

PAPER • OPEN ACCESS

Valuation of a financial claim contingent on the outcome of a quantum measurement

To cite this article: Lane P Hughston and Leandro Sánchez-Betancourt 2024 *J. Phys. A: Math. Theor.* **57** 285302

View the [article online](#) for updates and enhancements.

You may also like

- [Production of Crystal Oriented Neodymium Silicate Oxypatite Ceramics and Their Conductivity](#)

Yuichi Tsunoda, Kiyoshi Kobayashi, Tohru Higuchi et al.

- [Determinants of corporate dividend policy in Indonesia](#)

H S Lestari

- [Anodic Polarization Behavior of Iron in Phosphoric Acid Solution Under the Action of Magnetic Field](#)

Yuanjie Tang, Shuangwei Cai, Fei Ning et al.

Valuation of a financial claim contingent on the outcome of a quantum measurement

Lane P Hughston^{1,*} 
and Leandro Sánchez-Betancourt^{2,3,*} 

¹ Department of Computing, Goldsmiths University of London, New Cross, London SE14 6NW, United Kingdom

² Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, United Kingdom

³ Oxford-Man Institute of Quantitative Finance, Walton Well Road, Oxford OX2 6ED, United Kingdom

E-mail: L.Hughston@gold.ac.uk and leandro.sbetancourt@gmail.com

Received 14 November 2023; revised 1 May 2024

Accepted for publication 16 May 2024

Published 2 July 2024



CrossMark

Abstract

We consider a rational agent who at time 0 enters into a financial contract for which the payout is determined by a quantum measurement at some time $T > 0$. The state of the quantum system is given in the Heisenberg representation by a known density matrix $\hat{\rho}$. How much will the agent be willing to pay at time 0 to enter into such a contract? In the case of a finite dimensional Hilbert space \mathcal{H} , each such claim is represented by an observable \hat{X}_T where the eigenvalues of \hat{X}_T determine the amount paid if the corresponding outcome is obtained in the measurement. We prove, under reasonable axioms, that there exists a pricing state \hat{q} which is equivalent to the physical state $\hat{\rho}$ such that the pricing function Π_{0T} takes the linear form $\Pi_{0T}(\hat{X}_T) = P_{0T} \text{tr}(\hat{q}\hat{X}_T)$ for any claim \hat{X}_T , where P_{0T} is the one-period discount factor. By ‘equivalent’ we mean that $\hat{\rho}$ and \hat{q} share the same null space: that is, for any $|\xi\rangle \in \mathcal{H}$ one has $\hat{\rho}|\xi\rangle = 0$ if and only if $\hat{q}|\xi\rangle = 0$. We introduce a class of optimization problems and solve for the optimal contract payout structure for a claim based on a given measurement. Then we consider the implications of the Kochen–Specker theorem in

* Authors to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

this setting and we look at the problem of forming portfolios of such contracts. Finally, we consider multi-period contracts.

Keywords: quantum mechanics, quantum measurement, contingent claims, absence of arbitrage, density matrices, Gleason's theorem, Kochen-Specker theorem

1. Introduction

By 'quantum finance' we mean the valuation, optimization and risk management of financial contracts for which the outcomes (in the form of one or more payments made between the various parties involved) are contingent on the results of one or more quantum measurements. The financial contracts that we consider can be easily implemented in a suitable laboratory. Our investigations fall within the scope of standard quantum mechanics and we are not concerned here with modifications of the standard framework or with interpretive issues. The resulting theory of quantum financial contracts is distinctly non-Kolmogorovian, inheriting as it does the full generality of quantum probability.

The idea of forging connections between quantum theory and finance theory is not a new one. Previous attempts have tended to fall into two broad categories. In the first category one has theories that work with the suggestion that asset prices—and perhaps other economic variables as well—are somehow subject to the laws of quantum mechanics. Examples can be found in [3, 5, 17, 19, 25, 31, 32, 42, 49, 57], to mention but a few. The idea is not an unreasonable thing to think about, given the rather general notions of 'complementarity' promoted by Bohr and Heisenberg. Nonetheless, it is probably safe to say that little by way of real progress has been made. The problem with this line of thinking is that there is no evidence to suggest that asset prices are in any sense 'quantum-like' in nature, so work along these lines is speculative. Even if one were to admit the idea that an asset price (say, the price of a barrel of oil) is akin to a physical variable, such as position or energy, that can be quantized, it is not at all evident what form the associated complementary variable would take—and where would Planck's constant enter the discussion?

In the second category of connections between quantum theory and finance, less controversially, one sees mathematical techniques originating in quantum theory being applied to problems in finance. Recently, we have witnessed a flurry of activity in the use of such techniques to improve on the traditional methods of computational finance, paving the way for the use of quantum computers to perform large-scale financial calculations. There is an extensive literature involving applications of quantum methods and quantum computation to the solution of 'classical' problems in finance; we mention [2–4, 26, 34, 39, 55, 58–60] as representative of the multiplicity of ideas being pursued. It should be emphasized, nonetheless, that no quantum ideas *per se* are involved in the finance theory underlying such endeavours, so the term 'quantum finance' may be a misnomer in that context.

Our theory represents a departure from these approaches. On the one hand, we are not suggesting the existence of quantum properties in ordinary financial assets such as stocks and bonds. Nor are we concerned here with the use of ideas derived from quantum theory to speed up the risk management of conventional financial assets, or even new assets, such as those based on blockchain technology. In particular, we are not concerned here with the use of quantum computers to tackle the work currently being carried out by classical computers.

We are concerned, rather, with the pricing of financial instruments (equivalently, 'securities' or 'financial products') for which the payouts are directly linked, by design, to the outcomes of quantum measurements. Needless to say, such financial products do not exist at present.

But they may exist in the future, and that is why the theory is of interest from a scientific perspective.

Let us recall a few basic ideas from the theory of finance. Here we proceed on an informal basis, in anticipation of the more rigorous treatment given later in the paper in the context of quantum theory. By a financial instrument, in classical finance theory, we mean a contract whereby two or more agents agree to exchange cash flows over a period of time in accordance with certain rules. The ‘agents’ might be individuals, small businesses, corporations, municipalities, sovereign states, and so on—any entity that has the capacity for entering into a financial contract. By a ‘cash flow’ we mean an event where one agent transfers some money (‘cash’) to another agent. The mechanism of transfer is not relevant to the present discussion, but we assume that it is essentially instantaneous and that the cash will be transferred from an account of one entity to an account of another.

In a financial contract both the amounts and the timings of the cash flows will in general be ‘contingent’ – that is to say, not fixed in advance, but determined by one or more external events. For example, if company A purchases a bond from company B , then there is an initial cash flow when A pays B the price of the bond. This will be followed by several intermediate cash flows, where B makes periodic ‘coupon payments’ to A —these interest payments are usually a fixed percentage of the ‘principal’ of the bond—thus, we might refer, say, to a 4.5% coupon paid semiannually at fixed dates. Then finally, B pays back the agreed ‘principal’ to A on some pre-agreed date, and the contract is concluded. All bond issuers are credit risky to some degree, and hence B may for some reason fail to pay a coupon on the date on which it is due or B may fail to pay all or part of the principal when this becomes due. Further details of the contract spell out how the various parties will proceed in such situations and how the subsequent cash flows will be structured.

Another common situation is one where A decides to sell the bond to some third agent C before the term of the bond has finished. In that case, there is a cash flow from the account of C to the account of A . Agent B is not involved in that transaction, but the original agreement between A and B is then transferred to an arrangement between C and B . Again, the details of how this happens are set out in the original contract. Such arrangements are usually enforceable in law and are governed by various regulations. The example of the ordinary coupon bond as a financial instrument illustrates the fact that cash flows typically involve both deterministic elements and uncertain elements—the latter including the timings of any defaults, and the timings of any transfers of ownership of the bond before its term is concluded.

Similarly, if an individual A buys stock issued by company B , then the stock goes into A ’s securities account and A ’s cash goes into the account of the broker from whom A bought the stock, and thence to the account of the seller of the stock. Needless to say, the registers for stock going into someone’s securities account or cash going into someone’s cash account are purely electronic, even though as an aid to thinking we use the nostalgic language of stock certificates or banknotes moving from one vault to another. While A holds the stock, he may receive dividend payments, both the timing and the amounts of which will in general be uncertain. Then when A chooses to sell the stock there will be a cash flow into his account for the amount of the sale, and the stock will leave his securities account.

For each type of financial instrument, the contract and the associated system of cash flows will vary. In the case of a cash-settled European-style call option, for example, agent A buys the option from an options dealer B for an amount C_0 (the ‘option premium’). The option contract gives A the right (but not the obligation) to buy the stock at some pre-designated time T for the so-called strike price K . Since the contract is cash-settled, this means that at T agent A will receive a cash flow of $H_T = \max(S_T - K, 0)$, where S_T is the terminal stock price. This assumes implicitly that A acts rationally, and only ‘exercises’ the option if $S_T > K$. In this

example, there are two cash flows—the deterministic amount C_0 paid by A to B at time 0, and the amount H_T paid by B to A at the maturity of the option. We refer to H_T as the ‘payout’ of the option. In more detail, we can assume (say) that initially A has an empty cash account, so to purchase the option first he borrows the amount C_0 from some third agent C , promising to pay agent C back the amount C_0 at T plus an agreed amount of interest. If the interest is charged on a continuously compounded basis at the rate r over the period from time 0 to time T then the total paid back to agent C for the loan of C_0 is $C_0 e^{rT}$. Then the total profitability of the option transaction as measured at T would be $\max(S_T - K, 0) - C_0 e^{rT}$, where negative profitability implies agent A is left in debt.

Rather than working with interest rates, finance theorists often prefer to work with so-called discount bonds. By a T -maturity discount bond, one means a contract between two agents C and D where C loans D a certain amount of money P_{0T} at time 0 in exchange for a positive cash flow of one unit of cash at time T . Thus the loan is made on a discounted basis and the principal repaid at T is unity. For example, C might loan D \$0.94 today in exchange for a repayment of \$1.00 a year from now. If the interest rate is r per annum continuously compounded, then $P_{0T} = e^{-rT}$. The finance theorist will tend to regard the bond price P_{0T} as the more fundamental object, since bonds trade in bond markets, and the interest rate is a derived concept. Then given the bond price we can *define* the interest rate by setting $r = -T^{-1} \ln P_{0T}$. So in principle one can work entirely with prices, never mentioning interest rates. In practice, it is often convenient to work with rates and one switches back and forth between rates and prices when bond markets are being discussed.

This leads us to the rather general idea of a financial contract in accordance with which at time 0 agent A pays an amount H_0 to agent B and then agent B pays a random variable H_T to agent A at time T . Then we say that agent A is purchasing a ‘financial product’ from agent B . The option contract mentioned above is an example of such a product, where H_0 is the option premium and H_T is the option payout. In that case, we model the terminal stock price S_T is a random variable, and then define the payout as a function of this random variable. Another example would be a binary (or digital) option contract, in accordance with which at time T agent B pays agent A one unit of account (say, one dollar) if $S_T > K$ and zero otherwise. Many other examples of such products can be structured where the ‘underlying’ asset prices might be share prices, bond prices, foreign exchange rates, commodity prices, and so on.

The events determining the cash flows do not necessarily have to be market events as such. Physical events are also possible as the basis for determining a cash flow. For example, in insurance markets cash flows can be triggered by hazardous physicalities, such as fires, floods, tornados, hurricanes, earthquakes, and so on. In exchange for this random payout, the holder of the policy pays an ‘insurance premium’ at the beginning of the contract. The insurers then limit their own potential exposures via reinsurance markets.

