The Wisdom of Crowds in FinTech: Evidence from Initial Coin Offerings*

Jongsub Lee† Tao Li‡ Donghwa Shin§

First draft: May 2018
This draft: September 2018

Abstract

Certification by a crowd of online analysts and early investors can generate excitement among potential token investors, leading to successful initial coin offerings ("ICOs"). We test the general notion of "wisdom of crowds" using novel data on over 1,500 ICOs, including sequential investor subscriptions *during* token sales. We find that favorable analyst opinions on the underlying project are associated with aggressive initial token subscriptions, which predicts subsequent token sales. Analyst ratings also predict long-run token performance in the secondary market. Overall, our results suggest that the wisdom of crowds could substitute underwriters intermediary role in financing decentralized blockchain-based startups.

Keywords: ICO, FinTech, Wisdom of crowds, Information cascade, Fundrasing success, Underpricing, Long-run performance

---

*The authors have benefited tremendously from comments and suggestions from Brad Burnham, Chris Burniske, Mark Flannery, Campbell Harvey, Mark Jamison, Raghavendra Rau, Jay Ritter, Wei Xiong, and participants at the 18th FDIC Bank Research Conference and the Florida Blockchain Think Tank. We thank Daniel Levy, Chase Maxwell, Joon-Young Park, Kevin Quach, Hongyu Shan, and Chengshuo Zhang for their excellent research assistance.

†Associate Professor of Finance, Seoul National University (SNU) Business School, 1 Gwanak-ro, Gwanak-gu, Seoul 08826, South Korea, Phone: +82.2.880.69521, jongsub.lee@snu.ac.kr, webpage: www.jongsublee.com.

‡Assistant Professor of Finance, Warrington College of Business, University of Florida, PO Box 117168, Gainesville, FL, USA. Phone: +1 (352) 392-6654, fax: +1 (352) 392-0301, Tao.Li@warrington.ufl.edu, webpage: https://site.warrington.ufl.edu/tao-li.

§Ph.D. candidate, Department of Economics, Princeton University, Julis Romo Rabinowitz Building, Princeton, NJ, USA, donghwa@princeton.edu.

1 Introduction

In recent years, initial coin offerings ("ICOs") have emerged as a popular way to fund blockchain-related startups, raising $13.7 billion in the first five months of 2018 (PwC, 2018). Through this new form of crowdfunding, an entrepreneur raises capital by creating and selling a virtual currency or "token," which provides a set of rights to its holders, including access to a platform, and can be resold in the secondary market.

In such initial offerings, information asymmetry between an entrepreneur and outside investors and the resulting adverse selection are two inherent problems that could hinder successful fundraising (e.g., Akerlof, 1970; Leland and Pyle, 1977; Rock, 1986). To overcome asymmetric information problems, traditional primary market issuances, in particular initial public offerings ("IPOs"), are delegated to financial intermediaries that perform due diligence, determine the value and riskiness of the businesses, and price and sell the new securities (Diamond, 1984). For example, IPO underwriters conduct a bookbuilding process that helps gauge investor demand and price the new equity effectively. However, in blockchain-based crowdfunding, a fast-growing FinTech sector, such financial intermediation is absent by design thanks to decentralized bookkeeping enabled by blockchain networks. What makes ICOs so successful without financial intermediation? Rather than piggybacking on underwriter-driven intermediation, the ICO market relies heavily on the "wisdom of crowds" – the collective action of a large group of individuals rather than that of a single expert.

In this paper, we explore how the wisdom of crowds helps overcome the information asymmetry associated with an ICO. This mechanism works in two stages by (1) certifying the quality of the underlying venture before a token sale starts, and (2) harnessing the wisdom of crowds during the fundraising period. The diverse opinions from a number of online analysts can lead to an aggregate signal that closely reflects the true quality of the risky startup. Such market-based certification by individual analysts not only screens out "lemons" (e.g., Akerlof, 1970; Leland and Pyle, 1977), but is also associated with significant investor interest during initial periods of the sale, encouraging subsequent investors to invest regardless of their own information (Welch, 1992). Unlike in IPO bookbuilding, each token subscription in an ICO is broadcast to all potential investors through a blockchain network. With a critical mass of supporters, harnessed wisdom of crowds would quickly result in meeting the pre-specified funding target (Cong and Xiao, 2018; Li and Mann, 2018).

Our analysis, built on a novel and comprehensive sample of 1,549 completed ICOs from January 2016 to March 2018, aims to uncover how the crowds' wisdom affects fundraising success, an event in which an ICO reaches its minimum fundraising target ("soft cap").

Equally importantly, we examine whether the "wisdom of crowds" phenomenon is related to tokens' secondary market performance. Our study is unique in its reliance upon primary market subscription information *during* the fundraising period, and only this allows for an empirical analysis of investor subscription dynamics in ICOs (Cong and Xiao, 2018; Li and Mann, 2018). To the best of our knowledge, we are the first to utilize these sequential transaction records to study primary market offerings.[1]

First, the probability of a successful fundraising campaign increases by 19.8 percentage points (relative to the unconditional success probability of 42.7%) with every one-standard-deviation increase in the average analyst rating, controlling for ICO characteristics. This result supports the positive intermediary role these experts play in a market without traditional underwriters. Due to reputational concerns, the majority of our analysts from ICObench, the most comprehensive database on ICOs, are likely to be unbiased in issuing their ratings.[2] The result remains qualitatively similar, however, when we exclude analysts who are founders of other ICOs, and thus are more likely to be biased. Similarly, the number of experts covering a token sale also positively predicts fundraising success. These results suggest that outside investors tend to follow a large crowd of informed analysts when making risky investment decisions.

The fraction of tokens for sale to outside investors would be inversely related to management's skin in the game. Without formal governance and incentive mechanisms post-sale, such as voting power to oust directors and stock incentive plans for managers, a significant stake retained by management could play a crucial governance role to align incentives of the management team and token holders (Leland and Pyle, 1977). Consistent with this signaling theory, we find that a one-standard-deviation increase in the fraction of tokens for sale is associated with a 6.1 percentage point decrease in the marginal probability of ICO success.

In addition, various other ICO mechanisms could bolster investor participation by promoting consumer demand and governance. We find that arranging a (private) presale can boost ICO success likelihood by 15.2 percentage points, indicating that successful presales are interpreted by later investors as evidence that earlier investors held favorable information (Cong and Xiao, 2018). In the absence of underwriters, a presale is also a clever way for insiders to gauge demand from informed investors such that they can set a more informed price in the main sale. This in turn leads to a higher likelihood of success.

[1] In a concurrent paper, Howell, Niessner, and Yermack (2018) show summary statistics of investor subscriptions in a case study on Filecoin, which is included in our sample.

[2] Most of the experts in our sample are repeat analysts, who are likely concerned about their reputation in the long run. They use real names and disclose biographic information. Moreover, ICObench assigns each analyst an ICO Success Score, which is publicly accessible. See Section 3.2 for more details.

Although early-bird bonus tokens could attract investors early in the process, excessive bonuses or discounts can result in credit rationing by rational investors as many of these sales are believed to be scams (Stiglitz and Weiss, 1981; Sagar, 2017). We find that ICOs providing large bonuses or discounts are 10.9 percentage point *less* likely to conclude successfully. This echoes the U.S. Securities and Exchange Commission's (SEC) recent warnings against token sales with overly generous bonuses.

ICOs featuring multi-language whitepapers/websites or accepting multiple currencies are significantly more likely to succeed, reflecting the global nature of token sales and the ease of transactions with an expanded set of payment methods. However, anti-money laundering measures, such as a Know Your Customer policy, negatively predict fundraising success. This suggests that these investor governance proxies are associated with higher transaction costs, reducing the participation incentives of outside investors in token sales.

Consistent with our main analysis above, we find that favorable analyst ratings positively predict total gross proceeds, while the fraction of tokens for sale, high bonuses, and a Know Your Customer policy are associated with a lower amount raised. We also use a Cox proportional hazards model to perform a duration analysis, which largely yields consistent results. Interestingly, token sales that feature multi-language websites or whitepapers or accept multiple currencies take a longer time to complete. This could be explained by the fact that these ICOs mainly target (small) retail investors, who typically participate in later stages of the sale by "following the crowd."

To investigate the channel through which quality certification by informed analysts leads to successful fundraising, we resort to the unique primary market subscription data collected from Etherscan, a leading "block explorer" that provides information on every transaction for Ethereum-based ICOs. For each token sale, we download information for all the transactions taking place between the ICO start and end dates. We obtain a transactions sample of over 650 ICOs, and for each ICO transactions are aggregated to the daily frequency.

This novel data set reveals that successful token sales attracted more than 2,000 backers on average, compared to the 39 supporters in failed ICOs. In successful ICOs, early investors also purchase tokens more aggressively, with the first-day per capita subscription at nearly 0.5% of token supply. This highlights the importance of the certification role large investors play in helping coordinate between investors (Sockin and Xiong, 2018). Furthermore, investor subscriptions during initial periods (even at the block-level) strongly predict subsequent token sales and eventual fundraising success, which is consistent with an "up-cascade" phenomenon described by Cong and Xiao (2018). Importantly, better analyst ratings ex ante not only predict stronger first-day token sales, but are also associated with faster sales

in earlier periods. Decentralized blockchain-based fundraising is also in line with the notion of "crowd"-funding illustrated by Strausz (2017).

We also find that early token investors achieve better investment returns than their followers. This ability of earning superior returns corroborates their role in leading the crowd of subsequent investors, and justifies their participation decisions in early stages of an ICO.

Before turning our attention to whether and how our proposed certification mechanisms affect post-ICO token performance, we present key statistics about secondary market trading. Borrowing the terminology from the IPO literature, we first document that the first-day return or short-run underpricing (e.g., Ibboston, 1975; Ritter, 1984; Beatty and Ritter, 1986) is significantly positive for listed tokens. The average underpricing is nearly 160%, substantially higher than the average IPO underpricing of 18% for the U.S. (Ritter, 2004). However, the median ICO underpricing is just 24%. Although newly issued tokens on average do not appear to underperform in the long run (up to one year), the median long-run excess return over a cryptocurrency index is -125%, statistically significant at the 1% level. This is consistent with the well-documented phenomenon of long-run IPO underperformance in small stocks (Ritter, 1991). The average (median) first-day turnover is 6.7% (1.3%), considerably lower than that for IPOs (Ritter, 2018).

Regarding cross-sectional patterns, we find that ICOs with generous bonuses have significantly lower first-day returns, all else being equal. This is consistent with Habib and Ljungqvist (2001) as bonuses can be viewed as marketing tools to attract early investors. Rational secondary market investors will request a lower price that is close to the intrinsic value of tokens solyear returns, suggesting that market-based certification by large crowds effectively signals the genuine quality of the startups in the long run (beyond the primary market phase).

Overall, we find strong empirical support for market-based certification mechanisms that are in play behind the emerging success of ICOs. In particular, decentralized analysts' opinions appear to mitigate information asymmetry and predict profitable investments in the long run. These results are consistent with the general notion of wisdom of crowds in the FinTech era (Strausz, 2017; Cong and Xiao, 2018; Li and Mann, 2018).

We make the following three contributions to the literature. First, we carefully examine the determinants of successful ICO fundraising campaigns. Following industry convention, we consider an ICO successful if it meets the minimum fundraising target. Existing empirical research on ICOs, represented by Amsden and Schweizer (2018) and Momtaz (2018), use exchange listing as their main criterion for ICO success. Because it can take a startup up

to several months to list its token post-sale, this exchange trading-based definition of ICO success is likely to suffer from substantial measurement errors. In fact, among successful ICOs in our sample, only 60% saw their tokens listed as of May 31, 2018. Collecting information on soft caps and gross proceeds for all ICOs in our sample requires a tremendous amount of manual work, which enables our paper to make a significant contribution to the empirical literature on ICOs.[3]

Second, we analyze how our proposed two-tier mechanisms of wisdom of crowds affect both ICO fundraising and secondary market performance. We not only identify what explains ICO fundraising success, but also factors that affect tokens' short-run underpricing and long-run performance. The fact that analyst ratings predict long-term superior investment performance is a deep cross-sectional insight on token pricing.

Last but not least, we are the first to analyze primary market subscription patterns *during* token sales. Using novel data on sequential investor subscriptions, we find evidence that is consistent with the existence of up-cascaded wisdom of crowds as predicted in Cong and Xiao (2018). That is, first-day token sales strongly predict subsequent token purchases and eventual success. Relatedly, positive analyst ratings ex ante not only predict stronger first-day subscriptions, but is also related to faster sales in earlier periods.

In addition to Cong and Xiao (2018), Li and Mann (2018), Sockin and Xiong (2018), our paper is related to several other studies in the context of ICOs. Catalini and Gans (2018) show that by eliciting consumers' willingness to pay, ICOs may increase entrepreneurial returns beyond what can be achieved through traditional equity financing. Chod and Lyandres (2018) demonstrate that ICOs can facilitate risk-sharing without diluting control rights. Canidio (2018) derives an equilibrium with a positive probability that the entrepreneurs may not develop any products post-ICO.

Our analysis on post-ICO token performance relates to Cong, Li, and Wang (2018), who develop a dynamic asset-pricing model of tokens and that features inter-temporal feedback effects. Howell, Niessner, and Yermack (2018) find that liquidity is higher for listed tokens when the issuers offer voluntary disclosure. Studies in this area also include Gandal and Halaburda (2014), Gans and Halaburda (2015), Athey, Parashkevov, Sarukkai, and Xia (2016), Ciaian, Rajcaniova, and Kancs (2016), Fernandez-Villaverde and Sanches (2016), and Pagnotta and Buraschi (2018).

[3]Amsden and Schweizer (2018) use only hard cap information for presales, which can not help identify eventual fundraising success. Momtaz (2018) defines ICO success based on positive returns measured from the first-day opening price (not the offer price) to first-day closing price. His return-based definition of success misclassified nearly 25% of successful ICOs in our sample, including several most successful ICO, such as Dragon Coins (the largest ICO ever), EOS, and Status, among others.

Another related area is the burgeoning literature on the economics of the blockchain technology. Yermack (2017) considers how the blockchain technology can lead to changes in corporate governance. Harvey (2016) discusses the mechanics of cryptofinance and their applications including Bitcoin. Cong and He (2018) show that blockchain-based decentralization can mitigate information asymmetry and improve welfare. Malinova and Park (2017) show that the transparent trading environment enabled by blockchain technology improves investor welfare. Biais, Bisiere, Bouvard, and Casamatta (2018) and Eyal and Sirer (2014) analyze cryptocurrency mining games, while Easley, O'Hara, and Basu (2017), Huberman, Leshno, and Moallemi (2017), and Cong, He, and Li (2018) analyze the compensation and organization of miners.

Lastly, our paper relates to an emerging literature on the wisdom of crowds, which include Surowiecki (2005), Kovbasyuk (2011), Da and Huang (2015), Kremer, Mansour, and Perry (2014), and Dindo and Massari (2017).

2 Institutional Background

In this section, we briefly introduce the concept and process of an ICO and describe important events that take place after the ICO is completed. We proceed with a brief discussion of the current regulatory environment before concluding with two examples of ICOs. Our overarching goal is to highlight inherent information asymmetry and governance challenges associated with ICOs.

2.1 What is an ICO?

An ICO is a new fundraising method made possible by the development of blockchain technology and cryptographic tokens. Through an ICO, a technology startup creates and distributes its (decentralized) platform's digital tokens in exchange for cryptocurrencies, such as Ether ("ETH") or Bitcoin ("BTC"), or fiat currencies to raise public capital to fund their operations and product development. The token typically provides a specific set of rights to its holders, including access to a platform or network, rights to create or develop features for an ecosystem, the right to cast a vote on governance issues, among others.

This approach is radically different from the traditional corporate IPO. With an IPO, investors exchange money for equity shares and voting rights in a relatively established company. In the U.S. and abroad, the process is underwritten by an investment bank and tightly regulated by securities regulators. To begin the IPO process, a firm is required to

file a registration statement, which is a set of documents including a prospectus. Company management and the underwriter conduct a "road show" to meet potential investors and gauge demand for the stock. The underwriter then "builds a book" by accepting orders from investors, who indicate the number of shares they desire and the price they are willing to pay. After the offer price is determined, the management files a final prospectus with the regulatory authority and shares are allocated to investors. The issue is typically "closed" a few days later and shares begin their secondary market trading on a stock exchange. The underwriter also commits to making a liquid secondary market by assigning analysts to cover the stock, and when necessary, it will step in to support the price.

