



Universiteit
Leiden
The Netherlands

Algorithms for analyzing evolving networks on the Dark Web & in science

Boekhout, H.D.

Citation

Boekhout, H. D. (2026, March 17). *Algorithms for analyzing evolving networks on the Dark Web & in science*. Retrieved from <https://hdl.handle.net/1887/4297227>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4297227>

Note: To cite this publication please use the final published version (if applicable).

Early warning signals for predicting cryptomarket vendor success using dark net forum networks

In this chapter we focus on identifying key players in dark net cryptomarkets that facilitate online trade of illegal goods. Law enforcement agencies aim to disrupt criminal activity conducted through these markets by targeting key players vital to the market's existence and success. We particularly focus on detecting successful vendors responsible for the majority of illegal trade. Our methodology aims to uncover whether the task of key player identification should center around plainly measuring user and forum activity, or that it requires leveraging specific patterns of user communication. We focus on a large-scale dataset from the EVOLUTION cryptomarket, which we model as an evolving communication network. Results indicate that user and forum activity, measured through topic engagement, is best able to identify successful vendors. Interestingly, considering users with higher betweenness centrality in the communication network further improves performance, also identifying successful vendors with moderate activity on the forum. But more importantly, analyzing the forum data over time, we find evidence that attaining a high betweenness score comes before vendor success. This suggests that the proposed network-driven approach of modelling user communication might prove useful as an early warning signal for key player identification. This chapter is based on:

- H. D. Boekhout, A. A. Blokland, and F. W. Takes. Early warning signals for predicting cryptomarket vendor success using dark net forum networks. *Scientific Reports*, 14(1):16336, 2024

3.1 Introduction

The dark net, a part of the internet that requires specific software or authorization to access [62], hosts a myriad of online fora that are increasingly a hotbed for criminal behavior and radicalisation [109, 40]. Dark net fora can, both theoretically and empirically, be split into those functioning as meeting places for the exchange of criminal information and those where criminal goods and services are traded, i.e., criminal marketplaces. These fora and marketplaces can serve up to hundreds of thousands of users. They are often moderated and organized in a professional manner, with cryptocurrencies, such as Bitcoin, serving as currency and are therefore referred to as *cryptomarkets* [103, 128]. To efficiently coordinate its activities disrupting these cryptomarkets, law enforcement agencies aim to target key players that are vital to these market's existence and success [66, 128].

Key players include the administrators and moderators responsible for the existence and proper functioning of the cryptomarket. However, they also include the more successful vendors that are responsible for the majority of the trade conducted on the cryptomarket. Identifying which users function as administrators can often be as easy as looking at the titles assigned to them on the cryptomarkets' forum. Similarly, if sales statistics were shared on the cryptomarket, currently successful vendors would be easily identifiable. We study the EVOLUTION cryptomarket, a dataset with more than half a million posts and over four thousand vendors (see Chapter 2). It is moreover one of the few cryptomarkets that recorded sales information. However, many cryptomarkets do not record sales information and at best provide a label to vendors independent of their success. Furthermore, it is nearly impossible to identify those vendors whose success is yet to come. Yet, if law enforcement agencies wish to disrupt future sales, it is exactly these future successful vendors that they would need to identify in order to dissuade them from continued participation in the cryptomarket. Therefore, in this chapter we focus on identifying key players in the form of both current and future successful vendors.

Existing research studying the workings of cryptomarkets and aimed at assisting law enforcement agencies in identifying key players, often uses methods such as topic modeling or sentiment analysis [108, 137, 105, 2, 79, 95]. These methods rely on (combinations of) commonly used words and sentence structures in the forum message contents. However, the rise of the use of message encryption in criminal communication, calls for the development of methods not reliant on knowledge of message content. In this work, we aim to develop a method to identify key players based on the temporal structure of their communication network alone; thus, ignoring message content entirely.

Communication networks model the interaction between entities within commu-

nication systems, such as mobile phone [94, 39, 121], face-to-face [91], and social media communication [149, 98]; but also communication through online fora [66, 78]. Online fora, including those associated with cryptomarkets, usually consist of *topics*, which may be grouped by subject. Each topic is started by one user with a first message, also called a *post*, and allows the set of users with access to respond by placing their own posts. This activity can be considered a form of indirect communication from the posters to those users who placed posts on the same topic before them. We can model this indirect communication using what we call a user-to-user communication network that directly connects users that posted in the same topic. At the very least, a link in such a network represents a shared interest in the same topic as well as a level of familiarity with one another due to the likelihood of having seen each others' posts. At best, a link can signify direct communication between two users that are, by means of forum posts, responding to one another. Thus, links represent potential social ties formed on a dark net forum. In this work, we leverage the structure of these communication networks without relying on knowledge of message content, with the goal of identifying and predicting successful vendors.

To find important users in a (criminal) network, one of the most commonly used approaches is to apply network centrality measures, that rank users based on their position in the network [97, 66, 39, 121, 91, 149]. Different network centrality measures often imply different roles a given user plays within a network. In this chapter, we explore four different measures: degree, harmonic closeness centrality, betweenness centrality and PageRank. This allows us to grasp what type of role, as defined by the user's structural position in the network, may be more suited to the task of identifying key players in cryptomarkets. The nuances of the interpretation of centrality measures can vary depending on whether we account for edge weights, i.e., the strength of social ties, and edge directions, i.e., who responds to whom. Therefore, we consider for each measure whether the direction and strength of social ties matters, for identifying (successful) vendors for law enforcement applications.

Beside network measures, several intuitive straightforward measures can be obtained directly from the forum data. We consider three such measures: post activity, topics started, and (started) topic engagement. We henceforth refer to these measures as forum *activity indicators*. The rationale behind these three activity indicators, relies on vendors' tendency to start topics to promote their listings [95, 64] and the concept of *name recognition*. Name recognition, also called brand awareness in a marketplace context, has been linked to improved trust [42] and market outcomes [75] (e.g., more sales). Furthermore, Duxbury & Haynie [60] concluded that trustworthiness is a better predictor of vendor selection than product diversity or affordability.

In this chapter, we investigate to what level employing network measures computed on user-to-user communication networks are useful in identifying both current and

future successful vendors on cryptomarkets. We look at three law enforcement applications, each increasingly more useful to law enforcement practitioners. We investigate whether (1) network measures can be used to distinguish vendors and their level of success; if (2) rankings induced by network measures can narrow down the user base to a significantly smaller set of potentially relevant users for law enforcement agencies to investigate; and, (3) to what extent the top ranked users include successful vendors and other key players. Furthermore, we study the EVOLUTION cryptomarket at different points in time, i.e., we look at various *snapshots* of the communication network. By doing so we simulate law enforcement agencies investigating the state of the cryptomarket at those specific points in time, while subsequent data, i.e., at that point future data, shows how the cryptomarket would progress without intervention. Consequently, we propose a methodology with the potential to serve as an *early warning signal* for future vendor success on cryptomarkets.

The remainder of this chapter is structured as follows. In Section 3.2 we shortly describe the dataset and measures used before reporting on our results. The results and their implications for law enforcement are discussed in Section 3.3. Then, Section 3.4 provides more in-depth descriptions on the dataset and network extraction as well as the activity indicators, network measures, and evaluation metrics used in this work. Finally, Section 3.5 discusses several supplementary results.

3.2 Results

In this section we first discuss our dataset in Section 3.2.1 and the (network) measures for identifying key players in Section 3.2.2. Next, we report and interpret results for the task of distinguishing vendors from non-vendors and predicting the levels of vendor success in Section 3.2.3. Then in Section 3.2.4, we explore to what extent the rankings induced by (network) measures can reduce the set of users for law enforcement agencies to investigate, while still including the greatest share of successful vendors. Finally, we look at the set of top ranked users for the most promising network centrality measure and activity indicator at a specific point in time in Section 3.2.5. We do so to establish how well represented key players are among these top ranked users.

3.2.1 Data

In this study we focus on the cryptomarket EVOLUTION. EVOLUTION was active from January 2014 until March 2015, when it closed due to an exit scam. At the time, it was one of the most popular cryptomarkets [128]. It formed a combination of a carding forum, where card information (e.g., credit/debit/ID/etc.) is traded, and an underground drug market [56].

We obtained raw data of the EVOLUTION marketplace and forum from the Dark Net Market archives [34]. From this, we extracted a structured dataset, established a method of linking the market and forum data, and subsequently extracted a communication network(s). The extraction, linking, and data quality resolution processes as well as the network extraction process, the resulting dataset(s), and various statistics on the dataset and its completeness, are presented in Chapter 2. The dataset itself is directly available on Zenodo [20] at <https://zenodo.org/records/10171217>. Parameters of the network extraction procedure control respectively the bounds on when two posts constitute a social tie (δ_ϕ and δ_t) and the strength of the social tie (ω_{lower} , t_{lim} , and ω_{first}). For the communication network(s) studied in this work, the same extraction procedure and parameters were used as those used in Chapter 2, i.e., $\delta_\phi = 10$, $\delta_t = 1$ month, $\omega_{lower} = 0.2$, $t_{lim} = 7$ days, and $\omega_{first} = 0.5$. We demonstrate the robustness of our findings for each of these parameters in Section 3.5.1.

The cryptomarket EVOLUTION observed two notable changes in user and post activity. In the initial months up to May 2014, the cryptomarket underwent steady growth in terms of both post activity and the number of active users. However, monthly post activity stabilised from May until October (see Figure 3.7). Notably, May saw a change in the vendor ranking system, which assigns textual labels to vendors that are visible on the marketplace to potential customers and imply a level of success and trustworthiness. Obtaining a label representing greater success and trustworthiness as a vendor now required sufficient positive feedback, but most important for us, the new ranking system also reported on the exact number of sales a vendor had made up to that point. The second major change to the cryptomarket came in early November 2014, as a by-product of the closure of six cryptomarkets following the joint international law enforcement operation dubbed “Onymous” [128]. After this disruption, EVOLUTION showed a significant increase in overall activity until its closure.

Both the communication networks and current and future sales counts were extracted on a monthly basis using data up to the end of each month, including all data prior to the given month. Consequently, we obtained 15 network snapshots (starting from January 2014 up to March 2015). Note that we rely on all data prior to the given month and not only the most recent month(s), because a vendor’s reputation plays an important role in their success and is based not just on the most recent activity. In fact, the build-up reputation is such a vital aspect that it is predominantly the successful vendors with a large number of sales and high reputation who choose to migrate and maintain their identity in new cryptomarkets after market closures [112]. Details on the network extraction process and the computation of monthly sales statistics are provided in the Section 3.4.

3.2.2 Network measures and activity indicators

Each considered network measure captures a different role a user may play within the user-to-user communication network. To cover a wide range of user roles that may be important to vendor success, we report on four centrality measures: (1) in-degree; (2) bidirectional harmonic closeness centrality; (3) directed weighted betweenness centrality; and (4) directed weighted PageRank. The *in-degree* of a user indicates the number of different users that posted (shortly) after them on the same topic(s). Thus, it can serve as an approximation of how many users have seen one or more of their posts and thus to some extent their level of name recognition. The bidirectional *harmonic closeness centrality* [124] is a measure of a user's ability to reach the entire network, following paths regardless of link direction. High harmonic closeness centrality indicates that it should be relatively easy to reach and therefore potentially be visible to the entire user base. The directed weighted *betweenness centrality* [68, 33] computes how often a user lies on shortest paths connecting other nodes, taking into account both the direction and strength of social ties. High betweenness nodes often lie "between" communities. Therefore, it may be a good measure of how well a (potential) vendor reaches different, otherwise separated, communities of customers. Finally, the directed weighted *Pagerank* [114] computes the probability that a random walker that infinitely traverses a network ends up at a given node, taking into account both the direction and strength of social ties. High PageRank centrality is often an indicator of being well connected to other important users. Duxbury & Haynie [60] found that buyers were more likely to continue ordering with vendors within the same community. Thus, a close connection with other key players, as indicated by a high PageRank value, can be indicative of a high perceived trust, positively affecting sales. Finally, we note that links in the communication network are temporally independent unless they rely on the same post(s), i.e., the link (a, b) is not dependent on the existence of link (b, c) unless they were formed based on the same post by user b . Consequently, were we to consider only time-respecting paths (e.g., as introduced by Kempe et al. [80]), which require temporally dependent links, we would not adequately capture the social aspect of the network, i.e., the desired concepts of familiarity and shared interest. Therefore, we focus on "static" network measures.

To evaluate the network measures we compare them against three activity indicators, which serve as our baselines. These activity indicators can be computed directly from the forum data, so without aforementioned communication network extraction, are intuitively meaningful in the context of cryptomarket vendor success and also do not require knowledge of message content. We consider: (1) post activity; (2) topics started; and (3) topic engagement. *Post activity* refers to the number of posts a user has

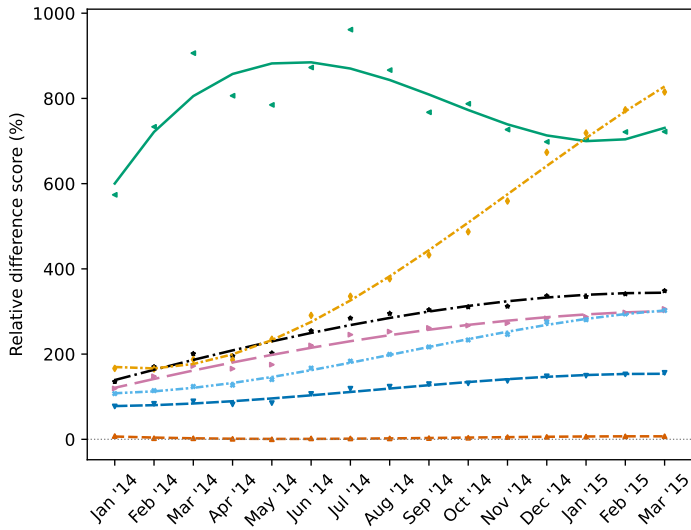
placed on the forum. It relies on the idea that greater activity means greater visibility, which in turn leads to greater name recognition. *Topics started* determines the number of topics a user started and *topic engagement* subsequently computes the total number of posts placed within those topics, regardless of who posted them. These measures rely on the fact that the more topics a user has started and the more engagement those topics received, the greater the likelihood that they are a (successful) vendor. This is supported by Armona [2], previously concluding that a similar measure of vendor forum sentiment could be indicative of higher demand for a vendor on the Agora cryptomarket. Whereas their measure relied entirely on forum post texts (and thread titles) for the selection of posts and computation of the sentiment, our activity indicators can be determined entirely independent of post content. Again, the increased visibility through starting topics also boosts name recognition.

