- **noun chunks**: base noun phrases that have a noun as their head[3];
- **noun heads**: the head of each noun chunk.

Since sentences are very short text samples, we aggregated them into longer pseudo-texts by incorporating the semantic knowledge from word embeddings that were pre-trained on Twitter data (glove.twitter.27B[4]). Using word embeddings allows for detecting semantic similarities between words based on their distributional properties in large corpora, which could

[2]We used NLTK's Lancaster Stemmer in our experiments (https://www.nltk.org/)

[3]These were extracted using Spacy's NLP model (https://spacy.io).

[4]https://nlp.stanford.edu/projects/glove/

boost the performance of our unsupervised topic detection model.

Tf-idf was used for feature weighting. As can be seen in Table III, the best SC was achieved when applying K-means using word embeddings, which resulted in 4 topics with an average SC score of 0.35. After closer inspection of the most informative words for each topic in the communications of 10 randomly selected users, we found that, compared to the word embeddings topic model, the other feature types generated more fine-grained sub-topics. Therefore, both the word embeddings and the nouns topic models were used to inform our qualitative analysis described in the next section. We provide a few examples in Table IV.

TABLE III
AVERAGE SILHOUETTE SCORE AND NUMBER OF TOPIC CLUSTERS PER FEATURE TYPE USING K-MEANS.

Feature Types	**Av. SC**	**# Topic Clusters (=*k*)**
BoW	0.09	98
Content words	0.11	97
Stemmed content words	0.11	92
Content words + bigrams	0.10	94
Noun chunks	0.17	91
Nouns	0.23	97
Noun chunks + content words	0.10	98
Word embeddings	0.35	4

V. QUALITATIVE ANALYSIS AND RESULTS

For our qualitative analysis we first performed a purposeful sampling on threads related to scamming support and shipping. Next, we applied Conversation Analysis (CA) to these threads. Conversation analysis is based on the following three main assumptions: (*i*) talk is structured, (*ii*) talk is forged contextually, and, (*iii*) analysis is grounded in data [20]. We chose CA for the following reasons. First, it allows us to study talk, as it occurs naturally in social settings and in local vernaculars or slang. It therefore has the potential to study informal and less codified language, which is not possible using other methods of analysis that privilege formal or codified language. Second, the approach privileges practical and common sense reasoning and sees social order as something constructed through conversation, rather than as a pre-existing given. Third, it is a data-driven and inductive approach. It is seen as a positivist-leaning approach due to its emphasis on rigour, validity, and replicability, which means that this is more of a text-reducing method (as opposed to text-enhancing methods, such as Discourse Analysis) [20], [21].

We applied the basic principles of CA to the examined threads. These are: (*i*) whether participants in a conversation take turns during the communication; (*ii*) whether there was evidence of adjacency pairs[5]; and (*iii*) whether users articulate preferences in any specific order (e.g., if the shipping of a product is delayed, would the buyer prefer to wait and be reassured by the vendor that the product will be delivered, or would the buyer prefer a refund instead).

[5]The idea that a question is followed by an answer, an invitation by an acceptance/rejection, etc.

Applying Conversation Analysis to thematically defined threads was useful overall. Specifically, threads about shipping and scamming could be relevant for exploring roles because different subjectivities become more pronounced within the context of grievances, or when users express frustration and discontent. Our key findings are as follows:

A. Cybercrime is not antisocial in nature

First, our CA confirmed that cybercrime is not anti-social in nature – this is evident by the possibility to study thematically defined threads where multiple users engage with each other on a particular topic. In each examined thread there were a handful of individuals who were dominating said thread, whilst the vast majority of users remained rather passive, although social interactions could occur through other means (see below). Interestingly, while there was a lot of communication activity, the textual interactions were not always conventional: users rarely take turns when communicating, a question is not always followed by an answer, there are more statements than questions, and not all questions or comments are addressed. There are plenty of lone comments and questions which remain unanswered. The lack of conventionality is interesting within the wider context of communication.

Furthermore, the examined threads also revealed an awareness of roles and rules and what needs to be done in order to progress from one role to another, or to "retire" (usually by means of an exit scam). The analysis indicated that exchange of important information and know-how is likely to happen via alternative communication channels, i.e. encrypted direct messages or offline.

