

Computer-Supported Cooperative Crime

Vaibhav Garg¹, Sadia Afroz²(✉), Rebekah Overdorf¹, and Rachel Greenstadt¹

¹ Drexel University, Philadelphia, USA

² UC Berkeley, Berkeley, USA

sadia.afroz@berkeley.edu

Abstract. This work addresses fundamental questions about the nature of cybercriminal organization. We investigate the organization of three underground forums: BlackhatWorld, Carders and L33tCrew to understand the nature of distinct communities within a forum, the structure of organization and the impact of enforcement, in particular banning members, on the structure of these forums. We find that each forum is divided into separate competing communities. Smaller communities are limited to 100–230 members, have a two-tiered hierarchy akin to a gang, and focus on a subset of cybercrime activities. Larger communities may have thousands of members and a complex organization with a distributed multi-tiered hierarchy more akin to a mob; such communities also have a more diverse cybercrime portfolio compared to smaller cohorts. Finally, despite differences in size and cybercrime portfolios, members on a single forum have similar operational practices, for example, they use the same electronic currency.

Keywords: Cybercrime · Economics · Social network analysis · Dunbar number

1 Introduction

The notion of what it means to be ‘organized’ is contentious, even in traditional crime [12]. Cybercrime complicates this debate through underground forums, where the cooperation can be described both as (vertically integrated) firms [20] and cybercrime commons [2]. Understanding the organization of such criminal networks, however, can help distinguish important actors on these forums, the economic efficiency of enforcement, and the comparative impact of distinct enforcement strategies [21, 28, 32].

Research in economics of security establishes the incentives for cybercriminals to organize (e.g. specialization [13]), the cost to cooperation (e.g. ripper tax [15]), and cybercriminals’ response in managing trust (e.g. banning misbehaving members [2, 32]). However, it does not examine the resulting nature of cybercriminal organizations as shaped by these distinct and often conflicting forces, e.g. preferential attachment [30].

This paper addresses three questions about the nature of cybercriminal organization on underground forums. *First*, is a single underground forum comprised of distinct cybercriminal communities? If so what are the similarities and differences across communities, specifically in terms of topics of communication between participants? *Second*, we compute and correlate various measures of centrality (or importance) for individual cybercriminals on a single underground forum. Centrally located cybercriminals may receive more responses to public posts [26], be more trusted by their peers [23], have access to more quality information [7], need fewer overall transactions and thus lower associated costs [25], and enjoy leadership positions [5]. *Finally*, we investigate the impact of community (rule) enforcement on underground forums. Specifically, we examine the impact of banning members on social networks metrics that are associated with sustainable trust management in cybercriminal online forums [2]. We make four contributions:

1. **We show that there are distinct sub-communities of cybercriminals on underground forums.** Smaller communities have 100-230 members, similar to the Dunbar number, and have a two-tiered hierarchy with centralized control similar to a gang [6]. Larger communities have flatter hierarchies, distributed control, and multiple tiers, similar to a mob [6] as well as a more diverse cybercrime portfolio.
2. **We note that most communities sub-specialize in specific crimes.** Communities on a single forum, however, have similar operational practices.
3. **We find that different measures of centrality correlate on all forums.** Some cybercriminals may enjoy disproportional advantage as they may simultaneously be more popular [26], imbue more trust [23], have access to better and more information [7], have lower transaction costs [25], and be considered leaders [5].
4. **We observe that banning misbehaving nodes can have a tangible and positive impact on the structure of the network.** When members with higher closeness/betweenness centrality are removed the change in small world characteristics may be greater. Thus, individuals who can propagate information over shorter paths are better at reducing trust in the network.

