Computational Results for Bidding Games

Srinivas Arun

August 2021

1 Preliminaries

Definition 1.1. A **Discrete Richman Game** is a combinatorial game in which

- Each player starts with some number of tokens.
- One player starts with the **tie-breaking advantage**.
- On each turn, the players simultaneously reveal bids.
  - If the bids are unequal, then the player with the larger bid pays the other player and moves.
  - If there is a tie, the player with the tie-breaking advantage can either use it and transfer it to the other player, or choose not to use it.

We denote a Discrete Richman Game by either $G(a^*, b)$ or $G(a, b^*)$, where a and b are the starting amounts of money, and $*$ is the tie-breaking advantage.

When working with bidding games in general, we usually represent a game by a directed graph, where vertices are subpositions and there is a directed edge $u \to v$ if some player can move from u to b. We designate two vertices v_l and v_r representing positions where Left and Right win immediately (if there are several, we can just treat them as one).

Definition 1.2. A **Classical Richman Game** is the same as a Discrete Richman Game, except 1) ties are broken randomly, and 2) players can withdraw any real amount of money.

In practice, Classical Richman Games are impossible to play. However, they avoid the somewhat unnatural concept of a tie-breaking advantage. Moreover, we can describe winning positions with just one number.

Definition 1.3. Let G be a Classical Richman game. The **Richman Value** $R(G)$ is smallest real number such that Left wins if her starting proportion of money is greater than $R(G)$, and loses if her starting proportion is less than $R(G)$.

Due to the following property, it turns out the concept of a Richman Value is not that important.

Theorem 1.4

$R(G) = 1 - P(G)$, where $P(G)$ is the probability that Left wins when turns are decided randomly.[1]

Proof. Induct on the birthday of G; the base case is clear.

If turns are decided randomly, players will obviously move to the position that gives them the highest probability of winning. This implies that

$$P(G) = \frac{1}{2}\left(\max_{x^L \in G^L}(P(x^L)) + \max_{x^R \in G^R}(P(x^R))\right).$$

Thus, by the inductive hypothesis, it suffices to show that

$$R(G) = \frac{1}{2}\left(\max_{x^L \in G^L}(R(x^L)) + \max_{x^R \in G^R}(R(x^R))\right) = \frac{1}{2}(R_A(G) + R_B(G)).$$

Suppose the total amount of money is 1, and Left starts out with $\frac{1}{2}(R_A(G) + R_B(G)) + \epsilon$ for some $\epsilon > 0$. Then, Left should bid $\frac{1}{2}(R_A(G) - R_A(G))$.

- If Left wins the bid, she moves to a position with Richman value $R_B(G)$, where she has $R_B(G) + \epsilon$.
- If Left loses the bid, Right moves to a position with Richman value $R_A(G)$, where Left has at least $R_A(G) + \epsilon$.

In either case, Left wins by the inductive hypothesis. Similarly, if Left starts out with less than $\frac{1}{2}(R_A(G) + R_B(G))$, then Right wins. Hence, $\frac{1}{2}(R_A(G) + R_B(G))$ is the Richman Value of G, so we are done. □

2 Computational Complexity

For Classical Richman Games, periodicity is trivially true; we can scale down the number of tokens without changing the outcome. Not surprisingly, a similar result holds for their discrete counterparts.

Theorem 2.1

For all Discrete Richman Games G, there exist constants c, d such that $G(a, b^*) = G(a + c, b + d^*)$ for all a, b.[1]

Proof. Let n be a sufficiently divisible positive integer. We claim that $c = n \cdot R(G), d = n \cdot (1 - R(G))$ works. It suffices to show that if Left wins $G(a, b^*)$, then she also wins $G(a + c, b + d^*)$.

Suppose Left's optimal first move in $G(a, b^*)$ is bidding x. Then, Left should start by bidding $x + n \cdot \frac{R_A(G) - R_B(G)}{2}$.

