Efficient Contextual Ontological Model of *n*-Qubit Stabilizer Quantum Mechanics

Christoffer Hindlycke[*,‡] and Jan-Åke Larsson[†,‡]
Department of Electrical Engineering, Linköping University, 581 83 Linköping, Sweden

(Received 11 February 2022; revised 23 August 2022; accepted 24 August 2022; published 19 September 2022)

The most well-known tool for studying contextuality in quantum computation is the n-qubit Stabilizer state tableau representation. We provide an extension that not only describes the quantum state but is also outcome deterministic. The extension enables a value assignment to exponentially many Pauli observables, yet it remains quadratic in both memory and computational complexity. Furthermore, we show that the mechanisms employed for contextuality and measurement disturbance are wholly separate. The model will be useful for investigating the role of contextuality in n-qubit quantum computation.

DOI: 10.1103/PhysRevLett.129.130401

Contextuality is an important nonclassical property of quantum mechanics (QM) that has been studied since the 1960s [1,2], whereas current progress in the area is connected to quantum information processing. One tool for studying this question is the stabilizer formalism [3]: in particular, the stabilizer state tableau representation (SSTR) [4], which captures the contextual behavior of the stabilizer subtheory of quantum theory. This is widely used, both in quantum error correction and as a starting point to study properties of the quantum advantage. A typical question is what needs to be added to stabilizer quantum theory to achieve the quantum advantage?

However, SSTR is not an ontological model but rather a representation of the quantum states in the stabilizer subtheory: quadratic in memory and computational complexity. An interesting question is if an ontological model (more specifically, an outcome-deterministic model) can be found that is also computationally efficient. This could then be used to study properties of the quantum advantage as compared to ontological models, rather than as compared to stabilizer quantum mechanics.

The presently known outcome-deterministic models are all either noncontextual or exponential in complexity. Perhaps the most well known is Spekkens' toy theory (STT) [5] from 2007 that models qubits as existing in one of four discrete ontic states, also linking predicted measurement outcomes of Y to those of X and Z. Although noncontextual, STT can still reproduce a number of quantum phenomena. This served as the stepping stone for the eight-state (cube) model [6,7], wherein an additional degree of freedom is introduced for each qubit, "decoupling" Y from X and Z. Another extension is quantum simulation logic (QSL) [8,9]; see below. In 2019, Lillystone and Emerson [10] proposed a contextual ψ-epistemic model of the stabilizer subtheory, which is outcome deterministic but exponential in memory complexity, owing to assigning an explicit phase value to each Pauli operator. An alternate model was also proposed that was quadratic in memory, but that model is no longer outcome deterministic. In this Letter, we draw upon these previous efforts in pursuit of our goal: an efficient, in terms of both computational and memory complexity, contextual outcome-deterministic model of the stabilizer subtheory.

We assume the reader is familiar with basics of linear algebra, the stabilizer formalism, and quantum computation [11]. The standard Pauli operators act on single qubits, on coordinate form

$$\mathbb{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$
$$Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{1}$$

The n-qubit Pauli group $\mathcal{P}_n$ consists of n-qubit Pauli operators and their respective global phase ± 1 or $\pm i$. Because $iXZ = Y$, any element of $\mathcal{P}_n$ can be written as $P = i^p \bigotimes_k i^{x_k z_k} X^{x_k} Z^{z_k}$, where (x, z) is a binary symplectic vector, which is so named because the two elements P and P' commute iff the symplectic product

$$P \cdot P' = \sum_k x_k z'_k - x'_k z_k \tag{2}$$

equals 0 mod 2. The noncommutative group operation $P + P' = P''$ gives, with $x + x' = x''$ and $z + z' = z''$,

$$\begin{aligned} P'' &= i^{p+p'} \bigotimes_k i^{x_k z_k + x'_k z'_k} X^{x_k} Z^{z_k} X^{x'_k} Z^{z'_k} \\ &= i^{p+p'-P\cdot P'} \bigotimes_k i^{x''_k z''_k} X^{x''_k} Z^{z''_k}. \end{aligned} \tag{3}$$

