

Redeemable Platform Currencies

Kenneth Rogoff and Yang You *

October 30, 2019

Abstract

Can massive online retailers such as Amazon and Alibaba issue digital tokens that potentially compete with bank debit accounts? We explore whether a large platform's ability to guarantee value and liquidity by issuing prototype digital tokens for in-platform purchases constitutes a significant advantage that could potentially be leveraged into wider use. Our central finding is that unless introducing tradability creates a significant convenience yield, platforms can potentially earn higher revenues by making tokens non-tradable. The analysis suggests that if platforms have any comparative advantage in issuing tradable tokens, it comes from other factors.

1 Introduction

As technology blurs the lines between finance and tech firms, and as innovation in transactions options continues to disrupt markets, many large platforms are issuing, or considering issuing, their own currencies, most famously Facebook's planned 2020 launch of its Libra coin.¹ One important advantage that some tech firms enjoy is an ability to ensure liquidity and value by guaranteeing that their tokens can be redeemed for within-platform purchases and, over the past few years, a number of digital currencies have appealed to this feature. Nevertheless, at this point, the theory of redeemable platform currencies remains relatively underdeveloped. The issue may be of broader significance in the future development of digital currencies in that

*krogoff@harvard.edu, yangyou@g.harvard.edu. Economics Department, Harvard University. We thank Dean Corbae, Robin Greenwood, Zoë Hitzig, Kathryn Holston, Shengwu Li, Laura Xiaolei Liu, Xiaosheng Mu, Julien Prat, Andrei Shleifer, Harald Uhlig, Jiaxi Wang, Glen Weyl, Randy Wright, David Zhang, and seminar participants at Harvard Business School Finance Lunch, Guanghua Peking University, PBCSF Tsinghua University, University of Wisconsin Madison and also the Molly and Dominic Ferrante Fund for research support.

¹Prior to Facebook, there have been many experiments on a smaller scale: for example, Tencent introduced QQ coins for consumers to purchase gaming and non-gaming service provided by Tencent QQ (https://en.wikipedia.org/wiki/Tencent_QQ);

a central issue to regulators is whether new kinds of digital assets offer functionality not embedded in traditional financial assets, and therefore potentially merit consideration for differential regulatory treatment.

Here we present a simple, tractable model of redeemable platform tokens that allows one to explore a number of issues related to their design, features, and supply policy. In principle, such tokens might constitute a prototype currency if they are made tradable (as opposed to non-tradable). However, this does not mean aiming for a prototype currency is necessarily an optimal strategy unless the tokens generate a considerable convenience yield to consumers (compared to bank accounts). We show that maintaining tradability turns out to imply a number of issuance and pricing constraints that can limit a platform's profits from token issuance. In addition, making tokens tradable constraints the ability of a platform to offer richer token price/quantity menus or to incorporate memory features.

To be clear, our stylized framework is partial equilibrium, and takes both the platform's customer base and their outside banking options as given. In the core model, the gains from trade derive from the platform's ability to earn a higher rate of return on its outside investments than can small retail consumers, though we will also discuss convenience yield.² We do not consider gains from trade due to pseudonymity, for example, crypto-currencies that can potentially be used for money laundering, tax evasion, and other illegal activities.³ Nor do we explore the case where platform currencies supplant government fiat money as a unit of account.⁴

Rather, we explore a narrower case where platforms issue redeemable digital tokens that are indexed to fiat currency. From the consumer's perspective, such tokens may be attractive because they are either offered at a discount (the primary focus of our analysis), pay interest, provide a money-like convenience yield, or some combination of the three. For platforms, the advantages include being able to directly tap low-interest retail consumers, to reduce transactions costs, and potentially to benefit from an array of indirect advantages such as strengthening consumer loyalty; these advantages alone can in some cases be significant enough to compensate for having to sell the tokens at a discount that our model endogenizes.

In general, in trying to persuade consumers to hold a significant number of tokens (and thereby garner large seigniorage profits), the core dilemma is this: If the only transaction use of the currency is within-platform, then beyond a relatively modest amount, the coins will

²Prat et al (2019) study ICO utility token pricing with a cash-in-advance constraint.

³Implicitly we are aiming to look beyond the time when regulation sharply circumscribes such uses. Regulators may also be concerned about potential vulnerabilities if and when crypto-currencies become more integral to the global financial system, see Budish (2018).

⁴For examples of the growing recent literature on Bitcoin and the potential for cryptocurrencies to compete with fiat money, see Biais et al (2018), Athey et al (2016), Sockin and Xiong (2018), as well as Schilling and Uhlig (2019).

have to be sold at a discount that is increasing in the number of tokens sold, or alternatively pay a rate of interest that diminishes the platform’s surplus as well as exacerbates fragility issues. Importantly, whether or not a token pays explicit interest can affect how it is taxed and regulated with significant international differences across jurisdictions.

The first part of the paper presents our simple partial equilibrium model of platform tokens and their liquidity. We begin by using the model to explore simple strategies where all tokens are sold for the same price in an initial one-time auction, examining both the case of non-tradable and tradable (“prototype currency”) tokens. A central result is that the non-tradable tokens can be sold at a higher price (for any given quantity) and yield higher profits to the platform. Essentially, tradability forces the issuer to compete with future resale markets and limits the power to charge a high price upfront. Conversely, and of potential significance in designing future regulation, consumers’ share of the gains from trade (due to differing discount rates) tends to be higher with traded tokens.

The next part of the paper explores more sophisticated issuance strategies in which platforms use a price menu approach in their initial coin offering, that is, “buy more and save more”. The advantage of a price menu is that the platform can potentially exploit all the potential gains from inter-temporal trade. But again, such an approach only can only work if the token is non-tradable. Indeed, for tradable tokens, introducing a price menu adds nothing to the platform’s options. Later, when we introduce tokens with memory, the potential advantages of non-tradability become even more apparent. Our analysis of embedded memory is reminiscent of Kocherlakota’s (1998) discussion of money as a crude form of memory.

We then go to look at the case where in addition to its “ICO” (initial coin offering), the firm commits to make “seasoned coin offerings” (SCO) sufficient to keep the outstanding stock of coins constant, that is replacing tokens that have been redeemed. Such an approach can enhance credibility, since the platform has an incentive to preserve its ongoing revenue stream. We show, however, that the prospect of future token sales again sharply discourages consumers from holding more than a very limited number of tokens, even if the issuer can credibly commit to its issuance policy (supported perhaps by a trigger strategy equilibrium.) Indeed, in this case there is actually no longer any advantage to making the tokens non-tradable.

The remainder of the paper goes on to relax a number of the simplifying assumptions of the core model, incorporating the possibility of runs, introducing non-zero cost to platform input goods, and allowing for a convenience yield. The most significant extension is to the case of heterogeneous agents. Allowing for heterogeneity creates a number of subtle pricing and issuance questions, for example, should platform token pricing be designed to peel off the

most active consumers? However, our main results, on tradability versus non-tradability, and on how appetite for token holdings can be extremely sensitive to future issuance policy, appear to generalize.

The final section concludes.

2 A Simple Model

In this section, we develop a simple model to capture how consumers value a token that is underpinned by future claims on platform consumption.

2.1 Consumer Demand

We assume that one unit of the (perishable) platform commodity costs one dollar (there is no inflation in the fiat currency), and provides one unit of consumption. In any given period t , the consumer demands one unit of the platform commodity with probability p , and zero units with probability $1 - p$. All infinitely-lived consumers are identical with time discount factor β . The fact that $p < 1$ captures that the consumer may not need platform consumption every period. The normalization of a single period's consumption to 1 captures limits to the consumer's period demand, but can be varied to study platforms that involve large lumpy expenditures; indeed all the main results here will go through.⁵

Consumers are risk-neutral and have a utility function that is linear in the consumption of the platform commodity given by

$$U_t = \sum_{s=t}^{\infty} \beta^{s-t} \theta_s C_s \quad (1)$$

where θ_s is the dummy of platform consumption shock in period s : $\theta_s = 1$ with probability p , and $\theta_s = 0$ with probability $1 - p$. θ_s is i.i.d.

2.2 Valuing the Marginal Claim

In all that follows, a critical issue is how a consumer values a credit that pays for her M^{th} unit of platform consumption, which will occur at some future date $N \geq M$, depending on the exact timing of the consumer's needs for the platform good. The probability that the consumer will use the M^{th} token in period N is given by

⁵Define Π as the fiat currency price of a platform good. The scale of a platform depends on $p\Pi$. In our analysis, the price of a unit good is normalized to one. But one can envision of a platform with low-frequency consumption (low p) but a high fiat currency price. One can easily show that all results go through with an arbitrary price of Π for the platform good. That is, platform scale is irrelevant, only the consumption probability matters.

$$X_{N,M} = \binom{N-1}{M-1} p^M (1-p)^{N-M} \quad (2)$$

where $\binom{N-1}{M-1}$ is the binomial coefficient $\frac{(N-1)!}{(N-M)!(M-1)!}$. Given consumers' linear utility function (1), expression (2) governs the value of the marginal claim which is a central input to how much a consumer is willing to pay for tokens.

2.3 Platform Currency and Issuance

We now introduce the possibility that platform can issue a “currency” in the form of non-interest-bearing tokens that can be converted to one unit of the platform commodity in any given period. Of course, given the assumed utility function, the consumer will never need more than one token in any given period. Importantly, the consumer is not required to use the platform token and can always pay one dollar of fiat currency (that is, one dollar). As with the consumer, the platform is risk-neutral.

The platform discounts the future at $\beta^* < \beta$, to capture that as a large platform, it has better outside investment opportunities than do small consumers.⁶ This wedge is the sole source of gains from inter-temporal trade to justify token issuance in this paper. It immediately follows that in an efficient equilibrium, with no other liquidity, capital constraints or credibility issues, the consumers would purchase the entire present value of future platform consumption in the initial period, with the allocation of the welfare surplus from trade depending on the relative bargaining power of the two parties, for example depending on consumers' outside options. As noted in the introduction, there may be many other reasons for gains from token issuance, but for the moment, we will focus exclusively on the discount wedge.

One critical issue is the extent to which the platform currency yields a flow of convenience services for transactions inside the platform, and potentially for trade outside, an assumption that is widely used to rationalize demand for currency that pays below the short-term market rate of interest.⁷ For now, we assume the convenience yield is zero in all transactions, that the token is only used for platform purchases, and that it does not effectively compete with fiat currency for trade outside the platform. We return to the convenience yield issue later.

⁶As Demirgüç-Kunt and Huizinga (2000) show, using bank data across 80 countries, the net interest margin banks are able to earn (the difference between their deposit and lending rates) depend on an enormous range of factors, including both explicit and implicit taxation, leverage market concentration, deposit insurance regulation, macroeconomic conditions and many other factors. Regulation can be expected to play a similarly large role in shaping the net interest margin for tech companies.

⁷Digital currencies clearly can yield substantial convenience: consumers do not need to enter a security code and wait for verification when they pay with Amazon credit. Alipay's convenience service makes Hangzhou a “cash-free” city, where 80% of people make payments with their smartphones, rather than cash or credit card.

2.4 Assumptions

Before proceeding to study token offerings, we initially make a number of assumptions to simplify the analysis, and later discuss what happens when we relax them.

1. Token issuance does not affect consumer demand for platform consumption. This abstracts from a number of possible benefits, for example, if currency issuance increases consumer time spent on the platform.
2. Zero production cost. This assumption not only abstracts from the cost of producing platform intermediation services, but also from the cost the platform pays in purchasing commodities to sell to consumers.
3. No platform failure or bankruptcy (otherwise a default premium is built into the token). Relatedly, if the platform issues tokens, these are assumed senior to any other debt the platform may issue.
4. The platform can make credible commitments to its future token issuance policy and to redeemability.
5. Any token issued by the platform is effectively a “stable coin” whose platform-use value is fixed in terms of fiat money, and we assume no inflation.
6. Platform tokens are tradable among consumers only if the platform allows it.

These assumptions can be modified to get more general results. In particular, we later allow for a convenience yield, a proportional cost of goods (relax assumption 2), and especially importantly, relax the assumption that the platform can make credible commitments (assumption 4). Other extensions are possible.

3 Introduction of Platform Currency through ICOs and SCOs

We now proceed to study the pricing and issuance strategies for a platform that introduces tokens either through a once and for all “initial coin offering” (ICO)⁸ or through a combination of an ICO and ongoing “seasoned coin offerings” (SCO). Note that if the platform did not engage in any financial offerings, its value (per consumer) would simply be the expected present value of sales:

$$\frac{\beta^*}{1 - \beta^*} P \tag{3}$$

⁸A more precise term would be “initial token offering”, however, we follow industry convention.

The first-best is that consumers transfer their entire willingness to pay to the platform in the first period. It is achievable by issuing a life-long membership which enables pay once and enjoy the free service for all time. The first-best discounted revenue:

$$\frac{\beta}{1-\beta}p \quad (4)$$

The present value of revenue after token issuance is bounded by $[\frac{\beta^*}{1-\beta^*}p, \frac{\beta}{1-\beta}p]$. We consider a range of issuance policies and compare policies from the standpoint of the issuer.

3.1 Non-tradable Initial Coin Offering

We first consider the case where the tokens issued by the platform are not tradable, and in which the platform announces a fixed (per capita) quantity of tokens that it is going to sell, M . Importantly, in order to sell the full quantity of tokens the platform has put up for sale, all the tokens must be priced at the value of marginal token M , which is the last to be spent. Making use of equation (2), we can solve for the token price in the non-traded case $P_{I,N}$ ⁹

$$P_{I,N} = \sum_{N \geq M} \beta^N X_{N,M} = \sum_{N \geq M} \beta^N \binom{N-1}{M-1} p^M (1-p)^{N-M} = \left[\frac{\beta p}{1-\beta(1-p)} \right]^M \quad (5)$$

One may view $\frac{\beta p}{1-\beta(1-p)}$ as the effective discount rate when the platform aims to issue an extra token. To sell an additional token, all tokens sold must depreciate $\frac{\beta p}{1-\beta(1-p)}$ which yields higher surplus for consumers. Note that we have assumed platform sets the issuance quantity M , but here it could equivalently set the token price P .¹⁰

⁹The last equation uses a combinatorial identity where $x = \beta(1-p)$

$$\sum_{k \geq 0} \binom{M-1+k}{k} x^k = \left(\frac{1}{1-x} \right)^M$$

An alternative and perhaps more intuitive approach to derive eq.(5) is through induction. Note that with one token, the present value to the consumer is $V(1) = \beta p(1 + \beta(1-p) + \beta^2(1-p)p + \dots) = \frac{\beta p}{1-\beta(1-p)}$. Then, we can solve $V(M)$ with an iterative iterative process. In a period when X^{th} token spent, the expected value of $(X+1)^{th}$ token is always $\frac{\beta p}{1-\beta(1-p)}$ where the value of X^{th} token is one. Then, we have

$$V(X+1) = \frac{\beta p}{1-\beta(1-p)} V(X)$$

Thus, the token price of M-token issuance is $V(M) = \left(\frac{\beta p}{1-\beta(1-p)} \right)^M$.