Further details on the computation and interpretation of the measures is provided in Section 3.4.

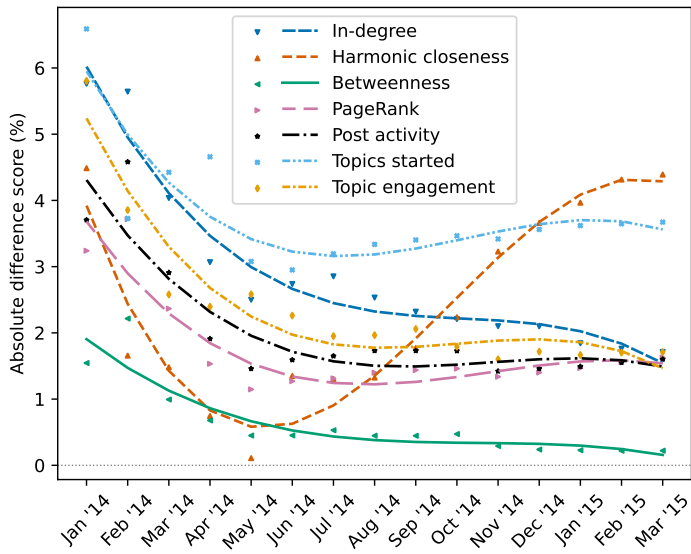
3.2.3 Distinguishing vendors and their level of success

To predict vendor success, we must determine if it is possible to distinguish between vendors and non-vendors, as well as between various levels of success. We look at the average network centralities and activity indicators for groups of users, in an attempt to distinguish groups with greater success. To this end, we divided, for each month, all active vendors, i.e., all users that are or will become vendors with at least one post already posted at that time, into five groups of success percentiles, each including respectively the top 0–20%, 20–40%, etc. of vendors in terms of sales. We refer to these groups as *vendor percentiles*. Separate vendor percentiles are formed for current and future success. We refer to the most and second most successful percentiles as the *top* and *sub-top percentile*, respectively. The *non-vendors*, consisting of regular forum users and those vendors with no recorded sales at all, form a separate sixth group.

First, we computed for each month the mean normalized value for each measure for the groups of all vendors and all non-vendors, using min-max normalization. From this the *relative* and *absolute difference scores* between vendors and non-vendors was computed for each of the four network measures and three activity indicators (see Section 3.4.5 for more details on their computation). The resulting scores are depicted in Figure 3.1. In these figures, lines give a third degree polynomial approximation of the trend based on the monthly centralities and activity indicators. Here, the third degree polynomial is used to try to account for the two aforementioned events that took place in the EVOLUTION cryptomarket [19]. Dashed lines are used for the network measures and dotted lines for the activity indicators. Figure 3.1 shows that, for all measures, vendors have higher network centralities and activity indicators than



(a) Relative difference score



(b) Absolute difference score

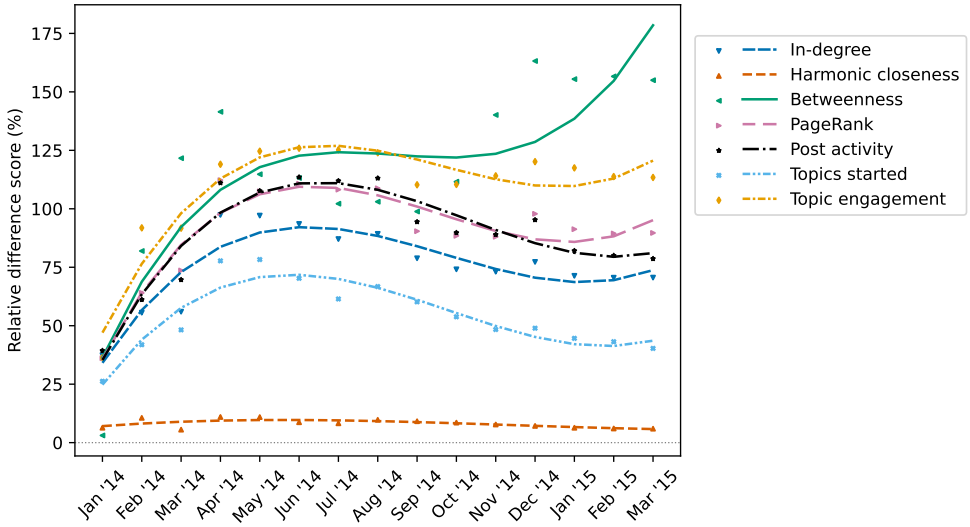
Figure 3.1: The relative (a) and absolute (b) difference score between vendors over non-vendors. Positive scores indicate that vendors achieve higher normalized network centralities or activity indicators than non-vendors on average.

non-vendors. Furthermore, they show that although the relative difference score for betweenness centrality of vendors over non-vendors is quite significant (600–1000%), the corresponding absolute difference score is the smallest of all these measures. This indicates that betweenness centrality has relatively small values overall with some extremely high outliers. On the contrary, harmonic closeness centrality has low relative difference scores but nominal absolute difference scores. Since these effects are expected to disappear when inducing a ranking from the actual values, it is less the size of the difference scores than the fact that they are positive that are an indicator of (useful) predictive power. After all, the ranking induced by the centralities and activity indicators is more useful to law enforcement practitioners than the actual values. Thus, the exclusively positive values in Figure 3.1 indicate the potential of all network measures and activity indicators to distinguish vendors from non-vendors.

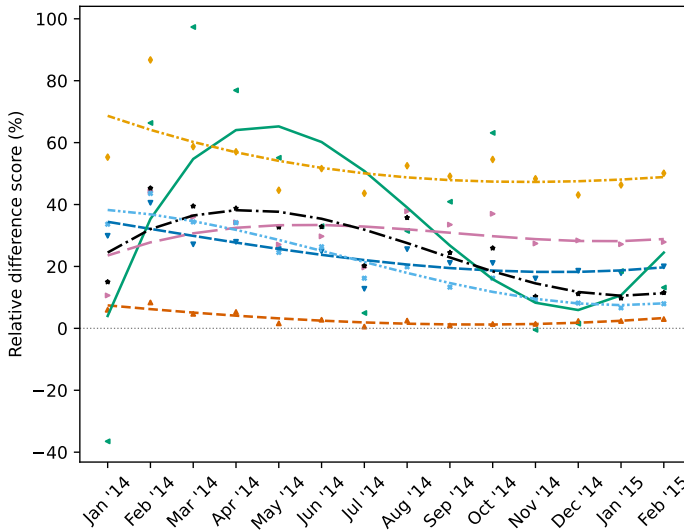
Next, we investigate whether these measures can also distinguish between the levels of success of vendors. To assess this, we looked at the relative difference scores between the top percentile and all vendors (Figure 3.2) and between the top and sub-top percentile (Figure 3.3) for both current and future success. Figure 3.2a shows that for all measures the currently most successful vendors have on average higher network centralities and activity indicators. After the first month and with the exception of July and August 2014 for betweenness centrality, Figure 3.3a demonstrates this also holds when comparing the top with the sub-top percentile. Interestingly, trend changes for most measures follow cryptomarket developments. For example, up until May the difference score increases monthly, similar to how the level of activity on the cryptomarket increased during this period. The following period, up to the November 2014 “Onymous” disruption [128], shows stable but slightly decreasing difference scores for most measures. Finally, after this disruption, we see a small increase in difference scores again.

When we consider future success, Figure 3.2b shows again positive difference scores between the top vendor percentile and all vendors. However, they are noticeably lower than for current success. Similarly, Figure 3.3b shows mostly positive difference scores when comparing with the sub-top percentile, but with lower scores. Thus, for both current and future success the network centralities and activity indicators show the potential to distinguish vendors’ level of success.

Notably, betweenness centrality shows trends that differ from all other measures. In particular, for current success we clearly see higher difference scores in the last months. On the contrary, for future success the final months show lower difference scores than before. This behaviour is likely due to the delay between successful vendors establishing themselves in the network and reaping the benefits in terms of sales. In other words, high betweenness centrality is expected to be more a prelude to than a consequence of vendor success. Thus, these results show the potential of betweenness

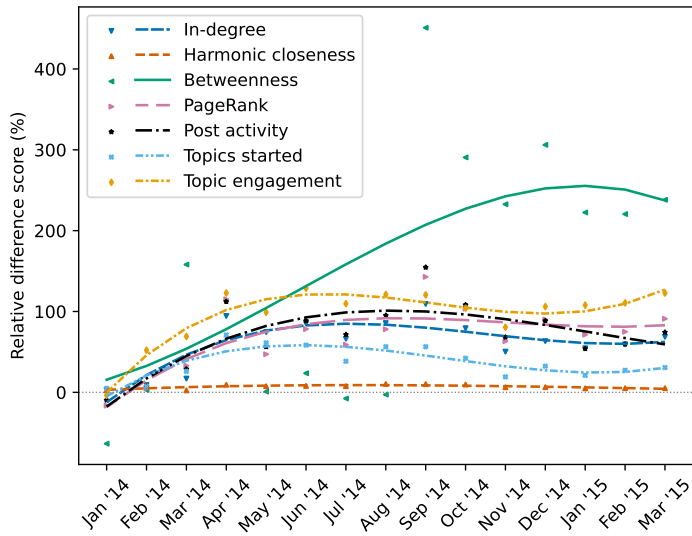


(a) Current success

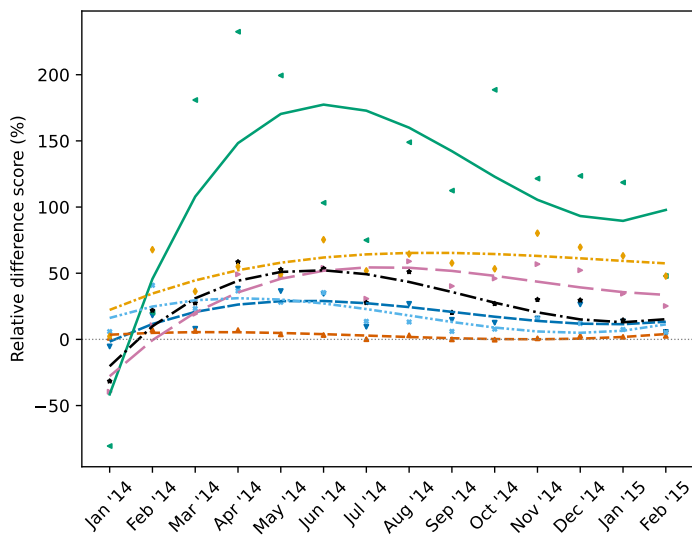


(b) Future success

Figure 3.2: The relative difference score between the top percentile and all vendors. Positive scores indicate on average higher normalized network centrality or activity indicators for the more “successful” group.



(a) Current success



(b) Future success

Figure 3.3: The relative difference score between the top and sub-top percentiles. Positive scores indicate on average higher normalized network centrality or activity indicators for the more “successful” group.

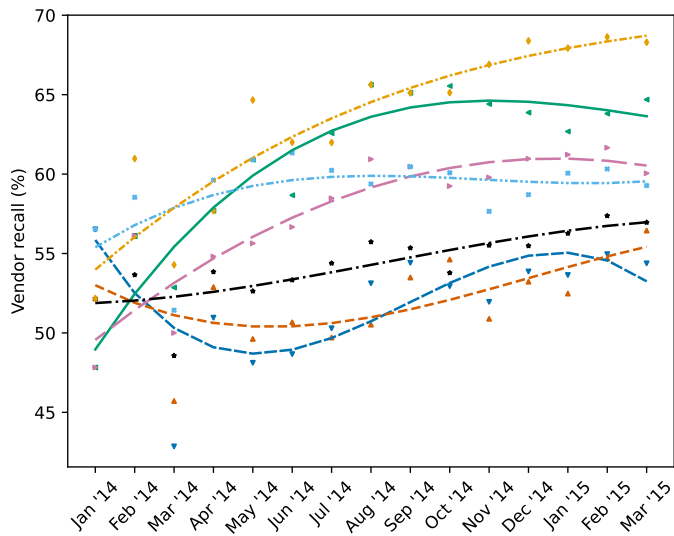
centrality as an early warning signal for future vendor success.

In short, for all measures vendors show positive difference scores over non-vendors and less successful vendors. Thus, rankings induced by these measures are expected to rank successful vendors (relatively) higher. Therefore, the induced rankings have the potential to assist law enforcement agencies by allowing them to focus investigative efforts on higher ranked users. Furthermore, betweenness centrality was shown to have potential as an early warning signal, as high betweenness appears to precede vendor success. Finally, among the remaining network measures and activity indicators, topic engagement consistently showed the highest difference scores. This suggests that topic engagement may provide the best predictions of vendor success.

