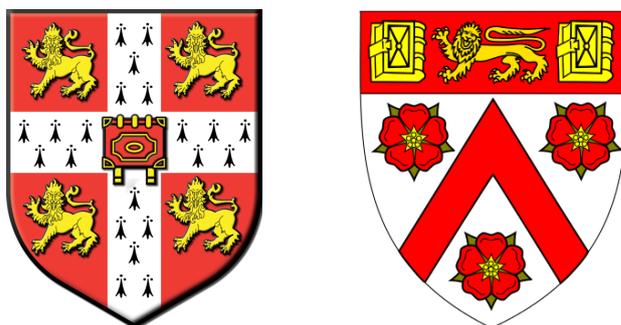


Relativistic Quantum Tasks

Emily Adlam

eadlam90@gmail.com



Thesis submitted to
The University of Cambridge

for the degree of

DOCTOR OF PHILOSOPHY

DAMTP
Wilberforce Road
Cambridge

September 2017

Contents

0.1	Abstract: Relativistic Quantum Tasks	7
0.2	Statement	9
0.3	Acknowledgements	9
1	Introduction	11
1.1	Special Relativity	12
1.2	Quantum Mechanics	14
1.2.1	Basic postulates	14
1.2.2	Entanglement	15
1.2.3	Monogamy	17
1.2.4	Mixed states and density operators	18
1.2.5	Quantum Contextuality	19
1.2.6	Quantum Information	21
1.2.7	Quantum Cryptography	26
1.2.8	Quantum Field Theory	30
1.2.9	Quantum Gravity	32
1.2.10	The Measurement Problem	32
2	Context	35
2.1	Tasks and Games	35
2.2	Paradoxes	39
2.3	Relativistic Quantum Cryptography	40
2.4	Zero-Knowledge-Proving	43

I	Distribution Tasks	47
3	Nonexistent Causal Diamonds	49
3.1	Introduction	49
3.2	Necessary and Sufficient Conditions	50
3.3	Summary	52
4	Quantum Paradox of Choice	53
4.1	Introduction	53
4.2	Necessary and Sufficient Conditions	55
4.3	Resolution	59
4.4	Finkelstein's Objection	60
4.5	Applications	62
4.6	Summary	63
5	Spooky Summoning	67
5.1	Introduction	67
5.2	Entanglement distribution	68
5.2.1	Necessary conditions	69
5.2.2	Entanglement vs Correlations	70
5.3	Entanglement summoning	71
5.3.1	Necessary conditions	72
5.3.2	Labelled calls	75
5.4	Summary	77
II	Relativistic Quantum Cryptography	81
6	Definitions	83
6.1	Relativistic Bit Commitment	83
6.1.1	Definition	84
6.1.2	Security definitions	86
6.2	KCEKQS	90

<i>CONTENTS</i>	5
7 Deterministic Bit Commitment	95
7.1 Introduction	95
7.2 Protocols	96
7.2.1 ETBC	97
7.2.2 ETRBC	99
7.3 Errors and Losses	102
7.4 Summary	103
8 Device-Independent Bit Commitment	107
8.1 Introduction	107
8.2 Protocols	110
8.2.1 CHSH 1	110
8.2.2 CHSH 2	114
8.2.3 CHSH 3	114
8.3 Extensions	116
8.3.1 Declining to commit	116
8.3.2 Errors and Losses	117
8.4 Discussion	117
8.4.1 RCCBC	119
8.4.2 Comparison	121
8.5 Summary	123
9 KCEKQS	125
9.1 Introduction	125
9.2 No-go theorems	126
9.3 Classical Protocols	127
9.3.1 CR1	129
9.3.2 CR2	131
9.4 Quantum A-to-B Protocols	133
9.4.1 QAB	133
9.5 Quantum B-to-A Protocols	134
9.5.1 QBA	135
9.6 Further Security Issues	138

9.7	Summary	144
10	Concluding Remarks	151
A	Supplementary Information for Chapter 9	157
A.1	Protocols with an abort option	157
A.2	No-Go Theorems	158
A.2.1	Zero-knowledge	158
A.2.2	Completeness vs soundness	159
A.3	Security Proofs for CR2	161
A.4	Security Proofs for QAB	162
A.5	Security Proofs for QBA	163

0.1 Abstract: Relativistic Quantum Tasks

Quantum mechanics, which describes the behaviour of matter and energy on very small scales, is one of the most successful theories in the history of science. Einstein's theory of special relativity, which describes the relationship between space and time, is likewise a highly successful and widely accepted theory. And yet there is a well-documented tension between the two theories, to the extent that it is still not clear that the two can ever be reconciled [1, 2].

This thesis is concerned with furthering the current understanding of the relationship between quantum mechanics and special relativity. In the first part of the thesis we study the behaviour of quantum information in relativistic spacetime. The field of quantum information arose from the realisation that quantum information has a number of crucial properties that distinguish it from classical information, such as the no-cloning property [3], quantum contextuality [4] and quantum discord [5]. More recently, it has been realised that placing quantum information under relativistic constraints leads to the emergence of further unique features which are not exhibited by either non-relativistic quantum information or relativistic classical information [6–8]; as part of this ongoing research programme we develop a new relativistic quantum ‘paradox’ which puts pressure on conventional views about the spatiotemporal persistence of quantum states over time. This part of the thesis is based on a paper co-authored with Adrian Kent and published in *Phys Rev A* (see ref [9]) and a second paper also co-authored with Adrian Kent (see ref [10]) published in response to Finkelstein's comments on our original paper (see ref [11]). We then study a new set of relativistic quantum protocols which involve the distribution of entangled states over spacetime, defining one task involving the distribution of the two halves of a known entangled state, and another task involving the distribution of the two halves of an unknown entangled state. This part of the thesis is based on unpublished work with Adrian Kent.

The second part of the thesis deals with relativistic quantum cryptography, a field which first began attracting serious attention when it was realised that a cryptographic task known as ‘bit commitment,’ can be implemented with perfect security under relativistic constraints [12, 13]. This result was highly significant,

since it is provably impossible to implement bit commitment with perfect security in a purely classical or purely quantum context [14, 15], and hence bit commitment is an ideal starting point for probing the power of relativistic quantum cryptography. In this thesis we propose several new relativistic quantum bit commitment protocols which have notable advantages over previously known protocols. This part of the thesis is based on two papers co-authored with Adrian Kent and published in *Phys Rev A* and the *International Journal of Quantum Information* (see refs [16, 17]). We then move to a related task, a generalization of zero-knowledge proving which we refer to as knowledge-concealing evidencing of knowledge of a quantum state; we prove no-go theorems concerning the possibility of implementing this task with perfect security, and then set out a relativistic protocol for the task which is asymptotically secure as the dimension of the state in question becomes large. These results have interesting foundational significance above and beyond their applications in the field of cryptography, providing a new perspective on the connections between knowledge, realism and quantum states. This part of the thesis is based on a paper co-authored with Adrian Kent (see ref [18]).

0.2 Statement

This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the preface and again in the text. It is not substantially the same as any that I have submitted, or am concurrently submitting for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution.

0.3 Acknowledgements

Thanks to my supervisor Adrian Kent for his insight and patience and for sharing his deep knowledge of the subject with me. Thanks also to Viswanathan Krishnan and Tzo Tze Ang for generously funding the studentship which made it possible for me to study at Cambridge.

I would like to thank my mother and Tom for constantly believing in me and for the many sacrifices they made to give me the opportunity to pursue my studies to the highest level.

And finally, thanks to Jamie for being beside me every step of the way. You are still the best person I know.

Chapter 1

Introduction

The twin revolutions of quantum theory and relativity in the early twentieth century led to an explosion of scientific progress whose implications have shaped not only modern physics but the entire modern world [19, 20]. And yet it is striking that these two contemporaneous theories have, to some extent, developed separately in parallel, coming into contact many times over the course of the century but never being fully reconciled [2, 21]. A unification of quantum theory with general relativity's account of gravity remains elusive, and is unclear whether it will ever be possible to achieve such a thing without profoundly altering the fundamental nature of at least one of the two theories [22]. We do now have formulations of quantum theory which are consistent with *special* relativity under appropriate conditions [23, 24], but there is still much to be learned about the way the two theories relate to one another.

This thesis will, in large part, be concerned with furthering the current understanding of the relationship between special relativity and quantum theory. We work principally in regimes where the theories are individually well understood and believed to be empirically consistent [25, 26], but where nonetheless interesting and surprising effects can result from their interaction. In this chapter, we will briefly introduce the basics of quantum mechanics and special relativity, then comment on some enduring difficulties.

1.1 Special Relativity

Special relativity, as initially introduced by Einstein in his beautiful 1905 paper [27], follows from two simple postulates:

1. The laws of physics are invariant in all inertial systems
2. The speed of light in a vacuum is the same for all observers, regardless of the motion of the light source

These simple ideas turn out to have a staggering range of consequences, including length contraction, time dilation, and mass-energy equivalence. But the consequence with which we will be most concerned in this thesis is the no-signalling principle, which states that no signal can travel faster than the speed of light.

Because the no-signalling principle will be the cornerstone of many of the results we introduce here, it is most perspicacious for us to regard special relativity as a theory about the causal structure of spacetime. Einstein's postulates lead naturally to a formulation of special relativity on a manifold of pointlike events with three spatial dimensions and one temporal dimension, equipped with an indefinite non-degenerate bilinear form (the Minkowski inner product) which defines a frame-invariant 'spacetime interval' between pairs of events in spacetime [28–30]. When the interval between two points is negative, or 'timelike,' there exists a frame of reference in which the corresponding events are simultaneous; when the interval is positive, or 'spacelike,' it is possible in principle for a signal travelling slower than the speed of light to travel from one event to the other; and when the interval is zero, or 'lightlike,' a signal can be sent from one to the other if and only if it travels at the speed of light. Two events can therefore stand in a cause-effect relation only if they take place at points separated by a spacelike or lightlike interval. Using this terminology we may define a past and future lightcone for every point in Minkowski space, where the *past lightcone* of x is the set of all points that are spacelike or lightlike related to x and in its past, and the *future lightcone* of x is the set of all points that are spacelike or lightlike related to x and in its future. Loosely speaking, the past lightcone contains all the points which could

possibly influence an event at x and the future lightcone contains all points that could possibly be influenced by an event at x [29, 30]. It is this causal structure which constitutes the empirical content of Minkowski spacetime, and indeed, it can be shown that the causal structure of Minkowski spacetime fully determines its topological and metrical structure [31], so one could equally well regard the causal structure as being the truly fundamental object of the theory.

For completeness, we recall that special relativity is now understood to be a limiting case of a more general theory which, aptly, is known as general relativity. Rather than the flat spacetime of special relativity, general relativity postulates a curved spacetime, with the curvature determined by the energy and momentum of the matter and radiation present in spacetime: this allows gravity to be described as a geometric property of spacetime [32]. The results described in this thesis are intended for use in standard terrestrial regimes where the approximation of flat spacetime is valid, and therefore we have not provided general relativistic formulations of our results. We note, however, that for a large class of physically reasonable general relativistic spacetimes,¹ the topological and metrical structure is determined up to a scale factor by the causal structure [31], and therefore since our theorems and proofs are couched directly in the language of causal structure, we conjecture that these results would still be valid in a large class of curved spacetimes, provided that participants in the protocols under consideration do not have the ability to manipulate the structure of spacetime. Of course all bets are off if one is faced with an adversary who is able to significantly alter the shape of spacetime during a protocol, but it seems safe to assume this will not be the case for the foreseeable future.

¹Specifically, those where the condition of *strong causality* is upheld - which is to say, loosely speaking, that every event has an arbitrarily small neighbourhood to which no causal curve returns after having left that neighbourhood. The spacetimes that fail to satisfy this condition are usually pathological in some way, and thus are not usually considered physically realistic [31]

1.2 Quantum Mechanics

1.2.1 Basic postulates

At the genesis of quantum mechanics, the new theory was met with resistance from many camps [33,34] - and understandably so, for the theory has a number of features which seem entirely at odds with common-sense intuitions about space, time and perhaps even the scientific process itself [35]. But quantum theory won out in the end, its empirical successes leaving its critics with little alternative but to concede defeat, although some never gave up hope that quantum mechanics might ultimately turn out to be derivable from a fundamental theory more consistent with their intuitions [36, 37].

Quantum mechanics is, in many ways, more a methodological prescription than a concrete scientific theory [38,39]: it sets out a mathematical framework for the construction of physical theories which must be supplemented with detailed experimental work to determine which specific mathematical objects represent the actual physical systems whose behaviour we would like to predict. However, in this thesis we will be more concerned with the abstract structure than with any specific realisation of it, and hence it is sufficient to regard quantum mechanics² as being characterised by the following four postulates [38]:

1. To every physical system we ascribe a Hilbert space \mathcal{H} known as the *state space* of the system.³ At any given time, the system is completely described by its *state vector*, which is a unit vector $|\psi\rangle$ in the state space.
2. Closed quantum systems evolve by unitary transformations.⁴ In particular, a closed quantum system can be associated with a fixed Hermitian⁵ operator H , and the time evolution of the state of the system is then given by $H|\psi(t)\rangle = i\hbar \frac{d|\psi(0)\rangle}{dt}$ (the *Schrödinger equation*).

²Throughout this thesis, we will use the terms ‘quantum mechanics’ and ‘quantum theory’ interchangeably to refer specifically to the non-relativistic theory.

³A Hilbert space is a complex vector space equipped with an inner product.

⁴A unitary transformation is a transformation that preserves the value of the inner product; unitary operators U satisfy $U^\dagger U = \mathbb{I}$, where U^\dagger denotes the conjugate transpose of U and \mathbb{I} denotes the identity operator.

⁵A Hermitian operator is defined as an operator which is equal to its own adjoint.

3. A measurement is described by a Hermitian operator M on the state space of the relevant system, with spectral decomposition $M = \sum_m m P_m$, where P_m is the projector onto the eigenspace of M with eigenvalue m . When a system is prepared in the state $|\psi\rangle$ and the measurement M is performed, the probability of obtaining the result m is equal to $\text{Tr}(P_m |\psi\rangle\langle\psi|)$, where $|\psi\rangle\langle\psi|$ denotes the outer product of the state vector $|\psi\rangle$ with itself; after this result has been obtained, the state of the system is $\frac{P_m |\psi\rangle}{\sqrt{\text{Tr}(P_m |\psi\rangle\langle\psi|)}}$.
4. When we combine two physical systems, the state space for the resulting composite system is the tensor product of the individual state spaces; if we combine n systems prepared in states $|\psi_1\rangle, |\psi_2\rangle, \dots, |\psi_n\rangle$, the resulting joint state is $|\psi_1\rangle \otimes |\psi_2\rangle \otimes \dots \otimes |\psi_n\rangle$.

1.2.2 Entanglement

Postulate 4 tells us that if we prepare the parts of a composite system independently and then simply combine them, the resulting joint state will be a tensor product of states on the individual parts of the system. But it is possible to produce other types of joint states by applying a unitary operation to or performing a measurement on some composite system, and in general, the result of such operations will be an *entangled* state, i.e. a state which can no longer be written as the tensor product of states on the individual subsystems. It seems that such composite systems have global properties that cannot be reduced to separate properties of the individual subsystems [38, 40, 41].

One might imagine that this feature is no more than a quirk of our choice of mathematical representation, but in fact, Bell showed in 1964 that the existence of entanglement has profound physical consequences [42, 43]. To do so, Bell studied ‘local hidden variable models’ - that is, models in which quantum systems have a set of ‘hidden’ properties in addition to their quantum state, and all correlations between measurements made at different spacetime points can be traced back to correlations between hidden variables that were established during a local interaction at some point in their common past. It is possible to prove that in any such model, given two systems S_a, S_b , and two different measurements $\{A, A'\}, \{B, B'\}$ for each of the two systems respectively, each of these four measurements

having two possible outcomes labelled by $+1$ or -1 , then if $\langle AB \rangle$ denotes the expectation value of the product of the two results when we perform measurements A, B on the two systems respectively, the following *CHSH inequality*⁶ must be satisfied:⁷

$$\langle AB \rangle + \langle AB' \rangle + \langle A'B \rangle - \langle A'B' \rangle \leq 2 \quad (1.1)$$

But there exist entangled quantum systems exhibiting correlations that violate this inequality, and therefore quantum mechanics cannot be fully explained by any local hidden variable model [44]. Moreover, the existence of these correlations has been verified by rigorous experiments successively eliminating larger and larger numbers of possible loopholes [45–47], and thus it seems likely that unless we wish to eschew a realist description of nature altogether, we will be forced to postulate some kind of ‘spooky action at a distance.’⁸ This is one of the features of quantum mechanics which has caused the most consternation over the last century [40, 41].

One might naturally wonder if the peculiar non-local character of quantum mechanics might violate the no-signalling principle we met in Special Relativity, but fortunately, it is not so! Direct conflict is averted by the quantum mechanical no-signalling theorem, which states that all operations on distinct quantum systems commute with each other [48], implying that although the outcomes of measurements on separate subparts of entangled systems can be correlated in a non-local way, no information can be transferred directly via these correlations. This theorem ensures that, despite the tension between quantum mechanics and

⁶Bell originally proved a related but different inequality, but the CHSH inequality originally proved by Clauser, Horne, Shimony and Holt [44] is more commonly used in modern studies, and hence we employ it here and throughout the thesis.

⁷This inequality applies to the case where the addition of the hidden variables makes the results of all measurements deterministic and also to the case where the hidden variables simply assign probabilities to outcomes.

⁸It should be acknowledged here that there are certain realist interpretations of quantum mechanics whose proponents would argue that this conclusion does not follow within their interpretation, most notably certain versions of the Everett approach; however, in most cases these interpretations save locality only by introducing something which is arguably just as counterintuitive, such as the postulation of multiple worlds in the Everett approach.

special relativity, the two do not make any predictions which are outright inconsistent. It is important to reinforce, however, that the quantum-mechanical and special relativistic no-signalling principles are not identical, and neither one implies the other: the relativistic principle does not rule out information transfer via measurements on entangled systems in cases where one measurement takes place in the future lightcone of the other, while the quantum-mechanical theorem has nothing to say about whether an object can physically travel faster than light.

1.2.3 Monogamy

One of the most intriguing properties of quantum entanglement is the fact that it is ‘monogamous’ - that is, there is a tradeoff between the amount of entanglement that a given quantum system can share with other systems. In particular, in chapter 5 we will make use of the fact that if a given system A is maximally entangled with some system B , then it cannot share entanglement with any other system [49].

The usual quantitative expression of the monogamy of entanglement employs the ‘concurrence’ (or ‘tangle’) to quantify the entanglement between various pairs of systems [50]. But while concurrence is a mathematically tractable measure of entanglement, its physical interpretation is not straightforward [51, 52], and therefore it is sometimes preferable to describe monogamy in a more operational way by stating it as a constraint not on entanglement itself but on the possible sets of correlations which can be obtained from measurements on quantum systems [52, 53]. This operational formulation also enables us to make comparisons to monogamy properties exhibited by other possible theories in the space of generalized probabilistic theories [54]: in particular, it is known that all non-signalling theories obey a monogamy bound ensuring that if some set of measurements (A, A') and (B, B') on two fixed systems S_a, S_b are capable of jointly violating the CHSH inequality as in equation 1.1, then no set of measurements (C, C') on any other system S_c combined with the *same* set of measurements (A, A') on system S_a can violate the analogous CHSH inequality 1.1 with B replaced everywhere by C [52]. Note the importance of using the same measurements on system A in both inequalities: without such a restriction we could trivially violate this monogamy bound in quantum theory by creating two maximally entangled

pairs $(\beta_1^1, \beta_2^1), (\beta_1^2, \beta_2^2)$, then naming β_2^1 as S_b , β_2^2 as S_c , and the combination of β_1^1 and β_1^2 as S_a , and measuring β_1^1 for the A, B inequality and β_1^2 for the A, C inequality.

It seems natural to hope that monogamy in quantum theory might be fully explained by the constraints arising from no-signalling. However, it can be shown that the monogamy bound on tripartite quantum correlations is stronger than the monogamy bound on non-signalling correlations, so the region of tripartite correlations which can be achieved by measurements on entangled quantum systems is actually smaller than the region which can be achieved in a general theory constrained only by no-signalling. [53] It remains an open question as to whether some other physical principle might explain the gap between the quantum and no-signalling bounds.

1.2.4 Mixed states and density operators

On the face of it, the postulates we have set out might seem to suggest that the state of a quantum system must always be describable by a state vector in some Hilbert space. But in fact there are two ways in which we can obtain different types of states within this formalism. First, we can draw a state from an ensemble of pure states $\{|\psi_i\rangle\}$ with associated probability distribution $\{p_i\}$, giving rise to a probabilistic mixture $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$; this is known as a proper mixture [55, 56]. Second, we can take two systems A, B in an entangled state ψ and then throw away the information about the state of A , leaving B in a reduced state $\rho = \text{Tr}_A(|\psi\rangle\langle\psi|)$; this is known as an improper mixture [55, 56]. Happily, it turns out that these two methods of preparation give rise to exactly the same type of mathematical object - to wit, a density matrix, a positive Hermitian matrix of trace one. Indeed, even pure states can be written as density matrices of the form $\rho = |\psi\rangle\langle\psi|$, with the set of pure states coinciding with the set of idempotent density matrices [38, 57, 58].

By applying the original four postulates consistently, we can derive analogous statements describing evolution and measurement for density matrices. The set of possible evolutions of a given system is now expanded to include all evolutions that can be obtained by appending an ancilla to the system, applying a unitary

evolution to the whole, and then tracing out the ancilla, which leads to the set of all completely-positive trace-preserving (CPTP) maps; the set of possible measurements on a given system is expanded to include all measurements that can be implemented by appending an ancilla to the system in question and applying a projective measurement to the whole, which leads to the set of all positive operator valued measures (POVMs), i.e. all sets of positive operators $\{K_n\}$ satisfying $\sum_n K_n = \mathcal{I}$, where the probability of obtaining the result indexed by n after performing the measurement $\{K_n\}$ on a system in the state ρ is given by $Tr(K_n\rho)$. Unfortunately there is no longer a unique formula for the post-measurement state after a POVM, because any given POVM can be implemented in a number of different ways and the post-measurement state depends on the particular implementation [38, 57, 58]. It should be noted that since the ideal of pure states, unitary operators and projective measurements cannot be perfectly realised in the laboratory, in real applications we are actually dealing with mixed states, CPTP maps and POVMs [59].

We pause at this point to reinforce that the pedagogical approach of deriving the existence of density matrices, CPTP maps and POVMs from the four postulates set out in section 1.2.1 - an approach known as the ‘Church of the Larger Hilbert Space,’ [60] - is not entirely uncontroversial. There are also advocates of the ‘Church of the Smaller Hilbert Space’ who contend that the density matrix should be thought of as the fundamental object of quantum mechanics and that we have no reason to suppose quantum systems cannot be in mixed states without being derived either from a larger pure state or a probabilistic mixture of pure states [61]. It is our view that a resolution to this disagreement must wait upon the solution to the problem set out in section 1.2.10, and thus we have adopted this mode of presentation simply for ease of exposition, without any intention of taking sides.

1.2.5 Quantum Contextuality

We have seen that in the most general picture, a quantum mechanical measurement is described mathematically by a POVM - that is, a set of positive semidefinite Hermitian operators on the Hilbert space of the measured system, where each

operator in this set is associated with one of the possible outcomes of the measurement. Prima facie it is tempting to imagine that measurement outcomes and hence also measurement operators are in one-to-one correspondence with underlying properties like ‘spin up’ which any quantum system must either have or not have, and this would suggest that it should always be possible to find a Non-Contextual Hidden Value (NCHV) model for any set of quantum measurements - that is, we should be able to assign either 1 or 0 to every measurement element (with 1 representing ‘the system has this property’ and 0 representing ‘the system does not have this property’), such that there is exactly one measurement element with the value 1 in every possible measurement, thus specifying the outcome that will be definitely be obtained if we perform this measurement on the given system. But a theorem due to Kochen and Specker tells us that there exist sets of quantum measurements for which it is impossible to find a NCHV model [62], and therefore it seems we cannot in general interpret quantum measurement elements in a straightforward way as giving us information about pre-existing properties of the world [63].

This original deterministic notion of contextuality has since been generalized by Spekkens [64], using the ontological models framework, where it is assumed that every system has a single real ‘ontic state,’ which determines the probabilities for the outcomes of any measurement on that system. An ontological model thus consists of a space Λ of ontic states λ , a set of probability distributions $\mu^P(\lambda)$ giving the probability that the system ends up in the state λ when we perform the preparation procedure P , a set of response functions $\vec{\xi}_{M,O}(\lambda)$ giving the probability that we obtain outcome O when we perform measurement M on a system whose ontic state is λ , and a set of column-stochastic matrices T^X representing the way in which the ontic state is transformed when some operation X is applied to the system.⁹ This allows us to say that a model is preparation non-contextual iff it represents every quantum state by a unique probability distribution $\mu(\lambda)$; trans-

⁹A column-stochastic matrix is a matrix whose entries are all non-negative and whose columns sum to one. Left-multiplication by a column-stochastic matrix is the most general possibility for the representation of quantum operations in an ontological model, since such an operation must map ontic state λ_i to λ_j with some probability, which is specified by entry (i, j) in the transformation matrix: since the entries are all probabilities, they must be nonnegative, and since each transformation must map λ_i to some state, the sum of the entries in a column must be one.

formation non-contextual iff it represents every possible quantum transformation by left multiplication by a unique transformation matrix; and measurement non-contextual iff it represents every possible quantum measurement M by a unique response function $\vec{\xi}_{M,X}(\lambda)$.¹⁰ Spekkens proved [64] that any ontological model of quantum mechanics which reproduces all the correct measurement statistics must exhibit preparation contextuality, but examples such as the Kochen-Specker model [65] show that this is not true for measurement contextuality.¹¹ In a different direction, the recent work of Cabello, Severini and Winter [66] has uncovered an intriguing connection between quantum contextuality and graph theory, and it has also been suggested that contextuality may be partly responsible for the computational power of quantum mechanics [4].

1.2.6 Quantum Information

The field of quantum information is the study of the storage, transmission and manipulation of information encoded in quantum systems, with a special focus on how these processes differ from comparable classical processes. Quantum information is an exciting and dynamic area of study, in part because it is a comparatively young field, first gaining momentum in the 1980s with the results of Bennett, Brassard, Wootters, Zurek, Jozsa, Deutsch, Feynman, Benioff and others [3, 38, 67–70]. It also has a number of important practical applications, encompassing the subfield of quantum cryptography, where quantum phenomena are employed to formulate improved protocols for traditional cryptographic tasks [71–77], and quantum computation, where quantum effects are deployed to achieve improvements on the best-known classical algorithms for some specific computational tasks [38, 68, 78, 79].

¹⁰The use of this definition of contextuality does not necessarily imply the endorsement of a description of quantum reality in terms of underlying states; the ontological models approach could simply be regarded as a helpful language in which we may express mathematical facts about the structure of quantum theory, such as its contextual character.

¹¹Transformation contextuality has not been studied so thoroughly, and most existing explicit ontological models omit transformations, so to our current knowledge the question of whether or not quantum mechanics necessarily exhibits transformation contextuality remains unresolved.

No-cloning

One of the cornerstones of quantum information theory is the no-cloning theorem, a result which will play a starring role in this thesis (in particular, see chapter 4). The theorem tells us that there can exist no quantum operation that (deterministically or probabilistically) creates perfect copies of an arbitrary unknown quantum state; this is a straightforward consequence of the linearity of quantum operations, and since it is really quite remarkable that such a profound and far-reaching result can be proven in just a few simple lines, we feel compelled to rehearse the proof here [3]:

Proof. Suppose there exists some unitary operation U that is capable of cloning an arbitrary quantum state: that is, there exists some fixed state $|\phi_o\rangle$ such that for any $|\psi\rangle$, we have that $U(|\psi\rangle \otimes |\phi_o\rangle) = |\psi\rangle \otimes |\psi\rangle$.

Thus for any distinct states $|\psi_1\rangle, |\psi_2\rangle$, we have that $U(|\psi_1\rangle \otimes |\phi_o\rangle) = |\psi_1\rangle \otimes |\psi_1\rangle$ and $U(|\psi_2\rangle \otimes |\phi_o\rangle) = |\psi_2\rangle \otimes |\psi_2\rangle$.

Now consider applying U to an equal superposition of the states $|\psi_1\rangle$ and $|\psi_2\rangle$. By linearity, $U((\frac{1}{\sqrt{2}}(|\psi_1\rangle + |\psi_2\rangle)) \otimes |\phi_o\rangle) = \frac{1}{\sqrt{2}}(|\psi_1\rangle \otimes |\psi_1\rangle + |\psi_2\rangle \otimes |\psi_2\rangle)$.

But from the definition of the cloning operation U , we ought to have $U((\frac{1}{\sqrt{2}}(|\psi_1\rangle + |\psi_2\rangle)) \otimes |\phi_o\rangle) = \frac{1}{\sqrt{2}}(|\psi_1\rangle + |\psi_2\rangle) \otimes (|\psi_1\rangle + |\psi_2\rangle)$.

Since $\frac{1}{\sqrt{2}}(|\psi_1\rangle \otimes |\psi_1\rangle + |\psi_2\rangle \otimes |\psi_2\rangle) \neq \frac{1}{\sqrt{2}}(|\psi_1\rangle + |\psi_2\rangle) \otimes (|\psi_1\rangle + |\psi_2\rangle)$ we have obtained a contradiction, and therefore no such cloning operation U exists. □

To be more precise, it can be shown that deterministic cloning (where two perfect copies are produced with probability one) is possible only for sets of mutually orthogonal quantum states [80], while probabilistic cloning (where two perfect copies are produced with probability strictly greater than zero but less than one) is possible only for sets of linearly independent quantum states [81]. Moreover, although in this thesis we will principally be concerned with cloning pure states, the result can also be generalized to mixed states: the ‘no-broadcasting’ theorem states that there is no quantum operation E that maps states on a Hilbert space \mathcal{H} to states on a Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ such that for any state ρ ,

$\text{Tr}_A(E(\rho)) = \text{Tr}_B(E(\rho)) = \rho$. We refer interested readers to the relevant literature for the proof of this theorem [82].

Quantum teleportation

We now turn to another stalwart of quantum information theory which will be an important tool for several of the protocols described in later chapters. *Quantum teleportation* [67] is a simple but powerful phenomenon which allows a generic quantum state ψ to be transmitted between two cooperating agents (traditionally named ‘Alice’ and ‘Bob’) who share a pair of maximally entangled particles. The process can be succinctly represented via the following identity, where $\Psi_{00}, \Psi_{01}, \Psi_{10}, \Psi_{11}$ are the four maximally entangled Bell states [38]:

$$|\psi\rangle_A \otimes |\Psi_{00}\rangle_{BC} = \frac{1}{2} \sum_{z,x \in \{0,1\}} |\Psi_{z,x}\rangle_{AB} \otimes X^x Z^z |\psi\rangle_C \quad (1.2)$$

This expression tells us that the combination of system A in state ψ and systems B, C in the maximally entangled state $|\Psi_{00}\rangle$ is mathematically identical to a linear superposition of four states of the form $|\Psi_{z,x}\rangle_{AB} \otimes X^x Z^z |\psi\rangle_C$, with z, x taking values in $\{0, 1\}$. Thus suppose that Alice is in possession of system A in the state ψ , and Alice and Bob also share a maximally entangled pair, with Alice in possession of system B and Bob in possession of system C . Alice can perform a measurement on systems A and B which has four possible outcomes, $\{|\Psi_{z,x}\rangle : z, x \in \{0, 1\}\}$, corresponding to the four branches of the superposition in equation 1.2. We know from postulate 3 that after this measurement, if the result obtained is $z, x \in \{0, 1\}$, the entire collection of three systems will be in the state $|\Psi_{z,x}\rangle_{AB} \otimes X^x Z^z |\psi\rangle_C$, and therefore when Bob learns Alice’s measurement result he can apply the inverse of the relevant operator $X^x Z^z$ to system C to recover the state ψ [67].

Now, it may seem that we are in imminent peril of violating either the no-cloning or no-signalling theorem with this procedure, but fear not. Teleportation does not achieve cloning, because system A is no longer in the state ψ at the end of the protocol. Nor does it achieve signalling, because immediately after

Alice performs her measurement, the state of system C is, from Bob's point of view, an average over the four possible measurement results; this average turns out to be indiscernible from the original state, so no information is obtainable by Bob until after he has received Alice's classical message. Nonetheless, it is true that something seemingly miraculous has happened: a quantum state, which contains an effectively infinite amount of information, has been transmitted by means of only a two-bit classical message. The reason that this is possible is that entanglement itself is a resource, and it is only by consuming this resource that we have achieved the transfer of the quantum state [38]. Moreover, although a quantum state in principle contains an infinite amount of information, only a limited amount of information can be obtained from it: Holevo's theorem puts a precise bound on the amount of information about a quantum state that can be retrieved by measurements on that state, and in the case of a qubit this limit turns out to be only one bit [38, 83].¹² Therefore in practice only a finite amount of information can be transferred in this process, so teleportation is not quite as paradoxical as it might first appear.

The von Neumann entropy

Our discussion now turns to a particularly important mathematical tool in the study of quantum information. To provide context we will first introduce a related classical concept known as the Shannon entropy,¹³ a measure of the 'uncertainty' associated with a probability distribution $\{p_i\}$, and hence also of the average information content obtained when one learns the value of a random variable A described by this probability distribution [85]. The Shannon entropy of a distribution $\{p_i\}$ associated with a random variable A is defined as $H(A) := \sum_i -p_i \log(p_i)$,

¹²We pause to note that it is possible to improve on this somewhat if entanglement is involved - for example, in a technique known as superdense coding, Alice and Bob share a pair of entangled qubits, and Alice sends a message to Bob by performing one of four operations on her qubit and then sending it to Bob, who can then determine which operation Alice performed and thus extract two bits of information from a single qubit. Superdense coding is thus, in a sense, the inverse of teleportation: rather than sending one quantum bit with two classical bits and some entanglement, we send two classical bits with one quantum bit and some entanglement. [38]

¹³The von Neumann entropy is, in fact, older than the Shannon entropy [84], and the related concept of thermodynamical entropy predates them both [84], but we present the von Neumann entropy as an analogue of the Shannon entropy rather than vice versa because this presentation sheds some light on the interpretation of the two concepts.

where the logarithm is used to ensure that the entropy associated with two independent processes is additive. Indeed, it can be shown that the Shannon entropy is the unique functional satisfying five conditions that we might reasonably expect for a measure of information content (continuity, unitary invariance, normalisation, additivity, and arithmetic mean [86]) and hence it is not just an arbitrary choice of measure but is, in a well-defined sense, the only possible such measure.¹⁴

The von Neumann entropy of a quantum system in the state ρ is defined as $S(A)_\rho := -\text{Tr}(\rho \log(\rho))$ [38]. This quantity is often treated as a quantum analogue of the Shannon entropy - and indeed, the von Neumann entropy is the unique functional that satisfies the quantum generalization of the five conditions defining the Shannon entropy [87], so the analogy is well founded. The relationship is reinforced by the fact that the von Neumann entropy reduces to the Shannon entropy of the probability distribution over measurement outcomes for density matrices that are diagonal, i.e. that are essentially classical systems [88].

Various derived quantities can be defined in terms of the Shannon entropy, and analogues can be defined using the von Neumann entropy. One such quantity is the mutual information. The classical mutual information $I(A : B)$, which quantifies the amount of information one obtains about variable A by learning the value of variable B or vice versa, is defined as [85]:

$$I(A : B) := \sum_{x,y} p(A = x, B = y) \log \frac{p(A = x, B = y)}{p(A = x)p(B = y)} = H(A) + H(B) - H(AB)$$

The quantum mutual information between two quantum systems, which can be interpreted similarly, is defined by analogy as [38]:

$$I(A : B)_{\rho_{AB}} := S(A)_{\rho_{AB}} + S(B)_{\rho_{AB}} - S(AB)_{\rho_{AB}}$$

Another useful quantity is the conditional entropy. The classical conditional

¹⁴However, we note in passing that it is possible to relax the ‘arithmetic mean’ property to give a family of entropies known as the ‘Renyi entropies,’ [86] which all have interesting and important properties, but which will not be employed in this thesis.

entropy of a random variable A relative to a random variable B , which can be interpreted as the amount of information needed to describe the outcome of a random variable A when the value of B is known, is defined as [85]:

$$H(A|B) := \sum_{a,b} p(a,b) \log\left(\frac{p(b)}{p(a,b)}\right) = H(A,B) - H(A)$$

The quantum conditional entropy of a quantum system A relative to a quantum system B is defined by analogy as [38]:

$$S(A|B)_{\rho_{AB}} := S(A,B)_{\rho_{AB}} - S(B)_{\rho_{AB}}$$

The interpretation of the quantum conditional entropy is less straightforward because, unlike the classical conditional entropy, it can be negative. This property is perhaps best understood via the operational interpretation of the quantum conditional entropy: if two agents share a set of bipartite systems A^i, B^i , each having overall state ρ_{AB}^i , then in the limit of a large number of systems, $S(A|B)$ quantifies the total amount of quantum information per state that Alice must send to Bob in order for him to be able to reconstruct the full state ρ_{AB}^i . When $S(A|B)$ is negative, this means that classical communication alone is sufficient for Bob to reconstruct the state, and furthermore, Alice and Bob will subsequently have the ability to transfer additional quantum information at no cost to them (because they will produce maximally entangled pairs which can be used to perform quantum teleportation) [89].

1.2.7 Quantum Cryptography

Classical cryptography is the study of secure communication in the presence of adversaries. In particular, the traditional task of classical cryptographers was to combine a message with some other information known as a ‘key’ in such a way that the message cannot be read by anyone who does not know the key; the resulting cryptosystem is ‘symmetric’ if both encoder and decoder use the same key,

and ‘asymmetric’ otherwise. Unfortunately no asymmetric cryptosystems have been shown to be secure without unproven assumptions, and while there do exist secure symmetric protocols, these have the disadvantage of relying on the possibility of generating and sharing a secure random key between participants, a cryptographic problem in and of itself for which, again, no classical protocol is known to be secure without unproven assumptions [90, 91]. Beyond these tasks, the field of classical cryptography has also expanded to include the study of other related types of secure computation and communication, including tasks such as bit commitment and zero-knowledge-proving which will play an important role in this thesis.

The idea of using quantum mechanics to obtain cryptographic advantages first arose in the 1970s with the work of Wiesner, Bennett and Brassard [90, 92, 93], when it was observed that a number of properties of quantum physics lend themselves particularly well to cryptographic applications. For example, the fact that generic quantum measurements disturb the state of the measured system gives us a means of checking whether someone has accessed or attempted to access information stored in the state of a quantum system, and this property is leveraged in the first major quantum cryptographic protocol, the BB84 quantum protocol for key distribution developed by Bennett and Brassard in 1984 [93].