¹⁰The idea that the value and liquidity of a token depend on an underlying matching probability is reminiscent of Kiyotaki and Wright's classic (1989) general equilibrium search-theoretic model of commodity money.

3.2 Optimal issuance

To calculate optimal issuance, it is necessary to take into account both the gross profit the platform gets from the ICO and the present value of foregone fiat money sales. As an intermediate step, and to help intuition, it is useful to first calculate the level of currency issuance that would maximize revenue ignoring foregone sales, in which case the platform maximizes $P_{I,N}M$, so that the first-order condition is given by.

$$\left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M + \ln \left(\frac{\beta p}{1 - \beta(1 - p)} \right) \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M M = 0$$

which implies that

$$M = \frac{1}{\ln \frac{1 - \beta + \beta p}{\beta p}}$$

which depends positively on both β and p . Notice that the platform's discount rate β^* does not enter this formula.

Of course, the full maximization problem for the firm involves taking into account that if a consumer purchases M tokens, then she will use tokens for her first M purchases instead of paying in fiat currency. Thus the platform's complete maximization problem is given by

$$\max_M \left\{ \underbrace{M \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M}_{\text{Token Revenue}} - \underbrace{\sum_{i=1}^M \left[\frac{\beta^* p}{1 - \beta^*(1 - p)} \right]^i}_{\text{Foregone Cash Revenue}} \right\}$$

Define $R_{I,N}$ as the total revenue for non-tradable token issuance. Then, we can rewrite the firm's profit of token issuance as

$$\underbrace{M \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M}_{\text{Revenue from Token Issuance}} + \underbrace{\left(\frac{\beta^* p}{1 - \beta^*(1 - p)} \right)^M \frac{\beta^* p}{1 - \beta^*}}_{\text{Revenue in Fiat Money}} - \underbrace{\frac{\beta^* p}{1 - \beta^*}}_{\text{No Token Issuance}} \quad (6)$$

where the first two terms are the total revenue from token issuance $R_{I,N}$ and fiat sales after all tokens are used.

$$M \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M + \left(\frac{\beta^* p}{1 - \beta^*(1 - p)} \right)^M \frac{\beta^* p}{1 - \beta^*} = R_{I,N}$$

The third term in eq.(6) represents the value of the firm in the absence of currency issuance. M^* is a local optimum if

$$\left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M \geq (M - 1) \left(\left[\frac{\beta p}{1 - \beta(1 - p)} \right]^{M-1} - \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^M \right) + \left[\frac{\beta^* p}{1 - \beta^*(1 - p)} \right]^M \quad (7)$$

and




Figure 1 shows the gain and loss by reducing one token (from three to two). The issuer gains from the price increase, but loses from revenue from the marginal token. Given three tokens are optimal in this example, this figure corresponds to optimal issuance constraint eq.(7).

$$\left[\frac{\beta p}{1 - \beta(1-p)} \right]^{M+1} < M \left(\left[\frac{\beta p}{1 - \beta(1-p)} \right]^M - \left[\frac{\beta p}{1 - \beta(1-p)} \right]^{M+1} \right) + \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{M+1} \quad (8)$$

Figures 1 and Figure 2 illustrate inequalities (7) and (8). The gray areas represent the net gain and loss from issuing one less (more) token (assuming the optimal number is 3). The blue bars represent the present value of the of the foregone fiat money revenue from the M^{th} token issued $\left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^M$. The dashed bars at the top represent the token price when M tokens are issued $\left[\frac{\beta p}{1 - \beta(1-p)} \right]^M$. (Here β appears instead of β^* .) For example, Figure 2 shows how if the platform issues one token above the optimum of 3, it gains additional revenue from the extra token issued, but suffers from the price decline on other tokens, as well as the present value of the foregone fiat revenue on the marginal redemption.

A few observations: First, clearly the optimal issuance level M is less than $\frac{1}{\ln \frac{1 - \beta + \beta p}{\beta p}}$, which maximizes the firm's gross revenue from token issuance without taking into account the foregone




Figure 2 shows the gain and loss from increasing tokens issued from three (optimal) to four. The issuer loses from the price decrease, but gains revenue from the extra token. Given three tokens are optimal, this figure corresponds to optimal issuance constraint, eq.(8).

future sales in fiat money. Second, Appendix 2.1 shows that eq.(7) and eq.(8) are necessary and sufficient conditions for the optimal issuance level. In another words, these two inequalities pin down a unique M as the revenue-maximizing issuance quantity. Third, It is straightforward to show that optimal issuance is monotonic in the key parameters β^* as long as the optimal issuance level M^* is larger than one. A low- β^* firm values present resources more and prefers to issue more tokens. (Also see Appendix 2.1 for the proof.)

Finally observe that with the pure ICO considered here, it does not matter if the platform announces a quantity or a price, since there is complete information, provided the firm is committed to selling all coins at the same price (perhaps because of regulation.) It is this very constraint that leaves consumers some surplus when $M > 1$, and allows them to enjoy some of the gains from token issuance.

3.3 Tradable ICO

We begin by noting that once all individuals are holding at most one token, the token price is governed by the willingness to pay of individuals who have fully depleted their token supply. If the price is higher than their willingness to pay (WTP), no one wants to buy, and selling pressure pushes the token price down. If the price is lower than the WTP, every consumer without a token wants to buy one and this bids up the price. Thus, the token price is unique when all individuals are holding at most one coin. Let \hat{P} denote this unique and steady-state price.

$$\hat{P} = \beta p + \beta(1-p)\hat{P}$$

The first term on the right-hand side denotes the present value of being able to consume the coin in the next period, and second term denotes the present value of being able to sell it. But this equation can be rearranged to yield

$$\hat{P} = \frac{\beta p}{1 - \beta(1-p)}$$

which is exactly the same as in the non-tradable case. Once all individuals have either zero or one token, there are no longer any gains from inter-consumer trade; a token has the same value to an individual whether she sells it or holds on until she has the first opportunity to use it. Inducing individuals to hold more than one coin, however, requires that they expect the price to appreciate at the rate β^{-1} every period, again assuming as we have so far that the convenience yield is zero.

Now suppose the platform wants to sell M tokens in an ICO, but where tokens are tradable; what is the price? The key observation is that if there are M tokens, it will take $(M-1)/p$ periods for the first $M-1$ coins (per capita) to be depleted. (This is much faster than would be the case without trade.) In period $1 + \frac{M-1}{p}$, the price must reach its steady-state value of \hat{P} . The ICO price for M tradable tokens must be given by

$$P_{I,T} = \beta^{\frac{M-1}{p}} \left(\frac{\beta p}{1 - \beta(1-p)} \right) \quad (9)$$

To compare the gross revenue from a non-traded ICO of M tokens with a traded ICO of the same size, we first observe that when $M = 1$, tradability does not matter since all agents are homogeneous. We then note from equation (5) that to issue one extra non-tradable token, the platform needs to discount token prices by $\frac{\beta p}{1 - \beta(1-p)}$ while in the case of tradable tokens,

equation (9) implies it would need to discount its price by $\beta^{\frac{1}{p}}$.¹¹ Thus to compare the price of tradable tokens with that of non-tradable tokens (for any equivalent-size ICO), we need only to compare the two discount factors. Proposition 1 answers this question.

Proposition 1 (Effective Discount Factor Dominance): The effective discount factor is higher (closer to 1) for non-tradable ICO tokens than for tradable ICO tokens. (See Appendix for the proof.)

$$\beta^{\frac{1}{p}} < \frac{\beta p}{1 - \beta(1 - p)} \quad (10)$$

Comments: Proposition 1 states that for any sale of $M > 1$ tokens in an ICO, the price will be higher if the tokens are non-tradable. What is the intuition for this result, given that the expected time to the redemption of the marginal token is greater in the case of non-tradability? The answer has to do with the fact that the consumer’s utility function is convex in time of consumption.¹² While commonly known that utility is concave in consumption for any given period, it is less known that the utility is convex in the time of consumption. To illustrate this convexity, consider the following two lotteries in time with the identical expected payoff.

Lottery 1 (Price P_C): One dollar in period $M + 2$.

Lottery 2 (Price P_D): One dollar in period $M + 1$ with 50% probability, and one dollar in Period $M + 3$ with 50% probability.

Lottery 1, sold at a price P_C , delivers payoff one dollar with a “compressed” distribution in time — with 100% certainty in period $M + 2$. Lottery 2, sold at a price P_D , delivers one dollar payoff with a “dispersed” distribution in time — with 50% probability in either period $M + 1$ or $M + 3$. As shown in Figure 3, if one dollar yields the same utility $u(1)$ for any period, then convexity implies that for a given expected cash flow, the more dispersed the distribution in time, the higher it will be priced:

$$P_C = \frac{1}{2}u(1)(\beta^{M+1} + \beta^{M+3}) > u(1)\beta^{M+2} = P_D$$

The initial ICO token price, in both the tradable and non-tradable case, depends on the willingness to pay for the marginal token. Tradability compresses the distribution of the time

¹¹We can immediately derive the ICO revenue maximizing $M = \frac{1}{-\log(\beta^{\frac{1}{p}})}$ for the tradable tokens case (ignoring the opportunity cost of lost future fiat money sales). Proposition 1 implies the optimal revenue-maximizing quantity of the tradable ICO is lower than that for non-tradable ICO.

¹²The fact that additive utility functions are convex in time has been previously noted and studied experimentally by DeJarnette et al. (2019).




Figure 3 plots the convex β -discounting function and shows the prices of the following two lotteries: Lottery 1 (price P_C): One dollar in period $M + 2$; Lottery 2 (price P_D): One dollar in period $M + 1$ with 50% probability, and one dollar in period $M + 3$ with 50% probability. Convexity implies $P_C > P_D$.

required to spend the marginal token. In Figure 4, we plot the probability distribution function of the period in which the marginal token is spent with $M = 10$ tokens and consumption shock probability $p = 0.5$. All non-tradable tokens might be spent in as few as 10 periods if consumption shocks arrive in every period ex-post, but will typically take a much longer time for most consumers. For tradable tokens, all consumers always use tokens in the first $\frac{M-1}{p} = 19$ periods. As shown in Figure 4, tokens start to deplete in the Period 19 ($\frac{M-1}{p} + 1$) with probability $p = 0.5$. The time distribution of expenditure for tradable tokens is compressed compared to non-tradable tokens. Given the convexity of the utility function in time, tradability thus lowers the token price.

An alternative interpretation is that tradability creates a resale market that pushes the platform to compete with itself. (This interpretation has a loose analogy to the Coase Conjecture albeit here consumers are homogenous.) The resale market introduces competition with the future and reduces the token price. One might ask how a tradable token can sell for less than a non-tradable token when the consumer always has the option of not trading it. The answer




Figure 4 plots the probability the period of the last token spent for tradable and non-tradable ICO respectively. We pick $M = 10$ tokens issued and the probability of consumption is $p = 0.5$. The solid line plots the probability distribution function of tradable tokens. The dashed line plots the probability distribution function of non-tradable tokens.

is that with tradability, the platform cannot command as high a price precisely because the consumer knows there is always an option of buying on the outside market in the future, and this drives the requirement that the market price of a tradable token must rise faster than the shadow price of a non-tradable token, as we have just proven. As we shall see in later sections where we look at richer pricing strategies and memory functions, the potential advantages of non-tradable tokens run much deeper than just this point.

Of course, the proceeds from the ICO do not capture the entire story, since whenever a consumer tenders a token for a later purchase, the platform has to forego fiat currency revenue that it would have enjoyed absent any token issuance. But, as we next demonstrate in Proposition 2, the present value of future fiat revenue sales is also higher when tokens are non-traded, so a non-traded token ICO is unambiguously more profitable than a traded token ICO.

Recall that $R_{I,N}$ is the total revenue from non-tradable token issuance, given below eq.(6). Define $R_{I,T}$ analogously.

Proposition 2 (Revenue Dominance): Tradability reduces the discounted revenue of the

issuer.

$$R_{I,N} > R_{I,T}$$

Proof of Proposition 2

The present value of firm revenue from a one-time tradable token ICO, including both the initial token sales revenue, and revenue from future fiat currency sales is given by

$$\begin{aligned} R_{I,T} &= M \times P_{I,T} + \beta^* \frac{M-1}{p} \frac{\beta^* p}{1 - \beta^*(1-p)} \left(\frac{\beta^* p}{1 - \beta^*} \right) \\ &< M \times P_{I,N} + \left(\frac{\beta^* p}{1 - \beta^*(1-p)} \right)^M \left(\frac{\beta^* p}{1 - \beta^*} \right) = R_{I,N} \end{aligned} \quad (11)$$

This inequality $\beta^* \frac{1}{p} < \frac{\beta^* p}{1 - \beta^*(1-p)}$ holds as the β^* version of the condition proved in Proposition 1. (That is, simply substitute β^* for β and the proof follows.) The only case where tradability does not affect the discounted revenue is where $p = 1$ or $\beta = 1$.

Comments: The logic is simple: The issuer starts to earn revenue in fiat money earlier with non-tradable tokens than with tradable tokens.¹³ For the first M periods, all agents under both types of ICOs have at least one token, and p percent of them use it each period. But starting after period M , a rising fraction of agents in the non-tradable ICO have no coins, and thus need to use fiat money for platform consumption. Under a tradable ICO, agents who hit zero coins can buy coins from agents who have two tokens or more; in fact, all agents have at least one coin for the first $\frac{M-1}{p}$ periods. Thus revenue from fiat money is more backward loaded with tradable issuance than non-tradable issuance, and hence has a lower present value.

For only in-platform use, tradable tokens are strictly dominated by non-tradable tokens in both revenues from token issuance and revenues from fiat money. To justify tradability, there must be additional benefits outside our model. Of course, it is true that in our setup, we are neglecting several potential merits of tradability. First, tradability makes the tokens liquid, and potentially would allow the platform to pay a lower return due to a liquidity premium, albeit one that is likely lower than on fiat currency. Second, we have been assuming risk neutrality; if agents are risk-averse in period utility, there would again be gains to tradability. Third, we have eliminated the possibility that the tradable token can be used at other platforms or peer-to-peer

¹³With tradable tokens, no consumer would pay fiat money to the platform until period $\frac{M-1}{p}$. If tokens are not tradable, a “lucky” consumer can spend all M tokens before period $\frac{M-1}{p}$ and thereafter pay fiat money for the platform consumption.

transfer. There are many crypto-exchanges that provide services for token trading, for example, Coinbase or Bitpanda. If tradability allows broader use of the token, which might be translated into a higher p in our model, this again could be an advantage of tradability that is outside the scope of our model.

3.4 Consumer Surplus

Token issuance yields gains from trade given that the platform's rate of return exceeds that of retail bank consumers. However, the quantity-price tradeoff implies that the issuer is unable to claim the entire surplus, and must share part of surplus from token issuance with consumers. For completeness, and since the issue might be of significance to regulators, here we explicitly derive consumer surplus.

