

Hierarchical Bayesian Estimation of Quantum Decision Model Parameters

Jerome R. Busemeyer¹, Zheng Wang², and Jennifer S. Trueblood³

¹ Indiana University, Bloomington, USA

² Ohio State University, Columbus, USA

³ University of California, Irvine, USA

Abstract. Quantum decision models have been recently proposed to account for findings that have resisted explanation by traditional decision theories. This paper compares quantum versus Markov models of decision making for explaining a puzzling empirical finding from human decision making called dynamic inconsistency – that is the failure of decision makers to carry out their planned decisions. A large data set that empirically investigated dynamic inconsistency was used to quantitatively evaluate the quantum and Markov models. In this application, the quantum model reduces to the Markov model when one of the parameters is set to zero. The parameters of the quantum model were estimated using Hierarchical Bayesian estimation. The distribution of the key quantum parameter was clearly located in the quantum regime and far below zero as predicted by the Markov model. These results provide further support for quantum models as compared to the traditional models of decision making.

1 Introduction

Several new quantum models of decision making have been introduced to account for decision making paradoxes that have resisted explanations by “classical” type of decision theories (Busemeyer, Wang, Lambert-Mogiliansky [3]; Lambert-Mogiliansky, Zamir, Zwirn [5]; Khrennikov and Haven [4]; Pothos & Busemeyer [6]; Yukalov & Sornette [8]). Perhaps quantum models succeed where classic models fail simply because quantum models are more complex and have greater model fitting flexibility (after all they are based on complex numbers). The purpose of this paper is to examine this issue by comparing a classic type of Markov model with a quantum model using Hierarchical Bayesian parameter estimation methods [2]. The model comparison is based on a large experiment designed to examine dynamic inconsistency in choices among two stage gambles [1]. Dynamic consistency is a principle of decision making required for backward induction when applied to decision trees. Dynamic consistency requires that a planned course of action for a future decision is implemented as planned when that decision is finally realized. Barkan and Busemeyer [1] observed systematic violations of dynamic consistency, and they used a random utility version of prospect theory to account for these findings. But more recently, Yukalov and

Sornette argued that quantum theory can also account for these findings [9]. Therefore, in this paper, two different types of models are proposed to explain these findings: a Markov model and a quantum decision model [6].

The paper is organized as follows. First we review the Barkan and Busemeyer [1] experimental methods and results. Second, we describe the two models that are being compared. Third, we present fits to the mean data for each model to get a rough idea about how well each model accounts for the findings (but this is not our main concern). Fourth, we present the results of the Hierarchical Bayesian parameter (which is our main concern). Finally, we draw some preliminary conclusions from this model comparison analysis.

2 Barkan and Busemeyer (2003)

A two stage gambling paradigm was used to study dynamic consistency, which was a modification of the paradigm used by Tversky and Shafir [7] to study the disjunction effect. A total of 100 people participated and each person played the 17 gambles involving real money shown in Table 1 twice except for the first one. Each gamble had an equal chance of producing a win or a loss. The columns labeled ‘win’ and ‘loss’ indicate the money that could be won or lost for each gamble (one unit was worth one cent). For each gamble in Table 1, the person was forced to play the first round, and then contingent on the outcome of the first round, they were given a choice whether or not to play the second round with the same gamble. On each trial the person was first asked to make a plan for the second play contingent on each possible outcome of the first play. In other words, during the planning stage they were asked two questions: “if you win the first play, do you plan to play the second gamble?” and “if you lose the first play, do you plan to play the second gamble?” Following the plan, the outcome of the first gamble was revealed, and then the person was given a final choice: decide again whether or not to play the second gamble after observing the first play outcome. To incentivize both plan and final choices, the computer randomly selected either the planned choice or the final choice to determine the real monetary payoff for each trial. The final payment for the trial was then shown to the person at the end of each trial. Participants were paid by randomly selecting four problems from the entire set, randomly selecting either their plan or final choice, and randomly selecting an outcome for each gamble to determine the actual payment.

