

Superselection rules and quantum protocols

Alexei Kitaev,¹ Dominic Mayers,^{1,2} and John Preskill¹

¹*Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA*

²*Département de Mathématiques et d'Informatique, Université de Sherbrooke, Québec, Canada*

(Received 19 October 2003; published 21 May 2004)

We show that superselection rules do not enhance the information-theoretic security of quantum cryptographic protocols. Our analysis employs two quite different methods. The first method uses the concept of a *reference system*—in a world subject to a superselection rule, unrestricted operations can be simulated by parties who share access to a reference system with suitable properties. By this method, we prove that if an n -party protocol is secure in a world subject to a superselection rule, then the security is maintained even if the superselection rule is relaxed. However, the proof applies only to a limited class of superselection rules, those in which the superselection sectors are labeled by unitary irreducible representations of a compact symmetry group. The second method uses the concept of the *format* of a message sent between parties—by verifying the format, the recipient of a message can check whether the message could have been sent by a party who performed charge-conserving operations. By this method, we prove that protocols subject to general superselection rules (including those pertaining to non-Abelian anyons in two dimensions) are no more secure than protocols in the unrestricted world. However, the proof applies only to two-party protocols. Our results show in particular that, if no assumptions are made about the computational power of the cheater, then secure quantum bit commitment and strong quantum coin flipping with arbitrarily small bias are impossible in a world subject to superselection rules.

DOI: 10.1103/PhysRevA.69.052326

PACS number(s): 03.67.Dd

I. INTRODUCTION

The central aim of modern cryptography is to formulate protocols that achieve cryptographic tasks with *computational security*, meaning that a dishonest party would need to perform a prohibitively difficult computation to break the protocol. A major goal of quantum cryptography is to formulate protocols, involving the exchange of quantum states, that achieve *information-theoretic security*, meaning that even an adversary with unlimited computational power would be unable to defeat the protocol [1]. Information-theoretic security (sometimes called “unconditional security”) has been established for quantum key distribution protocols [2–7] but it has also been shown that, even in the quantum world, information-theoretic security is not attainable for certain tasks. For example, unconditionally secure quantum bit commitment is impossible [8,9], as is (strong) quantum coin flipping with arbitrarily small bias [10,11].

Superselection rules are limitations on the physically realizable quantum operations that can be carried out by a local agent. For example, it is impossible to create or destroy an isolated particle that carries locally conserved charges, such as an electrically charged particle, a fermion, or (in a two-dimensional medium) an anyon. Recently, Popescu [12] has suggested that superselection rules might have interesting implications for the security of quantum cryptographic protocols. The intuitive idea behind this suggestion is that superselection rules could place inviolable limits on the cheating strategies available to the dishonest parties, thus enhancing security. Might, say, unconditionally secure bit commitment be possible in worlds (perhaps including the physical world that we inhabit) governed by suitable superselection rules? An affirmative answer could shake the foundations of cryptography.

The purpose of this paper is to answer Popescu’s intriguing question. Sadly, our conclusion is that superselection rules can never foil a cheater who has unlimited quantum-computational power.

In the case of quantum bit commitment, and other two-party protocols, our argument hinges on a quite simple observation. In a two-party protocol, one participant (Alice) has control of a local system A , and the other participant (Bob) has control of another local system B . In addition, there is a message system M that they pass back and forth. In each step of the protocol, one party performs a joint quantum operation on her/his local system and the message system, and then sends the message system to the other party. Suppose that in each step, any part of the full system ABM that is beyond Alice’s control is under Bob’s control and vice versa—no part of the full system is inaccessible or in the possession of a third party. Suppose further that the full system ABM has trivial total charge (belongs to the trivial superselection sector). Then at any stage of the protocol, the algebra of operations that Alice can perform is the *commutant* of the algebra of operations that Bob can perform; that is, Alice’s algebra contains *all* operations that commute with Bob’s algebra. Likewise, Bob’s algebra is the commutant of Alice’s. By a minor extension of the standard argument, it then follows that unconditionally secure quantum bit commitment is impossible *if* the total charge shared by the parties is trivial.

Now, if the total charge in *nontrivial*, then Alice’s algebra is surely a subalgebra of the commutant of Bob’s, but it may be a *proper* subalgebra; similarly, Bob’s algebra may be a proper subalgebra of Alice’s. This unusual property of the local operations seems to open new possibilities for the design of quantum protocols. Regrettably, though, there is no way for an honest party to ensure that the total charge is really nontrivial when the other party is dishonest. Though

the honest protocol may call for the parties to start out with nontrivial charges, we may always imagine that there are actually compensating charges beyond the grasp of Alice and Bob, so that the total charge of the world is really trivial. Furthermore, a cheater might seize control of the compensating charge, while for an honest party it makes no difference whether the compensating charge is present or not. It follows that a protocol that calls for the total charge to be nontrivial can be no more secure than one in which the total charge is actually trivial; we conclude again that unconditionally secure quantum bit commitment is impossible, irrespective of the value of the total charge shared by the parties in the honest protocol.

Aside from quantum bit commitment, we will also study the impact of superselection rules on the information-theoretic security of a broad class of other quantum protocols, using two different methods. We analyze in detail the important special case in which the superselection sectors can be identified with the unitary irreducible representations of a compact symmetry group. In that case, we argue that it is possible in principle to prepare a *reference state* that establishes a preferred orientation in the symmetry group. A party with access to the reference state can use it to perform operations that are ostensibly forbidden by the superselection rule. In particular, consider an n -party quantum protocol where up to $k < n$ of the parties are dishonest, and suppose that in a world with no superselection rules the dishonest parties have a cheating strategy that breaks the protocol. Then, even in a world with superselection rules, the dishonest parties, by sharing a suitable reference state, can simulate this cheating strategy faithfully. We conclude that if a quantum protocol is information-theoretically secure in a world with a superselection rule, the security will be maintained even if the superselection rule is relaxed, at least in the case where the superselection rule arises from a compact symmetry group.

Superselection rules arising from compact symmetry groups are not the most general possible ones. In particular, an especially rich variety of superselection rules are potentially realizable in two-dimensional systems such as those that admit non-Abelian anyons. However, even superselection rules of this more general kind cannot foil a cheater. We find that for any two-party protocol that is secure in a world subject to a superselection rule, the security is maintained when the superselection rule is relaxed.

Our analysis of these more general superselection rules does not rely on the concept of a reference system; rather it is founded on a completely different idea, the concept of the *format* of a message. A superselection rule can always be characterized by saying that there are charges that must be conserved by all local operations, and when we relax the superselection rule, in effect we are permitting a cheater to violate these conservation laws. For the purpose of assessing the security of a two-party protocol, we are interested in how the actions of the cheating party (Alice) affect the outcomes of measurements performed by the honest party (Bob). Potentially, if Alice is granted the power to violate conservation of “charge,” her ability to influence Bob’s measurements will be strengthened.

However, if the total charge shared by Alice and Bob is trivial (as we are entitled to assume in an analysis of secu-

urity), then if charge is conserved, Alice and Bob hold conjugate charges at each stage of the protocol. Therefore, Bob always knows what charge Alice is supposed to have, which constrains the type of message that Alice can send to Bob if she is honest. When Bob receives a message he can verify its format, checking whether the message could have been sent by a party who performed a charge-conserving operation, and he can abort the protocol if the verification fails. Therefore, if the protocol ends normally, Alice has been forced to respect charge conservation—her power to flout the superselection rule does not enhance her ability to fool Bob. This reasoning shows that superselection rules cannot thwart cheating, but because the argument relies on the property that Alice and Bob hold perfectly correlated charges, it works only for two-party protocols.

For cryptographic protocols with more than two parties, and for general superselection rules, new subtleties arise. In two spatial dimensions, general charges are not merely locally conserved, they may also have nontrivial *braiding* properties—the exchange of two charges may induce a nontrivial transformation on their joint Hilbert space. This means that the effect of sending a message from one party to another can depend on the path along which the message travels. It is an interesting problem to specify appropriate definitions of security for protocols in this setting, but we will not attempt to address this issue here. For the special case of charges labeled by unitary representations of compact groups, the braiding properties are trivial; therefore in that case we can analyze multiparty protocols without confronting such questions.

Verstraete and Cirac [13] recently discussed a data-hiding protocol whose security is premised on a superselection rule. However, as the authors recognized, the protocol is not unconditionally secure; it can be broken if the parties establish a suitable shared reference state via quantum communication. The notion that the naive implications of a superselection rule can be evaded through the use of a suitable reference system was emphasized long ago by Aharonov and Susskind [14]; see [15] for a recent discussion. A special case of our main result was reported earlier in [16].

The rest of this paper is organized as follows: We develop the concept of a reference system in Sec. II, first for Abelian, then for non-Abelian symmetries, and we explain how a reference system can be used to simulate unrestricted operations in a world subject to superselection rules arising from a symmetry group; this observation is applied in Sec. III to the analysis of the security of quantum protocols. In Sec. IV we explore the distinction between an *itinerant* reference system that is passed from party to party as needed during a protocol, and a *distributed* reference system that can be prepared and passed out to the parties before the protocol begins. Superselection rules arising from non-Abelian symmetries are further characterized in Sec. V, and we comment in Sec. VI on the data-hiding protocol of Verstraete and Cirac. Our analysis of the impact of superselection rules on the security of quantum bit commitment is in Sec. VII; we also show there that for the analysis of security of an n -party protocol, it suffices to consider the case in which the total charge held by the parties is trivial. Two-party protocols subject to

general superselection rules are investigated in Sec. VIII, and Sec. IX contains some concluding comments.

II. SUPERSELECTION RULES AND REFERENCE SYSTEMS

A superselection rule is a decomposition of Hilbert space into sectors that are preserved by local operations. The different sectors can be distinguished by attaching to each sector a label, which we refer to as the sector’s “charge.” Therefore, an equivalent way to characterize a superselection rule is to say that the charge is locally conserved. In the context of a cryptographic protocol, this means that when one of the parties (Alice, say) performs an operation, the charge in Alice’s laboratory is preserved.

An important special case arises if the Hilbert space \mathcal{H} transforms as a unitary representation of a compact group G , and the sectors are labeled by the irreducible representations of G . An equivalent way to describe the superselection rule in that case is to say that the allowed operations must commute with the action of G on \mathcal{H} . In fact, it has been shown by Doplicher and Roberts [17] that such superselection rules are almost the most general ones allowed under rather weak conditions that apply in particular to quantum field theories (without gravity) in three or more spatial dimensions. We say “almost” because there is an additional freedom to assign to a localized state an even or odd fermion number. This fermion number is more than just a conserved charge, because of the property that the wave function changes sign when two fermions are exchanged.

In two spatial dimensions, there is a richer classification of superselection rules, reflecting the exotic quantum numbers carried by pointlike non-Abelian anyons that occur in topological quantum field theories [18–20]. We will postpone further discussion of non-Abelian anyons until Sec. VIII, concentrating for now on the superselection rules associated with compact symmetry groups (and ignoring fermions).

An important example is the group $U(1)$ associated with conservation of the electric charge Q . An agent acting locally can create or annihilate pairs of particles that carry equal and opposite charges, but cannot change the total charge in her vicinity. In particular, this agent is unable to transform any eigenstate of Q into a coherent superposition of states with different charges, as emphasized by Wick, Wightman, and Wigner [21,22].

While we might readily accept that local creation of electric charge is physically impossible, other conservation laws impose superselection rules that do more violence to our intuition. Suppose, for example (in nonrelativistic quantum mechanics), that our agent’s actions are required to conserve the angular momentum \vec{J} locally. Are we to conclude that if the agent is presented with a spin- $\frac{1}{2}$ object polarized spin-up along the z axis, it is impossible for him to transform it to a coherent superposition of the spin-up and spin-down states? How are we to describe what happens when a magnetic field is turned on pointing in the x direction and the spin begins to precess? A partial resolution of this puzzle is attained by noting that the angular momentum of a classical magnet has

an uncertainty large compared to \hbar , so that conservation of angular momentum need not prevent the magnet from coherently exchanging $J_z = \hbar$ with the spin. But this explanation does not fully address how the existence of the classical magnet is itself compatible with the superselection rule.

Such issues were cogently discussed many years ago by Aharonov and Susskind [14]. They emphasized that even if the total angular momentum has a definite value (like zero), we can still speak sensibly of the *relative* orientation of two subsystems. Whenever an experimentalist observes the precession of a spin, it is implicit that a reference state has been established that in effect breaks the rotational symmetry, and that the precession is measured relative to this reference standard. Furthermore, Aharonov and Susskind [14] emphasized that just as conservation of angular momentum need not prevent us from measuring the relative angular orientation of two objects, so the charge superselection rule need not prevent us from measuring relative phases in superpositions of states of different charge.

A. Abelian case

Before we discuss the more general case in which the symmetry may be non-Abelian, it will be useful to consider the symmetry group $G=U(1)$. Then the charge operator Q (the generator of G) has eigenvalues $q \in \mathbb{Z}$, and we denote the corresponding orthonormal eigenstates by $|q\rangle$. Formal states of definite phase (with continuum normalization) can be constructed as

$$|\theta\rangle = \frac{1}{\sqrt{2\pi}} \sum_{q=-\infty}^{\infty} e^{-iq\theta} |q\rangle \quad (0 \leq \theta < 2\pi), \quad (1)$$

where

$$\langle \theta' | \theta \rangle = \frac{1}{2\pi} \sum_{q=-\infty}^{\infty} e^{-iq(\theta-\theta')} = \delta(\theta' - \theta) \quad (2)$$

and

$$|q\rangle = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} d\theta e^{iq\theta} |\theta\rangle. \quad (3)$$

The phase state $|\theta\rangle$ is the improper eigenstate with eigenvalue $e^{i\theta}$ of the unitary operator

$$U_+ = \sum_{q=-\infty}^{\infty} |q+1\rangle \langle q| \quad (4)$$

that increments the value of the charge by one unit. While the phase θ is physically unobservable due to the charge superselection rule, the relative phase of $\theta' - \theta$ of the two states $|\theta'\rangle$ and $|\theta\rangle$ commutes with the charge operator Q and so is measurable in principle. Indeed, the state

$$\int_0^{2\pi} d\theta' |\theta' + \theta'\rangle \otimes |\theta + \theta'\rangle = \sum_{q=-\infty}^{\infty} e^{-iq(\theta-\theta')} |-q\rangle \otimes |q\rangle \quad (5)$$

has a definite value of the relative phase $\theta' - \theta$ and total charge zero. That is, it is an (unnormalizable) eigenstate with

eigenvalue $e^{i(\theta-\theta')}$ of the charge-conserving operator $U_- \otimes U_+$, where $U_- = U_+^\dagger$.