B. Customer service prevails without moral labels

We uncovered that forum users did not necessarily think that what they were doing was essentially criminal. To be clear, there was awareness that it is illegal or punishable by law, but the examined conversations did not reveal any moral labels (right/wrong) attached to the activity. We found evidence of "customer service" type of interactions between vendors and buyers. There were discussions about refunds, discounts and raising formal disputes, which is an interesting indication that users might not necessarily view these interactions as criminal *per se*.

C. Entrepreneurs in action

Regarding the users who were most active and contributed the most to threads, we found that they were rather entrepreneurial: they were openly opportunistic in the pursuit of income and seemed to be into drugs, cryptocurrencies and other types of fraud (carding, fake Pay Pal, Ebay, Amazon refunds, etc.). These users also seemed more likely to try and diversify their cyber criminal portfolio, rather than just specialise in particular types of drugs.

D. Fluidity of roles

While on the surface it seemed that illicit drug markets, such as Evolution, were highly structured domains, we also

TABLE IV
EXAMPLES OF THE NOUN- AND WORD EMBEDDINGS-BASED TOPIC MODELS.

Feature Type	Top Content Words per Cluster
Nouns	Cluster 1: 'like', 'just', 'pretty', 'don', 'really', 'people', 'shit', 'actually', 'think', 've', 'drugs', 'good', 'time', 'know', 'drug'
	Cluster 2: 'account', 'vendor', 'just', 'use', 'like', 'banned', 'actually', 'post', 'buyer', 'time', 'good', 'forums', 'iban', 'number', 'definitely'
	Cluster 3: 'really', 'exactly', 'don', 'know', 'does', 'sorry', 'sure', 'going', 'problem', 'like', 'yeah', 'read', 'uhm', 'think', 'honest'
	Cluster 4: 'money', 'make', 'just', 'lot', 'making', 'don', 'like', 'people', 'pretty', 'way', 'trying', 'vendor', 'buy', 'think', 'account'
	Cluster 5: 'right', 'just', 'know', 'want', 'probably', 'don', 'sure', 'wait', 'really', 'think', 'looking', 'need', 'll', 'yeah', 'work'
	Cluster 6: 'way', 'like', 'pretty', 'just', 'really', 'good', 'best', 'don', 'great', 'make', 'bad', 'people', 'actually', 'things', 'experience'
	Cluster 7: 'don', 'think', 'really', 'know', 'sure', 'just', 'll', 'going', 've', 'actually', 'like', 'personally', 'pretty', 'didn', 'yeah'
	Cluster 8: 'vendor', 'just', 'order', 'like', 'good', 'time', 'really', 'fe', 'don', 'make', 'account', 'isn', 'know', 'think', 'message'
Word Embeddings	Cluster 1: 'vendor', 'moved', 'section', 'welcome', 'account', 'url', 'trashed', 'evo', 'wrong', 'fraud', 'vendors', 'forum', 'escrow', 'scam', 'open'
	Cluster 2: 'vendor', 'just', 'order', 'make', 'good', 'account', 'sure', 'dont', 'like', 'evo', 'open', 'people', 'fe', 'forum', 'time'
	Cluster 3: 'url', 'trashed', 'en', 'vous', 'pas', 'je', 'inappropriate', 'le', 'confirmed', 'onion', 'cest', 'ca', 'des', 'anglais', 'il'

found a certain flexibility with regard to the rules. One way of progressing is to buy a vendor's account. Apparently, these could either be externally funded or self-funded, which suggests that one's ability to change roles would depend on one's economic capital and offline networks.

E. Privacy everywhere

Important aspects of textual interactions take place via private messages, video calls, email, as well as offline (if the users knew each other). Moreover, users often went to great lengths to preserve their privacy, as well as the integrity and anonymity of their immediate networks.

F. Social trust relationships

Finally, we confirmed that trust was important for the communication and cohesion of Darknet communities (see also [22]). However, it gained an extra layer, because in these high-risk transactions, even reputable vendors could fail to deliver on occasion. This means that disappointed customers have a choice between confronting the vendor or trying their luck elsewhere. Judging by the volume of posts in terms of seeking support, it looks like users place substantial trust in the administrators of the forum; more so than on their vendors. This could indicate that administrators/moderators might perform a more important role than previously believed.

Overall, employing CA to study purposefully sampled threads has been analytically beneficial and could be an important methodological contribution, because, generally, there is lack of clear guidance on how to (qualitatively) make the most out of Darknet Fora data. It also revealed important underlying power relations between users. In order to investigate these power relations in a more systematic way, we applied the Social Network Analysis (SNA) as described in Section IV-A. The results are described in the next section.