The BB84 protocol requires two participants, traditionally named Alice and Bob. Denote by $\{\psi^{00}, \psi^{01}\}$ two states that form an orthogonal basis for a qubit, and let M_0 be a projective measurement in this basis; denote by $\{\psi^{10}, \psi^{11}\}$ two distinct states that also form an orthogonal basis for a qubit, and let M_1 be a projective measurement in this basis. To begin the protocol, Alice generates two random bit strings w and x of length N , then prepares a set of N qubits $\{Q_i\}$, with the state of Q_i given by $\psi^{w_i x_i}$. She then sends these N qubits to Bob, who generates a random bit string y of length N and measures all N qubits, performing measurement M_{y_i} on system Q_i . Bob thus generates a bit string z of length N containing the results of his measurements. Bob then announces the string y publicly and Alice tells Bob the set of indices S on which y matches w ; Alice forms a string x' by removing all the bits of x in positions with indices not belonging to S , and Bob likewise forms a string z' by removing all bits of z in positions with indices not belonging to S . Alice and Bob then randomly select some subset S' of the indices

in S and publicly compare the values of x' and z' on all indices in S' .¹⁵ If any of the values do not match, Alice and Bob abort the protocol¹⁶. Otherwise, they form strings x'' and z'' by dropping all the bits in positions with indices belonging to S' ; with high probability x'' and z'' will be identical and hence can be used as a shared secret key. In heuristic terms, the security of this protocol is derived from the fact that even if an eavesdropper intercepts and measures the quantum systems sent by Alice, the measurements will with high probability disturb the states of a significant proportion of these systems, and thus with high probability the comparison of bits in x' and z' will result in the protocol being aborted. Moreover, eavesdroppers cannot get around this by making copies of the states of the systems they intercept and measuring the copies instead, because the no-cloning theorem prohibits reliable copying of unknown quantum states.

The BB84 protocol represented a crucial advance for cryptography in that it was the first *unconditionally secure* key distribution protocol, meaning that its security depends only on the structure of quantum theory and on appropriate assumptions about the physical properties of the devices employed in the protocol [94, 95]. In particular, we do not have to assume anything about the hardness of certain computations - which is fortunate, because something else quantum mechanics is good for is speeding up certain types of computations, with the consequence that a number of classical cryptographic protocols whose security depends on computational complexity assumptions become insecure when one is faced with an adversary who could potentially have access to a quantum computer [96].

There have been many developments since BB84 was first proposed. Although we have endeavoured to give a simple intuition as to why BB84 might be secure, a fully rigorous security proof was not available until around 2000, when a series of proofs were produced by Mayers, Lo and Chau, Shor and Preskill, and Biham et al. [94, 97–99] Moreover, our simplified description of the protocol assumed zero probability of errors in Alice's preparations and Bob's measurements, but of

¹⁵For example, they could do this by computing the XOR of randomly selected pairs of bits and announcing the results

¹⁶This description assumes there is no probability that the preparation or measurement devices will make errors; if the possibility of errors is taken into account, Alice and Bob may have to allow some small number of mismatched bits.

course in real applications errors may arise, and error correction strategies have been developed to deal with this possibility [90]. There are also techniques known as *privacy amplification* designed to further lower the amount of information that an eavesdropper could potentially have about the secret key [90, 100]. Many variations on BB84 have been developed, such as Bennett's two-state variant [101] and a version due to Ekert which derives its security from the monogamy of entanglement [71].

More generally, it has been shown that quantum mechanics can also be applied to many other interesting cryptographic tasks. Another important feature of quantum mechanics is that it exhibits what many physicists currently believe to be genuine intrinsic randomness; this is useful for cryptographic purposes, because many cryptographic protocols (including BB84!) require the parties to generate random numbers that are secure in the sense that they cannot be predicted or guessed by eavesdroppers and/or other parties in the protocol. Several protocols have been developed to use quantum mechanics for (biased) coin-tossing [102, 103], and it is also possible to do quantum-mechanical randomness expansion, where a small random seed is turned into a much larger string of random numbers in an unconditionally secure way [103, 104]. Another useful technique which will see a brief cameo in this thesis is quantum secret sharing. In classical cryptography, (m, n) -secret sharing involves splitting a message into a number n of parts so that the message can be read by anyone in possession of at least m parts but someone who comes into possession of fewer than m parts will gain no information about the message. This protocol can be implemented classically but is vulnerable to eavesdropping, a problem which can be remedied by moving to (m, n) -quantum secret sharing, where a quantum state is divided into n shares such that the state can always be reconstructed from any number of shares greater than or equal to m , but any number of shares less than or equal to m gives no information at all about the state [105]. Note that in the quantum case, we must always have $m > \frac{n}{2}$, otherwise we would be able to violate the no-cloning theorem.

Thus quantum cryptography has become a large and active field, and it is now known that many cryptographic tasks which were once believed impossible can in principle be implemented using quantum techniques. On the other hand, quantum cryptography cannot solve *every* outstanding problem in classical cryptog-

raphy. For example, there is an important primitive of classical cryptography known as bit commitment which cannot be implemented with unconditional security in a classical context in non-relativistic Galilean spacetime; in the early days of quantum cryptography it was hoped that quantum techniques might enable us to perform unconditionally secure bit commitment, but it was eventually shown by Mayers [8], and Lo and Chau [14], that this is also impossible in the quantum context if we still assume that the protocols take place in non-relativistic Galilean spacetime. Thus it is necessary to look elsewhere for solutions, and later in this thesis we will do exactly that.

1.2.8 Quantum Field Theory

Because particles in scattering experiments are frequently accelerated almost to the speed of light, it is not possible to neglect relativistic effects in the description of scattering experiments, and therefore in order to do particle physics accurately it is necessary to make some adjustments to standard quantum mechanics. There exists a relativistic formulation of quantum mechanics where the Schrödinger equation is replaced by the Klein-Gordon and Dirac equations, but it turns out that this is not sufficient to allow us to study particle physics, because both non-relativistic and relativistic quantum mechanics are defined only for scenarios that can be described by a finite, constant number of degrees of freedom, whereas in particle physics it is necessary to describe fields with an infinite number of degrees of freedom and scattering processes in which particles may be created or destroyed. Thus in order to apply quantum mechanics to this realm it has been necessary to create an extension of quantum mechanics accommodating an infinite number of degrees of freedom. The extension, known as quantum field theory, is now able to model successfully almost all features of elementary particle physics - gravity alone, amongst the fundamental forces, still resists being cast in this form [23,24].

One might reasonably enquire at this juncture as to why the results of this thesis are couched within the framework of standard quantum mechanics rather than relativistic quantum mechanics or quantum field theory, both of which come with special relativity already built into their equations. However, we contend that stan-

Standard quantum mechanics is the most appropriate setting for our results, because the protocols we describe are operational in nature, characterising processes that take place in regimes known to be well-described by non-relativistic quantum mechanics. It would certainly be possible in principle to produce analogues of our proofs using relativistic quantum mechanics or quantum field theory, but the calculations would be inelegant, and are unnecessary given that quantum field theory and relativistic quantum mechanics are known to reduce to non-relativistic quantum mechanics in the low-energy, macroscopic limit. Furthermore, we note that although the ‘heuristic’ formulation of quantum field theory is well-developed, we are still lacking a mathematically precise formulation of it for interacting particles in $3 + 1$ dimensions [106, 107]. This fact, together with continuing difficulties over the treatment of gravity, gives good reason to be cautious about the existing framework of quantum field theory, and motivates us to continue working in a framework which does not presuppose it. Indeed, where possible in this thesis we even avoid presupposing the details of non-relativistic quantum mechanics and instead base our proofs on simple, general principles to ensure that our results have maximum generality. For example, the security proofs in chapter 8 depend *only* on the no-signalling principle, and since this is a highly robust physical principle which appears under various guises not only in standard quantum mechanics but also in relativity and quantum field theory, protocols whose security depends only on this principle are likely to remain secure even if both quantum mechanics and quantum field theory are eventually replaced by a more fundamental theory.

For clarity, we note that it is possible to derive analogues of many results from quantum information theory within quantum field theory, and some of these can even be extended in a qualified way to curved spacetime. The study of such extensions is a rapidly developing field known as relativistic quantum information [108–111]. In this thesis, however, our attention is confined to standard quantum mechanics applied on flat Minkowski spacetime, so while this work could be said to fall under the purview of ‘relativistic quantum information,’ it is somewhat distinct from the main concentration of research activity in the field.

1.2.9 Quantum Gravity

As we have noted, there exist underlying tensions between quantum mechanics and relativity - both special and general - and combining the two has turned out to be a remarkably difficult problem, with the quantization of gravity being a particularly intractable problem [22]. Nonetheless, a number of promising avenues toward ‘quantum gravity’ have been proposed, including loop quantum gravity [112] and string theory [113], which are the most prominent contenders at the present time. Both of these approaches attempt to resolve the tension by altering relativity in some way - for example, loop quantum gravity quantizes spacetime by deriving it out of an underlying quantum substratum, while string theory turns gravity into a mode of an oscillating string.

Quantum gravity is chiefly a problem about the behaviour of matter at very small scales and very high energies, regimes where we anticipate either quantum mechanics or special relativity may break down and therefore it is entirely up in the air as to which, if either, will ultimately prevail. By contrast, in this thesis we will be focusing on regimes in which the two theories are very well-defined and, modulo some lingering doubts related to the measurement process, seem to be consistent with one another. Thus, although we will be intimately concerned with the relation between quantum mechanics and special relativity, we will not be engaging with the disputes over quantum gravity.

1.2.10 The Measurement Problem

By many measures, quantum mechanics (or quantum field theory) is one of the most successful scientific theories ever known - and yet there is trouble in paradise. We now pause to consider a serious conceptual problem in the foundations of quantum theory which is yet to find a universally satisfactory resolution.

The third of the postulates of quantum mechanics (sometimes known as the ‘Born rule’), together with its generalisation to POVMs, tells us how to extract precise quantitative values for the probabilities with which various measurement outcomes will be obtained when a quantum measurement is performed, and these predictions have been verified to a very high degree of accuracy; however, the axiomatic formulation of quantum mechanics is conspicuously silent about the

physical interpretation of these probabilities. Several questions about the nature of quantum probabilities thus present themselves. First, quantum mechanics does not easily lend itself to an interpretation in which the probabilities can be understood as ignorance probabilities describing relative frequencies of occurrence for some underlying set of hidden variables [114], and therefore it is common to claim that quantum mechanics must be ‘inherently probabilistic,’ meaning that the probabilities it assigns are fundamental features of the theory which cannot be further analysed [115–117]. But this claim is a very extreme one: if true, it would mean that quantum probabilities are unique in nature, probabilities unlike any other forms of probability that we have encountered, and it is unclear that the evidence really justifies this step. Second, even if we put aside the interpretation of probability itself, there remain unanswered questions about what exactly the Born rule probabilities are probabilities *for*. Some schools of thought argue that they describe nothing other than probabilities for observers to make observations [118], but this is unpalatable from the point of view of realism about science and also seems to stymie any attempt to achieve greater understanding of quantum measurement.

Of realist views, two broad categories of response can be distinguished: ‘ ψ -ontic’ views, which take the quantum state to be an element of reality, and ‘ ψ -epistemic’ views, which regard the quantum state as merely a description of an observer’s knowledge of reality. ψ -ontic approaches may be further divided into ψ -complete models, where there is a one-to-one relation between quantum states and real (ontic) states, and ψ -supplemented models, where the ontic state space is parametrised by the quantum state together with some other supplementary variables (see Harrigan and Spekkens, 2007 [65]). An example of the former is the Everett interpretation [119], which claims that all outcomes of a measurement exist in separate branches of the wavefunction, and the Born rule probabilities (very loosely speaking!) may be thought of as the probabilities for an observer to find himself in one branch or another [119, 120]; an example of the latter is the de-Broglie-Bohm approach which postulates a quantum state that acts as a ‘pilot-wave’ for the set of particles on which the actual physical world of our experience supervenes, causing them to occupy positions with probabilities corresponding to the Born rule probabilities for those positions [37, 121]. There also exist a

number of ψ -ontic views that invoke some form of wavefunction collapse, the idea being that most of the time quantum states evolve unitarily, but on occasion - perhaps at random [122–125], perhaps due to measurement [126] or perhaps due to a gravitational interaction [127] - they undergo a non-unitary collapse into one out of a range of possible states, with probabilities for the various possible states corresponding to the Born rule probabilities. Some such models are ψ -complete, while others may arguably be regarded as ψ -supplemented.

Well-developed, realist ψ -epistemic views seem harder to come by, but perhaps the best example is Quantum Bayesianism, which claims that the quantum state represents nothing objective at all - it is merely a measure of the observer's subjective degrees of belief in various outcomes of the measurement [128]. Although the view is anti-realist about the quantum state itself, its proponents maintain that it should be nonetheless understood as a form of 'participatory realism,' [129]: we are to remain realists about physics while accepting that the world may not admit of a straightforward realist description [130].

We will largely put aside such questions of interpretation in this thesis, although we will briefly return to these matters in our concluding remarks.

Chapter 2

Context

Having set out a general picture of the state of play in quantum mechanics and special relativity, we now describe the more immediate context of the work presented in this thesis. The first part of the thesis draws on a long-standing tradition of advancing our understanding of quantum mechanics using games and tasks, as well as via the resolution of apparent paradoxes. In the second part of the thesis we move to relativistic quantum cryptography, where we will be particularly concerned with protocols for bit commitment and zero-knowledge proving.

2.1 Tasks and Games

The study of information-theoretic games and general operational tasks has been an important tool of quantum information theory ever since the birth of the field [131–135]. Practically speaking, the study of games is useful because it offers a precise way of quantifying the potential cryptographic advantages offered by quantum theory, and on the more theoretical side it improves our understanding of the difference between quantum and classical physics and may help us pin down the source of these quantum cryptographic advantages [136–138].

Perhaps the first major work on quantum game theory was Meyer’s 1999 analysis of quantum strategies [139], and since then his operational methodology has been employed in a wide range of contexts. It is of course very natural to parse cryptographic protocols in the form of a game, with the communicating parties

on one side and potential eavesdroppers on the other [71], but there are also many less obvious applications: it has been shown that quantum cloning [140] can usefully be cast as a game with a single party playing against nature itself [141], and that quantum computing algorithms can be understood as a game with classical agents on one side and quantum agents on the other [142]. A number of interesting classical games, such as the Prisoner's Dilemma [138, 143, 144], the battle of the sexes [145, 146] and the Monty Hall problem [147, 148], have been generalised to the quantum realm and implemented experimentally [144, 144, 149], and optimal strategies have been derived for general two-player games [138] and multiplayer games [150].

However, *relativistic* quantum games and tasks are still comparatively unexplored. The first example of a relativistic quantum task was *summoning*, introduced as a simple illustration of an operational task that distinguishes relativistic quantum theory from both relativistic classical theories and non-relativistic quantum theory [151]. In its original form, a summoning task involves two agencies, 'Alice' and 'Bob', each comprising collaborating networks of agents distributed throughout spacetime: Bob secretly prepares a random quantum state in some agreed Hilbert space and gives it to Alice at a point P , and at some later point, not originally known to Alice, Bob then asks Alice to return the state. To derive constraints on Alice's ability to respond successfully, we must employ results from both special relativity and quantum theory. Specifically, the (relativistic) no-signalling principle and the (quantum) no-cloning theorem together imply that no matter how densely Alice's agents are distributed, in general there will be no strategy which guarantees a successful response to Bob's request. This result is known as the 'no-summoning theorem' [151], and the result remains true for variations of the task in which time delays in returning and some loss of fidelity in the returned state are allowed [151]. Note that it is the relativistic version of no-signalling - the prohibition on faster-than-light information transfer - which is needed here: the quantum no-signalling theorem alone is not sufficient to derive the result, and therefore the no-summoning theorem is indeed intrinsically both quantum and relativistic.

Summoning was subsequently generalised by Hayden and May (HM) who introduced a version of the task defined by a spacetime point P_S and a set of N

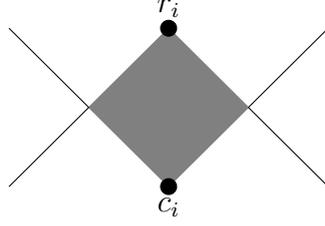


Figure 2.1: The causal diamond associated with the call-response pair (c_i, r_i) .

ordered pairs of spacetime points (c_i, r_i) such that every point r_i is in the future lightcone of the corresponding point c_i [7]. In this task, Alice is given a quantum system Q_B in some unknown state ψ at point P_s , and if a call is subsequently made at point c_i , Alice must return a quantum system in state ψ at the corresponding response point r_i . Let us say that a spacetime task is *feasible* if there exists any protocol that is possible according to special relativity and quantum mechanics such that if Alice employs this protocol, then she will be able to complete the task with a one hundred percent success rate. HM found an intriguing characterisation of the necessary and sufficient conditions for a given task to be feasible. Working in Minkowski space, write $x > y$ if the spacetime point x is in the causal future of y , and $x \geq y$ if either $x > y$ or $x = y$, then define the *causal diamond* D_i to be the set $\{p : r_i \geq p \geq c_i\}$, as shown in fig 2.1. Then, by using iterative applications of quantum teleportation and secret sharing, it can be shown that:

Theorem 1. [7] Consider a summoning task defined by a start point P_S and a set of N ordered call-response pairs (c_i, r_i) , such that $\forall i r_i \geq c_i$, in which Alice is given a quantum system Q_B in some unknown state ψ at P_S , and if a call is subsequently made at exactly one call point c_i , Alice must return a quantum system in the state ψ at the corresponding response point r_i . The task is feasible iff the following conditions hold:

1. Every response point $r_i \geq P_s$.
2. Every pair of causal diamonds D_i and D_j are causally related, meaning that there exists $x_i \in D_i$ and $x_j \in D_j$ with $x_i \geq x_j$, or vice versa.

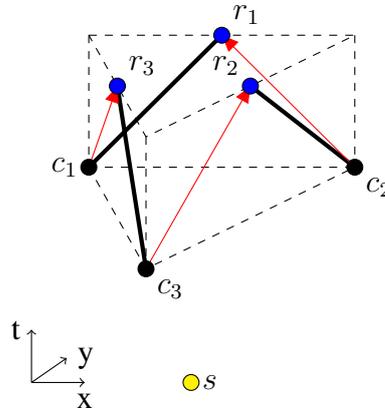


Figure 2.2: A $2 + 1$ dimensional example, taken from Ref [7], which is feasible even though there does not exist a continuous spatiotemporal path that starts from s and runs sequentially through each of the causal diamonds. The black lines are lightlike and in this case they represent the entirety of the causal diamond for each pair; the red arrows are also lightlike. We have that $c_1 < r_3, c_2 < r_1$ and $c_3 < r_2$.

These conditions for feasibility are considerably weaker than naive intuition might suggest. In particular, a task can be feasible even if there does not exist a continuous spatiotemporal path that starts from s and runs sequentially through the causal diamonds - fig 2.2 shows an example of such a task. This means that it is sometimes possible to complete a summoning task even if it is not possible to complete a comparable classical task in which an agent is given a single copy of a physical object that they cannot duplicate and is required to return it at the appropriate call point in response to the occurrence of a call. Consequently, summoning may indeed be regarded as a phenomenon that is both intrinsically quantum and intrinsically relativistic.

In addition to its theoretical interest, this summoning task has important practical applications. In particular, many distributed quantum computations over networks where relativistic signalling constraints are significant may usefully be modelled as summoning tasks: one can imagine quantum data that needs to be routed to one of a number of nodes, with the destination depending on classical data generated at other nodes during the computation. Indeed, in the teleportation model of distributed quantum computation [152], each round of adaptive compu-

tation is essentially a summoning task: the measurement result from the previous round determines the measurement to be made in the present round, and thus plays the role of the ‘call’, while the locations of the gates for the various possible measurements play the role of the ‘response points’.

However, HM’s definition of a summoning task is only one of many possible versions of the original summoning task. It is likely that other formalisations might be equally relevant, or even more relevant, to the types of practical situations that are usefully modelled as distribution tasks. These other formalisations might also shed new light on theoretical issues in the intersection between quantum mechanics and special relativity. Thus in the first part of this thesis we study two different generalisations, one where response points are not required to lie in the future lightcone of the associated call point and one where Bob can make multiple calls; our results lead us to question the interpretation proposed by HM for their theorem. We also define two new distribution tasks which require Alice to distribute *entangled* states, and prove some preliminary results about these tasks.

2.2 Paradoxes

There is a long-standing tradition of using apparent paradoxes to refine our understanding of quantum theory - an incomplete list of recent examples includes the quantum Cheshire Cat, the quantum Zeno paradox, the quantum Gibbs paradox, and the quantum pigeonhole effect [153–158]. This venerable tradition might be said to have begun with Einstein and his co-authors Podolsky and Rosen, who set out what has become known as the ‘EPR Paradox’ - the supposed paradox being that if distinct quantum systems satisfy separability and locality, quantum mechanics cannot be complete [33, 159]. Of course, the EPR effect is no longer regarded as a paradox now that we have accustomed ourselves to the realities of quantum physics, and indeed, the same is true of all the paradoxes we have listed above - with familiarity, the appearance of paradox greatly decreases and perhaps vanishes altogether. This is to be expected, because it is generally agreed that there are no actual *physical* paradoxes, if that term is used in the strong sense where a paradox is not merely puzzling but literally a contradiction: all apparent physical paradoxes arise from misunderstandings and/or attempts to extend a concept too

far beyond its realm of applicability.

Of course, this is precisely the reason why the study of paradoxes can be such a useful way to refine our understanding of a physical theory. Out of the furore over the supposed EPR paradox arose our modern understanding of entanglement [159, 160]; from the Cheshire cat and pigeonhole effects we can obtain proofs of contextuality [161]. These supposed paradoxes were signs that there was something wrong with our classical way of thinking and in resolving them we have developed new ways of thinking about the quantum world.

Likewise, there exist a number of relativistic ‘paradoxes’ that have been used to test the limits of special relativity and deepen our understanding of its implications - the ‘twin paradox’ is perhaps the best known, but the ‘ladder paradox,’ and ‘Bell’s spaceship paradox’ remain firm favourites of undergraduate lecturers and have had their part to play in the history of relativistic thought [162–164]. As in the quantum case, these paradoxes can be resolved by careful treatment and proper understanding, and thus they too cease to look like paradoxes once one is sufficiently familiar with the theory.

In light of the role played by paradox in advancing theoretical understanding, it is natural to wonder if we might similarly improve our understanding of the relationship between quantum theory and special relativity by uncovering apparently paradoxical effects that arise from the interaction between the two theories. In chapter 4 we demonstrate the existence of such a paradox; to our knowledge it is the first genuinely relativistic quantum paradox, in the sense that it can be formulated only in the framework of relativistic quantum theory. Our paradox does indeed seem to have implications for the current understanding of the nature of quantum states as spatiotemporal objects, and after presenting the effect we consider these implications carefully.

2.3 Relativistic Quantum Cryptography

The study of relativistic quantum tasks and games reveals that combining relativity with quantum mechanics can lead to surprising and counterintuitive phenomena that cannot be reproduced in either of these theories individually. It is this fact which is the starting point for the field of relativistic quantum cryptography, where

the combined power of relativistic causality and quantum information theory is exploited to achieve cryptographic advantages over classical cryptography or standard quantum cryptography. A variety of interesting tasks (e.g. [77, 165–169]) are now known to be achievable within relativistic quantum cryptography, either with unconditional security or with security significantly enhanced relative to classical protocols. On the other hand, there also exist a number of negative results defining some of the limits of the field - for example, Colbeck has shown that for many classes of function, unconditionally secure two-party classical computation is impossible even in the relativistic context [170], while Buhrman et al have shown that secure-position verification is impossible if adversaries are allowed to share an arbitrarily large entangled quantum state [171].

The first significant application of relativistic cryptography was to bit commitment [13, 165, 172, 173], a basic cryptographic primitive whose applications include coin tossing [174–176], electronic voting [177], zero-knowledge proofs [72, 178], oblivious transfer [179, 180], and secure two-party computation [8, 14, 15, 181–183]. A bit commitment protocol involves two mistrustful parties who control disjoint secure regions and exchange information: the committer (henceforth called Alice) carries out actions that commit her to a particular bit value, and later, if she chooses, she may give the receiver (henceforth called Bob) some classical or quantum information that unveils the committed bit. Ideally, the protocol should guarantee to Bob that Alice cannot change her mind about the value of the bit after the time of the commitment, but also guarantee to Alice that Bob can learn no information about the committed bit unless and until she unveils.

As we noted in section 1.2.7, it has long been known that it is impossible to provide such guarantees using only classical information theory in non-relativistic Galilean spacetime [184], and a no-go theorem due to Mayers, and to Lo and Chau demonstrates that unconditionally secure bit commitment is also impossible in a purely quantum context if it is assumed that the entire protocol takes place in non-relativistic Galilean spacetime [8, 14, 15, 182, 183]. However, it is possible to circumvent these obstacles by moving to *relativistic* bit commitment protocols couched in *Minkowski* spacetime, where ‘Alice’ and ‘Bob’ now represent two separate networks of collaborating agents distributed in spacetime. Each set of agents is assumed to be acting with perfect trust in one another, but because they are sep-

arated in space, and because quantum information cannot be broadcast, on any given spacelike hyperplane the collaborating agents may not all be in possession of the same information, and by leveraging this fact we can come up with relativistic bit commitments that are unconditionally secure [6, 134, 165, 173, 185, 186].

For example, in Kent's original relativistic bit commitment protocol [12], Alice has two separated agents A_0 and A_1 who share a list $m_0, m_1, m_2 \dots$ of independently chosen random numbers in the range $\{0, 1 \dots N\}$ with $N = 2^q$ for some integer q . To begin the protocol, Bob sends A_0 a pair (n_0, n_1) of random numbers in the range $\{0, 1 \dots N\}$ and A_0 replies with $n_0 + m_0$ if she wants to commit to bit value 0 and $n_1 + m_0$ if she wants to commit to bit value 1. At some spacelike separated point, one of Bob's agents sends A_1 a set of q labelled pairs $(n_0^1, n_1^1), (n_0^2, n_1^2) \dots (n_0^q, n_1^q)$ and A_1 similarly commits to the binary form of m_0 by replying with $n_0^1 + m_1$ if the first bit of the binary encoding of m_0 is 0 and $n_1^1 + m_1$ if the first bit is 1, and so on. At some later spacelike separated point A_0 will likewise commit to the binary forms of $m_1, m_2 \dots m_q$, and this sequence of communications continues until one of Alice's agents A_i chooses to unveil the commitment by revealing the set of random numbers used by the other agent $A_{i \oplus 1}$ in the most recent round. The security of this protocol is thought to depend only on Minkowski causality, and hence it is conjectured to be unconditionally secure [172].

There also exist relativistic *quantum* bit commitment protocols, such as Kent's 'flying qubits' scheme, which is defined by a commitment point (t_c, x_c) and two unveiling points $(t_0, x_0), (t_1, x_1)$ which are lightlike separated from (t_c, x_c) and spacelike separated from each other: at (t_c, x_c) one of Bob's agents gives one of Alice's agents a quantum system whose state she does not know, and she is required to send it to (t_0, x_0) if she wishes to commit to bit value 0 and to (t_1, x_1) if she wishes to commit to bit value 1. The security of this protocol depends on both the special relativistic no-signalling principle and the quantum no-cloning theorem, so the protocol uses both special relativity and quantum mechanics in non-trivial ways [187]. Another relativistic quantum bit commitment scheme, also due to Kent [188], uses an adaptation of the BB84 key distribution scheme that we introduced in section 1.2.7; the feasibility of this scheme has been demonstrated experimentally [186], and it seems likely that schemes of this kind will play an

important role in the next generation of cryptographic technologies.

In the second part of this thesis, we propose several new relativistic quantum bit commitment protocols that have a number of advantages over existing versions - those described in chapter 7 require very little randomness, or even none at all, while those described in chapter 8 are fully ‘device independent,’ [76, 103, 189–198]. These advantages may be useful for practical applications and are also of theoretical interest in terms of our understanding of the relationship between cryptographic protocols and physical laws.

2.4 Zero-Knowledge-Proving

Zero-knowledge-proving is a primitive of classical cryptography in which one agent (henceforth called Alice) proves a fact to another agent (henceforth called Bob) without giving away any information other than that the fact is true. It has a wide range of practical applications, particularly in electronic voting schemes [199] and digital signature schemes [200], and is also used for a variety of theoretical purposes, such as showing that a language is easy to prove [201]. Zero-knowledge proving of *knowledge*, where Alice is required to prove that she knows some fact without giving Bob any information about the fact, is a particularly useful version of this task which plays a key role in a number of identification protocols [202].

An ideal zero-knowledge-proving protocol has three key properties. First, completeness: if Alice does indeed know the relevant fact and proceeds honestly with the protocol, then her proof is accepted with probability one. Second, soundness: if Alice does not know the relevant fact, then for any possible strategy that she might employ, the probability that her proof is accepted is zero. And finally, zero-knowledge: if Alice performs the protocol correctly, then for any possible strategy that Bob might employ, he gains no information about the relevant fact other than that it is true. In practice it is not always possible to ensure that these properties hold exactly - in particular, in many zero-knowledge protocols Alice has some small probability of producing a correct proof even if she does not actually know the relevant fact. However, this probability can usually be made arbitrarily small by iterating the protocol a sufficiently large number of times, so one

can reasonably (modulo epsilonics) speak of classical zero knowledge *proofs*.

The possibility of a quantum generalisation of zero-knowledge proving was explored by Horodecki et al. [203] who studied what they called ‘zero knowledge convincing protocols on quantum bit’ [sic]. In their model, a verifier (henceforth called Bob) knows he has a single copy of a pure qubit, but has no other information about the state. A prover (henceforth called Alice) wishes to make a prediction that Bob can verify and that will hold with certainty only if she knows what the state is, but to do so without giving Bob any additional information about the identity of the state. Horodecki et al. showed that no non-relativistic protocol involving classical information exchanges and quantum Alice-to-Bob communications can implement this task securely [203]. They also discussed some protocols that implement very weak versions of the task, either giving Bob a great deal of information about the qubit, or giving him only weak evidence of Alice’s knowledge, or both.

Given that relativistic quantum cryptography was born out of the demonstration that it is possible to bypass certain classical and quantum no-go theorems by moving to the relativistic context, it is natural to wonder if the Horodecki et al. no-go theorem might not also be bypassed by relativistic techniques. A number of other interesting questions were also left open by the discussion of ref [203]. For example, how much evidence *can* Alice provide? Does it help to allow quantum communication from Bob to Alice rather than just from Alice to Bob? What bounds exist on the tradeoffs between the evidence Alice provides and the amount of knowledge she gives away? How do the possible protocols depend on the dimension d of the state space?

An important preliminary observation is that Alice cannot *prove* that she knows a precise classical description of a single quantum state, and this remains true even if she is not concerned about giving Bob information in the course of the proof. This is because the classical information about the state that can be extracted by measurement is bounded by Holevo’s theorem [83], so Alice always has a boundedly nonzero chance of guessing the outcome of any measurements Bob might make, even if she knows nothing about the state: for example, she can predict the outcome of a complete projective a on a qubit with probability $\frac{1}{2}$ even if she has no information about the state of the qubit. Alice may also have a high

chance of guessing the relevant outcome even if she has only partial information about the quantum state: for example, if she knows that the state of a qudit lies in some dimension 2 subspace, she can specify a complete projective measurement whose outcome she can predict with probability $\frac{1}{2}$. Moreover, Alice may have a high chance of guessing the outcomes even if she has *incorrect* information: if the state is η and Alice believes it is η' , where $\eta \neq \eta'$ but $\text{Tr}(\eta, \eta')$ is close to 1, then she is almost as likely to pass any protocol testing her knowledge of η as she would be if she actually knew η . And since we are considering a situation where Bob only has a single copy of the state in question, it is not possible to make up for these shortcomings by simply repeating the protocol as we would do in the case of classical zero-knowledge-proving.

In view of these constraints, in this thesis we will work not with traditional zero-knowledge-*proving*, but with a generalisation which we refer to as *knowledge-concealing evidencing of knowledge about a quantum state* (KCEKQS). In the second part of this thesis we give a formal definition of KCEKQS and generalize the no-go theorem of ref [203], before discussing some existing protocols for KCEKQS and then putting forward our own protocol which performs significantly better than any existing protocol.

Part I

Distribution Tasks

Chapter 3

Summoning Tasks with Nonexistent Causal Diamonds

The properties of quantum information in spacetime can be investigated by studying operational tasks, such as ‘summoning’, in which an unknown quantum state is supplied to an agent and must be returned at a specified point when a corresponding call is made. Hayden and May recently proved necessary and sufficient conditions for a summoning task to be feasible. We prove comparable necessary and sufficient conditions for a generalised version of the summoning task where the point where the state is to be returned does not necessarily lie in the causal future of the point where the call is made. This result has practical applications for distributed quantum computing and cryptography and also implications for our understanding of relativistic quantum information and its localisation in spacetime.

Based on a paper co-authored with Adrian Kent [9].

3.1 Introduction

Hayden and May’s result (theorem 1) assumes that summoning tasks are defined such that each response point lies in the future lightcone of its associated call point - after all, if this is not the case, then some of the causal diamonds referenced in

the theorem will not even exist.

But in fact this assumption is slightly restrictive, since a task can be feasible even when one of the causal diamonds does not exist. This is because in the case of only two call points, c_0 and c_1 , receiving the information that a call has not been made at call point c_0 is equivalent to receiving the information that either a call will be made at c_1 or no call will be made at all, since HM's formulation of the task comes with a guarantee that Bob will make no more than one call. It follows that if the response point r_1 lies in the future lightcone of some point in the causal diamond D_0 , we can construct a successful protocol even if r_1 does not lie in the future lightcone of c_1 .

Thus, as a warm-up to the study of generalised summoning tasks, we define an altered version HM's original task to allow for this possibility and generalise theorem 1 to cover this scenario.

3.2 Necessary and Sufficient Conditions

Here, and in all future discussion of distribution tasks, we adopt the approximation that quantum states may be effectively localised to a point (see Ref. [7] for further discussion of this approximation and its limitations), and for simplicity we work in the coordinate system defined by Alice's rest frame, using that coordinate system to write $P_s = (t_s, x_s)$, $c_i = (t_i^c, x_i^c)$, $r_i = (t_i^r, x_i^r)$.

Theorem 2. *Consider a summoning task defined by a start point P_S and a set of N ordered call-response pairs (c_i, r_i) , in which Alice is given a quantum system Q_B in some unknown state ψ at P_S , and if a call is subsequently made at exactly one call point c_i , Alice must return a quantum system in the state ψ at the corresponding response point r_i . The task is feasible iff the following conditions hold:*

1. *Every response point is in the future lightcone of the start point.*
2. *For every pair $(c_i, r_i), (c_j, r_j)$, either both of the response points lie in the future lightcone of c_i , or both of the response points lie in the future lightcone of c_j .*

Proof. The necessity of condition 1) follows from no-signalling.

To see that condition 2) is necessary, suppose there exists a feasible task such that this condition does not hold for some pair $(c_i, r_i), (c_j, r_j)$. Then when we employ a protocol that is guaranteed to succeed for this task, if calls are made at both c_i and c_j , it follows from no-signalling that the unknown state will be returned at both r_i and r_j , in violation of the no-cloning theorem. We have obtained a contradiction, so the condition must hold for any pair $(c_i, r_i), (c_j, r_j)$ if the task is feasible.

To see that the conditions are sufficient, note that HM's proof [7] shows that a protocol for any number of call points can be built recursively out of a protocol for two call points. Their argument extends to the more general configurations we consider, so we need only show that there exists a protocol that guarantees a successful response for any two call-response pairs that satisfy the conditions of theorem 2 above, but not the conditions given in theorem 1, i.e. two pairs $(c_i, r_i), (c_j, r_j)$ such that $r_i > c_i, r_j > c_i, r_j \not> c_j$. The protocol is as follows.

1. Before the protocol begins, Alice creates a Bell pair on the systems B_S, B_i , sending B_S to her agents at x_S and B_i to her agents at x_i^c .
2. At time t_s , when Alice's agents at x_s receive a system Q_B in the state ψ , they perform a Bell measurement on the joint systems Q_B, B_S , thus teleporting ψ to B_i . They then broadcast their measurement result in all directions.
3. If a call is made at c_i , Alice's agents at x_i^c send B_i to x_i^r , otherwise they send it to x_j^r .
4. If a call is made at c_i , Alice's agents at x_i^c receive the teleportation data from the start point and apply the appropriate operation on B_i to recover the state ψ , then hand B_i over to Bob at r_i .
5. If a call is not made at c_i , Alice's agents at x_j^c receive the teleportation data from the start point and apply the appropriate operation on B_i to recover the state ψ , then hand B_i over to Bob at r_j .

□

3.3 Summary

Examining theorem 2, it is straightforward to infer that for a feasible task there can be at most one response point that does not lie in the future lightcone of its associated call point, so clearly the generalisation we have made is a fairly minor one. However, it could nonetheless be of practical importance. For example, suppose we are modelling a distributed quantum computation as a summoning task: to increase efficiency, it might be important to design the task so as to minimise the average time taken per round, and placing one response point in a position that is spacelike separated or even in the causal past of the corresponding response point might be one way of making this average time smaller. If such tasks are implemented a large number of times in the course of a computation, even a small change such as this could potentially lead to a significant computational speedup.

We also suggest that the correction makes an important interpretational difference. HM comment that a qubit can be summoned from some call point to the corresponding response point ‘if and only if the qubit is localized in the causal diamond’ and therefore ‘the summoning task is possible if and only if the qubit’s information is replicated in each and every one of the causal diamonds.’ [7] But theorem 2 demonstrates that summoning can be completed successfully even in cases where the causal diamond does not even exist; moreover, theorem 2 will of course also apply to all cases covered by theorem 1, and therefore this result demonstrates that there exists an alternative - and, arguably, equally simple - way of characterising the set of feasible distribution tasks that does not make reference to causal diamonds. This puts some pressure on the special ontological role played by the causal diamonds in HM’s interpretation of their result, and thus makes it less plausible to think of the conditions obtained as a description of the regions of spacetime where quantum states may in some sense be ‘localised.’

Chapter 4

A Quantum Paradox of Choice

The properties of quantum information in spacetime can be investigated by studying operational tasks, such as ‘summoning’, in which an unknown quantum state is supplied to an agent and must be returned at a specified point when a corresponding call is made. Hayden-May recently proved necessary and sufficient conditions for a summoning task to be feasible. We prove comparable necessary and sufficient conditions for a generalised version of the summoning task where several calls may be made and a correct response to any one call constitutes a successful response to the task. Thus we demonstrate the existence of an apparent paradox: this extra freedom makes it strictly harder to complete the summoning task. This result has practical applications for distributed quantum computing and cryptography and also implications for our understanding of relativistic quantum information and its localisation in spacetime.

Based on two papers co-authored with Adrian Kent [9, 10].

4.1 Introduction

On a dark and stormy evening in an oddly named tavern in London, you come across a man named Harry who insists that he can perform magic. When pressed to prove his credentials, Harry the Magician (HM) agrees to a demonstration. He asks you to give him an object that you are sure he cannot copy, and you reluctantly hand over your prized signed first-edition of *Speakable and Unspeakable*

in *Quantum Mechanics* [204]. Harry then goes behind a curtain, and when he emerges, he presents you with N boxes and asks you to choose one. Opening your chosen box, he reveals your book inside!

Being sceptical by nature, you float the possibility that he might simply have been lucky, but Harry is happy to repeat the trick as many times as you like until you are convinced that the result cannot simply be coincidence. You are still suspicious, however. You imagine that perhaps Harry has arranged some concealed mechanism that passes your book sequentially through the boxes, allowing him to stop the book when it reaches the box you have selected. But then, on Harry's next attempt, you try pointing to two different boxes and asking Harry to open either one of them. Harry picks a box and opens it, but to your surprise, your book is *not* inside. Indeed, on repetition you discover that Harry is no longer reliably able to make the trick work when you select more than one box, even though you allow him to choose which of your selections to open. This argues against your mechanical explanation, and indeed seems to make any mechanistic explanation problematic - how can giving Harry more freedom make him unable to complete the task?

