

Digital Tulips? Returns to Investors in Initial Coin Offerings

Hugo Benedetti

Boston College

hugo.benedetti@bc.edu

Leonard Kostovetsky

Boston College

kostovet@bc.edu

Draft Date: May 20th, 2018

Abstract: Initial coin offerings (ICOs), sales of cryptocurrency tokens to the general public, have recently been used as a source of crowdfunding for startups in the technology and blockchain industries. We create a dataset on 4,003 executed and planned ICOs, which raised a total of \$12 billion in capital, nearly all since January 2017. We find evidence of significant ICO underpricing, with average returns of 179% from the ICO price to the first day's opening market price, over a holding period that averages just 16 days. Even after imputing returns of -100% to ICOs that don't list their tokens within 60 days and adjusting for the returns of the asset class, the representative ICO investor earns 82%. After trading begins, tokens continue to appreciate in price, generating average buy-and-hold abnormal returns of 48% in the first 30 trading days. We also study the determinants of ICO underpricing and relate cryptocurrency prices to Twitter followers and activity. While our results could be an indication of bubbles, they are also consistent with high compensation for risk for investing in unproven pre-revenue platforms through unregulated offerings.

1. Introduction

The traditional sources for seed and early-stage funding have recently been supplemented with *crowdfunding*: raising money from many small investors, in small amounts, over the Internet. Early on, crowdfunding was provided in exchange for future rewards or deals on products (e.g., Indiegogo, Kickstarter), and more recently for securities (equity crowdfunding). Advances in the blockchain technology have also led to a new hybrid form of crowdfunding: *token offerings*, also known as *initial coin offerings* (ICOs), which are the subject of this paper.

Tokens are *cryptocurrencies*, digital currencies for which all records and transaction data are protected by cryptographic methods. Entrepreneurs issue branded tokens to raise capital to create an online platform or ecosystem, in which all transactions require the use of that native token. In the 16 months since January 2017, over 1,000 startups successfully raised a total of about \$12 billion using ICOs. Soon after the ICO, tokens are usually listed on one or more online exchanges, providing liquidity for token-holders and a signal of the quality and future prospects of the platform. The token markets are fairly liquid, with \$3 million in average daily volume (for the average token) in the first thirty days of trading.

The closest analogue to the ICO is the Initial Public Offering (IPO) of equity. In addition to selling a different asset, two key differences between ICOs and IPOs are: (1) ICO firms are much younger and smaller, typically in the earliest stage of a firm's life cycle, and (2) ICO firms do not use an underwriter to help determine value and attract

buyers. As a result, it is not clear how two well-known characteristics of the IPO market, underpricing (Beatty and Ritter, 1986) and post-IPO underperformance (Ritter, 1991) translate to ICOs and listed tokens.

In this paper, we study the market for crypto-tokens, focusing on how entrepreneurs determine the price for tokens, the returns to investors from buying tokens during an ICO and selling them once they are listed on an exchange, and the returns to investors from investing in tokens on the listing date and holding them for various fixed time horizons. We also use data from Twitter accounts of cryptocurrency firms to investigate the relationship between Twitter followers and activity, and market prices, and to measure the attrition rate of crypto-companies after completion of the ICO. Our paper aims to provide a comprehensive analysis of how startups in this industry transition and perform from birth, through the offering, to the listing, and beyond.

Consistent with the IPO literature (directionally but not in degree), startups sell their tokens during the ICO at a significant discount to the opening market price, generating an average return for ICO investors of 179%, accrued over an average holding period of 16 days from the ICO end date to the listing date. Even after imputing large negative returns to tokens that are not listed within 60 days, the representative investor nearly doubles her money by investing in an ICO. Because some tokens are illiquid, we examine, and find our results are robust to, gradual liquidation of tokens after listing. ICO returns are positive and significant when the cryptocurrency asset class is performing

well but also when it is performing poorly, indicating that the results are not an artifact of the strong overall performance in cryptocurrencies over the last few years.

The degree of underpricing is much larger than that for IPOs but is not surprising considering the entrepreneur’s lack of expertise in determining market demand for the token/platform, greater uncertainty about the value of a startup company whose platform is typically still in the idea stage, and the urgency in distributing tokens to allow the platform to function. We examine the determinants of ICO underpricing and find that the returns to ICO investors have declined over time, suggesting that firms are learning from prior offerings about market demand for different kinds of platforms. Pre-sales, also called pre-ICOs, have gotten more popular over time and we find that they also act to reduce underpricing. We also find that the nominal ICO price, in dollars, is negatively correlated with underpricing and ICO returns, indicating a tendency of token prices to drift toward a “normal” nominal price when they are traded. On the other hand, the age of the firm (based on its Twitter account activation date), a proxy for information asymmetry, is not related to ICO underpricing, a significant difference between ICOs and IPOs.

Next, we use Twitter data to gain a better understanding of the life cycle of ICO firms and the relationship between market values and Twitter activity. The coefficient from a simple regression of (the natural logarithm of) cryptocurrency market capitalization on (the natural logarithm of) the number of Twitter followers is 1.2 (with a standard error of 0.033). The fact that this coefficient is significantly greater than one

is indicative of the positive network externalities that form when a platform reaches a critical mass of users.

We use intensity of tweets from the cryptocurrency official Twitter account after the ICO to estimate that the survival rate for startups after 120 days (from the end of the ICO) is only 44.2%, assuming that all firms inactive on Twitter in the fifth month did not survive. Breaking it down by category, 83% of the 694 ICOs that don't report capital and don't list on an exchange are inactive after 120 days. For the 420 ICOs that raise some capital but don't list, this figure falls to 52%, and for the 440 ICOs that list on an exchange, only 16% are inactive in the fifth month. We also show that in cryptocurrency markets, company announcements (as measured by Tweets) are good news, while no news is bad news. Daily market returns of tokens go up by about 0.3% for each Tweet that day and are also somewhat positively correlated with previous day's Twitter intensity. On the other hand, returns are negatively correlated with Twitter intensity from the prior month, suggesting some overreaction to news announced on Twitter, leading to reversals.

Finally, we calculate the average returns to investors from buying a token when it is listed and then holding it for different time horizons. Because of overall appreciation in cryptocurrencies, we adjust for asset class returns in three ways: by subtracting the return of Bitcoin (the most popular and actively traded cryptocurrency), by subtracting a value-weighted return to an index of all tokens, and by subtracting the return to a matched (by market capitalization) cryptocurrency that has traded for at least one year.

In contrast to IPOs, crypto-tokens continue to generate abnormal positive average returns after the ICO. The first day's average abnormal returns range from 14% to 16%, 30-day average abnormal returns range from 41% to 67%, and 180-day average abnormal returns range from 150% to 430%. Thus, even after the positive and significant returns from ICO to listing, token prices continue to drift higher (relative to the asset class) in the first six months after the listing date.

The literature on blockchains, cryptocurrencies, and ICOs has grown significantly in just the last year. A number of recent theory papers have focused on the value of this hybrid asset, simultaneously an instrument used to raise capital, and a currency that will be used as a medium of exchange once the platform is built. Using a dynamic model, Cong, Li, and Wang (2018) derive the token price as a function of user population and preferences, and platform utility, and show that tokens lead to accelerated platform adoption as investors also become customers, and vice versa. Li and Mann (2018) and Sockin and Xiong (2018) focus on the advantages of tokens in coordinating platform adoption, and aggregating information about a platform's fundamental value. ICOs could also be beneficial for the entrepreneur, in extracting information about how much value consumers derive from the platform (Catalini and Gans, 2018), and in transferring idiosyncratic venture risk without losing control rights (Chod and Lyandres, 2018).

The empirical literature on cryptocurrencies has also developed, with a number of papers looking at the return structure of cryptocurrency markets (Stoffels, 2017; Krueckeberg and Scholz, 2018) and transaction costs (Easley, O'Hara, and Basu, 2017).

Several studies have focused on ICOs, with a focus on determinants of ICO success (Fisch, 2018; Amsden and Schweizer, 2018). Momtaz (2018) examines token returns on the first day of trading, showing them to be on average positive and significant, as well as determinants of first day returns and ICO success. Our paper is the only one to date, that we are aware of, that looks at returns to investing in an ICO and exchange-traded tokens.

More broadly, our paper complements the growing literature on the value of blockchains as a development of *Fintech*. Yermack (2017) discusses the possible uses of blockchain technology for corporate governance, while Cong and He (2018) analyze how they can improve contracting through the use of “smart” self-enforcing contracts. Catalini and Gans (2017) provide an overview of the disruptive potential of blockchain technology. Saleh (2017) studies the proof-of-stake mechanism as an alternative to the electricity-reliant proof-of-work currently used by most public blockchains. Harvey (2014) provides a primer on cryptocurrencies. Outside of cryptocurrencies, our paper is also related to the recent literature on crowdfunding design including Ellman and Hurkens (2015), Chang (2016), Chemla and Tinn (2016), and Strausz (2017).

Our paper makes three main contributions. First, we merge multiple online databases with hand-collected data to give us an unbiased and comprehensive summary of ICO fundraising metrics, prices and returns, and social media activity. Therefore, our findings are less subject to critiques of survivorship, recency, and selection bias. Second, we show that ICO investors are receiving extraordinary compensation for providing capital to an unproven firm and product through unregulated means. On the one hand, this is

consistent with the tight relationship between risk and return that we would expect in markets with rational agents. On the other hand, it flies in the face of many regulators and governments who view token sales as scams that take advantage of unsophisticated investors, and thus must be restricted to sophisticated investors or even banned. Finally, we provide calculations for ICO returns which are directly comparable (and contrastable) with those on IPOs, which can be helpful in better understanding the causes of the well-documented IPO anomalies.

Our paper proceeds as follows. Section 2 explains the institutional details of the ICO market. Section 3 introduces the data and methodology for this paper. Section 4 presents results on ICO returns. Section 5 provides results for Twitter activity. Section 6 focuses on ICO returns after listing. Section 7 concludes.

2. Background on initial coin offerings and cryptocurrencies

Cryptocurrencies are digital currencies with all transactions kept on a public decentralized distributed ledger called the *blockchain*. They are typically created by a process called *mining*, in which computers run software that solves complicated math problems to verify transactions and are rewarded with new units of currency. For instance, on January 3, 2009, the first 50 Bitcoins were “mined” for the first (or genesis) block of the Bitcoin blockchain. From 2009 through 2013, a number of alternative cryptocurrencies, each with their own blockchain, were introduced with some unique features but generally the same functional form as Bitcoin.

