ON THE COMBINATORIAL VALUE OF HEX POSITIONS

Peter Selinger
Department of Mathematics and Statistics, Dalhousie University, Nova Scotia, Canada

Abstract

We develop a theory of combinatorial games that is appropriate for describing positions in Hex and other monotone set coloring games. We consider two natural conditions on such games: a game is *monotone* if all moves available to both players are good, and *passable* if in each position, at least one player has at least one good move available. The latter condition is equivalent to saying that if passing were permitted, no player would benefit from passing. Clearly every monotone game is passable, and we prove that the converse holds up to equivalence of games. We give some examples of how this theory can be applied to the analysis of Hex positions.

1. Introduction

Hex is a strategic perfect information game for two players, invented in 1942 by Piet Hein and later independently discovered by John Nash [10, 14]. Hex is attractive because its rules are extremely simple, yet the game possesses a surprising amount of strategic depth. Hex is played on a parallelogram-shaped board, typically of size $n \times n$, that is made of hexagonal fields called *hexes* or *cells*, as shown in Figure 1(a). One pair of opposing edges is colored black, and the other pair is colored white. The players, called Black and White, alternate placing a stone of their color on any empty hex, with Black starting. Once placed, these stones are never moved or removed. Each player's objective is to connect the two edges of that player's color with an unbroken chain of that player's stones. It is easy to see that if the board is completely filled with stones, exactly one player will have such a connection; to see this, consider the connected component of one of Black's edges. This either includes the other black edge, in which case Black has a connection, or it does not, in which case the white stones along its boundary form a connection for White. Consequently, there are no draws in the game of Hex; exactly one player will win. Also, the game always ends in a finite number of moves. Conventionally, the game ends as soon as one player has completed a connection. However, since the winner cannot change after that point, it would be equivalent, and sometimes simpler from

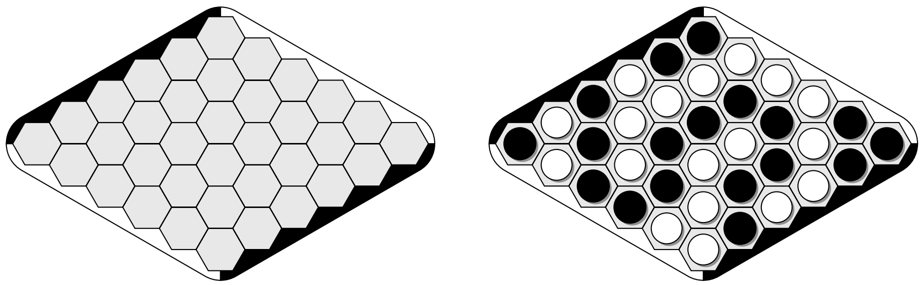

Figure 1: (a) A Hex board of size 6×6. (b) A winning position for Black.

a theoretical point of view, to require the game to continue until all hexes have been filled. Figure 1(b) shows an example of a completed game that has been won by Black.

Hex has many interesting properties. For example, there is an easy non-constructive proof that Hex on a board of size $n \times n$ is a win for the first player [14]; however, no actual winning strategy is known for $n \geqslant 11$. The proof is by a strategy stealing argument: if the second player had a winning strategy, the first player could mentally place an opponent's stone on the board and then play the same strategy. Determining whether a winning strategy exists in a given position on a board of arbitrary size is known to be a PSPACE-complete problem [16]. Of interest to us in this paper is another easy property: Hex is a monotone game. Informally, by this we mean that additional black stones on the board can only help Black, and additional white stones can only help White. A consequence is that making a move is always at least as good as making no move; although the ability to pass is not typically part of the rules of Hex, passing would never be to a player's advantage if it were permitted.

In practice, the rules of Hex, as stated above, give too much of an advantage to the first player. Although no explicit winning strategy for Black is known for large enough board sizes, Black still ends up winning a majority of games. For this reason, the additional *swap rule* is employed. It states that after Black's first move, White may choose to switch colors. This rule incentivizes Black to play a first move that is as fair as possible, and therefore leads to a more balanced game. The swap rule is not relevant for the rest of this paper and we will not consider it here.

Combinatorial game theory is a formalism for the study of sequential perfect information games that was introduced by Conway [5] and Berlekamp, Conway, and Guy [1]. Its roots go back further to the study of impartial games such as Nim. Combinatorial game theory was initially developed for *normal play* games, in which the first player who is unable to make a move loses the game. It has also been adapted to a variety of other conventions, such as *misère play*, in which the first player who is unable to move wins, or *scoring games*, in which the final outcome is a numerical score.

In this paper, we develop a variant of combinatorial game theory that is appropriate for Hex and other monotone set coloring games. The main difference between the games we describe here and other kinds of combinatorial games is the winning condition. We already mentioned that in normal play games, the loser is the first player who cannot move. In some games like Gomoku (also known as "Five in a Row"), the winner is the first player who achieves a winning condition, such as building a straight line of 5 of the player's stones. By contrast, although Hex also has a winning condition, it is immaterial *when* the winning condition is achieved. As we will see, this feature, along with monotonicity, gives rise to a particular family of combinatorial games with attractive mathematical properties.

An earlier version of this paper appeared on the arXiv at [17]. Several problems that were left open in the earlier version have since been solved, including in [7] and [6]. The present paper has been updated accordingly.

1.1. Related Work

Set coloring games were considered by van Rijswijck [20], and earlier, under the name "division games", by Yamasaki [21]. These works only considered games with two atomic outcomes (i.e., Black wins or White wins), rather than local games over a partially ordered set of outcomes as we do here. Monotone games, i.e., those where making a move is always at least as good as passing, are sometimes called "regular" games in the literature.

Some authors, such as van Rijswijck [20] and Henderson and Hayward [12], have applied ideas from combinatorial game theory to Hex and other set coloring games. For example, they considered notions of dominated and reversible moves appropriate to these games. However, while these works were in the "spirit" of combinatorial games, they did not develop the "letter" of an actual combinatorial game theory for Hex. For example, Henderson and Hayward explicitly state that "Hex is not a combinatorial game in the strictest sense", and consider only the *outcome class* of games (such as positive, negative, or fuzzy), rather than their full combinatorial value. The present paper remedies this situation by providing such a theory.

Play in local Hex regions, as opposed to on the entire board, has also been considered in the literature, though not at the same level of generality as we do here. For example, what we call an n-terminal region is called the carrier of a $2n$-sided decomposition by Henderson and Hayward [12].

The notion of combinatorial games we develop in this paper is closely related to the concept of passing, i.e., allowing a player not to make a move. In general, the notion of passing is problematic in combinatorial game theory, because it creates the possibility of games with infinite plays. The theory of *loopy games* [1, Ch. 11] was developed to deal with such infinite games. By contrast, the *passable games* we develop in this paper are not loopy: in fact, passing is not allowed in these games, all plays are finite, and therefore ordinary well-founded Conway induction can always be used on them. However, passable games have the property that they are *equivalent* to games in which passing is permitted. The defining property of such games is that no player has an incentive to pass, and therefore it would not alter the nature of the game if passing were permitted.

Another related concept from combinatorial game theory is *temperature*, a notion that quantifies, roughly speaking, how motivated the players are to make a move [1, Ch. 6]. A game, or a component of a game, is *hot* if the players can gain an advantage by making a move, and *cold* if the players would prefer not to move. (Thus, the hotter a game is, the more urgency the players feel to move in it. In cold games, players only make a move when they have no other choice.) In these terms, Hex and all monotone set coloring games are hot: making the next move is never disadvantageous. We do not explicitly use the concept of temperature in this paper, except to note that there is an upper limit on how hot a Hex position can be (Proposition 11.4).

For book-length treatments of modern Hex strategy, see [3] and especially [18]. Seymour has also created an excellent collection of Hex puzzles [19]. For a detailed scholarly study of the history of Hex, see [9].

1.2. Contents

The rest of this paper is organized as follows. In Section 2, we consider play in local Hex regions and define the poset of outcomes for such a region. In Section 3, we describe the class of set coloring games to which our combinatorial game theory applies. In Section 4, we describe a class of combinatorial games appropriate for Hex and other monotone set coloring games. These games are defined over a given poset of atomic outcomes. We motivate and define the order relations $\leqslant$ and $\lhd$, whose definition is subtle and is the main technical vehicle making these games work. We show that these game admit canonical forms, and we define the class of monotone games. In Section 5, we define relations of left and right order and equivalence, which are useful for technical reasons. In Section 6, we define passable games and prove the fundamental theorem of monotone games, which states that monotone games and passable games are the same up to equivalence. In Section 7, we prove certain special properties of games over linearly ordered sets of atoms; in particular, on this class of games, canonical forms of monotone games are monotone. In Section 8, we show that certain common operations on combinatorial games, and especially the sum operation, can be generalized to passable games, and we consider copy-cat strategies in this context. In Section 9, we introduce the concept of global decisiveness and show that otherwise non-equivalent games sometimes become equivalent in its presence. In Section 10, we enumerate all passable game

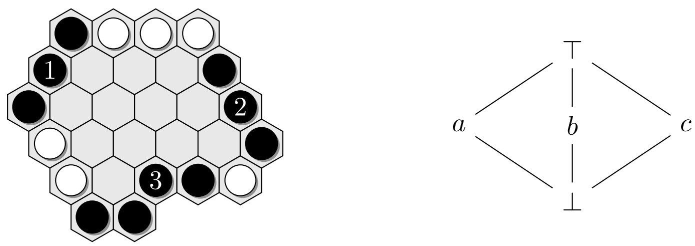

Figure 2: (a) A 3-terminal Hex region. (b) The outcome poset for this region.

values over certain small atom sets, and show that in all other cases, the set of passable game values is infinite. In Section 11, we show that many, but not all, abstract passable game values are realizable as Hex positions. In Section 12, we give an application of this theory by computing the size of the minimal virtual connection on Hex boards of size $4 \times n$, for all n. Finally, in Section 13, we list some open problems and point to avenues for future work.

2. Local Play in Hex

At the end of a game of Hex, there are only two possible outcomes, which we denote by $\top$ (Black wins) and $\bot$ (White wins). However, when we are concerned with play in some local *region* of the board, the set of possible outcomes can be richer. We first illustrate this idea with some examples.

Consider the board region shown in Figure 2(a). We imagine that this region is part of a larger game. The region is entirely surrounded by black and white stones. Among these boundary stones, there are three connected groups of black stones, which we have labelled 1, 2, and 3. (Dually, there are also three connected groups of white stones.) We refer to this type of region as a *3-terminal region*. Once the enclosed region has been completely filled with stones, it is clear that the only thing within the region that can affect the final outcome of the surrounding game is which of the three black terminals are connected to each other (or equivalently, which of the three white terminals are connected to each other). We refer to this as the *outcome* of the region. There are five possible outcomes for a 3-terminal region. We denote them by $\{\bot, a, b, c, \top\}$, and they are defined as follows:

- $\bot$: None of Black's terminals are connected to each other.
- a: Terminals 2 and 3 are connected to each other, but terminal 1 is not.
- b: Terminals 1 and 3 are connected to each other, but terminal 2 is not.
- c: Terminals 1 and 2 are connected to each other, but terminal 3 is not.

- $\top$: All of Black's terminals are connected to each other.

Here is an example of each outcome:

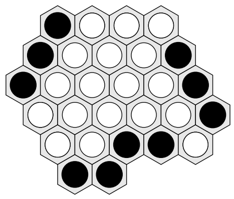 $= \bot$, 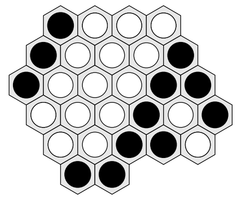 $= a$,

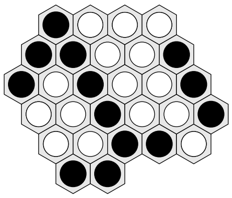 $= b$, 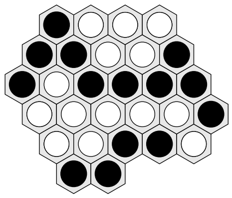 $= c$, 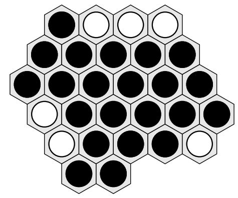 $= \top$.

There is a natural partial order on the set of outcomes of any region. Namely, if x and y are local outcomes, we say that $x \leqslant y$ if for every way of embedding the region in a larger board, and for every way of completely filling the rest of the board with stones, if the position with outcome x in the region is winning for Black, then so is the corresponding position with outcome y in the region. In case of a 3-terminal region, we have $\bot < a, b, c < \top$, where a, b, and c are incomparable. A Hasse diagram for this partially ordered set is shown in Figure 2(b). We emphasize that all 3-terminal regions have outcomes in the poset $\{\bot, a, b, c, \top\}$, no matter the shape or size of the region.

We note that in a 2-terminal region, there are only two possible outcomes: either Black or White connects their terminals. The entire Hex board forms a 2-terminal region, with the colored edges as the terminals. In a 4-terminal region, there are 14 possible outcomes, and in a 5-terminal region, there are 42 possible outcomes. More generally, the number of outcomes for an n-terminal region is equal to the number of non-crossing partitions of $\{1, \dots, n\}$, which is equal to the nth Catalan number $C(n) = \frac{(2n)!}{n!(n+1)!}$ [15].

As pointed out by Henderson and Hayward [12], it is not actually necessary for consecutive opposite-colored terminals in an n-terminal region to touch. It is sufficient for such terminals to be connected by the two-colored bridge pattern shown in Figure 3(a). This pattern serves to separate the cells marked a and b so that one is inside and the other outside the region. The reason for this is that we can, without loss of generality, consider a and b to be non-adjacent. Indeed, if they are both occupied by the same color, they are indirectly connected anyway, either via the black stone on the left or the white one on the right. For example, Figure 3(b) shows a 3-terminal region that is bounded by stones and two-colored bridges.

Another interesting phenomenon is that the boundary of a Hex region can include what we may call "invisible terminals". For example, in Figure 4(a), we have

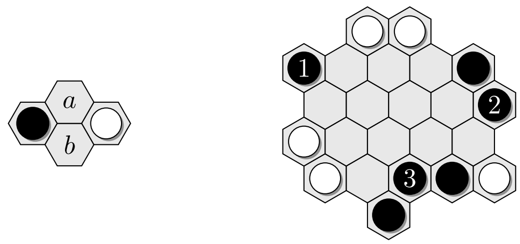

Figure 3: (a) A two-colored bridge. (b) A 3-terminal region surrounded by stones and two-colored bridges.

Figure 4: (a) An invisible white terminal next to a black stone. (b) A 4-terminal region with several invisible terminals.

inserted a thin white rectangle between two empty cells next to a black stone. Note that we may regard this rectangle as an additional (strangely shaped) board cell that is occupied by White. Indeed, doing so does not affect the connectivity of other cells: if a and b are both occupied by White, they are connected via the white rectangle, and if they are both occupied by Black, they are connected via the black stone. Since such thin rectangles are only imagined and not usually visible, we call them *invisible rectangles*. They can be white or black. When invisible rectangles form part of a region's boundary, as in Figure 4(b), we refer to them as *invisible terminals*. The two-colored bridge of Figure 3 is in fact a special case of an invisible terminal whose color does not matter.

