

2008

Nonclassicality without Entanglement Enables Bit Commitment

Howard Barnum

Los Alamos National Laboratory

Oscar C. O. Dahlsten

Institute for Theoretical Physics

Matthew S. Leifer

Chapman University, leifer@chapman.edu

Ben Toner

Centrum voor Wiskunde en Informatica

Follow this and additional works at: http://digitalcommons.chapman.edu/scs_articles

 Part of the [Quantum Physics Commons](#)

Recommended Citation

Barnum, H., Dahlsten, O.C.O., Leifer, M., Toner, B. 2008. Nonclassicality without Entanglement Enables Bit Commitment. *Proceedings of IEEE Information Theory Workshop*, 2008, 386—390.

This Conference Proceeding is brought to you for free and open access by the Science and Technology Faculty Articles and Research at Chapman University Digital Commons. It has been accepted for inclusion in Mathematics, Physics, and Computer Science Faculty Articles and Research by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Nonclassicality without Entanglement Enables Bit Commitment

Comments

This is an author-prepared, prepublication version of a paper that was presented at the IEEE Information Theory Workshop in 2008 and was later published in *Proceedings of IEEE Information Theory Workshop 2008*, pp. 386—390.

Copyright

IEEE

Nonclassicality without entanglement enables bit commitment

Howard Barnum

CCS-3: Information Sciences, and Quantum Institute
Los Alamos National Laboratory
Los Alamos, NM USA
Email: barnum@lanl.gov

Oscar C.O. Dahlsten

Institute for Theoretical Physics
ETH Zürich, Switzerland
Email: dahlsten@phys.ethz.ch

Matthew Leifer

Perimeter Institute for Theoretical Physics
and Institute for Quantum Computing
University of Waterloo
Waterloo, Ontario, Canada
Email: matt@mattleifer.info

Ben Toner

Centrum voor Wiskunde en Informatica
Amsterdam, The Netherlands
Email: Ben.Toner@cwi.nl

Abstract—We investigate the existence of secure bit commitment protocols in the convex framework for probabilistic theories. The framework makes only minimal assumptions, and can be used to formalize quantum theory, classical probability theory, and a host of other possibilities. We prove that in all such theories that are locally non-classical but do not have entanglement, there exists a bit commitment protocol that is exponentially secure in the number of systems used.

I. INTRODUCTION

In the 1984 paper [1] in which they introduced information-theoretically secure quantum key distribution, Bennett and Brassard also considered the possibility of information-theoretically secure bit commitment. Bit commitment is a basic primitive in classical cryptography, to which many practically important cryptographic tasks, such as secure function evaluation, can be reduced. In a bit commitment protocol, one party, usually called Alice, performs some act that is supposed to irrefutably convince another party, Bob, that she has irrevocably committed to a value, 0 or 1, of a bit, without leaking any information about the value of the bit to Bob. Later she can perform another act that reveals the value of the bit to Bob and enables him to perform some test that may be necessary for him to verify that she was indeed committed. Classically, bit commitment can be achieved with *computational* security, but not with information-theoretic security.

Bennett and Brassard showed that the bit commitment scheme they considered could be defeated by the use of entangled states. Attempts were made [2] to construct secure bit commitment protocols, but Lo and Chau [3], and independently Mayers [4], showed that an entangled attack akin to Bennett and Brassard's defeats all quantum bit commitment protocols, and there is now a solid consensus that this does indeed cover all reasonable schemes and attacks [5].

Soon after this development, Brassard [6] and Fuchs [7] asked whether the impossibility of bit commitment might be

a manifestation of a deep information-theoretic property of quantum mechanics, fit for a crucial role in an information-theoretic characterization, or reconstruction, of the formalism of quantum theory. Such a reconstruction, at its most ambitious, is envisioned as similar to Einstein's reconstruction of the dynamics and kinetics of macroscopic bodies on the basis of simple principles with clear operational meanings and experimental consequences. As argued in (for example) [8], [9], [7], such a reconstruction could lend force to the view that the foundations of quantum mechanics are properly couched in terms of information, a view which has received increasing attention with the rise of quantum information science. Short of this ambitious goal, there are still strong reasons to pursue an informational characterization of quantum mechanics. It should lead to a principled understanding of the features of quantum mechanics that account for its better-than-classical information processing power. Such an understanding could help guide the search for new algorithms and protocols, both positively, by providing conceptual tools to exploit in a variety of settings, and negatively by identifying information-processing tasks requiring properties that quantum mechanics lacks.

