DISCUSSION PAPER SERIES

DP16200
(v. 2)

CONFLICTED ANALYSTS AND INITIAL COIN OFFERINGS

Andreas Barth, Valerie Laturnus, Sasan Mansouri and Alexander F. Wagner

FINANCIAL ECONOMICS

CEPR

ISSN 0265-8003

CONFLICTED ANALYSTS AND INITIAL COIN OFFERINGS

Andreas Barth, Valerie Laturnus, Sasan Mansouri and Alexander F. Wagner

Discussion Paper DP16200
First Published 28 May 2021
This Revision 09 February 2023

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

CONFLICTED ANALYSTS AND INITIAL COIN OFFERINGS

Abstract

This paper studies the contribution of analysts to the functioning and failure of the market for Initial Coin Offerings (ICOs). The assessments of freelancing analysts exhibit biases due to reciprocal interactions of analysts with ICO team members. Even favorably rated ICOs tend to fail raising some capital when a greater portion of their ratings reciprocate prior ratings. 90 days after listing on an exchange the market capitalization relative to the initial funds raised is smaller for tokens with more reciprocal ratings. These findings suggest that the failure of ICOs is related to conflicts of interest.

JEL Classification: G14, G24, L26, D82, D83

Keywords: Analysts, Asymmetric information, Fintech, Initial coin offering (ico)

Andreas Barth - andreas.barth@finance.uni-frankfurt.de
Goethe University Frankfurt and SAFE

Valerie Laturnus - laturnus@finance.uni-frankfurt.de
Goethe University Frankfurt

Sasan Mansouri - mansouri@finance.uni-frankfurt.de
Goethe University Frankfurt

Alexander F. Wagner - alexander.wagner@bf.uzh.ch
University of Zurich, CEPR, ECGI, Swiss Finance Institute and CEPR

Conflicted Analysts and Initial Coin Offerings *

Andreas Barth, Valerie Laturnus, Sasan Mansouri and Alexander F. Wagner†

January 30, 2023

ABSTRACT

This paper studies the contribution of analysts to the functioning and failure of the market for Initial Coin Offerings (ICOs). The assessments of freelancing analysts exhibit biases due to reciprocal interactions of analysts with ICO team members. Even favorably rated ICOs tend to fail raising some capital when a greater portion of their ratings reciprocate prior ratings. 90 days after listing on an exchange the market capitalization relative to the initial funds raised is smaller for tokens with more reciprocal ratings. These findings suggest that the failure of ICOs is related to conflicts of interest.

JEL classification: G14, G24, L26, D82, D83

Keywords: Analysts, Asymmetric Information, Blockchains, Conflicts of Interest, FinTech, Initial Coin Offering (ICO)

---

*We are grateful to the editor (Bruno Biais) and three anonymous referees for excellent comments that have greatly improved the paper. We also thank Christian Leuz, Evgeny Lyandres, Arthur Petit-Romec, and David Yermack for helpful comments. Barth acknowledges financial support from the Chaire Fintech at University Paris Dauphine - PSL. Wagner acknowledges financial support from the University of Zurich Research Priority Program "Financial market regulation". This paper was initiated when Wagner was visiting the Center for Advanced Studies on the Foundations of Law and Finance at Goethe University Frankfurt. We declare that we have no relevant or material financial interests that relate to the research described in this paper.

†Barth, Laturnus and Mansouri are with Goethe University Frankfurt, Wagner is with University of Zurich, CEPR, ECGI and SFI.

✉: Andreas.Barth@finance.uni-frankfurt.de

✉: Laturnus@finance.uni-frankfurt.de

✉: mansouri@finance.uni-frankfurt.de

✉: alexander.wagner@bf.uzh.ch

1 Introduction

The question of how analysts contribute to the functioning of capital markets has been on the agendas of accounting and finance researchers for many years (Bradshaw et al., 2017). While professional analysts in traditional financial markets are heavily regulated, little is known about the role of freelancing analysts in unregulated financial markets.

This paper uses the setting of Initial Coin Offerings (ICOs) – an unregulated financial market that experienced a massive rise and fall in the late 2010s – to investigate determinants and consequences of the quantitative and qualitative aspects of investment ratings issued by human experts (henceforth referred to as ICO analysts). Strikingly, even among ICOs with an average rating in the top quartile, fewer than 50% succeed (in the sense of completing the token sale and collecting at least USD 1 in funding). Our analysis suggests that conflicts of interest in ICO analyst ratings can help explain the failure of ICOs. We find that ICO analysts tend to reciprocate favorable ratings for their own ventures; however, investors place lower emphasis on reciprocal ratings.

Initial Coin Offerings (ICOs) are token sale events on an own or existing blockchain that facilitate financing for an entrepreneurial venture. ICOs experienced an enormous boom in 2017-2018, but the volume of the market has declined massively since then. Token offerings are a potentially powerful instrument for new ventures to obtain crowdfunding-like resources (Goldstein et al., 2019; Li and Mann, 2020; Chod and Lyandres, 2021; Gryglewicz et al., 2021; Chod and Lyandres, 2022; Lee and Parlour, 2022; Lyandres et al., 2022). However, despite all the promises, the ICO market failed.

Understanding the workings and failures of this relatively new market and in particular studying the cross-section of ICOs and their analysts is of particular interest for at least three reasons. First, the ICO environment provides a relatively clean setup for investigating how analysts contribute to capital markets. The market is particularly interesting for a study of the role of information intermediaries because its regulation has only recently begun to clarify. Initial Exchange Offerings (IEOs) and Security Token Offerings (STOs) emerged recently as alternatives to ICOs. STOs need to be registered and approved by the U.S. Securities and Exchange Commission (SEC), for example, but like ICOs, they offer little investor protection. Understanding which ICOs failed despite potential monitoring by human professionals and the possible concomitant market discipline is important for clarifying the motivation for further regulation of these newer versions of FinTech markets.

Second, like financial analysts, ICO analysts potentially suffer from conflicts of interest.[1] However, the conflicts in this case (i) are potentially more extreme and (ii) can be more directly identified than in the case of the typical security analysts. As for (i), ICO analysts do not only *provide ratings* for ICOs, but may also *run their own* ICOs. Thus, whenever an ICO analyst i provides a rating for an ICO j, it is possible that he/she does so after a team member of this ICO j has previously rated an ICO of analyst i. As for (ii), most of the literature on financial analysts classifies analysts as "affiliated" (and thus potentially conflicted) if they belong to a bank that has or applies for an underwriting relationship with the firms on which they are reporting or if analysts want to get hired by the firm they analyze ("revolving door analysts"). These potential biases are largely hidden information, and particularly revolving door analysts can only be identified ex post their job change. By contrast, the ICO setting presents a situation where linkages are more direct and where investors can be aware of potential biases right away.

Third, non-professional analysts and their crowd forecasts have been shown to be important information intermediaries for equity investors (Chen et al., 2014; Jame et al., 2016; Drake et al., 2017; Campbell et al., 2019; Da and Huang, 2020; Farrell et al., 2020). However, we know little about the potential conflicts of interest that such analysts face and whether market participants consider the differential credibility and informativeness of these analyses in their investment decisions.

We collect data on 5,384 ICOs between 2017 and 2020 from the platform ICObench.com. We identify 539 experts who issued a total of 13,834 ratings.

[1]See, for example, Lin and McNichols (1998), Michaely and Womack (1999), and Chan et al. (2007) for evidence of biased financial analysts.

We begin by investigating determinants of analysts' ratings. We first find that analysts who, in the past, had issued very positive ratings for ICOs that did not succeed (i.e., analysts with large forecast errors) provide, on average, lower ratings in the future. ICObench.com also ranks ICO analysts, which gives an equivalent setting to all-star financial analysts (Leone and Wu, 2002). We observe that "star analysts" are less optimistic and their ratings are, on average, lower. In addition to quantitative ratings, we also consider the length and linguistic tone of the reviews that accompany the evaluation, i.e., the qualitative nature of ICO analyst ratings. We observe that lower ratings often accompany longer reviews with a more negative tone. In all of these analyses, we compare different ratings for the same ICO, which helps to rule out that these results are purely driven by the self-selection of analysts to certain ICOs.

Importantly, reciprocal ratings are special: the total rating score an analyst gives to an ICO j is higher if she received a rating in the past for her own ICO by any team member of coin j. This effect is stronger the higher the prior received rating was. These effects continue to hold when we compare analysts providing a rating to the same ICO in a given month. Comparing different assessments of the same analyst for virtually identical ICOs and different assessments for the same ICO by virtually identical analysts allows us to rule out that the assessment is due to the high difficulty of forecasting tasks or due to a non-random match between founders of good ICOs that also serve as analysts.

Next, we analyze the explanatory power of ICO analyst ratings on the outcomes of an ICO campaign. We first confirm the result of prior work that investors appear to value the fact that a human analyst provided a rating for the ICO. ICOs with any analyst coverage are more likely to complete the token sale and collect at least USD 1 in funding.[2] Moreover, a better average quantitative rating by human analysts translates into a higher probability that an ICO offering has been completed and received funding.

Interestingly, while the unconditional failure rate of ICOs is about 64%, even among ICOs with an average analyst rating in the top quartile 53.6% fail. Our main interest is in

[2] This is in line with several studies that document the benefits of financial analyst coverage (Sufi, 2009; Demiroglu and Ryngaert, 2010; Crawford et al., 2012; Mola et al., 2013).

the characteristics of analysts or of the ICO itself that lead to such disagreement between analysts' advice and the market outcome.

The share of reciprocal ratings is an important determinant of failure despite high ratings: if ICO j receives a rating from many reciprocal analysts, i.e., analysts whose rating is a response to a rating they received from a team member of ICO j, the market is more likely to disregard analyst recommendations. Moreover, even among successful ICOs, the market capitalization 90 days after listing on an exchange relative to the initial funds raised is smaller for ICOs with more reciprocal ratings. There are two possible interpretations of these results. First, it is conceivable that, even though we control for a wide variety of factors presumably capturing variation in ICO quality, reciprocal ratings occur with "objectively" bad ICOs; i.e., they pick up some additional variation in quality. Second, investors may trust ICOs with more reciprocal ratings less (even when they may potentially be worth funding).[3] Either way, the findings imply that investors do not blindly pile capital into highly rated ICOs.

Interesting patterns also emerge for the linguistic measures of the rating. The length and linguistic tone of the reviews that accompany the evaluation explain only little of the variation in the success of ICOs. However, the likelihood that an ICO will fail despite receiving high ratings increases with the positivity of the tone and complexity of the language in the reviews.

Finally, the quantitative and qualitative ratings by human analysts do not systematically differ on average for ICOs that prove to be fraudulent. A higher share of reciprocal ratings is not associated with a higher fraud probability, suggesting that criminal intentions do not typically drive reciprocity. ICOs exhibiting fraud do show a larger dispersion of both rating scores and rating review tones among analysts.

Overall, the results suggest that the failure of ICOs was not uniform but was related to measures of conflicts of interest. Having access to information about the track record and potentially conflicting activities of analysts allowed ICO investors to respond to quali-

[3]Several studies discuss whether or not investors are sophisticated enough to detect biased ratings (Ellis, 1998; Baker and Mansi, 2002; Livingston et al., 2010; Hirth, 2014; Badoer et al., 2019).

tative differences among analysts' ratings in a differentiated way. Even easier access would arguably have further enhanced efficiency of capital allocation in this market. Information intermediaries and platforms collecting data about crypto analysts play an important role in the functioning of market discipline in unregulated markets.[4]

These results add to the literature in four important ways. First, the literature on financial analysts suggests that a close link between analysts and firms leads to superior information and better assessments (Bae et al., 2008; Bradley et al., 2017), but also highlights the problem of conflicts of interest in a similar spirit of "affiliated" analysts (e.g. Lin and McNichols, 1998; Michaely and Womack, 1999; O'Brien et al., 2005; Malmendier and Shanthikumar, 2007; Agrawal and Chen, 2008; Kadan et al., 2009) or revolving door analysts (Lourie, 2019; Kempf, 2020).[5] However, there is scarce data on the direct interactions of analysts with the firms they analyze. The data on ICO analysts provide distinct advantages in that respect, and by showing that investors do take differences among analysts into account, we highlight that these data are of value to investors.

Second, the paper complements the literature on semi-professional analysts in equity markets (Chen et al., 2014; Drake et al., 2017). That literature recognizes the possibility of conflicts of interest if the semi-professional analyst is holding positions on the stock themselves, resulting in a subjective, distorted analysis (Campbell et al., 2019).[6] While these studies focus on equity markets in which semi-professional analysts complement the information produced by professional analysts, one particular advantage of the ICO market, besides very detailed and structured information, is the absence of professional analysts.[7]

Third, the paper adds to the growing literature on the relationship between machine-

[4]Transparency about the background, characteristics, or track records of information providers and intermediaries has been identified as critically important in other settings. For example, Law and Mills (2019) highlight the importance of the transparency provided by the Financial Industry Regulatory Authority (FINRA) about brokers' (criminal) backgrounds.

[5]A similar conflict of interest is present for rating agencies (e.g. Bolton et al., 2012; Bar-Isaac and Shapiro, 2013; Baghai and Becker, 2017; Chu and Rysman, 2019).

[6]Campbell et al. (2019) use non-professional analysts' disclosures of stock positions as an indicator of the analyst's position, which may not be reported truthfully.

[7]There are of course many stocks that professional analysts do not cover. This lack of coverage is the analysts' choice, however, and as such provides information to the market.

generated evaluations and human expert ratings.[8] In addition to human evaluations, many platforms set up machine-generated ratings. These ratings do not evaluate the content of an ICO, but are based on observable factors such as features of the ICO's campaign and team.[9] Importantly, we show that both ratings are informative regarding ICO success. However, many ICOs fail despite high ratings by human analysts, which is why we analyze this discrepancy.

Finally, ICOs are (or were) a potentially powerful way to fund new ventures, not least because of the underlying distributed ledger-based technology and the platform's special features (Bakos and Halaburda, 2019; Biais et al., 2019; Cong and He, 2019; Easley et al., 2019; Hinzen et al., 2022). This paper advances our knowledge of the failure of the ICO market. Usually, the sales of tokens or ICOs appear at a very early planning stage of a product or firm's life cycle and suffer from severe information asymmetry and adverse selection problems (Malinova and Park, 2018; Gan et al., 2020; Chod and Lyandres, 2021; Chod et al., 2022). As such, tokens have no intrinsic value at the time of the investment. Instead, they derive value from trust in future usage (Conley, 2017). Hence, the literature has investigated both the supply side, i.e., choices by ICO entrepreneurs (Adhami et al., 2018; Amsden and Schweizer, 2018; Deng et al., 2018; Ernst and Young, 2018; Cerchiello et al., 2019; Davydiuk et al., 2019; Fisch, 2019; PwC, 2019; Chakraborty and Swinney, 2020; Howell et al., 2020; Roosenboom et al., 2020; Benedetti and Kostovetsky, 2021), and the demand side, i.e. choices by investors (Fisch et al., 2019; Fahlenbrach and Frattaroli, 2020; Fisch and Momtaz, 2020).[10] Little attention has been paid to the information providing

[8]For example, Aubry et al. (2019) use data on paintings auctioned to study the accuracy and usefulness of valuations generated by using a pricing algorithm based on neural networks. With data from a leading startup accelerator, Catalini et al. (2018) show that artificial intelligence can help humans to screen and evaluate information when there is an information overload.

[9]Automated algorithms that simply count disclosed information are usually applied. For example, a high number of social media platforms on which an ICO is present or being listed on several rating websites automatically improves the rating for the respective ICO (Boreiko and Vidusso, 2019).

[10]There is also literature on the price dynamics of tokens (Li and Mann, 2020; Cong et al., 2021, 2022; Lee and Parlour, 2022) as well as studies of asset pricing properties of coins on secondary markets and post-ICO performance (Dittmar and Wu, 2019; Hu et al., 2019; Fisch and Momtaz, 2020; Lyandres et al., 2022). See Li and Mann (2021) for a review of recent literature advances in ICO research.

intermediaries between supply and demand, however, and the literature largely focuses on the governance role of whitepapers provided by the ICO team (Adhami et al., 2018; Feng et al., 2019; Giudici and Adhami, 2019; Zetzsche et al., 2019; Zhang et al., 2019; Samieifar and Baur, 2020; Florysiak and Schandlbauer, 2022).

To the best of our knowledge, only three previous papers examine ICO analysts (Aggarwal et al., 2019; Bourveau et al., 2022; Lee et al., 2022). All three papers document that ICOs with higher expert assessments are more successful. Closest to our work is Bourveau et al. (2022), who emphasize the positive role of information intermediaries to gauge ICO quality. Importantly, the information intermediary as defined in Bourveau et al. (2022) combines the machine-generated rating of disclosure quantity with human evaluations, whereas we differentiate between the two. As in Bourveau et al. (2022), we observe that both the ratings by human analysts and the machine-generated rating Benchy are predictive (though using a somewhat larger set of controls). We focus in particular on human analysts and the striking fact that more than 50% of the ICOs with the highest quartile of *human* ratings fail. We show that accounting for the heterogeneity among analysts is important. In particular, we exploit the specific feature of the market that ICO analysts provide ICO ratings, often while also running their own ICOs. We show that reciprocal ratings are biased, but also that investors discount such reciprocal ratings. In sum, our analysis highlights that, while analysts may provide important information, their information provision is subject to conflicts of interest.

The rest of the paper is organized as follows. Section 2 presents the data and descriptive statistics. Section 3 describes the results, and Section 4 concludes.

2 Data and descriptive statistics

2.1 Sample and data source

We collect data on ICOs, ICO ratings and ICO experts from the platform ICObench.com. Our sample consists of 5,384 ICOs (of which 2,378 were rated by at least one expert, and which thus constitute our main sample) and spans the time period from the start of ICObench.com in 2017 to February 2020.[11] According to the web traffic statistics from Alexa Internet, ICObench.com was an important source of rating information for investors, and was able to achieve a site visit rank of 3,644 during the peak of the ICO market (compared to a site visit rank of around 2,200 for the Financial Times in the same period).[12] ICOs in our sample were launched in 127 different countries, of which the USA, Singapore and the UK have the highest market shares.[13]

2.2 ICO analysts

In contrast to regulated financial analysts, ICO analysts are not certified. However, they have to apply for expert status on a platform, in our case ICObench.com. In their application, experts are required to describe their level of experience in crypto assets and motives to rating ICOs. The platform confirms the analysts after reviewing their credentials. The selection is relatively stringent. As of March 2020, the ICObench.com platform hosts more than 111,000 community members of which only 539 have expert status and thus the ability to provide ratings.