In 2013, J.R. Sweezy introduced MasterCoin (later rebranded as OMNI), a protocol built on top of Bitcoin. Unlike earlier cryptocurrencies, units of MasterCoin were created by a fundraiser in the month of August 2013, during which, interested parties could send Bitcoins to an account and receive MasterCoins back at a pre-established exchange rate. The fundraiser collected approximately \$500,000 worth of Bitcoin which was to be used for development and payment of bounties for important tasks. Other cryptocurrencies started using the fundraiser model for the creation of cryptocurrencies, with Ethereum raising over \$15 million in August 2014. In 2015, Ethereum's introduction of a standard for implementing tokens (ERC20) further streamlined the ICO process. In 2015, there were 9 such offerings, 74 in 2016, and more than 1000 ICOs in 2017.

ICOs typically begin when the organization issuing the cryptocurrency publishes a *White Paper*, a prospectus which details the project's goal, roadmap, team, and schedule for the offering. Interested buyers can then register for the offering and confirm their identities to avoid violations of Know Your Customer (KYC) and Anti-Money Laundering (AML) regulations, and local securities laws. ICOs are nowadays typically conducted in multiple stages, which can altogether last several months, with earlier investors getting better terms. Some stages can be restricted to preferred users, angel investors, venture capitalists, and/or accredited investors.

Because cryptocurrency accounts are anonymous¹, transactions are irrevocable, and ICOs are unregulated, there is a high incidence of scams and theft. One common scheme involves hacking the website or social media accounts of a legitimate ICO and changing instructions so buyers send money to the hackers rather than the token sellers. This happened to CoinDash/Blox in July 2017, resulting in \$7 million stolen in just half an hour. The token seller can also be hacked after the ICO, as happened to The DAO in June 2016, resulting in the theft of approximately \$60 million in cryptocurrency.

Sometimes, the organizers of the ICO are scammers themselves. Recently, a Vietnamese pyramid scheme used an ICO to raise \$650 million and then disappeared with the money. The co-founders of Centra, which had raised \$32 million in an ICO, were arrested in April 2018 in Florida for fabricating information about deals that their company was making and listing fictitious people on their website. More common than these obvious cases of criminality are soft scams, in which the entrepreneurs pretend to be using ICO proceeds for project development but instead slowly abandon the project and keep most of the ICO proceeds for themselves.

In response to these high-profile illegal activities, governments and regulators have tightened the vise on ICOs. China and South Korea banned them entirely in late 2017. In the United States, ICOs are currently in a state of legal limbo, with Securities and Exchange Commission (SEC) officials testifying that they believe (almost) all tokens are

¹ Technically, cryptocurrencies are “pseudonymous”, which means the addresses are public and uniquely identifiable and traceable, but the accountholder’s identity is anonymous.

securities and therefore need to be registered, but bringing enforcement actions only against schemes that are clearly fraudulent. The SEC even set up a website for a nonexistent coin called the Howey Coin (with a fake White Paper and fake Twitter accounts for celebrities who supposedly endorsed this coin), in order to educate investors about the risks associated with ICOs. In March 2018, the Praetorian Group became the first ICO to register with the SEC as a security offering. Other countries, including ICO safe havens like Switzerland and Singapore are also revising their regulations, making the regulatory climate for ICOs all around the world highly uncertain.

3. Data and methodology

We merge data for three different components of our analysis: ICO characteristics, daily market data on cryptocurrencies, and social media announcements on Twitter. A large number of ICO aggregator websites (over 15 as of the writing of this paper) have been launched due to increased interest by potential investors and the general public. We collect data from five aggregator websites that, when combined, give us largest number of ICOs, the most available characteristics, and the highest accuracy: *icodata.io* (1731 observations), *icobench.com* (3073 observations), *icorating.com* (2312 observations), *icodrops.com* (457 observations), and *ico-check.com* (533 observations). In addition to gathering ICO characteristics, we also collect four identifiers to help match the datasets together: token/platform name, ticker symbol, website URL, and twitter handle. When we merge the datasets together, we are left with 4003 ICOs, including 2390 ICOs which

were completed before April 30, 2018, and 1613 ICOs which are ongoing as of that date or planned for the future. We drop the latter from our sample and focus on the set of ICOs that have been completed.

One potential problem with our data collection process is that since the aggregator websites have not been around since the first ICO in 2013, they may be backfilling based on survivorship (or ICO success), creating a survivorship bias in our dataset of ICOs. Survivors are likely to be better performers which would bias upward the returns that we estimate. In order to mitigate this concern, we supplement the website aggregators data with hand-collected data on ICOs that were completed before December 31, 2016. This data is manually collected using Google Searches with time ranges (using “ICO” and common synonyms) and announcements on *bitcointalk.org*. We find that 50% of the 96 hand-collected ICOs (from 2013 through 2016) are not found in any of the five data aggregators. However, the aggregators do much better with more recent ICOs, with 83% of ICOs from the 4th quarter of 2016 being covered by at least one of the data aggregators.

We use *coinmarketcap.com* (CMC) for cryptocurrency market data. CMC is widely considered the best available data source for cryptocurrency volume and prices. It collects information daily from the most widely-used exchanges² and posts opening, closing, high, and low prices, as well as volume and market capitalization (in dollars) for most included

² On April 30, 2018, CMC provides Bitcoin prices from 88 different exchanges, and Ether prices from 91 different exchanges.

cryptocurrencies.³ It is a survivorship-bias-free dataset that currently includes data from approximately 1600 active and 1100 defunct cryptocurrencies. CMC does require organizations to submit a form in order to list their currencies so there is occasionally a small lag between the exchange listing date and the date when prices start appearing on the website.

Finally, for each Twitter account, we collect information from *Twitter.com* on the time/date and text of tweets, the time/date when the account was opened, and the number of followers (as of the collection date: May 8, 2018). Twitter only allows collection of the last 3200 tweets from each user, but only 2.6% of the twitter accounts in our sample have more than that number of tweets as of the collection date. We code the number of Tweets in a day as missing instead of zero, if there are 3200 more recent Tweets available.

3.1 Initial coin offerings (ICO) data

Each of the five web aggregators have their own advantages and disadvantages. In order to figure out which source to use to resolve disagreement in variable values, we go to the source data (company ICO announcements) for a sample of such discrepancies. Based on this sampling, we determined the following: *Icobench* has the most observations but often provides inaccurate values for certain variables (also true for *Icorating*). *Icodata* has fewer observations and small number of characteristics but the data is highly accurate.

³ Market capitalization does not necessarily capture the value of floated tokens as many are kept in reserve by the issuers for bounties, employee compensation, etc.

Icodrops has the highest-quality data on ICO prices. *Ico-check* provides the full schedule of ICO prices (including all bonuses for investing early). Based on this analysis, our hierarchy of data sources is (1) *ICodata*, (2) *ICOdrops*, (3) *ICObench*, (4) *ICOrating*, (5) *Icocheck*, except (1) and (2) are switched for price data. For any variable of interest, we use this hierarchy to determine which source to use across all datasets that provide that variable.

We collect the following ICO variables: *ICO start date*, *ICO end date*, *Capital raised*, *ICO price*, *Hardcap*, *Tokens % sold*, *Country of registration*, and *Pre-ICO* (dummy variable). We use the ICO start and end dates to generate *Length of ICO*. We define *Country Rule of law* to be the home country’s World Bank Rule of Law rating⁴ (as of 2016, the most recent year available), and *Country GDP per capita* to be the home country’s 2016 World Bank per-capita GDP (in US\$). *Twitter Age (months)* is the number of months between the date when the Twitter account was opened and the ICO was started (or one plus months to *ICO end date* if *ICO start date* is missing). *Pre-ICO Twitter intensity* is the average daily number of tweets in the six months before the ICO started, and *ICO Twitter intensity* is the average daily number of tweets during the ICO.

We report summary statistics for all ICO characteristics in Table 1. For each variable, Panel A shows the number of non-missing observations, along with the cross-

⁴ According to the World Bank, this rating “reflects perceptions of the extent to which agents have confidence in and abide by the rules of society, and in particular the quality of contract enforcement, property rights, the police, and the courts, as well as the likelihood of crime and violence.”

sectional mean, standard deviation, 10th, 50th, and 90th percentile values. In Panel B, we report mean values (and median values in brackets) for various subsamples. In Columns (1) through (3), we separately report summary statistics for those offerings that didn't report raising capital, those offerings that raised capital, which we call "successful" ICOs (*Successful* = 1), and for those offering that were listed on CMC, which we call "listed" ICOs (*Listed* = 1). Because the ICO market is relatively new and has been changing as it matures, we also separately report summary statistics in Columns (4) through (6) for subsamples of ICOs that were completed before July 2017, from July to December 2017, and in 2018.

Panel A of Table 1 shows that 48% of the 2390 ICOs had non-zero and non-missing values for capital raised. The remaining 52% of ICOs likely fall into one of four categories: (1) They raise capital and use it to continue with the project, but don't announce the amount raised, (2) They raise capital but don't reach their "softcap", the minimum required to go through with the project, so they refund the funds to investors, (3) They are scams used to steal the funds, (4) They are announced, along with start and end dates, but never actually take place. We can't know for sure which category each ICO falls into, but in Section 5, we will use their Twitter activity after the ICO to try to determine what fraction of these projects actually held an ICO and what fraction continued to be active after the completion of the ICO. Overall, only 26% of ICOs have listed their tokens on an exchange, with the important caveat that many ICOs have been completed recently and might be listed in the future.

The average (successful) ICO raised \$11.5 million, but the distribution is positively skewed due to a small number of “mega-ICOs”, so the median value raised is only \$3.8 million. Panel B unsurprisingly shows that listed ICOs raise about three times as much, on average, compared to non-listed ICOs. Furthermore, the average capital raised by successful ICOs has gone up over time from \$7.9 million (before July 2017) to \$11 million (in the second half of 2017) to \$14 million (in 2018). Most entrepreneurs choose a fairly low ICO price for their tokens, with a median value of \$0.30 making most of them akin to penny stocks, although several large outliers lead to a cross-sectional average ICO price of \$11.20. Because the minimum tick size is tiny on most cryptocurrency exchanges, the nominal price of a token is unlikely to have major liquidity or microstructure effects on trading, but it could be important in attracting behavioral investors who are attracted to low nominal prices (Birru and Wang, 2016). Panel B suggests this might be the case, with a slightly lower median ICO price for successful than unsuccessful ICOs, and a much lower median ICO price for listed ICOs.

The hardcap, the maximum amount of funds that the ICO is allowed to raise, is approximately \$43 million for the average ICO (median value is \$23 million), and the average percent of all tokens that are sold during the ICO is 60%. Panel B shows that these variables are generally not correlated with ICO “success” and have not changed much over time, except for a small increase in token percentage sold since the earliest ICOs. 40% of projects hold a pre-ICO, in which tokens are sold to preferred or accredited investors at substantial discounts, with the funds raised from the ICO used to pay for the

costs associated with the ICO. Pre-ICOs do not seem to be correlated with success, but they have become more popular over time, with an average incidence of 1% for ICOs completed before July 1, 2017, 29% for the second half of 2017 ICOs, and 57% for 2018 ICOs. The average ICO lasts 37 days (with a median length of 31 days) although this figure has recently been rising with an average of 41 days for 2018 ICOs.