Brassard and Fuchs' conjecture was that the impossibility of bit commitment might, in conjunction with the possibility of secure secret key distribution and the impossibility of instantaneous signaling between distinct physical systems, suffice to characterize quantum theory. Clifton, Bub, and Halvorson proved a result (the CBH theorem) [8], close to this conjecture in the framework of C^* -algebraic theories. They demonstrated the existence of a protocol related to the no-bit commitment theorem, but weaker, between two "local" algebras, whenever the local algebras are not commutative (not classical) and there are entangled states between the algebras. However, in finite dimensions, C^* -algebraic theories are essentially quantum mechanics with superselection rules, so in our view,

a much broader framework is desirable. Further evidence for this view is Halvorson’s demonstration [10] that no-bit-commitment follows from no-signaling and no-cloning within the C^* -algebraic framework. To obtain the most illuminating characterization of quantum mechanics in terms of information processing, one should work in a framework wide enough to include not only quantum and classical mechanics, but also a wide variety of other theories that can serve as foils to them; the C^* -algebraic framework is too restrictive.

It is therefore an open question whether non-classical theories without entanglement are ruled out by demanding the impossibility of secure bit commitment, in some appropriately broad framework. In this paper, we answer that question in the affirmative. We work in a framework that allows for a wide range of probabilistic theories, including not only quantum and classical theories, but also theories of Popescu-Rohrlich, or nonlocal, boxes [11], [12] that allow nonlocality stronger than that in quantum mechanics, as well as many other types of theory. For any nonclassical theory within the framework that does not permit entanglement between systems, we construct a bit-commitment protocol that is exponentially secure in the number of systems used.

We proceed as follows. First the framework of generalized probabilistic theories is introduced and our bit-commitment protocol is defined. We then prove that such a protocol always exists in a non-classical theory. Next, we prove it to be exponentially secure in all theories that don’t allow entanglement. Finally we give a summary and discussion.

II. THE FRAMEWORK

The framework is that of *convex operational* or *generalized probabilistic* theories, for which no-cloning and no-broadcasting theorems were proved in [13], [14], to which we refer for further background. The set of *normalized states* of a system is a compact convex set $\Omega \subseteq \mathbb{R}^d$. Embed Ω in \mathbb{R}^{d+1} , avoiding the origin, and let $\text{Cone}(\Omega)$ be the set of linear combinations of elements of Ω with nonnegative coefficients—the convex cone of *unnormalized* states. Its *dual cone*, $\text{Cone}(\Omega)^*$, consists of those linear functionals from \mathbb{R}^{d+1} to \mathbb{R} that are nonnegative on $\text{Cone}(\Omega)$. Measurement outcomes are represented as *effects*: functionals $e \in \text{Cone}(\Omega)^*$ satisfying $e(\omega) \leq 1$ for all $\omega \in \Omega$. $e(\omega)$ is the probability of outcome e for a system prepared in state ω . Equivalently, effects are elements of the interval $[0, u]$ in the dual cone, whose endpoints are the zero functional and the unit functional u that gives 1 on all normalized states. *Measurements* are sets $\{e_i\}$ of effects with $\sum_i e_i = u$ (i.e. $\forall \omega \in \Omega, \sum_i e_i(\omega) = 1$).

For two state spaces, Ω_A and Ω_B , a spectrum of possible “tensor products” is identified—these are candidates for describing a composite system built from subsystems with state spaces Ω_A and Ω_B . In this work we need only one:

Definition: The *minimal tensor product* $\Omega_A \otimes \Omega_B$ is the convex hull of the set of product states $(\omega_A, \omega_B) \in \Omega_A \times \Omega_B$.

This generalizes the quantum-mechanical construction of the unentangled or *separable* density matrices. The general framework requires only that a tensor product be convex,

contain the minimal tensor product, and be contained in what’s known as the *maximal tensor product*, of less interest here.