2 Background

Previous work observed two kinds of organizations in traditional crime [27]: gangs and mobs. Gangs have a two-tier hierarchy with a central leader and a group of followers that adhere to central command; mobs have a more complex command and control structure and typically specialize in specific crimes. Ethnographic accounts from the 1980's noted that cybercriminal organizations lacked characteristics of a mob, as they do not specialize [22]. While there is incentive for individual market participants to specialize [24], it is unclear whether the same is true for organized cybercrime entities, e.g. to leverage comparative advantage [18]. In this paper, we begin to explore the notion of 'organization' in

cybercrime as it applies to underground forums. We analyze three underground forums that were leaked anonymously and were publicly available. Our study is orthogonal to the previous studies on these forums that provided descriptions of the forums [26], analyzed cybercrime commons [2] and proposed an algorithm to identify duplicate identities of pseudonymous cybercriminals [1].

Trust: The main challenge to cybercriminal organization is the lack of trust among peers [32] and incentive to cheat [15]. Décary [7] notes the presence of small communities on IRC chat rooms; he argues that small communities allow each member to know everyone else, emphasizing the importance of direct ties. Humans may only have meaningful relationships with up to a 150 people (i.e. the Dunbar number), with a confidence interval of 100 to 230 [10], even online [9]. We examine whether or not an underground forum starts to divide into distinct communities as the size of an underground forum increases beyond the Dunbar number.

The Dunbar number should not necessarily limit cybercriminal membership in a single community, as it is possible to design mechanisms to scale trust [2]. For example, peer-produced ratings allow buyers to evaluate a seller for credibility [26]; forum members who do not comply with rules, e.g. by creating duplicate accounts, can be banned. We explore how trust and trust management strategies, e.g. banning members, shape the organization of cybercriminal networks on underground forums.

Importance of centrality measures: Individual criminals have a higher probability of pay-off depending on their ability to interpret market signals of quality (of goods, services, and individual traders) [8]. Thus, a cybercriminal’s ability to succeed or make profits may depend on their location in the network, which is measured by centrality. Examination of Russian malware writers noted that individuals with higher technical skills were more centrally located [16]. Simultaneously, Dupont examined a co-offending network of 10 cybercriminals and noted that the more popular criminal did not control the most botnets [11]. From an enforcement perspective, focusing on degree central criminals is efficient in the former case but not in the latter. Examining the correlation between various centrality measures on underground forums would illuminate the structural properties of the market and thereby inform deterrence measures [29].

3 Analyses

This study investigates the nature of organization in cybercrime as it manifests on underground forums. We analyze three underground forums: BlackhatWorld (BW), Carders (CC), and L33tCrew (LC) (Table 1).

To analyze the forums, we model the private message interactions of a forum as a weighted directed graph, $G = (V, E)$, where each node, $v \in V$, is a member of the forum, each edge, $e = A \rightarrow B$, is a non-trivial and non-administrative private message from member A to member B and weight w_{AB} is the edge weight

Table 1. Summary of forums

Forum	Language	Date covered	Users	Users with private msg	Banned users
BlackhatWorld	English	08/2005-03/2008	8718	1690 (19.38 %)	43
Carders	German	02/2009-12/2010	8425	4290 (50.92 %)	1849
L33tCrew	German	05/2007-11/2009	18834	7687 (40.81 %)	913

denotes the number of messages sent from A to B . We remove the administrative and automated messages from the private messages. If a member only had administrative messages during his entire time in the forum, that member is also removed from the network. The resulting graph is used as the social network of a forum in the following analyses.

3.1 Analysis 1: Identifying Communities

Our goal for this study is to see whether or not distinct communities exist within a forum and compare topics among these communities.

Methodology: The main challenge to cybercriminal organization is lack of trust. Trust may not scale beyond Dunbar limits; thus, as a forum gets larger it may begin to fragment into distinct communities. To find these communities we use the Louvain method which is a fast heuristic approach based on modularity optimization [4]. Modularity of a network is the fraction of the edges that fall within the given groups minus the expected fraction if edges were distributed at random. The range of values for modularity is $[-1, 1]$. Networks with high modularity have dense connections between the nodes within modules but sparse connections between nodes in different modules.

