Seeing is Not Always Believing: An Empirical Analysis of Fake Evidence Generators

Zhaojie Hu*
University of Central Florida
Orlando, USA
zhaojie.hu@ucf.edu

Jingzhou Ye*
University of Central Florida
Orlando, USA
jingzhou.ye@ucf.edu

Yifan Zhang
Indiana University Bloomington
Bloomington, USA
yz113@iu.edu

Xueqiang Wang
University of Central Florida
Orlando, USA
xueqiang.wang@ucf.edu

2024 IEEE 9th European Symposium on Security and Privacy (EuroS&amp;P) | 979-8-3503-5425-6/24/$31.00 ©2024 IEEE | DOI: 10.1109/EUROSP60621.2024.00037

Abstract—**Online scams pose a growing threat to the cyberspace, with cybercriminals frequently using fake evidence, such as identification and financial documents, to illicitly elevate their credibility in online activities. This deceptive trend is fueled by an emerging set of fake evidence generators (FeGens). These FeGens replicate the output of authoritative sources, such as official bank applications, to automatically generate large quantities of authentic-looking fake evidence. To the best of our knowledge, FeGens, as effective tools for cybercriminals, have not been systematically analyzed in terms of their supply chain, including development, promotion, and delivery, as well as the risks and impacts they pose to end users. In this paper, we present the first systematic empirical analysis of FeGens and related fake evidence. Our findings shed light on the FeGen ecosystem, particularly the tactics employed by FeGen developers and retailers to mimic authoritative sources and promote the use of FeGens. We also evaluate the effectiveness of FeGens and associated risks in cybercrime.**

1. Introduction

Internet users today face a growing threat from online scam, in which cybercriminals use deceptive tactics to exploit user vulnerabilities and manipulate trust in online activities. Typical examples include scams that use fake identity and financial documents (e.g., driver's licenses and bank statements) as supporting evidence to illicitly obtain mortgages and other loans [57], [58], [63]. This *fake evidence* mimics real evidence generated by authoritative sources, such as banks and government agencies. When abused, it can lead to an elevation of the scammers' credibility, thereby damaging trust and public confidence in cyberspace.

Fake evidence generator (FeGen). Traditional techniques for generating fake evidence often require professional graphics editing skills to manipulate real evidence [86]. However, they struggle to meet the growing demands of cybercriminals who target a large population of potential victims and a widening variety of online evidence. As a response to these challenges, new techniques, or tools, have emerged, which we refer to as *fake evidence generators* (FeGen). Using FeGens, cybercriminals can automatically replicate the output of authoritative sources through user-friendly interfaces with minimal effort, significantly enhancing their capabilities for online misconduct. An example user interface (UI) for a bank-related FeGen is shown in Figure 1. This UI closely resembles the bank transfer UI of the official Bank of China application [84]. However, the input fields of the UI, such as the transfer amount and payee information, can be specified with arbitrary values by cybercriminals, allowing them to generate fake bank transfer receipts. Importantly, the FeGen lacks backend logic, and all the above steps are completed without triggering an actual transfer. Cybercriminals may use these fake receipts to mislead potential victims into participating in illicit underground businesses, such as online gambling.

Eliminating such fake evidence in online scams is instrumental for restoring user confidence in cyberspace. However, thus far, little has been done to analyze and understand the ecosystem of fake evidence, in particular the supply chain of FeGens, including their development, distribution, and real-world impacts.

Understanding FeGen ecosystem. This paper seeks to fill the gap by presenting the first systematic empirical analysis of FeGens and the associated fake evidence, and evaluate their impacts to online users. To facilitate this analysis, we build the first datasets for real-world FeGens and the generated evidence. Based upon the datasets, we investigate the FeGen ecosystem from several important perspectives. Specifically, we examine the landscape to provide quantitative insights into the distribution of FeGens and fake evidence in the wild. Then, we explore how FeGens, which are often harmful and the target of security scrutiny, are disseminated to cybercriminals by characterizing the promotion and delivery channels. After that, we analyze the code of FeGens to unveil the tactics deployed by underground developers, in order to create FeGens that closely mimic the output of authoritative

* *The two lead authors contributed equally to this work and are listed alphabetically.*

DOI 10.1109/EuroSP60621.2024.00037

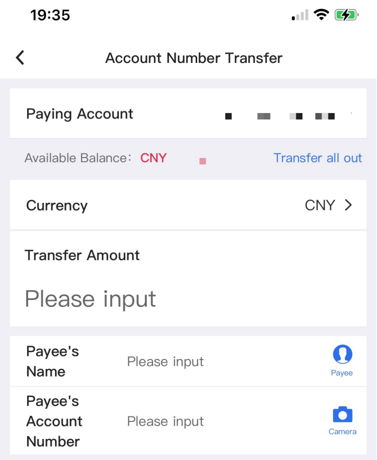

Account Number Transfer

Paying Account [redacted]

Available Balance: CNY [redacted] Transfer all out

Currency CNY >

Transfer Amount

Please input

Payee's Name Please input

Payee's Account Number Please input

Figure 1: The bank transfer user interface (UI) of a FeGen

software. We further evaluate whether FeGens, when misused by legitimate users, can potentially pose security and privacy risks to those users. Finally, we address the crucial question of whether the fake evidence generated by FeGens is indistinguishable to online users when compared to real evidence, based on a user-based study.

Findings. Looking into the ecosystem, we are surprised to find that FeGens are trending with a significant impact on today's development of online scams, which is less known to the security and privacy community. More specifically, by investigating known scam reports and online FeGen promotional messages, we found that FeGens are becoming promising tools for generating a broad range of fake evidence, covering at least six fake evidence types and 29 subtypes, such as fake identification documents (e.g., driver's licenses and passports), financial documents (e.g., pay stubs and digital wallet transactions), etc. From these fake evidence types, we built a FeGen dataset with 102 real-world FeGens (with 124 instances in the form of installable software tools across multiple platforms or as-a-service websites) that span across different languages, i.e., English and Chinese. Interestingly, through an analysis of FeGen in different languages and their active periods, we observed that FeGens and the associated fake evidence are not evenly distributed. For example, FeGens have long been popular in the English-speaking world with a focus on financial and identification documents, while they have gained popularity in the Chinese-speaking world in the past few years, trending towards fake social media and digital wallet evidence.

Further examination of the FeGens from a supply chain perspective allows us to make a series of new observations regarding the promotion, delivery, and development of FeGens, as well as their risks and impacts on end users. First, many FeGens offer premium services for generating fake evidence and extensively use cryptocurrencies (e.g., Bitcoin, USDT) for FeGen transactions, potentially to evade financial scrutiny. The activities in cryptocurrency accounts indicate not only a large volume of FeGen users (e.g., cybercriminals) but also the financial network that connects FeGen retailers through cryptocurrency. Second, FeGen retailers are promoting FeGens through various channels. Telegram groups are a popular choice that incorporates lesser-known promotional strategies, including collaborative promotional campaigns across multiple groups and groups that illicitly boost their popularity using fake group members. In terms of FeGen delivery, retailers are inclined to use distribution channels that are less monitored by content review, such as CowTransfer, tmp.link, MuseTransfer, and lanzoucloud, which are less known to the general public. Third, while most FeGen developers tend to reuse some UI elements from authoritative softwares to enhance the authenticity of their fake evidence, it's surprising to discover that some FeGens go as far as achieving complete impersonation of authoritative sources by stealing entire user interfaces (UIs). Notable examples of these FeGens are the group of mobile bank app simulators that closely mimic the interfaces of seven leading banks in China, including *ICBC* [10], *CMB* [11], and *BoCom* [9]. To appear more legitimate, many FeGens hide their developer identities through anonymization techniques or even disguise themselves as popular legitimate software, such as that from Google Inc. Fourth, we observed, through a user study, that the fake evidence generated by the FeGens is effective in deceiving users because online users find it challenging to clearly distinguish fake evidence from real evidence. What's even more concerning is that a non-negligible percentage (7.2%) of users have experienced financial losses as a result of scams involving fake evidence. In the last, according to VirusTotal [5], over a quarter (26.6%) of FeGens pose security and privacy risks to their users. A closer inspection of these reported FeGens illustrates the fact that FeGens are not only involved in cybercrime themselves but are also actively targeted by other malware, referred to as "FeGen predators", in the wild.

Contributions. The contributions of the paper are outlined as follows:

- We build the first dataset for a new type of cybercrime infrastructure – fake evidence generators (FeGens). Based upon the dataset, we conduct the first systematic empirical analysis of the FeGen ecosystem.
- We characterize FeGens by analyzing them from a supply chain perspective, which leads to a series of new findings that broaden our knowledge on the underground development, promotion, distribution and use of FeGens.
- We assess the indistinguishability of fake evidence generated by FeGens from real evidence, which indicates the effectiveness of FeGens in facilitating online scams.

2. Background

Digital evidence for online activities. Digital evidence can essentially be any piece of information that supports claims related to activities involving computing systems. It is commonly used in forensic analysis and cybercrime investigation [52], [53], [92], such as for using device and service logs as evidence for account activities. In this study, we examine digital evidences from a slightly different and narrower perspective – digital evidences for online activities. We define evidence as information supposedly generated by an authoritative source (e.g., an official bank application) and used to establish facts between an individual and the authoritative source (e.g., confirming the individual's account balance). Due to the significant trust and value people place in such evidence, online fraudsters are often motivated to create fake but authentic-looking evidence to appear more trustworthy in online activities, such as scams. For instance, a fraudster may use forged pay stubs or bank statements when applying for loans [58], or present them as evidence of income to deceive victims into engaging in illicit services, such as online gambling [69], [91].

Ecosystem of FeGen. Fake evidence can result in serious ethical and legal consequences when used in scams [66]. Exacerbating the problem is the emergence of FeGens, which can support the bulk generation of fake evidence. Compared to traditional techniques, like creating fake evidence using graphic editors on real evidence, FeGens are increasingly favored by miscreants due to their high throughput of generating authentic-looking evidence and their ease of use. For example, FeGens often offer one-click generation and are fully automated, potentially leaving fewer artificial traces for detection.

The use of FeGens is facilitated by an ecosystem that involves many parties. As depicted in Figure 2, this ecosystem begins with underground developers creating FeGens capable of mimicking the output of authoritative sources, such as banking websites or applications. These developers illicitly profit by transferring FeGens to FeGen retailers (① and ②). Then, the retailers promote FeGens by posting advertisements on widely accessible platforms, such as social media and websites (③). At the same time, FeGens are uploaded to distribution channels (③), which typically serve as storage and content delivery services for underground businesses. In the next step, potential FeGen users, or miscreants, may search for available FeGens and eventually visit the promotional channel (④) to download FeGen installation packages from the distribution channel (⑤). Following the successful downloads of the installation packages, FeGen users often pay retailers to gain access to FeGens (⑥). Afterward, FeGen users deploy FeGens on their required platforms (⑦) and use these instances to generate fake evidence ⑧. In the last step, FeGen users may use the fake evidence to deceive victims (⑨) and obtain illicit gains from these victim users (⑩), such as stealing money from their accounts or involving them in unwanted businesses. It is important to note that the above ecosystem serves as a simplified illustration of how FeGens are created and used. In reality, the ecosystem can be more complicated. For example, the flow above illustrates the distribution of FeGen installation packages (i.e., on-premises deployment). In many cases, FeGens are provided as a service (e.g., through websites) by FeGen retailers to users. Also, FeGen users may not only generate fake evidence for their own use but also profit from their FeGen instances by opening access to other online miscreants.

In this study, we conduct an empirical analysis of each component in this ecosystem in an attempt to offer the very first characterization of FeGens and an understanding of their risks and impacts.

3. Datasets

The very first step towards a comprehensive understanding of the generation of fake evidence is to gather

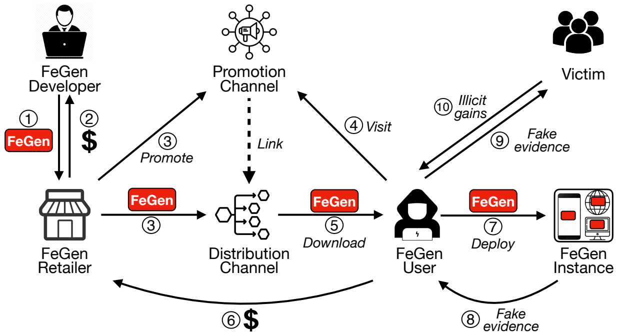

Figure 2: FeGen ecosystem

a FeGen dataset for analysis. This task is nontrivial because FeGens are often distributed through various underground channels by online miscreants, rather than centralized software marketplaces. Below, we elaborate on an effective methodology that we used to build the dataset. In this section, we use FeGens to refer to the entities that allow their customers to generate fake evidence through installable software tools or FeGen as-a-service (AAS) websites. We call these tools and websites FeGen instances. One FeGen can have several instances in the form of software tools running on various platforms, as well as multiple AAS websites.