We can also consider other types of regions besides n-terminal regions. For example, consider the region shown in Figure 5(a). This is a 2-terminal region whose boundary contains a gap marked "$*$". In this example, we do not consider the gap itself to be part of the region. Within the region, there are only three distinct outcomes:

- $\bot$: White's terminals are connected.
- a: Neither Black's nor White's terminals are connected within the region (but Black's terminals would be connected if the cell marked "$*$" were occupied by a black stone, and similarly for White's terminals).
- $\top$: Black's terminals are connected.

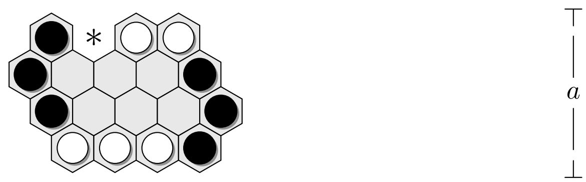

Figure 5: (a) A 2-terminal region with a gap. (b) The outcome poset for this region.

Here is an example of each outcome:

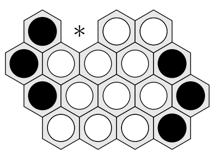 $= \bot$, 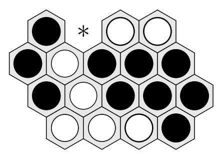 $= a$, 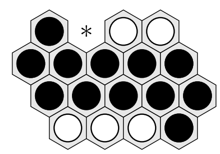 $= \top$.

The set of outcomes is linearly ordered; we have $\bot < a < \top$. The order is also shown in Figure 5(b).

Note that we could have alternatively regarded the region in Figure 5 as a 3-terminal region, by inserting two invisible terminals (a white one and a black one) between the region and the gap. However, this would have yielded a less precise outcome poset, because a generic 3-terminal region has five possible outcomes, whereas a 2-terminal region with gap only has three possible outcomes.

As another example, consider the region shown in Figure 6(a). This is a 3-terminal region, with the additional property that two of White's terminals are board edges (or connected to board edges). This edge condition affects the outcome poset: of the five possible outcomes $\{\bot, a, b, c, \top\}$ for a generic 3-terminal region, two now become equivalent. Namely, if Black's terminal 3 is connected to neither terminal 1 nor terminal 2, then White's two edges are connected, which means that White wins the surrounding game. In that case, it does not matter whether terminals 1 and 2 are connected to each other, so the outcomes $\bot$ (none of the Black's terminals are connected) and c (only terminals 1 and 2 are connected) are now equivalent. Consequently, this type of region has only four distinct outcomes $\{\bot, a, b, \top\}$, and their partial order is as shown in Figure 6(b). We call such a region a *fork*, because all that matters about the outcomes is what terminal 3 is connected to (left, right, both, or none).

In general, a Hex region consists of a set of cells, some of which may already be occupied by black or white stones (such as the boundaries in the above examples). We may or may not stipulate additional restrictions on how the region can be embedded in a larger game of Hex; for example, we may require certain stones to be connected to board edges. A *completion* of the region is an assignment of black and white stones to all of its empty cells. We define a preorder on the set of

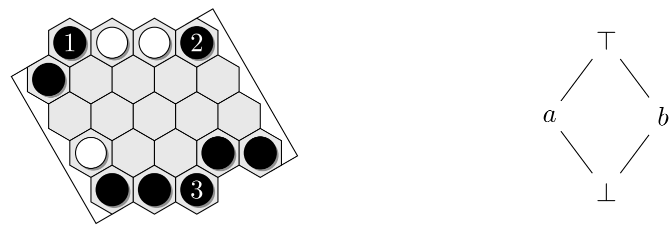

Figure 6: (a) A 3-terminal region with edges, or "fork". (b) The outcome poset for this region.

completions as outlined above, i.e., $x \leqslant y$ if for every allowable embedding of the region in a larger game of Hex, and for every completion of the remaining game, if x gives a winning position for Black then so does y. Two completions are *equivalent* is $x \leqslant y$ and $y \leqslant x$, and an *outcome* for the region is an equivalence class of completions.

3. Set Coloring Games

Before we develop a combinatorial game theory for Hex, we describe a more general class of games to which this theory can be applied. These are the set coloring games of van Rijswijck [20]. We will only be concerned with set coloring games with two players, each of whom uses one color. Rather than just considering games in which one player wins and the other loses, we will consider set coloring games over some arbitrary poset of outcomes. They are defined as follows. Let $\mathbb{B} = \{\bot, \top\}$ be the set of booleans. Here, $\bot$ denotes "false" or "bottom", and $\top$ denotes "true" or "top". If X and Y are sets, we write Y^X for the set of functions from X to Y.

Definition 3.1 (Set coloring game)**.** Let A be a partially ordered set, whose elements we call *outcomes*. A *set coloring game* over A is given by a finite set X of *cells*, and a function $\pi : \mathbb{B}^X \to A$ called a *payoff function*. It is played as follows: there is a game board, initially empty, whose cells are in one-to-one correspondence with the elements of X. The players, whom we call Black and White, take turns, with Black starting. Alternately, Black colors one cell with color $\top$, and White colors one cell with color $\bot$. This continues until all cells are colored. A fully colored game board corresponds to a function from X to $\mathbb{B}$, i.e., an element $f \in \mathbb{B}^X$. Then the outcome of the game is $\pi(f) \in A$.

We call a completely filled board an *atomic* position. An atomic position has no moves for either player. All other positions have at least one move for each player. Of particular interest to us are set coloring games that are monotone.

Definition 3.2 (Monotone set coloring game)**.** We equip the set $\mathbb{B}$ with the natural order $\bot < \top$. Note that this induces a partial order on $\mathbb{B}^X$ given pointwise, i.e., $f \leqslant g$ if for all $x \in X$, we have $f(x) \leqslant g(x)$. A set coloring game (X, π) is called *monotone* if its payoff function π is a monotone function with respect to this order.

Evidently, Hex is a monotone set coloring game over $\mathbb{B}$, since the winning condition (Black has a connection between her two edges) remains valid when changing any number of stones from white to black. Moreover, any Hex region is a monotone set coloring game over its outcome poset. For example, a 3-terminal region, as defined in Section 2, is a monotone set coloring game over $A = \{\bot, a, b, c, \top\}$.

We can also consider certain generalizations of Hex that are monotone set coloring games. A *planar connection game* is like Hex, except the board is tiled by arbitrary polygonal cells, not necessarily hexagons. To ensure that there is a unique winner, we must require that at most three cells meet at any point. See [4] for an example. Another, even more general example of a monotone set coloring game is the *vertex Shannon game*. It is played on an arbitrary undirected graph with two distinguished vertices called *terminals*. The players alternately color a non-terminal vertex in their color until all vertices are colored; Black wins if and only if in the resulting coloring, the two terminals are connected by a path of black vertices. (Unlike Hex and other planar connection games, the class of vertex Shannon games is not obviously self-dual, i.e., it is not obviously invariant under switching the roles of the two players.)

4. A Class of Combinatorial Games

We are now ready to develop a notion of combinatorial games that is appropriate for Hex and other monotone set coloring games. In many ways, the theory of these games is similar to standard combinatorial game theory, say for normal play games. As already mentioned in the introduction, the main difference is how the games end. Our games end when an atomic position is reached, and in any non-atomic position, there is at least one possible move for each player (so that the game can never end due to a player's inability to make a move). In addition, as we will see, our games are designed so that when a is an atomic outcome, the games a and $\{a \mid a\}$ are equivalent. This reflects the fact that in monotone set coloring games, once the outcome in a region is determined, it does not matter if the players are allowed additional useless moves. In this regard, our theory differs from the class of games, such as Gomoku, where the outcome is determined by which player *first* achieves a winning condition. In such games, $\{\top \mid \top\}$ is not equivalent to $\top$, because if one player needs one more move to achieve a win, the other player might win first by playing in another region of the game.

Our notion of games has different properties than normal play games. For example, in normal play games, $G \lhd H$ holds if and only if $H \nleqslant G$, whereas in our games, we often have $G \lhd G$. Thus, while many aspects of our proofs are the same as for other kinds of combinatorial games, some crucial details are different. Consequently, we give full proofs even of results that seem elementary. Where proofs have been omitted, they are completely routine.

4.1. Games over an Outcome Poset

Definition 4.1 (Game over a poset)**.** Let A be a partially ordered set, whose elements we call *atoms* or *outcomes*. Games over A are defined inductively:

- For every $a \in A$, $[a]$ is a game, called an *atomic game*. We usually write a instead of $[a]$ when no confusion arises.
- If L and R are non-empty sets of games, then $G = \{L \mid R\}$ is a game, called a *composite game*. As usual, L and R are called the sets of *left* and *right options* of G, respectively.

The fact that it is an inductive definition means that games are freely generated by the two rules above. When we say that two games are *equal*, we mean that they are literally the same game, i.e., they are both atomic with equal atoms, or they are both composite with equal sets of left options and equal sets of right options. This should not be confused with the notion of *equivalence* of games that we will define later. Note that we have required non-atomic games to have at least one left option and at least one right option. Such games are also called "all-small" or "dicot" in the context of normal-play and misère games, respectively [5, 13].

We will follow the usual notational conventions of combinatorial game theory. For example, we write $\{x, y \mid z, w\}$ instead of $\{\{x, y\} \mid \{z, w\}\}$. We often write G^L and G^R for a typical left and right option of G, respectively. We sometimes write $G = \{G^L \mid G^R\}$ to indicate a game that has (possibly more than one) typical left and right option. As is usual in combinatorial game theory, the two players are called Left and Right. In set coloring games, we identify Black with Left and White with Right.

Definition 4.2 (Well-founded relation and Conway induction)**.** On the collection of all games over A, we define $\preccurlyeq$ to be the smallest reflexive transitive relation such that for any composite game G, we have $G^L \preccurlyeq G$ for all left options G^L of G and and $G^R \preccurlyeq G$ for all right options G^R of G. Then $\preccurlyeq$ is a well-founded partial order, i.e., it has no infinite strictly descending sequences. If $G \preccurlyeq H$, we say that G is a *position* of H. In other words, the positions of H are H itself, the options of H, the options of options of H, and so on. If $G \preccurlyeq H$ and $G \neq H$, we say that G is a *smaller* game than H. Informally, it means that G was defined "before" H. We can prove statements about games by induction on this well-founded relation. This kind of induction is often called *Conway induction*.

We will often (but not always) assume that the outcome set A has top and bottom elements, which we denote by $\top$ and $\bot$, respectively.

4.2. Order

We now define relations $G \leqslant H$ and $G \lhd H$ on games. As in standard combinatorial game theory, the intuition is that $G \leqslant H$ means that from Left's point of view, being first to move in H is at least as good as being first to move in G, and being second to move in H is at least as good as being second to move in G. Also, $G \lhd H$ means that from Left's point of view, being first to move in H is at least as good as being second to move in G. As usual, we take Left's point of view, i.e., "better" means "better for Left" unless stated otherwise.

In combinatorial game theory, it is common to define these relations by first defining the sum and negation operations on games, and then to take $G \leqslant H$ to mean $0 \leqslant H - G$. For reasons that will become apparent later, this definition is not convenient in our setting — among other things, because there is no game "0", and because the sum of games is not well-defined until we consider monotone games, which requires the order $\leqslant$ to be defined first. We therefore define the relations $G \leqslant H$ and $G \lhd H$ directly, without reference to sums.

Except for the atomic cases, the definitions of $\leqslant$ and $\lhd$ are the same as in standard combinatorial game theory. Since an atomic game has no left or right options, by convention, when G is atomic, we take any statement of the form "for all left options G^L" to be vacuously true, and any statement of the form "there exists a left option G^L" to be trivially false.

Definition 4.3. For games over a partially ordered set A, the relations $\leqslant$ and $\lhd$ are defined by mutual recursion as follows.

- $G \leqslant H$ if all three of the following conditions hold:
  1. All left options G^L satisfy $G^L \lhd H$, and
  2. all right options H^R satisfy $G \lhd H^R$, and
  3. if G or H is atomic, then $G \lhd H$.
- $G \lhd H$ if at least one of the following conditions holds:
  1. There exists a right option G^R such that $G^R \leqslant H$, or
  2. there exists a left option H^L such that $G \leqslant H^L$, or
  3. $G = [a]$ and $H = [b]$ are atomic and $a \leqslant b$.

Since this definition is at the heart of this paper, and since it is not self-evident that it captures the correct notions, some further explanations are in order. The fact that it is the "right" definition will also be substantiated by the theorems and applications that follow later.

The first thing to note is that clauses 1. and 2. in the definition of $\leqslant$ and $\lhd$ are exactly the standard ones that can be found, for example, in normal play games. Thus, if both games are composite, there is nothing unusual in this definition. The novelty lies in the treatment of atomic games.

To understand the atomic cases, consider the postulate that the game a should be equivalent to $\{a \mid a\}$, or in symbols, $a \simeq \{a \mid a\}$. The reason for the postulate is that it is true in monotone set coloring games. Specifically, in a set coloring game, a cell is called *dead* (or sometimes *negligible*) if the color of that cell does not affect the outcome of the game [2, 21]. The atomic game a represents a completely filled board, and the game $\{a \mid a\}$ represents a board with one remaining dead cell. In a monotone set coloring game, adding a dead cell does not change the strategic value of the game, because playing there is not to either player's advantage, and the cell will eventually be filled anyway.

If we postulate $a \simeq \{a \mid a\}$ and plug this into the usual recursive clauses for $\leqslant$ and $\lhd$, we get the following properties:

$$\begin{array}{lcl} a \leqslant H & \iff & a \lhd H \text{ and } \forall H^R . a \lhd H^R. \\ a \lhd H & \iff & a \leqslant H \text{ or } \exists H^L . a \leqslant H^L. \end{array} \tag{1}$$

Note, however, that we cannot take (1) as a definition, because it is circular: $a \leqslant H$ is described in terms of $a \lhd H$ and vice versa. It turns out, however, that the clauses in (1) have both a largest and a smallest solution, i.e., among the pairs of relations $(\leqslant, \lhd)$ satisfying (1), there is a largest and a smallest one. It turns out that the smallest solution is the one that we need. More generally, consider a circular system of boolean equations of the form

$$\begin{array}{lcccc} P & \iff & Q & \text{and} & R, \\ Q & \iff & P & \text{or} & S. \end{array} \tag{2}$$

It is easy to check that this system has exactly five solutions: $(P, Q, R, S) = (\top, \top, \top, \top)$, $(\bot, \top, \bot, \top)$, $(\top, \top, \top, \bot)$, $(\bot, \bot, \top, \bot)$, and $(\bot, \bot, \bot, \bot)$. In particular, unless $(R, S) = (\top, \bot)$, P and Q are uniquely determined. In case $(R, S) = (\top, \bot)$, there are the two solutions $(P, Q) = (\top, \top)$ and $(P, Q) = (\bot, \bot)$. If we choose the second of these, the equations (2) simplify to

$$\begin{array}{lcccc} P & \iff & Q & \text{and} & R, \\ Q & \iff & & & S. \end{array}$$

Applying this process to (1) leads us to define

$$\begin{array}{lcl} a \leqslant H & \iff & a \lhd H \text{ and } \forall H^R . a \lhd H^R, \\ a \lhd H & \iff & \exists H^L . a \leqslant H^L. \end{array}$$

Note that $a \leqslant H$ nevertheless implies $a \lhd H$, so that (1) is also satisfied.