- If Left wins the bid, she moves to $G_L(a - x + n \cdot R_B(G), b + x + n(1 - R_B(G)))$.
- If Right wins the bid, he moves to $G_R(a + x + n \cdot R_A(G), b - x + n(1 - R_A(G)))$ (or a position more favorable to Left).

By the inductive hypothesis, both scenarios are at least as good for Left as $G(a - x, b + x)$, so Left wins. □

Corollary 2.2

For all Discrete Richman Games G, we can determine the outcome of the game for every possible pair of starting token amounts in constant time, albeit with a grotesquely large constant factor.

We can halve this constant factor with the following observation.

Theorem 2.3

In a Discrete Richman Game, the tie-breaking advantage has positive value but is worth less than a single token.[1]

Proof. Suppose Left wins $G(a, b^*)$. To prove that the tie-breaking advantage has positive value, it suffices to show that Left wins $G(a^*, b)$.

For convenience, suppose $A = G(a, b^*)$ and $B = G(a^*, b)$ are being played simultaneously, and we are playing for Left. Since Right does not have a winning strategy in $G(a, b^*)$, we may assume that we control Right's moves in A.

By copying Right's moves in B over to his moves in A, we can ensure that A and B progress identically until the first point at which the bids are tied. We can now use our tie-breaking advantage in A, and force Right NOT to use his tie-breaking advantage in B. Now both games are at the exact same position, and we can copy our winning strategy in A to win B. Thus, Left wins $G(a^*, b)$.

Next, suppose Left wins $G(c^*, d)$. To prove that the tie-breaking advantage is worth less than a single token, it suffices to show that Left wins $G(c + 1, d - 1^*)$. As before, assume $A = G(c^*, d)$ and $B = G(c + 1, d - 1^*)$ are being played simultaneously, we are playing for Left, and we can control Right's moves in A.

Assume WLOG that the first moves in A and B are different. Suppose Left's optimal first move in A is to bid x and use the tie-breaking advantage if necessary. Then, Left should bid $x + 1$ in B.

- If Left wins the bid, she moves B to $G_L(c - x, d + x^*)$. By making Right bid x in A, the same position is obtained from A. Now Left can copy her strategy in A to win B.
- If Right uses the tie-breaking advantage to win the bid, then Left can make Right bid $x + 1$ in A. Now Right pays $x + 1$ to move in both games, contradicting our assumption.
- If Right bids $x + y$ for $y \geq 2$, he wins the bid and moves B to $G_R(c + x + y + 1, d - x - y - 1^*)$. By making Right bid $x + 1$ in A, the position $G_R(c + x + 1^*, d - x - 1)$ is obtained from A. Now we are faced with the same problem we began with, so Left can repeatedly apply the above strategy to win.

Hence, Left has a winning strategy in $G(c + 1, d - 1^*)$. □

It turns out that computing outcomes of Classical Richman Games is much harder. However, we can obtain results for some special cases.

Theorem 2.4

If the graph associated with a game G is finite and acyclic, the Richman Values of all vertices can be calculated in linear time in the number of vertices and edges.[2]

Proof. Assign probabilities to vertices in reverse topological order (it is well-known that topological sort can be done in linear time). For any vertex v, we can calculate $P(v)$ by scanning through the probabilities of its neighbors (which have already been assigned). Now we are done by Theorem 0.5. □

When the graph has cycles, we have to be much more careful, as the equations we know for Richman values refer to themselves. In the following case, we can just directly solve the ensuing system of equations.

Theorem 2.5

If there are at most two moves from any subposition of a Classical Richman Game G, then the Richman Values of all subpositions can be calculated in polynomial time.[2]

Another result for graphs with cycles is below; the proof is a very nice local argument.