Consumers derive utility $\frac{\beta p}{1-\beta}$ from consumption but can pay for it in either tokens or fiat money. The consumer's spending consists of two parts: token spending today, and the expected expense in fiat money after depletion of tokens. For the non-tradable tokens, we can also write the consumer surplus as the total willingness to pay for the first M tokens minus the cost of purchasing the tokens.

For non-tradable ICO tokens,

$$\begin{aligned}
 CS_{I,N}(M) &= \underbrace{\frac{\beta p}{1-\beta}}_{\text{Consumption Utility}} - \underbrace{M \times P_{I,N}}_{\text{Token Spending}} - \underbrace{\left(\frac{\beta p}{1-\beta(1-p)}\right)^M \left(\frac{\beta p}{1-\beta}\right)}_{\text{Fiat Money Spending}} \\
 &= \underbrace{\sum_{i=1}^M \left[\frac{\beta p}{1-\beta(1-p)}\right]^i}_{\text{Consumption Utility from First M Goods}} - \underbrace{M \times P_{I,N}}_{\text{Token Spending}}
 \end{aligned}$$

For tradable ICO tokens,

$$CS_{I,T}(M) = \frac{\beta p}{1-\beta} - M \times P_{I,T} - \beta^{\frac{M-1}{p}} \left(\frac{\beta p}{1-\beta(1-p)}\right) \left(\frac{\beta p}{1-\beta}\right)$$

Almost parallel to the Proposition 2, we can easily show that tradable tokens are preferred by consumers: $CS_{I,T}(M) > CS_{I,N}(M)$.¹⁴ Consumers benefit from tradability from both lower token price paid today and the later fiat-money spending on expectation in future. In our model, consumers always prefer a free market while the issuer benefits from more restrictions on consumers for a more favorable split of welfare gain. $CS_{I,N}(M)$ and $CS_{I,N}(M)$ fully incorporate

¹⁴The only difference is that consumers discount the fiat-money spending with β , rather than β^* on the issuer's side. $\left(\frac{\beta p}{1-\beta(1-p)}\right)^M \left(\frac{\beta p}{1-\beta}\right) > \beta^{\frac{M-1}{p}} \left(\frac{\beta p}{1-\beta(1-p)}\right) \left(\frac{\beta p}{1-\beta}\right)$ holds by Proposition 1.

the β convexity effect we discussed earlier, implying the price effect dominates for consumers.

3.5 Non-tradable ICO with a Price Menu

We now consider the possibility that instead of selling all tokens at the same price, the platform is allowed to offer a menu that relates the total price paid to the number of tokens sold. Consumers are able to get a lower average price, the more tokens they buy. In this case, it is easy to show that the firm can garner all the gains from trade and leave zero consumer surplus.

To derive the its optimal price menu, the firm makes use of eq.(2) which gives the marginal value of the M^{th} coin. For M tokens, it charges

$$\sum_{i=1}^M \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^i$$

In this case, as $M \rightarrow \infty$, the platform gets the maximum possible discounted revenue (first-best) out of consumers

$$\lim_{M \rightarrow \infty} \sum_{i=1}^M \left[\frac{\beta p}{1 - \beta(1 - p)} \right]^i = \frac{\beta p}{1 - \beta}$$

Consumers get zero surplus since the platform can design a price menu so that consumers are indifferent along the menu. Thus, the design of price menu pushes up the average token price and therefore total platform revenue corresponding to any given token issuance.^{15,16}

$$P_{I,PM} = \frac{\beta p}{1 - \beta} \left[1 - \left(\frac{\beta p}{1 - \beta(1 - p)} \right)^M \right] \frac{1}{M} \quad (12)$$

It is quite straightforward to prove that a price menu approach adds nothing when the tokens are tradable: if the platform sells at a lower average price to bulk, token buyers, it cannot prevent the arbitragers from making a profit through resale. (Nor can it stop coalitions of consumers from buying in bulk to get a lower average price.)

¹⁵If consumers can share information efficiently and ship products at a low cost, one customer can arbitrage along the price menu by aggregating demand from other consumers, buying a large quantity from the issuer at a low price, and shipping products to others. In reality, customers are willing to accept price discrimination if it is costly to arbitrage in the product market or if the purchased service is not transferable at all. For example, users have to link their Uber accounts to their cell phone numbers. Thus the nature of some businesses can make a price menu approach feasible in reality. However, for some durable goods, say an iPhone, the shipment cost is quite small compared with the product value.

¹⁶ $\frac{\beta p}{1 - \beta} \left[1 - \left(\frac{\beta p}{1 - \beta(1 - p)} \right)^M \right] \frac{1}{M} > \left(\frac{\beta p}{1 - \beta(1 - p)} \right)^M = P_{I,N}$

3.6 Credibility of One-Time ICOs

It is important to emphasize that in the ICO case we have considered until now, the platform will be tempted to issue more tokens as the supply dwindles. Obviously, this could devalue the existing base in the case of tradable tokens, but the possibility also turns out to be relevant even in the non-tradable case. As we shall see in the next section, the consumer may regret having bought so many tokens in the ICO and wish instead she had earned interest on her savings and purchased more tokens later. Put differently, expectations of future issuance affects the shadow price at which the implicit value of tokens will rise. Thus in both cases, traded and non-traded ICOs, expectations of future issuance affects the current value of tokens. Our analysis, therefore, assumes not only that the platform is credibly able to commit to the purchasing power of the tokens when tendered, but also to its issuance strategy.

As a private sector business subject to national laws and courts, a platform may have at its disposal some devices for enhancing credibility not available to a sovereign currency issuer. For example, it can be subordinated to a regulator who ensures that the platform's "white paper" describing its token issuance policy cannot be violated without severe penalties. The ICO tokens can be made senior to other debt and to any tokens issued in the future. Similarly, the platform is legally bound to honor the fiat currency value of its tokens, so the exchange rate is fixed. Obviously, many subtleties and nuances are surrounding all these issues, and the issue of credibility is fundamental.

In the next section, we consider a class of richer issuance strategies, beyond a one-time ICO, that helps illustrate some of these points.

3.7 Non-tradable/tradable ICO+SCO (Price Only)

Until now, we have considered one-time token issuance strategies that, over time, lead to a shrinking supply of tokens as consumers redeem them for in-platform purchases. A true prototype currency would not self-extinguish, particularly if the issuer wants to maintain the possibility of eventual use outside the platform. In this section, we introduce the possibility that after the initial ICO, the platform commits to subsequently engaging in routine "SCO" (seasoned coin offerings) sufficient to maintain a constant steady-state supply of tokens. Although we are going to continue to assume that the platform can credibly commit to its issuance strategy, understanding how the expectation of ongoing sales affects the price of the initial ICO is also relevant to understanding how lack of credibility might affect initial issuance and price.

The introduction of SCOs turns out to change the calculus of token issuance quite fundamentally. In particular, we will demonstrate here the strong result that if SCOs are used to

maintain a constant token supply of tokens, then the maximum number of coins consumers will hold is one per person. This result is the same whether tokens are tradable or not, and in fact the tradable and non-tradable cases become equivalent. Importantly, this result applies only to the kind of “memoryless tokens we have been considering so far; later we will introduce the possibility that that platform can condition future token sales to individuals on past purchases.

We begin with the case of non-tradable tokens. One new question in this scenario is how to set the price in the SCO. In principle, there are three issuance strategies. (1) A no information policy, in which all consumers are offered the same price in every SCO regardless of their history in purchasing tokens and spending them. (In Section 5, we consider SCO issuance with memory.) (2) A history-dependent policy where the platform can charge a price for SCO tokens that is a function of the consumer’s entire history with the platform. (3) A Markov policy where issuance depends only on the consumer’s current account information (holdings of tokens) but is not path dependent. Policies (2) and (3) may seem very “un-money like” but actually incorporate the richer possibilities that digital currencies offer, ideas that are seldom considered in contemporary policy discussions say, about how retail central bank digital currencies might replace paper currency.

In this section, we will focus on the “no information” policy that is perhaps most likely to be regulatory-compliant and least likely to run afoul of privacy concerns, though later we will consider strategies with memory. The no information SCO strategy has very stark implications. The basic problem the platform faces is that for any steady state M it tries to sustain, consumers will only be willing to hold excess coins – more tokens that can be spent in one period – if they anticipate that the price will be rising over time at the consumer’s implicit interest rate (β^{-1}). But this is only possible in equilibrium if the quantity of tokens is falling over time – which is a contradiction unless the excess coins yield sufficient transactions convenience services, which we are abstracting from throughout most of this paper. Thus, the only steady-state coin holding has every consumer entering each period with exactly one token. At the end of the period, the platform will offer p tokens per person at price

$$P_S = \frac{\beta p}{1 - \beta(1 - p)} \quad (13)$$

which of course corresponds to the ICO price at $M = 1$. In a sense, we might refer to this equilibrium as a “token-in-advance” model, since the consumer is always using tokens for platform purchases. The same result holds with tradable tokens.

Proposition 3 (Token-in-advance Theorem): In any equilibrium with a constant supply M of tokens, and with memoryless issuance strategy, $M = 1$ regardless of tradability.

Proof of Proposition 3:

First, the token price needs to be constant in every period. If the token price is expected to appreciate indefinitely at the interest rate, one token will eventually worth more than the market value of the platform. However, the token price cannot exceed one because the value of underlying consumption is bounded (at one by assumption). With a constant token supply, the price must also be constant. Therefore consumers cannot get capital gains to substitute for explicit interest payments, and the equilibrium token supply (per capita) cannot exceed one in both tradable and non-tradable issuance.¹⁷

Define R_S as the total revenue of the ICO + SCO (“token-in-advance”):

$$R_S = \underbrace{\frac{\beta p}{1 - \beta(1 - p)}}_{\text{Token revenue today}} + \underbrace{\frac{\beta^* p}{1 - \beta^*} \frac{\beta p}{1 - \beta(1 - p)}}_{\text{Token revenue in future}}$$

The issuer nevertheless gains a higher discounted revenue from issuing one token than with no tokens.

$$R_S - \frac{\beta^* p}{1 - \beta^*(1 - p)} = \frac{(\beta - \beta^*)p}{(1 - \beta + \beta p)(1 - \beta^*)} > 0$$

Comments: Despite being able to earn ongoing revenues, the platform can only sell one token in the first period, and only $p < 1$ tokens per period thereafter. Of course, from the point of view of credibility, this ICO+SCO equilibrium might be easier to implement than the ICO. Also, the “token in advance” model might be more viable in an environment where consumers face liquidity constraints.

3.8 Comparison of ICO+SCO with Non-tradable ICO

From a revenue perspective, an ICO+SCO allows the issuer to secure ongoing token revenue from all future SCOs, rather than only from the one-time initial ICO. Also important is the fact that by releasing tokens more slowly, the platform will be able to get a higher (un-discounted) average price, that is, garners a larger share of the gains from trade. The disadvantage, of course, will be the expectations of future SCO issuance limits how much the platform can front-load revenue into the initial ICO. A natural question is whether a non-tradable ICO issuance mechanism can beat the simple ICO+SCO with the “token-in-advance” constraint. This involves comparing the discounted revenue of the non-tradable ICO:

¹⁷No consumer would buy a second token since any consumer would prefer to invest in the risk-free asset and wait to buy a new token from the market or from the issuer after the next consumption shock arrives.

$$R_{I,N} = \underbrace{M\left(\frac{\beta p}{1 - \beta(1-p)}\right)^M}_{\text{Revenue today}} + \underbrace{\frac{\beta^* p}{1 - \beta^*}\left(\frac{\beta^* p}{1 - \beta^*(1-p)}\right)^M}_{\text{Revenue in future}}$$

Proposition 4 (ICO versus ICO+SCO Dominance): Under optimal issuance, the non-tradable ICO dominates ICO+SCO if β^* is sufficiently low. When the consumption probability p is low ($p \rightarrow 0$) or β^* is high ($\beta^* \rightarrow \beta$), an ICO+SCO dominates the non-tradable ICO. (See Appendix for the proof.)

Comment: The tradeoff between ICO and ICO+SCO essentially depends on the size of gains from ICO token issuance. For any parameters, it is easy to compare numerically with the closed-form expression of $R_{I,N}$ and R_S . Proposition 4 studies the dominance in three extreme cases: When p is too small or the issuer's discount factor β^* is close to the consumer's discount factor β , the issuer may prefer to do the ICO+SCO issuance mechanism because the issuer can benefit from “token-in-advance” in every future period while the issuer cannot benefit a lot from the large-quantity issuance in the ICO. For example, one can show that the ICO+SCO strictly dominates the non-tradable ICO in the parameter space where the optimal issuance quantity is 2 (two) under the non-tradable issuance.¹⁸ The non-tradable ICO will dominate the ICO+SCO when the issuer is impatient enough (β^* is small). The benefit of the front-loading cash flow can be sufficiently large to offset the loss of future SCO revenue.

Certainly, the ICO + SCO with a constant steady supply of tokens and a constant steady-state price is much simpler than the optimal one-time ICO. It is also straightforward to show that for reasonable parameters, it can be supported as a trigger strategy equilibrium if we relax the no commitment assumption. Indeed, if the platform lacks credibility, the equilibrium can devolve to the ICO + SCO case (“token-in-advance”). However, absent credibility issues and when p and β are both close to one, the one-time non-tradable token ICO is much more profitable.

¹⁸ We compute the revenue of a non-tradable ICO minus the revenue of an ICO+SCO when the optimal token issuance is 2 under the non-tradable ICO. The revenue gap is strictly negative when $\beta^* < \beta$

Denote $a = \frac{\beta p}{1 - \beta(1-p)}$ and $a^* = \frac{\beta^* p}{1 - \beta^*(1-p)}$

$$h(\beta^*, \beta | M^* = 2) = 2a^2 + \frac{a^*}{1 - a^*} a^{*2} - a - \frac{a}{1 - a} a = a^{*2} + \frac{a^*}{1 - a^*} (a^{*2} - a) = \frac{a^*}{1 - a^*} (a^* - a) < 0$$

4 Assumptions Revisited

4.1 Runs and Interest Payments

As the model is constructed, the platform tokens are not subject to runs because agents tender their tokens if and only if a consumption shock hits, and the good is assumed non-storable. Of course, in reality, the offerings of platforms such as Alibaba and Amazon cover a wide range of durable goods, which opens the possibility of having a panic with say, consumers using their tokens to buy durable they do not yet need, despite storage costs. The platform can deal with runs in standard fashion, for example, by reserving the right to suspend sales temporarily, but the point is that even commodity-backed platform currencies are not immune to runs absent a fully-credible outside guarantee.¹⁹ Of course, in principle, the proceeds from token sales can be deposited in low return, but highly liquid, government securities. The platform could have a guaranteed refund in fiat currency if it were to temporarily stock out of goods in any given period; then, however, it would enjoy much smaller profits from token issuance.²⁰

Some crypto-currencies have indeed adopted a business model of setting a fixed exchange rate and claiming to hold all assets in treasuries, with the idea of making a profit by selling at par, paying zero interest, and then making a profit from the interest-bearing government assets. This approach, of course, has its fragilities. First and foremost, once international government regulation requires these assets to be easily traceable by governments and fully compliant with tax evasion and anti-money laundering laws, it is not at all clear that consumers will recognize any “convenience yield”. Second, even the most efficient crypto-currencies require considerable business costs to run. Last but not least, they are ultimately subject to the same kinds of fragilities as fixed exchange rate currencies and currency boards, where even slight temporary illiquidities or fiscal weakness can lead to an immediate attack. (See, for example, Obstfeld and Rogoff (1995).)²¹

Another approach for the platform would be simply to create an outside bank to handle its tokens, aiming to combine or leverage its token issuance business with a standard bank-like lending business. This approach would thereby create a chaebol or keiretsu-like structure which might allow the platform to use data across businesses to create synergies. The competition between tech companies and banks is a critical area but beyond the scope of this paper. Our narrow point here, though, is that the ability of chaebol and keiretsu to back tokens with

¹⁹The classic reference on pure multiple equilibrium bank runs is Diamond and Dybvig (1983).

