

Working Paper presented at the

Peer-to-Peer Financial Systems

2016 Workshop

July, 2016

Market Design for

Trading with Blockchain

Technology

Katya Malinova

University of Toronto

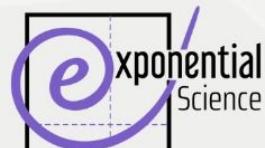
Andreas Park

University of Toronto



P2P Financial Systems

Powered by



Market Design for Trading with Blockchain Technology^{*}

Katya Malinova[†] Andreas Park[‡]

First Version: February 2, 2016

This Version: July 7, 2016

Abstract

An intrinsic feature of blockchain technology is the so-called distributed ledger, which involves maintaining a decentralized record of all transactions for a security. A distributed ledger allows anyone with access to it to identify asset holdings by (anonymous) identifiers. In this paper, we argue that the mapping between identifiers and end-investors should be a market design choice. We capture features of blockchain technology in a theoretical model of intermediated and peer-to-peer trading, and we study the impact of the usage of identifiers and the corresponding transparency of holdings on trading behavior, trading costs and investor welfare. We find that, despite the risk of front-running, the most transparent setting yields the highest investor welfare. In the absence of full transparency, for low levels of liquidity in the intermediated market, welfare is highest if investors are required to concentrate their holdings under single identifiers.

^{*}We thank Stephen Bain from RBC Capital Markets, Jacob Farber and Tim Swanson from R3, and Jeff Coleman from Ledgerlabs for inspiring and informative discussions on the topic.

[†]University of Toronto, Department of Economics, katya.malinova@utoronto.ca.

[‡]University of Toronto Mississauga, Department of Management and Rotman School of Management, andreas.park@rotman.utoronto.ca (corresponding author).

Since the second half of 2015, “Blockchain” has been the buzzword in the financial press; see Figure 1 for an illustration. The reason for the buzz is that people came to realize that blockchain technology has the capacity to be as disruptive to the financial industry as the internet was to offline commerce. The internet enables producers (e.g., airlines) to sell goods and services (e.g., flight tickets) to consumers by cutting out profit-seeking retailers (e.g., travel agencies). Blockchain technology enables people to exchange sensitive information and financial claims (e.g., money, stocks, bonds, property ownership) without involving rent-seeking intermediaries (e.g., banks) that act as keepers of information. Indeed, Blythe Masters, the inventor of credit default swaps, described the blockchain as the “financial challenge of our time” and urged bankers to pay as much attention to this technology as they should have to the internet in 1994.¹ During a panel at the 2016 AFA meetings, Michael Casey from MIT’s Media Lab described blockchain technology as the most significant change to the operation of financial institutions since the Medici in the 15th century.²

So what exactly is blockchain technology? The World Wide Web is the technology that enables the frictionless transfer of information; the Blockchain is the technology that enables the frictionless transfer of value. In essence, the blockchain is a ledger of all transactions for an item that is kept not in a (private) database, library, or filing cabinet, but in a public, distributed, commonly accessible network. Transactions are verified by linking ownership of transacting parties to so-called public keys, which can be thought of as (anonymous) identifiers. Verification of transactions occurs via so-called proof-of-work protocols where multiple parties compete to verify for a fee.³ Historically, non-cash

¹See: <http://tabbforum.com/videos/blythe-masters-on-the-significance-of-blockchain-the-financial-challenge-of-our-time>.

²In January 2016, the Financial Times reported that the “Blockchain debate eclipses Basel III at Davos.”

³For an excellent summary of the functioning of blockchain technology w.r.t. Bitcoins, see Harvey

financial transactions have required banks or investment dealers to serve as “trusted third parties”. These third parties ensured that the counterparties of a transaction had the cash and the security. Blockchain technology instead allows true “end-investor”-to-“end-investor” transactions, facilitated by decentralized proof-of-work verification protocols, without a “trusted third party,” which, after all, is a rent-seeking intermediary.

At this point, there is little knowledge as to the economic impact of this technological transformation, but there is wide consensus that this technology will be adopted in the near future.⁴ Blockchain technology has a number of intrinsic organizational features that have the capacity to change the nature of financial interactions and that raise interesting and important questions in particular with regard to optimal market design. In this paper, we discuss the critical organizational features that relate to counterparty transparency, and we provide a theoretical framework to help understand their impact on economic welfare. Since the technology has not been widely used with the exception of bitcoins and some small-scale experiments,⁵ there is no meaningful data to perform empirical studies.

Much of the discussion and the interest with regard to its implementation relate to the possibility of significant post-trade clearing-and-settlement efficiency gains because with the blockchain, the transaction and the settlement are the same. From a market design perspective, however, it is much more interesting how blockchain technology can affect the pre-trade market organization. One intrinsic feature of a distributed ledger is that it has the capacity (and one can argue the objective) to remove information frictions.

(2015) or the youtube video <https://youtu.be/Lx9zgZCMqXE>.

⁴See, for instance, Bali and Roche (2015): The TABB Group predicts that blockchain technology may be used for syndicated loans as early as Q2 2016.

⁵For instance, NASDAQ recently launched trading in the private-market company Chain.com, or R3, the consortium of 40+ global banks that aims to develop workable blockchain solutions, recently ran an experiment involving Blockchain transfers among 11 banks using Ethereum’s technology.

As part of the verification process of transactions, verifiers check that the selling party in a transaction owns the item (e.g., the share), and thus ownership is linked to a public identifier (or, in terms of the technology: a public key). These identifiers are anonymous in the sense that they cannot be traced back to a person or institution by anyone other than the identifier’s owner. The ability to directly attribute ownership to an (anonymized) identifier contrasts today’s world, where the trusted third party (and only this party) can make the connection to the seller and can verify ownership.

Alas, it is not clear how much of this ownership information will and should be revealed in the final implementation of blockchain technology. For Bitcoins, for instance, holdings are transparent to all members of the Bitcoin network. Financial institutions, however, are likely deeply skeptical about this degree of counterparty transparency, and they are working on blockchain solutions to limit knowledge of holdings to verifiers only.

Our goal in this paper is to provide insights that help understand and determine optimal market design of counterparty transparency, and we provide a theoretical framework to analyze settings with different levels of counterparty transparency. We model the trading of financial securities, and we differentiate between small “retail” investors and large, institutional investors, where our focus is on the latter. There are (at least) two ways to restrict transparency about the total size of one’s holdings. The first is to limit the investors’ ability to view the information that is attached to the public keys. The second is to allow investors to create multiple keys or, in the extreme case, implement a “one share – one key” system (instead of “one entity – one key”) so that the size of an investor’s holdings cannot be inferred from the holdings that are attributed to each key.

We consider three regimes for the peer-to-peer interaction. In the first two regimes, there is a one-to-one mapping between public keys and investors, in the third there is a

one-to-one mapping from public keys to shares. In the first and third settings, investors trade peer-to-peer not knowing the size of their counterparty, in the second setting investors know the size of their counterparty. Our benchmark is a setting where all trades go through market makers who charge fees for intermediation. We then compare the benchmark setting to markets where, in addition to trading with the intermediary, investors also have the option to trade directly among each other. We assume this peer-to-peer interaction is facilitated by a system that allows investors to contact one another to arrange transactions.⁶ While such a system does not exist in practice at this point (neither does security trading with blockchain technology), it would be a logical innovation to trading, and in particular can employ the so-called “smart contract” features.⁷

Before we get to the details of our model and our results, let us briefly discuss the trade-offs regarding transparency. Investors who want to trade need liquidity, meaning that they need to find a party that is willing to take the other side of their trade. Finding a counterparty can be challenging, in particular when the investors want to exchange a large quantity or if the security is traded infrequently; examples are the markets for corporate and off-the-run government bonds. Liquidity sourcing in the sense of knowing who is interested in trading the security is arguably useful for investors on either side of a trade, and this knowledge is critical for a dealer’s success. In current markets,

⁶For bond trading such systems exist: the firm Algomi has developed an intricate system that allows direct interactions, albeit without blockchain technology.

