Cryptographic Techniques for Transparent and Efficient Markets

Isaac Fry
Computer Science
Colorado School of Mines
ifry@mines.edu

Keenan Schott
Computer Science
Colorado School of Mines
keenanschott@mines.edu

Ethan Ko
Computer Science
Colorado School of Mines
ethanko@mines.edu

Garrett Weissert
Computer Science
Colorado School of Mines
gweissert@mines.edu

Hanuman Chu
Computer Science
Colorado School of Mines
hkchu@mines.edu

***Abstract*—This proposal aims to investigate the potential integration of Homomorphic Encryption (HE) and Secure Multiparty Computation (SMPC) in financial markets to strengthen the balance between market integrity and market transparency. Traditional markets, which are subject to regulations like the National Best Bid and Offer (NBBO) face a dilemma: promoting fairness and liquidity while also preventing the misuse of public information by market manipulators, like high-frequency traders (HFTs) and large institutional investors in dark pools. The combination of HE and SMPC offer a secure cryptographic solution by enabling transactions to be provably correct without any participating party knowing information about the transaction. This research surveys the existing literature to assess both the practicality of implementing the computationally intensive processes for HE and SMPC and the practicality of convincing governmental regulatory bodies to approve new technologies. We conclude that for the implementation of HE and SMPC into a market, the market must exhibit three characteristics: time unconstrained, a true double auction with few parties, and non-iterative. While modern stock exchanges do not exhibit these characteristics, there are other types of markets that would be excellent candidates for HE and SMPC.**

***Index Terms*—homomorphic encryption, secure multiparty computation, financial cryptography, high frequency trading, dark pools, market trust, market integrity**

I. Introduction

Financial market dynamics depend on the structured mechanism of orders. Many markets organize trades in order books, which act as double auctions that match and clear bids from multiple parties in a way that satisfies supply and demand [1, 2, 3]. Exchanges in the United States must adhere to the National Best Bid and Offer (NBBO) and precedence for auctions [4, 3, 5]. The public transparency of auctions ensures that transactions are fair, but the high quality and quantity of information available affects asset pricing and market liquidity, especially for large block orders from institutional traders [6, 7, 8, 9].

Auctioneers assume responsibility for correctly matching bids without unfairly skewing orders towards either supply or demand. When parties submit orders, they specify the price at which they will buy or sell a predetermined quantity [3, 8]. Trades occur at the point where buy and sell orders match in price. Auctioneers are generally required to clear higher bids and lower sells first, and in the case of multiple equal orders, they must then clear the orders that came first [3]. Furthermore, orders can be specified as either market orders or limit orders [3]. Market orders force the auctioneer to fulfill orders at the best current buy or sell price, while limit orders specify a minimum sell price or maximum buy price. Traditionally, the auctioneer, whether it be an exchange or otherwise, will need information on all current orders to calculate the correct clearing price, which in turn must reflect the NBBO [5].

Maintaining ideal transparency equilibrium is difficult. On one hand, dark pools are private exchanges that eliminate the public price-signaling effects of large block orders while still proving adherence to NBBO [8]. However, the lack of transparency creates opportunities for misuse, especially seeing as private exchanges act on incentives that can skew transactions in their favor [10, 4, 9]. On the other hand, full transparency can inform market anticipators, like high-frequency traders (HFTs), and enable market manipulation [7, 9]. Thus, markets are caught in a dilemma of needing to prove the correctness of transactions without revealing information about the transaction.

Homomorphic Encryption (HE) and Secure Multiparty Computation (SMPC) offer a solution to finding ideal transparency [11, 4]. HE enables computation on encrypted data via the mapping of homomorphic mathematical structures [11]. The goal of fully HE (FHE) is to perform any computable function on encrypted data an unlimited amount of times, though there are partial HE and somewhat HE techniques that restrict either the number of computations or the number of times computations can be done [4]. Leveraging homomorphic structures requires more computational power than performing the same operations on the unencrypted plaintext [12]. Using HE, a clearing price could be calculated by an auctioneer without knowing the pricing information of the order book, negating some of the concerns about excessive transparency in markets [13, 12].