Distinct communities may be similar or different. For example, different communities may compete for the same cybercrime; alternatively, they may specialize to leverage comparative advantage [17]. We use topic modeling to examine whether different communities specialize. We apply Latent Dirichlet Allocation (LDA) [3] on the private messages of community members to discover topics of their discussion and rank the topics based on their occurrence.

Results: The largest forum, L33tCrew, has the smallest number of non-trivial communities, but most of the communities on L33tCrew are larger than the other two forums (Table 3). The size of the communities on BlackhatWorld is smaller than Carders and L33tCrew and the communities are well separated (high modularity score) compared to that of the other forums (Table 2). We suspect this is because BlackhatWorld was less mature than the other two forums when the data was collected. Every forum has some common topics, usually the payment method or method of communication (Table 4). For example, the members of Carders use ukash/Paysafecard whereas on BlackhatWorld members use paypal. On L33tCrew the most common media for communication are ICQ/Jabber, but on BlackhatWorld members use Aim, Yahoo! and MSN instant messaging services. Details of the results are explained in the following subsections.

Table 2. Network structure of the forums. Here, ACC = Avg. Clustering Coeff. and LCC = Largest Connected Component. The # *communities* column shows the number of communities with the number of trivial communities (that have less than 4 members) are shown within bracket. We found that as forums get larger, the number of large communities decreases.

Forum	Density	ACC	LCC	# communities	Modularity	Largest community
BlackhatWorld	0.002	0.052	943	18 (+4)	0.46	212
Carders	0.003	0.103	2923	14 (+13)	0.29	800
L33tCrew	0.003	0.108	6116	8 (+16)	0.28	2348

BlackhatWorld: On BlackhatWorld, 1620 members, out of 8718, participated in private message interaction. The Louvain method discovers 22 communities with modularity score 0.46. 18 of the communities have at least 4 members. The largest community in BlackhatWorld has 212 members. All the communities have similar structures: two-tier organization with a few central members and the majority of the members are connected to the central members. Every

Table 3. Size and special topics of the communities.

C #	BlackhatWorld		Carders		L33tCrew	
	Memb	Special topic	Memb	Special topic	Memb	Special topic
1	212	Video upload	800	Drugs	2348	Cardable shops
2	203	Blogger generator	527	Gametimecards	1696	Anonymity services
3	142	Ebook	375	WebMoney	1447	Apple devices
4	138	Account creators	352	Bots	1419	Crypter
5	104	Invites	311	Packstation	393	Tickets
6	99	Keyword stuffing	284	Fake packstation	198	Accounts
7	97	Xrumer	253	Video game	116	Perfume
8	93	Article generator	245	ATM skimmer	35	Trojans
9	90	Account creators	237	Cardable shops		
10	81	Torrents	231	VPN, Wii		
11	79	Fantomaster	212	VPN		
12	77	Bulk email	197	Trojan		
13	60	Cloaking	111	Gamekeys		
14	59	Adsense	124	Jabber		
15	47	Cracked tools				
16	46	Stumblebot				
17	39	Tutorials				
18	16	Script				

Table 4. Common topics of the forums

BlackhatWorld	Carders	L33tCrew
Payment method (Paypal)	Payment method (PSC, Ukash, WMZ)	Payment (PSC, euro, WMZ)
Contact (AIM, Yahoo!)	Contact (ICQ, Jabber)	Contact (ICQ)
Blackhat seo tools	Carding, Stolen accounts	Carding, Stolen accounts
Make money online		

community has some special topics, for example, community 1 trades tools for automatic video uploading and CAPTCHA solving (Table 3).

Carders: On Carders, we found 27 communities with 0.29 modularity, out of which 14 communities have more than 100 members. Our result shows that smaller communities tend to have one central node and show a two-tier hierarchy (Appendix A). The largest community with 800 members has several central members instead of just one. The topics of this community are more varied than the other communities. These topics include selling Apple products (iPhone, iPad, macbook), crypting services and drugs, for example, MDMA. Other communities have their own specialized topics. For example, community 10 trades VPN services and handheld devices like Wii and iPod. Interestingly, although many communities sell similar types of products like drugs and accounts, there are differences in the actual product being traded, for example, community 1 sells ephedrone (ephe) and diazepam but community 5 sells Viagra.

