Plekra: A Post-Quantum Peer-to-Peer Electronic Cash System with Per-Device Lottery Issuance

Plekra
plekra.org

Abstract

A fair electronic cash system should allow anyone to take part in the distribution of new coins without first holding coins or expending large amounts of computation. Proof-of-work hands distribution to computing power and proof-of-stake hands it to capital; both shut participants out. We propose a scheme that treats distribution and security separately. Every 10 minutes new coins are awarded to one participant by a random draw, and half of all transaction fees are divided equally among the participants who did not win. The credential for participation is paid for with a physical security chip inside a computer or with continuously occupied disk space; each of the two kinds of credential receives half of the distribution, so the compromise of either affects at most half. The ledger is ordered on a directed acyclic graph by a stake-weighted Byzantine fault-tolerant committee, and the stake of misbehaving members is slashed, which makes the cost of an attack measurable. Draw numbers come from a sequential hash computation over on-chain state fixed in advance, so that anyone able to influence the input cannot compute the result in time. Accounts have no single master private key: everyday operations are performed by device keys within limits, and recovery relies on the user's password together with threshold verification by randomly selected guardian nodes. All signatures use post-quantum algorithms. The chain natively provides spot trading by batch auction, fully covered options, and perpetual contracts opened according to market depth, with liquidation independent of any external price.

1. Introduction

Bitcoin [1] showed that mutually distrustful participants can maintain a single ledger without a central authority. It assigns the right to record transactions, and the ownership of new coins, according to computing power: the more computing power, the larger the share. With the arrival of specialized mining hardware, the right to record and the new coins have in practice become concentrated among a few parties with cheap electricity and specialized hardware, and ordinary computers have long been unable to take part. Proof-of-stake eliminates the energy consumption but usually distributes new coins in proportion to holdings, so latecomers must first buy coins before they can participate.

Beyond distribution, existing systems have three structural problems. First, full control of an account is bound to a single private key: losing the key means losing the assets permanently, and leaking it means losing control of them completely. Second, the discrete logarithm problem on which ECDSA and EdDSA rely can be solved by Shor's algorithm on a fault-tolerant quantum computer; NIST published post-quantum signature and key-encapsulation standards in 2024 [2], and specifies that quantum-vulnerable public-key algorithms be deprecated after 2030 and disallowed after 2035 [3]. Third, derivatives trading depends on centralized exchanges holding custody of assets, while on-chain trading is exposed to front-running and sandwich attacks [4].

The system proposed in this paper responds to these problems as follows:

- **Separating distribution from security.** New coins are distributed at random according to participation credentials, independent of holdings, computing power and time online; ledger security is borne by stake,

and misbehaving stakers are slashed.

- **Physical resources as the cost of participation.** In an open network without an identity authority, if creating an identity costs nothing, one person can pose as arbitrarily many [5]. We use physical security chips and disk space as that cost; the weaknesses of the two complement each other.
- **Rule-based accounts.** An account is defined by a set of authorization rules rather than by a single private key, and the compromise of any single key is not enough to transfer all of its assets.
- **Native finance.** Spot, options and perpetual contracts are matched and settled on chain, completing atomically within the same batch as the asset ledger.

2. Block Net

Transactions are ordered by a committee. In each round every committee member proposes a block, which references blocks of the previous round carrying more than 2/3 of the cumulative voting power; the reference itself constitutes a vote, and no additional certification messages are required. The resulting structure is a directed acyclic graph (Figure 1), which we call the block net. The commit rule is deterministic: all honest members derive the same total order of transactions from the same graph. The construction is taken from Mysticeti [6], which has achieved sub-second commit in a production network of more than a hundred validators.