ICObench.com ranks the analysts based on several factors, including profile completeness and analysts' consistency in providing contributions to the platform.[14] This in turn provides an analogy to the widely used all-star rankings of financial analysts. We collect these rankings over time and flag whether an analyst is among the top 30 analysts, i.e., within approximately the top 5%. The dummy variable $StarAnalysts_{ij}$ equals one if analyst i is listed among the

[11] All ICOs in our sample are utility tokens with or without launchpad.

[12] Alexa Internet identifies "ICO rating" as the main 'Buyer Keyword' for ICObench.com, that is, those people who were searching in order to buy a product or service and landed on ICObench.com had searched primarily for "ICO rating".

[13] The compiled dataset is of comparable size to data used in other empirical ICO studies. For example, Benedetti and Kostovetsky (2021) use a sample of 2,390 ICO campaigns. Florysiak and Schandlbauer (2022) analyze 2,665 ICOs. Recently, Lyandres et al. (2022) cover the largest data set from the ICO universe with 5,450 ICO projects merged from various websites. Note that our sample period also covers the time after the collapse of the ICO market.

[14] The expert weight is calculated based on a profile score, a rating score, a time score, an acceptance score, and a contribution score. See https://icobench.com/faq for a detailed description.

top 30 list prior to evaluating ICO j.

Interestingly, many ICO analysts are involved in one or more ICO campaigns themselves.[15] Section 2.4 elaborates on how we empirically exploit this unique setting.

2.3 Ratings

We identify 539 experts on ICObench.com who rated 2,378 ICOs. Each analyst rated an average of 29.64 ICOs, resulting in 13,834 ratings overall. Experts can provide a rating for three dimensions of an ICO - team, vision and product - with each dimension being scored from 1 (poor) to 5 (best). The $AnalystRating_{ij}$ of analyst i for ICO j is defined as the sum of these three individual ratings,

$$\begin{aligned} AnalystRating_{ij} &= AnalystRating(Team)_{ij} + AnalystRating(Product)_{ij} \\ &\quad + AnalystRating(Vision)_{ij}, \end{aligned}$$

i.e., an integer in the interval $[3, 15]$.

For all ratings, we collect the date when the analyst issued the rating. The main analysis only considers ratings issued before ICO completion (or cancellation), which helps prevent look-ahead bias. However, our findings also hold when we include all ratings. Analysts have the opportunity to modify their ratings: when this happens, users can only see the updated rating score as well as two dates - the date of the first rating and the date of the update, but not the full history.[16] This paper considers the modification date as the date for the rating and flags a modified rating by analyst i to ICO j with a dummy variable $Modified_{ij}$.

The information about the timing of the rating allows us to construct for each rating by analyst i to ICO j a measure of the rating experience for the analyst up to this rating of ICO

[15] Note that if ICO analysts become part of the ICO project by advising the team members, they lose the ability to rate their own ICO. We found that 4 analysts rated an ICO project before becoming an advisory team member.

[16] There is a well-documented phenomenon of "walking down" forecasts in the literature on sell-side analysts. The absence of access to the rating history of analysts on ICObench.com prevents us from studying this phenomenon in the ICO context.

j, $AnalystExperience_i^{j-1}$, which is defined as the natural logarithm of 1 plus the number of ICOs that analyst i rated before providing a rating for ICO j.

When issuing a rating, the analyst gives a score and typically justifies the decision by writing a review. We collect all reviews and calculate linguistic measures from these texts. Based on the Loughran and McDonald (2011) dictionaries, we calculate the tone of the language, defined as the difference between positive and negative words to total words, as well as the uncertainty of the language, defined as the count of uncertain words divided by total words. We further control for the complexity of the reviews, measured by the Gunning (1952) Fog index, which is a function of the number of words per sentence (length of a sentence) and the share of complex words (words with more than two syllables) relative to total words.[17]

For some analyses, we aggregate the analyst-ICO information to the ICO level. More precisely, we count the total number of analysts who rated ICO j in $\# Analysts_j$. We further aggregate all analyst ratings for ICO j in the variable $AnalystRating_j$ by averaging all ratings that ICO j received from all analysts that cover this ICO. Finally, we proxy the lack of consensus among analysts that provide a rating for ICO j with $AnalystDispersion_j$, defined as the standard deviation of all ratings for ICO j. To maximize sample size, we set analyst dispersion equal to zero if there is only one rating. However, our main results on the role of reciprocal ratings do not depend on this choice.

Figure 1 presents the number of ratings in a given month over time of the newly announced ICOs, the number of ratings by analysts who registered in the same month, and the Bitcoin price in US dollars. While the number of new ratings went up hand-in-hand with the number of ICOs to the peak of Bitcoin's price in January 2018, the number of ratings exploded thereafter and has only recently converged again to the number of announced ICOs. Figure 1 further shows that the surging demand for information about crypto assets was met by an increase in the supply of analysts.

[17] The Gunning (1952) Fog index is defined as $Fog = 0.4 \cdot \left(\frac{TotalWords}{TotalSentences} + \frac{ComplexWords}{TotalWords} \right)$.

[Figure 1 about here]

Figure 2 shows the monthly averaged *AnalystRating* as well as analysts' rating dispersion, measured by the standard deviation of ratings within an ICO in a given month. We observe that the average total rating of experts was overall very positive with a small decrease in the rating score around the Bitcoin price drop in 2018. Analyst dispersion remained at a relatively constant level over the sample period. It only slightly increased around the time when the Bitcoin price was low at the end of 2018, but decreased as the Bitcoin price rose again at the end of 2019.

[Figure 2 about here]

Complementing the assessment of human experts, many platforms have set up machine-generated ratings. Instead of evaluating an ICO's quality directly, these ratings are based on the availability of information *about* the ICO. The idea is that more transparency indicates higher trustworthiness and quality of the ICO. Importantly, the machine-generated rating does not include any human assessment. For every ICO in our database, we collect the machine-generated rating by ICObench.com, which is called "Benchy". The Benchy bot provides a higher rating for higher transparency on team and event information. Moreover, Benchy uses factual data, such as "presence of the social media links" and "the level of activity on them", see `https://icobench.com/faq`. Benchy re-evaluates each ICO profile at least once daily and issues a rating ranging between 1 (poor) and 5 (best). Only the most recent evaluation is observable, not the history of Benchy ratings.

While all ICOs listed on the platform ICObench.com automatically receive a machine-generated rating from the Benchy bot, 2,378 out of 5,384 ICOs listed on this website were also rated by ICO analysts. On average, the ICOs with(out) an analyst rating have a Benchy rating of 3.2 (2.7) out of 5.

2.4 Reciprocal ratings

A specific feature of the market is that ICO analysts also participate in ICOs. We identify those experts that are involved in one or several ICO projects by collecting each expert's self-description of experiences and achievements from the 'About'-section of their profile pages on ICObench.com. Table 1 shows the distribution of ICO projects among analysts. Of the 539 experts in our sample, 319 have been involved in at least one ICO, with some analysts being very active in launching ICOs.

[**Table 1 about here**]

We use this information to flag whether a rating of analyst i for ICO j is a response to a rating that analyst i received for an ICO they were involved with from any team member of ICO j at any prior point in time. We generate the indicator variable $ReciprocalRating_{ij}$ as follows:

$$ReciprocalRating_{ij} = \begin{cases} 1, & \exists\ AnalystRating_{j'i'} \text{ before } AnalystRating_{ij} \text{ where } \mathbf{i} \in \Omega_{i'},\ j' \in \Omega_{\mathbf{j}}, \\ 0, & \text{else} \end{cases}$$

where $\Omega_{\mathbf{j}}$ refers to the set of all team members of the ICO j. Table 2 represents a hypothetical illustration of how we define this variable. $ReciprocalRating_{ij}$ thus flags whether any member of ICO j has provided a rating of any ICOs with which that expert i is associated.[18] Reciprocal ratings are not directly flagged by ICObench.com, but users can easily obtain the information given the available links to each analyst's associated ICOs and the timeline of the ratings provided on ICObench.com.

Whenever $ReciprocalRating_{ij}$ indicates reciprocity, we additionally identify the level of the rating that is reciprocated, i.e., the $AnalystRating$, as well as the three components

[18] One might argue that an analyst rating is already reciprocated if it is issued in expectation of receiving a rating for the analyst's own *future ICOs*. Further below, we employ a modified version of our reciprocal dummy, which is equal to one if an analyst launches his/her own ICO at a later stage, i.e., if the analyst may expect to receive a reciprocal rating in the future, and zero otherwise.

$AnalystRating(Team)$, $AnalystRating(Vision)$ and $AnalystRating(Product)$ that the ICO with which expert i is associated previously received from any member of ICO j. The level of the reciprocated rating is labeled $ReceivedRating_{ij}$.

[**Table 2 about here**]

2.5 ICO outcome variables

We consider multiple measures of ICO outcomes, some related to the initial completion, some related to medium-term performance. Specifically, as a short-term measure for ICO success, we construct a dummy variable *Success*, which takes the value of 1 if the ICO-related coin successfully completes the offering and receives funding. For these ICOs, we collect information on the dollar amount raised during the campaign from ICObench.com, ICOmarks.com, tokendata.io, and ICOdata.io. Tokens were classified as failed when we could not find the amount raised nor any information indicative of success on the above-mentioned web pages. In total, we identify 1,932 successful ICOs among our 5,384 ICOs.

Figure 3 shows the time trend of successful ICOs. ICOs became popular at the beginning of 2017. While only 29 ICO tokens were on sale before then, the number increased to 1,127 ICOs within one year with around 94 offerings per month and a 53% success rate. The market peaked in 2018, with 3,360 ICOs in total and a success rate of 33%. In 2019, around 64 ICOs were sold per month, of which 25% were successful on average. Thus, the flow of ICOs continues, albeit at a lower level, even after the sharp decline of cryptocurrency prices and the corresponding decline in enthusiasm towards ICOs.

[**Figure 3 about here**]

In addition, we generate a more medium-term success measure, which we label *Market-Performance*, defined as the market value of the token 90 days after its listing on an exchange (from CoinMarketCap.com) over the initial amount raised by the ICO, expressed in percent. We observe the market capitalization information only for a subset of ICOs in our sample, either because CoinMarketCap.com does not cover the exchange on which the token was listed or because the project failed. Therefore, we either use the market capitalization 90 days post exchange listing for the restricted sample of successfully listed ICOs, or set the market capitalization to zero for projects that raised funding during the campaign but without any information on CoinMarketCap.com, assuming a failure of these projects (Howell et al., 2020).[19]

Finally, we also collect information about scams, i.e., ICOs that were launched with the intention of defrauding investors. To do so, we use the marker 'Scam or Other Issues' for dead coins listed on Coinopsy.com, as well as information from Deadcoins.com, a message board where users post about scams. Some of these ICOs can also be found in U.S. Securities and Exchange Commission (SEC) press releases, especially when they fine ICO companies for fraudulent practices.[20] With this (likely conservative) method, 234 ICOs were flagged as scams in our data.

2.6 Forecast errors

Combining the ICO success variable and the analyst rating score allows us to construct an ex-post forecast error measure for each rating. As the outcome of an ICO is either success or failure, we define the forecast error of a rating as the distance to the highest (lowest) possible rating in case of success (failure):

$$ForecastError_{ij} = \begin{cases} 15 - AnalystRating_{ij}, & \text{if ICO succeeded} \\ AnalystRating_{ij} - 3, & \text{if ICO failed} \end{cases}$$

For example, if analyst i gives a successful ICO j a total rating of 15, their rating was fully precise, resulting in a $ForecastError_{ij}$ measure of 0. If that same ICO had failed, however,

[19] Note that the medium-term success measure is only available for ICOs with a non-zero amount raised during the campaign by construction. If an ICO raises funding during the campaign but fails within the first 90 days, this results in a $MarketPerformance$ of zero.

[20] See, e.g., `https://www.sec.gov/news/press-release/2019-259`.

the forecast error of this rating would flag a 12.[21]

Figure 4 shows the monthly averaged $AnalystRating$, the average forecast error and the number of successful ICOs (as a share of total ICOs) over time. In addition, we plot the monthly average forecast error separately for ratings where analysts were too optimistic and too pessimistic, respectively.[22] Interestingly, ratings become less precise over time. This is driven by overly optimistic analysts.

[**Figure 4 about here**]

In our regression analysis, we use an analyst-specific measure of the forecast error that takes the entire history of an analyst's ICO-specific $ForecastError_{ij}$ into account. We recursively average the $ForecastError_{ij}$ of analyst i over all of their issued ratings up to ICO j using an expanding window. We denote this variable $ForecastError_i^{j-1}$.

2.7 ICO characteristics

For every ICO in the sample, we collect data on the campaign characteristics that have been found in the literature to indicate the perceived quality of an ICO by investors (Howell et al., 2020; Bourveau et al., 2022; Lyandres et al., 2022). For many characteristics, we generate binary indicators that flag whether an ICO exhibits the respective feature. The dummy variable $Presale$ equals one if an ICO offers coins at the pre-sale stage and zero otherwise. The $Bonus$ and $Bounty$ dummies equal one if there were discounts on the token sale or incentives to boost social media presence, respectively. The dummy MVP flags the availability of a minimum viable product or whether a product prototype was in place. The dummy KYC equals one if investors need to validate their identity by signing up to a

[21]While this $ForecastError$ measure is not immediately available to investors on ICObench.com, one can easily view the entire timeline of an analyst's ratings with a link to detailed information on the rated ICO.

[22]An analyst is too optimistic (pessimistic) if their rating of a failing (successful) ICO was larger (smaller) than 3 (15). Thus, we calculate the monthly average optimistic (pessimistic) forecast error over the optimistic (pessimistic) cases of the $ForecastError_{ij}$ definition.

whitelist to access the token sale.[23] The dummy IEO indicates the use of a centralized token launch platform provided by a cryptocurrency exchange. The variable $HardCap$ equals one if the ICO discloses a maximum amount that the team is planning to raise. $VestingDisclosure$ is a dummy variable that flags one if the ICO provides vesting information in the whitepaper. $Facebook$ and $Bitcointalk$ are dummy variables that equal one if the ICOs generated a website on Facebook and Bitcointalk. The $RetentionRatio$ is the percentage of tokens that is retained by the team members; it captures the "skin in the game" of ICO team members. Similar to the ICO literature, we control for the overall advancement of the project by $GitHubCommits$, i.e. the number of code revisions that ICO team members have saved on GitHub.com. Finally, we also collect information on the number of team members and advisors of each ICO project.

In addition to these $VentureOffering\ Controls$, we derive useful information ratios from whitepapers and user posts on the social media platform Bitcointalk.org to proxy for the informative value in ICO whitepapers. By using the Loughran and McDonald (2011, 2020) and Lyandres et al. (2022) dictionaries, we compute the tone, the level of linguistic complexity, technology, and uncertainty in ICO whitepapers. Moreover, we include $WhitePaperLength$, the natural logarithm of (1 + total words of the whitepaper) to our regression models, and summarize all these variables as $WhitePaper\ Controls$.

We also repeat the textual analysis and construct the textual information ratios for all text messages on Bitcointalk.org published between the initial announcement and the end date of the ICO event. We add another dictionary-based ratio which measures extreme language usage on social media,[24] and classify these variables, together with the number of text messages and their length (1 + total words of all text messages on the Bitcointalk.org webpage of ICO j), as $SocialMedia\ Controls$.

[23]Note that ICObench.com provides information on two different KYC procedures. One KYC symbol means the identity verification of ICObench.com profiles, while the second flags the identification and registration process of investors to receive access to the token sale. We use the second piece of information on KYC throughout the paper.

[24]The dictionary for extreme language is taken from Bochkay et al. (2020).

Finally, we collect the year-month information of the date when the ICO was launched and calculate the BTC return during the campaign of the ICO as a proxy for the overall market sentiment.

2.8 Descriptive statistics

Table 3 shows descriptive statistics of the key variables of rating and ICO characteristics. All variables are defined in Table A1.

[**Table 3 about here**]

In our sample, the average *AnalystRating* is 11, where the *AnalystRating*(*Product*) is slightly more pessimistic than the evaluation of the other two dimensions $Team$ and $Vision$. Of all ratings, 12% are flagged *ReciprocalRating*, which are found to be somewhat more positive with an average $ReceivedRating_{ij}$ of 13. We observe a success rate of 36%. In terms of dollar amount raised, EOS, Telegram, and Bitfinex were the most successful ICOs in our sample. On average, ICOs in our sample have a market performance of 86% of the initial dollar amount raised 90 days after listing on an exchange.The scam rate is 4.3%. Each ICO is covered by 2.6 analysts, on average, and 44% of all ICOs are covered by at least one analyst. The ICOs for Sharpay (94), Truegame (82), and WePower (64) had the largest number of analysts covering them.

3 Empirical Analysis

Section 3.1 analyzes the determinants of ICO analyst ratings. Section 3.2 considers whether investors consider differences in the reliability of analyst ratings.

3.1 What determines analyst ratings?

3.1.1 Baseline results

We model the rating of analyst i for ICO j as a function of analyst characteristics, as indicated in the following equation:

$$\begin{aligned} Rating_{ij} = \beta_0 + \beta_1 \cdot Benchy_j + \beta_2 \cdot StarAnalysts_{ij} + \beta_3 \cdot ForecastError_i^{j-1} \\ + \beta_4 \cdot Modified_{ij} + \beta_5 \cdot X_j + Month_{ij} + \alpha_j + \epsilon_{ij}, \end{aligned} \tag{1}$$

$Rating_{ij}$ denotes the respective rating score that analyst i gives to ICO j for the different rating categories team, vision and product (on a scale from 1–5), as well as the total rating score ($AnalystRating_{ij}$) as the sum of the three categories (on a scale from 3–15). The vector X_j contains the ICO characteristics as described in Subsection 2.7. Time trends of ratings are absorbed by $Month_{ij}$ dummies. α_j denotes ICO fixed effects. We allow for a potential serial correlation of ratings within each analyst and within each ICO and employ two-way clustering of standard errors (Cameron et al., 2011) at the analyst and ICO dimensions.

[Table 4 about here]

Table 4 summarizes the results of this analysis. Column (1) shows that machine-generated and human expert ratings point in the same direction, i.e., ICOs with higher machine-generated ratings receive a higher rating score by human analysts on average. Moreover, in the cross-section of analysts, columns (2) to (4), we find a statistically significant negative coefficient on $ForecastError_i^{j-1}$, implying that analysts with historically higher forecast errors give on average lower ratings. The negative relationship between the past forecast errors and the rating also remains in the within ICO estimation, as column (5) shows. Analysts listed within the top 30 analysts on ICObench.com are more critical and issue lower ratings on average.