This makes $\mathcal{P}_n$ modulo phase a symplectic vector space for which a symplectic basis $\{M_k; C_k\}_{k=1}^n$ obeys $M_j \cdot M_k = C_j \cdot C_k = 0 \bmod 2$ and $M_j \cdot C_k = \delta_{jk} \bmod 2$. Expansion of $M \in \mathcal{P}_n$ in this basis uses $m_k = M \cdot C_k \bmod 2$, $c_k = M \cdot M_k \bmod 2$, and binary phases v and w:

$$M = (-1)^v i^w \left(\sum_k m_k M_k + \sum_k c_k C_k\right). \qquad (4)$$

An n-qubit stabilizer state $|\psi\rangle$ is uniquely determined by the subgroup $S(|\psi\rangle) \subset \mathcal{P}_n$ that stabilizes $|\psi\rangle$. Equivalently, a stabilizer state can be obtained from $|0\rangle^{\otimes n}$ using only Clifford-group gates (generated by Hadamard, Phase or "S," and CNOT), possibly also including Pauli-group measurements. The elements of a stabilizer subgroup are Hermitian so they can be written as $P = (-1)^v \otimes_k i^{x_k z_k} X^{x_k} Z^{z_k}$. Two such elements commute so that they give $P \cdot P' = 0 \bmod 2$ and

$$P + P' = P'' = (-1)^{v+v'-P\cdot P'/2} \otimes_k i^{x_k'' z_k''} X^{x_k''} Z^{z_k''}. \qquad (5)$$

Aim for the model.—The overall goal here is naturally to construct a model that reaches the known lower memory bound [12], which is a number of classical bits quadratic in the number of qubits, while being relatively simple to understand. We will take inspiration from STT and use elements of the representation of QSL. The latter is an efficient (linear complexity, i.e., constant overhead) classical simulation framework for quantum computation that implements one single additional resource available in quantum systems as compared to classical-bit computation: that of an additional degree of freedom of each elementary system. This allows for construction of quantumlike oracles, and QSL captures enough of the quantum behavior to run, for example, Simon's algorithm and the Deutsch-Jozsa algorithm within the oracle paradigm [9].

QSL (and STT) achieves this by keeping track of two classical bits for each qubit in the model. The two bits are associated with the computational degree of freedom z and the phase degree of freedom x: in effect, modeling a qubit using only four discrete states. Measuring X or Z returns the corresponding bit, whereas measuring Y returns the XOR (eXclusive-OR) of the x and z bits; and this makes the output deterministic given the internal state of the model. Randomization occurs as dictated by quantum mechanics: Measuring X randomizes the z bit to zero or one uniformly, and vice versa. Measuring Y randomizes the x and z bits in such a way that their XOR is unchanged ($= y$). Measurement outcomes are repeatable, and we obtain measurement disturbance as it occurs in quantum mechanics. Gates in QSL act on these bit values; for the Clifford-group gates,

$$H(z_h; x_h) = (x_h; z_h), \qquad S(z_s; x_s) = (z_s \oplus 1; x_s \oplus z_s),$$
$$\text{CNOT}(z_c, z_t; x_c, x_t) = (z_c, z_t \oplus z_c; x_c \oplus x_t, x_t). \qquad (6)$$

This makes phase kickback manifest in the CNOT gate, and many quantum-mechanical identities are obeyed, e.g., $HXH = Z$ and $HZH = X$. However, some identities fail; e.g., because the value of y is given by the XOR of x and z in QSL, we obtain $HYH = Y$ rather than the quantum-mechanical $HYH = -Y$. One effect of this is that QSL (and STT) are noncontextual. In this Letter, our aim is to add contextuality.

A contextual ontological model.—The main feature of QSL (and STT) is that it contains a value assignment to the symplectic basis $\{Z_k; X_k\}_{k=1}^n$, where Z_k and X_k are one-qubit Pauli operators acting on system k. QSL now gives the outcome of a measurement M by mod 2 summing the bit values of the symplectic basis elements contained in M.

Inspired by this, the new model will still contain a value assignment to a symplectic basis for $\mathcal{P}_n$ but not necessarily the basis used in QSL. We choose $\{M_k\}$ to be a basis for the stabilizer group of the quantum state of the system so that the phase (± 1) of the elements gives the predicted outcome of any Pauli measurement from that subgroup, corresponding to the value assignment. This is not so different from SSTR; but, for reasons that will become clear later, we will call this stabilizer group the measurement context $\mathcal{M}$.