The mathematical theory of how one prices financial instruments has an interesting history, with the first definitive results being obtained in [11], followed in subsequent decades by many significant further developments, with numerous practical applications. The important point, however, is that the terminal payout H_T of such an instrument is modelled in the standard theory as a random variable on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a set that represents the possible outcomes of chance, \mathcal{F} denotes the collection of subsets of Ω that can be assigned probabilities (so-called ‘events’), and \mathbb{P} denotes the system of probabilities assigned to these events. Let us write $\mathbb{E}[-]$ for expectation under \mathbb{P} (the usual ‘objective’ probability measure). The theory then shows that in order for the market to be arbitrage-free there must exist a positive random variable π_T on Ω satisfying $0 < \mathbb{E}[\pi_T] < \infty$ such that for any contract with

payout H_T satisfying $\mathbb{E}[\pi_T H_T] < \infty$ the price takes the form

$$H_0 = \mathbb{E}[\pi_T H_T]. \tag{1}$$

Thus, the absence of arbitrage ensures the existence of a universal ‘pricing kernel’ π_T that can be used to price all financial products with payouts that are known at T . An alternative way of putting this is to set $Q_T = \pi_T / \mathbb{E}[\pi_T]$ and write

$$H_0 = P_{0T} \mathbb{E}^{\mathbb{Q}}[H_T]. \tag{2}$$

Here we set $P_{0T} = \mathbb{E}[\pi_T]$ and for any positive random variable X_T whose outcome is known at T we define $\mathbb{E}^{\mathbb{Q}}[X_T] = \mathbb{E}[Q_T X_T]$. Then $\mathbb{E}^{\mathbb{Q}}[X_T]$ is the expectation of X_T with respect to a probability measure \mathbb{Q} that is ‘equivalent’ to \mathbb{P} in the sense that for any event $A \in \mathcal{F}$ it holds that $\mathbb{P}(A) = 0$ if and only if $\mathbb{Q}(A) = 0$. One says that Q_T is the Radon-Nikodym derivative of \mathbb{Q} with respect to \mathbb{P} . Thus, we see that the price H_0 is the discounted expectation of the payout H_T where the expectation is taken with respect to the measure \mathbb{Q} .

In short, a market is arbitrage-free if and only there exists a probability measure \mathbb{Q} which is equivalent to \mathbb{P} on events of probability zero such that for any time T and for any claim H_T the price of that claim is given by (2). Of course, we have skipped many details and we have made no attempt to define ‘arbitrage’. Nonetheless, the key point that emerges here is that the pricing formula involves a system of pricing probabilities, which we call \mathbb{Q} , that is distinct from the usual ‘objective’ or ‘real world’ probability system \mathbb{P} . This ‘change of measure’ idea in the determination of prices is one of the pillars of modern finance theory. We mention this since many physicists will assume (perhaps through some well-intended but misguided notion of ‘fairness’) that the price H_0 of an investment with payout H_T must surely be just the mean value of H_T over \mathbb{P} —but this is not correct! It is not surprising then that a similar ‘change of state’ idea arises when quantum theory is brought into play.

In the present paper we are concerned with the situation where agent A pays agent B a premium H_0 at time 0 in exchange for a cash flow of the amount H_T at time T , but now where H_T is determined by the outcome of a quantum measurement. The problem is that the outcomes of quantum measurements cannot be consistently modelled as random variables in a Kolmogorovian framework of the type outlined above on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ in the case of classical finance theory. So we need a more general framework for finance based on quantum probability in order to model the situations that we have in mind.

We have mentioned the case of a single payout H_T determined by the outcome of a single measurement. More generally, one can envisage the existence of a rather general type of financial product that delivers a sequence of cash flows determined by the outcomes of a sequence of experiments. In such contracts, the experiments themselves may be ‘adapted’, in the sense that experiments performed at a later stage of the sequence may depend in their design and implementation on the outcomes of experiments made at an earlier stage of the sequence. In this paper, we consider the most basic of such financial products.

In section 2 we set out what we mean by a financial contract in a one-period market for which the payout is determined by a quantum measurement. We show that such a contract can be represented by a quantum observable (a Hermitian operator) for which the eigenvalues represent the possible cash flows. We begin with the example of a two-dimensional Hilbert space, for which the underlying experiment triggering the outcome of the contract involves measuring the spin of a spin $\frac{1}{2}$ particle along a certain choice of axis. The state of the particle before the measurement is given by a known density matrix. The contract specifies the payments made for the two possible outcomes. The totality of these contracts constitute the ‘market’ associated with such spin measurements.

In section 3 we extend the discussion to the case of an n -dimensional Hilbert space \mathcal{H} and we introduce the notion of a one-period discount bond, which pays out one unit of account at time T regardless of the outcome of the experiment. In that case the associated financial observable is the identity operator on \mathcal{H} .

We also introduce the notion of a so-called Arrow–Debreu security, as it arises in the present context, for which the underlying experiment takes the form of a projection operator of rank unity. Such a contract either pays one unit of account or nothing, depending on which eigenvalue of the projection operator is attained when the measurement is performed.

The physical state of the underlying quantum system on which the measurement is taken is represented by a density matrix $\hat{\rho}$, an n by n positive semidefinite matrix with trace unity. We say that two density matrices $\hat{\rho}$ and \hat{q} are *equivalent* if for all $|\xi\rangle \in \mathcal{H}$ it holds that $\hat{\rho}|\xi\rangle = 0$ if and only if $\hat{q}|\xi\rangle = 0$, i.e. they share the same ‘null space’. The significance of this equivalence relation among density matrices becomes apparent later when we consider the pricing of contracts. We argue that if two contracts \hat{U}_T and \hat{V}_T depend on the outcome of *the same experiment*, and differ from one another only in the amounts paid for the various outcomes of the experiment, then the prices of these contracts should satisfy a linear relation

$$\Pi_{0T} [a\hat{U}_T + b\hat{V}_T] = a\Pi_{0T} [\hat{U}_T] + b\Pi_{0T} [\hat{V}_T], \tag{3}$$

where Π_{0T} denotes the pricing map. More precisely, if the operators \hat{U}_T and \hat{V}_T *commute*, then the prices should be additive.

In section 4 we present our main result, which is to show that under a certain set of axioms the pricing map necessarily takes the form

$$\Pi_{0T} [\hat{X}_T] = P_{0T} \text{tr} (\hat{q}\hat{X}_T), \tag{4}$$

for some density matrix \hat{q} , where tr denotes the trace, subject only to the condition that the physical state $\hat{\rho}$ and the pricing state \hat{q} are equivalent in the sense mentioned above. The axioms are surprisingly simple. The first is that the price of a non-negative contract should vanish if and only if the expectation value of the contract vanishes. The second is that the pricing function should act linearly on any set of mutually commuting contracts. And the third is that the price of the observable corresponding to the identity operator should be that of a unit discount bond, which we regard as an input to the model. It should be emphasized that (a) we do not assume that the pricing map is linear, and that (b) no portfolio arguments are involved. The key ingredient in the proof is Gleason’s theorem, which turns out to be surprisingly well adapted for applications in a financial context.

In sections 5 and 6 we look at a classical investment problem in a quantum context. The problem is that faced by an investor with a fixed budget who wishes to invest optimally in order to maximize the expected utility of the outcome of the investment. In the classical theory, we define a von Neumann–Morgenstern utility function as a concave, strictly increasing function of a strictly positive argument. The argument (in the present context) is the amount of cash received at some designated time T by some agent. The utility represents the degree of satisfaction received by the agent from that cash flow. The underlying concept of ‘utility’ (like ‘goodness’) is ultimately an undefinable that is characterized by its stated properties and its behaviour in various examples. Nonetheless, the notion of utility is highly intuitive and absolutely essential to many arguments in financial economics. The goal of an investment problem is to determine, for a given budget, the optimal investment over various investment opportunities, where the criterion for optimization is to maximize the expected utility of the payout over some designated time frame. We show that the quantum optimization problem, like the classical problem, can be solved exactly; and indeed, although the mathematical ideas run in

many respects in parallel in the two theories, it is not *a priori* obvious what form the solution of the quantum problem will take.

In particular, if the quantum investment problem is generalized in such a way as to involve a choice between several incompatible experiments to determine the payout of the contract, then the problem cannot even be formulated in the classical theory of finance, and yet admits a neat formulation and solution in the case of quantum securities as we have defined them. This is shown in section 7, where we discuss more generally the role of quantum probability in finance. The example we consider is based on a construction of the Kochen–Specker type due to Cabello *et al* [15, 16] involving a collection of nine incompatible observables. In this way we can formulate a quantum optimization problem, with no classical analogue, in which the investor faces a choice between nine different quantum financial contracts.

In section 8 we return to the problem of portfolios, which we treat as a class of financial products for which the payouts depend on the results of two or more experiments. In the case of a one-period market, we consider the situation where one carries out measurements simultaneously on a pair of particles. The particles are associated with distinct Hilbert spaces, so no incompatibilities arise between the measurements and results are obtained for each. As a consequence, financial contracts can be devised for which the payouts are made by totalling the results of each of the experiments. The two contracts can then consistently be regarded as part of the same portfolio in such a setup. The density matrix of the two-particle system as a whole can be entangled, allowing for correlations between the outputs of the individual constituents. A similar situation arises for portfolios involving any number of constituents. The surprising feature here is the relation between the ideas of entanglement in quantum mechanics and cointegration in portfolio theory. We conclude in section 9 with a brief discussion of multi-period markets.

2. Quantum measurements and contingent cash flows

Let time 0 be the present and T a fixed time in the future. We consider the situation where an agent A enters into a contract with another agent B in accordance with which A pays B an amount H_0 (‘the price’) at time 0 and then B pays A an amount H_T (the ‘payout’) at time T , where H_T is contingent in some specified way on the outcome of a quantum measurement. We refer to such a setup as a one-period market.

By a quantum measurement, we mean the measurement of an observable associated with a microscopic system, such as a particle, or an atom or a molecule. More elaborate setups can be considered, involving multiple measurements, multiple payments and multiple agents; but for simplicity we look at a one-period market involving two agents. As an example, suppose the payout is determined by a measurement of the spin of a spin one-half particle along the z -axis. The outcome of such a measurement either gives $+\frac{1}{2}\hbar$, corresponding to spin up along that axis, or $-\frac{1}{2}\hbar$, corresponding to spin down. Henceforth, we work with physical units such that $\hbar = 1$. For the basics of quantum theory, see, for instance, [38]. We fix a two-dimensional Hilbert space \mathcal{H}^2 and on it we introduce the usual observable for the spin along the z -axis, given by the Hermitian operator

$$\hat{S}_z = \frac{1}{2}|z_1\rangle\langle z_1| - \frac{1}{2}|z_2\rangle\langle z_2|, \quad (5)$$

where $|z_1\rangle$ is a unit Hilbert space vector representing the upward direction along the z -axis and $|z_2\rangle$ denotes a unit Hilbert space vector orthogonal to $|z_1\rangle$ representing the downward direction along the z -axis. Thus, $\langle z_1|z_1\rangle = 1$, $\langle z_2|z_2\rangle = 1$, $\langle z_1|z_2\rangle = 0$, $\langle z_2|z_1\rangle = 0$, and the possible outcomes of the measurement are the eigenvalues of \hat{S}_z , which are $+\frac{1}{2}$ and $-\frac{1}{2}$.

The probabilities of these outcomes are determined by the *state* of the system, which is represented by a density matrix $\hat{\rho}$. The density matrix in quantum theory has a status that is analogous in certain respects to that of the probability measure in classical probability theory. The density matrix is assumed to be a positive-semidefinite Hermitian operator with trace unity, which in the case of a two-dimensional Hilbert space takes the form

$$\hat{\rho} = p_1|\psi_1\rangle\langle\psi_1| + p_2|\psi_2\rangle\langle\psi_2|, \tag{6}$$

for some orthonormal basis $\{|\psi_1\rangle, |\psi_2\rangle\}$ in \mathcal{H}^2 , where $p_1 \geq 0, p_2 \geq 0$, and $p_1 + p_2 = 1$. In general, such a matrix will have rank two, but if $p_1 = 0$ or $p_2 = 0$ then it will have rank one. A state with rank one is called a ‘pure’ state.