L33tCrew: On L33tCrew, 7687 members participated in private message interactions. The Louvain method found 24 communities with modularity score 0.28, out of which 8 communities had at least 4 members. Communities in L33tCrew are much larger than BlackhatWorld and Carders. Some communities specialize in specific topics (Table 3), for example community 1 trades cardable shops list (online stores that accept stolen cards), stealer (malware for stealing accounts) and fake packstation.

3.2 Analysis 2: Identifying Central Members

Centrality measures enumerate distinct properties, i.e. each measure represents a separate notion of the node’s importance in the network. The *degree centrality* of a node indicates the total number of edges that connect it to other nodes. Degree central cybercriminals exude higher trust to peers [23] and receive higher responses to public posts [26]. *Betweenness centrality* enumerates the number of shortest paths that pass through a node. On IRC chat rooms, individuals with high betweenness centrality have access to more information both quantitatively and in terms of diversity [7]. Finally, *closeness centrality* indicates how far a node is from every other node in the network. High closeness centrality may lower transaction costs by reducing the number of overall transactions for a specific cybercriminal [25]. These centrality measures examine direct connections.

Eigenvector centrality indicates the importance of indirect connections by examining both the popularity of a node and the popularity of their connections [5]. Criminals with high eigenvector centrality may indicate leadership [5]. If centrality measures correlate it would indicate that the same criminals that exude higher trust also enjoy other advantages such as lower cost and access to higher quality information.

Methodology: We use Networkx [14] to compute six centrality measures (CM) on the social networks of the forums: degree (D), in degree (ID), out degree (OD), closeness (C), betweenness (B), and eigenvector centrality (E). We calculated the correlations between the various centrality measures for all three forums using SciPy statistics package [19] and report the Spearman’s ρ in Table 5.

Table 5. Intercorrelation (Spearman’s ρ) between the centrality measures, ranges from -1 to 1 where 1 indicates perfect positive correlation. Here E = Eigenvector, C = Closeness, B = Betweenness, ID = In-degree, OD = Out-degree, and D = Degree centrality. On BlackhatWorld all the centrality measures are positively correlated which means that some cybercriminals were simultaneously popular (degree), closer to other nodes (closeness), connected to other popular criminals (eigenvector) and had a higher proportion of shortest path going through them (betweenness). On Carders and L33tCrew, all but closeness centrality are positively correlated.

	BlackhatWorld					Carders					L33tCrew				
Cent	C	B	ID	OD	D	C	B	ID	OD	D	C	B	ID	OD	D
E	0.08	0.66	0.81	0.50	0.71	-0.43	0.79	0.91	0.62	0.77	-0.55	0.85	0.95	0.84	0.91
C		0.33	0.18	0.51	0.37		-0.19	-0.33	-0.11	-0.21		-0.39	-0.51	-0.35	-0.41
B			0.81	0.84	0.88			0.90	0.83	0.90			0.91	0.92	0.94
ID				0.56	0.85				0.71	0.88				0.88	0.96
OD					0.87					0.94					0.96

Results: Spearman’s ρ assesses how well the relationship between two variables can be described using a monotonic function. All the correlations were statistically significant for $p < 0.001$; however, the degree of correlation differs, as is evident from the ρ values that range from 0.08 to 0.96. Thus, some cybercriminals were simultaneously popular (degree), had a higher proportion of shortest path going through them (betweenness), closer to other nodes (closeness), and connected to other popular criminals (eigenvector).

Centrally located criminals have competitive advantage, e.g. through better access to market signals [7]. All four centrality measures were highly correlated across all forums. In addition the distribution of centrality was highly skewed, i.e. a few nodes had high centrality, while most were peripheral. This indicates that a majority of cybercriminals may receive a lower volume of responses to their posts [26], find it difficult to collaborate with the most technically adept cybercriminals [11] and have less access to quality information [7]. Thus, they are likely to be ripped off, with a handful of centrally located individuals who enjoy high profits and low transaction costs.