Furthermore, in line with the literature, the coefficients of the control variables suggest that analysts consider the characteristics of the underlying ICO (Deng et al., 2018; Roosenboom et al., 2020; Bourveau et al., 2022). In general, we find that ICOs with a pre-sale event, with a KYC feature and an IEO feature receive better ratings. Moreover, analysts perceive it as a good signal when founders retain a higher share of the tokens themselves, when projects have many advisors and team members and when the ICO whitepaper contains many technical words.

3.1.2 Reciprocal ratings

When ICO analysts issue new ratings, are they based on ratings that their own affiliated ICOs previously received? While descriptive evidence as shown in Figure 5 points to this, we aim to answer this question more formally by running regressions, as specified in the following equation:

$$
\begin{aligned}
Rating_{ij} = \beta_0 + \beta_1 \cdot ReciprocalRating_{ij} + Analyst \times Month_{\ ij} \\
+ ICO \times Month_{\ ij} + \epsilon_{ij},
\end{aligned}
\tag{2}
$$

where $ReciprocalRating_{ij}$ indicates a dummy that flags whether analyst i received a rating from a team member of ICO j. We include $Analyst \times Month$ and $ICO \times Month$ dummies to exploit only the analyst and ICO pairing within the month of the rating. These fixed effects detect the variation previously established in Table 4, and they help to rule out that the results were driven by a non-random match between founders of good ICOs that also serve as analysts. Comparing different assessments of the same analyst for virtually identical ICOs and different assessments for the same ICO by virtually identical analysts allows us to differentiate between whether analysts behave in a deliberately optimistic, biased manner, or whether the optimistic assessment is due to the high difficulty of forecasting tasks. For reciprocal ratings, we also analyze whether the level of the prior rating predicts the level of the reciprocal rating.[25] We again employ two-way clustering at the analyst and ICO

[25] We use $Analyst \times Quarter$ and $ICO \times Quarter$ dummies for these analyses because of the restricted sample size.

dimensions.

[Table 5 about here]

Table 5 shows that ratings do indeed contain a reciprocal element. Column (1) indicates a positive association between the total rating an analyst gives to an ICO and the $ReciprocalRating_{ij}$ dummy. More specifically, the total rating score is around 0.25 points higher when the analyst is in a position to respond to a prior rating. Additionally, within the sample of reciprocal ratings, column (2) shows that ratings are more positive the higher the previously received rating was. In other words, analysts reciprocate positive ratings. For example, column (2) shows that each one-unit (one standard deviation) increase for the previous rating leads an analyst to issue a total rating of around 0.08 (0.14, or 6% of a standard deviation) higher rating. Note that this result holds within ICO-time and analyst-time combinations, i.e., comparing ratings by two otherwise identical analysts, where one analyst previously received a rating by a team member of coin j and the other one did not. This reciprocal rating behavior is similar to the *quid pro quo* between hedge funds and sell-side equity analysts described in Klein et al. (2019).

Columns (3)–(8) analyze the three different rating categories team, vision and product separately. The coefficient for the $ReciprocalRating_{ij}$ dummy is positive and significant for all three categories, indicating that, on average, analysts give a higher rating for the team, vision and product of ICO j if any team member of ICO j rated them. However, the actual score is significant only for the team dimension. That is, an analyst rates the team component of ICO j more highly if they received a more favorable team rating from a team member of ICO j. These findings are intuitive, as the team category constitutes a "soft factor". The results also indicate the relatively personal nature of the reciprocity.

3.1.3 Linguistic characteristics of rating reviews

When issuing ratings, analysts often justify the rating scores with written reviews. We next analyze whether more optimistic ratings are special in terms of the linguistic nature of the written review. The literature on earnings conference calls uses the number of words spoken by analysts as a proxy for the question difficulty, so analysts who ask lengthier questions are regarded as more critical (Merkley et al., 2017). Correspondingly, we investigate whether the rating score correlates with the length of the written text or with the linguistic tone of the review. Moreover, we investigate whether the relationship between the rating score and the review length and tone differs for reciprocal versus non-reciprocal ratings. This idea follows Cohen et al. (2020), who document that biased analysts ask easier questions. We run regressions for the overall sample as well as for reciprocal and non-reciprocal ratings separately as specified in the following equation:

$$\begin{aligned} Linguistic\,Measure_{ij} = \beta_0 + \beta_1 \cdot TotalRating_{ij} + Analyst \times Time_{\ ij} \\ + ICO \times Time_{\ ij} + \epsilon_{ij}, \end{aligned} \tag{3}$$

where $Linguistic\,Measure_{ij}$ interchangeably indicates the length of the rating review measured by the (natural logarithm of the) number of words and the ratio of positive words minus negative words to total words in the review. As before, we employ two-way clustering by analysts and ICOs.

[**Table 6 about here**]

Table 6 shows the results. In Panel A column (1), we find a negative relationship between the rating score and the length of the review, suggesting that more negative ratings come with a more detailed explanation. In column (2), we include $Analyst \times Month$ and $ICO \times Month$ dummies to rule out the possibility that the results have been driven by a non-random match between analyst characteristics (e.g. mood) and the quality of the rated ICO. For the review tone in Panel B, columns (1) and (2), we find that analysts use more positive terminology when reviewing an ICO that they score higher.

When investigating whether the relationship between the rating score and the review length and tone differs for reciprocal versus non-reciprocal ratings, we find that lower rating scores are justified with even lengthier reviews for reciprocal ratings, with a statistically significant difference to the coefficient for non-reciprocal ratings.[26] The relationship between review tone and rating score does not differ noticeably between reciprocal and non-reciprocal ratings.

3.1.4 Order of ratings

The literature on security analysts has documented herding behavior among analysts and shows that their buy or sell recommendations have a significant positive influence on subsequent analysts' recommendations (Welch, 2000). Thus, reciprocal analysts' scores may impact investors as well as other analysts when they cover the ICO at an early stage. Therefore, we analyze whether analysts provide reciprocal ratings faster and move earlier for ICOs where they issue more positive ratings. We generate a variable that counts the rank of rating arrival per ICO j from analyst i, i.e., whether analyst i was the first, second, third, or ... last analyst who rated for ICO j. We relate the order of the rating coverage to the $ReciprocalRating_{ij}$ dummy, as indicated in the following equation:

$$\begin{aligned} OrderRank_{ij} = \beta_0 + \beta_1 \cdot AnalystRating_{ij} + \beta_2 \cdot ReciprocalRating_{ij} + \beta_3 \cdot StarAnalysts_j \\ + \beta_4 \cdot ForecastError_i^{j-1} + Month_{ij} + \alpha_i + \alpha_j + \epsilon_{ij}. \end{aligned} \quad (4)$$

We again absorb any ICO and analyst characteristics with ICO and analyst fixed effects and control for time trends by $Month_{ij}$. As before, we use two-way clustered standard errors at the analyst and ICO dimensions.

[Table 7 about here]

The results are shown in Table 7. In line with the literature on analyst coverage of stocks (Demiroglu and Ryngaert, 2010), we first find that analysts who give favorable ratings tend

[26]Because of the restricted sample size for the sample of reciprocal ratings, we use *Analyst* $\times$ *Quarter* and *ICO* $\times$ *Quarter* dummies for these analyses.

to issue their ratings early. Second, star analysts tend to move first and rate the same ICO earlier than their less experienced peers. Third, reciprocal ratings tend to be issued early. In particular, in the chronological sequence of ratings given to an ICO j, a reciprocal analyst appears to issue their rating, on average, 1.3 positions earlier than a non-reciprocal analyst.

3.2 Are ICOs with higher ratings more successful, and which ICOs fail despite high ratings?

So far, we have established several important determinants of ICO analyst ratings, with analyst-specific factors such as prior forecasting ability and reciprocal status playing a major role, in addition to objective differences among the ICOs. Now, we investigate whether ICOs with higher ratings are indeed more successful. First, we establish baseline results for (unconditional) ICO success, but our main interest is in explaining when investors deviate from the ICO analyst consensus, that is, the ICO success probability conditional on an extreme positive (or negative) rating outcome. We also consider whether the factors that explain such deviations predict scams.

3.2.1 Ratings and ICO success

Table 8 presents descriptive statistics for the relationship between ratings and ICO success. Panel A indicates that an ICO is more likely to be successful when it motivates analysts to rate it.[27] In Panel B, we tabulate success statistics for groups of the quantitative rating score. The probability of receiving funding, the market capitalization 90 days after listing divided by the capital raised (in percent), and the average dollar amount raised, are higher for ICOs with more positive ratings, though the relationship is not strictly monotonic.

While these results highlight that successful ICOs have higher ratings on average, there

[27] There are a few ICOs that were extremely successful in terms of market performance (most of them without analyst coverage). All of our results hold if we trim our data at the 99^{th} percentile of market performance, which corresponds to excluding ICOs with a market growth in the first 90 days of about 1500% of the amount raised.

are numerous cases in which ICOs were either unsuccessful despite positive ratings or successful despite negative ratings. To quantify this phenomenon, we define for each ICO j a $Disagreement_j$ dummy as a conditional success outcome. More precisely, the $Disagreement_j$ dummy equals one if (i) the average $AnalystRating_j$ of an ICO is greater than or equal to 13 but the ICO is unsuccessful, or (ii) the average total rating is less than or equal to 5 and the ICO is successful. In our sample, this $Disagreement_j$ dummy is one in 413 of 2,378 rated ICOs (17%). While the unconditional failure rate of ICOs is about 64%, even among ICOs with an average rating in the top quartile 53.6% fail.[28]

These mismatches between ratings and ICO success do not occur randomly. To illustrate, in Panel C, we tabulate the disagreement dummy against the occurrence of reciprocal ratings. We observe that the ICO outcome is less likely to correspond to what one would expect given the ratings level if reciprocal analysts cover the ICO. ICOs that receive very favorable recommendations fail much more frequently if the reciprocal rating share is positive than if none of the ratings is reciprocal. Moreover, Panel C also shows that the market performance 90 days after listing, is lower for ICOs with a reciprocal rating.

[**Table 8 about here**]

In order to formally analyze ICO success in a regression framework, we first explain the unconditional success of ICO j using characteristics of participating analysts and a large set of ICO characteristics in a logit regression:

$$\begin{aligned} Success_j = \beta_0 &+ \beta_1 \cdot AnalystRating_j + \beta_2 \cdot \#\, Analysts_j + \beta_3 \cdot ReciprocalRatingShare_j \\ &+ \beta_4 \cdot StarAnalysts_j + \beta_5 \cdot PreviousRatings_j + \beta_6 \cdot AnalystDispersion_j \\ &+ \beta_7 \cdot Benchy_j + \beta_8 \cdot X_j + \beta_9 \cdot Z_j + Month_j + \epsilon_j. \end{aligned} \tag{5}$$

[28]Note that disagreement most often concerns the case of a rating being high but in which the ICO fails. There are very few cases of successful ICOs with an average poor rating (N=44).

$Success_j$ indicates the success dummy as described in Section 2. Alternatively, we run OLS regressions with $MarketPerformance_j$ as the dependent variable.[29] X_j again represents the controls as in Equation 1. Additionally, we control for linguistic measures with Z_j, which contains the average tone, uncertainty, and complexity levels, as well as the length of all rating reviews written about ICO j. We further include a dummy for the month-end of the ICO, $Month_j$, to absorb time trends common to all ICOs, as well as the BTC return during the ICO campaign, to control for the overall market sentiment. In regressions with market capitalization as the dependent variable, these fixed effects and the BTC return are arguably particularly important to control for general market developments and focus on the cross-section of ICOs.

In addition to the unconditional success of ICOs, we investigate the success conditional on having received very high or very low ratings. Thus, we run the following logit regression on the ICO level:

$$\begin{aligned} Disagreement_j = {} & \beta_0 + \beta_1 \cdot ReciprocalRatingShare_j + \beta_2 \cdot \#\, Analysts_j \\ & + \beta_3 \cdot StarAnalysts_j + \beta_4 \cdot PreviousRatings_j \\ & + \beta_5 \cdot AnalystDispersion_j + \beta_6 \cdot Benchy_j \\ & + \beta_7 \cdot X_j + \beta_8 \cdot Z_j + Month_j + \epsilon_j, \end{aligned} \tag{6}$$

where X_j is the same set of controls as in Equation 1, Z_j includes linguistic measures of all rating reviews written about ICO j, and $Month_j$ dummies and the BTC return during the ICO campaign absorb common time trends and control for the overall market sentiment, respectively.

[Table 9 about here]

[Table 10 about here]

[29] In the Appendix, we alternatively use the dollar amount raised during the ICO campaign as a measure of success.

As a baseline result, regressions in Table 9 confirm that ratings on average help predict ICO success. Specifically, the likelihood of an ICO being successful as measured by its initial listing in columns (1) through (3) is higher if the number of analysts rating a given ICO is high.[30] This result holds even after controlling for a wide variety of ICO characteristics.[31]

Moreover, columns (1) through (3) show that ICOs that receive higher human ratings are more likely to succeed. We also find the machine-generated rating, Benchy, to be predictive, indicating that ICOs are, on average, more likely to be successful the more publicly available information there is about them.[32] These results are in line with Bourveau et al. (2022), who find that information intermediaries (which in their case are proxied by a combined measure of human analysts and the machine-generated rating) help mitigating the high asymmetric information environment of ICOs.

Our main interest is in the role of the heterogeneity among analysts, and in what predicts ICO failure despite high ratings. First, Table 9 shows that ICOs with a higher share of reciprocal ratings experience a lower growth in market capitalization in the first 90 days after being listed. In particular, column (5) suggests that the market capitalization relative to the amount raised during the campaign of an ICO with an average share of reciprocal ratings is around 5 percentage points lower compared to an ICO without any reciprocal rating.

[30]This finding is in line with the general literature on analysts and rating agencies, which indicates that the market appreciates analyst coverage (Demiroglu and Ryngaert, 2010) and the existence of ratings (Sufi, 2009).

[31]The control variables themselves are not of primary interest here, and the full tables are displayed in the Online Appendix Table OA4 and Table OA5 for space reasons. Some similar results, as in the prior literature, do emerge. For example, we find a positive coefficient for Bitcointalk and negative coefficients for the Bonus dummy. Successful ICOs tend to have longer whitepapers. Two control variables that received little prior attention in the literature are MVPs and IEOs. The use of crypto-exchange launchpads for Initial Exchange Offerings (IEOs) positively correlates with the two success variables. Somewhat surprisingly, ICOs with a minimum viable product (MVP) feature a lower probability of success. This unexpected result might be due to a non-regulated definition of minimum viable products. For example, drafts of codes on GitHub that are open to a discussion by other GitHub users were classified as MVP. We further find that ICOs with a large number of commits on GitHub are more likely to be successful. It is possible that more experienced analysts are likely to participate in many ICOs, and that reciprocal ratings may thus partly represent the analyst's expertise; however, including controls for analyst experience does not affect the findings. We do not find any significant effect on success of bitcoin return during the campaign.

[32]Note that Benchy is a rating only of the availability of the information, not of the ICO as such. The positive effect of Benchy is in line with the finding that investors value the dissemination of corporate news releases via robots, even when that information is in principle already available (Blankespoor et al., 2018).

While the reciprocal share does not explain the binary ICO success indicator unconditionally, descriptive evidence in Figure 5 suggests that it does correlate significantly with failure conditional on high ratings. This evidence is confirmed by regression results shown in Table 10. Specifically, the probability that markets disagree with a very positive analyst evaluation is 7.5 percent higher for an ICO with an average share of reciprocal ratings compared to an ICO without any reciprocal rating.

A few additional comments are in order. First, Table A3 shows that the effect emerges largely from failed ICOs despite high ratings (not from successful ICOs despite low ratings). Second, in Online Appendix OA.1, we document that only the non-reciprocal rating score predicts ICO success, and that the share of reciprocal ratings leads to a disagreement of the market only with reciprocal ratings. Moreover, we show in Online Appendix OA.2 that the effect is only present for actual reciprocal ratings (as defined so far), but not for ratings for which the analyst might expect a quid pro quo rating, because they themselves are doing an ICO at a later point in time. Thus, investors disagree only if they *observe* a reciprocal rating structure.

There are two possible interpretations of this negative association of the reciprocal share and the short-term and medium-success of ICOs. First, we note that we control for a wide variety of factors presumably capturing variation in ICO quality. However, it is still possible that reciprocal ratings are correlated with some additional unobserved variation in ICO quality. The second interpretation is that, as a matter of principle, investors trust ICOs with more reciprocal ratings less, even when these ratings do not suffer from a conflict of interest.

While these interpretations are not mutually exclusive, an additional test provides further insight. For each ICO, we calculate the difference between the average reciprocal and non-reciprocal ratings. We then divide the sample into cases in which the average reciprocal rating is higher than or equal to the average of non-reciprocal ratings, and cases in which reciprocal ratings are lower than the non-reciprocal ones. In the former case reciprocal ratings influence the overall ICO rating to a large extent, whereas in the latter case reciprocal ratings are less likely to bias the overall ICO rating. If investors dislike reciprocal ratings in general, we would expect the reciprocal rating share to be a significant determinant of disagreement in both cases. Columns 4 and 5 of Table 10 present the results. The share of reciprocal analysts matters only for the conditional success for those cases for which reciprocal ratings are at least as positive as non-reciprocal ratings. A caveat is that these regressions are based on relatively small samples (because they are only available for the subsample with reciprocal ratings).[33] That said, they provide some suggestive evidence that investors are not concerned with reciprocal ratings per se, but rather that positive reciprocal ratings provide an additional signal of the poor quality of an ICO.

Consider next the linguistic measures of the rating, which also do not predict ICO success per se (at least once controlling for the quantitative rating), but still provide insight. Table 10 shows that the likelihood of failure for a highly-rated ICO increases as the positivity of the tone and complexity of the language increase. Similarly, ICO failure despite high average ratings occurs more frequently when the analysts were more positive in ratings prior to their rating of ICO j.

Table 9 suggests that star analyst coverage is not predictive for ICO success. Again, however, Table 10 gives some indication that highly rated ICOs fail less frequently when many star analysts cover them.[34]

Finally, analyst dispersion is also relevant for ICO success only if the average view of analysts is very positive. Interestingly, and at first surprisingly, when analysts' ratings are highly dispersed and higher on average, ICOs are less likely to fail. Intuitively, the combination of high average ratings and high dispersion occurs when there are several extremely positive and some negative views. The very positive ratings then carry the day. This is sim-

[33]Due to the small sample and the large set of controls, we run linear regressions in these analyses.