The probability for a given outcome is the trace of the product of the density matrix $\hat{\rho}$ with the projection operator onto the Hilbert subspace associated to the eigenvalue corresponding to that outcome (the ‘Born rule’). Thus we have

$$\text{Prob}(S_z = \frac{1}{2}) = \text{tr}(\hat{\rho}|z_1\rangle\langle z_1|) = \langle z_1|\hat{\rho}|z_1\rangle, \tag{7}$$

$$\text{Prob}(S_z = -\frac{1}{2}) = \text{tr}(\hat{\rho}|z_2\rangle\langle z_2|) = \langle z_2|\hat{\rho}|z_2\rangle. \tag{8}$$

In the case of a contingent claim where the payout is determined by the result of such a spin measurement, it should be clear that the claim itself can also be represented by a Hermitian operator on \mathcal{H}^2 , in this case, an operator of the form

$$\hat{Z}_T = z_1|z_1\rangle\langle z_1| + z_2|z_2\rangle\langle z_2|, \tag{9}$$

where z_1 denotes the payment made to agent A in the case the measurement outcome is spin $+\frac{1}{2}$ and z_2 is the payment made to A when the measurement outcome is spin $-\frac{1}{2}$. One can think of such a contract as being an example of a so-called real option [21, 37, 62]. Payments are understood to be made in some fixed numeraire or unit of account. Thus, we conclude that *a contingent claim for which the payouts are determined by the result of a quantum measurement can be represented by an observable, in the usual quantum mechanical sense, whose eigenvalues correspond to the possible cash flows at time T .*

Among the various observables that can be represented in the form (9) there is a special observable that takes the form

$$\hat{P}_{0T} = 1|z_1\rangle\langle z_1| + 1|z_2\rangle\langle z_2|, \tag{10}$$

which pays one unit of account at time T , regardless of the outcome of the spin measurement. This is evidently a ‘risk-free’ asset, since the payout is fixed and guaranteed, and we write

$$\hat{P}_{0T} = \hat{\mathbf{1}}. \tag{11}$$

Here $\hat{\mathbf{1}}$ denotes the identity operator on \mathcal{H}^2 . The risk-free asset \hat{P}_{0T} represents a discount bond that pays one unit of account (e.g. one ‘dollar’) at maturity T . It has the property that it does not depend on the choice of axis along which the spin measurement is taken.

In addition to contracts of the form (9), we can more generally consider contracts of the same type, but where the measurement of the spin is taken along some other axis. Each such contract is characterized by (a) the choice of a basis in Hilbert space along which the spin measurement is made, together with (b) the payouts that take place as a consequence of the results of the measurement. Indeed, it is a theorem that any positive Hermitian operator \hat{Z}_T on \mathcal{H}^2 other than multiples of the identity can be expressed uniquely in the form (9) for some choice of the orthonormal basis $\{|z_1\rangle, |z_2\rangle\}$ in \mathcal{H}^2 , modulo multiplicative phase factors.

To complete the discussion we need to determine the *price* paid by agent A to agent B in exchange for the payout corresponding to \hat{X}_T . In short, we need a *pricing function* that maps each financial observable \hat{X}_T to a price X_0 , which we develop in section 4.

It should be emphasized that although we may question why an agent might wish to purchase such a security, there is nothing mysterious or obscure about the construction of such a market itself. As in all economic considerations, the issue of why there is supply and demand for a certain product is a matter quite distinct from the issue of how the market for that product will function, given that there is indeed supply and demand.

3. Financial observables

It will be useful going forward to generalize our considerations to the case of a Hilbert space \mathcal{H} of arbitrary finite dimension n . As usual, we can write $|\xi\rangle$ for a typical element of \mathcal{H} and $\langle\xi|$ for its complex conjugate. The observable that determines the payout will in the generic situation be a non-degenerate Hermitian operator \hat{X}_T on this space and hence admit n distinct real eigenvalues, each corresponding to a distinct cash flow.

For example, if the quantum system admits n different energy levels, and the underlying physical observable being measured is the energy of the system, then the contract will in general result in a different cash flow $\{x_j\}_{j=1,2,\dots,n}$ for each of the possible energy outcomes. For the financial observable representing such a contract we can write

$$\hat{X}_T = \sum_{j=1}^n x_j |x_j\rangle\langle x_j|, \quad \langle x_j | x_k \rangle = \delta_{jk} \quad (12)$$

for some orthonormal basis $\{|x_j\rangle\}_{j=1,2,\dots,n}$ in Hilbert space. More generally, the set of all financial observables associated with a given Hilbert space will include some that are degenerate in the sense that the same payout will result for two or more distinct values of the outcome j . Such a degeneracy can result either because there is a degeneracy in the spectrum of the underlying physical observable, or because two or more distinct eigenvalues of the physical observable are assigned the same cash flow. An example of the latter is a unit discount bond, for which $x_j = 1$ for all $j = 1, 2, \dots, n$ even though the underlying energy levels may be distinct. Then the identity operator on \mathcal{H} represents such a bond.

Another example of a degenerate observable is the analogue of a so-called Arrow–Debreu (A-D) security [1], which for each value of j has the payout $x_j = \mathbb{1}\{j = k\}$ for some fixed value of k . Here $\mathbb{1}\{E\}$ denotes the indicator function for the event E . Thus $x_j = 1$ if $j = k$, and $x_j = 0$ if $j \neq k$. The A-D securities are represented by pure projection operators, each with payout unity or zero, depending on the result of the underlying quantum measurement, whose outcome is also unity or zero. Thus the set of all Arrow–Debreu contracts is precisely the set of all pure projection operators on \mathcal{H} .

The state of a quantum system in n dimensions is represented by a positive semidefinite Hermitian matrix with trace unity. Such a matrix can be put in the form

$$\hat{p} = \sum_{j=1}^n p_j |\psi_j\rangle\langle\psi_j|, \quad (13)$$

for some orthonormal basis $\{|\psi_j\rangle\}_{j=1,2,\dots,n}$, with $p_j \geq 0$ for $j = 1, 2, \dots, n$ and $\sum_{j=1}^n p_j = 1$.

In the case of a density matrix of maximal rank with distinct eigenvalues, this basis is uniquely determined up to phase factors. If the density matrix is of maximal rank but with a degenerate spectrum, the basis is determined modulo unitary transformations on the degenerate subspaces. In the case of a density matrix of lower rank, the basis is determined at best only up to an arbitrary unitary transformation of the basis vectors that span the null space of the density matrix. Given two density matrices \hat{p} and \hat{q} , we say that \hat{q} is *absolutely continuous* with respect

to \hat{p} if the null space of \hat{p} is a subspace of the null space of \hat{q} . Thus, \hat{q} is absolutely continuous with respect to \hat{p} if and only if for all $|\psi\rangle \in \mathcal{H}$ such that $\hat{p}|\psi\rangle = 0$ it holds that $\hat{q}|\psi\rangle = 0$. We say that \hat{p} and \hat{q} are *equivalent* if each is absolutely continuous with respect to the other, that is to say, if they share the same null space. It is easy to see that ‘equivalence’ in this sense is an equivalence relation in the usual mathematical sense, and it follows that all density matrices of maximal rank are equivalent.

We say that a claim \hat{X}_T is *positive* if $\langle \psi | \hat{X}_T | \psi \rangle \geq 0$ for all $|\psi\rangle \in \mathcal{H}$ and *strictly positive* if $\langle \psi | \hat{X}_T | \psi \rangle > 0$ for all $|\psi\rangle \in \mathcal{H}$. By (12), \hat{X}_T is positive if and only if $x_j \geq 0$ for all j and strictly positive if and only if $x_j > 0$ for all j .

In fact, it should be evident that any claim \hat{X}_T can be split in a canonically minimal way into positive part \hat{X}_T^+ and a negative part \hat{X}_T^- such that $\hat{X}_T = \hat{X}_T^+ + \hat{X}_T^-$, where the positive eigenvalues of \hat{X}_T are those of \hat{X}_T^+ and the negative eigenvalues of \hat{X}_T are those of \hat{X}_T^- . It will thus suffice for our purpose to look at financial contracts with positive cash flows.

Let us now consider a one-period market represented by the set of all positive claims on an n -dimensional Hilbert space. The problem is to assign a value or price to each such claim \hat{X}_T on the basis of as few assumptions as possible. One possible approach would be to consider the so-called expectation value of \hat{X}_T in the state \hat{p} , given by

$$\langle \hat{X}_T \rangle_p = \text{tr}(\hat{p}\hat{X}_T). \tag{14}$$

The expectation value can be interpreted as the average value of the payoff when the average is calculated by taking numerous independent copies of the experimental setup, performing an identical measurement on each system, and averaging the results. One might think that this expectation gives a fair price for entering into the contract; but that is merely a guess. In fact, agents will typically pay less than the expectation value, in order to allow for a non-trivial rate of return on the investment in compensation for the risks involved, and in principle, the price could be any non-negative map $\Pi_{0T} : \hat{X}_T \mapsto \Pi_{0T}(\hat{X}_T) \in \mathbb{R}^+$, which need not necessarily be linear. On the other hand, we can be confident that if $\langle \hat{X}_T \rangle_p = 0$, then the price must be zero, since no rational agent would pay a strictly positive premium for an investment that paid zero with probability one. Thus we conclude that the price vanishes if and only if the expectation value of the payoff vanishes.

But we are still some distance from determining the form that the price takes. Since the expectation value is a linear function of the observable, this suggests that we look more closely at linear functionals. If \hat{X}_T and \hat{Y}_T are claims, then so is the linear combination

$$\hat{Z}_T = a\hat{X}_T + b\hat{Y}_T \tag{15}$$

for $a, b \geq 0$. Hence the space of positive claims has a convex structure. It should be clear that the experiments underlying the \hat{X}_T and \hat{Y}_T are in general different and that the experiment underlying \hat{Z}_T is different yet again. If we write these claims in their diagonalized forms

$$\hat{X}_T = \sum_{j=1}^n x_j |x_j\rangle \langle x_j|, \quad \hat{Y}_T = \sum_{j=1}^n y_j |y_j\rangle \langle y_j|, \tag{16}$$

with respect to the relevant basis vectors, one sees that the payouts and basis vectors associated with these claims are uniquely determined, up to the usual ambiguities associated with degeneracies and null spaces, and at the same time the payouts and basis vectors of (15) are represented by the decomposition

$$\hat{Z}_T = \sum_{j=1}^n z_j |z_j\rangle \langle z_j|. \tag{17}$$

Thus, if we have two contracts, each with positive payouts, depending on separate measurements, then any linear combination of the operators corresponding to the two contracts, with positive coefficients, will give rise to the operator corresponding to yet another contract, with a different set of payouts, depending on still another measurement.

Hence, a linear combination (15) is *not*, generally, to be understood as representing a ‘portfolio’ of its constituents (see section 8). This is because the payout of a portfolio is given by the totality of the payouts of its constituents. One is tempted, nonetheless, to conjecture that the price of the contract represented by a linear combination of two contracts should equal the corresponding linear combination of the prices of the constituents. But it is not obvious that this will be the case, since the new contract involves a different payout structure and a different experiment—so we do not wish to *assume* linearity in general.

Nevertheless, if \hat{U}_T and \hat{V}_T depend on the outcome of *the same experiment*, and differ only in the amounts paid for the various possible outcomes, then the price of $a\hat{U}_T + b\hat{V}_T$ should be equal to the corresponding linear combination of the prices of \hat{U}_T and \hat{V}_T . More precisely, if the \hat{U}_T and \hat{V}_T commute, then the prices should be additive. For if \hat{U}_T and \hat{V}_T commute, we can find an orthogonal basis $\{|w_j\rangle\}_{j=1,2,\dots,n}$ in which both are diagonalized:

$$\hat{U}_T = \sum_{j=1}^n u_j |w_j\rangle \langle w_j|, \quad \hat{V}_T = \sum_{j=1}^n v_j |w_j\rangle \langle w_j|. \quad (18)$$

Then if we form the linear combination $\hat{W}_T = a\hat{U}_T + b\hat{V}_T$ we obtain

$$\hat{W}_T = \sum_{j=1}^n (au_j + bv_j) |w_j\rangle \langle w_j|, \quad (19)$$

showing that the payouts for \hat{W}_T are given by linear combinations of the payouts of the constituents. Thus, *for commuting observables, the price of a linear combination of contracts should be the corresponding linear combination of the prices of the individual contracts*. But it is not obvious that linearity extends to non-commuting contracts.