In the freewheeling world of ICOs, however, none of these features exist. There is no investment bank to underwrite the token, conduct a bookbuilding, or support secondary market trading. Token sales usually are open to investors around the world, regardless of where the startup is based. Unlike such arrangement in an IPO, insiders' tokens are often not subject to a lock-up period of 90 to 180 days after an ICO. The vast majority of tokens are deemed as utility tokens as opposed to securities or equity stakes. They typically lack voting rights hence control. As of today, regulatory oversight is minimal in that blockchain startups are not required to file any regulatory documents. In most cases, the startups have no corporate track records or even products, although more established technology firms are increasingly using ICOs to fund their operations. Table A1in the Appendix compares the fundraising steps between ICOs and IPOs.

One major difference between ICOs and traditional crowdfunding is that the latter investment is less liquid. Even in equity crowdfunding in which investors obtain a financial stake in the company they support, it is difficult to resell their securities due to a lack of liquidity. However, when tokens generated in an ICO are listed on exchanges, they provide the buyers liquidity and a potentially positive rate of return when they are sold at a higher price.

2.2 The ICO process and post-ICO events

A typical ICO begins with the presentation of a whitepaper, which describes the business idea and model, the team, and the technical specifications of a project before the ICO. The entrepreneurs lay out a timeline for the project and describe how raised funds will be spent, such as on marketing, and research and development. They often specify a "soft cap" that is the minimum amount received at which the initial offering will be considered a success. Startups usually specify a "hard cap" as well, which is the maximum fundraising goal for a crowdsale. Most projects set very high hard caps that they are unlikely to achieve. In our

sample, only 12.2% of ICOs hit their hard caps.

An ICO fixes the number of tokens on offer before the sale. The whitepaper and/or the project website features a discussion of how the tokens will be distributed, including how many tokens are for sale, and how many tokens the insiders will keep.[4] Retaining a reasonably high fraction of tokens with the firm can send a signal to the market that the entrepreneurs have more skin in the game, thus are more likely to expend serious efforts in developing the project (Leland and Pyle, 1977; Downes and Heinkel, 1984).

Investors who purchase tokens early may be given preferential terms, in the form of an "early bird" bonus or discount. One purpose of the bonus or discount is to compensate for the higher risks early buyers bear. Some ICOs include a presale period or also known as a pre-ICO, a token sale event that startup enterprises run before the official crowdsale campaign goes live. Presales generally target larger investors, many of whom are institutional investors. The fundraising targets for presales are usually lower than those of the main sales and tokens are typically sold at a steep discount. Some companies run presales to collect funds to pay for the expenses incurred for launching their main ICOs. Investors and regulators may be wary of ICOs that provide extremely high bonuses, which sometimes exceed 100% (equivalent to a 50% discount). For example, the SEC has warned investors against token sales that offer high discounts to early buyers.[5] Offering too high a discount can lead to credit rationing from investors (Stiglitz and Weiss, 1981).

Due to the open-source nature of blockchain-based projects, before ICOs start, many startups choose to publish part or all of their initial codes that utilize smart contracts. Once the ICO period is set, the marketing campaign starts, which often begins with an announcement of the token sale on BitcoinTalk.org, a favorite social website of cryptocurrency enthusiasts. Other social networks, such as Medium, Steemit, Reddit, and Twitter, are often used as well.

By industry convention, an ICO is considered as a success if the amount it collects surpasses the soft cap. If a token sale does not reach its soft cap, funds are usually returned to investors. This is the "all-or-nothing" arrangement commonly used in ICOs. In rare cases, the team may decide to move forward regardless. If the hard cap is reached, additional subscriptions will be rejected and the funds will be returned.

After an ICO is successfully completed, the entrepreneurs typically begin to plan for an

[4]Some whitepapers also specify how many (reward) tokens are reserved for bounty programs. During a bounty program, the startup provides compensation for a number of tasks including marketing on social networks, bug reporting or even improving aspects of the cryptocurrency framework. The bounty tokens usually are a small percentage of the total supply of tokens.

[5]In May 2018, the SEC set up a website, HoweyCoins.com, that imitates a fake ICO to educate retail investors about red flags of a potential scam. This ICO features a "double 25% discount" for early investors.

exchange listing. Most cryptocurrency exchanges require an application and a listing fee, and depending on each case, listing can take from several days to several months. Secondary market trading starts immediately after listing. If the project is implemented successfully and more capital is needed, the startup may return to the ICO market for a seasoned offering. Figure 1 presents the timeline for a typical ICO.

[Insert Figure 1 here.]

2.2.1 A changing regulatory environment

During the past few years, ICOs and cryptocurrency exchanges have operated in a legal and regulatory grey area. The first regulatory warning came from the SEC in July 2013, in the form of an investor alert about Ponzi schemes that involved Bitcoin and other virtual currencies. Since then, the SEC has issued a series of warnings suggesting that many token sales may have violated U.S. securities laws, including a July 2017 Report of Investigation that determined that the Ethereum-based DAO tokens were securities, and offers and sales of the DAO tokens were subject to the federal securities laws. In addition to issuing dozens of subpoenas and information requests in February 2018 to technology startups involved in ICOs, the SEC has recently halted several high-profile ICO frauds, such as Centra and AriseBank. In May 2018, more than 40 state and provincial jurisdictions in the U.S. and Canada announced one of the largest coordinated series of enforcement actions to crack down on fraudulent ICOs, resulting in almost 70 open investigations and 35 pending or completed enforcement actions.

Through these regulatory actions, the SEC has made clear that (1) ICO issuers must be able to demonstrate that their tokens are not securities or that they follow securities laws, (2) market participants must ensure that their cryptocurrency activities do not undermine their anti-money laundering and "Know Your Customer" obligations, the latter of which refers to a process of identifying and verifying the identity of potential clients.

Among major economies, China appears to be the most stringent cryptocurrency regulator, banning ICOs and shutting down exchanges in September 2017. The crackdown has recently broadened to Bitcoin miners, forcing some of the industry's biggest players to shift operations overseas. In neighboring South Korea, securities officials in January 2018 disallowed anonymous accounts from trading cryptocurrencies. European Union ("EU") countries, together with Switzerland, Singapore, and Japan, have taken a relatively friendly stance toward cryptocurrency regulation. However, in April 2018 the EU approved a regulation proposed in December 2017 that requires cryptocurrency exchanges to register with authorities and apply due diligence procedures, including a Know Your Customer policy.

Due to such regulatory pressure and demand from cryptocurrency exchanges to combat money laundering, startups that launch ICOs increasingly ask their clients who participate in token sales to go through a Know Your Customer process. Many recent ICOs have also routinely prevented investors in the U.S., China and certain other countries from participating in their ICOs.

2.2.2 Two examples of ICOs

To give the reader a flavor of how an ICO actually works, we provide a description of two such cases. The first ICO illustrates features of successful ICOs, while the second highlights issues associated with a failed ICO.

The Aragon Token Sale

Founded by Luis Cuende and Jorge Izquierdo in Spain, the Aragon Network is a decentralized application built on the Ethereum Blockchain that allows users to create and manage decentralized companies. It enables users to implement basic features such as governance, fundraising, payroll and accounting, among other features. Aragon also includes a token (ticker ANT), which grants voting rights for making decisions about the direction of future development.

Aragon published a whitepaper in both English and Chinese on April 20, 2017, introducing its business model, functioning of the organization and features of the token.[6] Aragon is among the few ICOs that require a relatively long vesting period for founders, who will vest 25% of their tokens every six months after the sale (two-year vesting with six-month cliffs). Aragon is also a leading startup that publishes how it uses the funds raised, detailing each expenditure on its website, including the addresses of the company's accounts and the vendors'.[7]

On the same day, the token sale was officially announced in a blog post on Aragon's website. The sale was originally planned for four weeks, from May 17 to June 14, 2017. Aragon sought to sell 70% of tokens to investors, and accepted only ETH. In the first two weeks, one ANT token was priced at 0.01 ETH (equivalent to $0.90 on May 17, 2017), and the price would increase to 0.015 ETH per token in the remaining weeks. Aragon also implemented a hidden cap of 275,000 ETH (or roughly $25 million), which was not revealed at the time of the sale.

[6]Aragon's whitepaper is available through https://github.com/aragon/whitepaper.

[7]Each post-ICO expenditure Aragon incurs can be viewed through http://transparency.aragon.one/#/. Aragon stated that it would use the funds raised to further develop its software, implementing security audits, and hiring additional developers and operational staff.

Due to overwhelming demand, the hard cap was reached in about 30 minutes and the sale ended. There were 6,593 transactions from 2,616 unique addresses, spanning 134 Ethereum blocks. Proposed transactions valued over $8 million did not go through before the sale ended. Figure 2, Panel A plots minute-by-minute investor contributions and the cumulative contributions, which indicates that within seven minutes Aragon raised over half of the hard cap. Panel B shows the value of tokens held by top investors. The top 10% of holders purchased about 80% of sold tokens. ANT began trading the next day, May 18, 2017, with an opening price of $1.49 per token and closing price of $1.52. The closing price on May 18, 2018, one year later, was $3.99.

[Insert Figure 2 here.]

Ebitz's ICO

In November 2016, a group of self-described "ethical hackers" announced the launch of Ebitz cryptocurrency, a clone version of ZCash, the 21st largest cryptocurrency by market value. Both platforms aim to protect privacy by publishing only each transaction ID on a public blockchain, but information on the sender, recipient, and amount of the transaction remains private. Unlike ZCash, however, Ebitz did not support large rewards to the founders or the standard consensus-based mining algorithm. The Ebitz ICO went live on November 28, 2016 and would end on December 26, 2016 or when the hard cap of 500 BTC was reached.

Ebitz planned to sell 95% of the 21 million emitted tokens to participants, while allocating the remaining 5% to developers and bounty programs. The platform offered an annual interest of 3% to its token holders. The ICO accepted both BTC and ETH as valid currencies for payment. Participants who invested during the first two days were promised a 25% early-bird bonus, while it was fixed at 20% for the remainder of the week. Bonuses for the second and third weeks were 15% and 10%, respectively.

Two days after the sale started, an investor revealed on BitcoinTalk that the email server for Ebitz actually belonged to the domain of Opair, a dubious platform that promoted a decentralized debit card system using its own token. The Opair platform was shut down in the summer of 2016 after users discovered that the LinkedIn profiles of some of the team members were fabricated.

Ebitz's website was quickly removed. However, the ICO still managed to raise about 200 BTC which were valued at $156,000 at the time. There was some speculation that these BTC mostly came from the developers themselves in an attempt to start a cascade and entice outside investors to purchase their tokens.

3 Data Sources and Sample Overview

3.1 The ICO sample

Our sample of ICOs announced between January 1, 2016 and March 31, 2018 is constructed using ICObench.com.[8] ICObench is one of the oldest rating platforms on ICOs and arguably maintains the most comprehensive database on ICOs. From the website, we collect the following information: startup name, token ticker, country of incorporation, ICO status (completed, ongoing, or upcoming), start and end dates of an ICO, soft and hard caps, gross proceeds, types of currencies accepted for an ICO, bonus/discount terms, token price, the number and percent of tokens for sale, whether an ICO includes a presale, whether an ICO has a Know Your Customer policy, and headline and individual ratings.[9] Excluding seasoned token offerings, we obtain an initial sample of 2,633 ICOs, which include all completed, ongoing, and upcoming sales. We then update the ICO status variable through May 31, 2018 by checking startups' websites for ongoing and upcoming offerings as of March 31, 2018. Our final sample includes 1,549 completed ICOs.

For certain variables, most notably soft and hard caps, gross proceeds, token price, bonus terms, and the number and percent of tokens for sale, ICObench misses information for a large number of ICOs.[10] Therefore, studies solely based on ICObench's data are subject to severe sampling bias. To ensure sample completeness, we manually collect such information from the startups' official websites and their ICO whitepapers. For ICOs whose websites or whitepapers are not available, we rely on news searches on popular blogging and social networking websites such as Steemit, Medium and BitcoinTalk.org, and other credible data providers such as ICORating.com and TokenData.io.

Furthermore, we use the sources above to verify each data point that ICObench collects for these variables. In case of inconsistencies between data entries on ICObench and those from official websites and whitepapers, we use the latter.

As many ICOs seek to raise funds from investors around the globe, the startups often

[8]ICObench's revenue mostly comes from listing fees for new ICOs. An entrepreneurs can choose to publish her ICO on ICObench using regular listing or priority listing. The former is free of charge, while the latter costs 0.1 BTC (equivalent to $643.1 as of June 30, 2018). As the website has "a very long queue of pending ICOs to be listed on ICObench it might take several days" for an ICO to be listed, entrepreneurs are encouraged to apply for priority listing.

[9]Gross proceeds, soft and hard caps, token price are often quoted in ETH, BTC, another fiat currency or a combination of these currencies, we convert the figures to dollars as of the ICO end dates. Converting the figures to dollars using the ICO start dates yield similar qualitative results.

[10]For example, ICObench collected information on the number and percent of tokens for sale for just 829 ICOs, while we collected such information for another 684 ICOs. Figures on gross proceeds were available for 484 ICOs, and we added such data for another 302 ICOs.

make their websites and whitepapers available in different languages. We visit each ICO's website and record the number of languages available. If a website or its associated whitepaper is no longer available, we take a conservative approach and assume only the English language is available for that ICO.

Lastly, as ICObench assigns multiple industry codes to each ICO, we rely on ICORating's industry classification. Matching on startup name and ticker yields 982 ICOs with industry classification. For the remaining 667 ICOs, we follow ICORating's methodology and manually collect industry data from their official websites, whitepapers, and news searches.

3.1.1 The bonus schedule

To precisely measure individual investors' investment returns (Section 5.2), we manually collect from whitepapers bonus schedules for the 298 listed tokens that offered bonuses or discounts. For each listed token, we record the start and end dates for each bonus period and their associated bonus/discount. In 102 ICOs, bonuses/discounts are conditional on sales milestones. We therefore calculate the end dates for each bonus ourselves. For robustness analysis, we also obtain bonus schedules during the presale periods. Figure A1 in the Appendix shows one example for each type of bonus structure.

3.2 Analyst ratings

One unique feature of ICObench's ratings is that in addition to its assessment algorithm that uses more than 20 different criteria (nicknamed the "Benchy" rating), a large number of experts provide their own ratings on each ICO. The assessment algorithm divides evaluation into four different groups: team, ICO information, product representation, and marketing and social media. All ICOs are rated with a scale from 1 to 5. Regarding the team, the algorithm records the number of team members, photos, full names, and social media links. A team is considered more trustworthy if any member has participated in multiple ICOs, either as an advisor or team member. ICO information refers to basic variables such as token ticker, ICO start and end dates, soft and hard caps, among others. For product presentation, the algorithm checks the availability of such information through whitepapers, milestones, and video presentations. The algorithm also monitors activity on various social networks to determine whether an ICO team reaches potential investors.

Independently, each expert assigns a rating from 1 to 5 to an ICO for team, vision, and product. Analysts disclose their real names and bigraphic information such as job titles and employers. Like the assessment algorithm, analysts consider a team strong and trustworthy

if it keeps the community updated with project progress and/or has participated in other cryptocurrency-related projects. Vision mostly concerns what a platform aims to achieve in the mid-term and long-term future. When evaluating products, analysts consider the following aspects: (1) product maturity level – working products are easier to evaluate than concepts, (2) technology, both blockchain- and non-blockchain-related, (3) specific problems with their products/services, (4) product roadmap that shows short-term and long-term strategies and growth, and (5) projects' commitment to understanding the market environment.

The headline rating is a weighted average of all participants' ratings based on their expertise, years of experience in the field, and available publications.[11]

ICObench assigns an ICO Success Score to each analyst, which is the average score for all the ICOs they have rated. An analyst can receive up to 15 points for each ICO. Crucially, analysts receive up to 10 points if a token gets listed, which is an important milestone for ICO success. ICObench explains analysts can improve their success score by identifying successful ICOs.[12]

3.3 Primary market subscription data

One unique feature of any cryptocurrency is that each transaction needs to be broadcast to all participants on a blockchain network before it is validated. For example, each transaction involving ETH will be sent to every Ethereum node (a computer with an installed Ethereum program) on the Internet. Miners, a group of competitive bookkeepers, will validate the transaction by completing a cryptographic "proof of work," which involves a cryptographic hash function that takes an input and delivers an output. The purpose of the "proof of work" is to make sure that the transfer is genuine and there is no double spending or counterfeiting. If the majority of participants validate the transaction, it will be added to a "block" on the Ethereum public ledger, which is a decentralized database containing the entire history of every ETH transaction. All transactions on any Ethereum block are viewable by the public.