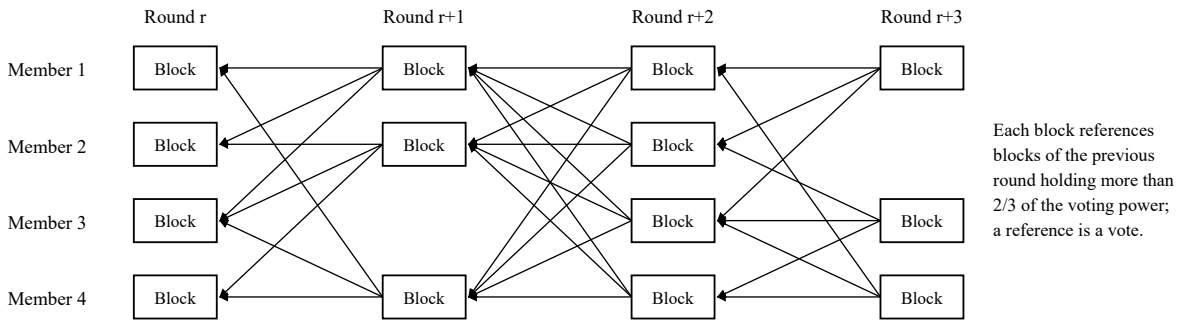


Figure 1. The block net. Members produce blocks in parallel in each round, without queuing behind a single block producer.

Blocks and votes are signed with ML-DSA-44 [2]. Every decision requires signatures from more than 2/3 of the committee's total voting power. Under a partially synchronous network [7] and the assumption that the adversary holds less than 1/3 of the voting power, the protocol satisfies both safety and liveness. Once the block containing a transaction is committed, the transaction's position in the order can no longer change; in steady state the confirmation target for ordinary transfers is 0.5 to 1 second.

3. Two-Lane Ledger

Totally ordered transactions are processed by two execution lanes. Lane A handles assets and distribution, including balances, accounts, participation credentials, draws, dividends, staking and governance; lane B handles finance and contracts, including spot trading, options, leverage, perpetual contracts and third-party WebAssembly contracts.

Every 1.5 to 2 seconds forms one batch. For each batch, deterministic execution rules produce a pair of state roots $(\mathrm{StateRoot}_A,\ \mathrm{StateRoot}_B)$, which are confirmed together with subsequent blocks. Operations in lane B that move funds, namely trades, option exercises and forced liquidations, belong to the same atomic transaction as the corresponding transfers in lane A; if any step fails, both lanes roll back together. The two lanes share the

same consensus backbone, so cross-lane operations need no cross-chain bridge and their security rests on no additional assumptions (Figure 2).

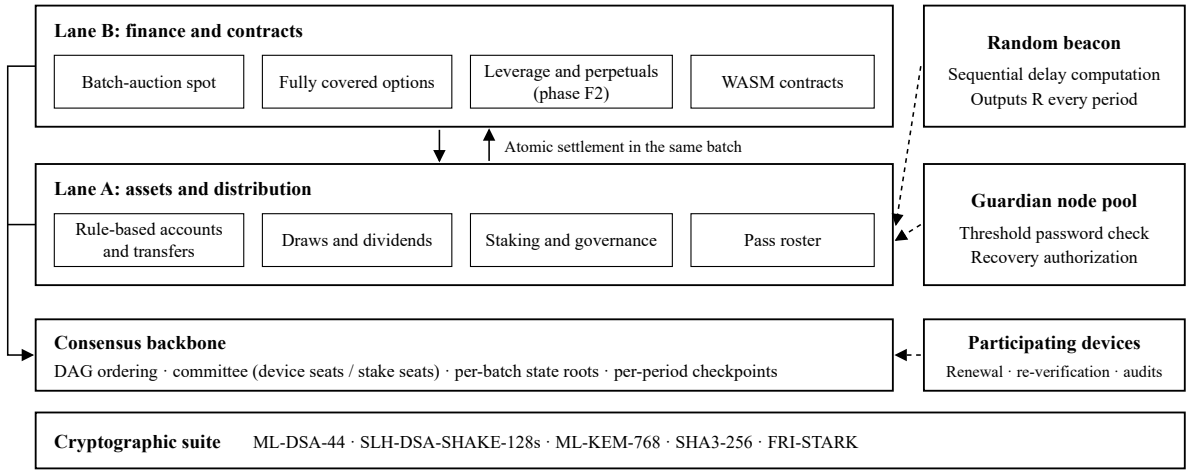


Figure 2. The two-lane ledger. Solid arrows denote state dependencies; dashed arrows denote inputs submitted to the ledger.

The execution capacity of each batch is divided into three classes, namely system operations, native finance and third-party contracts, each with its own fee market. Congestion caused by third-party contracts raises fees only in their own class and does not affect the execution of transfers or forced liquidations.

4. Participation Credentials

The credential for taking part in draws and dividends is called a pass, and passes come in two tracks.