4. Existence of pricing operator

Let us codify our assumptions somewhat more explicitly. As usual, we write $\mathbb{R}^+ = \{x \in \mathbb{R} : x \geq 0\}$. We fix a quantum system with state $\hat{\rho}$ on an n -dimensional Hilbert space \mathcal{H} and write \mathcal{V}^+ for the cone for positive contracts on \mathcal{H} . Thus our market is characterized by the triple $\{\mathcal{H}, \hat{\rho}, \mathcal{V}^+\}$. Let us write P_{0T} for the price of a unit discount bond. Our goal is to assign a price to each contract $\hat{X}_T \in \mathcal{V}^+$. By a *pricing function* on the market $\{\mathcal{H}, \hat{\rho}, \mathcal{V}^+\}$ in a one-period setting we mean a mapping $\Pi_{0T} : \mathcal{V}^+ \rightarrow \mathbb{R}^+$ satisfying the following:

- (i) For all $\hat{X}_T \in \mathcal{V}^+$ it holds that $\Pi_{0T}[\hat{X}_T] = 0$ if and only if $\text{tr}(\hat{\rho}\hat{X}_T) = 0$.
- (ii) If the m contracts represented by the Hermitian matrices $\{\hat{X}_T^k\}_{k=1,2,\dots,m}$ commute, then for all $\{a_k \geq 0\}_{k=1,2,\dots,m}$ one has

$$\Pi_{0T} \left[\sum_{k=1}^m a_k \hat{X}_T^k \right] = \sum_{k=1}^m a_k \Pi_{0T} [\hat{X}_T^k]. \quad (20)$$

- (iii) $\Pi_{0T}[\hat{\mathbf{1}}] = P_{0T}$.

The axioms can be interpreted as follows. Axiom (1) ensures the absence of arbitrage: the price of a positive contract vanishes if and only if the expected payout vanishes. Axiom (2) ensures that the pricing function is linear when it acts on a collection of contracts represented by commuting observables. Axiom (3) fixes the price of the risk-free asset. Then we obtain the following general characterization of the price of a contract:

Proposition 1. *If $n \geq 3$ then there exists a state \hat{q} on $\{\mathcal{H}, \hat{p}, \mathcal{V}^+\}$ that is equivalent to \hat{p} such that for any contract $\hat{X}_T \in \mathcal{V}^+$ the price of \hat{X}_T is given by*

$$\Pi_{0T}[\hat{X}_T] = P_{0T} \text{tr}(\hat{q}\hat{X}_T). \tag{21}$$

Proof. Consider the pricing of A-D securities. For each such contract, the measurement involves a projection operator $\hat{\Lambda} = |\lambda\rangle\langle\lambda|$ for some normalized vector $|\lambda\rangle \in \mathcal{H}$. The pricing function is a map from the space of pure projections on \mathcal{H} to \mathbb{R}^+ . It is well known that the space of pure projections on an n -dimensional Hilbert space is isomorphic to the complex projective space $\mathbb{C}\mathbb{P}^{n-1}$. Thus we obtain a function $\Pi_{0T} : \mathbb{C}\mathbb{P}^{n-1} \rightarrow \mathbb{R}^+$ with the property that for any n points $\{\lambda_j \in \mathbb{C}\mathbb{P}^{n-1}\}_{j=1,2,\dots,n}$ determining an orthogonal basis in \mathcal{H} one has

$$\sum_{j=1}^n \Pi_{0T}(\lambda_j) = P_{0T}. \tag{22}$$

This is because the projection operators associated with an orthonormal basis commute and hence by Axiom (2) the sum of the prices of the projection operators must equal the price of the sum of the projection operators. But the latter sum gives the identity operator, which offers a risk-free payout of unity. Thus we obtain a unit discount bond, for which the price is P_{0T} by Axiom (3). Gleason’s theorem [29] can now be applied to the problem and it follows that there exists a state \hat{q} such that the price of any claim of the form $\hat{\Lambda}$ is given by

$$\Pi_{0T}[\hat{\Lambda}] = P_{0T} \text{tr}(\hat{q}\hat{\Lambda}). \tag{23}$$

Now, any contract \hat{X}_T can be constructed as a linear combination of orthogonal pure projection operators with positive coefficients. Since these operators commute, Axiom (2) implies that the price of such a contract will be given by the sum of the prices of its elements, and this gives us (71). The fact that the ‘pricing’ operator \hat{q} must be equivalent to the ‘physical’ state \hat{p} then follows as a consequence of Axiom (1), which taken with (71) ensures that for any positive contract \hat{X}_T we have $\text{tr}(\hat{p}\hat{X}_T) = 0$ if and only if $\text{tr}(\hat{q}\hat{X}_T) = 0$. \square

The point here is that we do not assume *a priori* the existence of a pricing state. The idea rather is to prove the existence of such a state under the *prima facie* much weaker assumptions implicit in our axioms. The requirement that the pricing function is linear when it is applied to any commuting family of A-D securities coupled with the assumption that the price of a one-period discount bond is known allows us to deduce that the pricing function takes the form (71). In the case of a finite-dimension Hilbert space, the associated projective Hilbert space takes the form of a complex projective space $\mathbb{C}\mathbb{P}^{n-1}$ equipped with the Fubini-Study metric [12]. Gleason’s theorem shows for $n \geq 3$ that any map $f : \mathbb{C}\mathbb{P}^{n-1} \rightarrow [0, 1]$ with the property that $\sum_{j=1}^n f(\lambda_j) = 1$ for any set of n points $\{\lambda_j\}_{j=1,2,\dots,n} \in \mathbb{C}\mathbb{P}^{n-1}$ that are maximally distant from each other under the Fubini-Study metric necessarily takes the form $f(\lambda) = \langle\lambda|\hat{q}|\lambda\rangle/\langle\lambda|\lambda\rangle$ for some positive operator \hat{q} with trace unity. The principle of no arbitrage (‘no free lunch’) then implies that \hat{q} is equivalent to \hat{p} .

It should be noted that the physical state \hat{p} refers to the state of the quantum system upon which measurement of a given physical observable determines the payment made under the

terms of the financial contract. Thus \hat{p} can be used to calculate the probability distribution of the payout, but gives no information about the price, except that minimal statement which is mandated by the absence of arbitrage—namely, that the price should be zero if and only if the probability of a payout greater than zero is zero.

The operator $P_{OT}\hat{q}$ plays the role of a pricing kernel in our theory. In the case of an n -dimensional Hilbert space the prices of any $n^2 - 1$ linearly independent financial contracts, alongside the price of the unit discount bond, will be sufficient to completely calibrate the pricing kernel, which can then be used to price other contracts. It may seem surprising that the knowledge of such a system of prices gives no information about the physical state \hat{p} , except to determine its null space, but the analogue of this phenomenon is well established in the classical theory of finance [10, 22–24, 27]. At first glance, one might conclude that our proposition 1 has little content, since the pricing operator \hat{q} is arbitrary apart from its having the same null space as \hat{p} ; but such a conclusion would be incorrect—the point is that the existence of a pricing operator is not assumed but rather is *deduced* from the minimal axioms we have chosen to characterize a pricing function. Thus, beginning only with the assumed existence of a pricing function, which might in principle be nonlinear, one can whittle the candidates for such a map down to a linear function of the form (71).

5. Optimal investment

A well-known problem in classical finance theory is to determine, given a budget X_0 , the investment that maximizes the expectation of the utility gained by the investor at T when the proceeds of the investment are liquidated. It is reasonable to pose a similar problem in quantum finance. We assume that (a) agent A 's attitudes towards risk are expressed by a standard von Neumann-Morgenstern utility function $\{U(x)\}_{x>0}$, (b) the physical state \hat{p} of the quantum system is known, (c) the basis under which the physical measurement is being made is known, and (d) the pricing state is known. The investment is thus characterized by an observable of the form (12), where the basis $\{|x_j\rangle\}_{j=1,\dots,n}$ is fixed, and the cash flows $\{x_j\}_{j=1,\dots,n}$ must be determined so that the budget is saturated and the expected utility is maximized. What makes the problem interesting is that the expected utility of the payout is calculated by use of the physical state \hat{p} whereas the budget constraint involves the pricing state \hat{q} , and that neither \hat{p} nor \hat{q} necessarily has any relation to the measurement basis.

Definition 1. By a standard utility function we mean a map $U : \mathbb{R}^+ \setminus \{0\} \rightarrow \mathbb{R}$ that satisfies the following conditions: (i) $U \in C^2(\mathbb{R}^+ \setminus \{0\})$, (ii) $U'(x) > 0$ for all $x > 0$, (iii) $U''(x) < 0$ for all $x > 0$, (iv) $\lim_{x \rightarrow \infty} U'(x) = 0$, and (v) $\lim_{x \rightarrow 0} U'(x) = \infty$.

These requirements can be relaxed in various contexts, but the ‘standard’ conditions often lead to well-posed problems for which solutions can be shown to exist and hence prove to be natural as a basis for modelling. We see that a standard utility function is a strictly convex, strictly increasing map defined for all strictly positive values of its argument. We refer to the map $U' : \mathbb{R}^+ \setminus \{0\} \rightarrow \mathbb{R}^+ \setminus \{0\}$ as the *marginal utility*. The final two conditions of the definition ensure that there exists an inverse marginal utility function $\{I(y)\}_{y>0}$ such that $I(U'(x)) = x$ for all $x > 0$. The identity

$$U(I(y)) - I(y)y = \sup_{x>0} (U(x) - xy), \tag{24}$$

which holds for all $y > 0$, can be used to establish the so-called fundamental inequality

$$U(I(y)) - I(y)y \geq U(z) - yz, \tag{25}$$

which holds for all $y > 0$ and $z > 0$ in the case of a standard utility function.

Examples of standard utility functions are (a) logarithmic utility, for which $U(x) = \log(x)$ for $x > 0$, and (b) power utility with index $p \in (-\infty, 1) \setminus \{0\}$, for which $U(x) = p^{-1}x^p$ for $x > 0$. For logarithmic utility one finds that $I(y) = 1/y$ and for power utility $I(y) = y^{1/(p-1)}$.