3.1. Identifying FeGens

In this study, we identify FeGens and their instances with a four-step pipeline.

Collect evidence types. As we introduce in Section 2, digital evidence can appear in a wide variety of forms. However, there is no systematic knowledge about which types of evidence, if abused, are indeed harmful to online users. To address the challenge, we leverage the observation that there are a number of scam reports published by government agencies (such as FTC [14]) and leading security companies, which indicate the types of fake evidence being used by miscreants.

Specifically, to collect the evidence types biased in favor of online miscreants, we first used Google Search [12] to identify public scam reports by searching the combination of "`scam`" and "`fake`" in both English and Chinese, the two most spoken languages. We reviewed the top 100 query reports for each language, as ordered by Google Search by relevance, and narrowed the scope down to 19 English web pages and 27 Chinese web pages. We manually reviewed these web pages and identified 21 types of fake evidence, with English pages covering 8 types and Chinese pages covering 16 types (with overlaps). These evidence types span six broad types, such as identification documents (e.g., `driver's license`, `ID card`), financial documents (e.g., `pay stub`, `bank statement`), and legal documents (e.g., `court order`), etc. A complete list of evidence types can be found in columns 1-2 of Table 1.

Gather potential FeGens. FeGens are not available in public marketplaces like regular software. But in order to attract users, retailers (or miscreants) of FeGens must promote them on public channels, such as websites and social media, providing descriptions of supported fake

TABLE 1: The breakdown of the FEGEN dataset. There are 124 FEGEN instances belonging to 102 FEGENS. Note that one FEGEN (and its instance) may generate multiple types of fake evidence.

Types of Fake Evidence§		Number of FEGENS	Number of FEGEN Instances					
			AAS Website	Software Tools				Subtotal
				Windows	macOS	Android	iOS	
Identification docs	Driver's license	3	2	-	-	1	-	3
	National ID card	8	3	4	1	1	-	9
	Passport	2	2	-	-	-	-	2
Financial docs	Airline ticket	5	5	-	-	-	-	5
	Bank check	5	4	-	-	1	-	5
	Bank statement	28	6	19	-	3	-	28
	Credit report	4	4	-	-	-	-	4
	Digital wallet balance	22	9	7	-	8	1	25
	Digital wallet transaction history	22	9	7	-	8	1	25
	eCommerce order	4	4	-	-	4	4	12
	Pay stub	7	7	-	-	-	-	7
	Property deed	2	1	1	-	-	-	2
	Parking ticket	0	-	-	-	-	-	0†
	Tax document	2	2	-	-	-	-	2
	Utility bill	2	2	-	-	-	-	2
Legal docs	Court order	0	-	-	-	-	-	0†
	Employment document	1	1	-	-	-	-	1
	Letter of authorization	0	-	-	-	-	-	0†
	Business contract	1	1	-	-	-	-	1
Certificates/Licenses	Award certificate	0	-	-	-	-	-	0†
	Business license	5	-	5	-	-	-	5
	Caste certificate	0	-	-	-	-	-	0†
	Certificate of degree/diploma	6	1	4	-	1		6
	Transcript	4	4	-	-	1	1	6
	Marriage certificate	2	2	-	-	-	-	2
Social network	Account profile	32	17	3	1	16	9	46
	Post	9	6	1	1	4	2	14
	Chat history	34	19	3	1	16	10	49
Medical	Medical record	2	1	-	-	1		2
	Subtotal	102	50	27	2	31	14	124

§ Scam reports that mention fake evidence are at [45].

† "0" indicates that we found the use of this type of fake evidence in prior scam reports but were unable to locate any FEGEN instances that generate the fake evidence.

evidence and guidance on how users can access them. We refer to these channels as promotion channels. Hence, we can identify the potential FEGENS being promoted by these channels by searching for promotional messages online. Specifically, we searched for the coexistence of evidence types with "*generator*" or "*generation*" on Google Search [12], Facebook [16], and Douyin [36], all of which are leading platforms. For each type of evidence, we reviewed the top 10 search results on each platform and confirmed that they are indeed related to the generation of fake evidence, rather than the accidental coexistence of search keywords in lengthy documents.

In total, we located 81 potential FEGENS by reviewing 630 search results, with 81 from Google Search, 4 from Facebook, and 15 from Douyin. Notably, all the potential FEGENS from Facebook and Douyin are also covered by Google Search. Therefore, we focused on Google Search in the latter step of the pipeline.

Expand the set of potential FEGENS. The above potential FEGENS cover fake evidence found in a small number of public scam reports. We want to expand the set to enhance comprehensiveness. An observation enabling the expansion is that FEGENS often support additional types of fake evidence beyond those collected earlier in order to maximize illicit gains. Those types of fake evidence are highly suspicious, since they are used by the similar miscreants involved in reported scam cases. Therefore, in the process of dataset enrichment, we deploy a guilt-by-association strategy where we gather all evidence types supported by the 81 potential FEGENS by reviewing the FEGEN descriptions in the search results and repeat the search process for dataset creation to find FEGENS that support these additional evidence types. As a result, we identified eight more types of evidence. Reviewing the search results of these types (the same method as *dataset creation*) yielded 44 potential FEGENS that were not included in the previous step. This brings the total number of potential FEGENS to 125.

We manually checked the promotional websites or social media for each potential FEGEN, and followed the guidance provided in these channels to identify FEGEN websites offering FEGENS as-a-service (AAS), or download links for FEGEN software tools. The results include 148 instances with 72 AAS websites, and 76 software tools, corresponding to the 125 potential FEGENS.

Refine the set of potential FEGENS. Not all FEGENS designed to generate fake evidence can be exploited to harm users. For instance, a FEGEN may generate fake diplomas that do not appear to be from any authoritative sources, e.g., without seals and institution names. We believe that such low-credible evidence is less effective in online scams. Therefore, we review the descriptions, sample evidence and videos presented by the retailers in the promotion channels, and refine the dataset by excluding those that fail to mimic authoritative output or have visual cues indicating the evidence is fake. In this step, 23 out of the 125 FEGENS, covering 24 instances, are removed from the dataset. We include all remaining 102

FeGens with 124 instances in the FeGen dataset, since they have a high chance to generate authentic-looking evidence. While some of these FeGens may not function precisely as advertised by the retailers (such as containing malicious content), we believe they represent a subset of potential FeGens available in the real world to online users intending to use fake evidence.

3.2. FeGen Dataset

Among the 102 FeGens, 32 target English-speaking users and 70 target Chinese-speaking users. Table 1 displays the detailed dataset breakdown. Overall, the FeGen instances generate six types of fake evidence: *identification documents*, *financial documents*, *legal documents*, *certificates and licenses*, *social networks*, and *medical documents*, with 29 sub-types (columns 1 and 2 of the table). The instances run on different platforms: 50 FeGen instances are in the form of FeGen AAS websites, while 74 instances are in the form of on-premises software. Specifically, 31 instances run on Android, 27 instances run on Windows, 14 instances run on macOS, and 2 instances run on macOS. 14 FeGens contain at least two instances that run on different platforms. These instances are not evenly distributed across different languages. In the 32 FeGens targeting English-speaking users, 24 are found on websites, while the remaining eight FeGens have seven instances on Android and two instances on iOS. In contrast, the FeGens targeting Chinese-speaking users have a wider distribution across all platforms, especially on Windows. Additionally, we found that 50 FeGens are offered free of charge (including 66 instances), with 20 of them relying on advertisements to generate revenues. On the other hand, 52 FeGens (including 58 instances) require payment, with 45 requiring subscription-based payments from users and seven requiring a per-document fee. We did not find any FeGens that offer both free and paid instances simultaneously.

The dataset includes the FeGen instances (e.g., URLs of FeGen AAS websites or executables of FeGen software tools), along with a number of attributes related to the instances, such as their promotion and distribution channels, instance descriptions, types of supported fake evidence, etc. In the Appendix, we list the FeGen AAS websites in Table 8, and the FeGen software tools in Table 9.

3.3. Evidence Dataset

To assess the impact of fake evidence to end users, we need an evidence dataset with both fake and corresponding real evidence. For this purpose, we ran all the 50 free FeGens from the above dataset to generate fake evidence. Specifically, we used Google Chrome version 120.0 to load FeGen AAS websites. For on-premise FeGen software, we executed them in the following environments: Windows 10 Pro, macOS Ventura 13.4, Pixel 7 with Android 13.0, and iPhone 11 with iOS 16.6.1. In total, we successfully ran 44 FeGen instances across 32 free FeGens and dumped fake evidence from them, as labeled in Table 8 and 9 in the Appendix. We failed to run any instances in the other 18 FeGens mainly due to compatibility problems. This process resulted in the creation of 94 pieces of fake evidence. We did not pay for FeGens and did not run any instances from the 52 FeGens that have their main functionalities behind a paywall. On the other hand, for each piece of fake evidence, we collected the authoritative sources that can generate corresponding real evidence, including 22 applications and websites for social networks, financial or healthcare institutions. Running them allowed us to generate 40 pieces of real evidence. For example, we ran the official *Bank of America* website [8] with the authors' personal information to gather real evidence related to banks. In this process, we were unable to obtain real evidence for 5 FeGens (out of the 32 FeGens that generate fake evidence) due to limited access to their authoritative sources, such as obtaining business licenses issued by the Chinese government.

TABLE 2: Evidence dataset

	# Pieces of Evidence
Fake	94
Real	40
Total	134

Discussion and open access to the dataset. To facilitate future studies on defending against FeGens, we have provided open access to both the FeGen and evidence datasets on our website [45]. However, in order to maintain the authenticity of the real evidence, we unavoidably used our personal information when generating the real evidence, such as creating official bank statements. Opening access to them may lead to the exposure of personal information. Therefore, we have removed the sections that could potentially reveal our private information from the real evidence by blurring parts of the images. We also removed the corresponding sections in the fake evidence to align them with the structure of real evidence. We anticipate that this modification will not significantly reduce the usability of the dataset since the majority of the evidence remains intact. The datasets do not contain private information of any other parties.

4. FeGen Ecosystem Analysis

To understand the FeGen ecosystem, we conduct a combination of qualitative and quantitative study of the FeGen dataset. In particular, we aim to answer the below research questions:

- **Landscape** - What is the landscape of the FeGen ecosystem?
- **Promotion** - How are the FeGens promoted in order to acquire a large group of users (e.g., miscreants)?
- **Distribution** - How are the FeGens successfully delivered to their users given that they are illicit by nature?

4.1. Landscape

FeGen for different languages. In Section 3, we presented the various types of fake evidence generated by FeGens. An important perspective to consider is whether these types of evidence are equally exploited in different

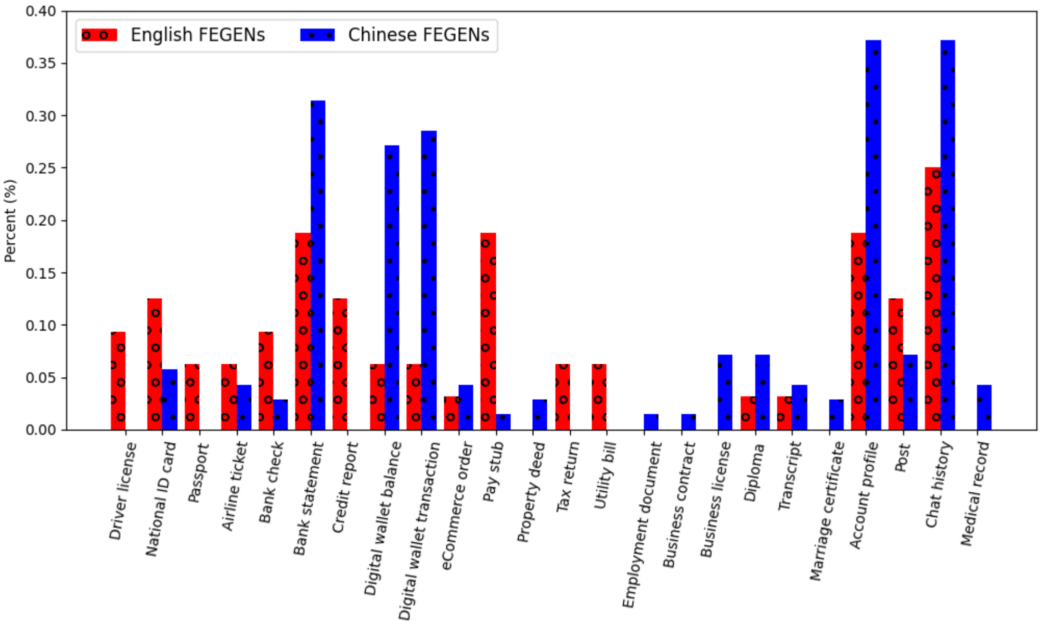

Figure 3: FeGen type distribution for different languages

regions. While our dataset only includes a small number of FeGens, we believe it allows us to develop a preliminary understanding of the distribution of fake evidence across regions, specifically, English and Chinese-speaking regions.