Table 1 displays the results obtained after averaging across the two replications for each person, and after averaging across all 100 participants. The probability of planning to take the gamble is shown under the column labeled “Plan.” There was little or no difference between the probabilities of taking the gamble, contingent on each planned outcome of the first gamble, and so the results shown here are averaged across the two hypothetical outcomes during the plan. See Barkan and Busemeyer [1] for the complete results listed separately for each contingent outcome. The probability of taking the gamble during the final stage is shown under the column labeled “Final.” The columns under the

label ‘‘Gamble’’ display the amount to win and lose for each gamble. Changes in probabilities down the rows of the Table show the effect of the gamble payoffs on the probability of taking the gamble. The difference between the planned and final columns indicates a dynamic inconsistency effect. Notice that following a win (the first 4 columns), the probability of taking the gamble at the final stage was always smaller than the probability of taking the gamble at the planning stage. In other words, participants changed their minds and became more risk averse after experiencing a win as compared to planning for a win. Notice that following a loss (the last 4 columns), the probability of taking the gamble at the final stage was always greater than the probability of taking the gamble at the planning stage. In other words, participants changed their minds and became more risk seeking after experiencing a loss as compared to planning for a loss.

Table 1. Barkan and Busemeyer (2003) Experiment

Gamble		Win First Play		Gamble		Lose First Play	
Win	Loss	Plan	Final	Win	Loss	Plan	Final
200	220	0.46	0.34	80	100	0.36	0.44
180	200	0.45	0.35	100	120	0.47	0.63
200	200	0.59	0.51	100	100	0.63	0.64
120	100	0.70	0.62	200	180	0.57	0.69
140	100	0.62	0.54	160	140	0.68	0.69
200	140	0.63	0.53	200	160	0.67	0.72
200	120	0.74	0.68	160	100	0.65	0.73
200	100	0.79	0.70	180	100	0.68	0.80
				200	100	.85	.82

3 Decision Models

3.1 Quantum Decision Model

The quantum model used to account for the dynamic inconsistency effect is the same model that was previously developed by Pothos and Busemeyer [6] to account for the disjunction effect. The essential idea is that the decision maker uses a consistent utility function for plans and final decisions and always incorporates the outcomes from the first stage into the decision for the second stage. The planned decision differs from the final decision, because the plan is based on a superposition over possible first stage outcomes that will be faced during the final stage.

The two stage game involves a set of four mutually exclusive and exhaustive outcomes $\{WT, WR, LT, LR\}$ where for example WT symbolizes the event ‘win the first stage’ and ‘take the second stage gamble,’ and LR represents the event ‘lose the first stage’ and ‘reject the second stage gamble.’ These four events correspond to four mutually exclusive and exhaustive basis states

$\{|WT\rangle, |WR\rangle, |LT\rangle, |LR\rangle\}$. The four basis states are represented in the quantum model as four orthonormal basis vectors that span a four dimensional vector space. The state of the decision maker is a superposition over these four orthonormal basis states.

$$|\psi\rangle = \psi_{WT} \cdot |WT\rangle + \psi_{WR} \cdot |WR\rangle + \psi_{LT} \cdot |LT\rangle + \psi_{LR} \cdot |LR\rangle, \quad (1)$$

$$\| |\psi\rangle \|^2 = 1.$$

The initial state is represented by a 4×1 matrix ψ_I containing elements ψ_{ij} $i = W, L$ and $j = T, R$ which is the amplitude distribution over the four basis states. Initially, during the planning stage, an equal distribution is assumed so that ψ_I has elements $\psi_{ij} = 1/2$ for all four entries. The state following experience of a win is updated to ψ_W which has $1/\sqrt{2}$ in the first two entries and zeros in the second two. The state following experience of a loss is updated to ψ_L which has $1/\sqrt{2}$ in the last two entries and zeros in the first two entries. Note that $(\psi_W^\dagger \cdot \psi_L) = 0$, and also we can write $\psi_I = \frac{1}{\sqrt{2}}\psi_W + \frac{1}{\sqrt{2}}\psi_L$.