The same procedure applies to any ICO, where each token transfer/subscription during the primary market sale is recorded on the public ledger associated with the underlying blockchain platform. As 78% of our sample projects use the Ethereum platform to run their ICOs, we collect primary market subscription data on all Ethereum-based ICOs from

[11]The reader is referred to https://icobench.com/ratings for a more detailed description of ICObench's rating methodology.

[12]Details about the success score can be accessed via https://icobench.com/faqq-6-4.

Etherscan, a leading "block explorer" that allows users to search for information about blocks and transactions on the Ethereum Blockchain. We manually search each Ethereum-based ICO's name and if there is a match, we record the ICO's contract hash address, a 42-character string. For each contract address, we download information for all transactions taking place between the ICO start and end dates. These include the transaction address, sender address, receiver address (all of which are 66-character strings), transaction time (e.g, May-20-2018 04:16:49 PM +UTC), quantity of tokens transferred, total dollar value of the transfer. We initially identify 952 ICOs that have primary market transactions available.

Our next step is to identify the addresses for ICO insiders and primary market investors. In most cases, it is straightforward to identify the insiders as all transfers are originated from one single address. The rest are primary market subscribers. However, when it is difficult to cleanly identify the insiders because multiple addresses are used to transfer tokens to investors, we take a conservative approach to exclude these ICOs in our main analysis. To facilitate the analysis on sequential investor subscriptions, we also drop the ICOs that distribute tokens *after* the token sales end. These criteria yield a transactions sample of 654 ICOs.

Our paper is unique in its reliance upon primary market subscription data, and only this allows for an empirical analysis of investor subscription dynamics during ICOs.

3.4 Secondary market prices, volumes, and transactions

Following a successful ICO, a token is expected to be listed on an exchange or several exchanges simultaneously. We collect its daily closing prices from CoinMarketCap.com, a website that is a top source for pricing data on hundreds of cryptocurrencies. For each token, CoinMarketCap aggregates pricing information from all major exchanges and produces one standard price quote in real time. It also publishes the 24-hour trading volume among major exchanges. We manually search each successful ICO's name on CoinMarketCap, and download its dollar price series, daily trading volumes, and circulating supply of tokens if available.[13] Our sample consists of 433 tokens that were listed following an ICO, and that had offer prices available.

To calculate individual token investors' investment returns, we collect from Etherscan secondary market sales data for each investor who purchased tokens in the primary market.

[13]To verify whether information from CoinMarketCap is accurate, we also download pricing and volume data from popular alternative pricing sites, such as Onchainfx.com, CryptoCompare.com, and CoinGecko.com. We find that for the vast majority of cryptocurrencies, the prices and volumes from CoinMarketCap are highly correlated with those from the alternative sites (the correlation is typically above 0.9).

Etherscan records both primary and secondary market transactions.

3.5 Sample overview

Figure 3 plots the frequency of ICO starts and the rate of fundraising success over our sample period. As in Mironov and Campbell (2018), we consider an ICO successful if its soft cap was reached or the project raised more than $0.5 million in the absence of a soft cap.[14] Just 19 ICOs opened in 2016. However, the market took off in 2017, enjoying a quarterly increase of 171% in 2017. The number of ICO starts continued to rise through the first quarter of 2018. The fundraising success rate was over 90% in the first half of 2017, and then dropped sharply in the second half of the year. The deteriorating success rate potentially reflects the "lemon's problem" that lower quality ICOs were created when the market became red hot (e.g., Akerlof, 1970). The decline in success rates also coincided with increasing regulatory scrutiny worldwide.

[Insert Figure 3 here.]

The top five largest ICOs to date are Dragon Coin, Huobi, HADC, Filecoin, and Tezos, all of which closed between July 2017 and March 2018.[15] Combined, these token sales raised nearly $1.4 billion, accounting for 15% of all ICO proceeds raised globally during the same period, according to EY Research. Appendix Table A2 shows the top 10 largest ICOs as of May 2018, with information on their fundraising periods and gross proceeds.

Table 1 reports the frequency of sales and fundraising success rate for ICOs from each of the top 10 countries and industries. As shown in Panel A, the U.S. is the most popular country for blockchain startups, followed by Russia. Interestingly, ICOs from Switzerland enjoy the highest fundraising success rate, assisted by "crypto-friendly" guidelines recently issued by Swiss regulators (Atkins, 2018). However, Russia-based ICOs are least likely to succeed, followed by ICOs from Canada.

[Insert Table 1 here.]

[14] Instead of using this industry convention to define ICO success, Amsden and Schweizer (2018) and Momtaz (2018) use exchange listing as the criterion for ICO success. After a successful fundraiser, it can take a startup several months to list their token on an exchange. Some entrepreneurs may choose not to list their tokens. Since ICO is a recent phenomenon, both Amsden and Schweizer (2018) and Momtaz (2018) likely miss a substantial number of successful ICOs. In our sample, only about 60% of successful ICOs were listed as of May 31, 2018.

[15] We do not count Telegram's record-breaking $1.7 billion token sale in the first quarter of 2018 because it was structured as a private sale.

Financial Services is the most popular industry for ICOs, attracting over 10% of all completed token sales. This is perhaps not surprising given that the original Bitcoin Blockchain was developed to replace the traditional centralized financial system. The average fundraising success rate for the top 10 industries is about 56%, substantially greater than the sample average of 45%. ICOs in the Blockchain Infrastructure industry achieve the highest success rate.

3.5.1 Analyst ratings

Our sample comprises 9,493 distinct ratings by 316 unique online analysts. Their average (median) ratings for team, vision, and product are 3.9 (4.0), 3.8 (4), and 3.6 (4.0), respectively. Of these analysts, 165 are founders or senior managers from blockchain-related companies, 96 are advisors to these firms, and 73 are blockchain analysts or followers. Some analysts are also investors, with 58 being cryptocurrency/blockchain investors and 16 being venture capitalists or angel investors. Notably, there are also 52 engineers and technicians, who are potentially able to provide valuable insights into the technical aspects of ICOs. Other analysts are from fields such as finance/business (39), marketing (30), consulting (21), among others. Note that these categories are not mutually exclusive as a given analyst can play multiple roles.

3.5.2 Patterns in primary market subscriptions

In a bookbuilding process commonly used in IPOs, the underwriter solicits investors' bids, which are used to construct a demand curve and allocate shares to the investors. However, in ICOs, no underwriter is building a book for the sale. Rather, the price and offering period are set ex ante by entrepreneurs, and the gross proceeds equals investors' cumulative subscriptions by the end of the offering period. To understand the path to fundraising success, we resort to our unique second-by-second subscription data, which are aggregated at the daily frequency.

Figure 4, Panel A plots the time series patterns of daily token sales for both successful and failed ICOs. The blue bars (line) represent daily (cumulative) token sales as a percentage of total token supply for successful ICOs, while the orange bars and line represent the corresponding figures for failed sales. In successful ICOs, investors purchase nearly 15% of token supply on the first day, while the 30-day cumulative demand is about 30% of the total supply. In contrast, in failed sales, investors on the first day buy fewer than 1% of all tokens for sale, and the cumulative sales are just about 2.5% of token supply.

The plot in Panel B further shows hourly token sales on the first day of a sale for successful and failed campaigns. In the first hour alone, successful ICOs sell about 7% of tokens offered, while failed campaigns sell few coins during this period. Volumes in successful sales go down as time passes. Panel C repeats the analysis for each block during the first hour of token sales. A block contains transactions within a 14-second interval. The time-series pattern is highly consistent with those in the daily and hourly plots. These patterns highlight the importance of "winning the battle" during the initial periods of an ICO, which often determines the outcome of the sale. Such a concave pattern of cumulative token sales is consistent with an information cascade mechanism proposed by Cong and Xiao (2018), Li and Mann (2018), and Welch (1992).

[Insert Figure 4 here.]

4 Fundraising Success and Campaign Duration

4.1 Successful fundraisers versus failed ICOs

Our first analysis examines the characteristics of successful ICO fundraising campaigns. Column (1) in Table 2 reports the attributes of ICOs that successfully raised funds, while column (2) compares the characteristics between successful and failed ICOs.

[Insert Table 2 here.]

Regarding ex ante ICO characteristics, most importantly, successful ICOs on average had a rating of 3.3 (out of 5) by online experts, 0.7 points higher than that for failed token sales. The difference is significant (at the 1% level), suggesting that analyst certification before an ICO goes live is an important predictor of fundraising success. In the absence of traditional underwriters who play a critical intermediary role in the IPO market, analysts fill the void and potentially help reduce information asymmetry in ICOs. These experts are likely unbiased as biased ratings may be uncovered in the long run, resulting in reputational damages. This could be a novel market solution for token sales, all of which feature decentralized fundraising platforms through blockchain technology. Successful token sales also attracted more analysts to initiate coverage, with the number doubling that for failed ones. This is indicative of the "wisdom of crowds" phenomenon observed in the ICO market. The pattern based on the medians is qualitatively similar.

The average soft cap or minimum funding goal for successful ICOs was $6.8 million, nearly identical to that set by unsuccessful fundraising campaigns (the difference is insignificant at

the 10% level). The hard cap or maximum goal for a successful ICO on average was $88 million, more than double the amount for failed ones (the difference is not significant at the 10% level). The median amounts tell a more consistent story, as both groups of ICOs have similar median soft and hard caps.

As an important governance indicator, the percent of tokens to be sold to investors measures management's skin in the firm. Successful ICOs sought to sell 57% of generated tokens to outsiders, compared to the target of over 61% in failed ones, with the difference being significant at the 1% level. This is a strong indication that investors embrace token sales in which management retains more stakes. Note that without formal governance and incentive mechanisms post-ICO, such as voting power to oust directors and performance compensation package commonly used after an IPO, management's stakes could play a vital governance role to align insiders' and outsiders' interests.[16] This is also consistent with Leland and Pyle's (1977) signaling hypothesis that firm value is positively related to the fraction of equity retained by the original stockholders.

Nearly 40% of successful ICOs included a presale before the main token sale, 18.2% higher than failed ICOs. Presales typically are open to only institutional or high-net wealth investors, and the proceeds raised are often used to cover the costs of launching the main ICOs. Successful initial sales are interpreted by subsequent investors as evidence that earlier investors held favorable information, encouraging later investors to invest regardless of their own information. We posit that in the absence of underwriters, a presale is a clever way for insiders to gauge demand from informed investors such that they can set a more informed price for the main ICO. Given the market power these early investors enjoy, it is not surprising that presales often provide a steeper discount than the main sale. This is analogous to the analysis of informed IPO investors by Biais, Bossaerts, and Rochet (2002) and Benveniste and Spindt (1989).

Interestingly and perhaps counterintuitively, high bonus offers, defined as 20% or more, are more prevalent in failed ICOs (the difference is significant at the 10% level).[17] Although generous bonuses can attract investor subscriptions in early periods of an ICO, many of these token sales provide extremely high bonuses that sometimes exceed 100%. Wary investors may conclude that such ICOs are potential lemons or scams, and are reluctant to extend credit to the entrepreneurs (Stiglitz and Weiss, 1981).

[16] Distinct from most IPOs in which management loses majority control of the firms (except some high-tech IPOs such as Facebook and Snap where management controls voting rights via dual class shares), in token sales management does not give away control as voting rights typically are not attached to tokens.

[17] Sagar (2017) considers ICO bonuses on offer exceed 20% as a red flag. Using an alternative threshold of 30% yields consistent results in our main analysis.

ICOs that adopted a Know Your Customer policy were less likely to conclude successfully. Only 12% of successful token sales asked for customer identification, compared to nearly 17% for failed ICOs (the difference is significant at the 1% level). Similarly, ICOs that required advance registration or restricted sales in certain countries (*Participation restriction*) were also less likely to succeed. This suggests that Know Your Customer and related restrictive measures adopted by many ICOs since Q3 2017 tend to dampen investor demand, the bulk of which likely comes from countries like China, South Korea, and the U.S.

Lastly, token sales featuring multi-language websites or whitepapers tended to be more successful, reflecting the fact that potential token purchasers are not based in a single country and language barriers exist. ICOs that accepted multiple (digital) currencies were more likely to succeed (the difference is significant at the 10% level), compared to ICOs that took just one currency. Given that major digital currencies such as BTC and ETH are drastically volatile, expanding the options of currencies, thus increased liquidity, can facilitate transactions.

Panel B compares key ex post ICO outcomes between successful and failed ICOs. On average, successful sales raised $18.7 million, far more than the $2.7 million failed ones raised, with the difference being significant at the 1%. In contrast, according to the Crowdfunding Center, successful crowdfunding campaigns on average raised just $29,900 in 2016, a tiny fraction of the amount raised in ICOs. Specialized crowdfunding platforms, such as Kickstarter, support even smaller fundraising campaigns (Xu, 2017). On the other hand, tech IPOs in 2017 grossed over $250 million on average (Ritter, 2018). Considering that most blockchain-related projects raise funds before actual launches and the companies are much smaller than IPO firms, ICOs have become an increasingly important source of alternative fundraising.

Perhaps more tellingly, successful ICOs on average achieved 59% of the hard cap, while unsuccessful ones obtained just 15.4% of the hard cap. Successful ICOs had more than 2,100 supporters on average, compared to the 39 subscribers in failed ICOs. Such a stark difference highlights the role of information cascades that can potentially lead to lopsided fundraising results.

Interestingly, successful ICOs on average took 30.0 days to complete, shorter than the 37.8 days failed fundraisers took (the difference is significant at the 1%). This is because although most ICOs are scheduled for about one month, successful ICOs often finish early (see the Aragon token sale in Section 2), usually at the time when they hit the hard cap.

4.2 Determinants of ICO success

Table 3, Panel A reports the results of predictive regressions where the dependent variable is ICO fundraising success, which equals one if an ICO reaches its soft cap or the project raises more than $0.5 million in the absence of a soft cap (Mironov and Campbell, 2018).[18] The set of regressors are the same as those presented in Table 2 with the critical difference that all variables in the regressions are measured at the time of an ICO announcement. The sample includes all 1,461 ICOs for which we have the required information. Column (1) displays the probit coefficients and their associated marginal probabilities. Column (2) reports coefficients from a linear probability model with country and quarter fixed effects. Since the results in the two columns are qualitatively similar, we rely on the probit model for our discussion.

[Insert Table 3 here.]

Consistent with results shown in Table 2, the average analyst rating, all else being equal, has a significantly positive effect (at the 1% level) on the likelihood of a successful fundraising campaign. A one-standard-deviation increase in the average rating is associated with an increase in the marginal probability of 19.8 percentage points. Relative to the unconditional probability of ICO success of 42.7%, the incremental probability is remarkable. This finding is consistent with the positive intermediary role these experts play in a market where traditional underwriters are absent. In robustness analysis, we exclude experts who are founders of other ICOs, and therefore are potentially biased, and the results are qualitatively similar. The number of analysts covering an ICO also positively predicts fundraising success, suggesting a "wisdom of crowds" phenomenon in which investors tend to follow a crowd of analysts when making investment decisions.

The coefficients associated with three more ICO characteristics support their governance or signaling roles. For a one-standard-deviation increase in the fraction of tokens for sale, there is a 6.1 percentage point decrease in the marginal probability of ICO success (significant at the 1% level). This suggests that investors favor ICOs in which management's and investors' incentives are more aligned through higher inside stakes (Leland and Pyle, 1977; Downes and Heinkel, 1984). On the other hand, token sales providing large bonuses or discounts are 10.9 percentage point less likely to successfully conclude the fundraising effort, reflecting credit rationing to ICOs with overly generous bonuses, many of which are believed to be potential scams (Sagar, 2017). Including a presale can boost the success likelihood by

[18]For ICOs that do not specify a soft cap, changing the target to $0.25 million or $0.75 million yields similar results.

15.2 percentage points (significant at the 1% level), suggesting that successful initial sales can promote the subsequent token sales by harnessing the widsom of crowds (Cong and Xiao, 2018; Li and Mann, 2018).

The probability of ICO success decreases by 5.7% percentage points (significant at the 1% level) when a token sale requires customer identification, a process that can deter some potential customers. Replacing this variable with *Whitelist*, an indicator equal to one if customers are required to register in advance for a sale, we obtain similar results (see Table A3 in the Appendix). Consistent with findings in Table 2, token sales that feature multi-language websites/whitepapers or accept multiple currencies are more likely to succeed, reflecting the global nature of token sales and the ease of transactions with an expanded set of currencies.

4.3 Gross proceeds and ICO success

Our analysis on fundraising success in subsection 4.1 ignores the degree of success, which can be measured by gross proceeds or gross proceeds as a percentage of the specified hard cap. In Table 3, Panel B, we repeat the same analysis by using these two alternative dependent variables. Our sample becomes smaller as this analysis requires that information on gross proceeds and/or hard cap is available.