Several of the ICO aggregator websites report a home country for each ICO. Figure 1 shows the distribution of ICOs by country, with darker colors indicating a greater number of ICOs. Although we collect and report statistics related to this variable (as reported in our data sources), we note that since these are token sales for decentralized digital platforms, they don't really have a single home country. Although Satoshi Nakamoto is supposed to be an alias for someone from Japan, no one would call Bitcoin a Japanese coin. In fact, the team of entrepreneurs and employees are often from many different countries, and the country of registration or incorporation is usually chosen for legal and tax reasons, which is why noted havens like British Overseas Territories, Singapore, Switzerland, Cyprus, and the Baltic States, especially Estonia, are overrepresented in Figure 1. The largest English-speaking countries and Russia are also popular locations for ICOs. Table 1 indicates that ICOs are mostly located in countries with above-average World Bank Rule of Law rankings and high standards of living, two metrics which are highly correlated with each other. These two measures are also positively associated with ICO success, with listed ICOs originating from countries that

are 0.2 points higher in their Rule of Law rating and have about \$4,000 more in GDP per capita, relative to the entire sample.

The last set of variables is related to activity characteristics based on Twitter data. The average length of time from the date that the platform's Twitter account is activated to the ICO start date is about 8 months, but the median value of this variable is only 3 months, so a large number open social media accounts just in time for their token sales. Accounts with a longer record of activity are slightly more likely to be successful, with listed ICOs have an average *Twitter age* of 9.4 months (median of 4 months). Crypto-companies are not extremely active in tweeting, with the average firm putting out 2.1 tweets per day in the six months before the start of the ICO (median firm tweets 0.7 times per day). This only rises slightly to 2.7 tweets per day (once per day for the median) during the ICO. Panel B provides some evidence that stronger activity before the ICO, and especially during the ICO, is correlated with success. However, increased tweeting during the ICO could be a result of ICO success rather than its cause, as entrepreneurs are more likely to share good news about strong token sales. In addition to Twitter, ICO platforms typically use multiple social media sites including Bitcointalk, Facebook, Telegram, Slack, Medium, Reddit, et al., so we are only capturing a portion, although hopefully a representative portion, of their social media and marketing footprint.

Figure 2 shows the distribution of ICOs over time (from February 2016 through April 2018) and how much each firm raised in capital. The ICO wave really started to build in March 2017, with significant growth in the number per month and how much

capital they were raising. Listing of tokens is depicted by the shape and color of each dot. Most of the ICOs that have not listed (shown with a red x) either didn't raise a lot of capital or happened only recently so they are likely to list in the future. Blue diamonds indicate tokens that listed but only after 60 days had passed from the end of the ICO. Since we don't know which tokens will list in the future, we do not have the full set of such "lagged" listed tokens. Therefore, to ensure that our sample is not biased, we generally exclude them, along with all ICOs that happened in the last 60 days, from our performance tests.

3.2 Market prices of cryptocurrencies data

We merge the ICO data from the previous section with daily market data on price, volume, and market capitalization from CMC, using the same four identifiers: token/platform name, ticker symbol, website URL, and twitter handle. We find matches for 627 of the 2390 ICO in the CMC database. Of those 627, 18 were listed before the start of the ICO (either due to a pre-ICO, an earlier offering, or a conversion) which we drop from our sample, leaving us with 609 listed ICOs.

We report summary statistics for these 609 tokens in Table 2. The average time between the end of the ICO and the first day of trading is 31 days, with a median time of 16 days. 44 were listed prior to the end of the ICO. The average ICO price for this sample is \$22.70 (median is \$0.14), while the average opening price is \$25.90 (median is \$0.19) suggesting a positive drift in price between the end of the ICO and listing dates.

Calculating a rate of return for the sample where ICO price and opening (and closing) market prices are available yields an average rate of return of 246% (median rate of return of 21%) to first open and 273% (median rate of return of 29%) to first close. It's important to highlight that this is an average return for "successful ICOs" that go on to be listed and therefore is not something that can be predicted at the time of the ICO and earned by an investor.

First day's market capitalization averages \$54.7 million, which is significantly higher than the \$17.2 million average capital raised. Some of this discrepancy is due to a fraction of tokens held by insiders and not sold in the offering, but most comes from the appreciation in price. Finally, we calculate the average daily volume in the first month of trading to see how liquid these tokens are for investors who bought them in the ICO and want to potentially liquidate them at a profit. The cross-sectional average value for first month volume is \$2.5 million (per day), but the distribution is highly skewed with a cross-sectional median value of \$140,000. A fairly large proportion of currencies are illiquid, with the 10th percentile having an average daily volume of \$5000 in the first month. In the next section of the paper, we investigate, in greater detail, the performance and feasibility of a trading strategy of buying cryptotokens at the ICO price and selling them when they are listed on an exchange.

4. ICO underpricing and returns to ICO investors

Our first set of tests examine the returns to token buyers from investing through ICOs. Our main variable of interest, *ICO->first open ROR*, is the rate of return from the price that the investor paid during the offering until the first available (opening) price on CMC. One concern with using this measure is that it is available only for the subset of ICOs which go on to list their tokens on an exchange, approximately 25% of the sample. Investors that buy tokens from the other 75% of ICOs own an illiquid and useless asset, at least until such time as the platform is built and the token can be exchanged for services on the platform. However, there are two points that mitigate this concern: (1) many of these “unlisted” ICOs don’t raise much or any capital or raise so little that they refund it, and (2) almost all the money actually invested in ICOs is invested in tokens that go on to list, so the ones with data are the ones that actually matter to the typical investor.

A second concern is that we don’t now know all the past ICOs that list because they might still list in the future. While the median time between the end of the ICO and listing is only 16 days, a few ICOs don’t list their tokens for many months or even years. In order to avoid any selection bias in our sample based on time-to-listing, we restrict our analysis to ICOs that were completed by March 1, 2018, and that listed in less than 60 days, since we currently know all the ICOs that satisfy these criteria.

We start with a graphical representation of ICO underpricing, scatter plotting returns against ICO end date on Figure 3 for all the ICOs in our (restricted) sample. Table 2 indicates that the returns variable has a high positive skewness due to several large outliers, so to make it easier to graph, we first perform a log-transformation and

then plot the natural logarithm of one plus *ICO*->*first open ROR* against date. All price calculations for this figure are denominated in Bitcoin which should remove any changes in value arising from the volatile fluctuations in the exchange rate between the cryptocurrency asset class and fiat currencies.

Figure 3 illustrates that most tokens were sold below their market price, but also, that many tokens were overpriced, and declined in value. The red-dashed line, which is the best fit line, is above the x-axis for the entire sample period, indicating that the average (log) return is positive, but it has a negative slope, suggesting that underpricing of tokens has declined over time (i.e., returns to ICO investors have been declining). We investigate this decline over time in the context of a multivariable regression, and present the results in Table 5, later in this section.

Table 3 shows the average returns to investing in an ICO. The first two rows are equal-weighted averages of all observations, with the prices denominated in dollars for the first row set of results and in Bitcoin for the second row set of results. Equal-weighted returns are the returns to a naïve investor, who puts the same amount of money in each ICO, not accounting for their merits. The last two rows are value-weighted returns, weighted by the amount of capital raised by each ICO. These are the returns to a sophisticated investor who, like the aggregate set of investors, realizes ex-ante that some projects have relatively less merit and are under-weighted in her portfolio, while others, with more potential, are over-weighted. This is also the return to a randomly-chosen dollar invested in ICOs, or, if each investor's ICO purchases were identical (say \$1000), the

return to a randomly chosen investor. In our opinion, the value-weighted returns are more informative and interesting in understanding the ICO market, but we report both sets of results for comparison in Table 3.

We start by calculating returns to investors in 416 ICOs that went on to list, in less than 60 days, and report the results in Column (1) of Table 3. The average of equal-weighted returns to investing in listed ICOs is a statistically significant 179% and 167% (in Bitcoin), with a very similar 173% and 162% (in Bitcoin) value-weighted average. From the sellers' point of view, crypto-companies are, on average, issuing tokens for less than half of their true market value, leaving significant money on the table.

For Columns (2) and (3), we also include (in addition to the 416 listed ICOs) another 471 ICOs that reported raising capital but did not list within 60 days. Since there are no available market values for these tokens in the aftermath of the ICO, we impute returns under two different scenarios. In Column (2), the average imputed return to unlisted tokens is -50%. Unlisted tokens investments are not a total loss if the raised capital is refunded due to inadequate funds, if there is an over-the-counter market for them, or if the tokens are listed on an exchange which is not included in CMC. With imputed returns of -50% to unlisted ICOs, average ICO returns are unsurprisingly lower than in Column (1), 57% and 52% (in Bitcoin) for equal-weighted averages and 105% and 98% (in Bitcoin) for value-weighted averages, but still positive and statistically significant. In Column (3), we look at worst-case scenario, imputing -100% to all ICOs that raised capital but did not list within 60 days. Under this scenario, the equal-weighted average

returns are 31% and 26% (in Bitcoin) and are no longer significant at the 5% level, but the value-weighted returns remain larger in magnitude and significant at 90% and 82% (in Bitcoin). The main reason for the divergence between equal and value weighting is the high correlation between raised capital and listing probability (which was illustrated in Figure 2) so most of the ICOs that don't list raised little capital and value-weighting minimizes the impact of their -100% return.

For Columns (4) and (5), we include an additional 732 ICOs that neither reported raising capital nor were listed within 60 days. Again, we calculate and report average equal-weighted investor returns after imputing -50% (in Column (4)) and -100% (in Column (5)) returns to unlisted ICOs. Since these ICOs raised little or no capital, they do not change the value-weighted returns we calculated and displayed in the last two rows of Columns (2) and (3). When including these ICOs, equal-weighted returns are reduced to 9% (6% in Bitcoin) with a -50% imputed return, and -28% (-31% in Bitcoin) with an imputed return of -100%. These are the returns to a naïve investor who invests across all ICOs, even those that didn't report raising capital, and they provide a lower bound to naïve investor returns. However, they are not at all a realistic estimate of returns, even for naïve investors, because many of the ICOs that don't report raising capital (and many of those that report raising capital but do not list) either refunded the capital they raised because of inadequate funds or they planned an ICO but never actually began collecting funds.

While a lower bound of 82%-90% for value-weighted returns is surprisingly high, there are two potential caveats that we need to worry about. First, we calculate the return by assuming ICO investors can sell their tokens at the first day's opening market price, even though the depth at that price might be negligible. Second, the overall cryptocurrency asset class has performed well since the first ICO and we do not know whether this explains the high ICO returns (even when denominated in Bitcoin). We examine each of these issues and present results in Table 4.