Similarly, the phases ϕ_q appearing in the expansion of the state $|\psi\rangle_A$ of a system A ,

$$|\psi\rangle_A = \sum_q \psi_q e^{-iq\phi_q} |q\rangle_A \quad (6)$$

(where the ψ_q 's are real and positive), are themselves unobservable, but they can be meaningfully compared to the phases appearing in the state $|\theta\rangle_R$ of a charge reservoir R . For example, by projecting $|\theta\rangle_R \otimes |\psi\rangle_A$ onto the sector with total charge zero, we obtain the state

$$\begin{aligned} |\psi\rangle_{RA} &= \frac{1}{\sqrt{2\pi}} \int d\theta' |\theta + \theta'\rangle_R \otimes e^{-iQ\theta'} |\psi\rangle_A \\ &= \sum_q \psi_q e^{-iq(\phi_q - \theta)} |-q\rangle_R \otimes |q\rangle_A \end{aligned} \quad (7)$$

which has measurable relative phases. A state like $|\theta\rangle_R$ of a charge reservoir R that provides a phase standard with which other states can be compared will be called a “reference state” or a “condensate.”

In the state $|\psi\rangle_{RA}$, the charge of the system A is compensated (“screened”) by the charge of the reservoir R . Therefore, the system and reservoir are entangled, and tracing out the reservoir destroys the coherence of the superposition of charge states for the system. While formally correct, this statement can be misleading if the reservoir remains accessible and is allowed to interact with the system during subsequent operations. For example, the operator $(U_+)_A$ that increases the charge of the system by one unit is disallowed by the superselection rule, but it can be accurately simulated by the allowed charge-conserving operator $(U_-)_R \otimes (U_+)_A$ acting on $|\psi\rangle_{RA}$ —this operator increases the charge of A by borrowing a unit of charge from R . If the reservoir remains accessible at all times, then an arbitrary (not necessarily charge-conserving) operation acting on A can be perfectly simulated by a charge-conserving operation acting on RA . Thus, at least as a matter of principle, the charge superselection rule places no inescapable restrictions on the allowed operations. This is the main point stressed by Aharonov and Susskind [14].

The phase reference state can be interpreted physically as a static piece of superconducting material with a definite value of the superconducting phase. While the phase itself is not gauge-invariant, the relative phase of the system and reservoir has observable consequences (like the Josephson effect) when the two are brought into contact. Similar issues, discussed in [23–27], arise when considering the physical content of relative phases in optical systems.

B. Non-Abelian case

Our discussion of the Abelian case has suggested that superselection rules are nullified if suitable reference systems are available. Now we consider the more general case, where the symmetry group is G , which may be either a finite group or a compact Lie group. The superselection rule dictates that allowed local operations must commute with G . But we may anticipate that if a condensate is accessible that completely

breaks the G symmetry, then in effect there is no operative symmetry at all, and the superselection rules place no restrictions on the allowed operations.

Formally, if the symmetry is completely broken, then the possible orientations of the condensate are in one-to-one correspondence with the elements of the symmetry group G . In a particular “fixed gauge,” the states of the condensate are denoted $|\phi\rangle$, where $\phi \in G$, and these states transform as the left regular representation of G . That is, a symmetry transformation $g \in G$ acting on the condensate is represented by the unitary $U(g)$ where

$$U(g)|\phi\rangle \rightarrow |g\phi\rangle. \quad (8)$$

These states can be expanded in the basis of irreducible representations of G as

$$|\phi\rangle = \sum_{q,i,a} \sqrt{\frac{n_q}{n_G}} D_{ia}^q(\phi) |q,i,a\rangle, \quad (9)$$

where n_q denotes the dimension of the irreducible representation $D^q(\phi)$ and n_G is the order of G . Inverting the Fourier transform we obtain

$$|q,i,a\rangle = \sum_{\phi \in G} \sqrt{\frac{n_q}{n_G}} D_{ia}^{q*}(\phi) |\phi\rangle. \quad (10)$$

Note that in Eqs. (9) and (10) we have used notation appropriate for a finite group; in the case of a compact Lie group, the sum over $\phi \in G$ would be replaced by an integral with respect to an invariant measure on the group. The states $|q,i,a\rangle$ transform under G as

$$U(g)|q,i,a\rangle = \sum_j |q,j,a\rangle D_{ji}^q(g). \quad (11)$$

In keeping with standard physics terminology, we will refer to the index $i=1,2,\dots,n_q$ in $|q,i,a\rangle$ as the “color index,” and to the action Eq. (11) of $U(g)$ on this index as a “gauge transformation.” The index $a=1,2,\dots,n_q$, distinguishing the n_q copies of the representation D^q that occur in the decomposition of the regular representation, will be called the “flavor” index. The physical “ G -invariant” operations are those that commute with all gauge transformations—these preserve q and act nontrivially only on the flavor, not the color. Therefore, by including the color we have chosen a redundant description of the physical Hilbert space. This redundancy, while not absolutely necessary, is quite convenient, and in particular will be useful for our discussion in Sec. III of the security of quantum protocols.

In addition to the G gauge symmetry, there is also a group G of “global” transformations that commute with $U(g)$, under which the states $|\phi\rangle$ transform as the right regular representation of G ; the element h of the global group is represented by $V(h)$, where

$$V(h)|\phi\rangle = |\phi h^{-1}\rangle \quad (12)$$

and

$$V(h)|q, i, a\rangle = \sum_b |q, i, b\rangle D_{ba}^{q*}(h). \quad (13)$$

Thus the global transformations act on the flavor index a of the states in the $\{|q, i, a\rangle\}$ basis—unlike the gauge transformations, they act nontrivially on the physical states.

In more geometric terms, a condensate may be interpreted as an asymmetric classical rigid body that can be rotated either “actively” or “passively.” What we have called the color (gauge) rotation is a passive rotation that acts on the space-fixed axes—it does not change the actual orientation of the body but only changes our mathematical description of the orientation. In contrast, what we have called the flavor (global) rotation is an active rotation that acts on the body-fixed axes and alters the physical orientation. A flavor rotation is G -invariant in the sense that it commutes with color rotations, and so is a physical operation, allowed by the superselection rule.

In contrast to the flavor orientation, the color orientation of an isolated system A has no invariant meaning, as it is modified by a color rotation. However, the orientation of A relative to the condensate R does have meaning, and an operator that rotates the relative orientation admits an invariant description. Suppose, for example, that system A is itself a condensate in the state ϕ_A , while the state of R is ϕ_R . The relative orientation

$$\phi_{\bar{R}A} \equiv \phi_R^{-1} \phi_A \quad (14)$$

is invariant if a common color rotation

$$U(h)_{RA}: \phi_A \rightarrow h\phi_A, \quad \phi_R \rightarrow h\phi_R \quad (15)$$

is applied to both objects. The transformation $U(g)_{RA}^{\text{inv}}$ that changes the relative orientation according to

$$U(g)_{RA}^{\text{inv}}: \phi_{\bar{R}A} \rightarrow g\phi_{\bar{R}A} \quad (16)$$

has an invariant meaning and commutes with the color rotation $U(h)_{RA}$. We may interpret the invariant rotation as one that rotates A while R is “held fixed,” acting as

$$U(g)_{RA}^{\text{inv}}(|\phi_R\rangle \otimes |\phi_A\rangle) = |\phi_R\rangle \otimes |\phi_R g \phi_R^{-1} \phi_A\rangle, \quad (17)$$

or equivalently

$$U(g)_{RA}^{\text{inv}} = \sum_{\phi \in G} (|\phi\rangle\langle\phi|)_R \otimes U(\phi g \phi^{-1})_A. \quad (18)$$

If system A is not a reference system but rather an object transforming as the irreducible representation q of G , then $U(\phi g \phi^{-1})$ can be expanded as

$$U(g)_{RA}^{\text{inv}} = \sum_{\phi \in G} (|\phi\rangle\langle\phi|)_R \otimes \left(\sum_{i,j,a,b} |q, i\rangle D_{ia}^q(\phi) D_{ab}^q(g) D_{bj}^q(\phi^{-1}) \langle q, j| \right)_A. \quad (19)$$

More generally, any transformation

$$M_A: |q, i\rangle_A \rightarrow \sum_j |q, j\rangle_A M_{ji} \quad (20)$$

acting on the color degree of freedom can be simulated by the invariant operation

$$M_{RA}^{\text{inv}} = \sum_{\phi \in G} (|\phi\rangle\langle\phi|)_R \otimes \left(\sum_{i,j,a,b} |q, i\rangle D_{ia}^q(\phi) M_{ab} D_{bj}^q(\phi^{-1}) \langle q, j| \right)_A. \quad (21)$$

M_{RA}^{inv} has an invariant meaning because it transforms the color of A relative to the color of the reference system R ; in effect, the color rotation is simulated by converting the color index into a flavor index (depending on ϕ), on which M may act with impunity. For fixed ϕ , the simulation is achieved via the isomorphism

$$|q, a\rangle_A \rightarrow |q, \phi, a\rangle_{RA} \equiv |\phi\rangle_R \otimes \sum_j |q, j\rangle_A D_{ja}^q(\phi), \quad (22)$$

such that

$$M_{RA}^{\text{inv}} |q, \phi, a\rangle_{RA} = \sum_b |q, \phi, b\rangle_{RA} M_{ba}. \quad (23)$$

Furthermore, this isomorphism can be extended to operators M that change the value of q as well as rotating the color for fixed q ; the operator

$$M_A: |q, i\rangle_A \rightarrow \sum_{q', j} |q', j\rangle_A M_{ji}^{q'q} \quad (24)$$

is simulated by

$$M_{RA}^{\text{inv}} |q, \phi, a\rangle_{RA} = \sum_{q', b} |q', \phi, b\rangle_{RA} M_{ba}^{q'q}, \quad (25)$$

which generalizes the result

$$M_{RA}^{\text{inv}} (|\theta\rangle_R \otimes e^{-iq\theta} |q\rangle_A) = |\theta\rangle_R \otimes \sum_{q'} e^{-iq'\theta} |q'\rangle M^{q'q} \quad (26)$$

that we found in the case of $G=U(1)$.

C. Properties of the simulation

We will refer to the world in which all operations are required to commute with the action of the symmetry group G as the “invariant world” or “ I world,” and we refer to the world in which arbitrary operations are allowed as the “unrestricted world” or “ U world.” What we have observed in Eqs. (22) and (25) is that the physics of the U world can be faithfully reproduced in the I world, as long as a suitable reference system is at our disposal.

Let us restate the main conclusion in a more succinct notation: Suppose A is an arbitrary system that transforms as some representation of the group G , and let R be a “reference system” that transforms as the left regular representation of G . Let M be an arbitrary transformation acting on A . Then there is a corresponding transformation M^{inv} acting on R and A defined as

$$M^{\text{inv}} = \sum_{\phi \in G} (|\phi\rangle\langle\phi|)_R \otimes [U(\phi) M U(\phi)^{-1}]_A. \quad (27)$$

M^{inv} is an invariant operator whose action on RA simulates the action of M on A .

That is, the operators M^{inv} have the following easily verified properties:

(i) M^{inv} is G -invariant. Proof: From the transformation properties of R and A we have

$$\begin{aligned} [U(g) \otimes U(g)]M^{\text{inv}}[U(g)^{-1} \otimes U(g)^{-1}] \\ = \sum_{\phi \in G} (|g\phi\rangle\langle g\phi|) \otimes [U(g\phi)MU(g\phi)^{-1}] = M^{\text{inv}}, \end{aligned} \quad (28)$$

where in the last step we have reparametrized the sum by replacing $\phi \rightarrow g^{-1}\phi$.

(ii) Invariant operators on RA provide a representation of operators on A . Proof: We have

$$\begin{aligned} M_1^{\text{inv}}M_2^{\text{inv}} &= \sum_{\phi_1, \phi_2 \in G} (|\phi_1\rangle\langle\phi_1| |\phi_2\rangle\langle\phi_2|) \\ &\quad \otimes [U(\phi_1)M_1U(\phi_1^{-1})U(\phi_2)M_2U(\phi_2^{-1})] \\ &= \sum_{\phi \in G} (|\phi\rangle\langle\phi|) \otimes [U(\phi)M_1M_2U(\phi)^{-1}] \\ &= (M_1M_2)^{\text{inv}}. \end{aligned} \quad (29)$$

(iii) If M is G -invariant, then $M^{\text{inv}} = I_R \otimes M_A$. Proof: If $U(\phi)$ commutes with M for each ϕ , then

$$M^{\text{inv}} = \sum_{\phi \in G} (|\phi\rangle\langle\phi|) \otimes M = I \otimes M. \quad (30)$$

(iv) If ρ is invariant and $\text{tr}(\rho_R) = 1$, then

$$\text{tr} M^{\text{inv}}(\rho_R \otimes \rho) = \text{tr} M\rho. \quad (31)$$

Proof: If $U(\phi)$ commutes with ρ for each ϕ , then

$$\begin{aligned} \text{tr} M^{\text{inv}}(\rho_R \otimes \rho) &= \sum_{\phi \in G} \langle\phi|\rho_R|\phi\rangle \text{tr}[M U(\phi)^{-1}\rho U(\phi)] \\ &= \text{tr}(\rho_R)\text{tr}(M\rho) = \text{tr}(M\rho). \end{aligned} \quad (32)$$

The properties (i) and (iv) mean that as long as the state ρ of A is G -invariant, then by making use of a reference system, measurements in the U world can be faithfully simulated by measurements in the I world. That is, given an arbitrary measurement performed on A (with operation elements that are not necessarily G -invariant), there is an invariant measurement performed on RA (with G -invariant operation elements) that has the same probability distribution of outcomes. Furthermore, it follows from property (ii) that the physics of the U world can be faithfully reproduced in the I world even if the measurement is preceded by a series of unitary transformations—applying V^{inv} in the I world has the same effect as applying V in the U world. Property (iii) tells us that, as expected, the reference system R is superfluous if the U -world transformation acting on A is already G -invariant.