⁷For Bitcoins, the main usage of blockchain technology is to record and transfer values in digital currencies. Yet, much more than quantities can be stored on the Blockchain using a so-called “smart contracts” system. Loosely speaking, smart contracts are computer protocols that specify rules governing transactions, and that are enforced by a network of peers (i.e., without the need of a middleman). Financial securities and instruments can potentially be represented directly on the distributed ledger, as smart contracts, without any reference to a central database. Since smart contracts are capable of embedding variety of information and are self-executing, a critical feature on trading securities that are represented by smart contracts, in our view, is that trading can be designed in a way that an investor is always present “in the market” and may be contacted (assuming that a communication system is enabled). With investors “always” being in the market, decentralized trading using blockchain technology arguably should not be subject to high search costs or concerns about fragmentation of liquidity.

investors often have to use the help of an intermediary who is willing to absorb their position or of a dealer who arranges a match. This service comes at a cost, borne by investors and, arguably, the issuer of the security. The cost's origin is, at its core, an information friction. We investigate in this paper whether distributed ledgers can mitigate these existing frictions, and we aim to identify new frictions that arise from the built-in transparency of blockchain technology.

Undoubtedly, transparency about investors' holdings and trades can have downsides. For instance, investors may copy portfolios and thus eliminate the copied investor's comparative advantage. In this paper, we focus on the trading costs of transparency in the sense that when trying to liquidate a position, owners of large positions may get front-run if so-called scalpers detect their trading interest.⁸

Trading in our model happens repeatedly in discrete time. Each period, exactly one of two large investors is hit with a liquidity shock. Our benchmark is a setting in the tradition of Biais (1993) where investors interact with risk-averse intermediaries in a setting with symmetric information. In this benchmark, the large investor clears his position by trading with the intermediaries. These are concerned about inventory risk and thus require risk compensation; the trade causes a (short-term) fluctuation in the price beyond the volatility warranted by fundamental risk.

We then compare the fully intermediated benchmark to three regimes with peer-to-peer interactions. In the first, the large investor can additionally contact the continuum of other traders; most of these are small and accept the large investor's offer with some positive probability. The other large trader is one of many in the continuum and is thus

⁸As Christoffersen, Danesh, and Musto (2015) document, mutual funds often delay publishing their holding information in 13-F forms for as long as possible. Christoffersen et al. argue that mutual fund managers are mostly concerned about being front-run by competitors. Danesh (2015) provides a theoretical model in the tradition of Kyle (1985) to analyze front-running behavior.

contacted with probability zero. Trades between the investors here occur at the fundamental value, but in approaching the continuum of other investors, the large investor incurs complexity costs.⁹ We show that in any equilibrium, large investors split their large trade between small investors and the intermediaries, payoffs for large investors are higher and price volatility is lower than in the benchmark, intermediated market.

In the second setting, investors see each others' holdings so that large investors are identifiable as large. For this setting we show that large investors only trade among each other. Even though front-running is in principle possible, it does not occur in equilibrium because of the repeated interactions of large investors. We further show that for a sufficiently large discount factor (i.e., when the future is sufficiently important), the cost of curbing front-running vanishes. This setting is also payoff dominant because large investors incur no complexity costs for contacting small investors and they do not interact with the intermediary; thus the latter's risk aversion does not impose a cost on end-investors. Notably, this setting displays fragmented trading in the sense that large and small investors do not interact. Such outcomes can occur in other settings. For instance, Pagano (1989) describes how the existence of multiple markets may lead, among other things, to fragmentation, when investors choose to concentrate their trades on a single market; in Lester, Rocheteau, and Weill (2015) investors segment based on the terms of trade offered by dealers, and investors with different asset valuations and different asset holdings trade at different speeds and different costs.

In our third setting, investors do not concentrate their holdings under a single public key but instead spread their holdings (in a one-share, one-key manner). Here, the equilibrium outcome depends on the model parameters. When future payoffs are discounted

⁹For the currently existing blockchain verification protocols, the complexity cost is effectively the fee that trading parties pay to the miners for each transaction.

heavily, front-running cannot be prevented by incentivizing potential front-runners. Instead, the large trader then avoid being front-run by what effectively amounts to “over-trading” with the intermediary. The payoffs for large investors in this latter outcome turn out to be dominated by those in the first and second regime, but are still superior to those with pure intermediation because large investors can trade at least some of their holdings with small investors. If the intermediated market is very liquid or if the discount factor is sufficiently large (the future is sufficiently important), then front-running can be curtailed and in this case, payoffs for large investors dominate those in the first regime — but not those in the setting with full transparency.

I. Related Literature

In addition to contributing to the fast-growing literature on blockchain technology, our paper relates to several strands of the literature in market microstructure and market design.

Literature on Blockchain Technology. The academic literature on blockchain technology is small but growing. As of March 19, 2016, SSRN lists only 37 working papers that use the term “Blockchain”, as of May 27, 2016, there are 56 papers. Most of these describe technological or legal issues. There are five notable exceptions. Harvey (2015) provides an overview of Bitcoin’s technology. Evans (2014) discusses concerns that could arise with regards to the verification incentives that are necessary to ensure that transactions can occur. Lee (2016) provides an overview of the functioning of blockchain technology, highlights security concerns, and she speculates qualitatively how trading using blockchain technology would affect short-selling¹⁰ and high-frequency

¹⁰At face value, it would appear that short-selling is intrinsically impossible with blockchain technology because shorting implies not-owning, and verifying ownership, a key part of trad-

trading. Finally, Brummer (2015) provides an overview of the effects of technological disruption on the regulation of financial markets. Yermack (2016) discusses the potential implications of blockchain-based trading on corporate governance.

A core feature of blockchain technology is that it, arguably, simplifies the clearing-and-settlement process; for theoretical work on the importance of these costs (albeit not in the context of blockchain technology) see Degryse, Van Achter, and Wuyts (2012).

Since this technology is brand-new, there is much room for many types of economic analyses, ranging from effects on labor markets, capital requirements (faster settlement frees capital), securitization, efficiency gains and cash-flow risks generated by smart contracts, and so on. The findings in this paper highlight the importance of the transparency afforded by blockchain technology in terms of affecting transaction costs; and we hope to provide a useful starting point for a debate on the optimal market design with these new technologies.

Literature on Over-the-Counter Markets. Peer-to-peer interactions have been extensively analyzed in studies of over-the-counter markets. This literature started with Diamond (1982), Ariel Rubinstein (1985), Gehrig (1993), and Yavas (1996), and developed into a related strand on asset pricing in search-based models, e.g., Weill (2002), Duffie, Garleanu, and Pedersen (2005), Miao (2006), Vayanos and Wang (2007), or Cujean and Praz (2015). The contribution of our work is to study transparency and its impact on trading decision and costs, focusing on the situations where the level of transparency is endogenous to the different possible mappings of ownership-to-public keys. Differently to the over-the-counter search models, which assume that a trader interacts with one other trader at a time, our setting allows investors to trade with multiple coun-

ing via the blockchain would thus not be possible. However, the company Overstock.com has developed the concept of a “short token” that they say would overcome the issue; see also <http://tabbforum.com/videos/transparent-shorting-via-blockchain>.

terparties at the same time. Furthermore, we introduce a setting where the large trader is able to achieve anonymity by trading through multiple accounts.