For Panel A of Table 4, we calculate returns to ICO investors using several different scenarios for the selling price. Column (1) reproduces the results from the first column of Table 3 (i.e., setting the sales price to the first day's opening price) to allow comparisons with the remaining columns. For Column (2), we assume the investor sells her tokens at the same prices as the first \$1 million in volume. For each token-day, we calculate that day's average price (averaging the open and close prices) and then calculate a volume-weighted average of the daily average price across the first N days where N is set to be the number of days until a cumulative \$1 million in trading has occurred. If N is greater than 60, meaning there was less than \$1 million in volume in the first 60 days, we assume the remaining fraction of tokens were sold at 50% of the closing price on the 60th day. The 50% discount represents the imputed liquidity penalty to the investor for selling all her remaining tokens on that day. For example, if only \$400,000 in tokens were traded over the first 60 days, we assume 40% of tokens were sold at the volume-weighted average price across those 60 days and the remaining 60% of tokens were sold at half of the 60th

day closing price. We repeat this exercise in Columns (3) and (4) of Table 4 using volume limits of \$10 million and \$50 million, essentially requiring the ICO investor to spread out her token sales over longer periods of time to ensure the existence of enough depth in the market for her sales.

The results across the last three columns of Panel A show that returns are not significantly affected by forcing slower liquidation of tokens by the ICO investor. The main reason is that tokens actually increase in price after their listing date (as we will show in greater detail in Section 6), so not immediately liquidating her position actually helps the investor get better returns. This leads to higher returns in Column (2) compared to Column (1). As we increase the volume requirements in Columns (3) and (4), the 60-day cutoff becomes more important and more tokens are sold by the investor at a 50% discount. This leads to declining average returns with higher volume limits. Still, even for investors who require \$50 million in volume over which to sell their tokens and also face a 60-day deadline at which time they have to sell their remaining tokens at a large discount, the value-weighted average return is still a significant 142% (116% in Bitcoin), indicating that the depth for token sales is not a major concern for calculating true return.

In Panel B of Table 4, we analyze the effect of the cryptocurrency asset class on ICO returns by dividing the sample into five groups, or quintiles, based on the average cumulative return of Bitcoin from the ICO end date to the day the token starts trading. This helps us to predict what ICO returns might look like in a different sample period when cryptocurrencies are not experiencing significant price appreciation. Column (1)

shows ICO returns for the lowest quintile, an average Bitcoin return of -22% and always less than -13%. Even for ICOs that occurred during this worst climate for cryptocurrencies, the equal-weighted average return, in dollars, was 124%, and the average value-weighted return, in dollars, was 62%. Not surprisingly, in better cryptocurrency climates in the remaining four columns, returns were even better. Still, across all five regimes, value-weighted average returns were no worse than 62%, suggesting that the positive returns to ICOs are not an artifact of the particular cryptocurrency wave that makes up our sample.

Next, for Table 5, we examine the cross-sectional determinants of ICO pricing, and try to understand why entrepreneurs sell their tokens below market prices during the offering. We are faced with a significant selection problem since we only know the market prices of ICOs whose tokens go on to be listed, and in order to go on to list your token, you need to raise money which is easier to do if you sell your token at a bigger discount. In order to handle this problem, we use the Heckman (1976, 1979) technique for correction of sample selection but calculate the efficient full information maximum likelihood estimator instead of the two-step estimator. The selection variable is the natural logarithm of the capital raised in the ICO since this is the primary determinant of whether an ICO is successful and the token goes on to be listed. Our dependent variable across all regressions is the natural logarithm of the ratio of *First opening price* to *ICO price*, which is also the natural logarithm of one plus the return.

In Column (1) of Table 5, we run an OLS regression of ICO returns on four variables: the date of the ICO, the natural log of the ICO price, an indicator for whether there was a presale, and the age of the Twitter account of the platform at the start of the ICO. In Columns (2) and (3), we instead run this specification using the Heckman technique, and in Column (3), calculate prices in Bitcoin units instead of dollars. Across all three specifications, we find a negative estimate on the *Date* variable, which is statistically significant in the Heckman tests. The estimate is -0.0007 (per day), so multiplying this figure by 365 yields approximately a 25-percentage point estimated reduction in ICO returns for an ICO that is one year more recent.

We also find a negative and significant coefficient on the nominal ICO price in all three specifications. This is consistent with ICO investors suffering from nominal price illusion which causes higher market demand (and thus higher returns relative to the ICO price) for tokens with low nominal prices and low demand (and thus lower returns) for high-priced tokens. Finally, there is a negative and weakly-significant (at the 10% level) coefficient on the pre-ICO indicator, suggesting some learning about the market value of their token during the pre-ICO, helping entrepreneurs properly price their tokens at the ICO. Interestingly, in contrast with IPOs, there is no indication that more established companies, with a longer track record as proxied by the length of time since their Twitter account was activated, suffer from less pricing. This could be because there is very little information available in general on these firms, or perhaps their Twitter account age is a weak proxy for their age or the information available about them.

In Columns (4) through (6) of Table 5, we add several additional explanatory variables for which we have worse data coverage or which might be endogenously determined with price. The *rule of law* rating does not significantly affect ICO returns, nor does their Twitter activity before or during the ICO. The length of the ICO as well as the capital raised are functions of the ICO's success (they end once they hit their hardcap so successful ICOs are shorter) which is a function of the token price discount relative to the market price, so it's hardly surprising to find a negative coefficient on length and a significant positive coefficient on capital raised. We study the relationship between Twitter activity and followers, and market prices in the next section of the paper.

5. Cryptocurrencies and information distribution through Twitter

Our next task is to understand how a platform's social media activity (on Twitter) ties in to the value of its token. Twitter activity also allows us to quantify the survival rate of crypto-companies after they finish raising capital, in order to get a better understanding of how many of the ICOs were failures or scams. In this section, we are using Twitter followers as a proxy for company users and Tweets as a proxy for company announcements and level of activity. First, we examine the relationship between number of Twitter followers and market capitalization to see if the theories about positive network externalities from a critical mass of users is actually borne out by the data. We only have data on the current number of followers, as of May 8, 2018, so we relate this to the market capitalization on the same date.

We have 1150 cryptocurrencies, including 632 that didn't have an ICO, for which we have both variables available and we plot these assets (after log transformations of both x and y variables) in Figure 4. Unsurprisingly, the estimated coefficient on the best-fit line in Figure 4 is positive with more users equating to a large market capitalization. More interestingly, the value of the coefficient is 1.2 (with a standard error of 0.033), indicating that it is also statistically greater than one. Thus, for each 1% increase in users, the market capitalization increases by 1.2%, which is consistent with the increasing returns to user adoption in peer-to-peer platforms where these tokens are to be used for economic activity.

Next, we use Twitter activity to estimate short-term survival rates for the three types of ICOs that we encounter: those that don't report raising any money and that don't list on an exchange, those that report raising capital but don't list, and those that list. We start by providing summary statistics on each of the categories in the first three columns of Table 6. Out of 694 ICOs in the first category, 51.4% never even opened a Twitter account and an additional 7.8% opened a Twitter accounts but had no activity during the ICO. This supports our hypothesis that most of the so-called ICOs in this class never made it past the earliest planning stage and should not be considered as losses when calculating average returns to ICO investors. Out of 420 ICOs in the category of those that raised capital but didn't list, 21.2% never opened a Twitter account and another 5.2% opened a Twitter account but had no activity during the ICO. These figures are much smaller than for the first category so most of these ICOs were actually legitimate

and should be included in calculating average returns, although many could have refunded their proceeds (due to inadequate funds) so they would not have created a loss for investors. Finally, only 2.7% of the 440 ICOs that listed never opened a Twitter account and another 6.8% weren't active during the ICO. Such low figures are consistent with these ICOs being successful enough for the tokens to be listed on an exchange.

In Columns (4) through (9) of Table 6, we look at Twitter intensity (and the number of inactive users) during the ICO and in 30-day increments for the 150 days after the end of the ICO. We restrict our analysis for these specifications to tokens that have a Twitter account and were active during the ICO. For ICOs that failed to raise capital but were active during the ICO activity, their Twitter activity decays very quickly from 1.48 Tweets per day during the ICO to 0.28 Tweets per day in the fifth month, and inactivity rises to 59.2% in the fifth month. Inclusion of the ICOs with no Twitter accounts or no ICO Twitter activity, implies an overall survival rate for this category after four months 17%.

Twitter intensity is higher for those ICOs that successfully raised capital and although it declines from the ICO period to the post-ICO periods, it doesn't decay with time, stabilizing at around 0.8 Tweets per day. Total inactivity rises to 35.2% in the second month, so the overall survival rate for this category in the fifth month is 48%. For listed ICOs, Twitter intensity is even higher than for the first two categories and also is stable over time, remaining at around 1.5 Tweets per day. Total inactivity rises to only 6.8% for an overall survival rate for this category in the fifth month of 83%. Since most

of the capital to crypto-companies goes to firms in the last category, our analysis indicates that the fraction of funds invested in firms that become inactive after the ICO (potential scams) is only 11%.

Finally, we investigate how daily Twitter activity is related to market returns by regressing daily returns (log-transformed to minimize effect of outliers) on daily Twitter intensity (number of Tweets) with leads and lags, and then report the results in Table 7. In Column (1), the main explanatory variables are same-day and prior-day Twitter intensity, and we also control for lagged market capitalization and two lags of the dependent variable, which is important in case there is short-term autocorrelation. Same-day Twitter intensity is positively associated with returns, with a coefficient of 0.2% (t-statistic of 7.56). Since firms are more likely to announce and amplify good news than bad news, it is not surprising that a Tweet is an indication of a good-news day, which explains the higher returns. The previous day's Twitter intensity is negatively associated with today's returns. This seems to indicate price overreaction to prior day's news and then a reversal.

In Column (2) of Table 7, we add additional lags for Twitter intensity over the last five days. As in Column (1), there is a positive coefficient on today's Twitter intensity and negative coefficients on most lags of Twitter intensity. However, the effect of previous day's Twitter intensity is no longer significant as countervailing effects from the past and present offset each other. In Column (3), we add two additional variables for the average Twitter intensity from 10 days to 6 days before today, and for 40 days to 11 days before

today, to control for the platform’s typical level of activity. Both new variables have negative and significant coefficients that are very similar to each other, an indication that Twitter activity continues to negatively predict returns as far ahead as 40 days. Even after controlling for these typical measures of activity, the two-day and three-day lags still have a negative coefficient while the one-day lag flips to a positive coefficient.

In Column (4) of Table 7, we add token fixed effects to absorb cross-sectional variation in Twitter intensity, while in Column (5), we replaced all Twitter intensity variables with log-transformations (natural logarithm of one plus the variable) to minimize the effect of outliers. In both specifications, the results are very similar to those in Column (3). Finally, in Column (6), we add five leads of Twitter intensity, and find positive coefficients that decline with time, likely an indication that some firms are announcing good news with a lag. In summary, we show that past high levels of social media activity are generally correlated with lower returns, an indication of overreaction to news and reversals. On the other hand, returns are positively associated with today’s activity (and that in the near future), likely due to a bias for good news in company announcements. Thus, for cryptocurrency returns, the motto should be: “no news is bad news.”