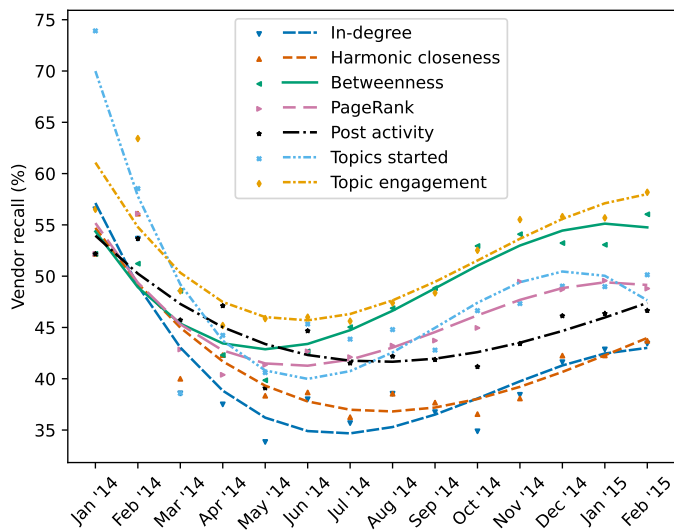
3.2.4 Detecting vendors in the user base

In their efforts to disrupt cryptomarkets, law enforcement agencies have access to limited personnel and resources. One method employed by law enforcement agencies to deal with this limitation, is to reduce the set of users to investigate based on a ranking induced by some measure. Rankings that after such a reduction still include many users of interest, are of course preferable. In the previous section, we established the predictive potential of the network measures and activity indicators for predicting (successful) vendors. Now, with the specific law enforcement perspective of aiming to find as many (hard to identify) vendors as possible, we want to explore how this predictive potential translates to the task of reducing the set of users to investigate. To do this, we consider what we call the *vendor recall*. The vendor recall computes what percentage of users among the top vendor percentile (the top 20% of vendors) is also among the top percentile of all users, i.e., among the top 20% of all users when ranked on a given network measure or activity indicator (see Section 3.4.5 for further details). Note that we focus on the top percentile, instead of the absolute top vendors, as this aligns with the law enforcement intervention method of dissuading continued participation in the cryptomarket. Since this intervention method is known to be ineffective for the absolute top vendors and comes at a relatively low cost to law enforcement agencies, it is more suited to targeting larger groups of vendors. Furthermore, although a vendor's sales volume is merely a proxy for their trade volume, we may reasonably expect those with the largest trade volume to be among the top vendor percentile in terms of sales. For these reasons we also prioritize reporting the recall of vendors over sales. Monthly vendor recalls are plotted in Figure 3.4 for current (a) and future (b) success, respectively. As noted before, lines in these plots are third degree polynomial approximations of the trend.

Figure 3.4 shows that, for both current and future success, degree and closeness centrality generally have a worse vendor recall than any of our activity indicators. From



(a) Current success



(b) Future success

Figure 3.4: Monthly vendor recall of top vendor percentile (top 0–20% vendors in terms of sales) among the top 20% of all users based on the network measures and activity indicators. Plots cover recall for both current and future success. Higher vendor recall indicates a greater portion of the top vendor percentile was found.

Table 3.1: Mean (and standard deviation) of the monthly overlap between network centrality based and activity indicator based detected vendors for the top vendor percentile (top 0–20% of vendors in terms of sales) as shown in Figure 3.4. The best values are shown in bold. (Abbreviations of activity indicators: pa = post activity, ts = topics started, and te = topic engagement.)

	Current success				Future success			
	pa	ts	te	pa \cup te \cup ts	pa	ts	te	pa \cup te \cup ts
	% of indicator detected vendors also found by centrality (higher is better)							
In-degree	90.3	77.9	78.4	72.3 \pm 4.49	83.8	70.4	72.1	66.0 \pm 4.60
Harmonic closeness	88.9	77.1	75.5	70.3 \pm 4.04	83.9	72.1	70.6	66.2 \pm 7.09
Betweenness	89.8	83.1	85.4	80.4 \pm 5.23	82.5	77.4	79.5	74.1 \pm 4.33
PageRank	95.7	83.2	86.5	79.6 \pm 2.78	90.9	76.4	82.4	74.6 \pm 4.14
	% of centrality detected vendors also found by indicator (lower is better)							
In-degree	94.1	88.1	94.8	98.7 \pm 2.04	93.6	83.1	91.9	97.8 \pm 4.04
Harmonic closeness	92.5	87.1	91.2	95.9 \pm 2.27	91.3	82.9	87.7	95.4 \pm 3.90
Betweenness	80.4	80.7	89.0	94.4 \pm 2.25	75.9	75.6	83.8	90.6 \pm 2.88
PageRank	90.6	85.4	95.0	98.8 \pm 1.14	89.0	79.1	92.3	97.0 \pm 2.02

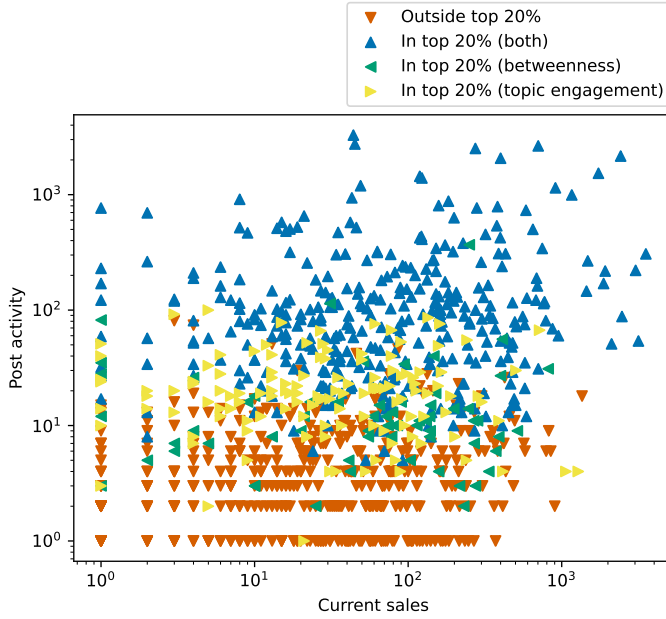
May onwards, PageRank outperforms post activity and performs on par with the topics started indicator. Meanwhile, from July onwards, betweenness centrality consistently outperforms both the post activity and topics started activity indicators and performs (nearly) on par with topic engagement. Overall, the topic engagement indicator most consistently achieves high performance in terms of vendor recall. These observations tell us two things. First, network centrality measures require the communication network to have developed and stabilised sufficiently before achieving reliable vendor recall. During the initial months the communication network and its structure are still undergoing significant changes. Consequently, we also see large fluctuations in vendor recall for the network measures between these months. Second, network measures do not strictly improve on our best activity indicator(s) in terms of vendor recall.

Despite achieving the best vendor recall, topic engagement is only able to detect up to 2/3rd of the most successful vendors for current success and even fewer for future success. Thus, there may still be a significant number of successful vendors that are not detected by the activity indicators, but that are included by network measures. To investigate this, we analyse the overlap of detected vendors between the network measures and activity indicators. Table 3.1 shows the average monthly overlap of each network measure with each individual activity indicator and the union of detected vendors by all activity indicators. We see that PageRank and betweenness centrality detect the greatest share of vendors also found by the activity indicators, detecting on average approximately 80% of all current vendors and 75% of all future vendors found. However, respectively nearly 99% and 97% of all vendors detected by PageRank are also found by the activity indicators. Thus, PageRank is not able to identify many new

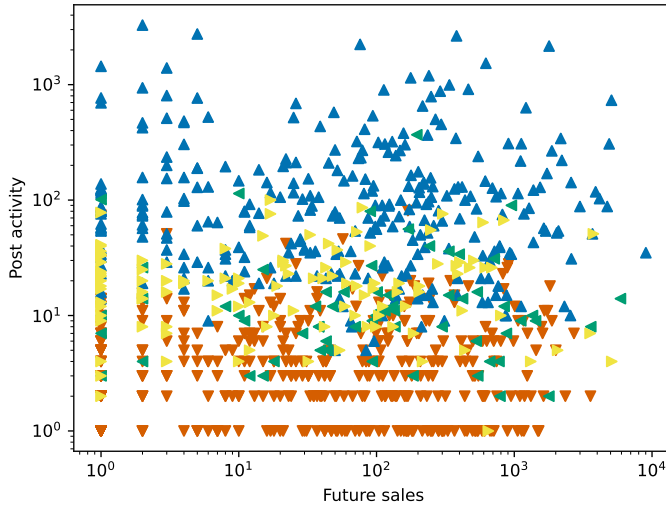
vendors. On the contrary, the activity indicators find respectively only 94% and 90% of the vendors included by betweenness centrality. Notably, individual indicators find far fewer. Thus, betweenness centrality is able to detect the largest share of successful vendors not included by any of the activity indicators. Therefore, reducing the set of users for law enforcement practitioners to investigate using betweenness centrality may provide a fresh perspective.

Despite finding additional vendors, the union of all successful vendors detected by betweenness centrality and all activity indicators only finds around 75% and 65% of the top percentile for current and future success, respectively. This means there is still a significant segment of the most successful vendors that would not be found for any of these measures. One possible explanation for scoring low on any of these measures is simply low posting activity. To assess whether this holds for the successful vendors that do not score high enough to be detected, we look at what we call the *post activity recall* of the top vendor percentile in Section 3.5.2. The post activity recall is the percentage of the top vendor percentile's total post activity, for a given month, that is associated with those vendors detected with vendor recall (see Section 3.4.5 for further details). We find that for both current and future success, the vast majority of post activity is associated with the vendors with high network centrality and activity indicators. Therefore, low post activity can be considered the main reason for the relatively low vendor recalls we observe. After all, although over 30% of successful vendors are not found, they are responsible for less than 10% of the post activity of the entire group (in most cases less than 4%). Indeed, Figure 3.5 shows that any vendor with activity above a certain threshold is always among the detected vendors, while most vendors with very few sales are not. Specifically, it demonstrates that for both topic engagement and betweenness centrality for September 2014, this threshold is below 100 posts (as confirmed by Figure 3.6). This also holds for the other centrality measures, as demonstrated in Section 3.5.3. We note that vendors with low post activity are also much less likely to be found using other methodologies. Therefore, applying the methods discussed in this chapter is likely not to miss vendors that other methodologies might have found. Thus, the relatively low vendor recall achieved by betweenness centrality and topic engagement should not discourage law enforcement practitioners from using them.

Figures 3.5 and 3.6 further indicate that vendors are overall more likely to be identified the greater their respective success. This is also demonstrated through *sales recall*, which measures what percentage of sales of the entire top percentile the detected vendors are responsible for (see Section 3.4.5 for further details), in Section 3.5.2. There we show that the sales recall is generally between 10–20% higher than the corresponding vendor recall. This indicates that the detected vendors are, on average, the more successful vendors. In Section 3.5.3 we show that this finding

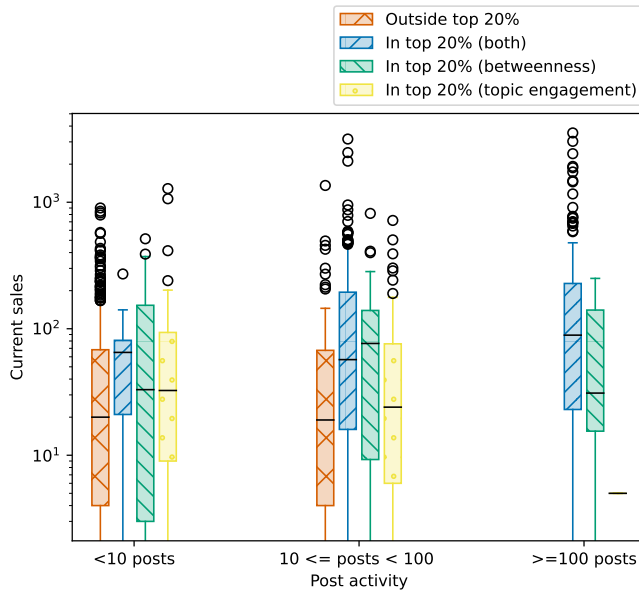


(a) Current success

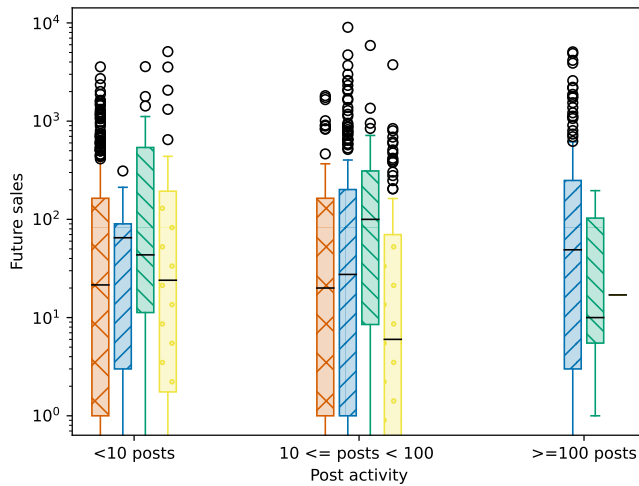


(b) Future success

Figure 3.5: Scatter plots of sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for topic engagement, betweenness centrality, and their intersection for September 2014, for current (a) and future (b) success.