For the case where both games are atomic, the postulates $a \simeq \{a \mid a\}$ and $b \simeq \{b \mid b\}$ immediately yield that $a \leqslant b$ if and only if $a \lhd b$. This makes sense because if a game is atomic, its value is already determined, so it no longer matters whose turn it is. The ordering of atomic games is given a priori by the poset structure on the set of atoms.

In summary, we arrive at the following desired properties:

- If both G and H are composite:
$$\begin{array}{lcl} G \leqslant H & \iff & (\forall G^L. G^L \lhd H) \text{ and } (\forall H^R. G \lhd H^R), \\ G \lhd H & \iff & (\exists G^R. G^R \leqslant H) \text{ or } (\exists H^L. G \leqslant H^L). \end{array}$$
- If G is composite and $H = [b]$ is atomic:
$$\begin{array}{lcl} G \leqslant [b] & \iff & (\forall G^L. G^L \lhd [b]) \text{ and } G \lhd [b], \\ G \lhd [b] & \iff & \exists G^R. G^R \leqslant [b]. \end{array}$$
- If $G = [a]$ is atomic and H is composite:
$$\begin{array}{lcr} [a] \leqslant H & \iff & [a] \lhd H \text{ and } \forall H^R. [a] \lhd H^R, \\ [a] \lhd H & \iff & \exists H^L. [a] \leqslant H^L. \end{array}$$
- If $G = [a]$ and $H = [b]$ are both atomic:
$$[a] \leqslant [b] \quad \iff \quad [a] \lhd [b] \quad \iff \quad a \leqslant b.$$

It is easy to check that Definition 4.3 is merely a more compact statement of these properties.

It is worth emphasizing that we are not saying that the games a and $\{a \mid a\}$ are actually equal, as this would give rise to games with infinite plays. Instead, we are merely saying that Definition 4.3 is motivated by the desire to make a and $\{a \mid a\}$ *equivalent*. As we will see in Lemma 4.11 below, this is indeed one of the consequences of Definition 4.3. Thus, we get the property $a \simeq \{a \mid a\}$ without having to consider infinite games or strategies.

4.3. Properties of the Order

The relations $\leqslant$ and $\lhd$ enjoy many of the usual properties that hold, say, in normal play games, but not all of them. For example, in normal play games, we have $G \lhd H$ if and only if $H \not\leqslant G$. This is not the case here: all atomic games satisfy $a \lhd a$ and $a \leqslant a$, and we will see that many non-atomic games satisfy $G \lhd G$ as well.

In this section, we prove some basic properties of the order relations. Throughout this section, we consider games over a fixed poset A of atoms.

Remark 4.4 (Duality)**.** For every statement about games, there is a dual statement, obtained by exchanging the roles of the left and right players (and replacing

$\leqslant$ by $\geqslant$ and so on). A property of games is valid if and only if its dual is valid, so in the following, we sometimes prove a lemma and then use both the lemma and its dual.

Lemma 4.5 (Reflexivity)**.** *The relation $\leqslant$ is reflexive.*

Proof. We prove $G \leqslant G$ by induction on G. For atomic games, we have $[a] \leqslant [a]$ by definition. Suppose that G is composite. To show $G \leqslant G$, first take any left option G^L. By the induction hypothesis, $G^L \leqslant G^L$, hence by definition of $\lhd$, we have $G^L \lhd G$. Similarly, take any right option G^R. By the induction hypothesis, $G^R \leqslant G^R$, hence by definition of $\lhd$, we have $G \lhd G^R$. It follows that $G \leqslant G$ as desired. □

Lemma 4.6 (Transitivity)**.** *For games G, H, K over A, we have:*

(a) $G \lhd H \leqslant K$ implies $G \lhd K$;

(b) $G \leqslant H \lhd K$ implies $G \lhd K$;

(c) $G \leqslant H \leqslant K$ implies $G \leqslant K$.

Proof. We prove all three properties by simultaneous induction. The base cases will be handled uniformly along with the non-base cases. Therefore, let G, H, K be any games (atomic or not), and assume that (a)–(c) are true for smaller triples of games. (Triples of games can be ordered componentwise with respect to the relation $\preccurlyeq$, which is again a well-founded relation.) We first prove (a), then (b), then (c) (so that the proof of (c) can use the results of (a) and (b), even for games of the same size).

(a) To show (a), assume $G \lhd H$ and $H \leqslant K$. We must show $G \lhd K$. From the definition of $G \lhd H$, there are three cases:

- Case 1: There exists some G^R such that $G^R \leqslant H$. Since G^R is smaller than G, by the induction hypothesis, from $G^R \leqslant H \leqslant K$, we get $G^R \leqslant K$. Therefore $G \lhd K$ as desired.
- Case 2: There exists some H^L such that $G \leqslant H^L$. By the definition of $H \leqslant K$, we have $H^L \lhd K$. Since H^L is smaller than H, by the induction hypothesis, from $G \leqslant H^L \lhd K$, we get $G \lhd K$ as desired.
- Case 3: $G = [a]$ and $H = [b]$ are atomic and $a \leqslant b$. From $H = [b] \leqslant K$, by definition, we get $[b] \lhd K$. By definition of $[b] \lhd K$, there are two possible cases:

  Case 3.1: There exists some K^L such that $[b] \leqslant K^L$. In that case, since K^L is smaller than K, by the induction hypothesis, from $G \leqslant [b] \leqslant K^L$, we get $G \leqslant K^L$, hence $G \lhd K$ as desired.

Case 3.2: $K = [c]$ is atomic and $b \leqslant c$. In this case, from $a \leqslant b \leqslant c$, we get $a \leqslant c$, hence $G = [a] \lhd [c] = K$ as desired.

(b) The proof of (b) is dual to that of (a).

(c) To show (c), assume $G \leqslant H$ and $H \leqslant K$. We must show $G \leqslant K$. Using the definition of $G \leqslant K$, there are three things we need to prove:

- We must show that every G^L satisfies $G^L \lhd K$. So let G^L be any left option of G. Since $G \leqslant H$ was assumed, we have $G^L \lhd H$. Since G^L is smaller than G, by the induction hypothesis, from $G^L \lhd H \leqslant K$, we get $G^L \lhd K$, as desired.
- We must show that every K^R satisfies $G \lhd K^R$. The proof is dual to the previous case.
- Finally, we must show that if either G or K is atomic, then $G \lhd K$. First assume G is atomic. Since G is atomic, from $G \leqslant H$, we get $G \lhd H$. By (a), from $G \lhd H \leqslant K$, we get $G \lhd K$, as desired. The case where K is atomic is dual. □

Corollary 4.7. *The relation $\leqslant$ forms a preorder on the collection of games over a poset A.*

At this point, it is customary to quotient out the equivalence relation induced by this preorder, i.e., to identify games G and H if $G \leqslant H$ and $H \leqslant G$. We will not do so here, because we will sometimes need to speak of properties of games that are not invariant under equivalence. We will continue to write $G = H$ to mean that G and H are literally equal, and we will say that G and H are *equivalent*, written $G \simeq H$, if $G \leqslant H$ and $H \leqslant G$. We will reserve the word *value*, or *combinatorial value*, to mean an equivalence class of games; thus, when we say two games have the same value, we mean that they are equivalent.

Lemma 4.8. *If $H \leqslant H'$, then $\{H, G^L \mid G^R\} \leqslant \{H', G^L \mid G^R\}$ and $\{G^L \mid G^R, H\} \leqslant \{G^L \mid G^R, H'\}$.*

Proof. This follows directly from the definitions. □

Lemma 4.9. *If G is composite, we have $G \leqslant \{H, G^L \mid G^R\}$. In other words, an additional left option can only help Left.*

Proof. This follows directly from the definitions. □

Lemma 4.10 (Gift horse lemma)**.** *If G is composite, we have $H \lhd G$ if and only if $G \simeq \{H, G^L \mid G^R\}$. Dually, if H is composite, we have $H \lhd G$ if and only if $H \simeq \{H^L \mid H^R, G\}$.*

Proof. It suffices to show the first claim, as the second one is its dual. For the right-to-left direction, assume $G \simeq \{H, G^L \mid G^R\}$. By definition of $\lhd$, we have $H \lhd \{H, G^L \mid G^R\}$. With $\{H, G^L \mid G^R\} \leqslant G$ and Lemma 4.6, this yields $H \lhd G$. For the left-to-right direction, assume $H \lhd G$. Since $G \leqslant \{H, G^L \mid G^R\}$ holds by Lemma 4.9, we only need to show $\{H, G^L \mid G^R\} \leqslant G$. But this follows directly from the definitions. □

Because of Lemma 4.10, when $H \lhd G$, we say that H is a *left gift horse* for G, and that G is a *right gift horse* for H.

Lemma 4.11 (All games are equivalent to composite games)**.** *For $a \in A$, we have $a \simeq \{a \mid a\}$.*

Proof. This follows directly from the definitions. □

Lemma 4.12. *If A has a top element $\top$, then all games G over A satisfy $G \leqslant \top$ and $G \lhd \top$. Dually, if A has a bottom element $\bot$, then all games G over A satisfy $\bot \leqslant G$ and $\bot \lhd G$.*

Proof. We show that $G \leqslant \top$ and $G \lhd \top$ by induction. If G is atomic, this is clear by definition. So assume G is composite. To show $G \lhd \top$, pick any right option G^R of G (recall that every composite game has at least one right option). By the induction hypothesis, $G^R \leqslant \top$, so by definition of $\lhd$, we have $G \lhd \top$. To show $G \leqslant \top$, first consider any left option G^L; by the induction hypothesis, we have $G^L \lhd \top$. Since $\top$ has no right options, and since we already showed $G \lhd \top$, the definition of $\leqslant$ then implies that $G \leqslant \top$ as desired. □

If the atom poset A has top and bottom elements, the partial order of values is particularly nice: it is a complete lattice, as the following proposition shows.

Proposition 4.13. *Assume the atom poset has a bottom element. Then any set of games has a least upper bound.*

Proof. The least upper bound of the empty set is $\bot$ by Lemma 4.12. Next, we show that any pair of games H, K has a least upper bound. By Lemma 4.11, we can assume without loss of generality that H and K are composite. Define

$$G = \{H^L, K^L \mid \{H, K \mid \bot\}\}.$$

We claim that G is the desired least upper bound. We first note that $G' = \{H, K \mid \bot\}$ is the smallest game such that $H, K \lhd G'$. Indeed, it is obvious that $H, K \lhd G'$. Suppose that G'' is some arbitrary game with $H, K \lhd G''$. Then $G' \leqslant G''$ follows directly from the definition of $\leqslant$.

Next, we show that $H, K \leqslant G$. Indeed, all their left options satisfy $H^L, K^L \lhd G$ by definition of G; conversely, the only right option of G satisfies $H, K \lhd \{H, K \mid \bot\}$.

Next, consider some arbitrary game M with $H, K \leqslant M$. We claim that $G \leqslant M$. Without loss of generality, assume that M is composite. First, consider any left option of G. They are H^L or K^L. Both satisfy $H^L, K^L \lhd M$ by the assumption that $H, K \leqslant M$. This shows the first part of $G \leqslant M$. Next, consider any right option M^R of M. We must show $G \lhd M^R$. From the assumption $H, K \leqslant M$, we get $H, K \lhd M^R$, and therefore $\{H, K \mid \bot\} \leqslant M^R$ by the "first note" above. This implies $G \lhd M^R$ as desired. We have shown two of the three conditions for $G \leqslant M$. Since neither G nor M is atomic, the third condition does not apply, so $G \leqslant M$. Therefore G is the least upper bound of H, K.

Finally, nothing in the above proof relies on the fact that we have exactly two games H, K. The same proof works for any finite or even infinite non-empty set of games. □

Note that by the dual of Proposition 4.13, greatest lower bounds also exist, provided that the atom poset has a top element. Also note that Proposition 4.13 only provides suprema of *sets* of games, whereas the collection of all games over a given atom poset is in general a proper class. For this reason, the usual trick of constructing a greatest lower bound as the least upper bound of all lower bounds does not work. In particular, we cannot use Proposition 4.13 to construct a top element by taking the supremum of all games.

4.4. Canonical Forms

The theory of canonical forms is similar to that for other kinds of combinatorial games, but some adjustments are needed to deal with atoms. The notion of bypassing a reversible option must be adjusted, and we need a new notion of passing option. As before, we consider games over a fixed atom poset A.

Definition 4.14 (Dominated option)**.** Suppose that H, K are distinct left options of G. We say that H *dominates* K if $K \leqslant H$. Dually, for distinct right options H, K of G, we say that H dominates K if $H \leqslant K$.

Lemma 4.15 (Removing dominated options)**.** *If G has a dominated left option K, then $G \simeq G'$, where G' is the result of removing the left option K from G. The dual statement for right options also holds.*

Proof. Suppose the left option K is dominated by another left option H. Since H is a left option of G', we have $H \lhd G'$. Since $K \leqslant H$, by Lemma 4.6, we have $K \lhd G'$. Therefore, K is a left gift horse for G' and $G \simeq G'$ by the Gift Horse Lemma 4.10. The proof for right options is dual. □

Definition 4.16 (Reversible option)**.** Suppose H is a left option of G. We say that H is *reversible* via K if H has a right option K such that $K \leqslant G$. (Therefore both G and H are necessarily composite.) Reversible right options are defined dually.

Lemma 4.17 (Bypassing reversible options)**.** *Suppose G is a game with a left option H that is reversible via K. Then $G \simeq G'$, where*

- $G' = \{K^L, G^L \mid G^R\}$*, when K is composite, and*
- $G' = \{K, G^L \mid G^R\}$*, when K is atomic.*

Here, G^L denotes all left options of G other than H, and K^L denotes all left options of K.

Proof. First, consider the case where K is composite. Since $K \leqslant G$, any left option of K is a gift horse for G, so that $G \simeq \{K^L, H, G^L \mid G^R\}$. To show that $\{K^L, H, G^L \mid G^R\} \simeq G'$, we need to show that H is a left gift horse for G', i.e., $H \lhd G'$. Since K is a right option of H, it suffices to show that $K \leqslant G'$. This requires showing three things:

- First, consider any left option K^L of K. We must show $K^L \lhd G'$. But K^L is, by definition, a left option of G', so $K^L \lhd G'$ as desired.
- Second, consider any right option G'^R of G'. We must show $K \lhd G'^R$. But G' has the same right options as G, so that $G'^R = G^R$ for some right option of G. By assumption, $K \leqslant G$, which implies $K \lhd G^R$ as desired.
- Third, assume K or G' is atomic. But this is not the case as we had assumed otherwise.

This concludes the proof in the case where K is composite. The case where K is atomic can be reduced to the previous case, because $K \simeq \{K \mid K\}$ by Lemma 4.11. □

The two cases of Lemma 4.17 are called *non-atomic* and *atomic reversibility*, respectively.

Definition 4.18 (Passing option)**.** Suppose H is a left option of G. We say that H is a *passing option* of G if $H \simeq G$.