²⁰The platform can also adopt a policy of suspending service in a stockout to discourage runs.

²¹For a discussion of the fragility of crypto-currencies, see Rogoff (2016).

platform goods does not necessarily constitute a significant advantage in itself.

Finally, we note that in principle, platform tokens can pay interest “in-kind” (in tokens) rather than in fiat currency. In particular, suppose tokens pay interest equal to $\frac{1-\beta}{\beta}$ on an ICO of $\frac{\beta p}{1-\beta}$ tokens, which could be tradable. This policy is sustainable since it involves paying out p tokens per period, exactly enough to replace tendered coins, assuming no runs. In a sense, this is a different implementation of lifetime memberships. Another important interpretation of the interest-bearing token we have just detailed is as a “security token” where effectively the consumer owns a share of the platform, with payments in services. We leave “security tokens” for future research.

Why then, shouldn’t the platform always make its tokens interest-bearing, or perhaps security tokens per the above example? There are at least a couple of reasons. First, and perhaps of the greatest concern in practice, is that the taxation and regulation of interest-bearing tokens may be very different than non-interest bearing tokens, bringing the platform issuance under banking and/or securities market regulation, with different results in different jurisdictions. Second, in a more general model with uncertainty, the required interest rate will fluctuate. And if the token market is relatively illiquid, it may be difficult to calibrate the interest rate required to fulfill the platform’s initial pledge to pay market interest. Uncertainty and risk also make the pledge of paying market interest challenging to accomplish, potentially opening up the platform to legal issues if not fully protected by (currently-nonexistent) regulation. In general, paying interest generates a different class of credibility issues, which are some cases may be more difficult to navigate. In sum, if the platform can pay market interest, either in fiat currency (using resources outside the platform) or in-kind (using only platform resources), it can potentially issue a tradable token that avoids some of the pitfalls identified here. However, this approach also has its own challenges, but a comparison is beyond the scope of this paper, where we focus mainly on non-interest bearing tokens.

4.2 Non-zero Cost of Input Goods

Assumption 2 posits zero cost of goods so that the entire revenue converts into platform profit. We relax this assumption by allowing X proportion of platform sales to be attributable to the input cost of goods. In this case, potential token demand is equal to gross sales by the platform each period, and not just net revenues.

The logic is straightforward: Token issuance adds financial income at the scale of gross revenue, which can be much larger than the size of profit from net platform revenue. For example, if an online retailer platform has a profit margin of 5%; the platform can issue tokens

with denomination 20 dollars for each one-dollar profit. If the platform can create an interest return wedge of 3%, the value-added from token issuance will be 0.6 dollars, which account for 60% increase in the platform profitability.

The pricing equations and issuance policy results remain the same as before when we relax the zero cost of goods assumption, except that token prices are proportional to gross sales, not net platform profit. The token prices only depend on the consumption probability, the sale price of the commodity, and the effective discount rate. Thus the breakdown in cost and profit does not affect the willingness to pay (WTP) for tokens. Thus, the value-added of token issuance is wholly determined by the revenue and not affected by the cost of goods.

The only change to the analysis from introducing non-zero input costs is to leverage up the present value of the platform's profits from token sales. We consider the maximum leverage effect from the benchmark platform value without token issuance to the first-best platform value.

The present value of platform without digital currency

$$\frac{\beta^*}{1 - \beta^*} pX$$

Under the first-best, the present value of platform profit is

$$\frac{\beta}{1 - \beta} p - \frac{\beta^*}{1 - \beta^*} p(1 - X)$$

The value to the platform of being able to leverage token issuance can be as high as

$$Leverage(X) = 1 + \frac{\beta - \beta^*}{\beta^*(1 - \beta)} \frac{1}{X}$$

where $Leverage(X)$ is monotonically decreasing in X and β^* (increasing in the platform investment return), and orthogonal to the consumption probability p . A low X in practice makes the token issuance to be spectacularly attractive for the online platforms with voluminous transactions but low profitability. Thus, in principle, token issuance has great potential when internet companies become financial service providers. Of course, as leverage increases, credibility problems become exacerbated and the platform becomes more vulnerable to runs per our earlier discussion.

4.3 Convenience Yield

One important potential merit of platform tokens is in providing a convenience yield for the token holders.²² In the money-in-the-utility model (Sidrauski 1967), utility is increasing and strictly concave in real money balances. In our model, a convenience yield would directly affect the token price by changing the effective discount rate. A larger convenience yield for consumers clearly benefits the platform which can then discount its tokens by less.

Convenience yield might be able to justify the issuance of tradable tokens if tradability brings greater convenience for transactions. For example, suppose tokens can be used as a digital unit to transfer money among consumers. In this case,

$$\beta(M, N) < \beta(M, T)$$

When tokens are more convenient to use, token holders are effectively more patient when holding tokens and willing to pay more fiat money in exchange for them. Where the government allows it, and where a single firm has dominant share across a large range of the economy, it is possible in principle that a platform currency could yield a significant enough convenience yield to compete with a government currency. For example, Alipay's success in China, particularly in Hangzhou, brings together payments for online shopping, restaurants, investment funds - even public transportation - into one unified digital payment system. This convenience has persuaded the younger generation to start keeping a large proportion of their savings in their Alipay's accounts. Tradable digital currency has great potential to be much more convenient than cash if the infrastructure is appropriately built.

Analytically, a convenience yield is quite straightforward to incorporate into our model if it is linear in token holdings (it simply modifies the consumer's discount factor β). A more general treatment, allowing for decreasing returns, would be more challenging. In any event, our read of the centuries-old history of money is that the government may initially allow or even foster private innovation in transaction technology, but eventually the government regulates and appropriates.²³

5 Money Memory

Up until now, we have shown that for a given level of token sales M in a one-time ICO, a platform will earn a larger profit from a non-tradable token than from a tradable token, and a

²²For example, payment with Amazon credit can be settled immediately, rather than going through credit card verification.

²³Rogoff, 2016.

larger profit still with a “buy more, save more” price-menu approach. If instead, the platform attempts to maintain a constant supply of outstanding coins with the ICO + SCO, the tradable and non-tradable cases turn out to be equivalent; whether or not the ICO + SCO can beat the one-time non-tradable ICO depends on β and p .

An important potential feature of a platform-backed currency – and in principle for any digital asset – is memory. A platform can fully observe a consumer’s account information, the full history of the account, and even the entire transaction history for each token. Making use of this information, a platform can design a mechanism for the SCO that induces consumers to hold $M > 1$ despite knowing that there will be future SCOs to replenish their stock.

We show that profitability of the ICO + SCO can be considerably enhanced if the platform can impose restrictions that tie a consumer’s ability to purchase future SCO tokens at a favorable price to her past behavior. We consider in turn two simple mechanisms, one where the platform can design a SCO based on the full history of consumer’s actions (a history-dependent mechanism), and a second where the platform can only design a SCO using current account information (a Markov mechanism), but cannot price discriminate using the individual’s historical records. In practice, it might be hard to implement a full history-dependent mechanism for many reasons: costly data storage and processing, the complexity of the issuance design for each history, violation of privacy, or the consumer’s sense of fairness. A Markov mechanism might be easier to implement and also more acceptable to consumers.

In the extreme, the platform can sell one token in the initial ICO for $\frac{\beta p}{1-\beta}$, which is the entire present value of future platform consumption to the consumer, but then commit to distributing free tokens to any agent who tenders their token for consumption in any given period. This is, of course, tantamount a membership system where the lifetime dues are paid once and for all upfront.²⁴

Formally, consider a specific class of issuance policies where in any future SCO, a platform only issues tokens to consumers with $M - 1$ tokens. Denote a as the ICO token price, b as the SCO token price, and define token issuance mechanism (X, Y, a, b) where X is the amount of tokens in the account, Y is amount of tokens to buy.²⁵ It is easy to check that “Buy M tokens in the ICO, and buy one token after a consumption shock in SCO” is an equilibrium strategy for consumers.²⁶ Using the account information, a platform can collect full future revenue with

²⁴In Appendix 2.2, we consider rolling memberships for limited time horizons and provide closed-form solutions. We also show that the discounted revenue from rolling memberships generates quantitatively similar revenue to the price-menu case.

²⁵The extreme case discussed above can be written as follows: The price scheme of the ICO is $(0, M, \frac{\beta p}{1-\beta} \frac{1}{M}, b)$, $(0, x, \infty, b)$ if $x \neq M$. The price scheme of the SCO is $(M - 1, 1, a, 0)$, (x, y, a, ∞) if $x \neq M - 1$ or $y \neq 1$.

²⁶First, consumers are indifferent between buying M tokens or never buying tokens. Thus, consumers have no incentive to deviate to “Not buying at all”. Second, consumers cannot benefit from buying more or fewer tokens in the ICO since it costs more. Third, consumers would take free tokens in a SCO. Otherwise, consumers need

a finite amount of tokens.

5.1 History-dependent Issuance

A history-dependent issuance can achieve the first-best in the sense that a platform can punish any possible deviation from its issuance proposal. Under the case of perfect information, a history-dependent issuance policy enables the platform to gain full control of consumer choices.

We show that history-dependent issuance allows the maximum flexibility in the cash flow arrangement. Consistent with Kocherlakota (1998), memory expands the set of feasible allocations. To illustrate this point, we expand the Markov SCO price menu to a history-dependent SCO: If a consumer did not buy a token after a consumption shock before, the platform stops selling tokens to the consumer (that is (x, y, a, ∞) for any (x, y) pair); If a consumer buys a token after each shock in the history, the platform offer one token at price b (that is, SCO: $(M - 1, 1, a, b)$, (x, y, a, ∞) if $x \neq M - 1$ or $y \neq 1$). In a richer framework, a platform can design more sophisticated contingent issuance policies, but we leave this for future research.

We start from the participation constraint

$$Ma + \frac{\beta p}{1 - \beta} b \leq \frac{\beta p}{1 - \beta}$$

binds the minimum ICO price a with the maximum SCO price b . The minimum ICO price must be higher than the ICO token price with price menu.²⁷

With history-dependent issuance, a consumer will be immediately excluded from the token market once she chooses not to purchase after any consumption shock.²⁸ Thus, the “now or never” inter-temporal constraint restricts consumers to buy one token right after a consumption shock if and only if

$$(1 + \frac{\beta p}{1 - \beta})b \leq [\frac{\beta p}{1 - \beta(1 - p)}]^{M-1} \frac{\beta p}{1 - \beta}$$

The constraint implies that the SCO token price cannot exceed the consumption value of the marginal M^{th} token: $b \leq [\frac{\beta p}{1 - \beta(1 - p)}]^M$.²⁹ The minimum ICO token price is equal to the price under information-free price menu.

$$a = \frac{1}{M} \frac{\beta p}{1 - \beta} (1 - [\frac{\beta p}{1 - \beta(1 - p)}]^M) \tag{14}$$

to pay fiat money for consumption.

²⁷ $a > \frac{1}{M} \frac{\beta p}{1 - \beta} [1 - (\frac{\beta p}{1 - \beta(1 - p)})^M] = P_{I,PM}$

²⁸Under a Markov policy, a consumer can still stay in the token market with probability $1 - p$.

²⁹It is impossible to set the SCO price higher than the consumption value. Otherwise, consumers would prefer to pay with fiat money rather than buying tokens.

A history-dependent issuance essentially incorporates memory into each token issued to consumers. Each token is contingent on the sequence of past actions. The account history helps the platform to achieve all possible cash flow allocations. A digital currency with memory thus further improves the welfare of issuers; that is, data is extremely valuable for the issuer.

5.2 Markov Issuance (ICO+SCO)

Under a Markov issuance policy, the issuer can only design issuance based on current holdings, but cannot retrieve the full history of the consumer’s behavior. Consumers may gamble by procrastinating the purchase of the SCO token because the issuer cannot punish a deviation based on their entire history. To incentivize consumers to buy the SCO token after a consumption shock, the issuer must design an issuance policy that satisfies a new “no procrastination” constraint:

$$(1 + \frac{\beta p}{1 - \beta})b \leq \beta \underbrace{[(1 - p)(1 + \frac{\beta p}{1 - \beta})b]}_{\text{No Consumption Shock: Still Use Tokens}} + \underbrace{p(\frac{\beta p}{1 - \beta(1 - p)})^{M-1} \frac{\beta p}{1 - \beta}}_{\text{Consumption Shock Arrives: Return to Fiat Money}}$$

The left-hand side is to “purchase” a token at a price b right after a consumption shock. The right-hand side is the payoff of procrastinating one period: without another consumption shock (probability $1 - p$), a consumer can still purchase a token at the SCO price b ; if another consumption shock arrives (with probability p), a consumer can never buy any token in the future and must make purchases with fiat money.³⁰ The “no procrastination” constraint pins down the maximum SCO price.³¹

$$a \leq \frac{1}{M} \frac{\beta p}{1 - \beta} (1 - \frac{\beta - \beta p}{1 - \beta p} [\frac{\beta p}{1 - \beta(1 - p)}]^M) \quad (15)$$

Use of account information provides additional value to the platform in two ways.³² First, the Markov issuance policy allows the platform to commit to a lower future SCO price in order to boost the ICO price. But this only works if the platform can condition future sales on holdings.