As shown in column (1), gross proceeds increases by $4.7 million when the average analyst rating increases by one point. This is substantial given that the sample average gross proceeds is about $15.2 million. Consistent with results in subsection 4.1, the fraction of tokens for sale and the availability of generous bonuses negatively predict the total amount raised in an ICO. The other covariates are not statistically significant at the 10% level, partly due to the smaller sample employed in this study.

Results reported in column (2) are largely consistent with those in column (1), with the coefficients on the number of analysts and Know Your Customer being statistically more significant.

4.4 ICO duration

Campaign duration is an alternative measure of ICO success in that highly successful ICOs often conclude at the time when the hard cap is reached. A longer fundraising campaign distracts management who need to focus on product development after an ICO. To assess the extent of such a cost, Table 4 reports results connecting ICO duration to key metrics identical to those shown in Table 3. In column (1), the dependent variable is the logarithm

of the number of days between an ICO's start date and the completion of the sale.

[Insert Table 4 here.]

As expected, a favorable analyst rating is associated with a quicker sale. Interestingly, ICOs that feature multi-language websites or whitepapers or accept multiple currencies take a longer time to consummate. This may be explained by the fact that such ICOs mainly target retail investors. In Li and Mann (2018), a fundraising process is formalized as a multi-stage game where heterogeneous investors with private signals of a project's quality decide whether and when to participate in the ICO. In equilibrium, investors with stronger signals participate early and those with weaker signals "follow the crowd." Assuming that retail investors generally possess weaker signals, they are more likely to participate in later stages of the sale. Therefore, if the entrepreneurs aim to achieve decentralized ownership, they can target a retail investor base by making information on the token sale available in multiple languages and accepting multiple currencies. These measures could potentially prolong the fundraising process. The same logic applies to ICOs that require a Know Your Customer policy.

The Cox (1972) proportional hazards model,[19] reported in column (2), yields qualitatively similar results. The estimated hazard ratio (equal to the exponentiated coefficient) associated with the dummy variable *Presale* implies that, conditional on an ICO being in process, the probability of a sale closure on a given day is 25.8% higher if the ICO has a presale. The coefficient estimates on other covariates are largely consistent with those in the OLS model.

5 Investor Subscriptions and the Path to ICO Success

5.1 Initial and subsequent subscriptions

Cong and Xiao (2018) show that the all-or-nothing mechanism used by crowdsales leads to uni-directional information cascades in which investors rationally ignore private signals and imitate preceding agents only if enough preceding investors decide to support the ICO. Li and Mann (2018) also feature such an information cascade in a simpler setting. We test whether this mechanism may exist in ICOs using a regression framework. Consistent with this prediction, in Table 5, we find that investor subscriptions on the first day strongly predict

[19] In the Cox model, the hazard function at a given time t (from initiation), conditional on the failure to complete an ICO, is characterized as $h_i(t) = h_0(t)e^{X_i\beta}$ where $h_0(t)$ is an unspecified (or nonparametric) function.

token sales on the second day, during the next four days or the next 14 days. Notably, other covariates only weakly predict subsequent token sales.

[Insert Table 5 here.]

As theories on information cascades emphasize an all-or-nothing threshold, we replace first-day subscription with an indicator equal to one if total subscriptions on the first day are greater than 5% of total supply. The results are qualitatively similar. Using the hourly or block-level frequency also yields consistent results. Overall, the evidence presented in this subsection is consistent with the "wisdom of crowds" phenomenon, a unique feature of fundraising in the era of FinTech.

5.2 Initial investors' investment returns

As early investors appear to harness the wisdom of crowds during the primary sale, it is natural to examine whether these investors achieve a higher return than later participants. Moreover, any investment success among this group of investors can justify their incentives in investing with the risky startups.

Following Frazzini (2006) and Brav, Jiang, and Li (2018), we assume that investors use the purchase price as the base for calculating gains and losses when a token is purchased multiple times during a primary sale, and partially sold in the secondary market on different dates. Investors use a cost-based mental accounting method (first in, first out) to compute the quantity of tokens for the corresponding reference price. For any investor on any date t, the aggregate cost basis equals

$$Basis_t = \lambda^{-1} \sum_{n=0}^{t} Tokens_{t,t-n} Price_{t-n},$$

in which $Tokens_{t,t-n}$ is the number of tokens an investor acquires on date $t-n$ that she still holds on date t, λ is a normalizing constant such that $\lambda = \sum_{n=0}^{t} Tokens_{t,t-n}$. $Price_{t-n}$ is the token price on $t-n$, taking into account discounts/bonuses offered at that time. The beginning date is the first day of primary sale.

Our main return measure, an investor's basis-adjusted return, is then computed as the percentage deviation of her aggregate cost basis from the current secondary market price. We calculate basis-adjusted returns at the end of the first day, first month, three months, and six months, respectively.

To measure the timing of investor subscriptions, we use the difference between the date of purchase and ICO open time. When an investor purchases a token on several different

dates, we value weight the purchase time using the number of tokens bought. For example, an investor buys 100 tokens three days after an ICO opens, and purchases another 200 tokens six days after the open time. Her effective purchase time is five days after the ICO opens ($3 \times \frac{100}{300} + 6 \times \frac{200}{300} = 5$).

Table 6, Panel A reports results connecting investors' basis-adjusted returns to effective time of token purchase. ICO-level fixed effects control for any time-invariant unobservable factors involved with each ICO. As shown in column (1), buying tokens one day earlier is associated with a 0.3 percentage-point increase in basis-adjusted returns at the end of first day after tokens are listed, all else being equal. This relationship might be mechanical as earlier investors, who are supposed to sell out on the first day, enjoy more bonuses (the adjusted R-squared is abnormally high at 0.93). However, we obtain similar results when examining basis-adjusted returns by the end of the first month, three months, and six months. Investors have much more flexibility in timing sales when the holding horizon extends to six months. Thus, our results suggest that early investors possess at least some superior investment skills.[20] The model fit also becomes more reasonable.

[Insert Table 6 here.]

Basis-adjusted returns are not defined if investors have completely sold out their holdings by time t (Frazzini, 2006). In our sample, 4.0% of investors sell out their positions on the first day in the secondary market, while the fraction increases to 22.7% in six months. For these investors, we create another variable, *Realized return*, which is defined below.

$$\textit{Realized return} = \frac{\textit{Total proceeds in the second market}}{\textit{Total primary market investment}} - 1,$$

In Panel B, we repeat the analysis in Panel A by replacing the dependent variable as *Realized return*. The results in column (2) show that subscribing tokens one day earlier is associated with a 0.4 percentage-point increase in realized returns during the first month of secondary market trading (significant at the 5% level). We obtain similar results when extending the window to three months or six months.

Overall, our findings suggest that early token investors achieve better investment returns than their followers. This ability of earning superior returns corroborates their role in leading the crowd of subsequent investors, and justifies their participation decisions in early stages of an ICO.

[20] In unreported results, we find that when the discount an investor receives increases by one percentage point, basis-adjusted returns increase by 1.4 to 4.0 percentage points, depending on the investment horizon. This confirms that early investors' superior returns are not entirely due to higher bonuses.

5.3 Predicting primary market subscriptions

To examine what factors may affect primary market subscriptions, we rely on the following empirical specification:

$$Subscription_{ijt} = \alpha + \beta Rating_{ijt} + \delta X_{ijt} + \lambda_t + \mu_j + \epsilon_{ijt} \quad (1)$$

in which $Subscription_{ijt}$ represents measures that gauge primary market subscriptions for token sale i that starts in quarter t in country j. $Rating_{ijt}$ is the average analyst rating. X_{ijt} is a vector representing ICO-level covariates. λ_t represent quarter fixed effects and μ_j are country fixed effects. ϵ_{ijt} is the error term. Standard errors are adjusted for heteroscedasticity and they are clustered along the quarter dimension.

In column (1) of Table 7, we examine the factors that predict first-day investor subscriptions as a fraction of token supply. The average rating from experts is a strong predictor of first-day subscriptions. A one-standard-deviation increase in the average rating is associated with an increase in first-day tokens sold of 2.3 percentage points (significant at the 1% level). Compared to the sample average first-day tokens sold of 3.9%, the incremental change is remarkable. This suggests that positive analyst ratings help harness demand in the absence of reputable underwriters in this decentralized market. On the other hand, generous bonuses are negatively correlated with first-day sales volumes, suggesting that the investor crowd regard such ICOs as potential frauds and thus are unwilling to extend credit to them.

Figure 4 suggests that due to higher demand in early periods of an ICO, the curve for cumulative daily token sales is concave. That is, tokens are sold faster in earlier periods. In column (2) of Table 7, we replace the dependent variable with a measure that captures the concavity of the cumulative demand curve. Our intuitive measure equals the cumulative token sales on the 15th day minus (1/2)×the cumulative token sales on the 30th day, with the latter date being the last day of an average ICO. In the benchmark of steady daily fundraising volumes, this concavity measure equals zero. However, it is strictly positive (negative) for a hot (cold) ICO. See Figure A2 in the Appendix for more details on how we construct this measure. Consistent with the results in column (1), a better analyst rating predicts faster token sales in earlier periods, all else being equal. The estimate is significant at the 1% level. ICOs with presales also see faster token sales early on, suggesting that successful initial sales can harness information cascades in a main sale. On the other hand, a Know Your Customer policy is found to slow down token sales, potentially causing delays for subscriptions by some customers.

An alternative measure for investor subscriptions during the ICO period is the number

of days until 10% of token supply is sold, which approximates the soft cap for many token sales. As a number of ICOs never sold 10% of all the tokens to investors, there exists right-censoring. To overcome this issue, we apply a Tobit model with an upper limit of 30 days, which is approximately the average duration for ICOs. Using a 60-day upper limit yields similar results. As shown in column (3), the average rating and the number of analysts negatively predict the time to sell 10% of token supply. Consistent with the results in column (2), ICOs with presales take a shorter time to sell 10% of their token supply.

5.4 The effect of early subscriptions on fundraising success

To examine how early investor subscriptions affect ICO success, we first use a univariate probit model that relates the first-day token subscription with ICO success. The results in column (1) of Table 8, Panel A indicate that a half-standard-deviation increase in first-day subscription is associated with an increase in the marginal probability of 57 percentage points. The order of magnitude is large as the unconditional probability of ICO success is 42.7%. To control for other covariates, we use equation (1) to obtain the residual for *first-day subscription*, which by definition is orthogonal to all the covariates. Then we include this residual in our main regression equation for ICO success. Results in column (2) confirms that higher first-day subscriptions strongly predict a higher likelihood of success.

[Insert Table 8 here.]

Replacing the success indicator with gross proceeds, we obtain qualitatively similar results, as shown in Panel B. Overall, results in this section shed light on an important question: how do ICO fundraisers become successful? Using novel data on each investor's subscription, we are the first to show that early contributes (i.e, the first-day subscriptions) are crucial to eventual successes of ICOs, possibly attributed to information cascade (Cong and Xiao, 2018; Li and Mann, 2018).

6 Post-ICO Performance

Given that ICO fundraising is just the first step for a successful blockchain-based project, it is crucial to analyze how major variables related to governance and certification mechanisms affect the short-run and longer-run token performance.

6.1 ICO Underpricing

IPO underpricing, a phenomenon that the share price jumps on the first day of trading relative to the offer price, has interested financial economists for decades (See Ritter and Welch (2002) and Ljungqvist (2007) for a review of theories and empirical findings on IPO underpricing). As explained in Section 2, although the market structure of ICOs is quite different from that of IPOs, it is useful to examine whether underpricing also exists in the ICO market.

[Insert Table 9 here.]

First-day return or underpricing equals $P_1/P_{ICO} - 1$, in which P_1 and P_{ICO} are first-day closing price and offer price, respectively. As shown in Table 9, the average ICO underpricing during our sample period is 158%, substantially higher than the average IPO underpricing of 18% for the U.S. (Ritter, 2018). However, the median ICO underpricing is only 24%, indicating that the distribution of underpricing is highly skewed to the right by very large first-day pops. Consequently, money left on the table is also highly right-skewed – the average figure is $42 million while the median is just about $1.3 million.

Since the average time from ICO completion to listing is 18.5 days, we adjust first-day returns by the returns on the value-weighted benchmark index of ETH and BTC during this "wait period." The average and median adjusted underpricing figures are slightly below the raw numbers. It is worth noting that to calculate first-day token returns, we rely on manual collection of primary market *offer prices* as the first-day return is measured from the token offer price to the first trading day closing price. This differentiates our study from Amsden and Schweizer (2018) and Momtaz (2018), the latter of which features a first-day return measured from the first trading day opening price to closing price.

First-day turnover, defined as the first-day trading volume divided by the number of tokens issued, has an average (median) value of 6.7% (1.3%). It is significantly lower than the average first-day turnover of 66.2% for IPOs between 2001 and 2016 (Ritter, 2018). This highlights the illiquid nature of many listed tokens.

6.1.1 Cross-sectional patterns in underpricing

To understand which factors may explain ICO underpricing, we start by regressing the first-day return on the wait-period benchmark return and the same covariates as in Table 3. As shown in column (1) of Table 10, a one-percent increase in the wait-period benchmark return is associated with a 2.1 percentage point increase in first-day returns, all else being equal

(the coefficient is significant at the 1% level). This suggests that a hot overall cryptocurrency market substantially boosts investors' sentiment in these newly listed "alt-coins."

[Insert Table 10 here.]

On the other hand, ICOs featuring generous bonuses have a significantly lower first-day return than that for token sales without large bonuses. This is consistent with Habib and Ljungqvist (2001), who find that promotional/marketing costs of an IPO can substitute for IPO underpricing. ICO bonuses can be viewed as marketing tools to attract early investors in the primary market. Knowing that early investors bought tokens at large discounts, rational secondary market investors would demand a lower price, which is closer to the intrinsic value these early investors' tokens represent.

Interestingly, analyst ratings do not appear to be correlated with underpricing. Any additional demand spurred by a good rating is likely to be fully absorbed in the primary market subscription process. This is reasonable as the majority of ICOs do not hit the hard cap. Another plausible reason is that unlike investment banks, who likely have incentives to underprice IPO stocks (Ritter and Welch, 2002; Ljungqvist, 2007), online analysts do not have financial stakes in startups undergoing an ICO and do not control the pricing process.

6.2 Long-run Performance

The median token return during the six month period after the first trading day significantly underperforms the benchmark of ETH and BTC by 37 percentage points, as shown in Table 9 (it is different from zero at the 1% level). However, the average six-month excess return is highly positive at 66%, reflecting the skewed nature of the underlying return distribution. To reduce risk mismeasurement, we use the top 10 digital currencies as an alternative benchmark, and obtain similar results.

The median one-year excess return is -125% (significant at the 1% level), while the average one-year excess return is not statistically significant from zero. This is consistent with Ritter (2018), who finds that in general, the longer is the horizon, the lower is the median return, due to the right-skewness of buy-and-hold returns.

It is worth noting that since ICOs are a recent phenomenon, many listed tokens still do not have a history of one year or more. Therefore, we have fewer observations for one-year returns than the shorter-term returns. We plan to analyze longer-term returns in the future.

6.2.1 Cross-sectional patterns in long-run performance

To investigate possible explanations for the long-run performance of ICOs, this subsection documents various cross-sectional patterns. Columns (2)-(5) in Table 10 report the results of a multiple regression using the raw 1-month, 3-month, 6-month, and 1-year total returns on ICOs as the dependent variable. The explanatory variables are the market-adjusted first-day return, the corresponding total return on the cryptocurrency market, the relevant covariates, and country fixed effects. This model is similar to Ritter (1991), who uses it to study long-run underperformance of IPOs. Because the dependent variables, 1-month, 3-month, 6-month, and 1-year total returns, are so skewed, the residuals are also highly non-normal. Consequently, bootstrapped t-statistics are reported.

For all the horizons, long-run returns are inversely related to initial returns, exhibiting potential mean reversions. The coefficient on the market return of 0.91-1.16 (except the 2.02 for the six month window) is somewhat surprising, as many would expect that the average alt-coin betas would be substantially above 1.

Most importantly, we find that a good analyst rating positively predicts both 3-month and 1-year token returns (the estimates are significant the 10% and 5% levels, respectively), but it is not associated with shorter-term returns. Analyst ratings focus on team, vision and product, all of which are long-term indicators of the startup quality. This implies that analyst ratings are informative, even beyond the primary market stage. It suggests that analyst ratings are likely unbiased as biased ratings may be uncovered in the long run, resulting in a reputational cost. This result is also consistent with a key finding in Jia, Ritter, Xie, and Zhang (2018), who use a sample of China's IPOs, and find that analyst coverage and earnings forecast optimism for an IPO before it starts is positively associated with IPO long-run returns.