Chip track. The Trusted Platform Module (TPM 2.0) in a computer carries an endorsement key certificate written by the manufacturer at the factory. The device generates a secret ς locally and derives

$$c = H(\varsigma \,\|\, 0), \qquad \mathrm{tag} = H(\varsigma \,\|\, 1)$$

The device registers its chip with the certificate, an attestation key generated inside the chip, and the commitment c; randomly assigned committee members challenge the chip and witness the registration. The chip's deduplication identifier is the hash of the endorsement public key, so each chip can be registered only once. To claim a pass, the device submits tag together with a zero-knowledge proof showing that it knows the ς of some registered chip and holds that chip's signature on the latest random value, without revealing which chip it is. The chain rejects duplicate values of tag. Chip registrations are public; the correspondence between chips and passes is not. A pass must present a chip signature again every 30 days.

The chip track accepts only physical chips. Virtual TPMs that cloud providers offer to virtual machines are not accepted, because their certificates can be issued on demand and a powered-off virtual machine incurs almost no cost. To limit the impact of a leaked manufacturer key, the effective share under each issuing certificate is capped at θ.

Disk track. Using the hash of the pass public key, PassID, as a seed, the holder generates 256 GiB of data following the proof-of-space constructions [8][9]. Given a challenge, the holder can answer within seconds simply by looking it up in the data; anyone who has not stored the data must redo most of the generation

computation. To renew, the holder submits a signed request, and the batch b in which the request is included determines the challenges

$$ch_j = H\big(\mathrm{StateRoot}_A(\mathsf{b}+1) \,\|\, \mathrm{PassID} \,\|\, j\big), \qquad j = 1, \dots, 4$$

The responses must be included in a block within about 30 seconds. The disk track relies only on hash functions; it depends on no manufacturer and is unaffected by quantum computing.

Complementarity of the two tracks. The chip track gives an ordinary computer the same share per unit as a large-scale operator, since each additional share requires an additional physical machine, but it depends on the chip manufacturers' factory signatures. The disk track depends on no commercial party, but storage has economies of scale that favor large operators. Draws alternate between the two tracks and the dividend pool is split equally between them; this ratio cannot be changed by governance. A leaked manufacturer key affects only the chip track and a concentration of storage affects only the disk track, so a failure on either side reaches at most half of the total distribution. Because disk passes are easily accumulated by a single operator, they do not enter the committee and do not vote in governance.

5. Random Beacon

A draw requires that no one able to influence its input can foresee the result. Let b_t be the last batch whose consensus time precedes the t-th draw by 8 periods. The seed and the draw number are

$$\mathrm{Seed}_t = H\big(t \,\|\, \mathrm{StateRoot}_A(\mathsf{b}_t) \,\|\, \mathrm{StateRoot}_B(\mathsf{b}_t)\big), \qquad R_t = H^{(n)}(\mathrm{Seed}_t)$$

where H is SHA3-256 and n is chosen so that the sequential computation takes about 50 minutes on the reference platform. The roster of participants is fixed at the same time as the seed (Figure 3). The last committee member able to influence the seed has only a few seconds to decide; to choose block contents in a directed way, it would have to complete one evaluation within that time, which requires a sequential speedup of more than 300 times. Hash chains have no known method of parallel speedup, and we assume that the adversary's sequential speedup does not exceed 100 times. Verifiable delay functions based on groups of unknown order [10] have succinct proofs but rely on assumptions that do not hold against quantum computers, and are therefore not used; the correctness of the hash chain is verified by committee members recomputing it in segments.

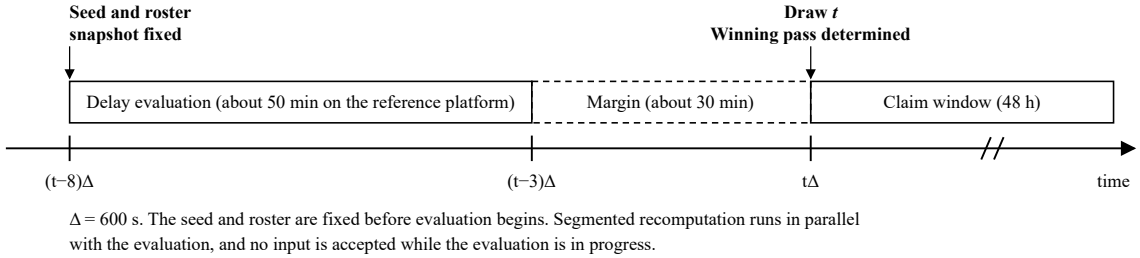


Figure 3. Timeline of the random beacon (the claim window is not drawn to scale).