To describe quantum theory in this framework, Ω is chosen to be isomorphic to the set of density operators on a Hilbert space and $\text{Cone}(\Omega)$ is the set of positive operators. The quantum tensor product lies strictly between the minimal and maximal tensor products. In classical theory, Ω is a simplex of probability distributions, i.e. the convex hull of $d + 1$ linearly independent points in \mathbb{R}^{d+1} , and the maximal and minimal tensor products coincide so there is no choice. Classical theories are, equivalently, characterized by the property that any state in Ω has a unique convex decomposition into pure (extremal) elements.

It is important to specify the dynamics of theories in this framework, because this specifies what Alice and Bob can do to their systems. In this framework, dynamics are *positive* linear maps $\mathcal{L} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^{d+1}$, i.e. ones that take $\text{Cone}(\Omega)$ to itself. Thus they take (not-necessarily-normalized) states to states. Further, they must be *norm-nonincreasing*: for all states $\omega \in \text{Cone}(\Omega)$, $u(\mathcal{L}(\omega)) \leq u(\omega)$; we use the term *operation*, standard for the quantum case, to denote these. The map $e_{\mathcal{L}} : \omega \mapsto u(\mathcal{L}(\omega))$ is an element of $[0, u]$, and is interpreted as an effect (measurement outcome) associated with the dynamics \mathcal{L} . Thus for normalized ω and positive \mathcal{L} , $e_{\mathcal{L}}(\omega)$ is interpreted as the probability with which the state undergoes \mathcal{L} . When $e_{\mathcal{L}} = u$, the map is *norm-preserving*; it is an *unconditional* dynamics not associated with obtaining a particular measurement outcome.

Early work on cryptography using stronger-than-quantum nonlocal correlations, including [15] and [16] where *entangled* correlations enabled bit commitment, did not situate these correlations in a unified framework describing dynamics, measurement, and state preparation such as the one we use here.

The assumptions embodied in this framework [14] are fairly minimal. Two are substantive: first, the “local observability” assumption effectively states that there are no “intrinsically nonlocal” degrees of freedom that cannot be determined by making repeated local measurements on the subsystems of identically prepared systems. Second, a “no-signaling” constraint, which it is reasonable to take as the definition of what we mean by an independent subsystem.

Our protocol uses the fact that any nonclassical state-space contains states that have more than one distinct convex decomposition into pure states. Alice encodes which bit she has committed to as a choice of one out of two such decompositions. The security analysis we give requires that the two sets of pure states used in the decompositions be *disjoint*, and that all the states be *exposed*, but this can be achieved in any nonclassical state space. A state is *exposed* if there is a measurement outcome whose probability is 1 in that state, and strictly less than 1 on any other state—an outcome that can be guaranteed by that state, and only by that state. We call such an effect the *distinguishing effect* for the state in question. It is immediate from the definitions that exposed states are pure.

We write $\text{cl}(S)$, $\text{conv}(S)$, and $\text{Exp}(S)$ for the topological closure, convex hull, and set of exposed points of a set S .

III. THE PROTOCOL

Let a system have a non-simplicial, convex, compact state space Ω of dimension d . The protocol uses a state μ that has two distinct convex decompositions $\{(p_i^0, \mu_i^0)\}, \{(p_j^1, \mu_j^1)\}$ into finite disjoint sets of exposed states, that is,

$$\mu = \sum_{i=1}^{N^0} p_i^0 \mu_i^0 = \sum_{j=1}^{N^1} p_j^1 \mu_j^1. \quad (1)$$

In the honest protocol, Alice first decides on a bit $b \in \{0, 1\}$ to commit to. She then draws n independent samples from the probability distribution $(p_1^b, p_2^b, \dots, p_{N^b}^b)$, obtaining a string $\mathbf{x} = (x_1, x_2, \dots, x_n)$. She sends the state $\mu_{\mathbf{x}}^b = \mu_{x_1}^b \otimes \mu_{x_2}^b \otimes \dots \otimes \mu_{x_n}^b$ to Bob.

