TRUST, BUT VERIFY: THE ECONOMICS OF SCAMS IN INITIAL COIN OFFERINGS^{*}

Kenny Phua

Bo Sang

Chishen Wei

Gloria Yang Yu

Abstract

Losses from frauds 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 less sophisticated investors are more likely to invest in misrepresented ICOs. We estimate that 40% of ICOs (U.S. \$12 billion) in our sample are scams. Overall, our findings uncover how screening strategies are used in financial scams and reinforce the importance of conducting due diligence.

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

*Wei (corresponding author): cswei@smu.edu.sg, Singapore Management University. Phua: kenny.phua@uts.edu.au, University of Technology Sydney. Sang: bo.sang.2017@pbs.smu.edu.sg, Singapore Management University. Yu: gloriayu@smu.edu.sg, Singapore Management University. We are grateful for insightful comments from Mykola Babiak, Stephen Dimmock, Hogyu Jhang, Tālis Putniņš, Kanis Saengchote, and seminar participants at Global AI Finance Research Conference, Hong Kong Polytechnic University, Massey University, Nanyang Technological University, Singapore Management University, University of Melbourne, 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.

TRUST, BUT VERIFY: THE ECONOMICS OF SCAMS IN INITIAL COIN OFFERINGS

Abstract

Losses from frauds 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 less sophisticated investors are more likely to invest in misrepresented ICOs. We estimate that 40% of ICOs (U.S. \$12 billion) in our sample are scams. Overall, our findings uncover how screening strategies are used in financial scams and reinforce the importance of conducting due diligence.

Keywords: Financial scams, Screening, Cryptocurrencies, Advisors 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 ICO¹

1 Introduction

Frauds and financial scams are estimated to exceed U.S. \$5 trillion annually. There are also significant psychic costs as victims often suffer depression, grief, shame, and suicidal thoughts.² While these welfare losses are substantial, empirical evidence on financial scams is relatively scarce. Data is often difficult to obtain because scammers work to evade detection and victims are often reluctant to step forward. To study the economics of scams, we exploit a unique setting from the market for initial coin offerings (ICOs) of cryptocurrencies. An ICO is a form of crowdfunding for a blockchain/cryptocurrency project. The ICO market has grown rapidly with almost no investor protection rules and mostly voluntary, unverified disclosures. ICOs have also become notorious for scams and frauds (Howell, Niessner, and Yermack, 2020). Investors' enthusiasm for ICOs, however, has not waned. ICOs have continued to successfully raise capital, with an estimated U.S. \$50 billion dollars raised through 2020 (PriceWaterhouseCoopers, 2020).

This paper investigates scams in the ICO market and demonstrates how malicious issuers target naïve investors. To perform our analysis, we collect 13 months of point-in-time snapshots of self-reported ICO data from five leading ICO listing websites. Listing websites are aggregators of past, current, and upcoming ICOs for prospective investors and are distinct from cryptocurrency exchanges. Manual collection is necessary for two important reasons. First, ICO data have no centralized repository and are scattered across listing websites. Second, the self-reported data are not reliable (Lyandres, Palazzo, and Rabetti, 2021) and can change over time. To analyze how an ICO was initially marketed to investors, we require data at the point-in-time in which the ICO occurred. For example, Figure 1 shows snapshots of the AdHive ICO on three websites. The ICODrops website reported a hardcap amount of \$17,490,000, but ICOBench and ICORating reported amounts of \$12,000,000. Dis-

¹In 2018, the U.S. Securities and Exchange Commission (SEC) created a mock ICO on HoweyCoins.com to educate the general public on the prevalence of scams and to urge investors to perform due diligence before investing in cryptocurrencies.

²See for example: Gee and Button (2019) and Button, Lewis, and Tapley (2009)

crepancies also occur in ICO ending dates, accepted payment methods, and total amounts raised.

- Figure 1 here -

The case of the AdHive ICO is not special. In our sample of 5,935 ICOs, 34% of tokens have discrepancies at their first appearances. The prevalence of these discrepancies is some-what puzzling. Why would so many issuers—who supposedly have the technical expertise to launch sophisticated blockchain projects—fail to accurately report ICO information on listing websites? A discrepancy implies that the issuer has misrepresented the offering because at least one of the reported material facts must be untrue. Investors may develop different perceptions and opinions of the offering depending on which website they happen to visit. Ideally, we would verify the accuracy of ICO information on listing websites against legal records. But, ICOs often avoid regulatory purview by sidestepping standard filing requirements for securities. Fewer than 1% of ICOs in our sample register with the SEC, although most are likely to qualify as securities under the Howey test (Gensler, 2021).

To better understand why misrepresentations are so prevalent, we model the behavior of a malicious ICO issuer who faces a pool of naïve and astute investors. Naïve investors are unsophisticated. They are unable to conduct proper due diligence and are likely to fall for an ICO scam. In contrast, astute investors carefully evaluate the offering and eventually refrain from funding it. Both naïve and astute investors may consume the issuer's time and resources by asking for more information or raising questions on public forums. From the issuer's perspective, astute investors are undesirable targets because they ultimately do not fund the scam. Indeed, a common tactic used by online vigilantes to disrupt techsupport scammers is to pose as victims and hold tedious, unfruitful conversations. In a recent interview, Kitboga (alias) said, "[...] *important for everyone to know* [...] *how much these scammers hate when you ask questions*".³ The former SEC chairman Jay Clayton also encouraged prospective investors to ask questions to ICO issuers (SEC, 2018). Because investor types are unobservable ex-ante, the issuer would optimally wish to screen out astute investors.

We hypothesize that malicious ICO issuers use misrepresentations, along with other suspicious actions, to screen out astute investors and target naïve investors. Astute investors will notice the cross-site discrepancies and immediately dismiss the offering without consuming the issuer's time and resources. In contrast, naïve investors overlook the cross-site

 $^{^{3}}$ Source: https://www.newsweek.com/laughter-death-threats-meet-kitboga-youtuber-exposing-tech-support-scams-938384

discrepancies and remain viable victims of the ICO scam. Ultimately, the investors who remain are likely to be naïve investors—the ideal targets of the malicious issuer. For centuries, this screening strategy has been observed in various advanced fee financial scams.⁴ A modern example is the Nigerian Prince email hoax, which solicits potential victims to send money to a fictitious Nigerian Prince in exchange for a large fortune. This incredible narrative is crafted to repel discerning individuals and target the most gullible victims (Herley, 2012). Using decentralized immutable blockchain data, we shed new light on this classic financial fraud scheme.

Our main finding is that ICOs with misrepresentations are significantly more likely to be scams. To identify ICO scams, we collect crowdsourced scam events from DeadCoin.com and corroborate these records with reports from news articles, message boards, and regulatory authorities. Estimates from our hazard regressions reveal that the presence of at least one misrepresentation more than triples the odds of a ICO scam. At the intensive margin, an additional misrepresentation raises the odds of a scam by 14.0%. To sharpen our analysis, we focus on misrepresentations of basic characteristics (i.e., ticker, start/end dates, duration of fundraising, country of origin, countries from which investors are banned, and acceptable payment modes). Such misrepresentations are a potent screen for investor naïvety because these characteristics are fundamental in performing basic due diligence and do not require investor expertise. Consistent with this view, we find that the odds of a scam increase by 24.0% per unit of such misrepresentations.

To assess our screening mechanism more carefully, we extract data from the Ethereum blockchain. First, we find the Ethereum block height corresponding to 10 days after the end date of every ICO. Next, we gather data on token holdings and transaction activities from wallets that hold its tokens as at that block height. Using these data, we characterize the sophistication of the typical token holders in ICOs and test whether misrepresentations are associated with lower investor sophistication. Consistent with this view, wallets that hold tokens of misrepresented ICOs (i) have lower portfolio values, (ii) are less diversified, and (iii) are less active. Overall, these findings lend further credence to our interpretation that malicious issuers use misrepresentations to screen for naïve investors.

An alternative interpretation of our findings is that misrepresentations are unintentional mistakes made by careless issuers. We design three sets of tests to evaluate this potential explanation. First, if the underlying motives are nefarious, we expect regulatory scrutiny

⁴For example, Eugène François Vidocq, a French private investigator, detailed in his 1832 memoirs a scam known as the "letters of Jerusalem". The scammer typically solicits the victims' (financial) help to recover fictitious treasures.

to reduce the use of misrepresentations. Consistent with this prediction, ICOs launched shortly after news of regulatory action in the cryptocurrency markets have significantly fewer misrepresentations. Second, misrepresentations may occur when low quality issuers fail to accurately market their listing. To the extent that these issuers have lower quality blockchain projects, misrepresentations should be negatively associated with ICO quality. However, using disclosure practices (Bourveau et al., 2021) and fundraising outcomes as proxies for ICO quality, we find no differences in quality between misrepresented and nonmisrepresented ICOs.

Third, we apply network analysis to detect suspicious patterns of misrepresentation behavior among ICO issuers. Short of conducting interviews with scammers, we cannot directly observe the true intentions behind misrepresentation behavior. But, we can examine whether the strategic use of misrepresentations leaves suspicious footprints throughout the ICO ecosystem. For this analysis, we exploit the prevalence of ICO advisors who are hired by issuers to launch token offerings. These advisors often work on multiple ICOs. If misrepresentation behavior is learned or passed through common advisors, the network position of an ICO should be related to its misrepresentation behavior.⁵ Consistent with this prediction, we find that ICOs with higher Katz centrality in the network have more misrepresentations. Surprisingly, we find that advisors of misrepresented ICOs are not penalized, but obtain more subsequent advisory opportunities. This finding suggests that there exists many malicious issuers who solicit the services of such advisors. Overall, our evidence indicates that misrepresentation behavior is systemic within the ICO ecosystem.

To complement their use of misrepresentations, malicious issuers may conduct other suspicious actions to target naïve investors. First, such issuers may use celebrity endorsements to attract less sophisticated investors. Consistent with the warnings of the SEC, we find that celebrity endorsements are strongly associated with ICO scam risk. Second, we conjecture that passive web traffic arising from paid advertisements, referral links, and search engines reflects visits from less sophisticated individuals. Using data on web traffic flows in our sample period, we find that malicious issuers prefer to promote their ICOs on listing websites with higher passive web traffic. These findings suggest that malicious ICO issuers use a variety of tactics to attract naïve investors to their offerings. Nevertheless, we find that misrepresentations retain a distinct predictive effect on ICO scam risk.