6. Draws and Dividends

Even-numbered draws are made from the chip track and odd-numbered draws from the disk track. The roster is stored sorted by PassID, and $u_1 = H(R_t \,\|\, 0)$ and $u_2 = H(R_t \,\|\, 1)$ are derived from R_t. In the disk track the winner is the pass at position $\lfloor u_1 N / 2^{256} \rfloor$ in the roster. In the chip track, u_1 first selects an issuing certificate g,

weighted by its effective count $\tilde{n}_g$, and u_2 then selects one of the n_g passes under it uniformly at random. The probability that pass i wins is therefore

$$\Pr[\,i \text{ wins}\,] = \frac{\tilde{n}_g}{\tilde{N}} \cdot \frac{1}{n_g}$$

When no issuing certificate has reached the share cap, this equals $1/N$. The whole procedure requires only one weighted selection and one lookup by index; there is no need to traverse the roster.

The winner must claim the reward within 48 hours using the pass key. Half of the reward is credited immediately, and the other half unlocks daily over the following 30 days. The winner does not take part in dividends during a cooldown of up to 90 days.

Of all fees, 50% enter the dividend pool, split equally between the two tracks. Each epoch (24 hours) releases 1/7 of the pool balance, divided equally among the passes of that track that are not in cooldown:

$$f_e = \frac{P_e / 7}{\sum_{i \in \mathcal{A}_e} w_i}$$

where P_e is the balance of that track's dividend pool, $\mathcal{A}_e$ is the set of passes on the roster and not in cooldown, and w_i is a weight determined by the share cap of the issuing certificate, equal to 1 when the cap has not been reached. Dividends are independent of time online and are credited automatically, without any claim transaction.

7. Incentives

Issuance. The total supply of PLK is capped at 10^9 coins. Supply is zero at genesis, and there is no pre-allocation of any kind. Issuance proceeds in tiers:

$$E_j = 10^8 \times 0.9^{\,j-1}, \qquad \rho_j = 2000 \times 0.9^{\,j-1}, \qquad \sum_{j=1}^{\infty} E_j = 10^9$$

where E_j is the allotment of tier j and ρ_j is the full reward per draw. Each tier contains 50,000 draws, about 347 days; cumulative issuance at the end of a tier is 10% for tier 1, 52% for tier 7 and 90% for tier 22. When a track has few participants, the reward per draw is scaled down to $\rho_j \cdot \min(1,\ \tilde{N}/10^5)$, so that a small number of early participants cannot obtain a disproportionate share in a short time.

Fees. Fees are not burned: 50% go to the dividend pool, 20% to the committee, 5% to guardian nodes, 20% to a protocol reserve whose principal cannot be withdrawn, and 5% to a public treasury spent by vote. Even a party that controls all passes and all stake can recover at most 75% of the fees it pays through self-trading, so self-trading is unprofitable.

Security subsidy. In the early period, fees are insufficient to attract enough stake. The fraction

$$s_e = \min\!\Big(15\%,\ \frac{\max(0,\ B_e - \tfrac{1}{2}F_{e-1})}{I_{e-1}}\Big), \qquad B_e = \frac{5\%}{365} \cdot \frac{2}{3} \cdot Q_e$$

of each draw's reward is credited to stakers, where Q_e is the amount issued so far, and F_{e-1} and I_{e-1} are, respectively, the fees received by stakers and the total draw rewards in the previous epoch. The subsidy tops up stakers' income to the level of "5% a year with two-thirds of the coins staked"; fees count against it at one half, and the subsidy falls to zero once fees reach twice the target. With 10^8 coins issued, the subsidy is about 3% of that year's issuance.

8. Consensus Security

Any address may stake PLK to become a candidate node, and other holders may delegate stake to it. The 128 nodes with the highest voting stake form the committee, and voting power is proportional to voting stake. A node that double-signs within the same epoch is slashed by the fraction

$$\max\left(5\%,\ \min(100\%,\ 3\varphi)\right)$$

where φ is the share of voting power that double-signed. When $\varphi < 1/10$ the event is treated as an isolated fault and only the node's own stake is slashed; causing a fork requires $\varphi \geq 1/3$, in which case own and delegated stake are both slashed in full. Anyone can prove a double signature by presenting the two signatures, so the cost of breaking safety is no less than about one third of the committee's total voting stake.

