

Discrete bidding games

Mike Develin

American Institute of Mathematics
360 Portage Ave., Palo Alto, CA 94306
develin@post.harvard.edu

Sam Payne*

Stanford University, Dept. of Mathematics
450 Serra Mall, Stanford, CA 94305
spayne@stanford.edu

Submitted: Feb 28, 2008; Accepted: May 26, 2010; Published: Jun 7, 2010
Mathematics Subject Classification: 91A46, 91B26, 91A60

Abstract

We study variations on combinatorial games in which, instead of alternating moves, the players bid with discrete bidding chips for the right to determine who moves next. We consider both symmetric and partisan games, and explore differences between discrete bidding games and *Richman games*, which allow real-valued bidding. Unlike Richman games, discrete bidding game variations of many familiar games, such as chess, Connect Four, and even Tic-Tac-Toe, are suitable for recreational play. We also present an analysis of Tic-Tac-Toe for both discrete and real-valued bidding.

Contents

1	Introduction	2
1.1	A game of bidding Tic-Tac-Toe	4
1.2	A game of bidding chess.	5
2	Preliminaries	7
2.1	Game model	7
2.2	Bidding	7
2.3	Tie-breaking	7
3	General theory	8
3.1	Value of the tie-breaking advantage	8
3.2	Using the tie-breaking advantage	9
3.3	Classical Richman calculus	10
3.4	Discrete Richman calculus	11

*Supported by the Clay Mathematics Institute.

3.5	Discrete bidding with large numbers of chips	14
3.6	Periodicity	16
4	Examples	17
4.1	Tug o' War	18
4.2	Ultimatum	19
5	A partial order on games	21
6	Bidding Tic-Tac-Toe	23
6.1	Optimal moves	23
6.2	Chip tables	32
7	Appendix: other tie-breaking methods	38

1 Introduction

Imagine playing your favorite two-player game, such as Tic-Tac-Toe, Connect Four, or chess, but instead of alternating moves you bid against your opponent for the right to decide who moves next. For instance, you might play a game of bidding chess in which you and your opponent each start with one hundred bidding chips. If you bid twelve for the first move, and your opponent bids ten, then you give twelve chips to your opponent and make the first move. Now you have eighty-eight chips and your opponent has one hundred and twelve, and you bid for the second move...

Similar bidding games were studied by David Richman in the late 1980s. In Richman's theory, as developed after Richman's death in [LLPU96, LLP⁺99], a player may bid any nonnegative real number up to his current supply of bidding resources. The player making the highest bid gives the amount of that bid to the other player and makes the next move in the game. If the bids are tied, then a coin flip determines which player wins the bid. The goal is always to make a winning move in the game; bidding resources have no value after the game ends. The original Richman theory requires that the games be *symmetric*, with all legal moves available to both players, to avoid the possibility of *zugzwang*, positions where neither player wants to make the next move. The theory of these real-valued bidding games, now known as Richman games, is simple and elegant with surprising connections to random turn games. The recreational games, such as chess, that motivated the work presented here, are partisan rather than symmetric, and it is sometimes desirable to force your opponent to move rather than to make a move yourself. However, the basic results and arguments of Richman game theory go through unchanged for partisan games, in spite of the remarks in [LLP⁺99, p. 260], provided that one allows the winner of the bid either to move or to force his opponent to move, at his pleasure. For the remainder of the paper, we refer to these possibly partisan real-valued bidding games as *Richman games*.

Say Alice and Bob are playing a Richman game, whose underlying combinatorial game is G . Then there is a critical threshold $R(G)$, sometimes called the *Richman value*

of the game, such that Alice has a winning strategy if her proportion of the total bidding resources is greater than $R(G)$, and she does not have a winning strategy if her proportion of the bidding resources is less than $R(G)$. If her proportion of the bidding resources is exactly $R(G)$, then the outcome may depend on coin flips.

The critical thresholds $R(G)$ have two key properties, as follows. We say that G is finite if there are only finitely many possible positions in the game, and we write \overline{G} be the game that is just like G except that Alice and Bob have exchanged roles. Let $P(G)$ be the probability that Alice can win G at random-turn play, where the player who makes each move is determined by the toss of a fair coin, assuming optimal play.

1. If G is finite, then $R(G)$ is rational and equal to $1 - R(\overline{G})$.
2. For any G , $R(G)$ is equal to $1 - P(G)$.

The surprising part of (1) is that, if G is modeled on a finite graph, which may contain many directed cycles, there is never a range of distributions of bidding resources in which both Alice and Bob can prolong the game indefinitely and force a draw. On the other hand, for infinite games $R(G)$ can be any real number between zero and one [LLP⁺99, p. 256], and $1 - R(\overline{G})$ can be any real number between zero and $R(G)$; in particular, the Richman threshold may be irrational and there may be an arbitrarily large range in which both players can force a draw. The connection with random turn games given by (2) is especially intriguing given recent work connecting random turn selection games with conformal geometry and ideas from statistical mechanics [PSSW07].

The discrete bidding variations on games that we study here arose through recreational play, as a way to add spice and interest to old-fashioned two player games such as chess and Tic-Tac-Toe. The real valued bidding and symmetric play in Richman's original theory are mathematically convenient, but poorly suited for recreational play, since most recreational games are partisan and no one wants to keep track of bids like $e^{\sqrt{\pi}} + \log 17$. Bidding with a relatively small number of discrete chips, on the other hand, is easy to implement recreationally and leads to interesting subtleties. For instance, ties happen frequently with discrete bidding with small numbers of chips, so the tie-breaking method is especially important. To avoid the element of chance in flipping coins, we introduce a deterministic tie-breaking method, which we call the *tie-breaking advantage*. If the bids are tied, the player who has the tie-breaking advantage has the choice either to declare himself the winner of the bid and give the tie-breaking advantage to the other player, or declare the other player the winner of the bid and keep the tie-breaking advantage. See Section 2.3 for more details. As mentioned earlier, partisan games still behave well under bidding variations provided that the winner of the bid has the option of forcing the other player to move in zugzwang positions.

Other natural versions of bidding in combinatorial game play are possible, and some have been studied fruitfully. The most prominent example is Berlekamp's "economist's view of combinatorial games" [Ber96], which is closely related to Conway's theory of thermography [Con76] and has led to important advances in understanding Go endgames.

Since this paper was written, Bidding Chess has achieved some popularity among fans of Chess variations [Bea08a, Bea08b]. Also, bidding versions of Tic-Tac-Toe and Hex have been developed for recreational play online, by Jay Bhat and Deyan Simeonov. Readers are warmly invited to play against the computer at

<http://bttt.bidding-games.com/online/>

and

<http://hex.bidding-games.com/online/>,

and to challenge friends through Facebook at

<http://apps.facebook.com/biddingttt> and <http://apps.facebook.com/biddinghex>.

The artificial intelligence for the computer opponent in Bidding Hex is based on the analysis of Random-Turn Hex in [PSSW07] and connections between random turn games and Richman games, and is presented in detail in [PR08]. Although we cannot prove that the algorithm converges to an optimal or near-optimal strategy, it has been overwhelmingly effective against human opponents.

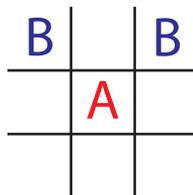
1.1 A game of bidding Tic-Tac-Toe

We conclude the introduction with two examples of sample bidding games. First, here is a game of Tic-Tac-Toe in which each player starts with four bidding chips, and Alice starts with tie-breaking advantage. In Tic-Tac-Toe there is no zugzwang, so the players are simply bidding for the right to move.

First move. Both players bid one for the first move, and Alice chooses to use the tie-breaking advantage, placing a red **A** in the center of the board.

Second move. Now Alice has three chips, and Bob has five chips plus the tie-breaking advantage. Once again, both players bid one. Bob uses the tie-breaking advantage and places a blue **B** in the upper-left corner.

Third move. Now Alice has four chips plus the tie-breaking advantage, while Bob has four chips. Alice bids two, and Bob also bids two. This time Alice decides to keep the tie-breaking advantage, and lets Bob make the move. Bob places a blue **B** in the upper-right corner, threatening to make three in a row across the top. The position after three moves is shown in the following figure.



Fourth move. Now Alice has six chips plus the tie-breaking advantage, and Bob has two chips. Bob is one move away from winning, so he bets everything, and Alice must give him two chips, plus the tie-breaking advantage, to put a red **A** in the top center and stop him.

Conclusion. Now Alice has four chips, and Bob has four chips plus the tie-breaking advantage. Alice is one move away from victory and bets everything, so Bob must also bet everything, plus use the tie-breaking advantage, to move bottom center and stop her. Now Alice has all eight chips, plus the tie-breaking advantage, and she coolly hands over the tie-breaking advantage, followed by a single chip, as she moves center left and then center right to win the game.

Normal Tic-Tac-Toe tends to end in a draw, and Alice and Bob started with equal numbers of chips, so it seems that the game should have ended in a draw if both players played well. But Alice won decisively. *What did Bob do wrong?*

1.2 A game of bidding chess.

Here we present an actual game of bidding chess, played in the common room of the mathematics department at UC Berkeley, in October 2006. Names have been changed for reasons the reader may imagine.

Alice and Bob each start with one hundred bidding chips. Alice offers Bob the tie-breaking advantage, but he declines. Alice shrugs, accepts the tie-breaking advantage, and starts pondering the value of the first move. Alice is playing black, and Bob is playing white.

First move. After a few minutes of thought on both sides, Alice bids twelve and Bob bids thirteen for the first move. So Bob wins the bid, and moves his knight to c6. Now Alice has one hundred and thirteen chips and the tiebreaking advantage, and Bob has eighty seven chips.

Second move. Alice figures that the second move must be worth no more than the first, since it would be foolish to bid more than thirteen and end up in a symmetric position with fewer chips than Bob. She decides to bid eleven, which seems safe, and Bob bids eleven as well. Alice chooses to use the tie-breaking advantage and moves her pawn to e3. Bob, who played chess competitively as a teenager, is puzzled by this conservative opening move.

Third move. Now Alice has one hundred and two chips, and Bob has ninety eight and the tie-breaking advantage. Sensing the conservative tone, Bob decides to bid nine. He is somewhat surprised when Alice bids fifteen. Alice moves her bishop to c4. The resulting position is shown below.



Fourth move. Now Alice has eighty seven chips, while Bob has one hundred and thirteen and the tie-breaking advantage. Since Alice won the last move for fifteen and started an attack that he would like to counter, Bob bids fifteen for the next move. Alice bids twenty two, and takes the pawn at f7. Bob realizes with some dismay that he must win the next move to prevent Alice from taking his king, so he bids sixty five, to match Alice's total chip count, and uses the tie-breaking advantage to win the bid and take Alice's bishop with his king. The resulting position after five moves is shown here.



Conclusion. Now Bob has a material advantage, but Alice has one hundred and thirty chips, plus the tie-breaking advantage. Pondering the board, Bob realizes that if Alice wins the bid for less than thirty, then she can move her queen out to f3 to threaten his king, and then bid everything to win the next move and take his king. So Bob bids thirty, winning over Alice's bid of twenty-five. Bob moves his knight to f6, to block the f-column, but Alice can still threaten his king by moving her queen to h4. Since Alice has enough chips so that she can now win the next two bids, regardless of what Bob bids, and capture the king. Alice suppresses a smile as Bob realizes he has been defeated. Head in his hands, he mumbles, *"That was a total mindf**k."*

Acknowledgments. I am grateful to the organizers and audience at the thirteenth BAD Math Day, in Fall 2006 at MSRI, where discrete bidding games first met the general public,

for their patience and warm reception. I also thank Elwyn Berlekamp, David Eisenbud, and Ravi Vakil for their encouragement, which helped bring this project to completion. And finally, I throw down my glove at bidding game masters Andrew Ain, Allen Clement, and Ed Finn. Anytime. Anywhere. —*SP*

2 Preliminaries

2.1 Game model

Let G be a game played by two players, Alice and Bob, and modeled by a colored directed graph. The vertices of the graph represent possible positions in the game, and there is a distinguished vertex representing the starting position. The colored directed edges represent valid moves. Red and blue edges represent valid moves for Alice and for Bob, respectively, and two vertices may be connected by any combination of red and blue edges, in both directions. Each terminal vertex represents a possible ending position of the game, and is colored red or blue if it is a winning position for Alice or Bob, respectively, and is uncolored if it is a tie. For any possible position v in the game, we write G_v for the game played starting from v .

