Title: Contextuality and quantum simulation

Date: Jul 27, 2017 10:00 AM

URL: http://pirsa.org/17070050

Abstract:

Contextuality and quantum simulation

Stephen Bartlett
Centre for Engineered Quantum Systems, School of Physics

Joint work with Hakop Pashayan, Angela Karanjai, Joel Wallman

THE UNIVERSITY OF SYDNEY

The power of quantum computation

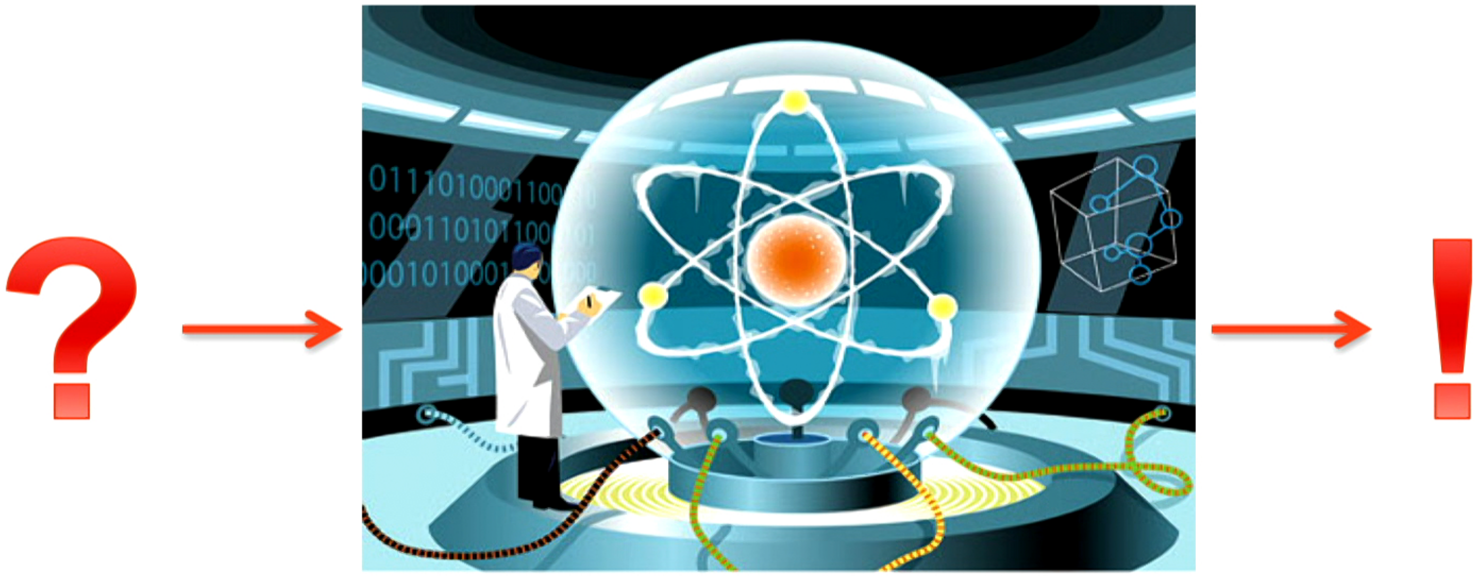

The power of quantum computation

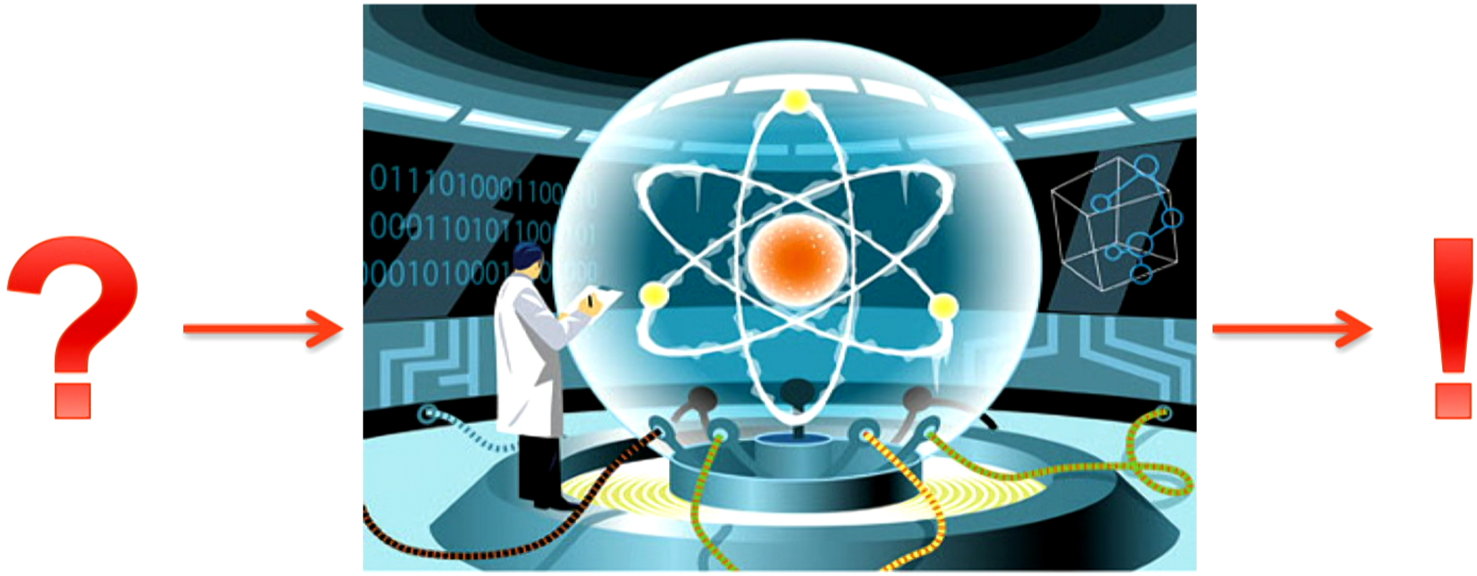

What makes quantum circuits/processes so hard to simulate?

- Exponentially large Hilbert space?
- Entanglement?
- Superposition of many ‘classical’ processes?

The power of quantum computation – a modern approach

- If the quantum process can be modelled efficiently by a classical stochastic process, then it is efficiently simulatable classically – not powerful

- Research program: partition quantum operations into two categories
  1. Those describable by a classical stochastic process – ***free operations***
  2. Those which cannot, which serve as ***resources***

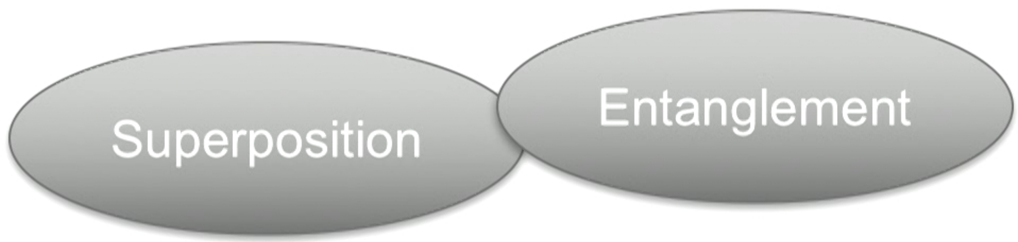

The power of quantum computation – a modern approach

- If the quantum process can be modelled efficiently by a classical stochastic process, then it is efficiently simulatable classically – not powerful

- Research program: partition quantum operations into two categories
  1. Those describable by a classical stochastic process – ***free operations***
  2. Those which cannot, which serve as ***resources***