The goal of agent A 's optimization problem is to determine the cash flows $\{x_j\}_{j=1,2,\dots,n}$ that maximize the expected value of the utility, providing that these cash flows can be realized with the specified budget. Thus, given a standard utility function $\{U(x)\}_{x>0}$ we set

$$\{x_j^*\}_{j=1,2,\dots,n} = \operatorname{argmax}_{\{x_j\}} \operatorname{tr} \left[\hat{p} \hat{U}(\{x_j\}) \right] \tag{26}$$

where $\hat{U}(\{x_j\}) = \sum_{j=1}^n U(x_j) |x_j\rangle \langle x_j|$ and the argmax is subject to the budget constraint

$$X_0 = P_{0T} \operatorname{tr} \left(\hat{q} \hat{X}_T \right), \quad \hat{X}_T = \sum_{j=1}^n x_j |x_j\rangle \langle x_j|. \tag{27}$$

Proposition 2. *Let the physical state of a quantum system on an n -dimensional Hilbert space be \hat{p} . Let the pricing state for a financial market based on measurements of the system be \hat{q} , with one-period discount factor P_{0T} . Let the risk preferences of the investor be represented by a standard utility function $U: \mathbb{R}^+ \setminus \{0\} \rightarrow \mathbb{R}$ and write I for the associated inverse marginal utility function. Then the optimal cash flow structure $\{x_j^*\}$ for an investment with budget X_0 paying out according to the measurement of a financial observable of the form*

$$\hat{X} = \sum_{j=1}^n x_j |x_j\rangle \langle x_j|, \tag{28}$$

for some fixed orthonormal basis $\{|x_j\rangle\}_{j=1,\dots,n}$, is given by

$$x_j^* = I \left[\lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle} \right], \tag{29}$$

where for any choice of $X_0 > 0$ the parameter λ is uniquely determined by the relation

$$P_{0T} \sum_{j=1}^n I \left[\lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle} \right] \langle x_j | \hat{q} | x_j \rangle = X_0. \tag{30}$$

Proof. The method of Lagrange multipliers can be used to obtain a candidate for the argmax. We introduce a Lagrange multiplier λ and seek a solution to the unconstrained problem

$$\{x_j^*\} = \operatorname{argmax}_{\{x_j\}} \left(\operatorname{tr} \left[\hat{p} \hat{U}(\{x_j\}) \right] - \lambda P_{0T} \operatorname{tr} \left(\hat{q} \hat{X}_T \right) \right), \tag{31}$$

or equivalently

$$\{x_j^*\} = \operatorname{argmax}_{\{x_j\}} \left(\sum_{j=1}^n U(x_j) \langle x_j | \hat{p} | x_j \rangle - \lambda P_{0T} \sum_{j=1}^n x_j \langle x_j | \hat{q} | x_j \rangle \right). \tag{32}$$

Differentiating with respect to x_j and setting the results to zero, we find that

$$U'(x_j) = \lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle} \tag{33}$$

for each value of j . Applying the inverse marginal utility function to each side of this equation, we are then led to (29) and the budget constraint (27) gives (30). That (30) admits a unique solution for λ for any $X_0 > 0$ follows from the fact that the monotonic decreasing map $I: \mathbb{R}^+ \setminus \{0\} \rightarrow \mathbb{R}^+ \setminus \{0\}$ is surjective, which is a consequence of the conditions (iv) and (v)

satisfied by a standard utility function. That the candidate solution is indeed a true solution can be checked by use of the fundamental inequality (25). It follows then from (29) that for any *alternative* choice of payout structure $\{x_j\}$ we have

$$U(x_j^*) - x_j^* \lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle} \geq U(x_j) - x_j \lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle}, \tag{34}$$

for each $j = 1, 2, \dots, n$. Multiplying by p_j and summing we obtain

$$\sum_{j=1}^n p_j U(x_j^*) - \sum_{j=1}^n p_j U(x_j) \geq \sum_{j=1}^n p_j x_j^* \lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle} - \sum_{j=1}^n p_j x_j \lambda P_{0T} \frac{\langle x_j | \hat{q} | x_j \rangle}{\langle x_j | \hat{p} | x_j \rangle}. \tag{35}$$

Then since $p_j = \langle x_j | \hat{p} | x_j \rangle$ we have

$$\sum_{j=1}^n p_j U(x_j^*) - \sum_{j=1}^n p_j U(x_j) \geq \lambda P_{0T} \left[\sum_{j=1}^n x_j^* \langle x_j | \hat{q} | x_j \rangle - \sum_{j=1}^n x_j \langle x_j | \hat{q} | x_j \rangle \right]. \tag{36}$$

Now, we know by (30) that λ has been chosen to ensure that the candidate solution $\{x_j^*\}$ satisfies the budget constraint

$$P_{0T} \sum_{j=1}^n x_j^* \langle x_j | \hat{q} | x_j \rangle = X_0. \tag{37}$$

If we require that the alternative choice of payout structure should also satisfy the budget constraint, or else operate under budget, so

$$P_{0T} \sum_{j=1}^n x_j \langle x_j | \hat{q} | x_j \rangle \leq X_0, \tag{38}$$

then the two terms on the right-hand side of (36) cancel, or else leave a difference that is positive (if the alternative choice is under budget), which gives

$$\sum_{j=1}^n p_j U(x_j^*) \geq \sum_{j=1}^n p_j U(x_j), \tag{39}$$

showing that the candidate solution for the optimal payout gives an expected utility that is no less than that of any alternative choice of payout structure with a budget no greater than that of the candidate solution. \square

6. Rate of return

As an example, we can look in detail at the case of logarithmic utility. Suppose we set $U(x) = \log x$ for $x > 0$. Then the inverse marginal utility function is given by $I(y) = 1/y$ for $y > 0$. It follows that for log utility the optimal payout structure takes the form

$$x_j^* = (\lambda P_{0T})^{-1} \frac{\langle x_j | \hat{p} | x_j \rangle}{\langle x_j | \hat{q} | x_j \rangle}. \tag{40}$$

Inserting this expression into the budget constraint (37) we obtain

$$\lambda^{-1} \sum_{j=1}^n \langle x_j | \hat{p} | x_j \rangle = X_0. \tag{41}$$

But the sum appearing in the expression above is unity since $\sum_{j=1}^n |x_j\rangle\langle x_j| = \hat{\mathbf{1}}$ and the trace of \hat{p} is one. Thus for log utility we deduce that $\lambda^{-1} = X_0$ and hence

$$x_j^* = (P_{0T})^{-1} X_0 \frac{\langle x_j | \hat{p} | x_j \rangle}{\langle x_j | \hat{q} | x_j \rangle}. \quad (42)$$

We observe that when the physical state and the pricing state are one and the same, the payouts of the optimal investment are identical for each outcome of chance, each giving $(P_{0T})^{-1} X_0$, the usual ‘future value’ of the initial investment. In that case, the optimal investment is to put the initial endowment into unit discount bonds, totalling X_0 in value. Then we have $\hat{X}_T = (P_{0T})^{-1} X_0 \hat{\mathbf{1}}$. It follows that if the pricing state is the physical state, the market assigns no premium to the return on a risky investment, ensuring that the optimal investment is in a discount bond and the rate of return is the interest rate.

The same conclusion applies, more generally, for any choice of the utility. This follows from (29) and (30), from which one concludes that if $\hat{p} = \hat{q}$ then $x_j^* = (P_{0T})^{-1} X_0$ for all j . It is interesting therefore to enquire what happens when the pricing state is different from the physical state. The expected return R_{0T} on an investment \hat{X}_T is given by the ratio of the expectation of \hat{X}_T under \hat{p} to the amount initially invested, namely X_0 . Thus, quite generally, we have

$$R_{0T} = (X_0)^{-1} \text{tr}(\hat{p}\hat{X}_T). \quad (43)$$

But $X_0 = P_{0T} \text{tr}(\hat{q}\hat{X}_T)$ by (71), so we deduce that

$$R_{0T} = (P_{0T})^{-1} \frac{\text{tr}(\hat{p}\hat{X}_T)}{\text{tr}(\hat{q}\hat{X}_T)}, \quad (44)$$

and it should be clear that if $\hat{p} = \hat{q}$, except possibly on the null space of \hat{X}_T , then the rate of return on the investment is the one-period interest rate.

Specializing now to the case of an optimal investment for an agent with logarithmic utility, let us calculate the rate of return. We have

$$\hat{X}_T = \sum_{j=1}^n x_j^* |x_j\rangle\langle x_j|, \quad (45)$$

where the optimal payout structure $\{x_j^*\}$ is given by (42). It follows then that

$$\begin{aligned} R_{0T} &= (X_0)^{-1} \text{tr}(\hat{p}\hat{X}_T) \\ &= (X_0)^{-1} \sum_{j=1}^n x_j^* \langle x_j | \hat{p} | x_j \rangle \\ &= (P_{0T})^{-1} \sum_{j=1}^n \frac{\langle x_j | \hat{p} | x_j \rangle^2}{\langle x_j | \hat{q} | x_j \rangle}. \end{aligned} \quad (46)$$

If we set $R_{0T} = e^{\mu T}$ then the rate of return μ can be split into two parts, namely a risk-free one-period interest rate and a so-called excess rate of return or risk premium, which is the part of the rate of return that exceeds the interest rate. We can represent this by writing $R_{0T} = e^{(r+\beta)T}$ where r is the interest rate and β is the excess rate of return. The interest rate is fixed by the relation $e^{rT} = (P_{0T})^{-1}$ and the excess rate of return is fixed by the relation

$$e^{\beta T} = \sum_{j=1}^n \frac{\langle x_j | \hat{p} | x_j \rangle^2}{\langle x_j | \hat{q} | x_j \rangle}. \quad (47)$$

Proposition 3. *The optimal investment in the case of an investor with logarithmic utility has a positive excess rate of return. The utility gained from such an investment in a market where the physical state and pricing state differ is greater than or equal to the utility gained from an investment in a risk-free bond.*

Proof. The expected utility gained from the payout of an optimal investment is

$$\text{tr} \left[\hat{p} \hat{U} \right] = \sum_{j=1}^n U(x_j^*) \langle x_j | \hat{p} | x_j \rangle. \tag{48}$$

Let us set $U(x_j^*) = \log x_j^*$ for logarithmic utility and insert (42). The result is

$$\sum_{j=1}^n U(x_j^*) \langle x_j | \hat{p} | x_j \rangle = \log \left[(P_{0T})^{-1} X_0 \right] + \sum_{j=1}^n \left[\langle x_j | \hat{p} | x_j \rangle \log \frac{\langle x_j | \hat{p} | x_j \rangle}{\langle x_j | \hat{q} | x_j \rangle} \right]. \tag{49}$$

The first term on the right-hand side of this equation isolates the part of the utility gain due to the interest rate. The second term can be interpreted as a *relative entropy*. In particular, if we set $p_j = \langle x_j | \hat{p} | x_j \rangle$ and $q_j = \langle x_j | \hat{q} | x_j \rangle$ then it is evident that $\{p_j\}_{j=1,2,\dots,n}$ and $\{q_j\}_{j=1,2,\dots,n}$ constitute a pair of absolutely continuous probability distributions. The second term on the right then takes the form of a Kullback–Liebler divergence [46]:

$$D_{KL}(p, q) = \sum_{j=1}^n p_j \log \left(\frac{p_j}{q_j} \right). \tag{50}$$

Thus, the utility thereby gained gives a measure of the divergence between the physical state and the pricing state. Now, it is well known that the Kullback–Liebler divergence is *non-negative*. It follows, then, that the utility gained from an optimal risky investment in a market where \hat{p} and \hat{q} are distinct will be greater than or equal to the utility gained from a risk-free bond investment, as claimed.

Moreover, we have the following. The standard logarithmic inequality $\log z \leq z - 1$, which holds for $z > 0$, implies that

$$\log \left(\frac{p_j}{q_j} \right) \leq \left(\frac{p_j}{q_j} \right) - 1 \tag{51}$$

for each j . Hence, multiplying by p_j and summing we obtain

$$\sum_{j=1}^n p_j \log \left(\frac{p_j}{q_j} \right) \leq \sum_{j=1}^n \left(\frac{p_j^2}{q_j} \right) - 1. \tag{52}$$

Thus, we have

$$e^{\beta T} = \sum_{j=1}^n \left(\frac{p_j^2}{q_j} \right) \geq 1 + D_{KL}(p, q), \tag{53}$$

and by the positivity of the Kullback–Liebler divergence we deduce that the excess rate of return β is positive for an optimal investment under logarithmic utility, as claimed. \square

7. Classical vs quantum probability

It is often maintained that quantum probability is more general than Kolmogorov's well-established 'classical' theory of probability [44] and that the latter is contained as a special case of the former. There is no doubt that quantum probability, when laid out as a mathematical theory, has a different look and feel when it is compared to Kolmogorov's theory; but despite the fact that numerous well-argued accounts of quantum probability can be found in the literature [18, 20, 30, 35, 45, 48, 56, 61], some even taking an axiomatic approach (see also [6, 20, 63]), it is not that easy to pinpoint the exact sense in which quantum theory is *essentially* non-Kolmogorovian—rather than, say, a reworking of Kolmogorov's theory in a different form. This issue is compounded by the fact that, except in the most loose terms, it is difficult to say what one means by 'probability' without embedding the concept in a mathematical framework.