Recall that we write \overline{G} for the game that is exactly like G except that Alice and Bob exchange roles. So \overline{G} is modeled by the same graph as G , but with all colors and outcomes switched.

2.2 Bidding

Alice and Bob each start with a collection of bidding chips, and all bidding chips have equal value, for simplicity. When the game begins, the players write down nonnegative integer bids for the first move, not greater than the number of chips in their respective piles. The bids are revealed simultaneously, and the player making the higher bid gives that many chips to the other, and decides who makes the first move. The chosen player makes a move in the game, and then the process repeats, until the game reaches an end position or one player is unable to continue.

2.3 Tie-breaking

One player starts, by mutual agreement, with the *tie-breaking advantage*. If Alice has the tie-breaking advantage, and the bids are tied, then she can either declare Bob the winner of the bid and keep the tie-breaking advantage, or she can declare herself the winner of the bid and give the tie-breaking advantage to Bob. Similarly, if Bob has the tie-breaking advantage, then he can either declare Alice the winner of the bid, or he can declare himself the winner of the bid and give the tie-breaking advantage to Alice. In each case, the winner of the bid gives the amount of the bid to the other, and decides who makes the next move.

One virtue of this tie-breaking method is that it is never a disadvantage to have the tie-breaking advantage (see Lemma 3.1 below). Another virtue is that the tie-breaking advantage is worth less than an ordinary bidding chip (Lemma 3.2). Other reasonable tie-breaking methods are possible, and many of the results in this paper hold with other methods. We discuss some other tie-breaking methods in the Appendix.

We write $G(a^*, b)$ for the bidding game in which Alice starts with a bidding chips and the tie-breaking advantage, and Bob starts with b bidding chips. Similarly, $G(a, b^*)$ is the bidding game in which Alice starts with a bidding chips and Bob starts with b bidding chips and the tie-breaking advantage.

3 General theory

As we make the transition from recreational play to mathematical investigation, one of the most basic questions we can ask about a game G is for which values of a and b does Alice have a winning strategy for $G(a^*, b)$ or for $G(a, b^*)$. Often it is convenient to fix the total number of chips, and simply ask how many chips Alice needs to win. And when Alice does have a winning strategy, we ask how to find it. The general theory that we present here shows some of the structure that the answers to these questions must have. For instance, if Alice has a winning strategy for $G(a^*, b)$, then she also has a winning strategy for $G((a+1)^*, b)$, and if she has a winning strategy for $G(a^*, b+1)$, then she also has a winning strategy for $G(a+1, b^*)$. Similar results hold when Bob starts with the tie-breaking advantage. In each case, she can just play as if her extra chip wasn't there, or as if Bob's missing chip wasn't missing. Some of the structural results that we present here, such as the periodicity result in Section 3.6 are less obvious, and can be used to greatly simplify computations for specific games. We apply this approach to solve several games, including Tic-Tac-Toe, in Sections 4 and 6.

For simplicity, we always assume optimal play, and say that Alice wins if she has a winning strategy, and that she does not win if Bob has a strategy to prevent her from winning.

3.1 Value of the tie-breaking advantage

Roughly speaking, we show that the value of the tie-breaking advantage is strictly positive, but less than that of an ordinary bidding chip.

Lemma 3.1. *If Alice wins $G(a, b^*)$, then she also wins $G(a^*, b)$.*

Proof. Alice's winning strategy for $G(a^*, b)$ is as follows. She plays as if she did not have the tie-breaking chip until the first time the bids are tied. The first time the bids are tied, Alice declares herself the winner of the bid, and gives Bob the tie-breaking advantage. The resulting situation is the same as if Bob had started with the tie-breaking advantage and declared Alice the winner of the bid. Therefore, Alice has a winning strategy for the resulting situation, by assumption. \square

Lemma 3.2. *If Alice wins $G(a^*, b + 1)$, then she also wins $G(a + 1, b^*)$.*

Proof. Alice's winning strategy for $G(a + 1, b^*)$ is as follows. She begins by playing as if she started with a chips and the tie-breaking advantage except that whenever her strategy for $G(a^*, b + 1)$ called for bidding k and using the tie-breaking advantage, she bids $k + 1$ instead. She continues in this way until either she wins such a bid for $k + 1$ or Bob uses the tie-breaking advantage.

Suppose that Alice's strategy for $G(a^*, b + 1)$ called for bidding k for the first move and using the tie-breaking advantage in case of a tie. Then Alice bids $k + 1$ for the first move. If Alice wins the bid, then the resulting situation is the same as if Alice had won the first bid in $G(a^*, b + 1)$ using the tie-breaking advantage, so Alice has a winning strategy by hypothesis. Similarly, if Bob wins the bid using the tie-breaking advantage, then the resulting situation is the same as if Bob had won the first bid in $G(a^*, b + 1)$ by bidding $k + 1$, so Alice has a winning strategy. Finally, if Bob bids $k + 2$ or more chips to win the bid, then the resulting situation is a position that could have been reached following Alice's winning strategy for $G(a + 1, b^*)$, except that Alice has traded the tie-breaking advantage for two or more chips, and Alice can continue with her modified strategy outlined above.

The analysis of the case where Alice's strategy for $G(a^*, b + 1)$ did not call for using the tie-breaking advantage for the first move is similar. \square

Although Lemma 3.2 shows that trading the tie-breaking chip for a bidding chip is always advantageous, giving away the tie-breaking chip in exchange for an extra bidding chip from a third party is not necessarily a good idea; for any positive integer n , there is a game G such that Alice has a winning strategy for $G(a^*, b)$, but not for $G(a + n, b^*)$, as the following example demonstrates.

Example 3.3. Let G be the game where Bob wins if he gets any of the next n moves, and Alice wins otherwise. Then for $n \geq 1$, $(k^*, 0)$ is an Alice win if and only if $k \geq 2^{n-1} - 1$, while $(k, 0^*)$ is an Alice win if and only if $k \geq 2^n - 1$.

3.2 Using the tie-breaking advantage

In order for the tie-breaking advantage to have strictly positive value, as shown in Lemma 3.1, it is essential that the player who has it is not required to use it. However, the following proposition shows that it is always a good idea to use the tie-breaking advantage, unless you want to bid zero.

Proposition 3.4. *Both players have optimal strategies in which they use the tie-breaking advantage whenever the bids are nonzero and tied.*

Proof. Suppose that Alice has an optimal strategy which involves bidding k , but letting Bob win the bid if the bids are tied. If k is positive, then Alice can do at least as well by bidding $(k - 1)^*$ instead. If Bob bids k or more, the resulting situation is unchanged, while if Bob bids $k - 1$ or less, then Alice pays $(k - 1)^*$ instead of k to win the bid, which is at least as good by Lemma 3.2. \square

If the bids are tied at zero, it is not necessarily a good idea to use the tie-breaking advantage, as the following example shows.

Example 3.5. Consider the game where the player who makes the second move wins. Suppose Alice and Bob are playing this game, and they both start with the same number of chips. Then the player who starts with the tie-breaking advantage has a unique winning strategy—bid zero for the first move, decline to use the tie-breaking advantage, and bid everything to make the second move and win.

Proposition 3.4 shows that when looking for an optimal strategy, we can always assume that the player with the tie-breaking advantage either bids 0 or $0^*, 1^*, \dots$. Furthermore, if the player with the tie-breaking advantage bids 0, then the second player wins automatically and does best to bid 0 as well. Otherwise, if the player with the tie-breaking advantage bids k^* , we may assume that the second player either bids k and gains k^* chips while letting the first player move, or else bids $k + 1$ and wins the bid. These observations significantly reduce the number of bids one needs to consider when searching for a winning strategy.

3.3 Classical Richman calculus

For the reader's convenience, here we briefly recall the classical methods for determining the critical threshold $R(G)$ between zero and one such that Alice has a winning strategy if her proportion of the bidding resources is greater than $R(G)$ and does not have a winning strategy if her proportion of the bidding resources is less than $R(G)$. This *Richman calculus* also gives a method for finding the optimal moves and optimal bids for playing G as a bidding game with real-valued bidding. See the original papers [LLPU96, LLP⁺99] for further details. In Section 3.4, we present a similar method for determining the number of chips that Alice needs to win a discrete bidding game with a fixed total number of chips, as well as the optimal bids and moves for discrete bidding.

First, suppose G is bounded. We compute the critical thresholds $R(G_v)$ for all positions v in G by working backwards from the end positions. If v is an end position then

$$R(G_v) = \begin{cases} 0 & \text{if } v \text{ is a winning position for Alice.} \\ 1 & \text{otherwise.} \end{cases}$$

Suppose v is not an end position. If Alice makes the next move, then she will move to a position w such that $R(G_w)$ is minimal. Similarly, if Bob makes the next move, then he will move to a position w' such that $R(G_{w'})$ is maximal. We define

$$R_A(G_v) = \min_{A:v \rightarrow w} R(G_w) \quad \text{and} \quad R_B(G_v) = \max_{B:v \rightarrow w'} R(G_{w'}),$$

where the minimum and maximum are taken over Alice's legal moves from v and Bob's legal moves from v , respectively. The critical threshold $R(G_v)$ is then

$$R(G_v) = \frac{R_A(G_v) + R_B(G_v)}{2}.$$

The difference $R_B(G_v) - R_A(G_v)$ is a measure of how much both players want to move (or to prevent the other player from moving). If this difference is positive, then both players want to move, and if the difference is negative then the position is zugzwang and both players want to force the other to move. In either case, an optimal bid for both players is $\Delta_v = |R_B(G_v) - R_A(G_v)|/2$.

Next, suppose G is locally finite, but not necessarily bounded. Let $G[n]$ be the truncation of G after n moves. So $G[n]$ is just like G except that the game ends in a tie if there is no winner after n moves. In particular, Alice wins $G[n]$ if and only if she has a strategy to win G in at most n moves. We can compute the critical threshold $R(G)$ when G is bounded by using the bounded truncations $G[n]$, as follows. First, $R(G[n])$ can be computed by working backward from end positions, since $G[n]$ is bounded. Now $R(G[n])$ is a nonincreasing function of n that is bounded below by zero, so these critical thresholds approach a limit as n goes to infinity. Furthermore, since G is locally finite, Alice has a winning strategy for G if and only if she has a winning strategy that is guaranteed to succeed in some fixed finite number of moves. It follows that

$$R(G) = \lim_{n \rightarrow \infty} R(G[n]).$$

For games that are not locally finite, Alice may have a winning strategy, but no strategy that is guaranteed to win in a fixed finite number of moves. In this case, $R(G)$ is not necessarily the limit of the critical thresholds $R(G[n])$, as the following example shows.

Example 3.6. Let \mathbf{A}^m be the game that Alice wins after m moves, and let G be the game in which the first player to move can choose between the starting positions of \mathbf{A}^m for all positive integers m . Then Alice is guaranteed to win G , so $R(G) = 0$, but the critical threshold of each truncation is $R(G[n]) = 1/2$. Indeed, if Bob wins the first move of $G[n]$, then he can move to the starting position of \mathbf{A}^n , which Alice cannot win in the remaining $n - 1$ moves.

3.4 Discrete Richman calculus

Here we return to discrete bidding and compute the number of chips that Alice needs to win a locally finite game, assuming that the total number of chips is fixed. Since Alice may or may not have the tie-breaking advantage, the total number of chips that Alice has is an element of $\mathbb{N} \cup \mathbb{N}^*$, which is totally ordered by

$$0 < 0^* < 1 < 1^* < 2 < \dots$$

If we fix the game G and the total number of ordinary chips k , then it follows from Lemmas 3.1 and 3.2 that there is a critical threshold $f(G, k) \in \mathbb{N} \cup \mathbb{N}^*$ such that Alice wins if and only if she has at least $f(G, k)$ chips. Note that Alice can have at most k^* chips, so if G is a game in which Alice never wins, then $f(G, k) = k + 1$, by definition.

The critical thresholds $f(G, k)$ can be computed recursively from end positions for bounded games, and the critical thresholds of locally finite games can be computed from the critical thresholds of their truncations, just like the critical thresholds $R(G)$ for real-valued bidding. However, one must account for the effects of rounding, since the bidding chips are discrete, as well as the tie-breaking advantage.