Furthermore, SMPC allows for multiple parties to collaborate on the computation so that an order is provably correct

even in cases of dishonest parties, which increases trust in the market [10, 4, 2, 3]. In this way, multiple parties create a "virtual auctioneer" that still computes clearing prices, but at no point does any bidding, selling, or auctioneering party ever possess complete information. SMPC also allows for semi-honest computations that are still correct even if some parties are intentionally malicious [4, 2]. By combining HE and SMPC, exchanges can prevent pricing manipulation while still proving the correctness of auctions.

II. Key Idea and Approach

This report explores how HE and SMPC could maintain market integrity while mitigating the adverse effects of excessive transparency. Via a survey of current literature, we seek to describe the practical feasibility of financial markets implementing both HE and SMPC into their trading microstructures. The rapid pace of auctions in active markets and the cryptographic demands of HE and SMPC would make these techniques unwieldy for practical use, but it is unclear to what extent these detrimental factors would be realized in true market scenarios - there seem to be no wide-scale implementations of this scheme in active markets. Compiling information from different pieces of contributing literature assist in building a definitive guide to how HE and SMPC could be used within markets.

The goal of this project is to analyze HE and SMPC in current financial markets, but to limit the scope in such a way as to exclude cryptocurrency markets. Blockchain currencies are inherently decentralized, which creates regulatory hurdles for governmental agencies that require the transaction of financial instruments to be centralized. In fact, the SEC has struggled to categorize cryptocurrency under the regulatory umbrella of financial securities in a way that both respects the decentralized nature of the currency and protects investors [14]. The use of HE and SMPC in this project is not to morph existing exchanges into a blockchain-esque structure, but instead to maintain both the current integrity and liquidity of the market microstructure while increasing the crytographic security of the market.

III. Research Methodology

We believe that a full auction simulation using HE and SMPC with both market and limit capabilities is outside the scope of this project. If we were to attempt to create a fully operational market system with HE and SMPC, we would face a variety of difficulties:

- Double auction markets are somewhat complex to implement, and running a simulation of a double auction market would require encoding true market heuristics, which are inherently vague and ambiguous.
- HE and SMPC are highly complex techniques, and it would require significant engineering to implement into a market simulation.
- Market pricing data is proprietary, and any market simulation without reliable market data is a poor representation of the entity we wish to simulate. Furthermore, accessing true market data or interacting with real exchanges is cost-prohibitive, and many financial firms closely guard their financial data because any information advantage can easily be flipped into a monetary advantage over other market participants (this issue leads to the original value proposition of dark pools).

Instead, we opt to compile existing research, which has covered discrete portions of the problem in substantial detail [1, 10, 4, 2, 12, 6].

For a full list of literature, see the References portion at the end of the paper. However, a piece of given literature generally falls into one of the following categories:

1) **Financial analysis**: seeing as this application of cryptographic techniques is incredibly domain-specific, we obtain general domain knowledge of financial markets and stock exchanges. This is where we are able to gain knowledge about high-frequency traders (HFTs), dark pools, regulations, and more.
2) **Cryptographic analysis**: the theory behind HE and SMPC is complex, and requires definitively technical literature that has little to do with financial systems.
3) **Case studies**: some literature applies HE and SMPC schemes to financial markets. For instance, the Danish sugar beet contracting auction or the money laundering investigation paper are both great literature examples of using HE and SMPC schemes, both in terms of the model mechanics and the practical limitations of the techniques.

Utilizing the data and simulations of existing literature still enables both quantitative and qualitative analysis without the time or cost demands of running full simulations. Thus, we are confident in our ability to draw conclusions on using HE and SMPC in financial markets without producing a technical simulation.

IV. Mathematical Background

A. Homomorphic Encryption (HE)

Homomorphic encryption is achieved by making a ciphertext space such that operations within it results in the encryption of the equivalent operations in plaintext space. A simple example of this would be an encryption function

$$E(M,k) = kM \tag{1}$$

This would would allow for addition to be done a number of times because

$$E(M_1,k)+E(M_2,k) = kM_1+kM_2 = E(M_1+M_2,k) \tag{2}$$

A similar example, using multiplication, would be

$$E(M,k) = M^k \tag{3}$$

Thus, multiplying two encrypted messages would be

$$E(M_1, k) * E(M_2, k) = M_1^k M_2^k \tag{4}$$
$$= (M_1 M_2)^k \tag{5}$$
$$= E(M_1 * M_2, k) \tag{6}$$

For fully homomorphic encryption (FHE), we need to be able to perform both addition and multiplication homomorphically, have secure encryption (unlike the above examples), and to limit persistent increases in noise as operations are performed; excess noise can result in a ciphertext losing the ability to be decrypted.