[34]ICO success might not only be driven by analysts outside the firm, but also by analysts inside the firm, who act as advisors. In untabulated results, we observe that top advisors indeed bring skill into the team, resulting in significantly higher ratings and a somewhat higher success probability (though not quite statistically significantly, t-stat=1.52).

ilar in spirit to the apparent anomaly that stocks with high dispersion of analyst opinions have high prices and, thus, lower future returns (Diether et al., 2002).[35]

Overall, the results highlight that even when a characteristic is not related unconditionally to ICO success, it is not irrelevant for understanding ICO success. Specifically, factors such as the reciprocity of ratings, analyst dispersion and the presence of star analysts explain deviations from the outcome given a very high level of ratings. Thus, while it is natural that average ratings predict the success of an ICO campaign, our key result is that the detailed characteristics of the ICO ratings and those who provide them contain important additional information.

3.2.2 Ratings and ICO scams

We have established that analyst ratings help to predict ICO success, but that investors tend to disregard reciprocal ratings. Does the latter result occur because ICOs with a higher fraction of reciprocal ratings are more likely to be fraudulent? To answer this question, we rerun the regression from Equation 5, but replace the success dummy with a dummy that equals one if the ICO was detected to intentionally defraud investors.

[Table 11 about here]

As Table 11 shows, we find no correlation between the share of reciprocal analysts and fraudulent ICOs. Also, neither the level of machine-generated ratings nor the level of human analyst ratings help to identify fraudulent ICOs. It still pays for investors to consider the human analyst assessments, however. In particular, ICOs with more dispersion among analysts both in quantitative and in qualitative ratings tend to be fraudulent.

[35] This interpretation of analyst dispersion has been challenged in equity markets. For example, Avramov et al. (2009) show that the analyst dispersion anomaly is driven by a small fraction of firms with very high credit risk.

4 Conclusion

The intersection of new technologies and financial markets (FinTech) holds great promise. One relatively recent phenomenon in this space is the opportunity for new ventures to engage in coin/token Offerings, a new form of financing. Yet despite the problem of asymmetric information looming large in these markets, there was a tremendous rise in ICOs, followed by a market crash. Motivated by this dramatic development, this paper studies the role of information intermediaries, human experts who may help to ameliorate this asymmetric information problem, in unregulated financial markets.

While the rise and fall of the ICO market is interesting in its own right, ICO analysts show many interesting parallels to equity analysts or rating agencies. Particularly noteworthy are potential conflicts of interest, and how investors interpret them. The advantage of the ICO setting is that detailed data on links between analysts and securities they rate are available. For example, we document that an ICO analyst i, when rating an ICO j, tends to issue a rating that depends on the rating that their own affiliated ICO had previously received from team members of ICO j. However, there is a higher probability that an ICO will fail, even when it has very favorable ratings, when more of those ratings are reciprocal. ICOs with a high share of reciprocal ratings also tend to have a lower market capitalization 90 days after their listing on an exchange, relative to the funds initially raised.

Thus, while the prior literature shows that information intermediaries predict the success of an ICO campaign, our key result is that conflicts of interest affect how effectively human analysts can mitigate the highly asymmetrical environment of the unregulated ICO market. The findings suggest that to ensure the functioning of newly emerging and largely unregulated financial markets, such as for example decentralized finance, IEOs, or STOs, there needs to be at least some regulation that prevents such conflicts of interest.

Understanding ICO success and failure requires looking beyond averages and studying the detailed characteristics of the ratings and those who provide them. A necessary precondition for investors to take such a differentiated approach to investments is the availability of information about the track record and potentially conflicting activities of analysts. Thus, the availability of information appears to support the ICO market's allocative role in society. While these results have been obtained on this largely unregulated market, the general insight that investors seem to value information about analysts is likely to be relevant for other markets as well.

References

Adhami, S., Giudici, G., and Martinazzi, S. (2018). Why do businesses go crypto? An empirical analysis of Initial Coin Offerings. *Journal of Economics and Business*, 100(C):64–75.

Aggarwal, R., Hanley, K. W., and Zhao, X. (2019). The role of information production in private markets: Evidence from Initial Coin Offerings. Working paper.

Agrawal, A. and Chen, M. A. (2008). Do analyst conflicts matter? Evidence from stock recommendations. *Journal of Law & Economics*, 51(3):503–537.

Amsden, R. and Schweizer, D. (2018). Are blockchain crowdsales the new 'gold rush'? Success determinants of Initial Coin Offerings. Working paper.

Aubry, M., Kräussl, R., Manso, G., and Spaenjers, C. (2019). Machine learning, human experts, and the valuation of real assets. Working paper.

Avramov, D., Chordia, T., Jostova, G., and Philipov, A. (2009). Dispersion in analysts' earnings forecasts and credit rating. *Journal of Financial Economics*, 91(1):83–101.

Badoer, D. C., Demiroglu, C., and James, C. M. (2019). Ratings quality and borrowing choice. *Journal of Finance*, 74(5):2619–2665.

Bae, K.-H., Stulz, R. M., and Tan, H. (2008). Do local analysts know more? A cross-country study of the performance of local analysts and foreign analysts. *Journal of Financial Economics*, 88(3):581–606.

Baghai, R. and Becker, B. (2017). Non-rating revenue and conflicts of interest. *Journal of Financial Economics*, 127.

Baker, H. K. and Mansi, S. A. (2002). Assessing credit rating agencies by bond issuers and institutional investors. *Journal of Business Finance & Accounting*, 29(9-10):1367–1398.

Bakos, Y. and Halaburda, H. (2019). The role of cryptographic tokens and ICOs in fostering platform adoption. Working paper.

Bar-Isaac, H. and Shapiro, J. (2013). Ratings quality over the business cycle. *Journal of Financial Economics*, 108(1):62 – 78.

Benedetti, H. and Kostovetsky, L. (2021). Digital tulips? Returns to investors in initial coin offerings. *Journal of Corporate Finance*, 66:101786.

Biais, B., Bisière, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *Review of Financial Studies*, 32(5):1662–1715.

Blankespoor, E., deHaan, E., and Zhu, C. (2018). Capital market effects of media synthesis and dissemination: Evidence from robo-journalism. *Review of Accounting Studies*, 23(1):1–36.

Bochkay, K., Hales, J., and Chava, S. (2020). Hyperbole or reality? Investor response to extreme language in earnings conference calls. *The Accounting Review*, 95(2):31–60.

Bolton, P., Freixas, X., and Shapiro, J. (2012). The credit ratings game. *Journal of Finance*, 67(1):85–111.

Boreiko, D. and Vidusso, G. (2019). New blockchain intermediaries: Do ICO rating websites do their job well? *Journal of Alternative Investments*, 21(4):67–79.

Bourveau, T., De George, E. T., Ellahie, A., and Macciocchi, D. (2022). The role of disclosure and information intermediaries in an unregulated capital market: Evidence from Initial Coin Offerings. *Journal of Accounting Research*, 60(1):129–167.

Bradley, D., Gokkaya, S., and Liu, X. (2017). Before an analyst becomes an analyst: Does industry experience matter? *Journal of Finance*, 72(2):751–792.

Bradshaw, M., Ertimur, Y., O'Brien, P., et al. (2017). Financial analysts and their contribution to well-functioning capital markets. *Foundations and Trends in Accounting*, 11(3):119–191.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business and Economic Statistics*, 29(2):238–249.

Campbell, J., DeAngelis, M., and Moon, J. (2019). Skin in the game: Personal stock holdings and investors' response to stock analysis on social media. *Review of Accounting Studies*, 24(C):731–779.

Catalini, C., Foster, C., and Nanda, R. (2018). Machine intelligence vs. human judgement in new venture finance. Working paper.

Cerchiello, P., Tasca, P., and Toma, A. M. (2019). ICO success drivers: A textual and statistical analysis. *Journal of Alternative Investments*, 21(4):13–25.

Chakraborty, S. and Swinney, R. (2020). Signaling to the crowd: Private quality information and rewards-based crowdfunding. *Manufacturing & Service Operations Management*, 1(C):1–15.

Chan, L. K. C., Karceski, J., and Lakonishok, J. (2007). Analysts' conflicts of interest and biases in earnings forecasts. *Journal of Financial and Quantitative Analysis*, 42(4):893–913.

Chen, H., De, P., Hu, Y. J., and Hwang, B.-H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *Review of Financial Studies*, 27(5):1367–1403.

Chod, J. and Lyandres, E. (2021). A theory of ICOs: Diversification, agency, and information asymmetry. *Management Science*, 67(10):5969–5989.

Chod, J. and Lyandres, E. (2022). Product market competition with crypto tokens and smart contracts. Working paper.

Chod, J., Trichakis, N., and Yang, S. A. (2022). Platform tokenization: Financing, governance, and moral hazard. *Management Science*.

Chu, C. S. and Rysman, M. (2019). Competition and strategic incentives in the market for credit ratings: Empirics of the financial crisis of 2007. *American Economic Review*, 109(10):3514–55.

Cohen, L., Lou, D., and Malloy, C. J. (2020). Casting conference calls. *Management Science*, 66(11):5015–5039.

Cong, L. W. and He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754–1797.

Cong, L. W., Li, Y., and Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies*, 34(3):1105–1155.

Cong, L. W., Li, Y., and Wang, N. (2022). Token-based platform finance. *Journal of Financial Economics*, 144(3):972–991.

Conley, J. P. (2017). Blockchain and the economics of crypto-tokens and Initial Coin Offerings. Working paper.

Crawford, S. S., Roulstone, D. T., and So, E. C. (2012). Analyst initiations of coverage and stock return synchronicity. *The Accounting Review*, 87(5):1527–1553.

Da, Z. and Huang, X. (2020). Harnessing the wisdom of crowds. *Management Science*, 66(5):1847–1867.

Davydiuk, T., Gupta, D., and Rosen, S. (2019). De-crypto-ing signals in Initial Coin Offerings: Evidence of rational token retention. Working paper.

Demiroglu, C. and Ryngaert, M. (2010). The first analyst coverage of neglected stocks. *Financial Management*, 39(2):555–584.

Deng, X., Lee, Y. T., and Zhong, Z. (2018). Decrypting coin winners: Disclosure quality, governance mechanism and team networks. Working paper.

Diether, K. B., Malloy, C. J., and Scherbina, A. (2002). Differences of opinion and the cross section of stock returns. *Journal of Finance*, 57(5):2113–2141.

Dittmar, R. and Wu, D. A. (2019). Initial Coin Offerings hyped and dehyped: An empirical examination. Working paper.

Drake, M. S., Thornock, J. R., and Twedt, B. J. (2017). The internet as an information intermediary. *Review of Accounting Studies*, 22(2):543–576.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109.

Ellis, D. (1998). Different sides of the same story: Investors? And issuers? Views of rating agencies. *Journal of Fixed Income*, 1:35–45.

Ernst and Young (2018). Initial Coin Offerings (ICOs): The class of 2017 – one year later. EY study.

Fahlenbrach, R. and Frattaroli, M. (2020). ICO Investors. *Financial Markets and Portfolio Management*, pages 1–59.

Farrell, M., Jame, R., and Qiu, T. (2020). The cross-section of non-professional analyst skill. Working paper.

Feng, C., Li, N., Wong, M., and Zhang, M. (2019). Initial Coin Offerings, blockchain technology, and white paper disclosures. Working paper.

Fisch, C. (2019). Initial Coin Offerings (ICOs) to finance new ventures. *Journal of Business Venturing*, 34(1):1–22.

Fisch, C., Masiak, C., Vismara, S., and Block, J. (2019). Motives and profiles of ICO investors. *Journal of Business Research*, forthcoming.

Fisch, C. and Momtaz, P. P. (2020). Institutional investors and post-ICO performance: An empirical analysis of investor returns in Initial Coin Offerings (ICOs). *Journal of Corporate Finance*, 64:101679.

Florysiak, D. and Schandlbauer, A. (2022). Experts or charlatans? ICO analysts and white paper informativeness. *Journal of Banking & Finance*, 139.

Gan, R. J., Tsoukalas, G., and Netessine, S. (2020). Initial Coin Offerings, speculators, and asset tokenization. Working paper.

Giudici, G. and Adhami, S. (2019). The impact of governance signals on ICO fundraising success. *Journal of Industrial and Business Economics*, 46(2):283–312.

Goldstein, I., Gupta, D., and Sverchkov, R. (2019). Utility tokens as a commitment to competition. Working paper.

Gryglewicz, S., Mayer, S., and Morellec, E. (2021). Optimal financing with tokens. *Journal of Financial Economics*, 142(3):1038–1067.

Gunning, R. (1952). *The technique of clear writing*. McGraw-Hill, New York.

Hinzen, F. J., John, K., and Saleh, F. (2022). Bitcoin's limited adoption problem. *Journal of Financial Economics*, 144(2):347–369.

Hirth, S. (2014). Credit rating dynamics and competition. *Journal of Banking & Finance*, 49(C):100 – 112.

Howell, S. T., Niessner, M., and Yermack, D. (2020). Initial Coin Offerings: Financing growth with cryptocurrency token sales. *Review of Financial Studies*, 33(9):3925–3974.

Hu, A. S., Parlour, C. A., and Rajan, U. (2019). Cryptocurrencies: Stylized facts on a new investible instrument. *Financial Management*, 48(4):1049–1068.

Jame, R., Johnston, R., Markov, S., and Wolfe, M. C. (2016). The value of crowdsourced earnings forecasts. *Journal of Accounting Research*, 54(4):1077–1110.

Kadan, O., Madureira, L., Wang, R., and Zach, T. (2009). Conflicts of Interest and Stock Recommendations: The Effects of the Global Settlement and Related Regulations. *Review of Financial Studies*, 22(10):4189–4217.

Kempf, E. (2020). The job rating game: Revolving doors and analyst incentives. *Journal of Financial Economics*, 135(1):41–67.

Klein, A., Saunders, A., and Wong, Y. T. F. (2019). Is there a quid pro quo between hedge funds and sell-side equity analysts? *Journal of Portfolio Management*, 45(6):117–132.

Law, K. K. F. and Mills, L. F. (2019). Financial gatekeepers and investor protection: Evidence from criminal background checks. *Journal of Accounting Research*, 57(2):491–543.

Lee, J., Li, T., and Shin, D. (2022). The wisdom of crowds in FinTech: Evidence from Initial Coin Offerings. *Review of Corporate Finance Studies*, 11(1):1–46.

Lee, J. and Parlour, C. A. (2022). Consumers as financiers: Consumer surplus, crowdfunding, and Initial Coin Offerings. *Review of Financial Studies*, 35(3):1105–1140.

Leone, A. J. and Wu, J. (2002). What does it take to become a superstar? Evidence from institutional investor rankings of financial analysts. Working paper.

Li, J. and Mann, W. (2020). Initial Coin Offerings and platform building. Working paper.

Li, J. and Mann, W. (2021). Initial Coin Offering: Current research and future directions. *The Palgrave Handbook of Technological Finance*, pages 369–393.

Lin, H. and McNichols, M. F. (1998). Underwriting relationships, analysts' earnings forecasts and investment recommendations. *Journal of Accounting and Economics*, 25(1):101–127.

Livingston, M. B., Wei, J. D., and Zhou, L. (2010). Moody's and S&P ratings: Are they equivalent? Conservative ratings and split rated bond yields. *Journal of Money, Credit and Banking*, 42(7).

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? *Journal of Finance*, 66(1):35–65.

Loughran, T. and McDonald, B. (2020). Textual analysis in finance. *Annual Review of Financial Economics*, 12:357–375.

Lourie, B. (2019). The revolving door of sell-side analysts. *The Accounting Review*, 94(1):249–270.

Lyandres, E., Palazzo, B., and Rabetti, D. (2022). Initial Coin Offering (ICO) success and post-ico performance. *Management Science*.

Malinova, K. and Park, A. (2018). Tokenomics: When tokens beat equity. Working paper.

Malmendier, U. and Shanthikumar, D. (2007). Are small investors naive about incentives? *Journal of Financial Economics*, 85(2):457–489.

Merkley, K., Michaely, R., and Pacelli, J. (2017). Does the scope of the sell-side analyst industry matter? An examination of bias, accuracy, and information content of analyst reports. *Journal of Finance*, 72(3):1285–1334.

Michaely, R. and Womack, K. L. (1999). Conflict of interest and the credibility of underwriter analyst recommendations. *Review of Financial Studies*, 12(4):653–686.

Mola, S., Rau, P. R., and Khorana, A. (2013). Is there life after the complete loss of analyst coverage? *The Accounting Review*, 88(2):667–705.

O'Brien, P., McNichols, M. F., and Hsiou-Wei, L. (2005). Analyst impartiality and investment banking relationships. *Journal of Accounting Research*, 43(4):623–650.

PwC (2019). 5th ICO / STO report. PwC study.

Roosenboom, P., van der Kolk, T., and de Jong, A. (2020). What determines success in Initial Coin Offerings? *Venture Capital*, 22(2):1–23.

Samieifar, S. and Baur, D. G. (2020). Read me if you can! An analysis of ICO white papers. *Finance Research Letters*, page 101427.

Sufi, A. (2009). The real effects of debt certification: Evidence from the introduction of bank loan ratings. *Review of Financial Studies*, 22(4):1659–1691.

Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics*, 58(3):369–396.

Zetzsche, D. A., Buckley, R. P., Arner, D. W., and Föhr, L. (2019). The ICO Gold Rush: It's a Scam, It's a Bubble, It's a Super Challenge for Regulators. *Harvard International Law Journal*, 60(2):267–315.

Zhang, S., Aerts, W., Lu, L., and Pan, H. (2019). Readability of token whitepaper and ICO first-day return. *Economics Letters*, 180(C):58–61.

Figure 1: Number of ratings in a month and the Bitcoin price in $

This figure presents the number of ratings in a given month over time, the number of the newly announced ICOs, the number of ratings by analysts who registered in the same month, as well as the Bitcoin price in $.