First, suppose that v is an end position. Then

$$f(G_v, k) = \begin{cases} 0 & \text{if } v \text{ is a winning position for Alice.} \\ k + 1 & \text{otherwise.} \end{cases}$$

Next, suppose v is not an end position. If Alice makes the next move, then she will move to a position w such that $f(G_w, k)$ is minimal. Similarly, if Bob makes the next move, then he will move to a position w' such that $f(G_{w'}, k)$ is maximal. We define

$$f_A(G_v, k) = \min_{A:v \rightarrow w} f(G_w, k) \quad \text{and} \quad f_B(G_v, k) = \max_{B:v \rightarrow w'} f(G_{w'}, k),$$

where the minimum and maximum are taken over Alice's legal moves from v and Bob's legal moves from v , respectively.

For an element $x \in \mathbb{N} \cup \mathbb{N}^*$, we write $|x|$ for the underlying integer, so $|a|$ and $|a^*|$ are both equal to a , for nonnegative integers a . We also define $a + * = a^*$. For a real number x , we write $\lfloor x \rfloor$ for the greatest integer less than or equal to x .

Theorem 3.7. *For any position v , the critical threshold $f(G_v, k)$ is given by*

$$f(G_v, k) = \left\lfloor \frac{|f_A(G_v, k)| + |f_B(G_v, k)|}{2} \right\rfloor + \varepsilon,$$

where

$$\varepsilon = \begin{cases} 0 & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is even, and } f_A(G_v, k) \in \mathbb{N}. \\ 1 & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is odd, and } f_A(G_v, k) \in \mathbb{N}^*. \\ * & \text{otherwise.} \end{cases}$$

Proof. Since the critical threshold for any locally finite game can be computed from its bounded truncations, it is enough to prove the theorem in the case where G is bounded. If the game starts at an end position, then the theorem is vacuously true. We proceed by induction on the length of the bounded game.

Suppose $|f_A(G_v, k)| + |f_B(G_v, k)|$ is even and $f_A(G_v, k) \in \mathbb{N}$. If Alice has at least $(|f_A(G_v, k)| + |f_B(G_v, k)|)/2$ chips, then she can bid

$$\Delta = \left| |f_A(G_v, k)| - |f_B(G_v, k)| \right| / 2$$

and guarantee that she will end up in a position w with at least $f(G_w, k)$ chips. Then Alice has a winning strategy, by induction, since G_w is a bounded game of shorter length than G . Similarly, if Alice has fewer than $(|f_A(G_v, k)| + |f_B(G_v, k)|)/2$ chips, then Bob can bid Δ and guarantee that he will end up in a position w' where Alice will have fewer than $f(G_{w'}, k)$ chips. Then Bob can prevent Alice from winning, by induction.

Therefore, Alice wins G if and only if she has at least $(|f_A(G_v, k)| + |f_B(G_v, k)|)/2$ chips, as was to be shown. The proofs of the remaining cases, when $|f_A(G_v, k)| + |f_B(G_v, k)|$ is odd, and when $f_A(G_v, k)$ is in \mathbb{N}^* , are similar. If $|f_A(G_v, k)| + |f_B(G_v, k)|$ is odd and $f_A(G_v, k)$ is in \mathbb{N}^* , then the ideal bid for both players is the round down

$$\Delta = \left\lfloor \left| |f_A(G_v, k)| - |f_B(G_v, k)| \right| / 2 \right\rfloor.$$

If $|f_A(G_v, k)| + |f_B(G_v, k)|$ is odd but $f_A(G_v, k)$ is in \mathbb{N} , then both players should try to make the smallest possible bid that is strictly greater than $\left\lfloor \left| |f_A(G_v, k)| - |f_B(G_v, k)| \right| / 2 \right\rfloor$. And if $|f_A(G_v, k)| + |f_B(G_v, k)|$ is even and $f_A(G_v, k)$ is in \mathbb{N}^* then both players should try to make the smallest possible bid that is strictly greater than $\left\lfloor \left| |f_A(G_v, k)| - |f_B(G_v, k)| \right| / 2 \right\rfloor - 1$. \square

Theorem 3.7 makes it possible to find both the critical threshold and the optimal strategy for a given bounded game by working backward from end positions.

Example 3.8. Suppose \mathbf{A} is a game that Alice is guaranteed to win and \mathbf{B} is a game that Bob wins. Then

$$f(\mathbf{A}, k) = 0 \quad \text{and} \quad f(\mathbf{B}, k) = k + 1.$$

Let \mathbf{E} be the game in which the first player to move wins. Then

$$f(\mathbf{E}, k) = \begin{cases} (k + 1)/2 & \text{if } k \text{ is odd.} \\ \lfloor (k + 1)/2 \rfloor^* & \text{if } k \text{ is even.} \end{cases}$$

As games become more complicated, it is more convenient to encode the possibilities in a table. For instance, the critical thresholds for the game \mathbf{E} could be put in a table as follows.

$k = 2n +$	0	1
$f(\mathbf{E}, k) = n +$	0^*	1

Let A^2 be the game that Alice wins if she makes either of the first two moves and Bob wins otherwise. Similarly, let B^2 be the game that Bob wins if he makes either of the first two moves and Alice wins otherwise. Then the critical thresholds for A^2 and B^2 are given by

$k = 4n +$	0	1	2	3	and	$k = 4n +$	0	1	2	3
$f(A^2, k) = n +$	0	0^*	0^*	1		$f(B^2, k) = 3n +$	1	1^*	2^*	3

See Section 6 for detailed computations using such tables in a more interesting situation.

3.5 Discrete bidding with large numbers of chips

When G is played as a discrete bidding game, the optimal moves for Richman play are not necessarily still optimal. This may be seen as a consequence of the effects of rounding and tie-breaking in the discrete Richman calculus. However, one still expects that as the number of chips becomes large, discrete bidding games should become more and more like Richman games. Roughly speaking, the effects of rounding should only be significant enough to affect the outcome when the number of chips is small or Alice's proportion of the total number of chips is close to the critical threshold $R(G)$. We think of these situations as "unstable."

Definition 3.9. *A strategy for Alice is stable if, whenever Alice makes a move following this strategy, she moves to a position w such that $R(G_w)$ is as small as possible.*

Similarly, we say that a strategy for Bob is stable if, whenever he makes a move following this strategy, he moves to a position w' such that $R(G_{w'})$ is as large as possible. We say that a discrete bidding game is stable if both Alice and Bob have stable optimal strategies. Note that the proofs of Lemmas 3.1 and 3.2 go through essentially without change when "Alice wins" is replaced by "Alice has a stable winning strategy." For instance, if Alice has a stable winning strategy for $G(a, b^*)$, then she also has a stable winning strategy for $G(a^*, b)$.

Theorem 3.10. *For any locally finite game G , and for any positive ϵ , Alice has a stable winning strategy for $G(a, b^*)$ provided that $a/(a + b)$ is greater than $R(G) + \epsilon$ and a is sufficiently large.*

Proof. First, we claim that it suffices to prove the theorem when G is bounded. Recall that $G[n]$ denotes the truncation of G after n moves. So $G[n]$ is bounded and Alice wins $G[n]$ if and only if she wins G in at most n moves. By [LLPU96, Section 2], $R(G)$ is the limit as n goes to infinity of $R(G[n])$. Therefore, replacing G by $G[n]$ for n sufficiently large, we may assume that G is bounded.

Suppose G is bounded and guaranteed to end after n moves. If $n = 1$, then the theorem is clear. We proceed by induction on n . Assume the theorem holds for games guaranteed to terminate after $n - 1$ moves. Alice's strategy is as follows. For the first move, she bids x such that $x/(a + b)$ is between the optimal real-valued bid Δ and $\Delta + \epsilon/2$, which is possible since a is large. If Alice moves, then she moves to a position w such that $R(G_w)$ is minimal. Otherwise, Bob moves wherever he chooses. Either way, Alice ends up in a game G_v that is guaranteed to terminate after $n - 1$ moves, holding a' chips where $a'/(a + b)$ is greater than $R(G_v) + \epsilon/2$, and hence has a winning strategy for a sufficiently large, by induction. Since G is locally finite, there are only finitely many possibilities for v . Therefore, a can be chosen sufficiently large for all such possibilities, and the result follows. \square

The conclusion of Theorem 3.10 is false if G is not locally finite; we give an example illustrating this in Section 4.

Theorem 3.11. *For any finite game G , and for any positive ϵ , Bob has a stable strategy for preventing Alice from winning $G(a^*, b)$ provided that $a/(a + b)$ is less than $R(G) - \epsilon$ and b is sufficiently large.*

Proof. Recall that \overline{G} is the game that is identical to G except that Alice and Bob exchange roles. If G is finite, then $R(\overline{G}) = 1 - R(G)$. Therefore, the result follows from Theorem 3.10 applied to $\overline{G}(b, a^*)$. \square

Under the hypotheses of Theorem 3.11, Bob actually has a winning strategy. For locally finite games that are not finite, $R(\overline{G})$ may be strictly larger than $1 - R(G)$, so Bob should not be expected to have a winning strategy for $G(a^*, b)$. With real-valued bidding he may only have a strategy to prolong the game into an infinite draw.

Next, we show that finite games always become stable when the number of chips becomes sufficiently large.

Theorem 3.12. *Let G be a finite game. Then $G(a^*, b)$ and $G(a, b^*)$ are stable when $a + b$ is sufficiently large.*

Proof. Assume the total number of chips is large. We will show that either

1. Alice has a stable winning strategy, or
2. Any unstable strategy for Alice can be defeated by a stable strategy for Bob.

A similar argument shows that either Bob has a stable winning strategy, or any unstable strategy for Bob can be defeated by a stable winning strategy for Alice, and the theorem follows.

Suppose Alice follows an unstable strategy that calls for her to move to a position w such that $R(G_w)$ is not as small as possible. Let δ be the discrepancy

$$\delta = R(G_w) - R(G_{w_0}),$$

where w_0 is a position that Alice could have moved to such that $R(G_{w_0})$ is as small as possible. If Alice's proportion of the bidding chips is greater than $R(G) + \delta/4$, then she has a stable winning strategy, by Theorem 3.10. Therefore, we may assume that Alice's proportion of the bidding chips is at most $R(G) + \delta/4$.

Suppose that the position is not zugzwang, so Bob is bidding for the right to move. Since the total number of chips is large, Bob can make a bid that is strictly between $\Delta - \delta/2$ and $\Delta - \delta/4$, where $\Delta = R(G) - R(G_{w_0})$ is the optimal real-valued bid. Then, if Alice wins the bid and moves to w , she finds herself with a proportion of chips that is less than $R(G_w) - \delta/4$ and hence Bob has a stable winning strategy (since the number of chips is large). Otherwise, Bob wins and moves to a position w' such that $R(G_{w'})$ is as large as possible. Then his proportion of the chips is greater than $R(\overline{G}_{w'}) + \delta/4$, so again he has a stable winning strategy.

The proof when the position is zugzwang is similar except that Bob should bid between $\Delta + \delta/4$ and $\Delta + \delta/2$ to force Alice to move. \square

3.6 Periodicity

Here we prove a periodicity result for finite stable games that allows one to determine the outcome of G for all possible chip counts for Alice and Bob by checking only a finite number of cases. We use this result extensively in our analysis of specific bidding games in Sections 4 and 6.

Fix a finite game G . Choose a positive integer M such that $M \cdot R(G_v)$ and $M \cdot \Delta_v$ are integers for all positions v in G . For instance, one can take M to be the least common denominators of $R(G_v)$ and Δ_v for all v . Let

$$m = M \cdot R(G) \quad \text{and} \quad m_v = M \cdot R(G_v).$$

Similarly, let $\bar{m} = R(\bar{G})$ and $\bar{m}_v = R(\bar{G}_v)$.

Theorem 3.13. *If Alice has a stable winning strategy for $G(a^*, b)$ then she also has a stable winning strategy for $G(a + m^*, b + \bar{m})$. Similarly, if Alice has a stable winning strategy for $G(a, b^*)$ then she also has a stable winning strategy for $G(a + m, b + \bar{m}^*)$.*

Proof. Since G is finite, Alice's stable winning strategy for $G(a^*, b)$ is guaranteed to succeed in some fixed number of moves. If the game starts at a winning position for Alice, then the theorem is vacuously true, so we proceed by induction on the number of moves.

Suppose Alice has a stable winning strategy for $G(a^*, b)$ in which she bids k for the first move. Then Alice can win $G(a + m^*, b + \bar{m})$ by bidding $k + M\Delta$ for the first move, and moving according to her stable strategy for $G(a^*, b)$. Regardless of whether she wins the bid, Alice ends up in a position v where, compared to her stable strategy for $G(a^*, b)$, she has at least m_v additional chips and Bob has at most \bar{m}_v additional chips. By induction, Alice has a stable winning strategy starting from v that is guaranteed to win in a smaller number of moves, and the result follows.