One encryption scheme that gets close to fulfilling the above requirements is CKKS. This scheme uses polynomial rings for the plaintext and ciphertext space and a random ternary polynomial with the same degree as the polynomial rings for the secret key, S[15]. Next, we choose a random polynomial from the polynomial rings, a, and choose a random error polynomial, e, with the same degree as the polynomial rings and bounded by some small value [15].

Our public key, P, is then built from this equation [15]

$$P = (-aS + e, a) \tag{7}$$

In this way, we can encrypt messages by creating a message unique key, u, in the same format as the secret key as well as 2 new random error polynomials, e_1 and e_2, and combining them resulting in the following equation for the ciphertext C [15]

$$C = P_0 u + e_1 + M, P_1 u + e_2 \tag{8}$$

We can then decrypt the message by adding the first part with the second multiplied with S because the following equivalence holds

$$C = P_0 u + e_1 + M + S(P_1 u + e_2) \tag{9}$$
$$= (-as + e)u + e_1 + M + Sau + Se_2 \tag{10}$$
$$= -Sau + eu + e_1 + M + Sau + Se_2 \tag{11}$$
$$= M + eu + e_1 + Se_2 \tag{12}$$

Since S and u are ternary and the error polynomials are very small, it is approximately the original message. Addition in CKKS is homomorphic because of the following

$$C = (P_0 u_a + e_1 a + M, P_1 u_a + e_2 a) \tag{13}$$
$$+ (P_0 u_b + e_1 b + M, P_1 u_b + e_2 b) \tag{14}$$
$$= (P_0(u_a + u_b) + e_1 a + e_1 b \tag{15}$$
$$+ M + M', P_1(u_a + u_b) + e_2 a + e_2 b) \tag{16}$$
$$= E(M_a + M_b) \tag{17}$$

To do multiplication, we have to create an additional term in the ciphertext for $c_1 a c_2 b + c_1 b c_2 a$. Afterwards, we then relinearize to reduce the size back to two terms and scale the result down to help reduce noise [15]. The full proof for homomorphousness would take too much space to prove in this paper because of the vast amount of terms and is generally outside the scope of a literature review.

Homomorphic encryption can be used for computing a clearing price in a number of schemes, one of which was developed by Oliver Baudron and Jacques Stern [16]. This scheme involves bidders encrypting their bid with the other parties' public keys and sending all those ciphertexts to a third party which then performs the necessary operations to find a clearing price and sends the result back to all the bidders who can then decrypt the results to find whether their bid was accepted [16]. This scheme is one of the best available, preventing any single party from seeing others' bids and stopping bidders from cheating by requiring a proof that their bid was encrypted fairly. However, even the state of the art isn't perfect, which can be seen in how the third party can decrypt all bids by colluding with one of the bidding parties [16].

B. Secure Multi-Party Computation (SMPC)

Secure multi-party computation allows participants to share inputs needed for a functional calculation while still keeping their inputs private.

To find the sum of the secrets, let m represent the number of computing parties. It is assumed that all parties' inputs are of equal length such that

$$|s_i| = |s_j| \tag{18}$$

Participants break up their secret input s into shares equal to the number of parties $s_1, s_2, \ldots, s_m$. This is commonly done with Shamir's Secret Sharing Scheme where participant i picks $m-1$ random numbers $a_{i1}, a_{i2}, \ldots, a_{im-1} \in \mathbb{Z}_q$ and evaluates the polynomial

$$p_i(x) = s_i + a_{i1}x + a_{i2}x^2 + \ldots + a_{im-1}x^{m-1} \tag{19}$$

Each polynomial is constructed such that $p_i(0) = s$ and is of degree $m-1$. Each party is then given a share $s_i = p_i(i)$ for $i = 1, 2, \ldots, m$. Each computing party then calculates the sum from the shares they were given and publishes their result. After all parties have published their results, polynomial interpolation is used on the published results to find $p_i(0)$ and get the final result [17].