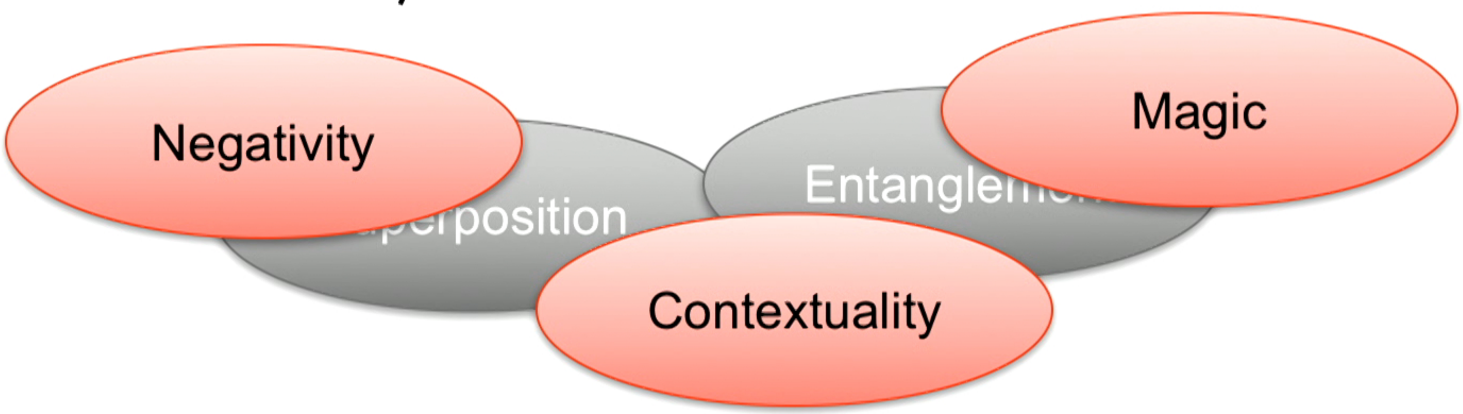

The power of quantum computation – a modern approach

- If the quantum process can be modelled efficiently by a classical stochastic process, then it is efficiently simulatable classically – not powerful

Ontological models

- Research program: partition quantum operations into two categories
  1. Those describable by a classical stochastic process – ***free operations***
  2. Those which cannot, which serve as ***resources***

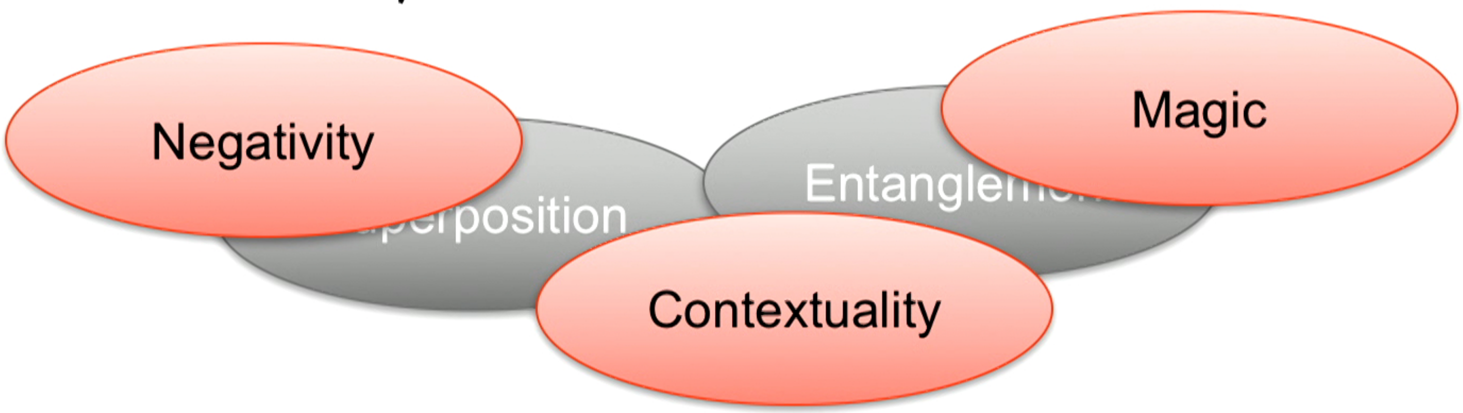

Simulation as a classical stochastic process

- Quantum state is associated with a probability distribution on a classical (phase) space
- Transformations associated with a stochastic map
- Measurements associated with conditional probability dist^{n}s

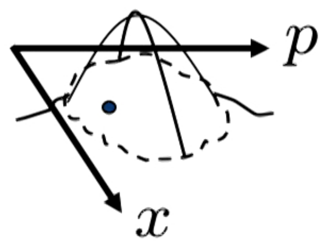

Simulation as a classical stochastic process

- Quantum state is associated with a probability distribution on a classical (phase) space
- Transformations associated with a stochastic map
- Measurements associated with conditional probability dist^{n}s

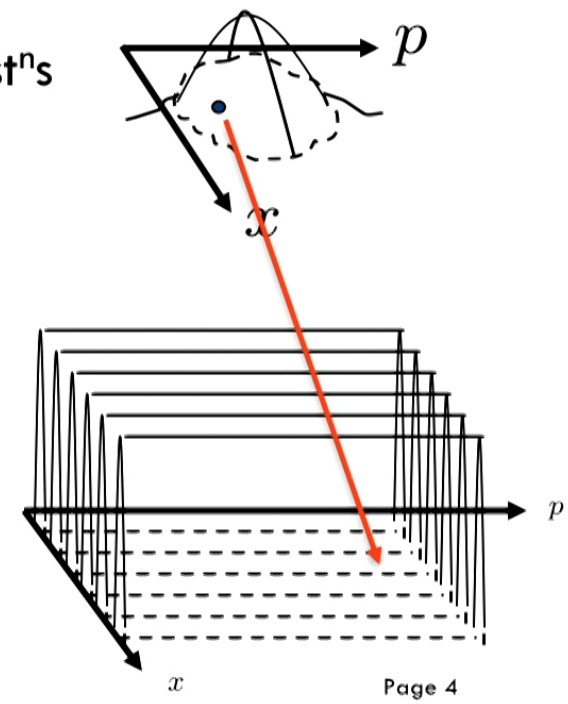

Simulation as a classical stochastic process

- Quantum state is associated with a probability distribution on a classical (phase) space
- Transformations associated with a stochastic map
- Measurements associated with conditional probability dist^{n}s
- Simulation through Monte Carlo sampling
- Corresponds to existence of an ontological model

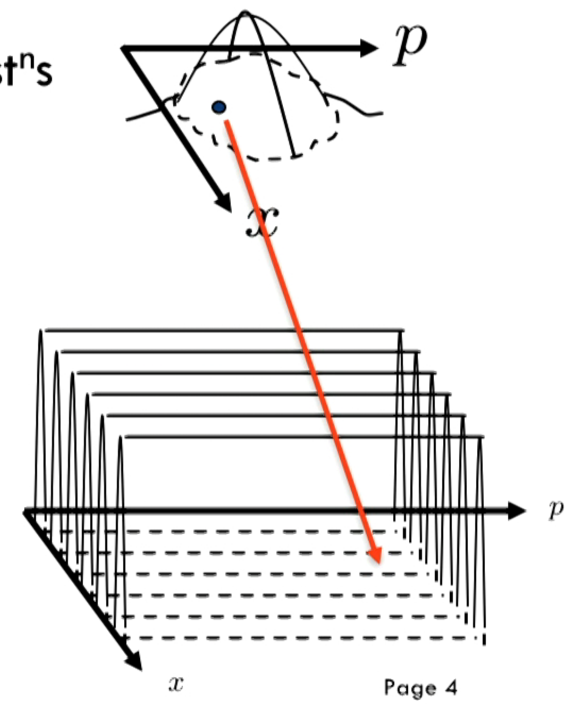

Ontological models for quantum simulation

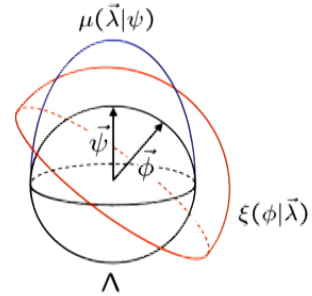

When does this work *exactly* and *efficiently*?