³⁰The present value of future spending in fiat money is

$$\sum_{i=M}^{\infty} (\frac{\beta p}{1 - \beta p})^i = (\frac{\beta p}{1 - \beta(1 - p)})^M (\frac{1}{1 - \frac{\beta p}{1 - \beta(1 - p)}}) = (\frac{\beta p}{1 - \beta(1 - p)})^{M-1} (\frac{\beta p}{1 - \beta})$$

³¹The upper bound of the SCO price is lower than the consumption value of the M^{th} token.

$$b \leq \frac{\beta - \beta p}{1 - \beta p} [\frac{\beta p}{1 - \beta(1 - p)}]^M < [\frac{\beta p}{1 - \beta(1 - p)}]^M$$

³²An important caveat is that the analysis here assumes the platform can commit, if it cannot then, of course, it may be tempted to sell in later periods to consumers who choose not buy tokens initially.

Second, the platform can continue to engage in short-term borrowing by selling tokens after every consumption shock.³³

6 Heterogeneous Agents

In our framework, the consumption probability p is the cornerstone for the token price. The assumption maps into the reality that many technology companies take “Daily active users” (DAU) as a significant parameter to focus on. In this section, we relax the assumption of homogeneity and address the following three questions: Does consumer heterogeneity encourage or discourage token issuance? Is it more profitable to only cater to frequent consumers or to be more inclusive? Most importantly, does introducing heterogeneity overturn our conclusion that if tradability does not produce sufficient convenience yield, then platforms may find the issuance of non-tradable tokens more profitable?

Heterogeneity raises a number of issues including for example, how a platform can price discriminate. Our illustrative analysis suggests that in principle, however, heterogeneity will not overturn our core results. Nevertheless, we show that heterogeneity reduces the benefits of token issuance if the platform cannot price discriminate among consumers.

For simplicity, we assume a society consisting of half frequent buyers p_H and half infrequent buyers p_L . A platform aims to issue M tokens in total to all consumers, M_L per infrequent consumer at price P_L , and M_H per frequent consumer at P_H respectively. We define a pooling equilibrium (both types of buyers purchase a positive number of tokens at the same price) as the case where $P_H = P_L$. In separating equilibrium (or price discrimination equilibrium) the two types of consumers buy tokens at different prices (or one type of consumers stay out the token market entirely). The issuance quantity and consumption frequency follows

$$\frac{M_L + M_H}{2} = M$$

$$\frac{p_L + p_H}{2} = p$$

6.1 Non-tradable ICO with Price Only

If a platform cannot price discriminate, all consumers coordinate in a pooling equilibrium where price $\widetilde{P}_{I,N}$ is the same for everyone. The willingness to pay for the last token equals to

³³If we write down the token sale revenue after each consumption shock, the amount of tokens purchased in ICO and SCO is $M, 0, \dots, 0, M, 0, \dots, 0$ with “no information” issuance of price menu. The Markov issuance policy allows the platform to front-load cash flow in all SCO periods by $M, 1, 1, 1, 1, \dots$

the

$$\left(\frac{\beta p_i}{1 - \beta(1 - p_i)}\right)^{M_i} = \widetilde{P}_{I,N} \quad i \in \{H, L\}$$

To issue M tokens, the corresponding price $\widetilde{P}_{I,N}$ must satisfy:

$$\frac{\log(\widetilde{P}_{I,N})}{\log\left(\frac{\beta p_L}{1 - \beta(1 - p_L)}\right)} + \frac{\log(\widetilde{P}_{I,N})}{\log\left(\frac{\beta p_H}{1 - \beta(1 - p_H)}\right)} = 2M$$

To simplify notation, we define function $f(p) = \frac{1}{\log\left(\frac{\beta p}{1 - \beta(1 - p)}\right)}$ and the non-tradable token price can be written as the following:

$$\widetilde{P}_{I,N} = e^{\frac{2}{f(p_L) + f(p_H)}} M$$

After introducing heterogeneity, the new effective discount factor of non-tradable tokens is $e^{\frac{2}{f(p_L) + f(p_H)}}$. Proposition 5 compares the new effective discount factor with the discount factor under heterogeneity $\frac{\beta p}{1 - \beta(1 - p)} (= e^{\frac{1}{f(p)}})$ as first defined in Section 3.1.

Proposition 5 (Heterogeneity of Non-tradable Tokens): The token price with agent heterogeneity is lower than the token price with homogeneous consumers of the same average consumption probability. (See Appendix 1 for the proof.)

$$\widetilde{P}_{I,N} < P_{I,N}$$

Comments: Proposition 5 illustrates that agent heterogeneity leads to a lower average price for the same token issuance. We note, however, that the magnitude of price sacrifice caused by heterogeneity is not necessarily large since the curvature of function $\log\left(\frac{\beta p}{1 - \beta(1 - p)}\right)\left(\frac{1}{f(p)}\right)$ is small for β and p near one (since $\log(x)$ is approximately linear around $x=1$). The effect of heterogeneity on revenue from fiat currency revenue (after the consumer uses up all her tokens) is ambiguous. Regardless, the magnitude of impact on cash revenue is only a second-order effect.³⁴

³⁴ Cash revenue with heterogeneity:

$$\frac{1}{2} \frac{\beta^*}{1 - \beta^*} [p_H \left(\frac{\beta^* p_H}{1 - \beta^*(1 - p_H)}\right)^{M_H} + p_L \left(\frac{\beta^* p_L}{1 - \beta^*(1 - p_L)}\right)^{M_L}]$$

Cash revenue with homogeneity (where $p = \frac{p_L + p_H}{2}$):

$$\frac{\beta^*}{1 - \beta^*} p \left(\frac{\beta^* p}{1 - \beta^*(1 - p)}\right)^M$$

To quantify the impact of heterogeneity, we pick a set of parameters $p_H = 0.8, p_L = 0.4, \beta = 0.9, \beta^* = 0.8$ and plot the difference. The cash revenue difference is no more than 0.0005, while the total discounted revenue

Corresponding to our discussion of the homogeneous case, we next extend our analysis to the impact of agent heterogeneity on the following four mechanisms: a non-tradable ICO, a tradable ICO, a non-tradable ICO+SCO, and a tradable ICO+SCO. Lastly, we study a price menu mechanism.

6.2 Tradable ICO

With tradability, all consumers must receive the same token price in the ICO. Moreover, the token price must be expected to appreciate to generate the risk-free return required to induce agents of either type to hold more than one token. Frequent consumers gain more welfare surplus since they are more likely to use the tokens. The token price under heterogeneity is given by (See Appendix 2.3 for the derivation of the closed-form solution.):

$$\widetilde{P}_{I,T} = \beta^{\frac{M-1}{p}} \left[(1 - \beta^\gamma (1 - P_L)^\gamma) \frac{\beta p_L}{1 - \beta(1 - p_L)} + \beta^\gamma (1 - P_L)^\gamma \frac{\beta p_H}{1 - \beta(1 - p_H)} \right]$$

where

$$\gamma = - \left\lfloor \frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)} \right\rfloor$$

One can show two results: First, similar to the non-tradable case, heterogeneity reduces the token price for tradable tokens. With tradability, the token price must appreciate at the rate of interest regardless of the distribution of consumption probabilities.

Proposition 6 (Heterogeneity of Non-tradable Tokens): When $M = 1$, the token price with heterogeneity is lower than the price with homogeneity. (See Appendix 1 for the proof.)

$$\widetilde{P}_{I,T} < P_{I,T}$$

Comments: Proposition 6 reveals that the token price under heterogeneity must be lower than the token price under homogeneity when only one token per person left in the economy. Heterogeneity can be viewed as a “friction” that limits the power of the platform to extract consumer surplus; frequent consumers retain positive surplus under token issuance. From the platform’s perspective, revenue from token issuance is reduced by the consumer surplus if price discrimination is not feasible. We will return to price discrimination later.

Second, the token price of the tradable ICO is still lower than non-tradable ICO with heterogeneity, even if prices are both lower than the case of agent homogeneity. Propositions 7 is 4 without token issuance. The difference from cash revenue is less than 1.25 basis points. We conclude that heterogeneity mildly discourages token issuance.

and 8 speak to the point that our conclusion about tradability is robust to introducing agent heterogeneity.

Proposition 7 (Effective Discount Factor Dominance with Heterogeneity): Under heterogeneity, the effective discount rate of non-tradable ICO tokens is still higher than that of tradable ICO tokens. (See Appendix 1 for the proof.)

$$\beta^{\frac{1}{p}} < e^{\frac{2}{f(p_L)+f(p_H)}}$$

Proposition 8 (ICO Price Dominance with Heterogeneity): When $M = 1$, the token price with tradability is lower than the non-tradable token price under heterogeneity (See Appendix 1 for the proof).

$$\widetilde{P}_{I,T} < \widetilde{P}_{I,N}$$

Comments: Proposition 7 is parallel to the Proposition 1 under the agent heterogeneity, implying that the tradable ICO price discounts faster than the non-tradeable ICO price as the quantity of tokens issued increases.³⁵ Proposition 8 proves that the tradable token price is lower than the non-tradable token price when $M = 1$ (Recall that token prices are the same for tradable and non-tradable when $M = 1$ under homogeneity). Proposition 8 is unique to the agent heterogeneity case since trading still occurs between high-type and low-type when less than one token circulates in the economy. Thus, as in the homogeneity case, the tradable token price is lower for any possible quantity of token issuance. Our core tradability result is robust to heterogeneity of consumption probabilities.

6.3 Tradable/ Non-tradable ICO+SCO

Similar to the homogeneous case, with tradability, there is no way to improve on a “token-in-advance” policy (selling a tokens one period ahead). Frequent consumers are willing to pay $\frac{\beta p_H}{1-\beta(1-p_H)}$ and infrequent consumers are willing to pay $\frac{\beta p_L}{1-\beta(1-p_L)}$. The new element here is having to choose between issuing to frequent consumers with a high price and issuing to everyone with a low price.

Pooling Equilibrium: Low price, broad consumer base: If the platform wants everyone to buy its tokens, the price needs to be $\frac{\beta p_L}{1-\beta(1-p_L)}$.³⁶

³⁵With tradability, the effective discount rate is always $\beta^{\frac{1}{p}}$ regardless of the consumption probabilities. Without tradability, the effective discount rate is lower with consumption heterogeneity as shown in Proposition 6. Thus, Proposition 7 is a tighter inequality than Proposition 1.

³⁶In the pooling equilibrium case, the platform may not want to issue the token one period ahead if

Under a pooling equilibrium, consumption heterogeneity makes the issuer worse off since the infrequent consumers drag down the token price.

$$\frac{\beta p_L}{1 - \beta(1 - p_L)} < \frac{\beta p}{1 - \beta(1 - p)}$$

Separating Equilibrium: High price, narrow consumer base: If the issuer only wants to cater frequent consumers only, the token price will be offered at $\frac{\beta p_H}{1 - \beta(1 - p_H)}$.³⁷ Intuitively, the platform should cater to frequent consumers (set a high price that only frequent consumers take up) only when there is a significant gap between the probabilities for high and low types. The issuer chooses the separating equilibrium if and only if

$$\frac{\beta p_L}{1 - \beta(1 - p_L)} < \frac{1}{2} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} + \frac{\beta^* p_L}{1 - \beta^*(1 - p_L)} \right)$$

Proposition 9 (ICO+SCO Revenue Dominance with Heterogeneity) : Heterogeneity reduces the discounted revenue of ICO+SCO issuance. (See Appendix 1 for the proof.)

$$\widetilde{R}_S < R_S$$

Comments: Proposition 9 verifies that the consumption probability heterogeneity causes the platform to earn strictly less revenue regardless of the issuance policy. Under a pooling equilibrium, infrequent consumers reduce the token price and make issuers unable to extract surplus from frequent consumers. Under a separating equilibrium, the issuer has to forgo half the population in the token issuance.

6.4 Price Menu Policies

The price menu mechanism enables the separating equilibrium where frequent consumers buy more tokens at a higher average price, and infrequent consumers buy fewer tokens at a lower

$$\frac{\beta p_L}{1 - \beta(1 - p_L)} < \frac{1}{2} \left(\frac{\beta^* p_L}{1 - \beta^*(1 - p_L)} + \frac{\beta^* p_H}{1 - \beta^*(1 - p_H)} \right)$$

³⁷The welfare gain from token issuance is only from frequent consumers,

$$\frac{1}{2} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} - \frac{\beta^* p_H}{1 - \beta^*(1 - p_H)} \right)$$

The revenue under separating equilibrium is

$$g(p_H, p_L) = \frac{1}{2} \left[\underbrace{\frac{1 - \beta^*(1 - p_H)}{1 - \beta(1 - p_H)} \frac{\beta p_H}{1 - \beta^*}}_{\text{Frequent Consumers with tokens}} + \underbrace{\frac{\beta^* p_L}{1 - \beta^*}}_{\text{Infrequent Consumers without tokens}} \right]$$

average price (or even excluded in the token market when p_L is small enough). We provide the condition for the separating equilibrium existence and its derivation in Appendix 2.4.

7 Conclusion

In this paper, we have studied to what extent large retailer platforms might have an advantage in issuing non-interest bearing digital tokens (currencies) by leveraging the fact that there are many consumers who are regular buyers, and who might find in-platform tokens appealing and convenient, while potentially both saving the platform fees paid to financial intermediaries as well as generating revenues of their own through a net interest margin.

Our core finding is that in many cases, it may be advantageous to the platform to issue non-tradable tokens rather than tradable ones, even if that means foregoing ideas of creating a prototype currency, unless the prototype currency can be expected to create significant convenience yield. Non-traded tokens give the platform the ability to implement more sophisticated pricing strategies (for example a price menu approach), and to incorporate memory features.

It is important to recognize that at the end of the day, a great deal depends on regulation, taxation, and other policy choices affecting not only technology companies but financial firms. Nevertheless, the simple benefit for platforms we look at here (net interest margin) is certainly an important one, especially if, as we assume, digital tokens give retail platforms access to the same kind of low interest-rate lenders that banks have so long profited from.

Our analysis has focused mainly on non-interest bearing tokens; if tokens can pay market interest, this can solve many of the problems we have analyzed, and this is certainly one solution. However, as discussed in the text, a pledge to pay market interest has its own issues, with implications for taxation, regulation, credibility, governance, and implementation. We have also illustrated the possibility of paying a market return by issuing security tokens, an important topic for future research.

The model presented here allows one to analyze a hierarchy of platforms depending on the frequency with which the consumer accesses them, and potentially also the size of transactions, and therefore how such differences might affect platform strategies when it comes to token/coin issuance. The huge range of crypto-currencies that have been issued to date, with ties to everything from social networking to real estate provide fertile ground for empirical analysis. The last part of our paper introduces a number of issues related to heterogeneity, which opens up a host of interesting questions for future research.