Lastly, Know Your Customer policies have a dampening effect on 1-month and 3-month returns. This is not unexpected, as such policies potentially curb demand from wary investors who do not want to reveal their identity. [21]

7 Conclusion

This is the first study that uses token subscription data to examine how the wisdom of crowds can help overcome information asymmetry in ICOs. We find that favorable ratings on the underlying project from a crowd of online experts are associated with aggressive

[21] We believe that not only illegal demand is curbed but true believers in decentralized blockchain may find this policy counterproductive or intrusive.

initial token subscriptions by investors, which in turn predicts subsequent token sales. Using blockchain-based subscription data, we find that successful token sales attract more than 2,000 investors on average, compared to the 39 investors in failed sales. Such "up-cascades" of wisdom of crowds is the key to successful ICO fundraising (Cong and Xiao, 2018). Analyst opinions also predict long-run token performance in the secondary market. This suggests that the wisdom of the two crowds – analysts and early token investors – can potentially substitute traditional underwriters' intermediary roles in financing blockchain-related ventures.

To prevent the "insanity of crowds," ICO analysts are required to be independent and possess a diversity of opinions (Sehra, Smith, and Gomes, 2017). Potential backdoor networks need be further investigated to ensure unbiasedness. How to set an optimal all-or-nothing threshold to ensure fundraising success without suffering from short-run underpricing is another important mechanism design problem. We aim to address these questions in future research.

References

Akerlof, George A., 1970, The Market for "Lemons": Quality Uncertainty and the Market Mechanism, *Quarterly Journal of Economics* 84(3), 488-500.

Amsden, Ryan, and Denis Schweizer, 2018, Are Blockchain Crowdsales the New "Gold Rush"? Success Determinants of Initial Coin Offerings, Working paper, Concordia University.

Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia, 2016, Bitcoin pricing, adoption, and usage: Theory and evidence, Working paper, Stanford University.

Atkins, Ralph, 2018, Switzerland Sets Out Guidelines to Support Initial Coin Offerings, Financial Times, February 16, 2018.

Beatty, Randolph P., and Jay R. Ritter, 1986, Investment Banking, Reputation, and the Underpricing OF Initial Public Offerings, *Journal of Financial Economics* 15, 213-232.

Benveniste, Lawrence M., and Paul A. Spindt, 1989, How Investment Bankers Determine the Offer Price and Allocation of New Issues, *Journal of Financial Economics* 24, 343-362.

Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta, 2018, The Blockchain Folk Theorem, Working paper, Toulouse School of Economics.

Biais, Bruno, Peter Bossaerts, and Jean-Charles Rochet, 2002, An Optimal IPO Mechanism, *Review of Economic Studies* 69, 117-146.

Brav, Alon, Wei Jiang, and Tao Li, 2018, Picking friends before picking (proxy) fights: How mutual fund voting shapes proxy contests, Working paper, Duke University, Columbia University and University of Florida.

Canidio, Andrea, 2018, Financial Incentives for Open Source Development: the case of Blockchain, Working paper, INSEAD.

Catalini, Christian, and Joshua S. Gans, 2018, Initial Coin Offerings and the Value of Crypto Tokens, Discussion paper, National Bureau of Economic Research.

Ciaian, Pavel, Miroslava Rajcaniova, and d'Artis Kancs, 2016, The Economics of Bitcoin Price Formation, *Applied Economics* 48, 1799-1815.

Chod, Jiri, and Evgeny Lyandres, 2018, A Theory of ICOs: Diversification, Agency, and Information Asymmetry, Working paper, Boston College and Boston University.

Cong, Lin W., and Zhiguo He, 2018, Blockchain disruption and smart contracts, *Review of Financial Studies*. In-principle Acceptance of FinTech Registered Report, Completion of the Paper in Progress.

Cong, Lin W., Zhiguo He, and Jiasun Li, 2018, Decentralized Mining in Centralized Pools, University of Chicago and George Mason University.

Cong, Lin W., Ye Li, and Neng Wang, 2018, Tokenomics: Dynamic Adoption and Valuation, University of Chicago, Ohio State University, and Columbia University.

Cong, Lin W., and Yizhou Xiao, 2018, Up-Cascaded Wisdom of the Crowd, Working paper, University of Chicago.

Cox, D.R.,1972, Regression Models and Life-Tables, *Journal of the Royal Statistical Society. Series B* 34(2), 187-220.

Da, Zhi, and Xing Huang, 2018, Harnessing the Wisdom of Crowds, Working paper, University of Notre Dame and Washington University.

Diamond, Douglas W., 1984, Financial Intermediation and Delegated Monitoring, *Review of Economic Studies* 51 (3), 393-414.

Dindo, Pietro, and Filippo Massari, 2017, The Wisdom of the Crowd Revisited, Working paper, University of New South Wales.

Downes, David, and Robert Heinkel, 1982, Signaling and the Valuation of Unseasoned New Issues, *Journal of Finance* 37, 1-10.

Easley, David, Maureen O'Hara, and Soumya Basu, 2017, From Mining to Markets: The Evolution of Bitcoin Transaction Fees, Working paper, Cornell University.

Eyal, Ittay, and Emin Guun Sirer, 2014, Majority is Not Enough: Bitcoin Mining is Vulnerable, in *International Conference on Financial Cryptography and Data Security* 436-454. Springer.

Fernandez-Villaverde, Jesus, and Daniel Sanches, 2016, Can Currency Competition Work? Discussion paper, National Bureau of Economic Research.

Frazzini, Andrea, The disposition effect and underreaction to news, *Journal of Finance* 61(4), 2017-2046.

Gandal, Neil, and Hanna Halaburda, 2014, Competition in the cryptocurrency market, Tel Aviv University and Bank of Canada.

Gans, Joshua S., and Hanna Halaburda, 2015, Some Economics of Private Digital Currency, in *Economic Analysis of the Digital Economy* 257-276, University of Chicago Press.

Habib, Michel A., and Alexander P. Ljungqvist, 2001 Underpricing and Entrepreneurial Wealth Losses in IPOs : Theory and Evidence, *Review of Financial Studies* 14(2), 433-458.

Harvey, Campbell R., 2016, Cryptofinance, Working paper, Duke University.

Howell, Sabrina, Marina Niessner, and David Yermack, 2018, Initial Coin Offerings: Financing Growth with Cryptocurrency Token Sales, Working paper, New York University.

Huberman, Gur, Jacob Leshno, and Ciamac C. Moallemi, 2017, Monopoly Without a Monopolist: An Economic Analysis of the Bitcoin Payment System, Working paper, Columbia University.

Ibbotson, Roger G., 1975, "Hot Issue" Markets, *Journal of Finance* 30(4), 1027-1042.

Jia, Chunxin, Jay R. Ritter, Zhen Xie, and Donghang Zhang, 2018, Pre-IPO Analyst Coverage: Hype or Information Production, Working paper, University of Florida.

Kovbasyuk, Sergei, 2011, Wisdom of the Crowd, Working paper, Einaudi Institute for Economics and Finance.

Kremer, Ilan, Yishay Mansour, and Motty Perry, 2014, Implementing the Wisdom of the Crowd, *Journal of Political Economy* 122(5), 988-1012.

Leland, Hayne E., and David H. Pyle, 1977, Informational Asymmetries, Financial Structure, and Financial Intermediation, *Journal of Finance* 32(2), 371-387.

Li, Jiasun, and William Mann, 2018, Initial Coin Offering and Platform Building, Working paper, George Mason University and UCLA.

Ljungqvist, Alexander, 2007, IPO Underpricing, in: B. Espen Eckbo, eds., Handbook of Empirical Corporate Finance, Vol. 1 (Elsevier, Amsterdam).

Mironov, Mikhail, and Steven Campbell, 2018, ICO Market Research Q1 2018, ICORating.com.

Momtaz, Paul P., 2018, Initial Coin Offerings, Working paper, UCLA.

Pagnotta, Emiliano, and Andrea Buraschi, 2018, An Equilibrium Valuation of Bitcoin and Decentralized Network Assets, Working paper, Imperial College.

PwC, 2018, Initial coin offerings: A strategic perspective, June 2018.

Ritter, Jay R., 1984, The "Hot Issue" Market of 1980, *Journal of Business* 57(2), 215-240.

Ritter, Jay R., 1991, The Long-Run Performance of Initial Public Offerings, *Journal of Finance* 47(1), 3-27.

Ritter, Jay R., 2018, Initial Public Offerings: Updated Statistics, University of Florida.

Ritter, Jay, and Ivo Welch, 2002, A Review of IPO Activity, Pricing, and Allocations, *Journal of Finance* 57(4), 1795-1828.

Rock, Kevin, 1986, Why New Issues Are Underpriced, *Journal of Financial Economics* 15, 187-212.

Sagar, Jayanand, 2017, Why Generous ICO Bonuses May Indicate Scam, NEWSBTC, November 23, 2017.

Sehra, Avtar, Philip Smith, and Phil Gomes, 2017, Economics of Initial Coin Offerings, Working paper.

Sockin, Michael and Wei Xiong, 2018, A Model of Cryptocurrencies, Working Paper

Stiglitz, Joseph E., and Andrew Weiss, 1981, Credit Rationing in Markets with Imperfect Information, *American Economic Review* 71(3), 393-410.

Strausz, Roland, 2017, A Theory of Crowdfunding: A Mechanism Design Approach with Demand Uncertainty and Moral Hazard, *American Economic Review* 107(6), 1430-1476.

Surowiecki, James, 2005, The wisdom of crowds, New York: Anchor Books.

Welch, Ivo, 1992, Sequential Sales, Learning and Cascades, *Journal of Finance* 47, 695-732.

Xu, Ting, 2017, Information Role of Crowdfunding, Working paper, University of Virginia.

Yermack, David, 2017, Corporate Governance and Blockchains, *Review of Finance* 21(1), 7-31.

Figure 1: An Illustration of the ICO Timeline
This timeline illustrates the timing of events for a typical ICO. A pre-announcement usually is a summary featuring the idea and team for a startup to the cryptocurrency community to gather interest and feedback. The documentation stage typically involves posting a whitepaper on the startup's website that describes the business model and technical specifications of the project. Many startups also publish initial codes for their ICOs. The ensuing marketing campaign often uses cryptocurrency forums and social network sites such as Medium, Steemit, Reddit, and Twitter. Before the official ICO goes live, there may be an optional presale of tokens. After the ICO, tokens are listed on exchanges.

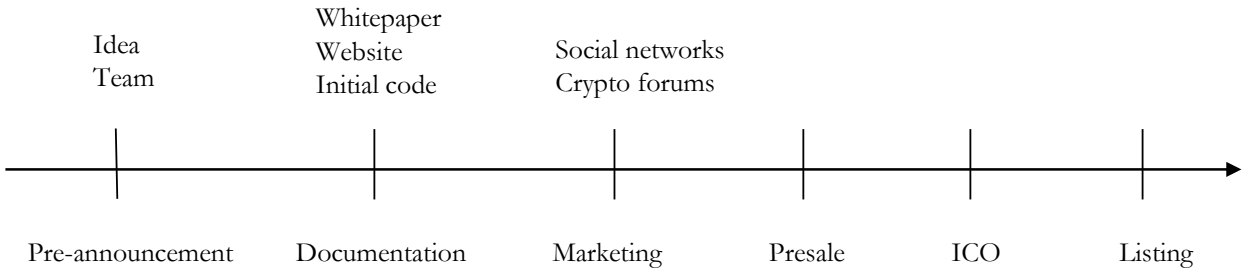

Figure 2: The Aragon Token Sale
Panel A shows time-series patterns of token subscriptions for Aragon Network, an ICO that concluded under 31 minutes. The orange bars plot the gross proceeds ($ million) by minute during the sale and the blue line plots the cumulative gross proceeds ($ million) by minute. The red line and red dotted line plot the hard cap and 50% of the hard cap, respectively. In Panel B, the blue curve plots the cumulative tokens held by percent of largest holders.

Panel A. Token subscription in Aragon Network

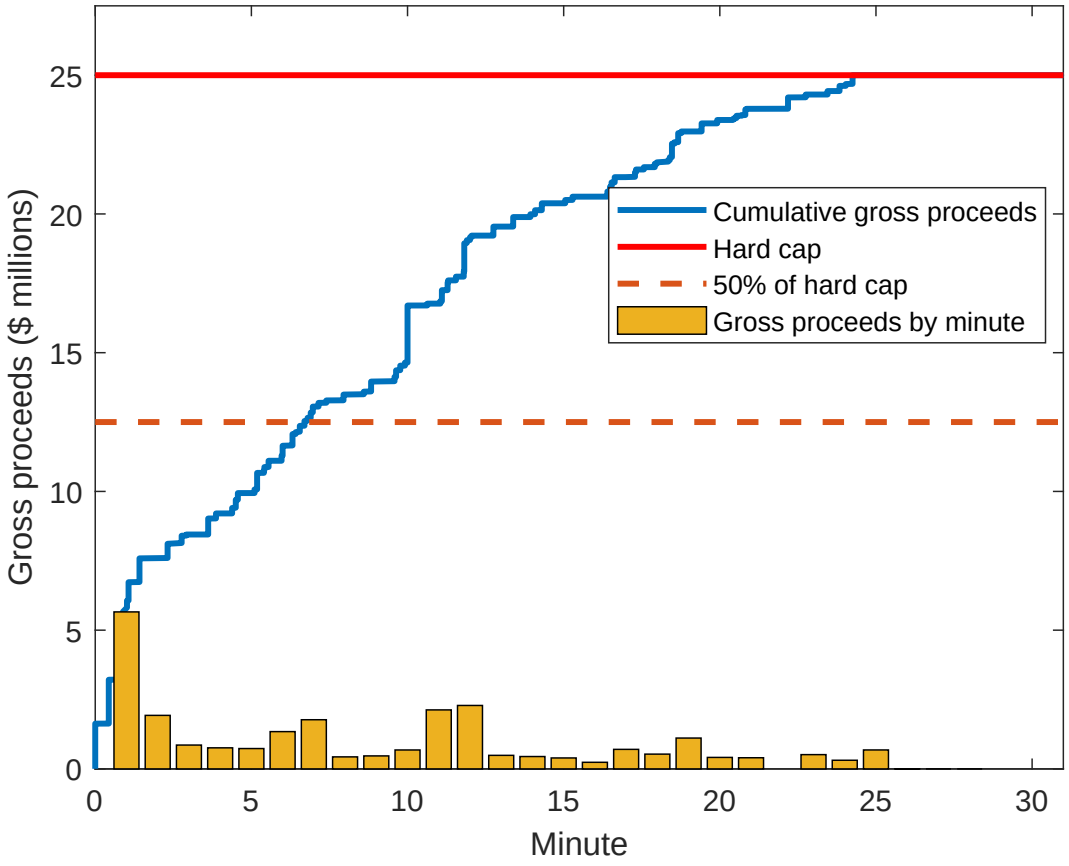

Panel B. Aragon token distribution by investor

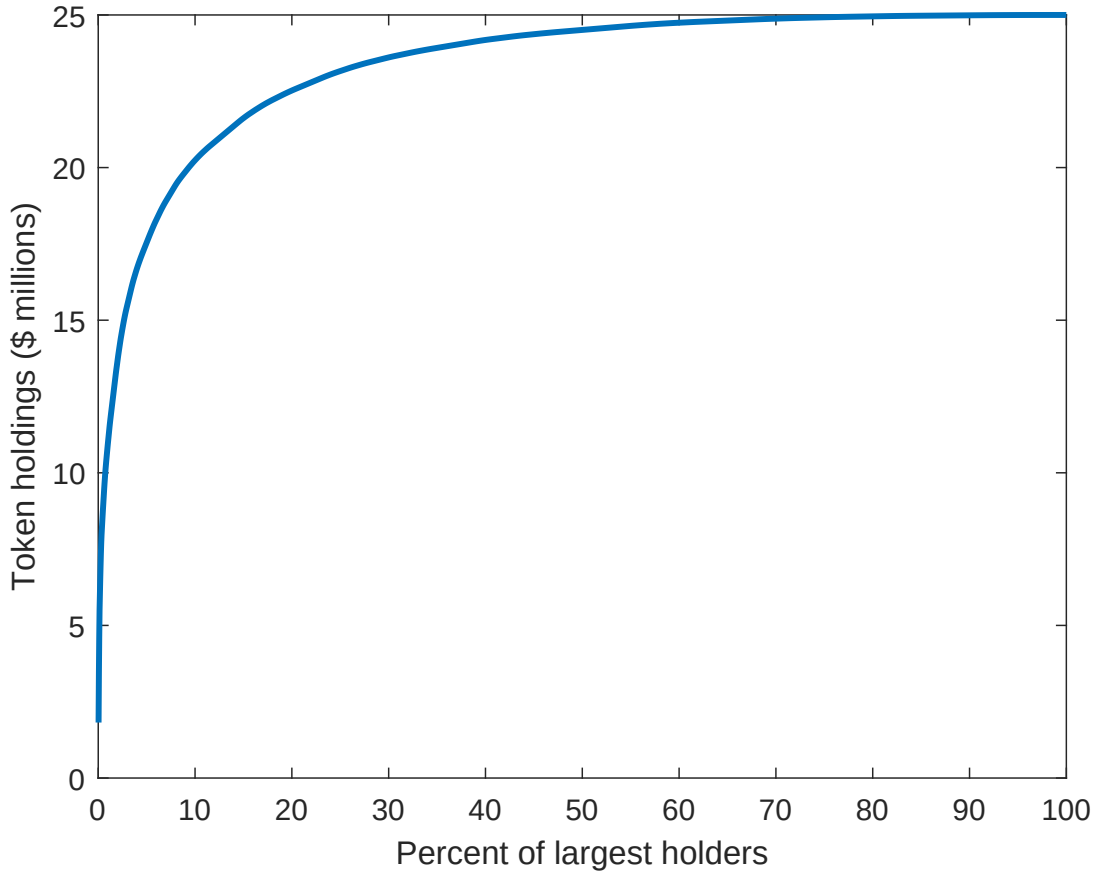

Figure 3: ICO Starts and Success
This figure features all ICOs that opened between Q1 2016 and Q1 2018. The blue bars (left axis) plot the number of opened ICOs in each quarter. The red line (right axis) plots the percentage of successful ICOs by quarter. We exclude ongoing ICOs as of May 31, 2018 when calculating the success rates. An ICO is considered successful if its soft cap was reached or the project raised more than $0.5 million in the absence of a soft cap.