1. A single qubit
2. Gaussian quantum optics

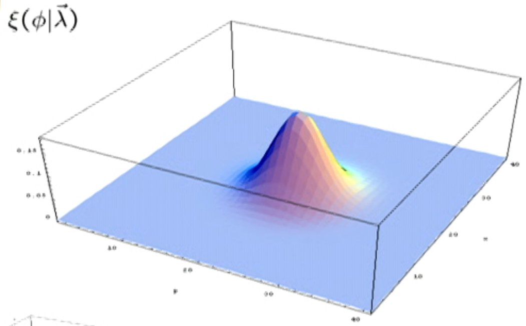

Bartlett, Nemoto, Sanders, Braunstein, PRL (2001)
Mari and Eisert, PRL (2012)
Veitch, Wiebe, Ferrie, Emerson, NJP (2013)
see also Bartlett, Rudolph, Spekkens, PRA (2012)

3. Stabilizer subtheory for odd-dimensional qudits

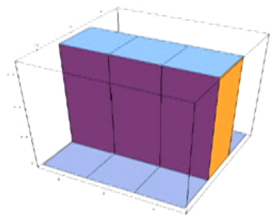

Veitch, Ferrie, Gross, Emerson, NJP (2012)
Mari and Eisert, PRL (2012)

All correspond to *noncontextual* ontological models,
(2) and (3) originating from quasiprobability representations

Statistical noise corresponds exactly to 'quantum noise'

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

States $\rho \longrightarrow W_\rho(\lambda)$ Like a probability distribution

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

States	$\rho \longrightarrow W_\rho(\lambda)$	Like a probability distribution
Unitaries	$U \longrightarrow W_U(\lambda\|\lambda')$	Like a conditional probability

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

States	$\rho \longrightarrow W_\rho(\lambda)$	Like a probability distribution
Unitaries	$U \longrightarrow W_U(\lambda\|\lambda')$	Like a conditional probability
Measurements	$E \longrightarrow W(E\|\lambda)$	Like a conditional probability

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

States	$\rho \longrightarrow W_\rho(\lambda)$	Like a probability distribution
Unitaries	$U \longrightarrow W_U(\lambda\|\lambda')$	Like a conditional probability
Measurements	$E \longrightarrow W(E\|\lambda)$	Like a conditional probability

- Real valued, normalized like probability distributions
- Born rule as you'd expect: $\mathrm{Tr}[EU\rho U^\dagger] = \sum_{\lambda,\lambda' \in \Lambda} W(E|\lambda) W_U(\lambda|\lambda') W_\rho(\lambda')$

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

- Classical hidden variables on a phase space Λ

States	$\rho \longrightarrow W_\rho(\lambda)$	Like a probability distribution
Unitaries	$U \longrightarrow W_U(\lambda\|\lambda')$	Like a conditional probability
Measurements	$E \longrightarrow W(E\|\lambda)$	Like a conditional probability

- Real valued, normalized like probability distributions
- Born rule as you'd expect: $\mathrm{Tr}[EU\rho U^\dagger] = \sum_{\lambda,\lambda'\in\Lambda} W(E|\lambda)W_U(\lambda|\lambda')W_\rho(\lambda')$
- **But can go negative!**

Quasiprobabilities

Quasiprobability representations: another way of describing quantum mech.

– Classi

Stat ion

Unit lity

Mea lity

– Re

– Bo

– **But**

Dual frames formalism

Two 'frames': $F(\lambda) : \lambda \in \Lambda$ and $G(\lambda) : \lambda \in \Lambda$

satisfying $A = \sum_{\lambda \in \Lambda} G(\lambda) \mathrm{Tr}[A F(\lambda)] \quad \forall\ A$

$$W_\rho(\lambda) = \mathrm{Tr}[F(\lambda)\rho]$$
$$W_U(\lambda'|\lambda) = \mathrm{Tr}(F(\lambda')UG(\lambda)U^\dagger)$$
$$W(E|\lambda) = \mathrm{Tr}[EG(\lambda)]\,.$$

Negativity and nonclassicality

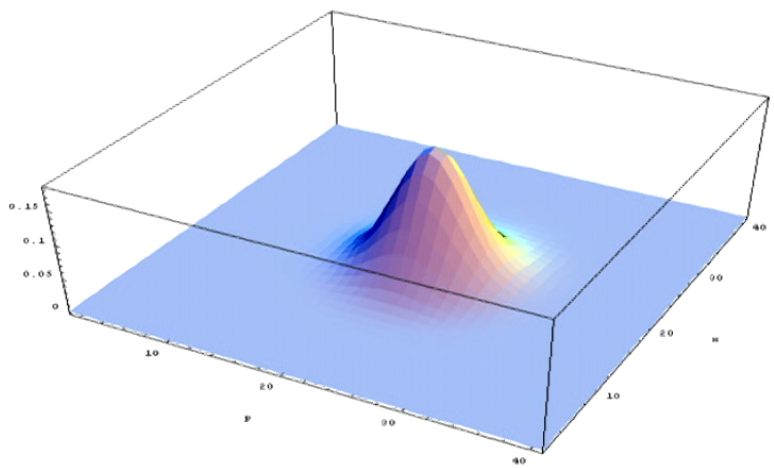

Classical

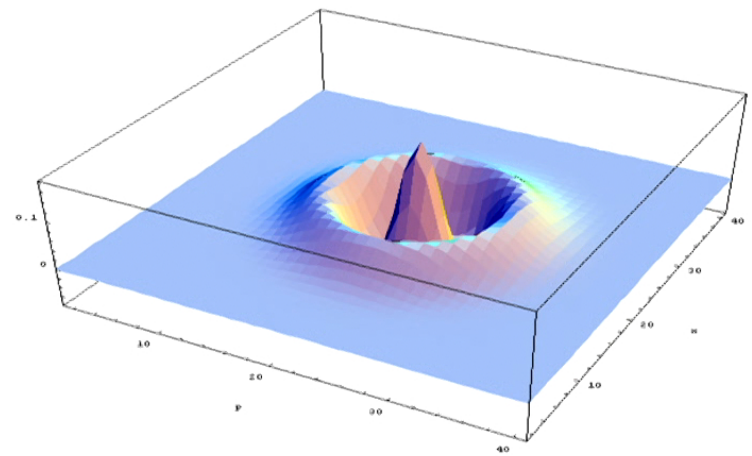

Quantum

Quasiprobabilities for finite quantum systems

Finite-dimensional quantum systems typically use a discrete phase space

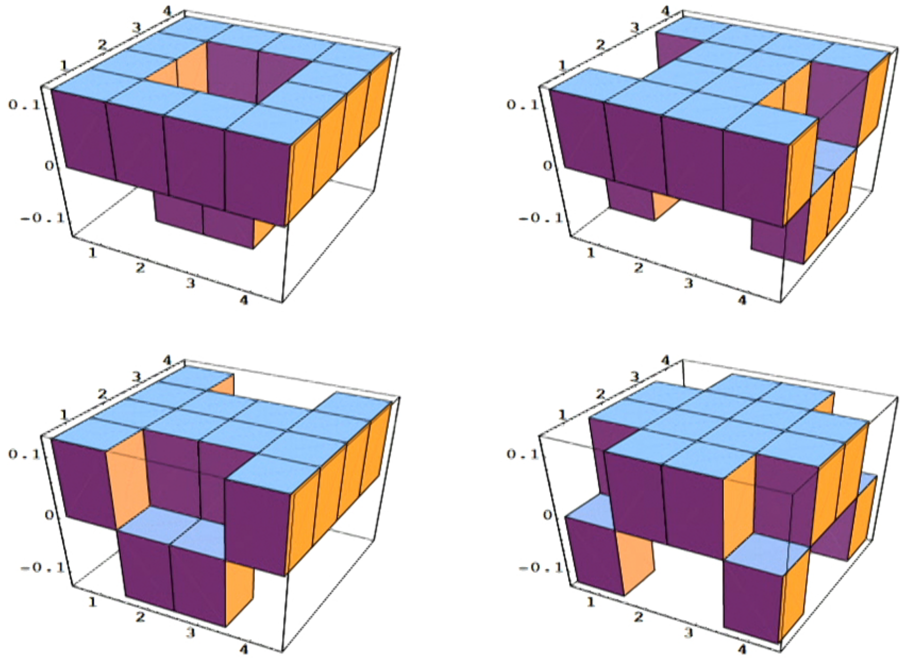

Gibbons, Hoffman, Wootters, PRA (2004); Gross, JMP (2006)

