Quantum coin hedging, and a counter measure

Maor Ganz[1] and Or Sattath[1,2]

[1]The Hebrew University
[2]MIT

Abstract

A quantum board game is a multi-round protocol between a single quantum player against the quantum board. Molina and Watrous [5] discovered quantum hedging. They gave an example for perfect quantum hedging: a board game with winning probability < 1, such that the player can win with certainty at least 1-out-of-2 quantum board games played in parallel. Here we show that perfect quantum hedging occurs in a cryptographic protocol – quantum coin flipping. For this reason, when cryptographic protocols are composed, hedging may introduce serious challenges into their analysis.

We also show that hedging cannot occur when playing two-outcome board games in sequence. This is done by showing a formula for the value of sequential two-outcome board games, which depends only on the optimal value of a single board game; this formula applies in a more general setting, in which hedging is only a special case.

The full version of this QCrypt extended abstract is available on `http://arxiv.org/abs/1703.03887`.

Quantum board games A quantum board game is a special type of an interactive quantum protocol. The protocol involves two parties: the player and the board. The board implements the rules of the game: in each round i of the protocol, applies some quantum operation O_i, sends a quantum message to the player, which can apply any operation it wants, and send a quantum message back to the board. At the final round of the board game, the board applies a two outcome measurement, which determines whether the player won or lost. We assume that the player knows the rules of the board game (the length of the messages, the operations O_i and the two outcome measurement). The player has the freedom to decide on his strategy – the protocol does not specify what

the player should do in each round; the only constraint posed on the player is that it must send a message of an appropriate length, as expected by the board.

Perfect hedging Molina and Watrous showed that hedging is possible in quantum board games[5]. Prefect hedging is best explained by an example: there exists a quantum board game for which no strategy can win with certainty, but it is possible for a player to guarantee winning 1-out-of-2 independent quantum board games, which are played in parallel. A formal definition of hedging is given in the full version, but for now, one can think of that example. In a follow up work, Arunachalam, Molina and Russo [2] analyzed a family of quantum board games, and showed a necessary and sufficient condition so that the player can win with certainty in at least 1-out-of-n board games. As discussed later, quantum hedging is known to be a purely quantum phenomenon.

One example where Hedging becomes relevant is when reducing the error (soundness) probability of quantum interactive proof protocols such as QIP(2): since the optimal strategy for winning t-out-of-n parallel repetitions is not necessarily an independent strategy, only Markov bound (and not the Chernoff bound) can be used to show soundness [3]. These aspects resembles the behavior that occurs in the setting of Raz's (classical) parallel repetition theorem[6]; the differences are that in the classical setting there are two players who want to win all board games, whereas in our setting, there is a single player, who wants to win at least t-out-of-n board games.

Coin flipping Quantum coin flipping is a two player cryptographic protocol which simulates a balanced coin flip. When Alice and Bob are honest, they both agree on the outcome, which is uniform on $\{0, 1\}$. Coin flipping comes in two flavors: Strong and weak. Perhaps the most intuitive one is *weak coin flipping*, which is the variant we will concentrate. Each player has an opposite desirable outcome: 0 implies that Alice wins, and 1 implies that Bob wins. An important parameter is the optimal winning probability for a cheating player against an honest player. In weak coin flipping we denote them by P_A and P_B. We define $P^* = \max\{P_A, P_B\}$ – the maximum cheating probability of both players. Mochon showed that there are families of weak coin flipping protocols for which P^* converges to $\frac{1}{2}$ ([4], see also [1]).

Is it possible to hedge in quantum coin flips? We show a weak coin flipping protocol where $P^* = \cos^2(\frac{\pi}{8})$, yet a cheating Bob can guarantee winning in at least 1-out-of-2 board games played in parallel. The coin flipping protocol is fairly simple: Alice sends Bob half of an EPR pair. Bob randomly chooses either the standard or Hadamard basis, and sends his choice to Alice. They both measure their halves of the EPR pair in the basis chosen by Bob, and the measurement outcome is the outcome of the protocol (where the

result $|+\rangle$ ($|-\rangle$) is interpreted as 0 (1)). To hedge, Bob plays the two games in parallel. Instead of picking the bases to measure at random, he measures the two halves of the EPR pair in some carefully chosen entangled basis, and sends the outcome to Alice. This strategy guarantees that Bob wins in *exactly* 1-out-of-2 games. Hedging is not necessarily symmetric: in this example, Alice cannot hedge at all (her optimal winning probability to win at least 1-out-of-2 coin flips played in parallel can be achieved by an independent strategy).

Avoiding hedging through sequential repetition Consider a cryptographic quantum protocol, which involves several uses of quantum two-outcome board games. For example, the protocol may use several occurrences of quantum coin flips played in parallel. As we have seen, the possibility of hedging makes it hard to analyze the resulting protocol, by simply analyzing each of the board games in it. We show that quantum hedging cannot happen when the two-outcome board games are played in sequence, even if the players are computationally unbounded.

We give a more generalized formulation for sequential board games. Suppose the player's utility for the outcome vector $a = (a_1, \ldots, a_n)$ is given by some target function $t(a)$, and the players goal is to maximize $\mathbb{E}[t(a)]$ over all possible strategies. We show that this maximal value is fully determined by the properties of each board game, and does *not* require an analysis of the entire system, which is the case when playing in parallel.

The authors are not aware of previous claims of that sort. The intuition for the proof is fairly simple and arguably not very surprising: if it is possible to hedge n games, then by simulating the board in the first game, and conditioning on some good event, allows the player to hedge $n-1$ games. But since hedging cannot occur in one game, we reach a contradiction.

References

[1] D. Aharonov, A. Chailloux, M. Ganz, I. Kerenidis, and L. Magnin. A simpler proof of the existence of quantum weak coin flipping with arbitrarily small bias. *SIAM J. Comput.*, 45(3):633–679, 2016.

[2] S. Arunachalam, A. Molina, and V. Russo. Quantum hedging in two-round prover-verifier interactions. *arxiv preprint arxiv:1310.7954*, 2013.

[3] R. Jain, S. Upadhyay, and J. Watrous. Two-message quantum interactive proofs are in PSPACE. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 534–543, 2009.

[4] C. Mochon. Quantum weak coin flipping with arbitrarily small bias. *arxiv preprint arxiv:0711.4114*, 2007.

[5] A. Molina and J. Watrous. Hedging bets with correlated quantum strategies. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, pages 2614–2629, 2012.

[6] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.