At genesis there are no coins to stake, so the network passes through three phases. In the bootstrap phase, the committee consists of up to 256 seats held by device nodes of the chip track. Once 10^7 coins have been issued and stake has reached 25% of the issued amount, the network enters the hybrid phase, in which every decision requires signatures from more than 2/3 of the device seats and, separately, from more than 2/3 of the stake seats; as long as either group is honest, the ledger does not fork. Once 10^8 coins have been issued and voting stake has remained no lower than 2/3 of the issued amount for 180 consecutive days, the network enters the staking steady state after a dual vote. If stake falls back substantially, the network automatically returns to the hybrid phase.

9. Accounts

An account is defined by a set of authorization rules that recognize three mutually independent classes of authorization factor: devices (limited session keys held inside the chip), a password (verified by a guardian cluster), and a cold key (an SLH-DSA key kept offline). Transfers within the limit require only a device; outgoing transfers above the limit and changes to the rules require factors of any two classes and pass through a waiting period of 24 to 72 hours. An operation in its waiting period can be revoked by a party presenting factors of at least as many classes. As long as the owner holds one more class of factor than the adversary, the owner can revoke every operation of the adversary; the compromise of a single device on its own causes a loss bounded by the limit.

A user who has lost all devices applies to the guardian cluster for recovery with a recovery name and the password. The cluster consists of 41 guardian nodes selected by the random beacon, and it is redrawn in its entirety every 90 days. The client shares the password verification key v among the cluster members using Shamir's scheme [11] with threshold 21; the shares of any 20 members reveal no information about v. For verification, the cluster uses verifiable multiparty computation with an honest majority [12][13] to make public the value $r(v - v')$, where r is a random number known to no member: the result is zero if and only if the password is correct, and is otherwise a random value that carries no information. Each attempt requires a deposit; after a failure the deposit goes to the account owner and the required deposit doubles, so offline exhaustive search is infeasible. Once verification succeeds, at least 21 members sign an authorization; the chain checks the attempt record, the commitment to the recovery record and the cluster membership, and the new device takes effect after a 72-hour waiting period.

10. Native Finance

Spot. Spot trading uses frequent batch auctions [14]: the orders within each batch are executed together at a single clearing price, with no ordering among them, so front-running has nothing to target. Orders are submitted as commitments before clearing and revealed afterwards.

Options. The seller must lock the full delivery assets before selling. The buyer may exercise at any time before expiry, and delivery completes atomically within the same batch, so default cannot occur.

Leverage and perpetual contracts. Leverage is capped at 20x for perpetual contracts and 3x for spot, and is opened pair by pair according to order-book depth. The total open interest that can be forcibly liquidated within a price move of δ_L may not exceed one quarter of the smaller of the robust depth and the current depth over the same range; the robust depth is the 10th percentile of resting order volume over the past 30 days. The liquidation price is the 5-minute time-weighted average of the on-chain spot clearing price, with its movement limited in each batch. Perpetual contracts and lending each have a segregated insurance vault; losses on perpetual positions that exceed their margin are absorbed by the insurance vault and by auto-deleveraging, without drawing on the protocol reserve. All liquidations use only on-chain prices and do not depend on external oracles.

11. Calculations

Device committee. The committee in the bootstrap phase has 256 seats. Let p be the probability that a given seat is controlled by the adversary; a failure requires the adversary to hold at least 86 seats:

$$\Pr(X \geq 86) = \sum_{x=86}^{256} \binom{256}{x} p^x (1-p)^{256-x}$$

p	10%	15%	20%	25%
Single epoch	1.0×10^{-24}	1.0×10^{-13}	2.4×10^{-7}	1.3×10^{-3}
One year (365 epochs)	3.8×10^{-22}	3.8×10^{-11}	8.9×10^{-5}	37%

The result is highly sensitive to p. The bootstrap phase is therefore the weakest stage of the entire lifecycle, and the network must move to a stake committee as early as possible.

