How cybercriminal communities grow and change: An investigation of ad-fraud communities

Jean-Loup Richet

► To cite this version:

Jean-Loup Richet. How cybercriminal communities grow and change: An investigation of ad-fraud communities. Technological Forecasting and Social Change, 2022, 174, 10.1016/j.techfore.2021.121282 . hal-03403899

HAL Id: hal-03403899

<https://hal.science/hal-03403899>

Submitted on 26 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How cybercriminal communities grow and change: An investigation of ad-fraud communities

Jean-Loup Richet^{a,*}

IAE Paris - Sorbonne Business School, Université Paris 1 Panthéon-Sorbonne, 12 Rue Jean Antoine de Baïf, 75013 Paris, France

ARTICLE INFO

Keywords:

Ad-fraud
Cybercrime
Online communities
Multimethod approach

ABSTRACT

Firms spend enormous resources on digital advertising and promoting their brand online. In the meantime, ad-fraud undertaken by cybercriminals cost \$42 billion in 2019 and could reach \$100 billion by 2023. However, while digital advertisers continue to wrestle with how to effectively counteract ad-fraud, the topic of advertising fraud itself has received little academic attention. Here, we investigate this gap between practice and research through an exploration of ad-fraud communities. Our research implemented a multimethod approach for data collection in a longitudinal (18 months, October 2017 to April 2019) online investigation of this phenomenon. Integrating qualitative and quantitative analysis, we examined (1) internal interactions within ad-fraud communities and (2) ad-fraud communities' performance and growth. Our online investigation extends our conceptual understanding of ad-fraud and explains how ad-fraud communities innovate. Our findings indicate that capabilities enacted by some communities foster requisite variety and enable the coordination of complex, iterative, and incremental dynamics (cocreation of artificial intelligence-based bots, customer involvement, and reinforcing capabilities). This research has both theoretical and practical implications for innovation in cybercriminal communities. Furthermore, we provide practical guidance for policy-makers and advertisers regarding how to improve their response to business threats. Indeed, a better understanding of how ad-fraud communities innovate enables organizations to develop countermeasures and intelligence capabilities.

1. Introduction

Advertising fraud has been recognized as one of the main challenges faced by digital advertisers (Perrin, 2019; White and Samuel, 2019). Ad-fraud is indeed growing each year: in 2019, advertisers lost \$42 billion in ad spending to fraudulent activities committed via online, mobile, and in-app advertising (Barker, 2019). Ad-fraud is planned and organized in the dark net's cybercriminal communities¹ (Benjamin et al., 2019), online places of gathering for cybercriminals specializing in ad-fraud. However, despite the significant consequences of ad-fraud for practitioners, there is almost no theoretical work and very little empirical evidence concerning the complex mechanisms and origins of ad-fraud communities (Kraemer-Mbula et al., 2013). While digital advertisers continue to wrestle with how to effectively counteract those committing ad-fraud, advertising fraud research is still in its infancy (Lamberton and Stephen, 2016, p. 158).

Early in our investigation of ad-fraud communities, we discovered

that such communities vary greatly with respect to the ways in which they address criminal needs and use innovative technologies, such as artificial intelligence (AI). Our research examined six ad-fraud communities with the aim of (1) assessing trends in ad-frauds and (2) identifying how ad-fraud communities perform and grow.

This article addresses important gaps in knowledge regarding ad-fraud communities. First, we examined how some ad-fraud communities develop innovative capabilities to improve customer growth, reinforce user-led innovation, and generate network externalities. In these communities, new developments in AI include the cocreation of AI-based bots to conduct large-scale ad-frauds campaigns, dynamic SEO frauds using AI, or deep faking Hollywood stars for the purpose of online ads. We examined how these developments in AI change the way the community performs and reconfigures itself. Second, we present the first empirical assessment of the dynamics, organization, and performance of ad-fraud communities. This endeavor speaks to a need for additional research on cybercrime (Benjamin et al., 2019; Kraemer-Mbula et al.,

* Corresponding author.

E-mail address: jean-loup.richet@iae.pantheonsorbonne.fr.

¹ A 'dark net' refers to a network within the internet that requires specific access authorization, configuration, or software – the most famous dark net is the Tor network (that requires the use of a dedicated software). It is opposed to the 'clear web', the publicly accessible Internet.

2013; Saridakis et al., 2016).

Our research also has strong policy implications: we highlight the social impact of cybercrimes and suggest policy initiatives to deter ad-fraud. We provide practical guidance for policy-makers and advertisers regarding how to implement strategic threat intelligence to better identify innovations in ad-fraud and improve responses to business threats. Indeed, a better understanding of how ad-fraud communities innovate enables organizations to detect emerging risks, identify potential targets, develop countermeasures and better control ad-fraud communities (Benjamin et al., 2019; Kraemer-Mbula et al., 2013).

We begin by identifying the theoretical/practical gap and a real-world issue, highlighting why advertising fraud is a significant problem for digital advertisers. Next, we develop our conceptualization of ad-fraud communities. We then describe our data collection and measures of community performance (in the form of community growth, popularity and innovation in ad frauds). The empirical section presents the results from our investigation. We then develop a theoretical elaboration from our empirical observations. Finally, we discuss the implications and limitations of our research and identify several avenues for future research in information systems.

2. Conceptual background: opening the blackbox of ad-fraud communities

Previous research in digital marketing has focused primarily on one type of ad-fraud: click fraud. This involves scripts or robots dedicated to (1) clicking on competitors' ads to make them exceed their daily advertising budget; or (2) clicking on a given website's ads so that the website owner generates ad revenues (pay per click). Research has been aimed at understanding how click fraud affects search engines' revenues (Wilbur and Zhu, 2009), online affiliate marketing revenues (Edelman and Brandi, 2015), and publisher networks (Asdemir et al., 2008). The information systems literature also focuses on cybercriminal behaviors (Hui et al., 2017) and improving the detection of botnets that could be used in various cybercriminal activities, including click fraud (Chen et al., 2017). Past research has emphasized the vulnerabilities of auction mechanisms in the fight against click fraud (Agarwal et al., 2009), the indirect benefit of delegating ad-fraud investigation to third parties (Min Chen et al., 2015), and the need to shift toward pay-per-action to deter click fraud (Nazerzadeh et al., 2013). However, very few studies have explored how ad-fraud communities organize, advertise, market their criminal businesses, and innovate, necessitating more research on cybercrime (Benjamin et al., 2019; Kraemer-Mbula et al., 2013).

Thus far, it appears that the academic community has been somewhat disinterested in the problem of ad-fraud community development (Lamberton and Stephen, 2016) while paradoxically highlighting tensions among accuracy, fraud and ethics inherent in the online ad industry (White and Samuel, 2019).