V. Financial and Regulatory Results

In an ideal world, the impact of HE and SMPC would positively impact financial markets in substantial ways. Theoretically, the combination of HE and SMPC would help achieve a transparency equilibrium - where trade secrecy is preserved without sacrificing informed liquidity - and increase the overall efficiency of the market [10, 4, 2]. Public-facing exchanges would be protected from HFTs, while dark pools would be provably correct in their transactions. Since financial markets involve such a variety of participants, including both governmental and private entities, at every point in the value chain, we would expect some aspects of HE and SMPC to

be more valuable to some participants than others, though all parties in the financial environment would likely benefit from better security.

Furthermore, in an ideal world, we would need to significantly compensate for the lack of informational liquidity in the market. We would expect that a decrease in the transparency of order books would need to be supplanted in some way by aggregated and anonymized supply and demand statistics to still encourage market liquidity [7]. This is possible, but would potentially change the trading tactics of some firms seeing as some trades take into account the identity of the other parties. However, anonymizing and aggregating the data is generally under the domain of regulating agencies, which considers a different set of constraints that cryptographic researchers.

Finally, the regulatory landscape would need to change in order to accommodate the encrypted data. For instance, financial crimes, like insider trading, might be more difficult to monitor due to the secrecy of pricing information [12]. Seeing as all exchanges, including dark pools, are subject to scrutiny by regulatory bodies like FINRA and SEC [18], the uptake of HE and SMPC in financial markets would require substantial governmental support, both from an accountability perspective and an infrastructure perspective. Politically, this is a difficult proposition to navigate, and it would likely depend on the policies of the current presidential administration and current session of Congress.

Overall, from a financial and regulatory perspective, the switch to HE and SMPC would be viable, though any regulatory changes would require the governing agency to give all affected parties plenty of due notice, probably on the timescale of multiple years.

VI. Cryptographic Results

While the financial and regulatory results are promising, the cryptography makes the implementation of HE and SMPC nearly infeasible for modern stock markets. The core mathematical grounding of HE and SMPC are both well-founded and very well researched, but these techniques are severely limited by their operating environment. These limitations are abundant in the highly interactive, rapidly evolving environment of stock exchanges. We note two primary limitations to the implementation of HE and SMPC.

A. Orders must be cleared in a timely manner

As one paper notes, the speed of modern financial markets essentially approaches the speed of light [19]. HFTs and other trading parties are pressing up against hard boundaries of physical limitations, and any sort of slow down is considered definitively detrimental to the party engaging in the exchange. In many ways, the main bottleneck for exchanges is the communication time; the computational time to calculate the optimal clearance is considered negligible and on the microsecond scale [9, 20]. Thus, adding a significant computational overhead to order clearance goes against conventional wisdom in the financial sector, and it would significantly disrupt the proceedings of the exchange.

Of course, there is debate as to whether HFTs, which would be the party most negatively affected by a computational overhead, are beneficial to stock exchanges. It could be argued that stock market orders are cleared in too much of a timely manner. Some researchers posit that a slowdown in the operating speed of the exchange is actually beneficial to the performance of the market [7].

Another paper discusses the viability of introducing intentional delays into stock markets, which is what HE would impose. The paper observes a natural experiment with the NYSE both during and after they artificially imposed a 350 microsecond delay into the clearing of orders. Comparing the state of the market during the active enforcement of the delay and after its subsequent removal, the authors conclude that restricting speed is likely to be damaging, both in terms of liquidity and overall market quality [20]. Thus, the computational overhead of HE would be damaging to the market to the extent that it negates the elimination of HFTs. This would be highly counterproductive.

Overall, both HE and SMPC contribute to a higher latency in the exchange environments. HE significantly increases the computational demand, to the point that it dwarfs the communication time [12]. SMPC would also increase the communication time, seeing as communication would not just take place within an exchange, but between multiple parties and, potentially, multiple exchanges. Considering that parties in these exchanges are already radically competitive in how they optimize their communication latency [9], it seems rather unlikely that they would agree to requiring communication between multiple parties.