3.3 Analysis 3: Impact of Enforcement

Underground forums are policed by moderators and admins who enforce forum rules by issuing warnings and banning users when these rules are violated. For example, users can be banned for spamming or having multiple accounts [26]. It is unclear if banning users has any impact on the functioning of the network, positive or negative. It has been noted that joining these forums is free [2] and cybercriminals often have duplicate accounts [1], in fact having duplicate accounts is the most frequent reason for individuals being banned [26]. After getting banned, banned users either simply rejoin the forum or use a potentially undetected duplicate account. Here we investigate the change in network topology when misbehaving nodes are removed and contrast it with the change witnessed due to the regular churn of users in the forum.

Methodology: For each banned user u_i , we calculate the corresponding node centrality, specifically betweenness, closeness, degree and eigenvector. Since the success of a network often corresponds with the *small world* characteristics [2, 31], we examine the change in average clustering coefficient (ACC) and average path length (APL) respectively; for disconnected graphs we consider the APL of the largest connected component. For each user u_i we construct two graphs G_{ib} and G_{ia} : G_{ib} from all of the private messages sent between all users in the 30 days before u_i was banned and a graph G_{ia} from all of the private messages sent between all users in the 30 days after u_i was banned. When multiple users are banned in the same time period we model them as one node. Thus, all the messages to and from all banned nodes are assigned to one node entity. If one banned node sends a message to another node banned in the same period, it would manifest as a loop in our graph. We calculate the centrality scores for u_i on G_{ib} ; we also compute ACC and APL on both G_{ib} and G_{ia} . ΔACC is given by $ACC_{ia} - ACC_{ib}$ and ΔAPL by $APL_{ia} - APL_{ib}$. We compute the correlations between network metrics, ΔACC and ΔAPL , and centrality measures to examine whether removing more central offenders has a higher impact.

Finally, ΔACC and ΔAPL should be significantly different when users are banned as compared to the change observed due to periodic churn in the underground forum. We partition the graph data into 30 day snapshots for the entire duration of the dataset. We compute the change in ACC and APL for these snapshots to get a vector of ΔACC_r and ΔAPL_r . We use Wilcoxon Test, a non-parametric test to compare the difference in means, to contrast the difference between ΔACC and ΔACC_r as well as ΔAPL and ΔAPL_r . The analysis is conducted using Networkx and R.

Results: We calculated the correlations between small world metrics, ΔACC and ΔAPL , and the various centrality measures (the Spearman's ρ is reported in Table 6). In general the results for BlackhatWorld and L33tCrew are not significant. This may be the result of fewer data points for those forums compared to Carders. BlackhatWorld only banned 43 members overall, while all the banning on L33tCrew happened in the last three months of its operational lifetime.

Table 6. Intercorrelation (Spearman’s ρ) of the centrality measures with ΔACC , ΔAPL .

CM	BlackhatWorld		Carders		L33tCrew	
	ΔACC	ΔAPL	ΔACC	ΔAPL	ΔACC	ΔAPL
Betweenness (B)	-0.39	0.32	-0.12***	-0.05*	-0.05	0.11
Closeness (C)	0.07	-0.12	-0.07**	-0.05*	-0.19*	0.11
Degree (D)	-0.15	0.22	-0.19***	-0.03	-0.06	0.10
Eigenvector (E)	0.07	-0.12	-0.14***	-0.04	-0.01	0.004
p-value: 0.05 > * > 0.01 > ** > 0.001 > ***						

For Carders, which banned 22% of its members, betweenness as well as closeness centrality correlated with small world characteristics ($p < 0.05$). Thus, banning individuals who can propagate information over shorter paths may be better for reducing trust in the network. From a deterrence perspective a potential solution for law enforcement is to hijack the accounts of cybercriminals with higher closeness/betweenness centrality to spread noise on the forum.