It is fortunate then that we have the results of Gleason [29], Bell [7–9], Kochen & Specker [33, 43], and others following in their footsteps, which add clarity to the matter. The point is that one has to work rather hard to come up with examples of situations in quantum probability that cannot be reduced to a classical probability model. But a number of such examples have been worked out involving finite-dimensional Hilbert spaces, so this creates the prospect of constructing financial models for claims based on the results of quantum measurements, in settings for which quantum probability is required in their analysis. Since most of what we know of modern finance theory is based explicitly on Kolmogorov's framework, it may be worthwhile to take note of a few examples of situations where quantum probability comes into play.

Among the numerous attempts that have been made to generalize or extend the Kochen–Specker construction [40, 41, 50–52], perhaps the simplest yet put forward is that of Cabello *et al* [15, 16], which entails the specification of a collection of nine different non-commuting observables on a four-dimensional Hilbert space.

In a financial context, one can think of this setup as involving a single quantum system being prepared in a state \hat{p} with nine different 'draft financial contracts' drawn up, each requiring measurement of one of the nine observables. The contracts specify the payments that will be made when one of the four possible outcomes occurs for the measurement associated with a specific contract. It is of the nature of quantum probability that only one of the nine contracts can be implemented, so we can envisage a rational agent being presented with the alternatives and choosing one optimally in accordance with their needs.

In any specific setting, there will only be one contract in play, namely the one chosen by the agent after careful consideration of their criteria for optimality. In each such specific setting the usual rules for Kolmogorovian probability apply. But for the setup and description of the problem as a whole—with the presentation and analysis of the nine contracts and the posing of the optimization problem, we require quantum probability. This example can be used to refute the claim of a skeptic who asks whether one is merely taking simple examples from classical finance and dressing them up in the language of quantum probability and calling the result quantum finance. The point is that completely tractable examples can be constructed within the context of quantum finance for which no classical analogue exists.

The setup is an elaborate although feasible one, and we can use the methods discussed to calculate the probabilities of the results for the nine different measurements and hence the expected utility gained from each choice. Each observable has four possible outcomes, thus determining an orthonormal tetrad in Hilbert space. These are the four eigenvectors of the Hermitian matrix corresponding to a given observable. The result of the measurement is to select one of these eigenvectors. Equivalently, each measurement measures four commuting




Figure 1. Diagram illustrating a result of the Kochen–Specker type in a 4-dimensional Hilbert space. Each of the 9 vertices are met by 4 lines and each of the 18 lines join 2 vertices. The 18 lines represent a set of normalized projection operators with the property that the 4 projection operators meeting a given vertex are mutually orthogonal and sum to the identity operator. It is easy to see that it is impossible to ‘colour’ the lines so that one blue line meets each vertex and 3 red lines meet each vertex. This illustrates the fact that in the standard Kolmogorov setup one cannot find a set of 18 random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, each taking values in the set $\{0, 1\}$, such that when the 18 random variables are assigned to the 18 lines, the sum of the 4 random variables meeting any given vertex will be one for all $\omega \in \Omega$.

projection operators, namely the projection operators associated with the four legs of the tetrad. The outcome of one of these four measurements will be unity and the rest nil.

The clever idea behind results of the Kochen–Specker type is to choose the observables so that some of the tetrads legs overlap when one moves from one observable to another. In the present situation, involving nine observables, the overlap structure is shown in figure 1. Alongside each vertex of the enneagon one sees the corresponding tetrad, where to ease the typography we write $\bar{1}$ for -1 . When two vertices are connected by a dotted line, this means that the associated tetrads share a vector in common. The analysis is simplified somewhat by the fact that the tetrads in this example can all be taken to be real.

If we label the nine observables $\{\hat{X}_r\}_{r=1,2,\dots,9}$ and if for each value of r the four projection operators associated with \hat{X}_r are denoted $\{\hat{\pi}_{rj}\}_{j=1,2,3,4}$, then the probability that outcome j will result, if contract r is chosen, is given by $\text{tr}(\hat{\pi}_{rj}\hat{p})$. The construction of an analogous setup within Kolmogorov’s system turns out to be impossible. Since this is a rather sweeping statement, let us be a little more precise about what is being claimed. The point is that in Kolmogorov’s theory, one would have to model the setup with 36 random variables on a single probability space. The 36 random variables are grouped into nine sets of four. Let’s call these

hypothetical random variables $\{X_{rj}\}_{r=1,2,\dots,9, j=1,2,3,4}$ (with no hats). Each random variable can take the value zero or one. Thus we have a total of 36 maps of the form

$$X_{rj} : \Omega \rightarrow \{0, 1\}. \tag{54}$$

There are two requirements that have to be satisfied to match the layout of the quantum setup. First, the sum of the four random variables for a given value of r must be unity. This means that one of them must be equal to one and the other three must be equal to zero for any given outcome of chance $\omega \in \Omega$. Secondly (this is where the rabbit goes into the hat) the 36 random variables have to be equal in pairs, in conformation with the structure of the diagram in figure 1. Thus, the 36 random variables are cut down in effect to 18 by the requirement that they must match in pairs.

Can one find such a set of 18 random variables? The answer, perhaps surprisingly, is no. This can be checked by a colouring argument. Given figure 1, can one colour each line red or blue in such a way that exactly one blue line meets each vertex? Suppose one finds a way of colouring four of the lines blue, no vertex being hit by more than one blue line. That would leave one vertex unmet by a blue line. Suppose then one tried to colour five lines blue. Well, that would mean at least one vertex was hit by more than one blue line. This shows that it is impossible to construct a set of 18 random variables on a probability space in such a way that the required properties are satisfied.

In financial terms, this means that we cannot model the payouts of the nine contracts as random variables on a probability space in such a way that the outcome of chance determines the payouts of all nine. A sceptic might ask, ‘Isn’t it unlikely in practice that one will come up against such a configuration of contracts?’ Well, that may be so, but the point is that quantum finance can handle such configurations whereas classical finance cannot.

8. Portfolios

Let us return now to the matter of portfolios. There are two rather distinct notions of portfolio that arise in quantum finance. The first notion involves a portfolio of contracts all depending in their payouts on the same experiment. In that case, we can fix the n axes of the n -dimensional Hilbert space determining the frame of the measurement and write $\{\hat{\pi}_j\}_{j=1,2,\dots,n}$ for the associated projection operators. Then, for a given outcome of the experiment one of these projection operators will give the result unity and the rest zero. The projection operators can be regarded as the A-D securities for that experiment and it should be evident that any contingent claim based on the outcome of the given experiment can be written as a portfolio of n such A-D securities. Thus, for such claims we can write

$$\hat{X} = \sum_{j=1}^n \theta_j \hat{\pi}_j, \tag{55}$$

where the $\{\theta_j\}_{j=1,2,\dots,n}$ represent the holdings in the various A-D securities. More generally, if we allow short positions in the A-D securities, then the resulting overall position can be expressed uniquely as the difference between two positive claims, with the understanding that we net claims involving long and short positions in the same A-D security.

Clearly, a linear combination of two portfolios in this setting gives another portfolio. Furthermore, it should be evident that the operator corresponding to the portfolio can be represented as the sum of a trace part, proportional to the identity operator, and a trace-free part. The trace part represents a position (long or short) in the risk-free asset, and the remainder consists of investments in risky assets. For example, in two dimensions, a portfolio of the form

$2|z_1\rangle\langle\bar{z}_1| + |z_2\rangle\langle\bar{z}_2|$ consists of a long position of three-halves of one unit of the risk free asset, a long position of one-half of a unit in the A-D security $\hat{\pi}_1$ and a short position of one-half of a unit in the A-D security $\hat{\pi}_2$, since we have $2\hat{\pi}_1 + \hat{\pi}_2 = \frac{3}{2}(\hat{\pi}_1 + \hat{\pi}_2) + \frac{1}{2}(\hat{\pi}_1 - \hat{\pi}_2)$. In this way, we can isolate the risk-free part of a portfolio. This first notion of a portfolio corresponds rather closely to the notion of a portfolio in a one-period market that arises in classical finance theory [1, 10, 22–24, 27] and can be pursued further in that spirit. The point is that once the measurement basis for the underlying experiment has been fixed, the various associated operators arising for positions with different portfolio weightings commute.

As we pointed out in section 3, however, it does not make sense to form a portfolio of several contracts each based on the same quantum system but with different measurement frames, since such measurements will in general be incompatible and cannot be simultaneously realized. In our approach to the problem, we consider portfolios of assets for which the payouts are based on separate measurements being made on two or more distinct quantum systems. Imagine, for example, a financial institution where in one room an experiment is carried out on Quantum System I, with certain results obtained, and another experiment is carried out in another room on Quantum System II, with certain results obtained. In each case, there are contracts leading to payouts depending on the results obtained.

Since the measurements do not interfere with one another (after all, they are carried out in different rooms) they can be carried out simultaneously, each delivering a certain number of units of account, so it makes sense to speak of holding a portfolio in the two assets, for which the payout is simply the totality of the payouts of the constituents of the portfolio, with appropriate weightings. Let us see how we model such a situation. To simplify the discussion, we stick with the case where there are two quantum systems involved, with measurements made on each of them.

The setup can then be easily generalized to the case where there are N such systems. The key idea is that to model a portfolio of two such contracts, we need to consider the tensor product of the Hilbert spaces of the individual systems. In fact, the two Hilbert spaces might even be of different dimensions.

The usual Dirac notation does not hold up so well in such a setting, so we use an *index notation* instead, which works quite smoothly [12, 28]. Thus, let \mathcal{H}_1 be a Hilbert space of dimension n and let \mathcal{H}_2 be a Hilbert space of dimension n' , where n and n' are not necessarily the same. We write ξ^a and $\xi^{a'}$ for typical elements of \mathcal{H}_1 and \mathcal{H}_2 respectively, where $a = 1, 2, \dots, n$ and $a' = 1, 2, \dots, n'$. Thus indices without dashes refer to the first Hilbert space and indices with dashes refer to the second Hilbert space. We write η_a and $\eta_{a'}$ for typical elements of the corresponding dual spaces \mathcal{H}_1^* and \mathcal{H}_2^* . The complex conjugates of ξ^a and $\xi^{a'}$ are denoted $\bar{\xi}_a$ and $\bar{\xi}_{a'}$ respectively. Then for the inner product between ξ^a and η_a we write $\xi^a \eta_a$ and for the inner product between $\xi^{a'}$ and $\eta_{a'}$ we write $\xi^{a'} \eta_{a'}$, with the usual summation convention.

We are interested in the tensor product Hilbert space $\mathcal{H}_{12} = \mathcal{H}_1 \otimes \mathcal{H}_2$, and we write $\xi^{aa'} \in \mathcal{H}_{12}$ for a typical element of this space. Then we write $\eta_{aa'}$ for a typical element of \mathcal{H}_{12}^* and $\bar{\xi}_{aa'}$ for the complex conjugate of $\xi^{aa'}$, and for the inner product of $\xi^{aa'}$ and $\eta_{aa'}$ we write $\xi^{aa'} \eta_{aa'}$. The state of a two-particle system takes the form of a density matrix $p_{bb'}^{aa'}$. Thus we require that it should be Hermitian, of unit trace, and positive, so

$$p_{bb'}^{aa'} = \bar{p}_{bb'}^{aa'}, \quad p_{cc'}^{cc'} = 1, \quad p_{bb'}^{aa'} \alpha^b \bar{\alpha}_a \beta^{b'} \bar{\beta}_{a'} \geq 0 \tag{56}$$

for all $\alpha^a, \beta^{a'}$. A two-particle density matrix is *pure* if $p_{bb'}^{aa'} = \xi^{aa'} \bar{\xi}_{bb'}$ for some state vector $\xi^{aa'}$. We say that the particles are *independent* if

$$p_{bb'}^{aa'} = p_b^a p_{b'}^{a'} \tag{57}$$

for some pair of one-particle states p_b^a and $p_{b'}^{a'}$. The state is said to be *separable* if it can be written in the form

$$p_{bb'}^{aa'} = \sum_{r=1}^k p_b^a(r) p_{b'}^{a'}(r), \quad (58)$$

for some collection of $2k$ one-particle states $\{p_b^a(r)\}_{r=1,2,\dots,k}$ and $\{p_{b'}^{a'}(r)\}_{r=1,2,\dots,k}$. But if the two-particle state is not separable then we say that the particles are *entangled*.