Negativity and nonclassicality

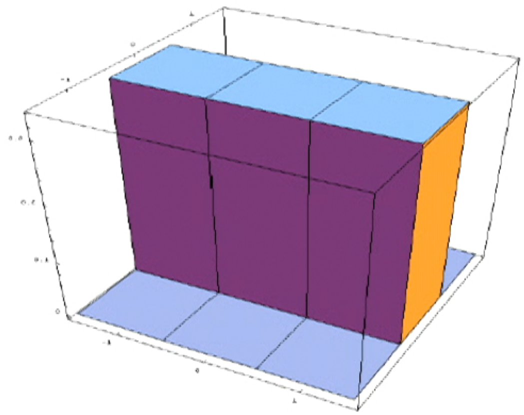

Classical

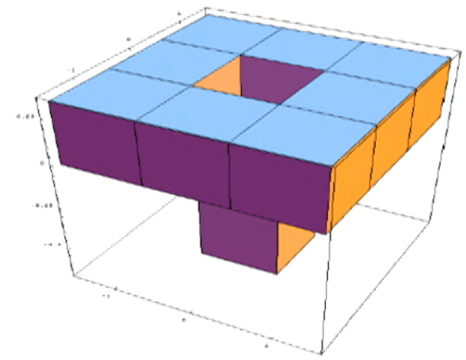

Quantum

The landscape

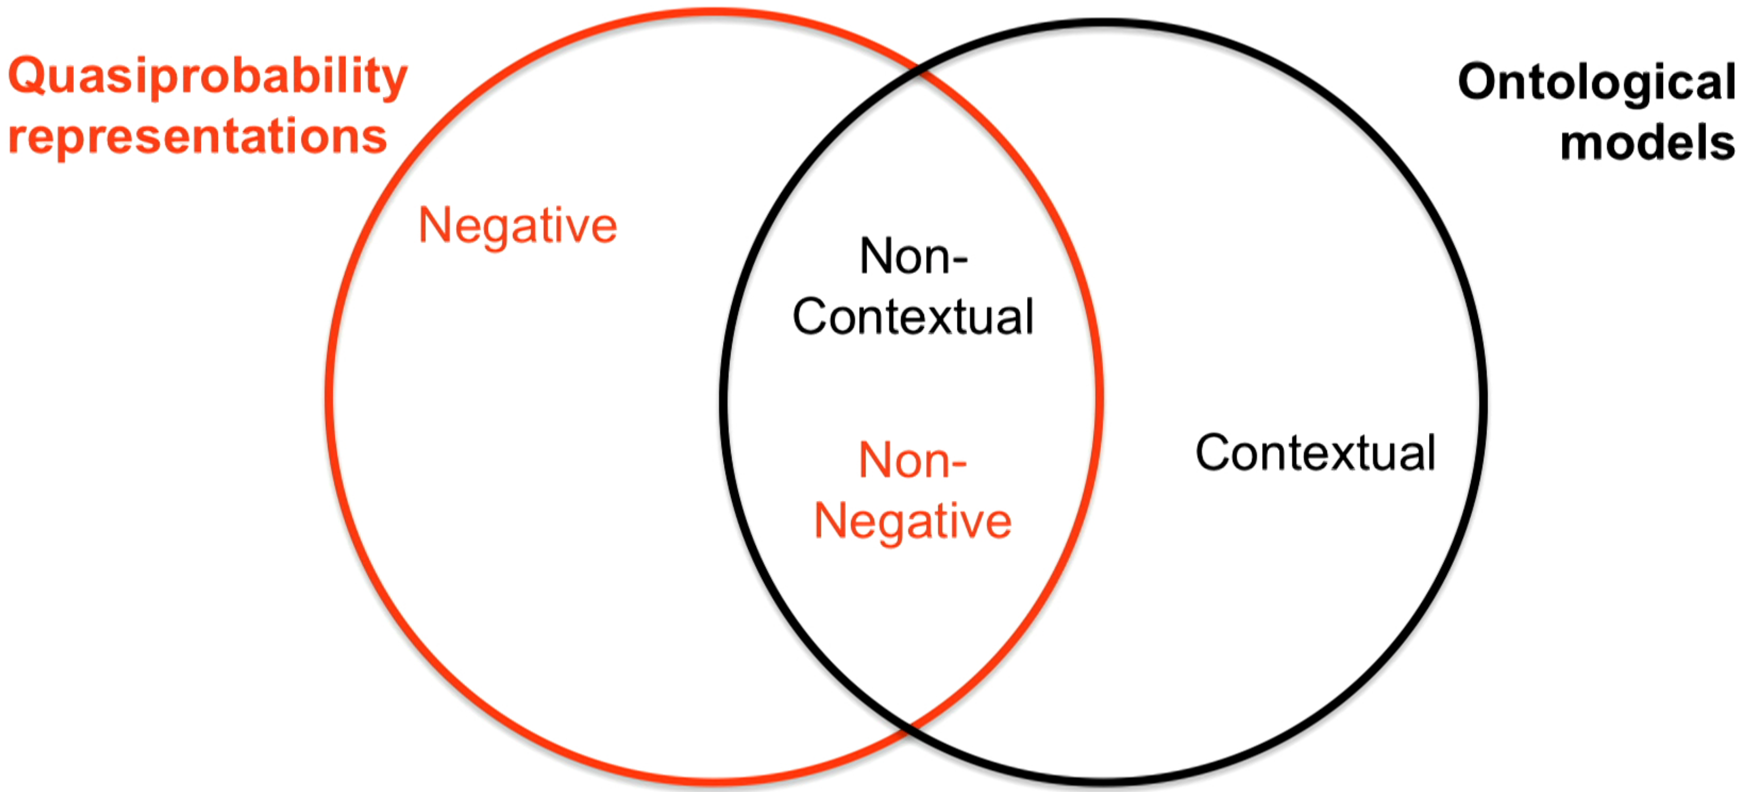

Spekkens, PRL (2008)

Operationalizing nonclassicality

Negativity in a quasiprobability can be related to notions of nonclassicality

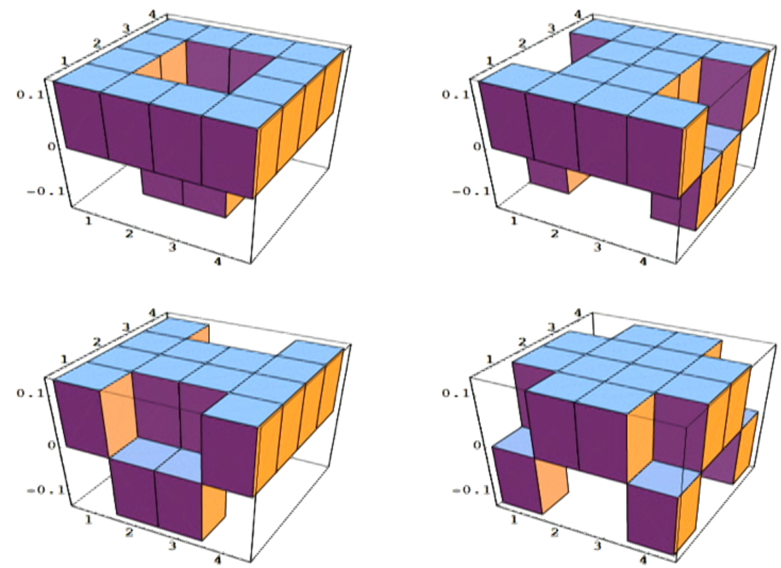

Operationalizing nonclassicality

Negativity in a quasiprobability can be related to notions of nonclassicality

- Negativity = contextuality

  Negativity in all quasiprobability representations is equivalent to a proof of contextuality

  Spekkens, PRL (2007)

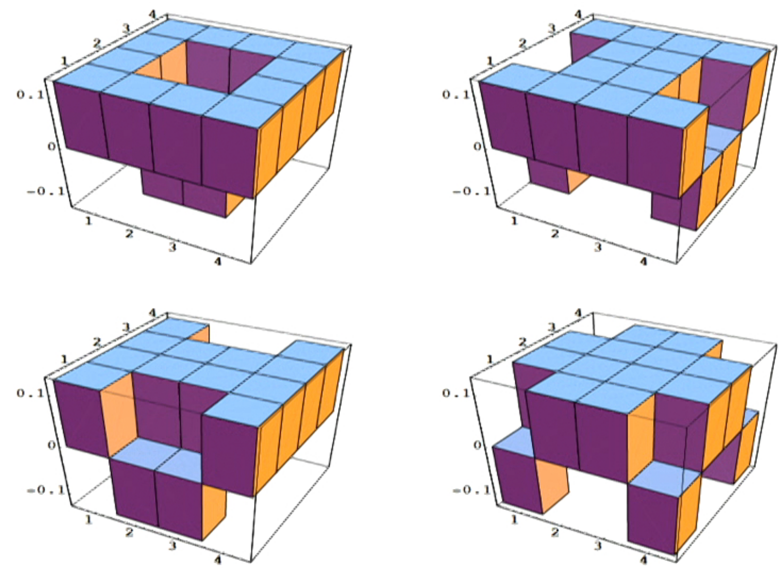

Operationalizing nonclassicality

Negativity in a quasiprobability can be related to notions of nonclassicality

- Negativity = contextuality

  Negativity in all quasiprobability representations is equivalent to a proof of contextuality
  Spekkens, PRL (2007)

- Negativity = simulation cost

  Negativity quantifies the rate of convergence of Monte Carlo methods
  Pashayan, Wallman, Bartlett, PRL (2015)