Theorem 2.6

If the graph associated with a Classical Richman Game G is finite and *undirected*, then the Richman values of all vertices can be calculated in polynomial time.[2]

Proof. We define $X = \{V, E\}$, the set of vertices and edges that we have processed so far. Initialize $R(v_l) = 0, R(v_r) = 1$, and $X = \{\{v_l, v_r\}, \{\}\}$. We will maintain the invariant that if an edge is in E, the vertices it connects are in V.

At each step of the algorithm, consider the set of paths $v_0 v_1 v_2 \dots v_n$ such that $v_0, v_n \in V$ and $v_1, \dots, v_{n-1} \notin V$, where we assume WLOG that $R(v_0) < R(v_n)$. Of these paths, choose the one such that $s = \frac{R(v_n) - R(v_0)}{n}$, which we call the *slope* of the path, is maximal. Now assign $R(v_i) = \frac{i \cdot R(v_n) + (n-i) \cdot R(v_0)}{n}$ for all i, and add the path to X.

Call a vertex *good* if its Richman cost is the average of the minimum and maximum Richman values of its neighbors in X.

- $R(v_i) = \frac{R(v_{i-1}) + R(v_{i+1})}{2}$ for $i = 1, \dots, n-1$, so $v_1, \dots, v_{n-1}$ are good.
- The neighbors of all vertices other than $v_0, \dots, v_n$ haven't changed, so they remain good.

Finally, we need to show that v_0, v_n are good. To show that v_0 is good, it suffices to show that s is at least as large in magnitude as the slopes of edges adjacent to v_0. This is equivalent to showing that the slopes of the chosen paths are nonincreasing.

Suppose, for the sake of contradiction, there is some path $u_0 u_1 \dots u_k$ with $u_0, u_k \in V$ and $u_1, \dots, u_{k-1} \notin V$ with slope $s' > s$.

- If $u_0 = v_i$ and $u_k = v_j$, then the path $v_0 \dots v_i u_1 \dots u_{k-1} v_j \dots v_n$ has slope $\frac{1}{i+k+n-j}(is + ks' + (n-j)s) > s$, contradicting our maximality assumption.
- If $u_0 = v_i$ and $u_k \notin \{v_0, \dots, v_n\}$, then the path $v_0 \dots v_i u_1 \dots u_k$ has slope $\frac{1}{i+k}(is + ks') > s$, contradicting our maximality assumption.
- If $u_0, u_k \notin \{v_0, \dots, v_n\}$, then $u_0 \dots u_k$ itself was a candidate in the first step of our process, and thus contradicts our maximality assumption.

Having reached a contradiction in all cases, we conclude that v_0 and v_n are good. This shows that our Richman Value labeling is valid for our new set X.

Repeating our algorithm as many times as possible, we will eventually reach a point at which no paths, as described, exist. Consider a vertex $v \notin V$.

- If v is connected to $x, y \in V$, then $x \to v \to y$ is a suitable path to continue our algorithm, contradiction.
- If v is not connected to any vertex in V, then it must be in a different component, contradiction.

Thus, v is connected to exactly one vertex v' in V. Now we can just let $R(v) = R(v')$. Repeating for all vertices not in V, we are done. □

3 Hex

Next, we turn our attention from computing outcomes quickly to determining good moves quickly. We will focus on the game HEX, where two players take turns coloring hexagons in an 11×11 rhombus, and the player who creates a path between their respective pair of opposite sides wins. For example, an end position where Left has won is shown below. For now, we will ignore bidding, and just assume turns are determined arbitrarily.

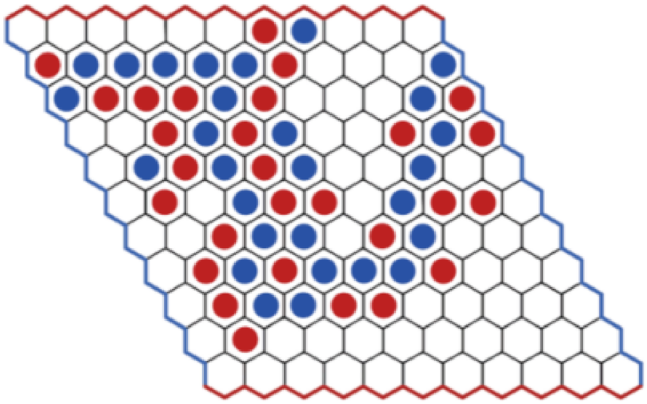

We have the following proposition; it is quite intuitive, so we have omitted the proof.