We compared the mean values for ΔACC and ΔAPL , for when users get banned vs. the regular churn in the network. The change in small world characteristics for all forums were the same for banned members as for the regular churn (p -value $\gg 0.05$). Thus, it appears that individuals currently being banned are not close to other nodes.

4 Discussion and Conclusion

In this paper we examine and evaluate the ‘organization’ of cybercriminals as it manifests on underground forums. Research on cybercrime often presupposes organization [24]. The nature and purpose of this organization, however, is seldom examined.

We noted the presence of distinct communities despite the focused nature of the forums. We found that smaller communities organize in a two-tier hierarchy akin to a gang and are limited in size to Dunbar number; larger communities can have thousands of members, manifest a multi-tiered complex hierarchy, and specialize in a more diverse portfolio of cybercrimes compared to smaller cohorts. We observed that some cybercriminals simultaneously had lower transaction costs, access to better information, and higher visibility in the network. Then it is likely that if law enforcement targets only the central members, it would both lower the overall profits and reduce trust within the carding community. Finally, we found that the impact of banning misbehaving cybercriminals is similar that of the periodic churn of the forum.

There are obvious limitations of this research in terms of generalization. The differences noticed between BlackhatWorld and the German carding forums might be an effect of localization. Future efforts need to repeat these analyses on

additional data sets of both specialized forums dedicated to specific topics and other general purpose underground forums. It is also important to examine the temporal development of trust and organization in these communities. Finally, given trust is a key element for the stability of the forums, it would be illuminating to investigate the strategic creation and positioning of fraudulent sybils to target the sustainability of these forums.

Acknowledgment. We thank the anonymous reviewers and our shepherd Jens Grossklags for their valuable feedback. We are grateful to Damon McCoy for providing us access to the dataset. This work is supported by Intel through the ISTC for Secure Computing and National Science Foundation CNS-1347151.

A Example of Social Networks

Figures 1 and 2 show the structure of the communities in Carders. Here, nodes are scaled according to their degree centrality.



Fig. 1. The largest community of Carders does not have any one central big node.



Fig. 2. Three communities of Carders: Community 12 (purple), 13 (green), 14 (brown)(Color figure online).

References

1. Afroz, S., Caliskan-Islam, A., Stoleran, A., Greenstadt, R., McCoy, D.: Doppelgänger finder: Taking stylometry to the underground. In: IEEE Symposium on Security and Privacy. IEEE (2014)
2. Afroz, S., Garg, V., McCoy, D., Greenstadt, R.: Honor among thieves: a commons analysis of cybercrime economies. In: eCrime Researcher's Summit. APWG, IEEE (2013)
3. Blei, D.M.: Probabilistic topic models. *Commun. ACM* **55**(4), 77–84 (2012)
4. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *J. Stat. Mech. Theory Exp.* **2008**(10), P10008 (2008)