Finally, we perform a welfare analysis of the financial losses from ICO scams in our

⁵Ballester, 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 level of this behavior.

sample. A key challenge in identifying financial scams is the reluctance of victims to report losses. Thus, many scams go unreported and undetected. To overcome this partial observability problem, we use detection-controlled estimation (DCE) methods (Feinstein, 1990) and estimate that the total financial losses exceed U.S. \$12 billion in our sample. As many as 40% of ICOs in our sample may be scams, but most go undetected. These large estimates imply that more stringent regulations and stronger enforcement actions may be justified to protect investor welfare.

Our paper contributes to a recent literature 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).⁶ A distinguishing feature of our study is the focus on the initial offering stage. While suspicions of ICO scams are widespread, evidence to date is relatively scarce. Using point-in-time data, we provide evidence on how unscrupulous actors target naïve investors and estimate the size of scams in the crypotcurrency market.

Our study also builds 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 on likelihood of a scam. Thus, our findings add a new perspective to existing studies that analyze the determinants of ICO success (Benedetti and Kostovetsky, 2021; Deng, Lee, and Zhong, 2018; Dittmar and Wu, 2019; Howell, Niessner, and Yermack, 2020). 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 ICO overview

An ICO allows entrepreneurs to raise capital via cryptographically secured tokens. Typically, an issuer resorts to an ICO when other sources of capital (e.g., venture capital and private equity) are prohibitively expensive or inaccessible. Thus, an ICO is a risky crowdfunding operation, in which the issuer sells tokens that will serve as the payment medium

⁶Aloosh and Li (2019) exploit individual accounts on the Mt. Gox crypto-exchange for direct evidence. Cong et al. (2020) applies Benford's Law to identify wash trading patterns for 29 exchanges.

for the products or services of the start-up. There are several stages in the ICO process. First, the issuer creates fundraising campaign materials. Next, the issuer sets pricing terms and markets the offering on listing websites. Finally, if the financing goals of the ICO are met, the issuer then creates and distributes tokens to the investors.

2.1 Fundraising campaign: Listing websites

The fundraising campaign entails (i) producing a whitepaper, (ii) hosting a website to provide additional information, (iii) maintaining an active social media presence, and (iv) listing the token on ICO listing websites. A whitepaper describes the goals, objectives, and development milestones of the project. But, whitepapers often lack details of business operations and rarely contain financial disclosures.

To list an ICO on a listing website, the issuer directly submits token information on the website and awaits approval. Listings are typically free, but for an additional fee, the website can feature and promote the ICO. The issuer may also hire advisors to advertise and market the ICO. These advisors usually have technical or marketing expertise, and may alleviate information asymmetry between the issuer and potential investors. However, celebrities with little or no blockchain expertise are also employed as advisors to promote the ICO. The SEC has warned that celebrity endorsements are often associated with ICO scams.

2.2 ICO pricing and listing on secondary markets

The pricing structure of ICOs are often opaque. On listing websites, issuers advertise a subscription price to the general public. But, many ICOs invite privileged investors to an earlier presale offering. While details on the presale pricing structure are not publicly available, Fahlenbrach and Frattaroli (2020) find that presales offer a significant discount to the subsequent public offering price. Presale funding rounds are controversial. They may signal strong demand from informed investors, but are also used to manipulate the sentiments of the general public. The SEC has also warned that presales are often associated with ICO scams.

The issuer may set funding goals in the ICO. The softcap is the minimum amount of funds raised to continue the project. An issuer may also specify a hardcap, which is the maximum number of tokens that can be sold in the ICO. The hardcap limits the amount of funds that can be raised in the ICO. If the softcap is met and the project is successful, the issuer will create and distribute the tokens to investors. Subsequently, investors may trade the tokens in the secondary market or use the tokens for its utility (e.g., access products or services funded by the ICO). Investors tend to have short holding periods and flip the tokens on cryptocurrency exchanges (Fahlenbrach and Frattaroli, 2020).

2.3 Regulatory environment

The ICO regulatory environment differs across countries. Some countries impose outright bans on ICOs (e.g., China and South Korea), while other countries adopt regulatory guidelines (e.g., Australia and the United States). The SEC of the United States uses the Howey Test framework to determine whether a digital asset qualifies as a security.⁷ Specifically, a digital asset is a security if (i) there is an investment of money and (ii) expectation of profits; (iii) the investment of money is in a common enterprise; and (iv) any profit comes from the efforts of a promoter or third party. The SEC Chairman Gary Gensler and his predecessor Jay Clayton believe most ICOs pass the Howey Test and are hence subject to U.S. securities laws.

Issuers of security tokens can register with SEC via form S-1 or apply for registration exemptions. Although most ICOs should arguably be classified as security offerings, fewer than 100 tokens in our sample are registered with the SEC potentially due to the high compliance costs. For exemptions, regulation D applies if funds are raised from only accredited investors; Regulation A and A+ apply if funds are raised from a broader set of investors but the offering is less than \$50 million; and issuers can also make token sales under Regulation Crowdfunding.

2.4 Are misrepresentations a violation of securities law?

ICOs classified as security offerings are subject to the Rule 10b-5, which specifies the conditions for securities fraud as follows⁸:

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

⁷See, https://www.sec.gov/corpfin/framework-investment-contract-analysis-digital-assets ⁸Rule 10b-5 is issued by the SEC under section 10(b) of the Securities Exchange Act of 1934.

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 are necessarily an untrue statement of material fact (violation of part (b)) because at least one of the reported characteristics is false. Moreover, such untrue statements can potentially mislead investors. If misrepresentation are purposely used to commit fraud or deceit, then they would also violate part (c) of the rule. As of January 2021, the SEC has taken regulatory actions against 68 ICOs and cryptocurrency offerings. The judgments from these regulatory actions totaled U.S. \$99.8 million, of which U.S. \$88.9 million were refunds and U.S. \$10.9 million were penalties. Additionally, 20 securities class action lawsuits have been filed against ICO issuers. However, websites that aggregate ICO information voluntarily reported by issuers have minimal disclosure requirements and are lightly regulated.

3 Main hypothesis

We develop a model to analyze how malicious issuers use cross-website discrepancies of the ICO attributes to screen for naïve investors. Our model shares similarities with frameworks that analyze the prevalence of other scams and hoaxes 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 the investor types is that naïve investors may not) fund the ICO scam while astute investors will not. We define d_i to be the number of misrepresentations tolerated by an investor i, above which the investor immediately dismisses an ICO scam. Some naïve investors could have lower d than astute ones. But, on average, naïve investors are more tolerant of misrepresentations such that the average d for naïve investors is higher than their astute counterparts. We structure the following description of our model around Figure 2.

In period one, the issuer sets the number of misrepresentations d^* , which acts as a cutoff (screen) for investors who are viable targets. In forming this targeting strategy, the issuer faces the risk of classification errors. For a given d^* , the fraction of naïve investors who immediately dismisses the ICO scam is $F_{d|type}(d^* | naïve)$. Conversely, the fraction of naïve investors who remain viable targets to the scam is the complementary conditional cumulative distribution function $\bar{F}_{d|type}(d^* | naïve) = 1 - F_{d|type}(d^* | naïve)$. Likewise, the fraction of astute investors targeted is $\bar{F}_{d|type}(d^* | astute)$. Because $\bar{F}(\cdot)$ is monotonically decreasing in d, a higher (lower) d^* leads the issuer to target lower (higher) fractions of both naïve and astute investors.

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

In the final period, naïve investors ultimately fund the scam and while astute investors do not. Targeting a naïve investor yields the issuer a net profit G = Q - C, where Q is the gross proceeds from the scam. Whereas, an astute investor refrains from funding the scam, hence yielding the issuer 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^* .

$$\mathbb{E}(\Pi) = m \left[z \cdot \bar{F}_{d|\text{type}}(d^* \mid \text{na\"ive}) \cdot G - (1-z) \cdot \bar{F}_{d|\text{type}}(d^* \mid \text{astute}) \cdot C \right],$$
where $z = n/m$
(1)

It is instructive to examine an indiscriminate targeting strategy that abandons the screening strategy. The issuer targets all investors by choosing $d^* = 0$, thereby setting $\overline{F}_{d|\text{type}}(\cdot) = 1$. Imposing these constraints and $\mathbb{E}(\Pi) > 0$, we obtain equation (2). When C > 0, equation (2) implies that that an indiscriminate targeting strategy is profitable if the fraction of naïve investors is greater than the ratio C/(C+G). For example, suppose 1% of investors are naïve and G = \$1,000, then C can at most be $0.01/(1-0.01) \times \$1,000 = \10.10 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 entertain investors' queries. Finally, targeting all investors is also

profitable in the technical case of $G \to \infty$, which is patently unrealistic. The prevalence of misrepresented ICOs suggests that the above conditions are unmet in our sample.

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

3.2 Misrepresentations as a screening device

We examine tradeoffs implied by the targeting strategies. Figure 3 presents probability density plots of d, conditional on investor types—astute (black) and naïve (red). Shaded areas in black and red represent the complementary conditional cumulative distributions $\bar{F}_{d|type}(d^* | \text{astute})$ and $\bar{F}_{d|type}(d^* | \text{naïve})$, respectively. In Subfigure 3a, the malicious issuer adopts a conservative targeting strategy by choosing a high number of misrepresentations (high d^*). Because $\bar{F}(\cdot)$ is monotonically decreasing in d, the conservative strategy avoids many costly astute investors. However, the issuer necessarily forgoes many profitable naïve investors in the population. In Subfigure 3b, the issuer sets an aggressive targeting strategy by choosing a low d^* . While this strategy captures more naïve investors, it also retains more costly astute investors hence eroding the issuer's profits. Thus, the issuer needs to strike a balance between extremely conservative and aggressive targeting strategies.