(a) Current success



(b) Future success

Figure 3.6: Box plots of sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for topic engagement, betweenness centrality, and their intersection for September 2014, for current (a) and future (b) success.

also holds for other months and network measures. Furthermore, from Figures 3.5 it appears that the vendors found by topic engagement, and not betweenness centrality, are generally slightly more active and less successful compared to those found by the inverse, i.e., by betweenness centrality and not topic engagement. Indeed, Figures 3.6 confirm that, for vendors with between 10 and 100 posts, those found exclusively by betweenness centrality are generally more successful. Notably, the effect seems to be even stronger for future success and this is moreover confirmed to hold for other months in Section 3.5.3. This observation once more highlights the potential of betweenness centrality as an early warning signal.

Throughout this section we have considered a single threshold at which we cut-off the rankings, namely 20% of all users. In Section 3.5.4, we investigate the performance of the measures at different thresholds. We observe that for low false positive rates, up to around 20%, our findings hold. For higher false positive rates, however, topic engagement clearly outperforms all other measures. However, given the limited resources of law enforcement agencies, it is unlikely that such large user samples would ever be considered for investigation. After all, the resources required to investigate even 20% of users would likely exceed those available to law enforcement agencies. Additionally, we find that topic engagement is the best measure for predicting vendors, regardless of their level of success.

To summarise, topic engagement provides the best single measure recall performance. Meanwhile, betweenness centrality identifies the greatest share of vendors that do not score high for any of the activity indicators. Additionally, betweenness centrality detects the most vendors of all network measures. Therefore, betweenness centrality is the network measure most likely to be of use to law enforcement practitioners for detecting vendors in the user base. Furthermore, betweenness centrality uniquely finds relatively more successful vendors among those with moderate activity. Notably, this effect is stronger for future success, further demonstrating its potential as an early warning signal.

3.2.5 Key player identification

In the previous section we determined that betweenness centrality and topic engagement are the measures with the greatest vendor recall performance. That is to say, they are likely to have the most successful vendors among the top ranked users when ranked on these measures. Here we look at the top scoring users to investigate to what extent the top scoring users are indeed key players in the cryptomarket. To this end, we report the top 25 users, their member title, and their current and future sales for September 2014 for these measures in Tables 3.2 and 3.3.

We see that among the top 25 users in betweenness centrality and topic engagement

Table 3.2: Top 25 users for betweenness for September 2014. Titles are determined as their most significant observed over the entire dataset (Administrator > Market Moderator > Moderator > Public Relations > Vendor > Banned > Troll > Member).

Rank	username	title	#sales	
			current	future
1	penissmith	Troll	45	5
2	wefinance	Banned	44	2
3	themostseekrit	Moderator	0	0
4	FRIM	Vendor	703	381
5	scrufffe	Member	0	0
6	Yasuo	Vendor	2420	1790
7	Scattermind	Public Relations	400	0
8	LudoTilMortem	Market Moderator	274	0
9	Kimble	Administrator	0	76
10	leon-trotsky	Troll	0	0
11	ScoobyJew	Moderator	0	0
12	elmachico777	Vendor	49	240
13	Cypher	Vendor	123	3
14	evilsmile	Banned	0	0
15	Trippy	Moderator	8	464
16	Grandeur	Vendor	1735	622
17	nerotic	Member	0	0
18	sinordos	Member	0	0
19	d33poutside	Administrator	43	2
20	moka	Vendor	0	0
21	johnjones	Member	0	0
22	misterbitcoin	Vendor	0	0
23	Gold	Vendor	21	216
24	Sportlife	Vendor	649	268
25	maaadcity	Member	0	0

there are ten (i.e., 40%) that occur in both rankings. Furthermore, we observe that for both measures over half of the top 25 users have current and/or future sales (56% and 64% respectively). The probabilities of this occurring for any random ranking of all users are more than a million times smaller (3.47×10^{-7} and 3.44×10^{-9} respectively). Note that not all users with sales also have the corresponding “Vendor” member title. The reason for this is twofold: first, more important titles such as “Administrator” and “Moderator” supersede the “Vendor” title; and second, the “Vendor” title did not exist before September leading to some older vendors with few future sales not to be labeled as such. This also illustrates a potential pitfall of relying too much on forum member titles for key player identification.

Of the users with sales, twelve are among the top percentile for current sales and eight are among the top percentile for future sales. Respectively three (*kalashnikov*, *Yasuo*, and *Grandeur*) and one (*SkypeMan*) of them are in fact in the top 10 current and future sales. This suggests that these two measures are suitable for predicting potential successful vendors. Notably, *Trippy*, who is included in the top 25 for betweenness centrality, is the only user that is a member of the top percentile for future sales, but

Table 3.3: Top 25 users for topic engagement for September 2014. Titles are determined as their most significant observed over the entire dataset (Administrator > Market Moderator > Moderator > Public Relations > Vendor > Banned > Troll > Member).

Rank	username	title	#sales	
			current	future
1	Yasuo	Vendor	2420	1790
2	themostseekrit	Moderator	0	0
3	Grandeur	Vendor	1735	622
4	First	Banned	119	1
5	SingularLee	Member	0	0
6	penissmith	Troll	45	5
7	kalashnikov	Vendor	3029	2214
8	wefinance	Banned	44	2
9	FRIM	Vendor	703	381
10	moka	Vendor	0	0
11	JoeBloggs	Member	17	0
12	highasakite	Vendor	158	3
13	ucard	Vendor	910	177
14	mountainhigh9	Vendor	0	0
15	Scattermind	Public Relations	400	0
16	misterbitcoin	Vendor	0	0
17	kesh	Vendor	16	4
18	alphawolf89	Vendor	378	94
19	SkypeMan	Vendor	585	5078
20	fbgduck55	Troll	0	0
21	DonaldTrump	Member	0	0
22	IronHeart	Vendor	144	148
23	fnufnu	Member	0	0
24	Verto	Administrator	1163	338
25	ScoobyJew	Moderator	0	0

not a member of the top percentile for current sales. Note that *Trippy*'s member title in September was still "Vendor". Additionally, betweenness centrality appears to include a greater proportion of vendors for whom the majority of their sales are yet to come. On the other hand, we observe that the inclusion of *kalashnikov* and *SkypeMan* for topic engagement means that it captures a substantially greater total of future sales among the top 25 users. If our goal were to identify the absolute top vendors specifically, these results may be interpreted to imply that topic engagement is the better choice of measure. However, recall that sales volume does not equate to trade volume, but is merely a proxy of it. After all, the trade volume associated with a single sale can differ between listings and we are not able to differentiate between which sales came from which listings. Therefore, 100 sales could represent a larger total trade volume than 1000 sales. Consequently, we can not conclusively say whether the inclusion of *SkypeMan* by topic engagement is indeed the better choice compared to the inclusion of *Trippy* by betweenness centrality. This uncertainty is another reason why we put a greater emphasis on vendor recall than sales recall in this research, and why we focus on the top vendor percentile instead of the pure top vendors in terms of

sales. Regardless, the results in Table 3.2 and 3.3 are a concrete example of how these measures can potentially serve as early warning signals for future vendor success.

In addition to vendors, we also find users with other important positions on the forum, such as “Administrator” and “Moderator”, among the top 25 for both measures. In fact, betweenness centrality and topic engagement combined include three out of four users to have held the title “Administrator” among their top users. Furthermore, the only missing administrator became inactive within a month of the founding of the cryptomarket. Thus, we can say that all active administrators were found. Additionally, betweenness centrality identifies five out of nine users to have held the title of “Moderator” and who registered before the end of September 2014 (four out of seven if we exclude users who obtained the title after September, including *d33poutside*). The probability of this happening randomly is more than 250 million times smaller (2.07×10^{-11} ; and 2.27×10^{-9} if we exclude users who obtained the “Moderator title after September). On the other hand, topic engagement includes two out of nine (two out of seven) with probabilities of this randomly occurring that are just over 700 times smaller (3.10×10^{-4} ; and 1.82×10^{-4}). Thus, these measures are suited to predicting key players beyond just successful vendors. Though neither measure perfectly identifies only key players, they provide an excellent way of identifying individuals to investigate further manually.

3.3 Discussion and conclusion

The identification of key players in cryptomarkets, such as successful vendors and administrators, is a vital step in law enforcement interventions. Whereas it can be easy to identify administrators due to titles given to these users, it may be harder to identify successful vendors. It is especially difficult to identify those vendors whose success is yet to come. These tasks might be further complicated when encryption is used for message contents. The results presented in this chapter show that network measures computed on the user-to-user communication network and three forum activity indicators, not reliant on knowledge of message content, are useful in predicting (future) successful vendors. Specifically, the topic engagement indicator and betweenness centrality show the best performance.

Our results show that, on average, it is possible to distinguish between vendors and non-vendors using both network centrality and the activity indicators. Additionally, we found that more successful vendors have on average higher centralities and activity indicators than less successful vendors. This holds for both current and (to a slightly lesser extent) future success. However, it is important to remember that these findings are about the average case; perfect delineations between successful vendors, less

successful vendors, and non-vendors cannot be made. Even so, they indicate that the rankings induced by the measures have predictive potential for vendor success and may be useful to law enforcement activities.

To reduce the workload for law enforcement agencies, it can be beneficial to reduce the set of users that need to be manually investigated. We found that the measures of betweenness centrality and topic engagement included the greatest proportion of successful vendors when applying such a reduction (up to two thirds of the successful vendors when reducing to 20% of the users). Additionally, results showed that the vast majority (up to 98%) of post activity of the most successful vendors was produced by those included and that they were the relatively more successful vendors. Thus, most successful vendors that are not retained by these measures are simply not very active on the forum. We note that the network centrality measures appear to require the communication network to have sufficiently developed and stabilised for good predictive performance. We found that betweenness centrality was the only network measure that was able to detect a substantial set of successful vendors that were not found by any of the activity indicators. Thus, there are vendors that may not be the most active, start the most topics or get the most engagement on their topics, but that are able to establish themselves in the structure of the communication network such that they lie on many shortest paths. High betweenness vendors may, for example, be connecting buyers of distant locations and/or diverse goods. However, the question of why (certain) vendors achieve high betweenness scores, remains an open question to be addressed in future work through methods such as topic modelling. In short, while topic engagement shows the best overall performance, betweenness centrality could provide the greatest added value to law enforcement activities for reducing the set of users to investigate.

The results highlight that the same measures are almost as effective at recognizing those that will do well in the future. This can partly be explained by those vendors that are already quite successful and will simply continue to do well. However, results indicate that the top ranked users by betweenness and topic engagement in fact include several vendors whose majority of sales are yet to come. Additionally, for vendors that are moderately active, betweenness centrality was shown to be more effective at finding vendors with high future sales. Furthermore, results suggest that high betweenness centrality may (often) precede sales success. Thus, beyond predicting current success, the proposed approach can provide early warning signals for future success. However, how early we may be able to predict future success remains an open question for future work.

Finally, we highlight some possible limitations of this work. First, this study focused on a single, somewhat older, dark web cryptomarket. Therefore, the extent to which our findings can be generalized to other cryptomarkets or dark web marketplaces, is

an open question, as other markets may show unique characteristics not represented in this study. Regardless, it is worth noting that most fora and marketplaces appear to be operated in a similar fashion, i.e., with the fora being used to advertise and discuss vendors and their listings. Moreover, betweenness centrality has been shown to perform well in similar criminal network related settings before. Therefore, we expect our findings may very well hold up for other cryptomarkets and dark web marketplaces. Second, although Section 3.5.1 demonstrates that our findings are generally robust with respect to variations in parameter choices during the communication network extraction, the performance of the network measures are sensitive to these parameter choices, implying that parameter tuning is likely needed depending on the considered data and precise setting. Since our findings would be applied in a setting where sales information is unknown, especially in the case of future sales, it would be infeasible to automatically search for optimal parameter values. Third, we note that, by virtue of not having access to hidden and missing data, this study focused only on the visible public communication on this cryptomarket. In Chapter 2 we estimated that roughly 8% of posts were on hidden parts of the forum or were otherwise missing from the scraped data. Additionally, any off-market private communication is not included in the analysis. These missing links may impact the extent to which key players can be identified through user-to-user communication networks. Finally, we note that our analysis is hardly exhaustive in terms of considered network measures. Although we did experiment with a wider selection of measures, none of which outperformed betweenness centrality, it is possible that another (specialized) network measure would provide better performance. Despite this, we note that the network measures reported on in this chapter cover a wide range of network interpretations relevant to the cryptomarket forum setting. Therefore, we believe that the results reported in this chapter are a good account of what can be achieved with network measures.

3.4 Methods

In this section we discuss our dataset in Section 3.4.1, followed by a description of how the communication networks were extracted in Section 3.4.2. Next, we discuss the rationale and computation of our activity indicators (Section 3.4.3) and the four network measures (Section 3.4.4) employed, in the context of finding key players in cryptomarkets. Finally in Section 3.4.5, we discuss how the metrics used in this chapter are computed.

3.4.1 Dataset

As previously discussed in the Section 3.2.1, we use the data presented in Chapter 2. This dataset consists of data on the forum and the market, as well as data that links forum users to market users, i.e., vendors. For the forum data, we rely almost exclusively on the post and user data, ignoring more general information about topics and fora. For the market data, we rely exclusively on the vendors data. This vendors data includes their sales statistics at specific moments in time. However, in most cases, these moments in time are not conveniently at the end of each month. Consequently, the *current sales* of a vendor at the end of a given month were estimated based on their average daily growth in the number of sales between the most recent sales information available before and after the change of month. For the months after the last available sales information, the final sales total is used. We note that, since sales statistics are only ever compared within the perspective of that same month, we can include as much future data as available instead of always looking ahead the same number of months. *Future sales* of a vendor were therefore determined as the difference between their current sales, for a given month, and the last available sales information.