The two most related papers from the literature on over-the-counter markets are Lester, Rocheteau, and Weill (2015) and Cujean and Praz (2015). Lester, Rocheteau, and Weill (2015) develop a competitive search model where traders observe posted prices of multiple dealers but face delay costs that are due to the dealers' limited processing capacity. In our model, trades with small investors occur at the expected value as the small investors do not post quotes; the existence of small investors and their potential willingness to trade are known and facilitated by the technology. In our setting, the per-trader cost is infinitesimal, but when contacting the continuum of small investors in our setting, large investors incur a non-zero aggregate processing cost; similar in spirit to the dealers' cost in Lester, Rocheteau, and Weill (2015). When contacting other large investors, investors' processing costs remain infinitesimal. This setting represents the presumed ideal situation when "natural" traders can easily find one another because of the revealed ownership that blockchain technology affords. The intermediated market in our setting serves as both a benchmark and as an outside option that allows us to determine the peer-to-peer equilibrium price.

Cujean and Praz (2015) model OTC trading in the tradition of Duffie, Garleanu, and Pedersen (2005). In their model, a dealer who is approached by the trader receives a signal about the trader's inventory (and hence the trader's likely liquidity needs). The dealer then proposes a price that takes the trader's liquidity need into account. Increased transparency (in the sense of more precise information about the inventory) in their model allows the trader to find a better counterparty with an off-setting liquidity need, but it exposes the trader to the risk of predatory pricing in the bargaining phase.

In the equilibrium, increased transparency increases inventory costs but leads to a more efficient asset allocation. In our framework, investors are also exposed to the possibility of predatory behavior by their trade counterpart, but they have a choice of who to interact (or not) with, in a repeated setting. We focus on comparing market designs that endogenously afford different degrees of transparency and that allow traders different degrees of control over their choices of counterpart.

Literature on Centralized vs. Decentralized Markets. Our model also touches upon the literature that compares centralized with de-centralized markets. This literature follows Biais (1993);¹¹ papers in this line of work compare transaction costs (usually in terms of bid-ask spreads) for trading systems where quotes are collected centrally with systems where investors must obtain quotes from intermediaries directly. Our model implicitly combines a central market, captured by the intermediary, with a decentralized market where investors trade peer-to-peer, without the involvement of an intermediary. We focus on the impact of transparency of investor holdings, and we assume that investors in the peer-to-peer market are able to avoid direct trading costs.

II. Model

Our model has three types of market participants: two large institutional investors, one of who is randomly selected to face the need of having to trade a large quantity; a continuum of small investors that in aggregate have the capacity to absorb the institutional order; a group of risk-averse intermediaries that can absorb order flow for a fee. We allow for two trading mechanisms: one is direct interaction between investors, the other is trading with intermediaries.

¹¹See also De Frutos and Manzano (2002), or Yin (2005).

The Asset has a fundamental value of v that is normally distributed with mean V and variance σ^2 . Information regarding the distribution of the fundamental value is public knowledge. The asset is infinitely divisible.

Large Investors. There are two large, risk neutral investors. At any time t , one of them is hit by a liquidity shock which requires them to trade a quantity Q . To simplify the exposition, we write the paper assuming $Q > 0$, but the arguments are symmetric. Investors discount future trading opportunities at rate $\delta < 1$. Each large investor has the capacity to absorb the other's shock without incurring a cost.

Small Investors can trade unit quantities, and they do so with probability ρ when approached. There is a continuum of αQ many of these small investors, with $\alpha \geq 1/\rho$. They are willing to trade at any price that is at or better than what they can get from the intermediary.

Intermediaries. There are N identical, risk averse intermediaries that have CARA utility function of wealth with risk aversion coefficient $\kappa > 0$. We assume, for simplicity, that at the beginning of each stage game, these intermediaries hold no inventory.¹² When an institution approaches the market, the intermediaries post supply schedules that maximize their utility, and then markets clear. Namely, they each specify for each price p how many shares they are willing to buy (or sell) $\forall p \in \mathbb{R} \quad q_i(\cdot) : p \rightarrow q_i(p)$. Markets clear at a uniform price.

Direct Trading Costs. Contacting mass q of investors is complex (e.g., data processing, or keeping track of offers) and costs $C(q)$, where $C', C'' > 0$.¹³ We assume, for simplicity, that $C(q) = \frac{\epsilon}{2}q^2$.

Ownership Concentration. As discussed in the introduction, blockchain technol-

¹²Loosely, we assume that between the arrivals of liquidity shocks for large investors, the intermediaries can manage their inventories by trading with the small investors.

¹³For Bitcoin or Ethereum transactions, “miners” receive a fee for each transaction that they verify.

ogy in principle enables different models of ownership concentration in the sense that investors can operate under a single or multiple public keys. We assume that small investors always operate under a single public key (and there is thus a continuum αQ of their public keys). For large investors, we consider two settings for ownership. In the first, each large investor operates under a single public key; in this setting, there are αQ public keys. In the second, each large investor owns a continuum Q of public keys and equally disperses their ownership over the keys, implying that in this second setting, there are $\alpha Q + Q$ many public keys that an investor may potentially trade with.

Transparency of Ownership and Contacting Protocols. We consider two settings for transparency: in the first, ownership is not revealed and contacts are memory-free, meaning that traders do not know who they interacted with in the past. Direct contact is possible in the sense that a large investor can “ping” the pool of public keys anonymously with an offer to trade; in this setting, we also allow “ping-back” in the sense that, after being contacted, investors can reject an offer and make a counter-offer. In the second setting, ownership and thus each investor’s size are visible, and investors can contact one another directly. We further assume that trading offers are binding and that they cannot be withdrawn.

Timing. Trading is organized as an infinitely repeated game in discrete time. The stage game timing is as follows: first, one of the large investors is randomly selected to be hit with a liquidity shock; we refer to this trader as the “liquidity trader.” This trader then approaches the other investors (small and/or large) and/or the intermediary. The other investors either accept the offer and the trade occurs, or they reject. The intermediary fills the large trader’s request immediately upon receipt. We allow large traders to front-run one another in the sense that before responding, they can build up

a position and make a counteroffer to the liquidity trader. If necessary, the liquidity trader fills the remainder of his position with the intermediary.

Indirect Trading Costs: Front-running. We model the indirect trading costs that arise from disclosing one’s trading intent through the cost of front-running. When a liquidity trader contacts another large investor, the liquidity trader may be front-run by the prospective counterparty. We model the mechanics as follows. Suppose a large investor is contacted by the liquidity trader who wishes to buy a quantity q . If he chooses to front-run, this investor buys a quantity q from the intermediary at a price, and then resells it to the liquidity trader at a higher price. We assume that the total price for the quantity q that the front-runner offers to the liquidity trader is the minimum cost that the liquidity trader has to pay to acquire q in the “public market” after the front-runner has moved the price by purchasing q .

Related Approaches. Our approach relates to the extensive economic literature on search and trading at decentralized exchanges in presence of intermediaries. This literature started with Diamond (1982), Ariel Rubinstein (1985), Gehrig (1993), and Yavas (1996), and developed into a related strand on asset pricing in search-based models, e.g., Weill (2002), Duffie, Garleanu, and Pedersen (2005), Miao (2006), or Vayanos and Wang (2007).

III. Benchmark: Pure Intermediation

We construct a benchmark of a purely intermediated market following Biais (1993). Assuming that there are $N > 0$ intermediaries, when asked to sell quantity q (i.e., when an investor wants to buy q units), each intermediary maximizes their expected utility by selling $q_i(p)$; in equilibrium the price clears the market so that $\sum_{i=1}^N q_i(p) = q$. At

the beginning of each stage game, intermediaries hold no inventory. However, when contacted by a front-runner, an intermediary may already hold a position I_i resulting from trading with the liquidity trader.