In the reveal phase, she sends b and \mathbf{x} to Bob. Bob then measures each subsystem of the state Alice sent in the commit phase. On the k -th subsystem, he performs a measurement containing the distinguishing effect for $\mu_{x_k}^b$ and aborts if the result is not the distinguishing effect. If he obtains the appropriate distinguishing effect for every subsystem, he accepts.

Example of protocol: If Ω is the state space of a qubit, we can transpose the one-qubit protocol of [1] to our setting. Ω can be visualised as the Bloch sphere in \mathbb{R}^3 with pure states on the surface and their mixtures inside the sphere. Let μ be the center of the sphere, i.e. the completely mixed state $\frac{1}{2}I = \frac{1}{2}|+\rangle\langle+| + \frac{1}{2}|-\rangle\langle-| = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|$, where $|0\rangle, |1\rangle$ is a basis and $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. Let $\mu_1^0 = |0\rangle\langle 0|$, $\mu_2^0 = |1\rangle\langle 1|$, $\mu_1^1 = |+\rangle\langle+|$, $\mu_2^1 = |-\rangle\langle-|$ and $p_i^b = \frac{1}{2} \forall i, b$. In the $n = 1$ case, if Alice decides to commit to $b = 0$ for example, she would send Bob either $|0\rangle$ or $|1\rangle$, each with probability $\frac{1}{2}$. Say she sends $|0\rangle$. To reveal she tells him “ $b = 0$ ” and that she sent $|0\rangle$. Bob would then measure in the $|0\rangle, |1\rangle$ basis, find $|0\rangle$ and accept. In [1], Bennett and Brassard considered this $n = 1$ protocol and showed it was completely nonbinding through an entangled attack.

IV. EXISTENCE OF THE PROTOCOL

The existence of the protocol just described in any non-classical theory follows from:

Theorem 1: Every nonsimplicial convex compact set Ω of dimension d contains a state μ with two convex decompositions into disjoint sets of exposed states, whose total cardinality is less than $d + 2$.

The theorem follows from two lemmas.

Lemma 1: Let Ω be a non-simplicial compact convex set of dimension d . Then the convex hull of any $d + 2$ pure states in Ω contains a state μ which has two convex decompositions,

$$\mu = \sum_{i=1}^{N^0} p_i^0 \mu_i^0 = \sum_{j=1}^{N^1} p_j^1 \mu_j^1, \quad (2)$$

into disjoint sets of pure states, with $N^0 + N^1 \leq d + 2$.

Proof: Let $\Gamma := \{\mu_1, \dots, \mu_{d+2}\}$ be an arbitrary set of $d + 2$ pure states. Then $\text{conv}(\Gamma)$ is non-simplicial because Ω has dimension d . Choose a state ω with two different convex decompositions $\{(p_i^0, \mu_i^0), i = 1, \dots, N^0\}$ and $\{(p_j^1, \mu_j^1), j \in$

$1, \dots, N^1\}$ into elements of Γ so that $N^0 + N^1$ is minimal. The sets $\{\mu_1^0, \dots, \mu_{N^0}^0\}$ and $\{\mu_1^1, \dots, \mu_{N^1}^1\}$ are then disjoint. For if they had a state in common, say (reindexing if necessary) $\mu_1^0 = \mu_1^1$, then the (unnormalized) state $\omega' := \omega - \min_b(p_1^b) \mu_1^b$ would also have two different convex decompositions, contradicting minimality. \square

To show there are $d + 2$ exposed states we'll use the following special case of Theorem 18.7 of [17].

Theorem 2: A compact convex set $\Omega \subseteq \mathbb{R}^d$ is the closure of the convex hull of its exposed points, i.e. $\Omega = \text{cl}(\text{conv}(\text{Exp}(\Omega)))$.

Lemma 2: A nonsimplicial convex compact set Ω of dimension d has at least $d + 2$ exposed points.