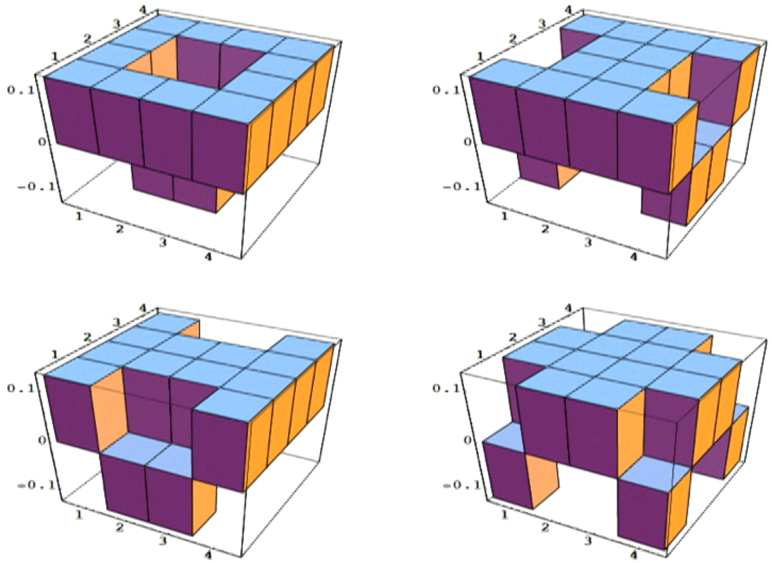

Operationalizing nonclassicality

Negativity in a quasiprobability can be related to notions of nonclassicality

- Negativity = contextuality

  Negativity in all quasiprobability representations is equivalent to a proof of contextuality

  Spekkens, PRL (2007)

- Negativity = simulation cost

  Negativity quantifies the rate of convergence of Monte Carlo methods

  Pashayan, Wallman, Bartlett, PRL (2015)

- Negativity = magic

  Negative states are those that can be distilled to magic states, that can supplement Clifford gates to allow universal quantum computation

  Veitch, Mousavian, Gottesman, Emerson, NJP (2014)
  Howard, Wallman, Veitch, Emerson, Nature (2014)

