Don't trust, verify: The economics of scams in initial coin offerings*

Kenny Phua Bo Sang Chishen Wei Gloria Yang Yu

Internet Appendix: https://bit.ly/3QY6tt8

Abstract

Losses from fraud and financial scams are estimated to exceed U.S. $5 trillion annually. To study the economics of financial scams, we investigate the market for initial coin offerings (ICOs) using point-in-time data snapshots of 5,935 ICOs. Our evidence indicates that ICO issuers strategically screen for naïve investors by misrepresenting the characteristics of their offerings across listing websites. Misrepresented ICOs have higher scam risk, and misrepresentations are unlikely to reflect unintentional mistakes. Using on-chain analysis of Ethereum wallets, we find that investors in misrepresented ICOs tend to be less sophisticated. Overall, our findings highlight the use of screening strategies to target victims in financial scams and reinforce the importance of conducting due diligence.

Keywords: Financial scams, Misconduct, Screening, Cryptocurrencies
JEL classification: D40, D84, G12, G14

*Yu (corresponding author): gloriayu@smu.edu.sg, Singapore Management University. Phua: kenny.phua@uts.edu.au, University of Technology Sydney. Sang: bo.sang@bristol.ac.uk, University of Bristol. Wei: chishen.wei@polyu.edu.hk, Hong Kong Polytechnic University. We are grateful for insightful comments from Mykola Babiak (discussant), Thomas Bourveau (discussant), Shaen Corbett (discussant), Stephen Dimmock, Hogyu Jhang (discussant), Leo Liu (discussant), Thomas Matthys, Marco Navone, Tālis Putniņš, Kanis Saengchote (discussant), Wanyi Wang (discussant), Jing Xu, and conference/seminar participants at Asian Bureau of Finance and Economic Research (ABFER) Annual Conference, Asian Finance Association Annual Conference, Australian National University, Boca Corporate Finance and Governance Conference, Financial Markets and Corporate Governance (FMCG) Conference, Global AI Finance Research Conference, Hong Kong Polytechnic University, Massey University, Monash University, Nanyang Technological University, Queensland University of Technology, Singapore Management University, University of Adelaide, University of Melbourne, University of Queensland, University of Sydney, University of Technology Sydney, UWA Blockchain and Cryptocurrency Conference, and Vietnam Symposium in Banking and Finance. This study is funded by Singapore Ministry of Education (MOE) Grant 18-C207-SMU-007. We acknowledge research support from Singapore Management University and University of Technology Sydney.

Don't trust, verify: The economics of scams in initial coin offerings

Internet Appendix: https://bit.ly/3QY6tt8

Abstract

Losses from fraud and financial scams are estimated to exceed U.S. $5 trillion annually. To study the economics of financial scams, we investigate the market for initial coin offerings (ICOs) using point-in-time data snapshots of 5,935 ICOs. Our evidence indicates that ICO issuers strategically screen for naïve investors by misrepresenting the characteristics of their offerings across listing websites. Misrepresented ICOs have higher scam risk, and misrepresentations are unlikely to reflect unintentional mistakes. Using on-chain analysis of Ethereum wallets, we find that investors in misrepresented ICOs tend to be less sophisticated. Overall, our findings highlight the use of screening strategies to target victims in financial scams and reinforce the importance of conducting due diligence.

Keywords: Financial scams, Misconduct, Screening, Cryptocurrencies
JEL classification: D40, D84, G12, G14

"We embrace new technologies, but we also want investors to see what fraud looks like. I encourage investors to do their diligence and ask questions."

— Former SEC Chairman Jay Clayton on the *HoweyCoin* ICO

1 Introduction

The monetary costs of fraud and financial scams globally are estimated to exceed U.S. $5 trillion annually. There is also significant social welfare loss as victims often suffer depression, shame, and unemployment.[1] To limit and hopefully prevent such harm, researchers have begun to investigate the prevalence of fraud and uncover the circumstances under which it arises. Recent studies discover widespread misconduct among financial advisors (Egan, Matvos, and Seru, 2019), which is driven by both professional and personal circumstances (Dimmock, Gerken, and Van Alfen, 2021). There is, however, less systematic evidence on financial scams because we rarely observe how perpetrators target their victims, and victims are often reluctant to step forward.

This paper studies the economics of financial scams by exploiting a unique setting in the market for initial coin offerings (ICOs). An ICO is a form of crowdfunding for blockchain projects with mostly self-reported unverified disclosures and lax regulations. While the ICO market was believed to be rife with scams, fraud, and abuse (Howell, Niessner, and Yermack, 2020; Gensler, 2021), investors' enthusiasm for ICOs and their potentially outsized returns remained robust. ICOs raised an estimated U.S. $50 billion dollars through 2020, mostly from retail investors (PriceWaterhouseCoopers, 2020). ICOs provide an ideal laboratory to study financial scams because (i) blockchain data is publicly available in immutable ledgers, and (ii) we can observe how issuers market their offerings to prospective investors.

To analyze how ICOs were sold to investors, we collect point-in-time snapshots of self-reported ICO data from five leading ICO listing websites. ICO data have no central repository and are scattered across listing websites for prospective investors.[2] Consistent

[1] Gee and Button (2019) provide estimates of the monetary losses in 2019. Button, Lewis, and Tapley (2009) examine how victims fare in the aftermath of scams.

[2] Listing websites host the ICO listing information and are distinct from cryptocurrency exchanges or brokerages.

with Lyandres, Palazzo, and Rabetti (2021), we find widespread cross-site discrepancies of ICO data. For example, the AdHive ICO was marketed on three websites but with conflicting material disclosure (see Figure 1). Hardcap was reported as $17,490,000 on the `ICODrops` website, but $12,000,000 on the `ICOBench` and `ICORating` websites, respectively. Notably, 34% of 5,935 ICOs have discrepancies at their first appearance in our sample. Such discrepancies may constitute a violation of securities law because at least one of the reported material facts must be untrue.

- Figure 1 here -

To understand the prevalence of misrepresentations, we model the behavior of a malicious ICO issuer who faces a pool of naïve and astute investors. Investor types are unobservable, ex ante. Naïve investors are unsophisticated in that they fail to conduct due diligence and hence fall for the ICO scam. In contrast, astute investors carefully evaluate the offering and eventually refrain from funding it. Moreover, they consume the issuer's time and resources by requesting information or raising questions on public forums. Thus, astute investors are undesirable targets because they impose costs on the issuer but ultimately do not fund the scam. Ideally, the issuer screens these investors out as early as possible.

Our main hypothesis is that malicious issuers use misrepresentations along with other suspicious activities to target naïve investors and screen out astute ones. Astute investors notice the misrepresentations, deduce that the ICO is fraudulent, and immediately dismiss the offering without consuming the issuer's time and resources. However, naïve investors overlook these misrepresentations and remain viable victims of the ICO scam. Thus, the remaining investors are likely to be naïve—the ideal targets of the malicious issuer. The cost of "servicing" victims in our model is a prevalent feature across many scams. In fact, vigilantes purposely exploit this cost by posing as victims and holding tedious, unfruitful conversations with tech-support scammers.[3]

We test the prediction that misrepresented ICOs have higher scam risk. To identify ICO scams, we collect crowdsourced scams from `DeadCoin.com` and corroborate these records with reports from news articles, message boards, and regulatory authorities. Our hazard regressions reveal that the odds of a ICO scam more than triples when there is at

[3]In a recent Newsweek interview, an online vigilante Kitboga (alias) said, "[...] *important for everyone to know* [...] *how much these scammers hate when you ask questions".* The former SEC chairman Jay Clayton also encouraged prospective investors to ask questions to ICO issuers (SEC, 2018).

least one misrepresentation. At the intensive margin, one more misrepresentation raises the odds of a scam by 14.0%. To sharpen our analysis, we focus on misrepresentations of basic, nondifferentiating ICO characteristics. These types of misrepresentations should be a potent screen for investor naïvety because these characteristics are fundamental in due diligence. Consistent with this idea, we find that the such misrepresentations more strongly predict scam risk.

The economic insights from our findings may generalize to the operations of other scams. For example, an infamous email hoax solicits victims to send money to a fictitious Nigerian prince in exchange for a large fortune (Herley, 2012).[4] Because this solicitation is time-consuming and labor-intensive, the perpetrator crafts an absurd narrative to repel astute individuals and target naïve victims. In contrast, online phishing scams (e.g., fake banking websites) are typically meticulous because victims directly input sensitive information without the need to interact with the scammer. Thus, phishing attacks aim to cast a wide net because the cost of attracting an astute individual is minimal.

To assess our screening mechanism more carefully, we perform on-chain analysis of wallets on the Ethereum network. We collect data on token holdings and characterize the sophistication of the typical token holder in every ICO. We find that wallets holding tokens of misrepresented ICOs (i) have lower portfolio values, (ii) are less diversified, and (iii) are less active. Thus, malicious issuers successfully target naïve investors and screen out astute ones. Furthermore, we find that `Reddit` message boards of misrepresented ICOs have fewer comments, questions, and unique users. Overall, our findings are consistent with a screening motive behind misrepresentations.

While the evidence is consistent with malicious intent, we address the lingering possibility that misrepresentation are unintentional mistakes. First, we link ICOs by their paid advisors, who are specialists hired by issuers to provide expertise in marketing, fundraising, and technical execution. If misrepresentations were unintentional, the distribution of misrepresentations should be randomly distributed around this network. Instead, we find that misrepresentation behavior is too concentrated and systematic to be explained by random chance. Using network analysis, we find suggestive evidence that misrepresentation behavior is learned from or passed through common advisors.[5]

[4]Variants of this scam date back several centuries. Eugène François Vidocq, a French private investigator, detailed in his 1832 memoirs a scam known as the "letters of Jerusalem". The scammer would write letters to solicit victims' financial help to recover fictitious treasures.

[5]Ballester, Calvó-Armengol, and Zenou (2006) show that when there are strategic complementarities in behavior, such as learning or social norms, agents who are more central in a network exhibit a higher

We briefly explore the labor market consequences of advisors who work on misrepresented ICOs. Rather than being penalized, such advisors obtain more subsequent advisory opportunities. Egan, Matvos, and Seru (2019) also find that financial advisors who commit misconduct can find future employment, albeit at less desirable firms with lower compensation. Advisors of misrepresented ICOs may be in high demand because there are many unsophisticated ICO investors and sufficiently many malicious issuers who actively solicit their advisory services.

Second, we address the possibility that misrepresentations are a symptom of low issuer quality not malice. Using disclosure practices and fundraising outcomes as proxies for ICO quality (Bourveau et al., 2021), we find no quality differences between misrepresented and non-misrepresented ICOs. Third, if the motives underlying misrepresentations are nefarious, regulatory scrutiny should deter malicious issuers from entering the ICO market. Consistent with this idea, we find that ICOs launched shortly after news of regulatory action in cryptocurrency markets have fewer misrepresentations. This pattern is unlikely driven by issuers being more careful in response to regulatory scrutiny.

We also suspect that other tactics are used alongside misrepresentations to target naïve investors. For example, malicious issuers could more precisely target less sophisticated investors by listing on websites that derive greater traffic from paid advertisements, referral links, and search engines. Using web traffic data, we find that malicious issuers prefer to promote their ICOs on listing websites with higher passive web traffic. Malicious issuers may also use celebrity endorsements to entice unsophisticated investors. Consistent with investor warnings issued by the SEC, we find that celebrity endorsements are strongly associated with ICO scam risk. Together, the results indicate that malicious issuers use a variety of methods, along with misrepresentations, to target their victims.

We conclude by performing a welfare analysis of the potential financial losses from ICO scams. A key challenge is that many scams go undetected because victims are reluctant to report losses. Thus, the socially optimal level of regulation depends on the prevalence and costs of ICO scams, which is balanced against the cost of regulation. To overcome the partial observability of ICO scams, we use detection-controlled estimation (DCE) methods (Feinstein, 1990). The DCE results indicate that nearly 40% of ICOs in our sample may be scams, but most go undetected. Total financial losses in our sample could exceed U.S. $12 billion. These large estimates could inform regulatory design for

level of this behavior.

cryptocurrency markets.

Our study contributes to a growing literature on how financial fraud is conducted. Egan, Matvos, and Seru (2019) find that the financial advisors who "specialize" in misconduct tend to target unsophisticated investors and work at firms that tolerate misconduct. Dimmock, Farizo, and Gerken (2018) find that misconduct behavior of financial advisors is learned or passed along through colleagues. Likewise, our analysis shows that victims of ICO scams are less sophisticated investors, and misrepresentation behavior appears to transmit through common ICO advisors. Like Egan, Matvos, and Seru (2019, 2022), we find that the labor market is such that advisors who work on misrepresented ICOs can obtain subsequent advisory opportunities. Overall, we demonstrate how scammers utilize a screening strategy to profitably target their victims.

Our paper also adds to evidence on the controversies surrounding cryptocurrencies (Yermack, 2015). For example, Griffin and Shams (2020) find that Tether, a digital currency pegged to the U.S. dollar, is used to manipulate bitcoin prices. Li, Shin, and Wang (2021) and Dhawan and Putniņš (2022) document choreographed pump-and-dump trading schemes in cryptocurrencies. Studies also find evidence of wash trading that artificially boosts trading volumes on crypto-exchanges (Aloosh and Li, 2019; Cong et al., 2020). Foley, Karlsen, and Putniņš (2019) find that Bitcoin facilitates a substantial amount of illicit activities. A distinguishing feature of our study is the focus on the initial offering stage. While suspicions of ICO scams abound, evidence to date is relatively scarce. Using point-in-time data, we demonstrate how unscrupulous actors target naïve investors and estimate the financial losses to scams in the cryptocurrency market.

Finally, we build on Lyandres, Palazzo, and Rabetti (2021) who document the limitations of available ICO data and the ways to characterize data quality. We find the data quality contains key information that identifies ICO scam risk. Thus, our findings add a new perspective on the determinants of ICO success (Benedetti and Kostovetsky, 2021; Deng, Lee, and Zhong, 2018; Dittmar and Wu, 2019; Howell, Niessner, and Yermack, 2020; Davydiuk, Gupta, and Rosen, 2022). Our findings may also be of interest to recent theoretical work on ICOs, which links token development to value and utility (Cong, Li, and Wang, 2020; Sockin and Xiong, 2020).

2 Institutional details

This section describes the primary features of ICO listing websites.[6] ICO issuers use these websites to market their offerings to the general public. To list an ICO, the issuer directly submits ICO information to the website for approval. Submissions require minimal technical sophistication. For example, Figure 2 contains a screenshot of the sign-up page on a representative listing website. Listings are typically free but the listing website will prominently feature and promote an ICO for an additional fee.

- Figure 2 here -

The SEC Chairman Gary Gensler and his predecessor Jay Clayton believe that most ICOs pass the Howey Test and are subject to U.S. securities laws. If so, ICOs must comply with § 240.10b-5, which specifies the conditions for securities fraud as follows (emphasis added):

> It shall be unlawful for any person, directly or indirectly, by the use of any means or instrumentality of interstate commerce, or of the mails or of any facility of any national securities exchange, (a) To employ any device, scheme, or artifice to defraud, (b) *To make any untrue statement of a material fact or to omit to state a material fact necessary in order to make the statements made, in the light of the circumstances under which they were made, not misleading,* or (c) To engage in any act, practice, or course of business which operates or would operate as a fraud or deceit upon any person, in connection with the purchase or sale of any security.

Misrepresentations of ICO characteristics violate condition (b) because at least one of the reported characteristics is false. Moreover, a misrepresentation can potentially "deceive" or mislead potential investors, which may violate condition (c) of the rule. As of December 2022, the SEC has taken regulatory actions against 64 ICOs and cryptocurrency offerings.[7] However, listing websites themselves have minimal disclosure requirements and are lightly regulated.