At the beginning of the 2000s, ad-fraud was one of the primary issues preventing practitioners from moving toward digital marketing (Advertising Age, 2006). Where are we now? A recent digital advertising report highlighted that fraud was perceived as the primary hindrance to ad budget growth by US digital media professionals (Perrin, 2019). Ad-fraud remains a significant problem faced by practitioners. Indeed, while US digital advertisers' spending increased by 10% in 2017, totaling \$72.5 billion, economic losses due to bot-based ad-fraud reached \$6.5 billion in the US that same year (ANA, 2017). However, these numbers do not take into account ad-fraud on Facebook and Google or fraud on display ads. In fact, in the US, nearly 20% of the total digital ad expenditure is wasted every year due to ad-frauds (Perrin, 2019). Research on worldwide ad spending forecasts that in 2023, advertisers will lose \$100 billion to fraudulent activities committed via online, mobile, and in-app advertising, based on the \$35 billion lost in 2018 and the \$42 billion in 2019 (Barker, 2019). Clouding the picture, given that ad-fraud involves reduced effort and risk but greater reward for fraudsters, ad-frauds have recently increased in both number and

scale and will continue to grow (ANA, 2017). The ad-fraud landscape is changing, as ad-fraud communities are becoming more sophisticated and organized (Benjamin et al., 2019; Hughes et al., 2017; Richet, 2013), following industry trends such as blockchain (Chang et al., 2020), big data (Kwon et al., 2015), Internet of Things and autonomous systems (Santoro et al., 2017; Shareef et al., 2021). The use of artificial intelligence bots is likely to drive future large-scale fraud, as bots are getting better at impersonating legitimate users' behaviors and circumventing fraud controls. Indeed, ad-fraud communities have adapted themselves to industry trends and adopted digital business models entirely dedicated to online value-added criminal activities or services, including the design and development of malware, online advertising fraud, massive distributed denial of service (DDoS) attacks, and bot networks for hire.

Ad-fraud communities are loosely coupled, favoring weak ties, temporary projects, anonymity, and technical expertise and have a transnational dimension (Brenner, 2002). Ad-fraud communities need convergence settings (Soudijn and Zegers, 2012), online meeting places where cybercriminals can exchange information, knowledge, and expertise, find new business partners, or apply for cybercrime job offers. Convergence settings facilitate cybercriminal cooperation, enabling it to persist, and ease interactions within the criminal network. The academic literature on organized cybercrime has highlighted the lack of conceptualization of online gathering places (Leukfeldt, 2014; Soudijn and Zegers, 2012). We suggest that ad-fraud communities are online gathering places that facilitate cybercriminal endeavors. Such communities need to develop specific capabilities (customer relationships and management, marketing, experimentation) to maintain cohesion and productivity. Indeed, ad-fraud communities engage in coproduction, requiring distributed forms of organization; they enable the interactions of a large and diverse number of stakeholders (Leukfeldt, 2014); communities can tap into the innovative capabilities of their members to nurture their growth and performance (Roma and Vasi, 2019).

We intend to investigate this gap between practice and research through our exploration of ad-fraud communities that specialize in ad-fraud. Our research examined six ad-fraud communities with the aim of (1) *assessing trends in ad-frauds* and (2) *identifying how ad-fraudsters communities perform and grow*. In these communities, we observed new developments in AI, such as the cocreation of AI-based bots to conduct large-scale ad-fraud campaigns, dynamic SEO fraud using AI, or deep faking Hollywood stars for the purpose of online ads; we assessed how some ad-fraud communities develop innovation capabilities to improve customer growth, reinforce user-led innovation, and generate requisite variety.

The following section describes the way in which we collected data on the clear web and on the dark net and how we measured and analyzed each ad-fraud community's characteristics. Then, we present the results from our online investigation and discuss the theoretical implications and limitations of our research.

3. Research methodology

Scholars have highlighted the need for alternative methodological approaches (Agerfalk, 2013; Benjamin et al., 2019). Herein, we adopted a novel method that integrates both quantitative and qualitative assessment, collecting a large volume of data on the characteristics of ad-fraud communities (1,294,861 threads, 12,245,156 messages, and 775,888 backlinks) and qualitatively analyzing samples (observations of conversations) from our dataset to triangulate the quantitative assessment (Walsham, 2006).

Indeed, our research uses a multimethod approach for data collection in a longitudinal (18 months, October 2017 to April 2019) online investigation of this phenomenon. This approach provided us with a unique opportunity to examine ad-fraud community performance, as we gathered a rich dataset of conversations from six ad-fraud communities and established our work in this novel context (Birks et al., 2013;

Walsham, 2006). This approach enabled us to make empirical observations about ad-fraud communities—online marketplaces and communities where cybercriminals exchange and discuss new techniques in ad-frauds, scams, and spam. Ad-fraud communities refer to themselves as *black hat* search engine optimization (SEO) communities. A black hat as opposed to a white hat hacker is a cybercriminal who violates computer security to destroy, modify, or steal data. *Black hat SEO* means the community focuses on ad-fraud (scamming search engines with fraudulent publishing practices but also spam, fraudulent impressions, clicks, conversions, etc.). Such online forums aim to attract sellers and buyers of ad-fraud services (from beginner cybercriminals to more organized temporary hacking groups). An expert panel from CEPOL and EUROPOL helped us select six criminal communities (see Table 1 for more details on each community). Our selection included four communities from the clear web and two from the dark net. Communities were accessible from both the clear web (through a search on Google, requiring registration) and the dark net (using the Tor network). Access to Tor requires the installation of software; access to ad-fraud communities requires knowledge of their addresses—those with the .onion suffix—registration, and eventual assessment of registrants.

We performed systematic and weekly data collection over a period of 18 months from October 2017 to April 2019 to assess development of each community. Communities log all user exchanges, so we were able to observe past and current discussions and exchanges from 2007 to 2019. We used a Python script to automatically scrape and duplicate social exchanges taking place in these six communities. A total of 1 million and 294,861 html pages resulted from this data collection, representing 1,294,861 threads scraped. The data set included chat room interactions, discussions, and messages sent and received ($n=12,245,156$). We were able to assess metrics such as threads, posts, and new members per day during the period covered (see Table 1). During our period of observation, we focused on building a subdataset of interactions on innovation and value creation. We captured relevant exchanges that we observed ($n=1150$) and documented field notes ($n=640$) from online observations as screenshots. We also collected backlinks related to these six ad-fraud communities ($n=775,888$) and analyzed them to measure cybercriminal community brands and popularity in the cybercriminal market (see Table 2). Observation data were triangulated with community characteristics (volume of posts and threads, number of new members, member activity, etc.) and backlinks (search engine metrics, volume of backlinks, their characteristics and origins). Moreover, we compared our online observations to external data provided by EUROPOL and INTERPOL research on cybercrime. Finally, the primary researcher on this project is an expert for EUROPOL: his in-depth knowledge of (criminal) organizational codes was

extremely useful as we tried to make sense of the empirical data during analysis and was essential for establishing the internal validity of our research (Baskerville and Myers, 2015; Patton, 2005).