The so-called paradox of choice, in which more choice makes consumers less happy, is a familiar concept in economics [205]. Harry's paradox, however, is much sharper: more freedom in choosing how to execute a task makes it *impossible*. Strange as this may sound, we show that such a situation can indeed arise when quantum mechanics is combined with classical relativity. Our paradox is derived from a different generalisation of HM's summoning task [7]: we alter the task so that Bob is no longer constrained to make exactly one call at exactly one call point. Since the r_i may be spacelike separated, the no-cloning theorem means that Alice cannot return the state several times; hence we define a *multiple-call summoning task* such that, if several calls are made at points c_i , Alice need only return the state at any one of the corresponding return points r_i in order to complete the task successfully. Using the no-signalling and no-cloning theorems we establish necessary and sufficient conditions on the geometric configuration of call and return points for a multiple-call summoning task to be feasible, and we show that these are strictly stronger conditions than those established by HM for the case where it is guaranteed that at most one call will be made.

4.2 Necessary and Sufficient Conditions

We use notation, terminology and approximations as defined in section 2.1 and chapter 3. Note that for a multiple-call summoning task we assume once again that every response point lies in the future lightcone of the associated call point, as in the original summoning task. We then have the following result:

Theorem 3. *Consider a multiple-call summoning task defined by a spacetime point P_s and a set of N ordered pairs of spacetime points (c_i, r_i) such that $\forall i r_i \geq c_i$, in which Alice is given a quantum system Q_B in some unknown state ψ at P_s , and if a call is subsequently made at exactly the set of call points with indices in the set K , Alice must return a quantum system in the state ψ at any response point $r_i : i \in K$. The task is feasible iff:*

1. Every response point $r_i \geq s$
2. For any subset K of $\{1, \dots, N\}$, there is at least one $k \in K$ such that $r_k \geq c_i$ for all $i \in K$.

Proof.

Necessity: The necessity of condition 1) follows from the no-signalling principle.

To see that condition 2) is necessary, suppose for the purpose of obtaining a contradiction there exists a feasible multiple-call summoning task that includes some subset \mathcal{M} of $M \leq N$ call-response pairs such that for every call-response pair $(c_i, r_i) \in \mathcal{M}$, there exists at least one call-response pair $(c_j, r_j) \in \mathcal{M}$ such that $r_i \not\geq c_j$.

Let us then consider the subtask defined by dropping from the original task all call-response points not in \mathcal{M} . Since the original task is feasible, this subtask is also feasible, so there exists some strategy S for Alice that is guaranteed to produce a successful response for this subtask. The probability of a successful response for a strategy containing probabilistic elements must be a convex combination of the probabilities of successful responses for various deterministic strategies, which means that there exists a protocol guaranteeing a successful response

to the task only if there exists a deterministic protocol guaranteeing a successful response to the task. Thus we henceforth confine our attention to deterministic strategies S .

Let Q_i^S be the total number of subsets $Q \subseteq \{c_1, c_2, \dots, c_M\}$ such that when Alice uses strategy S , ψ is returned at r_i if calls are made at all points in Q . Each return point is in the causal future of the corresponding call point, so Alice may decline to return anything at r_i if no call is made at c_i , and we therefore assume without loss of generality that strategy S returns no state anywhere if no calls are made. From the no-signalling principle, for any i, j such that $r_i \not\geq c_j$ and any fixed strategy for Alice, the response made at r_i when calls are made at some set of call points Q with $c_i, c_j \in Q$ is the same as the response made at r_i when calls are made at the set $\{Q \setminus c_j\}$. Thus for any r_i for which there exists some j such that $r_i \not\geq c_j$, Q_i must be even. In this subtask there exists some such j for every i and therefore $\sum_{i=1}^M Q_i^S$ must also be even.

Since S guarantees a successful response, Alice must respond at one or more points to any possible set of calls; from the no-cloning theorem, Alice can respond at no more than one point to each possible set of calls; and hence S must ensure that Alice responds at exactly one point whenever calls are made at some subset $Q \subseteq \{c_1, c_2, \dots, c_M\}$. Thus we must have $\sum_{i=1}^M Q_i^S = \sum_{j=1}^M \binom{M}{j} = 2^M - 1$, which is always odd.

We have defined a contradiction; thus we infer that there exists a successful protocol for the sub-task defined by the set \mathcal{M} only if there exists at least one call-response pair $(c_i, r_i) \in \mathcal{M}$ such that $\forall (c_j, r_j) \in \mathcal{M}$ we have $r_i \geq c_j$.

Since this reasoning may be applied to any sub-task, there exists a successful protocol for a multiple-call summoning task only if for any subset S of call-response pairs, at least one response point r_k belonging to a pair in S lies in the future lightcone of all the other call points belonging to pairs in S .

Sufficiency: We now show that the conditions are sufficient, by exhibiting a protocol that always succeeds for a multiple-call summoning task that satisfies the conditions of theorem 3.

Let d be the agreed dimension of the Hilbert space of the unknown state ψ . Define $S_N = \{1, \dots, N\}$. From condition 2) of theorem 3, there is at least one

$i \in S_N$ such that $r_i \geq c_j$ for all $j \in S_N$. Choose one such, i_N , and define $S_{N-1} = S_N \setminus \{i\}$. Similarly, choose $i_{N-1} \in S_{N-1}$ such that $r_{i_{N-1}} \geq c_j$ for all $j \in S_{N-1}$, and so on. We thus obtain an ordered sequence of call response pairs $(c_{i_1}, r_{i_1}), \dots, (c_{i_N}, r_{i_N})$ such that the return point of any pair in the sequence lies in the causal future of the call points of all previous pairs. We relabel the c_{i_j}, r_{i_j} , writing j for i_j .

Alice may now proceed as follows:

1. Before the protocol begins, she distributes maximally entangled pairs of states in $\mathcal{C}^d \otimes \mathcal{C}^d$ on the systems $(Q_s, Q_1^c), (Q_1^c, Q_2^c), \dots, (Q_{N-1}^c, Q_N^c)$ between agents at the spatial locations $(x_s, x_1^c), (x_1^c, x_2^c), \dots, (x_{N-1}^c, x_N^c)$.
2. At time t_s , when the state is given to Alice's agent at x_s , the agent immediately uses Q_s to teleport it to Q_1^c , broadcasting the classical teleportation data.
3. If a call is made at c_1 , Alice's agent at x_1^c immediately sends Q_1^c to x_1^r , where another agent uses the classical teleportation data from P_s to reconstruct the original state and return it to Bob at r_1 .
4. If a call is not made at c_1 , Alice's agent at x_1^c immediately teleports the state of Q_1^c to Q_2^c , broadcasting the classical teleportation data.
5. If a call is made at c_2 , Alice's agent at x_2^c immediately sends Q_2^c to x_2^r . If a call has already been made at c_1 , Alice's agent at x_2^r does nothing; otherwise she uses the classical teleportation data from P_s and c_1 to reconstruct and return the state at r_2 .
6. The process continues until either a call is made at some c_i , and the state is reconstructed and returned at r_i , or the protocol terminates without a call being made. If no call is made, Alice may reconstruct the state at r_N if she wishes, but she does not return it to Bob.

□

Comments

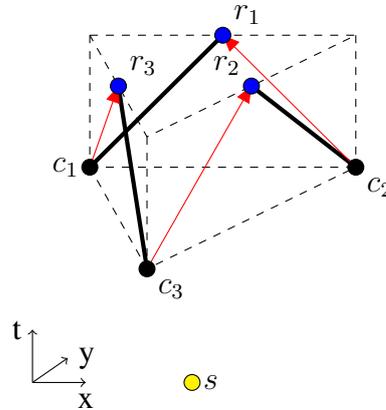


Figure 4.1: A $2 + 1$ dimensional example, taken from Ref [7], which is feasible if one call is guaranteed, but not if more than one call may arrive. The black lines are lightlike and in this case they represent the entirety of the casual diamond for each pair; the red arrows are also lightlike. We have that $c_1 < r_3, c_2 < r_1$ and $c_3 < r_2$.

1. Any task satisfying the conditions of theorem 3 also satisfies those of theorem 1. However, it is easy to construct tasks that satisfy the conditions of theorem 1 but not those of theorem 3. For example, Fig. 3 of [7], reproduced below as fig 4.1, describes one such set. Allowing the possibility of more than one call thus makes the summoning task strictly harder, as we claimed in our initial description of the quantum paradox of choice.
2. Nonetheless, the conditions of theorem 3 still do not imply that there is a causal path running from the start point through each causal diamond. An example is given in fig 2.
3. The ordered sequence of pairs used in the proof of sufficiency is not necessarily unique. For example, a nested set of two pairs, with $c_i \leq c_j < r_j \leq r_i$ (and appropriate relations to the other diamonds) may be taken in either order. More generally, one can construct examples including sets of n non-overlapping diamonds (c_i, r_i) (for $1 \leq i \leq n$) for which $c_i < r_j$ for all i, j ; Fig. 2 gives an example of this type for $n = 3$.

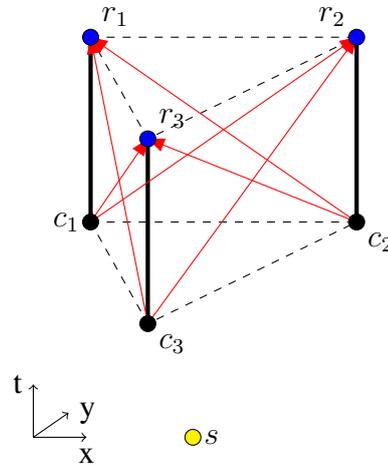


Figure 4.2: An example with $c_1, c_2, c_3 < r_1, r_2, r_3$. The task satisfies the conditions of both theorem 3 and 1, so summoning with any number of calls is feasible, even though there is no causal path running through all three causal diamonds. The timelike black lines run between the call points and their return points and define the centres of the diamonds at each time; the red arrows are lightlike.

4.3 Resolution

The resolution of the apparent paradox rests on a feature of summoning tasks whose importance is easy to overlook. In HM's original formulation of the summoning task, it seems *prima facie* as if the guarantee of at most one call plays no special role other than to ensure that Alice is never required to produce two copies of an unknown state, which of course would make the task impossible since the no-cloning theorem prohibits her from achieving such a thing. Thus it is initially surprising that summoning becomes more difficult when more than one call is allowed, even though we stipulate that Alice need only return the state at a single point and thus will never be asked to violate the no-cloning theorem.

But in fact, the guarantee of at most one call plays a much more significant role than previously appreciated. When Alice knows that no more than one call will occur, as soon as she learns that a call has been made at one call point she can assume that no call has been made at any other call point (or else Bob has made the task invalid so she is not required to produce a response), and this knowl-

edge serves as a non-local resource which improves her ability to coordinate the behaviour of her agents. On the other hand, in a multiple-call summoning task, learning that a call has been made at one point tells Alice nothing about the distribution of calls at other points. Thus although she has greater freedom, she also has fewer resources available to her in terms of coordinating the behaviour of her agents, and in some cases this loss of coordination prevents her from being able to complete tasks that she *is* able to complete when the guarantee of only one call is in force.

4.4 Finkelstein's Objection

In ref [11], Finkelstein put forward a putative classical analogue of our multiple summoning task. In Finkelstein's proposed task, Alice has agents at space points L and R , and Bob will either request for a signal to be sent from L to R at time T and no signal to be sent from R to L , submitting his request to agent L at time T , or for a signal to be sent from R to L and no signal to be sent from L to R , submitting his request to agent R at time T . Finkelstein argued that if only one request will be made, Alice can ensure compliance by simply having the two agents do as they are requested, but if two requests may be made then Alice cannot ensure compliance.

We pointed out in ref [10] that this example is not correct: Alice can ensure compliance when two requests are made by simply stationing an agent halfway between L and R who will intercept both signals and send only one of the two onwards to its final destination. However, we acknowledge that it is certainly possible to come up with classical tasks which do have the features that Finkelstein intends. For example, suppose that Alice has two agents, A_0 and A_1 separated by a distance D and Bob has agents B_0 and B_1 adjacent to A_0 and A_1 respectively, with all these agents stationary in some mutually agreed inertial frame. At time $t = 0$ in the agreed frame, each B_i sends the associated A_i a classical bit, with a guarantee that at least one bit value 1 will be sent, and Alice's task is to have her agents A_i return, effectively instantaneously, two *different* classical bits to the two agents B_i , ensuring that the agent who sent 1 also receives 1. The task is clearly trivial if there is an additional guarantee that only one bit value 1 will be

sent, but if it is allowed that both of Bob's agents B_i send bit value 1, then Alice cannot ensure that she always completes the task successfully, even though in the case that both B_i send bit value 1 there are two valid ways of completing the task rather than just one as in the trivial case.

Does the existence of such classical tasks threaten our description of our result as a paradox? First of all, it should be noted that we do not claim that our result represents a paradox in the sense that it is an actual physical *contradiction* - indeed, we have described the resolution to the apparent paradox in section 4.3. The term 'paradox' is intended only to highlight the fact that the effect we describe seems surprising from the point of view of a certain intuitive way of thinking about quantum states and their spatiotemporal localisation. As noted in section 2.2, a paradox is characterised as such in terms of the limitations of human cognition and of pre-existing mental models, which may of course be different for different readers: if there are people whose pre-existing intuitions about summoning relativistic quantum information assured them that theorem 3 must obviously be true, for such people our result will of course not seem to present much of a paradox, but most of our audiences have found the effect surprising and hence the term 'paradox' is appropriate.

A more serious concern is that the existence of classical tasks like the one set out above might threaten our contention that this result is a specifically relativistic quantum phenomenon that is not exhibited in classical physics, or in quantum physics alone, or relativistic physics alone. We of course acknowledge that the mere existence of a task where being given more options makes the task harder is not a specifically relativistic or quantum phenomenon, as demonstrated by the example above. However, we maintain that the specific effect we have described here is indeed irreducibly quantum, because the distinction between summoning tasks with and without guarantees, and indeed the very fact that there are interesting constraints on summoning tasks, relies on the no-cloning theorem; moreover theorem 3 also relies on the delocalisability of quantum information via quantum teleportation, which allows summoning even in configurations where sending the state along any given causal path cannot succeed. We also maintain that the effect is irreducibly relativistic, because the distinction between summoning tasks with and without guarantees, and the existence of interesting constraints on sum-

moning tasks, also relies on the impossibility of superluminal signalling in special relativity¹ Thus, whether or not one finds it illuminating to describe our result as a ‘paradox,’ it is certainly an intrinsically quantum and relativistic result, and thus represents an advance in our understanding of the behaviour of quantum systems in relativistic spacetime.

4.5 Applications

This generalisation of the original summoning task is a natural one to consider in the context of a distributed quantum computation. In general, we are likely to be interested in using such computations to perform calculations whose answers we cannot efficiently calculate using ordinary computers, and therefore if we have a scenario where calls may be made for a given state at any one out of several distinct nodes, with the choice of call depending on the results of previous rounds of computation, it might not always be possible to guarantee that only one call will be made at any given round. We would not want the whole computation to fail whenever more than one call was made, and therefore ideally we would like to arrange distribution protocols that are able to cope with the possibility of multiple calls. The results of this chapter show how to design distribution tasks for which there exists a protocol that is able to complete the task with one hundred percent success rate, and also provide an explicit protocol that is guaranteed to succeed whenever there exists any protocol that is guaranteed to succeed.

Moreover, the original no-summoning theorem has already led to new applications in relativistic quantum cryptography [16, 17, 165, 173, 186]. The stronger results reported here and in Ref. [7] suggest further ways of exploiting summoning as a general way of controlling the flow of quantum information, and therefore we expect these results to find application in future cryptographic protocols as well as in quantum network algorithms.

¹It is true that there exists a no-signalling theorem in quantum physics, but this is not sufficient to derive theorem 3; we specifically make use of the *relativistic* no-signalling principle, as indicated by the fact that theorem 3 is couched in the language of Minkowski spacetime. See section 1.2.2 for further comments.

4.6 Summary

The mysterious man named Harry disappears into the night before you are able to ask him for a more exciting demonstration of his powers, leaving behind only your copy of *Speakable and Unspeakable* and what appear to be a few feathers from a snowy owl. But no matter - we have learned a great deal from him already.

Like the original no-summoning theorem, our results rely crucially on both relativity and quantum theory. The paradox and its resolution thus allow us to probe intuitions about the nature of quantum states as spatiotemporal entities. There is a tradition in physics of describing quantum states and quantum information using the language of persisting physical objects [206]; for example, a widely cited review by Horodecki et al. [207] states:

‘This is the essence of teleportation: a quantum state is transferred from one place to another: not copied to other place, but moved to that place.’

Hayden and May’s discussion of their work follows the same linguistic tradition. In particular, they interpret their results as giving a characterization of the spatiotemporal location of the quantum information stored in a quantum state in between preparations and measurements:

‘We fully characterize which regions of spacetime can all hold the same quantum information. Because quantum information can be delocalized through quantum error correction and teleportation, it need not follow well-defined trajectories. Instead, replication of the information in any configuration of spacetime regions not leading to violations of causality or the no-cloning principle is allowed . . . This provides a simple and complete description of where and when a qubit can be located in spacetime, revealing a remarkable variety of possibilities.’ [7]

Yet the rationale for this common mode of description remains largely unanalysed. Although much ink has been spilled over the question of the reality of the quantum state, these arguments focus almost exclusively on *instantaneous*

facts [65, 208–214]; following Harrigan and Spekkens [65] it is usually held that a quantum state may be regarded as an element of reality iff no two quantum states (at a time) are compatible with the same underlying ontic state of the world (at that time), and this is the definition which occupies the attention of most of the community, so we seldom ask whether it is accurate or even useful to think of quantum states as *persisting* physical entities.

In particular, our results cast doubt on this way of talking about HM’s result, because our analysis of multiple-call summoning tasks makes it clear that the guarantee of a single call plays a pivotal role in Alice’s ability to complete various summoning tasks. This means that the special role of the future lightcones of the call points in a summoning task is explained mainly by the fact that they are regions where Alice is in possession of nontrivial information about which call has been made, suggesting that the causal diamonds are in large part privileged *epistemologically* rather than *ontologically*. Hence it is somewhat misleading to regard the causal diamonds as regions in which some specific quantum information is localised, since the possibility of this interpretation was an artifact of the extra information imparted by the guarantee and not really anything to do with the spatiotemporal history of a persisting quantum state.

Indeed, thinking about the quantum state as something akin to a persisting physical object is in part what gives rise to the kind of thinking that makes our result seem paradoxical in the first place, and hence one lesson to be learned from the apparent paradox and its resolution is that any supposed account of the spatiotemporal location of a quantum state in between preparations and measurements should be treated with great care. As ref [206] points out, quantum states and quantum information are *not* persisting physical entities in the ordinary everyday sense, and describing them as such can give rise to a number of serious fallacies.

Of course, we do not mean to suggest that one should never speak informally about the spatiotemporal history of quantum states or quantum information in between preparations and measurements, as this can be a helpful way of proving and synthesising operational results. But we reinforce that such language should not necessarily be taken literally, and that to guard against fallacies it is important to use it precisely. In particular, if the behaviour of quantum information in a

system depends on external events, such as the calls in a summoning task, any discussion of its localisation should reflect that external dependence.

For example, one way of parsing HM's statements about regions of spacetime 'hold(ing) the same quantum information' would be to say that *if* Alice follows their prescribed strategy in a summoning task, and *if* a call arrives at c_i and no call arrives at any other call point, *then* the unknown state's quantum information becomes localised within D_i and the state is reconstructed at r_i . Another option would be to observe that on a run of the protocol where the HM algorithm is followed but no call is actually made, then each causal diamond contains a spacelike hypersurface on which the unknown state could have been reconstructed, had a call been made at the relevant call point.² Both of these ways of summarising the HM result get across the key point about the accessibility of the relevant quantum information within the various causal diamonds, without misleadingly appearing to make ontological statements about the spatiotemporal history of a quantum state in between preparations and measurements.

²Indeed, this suggestion was made in the course of conversation between A. Kent and P. Hayden. It may not necessarily represent the position of any of the authors cited - as far as we are aware HM have not produced a definitive response to the discussion of ref [9] - but it merits a mention.

Chapter 5

Spooky Summoning

We introduce a new set of relativistic quantum protocols that involve distributing entangled states over spacetime. We define two related tasks: distributing two halves of a known entangled state, and distributing two halves of an unknown entangled state. We derive necessary conditions on the spacetime configurations of call and response points for a task to be feasible, and make some comparisons between the sets of feasible tasks in different contexts.

Based on unpublished work with Adrian Kent.

5.1 Introduction

The work presented in the two previous chapters, together with the results of Hayden and May [7], provides a comprehensive characterisation of the ways in which a single quantum state can be distributed over spacetime. But many relativistic quantum communication and cryptographic protocols [6, 7, 215, 216] and distributed quantum computing routines [217–219] also require participants to share *entanglement* with networks of agents at arbitrary points throughout space, and since in practice the creation and transmission of entanglement is a costly resource [215, 220], it is of immediate practical concern to understand the exact conditions under which it is possible to achieve various spacetime arrangements of entangled states. For example, in measurement-based distributed quantum computing schemes [221], certain entangled states may need to be shared between

different pairs of spatially separated locations depending on the results of measurements performed in the course of the protocol, so in order to achieve an optimal tradeoff between the number of entangled states that must be produced and the speed of the computation, it is helpful to characterize the circumstances under which agents can interactively distribute entanglement to various configurations of spacetime points.

Working in the tradition of the summoning tasks discussed in previous chapters, we consider the problem of disseminating entanglement as a quantum distribution task. We define two such tasks: distributing two halves of a known entangled state, and distributing two halves of an unknown entangled state. For each task we find necessary conditions on the spacetime configurations of call and response points for a feasible task; we then compare the two tasks and discuss the possibility of attaching labels to states and calls. Finally, we examine possibilities for further research on this topic and call attention to the important role played by asymmetry.

5.2 Entanglement distribution

Definition 1. An **entanglement distribution task** is defined by a set of $N \geq 3$ ordered pairs $\{(c_i, r_i)\}$ and a fixed maximally entangled bipartite state ξ on a joint Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, where ξ is known to both Alice and Bob. For a given run of the task, Bob makes calls at exactly two call points c_i and c_j ; in order to complete the task successfully, Alice must return systems V_i and V_j at r_i and r_j respectively, such that the composite system V_i, V_j is in a joint state ϕ with either $E_{iA}E_{jB}\phi = \xi$ or $E_{iB}E_{jA}\phi = \xi$, where for $x \in \{i, j\}, C \in \{A, B\}$ the map $E_{xC} : \mathcal{H}_x \rightarrow \mathcal{H}_C$ is some pre-agreed isomorphism from the Hilbert space of V_x to the Hilbert space \mathcal{H}_C .

Notes

1. In the most general case, one might allow that a response point r_x may not always lie in the future lightcone of the corresponding call point c_x , which would lead to a generalisation similar to that discussed in chapter

3. However, we will at present assume that for every x , r_x is in the future lightcone of c_x .
2. Alice and Bob must agree in advance on the way in which Alice will return the state ξ - for example, they might specify that Alice will return the state in some specific form with respect to a pre-agreed basis. This amounts to specifying the isomorphisms E_{xC} , which will be known to both Alice and Bob.
3. Since the state ξ is maximally entangled, it must always be pure.

Notation and terminology For any two call-reponse pairs (c_i, r_i) , (c_j, r_j) , let $i \perp j$ denote $r_i \not\prec c_j \cap r_j \not\prec c_i$.

As in the previous chapters, we say that a task is ‘feasible’ iff there exists a protocol that guarantees a successful response to the task.

5.2.1 Necessary conditions

Theorem 4. *An entanglement distribution task defined by a set of $N \geq 3$ ordered pairs $\{(c_i, r_i)\}$ is feasible only if, for any distinct $x, y, z \in \{1, 2, \dots, N\}$, if $c_x \not\prec r_y, c_x \not\prec r_z, c_y \not\prec r_x$ and $c_y \not\prec r_z$, then $(c_z < r_y \cup c_z < r_x) \cap (c_y < r_x \cup c_y < r_z)$*

Proof. Suppose there is a feasible task that includes a set of three response points x, y, z such that $c_x \not\prec r_y, c_x \not\prec r_z, c_y \not\prec r_x$ and $c_y \not\prec r_z$. Then the joint response at r_y, r_z is independent of the occurrence of a call at c_x and the joint response at r_x, r_z is independent of the occurrence of a call at c_y , so from no-signalling, when we apply a protocol that is guaranteed to succeed, and calls are made at all three call points c_x, c_y, c_z , the system returned at r_z will be maximally entangled with the systems at both r_x and r_y , in violation of the monogamy of entanglement. Hence such a task cannot be feasible.

The same applies if there exists a feasible task that includes a set of three response points x, y, z such that $c_x \not\prec y, c_x \not\prec z, c_z \not\prec x$ and $c_z \not\prec y$. \square

5.2.2 Entanglement vs Correlations

We have chosen to define this task in terms of the distribution of a maximally entangled state, and thus the proof for the necessary condition depends on the monogamy of quantum entanglement [50]. However, a similar task can also be defined for general non-signalling theories by moving from a description in terms of entanglement to a description in terms of correlations: instead of a maximally entangled state, Alice is required to return two sets of labelled systems at the relevant response points, and Bob subsequently tests to see if an appropriate set of measurements performed on these systems will violate the CHSH inequality. Unlike in the entanglement distribution task, no (quantum) protocol can ever *guarantee* success in this correlation task, since even if Alice does return an appropriate set of maximally entangled states, it is always possible that a statistically unlikely outcome will occur and the measurements performed on these states will fail to violate the chosen CHSH inequality. However this could be dealt with by some appropriate redefinition of a feasible task - for example, one might say such a task is feasible if there exists a protocol that ensures the probability of success goes to one as the number of pairs of systems returned by Alice goes to infinity.

It is then interesting to consider how the set of correlation tasks which are feasible in the context of a general non-signalling theory where the set of achievable correlations is restricted *only* by the no-signalling principle would relate to the set of feasible entanglement distribution tasks. First, note that clearly any set of sufficient conditions for the entanglement task would also be a set of sufficient conditions for the correlation task in the most general non-signalling theory. Moreover, although quantum correlations obey a stronger monogamy relation than general non-signalling correlations [53], the proof for the necessary condition in theorem 4 can be translated directly to the correlation task by simply replacing the monogamy of entanglement with the general non-signalling tripartite monogamy bound that we introduced in section 1.2.3, so the same necessary condition applies in the correlation task even for a theory restricted only by the no-signalling principle. Suppose the same translation can be made for the proofs of all the other necessary conditions required to make up a full set of necessary and sufficient conditions for the entanglement distribution task; then the set of feasible tasks for the

entanglement and correlation tasks will coincide. If indeed this turns out to be the case, then it would be an interesting area for future investigation to look for variants on this type of distribution task where as a result of the different monogamy relations that apply in quantum theory and the most general non-signalling theory, the set of tasks that are feasible according to quantum mechanics is strictly smaller than the set of tasks that are feasible in a theory restricted only by the no-signalling principle.

5.3 Entanglement summoning

Definition 2. An **entanglement summoning task** is defined by a start point P_s and a set of $N \geq 3$ ordered pairs $\{(c_i, r_i)\}$. For a given run of the task, at the start point Bob gives Alice a pair of systems Q_A, Q_B in some state ξ , where the state ξ is unknown to Alice, and then makes calls at exactly two call points c_i and c_j ; in order to complete the task successfully, Alice must return systems V_i and V_j at r_i and r_j respectively such that the composite system V_i, V_j is in a joint state ϕ with either $E_{iA}E_{jB}\phi = \xi$ or $E_{iB}E_{jA}\phi = \xi$, where for $x \in \{i, j\}, C \in \{A, B\}$ the map $E_{xC} : \mathcal{H}_x \rightarrow \mathcal{H}_C$ is a pre-agreed isomorphism from the Hilbert space of V_x to the Hilbert space of Q_C .

Notes

1. As before, we will assume that for every x the point r_x is in the future lightcone of c_x .
2. As before, Alice and Bob should agree in advance on the form in which Alice will return the state ξ . Since Alice does not know the actual state in this scenario, this will entail stipulating the ways in which she is allowed to transform the state initially handed to her by Bob and the types of systems she is allowed to transfer it to. This amounts to specifying the isomorphisms E_{xC} , which will be known to both Alice and Bob.
3. An entanglement summoning task will qualify as feasible only if there exists a protocol for Alice which guarantees a successful response for *any*

joint state ξ . The unknown state ξ is not guaranteed to be maximally entangled, and indeed, might even be separable. In particular, the protocol must still work in the case where ξ is *asymmetric*, i.e. $Tr_{Q_A}(\xi) \neq Tr_{Q_B}(\xi)$. For example, one could employ a state $\xi = |\psi\rangle\langle\psi|$ for $\psi = \sqrt{\frac{1}{2} - \delta}|01\rangle + \sqrt{\frac{1}{2} + \delta}|10\rangle$ for some $\delta \in (0, \frac{1}{2}]$; the two halves of this state are distinguishable since a system prepared in $Tr_{Q_B}(\xi)$ would be more likely than a system prepared in $Tr_{Q_A}(\xi)$ to be found in the state $|1\rangle$.

5.3.1 Necessary conditions

Theorem 5. *An entanglement summoning task defined by a start point P_s and a set of $N \geq 3$ ordered pairs (c_i, r_i) is feasible only if:*

1. *Every response point lies in the future lightcone of the start point.*
2. *For any set of three call-response pairs, (c_i, r_i) , (c_j, r_j) and (c_k, r_k) , at least one response point from r_i, r_j, r_k lies in the intersection of the three future lightcones of c_i, c_j and c_k .*
3. *There exists no more than one set of two call-response pairs $(c_x, r_x), (c_y, r_y)$ such that $x \perp y$, and if there exists such a pair $(c_x, r_x), (c_y, r_y)$, then for every $(c_z, r_z) : z \neq x, z \neq y$, the response point r_z lies in the future lightcone of c_x and also the future lightcone of c_y .*

Proof. The necessity of condition 1) follows from no-signalling.

To see that condition 2) is necessary, suppose there exists a feasible task which includes a set of three call-response pairs (c_i, r_i) , (c_j, r_j) and (c_k, r_k) for which none of r_i, r_j, r_k lies in the intersection of the three future lightcones of c_i, c_j and c_k . We denote the two mixed states obtained from ξ by tracing out Q_A and Q_B respectively as ξ_0 and ξ_1 . Then suppose that we employ a protocol which guarantees a successful response to the task; when calls are made at c_i and c_j , ξ_0 and ξ_1 must be returned at r_i and r_j or vice versa; likewise mutatis mutandis for calls made at c_i, c_k and calls made at c_j, c_k . Moreover, since none of r_i, r_j, r_k lies in the intersection of the three future lightcones of c_i, c_j and c_k , no pair of two

calls can jointly prevent a response at the third response point. Thus suppose that on some run, Bob makes calls at c_i, c_j and c_k ; from no-signalling, responses must be made successfully at all three points r_i, r_j and r_k , and therefore either ξ_0 or ξ_1 is returned at two distinct points, so Bob can retain these states in order to produce two copies of the same state on some spacelike hypersurface, in violation of the no-broadcasting theorem. Hence we have obtained a contradiction, so no such task can be feasible.

To see that condition 3) is necessary, suppose there exists a feasible task such that on some run of this task, a call is made at some call point c_x and a successful response is made at r_x . We are then left with a summoning task as in ref [7] for the second half of the entangled state, and from the conditions derived in [7], this subtask is feasible only if for every remaining pair $(c_j, r_j), (c_k, r_k)$, either $c_k < r_j$ or $c_j < r_k$. Since the subtask must be feasible if the larger task is feasible, it must be the case that in this task, $j \perp k$ implies either $j = x$ or $k = x$. So for a feasible entanglement summoning task, either there is no set of two call-response pairs $(c_i, r_i), (c_j, r_j)$ such that $i \perp j$, or there exists a call-response pair (c_x, r_x) such that $\forall j, k \quad j \perp k \implies (j = x) \cup (k = x)$.

Now suppose that for this unique pair (c_x, r_x) , $\exists y, z$ such that $x \perp y$ and also $c_x \not< r_z$ or $c_y \not< r_z$; then the triple $(c_x, r_x), (c_y, r_y), (c_z, r_z)$ fails to satisfy condition 2) which we have already derived above. Thus there can be at most one call-response pair (c_y, r_y) such that $x \perp y$, and for any $z \neq x, z \neq y$ we must have $c_x < r_z \cap c_y < r_z$.

□

Note that wherever a task is feasible for entanglement summoning, the task defined by removing the start point is also feasible for entanglement distribution, since Alice can simply create a copy of ξ at some point in the past lightcone of all the response points and then apply the protocol for entanglement summoning. However, the converse is not true: there exist configurations which define a feasible entanglement distribution task, but do not define a feasible summoning task for any possible choice of start point. Fig 5.1 shows such a configuration: to show that it defines a feasible entanglement distribution task, we describe a protocol

that guarantees a successful response below, but not a feasible entanglement summoning task, since we have both $1 \perp 2$ and $3 \perp 0$, in violation of condition 3) of theorem 5.

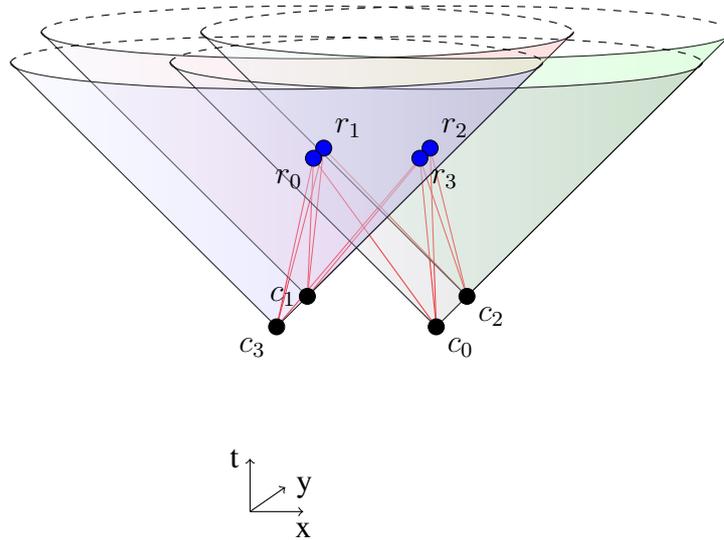


Figure 5.1: A spacetime configuration for a task defined by a set of four call-response pairs with $1 \perp 2$ and $3 \perp 0$. For the case of the entanglement summoning task, we take it that the start point is anywhere in the past lightcone of all four response points.

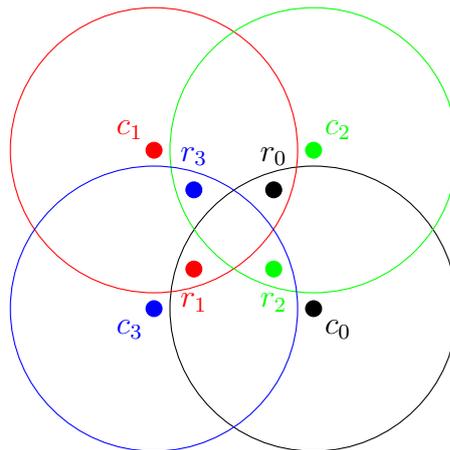


Figure 5.2: Projection on the x - y plane of the four lightcones for a task defined by a set of four call points with $1 \perp 2$ and $3 \perp 0$.

Protocol for Entanglement Distribution

1. Before the start of the protocol, Alice prepares six copies of the maximally entangled state ξ on the entangled pairs $(X_{01}^0, X_{01}^1), (X_{02}^0, X_{02}^2), (X_{13}^1, X_{13}^3), (X_{23}^2, X_{23}^3), (X_{03}^0, X_{03}^3), (X_{12}^1, X_{12}^2)$, sending $X_{01}^0, X_{02}^0, X_{03}^0$ to r_0 , $X_{01}^1, X_{12}^1, X_{13}^1$ to r_1 , $X_{02}^2, X_{12}^2, X_{23}^2$ to r_2 , $X_{03}^3, X_{13}^3, X_{23}^3$ to r_3 .
2. If a call is made at c_0 and also at c_1 , Alice's agent at r_0 returns X_{01}^0 ; if a call is made at c_0 and also at c_2 , Alice's agent at r_0 returns X_{02}^0 ; if a call is made at c_0 and no other call is made in the past lightcone of r_0 , Alice's agent at r_0 returns X_{03}^0 .
3. Likewise mutatis mutandis for c_1, c_2 and c_3 .

5.3.2 Labelled calls

An alternative way of formulating an entanglement summoning task would be for Bob to give Alice the unknown bipartite state with labels attached to the two halves of the state and then also to include labels with his calls telling Alice which half of the entangled state she is to return at the corresponding response point.

The necessary and sufficient conditions defining the set of feasible labelled tasks would then be just the same as the necessary and sufficient conditions for two separate summoning tasks, as set out in ref [7]. Comparing these conditions to the necessary conditions set out in theorem 5, we observe that the distinction between labelled and unlabelled tasks has some similar features to the paradox discussed in the previous chapter.

In one sense, the labelled task is more difficult than the unlabelled task, because there exists no protocol which guarantees success for a task with labelled calls if it includes any set of two call-response pairs $(c_x, r_x), (c_y, r_y)$ such that $x \perp y$, whereas there may exist a protocol that guarantees success for a task with unlabelled calls if it includes no more than one such set. Fig 5.3 thus gives an example of a spacetime configuration that defines a feasible unlabelled task (provided that the start point is somewhere in the past lightcone of all three response points) but not a feasible labelled task.

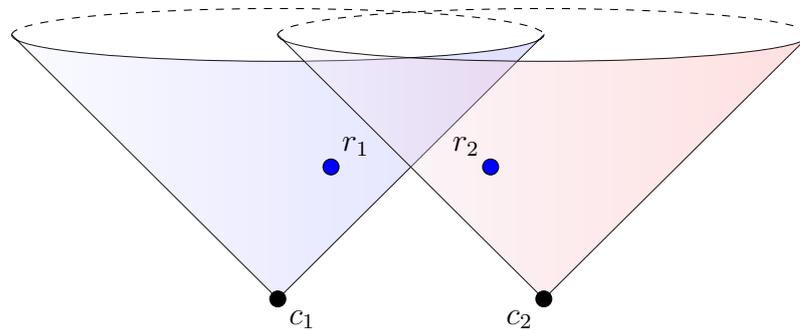


Figure 5.3: A spacetime configuration for an entanglement summoning task that is feasible with unlabelled calls but not with labelled calls. We take it that the start point is somewhere in the past lightcone of r_1 and r_2 .

Yet in another sense, the unlabelled task is more difficult than the labelled task, because the necessary conditions for a labelled task constrain only the relative configuration of any two call-response pairs, whereas the necessary conditions for the unlabelled task constrain the relative configuration of any three. Fig 5.4 thus gives an example of a spacetime configuration that defines a feasible labelled task (provided that the start point is somewhere in the past lightcone of all four response points) but not a feasible unlabelled task, since it violates condition 2) of theorem 5.