Lemma 4.19 (Simplifying passing options)**.** *If H is a passing option of G, then $G \simeq H$.*

Proof. This holds by definition. □

Definition 4.20 (Canonical form)**.** A game G is in *canonical form* if G has no dominated, reversible, or passing options, and all left and right options of G are in canonical form.

Lemma 4.21 (Uniqueness of canonical form)**.** *If $G \simeq G'$ and both G and G' are in canonical form, then $G = G'$.*

Proof. We prove this by induction. Assume $G \simeq G'$ and both are in canonical form. To show $G = G'$, we must show that G and G' have exactly the same left options and right options, and moreover, if G and G' are atomic, then they are equal. The latter is clear, for if $G = [a]$ and $G' = [a']$, then $G \simeq G'$, by definition, holds if and only if $a = a'$.

Now consider any left option H of G. We will first show that $H \leqslant H'$ for some left option H' of G'. Since H is a left option of G, we have $H \lhd G$. Since $G \simeq G'$, we have $H \lhd G'$. By definition of $\lhd$, there are three possibilities:

- Case 1: $H^R \leqslant G'$ for some right option H^R of H. But since $G \simeq G'$, we then have $H^R \leqslant G$, so that H is reversible via H^R in G, contradicting the assumption that G is in canonical form.
- Case 2: $H \leqslant G'^L$ for some left option G'^L of G'. Then let $H' = G'^L$.
- Case 3: $H = [b]$ and $G' = [a']$ are atomic and $b \leqslant a'$. From $G' \leqslant G$, since G' is atomic, it follows that $G' \lhd G$, i.e., $[a'] \lhd G$. Since G is not atomic, it follows that $[a'] \leqslant G^L$ for some left option G^L of G. Since $b \leqslant a'$, we also have $[b] \leqslant G^L$. Since both $[b]$ and G^L are left options of G and therefore cannot be dominated, we have $G^L = [b]$. Since $[a'] \leqslant G^L$, we have $a' \leqslant b$. But we also assumed $b \leqslant a'$, so $a' = b$, so $G' = H$, so $G \simeq H$, so that H is a passing option of G, contradicting the assumption that G is in canonical form.

We showed that for every left option H of G, there is some left option H' of G' such that $H \leqslant H'$. The same argument shows that there is some left option H'' of G such that $H' \leqslant H''$. By transitivity, this implies $H \leqslant H''$. Since both H and H'' are left options of G and G is in canonical form, H cannot be dominated, so it follows that $H = H''$. Then $H \leqslant H' \leqslant H'' = H$, so that $H \simeq H'$. Since H and H' are in canonical form, by the induction hypothesis, $H = H'$.

We have shown that every left option of G is a left option of G'. A symmetric argument shows that every left option of G' is a left option of G, so that G and G' have exactly the same left options. The analogous property for right options holds by duality. □

Definition 4.22. A game is *short* if it has only finitely many positions; equivalently, a game is short if it has finitely many left and right options, and all the left and right options are short.

Lemma 4.23 (Existence of canonical form)**.** *Every short game has a canonical form (i.e., for every G, there exists G' such that $G \simeq G'$ and such that G' is in canonical form).*

Proof. We prove this by induction on the total size of the game, i.e., the total number of positions in G. Repeatedly remove a dominated option, bypass a reversible option, or simplify a passing option anywhere in G. Since each of these steps strictly decreases the size, the procedure terminates. □

4.5. Monotone Games

Monotone set coloring games, including Hex, are combinatorial games in the sense of Definition 4.1, but they have additional properties. In particular, in a monotone set coloring game, an additional move can never hurt a player: adding some stones of a player's color to a position is never bad for that player. This motivates the following definitions.

Definition 4.24 (Good option, monotone game)**.** Let G be a game over a poset A. A left option G^L of G is said to be *good* if $G \leqslant G^L$. Dually, a right option G^R of G is *good* if $G^R \leqslant G$. A game G is called *locally monotone* if all of its left and right options are good. It is said to be *monotone* if it is locally monotone and recursively, all of its options are monotone.

Note that in this definition, an option is "good" if it is beneficial for the player to whom the option belongs. This deviates from our usual convention of always taking Left's point of view.

Example 4.25. Over the booleans $\mathbb{B}$, the games $\top$, $\bot$, and $\{\top \mid \bot\}$ are monotone. The game $\{\bot \mid \top\}$ is not monotone.

Let $* = \{\top \mid \bot\}$. This game is called $*$ or "star" because it is a first-player win for both players, and thus is is roughly analogous to the nimber $* = \{0 \mid 0\}$ from normal play games. (However, although these two games share a name, they should only be considered analogous, not actually equal, because they belong to different classes of games, where wins are propagated in a different way.) The game $*$ has sometimes appeared under this name in the Hex literature; for example, it inspired the term *star decomposition* in [12].

Example 4.26. The game $G = \{* \mid *\}$ is not monotone. Indeed, since $\top \not\lhd \bot$, we have $\top \not\leqslant \{\top \mid \bot\} = *$, and therefore $\top \not\lhd \{* \mid *\} = G$, and hence $* = \{\top \mid \bot\} \not\leqslant G$. Since $*$ is a right option of G, this shows that G is not monotone.

Remark 4.27. If $G = \{G^L \mid G^R\}$ is monotone, then G^L, G^R are monotone and $G^R \leqslant G^L$. This is obvious from the definition of monotonicity and by transitivity, since $G^R \leqslant G \leqslant G^L$. However, the converse is not true: if G^R and G^L are monotone, assuming $G^R \leqslant G^L$ is not sufficient to imply that $G = \{G^L \mid G^R\}$ is monotone. Example 4.26 is a counterexample with $G^L = G^R = *$.

Lemma 4.28. *If G is a monotone game, then $G \lhd G$.*

Proof. If G is atomic, this is trivial by definition. If G is composite, pick any left option G^L (there is at least one). Since G is monotone, we have $G \leqslant G^L$. Thus, by definition of $\lhd$, we have $G \lhd G$. □

4.6. The Problem with Monotonicity

The problem with the concept of monotonicity is that it is a property of games, not equivalence classes of games. As the following example shows, the canonical form of a monotone game is not in general monotone. The absence of canonical forms for monotone games makes it awkward to work with monotone games up to equivalence.

Example 4.29 (Canonical forms of monotone games need not be monotone)**.** Consider the atom poset $A = \{\bot, a, b, \top\}$ with top element $\top$, bottom element $\bot$, and a, b incomparable. Consider the following games:

$$J = \{b \mid \bot\}, \quad K = \{\top \mid a\}, \quad H = \{\top \mid J\}, \quad G = \{K, H \mid \bot\}, \quad G' = \{K, b \mid \bot\}.$$

One can check that G is monotone. Moreover, the left option H is reversible via J in G, so that $G \simeq G'$. In fact, G' is the canonical form of G. However, G' is not monotone, because $G' \not\leqslant b$.

Fortunately, this problem has a solution, which we develop in Section 6. We will define the notion of a *passable game*, which is closed under canonical forms, and we will prove a game is equivalent to a monotone game if and only if its canonical form is passable.

5. Left and Right Equivalence

In this section, we introduce some technical notions that will give us more insight into the structure of games. This will be useful in the characterization of monotone games in Section 6. The results of this section also have other applications, for example to the efficient enumeration of games up to equivalence, which we describe in greater detail in Section 10. Readers who are in a hurry can skip directly to Section 6 and come back here as necessary.

To better understand equivalence of games, we consider some coarser equivalence relations, which we call left and right equivalence. For example, we can say that games H and K are "left equivalent" if they are equivalent when used as left options, i.e., if $\{H, G^L \mid G^R\} \simeq \{K, G^L \mid G^R\}$ for all G^L, G^R. In a similar vein, it makes sense to consider left and right orderings; for example, we define $H \leqslant_l K$ if $\{H, G^L \mid G^R\} \leqslant \{K, G^L \mid G^R\}$ for all G^L, G^R.

5.1. The Left and Right Orders

In the following, we use the letters S and T to denote non-empty sets of games. The case where S and T are single games is a special case. For notational convenience, when S consists of a single game H, we write $S = H$ rather than $S = \{H\}$.

Definition 5.1. Let S and T be non-empty sets of games. The *left order* is defined as follows: We say that $S \leqslant_l T$ if for every (possibly empty) set of games L and every non-empty set of games R, we have $\{S, L \mid R\} \leqslant \{T, L \mid R\}$. The *right order* is defined dually: we say that $S \leqslant_r T$ if $\{L \mid R, S\} \leqslant \{L \mid R, T\}$ for all R and all non-empty L.

By Lemma 4.8, $H \leqslant K$ implies $H \leqslant_l K$. However, the converse is not true. For example, we have $\{\top \mid \{a \mid \bot\}\} \leqslant_l a$. This can be seen by noting that as a left option, $\{\top \mid \{a \mid \bot\}\}$ is always reversible, and bypassing it results in a. On the other hand, $\{\top \mid \{a \mid \bot\}\} \nleqslant a$.

The next two lemmas give useful characterizations of the left order. If S is a set of games and G is a game, we write $S \lhd G$ to mean that $H \lhd G$ holds for all $H \in S$. Similarly, we write $S \leqslant G$ to mean that $H \leqslant G$ holds for all $H \in S$.

Lemma 5.2. *Let S and T be non-empty sets of games over any poset A (it does not need to have a top or bottom element). Then the following are equivalent:*

(a) $S \leqslant_l T$.

(b) *For all G, $T \lhd G$ implies $S \lhd G$.*

In particular, $H \leqslant_l K$ if and only if every right gift horse for K is a right gift horse for H.

Proof. To show that (a) implies (b), assume $S \leqslant_l T$ and $T \lhd G$. We must show $S \lhd G$. Since each $K \in T$ is a left gift horse for G, we have $G \simeq \{T, G^L \mid G^R\}$. Since $S \leqslant_l T$, we have $\{S, G^L \mid G^R\} \leqslant \{T, G^L \mid G^R\}$, and therefore $\{S, G^L \mid G^R\} \leqslant G$ by transitivity. Then by definition of $\leqslant$, we have $H \lhd G$ for all $H \in S$, thus $S \lhd G$ as desired.

To show the opposite implication, assume (b) holds. We must show $S \leqslant_l T$. Consider any L, R. We must show $\{S, L \mid R\} \leqslant \{T, L \mid R\}$. We clearly have $T \lhd \{T, L \mid R\}$, so by assumption (b), we have $S \lhd \{T, L \mid R\}$. Then $\{S, L \mid R\} \leqslant \{T, L \mid R\}$ follows directly from the definition of $\leqslant$.

The final claim of the lemma is just the special case when $S = H$ and $T = K$ are singletons. □

Lemma 5.3. *Suppose the atom poset A has a bottom element, and let S, T be non-empty sets of games over A. Then the following are equivalent:*

(a) $S \leqslant_l T$,

(b) $\{S \mid \bot\} \leqslant \{T \mid \bot\}$,

(c) $S \lhd \{T \mid \bot\}$.

Proof. (a) $\Rightarrow$ (b) follows directly from the definition of $\leqslant_l$. (b) $\Rightarrow$ (c) follows directly from the definition of $\leqslant$. To prove (c) $\Rightarrow$ (a), assume $S \lhd \{T \mid \bot\}$. Consider any L, R. We must show $\{S, L \mid R\} \leqslant \{T, L \mid R\}$. But we have $\{T \mid \bot\} \leqslant \{T, L \mid R\}$, so the assumption $S \lhd \{T \mid \bot\}$ implies $S \lhd \{T, L \mid R\}$ by Lemma 4.6. Then $\{S, L \mid R\} \leqslant \{T, L \mid R\}$ follows directly from the definition of $\leqslant$. □

Note that the condition in the definition of the left order says that $\{S, L \mid R\} \leqslant \{T, L \mid R\}$ must hold for "all" L, R. One may wonder whether we would wind up with a different definition of left order if we restricted L and R, for example to sets of monotone games. Lemma 5.3 shows that this makes no difference, since it suffices to consider $L = \emptyset$ and $R = \bot$.

Lemma 5.4. *Let A be any poset (it does not need to have a top or bottom element). Let H, K be games over A such that $H \leqslant_l K$. Then half of the conditions in the definition of $H \leqslant K$ are satisfied:*

(a) All right options K^R satisfy $H \lhd K^R$, and

(b) if K is atomic, $H \lhd K$.

Proof. By assumption, $H \leqslant_l K$. To show (a), consider any right option K^R of K. Then $K \lhd K^R$, and therefore by Lemma 5.2, we have $H \lhd K^R$, as desired. To show (b), assume K is atomic. Then $K \lhd K$, and again by Lemma 5.2, we have $H \lhd K$, as desired. □

The converse implication of Lemma 5.4 is not true. For example, let $H = \{\top \mid a\}$ and $K = a$. Then conditions (a) and (b) are satisfied, but $H \not\leqslant_l K$.

Lemma 5.5. *We have $H \leqslant K$ if and only if both $H \leqslant_l K$ and $H \leqslant_r K$ hold.*

Proof. The left-to-right implication holds by Lemma 4.8. For the right-to-left implication, assume $H \leqslant_l K$ and $H \leqslant_r K$. By Lemma 5.4 applied to $H \leqslant_l K$, half of the conditions for $H \leqslant K$ hold. By the dual of Lemma 5.4 applied to $H \leqslant_r K$, the other half holds, therefore $H \leqslant K$. □

Corollary 5.6. *On the class of games over a given atom poset, the relations $\leqslant$, $\leqslant_l$, and $\leqslant_r$ are all uniquely determined by the relation $\lhd$ (i.e., without requiring knowledge of the internal structure of games).*

Proof. By Lemma 5.2, the relation $\leqslant_l$ is uniquely determined by $\lhd$; the relation $\leqslant_r$ is determined dually, and finally $\leqslant$ is determined by Lemma 5.5. □

Definition 5.7. Let A be a poset with top and bottom elements. Let S be a non-empty set of games over A. We define $\uparrow S = \{\top \mid \{S \mid \bot\}\}$. Dually, define $\downarrow S = \{\{\top \mid S\} \mid \bot\}$.