VI. A ROLE TYPOLOGY BEYOND SUPPLY-DEMAND LOGIC

The results of our Social Network Analysis showed that over 95% of all users in the Evolution dataset were passive. Within the active users, the analysis identified 56 sub-communities, of which only 4 contained users who yielded above average centrality measure scores. Therefore, it was useful to see what these users do to make themselves stand out. In other words, our understanding is that roles are based on social interactions via the medium of text, which supports the claim that Darknet interactions are a social activity. When applying a further cut-off of including users who produced at least 1,000 posts, our approach finally yielded 135 potential users of interest. Next, for each of these users, we extracted topic clusters using both the word embeddings and the nouns topic models presented in Section IV-B to gain a better understanding of the thematic scope of their communications.

Based on the automated analyses described above, we were able to formulate a suggested typology of the roles in Darknet fora. The typology of roles is based on (*i*) **power relations**: we account for these based on how active some members are in the forum (i.e., volume of posts) and how well-connected they are based on the SNA centrality measures scores; and, (*ii*) **thematic scope of their contributions**: based on the output of the topic models described above. The resulting role typology is displayed in Table V.

TABLE V
TYPOLOGY OF THE ROLES IN CYBERCRIMINAL DARKNET FORA.

Free riders	The passive majority of users who do not contribute (much) to the forum and display no desire or ambition to move up the ranks. Based on the SNA, this group consists of 95% of users in the Evolution dataset.
Entrepreneurs	Opportunistic users who are always looking for new ways to generate income. These users are the most willing to diversify their activities – drugs, cryptocurrencies, selling/buying vendor accounts, carding, low-level scams (refunds from Amazon, Ebay, PayPal, etc.). We focused on the 0.3% of users (87 users) who have transitioned from being a member to being a vendor, as per SNA results.
Influencers	Users who portray themselves as knowledgeable and as experts. The influencers are likely to provide advice, to be mentors, or to demonstrate technical expertise and experience. We identified the influencers by focusing on the volume of posts per user as per SNA. We focused on the 35 persons or 0.12% of Evolution members, who had posted 1,000 posts or more.
Gatekeepers	The very few individuals who have the highest centrality scores and without whom the network would fragment. Here, we have focused on the 0.04% (13 people in total) users with admin/moderator functions.

Figure 2 shows a social network graph of these users' connections within the Evolution dataset.

This suggested typology is not necessarily exhaustive, nor does it need to be. Our suggested typology has the following

benefits:

- Focusing on the three active roles (entrepreneurs, influencers and gatekeepers) will help optimise the volume of data, which in turn enables a more useful visualisation of the data.
- It allow us to account for status within the forum, as well as for cultural capital (variety of topics discussed, expertise).
- It is a dynamic typology, which goes beyond a basic supply-demand logic. In other words, we can gain an in-depth and qualitatively more interesting understanding of roles as they emerge through and are defined by linguistic interactions.
- It allows for comparative studies and replicating our findings by exploring textual dynamics in other online fora.
- These three roles are interesting both from an academic and law enforcement perspective, because they are typically embedded within many social interactions. Exploring them further can contribute to a better understanding of emerging trends in a forum and allow for the prioritisation of investigative targets according to different mission briefs.

The final part of our analysis focuses on determining the validity and credibility of the proposed typology, with particular emphasis on the three active roles: the entrepreneurs, the influencers, and the gatekeepers. This is done by applying our topic modelling approach described in Section IV-B to all communications produced by these users in the Evolution dataset. More specifically, we manually explored the most prominent topics of each selected user in each of the three categories, i.e., the topics that had been attributed at least 100 sentences by our topic modelling approach, including the top words for each topic and 25 randomly selected original sentences attributed to each topic. The main goal was to verify the validity of our proposed role typology. Thus, we explored data for 34 influencers, 31 entrepreneurs, and 8 gatekeepers. We provide a few examples of the topic modelling results for each role in Table VI. The top words are extracted using all BoW features after we clustered the messages using the word embeddings and noun topic models.