The above exercise conveys the intuition for (i) why misrepresentations are so widespread, and (ii) how they are used as a screening device. To complete our analysis, we formalize the intuition from Figure 3. Given that d^* affects the quantities of naïve and astute investors being targeted, we can solve for the optimal targeting strategy (henceforth, OTP) of the malicious issuer. Using the chain rule, the issuer maximizes profits in equation (1) by choosing d^* such that:⁹

$$\partial \bar{F}(d^* \mid \text{na\"ive}) / \partial \bar{F}(d^* \mid \text{astute}) = \frac{1-z}{z} \cdot \frac{C}{G}$$
 (3)

Under the OTP, equation (3) prescribes the 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 OTP prescribes a low rate, which translates to an aggressive targeting strategy (see, Subfigure 3b). If the issuer has an inferior

⁹We first write the first order conditions of $\mathbb{E}(\Pi)$ with respect to $\bar{F}(d^* \mid \text{naïve})$ and $\bar{F}(d^* \mid \text{astute})$: $\partial \mathbb{E}(\Pi) / \partial \bar{F}(d^* \mid \text{naïve}) = zGm$ and $\partial \mathbb{E}(\Pi) / \partial \bar{F}(d^* \mid \text{astute}) = (1-z)Cm$. Next, we use chain rule to express the OTP as a function of z, C, and G: $\partial \bar{F}(d^* \mid \text{naïve}) / \partial \bar{F}(d^* \mid \text{astute}) = (1-z)Cm/zGm = (1-z)/z \cdot C/G$.

technology to entertain investors' queries (high C), then the issuer optimally chooses a higher rate that is achieved by a higher and more conservative d^* . Above all, 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 our sample ICOs.

In the context of our above analyses, we discuss two candidate explanations of ICO misrepresentations. First, the malicious issuer is unlikely to use misrepresentations to maximize investor interest. If that were the goal, misrepresentations are counterproductive because cross-site verification of ICO information is easy. Put differently, maximizing investor interest is like an overly aggressive targeting strategy, attracting too many costly astute investors who eventually balk at funding the ICO. Second, 34% of ICOs have misrepresentations at their first appearances in our sample. The sheer number of misrepresented ICOs makes issuers' carelessness an unsatisfactory explanation. Absent the screening mechanism and intention to scam, it is puzzling that so many issuers fail to accurately provide ICO information on listing websites.

Instead, we propose that the malicious issuer uses misrepresentations to screen for investor sophistication. Because investor sophistication is unobservable, a good strategy is to get naïve investors to self-identify. ICO misrepresentations will induce suspicions in all but the most naïve investors. Any astute investor who performs due diligence would recognize the misrepresentations and ignore the ICO. Those who remain are the naïve investors—the ideal targets of the malicious issuer.¹⁰ The issuer increases her odds of profitability by targeting naïve investors and repelling their astute counterparts. Having established the modus operandi of malicious issuers, we hypothesize that ICO misrepresentations predict scam risk.

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.

¹⁰The 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 an incredible story, in which the scammer impersonates a member of the Nigerian royal family. The inclusion of these tell-tale signs is not accidental but 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.

4.1 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.¹¹ 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.¹²

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 allegation with several media sources.¹³ 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",

¹¹Based 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.

¹²The numbers of unique ICOs covered by the listing websites are: ICORating (4,166), ICOBench (4,021), ICOData (1,896), ICODrops (625), and ICOCheck (580).

¹³Notably, the **Deadcoin** website also prominently displays a form to contest scam allegations.

"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.¹⁴ Figure 4 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 4 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 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 remaining control variables are continuous. The *duration* of an ICO is the length of its fundraising period in days. Finally, the *enforcement* and *disclosure* indices

¹⁴The 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.

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 tests of 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, to assess our screening mechanism more carefully, we extract data from the Ethereum blockchain to characterize the sophistication of investors who hold tokens in misrepresented ICOs.

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.¹⁵ 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 within every group:

$$S(t) = \begin{cases} \frac{(r_t - f_t)}{r_t} \times S(t - 1), & \text{for } t > 0\\ 1, & \text{for } t = 0 \end{cases}$$
(4)

Figure 5 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 5 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

¹⁵Right-censored observations are not necessarily cleared of scams.

conditional on survival up to time t, and $h_0(t)$ is the baseline hazard when all covariates equal zero, we estimate specification (5).

$$h(t) = h_0(t) \exp\left(\beta_1 misrep + \mathbf{X}^\top \boldsymbol{\beta}\right)$$
(5)

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 3.¹⁶ 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.¹⁷ 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).¹⁸ 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

 $^{^{16}}$ 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.

¹⁷Basic ICO characteristics are *ticker*, *country*, *banned*, *start date*, *end date*, *duration*, and acceptable payment modes. Nonbasic ICO characteristics are *softcap*, *hardcap*, *whitelist*, and *presale*.

¹⁸In 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.

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 Misrepresentations and wallet characteristics

To assess our screening mechanism more carefully, we extract data from the Ethereum blockchain.¹⁹ This blockchain is a digitally distributed, decentralized, public ledger of all transactions that occur on the Ethereum network. This means that we can observe token holdings and transaction activities of cryptocurrency wallets (henceforth, wallets) on the network. For every ICO, we use these data to characterize the sophistication of wallet-users who hold its tokens. Thereafter, we examine the relation between misrepresentations in an ICO and the sophistication of its typical token holder.

We provide details on the data collection process and how we measure the sophistication of wallet-users. First, we find the contract addresses of our sample ICOs by manually matching them by name and ticker on the website Etherscan.io.²⁰ For every ICO, we then find the Ethereum block height corresponding to 10 days after its end date. Next, by querying the contract address of an ICO in the Covalent Unified Application Programming Interface (API), we find the wallet addresses that hold its tokens as at its corresponding block height. For this analysis, we focus on the top 100 wallet addresses of every ICO by token holdings, and exclude an ICO if it has fewer than 30 wallets holding its tokens. Finally, we again query the Covalent Unified API to extract granular data on 110,607 wallets holding tokens of 1,996 ICOs.²¹

$$\log(sophistication) = \alpha + \beta_1 \mathbb{1}(misrep > 0) + \mathbf{X}^\top \boldsymbol{\beta}$$
(6)

For every ICO, we characterize the sophistication of its typical token holder. Specifically, we compute—at the wallet level—(i) total portfolio value (in U.S. dollars) of all tokens held, (ii) number of distinct tokens held, and (iii) number of transactions. Then, we aggregate these

 $^{^{19}\}mathrm{Most}$ ICO tokens adopt the ERC-20 (Ethereum Request for Comments 20) standard, which facilitates interoperability with other tokens on the Ethereum network.

 $^{^{20}{\}rm Every}$ ICO token has a unique contract address on the Ethereum blockchain. The Internet Appendix contains further details of our matching process.

²¹Of the full sample of 5,935 ICOs, we unambiguously matched 4,611 ICOs to their contract addresses on Etherscan.io. The remaining attrition is due to our requirement that the ICO token must have at least 30 holders at 10 days after the end date.

measures at the ICO level by taking their medians to obtain *value*, *diversity*, and *activity*, respectively. We make three conjectures about naïve investors. First, to the extent that wealth positively correlates with sophistication, they have lower wallet values. Second, they are more reckless or uninformed, so they diversify less by holding fewer distinct ICO tokens. Third, they have weaker technical or trading expertise, so they make fewer transactions. Thus, we expect investor sophistication to correlate positively with *activity*, *diversity*, and *age*. To test whether malicious issuers successfully use misrepresentations to screen for naïve investors, we estimate Poisson regressions in specification (6).

- Table 4 here -

Our results in Table 4 suggest that investors who hold tokens of misrepresented ICOs are less sophisticated. The dependent variable is one of *value*, *diversity*, and activity. The key independent variable is 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 calendarquarter cohort fixed effects, and standard errors are clustered by these cohorts. For ease of interpretation, we express estimated coefficients as incidence rate ratios. Column 1 indicates that the typical investor in misrepresented ICOs have a 40.1% (t = 2.61) lower wallet *value*. This result supports our view that misrepresented ICOs tend to attract less sophisticated investors. In column 2, we find that switching on 1(misrep > 0) is associated with a 19.7% (t = 2.88) decline in *diversity*. This finding suggests that token holders in misrepresented ICOs are more reckless and less financially savvy, pointing to investor naïvety. Column 3 shows that 1(misrep > 0) is associated with a 9.0% (t = 2.62) decrease in transaction *activity*. Thus, wallets that hold misrepresented ICO tokens likely belong to naïve investors who—due to their weaker expertise or inexperience—make fewer transactions.

Overall, our findings suggest that malicious issuers successfully use misrepresentations to screen for naïve investors. Wallets that hold tokens of misrepresented ICOs have characteristics associated with a lack of investor sophistication—they are less wealthy, less diversified, and less active. A caveat of our findings here is that a single person may control multiple wallets. However, it is unclear how this feature necessarily biases our findings.

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 focus on the misrepresentation behavior of ICOs launched shortly after news of regulatory actions taken by U.S. authorities. Second, we examine the relation between misrepresentations and ICO quality. Third, we apply network analysis to assess systematic patterns of misrepresentation behavior in the ICO ecosystem.

6.1 Regulatory action and misrepresentations

We examine whether the threat of regulatory action deters the use of misrepresentations by malicious ICO issuers. To test the deterrence effect, we begin by collecting news of regulatory actions taken by the U.S. authorities. As Appendix A shows, these regulatory actions primarily involve ICO fraud and conflicts of interest. None of these actions targets inaccurate disclosures on listing websites. Under our screening hypothesis, the prospect of costly regulatory scrutiny should deter malicious issuers from using misrepresentations. Alternatively, under the unintentional-mistakes explanation, regulatory scrutiny should have no effect on misrepresentations.

$$\log\left(\frac{p}{1-p}\right) = \alpha + \beta_1 news + \mathbf{X}^{\top} \boldsymbol{\beta}$$
(7)

$$\log(misrep) = \alpha + \beta_1 news + \mathbf{X}^{\top} \boldsymbol{\beta}$$
(8)

We construct two variables based on the timings of regulatory news releases and the first appearances of ICOs in our sample. First, the indicator variable $1(regulatory \ action)$ switches on if regulatory news is released in the calendar month prior to the first appearance of the ICO in our sample. Second, *regulatory intensity* is the number of regulatory news articles released one month prior to the first appearance of the ICO in our sample. Subsequently, we test how these variables affect the use of ICO misrepresentations. We estimate logistic regressions in specification (7). The outcome variable in this specification is 1(misrep > 0), an indicator that equals one if the ICO has at least one misrepresentation at its first appearance in our sample. The term p is the corresponding probability that 1(misrep > 0) switches on. Because *misrep* is a strictly non-negative quantity, we also estimate Poisson regressions in specification (8). The vectors **X** and β represent vectors of control variables and their corresponding estimated coefficients, respectively.