Proof: By Theorem 2, the closure of the convex hull of $\text{Exp}(C)$ is equal to C , and therefore $\text{cl}(\text{Cone}(\text{Exp}(C))) = \text{Cone}(C)$. Taking the closure of a convex subset (compact or not) of \mathbb{R}^n can't increase the dimension of the subspace it spans, so the linear span of $\text{Exp}(C)$ must be \mathbb{R}^{d+1} , and we may pick a linearly independent subset of $\text{Exp}(C)$, consisting of $d + 1$ exposed points. There must be an exposed point not in the convex hull of these $d + 1$ points, for if not the convex hull of the exposed extreme points of C would be a simplex, whence, using Theorem 2 and the fact that a finite-dimensional simplex is closed, C itself would be a simplex. \square

Since exposed states are pure, Lemmas 1 and 2 immediately imply Theorem 1.

V. SECURITY OF THE PROTOCOL

We adapt our security definition from Ref. [18], simplifying to the setting where there is no communication from Bob to Alice. We start with the formal definition:

Definition: Let $\varepsilon \geq 0$. We say that a bit commitment protocol with one-way communication is ε -secure if it has the following properties:

- (ε -soundness) Assume that both parties are honest. Then the probability that Bob aborts is at most ε and, if he does not abort, then after the reveal phase he learns the bit b that Alice committed to.
- (ε -hiding) Assume that Alice is honest. Then for all cheating strategies of Bob aiming to guess the commitment before the reveal phase. $q_0 + q_1 \leq 1 + \varepsilon$, where q_b is the probability that Bob guesses correctly given that Alice committed b .
- (ε -binding) Assume that Bob is honest. Then for all commitments of Alice, $p_0 + p_1 \leq 1 + \varepsilon$, where p_b is the maximum probability that Alice successfully reveals b .

If any of the above hold for $\varepsilon = 0$, we say that the protocol satisfies that property *perfectly*.

Our protocol is perfectly *sound* because if Alice is honest, the distinguishing measurements that Bob makes based on Alice's claim give the correct answers with probability 1. In general, one would consider the probability of *either* honest participant accusing the other of cheating, but in a one-way protocol, there is no provision for Alice to abort.

The protocol is also perfectly *hiding*—there is no way for Bob to obtain information about the bit b during the post-commit, pre-revelation phase, as the state $\mu^{\otimes n}$ that (honest) Alice sent is independent of b .

The nontrivial part of the security analysis is to show that the protocol is ε -*binding*—to show that Alice can't cheat by choosing which bit to reveal after she is supposed to be committed to one or the other. In an ideal bit commitment protocol, Alice could use randomness to commit to 0 with probability p_0 and 1 with probabilities $p_1 = 1 - p_0$, so she can achieve any pair p_0, p_1 in the definition such that $p_0 + p_1 = 1$. Our protocol only allows her to do a little better. For example, if she wants to be able to reveal 0 with probability 1, then the probability that she can reveal 1 is at most ε . Although it is suitable for present purposes, we note that our definition of ε -binding is too weak to establish composable security [19].

We'll need a lemma about measurements.

Lemma 3: Suppose two exposed states $\mu \neq \nu$ have distinguishing effects a and b . Let

$$f(\mu, \nu) := \sup_{\omega \in \Omega} (a(\omega) + b(\omega)). \quad (3)$$

Then $1 \leq f(\mu, \nu) < 2$.

Proof: For the upper bound, the function $a + b$ is linear and the set Ω is convex and compact, so the supremum of $a + b$ is achieved on a pure state ω' . Suppose $a(\omega') + b(\omega') = 2$. Then we must have $a(\omega') = 1$ and $b(\omega') = 1$, which implies $\omega' = \mu = \nu$, a contradiction. The lower bound follows from considering $\omega = \mu$. \square

Now define $\delta := \min_{1 \leq i \leq N^0, 1 \leq j \leq N^1} (2 - f(\mu_i^0, \mu_j^1))$, where μ_i^0, μ_j^1 run over the states used in the protocol. Note that $\delta < 1$, since at least one pair of states μ_i^0, μ_j^1 is not perfectly distinguishable. This quantity δ will control the number n of systems we need to use to achieve ε -security.

The proof also uses the following description of an optimal set of strategies for a cheating Alice.