In terms of data analysis, our research unfolded through multiple incremental and iterative phases. First, we read through a large number of threads and messages, making general observations to gain an understanding of each community's philosophy, culture, and practices (Baskerville and Myers, 2015; Schultze, 2000). We created descriptive and analytical codes with the aim of developing an understanding of the characteristics of ad-fraud communities and the way they changed over time. Second, we narrowed our focus: we implemented selective coding as we set out to identify and document concepts related to empirical indicators (customer interaction and the components of community performance such as community growth, popularity, user-led knowledge creation on AI and innovation in ad-fraud) with emerging categories. Examining our data, we identified different types of relationships among the concepts that emerged (Charmaz, 2006). This process led us to develop the following intermediate inquiry (see Fig. 1 below):

Constant comparison between what was emerging from our data and existing theories enabled us to identify a theoretical gap regarding the nature of innovation in ad-fraud communities, which will be further detailed in the discussion section.

We present our findings in the next section, beginning with our assessment of the characteristics and performance of the ad-fraud communities we observed. We then describe the different types of ad-fraud. We end this empirical section with an account of the capabilities deployed by ad-fraud communities to achieve innovation. Finally, our theoretical elaboration is developed in the discussion section.

4. Findings from our empirical analysis

This section is organized as follows. In the first subsection, we explain how we determined market selection and market orientation and how we categorized communities based on the combination of their market selection and market orientation. In fact, we observed varying performance indicators (community growth, popularity, user-led knowledge creation on AI and innovation in ad-fraud) among the communities from our scope, depending on community selection (specialized or general) and community orientation (technical or customer oriented).

In the second subsection, we examine how each community fosters internal interactions (based on the evolution in the volume of threads and posts, triangulated with the analysis of the content exchanged (automated content vs real interactions)). In the third subsection, we observe community performance and growth (based on the evolution in

Table 1
Characteristics of the six ad-fraud communities observed.

Cybercriminal community	SPECTEC1	GENCUST1	SPECCUST	SPECTEC2	GENCUST2	GENTEC
Market selection	Specialized	General	Specialized	specialized	General	General
Market orientation	Technical	Customer	Customer	Technical	Customer	Technical
Creation date	2014	2012	2016	2014	2015	2007
# of threads collected	19,896	19,461	210,576	318,869	719,709	6350
# of posts collected	42,191	684,931	234,321	2,103,946	9,138,887	40,880
# of members	31,001	152,834	40,275	308,415	233,471	36,250
# of active members at the beginning of the period of coverage	235	2059	350	2087	3501	286
# of active members at the end of the period of coverage (baseline for ratio)	238	2304 (+11.9%)	356	2105 (+0.9%)	3899 (+11.4%)	281
	(+1.3%)		(+1.7%)			(−1.7%)
New members per day	2	31	3	11	47	3
Additional members during the period of coverage	912	16,987	1791	6032	25,524	1620
# of new threads (during the period of coverage)	1932	5433	4137	19,654	12,949	1370
Threads per day	4	10	8	36	24	3
# of new posts (during the period of coverage)	903	65,664	3010	15,711	133,736	607
Posts per day	2	120	5	29	244	1
Ratio # new threads/active member	8.1	2.4	11.6	9.3	3.3	4.9
Ratio # new posts/active member	3.8	28.5	8.5	7.5	34.3	2.2
Replies per thread	0.5	12.1	0.7	0.8	10.3	0.4
Location	clear web	clear web	clear web	dark net	dark net	clear web

Table 2

Popularity and brand awareness measures of the six ad-fraud communities observed.

Cybercriminal community # backlinks	SPECTEC1 1494	GENCUST1 14,554	SPECCUST 606,958	SPECTEC2 1078	GENCUST2 143,940	GENTEC 7864
Deleted backlinks (last 4 months)	1	790	1	85	20,199	98
Referring domain	10	1028	4	50	6949	85
Average backlinks/domain	149	14	151,740	22	21	93
Referring IP	10	916	4	38	5721	60
Average backlinks/IP	149	16	151 740	28	25	131
Referring subnets	10	808	4	34	4845	57
Average backlinks/subnet	149	18	151,740	32	30	138
Citation flow	22	41	46	17	44	20
Trust flow	0	16	0	1	17	0
Facebook Share	0	1828	0	917	1469	118
Backlinks with branded Anchor text	33%	12%	75%	28%	11%	37%

Fig. 1. A multimethod approach integrating qualitative and quantitative analysis.

the volume of active users and evolution of backlinks, triangulated with the analysis of the discussions on innovations and competition). In the final subsection, we highlight the differences among groups of communities by discussing differences in the way communities innovate. Indeed, specialized ad-fraud communities, such as SPECTEC1, SPECCUST, and SPECTEC2, provide a wealth of information on how to conduct traditional ad-fraud techniques and incremental improvements of these techniques, while general and customer-oriented communities such as GENCUST1 and GENCUST2 reveal emerging trends and new AI-based techniques.

4.1. Characteristics of ad-fraud communities: market selection and market orientation

The six ad-fraud communities we observed were communities aimed at attracting sellers and buyers of ad-fraud services (from beginner cybercriminals to more organized temporary hacking groups). Table 1 shows the key characteristics of each community.

The name of each community (SPECTEC, GENCUST, etc.) was anonymized by the research team with a pseudonym based on the combination of the community's market selection and market orientation (SPEC for *specialized* market selection; GEN for *general* market selection; TEC for *technical* market orientation; CUST for *customer* market orientation).

Market selection concerns the communities' targeting strategy—in

our case, the communities specified which groups of customers they targeted and their service domain (specialized or general). For example, some *specialized* communities focused on a certain type of ad-fraud (such as traffic fraud or click fraud). Specialized communities aimed to attract very specific and expert ad-fraud service providers (sellers) and buyers interested in these niche services (SPECTEC1, SPECCUST, SPECTEC2). Other communities were *general* communities aimed at gathering a large diversity of cybercriminals interested in fraud (including money launderers, click fraudsters, malware writers, botmasters, etc.) including GENCUST1, GENCUST2, GENTEC. Market selection criteria are frequently stated on the first home page of each community. Moreover, market selection is perceived through the architecture of the cybercriminal community. Indeed, general communities embed a large number of subcommunities loosely related to ad-fraud that appear in the community's categories and subforum boards: they range from *warez* (stolen and cracked software offered for free) to pornographic content and *how to* guides for making money from porn through malware development and diffusion.

Ad-fraud communities also adopt different **market orientations** (see Table 1). Half of the communities in our dataset (SPECTEC1, SPECTEC2, GENTEC) had *technical* and engineering orientations: they asked their potential customers to provide proof of their technical skills to register on the communities; they had large discussion sections on programming, scripting, and technical web design; discussions focused on technical topics, such as HTML, JavaScript, PHP, Perl, etc. The other communities

from our dataset (GENCUST1, SPECCUST, GENCUST2) had a strong customer orientation: they aimed to ease the use of the community for customers, with one-click registration, after sales service sections, sections for beginners, membership layers, discussion aimed at vulgarizing technical topics and technical slang, etc.