[6] A detailed overview of ICOs is available in the Internet Appendix.

[7] The judgments from these regulatory actions totaled U.S. $354.5 million, of which U.S. $295.8 million were refunds and U.S. $58.7 million were penalties. Additionally, 38 securities class action lawsuits have been filed against ICO issuers.

3 Why are misrepresentations so prevalent?

To understand the prevalence of misrepresentations, we develop a model to analyze how malicious issuers use discrepancies across listing websites to screen for naïve investors. Our model shares similarities with frameworks that analyze other scams in cyberspace (e.g., Herley, 2012).

3.1 The issuer's classification problem

There are three periods in the model. The malicious ICO issuer faces a mass of m investors, of which there are n naïve investors and $m - n$ astute investors. Individual investor types are ex ante unobservable. The key difference between investor types is that naïve investors will fund the ICO scam while astute investors will ultimately not. We define x_i to be the number of misrepresentations that investor i tolerates, above which the investor immediately dismisses an ICO scam. The probability density functions of x for naïve and astute investors are $\phi(x \mid \text{naïve})$ and $\phi(x \mid \text{astute})$, respectively, where d is the number of misrepresentations set by the issuer. On average, naïve investors are more tolerant than astute investors of misrepresentations: $\bar{x}^{\text{naïve}} > \bar{x}^{\text{astute}}$. But, some naïve investors may have lower x than astute ones.

- Figure 3 here -

We describe our screening model in Figure 3. In period one, the issuer sets the number of misrepresentations d, which acts as a cutoff to target a pool of potential victims.[8] For a given d, the fraction of naïve investors who immediately dismisses the ICO scam is $\int_0^d \phi(x \mid \text{naïve})\, dx$. The remaining fraction of naïve investors who remain potential victims is the complementary conditional CDF $\int_d^\infty \phi(x \mid \text{naïve})\, dx$. Likewise, the fraction of astute investors who remain is $\int_d^\infty \phi(x \mid \text{astute})\, dx$. Because the complementary

[8] To crystallize the screening mechanism, our model abstracts away from investors' incentives to participate in the ICO market. There are at least two reasons why investors may be willing to fund ICOs despite the prevalence of scams. First, investors may be attracted to the high skewness in the distribution of ICO returns. Conditional on successful listings on cryptocurrency exchanges, Lyandres, Palazzo, and Rabetti (2021) find that the average (maximum) ICO return on the first trading day is 384% (3,870%). These patterns imply that investors may also be willing to make many losing bets in hopes of capturing an investment that yields outsized returns. Second, overconfident investors (Daniel, Hirshleifer, and Subrahmanyam, 1998; Odean, 1998) may be willing to participate in the ICO market because they overestimate their abilities to evaluate ICOs and avoid scams.

conditional CDFs are nonincreasing in d, the targeting strategy is determined by d, such that a higher d targets lower (higher) fractions of both naïve and astute investors.

In period two, the remaining astute investors (i.e., those who have not dismissed the scam) may request more information from the issuer directly or raise questions about the ICO on public forums such as `Reddit`, `Twitter`, and `Bitcointalk`. The issuer cannot avoid these costs by ignoring investor queries without raising suspicion. Without loss of generality, we assume that naïve investors do not raise such queries. The issuer incurs a cost C per remaining astute investor that reflects the time and resources needed to address questions.

In the final period, naïve investors fund the scam while astute investors ultimately do not. Targeting a naïve investor yields the issuer a net profit G. An astute investor refrains from funding the scam, causing a net loss C. Astute investors are undesirable because they consume resources but provide no financial rewards to the issuer. The issuer's expected profits $\mathbb{E}(\Pi)$ can be expressed as a function of d.

$$\frac{\mathbb{E}(\Pi)}{m} = zG \overbrace{\int_d^\infty \phi(x \mid \text{naïve})\, dx}^{\text{frac. naïve investors targeted}} - (1-z)\, C \overbrace{\int_d^\infty \phi(x \mid \text{astute})\, dx}^{\text{frac. astute investors targeted}}, \tag{1}$$

$$\text{where} \quad z = \frac{n}{m}$$

Consider an indiscriminate targeting strategy that abandons the screening strategy. Because $\int_0^\infty \phi(x \mid \cdot)\, dx = 1$, the issuer targets all investors by choosing $d = 0$. Imposing these constraints and $\mathbb{E}(\Pi) > 0$, we obtain equation (2).

$$z = \frac{n}{m} > \frac{C}{G + C} \tag{2}$$

When $C > 0$, equation (2) implies that an indiscriminate targeting strategy is profitable only if the fraction of naïve investors in the investor mass is greater than the ratio $C/(C+G)$. For example, suppose 0.1% of investors are naïve and $G = \$500$, then C can at most be $0.001/(1-0.001) \times \$500 = \$0.50$ per investor. Indiscriminate targeting can also be profitable in the special case of $C = 0$. However, this case is unlikely given the threat of reputation loss and regulatory scrutiny, and resources required to service investors' queries. Finally, targeting all investors is also profitable in the edge case of $G \to \infty$.

3.2 Misrepresentations as a screening device

We illustrate the tradeoffs implied by various targeting strategies in Figure 4. The figure shows probability density plots of x for astute (black) and naïve (red) investors. Shaded areas in black and red represent the complementary conditional cumulative distributions $\int_d^\infty \phi(x \mid \text{astute})\, dx$ and $\int_d^\infty \phi(x \mid \text{naïve})\, dx$, respectively. In Subfigure 4a, the malicious issuer adopts a conservative targeting strategy that avoids many costly astute investors by choosing a high number of misrepresentations (high d). But, this strategy forgoes many profitable naïve investors in the population. In Subfigure 4b, the issuer sets an aggressive targeting strategy by choosing a low d. While this strategy captures more naïve investors, it retains more costly astute investors who erode the issuer's profits. For completeness, we solve for the issuer's optimal targeting strategy in Appendix A.

- Figure 4 here -

The model formalizes our argument that malicious issuers use misrepresentations to screen for investor sophistication. As investor sophistication is unobservable, issuers implement a strategy that induces naïve investors to self-identify. ICO misrepresentations will raise suspicions among astute investors, who perform due diligence and immediately dismiss the ICO. The remaining investors tend to be naïve—the ideal targets of the malicious issuer.[9] Thus, the issuer earns more profits by targeting naïve investors and screening out their astute counterparts.

The screening model also produces a testable hypothesis that ICO misrepresentations predict scam risk. We empirically test this prediction. The model provides a theoretical explanation for why 34% of ICOs have misrepresentations at their first appearances in our sample. It is tempting to dismiss this behavior as simply issuers' carelessness. But, the sheer quantity of misrepresentations, the amount of money at stake, and the minimal technical sophistication required to submit ICO information to listing websites (see, Figure 2) should indicate that other factors are at work.

[9]The use of misrepresentations as a screening device in ICO scams has parallels with other notorious scams such as the advance-fee scams. The advance-fee scammer promises prospective victims in e-mails a large sum of money in return for a small upfront administrative fee. These e-mails often contain grammatical errors and use outlandish language. In some cases, the emails also tell implausible stories, in which the scammer impersonates a member of the Nigerian royal family. The inclusion of these telltale signs is strategic (Herley, 2012). Astute people, who could waste the scammer's time and resources, recognize these signs and ignore the emails. Whereas, only the most gullible victims would respond to the emails, hence self-identifying their gullibility to the scammer.

4 Data, variables, and descriptive statistics

This section describes our data collection process, defines the main variables, and presents the descriptive statistics of our sample.

4.1 Data sources

We systematically collect point-in-time ICO data from five major websites that aggregate ICO listings—(i) `ICOBench` (ii) `ICOCheck` (iii) `ICOData` (iv) `ICODrops` (v) `ICORating`. We select these five listing websites based on (i) their popularity reported by Alexa Traffic Rank on August 15th 2018, (ii) the number of ICOs covered, and (iii) the technical feasibility of scraping the websites.[10] On the 15th of every month from August 2018 to August 2019, we scrape ICO data from these five websites. In total, we have 13 data collection events and a time-series of ICO characteristics for every ICO-website pair. Because ICO identifying information may vary across websites, we manually cross-check all ICOs and designate a set of unique identifiers to every ICO in our sample. To resolve residual conflicts in our collected data, we hand-check other Internet sources. Thus, we alleviate concerns of variation in ICO names, misspellings, and name changes. Overall, our sample contains 5,935 matched ICOs.[11]

We collect ICO scam allegations from a prominent crowdsourced anti-fraud project hosted on `DeadCoins.com`. The `DeadCoins` website curates a list of ICOs that are alleged scams, alongside a summary of every scam and corresponding information sources. Reasons behind scam allegations include charges by regulators for fraudulent activities, cancellation by exchanges, obvious technical flaws, disappearance of ICO issuers, and prolonged inactivity. For example, the Shopin token was marked as "dead" (i.e., inactive) on `Deadcoins` following a SEC complaint. Subsequently, the founders and company behind the Shopin token were charged with securities fraud and violations of registration processes.

To mitigate concerns of false positivity, we corroborate every `Deadcoin` scam alle-

[10] Based on the Alexa Traffic Rank on November 30th 2018, Lyandres, Palazzo, and Rabetti (2021) obtain ICO data from `ICOBench`, `ICODrops`, `ICORating`, `ICOMarks`, and `ICOData`. We replace `ICOMarks` with `ICOCheck` for the latter two considerations.

[11] The numbers of unique ICOs covered by the listing websites are: `ICORating` (4,166), `ICOBench` (4,021), `ICOData` (1,896), `ICODrops` (625), and `ICOCheck` (580).

gation with several media sources.[12] First, we check whether the ICO is reported by regulatory authorities (e.g., SEC, DoJ). Second, we search on Factiva for press coverage (e.g., news articles, website articles, journal articles) of the ICO scam. Third, we search popular online forums and social media (e.g., `Reddit`, `Cryptocompare`) for mentions of the ICO scam. We admit an alleged ICO scam into our sample only if it is found on at least one of the above three media channels. In total, we match 115 ICO scams to our sample.

We collect regulatory filings (Form D, Form 1-A, and Form C) of ICOs that are available on the SEC EDGAR database. We search the database using the keywords "token", "ICO", "initial coin offering", "coin", and "crypto". We then manually determine whether every filing is ICO-related. We first read the filing document and check whether it pertains to an initial coin offering or other types of offering. If this information is not stated, we then use the firm name written in the document combined with the keywords "ICO", "offering", "token" to perform a search on SEC EDGAR. All else failing, we use the names of persons (i.e., founders, CEOs, and directors) in the filing combined with the above keywords to perform another search on SEC EDGAR. In our sample, 77, two, and eight ICOs have filed for a Form D, Form 1-A, and Form C, respectively.

4.2 Variables

Our key independent variable is the *misrep* of an ICO—the total number of cross-website discrepancies of 13 commonly reported characteristics at its first appearance in our sample.[13] Figure 5 visualizes the proportion of ICOs with at least one cross-website discrepancy by these characteristics at first appearances in our sample. The most common misrepresented characteristic is *whitelist* (36.9%). Other commonly misrepresented characteristics are *start date* (25.9%), *end date* (26.12%), *presale* (20.7%), and *banned* (16.6%). Misrepresentations in *softcap*, *ticker*, and *country* are uncommon.

- Figure 5 here -

In our empirical tests, we control for a suite of variables that describes the fundraising structure and regulatory environment of an ICO. The following control variables are coded

[12] Notably, the `Deadcoin` website also prominently displays a form to contest scam allegations.

[13] The 13 characteristics used to construct *misrep* are *banned*, *whitelist*, *presale*, *hardcap*, *softcap*, *accept BTC*, *accept ETH*, *accept USD*, *ticker*, *start date*, *end date*, *duration*, and *country*.

as indicators that switch on if the ICO has the corresponding features. An ICO is *banned* if it is banned by at least one regulatory authority. A *whitelist* allows an ICO issuer to limit the sale of tokens to a selected group of registered investors. An ICO can hold a *presale* round to sell tokens before the public fundraising campaign is set up. The *hardcap* is the upper limit on the number of tokens that can be sold in an ICO. The *softcap* is the minimum amount of funds that must be raised in an ICO, or else funds are returned to investors and the project is discontinued. We control for payment options in the ICO with *accept BTC* (*ETH*, *USD*). The last indicator is *SEC filing*, which switches on if the ICO has regulatory filings with the SEC. The *duration* of an ICO is the length of its fundraising period in days. Finally, the *enforcement* and *disclosure* indices from La Porta et al. (2000) control for the regulatory environment in the ICO's country of registration.

4.3 Descriptive statistics

Table 1 reports summary statistics of our sample. Panel A reports that the average ICO has 1.28 *misrep*, and 34% of ICOs have at least one *misrep*. 95% of ICOs are banned in at least one country, which is unsurprising as ICOs are illegal in several countries (e.g., China, Egypt, Morocco). About half of ICOs impose selectivity in their investor clientele or fundraising structures; 55% of ICOs have an investor *whitelist*, and 47% of them have *presale* rounds. Most ICOs (70%) have a *hardcap* in their fundraising structures, but only a minority (26%) have a *softcap*. ETH (USD) is the most (least) popular payment currency among ICO issuers. Fewer than 1% of ICOs in our sample have regulatory filings with the SEC. The fundraising period for the average (median) ICO is 54 (37) days. Panel B reports the Pearson pairwise correlations among our variables. Our key variable *misrep* is weakly correlated with most variables, except for *presale* (0.31), *hardcap* (28%), and *accept ETH* (31%).

- Table 1 here -

Table 2 reports differences in ICO scam rates and characteristics between (i) ICOs with at least one *misrep* and (ii) ICOs with no *misrep*. We observe significant differences across the two groups. ICOs with at least one *misrep* are more likely to incur a scam allegation (4% vs. 1%). Such ICOs also have weaker governance—they are less likely to have a investor *whitelist* (46% vs. 60%) and are more likely to hold a *presale* funding round (68% vs. 36%). These ICOs are also more likely to have salient attributes that imply limited supply—misrepresented ICOs have shorter fundraising periods (*duration* of

48 days vs. 58 days) and are more likely to have a *hardcap* (89% vs. 60%). Misrepresented ICOs also accept a wider range of payment options.

- Table 2 here -

5 Misrepresentations and ICO scams

We design two sets of tests for our hypothesis that malicious ICO issuers use misrepresentations to screen for naïve investors. First, we perform survival analysis to examine whether misrepresented ICOs are more likely to be scams. Second, we use data from the Ethereum blockchain and `Reddit` message boards to assess our screening mechanism more carefully.

5.1 Survival analysis: ICO scam risk

We perform survival analysis to test the hypothesis that ICOs with more misrepresentations are more likely to be scams. Our objective is to track the survival time of an ICO—the time elapsed between its entry into our sample and occurrence of a scam allegation. There are three notable features of our empirical setting that are well accommodated by survival analysis. First, ICOs can enter and exit our sample at different points in time. Second, we only have information about which ICOs survive (i.e., remain in our sample) at any point in time. An ICO exits our sample when it incurs a scam allegation. Otherwise, it is right-censored. Right-censoring occurs if an ICO (i) becomes unlisted on listing websites, or (ii) survives till the end of our 13-month observation window without a scam allegation.[14] Third, survival times usually do not have normal distributions.