With negative exponential (i.e., CARA) utility of wealth w , $u(w) = -e^{-\kappa w}$, and $w = -(v - p)q_i + I_i$, the intermediary chooses quantity q_i given price p , in order to maximize his expected utility, $\max_{q_i} \mathbb{E}U[-(v - p)q_i(p) + I_i \times v]$. For CARA-normal frameworks, this task is akin to maximizing the certainty equivalent for each price p :

$$\max_{q_i} [I_i \times V - (V - p)q_i] - \frac{\kappa}{2}\sigma^2[-q_i + I_i]^2,$$

which leads to first order condition

$$V - p - \kappa\sigma^2 \times I_i + \kappa\sigma^2 q_i = 0.$$

Solving this for q_i , yields the demand schedule

$$q_i(p) = -\frac{V - p}{\kappa\sigma^2} + I_i.$$

The market clearing condition

$$\sum_{i=1}^N q_i(p) = q$$

implies, substituting for q_i , and simplifying, that

$$\sum_{i=1}^N \left(-\frac{V - p}{\kappa\sigma^2} + I_i \right) = q \Leftrightarrow p^{\text{mm}}(I, q) = V + \frac{\kappa\sigma^2}{N} (-I + q), \quad (1)$$

where I denotes the combined inventory of the intermediaries: $I = \sum_{i=1}^N I_i$. We further

simplify the exposition by defining the (il-)liquidity factor ℓ as follows

$$\ell := \frac{2\kappa\sigma^2}{N}.$$

Price changes in this model occur for two reasons: changes in the fundamental, and trades due to liquidity shocks. When an investor approaches the intermediaries who hold total inventory I in order to buy q units, and the investor's payoff is

$$\pi^{\text{mm}}(I, q) = q \times (V - p^{\text{mm}}(I, q)) = -\frac{\ell}{2} q(q - I).$$

We henceforth make the following assumption

Assumption 1: *The expected value of the asset is $V = 0$.*

Trades among investors are zero-sum in a welfare sense: if one investors charges the other a price in excess of the fundamental, then the latter's loss is the former's gain. Trades with the intermediary, however, are welfare reducing because the intermediary is risk-averse whereas all other traders are risk-neutral. Trades with the intermediary also cause prices to deviate from the fundamental value and thus, arguably, to lower price efficiency.

IV. Non-Transparent Concentrated Ownership

In this section, we study a setting in which the liquidity trader contacts other traders without knowing the counterparty's size and where each investor trades using a single public key.

Since the large trader is infinitesimal, the chance of contacting him is zero. Small

investors accept an offer to trade with probability ρ , thus when approaching measure x of them, the liquidity trader trades quantity $x\rho$. Assuming that the intermediary has an inventory of I , the liquidity trader pays price $x\rho\frac{\ell}{2} \times (-I)$ to the small investors. Moreover, approaching mass x of small traders, the liquidity trader incurs complexity cost $C(x)$. When purchasing q from the intermediary, the liquidity trader pays the price $p^{\text{mm}}(I, q)$ defined in (1), which is higher than the price incurred when trading with small investors. Suppose the liquidity trader wishes to trade a total quantity of q . If he approaches a mass x of small investors, he trades quantity $\rho \times x$ with them, and he trades a quantity, $q - \rho x$ with the intermediaries. The risk-neutral large trader chooses x to maximize

$$\max_x -\frac{c}{2}x^2 - x\rho\frac{\ell}{2} \times (-I) - (q - \rho x)(q - \rho x - I) \times \frac{\ell}{2} \Rightarrow x^c(q) = q \times \frac{\ell\rho}{\ell\rho^2 + c}, \quad (2)$$

where henceforth superscript c signifies the anonymous **concentrated** ownership trading outcomes. Note that the optimal quantity choice does not depend on the intermediaries inventory. The liquidity trader's payoff is:

$$\pi^c(q, I) = -q^2 \times \frac{\ell}{2} \frac{c}{\ell\rho^2 + c} - \frac{\ell}{2} q(-I). \quad (3)$$

In equilibrium, the liquidity trader wishes to trade $q = Q$, and since at the beginning of each stage game the intermediaries have no inventory, the payoff would be $\pi^c(Q, 0)$. By trading approaching $x = Q/\rho$ small traders, the liquidity trader could fill his entire position with the small investors. However, in the absence of front-running, the marginal cost of trading an arbitrarily small quantity with intermediaries is arbitrarily close to zero. The large investor therefore always trades with both small investors and

intermediaries, and he chooses the equilibrium quantities to trade with each entity so that the marginal cost in the intermediated market coincides with the marginal cost of contacting small investors.

Proposition 1 (Non-Transparent Concentrated Ownership): *Large investors always split their position among small investors and intermediaries. They trade more with small traders if intermediaries are more risk-averse or if fundamental risk increases, and they trade less with small traders if complexity costs increase.*

Components of our setting in this section are similar to Pagano (1989). He studies a setting where traders can choose between a centralized market (where illiquidity may cause price dislocations and thus high transaction costs) and a search market, where they incur search costs. The equilibrium in Pagano is described by the number of investors who gravitate towards the centralized market relative to those that search. In our equilibrium, all investors “search”; the critical relation for the equilibrium are the search costs in the decentralized market relative to the price impact on the centralized market.

V. Transparent Concentrated Ownership

In this section, we continue to assume that each investor owns a single public key, but we assume that holdings are transparent and large investors can be identified as large. We search for an equilibrium where the large investors only trade among themselves and not with small investors or the intermediary, so that direct trading costs are zero, and we focus on the existence of equilibrium with no front-running. In describing the equilibrium of the infinitely repeated stage game, we consider only so-called “trigger”-strategies. That is, if any participant observes behavior that is “off the equilibrium

path”, e.g. he gets offered a price other than the equilibrium price, a liquidity provider front-runs, or a liquidity trader trades with the intermediary and the small investors, then the large investors do not contact one another any longer. If the trigger strategy is invoked, then the equilibrium behavior from there on is as described in Section IV.

The timing of the game is as follows. After the liquidity trader is hit with a shock, he contacts the other large investor, which we refer to as the liquidity provider, and makes a take-it-or-leave it offer to trade at a price $p^{p2p} \geq 0$. ¹⁴ Subject to non-negative prices, in describing the equilibrium, we further restrict attention to the case where we give the liquidity provider full bargaining power, that is, if there are multiple feasible equilibrium prices, the smallest possible one obtains.

Each stage game resembles a prisoner’s dilemma. Large investors can choose to “cooperate” by offering the other large investor a price concession when hit with a shock and by not front-running when approached by a “shocked” investor. The liquidity trader may deviate by instead trading with the intermediary and with the continuum of small investors, as in Section IV. The liquidity provider may deviate by front-running. After a deviation, in each stage game, when hit with a liquidity shock the trader earns $\pi^c(Q, 0)$ (as defined in (3), where by assumption the intermediaries’ inventories are 0 at the beginning of the stage game) and 0 otherwise. Thus if they could agree to trade at price 0, both investors will be better off. However, payoffs from front-running are attractive:

Lemma 1 (Front-Running Profits): *The maximum stage payoff that the liquidity provider may achieve through front-running is $-\pi^c(Q, 0)$.*

Proof: To extract rents from the liquidity trader, the front-runner first accumulates

¹⁴Negative equilibrium prices are, in principle, possible because liquidity providers “pay it forward”, meaning that they accept a low price today in return for getting a better price in the future when they are hit with a shock. In the proof of Proposition 2 we will show that whenever $p^{p2p} < 0$ is an admissible equilibrium outcome, $p^{p2p} = 0$ is, too.

the position by trading with the intermediary.¹⁵ We assume that he is then able to resell this quantity to the liquidity trader at a price that equals the minimum possible price that the liquidity trader would need to pay to build this position on the open market (i.e., with the continuum of small traders and with the intermediary).