In Figure 3, we first divide the FeGens into two groups based on their language. Then, for each group, we check the advertised types of fake evidence supported by the FeGens in that group and plot the distribution of these fake evidence types. As the figure suggests, there are significant disparities in the support of fake evidence types between Chinese and English FeGens, let alone even distribution. More specifically, we notice that English FeGens are more inclined to generate fake identification and financial documents, such as driver's license, passport, credit report, tax documents, and utility bill, etc. This finding is aligned with the FTC's recent actions in combating "*identity theft, tax fraud, and similarly unsavory conduct*" [63], indicating that FeGens indeed provide infrastructural support for cybercriminals involved in such activities. A reasonable guess that might explain why Chinese FeGens do not target these types of evidence is that such evidence is less commonly used by Chinese users. For example, there is no nationally coordinated credit system, and as a result, there is less use and trust in credit reports in China compared to English-speaking countries [73]; Chinese individuals are not required to file tax returns, leading to fewer use of tax documents [50].

On the other hand, a larger portion of Chinese FeGens generate fake evidence for social network platforms and digital wallets. For example, at least 19 systems allow users to generate fake social network accounts, chat histories, and associated digital wallet transactions on two leading platforms, i.e., `WeChat` [23] and `Alipay` [15]. This is perhaps caused by the critical role these platforms play and the increased trust that Chinese users place in them. They not only infiltrate every aspect of individual lives but also have a significant influence on how many Chinese government agencies operate, e.g., the agencies may provide services, share updates, and communicate information with the public via these platforms. As a result, it appears reasonable that these platforms have become one of the most favorable targets for FeGens in China.

> *Finding 1:* FeGens *of different languages are inclined to generate different types of fake evidence, potentially indicating an uneven distribution of scams in these regions.*

Active time of FeGens. Another aspect to investigate is how long the FeGens have been active in the wild. To this end, we estimated the initial appearance of the FeGen instances with the below approaches. For FeGen AAS websites, we checked the creation date of their domains using the WHOIS domain lookup tool [98]. For the AAS websites that host FeGens along with other content, we checked the timestamp of the first snapshot of the FeGen-related pages in the Wayback Machine [77]. For FeGen on-premises software, we checked the earlier date of when the software was published on the software marketplace, and when it was first seen on VirusTotal [5].

With the above approach, we estimated the first appearance of all FeGen instances. Figure 4 presents the distribution of the instances over time. In total, the FeGens were active between 2011-06-14, and 2023-07-09. The earliest one, *IDCreator*, generates fake identification documents, while the latest one, *Paper Work master*, generates fake financial documents. When comparing the distribution of English (●) and Chinese (x) FeGen instances, we noticed that English instances became available to the public relatively earlier and span a longer time window. Specifically, English instances have a mean first-appear time of 2018-02-07, with the first quartile (Q1) at 2015-10-28 and the third quartile (Q3) at 2020-09-05. On the other hand, Chinese instances have a mean first-appear time of 2020-11-15, with Q1 at 2020-01-03 and Q3 at 2022-06-18. Based on earlier observations of fake evidence types, we found that fake identity and financial documents has become a long-standing need among cybercriminals in English-speaking regions. Although these documents are widely believed to be illicit and harmful [55], no effective measures have been implemented to eliminate them over the years. Interestingly, some FeGen retailers may have noticed users' concerns about the legitimacy of using FeGens. However, they seem to try address these concerns by convincing users that they are not spam, rather than directly addressing their legitimacy. For example, a FeGen called `Fake USA Utility Service` [17] responds to users' questions as follows:

Question*: How do I know your site is legit?*
Answer*: Before you make any purchase, we are more than happy to send you any samples of documents. If you are still not satisfied by this then we can put you in touch with some of our existing customers.*

In contrast, Chinese FeGens were introduced at a later stage (Q1 at 2020-01-03). But they are in large numbers, and primarily concentrate on generating fake evidence related to social networks and digital wallets. Also, there has been a noticeable upward trend in the number of FeGens introduced in the past few years, as indicated by the dense clusters of ● markers on the right side of the figure.

> *Finding 2:* FeGen*s have long been popular in the English-speaking world and are gaining popularity in the Chinese-speaking world in recent years. New* FeGen*s are trending toward fake social media and digital wallet evidence.*

Financial gains of FeGen retailers. In the FeGen supply chain, retailers make illicit gains by offering fake evidence services to FeGen users. To understand this financial aspect, we either attempted to generate fake evidence using the FeGens or contacted the retailers to request their payment information. The results show that 52 (51.0%) FeGens offer premium services for generating high-authentic evidence. Specifically, on average, FeGen users need to pay $33.7 for a yearly or unlimited subscription, with a median of $28.6. The minimum cost is $0.3 for generating fake social media chat history, and the highest is $191.6 (or 1399 CNY) for generating fake bank statements (with `http://samsr.com`). Figure 7 illustrates the detailed payment distribution grouped by types of fake evidence. Most English FeGens that provide premium services focus on financial and identification documents, aligning with the overall distribution of FeGens. In contrast, Chinese FeGens that provide premium services are scattered across all types of FeGens.

Table 3 summarizes these payment methods, grouped by the languages of FeGens. The 52 FeGens accept various payment methods with more than one-third (19) accepting multiple methods. In more detail, there are 14 FeGens in English and 38 FeGens in Chinese that require payment. The preferred payment methods differ between English and Chinese FeGens: credit/debit cards and cryptocurrencies are most commonly accepted for English FeGens, while WeChat Pay and Alipay are widely used for Chinese FeGens due to their popularity as digital wallets in China. Interestingly, we noticed that FeGen retailers are promoting the use of cryptocurrencies, including Bitcoin [22] and Ethereum [35], and USDT [34], etc. For instance, at least four FeGens accept multiple payment methods that include cryptocurrencies, credit cards, and PayPal, and they offer discounts of 10% to 30% for cryptocurrency payments. We believe this trend is driven by the anonymity provided by cryptocurrencies, indicating the intent of FeGen retailers to bypass potential scrutiny of financial transactions.

> *Finding 3: Over half of* FeGen*s offer premium fake evidence generation services, and there is a trend in promoting the use of cryptocurrencies for* FeGen *transactions, potentially to evade scrutiny.*

For most payment methods listed in Table 3, assessing the number of transactions happening in the real world is challenging. However, due to the transparency inherent in cryptocurrencies, we were able to calculate the potential illicit gains acquired by some FeGen retailers that use cryptocurrency for transactions. After requesting payment information from FeGen retailers, we identified a total of 14 cryptocurrency addresses including eight addresses on Bitcoin (the most well-known cryptocurrency) and six addresses on USDT (a stable currency pegged to the US dollar and favored by cybercriminals). Seven of the Bitcoin addresses are one-time addresses, making it hard to check the transaction histories of the corresponding FeGens. Therefore, we investigated the transactions of the remaining one Bitcoin addresses and the six USDT addresses. In summary, these addresses received a total of $2.94 million from over 14,048 transactions. We cannot validate whether these transactions are indeed related to purchases of FeGens, as retailers may request variable payment amounts, such as with discounts. However, these figures indicate the potential widespread use of fake evidence generation in the real world, with a substantial number of users might have paid for such services.

Figure 5 displays the account activities of seven cryptocurrency addresses: one Bitcoin address and six USDT addresses (these addresses are listed in Table 10 in the Appendix). Each colored line in the figure represents changes in one cryptocurrency account balance over time. Upon comparing the lines, we found that the distribution of FeGen-related transactions is not uniform across different accounts. For example, one USDT account, i.e., `Account-TMkou`, has a large number of transactions and thus higher account balance, whereas the other accounts have smaller volumes. Interestingly, some USDT accounts appear to interact with each other, as evidenced by transferring funds between them (e.g.,

TABLE 3: FeGen payment methods

Payment Method	# of FeGens (EN)	# of FeGens (CN)
WeChat		19 (50.0%)
Alipay		15 (39.5%)
Credit/Debit card	8 (57.1%)	
Cryptocurrency	8 (57.1%)	6 (15.8%)
App in-app purchase		7 (18.4%)
Bank transfer	2 (14.3%)	1 (2.6%)
PayPal	2 (14.3%)	
Google Pay	1 (7.1%)	
Taobao		1 (2.6%)
Subtotal	14	38

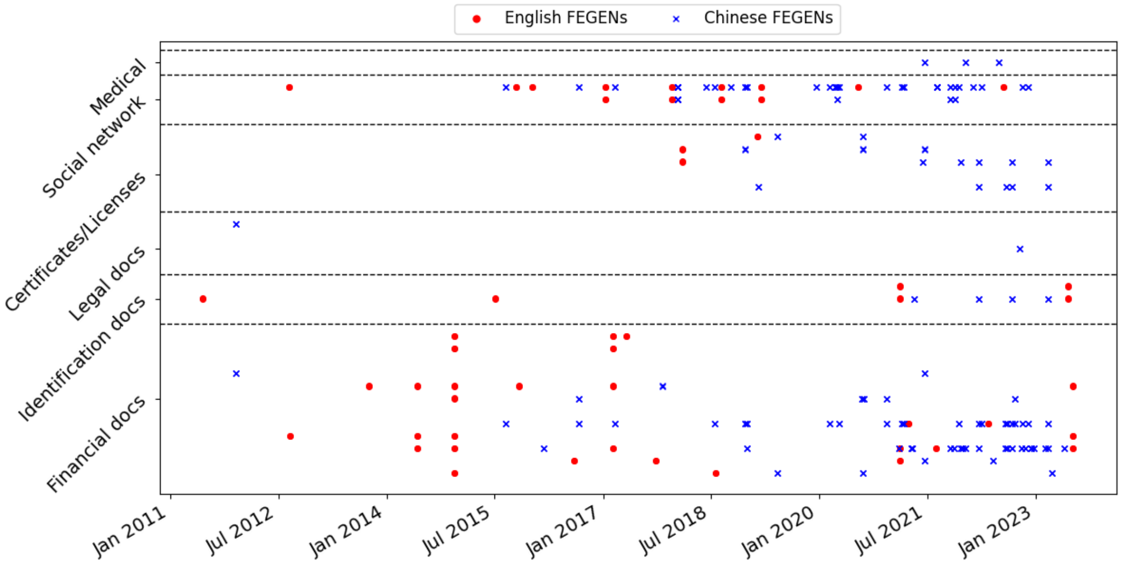

Figure 4: Time distribution of FeGen

Account-TNpDY and Account-TMagt), or by transferring funds to common accounts. Upon closer examination, we found that these USDT accounts correspond to FeGens that offer highly similar services, leading us to suspect that they originate from the same developer or upstream retailers. Furthermore, we noticed that a user of FeGen may utilize multiple FeGens, as indicated by payments to multiple FeGen-related accounts from a same account. In Figure 5, we show the interactions between the FeGen-related accounts and other accounts (i.e., circled), and the arrows indicate the direction of money flow.

Finding 4: Cryptocurrency account activities indicate that FeGens *are in extensive use in the real world, and potentially there exists a retailer network actively promoting them.*

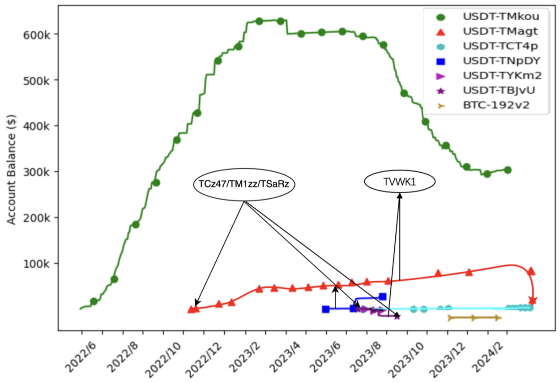

Figure 5: Balances of cryptocurrency addresses involved in FeGens

4.2. FeGen Promotion Channels

As mentioned in Section 2, FeGen retailers are motivated by illicit gains to promote FeGens and increase their deployment among miscreants. We use the FeGen-related search results collected in Section 3.1 as the initial promotion channels. Then, we follow the referred resources in these results to identify other channels through which FeGens are advertised. Based on these channels, we provide a qualitative analysis of how FeGens are promoted.