Lemma 5.8. *We have $S \leqslant_l T$ if and only if $S \leqslant {\uparrow}T$.*

Proof. For the right-to-left implication, note that $S \leqslant \{\top \mid \{T \mid \bot\}\}$ implies $S \lhd \{T \mid \bot\}$ by the definition of $\leqslant$. Hence $S \leqslant_l T$ by Lemma 5.3. For the left-to-right implication, assume $S \leqslant_l T$, and consider some $H \in S$. To show $H \leqslant \{\top \mid \{T \mid \bot\}\}$, we must show three things. First, clearly any left option H^L satisfies $H^L \leqslant \top$ by Lemma 4.12. Second, we have $H \lhd \{T \mid \bot\}$ by Lemma 5.3. Third, if H is atomic, we must show that $H \lhd \{\top \mid \{T \mid \bot\}\}$. But this is true for all H (atomic or not), because $H \leqslant \top$. □

Corollary 5.9. *We have $S \leqslant_l T$ if and only if ${\uparrow}S \leqslant {\uparrow}T$.*

Proof. Left-to-right: $S \leqslant_l T$ implies, by definition, $\{S \mid \bot\} \leqslant \{T \mid \bot\}$, which implies ${\uparrow}S \leqslant {\uparrow}T$ by Lemma 4.8. Right-to-left: Assume ${\uparrow}S \leqslant {\uparrow}T$. From $S \leqslant_l S$, we get $S \leqslant {\uparrow}S$ by Lemma 5.8, hence $S \leqslant {\uparrow}T$ by transitivity, hence $S \leqslant_l T$ by Lemma 5.8. □

We remark that if the atom poset has top and bottom elements, both the left order $\leqslant_l$ and the right order $\leqslant_r$ admit least upper bounds (and by duality, greatest lower bounds). Let $H \vee K = \{H^L, K^L \mid \{H, K \mid \bot\}\}$ denote the least upper bound of H and K with respect to the order $\leqslant$, as in Proposition 4.13. Then using Lemma 5.8, we easily see that $H \vee K$ is also a least upper bound for $\leqslant_l$. Namely, for any G, we have $H, K \leqslant_l G$ if and only if $H, K \leqslant {\uparrow}G$ if and only if $H \vee K \leqslant {\uparrow}G$ if and only if $H \vee K \leqslant_l G$. With respect to the right order, the least upper bound of H and K is not $H \vee K$, but ${\downarrow}H \vee {\downarrow}K$. To see why, first note that ${\downarrow}({\downarrow}H \vee {\downarrow}K) \simeq {\downarrow}H \vee {\downarrow}K$. Then by the dual of Lemma 5.8, for any G, we have $H, K \leqslant_r G$ if and only if ${\downarrow}H, {\downarrow}K \leqslant G$ if and only if ${\downarrow}H \vee {\downarrow}K \leqslant G$ if and only if ${\downarrow}({\downarrow}H \vee {\downarrow}K) \leqslant G$ if and only if ${\downarrow}H \vee {\downarrow}K \leqslant_r G$. The same argument of course also applies to least upper bounds of more than two games.

5.2. Left and Right Equivalence

Definition 5.10. Let S and T be non-empty sets of games. We say that S, T are *left equivalent*, in symbols $S \simeq_l T$, if $S \leqslant_l T$ and $T \leqslant_l S$. In other words, S and T are left equivalent if $\{S, L \mid R\} \simeq \{T, L \mid R\}$ holds whenever L and R are a set of games and a non-empty set of games, respectively. Dually, we say that S, T are *right equivalent*, in symbols $S \simeq_r T$, if $S \leqslant_r T$ and $T \leqslant_r S$.

Lemma 5.11. *Let S be a set of games over an atom poset with top and bottom. Then S is left equivalent to ${\uparrow}S$.*

Proof. Applying the left-to-right direction of Lemma 5.8 to $S \leqslant_l S$, we get $S \leqslant {\uparrow}S$, which implies $S \leqslant_l {\uparrow}S$. Conversely, applying the right-to-left direction of Lemma 5.8 to ${\uparrow}S \leqslant {\uparrow}S$, we get ${\uparrow}S \leqslant_l S$. Thus, we have $S \simeq_l {\uparrow}S$, as claimed. □

Corollary 5.12. *The game* $\uparrow S$ *is the maximal element (with respect to the order* $\leqslant$*) in the left equivalence class of* S*.*

Proof. By Lemma 5.11, $\uparrow S$ is in the left equivalence class of S; if T is any other member of this left equivalence class, then $T \simeq_l S$, hence $T \leqslant_l S$, hence $T \leqslant \uparrow S$ by Lemma 5.8. ☐

Lemma 5.13. *We have* $H \simeq K$ *if and only if both* $H \simeq_l K$ *and* $H \simeq_r K$ *hold.*

Proof. This follows immediately from Lemma 5.5. ☐

The following lemma gives a remarkable equivalence between the collection of all left equivalence classes and the right equivalence class of $\top$.

Lemma 5.14. *Let* A *be a poset with top and bottom. Then for games over* A*, each left equivalence class has a unique maximal element (up to equivalence), and moreover, these maximal elements are precisely the members of the right equivalence class of* $\top$*.*

Proof. Each left equivalence class has a unique maximal element by Corollary 5.12. Since this maximal element is of the form $\{\top \mid G\}$, it is reversible when used as a right option, and therefore easily seen to be right equivalent to $\top$. Conversely, let G be any game that is right equivalent to $\top$. Then $G \simeq_l \uparrow G$ (by Lemma 5.11) and $G \simeq_r \uparrow G$ (because they are both right equivalent to $\top$). Hence by Lemma 5.13, $G \simeq \uparrow G$, i.e., G is the maximal element of its left equivalence class, as claimed. ☐

For a visualization of Lemmas 5.13 and 5.14, see Figures 8 and 9 in Section 10 below.

6. Passable Games and the Fundamental Theorem

6.1. Passable Games

Recall from Lemma 4.28 that every monotone game satisfies $G \lhd G$. The converse is clearly not true; for example, the game $G' = \{\{\top \mid a\}, b \mid \bot\}$ from Example 4.29 satisfies $G' \lhd G'$ (because G' is equivalent to a monotone game), but it is not itself monotone. It turns out that games satisfying $G \lhd G$ are an important class of games, which we will investigate in this section, culminating in the fundamental theorem of monotone games.

We start with an obvious observation. Recall from Definition 4.24 that a left option G^L is called *good* if $G \leqslant G^L$, and dually, a right option G^R is *good* if $G^R \leqslant G$.

Proposition 6.1. *Consider a game G over some atom poset A. The following are equivalent:*

(a) $G \lhd G$.

(b) *G is atomic or has at least one good (left or right) option.*

(c) *G is a left gift horse for itself.*

(d) *G is a right gift horse for itself.*

(e) *If Left had the option to pass in the position G, it would not benefit Left.*

(f) *If Right had the option to pass in the position G, it would not benefit Right.*

Proof. The equivalence of (a) and (b) follows directly from the definition of $\lhd$. Indeed, by definition, for $G \lhd G$ to be satisfied, one of three conditions must hold: either $G^R \leqslant G$ (in which case G has a good right option), or $G \leqslant G^L$ (in which case G has a good left option), or G is atomic. That is precisely what (b) is stating. The equivalence of (a), (c), and (d) is obvious, since this is the definition of a gift horse. To see why (c) and (e) are equivalent, assume that $G = \{G^L \mid G^R\}$ is composite (this is without loss of generality by Lemma 4.11). Consider the game $G' = \{G, G^L \mid G^R\}$. This game has the same moves as G, except that Left also has the option to pass (i.e., move to G). Clearly $G \leqslant G'$ by Lemma 4.9. To say that "passing does not benefit Left" means that $G \simeq G'$, which holds, by the Gift Horse Lemma 4.10, if and only if G is a left gift horse for itself. The equivalence between (d) and (f) is dual. $\square$

This motivates the following definition.

Definition 6.2. A game G is *locally passable* if $G \lhd G$. A game G is *passable* if it is locally passable and recursively, all of its options are passable.

The term "passable" should not be understood to mean that players are allowed to pass in these games. Indeed, none of our games literally allow passing, as that would potentially lead to infinite plays. Rather, it means that players *do not want to pass* in these games, and therefore, it does not matter whether passing is allowed or not. Alternatively, we can understand a passable game to be a game in which passing is "almost" allowed, in the sense that such a game $G = \{G^L \mid G^R\}$ is *equivalent* to a game $\{G, G^L \mid G^R, G\}$ that has G itself as a left option and a right option.

Lemma 6.3 (Canonical form of passable game)**.** *The canonical form of a passable game is passable.*

Proof. First, note that the class of locally passable games is closed under equivalence; indeed, if $G \lhd G$ and $G' \simeq G$, then $G' \lhd G'$; this follows from Lemma 4.6. Now assume G is some passable game and G' is its canonical form. Because of the way canonical forms can be computed by repeatedly removing dominated options, simplifying passing options, and bypassing reversible options, every position occurring in G' (including G' itself) is equivalent to some position occurring in G. Therefore, every position in G' is locally passable, and it follows that G' is passable. □

We also note that the class of passable games is closed under least upper bounds (and dually, under greatest lower bounds).

Lemma 6.4. *If the atom poset has a bottom element, the least upper bound of any set of passable games is equivalent to a passable game.*

Proof. Note that the concept of a least upper bound is only well-defined up to equivalence, hence the lemma states "equivalent to a passable game". We will show that the particular least upper bound that was constructed in the proof of Proposition 4.13 is passable.

For simplicity, we show this in the case of two games, but the same proof applies to arbitrary non-empty sets of games. Also, the least upper bound of the empty set of games is of course $\perp$ and is passable. So consider passable games H, K. As in Proposition 4.13, we may assume without loss of generality that H, K are composite. Then the least upper bound is $G = \{H^L, K^L \mid \{H, K \mid \perp\}\}$. The options of G are easily seen to be passable, so it suffices to show that G is locally passable. From the upper bound property, we have $H, K \leqslant G$. Since H and K are passable, we have $H \lhd H$ and $K \lhd K$. By transitivity, we get $H, K \lhd G$. This implies $\{H, K \mid \perp\} \leqslant G$. Since $\{H, K \mid \perp\}$ is a right option of G, this implies $G \lhd G$ as desired. □

6.2. The Fundamental Theorem of Monotone Games

By Lemma 4.28, we know that every monotone game is passable. Perhaps surprisingly, if A has top and bottom elements, the converse is true up to equivalence of games. The purpose of the rest of this section is to prove the following theorem.

Theorem 6.5 (Fundamental theorem of monotone games)**.** *Assume A has top and bottom elements. Then every passable game over A is equivalent to a monotone game over A.*

Together with Lemma 6.3, we immediately get the following characterization of the canonical form of monotone games:

Corollary 6.6. *Assume the atom poset has top and bottom elements and G has a canonical form. Then G is equivalent to a monotone game if and only if the canonical form of G is passable.*

The corollary gives an easy method for enumerating equivalence classes of short monotone games: we can simply enumerate all short passable games in canonical form.

6.3. Proof of the Fundamental Theorem

Recall from Definition 4.24 and Proposition 6.1 that a composite game is:

- *locally monotone* if all of its left and right options are good, and
- *locally passable* if it has at least one good left or right option.

To prove the fundamental theorem, we introduce an intermediate notion: we say that a composite game is

- *locally semi-monotone* if it has at least one good left option and at least one good right option.

By convention, atomic games are locally semi-monotone as well. Naturally, we then say that a game is *semi-monotone* if it locally semi-monotone and recursively, all of its options are semi-monotone. Our strategy for proving the fundamental theorem is: we will first show that every passable game is equivalent to a semi-monotone game, and then that every semi-monotone game is equivalent to a monotone game. Throughout this section, we assume that all games are over a fixed poset A of atoms with top and bottom elements.

Lemma 6.7. *Let S be a set of monotone (respectively semi-monotone, passable) games and assume that S forms a $\lhd$-clique in the sense that $H \lhd K$ for all $H, K \in S$. Then $\{\top \mid S\}$, $\{S \mid \bot\}$, $\uparrow S$, and $\downarrow S$ are monotone (respectively semi-monotone, passable). In particular, if H is a monotone game, then so are $\{\top \mid H\}$, $\{H \mid \bot\}$, $\uparrow H$, and $\downarrow H$.*

Proof. Let $G = \{\top \mid S\}$. Since all options of G are monotone (resp. semi-monotone, passable) by assumption, it suffices to show that G is locally monotone. Clearly, we have $G \leqslant \top$ by Lemma 4.12. So all that is left to show for monotonicity is that $H \leqslant G$ for all $H \in S$. So consider some $H \in S$. To show $H \leqslant G$, first consider any left option H^L. We must show $H^L \lhd G$, but this holds because $H^L \leqslant G^L = \top$. Second, consider any right option G^R. We must show $H \lhd G^R$. But both H, G^R are members of S, so the claim follows from the assumption that S is a $\lhd$-clique. Finally, in case H is atomic, we must show $H \lhd G$. But this is once again obvious since $\top$ is a left option of G. Therefore we have proved that $\{\top \mid S\}$ is monotone (resp. semi-monotone, passable).

The remaining claims of the lemma are easy: the claim for $\{S \mid \bot\}$ holds by duality, and the claims for $\uparrow S = \{\top \mid \{S \mid \bot\}\}$ and $\downarrow S$ follow by applying the previous claims twice. The claims in the lemma's final sentence follow because if H is a single monotone game, it is automatically a $\lhd$-clique. $\square$

Lemma 6.8. *Every passable game is equivalent to a semi-monotone game.*

Proof. We prove this by induction. Suppose G is a passable game. If G is atomic, then G is semi-monotone by definition, and there is nothing to show. If G is composite, then each of its left and right options is equivalent to a semi-monotone game by the induction hypothesis; therefore, G is equivalent to some composite game G' all of whose options are semi-monotone. Since G' is passable, G' either has a good left option or a good right option. We assume without loss of generality that G' has a good left option H (the case where G' has a good right option is dual).

Because G' may not already have a good right option, we will add one. We are therefore looking for a game K that is (1) a right gift horse for G', so that adding it as a right option to G' will not change the value of G', (2) semi-monotone, and (3) good, i.e., $K \leqslant G'$. If K satisfies all three conditions, then $G'' = \{G'^L \mid G'^R, K\}$ is semi-monotone and equivalent to G', hence to G. So the only thing left to do is to find a game K satisfying the above three conditions.

We claim that $K = \{H \mid \bot\}$ is the desired game. First, since H is a good left option of G', we have $G' \leqslant H$ and therefore $G' \lhd \{H \mid \bot\} = K$, so that K is a right gift horse for G'. Second, K is semi-monotone by Lemma 6.7. Third, we have $K \leqslant G'$ since the only left option of K is also a left option of G', and $\bot$ is below any right option of G'. So indeed, K satisfies properties (1)–(3), which finishes the proof. $\square$

Lemma 6.9. *Every semi-monotone game is equivalent to a monotone game.*

Proof. We prove this by induction. Let G be a semi-monotone game. If G is atomic, then G is monotone and there is nothing to show. If G is composite, then each of its left and right options is equivalent to a monotone game by the induction hypothesis; therefore, G is equivalent to some composite semi-monotone game G' all of whose options are monotone. Let L and R be the set of left and right options of G', respectively. Define $L' = \{\uparrow H \mid H \in L\}$ and $R' = \{\downarrow K \mid K \in R\}$. Let $G'' = \{\uparrow L' \mid \downarrow R'\}$. We claim that G'' is the desired game, i.e., we claim that $G'' \simeq G$ and G'' is monotone.

To show $G'' \simeq G$, recall that by Lemma 5.11 and its dual, each $\uparrow H$ is left equivalent to H and each $\downarrow K$ is right equivalent to K. We therefore have $G' \simeq$

$\{L' \mid R'\}$. Using Lemma 5.11 a second time, we also know that L' is left equivalent to $\uparrow L'$ and R' is right equivalent to $\downarrow R'$. It follows that $G \simeq G' \simeq \{\uparrow L' \mid \downarrow R'\} = G''$.

To show that G'' is monotone, first note that each $\uparrow H$ and each $\downarrow K$ is monotone by Lemma 6.7. Moreover, since $\uparrow H \lhd \uparrow H'$ for all H, H', the set L', and dually R', forms a $\lhd$-clique. Using Lemma 6.7 a second time, it follows that $\uparrow L'$ and $\downarrow R'$ are monotone. So all the options of G'' are monotone. The only thing left to prove is that G'' is locally monotone, i.e., $\downarrow R' \leqslant G'' \leqslant \uparrow L'$.