Figure 3.7 shows the total and monthly post activity and number of active users and vendors. Here, *active* users and vendors are those with at least one post up to and including the given month, where for the monthly active users we require at least one post that month. Throughout Section 3.2 we relied on the total sets of active users and vendors for each month.

3.4.2 Network extraction

Along with the dataset, we also utilise the communication network extraction method proposed in Chapter 2. This extraction method creates nodes for all active users and adds an edge connecting nodes for any posts by a pair of users that are in the same topic and adhere to certain parameters. The direction of these edges are from the user who placed the later post to the user who placed to earlier post. Additionally, edges are formed from every user who placed a post in a topic to the user who placed the first post in the topic. All edges are weighted to indicate the strength of the social tie implied by the edge. Here, an exponentially decaying weighting function is intended to model the probability of the edge representing a direct response.

As mentioned in the Section 3.2.1, we used the following parameters for network extraction: $\delta_\phi = 10$, $\delta_t = 1$ month, $\omega_{lower} = 0.2$, $t_{lim} = 7$ days, and $\omega_{first} = 0.5$. The first two parameters, i.e., $\delta_\phi = 10$, $\delta_t = 1$ month, set limitations on the existence of an edge. Specifically, they allow the formation of an edge for users of posts that are ten or fewer posts apart and were placed at most one month apart. These limitations may lead to some information loss with respect to connections to older posts. However, the

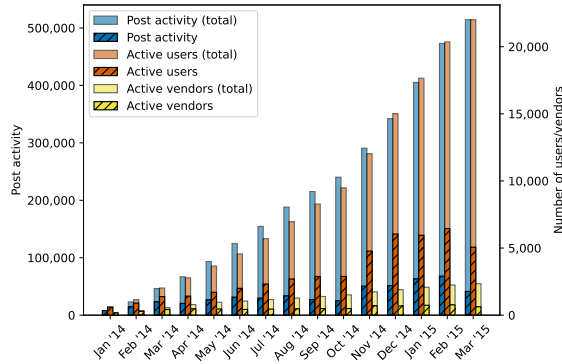


Figure 3.7: Post activity and active users over time.

likelihood of links to such older posts representing a meaningful connection is very low, while including them would unfairly favor the users posting in long (running) topics due to the sheer number of links they would receive. The parameters $\omega_{lower} = 0.2$, $t_{lim} = 7$ days, determine the scope and decay of the exponential weighting function applied to “regular” edges, i.e., they determine the strength of the implied social tie. Specifically, ω_{lower} sets the minimum weight at 0.2, while t_{lim} determines that this minimum weight applies for all pairs of posts at least seven days apart. The resulting exponential weighting function is shown in Figure 2.4 in Chapter 2. Thus, ω_{lower} and t_{lim} determine the likelihood that a post was placed in response to or after having at least seen a specific earlier post, while δ_ϕ and δ_t determine at what point we consider this likelihood too low to imply a social tie. The final parameter, $\omega_{first} = 0.5$, sets the weight for all other edges, i.e., edges formed from linking posts to the initial post. Robustness of our results for these parameters is investigated in Section 3.2.1.

Monthly communication networks were extracted based on all posts up to the end of the given month, thus including posts from previous months. Additionally, we simplify the networks by merging all parallel edges, i.e., all edges connecting the same two nodes in the same direction, into single edges. The weights of the resulting edges are exactly the sum of the parallel edges that were merged. In other words, the resulting weights represent the combined likelihood of a meaningful social tie connecting two users. As a result, we obtain 15 simplified monthly weighted directed networks $G = (V, E)$, where each node $u \in V$ represents an active user and each weighted edge $(u, v) \in E$ represents the inferred weight of the social tie from user $u \in V$ to user $v \in V$. It is on these monthly weighted directed networks that the network measures were computed.

3.4.3 Activity indicators

To evaluate the performance of predicting vendor success using network measures, we compare against activity indicators that can be directly computed from the forum data as a baseline. Similar to the rationale for our use of network measures, these activity indicators must also adhere to the requirement that we lack knowledge of message content. We considered three activity indicators in this chapter: post activity, topics started and topic engagement. Each of these indicators can be computed with little computational cost and independent of any knowledge of message content. Below, we discuss why we believe these are appropriate indicators and how they are computed.

3.4.3.1 Post activity

Post activity refers to the number of posts a user has posted on the forum up to a given moment in time. A straightforward link can be made between a user's visibility on a forum and their post activity: the more often someone posts, the more likely it is that another user will come across one of them. This increased visibility leads to greater name recognition, which has been linked to improved trust [42] and market outcomes [75] (e.g., more sales); and trustworthiness has been shown to be a better predictor of vendor selection than product diversity or affordability [60]. Therefore, post activity can be used as an indicator of the likelihood of vendor success.

3.4.3.2 Topics started

Forums that accompany cryptomarkets are intended to allow vendors and their customers to interact. It is common practice for vendors to promote products listed for sale by starting a topic promoting their listings [95]. The number of topics a user has started is therefore a potential indicator of being a vendor. As a greater number of topics started may lead to greater visibility, greater name recognition, and simply a greater reach, it may also lead to increased success for vendors [64].

3.4.3.3 Topic engagement

Topic engagement is the total number of responses to all topics started by a user. Topic engagement combines the fact that starting topics is a good indicator of being a vendor with the fact that when topics receive considerable engagement they are naturally also more visible. Additionally, engagement in any topics about a specific listing is likely to be associated to that listing or the vendor. For example, a post may concern feedback on the particular listing or on the vendor themselves. Either way, engagement on these topics is also highly probable to be associated with actual sales. Thus, where the topics

started baseline is more likely to be a good indicator of being a vendor or not, topic engagement is more likely to be a good indicator of the success of any such vendor.

3.4.4 Network centrality measures

In this subsection we discuss the various network measures utilised in this chapter. We discuss their computation and interpret their meaning within the context of cryptomarket communication networks. Recall from Section 3.1 that a link can represent two types of phenomena, either they represent a certain familiarity or shared interest, i.e., a passive relationship, or they represent the more active relationship of responding to one another, i.e., communication. The weighting and direction of edges in our communication networks are closely associated with the latter active relationship. After all, the direction of a link implies who posted first and who responded, while the weighting indicates how quickly the response was made. Therefore, the choice of measure variant is closely related with how we choose to interpret the links in the network.

All network measures were computed using the *igraph* package [51].

3.4.4.1 Degree

The degree of a node is a measure of the number of distinct neighbors connected to that node. While degree captures this regardless of edge direction, in- and out-degree count only the neighbors connected through incoming and outgoing edges, respectively. Furthermore, the weighted degree variants sum the weights of the connections with the neighbors.

The degree can be interpreted as the number of different users that a (potential) vendor responds to or receives responses from. The weighted variant also takes into account how strong the relation to these users is. Thus, a high in-degree in our networks indicates many different users responding within a relatively short time frame or responding in a topic they started. Since it is likely that those that respond shortly after a post have seen that post, a high in-degree implies visibility to many different users, thereby improving the aforementioned brand awareness. As brand awareness promotes trust and sales [42, 75] and trust is a good predictor of vendor selection [60], a high in-degree might serve as a good predictor of vendor success.

Unlike incoming edges used for in-degree, outgoing edges do not imply visibility of the user to the neighbors these edges connect to, since those neighbors posted before the user. Furthermore, in-degree combines the visibility of responses with the visibility of starting topics. Thus, the in-degree is similar to the activity indicators, but focused on the number of individuals engaging with the user rather than the volume of engagement. For these reasons we focus on the in-degree. We report results for

the unweighted in-degree, as we believe the number of neighbors, i.e., the number of potential customers, to be a better predictor of vendor success than the combined strength of the social ties to these neighbors. Note that the weighted in-degree showed similar results, but with slightly fewer detected vendors that were not found by the activity indicators.

3.4.4.2 Harmonic closeness centrality

Closeness centrality [124] is a measure of how easily a node can reach every other node in the network. Essentially, it determines whether a node is central based on its distances, i.e., shortest path lengths, to all other nodes. In other words, where degree was a measure of how well someone is connected locally, closeness is a measure of how well connected a node is globally, i.e., to the entire network. Harmonic closeness centrality behaves essentially the same as standard closeness centrality and extends properly to directed and disconnected networks, i.e., networks with node pairs that are not connected by any (directed) path [29].

Let $d_G(u, v)$ be the shortest distance connecting nodes $u, v \in V$, where if no path exists $d_G(u, v) = \infty$. Interpreting $\frac{1}{\infty}$ as 0, we can define the harmonic closeness centrality ($hcc_G(u)$) of node u as:

$$hcc_G(u) = \sum_{v \in V, u \neq v} \frac{1}{d_G(u, v)}. \quad (3.1)$$

For bidirectional harmonic closeness centrality, the shortest paths can be determined following edges regardless of their direction. However, for incoming and outgoing harmonic closeness centrality the paths may follow edges only in one direction, either following the direction of the edges (outgoing) or going against the direction of the edges (incoming). The weighted variants of these measures use the inverse of the edge weights during shortest distance computation, so that stronger connections equate to shorter distances.

The interpretation of distance for nodes that are more than a single edge apart, i.e., that are connected by a path of at least length two, with respect to vendor success in cryptomarket communication networks is not straightforward. After all, other than having a shared author for one of the involved posts, the posts that were responsible for the formation of subsequent links in a path may be wholly unrelated to each other. For the weighted and directed variant, i.e., links representing responding to one another, the interpretation of paths (and their lengths) that start or end at a specific user are unclear as the path may not represent a single “conversation”. Perhaps at best, shorter paths would imply closer familiarity. Therefore, we are likely better served by relying on the link representation of familiarity directly, i.e., using the undirected and

unweighted variant. Because we also detected the largest share of vendors not found by any of the activity indicators with the unweighted bidirectional variant of harmonic closeness centrality, we report on this variant in the Section 3.2.

For the interpretation of a link as familiarity and shared interest, one can interpret a smaller distance as it being more likely for a user's posts to be visible to other users. This interpretation would not be dissimilar to that of the "friendship" relation in a social media network such as Facebook. Even so, it is unknown how the topics that are responsible for forming the edges that make up the connecting paths are related. They may originate from the same or a highly similar topic, increasing the odds of being visible, or they may differ greatly, making it unlikely that these connections truly form a meaningful path. Thus, a high closeness centrality does not intuitively imply a successful vendor. Regardless, closeness centrality has often proven to capture users at important positions in a network [66, 41]; its incorporation of global network information in a substantially different manner than betweenness, convinced us that it should be included in our analyses.

3.4.4.3 Betweenness centrality

Betweenness centrality [68, 33] measures the extent to which a node is on shortest paths connecting pairs of nodes in the network. In other words, it measures how important a node is with respect to connecting various communities in the network. In the context of cryptomarkets, this makes it a good measure of how well a (potential) vendor reaches different communities of potential buyers. A vendor with a high betweenness is more likely to have a larger pool of buyers as they may be able to draw from more relatively distinct communities of buyers. Additionally, betweenness centrality has been shown to perform well in identifying key players in criminal networks [97, 39].

The betweenness centrality of node $u \in V$ is determined by computing the sum of the fraction of shortest paths connecting nodes $v, w \in V$ that pass through u . Let σ_{vw} indicate the number of shortest paths connecting nodes $v, w \in V$, and let σ_{vuw} indicate the number of those shortest paths that pass through node $u \in V$. Then, interpreting $\frac{x}{0}$ as 0, the betweenness centrality $bc(u)$ of node u can be defined as:

$$bc(u) = \sum_{v, w \in V, u \neq v \neq w \neq u} \frac{\sigma_{vuw}}{\sigma_{vw}} \quad (3.2)$$

For directed betweenness centrality, paths must follow the direction of the edges, while undirected betweenness can follow edges in either direction. Like for harmonic closeness centrality, the weighted variants use the inverse of the edge weights during shortest path computation, so that stronger connections equate to shorter distances. As direction and weighting can have a large impact on the probability of lying on a

shortest path, we use the directed weighted betweenness in this study. We note that where the interpretation of paths was unclear for harmonic closeness centrality, this is less of an issue for betweenness centrality. After all, for betweenness centrality we rely on the fact that the user exists on many shortest paths, and is less reliant on its specific position in the path. When viewed on the scale of the whole network, occurring on many shortest paths implies being a central figure in the various conversations on the forum, regardless of whether the individual paths relate to the same conversation. For a vendor, this in turn might imply that they or their products are often discussed and that the vendor actively engages with their clientele, all of which may promote their sales. Thus, a high directed weighted betweenness centrality can be interpreted here as a good indicator of vendor success. We note that preliminary results showed this variant to also have the best performance.

3.4.4.4 PageRank

The final measure we consider is PageRank [114]. PageRank computes the probability that a random walker, that follows one of the available neighboring edges or jumps to a random node with a particular probability, ends up at a given node. For the directed variant the choice of edge is restricted to following the direction of the edges and adding weights impacts the odds of following any given edge. Similar as for betweenness centrality, we report the results for the variant taking both direction and weighting into account as this provides the random walker with more context. Note that we found this variant of PageRank to indeed have the best performance.

High PageRank values often follow from having paths (or edges) incoming from (many) other important (i.e., high value) nodes in the network. Therefore, we can interpret a high PageRank value as being closely connected to other key players. As previously stated, Duxbury & Haynie [60] found that buyers were more likely to continue ordering with vendors within the same community. This means that the close connection between users with high PageRank value can be indicative of a boost in their perceived trust and may stimulate their sales. Thus, a high PageRank value may be able to predict successful vendors.

3.4.5 Evaluation metrics

In this subsection we discuss the normalization and various evaluation metrics employed in this work.