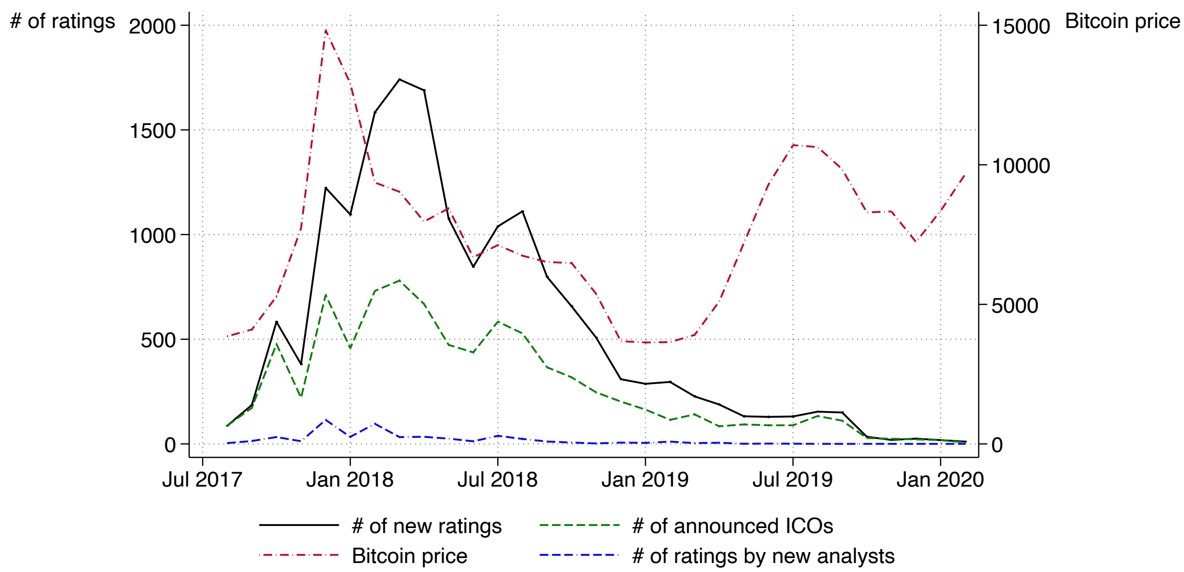

Figure 2: ICO analyst ratings and rating dispersion over time

This figure plots the average total rating and the analysts' rating dispersion (left axis) as well as the Bitcoin price in $ (right axis).

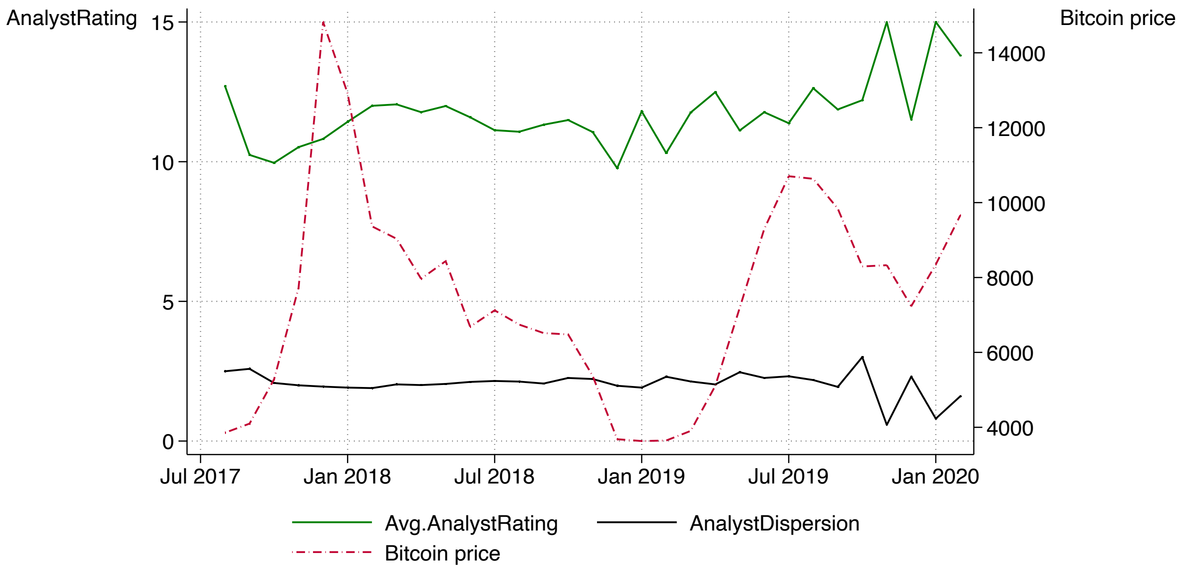

Figure 3: Successful and unsuccessful ICOs over time

The figure shows the number of ICOs over time, distinguishing between successful and failed ICOs. An ICO is labeled successful if the related coin successfully completes the offering and receives funding. In total, we identify 5,384 ICOs of which 1,932 ICOs succeeded.

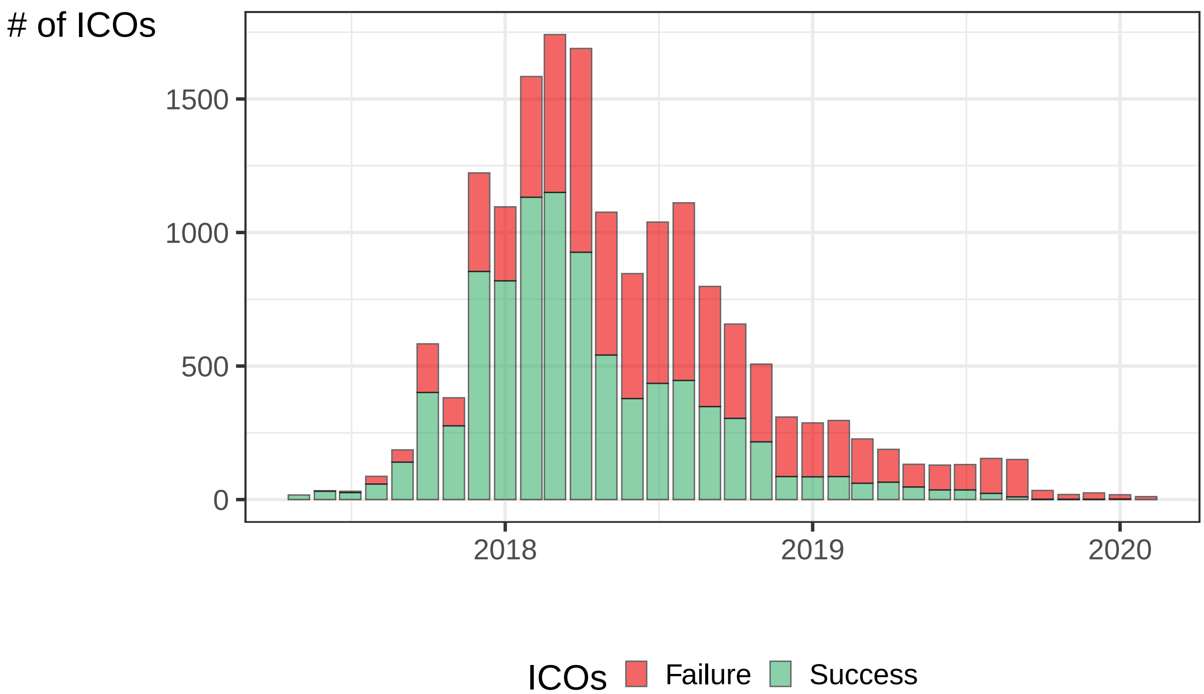

Figure 4: Forecast error over time

This figure shows the average forecast error and the number of successful ICOs (as a share of total ICOs) in a given month. The average forecast error is further split into $ForecastErrorOptimistic_i$ and $ForecastErrorPessimistic_i$ to capture the monthly averaged forecast error separately for the ratings when the analyst was too optimistic and pessimistic, respectively.

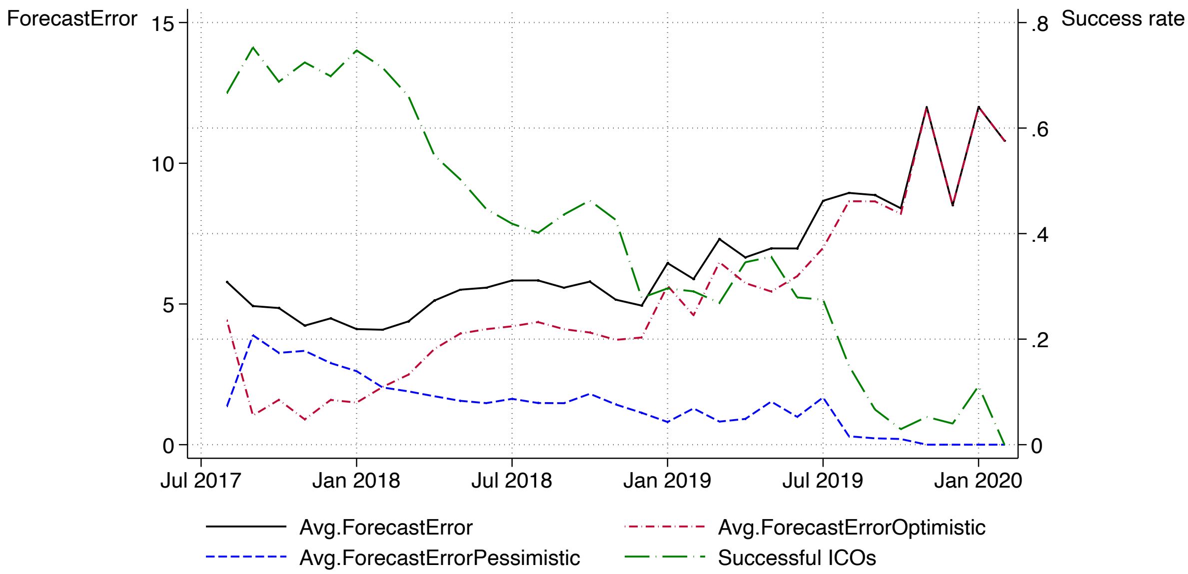

Figure 5: Reciprocal Ratings

The figure shows binned scatter plots summarizing the main results. Panel (a) uses within ICO variation, i.e., ICO fixed effects are absorbed. All variables are defined in Table A1.

(a) Reciprocal ratings are more favorable

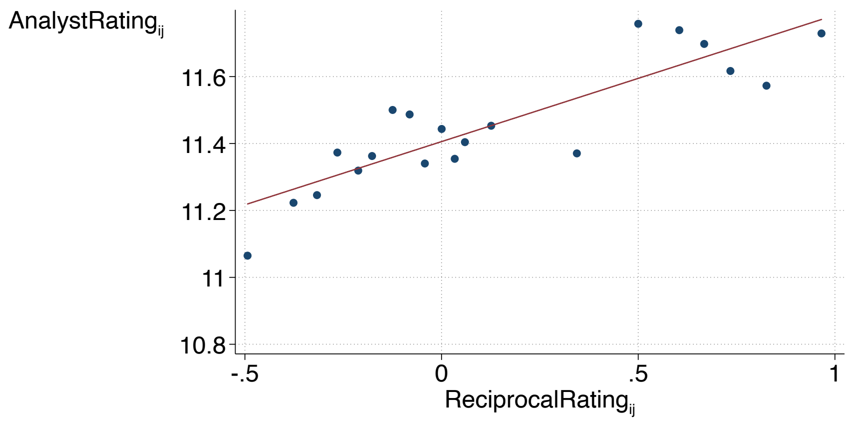

(b) Even ICOs with high average ratings fail frequently, and especially so when many ratings are reciprocal

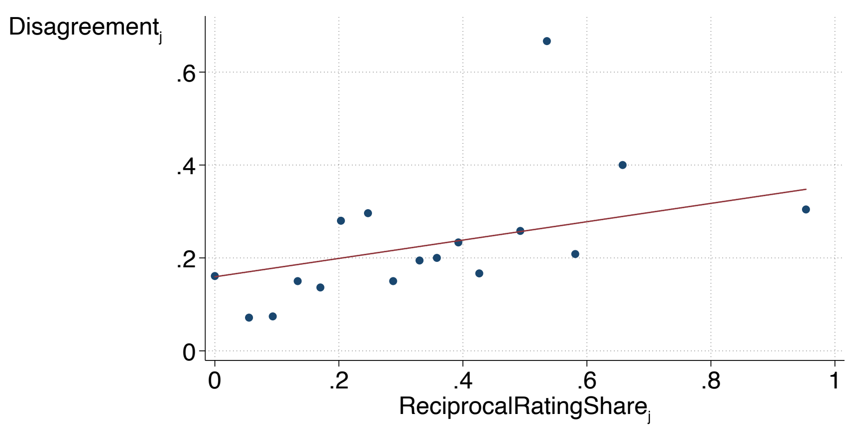

Table 1: ICO affiliation of analysts from the platform ICObench

This table tabulates the distribution of ICO projects among analysts. The total number of analysts in our sample is 539. The list of associated ICOs for each analyst is available on their webpage in ICObench.com.

Number of associated ICOs	Count
0	220
1	121
2	50
3	28
4	23
5	15
6	9
7	13
8	3
9	8
>=10	49
Total number of analysts	539

Table 2: ICO analyst networks: An example

This table presents a hypothetical example of our data set. In Panel A, we show the team members of the three ICOs in the sample, namely, "A-Tokens" where Adam and Ashley are among the team members, "Bethereum" where the team includes Barbara and Benjamin, and "CryptoPay" with Cora and Chris in the team. In Panel B, we outline a hypothetical rating history. For example, in October 2017, Ashley (member of A-Tokens) provides a rating of 12 for Bethereum. In December 2017, Chris (member of CryptoPay) provides a rating of 15 for Bethereum. For this rating, we set *ReciprocalRating* equal to 1 because, in a month before that, in November 2017, Benjamin (member of Bethereum) gave a rating of 14 for CryptoPay, with which Chris is affiliated. Hence, we consider the rating given in December 2017 as a reply to the rating received in November 2017.

A. ICOs and members:

A-Tokens	Bethereum	CryptoPay	...
Adam Ashley	Barbara Benjamin	Cora Chris	

B. Ratings:

Date	Analyst	provides a rating for:	AnalystRating	ReciprocalRating	ReceivedRating if *ReciprocalRating* = 1
1)Oct 2017	Ashley	Bethereum	12	0	-
2)Nov 2017	Benjamin	CryptoPay	14	0	-
3)Dec 2017	Chris	Bethereum	15	1	14
4)Jan 2018	Adam	CryptoPay	9	0	-
...					

Table 3: Descriptive statistics

This table shows descriptive statistics of the variables used in the analysis. The variables are sorted alphabetically within each panel. The sample consists of 5,384 ICOs listed in ICObench.com of which 2,378 received 13,834 ratings in total. All variables are defined in Table A1.

	N	Min	P25	Mean	P50	P75	Max	Std. Dev.
A. Rating characteristics								
$AnalystExperience_i^{j-1}$	13,834	0	2.1	3.1	3.3	4.3	6.2	1.6
$AnalystRating(Product)_{ij}$	13,834	1	3	3.6	4	5	5	1.1
$AnalystRating(Team)_{ij}$	13,834	1	3	3.9	4	5	5	1.1
$AnalystRating(Vision)_{ij}$	13,834	1	3	3.9	4	5	5	1.1
$AnalystRating_{ij}$	13,834	3	10	11	12	14	15	3
$ForecastError_i^{j-1}$	12,460	0	4.1	4.8	4.8	5.6	12	1.5
$Modified_{ij}$	13,834	0	0	.13	0	0	1	.33
$OrderRank_{ij}$	7,639	1	6	14	11	18	94	12
$ReceivedRating(Product)_{ij}$	1,754	1	4	4.1	4	5	5	.74
$ReceivedRating(Team)_{ij}$	1,754	1	4	4.3	4	5	5	.67
$ReceivedRating(Vision)_{ij}$	1,754	1	4	4.2	4	5	5	.7
$ReceivedRating_{ij}$	1,754	3	12	13	12	14	15	1.8
$ReciprocalRating_{ij}$	13,834	0	0	.13	0	0	1	.33
$ReviewLength_{ij}$	9,165	1.1	3.4	3.8	3.9	4.3	7.9	.97
$ReviewTone_{ij}$	9,165	-.75	-.043	-.014	-.0086	.018	.67	.075
$StarAnalyst_{ij}$	13,834	0	0	.27	0	1	1	.44
B. ICO characteristics								
$AmountRaised_j$	5,339	0	0	5.3	0	14	22	7.3
$MarketPerformance_j$	1,892	0	0	.86	0	.039	76	4.7
$Scam_j$	5,339	0	0	.043	0	0	1	.2
$Success_j$	5,339	0	0	.35	0	1	1	.48
Rating Controls:								
$AnalystDispersion_j$	5,339	0	0	.61	0	.96	8.5	1.2
$AnalystExperience_j$	5,339	.69	.69	2.1	.69	4	6.1	1.7
$AnalystRating_j$	2,378	3	9	10	11	13	15	3
$Benchy_j$	5,339	.1	2.4	2.9	2.9	3.5	5	.75
$PreviousRatings_j$	2,322	3	11	11	11	12	15	1.4
$ReciprocalRatingShare_j$	2,378	0	0	.072	0	0	1	.19
$ReviewComplexity_j$	5,339	0	0	4.4	0	11	59	6.2
$ReviewLength_j$	1,883	1.1	3.6	4	4	4.5	7.3	.81
$ReviewToneDispersion_j$	5,339	0	0	.013	0	0	.55	.033
$ReviewTone_j$	5,339	-.67	0	-.007	0	0	.29	.034