Public promotion channels. Unlike clearly malicious malware, FeGens can sometimes serve benevolent purposes. For instance, a Fake Bitcoin Wallet FeGen [85] may impersonate an official digital wallet application by mimicking its user interfaces, but claims to be "*for entertainment and pranking your friends only*". As a result, many FeGen retailers have attempted to take advantage of this aspect and promote their FeGens through public channels, such as social media and websites. In particular, with the proliferation of short video platforms, we observed that retailers are increasingly marketing the FeGens by abusing short videos, as seen on platforms like Douyin [36]. In these videos, retailers often emphasize the authenticity and ease of use of the FeGens, encouraging potential users to access the FeGen delivery channels by leaving "*request to access*" comments under the short videos. This behavior enables us to estimate the number of potential FeGen users by examining the volume of comments on the short videos. For example, when we searched for the types of fake evidence on Douyin, we found short videos that correspond to at least 15 unique FeGens, which have a total of 8,838 "*request to access*" comments. This finding confirms our initial hypothesis that FeGens are actively exploited by a significant number of online users. The purpose of "*entertainment and pranking*" with FeGens may warrant a thorough ethical and legal discussion. However, considering that FeGens, if abused, can become effective tools for disrupting trust in cyberspace, a reasonable suggestion for leading short video platforms may be to discourage their customers to use such platforms as an ad hoc promotion channel but instead encourage the use of more authoritative channels, such as official software marketplaces, which are subject to centralized and more robust reviews.

Underground promotion networks. In addition to public channels, FeGens are actively promoted through underground channels. Telegram [30] is a prominent underground channel favored by cybercriminals for its end-to-end encryption and user anonymization capabilities, which makes tracking cybercriminals hard. To quantify the FeGens promoted through this channel, we collected the titles of 81 promotional messages from Google Search (Section 3.1). Then, we searched for keywords from these titles using three Telegram indexing services, i.e., *SuperIndex News* [32], *Telegram Channels* [31], and *TGStat* [33], with the hope of identifying promotional messages related to the same FeGens on Telegram. We successfully identified 19 private Telegram groups promoting at least one FeGen. These groups cover 14 distinct FeGens and host a total of 81,300 group members. Interestingly, 16 groups are dedicated to promoting one specific FeGen, while the remaining three groups (`@hongyegongzuoshi`, `@YSGNB`, and `WL222222`) form a campaign that promotes each other's FeGens. Communication with the group owners revealed that these FeGens are not technically interconnected (since they cover fake bank evidence on different platforms such as smartphone bank applications and desktop clients), but the retailers are collaborating to effectively expose the FeGens to all members of the three groups. Furthermore, we observed that some private groups have a non-negligible number of users who appear to be fake members (e.g., users without a customized profile). For example, although we were unable to validate, we found that some groups, such as `@hao1234ghjgf`, have over 75% of accounts likely to be bulk-registered fake accounts. This is indicated by common account patterns such as having no profile, login history, or any other activities [71]. We believe this practice is used to influence the choices of FeGen users (or potential cybercriminals) because these groups seem more popular and trustworthy due to the presence of fake members.

Finding 5: Telegram has become an important platform for promoting FeGen*s, with collaborative promotional campaigns that provide "store in store" experience, and potentially fake group members to manipulate popularity and trustworthiness of the groups.*

4.3. FeGen Delivery Channels

Characterizing the delivery channels of FeGens can potentially assist in reducing their distribution. In this study, we explored how regular users may access FeGens, e.g., through manual analysis of the FeGen websites and reaching out to FeGen retailers. The aggregated delivery channels are presented in Table 4. Out of the FeGens, 50 are delivered as a service via websites. The remaining FeGens are provided as on-premise software applications through various channels, including official and third-party (3P) stores, retailer websites, and cloud storage services. More specifically, FeGen instances targeting English-speaking users are primarily delivered through websites and official stores, whereas instances for Chinese-speaking users are mainly distributed through cloud storage services, websites, and third-party stores. Notably, most of the cloud storage services delivering the FeGen instances for Chinese users are less well-known to the general public, such as `Telegram cloud storage` [37] (18), `CowTransfer` [38] (3), `tmp.link` [39] (3), `MuseTransfer` [42] (2), and `lanzoucloud` [41] (2). We suspect that these services are also less monitored and that FeGen retailers may have chosen them as a means to potentially evade detection.

TABLE 4: Delivery channels of FeGen instances

Delivery Channel		# of FeGen Instances (EN)	# of FeGen Instances (CN)
As-a-service	FeGen website	24	26
On-premises software tools	Google Play	5	0
	3P Android store	2	16
	Apple App Store	2	10
	3P Windows store	0	5
	Retailer website	0	9
	Cloud storage service	0	31

*A cross-platform FeGen can be distributed through different channels.

Furthermore, 18 FeGens targeting mobile platforms are distributed via 3P stores such as Mi Store [27] and Tencent App Store [29]; in contrast, five FeGens are found on Google Play [21], and 12 FeGens are on the Apple App Store [19]. We are interested in understanding why the official application marketplaces, i.e., Google Play and Apple App Store, did not flag these FeGens as risks. To answer this question, we conducted a manual analysis of the descriptions and user interfaces of these FeGens. These applications essentially fall into two categories: they are either advertised as pranking applications or they bury the FeGen functionality deep within the application. An example from the latter category is the "`Weimai Watermark Camera`" (ID: `1030178866`) on the Apple App Store. The majority of the application provides image and video editing functionalities, while the auto-generation of social network chat history is embedded in a number of "video utilities". This makes it challenging for an application reviewer to cover this hidden functionality manually.

We also want to highlight the difficulties in automatically identifying these delivery channels. An example worth noting is a FeGen named "`chat history generator`". To access its delivery channel, a regular user first needs to request access in its promotional messages (e.g., a blog post or short video) by leaving a comment following a specific format and content. Then, the user will be advised to subscribe to a WeChat Official Account [24] (which was designed to share e-books) and enter a *hidden command* to retrieve the delivery link on Baidu Netdisk [25] with a provided passcode. This process requires the user to interact with at least three parties: the promotional platform, the WeChat Official Account, and Baidu Netdisk, all of which are protected by a passcode or hidden command. This cross-party process

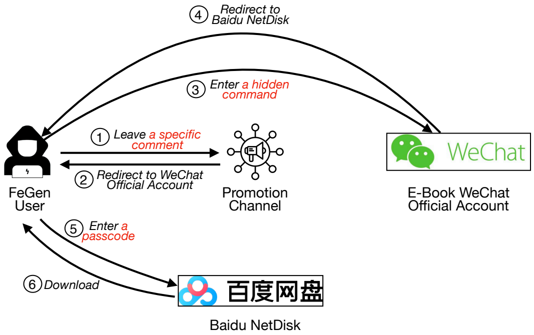
④ Redirect to Baidu NetDisk
③ Enter a hidden command
① Leave a specific comment
② Redirect to WeChat Official Account
FeGen User
Promotion Channel
WeChat
E-Book WeChat Official Account
⑤ Enter a passcode
⑥ Download
百度网盘
Baidu NetDisk

Figure 6: An example of FEGEN delivery process.

poses a serious challenge for us and, we believe, other investigators in designing automated tools to collect and investigate FEGEN delivery channels. Figure 6 illustrates this process by presenting the interactions between these parties.

Finding 6: FEGENs are delivered through a variety of channels, with many of them being less-known, less-monitored by content review, and even evasive.

5. FEGEN Development Analysis

The goal of this section is to report how the FEGENs are developed such that they can generate fake evidence that looks highly authentic. To meet this goal, we either analyzed the FEGENs individually or compared them to each other or to the authoritative software applications that generate real evidence. Particularly, different from earlier reports about FEGEN retailers (the entities managing promotion channels and delivery channels, such as Telegram groups and FEGEN website owners), this section focuses on the practices of FEGEN developers as identified by the software and website author information. It's worth noting that we report on both separately, although the retailers and developers may be the same entity, especially for FEGEN websites.

UI element reuse. The fake evidence in this study usually takes the form of images, i.e., UI screenshots that mimic real evidence. Apparently, many FEGEN developers may choose to reuse certain UI elements, such as icons and layouts, of authoritative software applications and then manipulate these elements in the design of FEGENs. To test this hypothesis, we conducted a cross-comparison between the UI elements of FEGENs and those of authoritative software applications. This comparison includes various types of UI elements, such as images, layouts, fonts, and style sheet files, for all four major platforms on which FEGENs run, i.e., web, Windows, Android, and iOS. We matched either the name or the hash (i.e., SHA256) of the elements to confirm whether they are used by FEGENs.

More specifically, we took all the 8 authoritative software applications used to generate real evidence in Section 3. We unpacked these applications using 7-zip [13], and extracted the above-mentioned resources, and compared to those of the 102 FEGENs (including both on-premises and as-a-service FEGENs). The result shows that most (86.27%) FEGENs share at least one UI element with an authoritative software application. For instance, consider a FEGEN on the Android platform, `com.weijixiang.jietubao`, that generates fake WeChat payment records and chat history. This FEGEN reuses the fonts from the official WeChat application (e.g., `WeChatSansSS-Regular.ttf`), while designing the other parts of the UIs itself. This choice to reuse certain UI elements helps create highly authentic fake evidence without requiring FEGEN developers to reverse engineer the entire authoritative software. Additionally, we noticed that many FEGENs even reuse the icons of underlying operating systems, such as *cellular/Wi-Fi signal strength indicators, battery level indicators, and cellular carrier names*, etc. These FEGENs enable users to generate fake evidence with greater dynamism, using a flexible combination of these icons.

Finding 7: Most FEGENs reuse UI elements of authoritative software, and they support flexible configuration of fake evidence to make it more convincing.

Complete UI impersonation. In addition to reusing certain UI elements, another option for FEGEN developers is to replicate the entire UI design of authoritative software applications. This complete UI impersonation is not commonly observed in FEGENs. We believe the reason is that many authoritative applications involve complex interactions between their UIs and the applications' code, making it difficult to separate the UIs from the applications. Some applications even deploy protective measures, such as obfuscation, to obscure their UIs and code.

However, our study reveals that determined FEGENs targeting mission-critical domains, such as banks, have managed to completely impersonate authoritative UIs. An example of this is a group of seven FEGENs known as "`mobile bank app simulators`". These simulators run on Android, and support the generation of financial evidence for seven leading banks in China, including *ICBC* [10], *CMB* [11], and *BoCom* [9]. A closer look of the simulators reveals that they reuse the exact and entire UI design of the official apps, including UI interactions, and they allow users to edit everything on the UI, such as account balances. As a result, miscreants who use these simulators may create both static evidence in the form of screenshots and record fake video evidence.

Interestingly, when these simulators are installed on the devices of miscreants, they register the same launcher activities [61] as the official applications. In other words, the simulators have the same entries as the official apps. As a result, the simulators can be launched from the app stores directly, making the resulting fake video evidence even more convincing to victims. A demo video of the simulators is available on our website [45]. According to the end-user license agreements (EULA) of these official apps, the behavior of the simulators, which involve *reverse engineering and reusing the app's internal assets* (such as layout and icons), constitutes a serious infringement of the apps' copyright.

Finding 8: Bank-related FEGENs achieve complete impersonation of authoritative bank apps by reusing their entire user interfaces (UIs).

FEGEN templates. FEGEN developers also intend to minimize the development effort by leveraging code templates. We perform a cross-FEGEN analysis to reveal the common code that different FEGEN are built upon. This leads to the discovery of six FEGENs that are built from some website templates, similar to two open-source projects on GitHub [20] and Gitee [28]. We also noticed that four mobile FEGENs leverage a cross-platform UI framework supported by DCloud [26]. These findings confirmed that underground FEGEN developers are indeed benefiting from online code templates and frameworks when creating illicit FEGENs.

Finding 9: Underground developers make use of open-source templates and framework for developing FEGEN*s.*

Anonymization of FEGEN developers. Similar to malware authors who hide their identities to avoid being traced, many FEGEN developers also take steps to maintain anonymity. To quantify this behavior, we analyzed the metadata of all FEGENs to identify potential information that indicates developer identities. Specifically, for on-premises FEGEN, we used off-the-shelf tools (such as `exiftool` and `apksigner`) to read metadata and signing information of executables. For FEGEN websites (i.e., as-a-service), we inspected the contact information on the websites as well as the registrant information of the corresponding domains by querying WHOIS [60]. We consider that the identity of developers is revealed when any piece of personally identifiable information (as directed by [80]), such as a *full name*, *phone number*, *physical address*, and *personal address*, etc., and information that identifies a registered business (e.g., the name and registration of a company). Note that we do not consider a business email as identifiable information (e.g., `support@fegen.com`) since it may not be associated with any specific identities.