Reference

- Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia. "Bitcoin pricing, adoption, and usage: Theory and evidence." (2016).
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, Catherine Casamatta, and Albert J. Menkveld. "Equilibrium bitcoin pricing." Available at SSRN (2018).
- Budish, Eric. "The economic limits of bitcoin and the blockchain." No. w24717. National Bureau of Economic Research Working Paper (2018).
- Diamond, Douglas W., and Philip H. Dybvig. "Bank runs, deposit insurance, and liquidity." *Journal of Political Economy* 91, no. 3 (1983): 401-419.
- Demirg-Kunt, Asli, and Harry Huizinga. "Financial structure and bank profitability." *World Bank Policy Research Working Paper* 2430 (2000).
- DeJarnette, Patrick, David Dillenberger, Daniel Gottlieb, and Pietro Ortoleva. "Time lotteries and stochastic impatience." Working Paper (2019).
- Kiyotaki, Nobuhiro, and Randall Wright. "On money as a medium of exchange." *Journal of Political Economy* 97, no. 4 (1989): 927-954.
- Kocherlakota, Narayana R. "Money is memory." *Journal of Economic Theory* 81, no. 2 (1998): 232-251.
- Obstfeld, Maurice, and Kenneth Rogoff. "The Mirage of Fixed Exchange Rates." *Journal of Economic Perspectives* 9 (1998): 73-96.
- Prat, Julien, Vincent Danos, and Stefania Marcassa, "Fundamental pricing of utility tokens", Working Paper (2019)
- Rogoff, Kenneth S. *The Curse of Cash*. Princeton University Press, 2016.
- Schilling, Linda, and Harald Uhlig. "Some simple bitcoin economics." No. w24483. National Bureau of Economic Research, (2018).
- Sidrauski, Miguel. "Rational choice and patterns of growth in a monetary economy." *The American Economic Review* (1967): 534-544.
- Sockin, Michael, and Wei Xiong. "A model of cryptocurrencies." Unpublished manuscript, Princeton University (2018).

Appendix: Proposition Proofs

Proof of Proposition 1: Effective Discount Factor Dominance

To show $\beta^{\frac{1}{p}} < \frac{\beta p}{1-\beta(1-p)}$, we rewrite the inequality linearly as

$$\iff \beta p > \beta^{\frac{1}{p}} - \beta^{1+\frac{1}{p}}(1-p)$$

Then, we define a function $\omega(\beta)$ and show $\omega(\beta) > 0$ in the range of $p \in (0, 1)$:

$$\omega(\beta) = \beta p - \beta^{\frac{1}{p}} + \beta^{1+\frac{1}{p}}(1-p)$$

First, it is easy to find that $\omega(0) = 0$ and $\omega(1) = 0$. Then, we characterize $\omega(\beta)$ with the first-order and second-order derivatives:

$$\omega'(\beta) = p - \frac{1}{p}\beta^{\frac{1}{p}-1} + (1 + \frac{1}{p})\beta^{\frac{1}{p}}(1-p)$$

$$\omega''(\beta) = \frac{1}{p}\left(\frac{1}{p} - 1\right)\beta^{\frac{1}{p}-2} + (1 + \frac{1}{p})\frac{1}{p}\beta^{\frac{1}{p}-1}(1-p) = \beta^{\frac{1}{p}-2}\frac{1}{p}\left(\frac{1}{p} - 1\right)[1 - (1+p)\beta]$$

Note that $\omega'(0) = p > 0$ and $\omega'(1) = 0$. From the second-order derivative ($\omega'' = 0 \iff \beta = \frac{1}{1+p}$), we find that $\omega'(\beta)$ is monotonically decreasing when $\beta < \frac{1}{1+p}$ but increasing when $\beta > \frac{1}{1+p}$.

Then, we show the existence of a unique β such that $\omega'(\beta) = 0$. Existence: $\omega'(0) = p > 0$. $\omega''(1) < 0$ implies that $\omega'(1 - \epsilon) < 0$ where ϵ is a positive infinitesimal. By the continuity, there must exist a β such that $\omega'(\beta) = 0$. Uniqueness: If there is more than one root, say $0 < \beta_1 < \beta_2 < 1$, then $\omega'(\beta_1) = \omega'(\beta_2) = \omega'(1) = 0$. By the continuity, there must exist $\hat{\beta}_1$ and $\hat{\beta}_2$ so that $\omega''(\hat{\beta}_1) = \omega''(\hat{\beta}_2) = 0$ and $\beta_1 < \hat{\beta}_1 < \beta_2 < \hat{\beta}_2 < 1$. However, we know that $\beta = \frac{1}{1+p}$ is the only root for $\omega''(\beta) = 0$ in the range of $(0, 1)$. This violation implies a unique solution to $\omega'(\beta) = 0$.

Last, we show that $\omega(\beta) > 0$ holds when $\beta \in (0, 1)$. $\omega'(0) = p > 0$ implies $\omega(\epsilon) > 0$ for a positive infinitesimal ϵ . If there is a β_3 where $\omega(\beta_3) \leq 0$, we can find a $\beta_4 \in (\epsilon, \beta_3)$ so that $\omega(\beta_4) = 0$. $\omega(0) = \omega(\beta_4) = \omega(1) = 0$ implies at least two roots for $\omega(\beta) = 0$. It violates the uniqueness of the solution to $\omega(\beta) = 0$.

To give a graphical illustration, we plot the gap between the two effective discount factors as a function of β with $p = 0.9$ as in Figure 5.




Figure 5 plots the difference between effective discount factors for the non-tradable and tradable tokens as a function of β . The probability of consumption shock is set as $p = 0.9$.

Proof of Proposition 4: Non-tradable ICO dominance over ICO+SCO

We consider three corner cases $\beta^* \rightarrow 0$, $p \rightarrow 0$, and $\beta^* \rightarrow \beta$:

Case 1 $\beta^* \rightarrow 0$: $\max_M R_{I,N} - R_S = \max_M M \left(\frac{\beta p}{1-\beta(1-p)} \right)^M - \frac{\beta p}{1-\beta(1-p)} \geq [M \left(\frac{\beta p}{1-\beta(1-p)} \right)^M] | (M=1) - \frac{\beta p}{1-\beta(1-p)} = 0$

When $\beta^* \rightarrow 0$, ICO dominates because the issuer does not value future revenue at all. “Token-in-advance” limits the issuer’s ability to collect token revenue in the ICO period.

Case 2 $p \rightarrow 0$: The gain to issue the second token is only second-order (a constant multiplies p^2):

$$\left(\frac{\beta p}{1-\beta(1-p)} \right)^2 - \left(\frac{\beta^* p}{1-\beta^*(1-p)} \right)^2 \approx \left[\left(\frac{\beta}{1-\beta} \right)^2 - \left(\frac{\beta^*}{1-\beta^*} \right)^2 \right] p^2$$

But the loss is first order (a constant multiplies p):

$$\left(\frac{\beta p}{1-\beta(1-p)} \right)^2 - \left(\frac{\beta p}{1-\beta(1-p)} \right) \approx -\frac{\beta}{1-\beta} p$$

The first-order loss is larger than the second-order gain. Thus, the optimal non-tradable ICO token issuance is one. ICO+SCO strictly dominates non-tradable ICO with one token outstanding because SCO can frontload cash flow and increase revenue by $p \left[\frac{\beta p}{1-\beta(1-p)} - \frac{\beta^* p}{1-\beta^*(1-p)} \right]$ in every future period.

Case 3 $\beta^* \rightarrow \beta$: The gain from the second token issuance is close to zero. The loss is $\left(\frac{\beta p}{1-\beta(1-p)} \right)^2 - \left(\frac{\beta p}{1-\beta(1-p)} \right) \approx -\frac{\beta}{1-\beta} p < 0$. Thus, the optimal non-tradable ICO token issuance is also one. Similarly, the ICO+SCO dominates the non-tradable ICO in this case.

Appendix 1: Additional Proofs (Online Only)

Kenneth Rogoff and Yang You

October 30, 2019

Proof of Proposition 5: Negative Value of Heterogeneity in a Non-Tradable ICO

In this section of Appendix 1, we prove that heterogeneity reduces the token price of non-tradable ICO tokens. See Section 6.1.

Define $g(x) = \log(\frac{\beta x}{1-\beta(1-x)})$ and it is easy to show $g(x) < 0$, $g'(x) = \frac{1-\beta}{x(1-\beta+\beta x)}$, and $g''(x) = -\frac{(1-\beta)(1-\beta+2\beta x)}{x^2(1-\beta+\beta x)^2} < 0$

$f(x) = \frac{1}{g(x)}$ (defined in Section 6.1) and we have $f''(x) = -\frac{g''(x)g(x)^2 - 2g'(x)^2g(x)}{g(x)^4}$. Then, we show $f''(x) > 0$.

$$f''(x) \propto g''(x)g(x) - 2g'(x)^2 \propto -(1-\beta+2\beta x)\log(\frac{\beta x}{1-\beta+\beta x}) - 2(1-\beta)$$

Then, the convexity of $f(x)$ is equivalent to the inequality:

$$f''(x) > 0 \iff -\log(\frac{\beta x}{1-\beta+\beta x}) > \frac{2(1-\beta)}{1-\beta+2\beta x}$$

First, we show that the inequality holds for $x = 1$, that is

$$-\log\beta > \frac{2(1-\beta)}{1+\beta}$$

We use monotonicity to prove the inequality:

$$(\frac{2(1-\beta)}{1+\beta} + \log\beta)' = \frac{-2(1+\beta) - 2(1-\beta)}{(1+\beta)^2} + \frac{1}{\beta} = \frac{(1-\beta)^2}{(1+\beta)^2\beta} > 0$$

Thus $f''(x) > 0$ holds when $x = 1$:

$$\frac{2(1-\beta)}{1+\beta} + \log\beta < (\frac{2(1-\beta)}{1+\beta} + \log\beta)|_{\beta=1} = 0$$

Then, we show $-\log(\frac{\beta x}{1-\beta+\beta x}) > \frac{2(1-\beta)}{1-\beta+2\beta x}$ by proving the following function is monotonically

increasing:

$$\frac{d\left[\frac{2(1-\beta)}{1-\beta+2\beta x} + \log\left(\frac{\beta x}{1-\beta+\beta x}\right)\right]}{dx} = \frac{(1-\beta)^3}{(1-\beta+2\beta x)^2 x(1-\beta+\beta x)} > 0$$

Combining monotonicity and $f''(x=1) > 0$, we have

$$\frac{2(1-\beta)}{1-\beta+2\beta x} + \log\left(\frac{\beta x}{1-\beta+\beta x}\right) < \frac{2(1-\beta)}{1+\beta} + \log\beta < 0$$

$f(x)$ is a convex function on $p \in [0, 1]$. The convexity of $f(x)$ implies that heterogeneity reduces the token price:

$$\widetilde{P}_{N,I} = (e^{\frac{2}{f(p_L)+f(p_H)}})^M < \left(\frac{\beta p}{1-\beta(1-p)}\right)^M = P_{N,I}$$

Proof of Proposition 6: Negative Value of Heterogeneity in a Tradable ICO

In this Appendix 1 section, we demonstrate that heterogeneity lowers the price of non-tradable tokens. This proposition serves the analysis in Section 6.2. To prove the following inequality:

$$(1-\beta(1-P_L))^{-\frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)}} \frac{\beta p_L}{1-\beta(1-p_L)} + \beta(1-P_L)^{-\frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)}} \frac{\beta p_H}{1-\beta(1-p_H)} < \frac{\beta p}{1-\beta(1-p)}$$

We simplify the inequality by defining $p_H - p = p - p_L = \epsilon$,

$$\iff \beta(1-p_L)^{-\frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)}} \frac{\beta(1-\beta)2\epsilon}{(1-\beta(1-p_L))(1-\beta(1-p_H))} < \frac{\beta(1-\beta)\epsilon}{(1-\beta(1-p))(1-\beta(1-p_L))}$$

$$\iff 2\beta(1-p_L)^{-\frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)}} < 1 + \frac{\beta\epsilon}{1-\beta(1-p)}$$

$$\iff \log 2 - \log(\beta(1-p_L)) \frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)} < \log\left(1 + \frac{\beta\epsilon}{1-\beta(1-p)}\right) \quad (1)$$

Then, we define a function

$$w(\beta, p_L, p_H) = \log(\beta(1-p_L)) \frac{\log(1+\frac{p_L}{2p_H})}{\log(1-\frac{1}{2}p_L)} + \log\left(1 + \frac{\beta\epsilon}{1-\beta(1-p)}\right)$$

We prove inequality (1) above in following three steps:

Step 1: We show that function w is monotonically decreasing in β .

$$\frac{\partial w}{\partial \beta} = \underbrace{\frac{1}{\beta} \frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)}}_{\text{negative}} + \underbrace{\frac{\beta^2 \epsilon}{(1 - \beta(1 - p))(1 - \beta(1 - p_H))}}_{\text{positive}} < \frac{1}{\beta} \left(-\frac{1}{p + \epsilon}\right) + \frac{\beta^2 \epsilon}{\beta^2 p(p + \epsilon)} < 0$$

$$\text{Thus } w(\beta, p_L, p_H) > w(1, p_L, p_H) = \log(1 - p_L) \frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)} + \log(1 + \frac{\epsilon}{p})$$

Step 2: we show that function w is monotonically increasing in p_H , that is, equivalently we prove the first-order derivative is positive:

$$\begin{aligned} \frac{\partial w}{\partial p_H} &= -\frac{\log(1 - p_L)}{\log(1 - \frac{p_L}{2})} \frac{p_L}{(2p_H + p_L)p_H} + \frac{p_L}{(p_H + p_L)p_H} > 0 \\ &\iff \frac{\log(1 - p_L)}{\log(1 - \frac{p_L}{2})} > \frac{2p_H + p_L}{p_H + p_L} \end{aligned}$$

We know $\frac{2p_H + p_L}{p_H + p_L}$ is monotonically increasing in p_H . A sufficient condition for the positive first-order derivative is the following:

$$\begin{aligned} \frac{\log(1 - p_L)}{\log(1 - \frac{p_L}{2})} &> \frac{2 + p_L}{1 + p_L} \\ &\iff (1 + p_L)\log(1 - p_L) - (2 + p_L)\log(1 - \frac{p_L}{2}) > 0 \end{aligned}$$

Take the first-order derivative of the equation above,

$$\left[\log(1 - p_L) - \frac{1 + p_L}{1 - p_L}\right] - \left[\log(1 - \frac{p_L}{2}) - \frac{1 + \frac{p_L}{2}}{1 - \frac{p_L}{2}}\right]$$

To sign the first-order derivative, we define $h(x) = \log(1 - x) - \frac{1+x}{1-x}$, $x \in [0, 1]$ and study its monotonicity:

$$h'(x) = -\frac{1}{1-x} - \frac{2}{(1-x)^2} < 0$$

$h(p_L) > h(\frac{p_L}{2})$ implies that w is monotonically increasing in p_H . Then, we only need to show the inequality holds when $p_H \rightarrow p_L$, that is:

$$w(1, p_L, p_H) \geq w(1, p_L, p_L) = \log(1 - p_L) \frac{\log(\frac{3}{2})}{\log(1 - \frac{1}{2}p_L)}$$

Step 3: $\frac{\log(1-p_L)}{\log(1-\frac{p_L}{2})}$ is monotonically increasing in p_L .

$$\frac{\frac{1}{2}}{1-\frac{p_L}{2}} \log(1-p_L) > \frac{1}{1-p_L} \log(1-\frac{p_L}{2})$$

$$\iff (1-p_L) \log(1-p_L) > \frac{1-\frac{p_L}{2}}{\frac{1}{2}} \log(1-\frac{p_L}{2})$$

We show $k(x) = \frac{(1-xp_L)}{x} \log(1-xp_L)$ is monotonically increasing in $x \in [0, 1]$,

$$k'(x) = -\frac{1}{x^2} \log(1-xp_L) - \frac{p_L}{x} > \frac{xp_L}{x^2} - \frac{p_L}{x} = 0$$

$k(1) > k(\frac{1}{2})$ implies monotonicity of function $\frac{\log(1-p_L)}{\log(1-\frac{p_L}{2})}$,

$$\frac{\log(1-p_L)}{\log(1-\frac{p_L}{2})} > \lim_{p_L \rightarrow 0} \frac{\log(1-p_L)}{\log(1-\frac{p_L}{2})} = \frac{\frac{-1}{1-p_L}}{\frac{-\frac{1}{2}}{1-\frac{p_L}{2}}} = 2$$

Combining all three steps above, we can show that Appendix 1 eq.(1) holds:

$$w(1, p_L, p_L) \geq 2 \log\left(\frac{3}{2}\right) = \log\left(\frac{9}{4}\right) > \log(2)$$

Appendix 1 eq.(1) is equivalent to the inequality of Proposition 6, which implies the tradable ICO token price with heterogeneity is lower than the price with homogeneity.