We plot the proportion of surviving ICOs—the survival function $S(t)$—with respect to survival time t. First, we sort ICOs by their *misrep* into four groups. Where r_t is the number of surviving and uncensored ICOs instantaneously before time t, and f_t is the number of ICOs that incur scam allegations, we next compute the survival function

[14] Right-censored observations are not necessarily cleared of scams.

within every group:

$$S(t) = \begin{cases} \frac{(r_t - f_t)}{r_t} \times S(t-1), & \text{for } t > 0 \\ 1, & \text{for } t = 0 \end{cases} \tag{3}$$

Figure 6 shows that all four groups begin with $S(0) = 1$ because our sample precludes ICOs that are known to be scams. As time progresses, the survival functions of all four groups decline as ICO scams are flagged on the `DeadCoin` website. However, we find that the survival function in the high-*misrep* group declines most quickly. In comparison, the decline in survival function of the low-*misrep* group is substantially slower. This difference in trends is first evidence that *misrep* is positively associated with the incidence of ICO scams.

- Figure 6 here -

We now estimate the effect of *misrep* on the incidence of ICO scams with Cox regression models. Where $h(t) = -\frac{\delta}{\delta t} \log S(t)$ is the expected hazard that denotes the rate of ICO scams conditional on survival up to time t, and $h_0(t)$ is the baseline hazard when all covariates equal zero, we estimate specification (4).

$$h_i(t) = h_0(t) \exp\left(\beta_1 misrep_i + \mathbf{X}_i^\top \boldsymbol{\beta}\right) + \epsilon_i \tag{4}$$

The vectors $\mathbf{X}$ and $\boldsymbol{\beta}$ represent vectors of control variables and their corresponding estimated coefficients, respectively. For ease of interpretation, we express estimated coefficients as hazard ratios. A hazard ratio that equals one implies that an increase in the covariate has no effect on the hazard of ICO scams. If the hazard ratio is above (below) one, then the covariate is associated with an increase (decrease) in the hazard of ICO scams.

- Table 3 here -

Our estimates in Table 3 show that ICOs with higher *misrep* are more likely to be scams. Column 1 shows that the presence of *misrep* more than triples ($t = 5.46$) the hazard ratio of ICO scams. At the intensive margin, we find in column 2 that an additional *misrep* is associated with a 25.3% ($t = 6.71$) rise in hazard of ICO scams. We further add coverage quartile fixed effects and stratify our ICOs by their calendar-quarter

cohorts in column 3.[15] These augmentations address two concerns. First, the coverage fixed effects alleviate the concern that *misrep* is mechanically driven by the number of websites that an ICO is listed on. Second, the stratification allows ICOs to have cohort-specific baseline hazards $h_0(t)$—this absorbs heterogeneity in hazard of ICO scams across cohorts. In this augmented specification, we find that an additional *misrep* increases the hazard of ICO scams by 14.0%. ($t = 2.18$). To add color to our findings, we focus on misrepresentations in a subset of basic ICO characteristics.[16] Basic ICO characteristics are salient, requires little expertise to understand, and should be fundamental to investors' due diligence. In column 4, we find that an additional $misrep^{basic}$ increases the hazard of ICO scams by 24.0% ($t = 4.86$).[17] This finding reinforces our screening hypothesis—investors who fail to notice discrepancies in the most basic ICO characteristics likely also fail to perform due diligence. Thus, such discrepancies are particularly potent screens for investor sophistication.

Overall, we find that misrepresentations of ICO attributes on listing websites are a powerful ex-ante predictor of scams. Consistent with our screening hypothesis, the predictive effect is primarily driven by misrepresentations of basic ICO information. Our findings suggest that simple cross-website verification of ICO attributes is an effective form of due diligence for prospective investors.

5.2 Assessing the screening mechanism

To assess our screening mechanism more carefully, we extract data from the Ethereum blockchain.[18] The data contain token holdings and transaction activities of cryptocurrency wallets (henceforth, wallets). Using wallet-level data, we examine the relation between the sophistication of the typical token holder and misrepresentations in an ICO. The Internet Appendix contains details of data collection in this test.

We characterize the (lack of) sophistication of a typical token holder by computing

[15] Coverage is the number of listing websites that an ICO is listed on. Two ICOs are in the same cohort if their ICO start dates are in the same calendar quarter.

[16] Basic ICO characteristics are *ticker*, *country*, *banned*, *start date*, *end date*, *duration*, and acceptable payment modes. Nonbasic ICO characteristics are *softcap*, *hardcap*, *whitelist*, and *presale*.

[17] In contrast, we find in untabulated results that misrepresentations of nonbasic characteristics has a negligible predictive effect ($-3.3\%, t = 0.40$) on ICO scam risk.

[18] The Etherium blockchain is a digitally distributed, decentralized, public ledger of all transactions that occurred on the network. Most ICO tokens adopt the ERC-20 (Ethereum Request for Comments 20) standard, which facilitates interoperability with other tokens on the Ethereum network.

three wallet level measures. First, we define the *value* of a wallet by computing the total portfolio value in U.S. dollars of all tokens held. To the extent that wealth positively correlates with sophistication, we expect that unsophisticated investors have lower wallet values. Second, we define *diversity* as the number of distinct tokens held. Unsophisticated investors may possess less diversified wallets with fewer distinct ICO tokens. Third, we define *activity* as the number of wallet transactions. Unsophisticated investors with less technical or trading expertise may make fewer transactions. We aggregate these measures at the ICO level by taking the medians of every measure.

To test whether malicious issuers successfully use misrepresentations to screen for naïve investors, we estimate Poisson regressions in specification (5) because our outcome variables are non-negative (Cohn, Liu, and Wardlaw, 2021). The dependent variable is *value*, *diversity*, or *activity*. The key independent variable is $\mathbb{1}(misrep > 0)$—an indicator that switches on if the ICO has at least one *misrep* at its first appearance in our sample. Our models include ICO calendar-quarter cohort fixed effects, and standard errors are clustered by these cohorts. For ease of interpretation, we express estimated coefficients as incidence rate ratios.

$$\{\log(value_i), \log(diversity_i), \log(activity_i)\} = \alpha + \beta_1 \mathbb{1}_i(misrep_i > 0) + \mathbf{X}_i^\top \boldsymbol{\beta} + \epsilon_i \quad (5)$$

Panel A of Table 4 shows that less sophisticated investors are more likely to hold tokens of misrepresented ICOs. Column 1 indicates that the typical investor in a misrepresented ICO has a 60.1% ($t = 2.61$) lower wallet *value*. In column 2, switching on $\mathbb{1}(misrep > 0)$ is associated with a 19.7% ($t = 2.88$) decline in *diversity*. Column 3 shows that transaction *activity* of investors in misrepresented ICOs is 9.0% ($t = 2.62$) lower. Overall, the results support the view that misrepresented ICOs attract less sophisticated investors as measured by wallet value, diversity, and transaction frequency.[19]

- Table 4 here -

We further test our screening mechanism by hand-matching ICOs to their `Reddit` message boards and tracking user activity with the `Pushshift` API. Our model predicts that the average investor in a misrepresented ICO should have fewer queries. Con-

[19] These patterns are inconsistent with the alternative view that sophisticated investors in an exuberant ICO market "ride the bubble" (Abreu and Brunnermeier, 2003; Brunnermeier and Nagel, 2004; Griffin et al., 2011).

sistent with this prediction, Panel B of Table 4 shows that `Reddit` message boards of misrepresented ICOs have fewer (i) comments ($\downarrow$ 48.8%, $t = 4.67$), (ii) questions ($\downarrow$ 46.6%, $t = 2.95$), and (iii) unique users ($\downarrow$ 20.0%, $t = 1.96$). As a robustness check, we show in the Internet Appendix that our findings in Table 4 also hold on the intensive margin with *misrep*. Overall, our findings are consistent with a screening motive behind misrepresentations.

6 Are misrepresentations unintentional mistakes?

The evidence in the previous section indicates that misrepresented ICOs are more likely to be scams. This finding is consistent with our main hypothesis that malicious issuers use misrepresentations to screen for naïve investors. Nevertheless, it is difficult to know the true motives behind misrepresentation behavior. An alternative explanation is that ICO misrepresentations could simply be unintentional mistakes. We design three sets of tests to address this explanation. First, we apply network analysis to assess systematic patterns of misrepresentation behavior in the ICO ecosystem. Second, we examine the relation between misrepresentations and ICO quality. Third, we focus on the misrepresentation behavior of ICOs launched shortly after news of regulatory actions taken by U.S. authorities.

6.1 Systematic patterns of misrepresentation behavior

To substantiate our view that the use of misrepresentation is strategic, we apply network analysis to assess unusual patterns of this behavior among ICO issuers. If misrepresentations are intentionally and strategically deployed, they should leave systematic footprints throughout the ICO ecosystem. Specifically, we examine whether ICO advisers (henceforth, advisors) play a role in promoting misrepresentation behavior. Advisors are hired by ICO issuers to provide technical, marketing, and economic expertise. About 60% of ICOs in our sample hire an advisor. Advisors are also controversial—some have been convicted of illegal touting and tax evasion, while others have allegedly failed to perform basic due diligence on client ICOs.

Because advisors often work on multiple ICOs, they could play a role in promoting misrepresentation behavior. We hypothesize that misrepresentation behavior is correlated among ICOs that share common advisors. This correlation could arise from strategic

complementarities that are typical in criminal behavior. Complementarities in misrepresentation behavior can materialize in two ways. First, there is no formal way to learn the effective use of misrepresentations as a screening device. So, malicious issuers may have to learn from their peers via common advisors who convey know-how about the use of ICO misrepresentations. This learning channel implies that a malicious issuer's payoffs from misrepresentations are higher with technological transfers from other issuers of misrepresented ICOs. Second, misrepresentation behavior may be viewed as an acceptable norm among ICOs that share common advisors. An issuer who observes the use of misrepresentations by other issuers may infer that this behavior is commonplace. In response, the issuer is likely to use more misrepresentations, which symmetrically leads other issuers to the same inference and to do likewise.

- Figure 7 here -

Ballester, Calvó-Armengol, and Zenou (2006) provide a network model of behavior under strategic complementarities. We apply the theoretical insights of that model to a network of ICOs linked by common advisors. If advisors play a role in promoting misrepresentation behavior, ICOs with higher Katz centrality in the network should exhibit more *misrep*. The Internet Appendix contains details on the definition of Katz centrality and the network model. To construct the ICO network, we manage to match 2,110 advisors with 2,271 ICOs using data extracted from the `ICOBench` listing website.[20] In this network, we link two ICOs if they share at least one common advisor. We present a circular layout of this network in Figure 7. ICOs are arranged according to their *misrep* on the circumference of the circle. As we move along the circumference in the clockwise direction, the ICOs have more *misrep*. Lines inside the circle represent links between ICOs. We observe that ICOs with more *misrep* tend to locate in regions with higher densities of links. Generally, such ICOs are also more central in the network.

- Table 5 here -

To examine the relation between Katz centrality and *misrep* more rigorously, we estimate Poisson regressions in Table 5. Estimated coefficients are presented as incidence rate ratios. Consistent with our model predictions, column 1 shows that a 10% increase

[20]This test has a smaller sample because we must exclude ICOs that either have no advisors or are unlinked to any ICOs.

in Katz centrality is associated with a 4.6% ($t = 2.27$) rise in *misrep*.[21] Next, we conjecture that transmissions of misrepresentation behavior is stronger between two ICOs if they share more common advisors. Thus, we also construct a weighted ICO network, in which links are weighted by the number of common advisors. In column 2, we find a quantitatively similar effect using weighted links—a 10% increase in Katz centrality is associated with a 5.4% rise ($t = 2.17$) in *misrep*. In the next two columns, we use as our key independent variable an indicator $\mathbb{1}(\textit{high centrality})$ that switches on if an ICO has an above-median Katz centrality. Columns 3 and 4 report that central ICOs have 6.1% ($t = 1.96$) and 6.7% ($t = 2.25$) higher *misrep* than peripheral ICOs, respectively.

Our empirical results in Table 5 support predictions from a network model—central ICOs use more misrepresentations. Owing to strategic complementarities, we find systematic patterns of misrepresentation behavior among advisor-linked ICOs. These patterns reject the idea that misrepresentations are merely idiosyncratic, random, unintentional mistakes. Overall, while advisors could be valuable information and service intermediaries in the ICO market, some may facilitate the promotion of malignant behaviors.

6.2 Misrepresentations and ICO quality

Misrepresentations may simply be unintentional mistakes. Suppose low quality issuers fail to exert the necessary effort to accurately market their offerings on listing websites. Then, to the extent that such issuers produce poorer blockchain projects, *misrep* should be negatively associated with ICO quality. While the lack of disclosure verifiability and regulatory oversight raises concerns of cheap talk, theory suggests that voluntary disclosures can still be informative (e.g., Crawford and Sobel, 1982; Gigler, 1994; Stocken, 2000). Thus, high quality ICOs may voluntarily disclose more to distinguish themselves from low-quality ICOs (Bourveau et al., 2021). First, ICO issuers may voluntarily disclose the source code of their smart contracts on blockchain explorer services such as `Etherscan.io`. Second, issuers may also post on `Etherscan` the security audits of their source code.

To test whether misrepresentations merely reflect poor ICO/issuer quality, we examine the relation between *misrep* and the code disclosure practices of ICOs. To operationalize this test, we define the indicator $\mathbb{1}(\textit{code posted})$ to equal one if the ICO

[21] We calculate this economic magnitude as follows: $\log(1.1) \times (1.485 - 1) = 0.046$.

discloses its source code on `Etherscan.io` and equals zero otherwise. Likewise, the indicator $\mathbb{1}(\textit{code audited})$ switches on if the ICO posts a security audit of its source code on `Etherscan.io`. We estimate logistic regressions following specification (6). The term p is the probability that $\mathbb{1}(\textit{code posted})$ (or, $\mathbb{1}(\textit{code audited})$) switches on in an ICO. The vectors $\mathbf{X}$ and $\boldsymbol{\beta}$ represent vectors of control variables and their corresponding estimated coefficients, respectively. For ease of interpretation, we express estimated coefficients as odds ratios.

$$\log\left(\frac{p_i}{1-p_i}\right) = \alpha + \beta_1 misrep_i + \mathbf{X}_i^\top \boldsymbol{\beta} + \epsilon_i \tag{6}$$

We estimate the relation between ICO quality and misrepresentations in Table 6. In column 1, we find that an additional *misrep* is weakly associated with 1.6% ($t = 0.31$) lower odds of the ICO disclosing its code on `Etherscan.io`. This finding fails to support the idea that misrepresentations are symptomatic of poor issuer quality. We corroborate this finding in column 2, which shows a weak relation between *misrep* and odds of the ICO posting a security audit of its source code ($+1.1\%, t = 0.26$). Our findings suggest that misrepresented and non-misrepresented ICOs are indistinguishable in quality from an investor's perspective.

As an additional test, we adopt a market-based measure of ICO quality. Bourveau et al. (2021) find that market participants can effectively gauge the quality of ICOs. If misrepresentations reflect low quality, then we expect misrepresented ICOs to raise less funds. Because the amount of funds *raised* is a strictly non-negative quantity, we estimate a Poisson regression in column 3. Here, we find that the link between *misrep* and the amount of funds raised in the ICO campaign is statistically insignificant ($+5.8\%, t = 1.04$). This finding further supports our view that misrepresentations are unrelated to ICO quality.

- Table 6 here -

Overall, we find that misrepresentations do not meaningfully vary with ICO quality. Thus, our findings reject the view that misrepresentations are merely unintentional mistakes, reflecting low issuer quality. Instead, malicious issuers strategically use misrepresented ICO information to target naïve investors.