- Table 5 here -

Our results in Table 5 show that ICOs that shortly follow regulatory news have fewer misrepresentations. We estimate logistic (Poisson) regressions in columns 1 and 2 (3 and 4). where the dependent variable is 1(misrep > 0) (misrep). Estimated coefficients in the first (last) two columns are expressed as odds (incidence rate) ratios. On the extensive margin, we find in column 1 that the odds of an ICO using misrepresentations is 46.0% (t = 3.23) lower following releases of regulatory news. On the intensive margin in column 2, we find 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). Our results in columns 3 and 4 corroborate the view that the threat of regulatory action deters misrepresentation behavior. In column 3, we find that ICOs launched after releases of regulatory news have 35.6% (t = 3.90) fewer misrep. Column 4 shows that, following the release of an additional news article, ICOs have 16.2%(t = 2.91) fewer misrepresentations.

Overall, we find that the threat of regulatory action is correlated with misrepresentation behavior. Thus, misrepresentations are unlikely to be unintentional mistakes. Rather, there likely are elements of malice and criminality in the use of misrepresentations. Remarkably, we find a link between regulatory news and misrepresentation behavior although our sample of news articles does not mention the latter. Our preferred interpretation is that the threat of regulatory scrutiny deters malicious issuers from the strategic use of ICO misrepresentations. Under our screening model framework, this effect amounts to the issuer adopting a more aggressive targeting strategy, which may hurt the profitability of the ICO scam.

An alternative interpretation of our empirical results is that malicious issuers merely delay the launches of their ICO scams. If malicious issuers tactically time their launches, we expect *misrep* to have a stronger predictive effect on ICO scam risk when the threat of regulatory scrutiny is weaker. To assess this interpretation, we estimate Cox regressions of ICO scams on the interaction terms $1(regulatory action) \times misrep$ and *regulatory intensity* × *misrep*.²² We find that the loadings on these interaction terms are statistically insignificant. Thus, our findings in Table 5 are unlikely to reflect an issuer timing effect.

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

²²The Internet Appendix contains detailed results of these estimations.

be negatively associated with ICO quality. High quality ICOs may choose higher levels of voluntary disclosure to signal their quality and separate 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 $1(code \ posted)$ to equal one if the ICO discloses its source code on Etherscan.io and equals zero otherwise. Likewise, the indicator $1(code \ audited)$ switches on if the ICO posts a security audit of its source code on Etherscan.io. We estimate logistic regressions following specification (9). The term p is the probability that $1(code \ posted)$ (or, $1(code \ audited)$) switches on. 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.

$$\log\left(\frac{p}{1-p}\right) = \alpha + \beta_1 misrep + \mathbf{X}^\top \boldsymbol{\beta}$$
(9)

Our results in Table 6 suggest that ICO quality is not significantly different between misrepresented and non-misrepresented ICOs. In column 1, we find that an additional *misrep* is 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 a reflection of issuer quality and are unintentional mistakes. Column 2 shows a weak relation between *misrep* and odds of the ICO posting a security audit of its source code (+1.1%, t = 0.26). Again, this pattern is inconsistent with the alternative story that misrepresentations point to lower ICO quality.

As a robustness check, we adopt a market-based measure of ICO quality in column 3. Suppose that the market places a lower value on low-quality blockchain projects. Then, under the alternative explanation, we should observe that misrepresented ICOs attract 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 is also inconsistent with a quality-based explanation of ICO misrepresentations.

- Table 6 here -

Overall, we find that measures of ICO quality do not significantly vary with ICO misrepresentations. Thus, our findings reject the view that misrepresentations are merely unintentional mistakes, reflecting low issuer quality. Instead, our results thus far point to the strategic motives of issuers to target naïve investors with misrepresented ICO information.

6.3 Systematic patterns of misrepresentation behavior

To further 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 (e.g., Ballester, Calvó-Armengol, and Zenou, 2006). 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 6 here -

We formalize the above hypothesis in a simple network model of misrepresentations with strategic complementarities. Appendix B contains details of this model. Our model predicts that—in a network of ICOs linked by common advisors—ICOs with higher Katz centrality in the network exhibit more *misrep*. We empirically test this prediction. To construct the ICO network, we manage to match 2,110 advisors with 2,271 ICOs using data extracted

from the ICOBench listing website.²³ 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 6. 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 7 here -

To examine the relation between Katz centrality and *misrep* more rigorously, we estimate Poisson regressions in Table 7. Estimated coefficients are presented as incidence rate ratios. Consistent with our model predictions, column 1 shows that a 10% increase in Katz centrality is associated with a 4.6% (t = 2.27) rise in *misrep*.²⁴ 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 1(*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 7 support predictions from our network model—central ICOs have more misrepresentations. Due to strategic complementarities, we find systematic patterns of misrepresentation behavior among advisor-linked ICOs. In additional tests, we show that advisors of misrepresented ICOs are not penalized, but obtain more subsequent advisory opportunities. This surprising finding suggests that there exists many malicious issuers who solicit the services of such advisors. The Internet Appendix contains details and results of these tests. Overall, while advisors are valuable information and service intermediaries in the ICO market, some may play a role in the promotion of malignant behaviors.

 $^{^{23}{\}rm This}$ test has a smaller sample because we must exclude ICOs that either have no advisors or are unlinked to any ICOs.

²⁴We calculate this economic magnitude as follows: $\log(1.1) \times (1 - 1.485) = 0.046$.

7 Other suspicious actions

While malicious ICO issuers use misrepresentations to target naïve investors, such issuers may also engage in other suspicious actions 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.²⁵ Celebrity endorsements may be a potent screening device because naïve investors are 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 examine 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 1(celebrity)—an indicator that switches on if an ICO is endorsed by a celebrity. Here, we find that the scam risk of an ICO

²⁵Source: https://www.investor.gov/ico-howeycoins

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 1(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 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 our 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 30 days 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 one standard deviation increases in *BTC search* and *BTC returns* raise the odds of ICO scams by 62.8% (t = 4.74) and 41.9% (t = 4.63), respectively.²⁶ This pattern supports our idea that malicious issuers opportunistically time their ICOs to ride on periods of strong sentiment in cryptocurrency markets. Crucially, misrepresentations continues to predict ICO scams—an additional misrep increases the odds of ICO scams by 11.3% (t = 6.16). 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.

As a robustness check, we set up Model B, which uses *altcoin search* (i.e., Google Trends search volume index for the word "ICO") and *altcoin returns* as 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 search* and *altcoin returns* as instruments, our results in columns 3 and 4 are also consistent with our prior conclusions. ICOs coinciding with stronger sentiment in the alt-coin market are subsequently more likely to be scams.²⁷ In addition, we continue to find that misrepresented ICOs are more likely to be scams and detected as such.

8.2 Welfare analysis of ICO scams

Using estimates from our DCE models, we fit the models in columns 1 and 3 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 re-compute the proportion of probable scams. Model A and Model B in Table 9 estimate that 38.6% ($\hat{\sigma} = 29.0\%$) and 40.4% ($\hat{\sigma} = 26.8\%$) of ICOs in our sample are scams, respectively. Thus, there are potentially many ICO scams that are undetected. For additional context, the ICO advisory firm Satis Group estimates in an industry report that 78% of ICOs are scams (Dowlat, 2018).²⁸

²⁶We calculate economic magnitudes in column 1 as follows. $\sigma(BTC \ search) = 20.95; \ 20.95 \times (1.030 - 1) = 0.6285. \ \sigma(BTC \ returns) = 29.4\%; \ 29.4\% \times (2.428 - 1) = 41.98\%.$

²⁷We calculate economic magnitudes in column 3 as follows. $\sigma(altcoin \ search) = 20.93; 20.93 \times (1.023 - 1) = 0.4814.$ $\sigma(altcoin \ returns) = 82.7\%; 82.7\% \times (1.362 - 1) = 29.94\%.$

²⁸The Satis Group report uses a smaller and earlier sample and a different definition of ICO scams.

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. 5.07million in our sample. Suppose 40% of the 5,935 ICOs are scams. Then, ICO investors may be facing a loss of U.S. 5.07 million $\times 0.4 \times 5,935 = U.S.$ 12.03 billion. Thus, given the prevalence of ICO scams, more stringent regulations and enforcement—although costly—may be justified to protect investors. Specifically, because misrepresentations remain a powerful predictor of ICO scams, educating investors to perform simple cross-site verification of ICO characteristics may yield large benefits.

On the other hand, individuals may view risky ICO investments and traditional gambling devices in the same light.²⁹ For example, the U.S. Census Bureau reports that stateadministered 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, while our paper is agnostic on the net welfare effects, our estimated scale of ICO scams and its associated financial impact may inform cost-benefit tradeoffs of future regulatory policies. Specifically, given the role of misrepresentations in ICO scams, investments in regulatory scrutiny of ICO listing websites and in investor education could be particularly beneficial.

9 Conclusions

In this paper, we analyze how malicious issuers target naïve investors in ICO scams. 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

²⁹Anecdotal evidence from social media, such as the Reddit forums, supports this consideration.

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.

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. Thus, culpable advisors, instead of being penalized, can continue 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. Against the backdrop of these estimates, more stringent regulations and stronger enforcement actions may be justified to protect investor welfare. We believe that increased regulatory scrutiny of ICO listing websites may be particularly beneficial. Regulators can also educate the general public on how these frauds are 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 News of regulatory actions taken by U.S. authorities

Date	Title	News summary
16 th Jun 2018	SEC: Fraud surrounds initial coin offerings, blockchain security notwithstanding.	SEC has a unit that monitors ICO scams.
21 st Jun 2018	Members of the House will now be required to disclose bitcoin, other cryptocurrency holdings; Ethics Committee strongly encourage House members who are consid- ering investing in an ICO to seek guidance.	Ethics Committee have taken ac- tions to regulate House members in ICO investments.
27 th 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. Pro- motional efforts for cryptocurren- cies have come under fire from fed- eral and state regulators.
15 th Aug 2018	Even free tokens face regulatory heat as coin offerings scrutinized; SEC punishes company that didn't sell any tokens, saying potential in- vestors 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.
12 th Sep 2018	SEC takes first action against hedge fund over cryptocurrency in- vestments; In a separate case that's another first, agency penalizes bro- kers 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 web- site that connects investors with initial coin offerings.
12 th 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.
11 th Oct 2018	SEC says stop ICOs that falsely claimed SEC approval.	SEC's complaint charges Blockvest and Ringgold with violating federal securities laws.