6. Returns after tokens are exchange-listed

In this section, we calculate average returns to investors who purchase tokens (that had been originally sold through an ICO) at the opening price on the day that each token is first listed, and hold each token for anywhere from one day to one year. Our main goal

is to understand whether the high returns to ICO investors are reversed after tokens are traded in the market, as happens to stocks after their initial public offering. Figure 5 shows the average cumulative returns to investing in tokens at their initial price, in comparison to a Token Index (in gray) and Bitcoin (in orange). In Table 8, for each of seven holdings periods, we report raw buy-and-hold returns and abnormal returns after three different adjustments: Bitcoin-adjusted returns, token index-adjusted returns, and match-adjusted returns.

The first row of Table 8 shows that the average first-day raw return (from open to close) is 16% (t-statistic of 9.17), and ranges from 14.1% to 16.1% after adjustments. In the second row, at a one-week horizon, raw return averages 33.2% (t-statistic of 4.65), and ranges from 29.3% to 30.7% after adjustments. Two-week average returns are not much different from one-week returns, but at the 30-day horizon, raw return averages 67.1% (t-statistic of 3.33), ranging between 41.4% and 57.7% after adjustments. Token average returns continue to increase with longer holding periods. As holding periods get longer, the number of tokens that started trading long enough ago to have reached that age by the time our sample period ends declines and the volatility of cryptocurrencies increases standard deviations, both acting to reduce statistical significance. For the 180-day holding period, adjusted returns range from 152.8% to 242.5%, but with t-statistics all less than two. Unlike ICO underpricing, the results for long-run performance (at least in the first year) run completely counter to what prior research has shown for initial public offerings. This could be due to the lack of a lockup period for most ICOs, which would

cause the opening price to already reflect the supply from insider sellers, while for IPOs, the supply would be released after the lockup leading to lower returns. Another explanation is that ICOs are much younger and riskier and thus need to provide a high expected rate of return to induce investor demand.

7. Conclusion

ICOs have the potential to change how startup companies raise money, providing more control to entrepreneurs, greater liquidity to investors, and additional investment opportunities to early adopters. In this paper, we document that tokens are sold in ICOs at a significant discount to their market price (and at a much greater discount than IPOs), generating at least an 82% average abnormal return for the representative (i.e. weighted by capital invested) ICO investor. We analyze the determinants of this underpricing, and find that it has declined over time, and is less of an issue when firms do a preliminary offering before the ICO or charge a higher nominal price. Unlike with IPOs, it does not seem to be related to the level of information asymmetry, as proxied by the age of the company at the ICO.

We calculate token returns after the asset is listed on an exchange, and find that prices continue to drift higher, generating an abnormal return of approximately 50% in just the first 30 days. We also show that there is a positive and convex relationship between (log) market cap and (log) number of Twitter users, that nearly all ICO capital is raised by crypto-companies that continue to be active (on Twitter) after 120 days, and

that daily Twitter intensity is associated with positive returns that day but negative returns in the future, suggesting overreaction and reversals.

Our paper shows that ICOs investors are compensated handsomely for investing in new unproven platforms through unregulated offerings. It suggests that scams, while plentiful in number, are not as important in terms of stolen capital because investors are shrewd enough to spot (and underfund) them. It also shows how ICOs are both similar to and different from IPOs. Regulatory uncertainty in the United States and around the world has recently slowed the explosive growth in ICOs, but our findings suggest that while regulators should continue to deter fraudulent activities, they need to be careful not to throw out the baby with the bathwater.

References

- Amsden, Ryan, Schweizer, Denis, 2018. Are blockchain crowdsales the new “gold rush”? Success determinants of initial coin offerings, Concordia University working paper.
- Beatty, Randolph P., Ritter, Jay R., 1986. Investment banking, reputation, and the underpricing of initial public offerings, *Journal of Financial Economics* 15, 213-232.
- Birru, Justin, Wang, Baolian, 2016. Nominal price illusion, *Journal of Financial Economics* 119, 578-598.
- Catalini, Christian, Gans, Joshua S., 2017. Some simple economics of the blockchain, MIT working paper.
- Catalini, Christian, Gans, Joshua S., 2018. Initial coin offerings and the value of crypto tokens, MIT working paper.
- Chang, Jen-Wen, 2016. The economics of crowdfunding, CSU Fullerton working paper.
- Chemla, Gilles, Tinn, Katrin, 2016. Learning through crowdfunding, Imperial College working paper.
- Chod, Jiri, Lyandres, Evgeny, 2018. A theory of ICOs: diversification, agency and information asymmetry, Boston College working paper.
- Cong, Lin William, Li, Ye, Wang, Neng, 2018. Tokenomics: dynamic adoption and valuation, University of Chicago working paper.
- Cong, Lin William, He, Zhiguo, 2018. Blockchain disruption and smart contracts, University of Chicago working paper.
- Easley, David, O’Hara, Maureen, Basu, Soumya, 2017. From mining to markets: the evolution of Bitcoin transaction fees, Cornell University working paper.
- Ellman, Matthew, Hurkens, Sjaak, 2015. Optimal crowdfunding design, Barcelona Graduate School working paper.
- Fisch, Christian, 2018. Initial coin offerings (ICOs) to finance new ventures: an exploratory study, Trier University working paper.

- Harvey, Campbell R., 2014. Cryptofinance, Duke University working paper.
- Heckman, James J. 1976. The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models, *Annals of Economic and Social Measurement* 5, 475-492.
- Heckman, James J., 1979. Sample selection bias as a specification error, *Econometrica* 47, 153-161.
- Krueckeberg, Sinan, Scholz, Peter, 2018. Cryptocurrencies as an asset class? Hamburg School of Business Administration working paper.
- Li, Jiasun, Mann, William, 2018. Initial coin offering and platform building. UCLA working paper.
- Momtaz, Paul, 2018, Initial coin offerings, UCLA working paper.
- Ritter, Jay R., 1991. The long-run performance of initial public offerings, *Journal of Finance* 46, 3-27.
- Saleh, Fahad, 2017. Blockchain without waste: proof-of-stake, NYU working paper.
- Sockin, Michael, Xiong, Wei, 2018. A model of cryptocurrencies. University of Texas working paper.
- Stoffels, Jaron, 2017. Asset pricing of cryptocurrencies and momentum based patterns, Erasmus School of Economics working paper.
- Strausz, Roland, 2017. A theory of crowdfunding: a mechanism design approach with demand uncertainty and moral hazard, *American Economic Review* 107, 1430-76.
- Yermack, David, 2017. Corporate governance and blockchains, *Review of Finance* 21, 7-31.

Figure 1: World map of initial coin offerings

Figure 1 colors countries based on the number of ICOs that were completed in each country prior to April 30, 2018. Dark red indicates more than 100 ICOs, crimson indicates 51 to 100 ICOs, lighter red indicates 26 to 50 ICOs, pink indicates 11 to 25 ICOs, and yellow indicates 6-10 ICOs. All other countries had five or fewer ICOs.

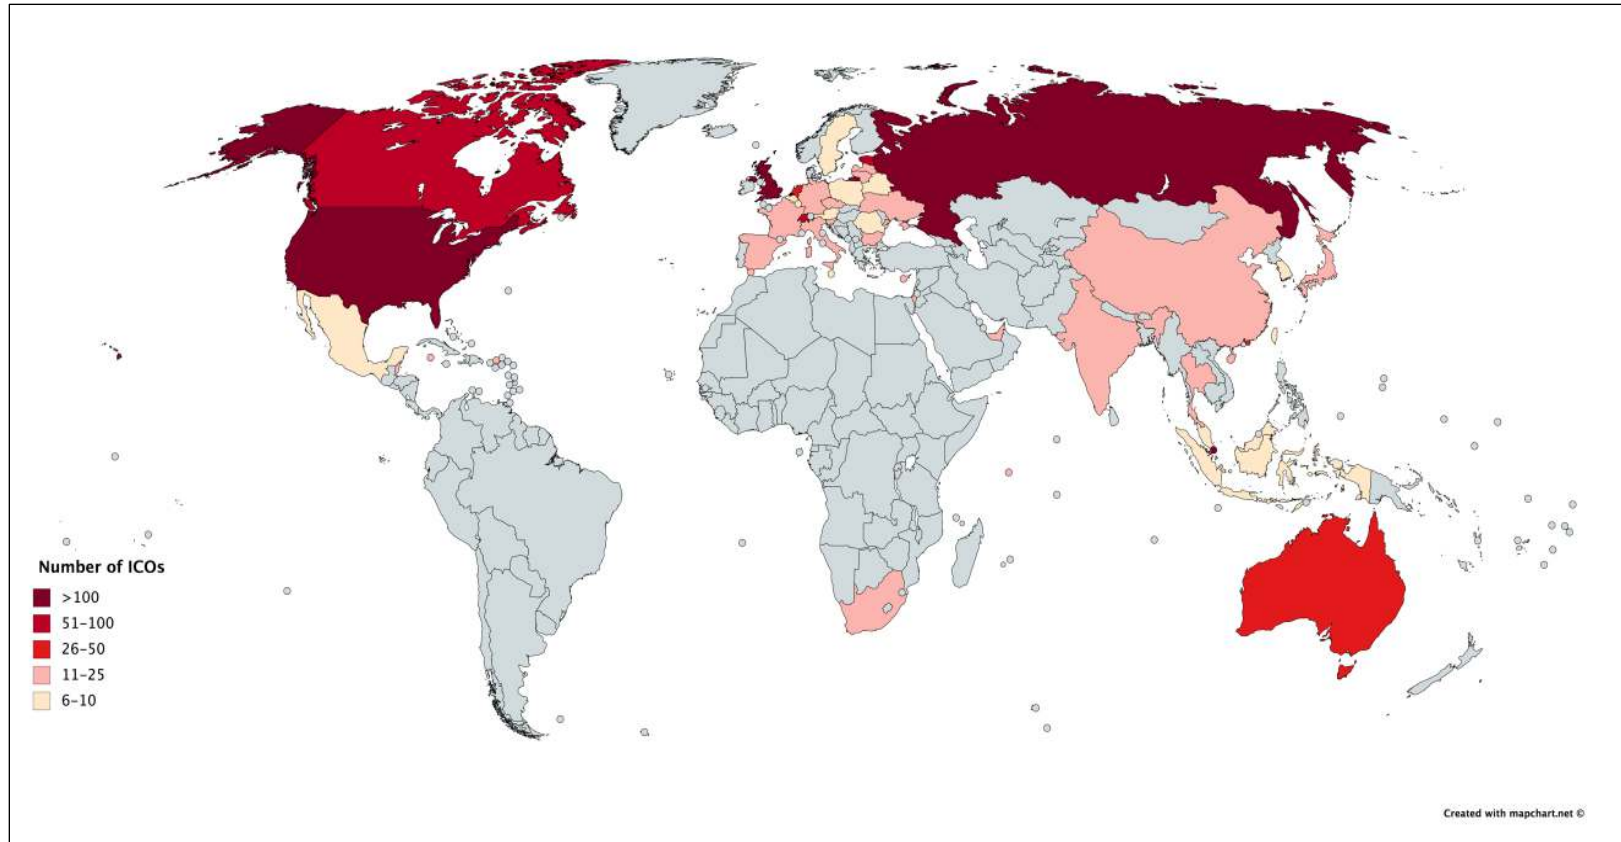


Figure 2: Initial coin offerings by date and capital raised

Figure 2 is a scatter plot where each point represents an ICO. The x-value is the date that the ICO ended and the y-value is how much the ICO raised. Red (x) indicates ICOs that have not listed, Green (circle) indicates listings within 60 days, and Blue (diamond) are listings after 60 days.