B. Results of order matching creates iterative responses to economic signals

Modern stock exchanges create economic pricing signals as orders are met and cleared. In essence, this is signaling the presence of demand and the corresponding presence of supply [7]. For markets to become truly competitive and to approach any semblance of balance, orders are submitted and responded to on an iterative basis, thus requiring substantial levels of computation and information clarity.

For instance, observe the quantitative finance institution of Renaissance Technologies, which is widely regarded to be the most successful quantitative trading firm in the world. Unique among other quantitative trading firms, Renaissance Technologies primarily employs non-finance academics that specialize in signal processing. At its core, the outsized success of Renaissance Technologies is dependent on the capture and analysis of massive amounts of data [21]. Some researchers speculate that Renaissance Technologies created a model that meaningfully analyzes both financial data and financial behavior. Reducing the iterative response capability by anonymizing the data would severely limit any practical insights for trading groups that are meaningfully engaging in the financial sector.

In modern stock exchanges, adding SMPC would indeed ensure the fulfillment of NBBO standards and regulations, but any sort of HE encryption scheme used within the SMPC

would reduce price discovery and market efficiency. While this is linked to the speed of price discovery, it's a wholly different issue: iteratively responding and reacting to economic signals within the market means that information about the financial instrument or tradable good must be entirely known, and fragmenting the market across multiple parties could plausibly lead to degradation of the market itself.

Note that HE cannot be removed from the notion of SMPC in any exchange-related scenario. A regulator would primarily enforce the adoption of SMPC to ensure that dark pools are truly adhering to NBBO. In order to distribute the computation among multiple parties, a dark pool would likely want to hide the market signals indicated in the information that is to be computed on (otherwise, the dark pool would be operating as a classical exchange with clear price discovery). In this way, HE would inherently become a part of the SMPC model.

C. Potential solutions

In order to negate the concerns noted above, there could be some steps taken to ensure the feasibility of these techniques, but they would be difficult to implement widely in modern stock exchanges:

- Adopting a slower computational time step: if an exchange wanted to slow down the market in hopes that market participants would realize the benefit of certifiably correct and yet transparent transactions, an exchange would be well within its rights to do this [1, 12, 7, 20]. However, as noted above, this could degrade the quality of the market.
- Using hardware accelerators for computation: the availability of FGPAs to become domain-specific accelerators for HE is promising, and it is an active area of research [22, 23]. Presumably, exchanges could implement accelerators to lower the time needed to clear a bid. However, performing computations on ciphertext via HE on an accelerator will likely never outperform computations on plaintext on the exchange's current state of the art. Further, hardware accelerators cannot break laws of physics, and communication times for SMPC would still be a significant bottleneck.
- Calculating compliance after clearance: As discussed below, HE and SMPC schemes could be used to prove correctness after a transaction takes place. The literature example of money laundering is a prescient example of this idea, and it certainly has strong applications [24]. In a limited scope, SMPC could technically be used to verify the NBBO compliance of dark pools after orders have cleared, which might be beneficial for general regulatory compliance. However, in many cases, back-calculations are not useful given the current real-time operating conditions of stock exchanges.

Thus, this combination of limitations and low feasibility makes a HE and SMPC scheme unlikely to be adopted by any modern stock exchange. Generally speaking, using these sorts of cryptographic techniques would not promote transparent and efficient market exchanges in the way that some literature aspired to argue.

On the other hand, there are scenarios where a HE and SMPC scheme would be rather useful and viable, which are discussed below.

VII. Discussion

Evidently, SMPC and HE in financial markets require a specific use case. In general, the application of a HE-SMPC scheme requires these three characteristics to be satisfied:

1) The market must be time unconstrained: there must be no pressure for orders to be cleared at a relatively rapid frequency
2) The market must be a double auction with few parties: an order book must be present for the benefits of HE and SMPC to be realized.
3) The market must be non-iterative: any auction that allows bidders to respond to price signals by other bidders will be entirely dysfunctional in an encrypted, information-sparse environment.