Front-running is costly to the liquidity trader because the trade of the front-runner with the intermediary would (*i*) move the public price, which the liquidity trader trades at with the small traders, and (*ii*) result in the positive inventory of the intermediary, implying a higher price when trading with the intermediary. Suppose the front-runner trades quantity q^{fr} with the intermediary. The front-runner pays the intermediary $q^{\text{fr}} \times p^{\text{mm}}(0, q^{\text{fr}}) = \ell/2 \times (q^{\text{fr}})^2$. After being front-run, the liquidity trader could solve the same problem as described in Section IV. for $I = -q^{\text{fr}}$, and earn $\pi^c(-q^{\text{fr}})$. The front-runner could then offer the liquidity trader quantity q^{fr} at a price such that the liquidity trader's payoff is at $\pi^c(Q, -q)$. Specifically, the front-runner would charge $\pi^c(Q, -q^{\text{fr}}) - \pi^c(Q - q^{\text{fr}}, -q^{\text{fr}})$ at a cost to him of $q^{\text{fr}} \times p^{\text{mm}}(0, q^{\text{fr}})$. The front-runner's profits are increasing in Q and he will thus choose $q^{\text{fr}} = Q$ so that his payoff would be $-\pi^c(Q, 0)$. \square

Consequently, in a static game, large traders would not contact one another as they fear being front-run, and rents are lost because the intermediary gets involved every time a large investor receives a shock.

However, investors do not solve a single-shot game, but they interact repeatedly. Consequently, a front-runner has to take into account that by front-running, he foregoes future trading opportunities. We discuss the repeated payoffs in what follows.

Liquidity Trader. When offering p to the other trader for quantity Q , in the absence of front-running, the liquidity trader pays $-pQ$ today. In the next period, with probability

¹⁵We assume that front-running by trading with the continuum of small investors is not possible, e.g., because trades take time to clear.

1/2, the trader receives another liquidity shock and has to pay $-pQ$, and with probability 1/2 the other trader receives a shock in which case the trader receives payment pQ . Taken together, the continuation payoff in the absence of front-running, is 0. Suppose the trader deviates and contacts the continuum. In this case, today the liquidity trader earns profit $\pi^c(Q, 0)$, and for all future periods, the trigger strategy is invoked. This implies that, when hit with a shock, the trader has profit $\pi^c(Q, 0)$, and when the other gets shocked, he has payoff 0. The continuation value when playing the trigger strategy is thus:

$$FV(\text{trigger}) = \frac{1}{2}\pi^c(Q, 0) + \delta \times FV(\text{trigger}) \Leftrightarrow FV(\text{trigger}) = \frac{1}{2}\frac{1}{1-\delta}\pi^c(Q, 0).$$

Taken together, we obtain the following payoff for the “shocked” liquidity trader who, instead of paying the concession to the large trader, deviates by trading with the continuum of small investors and the intermediaries:

$$\Pi^{\text{LT}}(\text{deviate}) = \pi^c(Q, 0) + \delta \times FV(\text{trigger}) = \pi^c(Q, 0) \left(1 + \frac{1}{2}\frac{\delta}{1-\delta}\right).$$

For the liquidity trader to continue to play on-path, the price p must be such that

$$-pQ \geq \pi^c(Q, 0) \left(1 + \frac{1}{2}\frac{\delta}{1-\delta}\right) \Leftrightarrow p \leq -\frac{\pi^c(Q, 0)}{Q} \left(1 + \frac{1}{2}\frac{\delta}{1-\delta}\right). \quad (4)$$

Liquidity Provider. When offered p by the other trader, the liquidity trader receives $-pQ$ today. His equilibrium (no-front-running) continuation payoff is the same as that of the liquidity trader, 0. Suppose the trader deviates and front-runs. In this case, today the liquidity trader earns profit $-\Pi^a$ today, and for all future periods, the trigger

strategy is invoked. The continuation value here is the same as for the liquidity trader. Taken together, we obtain

$$\Pi^{\text{LP}}(\text{deviate}) = -\Pi^{\text{a}} + \delta \times \text{FV}(\text{trigger}) = \pi^{\text{c}}(Q, 0) \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right).$$

For the liquidity provider not to front-run, the price p must be such that

$$pQ \geq \pi^{\text{c}}(Q, 0) \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right) \Leftrightarrow p \geq \frac{\pi^{\text{c}}(Q, 0)}{Q} \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right). \quad (5)$$

Proposition 2 (Peer-to-Peer with Transparency and Concentrated Ownership): *For all parametric configurations, there exists a price $p \geq 0$ such that the large investors can agree on a price such that the trade occurs and no investor engages in front-running. For $\delta \geq 2/3$, $p = 0$ is a always an equilibrium price.*

Proof: To see why the result is true, observe that for an equilibrium price p to exist, the conditions (4) and (5) must hold simultaneously,

$$\frac{\pi^{\text{c}}(Q, 0)}{Q} \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right) \leq p \leq -\frac{\pi^{\text{c}}(Q, 0)}{Q} \left(1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right).$$

Since $\pi^{\text{c}}(Q, 0) < 0$, we can simplify

$$-1 - \frac{1}{2} \frac{\delta}{1 - \delta} \leq -1 + \frac{1}{2} \frac{\delta}{1 - \delta}.$$

This relation always holds. Further, for $\delta > 2/3$, $\frac{\delta}{2} > 1 - \delta$, thus $\frac{\pi^{\text{c}}(Q, 0)}{Q} \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right) < 0$ so that $p = 0$ is an equilibrium. \square

Since, by assumption, we allow the liquidity trader to choose his preferred price, the

equilibrium price in this setting is

$$p^{\text{p2p}} = \max \left\{ 0, \frac{\pi^c(Q, 0)}{Q} \left(-1 + \frac{1}{2} \frac{\delta}{1 - \delta} \right) \right\}. \quad (6)$$

VI. Non-Transparent Dispersed Ownership

In this section, we study a setting where each large trader owns a continuum of public keys (“one share, one public key”). We further assume that after the round of trading, there is no memory — otherwise, since large investors get contacted with positive probability by a liquidity trader, they would be able to learn (or extract from the data) the identify of the other investor. In practice, this can be achieved by generating numerous public keys. Thus once both large investors were hit by a liquidity shock once, their identities are known and the situation is akin to trading in Section V. (except that traders incur complexity costs).

The situation is thus that it is not possible for a liquidity trader to contact only the large trader as he is indistinguishable from small traders. We are looking for the existence of an equilibrium in which the large trader contacts the continuum and possibly the intermediary, offers *the continuum* a price concession, and small traders accept with probability ρ and large traders always accept the offer. In this case, the liquidity trader will also contact the intermediaries for a portion of their trade, for the same reason as in Section IV. (the marginal cost of trading a close-to-zero quantity with the intermediary is arbitrarily close to zero). As in the last section, we assume that deviations are followed by a trigger strategy whereby large investors reject any offer from the liquidity trader and liquidity traders submit their orders taking this reaction into account.

When the large investor disperses his size Q ownership (or willingness to trade), the

total mass of the continuum of traders is $Q + \alpha Q$. When contacting x investors, fraction $\alpha\rho/(1 + \alpha) \times x$ of small investors accepts. We look for an equilibrium in pure strategies, and we assume that public keys that belong to the large investor either all accept or all reject. If the large investor accepts the offer in equilibrium, then when contacting x investors, the liquidity trader will trade quantity $(\alpha\rho+1)/(1+\alpha) \times x$ with the continuum.