Evaluation of the payoffs causes the initial state ψ_I to be "rotated" by a unitary operator U into the final states used to make a choice about taking or rejecting the second stage gamble:

$$\psi_F = U \cdot \psi_I \quad (2)$$

$$U = \exp\left(-i \cdot \frac{\pi}{2} \cdot (H_1 + H_2)\right)$$

where

$$H_1 = \begin{bmatrix} \frac{h_W}{\sqrt{1+h_W^2}} & \frac{1}{\sqrt{1+h_W^2}} & 0 & 0 \\ \frac{1}{\sqrt{1+h_W^2}} & \frac{-h_W}{\sqrt{1+h_W^2}} & 0 & 0 \\ 0 & 0 & \frac{h_L}{\sqrt{1+h_L^2}} & \frac{1}{\sqrt{1+h_L^2}} \\ 0 & 0 & \frac{1}{\sqrt{1+h_L^2}} & \frac{-h_L}{\sqrt{1+h_L^2}} \end{bmatrix}, \quad H_2 = \frac{-\gamma}{\sqrt{2}} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}. \quad (3)$$

The upper left corner of H_1 is defined by the payoffs given a win; and the bottom right corner of H_1 is defined by the payoffs given a loss (this is described in more detail below). The matrix H_2 aligns beliefs and actions by amplifying the potentials for states WT, LR and and attenuating potentials for states WR, LT . The parameter γ is a free parameter that allows changes in beliefs during the decision process.

The utilities for taking the gamble or not are mapped into the parameters h_W and h_L in H_1 , and the latter must be scaled between -1 to $+1$. To accomplish this, the parameter h_W used to define H_1 is defined as

$$h_W = \frac{2}{1 + e^{-D_W}} - 1, \quad (4)$$

$$D_W = u(G|Win) - x_W^a,$$

$$u(G|Win) = (.50) \cdot (x_W + x_W)^a + (.50) \cdot (x_W - x_L)^a, \text{ if } (x_W - x_L) > 0$$

$$u(G|Win) = (.50) \cdot (x_W + x_W)^a - (.50) \cdot b \cdot |(x_W - x_L)|^a, \text{ if } (x_W - x_L) < 0$$

where x_W represents the amount won on gamble G . The parameter h_L used to define H_1 is defined as

$$h_L = \frac{2}{1 + e^{-D_L}} - 1 \quad (5)$$

$$D_L = u(G|Loss) - (-b \cdot x_L^a)$$

$$u(G|Loss) = (.50) \cdot (x_W - x_L)^a - (.50) \cdot b \cdot (x_L + x_L)^a, \text{ if } (x_W - x_L) > 0$$

$$u(G|Loss) = -(.50) \cdot b \cdot |(x_W - x_L)|^a - (.50) \cdot b \cdot (x_L + x_L)^a, \text{ if } (x_W - x_L) < 0$$

where x_L represents the amount lost on gamble G . Parameters a and b are risk aversion and loss aversion parameters respectively. The projection matrix

$$M = \begin{bmatrix} T & 0 \\ 0 & T \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}. \quad (6)$$

is used to map states into the response for taking the gamble on the second stage. The probability of planning to take the second stage gamble equals

$$p(T|Plan) = \|M \cdot U \cdot \psi_I\|^2. \quad (7)$$

The probability of taking the second stage game following the experience of a win equals

$$p(T|Win) = \|M \cdot U \cdot \psi_W\|^2. \quad (8)$$

The probability of taking the second stage game following the experience of a loss equals

$$p(T|Loss) = \|M \cdot U \cdot \psi_L\|^2. \quad (9)$$

If $\gamma \neq 0$ then we find that the quantum model produces interference that helps account for the observed dynamic inconsistency effects:

$$\begin{aligned} \|M \cdot U \cdot \psi_I\|^2 &= \frac{1}{2} \cdot \|M \cdot U \cdot (\psi_W + \psi_L)\|^2 & (10) \\ &= \frac{1}{2} \cdot \|M \cdot U \cdot \psi_W + M \cdot U \cdot \psi_L\|^2 \\ &= \frac{1}{2} \cdot \|M \cdot U \cdot \psi_W\|^2 + \frac{1}{2} \cdot \|M \cdot U \cdot \psi_L\|^2 \\ &\quad + \frac{1}{2} \cdot (\psi_W^\dagger \cdot U \cdot M) \cdot (M \cdot U \cdot \psi_L) \\ &\quad + \frac{1}{2} \cdot (\psi_L^\dagger \cdot U \cdot M) \cdot (M \cdot U \cdot \psi_W). \end{aligned}$$

In sum, this quantum model has only three parameters: a and b are used to determine the utilities; the third is the parameter γ for changing beliefs to align with actions. These three parameters were fit to the 33 data points in Table 1 (each gamble played twice expect for the first), and the best fitting parameters (minimizing sum of squared error) are $a = .7101$, $b = 2.5424$, and $\gamma = -4.4034$. The risk aversion parameter is a bit below one as expected, and the loss parameter b exceeds one, as it should be. The model produced an $R^2 = .8234$ and an *adjusted* $R^2 = .8120$ (the adjusted R-square includes a penalty that depends on the number of model parameters fit to the data).