To show $G'' \leqslant \uparrow L'$, first consider any left option G''^L. We must show $G''^L \lhd \uparrow L'$, but this is plainly true since $G''^L = \uparrow L'$ (and $\uparrow L'$ is monotone, hence passable). Next, consider the only right option $\{L' \mid \bot\}$ of $\uparrow L'$. We must show $G'' \lhd \{L' \mid \bot\}$. Since G' is semi-monotone, G' has at least one good left option, say $H \in L$. By goodness, $G' \leqslant H$, which implies $G'' \leqslant \uparrow H$ since $G'' \simeq G'$ and $H \leqslant \uparrow H$. Since $\uparrow H \in L'$, it follows that $G'' \lhd \{L' \mid \bot\}$ as claimed. Since neither G'' nor $\uparrow L'$ is atomic, this finishes the proof of $G'' \leqslant \uparrow L'$.

The proof of $\downarrow R' \leqslant G''$ is dual. □

Proof of Theorem 6.5. Theorem 6.5 follows directly from Lemmas 6.8 and 6.9. □

Example 6.10. Consider the poset $A = \{\bot, a, b, \top\}$, with a and b incomparable. Let $G = \{a, b \mid a\}$. One can easily check that $a \leqslant G$, and hence $G \lhd G$, so G is passable. However, G is certainly not monotone; for example, we neither have $G \leqslant a$ nor $G \leqslant b$. By Theorem 6.5, G is equivalent to a monotone game. The particular game constructed in the proof of Theorem 6.5 is

$$\begin{aligned} G'' &= \{\uparrow(\uparrow a, \uparrow b, \uparrow\{\top \mid a\}) \mid \downarrow\downarrow a\} \\ &= \{\{\top \mid \{\{\top \mid \{a \mid \bot\}\}, \{\top \mid \{b \mid \bot\}\}, \{\top \mid \{\{\top|a\} \mid \bot\}\} \mid \bot\}\} \\ &\qquad \mid \{\{\top \mid \{\{\top \mid a\} \mid \bot\}\} \mid \bot\}\}. \end{aligned}$$

Indeed, one can check that G'' is monotone and equivalent to G. Note that G'' is not the simplest monotone game equivalent to G. We can obtain a simpler one by removing dominated options and bypassing only those reversible options for which bypassing does not break monotonicity. By doing so, we obtain the following simpler monotone game equivalent to G:

$$G''' = \{\{\top \mid \{\{\top \mid a\}, \{\top \mid \{b \mid \bot\}\} \mid \bot\}\} \mid a\}.$$

Although G''' is (perhaps) the simplest monotone game equivalent to G, it is not in canonical form, since it still has some reversible options. The canonical form of G''' is of course $G = \{a, b \mid a\}$.

Remark 6.11. In the proof of Lemma 6.9, we applied two levels of $\uparrow$ operations: first to each option of G', and then to the sets of all left options and of all right options. For example, if $G' = \{G_1, G_2 \mid G_3, G_4\}$, then $G'' = \{\uparrow(\uparrow G_1, \uparrow G_2) \mid \downarrow(\downarrow G_3, \downarrow G_4)\}$. Both steps are necessary; in general, neither $\{\uparrow G_1, \uparrow G_2 \mid \downarrow G_3, \downarrow G_4\}$

nor $\{\uparrow(G_1, G_2) \mid \downarrow(G_3, G_4)\}$ is monotone. Indeed, for the game from Example 6.10, one can check that neither $\{\uparrow a, \uparrow b, \uparrow\{\top \mid a\} \mid \downarrow a\}$ nor $\{\uparrow(a, b, \{\top \mid a\}) \mid \downarrow a\}$ is monotone, but $\{\uparrow(\uparrow a, \uparrow b, \uparrow\{\top \mid a\}) \mid \downarrow\downarrow a\}$ is.

6.4. A Useful Reasoning Principle for Passable Games

We will prove the following property of passable games, which is sometimes useful: if all left options satisfy $G^L \leqslant a$, then $G \leqslant a$. Informally, this is true because if it is Left's turn, Left can achieve at most outcome a. If it were possible for Left to do better when it is Right's turn, then Right would prefer to pass, contradicting the fact that Right does not prefer to pass in a passable game.

The formal proof requires a lemma.

Lemma 6.12. *Let G be a passable composite game, a an atom, and assume that all left options G^L satisfy $G^L \leqslant a$. Then the following hold for all H:*

(a) If $H \lhd G$ then $H \lhd a$.

(b) If $H \leqslant G$ then $H \leqslant a$.

Proof. We prove (a) and (b) by joint induction on H. In each inductive case, we prove (a) before (b). To prove (a), assume $H \lhd G$. By definition of $\lhd$, we either have $H \leqslant G^L$ for some G^L, or $H^R \leqslant G$ for some H^R. If $H \leqslant G^L$, we use the assumption $G^L \leqslant a$ to conclude $H \leqslant a$ by transitivity, and therefore also $H \lhd a$ by definition of $\leqslant$. If $H^R \leqslant G$, then by the induction hypothesis (b) we have $H^R \leqslant a$, therefore $H \lhd a$ as claimed.

To prove (b), assume $H \leqslant G$. We must show $H \leqslant a$. The definition of $\leqslant$ requires us to prove two things: all $H^L \lhd a$, and $H \lhd a$. To prove the first claim, consider any H^L. Since $H \leqslant G$, we have $H^L \lhd G$, and therefore $H^L \lhd a$ by the induction hypothesis (a). To prove the second claim, note that G is passable, so $G \lhd G$. From the assumption $H \leqslant G$, we get $H \lhd G$ by Lemma 4.6. Therefore, by part (a), which we already proved for H, we have $H \lhd a$ as claimed. □

Proposition 6.13. *Let G be a passable composite game, a an atom, and assume all left options G^L satisfy $G^L \leqslant a$. Then $G \leqslant a$.*

Proof. This follows by Lemma 6.12(b), with $H = G$. □

Note: not only does Lemma 6.12 imply Proposition 6.13, but also the other way round. However, since the lemma is proved by induction on H, and not by induction on G, we cannot prove the proposition directly without the lemma.

7. Games over Linearly Ordered Sets of Atoms

As we have seen in Section 4.6, the canonical form of a monotone game is not in general monotone; in fact, this was the reason we introduced the more general class of passable games. However, perhaps surprisingly, when the set A of atoms happens to be linearly ordered, the class of monotone games does turn out to be closed under canonical forms. The purpose of this section is to prove it. We start with two lemmas that state another useful property of games over linearly ordered posets.

Lemma 7.1. *Suppose A is a linearly ordered poset, $a \in A$ is an atom, and H is a passable game over A. Then we have:*

(a) $[a] \lhd H$ *or* $H \leqslant [a]$.

(b) $H \lhd [a]$ *or* $[a] \leqslant H$.

Proof. We prove this by induction. We first prove (a). If $H = [b]$ is atomic, we have $[a] \lhd H$ if and only if $a \leqslant b$ and $H \leqslant [a]$ if and only if $b \leqslant a$, so the claim follows from the fact that A is linearly ordered.

If H is not atomic, assume $[a] \ntriangleleft H$. To show $H \leqslant [a]$, first consider any left option H^L of H. Since by assumption, $[a] \ntriangleleft H$, we have $[a] \nleqslant H^L$, so by the induction hypothesis, $H^L \lhd [a]$. Therefore, the first part of the definition of $H \leqslant [a]$ is satisfied. The second part is trivially satisfied, since $[a]$ is atomic. For the third part, we must show $H \lhd [a]$. Since H is passable, H either has a good left option or a good right option. If H has a good left option H^L, then $H \leqslant H^L$. Since we already showed $H^L \lhd [a]$ above, it follows by Lemma 4.6 that $H \lhd [a]$. On the other hand, if H has a good right option H^R, then $H^R \leqslant H$. Since by assumption, $[a] \ntriangleleft H$, by Lemma 4.6, we have $[a] \ntriangleleft H^R$, therefore by the induction hypothesis, $H^R \leqslant [a]$, which implies $H \lhd [a]$ as desired. It follows that $H \leqslant [a]$.

The proof of (b) is dual. □

Lemma 7.2. *Suppose A is a linearly ordered poset. Then for all passable games G, H over A, we have $G \lhd H$ or $H \leqslant G$.*

Proof. We prove this by induction. If G or H is atomic, then the result holds by Lemma 7.1. Therefore, assume both G and H are composite. Suppose that $G \ntriangleleft H$. To show $H \leqslant G$, first consider an arbitrary left option H^L of H. We must show $H^L \lhd G$. From the definition of $G \ntriangleleft H$, we know that $G \nleqslant H^L$, therefore by the induction hypothesis, $H^L \lhd G$ as desired. The dual argument shows that $H \lhd G^R$ holds for all G^R. Since G and H are composite, this proves that $H \leqslant G$, as desired. □

Remark 7.3. If the games are not passable and A has at least 3 elements, Lemmas 7.1 and 7.2 are not true. The simplest counterexample is $A = \{\bot, a, \top\}$, $G = [a]$, $H = \{\bot \mid \top\}$. We neither have $[a] \lhd \{\bot \mid \top\}$ nor $\{\bot \mid \top\} \leqslant [a]$.

If $A = \{\bot, \top\}$, then Lemmas 7.1 and 7.2 are valid even if the games are not passable. Because in that case, Lemma 7.1 only has the two cases $a = \bot$ and $a = \top$, and both happen to be true because $\bot \lhd H \leqslant \top$ and $\bot \leqslant H \lhd \top$. Note that the proof of Lemma 7.2 goes through in that case as well.

We now come to the main result of this section.

Theorem 7.4. *Suppose A is a linearly ordered poset, that G is a passable game over A, and that G is in canonical form. Then G is monotone.*

Proof. We prove this by induction. Since all options of G are passable and in canonical form, they are monotone by the induction hypothesis. Therefore, all we have to show is that G is locally monotone, i.e., all options of G are good. We note that, since G is in canonical form, none of its options are dominated or reversible. We prove a number of claims in turn.

(a) For any two left options H, H' of G, we have $H \lhd H'$.

Proof: If $H = H'$, then the claim holds because H is passable. If $H \neq H'$, then since H' is not dominated, we have $H' \not\leqslant H$. Then $H \lhd H'$ holds by Lemma 7.2.

(b) For any left option H of G and any right option H^R of H, we have $G \lhd H^R$.

Proof: Since H is not reversible, we must have $H^R \not\leqslant G$. Then $G \lhd H^R$ holds by Lemma 7.2.

(c) All composite left options of G are good.

Proof: Let H be a left option of G and assume that H is composite. We must show $G \leqslant H$. From (a), we know that $G^L \lhd H$ for all G^L, and from (b), we know that $G \lhd H^R$ for all H^R. Since both G and H are composite, the third condition in the definition of $\leqslant$ does not apply, and we have $G \leqslant H$ as claimed.

(d) G has at most one atomic left option.

Proof: If a, b are distinct atomic left options of G, then we must either have $a \leqslant b$ or $b \leqslant a$ since the poset A is linearly ordered. But then either a or b would be dominated, contradicting the fact that G is in canonical form.

(e) All atomic left options of G are good.

Proof: Let a be an atomic left option of G. We must show $G \leqslant a$. We distinguish two cases: either G has some composite left option or it does not.

In case G has some composite left option H, then by (c), we have $G \leqslant H$. Since a is not dominated, we have $a \nleqslant H$, and therefore $a \nleqslant G$. Then $G \lhd a$ holds by Lemma 7.2. The other condition in the definition of $G \leqslant a$ already holds by (a). Therefore $G \leqslant a$ as claimed. In case G does not have any composite left options, then all left options are atomic. By (d), a is the only left option of G. Then $G \leqslant a$ holds by Proposition 6.13.

Together, (c) and (e) imply that all left options of G are good. The proof for right options is dual. This finishes the proof of the monotonicity of G. □

Corollary 7.5. *For games over a linearly ordered poset A, the canonical form of a monotone game is monotone.*

8. Operations on Games

8.1. Sum

As already mentioned in Section 4.2, in the standard treatment of combinatorial game theory, it is common to define the negation $-G$ and sum $G + H$ of games before defining the order, because $G \leqslant H$ can then be conveniently defined to mean that $H - G$ is a second-player win for Left. We did not take this route here, because in the games we consider in this paper, the sum of games works a bit differently than usual. Indeed, as we will see, the sum is well-behaved only for passable games, so it was necessary to define passable games, and therefore the ordering on games, before we could consider sums.

Unsurprisingly, the sum of games is defined in the same way as in other branches of combinatorial game theory when the games are composite. As usual, the difference lies in the treatment of atoms. When G and H are atomic games with respective outcomes a and b, we define their sum to be an atomic game whose outcome is the pair (a, b). This appropriately reflects what goes on when one plays in multiple disconnected regions in a monotone set coloring game such as Hex. This leads us to the following definition.

Definition 8.1 (The sum of games)**.** Suppose G and H are games over outcome posets A and B, respectively. The *sum* of G and H, written $G + H$, is a game over $A \times B$. It is defined as follows.

- $G + H = \{G^L + H, G + H^L \mid G^R + H, G + H^R\}$, when at least one of G, H is composite, and
- $[a] + [b] = [(a, b)]$, when $G = [a]$ and $H = [b]$ are both atomic.

We note that in this definition, we have used our usual convention that atomic games have no left and right options. Thus, in the cases where one of G, H is atomic and the other is composite, the definition specializes to:

- $[a] + H = \{[a] + H^L \mid [a] + H^R\}$,
- $G + [b] = \{G^L + [b] \mid G^R + [b]\}$.

Next, one would expect that we prove that the sum is a monotone operation, i.e., that $G \leqslant G'$ and $H \leqslant H'$ imply $G + H \leqslant G' + H'$. However, we will not prove this, because it is not true. In a nutshell, we built into the definition of the order that $a \simeq \{a \mid a\}$, so that a player can always, up to equivalence, pass in any atomic component of a larger game. This is justified when all games are passable, because in this case, passing is not to any player's advantage. But it is not justified for non-passable games, where passing may actually be advantageous. Since we are only interested in passable games, defining the order in this way is the right thing to do. But the price to pay is that monotonicity of sum does not hold for arbitrary games.

Example 8.2 (Non-monotonicity of sum)**.** Let $G = [a]$ and $H = \{\bot \mid \top\}$. Also consider the game $G' = \{a \mid a\}$. Then

$$\begin{aligned} G + H &= \{(a, \bot) \mid (a, \top)\}, \\ G' + H &= \{\{(a, \bot) \mid (a, \bot)\}, \{(a, \bot) \mid (a, \top)\} \mid \{(a, \top) \mid (a, \top)\}, \{(a, \bot) \mid (a, \top)\}\}. \end{aligned}$$

If we consider the outcome $(a, \top)$ to be winning for Left and $(a, \bot)$ to be winning for Right, the game $G + H$ is a second-player win for both players, whereas $G' + H$ is a first-player win for both players. The example shows that although $G \leqslant G'$ and $G' \leqslant G$, we have $G + H \not\leqslant G' + H$ and $G' + H \not\leqslant G + H$. In particular, the sum operation is not monotone on these games.