The second half of the symplectic basis is now needed to generate $\mathcal{P}_n$. In SSTR, this is called destabilizer [4] and is used to identify measurements for which the outcome should be random. This is where our ontological model will deviate from SSTR. Similar to QSL, we here choose C_k conjugate to M_k, filling out the symplectic basis, under the name conjugate context $\mathcal{C}$ and use the same value assignment to its elements, associating the phase to a (predicted) outcome of any Pauli measurements from that subgroup. Measurement in the model will use three distinct steps: (A) *Retrieve the measurement outcome v.*—Expand M in the symplectic basis as in Eq. (4), use v as outcome, and ignore w because M is Hermitian. (B) *Store $(-1)^v M$ as a basis element of $\mathcal{M}$.*—Find k so that $M \cdot M_k = c_k \neq 0 \bmod 2$. (i) If successful ($M \notin \mathcal{M}$), update the elements M_j ($j \neq k$) for which $M \cdot M_j = c_j \neq 0 \bmod 2$ to $M_j + M_k$, and replace C_k with M_k. (ii) Otherwise ($M \in \mathcal{M}$), find k so that $m_k \neq 0$. Then, replace M_k with $(-1)^v M$, and update the elements C_j ($j \neq k$) for which $M \cdot C_j = m_j \neq 0 \bmod 2$ to $C_j + C_k$. (C) *Perform measurement disturbance.*—Randomize the phase for the possibly new C_k. Step (A) gives a well-defined deterministic map from bit values in the model to the outcome v. Step (B) ensures that the measurement and conjugate contexts remain a symplectic basis, having updated $M_k = (-1)^v M$. This makes step (C) implement measurement disturbance with minimal complexity because only one fair coin toss is needed, mirroring the measurement disturbance as it occurs in QM.

We turn now to Clifford-group gate implementation, which is straightforward: Apply the gates to all elements of the symplectic basis, including the phase according to QM identities. Here, in contrast to QSL, the Hadamard gate acting on Y will indeed result in $-Y$. Clifford-group gates preserve the commutation relations between Pauli operators, and so the symplectic basis will remain a symplectic basis. In coordinates [4],

$$H(z_h;x_h;r)=(x_h;z_h;r\oplus x_h z_h),$$
$$S(z_s;x_s;r)=(z_s\oplus x_s;x_s;r\oplus x_s z_s),$$
$$\mathrm{CNOT}(z_c,z_t;x_c,x_t;r)=[z_c,z_t\oplus z_c;x_c\oplus x_t,x_t;r\oplus x_c z_t(x_t\oplus z_c\oplus 1)]. \tag{7}$$

The final part of the model is state preparation. First, choose M_k so that they stabilize the initial state and mutually commute. Second, choose mutually commuting C_k, with random phase, that anticommute with the corresponding M_k and commute with M_j, $j\neq k$.

The model construction obeys the knowledge balance principle of STT [5]: "If one has maximal knowledge, then for every system, at every time, the amount of knowledge one possesses about the ontic state of the system at that time must equal the amount of knowledge one lacks." Step (C) of the measurement procedure ensures that this balance is maintained.

State preparation can also be done using Clifford-group gates on $|0\rangle^{\otimes n}$, which is stabilized by $M_k = Z_k$; and one good choice of conjugate context basis with random phases r_k (fair coin tosses) is $C_k = (-1)^{r_k} X_k$. Alternatively, pick a completely random initial state and perform measurement and transformations to create the desired state. This latter method reproduces the standard quantum-mechanical statement "preparation is measurement." ("Any measurement in quantum theory can in fact only refer either to a fixation of the initial state or to the test of such predictions, and it is first the combination of measurements of both kinds which constitutes a well-defined phenomenon" [13].) Any stabilizer state can be prepared using either method.

Theorem 1. The model presented above is an ontological model of the n-qubit stabilizer subtheory.

Proof.—It suffices to show that our model gives the same predictions as SSTR [4]. As already observed, we can use $|0\rangle^{\otimes n}$ [i.e., $\{Z_k;(-1)^{r_k}X_k\}_{k=1}^n$] as the canonical initial state. The only difference to the standard initial tableau of SSTR is that our model uses random r_k, whereas SSTR sets $r_k = 0$ and then never uses these values. The application of gates is identical to SSTR [see Eq. (7)], also implying that basis elements C_k that have independent random phases before a gate array have independent random phases after the gate array.