Our combination of manual and automated analyses suggest that the proposed typology is valid and that there is analytical merit in exploring textual interactions within Darknet fora by focusing on power relations. However, there is a certain fluidity that defines the roles, and it is possible that there will be spill overs between them. More specifically, our key findings are as follows:

- **Brusque influencers**. The language used by the influencers was rather brusque, to the point, at times rude or even insulting to other users. The explored rhetorical data suggested that influencers have established good reputations and so they used language in a more functional way (e.g., no chit-chat). We also found evidence that influencers seek to project assertiveness and confidence, bordering on cynicism. Unlike with other roles, influencers appear particularly emotionally detached and lacking empathy, which could possibly stem from their technical expertise, as well as knowledge of what the rules are and how to break them:
  "*I got an idea, big boy: how about you try doing it yourself, and then maybe think twice about how much bullshit it is to sell counterfeit money to a bunch of daydreaming, inexperienced beginners?*"(themostseekrit)
  "*I'm allowed to my person opinions and anyone who says otherwise can get fucked as-well.*" (nswgreat)
- **Polite entrepreneurs**. The examined data suggests that entrepreneurs are the politest of the three roles, and the most likely to display attention to customer service and customer satisfaction. Greetings and expressions of agreement and gratitude were featuring more prominently in the data describing this role. Entrepreneurs' enthusiasm for selling product and customer satisfaction is not unlike common marketing strategies used by regular retailers: discounts, money-back-guarantees, and free samples are some of the tools that entrepreneurs employ in order to establish their reputation within the Evolution community:
  "*Check me out, ask me questions, let's getrich together:) Included my link now, still waiting on certain people with a lot of posts and rep to answer back with how everything worked for them!*" (ohman1988)
  "*To those of you that did not get to try our free samples, we will soon have another exciting new product which we will need reviews on too!*" (InsideTheWhale)
- **Enigmatic gatekeepers**. there did not seem to be a coherent theme that defined the role of gatekeepers. Overall, this role posts the least in terms of volume and

Fig. 2. Social Network Analysis graph of the Evolution dataset depicting free riders (grey), entrepreneurs (red), influencers (green) and gatekeepers (blue).

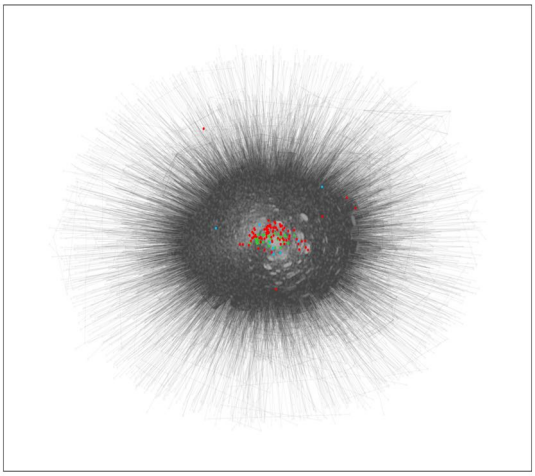

TABLE VI
EXAMPLES OF THE TOPIC MODELLING RESULTS FOR EACH ROLE IN THE TYPOLOGY.

Role	Top words per topic cluster	Example message of the topic cluster
entrepreneur	'just', 'copies', 'like', 'guys', 'good', 'left', 'little', 'free', 'folks', 'people', '30', 'use', 'know', 'price'	*Have 4 more copies left at the $15 Intro price, then sadly I'll be moving up to $25-30 haven't decided yet;p Get it before it's gone and at a cheap cheap price!!*
entrepreneur	'uk', 'bank', 'million', '10', 'price', 'people', 'account', 'just', 'vendor', 'transfer', 'address', '100', 'gbp', 'good'	*My price is accurate Source: I have a database of over 105Million unique people Acquired over the last 3 years.*
influencer	'glad', 'decided', 'feedback', 'just', 'afraid', 'like', 'actually', 'working', 'try', 'appreciate', 'vendors', 'buyers'	*It would be nice to see a wider adoption of PGP, but we can't really force it on buyers.*
influencer	'vendor', 'account', 'order', 'marked', 'section', 'url', 'forum', 'mandatory', 'open', 'shipped', 'fe', 'cancel', 'tag', 'new', '50'	*SOmetimes vendor list the same item twice and one fo the two is slightly cheaper, this one is FE.*
gatekeeper	'bought', '556', 'number', 'digits', 'thanks', 'forum', 'market', 'ids', 'unconnected', 'good', 'youve', 'fnufnu', 'just', 'parts', 'left'	*Thanks for confirmation.On BMR it wasnt a concern as forum and market ids were unconnected so nobody could possibly tell what you had bought but with the evolution way of having them conjoined its good to get it confirmed by someone in the know that noone can see what youve bought.I think its a good idea to establish separate unconnected ids for market and forum anyway - theres no extra marks for being a hero.*
gatekeeper	'wallet', 'bitcoins', 'think', 'good', 'works', 'tor', 'blockchain', 'thank', 'tails', 'use', 'today', 'mail', 'time', 'thanks', 'hope'	*I would be content to make payments western union.Is it possible on the way soon therm today to its transactions with localbitcoin with a wallet blockchain, mix with helix, before sending another wallet and finally transferring on Evolution.What do you think of this method for a single buyer.This is anything?*