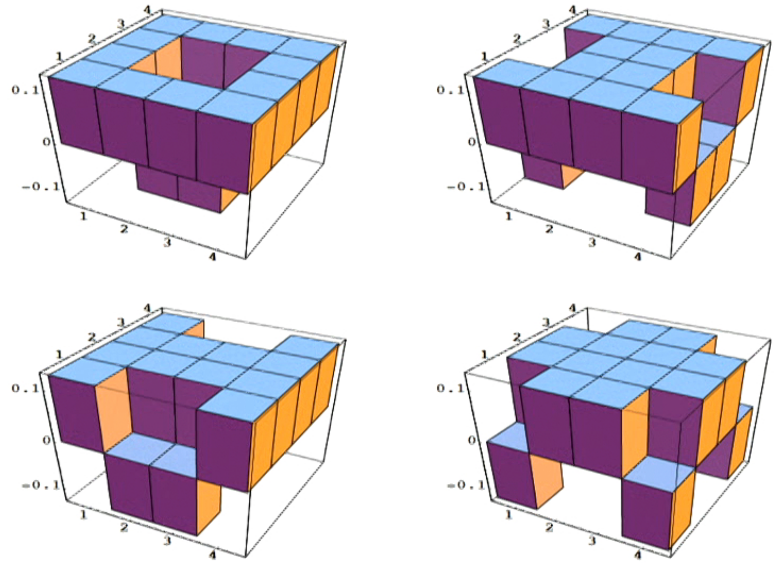

The landscape

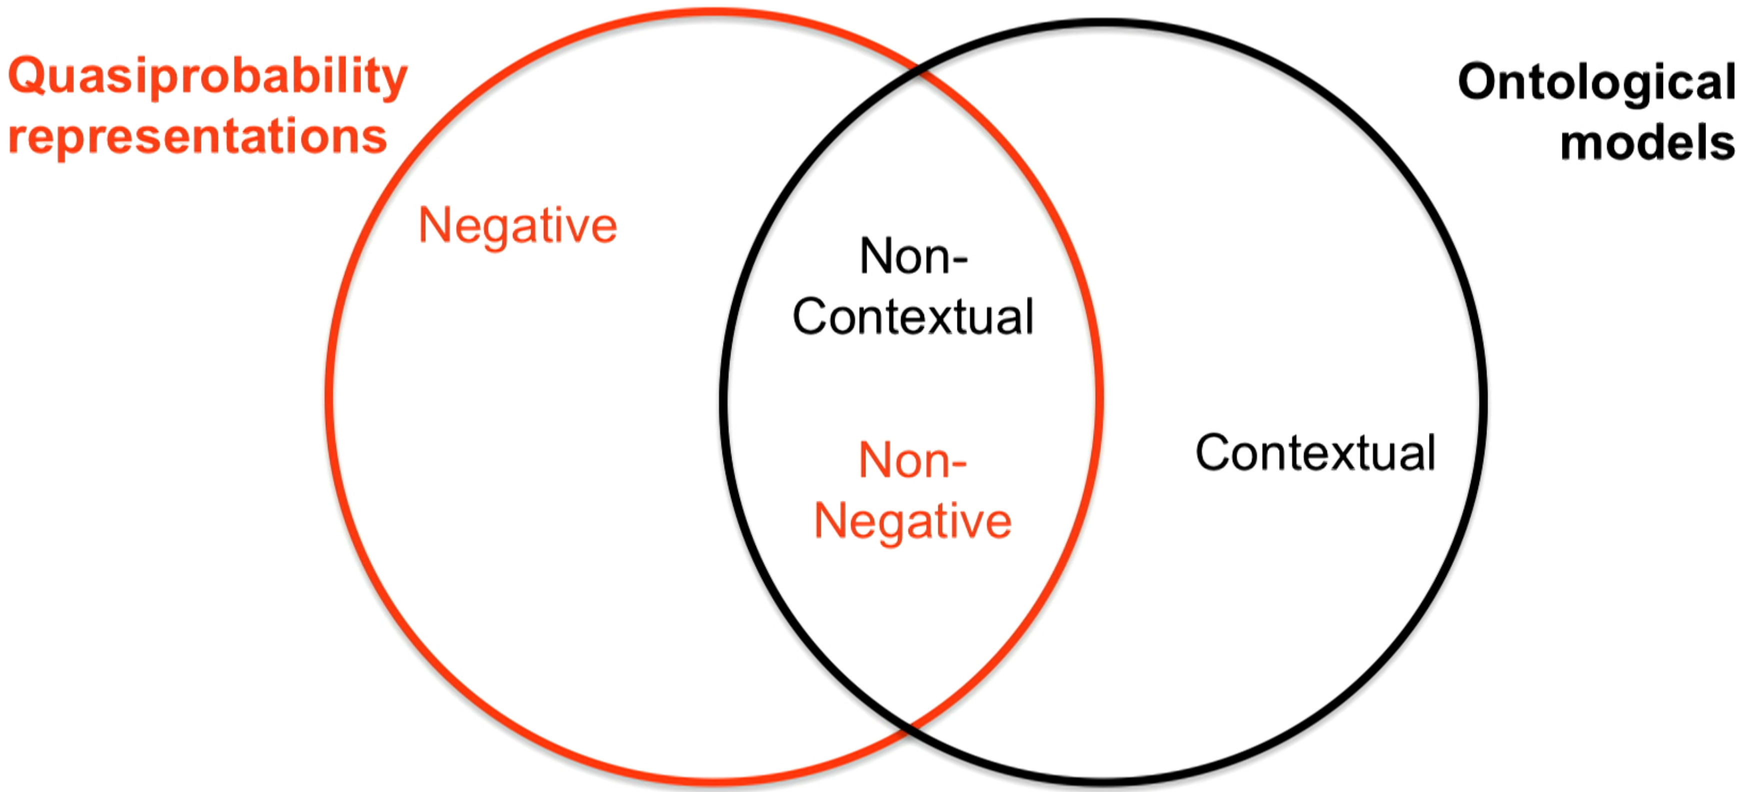

The landscape

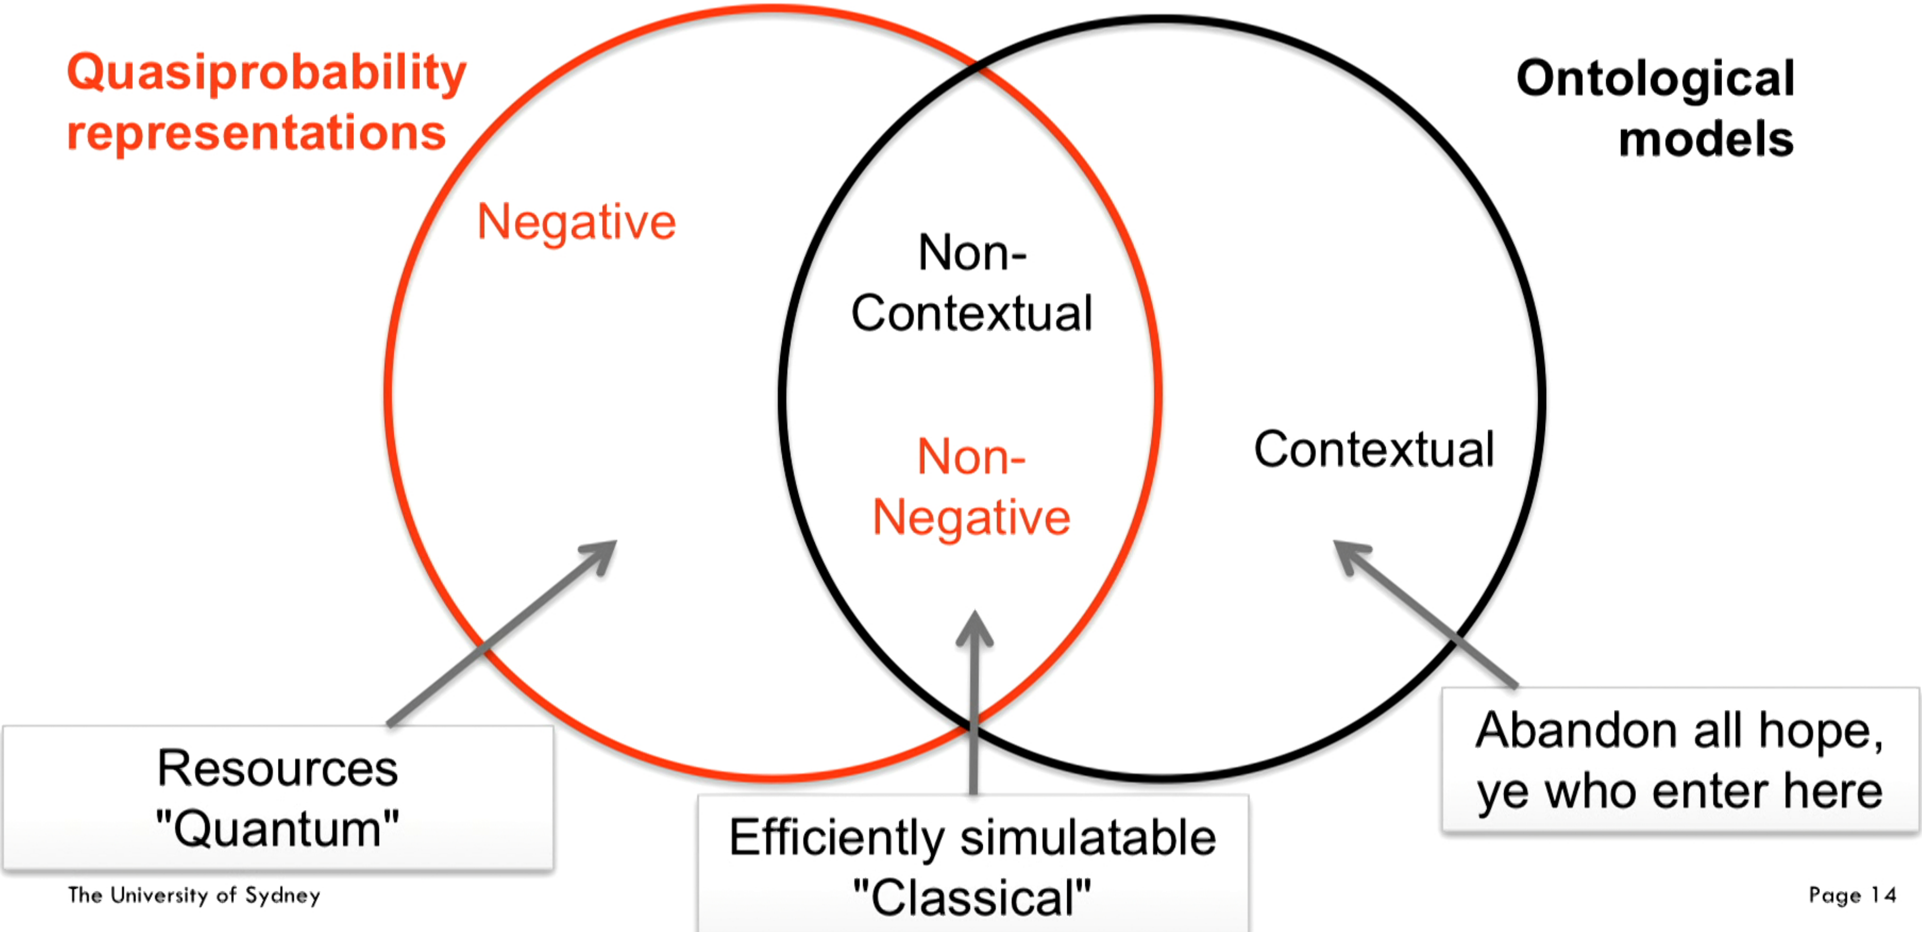

Estimating outcome probabilities of quantum circuits using quasiprobabilities

Phys. Rev. Lett. 115, 070501 (2015)
arXiv:1503.07525

Hakop Pashayan, Joel Wallman, and Stephen Bartlett

THE UNIVERSITY OF SYDNEY

Structure of our result

› Can we push the boundary on simulatability?

1. Quantify negativity – review
2. Poly-precision estimators for Born rule probabilities
3. Born rule probabilities as quasiprobabilistic sum over trajectories
4. Construct a true probability distribution of trajectories as a Markov chain
5. Construct an unbiased estimator
6. Bound convergence of this estimator in terms of the amount of negativity

Structure of our result

› Can we push the boundary on simulatability?

1. Quantify negativity – review
2. Poly-precision estimators for Born rule probabilities
3. Born rule probabilities as quasiprobabilistic sum over trajectories
4. Construct a true probability distribution of trajectories as a Markov chain
5. Construct an unbiased estimator
6. Bound convergence of this estimator in terms of the amount of negativity

Main Result

Estimator converges to true quantum mechanical probability at a rate determined by the amount of negativity in the circuit

If the negativity is polynomially bounded -> efficiently yields a poly-precision estimate

Pashayan, Wallman, Bartlett (2015)

Quantifying negativity

Veitch, Mousavian, Gottesman, Emerson, NJP (2014)

Quantifying negativity

Veitch, Mousavian, Gottesman, Emerson, NJP (2014)

Define the *negativity* of a state: the 1-norm of its quasiprobability representation

$$\mathcal{M}_\rho = \sum_{\lambda \in \Lambda} |W_\rho(\lambda)|$$

Quantifying negativity

Veitch, Mousavian, Gottesman, Emerson, NJP (2014)

Define the *negativity* of a state: the 1-norm of its quasiprobability representation

$$\mathcal{M}_\rho = \sum_{\lambda \in \Lambda} |W_\rho(\lambda)|$$

Negativity is multiplicitive, not additive (could take the log of this quantity)

Quantifying negativity

Veitch, Mousavian, Gottesman, Emerson, NJP (2014)

Define the *negativity* of a state: the 1-norm of its quasiprobability representation

$$\mathcal{M}_\rho = \sum_{\lambda \in \Lambda} |W_\rho(\lambda)|$$

Negativity is multiplicitive, not additive (could take the log of this quantity)

If W_ρ is nonnegative, then $\mathcal{M}_\rho = 1$

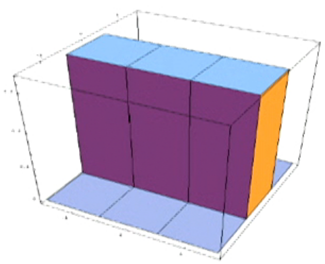

$\mathcal{M}_\rho = 1$

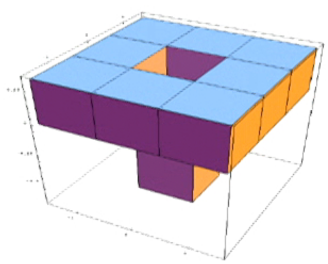

$\mathcal{M}_\rho > 1$

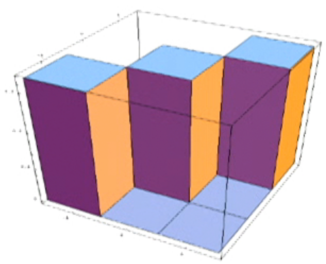

$\mathcal{M}_\rho = 1$

Negativity for states, unitaries, measurements

Quantum States

$$\mathcal{M}_\rho = \sum_{\lambda\in\Lambda} |W_\rho(\lambda)|$$

Measurements (POVM elements)

$$\mathcal{M}_E = \sum_{\lambda\in\Lambda} |W(E|\lambda)|$$

Unitaries

Point negativity: $$\mathcal{M}_U(\lambda) = \sum_{\lambda'\in\Lambda} |W_U(\lambda'|\lambda)|$$