Therefore, step (A) of the measurement procedure gives the same predictions as SSTR: if $M\in\mathcal{M}$, the outcome v obtained from Eq. (4) equals the total `rowsum` of SSTR because both realize the group operation in $\mathcal{P}_n$; and, if $M\notin\mathcal{M}$, the outcome v will be random because it contains one or more independent fair coin tosses. Step (B) updates the basis $\{M_k;C_k\}_{k=1}^n$. No update is done in SSTR if $M\in\mathcal{M}$, whereas our model changes basis elements but neither $\mathcal{M}$ nor the value assignment for $\mathcal{P}_n$; so, future predictions remain unchanged. If $M\notin\mathcal{M}$, the state update of step (B) is identical to SSTR, with the caveat that SSTR only handles one-qubit Z measurements (see the update rules for case 1 on p. 4 of [4]), but this restriction can be removed. The final step [step (C)] implements measurement disturbance, which is needed in our model to maintain random independent phases for all C_k so that predictions for later measurement outcomes are also exactly the same as for SSTR. ∎

Memory and computational complexity.—Storing the two contexts requires $4n^2+2n$ bits. Keeping track of interim operators and indices during measurement updating requires, at most, $6n+2n\log n+\log n+4$ bits for a maximum concurrent memory cost of $4n^2+8n+2n\log n+\log n+4$ bits. The model is quadratic in memory complexity, reaching the lower bound in Ref. [12] for classical models that simulate quantum contextuality.

Initializing the model and applying k gates require, at most, $4n+2+16kn$ operations. Expanding M according to Eq. (4) requires $6n^2+4n$ operations. Updating the symplectic basis requires $4n^2-n$ operations because we may make use of many of the calculations carried out when expanding M. Finally, randomizing the phase of one operator requires two operations. Thus, for k gates and l measurements, the number of operations required is equal to $4n+2+16kn+l(10n^2+3n+2)$: The model is computationally efficient. Note that, for algorithms that reduce to a decision problem (where we can encode the phase value of $n-1$ qubits into ancilla qubits using consecutive CNOT gates), the model is indeed quadratic in computational complexity, in the same way as SSTR.

Examples of contextual behavior.—From here on, we suppress the tensor notation; i.e., $XXYZ$ should be read as $X\otimes X\otimes Y\otimes Z$. The standard example is the Peres-Mermin (PM) square [2,14–18]:

$$\begin{matrix} ZI & IZ & ZZ \\ IX & XI & XX \\ ZX & XZ & YY \end{matrix} \tag{8}$$

A model that assigns noncontextual values to phases will give an even number of rows and columns that yield measurement outcomes that sum to 1 mod 2, whereas QM predicts an odd number of such rows and columns, namely, the rightmost column only. A value assignment therefore needs to be contextual (depend on measurement context; here, meaning row or column) to give QM behavior.

The PM square is state independent; but, for purposes of demonstration, let us here assume we begin in the state $|00\rangle$ so that state preparation in our model gives the symplectic basis $\{ZI,IZ;XI,-IX\}$ (random phases 0, 1 drawn by the authors). From this starting state, let us look at measurement sequences $ZZ;XX;YY$ and $ZX;XZ;YY$; the first sequence starts with ZZ. (A) We have $M_1+M_2=ZI+IZ=ZZ=M$, and so $v=0$. (B) Case ii. All $c_k=0$ and $m_1=1$; so, update the basis to $\{ZZ,IZ;XI,-IX+XI=-XX\}$. (C) Randomize the phase of C_1: $\{ZZ,IZ;\pm XI,-XX\}$.

Then, measure XX. (A) We have $C_2 = -XX = -M$, and so $v = 1$. (B) Case i. $c_2 = 1$; update to $\{ZZ, -XX; \pm XI, IZ\}$. (C) Randomize the phase of C_2: $\{ZZ, -XX; \pm XI, \pm IZ\}$.