Figure 3: Initial coin offerings by date and investor returns

Figure 3 is a scatter plot where each point represents an ICO. The x-value is the date that the ICO ended and the y-value is the natural logarithm of the ratio of the first day's opening price of the token to its ICO price (prices are in Bitcoin). The red dotted line is the best fit line.




Figure 4: Plot of current market capitalization on current number of followers

Figure 4 is a scatter plot where each point represents a traded cryptocurrency. The x-value is the natural logarithm of the number of Twitter followers on May 6, 2018 and the y-value is the natural logarithm of the market cap (in \$1000s) on that date. The red dotted line is the best fit line.

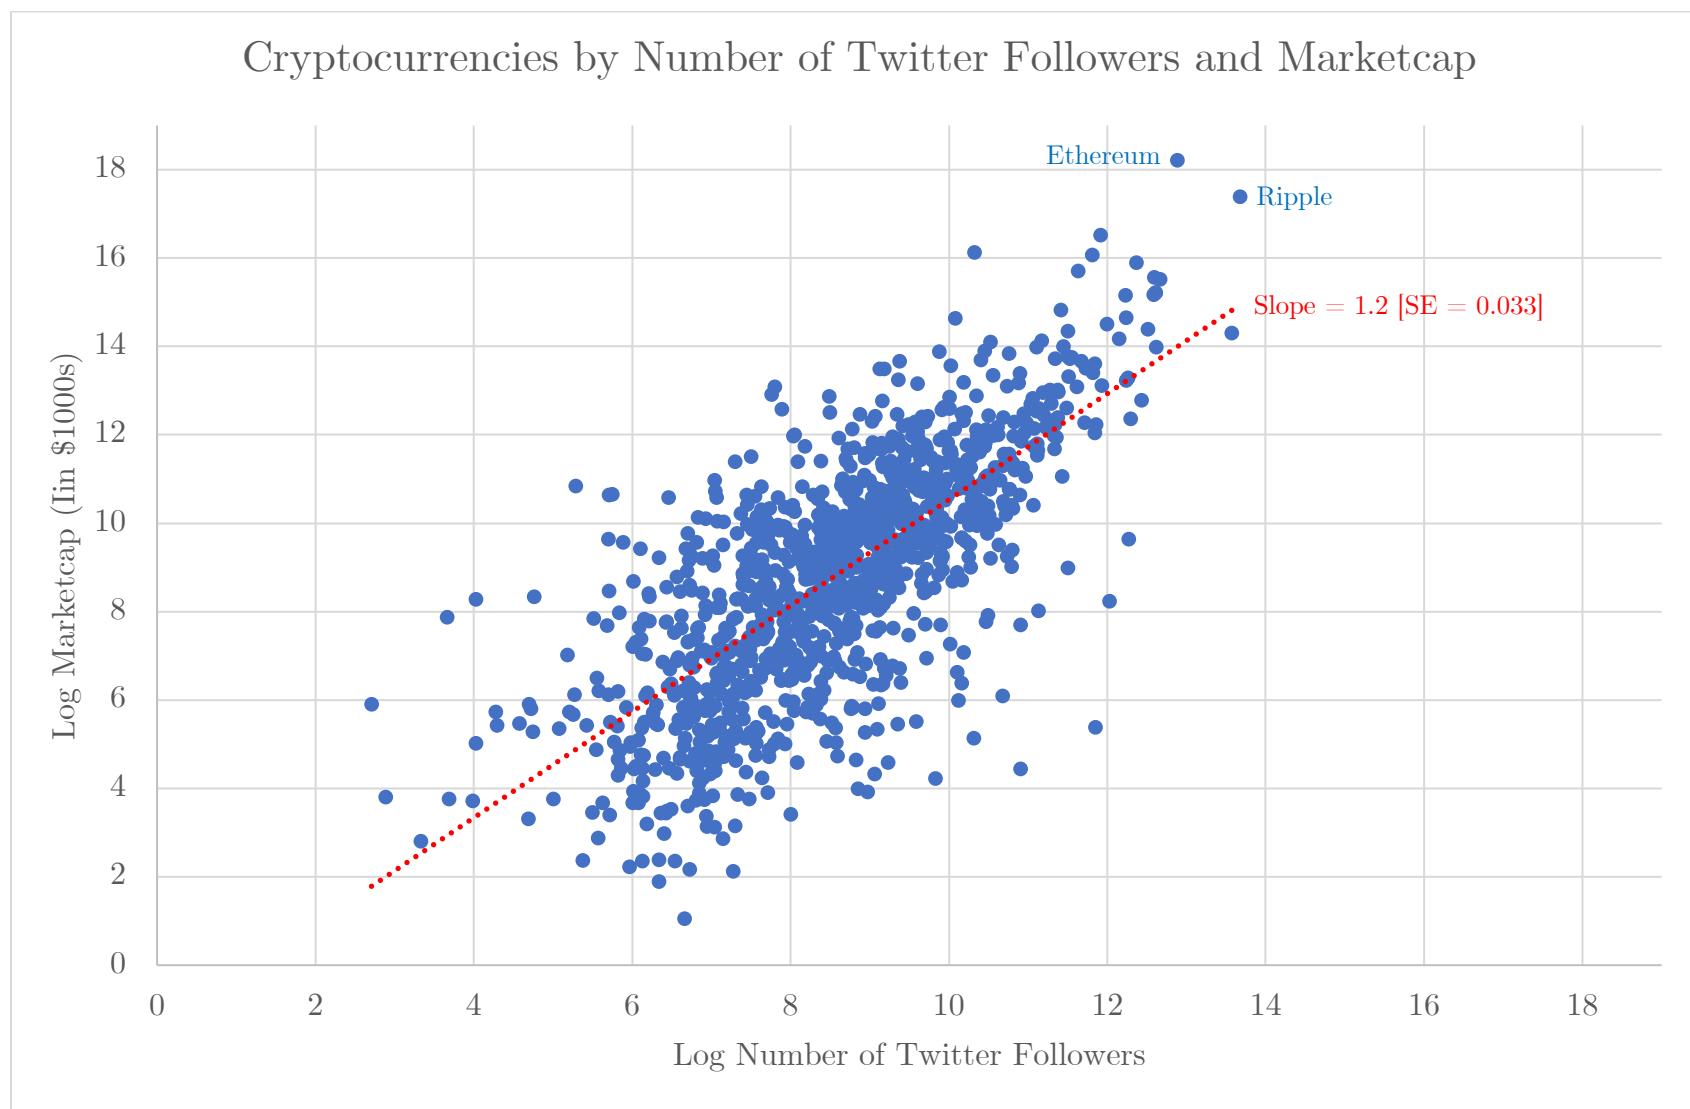


Figure 5: Initial coin offering post-listing cumulative returns

Figure 5 shows (in blue) the average buy-and-hold cumulative returns (for horizons less than 250 days) from investing in ICOs after they are listed on an exchange. For comparison, it also shows average returns to investing instead in a token index, which is a value-weighted average of all tokens, (in gray) or investing in Bitcoin (in yellow) at the time each ICO is listed.




Table 1: ICO summary statistics

Table 1 presents summary statistics on variables related to Initial Coin Offerings (ICOs). Panel A presents summary statistics for 2,390 ICOs that were completed on or before April 30, 2018. For each variable, it shows the number of non-missing observations, along with the cross-sectional mean, standard deviation, 10th, 50th, and 90th percentile values. Panel B reports means (and medians in brackets) for the same variables as Panel A but with six subsamples. Column (1) shows summary statistics for all ICOs that have no reported amount for raised capital, Column (2) shows them for ICOs that raised capital but were not listed on an exchange, Column (3) shows them for ICOs that raised capital and were listed on an exchange. Column (4) provides summary statistics for all ICOs that were completed on or before June 30, 2017, Column (5) for all ICOs that were completed between July 1, 2017 and December 31, 2017, and Column (6) for all ICOs that were completed on or after January 1, 2018. *Successful* is a dummy variable that equals one for all ICOs that raised capital, and zero otherwise. *Listed* is a dummy variable that equals one for all ICOs whose tokens were listed on an exchange, and zero otherwise. *Capital raised* is the amount (in millions of dollars) that was raised during the ICO. *ICO price* is the nominal price (in dollars) of the token during the ICO. *Hardcap* is the maximum amount (in millions of dollars) that could be raised during the ICO, which, when hit, would cause the offering to end. *Tokens % sold* is the percent of all tokens that would be sold during the ICO. *Pre-ICO* is a dummy variable that equals one if there was an earlier presale or pre-ICO for the token, and zero otherwise. *Length of ICO* is the number of days between the start and end of the ICO. *Country Rule of law* is the World Bank “rule of law” rating (median country has a value of zero) for the country from which the ICO was launched. *Country GDP per capita* is the GDP-per-capita (as of 2016, in dollars) for the country from which the ICO was launched. *Twitter age* is the number of months between the date when the platform’s Twitter account was activated and the ICO start date. *Pre-ICO Twitter intensity* is the average number of tweets per day in the six months before the ICO. *ICO Twitter intensity* is the average number of tweets per day during the ICO.