4.2. Internal interactions within ad-fraud communities

Threads in the context of Table 1 are initiated conversation threads on ad-fraud techniques, information on fraud, cracked software related to SEO, or ad-fraud guides offered for download, questions, debates, etc. Posts are discussions and replies to threads. The low number of threads compared to the number of posts could be explained by the frequency of pruning, i.e., yearly deletion of old threads and some posts—this was the case with GENCUST1, for instance (only 19,461 threads for a community created in 2012). Conversely, the huge number of posts on SPECTEC2 and GENCUST2 could be explained by a lack of pruning practices—GENCUST2, for instance, archives old threads in a special section but never deletes them (719,709 collected in total, for a community created in 2015). Thus, the total number of posts or threads is not an accurate indicator of a community's engagement with its cybercriminal customer base. Hence, we moved toward metrics related to internal engagement within each community—threads per day, posts per day, and corresponding ratios (threads and posts per active member, replies per threads). The number of new threads per active member hinted at a phenomenon we observed during our qualitative investigation. Some communities boast a huge number of new threads per active member. However, this is not an indicator of interaction. For instance, in the SPECTEC2 community, many threads go unanswered (19,654 new threads created in the 18-month period of coverage). Most of the threads were information diffusion or cracked SEO software and tools that do not necessitate a reply. The name of the most active cybercriminal was *Releases* (in fact, it is a bot that provides automated press releases and information on the cybercrime market).

In contrast, *replies per thread* were an interesting metric for internal engagement: the greater the number of threads, the more debates and discussions they generated. This could be an indicator of the internal engagement of active members and interactions. GENCUST1 and GENCUST2 are both customer oriented and general communities (aiming to attract a large diversity of criminal members interested in ad-fraud). In these communities, there were fewer threads per day per active member (10 and 24, respectively, compared to 36 for SPECTEC2, for instance). However, these threads triggered much more debate: GENCUST1 and GENCUST2 boasted higher numbers of replies per thread (12.1 and 10.3, compared to 0.8 for SPECTEC2 or 0.4 for GENTEC), a sign of intense internal activity.

4.3. Performance of ad-fraud communities

In most ad-fraud communities, activity is required to access VIP and special sale sections. Individual members are considered active if they have created at least one post in a period of 12 months. Because there was no pruning of nonactive members undertaken by the ad-fraud communities we observed, member activity could be considered a proxy for member retention and loyalty (if an individual member stopped posting, we hypothesized that he/she was no longer interested in engaging with peers or being part of the community). We observed that two communities that adopted both a customer orientation and general market selection (GENCUST1 and GENCUST2) exhibited a significant increase in active members (+11.9% for GENCUST1 and 11.4% for GENCUST2), while GENTEC (general market selection but technical orientation) presented a decrease in active members (−1.7%). The other communities had similar low baselines (+1.3% for SPECTEC1, +1.7% for SPECCUST, and +0.9% for SPECTEC2).

Customer-oriented ad-fraud communities (GENCUST1, SPECCUST, GENCUST2) showcased forms of market knowledge. A SPECCUST

member commented on June 20, 2017, “Most service providers I’ve seen here refuse to accept pharma, gambling, or adult sites. Check the replies in the thread but if you have no luck here there’s another forum called, like, [URL of GENCUST2] and the people there are the really dodgy pharma/hacked links types... Pretty much every topic on that forum deals with adult niches and stuff like that.” We found similar exchanges on GENCUST1 and GENCUST2 where members discussed competing communities. For instance, on GENCUST1, a discussion was found concerning the SPECTEC1 community:

“Yes, they did have problems with admin rights until a few days ago, but they fixed it for me at least. I have their premium package which I mainly got for Xrumer² and SeNuke³ and they work fine for me, so I am happy for now. Can’t comment on their VIP offer because I haven’t tried it.”

“Don’t even go there. The mods are all rude and unhelpful.”

“I am getting the same poor service. Admin ignore you or say they will sort it but never do. Do they expect people are just going to accept it.”

Collected backlinks (links from a referrer website to a referent) related to these six ad-fraud communities ($n = 775,888$) reinforced this finding on the differences between ad-fraud communities and how they developed. Backlinks (see Table 2 for their characteristics) are an interesting indicator of online popularity (the more referrals a community has, the more popular it is considered by search engines). Correlated with increases in active members, online popularity was considered an indicator of community performance. Indeed, we used attractiveness as a proxy for financial performance: the more customers (defined as sellers and buyers of ad-fraud services) a community has, the more it benefits from transaction fees.

When the values for the referring domain, referring IP, and referring subnets are the same, this indicates that each backlink comes from the same IP and subnet. This is a strong hint of a non-natural backlink strategy. The cybercriminal community cheated and most likely undertook a black hat SEO campaign to artificially create fake backlinks: this was the case for SPECTEC1 (10) and SPECCUST (4). Backlinks with many branded anchor texts could also indicate sponsored backlinks (buying thousands of backlinks to artificially improve search engine rankings). This was the case for SPECCUST (75%) and GENTEC (37%). On the other hand, few branded anchors may indicate a more organic backlinking strategy (meaning customers discuss the cybercriminal community on their blogs, websites, etc.). For that criterion, GENCUST1 (12%) and GENCUST2 (11%) were the most efficient. We observed that even communities only available on the dark net (SPECTEC2 and GENCUST2) are *shared* on Facebook and discussed on the clear web. This paradox (hidden communities advertising on the clear web) could be explained by the black hat SEO culture of ad-fraud communities. Indeed, tricking search engines is part of the ad-fraud community’s DNA. Even hidden communities boast about being the best at SEO and search engine fraud. The large number of deleted backlinks for GENCUST2 (14% of the total number of backlinks) suggests the use of black hat techniques for backlink creation, such as creating thousands of comment spams for temporary SEO boosts.

Citation and trust flow (Table 2) are ranking algorithms that are widely used in the SEO industry to assess the quality and quantity of backlinks. Citation flow accounts for how influential a URL might be based on the number of sites linked to it. This algorithm takes into account both the direct and indirect numbers of backlinks (backlinks of the referrer website). GENCUST1, SPECCUST and GENCUST2 were judged as the most influential communities from our dataset (higher citation flow). Trust flow aims to assess the quality of backlinks—it matches their quality with a paid database of 8 trillion backlinks (Jones, 2012) from

² Software that spams online forums and comment sections.

³ Software used to undertake large-scale mass emailing spam campaigns.

July 2013–November 2018 and assesses their proximity to a set of trusted and mainstream websites (.gov and .edu websites, Wikipedia, Google, etc.). The further away a given URL is from the seed set of known trusted sites, the less trustworthy it will be. A trust ratio (trust flow divided by citation flow) below 1 indicates a nontrustworthy website (backlinks from pornographic websites, infected and compromised websites, artificial links, etc.). SPECTEC1, SPECCUST, and GEN-TEC have a trust ratio of zero and SPECTEC2 of 0.06: this is an indicator of the very poor backlink quality of these websites (mostly artificial backlinks). GENCUST1 and GENCUST2 have a trust ratio of 0.39, meaning they have better backlinks than their competitors in terms of quality (perhaps more organic backlinks from other ad-fraud communities and customers). The following section summarizes our qualitative assessment of the presence of threads discussing competition (market knowledge) and innovations in ad-fraud in each community.