The proof for the situation where Bob starts with tie-breaking advantage is similar. \square

Theorem 3.14. *If Bob has a stable strategy to prevent Alice from winning $G(a^*, b)$ then he also has a stable strategy to prevent Alice from winning $G(a + m^*, b + \bar{m})$. Similarly if Bob has a stable strategy to prevent Alice from winning $G(a, b^*)$, then he also has a stable strategy to prevent Alice from winning $G(a + m, b + \bar{m}^*)$.*

Proof. Similar to proof of Theorem 3.13. \square

Using these periodicity results, we can determine the exact set of chip counts for which Alice can win G by answering the following two questions for finitely many x and y in $\mathbb{N} \cup \mathbb{N}^*$.

- If Alice starts with x chips then how many chips does Bob need to prevent her from winning?
- If Bob starts with y chips then how many chips does Alice need to win?

There is a unique minimal answer in $\mathbb{N} \cup \mathbb{N}^*$ to each such question, and the answer is generally not difficult to find if G is relatively simple and x or y is small. Furthermore, by Theorem 3.12, there is an integer n such that $G(a^*, b)$ and $G(a, b^*)$ are stable provided that $a + b$ is at least n . Then, if we know the answers to the above questions for $x < m + n$ and $y < \bar{m} + n$, we can easily deduce whether Alice wins for any given chip counts using Theorems 3.13 and 3.14.

We conclude this section with some open problems that ask to what extent, if any, these results extend from finite games to locally finite games. For fixed $\epsilon > 0$ and a large number of chips, by Theorem 3.10 we know that Alice has a stable winning strategy if her chip count is at least $R(G) + \epsilon$ and Bob has a winning strategy if his proportion of the chips is at least $R(\bar{G}) + \epsilon$. However, if $R(G) + \epsilon$ is less than $1 - R(\bar{G}) - \epsilon$, then there is a gap where the outcome is unclear, even when the number of chips is large.

Problem 3.15. *Is there a locally finite game G and a positive number ϵ such that Alice has a winning strategy for infinitely many chip counts $G(a^*, b)$ such that $a/(a + b)$ is less than $R(G) - \epsilon$?*

Roughly speaking, Problem 3.15 asks whether strategies to force a draw in a locally finite game with real-valued bidding can always be approximated sufficiently well by discrete bidding with sufficiently many chips. However, it is not clear whether one should follow stable strategies in locally finite games with large numbers of chips.

Problem 3.16. *If G is locally finite, are $G(a^*, b)$ and $G(a, b^*)$ stable for a and b sufficiently large?*

If the answer to Problem 3.15 is negative, then the answer to Problem 3.16 is negative as well. To see this, suppose Alice wins a game G with a proportion of chips less than $R(G) - \epsilon$ and an arbitrarily large total number of chips. Let H be a stable game with Richman value between $R(G) - \epsilon$ and $R(G)$ (which is not difficult to construct), and let $G \wedge H$ be the game in which the first player to move gets to choose between the starting position of G and the starting position of H . In $G \wedge H$, Alice's optimal first move for large chip counts would be to move to the starting position for G , while her stable strategy (i.e. her optimal strategy for real-valued bidding) would be to move to the starting position for H . In the Appendix, we show that the answer to Problem 3.16 is negative for different tie-breaking methods.

4 Examples

In this section, we analyze discrete bidding play for two simple combinatorial games, Tug o' War and what we call *Ultimatum*. We use these examples to construct games with strange behavior under discrete bidding play.

4.1 Tug o' War

The Tug o' War game of length n , which we denote by Tug^n , is played on a path of length $2n$, with vertices labeled $-n, \dots, -1, 0, 1, \dots, n$, from left to right. The game starts at the center vertex, which is marked 0. Adjacent vertices are connected by edges of both colors in both directions. Alice's winning position is the right-most vertex of the path, and Bob's winning position is the left-most vertex of the path. So Alice tries to move to the right, Bob tries to move the left, and the winner is the first player to reach the end. In particular, Tug o' War is stable, since the optimal moves do not depend on whether bidding is discrete or real-valued, and since it is also symmetric the critical threshold is $R(\text{Tug}^n) = 1/2$.

Proposition 4.1. *Suppose Bob's total number of chips is less than n . Then Alice wins Tug^n if and only if her total number of chips is at least $(n - 1)^*$.*

Proof. We define the weight of a position in the bidding game to be the number of the current vertex plus the number of chips that Alice has, including the tie-breaking chip. Note that, in order to reach a winning position, Alice must first reach a position of weight at least n .

Suppose Alice's chip total is at most $n - 1$. Then Bob can force a draw by bidding zero every time, and using the tie-breaking advantage to win whenever possible. Indeed, if Bob does this, then the weight of the position in the game never exceeds its starting value, so Alice cannot win.

Suppose Alice's chip total is at least $(n - 1)^*$. Then Alice can win with the following strategy. Whenever she has the tie-breaking advantage, she bids zero and uses the advantage if possible. Whenever she does not have the tie-breaking advantage, she bids one. With this strategy, the weight of the position never decreases, and it follows that Bob cannot win. Bob may win the first several moves, but eventually he will run out of chips, and Alice will win a move for zero, using the tie-breaking advantage. Then Alice may win a certain number of moves for one chip each. Since the weight of the position is at least n , if Bob lets her win moves for one chip indefinitely, then Alice will win. So eventually Bob must bid either two chips or one plus the tie-breaking advantage, and the weight of Alice's position increases by one. It follows that Alice can raise the weight of her position indefinitely, until eventually she must win. \square

The Richman game version of Tug o' War was studied in [LLP⁺99, p. 252], where the critical threshold of the vertex labeled k was determined to be $(k + n)/2n$. Therefore, the periodicity results of Section 3.6 hold with $M = 2n$ and with m and \bar{m} both equal to n . The cases covered by Proposition 4.1 then completely determine the outcome of Tug o' War for all possible chip counts.

Corollary 4.2. *Let a, b, k , and k' be nonnegative integers, with a and b less than n . Then Alice wins $\text{Tug}^n(kn + a, (k'n + b)^*)$ if and only if k is greater than k' . Furthermore, Alice wins $\text{Tug}^n(kn + a^*, k'n + b)$ if and only if either k is greater than k' or k is equal to k' and a is equal to $n - 1$.*

Tug o' War game is perhaps the simplest game that is not bounded, and yet we can use it to construct some interesting examples of bidding game phenomena for games that are not locally finite.

Proposition 4.3. *Let G be the game in which the first player to move can go to the starting position of Tug^n for any n . Then $R(G) = 1/2$, but $G(3a, a^*)$ is a draw for all values of a greater than one.*

Proof. Since all of the possible first moves lead to positions v with $R(G_v) = 1/2$, the critical threshold is $R(G) = 1/2$. Consider $G(3a, a^*)$ for any a . Bob's strategy to force a draw is as follows. Bob bets all of his chips for the first move. Then Alice can either bet $a + 1$ and win the bet, getting to play any $\text{Tug}^n(2a - 1, 2a + 1^*)$ which is at best a draw for Alice. Otherwise, Bob wins the bet, and chooses to play Tug^n for some n greater than $4a + 1$, which leads to a draw by Proposition 4.1. So Bob can force a draw. \square

This type of behavior, where Bob can force a draw even though Alice's proportion of the chips is much greater than $R(G)$ and the total number of chips is large, is impossible for locally finite games by Theorem 3.10.

The following example shows that non locally finite games may also be unstable even for large numbers of chips.

Proposition 4.4. *Let G be the game in which the first player to move can either play Tug^1 , or Tug^n starting at the node labeled -1 for any $n \geq 10$. Then $G(3a^*, 2a)$ fails to be stable for all a .*

Proof. We claim that Bob has no optimal stable strategy. For the first move, Bob's only stable strategy is to move to Tug^1 , since its Richman value is $1/2$ and the Richman values of all Tug^n starting at -1 is less than $1/2$. However, if Bob wins the first bid and moves to Tug^1 then he will lose. Nevertheless, we claim that Bob has a strategy to force a draw. This strategy is as follows.

Bob bets all of his chips on the first turn. If Alice lets him win the bid, he can move to Tug^n for n large, which leads to a draw. Otherwise, if Alice bets $2a^*$, she may choose to play either $\text{Tug}^1(a, 4a^*)$ or $\text{Tug}^n(a, 4a^*)$ starting from the node labeled -1 . If Alice chooses Tug^1 then Bob wins. Otherwise, Bob can win the next move for a^* , leading to $\text{Tug}^n(2a^*, 3a)$, which is at worst a draw for Bob. Therefore Bob has a nonstable strategy that is better than any stable strategy. \square

4.2 Ultimatum

We now describe the Ultimatum game of degree n , which we denote by Ult^n . It is played on a directed graph with vertices labeled $B, -n, \dots, -1, 0, 1, \dots, n - 1, n, A$. There are red edges from 0 to n , from k to A for $k > 0$, and from k to $k + 1$ for $k < 0$. Similarly, there are blue edges from 0 to $-n$, from k to B for $k < 0$, and from k to $k - 1$ for $k > 0$. The game starts at 0 . In other words, when the game starts, the first player to move gives the other an ultimatum—the other player must make each of the next n moves (in which case the game reverts to the beginning state), or else lose the game.

Since Ult^n is finite and symmetric, the critical threshold $R(U^n)$ is $1/2$.

Proposition 4.5. *Suppose b is less than 2^n . Then Alice has a winning strategy for $\text{Ult}^n(a^*, b)$ if and only if a is greater than b , or a is equal to b and $b \neq 2^{n-1} - 1$.*

Proof. Suppose a is greater than b . We claim that Alice can win by bidding b chips on the first move and using the tie-breaking advantage. Then Alice still has at least one chip left, so Bob must bid at least one chip plus the tie-breaking advantage for the second move, three chips for the third move, and $3 \cdot 2^k$ for move number $k + 3$. It follows that if Bob is able to make n moves in a row, then Alice receives at least $3 \cdot (2^{n-1} - 1)$ chips from Bob before returning to 0. In particular, Alice returns to the starting position with strictly more chips than she started with, and hence must eventually win.

Suppose a is equal to b and less than $2^{n-1} - 1$. Then Alice can win by bidding all of her chips on the first move and using the tie-breaking advantage. Bob must give her the tie-breaking advantage to take the second move, and one chip to take the third move, and 2^k chips to take move number $k + 3$. It follows that if Bob is able to make n moves in a row, then Alice will have collected $2^{n-1} - 1$ chips from Bob by the time they return to the starting position. Now Alice has more chips than Bob, plus the tie-breaking advantage, so she has a winning strategy.

Suppose a is equal to b and greater than $2^{n-1} - 1$. Then Alice bids $a - 1$; if Bob bids all of his chips to win, then Alice can pay him $2^{n-1} - 1$ chips, plus the tie-breaking advantage, to return to the starting position. Now Alice has at least two more chips than Bob, so she can bid Bob's number of chips plus one to move to vertex n . Then, if Bob has enough chips to make the next n moves, Alice will return to the starting vertex with more chips than Bob, plus the tie-breaking advantage, and will therefore win.

Finally, suppose a and b are both equal to $2^{n-1} - 1$. Then Bob can prevent Alice from winning by bidding all of his chips for the first move. If Alice bids all of her chips plus the tie-breaking advantage to take the first move, then Bob has exactly enough chips to return to the starting position with $2^{n-1} - 1$ chips left. Otherwise, Bob makes the first move, and Alice has just enough chips to return to the starting position with $2^{n-1} - 1$ chips left. Since the tie-breaking advantage is always an advantage, by Lemma 3.1, Bob's position is no worse than when the game began, so Alice cannot win. \square

Proposition 4.6. *Suppose $b < 2^n$. Then Alice has a winning strategy for $\text{Ult}^n(a, b^*)$ if and only if a is greater than $b + 1$, or a is equal to $b + 1$ and $b \neq 2^{n-1} - 1$.*

Proof. The proof is essentially identical to the previous proposition's. If Alice has more than $b + 1$ chips, she just bets $b + 1$ and wins. If she has $b + 1$ chips, then unless the congruence condition holds, she can win by betting $b + 1$ chips if b is less than $2^{n-1} - 1$ and b chips if b is greater than $2^{n-1} - 1$. If b is equal to $2^{n-1} - 1$ then Bob can bet all his chips to return with either b^* chips or $b + 1$ chips, thus forcing a draw. \square

Now Ult^n is clearly stable, since there is only one move available to each player from each position, so the above cases can be used to determine the outcome of Ult^n for all possible chip counts using the periodicity results of Section 3.6.