Now we are in a position to discuss the idea of measurements on a two-particle system and the contracts one can associate with such measurements. A generic contract based on the outcome of a measurement made on a two-particle system is described by a Hermitian operator $X_{bb'}^{aa'}$. We are interested in the case when the measurement splits into a measurement on System I and a measurement on System II and one adds the results to give the payout of the contract. Such a contract takes the form

$$X_{bb'}^{aa'} = U_b^a \delta_{b'}^{a'} + \delta_b^a V_{b'}^{a'}, \quad (59)$$

where δ_b^a and $\delta_{b'}^{a'}$ denote the identity operators on \mathcal{H}_1 and \mathcal{H}_2 respectively. The eigenstates of such an operator are of the form

$$p_{bb'}^{aa'} = \alpha^a \bar{\alpha}_b \beta^{a'} \bar{\beta}_{b'}, \quad (60)$$

where α^a is an eigenvector of U_b^a and $\beta^{a'}$ is an eigenvector of $V_{b'}^{a'}$. Thus $U_b^a \alpha^b = u \alpha^a$ and $V_{b'}^{a'} \beta^{b'} = v \beta^{a'}$ for $u, v \in \mathbb{R}^+$ and the sum $u + v$ gives the overall payout of the contract. Such a contract represents a portfolio consisting of one unit of a contract based on System I and one unit of a contract based on System II. More generally, for a portfolio consisting of θ_1 units of the first contract and θ_2 units of the second contract we have

$$X_{bb'}^{aa'}(\theta_1, \theta_2) = \theta_1 U_b^a \delta_{b'}^{a'} + \theta_2 \delta_b^a V_{b'}^{a'} \quad (61)$$

and the payout will be of the form $\theta_1 u + \theta_2 v$. The setup for a portfolio of arbitrary size can be constructed analogously. In particular, one can check that the expected payout of a portfolio is equal to the sum of the expectations of the constituents. This is because whenever the density matrix of the two-particle state hits one of the identity operators in the portfolio operator, all but one of the systems gets traced out and one is left with the trace of the product of a single particle density operator and the observable associated with that system. For example, in the case of a two-particle system one finds that

$$\begin{aligned} p_{aa'}^{bb'} X_{bb'}^{aa'}(\theta_1, \theta_2) &= p_{aa'}^{bb'} \left(\theta_1 U_b^a \delta_{b'}^{a'} + \theta_2 \delta_b^a V_{b'}^{a'} \right) \\ &= \theta_1 p_{aa'}^{bb'} U_b^a \delta_{b'}^{a'} + \theta_2 p_{aa'}^{bb'} \delta_b^a V_{b'}^{a'} \\ &= \theta_1 p_a^b U_b^a + \theta_2 p_a^{b'} V_{b'}^{a'}, \end{aligned} \quad (62)$$

where $p_a^b = p_{ac'}^{bc'}$ and $p_a^{b'} = p_{ca'}^{cb'}$. Likewise one can check that the price of a portfolio is equal to the weighted sum of the prices of its constituents. The point is that the two-particle system is itself a quantum system with a financial observable based on it, of the form (59), so by proposition 1 there exists a pricing operator $q_{aa'}^{bb'}$ such that

$$q_{aa'}^{bb'} X_{bb'}^{aa'}(\theta_1, \theta_2) = \theta_1 q_a^b U_b^a + \theta_2 q_a^{b'} V_{b'}^{a'}, \quad (63)$$

where the traced-out operators $q_a^b = q_{ac'}^{bc'}$ and $q_a^{b'} = q_{ca'}^{cb'}$ are the pricing operators associated with the respective individual systems.

There is one further aspect of the portfolio problem that can be analyzed and this concerns the matter of correlations. If the state of the two-particle system is of the form (57), so the

two particles are independent, then the outcomes of the experiments on the two systems will be uncorrelated. But if the systems are entangled, then the correlation will in general be non-vanishing, leading to relations such as

$$p_{aa'}^{bb'} [U_b^a - \delta_b^a p_c^d U_d^c] [V_{b'}^{a'} - \delta_{b'}^{a'} p_{c'}^{d'} V_{d'}^{c'}] \neq 0. \tag{64}$$

The point about entanglement is that even if the two systems are in separate rooms (or different cities) the outcomes may be correlated, owing to the original construction of the state of the two-particle system to which they belong. The same is true of the prices: if $q_{bb'}^{aa'}$ is entangled, then there will be correlations in the prices, as shown in relations such as

$$q_{aa'}^{bb'} [U_b^a - \delta_b^a q_c^d U_d^c] [V_{b'}^{a'} - \delta_{b'}^{a'} q_{c'}^{d'} V_{d'}^{c'}] \neq 0. \tag{65}$$

Thus, in the general situation we see that when there is a market based on contracts associated with measurements being made on a number of different quantum systems, there will be correlations between outcomes of measurements and correlations between prices, where the former are determined by the structure of physical density operator for the market as a whole and the latter by the structure of the pricing operator for the market as a whole.

Since our work overlaps with that of [5] in some respects, we comment briefly on where we agree and where we differ. The differences are perhaps most apparent in the treatment of portfolios. The authors of [5] consider a market described by a quantum density operator rather than by classical probabilities, then introduce so-called quantum assets. A quantum asset is defined to be a positive semidefinite Hermitian matrix. The assets are given a ‘financial interpretation,’ namely that, ‘Each eigenstate can be considered a natural event for the quantum asset. Each eigenvalue is the outcome or payoff of this asset when the corresponding event happens.’ We are generally in agreement with this idea, though in our approach there is no abstraction: a quantum asset is a financial contract with a well-defined structure involving a payoff contingent on an experiment. In contrast, of their definition the authors of [5] say, ‘This definition leaves questions about the existence/validity of such assets and their intrinsic value for future work.’ A ‘toy example’ is given involving a market maker with access to a quantum computer. The initial state is known to the market along with the current bid and offer prices for a certain asset. Then a unitary transformation, also known to the market, is applied to the state and a measurement is made, the random outcome of which determines the market maker’s new bid and offer. Investors at the initial prices may or may not make a profit by trading again at the new prices. In our view, this example is defective, since the algorithms used by quantum computers are typically designed to give a definite or near-definite result, not a probabilistic result; in short, the introduction of a quantum computer in the toy example is spurious. To complicate matters, the authors fail to make a distinction between the price and the payout of an asset, leading to the confused idea that a market maker quotes a random price based on the outcome of a measurement. We also differ from [5] in our treatment of the no-arbitrage condition. The authors of [5] introduce the idea of a portfolio of quantum assets as a linear combination of the matrices associated with the various assets, each matrix being weighted by the number of units held in that asset. The ‘expected value’ of the portfolio is then given in [5] by the trace of the product of the market density operator and the weighted sum of the matrices. The problem is that the weighted sum of the assets is indeed an asset itself, but its payoffs are not given by weighted sums of the payoffs of the individual assets. In this respect, our approach differs completely from that of paper [5]. In our view, the attempt set out in [5] to propose an analogue of the fundamental theorem of asset pricing on this basis is ill-posed.

In contrast, in the approach of the present paper the assumption of no arbitrage is that of Axiom 1; the relation to classical finance is embodied in Axiom 2; and the risk-free rate is

specified in Axiom 3. The physical density operator \hat{p} arises in the specification of the experiment that underlies the contract. The density operator \hat{q} is shown to exist in proposition 1 and its equivalence to \hat{p} follows as a consequence of Axiom 1. Portfolios then arise naturally as multi-particle systems.

9. Conclusion

The physical density operator is objective in nature, the only limitations in its determination being in the usual practicalities of the laboratory settings where the states are manufactured. The pricing operator, on the other hand, if classical finance theory is any guide in the matter [1, 10, 22–24, 27], will be determined by the collective appetite for risk and reward among market participants. Hence, as in all markets, prices will be subject to fluctuation and change over time and may even be amenable to a Bayesian treatment. In the one-period setting that we have developed here, all we can say *a priori* of a definite nature about the pricing operator is that it exists and that the physical density operator and the pricing operator are ‘equivalent’, as we have seen in proposition 1.

In the one-period version of the theory, one can be somewhat agnostic on the matter of dynamics. This is because \hat{p} , \hat{q} and P_{0T} are specified at time 0 and no further data are needed apart from the observable \hat{X}_T being measured at time T . From a dynamical perspective it is convenient to work in the Heisenberg representation. Then \hat{p} and \hat{q} are fixed and $\{\hat{X}_t\}_{t \geq 0}$ is dynamical, given by

$$\hat{X}_T = e^{-i\hat{H}T} \hat{X}_0 e^{i\hat{H}T}, \tag{66}$$

where \hat{X}_0 denotes the initial value of the observable being measured and \hat{H} is the Hamiltonian of the underlying physical system. Since T is fixed, it suffices to specify \hat{X}_T , and we let \hat{X}_0 and \hat{H} drop out of the picture.

The Heisenberg representation is also convenient when interventions are taken into account in a multi-period model. Suppose, for example, we consider a two-period model involving a pair of systems defined on the product of two Hilbert spaces. We write $0 < t < u$ and let $X_b^a(t)$ and $Y_{b'}^{a'}(u)$ be a pair of observables, one for a measurement acting on the first particle at time t and another for a measurement acting on the second particle at time u . If we assume that the two particles are non-interacting, then the two observables evolve independently, each according a law of the form (66), with distinct Hamiltonians. The physical state of the system can be represented in line with the scheme outlined in the previous section by a tensor of the form $p_{aa'}^{bb'}(0)$ and for the pricing state we write $q_{aa'}^{bb'}(0)$. Note that although the two particles are non-interacting, we allow for the possibility that they may have been prepared in an entangled state, so the two density matrices need not be separable.

Then for the valuation of the contract defined by the measurement of $X_b^a(t)$ alone, one can work with the reduced density matrices defined by $p_a^b(0) = p_{ac'}^{bc'}$ and $q_a^b(0) = q_{ac'}^{bc'}$. But for a payout involving both a measurement of $X_b^a(t)$ and a measurement of $Y_{b'}^{a'}(u)$, matters are a little more complicated. This is because once the result of the first measurement is known, the market may change its assessment of the pricing state, in line with classical idea that the pricing kernel is an adapted process, so that market participants will adjust their attitudes towards risk following a movement in the market. We refer to such state changes in the Heisenberg representation as ‘interventions.’ Now, the change in the physical state is relatively straightforward:

this is the usual Lüders state-reduction rule [47], which depends on the outcome of the measurement of $X_b^a(t)$. The transformation is thus given by

$$p_{aa'}^{bb'}(0) \rightarrow p_{aa'}^{bb'}(t) = L_a^c L_d^b p_{ca'}^{db'}(0) / L_e^c L_d^e p_{cf'}^{df'}(0), \quad (67)$$

where L_a^b denotes the projection operator onto the Hilbert subspace defined by the random outcome of the measurement of $X_b^a(t)$.

But the pricing state need not follow the Lüders rule: it is constrained only by the requirement that the new pricing state arising after the measurement should be equivalent to the new physical state. Thus, in general, we have a transformation of the form

$$q_{aa'}^{bb'}(0) \rightarrow q_{aa'}^{bb'}(t), \quad (68)$$

where $q_{aa'}^{bb'}(t)$ and $p_{aa'}^{bb'}(t)$ share the same null space. One way of achieving this is by taking any alternative state $r_{aa'}^{bb'}(0)$ which is equivalent to $p_{aa'}^{bb'}(0)$ and then passing it through the Lüders projection sieve on the first Hilbert space to give

$$q_{aa'}^{bb'}(t) = L_a^c L_d^b r_{ca'}^{db'}(0) / L_e^c L_d^e r_{cf'}^{df'}(0). \quad (69)$$

Then for each possible result of the first measurement we obtain a new pricing state, which can be used to form a time- t conditional valuation of the payout triggered by the later time- u measurement. This can be compared with the time-0 valuation of the second payout, which is obtained by using the original pricing state but tracing out the first Hilbert space.