4.4. Incremental innovation vs innovation fueling engagement

4.4.1. Incremental innovations in specialized ad-fraud communities

Specialized ad-fraud communities, such as SPECTEC1, SPECCUST, and SPECTEC2, provide a wealth of information on how to conduct traditional techniques and incremental innovations in identity, attribution, and services fraud. We classified the large variety of ad-fraud that we observed into three categories: (1) identity fraud; (2) attribution fraud; and (3) ad-fraud services.

Identity fraud aims to impersonate real users and inflate audience numbers. Several ad-fraud techniques relate to this category and include traffic from bots (coming from a hosting company or a data center, or from compromised devices); cookie stuffing; falsifying user characteristics, such as location and browser type; fake social traffic (misleading users on social networks into visiting the advertised website); and the creation of fake social signals to make a bot look more legitimate, for instance by opening a Twitter or Facebook account. Specialized communities provide a large diversity of identity fraud resources and discussions on their trends. On May 1, 2018, a cybercriminal asked on SPECTEC2: “Are Social Signals that are fired using bots still working?” Other cybercriminal members replied, providing more details on the evolution of this technique: “Working as in, giving same results as organic social signals [...] Accounts that have very high interaction or user input, are given heavier authority/weight when sending a social signal to your site, whereas those accounts that have very little friends or rather ‘empty’ accounts, are quickly determined as the least priority given, thus they command very little weight when being used on your site. Accounts with many friends or circles, very active, have high interaction—those would be good sources of social signals.”

Attribution fraud aims to impersonate real users’ behaviors (clicks, activities, conversations, etc.). Multiple ad-fraud techniques belong to this category: hijacked devices and the use of infected users (through a malware) as part of a botnet to participate in ad fraud campaigns; click farms (companies where low-wage employees are paid to click or engage in conversations and affiliates’ offers); incentivized browsing; video placement abuse (delivered in display banner slots); hidden ads (that will never be viewed by real users); domain spoofing (ads served on a website other than the advertised real-time bidding website); and clickjacking (user is forced to click on the ad). SPECCUST has a particularly lengthy section on attribution fraud (click and affiliate frauds through the use of bots). On a sale listing created in April 2019, a criminal member asked whether the “bot for sale” was able to “go to like ‘google.com’ and search for your website on the given keyword to increase CTR in Google search console?” To which the sales technician replied, highlighting recent progress in attribution frauds enabled by the bot: “Indeed, it can search your keywords on search engines (Google, Yahoo, and Bing) and shopping sites (Amazon and Ebay), then find your site or product and click into your site or product page to view. [Name of the bot] can simulate many different people to do this based on your settings. [...] People use [name of the bot] to improve their Alexa standing... reduce bounce

rates... sell website on filippa... sell traffic, etc.”

Ad-fraud services are related to all online infrastructure and hosting services that might be needed to undertake identity or attribution fraud. Services can involve the creation of spam websites (fake networks of websites created to provide artificial backlinks); link building services; hosting services; creation of fake and scam pages impersonating a famous brand and used as part of an ad fraud campaign. For instance, a GENTEC advertisement promoted the following on January 6, 2019: “Selling RDPs and webhosting for fraud, scampages,⁴ etc. Servers located in France. Payment will be in BTC and no registration, I just need a domain (or you can get free subdomain) and you are set. Fraud RDP starting at BTC \$10 monthly and webhosting (with email) starting at BTC \$4 monthly.”

A successful ad-fraud campaign involves a sophisticated combination of these three types of ad-fraud—sending fake traffic through bots using fake social accounts and falsified cookies; bots will click on the ads available on a scam page that is faking a famous brand. Thus, technical communities provide knowledge on existing techniques and their incremental developments.

4.4.2. Innovation fueling engagement: general and customer-oriented ad-fraud communities

Contrary to other ad-fraud communities, customer-oriented communities, such as GENCUST2 and GENCUST1, showcased a particular focus on innovation through artificial intelligence. The latter triggered further changes in the way the cybercriminal community engaged with its customers, aiming at fostering customer interactions and engagement related to innovation and AI. For instance, on GENCUST2, on December 19, 2017, a programmer created the following thread:

“I am a programmer, write here you bots ideas”

“I made a bot that scrape random Wikihow page and gets the content of it (pictures and text), and then it creates a video and uploads it to my YT channel. The problem with this bot is that YT will not allow me to monetize the copyright content. So now I am searching for new ideas, if you have good idea for bot that can make money, I will do it and publish here for free. Feel free to suggest ideas!”

This post generated 94 replies in just one month, with members providing tips on how to dynamically rotate IP addresses thanks to machine learning as part of an “intelligent botnet” that could commit more sophisticated ad-fraud. In reaction to this popular thread, an entire section was created by the community owner dedicated to the use of AI.

Discussions on how to undertake innovative ad-fraud using AI are not limited to this section and span the entire GENCUST2 community. For instance, in another section on content marketing, a user explains the concept of deep fake technology and how he applies it to his porn website business by “making celebrity fake porn by swapping the celebrity faces onto porn videos” through deep fake. Discussing neural networks (autoencoders) and machine learning (generative adversarial networks), conversations flourished on the applications, such as using famous Hollywood movies stars as part of fake branded content creation. One of the cybercriminals boasted having fooled a company into a wire transfer using an AI-powered deep fake of a chief executive’s voice. Another commented a few months later: “Anyone working on projects involving this deep fake niche? I have a streaming community and forum up and running and want to grow it with someone. Let me know if you’re interested.”

On GENCUST1, we also observed a similar phenomenon. It started on February 10, 2017, with a discussion on the future of e-whoring. E-whoring is a technique of money-making used by black hat marketers that involves usurping the identity of a young woman on social networks and chat rooms to scam people into signing up for adult websites. This frequently involves ‘identity theft’ (stealing profile pictures of women on the internet and using them as part of this scam). The cybercriminal

⁴ Slang for a scam website impersonating a brand.

gets paid through adult affiliate programs. The cybercriminal discussion was “into computer science” and highlighted that “I’m currently growing several fake FB profiles (females) and everyday responding to tons of messages from creepy dudes is not truly filling me with optimism and nice thoughts about humanity [...] have you wondered how bots could be used for e-whoring? I have done split test playing (scripting) the virginal school girl vs slutty wench [...] AI-based chatbot is the future of ewhoring”. Following the rich discussion that unfolded over the months, the GENCUST1 community created a new section on automation and ad-fraud in October 2017. Topics ranged from dynamic SEO fraud using AI to discussions on how to process these technologies into actual fraud. This also contributed to increased engagement within the GENCUST1 community and contributed to creating new knowledge on the future of ad-fraud.