Corollary 4.7. *Alice wins $\text{Ult}^n(a^*, b)$ if and only if either a is greater than b or a and b are equal but not congruent to $2^{n-1} - 1 \pmod{2^n}$. Similarly, Alice wins $\text{Ult}^n(a, b^*)$ if and only if either a is greater than $b + 1$, or a is equal to $b + 1$ and not congruent to $2^{n-1} \pmod{2^n}$.*

Using these computations for the Ultimatum games, we can construct another interesting example: a discrete bidding game G whose critical threshold $R(G)$ is rational, but when Alice's proportion of the total number of chips is exactly equal to $R(G)$ the outcome is not periodic in the total number of chips. For finite games, such behavior is impossible by the periodicity results of Section 3.6.

Proposition 4.8. *Let G be the game where on the first move, either player can choose one of $\text{Ult}^1, \text{Ult}^3, \text{Ult}^5, \dots$, and then the players play the chosen game with the current chip counts. Then $R(G) = 1/2$, but the sequence of results of $G(k^*, k)$ is aperiodic.*

Proof. Since the game is symmetric, it cannot be a win for Bob. If Alice either bets 1 or uses the tiebreaking advantage in a 0-0 tie, then she will be playing an ultimatum game with fewer chips and cannot win. So her only chance at forcing a win is to let Bob win for 0 chips. So Bob picks any ultimatum game. If he can pick a game for which $\text{Ult}^n(k^*, k)$ is a draw, $G(k^*, k)$ is a draw; otherwise, $G(k^*, k)$ is an Alice win.

However, the values of k for which $\text{Ult}^n(k^*, k)$ is a draw are those for which $k \equiv 2^{n-1} - 1 \pmod{2^n}$. Only games Ult^n for odd n are in play, so Bob can draw whenever k is congruent to 0 mod 2, or 3 mod 8, or 15 mod 32, and so on. It is easily seen that the set of k which satisfy these congruence conditions is not a finite union of arithmetic progressions and hence the outcome is not periodic in k . \square

5 A partial order on games

If we know the optimal moves for Alice and Bob, or if both the game G and the total number of chips k are fixed, then the critical threshold $f(G, k)$ is easy to compute using the formulas in Section 3.4. On the other hand, for fixed G , the tree of optimal moves may vary as k varies. In this section, we discuss situations in which one can find moves that are optimal for all k .

We define a partial order on games by setting $G \leq G'$ if $f(G, k) \leq f(G', k)$ for all k . In other words, $G \leq G'$ means that, regardless of the total number of chips, if Alice has enough chips to win G' then she also has enough chips to win G . Then G and G' are equivalent in this partial order if $f(G, k) = f(G', k)$ for all k . The point of this partial order is that if $G \leq G'$, then Alice will always prefer to move to a position v such that G_v is equivalent to G rather than a position v' such that $G_{v'}$ is equivalent to G' , regardless of the total number of chips.

Example 5.1. Let \mathbf{A} be the equivalence class of games in which Alice always wins, and let \mathbf{B} be the equivalence class of games in which Bob always wins. Then

$$\mathbf{A} \leq G \leq \mathbf{B}$$

for every game G .

If G and G' are games with many legal moves available from each position, then it may be very difficult to tell whether G and G' are comparable in this partial order. In practice, we can sometimes compare relatively complicated games G and G' by relating them to games with fewer legal moves.

Example 5.2. For a positive integer n , let A^n be the game in which Alice wins if she makes any of the first n moves, and Bob wins otherwise. Roughly speaking, we think of A^n as a sudden-death game in which Alice has an advantage of order n . Similarly, let B^n be the game in which Bob wins if he makes any of the first n moves and Alice wins otherwise. We write \mathbf{E} for the equivalence class of games in which the player with more chips always wins. One game in the class of \mathbf{E} is the game in which the first player to move wins. Then A^1 and B^1 are both equivalent to \mathbf{E} , and

$$\mathbf{A} < \dots < A^3 < A^2 < \mathbf{E} < B^2 < B^3 < \dots < \mathbf{B}.$$

The games A^n and B^n are useful for comparing positions, since $A^n \leq G$ for any game G in which Bob wins if he makes all of the first n moves, and $G' \leq A^n$ for any game G' in which Alice wins if she makes any of the first n moves, and similarly for B^n .

Example 5.3. Games that are equivalent to \mathbf{E} in this partial order are not necessarily obviously symmetric. For instance, Tic-Tac-Toe played starting from any of the following six positions is equivalent to \mathbf{E} .



From each of these positions, Alice wins if she makes two moves before Bob does, and Bob can prevent Alice from winning if he makes two moves before Alice does. It follows that each of these positions is equivalent to $A^2 \wedge B^2$, which is equivalent to \mathbf{E} .

Recall that the wedge sum $G \wedge H$ of two games G and H is the game where the first player to move can choose between the starting position for G and the starting position for H . Although $G \wedge H$ is the same as $H \wedge G$, when one game is clearly better for each player we write the game that Alice prefers on the left. For example, $A^2 \wedge B^3$ is a game in which Alice can move to a position equivalent to A^2 and Bob can move to a position equivalent to B^3 . This notation is convenient for working with equivalence classes of games and our partial order, since $G \wedge H$ is equivalent to $G' \wedge H$ if G is equivalent to G' , and $G \wedge H \leq G' \wedge H'$ if $G \leq G'$ and $H \leq H'$. Simple relations between games are also easily expressible in this notation. For instance, we have equivalences of games

$$A^n \equiv \mathbf{A} \wedge A^{n-1} \quad \text{and} \quad \mathbf{E} \equiv G \wedge \overline{G},$$

for any integer $n \geq 2$ and any game G .

In Section 6, we completely analyze bidding Tic-Tac-Toe by comparing positions in Tic-Tac-Toe with iterated wedge sums of games \mathbf{A} , \mathbf{B} , A^m and B^n . For instance, we show that Tic-Tac-Toe played from the position

A	B
A	B

is equivalent to $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$.

6 Bidding Tic-Tac-Toe

In this section, we use comparisons with bivalent games and the partial orders defined in Section 5 to completely analyze Tic-Tac-Toe for both real-valued bidding and discrete bidding with arbitrary chip counts.

Since Tic-Tac-Toe is symmetric, to determine whether or not Alice wins Tic-Tac-Toe(a, b) for all possible chip counts

$$(a, b) \in (\mathbb{N} \times \mathbb{N}^*) \cup (\mathbb{N}^* \times \mathbb{N}).$$

Indeed, Bob wins Tic-Tac-Toe(a, b) if and only if Alice wins Tic-Tac-Toe(b, a), and the outcome is a tie if and only if neither Alice nor Bob wins. Therefore, we will focus our analysis on the game TTT that is played just like Tic-Tac-Toe except that Bob is declared the winner of any game that would normally be a tie. In particular, Alice wins Tic-Tac-Toe if and only if she wins TTT , so these games have the same critical thresholds.

6.1 Optimal moves

We begin by determining the optimal moves in Tic-Tac-Toe, as much as possible, using the partial order defined in Section 5. This simple approach determines all of the optimal moves except Alice's optimal move from an empty board, which we reduce to two possibilities. In Section 6.2 we compute the number of chips that Alice needs to win for every possible chip total for each of these two possibilities, using the recursion from Theorem 3.7 for small chip totals and the periodicity results in Theorems 3.13 and 3.14 for higher chip totals.

Theorem 6.1. *The tree of moves for TTT shown in Figure 1 is optimal for real-valued bidding. These moves are also optimal for discrete bidding if the total number of chips is not equal to five. Furthermore, the critical thresholds of each position reached through these optimal moves is as indicated in the figure. In particular,*

$$R(TTT) = 133/256.$$

For discrete bidding with five chips, a tree of optimal moves is shown in Figure 2, below.

The critical threshold $R(TTT) = 133/256$ was computed independently by Theodore Hwa [Hwa06].

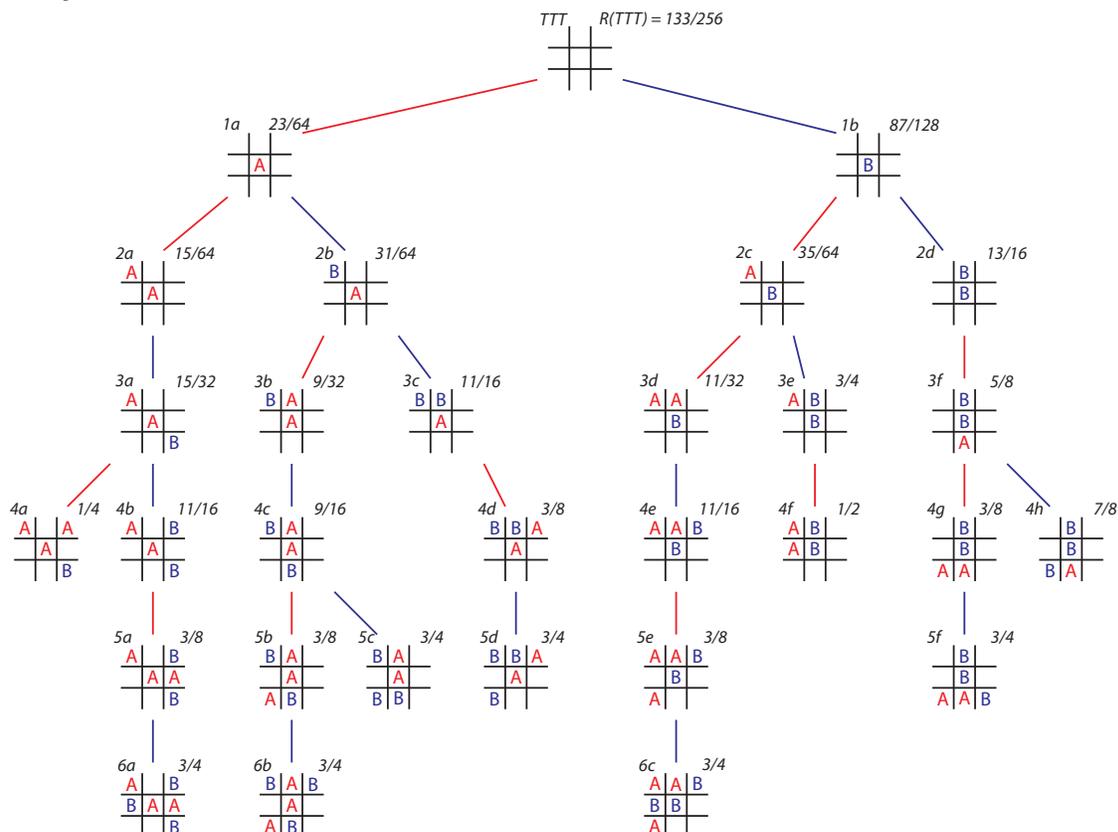


Figure 1: Optimal moves and critical thresholds for TTT with real-valued bidding

Our general approach to studying TTT is to work our way backward from end positions. Knowing which sequence of moves is optimal late in the game greatly simplifies the analysis of earlier positions, by reducing to comparisons with wedge sums of games equivalent to \mathbf{A} , \mathbf{B} , A^n or B^n .

For discussing TTT and other games, it is useful to have some basic language for describing positions, although we attempt to keep jargon to a minimum. We say that Alice has a *threat* if she can win on the next move, and Bob has a *threat* if he can prevent Alice from winning with his next move. If Alice has a threat from a position v then, in the partial order on games discussed in Section 5,

$$G_v \leq \mathbf{E},$$

and if Bob has a threat from v' then $\mathbf{E} \leq G_{v'}$, where \mathbf{E} is the game in which the first player to move wins. We say that Alice has a *double threat* if she can win if she gets either of the next two moves, and Bob has a *double threat* or *triple threat* if he can prevent Alice from winning if he gets any of the next two moves or three moves, respectively.

Example 6.2. Bob has a triple threat from the position

	B	
	B	
B	A	

since he can prevent Alice from winning if he moves anywhere in the right column. This position is equivalent to B^3 , since Alice wins if she gets three moves in a row. On the other hand, Alice has a double threat from the position

A	B	A
	A	

but this position is strictly better for Alice than A^2 , since Bob cannot prevent Alice from winning if he gets two moves in a row.

We say that a move is a *block* if it goes from a position in which the opponent has a threat to a position in which the opponent does not have a threat. A move is a *counter-attack* if it is a block that also creates a threat.