(To be continued)

Date	Title	News summary
22 nd Oct 2018	SEC suspends trading in company for making false cryptocurrency- related claims about SEC regula- tion and registration.	SEC suspended trading in the secu- rities of a company for making false cryptocurrency-related claims.
16 th 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 to- kens.
30 th Nov 2018	Boxer Mayweather Jr., producer DJ Khaled agree to settle SEC crypto charges.	Celebrity endorsements of coin of- ferings may be illegal if the promot- ers fail to disclose the source and amount of their compensation.
21 st May 2019	SEC obtains emergency order halt- ing alleged diamond-related ICO Scheme targeting hundreds of in- vestors.	SEC halted a Ponzi scheme, which was purportedly a cryptocurrency business.
5 th 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 A.1. News of regulatory actions taken by U.S. authorities (Aug '18–Aug '19)

Appendix B Details of network model

Consider a set of ICOs $N = \{1, 2, ..., n\}$ that are members of a network g. In this network, two ICOs share a link if they are advised by at least one common advisor. Formally, for two ICOs i and j, we write:

$$g_{ij} = \begin{cases} 1, & \text{share a direct link} \\ 0, & \text{do not share a direct link or } i = j \end{cases}$$
(B.1)

A square symmetric matrix $\mathbf{G} = [g_{ij}]$ represents this network and tracks the direct links among ICOs. The matrix \mathbf{G} is also known as the adjacency matrix of the network. We will also consider a weighted network, in which g_{ij} 's are not necessarily binary and can take on numeric weights.

We use Katz centrality, which measures the network prominence of an ICO as the weighted sum of walks that emanate from it to other ICOs in the network. This implies that we need to account for indirect links of ICOs. To track indirect links in networks, we use the k-th power of the adjacency matrix— $\mathbf{G}^k, k \in \mathbb{Z}$.³⁰ An element $g_{ij}^{[k]}$ in \mathbf{G}^k gives the number of walks of length $k \geq 1$ from i to j in the network.³¹ In the special case of k = 0, \mathbf{G}^k is defined as the identity matrix \mathbf{I} .

To operationalize the Katz centrality measure, consider a matrix \mathbf{M} that tracks the number of walks of all lengths between any two ICOs.

$$\mathbf{M} = \sum_{k=0}^{+\infty} \theta^k \mathbf{G}^k \quad \text{with element } m_{ij} = \sum_{k=0}^{+\infty} \theta^k g_{ij}^{[k]}$$
(B.2)

The term θ^k is the decay factor applied to walks of length k. Economically, the decay factor controls how much influence an ICO has on another ICO in the network. This influence increasingly wanes if two ICOs are further away (i.e., k > 1) from each other. We can also

³⁰While an ICO plays its equilibrium number of misrepresentations based on its direct network neighbors', these neighbors respond to their own set of neighbors, so on and so forth. Thus, the equilibrium response of an ICO also depends on other indirectly linked ICOs.

³¹This is an established result in graph theory.

derive an equivalent expression of \mathbf{M} below.³²

$$\mathbf{M} = [\mathbf{I} - \theta \mathbf{G}]^{-1} \tag{B.3}$$

By definition, the Katz centrality of ICO *i*—denoted as $b_i(g, \theta)$ —is the sum of elements of the *i*-th row in **M**.

$$b_i(g,\theta) = \sum_{j=1}^n m_{ij} = \sum_{j=1}^n \sum_{k=0}^{+\infty} \theta^k g_{ij}^{[k]}$$
(B.4)

Following equation (B.3), the $(n \times 1)$ vector of Katz centralities is thus:

$$\mathbf{b}(g,\theta) = \mathbf{M} \cdot \mathbf{1} = [\mathbf{I} - \theta \mathbf{G}]^{-1} \cdot \mathbf{1}$$
(B.5)

We specify a linear-quadratic utility function of ICOs (or equivalently, issuers) that captures both ICO-specific and complementary components of misrepresentation behavior. This formulation is popular in network economics because it admits a tractable solution and cleanly characterizes the equilibrium as a function of network structure. For $\alpha_i > 0$ and $\theta > 0$, we write equation (B.6).

$$u_i(d_i, d_{-i}, g) = \alpha_i d_i - \frac{1}{2} d_i^2 + \theta \sum_{j=1}^n g_{ij} d_i d_j$$
(B.6)

Let us emphasize the perspective here—the network has already formed. The ICO observes its advisor-linked ICO peers and chooses d_i to maximize utility in equation (B.6). We are agnostic about the network formation process. Rather, our model takes the network structure as given and focuses on the complementarities in misrepresentation behavior among ICOs.

We obtain a tractable solution of this game. The ICO network has formed, and the issuer chooses d_i to maximize utility. The first-order condition of equation (B.6) gives the following best-response function.

$$d_i^* = \alpha_i + \theta \sum_{j=1}^n g_{ij} d_j^*, \quad \forall i = 1, 2, ..., n$$
 (B.7)

³²Following equation (B.2), we first express $\mathbf{M} = \mathbf{I} + \theta \mathbf{G} + \theta^2 \mathbf{G}^2 + \theta^3 \mathbf{G}^3 + \dots$ Next, we multiply this expression by $\theta \mathbf{G}$ to get $\theta \mathbf{G} \mathbf{M} = \theta \mathbf{G} + \theta^2 \mathbf{G}^2 + \theta^3 \mathbf{G}^3 + \dots$ Finally, the difference of these expressions yields equation (B.3).

The best-response function can be equivalently expressed in matrix form. Solving equation (B.8) and using equation (B.5), we show that the Nash equilibrium vector \mathbf{d}^* is proportional to the vector of Katz centralities \mathbf{b} .

$$\mathbf{d}^* = \boldsymbol{\alpha} + \theta \mathbf{G} \mathbf{d}^* = [\mathbf{I} - \theta \mathbf{G}]^{-1} \boldsymbol{\alpha} = \mathbf{M} \boldsymbol{\alpha}$$
(B.8)

References

- Aloosh, A. and Li, J. (2019). "Direct evidence of Bitcoin wash trading". Available at SSRN 3362153.
- 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.
- Button, M., Lewis, C., and Tapley, J. (2009). "A better deal for fraud victims". National Fraud Authority, United Kingdom.
- 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.
- 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.
- Dittmar, R. and Wu, D. ((2019). "Initial coin offerings hyped and dehyped: An empirical examination". Available at SSRN 3259182.
- Dowlat, S. (July 2018). Cryptoasset market coverage initiation: Network creation. Tech. rep. Satis Group.
- Fahlenbrach, R. and Frattaroli, M. (2020). "ICO investors". Financial Markets and Portfolio Management.
- Feinstein, J. (1990). "Detection controlled estimation". Journal of Law and Economics 33, 233–276.
- 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.
- Griffin, J. and Shams, A. (2020). "Is Bitcoin really unterhered?" 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.
- 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".
- 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.

	Ended	
AdHive Al-controlled influencer marketing platform AdHive is the first Al-controlled influencer Marketing platform with Blockchain technological solutions. The AdHive platform fully automates all steps of interaction with Influencers in order to save a huge amount of time and effort for advertises. The factor will offic heads the moonchube in order to an adverse at one an uniform and effort for advertises. The factors of the one of the value of a one an influence of the moon of the number of	Token Type Price in ICO Country Whitelist/KYC Restricted areas preICO start preICO start	ADH Utility 0.1369 USD Estonia KYC & Whitelist USA, China 30th Jan 2018 30th Jan 2018
nfluencer channels without having to worry about proper execution. Native video advertising will become easy to run, and new opportunities for blog monetization will power community development and increase audience and divertising capacity.	ICO start ICO end	28th Feb 2018 14th Mar 2018
Entertainment Communication Business services Artificial Intelligence Internet Media Other Platform	VISIT I	CO WEBSITE

Token info		Investment info	þ
Token	ADH	Min. investment	0.05 ETH, 0.005 BTC
Platform	Ethereum	Accepting	ETH, BTC, Fiat
Туре	ERC20	Distributed in ICO	60%
		Soft cap	2,000,000 USD
Price in ICO	0.1369 USD	Hard cap	12,000,000 USD
BONUS Pre-sale: 15%-30%	Bonus Token Sale Phase #1:	Raised	\$12,000,000
0%-15% Bonus			

I C O R A T I N G	Ratings	Crowdsales	Articles	Analytics	Reports	News	Ad
		7	1 Independent C	rypto Opinions and	d Ratings		
AdHive							
Marketing & Advertising							
Crowdsale							
Pre-sale		Toke	n details				
Pre-sale start date 3	0 Jan 2018	Ticke	r			А	DH
Pre-sale end date 0	6 Feb 2018	Туре				Utility-to	ken
		Addit	tional Toker	Emission			No
Tokon Sala		Acce	pted Curre	ncies		1	ETH
loken Sale		Toker	n distributio	on 60)% – Token S	Sale	
ICO start date 2	8 Feb 2018			10	5% - Netwo 1.5% - AdHi	rk Growth ive Founde	rs
ICO end date 2	8 Feb 2018			65 3.	% - Advisor 5% - Comm	y Board nunity gran	ts
Hard cap size 12,000,000) USD (fiat)			25 15	% - Reserve % - Legal Co	Fund ompliance	
Raised 12,00	0,000 USD						

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

ICODR OPS Q Search ICO

Available for Token Sale: 30%



AdHive (Advertising)

World's first Al-controlled Influencer Marketing platform. Our service offers a fully automated, blockchainbased solution for mass placement of native video ads on influencers' channels.

ACTIVE ICO UPCOMING ICO ENDED ICO WHITELIST ICO STATS



Figure 1. (continued)

Accepts: ETH, BTC

Issuer selects d^* in targeting strategy	Targeted investors impose costs	Gross funding proceeds
 Targeted investors mz · F̄_{d type}(d* naïve) m(1-z) · F̄_{d type}(d* astute) Dismissed investors mz · F_{d type}(d* naïve) m(1-z) · F_{d type}(d* astute) 	$mz \cdot \bar{F}_{d \text{type}}(d^* \mid \text{na\"ive}) \times C$ $m(1-z) \cdot \bar{F}_{d \text{type}}(d^* \mid \text{astute}) \times C$	$mz \cdot \bar{F}_{d type}(d^* \mid \text{na\"ive}) \times Q$
Period (1): ICO launches	Period (2): ICO in progress	Period (3): ICO completes

Figure 2. 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), these targeted investors 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 ultimately refrain from funding the scam.