3.2 Markov Decision Model

The Markov model is a special case of the quantum model when the key parameter, γ , is set to zero. In this case ($\gamma = 0$) there are no interference effects:

$$U = \exp\left(-i \cdot \frac{\pi}{2} \cdot H_1\right) = \begin{bmatrix} U_1 & 0 \\ 0 & U_2 \end{bmatrix}, \quad (11)$$

$$M \cdot U \cdot \psi_W = \frac{1}{\sqrt{2}} \begin{bmatrix} T \cdot U_1 \\ 0 \end{bmatrix},$$

$$M \cdot U \cdot \psi_L = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ T \cdot U_2 \end{bmatrix},$$

$$(\psi_W^\dagger \cdot U \cdot M) \cdot (M \cdot U \cdot \psi_L) = 0.$$

So if we force $\gamma = 0$, then the quantum model no longer produces 'quantum like' interference effects. Instead, the choice probability for the plan is an equal weight average of the two choice probabilities produced after either winning or losing the first stage: $p(T|plan) = (.50) \cdot p(T|Win) + (.50) \cdot p(T|loss)$, where $p(T|Win)$ is defined by Equation 8 with $\gamma = 0$ and $p(T|loss)$ is defined by Equation 9 with $\gamma = 0$. This model was fit to the results in Table 1 by using only two parameters a and b for the quantum model (with $\gamma = 0$), and it produced an $R^2 = .7854$ and an *adjusted* $R^2 = .7787$ which still falls below the adjusted R^2 for the three parameter quantum model, and so the γ parameter is making a useful contribution in this application.

In summary, comparing the two key models on the basis of fitting the means, we find that the quantum model with $\gamma \neq 0$ produces an increase in adjusted R-square over the Markov model when the two models are fit to the means. However, the next section provides a Hierarchical Bayesian estimation of the key quantum parameter to determine whether or not its posterior distribution lies near zero or within a quantum regime.

4 Hierarchical Bayesian Model Comparison

4.1 Log Likelihood for Each Person

The Bayesian model estimation was computed using the 33 choice trials observed from each person. On each trial, a gamble was presented and the person made both a plan for an outcome and a final choice after observing that same outcome. For person i on trial t we observe a data pattern $X_i(t) = [x_{TT}(t), x_{TR}(t), x_{RT}(t), x_{RR}(t)]$ defined by $x_{ij}(t) = 1$ if event (i, j) occurs and otherwise zero, where TT is the event "planned to take gamble and finally did take the gamble," TR is the event "planned to take gamble but changed and finally rejected gamble," RT is the event "planned to reject the gamble but changed and finally did take the gamble" and RR is the event "planned to reject gamble and finally did reject the gamble." The data for the 33 trials from a single person is represented by the 33 tuple $X_i = [X_i(1), \dots, X_i(33)]$. Finally, the data for all 100 participants is defined by the 4×100 tuple $X = [X_1, \dots, X_{N=100}]$.

Two allow for possible dependencies between a pair of choices within a single trial, an additional memory recall parameter was included in each model. For both models, it was assumed that there is some probability m , $0 \leq m \leq 1$ that the person simply recalls and repeats the planned choice for the final choice, and there is some probability $1 - m$ that the person forgets or ignores the planned choice when making the final choice. After including this memory parameter, the prediction for each event becomes

$$\begin{aligned} p_{TT} &= p(T|plan) \cdot (m \cdot 1 + (1 - m) \cdot p(T|final)) \\ p_{TR} &= p(T|plan) \cdot (1 - m) \cdot p(R|final) \\ p_{RT} &= p(R|plan) \cdot (1 - m) \cdot p(T|final) \\ p_{RR} &= p(R|plan) \cdot (m \cdot 1 + (1 - m) \cdot p(R|final)) \end{aligned} \quad (12)$$

Using these definitions for each model, the log likelihood function for the 33 trials from a single person can be expressed as

$$\begin{aligned} \ln L(X_i(t)) &= \sum x_{jk}(t) \cdot \ln(p_{jk}) \\ \ln L(X_i) &= \sum_{i=1}^{33} \ln L(X_i(t)). \end{aligned} \quad (13)$$

The predictions p_{jk} used in the formulas shown above depend on the four model parameters $\theta_i = [a_i, b_i, m_i, \gamma_i]$ for person i . Therefore, the likelihood of the data for person i given the model parameters is then equal to $L(X_i|\theta_i) = \exp(\ln L(X_i))$.