Panel A: Summary Statistics						
<u>Variables</u>	<u>Obs.</u>	<u>Mean</u>	<u>SD</u>	<u>p10</u>	<u>p50</u>	<u>p90</u>
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Successful (dummy)</i>	2390	0.48	0.50	0	0	1
<i>Listed (dummy)</i>	2390	0.26	0.44	0	0	1
<i>Capital raised, (\$MIL) (>0)</i>	1147	11.5	24.5	0.042	3.8	30.0
<i>ICO price (\$)</i>	1684	11.2	203	0.020	0.30	2.25
<i>Hardcap (\$MIL)</i>	924	42.9	90.9	5.0	23.0	75.0
<i>Tokens % sold</i>	1193	0.57	0.21	0.30	0.60	0.83
<i>Pre-ICO (dummy)</i>	2390	0.40	0.49	0	0	1
<i>Length of ICO (days)</i>	2324	37	28	8	31	68
<i>Country Rule of law</i>	1624	1.06	0.95	-0.80	1.63	1.84
<i>Country GDP per capita (\$000)</i>	1609	36.6	21.6	8.1	40.4	57.6
<i>Twitter age (months)</i>	1751	7.9	12.1	1	3	26
<i>Pre-ICO Twitter intensity</i>	1549	2.1	25.5	0.09	0.7	2.9
<i>ICO Twitter intensity</i>	1573	2.7	8.8	0.03	1.0	6.0

Panel B: Mean and [median] values of ICO characteristics for six subsamples of ICOs

Variables	Subsamples:					
	<i>Successful = 0</i>	<i>Successful = 1, Listed = 0</i>	<i>Successful = 1, Listed = 1</i>	<i>ICO ended before July 2017</i>	<i>ICO ended July-Dec 2017</i>	<i>ICO ended in 2018</i>
	Mean [Median]					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Capital raised, (\$MIL) (>0)</i>	0	6.1 [1.3]	17.0 [8.7]	7.9 [1.3]	11 [2.4]	14.0 [6.5]
<i>ICO price (\$)</i>	4.6 [0.5]	8.0 [0.4]	23 [0.1]	56 [0.1]	8.2 [0.3]	6.3 [0.3]
<i>Hardcap (\$MIL)</i>	55 [22]	34 [21]	38 [25]	88 [25]	39 [26]	43 [20]
<i>Tokens % sold</i>	0.6 [0.6]	0.6 [0.6]	0.5 [0.5]	0.5 [0.3]	0.6 [0.6]	0.6 [0.6]
<i>Pre-ICO (dummy)</i>	0.38 [0]	0.48 [0]	0.36 [0]	0.01 [0]	0.29 [0]	0.57 [1]
<i>Length of ICO (days)</i>	38 [31]	45 [32]	26 [28]	30 [30]	33 [30]	41 [31]
<i>Country Rule of law</i>	1.0 [1.6]	1.0 [1.6]	1.2 [1.7]	1.2 [1.7]	0.9 [1.6]	1.1 [1.6]
<i>Country GDP per capita (\$000)</i>	36 [40]	33 [40]	41 [43]	42 [42]	35 [40]	37 [40]
<i>Twitter age (months)</i>	6.6 [2]	7.4 [3]	9.4 [4]	8.0 [4]	8.3 [3]	7.2 [3]
<i>Pre-ICO Twitter intensity</i>	2.6 [0.5]	1.2 [0.7]	2.2 [1.0]	1.2 [0.8]	1.4 [0.7]	2.8 [0.7]
<i>ICO Twitter intensity</i>	1.2 [0.4]	2.0 [0.9]	5.3 [2.6]	5.6 [3.5]	2.6 [1.0]	2.4 [0.8]

Table 2: Summary statistics for listed firms

Table 2 presents summary statistics for 609 cryptocurrencies that were sold through an ICO and listed on an exchange. For each variable, it shows the number of non-missing observations, along with the cross-sectional mean, standard deviation, 10th, 50th, and 90th percentile values. If a token was already listed before the start of the ICO, it is not included in this sample. *Days to listing* is the number of days between the completion of the ICO and the first trading day. It is negative for 44 tokens which were exchange-listed before the ICO was officially completed. *ICO price* is the nominal price (in dollars) of the token during the ICO, as in Table 1. *First opening price* is the first (opening) price on the first day of trading, and *First closing price* is the last price on the first day of trading. Cryptocurrencies exchanges are open all the time so the “close” is based on midnight Coordinated Universal Time (UTC). *ICO->first open ROR* is the rate of return from the ICO price to the first day’s opening price, while *ICO->first close ROR* is the rate of return from the ICO price to the first day’s closing price. *First day marketcap* is the market capitalization of the cryptocurrency based on its price at the close of the first day of trading. For tokens which are missing values for their first day’s market capitalization, we impute number of tokens (and use them to calculate market-cap) using the first available market capitalization. *First month volume* is the average daily volume (in millions of dollars) in the first 30 days of trading. *Capital raised* is the amount (in millions of dollars) that was raised during the ICO, as in Table 1.

<u>Variables</u>	<u>Obs.</u>	<u>Mean</u>	<u>SD</u>	<u>p10</u>	<u>p50</u>	<u>p90</u>
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Days to listing</i>	609	30.5	49.4	1	16	76
<i>ICO price</i>	523	22.7	356.8	0.011	0.14	1.35
<i>First opening price</i>	609	25.9	417.2	0.014	0.19	2.45
<i>First closing price</i>	609	26.8	429.0	0.015	0.20	2.84
<i>ICO->first open ROR</i>	523	2.46	10.2	-0.67	0.21	4.39
<i>ICO->first close ROR</i>	523	2.73	10.5	-0.64	0.29	4.82
<i>First day marketcap</i>	548	54.7	201	0.9	12.2	107.0
<i>First month volume (avg daily, \$MIL)</i>	609	2.5	10.1	0.005	0.14	4.9
<i>Capital raised (\$MIL)</i>	557	17.2	30.0	0.4	9.1	38.2

Table 3: Returns to investing at ICO price and selling at first open

Table 3 presents average returns from investing in tokens at their ICO price, holding them until they are listed, and then selling them at the first day's opening price. The first two rows show average equal-weighted returns, assuming a naive investor who provides the same amount of funds for each project, while the last two rows show average value-weighted returns (weighted by capital raised), which assume a representative investor who uses the same factors that determine how much each project raises to decide the amount of funds to invest in that project. First and third rows show the returns in dollars, while the second and fourth rows convert prices to bitcoin (at the exchange rate for that day) before calculating returns. Column (1) shows returns for the sample of ICOs that were listed within 60 days. Columns (2) and (3) show returns for the sample of ICOs that raised capital and imputes returns of -50% (and -100% in Column (3)) to all tokens that were not listed or were listed after 60 days. Columns (4) and (5) show returns for the entire sample of ICOs and imputes returns of -50% (and -100% in Column (5)) to all tokens that were not listed or were listed after 60 days. All ICOs that ended on or after March 1, 2018 are dropped from the sample. T-statistics, using heteroskedasticity-consistent standard errors, are shown in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

	Sample: <i>Listed</i> = 1	<i>Successful</i> = 1		All ICOs	
Imputed ROR for unlisted:		-50%	-100%	-50%	-100%
	(1)	(2)	(3)	(4)	(5)
<u>Equal-Weighted Returns</u>					
<i>ICO->first open ROR</i>	1.79***	0.57***	0.31*	0.09	-0.28***
in dollars, EW	[5.38]	[3.58]	[1.90]	[1.00]	[3.12]
<i>ICO->first open ROR</i>	1.67***	0.52***	0.26	0.06	-0.31***
in bitcoin, EW	[5.14]	[3.32]	[1.60]	[0.68]	[3.52]
Observations	416	887	887	1619	1619
<u>Value-Weighted Returns</u>					
<i>ICO->first open ROR</i>	1.73***	1.05***	0.90**		
in dollars, VW	[3.52]	[2.99]	[2.51]		
<i>ICO->first open ROR</i>	1.62***	0.98***	0.82**		
in bitcoin, VW	[3.38]	[2.84]	[2.36]		
Observations	397	859	859		

Table 4: Returns to ICO investing – Liquidity and sensitivity to aggregate crypto-market

Table 4 presents average returns from investing in tokens at their ICO price, as in Table 3, but using different methods for calculating selling price (in Panel A) and for different subsamples of ICOs sorted on Bitcoin performance between the ICO end date and the listing date (in Panel B). In Column (1) of Panel A, returns to the ICO investor are calculated assuming the investor is able to sell all tokens at the first day's opening price (duplicating Column (1) of Table 3). In Columns (2) through (4), the sales price is now the volume-weighted average price (daily price is now the average of each day's open and close) over the first N days, when N is the minimum of the number of days by which the cumulative dollar volume reaches the defined volume limit or 60 days. The volume limit is defined to be \$1 million in Column (2), \$10 million in Column (3), and \$100 million in Column (4). If the cumulative volume does not hit the volume limit by the end of the 60th day, we assume the remaining fraction of the tokens are sold at the closing price of the 60th day with a 50% liquidity penalty. In Panel B, we split the sample of ICOs into five groups (quintiles) based on the return of Bitcoin between the ICO end date and the listing date, and report returns from investing at the ICO price and selling at the first day's opening price, separately for each quintile. All ICOs that ended on or after March 1, 2018 are dropped from the sample. T-statistics, using heteroskedasticity-consistent standard errors, are shown in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Panel A: Returns using volume-weighted selling price calculations over first 60 days				
Volume Limit:	None	\$1 million	\$10 million	\$50 million
	(1)	(2)	(3)	(4)
<u>Equal-Weighted Returns</u>				
<i>ICO->Exchange ROR</i>	1.79***	2.19***	1.79***	1.48***
in dollars, EW	[5.38]	[3.50]	[5.35]	[6.16]
<i>ICO->Exchange ROR</i>	1.67***	1.79***	1.33***	1.02***
in bitcoin, EW	[5.14]	[4.51]	[5.30]	[5.31]
Observations	416	416	416	416
<u>Value-Weighted Returns</u>				
<i>ICO->Exchange ROR</i>	1.73***	2.59**	2.01***	1.42***
in dollars, VW	[3.52]	[2.56]	[2.78]	[4.50]
<i>ICO->Exchange ROR</i>	1.62***	2.20***	1.74**	1.16***
in bitcoin, VW	[3.38]	[2.72]	[2.57]	[4.14]
Observations	397	397	397	397

Panel B: Returns for quintiles based on bitcoin performance from end of ICO to listing					
Bitcoin Perf. Quintile:	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
	(1)	(2)	(3)	(4)	(5)
<u>Equal-Weighted Returns</u>					
<i>ICO->first open ROR</i> in dollars, EW	1.24*** [2.58]	1.78** [2.25]	2.51** [2.18]	1.10*** [4.60]	2.33*** [3.14]
<i>ICO->first open ROR</i> in bitcoin, EW	2.01*** [3.07]	2.03** [2.31]	2.46** [2.16]	0.89*** [4.10]	1.00*** [2.77]
Observations	83	83	83	83	84
<u>Value-Weighted Returns</u>					
<i>ICO->first open ROR</i> in dollars, VW	0.62*** [3.44]	2.54 [1.44]	1.38*** [3.97]	1.21*** [2.75]	3.17*** [2.70]
<i>ICO->first open ROR</i> in bitcoin, VW	1.20*** [4.57]	2.78 [1.50]	1.36*** [3.95]	0.98** [2.49]	1.44** [2.51]
Observations	80	80	82	78	77