4.4.3. Summary of empirical findings

The following summarizes our key analysis findings across the six ad-fraud communities: general and customer-oriented communities aimed to attract a large diversity of criminal members interested in a variety of ad-fraud, contrary to specialized communities that focused on a niche. Indeed, specialized ad-fraud communities, such as SPECTEC1, SPECTEC2, and SPECTEC3, provide a wealth of knowledge on ad-fraud, most of the time though automated content, and incremental innovations on ad-fraud. However, communities that were both general AND customer-oriented (GENCUST1, GENCUST2) showcased the most internal interactions (see Table 1) and better performance and growth (as demonstrated by measures of brand awareness and popularity in Table 2). We observed that the value offering (value proposition) of these communities was positively perceived by their criminal customers—in turn, these improved perceptions favorably altered customers’ buying behaviors and engagement in the marketplace. In particular, new content (AI-based bots, deep fake, etc.) was driven by the development of a collaborative environment and ignited network effects attracting further customers while reinforcing generativity and interactions. A major contribution of our study is a theorization of the role played by specific capabilities discussed below that identifies and explains the nature of innovation in ad-fraud communities. In the next section, we discuss and theoretically integrate our key findings.

5. Discussion

As part of our multimethod approach integrating qualitative and quantitative analysis, we first examined (a) **internal interactions** within ad-fraud communities through a quantitative assessment of the number of new threads and posts created compared to the number of active users (out of a dataset of 1,294,861 threads and 12,245,156 messages posted in the six ad-fraud communities studied) and a qualitative assessment of the quality of the content exchanged in each ad-fraud community (automated content vs real interactions).

We then examined the (b) **performance** of ad-fraud communities through a quantitative assessment of the number of new active users and retention of users AND the assessment of each community’s dataset of backlinks ($N = 775,888$) to measure online popularity (the more backlinks, the more popular a community is). We also performed a qualitative assessment of the presence of threads discussing competition (market knowledge) and innovations in ad-fraud in each community.

Our empirical observations concerning innovations in ad-fraud communities led us to develop the following propositions, highlighting causal reinforcing loops favoring the development of ad-fraud communities. Our research is one of the first studies to document the way ad-fraud communities innovate and create value for their criminal customers. This has significant theoretical and managerial implications.

5.1. Reinforcing loops between community-specific capabilities and community performance

We observed GENCUST1 and GENCUST2’s support for cybercriminal

customers. We also noted efficient community management aimed at fostering customer engagement through customer journey and aimed at favoring interactions for value creation (see Section 4.4.2). Specific capabilities were enacted by the community to develop a vibrant community with user engagement, collaboration and experimentation. These specific capabilities represent the understanding of the broader environment and the understanding of customers (Moorman and Day, 2016). This process of transforming internal and external resources into value for customers is associated with two capabilities:

- (1) Market learning capabilities, aimed at anticipating broader marketplace changes and focused on knowledge generation and dissemination (Morgan, 2012), and including market planning and responsiveness capabilities, such as customer and competitor analysis (Moorman and Day, 2016).
- (2) Customer and relational capabilities aimed to improve the way the community delivers value to customers (Ceccagnoli et al., 2012; Vorhies and Morgan, 2005), reinforce marketing communications (social media, online advertising), and develop the community’s brand assets (Morgan et al., 2009).

Our research suggests that performance (in the form of community growth, popularity, user-led knowledge creation on AI) is both reinforced by and reinforces the capabilities of ad-fraud communities. First, customers cocreating an AI-based ad bot, organizing innovative collaboration in ad-fraud campaigns, discussing disruptive and fraudulent uses of legitimate big data or machine learning software, contributed to improving the cybercriminal community assessment of market evolution and thus reinforced market knowledge and learning capabilities. Second, we observed that some of the ad-fraud communities actively fostered internal interactions. For instance, the GENCUST2 community encouraged its customers to provide insights into new trends and innovations, which have indeed led to innovations being proposed and discussed in these communities. Thus, further innovations were an outcome of active customer management capabilities. We suggest the following theoretical proposition:

Proposition 1. *Market learning AND customer and relational capabilities strengthen and are strengthened by community performance (user-led knowledge creation on AI and innovation in ad-fraud; growth and popularity).*

Market learning AND customer and relational capabilities enable communities to better understand and forecast customers’ needs and achieve superior product differentiation and performance (Vorhies and Morgan, 2005). These capabilities are dynamic and fuel network externalities (Helfat et al., 2009); they involve the deployment of complex coordination to match marketing conditions (Morgan et al., 2009; Teece, 2007). However, to our knowledge, such capabilities in digital contexts have not been carefully addressed (Moorman and Day, 2016).

5.2. Requisite variety and its reinforcing feedback loops

The six ad-fraud communities we observed aimed to attract sellers and buyers of ad-fraud services (from beginner cybercriminals to more organized temporary hacking groups). Some communities were engaged in sensing activities: (1) they were sensing the competition offers (see page 16 for instance, SPECCUST and GENCUST1 demonstration of market sensing). (2) And/or they were open to exploring new opportunities (AI-based bots, for instance), enabling them to develop new services (adapting the community’s architecture to customers’ demands and needs), as in the case of GENCUST1 and GENCUST2. This sensing behavior is part of a process aimed at requisite variety (Weick, 1979), a matching process between environmental demand and the characteristics of the community. This requisite variety makes it possible to ignite the network effect—once the community reaches this critical mass, or tipping point, network effects become noticeable and self-reinforcing

(Evans et al., 2011). Our findings indicate that ad-fraud communities' capabilities also nurture network effects and increase community performance.

Requisite variety (diversity of customers) reinforced the capabilities deployed by these ad-fraud communities, creating value for the community (internal interaction, innovation leading to customer growth) and creating value for existing customers (engaging in customer-led innovation, fostering customer interests and retention). This highlights the central role of requisite variety (Orton and Weick, 1990) in reinforcing capabilities and enabling performance in ad-fraud communities. Hence, we suggest the following theoretical proposition:

Proposition 2. *Requisite variety, enabled by market learning AND customer and relational capabilities, is necessary for community performance and enables a second-order reinforcing effect toward market learning AND customer and relational capabilities.*

Although we studied communities in the area of ad-fraud, some of our findings might be generalized to other digital communities. We respond to a call for more research on how digital communities innovate (Hossain, 2018; Roma and Vasi, 2019).

In our case, general ad-fraud communities that were both customer-oriented and focused on generating interaction showed the best performance (innovation, customer growth, and loyalty). Our findings suggest that communities with a clear market selection, value-added orientation and reliance on user-led innovations have superior performance. Our study takes initial steps to develop our understanding of collective innovation and value cocreation in the context of ad-fraud communities.

5.3. Implications for practice: cybercrime, law enforcement and threat intelligence

In light of cybercrime's impacts on brands and digital advertisement initiatives, our study responds to a call for more research on cybercriminal undergrounds. This lack of research seems paradoxical given the important societal impacts of cybercrime and the increased need for a better understanding of how to tackle it (Hughes et al., 2017; Kraemer-Mbula et al., 2013). Therefore, from a practical perspective, our study improves the understanding of cybercriminal communities at the community level.