Thus, one can think of the financial product under consideration as a contract with two cash flows, one at t and one at u . The value of the contract at time 0 is

$$S_0 = P_{0t} q_{aa'}^{bb'}(0) X_b^a(t) \delta_{b'}^{a'} + P_{0u} q_{aa'}^{bb'}(0) \delta_b^a Y_{b'}^{a'}(u). \quad (70)$$

Then at time t the contract delivers its first cash flow and goes ex-dividend; and its new value, conditional on the outcome of the first measurement, is

$$S_t = P_{tu} q_{aa'}^{bb'}(t) \delta_b^a Y_{b'}^{a'}(u), \quad (71)$$

where $P_{st} = P_{0t} / P_{0s}$ is the usual forward discount factor. Note that $q_{aa'}^{bb'}(t)$ depends on the random outcome of the first measurement. Finally, the second cash flow kicks in at time u and the asset goes ex-dividend once more, so we have $S_u = 0$. In this way we obtain stochastic processes for the value of the asset and its dividend flow. The scheme can easily be generalized to a market of any number of periods, each involving a new measurement.

That the non-Kolmogorovian character of quantum probability may have implications for the development of quantum technologies is widely appreciated—see [36] and references cited therein. And indeed, if quantum computers eventually replace the classical computers currently used for algorithmic trading by financial institutions, as they no doubt will, then the role of valuations of the type we have considered here may be important in that context. There is also a widely held view that quantum probability may play a part in cognitive science and hence behavioural finance as well—see [13, 14, 32, 42, 53, 54, 64] and references cited therein. The suggestion is that the brain uses quantum probability in a crucial way in its decision-making apparatus. In that respect, quantum cognition and quantum psychology can be viewed as a possible basis through which asset prices might be subject to quantum-like laws. It would be outside of the scope of the present discussion to look at such proposals in detail here, but if judgements and decisions are made on the basis of quantum probability, then in some situations these assessments will involve *valuations*, rather than probability estimates, and it would be the pricing operator, rather than the physical density operator, that would come into play in these valuations. In such cases, external intervention in the form of Bayesian updating could be modelled, e.g. as in [13]. This is consistent with the point we made earlier about the pricing

operator being specific to the risk and reward profiles of market operatives and in a state of flux as new information arrives. These and other further developments of the theory we hope to explore elsewhere.

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

The authors wish to thank J R Boland, D C Brody, B K Meister and the anonymous referees for helpful comments and suggestions.

ORCID iDs

Lane P Hughston  <https://orcid.org/0000-0003-1185-934X>

Leandro Sánchez-Betancourt  <https://orcid.org/0000-0001-6447-7105>

References

- [1] Arrow K J and Debreu G 1954 Existence of an equilibrium for a competitive economy *Econometrica* **22** 265–90
- [2] Assouel A, Jacquier A and Kondratyev A 2022 A quantum generative adversarial network for distributions *Quantum Mach. Intell.* **4** 28
- [3] Baaquie B E 2013 Financial modelling and quantum mathematics *Comput. Math. Appl.* **65** 1665–73
- [4] Baaquie B E 2018 *Quantum Field Theory for Economics and Finance* (Cambridge University Press)
- [5] Bao J and Rebstrost P 2023 Fundamental theorem for quantum asset pricing (arXiv:2212.1381v2)
- [6] Barnum H, Caves C M, Finkelstein J, Fuchs C A and Schack R 2000 Quantum probability from decision theory? *Proc. R. Soc. A* **456** 1175–82
- [7] Bell J S 1964 On the Einstein Podolsky Rosen paradox *Physics* **1** 195–200
- [8] Bell J S 1966 On the problem of hidden variables in quantum mechanics *Rev. Mod. Phys.* **38** 447–52
- [9] Bell J S 1987 *Speakable and Unsayable in Quantum Mechanics* (Cambridge University Press)
- [10] Bingham N H and Kiesel R 2004 *Risk Neutral Valuation* 2nd edn (Springer)
- [11] Black F and Scholes M 1973 The pricing of options and corporate liabilities *J. Political Econ.* **81** 637–54
- [12] Brody D C and Hughston L P 2001 Geometric quantum mechanics *J. Geom. Phys.* **38** 19–53
- [13] Brody D C 2023 Quantum formalism for cognitive psychology *Sci. Rep.* **13** 16104
- [14] Busemeyer J R and Bruza P D 2012 *Quantum Models of Cognition and Decision* (Cambridge University Press)
- [15] Cabello A, Estebaranz J M and García-Alcaine G 1996 Bell-Kochen-Specker theorem: a proof with 18 vectors *Phys. Lett. A* **4** 183–7
- [16] Cabello A 1997 A proof with 18 vectors of the Bell-Kochen-Specker theorem *New Developments on Fundamental Problems in Quantum Physics* ed M Ferrero and A van der Merwe (Kluwer Academic) pp 59–62
- [17] Chen Z 2004 Quantum theory for the binomial model in finance theory *J. Syst. Sci. Complex.* **17** 567–73
- [18] Davies E B 1976 *Quantum Theory of Open Systems* (Academic)
- [19] D’Hooghe B, Aerts D and Haven E 2008 Quantum formalisms in non-quantum physics situations: historical developments and directions for future research *9th Biennial IQSA Meeting Quantum Structures Brussels-Gdansk’08*
- [20] Deutsch D 1999 Quantum theory of probability and decisions *Proc. R. Soc. A* **455** 3129–37
- [21] Dixit A and Pindyck R 1994 *Investment Under Uncertainty* (Princeton University Press)
- [22] Dothan M U 1990 *Prices in Financial Markets* (Oxford University Press)

- [23] Duffie D 2001 *Dynamic Asset Pricing Theory* 3rd edn (Princeton University Press)
- [24] Etheridge A 2002 *A Course in Financial Calculus* (Cambridge University Press)
- [25] Focardi S, Fabozzi F J and Mazza D 2020 Quantum Option Pricing and Quantum Finance *J. Derivatives* **28** 79–98
- [26] Fedorov A K, Gisin N, Belousov S M and Lvovsky A I 2022 Quantum computing at the quantum advantage threshold (arXiv:2203.17181v1)
- [27] Föllmer H and Schied A 2010 *Stochastic Finance* 3rd edn (De Gruyter)
- [28] Geroch R 2013 *Quantum Field Theory* (Minkowski Institute Press)
- [29] Gleason A M 1957 Measures on the closed subspaces of a Hilbert space *J. Math. Mech.* **6** 885–94
- [30] Gudder S P 1988 *Quantum Probability* (Academic)
- [31] Haven E 2005 Pilot-wave theory and financial option pricing *Int. J. Theor. Phys.* **44** 1957–62
- [32] Haven E and Khrennikov A 2016 Statistical and subjective interpretations of probability in quantum-like models of cognition and decision making *J. Math. Psychol.* **74** 82–91
- [33] Held C 2022 The Kochen-Specker theorem *Stanford Encyclopedia of Philosophy* Fall 2022 edn, ed E N Zalta and U Nodelman
- [34] Herman D, Googin C, Liu X, Sun Y, Galda A, Safro I, Pistoia M and Alexeev Y 2023 Quantum computing for finance *Nat. Rev. Phys.* **5** 450–65
- [35] Holevo A S 1982 *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland)
- [36] Holik F H 2022 Non-Kolmogorovian probabilities and quantum technologies *Entropy* **24** 1666
- [37] Hughston L P and Zervos M 2001 Martingale approach to the pricing of real options *Disordered and Complex Systems (AIP Conf. Proc.)* vol 553, ed P Sollich, A Coolen, L P Hughston and R F Streater pp 325–30
- [38] Isham C J 1995 *Lectures on Quantum Theory* (Imperial College Press)
- [39] Jacquier A and Kondratyev O 2022 *Quantum Machine Learning and Optimisation in Finance* (Packt Publishing Ltd)
- [40] Kernaghan M 1994 Bell-Kochen-Specker theorem for 20 vectors *J. Phys. A: Math. Gen. A* **27** L829–30
- [41] Kernaghan M and Peres A 1994 Kochen-Specker theorem for eight-dimensional space *Phys. Lett. A* **198** 1–5
- [42] Khrennikov A 2010 *Ubiquitous Quantum Structure: From Psychology to Finance* (Springer)
- [43] Kochen S and Specker E 1967 The problem of hidden variables in quantum mechanics *J. Math. Mech.* **17** 59–87
- [44] Kolmogorov A N 1956 *Foundations of the Theory of Probability Theory* (Chelsea Publishing Co.) Kolmogorov A N 1933 *Grundbegriffe der Wahrscheinlichkeitsrechnung, Ergebnisse Der Mathematik* (Engl. transl.)
- [45] Kraus K 1983 *States, Effects and Operations* (Springer)
- [46] Kullback S and Leibler R A 1951 On information and sufficiency *Ann. Math. Stat.* **22** 79–86
- [47] Lüders G 1951 Über die Zustandsänderung durch den Messprozess *Ann. Phys., Lpz.* **8** 322–8
- [48] Mackey G W 1963 *Mathematical Foundations of Quantum Mechanics* (W A Benjamin)
- [49] McCloud P 2018 Quantum bounds for option prices SSRN: 3082561 (arXiv:1712.01385)
- [50] Mermin N D 1990 Simple unified form for the major no-hidden-variables theorems *Phys. Rev. Lett.* **65** 3373–6
- [51] Penrose R 2000 On Bell non-locality without probabilities: some curious geometry *Quantum Reflections* ed J Ellis and D Amati (Cambridge University Press) pp 1–27
- [52] Peres A 1991 Two simple proofs of the Kochen-Specker theorem *J. Phys. A* **24** L175–8
- [53] Pothos E M and Busemeyer J R 2013 Can quantum probability provide a new direction for cognitive modelling? *Behav. Brain Sci.* **36** 255–327
- [54] Pothos E M and Busemeyer J R 2022 Quantum cognition *Ann. Rev. Psychol.* **73** 749–78
- [55] Rebentrost P, Gupt B and Bromley T R 2018 Quantum computational finance: Monte Carlo pricing of financial derivatives *Phys. Rev. A* **98** 022321
- [56] Segal I E 1947 Postulates for general quantum mechanics *Ann. Math.* **48** 930–48
- [57] Segal W and Segal I E 1998 The Black-Scholes pricing formula in the quantum context *Proc. Natl. Acad. Sci.* **95** 4072–5
- [58] Stamatopoulos N, Egger D J, Sun Y, Zoufal C, Iten R, Shen N and Woerner S 2020 Option pricing using quantum computers *Quantum* **4** 291

- [59] Stamatopoulos N, Mazzola G, Woerner S and Zeng W J 2022 Towards quantum advantage in financial market risk using quantum gradient algorithms *Quantum* **6** 770
- [60] Stamatopoulos N, Clader B D, Woerner S and Zeng W J 2024 Quantum risk analysis in financial derivatives (arXiv:[2404.10088](https://arxiv.org/abs/2404.10088))
- [61] Streater R F 2000 Classical and quantum probability *J. Math. Phys.* **41** 3556–603
- [62] Trigeorgis L 1996 *Real Options* (MIT Press)
- [63] Wallace D 2010 How to prove the Born rule *Many Worlds? Everett, Quantum Theory and Reality* ed S Saunders, J Barrett, A Kent and D Wallace (Oxford University Press)
- [64] Yukalov V I and Sornette D 2017 Quantum probabilities as behavioral probabilities *Entropy* **19** 112