Since the game H is not passable, Example 8.2 is nothing to worry about. In Corollary 8.5, we will show that the sum operation is well-behaved on passable games. The following lemma holds for all games (passable or not). The proof is straightforward and we omit it.

Lemma 8.3. *Let $a, b \in A$ be atoms, and let G, H by any games over B.*

(a) If $a \leqslant b$ and $G \lhd H$, then $[a] + G \lhd [b] + H$.

(b) If $a \leqslant b$ and $G \leqslant H$ then $[a] + G \leqslant [b] + H$.

The symmetric properties, about $G + [a]$ and $H + [b]$, also hold.

Proposition 8.4 (Monotonicity of sum on passable games)**.** *Let G, G' be games over a poset A, and let H be a passable game over a poset B. Then:*

(a) $G \lhd G'$ *implies* $G+H \lhd G'+H$.

(b) $G \leqslant G'$ *implies* $G+H \leqslant G'+H$.

Proof. We prove (a) and (b) by joint induction. In each inductive case, we prove (a) before (b). Assume the proposition is true for all smaller triples of games (G, G', H).

To prove (a), assume $G \lhd G'$. We must show $G+H \lhd G'+H$. By the definition of $G \lhd G'$, there are three cases. Case 1: $G^R \leqslant G'$ for some right option G^R of G. By the induction hypothesis, $G^R+H \leqslant G'+H$. Since G^R+H is a right option of $G+H$, it follows that $G+H \lhd G'+H$, as desired. Case 2: $G \leqslant G'^L$ for some left option G'^L of G'. This case is analogous. Case 3: $G=[a]$ and $G'=[a']$ are atomic and $a \leqslant a'$. Here we use the assumption that H is passable (and this is the only place in the proof where it is used, apart from inductively using the fact that the options of H are passable). Since H is passable, we have $H \lhd H$, and therefore by Lemma 8.3, we have $[a]+H \lhd [a']+H$, i.e., $G+H \lhd G'+H$, as claimed.

To prove (b), assume $G \leqslant G'$. To show $G+H \leqslant G'+H$, first consider any left option K of $G+H$. We must show $K \lhd G'+H$. There are two cases, depending on what kind of left option K is. Case 1: $K = G^L+H$ for some left option G^L of G. From $G \leqslant G'$, we get $G^L \lhd G'$, and therefore by the induction hypothesis, $G^L+H \lhd G'+H$ as desired. Case 2: $K = G+H^L$ for some left option H^L of H. By the induction hypothesis, since H^L is passable, $G+H^L \leqslant G'+H^L$. Since $G'+H^L$ is a left option of $G'+H$, by definition of $\lhd$, we have $G+H^L \lhd G'+H$, as desired. This proves the first property required for $G+H \leqslant G'+H$. Second, we must show that $G+H \lhd J$ for every right option J of $G'+H$. This case is analogous to the previous one. The remaining thing to show is that if $G+H$ or $G'+H$ is atomic, then $G+H \lhd G'+H$. But by definition of $+$, it follows that G or G' is atomic. Since we assumed $G \leqslant G'$, it follows that $G \lhd G'$. Then by (a), we have $G+H \lhd G'+H$, as desired. □

Of course, the symmetric version of Proposition 8.4 also holds, i.e., if G is passable, then $H \lhd H'$ implies $G+H \lhd G+H'$ and $H \leqslant H'$ implies $G+H \leqslant G+H'$. Using transitivity, we thus obtain the monotonicity properties of the sum of passable games that are summarized in the following corollary.

Corollary 8.5. *Let* G, G' *be passable games over* A*, and let* H, H' *be passable games over* B*. Then*

(a) $G \lhd G'$ *and* $H \leqslant H'$ *imply* $G+H \lhd G'+H'$.

(b) $G \leqslant G'$ *and* $H \lhd H'$ *imply* $G+H \lhd G'+H'$.

(c) $G \leqslant G'$ *and* $H \leqslant H'$ *imply* $G+H \leqslant G'+H'$.

In particular, on passable games, the sum is well-defined up to equivalence: if $G \simeq G'$ *and* $H \simeq H'$*, then* $G+H \simeq G'+H'$.

Note that Corollary 8.5 also implies that the sum of passable games is passable. The sum also enjoys other expected properties: up to obvious isomorphisms of atom sets, the sum is symmetric and associative, and it has a unit, which is the unique game over one atom. These properties are straightforward and we omit the proofs.

From here on for the rest of the paper, all games will be assumed to be passable. We will usually state this explicitly (for example for the benefit of readers who skipped the current paragraph), but should it ever not be stated, the games are assumed to be passable anyway.

8.2. The Opposite Game

In standard combinatorial game theory, the negation of a game G, usually written $-G$, is the game obtained by exchanging the roles of the players. We have a similar operation here, but in addition to exchanging the roles of the players, we must also invert the outcome poset. To emphasize that this operation is not an additive inverse for the sum operation, we call it the opposite game, rather than the negation, and we denote it by G^{op}.

If A is a poset, let A^{op} denote the opposite poset, which has the same elements as A, but with the opposite order, i.e., $a \leqslant_{A^{\mathrm{op}}} b$ if and only if $b \leqslant_A a$.

Definition 8.6 (Opposite game)**.** Let G be a game over a poset A. The *opposite game* of G, denoted by G^{op}, is a game over A^{op}. It is obtained by exchanging the roles of Left and Right. More formally, we define $[a]^{\mathrm{op}} = [a]$ and $\{G^L \mid G^R\}^{\mathrm{op}} = \{(G^R)^{\mathrm{op}} \mid (G^L)^{\mathrm{op}}\}$.

Note that G is monotone (respectively, passable) if and only if G^{op} is monotone (respectively, passable). All notions dualize; for example, we have $G \leqslant H$ if and only if $H^{\mathrm{op}} \leqslant G^{\mathrm{op}}$, $G \simeq_l H$ if and only if $G^{\mathrm{op}} \simeq_r H^{\mathrm{op}}$, and so on.

8.3. The Map Operation

We can use a monotone function $f : A \to B$ to turn a game over A into a game over B.

Definition 8.7 (Map operation)**.** Let A, B be posets and $f : A \to B$ a monotone function. Given a game G over A, we define a game $f(G)$ over B as follows:

- $f([a]) = [f(a)]$ for atomic games;
- $f(G) = \{f(G^L) \mid f(G^R)\}$ when G is composite.

We say that the function f *maps* the game G to the game $f(G)$. We will often apply maps to sums of games, in which case we write $G +_f H$ instead of $f(G + H)$.

The map operation satisfies some obvious properties; for example, it is functorial: if $f : A \to B$ and $g : B \to C$, we have $g(f(G)) = (g \circ f)(G)$, and if $\mathrm{id} : A \to A$ is the identity function, we have $\mathrm{id}(G) = G$. Also, if $f : A \times A \to A$ is an associative operation, then $(G +_f H) +_f K = G +_f (H +_f K)$. Finally, as expected, the map operation is monotone in both arguments, i.e., $f \leqslant f'$ and $G \leqslant G'$ imply $f(G) \leqslant f'(G')$, and similarly for $\lhd$. Also, if G is passable, then so is $f(G)$.

Remark 8.8. Any monotone set coloring game (X, π) is actually $\pi(G)$, where G is a direct sum of $|X|$ copies of the game $\{\top \mid \bot\}$, and $\pi : \mathbb{B}^X \to A$ is the payoff function.

8.4. Copy-Cat Strategies

One difference between the sum operation in our games and in other kinds of combinatorial games is that there are actually many different sum operations, depending on what winning condition we want to impose on the combined game. In a sense, each different monotone function $f : A \times B \to C$ defines a different sum operation $G +_f H$. Many of these are not self-dual. As a simple example, consider the boolean functions "and" and "or" : $\mathbb{B} \times \mathbb{B} \to \mathbb{B}$. If G, H are games over $\mathbb{B}$, then to win the game $G +_{\text{and}} H$, Left must win both components, but Right only needs to win one component. The game $G +_{\text{or}} H$ is its dual: here Left only needs to win one component, but Right must win both components. In Hex, such games might look like this:

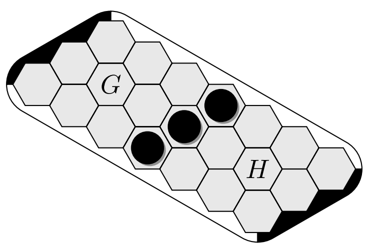

Black must win both components.

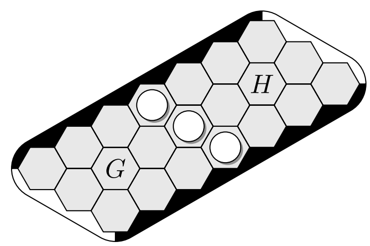

Black must win one component.

Games like these were called *conjunction* and *disjunction* games in [20].

We can generalize this idea to an arbitrary outcome poset. Given a poset A, define the monotone functions $\lambda, \rho : A \times A^{\mathrm{op}} \to \mathbb{B}$ as follows:

$$\lambda(a, b) = \begin{cases} \top & \text{if } b \leqslant_A a, \\ \bot & \text{otherwise.} \end{cases} \qquad \rho(a, b) = \begin{cases} \top & \text{if } a \not\leqslant_A b, \\ \bot & \text{otherwise.} \end{cases}$$

These functions are each other's duals, in the sense that $(G +_\lambda H)^{\mathrm{op}} = H^{\mathrm{op}} +_\rho G^{\mathrm{op}}$. They are called "<u>l</u>ambda" and "<u>r</u>ho" because they favor the <u>L</u>eft and <u>R</u>ight player, respectively. We have the following properties:

Lemma 8.9. *For all passable games G, H:*

(a) $G +_\lambda G^{op} \simeq \top$ is a second-player win for Left.

(b) $G +_\rho G^{op} \simeq \bot$ *is a second-player win for Right.*

(c) $G \leqslant H \iff \top \leqslant H +_\lambda G^{op} \iff G +_\rho H^{op} \leqslant \bot$.

(d) $G \lhd H \iff \top \lhd H +_\lambda G^{op} \iff G +_\rho H^{op} \lhd \bot$.

We omit the proofs, which are straightforward inductions on G and H, and basically the same as the usual proofs of $G \leqslant H$ if and only if $0 \leqslant H + (-G)$ and related properties in normal play games. In particular, the second-player winning strategies employed by Left in (a) and Right in (b) are the copy-cat strategies, which consist of always copying the other player's last move in the opposite component of the game.

9. Contextual Order and Global Decisiveness

9.1. Contexual Order

As motivated in Section 2, our games over an outcome poset A are intended to represent "local" play, such as play in a particular region of a Hex board. This local play is a component of a larger "global" game. In the global game, success or failure is measured by winning or losing, i.e., global games are always over the outcome poset $\mathbb{B} = \{\bot, \top\}$. Now that we have introduced sums and the map operation, we can briefly describe how the local and global notions hang together. This is not unlike what happens in other branches of combinatorial game theory, but appropriately adjusted to accommodate games over outcome posets.

Definition 9.1 (Context)**.** Let A be a poset. A *context* for A is a triple (B, f, K), where B is a poset, $f : A \times B \to \mathbb{B}$ is a monotone function, and K is a passable game over B.

Definition 9.2 (Contextual order and equivalence)**.** Fix a poset A. We define several different *contextual order* relations on games over A, as follows. For passable games G, H over A, we say:

- $G \leqslant_1^c H$ if for all contexts (B, f, K), if Left has a first-player winning strategy in $G +_f K$, then Left has a first-player winning strategy in $H +_f K$.
- $G \leqslant_2^c H$ if for all contexts (B, f, K), if Left has a second-player winning strategy in $G +_f K$, then Left has a second-player winning strategy in $H +_f K$.
- $G \lhd^c H$ if for all contexts (B, f, K), if Left has a second-player winning strategy in $G +_f K$, then Left has a first-player winning strategy in $H +_f K$.

We also write $G \leqslant^c H$ if $G \leqslant_1^c H$ and $G \leqslant_2^c H$, and $G \simeq^c H$ if $G \leqslant^c H$ and $H \leqslant^c G$. The latter relation is called *contextual equivalence.*

We note that up to equivalence, there are only three passable games over $\mathbb{B}$: $\top$, which is a first- and second-player win for Left; $*$, which is a first-player win but not a second-player win for Left; and $\bot$, which is neither a first- nor a second-player win for Left. Thus, Left has a first-player winning strategy in some game X over $\mathbb{B}$ if and only if $* \leqslant X$, or equivalently, $\top \lhd X$, and Left has a second-player winning strategy if and only if $\top \leqslant X$.

The following proposition shows that the contextual order relations coincide with the relations we defined in Section 4.2.

Proposition 9.3. *For passable games G, H over A, we have:*

(a) $G \leqslant_1^c H \iff G \leqslant H$.

(b) $G \leqslant_2^c H \iff G \leqslant H$.

(c) $G \lhd^c H \iff G \lhd H$.

Proof. In all three cases, the right-to-left implication follows from the monotonicity properties of the sum and map operations. For example, $G \leqslant H$ implies $G +_f K \leqslant H +_f K$ for all contexts. We must show the left-to-right implications. For (a), assume $G \leqslant_1^c H$. Consider the context $(A^{\mathrm{op}}, \rho, H^{\mathrm{op}})$. We know from Lemma 8.9(b) that Right has a second-player win in the game $H +_\rho H^{\mathrm{op}}$; therefore Left does not have a first-player win in that game. Since $G \leqslant_1^c H$, it follows that Left does not have first-player win in $G +_\rho H^{\mathrm{op}}$, so we must have $G +_\rho H^{\mathrm{op}} \simeq \bot$. Then by Lemma 8.9(c), we get $G \leqslant H$ as claimed. The proof of (b) is dual. Here we assume $G \leqslant_2^c H$ and use the context $(A^{\mathrm{op}}, \lambda, G^{\mathrm{op}})$. By Lemma 8.9(a), Left has a second-player winning strategy in $G +_\lambda G^{\mathrm{op}}$, and since $G \leqslant_2^c H$, Left also has a second-player winning strategy in $H +_\lambda G^{\mathrm{op}}$, so $H +_\lambda G^{\mathrm{op}} \simeq \top$, so $G \leqslant H$ by Lemma 8.9(c). The proof of (c) is also very similar. Here, we assume $G \lhd^c H$ and use the context $(A^{\mathrm{op}}, \lambda, G^{\mathrm{op}})$. By Lemma 8.9(a), Left has a second-player winning strategy in $G +_\lambda G^{\mathrm{op}}$. Therefore, since $G \lhd^c H$, Left has a first-player winning strategy in $H +_\lambda G^{\mathrm{op}}$. This means that $\top \lhd H +_\lambda G^{\mathrm{op}}$. Thus, by Lemma 8.9(d), $G \lhd H$. ☐

Proposition 9.3 is perhaps not very surprising; similar results hold in other branches of combinatorial game theory, by and large with the same proofs. The fact that our a priori notions of order and equivalence coincide with the contextual order and equivalence does provide an additional measure of evidence that our definitions are good. (Note that the order and equivalence of games over $\mathbb{B}$ is likely uncontroversial, since there are only three passable game values $\{\bot, *, \top\}$, and they are completely determined by the existence of winning strategies for the players.)