Measurement of YY will find $M_1 + M_2 = ZZ + (-XX) = YY = M$; so, $v = 0$, making the outcomes from the rightmost column of Eq. (8) total $0 \oplus 1 \oplus 0 = 1$ as QM predicts.

Restarting from the initial state $\{ZI, IZ; XI, -IX\}$, the second sequence starts with ZX. (A) We have $M_1 + C_2 = ZI + (-IX) = -ZX = -M$, and so $v = 1$. (B) Case i. $c_2 = 1$; update to $\{ZI, -ZX; XI + IZ = XZ, IZ\}$. (C) Randomize the phase of C_2: $\{ZI, -ZX; XZ, \pm IZ\}$.

Then, measure XZ. (A) We have $C_1 = XZ = M$, and so $v = 0$. (B) Case i. $c_1 = 1$; update to $\{XZ, -ZX; ZI, \pm IZ\}$. (C) Randomize the phase of C_1: $\{XZ, -ZX; \pm ZI, \pm IZ\}$. Here, measurement of YY will find $M_1 + M_2 = XZ + (-ZX) = -YY = -M$; so, $v = 1$, making the outcomes from the bottom row of Eq. (8) total $1 \oplus 0 \oplus 1 = 0$ as QM predicts.

The measurement outcomes of ZZ, XX, ZX, and XZ are as one would expect from the initial state. But, importantly, the measurement outcome of YY depends deterministically on what measurements are performed together with YY: the so-called measurement context. The model stores performed measurements in $\mathcal{M}$, hence the name. The map to the measurement outcome of YY is completely deterministic, given the initial state, but depends on what measurements are performed before YY; so, the model is contextual, which is what enables it to reproduce the QM contextual behavior. Note that although the chosen order of measurements may influence the outcomes, this influence is deterministic; and for commuting measurements, the associated measurement disturbances do not change the outcomes.

Another example is the Greenberger-Horne-Zeilinger (GHZ) paradox that uses an entangled state of three qubits with stabilizer-group generators, e.g., $-XYY$, $-YXY$, and $-YYX$; another stabilizer is $XXX = (-XYY) + (-YXY) + (-YYX)$. These encode the correlations of the GHZ paradox, which are such that an ontological model (in the terminology used in this Letter) can only reproduce these correlations if the measurement outcome at one qubit depends on what measurements are performed on the other qubits [19]. In this situation, such influences are usually called nonlocal. In our model, the GHZ state below uses three random phases $a = (-1)^r$, $b = (-1)^s$, and $c = (-1)^t$; and single system measurements give, e.g.,

$$\begin{aligned} &\{-XYY, -YXY, -YYX; aYII, bIYI, cIIY\} \\ \xrightarrow{M=Y_1}&\{aYII, -YXY, -YYX; \pm XYY, bIYI, cIIY\} \\ \xrightarrow{M=Y_2}&\{aYII, bIYI, -YYX; \pm XYY, \pm YXY, cIIY\} \\ \xrightarrow{M=X_3}&\{aYII, bIYI, -abIIX; \pm cXYI, \pm cYXI, \pm IIY\}. \end{aligned} \tag{9}$$

The binary outcomes sum to $r \oplus s \oplus (1 \oplus r \oplus s) = 1$, and so give the expected anticorrelation. Another choice of measurement sequence gives

$$\begin{aligned} &\{-XYY, -YXY, -YYX; aYII, bIYI, cIIY\} \\ \xrightarrow{M=X_1}&\{-XYY, -bcXII, IZZ; -aIXY, \pm YXY, cIIY\} \\ \xrightarrow{M=X_2}&\{-acIXI, -bcXII, XXX; \pm XYY, \pm YXY, cIIY\} \\ \xrightarrow{M=X_3}&\{-acIXI, -bcXII, abIIX; \pm cXYI, \pm cYXI, \pm IIY\}. \end{aligned} \tag{10}$$

The outcomes sum to $(1 \oplus r \oplus t) \oplus (1 \oplus s \oplus t) \oplus (r \oplus s) = 0$, and so give the expected correlation. The model is nonlocal because the measurement X_3 gives the outcome $1 \oplus r \oplus s$ in the first case but $r \oplus s$ in the second.