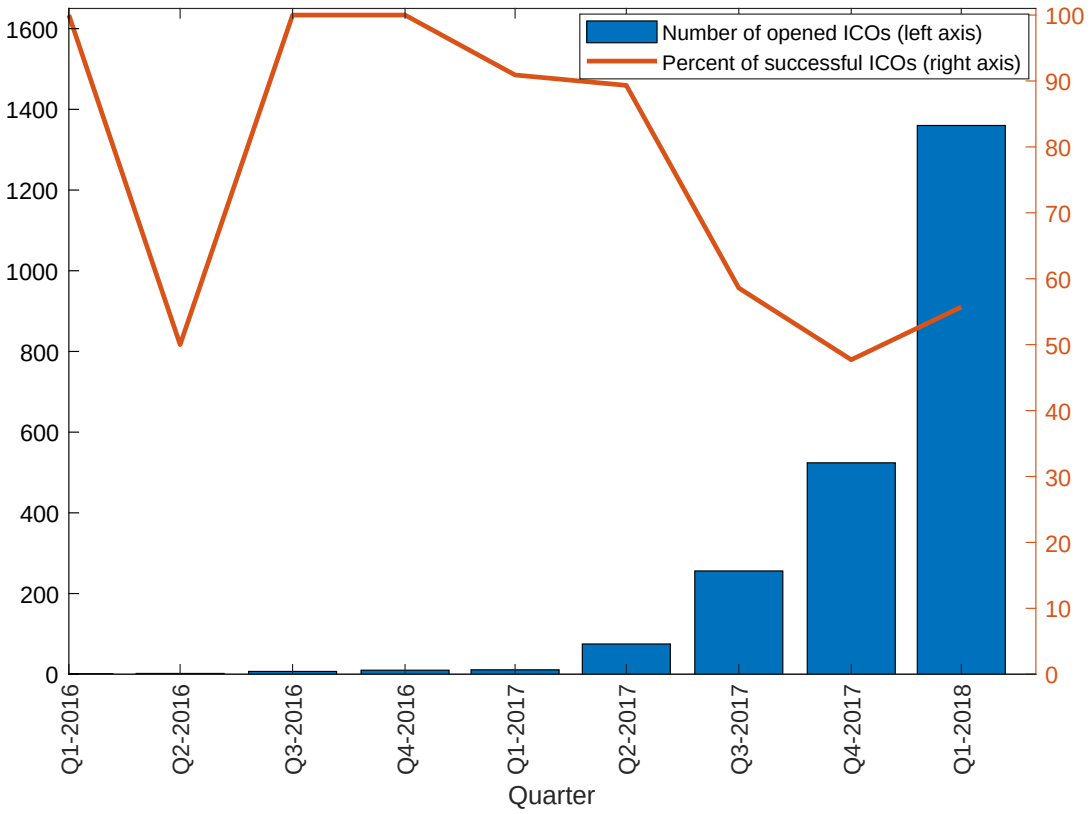

Figure 4: Primary Market Subscriptions in ICOs
This figure shows time-series patterns of token subscriptions during ICOs. Our sample includes all Ethereum-based ICOs that sold a positive number of tokens. In Panel A, the blue (red) bars plot the average daily token sales as a percentage of total tokens for sale in successful (failed) ICOs. The blue (red) dotted line plots the cumulative daily token sales as a percentage of total tokens for sale in successful (failed) ICOs. The plot in Panel B shows the average hourly token sales on the first day of ICO for successful and failed ICOs, while Panel C repeats the analysis for each block during the first hour of token sales. A block contains transactions within a 14-second interval. An ICO is considered successful if its soft cap was reached or the project raised more than $0.5 million in the absence of a soft cap.

Panel A. Daily token sales

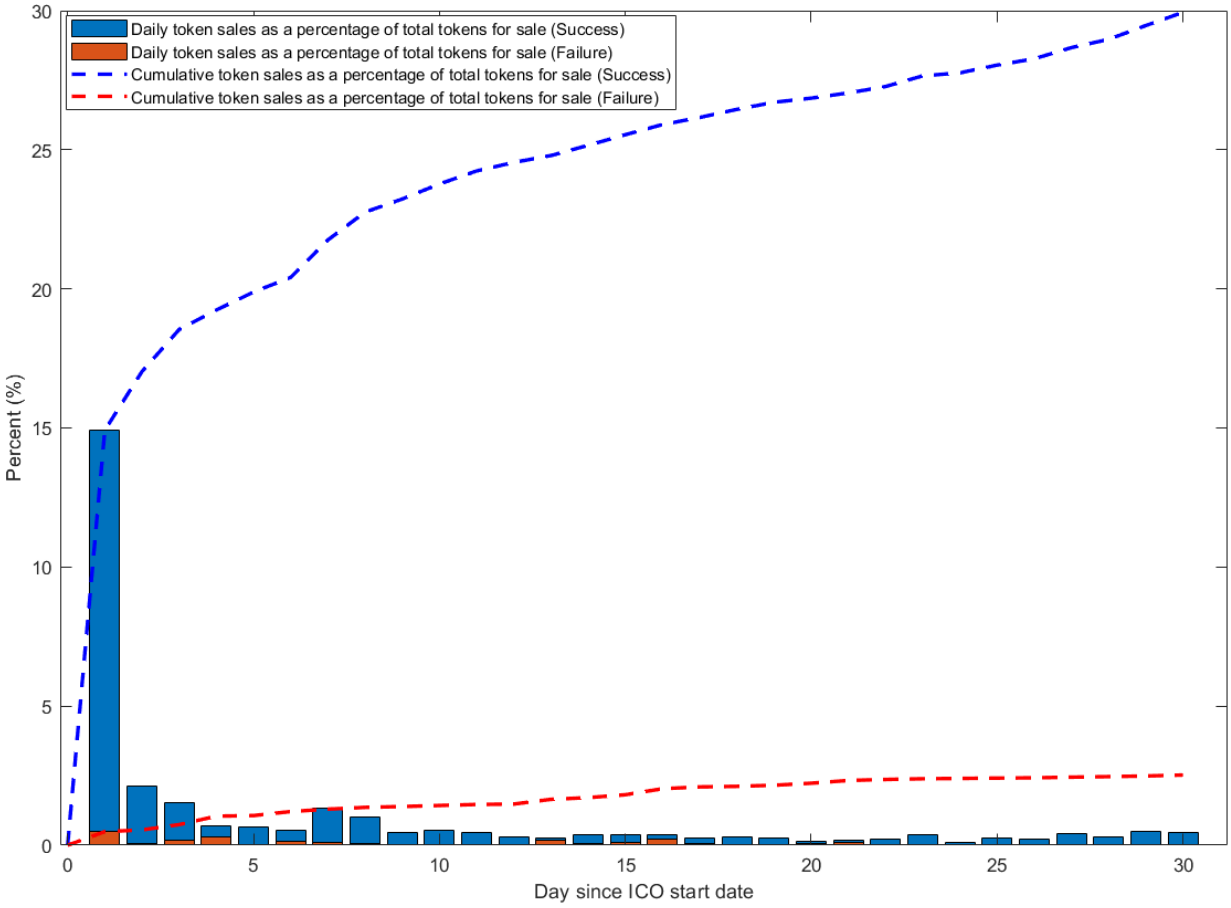

Panel B. Hourly token sales on the first day of an ICO

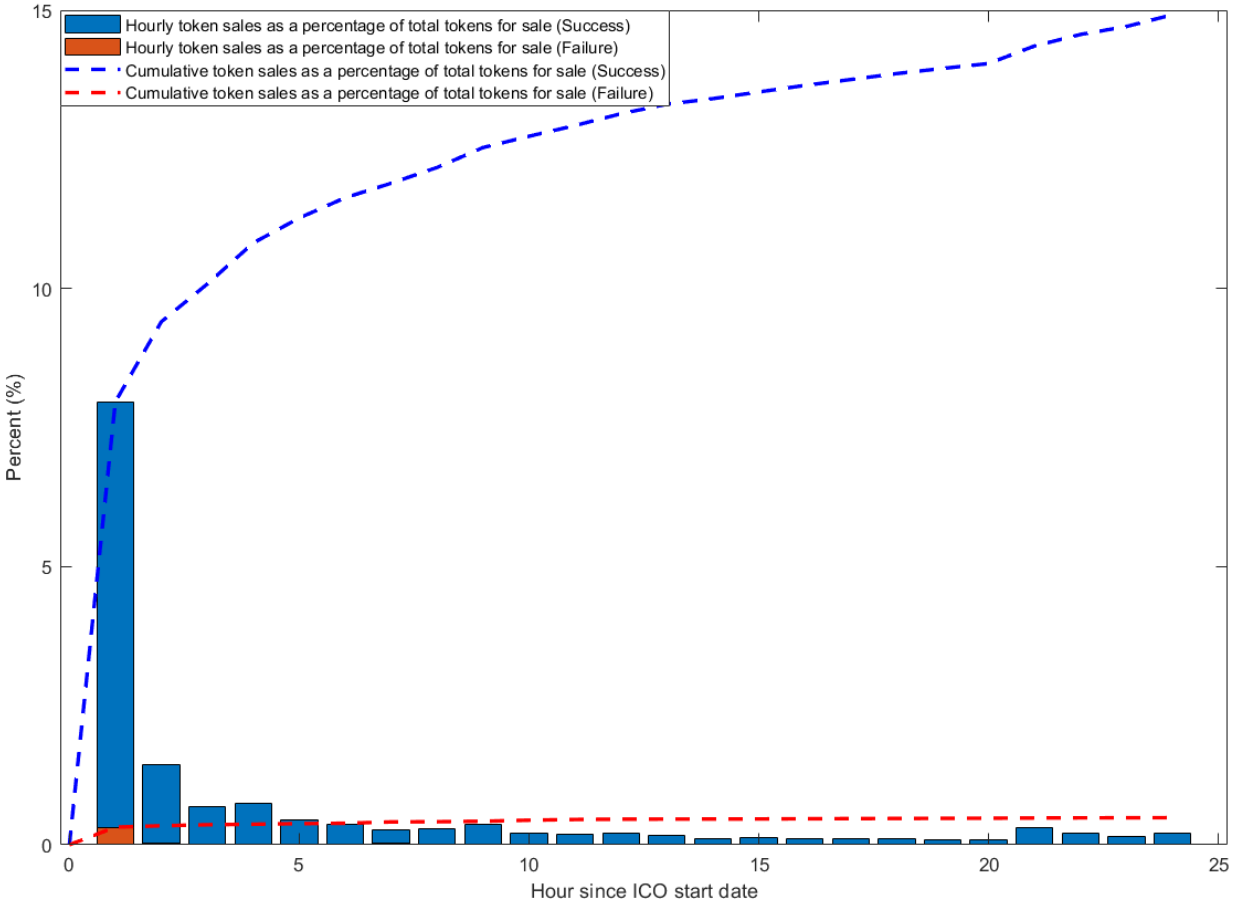

Panel C. Block-by-block sales during the first hour of an ICO

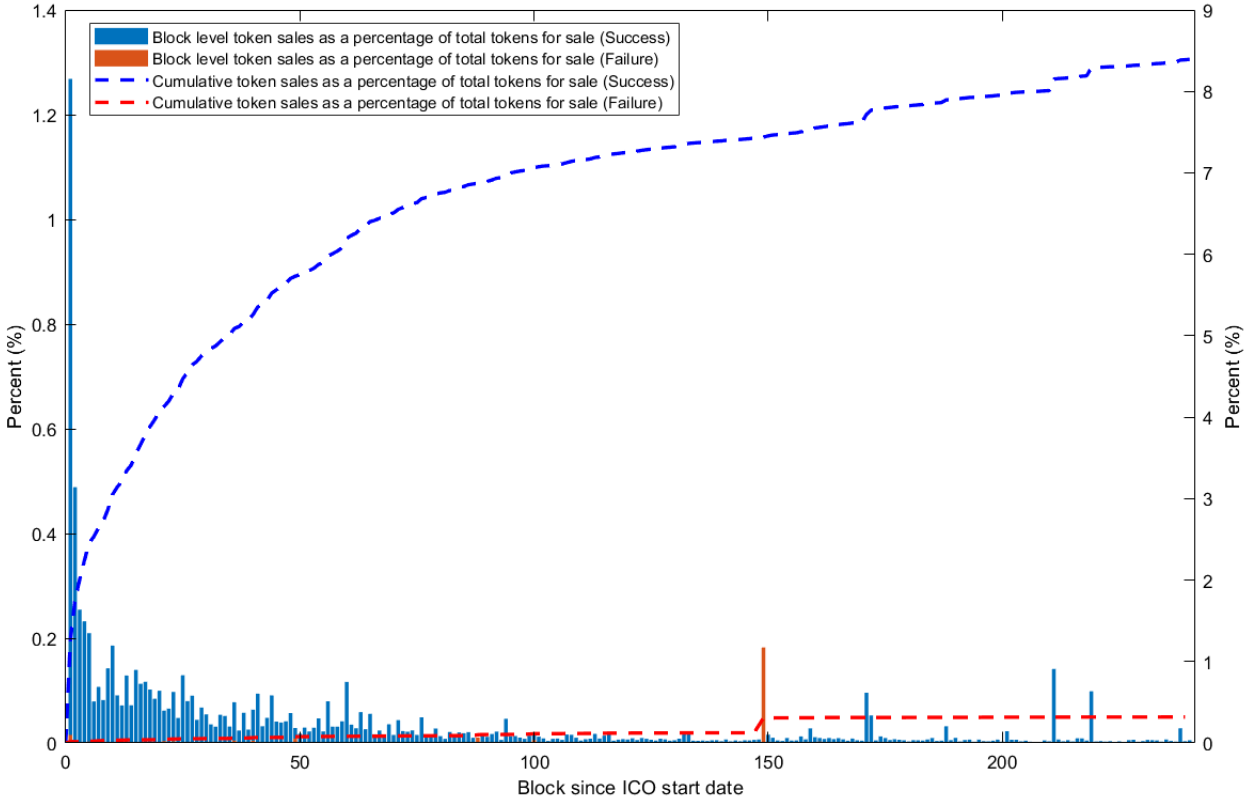

Table 1: ICOs by Country and Industry
This table provides descriptive statistics on ICOs from top 10 countries of incorporation in Panel A, and from top 10 industries in Panel B. We identify ICOs through ICObench, a data provider that specializes in ICO analytics. Our sample includes 1,549 ICOs that started between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. In both panels, we report the number and proportion of ICOs within each country/industry, and the associated fundraising success rate. An ICO is considered successful if its soft cap was reached or the project raised more than $0.5 million in the absence of a soft cap.