One particularly fascinating case was found that satisfied all three characteristics: agricultural bidding [2]. In Denmark, Danisco, the only processor in the Danish market, buys sugar beets from several thousand farmers. All bids go to an auctioneer, who computes, for each price, the total supply and demand in the market, as well as the price where total supply equals total demand. Since these bids reveal information on a farmer's productivity and business, Danisco would not be a trustworthy auctioneer given their position. On the other hand, Danisco uses contracts as security for farmers' debt. In brief, the auction cannot be run independently of Danisco. A third-party, such as a consultancy house, also proves to be an expensive solution. Thus, the auctioneer was decided to be three separate parties: Danisco, the sugar beet growers' association, and the researchers themselves. Each bidder sent their bid in an encrypted form to these three parties, who then performed computations on the data while it's still encrypted. The above scenario matches every requirement for a joint HE-SMPC use case.

In financial markets, particularly in the age of information, time is of the essence. Given that the need to computing sugar beet contracts was in a time unconstrained context, SMPC, despite its communication overhead, proved to be a valid solution. Furthermore, it was the correct auction type: in a double auction, SMPC provides a solution for matching supply and demand order books. There's also no iterative response - participants don't respond to other bids and attempt to match. Agricultural bidding only requires one round of SMPC usage, where a round is a sequence of uninterrupted computations with no input from seller or buyer (for an example of an iterative response auction, think of eBay: bidders seek to outbid other participants, which in turn signals demand). Moreover, there are only three parties serving as auctioneer - in a scenario where there are hundreds of parties, the overhead of SMPC would be entirely infeasible and undesirable.

In this way, the cryptographic techniques of HE and SMPC serve to create more transparent and efficient markets in non-iterative, double auction, time unconstrained scenarios. Certainly, these scenarios are limited in nature and not probabilistic given the rapid pace of modern information technology, but there are specific cases, like agricultural bidding, where it could be of use.

Another example where SMPC is put to good use is anti-money laundering algorithms [24]. Money laundering is a crime where the illegal source of money is concealed via a series of transactions, which is primarily accomplished by laundering transactions across multiple banks or financial institutions. Banks, however, only see a small subset of these transactions, specifically, the ones that occur through their organization. Due to consumer privacy law, banks cannot share this information with each other, and it would likely be disadvantageous to provide competitors with specific financial information from outside their institution.

In brief, banks wish to collaborate with each other without violating privacy. As with the Danish sugar beet market, a trusted third party could be used, but the security concern is still present: the data would simply sit on a third party server somewhere else, meaning that there is still a violation of banking privacy laws. The researchers argue that this is an optimal scenario for HE-SMPC schemes because each bank remains in control of their own data while allowing other banks to compute on their data.

The researchers propose a "risk propagation algorithm", where the input can be modeled as a directed graph with a weight matrix assigning a weight to every edge representing the aggregate transactions between two parties. Each node has an attribute value r_j between 0.0 and 1.0; this value can initially be determined or can be the result of the algorithm performed on a different network with the same nodes. The goal of the algorithm is to update the attribute value using the values of adjacent nodes.

Now, assuming the set of nodes is distributed over many parties, this is where SMPC is particularly valuable. If a bank wants to update the attribute value of one of its nodes, its neighbors need to send their corresponding attribute values in homomorphic encrypted form. This offers the benefit of the security discussed earlier and the simplicity of utilizing neighboring nodes in the graph to transmit data. In fact, an individual bank does not even need to be aware of the full graph and can receive information about the entire graph only via signals from neighboring nodes.

Similar to the previous example, this is a use case perfect for SMPC and HE. Given that a bank will typically be directly connected with only a few other banks, and thus nodes in the graph have few edges, the computation can be fairly trivially conducted. Although the researchers did not achieve real-time propagation, they achieved their graph model in intervals. In their experimentation, the authors observe that this scheme grows linearly with the number of nodes computed upon [24]. Computation can be performed across a graph of 200,000 nodes in just under 3 hours, and they believe that operations could be computed on a graph with millions of nodes in a reasonable amount of time, saying, it would only take a "a few days."