Proof of Proposition 7: Effective Discount Factor Dominance with Heterogeneity

This Appendix 1 section revisits the effective discount factor dominance in Proposition 1 and proves that the effective discount factor is still lower when tokens are tradable. This proposition supports the analysis in Section 6.2. With heterogeneity, the effective discount factor of non-tradable tokens is $e^{\frac{2}{f(p_L)+f(p_H)}}$ while the discount factor of tradable tokens remains the same $\beta^{\frac{1}{p}}$: $e^{\frac{2}{f(p_L)+f(p_H)}} > \beta^{\frac{1}{p}}$ where $\frac{p_H+p_L}{2} = p$.

Case 1 $p < 0.5$: the convexity of $f(x)$ implies that the non-tradable discount factor is higher than the case where $p_L = 0, p_H = 2p$.

$$e^{\frac{2}{f(p_L)+f(p_H)}} \geq e^{\frac{2}{f(0)+f(2p)}} = e^{2 \log\left(\frac{2\beta p}{1-\beta(1-2p)}\right)} = \left(\frac{2\beta p}{1-\beta(1-2p)}\right)^2$$

Applying the formula in Proposition 1,

$$\left(\frac{2\beta p}{1 - \beta(1 - 2p)}\right)^2 > (\beta^{\frac{1}{2p}})^2 = \beta^{\frac{1}{p}}$$

Case 2 $p \geq 0.5$: the convexity of $f(x)$ implies that the non-tradable discount factor is higher than the case where $p_L = 2p - 1$, $p_H = 1$:

$$e^{\frac{2}{f(p_L)+f(p_H)}} \geq e^{\frac{2}{f(2p-1)+f(1)}}$$

After taking logs, we need to prove:

$$\frac{2}{f(2p-1) + f(1)} > \frac{1}{p} \log \beta$$

$$\iff (2p-1) \log\left(\frac{\beta(2p-1)}{1 - \beta + \beta(2p-1)}\right) > \log \beta$$

Similarly, applying Proposition 1,

$$(2p-1) \log\left(\frac{\beta(2p-1)}{1 - \beta + \beta(2p-1)}\right) > (2p-1) \log(\beta^{\frac{1}{2p-1}}) = \log \beta$$

For $p_L, p_H \in [0, 1]$ and $p = \frac{p_L + p_H}{2}$, the effective discount factor of non-tradable tokens still dominates that of tradable tokens in the heterogeneity case, that is

$$\beta^{\frac{1}{p}} < e^{\frac{2}{f(p_L)+f(p_H)}}$$

Proof of Proposition 8: ICO Price Dominance with Heterogeneity

This Appendix 1 section compares the tradable and non-tradable ICO price when only one token is issued ($M=1$). This proposition serves the analysis in Section 6.2. When only one token outstanding, the non-tradable token price is higher than the tradable token price if and only if:

$$(1 - \beta(1 - p_L))^{\frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)}} \frac{\beta p_L}{1 - \beta(1 - p_L)} + \beta(1 - p_L)^{\frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)}} \frac{\beta p_H}{1 - \beta(1 - p_H)} < e^{\frac{2}{f(p_L)+f(p_H)}}$$

Define $x = \frac{\beta p_L}{1 - \beta(1 - p_L)}$ and $y = \frac{\beta p_H}{1 - \beta(1 - p_H)}$. We know $0 < x < y < \beta < 1$.

We aim to rewrite the inequality as a function of x, y, β . Thus, we express p_L and p_H in x and y : $p_L = \frac{1 - \beta}{\beta} \frac{x}{1 - x}$, $p_H = \frac{1 - \beta}{\beta} \frac{y}{1 - y}$, $\frac{p_L}{p_H} = \frac{x(1 - y)}{(1 - x)y}$. The non-tradable token price is rewritten in x and y as the following:

$$\iff x + \beta(1 - p_L)^{-\frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)}} (y - x) < e^{\frac{2}{f(p_L) + f(p_H)}}$$

We subtract x from both side, take log, and replace p_L, p_H with x, y . Then, we move two terms with x only to the left-hand side, and the other two terms depending on both x and y to the right-hand side:

$$\iff \underbrace{\frac{\log(\beta(1 - \frac{1-\beta}{\beta} \frac{x}{1-x}))}{\log(1 - \frac{1}{2} \frac{1-\beta}{\beta} \frac{x}{1-x})}}_{LHS} > \underbrace{\frac{\log(y - x) - \log(e^{\frac{2}{\frac{1}{\log(x)} + \frac{1}{\log(y)}}} - x)}{\log(1 + \frac{1}{2} \frac{x(1-y)}{(1-x)y})}}_{RHS}$$

The *LHS* (left-hand side of the above inequality) is a function of (β, x) and the *RHS* (right-hand side of the above inequality) is a function of (x, y) . We first show the *RHS* is monotonically decreasing in y . However, it is hard to show it with algebra. We use a graphical proof with the three-dimensional surface in Appendix 1 Figure 1 and the two-dimensional plots in Appendix 1 Figure 2 to illustrate monotonicity.

The upper bound of *RHS* is reached with $y \rightarrow x$

$$\begin{aligned} \lim_{y \rightarrow x} \frac{\log(y - x) - \log(e^{\frac{2}{\frac{1}{\log(x)} + \frac{1}{\log(y)}}} - x)}{\log(1 + \frac{1}{2} \frac{(1-x)y}{x(1-y)})} &= \frac{\log(\lim_{y \rightarrow x} \frac{y-x}{e^{\frac{2}{\frac{1}{\log(x)} + \frac{1}{\log(y)}}} - x})}{\log(\frac{3}{2})} \\ &= \frac{\log(\lim_{y \rightarrow x} \frac{1}{(\frac{2}{(\frac{1}{\log(x)} + \frac{1}{\log(y)})^2}) (\frac{1}{\log(y)^2})^{\frac{1}{y}} e^{\frac{2}{\frac{1}{\log(x)} + \frac{1}{\log(y)}}}})}{\log(\frac{3}{2})} = \frac{\log(\frac{1}{\frac{\log(x)^2}{2} \frac{1}{\log(x)^2} \frac{1}{x}})}{\log(\frac{3}{2})} = \frac{\log(2)}{\log(\frac{3}{2})} \end{aligned}$$

The rest of the proof shows that $LHS > \frac{\log(2)}{\log(\frac{3}{2})}$. *LHS* is only about p_L and the minimum of *LHS* is achieved at p_L^* where p_L^* satisfies the F.O.C:

$$\frac{1}{\beta(1 - p_L^*)} \beta \log(1 - 0.5p_L^*) = \frac{0.5}{(1 - 0.5p_L^*)} \log(\beta(1 - p_L^*)) \rightarrow \frac{\log(\beta(1 - p_L^*))}{\log(1 - 0.5p_L^*)} = \frac{2 - p_L^*}{1 - p_L^*}$$

We substitute the F.O.C into the *LHS* to find the lower bound:


$$\frac{\log(\beta(1 - \frac{1-\beta}{\beta} \frac{x}{1-x}))}{\log(1 - \frac{1}{2} \frac{1-\beta}{\beta} \frac{x}{1-x})} = \frac{\log(\beta(1 - p_L))}{\log(1 - 0.5p_L)} \geq \frac{\log(\beta(1 - p_L^*))}{\log(1 - 0.5p_L^*)} = \frac{2 - p_L^*}{1 - p_L^*} \geq 2$$

A trivial fact is that


$$\frac{9}{4} > 2 \longrightarrow 2\log\left(\frac{3}{2}\right) > \log(2) \longrightarrow 2 > \frac{\log(2)}{\log\left(\frac{3}{2}\right)}$$

To summarize, all proofs above show the following inequalities hold for any pair (p_L, p_H) . It implies that the tradable ICO token price is still lower than the non-tradable ICO token price with agent heterogeneity.

$$LHS \geq 2 > \frac{\log(2)}{\log\left(\frac{3}{2}\right)} = \lim_{y \rightarrow x} RHS \geq RHS$$



Appendix 1 Figure 1



Appendix 1 Figure 2

Proof of Proposition 9: ICO+SCO Revenue Dominance with Heterogeneity

This Appendix 1 section shows that the platform earns less revenue with agent heterogeneity regardless of issuance policy. This proposition serves the analysis in Section 6.3.

Consider a small perturbation

$$g(p_H + \epsilon, p_L - \epsilon) - g(p_H, p_L) \propto \frac{1 - \beta^*(1 - p_H)}{1 - \beta(1 - p_H)} \beta - \beta^* > \frac{1 - \beta^*(1 - 2)}{1 - \beta(1 - 2)} \beta - \beta^* = \frac{\beta - \beta^*}{1 + \beta} > 0$$

That says, given only catering to frequent consumers, an issuer would prefer a more dispersed consumption probabilities.

$$g(p_H, p_L) < g(2p, 0) = \frac{1 - \beta^*(1 - 2p)}{1 - \beta(1 - 2p)} \frac{\beta p}{1 - \beta^*} < \underbrace{\frac{1 - \beta^*(1 - p)}{1 - \beta(1 - p)} \frac{\beta p}{1 - \beta^*}}_{\text{Revenue under homogeneity}}$$

The revenue from “Separating Equilibrium: High price, narrow consumer base” is strictly lower than the revenue when agents are homogeneous. As we have already showed, the issuer also earns a lower revenue from the “Pooling Equilibrium: Low price, broad consumer base”. Thus, the issuer is worse off with consumption frequency heterogeneity.

Appendix 2 (Online Only)

Kenneth Rogoff and Yang You

October 30, 2019

Appendix 2.1: Monotonicity of Non-tradable ICO Optimal Issuance

Appendix 2.1 shows that the optimal issuance quantity M^* of non-tradable ICO tokens is weakly decreasing in β^* when the optimal issuance is larger than one. We consider a local optimum M_l which satisfies eq.(7) and eq.(8). M^* is a local optimum that achieves maximum discounted revenue.

Define $a = \frac{\beta p}{1-\beta(1-p)}$. First, we show $F(M) = Ma^M - (M-1)a^{M-1}$ is monotonically decreasing for any integer in the range $M \in [1, \bar{M}]$ where \bar{M} is the largest integer so that $F(\bar{M}) \geq 0$.¹ We prove the monotonicity by showing $F(M) - F(M-1) < 0$:

$$F(M) - F(M-1) = a^{M-2}[Ma^2 - 2(M-1)a + (M-2)] = a^{M-2}[Ma - (M-2)](a-1)$$

$F(\bar{M}) \geq 0$ implies that $a \geq \frac{\bar{M}-1}{\bar{M}}$. We can show

$$[Ma - (M-2)] \geq M \frac{\bar{M}-1}{\bar{M}} - (M-2) \geq M \frac{M-1}{M} - (M-2) = 1 > 0$$

$(a-1) < 0$ by construction. Thus, we know $F(M) < F(M-1)$ for any $M \in [1, \bar{M}]$.

The optimality conditions eq.(7) and eq.(8) can be rewritten as

$$F(M_l) > \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{M_l}$$

$$F(M_l + 1) < \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{M_l + 1}$$

Given M_l , eq.(7) must still hold for any lower β^* . Eq.(8) might be violated, that is, issuing an extra token (the $(M_l + 1)^{th}$ token) is profitable. If eq.(8) holds, M_l is still a local optimum

¹ $F(M) > 0$ says that issuing an extra token generates a higher token revenue.

for the lower β^* . If violated, there must exist a $\widehat{M}_l \geq M_l + 1$ so that²

$$F(\widehat{M}_l) > \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{\widehat{M}_l}$$

$$F(\widehat{M}_l + 1) < \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{\widehat{M}_l + 1}$$

For any lower β^* , we can always find a larger local optimum \widehat{M}_l for each local optimum M_l .

In the last step, we show $M_l = M^*$, that is local optimum is unique and thus equivalent to the global optimum. In other word, eq.(7) and eq.(8) are the necessary and sufficient conditions for the optimal token issuance quantity.

The uniqueness of the local optimum is equivalent to the uniqueness of the root for $F(x) = [a^*]^x$ where $a^* = \frac{\beta^* p}{1 - \beta^*(1-p)} < a$ and $x \geq 1$. Then $M^* = M_l = \lfloor x \rfloor$. We know that $F(1) = a > a^*$ and $\lim_{x \rightarrow \infty} F(x) \approx 0 \approx \lim_{x \rightarrow \infty} a^x$. Thus, there is at least one root for $F(x) = [a^*]^x$.

Take logs on both sides:

$$\log(xa^x - (x-1)a^{x-1}) = x \log(a^*)$$

If there are two roots $x_1 < x_2$, we can find two roots x'_1 and x'_2 for the first-order derivative equality $\frac{d \log(xa^x - (x-1)a^{x-1})}{dx} = \frac{d(x \log(a^*))}{dx}$ where $x_1 < x'_1 < x_2 < x'_2 < \infty$.

Then we compute the first-order derivatives of both sides and illustrate that the equality has one root at most:

$$\frac{\log(a)(xa^x - (x-1)a^{x-1}) + a^x - a^{x-1}}{xa^x - (x-1)a^{x-1}} = \log(a^*)$$

$$\Leftrightarrow \log(a) + \frac{a-1}{xa - (x-1)} = \log(a^*)$$

$$\Leftrightarrow x = \frac{1}{1-a} - \frac{1}{\log(a) - \log(a^*)}$$

The equality of derivatives is essentially linear in x . Thus, it is obvious that the above equality has at most one root. This contradiction implies that $F(x) = [a^*]^x$ can not have two roots x_1 and x_2 . Thus, any M_l satisfies eq.(7) and eq.(8) must be the global optimum M^* .