With the above methods, we found that 59.7% of FEGENs do not release their developer information, including 76.0% FEGEN websites, 67.7% Android and 55.6% Windows applications. The detailed numbers are shown in Table 5. This percentage is alarmingly high compared to legitimate websites and software applications that are encouraged to disclose their developer identities, e.g., via website "Contact us" [76] or the signer certificates [62]. Notably, our observation indicates that the high percentage is not unintentional. Rather, developers are aggressively concealing their identity. For example, 38 (76.0%) website FEGENs fail to provide any contact information and, at the same time, use *domain privacy services* to hide the registrant information of the website domains, such as `Domains by Proxy` by `GoDaddy`. Some of the on-premises FEGENs use non-identifiable names in their certificate CN (common name) field, or as the company name in the metadata. For example, a FEGEN named `"Fake Bitcoin Wallet"` lists its developer name as `"Android"` with *Google Inc.* Even worse, some FEGENs disguise themselves by using the name of leading companies. For instance, a bank app simulator disguises itself as the 360 antivirus application by *360.cn Inc.*

Finding 10: FEGEN *developers often hide their identity and even disguise their* FEGEN*s as popular legitimate software.*

6. FEGEN Impact Analysis

An important aspect that we haven't discussed is the real-world impact of FEGENs, which is challenging to assess due to the limited availability of public information. To gain preliminary insights into this aspect, we conducted a user study involving hundreds of regular online users. We aim to answer two specific research questions:

- How many online users have been victims of online scams involving fake evidence?
- How effective is fake evidence in enhancing the credibility of online scams? In other words, can end users distinguish it from real evidence?

6.1. Design of User Study

The user study is based on a questionnaire with 12 questions. At the beginning of the questionnaire, we describe the definition of fake evidence in order to help participants understand the context of our survey. The first part of the questionnaire (Q1-Q4) focuses on gathering participants' past experiences with fake evidence. Specifically, we inquire whether participants believe they have ever come across fake evidence online (Q1). If they have, we ask about how often they have seen it (Q2) and the most common types of fake evidence that they have seen (Q3). Then, we assess whether the use of fake evidence has resulted in financial losses for the participants and, if so, the extent of those losses (Q4).

In the next part (Q5-Q9), we determine the distinguishability of fake evidence from real evidence. We first show a pair of fake and real evidence randomly selected from our evidence dataset to the participants, and asked them to identify the real one (Q5) and explain their decision (Q6). Then, we present three individual pieces of fake evidence and ask the participants' confidence (with a range from 1 to 5) to tell they are fake (Q7-Q9). Note that since the fake evidence in the dataset is only from free FEGENs, the results may not fully represent the overall user perception with all types of fake evidence found in the wild.

The third part of the questionnaire collects participants' opinions on the necessity of additional measures, such as legislation, to prevent the dissemination of fake evidence (Q10). In the end, we collect the participants' demographics such as their age range and educational background (Q11-Q12). Such information allows us to better evaluate the impact of fake evidence on the general population. This study was reviewed and approved by the Institutional Review Board (IRB) of our institution.

TABLE 5: Anonymous FEGEN developers

FEGEN Platform	% of Anonymized FEGEN Instances
AAS Website	76.0% (38/50)
Windows	55.6% (15/27)
MacOS	0.0% (0/2)
Android	67.7% (21/31)
iOS	0% (0/14)
Total	59.7% (74/124)

6.2. Results

The survey was conducted anonymously through Wenjuanxing [18], the most popular online survey platform in China, for one week in October 2023. To maintain response quality, we implemented a completion time threshold (i.e., one minute) to filter out invalid responses from bots and reckless participants. In total, we received 208 responses, and fortunately all of them appear valid. Out of the 208 participants, 102 are aged 18-30, 104 are aged 31-60, and 2 are aged over 60. Most of them (87.5%) hold a bachelor's degree or higher. Each participant received 2 CNY as compensation for completing the survey.

According to the responses, the majority of the participants (76.0%) believe that they came across some kind of fake evidence before, with another 14.42% unsure (Q1). 189 (90.9%) participants state that they come across fake evidence at least once per month, i.e., 44 weekly and 145 monthly (Q2). Among those participants, 15 (7.2%) confirmed that fake evidence indeed caused financial loss to them, with an average loss of 575.4 CNY or 78.8 USD (Q4). The types of fake evidence participants encountered roughly align with the distribution of FeGens, with fake social network chat history (81.7%) and fake bank transactions (73.6%) topping the list (Q3).

Finding 11: Online users frequently encounter fake evidence, and a non-negligible number (7.2%) of them experience financial losses.

When presented with both fake and real evidence, 112 participants correctly recognized the real evidence, while the other 96 participants failed to do so (Q5). We used a *Chi-Square test* to determine whether this distribution is significantly different from a random choice. In this case, the expected frequency E for a random choice would be (104, 104), and the observed frequency O is (96, 112). Therefore, the chi-squared statistic ($\chi^2 = \sum \frac{(O-E)^2}{E}$) is 1.23. With one degree of freedom, the associated p-value is approximately 0.27. This p-value exceeds the significance level of 0.05, indicating that we do not have enough evidence to reject the null hypothesis. In other words, the distribution of the answers is not significantly different from a random choice, suggesting that participants are unable to effectively distinguish between fake and real evidence. This result is also confirmed by the reasons behind the participants' choices (Q6). Participants provided various ad-hoc reasons for their answers, such as "*based on my personal feeling*" and "*no reason*", with no single reason agreed upon by more than two participants. Furthermore, when presented with individual pieces of fake evidence (Q7-Q9), participants find it challenging or very difficult to tell the evidence's authenticity, with an average difficulty level of 3.64 on a scale ranging from 1 to 5 (i.e., easy, moderate, challenging, very difficult, and impossible).

Finding 12: It is very difficult for participants to distinguish fake evidence from real evidence.

Finally, the majority of participants (79.81%) believe that further measures should be implemented to restrict the spread of fake evidence (Q10). Notably, although the survey accepts any participants aged 18 and older, the results may still be biased due to the large group of well-educated participants. These individuals tend to be more discerning when it comes to fake evidence and may be less affected. As such, we believe that fake evidence may have more significant impacts on the general public in reality.

7. FeGen Risk Analysis

An imperative aspect to analyze is whether FeGens incur additional security and privacy costs for their users, particularly when considering that users may be benign and innocent, using FeGens for entertainment purposes (as we will discuss in Section 8).

7.1. Overall Risks

To evaluate the security and privacy risks associated with FeGens, we scanned the FeGen instances using the off-the-shelf service VirusTotal [5] – a tool that incorporates over 70 antivirus and URL/domain scanners and is capable of analyzing instances on multiple platforms.

Table 6 displays the distribution of potentially malicious FeGens flagged by VirusTotal, categorized by languages and platforms. In total, 33 (26.6% out of 124) FeGen instances are reported as malicious. Notably, these instances are from 26 FeGens that exclusively target Chinese-speaking users. This finding suggests that Chinese users face a much greater security threat when using FeGens compared to English-speaking users. In particular, the malicious instances are primarily found on the Android and Windows platforms, followed by websites. No instances on MacOS and iOS are reported as malicious. We attribute the discrepancies between Android/Windows and iOS to the fact that iOS instances are exclusively distributed through the Apple App Store, which undergoes rigorous reviews. In contrast, Android/Windows instances are freely distributed online with limited or even missing security scrutiny. Regarding the number of scanners, the majority (28, 84.8% out of 33) of the instances are flagged by at least two scanners, with a maximum of 44 scanners. The remaining five instances are detected by only one scanner, which can be false positives due to the lack of cross-validation of different scanners. We include all these instances in our discussion for the sake of completeness.

The most common threat label is "*trojan*", which applies to 18 instances (12 Windows and 6 Android) that may contain malicious code alongside their advertised FeGen functionalities. Another common label is "*jiagu*", which covers six Android instances, indicating that they are riskware packed by untrusted packers. One of the four website instances is reported as phishing that illicitly collects user sensitive information. The other instances are flagged by generic labels such as *malware*, *grayware*, *PUP (Potentially Unwanted Program)*, etc.

7.2. Case Study: FeGen Predators

A single instance in our FeGen dataset is flagged by an exceptionally high number (44) of antivirus scanners. To figure out the reason, we inspected the promotion and delivery channels, and the instance itself. This results in

TABLE 6: FeGen instances flagged by VirusTotal

Running Platform	# of FeGen Instances (EN)	# of FeGen Instances (CN)
AAS Website	0/24 (0.0%)	4/26 (15.4%)
Windows		13/27 (48.1%)
MacOS		0/2 (0.0%)
Android	0/7 (0.0%)	16/24 (66.7%)
iOS	0/2 (0.0%)	0/12 (0.0%)
Subtotal	0/33	33/91 (36.3%)

the discovery of a security threat that specifically targets FeGens, which we call FeGen predators.

The predator instance, named "*Three Kingdoms*", masquerades as an online banking transfer generator for the Windows platform. VirusTotal classifies it as a "trojan dropper", which is utilized in phishing campaigns by Mustang Panda [82], a Chinese-based threat actor active since 2012. Our case study reveals that this instance is not the official generator but rather a close clone. It leverages the following tactics to mimic the authentic generator. First, the authentic generator is promoted and distributed in a Telegram channel, `@zuotu222`, with a customer service account, `@WL222222`. The predator managed to duplicate all messages in this channel in a fake channel, `@sanguoyanyi222`, and created a deceptive and squatting customer service account, `@WI22222`. Ironically, the predator further claims that all other promotion channels are fake and malicious. To avoid detection, the predator advises its victims to disable on-device antivirus scanners before downloading and installing the predator instance. Second, to make the predator instance appear more similar to the authentic generator, the predator's author adds padding to the dropper, matching the size of the authentic generator (811.1MB). This was confirmed via entropy analysis on the potential padding file named "*data_updating_system.rar*". We suspect that the predator chooses to clone FeGens to ride on their popularity. At the same time, the predator can potentially gain amplified power to reach more victims, thanks to the fact that many FeGen users are online miscreants who possess large amounts of victim data. Telegram data shows that such a predator may have already infected a substantial number of FeGen users, with 19,819 channel members as of this case study in August 2023.

Finding 13: 26.6% of FeGen*s may pose security threats to their users. There is evidence of* FeGen*s being targeted by malware (i.e.,* FeGen *predator) in the wild.*

Discussion. The above empirical analysis of FeGen risks primarily focuses on the malicious aspects of FeGens. Nonetheless, it's important to note that there are other types of risks, such as vulnerabilities in FeGens. For instance, 17 FeGens use HTTP rather than HTTPS for delivering on-premises FeGen software tools or servicing their websites. We don't cover these aspects due to the lack of systematic techniques for evaluating vulnerabilities.

8. Discussion

Ethical Considerations. We pay special attention to ensure that we stay within ethical and legal boundaries. Our work was reviewed and approved by IRB before involving any human subjects in this study (Section 6). We did not pay potential cybercriminals to download any FeGen instances. We located the FeGens by following promotional messages posted on the open Internet, and downloaded the instances that were made freely downloadable through FeGen distribution channels. Also, we did not pay to run FeGen for generating fake evidence in any form, whether through subscriptions or per-document fees.

We reached out to some FeGen retailers pretending to be potential customers in order to obtain more information about FeGens, e.g., through Telegram and QQ. For example, we contacted the retailers to acquire payment account details which enables us to estimate their financial gains. We believe this practice carries low ethical risk, considering the ongoing industrial and research efforts [87], [95] in combating cybercrime which also involve interacting with potential cybercriminals.

Legitimacy of fake evidence and FeGens. Both our user study (Section 6.2), online discussions [54], [83], [88], [93], [99], and recent law enforcement [57], [58], [90], [96] have confirmed that using fake evidence in online activities is harmful. However, there is an ongoing debate about its legitimacy among those involved in fake evidence. FeGen retailers often argue that fake evidence can be used for "entertainment" without causing harm, which may hold some truth although it is disallowed by major software marketplaces such as Google Play [68]. Nevertheless, it is essential to note that, in practice, there are no restrictions to prevent online miscreants from abusing fake evidence. For example, 17 FeGens claim to be for "entertainment purpose only" in their terms of service (ToS), but none of them impose any restrictions on the actual use or present their ToS to users, let alone ask for users' consent. Another common argument by Chinese FeGen retailers is that FeGens can be used by micro-merchants to promote their products and services without violating any laws, such as using fake eCommerce orders to attract users. However, this argument does not hold as such behaviors clearly violate Article 287 of the Criminal Law regarding "*illegal use of information networks*" [1]. Regarding FeGen development, it often involves the unethical practice of reverse engineering authoritative software for monetary or even malicious purposes, which usually violates the end-user agreements of software applications, such as WeChat [23], Alipay [15], and Chase [7].