Our final example is the quantum shallow circuits algorithm [20], which always succeeds when run by our model: a fact which follows immediately from Theorem 1 because the algorithm only uses (a subset of) the Clifford gates. We demonstrate the behavior for the problem instance

$$f(x) = x^T A x \mod 4, \qquad \text{with } A = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}. \tag{11}$$

The task is to find z so that $f(x) = 2z \cdot x \mod 4$ on the subset of vectors where $Ax = 0 \mod 2$. The algorithm uses the circuit in Fig. 1, and our model gives

$$\begin{aligned} &\{ZII, IZI, IIZ; aXII, bIXI, cIIX\} \\ \xrightarrow{HHH}&\{XII, IXI, IIX; aZII, bIZI, cIIZ\} \\ \xrightarrow{CZ_{12}}&\{XZI, ZXI, IIX; aZII, bIZI, cIIZ\} \\ \xrightarrow{CZ_{13}}&\{XZZ, ZXI, ZIX; aZII, bIZI, cIIZ\} \\ \xrightarrow{ISS}&\{XZZ, ZYI, ZIY; aZII, bIZI, cIIZ\} \\ \xrightarrow{HHH}&\{ZXX, -XYI, -XIY; aXII, bIXI, cIIX\} \\ \xrightarrow{M=Z_1}&\{ZXX, bcZII, IYY; -aIYI, \pm XYI, cIIX\} \\ \xrightarrow{M=Z_2}&\{-abIZI, bcZII, -ZZZ; \pm IYI, \mp aXII, cIIX\} \\ \xrightarrow{M=Z_3}&\{-abIZI, bcZII, acIIZ; \pm IYI, \mp aXII, \pm IIX\}. \end{aligned} \tag{12}$$

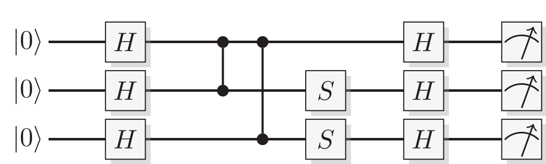

FIG. 1. The quantum shallow circuits algorithm for the problem instance of Eq. (11).

TABLE I. Comparison between models. Note that Spekkens' toy theory is not given in efficient form in Ref. [5] but can be cast in that form.

Model	Efficient	Contextual	Outcome-deterministic
Stabilizer state tableau representation [4]	✓	✓	✗
Spekkens' toy theory [5]	✓	✗	✓
Quantum simulation logic [9]	✓	✗	✓
Lillystone-Emerson [10]	✗	✓	✓
Lillystone-Emerson alternate [10]	✓	✓	✗
This work	✓	✓	✓

Note that gates have a bounded fan-in in our model. The measurement output, both from our model and from QM, is one of the solutions with equal probability:

$$z = \begin{pmatrix} s \oplus t \\ 1 \oplus r \oplus s \\ r \oplus t \end{pmatrix} \in \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}. \quad (13)$$

Conclusion.—We have presented an efficient contextual ontological model of stabilizer quantum mechanics. Previously proposed models all lack at least one of the properties of efficiency, contextuality, and outcome determinism; see Table I for a comparison. In addition, our model is ψ-ontic. Unlike Spekkens' toy theory [5] and quantum simulation logic [9], our model implements contextuality for the stabilizer subtheory, and is thus able to successfully run algorithms relying on that quantum resource, such as the quantum shallow circuits algorithm as shown above. In contrast to the models by Lillystone and Emerson [10], our model combines outcome determinism and efficiency.

Outcome determinism is an important difference to the stabilizer state tableau representation [4], but note that this is more than a mere philosophical issue because it can also be utilized in the analysis of quantum algorithms. The stabilizer state tableau representation efficiently stores the stabilizer group of a single stabilizer state and enables efficient use of Clifford-group gates and Pauli measurements so that we can follow a single quantum state as it is transformed, one gate after another, and subsequently measured. Our model additionally treats the conjugate context on almost the same footing, storing it alongside the measurement context (that stores the stabilizer group of some selected state). There are then several choices of stabilizer group possible in our model using elements from both contexts so that our model enables us to simultaneously follow the behavior of all of these exponentially many quantum states as they are transformed, one gate after another, and subsequently measured.