Lemma 6.3. *If a counterattack is available, then the optimal move must be a block.*

Proof. A counterattack moves to a position that is at least as good as \mathbf{E} , and any move that is not a block moves to a position that is no better than \mathbf{E} . \square

In Tic-Tac-Toe, if a player has a block available then it is always unique. Therefore, counterattacks are always optimal in Tic-Tac-Toe. Note however, that blocks are not always optimal in Tic-Tac-Toe. See the analysis of Alice's move from position 3e, below, for an example where blocking is not optimal.

Proposition 6.4. *Every move in Figure 1 other than Alice's first move is optimal for all chip counts.*

Proof. There are twenty-nine moves in the stable game tree above, so we must analyze each of these twenty-nine moves. The moves from positions 5a, 5b, 5e, 4a, 4d, 4e, 4f, 3d, and 3e are counterattacks, and by Lemma 6.3, counterattacks are always optimal. We now analyze each of the remaining moves other than Alice's first move, moving from bottom to top and from left to right in Figure 1.

Alice's move from 4c: Wherever Alice moves, Bob can create a double threat by moving top right or bottom left. Therefore, the best game Alice can hope to reach is $\mathbf{A} \wedge B^2$, which she reaches by moving bottom left.

Bob's move from 4c: Wherever Bob moves, Alice can win if she gets two moves in a row. Therefore, the best game Bob can hope to move to is B^2 , which he reaches by moving bottom left.

Alice's move from 3a: Wherever Alice moves, Bob can win if he gets two moves in a row by filling either the right column or the bottom row. Therefore, the best game Alice can hope to move to is A^2 , which she reaches by moving top right.

Bob's move from 3a: Bob's move in either corner creates $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$, which is better for him than $A^2 \wedge \mathbf{B}$ and hence better than any position from which Alice can produce a double threat. If Bob moves on a side, then Alice can create a double threat by moving in a corner. Therefore, Bob's corner move is optimal.

Bob's move from 3b: If Bob does not block, then the best game he can reach is \mathbf{E} , which is equivalent to $A^2 \wedge B^2$. Blocking gives $(\mathbf{A} \wedge B^2) \wedge B^2$. Since $\mathbf{A} \wedge B^2$ is strictly better than A^2 , it follows that blocking is optimal for Bob.

Alice's move from 3e: This is an interesting move, since Alice's optimal strategy is not to block, but instead to create a threat of her own to reach a game equivalent to \mathbf{E} . If Alice blocks, then the resulting position is $A^2 \wedge B^3$, since Bob can move lower left to create a triple threat, and this is strictly worse for Alice than \mathbf{E} . Of course, if Alice moves in the right column, then Bob has a threat but Alice does not, which is again worse than \mathbf{E} . Therefore, moving in the left column is optimal for Alice.

Alice's move from 3f: Wherever Alice moves, Bob can produce a double threat by moving in a bottom corner. Therefore Alice's bottom corner move, which gives $\mathbf{A} \wedge B^2$, is optimal.

Bob's move from 3f: Wherever Bob moves, Alice can win if she gets the next three moves in a row. Therefore, Bob's move in a bottom corner, which produces a triple threat, is optimal.

Bob's move from 2a: If Bob does not block, then he must get the next move as well. And if blocking is not optimal, then his optimal next move must also be nonblocking (since otherwise he may as well have blocked the first time). If Bob makes two moves without blocking, then the best he can do is create a symmetric threat, reaching \mathbf{E} . However, if he does block, then in two moves he can reach a position in which he has a threat and Alice does not, which is strictly better than \mathbf{E} . Therefore, blocking is optimal.

Alice's move from 2b: Suppose Alice moves upper right or lower left. Then Bob's next move must be a counterattack, and Alice's move after that must be a counterattack as well, and the resulting position after five moves is the one labeled 5a in the game tree. Therefore either of these moves leads to $\mathbf{A} \wedge (TTT_{5a} \wedge \mathbf{B})$.

Suppose Alice moves lower right. Then Bob can move upper right forcing Alice to counterattack in the top center, leading to position 5a. Since this move does not create a threat, it is strictly worse than $\mathbf{A} \wedge (TTT_{5a} \wedge \mathbf{B})$.

Finally, suppose Alice moves top or left. Then we have seen in the analysis of position 4c that Bob's optimal response is to block. It follows that moving top or left leads to $\mathbf{A} \wedge (TTT_{5b} \wedge B^2)$. Since TTT_{5a} and TTT_{5b} are both equivalent to $\mathbf{A} \wedge B^2$, it follows that any move other than lower right is optimal. In particular, moving top, as shown in the

game tree, is optimal.

Bob's move from 2b: If Bob moves top or left, then the resulting position is equivalent to $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$.

Suppose Bob moves upper right or lower left. Then Alice must counterattack. No matter what Bob does next, Alice will still be able to win if she gets the next two moves. Therefore, this is no better for Bob than $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$.

Finally, suppose Bob moves lower right. Then Alice can move in the top, forcing Bob to counterattack, and again Alice can win if she gets the next two moves. So this is still no better than $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$. Therefore, moving top or left is optimal.

Alice's move from 2c: If Alice moves top or left, a series of counterattacks ensues, as shown in the stable game tree. This position gives $\mathbf{A} \wedge ((\mathbf{A} \wedge B^2) \wedge \mathbf{B})$.

If Alice moves upper right or lower left, then Bob counterattacks and then best Alice can do in response is to reach a symmetric position. Therefore, these moves give $\mathbf{A} \wedge (\mathbf{E} \wedge \mathbf{B})$ which is strictly worse for Alice than moving top or left.

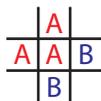
Finally, suppose Alice moves lower right. Then Bob can move on a side and force a series of counterattacks. After such a move from Bob, this position is then $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$, which is the same as if Alice moves top or left and then Bob counters. However, since lower right does not create a threat for Alice, this is strictly worse than moving top or left, and it follows that top and left are Alice's only optimal moves.

Bob's move from 2c: Wherever Bob moves, Alice can create a threat on the next move. Therefore, the best game Bob can hope to reach is $\mathbf{E} \wedge \mathbf{B}$, which he gets to by moving top or left.

Alice's move from 2d: If Alice does not block, then she must get two moves in a row, and if she gets two moves in a row then the best game she can reach without blocking is \mathbf{E} . Since she can reach position 4g in two moves, which is strictly better than \mathbf{E} , it follows that blocking is her only optimal move.

Alice's move from 1a: If Alice moves in a corner, then the resulting game is $\mathbf{A} \wedge (A^2 \wedge TTT_{3a})$. We must show that Alice cannot do better by moving on the top.

Suppose Alice moves on the top. Then Bob can block on the bottom. If Bob gets the next move, then he can force a series of counterattacks, guaranteeing that this is no worse for him than 3a. Therefore, it will suffice to show that if Alice gets the next move after Bob blocks, then the best she can reach is a game equivalent to A^2 . If Alice moves in the top row, then she has a double threat and Bob can win if he gets two moves in a row, so this is exactly A^2 . If Alice moves in the middle row, then Bob can block, reaching a position



that is symmetric by Example 5.3. And if Alice moves in the bottom row, then Bob can block, reaching a position

	A	B
	A	
A	B	

that is symmetric by Example 5.3. It follows that Alice cannot do better than moving in a corner.

Bob's move from 1a: If Bob moves in a corner, then the resulting game is $(\mathbf{A} \wedge TTT_{4c}) \wedge (TTT_{4d} \wedge \mathbf{B})$. We must show that Bob cannot do better by moving on top.

Suppose Bob moves on top and Alice wins the next move. Then Alice can move on the left side, and we claim that the resulting position is equivalent to $\mathbf{A} \wedge \mathbf{E}$. If Bob does not block, then the best he can reach is a symmetric position. If Bob does block, then the resulting position is still symmetric, since Alice's best move is bottom left, which gives A^2 , and Bob's best move is top right, which gives B^2 . Now $\mathbf{A} \wedge \mathbf{E}$ is strictly worse for Bob than $\mathbf{A} \wedge TTT_{4c}$, which results if Bob moves in a corner and Alice wins the next move.

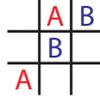
Suppose Bob moves on top and gets the next move as well. The only way he can do better than starting in a corner is by moving on a side again. Suppose he moves left on his second move. Then Alice can move in the upper left, and Bob's counterattack leads to a position B^2 . If Bob moves on the bottom on his second move, then Alice again moves in the upper left, and Bob's counterattack is worse for him than B^2 . It follows that after making two consecutive moves on sides, Bob's position is strictly worse than the position $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$ that he would reach by moving first in a corner and then on an adjacent side. Therefore, the corner move from position 1a is optimal for Bob.

Alice's move from 1b: Suppose Alice moves on the top instead. If Alice gets the following move as well, she could only do better by moving on another side, either left or bottom. If she plays on the left, then Bob can move upper left, and Alice will need to block in order to win. Therefore, Alice blocks and Bob's next move creates B^2 . In other words, after Alice moves top and left, the position is strictly worse for Alice than $\mathbf{A} \wedge ((\mathbf{A} \wedge B^2) \wedge \mathbf{B})$. On the other hand, if Alice moves top and then bottom, then Bob can move left and the best Alice can do is to reach the position

	A	
B	B	A
	A	

that is symmetric by Example 5.3. Therefore, if Alice moves top and then bottom, the resulting position is strictly worse for Alice than $\mathbf{A} \wedge (\mathbf{E} \wedge \mathbf{B})$. Both of these situations are worse than the position 3d, which is equivalent to $\mathbf{A} \wedge ((\mathbf{A} \wedge B^2) \wedge \mathbf{B})$, which Alice gets to by playing top left and then top.

Next, suppose Alice moves on top and Bob makes the next move. Bob can play top right, and then Alice needs to take the lower left corner to have any chance of winning. So Alice blocks, and reaches a position



that is symmetric, by Example 5.3. In particular, if Alice moves top and then Bob makes the next move, the result is no better for Alice than $\mathbf{E} \wedge \mathbf{B}$. It follows that Alice cannot do better than reaching the position $TTT_{4d} \wedge (\mathbf{E} \wedge \mathbf{B})$ that she gets by playing in the corner.

Bob's move from 1b: Suppose Bob moves in the upper left corner instead. Then Alice can block in the lower right. If Bob gets the next move, he can do no better than to reach B^2 , while Alice can move on the bottom to reach a position $\mathbf{A} \wedge B^2$. Therefore, if Bob moves in a corner, his position is no better than $((\mathbf{A} \wedge B^2) \wedge B^2) \wedge \mathbf{B}$, which is worse for him than TTT_{2d} , which is equivalent to $((\mathbf{A} \wedge B^2) \wedge B^3) \wedge \mathbf{B}$.

Bob's first move: Suppose Bob makes the first move, but not in the center. If Alice makes the second move, she can go in the center. By our analysis of Bob's move from 1a,



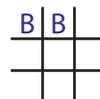
and by inspection of the tree of optimal moves from positions 2b and 2c in Figure 1, it is straightforward to check that $TTT_{2c} \geq TTT_{2b}$. Therefore, if Alice makes the second move, Bob would have done best to move in the center.

Suppose Bob makes the second move as well. If Bob's optimal first move is not in the center, then his optimal second move must be not in the center as well. Suppose Bob starts with any of the following two moves:



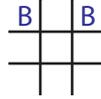
Then Alice can respond by moving in the center. After this, Bob can do no better than to create a double threat on his next move. Therefore each of the four positions above is no better than $((\mathbf{A} \wedge B^2) \wedge B^2) \wedge \mathbf{B}$.

Similarly, if Bob starts with



and then Alice blocks, the resulting position is still no better for Bob than $(\mathbf{A} \wedge B^2) \wedge B^2$.