From a policy and law enforcement perspective, our findings contribute to the literature on enforcement and cybercrime deterrence (Antia et al., 2006; Hui et al., 2017; Kaur et al., 2021). Even if online advertising is borderless, a regime regulating click fraud would best serve advertisers' interests (compared to ad networks and search engines interested in maximizing their profit and safeguarding their self-interest): there is a serious lack of an international standard for defining, detecting, and controlling ad-fraud. We provide practical recommendations for law enforcement agencies: there is a need to fight cybercrime by hindering how ad-fraud communities define, develop, and deliver value to their criminal members and how communities combine and deploy their marketing resources. Examples of these possible law enforcement strategies include corrupting customer perception of the cybercriminal community and its brand; polluting knowledge creation (covert operations); reducing network effects by attacking requisite variety (detering malware as a service with corrupted files, prosecuting developers and every link of the criminal value chain; playing on the risk perception of weaker actors in the criminal value chain such as warez hackers)—and so on. Our investigation provides counterintuitive insight into the fight against cybercrime. Indeed, law enforcement activities have been designed to tackle specialized ad-fraud communities, hypothesized to be more innovative. Conversely, we show that general ad-fraud communities innovate by capitalizing on a diversity of internal skills. Law enforcement agencies should focus their investigations on new trends in ad-fraud communities by fostering diversity.

Finally, our research has practical implications for publishers and digital advertisers. Indeed, we recommend digital advertisers adopt a critical stance and demand transparency: traffic sources, audience extension practices, narrow and cheap reach might in fact originate from specialized ad-fraud communities. Advertisers and publishers need to develop cyberthreat intelligence capabilities to monitor innovations in general ad-fraud communities to protect investments—this is particularly interesting from a brand management perspective (detecting and fighting bogus websites impersonating and harming the brand). However, they also need to monitor traffic sources and techniques promoted in specialized ad-fraud communities.

6. Limitations and conclusions

Although our study is one of the first to investigate ad-fraud communities using actual cybercriminal behavioral data, it is not without limitations.

A major limitation of our research is related to the nature of the research subject: verifying the authenticity of research subjects in the context of ad-fraud communities and cybercrime is impossible—the anonymity of these communities enables participants to conceal their identity and personal characteristics and deceive others. One individual can manage several personas in the community. However, in the context of our research, misrepresentation and concealment are part of our subjects' social life (Adler, 1998). Another limitation is related to the context of our investigation. Despite the extensiveness of this study, our results are based on an investigation of six ad-fraud communities. Therefore, even if we followed best practices for internal validity and external credibility (Patton, 2005; Walsham, 2006), we would caution generalizing from our theoretical contribution beyond the context of ad-fraud communities. Finally, our research does not investigate the impacts of ad-fraud: our investigation focuses on ad-fraud origins and how ad-fraud communities innovate in this area. Future research could establish direct complementary evidence of impacts from advertisers' data; this would represent an interesting theoretical and practical complement to our research. In particular, future research could develop the impact of deep fake (ad-frauds based on AI techniques such as deep fake to impersonate a brand or a person) on brand equity and revenues. For instance, the impact of ad-fraud campaigns on the unfolding of brand crises and social media firestorms—how might such fraudulently initiated online firestorms be mitigated?

Our investigation explains how ad-fraud communities innovate and perform. In particular, we demonstrated how user-generated innovative uses of AI and technology have changed the way some ad-fraud communities perform: initial discussions on AI-based innovations triggered the development of a collaborative environment, driving the realization of innovative content (AI-based bots, deep fake, etc.). This content provides an impulse and ignites network effects that attract additional customers while reinforcing generativity and interactions. This process requires the development of specific capabilities (market learning, customer and relational capabilities) and a well-balanced community ecosystem to attract and maintain customers in the community.

Author statement

The conceptualization, methodology, writing, revising and all other relevant CRediT roles have been performed by the single author of the paper.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.techfore.2021.121282](https://doi.org/10.1016/j.techfore.2021.121282).