The model can be implemented and used in practical applications for thousands of qubits on a modern classical computer, for example, using Python [21]. That the model can follow exponentially many quantum states using quadratic classical resources is a direct consequence of the model structure, the many possible stabilizer choices, and outcome determinism. It is our belief that this remarkable property should prove quite helpful in enhancing our understanding of quantum algorithms.

A second property of the model is equally intriguing to us: The mechanism governing contextuality is entirely separated from that ensuring measurement disturbance. They are two distinct steps in the measurement update process, with no interaction between them. The exact ramifications of this are, at least to us, difficult to foresee; but, we strongly believe this provides a very promising venue to explore further.

Finally, because our model successfully reproduces the contextual behavior of the stabilizer subtheory while reaching the theoretical lower memory bound, it severely limits how much of the quantum advantage can arise from stabilizer contextuality alone. At the very least, it suggests that to attribute the quantum advantage to contextuality, one will need to delve further into the structure of contextuality itself, beyond the stabilizer subtheory.

[*]christoffer.hindlycke@liu.se
[†]jan-ake.larsson@liu.se
[‡]These authors contributed equally to this work.

[1] S. Kochen and E. P. Specker, The problem of hidden variables in quantum mechanics, J. Math. Mech. **17**, 59 (1967), https://www.jstor.org/stable/24902153.
[2] Costantino Budroni, Adán Cabello, Otfried Gühne, Matthias Kleinmann, and Jan-Åke Larsson, Kochen-Specker contextuality, arXiv:2102.13036 [Rev. Mod. Phys. (to be published)].
[3] Daniel Gottesman, Theory of fault-tolerant quantum computation, Phys. Rev. A **57**, 127 (1998).
[4] Scott Aaronson and Daniel Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).
[5] Robert W. Spekkens, Evidence for the epistemic view of quantum states: A toy theory, Phys. Rev. A **75**, 032110 (2007).
[6] Joel J. Wallman and Stephen D. Bartlett, Nonnegative subtheories and quasiprobability representations of qubits, Phys. Rev. A **85**, 062121 (2012).
[7] Pawel Blasiak, Quantum cube: A toy model of a qubit, Phys. Lett. A **377**, 847 (2013).
[8] Niklas Johansson and Jan-Åke Larsson, Efficient classical simulation of the Deutsch–Jozsa and Simon's algorithms, Quantum Inf. Process. **16**, 1 (2017).
[9] Niklas Johansson and Jan-Åke Larsson, Quantum simulation logic, oracles, and the quantum advantage, Entropy **21**, 800 (2019).
[10] Piers Lillystone and Joseph Emerson, A contextual ψ-epistemic model of the n-qubit stabilizer formalism, arXiv:1904.04268.

[11] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information, volume 10th Anniversary Edition* (Cambridge University Press, Cambridge, England, 2010).
[12] Angela Karanjai, Joel J. Wallman, and Stephen D. Bartlett, Contextuality bounds the efficiency of classical simulation of quantum processes, arXiv:1802.07744.
[13] N. Bohr, The causality problem in atomic physic, *The Philosophical Writings of Niels Bohr*, Causality and Complementarity Vol. 4, edited by J. Faye and H. J. Folse (Ox Bow Press, Woodbridge, 1999), p. 94.
[14] Asher Peres, Incompatible results of quantum measurements, Phys. Lett. A **151**, 107 (1990).
[15] N. David Mermin, Simple Unified Form for the Major No-Hidden-Variables Theorems, Phys. Rev. Lett. **65**, 3373 (1990).
[16] A. Peres, Two simple proofs of the Kochen-Specker theorem, J. Phys. A **24**, L175 (1991).
[17] Asher Peres, *Quantum Theory: Concepts and Methods*, Fundamental Theories of Physics Vol. 57 (Kluwer, Dordrecht, 1993).
[18] N. David Mermin, Hidden variables and the two theorems of John Bell, Rev. Mod. Phys. **65**, 803 (1993).
[19] Daniel M. Greenberger, Michael A. Horne, Abner Shimony, and Anton Zeilinger, Bell's theorem without inequalities, Am. J. Phys. **58**, 1131 (1990).
[20] Sergey Bravyi, David Gosset, and Robert König, Quantum advantage with shallow circuits, Science **362**, 308 (2018).
[21] Link to a Python implementation of our model, https://gitlab.liu.se/icg/efficient-contextual-ontological-model.