SMPC provides a clear advantage for this money laundering detection scenario, seeing as the computation is split across multiple institutions (i.e. parties) and the data is homomorphically encrypted before being shared. In fact, a HE-SMPC makes this type of analysis possible in an environment where regulations and norms of commerce essentially prohibit any sharing of information. In this way, the cryptographic techniques of HE and SMPC serve to create a more transparent and efficient market [6].

While these two literature examples paint an optimistic picture for specific use-cases of HE-SMPC schemes, it's clear that they cannot be used in stock exchanges. Therefore, it is unlikely for the cryptographic techniques of HE and SMPC to be used in making stock exchange markets more transparent and efficient, though it's theoretically possible given a number of concessions from the market [1].

VIII. Conclusion

In sum, this report investigated the feasibility and advantages of implementing homomorphic enryption (HE) and Secure Multiparty Computation (SMPC). Our literature survey shows that while HE and SMPC possess the potential to enhance market integrity and transparency, their practical implementation into rapidly evolving, high-speed financial environments like modern stock exchanges is currently infeasible, primarily due to significant computational overhead and the real-time demands of these exchanges.

However, these limitations should not exclude the possibility of HE and SMPC from being implemented into specific scenarios. We propose that markets wishing to implement HE and SMPC must exhibit the characteristics of being time unconstrained, double auctions, and non-iterative. In certain types of commodity auctions, like agricultural contract bidding, HE and SMPC offer substantial benefits, mainly a higher level of both regulatory compliance and security.

Future work regarding HE and SMPC may include investigating ways to increase the speed of both schemes. Similarly, an investigation into what extent modern market participants are using market data would be beneficial in order to draw conclusions on how the removal, aggregation, or anonymization of this data through the use of HE would affect markets. Another topic could assess how viable it is to perform historical validation of market data, as seen in the money laundering literature. In this way we may be able to harness the transparency desired changing any fundamental market dynamics.

Ultimately, this report surveys the existing literature on the use of HE and SMPC, as well as covers related topics to verify and conclude whether or not these cryptographic techniques can be utilized in markets to establish trust, transparency, and efficiency.

References

[1] C. S. Jutla and B. Mishra, "Improved stock market structure using cryptography," Cryptology ePrint Archive, Paper 2022/451, May 2022, [Online]. Available: https://ia.cr/2022/451. [Accessed March 14, 2024].

[2] P. Bogetoft, D. L. Christensen, I. Damgård, M. Geisler, T. Jakobsen, M. Krøigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. Schwartzbach, and T. Toft, "Secure multiparty computation goes live," in *FC 2009: Financial Cryptography and Data Security*, 2009, pp. 325–343, [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-642-03549-4_20. [Accessed March 14, 2024].

[3] L. Harris, "Trading and exchanges: Market microstructure for practitioners - draft copy," University of Southern California: Oxford University Press, 2022, [Online]. Available: https://www.acsu.buffalo.edu/~keechung/MGF743/Readings/Trading-Exchanges-Market-Microstructure-Practitioners/%20Draft/%20Copy.pdf.

[4] D. C. Parkes, C. Thorpe, and W. Lei, "Achieving trust without disclosure: Dark pools and a role for secrecy-preserving verification," Harvard EconCS Group, 2015, [Online]. Available: https://econcs.seas.harvard.edu/files/econcs/files/parkes_amma15.pdf. [Accessed March 14, 2024].

[5] A. Hayes and A. Ganti, "What is the national best bid and offer (nbbo)? how quote works," Investopedia, 2022, [Online]. Available: https://www.investopedia.com/terms/n/nbbo.asp.

[6] Z. Zhu and R. Huang, "Efficient smc protocol based on multi-bit fully homomorphic encryption," *Journal of Applied Sciences*, vol. 11, no. 21, Nov 2021, [Online Serial]. Available: https://doi.org/10.3390/app112110332. [Accessed March 13, 2024].

[7] D. Easley and M. O'hara, "Information and the cost of capital," *The Journal of Finance*, vol. 59, no. 4, pp. 1553–1583, 2004, [Online Serial]. Available: https://doi.org/10.1111/j.1540-6261.2004.00672.x. [Accessed March 13, 2024].

[8] E. Picardo, C. Potters, and M. Kazel, "An introduction to dark pools," Investopedia, 2022, [Online]. Available: https://www.investopedia.com/articles/markets/050614/introduction-dark-pools.asp.