6.3 Regulatory scrutiny and misrepresentations

If misrepresentations are nefarious and not simply careless mistakes, then the threat of regulatory action should deter malicious issuers from entering the ICO market. We expect ICOs launched after periods of higher regulatory scrutiny to have fewer *misreps*, on average.[22] To test the deterrence effect, we collect news of regulatory actions taken by the U.S. authorities. As Appendix B shows, these regulatory actions primarily involve ICO fraud and conflicts of interest. None of these actions specifically mention inaccurate disclosures on listing websites.

$$\left\{\log\left(\frac{p_i}{1-p_i}\right), \log\left(misrep_i\right)\right\} = \alpha + \beta_1 regulatory\ scrutiny_i + \mathbf{X}_i^\top \boldsymbol{\beta} + \epsilon_i \tag{7}$$

We first measure *regulatory scrutiny* as the number of regulatory news articles released in the month prior to the first appearance of every ICO in our sample. Next, we test the effect of *regulatory scrutiny* on the use of misrepresentations. We estimate logistic and Poisson regressions according to specification (7). The first outcome variable in this specification is the logit of p, which is the probability that the ICO has at least one misrepresentation at its first appearance in our sample. To test the deterrence effect on the intensive margin, we use *misrep* as the second outcome variable. Because *misrep* is a strictly non-negative quantity, we estimate Poisson regressions (Cohn, Liu, and Wardlaw, 2021). The vectors $\mathbf{X}$ and $\boldsymbol{\beta}$ represent vectors of control variables and their corresponding estimated coefficients, respectively.

- Table 7 here -

Our results in Table 7 show that ICOs that launch immediately after regulatory scrutiny have fewer misrepresentations. We estimate a logistic (Poisson) regression in column 1 (2) where the dependent variable is $\mathbb{1}(misrep > 0)$ ($misrep$). Column 1 shows that the release of an additional regulatory news article decreases the odds of a misrepresented ICO in the next month by 20.5% ($t = 2.13$). On the intensive margin, column 2 shows that ICOs have 16.2% ($t = 2.91$) fewer misrepresentations per regulatory news article.

[22] Despite the heightened regulatory scrutiny, malicious issuers may still choose to enter the ICO market. In that case, the additional threat of regulatory penalties may lead to a higher C in our model. Equation (A.1) shows that a higher C leads the malicious issuer to pursue a more conservative screening strategy by choosing a higher d^*. Thus, our empirical findings in this section likely reflect a lower bound of the deterrence effect.

These patterns suggest that the use of misrepresentations likely reflects strategic and malicious behavior. Notably, we find a link between *regulatory scrutiny* and misrepresentation behavior although our sample of news articles does not mention the latter.

Alternatively, our results may merely reflect greater care taken by ICO issuers in the face of regulatory scrutiny. If regulatory scrutiny simply spurs greater care among issuers, corrections of prior misrepresentations should also be more likely after the news events. Using our point-in-time data snapshots, we track whether issuers correct their misrepresentations from month to month. We test this alternative story with $\mathbb{1}(\Delta misrep < 0)$—an indicator that switches on when an ICO has a decline in misrepresentations from the previous month. Column 3 shows that there is no statistically significant link between *regulatory scrutiny* and $\mathbb{1}(\Delta misrep < 0)$. This finding fails to support the view that regulatory scrutiny merely spurs greater care taken by issuers. The Internet Appendix shows that our findings hold with a binary measure of regulatory scrutiny.

7 Other suspicious actions

While malicious ICO issuers use misrepresentations to target naïve investors, such issuers may also use other tactics to screen for investor sophistication. We collect data on two examples of such actions—celebrity endorsements and choice of listing websites—and test their predictive effects on ICO scam risk.

First, the U.S. SEC warns on an investor education website that celebrity endorsements of ICOs are prominent red flags of investment scams.[23] Celebrity endorsements may be a potent screening device because naïve investors are more likely to act on financial advice offered on social media, particularly when it comes from famous individuals. To collect data on celebrity endorsements, we conduct web searches using combinations of these keywords: "celebrity"/"promoter"/"influencer" and "ICO"/"initial coin offering"/"token". Next, we read all relevant search results and identify ICOs that are promoted by celebrities. To ensure completeness of our search efforts, we also search for the same combinations of keywords on the Factiva database. Our sample includes celebrities who span the entertainment, sports, business and media sectors.

Second, most ICOs are promoted on multiple, but not all, listing websites. We exam-

[23] Source: `https://www.investor.gov/ico-howeycoins`

ine whether malicious issuers choose listing websites based on the characteristics of their web traffic. Using data from SEMrush—a web traffic analytics vendor—we measure the quantities of passive and active web traffic in each of the five listing websites. Specifically, passive web traffic counts visitors referred to a listing website via paid advertisements, third-party referral links, and search engines. Whereas, active web traffic counts visitors who access a listing website by directly typing its Uniform Resource Locator (URL) in browsers or through the use of saved browser bookmarks. Then, we define the *web traffic ratio* of an ICO as the ratio of passive traffic to active traffic, aggregated across the listing websites that list it in the month prior to its start date. We conjecture that active web traffic reflects a purposeful and targeted pattern of information acquisition, which is typical of more sophisticated investors.

- Table 8 here -

To test whether celebrity endorsements and strategic choices of listing websites predict ICO scams, we estimate Cox regressions in Table 8. We express estimated coefficients as hazard ratios. The key independent variable in column 1 is $\mathbb{1}(\textit{celebrity})$—an indicator that switches on if an ICO is endorsed by a celebrity. Here, we find that the scam risk of an ICO with a celebrity endorsement is more than 25 times ($t = 10.64$) that of an ICO without one. This finding supports the warning issued by the SEC that celebrity endorsements are red flags of investment scams. In column 2, we examine whether celebrity endorsements subsume the predictive effect of *misrep* on ICO scam risk. They do not. While $\mathbb{1}(\textit{celebrity})$ remains a strong predictor of ICO scam risk, we find that an additional *misrep* raises the odds of a scam by 14.5% ($t = 2.04$). This result suggests that misrepresentations and celebrity endorsements are distinct screening devices in the malicious issuer's repertoire. Because only a minority of ICOs are endorsed by celebrities, keeping a lookout for misrepresentations remains incrementally useful.

Column 3 shows that a unit increase in *web traffic ratio* is associated with a 26.5% ($t = 2.23$) higher odds of an ICO scam. This pattern suggests that malicious issuers strategically choose listing websites that receive a relatively larger share of passive web traffic. Through the lens of our theoretical framework in Section 3, this strategic choice has a similar effect to choosing an investor mass with a higher density z of naïve investors. In turn, a higher z increases the issuer's expected profits, ceteris paribus. In column 4, we find that *misrep* remains a positive and statistically significant predictor of ICO scam risk. Thus, misrepresentations have a screening effect incremental to that from the strategic choice of listing websites.

Overall, to complement their use of misrepresentations, malicious issuers may use other strategies to target naïve investors. We find that celebrity endorsements and the choice of listing websites are two such strategies. Nevertheless, misrepresentations have a distinct predictive effect on ICO scam risk. To identify ICO scams, investors could use simple cross-site verification—alongside these red flags—to look for misrepresentations.

8 Partial observability of ICO scams

We account for the partial observability of ICO scams and discuss its econometric implications. Specifically, we face an inherent data limitation—our sample of ICO scams detected on the `DeadCoins` website may be incomplete. First, we discuss and address incomplete detection of ICO scams. Next, we estimate the proportion of ICOs that are scams, including those that go undetected. Finally, we discuss welfare effects from our findings.

8.1 Detection controlled estimation

To motivate our discussion, consider this scenario: (i) Unsophisticated ICO scams tend to have more misrepresentations, and (ii) such scams are more prone to detection on the `DeadCoins` website. Two econometric issues ensue. First, we may overestimate the effect of *misrep* on ICO scam risk because we cannot directly observe the sophistication of ICO scams. Second, we may underestimate the prevalence of ICO scams because we inadequately detect sophisticated scams. By reducing ICO scams, tighter regulations may improve investor welfare. However, these improvements must be balanced against the cost of regulations. Thus, the socially optimal level of regulations is a function of the prevalence of ICO scams, which we need to carefully assess.

To account for incomplete detection, we use detection controlled estimation (DCE) methods (Wang, Winton, and Yu, 2010; Comerton-Forde and Putniņš, 2014; Foley, Karlsen, and Putniņš, 2019). In our DCE model, we simultaneously estimate a system of two equations: one models ICO scams, while the other models detection conditional on the occurrence of ICO scams. Thereafter, we estimate the DCE model using the maximum likelihood method. The Internet Appendix contains full details of the DCE model and a derivation of its likelihood function.

To identify our DCE model, we require instrumental variables that are uniquely

associated with either the scam or detection stage. In selecting our instruments, we hypothesize that malicious issuers opportunistically perform ICOs during periods of strong sentiment in cryptocurrency markets to capture more funds. Operationally, we measure market sentiment with *BTC returns* (*BTC search*), which is the cumulative returns of Bitcoin (cumulative Google Trends search volume index of the word "Bitcoin") in the one month prior to ICO start dates.

Both instruments are arguably unassociated with detection probabilities for three reasons. First, to the extent that detection is idiosyncratic (i.e., ICO-specific), our Bitcoin-based measure of marketwide sentiment should be orthogonal to detection probabilities. Second, if ICO scams were primarily detected on the basis of our sentiment-timing mechanism, then we should expect detection to be quick. However, we find that several months elapse between the end date of the average ICO scam and its subsequent detection on the `DeadCoins` website. Third, we manually verify that reasons behind scam allegations on the `DeadCoins` website do not allude to sentiment-timing.

- Table 9 here -

Table 9 reports estimates from our DCE models. Estimated coefficients are expressed as odds ratios. The first two columns belong to Model A, which uses *BTC search* and *BTC returns* as instruments in the scam stage. We find in column 1 that increases in *BTC search* and *BTC returns* are positively and significantly associated with ICO scams. This pattern supports our idea that malicious issuers time their ICOs to ride on periods of strong sentiment in cryptocurrency markets. Crucially, misrepresentations continue to predict ICO scams. Column 2 shows that an ICO scam with more *misrep* is more likely to be detected, suggesting that misrepresentations also draw scrutiny from market participants. This finding is consistent with our screening mechanism in which the malicious issuer's objective is not necessarily to avoid detection.[24] In fact, the screening strategy involves the use of "tell-tale signs" (e.g., misrepresentations, celebrity endorsements) that are obvious to many people but which some naïve investors are oblivious to.

We next set up Model B, which uses *altcoin search* (i.e., Google Trends search volume index for the word "ICO") and *altcoin returns* as alternative instruments in the scam stage. These instruments are constructed similarly to our Bitcoin-based instruments, but are based on alternative coins—all cryptocurrencies excluding Bitcoin. Using *altcoin*

[24]Unlike traditional markets, participants on markets for digital assets are often pseudonymous so the reputational and regulatory costs from detection are lower.

search and *altcoin returns* as instruments in columns 3 and 4, our conclusions remain unchanged. ICOs launched during stronger sentiment in the alt-coin market are subsequently more likely to be scams.[25] In addition, we continue to find that misrepresented ICOs are more likely to be scams and detected as such.

As robustness checks, we use alternative instruments in the scam stages of Models C and D. The *app download* instrument is the log-transformed number of downloads of cryptocurrency exchange mobile applications in the month prior to the ICO start date (Auer et al., 2022). A high *app downloads* reflects a large increase in retail cryptocurrency investors, many of whom may be naïve. In the same vein, more visits to the "Initial coin offering" page on Wikipedia (i.e., higher *wikipedia search*) indicate stronger retail interest in the ICO market (Focke, Ruenzi, and Ungeheuer, 2020). We find that ICO scams are more likely to be launched during periods of high *app downloads* and high *wikipedia search*. Nevertheless, *misrep* continues to predict the incidence of ICO scams.

8.2 Welfare analysis of ICO scams

Using estimates from our DCE models, we fit the models in columns 1, 3, 5, and 7 of Table 9 to probabilistically identify ICO scams. To obtain an empirical distribution of the proportion of probable scams, we perform a stratified bootstrap (`DeadCoins` sample vs. all other ICOs) over 500 iterations. In every iteration, we re-estimate our DCE models and recompute the proportion of probable scams. Models A through D in Table 9 estimate that 38.3% ($\hat{\sigma} = 1.3\%$), 38.7% ($\hat{\sigma} = 1.2\%$), 28.7% ($\hat{\sigma} = 1.3\%$), and 20.6% ($\hat{\sigma} = 1.4\%$) of ICOs in our sample are scams, respectively. Thus, many ICO scams potentially go undetected. As a benchmark, the ICO advisory firm Satis Group estimates in an industry report that 78% of ICOs are scams (Dowlat, 2018).[26]

We discuss welfare considerations from our empirical exercise. Should policymakers be concerned about harm to ICO investors? This is an important question, to which there is no obvious answer. On one hand, the potential financial losses to ICO investors are substantial based on a back-of-envelope calculation. On average, an ICO raises U.S. $5.07 million in our sample. Suppose 40% of the 5,935 ICOs are scams. Then, ICO

[25] We calculate economic magnitudes in column 3 as follows. $\sigma(\textit{altcoin search}) = 20.93$; $20.93 \times (1.023 - 1) = 0.4814$. $\sigma(\textit{altcoin returns}) = 82.7\%$; $82.7\% \times (1.362 - 1) = 29.94\%$.

[26] The Satis Group report uses a smaller and earlier sample, a different definition of ICO scams, and a different estimation methodology.

investors may be facing a loss of U.S. \$5.07 million $\times$ 0.4 $\times$ 5,935 = U.S. \$12.03 billion. These large estimates of the prevalence of scams and the consequent financial losses could inform discussions on regulatory design for cryptocurrency markets.

On the other hand, individuals may view risky ICO investments and traditional gambling devices in the same light.[27] For example, the U.S. Census Bureau reports that state-administered lottery funds alone generated U.S. \$76.4 billion in sales in 2018. To the extent that the average skewness-loving individual substitutes between ICO investments and traditional gambling devices, the net welfare loss to her from ICO scams would be smaller. From this perspective, more choices of gambling devices offered by the multitude of ICOs on the market may even increase individual welfare.

Overall, our paper is agnostic on the net welfare effects. However, our estimated scale of ICO scams and its associated financial impact may inform cost-benefit tradeoffs of future regulatory policies.

9 Conclusions

In this paper, we analyze how malicious actors may target their victims in financial scams and fraud. Using point-in-time snapshots of data extracted from ICO listing websites, we find widespread cross-site discrepancies in ICO characteristics. The results suggest that malicious ICO issuers strategically use cross-site misrepresentations to screen for naïve investors. Astute investors conduct due diligence and immediately dismiss the ICO scam. However, naïve investors overlook these misrepresentations, fall for the scam, and eventually fund the ICO. Ultimately, the investors who remain are likely to be naïve—the ideal targets of the malicious issuer. Our evidence indicates that the use of misrepresentations is nefarious—an additional misrepresentation raises the hazard of ICO scams by 14.0%. This effect is concentrated in the misrepresentations of basic ICO characteristics that are fundamental to investors' due diligence. Using wallet information from the Ethereum blockchain, we find that cryptocurrency wallets holding tokens of misrepresented ICOs (i) have less total values, (ii) are less diversified, and (iii) are less active. These patterns support our view that malicious issuers (successfully) use misrepresentations to screen for naïve or unsophisticated investors.

[27] Anecdotal evidence from social media, such as the `Reddit` forums, supports this consideration.