the references are too thematically random, which made it difficult to infer an overarching theme. Whilst some sentences suggested that gatekeepers effectively liaise between vendors and buyers, trying to resolve disputes and providing some form of customer service support, others indicated that forum admins/moderators could be vendors or buyers themselves:
"*We have a team of active and dedicated moderators, all of which either have been a vendor, or still are a vendor.*" (BlueHighSky)

VII. DISCUSSION AND CONCLUSION

The present study confirms that combining quantitative and qualitative approaches is essential when dealing with conversational (textual) data collected from Darknet fora, which lack any form of pre-categorisation or labelling information, and in which users often go to great lengths to remain anonymous and preserve the integrity and anonymity of their immediate networks. Moreover, it also presents a corpus-based approach that enables a shift of the research focus from self-reporting surveys, leading to a limited view of the dynamics of different user roles in such fora, to a more systematic approach that incorporates all users and their communications in a selected dataset. Until now, this kind of inclusive approach is missing in the research on cybercriminal roles in Darkweb markets. Additionally, it is worth noting that CA is traditionally used to study verbal communications, which are then transcribed in studies for patterns and irregularities. However, it is possible to adapt the method for the study of online forum discussions. This is an important methodological contribution, because generally there is lack of clear guidance on how to (qualitatively) make the most of online discussion fora data.

Our qualitative analysis confirms that cybercrime is not anti-social in nature. Multiple users tend to engage with each other on a particular topic. We noted that high-activity threads are usually dominated by a handful of individuals — either forum admins, more experienced users, or persons who want to establish their reputation through active engagement. The analysis also shows that users often go to great lengths to preserve their privacy, with trust playing an important role in the communication and cohesion of Darknet communities.

In an extra layer of complexity, in high-risk transactions (e.g., sale of illegal products) even reputable vendors can fail to deliver on occasion. In such cases, it seems that users place substantial trust in the administrators of a forum, rather than on the vendors, indicating that administrators might play a more important role than previously attested.

We have developed a dynamic typology that goes beyond a basic supply- demand logic, drawing together our understanding of roles and power relations between roles as they emerge through, and are defined, by linguistic interactions. Significant elements of these interactions include the power relationships between users — "novices" and "internals" being a distinctive categorisation in most cases — which in some venues are expressed as "customer service" interactions between the vendors and buyers of illegal products. This suggests that, although users are aware of the illegal nature of their activities, they might not necessarily view them as criminal or morally wrong.

We attested a certain fluidity that defines the different roles and progression to a different role. For example, one way of progressing from a member to a vendor is to buy an existing vendor's account, using own or external funding, which suggests that a user's ability to change roles also depends on their economic capital and offline networks. Additionally, it became clear that forum admins or moderators can easily be vendors or buyers themselves.

Finally, to support our qualitative analysis summarised above, we developed novel techniques for automatically categorising offenders within this dynamic typology. More specifically, we applied a Social Network Analysis approach to identify the most important users in the Evolution forum based on their contributions to different forum conversations. By combining the output of three different centrality scores (Degree Centrality, Local Clustering Coefficient and Eigenvector Centrality) with a community detection method, we

showed that it is feasible to automatically detect users with a strategic position within the a user network and, hence, could play a significant role in influencing the communication or information flow within this forum. Furthermore, we developed an unsupervised learning approach to automatically reveal linguistic patterns from these users' messages, which provided us with an understanding of the thematic scope that was present in their communications. The modules developed for supporting this analysis have been implemented into a software package that is designed to assist law enforcement agencies in their investigations pertaining to cybercriminal activities on the Darknet.