Lemma 4: An optimal strategy for Alice is as follows: she tosses some coins and generates randomness λ with probability weight $p(\lambda)$. She then prepares an arbitrary string of pure states $\omega_1^\lambda \otimes \omega_2^\lambda \otimes \dots \otimes \omega_n^\lambda$. She sends them to Bob. In the reveal phase, she can send an arbitrary bit b and an arbitrary “claim sequence” $\mathbf{x}^{\lambda, b}$, that depends on the bit she wants to claim and the randomness.

The state claim $\mathbf{x}^{\lambda, b}$, which is classical information, is encoded in perfectly distinguishable states of some systems Γ in the theory; it is easily shown that doing otherwise can't help Alice.

Proof of Lemma 4: A general cheating strategy for Alice is to prepare an arbitrary state in $\Upsilon \otimes \Gamma^{\otimes n} \otimes \Omega^{\otimes n}$ (where Υ is some state space in the theory), and then do a b -dependent positive map \mathcal{L}^b on $\Upsilon \otimes \Gamma^{\otimes n}$ just before sending $\Gamma^{\otimes n}$ to Bob, in an attempt to reveal b . Letting r^l be probabilities, $\tau^l \in \Upsilon$, $\gamma_k^l \in \Gamma$, $\omega_k^l \in \Omega$, the state before revelation is:

$$\sum_l r^l \tau^l \otimes \gamma_1^l \otimes \dots \otimes \gamma_n^l \otimes \omega_1^l \otimes \dots \otimes \omega_n^l. \quad (4)$$

After Alice attempts to reveal b it is:

$$\psi^b := \sum_l r^l \mathcal{L}^b(\tau^l \otimes \gamma_1^l \otimes \dots \otimes \gamma_n^l) \otimes \omega_1^l \otimes \dots \otimes \omega_n^l. \quad (5)$$

Let $\phi^{lb} \equiv \sum_m t^{lmb} \gamma_1^{lmb} \otimes \dots \otimes \gamma_n^{lmb}$ be the marginal state on $\Gamma^{\otimes n}$ induced by the state $\mathcal{L}^b(\tau^l \otimes \gamma_1^l \otimes \dots \otimes \gamma_n^l)$. Then the state of $\Gamma^{\otimes n} \otimes \Omega^{\otimes n}$ is

$$\chi^b := \sum_l r^l \sum_m t^{lmb} \gamma_1^{lmb} \otimes \dots \otimes \gamma_n^{lmb} \otimes \omega_1^l \otimes \dots \otimes \omega_n^l. \quad (6)$$

Bob will subject each copy of Γ to a standard measurement to read Alice's claim; the k -th system will yield a value x with probability $p_k^{lm}(x)$ determined by γ_k^{lmb} . Alice could achieve the same result by sampling the distribution of measurement results p_k^{lm} . Bob would obtain from γ_k^{lmb} , perhaps keeping a record q of the result of sampling, and sending a *definite* string $x_1^{lmqb} \otimes \dots \otimes x_n^{lmqb}$ for the claim, encoded as distinguishable states that will definitely give this string of outcomes. Letting λ stand for lmq , we see that an optimal strategy for Alice is as described in the Lemma. \square

Theorem 3: Our bit commitment protocol is ε -binding with $\varepsilon = (1 - \delta)^n$.

Proof: Let a_i^b be the distinguishing effect for μ_i^b ($b \in \{0, 1\}, i \in \{1, \dots, N^b\}$). Define $q_k^b(\lambda) := a_{x_k^{\lambda, b}}^b(\omega_k^\lambda)$; this is the probability that ω_k^λ passes the test Bob performs on the k -th system in the reveal phase when Alice tries to reveal b . Then $q_k^0(\lambda) + q_k^1(\lambda) \leq 2 - \delta$, by our choice of δ .