Negativity: $$\mathcal{M}_U = \max_{\lambda\in\Lambda}\mathcal{M}_U(\lambda)$$

Estimating measurement probabilities

Trajectories in phase space

What do quasiprobabilities tell us about the probabilities of measurement outcomes?

$$p = \sum_{\lambda_0,\lambda_1,\ldots,\lambda_L} W(E|\lambda_L)W_{U_L}(\lambda_L|\lambda_{L-1})\cdots W_{U_1}(\lambda_1|\lambda_0)W_\rho(\lambda_0)$$

Trajectories in phase space

What do quasiprobabilities tell us about the probabilities of measurement outcomes?

$$p = \sum_{\lambda_0, \lambda_1, \ldots, \lambda_L} W(E|\lambda_L) W_{U_L}(\lambda_L|\lambda_{L-1}) \cdots W_{U_1}(\lambda_1|\lambda_0) W_\rho(\lambda_0)$$

Trajectories through phase space

Trajectories in phase space

What do quasiprobabilities tell us about the probabilities of measurement outcomes?

$$p = \sum_{\lambda_0,\lambda_1,\ldots,\lambda_L} \underbrace{W(E|\lambda_L)W_{U_L}(\lambda_L|\lambda_{L-1})\cdots W_{U_1}(\lambda_1|\lambda_0)W_\rho(\lambda_0)}$$

Trajectories through phase space

Quasiprobability associated to each trajectory

If these were all nonnegative, it provides a natural estimation algorithm

Veitch, Ferrie, Gross, Emerson, NJP (2012)
Mari and Eisert, PRL (2012)

But what if they are negative?

Can we estimate *p* by sampling from some true probability distribution?

What's a good estimator?

What would make a good estimator of a probability associated with a measurement outcome?

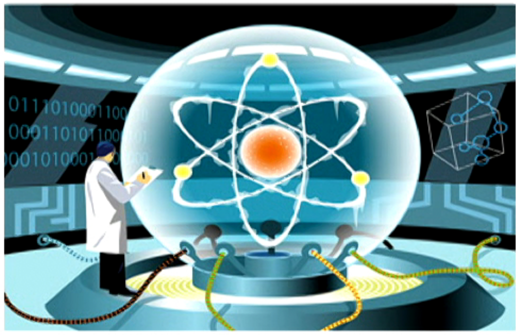

Poly-precision estimator: for any fixed confidence, yields an estimate within ε of the true Born rule probability using resources that scale polynomially in $1/\varepsilon$.

True probabilities from quasiprobabilities

Quasiprobability for a trajectory

$$W(\vec{\lambda}) = W(E|\lambda_L)W_{U_L}(\lambda_L|\lambda_{L-1})\cdots W_{U_1}(\lambda_1|\lambda_0)W_\rho(\lambda_0)$$

May be negative, so how do we sample?

True probabilities from quasiprobabilities

Quasiprobability for a trajectory

$$W(\vec{\lambda}) = W(E|\lambda_L)W_{U_L}(\lambda_L|\lambda_{L-1})\cdots W_{U_1}(\lambda_1|\lambda_0)W_\rho(\lambda_0)$$

May be negative, so how do we sample?

First attempt: sample from $\mathrm{Pr}(\vec{\lambda}) = \frac{|W(\vec{\lambda})|}{\mathcal{M}_c}$ $\quad \mathcal{M}_c = \sum_{\vec{\lambda}} |W(\vec{\lambda})|$

Estimate of the probability for each trajectory is $\hat{q}_1 = \mathcal{M}_c \mathrm{Sign}[W(\vec{\lambda})]$

This gives an unbiased estimator, minimizes the range, and has the smallest variance of all estimators over the space of trajectories…

True probabilities from quasiprobabilities

Quasiprobability for a trajectory

$$W(\vec{\lambda}) = W(E|\lambda_L)W_{U_L}(\lambda_L|\lambda_{L-1})\cdots W_{U_1}(\lambda_1|\lambda_0)W_\rho(\lambda_0)$$

May be negative, so how do we sample?

First attempt: sample from $\Pr(\vec{\lambda}) = \frac{|W(\vec{\lambda})|}{\mathcal{M}_c}$ $\quad \mathcal{M}_c = \sum_{\vec{\lambda}} |W(\vec{\lambda})|$

Estimate of the probability for each trajectory is $\hat{q}_1 = \mathcal{M}_c \mathrm{Sign}[W(\vec{\lambda})]$

This gives an unbiased estimator, minimizes the range, and has the smallest variance of all estimators over the space of trajectories…

But is impossible to sample from!

Our algorithm

Circuit with an efficient description (product input + output, local unitaries)

1. Sample initial point in trajectory from modified distribution

$$\Pr(\lambda_0) = |W_\rho(\lambda_0)|/\mathcal{M}_\rho$$

Our algorithm

Circuit with an efficient description (product input + output, local unitaries)

1. Sample initial point in trajectory from modified distribution

$$\Pr(\lambda_0) = |W_\rho(\lambda_0)|/\mathcal{M}_\rho$$

2. At each timestep $l=0,\ldots,L$, sample from conditional distribution

$$\Pr(\lambda_l|\lambda_{l-1}) = |W_{U_l}(\lambda_l|\lambda_{l-1})|/\mathcal{M}_{U_l}(\lambda_{l-1})$$

3. Estimate based on single trajectory

$$\hat{p}_1(\lambda) = \mathcal{M}_\rho \mathrm{Sign}[W_\rho(\lambda_0)] \prod_{l=1}^{L} \left[\mathcal{M}_{U_l}(\lambda_{l-1}) \mathrm{Sign}[W_{U_l}(\lambda_l|\lambda_{l-1})]\right] W_E(\lambda_L)$$

Properties of this estimate

Properties of estimator $\hat{p}_1(\lambda) = \mathcal{M}_\rho \mathrm{Sign}[W_\rho(\lambda_0)] \prod_{l=1}^{L} [\mathcal{M}_{U_l}(\lambda_{l-1}) \mathrm{Sign}[W_{U_l}(\lambda_l|\lambda_{l-1})]] W_E(\lambda_L)$

Properties of this estimate

Properties of estimator $\hat{p}_1(\lambda) = \mathcal{M}_\rho \mathrm{Sign}[W_\rho(\lambda_0)] \prod_{l=1}^{L} [\mathcal{M}_{U_l}(\lambda_{l-1}) \mathrm{Sign}[W_{U_l}(\lambda_l|\lambda_{l-1})]] W_E(\lambda_L)$

- Efficiently computable
- Unbiased estimator of Born rule probability

$$
\begin{aligned}
\langle \hat{p}_1(\vec{\lambda}) \rangle &= \sum_{\vec{\lambda}} \hat{p}_1(\vec{\lambda}) \mathrm{Pr}(\vec{\lambda}) \\
&= \sum_{\vec{\lambda}} \hat{p}_1(\lambda) \frac{|W_\rho(\lambda_0)|}{\mathcal{M}_\rho} \prod_{l=1}^{L} \frac{|W_{U_l}(\lambda_l|\lambda_{l-1})|}{\mathcal{M}_{U_l}(\lambda_{l-1})} \\
&= \sum_{\vec{\lambda}} W_\rho(\lambda_0) \prod_{l=1}^{L} W_{U_l}(\lambda_l|\lambda_{l-1}) W_E(\lambda_L) \\
&= \mathrm{Pr}(E|\rho, U)
\end{aligned}
$$

Sampling and convergence

Compute $\hat{p}_1(\lambda)$ for *s* independent trajectories, take the average

- Unbiased, and bound to the interval $[-\mathcal{M}, +\mathcal{M}]$
- Use Hoeffding inequality for upper bound on convergence:

Average of *s* samples will be within ϵ of the quantum probability with probability $1 - \delta$ if the total number of samples taken is

$$s(\epsilon, \delta) = \frac{2}{\epsilon^2}\mathcal{M}^2 \ln(2/\delta)$$

Sampling and convergence

Compute $\hat{p}_1(\lambda)$ for *s* independent trajectories, take the average

- Unbiased, and bound to the interval $[-\mathcal{M}, +\mathcal{M}]$
- Use Hoeffding inequality for upper bound on convergence:

Average of *s* samples will be within ϵ of the quantum probability with probability $1 - \delta$ if the total number of samples taken is

$$s(\epsilon, \delta) = \frac{2}{\epsilon^2}\mathcal{M}^2 \ln(2/\delta)$$

If the total negativity grows at most polynomially in *N*, we have an efficient estimate of the quantum probability to within $\epsilon = 1/poly(N)$, with an exponentially small failure probability

The good…

- Can simulate *any* quantum process, perhaps inefficiently
- Quantifies the efficiency using a measure of 'amount of contextuality' (negativity)
- It's actually useful!