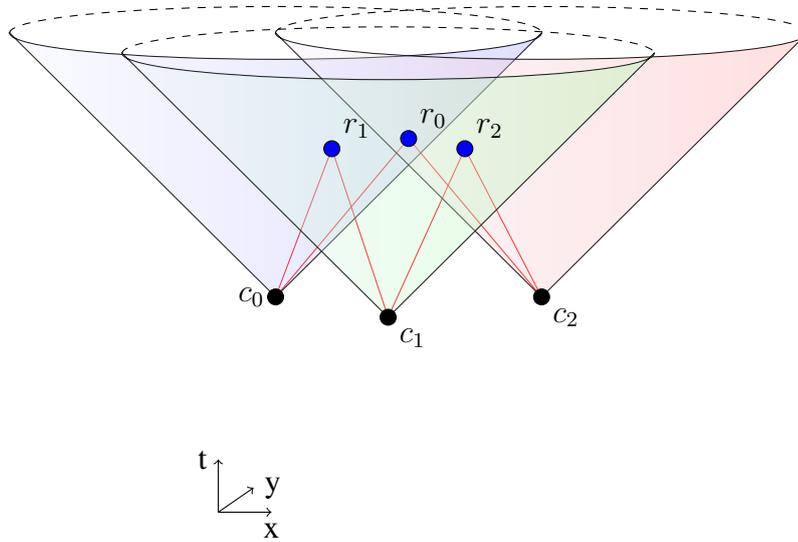


Figure 5.4: A spacetime configuration for a task that is feasible with labelled calls but not with unlabelled calls. Every response point lies in the future lightcone of two call points, but no response point lies in all three. We take it that the start point is somewhere in the past lightcone of r_0 , r_1 and r_2 .

The reason that the unlabelled task is sometimes easier is that Alice has greater freedom about how to respond to certain tasks. Conversely, the reason the labelled task is sometimes easier is that, as in chapter 4, giving Alice greater freedom means she has less information to coordinate the actions of her distant agents: the guarantee that two calls with different labels will be made provides Alice with a nonlocal resource in much the same way as the guarantee that only one call will be made provides a nonlocal resource. As a result of the two competing effects, the set of feasible labelled tasks is not a proper subset of the set of feasible unlabelled tasks and nor is the converse true, so neither form of the task is strictly speaking ‘easier’ than the other.

5.4 Summary

The work in this chapter is a starting point for the study of entanglement distribution and summoning tasks, but clearly much remains to be done. Ideally one would like to derive a full set of necessary and sufficient conditions defining the

exact set of feasible tasks in each case, and for practical purposes it would also be useful to have an explicit protocol that is guaranteed to succeed for any feasible task, as we have for the original summoning task and the variants discussed in preceding chapters.

Nonetheless, the necessary conditions we have set out already rule out a large class of configurations and allow us to make meaningful comparisons between the two tasks and also between labelled and unlabelled tasks. The latter comparison has some interesting features. In the case of our ‘quantum paradox of choice,’ when the guarantee of only one call was removed, the loss of global coordination outweighed any advantage that might have arisen from the extra freedom, and hence the conditions for the task with ‘more freedom’ turned out to be strictly stronger than the conditions for the original task. But when we remove labels from a labelled entanglement distribution task, we find that the advantage of greater freedom and the disadvantage of losing global coordination are both significant, and as a result of the two competing effects the conditions for the unlabelled task are neither strictly stronger nor strictly weaker than the conditions for the labelled task. This gives a more nuanced picture of the relationship between freedom and global coordination in distribution tasks, and vindicates the intuition discussed in chapter 4 that giving an agent greater choice about how to complete a task should, all else being equal, make the task easier.

We also call attention to some interesting points raised by our choice of definition for entanglement summoning. We defined the entanglement distribution task with respect to the distribution of a *maximally entangled* state¹, but in the entanglement summoning task (ES1) we did not limit Bob to giving Alice maximally entangled states. Indeed, had we done so, the proofs for conditions 2) and 3) of theorem 5 would no longer have applied, because these proofs rely on the fact that the two reduced states $Tr_{Q_A}(\xi)$ and $Tr_{Q_B}(\xi)$ are not known to Alice, whereas if the state ξ were maximally entangled, both reduced states would be equal to the maximally mixed state and hence would be known to Alice. One might of course define an alternative entanglement summoning task (ES2) in which Alice is given

¹Note that if we do not insist on a maximally entangled state for the entanglement distribution task, the necessary condition set out in theorem 4 will not necessarily hold, because the monogamy constraints for general entangled systems are more complex than the simple constraint for maximally entangled states.

a guarantee that the state provided by Bob will always be maximally entangled, and it would be interesting to establish whether or not the set of feasible tasks for the two variants would coincide. Clearly any configuration which defines a feasible ES1 task would also define a feasible ES2 task, since by definition any protocol which guarantees success for an ES1 task must also guarantee success in the case where Bob chooses a maximally entangled state for ξ , but it is an open question as to whether there might exist configurations which define feasible ES2 tasks but not feasible ES1 tasks. One argument suggesting that there may *not* exist such configurations is that any strategy for Alice in an ES2 task must essentially come down to sending the state ξ obtained from Bob through some set of quantum channels to the relevant response points, and if there exists a set of quantum channels that faithfully transmit the two halves of any maximally entangled bipartite quantum state while preserving the entanglement between them, then the channels must be noiseless [222], so they will also faithfully transmit the two halves of any other bipartite quantum state while preserving the entanglement between them, which means that if Alice has a protocol that is guaranteed to succeed for some ES2 task, the same protocol would also be guaranteed to succeed for the corresponding ES1 task. However, this argument is informal and does not do enough to rule out the possibility that there exists some special strategy that always succeeds for maximally entangled states but that sometimes fails for states which are not maximally entangled. Resolving this issue one way or another would not only be of practical value in terms of our understanding of the ways in which entangled states can be used in distributed quantum computing, but might further elucidate some intriguing connections between entanglement, symmetry and indiscernibility for quantum states.

Part II

Relativistic Quantum Cryptography

Chapter 6

Definitions

We provide formal definitions of relativistic bit commitment and knowledge-concealing evidencing of knowledge, and describe the criteria we will use to measure the security of these protocols.

Based in part on papers co-authored with Adrian Kent [16–18].

6.1 Relativistic Bit Commitment

Relativistic bit commitment protocols are based on a simple principle: the data that fixes Alice’s commitment is revealed at a set of spacetime points that are spacelike or lightlike separated from the commitment point. Consequently all this data is in the possession of Bob’s agents on some spacelike hyperplane Σ passing through the commitment point, so nothing Alice does at any point to the future of this hyperplane will change her commitment. However, Bob cannot find out Alice’s commitment until he can bring together information from the commitment point and the other points where data was unveiled, and therefore there is some region of spacetime in which Alice is committed but neither Bob nor any of his agents know which bit she is committed to.

Idealisations As is usual in quantum cryptography, we initially present our protocols in an idealised form assuming perfect quantum state preparations, transmissions, measurements and computations. However, the protocols are tolerant to

errors and losses, as we discuss later in chapters 7 and 8.

We also make standard idealisations about the background geometry and signalling speed. We suppose that spacetime is Minkowski, that signals are sent at precisely light speed, and that all information processing is instantaneous; hence we also assume that all actions are localised at some single spacetime point, though of course in reality agents would be acting within a spatially small secure laboratory during a small time interval.

Again, these assumptions can be relaxed. The protocols remain secure in realistic implementations with finite separations and near light speed communication. If these corrections are small, the only significant change to the security is that Bob is guaranteed that Alice's commitment is binding from some point in the near causal future of the 'commitment point' rather than from the commitment point itself [165]. Allowing for small deviations from Minkowski geometry also requires small corrections to the geometry when stating the security guarantees, but does not essentially affect security beyond that [13].

In addition, we will assume agents are located at fixed points in space throughout the protocol, relative to some chosen inertial frame. Since we allow arbitrary numbers of agents, this assumption involves no loss of generality, so long as we assume that Alice's agents have secure classical and quantum communication channels.¹

6.1.1 Definition

Our formal definition of relativistic bit commitment is as follows. All measurements of space and time are given relative to some appropriate fixed reference frame, such as Bob's rest frame.

Definition 3. Relativistic bit commitment protocol: A protocol defined by two verification functions V_0 and V_1 and a set of spacetime points $\{(t_c, x_c), (t_0, x_0), (t_1, x_1,)\}$, with the points (t_0, x_0) and (t_1, x_1) spacelike or lightlike separated from (t_c, x_c) , in which:

¹Note, however, that this assumption may not always be justified; if not, the possibility of mobile agents should be kept in mind.

1. At the commitment point (t_c, x_c) , Bob gives Alice a classical bit string B_c and a quantum system X_c in a state $\psi(X_c)$, and Alice instantaneously returns a classical bit string C_c and a quantum system Q_c .
2. For $y \in \{0, 1\}$, at the unveiling point (t_y, x_y) Bob gives Alice a classical bit string B_y and a quantum system X_y in a state $\psi(X_y)$, and if Alice wishes to commit to bit value y she instantaneously returns a classical bit string C_y and a quantum system Q_y .
3. For $y \in \{0, 1\}$, in order to verify a commitment to bit value y Bob performs a measurement on Q_c, Q_y and/or other quantum systems he may have in his possession, obtaining a joint outcome M_y .
4. For $y \in \{0, 1\}$, the verification function V_y maps $B_c, B_y, C_c, C_y, M_y, \psi(X_c)$ and $\psi(X_y)$ to $\{0, 1\}$. For any given run of the protocol, Bob will accept that Alice made a valid commitment to bit value y iff the verification function V_y returns 1.

Notes

1. Any classical data can be written as a bit string of some fixed length, so for convenience we will assume that all classical data is given in the form of bit strings unless otherwise indicated.
2. The unveiling points (t_0, x_0) and (t_1, x_1) need not always be distinct - for example, the two points coincide in Kent's protocol [13] as described in section 2.3.
3. We allow that any of the classical strings or quantum systems may be null, if there are points where Bob hands over no classical strings and/or no quantum systems, or Alice returns no classical strings and/or no quantum systems.
4. In condition 3) we specify only a single variable to represent Bob's measurement results, in order to allow for the possibility of joint measurements. The case where separate measurements are performed on Q_c and Q_y can be

subsumed under this case, since two separate measurements can always be regarded as a single measurement with a separable outcome space.

5. The verification functions would be written in full as $V_y(B, B_y, C, C_y, M, \psi(X_c), \psi(X_y))$, but we sometimes omit the arguments for notational convenience.
6. The verification functions are allowed to depend on the states $\psi(X_c), \psi(X_y)$, since these states are known to Bob, but can depend on the states of quantum systems handed over by Alice only via the measurement result M_y .

6.1.2 Security definitions

One needs to be careful about what, precisely, a bit commitment protocol is intended to guarantee in relativistic scenarios. Here we follow the physically motivated definition first set out in Ref. [6], which requires that a bit commitment should guarantee that the committed data was available to and input by Alice's committing agent A_c at the spacetime point (t_c, x_c) . This definition allows for the possibility of A_c inputting a quantum superposition of the values 0 and 1² but excludes protocols in which the unveiling agents could influence the value of the unveiled bit by using correlated information that they acquired independently of A_c [6].³

For any strategies S, S' and for $i \in \{0, 1\}$, let $p_i(S, S')$ be the probability that if Alice employs a fixed strategy S up to and including the time of the commitment (in the relevant fixed reference frame), and her agents at (t_i, x_i) subsequently attempt to unveil a commitment to bit value i using strategy S' , a valid commitment to bit value i is indeed unveiled. Define $p_i^\Sigma(S) = \max_{S'} p_i(S, S')$ where the

²Like all technologically unconstrained quantum bit commitment protocols [223, 224], the protocols we describe in chapters 7 and 8 do not prevent Alice from committing to a quantum superposition of bits. This gives her no advantage in stand-alone applications of bit commitment, such as making a secret prediction, but it does mean that one cannot assume that in a task involving bit commitment subprotocols, any unopened bit commitments necessarily had definite classical bit values, even if all unveiled bit commitments produced valid classical unveilings.

³Following Ref. [6], another discussion of security definitions from a somewhat different perspective was given in Ref. [134].

maximum is taken over all strategies S' that are possible according to some set of physical principles Σ .

Thus we will say that a general relativistic bit commitment protocol is Σ -secure against both Alice and Bob, for some set of physical principles Σ , iff the protocol has the following properties:

1. For any strategy S for Alice that is possible according to Σ , $p_0^\Sigma(S) + p_1^\Sigma(S) < 1 + \epsilon(N)$ and $\epsilon(N) \rightarrow 0$ as $N \rightarrow \infty$,⁴ where N is a variable security parameter of the protocol.⁵
2. For any strategy for Bob that is possible according to Σ , if Alice's agents choose not to unveil, the probability of any of Bob's agents correctly guessing the committed bit at any point in spacetime is $\frac{1}{2}$.

Notes

1. Since we are dealing with relativistic protocols, we will assume that Σ always includes, at minimum, the relativistic no-signalling principle. Thus for any allowed Σ , given a protocol is Σ -secure against Bob, may invoke the no-signalling principle to conclude that there is no strategy that is possible according to Σ such that when Alice does choose to unveil, Bob can guess Alice's commitment at some point that does not lie in the future lightcone of the unveiling points.
2. We will say that a protocol is *unconditionally* secure if Σ includes only:
 - (a) Standard cryptographic assumptions, including the possibility of secure classical computation and the existence of secure laboratories

⁴We need only consider the limit as N goes to infinity, because if the protocol is secure in some other limit, either we can simply set the parameter equal to the (finite) limit as part of the definition of the protocol, or we can redefine the parameter so that unconditional security is obtained in the limit as the new parameter goes to infinity - for example, if the protocol is unconditionally secure in the limit as some quantity N' goes to zero, either we just set that quantity to zero in the definition of the quantity, or if it is not possible to set N' exactly equal to zero, we define a new parameter $N'' = \frac{1}{N'}$ so that the protocol is unconditionally secure in the limit as N'' goes to infinity.

⁵Usually, the security parameter will be the length of the bit strings C_0 and C_1 and/or be the size of the Hilbert space of the systems Q_0 and Q_1 , but we allow for other possibilities. One could also imagine protocols that employ more than one security parameter, but most existing protocols require only a single parameter and for simplicity we will conform to that standard here.

whose resident cryptographer controls all inward and outward information flows.

- (b) Known physical laws (in particular, the laws of quantum mechanics and special relativity).
 - (c) Stipulations about the constitution of the devices used in the protocol.
3. We will say that a protocol has *device-independent* security if Σ includes only a) and b) above.
 4. In the protocols we consider in chapters 7 and 8 Alice has one committing agent, A_c , and two unveiling agents, A_0 and A_1 , who are responsible for unveiling commitments to 0 and 1 respectively. In certain contexts it may be desirable for Alice to have the option of committing to neither bit value, and hence an additional security criterion may be required for protocols of this type: if A_c does not make a valid commitment to bit value b , and A_b follows the unveiling protocol and $A_{\bar{b}}$ does not, then Bob's agents, at any point in spacetime, should gain no information about whether A_c committed to bit value \bar{b} or declined to make a valid commitment. With simple modifications, the protocols set out in chapters 7 and 8 also satisfy this stronger criterion.
 5. One might reasonably choose to relax the requirement for security against Bob so that the probability of any of Bob's agents guessing the committed bit need only lie in some range $[\frac{1}{2} - \chi(N'), \frac{1}{2} + \chi(N')]$ such that $\chi(N') \rightarrow 0$ as $N' \rightarrow \infty$, where N' is a variable security parameter of the protocol. However, all the protocols we describe are able to meet the security requirement given above and so we have chosen to define security with respect to this stronger criterion. An advantage of this definition is that for protocols which satisfy it, Bob learns strictly no information in the course of a single bit commitment protocol, and therefore we do not need to be concerned about accumulation of information in protocols where large numbers of bit commitments are composed, such as bit string commitments or extended bit commitment protocols as in ref [13].
 6. In addition to the security features we have just described, note that any useful bit commitment protocol must also have the property that Alice can

reliably make her desired commitment - a property sometimes described as ‘completeness.’

Definition 4. Completeness: For any $y \in \{0, 1\}$, for any choice of B_c and $\psi(X_c)$, there exists a set of actions Alice can take at or before the point (t_c, x_c) such that for any allowed choice of B_y and $\psi(X_y)$, she can return a string C_y and a system Q_y ensuring that with probability close to 1 Bob obtains a measurement result M_y which leads to $V_y = 1$.

Geometry Obviously, in any bit commitment protocol the unveiling should take place ‘later’ than the commitment. There are several ways for this desideratum to be realised in the relativistic context. One possibility is that the unveiling points are in the lightlike causal future of the commitment point (e.g. see [165, 173]), and hence the statement that the unveilings are later than the commitment is true independent of the frame. We call these *lightlike causal* (LC) relativistic bit commitments.

However, we also wish to consider protocols in which the unveiling points are spacelike separated from the commitment point. The most obviously interesting case is that in which all unveiling points are later than the commitment point with respect to some fixed frame F ; we call such protocols *fixed frame positive duration* (FFPD) relativistic bit commitments. One motivation for considering this case is that it allows us to define sequences of protocols in which the unveiling points tend towards the future light cone, and so to relate LC and FFPD commitments. Another is that there are many practical situations – such as protocols carried out on terrestrial computer networks – in which there is a generally agreed (approximately) inertial frame and time coordinate. In such scenarios, commitments are potentially useful provided they have a positive duration with respect to this coordinate.

All the protocols that we set out in chapters 7 and 8 will use the same arrangement of agents in spacetime: we have two distinct unveiling points (t_0, x_0) and (t_1, x_1) , both spacelike separated from (t_c, x_c) and from each other (see fig 6.1). The distance between x_c and x_0 is equal to d and the distance between x_c and x_1 is also equal to d . For these protocols both Alice and Bob must have agents located in separate secure laboratories adjacent to each of the points (t_c, x_c) , (t_0, x_0) and

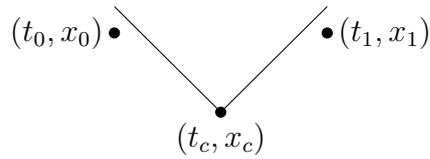


Figure 6.1: Diagram showing the arrangement in spacetime of the commitment and unveiling points for the protocols described in chapters 7 and 8. Diagonal lines indicate the boundary of the future lightcone of the point (t_c, x_c) .

(t_1, x_1) ; we refer to the agents adjacent to (t_c, x_c) as A_c and B_c , and those adjacent to (t_i, x_i) as A_i and B_i . Although it is not necessary for much of our discussion, we assume that t_0 and t_1 are strictly greater than t_c , so these are FFPD relativistic bit commitments with respect to the inertial frame used to define the coordinate system.

6.2 KCEKQS

As noted in section 2.4, we will work not with traditional zero-knowledge-proving, but with a generalisation that we refer to as *knowledge-concealing evidencing of knowledge about a quantum state* (KCEKQS). A KCEKQS protocol requires Alice to give Bob evidence that she has some form of knowledge about a quantum state whilst concealing as much information about the state as possible. Ideally, a successfully completed protocol should give Bob a significant amount of evidence without assuming Alice's honesty. Ideally, too, the protocol should ensure a small bound on the information about the state obtainable by Bob without assuming Bob's honesty or the successful completion of the protocol. At present we will define the protocol so that Alice has no option to abort the protocol, although we will return to that possibility in chapter 9 and in the supplementary information.

We assume that each party has trusted error-free devices in their own laboratory and that there are trusted error-free classical and quantum communication channels between the laboratories.

Definition 5. Non-relativistic Knowledge-Concealing Evidencing of Knowl-

edge of a Quantum State (Non-relativistic KCEKQS): A protocol involving two mistrustful parties, Alice and Bob, occupying disjoint secure laboratories. Bob begins in possession of a quantum system Q_B prepared in some pure state $\eta = |\eta\rangle\langle\eta|$ which is drawn uniformly at random from the Hilbert space of Q_B and which is not known to Bob. The protocol requires Alice and Bob to act on alternate rounds and terminates after a fixed finite number of rounds; at each round, Alice and/or Bob may be required to carry out unitary operations and/or measurements on a quantum system in their possession, and/or to send classical and/or quantum information to the other party. The final round of the protocol requires Bob to generate a bit value b from the classical and quantum information in his possession, where $b = 1$ means Bob accepts that Alice has provided evidence of knowledge about η , and $b = 0$ means Bob rejects Alice's evidence.

A *relativistic* KCEKQS protocol is defined similarly, but in the relativistic context each party may have *several* trusted agents occupying separate secure laboratories, with secure communications between them, and the protocol will stipulate that these agents must carry out their unitary operations and/or measurements on quantum systems and/or classical communications and/or quantum communications within their agreed location regions and within agreed time intervals.

Security Definitions As in the case of bit commitment, one must be careful to specify precisely what guarantees a KCEKQS protocol is intended to provide to the two participants. We will characterise the effectiveness of a KCEKQS protocol by three parameters ϵ_C , ϵ_K and ϵ_S ; we discuss other relevant features of KCEKQS protocols in chapter 9. When evaluating these parameters for specific protocols, we will mostly consider the ideal case of error-free devices and channels. In realistic implementations, channel noise, device errors and losses may alter the parameter values,⁶ but the no-go theorems we set out in chapter 9 still hold in reasonable models of noise, errors and losses, so long as these are uncorrelated with η and with any knowledge Alice may have of or about η .

Let $p(1)$ and $p(0)$ be the probabilities for the two values of the bit generated by Bob at the final round of the protocol, where the probabilities are calculated by

⁶Typically, uncorrected noise, errors and losses will increase ϵ_C and decrease ϵ_K , and uncorrected losses will decrease ϵ_S .

averaging over η under the stated conditions. Then we define our parameters as follows:

- **Completeness:** ϵ_C is the value of $1 - p(1)$ in the case where Alice has a precise classical description of η , assuming that both parties perform the protocol correctly.
- **Soundness:** ϵ_S is the supremum of $p(1)$ over all possible (honest or dishonest) strategies for Alice in the case where she has no classical or quantum information about the state η , assuming that Bob performs the protocol correctly.
- **Knowledge-concealing:** ϵ_K is the supremum of the expected squared fidelity $F^2(\eta, \phi) = |\langle \eta | \phi \rangle|^2$ over (honest or dishonest) strategies that give Bob the value of a pure state ϕ as a guess for η in the case where at the start of the protocol Bob knows nothing about the state η , assuming that Alice performs the protocol correctly.

Let ϵ_M be the supremum of the expected squared fidelity obtainable by Bob if he does not take part in the protocol and instead simply carries out quantum operations and measurements on Q_B . ϵ_M depends only on the dimension of the state in question and hence is not a parameter of the specific protocol: from previous work on quantum state estimation, we have that $\epsilon_M = \frac{2}{d+1}$ [225]. We call $\epsilon_K - \epsilon_M$ the *knowledge gain* available from a given protocol for a dishonest Bob, and we say a protocol is *zero-knowledge* if $\epsilon_K = \epsilon_M$. We say a protocol is *non-trivial* if $1 - \epsilon_C > \epsilon_S$. This definition of non-triviality is intended to characterise what might reasonably be considered to be a useful knowledge-evidencing protocol. To justify this choice, consider two possible hypotheses: the first is that Alice has no classical or quantum information correlated with η and the second is that Alice knows the classical value of η precisely and follows the protocol honestly. If the protocol allows Alice to attain a higher or equal value of $p(1)$ when the first hypothesis holds than the value attained when the second holds, then outcome 1 gives Bob no evidence to prefer the second hypothesis over the first. Our non-triviality condition excludes this possibility, meaning that outcome 1 gives Bob at least some evidence to prefer the second hypothesis over the first. Of course,

there are other possible hypotheses; for example, Alice could have some incomplete classical information about η , or some beliefs about η that she expresses in a probability distribution, or some quantum information correlated with η . Non-triviality does not necessarily imply that the outcome 1 gives Bob evidence that favours the hypothesis that Alice knows η precisely over any of these other hypotheses; see section 9.6 for further discussion of this issue.

Classification As defined, general KCEKQS protocols allow both classical and quantum communications in both directions. However, the protocols we discuss will all employ some more restricted class of communications, which provides a useful classification scheme. We will consider *classical* protocols, in which Alice and Bob employ only classical communication, *quantum A-to-B* protocols, which additionally allow quantum communications from A to B, and similarly *quantum B-to-A* protocols. Classical and quantum A-to-B protocols were discussed in ref [203], but to our knowledge our quantum B-to-A protocol is the first of its type, and we will show that (under certain assumptions) it performs significantly better than any of the existing classical or quantum A-to-B protocols.

Another common classification scheme describes zero-knowledge proving protocols in terms of how close they are to true zero-knowledge. Let p_x denote the probability distribution over the set of messages exchanged by Alice and Bob which is obtained if the fact to be proved is indeed true and a correct proof is indeed produced; let p_y be the probability distribution over the set of messages produced by a probabilistic polynomial-time machine which simulates the protocol conditional on the relevant fact being true. In a *perfect* zero-knowledge protocol, these distributions are identical; in a *statistical* zero-knowledge protocol, the distributions have negligible statistical difference between them; in a *computational* zero-knowledge protocol, samples from the two distributions are indistinguishable by any polynomial-time machine [201, 226]. Perfect and statistical zero-knowledge protocols may also be described as ‘information-theoretically secure,’ since their security depends only on information theory and not on any assumptions about computational hardness.

Chapter 7

Deterministic Relativistic Quantum Bit Commitment

We describe new unconditionally secure bit commitment schemes whose security is derived from Minkowski causality and the monogamy of quantum entanglement. We first describe an ideal scheme that is purely deterministic, in the sense that neither party needs to generate any secret randomness at any stage. We also describe a variant that is purely deterministic for the committer, requires only local randomness generation from the receiver, and allows the commitment to be verified in the neighbourhood of the unveiling point. We show that these schemes still offer near-perfect security in the presence of losses and errors, which can be made perfect if the committer uses an extra single random secret bit. We discuss scenarios where these advantages are significant.

Based on a paper co-authored with Adrian Kent [16]

7.1 Introduction

Existing relativistic classical and quantum bit commitment protocols [13,165,173] require at least one party to locally generate and then securely store and/or distribute secret classical random strings. While this is a reasonable capability to assume in many cryptographic contexts, it may not always be practical - for example, if protocols are being implemented over a network of many sites, it may

not be desirable to set up random number generators or secure classical memories at every site. One might at first think that quantum protocols cannot have any advantage here, since if a party can securely and reliably prepare, distribute and measure entangled quantum states, they can obtain secure classical random strings from those states as and when required. In many scenarios this argument may indeed apply. However, quantum information has security advantages compared to classical information, particularly when one considers a protocol as part of a larger cryptographic exchange. For example, if a party is concerned that there has been a security breach at one of their sites, they can check whether a distributed quantum state remains in the correct form, whereas they cannot tell for sure whether a purportedly secret distributed classical random string has been read at some location by an adversary.

This provides motivation for the development of relativistic quantum bit commitment protocols that do not depend so heavily for their security on trusted random numbers. In this chapter we set out two new quantum protocols that do indeed minimise the use of randomness - in fact, one of these protocols, in its ideal form, requires no randomness at all. In addition to the practical utility of these protocols, they are also theoretically interesting: the fact that relativistic quantum bit commitment can be implemented without the need for trusted randomness adds new nuance to our understanding of the relationship between physics and cryptographic primitives.

7.2 Protocols

For the sake of clarity, in the following protocols we describe a procedure in which Alice and her agents exchange qubits by secure physical transportation in the preparation phase. However, they may alternatively employ teleportation or a secure quantum channel without significantly altering the protocols' security. Likewise, Bob and his agents may exchange qubits by any secure means. Bob may also arrange to combine his qubits at a variety of locations, depending on where he wishes to verify the unveiled bit.

7.2.1 ETBC: Simple entanglement transfer protocol

1. A_c prepares a total of $2N$ Bell pairs in the state $\Psi^- = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$, labelling the first N pairs as (W_{0P}^j, W_{0Q}^j) and the second N pairs as (W_{1P}^j, W_{1Q}^j) , where $j \in [1, N]$. She gives the qubits W_{0Q}^j to A_0 , gives the qubits W_{1Q}^j to A_1 , and retains the qubits W_{0P}^j, W_{1P}^j .
2. A_0 and A_1 travel to the spatial locations x_0 and x_1 respectively. We assume that A_c, A_0 and A_1 have secure laboratories that protect their qubits, so Bob cannot interfere with them in any way after the initial preparation, including during the transportation phase.
3. **Commitment:** At time t_c , if A_c wishes to commit to bit value b , she gives B_c the qubits W_{bP}^j , for $j \in [1, N]$, labelled in sequence. Bob does not know whether he has received $\{W_{0P}^j\}$ or $\{W_{1P}^j\}$, so he simply labels these qubits as W^j .
4. **Unveiling:** For $i \in \{0, 1\}$, if the agent A_i wishes to unveil, then at time t_i she gives the labelled qubits W_{iQ}^j to Bob's agent B_i .

If the agent A_c wishes to unveil, then at some time later than t_c she sends to Bob's neighbouring agent a classical message stating the bit value b . (If preferred, one or both of the A_i may be assigned to play this role instead).

Note that in principle the agents A_c, A_0 and A_1 may make these decisions independently. To coordinate them and ensure that all or none unveil, Alice needs to give them instructions in advance. These instructions could depend on separate events in the past lightcones of the unveiling points, if Alice knows these events will be correlated.

5. **Verification:** Once at least one of Bob's agents knows the bit value b to which Alice purports to have made a commitment, Bob's agents share this information amongst themselves and securely transmit to a single agent all the qubits given to B_c and B_b . The receiving agent then carries out projective measurements in the Bell basis on the qubits (W^j, W_{bQ}^j) for each $j \in [1, N]$. The verification function returns 1 iff all the outcomes correspond to the Bell state Ψ^- for all j . (This verification step can be carried

out at a location of Bob's choice: for example, it could be performed by an agent half-way between x_c and x_b).

We prove that ETBC is Σ -secure against Alice and Bob, where Σ includes the validity of quantum mechanics and special relativity and the assumption that Bob's measuring devices are reliable, together with standard cryptographic assumptions as set out in chapter 2.1. ETBC is therefore unconditionally secure, but since Σ includes an assumption about the functioning of Bob's measuring devices, ETBC does not have device-independent security.

Security against Alice Write the Hilbert spaces for the N qubits held by B_0 , B_1 and B_c as \mathcal{H}_0 , \mathcal{H}_1 and \mathcal{H}_c respectively, and write $\mathcal{H} = \mathcal{H}_0 \otimes \mathcal{H}_1 \otimes \mathcal{H}_c$. Bob tests for a purported commitment to zero by a measurement defined by the projection

$$P_0 = \otimes_{j=1}^N (\mathbb{I}_1^j \otimes |\Psi_{-}\rangle_{0c}^j \langle \Psi_{-}|_{0c}^j).$$

Bob tests for a purported commitment to one by a measurement defined by the projection

$$P_1 = \otimes_{j=1}^N (|\Psi_{-}\rangle_{1c}^j \langle \Psi_{-}|_{1c}^j \otimes \mathbb{I}_0^j).$$

Here \mathbb{I}_k^j is the identity operator on the j -th qubit in \mathcal{H}_k and $|\Psi_{-}\rangle_{kl}^j$ is a Bell state of the j -th qubits in $\mathcal{H}_k \otimes \mathcal{H}_l$. The operator $Q = P_0 P_1$ can be written as $Q = \otimes_{j=1}^N Q_j$, where Q_j acts on the triple of j -th qubits from each Hilbert space and has operator norm $\|Q_j\|_{op} = \frac{1}{2}$; hence Q has operator norm $\|Q\| = 2^{-N}$.¹

For any state $|\psi\rangle$ defining triples of N qubits that Alice might hand over to B_c , B_0 and B_1 , we thus have

$$\begin{aligned} \|Q|\psi\rangle\|_{op} &= \|P_0|\psi\rangle - P_0(1 - P_1)|\psi\rangle\|_{op} \\ &\geq \|P_0|\psi\rangle\|_{op} - \|P_0(1 - P_1)|\psi\rangle\|_{op} \\ &\geq \|P_0|\psi\rangle\|_{op} - \|(1 - P_1)|\psi\rangle\|_{op} \\ &\geq (p_0^{1/2} - (1 - p_1)^{1/2}) \end{aligned}$$

¹The operator norm of an operator C on a normed space V is defined by $\|C\|_{op} = \max_{x \in V} \frac{|Cx|}{|x|}$ where $|\cdot|$ denotes the norm associated with the space V [227]. In the quantum context, the operator norm is equal to the greatest probability with which an outcome associated with the operator in question can occur in any measurement performed on any quantum state.

where p_0 and p_1 are the respective probabilities that Alice successfully persuades Bob that 0 and 1 was unveiled if she prepares the state $|\psi\rangle$.

It follows that $p_0 + p_1 \leq 1 + 2^{-N+1} + 2^{-2N}$. As this inequality holds for any possible state $|\psi\rangle$, it implies that the protocol is Σ -secure against Alice with security parameter N .

Security against Bob At commitment, Bob receives a set of N qubits entangled with another N qubits not in his possession. They have the same reduced state (a uniform mixture) regardless of the committed bit. Thus he cannot obtain any information about the bit before unveiling, so the protocol is Σ -secure against Bob.

7.2.2 ETRBC: Entanglement transfer protocol with randomisation

A possible disadvantage of the protocol we have just set out is that since Bob does not initially know whether Alice will choose to unveil a commitment to 0 or to 1, and the no-summoning theorem [228] prevents him from having the qubits W^j available at spacelike separated points along the different directions associated with 0 and 1, the time between Alice's unveiling and the earliest time at which Bob can *verify* her commitment is twice as long for this protocol as compared to other protocols with a similar geometry, such as the protocol described in ref [165]. In time-sensitive situations this may be a disadvantage.

This is what motivates the second version of our protocol. It eliminates this potential drawback by allowing each B_i to perform a verification test at the earliest possible point, i.e. as soon as a light signal from (t_c, x_c) can reach x_i . In order to achieve this, Alice proceeds just as above, but now B_c randomly selects half the qubits given to him to send securely to B_0 , sending the other half to B_1 . This allows both B_0 and B_1 to directly test the bit value as soon as they receive these qubits, rather than waiting for Alice to confirm her choice of bit:

1. A_c prepares $2N$ Bell pairs, (W_{0P}^j, W_{0Q}^j) and (W_{1P}^j, W_{1Q}^j) with $j \in [1, N]$, in the state Ψ^- . She gives the qubits W_{0Q}^j to A_0 and the qubits W_{1Q}^j to A_1 .

We take N even for simplicity. (The protocol can easily be varied to also allow for odd N).

2. A_0 and A_1 travel to the spatial locations x_0 and x_1 respectively. We assume that A_c , A_0 and A_1 have secure laboratories that protect their qubits, so Bob cannot interfere with them in any way after the initial preparation, including during the transportation phase.
3. **Commitment:** At time t_c , if A_c wishes to commit to bit value b , she gives B_c the qubits W_{bP}^j , for $j \in [1, N]$, labelled in sequence. Bob does not know whether he has received $\{W_{0P}^j\}$ or $\{W_{1P}^j\}$, so he simply labels these qubits as W^j .
4. **Distribution:** B_c sends a randomly selected size $N/2$ subset J_0 of the qubits $\{W^j\}$ to B_0 and the remaining subset, J_1 , to B_1 . All qubits are sent with the corresponding labels j .
5. **Unveiling:** For $i \in \{0, 1\}$, if the agent A_i wishes to unveil, then at time t_i she gives the labelled qubits W_{iQ}^j to Bob's agent B_i . (A_c and/or either or both of the A_i may also send to Bob's neighbouring agent a classical message stating the bit value b , although it is not necessary in this protocol. In any case, as in the previous protocol, some advance instructions from Alice are needed to ensure any unveiling decisions are coordinated.)
6. **Verification:** Once he has received the qubits sent by B_c , each B_i carries out projective measurements in the Bell basis on the qubits (W^j, W_{iQ}^j) for each $j \in J_i$. The verification function V_i returns 1 iff all the outcomes correspond to the Bell state Ψ^- for all $j \in J_i$.

We prove that this protocol is Σ -secure against Alice and Bob, using the same Σ as for ETBC. Thus ETRBC is likewise unconditionally secure but does not have device-independent security.

Security against Alice Write the Hilbert spaces for the N qubits held by B_0 , B_1 and B_c as \mathcal{H}_0 , \mathcal{H}_1 and \mathcal{H}_c respectively, and write $\mathcal{H} = \mathcal{H}_0 \otimes \mathcal{H}_1 \otimes \mathcal{H}_c$. B_0

tests for a commitment of zero by a measurement defined by the projection

$$P_0^{J_0} = \otimes_{j \in J_0} (\mathbb{I}_1^j \otimes |\Psi_-\rangle_{0c}^j \langle \Psi_-\rangle_{0c}^j).$$

B_1 tests for a commitment of one by a measurement defined by the projection

$$P_1^{J_1} = \otimes_{j \in J_1} (|\Psi_-\rangle_{1c}^j \langle \Psi_-\rangle_{1c}^j \otimes \mathbb{I}_0^j).$$

Suppose that Alice prepares a state $|\psi\rangle$ such that the probability of passing the test for zero is $p \geq p_0$. Then there must be at least one subset J_0 for which:

$$p_0^{J_0} = \langle \psi | P_0^{J_0} | \psi \rangle \geq p_0$$

Consider any subset J'_0 such that $J_0 \cap J'_0 \leq N/3$.

By a similar argument to that above, we obtain

$$\|P_0^{J_0} P_1^{J'_0}\|_{op} \leq 2^{-N/6}$$

and

$$p_1^{J'_0} \leq 1 + 2^{-N/6+1} + 2^{-N/3} - p_0^{J_0} \leq 1 - p_0 + 2^{-N/6+1} + 2^{-N/3}.$$

The proportion of subsets J'_0 with $J_0 \cap J'_0 > N/3$ falls off exponentially with N : to leading order it is bounded by $(N/6)(2^{-10/6}3)^N$. Hence the probability p_1 of passing the test for bit value 1 is bounded by $p_1 \leq 1 - p_0 + 2^{-N/6+1} + 2^{-N/3} + O(N/6(2^{-10/6}3)^N)$, so the protocol is Σ -secure against Alice with security parameter N .

Security against Bob As before, at (t_c, x_c) Bob receives a set of N qubits entangled with another N qubits not in his possession. They have the same reduced state (a uniform mixture) regardless of the committed bit. He thus cannot obtain any information about the bit before he receives data from the unveiling points, so the protocol is Σ -secure against Bob.

7.3 Errors and Losses

In any realistic implementation, Alice's state preparation and Bob's measurements will be imperfect and their communication channels and storage devices will have some noise and losses. To show that the protocols remain feasible with imperfect technology, we need versions adapted to allow for some non-zero level of errors and losses.

We first assume that Bob follows the protocol and measures each singlet separately, and that the errors and losses for each singlet are small and statistically independent. It then follows that for protocol ETBC Bob can test for a purported commitment of zero, with negligible probability of getting a false negative result, by checking that he gets positive answers for a proportion $(1 - \epsilon)N$ of tests for the singlet $|\Psi_{-}\rangle_{20}$, where $\epsilon > 0$ is small. The error model implies that the probability of a state $|\psi\rangle$ passing the test is no more than $|P_0^\delta |\psi\rangle|^2 + \gamma(\delta, N)$, where $P_0^\delta = \sum_{m=(1-\delta)N}^N P_m^0$, where the function γ and the parameter $\delta > \epsilon$, which is also small, is chosen so that $\gamma(\delta, N) \rightarrow 0$ as $N \rightarrow \infty$. The operator P_m^0 is the projection onto the subspace of states spanned by states of the form $\otimes_{i=1}^N |\Psi_i\rangle_{20} |\Phi_i\rangle_1$, where the $|\Psi_i\rangle_{20}$ are Bell states, of which precisely m are $|\Psi_{-}\rangle$, and the $|\Phi_i\rangle_1$ are arbitrary qubits in \mathcal{H}_1 .