Panel A. Most popular countries of incorporation

Country	No. of ICOs	Percent of total	Fundraising success rate
United States	278	18.0%	43.2%
Russia	176	11.4%	33.0%
Worldwide or multiple	160	10.3%	41.3%
United Kingdom	137	8.8%	43.1%
Singapore	113	7.3%	57.5%
Switzerland	62	4.0%	66.1%
China (includes Hong Kong)	52	3.4%	51.9%
Canada	50	3.2%	34.1%
Estonia	46	3.0%	47.8%
Australia	31	2.0%	41.9%
Sum	1,105	71.3%	44.2%
Total (all countries)	1,549	100%	45.4%

Panel B. Most popular industries

Industry	No. of ICOs	Percent of total	Fundraising success rate
Financial Services	159	10.3%	61.4%
Gaming and Virtual Reality	101	6.5%	59.3%
Investment	99	6.4%	50.0%
Exchanges and Wallets	99	6.4%	57.9%
Blockchain Infrastructure	82	5.3%	77.8%
Social Media and Communication	69	4.5%	36.1%
Trading	68	4.4%	54.1%
Business Services and Consulting	66	4.3%	51.4%
Marketing and Advertising	60	3.9%	58.8%
Commerce and Retail	58	3.8%	45.5%
Sum	863	55.7%	56.2%
Total (all industries)	1,549	100%	45.4%

Table 2: ICO Characteristics

This table reports characteristics of the 704 successful ICOs, and compares them with the 845 failed token sales. Our sample includes all ICOs on ICObench that started between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. An ICO is considered successful if its soft cap was reached or the project raised more than $0.5 million in the absence of a soft cap. *Analyst rating* is the average rating (scale 1-5) for an ICO by independent experts on ICObench. *No. of analysts* is the number of analysts that rate an ICO on ICObench. *Soft cap* is the minimum amount of funds needed and aimed at by the startup to proceed as planned, and *Hard cap* is the maximum amount of capital that it aims to gather. *Presale* is an indicator equal to 1 if an ICO runs a token sale event before the official crowdsale goes live. *High bonus* equals 1 if an ICO offers a bonus over 20% (equivalent to a discount of 16.7%), and 0 otherwise. *Fraction of tokens for sale* is the number of tokens for sale divided by the total number of tokens generated. *Know Your Customer (KYC)* is an indicator equal to 1 if clients are required to provide information to confirm their identity. *Whitelist* is a dummy variable equal to 1 if customers have to register in advance to participate in an ICO. *Participation restriction* equals 1 if an ICO is restricted in certain countries, and 0 otherwise. *Multiple languages* is an indicator equal to 1 if the whitepaper or website for an ICO features more than one language. *Multiple currencies* equals 1 if an ICO accepts multiple currencies (digital or fiat), and 0 otherwise. *Gross proceeds* is the amount raised from investors in millions. *No. of subscribers* is the number of token buyers in an ICO. *Duration of offering* is the number of days between the ICO start and end dates. Column (1) reports the average, median and standard deviation of characteristics for successful ICOs. Column (2) shows the difference in average between successful and failed ICOs and its associated t-statistics. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Ex ante ICO characteristics

	Successful ICOs			Difference between successful and failed ICOs	
	Average (1a)	Median (1b)	Std. Dev. (1c)	Diff. in Avg. (2a)	t-stat. (2b)
Analyst rating	3.287	3.500	0.816	0.665***	15.49
No. of analysts	5.805	3.000	8.361	3.087***	9.45
Soft cap ($ million)	6.748	2.680	13.537	-0.046	-0.04
Hard cap ($ million)	88.204	21.755	777.613	48.385	1.02
Fraction of token for sale	0.572	0.600	0.224	-0.042***	-3.68
Presale	0.395	0	0.489	0.182***	8.26
High bonus	0.347	0	0.472	-0.045*	-1.76
Know Your Customer (KYC)	0.117	0	0.322	-0.048***	-2.67
Whitelist	0.082	0	0.274	-0.051***	-3.25
Participation Restriction	0.105	0	0.306	-0.037**	-2.20
Multiple languages	0.413	0	0.493	0.201***	8.79
Multiple currencies	0.320	0	0.467	0.040*	1.69

Panel B. Ex post ICO characteristics

	Successful ICOs			Difference between successful and failed ICOs	
	Average (1a)	Median (1b)	Std. Dev. (1c)	Diff. in Avg. (2a)	t-stat. (2b)
Gross proceeds ($ million)	18.733	9.500	39.735	16.07***	5.28
Gross proceeds/Hard cap	0.590	0.584	0.385	0.436***	11.23
No. of subscribers	2,198.468	658	4,321.582	2,159.710***	10.58
Duration of offering (days)	30.033	29.5	26.560	-7.752***	-5.17

Table 3: Determinants of Fundraising Success

This table examines the determinants of fundraising success for all ICOs that opened between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. The sample includes a total of 1,461 ICOs that have all the required information. In Panel A, the dependent variable is an indicator equal to 1 if an ICO reaches its soft cap or the project raises more than $0.5 million in the absence of a soft cap, and 0 otherwise. In Panel B, the dependent variables are gross proceeds in millions of dollars and gross proceeds divided by the hard cap. Our sample is reduced to 727 (543) ICOs that have information on gross proceeds (both gross proceeds and hard cap). All independent variables are as defined in Table 2, and are measured immediately before the ICO start date. In each column, we report coefficient estimates, their heteroscedasticity-robust t-statistics and, when applicable, the corresponding marginal probability change induced by a one-unit change in the value of a specific covariate from its sample average. Standard errors are clustered at the quarter level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Fundraising success

Dependent variable: ICO success	Probit model			Linear probability model	
	Coefficient	t-stat.	Marg. Prob.	Coefficient	t-stat.
ICO characteristics	(1a)	(1b)	(1c)	(2a)	(2b)
Analyst rating	0.564***	5.17	22.2%	0.164***	4.75
No. of analysts	0.035***	5.16	1.4%	0.012***	8.22
Fraction of tokens for sale	-0.735***	-4.26	-29.0%	-0.160***	-3.24
Presale	0.385***	7.03	15.2%	0.119***	7.95
High bonus	-0.279***	-3.42	-10.9%	-0.083***	-3.12
Know Your Customer	-0.147***	-3.29	-5.7%	-0.057**	-2.27
Multiple languages	0.253*	1.88	10.0%	0.095**	2.35
Multiple currency	0.022	0.46	0.9%	0.015	1.20
Quarterly dummies		Yes		Yes	
Country fixed effects		No		Yes	
Observations		1,461		1,461	
Pseudo R-squared		0.26			
Adj. R-squared				0.32	
% (Dep variable = 1)		45.2%		45.2%	

Panel B. Gross proceeds

Dependent variable	Gross proceeds ($ million)		Gross proceeds/Hard cap	
	Coefficient	t-stat.	Coefficient	t-stat.
ICO characteristics	(1a)	(1b)	(2a)	(2b)
Analyst rating	4.680**	1.98	0.055***	3.23
No. of analysts	0.068	0.41	0.006***	3.25
Fraction of tokens for sale	-17.569***	-2.87	-0.238*	-1.79
Presale	-5.618*	-1.82	-0.023	-0.50
High bonus	-4.747**	-2.13	-0.108**	-2.18
Know Your Customer	-3.260	-1.14	-0.260***	-5.85
Multiple languages	5.450	1.04	-0.014	-0.92
Multiple currencies	0.830	0.46	-0.012	-0.62
Quarterly dummies	Yes		Yes	
Country fixed effects	Yes		Yes	
Observations	727		543	
Adj. R-squared	0.03		0.14	

Table 4: ICO Duration

This table relates token sale completion/resolution to ICO characteristics. The sample consists of a total of 1,115 ICOs from Q1 2016 to Q1 2018 that have all the required information. In column (1), the dependent variable is the logarithm of the number of days between ICO kickoff and completion. All independent variables are as defined in Table 2, and are measured immediately before the ICO start date. Column (1) reports results of an OLS model while column (2) applies a Cox (1972) proportional hazards model to estimate the hazard rate on a daily frequency for ICO completion. In both specifications, we report coefficient estimates and their heteroscedasticity-robust t-statistics. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent variable:					
# Days to resolution	OLS		Cox model		
	Coefficient	t stat.	Coefficient	t stat.	Hazard rate
	(1a)	(1b)	(2a)	(2b)	(2c)
Analyst rating	-0.111***	-3.50	0.101***	2.49	1.11
No. of analysts	0.001	0.07	-0.006	-1.20	0.99
Fraction of tokens for sale	0.088	0.68	-0.079	-0.56	0.92
Presale	-0.089	-0.80	0.229***	3.56	1.26
High bonus	0.112***	4.49	-0.017	-0.27	0.98
Know Your Customer	-0.122**	-2.33	0.032	0.28	1.03
Multiple languages	0.140*	1.75	-0.169***	-2.58	0.84
Multiple currencies	0.124***	3.39	-0.225***	-3.36	0.80
Quarterly dummies	Yes			Yes	
Country fixed effects	Yes			No	
Observations	1,115			39,850	
Adj. R-squared	0.07				
Wald Chi-squared				90.86	

Table 5: Primary Market Token Subscriptions

This table examines how first-day primary market subscriptions affect subsequent investor subscriptions for all Ethereum-based ICOs that opened between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. The sample includes a total of 544 ICOs that have all the required information. *First-day subscription* is the number of tokens subscribed on the first day divided by the number of tokens for sale. *Second-day, Subscription between 2nd and 5th days, Subscriptions between 2nd and 15th days* are similarly defined. All other independent variables are as defined in Table 2, and are measured immediately before the ICO start date. In each column we report coefficient estimates and their heteroscedasticity-robust t-statistics. Standard errors are clustered at the country level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Dependent variable		
	Second-day subscription (1a)	Subscription b/t 2^{nd} and 5^{th} days (1b)	Subscription b/t 2^{nd} and 15^{th} days (1c)
First-day subscription	0.114***	0.143***	0.166***
	[4.76]	[4.36]	[3.43]
Analyst rating	0.004	0.009*	0.015*
	[1.30]	[1.69]	[1.75]
No. of analysts	0.001	-0.001	0.001
	[0.25]	[-0.35]	[0.70]
Fraction of tokens for sale	0.003	0.014	-0.002
	[0.48]	[1.22]	[-0.11]
Presale	0.003	0.011	0.019
	[0.62]	[1.33]	[1.57]
High bonus	-0.004	-0.010*	-0.010
	[-1.41]	[-1.69]	[-0.91]
Know Your Customer	-0.005	-0.011	-0.034**
	[-1.43]	[-1.45]	[-2.25]
Multiple languages	0.005	0.006	0.001
	[1.09]	[0.74]	[0.04]
Multiple currencies	0.007*	0.016*	0.016
	[1.88]	[1.72]	[1.08]
Quarterly fixed effects	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	544	544	544
Adj. R-squared	0.20	0.17	0.22

Table 6: Initial Investors' Investment Returns

This table reports results on the relationship between investor subscription timing and investment returns. The sample includes 298 listed Ethereum-based ICOs that opened between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. In Panel A, the dependent variable is an investor's basis-adjusted return, computed as the percentage deviation of her aggregate cost basis from the current secondary market price (Frazzini, 2006). For any investor at any date t, the aggregate cost basis equals $\lambda^{-1}\sum_{n=0}^{t} Tokens_{t,t-n}Price_{t-n}$, in which $Tokens_{t,t-n}$ is the number of tokens an investor acquired on date $t-n$ that she still holds on date t, λ is a normalizing constant such that $\lambda = \sum_{n=0}^{t} Tokens_{t,t-n}$. $Price_{t-n}$ is the token price on $t-n$, taking into account discounts/bonuses offered at that time. The beginning date is the first day of primary sale. *# days waited until subscription* is the difference between the value-weighted date of purchase and ICO open time. In Panel B, the dependent variable is *Realized return*, which equals $\frac{Total\ proceeds\ in\ the\ second\ market}{Total\ primary\ market\ investment} - 1$. We calculate basis-adjusted returns and realized returns at the end of the first day, first month, three months, and six months, respectively., In each column we report coefficient estimates and their heteroscedasticity-robust t-statistics. Standard errors are clustered at the ICO level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Subscription timing and basis-adjusted returns

	Dependent variable: Basis-adjusted return (%)			
	First day	One month	Three months	Six months
	(1)	(2)	(3)	(4)
# days waited until subscription	-0.320***	-0.289**	-0.339**	-0.581**
	[-2.84]	[-2.37]	[-2.25]	[-2.28]
ICO fixed effects	Yes	Yes	Yes	Yes
Observations	482,126	426,552	345,567	195,936
Adj. R-squared	0.93	0.67	0.57	0.72

Panel B. Subscription timing and realized returns

	Dependent variable: Realized return (%)			
	First day	One month	Three months	Six months
	(1)	(2)	(3)	(4)
# days waited until subscription	-0.695**	-0.388**	-0.616***	-0.576***
	[-2.53]	[-2.39]	[-2.61]	[-2.78]
ICO fixed effects	Yes	Yes	Yes	Yes
Observations	22,608	70,791	115,113	123,150
Adj. R-squared	0.99	0.71	0.61	0.57

Table 7: Determinants of Primary Market Subscriptions

This table examines the determinants of primary market subscriptions for all Ethereum-based ICOs that opened between Q1 2016 and Q1 2018 and were completed as of May 31, 2018. The sample includes a total of 544 ICOs that have all the required information. We study the determinants of first-day token subscriptions, the concavity of investor subscriptions, and the number of days to sell 10% of token supply. *First-day subscription* is the number of tokens subscribed on the first day divided by the number of tokens for sale. Concavity of subscriptions equals the cumulative token sales on the 15th day minus 1/2 the cumulative token sales on the 30th day, with the latter date being the last day of an average ICO. # days to sell 10% of token supply is the number of days until 10% of token supply is sold. All other independent variables are as defined in Table 2, and are measured immediately before the ICO start date. In column (3), we apply a Tobit model with an upper limit of 30 days, which is approximately the average duration for ICOs. In each column we report coefficient estimates and their heteroscedasticity-robust t-statistics. Standard errors are clustered at the country level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	OLS		Tobit
Dependent variable	First-day subscription	Concavity of subscriptions	# days to sell 10% of token supply
	(1)	(2)	(3)
Analyst rating	0.026***	0.024**	-9.516***
	[3.08]	[2.27]	[-5.18]
No. of analysts	-0.001	-0.001	-0.490***
	[-0.62]	[-1.34]	[-5.78]
Fraction of tokens for sale	-0.004	-0.008	7.899
	[-0.17]	[-0.30]	[1.28]
Presale	0.012	0.024**	-8.113***
	[1.02]	[2.51]	[-2.81]
High bonus	-0.017**	-0.013	3.384
	[-2.16]	[-1.55]	[0.84]
Know Your Customer	-0.011	-0.028**	4.485
	[-1.06]	[-2.33]	[1.61]
Multiple languages	-0.008	0.002	-0.124
	[-0.82]	[0.13]	[-0.03]
Multiple currencies	-0.012	0.003	-0.525
	[-1.58]	[0.28]	[-0.71]
Quarterly fixed effects	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	544	544	544
Adj. R-squared	0.18	0.13	
Pseudo R-squared			0.14

Table 8: Investor Subscriptions and ICO Success

This table replicates Table 3 except that column (1) replaces the covariates with *First-day subscription*, while the covariates in column (2) further include the residual from the regression in equation (1). In each column, we report coefficient estimates, their heteroscedasticity-robust t-statistics, and the marginal probability change induced by a one-unit change in the value of a specific covariate from its sample average. Standard errors are clustered at the quarter level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Fundraising success

Dependent variable: ICO success	Probit model					
	Coefficient	t-stat.	Marg. Prob.	Coefficient	t-stat.	Marg. Prob.
ICO characteristics	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)
First-day subscription	18.959***	2.89	762.8%			
Residual of first-day subscription				11.306***	5.47	402.7%
Analyst rating				0.813***	6.65	29.0%
No. of analysts				0.053***	5.14	1.9%
Fraction of tokens for sale				-0.588***	-4.42	-20.9%
Presale				1.037***	13.42	37.8%
High bonus				-0.716***	-3.93	-24.6%
Know Your Customer				-0.456**	-2.32	-14.7%
Multiple languages				0.125	1.13	4.5%
Multiple currencies				-0.241	-0.70	-8.3%
Quarterly dummies		Yes			Yes	
Observations		635			540	
Pseudo R-squared		0.22			0.40	
% (Dep variable = 1)		31.2%			32.7%	

Panel B. Gross proceeds

	Dependent variable : Gross proceeds ($ million)			
ICO characteristics	Coefficient (1a)	t-stat (1b)	Coefficient (2a)	t-stat (2b)
First-day subscription	17.752***	3.15		
Residual of first-day subscription			15.632***	2.34
Analyst rating			1.812*	1.65
No. of analysts			0.148***	3.74
Fraction of tokens for sale			-15.124***	-3.21
Presale			3.483***	2.09
High bonus			-4.379***	-2.86
Know Your Customer			-3.966***	-4.83
Multiple languages			3.020***	2.19
Multiple currencies			2.623***	3.30
Quarterly fixed effects	Yes		Yes	
Country fixed effects	Yes		Yes	
Observations	247		247	
Adj. R-squared	0.22		0.31	

Table 9: Token Returns and Turnover

This table reports statistics on returns and first-day turnover for all listed tokens that were sold through an ICO between Q1 2016 and Q1 2018. The number of observations varies depending on information availability. *First-day return* is measured from the token offer price to the first trading day closing price. *One-month, Three-month, Six-month, and One-year returns* are measured from the first after-market closing price to the 30th, 91th, 182th, and 365th trading closing prices, respectively. One-month, three-month, six-month, and one-year returns all exclude first-day returns. An excess return is calculated as the raw return minus the corresponding compounded daily return on the value-weighted index of Ethereum and Bitcoin. *Gross proceeds* is the amount raised from investors in millions. *Money left on the table* is calculated as the number of tokens issued times the change from the offer price to the first-day closing price. *First-day turnover* is the first-day volume divided by tokens issued. Columns (2), (4), and (6) report the average, median, and standard deviation of the variables of interest, while columns (3) and (5) report the p-value from the t-test and the Wilcoxon signed rank test, which is asymptotically normal, for testing whether the average and median are different from zero, respectively.