Limitations. The efforts to minimize ethical risks may limit our capability to perform certain analysis tasks thoroughly and fairly. For example, the fake evidence in the dataset originates solely from free FeGens and may not accurately represent the distribution of fake evidence found in the wild. Hence, our user study (Section 6) only demonstrates the impacts of fake evidence generated by free FeGens, i.e., indistinguishable from real evidence for users. A more thorough analysis would require paying FeGen retailers to generate a more comprehensive set of fake evidence. Similarly, the analysis of fake evidence types (Section 4.1) may not be accurate, as it relies on the advertised functionalities of FeGens on the websites rather than confirmed fake evidence types obtained by actually running the FeGen instances (many of which require payment). Additionally, this study offers only a singular perspective on FeGens that reflects the ecosystem at the time of the study. A longitudinal study over different time points may show more insights into the evolution of FeGens, such as how adversaries update FeGens to keep up with real evidence, which we leave for future investigation.

Other methods for creating fake evidence. In addition to using FeGens, online miscreants may use various methods to create fake evidence. For example, they might run authoritative applications to generate real evidence and then apply graphics editing to create fake evidence. These methods are beyond the scope of this study, and we focus on FeGens, which offer better usability and efficacy.

With the advancement of generative AI models, a natural question arises: can the fake evidence discussed in this study be directly generated using generative AI models for images? To answer this question, we explored several popular online text-to-image generators, including Canva [2], Muse [3], Tiamat [4], and Writesonic [6]. These generators are built upon different AI models, such as DALL-E [40], Midjourney [43], Diffusion model [47], and Photosonic AI [44], known for creating realistic images. Specifically, we used three examples of fake evidence – fake bank statements, fake Alipay transfers, and fake Facebook profiles – to query these generators and manually reviewed the generated images. The results show that none of these generators are capable of producing authentic-looking evidence. Most of them provide less related images (e.g., official icons of a bank rather than a bank statement), while others prevent users from generating fake evidence through content filtering. A detailed result is in Table 7.

TABLE 7: Generating fake evidence with AI image generators

Generative AI	Canva	Muse	Tiamat	Writesonic
Model	DALL.E	Midjourney V4	Diffusion	Photosonic
Facebook profile	✗*	✗	✗	✗
Alipay transfer	✗	✗	✗	✗
Bank statement	✗*	✗	✗	✗

* Failed to generate in the presence of content filtering.

9. Related Work

Generation of fake documents. This study illustrates how cybercriminals can enhance their credibility using fake documents. In addition to that, fake documents can also be helpful in defending against malicious online activities. A common research topic in this direction involves using fake but authentic-looking documents in a decoy system to prevent attackers from obtaining valuable information during a compromise [104]. With advancements in techniques like AI and NLP, a multitude of research studies explore the automated generation of fake documents. For example, early research studies [49], [94], [97] use real documents as input and replace sensitive information (such as email addresses and login accounts) with fake or "bait" information. To safeguard intellectual properties, [56], [72], [100] extend the replacement to content in technical documents, such as textual components, tables and equations, based upon graph representations of document content. Recent studies aim to eliminate the limitations of the above approaches by introducing methods for generating more context-sensitive and realistic documents [75], or without the need for an ontology of bait information [46], [65], [78], with the integration of genetic and NLP algorithms. Compared to the above approaches, the FeGens in this study have two distinctions: they apply to more specific domains related to miscreants' interests, instead of generic technical documents, and are expected to offer greater customization of fake information. As a result, developers of these FeGens tend to use templates for generating highly crafted and customizable documents (Section 5), rather than deploying advanced techniques similar to those used in the above approaches.

Detection of fake evidence. The fake evidence in this study represents a subset of fake online content that mimics output of authoritative sources. Hence, we look into prior studies for detecting fake online content in general. Most prior studies focus on the detection of fake social media accounts (e.g., [48], [51], [59], [79], [103]), fake news (e.g., [67], [81], [89], [101]), and fake images (e.g., [64], [70], [74], [102], [105], [106]). For example, [81] extracts textual features from news articles and uses deep learning models to identify whether the articles are fake. [103] models the connections between online accounts as a graph and uses graph-based algorithms to detect fake accounts. [70] detects deepfakes by modeling the convolutional generative process of generative AI models. With a sufficient amount of fake evidence, the aforementioned solutions have the potential to detect all types of fake evidence. In this study, our focus is on characterizing and understanding the ecosystem of FeGens – the underlying tools for generating fake evidence. We leave the detection of fake evidence for future work. Actually, we have already taken a step towards this direction by creating and sharing a small-sized evidence dataset (Section 3), and are actively enriching it to facilitate future detection solutions.

10. Conclusion

In this paper, we present our systematic analysis of FeGens, including their ecosystem and their impacts on online users. We base our analysis on new datasets that contain over a hundred real-world FeGens and the associated evidence. Through empirical analysis, we characterize the FeGens from a supply chain perspective, leading to a series of new findings, including the tactics used by cybercriminals for development, promotion, delivery, and the risks associated with FeGens. Furthermore, we explore the capabilities of FeGens in generating authentic-looking fake evidence by evaluating whether the evidence is distinguishable from real evidence for normal online users. The results suggest that FeGens are effective tools for enhancing the credibility of cybercriminals in online scams.

References

[1] Article 287 of the chinese criminal law. https://lvlin.baidu.com/question/2084840846971015508.html/.

[2] Canva. https://www.canva.com/.

[3] Muse. https://www.midjourneyai.ai/zh-CN/.

[4] Tiamat. https://www.tiamat.com/.

[5] Virustotal. https://www.virustotal.com/gui/.

[6] Writesonic. https://app.writesonic.com/photosonic/.

[7] Chase. https://www.chase.com/, 1799.

[8] Bank of american. https://www.bankofamerica.com/, 1904.

[9] Bank of communications. https://www.bankcomm.com/, 1908.

[10] Industrial and commercial bank of china. https://www.icbc.com.cn/, 1984.

[11] China merchants bank. https://www.cmbchina.com/, 1987.

[12] Google search. https://www.google.com/, 1996.

[13] 7 zip. https://sparanoid.com/lab/7z/, 1999.

[14] Federal trade commission. https://www.ftc.gov/, 2003.

[15] Alipay. https://global.alipay.com/platform/site/ihome/, 2004.

[16] Facebook. https://www.facebook.com/, 2004.

[17] Fake usa utility service. https://fakeutilities.com/utility-bills/USA-utility-bills/, 2006.

[18] Wenjuanxing. https://www.wjx.cn/, 2006.

[19] Apple app store. https://www.apple.com/app-store/, 2008.

[20] Github. https://github.com/, 2008.

[21] Google play. https://play.google.com/store/games?device=windows&pli=1/, 2008.

[22] Bitcoin. https://bitcoin.org/en/, 2009.

[23] Wechat. https://www.wechat.com/, 2011.

[24] Wetchat official. https://mp.weixin.qq.com/cgi-bin/loginpage?t=wxm2-login&lang=en_US/, 2011.

[25] Baidu netdisk. https://pan.baidu.com/, 2012.

[26] Dcloud. https://www.dcloud.io/, 2012.

[27] Mi store. https://m.app.mi.com/, 2012.

[28] Gitee. https://gitee.com/, 2013.

[29] Tecent app store. https://sj.qq.com/, 2013.

[30] Telegram. https://telegram.org/, Aug 2013.

[31] Telegram channels. https://telegram.org/tour/channels/, 2013.

[32] Telegram superindex. https://telegram.me/s/SuperIndexNews?before=558/, 2013.

[33] Telegram tgstat. https://tgstat.com/, 2013.

[34] Usdt. https://tether.to/en/, 2013.

[35] Ethereum. https://ethereum.org/en/, 2014.

[36] Douyin. www.douyin.com, 2016.

[37] Telegram cloud storage. https://tgstorage.com/, 2016.

[38] Cowtransfer. https://cowtransfer.com/, 2017.

[39] tmp.link. https://www.tmp.link/, 2018.

[40] Dall-e. https://openai.com/dall-e-2, 2021.

[41] lanzoucloud. https://github.com/topics/lanzoucloud/, 2021.

[42] Musetransfer. https://musetransfer.com/, 2021.

[43] Midjournery. https://www.midjourney.com/, 2022.

[44] Photosonic ai. https://photosonic.pro/, 2022.

[45] Supplement materials: Fake evidence generators (fegens). https://sites.google.com/view/fegen/home, 2023.

[46] Almas Abdibayev, Dongkai Chen, Haipeng Chen, Deepti Poluru, and VS Subrahmanian. Using word embeddings to deter intellectual property theft through automated generation of fake documents. *ACM Transactions on Management Information Systems (TMIS)*, 12(2):1–22, 2021.

[47] Rajae Aboulaich, Driss Meskine, and Ali Souissi. New diffusion models in image processing. *Computers & Mathematics with Applications*, 56(4):874–882, 2008.

[48] Yazan Boshmaf, Dionysios Logothetis, Georgos Siganos, Jorge Lería, Jose Lorenzo, Matei Ripeanu, and Konstantin Beznosov. Integro: Leveraging victim prediction for robust fake account detection in osns. In *NDSS*, volume 15, pages 8–11. Citeseer, 2015.

[49] Brian M Bowen, Shlomo Hershkop, Angelos D Keromytis, and Salvatore J Stolfo. Baiting inside attackers using decoy documents. In *Security and Privacy in Communication Networks: 5th International ICST Conference, SecureComm 2009, Athens, Greece, September 14-18, 2009, Revised Selected Papers 5*, pages 51–70. Springer, 2009.

[50] Bert Brys, Stephen Matthews, Richard Herd, and Xiao Wang. Tax policy and tax reform in the people's republic of china. 2013.

[51] Qiang Cao, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. Aiding the detection of fake accounts in large scale social online services. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pages 197–210, 2012.

[52] Eoghan Casey. *Digital evidence and computer crime: Forensic science, computers, and the internet*. Academic press, 2011.

[53] Eoghan Casey and Curtis W. Rose. Chapter 2 - forensic analysis. In Eoghan Casey, Cory Altheide, Christopher Daywalt, Andrea de Donno, Dario Forte, James O. Holley, Andy Johnston, Ronald van der Knijff, Anthony Kokocinski, Paul H. Luehr, Terrance Maguire, Ryan D. Pittman, Curtis W. Rose, Joseph J. Schwerha, Dave Shaver, and Jessica Reust Smith, editors, *Handbook of Digital Forensics and Investigation*, pages 21–62. Academic Press, San Diego, 2010.

[54] CECN. Fake bank statement sales online: A six-month statement costs less than 200 yuan, bank staff claim unable to discern authenticity. http://finance.ce.cn/bank12/scroll/201807/26/t20180726_29856464.shtml, 2018. In Chinese.

[55] Andrew Ceresney. Speech financial reporting and accounting fraud. https://www.sec.gov/news/speech/spch091913ac/, Sep 2013.

[56] Tanmoy Chakraborty, Sushil Jajodia, Jonathan Katz, Antonio Picariello, Giancarlo Sperli, and VS Subrahmanian. A fake online repository generation engine for cyber deception. *IEEE Transactions on Dependable and Secure Computing*, 18(2):518–533, 2019.

[57] Federal Trade Commision. Ftc shuts down purveyors of fake documents used for fraud, identity theft. https://www.ftc.gov/news-events/news/press-releases/2018/09/ftc-shuts-down-purveyors-fake-documents-used-fraud-identity-theft/, Sep 2018.

[58] Federal Trade Commission. Federal trade commission 2018 privacy and data security update. https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2018/2018-privacy-data-security-report-508.pdf/, Dec 2018.

[59] Mauro Conti, Radha Poovendran, and Marco Secchiero. Fakebook: Detecting fake profiles in on-line social networks. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 1071–1078. IEEE, 2012.

[60] Leslie Daigle. Whois protocol specification. Technical report, 2004.

[61] Android Developers. Android launcher activity. https://developer.android.com/reference/android/app/LauncherActivity.

[62] Android Developers. Sign your app: Keystores, keys, and certificates. https://developer.android.com/studio/publish/app-signing#certificates-keystores.