As in the preceding section, we analyze the decisions of liquidity traders and liquidity providers separately. We are looking for an equilibrium where the large investor accepts the offer.

Liquidity Trader. Assume the liquidity trader offers p to the continuum of mass x and the other large trader accepts. Then the liquidity trader faces a maximization problem similar to that in (2) for inventory 0

$$\pi^{\text{da}}(p) = \max_x -p \cdot x \frac{\alpha\rho + 1}{1 + \alpha} - \frac{c}{2}x^2 - \frac{\ell}{2} \left(Q - x \frac{\alpha\rho + 1}{1 + \alpha} \right)^2, \quad (7)$$

where henceforth superscript da signifies *dispersed* ownership and large traders *accepting* the offer. Assuming on-path play, in the next period, with chance 1/2, the trader receives a shock and earns payoff $\pi^{\text{da}}(p)$, and with chance 1/2 the other receives a shock, in which case the trader receives payment p for the fraction of shares that he gets to trade. This fraction is $x^{\text{da}}/(1 + \alpha)$, where x^{da} is the quantity that maximized profits in (7). The per-period payoff in equilibrium where the offer of price p is accepted is thus:

$$\frac{1}{2} \frac{px^{\text{da}}}{1 + \alpha} + \frac{1}{2} \pi^{\text{da}}(p).$$

The equilibrium payoff from offering p is therefore:

$$\Pi^{\text{da}}(p) = \pi^{\text{da}}(p) + \frac{\delta}{1-\delta} \left(\frac{1}{2} \frac{px^{\text{da}}}{1+\alpha} + \frac{1}{2} \pi^{\text{da}}(p) \right). \quad (8)$$

Alternatively, the liquidity trader can deviate, offer a price of 0, which would still be accepted by small investors with probability ρ but not by the public keys that belong to the large investor. In this case, the liquidity trader solves the same problem as in (2) for inventory 0 except that the acceptance probability is not ρ but $\alpha\rho/(1+\alpha)$. To maximize the stage payoff from this deviation, he chooses x to maximize:

$$\pi^{\text{dr}}(0) = \max_x -\frac{c}{2}x^2 - \frac{\ell}{2} \left(Q - x \frac{\alpha\rho}{1+\alpha} \right)^2, \quad (9)$$

where superscript dr signifies *dispersed* ownership and large traders *rejecting* the offer.

The trigger strategy prescribes that the liquidity trader trades in this manner (with small traders and the intermediary only) in every period, by offering the price of 0 to the continuum, and that public keys that belong to the other large (“non-shocked”) investor reject this offer. Since each large investor is shocked with probability 1/2 each period, the stage payoff is $\pi^{\text{dr}}(0)/2$.

The payoff to deviating (offering a price of 0) is therefore:

$$\Pi^{\text{dr}}(0) = \pi^{\text{dr}}(0) + \frac{1}{2} \frac{\delta}{1-\delta} \pi^{\text{dr}}(0).$$

For the liquidity trader to be willing to offer $p \geq 0$, the equilibrium price must satisfy:

$$\pi^{\text{da}}(p) + \frac{\delta}{1-\delta} \left(\frac{1}{2} \frac{px^{\text{da}}}{1+\alpha} + \frac{1}{2} \pi^{\text{da}}(p) \right) \geq \pi^{\text{dr}}(0) + \frac{1}{2} \frac{\delta}{1-\delta} \pi^{\text{dr}}(0). \quad (10)$$

Liquidity Provider. Assume the liquidity provider is offered p . Then the stage payoff is $px^{\text{da}}/(1 + \alpha)$, because, in equilibrium, he will receive a request for trading quantity $x^{\text{da}}/(1 + \alpha)$. His future on-path payoff is the same as for the liquidity trader, and thus his on-path equilibrium payoff is

$$\Pi^{\text{LP}}(\text{accept the equilibrium } p) = \frac{px^{\text{da}}}{1 + \alpha} + \frac{\delta}{1 - \delta} \left(\frac{1}{2} \frac{px^{\text{da}}}{1 + \alpha} + \frac{1}{2} \pi^{\text{da}}(p) \right).$$

The liquidity provider could also deviate and front-run for $x^{\text{da}}/(1 + \alpha)$ (we assume in this version that he front-runs for the entire quantity). We further assume in this version that if the liquidity trader has already contacted the continuum of small traders as part of the equilibrium, then he will not go back to the continuum but must trade either with the front-runner or with the intermediary.¹⁶ The liquidity provider's payoff to front-running is then similar to that computed in Lemma 1,

$$\pi^{\text{fr}} = -\pi^{\text{c}}(x^{\text{da}}/(1 + \alpha), 0) = \frac{\ell}{2} (x^{\text{da}}/(1 + \alpha))^2.$$

Note that in this setting, although the intermediary has a non-zero inventory at the time when he is approached by the front-runner, the additional cost is a wash between the front-runner and the liquidity trader. Once this trader performed front-running, the trigger strategy is invoked; the continuation payoff for the liquidity provider is the same as for the liquidity trader.

Taken together, the payoff to the liquidity provider if he deviates and front-runs the

¹⁶For fixed parameters one can find a cost parameter c that is sufficiently high such that this assumption is warranted, provided that after contacting x_0 , the incremental cost of contacting an additional mass of x_1 within the same stage game is $C(x_0 + x_1) - C(x_0)$.

liquidity trader is:

$$\Pi(\text{front-running}) = \frac{\ell}{2} \left(\frac{x^{\text{da}}}{1+\alpha} \right)^2 + \frac{1}{2} \frac{\delta}{1-\delta} \pi^{\text{dr}}(0).$$

For the liquidity provider not to front-run, he must be offered $p \geq 0$ such that

$$\frac{px^{\text{da}}}{1+\alpha} + \frac{\delta}{1-\delta} \left(\frac{1}{2} \frac{px^{\text{da}}}{1+\alpha} + \frac{1}{2} \pi^{\text{da}}(p) \right) \geq \frac{\ell}{2} \left(\frac{x^{\text{da}}}{1+\alpha} \right)^2 + \frac{1}{2} \frac{\delta}{1-\delta} \pi^{\text{dr}}(0). \quad (11)$$

We provide the following equilibrium characterization.

Proposition 3 (Equilibrium with Dispersed Ownership): *When the intermediated market is sufficiently liquid (ℓ is sufficiently small), or when the discount factor δ is sufficiently large (the future is important), there exists an equilibrium with no front-running where the large liquidity trader trades with the intermediaries, and he trades with the continuum of small and large traders at price p that is arbitrarily close to 0.*

If the market is very liquid (ℓ small), then front-running is not profitable. Both the stage payoff to deviating and the future cost decline as the market is more liquid (ℓ declines), however, the costs decline proportional to ℓ^2 whereas the benefit declines at rate ℓ^3 . When future interactions and payoffs are sufficiently important (δ is large), e.g., when investors interact sufficiently frequently, front-running can also be avoided because future benefits of being able to trade with the other large investor when hit by a liquidity shock outweigh the one-time profits that can be obtained by front-running. This latter result is a standard Folk Theorem.

Finally, for very low values of δ , numerically we find that even when there exist prices such that front-running can be curbed ((11) holds), these prices may be too high for the liquidity trader to be willing to offer them ((10) is violated). Instead, the liquidity

trader prefers to over-trade with the intermediary. Figure 2 provides an example for a setting where $Q = 10$, $c = .01$, $\rho = .1$, $\ell = 2$ (e.g., $\sigma = 1$, $N = 2$, and $\kappa = 2$). The horizontal axis is for values $\delta \in (0.001, 0.02)$, the vertical axis is for $\alpha \in (10, 100)$. The red area signifies the combinations of δ and α such that the minimum price that would make large traders not front-run one another is too high. Small values of δ capture high discounting of future payoffs, and can be viewed, e.g., to describe situations where liquidity shocks are rare.