Partial storage. A holder who stores only a fraction s of the disk data passes the 4 challenges of each renewal with probability s^4. A renewal remains valid across two windows, so the fraction of time the pass stays on the roster is about $1-(1-s^4)^2$:

s	50%	80%	90%	95%
Fraction of time on the roster	12%	65%	88%	97%
Return per unit cost relative to full storage	0.24	0.81	0.98	1.02

Storing less data does not pay; near full storage the advantage is no more than about 2%.

Draw returns. For a track in tier 1 with N active passes, the expected return of a single pass is:

N	Reward per draw	Expected daily return	Mean interval between wins	Wins at least once in a year
10^4	200	1.44	139 days	92.8%
10^5	2,000	1.44	3.8 years	23.1%
10^6	2,000	0.144	38 years	2.6%

As the network grows, the variance of draw returns increases rapidly; dividends give those who do not win a continuing, low-variance return.

Guardian cluster. Let q be the adversary's share of the guardian node pool. The probability that a single cluster contains at least 21 adversarial members is 3.6×10^{-11} at $q = 10\%$ and 8.4×10^{-6} at $q = 20\%$. Even if a

cluster is compromised, a recovery initiated by the adversary must still pass the 72-hour waiting period, and can be revoked by the owner with the remaining factors.

12. Conclusion

We have proposed an electronic cash system that relies on neither mining nor prior holdings. New coins are distributed to participants by random draw, and the cost of participation is a physical chip or disk space; each of the two tracks receives half, so the failure of either kind of cost cannot take away the whole distribution. Ledger security is provided by slashable stake and is established through a three-phase transition from a device committee, through a hybrid committee, to a stake committee. Draw numbers come from a sequential hash computation over inputs fixed in advance, accounts have no single master private key, all signatures use post-quantum algorithms, and financial trades and asset settlement complete atomically within the same batch.

The design has its limits. The unit of distribution is the device rather than the natural person, and whoever controls more devices obtains a larger share; the chip track depends on the factory signatures of chip manufacturers; the disk track may concentrate toward large storage operators, which may in turn affect the distribution of stake; security in the bootstrap phase is weaker than in the staking steady state; and a leveraged market that does not rely on external prices may suffer losses beyond margin under violent market moves. These trade-offs and their full analysis, together with all parameters, are set out in the Plekra Technical Specification.

References

1. S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008.
2. NIST, *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*; *FIPS 204: Module-Lattice-Based Digital Signature Standard*; *FIPS 205: Stateless Hash-Based Digital Signature Standard*, Aug. 2024.
3. NIST, "Transition to Post-Quantum Cryptography Standards," *NIST IR 8547*, Oct. 2026.
4. P. Daian et al., "Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability," in *Proc. IEEE S&P*, 2020.
5. J. R. Douceur, "The Sybil Attack," in *Proc. IPTPS*, 2002.
6. K. Babel et al., "Mysticeti: Reaching the Latency Limits with Uncertified DAGs," in *Proc. NDSS*, 2025.
7. C. Dwork, N. Lynch, and L. Stockmeyer, "Consensus in the Presence of Partial Synchrony," *Journal of the ACM*, vol. 35, no. 2, 1988.
8. S. Dziembowski, S. Faust, V. Kolmogorov, and K. Pietrzak, "Proofs of Space," in *Proc. CRYPTO*, 2015.
9. H. Abusalah, J. Alwen, B. Cohen, D. Khilko, K. Pietrzak, and L. Reyzin, "Beyond Hellman's Time-Memory Trade-Offs with Applications to Proofs of Space," in *Proc. ASIACRYPT*, 2017.
10. D. Boneh, J. Bonneau, B. Bünz, and B. Fisch, "Verifiable Delay Functions," in *Proc. CRYPTO*, 2018.
11. A. Shamir, "How to Share a Secret," *Communications of the ACM*, vol. 22, no. 11, 1979.
12. T. Rabin and M. Ben-Or, "Verifiable Secret Sharing and Multiparty Protocols with Honest Majority," in *Proc. STOC*, 1989.
13. E. Ben-Sasson, S. Fehr, and R. Ostrovsky, "Near-Linear Unconditionally-Secure Multiparty Computation with a Dishonest Minority," in *Proc. CRYPTO*, 2012.
14. E. Budish, P. Cramton, and J. Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, vol. 130, no. 4, 2015.