[9] M. Lewis, *Flash Boys*. New York: Norton, 2015.

[10] J. Cartlidge, N. P. Smart, and Y. T. Alaoui, "Multi-party computation mechanism for anonymous equity block trading: A secure implementation of turquoise plato uncross," *Intelligent Systems in Accounting, Finance and Management*, Nov 2021, [Online serial]. Available: Wiley, https://doi.org/10.1002/isaf.1502. [Accessed March 14, 2024].

[11] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys*, vol. 51, no. 4, pp. 1–35, July 2018, [Online serial]. Available: https://doi.org/10.1145/3214303. [Accessed March 14, 2024].

[12] J. Wahlman, "Pet-exchange: A privacy enhanced trading framework," Master's thesis, Dept. of Computer Science, Linköping University, Linköping, Sweden, 2022, [Online]. Available: https://www.diva-portal.org/smash/get/diva2:1679097/FULLTEXT01.pdf. [Accessed March 13, 2024].

[13] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *STOC '09: Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 169–178, [Online]. Available: https://doi.org/10.1145/1536414.1536440. [Accessed March 14, 2024].

[14] Security and Exchange Commission (SEC), "Cryptocurrency/icos," SEC, 2024, [Online]. Available: https://www.sec.gov/securities-topics/ICO.

[15] I. Inc., "Introduction to the ckks/heaan fhe scheme," Inferati Inc., Washington USA, Nov. 2022, [Online]. Available: https://inferati.azureedge.net/docs/inferati-fhe-ckks.pdf.

[16] O. Baudron and J. Stern, "Non-interactive private auctions," Financial Cryptography pp. 364–377, 2002, [Online]. Available: https://www.di.ens.fr/~stern/data/St92.pdf.

[17] S. Arora, "Lecture 21: Taste of cryptography: Secret sharing and secure multiparty communication," Princeton University, Computer Science Department, Fall 2016, [Lecture notes]. Available: https://www.cs.princeton.edu/courses/archive/fall16/cos521/Lectures/lec21.pdf.

[18] Financial Industry Regulatory Authority (FINRA), "Can you swim in a dark pool?" FINRA, 2023, [Online]. Available: https://www.finra.org/investors/insights/can-you-swim-dark-pool.

[19] J. J. Angel, "When finance meets physics: The impact of the speed of light on financial markets and their regulation," *Financial Review*, vol. 49, pp. 271–281, 2014. [Online]. Available: https://doi.org/10.1111/fire.12035

[20] Y. Ait-Sahalia and M. Saglam, "High frequency market making: The role of speed," Available at SSRN, June 2017, sSRN: https://ssrn.com/abstract=2331613 or http://dx.doi.org/10.2139/ssrn.2331613.

[21] B. Gilbert and D. Rosenthal, "The complete history and strategy of renaissance technologies," Acquired, 2024, acquired, [Podcast]. Available: https://www.acquired.fm/episodes/renaissance-technologies. [Accessed: March 19, 2024].

[22] R. Agrawal, L. de Castro, G. Yang, C. Juvekar, R. Yazicigil, A. Chandrakasan, V. Vaikuntanathan, and A. Joshi, "Fab: An fpga-based accelerator for bootstrappable fully homomorphic encryption," 2022.

[23] Y. Meng, F. D. de Souza, S. Butt, H. de Lassus, Y. Zhou, Y. Wang, T. A. González, J. Jin, and F. Bergamaschi, "Accelerating fully homomorphic encryption with an open source fpga library,"

April 2022, [Accessed: April 30, 2024]. [Online]. Available: https://www.intel.com/content/www/us/en/developer/articles/technical/homomorphic-encryption/accelerating-homomorphic-encryption-for-fpga.html

[24] M. B. van Egmond, V. Dunning, S. van den Berg, T. Rooijakkers, A. Sangers, T. Poppe, and J. Veldsink, "Privacy-preserving anti-money laundering using secure multi-party computation," Cryptology ePrint Archive, Paper 2024/065, 2024, https://eprint.iacr.org/2024/065. [Online]. Available: https://eprint.iacr.org/2024/065