To derive these properties, we require that the reference system transform as the regular representation of G , but no condition is needed on the state ρ_R of the reference system.

Loosely speaking, the reference system is needed so that when a noninvariant operation acts on A , the change in the charge of A can be balanced by a compensating change in the charge of R . But if the state ρ of A is invariant, then only the charge-conserving part of M contributes to the expectation value $\text{tr}(M\rho)$ anyway. In the simulation of this charge-conserving part of M , the reference system is superfluous and its state irrelevant.

Note that if G is a Lie group rather than a finite group, then the regular representation is infinite-dimensional, and our formal arguments require R to be an infinite-dimensional system. How is the fidelity of the simulation affected if R is truncated to a finite-dimensional system? In fact, the fidelity will still be perfect if the charge remains bounded in the process to be simulated. Consider, for example, the case $G = U(1)$, for which Eq. (27) becomes, e.g.,

$$(|q-r\rangle\langle q|)^{\text{inv}} = \sum_{q'} (|q'+r\rangle\langle q'|)_R \otimes (|q-r\rangle\langle q|)_A; \quad (33)$$

in the I world, a process in which r units of charge are removed from A is simulated by adding the r units to R . Suppose we are assured that the total charge added to or removed from A will never exceed r units. Then we may choose the initial state of R to carry charge zero, and we can limit R to the $(2r+1)$ -dimensional space spanned by the states $|q_R\rangle$, $q_R = -r, -r+1, \dots, r-1, r$. This truncated reference system suffices because states with $|q_R| > r$ will never be accessed in the simulation anyway. A similar remark applies if G is an arbitrary compact Lie group.

III. REFERENCE SYSTEMS AND QUANTUM PROTOCOLS

We have concluded that in the presence of a suitable reference system, superselection rules place no inescapable restrictions on the allowed operations. We may anticipate, therefore, that a cryptographic protocol is secure in the invariant “ I world” (governed by the superselection rule) if and only if it is secure in the unrestricted “ U world.” If we faithfully adhere to the usual stringent principles of quantum cryptology and place no restrictions on the resources available to our adversaries, then we must admit the possibility that the dishonest parties could share access to a reference system during the execution of the protocol. For the case of superselection rules arising from compact symmetry groups, this observation suffices to answer Popescu’s question about the impact of superselection rules on the security of quantum protocols.

Let us now discuss this point in greater detail. To be explicit, consider at first a protocol involving two parties, Alice and Bob. Alice holds a private local system A that is beyond Bob’s control, and Bob holds a private local system B that is beyond Alice’s control. In addition, there is a message system M that they can pass back and forth. At the beginning of the protocol, they share a product state $\rho_A \otimes \rho_B \otimes \rho_M$. In each round of the protocol, one of the parties performs a joint quantum operation on her/his local system and the message, and then sends the message system to the other party. Finally, after all quantum communication is completed, both parties perform local measurements. (See Fig. 1.)




FIG. 1. A two-player quantum game. Alice and Bob have private systems, and a message system that they pass back and forth. At the end of the game, Alice and Bob measure their private systems.

For example, the goal of the protocol might be to flip an unbiased coin. In that case, the final measurement performed by each party has two possible outcomes, 0 or 1. If both parties follow the protocol, then both obtain the same outcome. Furthermore, the two outcomes are equiprobable. A coin-flipping protocol is *secure* if neither party, by departing from the protocol, can bias significantly the outcome of the other party’s measurement.

We say that a *strong* coin-flipping protocol has bias ϵ if neither party by cheating can force *either* outcome to occur with probability greater than $\frac{1}{2} + \epsilon$. In a *weak* coin-flipping protocol, Alice wins if the outcome is 0 and Bob wins if the outcome is 1, and we say that the bias is ϵ if neither can force a *win* with probability greater than $\frac{1}{2} + \epsilon$. (Thus, in a weak protocol with bias ϵ , a cheater might be able to *lose* on purpose with a probability exceeding $\frac{1}{2} + \epsilon$.) Note that the protocol might abort if cheating is detected; by “the probability of outcome 0” we mean the joint probability that the protocol does not abort and the outcome is 0. Kitaev [10,11] has shown that, if no superselection rules are imposed, then strong quantum coin flipping is impossible with bias $\epsilon < (1/\sqrt{2}) - \frac{1}{2} = 0.207$. Ambainis [28] has shown that a weak coin-flipping protocol with bias ϵ requires at least $\Omega[\log \log(1/\epsilon)]$ rounds of communication.

We are interested in whether these conclusions about coin-flipping in the U world remain valid in the I world. For a coin-flipping protocol in the I world, we may assume that the initial state shared by Alice and Bob is a tensor product of invariant states $\rho_A \otimes \rho_B \otimes \rho_M$. In the honest protocol, Alice and Bob take turns applying G -invariant operations to the system that they share, then measure invariant observables. In fact, without loss of generality, we may assume [8] that each operation applied by Alice or Bob is an invariant unitary transformation, and that the final measurement is an invariant projective measurement.

If Alice and Bob play the game honestly, then the probability $P_B(b)$ that Bob’s measurement yields the particular outcome b can be expressed as

$$P_B(b) = \text{tr}[E_{B,b} V(\rho_A \otimes \rho_B \otimes \rho_M) V^\dagger], \quad (34)$$

where

$$V = V_{B_n} V_{A_n} \cdots V_{B_2} V_{A_2} V_{B_1} V_{A_1}. \quad (35)$$

Here the V_{A_j} are unitary transformations applied to AM (we have assumed that Alice makes the first move in the game), the V_{B_j} are unitary transformations applied to BM , and the $E_{B,b}$ are the projectors defining Bob’s final measurement. Furthermore, in the I -world protocol, V_{A_j} , V_{B_j} , and $E_{B,b}$ are

G -invariant. In effect, then, Bob measures the invariant operator

$$F_{B,b} = V^\dagger E_{B,b} V \quad (36)$$

in the invariant state $\rho_A \otimes \rho_B \otimes \rho_M$.

Of course, a protocol in the I world can be regarded as a special case of a protocol in the U world, where the initial state is a product state, and Kitaev’s result applies to this U -world protocol. Therefore, one of the parties (Alice, say) can force one of the outcomes (0, say) with probability at least $1/\sqrt{2}$. However, Alice’s cheating strategy that achieves this result might employ operations that are not G -invariant. To show that Kitaev’s result also applies to the original I -world protocol, we must show that Alice’s cheating strategy in the U world can be faithfully simulated in the I world by making use of a suitable reference system. For this purpose, we apply the properties of the invariant operator M^{inv} that were discussed in Sec. II C.

When Alice cheats in the U world, she replaces the operator V_{A_j} called for in the honest protocol with an arbitrary operator V'_{A_j} applied to AM , where V'_{A_j} is not necessarily G -invariant. Then Bob’s measurement yields the outcome b with probability

$$P'_B(b) = \text{tr}[F'_{B,b}(\rho_A \otimes \rho_B \otimes \rho_M)], \quad (37)$$

where

$$F'_{B,b} = V'^\dagger E_{B,b} V' \quad (38)$$

and

$$V' = V_{B_n} V'_{A_n} \cdots V_{B_2} V'_{A_2} V_{B_1} V'_{A_1}. \quad (39)$$

This cheating strategy in the U world can be simulated in the I world if Alice has a reference system R —instead of applying the noninvariant operator V'_{A_j} to the system AM , she applies the invariant operator $V'^{\text{inv}}_{A_j}$ to RAM . Note that since Bob follows the honest protocol, which requires V_{B_j} to be G -invariant, applying V_{B_j} to BM is equivalent to applying $V^{\text{inv}}_{B_j}$ to RBM , by property (iii) in Sec. II C. Therefore, when Alice adopts the I -world strategy, Bob obtains outcome b with probability

$$\tilde{P}'_B(b) = \text{tr}[\tilde{F}'_{B,b}(\rho_R \otimes \rho_A \otimes \rho_B \otimes \rho_M)], \quad (40)$$

where

$$\tilde{F}'_{B,b} = \tilde{V}'^\dagger E_{B,b} \tilde{V}' \quad (41)$$

and

$$\tilde{V}' = V^{\text{inv}}_{B_n} V'^{\text{inv}}_{A_n} \cdots V^{\text{inv}}_{B_2} V'^{\text{inv}}_{A_2} V^{\text{inv}}_{B_1} V'^{\text{inv}}_{A_1}. \quad (42)$$

But since the invariant operators provide a representation [property (ii)], we may write $\tilde{V}' = V'^{\text{inv}}$, and since $E_{B,b} = E^{\text{inv}}_{B,b}$ as well, we have

$$\tilde{F}'_{B,b} = F'_{B,b}{}^{\text{inv}}. \quad (43)$$

Finally, the initial state $\rho_A \otimes \rho_B \otimes \rho_M$ shared by Alice and Bob is G -invariant; therefore, by property (iv),

$$\tilde{P}'_B(b) = P'_B(b); \quad (44)$$

the measurement outcome b in the I -world protocol occurs with the same probability as the outcome b in the U -world protocol.

Therefore, Alice's simulated cheating strategy in the I world perfectly reproduces the probability distribution for Bob's measurement outcome that is achieved by her cheating strategy in the U world. The same is true if Bob makes the first move in the game instead of Alice. Similarly, if Bob is the cheater, Bob has a strategy in the I world that simulates his U world cheating strategy. We conclude that if Alice (or Bob) can cheat in the U world, then she (he) can cheat just as successfully in the I world. Thus, Kitaev's proof of the impossibility of strong coin flipping with bias $\epsilon < (1/\sqrt{2}) - \frac{1}{2}$, originally formulated in the U world, also applies to the I world. Similarly, Ambainis's lower bound on the number of rounds of communication needed for weak coin flipping also applies to the I world.

This conclusion that cheating in the U world can be successfully simulated in the I world applies not just to coin flipping protocols, but to any two-party protocol in which the goal of a cheating Alice is to bias the outcome of a measurement performed by an honest Bob. Furthermore, it is straightforward to generalize the argument to an n -party protocol, in which k cheating parties wish to bias the outcomes of measurements performed by the $n-k$ honest parties. For such a protocol in the I world, where the initial state is a product of invariant states, any cheating strategy that can be executed in the U world can be simulated perfectly in the I world if the k cheating parties share access to a reference system. Therefore, the protocol can be no more secure in the I world than in the U world.

To summarize, let us refer to an n -party quantum game as an I -world game if the initial state is a product of invariant states, and if in the honest protocol all operations performed by the parties are invariant operations. If $k < n$ parties are cheaters, we say that their cheating strategy is an I -world cheating strategy if the cheaters are required to perform invariant operations, and we say that their cheating strategy is a U -world cheating strategy if the operations performed by the cheaters are unrestricted. Let us say that an I -world cheating strategy is *equivalent* to a U -world cheating strategy if both strategies produce the same probability distributions for the outcomes of the measurements performed by the $n-k$ honest parties. We have proved the following.

Theorem 1. Suppose that in the I world all quantum operations are required to be G -invariant, where G is a compact Lie group, and that in the U world quantum operations are unrestricted. Consider an n -party I -world quantum game, and a U -world cheating strategy A' in which $k < n$ parties cheat. Then there is an I -world cheating strategy \tilde{A}' that is equivalent to A' .

As we observed in Sec. II C, the reference system required by the cheaters in the I world can be finite-dimensional, as long as the cheaters in the U world apply operations that change the "charge" by a bounded amount.

IV. DISTRIBUTED REFERENCE SYSTEMS

The key ingredient in our discussion of I -world quantum protocols is the observation that G -noninvariant operations can be faithfully simulated through the use of a reference system. Suppose, for example, that Alice and Bob take turns acting on a system C that they pass back and forth. Then Alice and Bob in the I -world can simulate an arbitrary U -world protocol in which the initial state of C is G -invariant. They carry out the simulation by passing the reference system R back and forth along with C , each taking turns applying invariant operations to RC . Similarly, in our analysis of cheating in Sec. III, we allowed the k cheaters to pass the reference system R among themselves as needed during the execution of the protocol. A reference system that travels from place to place might be called *itinerant*.

Here we will briefly discuss an alternative scenario, in which the parties share a *distributed* reference system—each party holds a fixed portion of this system throughout the execution of the protocol. This discussion is not actually needed for our analysis of security, but it is helpful nonetheless for understanding the physics of superselection rules. Indeed, in many physical situations in which reference systems are used (e.g., in optical physics), the system is distributed rather than itinerant.

Let A denote Alice's part of the reference system, B denote Bob's part, and suppose that at the start of the protocol AB is prepared in the state

$$|0\rangle_{AB} = \frac{1}{\sqrt{n_G}} \sum_{\phi \in G} |\phi\rangle_A \otimes |\phi\rangle_B. \quad (45)$$

This state has trivial total charge; indeed, when expressed in the Fourier-transformed charge-eigenstate basis, it is

$$|0\rangle_{AB} = \frac{1}{\sqrt{n_G}} \sum_{q,i,a} |\bar{q},i,a\rangle_A \otimes |q,i,a\rangle_B. \quad (46)$$

Thus, in principle Alice (say) could prepare $|0\rangle_{AB}$ in her lab and then ship half of it to Bob. (The state $|0\rangle_{AB}$ is unnormalizable and unphysical if G is a Lie group. For now we will suppose that G is a finite group, but we will comment on the case of a Lie group below.)