Figure 3. 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 $\bar{F}_{d|\text{type}}(d^* | \text{astute})$ and $\bar{F}_{d|\text{type}}(d^* | \text{naïve})$, respectively. Sub-figures 3a and 3b visualize a conservative targeting strategy (high d^*) and an aggressive targeting strategy (low d^*), respectively.



Figure 4. 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 5. 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 6. This figure presents a circular layout of the advisor-linked ICO network described in Section 6.3. 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.

	v					
	Ν	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	0.89	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 A. Summary statistics

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	(\mathbf{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	(1)	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*, 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 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 characteristics at its first appearance in our sample. The indicator 1(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.

	(1)	(2)	(3)	(4)
1(Misrep > 0)	3.740			
	(5.46)			
Misrep		1.253	1.140	
*		(6.71)	(2.18)	
Misrep ^{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	Ń	Ń	Ý	Ý
Coverage-quartile FE	Ν	Ν	Υ	Υ
Clustered SE	Ν	Ν	Υ	Υ

Table 4. On-chain analysis: Misrepresentations and wallet characteristics

This table presents estimates from Poisson regressions. Estimated coefficients are expressed as incidence rate ratios. For every ICO, we first analyze individual cryptocurrency wallets that hold its tokens at 10 days after the ICO end date. 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 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 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.

	(1)	(2)	(3)
Dependent variable:	Value	Diversity	Activity
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
D . 1	(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	Ý	Ý	Ý
Clustered SE	Y	Y	Y

Table 5. Regulatory action and misrepresentations

Columns 1 and 2 (3 and 4) of this table present estimates from logistic (Poisson) regressions. Estimated coefficients in columns 1 and 2 (3 and 4) are expressed as odds (incidence rate) ratios. The dependent variable in columns 1 and 2 is 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 columns 3 and 4 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 1(regulatory action) and regulatory intensity. The variable 1(regulatory action) is an indicator that equals one if regulatory news is released within the calendar month prior to the first appearance of the ICO in our sample, and equals zero otherwise. The variable regulatory intensity is the number of regulatory news articles released within the calendar month prior to the first appearance of the ICO in our sample. 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.

Dependent variable: Misrep								
	(1)	(2)	(3)	(4)				
Dependent variable:	1(Mis	$\operatorname{rep} > 0)$	M	isrep				
1(Regulatory action)	0.540		0.644					
	(3.23)		(3.90)					
Regulatory intensity	· · · ·	0.795	· · ·	0.838				
		(2.13)		(2.91)				
Banned	0.763	0.772	0.921	0.926				
	(1.53)	(1.41)	(1.72)	(1.58)				
Whitelist	0.509	0.506	0.940	0.938				
	(4.42)	(4.47)	(1.49)	(1.53)				
Duration	0.998	0.998	0.998	0.998				
	(2.29)	(2.25)	(3.30)	(3.30)				
Presale	4.270	4.277	2.425	2.432				
	(8.71)	(8.71)	(8.99)	(9.03)				
Hardcap	4.471	4.479	3.239	3.253				
~ ^	(10.13)	(10.19)	(22.39)	(22.11)				
Softcap	0.811	0.818	0.990	0.993				
	(1.85)	(1.78)	(0.30)	(0.21)				
Accept BTC	1.268	1.270	1.140	1.141				
	(2.35)	(2.29)	(3.48)	(3.45)				
Accept ETH	4.879	4.867	2.444	2.447				
	(6.30)	(6.29)	(8.89)	(8.83)				
Accept USD	0.881	0.882	0.996	0.997				
	(1.23)	(1.23)	(0.11)	(0.10)				
Enforcement	1.541	1.538	1.161	1.161				
D • 1	(3.29)	(3.26)	(4.26)	(4.26)				
Disclosure	1.373	1.373	1.146	1.147				
	(4.21)	(4.19)	(6.73)	(6.76)				
SEC filing	1.000	1.014	0.990	0.995				
	(0.00)	(0.04)	(0.08)	(0.04)				
# ICOs	5,935	$5,\!935$	$5,\!935$	$5,\!935$				
Cohort FE	Y	Y	Υ	Y				
Clustered SE	Y	Y	Υ	Υ				

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 $1(code \ posted)$, $1(code \ audited)$, and raised. The indicator $1(code \ posted)$ equals one if the ICO posts the source code of its smart contract on Etherscan.io and equals zero otherwise. The indicator $1(code \ audited)$ equals one if the ICO posts a security audit of its source code on Etherscan.io and equals zero otherwise. The indicator $1(code \ audited)$ 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:	1(Code posted)	$\mathbb{1}(\text{Code audited})$	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	Υ	Υ

Table 7. 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. See Appendix B for 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. See Appendix B for details on Katz centrality. 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	Ν	Y
$\log(\text{Centrality})$	1.485	1.567		
	(2.27)	(2.17)		
1(High centrality)			1.061	1.067
D	0 0 - 1	0 0 - ((1.96)	(2.25)
Banned	0.974	0.974	0.974	0.974
TT TI 4. 14 .	(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
D	(1.56)	(1.56)	(1.56)	(1.57)
Presale	(1.590)	1.591	1.588	1.587
TT 1	(7.47)	(7.49)	(7.60)	(7.62)
Hardcap	1.598	1.599	1.596	1.597
a c	(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
D . 1	(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	Υ	Υ	Y	Υ
Clustered SE	Υ	Υ	Υ	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 $1(\ celebrity)$, web traffic ratio, and misrep. The indicator $1(\ celebrity)$ 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 in some models are clustered by ICO cohorts. t-statistics are reported in parentheses.

Event: ICO scam				
	(1)	(2)	(3)	(4)
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)
Banned	1.062	1.058	1.059	1.048
TTT1 + . 1+ .	(0.18)	(0.16)	(0.18)	(0.14)
Whitelist	1.497	1.374	1.507	1.423
	(2.11)	(1.55)	(1.87)	(1.54)
Duration	(0.999)	(0.999)	1.000	1.000
Dregale	(0.21)	(0.23)	(0.02)	(0.01)
Fresale	(1.65)	1.042	(0.74)	(0.940)
Hardcan	(1.05) 1 758	(0.20)	(0.74)	(0.37) 1.651
Hardcap	(2.77)	(1.88)	(2.67)	(1.85)
Softcap	0.917	0.919	(2.01) 0.947	(1.00)
Solicap	(0.48)	(0.48)	(0.39)	(0.45)
Accept BTC	1.376	1.361	1.339	1.303
F	(1.38)	(1.38)	(1.30)	(1.17)
Accept ETH	1.110	1.034	1.383	1.286
-	(0.38)	(0.14)	(1.14)	(0.98)
Accept USD	1.321	1.321	1.308	1.298
	(0.89)	(0.91)	(0.87)	(0.84)
Enforcement	0.592	0.582	0.582	0.586
	(2.20)	(2.18)	(2.13)	(2.09)
Disclosure	0.927	0.919	0.898	0.892
	(1.08)	(1.25)	(1.64)	(1.77)
SEC filing	0.591	0.602	0.550	0.560
	(0.81)	(0.67)	(0.93)	(0.77)
# ICOs	5,935	5,935	5,935	5,935
Cohort strata	Ý	Ý	Y	Ý
Coverage-quartile FE	Υ	Y	Υ	Y
Clustered SE	Υ	Y	Υ	Y

Table 9. Partial observability of ICO scams

This table presents estimates from detection controlled estimation (DCE) models. Estimated coefficients are expressed as odds ratios. We simultaneously model the scam and detection processes of ICO scams. The instruments for the scam process in Model A (Model B) are *BTC search* and *BTC returns* (altcoin search and altcoin returns). The variable *BTC search* (altcoin search) is the cumulative search volume index of the word "Bitcoin" ("ICO") on Google Trends 30 days prior to the ICO start date. The variable *BTC returns* (altcoin returns) is the cumulative returns of Bitcoin (non-Bitcoin cryptocurrencies) 30 days 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)				
-	Model A		Model B					
	Scam	Detection	Scam	Detection				
BTC search	1.030							
BTC returns	(4.74) 2.428 (4.63)							
Altcoin search	(1.00)		1.023					
			(5.20)					
Altcoin returns			1.362					
Misren	1 1 1 2	1 110	(5.06) 1.130	1 1 1 6				
wistep	(6.16)	(6.32)	(6.65)	(6.60)				
Banned	0.716	1.103	0.752	1.158				
	(1.51)	(0.48)	(1.33)	(0.72)				
Whitelist	2.902	0.842	1.786	0.915				
	(4.24)	(1.74)	(4.13)	(0.95)				
Duration	1.000	1.000	1.004	0.999				
ות	(0.06)	(0.41)	(2.79)	(1.45)				
Presale	(2.59)	1.157	(4.22)	1.214				
Hardcan	(3.38) 2 418	(1.40) 1 106	(4.33) 1 /12	(1.80) 1 301				
Hardcap	(4.35)	(0.78)	(2.69)	(2.12)				
Softcap	0.649	1.006	0.580	1.031				
Sorroup	(3.39)	(0.06)	(3.96)	(0.31)				
Accept BTC	1.492	1.095	0.976	1.178				
*	(3.00)	(0.91)	(0.23)	(1.64)				
Accept ETH	2.265	0.846	4.859	0.632				
	(3.79)	(1.47)	(4.50)	(3.40)				
Accept USD	2.043	0.967	0.625	1.260				
	(3.46)	(0.25)	(2.60)	(1.64)				
Enforcement	(4.46)	1.1(3)	(4.85)	1.131				
Dicelocure	(4.40) 1.220	(1.31)	(4.85) 1 277	(1.02)				
Disclosure	(3.45)	(1.90)	(4.91)	(2, 30)				
SEC filing	(3.43)	(1.20) 2 107	(4.21) 0.353	(2.30) 0.887				
	(3.12)	(1.21)	(2.08)	(0.25)				
	(0.1-)	()	(=.00)	(0.20)				
# ICOs	$5,\!935$	$5,\!935$	$5,\!935$	$5,\!935$				

Detection controlled estimation (DCE)

Internet Appendix to: TRUST, BUT VERIFY: THE ECONOMICS OF SCAMS IN INITIAL COIN OFFERINGS^{*}

Kenny Phua Bo Sang Chishen Wei Gloria Yang Yu

The Internet Appendix contains supplementary information and additional tests for the paper "TRUST, BUT VERIFY: THE ECONOMICS OF SCAMS IN INITIAL COIN OFFERINGS". Contents of the Internet Appendix are organized as follows.