We find that ICO misrepresentations are unlikely to be unintentional mistakes. First, the threat of regulatory scrutiny deters the use of misrepresentations. This finding implies that there are likely to be elements of malice and criminality in the use of misrepresentations. Second, misrepresented ICOs and their non-misrepresented counterparts do not have significantly different disclosure practices and fundraising outcomes. To the extent that issuer quality is positively correlated with these proxies, our findings are inconsistent with a quality-based explanation. Third, we use network analysis to show that misrepresentation behavior is likely to be deliberate in the ICO ecosystem. We present a simple network model that captures complementarities (e.g., learning and social norms) in misrepresentation behavior. Due to complementarities facilitated by advisors, the model predicts that ICOs with higher Katz centrality use more misrepresentations. Our empirical results support this prediction. Furthermore, we find that advisors of misrepresented ICOs are more likely to obtain subsequent advisory opportunities. The absence of penalties in the advisory labor market implies that culpable advisors have incentives, or at least fewer qualms, to promote malignant behaviors in the ICO ecosystem.

A welfare analysis of the financial losses from ICO scams in our sample shows that around 40% of ICOs are potentially scams, but most go undetected. Based on this estimate, the financial losses to ICO investors due to ICO scams could exceed U.S. $12 billion. Our estimates of the true prevalence of scams and the consequent financial losses could inform discussions on regulatory design for cryptocurrency markets. Social planners may also educate the general public on how fraud is conducted by bringing attention to red flags such as misrepresentations. Even in an environment with limited regulations and investor protection, simple and low-cost due diligence can help investors avoid scams. Specific to our setting, our analysis also highlights two important issues hindering the adoption of ICOs as a financing vehicle—(i) unreliability of self-reported ICO information and (ii) widespread scams.

Appendix A The issuer's optimal targeting strategy

To complete our analysis in Section 3, we formalize the intuition from Figure 4 by solving for the optimal targeting strategy (henceforth, OTS) of the malicious issuer under perfect information. With $\Phi^{\text{naïve}}(x) \coloneqq \int_d^\infty \phi_{d|\tau}(x \mid \text{naïve})\, dx$, and $\Phi^{\text{astute}}(x) \coloneqq \int_d^\infty \phi_{d|\tau}(x \mid \text{astute})\, dx$, we first rewrite equation (1) as follows.

$$\frac{\mathbb{E}(\Pi)}{m} = mzG \left[\underbrace{\Phi^{\text{naïve}}(\infty)}_{=1} - \Phi^{\text{naïve}}(d) \right] - m(1-z)C \left[\underbrace{\Phi^{\text{astute}}(\infty)}_{=1} - \Phi^{\text{astute}}(d) \right] \tag{A.1}$$

The issuer maximizes profits by choosing d, yielding this first-order condition.

$$\frac{\partial \mathbb{E}(\Pi)}{\partial d} = -mzG \cdot (d)'\phi(d \mid \text{naïve}) + m(1-z)C \cdot (d)'\phi(d \mid \text{astute}) = 0 \tag{A.2}$$

From equation (A.2), we express the OTS as the ratio of naïve investors targeted (true positives) to astute investors targeted (false positives). This expression aligns with the intuition in receiver operating characteristic curves, which are used to assess the quality of binary classifiers. The OTS occurs at the slope:

$$\frac{\phi(d^* \mid \text{naïve})}{\phi(d^* \mid \text{astute})} = \frac{1-z}{z} \cdot \frac{C}{G} \tag{A.3}$$

Under the OTS, equation (A.3) prescribes the optimal rate of naïve investors targeted per astute investor. This rate is a function of z, C, and G. For example, suppose the issuer believes that there are many naïve investors (high z). Then, the OTS prescribes a low rate, which translates to an aggressive targeting strategy (see, Subfigure 4b). If the issuer has an inferior technology to entertain investors' queries (high C), then the issuer optimally chooses a higher rate that is achieved by a conservative targeting strategy (see, Subfigure 4a).

In practice, however, issuers cannot observe the parameters—z, C, and G—and may form heterogeneous beliefs about them. In turn, these heterogeneous beliefs may lead to heterogeneity in misrepresentation behavior across ICOs.

Appendix B News of regulatory actions taken by U.S. authorities

Date	Title	News summary
16th Jun 2018	SEC: Fraud surrounds initial coin offerings, blockchain security notwithstanding.	SEC has a unit that monitors ICO scams.
21st Jun 2018	Members of the House will now be required to disclose bitcoin, other cryptocurrency holdings; Ethics Committee strongly encourage House members who are considering investing in an ICO to seek guidance.	Ethics Committee have taken actions to regulate House members in ICO investments.
27th Jun 2018	Facebook to accept cryptocurrency ads again; January's blanket ban is reversed, though crypto firms will have to get case-by-case approval.	Tech companies such as Facebook banned cryptocurrencies ads. Promotional efforts for cryptocurrencies have come under fire from federal and state regulators.
15th Aug 2018	Even free tokens face regulatory heat as coin offerings scrutinized; SEC punishes company that didn't sell any tokens, saying potential investors were misled about details of oil-drilling project.	The SEC punished a firm that did not sell any tokens to crack down on fraud in the market for initial coin offerings.
12th Sep 2018	SEC takes first action against hedge fund over cryptocurrency investments; In a separate case that's another first, agency penalizes brokers who ran an "ICO superstore".	The SEC fined a hedge fund manager who falsely advertised his cryptocurrency fund as the first regulated crypto-fund in the United States. Separately, the SEC also fined two men who ran a website that connects investors with initial coin offerings.
12th Sep 2018	Judge lets cryptocurrency fraud case go forward, in win for SEC; For first time a federal court weighs in on the government's jurisdiction over ICOs in a criminal case.	The SEC scored a victory in their crackdown on cryptocurrency fraud as a judge ruled that initial coin offerings are subject to U.S. securities laws.

(To be continued)

Date	Title	News summary
11th Oct 2018	SEC says stop ICOs that falsely claimed SEC approval.	SEC's complaint charges Blockvest and Ringgold with violating federal securities laws.
22nd Oct 2018	SEC suspends trading in company for making false cryptocurrency-related claims about SEC regulation and registration.	SEC suspended trading in the securities of a company for making false cryptocurrency-related claims.
16th Nov 2018	SEC settles enforcement actions over two initial coin offerings	Two startups agreed to comply with investor protection rules and offer money back to thousands of people who bought their digital tokens.
30th Nov 2018	Boxer Mayweather Jr., producer DJ Khaled agree to settle SEC crypto charges.	Celebrity endorsements of coin offerings may be illegal if the promoters fail to disclose the source and amount of their compensation.
21st May 2019	SEC obtains emergency order halting alleged diamond-related ICO Scheme targeting hundreds of investors.	SEC halted a Ponzi scheme, which was purportedly a cryptocurrency business.
5th Jun 2019	SEC challenges Canada firm's coin offering	SEC sued Kik for not providing investors with full and fair disclosure about its token and its business.

Table B.1. News of regulatory actions taken by U.S. authorities (Aug '18–Aug '19)

References

Abreu, D. and Brunnermeier, M. (2003). "Bubbles and crashes". Econometrica 71, 173–204.

Aloosh, A. and Li, J. (2019). "Direct evidence of Bitcoin wash trading". Available at SSRN 3362153.

Auer, R., Cornelli, G., Doerr, S., Frost, J., and Gambacorta, L. (2022). "Crypto trading and Bitcoin prices: Evidence from a new database of retail adoption". Bank for International Settlements Working Paper.

Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2006). "Who's who in networks. Wanted: The key player". Econometrica 74, 1403–1417.

Benedetti, H. and Kostovetsky, L. (2021). "Digital tulips? Returns to investors in initial coin offerings". Journal of Corporate Finance 66.

Bourveau, T., De George, E., Ellahie, A., and Macciocchi, D. (2021). "The role of disclosure and information intermediaries in an unregulated capital market: Evidence from initial coin offerings". Journal of Accounting Research 60, 129–167.

Brunnermeier, M. and Nagel, S. (2004). "Hedge funds and the technology bubble". Journal of Finance 59, 2013–2040.

Button, M., Lewis, C., and Tapley, J. (2009). "A better deal for fraud victims". National Fraud Authority, United Kingdom.

Cohn, J., Liu, Z., and Wardlaw, M. (2021). "Regression with skewed, non-negative outcome variables in finance". Working paper.

Comerton-Forde, C. and Putniņš, T. (2014). "Stock price manipulation: Prevalence and determinants". Review of Finance 18, 23–66.

Cong, L. W., Li, X., Tang, K., and Yang, Y. (2020). "Crypto wash trading". Available at SSRN 3530220.

Cong, L. W., Li, Y., and Wang, N. (2020). "Tokenomics: Dynamic adoption and valuation". Review of Financial Studies 00, 1–51.

Crawford, V. and Sobel, J. (1982). "Strategic information transmission". Econometrica, 1431–1451.

Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (1998). "Investor psychology and security market under- and overreactions". Journal of Finance 53, 1839–1885.

Davydiuk, T., Gupta, D., and Rosen, S. (2022). "De-crypto-ing signals in initial coin offerings: Evidence of rational token retention". Management Science, forthcoming.

Deng, X., Lee, Y. T., and Zhong, Z. (2018). "Decrypting coin winners: Disclosure quality, governance mechanism and team networks". Available at SSRN 3247741.

Dhawan, A. and Putniņš, T. (2022). "A new wolf in town? Pump-and-dump manipulation in cryptocurrency markets". Review of Finance, forthcoming.

Dimmock, S. G., Gerken, W. C., and Van Alfen, T. (2021). "Real estate shocks and financial advisor misconduct". The Journal of Finance 76, 3309–3346.

Dimmock, S. G., Farizo, J., and Gerken, W. C. (2018). "Misconduct and fraud by investment managers". Available at SSRN 3228688.

Dittmar, R. and Wu, D. A. (2019). "Initial coin offerings hyped and dehyped: An empirical examination". Available at SSRN 3259182.

Dowlat, S. (2018). "Cryptoasset market coverage initiation: Network creation".

Egan, M., Matvos, G., and Seru, A. (2019). "The market for financial adviser misconduct". Journal of Political Economy 127, 233–295.

— (2022). "When Harry fired Sally: The double standard in punishing misconduct". Journal of Political Economy 130, 000–000.

Feinstein, J. (1990). "Detection controlled estimation". Journal of Law and Economics 33, 233–276.

Focke, F., Ruenzi, S., and Ungeheuer, M. (2020). "Advertising, attention, and financial markets". Review of Financial Studies 33, 4676–4720.

Foley, S., Karlsen, J., and Putniņš, T. (2019). "Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?" Review of Financial Studies 32, 1798–1853.

Gee, J. and Button, M. (2019). "The financial cost of fraud 2019: The latest data from around the world".

Gensler, G. (2021). "Remarks before the Aspen Security Forum". Public Statement. URL: https://www.sec.gov/news/public-statement/gensler-aspen-security-forum-2021-08-03.

Gigler, F. (1994). "Self-enforcing voluntary disclosures". Journal of Accounting Research 32, 224–240.

Griffin, J., Harris, J., Shu, T., and Topaloglu, S. (2011). "Who drove and burst the tech bubble?" Journal of Finance 66, 1251–1290.

Griffin, J. and Shams, A. (2020). "Is Bitcoin really untethered?" Journal of Finance 75, 1913–1964.

Herley, C. (2012). "Why do Nigerian scammers say they are from Nigeria?" *WEIS*.

Howell, S., Niessner, M., and Yermack, D. (2020). "Initial coin offerings: Financing growth with cryptocurrency token sales". Review of Financial Studies 33, 3925–3974.

La Porta, R., Lopez-De-Silanes, F., Shleifer, A., and Vishny, R. (Feb. 2000). "Agency problems and dividend policies around the world". Journal of Finance 55, 1–33.

Li, T., Shin, D., and Wang, B. (2021). "Cryptocurrency pump-and-dump schemes". Available at SSRN 3267041.

Lyandres, E., Palazzo, B., and Rabetti, D. (2021). "ICO success and post-ICO performance". Management Science, forthcoming.

Odean, T. (1998). "Volume, volatility, price, and profit when all traders are above average". Journal of finance 53, 1887–1934.

PriceWaterhouseCoopers (2020). "6th ICO/STO Report: A strategic perspective".

SEC (2018). "The SEC has an opportunity you won't want to miss: Act now!" Press Release. URL: https://www.sec.gov/news/press-release/2018-88.

Sockin, M. and Xiong, W. (2020). "A model of cryptocurrencies".

Stocken, P. (2000). "Credibility of voluntary disclosure". RAND Journal of Economics, 359–374.

Wang, T. Y., Winton, A., and Yu, X. (2010). "Corporate fraud and business conditions: Evidence from IPOs". Journal of Finance 65, 2255–2292.

Yermack, D. (2015). "Is Bitcoin a real currency? An economic appraisal". *Handbook of digital currency*. Elsevier, 31–43.

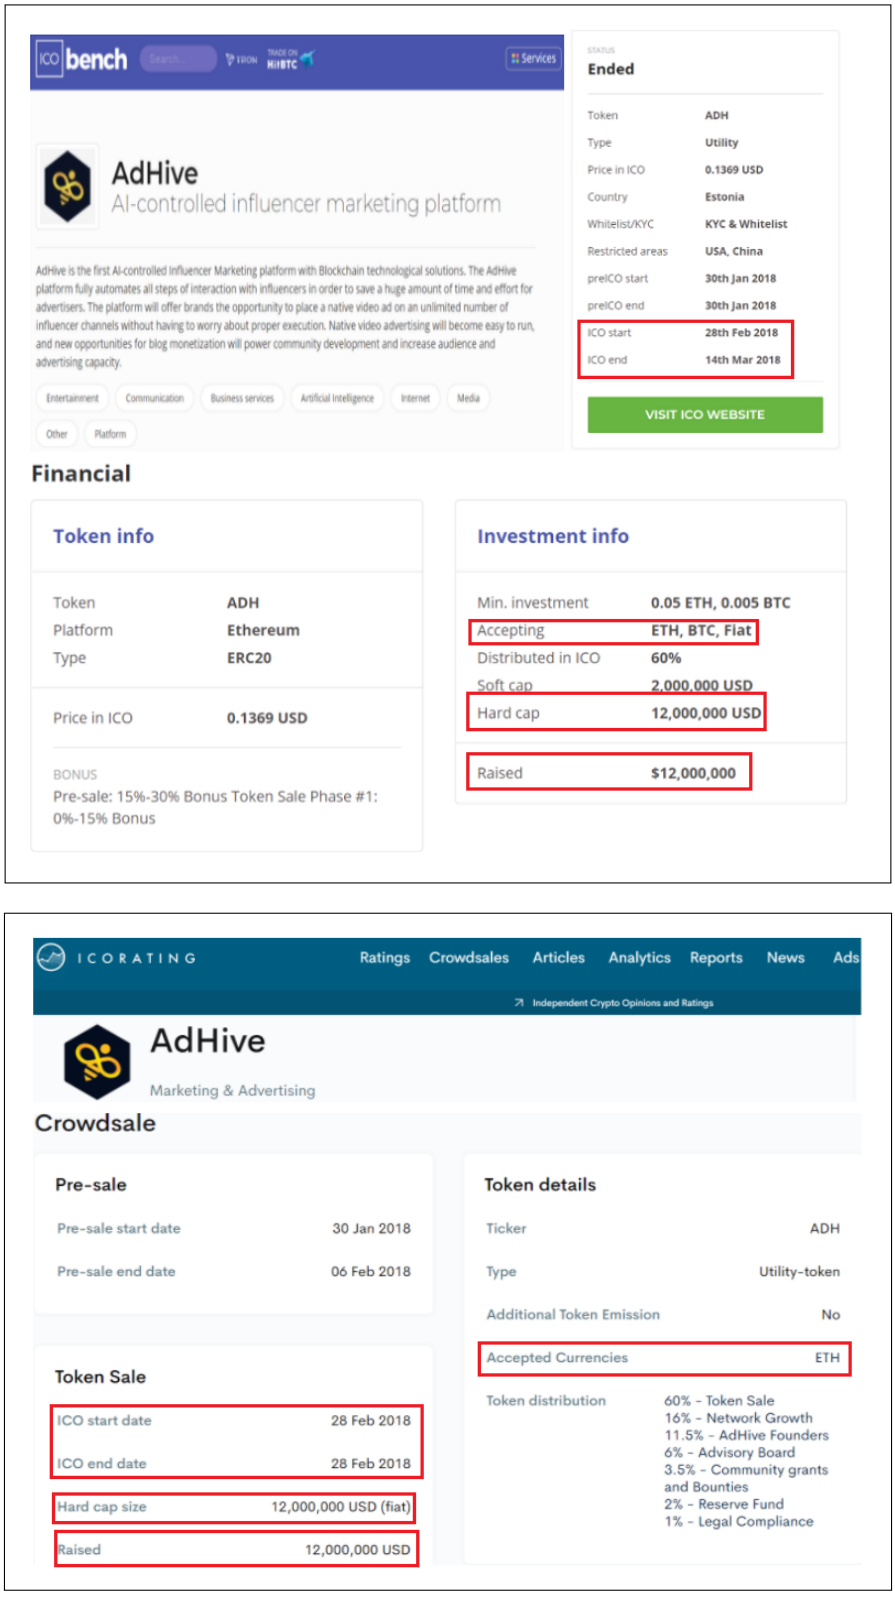

Figure 1. This figure presents screenshots of the AdHive ICO information pages on three ICO listing websites—`ICOBench.com`, `ICORating.com`, and `ICODrops.com`.