4.2 Grid Analysis of Log Likelihood Function

Each model has four parameters $\theta_i = (a, b, m, \gamma)$, a risk aversion parameter, a loss aversion parameter, a memory parameter, and a choice model parameter. The first three parameters were common across both models and they only differ with respect to the fourth parameter. We used a fine grid of 21 points per parameter.

$$a \in [.400, .45, \dots, .85, .90, .95, \dots, 1.35, 1.40], \quad (14)$$

$$b \in [.50, .60, \dots, 1.40, 1.50, 1.60, \dots, 2.40, 2.50], \quad (15)$$

$$m \in [.00, .05, \dots, .45, .500, .55, \dots, .95, 1.00], \quad (16)$$

$$\gamma \in [-5.00, -4.5, \dots, -.5, 0.0, .5, \dots, 4.5, 5.00] \text{ (quantum)}. \quad (17)$$

This grid generated 21^4 combinations, and we evaluated the log likelihood function for each model at each combination. These ranges were chosen on the basis of past fits of these models. The risk aversion parameter ranges from risk aversion

to risk seeking; the loss aversion parameter ranges across loss insensitivity to loss sensitivity; and the memory parameter ranges from no recall to perfect recall. The key γ parameter ranges from positive to negative values for the quantum model. Define $[a_i, b_i, m_i, \gamma_i] = [\theta_{i1}, \theta_{i2}, \theta_{i3}, \theta_{i4}] = \theta_i$ as the 4-tuple of parameters from a single person i , and define $\theta = [\theta_1, \dots, \theta_N]$ as the $4 \cdot N$ tuple of the four parameters for $N = 100$ participants.

4.3 Hierarchical Parameters

The hierarchical parameters are used to determine the distribution of θ_i across individuals. Define $\pi = [\pi_1, \pi_2, \pi_3, \pi_4]$ as a 4-tuple containing four hierarchical parameters, where π_j is the hierarchical parameter used to determine the distribution of θ_{ij} across the individuals i . Each hierarchical parameter was evaluated by a grid of 19 points $\pi_j \in [.05, .10, \dots, .90, .95]$ which generated a grid of 19^4 combinations.

Define $r(\pi)$ as the prior distribution over the hierarchical parameters. We assumed an independent uniform so that $r(\pi) = 19^{-4}$. Define $q(\theta_i|\pi_i)$ as the prior distribution over model parameter θ_i given the hierarchical parameter π_i . For this prior we assumed an independent binomial distribution across the 21 values of each model parameter

$$q(\theta_i|\pi) = \prod_{j=1}^4 q(\theta_{ij}|\pi_j), \quad q(\theta_{ij} = \theta_k|\pi_j) = \binom{21}{k} \cdot \pi_j^k \cdot (1 - \pi_j)^{21-k}. \quad (18)$$

The joint distribution of data and parameters then equals

$$p(\pi, \theta, X) = r(\pi) \cdot \prod_{i=1}^{N=100} q(\theta_i|\pi) \cdot L(X_i|\theta_i). \quad (19)$$

We marginalize over θ to obtain the joint distribution of hierarchical parameters and data

$$p(\pi, X) = \sum_{\theta} p(\pi, \theta, X). \quad (20)$$

Finally, we obtain the posterior distribution over the hierarchical parameters

$$p(\pi|X) = \frac{p(\pi, X)}{\sum_{\pi} p(\pi, X)}. \quad (21)$$

The posterior distribution for each hierarchical parameter is plotted in Figure 1 shown below. The top left distribution indicates that the risk aversion hierarchical parameter distribution is located below .50, which implies that the mean of the risk aversion parameter equals .6518, indicating somewhat strong risk aversion, which is a common finding in the literature. The top right distribution indicates that the loss aversion hierarchical parameter distribution is located

above .50, which implies that the mean of the loss aversion equals 1.97, higher sensitivity to losses, which is also a common finding in the literature. The bottom left distribution indicates that the hierarchical memory parameter is slightly above .50, which implies that the mean of the memory parameter equals .5932, so that a little more than half the time people were simply recalling their previous choices. The bottom right distribution shows the hierarchical distribution for the key quantum parameter. According to the Markov model, this should be located around .50 to produce a mean value equal to zero. Contrary to this expectation, the entire distribution lies below .50, which implies a mean value equal to -2.67 .