Variable	N	Average	p-value from t-test	Median	Wilcoxon p-value	Std. Dev.
	(1)	(2)	(3)	(4)	(5)	(6)
First-day return	432	158.2%	0.00	24.4%	0.00	485.0%
One-month return	430	63.3%	0.00	-24.5%	0.11	394.8%
Three-month return	420	71.6%	0.00	-30.0%	0.40	276.8%
Six-month return	370	143.6%	0.00	-43.3%	0.43	572.2%
One-year return	133	334.9%	0.00	-15.9%	0.02	1228.3%
First-day excess return	432	143.1%	0.00	24.1%	0.00	480.1%
One-month excess return	430	45.7%	0.03	-28.9%	0.00	390.5%
Three-month excess return	420	26.0%	0.04	-34.8%	0.00	257.1%
Six-month excess return	370	66.3%	0.02	-37.4%	0.00	529.1%
One-year excess return	133	21.7%	0.79	-125.3%	0.00	959.6%
Gross proceeds ($ million)	410	21.44	0.00	11.14	0.00	45.04
Money left on the table ($ million)	410	42.20	0.00	1.26	0.00	288.50
First-day turnover	405	6.7%	0.00	1.3%	0.00	27.1%
Days from ICO completion to listing	432	18.46	0.00	14.00	0.00	35.00

Table 10: Short-Run and Long-Run Token Returns

This table examines return patterns for all listed tokens that were sold through an ICO between Q1 2016 and Q1 2018. The number of observations varies depending on information availability. *First-day return, One-month, Three-month, Six-month, and One-year returns* are defined as in Table 9. *First-day excess return* is calculated as the raw first-day return minus the return on the value-weighted index of Ethereum and Bitcoin. *Market returns* are the value-weighted Ethereum and Bitcoin index return for the same return intervals as the dependent variables. All other independent variables are as defined in Table 2, and are measured immediately before the ICO start date. In each column we report coefficient estimates and their bootstrapped t-statistics. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

		Excluding first-day returns			
	First-day return	One-month return	Three-month return	Six-month return	One-year return
	(1)	(2)	(3)	(4)	(5)
First day market return	2.142**				
	[1.86]				
First-day excess return		-0.038*	-0.046**	-0.139***	-0.107
		[-1.93]	[-2.41]	[-3.16]	[-0.18]
One-month market return		1.159***			
		[2.85]			
Three-month market return			0.908***		
			[7.43]		
Six-month market return				2.020***	
				[5.85]	
One-year market return					1.102***
					[3.69]
Analyst rating	-0.002	0.108	0.603*	0.678	2.012**
	[-0.01]	[0.74]	[1.74]	[0.95]	[2.16]
No. of analysts	0.034	-0.011	-0.023	-0.020	2.286
	[0.56]	[-0.95]	[-0.89]	(-0.64]	[0.80]
Fraction of tokens for sale	1.130	-0.372	0.607	1.202	-3.661
	[1.35]	[-0.98]	[0.81]	[1.07]	[-0.38]
Presale	0.218	0.155	0.146	0.792	-1.989
	[0.450]	[0.81]	[0.67]	[1.46]	[-0.75]
High bonus	-1.335**	-0.452**	-0.136	-0.183	-2.577
	[-2.21]	[-2.26]	[-0.38]	[-0.19]	[-1.17]
Know Your Customer	-0.708	-0.490*	-0.761*	-0.320	-0.507
	[-0.77]	[-1.85]	[-1.83]	[-0.46]	[-0.15]
Multiple languages	0.582	-0.277	-0.613*	-0.820	-0.887
	[1.38]	[-1.15]	[-1.76]	[-1.28]	[-0.28]
Multiple currencies	0.323	-0.180	0.089	1.102	-0.855
	[0.91]	[-0.52]	[0.17]	[1.52]	[-0.22]
Country fixed effects	Yes	Yes	Yes		
Observations	423	410	366	356	131
Adj. R-squared	0.06	0.06	0.25	0.34	0.42

Appendix

Figure A1: An Illustration of ICO Bonus Schedules

Panel A shows the bonus structure for Hash Rush, which ran an ICO from September 20 to October 20, 2017. HOQU's bonus structure is reported in Panel B. During HOQU's token sale from November 27, 2017 to February 26, 2018, bonuses were conditional on sales milestones (in HOQU's tokens).

Panel A. Hash Rush's bonus structure

Panel B. HOQU's bonus structure

Bonus system during the main sale period

Bonus	Number of tokens
20%	13,777,764 HQX
15%	13,777,764 HQX
13%	13,777,764 HQX
11%	13,777,764 HQX
9%	13,777,764 HQX
7%	13,777,764 HQX
5%	13,777,764 HQX
3%	13,777,764 HQX
0%	27,555,528 HQX

Figure A2: An Illustration of the Concavity of Investor Subscriptions
We attempt to measure the concavity of the cumulative token demand curve. Our intuitive concavity metric equals the cumulative token sales on the 15th day minus 1/2the cumulative token sales on the 30th day, with the latter date being the last day of an average ICO. In the benchmark of steady daily fundraising volumes, this concavity measure equals zero (the orange line). However, it is strictly positive for a hot ICO that sell tokens faster in earlier periods (the red curve). The concavity measure is strictly negative for a cold ICO that sell tokens faster in later periods (the green curve).

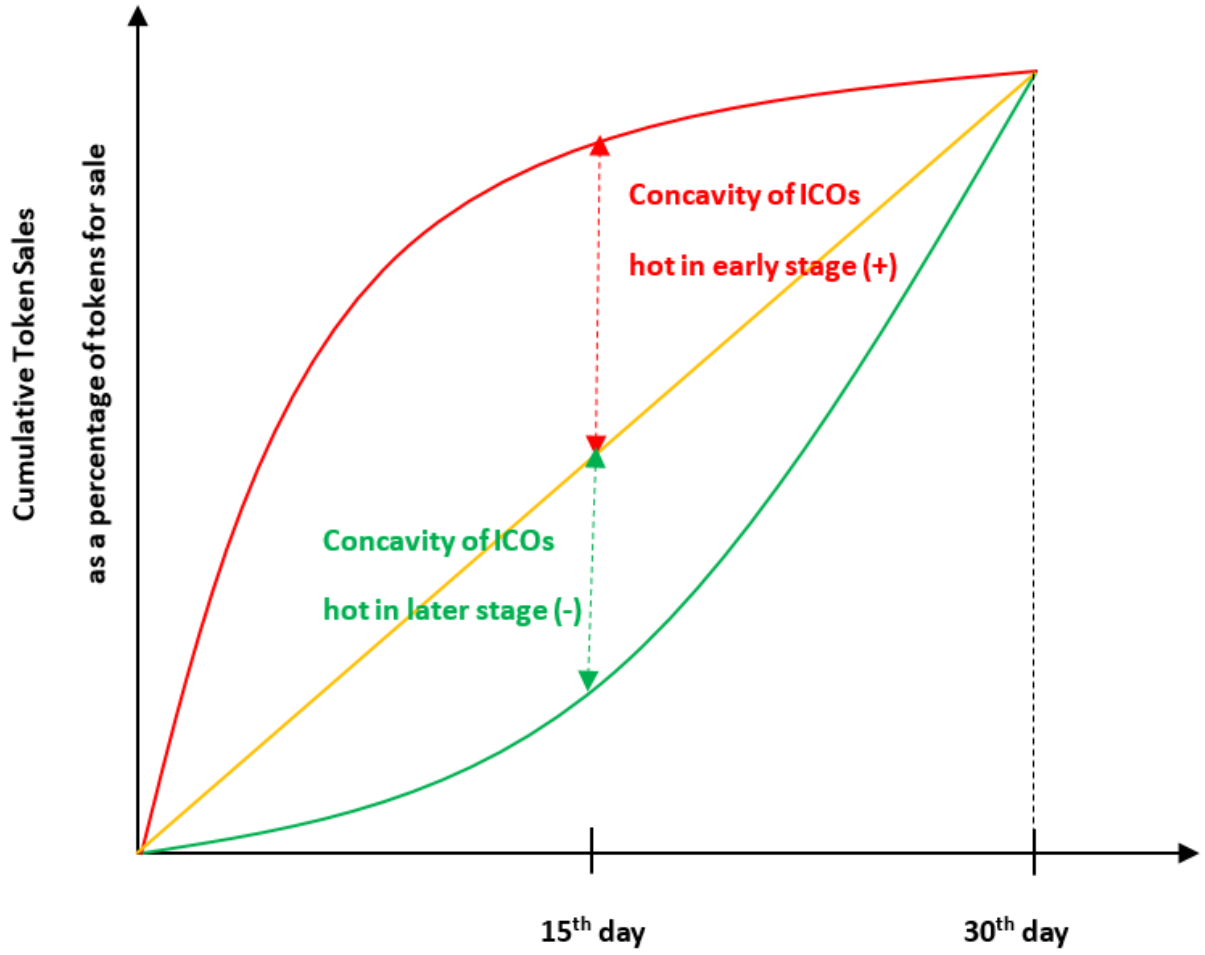

Table A1: Differences between ICOs and IPOs
This table describes the main differences between ICOs and IPOs at each stage of the fundraising process.

	ICO	IPO
1. Pre-announcement	Announce a summary featuring the idea and team to the cryptocurrency community to gather interest and feedback	Hire an investment bank to underwrite the IPO
2. Documentation	- Whitepaper - Website - Initial code	Filings with the securities regulator - Registration statement - Prospectus
3. Marketing	Public relations campaign - Crypto forums - Social network sites such as Medium, Steemit, Reddit, and Twitter	Road show - Meeting with potential investors -Bookbuilding by the underwriter - Offer price set
4. The sale	Subscribers send cryptocurrencies and/or fiat currencies to a digital address. Smart contracts issue tokens based on the exchange ratio.	Shares are allocated to investors
5. Listing	Tokens are listed on a cryptocurrency exchange	Shares are listed on a stock exchange

Table A2: Top 10 ICOs by Gross Proceeds

This table lists the top 10 ICOs by gross proceeds as of May 31, 2018. The ticker for each token is shown in the parenthesis next to the startup name. Information for each ICO is collected from ICObench, and the whitepaper and website for the token sale.

	Startup Name	ICO Start - End	Gross Proceeds ($)
	Dragon Coin (DRG)	2/15/2018 - 3/15/2018	320,000,000
	Huobi (HT)	1/24/2018 - 2/28/2018	300,000,000
	HADC (DAC)	11/27/2017 - 12/22/2017	258,000,000
	Filecoin (FIL)	8/10/2017 - 9/10/2017	257,000,000
	Tezos (XTZ)	7/1/2017 - 7/13/2017	232,000,000
	Sirin Labs (SRN)	12/12/2017 - 12/26/2017	157,886,000
	Styras (STY)	11/28/2017 - 12/30/2017	154,000,000
	The Bancor Protocol (BNT)	6/12/2017 - 6/12/2017	153,000,000
	Bankera (BNK)	11/27/2017 - 2/27/2018	151,800,000
	The DAO (DAO)	4/30/2016 - 5/28/2016	150,000,000

Table A3: Determinants of fundraising success

This table replicates Table 3 except that we replace the independent variable, *Know Your Customer*, with *Whitelist*. All independent variables are as defined in Table 2. In each column, we report coefficient estimates, their heteroscedasticity-robust t-statistics, and the marginal probability change induced by a one-unit change in the value of a specific covariate from its sample average. Standard errors are clustered at the quarter level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Fundraising success

Dependent variable:					
ICO success	Probit model			Linear probability model	
	Coefficient	t-stat.	Marg. Prob.	Coefficient	t-stat.
ICO characteristics	(1a)	(1b)	(1c)	(2a)	(2b)
Analyst rating	0.564***	4.96	22.2%	0.162***	4.66
No. of analysts	0.035***	5.34	1.4%	0.012***	9.04
Fraction of tokens for sale	-0.734***	-4.12	-28.9%	-0.159***	-3.24
Presale	0.383***	7.74	15.1%	0.118***	9.07
High bonus	-0.275***	-3.44	-10.7%	-0.082***	-3.05
Whitelist	-0.250**	-2.39	-9.6%	-0.091**	-2.21
Multiple languages	0.265**	2.06	10.5%	0.099**	2.49
Multiple currency	0.020	0.39	0.8%	0.014	1.26
Quarterly dummies		Yes		Yes	
Country fixed effects		No		Yes	
Observations		1,461		1,461	
Pseudo R-squared		0.26			
Adj. R-squared				0.32	
% (Dep variable = 1)		45.3%		45.3%	

Panel B. Gross proceeds

Dependent variable	Gross proceeds ($ million)		Gross proceeds/Hard cap	
	Coefficient	t-stat.	Coefficient	t-stat.
ICO characteristics	(1a)	(1b)	(2a)	(2b)
Analyst rating	4.834**	2.09	0.054***	3.35
No. of analysts	0.094	0.69	0.006***	3.03
Fraction of tokens for sale	-17.879***	-2.99	-0.227*	-1.80
Presale	-5.594*	-1.69	-0.019	-0.41
High bonus	-5.002**	-2.13	-0.107**	-2.25
Know Your Customer	-7.515***	-4.14	-0.224***	-3.00
Multiple languages	5.426	1.07	-0.007	-0.54
Multiple currencies	0.565	0.30	-0.021	-1.06
Quarterly dummies	Yes		Yes	
Country fixed effects	Yes		Yes	
Observations	727		543	
Adj. R-squared	0.03		0.13	

Table A4: Initial Investors' Investment Returns

This table replicates Table 6 except that the beginning date for both basis-adjusted returns and realized returns is the first day of presale. All variables are defined as in Table 6. In each column we report coefficient estimates and their heteroscedasticity-robust t-statistics. Standard errors are clustered at the ICO level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Subscription timing and basis-adjusted returns

	Dependent variable: Basis-adjusted return (%)			
	First day	One month	Three months	Six months
	(1)	(2)	(3)	(4)
# days waited until subscription	-0.435***	-0.356**	-0.378**	-0.530**
	[-2.71]	[-2.32]	[-2.20]	[-2.21]
ICO fixed effects	Yes	Yes	Yes	Yes
Observations	538,953	479,893	395,021	235,506
Adj. R-squared	0.93	0.69	0.58	0.73

Panel B. Subscription timing and realized returns

	Dependent variable: Realized return (%)			
	First day	One month	Three months	Six months
	(1)	(2)	(3)	(4)
# days waited until subscription	-1.081**	-0.537**	-0.848***	-0.840***
	[-2.43]	[-2.75]	[-2.81]	[-3.23]
ICO fixed effects	Yes	Yes	Yes	Yes
Yes	Yes			
Observations	23,333	74,293	126,880	130,423
Adj. R-squared	0.99	0.73	0.60	0.57