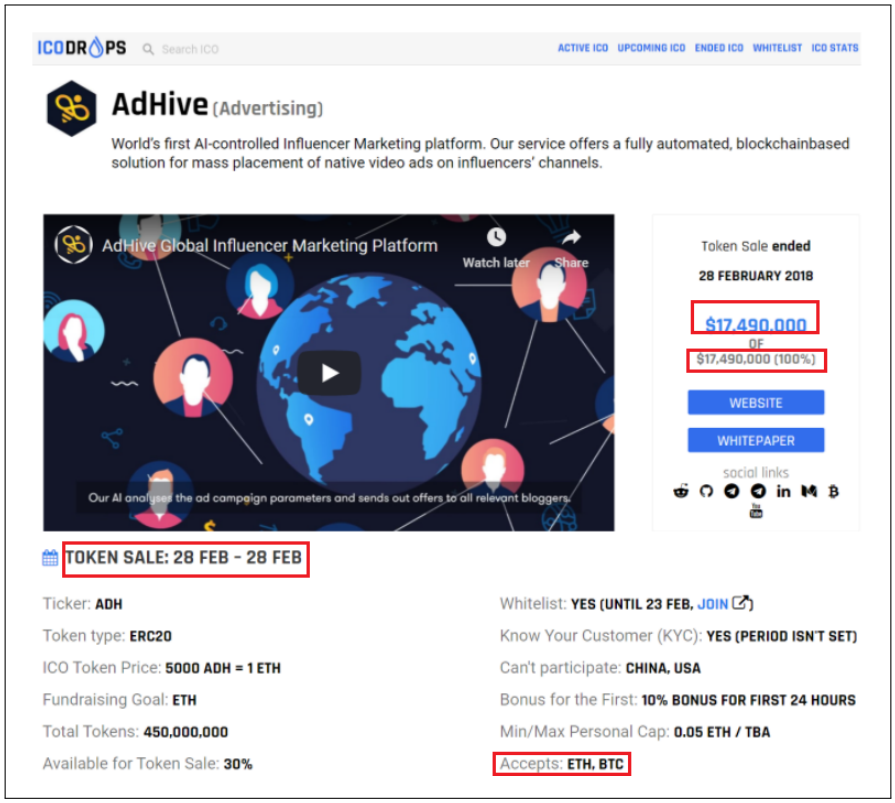

Figure 1. (continued)

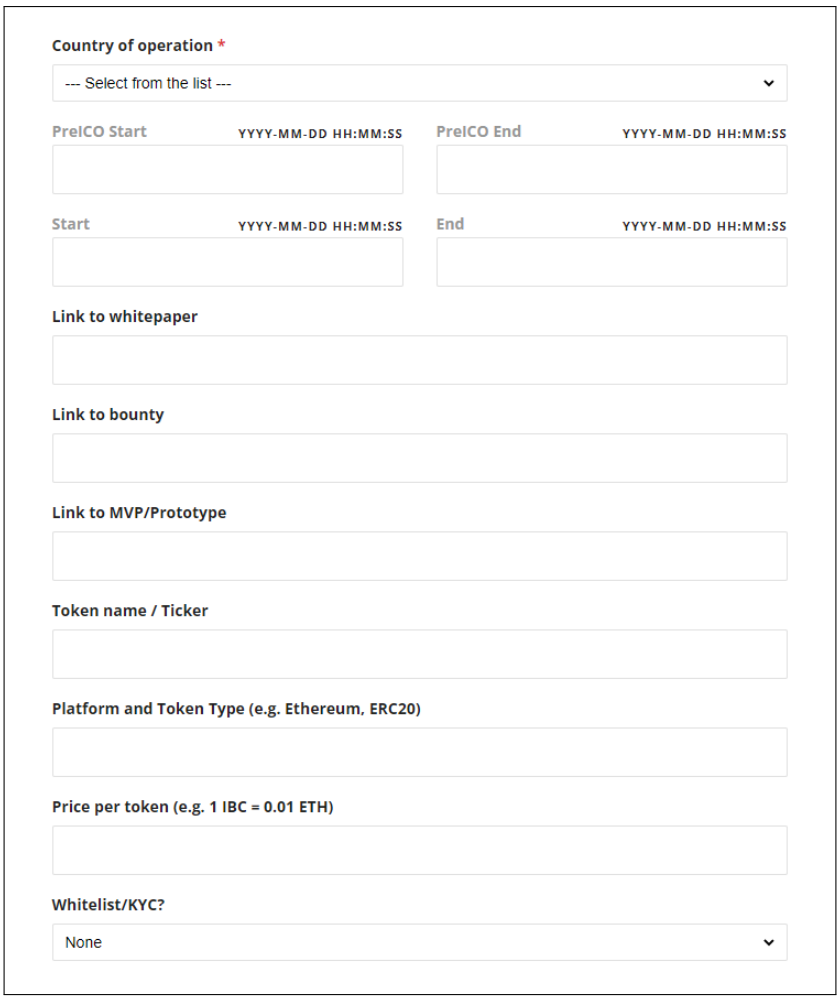

Figure 2. This figure presents a partial screenshot of the `ICOBench.com` webpage on which issuers self-report ICO data.

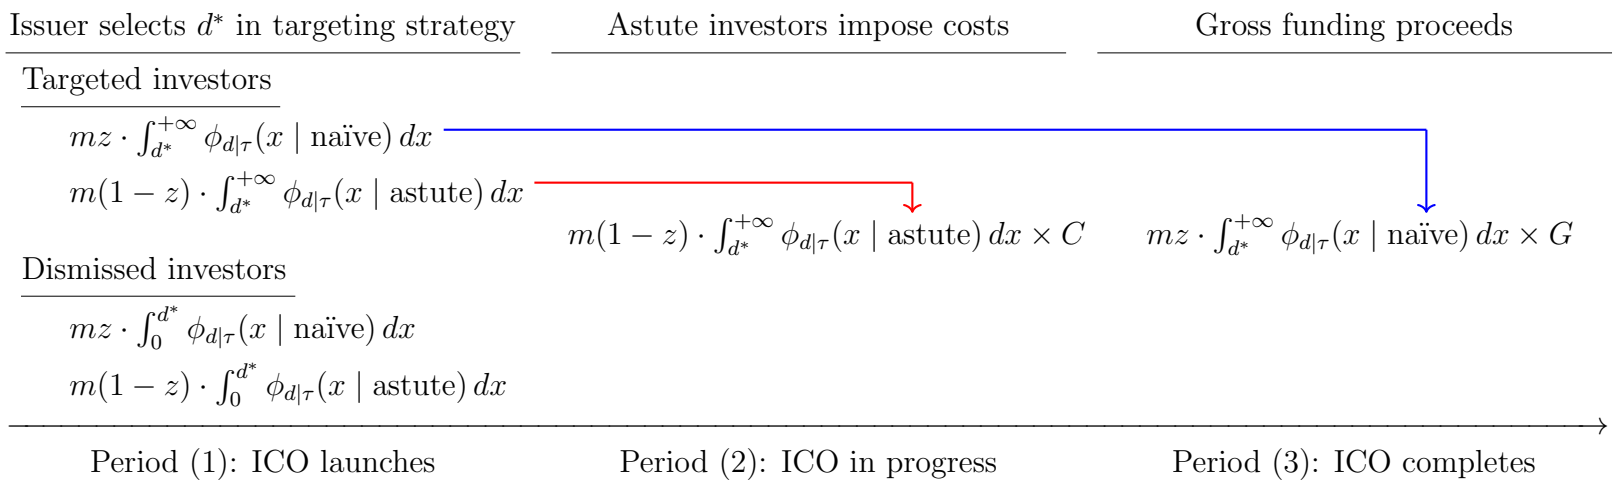

Figure 3. This figure visualizes the three periods of the model described in Section 3. The ICO launches in Period (1), and the issuer selects d^* in the targeting strategy. Some naïve and astute investors immediately dismiss the ICO. The remaining investors are targeted. In Period (2), astute investors who are targeted impose costs on the issuer by seeking additional information and asking questions on public forums. In Period (3), only naïve investors proceed to fund the completed ICO scam. Astute investors, even if targeted, ultimately refrain from funding the scam.

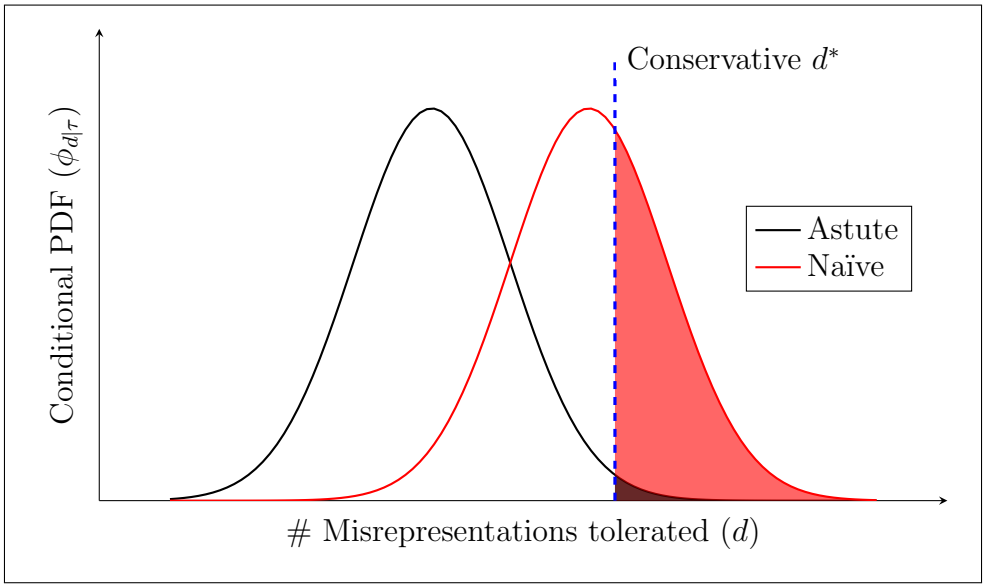

(a) Conservative targeting strategy

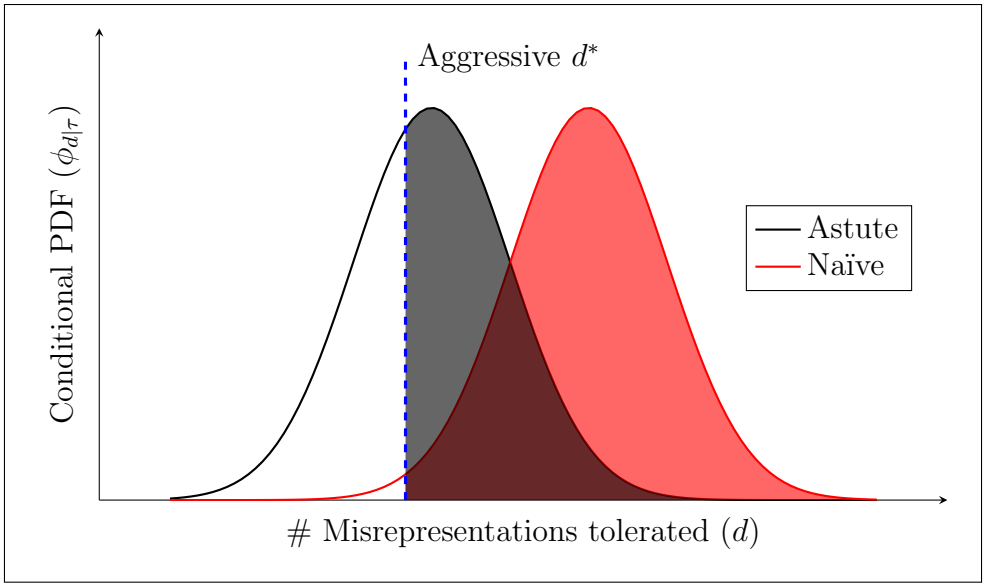

(b) Aggressive targeting strategy

Figure 4. This figure presents probability density plots of d, conditional on two investor types—astute (black) and naïve (red). Shaded areas in black and red represent the complementary conditional cumulative distributions $\bar{F}_{d|\text{type}}(d^* \mid \text{astute})$ and $\bar{F}_{d|\text{type}}(d^* \mid \text{naïve})$, respectively. Subfigures 4a and 4b visualize a conservative targeting strategy (high d^*) and an aggressive targeting strategy (low d^*), respectively.

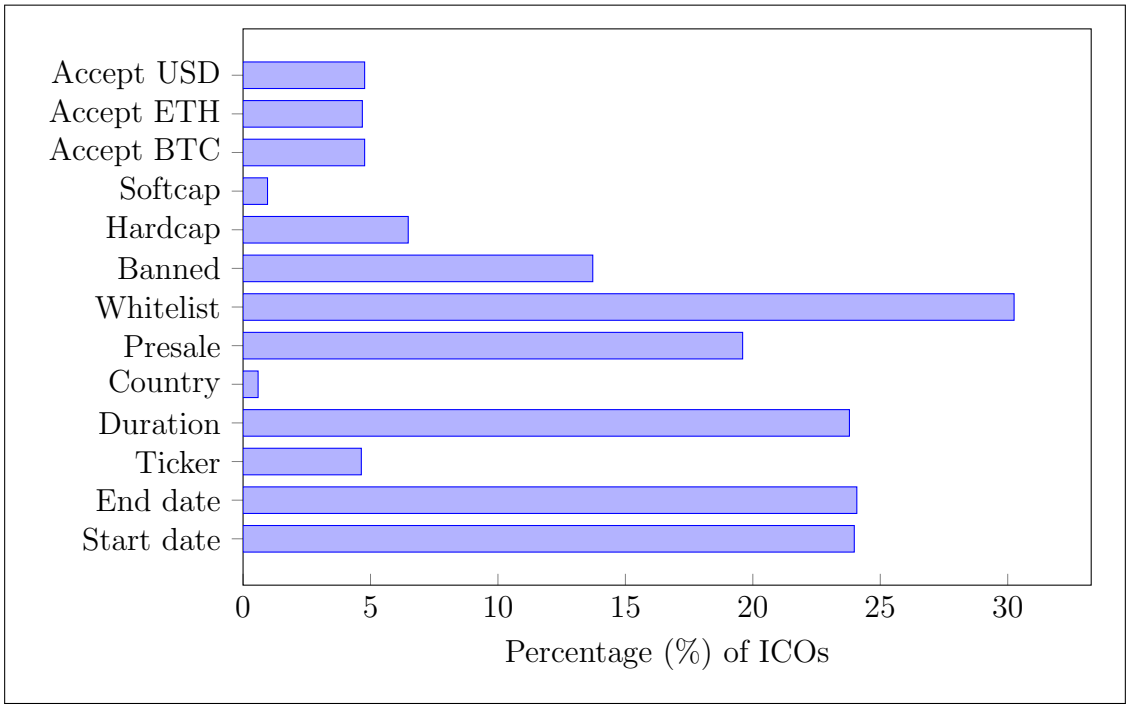

Figure 5. This figure presents the proportion of ICOs with at least one cross-website discrepancy in a particular characteristic at first appearances in our sample.