The good…

- Can simulate *any* quantum process, perhaps inefficiently
- Quantifies the efficiency using a measure of 'amount of contextuality' (negativity)
- It's actually useful!

the bad…

- Estimating outcomes, not simulating processes
  (a challenge when there are many possible outcomes)

The good…

- Can simulate *any* quantum process, perhaps inefficiently
- Quantifies the efficiency using a measure of 'amount of contextuality' (negativity)
- It's actually useful!

the bad…

- Estimating outcomes, not simulating processes
  (a challenge when there are many possible outcomes)

and the ugly

- Individual runs are not sampled from the correct distribution, only converges on average

The good…

- Can simulate *any* quantum process, perhaps inefficiently
- Quantifies the efficiency using a measure of 'amount of contextuality' (negativity)
- It's actually useful!

the bad…

- Estimating outcomes, not simulating processes
  (a challenge when there are many possible outcomes)

and the ugly

- Individual runs are not sampled from the correct distribution, only converges on average

We should be able to do better in some cases
e.g., the qubit stabilizer subtheory

The landscape

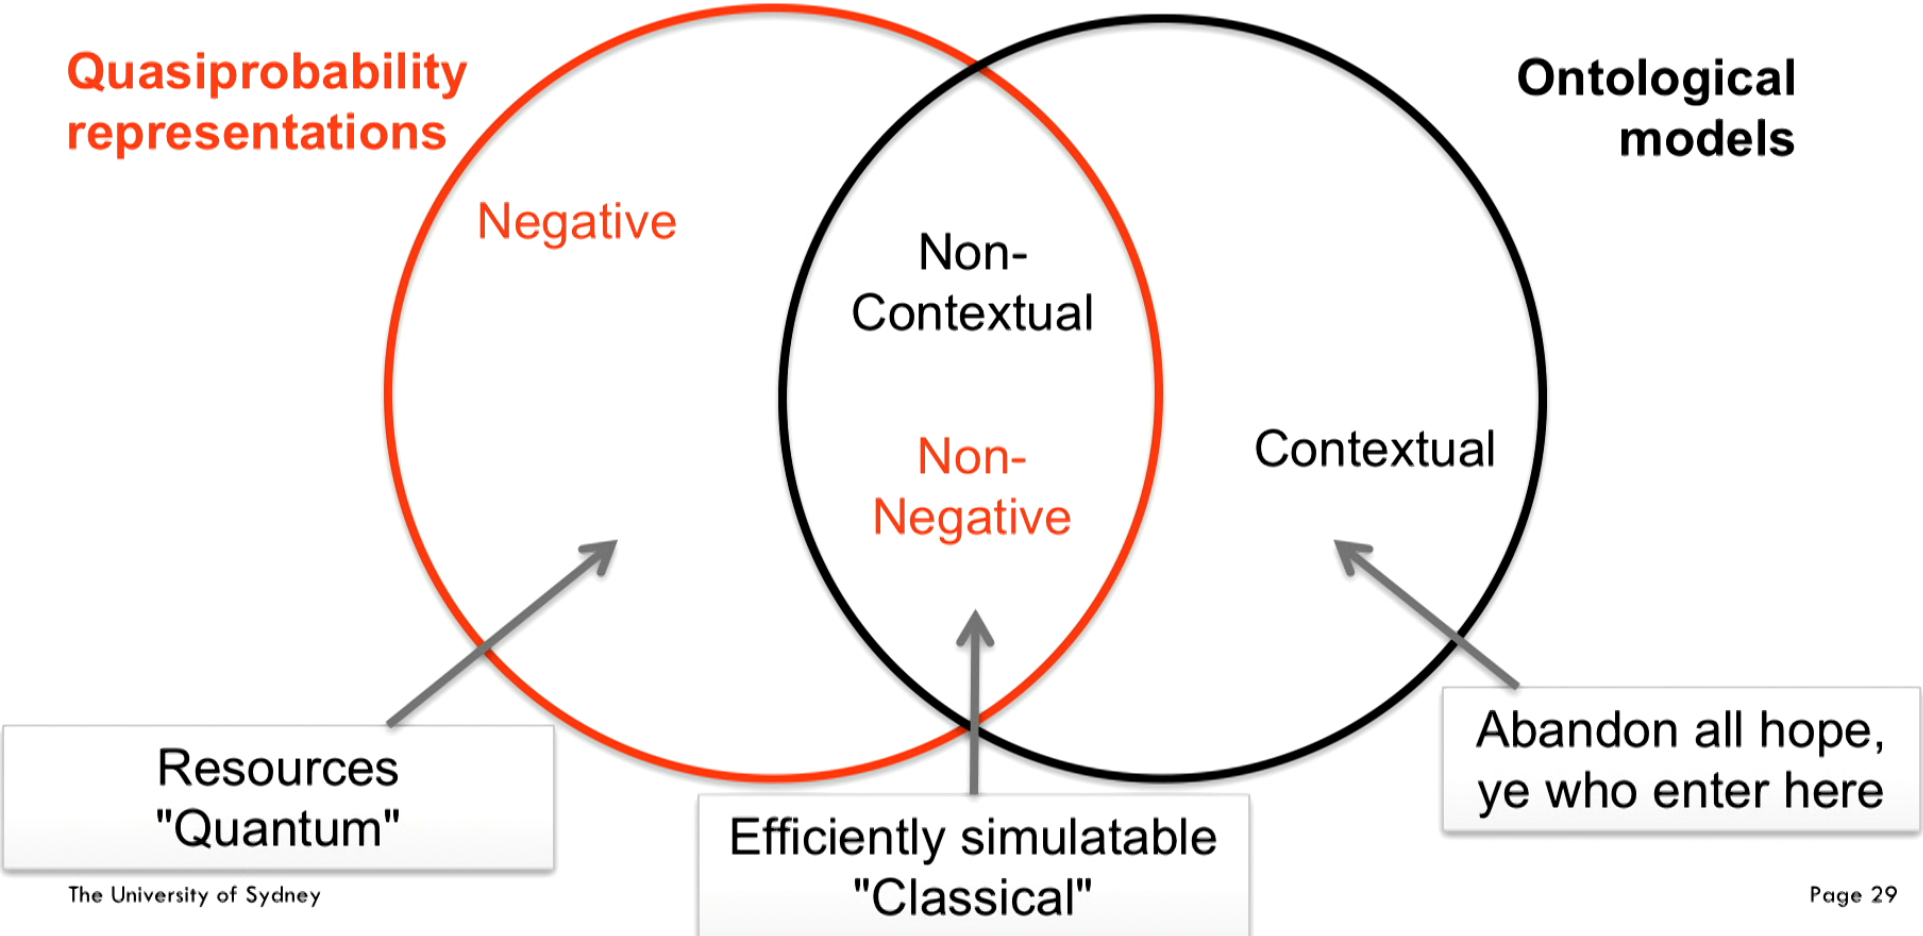

Understanding contextuality in the qubit stabilizer formalism

Qubit stabilizer theory is efficiently simulatable, just like qudit stabilizers

It is basically classical

But the theory is negative in any quasiprobability representation

It allows state-independent proofs of contextuality (Peres-Mermin, GHZ)

It is very nonclassical

A case for contextual ontological models

- What to do about contextuality?
  1. Resource theory approach: ban it
  2. Simulation approach: embrace it!

A case for contextual ontological models

- What to do about contextuality?
  1. Resource theory approach: ban it
  2. Simulation approach: embrace it!

- 'Better' simulations via exact Monte Carlo sampling
- Adding a context to the statistical model increases complexity of simulation
  - Analogy: simulating a system coupled to a non-Markovian environment

The landscape