VII. Comparing Transparency Regimes

We have described trading under four market design regimes: (1) all trades go through an intermediary; (2) there is a one-to-one mapping from ownership to public keys and traders interact anonymously in the sense that they cannot see the other traders' size; (3) there is a one-to-one mapping from ownership to public keys and investors can see each other's size, and (4) there is a one-to-one mapping from shares to public keys and traders therefore cannot infer the size of others.

Setting (1) captures the world as it is where peers cannot interact directly; arguably peer-to-peer trading is possible in limit order markets where investors can post orders, but the reality of public equity markets is that a large fraction of trades involves high-frequency traders who act as intermediaries. Moreover, trades in markets such as those for corporate bonds and options almost always involve intermediaries.

In our opinion, settings (2)-(4) capture the types of market organization that are broadly compatible and possible with Blockchain technology. These settings share aspects with other settings in the literature, but they also contain features that are unique to distributed ledger technology. The contribution of our work is to study transparency

and its impact on trading decisions and costs, focusing on the situations where the level of transparency is endogenous to the different possible mappings of ownership-to-public keys.

Setting (2) describes a situation where there is a system that allows peer-to-peer transactions, but where there is either full (chosen) anonymity in the sense that investors do not know each others' size or where the distributed ledger of the blockchain is kept behind closed doors; following the discussion among market participants, such an arrangement (sometimes referred to as a “private” blockchain) is currently what financial institutions envision for the implementation of the technology. Arguably, this setup resembles a stylized version of today’s equity markets where investors who trade large orders trade on a number of trading venues. Front-running plays no role in this setup because the chance of meeting a potential front-runner is zero and because large traders are not identifiable as such (all peer-to-peer contacts involve small trades only).

Setting (3) captures full transparency with regard to investor size: counterparties know the other “large” players in the market (even though they will not know which institution is behind a public key) and they can (and will) use technology to contact prospective counterparties directly. For large investors, identifying a prospective counterparty is essentially costless, and the focus is on possible front-running costs. This setting captures an idealized situation where liquidity sourcing for large orders is cheap, and where large investors can easily identify, trade with and remember one another.¹⁷

Setting (4) resembles the setting of (2) where investors contact each other via a black box mechanism, except that large investors split their ownership across multiple public keys. Whenever a large investor tries to access the peer-to-peer market in this setup,

¹⁷One can also imagine implementing “memory” with a counterparty rating system akin to those in place at online market platforms such Amazon or Ebay.

the other large investor receives multiple trading requests with a positive probability – but it is not possible for the liquidity trader to only contact the other large investor. Having multiple public keys allows the investor on the liquidity providing side to gather information about another investor’s trading interest. In this setup as in setting (3), there is a risk of being front-run. In contrast to setting (3), where front-running may be countered by a poison-the-well trigger strategy according to which no large-to-large peer-to-peer would ever occur again, here the trigger strategy involves large investors trading an excessively large portion with the intermediary.

End-investors in our model are risk-neutral. Therefore the more trades occur with the risk-averse intermediary, the lower is overall welfare. Interactions with the intermediary also lead to short-term deviations of the market price from the fundamental value. Interactions between large traders are zero-sum. An additional cost in the model is finding counterparties among a continuum of traders, a process that is assumed to be subject to complexity costs. Finally, in setting (4) it is possible that in equilibrium, traders must pay a premium (price concession over the expected fundamental value) to all traders who they approach. Thus small retail-type traders may obtain rents from large traders in the absence of full transparency.

Proposition 4 (Payoffs for Liquidity Traders):

1. *Payoffs in all three settings with peer-to-peer interaction dominate those under pure intermediation.*
2. *Payoffs in a setting with transparent, concentrated ownership dominate payoffs in the two settings where investor holdings are not publicly known.*
3. *The ranking of payoffs for the two settings with non-transparent ownership depends*

on whether the model parameters admit an equilibrium where large investors trade with each other in the dispersed ownership setting:

- (i) *When large investors do not trade with one another in the setting with dispersed holdings, payoffs in the setting with concentrated holdings dominate those with dispersed holdings.*
- (ii) *When large investors trade with one another in equilibrium and the equilibrium price concession p is sufficiently small, the payoffs in the setting with dispersed holdings dominate those with concentrated holdings.*

The best-case scenario for large investors in our setting is concentrated and transparent ownership. In this case, large investors only trade with each other, avoiding all the direct costs of trading (large investors incur no complexity costs for contacting small investors and do not have to interact with the intermediary). The indirect (front-running) costs of trading are also zero because of repeated interactions.

Payoffs for large investors in the non-transparent setting where each investor concentrates his ownership under a single key, but individual holdings are not known, are lower than in the transparent setting, but still dominate those under pure intermediation. Front-running in this setting plays no role as it is a measure-zero event. Fewer trades occur with the intermediary compared to pure intermediation and, by construction, this behavior yields higher payoffs to large traders.

When investors do not concentrate their holdings under a single public key but instead spread their holdings (in a one-share, one-key manner), the analysis becomes more complex. First, when liquidity in the intermediated market is low, front-running cannot be prevented by incentivizing potential front-runners. Instead, in equilibrium, large investors offer $p = 0$ to the continuum and reject each other's offers. The payoffs

for large investors in this latter outcome are dominated by those in the non-transparent, concentrated ownership setting: large investors face lower probability of acceptance when trading with the continuum of public keys (since large investor public keys are “dispersed” in this continuum but reject all offers), and as a consequence, large investors “overuse” the intermediaries. Still, even this setting is superior to pure intermediation because large investors can trade at least some of their holdings with small investors.

If the intermediated market is very liquid (ℓ small), if investors are patient (δ is close to 1), or if there are many small investors relative to the trade-size of large investors (α is large), then front-running can be curbed at a minimal price concession p . In equilibrium, large investors accept each other’s offers, therefore the probability of acceptance when trading with the continuum and (for small enough p) payoffs for large investors are higher than in the setting with non-transparent but concentrated ownership.

In sum, in this stylized setting, it is optimal for large investors to interact repeatedly and transparently because this scenario permits large investors to capitalize on the discipline afforded by repeated interactions. In the absence of full transparency, the mapping between public keys and investor holdings plays a key role in minimizing the costs of intermediation. If the intermediated market is illiquid or if there are only few small investors relative to large investors, then the cost of front-running outweighs the benefits of large investors finding each other, and it is best to require concentrated ownership so that investors cannot split their holdings across multiple public keys. If the intermediated market is liquid or if there are only many small investors relative to large investors, then large investors are better off if they all split their holdings across multiple public keys, in order to maximize their chances of trading with each other.

VIII. Conclusion

In our opinion, blockchain technology should be viewed as a transformative technology that has the capacity to fundamentally change the interactions among market participants. With this paper, we hope to contribute to the debate on how to best structure and design markets with the advent of this new technology. We believe that there are three features of a distributed ledger that are critical and that allow new market designs relative to the current world of security trading. The first is the electronic nature of blockchain securities where investors can program trading rules directly into the security (so that they become “smart contracts”). Current market arrangements may, in principle, allow investors to specify such rules, but they likely require very complicated and thus costly system setups, including the integration within the brokers’ existing systems. The second is the possibility and simplicity of peer-to-peer interactions. Arguably, current markets also allows such interactions, e.g., for block-trades, but, again, these trades are few, expensive to arrange, and they are only available to select parties. The third is the possibility of linking ownership to a public key. This feature potentially allows market participants to contact one another directly — they don’t need to go through a third party or a separate system. Thus instead of having to search for liquidity by scanning different systems, hoping for a counterparty to show up at the right time and at the right marketplace, investors can contact other investors (or, rather, a set of public keys that are anonymously linked to other investors) in a directed search and thus source liquidity directly. In current markets, such liquidity sourcing does exist,¹⁸ but it usually occurs at the broker level, which, arguably, is a friction. In

¹⁸Examples are Liquidnet’s dark pool or ITG’s POSIT Alert system. For corporate bonds, Algomi operates a peer-to-peer matching system.

this paper, we study several possible market designs that incorporate these features, focussing on the mapping between public keys and end-investors. We argue that this mapping should be a market design choice.