In the state $|0\rangle_{AB}$, Alice's condensate and Bob's, have values that are distributed uniformly over the group G , but these values are locked together. Therefore, if $|\psi\rangle_C$ is any pure state of C , then M_{AC}^{inv} and M_{BC}^{inv} act on $|0\rangle_{AB} \otimes |\psi\rangle_C$ in the same way,

$$\begin{aligned} M_{AC}^{\text{inv}}(|0\rangle_{AB} \otimes |\psi\rangle_C) &= M_{BC}^{\text{inv}}(|0\rangle_{AB} \otimes |\psi\rangle_C) \\ &= \frac{1}{\sqrt{n_G}} \sum_{\phi \in G} |\phi\rangle_A \otimes |\phi\rangle_B \\ &\quad \otimes [U(\phi)MU(\phi)^{-1}]|\psi\rangle_C. \end{aligned} \quad (47)$$

Furthermore M_{AC}^{inv} and M_{BC}^{inv} act identically on any state of the form

$$|\Psi\rangle_{ABC} = \frac{1}{\sqrt{n_G}} \sum_{\phi \in G} |\phi\rangle_A \otimes |\phi\rangle_B \otimes |\psi_\phi\rangle_C, \quad (48)$$

where $|\psi_\phi\rangle_C$ might depend on ϕ , a form that is maintained as successive invariant operations are applied to AC and to BC . Therefore, the outcome of the protocol would be the same if each invariant operation M_{BC}^{inv} applied to BC were replaced by the corresponding invariant operation M_{AC}^{inv} applied to AC . We conclude that the simulation in which the distributed reference system AB is prepared in the initial state $|0\rangle_{AB}$ is equivalent to a simulation that uses an itinerant reference system A . Since this latter simulation has all of the properties listed in Sec. II C, we find that a bipartite I -world protocol using the distributed reference system can faithfully simulate an arbitrary U -world protocol.

Note that the distributed state can serve the same purpose if there is a fixed offset of Bob's condensate relative to Alice's, as long as the offset is known. That is, if Alice and Bob share the state

$$\begin{aligned} |0, \tilde{\phi}\rangle_{AB} &= \frac{1}{\sqrt{n_G}} \sum_{\phi \in G} |\phi\rangle_A \otimes |\phi \tilde{\phi}\rangle_B \\ &= \frac{1}{\sqrt{n_G}} \sum_{q,a,b} D_{ab}^q(\tilde{\phi}) \\ &\quad \times \left(\sum_i |\bar{q}, i, a\rangle_A \otimes |q, i, b\rangle_B \right), \end{aligned} \quad (49)$$

then the invariant operations M_{BC}^{inv} and $[U(\tilde{\phi})MU(\tilde{\phi})^{-1}]_{AC}^{\text{inv}}$ act in the same way. If Bob knows $\tilde{\phi}$, then he can participate successfully in the simulation by "twisting" his operations appropriately.

Similarly, in a protocol with k parties, the distributed reference state

$$|0\rangle_{k \text{ parties}} = \frac{1}{\sqrt{n_G}} \sum_{\phi \in G} |\phi\rangle_{R_1} \otimes |\phi\rangle_{R_2} \cdots \otimes |\phi\rangle_{R_k} \quad (50)$$

provides a common "phase standard" for all the participants, allowing them to simulate a U -world protocol in the I world—the ℓ th party simulates the noninvariant operation M by applying M^{inv} to the target system and her part R_ℓ of the reference system. Again, the parties can twist their local operations to compensate for known relative offsets of their condensates, if necessary.

In the state $|0\rangle_{AB}$, there is a quantum correlation between Alice's condensate and Bob's. A common reference standard can be provided instead by a classically correlated state such as

$$\rho_{AB} = \frac{1}{n_G} \sum_{\phi \in G} (|\phi\rangle\langle\phi|)_A \otimes (|\phi\rangle\langle\phi|)_B. \quad (51)$$

If Alice and Bob are equipped with the state ρ_{AB} , then again M_{AC}^{inv} and M_{BC}^{inv} act in the same way, hence they can use this distributed reference state to simulate a U -world protocol in

the I world. The state is G -invariant, but unlike $|0\rangle_{AB}$ it is not a charge eigenstate; rather it is a mixture of (invariant) states with various charges. For example, in the case $G=U(1)$, $|0\rangle_{AB}$ is the (unnormalizable) state

$$|0\rangle_{AB} = \int_0^{2\pi} |\theta\rangle_A \otimes |\theta\rangle_B = \sum_{q=-\infty}^{\infty} | -q\rangle_A \otimes |q\rangle_B; \quad (52)$$

Alice's charge and Bob's charge are perfectly anticorrelated. In contrast, ρ_{AB} is

$$\begin{aligned} \rho_{AB} &\propto \int d\theta (|\theta\rangle\langle\theta|)_A \otimes (|\theta\rangle\langle\theta|)_B \\ &\propto \sum_{q_A, q_B, q} |q_A, q_B\rangle\langle q_A - q, q_B + q|. \end{aligned} \quad (53)$$

Formally, this state appears to be separable, as it is a mixture of the product states $|\theta\rangle \otimes |\theta\rangle$, but this is deceptive, because $|\theta\rangle \otimes |\theta\rangle$ is not G -invariant and is therefore incompatible with the superselection rule. On the other hand, in the charge-eigenstate basis, ρ_{AB} can be expressed as a mixture of G -invariant pure states, each with a definite total charge; however, these pure states are highly entangled, with an indefinite value of Alice's (and Bob's) local charge. The state ρ_{AB} is not a mixture of invariant product states, and therefore cannot be prepared without quantum communication between Alice and Bob. Classical communication alone is insufficient for Alice and Bob to establish their common phase standard.

Now let us return to the question we postponed earlier: what if G is a Lie group, so that the states $|0\rangle_{AB}$ and ρ_{AB} are unnormalizable? To be specific, consider again the case $G=U(1)$, and suppose that Alice and Bob are instructed to perform this protocol: Alice is presented with a charge-zero state $|0\rangle$. She is instructed to rotate this state to the superposition of charge eigenstates $(|0\rangle + |1\rangle)/\sqrt{2}$ and to send the resulting state to Bob. Bob is to perform an orthogonal measurement in the basis $(|0\rangle \pm |1\rangle)/\sqrt{2}$ and so verify that Alice prepared the correct state. To make sense of this procedure, Alice and Bob must share a common reference state that serves to lock together their phase conventions; for example, this state could be a shared pure state $|\psi\rangle_{AB}$ with definite total charge. Alice's coherent operation on system C acts as

$$|\psi\rangle_{AB} \otimes |0\rangle_C \rightarrow \frac{1}{\sqrt{2}} [|\psi\rangle_{AB} \otimes |0\rangle_C + (U_-)_A |\psi\rangle_{AB} \otimes |1\rangle_C]; \quad (54)$$

that is, Alice simulates the charge-nonconserving operator $(U_+)_C$ by applying the invariant operator $(U_-)_A \otimes (U_+)_C$ to AC . When Bob receives system C , he performs his measurement by first simulating the transformation

$$|0\rangle_C \rightarrow \frac{1}{\sqrt{2}} (|0\rangle_C + |1\rangle_C), \quad (55)$$

$$|1\rangle_C \rightarrow \frac{1}{\sqrt{2}}(|0\rangle_C - |1\rangle_C),$$

and then measuring the charge of C . After Bob's first step, the state of ABC has become

$$\begin{aligned} & \frac{1}{2}[I_A \otimes I_B + (U_-)_A \otimes (U_+)_B]|\psi\rangle_{AB} \otimes |0\rangle_C \\ & + \frac{1}{2}[I_A \otimes (U_-)_B - (U_-)_A \otimes I_B]|\psi\rangle_{AB} \otimes |1\rangle_C. \end{aligned} \quad (56)$$

When Bob measures the charge, the probability that he obtains the outcome $|1\rangle_C$ and fails to verify Alice's state is

$$P_1 = \frac{1}{2}[1 - \text{Re}_{AB} \langle \psi | (U_-)_A \otimes (U_+)_B | \psi \rangle_{AB}]. \quad (57)$$

If, for example, the shared reference state is

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{N}} \left(\sum_{q=0}^{N-1} |-q\rangle_A \otimes |q\rangle_B \right), \quad (58)$$

a normalizable approximation to the state $|0\rangle_{AB}$, our expression for P_1 becomes

$$P_1 = \frac{1}{2N}. \quad (59)$$

Thus, for finite N , the state received by Bob does not match perfectly with the state prepared by Alice—the superposition of charge eigenstates decoheres slightly. But this decoherence becomes negligible in the limit $N \rightarrow \infty$, where the “charge fluctuations” of the shared condensate are large.

The lesson we learn from this example generalizes to non-Abelian compact Lie groups. We can replace the unnormalizable state

$$|0\rangle_{AB} = \frac{1}{\sqrt{n_G}} \sum_{q,i,a} |\bar{q}, i, a\rangle_A \otimes |q, i, a\rangle_B \quad (60)$$

by a normalizable state with a truncated sum over the charge q . If Alice and Bob use this truncated distributed reference state to simulate a U -world protocol, their simulation will not have perfect fidelity. But as long as all operations applied by Alice and Bob change the charge by a bounded amount, the fidelity can be arbitrarily close to 1 if the reference state is chosen appropriately. If Alice and Bob are permitted to use a truncated *itinerant* reference system rather than a distributed one, then perfect fidelity can be achieved, as observed in Sec. II C.

V. INVARIANT OPERATIONS AND COMMUTANTS

Our observations in Sec. II B emphasized the similarities between Abelian and non-Abelian superselection rules, enabling us to formulate a security analysis in Sec. III that applies to both Abelian and non-Abelian symmetry groups. But in several respects the arguments in Sec. III are still not adequate. For one thing, so far we have treated only the special case of superselection sectors labeled by unitary irre-

ducible representations of compact groups. For another, while it is possible to formulate a security analysis of quantum bit commitment within the framework of our argument in Sec. III, it is more natural to structure the argument differently, following more closely the standard analysis of quantum bit commitment.

In this section, we will emphasize the essential differences between superselection rules arising from non-Abelian symmetry groups and those arising from Abelian groups. The discussion will pave the way for our analysis of quantum bit commitment in Sec. VII and of general two-party protocols in Sec. VIII.

A crucial difference between Abelian and non-Abelian charges is that non-Abelian charges are nonadditive: the charges of two subsystems A and B do not necessarily determine the charge of the composite system AB . This feature can be restated as a property of the algebra of observables of the bipartite system. Let \mathcal{A} denote the algebra of local operators (an associative algebra, closed under Hermitian conjugation, that commutes with all locally conserved charges) acting on subsystem A , and let \mathcal{B} denote the algebra of local operators acting on B . The commutant of \mathcal{A} , denoted \mathcal{A}' , is the algebra of operators acting on the composite system AB that commute with everything in \mathcal{A} , and similarly for \mathcal{B}' . Now, if all superselection rules are Abelian, then $\mathcal{A}' = \mathcal{B}$ and $\mathcal{B}' = \mathcal{A}$. But if the superselection rules are non-Abelian, the theory has sectors with nontrivial total charge in which this relation does not hold. This unusual structure of the local observables has potential implications for the security of quantum protocols.

To be more explicit, suppose that the superselection rules arise from a non-Abelian symmetry group G , and the operations that Alice (or Bob) can perform must commute with G . A state $|\psi\rangle$ in Alice's (or Bob's) Hilbert space can be decomposed into irreducible representations of G , as

$$|\psi\rangle = \sum_{q,i,a} \psi_{i,a}^q |q, i, a\rangle; \quad (61)$$

here q labels the irreducible representation (or “charge”), i is the “color” index acted upon by the representation of G , and a is the “flavor” index that distinguishes among the various copies of the irreducible representation q appearing in the decomposition. Note that since we are no longer assuming that Alice's system transforms as the regular representation of G , there need be no connection between the number of flavors and the number of colors associated with q . The action of a color gauge rotation representing $g \in G$ on $|\psi\rangle$ is

$$U(g)|\psi\rangle = \sum_{q,i,j,a} \psi_{i,a}^q |q, j, a\rangle D_{ji}^q(g). \quad (62)$$

An operator M allowed by the superselection rule, which must commute with each $D^q(g)$, preserves the charge q and acts only on the flavor index according to

$$M|\psi\rangle = \sum_{q,i,a,b} \psi_{i,a}^q |q, i, b\rangle M_{ba}^q. \quad (63)$$

Since allowed operations act nontrivially only on the flavor index, it is convenient to use a notation that suppresses the color index i . We denote by \mathcal{H}_q the *invariant* Hilbert space in

the charge- q sector, spanned by states $|q, a\rangle$ that are labeled only by the flavor a within the sector. The corresponding operator algebra respecting the superselection rule is $\mathcal{L}(\mathcal{H}_q)$, spanned by linear operators acting on this invariant space. Thus Alice's invariant Hilbert space is

$$\mathcal{H}_A = \bigoplus_q \mathcal{H}_{A,q} \quad (64)$$

and Alice's local operator algebra is

$$\mathcal{A} = \bigoplus_q \mathcal{L}(\mathcal{H}_{A,q}). \quad (65)$$

Similarly, Bob's operator algebra is

$$\mathcal{B} = \bigoplus_q \mathcal{L}(\mathcal{H}_{B,q}). \quad (66)$$

Now consider the composite system AB . Its invariant Hilbert space too can be expressed as a direct sum over charge sectors

$$\mathcal{H} = \bigoplus_q \mathcal{H}_q, \quad (67)$$

while the full operator algebra is $\bigoplus_q \mathcal{L}(\mathcal{H}_q)$. But we should consider how \mathcal{H}_q is related to the invariant Hilbert spaces of the subsystems. The charge- q Hilbert space of the joint system can be expressed as

$$\mathcal{H}_q = \bigoplus_{q_A, q_B} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{B, q_B} \otimes V_q^{q_A, q_B}, \quad (68)$$

where $V_q^{q_A, q_B}$ denotes the space of invariant linear maps from the irreducible representation q to the tensor product of irreducible representations $q_A \otimes q_B$. This space can be nontrivial (of dimension greater than 1) if the tensor product contains the representation q more than once.