- Section I Details of the process to match ICOs in our sample to their contract addresses on the Ethereum network.
- Section II Testing an alternative interpretation of the relation between misrepresentations and regulatory action.
- Section III Examining the subsequent advisory opportunities of advisors in misrepresented ICOs.
- Section IV Details of our detection controlled estimation (DCE) model and derivation of its likelihood function.

^{*}Wei: cswei@smu.edu.sg, Singapore Management University. Phua: kenny.phua@uts.edu.au, University of Technology Sydney. Sang: bo.sang.2017@pbs.smu.edu.sg, Singapore Management University. Yu: gloriayu@smu.edu.sg, Singapore Management University.

I Matching ICOs to Ethereum contract addresses

We describe how we match an ICO token to its contract address on the Ethereum network. To begin our matching process, we search for either the name or ticker of the ICO on the Etherscan.io/tokens website. In the ideal case, we would find only one match from this search. In that case, we collect the contract address stated on the Etherscan page. To increase the likelihood that we capture the token contract, we check that there is a number of "Holders" stated on the page. If we find multiple matches, we pick the contract address with the highest number of "Holders". Sometimes, the name/ticker of an ICO is exceedingly common and will match with many ICO projects. Due to time and resource constraints, we abandon a token match in such cases where our search turns up 100 or more potential matches. For our analysis in Section 4.2 of the main text, we also exclude from our sample ICOs that are unmatched on the Etherscan website. Figure I contains screenshots of the Etherscan page for the BNB ICO token.

- Figure I here -

II Do malicious issuers delay ICO launches around regulatory action?

Following our analysis in Section 5.1 of the main text, we test whether our results are explained by malicious issuers delaying ICO launches around regulatory action. If malicious issuers tactically time their launches, we expect *misrep* to have a stronger predictive effect on ICO scam risk when the threat of regulatory scrutiny is weaker. To assess this explanation, we adopt the empirical setup of Table 3 in the main text, and augment the models with interaction terms between *misrep* and 1(regulatory action) (or *regulatory intensity*). We estimate specification (I), where the vectors **X** and β represent vectors of other variables (i.e., control variables and main effects) and their corresponding estimated coefficients, respectively.

$$h(t) = h_0(t) \exp\left(\beta_1 misrep + \beta_2 news \times misrep + \mathbf{X}^\top \boldsymbol{\beta}\right)$$
(I)

We present our estimates in Table I. In column 1, we find that the estimated loading on $1(regulatory action) \times misrep$ is statistically insignificant (t = 0.69). In column 2, we stratify our hazard model by ICO cohorts and saturate it with coverage-quartile fixed effects. Here, we continue to find that the interaction term $1(regulatory action) \times mis$ - rep has a statistically insignificant loading (t = 0.43). Using regulatory intensity in lieu of 1(regulatory action), column 3 also shows that the predictive effect of misrep on ICO scam risk does not significantly vary around regulatory news, or lack thereof. Overall, our findings suggest that malicious issuers are not simply delaying their ICO launches around regulatory news. Rather, the threat of regulatory scrutiny deters malicious issuers from the strategic use of ICO misrepresentations.

- Table I here -

III Subsequent advisory opportunities

Our findings in Section 5.3 of the main text suggest that advisors facilitate the promotion of misrepresentation behavior in the ICO market. We now analyze the subsequent advisory opportunities of advisors in misrepresented ICOs—*misrep advisors*. On one hand, market participants may penalize *misrep advisors*, so these advisory opportunities may diminish. On the other hand, if there are enough malicious issuers who actively solicit the services of *misrep advisors*, we expect these advisory opportunities to persist or increase.

To perform this analysis, we track ICOs that are subsequently advised by one or more advisors of every ICO. As an illustration, suppose ICO P starts on January 2019 and is advised by Alice, Bob, and Carol. Subsequently, Alice advises ICO Q, which starts on April 2019. Later, Bob and Carol advise ICO R with a start date on May 2019. Under our empirical framework, ICOs Q and R are *subsequent ICOs* to ICO P. In a nutshell, conditional on an ICO, higher incidence rates of *subsequent ICOs* translate to more subsequent advisory opportunities for advisors.

To test whether misrepresentations of an ICO affect subsequent advisory opportunities of its advisors, we examine both the likelihood of future advisory opportunities and the length of time intervals between ICOs. Time intervals between ICOs are economically meaningful because shorter time intervals may indicate higher demand for the advisors' services while longer time intervals suggest that the advisors are out of favor. We estimate Cox regression models, which explicitly account for the time dimension of advisory opportunities. Logistic regressions are also a viable econometric alternative, but they ignore the time dimension of subsequent advisory opportunities.

- Table II here -

Table II shows that *misrep advisors* are significantly more likely to be hired to advise

future ICOs. In columns 1 and 2, we estimate standard Cox models, which do not consider recurrent events. That is, standard Cox models consider time to the most immediate *subsequent ICOs* but ignore the other ones.¹ Column 1 shows that an additional *misrep* increases the hazard of *subsequent ICOs* by 2.5% (t = 2.94). To assess economic significance, we construct a back-of-envelope benchmark. In our sample, there are 2,271 ICOs distributed over 13 months, 2,110 advisors, and an average of 2.89 advisors per ICO. Assuming ICOs are uniformly distributed across time, the unconditional odds of having at least one *subsequent ICO* in the next month is 26.9%.² This implies that the economic effect of an additional *misrep* is about one-tenth the unconditional odds of *subsequent ICOs*.³ In column 2, we find that the presence of at least one misrepresentation increases the hazard of *subsequent ICOs* by 17.2% (t = 2.42). Our findings suggest that *misrep advisors*, instead of being penalized, enjoy more subsequent advisory opportunities.

For robustness, we next estimate Prentice, Williams, and Peterson Total Time (PWP-TT) models. The PWP-TT models extend the standard Cox models and allow us to consider recurrent subsequent ICOs events. This feature allows us to use information of all subsequent ICOs—not only the immediate ones—in our estimation. As the track record of prior advisory opportunities becomes more established, subsequent ICOs may become more likely. PWP-TT models accommodate this feature in our setting by allowing likelihoods of subsequent ICOs to vary with event order. In column 3, we continue to find that misrep is positively associated with subsequent ICOs (+0.8% increase in hazard, t = 2.79). Because we account for recurrent subsequent ICOs, the overall effect of misrep may be partly displaced by the track-record effect of prior advisory opportunities, hence the smaller economic magnitude. Our conclusions remain unchanged with our results in column 4.

Overall, our analysis in this section helps us better understand the complexion of misrepresentation behavior in the ICO market. Particularly, we find that *misrep advisors*, far from being penalized by the market, enjoy better subsequent advisory opportunities. Thus, advisors may have incentives to continue adopting and promoting the strategic use of misrepresentations across ICOs.

¹Using the Alice-Bob-Carol example as an illustration, standard Cox models focus on time to ICO Q (April 2019) but discard information about ICO R (May 2019).

²The probability of an advisor being hired by an ICO is $2.89/2,110 = 1.37 \times 10^{-3}$. Assuming ICOs are uniformly distributed across time, there is an average of 2,271/13 = 175 ICOs per month. The binomial distribution probability of having at least one *subsequent ICO* in the next month is $1 - f_{\text{binom}}(0, 175, 1.37 \times 10^{-3}) = 21.2\%$. Thus, the odds are 21.2%/(1 - 21.2%) = 26.9%.

 $^{^{3}}$ We make this approximate comparison only for benchmarking purposes. Technically, hazard ratios are not exactly comparable to odds ratios. The former has a condition of non-event up till the current time.

IV Details of the Detection Controlled Estimation (DCE) model

In this section, we set up the DCE model and derive its likelihood function.

IV.I Setting up the DCE model

To account for incomplete detection, we use detection controlled estimation (DCE) methods developed by Feinstein (1990). 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. Equation (II) characterizes ICO scams. The latent variable Y_{Si} is a function of ICO-specific attributes, which are represented by vector \mathbf{X}_{Si} . The binary variable S_i equals one if ICO *i* is a scam, and equals zero otherwise. Notably, we cannot directly observe S_i under incomplete detection. Instead, we only observe ICOs that are both scams and detected.

$$Y_{Si} = \mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S} + \epsilon_{Si} \qquad S_{i} = \begin{cases} 1, & \text{if } Y_{Si} \ge 0\\ 0, & \text{if } Y_{Si} < 0 \end{cases}$$
(II)

To complete the DCE model, we characterize the detection process in equation (III). Conditional on ICO *i* being a scam, the probability of detection on the **DeadCoins** website Y_{Di} is a function of ICO-specific attributes represented by vector \mathbf{X}_{Di} . The binary variable D_i equals one if an ICO *i* is jointly a scam and detected, and equals zero otherwise.⁴ Under incomplete detection, we again do not observe D_i because the **DeadCoins** website may not detect some ICO scams.

$$Y_{Di} = \mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D} + \epsilon_{Di} \qquad D_{i} = \begin{cases} 1, & \text{if } Y_{Di} \ge 0\\ 0, & \text{if } Y_{Di} < 0 \end{cases}$$
(III)

We estimate our DCE models using the maximum likelihood method. First, we define

⁴The DCE model implicitly assumes that there are no false positives. That is, the model assumes there are no ICOs detected on the **DeadCoins** website that are in fact not scams. Such false positives are improbable because maligned ICO issuers should have strong reputational incentives to clear their names. Indeed, the **Deadcoins** website prominently displays a form to contest scam allegations. Moreover, we manually verify the legitimacy of every scam allegation with reports of regulatory actions, media articles, and message board discussions.