A key limitation of the employment of CA was the often unconventional mode of communication (see Section V), which precluded us from employing the method on a larger scale. Additionally, our method was applied on only one dataset —although vast — of cyber offender communications on DNMs, focusing on illegal drug trade, without any ground truth data available to evaluate our unsupervised learning approach. Therefore, future steps include assessing the validity of our role typology and methodology on other Darknet fora that focus on, for example, on cyber dependent crimes, radicalisation or child sexual exploitation. Additionally, we aim to manually annotate an extensive amount of DNM conversations to enable a more detailed evaluation of our topic models. Finally, word embeddings that were pre-trained on Twitter data might include semantic similarity assumptions that do not uphold when applied to Darknet communications between cyber offenders. For example, offenders can use guarded language or specialized vocabulary in order to hide their illegal activities from law enforcement investigators (e.g. "snow" instead of "cocaine"). Therefore, we intend to train new word embedding models on the entire DNM dataset and include them in our experiments.

REFERENCES

[1] National Cyber Crime Unit, Prevent Team, "Intelligence assessment: Pathways into cyber crime," 2017, available at http://www.nationalcrimeagency.gov.uk/publications/791-pathways-into-cyber-crime/file.

[2] T. Hyslip and T. Holt, "Defining the profile of potential cybercriminals," *Homeland Defence & Security Information Analysis Center (HDIAC) Journal*, vol. 5, pp. 25–30, 2018.

[3] H.-J. Woo, "The hacker mentality: exploring the relationship between psychological variables and hacking activities," Ph.D. dissertation, University of Georgia, 2003.

[4] A. Hutchings and R. Clayton, "Exploring the provision of online booter services," *Deviant Behavior*, vol. 37, no. 10, pp. 1163–1178, 2016.

[5] R. Liao, S. Balasinorwala, and H. R. Rao, "Computer assisted frauds: An examination of offender and offense characteristics in relation to arrests," *Information Systems Frontiers*, vol. 19, no. 3, pp. 443–455, 2017.

[6] N. Arnold, M. Ebrahimi, N. Zhang, B. Lazarine, M. Patton, H. Chen, and S. Samtani, "Dark-net ecosystem cyber-threat intelligence (cti) tool," in *2019 IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE, 2019, pp. 92–97.

[7] S. A. Ríos and R. Muñoz, "Dark web portal overlapping community detection based on topic models," in *Proceedings of the ACM SIGKDD workshop on intelligence and security informatics*, 2012, pp. 1–7.

[8] B. R. Lane, D. Lacey, N. A. Stanton, A. Matthews, and P. M. Salmon, "The dark side of the net: Event analysis of systemic teamwork (east) applied to illicit trading on a darknet market," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 62, no. 1. SAGE Publications Sage CA: Los Angeles, CA, 2018, pp. 282–286.

[9] S.-Y. Huang and H. Chen, "Exploring the online underground marketplaces through topic-based social network and clustering," in *2016 IEEE Conference on Intelligence and Security Informatics (ISI)*. IEEE, 2016, pp. 145–150.

[10] A. Singh, "The underground economy of credit-card fraud." https://www.peerlyst.com/posts/the-underground-economy-of-creditcard-fraud, 2016.

[11] Gwern Branwen. (2015, July) Dark net market archives, 2011-2015. [Online]. Available: www.gwern.net/DNM-archives

[12] D. Knoke and S. Yang, *Social network analysis*. Sage Publications, 2019.

[13] A. Hagberg, P. Swart, and D. S Chult, "Exploring network structure, dynamics, and function using networkx," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2008.

[14] A. Alamsyah, B. Rahardjo *et al.*, "Community detection methods in social network analysis," *Advanced Science Letters*, vol. 20, no. 1, pp. 250–253, 2014.

[15] A. Clauset, M. E. Newman, and C. Moore, "Finding community structure in very large networks," *Physical review E*, vol. 70, no. 6, p. 066111, 2004.

[16] C. C. Aggarwal and C. Zhai, "A survey of text clustering algorithms," in *Mining text data*. Springer, 2012, pp. 77–128.

[17] A. K. Jain, "Data clustering: 50 years beyond k-means," *Pattern recognition letters*, vol. 31, no. 8, pp. 651–666, 2010.

[18] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.

[19] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of computational and applied mathematics*, vol. 20, pp. 53–65, 1987.

[20] A. Bryman, *Social research methods*. Oxford university press, 2016.

[21] R. M. Smedley and N. S. Coulson, "A practical guide to analysing online support forums," *Qualitative research in psychology*, vol. 18, no. 1, pp. 76–103, 2021.

[22] L. Norbutas, S. Ruiter, and R. Corten, "Believe it when you see it: Dyadic embeddedness and reputation effects on trust in cryptomarkets for illegal drugs," *Social Networks*, vol. 63, pp. 150–161, 2020.