5. Bonacich, P.: Technique for analyzing overlapping memberships. *Sociol. Methodol.* **4**, 176–185 (1972)
6. Brenner, S.W.: Organized cybercrime-how cyberspace may affect the structure of criminal relationships. *North C. J. Law Technol.* **4**(1), 1–50 (2002)
7. Décary-Héту, D.: Information exchange paths in irc chat rooms. In: Morselli, C. (ed.) *Crime and Networks*, pp. 218–230. Taylor & Francis Group, New York (2014)
8. Décary-Héту, D., Leppänen, A.: Criminals and signals: An assessment of criminal performance in the carding underworld. *Security Journal* (2013)
9. Dunbar, R.: You’ve got to have (150) friends. *The New York Times*, The Opinion Pages (2010)
10. Dunbar, R.I.: Neocortex size as a constraint on group size in primates. *J. Hum. Evol.* **22**(6), 469–493 (1992)
11. Dupont, B.: Skills and trust: a tour inside the hard drives of computer hackers. In: Morselli, C. (ed.) *Crime and Networks*, pp. 195–217. Taylor & Francis Group, New York (2014)
12. Finckenauer, J.O.: Problems of definition: what is organized crime? *Trends Organized Crime* **8**(3), 63–83 (2005)
13. Franklin, J., Perrig, A., Paxson, V., Savage, S.: An inquiry into the nature and causes of the wealth of internet miscreants. In: *ACM Conference on Computer and Communications Security*, pp. 375–388 (2007)
14. Hagberg, A., Swart, P., S Chult, D.: Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL) (2008)
15. Herley, C., Florêncio, D.: Nobody sells gold for the price of silver: dishonesty, uncertainty and the underground economy. In: Moore, T., Pym, D., Ioannidis, C. (eds.) *Economics of Information Security and Privacy*, pp. 33–53. Springer, US (2010)
16. Holt, T.J., Strumsky, D., Smirnova, O., Kilger, M.: Examining the social networks of malware writers and hackers. *Int. J. Cyber Criminol.* **6**(1), 891–903 (2012)
17. Hunt, S.D., Morgan, R.M.: The comparative advantage theory of competition. *J. Mark.* **59**(2), 1–15 (1995)
18. Jennings, W.P.: A note on the economics of organized crime. *East. Econ. J.* **10**(3), 315–321 (1984)
19. Jones, E., Oliphant, T., Peterson, P.: *Scipy: Open source scientific tools for python* (2001). <http://www.scipy.org/>
20. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G.M., Paxson, V., Savage, S.: Spamalytics: an empirical analysis of spam marketing conversion. In: *Proceedings of the 15th ACM Conference on Computer and Communications Security*, pp. 3–14. ACM (2008)
21. McCoy, D., Pitsillidis, A., Jordan, G., Weaver, N., Kreibich, C., Krebs, B., Voelker, G.M., Savage, S., Levchenko, K.: Pharmaleaks: understanding the business of online pharmaceutical affiliate programs. In: *Proceedings of the 21st USENIX Conference on Security Symposium*, pp. 1–16. USENIX Association (2012)
22. Meyer, G.R.: The social organization of the computer underground. Technical report, DTIC Document (1989)
23. Monsma, E., Buskema, V., Soudijn, M., Nieuwbeerta, P.: Partners in Cybercrime. An Online Cybercrime Forum Evaluated From a Social Network Perspective. Ph.D. thesis, Thesis in Sociology and Social Research, Universiteit Utrecht, Netherlands (2010)
24. Moore, T., Clayton, R., Anderson, R.: The economics of online crime. *J. Econ. Perspect.* **23**(3), 3–20 (2009)

25. Morselli, C., Giguère, C., Petit, K.: The efficiency/security trade-off in criminal networks. *Soc. Netw.* **29**(1), 143–153 (2007)
26. Motoyama, M., McCoy, D., Levchenko, K., Savage, S., Voelker, G.M.: An analysis of underground forums. In: Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, pp. 71–80. ACM (2011)
27. Peretti, K.K.: Data breaches: what the underground world of carding reveals. *Santa Clara Comput. High Technol. Law J.* **25**(2), 375–413 (2008)
28. Sparrow, M.K.: The application of network analysis to criminal intelligence: an assessment of the prospects. *Soc. Netw.* **13**(3), 251–274 (1991)
29. Xu, J.J., Chen, H.: Crimenet explorer: a framework for criminal network knowledge discovery. *ACM Trans. Inf. Syst. (TOIS)* **23**(2), 201–226 (2005)
30. Yip, M., Shadbolt, N., Webber, C.: Structural analysis of online criminal social networks. In: IEEE International Conference on Intelligence and Security Informatics (ISI), pp. 60–65. IEEE (2012)
31. Yip, M., Shadbolt, N., Webber, C.: Why forums?: an empirical analysis into the facilitating factors of carding forums. In: Proceedings of the 5th Annual ACM Web Science Conference, pp. 453–462. ACM (2013)
32. Yip, M., Webber, C., Shadbolt, N.: Trust among cybercriminals? Carding forums, uncertainty and implications for policing. *Policing Soc.* **23**(4), 516–539 (2013)