When expressed in terms of a particular color basis for the irreducible representations q , q_A , and q_B , the components of $V_q^{q_A, q_B}$ are the Clebsch-Gordon coefficients ($3j$ symbols), of the group G . Let $\{|q_A, i\rangle\}$ denote an orthonormal basis for the representation q_A , $\{|q_B, j\rangle\}$ a basis for q_B , and $\{|q(\alpha), k\rangle\}$ a basis for $q(\alpha)$, where the index α labels the various copies of the representation q that may be contained in $q_A \otimes q_B$. Then the components of $V_q^{q_A, q_B}$ are

$$[V_q^{q_A, q_B}(\alpha)]_k^{ij} = (\langle q_A, i | \otimes \langle q_B, j |) |q(\alpha), k\rangle. \quad (69)$$

These components comprise a G -invariant tensor with the property

$$[V_q^{q_A, q_B}(\alpha)]_k^{ij} = \sum_{i', j', k'} D_{ii'}^{q_A}(g) D_{jj'}^{q_B}(g) [V_q^{q_A, q_B}(\alpha)]_{k'}^{i' j'} D_{k' k}^q(g). \quad (70)$$

Invariant operations act not on the color indices of $[V_q^{q_A, q_B}(\alpha)]_k^{ij}$, but rather on the index α that distinguishes the flavors of q contained in $q_A \otimes q_B$. Furthermore, the invariant operations can also alter the charges q_A and q_B appearing in Eq. (68), while preserving the total charge q .

The notation of Eq. (68) and its implications may be clarified by discussing specific examples. The trivial representation ($q=1$) is contained only in the tensor product of q_A with

its conjugate representation \bar{q}_A , and it occurs only once in this product. Therefore, in the case where the total charge is $q=1$, Eq. (68) reduces to

$$\mathcal{H}_1 = \bigoplus_q \mathcal{H}_{A, q} \otimes \mathcal{H}_{B, \bar{q}}, \quad (71)$$

in this case, the factor $V_q^{q_A, q_B}$ is superfluous. Now, the joint operator algebra contains operations that cannot be executed by Alice and Bob locally—these operations change Alice's charge and Bob's while preserving the total charge (of course, this can happen even if G is Abelian). But any operation that commutes with Alice's algebra \mathcal{A} must preserve Alice's charge q , and act trivially in each of Alice's charge sectors; such operations preserve Bob's charge \bar{q} as well, and thus are in Bob's algebra \mathcal{B} . Therefore, \mathcal{A} and \mathcal{B} are commutants of one another.

However, if the total charge is nontrivial, then \mathcal{B} need not be the commutant of \mathcal{A} . To illustrate this phenomenon, consider the case $G=\text{SU}(2)$, where the irreducible representation is labeled by the spin j . For $\text{SU}(2)$, $V_j^{j_A, j_B}$ is always one- (or zero-) dimensional, and Eq. (68) reduces to

$$\mathcal{H}_j = \bigoplus_{j_A, j_B} \mathcal{H}_{A, j_A} \otimes \mathcal{H}_{B, j_B}, \quad (72)$$

where it is implicit that each product of representations appearing on the right-hand side transforms as spin j . To be concrete, suppose that Alice's system has spin $\frac{1}{2}$, Bob's contains both a spin-0 and a spin-1 component, and the total spin is $\frac{1}{2}$; then

$$\mathcal{H}_{1/2} = \mathcal{H}_{A, 1/2} \otimes (\mathcal{H}_{B, 0} \oplus \mathcal{H}_{B, 1}). \quad (73)$$

Note that in this case, contrary to the case in which the total charge is trivial, a single value of j_A can be combined with either of two different values of j_B to obtain the same total charge j . Therefore, there are invariant operations acting on the joint system that preserve Alice's charge and the total charge, but change Bob's charge. These operations are in the commutant of \mathcal{A} but not in \mathcal{B} ; hence $\mathcal{A}' \neq \mathcal{B}$.

We arrive at another way of looking at this property of $\mathcal{H}_{1/2}$ if we imagine that there is a third party Charlie who holds a compensating charge, so that the total charge is trivial. Now

$$\mathcal{H}_0 = \mathcal{H}_{A, 1/2} \otimes (\mathcal{H}_{B, 0} \otimes \mathcal{H}_{C, 1/2} \oplus \mathcal{H}_{B, 1} \otimes \mathcal{H}_{C, 1/2}); \quad (74)$$

an operation in \mathcal{A}' can be performed by Bob and Charlie acting together, but not by Bob alone.

In order that $\mathcal{A}' \neq \mathcal{B}$, it is not necessary for one of the parties to possess a state with indefinite charge. For example, in the case $G=\text{SU}(3)$, the tensor product of the irreducible octet representation 8 with itself contains two copies of 8, one symmetric and one antisymmetric under interchange of the factors,

$$8_A \otimes 8_B \supseteq 8_{\text{sym}} \oplus 8_{\text{anti}}. \quad (75)$$

Thus, in the decomposition

$$\mathcal{H}_8 = \mathcal{H}_{A,8} \otimes \mathcal{H}_{B,8} \otimes V_8^{8,8}, \quad (76)$$

the joint invariant Hilbert space is two-dimensional, while Alice and Bob both have one-dimensional Hilbert spaces and trivial invariant operator algebras. Then \mathcal{A}' is the full operator algebra, clearly different from \mathcal{B} , and similarly \mathcal{B}' is different from \mathcal{A} . Again, an alternative description of the invariant space is to note that Charlie could hold a compensating 8 charge, in which case the total charge is trivial and

$$\mathcal{H}_1 = (\mathcal{H}_{A,8} \otimes \mathcal{H}_{B,8} \otimes \mathcal{H}_{C,8}) \otimes V_1^{8,8,8} \quad (77)$$

is two-dimensional.

For the purpose of describing G -invariant operations, it is always legitimate to introduce a compensating charge without incurring any loss of generality. To see this, first note that if \mathcal{E} is a G -invariant quantum operation, then

$$\mathcal{E}[U(g)\rho U(g)^{-1}] = U(g)\mathcal{E}(\rho)U(g)^{-1} \quad (78)$$

for any $g \in G$ and any state ρ . In particular, then,

$$\mathcal{E}[\mathcal{G}(\rho)] = \mathcal{G}[\mathcal{E}(\rho)], \quad (79)$$

where \mathcal{G} is the map

$$\mathcal{G}(\rho) = \frac{1}{n_G} \sum_{g \in G} U(g)\rho U(g)^{-1}, \quad (80)$$

which induces decoherence of a superposition of distinct irreducible representations of G ,

$$\mathcal{G}(|q,i,a\rangle\langle q',j,b|) = \delta^{qq'} \delta_{ij} \left(\frac{1}{n_q} \sum_l |q,l,a\rangle\langle q,l,b| \right). \quad (81)$$

Equation (79) means [26] that the state

$$|\psi\rangle = \sum_{i,a} \psi_{i,a}^q |q,i,a\rangle \quad (82)$$

cannot be distinguished by any G -invariant operation from the state

$$\mathcal{G}(|\psi\rangle\langle\psi|) = \sum_{q,a,b,i} \psi_{i,a}^q \psi_{i,b}^{q*} \left(\frac{1}{n_q} \sum_j |q,j,a\rangle\langle q,j,b| \right). \quad (83)$$

Now, consider a system A whose charge is screened by a system C , so that the state of the joint system has trivial total charge,

$$|\psi\rangle_{AC} = \sum_{q,a,i} \psi_a^q |q,i,a\rangle_A \otimes |\bar{q},i\rangle_C. \quad (84)$$

Tracing over system C produces the state

$$\text{tr}_C(|\psi\rangle\langle\psi|)_{AC} = \sum_{q,a,b} \psi_a^q \psi_b^{q*} \left(\frac{1}{n_q} \sum_j |q,j,a\rangle\langle q,j,b| \right). \quad (85)$$

But the state Eq. (83) is just a convex combination of states of the form Eq. (85). Therefore, if only G -invariant operations are to be considered, it is always harmless to replace system A by half of a bipartite state that carries trivial total charge.

Up until now, we have explicitly discussed only the case of superselection sectors arising from a compact symmetry group, but much of the formalism we have outlined in this section can be extended to a more general setting. Whatever the origin of the superselection rule, the allowed operations act on a suitable invariant space. Sectors can still be classified by conserved charges, but in the general case, the space $V_q^{q_A, q_B}$ is defined more abstractly, rather than in terms of group representations. One important property that continues to hold in the general setting (which will play a central role in our analysis of quantum bit commitment in Sec. VII and of general two-party games in Sec. VIII) is that for each value q of the charge, there is a unique conjugate charge \bar{q} such that the fusion of the charges contains the trivial charge sector.

VI. DATA HIDING

Verstraete and Cirac [13] described a data-hiding protocol whose security is founded on the charge superselection rule for $G=U(1)$. Suppose that a trusted third party Charlie prepares one of the two orthogonal states,

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle), \quad (86)$$

where $|0\rangle$ and $|1\rangle$ denote states of charge 0 and 1, respectively, and distributes half to Alice and half to Bob. If Alice and Bob could each measure the Pauli operator X that interchanges $|0\rangle$ and $|1\rangle$, they could distinguish the states $|+\rangle$ and $|-\rangle$ by performing these measurements and comparing their outcomes. However, X does not commute with the electric charge Q ; if Alice and Bob are permitted only to perform local charge-conserving operations and to communicate classically, then they will be powerless to distinguish the two possible states.

On the other hand, if Alice and Bob share access to a common phase reference state, their activities will be unrestricted and nothing will prevent them from performing the X measurements that unlock the classical bit stored in the state prepared by Charlie [aside from the small loss of fidelity that arises if the reference state has large but finite charge fluctuations, as in Eq. (58)]. In Bloch sphere language, Alice and Bob have no *a priori* means of orienting their measurement axes in the x - y plane, but a shared phase standard enables them to lock their axes together and compare their measurements. Since the state prepared by Charlie is invariant under rotations about the z axis, the overall orientation in the x - y plane is irrelevant; only the relative orientation needs to be fixed to identify Charlie's state.

To be more explicit, while X does not commute with the charge,

$$X_{AA'}^{\text{inv}} = (U_-)_A \otimes \sigma_{A'}^+ + (U_+)_A \otimes \sigma_{A'}^-, \quad (87)$$

commutes with Q , as does $X_{BB'}^{\text{inv}}$. If Alice and Bob share a distributed reference state $|\psi\rangle_{AB}$ that is an eigenstate of $(U_-)_A \otimes (U_+)_B$ with eigenvalue 1, then

$$|\psi\rangle_{AB} \otimes |\pm\rangle_{A'B'} \quad (88)$$

is an eigenstate of

$$X_{AA'}^{\text{inv}} \otimes X_{BB'}^{\text{inv}} \quad (89)$$

with eigenvalue ± 1 . Therefore, Alice and Bob can unlock the hidden bit by each measuring X^{inv} and comparing their results. The same holds, of course, if the shared reference state ρ_{AB} is a mixture of eigenstates of $(U_-)_A \otimes (U_+)_B$, each with eigenvalue 1, as in Eq. (53). As Verstraete and Cirac observed [13], quantum communication is needed to establish this shared phase standard.

In the absence of a shared phase standard, neither Alice nor Bob can detect the bit encoded in the state $|\pm\rangle$ of Eq. (86); however, either Alice or Bob can manipulate the bit. Each can measure the charge q , and either can apply a phase to the state conditioned on the charge, flipping $|+\rangle \leftrightarrow |-\rangle$. But the property that $B' \neq A$ indicates that the situation can be more subtle in the non-Abelian case (with nontrivial total charge). Suppose, for example, that $G = \text{SU}(2)$ with total charge $j = \frac{1}{2}$ as in Eq. (73). Two states with the same value of the total charge and of Alice's charge, but different values of Bob's charge, are $|j = \frac{1}{2}, j_A = \frac{1}{2}, j_B = 0\rangle$ and $|j = \frac{1}{2}, j_A = \frac{1}{2}, j_B = 1\rangle$. Charlie might prepare either of the linear combinations

$$|\pm\rangle = \frac{1}{\sqrt{2}} \left(\left| j = \frac{1}{2}, j_A = \frac{1}{2}, j_B = 0 \right\rangle \pm \left| j = \frac{1}{2}, j_A = \frac{1}{2}, j_B = 1 \right\rangle \right), \quad (90)$$

and then distribute the AB system to Alice and Bob. Again, neither Alice nor Bob can detect the hidden bit, but now there is a notable asymmetry between Alice's power and Bob's. Since Bob has a superposition of two different charge states, he can tamper with the hidden bit by applying a phase controlled by the charge. Alice, on the other hand, has a trivial invariant operator algebra, and has no control over the shared state.

We may take this observation a step further. Suppose, for example, that $G = \text{SU}(3)$ with total charge $q = 8$ as in Eq. (76). Charlie might prepare either of the linear combinations

$$|\pm\rangle = \frac{1}{\sqrt{2}} (|q = 8_{\text{sym}}, q_A = 8, q_B = 8\rangle \pm |q = 8_{\text{anti}}, q_A = 8, q_B = 8\rangle), \quad (91)$$

and then distribute the AB system to Alice and Bob. Again, neither Alice nor Bob can detect the hidden bit, but furthermore, neither one can tamper with the bit's value.

However, in the non-Abelian case as in the Abelian case, the hidden bit can be opened via local operations and classical communication between Alice and Bob if they are provided with correlated reference systems that effectively remove the restrictions imposed by the superselection rule.

VII. QUANTUM BIT COMMITMENT AND SUPERSELECTION RULES

During the commitment stage of quantum bit commitment, Alice encodes a classical bit by preparing one of two

distinguishable quantum states with density operators ρ_0 or ρ_1 , and then she sends half of the state to Bob. In the unveiling stage, Alice sends the other half of the state to Bob, so that he can verify whether the state is ρ_0 or ρ_1 . The protocol is binding if, after commitment, Alice is unable to change the value of the bit. The protocol is concealing if, after commitment and before unveiling, Bob is unable to discern the value of the bit. The protocol is secure if it is both binding and concealing.

In the absence of superselection rules, unconditionally secure quantum bit commitment is impossible [8,9]. If we imagine that the states ρ_0 and ρ_1 are pure states shared by Alice and Bob, then if the protocol is concealing, Bob's density operator (obtained by tracing over Alice's system) must be the same in both cases: $\rho_{0,B} = \rho_{1,B}$. But then by the HJW theorem [29] Alice can apply a unitary transformation to her half of the state that transforms ρ_0 to ρ_1 , so that the protocol is not binding.

A. Bit commitment with mixed states

We reached this conclusion under the assumption that ρ_0 and ρ_1 are pure states, but we can extend the argument to the case where the states are mixed by appealing to the concept of a *purification* of a mixed state. We will describe this extension in detail, as we will follow very similar reasoning in our discussion in Sec. VII C of bit commitment with nontrivial total charge.