$ReviewUncertainty_j$	5,339	0	0	.0055	0	0	.33	.015
$StarAnalysts_j$	2,378	0	0	.31	.22	.5	1	.34
$\#Analysts_j$	5,339	0	0	2.6	0	2	94	6.1
VentureOffering Controls:								
$Bonus_j$	5,339	0	0	.14	0	0	1	.35
$Bounty_j$	5,339	0	0	.28	0	1	1	.45
$GitHubCommits_j$	5,339	0	0	1.5	0	1.6	13	2.9
$HardCap_j$	5,339	0	0	.59	1	1	1	.49
IEO_j	5,339	0	0	.052	0	0	1	.22
KYC_j	5,339	0	0	.49	0	1	1	.5
MVP_j	5,339	0	0	.2	0	0	1	.4
$Presale_j$	5,339	0	0	.53	1	1	1	.5
$RetentionRatio_j$	4,224	0	30	46	45	60	100	21
$VestingDisclosure_j$	5,339	0	0	.27	0	1	1	.45
$\#Advisors_j$	5,339	0	0	1.1	1.1	2.1	4.3	1.1
$\#TeamMembers_j$	5,339	0	1.4	1.7	1.9	2.4	4	.99
WhitePaper Controls:								
$WhitePaperComplexity_j$	5,339	0	0	.0069	.0062	.011	.069	.0074
$WhitePaperLength_j$	5,339	0	0	5.5	8.2	8.9	11	4.2
$WhitePaperTechnicalWords_j$	5,339	0	0	.096	.12	.16	.35	.078
$WhitePaperTone_j$	5,339	-.11	-.00055	.00088	0	.0039	.057	.0084
$WhitePaperUncertainty_j$	5,339	0	0	.0086	.0083	.014	.07	.0086
SocialMedia Controls:								
$Bitcointalk_j$	5,339	0	0	.57	1	1	1	.49
$Facebook_j$	5,339	0	1	.78	1	1	1	.41
$SocialMediaComplexity_j$	5,339	0	0	13	1.7	8.1	1,634	52
$SocialMediaCount_j$	5,339	0	0	3.1	3.3	5.3	9.6	2.6
$SocialMediaExtremeWords_j$	5,339	0	0	.0002	0	.00015	.012	.00069
$SocialMediaLength_j$	5,339	0	0	5.8	7.3	9	14	4.1
$SocialMediaTechnicalWords_j$	5,339	0	0	.0095	.002	.0081	.33	.024
$SocialMediaTone_j$	5,339	-.041	-3.6e-06	.00034	0	.00017	.081	.0031
$SocialMediaUncertainty_j$	5,339	0	0	.0007	.000064	.00059	.064	.0022
MarketSentiment:								
$MarketSentiment_j$	5,338	-.057	-.0056	-.001	-.00098	.0036	.057	.0091

Table 4: Rating determinants

This table presents linear regression results for Equation 1. The dependent variable is the total rating score that an analyst gave to an ICO. Control variables include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). The specification in column (5) includes $Month$ and ICO fixed effects. All variables are defined in Table A1. t-statistics are given in parentheses. Standard errors are clustered at the ICO and analyst level. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$AnalystRating_{ij}$				
	(1)	(2)	(3)	(4)	(5)
$Benchy_j$	1.518***	1.599***		0.706***	
	(10.43)	(11.42)		(5.11)	
$StarAnalyst_{ij}$		-0.802***	-0.721***	-0.735***	-0.442***
		(-3.80)	(-3.43)	(-3.51)	(-2.68)
$ForecastError_i^{j-1}$		-0.161***	-0.130***	-0.143***	-0.080**
		(-3.66)	(-3.09)	(-3.49)	(-2.27)
$AnalystExperience_i^{j-1}$		0.063	0.037	0.044	-0.009
		(0.83)	(0.51)	(0.58)	(-0.12)
Observations	13834	12460	11257	11257	11698
R^2	0.121	0.146	0.109	0.123	0.528
VentureOffering Controls	No	No	Yes	Yes	Implied
WhitePaper Controls	No	No	Yes	Yes	Implied
SocialMedia Controls	No	No	Yes	Yes	Implied
MarketSentimet	No	No	Yes	Yes	Implied
Month Dummies	No	No	No	No	Yes
ICO FE	No	No	No	No	Yes

Table 5: Reciprocal ratings

This table presents linear regression results for Equation 2. The dependent variable is the total rating score that an analyst gave to an ICO for their team, vision, product, and the sum of all three categories ($AnalystRating$). In columns (1), (3), (5) and (7), regressions include all the ratings in the sample. In columns (2), (4), (6) and (8), we restrict the sample to the reciprocal ratings ($ReciprocalRating = 1$). All specifications include $Analyst$ and ICO fixed effects multiplied by dummies for the time of the rating (i.e., $Analyst \times Month$ and $ICO \times Month$ fixed effects in odd columns and $Analyst \times Quarter$ and $ICO \times Quarter$ fixed effects in even columns). All variables are defined in Table A1. t-statistics are given in parentheses. Standard errors are clustered at the ICO and analyst level. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$AnalystRating_{ij}$		$AnalystRating$ $(Team)_{ij}$		$AnalystRating$ $(Vision)_{ij}$		$AnalystRating$ $(Product)_{ij}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$ReciprocalRating_{ij}$	0.252**		0.062*		0.074*		0.117***	
	(2.46)		(1.78)		(1.75)		(2.83)	
$ReceivedRating_{ij}$		0.080*						
		(1.73)						
$ReceivedRating(Team)_{ij}$				0.117***				
				(3.12)				
$ReceivedRating(Vision)_{ij}$						0.065		
						(1.14)		
$ReceivedRating(Product)_{ij}$								0.000
								(0.01)
$Modified_{ij}$		-1.114***		-0.381***		-0.298**		-0.427***
		(-3.54)		(-3.02)		(-2.53)		(-3.46)
Observations	10354	1302	10354	1302	10354	1302	10354	1302
R^2	0.757	0.682	0.717	0.621	0.692	0.666	0.708	0.647
Analyst FE × Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ICO FE × Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: Linguistic nature of rating reviews

This table presents linear regression results for Equation 3. The dependent variable in Panel A is *ReviewLength*, defined as the natural logarithm of the total number of words in a review, and in Panel B *ReviewTone*, defined as the ratio of positive words minus negative words to total words in the review. We restrict the sample to reciprocal ratings ($ReciprocalRating_{ij} = 1$) in column (3) and to non-reciprocal ratings ($ReciprocalRating_{ij} = 0$) in column (4). We include *Analyst* and *ICO* fixed effects multiplied by dummies for the month of ratings (i.e., $Analyst \times Month$ and $ICO \times Month$ fixed effects) in column (2). As in Table 5, we can only include the interaction of ICO (analyst) and quarter dummies when restricting the sample to reciprocal ratings in column (3), i.e., $Analyst \times Quarter$ and $ICO \times Quarter$ fixed effects. In order to compare the coefficients for reciprocal and non-reciprocal ratings, we also include $Analyst \times Quarter$ and $ICO \times Quarter$ fixed effects in column (4). All variables are defined in Table A1. t-statistics are given in parentheses. Standard errors are clustered at the ICO and analyst level. ***, **, * indicate significance at the 1%, 5% and 10% levels.

Panel A

	$ReviewLength_{ij}$			
	(1)	(2)	(3)	(4)
$AnalystRating_{ij}$	-0.056***	-0.041***	-0.090***	-0.036***
	(-5.42)	(-5.41)	(-5.08)	(-5.36)
Observations	9165	6206	866	6119
R^2	0.033	0.825	0.800	0.786
Analyst FE × Time Dummies	No	Yes	Yes	Yes
ICO FE × Time Dummies	No	Yes	Yes	Yes

Panel B

	$ReviewTone_{ij}$			
	(1)	(2)	(3)	(4)
$AnalystRating_{ij}$	0.006***	0.006***	0.009***	0.005***
	(10.09)	(7.49)	(3.33)	(8.13)
Observations	9165	6206	866	6119
R^2	0.062	0.537	0.522	0.477
Analyst FE × Time Dummies	No	Yes	Yes	Yes
ICO FE × Time Dummies	No	Yes	Yes	Yes

Table 7: Order of rating issuance

This table presents linear regression results for Equation 4. The dependent variable is the order rank of the rating for an ICO. A lower value of the variable indicates that analyst i issued the rating for ICO j earlier. All specifications include month dummies of the analyst rating and *ICO* fixed effects. The specifications in columns (4), (5), and (6) include also *Analyst* fixed effects. The sample is restricted to ICOs with more than ten ratings. All variables are defined in Table A1. t-statistics are given in parentheses. Standard errors are clustered at the ICO and analyst level. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$OrderRank_{ij}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$AnalystRating_{ij}$	-0.104		-0.085	-0.182***		-0.177***
	(-1.65)		(-1.36)	(-2.81)		(-2.75)
$ReciprocalRating_{ij}$		-1.764***	-1.725***		-1.249**	-1.213**
		(-3.17)	(-3.08)		(-2.56)	(-2.48)
$StarAnalyst_{ij}$	-1.109***	-0.855***	-0.886***			
	(-3.00)	(-2.60)	(-2.69)			
$ForecastError_i^{j-1}$	-0.028	-0.027	-0.036	-0.130	-0.133	-0.122
	(-0.32)	(-0.31)	(-0.42)	(-1.18)	(-1.21)	(-1.12)
Observations	6829	6829	6829	6767	6767	6767
R^2	0.672	0.674	0.674	0.709	0.709	0.709
Month Dummies	Yes	Yes	Yes	Yes	Yes	Yes
ICO FE	Yes	Yes	Yes	Yes	Yes	Yes
Analyst FE	No	No	No	Yes	Yes	Yes

Table 8: Ratings and ICO success: Descriptive evidence

This table presents descriptive statistics for the relationship between ratings and ICO success. Panel A shows the success of ICOs that human analysts did or did not cover. Panel B links ICO success to the quantitative rating score. Panel C shows investor disagreement for ICOs with or without any reciprocal rating. Market Performance displays the ICO value on the 90th day post listing on CoinMarketCap divided by the amount raised during the ICO campaign, expressed in percent.

Panel A

Analyst Coverage	Total #	Funded #	Funded in %	MarketPerformance avg. in %	AmountRaised avg. in $	Ln(AmountRaised) avg. in $
No	2,961	858	28.98	116.53	19,326,718	4.27
Yes	2,378	1,033	43.44	45.68	12,692,225	6.60
Total	5,339	1,891	35.42	77.82	15,704,397	5.31

Panel B

AnalystRating Score	Total #	Funded #	Funded in %	MarketPerformance avg. in %	AmountRaised avg. in $	Ln(AmountRaised) avg. in $
3	211	46	21.80	5.58	23,140,030	3.10
4–6	250	52	20.80	42.71	6,313,606	2.98
7–9	486	174	35.80	60.53	14,316,201	5.44
10–12	1,071	545	50.89	57.01	12,164,440	7.75
13–15	754	392	51.99	105.84	31,994,976	8.14

Panel C

Reciprocal Rating	MarketPerformance avg. in %	Total #	Disagreement #	Disagreement in %	Disagreement with Avg. Rating ≥ 13 #	Disagreement with Avg. Rating ≥ 13 in %
Yes*	8.39	415	97	23.37	96	23.13
No**	55.79	1,963	316	16.10	272	13.86

*$ReciprocalRatingShare_j > 0$

** $ReciprocalRatingShare_j = 0$

Table 9: Ratings and ICO success

This table presents, in columns (1)-(3), marginal effects of logit regressions of Equation 5, where the dependent variable is the *Success* dummy. In columns (4)-(5), it presents coefficients of linear regressions of $MarketPerformance$. The controls for which coefficients are not shown for space reasons include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. The full table is available in the Online Appendix OA.2. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Success_j$			$MarketPerformance_j$	
	(1)	(2)	(3)	(4)	(5)
$\# Analysts_j$	0.045***	0.043***	0.036***	-0.005	-0.005
	(6.00)	(4.89)	(4.06)	(-0.81)	(-0.78)
$Benchy_j$	0.788***	0.717***	0.633***	0.208*	0.238*
	(8.94)	(5.31)	(4.61)	(1.79)	(1.84)
$AnalystRating_j$	0.112***	0.070**	0.072**	-0.006	-0.008
	(5.42)	(2.32)	(2.30)	(-0.19)	(-0.27)
$ReciprocalRatingShare_j$		0.018	0.038	-0.495**	-0.483**
		(0.06)	(0.13)	(-2.38)	(-2.33)
$PreviousRatings_j$		0.069	0.079	0.108	0.103
		(1.22)	(1.35)	(1.47)	(1.30)
$StarAnalysts_j$		-0.266	-0.263	0.246	0.280
		(-1.02)	(-0.99)	(0.87)	(0.99)
$AnalystDispersion_j$		0.009	-0.000	-0.011	-0.011
		(0.20)	(-0.01)	(-0.24)	(-0.26)
$AnalystExperience_j$		0.136	0.120	-0.037	-0.050
		(1.43)	(1.26)	(-0.38)	(-0.50)
$ReviewToneDispersion_j$		0.383	0.102	3.571	3.498
		(0.27)	(0.07)	(0.97)	(0.98)
$ReviewTone_j$		-1.233	-1.262	-1.026	-0.883
		(-1.05)	(-1.06)	(-0.90)	(-0.82)
$ReviewUncertainty_j$		-4.779*	-4.597	2.872	2.494
		(-1.67)	(-1.57)	(0.71)	(0.61)
$ReviewComplexity_j$		0.031	0.034*	-0.004	-0.005

		(1.54)	(1.65)	(-0.26)	(-0.34)
$ReviewLength_j$		0.096	0.107	-0.097*	-0.110*
		(1.10)	(1.18)	(-1.81)	(-1.90)
Observations	2330	1590	1590	717	717
R^2				0.158	0.164
Pseudo R^2	0.155	0.219	0.236		
VentureOffering Controls	No	Yes	Yes	Yes	Yes
WhitePaper Controls	No	Yes	Yes	Yes	Yes
SocialMedia Controls	No	No	Yes	No	Yes
MarketSentimet	No	No	Yes	No	Yes
Month Dummies	Yes	Yes	Yes	Yes	Yes

Table 10: ICO outcomes that deviate from what ratings predict

This table presents marginal effects of logit regressions in columns (1) to (3) and coefficients of linear regressions in columns (4) and (5) for Equation 6. The dependent variable is the *Disagreement* dummy which equals one if (i) analysts give an average $AnalystRating_j \geq 13$ and the ICO fails, or if (ii) analysts give an average $AnalystRating_j \leq 5$ and the ICO succeeds. In column (4), we restrict the sample to cases where the reciprocal ratings are on average greater than or equal to the average of non-reciprocal ratings for the same ICO. In column (5), we restrict the sample to ICOs where the average reciprocal rating is lower than the average of non-reciprocal ratings. All analyst variables are average values of every analyst that rates the ICO. The controls for which coefficients are not shown for space reasons include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. The full table is available in the Online Appendix OA.2. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Disagreement_j$				
	(1)	(2)	(3)	(4)	(5)
$\#\,Analysts_j$	-0.001	-0.006	-0.005	-0.002	-0.001
	(-0.12)	(-0.66)	(-0.61)	(-0.95)	(-0.21)
$StarAnalysts_j$	-0.432**	-0.230	-0.221	-0.128	0.526
	(-2.14)	(-0.80)	(-0.76)	(-0.66)	(1.49)
$PreviousRatings_j$	0.292***	0.299***	0.309***	0.026	0.078
	(5.09)	(4.24)	(4.32)	(0.75)	(0.86)
$ReciprocalRatingShare_j$	0.906***	0.799**	0.778**	0.379**	0.275
	(3.09)	(2.39)	(2.32)	(2.13)	(0.97)
$Benchy_j$	0.131	-0.144	-0.128	-0.205**	0.051
	(1.33)	(-1.07)	(-0.91)	(-2.33)	(0.41)
$AnalystDispersion_j$	-0.306***	-0.338***	-0.342***	-0.059*	-0.065*
	(-6.46)	(-5.69)	(-5.73)	(-1.76)	(-1.77)
$AnalystExperience_j$		-0.027	-0.038	0.021	-0.221**
		(-0.25)	(-0.34)	(0.26)	(-2.27)
$ReviewToneDispersion_j$		3.825**	4.199**	2.022*	1.556
		(2.12)	(2.27)	(1.78)	(1.39)
$ReviewTone_j$		4.615**	4.683**	0.430	2.632**
		(2.42)	(2.35)	(0.48)	(2.23)

$ReviewUncertainty_j$		-4.303	-3.971	5.134	-3.123
		(-1.01)	(-0.93)	(1.39)	(-0.76)
$ReviewComplexity_j$		0.049**	0.049**	0.028	-0.006
		(2.19)	(2.25)	(1.18)	(-0.36)
$ReviewLength_j$		-0.055	-0.068	-0.054	0.267**
		(-0.47)	(-0.57)	(-0.75)	(2.37)
Observations	2320	1592	1592	212	134
R^2				0.345	0.534
Pseudo R^2	0.148	0.173	0.180		
VentureOffering Controls	No	Yes	Yes	Yes	Yes
WhitePaper Controls	No	No	Yes	Yes	Yes
SocialMedia Controls	No	No	Yes	Yes	Yes
MarketSentimet	No	No	Yes	Yes	Yes
Month Dummies	Yes	Yes	Yes	Yes	Yes

Table 11: ICO scams

This table presents marginal effects of logit regressions analogous to Equation 5, with the dependent variable being the *Scam* dummy. The controls for which coefficients are not shown for space reasons include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Scam_j$		
	(1)	(2)	(3)
$\#Analysts_j$	0.037***	0.035***	0.034**
	(3.67)	(3.46)	(2.37)
$Benchy_j$	-0.195	-0.339*	-0.479
	(-1.08)	(-1.91)	(-1.49)
$AnalystRating_j$	0.057	0.060	0.063
	(1.39)	(1.19)	(0.66)
$StarAnalysts_j$		0.642*	0.021
		(1.71)	(0.02)
$PreviousRatings_j$		0.100	0.685***
		(1.06)	(3.60)
$ReciprocalRatingShare_j$		0.230	0.014
		(0.36)	(0.02)
$AnalystDispersion_j$		0.278***	0.569***
		(3.32)	(4.72)
$AnalystExperience_j$			0.145
			(0.60)
$ReviewToneDispersion_j$			5.810**
			(2.31)
$ReviewTone_j$			-1.537
			(-0.63)
$ReviewUncertainty_j$			-11.894
			(-1.59)
$ReviewComplexity_j$			0.005
			(0.09)
$ReviewLength_j$			-0.313

			(-1.36)
Observations	2114	2072	1371
Pseudo R^2	0.062	0.082	0.228
VentureOffering Controls	No	No	Yes
WhitePaper Controls	No	No	Yes
SocialMedia Controls	No	No	Yes
MarketSentimet	No	No	Yes
Month Dummies	Yes	Yes	Yes

Appendix

Table A1: Variable definitions

Variable	Definition
$\# Advisors_j$	Number of advisors who support an ICO.
$\# Analysts_j$	Number of analysts who rate an ICO.
$\# TeamMembers_j$	Number of team members in an ICO.
$AmountRaised_j$	Natural logarithm of 1 + \$ amount raised during the ICO campaign.
$AnalystDispersion_j$	Standard deviation of ratings within an ICO.
$AnalystExperience_i^{j-1}$	Natural logarithm of 1 + the number of ICOs that analyst i rated before providing a rating for ICO j.
$AnalystExperience_j$	Average experience of all analysts who rated ICO j.
$AnalystRating_{ij}$	The sum of team, vision, and product ratings for the respective ICO, ranging from 3 to 15.
$AnalystRating_j$	Average rating of ICO j by all analysts.
$AnalystRating(Team)_{ij}$/ $AnalystRating(Vision)_{ij}$/ $AnalystRating(Product)_{ij}$	Rating score for team/ vision/ product of an ICO, ranging from 1 (lowest) to 5 (highest).
$Benchy_j$	Machine-generated rating created by ICObench.com.
$Bitcointalk_j$	Dummy variable that equals 1 if the ICO is discussed on the forum Bitcointalk.org.
$Bonus_j$	Dummy variable that equals 1 for ICOs with a quantity discount at the token sale or a discount program for early-bird investors.
$Bounty_j$	Dummy variable that equals 1 for ICOs with incentives to promote social media presence.
$Disagreement(Buy)_j$	Dummy variable that equals 1 if an analyst gives a buy recommendation (*AnalystRating_j* ≥ 13) and the ICO fails.
$Disagreement(Sell)_j$	Dummy variable that equals 1 if an analyst gives a sell recommendation (*AnalystRating_j* ≤ 5) and the ICO succeeds.
$Disagreement_j$	Dummy variable that equals 1 if, (i) on average, analysts recommend buying (*AnalystRating_j* ≥ 13) and the ICO fails, or (ii) on average, analysts recommend selling (*AnalystRating_j* ≤ 5) and the ICO succeeds.
$Facebook_j$	Dummy variable that equals 1 if an ICO has a Facebook page.
$ForecastError_{ij}$	The distance of the $AnalystRating_{ij}$ from the highest (lowest) possible rating in the case of ICO success (failure).
$ForecastError_j$	A recursive average of the previous forecast errors of all analysts covering ICO j up to the rating issuance date.