At this point in time, market participants and regulators still do not know how the transformation of the current trading arrangement to blockchain securities will manifest itself. A simplistic view, in our opinion, would be to merely see it as an efficiency enhancement in the sense that blockchain-based trading simplifies and speeds-up the trade-and-settlement process. While the simplification is certainly part of the technology, there is more to it. First, the intrinsic verification of ownership in principle eliminates the necessity of investors to hold accounts at brokerages (either for cash or for securities). Second, the nature of a distributed ledger is that information on ownership is available to entities other than the account manager or owner. Third, the electronic nature of blockchain securities allows investors to build-in liquidation rules in the sense that investors could design their holding as a standing limit order that gets executed as soon as someone offers the right price. In other words, it is imaginable that (some) investors are always in the market, with implications for portfolio management and asset pricing.

REFERENCES

Ariel Rubinstein, Asher Wolinsky, 1985, Equilibrium in a market with sequential bargaining, *Econometrica* 53, 1133–1150.

Bali, Shagun, and Terry Roche, 2015, Blockchain technology: Pushing the envelope in fintech, Industry report TABB Forum.

Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *The Journal of Finance* 48, 157–185.

Brummer, Chris, 2015, Disruptive technology and securities regulation, *Fordham Law Review* forthcoming.

Christoffersen, Susan Kerr, Erfan Danesh, and David K. Musto, 2015, Why do institutions delay reporting their shareholdings? Evidence from form 13f, Working Paper No. 2661535 Rotman School of Management Working Paper.

Cujean, Julien, and Remy Praz, 2015, Asymmetric information and inventory concerns in over-the-counter markets, *Working Paper*.

Danesh, Erfan, 2015, Strategic trading and delayed disclosure by informed traders, Working paper Rotman School of Management, University of Toronto.

De Frutos, M. Ángeles, and Carolina Manzano, 2002, Risk aversion, transparency, and market performance, *The Journal of Finance* 57, 959–984.

Degryse, Hans, Mark Van Achter, and Gunther Wuyts, 2012, Internalization, clearing and settlement, and liquidity, Working paper European Banking Center Discussion Paper No. 2012-001 Available at SSRN: <http://ssrn.com/abstract=1983197>.

Diamond, Peter A, 1982, Aggregate Demand Management in Search Equilibrium, *Journal of Political Economy* 90, 881–94.

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–1847.

Evans, David S., 2014, Economic aspects of bitcoin and other decentralized public-ledger currency platforms, Discussion Paper, No. 685 University of Chicago Coase-Sandor Institute for Law & Economics Research Paper.

Gehrig, Thomas, 1993, Intermediation in search markets, *Journal of Economics & Management Strategy* 2, 97–120.

Harvey, Campbell R., 2015, Cryptofinance, Working paper Fuqua School of Business.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.

Lee, Larrisa, 2016, New kids on the blockchain: How bitcoin's technology could reinvent the stock market, *Hastings Business Law Journal* forthcoming.

Lester, Benjamin, Guillaume Rocheteau, and Pierre-Olivier Weill, 2015, Competing for order flow in otc markets, *Journal of Money, Credit and Banking* 47, 77–126.

Miao, Jianjun, 2006, A search model of centralized and decentralized trade, *Review of Economic Dynamics* 9, 68 – 92.

Pagano, Marco, 1989, Trading volume and asset liquidity, *The Quarterly Journal of Economics* 104, 255–274.

Vayanos, Dimitri, and Tan Wang, 2007, Search and endogenous concentration of liquidity in asset markets, *Journal of Economic Theory* 136, 66 – 104.

Weill, Pierre-Olivier, 2002, The liquidity premium in a dynamic bargaining market, Working paper Stanford University.

Yavas, Abdullah, 1996, Search and Trading in Intermediated Markets, *Journal of Economics & Management Strategy* 5, 195–216.

Yermack, David, 2016, Corporate governance and blockchains, Working Paper No. 21802 NBER.

Yin, Xiangkang, 2005, A comparison of centralized and fragmented markets with costly search, *The Journal of Finance* 60, 1567–1590.

Figure 1
Google Search Terms and Articles Containing the Word “Blockchain”

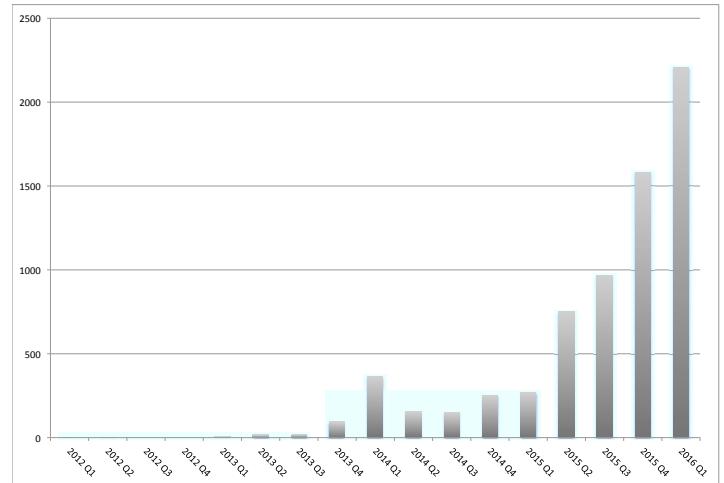
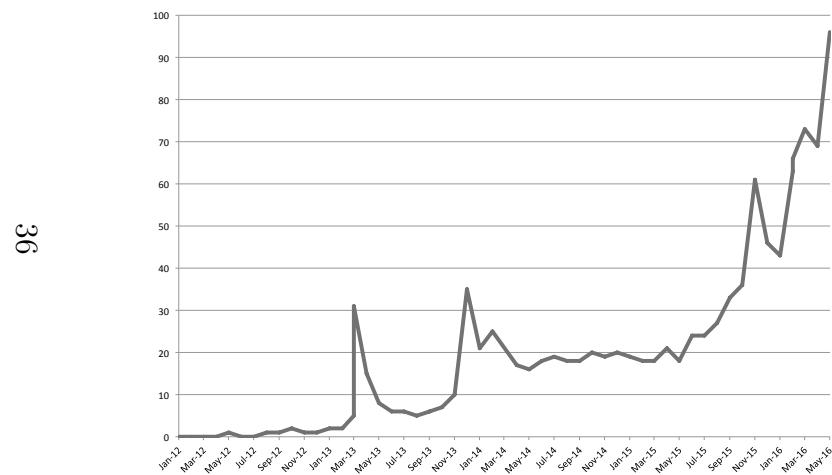


Figure 2
Illustration of Parametric Configurations when Costs of Preventing Front-Running are too high

Figure 2 provides an example for a setting where $Q = 10$, $c = .01$, $\rho = .1$, $\ell = 2$ (e.g., $\sigma = 1$, $N = 2$, and $\kappa = 2$). The horizontal axis is for values $\delta \in (0.001, 0.02)$, the vertical axis is for $\alpha \in (10, 100)$. The red area signifies the combinations of δ and α such that the minimum price that would make large traders not front-run one another is too high.