3.4.5.1 Relative/absolute difference score

Let $s_{u,t,i}$ indicate the activity indicator or network measure value for user/node u , month t , and measure i . We first apply min-max normalization to these values, i.e.,

$$sn_{u,t,i} = \frac{s_{u,t,i} - \min_{v \in V} s_{v,t,i}}{\max_{v \in V} s_{v,t,i} - \min_{v \in V} s_{v,t,i}}. \quad (3.3)$$

Let $V_a, V_b \subseteq V$ indicate two groups of users, for example, the group of all vendors and all non-vendors for a given month. The *absolute* and *relative difference score* for a given month t and measure i can then be computed as:

$$abs_diff_score(a, b, t, i) = \frac{\sum_{u \in V_a} sn_{u,t,i}}{|V_a|} - \frac{\sum_{u \in V_b} sn_{u,t,i}}{|V_b|}; \quad (3.4)$$

$$rel_diff_score(a, b, t, i) = \frac{abs_diff_score(a, b, t, i)}{\sum_{u \in V_b} sn_{u,t,i}/|V_b|}. \quad (3.5)$$

3.4.5.2 Recall metrics

Let $TV_t \subseteq V$ indicate the top vendor percentile for a given month t and let $TU_{t,i} \subseteq V$ indicate the top 20% users based on measure i for a given month t . Then the *vendor recall* is computed as

$$vendor_recall(t, i) = \frac{|TV_t \cap TU_{t,i}|}{|TU_{t,i}|} \times 100\%. \quad (3.6)$$

The monthly overlap between detected vendors (for which the mean and standard deviation are presented in Table 3.1) for given measures i, j and month t is computed as follows:

$$overlap_{t,i,j} = \frac{|TV_t \cap TU_{t,i} \cap TU_{t,j}|}{|TV_t \cap TU_{t,i}|} \times 100\%. \quad (3.7)$$

Thus, $overlap_{t,i,j}$ computes the percentage of (top vendor percentile) vendors detected by measure i that were also found by measure j . Note that in general $overlap_{t,i,j} \neq overlap_{t,j,i}$.

Let $pa_{u,t}$ indicate the post activity of user/node u up to and including month t ; and let $sales_{u,t}$ indicate their sales. Then we can compute the *post activity recall* and *sales recall* as follows:

$$post_activity_recall(t, i) = \frac{\sum_{u \in TV_t \cap TU_{t,i}} pa_{u,t}}{\sum_{u \in TV_t} pa_{u,t}} \times 100\%; \quad (3.8)$$

$$sales_recall(t, i) = \frac{\sum_{u \in TV_t \cap TU_{t,i}} sales_{u,t}}{\sum_{u \in TV_t} sales_{u,t}} \times 100\%. \quad (3.9)$$

3.5 Supplementary results

In this section we present several supplementary results. First we investigate the robustness of our results from Section 3.2 in Section 3.5.1. Then, we look at what percentages of all post activity and sales, the top percentile of detected vendors are responsible for in Section 3.5.2. Next, in Section 3.5.3 we look at the relationship between sales and post activity beyond September 2014, to verify if our results from Section 3.2 are robust for other months. Finally, in Section 3.5.4 we consider whether the results presented in Section 3.2 hold up for different choices of detection thresholds.

3.5.1 Robustness of results

The results presented in Section 3.2 rely on the network extracted based on a single set of parameter values. Here, we explore the robustness of those results for different parameter values (see Section 3.4 for a description of these parameters). We do so by changing the value of a single parameter at a time while maintaining the same values for all other parameters. Figures 3.8 and 3.9 shows vendor recall trends (similar to main paper Figure 3.4a) for each set of parameter values. Each row of plots in these figures shows the results for changing one specific parameter, with the middle column corresponding to the “default” values. Note that since we are only changing the network formation process, the vendor recall results for the forum activity indicators do not change. Thus, they provide a visual aid in analysing the changes in network measure performance.

Figure 3.8 shows vendor recall trends for varying values for the parameters influencing the formation of edges. First, δ_ϕ determines how many posts two posts may be apart at most, to still form an edge connecting the users who placed them. Figures 3.8a–e show that the in-degree performs better at lower values of δ_ϕ . Similarly, bidirectional harmonic closeness centrality performs slightly better at lower values. However, even then neither comes close to outperforming betweenness centrality. Some variation in performance can be observed for betweenness centrality as well. Slightly better performance than default is observed for $\delta_\phi = 5$ and 20. As such, there is no indication that specifically using a smaller or larger value of δ_ϕ would improve performance. Second, δ_t determines how much time between two posts may have elapsed at most, to still form an edge connecting the users who placed them. Figures 3.8f–j show that changing δ_t between 7 days and 3 months hardly affects the vendor recall at all. Thus, our findings in the main paper can be considered robust for both δ_ϕ and δ_t .

Figure 3.9 shows vendor recall trends for varying values for the parameters influencing the weight of edges. First, ω_{lower} and t_{lim} determine the scope and rate of decay of the exponential weighting function applied to “regular” edges, i.e., the im-

plied social ties. Specifically, ω_{lower} sets a minimum weight, and t_{lim} determines after how much time (between the placement of posts) this minimum weight is reached. For both of these parameters we can see no meaningful change in the vendor recall performance for either lower or higher values (see Figures 3.9a–j). Therefore, our findings in Section 3.2 can be considered robust for both ω_{lower} and t_{lim} . Second, ω_{first} sets the standard weight for all edges formed by linking to the initial post of topics. Figures 3.9k–o show slight improvement of vendor recall for weighted directed PageRank as this parameter gets closer to one, the maximum weight. However, it remains the case that at no point does it outperform betweenness centrality. We also observe a slight variation in the performance of betweenness centrality. Specifically, we see movement in when it performs well. It appears that betweenness centrality performs slightly better in the early months for lower values of ω_{first} , whereas it performs slightly better in the later months for higher values. Consequently, it may be advantageous to choose the ω_{first} value based on the age of the cryptomarket one is dealing with. Regardless, the change in performance is small and we can conclude that our findings in the main paper can also be considered robust for ω_{first} .

Having established the robustness of changing a single parameter at a time, we finally consider changing several parameters simultaneously. For this we first combined parameter values which showed slight improvements compared to default parameter values in Figures 3.8 and 3.9 and plot the vendor recall trends for various combinations of these values in Figure 3.10a–j. Although some combinations show relatively improved performance for the centrality measures, none show sufficient improvement to become inconsistent with our findings.

Next, we investigate how potential loss of information caused by parameters δ_ϕ and δ_t , as described in Section 3.4.2, impacts vendor recall performance if chosen differently in Figures 3.10k–n (i.e., the bottom row of Figure 3.10). Specifically, we consider each variation of including fewer or more posts and allowing for less or more maximum time between posts. The figures show that the performance changes are largely dominated by the δ_ϕ parameter, showing the same performance changes as those we observed for the default δ_t . Specifically, we see a slight improvement in the performance of betweenness centrality at both a lower and higher value of δ_ϕ than default. Thus, there does not seem to be a clear connection between the performance and the information loss caused by smaller δ_ϕ values. On the contrary, Figures 3.10k,m show that when we decrease δ_t to 14 days, the performance drops slightly for all tested δ_ϕ parameter values. However, Figures 3.10l,n show that when we extend δ_t to 3 months, we observe the same performance as at 1 month for the same δ_ϕ parameter value. As this was also the case for the default value of δ_ϕ , we may conclude that although performance can be compromised by the information loss caused by δ_t , there is a point beyond which extending δ_t no longer improves performance. In other words,

there is a value of δ_t at which the performance stops improving. It appears that for the EVOLUTION cryptomarket this aligns with our chosen default value of 1 month. However, note that even at a δ_t as low as 7 days (Figure 3.8f), the performance drop observed is very small. Thus, it appears that the vast majority of the added-value to performance lies in the responses that are written within days of a post.

In short, we observe that in general the vendor recall performance varies little for changes in parameter values determining the formation of the network. Thus, we can consider our findings in Section 3.2 to be robust for these parameters.

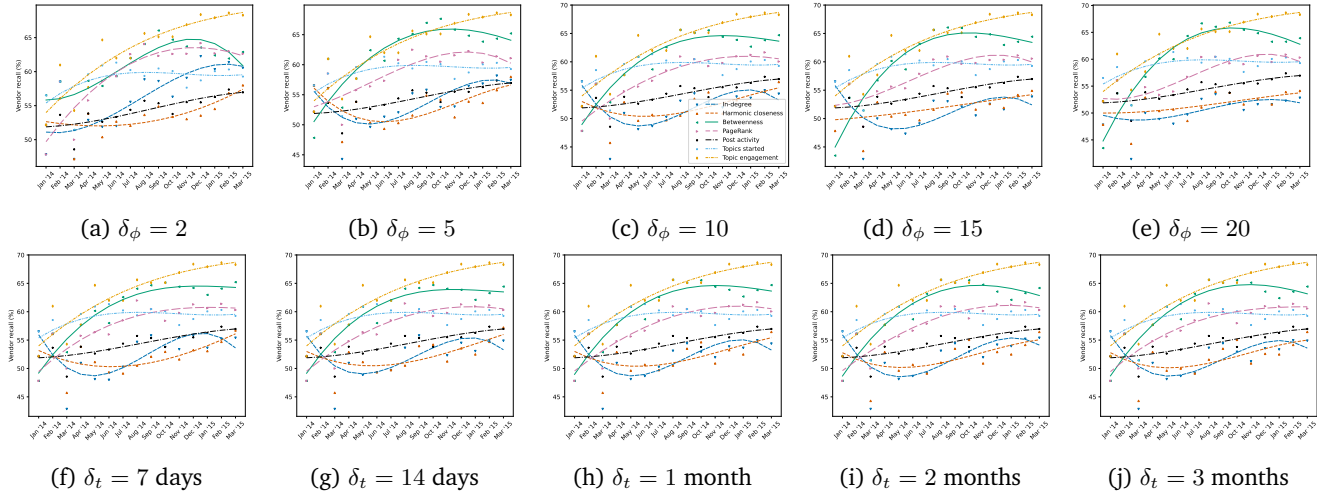


Figure 3.8: Results for varying edge formation parameters. Monthly vendor recall of top vendor percentile (top 0–20% vendors in terms of sales) among the top 20% of all users based on the network measures and activity indicators for current success. Each plot displays monthly vendor recall for a different set of parameter values used for generating the network. Each row of plots varies a single parameter influencing edge formation, with the remaining parameters at their default values. The center column always corresponds exactly with the default parameters.

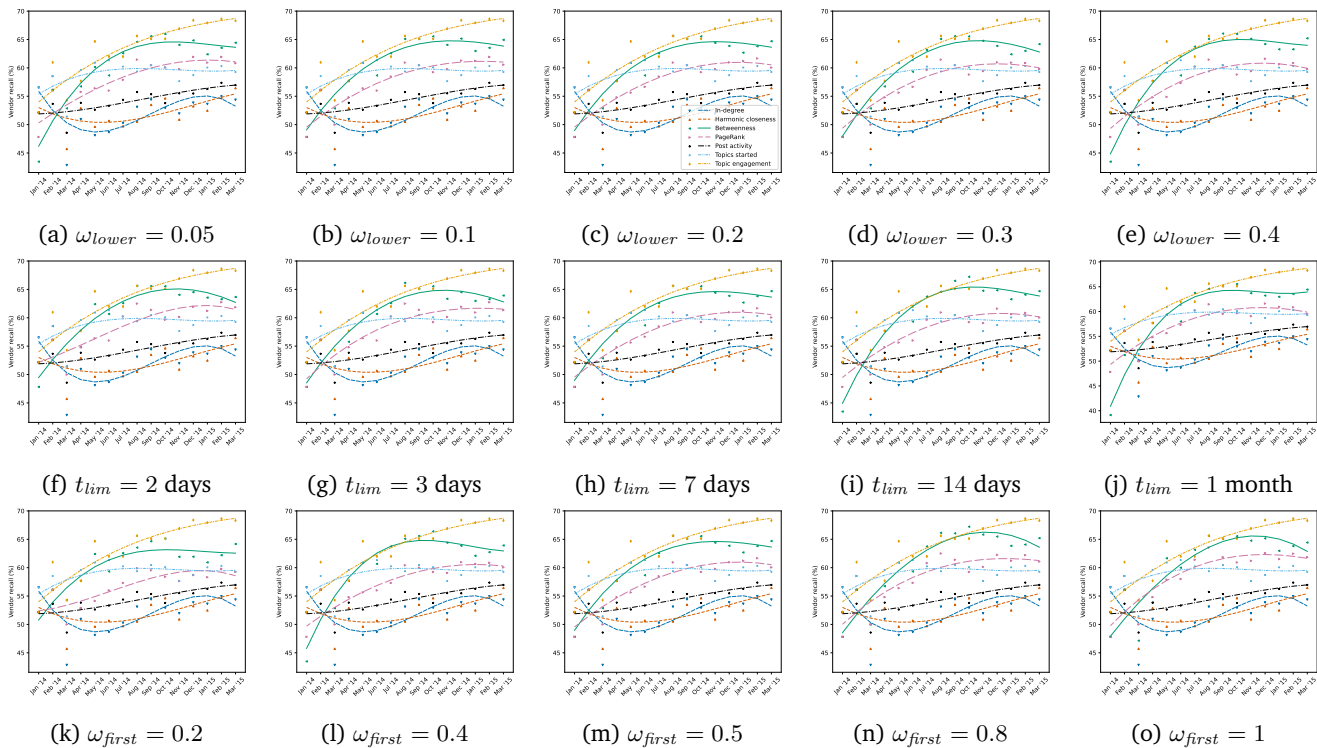


Figure 3.9: Results for varying edge weighting parameters. Monthly vendor recall of top vendor percentile (top 0–20% vendors in terms of sales) among the top 20% of all users based on the network measures and activity indicators for current success. Each plot displays monthly vendor recall for a different set of parameter values used for generating the network. Each row of plots varies a single parameter influencing edge weighting, with the remaining parameters at their default values. The center column always corresponds exactly with the default parameters.