$ForecastError_i^{j-1}$	A recursive average of all previous forecast errors for any analyst i up to the rating issuance date for ICO j.
$ForecastErrorOptimistic_i$	The distance of the highest possible rating score to the $AnalystRating_{ij}$, defined as $15 - AnalystRating_{ij}$, if the ICO was unsuccessful, and averaged over all ICOs j.
$ForecastErrorPessimistic_i$	The distance of the $AnalystRating_{ij}$ from the lowest possible rating score, defined as $AnalystRating_{ij} - 3$, if the ICO was successful, and averaged over all ICOs j.
$GitHubCommits_j$	The total amount of commits (project updates or code changes) on GitHub.com of ICO j before the ICO event ended.
$HardCap_j$	Dummy variable that equals 1 for ICOs that disclose a hard cap (a maximum amount of funds that the ICO is planning to raise).
IEO_j	Dummy variable that equals 1 for ICOs conducted on the platform of a cryptocurrency exchange (Initial Exchange Offerings).
KYC_j	Dummy variable that equals 1 for ICOs where investors are required to sign up to a whitelist using their wallet address to receive access to the ICO sale (Know Your Customer).
$MarketPerformance_j$	The value of market capitalization 90 days after listing on an exchange from CoinMarketCap.com divided by the amount raised during the campaign of ICO j. The variable is expressed in percent.
$MarketSentiment_j$	The average daily BTC return during the campaign of the ICO.
$Modified_{ij}$	Dummy variable that equals 1 if the rating for ICO j was modified by analyst i at any point in time.
$Month_{ij}$	Dummy variable for each month, indicating the month when a rating was given.
$Month_j$	Dummy variable for each month, indicating the month when an ICO was launched.
MVP_j	Dummy variable that equals 1 for ICOs with a prototype. This can be a version of a new product with sufficient features to satisfy early adopters (minimum viable product) or drafts of code on GitHub.com that are open to discussion by other GitHub users.
$OrderRank_{ij}$	The order rank of the rating by analyst i issued for ICO j in a given month.
$Presale_j$	Dummy variable that equals 1 if an ICO featured a token sale event that ran prior to the official ICO campaign.
$PreviousRatings_j$	Average past *AnalystRating* of all analysts that provide a rating for ICO j

$ReceivedRating_{ij}$/ $ReceivedRating(Team)_{ij}$/ $ReceivedRating(Vision)_{ij}$/ $ReceivedRating(Product)_{ij}$	Level of the rating when ReciprocalRating dummy equals 1, i.e., level of rating that the analyst of ICO j received for their own ICO from any team member of ICO j prior to the rating issuance date.
$ReciprocalRating_{ij}$	Dummy variable that equals 1 for reciprocal ratings. A rating is reciprocal when the corresponding analyst was a team member of another ICO project that previously received a rating by one of the team members of this new *ICO*. Table 2 represents a hypothetical illustration of our variable composition.
$ReciprocalRatingShare_j$	Share of reciprocal analysts who provide a rating for ICO j.
$RetentionRatio_j$	The percentage of tokens that is retained by the ICO members.
$ReviewComplexity_j$	The complexity of an analyst's review text, measured by the Gunning (1952) Fog index, and averaged together on the ICO level.
$ReviewLength_{ij}$	Natural logarithm of the number of total words in an analyst review. For the $ReviewLength_j$, we measure the natural logarithm of the average review text lengths for ICO j.
$ReviewTone_{ij}$	The tone of the analyst review text. Using the Loughran and McDonald (2011) *Positive* and *Negative* word-lists, the tone of a text is defined as the difference between the count of positive and negative words divided by the total number of words.
$ReviewTone_j$	The tone averaged across all analysts' review texts for ICO j.
$ReviewToneDispersion_j$	The standard deviation of $ReviewTone_{ij}$ within an ICO.
$ReviewUncertainty_j$	The uncertainty of the analysts' review texts, averaged together the on ICO level. Using the Loughran and McDonald (2011) *Uncertainty* word-list, the uncertainty of a text is defined as the count of uncertain words divided by the total number of words.
$Scam_j$	Dummy variable that equals 1 for ICO projects that intentionally defraud investors.
$SocialMediaComplexity_j$	The average of dictionary-based ratios that evaluate the use of complex language in all text messages on Bitcointalk before the ICO event ended. Using the Loughran and McDonald (2011) *Complexity* word-list, the complexity of a text is defined as the count of complex words divided by the total number of words.
$SocialMediaCount_j$	The total number of text messages on Bitcointalk before the ICO event ended.
$SocialMediaExtremeWords_j$	The average of dictionary-based ratios that evaluate the use of extreme language in all text messages on Bitcointalk before the ICO event ended. Using the Bochkay et al. (2020) *extreme* word-list, the extreme language of a text is defined as the count of extreme words divided by the total number of words.

$SocialMediaLength_j$	The total number of words on Bitcointalk before the ICO event ended.
$SocialMediaTechnicalWords_j$	The average of dictionary-based ratios that evaluate the use of technical language in all text messages on Bitcointalk before the ICO event ended. Using the Lyandres et al. (2022) *tech* word-list, the technical language of a text is defined as the count of technical words divided by the total number of words.
$SocialMediaTone_j$	The average of scores between -1 and 1 that evaluate the tone of all text messages on Bitcointalk before the ICO event ended. Using the Loughran and McDonald (2011) *Positive* and *Negative* word-lists, the tone of a text is defined as the difference between the count of positive and negative words divided by the total number of words.
$SocialMediaUncertainty_j$	The average of dictionary-based ratios that evaluate the use of uncertain language in all text messages on Bitcointalk before the ICO event ended. Using the Loughran and McDonald (2011) *Uncertainty* word-list, the uncertainty of a text is defined as the count of uncertain words divided by the total number of words.
$StarAnalysts_{ij}$	Dummy variable that equals 1 when ICO j was rated by one of the top 30 analysts i according to a ranking on ICObench.com.
$StarAnalysts_j$	Share of the top 30 analysts that provide a rating for ICO j.
$Success_j$	Dummy variable that equals 1 for ICOs that completed the token sale and collected (at least $1 in) funding.
$VestingDisclosure_j$	Dummy variable that equals 1 for ICOs that disclose vesting information in their whitepapers.
$WhitePaperComplexity_j$	A dictionary-based ratio that evaluates the use of complex language in a whitepaper. Using the Loughran and McDonald (2011) *Complexity* word-list, the complexity of a text is defined as the count of complex words divided by the total number of words.
$WhitePaperLength_j$	The natural logarithm of (1 + total words of the white paper), set to 0 if no whitepaper could be found.
$WhitePaperTechnicalWords_j$	A dictionary-based ratio that evaluates the use of technical language in a whitepaper. Using the Lyandres et al. (2022) *tech* word-list, the technical language of a text is defined as the count of technical words divided by the total number of words.
$WhitePaperTone_j$	A score between -1 and 1 that evaluates the tone of the whitepaper. Using the Loughran and McDonald (2011) *Positive* and *Negative* word-lists, the tone of a text is defined as the difference between the count of positive and negative words divided by the total number of words.

$WhitePaperUncertainty_j$	A dictionary-based ratio that evaluates the use of uncertain language in a whitepaper by using the Loughran and McDonald (2011) *Uncertainty* word-list.

Table A2: Ratings and ICO success: An alternative success measure

This table presents linear regression results for Equation 5. The dependent variable is the natural logarithm of (1 + the amount raised) by an ICO during the campaign. The controls for which coefficients are not shown for space reasons include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$AmountRaised_j$		
	(1)	(2)	(3)
$\#Analysts_j$	0.142***	0.133***	0.097***
	(7.81)	(7.20)	(4.92)
$Benchy_j$	2.120***	2.005***	1.339***
	(9.66)	(8.81)	(4.06)
$AnalystRating_j$	0.320***	0.327***	0.200***
	(6.06)	(5.60)	(2.73)
$StarAnalysts_j$		-0.263	-0.534
		(-0.56)	(-0.78)
$PreviousRatings_j$		0.070	0.144
		(0.64)	(1.01)
$ReciprocalRatingShare_j$		-0.010	0.303
		(-0.01)	(0.35)
$AnalystDispersion_j$		0.206*	0.046
		(1.85)	(0.35)
$AnalystExperience_j$			0.299
			(1.15)
$ReviewToneDispersion_j$			2.318
			(0.54)
$ReviewTone_j$			-3.117
			(-0.97)
$ReviewUncertainty_j$			-12.581
			(-1.64)
$ReviewComplexity_j$			0.079
			(1.48)
$ReviewLength_j$			0.274

			(1.08)
Observations	2378	2322	1633
R^2	0.209	0.214	0.302
VentureOffering Controls	No	No	Yes
WhitePaper Controls	No	No	Yes
SocialMedia Controls	No	No	Yes
MarketSentimet	No	No	Yes
Month Dummies	Yes	Yes	Yes

Table A3: ICO outcomes that deviate from what ratings predict

This table presents marginal effects of logit regressions for Equation 6. The dependent variable is the *Disagreement*(*buy*) dummy, which equals one if analysts recommend buying (average $AnalystRating_j \geq 13$) and the ICO fails in columns (1) and (2), and *Disagreement*(*sell*) dummy, which equals one if analysts recommend not buying (average $AnalystRating_j \leq 5$) and the ICO succeeds in columns (3) and (4). All analyst variables are average values over all analysts that rate the ICO. Control variables for which coefficients are not shown for space reasons include Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. All variables are defined in Table A1. *t*-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Disagreement(Buy)_j$		$Disagreement(Sell)_j$	
	(1)	(2)	(3)	(4)
$\#Analysts_j$	-0.007	-0.006	-0.000	-0.004
	(-0.77)	(-0.69)	(-0.01)	(-0.11)
$Benchy_j$	-0.138	-0.102	-0.204	-0.628
	(-0.95)	(-0.67)	(-0.41)	(-1.14)
$StarAnalysts_j$	-0.178	-0.145	0.216	0.404
	(-0.57)	(-0.45)	(0.29)	(0.44)
$PreviousRatings_j$	0.388***	0.407***	-0.159	-0.177
	(4.77)	(4.93)	(-1.17)	(-1.18)
$ReciprocalRatingShare_j$	0.856**	0.825**	-1.026	-1.221
	(2.44)	(2.34)	(-0.54)	(-0.52)
$AnalystDispersion_j$	-0.356***	-0.367***	-0.208	-0.259
	(-5.74)	(-5.86)	(-0.90)	(-1.12)
$AnalystExperience_j$	-0.038	-0.041	-0.229	-0.434
	(-0.33)	(-0.34)	(-0.53)	(-0.94)
$ReviewToneDispersion_j$	5.721***	6.384***	-2.280	1.191
	(2.91)	(3.23)	(-0.37)	(0.19)
$ReviewTone_j$	8.339***	8.626***	-12.782***	-16.791***
	(4.72)	(4.74)	(-3.17)	(-3.22)
$ReviewUncertainty_j$	-6.529	-6.295	8.750*	10.776**
	(-1.39)	(-1.33)	(1.69)	(2.03)

$ReviewComplexity_j$	0.049**	0.050**	-0.020	0.013
	(2.06)	(2.17)	(-0.24)	(0.18)
$ReviewLength_j$	0.028	0.014	0.072	-0.034
	(0.22)	(0.10)	(0.25)	(-0.09)
Observations	1569	1569	1003	1003
Pseudo R^2	0.207	0.217	0.258	0.342
VentureOffering Controls	Yes	Yes	Yes	Yes
WhitePaper Controls	No	Yes	No	Yes
SocialMedia Controls	No	Yes	No	Yes
MarketSentimet	No	Yes	No	Yes
Month Dummies	Yes	Yes	Yes	Yes

Online Appendix

OA.1 Reciprocal vs. non-reciprocal ratings

In this Online Appendix, we calculate a separate average rating score for reciprocal and non-reciprocal ratings. In particular, we calculate for each ICO an average rating based on non-reciprocal ratings and an average rating based on reciprocal ratings. Naturally, because only a minority of ICOs have reciprocal ratings, the number of observations for the latter is lower. Based on the average non-reciprocal rating score and the average reciprocal rating score, we also redefine the $Disagreement$ dummy. In particular, the $Disagreement(NonReciprocal)$ dummy ($Disagreement(Reciprocal)$ dummy) equals one if (i) nonreciprocal analysts (reciprocal analysts) give an average $NonReciprocalRating_j \geq 13$ ($ReciprocalRating_j \geq 13$) and the ICO fails, or if (ii) nonreciprocal analysts (reciprocal analysts) give an average $NonReciprocalRating_j \leq 5$ ($ReciprocalRating_j \leq 5$) and the ICO succeeds.

We present the results in Table OA1. We find that the average non-reciprocal rating score predicts ICO success. The average reciprocal rating score does not predict ICO success even when not controlling for ICO characteristics. Moreover, we find that a market disagreement with non-reciprocal ratings is not correlated with the share of reciprocal ratings, but there is a strong correlation between the share of reciprocal ratings and market disagreement with reciprocal ratings.

Table OA1: Reciprocal vs. non-reciprocal ratings

This table presents marginal effects of logit regressions for Equation 5 and Equation 6. The dependent variable in Panel A is the *Success* dummy, which equals one if the ICO was successful in obtaining some funding. In Panel B, it is the $Disagreement(Reciprocal)$ dummy, which equals one if (i) analysts give a *reciprocal* $AnalystRating_j \geq 13$ and the ICO fails, or if (ii) analysts give a *reciprocal* $AnalystRating_j \leq 5$ and the ICO succeeds. The variable $Disagreement(NonReciprocal)$ is likewise based on *non-reciprocal* analyst ratings. The analyst variables are average values over all analysts that rate the ICO. The controls for which coefficients are not shown for space reasons include AnalystRating, Benchy, PreviousRating, StarAnalyst, #Analysts, AnalystDispersion, AnalystExperience, ReviewToneDispersion, ReviewTone, ReviewUncertainty, ReviewComplexity, ReviewLength (denoted as **Rating Controls**), Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

Panel A

	$Success_j$			
	(1)	(2)	(3)	(4)
$NonReciprocalRating_j$	0.109***	0.071**		
	(5.32)	(2.30)		
$ReciprocalRating_j$			0.065	0.031
			(1.18)	(0.38)
Observations	2298	1565	391	359
Pseudo R^2	0.156	0.237	0.140	0.284
Ratings Controls	No	Yes	No	Yes
VentureOffering Controls	No	Yes	No	Yes
WhitePaper Controls	No	Yes	No	Yes
SocialMedia Controls	No	Yes	No	Yes
MarketSentimet	No	Yes	No	Yes
Month Dummies	Yes	Yes	Yes	Yes

Panel B

	Disagreement $(NonReciprocal)_j$		*Disagreement* $(Reciprocal)_j$	
	(1)	(2)	(3)	(4)
$ReciprocalRatingShare_j$	-0.204	-0.242	6.120***	6.438***
	(-0.61)	(-0.72)	(11.34)	(10.96)
Observations	1592	1592	1400	1400
Pseudo R^2	0.164	0.172	0.385	0.407
Ratings Controls	Yes	Yes	Yes	Yes
VentureOffering Controls	Yes	Yes	Yes	Yes
WhitePaper Controls	No	Yes	No	Yes
SocialMedia Controls	No	Yes	No	Yes
MarketSentimet	No	Yes	No	Yes
Month Dummies	Yes	Yes	Yes	Yes

OA.2 Actual versus expected quid pro quo

In this Online Appendix, we generate a modified version of our reciprocal dummy, which equals one if an analyst launches their own ICO at a later stage, i.e. expecting a quid pro quo in the future, and zero otherwise. Accordingly, we define an $ExpectedOwnICORatingShare$ for each ICO. While both (i) actual reciprocal ratings (as measured in the main text) and (ii) expected reciprocal ratings (as measured with the modified dummy) might be biased, there is an important difference between the two: in case of (i), the information of reciprocity is available to the market, while it is not for case (ii). In fact, any rating could potentially be biased due to an analyst's hope of a quid pro quo in the future. We present results of the analysis of the three success measures – (unconditional) success, long-term success, and conditional success – below in Table OA2. We show the most saturated model. We find that markets understand the potential bias of actual reciprocal ratings (that could be easily identified as being reciprocal). Investors do not seem to discount ICOs with a higher share of ratings that have no actual reciprocity, but only a "perfect foresight" reciprocity with the future actions of an analyst. Potentially, these analysts might be seen as very informed agents who do not reciprocate ratings, but run their own ICOs.