Bob similarly tests for a purported commitment to bit value 1 by checking that he gets positive answers for a proportion $(1 - \epsilon)N$ of tests for the singlet $|\Psi_{-}\rangle_{12}$. The probability of a state $|\psi\rangle$ passing this test is, up to negligible quantities, no more than $|P_1^\delta |\psi\rangle|^2 + \gamma(\delta, N)$, where $P_1^\delta = \sum_{m=(1-\delta)N}^N P_m^1$ is defined similarly to P_0^δ .

The operator P_0^δ can be written as a sum of $\sum_{x=0}^{N\delta} \binom{N}{N-x} 3^x$ terms involving one-dimensional projectors onto tensor products of Bell states in $\mathcal{H}_0 \otimes \mathcal{H}_c$, tensored with the identity on \mathcal{H}_1 . The operator P_1^δ can be written similarly, using Bell state projections on $\mathcal{H}_1 \otimes \mathcal{H}_c$. The operator $Q^\delta = P_0^\delta P_1^\delta$ can thus be written as a sum of $(\sum_{x=0}^{N\delta} \binom{N}{N-x} 3^x)^2$ rank one operators, each of which has operator norm no more than $2^{-N+2\delta N}$. This gives the (weak, but adequate for our purpose) bound $\|Q^\delta\|_{op} \leq 2^{-N+2\delta N} 3^{2\delta N} (N\delta + 1)^2 (C_{N-N\delta}^N)^2$, which tends to zero for large N and fixed small δ . The security argument then runs as before. Moreover, a similar analysis applies to protocol ETRBC, so we can reasonably say that both

protocols remain Σ -secure even when we take account of small errors and losses.

For completeness, we should note another possible security issue. If the errors in Alice's singlet state preparations vary over time in some predictable way, then the reduced density matrices for the states handed over to B_c by A_c may also vary predictably. Given a deterministic protocol, we have to assume that the order in which A_c labels the singlets after producing them is public information. B_c might then be able to infer some information about the committed bit by measuring these states without waiting to combine them with states returned by the A_i . This may not seem a significant practical worry, since in practice one might reasonably expect the predictable component of any variation in Alice's preparation devices to be very small. Moreover, Alice might employ various strategies to reduce it further: for example, the information revealed by a monotonic drift of some parameter over time could be greatly reduced by taking the odd time ordered singlets produced (the 1st, 3rd, and so on) to be the first N for the protocol, and the even ordered to be the second N . Nonetheless, in a context where this possibility were regarded as a serious concern, the problem could be eliminated by requiring agent A_c to group the states into two batches of N singlets by some deterministic method, then decide randomly which batch is labelled from 1 to N and which from N to $2N$. Note that this would require her to generate and store a single random bit, which would have a small impact on the security advantages described in the next section.

7.4 Summary

The first protocol we have described here has a theoretically interesting advantage over any previous relativistic bit commitment protocol in that it is deterministic: neither party needs to make any random choices of classical data or quantum states. It thus satisfies the strongest possible form of Kerckhoff's cryptographic principle, which states that a cryptographic system should be secure even if everything about it except the choice of key is public knowledge [229]. This is desirable because generating secure random numbers is itself a nontrivial cryptographic problem [104] and the need to generate large strings of random numbers is therefore a significant security loophole for many existing bit commitment pro-

protocols [6, 134, 165, 173, 185, 186]. Eliminating the need for trusted randomness removes some potential security issues and may result in a protocol that requires fewer resources.

As noted above, however, these advantages come at a price in ETBC: Bob cannot verify Alice's commitment as soon as his agents B_0 and B_1 receive Alice's unveiling data, but must wait until that data can be brought together with the commitment data initially handed over at P_c . In scenarios where there are significant time constraints, it may therefore be preferable to use the protocol ETRBC. This second version of the protocol is no longer entirely deterministic: B_c needs to be able to generate a classical random string that is secure, at least in the sense that Alice cannot predict it in advance, although it does not matter if Alice learns the string immediately. However, this is still less demanding than requiring Bob to generate a secure random quantum state or sequence of states and keep its classical description secure [165, 173]. The protocol also has an advantage over purely classical relativistic protocols [13] in that Alice does not need to generate any secure random data.

In addition to these practical advantages, it is of theoretical interest to see that there exists a secure relativistic bit commitment protocol that does not depend on generating secure random numbers. To the best of our knowledge, all previous relativistic bit commitment protocols required at least one agent to generate strings of random numbers, and therefore one might have been forgiven for supposing that the use of trusted randomness was an essential feature of any secure relativistic bit commitment protocol. Our work demonstrates that this is not so. However, there is an important nuance to be kept in mind: although our protocols do not require the generation of trusted random numbers, they *do* require the production of trusted entangled quantum states, and if an agent can generate trusted entangled quantum states they can always use them to generate trusted randomness. Thus although we have shown that it is possible to achieve secure bit commitment without generating trusted random numbers, we have not shown that it is possible to achieve secure bit commitment without *being able* to generate trusted random numbers. The relationship between relativistic bit commitment and randomness is certainly a subtle one that would merit further investigation. In particular, it would be interesting to see if it is possible to come up with some relativistic bit

commitment protocol - whether classical or quantum - that can be implemented even if neither party has access to any resources that would allow them to generate trusted random numbers.

Chapter 8

Device-Independent Relativistic Quantum Bit Commitment

We examine the possibility of device-independent relativistic quantum bit commitment. We describe relativistic quantum bit commitment schemes that offer device-independent security and are even secure against hypothetical post-quantum adversaries subject only to the no-signalling principle. We compare our protocols to a relativistic classical bit commitment scheme with similar features, and note some possible advantages of the quantum schemes.

Based on a paper co-authored with Adrian Kent [17].

8.1 Introduction

The protocols we described in chapter 7 require Bob to rely on his devices to implement projective measurements for Bell states, up to known small levels of losses and errors, and therefore these protocols are secure only insofar as Bob can be confident that his devices do indeed implement the correct measurements at all times. But if quantum cryptography were to go into widespread use, the reality is that most people taking part in cryptographic protocols would not be able to build their own quantum measuring instruments, so Bob would have to trust some outside manufacturer to produce these devices for him. What if the manufacturer conspires with Alice to include hidden features designed to give the

appearance of normal functioning while actually compromising security? What if the manufacturer is simply incompetent and his measuring devices are wildly inaccurate? And what is Bob to do if he dares have the heretical thought that quantum mechanics might not be completely correct and Alice might be using some more advanced theory to cheat him?

In view of these potential security loopholes, we have good reason to aspire to ‘device-independent’ protocols, i.e. protocols whose security does not depend on any assumptions about the properties of the devices employed to prepare states and perform measurements, except that they are constrained by some set of physical laws. In particular, there has been significant progress on the development of quantum cryptographic protocols where the only assumption made about the quantum devices employed is that they obey no-signalling constraints. [76, 103, 189–198].¹ To achieve even higher levels of security, the no-signalling requirement may be enforced by designing protocols ensuring that relevant operations take place at a spacelike separation, so security is essentially guaranteed by the causal structure of spacetime. Since the no-signalling principle is one of the most robust principles in all of science, appearing under various guises in special relativity, quantum mechanics and quantum field theory, participants may have a high level of confidence in the security of these protocols, even if they do not trust the manufacturers of the devices they are employing, and even if they are not confident in the veridicality of any one of these theories individually.²

In this chapter, we propose several new relativistic quantum bit commitment protocols that are device-independent in this strong sense. Full device-independence is an even more stringent constraint in relativistic quantum cryptography than in standard quantum cryptography, because many device-independent quantum cryptographic protocols rely on performing a large number of tests to establish the typical behaviour of the devices employed and then randomly selecting which tests will actually be used in the protocol. This method cannot straightfor-

¹Of course, these protocols still require standard classical cryptographic assumptions as set out in chapter 2.1

²We note that certain types of protocols may not translate well to this context. For example, if one assumes that one’s devices may be sending out messages to one’s adversaries, limited only by the speed of light, then it is never going to be possible to generate a secret shared key whose secrecy can be guaranteed for longer than some short finite period.

wardly be applied to relativistic cryptographic protocols, because these protocols require specific configurations of agents and devices in spacetime, arranged to ensure that the agents cannot exchange data during critical parts of the protocol. Thus if two separated agents are using devices that are supposed to produce outputs that are correlated in some way, they may be unable to check this behaviour in the course of the protocol; performing tests of the devices before the protocol is not sufficient to guard against this, because we must take account of the possibility of spacetime *location attacks*, where devices that are able to track their own space and time coordinates are programmed to reproduce the expected behaviour when they are originally tested but not when they are implementing a critical part of the protocol at specified spacetime locations. For example, if a particularly critical bit commitment is going to be initiated at noon GMT by devices in Cambridge and Auckland, one of Alice's devices might be programmed so that it usually behaves as Alice expects but, if and only if it is in Cambridge at noon, it gives her an output that should encrypt her committed bit but actually immediately reveals it to Bob's Cambridge agent. Moreover, comparing outcomes after the protocol has terminated may suffice to *detect* that a location attack has taken place, but this will usually be too late to prevent sensitive information from being revealed to outsiders. Thus location attacks are potentially a very serious problem for relativistic device-independent quantum cryptography; nonetheless, the protocols we describe in this chapter resist such attacks.

The existence of device-independent quantum relativistic bit commitment protocols is interesting theoretically, but from a practical point of view it is also important to determine whether our protocols have significant advantages over existing protocols. Since we have defined device-independence with respect to *quantum* preparation and measurement devices, all *classical* relativistic bit commitment protocols, such as those set out in refs [12, 13], are automatically device-independent in this sense simply in virtue of not employing any quantum preparation or measurement devices. Thus device-independence alone does not give our protocols any advantages over classical relativistic bit commitment protocols, so we must consider whether the quantum properties of these protocols offer any further cryptographic advantages. Existing classical protocols involve configurations of agents in spacetime somewhat different from the device-independent

protocols we consider here, but in order to facilitate comparison, in section 8.4.1 we describe a new, unconditionally secure classical relativistic bit commitment protocol that uses the same configuration of agents as our quantum protocols. We then demonstrate that our quantum protocols have several potential advantages over this particular classical protocol.

8.2 Protocols

Definitions, notation, geometry and idealisations are as set out in section 2.1. In addition, for the security proofs we employ a binary notation, denoting measurement directions X, Y by 0, 1, measurement directions X', Y' also by 0, 1, and measurement outcomes $+1, -1$ also by 0, 1. A single Bell experiment is then described by two measurement settings $(x, y) \in \{0, 1\}$ and two measurement results $(t, s) \in \{0, 1\}$. We are interested in the *CHSH game score* $R(t \oplus s = xy)$, which is the number of experiments in the set for which the settings (x, y) and outcomes (t, s) satisfy the condition $t \oplus s = xy$. In the case where the settings are random and independent, the CHSH game score is asymptotically related to the average value of the CHSH observable $(XX' + XY' + YX' - YY')$: for a set of N experiments, this value is equal to is $\frac{4}{N}(2R(t \oplus s = xy) - N)$ to leading order in N . However, we do not assume independent random settings in all of our protocols.

8.2.1 CHSH 1: CHSH test protocol with fixed directions

1. Alice and Bob agree on a set of four directions X, Y, X', Y' such that all four directions lie in a single plane, X is orthogonal to Y , X' is orthogonal to Y' , X' is separated from X by $\frac{\pi}{4}$ and Y' is separated from Y by $\frac{\pi}{4}$. They also agree on a bit string L^0 with bitwise complement L^1 .
2. A_c instructs her devices to prepare $2N$ Bell pairs (W_a^i, W_b^i) in the singlet state Ψ^- , then randomly draws an integer $w \in \{0, 1\}$. If $w = 0$ she gives the systems $[W_b^1, \dots, W_b^N]$ to an agent A_0 and the remaining systems $[W_b^{N+1}, \dots, W_b^{2N}]$ to an agent A_1 ; if $w = 1$ she instead gives the systems $[W_b^{N+1}, \dots, W_b^{2N}]$ to A_0 and the systems $[W_b^1, \dots, W_b^N]$ to A_1 .

3. A_0 and A_1 then travel to spatial locations x_0 and x_1 respectively. We assume that A_c , A_0 and A_1 have secure laboratories that protect their qubits, so B_c cannot interfere with them in any way after the initial preparation. In particular, A_0 and A_1 travel within secure laboratories. We also assume that the first and second sets of states (i.e. the purported half-singlet states labelled a and b respectively) are securely separated in sub-laboratories before A_c chooses the value of w . Thus, even if the a states are actually malicious devices operating within the laboratory, they obtain no information about the value of w or the destinations of the b states, and vice versa.
4. **Commitment:** At time t_c , B_c gives A_c a bit string L drawn uniformly at random from $\mathbb{Z}_2^{\otimes N}$. If A_c wants to commit to 0, she instructs her devices to measure each qubit $W_a^{j+wN} \in [W_a^{1+wN}, \dots, W_a^{N+wN}]$ in direction $(X)^{L_j}(Y)^{1-L_j}$. If A_c wants to commit to 1 she instead instructs her devices to measure each qubit $W^{N+j-wN} \in [W_a^{N+1-wN}, \dots, W_a^{2N-wN}]$ in direction $(X)^{L_j}(Y)^{1-L_j}$. In either case A_c immediately tells B_c her outcomes.
5. **Unveiling:** For $i \in \{0, 1\}$:

At time t_i , if agent A_i wishes to unveil, she measures each qubit $W_b^{j+wN+iN-2iwN} \in [W_b^{1+wN+iN-2iwN} \dots W_b^{N+wN+iN-2iwN}]$ in direction $(X')^{L_j^i}(Y')^{1-L_j^i}$. At time t_i she tells B_i her measurement outcomes.

Note that in principle the agents A_0 and A_1 may make their decisions about whether or not to unveil independently. If Alice wishes to coordinate them and ensure that all or none unveil, she needs to give them instructions in advance. These instructions could depend on separate events in the past lightcones of their unveiling decision points, if Alice believes these events will be correlated.

6. **Verification:** For $i \in \{0, 1\}$:

Bob's agents between x_c and x_i wait for the data from A_c and A_i , and then calculate the CHSH game score for the received data. The verification function returns 1 iff the score is greater than $N((2 + \sqrt{2})/4 - \xi)$. Here ξ is some predetermined small security parameter, chosen such that $\xi \gg N^{-1/2}$; this

ensures that the probability that N correctly prepared and measured singlets will fail Bob's test is suitably small.

We prove that CHSH1 is Σ -secure against Alice and Bob, where Σ includes the relativistic no-signalling principle, as well as the standard cryptographic assumptions set out in chapter 2.1. CHSH1 is therefore unconditionally secure and also has device-independent security.

Security against Alice A_0 , A_c and A_1 announce their outcomes at spacelike separation from one another during the protocol. They may pre-agree a collective strategy S , which may rely on shared quantum or (hypothetical) post-quantum no-signalling resources, but cannot involve signalling to one another during the protocol. We are interested in bounding the probability that A_0 and A_1 both succeed in passing Bob's tests for a valid commitment of 0 and 1 respectively.

Let $p_0^\Sigma(S)$ and $p_1^\Sigma(S)$ be defined as in chapter 6, and let $p_0^\Sigma(S) + p_1^\Sigma(S) = 1 + \epsilon_S(N)$. Note that the rules of the protocol allow for the possibility that Bob accepts valid commitments to both 0 and 1 if both tests are passed.

For any strategy S which Alice may employ up to and including the time of the commitment in the relevant fixed reference frame, and any subsequent choice of strategy S' by Alice's agents, $\epsilon_S(N) \leq \sum_L \frac{p_{L,L_0}^S}{2^N}$ where p_{L,L_0}^S is the probability that on a run of the protocol when Alice employs strategies S and S' and Bob chooses the bit string L , Alice and her agents produce three sets of outcomes O , O^0 and O^1 such that:

$$d(O^0 \oplus O, L^0 L) \leq N \left(\frac{1}{2} - \frac{1}{2\sqrt{2}} + \xi \right) \quad (8.1)$$

$$d(O^1 \oplus O, L^1 L) \leq N \left(\frac{1}{2} - \frac{1}{2\sqrt{2}} + \xi \right) \quad (8.2)$$

Here $d(x, y)$ denotes the Hamming distance between the bit strings x and y , which is the number of positions on which the strings differ [85]; $O^i \oplus O$ is the

string given by element-wise modular addition of O^i and O ; and $L^i L$ is the string given by element-wise multiplication of L^i and L .

These equations can be simultaneously satisfied only if $d(O^0 \oplus O^1, L(L^0 \oplus L^1)) \leq N(1 - \frac{1}{\sqrt{2}} + 2\xi)$. Since L^0 is the bitwise complement of L^1 , this implies that $d(O^0 \oplus O^1, L) \leq N(1 - \frac{1}{\sqrt{2}} + 2\xi)$.

Thus for given $O^0 \oplus O^1$, equations 8.1 and 8.2 may be satisfied simultaneously only if L lies within the Hamming ball H of radius $r = N(1 - \frac{1}{\sqrt{2}} + 2\xi)$ centred on $O^0 \oplus O^1$. For $r \leq \frac{N}{2}$, the volume of this ball is less than or equal to $2^{NH(r/N)}$ where H is the binary entropy, defined as $H(x) = -x \log(x) - (1-x) \log(1-x)$ [230].

$H(x) \leq 1$ with equality if and only if $x = \frac{1}{2}$, so for any $\xi < \frac{1}{2\sqrt{2}} - \frac{1}{4}$ we have $H(r/N) < 1$.

By the no-signalling principle, the probability distribution for $O^0 \oplus O^1$ is independent of L . Thus since L is chosen uniformly at random from $\mathbb{Z}_2^{\otimes N}$, we must have $\epsilon_S(N) \leq 2^{-N(1-H(r/N))}$, which for $H(r/N) < 1$ goes to zero as N goes to infinity. Hence the protocol is Σ -secure against Alice.

Security against Bob Alice's devices may perhaps have been designed by Bob in an attempt to cheat the protocol. If the devices do something other than performing CHSH measurements on a shared quantum singlet, A_c 's outputs may give B_c information about whether A_c measured the qubits $[W_a^1 \dots W_a^N]$ or $[W_a^{N+1} \dots W_a^{2N}]$. This may allow B_c to update his prior values (which are originally equiprobable) for $P(b|w)$, the conditional probability of the committed bit value b given the value of A_c 's randomly chosen bit w . However, by assumption, w is random and kept secret throughout the protocol. Before receiving unveiling data, Bob's estimate of $P(b) = \frac{1}{2}(P(b|0) + P(b|1))$ thus remains unaltered.

Security against Bob therefore relies on Alice being able to generate one unconditionally secure random bit per committed bit, and to store this bit securely in A_c 's laboratory during the protocol.

8.2.2 CHSH2: CHSH protocol with secret complementary bit strings

CHSH1 may straightforwardly be varied so that Bob keeps the bit strings L^0 and L^1 secret from Alice until the points Q^0 and Q^1 . We simply add an extra preparation step:

1. B_c chooses a length N bit string L^0 drawn uniformly at random from $\mathbb{Z}_2^{\otimes N}$. He communicates string L^0 to B_0 and its bitwise complement L^1 to B_1 .

We also alter the unveiling step:

1. **Unveiling:** For $i \in \{0, 1\}$:

At time t_i , agent B_i communicates the string L^i to agent A_i . If A_i wishes to unveil, she measures each qubit $W_b^{j+wN+iN-2iwN} \in [W_b^{1+wN+iN-2iwN} \dots W_b^{N+wN+iN-2iwN}]$ in direction $(X')^{L_j^i}(Y')^{1-L_j^i}$. She then broadcasts her measurement outcomes.

The remaining steps are identical to those for CHSH1.

We prove that CHSH2 is Σ -secure against Alice and Bob, using the same Σ as for CHSH1. CHSH2 is therefore unconditionally secure and also has device-independent security.

Security against Alice The only change in this protocol is that Alice is given less information, which can only decrease her probability of cheating successfully. Hence the security of this protocol against Alice follows immediately from the security of CHSH1 against Alice.

Security against Bob The proof of security against Bob is the same as for CHSH1.

8.2.3 CHSH 3

CHSH1 may also be varied so that the measurement directions for A_0 and A_1 are not fixed in advance. Instead, each B_i randomly selects a set of N measurement

directions to be used by A_i , so the string L^1 is no longer guaranteed to be the bitwise complement of L^0 .

The steps in this protocol are identical to those for CHSH1 except for the unveiling stage:

1. **Unveiling:** For $i \in \{0, 1\}$:

At time t_i , B_i gives A_i a bit string L^i drawn uniformly at random from $\mathbb{Z}_2^{\otimes N}$. If agent A_i believes Alice wishes to unveil, she immediately measures each qubit $W_a^{j+wN+iN-2iwN} \in [W_a^{1+wN+iN-2iwN} \dots W_a^{N+wN+iN-2iwN}]$ in direction $(X')^{L_j^i} (Y')^{1-L_j^i}$. She then broadcasts her measurement outcomes.

We prove that CHSH3 is Σ -secure against Alice and Bob, using the same Σ as for CHSH1. CHSH3 is therefore unconditionally secure and also has device-independent security.

Security against Alice As before, for any strategy S , let $p_0(S)$ and $p_1(S)$ be the probabilities that when Alice employs the strategy, Alice's agents convince Bob that they have validly unveiled 0 or 1 respectively according to the rules of the protocol, and suppose $p_0(S) + p_1(S) = 1 + \epsilon_S^2(N)$.

By the no-signalling principle, A_0 's success probability is independent of the choice of L_1 , and A_1 's success probability is independent of the choice of L_0 . Hence we may calculate $p_0(S)$ and $p_1(S)$ assuming that L_0 and L_1 are bitwise complements. We thus obtain the same bounds on $p_0(S) + p_1(S)$ as for the first two protocols.

Security against Bob The proof of security against Bob is the same as for CHSH1.

8.3 Extensions

8.3.1 Declining to commit

One possible security loophole in our analysis is that although we have shown that Bob will be unable to learn the value of Alice's commitment when she follows the instructions set out above, we have not considered what happens if Alice simply does not make a commitment. As we noted in chapter 2.1, in certain cryptographic contexts, it might be desirable for Alice to have the option of refraining from making a commitment without Bob being able to tell that she has done so. As written our protocols do not guarantee device-independent security in this scenario: A_c might decline to commit by simply returning a random string of N bits to B_c , but because the protocols is supposed to be device-independent, we must assume that Bob may have designed Alice's preparation devices, so he may be able to distinguish between randomly chosen bit strings and strings produced by measurements on Alice's Bell pairs, meaning that he may be able to detect if A_c tries this tactic to refrain from making a commitment.

However, our protocols can straightforwardly be altered to provide device-independent security in this scenario as well. To do this, we simply require that Alice and Bob simultaneously perform two runs of CHSH1, with the rule that Alice commits to bit value b in both protocols if she wishes to produce a valid commitment to b , but commits to different values in each protocol if she wishes to refrain from commitment. Bob will then accept that Alice has made a valid commitment to bit value b only if she unveils valid commitments to bit value b for both protocols. CHSH2 or CHSH3 may similarly be altered in this way.

Device-independent security is then assured, both in the case where Alice makes a commitment and in the case where she refrains from committing, because Bob obtains no information about the values of w used in either of the two protocols unless and until Alice unveils, so he cannot tell whether or not she made a valid commitment unless and until she unveils.

8.3.2 Errors and Losses

We now consider how to allow for the possibility of small errors in preparations, transmissions and measurements of quantum states. We consider an error model in which any such errors occur randomly and independently, and where the combined rate of errors and losses is bounded by a small parameter δ . In this model, the expectation values of the distances $d(O^i \oplus O, L^i L)$ are altered by no more than δN from the theoretically expected values, with variance no larger than $\delta^{1/2} N^{1/2}$.

Bob's verification tests involve a security parameter ξ , and the security proof for CHSH1 holds for

$$\xi < \frac{1}{2\sqrt{2}} - \frac{1}{4}. \quad (8.3)$$

In particular, if δ is agreed to be the maximum tolerable rate of errors and losses, and δ is small, Bob can use any value $\xi > \delta$ in his verification tests, consistent with (8.3), and still ensure security against Alice.

The security proofs of CHSH2 and CHSH3 do not depend on additional assumptions about errors and losses, and so similarly remain secure for small δ . Thus we can ensure that both protocols are secure against Alice provided the expected rate of errors and losses is small.

The proofs of security against Bob do not depend on any assumptions about errors or losses, and therefore no alteration is required for these proofs.

8.4 Discussion

Device memory attacks A general concern in device-independent quantum cryptography is the possibility that maliciously designed devices may keep records of their inputs and outputs and may make their future outputs depend on these data [196]. Our protocols are secure against these attacks for a single bit commitment, but if Alice reuses her devices for a sequence of commitments, our protocols do not protect her from the possibility that data released in a later commitment may give information about earlier commitments. This is not a concern

if each one of her commitments is unveiled before the next commitment phase begins, but in many common scenarios, some commitment data is unveiled and some should be kept secret indefinitely - for example, this is likely to be the case in protocols where a relativistic bit commitment protocol is extended to produce a bit commitment protocol of longer duration by performing a series of protocols over time [13]. Our protocols do not guarantee security for device reuse against completely general attacks in such scenarios.

However, our protocols *do* guarantee security for device reuse against memory attacks under two additional conditions: first, the devices are not sensitive to location and hence not able to carry out location attacks, and second, the agents A_i make coordinated decisions about whether to unveil. To see this, observe that if neither A_0 and A_1 accept any device input or supply any output unless they intend to unveil a bit, and if the devices are not sensitive to their location, then the devices can carry no information correlated with the choice of bit after a protocol in which nothing is unveiled, so they have no ability to leak information about b in future protocols.

Randomness We note that the unconditional, device-independent security of both protocols depends on the assumption that Bob has access to a device independent method of generating random numbers, which is secure in the sense that Alice cannot predict its output in advance. As noted in the previous chapter, generating trusted randomness is not a trivial problem, and thus the need for random numbers could be regarded as undermining our claim that these protocols have device-independent security. However, although there exists no device independent method of generating random numbers without any random seed, there do exist device independent methods of randomness expansion [103, 197], which can create a long random string from some small random seed. Since the random seed used in these protocols can be of a low quality of randomness (i.e. the expansion still works if the ‘random’ input is partially correlated with variables which may be known to the adversarial party [231]), it could be obtained by some classical quasi-random process that Alice is unlikely to be able to predict or control perfectly, such as a coin flip. Therefore device-independent randomness expansion suffices for effectively device-independent security of these protocols.

The CHSH2 protocol requires Bob to generate $2N$ random bits and the CHSH3 protocol requires him to generate $3N$ random bits, in contrast with only N bits in CHSH1. However there may be contexts in which CHSH2 and CHSH3 are preferred - for example, if Bob wishes to control the time of Alice's measurements, which might be particularly important for chained protocols. In the CHSH1 protocol, Alice may make the measurements whose results she unveils at t_0 and t_1 at any time, even before the start of the protocol, but in CHSH2 and CHSH3 she must wait until Bob tells her which measurement directions to use. In such contexts CHSH3 offers greater security to Bob, because in CHSH2, L_1 must be the bitwise complement of L_0 and therefore it follows from the no-signalling principle that B_0 and B_1 must agree on these strings at some time in the causal past of P . This is a potential security weakness, since if Alice can find out the strings L_0 and L_1 in advance of the commitment, then A_0 and A_1 can perform their measurements earlier than Bob expects.

8.4.1 Random code classical bit commitment protocol (RC-CBC)

In order to make a fair comparison between our device independent quantum protocols and classical relativistic protocols, we now describe a classical bit commitment protocol that uses the same geometry as our protocols:

1. Alice and Bob agree on a security parameter N ; for simplicity we take N to be even.
2. A_c generates two independent N -bit random classical strings, S^0 and S^1 , and for $i \in \{0, 1\}$ she securely shares string S^i with agent A_i .
3. **Commitment:** At t_c , B_c gives A_c a randomly chosen size $N/2$ subset J of $\{1, \dots, N\}$. A_c immediately responds by giving B_c the string subset $\{S_j^i : i = 0, 1; j \in J\}$. These string elements are sent with their labels i and j . To commit to bit value b , A_c also sends the complementary string subset $S_j^b = \{S_j^b : j \in \bar{J}\}$. These string elements are sent with their labels j , but not the value b . Hence B_c refers to this string as $S_{\bar{J}}$, with b unknown to him.

4. **Unveiling:** For $i \in \{0, 1\}$:

At time t_i , if agent A_i wishes to unveil, she gives B_i the complete string S^i .

5. **Verification:** On receiving the two substrings S_J^0 and S_J^1 , B_c checks that their Hamming distances satisfy $|d(S_J^0, S_J^1) - N/4| < CN^{3/4}$, where $C > 0$ is an agreed parameter independent of N . He shares this information with his agents. Bob's agent between x_c and x_i waits for the data from A_c and A_i . When he receives it, he applies a verification function that returns 1 iff B_c 's check gave a positive outcome and in addition, the string S^i reported by B_i is the union of the strings S_J^i and $S_{\bar{J}}$ reported by B_c , with matching labels for each element.

We prove that RCCBC is Σ -secure against Alice and Bob, using the same Σ as for CHSH1. RCCBC is therefore unconditionally secure and also has device-independent security.

Security against Alice A_0 , A_c and A_1 announce their outcomes at spacelike separation from one another during the protocol. They may pre-agree a collective strategy S , which may rely on shared quantum or (hypothetical) post-quantum no-signalling resources, but cannot signal to one another during the protocol. We are interested in bounding the probability that A_0 and A_1 both succeed in passing Bob's tests for a valid commitment of 0 and 1 respectively.

Let $p_0^\Sigma(S) + p_1^\Sigma(S) = 1 + \epsilon_S(N)$.

Hence for any strategy S , $\epsilon_S(N) \leq \sum_L \frac{p_J^S}{2^N}$ where p_J^S is the probability that when Alice employs strategy S and Bob chooses the subset J , Alice's agents produce strings S^0 and S^1 such that $|d(S_J^0, S_J^1) - N/4| < CN^{3/4}$ and also that $S_j^0 = S_j^1 = S_j$ for $j \in \bar{J}$.

If Alice's agents succeed in producing strings with these properties, they can infer that the complement \bar{J} of B_c 's chosen subset J lies within a subset of size no more than $3N/4 + CN^{3/4}$ of $\{1, \dots, N\}$. By the no-signalling principle, it follows that $\epsilon_S(N) \rightarrow 0$ as $N \rightarrow \infty$.

Security against Bob Alice's strings are randomly generated. Bob thus receives no information correlated with the bit value i unless and until at least one of the

A_i unveils.

Note that this security argument relies on Alice being able to generate $2N$ unconditionally secure classical random bits per committed bit, to share these bits between A_c and the A_i before the protocol, and to keep them securely in the agents' laboratories before and during the protocol.

8.4.2 Comparison

The classical and quantum relativistic bit commitment protocols described in this chapter all have the same configuration of agents for the two parties, and are all unconditionally secure. All are 'device independent' in the sense we use that term: while the quantum protocols do not require either party to trust the quantum devices used, they still require both parties to have trusted classical computing and memory devices, and thus have no advantage over the classical protocol in this respect.

The classical protocols and the CHSH1 quantum protocol both require B_c to generate a random string that Alice cannot predict in advance. The string may be generated at P_c , and immediately handed over to A_c , so it does not need to be kept secure or distributed securely to Bob's other agents. The CHSH2 and CHSH3 protocols both require Bob to generate some additional random bits, and in the case of CHSH2 some random bits must be distributed securely, so both protocols are more demanding in this sense than the classical one. Thus from Bob's point of view, the quantum protocols have no advantage over the classical one in terms of the amount of randomness required.

On the other hand, while the classical protocol requires Alice to generate two N -bit random strings, to share them securely between A_c and A_0 or A_1 , and to store them securely in the various agents' laboratories, the CHSH protocols each require A_c to generate and keep secure only a single random classical bit. If the devices are distinguishable, this bit is also shared and stored securely in A_0 and A_1 's laboratories, since the set of physical devices given to these agents depends on the bit value. Thus from Alice's point of view the quantum protocols do have a security advantage over the classical one in the sense of requiring less classical randomness.

One might worry at this juncture that there is no real advantage here, because, as in the previous chapter, as long as Alice is in possession of the large number of maximally entangled qubit pairs that is required to implement these protocols, she can always choose to use these qubit pairs to produce random numbers, and thus if Alice has the ability to implement CHSH1, CHSH2 or CHSH3, she also has the ability to generate the random numbers needed to implement *RCCBC*. And indeed, it is true that if the devices shared by A_0 and A_1 do actually contain sets of qubits, and each qubit is actually maximally entangled with a another qubit in another appropriately located device, and the devices do actually carry out projective CHSH measurements on these qubits, Alice could use them to generate classical random strings shared between A_c and each of the A_i in order to implement the classical protocol *RCCBC*. However, this approach would no longer be device-independent, because Alice would be making assumptions about the quantum properties of the devices. To get around this, Alice might choose to employ the purported qubits to perform device-independent randomness expansion [104], turning a small amount of trusted randomness into a long string of random numbers, but this would still require more than one bit of randomness and thus the amount of trusted randomness required for Alice to implement *RCCBC* under these circumstances would necessarily be greater than the amount required by CHSH1, CHSH2 and CHSH3.

Moreover, although it may be harder to keep entangled quantum states secure than to keep shared random classical strings secure, quantum states do offer a potentially valuable security advantage. Imagine a scenario in which Alice has a large network of agents who carry out commitments, generally relying on untrusted quantum devices. For example, Alice may not have the financial resources to construct the measurement and preparation devices needed for all of her agents, so she must rely on externally supplied devices, and cannot rule out the possibility that the manufacturer of the devices might be conspiring with Bob. However, it might be possible for Alice to get hold of just a few trusted devices - perhaps she can manufacture some personally, or buy them from a more expensive but more reliable supplier. Even simply obtaining a few devices from a a different but equally dubious source might be sufficient, as the probability that Bob will be able to suborn two different manufacturers will usually be lower than the probability

that any one manufacturer is in his pay. Alice could then use these additional devices to verify some of the purported singlets used by her agents, and hence get information about the extent of any potential security breaches. In the classical protocol, on the other hand, there is no reliable way for her to verify the privacy of the shared classical random strings she uses, so if Bob is able to breach her security, he can gain information about the strings and hence about her bit commitments without altering them or leaving any other physical trace.

8.5 Summary

The quantum relativistic bit commitment protocols that we have described in this chapter eliminate all but very general and robust assumptions about the properties of the devices employed in the protocols to prepare states and perform measurements, so participants may perform them with untrusted devices safe in the knowledge that their adversaries cannot cheat by interfering with their instruments in advance. Indeed, due to the high level of confidence we have in the impossibility of sending signals faster than light, it is likely that these protocols would be secure even against adversaries who had knowledge of a theory more powerful than quantum mechanics and/or special relativity.

Of course, the classical protocol we have described (and any *classical* relativistic bit commitment protocol), is trivially also device-independent in the sense in which we have used that term. However, the device-independent quantum protocols have subtle advantages that might potentially be valuable in certain scenarios, so the choice between classical and quantum protocols will depend which technology is most readily available and which potential security loopholes are deemed most serious in the context of implementation. That said, we underline that we have compared specific device-independent protocols against a specific classical protocol; it remains an open task to identify the optimal protocols of each type, for some natural measure of optimality, and without such a characterisation it is not possible to state definitively that there exists a quantum advantage here.

In addition to the practical advantages offered by our protocols, the results presented in this chapter are also interesting from a theoretical point of view. Al-

though this is not the first ‘device-independent’ relativistic bit commitment protocol, it is the first protocol that is device-independent in a non-trivial sense, in that the agents involved *are* required to use quantum devices, but the security does not depend on any assumptions about the quantum-mechanical properties of these devices. This means the protocol exhibits an asymmetry between completeness and security: the correlations that Alice is supposed to exhibit cannot be achieved using only classical methods, so in order to demonstrate completeness we must assume the validity of quantum mechanics, but the proofs of security against Alice and Bob depend only on the no-signalling principle together with standard cryptographic assumptions and thus these security proofs would be convincing even to an observer who trusted only the no-signalling principle and not the veridicality of quantum mechanics. Thus CHSH1, CHSH2 and CHSH3 are interesting examples of protocols that use quantum mechanics to achieve cryptographic advantages over classical protocols but also do not depend for their security on the correctness of quantum mechanics.

Chapter 9

Knowledge-Concealing Evidencing of Knowledge of a Quantum State

Bob has a black box that emits a single pure state qudit which is, from his perspective, uniformly distributed. Alice wishes to give Bob evidence that she has knowledge about the emitted state while giving him little or no information about what the state is. We prove two no-go theorems demonstrating that such zero-knowledge evidencing of knowledge is impossible in quantum relativistic protocols, extending a previous result of Horodecki et al. We then study a weaker version of the task which we refer to as ‘Knowledge-Concealing Evidencing of Knowledge of a Quantum State,’ and present a new relativistic quantum protocol for this task that significantly outperforms all existing protocols.

Based on a paper co-authored with Adrian Kent [18].

9.1 Introduction

Since relativistic quantum cryptography is still a very new field, it seems likely that there may be many valuable applications of relativistic quantum cryptography that are yet to be discovered. In this chapter we explore a new type of application, taking inspiration from the zero-knowledge proving task introduced in ref [203]. Although Horodecki et al. proved that zero-knowledge proving of knowledge of a quantum state is impossible in a purely quantum context, they did not consider

relativistic contexts, leaving open the possibility that relativistic methods might allow us to bypass the no-go theorems. In this chapter we address this possibility by studying KCEKQS tasks.

We first prove a stronger no-go theorem showing that no protocol that provides non-trivial evidence of Alice’s knowledge about a pure quantum state of finite dimension can prevent Bob from acquiring some additional knowledge about the state, even in the setting of relativistic quantum cryptography. We also prove a bound on the strength of evidence Alice can provide. This settles the issue of whether it is possible to circumvent the no-go theorem of Horodecki et al. [203] by moving to the relativistic setting.

However, this still leaves the possibility that there may exist relativistic protocols that, while not achieving perfect zero-knowledge-proving, are still a significant improvement on all existing purely quantum protocols - and indeed this turns out to be the case. After presenting our no-go theorems, we will discuss some simple protocols based on protocols previously considered by Horodecki et al. [203], and show that they are relatively weak in knowledge-concealment, or in evidencing knowledge, or both. We then propose a new relativistic quantum protocol and show that for quantum states of large dimension it is secure against restricted attacks by Alice and general attacks by Bob, in a sense we make precise below.

9.2 No-go theorems

Horodecki et al. [203] showed that no non-relativistic KCEKQS classical or quantum A-to-B protocol for an unknown qubit has $\epsilon_C = 0$, $\epsilon_S < 1$, and is zero-knowledge. (Please refer back to section 6.2 for the definitions of the parameters ϵ_S , ϵ_C and ϵ_K , which will appear throughout this chapter). We establish here a considerably more general result, applying to relativistic KCEKQS protocols for qudits with two-way classical and/or quantum communications and general parameter values. Our statements apply to protocols whose security is based only on quantum theory and special relativity, i.e. within the standard scenario for unconditionally secure relativistic quantum cryptography [13]. See the supplementary information for the proofs.

Theorem 6. *There exists no non-trivial zero-knowledge KCEKQS protocol.*

We also establish a tradeoff between completeness and soundness that bounds the amount of evidence Alice can provide. Again, this result applies to relativistic KCEKQS protocols for qudits with two-way classical and/or quantum communications and general parameter values.