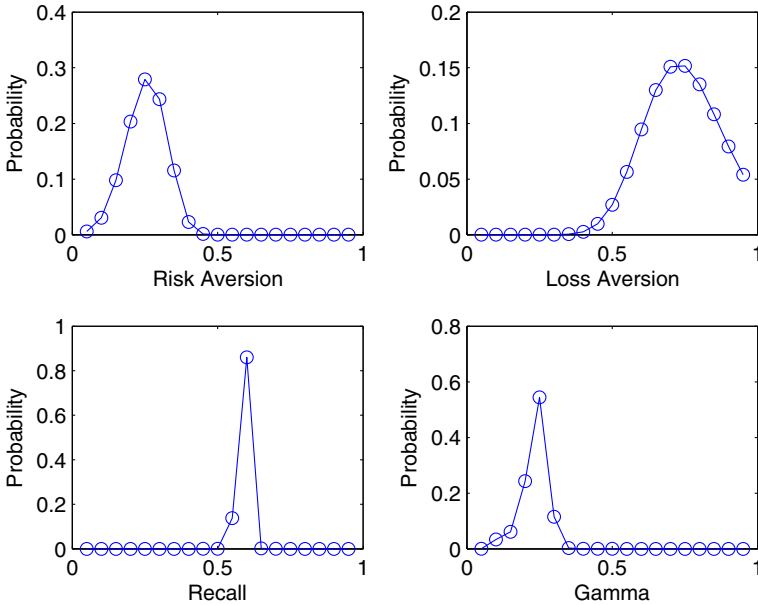


Fig. 1. Posterior distribution of hierarchical parameters of the quantum decision model

5 Conclusions

This paper presents the first hierarchical Bayesian estimation of the parameters used in a quantum decision model. A classic Markov model is a special case of the quantum model when the key quantum parameter is zero. The posterior distribution of the key quantum parameter was entirely below the value expected by the Markov model, providing strong evidence that the parameter lies within a quantum regime. Of course, it is much too soon to conclude that the quantum model is always superior to a Markov model. The models need to be compared

using other data sets from various other experiments. Even within the same data set, various other prior distributions need to be examined. Further, with the two stage gambling paradigm, the learned model could be used to predict the next result with cross-validation methods. But the surprising lesson learned from this model comparison exercise was that contrary to expectations, there is clear evidence for the quantum model parameter.

References

1. Barkan, R., Busemeyer, J.R.: Modeling Dynamic Inconsistency with a Changing Reference Point. *Journal of Behavioral Decision Making* 16, 235–255 (2003)
2. Berger, J.O.: *Statistical Decision Theory and Bayesian Analysis*. Springer, New York (1985)
3. Busemeyer, J.R., Wang, Z., Lambert-Mogiliansky, A.: Comparison of Markov and Quantum Models of Decision Making. *Journal of Mathematical Psychology* 53(5), 423–433 (2009)
4. Khrennikov, A.Y., Haven, E.: Quantum Mechanics and Violations of the Sure Thing Principle: The use of Probability Interference and Other Concepts. *Journal of Mathematical Psychology* 53(5), 378–388 (2009)
5. Lambert-Mogiliansky, A., Zamir, S., Zwirn, H.: Type Indeterminacy: A Model of the ‘KT’ (Kahneman-Tversky)-Man. *Journal of Mathematical Psychology* 53(5), 349–361 (2009)
6. Pothos, E., Busemeyer, J.: A Quantum Probability Explanation for Violations of ‘Rational’ Decision Theory. *Proceedings of the Royal Society B* 276(1165), 2171–2178 (2009)
7. Tversky, A., Shafir, E.: The Disjunction Effect in Choice Under Uncertainty. *Psychological Science* 3, 305–309 (1992)
8. Yukalov, V., Sornette, D.: Decision Theory with Prospect Interference and Entanglement. *Theory and Decision* 70, 283–328 (2010)
9. Yukalov, V., Sornette, D.: Processing Information in Quantum Decision Theory. *Entropy* 11(4), 1073–1120 (2009)