Suppose that at the start of the bit commitment protocol, Alice and Bob share a product state $\rho_A \otimes \rho_B$, where the states ρ_A and ρ_B are mixed. An equivalent way to describe Alice's initial state is to introduce the ancilla system C and a pure state $|\psi\rangle_{AC}$ (a purification of ρ_A), such that the density operator ρ_A is obtained from $|\psi\rangle_{AC}$ by tracing over system C :

$$\rho_A = \text{tr}_C(|\psi\rangle\langle\psi|)_{AC}. \quad (92)$$

Similarly, to describe ρ_B we can introduce the ancilla D and a state $|\varphi\rangle_{BD}$ that purifies ρ_B . Without loss of generality, we may assume that in each step of the protocol, Alice or Bob applies a unitary transformation, so that the state of the full system $ABCD$ remains pure. (A general quantum operation performed by Alice, say, can be realized as a unitary transformation applied jointly to Alice's system and to an appropriate ancilla; therefore, the operation is unitary provided that we include this ancilla as part of the system.) In particular, after the bit is committed, the state of the full system is one of the two pure states $|\psi_0\rangle_{ABCD}$ or $|\psi_1\rangle_{ABCD}$.

If both parties are honest, the ancillas C and D are off limits—Alice can manipulate only A and Bob can manipulate only B —and in that case the mixed state protocol and its purification are completely equivalent. Furthermore, if one party cheats, whether the other party starts out with a mixed state or its purification has no impact on the effectiveness of the cheating strategy, because the honest party never touches the purifying ancilla anyway.

Now let us see that in any quantum bit commitment protocol, one of the players can cheat successfully. First suppose that Bob cheats. Though the honest protocol calls for Bob to start out with the mixed state ρ_B , a cheating Bob can throw

this state away, and replace it with the purification $|\varphi\rangle_{BD}$, where D is now an ancilla system that Bob controls. Therefore, if the protocol is perfectly concealing (even when Bob cheats), then

$$\rho_{0,BD} \equiv \text{tr}_{AC}(|\psi_0\rangle\langle\psi_0|)_{ABCD} = \rho_{1,BD} \equiv \text{tr}_{AC}(|\psi_1\rangle\langle\psi_1|)_{ABCD}; \quad (93)$$

Bob is unable to collect any information about the committed bit through any joint measurement on BD .

Similarly, a cheating Alice could throw away her initial state and replace it by its purification; then Alice could control both A and the ancilla C . Applying the HJW theorem as before, we conclude that if $\rho_{0,BD} = \rho_{1,BD}$, then Alice can apply a unitary transformation to AC that transforms $|\psi_0\rangle_{ABCD}$ to $|\psi_1\rangle_{ABCD}$. We conclude that if the protocol is concealing, then it is not binding. Unconditionally secure quantum bit commitment is impossible, even with mixed states. That quantum bit commitment is impossible even when mixed strategies are used was proved in [8] using a slightly different approach.

B. Trivial total charge

The argument in Sec. VII A shows that for an analysis of the security of quantum bit commitment, we may assume that Alice and Bob share a pure state. But how is the security affected if superselection rules constrain Alice's and Bob's operations? We will first consider the special case in which the total charge that Alice and Bob share is trivial. After commitment, then, Alice and Bob share one of the two pure states $|\psi_0\rangle$ or $|\psi_1\rangle$, each with trivial total charge. Choosing the Schmidt basis in each charge sector, the state $|\psi_0\rangle$ can be expanded as

$$|\psi_0\rangle_{AB} = \sum_q \sqrt{p_q} \sum_b \sqrt{\lambda_{q,b}} |\bar{q}, b\rangle_A \otimes |q, b\rangle_B, \quad (94)$$

where Bob's density operator is

$$\rho_{0,B} = \text{tr}_A(|\psi_0\rangle\langle\psi_0|) = \sum_q p_q \rho_{0,B,q} \quad (95)$$

and

$$\rho_{0,B,q} = \sum_b \lambda_{q,b} |q, b\rangle\langle q, b|. \quad (96)$$

Bob can measure the probability p_q that his charge is q ; therefore if the protocol is concealing, then the distribution $\{p_q\}$ must be the same for $|\psi_1\rangle$ as for $|\psi_0\rangle$. Furthermore, Bob's density operator in the charge- q sector must not depend on whether the state is $|\psi_0\rangle$ or $|\psi_1\rangle$; therefore $|\psi_1\rangle$ can be expanded as

$$|\psi_1\rangle_{AB} = \sum_q \sqrt{p_q} \sum_b \sqrt{\lambda_{q,b}} |\bar{q}, \tilde{b}\rangle_A \otimes |q, b\rangle_B, \quad (97)$$

where $\{|\bar{q}, \tilde{b}\rangle_A\}$ is another basis for Alice's charge- \bar{q} sector. But now Alice can apply a unitary transformation conditioned on the charge that rotates one basis to the other:




FIG. 2. “Purification” of a two-party game with nontrivial total charge. At the beginning of the game, the charge of C (hidden behind a brick wall) compensates for Alice's charge q_A , and the charge of D (also hidden) compensates for Bob's charge q_B . Honest players never touch the compensating charges, but a cheating Alice might manipulate C and a cheating Bob might manipulate D .

$$U_{\bar{q}}: |\bar{q}, b\rangle \rightarrow |\bar{q}, \tilde{b}\rangle, \quad (98)$$

which transforms $|\psi_0\rangle$ to $|\psi_1\rangle$. Therefore, the protocol is not binding.

Obviously, the same argument applies, in the Abelian case, even if the total charge is nontrivial [16]. The key property of the states that is used in the argument is that Alice's charge is perfectly correlated with Bob's, so that $\mathcal{B}' = \mathcal{A}$.

C. Nontrivial total charge

The property that $\mathcal{B}' \neq \mathcal{A}$ in the non-Abelian case (with nontrivial total charge) encourages one to hope that a bit commitment protocol can be formulated whose security is founded on a non-Abelian superselection rule. Indeed, consider again the case $G = \text{SU}(2)$ with total charge $j = \frac{1}{2}$ as in Eq. (73). When Alice has control of the full AB system, she can prepare either of the states $|\pm\rangle_{AB}$ shown in Eq. (90), and then she can send the B system to Bob. Now Bob is unable to distinguish the two states, because he cannot measure the relative phase in a superposition of two states of different charge. Furthermore there is no invariant operation Alice can apply that changes $|+\rangle$ to $|-\rangle$ or vice versa. It seems, then, that the protocol is both concealing and binding. At any rate, quantum bit commitment in a world with non-Abelian superselection rules seems fundamentally different from quantum bit commitment in a world in which all superselection rules are Abelian.

But, as always in a discussion of information-theoretic security, we must be sure to consider the most general possible cheating strategies. And in fact, we can argue that for the security analysis, there is no loss of generality if we assume that the charge shared by the parties is trivial, the case we have already dealt with in Sec. VII B. This reduction to the case of trivial total charge follows closely our discussion in Sec. VII A, where we showed that it suffices to assume that the parties share a pure state.

Consider a general two-party quantum bit commitment protocol in which the initial state shared by Alice and Bob is a tensor product $\rho_A \otimes \rho_B$ of invariant states. The state ρ_A can be purified if we introduce an ancilla C ; furthermore, the pure state of AC can be chosen to have trivial total charge. Similar, we can purify ρ_B using the ancilla D , in such a way that the pure state of BD has trivial total charge. (See Fig. 2.) Each operation performed by Alice or Bob can be taken to be a charge-conserving unitary transformation; therefore, at each stage of the protocol, the state of the full system $ABCD$ is a pure state with trivial total charge.

In the honest protocol, the ancillas C and D are inaccessible. But if Bob cheats, he can throw away the initial invariant state ρ_A called for in the protocol, and replace it by a trivially charged pure state of BD , where D is now an ancilla that Bob controls. Therefore, if the bit commitment protocol is concealing, then $\rho_{0,BD} = \rho_{1,BD}$ —Bob cannot learn anything about the committed bit from any invariant joint measurement on BD . Since the state of the full system $ABCD$ is a pure state with trivial charge, the argument of Sec. VII B suffices to show that Alice can transform $|\psi_0\rangle$ to $|\psi_1\rangle$ with an invariant local operation applied to AC . Hence, the protocol is not binding. We have proved, then, that, even when the protocol calls for a nontrivial total charge, if Bob is unable to cheat then Alice can cheat—unconditionally secure quantum bit commitment is impossible. We have the following.

Theorem 2. Consider a quantum bit commitment protocol in the I world, where at the beginning of the protocol Alice and Bob share a product of invariant states. Then if the protocol is concealing, it is not binding.

Our proof, which reduces the case of nontrivial total charge to the case of trivial total charge, is really just a minor variant of the argument in Sec. VII A that reduces the case of a protocol where Alice and Bob share a mixed state to the case where they share a pure state.

In the case of our bit commitment protocol in which the total charge of AB is $j = \frac{1}{2}$, if Alice is unable to access the compensating charge in \bar{C} , then she cannot cheat successfully. But if Alice controls the whole AC system, then Alice's charge $j_{AC} = 0, 1$ is perfectly correlated with Bob's, and she can rotate the relative phase of the $j_{AC} = 0$ and $j_{AC} = 1$ components of her state, transforming $|+\rangle$ to $|-\rangle$.

This reduction of a protocol with nontrivial total charge to a protocol with trivial total charge can be generalized. In the I world, consider an n -party protocol in which up to $k < n$ of the parties might cheat, where the initial state is the product of invariant states $\otimes_{i=1}^n \rho_i$, and where all operations performed by the parties are required to conserve the local charge. Then we may imagine that each party is issued a compensating charge at the beginning of the protocol, so that each party actually starts out with trivial charge. The honest parties will never touch their compensating charges, but a cheating party cannot be prevented from performing arbitrary joint operations on her system and her compensating charge. This strategy is realizable because the cheater might throw away the invariant state she holds at the beginning of the protocol, and replace it by a charge-zero state that she controls fully. Furthermore, if an attack by the cheaters is successful in the protocol where the honest players start out with trivial charge, then it will also be successful if the honest players start out with a product of charged invariant states; since honest players never make use of the compensating charges, their presence can have no impact on the effectiveness of the attack. Therefore, we have the following.

Theorem 3. Let P be an n -party quantum protocol in the I world that securely realizes a task Π , where the initial state in P is a product of n invariant states. Then there is an I -world protocol P' that also securely realizes Π , where the initial state in P' is a product of n pure states, each with trivial charge.

In other words, in a security analysis, we may assume without any loss of generality that each party holds a pure

state with trivial charge at the start of the protocol.

Note that for the proofs of Theorems 2 and 3, our observations from Secs. II and III on the use of reference systems are not needed. Rather, to prove Theorems 2 and 3, we use only two properties of the I -world superselection sectors: first, that for each charge sector \mathcal{H}_q there is a unique conjugate charge sector $\mathcal{H}_{\bar{q}}$ such that the trivial sector \mathcal{H}_1 is contained in $\mathcal{H}_q \otimes \mathcal{H}_{\bar{q}}$, and second, that any invariant state has a purification with trivial total charge. These properties hold not just for the case of superselection rules arising from a symmetry group G , but also for the more general superselection rules considered in Sec. VIII. Therefore, Theorems 2 and 3 apply in this more general setting.

VIII. TWO-PARTY PROTOCOLS IN GENERAL

A. Overview

We will now analyze the impact of superselection rules on the security of general two-party protocols. We will show that for any protocol P in the invariant world (I world) subject to the superselection rule, there is a corresponding protocol \tilde{P} in the unrestricted world (U world), where \tilde{P} simulates P in the following sense: First, when performed honestly, \tilde{P} and P accomplish the same task. And second, for any cheating strategy that can be adopted by a dishonest party in \tilde{P} , there is a corresponding cheating strategy in P that is just as effective. In particular then, if \tilde{P} is insecure, then so is P . We conclude, therefore, that superselection rules cannot enhance the (information-theoretic) security of two-party protocols. The methods we will use to establish this result are quite different from those used in Sec. III to treat the case of superselection rules arising from a symmetry group.

Before going into detail, we will briefly describe the main ideas used in our argument. First of all, we will restrict our attention to a protocol in which the total charge shared by the two parties is trivial (belongs to the trivial superselection sector). We know from Theorem 3 in Sec. VII C that it suffices to treat this special case in an analysis of security. A protocol with trivial total charge has this useful property: if Alice knows that she holds charge q after sending a message to Bob, then Alice also knows that Bob will hold the conjugate charge \bar{q} upon receiving the message. Similarly, Bob knows what Alice's charge will be after she receives a message sent by Bob. Our analysis of security relies on the property that Bob has a definite charge if Alice does, and therefore it applies only to two-party protocols.

In the I world, charge is conserved, so that the total charge shared by Alice and Bob is trivial at each stage of the protocol; furthermore, local operations performed by Alice or Bob must preserve the conserved charge. In the U world, charge need not be conserved, but the protocol \tilde{P} that simulates the I -world protocol P can be chosen to respect conservation of a fictitious “charge” that behaves like the actual conserved charge of the I world. However, a dishonest party who is not bound to follow the protocol \tilde{P} can perform operations that violate “charge” conservation. Our task is to

ensure that the greater freedom enjoyed by a dishonest party in the U world does not enhance her ability to cheat successfully.

For this purpose, our argument relies on the concept of the *format* of a message exchanged between the parties. In the U world, the format is simply the Hilbert space containing the message. In the protocol \tilde{P} , the recipient of a message always checks that the format of the message is valid, and aborts the protocol if the message is invalid. A valid message corresponds to one that could have been sent in the I world, while a message is invalid only if the sender violated the local conservation of “charge” before sending it. Thus, a message that upon receipt is found to be in the proper format could have been sent by a party who performed a charge-conserving local operation—in effect the sender is unable to play a charge-nonconserving strategy without being detected. Since effective charge conservation is enforced by halting the protocol when a charge nonconservation is detected, it will be essential for our argument to consider games that can be aborted at any stage by either party. A cheating strategy for the I -world protocol P and the corresponding cheating strategy for its U -world counterpart \tilde{P} will cause the game to halt prematurely with the same probability, as well as produce the same probability distribution of outcomes in the event that the game ends normally, without being aborted.