[63] Lesley Fair. “authentic fake” financial documents challenged in three genuine ftc cases. https://www.ftc.gov/business-guidance/blog/2018/09/authentic-fake-financial-documents-challenged-three-genuine-ftc-cases/, Sep 2018.

[64] Hany Farid. Image forgery detection. *IEEE Signal processing magazine*, 26(2):16–25, 2009.

[65] Yun Feng, Baoxu Liu, Yue Zhang, Jinli Zhang, Chaoge Liu, and Qixu Liu. Automated honey document generation using genetic algorithm. In *Wireless Algorithms, Systems, and Applications: 16th International Conference, WASA 2021, Nanjing, China, June 25–27, 2021, Proceedings, Part III 16*, pages 20–28. Springer, 2021.

[66] Roger Allan Ford. Data scams. *Hous. L. Rev.*, 57:111, 2019.

[67] Sherry Girgis, Eslam Amer, and Mahmoud Gadallah. Deep learning algorithms for detecting fake news in online text. In *2018 13th international conference on computer engineering and systems (ICCES)*, pages 93–97. IEEE, 2018.

[68] Google. Policy center - deceptive behavior. https://support.google.com/googleplay/android-developer/answer/9888077.

[69] Mark Griffiths and Adrian Parke. Internet gambling. In *Encyclopedia of Internet Technologies and Applications*, pages 228–234. IGI Global, 2008.

[70] Luca Guarnera, Oliver Giudice, and Sebastiano Battiato. Deepfake detection by analyzing convolutional traces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 666–667, 2020.

[71] HAMZA_X_PRO. How to identify a fake telegram account. https://www.hamzaxpro.com/post/how-to-identify-a-fake-telegram-account.

[72] Qian Han, Cristian Molinaro, Antonio Picariello, Giancarlo Sperli, Venkatramanan S Subrahmanian, and Yanhai Xiong. Generating fake documents using probabilistic logic graphs. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2428–2441, 2021.

[73] Kevin Hong. The complicated truth about china’s social credit system, 2019.

[74] Chih-Chung Hsu, Chia-Yen Lee, and Yi-Xiu Zhuang. Learning to detect fake face images in the wild. In *2018 international symposium on computer, consumer and control (IS3C)*, pages 388–391. IEEE, 2018.

[75] Yibo Hu, Yu Lin, Erick Skorupa Parolin, Latifur Khan, and Kevin Hamlen. Controllable fake document infilling for cyber deception. *arXiv preprint arXiv:2210.09917*, 2022.

[76] iDreamBiz Team. The importance of ’about us’ and ’contact us’ pages for a website. https://www.idreambiz.com/blog/the-importance-of-about-us-and-contact-us-pages-for-a-website.html.

[77] Hadi Jahanshahi, Mucahit Cevik, José Navas-Sú, Ayşe Başar, and Antonio González-Torres. Wayback machine: A tool to capture the evolutionary behavior of the bug reports and their triage process in open-source software systems. *Journal of Systems and Software*, 189:111308, 2022.

[78] Prakruthi Karuna, Hemant Purohit, Sushil Jajodia, Rajesh Ganesan, and Ozlem Uzuner. Fake document generation for cyber deception by manipulating text comprehensibility. *IEEE Systems Journal*, 15(1):835–845, 2020.

[79] Sarah Khaled, Neamat El-Tazi, and Hoda MO Mokhtar. Detecting fake accounts on social media. In *2018 IEEE international conference on big data (big data)*, pages 3672–3681. IEEE, 2018.

[80] U.S. DEPARTMENT OF LABOR. Guidance on the protection of personal identifiable information. https://www.dol.gov/general/ppii/.

[81] Jun Lin, Glenna Tremblay-Taylor, Guanyi Mou, Di You, and Kyumin Lee. Detecting fake news articles. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 3021–3025. IEEE, 2019.

[82] Malpedia. Mustang panda - actor profile. https://malpedia.caad.fkie.fraunhofer.de/actor/mustang_panda, 2023. Accessed: 2023-10-26.

[83] People’s Daily Online. "entertainment" software becomes an illegal wechat business fishing tool. http://it.people.com.cn/n/2015/0327/c1009-26758213.html, 2015. In Chinese.

[84] Google Play. App - bank of china. https://play.google.com/store/apps/details?id=com.boc.bocsoft.bocmbovsa.buss&hl=en_US&gl=US.

[85] Google Play. App - fake bitcoin wallet. https://play.google.com/store/apps/details?id=co.za.binarymatter.bitcoinwalletfake&pli=1.

[86] Sean Quagliani. Four common ways fraudsters create fake income documents. https://www.linkedin.com/pulse/four-common-ways-fraudsters-create-fake-income-sean-quagliani.

[87] Tom Simonite. Microsoft chatbot trolls shoppers for online sex. https://www.wired.com/story/microsoft-chatbot-trolls-shoppers-for-online-sex/.

[88] Sina. Exposure of the black industry chain of fake bank statements: Just 180 yuan can deceive millions in housing loans. https://gd.sina.com.cn/jingji/news/2018-07-27/detail-ihfvkitx4513183.shtml, 2018. In Chinese.

[89] Vivek K Singh, Isha Ghosh, and Darshan Sonagara. Detecting fake news stories via multimodal analysis. *Journal of the Association for Information Science and Technology*, 72(1):3–17, 2021.

[90] Sohu. Fake transfer screenshots exposed: Man convicted of deceiving others in transactions involving fish, dogs, and a bracelet. https://www.sohu.com/a/481228299_121080990, 2021. In Chinese.

[91] Toine Spapens. Illegal gambling. *The Oxford handbook of organized crime*, pages 402–418, 2014.

[92] Jia-Rong Sun, Mao-Lin Shih, and Min-Shiang Hwang. A survey of digital evidences forensic and cybercrime investigation procedure. *Int. J. Netw. Secur.*, 17(5):497–509, 2015.

[93] VOGUE. The dark true story behind princess diana’s explosive bbc interview. https://www.vogue.com/article/the-dark-true-story-behind-princess-dianas-explosive-bbc-interview, 2022.

[94] Lei Wang, Chenglong Li, QingFeng Tan, and XueBin Wang. Generation and distribution of decoy document system. In *Trustworthy Computing and Services: International Conference, ISCTCS 2013, Beijing, China, November 2013, Revised Selected Papers*, pages 123–129. Springer, 2014.

[95] Peng Wang Wang, Xiaojing Liao Liao, Yue Qin, and XiaoFeng Wang. Into the deep web: Understanding e-commercefraud from autonomous chat with cybercriminals. In *Proceedings of the ISOC Network and Distributed System Security Symposium (NDSS), 2020*, 2020.

[96] weixin_39535287. Man sentenced for developing and selling fraud-assistance software. https://blog.csdn.net/weixin_39535287/article/details/112201340, 2021. In Chinese.

[97] Jonathan White and Dale Thompson. Using synthetic decoys to digitally watermark personally-identifying data and to promote data security. In *Security and Management*, pages 91–99. Citeseer, 2006.

[98] Whois.com. Whois domain lookup. https://www.whois.com/whois/.

[99] Xinmin. Transaction records, id cards created at the click of a button: "fake generator" nets 300,000 yuan per month. https://wap.xinmin.cn/content/31766083.html, 2020. In Chinese.

[100] Yanhai Xiong, Giridhar Kaushik Ramachandran, Rajesh Ganesan, Sushil Jajodia, and VS Subrahmanian. Generating realistic fake equations in order to reduce intellectual property theft. *IEEE Transactions on Dependable and Secure Computing*, 19(3):1434–1445, 2020.

[101] Junxiao Xue, Yabo Wang, Yichen Tian, Yafei Li, Lei Shi, and Lin Wei. Detecting fake news by exploring the consistency of multimodal data. *Information Processing & Management*, 58(5):102610, 2021.

[102] Jiachen Yang, Shuai Xiao, Aiyun Li, Guipeng Lan, and Huihui Wang. Detecting fake images by identifying potential texture difference. *Future Generation Computer Systems*, 125:127–135, 2021.

[103] Dong Yuan, Yuanli Miao, Neil Zhenqiang Gong, Zheng Yang, Qi Li, Dawn Song, Qian Wang, and Xiao Liang. Detecting fake accounts in online social networks at the time of registrations. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 1423–1438, 2019.

[104] Jim Yuill, Mike Zappe, Dorothy Denning, and Fred Feer. Honeyfiles: deceptive files for intrusion detection. In *Proceedings from the Fifth Annual IEEE SMC Information Assurance Workshop, 2004.*, pages 116–122. IEEE, 2004.

[105] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-attentional deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2185–2194, 2021.

[106] Yipin Zhou and Ser-Nam Lim. Joint audio-visual deepfake detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14800–14809, 2021.

A. Appendix

List of FeGens. Table 8 lists the FeGens that have instances in the form of AAS websites, and Table 9 lists the FeGens that have installable software tools.

Cryptocurrency addresses of FeGen retailers. Table 10 displays the list of cryptocurrency addresses used by FeGen retailers, with two Bitcoin addresses and six USDT addresses.

Payment distribution. Figure 7 shows the distribution of payments for FeGens that provide premium fake evidence generation services.

TABLE 8: Information about the FeGens and their instances in the form of AAS websites

FeGEN Name	URL of FeGen AAS Websites
12tool§	http://www.12tool.com
Airline Ticket Generator†	https://app.letongkj.com/index/Core/index.html?id=133
Allbusinesstemplates†	https://www.allbusinesstemplates.com
Appbs†	https://sc.appbs.cn/
Baituling§	https://www.baituling.com/
Bank Check Generator 1†	http://www.mtu9.com/zb/zhipiao/
Bank Check Generator 2†	https://www.hashemian.com/tools/check-generator.php
Bank Statement Fake	https://bankstatementfake.com
BANKDOCS	https://fakeutilities.com
BANKSY	https://www.fakebankstatement.co.uk
BANKUS	http://www.banknovelties.com
ChatTree§	https://www.chatree.cn/
D and More	https://www.diplomasandmore.com
Dvgod†	https://tool.dvgod.com/
Fake Details Generator†	https://fakedetail.com/
Fake Documents Online	https://www.fakedocuments.online
Fake Flight Tickets§	https://www.fakeflighttickets.com
Fake iPhone Text Messages†	http://iphonefaketext.com/
Form Pros	https://www.formpros.com
Generate Status†	https://generatestatus.com
Haozhengming§	https://www.haozhengming.cn/
Hixiaopa†	https://c.tianhezulin.com/
Hongbao	https://show.verydog.cn/hongbao
IDCreator§	https://www.idcreator.com
iFake Text Message†	https://ifaketextmessage.com/
Jietubao	https://www.jietujie.com/
Liaotiantu§	https://www.liaotiantu.com/zb/8
Marriage Certificate Generator 1†	http://zb.yanluwei.cn/jiehunzheng/jieguo.php
Marriage Certificate Generator 2†	http://www.zuiwuliao.cn/funny/41.html
Medical Record†	http://www.zuiwuliao.cn/funny/342.html
Online WeChat and Alipay Generator†	https://www.goodsunlc.com/status/screenshots/
Paper Work Master	https://paperworkmaster.com
Paystub Generator	https://www.thepaystubs.com
Phone Gaps†	https://www.phonegags.com
Replace Your Docs	https://www.replaceyourdoc.com
Runner Toolbox†	https://runnerstool.newrathon.com/cert-calculator
Sozz†	http://www.sozz.cc
Taobao Order Generator†	https://shengcheng.6cm.co/taobao-order.html
Tweetgen†	https://www.tweetgen.com
ValidGrad§	https://validgrad.com
Verif Tools§	https://verif.tools/en/
Vjietu†	https://www.vjietu.com
Vjietu Pro	https://vjietu.pro
WeChat Chat History Generator†	https://zixiwangluo.github.io/wxdh/
Weixin Duihuaqi	https://weixinduihuaqi.com/
WhatsApp Fake Chat†	https://www.fakewhats.com
Xiaobeizi†	https://app.ippapp.com/screenchat/
Yijietu†	https://1jietu.com/
Zeoob†	https://zeoob.com
Zjietu†	https://www.zjietu.com/

†FeGen instances that we ran and exported fake evidence successfully.
§FeGen instances that can generate fake evidence but don't allow export.