Thus, a lower β^* (better outside investments) incentivizes the platform to issue more tokens.

²The violation of eq.(8) implies $F(M_l + 1) \geq \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^{M_l + 1}$. But for $M \rightarrow \infty$, the revenue gain from issuing an extra token must be negative, that is, $F(M) < 0 < \left[\frac{\beta^* p}{1 - \beta^*(1-p)} \right]^M$. From M_l to ∞ , another local optimum must exist.

Appendix 2.2: Rolling membership

In our framework, membership, as a long-term commitment, can also be used to back a platform token. The membership allows token holders to consume platform goods within the membership period. We have already considered the case of life membership in Section 3, which allows the platform to extract the maximum surplus from the consumer, but entails credibility issues. In Appendix 2.2, we consider shorter membership periods. A familiar example is Spotify. Spotify buys copyrights from music producers and provides membership to consumers. Consumers can listen to a range of music whenever they want within a limited period. One major advantage of membership is that a platform can lock in consumption for a certain period without having to provide an implicit interest rate in the form of a rising price. We show that membership issuance has similar properties to a price menu, albeit with certain differences we shall detail below.

Denote $P_{I,Mem}(N)$ as the price of membership token valid for N periods. $P_{I,Mem}(N)$ can be written as a combination of marginal claims as well:

$$P_{I,Mem}(N) = \sum_{n=1}^N \sum_{M \leq n} P_{n,M} = \sum_{n=1}^N \sum_{M \leq n} \beta^n \binom{n-1}{M-1} p^M (1-p)^{n-M} = \sum_{n=1}^N \beta^n p = \beta p \frac{1-\beta^N}{1-\beta} \quad (1)$$

One unit of membership token can satisfy all individual demand for N periods. We set $N = \frac{M}{p}$ to make each membership delivers M commodities in total, comparable to M tokens with an ICO or ICO+SCO. We consider two cases: a. A one-time fixed, finite-term ICO membership token, that is to issue membership tokens available for $\frac{M}{p}$ periods and then switch back to fiat money after. b. An ICO+SCO membership token, that is to issue membership tokens available for $\frac{M}{p}$ periods and then issue again whenever senior tokens expire.

For the ICO membership issuance, the total revenue is

$$R_{I,Mem} = \beta p \frac{1-\beta^{\frac{M}{p}}}{1-\beta} + \beta^{\frac{M}{p}} \frac{\beta^* p}{1-\beta^*}$$

For the ICO+SCO membership issuance, the total revenue is

$$R_{S,Mem} = \frac{P_{N,\infty}}{1-\beta^* \frac{M}{p}} = \frac{\beta p}{1-\beta} \frac{1-\beta^{\frac{M}{p}}}{1-\beta^* \frac{M}{p}}$$

The discounted values of ICO and ICO+SCO membership are close to the non-tradable ICO and ICO+SCO that employ a price menu mechanism correspondingly. With the ICO issuance, the membership token gains slightly more revenue from the token issuance, but slightly less

revenue in fiat money.³ With the ICO+SCO issuance, the revenue ratio of the membership issuance to the price-menu issuance is approximately one.⁴

$$R_{I,Mem} \approx R_{I,PM}$$

$$R_{S,Mem} \approx R_{S,PM}$$

The most prominent advantage of membership is the simplicity: a single price is sufficient to claim the entire surplus from consumers. Tradability does not matter in this scenario since everyone shares the same willingness to pay to depend solely on the valid periods remaining. The main caveat is the moral hazard; that is, consumers may over-consume with the membership or claim the goods and resale to other consumers without the membership. The membership currency is only applicable to industries with modest or zero marginal cost or non-transferable products, e.g., music and movies.

Appendix 2.3: Solution of Tradable ICO Price with Heterogeneity

This section solves the closed-form solution for the tradable ICO token price under heterogeneity as in section 6.2. The price path would be the following: Before period $\frac{M-1}{p}$, the token price appreciates at the interest rate. From period $\frac{M-1}{p}$ to the period when only frequent consumers hold tokens, infrequent consumers have at most one token left in hands and start to pay for the platform consumption with fiat money (The token usage speed also slows down). When only frequent consumers hold tokens ($\frac{1}{2}$ tokens remain in the economy), the token price will be $\frac{\beta p_H}{1-\beta(1-p_H)}$. The price path described above is the unique equilibrium for stable tokens under agent heterogeneity.⁵

³We compute the revenue difference:

$$R_{I,Mem} - R_{I,PM} = \left[\frac{\beta^* p}{1-\beta^*(1-p)} \right]^M - \beta^* \frac{M}{p} \left] \frac{\beta^* p}{1-\beta^*} - \left[\frac{\beta p}{1-\beta(1-p)} \right]^M - \beta \frac{M}{p} \right] \frac{\beta p}{1-\beta}$$

Define $h(\beta) = \frac{\beta p}{1-\beta(1-p)} \frac{M}{p} - \beta \frac{M}{p}$ and pick $p = 0.5$ and $m = 10$. The range of $h(\beta)$ is $(-0.015, 0]$. In this case, the revenue gap is no more than 1.5 percent of the total revenue difference between the first-best ($\frac{\beta p}{1-\beta}$) and the revenue with no token issuance ($\frac{\beta^* p}{1-\beta^*}$).

⁴We calculate the revenue ratio of the ICO + SCO tokens:

$$\frac{R_{S,Mem}}{R_{S,PM}} = \frac{1 - \beta \frac{M}{p} \left[1 - \frac{\beta^* p}{1-\beta^*(1-p)} \right]^M}{1 - \beta^* \frac{M}{p} \left[1 - \frac{\beta p}{1-\beta(1-p)} \right]^M}$$

The ratio $\frac{R_{S,Mem}}{R_{S,PM}} \approx 1$. Define a function $g(\beta) = \frac{\beta p}{1-\beta(1-p)} \frac{M}{p}$. For example, we pick $p = 0.5$ and $M = 10$, the ratio ranges from 1 to 1.02 when $\beta \in (0, 1)$. The narrow range of $g(\beta)$ bounds $\frac{R_{S,Mem}}{R_{S,PM}}$ close to 1.

⁵The infrequent consumers do not sell the last token to the frequent consumers before the period $\frac{M-1}{p}$ because they know the token will appreciate at the interest rate and they can get an extra benefit if a consumption shock hits in the next period. As long as the token is tradable, there is no reason for any consumer to stay out of the market. In the period $\frac{M-1}{p}$, is it possible that infrequent consumers have already sold all tokens to frequent

From period $\frac{M-1}{p}$ to the period when only frequent consumers hold tokens, infrequent consumers are indifferent between selling the token to frequent consumers or holding the token for future personal consumption, that is,

$$P_t = \beta[(1 - p_L)P_{t+1} + p_L] \quad (2)$$

The next step is to compute the number of periods until $\frac{1}{2}$ of all tokens are depleted after which no infrequent consumer holds any tokens. Define $x(t) \in (0.5, 1)$ as the quantity of tokens left in the economy. As long as there are any tokens being held by infrequent consumers, there must be $\frac{1}{2}$ held by frequent consumers and $x(t) - \frac{1}{2}$ left in the hands of infrequent consumers (Otherwise, an infrequent consumer will sell her token to a frequent consumer in the trading phase).

In period $\frac{M-1}{p}$, the quantity of tokens outstanding $x(0) = 1$. From period t to $t + 1$, there are $\frac{1}{2}p_H + \frac{1}{2}(x(t) - \frac{1}{2})p_L$ being used in this period. Thus,

$$x(t+1) = x(t) - \frac{1}{2}p_H - \frac{1}{2}[x(t) - \frac{1}{2}]p_L$$

We can solve the expression of $x(t)$,

$$x(t+1) + \frac{p_H}{p_L} - \frac{1}{2} = (1 - \frac{1}{2}p_L)(x(t) + \frac{p_H}{p_L} - \frac{1}{2})$$

$$x(t) = (1 - \frac{1}{2}p_L)^t (\frac{p_H}{p_L} + \frac{1}{2}) - \frac{p_H}{p_L} + \frac{1}{2}$$

Denote γ as the number of periods to deplete tokens among infrequent consumers. Then γ should be the smallest integer so that

$$x(\gamma) \leq \frac{1}{2}$$

$$\gamma = -\lfloor \frac{\log(1 + \frac{p_L}{2p_H})}{\log(1 - \frac{1}{2}p_L)} \rfloor \quad (3)$$

Combining Appendix 2 eq.(2) and eq.(3), we can write the price in period $\frac{M-1}{p}$ as a weighted average willingness to pay of frequent and infrequent consumers:

consumers? This is not an equilibrium either. When an infrequent consumer knows at least one frequent consumer is holding more than one token, she knows that the token price will continue to appreciate by the interest rate over the next period. Thus, the only equilibrium is the one where infrequent consumers are indifferent between holding the last token or not.

$$\widetilde{P}_{I,T}(t = \frac{M-1}{p}) = (1 - \beta^\gamma(1 - p_L)^\gamma) \frac{\beta p_L}{1 - \beta(1 - p_L)} + \beta^\gamma(1 - p_L)^\gamma \frac{\beta p_H}{1 - \beta(1 - p_H)}$$

Where γ represents the number of periods until only frequent consumers hold tokens solved above. The token price under heterogeneity is

$$\widetilde{P}_{I,T} = \beta^{\frac{M-1}{p}} [(1 - \beta^\gamma(1 - p_L)^\gamma) \frac{\beta p_L}{1 - \beta(1 - p_L)} + \beta^\gamma(1 - p_L)^\gamma \frac{\beta p_H}{1 - \beta(1 - p_H)}]$$

Appendix 2.4: Separating Equilibrium for Price Menu with Heterogeneity

This appendix section formalizes our discussion in Section 6.4 and derives the condition for when a separating equilibrium strictly dominates the pooling equilibrium; that is, if price menu is allowed, the platform offers a price menu so that frequent consumer buy more tokens at a higher price while the infrequent consumers buy fewer tokens at a lower price, rather than the same price and quantity.

To maintain a separating equilibrium, the following two incentive constraints and two participation constraints need to be satisfied:

For frequent consumers,

$$\underbrace{\sum_{i=1}^{i=M_H} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)}\right)^i}_{\text{WTP for } M_H \text{ Tokens}} - M_H P_H \geq \underbrace{\sum_{i=1}^{i=M_L} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)}\right)^i}_{\text{WTP for } M_L \text{ Tokens}} - M_L P_L \quad (\text{IC})$$

$$M_H P_H \leq \sum_{i=1}^{i=M_H} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)}\right)^i \quad (\text{PC})$$

For infrequent consumers,

$$\underbrace{\sum_{i=1}^{i=M_L} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)}\right)^i}_{\text{WTP for } M_L \text{ Tokens}} - M_L P_L \geq \underbrace{\sum_{i=1}^{i=M_H} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)}\right)^i}_{\text{WTP for } M_M \text{ Tokens}} - M_H P_H \quad (\text{IC})$$

$$M_L P_L \leq \sum_{i=1}^{i=M_L} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)}\right)^i \quad (\text{PC})$$

The total discounted revenue is

$$Rev = \frac{1}{2} \left[\underbrace{M_H P_H + M_L P_L}_{\text{Token Revenue}} + \underbrace{\frac{\beta^* p_H}{1 - \beta^*} \left(\frac{\beta^* p_H}{1 - \beta^*(1 - p_H)} \right)^{M_H} + \frac{\beta^* p_L}{1 - \beta^*} \left(\frac{\beta^* p_L}{1 - \beta^*(1 - p_L)} \right)^{M_L}}_{\text{Fiat Money Revenue}} \right]$$

The optimal token prices for frequent and infrequent consumers are

$$P_H = \frac{M_L P_L + \sum_{i=M_L+1}^{i=M_H} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} \right)^i}{M_H}$$

$$P_L = \frac{\sum_{i=1}^{i=M_L} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)} \right)^i}{M_L}$$

We plug in $M_H = 2M - M_L$ and rewrite the revenue as a function of M_L ,⁶

$$R(M_L) = 2 \sum_{i=1}^{i=M_L} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)} \right)^i + \sum_{i=M_L+1}^{i=2M-M_L} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} \right)^i + \frac{\beta^* p_H}{1 - \beta^*} \left(\frac{\beta^* p_H}{1 - \beta^*(1 - p_H)} \right)^{2M-M_L} + \frac{\beta^* p_L}{1 - \beta^*} \left(\frac{\beta^* p_L}{1 - \beta^*(1 - p_L)} \right)^{M_L}$$

If we neglect the discrete choice of M_L and M_H , we write the first-order derivative at point $M_L = M$. The first term is the loss of token issuance by pooling two types in the pooling equilibrium.⁷ The second term is the gain in fiat money revenue by persuading infrequent consumers to buy more tokens (that is, the platform can get fiat money from frequent consumers sooner).

$$\frac{dR(M_L)}{dM_L} \Big|_{M_L=M} = 2 \left[\underbrace{\left(\frac{\beta p_L}{1 - \beta(1 - p_L)} \right)^M - \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} \right)^M}_{\text{Loss from Token Revenue}} \right] + \left[\underbrace{\left(\frac{\beta^* p_H}{1 - \beta^*(1 - p_H)} \right)^M - \left(\frac{\beta^* p_L}{1 - \beta^*(1 - p_L)} \right)^M}_{\text{Gain from Fiat Money Revenue}} \right]$$

The platform would prefer to choose a separating equilibrium if the platform can gain

⁶ The optimal issuance policy (M_L, M_H, P_L, P_H) can be solved from the revenue maximization problem with incentive constraints and participation constraints.

For frequent consumers,

$$M_H P_H - M_L P_L \leq \sum_{i=M_L}^{i=M_H} \left(\frac{\beta p_H}{1 - \beta(1 - p_H)} \right)^i$$

For infrequent consumers,

$$M_H P_H - M_L P_L \geq \sum_{i=M_L}^{i=M_H} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)} \right)^i$$

The optimal M_L is pinned down by the F.O.C

$$0 = \frac{dRev(M_L)}{dM_L}$$

⁷In the pooling equilibrium, every consumer chooses the price quantity-price pair $(M, \frac{\sum_{i=1}^{i=M} \left(\frac{\beta p_L}{1 - \beta(1 - p_L)} \right)^i}{M})$.

sufficiently large profit from high-frequency consumers. The condition is⁸

$$\frac{\left(\frac{\beta^* p_H}{1-\beta^*(1-p_H)}\right)^M - \left(\frac{\beta^* p_L}{1-\beta^*(1-p_L)}\right)^M}{\left(\frac{\beta p_H}{1-\beta(1-p_H)}\right)^M - \left(\frac{\beta p_L}{1-\beta(1-p_L)}\right)^M} < 2$$

⁸A platform would choose the pooling equilibrium for sure if $\beta^* < \beta = 1$.