B. Superselection rules and charges

Before proceeding to our proof, we should recall the properties of superselection rules and charges that will be invoked in the argument. These properties have been explored already in Sec. V, for the special case of super-selection sectors labeled by irreducible unitary representations of compact groups. Here we wish to emphasize that some of the same ideas can be extended to a more general setting, and we will indicate how a two-party protocol in which conserved charges are exchanged can be simulated using ordinary qubits.

In general, a superselection rule is a decomposition of Hilbert space into a direct sum of sectors such that each sector is preserved by the allowed operations. The charge q is a label that distinguishes the distinct sectors, and we may say that the operations allowed by the superselection rule conserve the charge. Thus, the Hilbert space is expressed as

$$\mathcal{H} = \bigoplus_q \mathcal{H}_q, \quad (99)$$

and the allowed operations belong to the algebra

$$\bigoplus_q \mathcal{L}(\mathcal{H}_q), \quad (100)$$

where $\mathcal{L}(\mathcal{H}_q)$ denotes linear operators acting on \mathcal{H}_q .

Depending on the particular form of the superselection rule, there are specific rules governing how the charge behaves when a system splits into two subsystems, or when two systems fuse to become a single system. These rules can be encoded in vector spaces $V_c^{a,b}$ defined by

$$\mathcal{H}_c = \bigoplus_{a,b} \mathcal{H}_a \otimes \mathcal{H}_b \otimes V_c^{a,b}. \quad (101)$$

The space $V_c^{a,b}$ is n -dimensional if there are n distinguishable ways that a charge c object can arise when objects with charges a and b fuse. Consistency of Eq. (101) with associativity of the tensor product requires the $V_c^{a,b}$'s to obey certain identities, but we will not discuss these further as they will not be needed for our proof.

There is a trivial-charge sector, denoted \mathcal{H}_1 , that behaves as the identity under fusion,

$$\mathcal{H}_c \otimes \mathcal{H}_1 = \mathcal{H}_c. \quad (102)$$

Furthermore, there is a unique charge \bar{q} , the conjugate of q , that can fuse with q to yield the identity

$$\mathcal{H}_1 = \bigoplus_q \mathcal{H}_q \otimes \mathcal{H}_{\bar{q}}. \quad (103)$$

Now, in the I world, consider a bipartite system shared by Alice and Bob. The Hilbert space decomposes as

$$\mathcal{H} = \bigoplus_q \mathcal{H}_q,$$

$$\mathcal{H}_q = \bigoplus_{q_A, q_B} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{B, q_B} \otimes V_q^{q_A, q_B}, \quad (104)$$

where q is the total charge, q_A is the charge of Alice's system, and q_B is the charge of Bob's system. The physical operations, allowed by the superselection rule, conserve the total charge, and hence belong to the algebra

$$\mathcal{O} = \bigoplus_q \mathcal{L}(\mathcal{H}_q). \quad (105)$$

The operations Alice can perform, which conserve Alice's charge and act trivially on Bob's system, belong to

$$\mathcal{A} = \bigoplus_{q, q_A, q_B} \mathcal{L}(\mathcal{H}_{A, q_A}) \otimes I_{B, q}^{q_A, q_B}, \quad (106)$$

where $I_{B, q}^{q_A, q_B}$ denotes the identity acting on $\mathcal{H}_{B, q_B} \otimes V_q^{q_A, q_B}$. Similarly, the algebra of operations that Bob can perform is

$$\mathcal{B} = \bigoplus_{q, q_A, q_B} I_{A, q}^{q_A, q_B} \otimes \mathcal{L}(\mathcal{H}_{B, q_B}), \quad (107)$$

where $I_{A, q}^{q_A, q_B}$ denotes the identity acting on $\mathcal{H}_{A, q_A} \otimes V_q^{q_A, q_B}$. In contrast, the commutant \mathcal{B}' of \mathcal{B} , which conserves the total charge and Bob's charge but need not conserve Alice's, is

$$\mathcal{B}' = \bigoplus_{q, q_B} \mathcal{L}\left(\bigoplus_{q_A} \mathcal{H}_{A, q_A} \otimes V_q^{q_A, q_B}\right) \otimes I_{B, q_B}, \quad (108)$$

where I_{B, q_B} is the identity on \mathcal{H}_{B, q_B} , and similarly

$$\mathcal{A}' = \bigoplus_{q, q_A} I_{A, q_A} \otimes \mathcal{L}\left(\bigoplus_{q_B} \mathcal{H}_{B, q_B} \otimes V_q^{q_A, q_B}\right). \quad (109)$$

Thus $\mathcal{A}' = \mathcal{B}$ and $\mathcal{B}' = \mathcal{A}$ if and only if the charges q_A and q_B are perfectly correlated (there is a unique q_B corresponding to each q_A and vice versa). This condition holds, in particular, if the total charge is trivial, in which case our formulas simplify to

$$\begin{aligned}
\mathcal{H} &= \mathcal{H}_1 = \bigoplus_q \mathcal{H}_{A,q} \otimes \mathcal{H}_{B,\bar{q}}, \\
\mathcal{A} &= \mathcal{B}' = \bigoplus_q \mathcal{L}(\mathcal{H}_{A,q}) \otimes I_{B,\bar{q}}, \\
\mathcal{B} &= \mathcal{A}' = \bigoplus_q I_{A,q} \otimes \mathcal{L}(\mathcal{H}_{B,\bar{q}}). \tag{110}
\end{aligned}$$

C. Simulating charge exchange

A novelty of a two-party protocol in the I world is that when Alice (for example) sends a message to Bob, she may choose to split the charge she possesses into two parts—the charge she retains and the charge of the message that she sends. If the total charge is trivial, then the full Hilbert space comprising Alice’s system A , Bob’s system B , and the message system M can be expressed as

$$\mathcal{H}_1 = \bigoplus_{q_A, q_B, q_M} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{B, q_B} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M}. \tag{111}$$

The isomorphisms

$$V_1^{q_A, q_B, q_M} \cong V_B^{q_A, q_M} \cong V_A^{q_B, q_M} \tag{112}$$

invite us to interpret Eq. (111) in complementary ways—namely, the charge \bar{q}_B of AM is conjugate to the charge q_B of B , and the charge \bar{q}_A of BM is conjugate to the charge q_A of A . Thus, Eq. (111) describes the splitting of Alice’s initial charge \bar{q}_B into the charge q_A that she retains and the charge q_M of the message, as well as the fusion of the charge q_M of the message with Bob’s initial charge q_B to yield Bob’s final charge \bar{q}_A . Furthermore, if $V_1^{q_A, q_B, q_M}$ is of dimension greater than 1, then a vector in $V_1^{q_A, q_B, q_M}$ describes the particular manner in which Alice performs the splitting, which in turn determines the result of Bob’s fusion.

While the information encoded in $V_1^{q_A, q_B, q_M}$ is an intrinsic property in the I world, if we are to simulate the process of charge exchange in the U world, then this information must be carried by ordinary qubits. In such a simulation, the Hilbert space of Alice’s system, Bob’s system, and the message is expanded to

$$\tilde{\mathcal{H}} = \bigoplus_{q_1, q_2, q_A, q_B, q_M} \mathcal{H}_{A, q_1} \otimes \mathcal{H}_{B, q_2} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M}, \tag{113}$$

but where now $V_1^{q_A, q_B, q_M}$ is to be regarded as an explicit part of the message. If the conditions $q_1 = q_A$ and $q_2 = q_B$ were imposed, then the “format” of this message would coincide perfectly with the information content of a message sent in the I world. But while in the I world these conditions arise from the intrinsic physics of the superselection rule, in the U world they must be imposed by hand through proper design of the protocol.

Thus, in the U -world protocol \tilde{P} that simulates the I -world protocol P , we will require the recipient of a message to verify its format—Alice checks that $q_1 = q_A$ and Bob checks that $q_2 = q_B$. Of course, at a given stage of the protocol

P , Alice or Bob might hold a coherent superposition of different charges, even though the total charge is always guaranteed to be trivial. Therefore, the verification step in \tilde{P} must be performed coherently; Alice, for example, checks that q_1 and q_A match without learning the value of q_1 or q_A . If verification fails, then the message recipient has detected cheating by the other party and aborts the protocol. If verification succeeds, then the message has been projected onto the valid format, and as far as the recipient is concerned, it is just as though the message had been sent in the right format to begin with.

Whenever Alice cheats in the U -world protocol \tilde{P} by modifying her charge, she risks detection, and if her cheating is undetected, then her operation is equivalent to a charge-conserving one. Therefore, Alice has an equivalent strategy in the I -world protocol P , in which she either halts the game herself with some probability before sending her message, or if the game does not halt, performs an operation allowed by the superselection rule. This observation suffices to establish that \tilde{P} simulates P , and thus that the superselection rule cannot thwart cheating.

To summarize, for the purpose of characterizing Alice’s ability to cheat, we are only interested in how Alice’s activities will affect Bob’s measurements. Although in the U world Alice has the power to violate conservation of “charge,” she is unable to fool Bob into accepting a message that is not isomorphic to one that could have been created in the I world. Therefore, Alice’s elevated power in the U world gives her no advantage.

D. Definitions

Having explained the main ideas, we will now present a more formal proof of our result. To begin, we must define the general notions of “protocol” and “simulation” in accord with our goals. The definitions are quite natural, but there are some technicalities that are necessary for the proof to work.

We consider quantum games between two parties, Alice and Bob. We assume that Alice sends the first message and the players alternate. The *protocol* of a game specifies the total number of messages, their format, the strategies for honest players, and a way to determine the game outcome. By “format” in the U world we mean the Hilbert space \mathcal{H}_M of a given message. In the I world, we specify the space \mathcal{H}_{M, q_M} for each value of the message charge q_M .

To define an honest strategy in the I world, we specify for each value of Alice’s charge q_A her corresponding space \mathcal{H}_{A, q_A} ; likewise, we specify Bob’s space \mathcal{H}_{B, q_B} for each q_B . The game starts with a pure state

$$|\xi_A\rangle \otimes |\xi_B\rangle \in \mathcal{H}_{A,1} \otimes \mathcal{H}_{B,1}, \tag{114}$$

where 1 stands for the trivial charge. If one of the players (say, Alice) cheats, she may use a different set of private spaces \mathcal{H}'_{A, q_A} , but the initial state still must be of the form $|\xi'_A\rangle \otimes |\xi_B\rangle$, where $|\xi'_A\rangle \in \mathcal{H}'_{A,1}$.

Alice’s and Bob’s actions in the k th step are described by operators W_{A_k}, W_{B_k} . The final outcome is determined by a pair of measurements that are performed independently on Al-

	<i>I</i> -world	<i>U</i> -world
Protocol	P	\tilde{P}
Alice strategy	A'	\tilde{A}'
Bob strategy	B'	\tilde{B}'

FIG. 3. The *U*-world protocol \tilde{P} simulates the *I*-world protocol P if the honest protocols realize the same task, and if for any cheating strategy in \tilde{P} there is an equivalent cheating strategy in P .

ice’s and Bob’s subsystems at the end of the game. We are interested in the joint probability distribution of the measurement results. However, if one of the players cheats, only the honest player’s subsystem is measured.

For the reasons explained in Sec. VIII A, we will assume that the game can be aborted by either player. If the game is aborted, we will not need to keep track of who ends the game or when it ends—we will only be interested in whether the game ends normally and if so what is the outcome. For this purpose, the quantum state can be characterized by a vector $|\psi\rangle$ such that $\langle\psi|\psi\rangle$ is the probability that the game has not been aborted. Operations performed by each player may then be described by contracting maps, i.e., operators W such that $W^\dagger W \leq I$. We assume that the game is never aborted if both players are honest, so that the probabilities of different outcomes add up to 1 in the honest game. If one of the players cheats, the total probability of all outcomes is generally less than 1.

Now we define what it means for one protocol to *simulate* another (see Fig. 3).

Definition. A protocol \tilde{P} simulates the protocol P if the following conditions are fulfilled:

- (i) The honest strategies in P and \tilde{P} give rise to the same probability distribution of the outcomes.
- (ii) For any cheating strategy \tilde{A}' by Alice compatible with the protocol \tilde{P} , there exists an equivalent strategy A' for the protocol P . (“Equivalent” means that Bob’s measurement result has the same probability distribution in both cases.)
- (iii) For any cheating strategy \tilde{B}' by Bob compatible with the protocol \tilde{P} there is an equivalent strategy B' for the protocol P .

Note that when we say that the two cheating strategies are equivalent, we mean in particular that the probability that the game ends normally is the same for both strategies.

To better understand our concept of simulation, it is very helpful to consider this simple example: Suppose that the message space \mathcal{H}_M of P is embedded in a larger space $\tilde{\mathcal{H}}_M$ of \tilde{P} . Honest players follow the same strategies in \tilde{P} as in P , so that condition (i) is obviously satisfied. However, the players in \tilde{P} must be prepared to receive messages that do not obey the format of P , i.e., do not fit into the subspace \mathcal{H}_M . In \tilde{P} such messages are rejected, and the game is aborted. This rule prevents a dishonest player from gaining any advantage

(relative to simply quitting the game) by sending an invalid message. More formally, suppose that Alice cheats using some strategy \tilde{A}' . In the corresponding strategy A' , Alice projects her message system $\tilde{\mathcal{H}}_M$ onto the subspace \mathcal{H}_M , before sending each message. Thus if the strategy \tilde{A}' calls for Alice to apply the operator \tilde{W}'_{A_k} in the k th round, then in the strategy A' Alice applies the contracting map $W'_{A_k} = \Pi \tilde{W}'_{A_k}$, where Π is the orthogonal projector onto \mathcal{H}_M . The strategies \tilde{A}' and A' are equivalent: whenever a message sent according to \tilde{A}' causes Bob to abort the game, the strategy A' requires Alice to abort the game herself. Similarly, given any cheating strategy \tilde{B}' for Bob in the game \tilde{P} , there is an equivalent cheating strategy B' in P . Thus, conditions (ii) and (iii) are satisfied, and \tilde{P} simulates P .

Our analysis of superselection rules in Sec. VIII E will be based on a closely related method of simulation.