Since Bob only accepts if he accepts the state ω_k^λ of each subsystem, we have:

$$p_0 + p_1 = \sum_\lambda p(\lambda) \left[\prod_{k=1}^n q_k^0(\lambda) + \prod_{k=1}^n q_k^1(\lambda) \right]. \quad (7)$$

By convexity, we can fix some best choice for the randomness λ and drop the label. An upper bound on

$$p_0 + p_1 = \prod_{k=1}^n q_k^0 + \prod_{k=1}^n q_k^1, \quad (8)$$

is obtained by maximizing it subject to $0 \leq q_k^0, q_k^1 \leq 1$ and $q_k^0 + q_k^1 \leq 2 - \delta$. We should saturate the second inequality, since adding to q_k^0 or q_k^1 can only increase the right-hand side of Eq. (8). Now let $Q_k^b := \prod_{k \neq k} q_k^b$, so that $p_0 + p_1 = Q_k^0 q_k^0 + Q_k^1 q_k^1$. Since this expression is affine in q_k^0 , it's clear that if $Q_k^0 > Q_k^1$, we should take $q_k^0 = 1$ and $q_k^1 = 1 - \delta$, and vice versa if $Q_k^1 > Q_k^0$. If $Q_k^1 = Q_k^0$, then we can take either $q_k^0 = 1$ and $q_k^1 = 1 - \delta$ or use the opposite assignment. Therefore,

$$p_0 + p_1 \leq \max_{m=0, \dots, \lfloor n/2 \rfloor} (1 - \delta)^m + (1 - \delta)^{n-m}. \quad (9)$$

If $0 < m < n/2$, then we can increase the sum by moving a $1 - \delta$ term from $(1 - \delta)^m$ to $(1 - \delta)^{n-m}$, from which it follows that

$$p_0 + p_1 \leq \begin{cases} 1 + (1 - \delta)^n & \text{if } n \text{ is odd;} \\ \max(1 + (1 - \delta)^n, 2(1 - \delta)^{n/2}) & \text{if } n \text{ is even.} \end{cases}$$

For even n , note that $1 + (1 - \delta)^n - 2(1 - \delta)^{n/2} = (1 - (1 - \delta)^{n/2})^2 \geq 0$, so the maximum is always achieved by the first term. This proves the theorem. \square

VI. RELATED WORK

Winter, Nascimento, and Imai [20] found the optimal rate at which a discrete memoryless classical channel from Alice to Bob can be used to commit bits. Because the set of achievable output distributions may be a nonsimplicial compact convex body Ω , and the channel allows Alice to prepare any distribution of products of states in this convex body, their setting has similarities with ours. But it permits only a fixed output measurement whereas ours permits any measurement of effects in the cone dual to this convex body. Our setting also differs by permitting unentangled nonclassical processing by Alice and Bob. Also, the discreteness of the classical channel implies that the set of possible output distributions for the channel is a polytope, whereas in our theories Ω can be an arbitrary compact convex body. Finally, we do not calculate rates, but demonstrate exponentially secure commitment of a single bit; bounding the rate in our theories would be interesting, but it is not obvious what good analogues of the bounding entropic expressions in [20] would be.

Wolf and Wullschleger (WW) [21] reach a conclusion qualitatively similar to ours, that in a setting more general than quantum theory, assumptions that rule out entanglement can provide a secure protocol. They have told us that their result will be strengthened in [22]. [21] assumes Alice and Bob have access to many independent uses of the same trusted bipartite box-pair, initially uncorrelated with anything else. The boxes have binary inputs and outputs, but WW state that extension to larger finite sets of inputs and outputs is straightforward. Under the very weak condition that one party's conditional state depends on the other's input, they provide a bit commitment protocol and a security proof. Our setting is more general as it does not assume a trusted joint Alice-Bob state.

VII. CONCLUSION AND DISCUSSION

In [13], [14], [23], it was shown that the no-broadcasting and no-cloning theorems, and the tradeoff between information gain and state disturbance, are generic in non-classical theories in our framework. For the project of characterizing quantum mechanics this focuses attention on properties, like the impossibility of bit commitment and the possibility of teleportation, that may *not* be generically non-classical.

Within our framework, if one makes the plausible assumption that an information-disturbance tradeoff (which is equivalent to nonclassicality) allows secure key distribution, we may paraphrase the Brassard-Fuchs conjecture as saying that the impossibility of bit commitment characterizes quantum mechanics from among the nonclassical theories in our framework. We have shown that nonclassical theories in which bit commitment is impossible must have entanglement, but in contrast to the situation for the C^* -algebraic framework, in the general framework that is very far from narrowing us down to quantum theory. An important open question, then, is

what, if any, sorts of theories in our framework that *do* have entanglement, nevertheless permit bit commitment.