Proposition 3.1

At any subposition of the game, Left's set of optimal moves is the same as Right's set of optimal moves.[3]

When the board is completely filled, call a hexagon *pivotal* if changing its color would change the outcome of the game.

Theorem 3.2

At any given point in the game, the optimal move (for both Left and Right) is to play in a hexagon with the highest chance of being pivotal, over all possible colorings of the remaining hexagons.[4]

Proof. WLOG assume Left starts, and suppose she colors a hexagon H out of the set S of all hexagons. Since turns are chosen randomly and both players have the same strategy, the remaining turns will color $S \setminus \{H\}$ randomly. Therefore, H should have the largest possible chance of being pivotal over all random colorings of $S \setminus \{H\}$. □

For the random-turn game, this is the best we can do. We can check a large number of random colorings of the remaining board, and determine which hexagons are most often pivotal. You can play against this strategy at Hexamania.

Fortunately, we can apply a similar strategy for the discrete bidding game. We will need the following theorem.

Theorem 3.3

The probability that a hexagon is *not* pivotal is twice the probability that it will have the losing color.[3]

Proof. Suppose a hexagon H is in the losing color. If we switch the color of H, then H is in the winning color and not pivotal. Hence, the probabilities of these two events are equal.

However, if H is losing, then it cannot be pivotal. Thus,

$$P(\text{not pivotal}) = P(\text{winning and not pivotal}) + P(\text{losing and not pivotal}) = 2P(\text{losing}),$$

as desired. □

This shows that the optimal move is to play in the hexagon with the least chance of being in the losing color. But how should we bid?

Theorem 3.4
The optimal bid is the floor of $\frac{1}{2}P(\text{pivotal})$ times the total number of tokens.[3]

Proof. The probability that Left wins if she makes the first move is $1 - P(\text{losing})$ and the probability that Right wins if he makes the first move is $P(\text{losing})$. Therefore, by our work in Theorem 1.4, the best bid is $\frac{1}{2}(1 - 2P(\text{losing})) = \frac{1}{2}P(\text{pivotal})$ of the total amount of resources. □

Now, to compute optimal moves and bids, we can just check a large number of colorings of the remaining hexagons on the board (say ≈ 500000), chosen uniformly at random, on each turn. One algorithm that does this is able to beats human opponents consistently; to see the code, you can contact the authors of [3].

4 Future Work

There are several interesting variants on discrete and classical Richman games, which behave very differently. In general, very little is known about these variants. Some examples include

- Poorman bidding, in which players pay the bank instead of the other player
- Taxman bidding, a generalization of Poorman and Richman bidding, in which players play a certain portion to the bank and the rest to the other player
- All-pay bidding, in which all players pay the bank the amount of their bid
- Scoring, in which players bid to increase their score and the highest score wins
- Infinite-duration games with the same rules

In addition, analyzing well-known games like Nim and Chess under auction player would be interesting and original.

Bibliography

[1] Mike Develin and Sam Payne. Discrete bidding games, 2010.

[2] Andrew J Lazarus, Daniel E Loeb, James G Propp, Walter R Stromquist, and Daniel H Ullman. Combinatorial games under auction play, Sep 1997.

[3] Sam Payne and Elina Robeva. Artificial intelligence for bidding hex, 2008.

[4] Yuval Peres, Oded Schramm, Scott Sheffield, and David B. Wilson. Random-turn hex and other selection games, Apr 2006.