Theorem 7. *For any qudit KCEKQS protocol, $\frac{\epsilon_S}{1-\epsilon_C} \geq \frac{1}{d}$.*

In particular, theorem 7 means that for small d , regardless of the value of ϵ_K , ϵ_S and ϵ_C cannot both be close to 0. This makes the case of large d particularly interesting to explore, as only for this case is there the possibility of achieving good values for both soundness and completeness.

We observe that the bound of theorem 7 is tight. For example, it is attained by a protocol in which Alice predicts to Bob the outcome of a projective measurement that includes η on the system Q_B : this has $\epsilon_S = \frac{1}{d}$ and $\epsilon_C = 0$. More generally, it is attained for a protocol in which Alice is required to predict this outcome and also predict the outcome of some independent random event where the probability for the most likely outcome is p : such a protocol would have $\epsilon_S = \frac{p}{d}$ and $\epsilon_C = 1-p$. As a result of this observation, we will say that a KCEKQS qudit protocol is *CS-optimal* if $\epsilon_S = \frac{1}{d}$, $\epsilon_C = 0$.

9.3 Classical Protocols

As noted above, Horodecki et al. [203] argue that non-trivial non-relativistic zero-knowledge protocols involving only classical communication from Alice to Bob with $\epsilon_C = 0$ are impossible for a qubit. Simply stated, any such protocol essentially comes down to having Alice predict some measurement outcome, and any prediction that holds with certainty for a pure qubit state η and is not certain for a random qubit state will always allow Bob to identify η exactly, leading to a value $\epsilon_K = 1$.

However, Horodecki et al. considered only protocols where the evidence provided by a honest Alice is certain to be accepted, i.e. protocols having $\epsilon_C = 0$; it is possible to achieve some improvement on the knowledge-concealing property

for classical protocols at the cost of accepting a non-zero value for ϵ_C . For example, Alice could choose a projective measurement which includes a randomly chosen projector P from those with $\langle \eta | P | \eta \rangle \geq 1 - \epsilon_C$ and specify the measurement to Bob, together with her prediction for the most likely outcome. Clearly, however, any protocol in which Alice simply sends Bob a classical prediction of a measurement outcome has similar issues to the protocols that Horodecki et al. analysed [203]: whenever the prediction is highly likely if Q_B is in state η and not so likely if Q_B is in a random state, Bob can obtain a significant amount of information about η simply by examining the prediction and calculating the set of states for which it is highly likely. Moreover, if Bob assumes that Alice is honest, he may choose to refrain from performing the measurement necessary to confirm her prediction and instead carry out some other measurement, thus gaining additional information.

Horodecki et al. also considered only non-relativistic protocols; it is possible to achieve some further improvements by moving to the setting of relativistic quantum cryptography, since Alice can use secure relativistic bit commitments [13, 16, 17, 165] to make commitments to her predictions, conferring two potential advantages. First, we can allow Alice to commit to more than one outcome and subsequently reveal to Bob only the outcome corresponding to the result that he has actually obtained, thus decreasing the amount of information that Bob obtains simply from the fact that she has made a certain prediction. Second, Alice need not reveal her prediction to Bob unless he first tells her the predicted outcome. The intuition is that this essentially forces Bob to carry out something close to the specified measurement if he wishes to have a reasonable chance of getting information from Alice, which prevents him from carrying out a different measurement that gives him a significant amount of additional information.

We now describe two classical protocols which use relativistic bit commitment protocols for one or both of these purposes. We describe these protocols under the assumption that Kent's procedure [13], which we described briefly in section 2.3, is used to perform the bit commitments, though in principle one could straightforwardly alter the protocols to use other bit commitment protocols; see section 9.6 for further discussion.

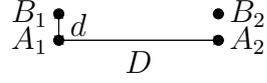


Figure 9.1: Diagram showing the arrangement in space of Alice's agents and Bob's agents for our the protocols CR1, CR2 and QBA.

9.3.1 CR1: Classical Relativistic KCEKQS protocol 1

Here and in the remaining protocols described in this chapter we take it that there is some agreed frame in which the agents remain at approximately the same position coordinates throughout the protocol; we take space and time coordinates with respect to this frame, and write $d(A_i, B_j)$ to denote the spatial distance between the stationary agents A_i and B_j in this frame. (Note that the time coordinates given for this and later relativistic protocols are merely examples of possible timings. The key requirement is that the timings should ensure that at each step, relativistic signalling constraints ensure that the relevant agent of Alice can have no information about any data supplied to the other agent by Bob's corresponding agent at the previous commitment round).

1. Alice and Bob each have two agents, A_1, A_2 and B_1, B_2 , configured so that $d = d(A_1, B_1) \approx d(A_2, B_2) \ll D = d(A_1, B_2) \approx d(A_2, B_1)$, as in the relativistic bit commitment protocol of Ref. [13].
2. Alice chooses a projective measurement $\{P_i\}$ such that $\exists x : \text{Tr}(P_x \eta) \geq 1 - \epsilon_C$. The value of x is secretly shared by both her agents.
3. At $t = 0$ Alice's agent A_1 tells Bob's agent B_1 the measurement $\{P_i\}$. Also at $t = 0$, Alice's agent A_2 and Bob's agent B_2 initiate a relativistic bit string commitment protocol committing A_2 to the binary encoding of the index x .
4. At $t = \delta \ll D$, B_1 performs the measurement $\{P_i\}$ on Q_B and reports his result to A_1 .
5. At $t = \delta'$, where $\delta' > \delta$ and $\delta' \ll D$, if B_1 reported the result P_x , A_1 unveils the commitment made by A_2 . If B_1 reported some other result or did not report any result, A_1 does not unveil the commitment.

6. If A_1 unveils a commitment that matches P_x , Bob accepts, after verifying the unveiled commitment by collecting data from his two agents. Otherwise, he rejects.

Soundness First, note that there is no exchange of quantum information in this protocol. For the purposes of security analysis, it can be treated as a relativistic bit string commitment protocol, in which Alice commits to the bit string defining x via parallel repetition of the protocol of Ref. [13]. For Alice to cheat, she must be able to unveil the possible bit strings y with success probabilities p_y for which $\sum_y p_y$ is significantly greater than 1. Conversely, proving the protocol secure against Alice requires showing that the relativistic bit string commitment protocol ensures that $\sum_y p_y \leq 1 + \epsilon(N, d)$, where N is the protocol security parameter and $\epsilon(N, d) \rightarrow 0$ as $N \rightarrow \infty$ for each finite d .

A full proof of security against Alice for this protocol thus requires analysis of the behaviour of the bit commitment protocol of ref [13] under parallel repetition. The security proof of [13] for a single classical relativistic bit commitment holds for both classical and quantum attacks; the extension of this argument to bit strings is an interesting topic in its own right, and we leave discussion of this for future work. Thus in the present discussion we simply assume without proof that, in the limit as the security parameters for the bit commitments become large, Alice can effectively commit to only one outcome of the measurement (or to some convex combination of outcomes if she makes a probabilistic commitment). Under this assumption, if Alice has no information about the state η , her optimal strategy is simply to commit to a randomly chosen outcome. Hence the value of ϵ_S is $\frac{1}{d}$, and so $\epsilon_S \rightarrow 0$ as $d \rightarrow \infty$, indicating that the protocol performs well with respect to soundness.

Completeness and Knowledge-Concealing On the other hand, the protocol does not provide a good trade-off between completeness and knowledge concealment. Assuming Alice does indeed know the state η , she produces a valid proof with probability at least $(1 - \epsilon_C)$. In this case Bob learns the value of a projection P_x with $\text{Tr}(P_x \eta) \geq 1 - \epsilon_C$.

Hence

$$\epsilon_K \geq (1 - \epsilon_C)^2. \quad (9.1)$$

This is a very poor result; in particular $\epsilon_C \approx 0$ implies $\epsilon_K \approx 1$. Since $\epsilon_M = \frac{2}{d+1}$ [225], for $\epsilon_C \approx 0$ we have $\epsilon_K - \epsilon_M \gtrsim \frac{d-1}{d+1} \geq \frac{1}{3}$, and for d large we have $\epsilon_K - \epsilon_M \approx 1$.

One way to improve on the bound 9.1), at the price of increasing ϵ_S , is to allow Alice to commit to more than one outcome of the projective measurement. We now consider a protocol of this type.

9.3.2 CR2: Classical Relativistic KCEKQS protocol 2

1. Alice and Bob each have two agents, A_1, A_2 and B_1, B_2 , configured so that $d = d(A_1, B_1) \approx d(A_2, B_2) \ll D = d(A_1, B_2) \approx d(A_2, B_1)$, as in the relativistic bit commitment protocol of Ref. [13].
2. Alice chooses a projective measurement $\{P_i\}$ for which there exists a set S consisting of q measurement elements such that $\text{Tr}((\sum_{i \in S} P_i)|\eta\rangle\langle\eta|) = 1 - \epsilon_C$ for some agreed value of ϵ_C . The indices $x \in S$ are secretly shared by both her agents.
3. At $t = 0$, Alice's agent A_1 tells Bob's agent B_1 the measurement $\{P_i\}$. Also at $t = 0$, Alice's agent A_2 and Bob's agent B_2 initiate relativistic bit string commitment protocols committing Alice to the binary encoding of each index $x \in S$.
4. At $t = \delta \ll D$, B_1 performs the measurement $\{P_i\}$ on Q_B and reports his result to A_1 .
5. At $t = \delta'$, where $\delta' > \delta$ and $\delta' \ll D$, if B_1 reported a result P_x with $x \in S$, A_1 unveils the commitment to x made by A_2 , but does not unveil any other commitments. If B_1 reported a result P_x with $x \notin S$, or did not report any result, A_1 does not unveil any commitments.

6. If A_1 unveils a commitment that matches P_x , Bob accepts, after verifying the unveiled commitment by collecting data from his two agents. Otherwise, he rejects.

Soundness In the supplementary information we show that, conditional on our previously stated assumptions about the security of relativistic bit commitment protocols, we have $\epsilon_S = \frac{q}{d}$. This is an integer multiple of the optimal value, but for small q and large d we may still obtain ϵ_S close to zero, so we can still achieve a good level of soundness in the limit of large d .

Completeness and Knowledge-Concealing In the supplementary information we show that, conditional on our previously stated assumptions about the security of relativistic bit commitment protocols, we have:

$$\epsilon_K \geq q \left(\frac{1 - \epsilon_C}{q} \right)^2$$

Thus for near CS-optimality, with $q \ll \frac{d}{2}$ and $\epsilon_C \approx 0$, we have $\epsilon_K \gg \epsilon_M$, so Bob is able to gain a large amount of information under these circumstances.¹

In summary, although it is possible to achieve small improvements on the classical protocols described in ref [203] by relaxing the values of the security parameters and/or moving to the relativistic context, all the generalisations considered above necessarily have either ϵ_K or ϵ_C large, and thus it seems unlikely that any of these generalisations would be sufficiently secure for any sensitive practical applications.

¹A variation on this protocol is for Alice to randomly choose a projective measurement $\{P_i\}$ from among those for which there exists a set S consisting of q measurement elements such that $\text{Tr}((\sum_{i \in S} P_i)\eta) \geq 1 - \epsilon$. She could, for example, use the uniform measure on the complete projective decompositions satisfying this criterion. This variation has $\epsilon_C < \epsilon$, since Alice's average success probability, when she knows η , is greater than $1 - \epsilon$. Similarly, it has $\epsilon_K > (1 - \epsilon)^2 q^{-1} > (1 - \epsilon_C)^2 q^{-1}$.

9.4 Quantum A-to-B Protocols

Horodecki et al. [203] also consider a protocol where Alice gives Bob a copy of η ; as they note, such protocols can achieve $\epsilon_C = 0$ and $\epsilon_K < 1$. Indeed it is possible to achieve $\epsilon_K \ll 1$ for states of large dimension, and in this sense the quantum protocol outperforms the classical A-to-B protocols just discussed. However, it is also important to consider the tradeoffs between ϵ_K , ϵ_C and ϵ_S , and we will see below that these tradeoffs are not favourable for this type of protocol. It is also worth highlighting that this protocol gives no evidence favouring the hypothesis that Alice has classical information about η over the hypothesis that Alice simply possesses relevant quantum information about η - for example, Alice can ensure $p(1) = 1$ even if she only has a black box that will make only a single copy of η and has no other classical or quantum information about η .

9.4.1 QAB: Quantum A-to-B KCEKQS protocol

To make our discussion more concrete, we now extend the discussion of Ref. [203] by considering a generalisation of the protocol suggested by Horodecki et al. in which Alice gives Bob N copies of η and Bob tests these copies using a projection on the symmetric subspace:

1. Alice prepares N systems $\{S_i\}$ in the state η and gives them to Bob.
2. Bob performs a measurement $\{\Pi_S, \mathbb{I} - \Pi_S\}$ where Π_S is the projector onto the symmetric subspace of the joint state space of the $\{S_i\}$ and Q_B .²
3. If the result is Π_S , Bob accepts; otherwise he rejects.

Completeness If Alice knows η and follows the protocol correctly, Bob will always accept her evidence, so this protocol achieves $\epsilon_C = 0$.

²A motivation for this choice is that, given a system in the state $\psi^{\otimes n}$ and another system in the state $\phi^{\otimes m}$, for some integers m, n , the measurement $\{\Pi_S, \mathbb{I} - \Pi_S\}$ (i) always gives outcome 1 if $\psi = \phi$ (ii) maximises the probability of outcome 0 if $\psi \neq \phi$, among measurements satisfying (i) [232, 233].

Soundness In the supplementary information, we show that $\epsilon_S = \frac{1}{N+1} + \frac{N}{d(N+1)}$.

Thus for $N = 1$ we have $\epsilon_S = \frac{1}{2} + \frac{1}{2d}$, which is a very poor result for soundness - even if Alice has no information whatsoever about the state she has at least probability $\frac{1}{2}$ of producing evidence which is accepted by Bob.

However, for N large, $\epsilon_S \rightarrow \frac{1}{d}$, so the protocol tends to CS-optimality in this limit.

Knowledge-Concealing We also show in the supplementary information that $\epsilon_K = \frac{N+2}{N+d+1}$, which tends to 1 for large N . Thus $\epsilon_K > \frac{1}{d\epsilon_S}$ and $\epsilon_K - \epsilon_M > \frac{d-1}{d+1} \frac{N}{N+2} \frac{1}{d\epsilon_S}$ for all d, N , which are relatively poor tradeoffs. In particular, if the relevant parameters are chosen to ensure near CS-optimality, with $\epsilon_S \approx \frac{1}{d}$, then Bob's knowledge gain is necessarily significant, with $\epsilon_K - \epsilon_M \approx \frac{d-1}{d+1}$. Note that Bob can both follow this protocol honestly in order to gain evidence about Alice's knowledge of the quantum state, and then also afterwards attempt to estimate η from the $N + 1$ states in his possession since the states of these copies of η are not changed in the course of a protocol in which both parties are honest. Thus, if Alice honestly follows the protocol, Bob can obtain evidence of her knowledge of η and still attain the value of $\epsilon_K = \frac{N+2}{N+1+d}$.

In summary, QAB forces us to choose between a high value of ϵ_K and a high value of ϵ_S , and thus once again, it seems unlikely that QAB would be sufficiently secure for any sensitive practical applications.

9.5 Quantum B-to-A Protocols

Ref [203] considered only classical and quantum A-to-B protocols, but here we describe a new type of protocol which uses quantum communication from Bob to Alice, as well as two-way classical communication. The basic idea is that Bob prepares a set of N systems in random states, then puts the original system Q_B in amongst these other systems and gives the whole collection to Alice, who is required to guess which system is Q_B . If Alice does know the state η , her best guessing strategy is to perform a measurement projecting on η , but this measurement is

likely to return a positive answer on a proportion $\frac{1}{d}$ of the systems, and so Alice will not usually be able to identify Q_B uniquely. Thus to ensure that honest Alice has a reasonable chance of having her evidence accepted, she must be allowed to make approximately $\frac{N+1}{d}$ guesses. However, if she reveals all these guesses to Bob he may gain a significant amount of information from them, and this is where the protocol becomes relativistic - we require Alice to perform relativistic bit commitments to the indices associated with her guesses and subsequently to unveil only the guess which is confirmed by Bob to be correct.

9.5.1 QBA: Quantum B-to-A KCEKQS protocol

As in the section on classical protocols, we will present an implementation of the protocol using Kent's bit commitment protocol [13]; as before, in principle one could straightforwardly alter the protocols to use other bit commitment protocols.

1. Alice and Bob each have two agents, A_1, A_2 and B_1, B_2 , configured so that $d = d(A_1, B_1) \approx d(A_2, B_2) \ll D = d(A_1, B_2) \approx d(A_2, B_1)$, as in the relativistic bit commitment protocol of ref [13] (see fig 9.1).
2. Bob gives Q_B to B_1 , who prepares N additional quantum systems $\{S_i\}$ in states chosen uniformly at random.
3. B_1 randomly permutes the systems $\{S_i\}$ and the system Q_B , assigns them all indices from 1 to $N + 1$, and then gives all $N + 1$ systems, labelled by their indices, to A_1 .
4. A_1 carries out the projective measurement $\{\eta, \mathbb{I} - \eta\}$ on each of the $N + 1$ systems that B_1 gave her. Write C' for the list of indices for which she obtains outcome η ; let $|C'| = q'$. If $q' \leq q$, she forms a list $C = C' \cup D$, where D is a list of $(q - q')$ copies of the dummy index 0.³ If $q' > q$, she picks a random size q sublist C of C' .
5. A_1 randomly permutes C . At $t = 0$, B_1 and A_1 initiate q relativistic bit string commitments [13] committing A_1 to each of the indices in the per-

³The use of the dummy index prevents cheating strategies in which Bob uses the number of Alice's commitments, made at the next step, to extract additional information about the state.

muted list. Each bit string commitment is set up so that Alice can commit to any index in $\{1, 2, \dots, N + 1\}$. This commitment is sustained by B_2 and A_2 at time $t = \delta$, where $0 < \delta \ll D$. These commitments involve a further security parameter N' .

6. At $t = \delta'$, where $0 < \delta' \ll D$ (and for definiteness we may take $\delta' > \delta$) B_1 tells A_1 the index $x \in \{1, \dots, N + 1\}$ that he assigned to Q_B .
7. If $x \in C$, A_1 unveils her commitment to that index (which she initiated and A_2 sustained). Otherwise she announces failure (or aborts, if the protocol includes an abort option) and Bob rejects.⁴
8. If Alice's unveiled commitment is indeed x , Bob accepts, once he is able to verify the commitment by comparing data from his agents. Otherwise he rejects.

Knowledge-Concealing In the supplementary information we show that $\epsilon_K \leq \frac{4}{d+1}$, so $\epsilon_K \rightarrow 0$ and $\epsilon_K - \epsilon_M \rightarrow 0$ for large d . Thus for large-dimensional quantum states Bob's knowledge gain in QBA is close to zero.

Completeness and Soundness In the supplementary information we show that if it is assumed that the relativistic bit commitments employed remain secure under parallel composition, so Alice is restricted to strategies in which she commits to q classical values chosen from $\{0, \dots, N + 1\}$ and unveils one of these committed values, then for an appropriate choice of the parameter q , QBA asymptotically tends to CS-optimality in the limit where the security parameter N is large. We conjecture that this remains true without the assumption, but a full security analysis requires a complete analysis of general quantum operations Alice could carry out to produce unveiling data, which we leave for future work.

In summary, when applied to quantum states of large dimension QBA performs well in terms of knowledge-concealment, and under our assumptions about

⁴In the ideal error-free case, if Alice knows η precisely and both parties honestly perform the protocol, failure is possible if and only if $q' > q$.

the composability of bit commitments it is also close to CS-optimality in this limit. Thus insofar as these assumptions are valid, it seems that QBA performs significantly better than existing classical and Quantum Bob-to-Alice protocols and therefore may be secure enough to be used in practical applications.

Abort option QBA could be altered to achieve $\epsilon_C = 0$ by allowing Alice to abort the protocol whenever she obtains a positive outcome more than q times at step 5.⁵ Under these conditions, an honest Alice with perfect knowledge will abort whenever the binomial random variable $X_N > q$, and thus for $q = \frac{N+1}{d}$ and large N , this implies an abort probability of roughly $\frac{1}{2}$. Taking $q = \frac{N}{d} + \epsilon N$, with $\epsilon > 0$, we see from equation (A.4) that the abort probability can be bounded by $\exp(-2\epsilon^2 N)$. This gives a protocol with $\epsilon_C = 0$, $\epsilon_S = \frac{1}{d} + \epsilon$ and with abort probability that tends to zero for large N .

These parameters may represent a reasonable tradeoff in some circumstances. However, introducing an abort option does not eliminate the possibility of unjustified mistrust. Without an abort option, it is possible that Alice and Bob may both honestly follow the protocol, and that Alice may know η precisely and thus correctly identify Q_B as a candidate, but that she may be unable to persuade Bob of this because she had more than q candidates and her random choice of a size q subset did not include the index of Q_B . Introducing an abort option removes the possibility of honest Alice being falsely suspected of cheating by honest Bob *for this specific reason*. However, an honest Bob may now unfairly suspect an honest Alice of cheating if she honestly aborts, since a dishonest Alice might abort because she had no information about η , carried out no measurements, made random or invalid commitments, and thus effectively used the protocol to steal Bob's copy of η ; Bob has no way to tell whether or not the protocol has been honestly aborted.

Indeed, no protocol with nonzero probability of either failure or abort can

⁵One could also include an abort option for the case where Alice does not obtain a positive outcome for any one of these measurements. However, there is a significant distinction between these cases. If Alice obtains no positive outcome, then in the error-free model she knows for certain that Bob is trying to cheat, whereas in the case where she obtains more than q positive outcomes, it is possible that Bob was honest and a statistically unlikely outcome was obtained. Thus in order to distinguish between an accusation of cheating and an instance of protocol failure for which no one is to be blamed, it might be preferable to have two types of abort in situations where this possibility is relevant.

guarantee that an honest Bob accepts the evidence provided by an honest Alice. This is because KCEKQS is a one-shot procedure, i.e. Bob has only one copy of η , so if he gives that copy to Alice and the protocol fails or is aborted, he has no further opportunity to learn anything about η or about Alice's knowledge of it. Of course, the parties might be able to repeat the protocol using a new state, but even if this next protocol succeeds, Alice will have proved only that she knows the new state, which gives Bob no direct evidence about whether she knew the previous state. Thus in most scenarios, it seems to us that the abort option version of the protocol offers no clear advantage.

9.6 Further Security Issues

Here we discuss several potential security issues and possible weaknesses for KCEKQS protocols that are not covered by our earlier security definitions.

Bit String Commitment We have reinforced throughout this chapter that our security proofs for CR1, CR2 and QBA depend on the assumption that the relativistic bit commitment protocols employed are secure under parallel repetition and hence can be used to perform secure bit string commitment. However, the security of Kent's procedure [13] under parallel repetition has not yet been proven. If in fact it turns out that this particular bit commitment protocol is not secure under parallel repetition, in principle one might replace it with some other relativistic bit commitment procedure that *is* secure under parallel repetition, but it is important to keep in mind that not every bit commitment geometry is suitable for use in CR1, CR2 and QBA. For example, the protocols we described in chapters 7 and 8 cannot straightforwardly be employed in this context because there would be nothing to stop Bob from giving the agents A_0 and A_1 two different indices and thus getting them to reveal two different pieces of information, so he would potentially obtain twice as much information as we have assumed in our security arguments.

Entanglement We have thus far made the assumption that the state η is pure, and therefore we have not considered the possibility that Alice might have access

to one or more systems that are entangled with Q_B . However, another natural scenario in which one might wish to employ a KCEKQS protocol is where Q_B is maximally entangled with a system Q_A in Alice's possession and Alice knows the joint state of (Q_A, Q_B) . Let us refer to this scenario as Sc_1 .

In this scenario, Alice can always proceed by performing a projective measurement on Q_A before the start of the protocol, so the system Q_B is subsequently in a pure state η that is known exactly to Alice. This recreates our original scenario, which we henceforth refer to as Sc_0 . It follows that if Alice can succeed in some KCEKQS protocol with probability $(1 - \epsilon_C)$ in the original scenario Sc_0 , she can also succeed with probability at least $(1 - \epsilon_C)$ in scenario Sc_1 . Conversely, while there are clearly protocols that give Bob evidence favouring Sc_1 over Sc_0 , no KCEKQS protocol can give Bob evidence favouring Sc_0 over Sc_1 . That is, the evidence provided by any KCEKQS protocol will be equally compatible with the hypothesis that Q_B was always in a pure state that Alice happened to have classical knowledge about, and the hypothesis that Q_B was originally in a maximally entangled state with some system in Alice's possession and Alice had classical knowledge about the whole maximally entangled state.

What precisely does Alice give evidence of? We have defined security for KCEKQS in terms of two extremes: the value of ϵ_S gives the probability that Alice's proof is accepted if she knows nothing at all about the state, and the value of $1 - \epsilon_C$ gives the probability that Alice's proof is accepted if she knows the state exactly. But of course there are intermediate possibilities. For example, Alice could have some classical knowledge about the quantum state without knowing it exactly, or she could have quantum information correlated with the state, such as one or more copies of it, or she could have beliefs about the state encoded in a probability distribution.

In each case, if Alice's information or beliefs allow her to produce a guess η' at the classical description of η such that $\text{Tr}(\eta\eta') \approx 1$, she can produce what we have called 'evidence of knowledge,' and in all of the KCEKQS protocols described above, Bob will accept this 'evidence of knowledge' with probability close to $1 - \epsilon_C$, even though Alice does not actually have perfect knowledge of η . Indeed, there may be a high probability that Bob accepts Alice's evidence

even when she actually has *false* beliefs about its classical value - for example, if she believes Q_B is in the state η' and it is actually in the state $\eta \neq \eta'$, where $\text{Tr}(\eta\eta') \approx 1$. As this illustrates, the notion of knowledge of a quantum state needs careful analysis. One might be inclined to frame a definition so that if Alice believes Q_B is in the pure state η' , and it is actually in the state $\eta \neq \eta'$, then she has no knowledge about the state, no matter how close $\text{Tr}(\eta\eta')$ is to 1. But this would mean accepting that someone who has no knowledge about a state can nonetheless appear to give strong evidence of knowledge of it! Alternatively, one might frame a definition so that, if Alice has probabilistic beliefs about the state η encapsulated in the density matrix ρ , then $\text{Tr}(\rho\eta)$ is a measure of her knowledge. If so, one has to accept that false (not merely uncertain) beliefs are consistent with a high degree of knowledge, which is in tension with the popular view that only true beliefs can count as knowledge [234]. Either way the consequences are somewhat counterintuitive, and this is perhaps best regarded a sign that the *type* of knowledge one has when one knows a quantum state is not quite the same as the types of classical knowledge that we are more familiar with.

We will not here attempt to provide a general definition of knowledge of a quantum state, and instead will simply observe that our security definitions do establish that the protocols give evidence of knowledge of η in an interesting, if restricted, sense: namely, they strengthen the evidence for the hypothesis that Alice knows the precise classical value of η compared to the hypothesis that she has no classical or quantum information correlated with η . One might frame additional security definitions that allow more to be established. For example, one could require *strong non-triviality*: if Alice does not have and is not able to obtain a precise classical description of η , then the probability that her proof is accepted is strictly less than $(1 - \epsilon_C)$. This requirement holds for QBA, but not for QAB, in which Alice can ensure acceptance probability $(1 - \epsilon_C)$ even if she only has some way of producing precisely N copies of η (and no more), and no other classical or quantum information correlated with η .

Composability and Accumulation of Information Classically, zero-knowledge proofs are often used as sub-protocols in the construction of more complex protocols, such as electronic voting schemes and digital signature

schemes. One might hope that our knowledge-concealing protocols could be used as building blocks for quantum protocols of this type. However, it is known that relativistic bit commitment protocols are not secure under general compositions [235], and since our knowledge-concealing protocols depend on the use of bit commitment, they are presumably also insecure under general compositions. Furthermore, since the amount of information obtained by Bob in a KCEKQS protocol will never be *exactly* zero for any finite-dimensional quantum state, one would have to consider carefully whether Bob's information gains could accumulate in such a way as to undermine the overall security. This does not mean that all compositions using such protocols are necessarily insecure, but it does mean that the security of such compositions would have to be assessed on a case-by-case basis, and possible attacks on the compositions would have to be addressed separately from attacks on single-run protocols.

Alternative Measures of Bob's information gain We have chosen ϵ_K as an appropriate measure of Bob's knowledge gain for a single shot protocol because it quantifies Bob's ability to predict the outcome of measurements on a system in the state η and/or to prepare a state that successfully simulates the state η . These operationally defined forms of knowledge about a state are relevant in most cryptographic contexts, and indeed are the only relevant factors in many scenarios.

However, it is important to reinforce that density matrices do not completely characterise an agent's knowledge of how a system was prepared, which means that our chosen knowledge-concealing criterion does not capture all the information Bob might learn in the course of any conceivable protocol. As a hypothetical illustration, imagine that in a KCEKQS protocol for the state of a qubit, Bob somehow learns in the course of the protocol that the original state was definitely either the state $|+\rangle$ or the state $|-\rangle$. Clearly there is a sense in which Bob has gained significant information in the course of this protocol - he has narrowed the set of possible states down from infinity to two! However, the density matrix ρ describing his new state of knowledge and the density matrix describing his knowledge before the protocol are both the maximally mixed state, so our knowledge-concealment criterion would suggest that he gains *no* information in the course of such a protocol.

This is a somewhat contrived example, but it does make the point that there is no unique measure of information about a quantum state that characterises every property that is relevant in every possible scenario, and therefore we cannot guarantee that our knowledge-concealing protocols will always hide from Bob the particular type of information that is considered valuable in a given context of application. Even assuming that Bob's aim is to produce a guess η' at η that optimizes some cost function, and assuming also that the cost function depends only and monotonically on $\text{Tr}(\eta\eta')$, there are still infinitely many cost functions that may be considered. As a simple example, Bob might be given a fixed reward if and only if $\text{Tr}(\eta\eta') > 1 - \epsilon$ for some small $\epsilon > 0$. In this case, a sensible parameter might be ϵ'_K , defined as the probability of this condition holding after the protocol if Bob follows an optimal strategy, and a sensible comparison would be to ϵ'_M , the maximum probability that the condition holds if Bob does not participate in the protocol at all and instead simply carries out a strategy that involves operations only on the unknown state.

We would therefore recommend that before deciding to employ the protocol QBA (or any other KCEKQS protocol) in some context, one should identify exactly what type of information it is important to hide in this particular context and establish whether that information is well-captured by the parameter ϵ_K . Even if this is not the case, the protocol QBA may nonetheless be an appropriate choice - our informal investigations suggest that it still conceals a significant amount of 'knowledge' even if knowledge is measured using some other reasonable measures of information content, such as the Kullback-Leibler divergence between two appropriately chosen probability distributions - but additional security analysis will be required to make certain of this.

What does Alice learn? We have considered in detail the amount of information that Bob gains about the state η in the course of our protocols, but we have not thus far considered how much *Alice* could potentially learn about η if she does not in fact know η . In some applications it might not matter if Alice gains information about η , but if we are using a KCEKQS protocol precisely because knowledge of the state η has value in some context (such as quantum money [236] or quantum voting [199]), then it may be important to limit the amount of information that

might potentially be gained by a dishonest Alice.

In CR1 and CR2, assuming that Bob performs the protocol correctly and honestly, Alice gains some information about the state, since Bob must tell Alice the result of a measurement that he has performed on the system Q_B . But clearly since Alice obtains information about the state only via Bob's classical communication, she learns no more than Bob does in the course of such protocols, if she begins with no knowledge about the state.⁶

The protocol QAB performs best out of all the protocols we have considered with regard to limiting Alice's knowledge gain: Bob need not tell Alice whether or not his measurement obtained a positive result, so Alice learns nothing whatsoever from the protocol. On the other hand, the protocol QBA is significantly vulnerable to this potential issue. If Bob performs QBA honestly he learns nothing about η from the protocol, whereas a dishonest Alice can store all the quantum states sent to her and make commitments to random indices, or fail to make valid commitments at all, then perform a measurement on Q_B after Bob has told her the index of Q_B . Effectively, she can steal η from Bob and obtain as much information about it as he could have obtained if he had not participated in the protocol. Clearly this strategy has a low probability of successfully persuading Bob that Alice has honestly followed the protocol, particularly for states of large dimension, and hence this behaviour is likely to be detected if Alice tries it repeatedly over some number of protocols, but nonetheless, on a single run of the protocol Alice is able to dishonestly gain some information about η , whereas an honest Bob sacrifices his copy of η and gains no information at all. For large d , the information a dishonest Alice can gain is small, but in contexts where even limited information about η is highly valuable, this might be an undesirable weakness.

A related issue is that both Alice and Bob may want to limit the information that could be gained by an eavesdropper who intercepts their classical or quantum communications during the protocol. In principle, eavesdropping can be completely prevented by using secure classical and quantum channels, but we may sometimes need to use KCEKQS protocols in scenarios where secure channels are not available. Moreover, although secure classical channels can be ensured

⁶Alice may however learn more than Bob does if she begins with some information about the state.

by using one time pads [237], secure quantum channels require distributing and storing perfect entangled states, so at present quantum eavesdropping remains a practical concern.

None of the protocols we have studied are completely immune to eavesdroppers. In CR1 and CR2, an eavesdropper who can listen in on the classical communication channels will learn all the same information as Alice and Bob themselves. In QAB, an eavesdropper who can intercept the copies of η sent from Alice to Bob can gain information about η by state estimation. In QBA an eavesdropper who can intercept both the quantum and classical information sent by Bob can gain the same amount of information as Alice could potentially gain, although an eavesdropper who can intercept only the quantum information learns significantly less and an eavesdropper who can intercept only the classical information learns nothing at all.

What if Bob has additional information? We have calculated bounds on the information that Bob gains during each protocol assuming he starts with strictly zero knowledge about the state of Q_B and only ever has a single copy of the state of Q_B .

In other possible scenarios, Bob might also have some limited classical information about η , or have additional correlated quantum information (for example further copies of η), or both. Our bounds do not necessarily apply in such scenarios. For example, suppose that Bob has Mq copies of η , for some large M , and that he and Alice perform CR2. Bob can apply the projective measurement specified by Alice on every copy in his possession. If Alice is honest, and $\epsilon_C \ll (Mq)^{-1}$. Bob is then likely to obtain positive outcomes for r elements of the measurement basis, where $r \leq q$, which allows him to identify a subspace V of dimension r such that $\text{Tr}(P_V \eta) \approx 1$.

9.7 Summary

We have proven two no-go theorems demonstrating that even in the relativistic setting there is no perfect KCEKQS protocol for quantum states of finite dimension, and quantifying the inevitable tradeoff between soundness and complete-

ness. We have also described a new protocol involving quantum Bob-to-Alice communications and relativistic signalling constraints, which, subject to certain assumptions about the composability of bit commitment protocols, represents a significant improvement on existing protocols. Although QBA is not zero knowledge for finite d , we showed that it reveals little extra information to Bob for large d . We also showed that based on our composability assumptions, the protocol is asymptotically CS-optimal, i.e. offers essentially optimal security against Alice; we conjecture that this will turn out to be true even without the assumptions, once a suitable analysis of bit commitment under parallel repetition is available. Thus we anticipate that QBA may be a valuable quantum cryptographic primitive in contexts where marginal revelations of information to Bob are acceptable.

Some intriguing possibilities for future work are also suggested by our results. First, we note also that according to the standard definition set out in section 6.2, the protocol QBA is a perfect zero-knowledge protocol, because if both Alice and Bob are honest and a valid proof is produced, then the transcript simply consists of Bob telling Alice the value of an index and Alice unveiling a commitment to that index, and the probability distribution p_x then simply reduces to a probability distribution over Bob's choice of index, exactly the same as the distribution p_y . However, we reinforce that this does not provide any guarantee that the protocol will not reveal information to Bob, because, as we have shown, Bob may gain information by various dishonest strategies, such as telling Alice the wrong index; furthermore, even an honest Bob may gain information in the case where Alice fails to produce a correct proof due to obtaining more than q positive measurement results and being unlucky in her random choice of commitments. Although it is known that the classical context, any protocol with the statistical zero-knowledge property against honest verifiers can be transformed into one with the statistical zero-knowledge property against dishonest verifiers [238]; the result does not immediately translate to our context, where relativistic constraints are used and quantum as well as classical information may be exchanged; nonetheless it seems natural to ask if some quantum and/or relativistic generalization is possible, and if such a thing could be found this might suggest ways to strengthen our proposed protocol. Second, Sahai and Vadan have shown that a certain problem, known

as *statistical difference*, is complete⁷ for the class of all problems which possess statistical zero-knowledge proofs (SZP), and their work on the related circuit polarization problem has been extended by Holenstein and Renner [239]. One could define a generalization of the class SZP to cover quantum problems and quantum and/or relativistic proofs, and if it were possible to find a complete problem for this more general class, it would be interesting to see how closely this problem were related to statistical difference and circuit polarization.

We also pause to observe that our results have some relevance from a theoretical perspective in the context of the long-standing and contentious debate about whether or not the quantum state is an element of reality. KCEKQS is of interest in this regard because in philosophy, realism about theoretical entities is usually said to include a commitment to the idea that statements about theoretical entities constitute knowledge about the world above and beyond descriptions of observables [240, 241], and one way of giving content to this sort of claim would be via the medium of zero-knowledge proofs. For example, in one famous toy example of zero-knowledge proving [242] we have a cave with a secret passageway whose location is known to Alice, and Alice has been set the task of proving to Bob that she knows where the passageway is without revealing its location to him. In this scenario, Alice does indeed know a description of the way out of the cave, and clearly knowing a sufficiently detailed description of the way out of the cave is equivalent to knowing the way out of the cave simpliciter. But nonetheless Alice's knowledge is not just about the description - it is knowledge about the *actual physical* cave, and hence by demonstrating her ability to pass through the passage, she can provide evidence of her knowledge which Bob can verify even if he does not know the description himself. Conversely, if what Alice knew were only a description which did not pertain to any actual physical cave, then giving evidence of her knowledge would essentially come down to stating all or part of the description, and this evidence would be verifiable only by someone else who knew the same description, making it impossible to perform effective knowledge-concealing evidencing of knowledge.

In the case of the quantum state, Gleason's theorem [243] tells us that know-

⁷A problem p is said to be 'complete' for some complexity class if it belongs to the class and every problem in the class can be algorithmically transformed to it.

ing a sufficiently detailed description of the experimental statistics predicted by some quantum state is equivalent to knowing the quantum state simpliciter, and therefore it is reasonable to ask whether knowledge of a quantum state is really just knowledge about the associated description of observables, or whether it constitutes knowledge about some actual physical entity over and above this description. One approach to answering this question would be to observe from our no-go theorems that zero-knowledge proving of knowledge of a quantum state is impossible, and even knowledge-concealing evidencing of knowledge cannot be done very effectively for quantum states of small dimension. That is, for quantum states of small dimension there is not much that Alice can do to provide evidence of her knowledge of the quantum state that is verifiable by someone who does not know the corresponding description of observables, and this suggests that Alice's knowledge of the quantum state is in some important way disanalogous to Alice's knowledge of the actual physical cave in the earlier case. Hence these no-go theorems might be taken to mitigate against the view that knowledge about quantum states constitutes knowledge about the world above and beyond descriptions of observables, at least for states of small dimension.