ACKNOWLEDGMENTS

Part of this work was completed at the workshop ‘‘Operational probabilistic theories as foils to quantum theory’’, July 2-13 2007 at the University of Cambridge, funded by The Foundational Questions Institute (FQXi) and SECOQC. At IQC, ML was supported in part by MITACS and ORDCF. ML and OD were supported in part by grant RFP1-06-006 from FQXi. Research at Perimeter Institute for Theoretical Physics is supported in part by the Government of Canada through NSERC and by the Province of Ontario through MRI. BT is supported by the EU FP6-FET Integrated Project QAP CT-015848, NWO VICI project 639-023-302, and the Dutch BSIK/BRICKS project. HB was supported by the US Department of Energy through the LDRD program at LANL.

REFERENCES

- [1] C. H. Bennett and G. Brassard, in Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing, Bangalore, India. IEEE Press, 1985, pp. 175–179.
- [2] G. Brassard, C. Crepeau, R. Jozsa, and D. Langlois, *Proceedings of the 34th IEEE Symposium on the Foundations of Computer Science (FOCS)*, pp. 42–52, 1993.
- [3] H.-K. Lo and H. F. Chau, *Phys. Rev. Lett.*, vol. 78, pp. 3410–3413, 1997.
- [4] D. Mayers, *Physical Review Letters*, vol. 78, p. 3413, 1997.
- [5] G. M. D’Ariano, D. Kretschmann, D. Schlingemann, and R. Werner, *Physical Review A*, vol. 76, p. 032328, 2007.
- [6] G. Brassard, *Nature Physics*, vol. 1, pp. 2–4, 2005.
- [7] C. A. Fuchs, *J. Mod. Opt.*, vol. 50, p. 987, 2003.
- [8] R. Clifton, J. Bub, and H. Halvorson, *Found. Phys.*, vol. 33, pp. 1561–1591, 2003, arXiv.org e-print quant-ph/0211089.
- [9] A. Grinbaum, 2007, to *British Journal for the Philosophy of Science*, vol. 58, pp. 387–408, 2007.
- [10] H. Halvorson, *Studies in History and Philosophy of Modern Physics*, vol. 35, pp. 277–293, 2004.
- [11] S. Popescu and D. Rohrlich, *Found. Phys.*, vol. 24, pp. 379–385, 1994.
- [12] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts, *Phys. Rev. A*, vol. 71, p. 022101, 2005.
- [13] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, 2007, arXiv.org e-print quant-ph/0611295.
- [14] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, *Phys. Rev. Lett.*, vol. 99, 240501, 2007.
- [15] H. Burhman, M. Christandl, F. Unger, S. Wehner and A. Winter, 2005, arXiv.org e-print quant-ph/0504133.
- [16] T. Short, N. Gisin and S. Popescu, *Quant. Inf. Proc.*, vol. 5, pp. 131-138, 2006
- [17] R. T. Rockafellar, *Convex Analysis*. Princeton: Princeton University Press, 1970.
- [18] H. Burhman, M. Christandl, P. Hayden, H.-K. Lo, and S. Wehner, *Proc. Roy. Soc. A*, vol. 462, 2071, pp. 1919-1932.
- [19] I. B. Damg ard, S. Fehr, R. Renner, L. Salvail, and C. Schaffner, in *Advances in Cryptology (CRYPTO 2007)*, LNCS 4622, pp. 360–378. Berlin: Springer, 2007.
- [20] A. Winter, A. C. A. Nascimento, and H. Imai, in *Proceedings of the 9th Cirencester Crypto and Coding Conference*, LNCS. Berlin: Springer, 2003.
- [21] S. Wolf and J. Wullschleger, *IEEE Information Theory Workshop on Theory and Practice in Information-Theoretic Security (Awaji Island, October 2005)*, 2005, e-print arXiv:quant-ph/0508233.
- [22] S. Winkler, J. Wullschleger, S. Wolf, in preparation.
- [23] J. Barrett, 2005, arXiv.org e-print quant-ph/0508211. To appear in *Physical Review A*.