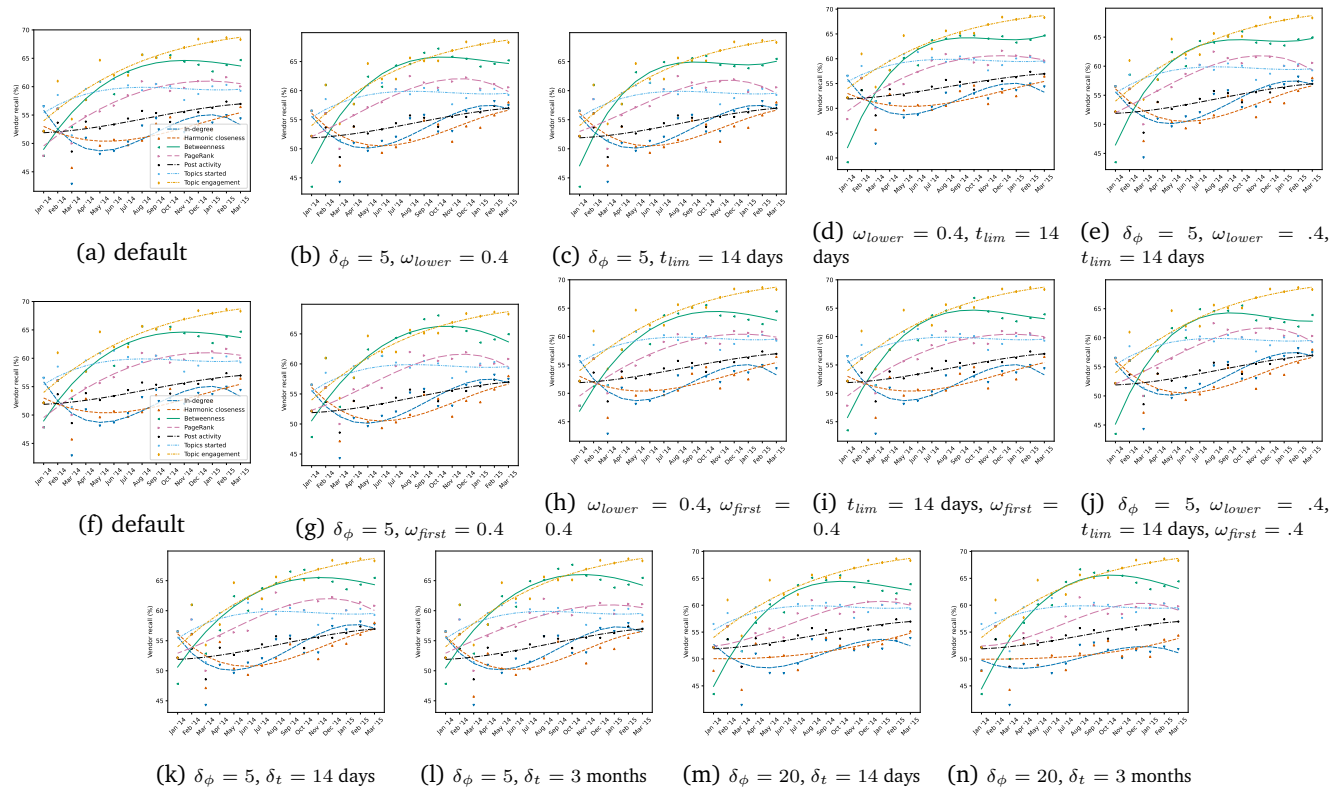


Figure 3.10: Monthly top vendor percentile recall results for varying multiple edge weighting parameters. For the top two rows, each set of varied parameters consists of those that showed some improvement in Figures 3.8 and 3.9. The bottom row specifically compares different variations for δ_ϕ and δ_t , to investigate the effects for different scopes of information loss.

3.5.2 Post activity and sales recall

In this section, we consider the *post activity recall* and *sales recall*, which measure respectively what percentage of post activity and sales of the entire top percentile the detected vendors are responsible for. We plot the monthly post activity recall in Figure 3.11 and current and future sales recall in Figure 3.12. These figures show the recall in terms of what percentage of post activity/sales the recalled vendors are responsible for.

From Figure 3.11 we see that the vast majority of post activity by the top vendor percentile is associated with vendors with high network centrality and activity indicators. Thus, many of the non-recalled vendors are likely to be those with very few posts. Not unexpectedly, the post activity indicator often has the highest post activity recall, while the topics started indicator often captures the least post activity.

Figure 3.12 shows that for current success most of our observations for vendor recall hold up. Perhaps the most significant change is that the differences between PageRank and topics started and between betweenness centrality and topic engagement are more prominent. Similarly, for future success PageRank now outperforms the topics started baseline more consistently. For both current and future success, we observe that the sales recall is generally between 10–20% higher than the corresponding vendor recall. This indicates that the detected vendors are, on average, the more successful vendors among the top percentile.

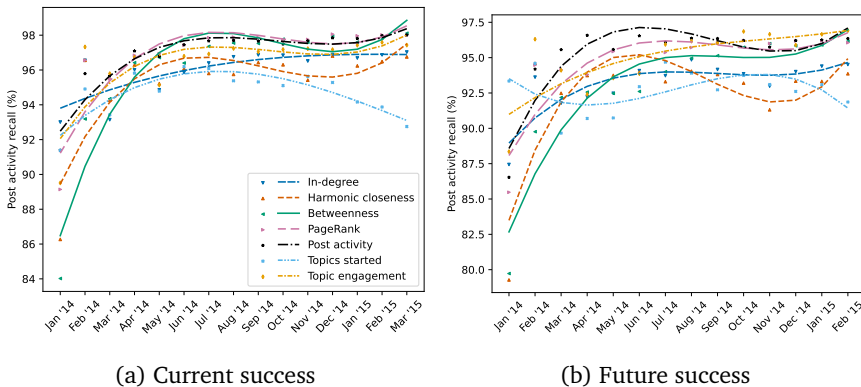


Figure 3.11: Monthly post activity recall for both current (a) and future (b) success. Higher post activity recall indicates that recalled vendors placed a relatively larger share of the top vendor percentile’s total post activity.

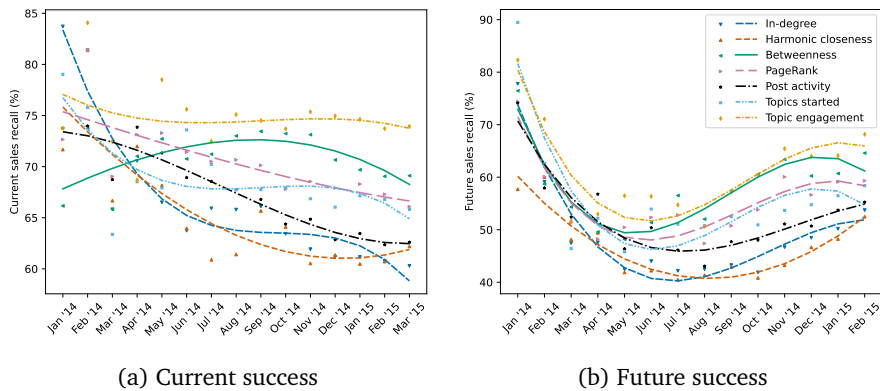


Figure 3.12: Monthly sales recall for both current (a) and future (b) success. Higher sales recall indicates a greater portion of the top vendor percentile’s total sales was attributed to the recalled vendors.

3.5.3 Sales and post activity of recalled and non-recalled users

Here, we explore the sales and post activity of recalled and non-recalled vendors for the various network centrality measures (see Figure 3.13) and months (see Figures 3.14, 3.15 and 3.16). We do so, to check if the conclusions drawn from the case study in Section 3.2, specifically with respect to Figures 3.5 and 3.6, are robust for the other months and if those conclusions also apply to the other centrality measures.

Figure 3.13 confirms that, for all centrality measures, there exists a threshold of activity (around 100 posts for September 2014) beyond which users are always included. Furthermore, we observe that in-degree, harmonic closeness centrality, and PageRank are far less inclined to include low activity vendors within the top 20% of their rankings. Although less obvious, for these centrality measures it does still appear to hold that greater success leads to greater odds of being included at lower activity. Whereas for betweenness centrality we found that the vendors uniquely found were relatively less active and more successful, we find that only the latter holds for the other measures. This lines up with our earlier finding from Table 3.1, that in-degree, harmonic closeness, and PageRank uniquely find very few vendors that are not likely to be found through their forum activity.

Figure 3.14 shows that, throughout all months, a threshold of activity exists beyond which users are always included. This threshold seems to slowly increase as the forum and the overall activity increases. However, after October 2014 this increase seems marginal at best. Furthermore, each month we observe the tendency for moderate activity vendors to be more likely to be included the more successful they are. We confirm this in Figure 3.15 by grouping vendors by their activity. Here, we have

excluded the first few months which we found in Section 3.2 to be a development period for the network, before achieving reliable vendor recall. We observe that indeed each month, though to a less significant extent for some, the average success of vendors uniquely found by betweenness exceeds that of those uniquely found by topic engagement. Furthermore, Figure 3.16 confirms our finding that this difference in average success is larger for future success. Thus, we can conclude that our findings in Section 3.2, with regards to Figures 3.5 and 3.6, are robust beyond only September 2014.

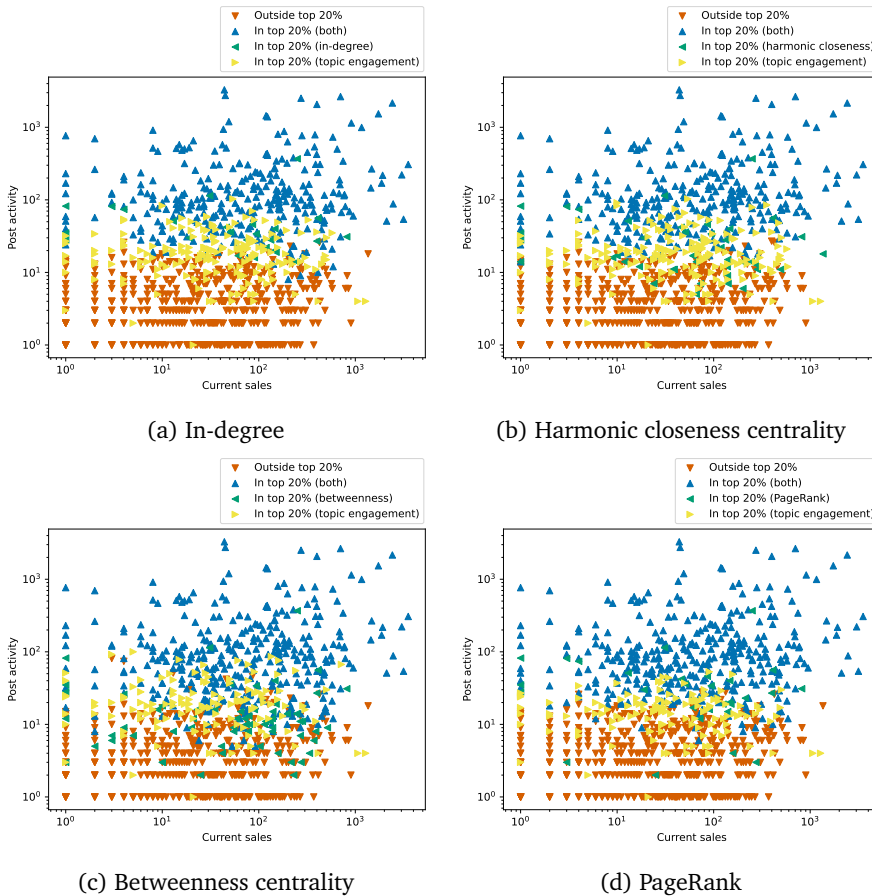


Figure 3.13: Sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for September 2014, for each of the network centrality measures compared with topic engagement and their intersection, respectively.