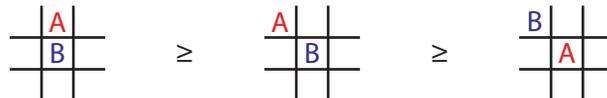
Finally, if Bob starts with



then Alice can block. On the next move, Bob can do no better than to move in the center and create a B^3 , while Alice can move in the center to create a position that is worse for Bob than $\mathbf{A} \wedge B^2$, since Alice can win on the next move, but Bob cannot create a double threat. In particular, all of these scenarios are worse for Bob than $(\mathbf{A} \wedge B^2) \wedge B^3) \wedge \mathbf{B}$, which is equivalent to TTT_{2d} . It follows that Bob's only optimal first move is in the center. \square

Lemma 6.5. *If center is not an optimal first move for Alice then Alice's optimal first two moves are in adjacent corners.*

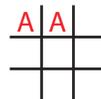
Proof. Suppose Alice makes an optimal first move, but not in the center. If Bob moves second then he can move in the center. However, by our analysis of Alice's move from position 1b, and by inspection of the trees of optimal moves from positions 2b and 2c, we have



Therefore, Alice must reach a position better than 2a by making the first two moves, neither in the center. We consider all possible positions that Alice can reach after making the first two moves. If Alice's first two moves are any of the following



then Bob moves in the center. Then Alice cannot reach a position better than A^2 on her next move, while Bob can move in either the top left corner or the top side to reach a position that is worse for Alice than $4b \equiv (\mathbf{A} \wedge B^2) \wedge \mathbf{B}$. Similarly, if Alice makes her first two moves



then Bob can block. After Bob blocks, if Alice moves next she can get no better than A^2 , while Bob can move in the center to reach position 4e, which is equivalent to 4b. In particular, none of these scenarios is better for Alice than the positions that she would reach by starting with 2a, and the only remaining possibility is that Alice could make her first two moves in adjacent corners. This proves the lemma. \square

Lemma 6.6. *Suppose center is not an optimal first move for Alice, for some fixed total number of chips. Then every move in Figure 2 is optimal.*

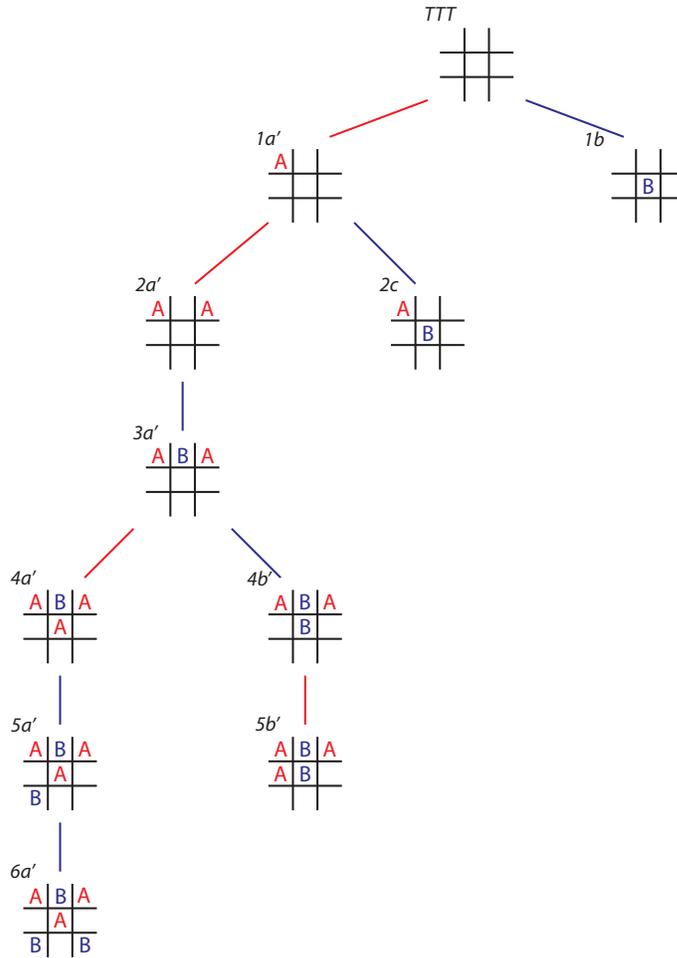


Figure 2: Optimal moves for TTT with five chips

Proof. Bob's move from $5a'$ is a counterattack, Alice's first move and her move from $1a'$ are optimal by Lemma 6.5, and Bob's first move is optimal by Theorem 6.1. We analyze the remaining six moves as follows.

Bob's move from $4a'$: If Bob does not block one of Alice's threats then he must win the following two moves as well. However, if he gets three moves in a row then he can also win by blocking with the first one.

Alice's move from $4b'$: Wherever Alice moves, the resulting position is equivalent to \mathbf{E} .

Alice's move from $3a'$: Alice's move in the center creates a position equivalent to $\mathbf{A} \wedge (\mathbf{A} \wedge B^2)$. If Alice plays anywhere else, then Bob moves in the center to create a position that is no better for Alice than \mathbf{E} , and hence worse than $\mathbf{A} \wedge B^2$. Indeed, if Alice does not move in the center column, then Bob's move in the center creates a threat, and if Alice moves in the bottom then Bob's move in the center creates a position

A	B	A
	B	
		A

that is equivalent to \mathbf{E} by Example 5.3.

Bob's move from 3d': Moving center creates a position equivalent to B^2 , or $\mathbf{E} \wedge \mathbf{B}$. If Bob moves anywhere else, then Alice can move center to reach a position in which she has a threat and Bob does not, which is worse for Bob than \mathbf{E} .

Bob's move from 2a': If Bob does not block then he must get the next move as well. If he blocks on the next move, he may as well have blocked the first time, and if he doesn't block on his second move, then he has to get a third move in a row. With three moves in a row, he can win after blocking with his first move.

Bob's move from 1d': If Bob plays anywhere other than center or bottom right then Alice can play in the center to reach a position that is no better for Bob than 3d, Meanwhile, if Bob gets two moves in a row he cannot prevent Alice from winning if she gets the next two moves. Therefore center is better than any other option except possible bottom right.

Suppose Bob plays bottom right. Then Alice can play top right and the best Bob can do is to block to reach a position

A	B	A
		B

that is equivalent to \mathbf{E} (Example 5.3). Again this is worse for Bob than if he had played center. □

6.2 Chip tables

Having determined all of the optimal moves for bidding Tic-Tac-Toe, except for Alice's first move which we have reduced to two possibilities, we now compute the total number of chips that Alice needs to win from each position that can be reached by these optimal moves. These computations are straightforward applications of the recursion given by Theorem 3.7 and the periodicity results in Theorem 3.13 and 3.14. As a consequence of these computations, we find that Alice's optimal first move depends on the total number of chips, as follows.

Theorem 6.7. *Center is an optimal first move for Alice if and only if the total number of chips is not five. Corner is an optimal first move for Alice if and only if the total number of chips is 0, 1, 2, 3, 4, 5, 6, 7, 9, 11, 12, 13, 14, 19, 20, 22, or 26.*

Proof. We prove the theorem by explicitly computing the number of chips that Alice needs to win assuming a first move in the center or a first move in the corner, for all possible

chip totals. Let $f'(TTT, k)$ be the number of chips that Alice needs to win, assuming a first move in the center, and $f''(TTT, k)$ be the number of chips that Alice needs to win, assuming she makes the first move in the corner, if she wins the first bid.

$256n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	1	1	1*	2*	2*	3*	4	4	4*	5*	5*	6*
12+	7	7	8	8*	9	9	10	10*	11	11*	12*	12*
24+	13	13	14	14*	15	15*	16	17	17*	17*	18	18*
36+	19*	20	20*	20*	21*	21*	22	23	23*	24	24*	25
48+	25*	26	26	27	27*	28*	28*	29	29*	30	30*	31*
60+	31*	32	33	33	34	34*	34*	35*	36	36*	37	37*
72+	37*	38*	39	40	40	40*	41	41*	42	42*	43	44
84+	44*	44*	45*	45*	46	46*	47*	48	48*	48*	49	50
96+	50*	51	51*	52	53	53	53*	53*	54*	55	55*	56
108+	57	57	57*	58	59	59	59*	60*	60*	61*	62	62
120+	62*	63*	63*	64*	65	65	66	66*	67*	67*	68	69
132+	69	70	70*	70*	71	72	72	73	73*	73*	74*	75
144+	75*	75*	76*	77	77*	78	79	79	79*	79*	80*	81
156+	81*	82	82*	83*	84	84	84*	85	86	86*	87	87
168+	88	88	88*	89*	90	90*	91	91*	92	92*	92*	93*
180+	94	95	95	95*	96	96*	97	98	98	98*	99*	99*
192+	100*	101	101	102	102*	103	103*	104	104	105	105*	106*
204+	106*	107	107*	108	108*	109	109*	110*	111	111	112	112
216+	112*	113	114	114*	115	115	115*	116*	117	117*	118	118*
228+	119*	119*	120	120	121	121*	122	122*	123*	123*	124	124*
240+	125*	125*	126	127	127	128	128*	128*	129	130	130	131
252+	131*	131*	132*	133								

Figure 3: Critical thresholds if Alice moves in the center; $f'(TTT, 256n+) = 133n+ \dots$

The entries in the interior of the preceding table are critical thresholds, while the position of the entry determines the total number of chips. For instance, the last entry in the second row is 12^* , which means that the critical threshold $f'(TTT, 256n + 23)$ is $133n + 12^*$, for any nonnegative integer n .

$256n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	1	1	1*	2*	2*	3	4	4	5	5*	6	6*
12+	7	7	8	9	9*	9*	10*	10*	11	12	12*	13
24+	13*	13*	14	15	15*	16	16*	17*	18	18	18*	19*
36+	20	20*	21	21*	22	22*	23	23*	24	24*	25	26
48+	26*	27	27	27*	28*	29	29*	30	31	31	31*	32*
60+	32*	33	34	34	35	35*	35*	36*	37	37	38	38*
72+	39	39*	40*	41	41	41*	42	43	43*	44	44*	45
84+	45*	46	46*	47	47*	48	48*	49*	50	50	50*	51*
96+	52	52*	53	54	54*	54*	55	55*	56	57	57*	57*
108+	58*	58*	59	60	61	61	61*	62	62*	63	64	64
120+	65	65*	65*	66*	67	67	68	68*	69*	69*	70	71
132+	71	71*	72*	72*	73*	74	74*	75	75*	75*	76*	77*
144+	78	78	79	79	79*	80*	81	81*	82	82	82*	84*
156+	84	84*	85	86	86*	86*	87	88	88*	89	89*	90
168+	90*	91	91*	92	92*	93	93*	94*	95	95*	95*	96
180+	97	97*	98	98*	99*	99*	100	101	101	101*	102*	102*
192+	103*	104	104	105	105*	105*	106*	107	107*	108	109	109*
204+	109*	110	110*	111*	112	112*	113	113*	114	114*	115	115*
216+	116	116*	117	118	118*	118*	119	120	120*	121	121*	122*
228+	123	123	123*	124	124*	125*	126	126	127	127	127*	128*
240+	129*	129*	130	130*	131	131*	132*	132*	133*	134	134	135
252+	135*	135*	136*	137								

Figure 4: Critical thresholds if Alice moves in the corner; $f''(TTT, 256n+) = 137n+ \dots$

It is straightforward to compare the tables in Figures 3 and 4 to see that $f'(TTT, k)$ is greater than or equal to $f''(TTT, k)$ unless $k = 5$, with equality only for $k = 0, 1, 2, 3, 4, 6, 7, 9, 11, 12, 13, 14, 19, 20, 22$, or 26 , which proves Theorem 6.7. These tables are computed by working backwards from ending positions. For completeness, we include tables of critical thresholds for the other positions in the trees of possible optimal moves in Figures 1 and 2. Positions $4f$ and $5b'$ are equivalent to \mathbf{E} , position $4a$ is equivalent to A^2 , positions $3e, 5c, 5d, 5f, 6a, 6b, 6c$, and $4b'$ are equivalent to B^2 ; see Example 3.8 for tables of critical thresholds for these positions. The critical thresholds for the remaining positions are as follows.

Positions $5a, 5b, 5e, 4d, 4g$, and $5a'$ are all equivalent to $\mathbf{A} \wedge B^2$.

$k = 8n+$	0	1	2	3	4	5	6	7
$f(\mathbf{A} \wedge B^2, k) = 3n+$	0*	0*	1	1*	2	2	2*	3

Positions $4b, 4e, 3c$ and $5a'$ are equivalent to $(\mathbf{A} \wedge B^2) \wedge \mathbf{B}$. From each such position v , we have $f(TTT_v, 16n+) = 11n+ \dots$

$16n+$	0	1	2	3	4	5	6	7
$0+$	1	1^*	2	3	3^*	4	5	5^*
$8+$	6^*	7	7^*	8^*	9	9^*	10^*	11

Position 4c is equivalent to $(\mathbf{A} \wedge B^2) \wedge B^2$, and $f(TTT_{4c}, 16n+) = 9n+$:

$16n+$	0	1	2	3	4	5	6	7
$0+$	1	1	1^*	2^*	3	3	4	4^*
$8+$	5^*	5^*	6	7	7^*	7^*	8^*	9

Position 4h is equivalent to B^3 .