 $P^{S}(\cdot)$ and $P^{D}(\cdot)$ to be monotonic link functions that map $\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}$ and $\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D}$ to latent probabilities of (i) an ICO being a scam and (ii) an ICO scam being detected, respectively.

$$P^{S}(\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}) = \operatorname{Prob}(S_{i} = 1)$$

$$P^{D}(\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D}) = \operatorname{Prob}(D_{i} = 1 \mid S_{i} = 1)$$
(IV)

Next, we assign ICOs identified on the DeadCoins website to set Z and all other ICOs to its complement Z^c. Finally, we select $\hat{\beta}_S$ and $\hat{\beta}_D$ to maximize the likelihood function in equation (V). The next subsection contains a detailed derivation of this likelihood function.

$$\log \ell = \sum_{i \in Z} \log \left[P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right] + \sum_{i \in Z^{c}} \log \left[1 - P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right]$$
(V)

$$\left(\hat{\boldsymbol{\beta}}_{S}, \hat{\boldsymbol{\beta}}_{D}\right) = \underset{\boldsymbol{\beta}_{S}, \boldsymbol{\beta}_{D}}{\arg\max} \log \ell(\boldsymbol{\beta}_{S}, \boldsymbol{\beta}_{D})$$
(VI)

IV.II Derivation of the DCE likelihood function

We define $P^{S}(\cdot)$ and $P^{D}(\cdot)$ to be monotonic link functions that map $\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}$ and $\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D}$ to latent probabilities of (i) an ICO being a scam and (ii) an ICO scam being detected, respectively. Formally, we write:

$$P^{S}(\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}) = \operatorname{Prob}(S_{i} = 1)$$
(VII)

$$P^{D}(\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D}) = \operatorname{Prob}(D_{i} = 1 \mid S_{i} = 1)$$
(VIII)

The following matrix presents the probabilities of various joint outcomes. These joint outcomes are mutually exclusive and exhaustive. So, the probabilities in the matrix sum to unity.

	Scam	Honest
Detected	$P^{S}(\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D})$	0
Undetected	$P^{S}(\mathbf{X}_{Si}^{\top}\boldsymbol{\beta}_{S}) \times \left[1 - P^{D}(\mathbf{X}_{Di}^{\top}\boldsymbol{\beta}_{D})\right]$	$1 - P^S(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_S)$

We can only observe ICOs that are both scams and detected (i.e., North-West quadrant of matrix). For ease of notation, we assign ICOs that are jointly scams and detected to set Z. That is, the sample of ICO scams detected on the DeadCoins website belong to the set

The log-likelihood function of ICOs that belong to set Z is the log of sum of the probabilities of the joint outcome $S_i \times D_i$. We thus write the following.

$$\log \ell_Z = \sum_{i \in Z} \log \left[P^S(\mathbf{X}_{Si}^\top \boldsymbol{\beta}_S) \times P^D(\mathbf{X}_{Di}^\top \boldsymbol{\beta}_D) \right]$$
(IX)

Similarly, the log-likelihood function of ICOs that belong to the complement of set Z is the log of sum of the probabilities of (i) an honest ICO and (ii) an ICO scam that is undetected. Formally, we write the following.

$$\log \ell_{Z^{\mathsf{c}}} = \sum_{i \in Z^{\mathsf{c}}} \log \left\{ P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times \left[1 - P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right] + 1 - P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \right\}$$
$$= \sum_{i \in Z^{\mathsf{c}}} \log \left[1 - P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right]$$
(X)

Set Z and its complement Z^{c} constitute the full sample of ICOs. Thus, the full-sample log-likelihood function is:

$$\log \ell = \sum_{i \in \mathbb{Z}} \log \left[P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right] + \sum_{i \in \mathbb{Z}^{c}} \log \left[1 - P^{S}(\mathbf{X}_{Si}^{\top} \boldsymbol{\beta}_{S}) \times P^{D}(\mathbf{X}_{Di}^{\top} \boldsymbol{\beta}_{D}) \right]$$
(XI)

1 Etherscan			All Filters ~	Filters v Search by Address / Txn Hash / Block / Token / Ens					٩		
Eth: \$4,449.	.13 (+5.28%) <u>∎</u> ∛ 136	Gwei			Home	Blockchain ~	Tokens 🗸	Resources ~	More -	JPHUA005	~ *
Token T	Tracker ERC-20										
ERC-20	Tokens										٩
A total of	472,460 Token Cont	tracts found							First <	Page 1 of 19	Last
#	Token		Price	Change (%)	Volume (24H)		∽ <u>Market</u>	<u>Cap</u> ⑦		Holders	
1	BNB (BNB) Binance aims to crypto exchang crypto finance.	b build a world-class e, powering the future of	\$622.4536 0.010696 Btc 0.139905 Eth	▲ 3.05%	\$2,465,283,527.00	0	\$103,825,	972,790.00		319,179 • 0.017%	
Eth: \$4,444.8	herscan 85 (+3.41%) 🔜 142 (Gwei			All Filters ~ Home	Search by Addr Blockchain ~	ess / Txn Hash Tokens 🗸	/ Block / Token / Resources ~	Ens More ~	9 JPHUA005	م ب
Sinance (1)	Exchange Mainnet L	aunched						Buy	• Exchan	ge 🗸 🛛 Earn 🗸	Gaming 🗸
Overview	V [ERC-20]				Profile Summa	ary [Edit]					More 🗸
PRICE \$621.75 @	@ 0.139881 Eth (+2.85%)		FULLY DILUTED MARKI \$10,308,307,778.64	ET CAP 🕲	Contract:		0xB8c77482	2e45F1F44dE174	5F52C7442	6C631bDD52	
Max Total	Max Total Supply: 16,579,517.0552533487 BNB ①			Decimals:		18					
Holders: 319,183 0.018%)			Official Site:		https://www.binance.com/						
Transfers		959,065			Social Profiles:		⊠ ⊡ f !	1 8 0			
Announcement: Binance Chain Mainnet Swap ×											
Transfers	Transfers Holders Info Exchange DEX Trades Contract Analytics Comments Q										
More tha (Showing th	an > 959,065 transac ne last 100k records)	tions found						FI	rst <	Page 1 of 4000 >	Last

Figure I. This figure presents screenshots of the BNB ICO token page on Etherscan.

	\mathbf{Ta}	bl	e I.	. N	fisre	present	tations	, regu	latory	news,	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 1(news) and news intensity, and their respective interactions with 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 variable 1(news) is an indicator that equals one if regulatory news is released within the calendar month prior to the first appearance of the ICO in our sample, and equals zero otherwise. The variable news intensity is the number of regulatory news articles released within the calendar month prior to the first appearance of the ICO in our sample. Section 3.2 of the main text contains variable definitions. Some models contain coverage-quartile fixed effects and are stratified by ICO cohorts. We include control variables found in Table 4 of the main text but suppress their estimated coefficients for brevity. Standard errors are clustered by ICO cohorts. t-statistics are reported in parentheses.

	(1)	(2)	(3)
$1(News) \times Misrep$	0.927	0.951	
	(0.69)	(0.43)	1
News intensity \times Misrep			1.030
Misrop	1 350	1 200	(0.40) 1.266
wistep	(2.87)	(1.34)	(6.70)
1(News)	1.308	0.769	(0110)
	(0.83)	(0.40)	
News intensity			0.933
			(0.43)
Controls	Y	Y	Y
# ICOs	$5,\!\bar{935}$	$5,\!\bar{935}$	5,935
Cohort strata	N	Y	Y
Coverage-quartile FE	N	Y	Y
Clustered SE	IN	Ŷ	Ŷ

Event:	ICO	scam
--------	-----	-----------------------

Table II. Misrepresentations and advisors' subsequent ICOs

This table presents estimates from Cox regressions. We estimate standard Cox hazard models in columns 1 and 2. In columns 3 and 4, we estimate Prentice, Williams, and Peterson Total Time models (extensions of the standard Cox model), which accommodate recurrent events. Estimated coefficients are expressed as hazard ratios. The failure event in these regressions is *subsequent ICO*. If one or more advisors of an ICO *i* subsequently advise another ICO *j*, then ICO *j* is a *subsequent ICO* of ICO *i*. The key independent variables in our regressions are *misrep* and 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 1(misrep > 0) equals one if the ICO has at least one *misrep*, and equals zero otherwise. Section 3.2 of the main text contains variable definitions. All models contain coverage-quartile fixed effects and are stratified by ICO cohorts. In columns 3 and 4, we further stratify models by the order of events. Standard errors are clustered by ICO cohorts. *t*-statistics are reported in parentheses.

Event: Subsequent ICOs				
	(1)	(2)	(3)	(4)
Consider recurrent events	N	Ν	Y	Y
Misrep	1.025		1.008	
_	(2.94)		(2.79)	
1(Misrep > 0)		1.172	()	1.090
		(2.42)		(1.75)
Banned	1.042	1.041	1.011	1.008
	(0.69)	(0.65)	(0.23)	(0.17)
Whitelist	0.947	0.948	0.971	0.970
	(1.26)	(1.18)	(1.10)	(1.19)
Duration	0.998	0.998	1.002	1.002
	(1.62)	(1.66)	(13.48)	(13.72)
Presale	1.210	1.206	1.081	1.078
	(3.46)	(3.69)	(2.91)	(2.36)
Hardcap	1.051	1.061	1.024	1.025
1	(1.03)	(1.16)	(0.76)	(0.81)
Softcap	1.025	1.023	1.007	1.006
1	(0.38)	(0.35)	(0.36)	(0.33)
Accept BTC	0.930	0.932	0.973	0.974
-	(1.39)	(1.33)	(0.95)	(0.94)
Accept ETH	1.095	1.102	1.102	1.103
*	(1.14)	(1.20)	(3.81)	(3.99)
Accept USE	1.136	1.134	1.014	1.014
*	(1.75)	(1.74)	(0.47)	(0.47)
Enforcement	0.960	0.957	0.972	0.970
	(1.06)	(1.12)	(0.75)	(0.75)
Disclosure	0.995	0.995	0.992	0.992
	(0.24)	(0.24)	(0.86)	(0.96)
SEC filing	1.199	1.191	0.959	0.957
0	(1.55)	(1.54)	(0.28)	(0.30)
# ICOs	9 971	9 971	9 971	9 971
Cohort strata	$\overset{\scriptscriptstyle{\scriptstyle \Delta},{\scriptstyle \Delta}}{\mathrm{V}}$	$\overset{\scriptscriptstyle 2,211}{\mathrm{V}}$	$\overset{2,211}{\mathrm{V}}$	$\overset{2,211}{\mathrm{V}}$
Event order strate	ı N	I N	I V	I V
Coverage quartile FF			I V	I V
Clustered SE	I V	I V	I V	
Olusiered SE	ĩ	I	I	ĩ

Event: Subsequent ICOs