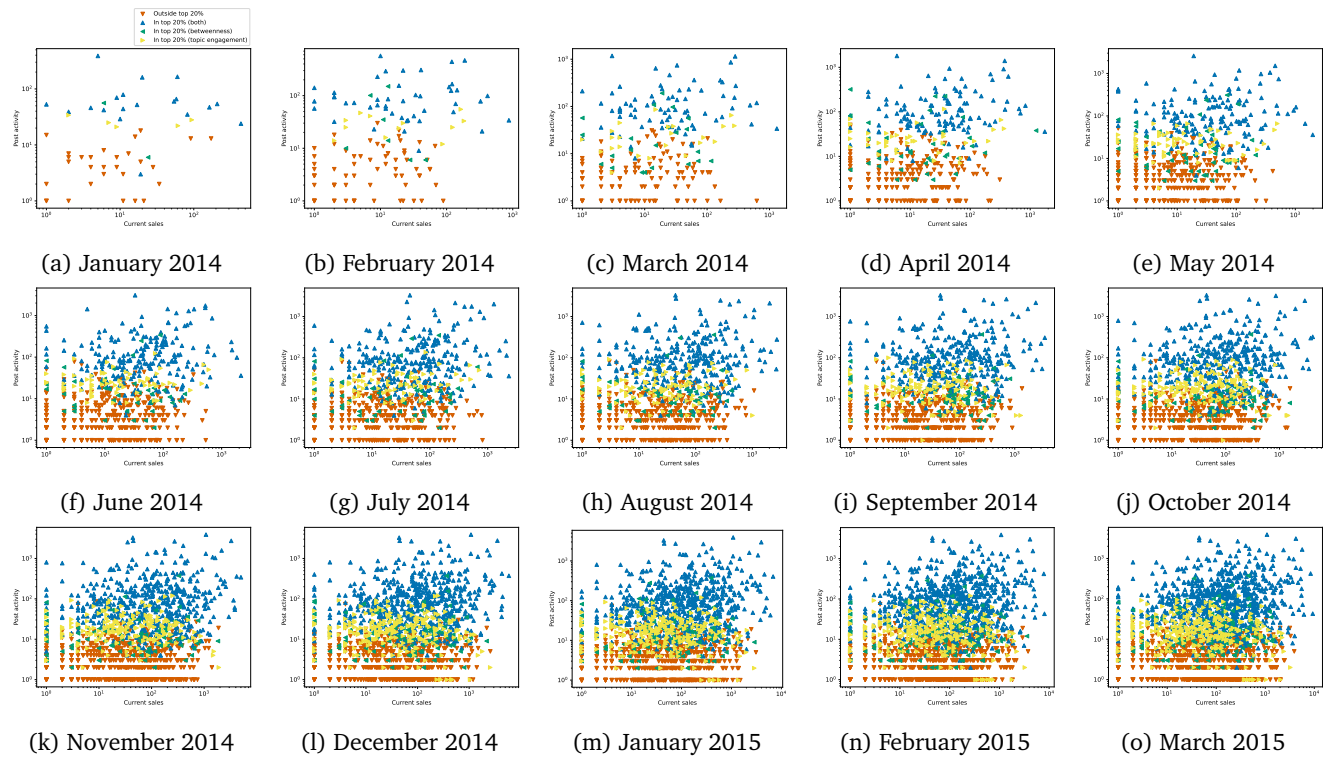


Figure 3.14: Current sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for topic engagement, betweenness centrality, and their intersection, for each month.

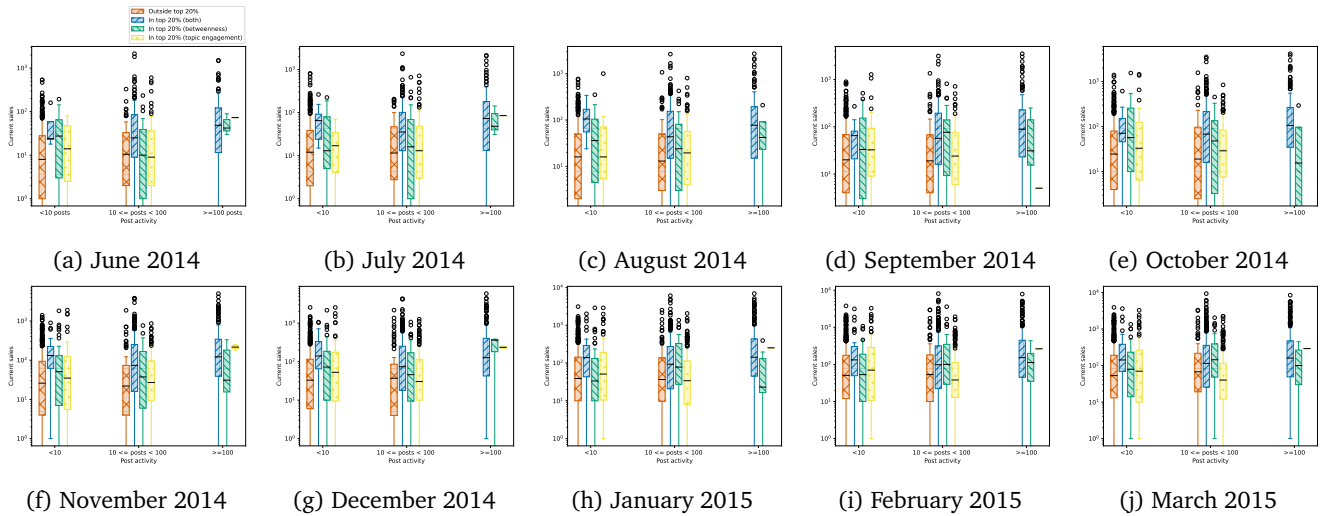


Figure 3.15: Current sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for topic engagement, betweenness centrality, and their intersection, for each month.

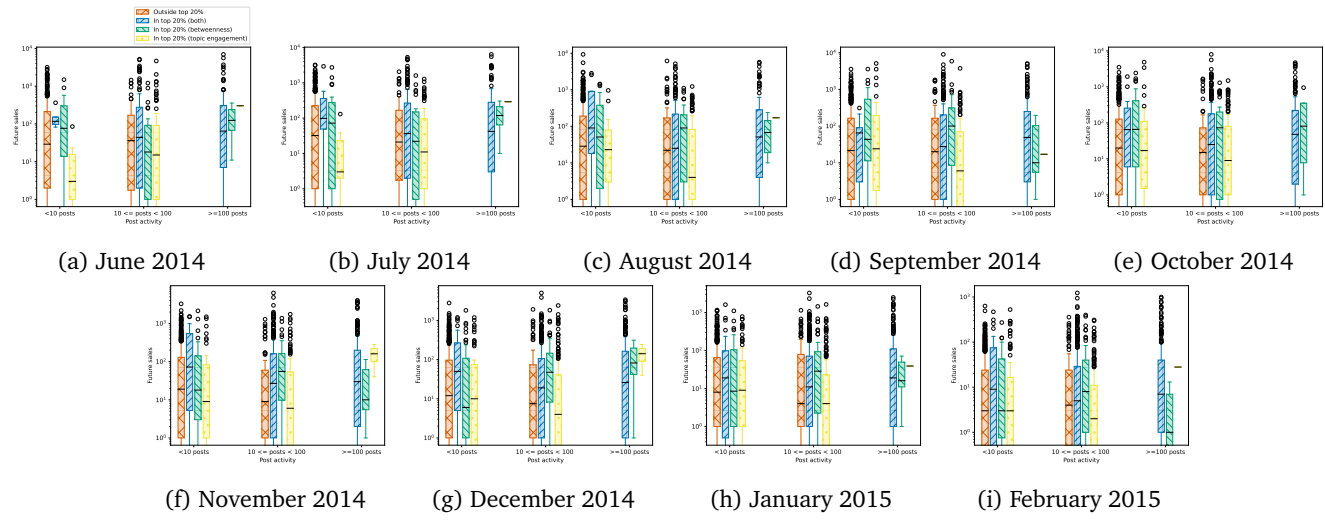


Figure 3.16: Current sales and post activity of recalled (in top 20%) and non-recalled (outside top 20%) users for topic engagement, betweenness centrality, and their intersection, for each month.

3.5.4 Performance at different thresholds

In Section 3.2, we presented evaluation metric results for a specific threshold for the measure rankings, namely 20% of all users. Here, we use ROC curves (receiver operating characteristic curves) to investigate the performance of the activity indicators and network measures at different thresholds. Specifically, we computed statistics at each interval of 5% in the range 0–100%, checking the performance for predicting groups of vendors for each of our measures. We consider the entire top vendor percentile, “less active” vendors among the top vendor percentile, all vendors, and all “less active” vendors. We set the activity threshold for “less active” at fewer than 100 posts, above which all vendors were shown to be found for all centrality measures. The resulting ROC curves, for September 2014, are shown in Figure 3.18.

For the top vendor percentile (Figures 3.18a,b), we see that topic engagement and betweenness centrality achieve a similar true positive rate up to a false positive rate of 20%. Afterwards, the topic engagement clearly outperforms all centrality measures. One factor contributing to the poor performance of betweenness centrality at false positive rates above 20% is that many users get the same lowest betweenness value of zero. Since identically scoring users are essentially randomly ordered, their ordering does not add any predictive power. In Figure 3.17 we show that for each monthly snapshot around 70% of users end up with the same lowest value for betweenness centrality, while these percentages are far lower for the remaining measures. Therefore, betweenness centrality is not suited for higher thresholds as such thresholds would include increasingly more users that are essentially randomly ordered. However, in practice, given the limited resources of law enforcement agencies, we are far more interested in the performance at low false positive rates. After all, the higher the false positive rate, the more resources would be wasted on non-vendors. Figures 3.18a,b confirm that, at least for any lower thresholds, our findings with regards to vendor recall hold up.

When we consider all vendors, not just the top vendor percentile, Figures 3.18c,d paint a different picture. Again, topic engagement provides the best overall performance, but now closely followed by the topics started indicator. On the contrary, though still outperforming the remaining measures at low thresholds, betweenness centrality has clearly worse performance. However, this is in line with our conclusions drawn from the results in Section 3.2.4, that topic engagement is the best overall predictor and that betweenness centrality performs particularly well for the more successful vendors. The performance of the topics started indicator for all vendors compared to for successful vendors, indicates that starting topics on the EVOLUTION forum was quite indicative of being a vendor in general.

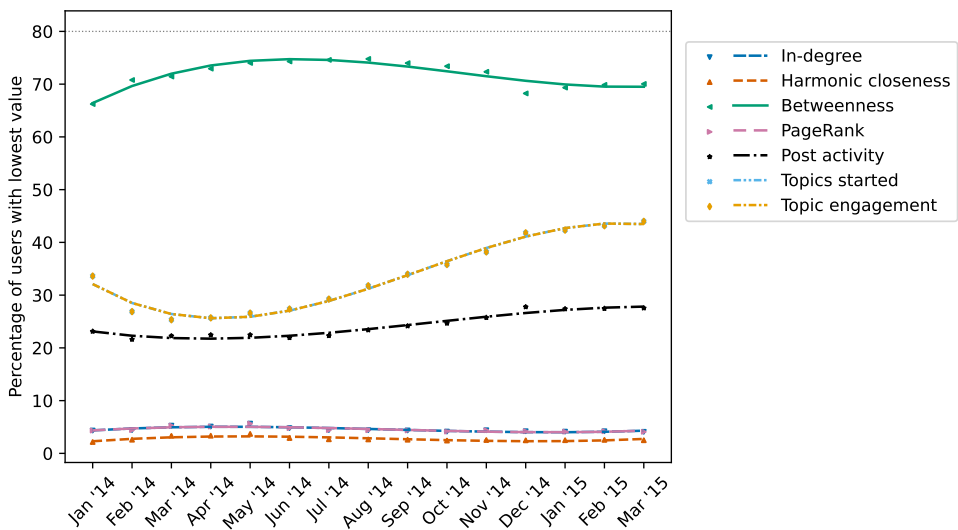
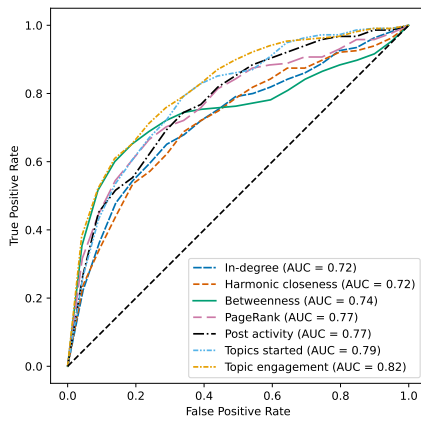
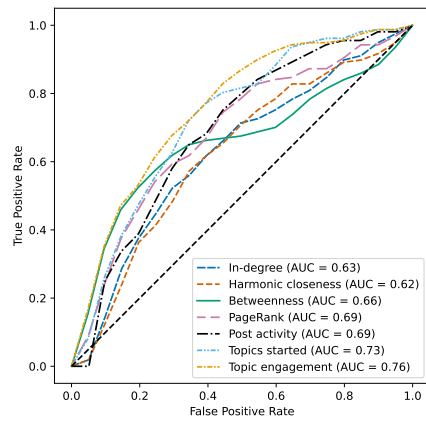


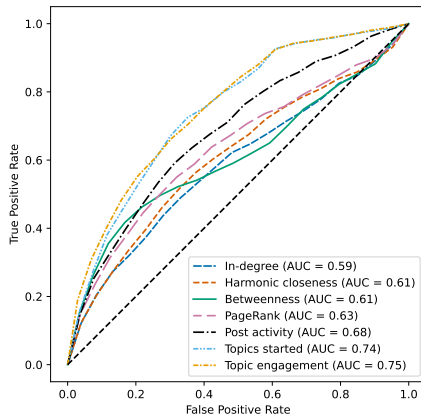
Figure 3.17: Percentage of users with the lowest value. Identical values indicate these users may have any random ordering, thus providing no additional predictive power. The grey dotted line at 80% indicates the threshold that would need to be exceeded for these random orderings to impact our standard threshold for the measure rankings of 20% of all users. None of the measures exceed this threshold for any of the snapshots.



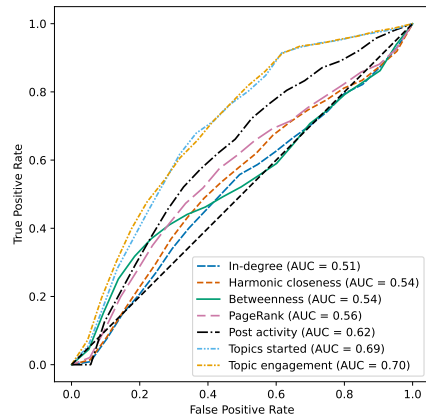
(a) Top percentile vendors regardless of activity.



(b) Top percentile vendors with fewer than 100 posts.



(c) All vendors regardless of activity.



(d) All vendors with fewer than 100 posts.

Figure 3.18: ROC curves for September 2014, predicting groups of vendors.