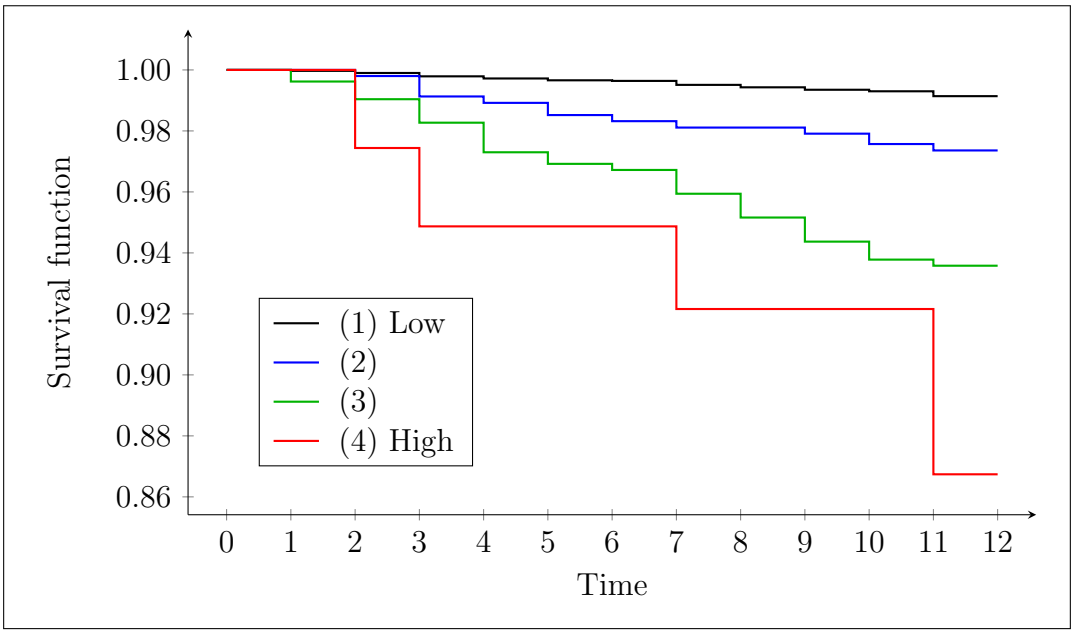

Figure 6. This figure presents the survival functions of ICOs in our sample. We assign every ICO into one of four groups based on its number of cross-website discrepancies in its characteristics at its first appearance in our sample (*misrep*). The x-axis is the time-to-event—months elapsed from the time of entry into our sample. The y-axis is the groupwise proportion of ICOs that are not identified as scams on `DeadCoin.com` (i.e., survive) at a given time.

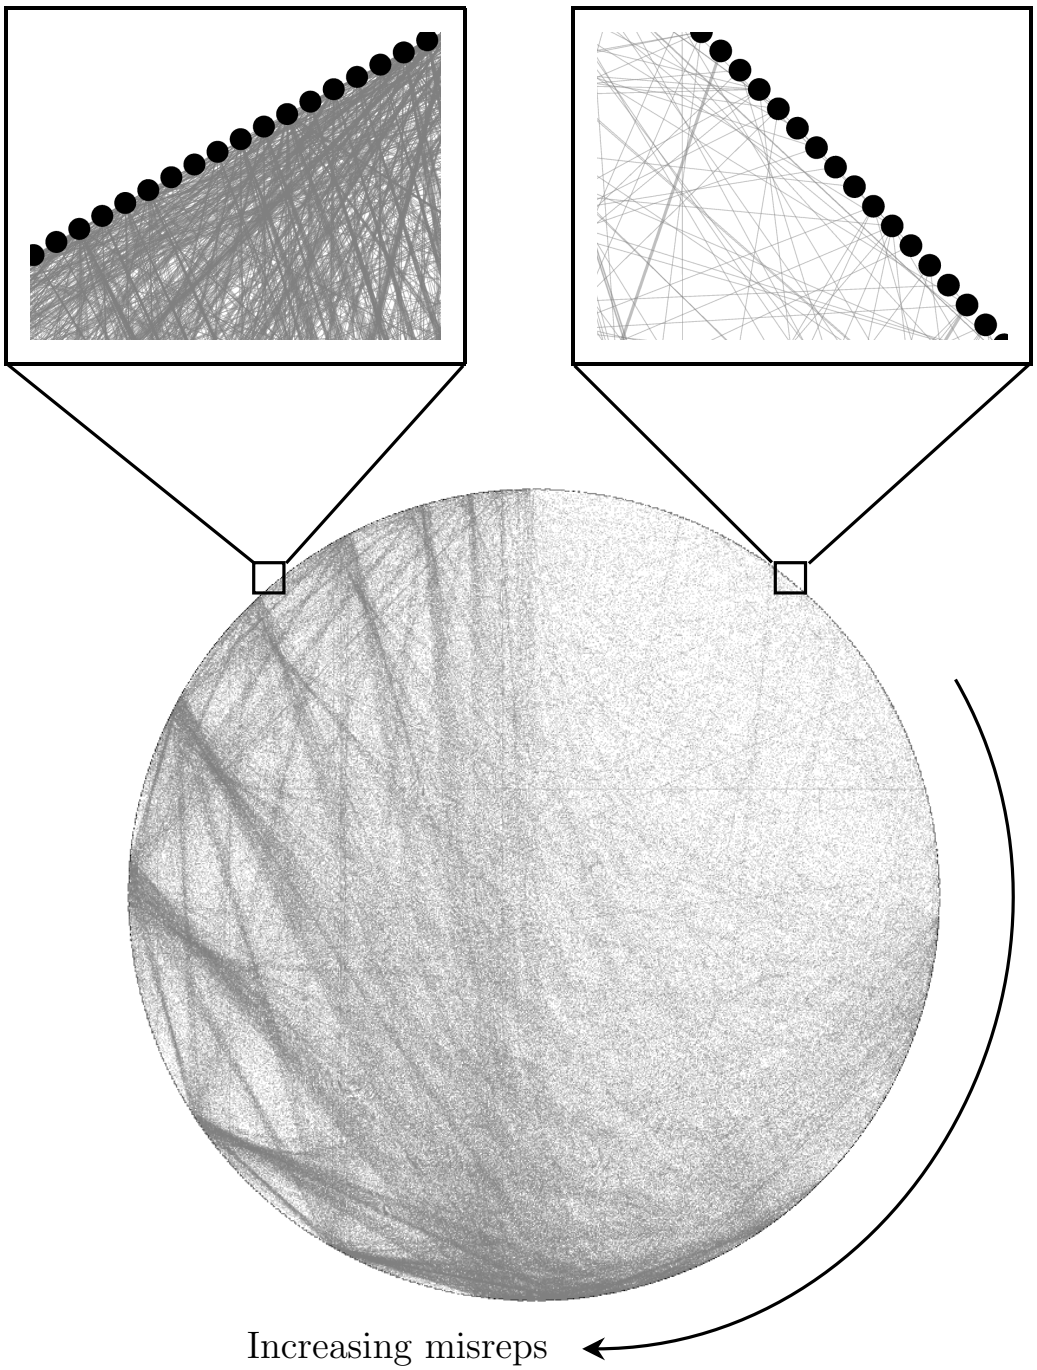

Figure 7. This figure presents a circular layout of the advisor-linked ICO network described in Section 6.1. The ICOs are arranged according to their *misrep* on the circumference of the circle. The ICO at the 12 o'clock position has the fewest *misrep*. As we move along the circumference in the clockwise direction, the ICOs have more *misrep*. Lines inside the circle represent network links between ICOs.

Table 1. Descriptive statistics

This table presents descriptive statistics of our sample at the ICO level. The variables presented in this table are extracted from the first appearance of each ICO in our 13-month observation window. Panel A reports the summary statistics of ICO characteristics and the misrepresentation measures. Panel B presents Pearson pairwise correlations between variables. Section 4.2 contains definitions of variables presented in this table.

Panel A. Summary statistics

	N	μ	σ	p10	p50	p90
Misrep	5,960	1.26	2.16	0	0	4
$\mathbb{1}_{\text{Misrep}>0}$	5,960	0.34	0.48	0	0	1
Banned	5,960	0.95	0.22	1	1	1
Whitelist	5,960	0.55	0.50	0	1	1
Presale	5,960	0.47	0.50	0	0	1
Hardcap	5,960	0.70	0.46	0	1	1
Softcap	5,960	0.26	0.44	0	0	1
Accept BTC	5,960	0.28	0.45	0	0	1
Accept ETH	5,960	0.58	0.49	0	1	1
Accept USD	5,960	0.10	0.30	0	0	0
SEC filing (%)	5,960	1.46	9.38	0	0	0
Enforcement	5,960	0.26	0.42	0	0	1
Disclosure	5,960	1.20	1.23	0	0.73	2.92
Duration (days)	5,960	54.38	50.25	15	37	109

Panel B. Pairwise correlations

		(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)	(j)	(k)	(l)
Misrep	(a)												
Banned	(b)	−0.01											
Whitelist	(c)	−0.07	0.10										
Duration	(d)	−0.12	−0.04	−0.03									
Presale	(e)	0.31	−0.01	0.17	−0.06								
Hardcap	(f)	0.28	0.02	−0.20	−0.06	0.12							
Softcap	(g)	0.03	−0.04	0.06	0.10	0.16	0.36						
Accept BTC	(h)	0.16	0.00	0.08	0.06	0.24	0.09	0.18					
Accept ETH	(i)	0.31	0.01	0.14	−0.01	0.43	0.17	0.16	0.44				
Accept USD	(j)	0.05	0.00	0.07	0.06	0.15	0.07	0.13	0.38	0.23			
SEC filing	(k)	0.04	0.00	0.02	0.01	0.04	0.02	0.01	0.04	0.03	0.05		
Enforcement	(l)	0.11	−0.02	−0.04	−0.01	0.05	0.05	0.07	−0.01	0.02	0.02	−0.03	
Disclosure	(m)	0.13	−0.11	−0.04	0.02	0.04	0.02	0.08	−0.03	−0.01	0.01	0.06	0.31

Table 2. Differences in means

This table presents differences in ICO scam rates and characteristics between misrepresented ICOs and non-misrepresented ICOs. Column (1) contains ICOs with at least one misrepresentation. Column (2) contains ICOs with no misrepresentations. We report differences in means (Δ) and their associated t-statistics. Section 4.2 contains definitions of variables presented in this table.

	(1)	(2)	$\Delta_{(1)-(2)}$	t
ICO scam	0.04	0.01	0.03	6.88
Banned	0.95	0.95	−0.01	0.90
Whitelist	0.46	0.60	−0.15	10.96
Presale	0.68	0.36	0.32	25.15
Hardcap	0.89	0.60	0.29	27.58
Softcap	0.29	0.25	0.04	3.16
Accept BTC	0.39	0.22	0.16	12.99
Accept ETH	0.80	0.46	0.34	28.82
Accept USD	0.12	0.09	0.04	4.21
SEC filing (%)	1.21	0.72	0.49	1.79
Duration (days)	47.71	57.91	−10.20	8.29
Enforcement	0.33	0.22	0.11	9.52
Disclosure	1.44	1.07	0.37	11.11

(1): ICOs with at least one misrepresentation
(2): ICOs with no misrepresentations

Table 3. Misrepresentations and ICO scams

This table presents estimates from Cox regressions. Estimated coefficients are expressed as hazard ratios. The failure event in these regressions is *ICO scam*. An ICO triggers the event if the `DeadCoin` site identifies it as a scam. Otherwise, it is right-censored. The key independent variables in our regressions are *misrep*, $\mathbb{1}(misrep > 0)$, and $misrep^{basic}$. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The indicator $\mathbb{1}(misrep > 0)$ equals one if the ICO has at least one *misrep*, and equals zero otherwise. The $misrep^{basic}$ of an ICO is the number of cross-site discrepancies of its basic characteristics at its first appearance in our sample. Section 4.2 contains variable definitions. Some models contain coverage-quartile fixed effects and are stratified by ICO cohorts. Standard errors in some models are clustered by ICO cohorts. t-statistics are reported in parentheses.

Event: ICO scam

	(1)	(2)	(3)	(4)
$\mathbb{1}$(Misrep >0)	3.740			
	(5.46)			
Misrep		1.253	1.140	
		(6.71)	(2.18)	
Misrep$^{\text{basic}}$				1.240
				(4.86)
Banned	0.992	0.984	1.015	1.015
	(0.02)	(0.04)	(0.04)	(0.04)
Whitelist	1.439	1.196	1.402	1.470
	(1.71)	(0.85)	(1.47)	(1.85)
Duration	0.999	1.000	1.000	1.000
	(0.45)	(0.18)	(0.08)	(0.00)
Presale	0.951	0.881	0.967	1.020
	(0.22)	(0.54)	(0.21)	(0.15)
Hardcap	1.709	1.653	1.619	1.625
	(1.76)	(1.62)	(1.72)	(1.93)
Softcap	0.873	0.879	0.985	0.951
	(0.61)	(0.58)	(0.12)	(0.35)
Accept BTC	1.355	1.331	1.291	1.210
	(1.36)	(1.27)	(1.14)	(0.84)
Accept ETH	1.024	1.081	1.159	1.066
	(0.10)	(0.30)	(0.61)	(0.26)
Accept USD	1.224	1.238	1.287	1.316
	(0.68)	(0.72)	(0.80)	(0.82)
Enforcement	0.635	0.643	0.625	0.603
	(1.77)	(1.72)	(1.95)	(2.11)
Disclosure	0.934	0.939	0.922	0.907
	(0.83)	(0.77)	(1.29)	(1.59)
SEC filing	0.674	0.587	0.559	0.552
	(0.39)	(0.53)	(0.78)	(0.76)
# ICOs	5,935	5,935	5,935	5,935
Cohort strata	N	N	Y	Y
Coverage-quartile FE	N	N	Y	Y
Clustered SE	N	N	Y	Y

Table 4. Assessing the screening mechanism

Panel A presents estimates from Poisson regressions. Estimated coefficients are expressed as incidence rate ratios. For every ICO, we first identify individual cryptocurrency wallets that hold its tokens. Next, we compute wallet characteristics by extracting data from the Ethereum blockchain. Finally, we aggregate wallet-level measures at the ICO level by taking medians. The dependent variables *value* (column 1), *diversity* (column 2), and *activity* (column 3). The *value* of an ICO is the median portfolio value (in U.S. dollars) of wallets that hold its tokens. The *diversity* of an ICO is the median number of distinct tokens held in wallets that hold its tokens. The *activity* of an ICO is the median number of blockchain transactions performed by wallets that hold its tokens. The key independent variable in our regressions is $\mathbb{1}(misrep > 0)$. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The indicator $\mathbb{1}(misrep > 0)$ equals one if the ICO has at least one *misrep*, and equals zero otherwise. Section 4.2 contains variable definitions. Models contain ICO cohort fixed effects. Standard errors are clustered by ICO cohorts. *t*-statistics are reported in parentheses.