We also remark that Theorem 1 proved in Sec. III can be restated: for a multiparty protocol P in the G -invariant world, there is a *U*-world protocol \tilde{P} that simulates P . In that case, we implicitly adopt a redundant description of the physical states appearing in P , admitting fictitious color degrees of freedom. Then \tilde{P} is exactly the same protocol as P , but with the color now reinterpreted as a physical variable. Similarly, Theorem 3 in Sec. VII C can be stated: any n -party *I*-world protocol in which the initial state is a product of n invariant states can be simulated by an *I*-world protocol in which the initial state is a product of n pure states, each with trivial charge.

E. Proof

Our goal is to prove the following.

Theorem 4. Let P be a two-party game in the *I* world, such that both parties hold trivial charges at the beginning of the game. Then there is a *U*-world game \tilde{P} that simulates P .

In the proof, we construct the *U*-world protocol \tilde{P} that simulates the *I*-world protocol P , and explain how the cheating strategy A' that is equivalent to \tilde{A}' is formulated. We achieve this by applying the procedure for simulating charge exchange in the *U* world that was described in Sec. VIII C.

Consider the *I*-world protocol P . If the total charge is trivial, then the full Hilbert space including Alice’s system A , Bob’s system B , and the message M is

$$\mathcal{H} = \bigoplus_{q_A, q_B, q_M} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{B, q_B} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M}. \tag{115}$$

Without loss of generality, we assume that the spaces \mathcal{H}_{A, q_A} , \mathcal{H}_{B, q_B} , \mathcal{H}_{M, q_M} are the same in each step of the protocol. We may also assume that the message is present at the beginning and at the end of the game and that the initial state has the form $|\xi_A\rangle \otimes |\xi_B\rangle \otimes |0\rangle$, where $|0\rangle \in \mathcal{H}_{M, 1}$.

Each time Alice receives one message and sends another, she applies an operator to AM that preserves Bob’s charge q_B ; this is a contracting map belonging to the algebra

$$\oplus_B \mathcal{L} \left(\oplus_{q_A, q_M} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M} \right). \quad (116)$$

Alice's honest strategy consists of a sequence of such operators—in the k th step she applies an operator W_{A_k} . Similarly, Bob's honest strategy is defined by operators W_{B_k} .

Now consider the U -world protocol \tilde{P} that simulates P . The Hilbert space of \tilde{P} is

$$\tilde{\mathcal{H}} = \tilde{\mathcal{H}}_A \otimes \tilde{\mathcal{H}}_B \otimes \tilde{\mathcal{H}}_M, \quad (117)$$

where

$$\begin{aligned} \tilde{\mathcal{H}}_A &= \oplus_{q_1} \mathcal{H}_{A, q_1}, & \tilde{\mathcal{H}}_B &= \oplus_{q_2} \mathcal{H}_{B, q_2}, \\ \tilde{\mathcal{H}}_M &= \oplus_{q_A, q_B, q_M} \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M}. \end{aligned} \quad (118)$$

Thus the space \mathcal{H} of the protocol P can be embedded in $\tilde{\mathcal{H}}$ by requiring $q_1 = q_A$ and $q_2 = q_B$. In \tilde{P} , these constraints are enforced by checks performed by both parties. A dishonest player's attempt to break the constraints will be detected immediately by the other party, in which case the game will halt.

Let us describe Alice's honest strategy in \tilde{P} . When Alice receives a message, she gains control of the space $\tilde{\mathcal{H}}_A \otimes \tilde{\mathcal{H}}_M$. First she verifies that $q_1 = q_A$ (without determining the value of q_1 or q_A); if verification fails, she aborts the game. Thus Alice effectively projects her input state onto the subspace

$$\mathcal{H}_{AM} = \oplus_{q_A, q_B, q_M} \mathcal{H}_{A, q_A} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M} \subseteq \tilde{\mathcal{H}}_A \otimes \tilde{\mathcal{H}}_M. \quad (119)$$

Then she applies the operator W_{A_k} (from the protocol P), which acts on \mathcal{H}_{AM} and preserves q_B . Thus Alice's strategy is defined by the contracting maps

$$\tilde{W}_{A_k} = F W_{A_k} F^\dagger, \quad (120)$$

where F denotes the embedding $\mathcal{H}_{AM} \rightarrow \tilde{\mathcal{H}}_A \otimes \tilde{\mathcal{H}}_M$. Bob's honest strategy is defined similarly.

If both players play the game \tilde{P} honestly, then the verification always succeeds and the conditions $q_1 = q_A$ and $q_2 = q_B$ are maintained throughout the game. Thus the honest strategies for \tilde{P} and P are clearly equivalent. Note that in \tilde{P} some information is encoded redundantly—for example, Alice can access the value of q_A by examining either the charge label of \mathcal{H}_{A, q_A} or one of the slots of the tensor $V_1^{q_A, q_B, q_M}$; similarly q_M is encoded both in \mathcal{H}_{M, q_M} and in $V_1^{q_A, q_B, q_M}$. However, this redundancy has no deleterious effect on the fidelity of the simulation.

Now suppose that Alice cheats in the game \tilde{P} . Then she may use an arbitrary Hilbert space $\tilde{\mathcal{H}}'_A$ and operators \tilde{W}'_{A_k} acting on

$$\tilde{\mathcal{H}}'_{AM} = \tilde{\mathcal{H}}'_A \otimes \tilde{\mathcal{H}}_M = \tilde{\mathcal{H}}'_A \otimes \left(\oplus_{q_A, q_B, q_M} \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M} \right). \quad (121)$$

In particular, when Alice cheats, her action on the message need not respect the condition $q_B = q_2$. To prove the theorem, we are to define an equivalent cheating strategy for the game P .

When Alice cheats in P , she uses an arbitrary Hilbert space \mathcal{H}'_{A, q_A} for each value of her charge q_A , and she applies operators W'_{A_k} that conserve Bob's charge q_B to the space

$$\mathcal{H}'_{AM} = \oplus_{q_A, q_B, q_M} \mathcal{H}'_{A, q_A} \otimes \mathcal{H}_{M, q_M} \otimes V_1^{q_A, q_B, q_M}. \quad (122)$$

The spaces $\tilde{\mathcal{H}}'_{AM}$ and \mathcal{H}'_{AM} seem to be distinct—in $\tilde{\mathcal{H}}'_{AM}$ the charge label carried by \mathcal{H}'_{A, q_A} matches the label in one of the slots of $V_1^{q_A, q_B, q_M}$, while in \mathcal{H}'_{AM} there is no such correlation. However, in the U world the variable q_A would be encoded redundantly if it appeared in both \mathcal{H}'_{A, q_A} and $V_1^{q_A, q_B, q_M}$, and it is not necessary to adopt this redundant encoding in order to emulate the physics of the I world. Instead, let us specify $\mathcal{H}'_{A, q_A} = \tilde{\mathcal{H}}'_A$ for each q_A —then \mathcal{H}'_{AM} and $\tilde{\mathcal{H}}'_{AM}$ are of the same form, but where it is understood in Eq. (121) that the information about the charge q_A is carried only by $V_1^{q_A, q_B, q_M}$. With this choice Alice's operator \tilde{W}'_{A_k} in \tilde{P} and her operator W'_{A_k} in P act on isomorphic spaces; however, W'_{A_k} must conserve Bob's charge q_B , while \tilde{W}'_{A_k} need not conserve charge.

Therefore, we define the corresponding cheating strategy in P by specifying

$$W'_{A_k} = \sum_{q_B} \Pi_{q_B} \tilde{W}'_{A_k} \Pi_{q_B}, \quad (123)$$

where Π_{q_B} is the projector onto the subspace with the given value of q_B . That is, Π_{q_B} projects $\tilde{\mathcal{H}}_M$ onto the space in which $V_1^{q_A, q_B, q_M}$ has the value q_B in the appropriate slot. The contracting map W'_{A_k} preserves q_B and therefore is admissible in the protocol P . Applying this W'_{A_k} causes Alice to abort the game P in the case where q_B would change in the game \tilde{P} . But in that case the new value of q_B would not match Bob's variable q_2 ; therefore, Bob would reject Alice's message and abort the game \tilde{P} . Hence the two games P and \tilde{P} are aborted with the same probability; furthermore, the final state that Bob measures in \tilde{P} , if \tilde{P} does not abort, is identical to the final state that Bob measures in P , if P does not abort. Therefore, when Alice cheats, Bob's measurement outcome has the same probability distribution in \tilde{P} as in P . The same is true for Alice's measurement when Bob cheats. Therefore, \tilde{P} simulates P , which completes the proof of Theorem 4.

IX. CONCLUSIONS

Recent progress in the theory of quantum computation and quantum cryptography highlights the importance of adopting a computational model compatible with fundamental physics—tasks that would be impossible in a classical

world may be physically realizable because Nature is quantum-mechanical. Further refinements of the model could lead to further insights regarding what information-processing tasks are achievable. Therefore, as Popescu [12] emphasized, the impact of superselection rules on the security of quantum protocols is of considerable potential interest. However, our disappointing conclusion is that superselection rules cannot foil a cheater who has unlimited quantum-computational power.

Contemplating this issue has led us to consider how physics in the invariant world can simulate physics in the unrestricted world, and vice versa. We feel that the simulation schemes we have devised offer fruitful insights into the physical meaning of superselection rules.

Our results do not address whether the security of protocols with more than two parties can be enhanced by superselection rules that do not arise from compact symmetry groups. New issues arise in this setting, because of the non-trivial braiding properties of non-Abelian anyons. For example, in the case of three parties (Alice, Bob, and Charlie), Alice can split her charge into two parts, and send one part on a voyage that circles Bob's lab and then returns to Alice's lab. This action can induce a change in the charge held by

Alice, accompanied by a compensating change in the total charge held by Bob and Charlie, even though the local charge in Bob's lab, and in Charlie's, is unaltered. Though strictly speaking Alice's operation is not "local," she can carry it out surreptitiously, without any cooperation from Bob and Charlie. Such new possibilities enhance the potential power of cheaters, but may also provide the honest parties with new methods for detecting cheating. Addressing the security of multiparty quantum protocols subject to general superselection rules will require different methods from those we have used in this paper, and might provide further enlightenment concerning the physics of non-Abelian anyons.

ACKNOWLEDGMENTS

We thank Stephen Bartlett, Michael Ben-Or, and Sandu Popescu for discussions. This work has been supported in part by the Department of Energy under Grant No. DE-FG03-92-ER40701, by the National Science Foundation under Grant No. EIA-0086038, and by the Caltech MURI Center for Quantum Networks under ARO Grant No. DAAD19-00-1-0374.

-
- [1] C. H. Bennett and G. Brassard, *Quantum Cryptography: Public Key Distribution and Coin Tossing*, in Proceedings of IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India (IEEE, New York, 1984), pp. 175–179.
 - [2] D. Mayers, *Quantum Key Distribution and String Oblivious Transfer in Noisy Channels*, in Advances in Cryptography—Proceedings of Crypto'96 (Springer-Verlag, New York, 1996), pp. 343–357; *J. ACM* **48**, 351 (2001).
 - [3] H.-K. Lo and H. F. Chau, *Science* **283**, 2050 (1999).
 - [4] E. Biham, M. Boyer, P. O. Boykin, T. Mor, and V. Roychowdhury, *A Proof of the Security of Quantum Key Distribution*, in Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (ACM Press, New York, 2000), pp. 715–724.
 - [5] P. W. Shor and J. Preskill, *Phys. Rev. Lett.* **85**, 441 (2000).
 - [6] M. Koashi and J. Preskill, *Phys. Rev. Lett.* **90**, 057902 (2003).
 - [7] K. Tamaki, M. Koashi, and N. Imoto, *Phys. Rev. Lett.* **90**, 167904 (2003).
 - [8] D. Mayers, *Phys. Rev. Lett.* **78**, 3414 (1997).
 - [9] H.-K. Lo and H. F. Chau, *Phys. Rev. Lett.* **78**, 3410 (1997).
 - [10] A. Kitaev, <http://www.msri.org/publications/ln/msri/2002/qip/kitaev/1/>
 - [11] A. Ambainis, H. Buhrman, Y. Dodis, and H. Roehrig, e-print quant-ph/0304112.
 - [12] S. Popescu (unpublished).
 - [13] F. Verstraete and J. I. Cirac, e-print quant-ph/0302039.
 - [14] Y. Aharonov and L. Susskind, *Phys. Rev.* **155**, 1428 (1967).
 - [15] S. D. Bartlett and H. M. Wiseman, e-print quant-ph/0303140.
 - [16] D. Mayers, e-print quant-ph/0212159.
 - [17] S. Doplicher and J. E. Roberts, *Commun. Math. Phys.* **131**, 51 (1990).
 - [18] K.-H. Rehren, *Braid Group Statistics and their Superselection Rules*, in The Algebraic Theory of Superselection Sectors, edited by D. Kastler (World Scientific, Singapore, 1990), pp. 333–355.
 - [19] A. Yu Kitaev, *Ann. Phys.* **303**, 2 (2003).
 - [20] M. H. Freedman, A. Kitaev, and Z. Wang, *Commun. Math. Phys.* **227**, 587 (2002).
 - [21] G.-C. Wick, A. S. Wightman, and E. P. Wigner, *Phys. Rev.* **88**, 101 (1952).
 - [22] G.-C. Wick, A. S. Wightman, and E. P. Wigner, *Phys. Rev. D* **1**, 3267 (1970).
 - [23] K. Mølmer, *Phys. Rev. A* **55**, 3195 (1997).
 - [24] T. Rudolf and B. C. Sanders, *Phys. Rev. Lett.* **87**, 077903 (2001).
 - [25] S. J. van Enk and C. A. Fuchs, *Quantum Inf. Comput.* **2**, 151 (2002).
 - [26] S. D. Bartlett, T. Rudolf, and R. W. Spekkens, e-print quant-ph/0302111.
 - [27] S. J. van Enk and T. Rudolf, e-print quant-ph/0303096.
 - [28] A. Ambainis, e-print quant-ph/0204022.
 - [29] L. P. Hughston, R. Jozsa, and W. K. Wootters, *Phys. Lett. A* **183**, 14 (1993).