$k = 8n+$	0	1	2	3	4	5	6	7
$f(B^3, k) = 7n+$	1	2	3	3^*	4^*	5^*	6^*	7

From position 3a, we have $f(TTT_{3a}, 32n+) = 15n+$:

$32n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
$0+$	0^*	1	1^*	2	2	3	3^*	3^*	4	5	5	5^*
$12+$	6	6^*	7	7^*	8	8^*	9	9^*	9^*	10^*	11	11
$24+$	11^*	12^*	12^*	13	13^*	14	14^*	15				

From position 3b, we have $f(TTT_{3b}, 32n+) = 9n+$:

$32n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
$0+$	0^*	0^*	0^*	1	1^*	1^*	2	2	2^*	2^*	3	3^*
$12+$	3^*	3^*	4	4^*	5	5	5	5^*	6	6	6^*	6^*
$24+$	7	7	7^*	8	8	8	8^*	9				

From position 3d, we have $f(TTT_{3d}, 32n+) = 11n+$:

$32n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
$0+$	0^*	0^*	1	1^*	1^*	2	2^*	2^*	3	3^*	3^*	4
$12+$	4^*	4^*	5	5^*	6	6	6^*	7	7	7^*	8	8
$24+$	8^*	9	9	9^*	10	10	10^*	11				

Position 3f is equivalent to $(\mathbf{A} \wedge B^2) \wedge B^3$.

$k = 8n+$	0	1	2	3	4	5	6	7
$f((\mathbf{A} \wedge B^2) \wedge B^3, k) = 5n+$	1	1^*	2	2^*	3	3^*	4^*	5

From position $2a$, we have $f(TTT_{2a}, 64n+) = 15n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	0	0*	0*	1	1	1*	1*	1*	2	2*	2*	2*
12+	3	3	3*	3*	4	4	4*	4*	4*	5	5*	5*
24+	5*	6	6	6*	6*	7	7	7*	7*	8	8	8*
36+	8*	9	9	9	9*	10	10	10	10*	10*	11	11
48+	11*	11*	12	12	12	12*	13	13	13	13*	13*	14
60+	14	14*	14*	15								

From position $2b$, we have $f(TTT_{2b}, 64n+) = 31n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	1	1	1*	2	2*	3	3*	3*	4*	5	5	6
12+	6*	6*	7	8	8*	8*	9	10	10	10*	11*	11*
24+	12	12*	13	13*	14	14	15	15*	16*	16*	17	17*
36+	18	18*	19	19	20	20*	20*	21*	22	22	22*	23*
48+	24	24	24*	25*	25*	26	27	27	27*	28	28*	29
60+	29*	29*	30*	31								

From position $2c$, we have $f(TTT_{2c}, 64n+) = 35n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	1	1	1*	2*	3	3	4	4*	5	5*	6	6*
12+	7*	7*	8	9	9*	9*	10*	11	11*	12	12*	13
24+	14	14	14*	15*	16	16	17	17*	18*	18*	19	20
36+	20*	20*	21*	22	22*	23	23*	24	25	25	25*	26*
48+	27	27	28	28*	29	29*	30	30*	31*	31*	32	33
60+	33*	33*	34*	35								

From position $2d$, we have $f(TTT_{2d}, 16n+) = 13n+$:

$16n+$	0	1	2	3	4	5	6	7
0+	1	2	2*	3*	4	5	6	6*
8+	7*	8*	9	10	10*	11*	12*	13

From position $1a$, we have $f(TTT_{1a}, 64n+) = 23n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	0*	1	1	1*	1*	2*	2*	2*	3	4	4	4*
12+	4*	4*	5*	6	6	6	7	7*	7*	7*	8*	8*
24+	9	9	9*	10	10*	10*	11	11*	12	12	12*	13
36+	13*	13*	14	14	15	15	15	15*	16*	16*	16*	17
48+	18	18	18	18*	18*	19*	20	20	20	21	21	21*
60+	21*	22	22*	23								

From position $1b$, we have $f(TTT_{1b}, 128n+) = 87n+$:

$128n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	1	1*	2	3	3*	4	5	5*	6	7	7*	8*
12+	9	9*	10	11	12	12*	13	13*	14*	15	16	16
24+	17	17*	18*	19*	19*	20	21	22	23	23*	23*	24*
36+	25*	26	27	27	28	28*	29*	30	30*	31	32	33
48+	33*	34	34*	35*	36	37	37*	38	39	39*	40	41
60+	41*	42	43	43*	44*	45	45*	46*	47	47*	48*	49
72+	49*	50*	51	52	52*	53	53*	54*	55*	56	56*	57
84+	58	58*	59*	59*	60*	61	62	63	63	63*	64*	65*
96+	66*	67	67	68	69	69*	70*	70*	71*	72	73	73*
108+	74	74*	75*	76*	77	77*	78	79	79*	80*	81	81*
120+	82*	83	83*	84*	85	85*	86*	87				

From position $4a'$, we have $f(TTT_{4a'}, 16n+) = 3n+$:

$16n+$	0	1	2	3	4	5	6	7
0+	0	0	0*	0*	1	1	1	1*
8+	1*	1*	2	2	2*	2*	2*	3

From position $3a'$, we have $f(TT_{3a'}, 32n+) = 15n+$:

$32n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	0*	0*	1*	2	2*	2*	3	4	4*	4*	5	5*
12+	6*	6*	7	7*	8	8	9	9*	10	10	10*	11*
24+	12	12	12*	13	14	14	14*	15				

From position $2a'$, we have $f(TTT_{2a'}, 64n+) = 15n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	0	0	0*	1	1	1	1*	2	2	2	2*	2*
12+	3	3	3*	3*	4	4	4*	4*	5	5	5	5*
24+	6	6	6	6*	7	7	7	7*	7*	7*	8	8*
36+	8*	8*	9	9*	9*	9*	10	10	10*	10*	11	11
48+	11*	11*	12	12	12*	12*	12*	13	13*	13*	13*	14
60+	14*	14*	14*	15								

From position $1a'$, we have $f(TTT_{1a'}, 64n+) = 25n+$:

$64n+$	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9	+10	+11
0+	0*	0*	1	1*	2	2	3	3	3*	3*	4*	4*
12+	5	5	6	6*	6*	6*	7*	8	8	8*	8*	9*
24+	10	10	10	11	11*	11*	12	12*	13	13	13*	14*
36+	14*	14*	15	16	16	16*	16*	17	18	18	18	18*
48+	19*	19*	20	20	21	21	21*	21*	22*	22*	23	23*
60+	24	24	24*	25								

This concludes our analysis of bidding Tic-Tac-Toe. □

7 Appendix: other tie-breaking methods

Throughout this paper we have used the tie-breaking method introduced in Section 2.3, in which the player who holds the tie-breaking chip can either give the chip to his opponent to win a tie, or keep the tie-breaking chip and lose the tie. We chose this method because it seemed simple and natural, with the tie-breaking advantage passing back and forth between the players just like the ordinary bidding chips, and because it has the following convenient properties.

1. It is always advantageous to have the tie-breaking chip (Lemma 3.1).
2. The tie-breaking chip is worth less than an ordinary chip (Lemma 3.2).

Other natural tie-breaking methods are possible, and here we briefly consider a few alternatives.

Loser's Ball (the ϵ -chip). We could make a rule that the player who loses one bid wins any tie on the next bid. One way to think about this tie-breaking method is by introducing an ϵ -chip whose value is strictly between zero and one. The player who holds the ϵ -chip is required to bid it, so the bids are never tied. Whichever player loses the bid takes all of the chips that were bid, and hence has the ϵ -chip for the next round. This method has the mildly unpleasant feature that holding the ϵ -chip is usually but not always an advantage. For instance, in the game where the second player to move wins,

if both players start with an equal number of chips then the player who starts without the ϵ -chip wins. In particular, the possible chip counts in $\mathbb{N} \cup \mathbb{N} + \epsilon$ with respect to this tie-breaking method are not totally ordered.

Make-it Take-it (the $-\epsilon$ -chip). A better idea than Loser's Ball is to make a rule that the player who wins one bid also wins any ties on the next bid. This method is satisfying, since it penalizes players for losing bids and often leads to taunting. One way to think about this tie-breaking method is by introducing a $-\epsilon$ chip whose value is strictly between zero and minus one. The player who holds the $-\epsilon$ -chip is required to bid it, so the bids are never tied. Whichever player loses the bid takes the $-\epsilon$ -chip for the next round, and hence loses any tie. Arguments similar to those given in Section 3 show that it is always an advantage *not* to have the $-\epsilon$ -chip and that it is always a good idea to accept the $-\epsilon$ -chip from your opponent together with an ordinary chip. Therefore, the possible chip counts in $\mathbb{N} \cup \mathbb{N} - \epsilon$ are totally ordered, and results analogous to those in Section 3 go through without major changes, except that the analogue of the recursion in Theorem 3.7 is given by

$$f(G_v, k) = \left\lfloor \frac{|f_A(G_v, k)| + |f_B(G_v, k)|}{2} \right\rfloor + \delta,$$

where

$$\delta = \begin{cases} 0 & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is even, and } f_B(G_v, k) \in \mathbb{N}. \\ -\epsilon & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is even, and } f_B(G_v, k) \in \mathbb{N} - \epsilon. \\ 1 - \epsilon & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is odd, and } f_A(G_v, k) \in \mathbb{N}. \\ 0 & \text{if } |f_A(G_v, k)| + |f_B(G_v, k)| \text{ is odd, and } f_A(G_v, k) \in \mathbb{N} - \epsilon. \end{cases}$$

The main disadvantage to the Make-it Take-it tie-breaking method seems to be that this description involves a lot of minus signs.

Ladies First. Suppose Alice wins all ties. The longer Alice and Bob play, the more the effects of Alice's advantage accumulate, so we should expect that Bob will have trouble if the game goes too long.

Let G be the game that Alice wins if she wins the first move, or both of the next two moves, or all of the next three moves, or the i -th move for all $n(n-1)/2 < i \leq n(n+1)/2$, for any n . Then Bob cannot win, but we can hope to prolong the game into an infinite draw. If G is truncated after $n(n+1)/2$ moves, then the critical threshold is

$$R(G[n(n+1)/2]) = \prod_{j=1}^n \left(1 - \frac{1}{2^j}\right).$$

In particular, the critical threshold for Richman play, with real-valued bidding is given by the infinite product $R(G) = \prod_{j>0} (1 - 1/2^j)$, which converges to a positive number. However, for discrete bidding with Ladies First tie-breaking, Alice wins no matter what the chip count. Alice can just bid zero indefinitely. Bob will have to give her a chip sometime between the moves numbered $n(n-1)/2$ and $n(n+1)/2$ for every n , until he runs out of chips. When Bob runs out of chips, Alice keeps bidding zero and winning all

ties, so she eventually wins. In particular, Theorem 3.11 does not extend to locally finite games with Ladies First tie-breaking.

References

- [Bea08a] J. Beasley, *Bidding Chess*, *Variant Chess* **7** (2008), no. 57, 42–44.
- [Bea08b] ———, *Isolated pawns*, *Variant Chess* **7** (2008), no. 58, 71.
- [Ber96] E. Berlekamp, *The economist's view of combinatorial games*, *Games of no chance* (Berkeley, CA, 1994), Math. Sci. Res. Inst. Publ., vol. 29, Cambridge Univ. Press, Cambridge, 1996, pp. 365–405.
- [Con76] J. H. Conway, *On numbers and games*, Academic Press [Harcourt Brace Jovanovich Publishers], London, 1976, London Mathematical Society Monographs, No. 6.
- [Hwa06] T. Hwa, personal communication by email, October 26, 2006.
- [LLP⁺99] A. Lazarus, D. Loeb, J. Propp, W. Stromquist, and D. Ullman, *Combinatorial games under auction play*, *Games Econom. Behav.* **27** (1999), no. 2, 229–264.
- [LLPU96] A. Lazarus, D. Loeb, J. Propp, and D. Ullman, *Richman games*, *Games of no chance* (Berkeley, CA, 1994), Math. Sci. Res. Inst. Publ., vol. 29, Cambridge Univ. Press, Cambridge, 1996, pp. 439–449.
- [PR08] S. Payne and E. Robeva, *Artificial intelligence for Bidding Hex*, To appear in *Games of no chance*. arXiv:0812.3677v1, 2008.
- [PSSW07] Y. Peres, O. Schramm, S. Sheffield, and D. Wilson, *Random-turn hex and other selection games*, *Amer. Math. Monthly* **114** (2007), no. 5, 373–387.