A range of other facts about quantum theory may be adduced in support of this point; for example, the counterintuitive consequences we encountered in attempting to find a sensible way to define Alice's knowledge of the quantum state η might similarly be taken to suggest that knowledge of quantum states is qualitatively different from familiar types of knowledge about physical things in the world. More generally, we saw in section 1.2.5 that quantum contextuality is usually interpreted as indicating that the results of quantum measurements do not give us knowledge about pre-existing properties of the world, and thus, insofar as the quantum state of a system may be regarded as a codification of everything that can be inferred about the results of future measurements on that system, one might be tempted to conclude that knowledge of the quantum state does not represent knowledge about occurrent properties of the world. Likewise, the fact that the quantum conditional entropy can be negative may also be interpreted as a sign that the knowledge encoded in the quantum state is not knowledge about underlying occurrent properties, because if it were, then we would have to accept that it is possible to know less about the properties of a part than about the properties

of a whole, which in classical logic represents a contradiction. Taken together, these observations might be regarded as an argument against realism about quantum states of small dimension, at least in the specific epistemic sense in which that term is sometimes used in the philosophy of science. Of course, the argument is much too informal to support any grand ontological conclusions, but it might at least help to explain why so many people have had the intuition that the quantum state may not be real.

On the other hand, the quantum *B-to-A* protocol that we set out in this chapter demonstrates that it does become possible to produce something close to a zero-knowledge proof of knowledge of a quantum state in the limit as the dimension of the quantum state concerned becomes very large. Our proposed epistemic criterion for reality, then, would suggest that it becomes more natural to be a realist about quantum states as their dimension becomes large. This might help explain why in classical physics it was not really controversial to suppose that systems have underlying ontological states that are responsible for the observed measurement statistics: classical systems are continuous and thus infinite-dimensional, and in the infinite limit it is indeed possible to provide evidence of knowledge of a (quantum) state that can be verified by someone who does not know a description of the state, thus meeting our epistemic criterion for reality. Such an effect might well foster the illusion that there exists an underlying state which is an entity in and of itself rather than merely a codification of predictions about future behaviour, and thus some of our classical intuitions about fundamental ontology could be thought of as a consequence of the infinite-dimensionality of classical systems.

Chapter 10

Concluding Remarks

The first two parts of this thesis have clear practical implications for quantum computing and relativistic quantum cryptography.

As we pointed out in chapter 4, many known paradigms for quantum computing allow us to recast a computation as a summoning task where the computer must produce stored quantum states at various spacetime points that depend on the results of previous rounds of computation. For example, in a measurement-based quantum computation we take an entangled multiqubit state and apply a sequence of measurements in different bases, where the basis choices may depend on the results of measurements at earlier times; since the gates for different measurements may be in different spatial locations, we can regard the result of the earlier measurement as a ‘call’ for some part of the state to be moved to the spatial location of a certain gate, which can be regarded as a ‘response point.’ Thus it is valuable to have an explicit formulation of the necessary and sufficient conditions characterising the set of feasible summoning tasks, and particularly helpful to have access to a generic prescription for a protocol that will succeed with certainty for any task belonging to this set. Previous work by Hayden and May fulfilled these desiderata for the special case in which it is guaranteed that exactly one call will be made, but since the answers to computations are not generally known in advance, it may be difficult to provide such a guarantee in realistic computing scenarios, and therefore the generalisation we provide in the first part of this thesis is likely to be of use as the field of distributed quantum computing develops. Likewise,

many quantum computations require the distribution not only of individual pure states but of entangled states, and thus our generalization of summoning tasks to the case of entangled pairs is also likely to be of use to the field.

In the second part of the thesis we set out new relativistic bit commitment protocols that have some advantages over existing protocols, making it possible to perform secure bit commitment either without using randomness as a resource or without relying on assumptions about the properties of preparation and measurement devices. Since bit commitment has a large variety of applications, it is helpful to have alternate possibilities that may be preferable to the existing protocols in certain contexts of applications. This is particularly relevant in view of the fact that no bit commitment protocol is guaranteed to be secure under all possible compositions [235] - due to the security advantages of our new protocols, it is likely that each of them may be secure under certain desirable compositions where older protocols are not secure. Finally, we also developed a relativistic quantum protocol for knowledge-concealing evidencing of knowledge of a quantum state, a task which we anticipate could play a similar role in future quantum cryptographic protocols as zero-knowledge proving plays in classical cryptography, thus potentially leading the way to new applications for relativistic quantum cryptography.

In addition to these practical applications, the results of Part I and II offer an interesting theoretical perspective on the long-standing argument about the reality of the quantum state. The current convention in quantum foundations is to reduce this argument to a choice between ' ψ -ontic' and ' ψ -epistemic' views, where a model is defined formally as ψ -ontic if and only if it always associates distinct quantum states with non-overlapping probability distributions over ontic states of the world; so for example much discussion has been generated by the recent PBR theorem [208] which shows that no theory in which the quantum state is not real in the ψ -ontic sense can reproduce all the predictions of quantum theory. And yet one might worry that the formal definition of ψ -ontic does not capture the full intuitive content of the statement that the quantum state is 'real.' As an example of the possible connotations of this terminology, consider the distinction made in metaphysics between 'concrete' and 'abstract' objects, which is commonly cashed out by saying that concrete objects are those that can in principle be assigned a spa-

tiotemporal location and/or are capable of persisting through time [244]. We saw in the first section of this thesis that Hayden and May's approach to characterizing the spatiotemporal history of a quantum state seems questionable in view of our results on multiple-call summoning tasks, and we conjectured that it may well be the case that no such approach holds up to careful analysis, suggesting that quantum states perhaps do not fall on the 'concrete' side of this particular distinction. The difficulty of saying anything meaningful concerning the spatiotemporal location of quantum states in between preparations and measurements suggests that the quantum state is, at best, real only in a qualified sense, and not in the same way as paradigmatic examples of concrete objects like rocks, trees and human beings.

Similarly, consider the usage of the term 'realism' in philosophy of science, where realism about a theoretical entity is often defined as entailing a commitment to the claim that statements about this entity constitute knowledge about the world above and beyond descriptions of observables [240, 241]. This suggests an epistemic criterion for the reality of the quantum state: the state should be regarded as 'real' in this sense only if it is possible to provide evidence of knowledge of a quantum state which can be verified even by someone who does not know the corresponding description of experimental statistics. The no-go theorems of chapter 9 suggest that quantum states - or at least, quantum states of small dimension - do not satisfy this particular criterion of reality. Furthermore, the fact that the no-go theorems apply specifically to systems with finite Hilbert spaces might go some way towards explaining why it seems natural to be a realist about classical states yet it is much more controversial to be a realist about the quantum state, since classical systems have continuous (infinite-dimensional) state spaces and hence are not covered by these no-go theorems.

We do not of course claim to have *proven* anything about the reality of the quantum state, and nor would we attempt to do so - the word 'real' is ambiguous and highly context-dependent and probably not well-suited to be the subject of a mathematical argument. Our purpose here is simply to provide an opportunity to reflect on what exactly one might mean by the claim that the quantum state is real and to understand why the PBR theorem might not completely settle the issue. Indeed, we would suggest that a more nuanced understanding of the notion of reality might help overcome the long-standing deadlock between the ψ -ontic

and ψ -epistemic views: we have seen that the quantum state may be ‘real’ in the ψ -ontic sense, but nonetheless not ‘real’ according to other common uses of that term, it may turn out that the ψ -ontic and ψ -epistemic pictures, more broadly construed, are not genuinely in disagreement, and there might well be interpretative approaches according to which one can reasonably say that the quantum state is *both* ontic and epistemic.

We also take note of a number of interesting directions for future research. We have explored several generalisations of Kent’s original summoning task, but many possibilities remain. Another useful generalisation would be to consider relaxing the requirement that a protocol for a distribution task must always succeed, which would allow us to establish necessary and sufficient conditions for the existence of protocols that ensure the unknown state is returned with probability above some threshold value, and to study how the sets of feasible tasks induced by different values of the threshold are related to one another. In addition, there is scope for much more to be said about entanglement distribution: ideally one would like, as in the case of the summoning tasks, to be able to specify a full set of necessary and sufficient conditions for a task to be feasible for entanglement distribution and entanglement summoning respectively, together with an explicit prescription for a protocol that in each instance guarantees a successful response whenever there exists any protocol that guarantees a successful response.

With regard to bit commitment, there are a number of questions to be addressed regarding the composability of relativistic protocols. Since we can never guarantee that a relativistic bit commitment protocol will be secure under all possible compositions [235], it is particularly important to establish conditions under which the bit commitment protocols set out here and elsewhere [165,173,245,246] *can* be composed securely. For example, one important type of composition involves using sets of simultaneous bit commitments to commit to bit strings, as required by some of the protocols in chapter 9. We should also address the fact that relativistic bit commitment schemes are usually very time-sensitive; in particular, the commitment and unveiling typically take place at a spacelike or lightlike separation, with the duration of the protocol upper bounded by the time taken for light to travel from the commitment point to the unveiling point. Since there are practical constraints on the distances by which agents can be separated in real

applications, and we frequently have need of protocols with a non-negligible duration, there is a strong motivation for the development of protocols whose lifetime can be extended by employing a sequence of communications to maintain security indefinitely. Ref [172] sets out a method of extending a particular classical protocol in this way, but this method is not suitable for extending the bit commitments described in ref [165] since they employ a different geometrical arrangement of agents in spacetime, and hence a priority for future work is to develop new extension procedures for protocols of this type. We have in fact begun developing two such protocols: one requires Alice to initiate new commitments at the unveiling point(s) committing her to the data she would have handed over had she unveiled, while the other uses Rudich linkings [172,247] to allow Alice to make a series of commitments over time and prove to Bob that they are all commitments to the same bit. The former employs a number of commitments per round that increases exponentially with the number of rounds, which makes it impractical for most real-world applications, but the latter requires only a constant number of commitments per round, making it potentially feasible for real applications.

Our results on knowledge-concealing evidencing of knowledge also point to some interesting areas for future investigation. First, our security proofs were based in part on some assumptions about the composition of certain bit commitment protocols, and thus it would be useful to reexamine these proofs once more results have been established concerning the composability of such protocols. Second, although we have shown that under certain assumptions our new protocol can be regarded as ‘optimal’ in an asymptotic sense, we have not attempted to show optimality for any finite case, so a logical next step would be to either prove stronger optimality results or to develop a better protocol. Finally, with a practical protocol for something akin to zero-knowledge proving of knowledge of a quantum state in hand, it would be interesting to examine some important classical cryptographic protocols that depend on zero-knowledge proofs (for example, see refs [199,200,248,249]) to see if it is possible to find quantum generalisations that might have some useful advantages over the classical versions.

Appendix A

Supplementary Information for Chapter 9

A.1 Protocols with an abort option

Before proving our main results, we extend our definitions to allow for the possibility that Alice may abort the protocol.

In a KCEKQS protocol with abort option, after each round of receiving data each of Alice's agents generates one of two possible outcomes, 0 and 1, from the classical and quantum information in her possession. If she gets outcome 0, she announces to Bob's neighbouring agents, within a pre-agreed time interval, that the protocol is aborted. Bob's agent communicates this to Bob's other agents, and Alice's agent also announces the abort to all of Alice's other agents. Any agent who receives an abort message stops participating in the protocol from that point. If there is no abort, the protocol terminates after a fixed finite number of communications between Alice's and Bob's agents, as in the original definition. The allowed timings of abort announcements and of Bob's final announcement are fixed so that, in the event of a valid abort announcement, no outcome to the process is calculated by Bob. The possible outcomes are thus disjoint events 1, 0, abort and we denote the probabilities of two outcomes by $p(1)$, $p(0)$ and $p(\text{abort})$. We define non-triviality for protocols with an abort option just as before.¹

¹In fact, a weaker definition of non-triviality suffices for our no-go theorem below. Let

A.2 No-Go Theorems

A.2.1 Zero-knowledge

Theorem 8. *There exists no non-trivial KCEKQS protocol which is zero-knowledge.*

Proof. Consider a non-trivial KCEKQS protocol \mathcal{P} applied to a system Q_B with Hilbert space \mathcal{H}_B . In any such protocol Bob may replace the system Q_B by another system prepared in any state of his choice. Suppose that he does so, but that the protocol is otherwise honestly performed by both parties (where Alice believes, correctly, that the state of Q_B is η). The protocol then defines a quantum measurement operation P^η on \mathcal{H}_B . Now if Bob inputs a randomly chosen pure state $\phi \in \mathcal{H}_B$, non-triviality and continuity imply that there is a neighbourhood N of ϕ with nonzero measure δ (with respect to the uniform measure), such that

$$P(\eta \in N \mid \text{outcome } 1) > \delta.$$

Nontriviality and continuity also imply that $\epsilon_C < 1$ and that $P(\text{outcome } 1) > 0$. Hence, in any non-trivial KCEKQS protocol, with nonzero probability, Bob gains some information about η from this strategy. Since he retains Q_B , he can also carry out any measurement he wishes on Q_B , and combining information from the two processes gives him on average strictly more information about η than is available from the measurement alone. Hence any non-trivial KCEKQS protocol must have $\epsilon_K > \epsilon_M$ and therefore cannot be zero-knowledge. □

$p(x|\psi; \eta)$ be the probability distribution for the three outcomes $x \in \{0, 1, \text{abort}\}$ when the state of Q_B is ψ , Bob performs the protocol correctly, and Alice performs the version of the protocol that would be correct if she knew that the state of Q_B were η . Then it is sufficient that $p(x|\psi; \eta)$ depends non-trivially on ψ for fixed η . However we use the stronger definition of non-triviality here as its relevance is more intuitively clear.

A.2.2 Completeness vs soundness

Theorem 9. *In a KCEKQS protocol for a pure quantum state in a Hilbert space of dimension d , the completeness and soundness parameters obey*

$$\frac{\epsilon_S}{1 - \epsilon_C} \geq \frac{1}{d}.$$

Proof. A protocol P for KCEKQS may require either party to carry out measurements, to make random choices from a classical probability distribution, to introduce quantum states, and/or to send classical data. For any such protocol, we can define a related fully quantum protocol QP in which all data are introduced as quantum ancillae at the start and kept at the quantum level until the outcomes are obtained. Thus, to define QP from P , measurements are replaced by unitary measurement interactions (without extracting measurement data), classical random choices are replaced by interactions with entangled ‘quantum dice’ (without extracting data about the dice outcome) and classical communications are replaced by quantum communications of states in a pre-agreed orthonormal basis (for instance the computational basis). We take the protocol QP to proceed thus until the final stage. At this point, all the information in Alice’s possession is sent to one of Alice’s agents, who carries out a measurement giving her the outcome abort or not abort. She communicates this outcome to all of Alice’s and Bob’s agents, within an agreed time window, so that Bob’s agents all know the outcome after a prespecified coordinate time. If no abort is communicated to any of Bob’s agents within the prescribed time, they send all the information in their possession to one of Bob’s agents, who carries out a single measurement giving him the outcome 1 or 0.

We take the Hilbert spaces under Alice’s and Bob’s control initially to be \mathcal{H}_A and \mathcal{H}_B respectively, and the Hilbert spaces under Alice’s and Bob’s control at the end of the protocol to be \mathcal{H}_{A_f} and \mathcal{H}_{B_f} . The initial and final Hilbert spaces for each party are not necessarily identical, since the protocol may require states to be sent from one party to another. By introducing ancillae as necessary, we may take the final measurements by Alice and Bob to be projective measurements $(P_{A_f}, I_{A_f} - P_{A_f})$ and $(P_{B_f}, I_{B_f} - P_{B_f})$.

If both parties are honest, the probabilities of all outcomes are the same in QP as in P . If Bob is honest, then it cannot be disadvantageous to Alice to replace P by QP : all strategies available to her in the former can be replicated in the latter, and the latter also generally offers her further strategies. In particular, the strategy we define below for Alice has the same success probabilities in QP and P . We may therefore without loss of generality assume a fully quantum protocol.

Let S_{AB} be the initial state of all the ancillae introduced by Alice and Bob in QP , and let U^η be the unitary operation defined by the protocol up to the final outcome measurements when Alice honestly follows the protocol and believes the state is η . Here U^η includes any state transfers between the parties, as well as local unitaries applied by each party. Thus U^η maps $\mathcal{H}_A \otimes \mathcal{H}_B$ to $\mathcal{H}_{A_f} \otimes \mathcal{H}_{B_f}$.² We will write η_B for Bob's unknown state when we wish to emphasise that it is initially under Bob's control.

Thus we have:

$$\int \text{Tr}(((I_{A_f} - P_{A_f}) \otimes P_{B_f})U^\eta(\eta_B \otimes S_{AB})(U^\eta)^\dagger)d\mu(\eta) = 1 - \epsilon_C$$

where $d\mu(\cdot)$ denotes the uniform measure over quantum states, and the integral is performed over the entire Hilbert space of Q_B .

Now suppose Alice does not in fact know η . Then she may always adopt the strategy of choosing a random state ϕ from the Hilbert space of Q_B and proceeding with the protocol as if she knows that Q_B is in the state ϕ . Since aborting cannot increase $p(1)$, an optimum strategy to maximise $p(1)$ is never to abort, i.e. to take $P_{A_f} = 0$. Assuming that Bob always performs his part of the protocol honestly, the expected value of $p(1)$ is then

$$\int \int \text{Tr}((I_{A_f} \otimes P_{B_f})U^\phi(\eta_B \otimes S_{AB})(U^\phi)^\dagger)d\mu(\eta)d\mu(\phi) \leq \epsilon_S.$$

Moving the integral inside the trace, and noting that $\int \psi d\mu(\psi) = \frac{1}{d}I_B$, where

²We assume the protocol does not require any states to be discarded, since neither party can trust that the other will in fact discard states as required. Thus $\mathcal{H}_A \otimes \mathcal{H}_B$ and $\mathcal{H}_{A_f} \otimes \mathcal{H}_{B_f}$ are isomorphic, although the factors may be different.

I_B is the d -dimensional identity matrix on Q_B , we obtain

$$\frac{1}{d} \int \text{Tr}((I_{A_f} \otimes P_{B_f})U^\phi(I_B \otimes S_{AB})(U^\phi)^\dagger)d\mu(\phi) \leq \epsilon_S.$$

The left hand side is bounded below by

$$\frac{1}{d} \int \text{Tr}((I_{A_f} \otimes P_{B_f})U^\phi(\phi_B \otimes S_{AB})(U^\phi)^\dagger)d\mu(\phi),$$

which is bounded below by

$$\frac{1}{d} \int \text{Tr}(((I_{A_f} - P_{A_f}) \otimes P_{B_f})U^\phi(\phi_B \otimes S_{AB})(U^\phi)^\dagger)d\mu(\phi) = \frac{1}{d}(1 - \epsilon_C).$$

Hence $\frac{\epsilon_S}{1 - \epsilon_C} \geq \frac{1}{d}$.

□

A.3 Security Proofs for CR2

Soundness As in the proofs for CR1, we assume that the bit commitment protocol remains secure under parallel repetition. Hence we limit our discussion to strategies where Alice commits honestly to a single set of q out of d possible measurement outcomes in the bit commitment protocol.

If Alice has no information about the state η , then she has no better strategy than to choose q random orthogonal one-dimensional projections. Hence $\epsilon_S = \frac{q}{d}$, i.e. $q = \epsilon_S d$.

Knowledge-Concealing and Completeness The relativistic bit commitment protocol [13] is perfectly secure against Bob both for a single bit commitment and when composed to define a bit string commitment: in either case, Bob obtains no information about Alice's commitments unless and until she unveils them. Suppose then that Alice knows η and randomly chooses a projective measurement $\{P_i\}$ for which there exists a set S consisting of q measurement elements such that $\text{Tr}((\sum_{i \in S} P_i)\eta) = 1 - \epsilon_C$. If Bob follows the protocol and carries out this measurement, he obtains outcome P_i with probability $p_i = \text{Tr}(P_i\eta)$. If Alice's

unveiled commitment confirms that $i \in S$, Bob's maximum expected squared fidelity guess is $\eta = P_i$, which has squared fidelity p_i . Hence

$$\epsilon_K \geq \sum_{i \in S} (p_i)^2 \geq q \left(\frac{1 - \epsilon_C}{q} \right)^2$$

A.4 Security Proofs for QAB

Throughout this section, we use the notation $w(a, b) = \binom{a+b-1}{a}$.

Soundness If Alice does not in fact know η , the probability that Bob's measurement obtains a positive result is maximized if she chooses a single state ϕ and prepares all N systems $\{S_i\}$ in the state ϕ .

We use the fact that $\text{Tr}(\Pi_S(|\phi\rangle\langle\phi|^{\otimes N} \otimes \mathbb{I}^{\otimes M})\Pi_S) = \frac{w(M+N, d)}{w(N, d)}$ [250]. Thus when Alice employs the optimum strategy, the probability that Bob's measurement obtains a positive outcome is given by:

$$\begin{aligned} \epsilon_S &= \int \text{Tr}(\Pi(|\phi\rangle\langle\phi|^{\otimes N} \otimes \psi)\Pi) d\mu(\psi) = \text{Tr}(\Pi(|\phi\rangle\langle\phi|^{\otimes N} \otimes \frac{\mathbb{I}}{d})\Pi) = \frac{w(N+1, d)}{w(N, d)d} \\ &= \frac{1}{N+1} + \frac{N}{d(N+1)} \end{aligned} \tag{A.1}$$

In particular, for $N = 1$, we have

$$\epsilon_S = \frac{1}{2} + \frac{1}{2d}$$

and $\epsilon_S > \frac{1}{2}$ for any d .

Knowledge-Concealing If Alice follows the protocol honestly, Bob has available $N + 1$ copies of the state η , so the average squared fidelity between the true state η and his best possible guess from measuring $N + 1$ states is $\frac{N+2}{N+1+d}$

[225]. Thus $\epsilon_K = \frac{N+2}{N+1+d}$. From equation (A.1), we obtain $\epsilon_K > \frac{1}{d\epsilon_S}$ and $\epsilon_K - \epsilon_M > \frac{d-1}{d+1} \frac{N}{N+2} \frac{1}{d\epsilon_S}$. In particular, near CS-optimality ($\epsilon_S \approx \frac{1}{d}$) requires N large, which implies near-zero concealment ($\epsilon_K \approx 1$) and significant knowledge gain ($\epsilon_K - \epsilon_M \approx \frac{d-1}{d+1}$). For large d , this also implies near-complete knowledge gain ($\epsilon_K - \epsilon_M \approx 1$).

A.5 Security Proofs for QBA

Soundness Alice's most general strategy starts by pre-sharing some quantum state between A_1 and A_2 . A_1 then receives $(N+1)$ qudits from B_1 . She can then carry out any quantum operation on the quantum systems in her control, where the choice of operation may depend on the classical data sent by B_1 , to generate responses to B_1 that purport to initiate the bit commitment protocols. A_2 can similarly carry out any quantum operation on the quantum systems in her control, where the choice of operation may depend on the classical data sent by B_2 , to generate responses to B_2 that purport to sustain the protocols. A_1 can then carry out further quantum operation on the quantum systems under her control, where the choice of operations may depend on the index x sent by B_1 , to generate data that purport to unveil x as one of her committed strings.

A full security analysis against Alice thus requires a discussion of security for a protocol composed of two rounds of relativistic bit commitment sub-protocols together with the remaining steps of the protocol above. We leave this discussion for a future more general analysis of protocols within which relativistic bit commitments are suitably composable. For the present discussion, we will make the restrictive assumption that Alice honestly follows each relativistic bit commitment protocol [13] for the first two rounds and uses no quantum information in these relativistic bit commitment protocols. We now show that, under this assumption, if the protocol involves an unknown d -dimensional state, and we take $(N+1) = Md$ for integer M , and $q = \frac{N+1}{d} = M$, then $\epsilon_S = \frac{1}{d}$.

Proof. Alice can achieve a success probability of $\frac{q}{N+1}$ by simply committing to q random distinct indices, and hence $\epsilon_S \geq \frac{q}{N+1}$. We now show that also $\epsilon_S \leq \frac{q}{N+1}$, and hence $\epsilon_S = \frac{q}{N+1}$.

Suppose that A_1 and A_2 begin the KCEKQS protocol with no information about η , and that Bob honestly follows the protocol. From Alice's perspective, she simply receives $(N + 1)$ random pure qudits, since η is a random qudit and the remaining qudits are independently randomly chosen by Bob. Thus until B_1 gives A_1 classical information about the index assigned to Q_B , A_1 and A_2 's state of knowledge is exactly symmetrical with respect to all of the $N + 1$ qudits sent by B_1 . A_1 is required to initiate commitments, and A_2 to sustain them, before they receive the index assigned to Q_B , and therefore their commitment strategy on these rounds cannot depend on that index.

Thus we may without loss of generality calculate Alice's success probability by considering some fixed commitment strategy and averaging over all $N + 1$ possible values for the index of Q_B . For simplicity, we will analyse a related protocol T in which Alice always makes q commitments and subsequently unveils all q commitments. We will say that Alice is successful in this task if and only if at least she unveils at least one valid commitment to the correct index for Q_B (whether or not other unveiled commitments turn out to be valid commitments to any index). Alice's success probability ϵ_S in the KCEKQS protocol is no greater than her success probability in T , since in either case Alice succeeds if and only if she can unveil at least one commitment to the correct index.

If the probability distribution over the values of Alice's unveiled commitments in task T depends only on her initial commitment strategy and not on the classical information she subsequently obtains from Bob, then the values of her q unveiled commitments must be uncorrelated with the index of Q_B . Therefore the probability that one of these q unveiled commitments was originally a commitment to the correct index for Q_B can be no greater than $\frac{q}{N+1}$, with this bound saturated whenever Alice uses a strategy which always produces q valid commitments to different bit values in $\{1, 2, \dots, N + 1\}$. Thus Alice's success probability can be greater than $\frac{q}{N+1}$ only if the probability distribution over the values of her unveiled commitments fails to be independent of the classical information that she receives from Bob.

If so, under our assumptions about Alice's restricted strategy, Alice must have some freedom to choose whether to unveil 0 or 1 for at least one bit i out of the $\approx q \log(N + 1)$ bits to which she commits. More precisely, for at least two

different indices q, r in $\{1, 2, \dots, N + 1\}$, if $p_i(0|q)$ is the probability that Alice unveils a 0 for bit value i when Bob tells her the index of Q_B is q , and $p_i(1|r)$ is similarly defined, then $p_i(0|q) + p_i(1|r) > 1$. But this contradicts the security of the relativistic bit commitment protocol [13], under our assumptions about Alice's restricted strategy. Hence, given those assumptions, Alice cannot succeed in protocol T with probability greater than $\frac{q}{N+1}$ and thus we have $\epsilon_S \leq \frac{q}{N+1}$. \square

Completeness We now show that $\epsilon_C \rightarrow 0$ as $N \rightarrow \infty$.

Proof. If Alice does know the state η , she will get a positive outcome on Q_B . The probability that her commitment fails to be accepted is thus given by:

$$\epsilon_C = \sum_{x=q}^N P(X_N = x) \frac{x + 1 - q}{x + 1}, \quad (\text{A.2})$$

where $p(X_N = x)$ is the probability that x out of Alice's N measurements on the N systems S_i obtain the result η .

If Bob is honest and chooses the states of the systems S_i at random, the result η is obtained on each run with probability $\frac{1}{d}$. Hence X_N is binomially distributed:

$$P(X_N = x) = \binom{N}{x} \frac{1}{d^x} \left(1 - \frac{1}{d}\right)^{N-x}. \quad (\text{A.3})$$

The distribution X_N has mean $\frac{N}{d}$. Hoeffding's inequality implies that

$$P(X_N \geq \frac{N}{d} + \epsilon N) \leq \exp(-2\epsilon^2 N). \quad (\text{A.4})$$

Taking $q = \frac{N+1}{d}$, it follows from equation (A.2) that

$$\epsilon_C \leq \epsilon d p_{\text{mode}} + \exp(-2\epsilon^2 N) \quad (\text{A.5})$$

for any $\epsilon > 0$, where p_{mode} is the maximum over x of the binomial distribution

$P(X_N = x)$. We can obtain an adequate bound simply by using $p_{\text{mode}} \leq 1$.³

Taking, for example, $\epsilon = \frac{1}{2}N^{-\frac{1}{2}}(\log N)^{\frac{1}{2}}$, we see that $\epsilon_C \rightarrow 0$ as $N \rightarrow \infty$. □

Knowledge-Concealing Bob begins with one copy of η . If dishonest, he may combine this with any ancillae he wishes, carry out any quantum operations he wishes, and produce a (perhaps highly entangled) state including $(N + 1)$ qudits that he sends to Alice, together with a system that he retains. He may then carry out any measurement he wishes to produce an index x . Alice responds, effectively, with either a 1 (if she unveils a commitment to x) or a 0 (if she fails to unveil a commitment to x). Bob may then carry out any measurement he wishes on his retained system to produce his maximum possible expected squared fidelity estimate of η . This measurement choice may depend on Alice's response.

Effectively, Bob's task is to optimize his state estimation in this scenario. His overall strategy S is fixed up to Alice's response, but then may involve different state estimation strategies depending on the one bit of information supplied by Alice. For any given overall strategy S , the overall expected squared fidelity of his estimate is

$$p_S(0)f_S(0) + p_S(1)f_S(1), \quad (\text{A.6})$$

where $p_S(b)$ is the probability of outcome b given strategy S and $f_S(b)$ is the expected squared fidelity obtained from S conditioned on outcome b .

Suppose that Alice does not respond at all. Bob may simply follow the fixed strategy S_b given by following S and assuming outcome b ; this produces an expected squared fidelity of at least $p_S(b)f_S(b)$. Since S_b is a possible strategy for the standard task of state estimation given one copy of an unknown qudit (and no further information), we have that

$$p_S(b)f_S(b) \leq f_{\text{max}} = \frac{2}{d+1} \quad (\text{A.7})$$

³A tighter bound for $N \gg d$ follows from the normal approximation to the binomial distribution, which gives $p_{\text{mode}} \approx \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{N\frac{1}{d}(1-\frac{1}{d})}}$.

where the right hand side is the maximum expected square fidelity obtainable from any state estimation strategy on an unknown qudit [225]. Hence Bob's overall expected squared fidelity is given by:

$$p_S(0)f_S(0) + p_S(1)f_S(1) \leq \frac{4}{d+1}. \quad (\text{A.8})$$

That is, $\epsilon_K \leq \frac{4}{d+1}$.

Bibliography

- [1] Bell, J. S. *Speakable and Unspeakable in Quantum Mechanics* (Cambridge University Press, 1987).
- [2] Seevinck, M. P. Can quantum theory and special relativity peacefully co-exist? *ArXiv e-prints* (2010). [quant-ph/1010.3714](https://arxiv.org/abs/quant-ph/1010.3714).
- [3] Wootters, W. & Zurek, W. A single quantum cannot be cloned. *Nature* **299**, 802 – 803 (1982).
- [4] Howard, M., Wallman, J., Veitch, V. & Emerson, J. Contextuality supplies the ‘magic’ for quantum computation. *Nature* **510**, 351–355 (2014). [quant-ph/1401.4174](https://arxiv.org/abs/quant-ph/1401.4174).
- [5] Modi, K., Brodutch, A., Cable, H., Paterek, T. & Vedral, V. The classical-quantum boundary for correlations: Discord and related measures. *Rev. Mod. Phys.* **84**, 1655–1707 (2012). URL <http://link.aps.org/doi/10.1103/RevModPhys.84.1655>.
- [6] Kent, A. Quantum tasks in minkowski space. *Classical and Quantum Gravity* **29**, 224013 (2012).
- [7] Hayden, P. & May, A. Summoning information in spacetime, or where and when can a qubit be? *Journal of Physics A: Mathematical and Theoretical* **49**, 175304 (2016). URL <http://stacks.iop.org/1751-8121/49/i=17/a=175304>.
- [8] Mayers, D. Unconditionally secure quantum bit commitment is impossible. *Physical review letters* **78**, 3414 (1997).

- [9] Adlam, E. & Kent, A. Quantum paradox of choice: More freedom makes summoning a quantum state harder. *Physical Review A* **93**, 062327 (2016). 1509.04226.
- [10] Adlam, E. & Kent, A. A Quantum Paradox of Choice and Purported Classical Analogues. *ArXiv e-prints* (2015). quant-ph/1509.08094.
- [11] Finkelstein, J. A classical paradox of choice. *ArXiv e-prints* (2015). 1509.06692.
- [12] Kent, A. Unconditionally Secure Bit Commitment. *Physical Review Letters* **83**, 1447–1450 (1999). quant-ph/9810068.
- [13] Kent, A. Secure classical bit commitment using fixed capacity communication channels. *J. Cryptology* **18**, 313–335 (2011).
- [14] Lo, H.-K. & Chau, H. Why quantum bit commitment and ideal quantum coin tossing are impossible. *Physica D: Nonlinear Phenomena* **120**, 177 – 187 (1998). URL <http://www.sciencedirect.com/science/article/pii/S0167278998000530>.
- [15] Lo, H.-K. & Chau, H. F. Is Quantum Bit Commitment Really Possible? *Physical Review Letters* **78**, 3410–3413 (1997). quant-ph/9603004.
- [16] Adlam, E. & Kent, A. Deterministic relativistic quantum bit commitment. *International Journal of Quantum Information* **13**, 1550029 (2015). quant-ph/1504.00943.
- [17] Adlam, E. & Kent, A. Device-independent relativistic quantum bit commitment. *Physical Review A* **92**, 022315 (2015). quant-ph/1504.00944.
- [18] Adlam, E. & Kent, A. Knowledge-Concealing Evidencing of Knowledge about a Quantum State. *ArXiv e-prints* (2017). 1706.06963.
- [19] Kragh, H. *Quantum Generations: A History of Physics in the Twentieth Century* (Princeton University Press, 2002). URL <https://books.google.co.uk/books?id=ELrFDI1dlawC>.

- [20] Brown, L., Pippard, B. & Pais, A. *Twentieth Century Physics* (CRC Press, 1995). URL <https://books.google.co.uk/books?id=oQn5ybiQKAoC>.
- [21] Redhead, M. Relativity and quantum mechanics - conflict or peaceful coexistence? *Annals of the New York Academy of Sciences* **480**, 14–20 (1986). URL <http://dx.doi.org/10.1111/j.1749-6632.1986.tb12405.x>.
- [22] Wallace, D. The quantization of gravity - an introduction. *ArXiv General Relativity and Quantum Cosmology e-prints* (2000). [gr-qc/0004005](https://arxiv.org/abs/gr-qc/0004005).
- [23] Peskin, M. & Schroeder, D. *An Introduction To Quantum Field Theory*. Frontiers in physics (Westview Press, 1995). URL <https://books.google.co.uk/books?id=EVeNNcslvX0C>.
- [24] Lancaster, T. & Blundell, S. *Quantum Field Theory for the Gifted Amateur* (OUP Oxford, 2014). URL <https://books.google.co.uk/books?id=nIk6AwAAQBAJ>.
- [25] Shimony, A. *The Search for a Naturalistic World View* (Cambridge University Press, 1993). URL <https://books.google.co.uk/books?id=Aee0XfBUDZQC>.
- [26] Peres, A. & Terno, D. R. Quantum information and relativity theory. *Rev. Mod. Phys.* **76**, 93–123 (2004). URL <http://link.aps.org/doi/10.1103/RevModPhys.76.93>.
- [27] Einstein, A. On the electrodynamics of moving bodies. *Annalen der Physik* **17**, 891 – 921 (1905).
- [28] Minkowski, H. Raum und ziet. In Sommerfeld, A. (ed.) *The Principle of Relativity* (Dover, New York, 1909).
- [29] Rindler, W. Length Contraction Paradox. *American Journal of Physics* **29**, 365–366 (1961).

- [30] French, A. *Special Relativity* (Taylor & Francis, 1968). URL <https://books.google.co.uk/books?id=8jPVRXNXj28C>.
- [31] Winnie, J. A. The causal theory of space-time. In Earman, J., Glymour, C. & Stachel, J. (eds.) *Foundations of Space-Time Theories* (University of Minnesota Press, 1977).
- [32] Wald, R. *General Relativity* (University of Chicago Press, 2010). URL <https://books.google.co.uk/books?id=9S-hzg6-moYC>.
- [33] Einstein, A., Podolsky, B. & Rosen, N. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.* **47**, 777–780 (1935). URL <http://link.aps.org/doi/10.1103/PhysRev.47.777>.
- [34] de Regt, H. W. Mara beller, quantum dialogue – the making of a revolution. *Erkenntnis* **56**, 247–252 (2002). URL <http://dx.doi.org/10.1023/A:1015612805156>.
- [35] Einstein, A. Quantum mechanics and reality. *Dialectica* **2 3 - 4**, 320 – 324 (1948).
- [36] Pais, A. Einstein and the quantum theory. *Rev. Mod. Phys.* **51**, 863–914 (1979). URL <http://link.aps.org/doi/10.1103/RevModPhys.51.863>.
- [37] Goldstein, S. Bohmian mechanics. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2016), fall 2016 edn.
- [38] Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, New York, NY, USA, 2011), 10th edn.
- [39] Gell-Mann, M. Questions for the future. In *Wolfson College Lectures* (Oxford University Press, 1980).

- [40] Wiseman, H. M. From Einstein's theorem to Bell's theorem: a history of quantum non-locality. *Contemporary Physics* **47**, 79–88 (2006). [quant-ph/0509061](https://arxiv.org/abs/quant-ph/0509061).
- [41] Maudlin, T. *Quantum Non-Localilty and Relativity: Metaphysical Implications of Modern Physics* (Wiley, 2011). URL <https://books.google.co.uk/books?id=grNxuUFyO78C>.
- [42] Bell, J. S. On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics* **38**, 447 (1966).
- [43] Shimony, A. Bell's theorem. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2013), winter 2013 edn.
- [44] Clauser, J. F., Horne, M. A., Shimony, A. & Holt, R. A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**, 880–884 (1969). URL <http://link.aps.org/doi/10.1103/PhysRevLett.23.880>.
- [45] Aspect, A., Grangier, P. & Roger, G. Experimental Tests of Realistic Local Theories via Bell's Theorem. *Phys. Rev. Lett.* **47**, 460–463 (1981). URL <http://link.aps.org/doi/10.1103/PhysRevLett.47.460>.
- [46] Kielpinski, D. *et al.* Experimental violation of a Bell's inequality with efficient detection. *Nature* **409**, 791–794 (2001).
- [47] Hensen, B. *et al.* Experimental loophole-free violation of a Bell inequality using entangled electron spins separated by 1.3 km. *Nature* **526**, 682 – 686 (2015). [1508.05949](https://doi.org/10.1038/526682a).
- [48] Bohm, D. *Quantum Theory*. Dover Books on Physics Series (Dover Publications, 1951). URL <https://books.google.com.au/books?id=9DWim3Rhym5C>.