References

- Adler, P., 1998. *Ethnography At the edge: Crime, deviance, and Field Research*. UPNE.
- Advertising Age, 2006. Search Marketing Fact Pack 2006. AdAge.
- Agarwal, N., Athey, S., Yang, D., 2009. Skewed bidding in pay-per-action auctions for online advertising. *Am. Econ. Rev.* 99, 441–447.
- Agerfalk, P.J., 2013. Embracing diversity through mixed methods research. *Eur. J. Inf. Syst.* 22, 251–256.
- ANA, 2017. *Fraud in Digital advertising: Bot Baseline 2017*. Association of National Advertisers.
- Antia, K.D., Bergen, M.E., Dutta, S., Fisher, R.J., 2006. How does enforcement deter gray market incidence? *J. Mark.* 70, 92–106. <https://doi.org/10.1509/jmkg.70.1.092.qxd>.
- Asdemir, K., Yurtseven, Ö., Yahya, M.A., 2008. An economic model of click fraud in publisher networks. *Int. J. Electron. Commer.* 13, 61–90.
- Barker, S., 2019. *Digital Advertising Fraud* - Juniper Research. Juniper.
- Baskerville, R.L., Myers, M.D., 2015. Design ethnography in information systems. *Inf. Syst. J.* 25, 23–46.
- Benjamin, V., Valacich, J.S., Chen, H., 2019. DICE-E: a framework for conducting darknet identification, collection, evaluation with ethics. *MIS Q.* 43.
- Birks, D.F., Fernandez, W., Levina, N., Nasirin, S., 2013. Grounded theory method in information systems research: its nature, diversity and opportunities. *Eur. J. Inf. Syst.* 22, 1.
- Brenner, S.W., 2002. Organized cybercrime-how cyberspace may affect the structure of criminal relationships. *NCJL & Tech* 4, 1.
- Ceccagnoli, M., Forman, C., Huang, P., Wu, D.J., 2012. Cocreation of value in a platform ecosystem: the case of enterprise software. *MIS Q.* 36.
- Chang, V., Baudier, P., Zhang, H., Xu, Q., Zhang, J., Arami, M., 2020. How Blockchain can impact financial services – The overview, challenges and recommendations from expert interviewees. *Technol. Forecast. Soc. Change* 158, 120166. <https://doi.org/10.1016/j.techfore.2020.120166>.
- Charmaz, K., 2006. *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*, 1 Edition. SAGE Publications Ltd, London, UK.
- Chen, Y., Kintis, P., Antonakakis, M., Nadji, Y., Dagon, D., Farrell, M., 2017. Measuring lower bounds of the financial abuse to online advertisers: a four year case study of the TDSS/TDL4 Botnet. *Comput. Secur.* 67, 164–180.
- Edelman, B., Brandi, W., 2015. Risk, information, and incentives in online affiliate marketing. *J. Mark. Res.* 52, 1–12.
- Evans, D.S., Schmalensee, R., Noel, M.D., Chang, H.H., Garcia-Swartz, D.D., 2011. *Platform Economics: Essays on Multi-Sided Businesses* (SSRN Scholarly Paper No. ID 1974020). Social Science Research Network, Rochester, NY.
- Helfat, C.E., Finkelstein, S., Mitchell, W., Peteraf, M., Singh, H., Teece, D., Winter, S.G., 2009. *Dynamic capabilities: Understanding strategic Change in Organizations*. John Wiley & Sons, Hoboken, NJ.
- Hossain, M., 2018. Motivations, challenges, and opportunities of successful solvers on an innovation intermediary platform. *Technol. Forecast. Soc. Change* 128, 67–73. <https://doi.org/10.1016/j.techfore.2017.10.018>.
- Hughes, B.B., Bohl, D., Irfan, M., Margolese-Malin, E., Solórzano, J.R., 2017. ICT/Cyber benefits and costs: reconciling competing perspectives on the current and future balance. *Technol. Forecast. Soc. Change* 115, 117–130. <https://doi.org/10.1016/j.techfore.2016.09.027>.
- Hui, K.-L., Kim, S.H., Wang, Q.-H., 2017. Cybercrime deterrence and international legislation: evidence from distributed denial of service attacks. *Mis Q.* 41, 497.
- Jones, D., 2012. Flow Metrics [WWW Document]. Majestic Blog. URL: <https://blog.majestic.com/development/flow-metrics/> (accessed 6.22.19).
- Kaur, P., Dhir, A., Tandon, A., Alzeiby, E.A., Abohassan, A.A., 2021. A systematic literature review on cyberstalking. An analysis of past achievements and future promises. *Technol. Forecast. Soc. Change* 163, 120426. <https://doi.org/10.1016/j.techfore.2020.120426>.
- Kraemer-Mbula, E., Tang, P., Rush, H., 2013. The cybercrime ecosystem: online innovation in the shadows? *Technol. Forecast. Soc. Change* 80, 541–555. <https://doi.org/10.1016/j.techfore.2012.07.002>.
- Kwon, T.H., Kwak, J.H., Kim, K., 2015. A study on the establishment of policies for the activation of a big data industry and prioritization of policies: lessons from Korea. *Technol. Forecast. Soc. Change* 96, 144–152. <https://doi.org/10.1016/j.techfore.2015.03.017>.
- Lamberton, C., Stephen, A.T., 2016. A thematic exploration of digital, social media, and mobile marketing: research evolution from 2000 to 2015 and an agenda for future inquiry. *J. Mark.* 80, 146–172.
- Leukfeldt, E.R., 2014. Cybercrime and social ties. *Trends Organ. Crime* 17, 231–249.
- Chen, Min, Jacob, V.S., Radhakrishnan, S., Ryu, Y.U., 2015. Can payment-per-click induce improvements in click fraud identification technologies? *Inf. Syst. Res.* 26, 754–772.
- Moorman, C., Day, G.S., 2016. Organizing for marketing excellence. *J. Mark.* 80, 6–35.
- Morgan, N.A., 2012. Marketing and business performance. *J. Acad. Mark. Sci.* 40, 102–119.
- Morgan, N.A., Vorhies, D.W., Mason, C.H., 2009. Market orientation, marketing capabilities, and firm performance. *Strateg. Manag. J.* 30, 909–920. <https://doi.org/10.1002/smj.764>.
- Nazerzadeh, H., Saberi, A., Vohra, R., 2013. Dynamic pay-per-action mechanisms and applications to online advertising. *Oper. Res.* 61, 98–111.
- Orton, J.D., Weick, K.E., 1990. Loosely coupled systems: a reconceptualization. *Acad. Manage. Rev.* 15, 203–223. <https://doi.org/10.2307/258154>.
- Patton, M.Q., 2005. *Qualitative Research*. John Wiley & Sons, Hoboken, NJ.
- Perrin, N., 2019. *eMarketer: Digital Ad Fraud 2019*.
- Richet, J.-L., 2013. *Laundering Money Online: a Review of Cybercriminals Methods. Tools and Resources for Anti-Corruption Knowledge (UNODC)*.
- Roma, P., Vasi, M., 2019. Diversification and performance in the mobile app market: the role of the platform ecosystem. *Technol. Forecast. Soc. Change* 147, 123–139. <https://doi.org/10.1016/j.techfore.2019.07.003>.
- Santoro, G., Vrontis, D., Thrassou, A., Dezi, L., 2017. The Internet of Things: building a knowledge management system for open innovation and knowledge management capacity. *Technol. Forecast. Soc. Change*.
- Saridakis, G., Benson, V., Ezingard, J.-N., Tennakoon, H., 2016. Individual information security, user behaviour and cyber victimisation: an empirical study of social networking users. *Technol. Forecast. Soc. Change* 102, 320–330. <https://doi.org/10.1016/j.techfore.2015.08.012>.
- Schultze, U., 2000. A confessional account of an ethnography about knowledge work. *MIS Q* 3–41.
- Shareef, M.A., Kumar, V., Dwivedi, Y.K., Kumar, U., Akram, M.S., Raman, R., 2021. A new health care system enabled by machine intelligence: elderly people's trust or losing self control. *Technol. Forecast. Soc. Change* 162, 120334. <https://doi.org/10.1016/j.techfore.2020.120334>.
- Soudijn, M.R., Zegers, B.C.T., 2012. Cybercrime and virtual offender convergence settings. *Trends Organ. Crime* 15, 111–129.
- Teece, S.D.J., 2007. Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strateg. Manag. J.* 28, 1319–1350.
- Vorhies, D.W., Morgan, N.A., 2005. Benchmarking marketing capabilities for sustainable competitive advantage. *J. Mark.* 69, 80–94.
- Walsham, 2006. Doing interpretive research. *Eur. J. Inf. Syst.* 15, 320–330.
- Weick, K., 1979. *The Social Psychology of Organizing*. Wesley, Reading, MA.
- White, G.R.T., Samuel, A., 2019. Programmatic advertising: forewarning and avoiding hype-cycle failure. *Technol. Forecast. Soc. Change* 144, 157–168. <https://doi.org/10.1016/j.techfore.2019.03.020>.
- Wilbur, K.C., Zhu, Y., 2009. Click fraud. *Mark. Sci.* 28, 293–308.

Expert in cybersecurity, **Jean-Loup Richet** is associate Professor of Information Systems Management and co-director of the Risk Chair at Sorbonne Business School, IAE Paris, Université Paris I Panthéon-Sorbonne, France. Jean-Loup Richet's work explore the boundaries of cybercrime and cybersecurity, focusing on trends in online money laundering or new frauds enabled by Artificial Intelligence and Machine Learning. He has published numerous papers in trade and academic journals (European Journal of Information Systems; IEEE Transactions on Engineering Management); his work was featured in The Wall Street Journal, Wired, CBS, MIT Technology Review, Computer World, and many other media outlets.