Table 5: Cross-sectional determinants of ICO pricing

Table 5 present estimated coefficients from cross-sectional regressions using OLS (in Columns (1) and (4)) and Heckman (1979) procedure using maximum likelihood estimation. The dependent variable in each regression is the natural logarithm of the ratio of the first day's opening price of the token to its ICO price. In Columns (1) and (2), the prices are denominated in dollars, while in columns (3) through (6), they are denominated in Bitcoins to adjust for changes in cryptocurrency values from the ICO to listing. Predictive variables are defined in Table 1. T-statistics, using heteroskedasticity-consistent standard errors, are shown in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

<u>Predictive Variables</u>	Dependent Variable: <i>Log (First Opening Price/ICO Price)</i>					
	OLS (1)	Heckman (2)	Heckman (3)	OLS (4)	Heckman (5)	Heckman (6)
<i>Date</i>	-0.0004 [1.38]	-0.0007** [2.04]	-0.0007** [2.04]	-0.0012** [2.41]	-0.0015*** [3.16]	-0.0016*** [3.40]
<i>Log ICO price</i>	-0.055** [2.00]	-0.059** [2.09]	-0.061** [2.11]	-0.037 [1.30]	-0.034 [1.20]	-0.028 [0.95]
<i>Pre-ICO (dummy)</i>	-0.217* [1.71]	-0.243* [1.94]	-0.212* [1.69]	-0.092 [0.63]	-0.083 [0.59]	-0.075 [0.54]
<i>Twitter age (months)</i>	-0.001 [0.28]	-0.003 [0.79]	-0.005 [1.50]	-0.002 [0.57]	-0.004 [0.94]	-0.004 [1.17]
<i>Country Rule of law</i>				0.084 [0.97]	0.056 [0.68]	0.043 [0.53]
<i>Pre-ICO Twitter intensity</i>				-0.009 [0.60]	-0.006 [0.39]	-0.004 [0.25]
<i>ICO Twitter intensity</i>				0.006 [1.35]	0.005 [1.21]	0.004 [1.01]
<i>Length of ICO (days)</i>						-0.005 [1.32]
<i>Log Capital raised</i>						0.170*** [2.73]
Observations	408	1611	1611	302	1505	1505
Prices in	Dollars	Dollars	Bitcoin	Bitcoin	Bitcoin	Bitcoin
R-squared	0.035			0.05		
rho		-0.27	-0.31		-0.30	0.01
LR test for rho=0 p-value		0.005	0.001		0.002	0.89

Table 6: Twitter official account activity during and after ICOs

Table 6 presents the Twitter official account activity for three different categories of ICOs: those that didn't raise capital and didn't list, those that raised capital and didn't list, and those that listed. For each ICO category, we present the number observations for that category in Column (1), the fraction of ICOs that don't have a Twitter account in Column (2), and the fraction of ICOs that have an account but had no activity during the ICO in Column (3). For the remaining calculations, we exclude all the ICOs with no account or no Twitter activity during the ICO. In Columns (4) through (9), we report the cross-sectional average of Twitter intensity (average number of Tweets per day) for different periods of time (during and after the ICO), along with the percentage of tokens that had no Twitter activity during each period. In Column (4), the period of interest is the ICO itself. The % of observations with 0 tweets is missing for this column because we dropped all ICOs that didn't have Twitter intensity during the ICO, as explained above. In Column (5), the period of interest is the 30 days after the ICO; in Column (6), it is the period from 31 days to 60 days after the ICO; in Column (7), it is the period from 61 days to 90 days after the ICO; in Column (8), it is the period from 91 days to 120 days after the ICO; in Column (9), it is the period from 121 days to 150 days after the ICO.

<u>Sample</u>	<i>Number of ICOs</i> (1)	<i>No Twitter Account</i> (2)	<i>ICO Twitter intensity = 0</i> (3)	<u>Statistic</u>	<i>During ICO</i> (4)	<i>1-30 Days After ICO</i> (5)	<i>31-60 Days After ICO</i> (6)	<i>61-90 Days After ICO</i> (7)	<i>91-120 Days After ICO</i> (8)	<i>121-150 Days After ICO</i> (9)
Capital = 0, Listed = 0	694	51.4%	7.8%	<i>Twitter intensity</i>	1.48	0.53	0.49	0.41	0.29	0.28
				% Obs w/ 0 Tweets		37.5%	45.6%	54.1%	58.3%	59.2%
Capital > 0, Listed = 0	420	21.2%	5.2%	<i>Twitter intensity</i>	2.65	0.86	0.73	0.79	0.91	0.87
				% Obs w/ 0 Tweets		15.9%	21.4%	27.2%	28.6%	35.2%
Listed = 1	440	2.7%	6.8%	<i>Twitter intensity</i>	5.48	1.76	1.58	1.64	1.59	1.47
				% Obs w/ 0 Tweets		1.0%	5.5%	6.5%	7.0%	6.8%

Table 7: Regressions of daily cryptocurrency returns on Twitter intensity (with lags and leads)

Table 7 presents estimated coefficients from OLS regressions of daily cryptocurrency returns on daily measures of Twitter intensity (number of Tweets from the cryptocurrency’s official Twitter account per day), with various lags and leads for the independent variable, along with controls. The dependent variable in all specifications is the natural logarithm of one plus the daily return (using closing prices), where the log transformation is done to minimize the effect of large outliers. In Column (1), the explanatory variables are today’s and yesterday’s Twitter intensity, along with lagged *Log Marketcap*, the natural logarithm of the market capitalization from the last period. In Column (2), we also add the 2nd, 3rd, 4th, and 5th lags of Twitter intensity. In Column (3), we also add *Twitter intensity* [-10,-6], the average Twitter intensity from 10 days to 6 days earlier, and *Twitter intensity* [-40,-11], the average Twitter intensity from 40 days to 11 days earlier. In Column (4), we add fixed effects for each asset. In Column (5), we replace each instance of Twitter intensity with its log transformation, i.e., natural logarithm of one plus the number of Tweets per day. In Column (6), we add five leads of *Twitter intensity*, for the next 5 days after the day that we measure the returns. All specifications also include time dummies, and various number of lags of the dependent variable (up to 10). T-statistics, using heteroskedasticity-consistent standard errors, are shown in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Dependent Variable:		<i>Log (1+Returns) (t=0)</i>					
<u>Predictive Variables</u>	(1)	(2)	(3)	(4)	(5)	(6)	
<i>Twitter intensity,</i> t=0	0.0020*** [7.56]	0.0026*** [8.44]	0.0029*** [11.26]	0.0028*** [11.23]	0.0205*** [19.10]	0.0022*** [10.23]	
<i>Twitter intensity,</i> lag 1	-0.0006*** [4.03]	0.0001 [1.34]	0.0004*** [3.15]	0.0004*** [2.90]	0.0047*** [6.55]	0.0001 [0.42]	
<i>Twitter intensity,</i> lag 2		-0.0004*** [2.88]	-0.0003** [2.26]	-0.0003** [2.43]	-0.0005 [0.81]	-0.0005*** [3.70]	
<i>Twitter intensity,</i> lag 3		-0.0006*** [4.35]	-0.0004*** [3.14]	-0.0005*** [3.30]	-0.0019*** [2.82]	-0.0006*** [4.48]	
<i>Twitter intensity,</i> lag 4		-0.0002** [2.01]	0.0001 [0.36]	0.0000 [0.18]	-0.0002 [0.34]	-0.0001 [0.94]	
<i>Twitter intensity,</i> lag 5		-0.0006*** [2.64]	-0.0001 [0.73]	-0.0002 [0.85]	-0.0014** [2.17]	-0.0003 [1.56]	
<i>Twitter intensity,</i> [-10,-6]			-0.0010*** [3.44]	-0.0011*** [3.65]	-0.0083*** [6.32]	-0.0014*** [4.30]	
<i>Twitter intensity,</i> [-40,-11]			-0.0009*** [3.12]	-0.0013*** [3.26]	-0.0094*** [7.74]	-0.0014*** [3.71]	
<i>Twitter intensity,</i> lead 1						0.0014*** [8.12]	
<i>Twitter intensity,</i> lead 2						0.0005*** [3.34]	
<i>Twitter intensity,</i> lead 3						0.0006*** [4.03]	
<i>Twitter intensity,</i> lead 4						0.0005*** [3.37]	
<i>Twitter intensity,</i> lead 5						0.0001 [1.22]	
<i>Log Marketcap,</i> lag 1	-0.0013*** [5.61]	-0.0008*** [3.65]	-0.0008*** [3.74]		-0.0007*** [3.29]	-0.0007*** [2.92]	
Observations	491683	477590	456713	456708	456713	453132	
Time dummies	Y	Y	Y	Y	Y	Y	
Currency fixed effects	N	N	N	Y	N	N	
Log return lags	2	5	10	10	10	10	
Log (1+Tw. Intensity)	N	N	N	N	Y	N	

Table 8: Cumulative post-listing returns for ICOs

Table 8 presents the average cumulative buy-and-hold returns from investing in tokens that had an ICO at the opening price on the first day they are trading on an exchange, and then holding them for seven different horizons: 1 day, 7 days, 14 days, 30 days, 90 days, 180 days, and 365 days. Column (1) shows just the raw returns without adjusting for the cryptocurrency market. Column (2) adjusts each token's return for the asset class returns by subtracting the cumulative return in the same period for Bitcoin. For Column (3), we create an index which is the value-weighted return of all cryptocurrency tokens (not coins like Bitcoin) and adjusts each token's return by subtracting the cumulative return from investing in this index. For Column (4), we match each token on the listing date to the cryptocurrency that has the closest market capitalization to that token and that has been listed for at least one year. We then adjust each token's return by subtracting the cumulative return from investing in the matched asset. T-statistics, using heteroskedasticity-consistent standard errors, are shown in brackets. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

	Raw Returns	Bitcoin Adj.	Tokens Adj.	Match Adj.
<u>Horizon</u>	(1)	(2)	(3)	(4)
1st Day:	0.160***	0.151***	0.161***	0.141***
	[9.17]	[8.74]	[9.12]	[7.68]
Obs:	609	609	600	588
7 Days:	0.332***	0.293***	0.294***	0.307***
	[4.65]	[4.12]	[4.11]	[4.12]
Obs:	591	591	582	571
14 Days:	0.302***	0.230***	0.215***	0.259***
	[4.65]	[3.56]	[3.39]	[3.73]
Obs:	578	578	569	559
30 Days:	0.671***	0.483**	0.414**	0.577***
	[3.33]	[2.40]	[2.09]	[2.78]
Obs:	561	561	552	543
90 Days:	1.405***	0.778***	0.224	-0.056
	[6.06]	[3.40]	[1.02]	[0.14]
Obs:	467	467	458	453
180 Days:	4.309***	2.425*	1.528	1.873
	[3.24]	[1.83]	[1.14]	[1.35]
Obs:	293	293	284	283
365 Days:	18.802**	12.936*	2.910	-0.763
	[2.39]	[1.66]	[0.34]	[0.08]
Obs:	80	80	71	76