- [49] Koashi, M. & Winter, A. Monogamy of quantum entanglement and other correlations. *Physical Review A* **69**, 022309 (2004). [quant-ph/quant-ph/0310037](#).
- [50] Coffman, V., Kundu, J. & Wootters, W. K. Distributed entanglement. *Physical Review A* **61**, 052306 (2000). [quant-ph/quant-ph/9907047](#).
- [51] Seevinck, M. Monogamy of Correlations vs. Monogamy of Entanglement. *ArXiv e-prints* (2009). [quant-ph/0908.1867](#).
- [52] Toner, B. Monogamy of non-local quantum correlations. *Proceedings of the Royal Society of London Series A* **465**, 59–69 (2009). [quant-ph/quant-ph/0601172](#).
- [53] Toner, B. & Verstraete, F. Monogamy of Bell correlations and Tsirelson’s bound. *eprint arXiv:quant-ph/0611001* (2006). [quant-ph/quant-ph/0611001](#).
- [54] Barrett, J. Information processing in generalized probabilistic theories. *eprint arXiv:quant-ph/0508211* (2005). [quant-ph/quant-ph/0508211](#).
- [55] d’Espagnat, B. *Conceptual Foundations of Quantum Mechanics* (Addison-Wesley, Advanced Book Program, 1971).
- [56] Busch, P., Lahti, J. & Mittelstaedt, P. *The Quantum Theory of Measurement* (Springer-Verlag Berlin Heidelberg, 1996).
- [57] Landau, L. & Lifshitz, E. *Quantum Mechanics: Non-Relativistic Theory*. Teoreticheskaia fizika (Elsevier Science, 2013). URL <https://books.google.co.uk/books?id=neBbAwAAQBAJ>.
- [58] Paris, M. G. A. The modern tools of quantum mechanics. A tutorial on quantum states, measurements, and operations. *European Physical Journal Special Topics* **203**, 61–86 (2012). [1110.6815](#).
- [59] de Muynck, W. M. POVMs: a Small but Important Step Beyond Standard Quantum Mechanics. In Nieuwenhuizen, T. M., Mehmani, B., Špička, V.,

Aghdami, M. J. & Khrennikov, A. Y. (eds.) *Beyond the Quantum*, 69–79 (World Scientific, 2007). [quant-ph/0608087](https://arxiv.org/abs/quant-ph/0608087).

- [60] Rickles, D. *The Ashgate Companion to Contemporary Philosophy of Physics*. Ashgate companion (Ashgate Pub. Limited, 2008). URL <https://books.google.co.uk/books?id=qQZugqxWAcC>.
- [61] Weinberg, S. Quantum mechanics without state vectors. *Physical Review A* **90**, 042102 (2014). [1405.3483](https://arxiv.org/abs/1405.3483).
- [62] Kochen, S. & Specker, E. The problem of hidden variables in quantum mechanics. In Hooker, C. (ed.) *The Logico-Algebraic Approach to Quantum Mechanics*, vol. 5a of *The University of Western Ontario Series in Philosophy of Science*, 293–328 (Springer Netherlands, 1975). URL http://dx.doi.org/10.1007/978-94-010-1795-4_17.
- [63] Held, C. The Kochen-Specker Theorem. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2014), winter 2014 edn.
- [64] Spekkens, R. W. Contextuality for preparations, transformations, and unsharp measurements. *Physical Review A* **71**, 052108 (2005). [quant-ph/0406166](https://arxiv.org/abs/quant-ph/0406166).
- [65] Harrigan, N. & Spekkens, R. W. Einstein, Incompleteness, and the Epistemic View of Quantum States. *Foundations of Physics* **40**, 125–157 (2010). [quant-ph/0706.2661](https://arxiv.org/abs/quant-ph/0706.2661).
- [66] Cabello, A., Severini, S. & Winter, A. (Non-)Contextuality of Physical Theories as an Axiom. *ArXiv e-prints* (2010). [quant-ph/1010.2163](https://arxiv.org/abs/quant-ph/1010.2163).
- [67] Bennett, C. H. *et al.* Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys. Rev. Lett.* **70**, 1895–1899 (1993). URL <http://link.aps.org/doi/10.1103/PhysRevLett.70.1895>.

- [68] Deutsch, D. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London* **400**, 97–117 (1985).
- [69] Feynman, R. P. Simulating physics with computers. *International journal of theoretical physics* **21**, 467–488 (1982).
- [70] Benioff, P. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics* **22**, 563–591 (1980). URL <https://doi.org/10.1007/BF01011339>.
- [71] Ekert, A. K. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.* **67**, 661–663 (1991). URL <http://link.aps.org/doi/10.1103/PhysRevLett.67.661>.
- [72] Brassard, G. *Modern cryptology: A tutorial*, vol. 325 (Springer-Verlag New York, 1998).
- [73] Bennett, C., Brassard, G., Crépeau, C. & Skubiszewska, M.-H. Practical quantum oblivious transfer. In Feigenbaum, J. (ed.) *Advances in Cryptology CRYPTO 91*, vol. 576 of *Lecture Notes in Computer Science*, 351–366 (Springer Berlin Heidelberg, 1992). URL http://dx.doi.org/10.1007/3-540-46766-1_29.
- [74] Scarani, V. Feats, Features and Failures of the PR-box. In Bassi, E. B. A., Dürr, D., Weber, T. & Zanghi, N. (eds.) *Quantum Mechanics: Are there Quantum Jumps? and On the Present Status of Quantum Mechanics*, vol. 844 of *American Institute of Physics Conference Series*, 309–320 (2006). quant-ph/quant-ph/0603017.
- [75] Acín, A. *et al.* Device-Independent Security of Quantum Cryptography against Collective Attacks. *Physical Review Letters* **98**, 230501 (2007). quant-ph/quant-ph/0702152.

- [76] Masanes, L., Renner, R., Christandl, M., Winter, A. & Barrett, J. Full security of quantum key distribution from no-signaling constraints. *IEEE Transactions on Information Theory* **60**, 4973–4986 (2014).
- [77] Colbeck, R. & Kent, A. Variable-bias coin tossing. *Physical Review A* **73**, 032320 (2006). [quant-ph/0508149](https://arxiv.org/abs/quant-ph/0508149).
- [78] Deutsch, D. & Jozsa, R. Rapid solution of problems by quantum computation. *Proceedings: Mathematical and Physical Sciences* **439**, 553–558 (1992). URL <http://www.jstor.org/stable/52182>.
- [79] Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* **26**, 1484–1509 (1997). URL <https://doi.org/10.1137/S0097539795293172>. <https://doi.org/10.1137/S0097539795293172>.
- [80] Yuen, H. Amplification of quantum states and noiseless photon amplifiers. *Physics Letters, Section A: General, Atomic and Solid State Physics* **113**, 405–407 (1986).
- [81] Duan, L. & Guo, G. Probabilistic Cloning and Identification of Linearly Independent Quantum States. *Physical Review Letters* **80**, 4999–5002 (1998). [quant-ph/9804064](https://arxiv.org/abs/quant-ph/9804064).
- [82] Barnum, H., Caves, C. M., Fuchs, C. A., Jozsa, R. & Schumacher, B. Non-commuting Mixed States Cannot Be Broadcast. *Physical Review Letters* **76**, 2818–2821 (1996). [quant-ph/9511010](https://arxiv.org/abs/quant-ph/9511010).
- [83] Holevo, A. S. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii* **9**, 3–11 (1973).
- [84] Petz, D. Entropy, von Neumann and the von Neumann entropy. In Re-dei, M. & Stolzner, M. (eds.) *John von Neumann and the Foundations of Quantum Physics* (Kluwer, 2001). [math-ph/0102013](https://arxiv.org/abs/math-ph/0102013).

- [85] Thomas, J. & Cover, T. *Elements of Information Theory* (Wiley, 2006).
- [86] Rényi, A. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, 547–561 (University of California Press, Berkeley, Calif., 1961). URL <http://projecteuclid.org/euclid.bsmsp/1200512181>.
- [87] Müller-Lennert, M., Dupuis, F., Szehr, O., Fehr, S. & Tomamichel, M. On quantum Rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics* **54**, 122203–122203 (2013). 1306.3142.
- [88] Bengtsson, I. & Życzkowski, K. *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge University Press, 2007). URL <https://books.google.co.uk/books?id=aA4vXMbuOTUC>.
- [89] Horodecki, M., Oppenheim, J. & Winter, A. Partial quantum information. *Nature* **436**, 673–676 (2005). [quant-ph/0505062](https://arxiv.org/abs/quant-ph/0505062).
- [90] Gisin, N., Ribordy, G., Tittel, W. & Zbinden, H. Quantum cryptography. *Reviews of Modern Physics* **74**, 145–195 (2002). [quant-ph/0101098](https://arxiv.org/abs/quant-ph/0101098).
- [91] Singh, S. *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography* (Knopf Doubleday Publishing Group, 2011). URL <https://books.google.co.uk/books?id=o3YbiVuTg70C>.
- [92] Bennett, C. H., Brassard, G., Breidbart, S. & Wiesner, S. Quantum cryptography, or unforgeable subway tokens. In Chaum, D., Rivest, R. L. & Sherman, A. T. (eds.) *Advances in Cryptology: Proceedings of Crypto 82*, 267–275 (Springer US, Boston, MA, 1983). URL https://doi.org/10.1007/978-1-4757-0602-4_26.
- [93] Bennett, C. H. & Brassard, G. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science* **560**, 7 – 11 (2014). URL <http://www.sciencedirect.com/science/>

article/pii/S0304397514004241. Theoretical Aspects of Quantum Cryptography celebrating 30 years of BB84.

- [94] Shor, P. W. & Preskill, J. Simple proof of security of the BB84 quantum key distribution protocol. *Physical review letters* **85**, 441 (2000).
- [95] Mayers, D. Unconditional security in quantum cryptography. *Journal of the ACM (JACM)* **48**, 351–406 (2001).
- [96] Gerjuoy, E. Shor’s factoring algorithm and modern cryptography. An illustration of the capabilities inherent in quantum computers. *American Journal of Physics* **73**, 521–540 (2005). [quant-ph/0411184](http://doi.org/10.1119/1.1841118).
- [97] Biham, E., Boyer, M., Boykin, P. O., Mor, T. & Roychowdhury, V. A proof of the security of quantum key distribution (extended abstract). In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing, STOC ’00*, 715–724 (ACM, New York, NY, USA, 2000). URL <http://doi.acm.org/10.1145/335305.335406>.
- [98] Mayers, D. Unconditional security in quantum cryptography. *Journal of the ACM (JACM)* **48**, 351–406 (2001).
- [99] Lo, H.-K. & Chau, H. F. Unconditional security of quantum key distribution over arbitrarily long distances. *science* **283**, 2050–2056 (1999).
- [100] Bennett, C. H., Brassard, G. & Robert, J.-M. Privacy amplification by public discussion. *SIAM Journal on Computing* **17**, 210–229 (1988). URL <https://doi.org/10.1137/0217014>. <https://doi.org/10.1137/0217014>.
- [101] Bennett, C. H. Quantum cryptography using any two nonorthogonal states. *Phys. Rev. Lett.* **68**, 3121–3124 (1992). URL <https://link.aps.org/doi/10.1103/PhysRevLett.68.3121>.
- [102] Aharonov, D., Ta-Shma, A., Vazirani, U. V. & Yao, A. C. Quantum bit escrow. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing, STOC*, 705–714 (ACM, New York, 2000). URL <http://doi.acm.org/10.1145/335305.335404>.

- [103] Colbeck, R. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. Ph.D. thesis (2009).
- [104] Vazirani, U. & Vidick, T. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 61–76 (ACM, 2012).
- [105] Hillery, M., Bužek, V. & Berthiaume, A. Quantum secret sharing. *Physical Review A* **59**, 1829–1834 (1999). [quant-ph/9806063](https://arxiv.org/abs/quant-ph/9806063).
- [106] Fredenhagen, K. & Rejzner, K. Perturbative algebraic quantum field theory. In *Mathematical Aspects of Quantum Field Theories*, 17–55 (Springer, 2015).
- [107] Wallace, D. Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **42**, 116 – 125 (2011). URL <http://www.sciencedirect.com/science/article/pii/S1355219810000808>. Philosophy of Quantum Field Theory.
- [108] Martín-Martínez, E. & Menicucci, N. C. Entanglement in curved spacetimes and cosmology. *Classical and Quantum Gravity* **31**, 214001 (2014). [1408.3420](https://arxiv.org/abs/1408.3420).
- [109] Zych, M., Pikovski, I., Costa, F. & Brukner, Č. General relativistic effects in quantum interference of clocks . In *Journal of Physics Conference Series*, vol. 723 of *Journal of Physics Conference Series*, 012044 (2016). [1607.04022](https://arxiv.org/abs/1607.04022).
- [110] Ahmadi, M., Bruschi, D. E. & Fuentes, I. Quantum metrology for relativistic quantum fields. *Physical Review D* **89**, 065028 (2014). [1312.5707](https://arxiv.org/abs/1312.5707).
- [111] Bruschi, D. E., Friis, N., Fuentes, I. & Weinfurter, S. On the robustness of entanglement in analogue gravity systems. *New Journal of Physics* **15**, 113016 (2013). [1305.3867](https://arxiv.org/abs/1305.3867).

- [112] Rovelli, C. Loop quantum gravity. *Living Reviews in Relativity* **11**, 5 (2008). URL <http://dx.doi.org/10.12942/lrr-2008-5>.
- [113] Blau, M. & Theisen, S. String theory as a theory of quantum gravity: a status report. *General Relativity and Gravitation* **41**, 743–755 (2009). URL <http://dx.doi.org/10.1007/s10714-008-0752-z>.
- [114] Conway, J. & Kochen, S. The free will theorem. *Foundations of Physics* **36**, 1441–1473 (2006). URL <http://dx.doi.org/10.1007/s10701-006-9068-6>.
- [115] Colbeck, R. & Renner, R. Is a system’s wave function in one-to-one correspondence with its elements of reality? *Phys. Rev. Lett.* **108**, 150402 (2012). URL <http://link.aps.org/doi/10.1103/PhysRevLett.108.150402>.
- [116] Srednicki, M. Subjective and objective probabilities in quantum mechanics. *Physical Review A* **71**, 052107 (2005). [quant-ph/0501009](http://arxiv.org/abs/quant-ph/0501009).
- [117] Maudlin, T. What could be objective about probabilities? *Studies in History and Philosophy of Science Part B* **38**, 275–291 (2007).
- [118] Myrvold, W. Philosophical issues in quantum theory. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2017), spring 2017 edn.
- [119] Wallace, D. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation* (OUP Oxford, 2012). URL <https://books.google.co.uk/books?id=jNaSBAAAQBAJ>.
- [120] Adlam, E. The Problem of Confirmation in the Everett Interpretation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **47**, 21 – 32 (2014). URL <http://www.sciencedirect.com/science/article/pii/S1355219814000276>.

- [121] Holland, P. *The Quantum Theory of Motion: An Account of the de Broglie-Bohm Causal Interpretation of Quantum Mechanics* (Cambridge University Press, 1995). URL <https://books.google.co.uk/books?id=BsEfVBzToRMC>.
- [122] Frigg, R. *GRW Theory (Ghirardi, Rimini, Weber Model of Quantum Mechanics)*, 266–270 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009). URL http://dx.doi.org/10.1007/978-3-540-70626-7_81.
- [123] Ghirardi, G. C., Rimini, A. & Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **34**, 470–491 (1986). URL <http://link.aps.org/doi/10.1103/PhysRevD.34.470>.
- [124] Tumulka, R. A Relativistic Version of the Ghirardi Rimini Weber Model. *Journal of Statistical Physics* **125**, 821–840 (2006). [quant-ph/0406094](http://arxiv.org/abs/quant-ph/0406094).
- [125] Tumulka, R. On spontaneous wave function collapse and quantum field theory. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **462**, 1897–1908 (2006).
- [126] Ghirardi, G. Collapse theories. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2016), spring 2016 edn.
- [127] Penrose, R. On gravity’s role in quantum state reduction. *General Relativity and Gravitation* **28**, 581–600 (1996). URL <http://dx.doi.org/10.1007/BF02105068>.
- [128] Fuchs, C. A. QBism, the Perimeter of Quantum Bayesianism. *ArXiv e-prints* (2010). [1003.5209](http://arxiv.org/abs/1003.5209).
- [129] Fuchs, C. A. On participatory realism. In *Information and Interaction*, 113–134 (Springer, 2017).
- [130] Timpson, C. G. Quantum Bayesianism: A study. *Studies in the History and Philosophy of Modern Physics* **39**, 579–609 (2008). [0804.2047](http://arxiv.org/abs/0804.2047).

- [131] Almeida, M. L. *et al.* Guess Your Neighbor's Input: A Multipartite Non-local Game with No Quantum Advantage. *Physical Review Letters* **104**, 230404 (2010). [quant-ph/1003.3844](https://arxiv.org/abs/quant-ph/1003.3844).
- [132] Vidick, T. The complexity of entangled games. *Phd Thesis, University of California, Berkeley* (2011).
- [133] Mančinska, L., Roberson, D. E. & Varvitsiotis, A. Deciding the existence of perfect entangled strategies for nonlocal games. *Chicago Journal of Theoretical Computer Science* 1–16 (2016). [quant-ph/1506.07429](https://arxiv.org/abs/quant-ph/1506.07429).
- [134] Tomamichel, M., Fehr, S., Kaniewski, J. & Wehner, S. One-Sided Device-Independent QKD and Position-Based Cryptography from Monogamy Games. In Johansson, T. & Nguyen, P. (eds.) *Advances in Cryptology EUROCRYPT 2013*, vol. 7881 of *Lecture Notes in Computer Science*, 609–625 (Springer Berlin Heidelberg, 2013). URL http://dx.doi.org/10.1007/978-3-642-38348-9_36.
- [135] Gawron, P., Miszczak, J. A. & Sładkowski, J. Noise Effects in Quantum Magic Squares Game. *International Journal of Quantum Information* **6**, 667 – 673 (2008). [quant-ph/0801.4848](https://arxiv.org/abs/quant-ph/0801.4848).
- [136] Piotrowski, E. W. & Sładkowski, J. An invitation to quantum game theory. *International Journal of Theoretical Physics* **42**, 1089–1099 (2003). URL <https://doi.org/10.1023/A:1025443111388>.
- [137] Flitney, A. P. & Abbott, D. An introduction to quantum game theory. *Fluctuation and Noise Letters* **02**, 175– 187 (2002). URL <http://www.worldscientific.com/doi/abs/10.1142/S0219477502000981>. <http://www.worldscientific.com/doi/pdf/10.1142/S0219477502000981>.
- [138] Eisert, J., Wilkens, M. & Lewenstein, M. Quantum Games and Quantum Strategies. *Physical Review Letters* **83**, 3077–3080 (1999). [quant-ph/9806088](https://arxiv.org/abs/quant-ph/9806088).

- [139] Meyer, D. A. Quantum strategies. *Phys. Rev. Lett.* **82**, 1052–1055 (1999). URL <http://link.aps.org/doi/10.1103/PhysRevLett.82.1052>.
- [140] Werner, R. F. Optimal cloning of pure states. *Phys. Rev. A* **58**, 1827–1832 (1998). URL <http://link.aps.org/doi/10.1103/PhysRevA.58.1827>.
- [141] Lee, C. F. & Johnson, N. F. Game-theoretic discussion of quantum state estimation and cloning. *Physics Letters A* **319**, 429–433 (2003). [quant-ph/0207139](https://arxiv.org/abs/quant-ph/0207139).
- [142] Meyer, D. A. Quantum games and quantum algorithms. In Lomonaco, S. & Brandt, H. (eds.) *Contemporary Mathematics: Quantum Computation and Information*, vol. 305 (AMS, 2000).
- [143] Iqbal, A. & Toor, A. H. Quantum repeated games. *Physics Letters A* **300**, 541–546 (2002). [quant-ph/0203044](https://arxiv.org/abs/quant-ph/0203044).
- [144] Du, J. *et al.* Experimental Realization of Quantum Games on a Quantum Computer. *Physical Review Letters* **88**, 137902 (2002). [quant-ph/0104087](https://arxiv.org/abs/quant-ph/0104087).
- [145] Du, J. *et al.* Remark On Quantum Battle of The Sexes Game. *eprint arXiv:quant-ph/0103004* (2001). [quant-ph/0103004](https://arxiv.org/abs/quant-ph/0103004).
- [146] Marinatto, L. & Weber, T. A quantum approach to static games of complete information. *Physics Letters A* **272**, 291–303 (2000). [quant-ph/0004081](https://arxiv.org/abs/quant-ph/0004081).
- [147] Flitney, A. P. & Abbott, D. Quantum version of the Monty Hall problem. *Physical Review A* **65**, 062318 (2002). [quant-ph/0109035](https://arxiv.org/abs/quant-ph/0109035).
- [148] D’Ariano, G. M. *et al.* The Quantum Monty Hall Problem. *Quantum Info. Comput.* **2**, 355–466 (2002). URL <http://dl.acm.org/citation.cfm?id=2011483.2011486>.

- [149] Schmid, C. *et al.* Experimental implementation of a four-player quantum game. *New Journal of Physics* **12**, 063031 (2010). 0901.0063.
- [150] Benjamin, S. C. & Hayden, P. M. Multiplayer quantum games. *Physical Review A* **64**, 030301 (2001). quant-ph/0007038.
- [151] Kent, A. A no-summoning theorem in relativistic quantum theory. *Quantum Information Processing* **12**, 1023–1032 (2013). URL <http://dx.doi.org/10.1007/s11128-012-0431-6>.
- [152] Gottesman, D. & Chuang, I. L. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* **402**, 390–393 (1999). quant-ph/quant-ph/9908010.
- [153] Aharonov, Y. & Vaidman, L. *The Two-State Vector Formalism of Quantum Mechanics*, 369–412 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2002). URL https://doi.org/10.1007/3-540-45846-8_13.
- [154] Schrödinger, E. Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften* **23**, 807–812 (1935). URL <http://dx.doi.org/10.1007/BF01491891>.
- [155] Aharonov, Y., Popescu, S., Rohrlich, D. & Skrzypczyk, P. Quantum Cheshire Cats. *New Journal of Physics* **15**, 113015 (2013). quant-ph/1202.0631.
- [156] Itano, M. Perspectives on the Quantum Zeno paradox. *Journal of Physics: Conference Series* **196**, 012018 (2009). URL <http://stacks.iop.org/1742-6596/196/i=1/a=012018>.
- [157] Dieks, D. & van Dijk, V. Another look at the quantum mechanical entropy of mixing. *American Journal of Physics* **56** (1988).
- [158] Rae, A. & Forgan, T. On the implications of the Quantum-Pigeonhole effect. *ArXiv e-prints* (2014). quant-ph/1412.1333.

- [159] Fine, A. The Einstein-Podolsky-Rosen Argument in Quantum Theory. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2016), fall 2016 edn.
- [160] Schrödinger, E. Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society* **31**, 555563 (1935).
- [161] Pusey, M. F. & Leifer, M. S. Logical pre- and post-selection paradoxes are proofs of contextuality. *Electron. Proc. Theor. Comput. Sci.* **195**, 295 (2015). 1506.07850.
- [162] Gu, Y.-Q. Some paradoxes in special relativity and the resolutions. *Advances in Applied Clifford Algebras* **21**, 103–119 (2011). URL <https://doi.org/10.1007/s00006-010-0244-6>.
- [163] Taylor, E. & Wheeler, J. *Spacetime Physics* (W. H. Freeman, 1992). URL https://books.google.co.uk/books?id=PDA8YcvMc_QC.
- [164] Hogg, D. W. Special relativity. *Lecture Notes* (1997).
- [165] Kent, A. Unconditionally secure bit commitment with flying qudits. *New Journal of Physics* **13**, 113015 (2011). URL <http://stacks.iop.org/1367-2630/13/i=11/a=113015>.
- [166] Kent, A. Location-oblivious data transfer with flying entangled qudits. *Physical Review A* **84**, 012328 (2011).
- [167] Buhrman, H. *et al.* *Position-Based Quantum Cryptography: Impossibility and Constructions*, 429–446 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011). URL https://doi.org/10.1007/978-3-642-22792-9_24.
- [168] Kent, A., Munro, W. J. & Spiller, T. P. Quantum tagging: Authenticating location via quantum information and relativistic signaling constraints. *Physical Review A* **84**, 012326 (2011).

- [169] Kent, A. Quantum tagging for tags containing secret classical data. *Physical Review A* **84**, 022335 (2011).
- [170] Colbeck, R. Impossibility of secure two-party classical computation. *Physical Review A* **76**, 062308 (2007). 0708.2843.
- [171] Buhrman, H. *et al.* Position-Based Quantum Cryptography: Impossibility and Constructions. *ArXiv e-prints* (2010). 1009.2490.
- [172] Kent, A. Unconditionally Secure Bit Commitment. *Physical Review Letters* **83**, 1447–1450 (1999). quant-ph/quant-ph/9810068.
- [173] Kent, A. Unconditionally Secure Bit Commitment by Transmitting Measurement Outcomes. *Physical Review Letters* **109**, 130501 (2012). quant-ph/1108.2879.
- [174] Blum, M. Coin flipping by telephone. In Gersho, A. (ed.) *CRYPTO U. C. Santa Barbara, Dept. of Elec. and Computer Engineering, ECE Report No 82-04* (1981).
- [175] Damgard, I. Commitment schemes and zero-knowledge protocols. In Damgard, I. (ed.) *Lectures on Data Security*, vol. 1561 of *Lecture Notes in Computer Science*, 63–86 (Springer Berlin Heidelberg, 1999). URL http://dx.doi.org/10.1007/3-540-48969-X_3.
- [176] Molina-Terriza, G., Vaziri, A., Ursin, R. & Zeilinger, A. Experimental quantum coin tossing. *Phys. Rev. Lett.* **94**, 040501 (2005). URL <http://link.aps.org/doi/10.1103/PhysRevLett.94.040501>.
- [177] Broadbent, A. & Tapp, A. *Information-Theoretic Security Without an Honest Majority*, 410–426 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007). URL https://doi.org/10.1007/978-3-540-76900-2_25.
- [178] Kobayashi, H. *General Properties of Quantum Zero-Knowledge Proofs*, 107–124 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2008). URL http://dx.doi.org/10.1007/978-3-540-78524-8_7.

- [179] Erven, C. *et al.* An experimental implementation of oblivious transfer in the noisy storage model. *Nature Communications* **5**, 3418 (2014). 1308.5098.
- [180] Li, Y.-B., Wen, Q.-Y., Qin, S.-J., Guo, F.-Z. & Sun, Y. Practical quantum all-or-nothing oblivious transfer protocol. *Quantum Information Processing* **13**, 131–139 (2014). URL <http://dx.doi.org/10.1007/s11128-013-0550-8>.
- [181] Lindell, Y. & Pinkas, B. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Naor, M. (ed.) *Advances in Cryptology - EUROCRYPT 2007: 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007. Proceedings*, 52–78 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007). URL http://dx.doi.org/10.1007/978-3-540-72540-4_4.
- [182] Kitaev, A., Mayers, D. & Preskill, J. Superselection rules and quantum protocols. *Physical Review A* **69**, 052326 (2004).
- [183] D'Ariano, G. M., Kretschmann, D., Schlingemann, D. & Werner, R. F. Re-examination of quantum bit commitment: The possible and the impossible. *Physical Review A* **76**, 032328 (2007).
- [184] Ben-Or, M., Goldwasser, S., Kilian, J. & Wigderson, A. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing, STOC '88*, 113–131 (ACM, New York, NY, USA, 1988). URL <http://doi.acm.org/10.1145/62212.62223>.
- [185] Croke, S. & Kent, A. Security details for bit commitment by transmitting measurement outcomes. *Physical Review A* **86**, 052309 (2012).
- [186] Lunghi, T. *et al.* Experimental Bit Commitment Based on Quantum Communication and Special Relativity. *Phys. Rev. Lett.* **111**, 180504 (2013). URL <http://link.aps.org/doi/10.1103/PhysRevLett.111.180504>.

- [187] Kent, A. Unconditionally secure bit commitment with flying qudits. *New Journal of Physics* **13**, 113015 (2011). [quant-ph/1101.4620](#).
- [188] Kent, A. Unconditionally Secure Bit Commitment by Transmitting Measurement Outcomes. *Physical Review Letters* **109**, 130501 (2012). [1108.2879](#).
- [189] Mayers, D. & Yao, A. Quantum cryptography with imperfect apparatus. *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS-98)* 503–509 (1998).
- [190] Barrett, J., Hardy, L. & Kent, A. No signaling and quantum key distribution. *Physical review letters* **95**, 010503 (2005).
- [191] Acín, A., Gisin, N. & Masanes, L. From Bells theorem to secure quantum key distribution. *Physical review letters* **97**, 120405 (2006).
- [192] Hänggi, E., Renner, R. & Wolf, S. Quantum Cryptography Based Solely on Bell’s Theorem. *EUROCRYPT 2010* 216–234 (2010). [quant-ph/0911.4171](#).
- [193] Hänggi, E. & Renner, R. Device-Independent Quantum Key Distribution with Commuting Measurements. *ArXiv e-prints* (2010). [quant-ph/1009.1833](#).
- [194] Masanes, L., Pironio, S. & Acín, A. Secure device-independent quantum key distribution with causally independent measurement devices. *Nature Communications* **2**, 238 (2011). [quant-ph/1009.1567](#).
- [195] Barrett, J., Kent, A. & Pironio, S. Maximally Nonlocal and Monogamous Quantum Correlations. *Physical Review Letters* **97**, 170409 (2006). [quant-ph/quant-ph/0605182](#).
- [196] Barrett, J., Colbeck, R. & Kent, A. Memory Attacks on Device-Independent Quantum Cryptography. *Physical Review Letters* **110**, 010503 (2013). [quant-ph/1201.4407](#).

- [197] Pironio, S. *et al.* Random numbers certified by Bell's theorem. *Nature* **464**, 1021–1024 (2010). [quant-ph/0911.3427](https://arxiv.org/abs/quant-ph/0911.3427).
- [198] Barrett, J., Colbeck, R. & Kent, A. Unconditionally secure device-independent quantum key distribution with only two devices. *Physical Review A* **86**, 062326 (2012). [quant-ph/1209.0435](https://arxiv.org/abs/quant-ph/1209.0435).
- [199] Fouard, L., Duclos, M. & Lafourcade, P. Survey on electronic voting schemes. *Supported by the ANR project AVOTE* (2010).
- [200] Nguyen, K., Bao, F., Mu, Y. & Varadharajan, V. Zero-knowledge proofs of possession of digital signatures and its applications. In Varadharajan, V. & Mu, Y. (eds.) *Information and Communication Security*, vol. 1726 of *Lecture Notes in Computer Science*, 103–118 (Springer Berlin Heidelberg, 1999). URL http://dx.doi.org/10.1007/978-3-540-47942-0_9.
- [201] Fortnow, L. The complexity of perfect zero-knowledge. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing, STOC '87*, 204–209 (ACM, New York, NY, USA, 1987). URL <http://doi.acm.org/10.1145/28395.28418>.
- [202] Feige, U., Fiat, A. & Shamir, A. Zero-knowledge proofs of identity. *Journal of Cryptology* **1**, 77–94 (1988). URL <http://dx.doi.org/10.1007/BF02351717>.
- [203] Horodecki, P., Horodecki, M. & Horodecki, R. Zero knowledge convincing protocol on quantum bit is impossible. *eprint arXiv:quant-ph/0010048* (2000). [quant-ph/quant-ph/0010048](https://arxiv.org/abs/quant-ph/0010048).
- [204] Bell, J. *Speakable and unspeakable in quantum mechanics* (Cambridge University Press, 2004).
- [205] Schwartz, B. *The Paradox of Choice: Why More is Less* (Ecco/HarperCollins Publishers, 2004).

- [206] Timpson, C. G. & Brown, H. R. Proper and Improper Separability. *International Journal of Quantum Information* **3**, 679 – 690 (2005). [quant-ph/quant-ph/0402094](#).
- [207] Horodecki, R., Horodecki, P., Horodecki, M. & Horodecki, K. Quantum entanglement. *Reviews of Modern Physics* **81**, 865–942 (2009). [quant-ph/quant-ph/0702225](#).
- [208] Pusey, M. F., Barrett, J. & Rudolph, T. On the reality of the quantum state. *Nature Physics* **8**, 476–479 (2012). [quant-ph/1111.3328](#).
- [209] Leifer, M. Is the quantum state real? An extended review of ψ -ontology theorems. *Quanta* **3**, 67 – 155 (2014). URL <http://quanta.ws/ojs/index.php/quanta/article/view/22>.
- [210] Hardy, L. Are quantum states real? *International Journal of Modern Physics B* **27**, 45012 (2013). [quant-ph/1205.1439](#).
- [211] Colbeck, R. & Renner, R. A system’s wave function is uniquely determined by its underlying physical state. *New Journal of Physics* **19**, 013016 (2017). URL <http://stacks.iop.org/1367-2630/19/i=1/a=013016>.
- [212] Maroney, O. J. E. How statistical are quantum states? *ArXiv e-prints* (2012). [quant-ph/1207.6906](#).
- [213] Barrett, J., Cavalcanti, E. G., Lal, R. & Maroney, O. J. E. No ψ -Epistemic Model Can Fully Explain the Indistinguishability of Quantum States. *Physical Review Letters* **112**, 250403 (2014). [quant-ph/1310.8302](#).
- [214] Montina, A. Exponential complexity and ontological theories of quantum mechanics. *Phys. Rev. A* **77**, 022104 (2008). URL <http://link.aps.org/doi/10.1103/PhysRevA.77.022104>.
- [215] Fedrizzi, A. *et al.* High-fidelity transmission of entanglement over a high-loss free-space channel. *Nature Physics* **5**, 389–392 (2009). [quant-ph/0902.2015](#).

- [216] Kent, A., Massar, S. & Silman, J. Secure and Robust Transmission and Verification of Unknown Quantum States in Minkowski Space. *Scientific Reports* **4**, 3901 (2014). [quant-ph/1208.0745](https://arxiv.org/abs/quant-ph/1208.0745).
- [217] D'Hondt, E. *Distributed quantum computation: a measurement-based approach*. Ph.D. thesis, Vrije Universiteit Brussel (2005).
- [218] Beals, R. *et al.* Efficient distributed quantum computing. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **469**, 20120686 (2013).
- [219] Runser, R. J. *et al.* Progress toward quantum communications networks: opportunities and challenges. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 6476 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 0 (2007).
- [220] Boone, K. *et al.* Entanglement over global distances via quantum repeaters with satellite links. *Phys. Rev. A* **91**, 052325 (2015). URL <http://link.aps.org/doi/10.1103/PhysRevA.91.052325>.
- [221] Jozsa, R. An introduction to measurement based quantum computation. *eprint arXiv:quant-ph/0508124* (2005). [quant-ph/quant-ph/0508124](https://arxiv.org/abs/quant-ph/0508124).
- [222] Hayashi, M. *Quantum Information: An Introduction* (Springer Berlin Heidelberg, 2006). URL <https://books.google.co.uk/books?id=UjDL15sP7vEC>.
- [223] Kent, A. Impossibility of unconditionally secure commitment of a certified classical bit. *Physical Review A* **61**, 042301 (2000). [quant-ph/quant-ph/9910087](https://arxiv.org/abs/quant-ph/quant-ph/9910087).
- [224] Kent, A. Why classical certification is impossible in a quantum world. *Quantum Information Processing* **11**, 493–499 (2012). URL <https://doi.org/10.1007/s11128-011-0262-x>.

- [225] Bruß, D. & Macchiavello, C. Optimal state estimation for d-dimensional quantum systems. *Physics Letters A* **253**, 249–251 (1999). [quant-ph/9812016](https://arxiv.org/abs/quant-ph/9812016).
- [226] Sahai, A. & Vadhan, S. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)* **50**, 196–249 (2003).
- [227] Atkinson, K. & Han, W. *Theoretical Numerical Analysis: A Functional Analysis Framework*. Texts in Applied Mathematics (Springer New York, 2007). URL <https://books.google.co.uk/books?id=8vpLPun8o-wC>.
- [228] Kent, A. A no-summoning theorem in relativistic quantum theory. *Quantum Information Processing* **12**, 1023–1032 (2013). [quant-ph/1101.4612](https://arxiv.org/abs/quant-ph/1101.4612).
- [229] Rudolph, T. The Laws of Physics and Cryptographic Security. *eprint arXiv:quant-ph/0202143* (2002). [quant-ph/0202143](https://arxiv.org/abs/quant-ph/0202143).
- [230] Rudra, A. Limits to list decoding random codes. In Ngo, H. (ed.) *Computing and Combinatorics*, vol. 5609 of *Lecture Notes in Computer Science*, 27–36 (Springer Berlin Heidelberg, 2009). URL http://dx.doi.org/10.1007/978-3-642-02882-3_4.
- [231] Colbeck, R. & Renner, R. Free randomness can be amplified. *Nature Physics* **8**, 450–454 (2012). [1105.3195](https://arxiv.org/abs/1105.3195).
- [232] Barnett, S. M., Chefles, A. & Jex, I. Comparison of two unknown pure quantum states. *Physics Letters A* **307**, 189–195 (2003). [quant-ph/0202087](https://arxiv.org/abs/quant-ph/0202087).
- [233] Sedlák, M., Ziman, M., Bužek, V. & Hillery, M. Unambiguous comparison of ensembles of quantum states. *Physical Review A* **77**, 042304 (2008). [0712.1616](https://arxiv.org/abs/0712.1616).
- [234] Ichikawa, J. J. & Steup, M. The analysis of knowledge. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2016), winter 2016 edn.

- [235] Vilasini, V., Portmann, C. & del Rio, L. Composable security in relativistic quantum cryptography. *ArXiv e-prints* (2017). 1708.00433.
- [236] Aaronson, S. Quantum Copy-Protection and Quantum Money. *Proceedings of IEEE Conference on Computational Complexity* 229–242 (201109). 1110.5353.
- [237] Buchmann, J. *Introduction to Cryptography*. Undergraduate Texts in Mathematics (Springer New York, 2013). URL <https://books.google.co.uk/books?id=BuQ1BQAAQBAJ>.
- [238] Goldreich, O., Ostrovsky, R. & Petrank, E. Computational complexity and knowledge complexity. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, 534–543 (ACM, 1994).
- [239] Renner, R. Security of quantum key distribution. *International Journal of Quantum Information* **6**, 1–127 (2008).
- [240] Chakravartty, A. Scientific realism. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2014), spring 2014 edn.
- [241] Sankey, H. Scientific realism: An elaboration and a defence. *Theoria: A Journal of Social and Political Theory* 35–54 (2001). URL <http://www.jstor.org/stable/41802172>.
- [242] Quisquater, J.-J., Guillou, L., Annick, M. & Berson, T. How to explain zero-knowledge protocols to your children. In *Proceedings on Advances in Cryptology, CRYPTO '89*, 628–631 (Springer-Verlag New York, Inc., New York, NY, USA, 1989). URL <http://dl.acm.org/citation.cfm?id=118209.118269>.
- [243] Pitowsky, I. *Quantum Mechanics as a Theory of Probability*, 213–240 (Springer Netherlands, Dordrecht, 2006). URL https://doi.org/10.1007/1-4020-4876-9_10.

- [244] Rosen, G. Abstract objects. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University, 2017), spring 2017 edn.
- [245] Lunghi, T. *et al.* Practical Relativistic Bit Commitment. *Physical Review Letters* **115**, 030502 (2015). 1411.4917.
- [246] Kaniewski, J. *Relativistic quantum cryptography*. Ph.D. thesis (2015). quant-ph/1512.00602.
- [247] Rudich, S. unpublished .
- [248] Naor, M. & Yung, M. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing, STOC '90*, 427–437 (ACM, New York, NY, USA, 1990). URL <http://doi.acm.org/10.1145/100216.100273>.
- [249] Bellare, M. & Goldwasser, S. *New Paradigms for Digital Signatures and Message Authentication Based on Non-Interactive Zero Knowledge Proofs*, 194–211 (Springer New York, New York, NY, 1990). URL https://doi.org/10.1007/0-387-34805-0_19.
- [250] Scarani, V., Iblisdir, S., Gisin, N. & Acín, A. Quantum cloning. *Reviews of Modern Physics* **77**, 1225–1256 (2005). quant-ph/0511088.