Table OA2: Actual versus expected quid pro quo

This table presents coefficients of linear regressions and marginal effects of logit regressions for Equation 5 and Equation 6. The dependent variables are the *Success* dummy, which equals one if the ICO was successful in obtaining any funding (columns (1) and (2)), *MarketPerformance*, defined as the value of market capitalization 90 days after listing on an exchange relative to the amount raised during the campaign (columns (3) and (4)), and the *Disagreement* dummy, which equals one if (i) analysts give an average $AnalystRating_j \geq 13$ and the ICO fails, or if (ii) analysts give an average $AnalystRating_j \leq 5$ and the ICO succeeds (columns (5) and (6)). The main explanatory variables are ExpectedOwnICORatingShare (the share of analysts that launch their own ICO at a later stage and therefore may expect reciprocity) and ReciprocalRatingShare (the share of actual reciprocal ratings relative to all ratings in ICO j). All analyst variables are average values over all analysts that rate the ICO. The controls for which coefficients are not shown for space reasons include AnalystRating, Benchy, PreviousRating, StarAnalyst, #Analysts, AnalystDispersion, AnalystExperience, ReviewToneDispersion, ReviewTone, ReviewUncertainty, ReviewComplexity, ReviewLength (denoted as **Rating Controls**), Presale, Bounty, MVP, KYC, Bonus, IEO, RetentionRatio, GitHubCommits, HardCap, VestingDisclosure, #Advisors, and #TeamMembers (denoted as **VentureOffering Controls**), whitepaper tone, whitepaper uncertainty, whitepaper complexity, whitepaper tech ratio, and the length of the white paper (denoted as **WhitePaper Controls**), Bitcointalk, Facebook, the number of social media messages, the length of social media messages, and textual analysis of social media messages (incl. tone, uncertainty, complexity, technical, and extreme language) (denoted as **SocialMedia Controls**), and the average daily BTC return during the campaign of the ICO (denoted as **MarketSentiment**). All specifications include month dummies. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Success_j$		$MarketPerformance_j$		$Disagreement_j$	
	(1)	(2)	(3)	(4)	(5)	(6)
$ExpectedOwnICORatingShare_j$	-0.129		0.053		-0.236	
	(-0.64)		(0.39)		(-1.09)	
$ReciprocalRatingShare_j$		0.038		-0.483**		0.778**
		(0.13)		(-2.33)		(2.32)
Observations	1590	1590	717	717	1592	1592
R^2			0.160	0.164		
Pseudo R^2	0.236	0.236			0.177	0.180
Ratings Controls	Yes	Yes	Yes	Yes	Yes	Yes
VentureOffering Controls	Yes	Yes	Yes	Yes	Yes	Yes
WhitePaper Controls	Yes	Yes	Yes	Yes	Yes	Yes
SocialMedia Controls	Yes	Yes	Yes	Yes	Yes	Yes
MarketSentimet	Yes	Yes	Yes	Yes	Yes	Yes
Month Dummies	Yes	Yes	Yes	Yes	Yes	Yes

Table OA3: Rating determinants - Full view

This table presents linear regression results for Equation 1. The dependent variable is the total rating score that an analyst gave to an ICO. The specification in column (5) includes month dummies and *ICO* fixed effects. All variables are defined in Table A1. t-statistics are given in parentheses. Standard errors are clustered at the ICO and analyst levels. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$AnalystRating_{ij}$				
	(1)	(2)	(3)	(4)	(5)
Rating Controls					
$Benchy_j$	1.518***	1.599***		0.706***	
	(10.43)	(11.42)		(5.11)	
$StarAnalyst_{ij}$		-0.802***	-0.721***	-0.735***	-0.442***
		(-3.80)	(-3.43)	(-3.51)	(-2.68)
$ForecastError_i^{j-1}$		-0.161***	-0.130***	-0.143***	-0.080**
		(-3.66)	(-3.09)	(-3.49)	(-2.27)
$AnalystExperience_i^{j-1}$		0.063	0.037	0.044	-0.009
		(0.83)	(0.51)	(0.58)	(-0.12)
VentureOffering Controls					
$Presale_j$			0.301**	0.257**	
			(2.41)	(2.13)	
$Bounty_j$			0.265**	0.209*	
			(2.13)	(1.73)	
MVP_j			0.214	-0.019	
			(1.52)	(-0.14)	
KYC_j			0.722***	0.509***	
			(4.46)	(3.24)	
$Bonus_j$			0.158	0.143	
			(1.39)	(1.28)	
IEO_j			1.256***	0.959***	
			(4.76)	(3.66)	
$RetentionRatio_j$			0.007**	0.006**	
			(2.30)	(1.97)	
$GitHubCommits_j$			0.025	0.007	
			(1.64)	(0.45)	

$HardCap_j$	0.278	0.275
	(1.59)	(1.60)
$VestingDisclosure_j$	-0.139	-0.120
	(-1.16)	(-1.05)
$\#Advisors_j$	0.448***	0.340***
	(6.80)	(5.17)
$\#TeamMembers_j$	0.193***	0.106**
	(3.58)	(2.03)
WhitePaper Controls		
$WhitePaperLength_j$	0.007	0.005
	(0.19)	(0.15)
$WhitePaperTone_j$	-3.734	-3.518
	(-0.63)	(-0.62)
$WhitePaperUncertainty_j$	-1.244	-2.521
	(-0.14)	(-0.28)
$WhitePaperComplexity_j$	2.710	3.173
	(0.32)	(0.37)
$WhitePaperTechnicalWords_j$	2.887*	2.931*
	(1.75)	(1.87)
SocialMedia Controls		
$Bitcointalk_j$	-0.264	-0.263
	(-1.47)	(-1.45)
$Facebook_j$	0.040	0.042
	(0.17)	(0.19)
$SocialMediaCount_j$	0.226**	0.202**
	(2.58)	(2.35)
$SocialMediaLength_j$	-0.096	-0.100*
	(-1.56)	(-1.66)
$SocialMediaTone_j$	7.212	5.878
	(0.21)	(0.16)
$SocialMediaUncertainty_j$	23.483	21.977
	(0.43)	(0.41)
$SocialMediaComplexity_j$	0.003	0.003
	(1.16)	(1.18)
$SocialMediaTechnicalWords_j$	-0.944	-0.708
	(-0.13)	(-0.10)

$SocialMediaExtremeWords_j$			-18.318	-26.537	
			(-0.11)	(-0.16)	
MarketSentiment					
$MarketSentiment_j$			8.915	7.350	
			(1.30)	(1.10)	
Observations	13834	12460	11257	11257	11698
R^2	0.121	0.146	0.109	0.123	0.528
Month Dummies	No	No	No	No	Yes
ICO FE	No	No	No	No	Yes

Table OA4: Ratings and ICO success - Full view

This table presents, in columns (1)-(3), marginal effects of logit regressions of Equation 5, where the dependent variable is the *Success* dummy. In columns (4)-(5), it presents coefficients of linear regressions of $MarketPerformance$. All analyst variables are average values over all analysts that rate the ICO. All specifications include month dummies. As the logit model predicts failure perfectly in some months, we lose a few observations from the inclusion of month fixed effects. All variables are defined in the paper's appendix, in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Success_j$			$MarketPerformance_j$	
	(1)	(2)	(3)	(4)	(5)
Rating Controls					
$\#Analysts_j$	0.045***	0.043***	0.036***	-0.005	-0.005
	(6.00)	(4.89)	(4.06)	(-0.81)	(-0.78)
$Benchy_j$	0.788***	0.717***	0.633***	0.208*	0.238*
	(8.94)	(5.31)	(4.61)	(1.79)	(1.84)
$AnalystRating_j$	0.112***	0.070**	0.072**	-0.006	-0.008
	(5.42)	(2.32)	(2.30)	(-0.19)	(-0.27)
$ReciprocalRatingShare_j$		0.018	0.038	-0.495**	-0.483**
		(0.06)	(0.13)	(-2.38)	(-2.33)
$PreviousRatings_j$		0.069	0.079	0.108	0.103
		(1.22)	(1.35)	(1.47)	(1.30)
$StarAnalysts_j$		-0.266	-0.263	0.246	0.280
		(-1.02)	(-0.99)	(0.87)	(0.99)
$AnalystDispersion_j$		0.009	-0.000	-0.011	-0.011
		(0.20)	(-0.01)	(-0.24)	(-0.26)
$AnalystExperience_j$		0.136	0.120	-0.037	-0.050
		(1.43)	(1.26)	(-0.38)	(-0.50)
$ReviewToneDispersion_j$		0.383	0.102	3.571	3.498
		(0.27)	(0.07)	(0.97)	(0.98)
$ReviewTone_j$		-1.233	-1.262	-1.026	-0.883
		(-1.05)	(-1.06)	(-0.90)	(-0.82)
$ReviewUncertainty_j$		-4.779*	-4.597	2.872	2.494
		(-1.67)	(-1.57)	(0.71)	(0.61)
$ReviewComplexity_j$		0.031	0.034*	-0.004	-0.005
		(1.54)	(1.65)	(-0.26)	(-0.34)
$ReviewLength_j$		0.096	0.107	-0.097*	-0.110*
		(1.10)	(1.18)	(-1.81)	(-1.90)
VentureOffering Controls					
$Presale_j$		-0.004	-0.026	0.008	0.034
		(-0.03)	(-0.19)	(0.06)	(0.28)

$Bounty_j$	-0.198	-0.313**	-0.193	-0.157
	(-1.46)	(-2.22)	(-1.48)	(-0.87)
MVP_j	-0.325**	-0.324**	-0.102	-0.109
	(-2.11)	(-2.06)	(-1.04)	(-1.09)
KYC_j	-0.372**	-0.325*	-0.192	-0.165
	(-2.24)	(-1.87)	(-1.30)	(-1.12)
$Bonus_j$	-0.799***	-0.832***	0.109	0.131
	(-5.27)	(-5.36)	(0.76)	(0.91)
IEO_j	1.184***	1.230***	1.461**	1.294*
	(3.88)	(3.84)	(2.06)	(1.78)
$RetentionRatio_j$	0.002	0.004	0.005	0.004
	(0.71)	(1.38)	(1.29)	(1.07)
$GitHubCommits_j$	0.045**	0.043**	0.006	0.003
	(2.39)	(2.27)	(0.30)	(0.15)
$HardCap_j$	0.740***	0.635***	-0.135	-0.100
	(4.94)	(3.74)	(-0.91)	(-0.73)
$VestingDisclosure_j$	0.237	0.208	-0.338***	-0.337***
	(1.60)	(1.40)	(-2.76)	(-2.82)
$\#Advisors_j$	0.152**	0.137**	-0.064	-0.068
	(2.29)	(2.01)	(-1.34)	(-1.41)
$\#TeamMembers_j$	0.150**	0.141**	-0.105	-0.101
	(2.45)	(2.26)	(-1.46)	(-1.32)
WhitePaper Controls				
$WhitePaperLength_j$	0.013	0.013	-0.081*	-0.083*
	(0.29)	(0.28)	(-1.92)	(-1.80)
$WhitePaperTone_j$	-3.906	-3.468	11.317*	10.778
	(-0.49)	(-0.43)	(1.74)	(1.60)
$WhitePaperUncertainty_j$	-5.283	-3.577	15.894	15.364
	(-0.44)	(-0.30)	(1.51)	(1.47)
$WhitePaperComplexity_j$	7.818	9.881	2.280	2.386
	(0.64)	(0.79)	(0.30)	(0.32)
$WhitePaperTechnicalWords_j$	0.426	0.047	6.024*	6.317*
	(0.20)	(0.02)	(1.88)	(1.86)
SocialMedia Controls				
$Bitcointalk_j$		0.417**		-0.016
		(2.41)		(-0.08)
$Facebook_j$		-0.034		-0.176
		(-0.16)		(-0.76)
$SocialMediaCount_j$		0.399***		0.036
		(3.63)		(0.48)
$SocialMediaLength_j$		-0.220***		-0.047
		(-2.93)		(-0.82)
$SocialMediaTone_j$		-0.330		9.993
		(-0.02)		(0.35)

$SocialMediaUncertainty_j$			46.556		18.753
			(0.89)		(0.63)
$SocialMediaComplexity_j$			-0.006		0.002
			(-1.54)		(0.41)
$SocialMediaTechnicalWords_j$			10.885		-5.938
			(1.55)		(-1.14)
$SocialMediaExtremeWords_j$			7.789		114.718
			(0.07)		(0.71)
MarketSentiment					
$MarketSentiment_j$			9.811		-8.513
			(0.98)		(-1.50)
Observations	2330	1590	1590	717	717
R^2				0.158	0.164
Pseudo R^2	0.155	0.219	0.236		
Month Dummies	Yes	Yes	Yes	Yes	Yes

Table OA5: ICO outcomes that deviate from what ratings predict - Full view

This table presents marginal effects of logit regressions in columns (1) to (3) and coefficients of linear regressions in columns (4) and (5) for Equation 6. The dependent variable is the *Disagreement* dummy which equals one if (i) analysts give an average $AnalystRating_j \geq 13$ and the ICO fails, or if (ii) analysts give an average $AnalystRating_j \leq 5$ and the ICO succeeds. In column (4), we restrict the sample to cases where the reciprocal ratings are on average greater than or equal to the average of non-reciprocal ratings for the same ICO. In column (5), we restrict the sample to ICOs where the average reciprocal rating is lower than the average of non-reciprocal ratings. All analyst variables are average values of every analyst that rates the ICO. All specifications include month dummies. All variables are defined in Table A1. t-statistics based on robust standard errors are reported in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels.

	$Disagreement_j$				
	(1)	(2)	(3)	(4)	(5)
Rating Controls					
$\#Analysts_j$	-0.001	-0.006	-0.005	-0.002	-0.001
	(-0.12)	(-0.66)	(-0.61)	(-0.95)	(-0.21)
$StarAnalysts_j$	-0.432**	-0.230	-0.221	-0.128	0.526
	(-2.14)	(-0.80)	(-0.76)	(-0.66)	(1.49)
$PreviousRatings_j$	0.292***	0.299***	0.309***	0.026	0.078
	(5.09)	(4.24)	(4.32)	(0.75)	(0.86)
$ReciprocalRatingShare_j$	0.906***	0.799**	0.778**	0.379**	0.275
	(3.09)	(2.39)	(2.32)	(2.13)	(0.97)
$Benchy_j$	0.131	-0.144	-0.128	-0.205**	0.051
	(1.33)	(-1.07)	(-0.91)	(-2.33)	(0.41)
$AnalystDispersion_j$	-0.306***	-0.338***	-0.342***	-0.059*	-0.065*
	(-6.46)	(-5.69)	(-5.73)	(-1.76)	(-1.77)
$AnalystExperience_j$		-0.027	-0.038	0.021	-0.221**
		(-0.25)	(-0.34)	(0.26)	(-2.27)
$ReviewToneDispersion_j$		3.825**	4.199**	2.022*	1.556
		(2.12)	(2.27)	(1.78)	(1.39)
$ReviewTone_j$		4.615**	4.683**	0.430	2.632**
		(2.42)	(2.35)	(0.48)	(2.23)
$ReviewUncertainty_j$		-4.303	-3.971	5.134	-3.123
		(-1.01)	(-0.93)	(1.39)	(-0.76)
$ReviewComplexity_j$		0.049**	0.049**	0.028	-0.006
		(2.19)	(2.25)	(1.18)	(-0.36)
$ReviewLength_j$		-0.055	-0.068	-0.054	0.267**
		(-0.47)	(-0.57)	(-0.75)	(2.37)
VentureOffering Controls					
$Presale_j$		-0.287*	-0.264*	-0.162**	-0.162

	(-1.89)	(-1.71)	(-2.18)	(-1.49)
$Bounty_j$	0.100	0.195	0.070	-0.011
	(0.65)	(1.18)	(0.98)	(-0.10)
MVP_j	0.350**	0.372**	0.104	-0.004
	(2.05)	(2.17)	(1.38)	(-0.04)
KYC_j	0.355*	0.332	0.171**	-0.212
	(1.74)	(1.63)	(2.18)	(-1.20)
$Bonus_j$	0.690***	0.696***	-0.050	0.217**
	(4.33)	(4.27)	(-0.61)	(2.13)
IEO_j	-0.484*	-0.378	0.077	0.071
	(-1.71)	(-1.23)	(0.40)	(0.29)
$RetentionRatio_j$	0.010***	0.009***	0.001	0.003
	(2.72)	(2.59)	(0.30)	(1.39)
$GitHubCommits_j$	-0.036	-0.036	0.007	0.002
	(-1.58)	(-1.57)	(0.70)	(0.12)
$HardCap_j$	-0.367**	-0.274	-0.132	-0.238
	(-2.11)	(-1.42)	(-1.12)	(-1.58)
$VestingDisclosure_j$	-0.027	-0.024	-0.054	-0.060
	(-0.18)	(-0.14)	(-0.67)	(-0.64)
$\#Advisors_j$	0.004	0.013	-0.067	-0.051
	(0.06)	(0.18)	(-1.45)	(-0.94)
$\#TeamMembers_j$	0.108	0.112	-0.002	-0.017
	(1.49)	(1.52)	(-0.06)	(-0.35)
WhitePaper Controls				
$WhitePaperLength_j$		0.008	-0.008	-0.056*
		(0.17)	(-0.35)	(-1.97)
$WhitePaperTone_j$		-5.463	-4.253	0.535
		(-0.61)	(-1.07)	(0.09)
$WhitePaperUncertainty_j$		3.649	0.941	20.241**
		(0.28)	(0.16)	(2.44)
$WhitePaperComplexity_j$		-14.042	-3.364	-1.177
		(-0.92)	(-0.46)	(-0.18)
$WhitePaperTechnicalWords_j$		0.325	-0.014	1.339
		(0.14)	(-0.01)	(0.89)
SocialMedia Controls				
$Bitcointalk_j$		-0.391*	-0.052	0.007
		(-1.91)	(-0.40)	(0.05)
$Facebook_j$		0.037	0.074	-0.046
		(0.15)	(0.51)	(-0.19)
$SocialMediaCount_j$		-0.075	0.015	-0.081
		(-0.65)	(0.26)	(-0.99)
$SocialMediaLength_j$		0.062	-0.010	0.052
		(0.81)	(-0.21)	(0.84)
$SocialMediaTone_j$		32.952*	-10.335	-33.211

			(1.88)	(-0.27)	(-0.53)
$SocialMediaUncertainty_j$			17.404	37.631	-282.573***
			(0.33)	(0.46)	(-2.73)
$SocialMediaComplexity_j$			-0.000	-0.001	0.002
			(-0.22)	(-0.24)	(1.42)
$SocialMediaTechnicalWords_j$			1.525	1.434	0.204
			(0.30)	(0.50)	(0.03)
$SocialMediaExtremeWords_j$			-238.776**	-9.831	43.790
			(-2.48)	(-0.06)	(0.16)
MarketSentiment					
$MarketSentiment_j$			7.708	-0.328	-0.120
			(0.64)	(-0.07)	(-0.02)
Observations	2320	1592	1592	212	134
R^2				0.345	0.534
Pseudo R^2	0.148	0.173	0.180		
Month Dummies	Yes	Yes	Yes	Yes	Yes