Panel A: Misrepresentations and wallet characteristics

	(1)	(2)	(3)
Dependent variable:	Value	Diversity	Activity
$\mathbb{1}$(Misrep > 0)	0.399	0.803	0.910
	(2.61)	(2.88)	(2.62)
Banned	14.899	1.093	0.863
	(2.78)	(0.57)	(2.51)
Whitelist	1.080	0.995	1.076
	(0.23)	(0.04)	(1.37)
Duration	0.992	0.998	0.998
	(2.38)	(1.93)	(5.49)
Presale	0.602	0.812	0.946
	(0.80)	(1.27)	(0.85)
Hardcap	2.019	0.860	1.001
	(1.90)	(1.82)	(0.02)
Softcap	1.233	1.042	0.965
	(0.57)	(0.28)	(0.92)
Accept BTC	1.802	1.012	0.917
	(0.89)	(0.14)	(1.57)
Accept ETH	1.605	0.982	0.951
	(1.26)	(0.10)	(0.90)
Accept USD	0.325	0.802	0.793
	(0.91)	(0.73)	(2.22)
Enforcement	1.010	1.031	0.999
	(0.03)	(0.23)	(0.01)
Disclosure	1.032	0.961	0.973
	(0.22)	(1.06)	(2.34)
SEC filing	0.000	0.666	0.962
	(12.23)	(1.39)	(0.17)
# ICOs	1,996	1,996	1,996
Cohort FE	Y	Y	Y
Clustered SE	Y	Y	Y

Table 4. (continued)

Panel B presents estimates from Poisson regressions. Estimated coefficients are expressed as incidence rate ratios. The dependent variables relate to investors' activity on `Reddit` subforums (i.e., subreddits) of ICOs up until the ICO end date. The *avg. # comments per post* is the number of user comments, divided by the number of posts on the subreddit. The *avg. # questions per post* is the number of questions (identified by the presence of a question mark), divided by the number of posts on the subreddit. The *avg. # users per post* is the number of unique users who made at least one comment, divided by the number of posts on the subreddit. The key independent variable in our regressions is $\mathbb{1}(misrep > 0)$. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The indicator $\mathbb{1}(misrep > 0)$ equals one if the ICO has at least one *misrep*, and equals zero otherwise. Section 4.2 contains variable definitions. Models contain ICO cohort fixed effects. Standard errors are clustered by ICO cohorts. t-statistics are reported in parentheses.

Panel B: Misrepresentations and Reddit activity

Dependent variable:	(1)	(2)	(3)
	Avg. per post		
	# Comments	# Questions	# Users
$\mathbb{1}$(Misrep > 0)	0.512	0.534	0.800
	(4.67)	(2.95)	(1.96)
Banned	2.487	2.885	1.367
	(2.87)	(2.45)	(1.06)
Whitelist	1.142	1.402	1.037
	(0.45)	(2.20)	(0.18)
Duration	1.004	1.005	1.002
	(2.74)	(4.49)	(1.34)
Presale	0.839	0.936	0.930
	(1.64)	(0.41)	(0.59)
Hardcap	1.573	1.295	1.622
	(2.14)	(0.91)	(2.13)
Softcap	1.325	1.204	1.291
	(1.53)	(1.28)	(1.26)
Accept BTC	1.508	1.335	1.343
	(1.28)	(1.42)	(1.23)
Accept ETH	1.539	1.670	1.161
	(1.69)	(3.59)	(0.88)
Accept USD	0.325	0.355	0.648
	(2.89)	(3.39)	(1.09)
Enforcement	0.670	0.645	0.825
	(1.64)	(2.65)	(0.82)
Disclosure	1.251	1.211	1.161
	(2.03)	(2.43)	(1.70)
SEC filing	8.998	9.105	2.694
	(5.53)	(7.29)	(2.34)
log (# Posts)	0.920	0.811	1.042
	(0.83)	(4.30)	(0.33)
log (Community size)	1.347	1.251	1.273
	(4.38)	(3.97)	(3.99)
# ICOs	541	541	541
Cohort FE	Y	Y	Y
Clustered SE	Y	Y	Y

Table 5. Central ICOs and misrepresentations

This table presents estimates from Poisson regressions. Estimated coefficients are expressed as incidence rate ratios. The dependent variable is *misrep*. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The key independent variables are log (*centrality*) and $\mathbb{1}$(*high centrality*). The variable log (*centrality*) is the log-transformed Katz centrality of the ICO. The Internet Appendix contains details on the Katz centrality measure. The variable $\mathbb{1}$(*high centrality*) is an indicator that equals one if the ICO has a higher Katz centrality than the median Katz centrality in the sample, and equals zero otherwise. Section 4.2 contains variable definitions. Models contain cohort fixed effects. Standard errors are clustered by ICO cohorts. *t*-statistics are reported in parentheses.

Dependent variable: Misrep

	(1)	(2)	(3)	(4)
Weighted links	N	Y	N	Y
log (Centrality)	1.485	1.567		
	(2.27)	(2.17)		
$\mathbb{1}$(High centrality)			1.061	1.067
			(1.96)	(2.25)
Banned	0.974	0.974	0.974	0.974
	(0.48)	(0.47)	(0.45)	(0.46)
Whitelist	1.134	1.134	1.133	1.133
	(1.85)	(1.85)	(1.82)	(1.82)
Duration	0.999	0.999	0.999	0.999
	(1.56)	(1.56)	(1.56)	(1.57)
Presale	1.590	1.591	1.588	1.587
	(7.47)	(7.49)	(7.60)	(7.62)
Hardcap	1.598	1.599	1.596	1.597
	(6.75)	(6.77)	(6.98)	(6.90)
Softcap	0.996	0.996	0.996	0.997
	(0.29)	(0.26)	(0.30)	(0.22)
Accept BTC	1.065	1.065	1.067	1.067
	(1.31)	(1.31)	(1.32)	(1.35)
Accept ETH	1.249	1.249	1.245	1.243
	(2.31)	(2.31)	(2.25)	(2.26)
Accept USD	1.033	1.034	1.036	1.035
	(0.76)	(0.77)	(0.81)	(0.79)
Enforcement	1.023	1.022	1.023	1.025
	(0.73)	(0.72)	(0.74)	(0.76)
Disclosure	1.001	1.001	1.000	1.000
	(0.08)	(0.08)	(0.02)	(0.02)
SEC filing	0.947	0.946	0.942	0.944
	(0.62)	(0.62)	(0.66)	(0.63)
# ICOs	2,271	2,271	2,271	2,271
Cohort FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y

Table 6. Misrepresentations and ICO quality

This table present estimates from logit (columns 1 and 2) and Poisson (column 3) regressions. Estimated coefficients are expressed as odds ratios (incidence rate ratios) in columns 1 and 2 (column 3). The dependent variables are $\mathbb{1}$(*code posted*), $\mathbb{1}$(*code audited*), and *raised*. The indicator $\mathbb{1}$(*code posted*) equals one if the ICO posts the source code of its smart contract on `Etherscan.io` and equals zero otherwise. The indicator $\mathbb{1}$(*code audited*) equals one if the ICO posts a security audit of its source code on `Etherscan.io` and equals zero otherwise. The variable *raised* is the amount of capital (in U.S. dollars) raised by the ICO. The key independent variables in our regressions is *misrep*. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. Section 4.2 contains variable definitions. Models contain cohort fixed effects. The sample sizes here are smaller than those in Table 3 because of data limitations. Standard errors are clustered by ICO cohorts. t-statistics are reported in parentheses.

	(1)	(2)	(3)
Dependent variable:	$\mathbb{1}$(Code posted)	$\mathbb{1}$(Code audited)	Raised
Misrep	0.984	1.011	1.058
	(0.31)	(0.26)	(1.04)
Banned	1.419	0.940	0.948
	(0.74)	(0.19)	(0.25)
Whitelist	0.942	0.953	2.300
	(0.48)	(0.17)	(4.72)
Duration	0.998	0.996	1.004
	(1.07)	(1.45)	(0.56)
Presale	0.988	0.790	0.748
	(0.08)	(0.82)	(1.53)
Hardcap	1.313	1.664	0.891
	(2.64)	(2.53)	(0.35)
Softcap	0.853	0.865	0.800
	(1.50)	(0.71)	(1.20)
Accept BTC	1.200	1.312	0.816
	(0.89)	(1.30)	(0.85)
Accept ETH	1.035	1.265	1.625
	(0.29)	(0.92)	(1.66)
Accept USD	0.811	0.969	1.594
	(0.78)	(0.10)	(1.32)
Enforcement	1.062	0.847	0.734
	(0.55)	(0.76)	(2.52)
Disclosure	1.110	1.130	0.980
	(2.47)	(1.77)	(0.24)
SEC filing	0.299	1.00	1.182
	(1.40)	(0.00)	(0.53)
# ICOs	4,604	4,604	2,985
Cohort FE	Y	Y	Y
Clustered SE	Y	Y	Y

Table 7. Regulatory scrutiny and misrepresentation behavior

Columns 1 and 3 (2) of this table present estimates from logistic (Poisson) regressions. Estimated coefficients in columns 1 and 3 (2) are expressed as odds (incidence rate) ratios. The dependent variable in column 1 is $\mathbb{1}(misrep > 0)$—an indicator that equals one if the ICO has at least one cross-site discrepancies of its characteristics at its first appearance in our sample, and equals zero otherwise. The dependent variable in column 2 is *misrep*. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The dependent variable in column 3 is $\mathbb{1}(\Delta misrep < 0)$—an indicator that equals one if the ICO has a reduction in cross-site discrepancies from the previous month, and equals zero otherwise. The key independent variable is *regulatory scrutiny*—the number of regulatory news articles released within the prior calendar month. Section 4.2 contains variable definitions. Models contain ICO cohort fixed effects. Standard errors are clustered by ICO cohorts. t-statistics are reported in parentheses.

	(1)	(2)	(3)
Dependent variable:	$\mathbb{1}(\text{Misrep} > 0)$	Misrep	$\mathbb{1}(\Delta\text{Misrep} < 0)$
Regulatory scrutiny	0.795	0.838	0.964
	(2.13)	(2.91)	(0.89)
Banned	0.772	0.926	1.379
	(1.41)	(1.58)	(3.41)
Whitelist	0.506	0.938	0.608
	(4.47)	(1.53)	(5.41)
Duration	0.998	0.998	0.999
	(2.25)	(3.30)	(1.11)
Presale	4.277	2.432	4.736
	(8.71)	(9.03)	(5.43)
Hardcap	4.479	3.253	2.225
	(10.19)	(22.11)	(4.59)
Softcap	0.818	0.993	0.867
	(1.78)	(0.21)	(3.51)
Accept BTC	1.270	1.141	1.039
	(2.29)	(3.45)	(0.53)
Accept ETH	4.867	2.447	0.433
	(6.29)	(8.83)	(4.46)
Accept USD	0.882	0.997	1.053
	(1.23)	(0.10)	(0.32)
Enforcement	1.538	1.161	1.214
	(3.26)	(4.26)	(1.40)
Disclosure	1.373	1.147	1.141
	(4.19)	(6.76)	(3.23)
SEC filing	1.014	0.995	1.607
	(0.04)	(0.04)	(1.28)
Unit of observation	ICO	ICO	ICO-month
# observations	5,935	5,935	56,991
Cohort FE	Y	Y	Y
Clustered SE	Y	Y	Y

Table 8. Other suspicious actions

This table presents estimates from Cox regressions. Estimated coefficients are expressed as hazard ratios. The failure event in these regressions is *ICO scam*. An ICO triggers the event if the `DeadCoin` site identifies it as a scam. Otherwise, it is right-censored. The key independent variables in our regressions are $\mathbb{1}(\textit{celebrity})$, *web traffic ratio*, and *misrep*. The indicator $\mathbb{1}(\textit{celebrity})$ equals one if an ICO is endorsed by a celebrity, and equals zero otherwise. To compute *web traffic ratio* of an ICO, we first classify web traffic to listing websites into two categories—passive and active. Passive web traffic counts visitors referred to a listing website via third-party referral links, paid advertisements, and search engines. Active web traffic counts visitors who access a listing website by directly typing its Uniform Resource Locator (URL) or through the use of saved browser bookmarks. Next, we define the *web traffic ratio* of an ICO as the ratio of passive traffic to active traffic, aggregated across the listing websites that list it in the month prior to its start date. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. Section 4.2 contains variable definitions. Models contain coverage-quartile fixed effects and are stratified by ICO cohorts. Standard errors are clustered by ICO cohorts. t-statistics are reported in parentheses.

Event: ICO scam

	(1)	(2)	(3)	(4)
$\mathbb{1}$(Celebrity)	25.780	27.027		
	(10.64)	(9.37)		
Web traffic ratio			1.265	1.254
			(2.23)	(2.07)
Misrep		1.145		1.136
		(2.04)		(2.12)
Controls	Y	Y	Y	Y
# ICOs	5,935	5,935	5,935	5,935
Cohort strata	Y	Y	Y	Y
Coverage-quartile FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y

Table 9. Partial observability of ICO scams

This table presents estimates from detection controlled estimation (DCE) models, which are implemented as bivariate probit models. The Internet Appendix contains details of the DCE framework. We simultaneously model the scam and detection processes of ICO scams. The *misrep* of an ICO is the total number of cross-site discrepancies of its characteristics at its first appearance in our sample. The instruments for the scam processes are *BTC search*, *altcoin search*, *BTC returns*, *altcoin returns*, *app downloads*, and *wikipedia search*. The variable *BTC search* (*altcoin search*) is the cumulative search volume index of the word "Bitcoin" ("ICO") on Google Trends in the one month prior to the ICO start date. The variable *BTC returns* (*altcoin returns*) is the cumulative returns of Bitcoin (non-Bitcoin cryptocurrencies) in the one month prior to the ICO start date. The variable *app downloads* is the log-transformed number of downloads of cryptocurrency exchange mobile applications in the month prior to the ICO start date. The variable *wikipedia search* is the log-transformed number of visits to the "Initial coin offering" page on Wikipedia in the one month prior to the ICO start date. Section 4.2 contains variable definitions. *t*-statistics are reported in parentheses.

Detection controlled estimation (DCE)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Model A		Model B		Model C		Model D	
	Scam	Detect	Scam	Detect	Scam	Detect	Scam	Detect
Misrep	0.144	0.100	0.136	0.111	0.150	0.069	0.147	0.017
	(7.75)	(6.00)	(6.99)	(6.78)	(8.28)	(3.60)	(8.25)	(0.65)
BTC search	0.027							
	(4.80)							
BTC returns	0.794							
	(4.48)							
Altcoin search			0.065					
			(3.87)					
Altcoin returns			0.416					
			(3.85)					
App downloads					0.115			
					(3.80)			
Wikipedia search							5.555	
							(3.58)	
Controls	Y	Y	Y	Y	Y	Y	Y	Y
# ICOs	5,935	5,935	5,935	5,935	5,935	5,935	5,935	5,935