TABLE 9: Information about the FeGens and their instances in the form of software tools

FeGEN Name	Platform	App ID / Package Name	Version Number	Hash‡
2345 Kantuwang	Windows		9.0	8c9b39e2754571e11996978ea7c7c2f9b4139daf82c8d4284b767874cde5c937
Aizimu	Android†	cn.nineton.sayingwrod	3.0.7	3a8c4755fc7188aa1388e34952e38ce04764867e1156966943f35159fe13d118
	iOS†	1457476705	3.1.8	
Bank USDT Generator	Windows		na	f5aee9f19397cecc99d75b899f94e8668a16a6ef4cc2ad2110982637bc44b857
Bigcircle	Android	com.huanhuatec.bigcircle	2.0	e77489a0235378b2c5de14f7896b8929e6dbf6c35aeaf7a617d3108dccee4d83
Bitbit	Android	com.bit.bitebit	1.0.7	5b82760d33d465f37587c6c4af289f7c2afaa51b9aa18f6b519d122de1de40c6
	Android	com.ljm.vipjt	1.1.6	f61b7256427412fa98199bd90725e12a32526e9906332d787f1d3eda22eb6b41
Business License Generator	Windows†		8	f876b160c13cd23bf114821f0fe2270178e2cd0d7fe7c0abde60ce5fceff9a59
Cash Prank Maker	Android	com.fakecompany.cashapppayment	20.4	2984b014698faf3120e7a4e813da67a31f67593d790b3e3cc7f769e79b369182
Chat Generator	Android	com.hehax.chat_create_shot	6.4.9	a38880857167361 7438f69122ba822fe00945ec94970768469147abf644a9403
Chatlike	Android	com.jinshengyuan.chatlike	3.4.9	18a3918c11edc7a1786702b3ceace98bd8540306cff4603ee9092e8b00200709
ChatTree	Android§	com.chatree.chat	1.5.2	a8b3d1392e867a18a9582a2d27189c35030140668452c6e0297686c922776af4
	Windows§		2.6.1	f0d615681b554e07de07f40c0d751eac1f4514dbc4750b1e868ead0c329ddf89
	macOS§		2.6.1	93bfaf4a56f0087e4d82a76c8eaea107600cf8811e239ae839a1bbd17f153153
Diploma Generator	Windows§		8.5	e115729f9190ac7bafbe4b8413f8bbc73cc4e0d97b412f4158dc9f74fbed726d
Driver License Generator	Android†	com.fillevoss.fake.driver.license.generator	1.1	65567fb01f8de4f6b246ace685d23a2f2882de996c9e139fb1d68e94cb4f962c
DY Studio	Windows		4.23	6a56aa8bde875261c211e4365c03f29f44cb90cfcb765d4121dd29ff36b373d0
Fake All	iOS	1518594150	1.1	
Fake Bank Checks	Android	com.christapps.fakebank.checks.cheque.pro	1.0.4	51a7caa6d27705d51180c4a06ad059496b9b6f8c814d9b23f392c3fa07131dec
Fake Bitcoin Wallet	Android†	co.za.binarymatter.bitcoinwalletfake	1.3	943dfea6cd1aa74a4e1c1acdba214bb20b2aca410eb13dc3fab1bb97e028cb57
Fake Call	Android†	com.fungame.fakecall.prankfriend	2.8.0	c7e4bf81560f6421c2a4059178b7f1dde5f3eb35fbd8d1e30e85ab4aba4cd2e8
Fake Call and Sms	Android	dev.qt.hdl.fakecallandsms	9.9.7	9bea608e01a5077a46ab93cc6f8faea6dc8462d8ea49ef418d14fd8befd56aad
Fake Call-Prank Caller ID	iOS†	1081494198	4.6	
Feisu Zengzhishui	Windows		2.6	8cf877054ea1cb60d2f2e825596de4ab2b2b74c3e586a8ff599472d953655c8d
Hixiaopa	Android†	scq.myapp	1.3.6	154da4bba678fa81e9e642016bd16965d203052935d0b54317a48975076b97ea
	iOS†	1571465759	1.3.6	
Hongye	Android	com.chinamworld.main	9.0.1	4f1785af88aa3af97169fe5621daabfa38b3bf808465ad9f37222f2d73e6ffb3
Huijietu	Windows		4.3	929898be05fb9b5e35a7d9dd1c9aa4108b6b07172c8132d0de1afef6e23c62ae
ID Card Generator	Windows†		1.1.0	abcbed46dcfab9a3b578b431b0dbf2e6e531de6b3cb2f19a1a91fc779af3a190
	macOS†		1.1.0	52e8bd2c795ff63e60498e9134cd842b93d74eb23ad0402bff9e22c5a7fad660
iMMcque	Android§	com.iMMcque.VCore	5.0.2	b744cc1156f86ca8c541229b7bf940e9187fc37f630acfa2fc5bf2886547c3ee
	iOS§	1034388318	6.4.7	
Jietubao	Android	com.weijixiang.jietubao	2.2.6	44d380a39a4b4ffd9892eec39fbe698556665be0c61c206eed3fcec4a959c715
	iOS	1591871897	1.0.4	
Leidian	Windows		9.0	251b74a49bf1c68d54c28b5fe360a0e60da83cee427b95a69a8051d5b4dd004d
Leidian Zuotuwang	Windows		5.1.2	19e0380e591370a37812131fb4235a4198e4162a73f6c659e24eb94e79e7fa64
Medical Record Generator	Android	com.jshare5.zdbg	1.0	f3bdaed169376e1628e9357f7ebba2361402a636c52487476810b327efb6f0b3
Niupian	Android†	cn.niupian.niupianwang	2.0.6	b5b85da26d7b769b043be4a1e3410a709006175fd0649daf2b30ac7b4b928e54
	iOS	1514916687†	2.0.6	
Online Bank Generator	Windows§		2.1	e5b1fe278ba9884fc093f88938f5dec50b5d8bdefcbf633ec1daadfc3aba10a5
Qnwsjtw	Android	com.kuowendianzi.qnwsjtw	3.8.6	5ca33d472f21fbd993691f2a691de664259f038b23c305a14de6fa86b47888f3
Qukadian	Android	com.xmy.videoclip	1.1.16	e62254cc9103f1fbcd20d16e520fdd031d65fd4c0e49f4c06c4eae7ead704fca
Quwan	Windows		3.6	ad9c65a02bb530015295d3f3715783e002ef8813107e9daf4b1aa5257066dc95
Sanguo Yanyi Daqiao	Windows		5.6	5003931f7870ef42418f6de2727230412ef2ded90e7502d97ef2d637c1d6b960
Simulated Business Licenses	Windows		na	b1758007633237f24f93026686d75d48494efbc65a03a8bd77927ee8b371bb4a
Three Kingdoms	Windows		5.0	e9f20d634cc8f4831ed0b2a8347d03e332b5529be8bb79e44e275d48781d4bec
Three Kingdoms 1	Windows		9.1	f3e285131b189368852333a81b31294b8334c1aa00fed9b176ad0ab67e0694d7
Three Kingdoms 2	Windows		9.1	77e48a9ce0b3d1588ed724142f07d32b4909c8ac879a7b42299fdd57005227c0
Three Kingdoms 3	Windows		9.1	cd85c6216673533f6f3362f69f80c8977258a01b25e7a79954ffb4aa9e0737c2
Three Kingdoms 4	Windows		8.4.0	c2efb46d98ae2504a8c4daa6e57231b7bf3148788ce7ff79f5944c57715546be
Vjietu	iOS†	1633186528	1.2.2	
	Android†	com.superman.yijietu	1.2.2	be6712fffcb9f9da4e9bd28ac5b410c379e44dc5303a57912154ac3fe715ce85
Vjietu Pro	Android	com.vjietu.pro	1.0.8	70e8005847aaa94605701ef889a02cc22a5be345e120fad4bbf19867147196c4
	iOS	6446059507	1.0.3	
Wanhua	Android	com.xishuai.wanhua	2.4.6	330cb9de34fefa9b4a7a1f8a9a951b40929d4bd40e550c94c90f1364dc5d069f
WeChat Generator	Android	com.wanjietu.wxdh	5.3.0	770466f90a9d14ec947d3c2f3b7c35fbbfb82453807590e5e48df53eb2a311bd
Weimai	Android§	com.dingYZ.dingyzwater	5.3.80	aa78361658e6190b4892db939e02438bcc95a331eac1e0e0e7babf06d8fe0c60
	iOS§	1030178866	5.3.80	

FeGEN Name	Platform	App ID / Package Name	Version Number	Hash
Weishang Jietuwang	iOS†	1478554328	1	
Weishangxiu	Android	quan.the.puzzle	1.3	a56054741f7c8373654fb39c778705a5a8833110985477a3d61f25f1e4437df7
	iOS	1537484457	2.2.9	
Xiaowai Weishang	Android	com.YiGeTechnology.XiaoWai	1.2.1	faecbfd1d154ad32d0c1efe5fe7919347a787cd9bb462f88703f8c297e4c2664
	iOS	1544742237	1.0.9	
Xiongmao Jieji	Windows		2.5	6999830a6ef48a7cc1b214f45a3a4ceffb7bfa8a2b0fec88e420b7229dee62e3
Xiongmao Kantuwang	Windows		3.1	0a6dc67c3fec61b6b3c6ae7a52510298fa1052e0eef7bd3df7818856937f1cf1
Xiongmao Zuotuwang	Windows		6.9	eff21a2ad8fb9b6fda0e372d9ffa947e04de2100314aaceed5610cb7f895c763
Xiongying Jieji	Windows		6.3.5	3ddbeb239fb613dd85bb1bd17ac5975c4167a42bc6bfdb230529df695a844380
Yijietu	Android†	com.superman.yijietu	1.2.2	be6712fffcb9f9da4e9bd28ac5b410c379e44dc5303a57912154ac3fe715ce85
	iOS†	1633186528	1.2.2	
YintuPro	Windows		2.2	97311f3aa5f74eecd723f55e91e782a9e918b4c502c0050755479d5c959caf5d
Yitu	Android	com.ieba5b69d51cea3a6	1.3.5	8c68f59b0ab5b2ff61036c9b7e114dfba78fbbf79498dded28e1b1de807bd810
Zhuantu	Windows		na	587c17ae88db3a3e12c32f5d35cf20b6b4fde68a404f8ea76e66d5e1fad29851
Zjietu	Android†	com.screenshot.making	2.2.1	a9b83d7fc111f447734e846cfd00de790e977417b2a7ce9f3f5763034aff0fdb
Zuotuwang	Windows		na	a987f25d587d1fbc8317a73f946d8d091ca9bf2af6caca5c638bdd9bdcbf8b88

†FeGen instances that we ran and exported fake evidence successfully.

§FeGen instances from which we can generate fake evidence for review, but failed to export the evidence.

‡We do not include the hashes for iOS apps since the hashes vary with the purchaser of the app.

TABLE 10: Cryptocurrency Addresses of FEGEN Retailers

	Cryptocurrency	Address
FEGEN-related accounts	USDT-TRC20	TMagtCZBCiX3Azc6kpsh39XTXtxbKvJ9Bz TCT4pmo3XDia8w47AvW5NnDCwtPvnuC4Et TMkouC8NFbWuQGbzonnq1rjrEieYRV92sz TNpDYk1C8MN4wQoygzKHGS46FCMGfQRbKb TYKm2GCFqZF42PUFx3QB4mg9MBehbm6ZY3 TBJvU63MiJehyWAPFGiQLoMBRt8wqrFZE2
	BTC	192v2JZbdGrYiw6BWAm6yenX9pRLv85dCf

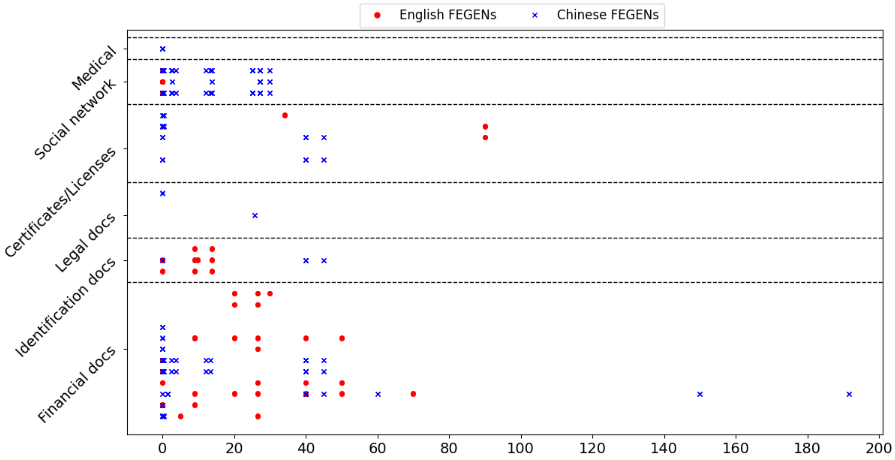

Figure 7: Payment distribution of FEGEN