But the real reason we stated Definition 9.2 is to be able to generalize it to a setting where the analog of Proposition 9.3 is false. We do so in Section 9.2.

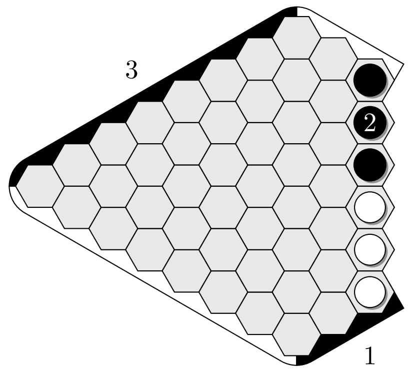

Figure 7: A 3-terminal region with globally decisive $\top$ and $\perp$, or "one-sided fork".

9.2. Games with Globally Decisive Moves

Consider the Hex region in Figure 7. It is a 3-terminal region in which two of Black's terminals are opposing board edges, and two of White's terminals are opposing board edges. We call such a region a *one-sided fork*, because it is a fork (Figure 6) with the additional property that connecting terminal 3 to terminal 1 is strictly better than connecting it to terminal 2, and in fact as good as connecting all three terminals.

Because four of the one-sided fork's terminals are edges of the global Hex board, it is clear that if Black connects her two edges within the region, then she immediately wins the global game, rather than just winning in the local region. In other words, there is no move whatsoever that White could play outside the region that would trump such a connection by Black. In this case, we say that the outcome $\top$ of the local region is "globally decisive". Similarly, if White connects his two edges within the region, he immediately wins the global game and trumps anything that Black could do outside the region. Therefore, the outcome $\perp$ is also globally decisive. It turns out that the theory of regions with globally decisive $\top$ and $\perp$ is slightly different than that with ordinary $\top$ and $\perp$; it sometimes (but not always) happens that some games that would otherwise be inequivalent can become equivalent due to global decisiveness.

In this section, we give a precise definition of global decisiveness and some examples of how globally decisive equivalence does not coincide with ordinary equivalence. Much of the theory of global decisiveness has not yet been worked out and is left for future work.

Definition 9.4. Let A and B be posets, and assume A has top and bottom elements. A monotone function $f : A \times B \to \mathbb{B}$ is called *(left) strict* if $f(\top, b) = \top$ and $f(\perp, b) = \perp$ for all $b \in B$. A context (B, f, K) for A is called *strict* if f is left strict.

Lemma 9.5. *Let (B, f, K) be a strict context over A. Then $\top +_f K \simeq \top$ and $\bot +_f K \simeq \bot$.*

Proof. By the strictness of f, every atom occurring in $\top +_f K$ is $\top$, and therefore $\top +_f K \simeq \top$ by repeated application of Lemma 4.11. The second claim is dual. $\square$

We say that a region of a game has *globally decisive* $\top$ and $\bot$ if it can only be played in strict contexts. An example of this is the one-sided fork in Figure 7. It then makes sense to define an ordering and equivalence on games by taking only strict contexts into account. The following definition does this.

Definition 9.6 (Globally decisive order and equivalence)**.** Let A be a poset with top and bottom elements. The *globally decisive order* relations $\leqslant_1^g$, $\leqslant_2^g$, and $\lhd^g$ are relations on games over A that are defined in exactly the same way as the relations $\leqslant_1^c$, $\leqslant_2^c$, and $\lhd^c$ in Definition 9.2, except that all contexts are restricted to *strict* contexts. Specifically:

- $G \leqslant_1^g H$ if for all strict contexts (B, f, K), if Left has a first-player winning strategy in $G +_f K$, then Left has a first-player winning strategy in $H +_f K$.
- $G \leqslant_2^g H$ if for all strict contexts (B, f, K), if Left has a second-player winning strategy in $G +_f K$, then Left has a second-player winning strategy in $H +_f K$.
- $G \lhd^g H$ if for all strict contexts (B, f, K), if Left has a second-player winning strategy in $G +_f K$, then Left has a first-player winning strategy in $H +_f K$.

We also write $G \leqslant^g H$ if $G \leqslant_1^g H$ and $G \leqslant_2^g H$, and $G \simeq^g H$ if $G \leqslant^g H$ and $H \leqslant^g G$. The latter relation is called *globally decisive equivalence*.

We note that these relations are at least as coarse as $\leqslant$ and $\lhd$; in other words, $G \leqslant H$ implies $G \leqslant_1^g H$ and $G \leqslant_2^g H$, and $G \lhd H$ implies $G \lhd^g H$. This is trivial, because if some condition holds for *all* contexts, then it certainly holds for all *strict* contexts. But in contrast to the situation in Section 9.1, the globally decisive order does not usually coincide with $\leqslant$. In fact, the following example shows that $\leqslant_1^g$ and $\leqslant_2^g$ do not even coincide with each other.

Example 9.7. Let $G = \top$ and $H = * = \{\top \mid \bot\}$. Consider any strict context (B, f, K). By Lemma 9.5, $G +_f K \simeq \top$, so that Left has first- and second-player winning strategies in $G +_f K$, regardless of K (effectively, Left has already won $G +_f K$). Also, the game $H +_f K$ has $\top +_f K \simeq \top$ as a left option, and it has $\bot +_f K \simeq \bot$ as a right option. Therefore, Left has a first-player winning strategy, but no second-player winning strategy, in $H +_f K$, again regardless of K. It follows that $G \leqslant_1^g H$ but $G \not\leqslant_2^g H$. Also, clearly $G \not\leqslant H$, which shows that $\leqslant_1^g$ and $\leqslant$ do not coincide. A dual argument shows that $\leqslant_2^g$ implies neither $\leqslant$ nor $\leqslant_1^g$.

The previous example shows that $\leqslant_1^g$ and $\leqslant_2^g$ do not coincide with $\leqslant$, but one may wonder if the relation $\leqslant^g$, defined as the intersection of $\leqslant_1^g$ and $\leqslant_2^g$, coincides with $\leqslant$. The following example shows that this is not the case. In other words, even $G \leqslant_1^g H$ and $G \leqslant_2^g H$ together do not imply $G \leqslant H$.

Example 9.8. Consider the linearly ordered poset $A = \{\bot, a, b, \top\}$, with $\bot < a < b < \top$. Let $G = \{\{\top \mid \{b \mid \{a \mid \bot\}\}\} \mid \bot\}$ and $H = \{b, \{\top \mid a\} \mid \bot\}$. One can check that $G \not\leqslant H$. However, we have $G \leqslant_1^g H$ and $G \leqslant_2^g H$. To see $G \leqslant_2^g H$, note that Left never has a second-player win in G in any strict context, because Right can always move to $\bot$. So $G \leqslant_2^g H$ is vacuously true.

Proving $G \leqslant_1^g H$ is more tricky and more fun, because it requires actually playing the game. Let $G' = \{b \mid \{a \mid \bot\}\}$. We first claim that $G' \leqslant_1^g H$. To that end, let (B, f, K) be some strict context and assume that Left has a first-player winning strategy in $G' +_f K$. We must show that Left has a first-player winning strategy in $H +_f K$. Left's first-player winning strategy in $G' +_f K$ must start with a winning move. There are two cases, depending on what this move is. Case 1: Left's winning move is $b +_f K$. But since $b +_f K$ is also a left option of $H +_f K$, Left then has a first-player win in $H +_f K$, which is what we had to prove. Case 2: Left's winning move is $\{b \mid \{a \mid \bot\}\} +_f K^L$ for some left option K^L of K. Since $\{b \mid \{a \mid \bot\}\} +_f K^L$ is a second-player win for Left, Left must have a first-player win in all of its right options, including in $\{a \mid \bot\} +_f K^L$. Left's winning move in $\{a \mid \bot\} +_f K^L$ can certainly not be in K^L, because then Right can play $\bot$ and win by strictness. Therefore, Left's winning move must be $a +_f K^L$, which means that Left has a second-player win in $a +_f K^L$. Now we will show that Left has a first-player win in $H +_f K$. We claim that the left option $\{\top \mid a\} +_f K$ is Left's winning move, i.e., that the game $\{\top \mid a\} +_f K$ is a second-player win for Left. So consider any right move. If Right moves in K, Left wins immediately by playing $\top$ and by strictness. If Right plays in the other component, the move is $a +_f K$, from which Left has the move $a +_f K^L$, which we already showed is winning for Left. Therefore, Left indeed has a first-player win in $H +_f K$. This concludes the proof that $G' \leqslant_1^g H$. Finally, it is easy to check that $G \leqslant G'$, which implies $G \leqslant_1^g G' \leqslant_1^g H$. The relation $\leqslant_1^g$ is transitive almost by definition, so $G \leqslant_1^g H$ holds.

Remark 9.9. The games G and H in Example 9.8 actually satisfy $G \simeq^g H$, but not $G \simeq H$. Indeed, it is easy to check that $H \leqslant G$, hence $H \leqslant^g G$. Since we proved $G \leqslant^g H$, we have $G \simeq^g H$. Thus, the example shows that globally decisive equivalence does not coincide with ordinary equivalence of games. We will see a Hex example of this in Section 11.1.

Remark 9.10. In Example 9.8, the proof that $G \leqslant^g H$ does not imply $G \leqslant H$ uses the global decisiveness of both $\top$ and $\bot$. Indeed, it is necessary to use them both, because we can show: if $\leqslant^g$ were defined with respect to contexts that are strict only for $\top$, or only for $\bot$, then $\leqslant^g$ would coincide with $\leqslant$. To see why, it suffices to note that the function λ defined in Section 8.4 is left strict with respect to $\top$ (but not $\bot$), and the function ρ is left strict with respect to $\bot$ (but not $\top$). Now if $G +_f K \leqslant H +_f K$ holds for all f that are strict for $\top$, then $G +_\lambda G^{\mathrm{op}} \leqslant H +_\lambda G^{\mathrm{op}}$ holds, which implies $G \leqslant H$ as in Proposition 9.3. Dually, if $G +_f K \leqslant H +_f K$ holds for all f that are strict for $\bot$, then $G +_\rho H^{\mathrm{op}} \leqslant H +_\rho H^{\mathrm{op}}$ holds, which also implies $G \leqslant H$.

Having given a definition of global decisiveness and some examples, it would be nice to know more of its properties. For example, it is not clear from the definition whether the relation $\leqslant^g$ is decidable, since it potentially requires looking at all possible contexts, of which there can be infinitely many. Even better would be to have a recursive definition for $\leqslant^g$, along similar lines as Definition 4.3, and a canonical form for games up to globally decisive equivalence. However, this is left for future work.

10. Enumeration of Game Values

10.1. An Efficient Algorithm for Enumerating Game Values

Given a finite poset A, there is an easy, but inefficient, method for enumerating all canonical forms of passable games of finite depth over A. First we enumerate all games of depth 0; these are the atomic games. Now given the set $\mathcal{G}_n$ of all canonical passable games of depth up to n, which is finite, it is easy to construct the set of all passable games of depth up to $n + 1$; each such game is either atomic, or it has a set of left options and a set of right options that are subsets of $\mathcal{G}_n$. Since there are only finitely many such subsets of $\mathcal{G}_n$, there are only finitely many potential games at depth $n + 1$ to consider. We can disregard all games that are not passable or are not in canonical form. What is left is the set $\mathcal{G}_{n+1}$ of all canonical passable games of depth up to $n + 1$.

The enumeration method described in the previous paragraph is extremely inefficient. Suppose there are 100 canonical passable games at depth n. Then we must consider $2^{100} - 1$ possible sets of left options and $2^{100} - 1$ possible sets of right options, given an astronomical set of games to consider (most of which turn out not to be in canonical form). The enumeration can be made far more efficient using the concepts of left and right equivalence from Section 5. The crucial insight is that we do not need to consider *all* games $\{L \mid R\}$, where L and R are sets of games of smaller depth. Instead, it suffices to consider just one representative L of each *left equivalence class* of sets of games, and one representative R of each right equivalence class. When L ranges over such representatives, all of the games $\{L \mid \bot\}$ are passable and distinct. Therefore, the number of left equivalence classes at each depth is not larger than the number of games to be enumerated at the next depth. The dual statement holds for right equivalence classes as well. This means that the number of games $\{L \mid R\}$ that must be considered at depth $n+1$ is at most the square of the number of distinct games that will ultimately be output at depth $n+1$. The enumeration is therefore reasonably efficient (relative to the amount of output produced), provided that we can efficiently enumerate left and right equivalence classes of games.

The enumeration of left (or dually, right) equivalence classes can be done efficiently as follows. Given a set $X = \{G_1, \dots, G_k\}$ of games, we first recursively compute representatives $L_1, \dots, L_m$ for all the left equivalence classes of non-empty subsets of $\{G_1, \dots, G_{k-1}\}$. Then we consider the $2m+1$ sets $L_1, \dots, L_m$ and $\{G_k\}, L_1 \cup \{G_k\}, \dots, L_m \cup \{G_k\}$. By eliminating duplicates, we obtain the set of all representatives of left equivalence classes of subsets of X. If the set X has k members, we need to repeat this step k times, and in each step, the number of representatives potentially doubles. But since we eliminate duplicates after each step, rather than only at the end, m never exceeds the final number of representatives computed. While this number may (or may not) be exponential as a function of k, the runtime is polynomial as a function of the amount of output produced.

10.2. Game Values over Specific Atom Sets

We compute the set of passable game values over various atom sets. Let us write L_n for the linearly ordered set with n elements.

- Over the set $L_1 = \{\bot\}$, there is only one passable game value, and it is $\bot$.
- Over the set $L_2 = \{\bot, \top\}$, there are exactly 3 distinct passable game values, and they are $\bot$, $* = \{\top \mid \bot\}$, and $\top$.
- Over the set $L_3 = \{\bot, a, \top\}$, there are exactly 8 distinct passable game values. They are shown, along with a Hasse diagram for the partial order $\leqslant$, in Figure 8.
- Over the set $L_4 = \{\bot, a, b, \top\}$, with $\bot < a < b < \top$, there are exactly 31 distinct passable game values, and they are shown, along with a Hasse diagram for the partial order $\leqslant$, in Figure 9.

Note that by Theorem 7.4, since all of the games enumerated above are over linear atom sets and in canonical form, these games are not just passable but also monotone.

Remark 10.1. Figures 8 and 9 also show the left and right equivalence classes of games. These figures confirm that the intersection of each left equivalence class with each right equivalence class is at most a singleton, as proved in Lemma 5.13. They also illustrate that each left equivalence class has a unique maximal element,

Depth 0:
$G_0 = \bot$
$G_1 = a$
$G_2 = \top$

Depth 1:
$G_3 = \{a \mid \bot\}$
$G_4 = \{\top \mid \bot\} = *$
$G_5 = \{\top \mid a\}$

Depth 2:
$G_6 = \{G_5 \mid \bot\}$
$G_7 = \{\top \mid G_3\}$

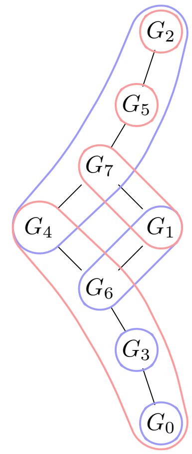

Figure 8: The passable game values over $L_3 = \{\bot, a, \top\}$ and their partial order. Left equivalence classes are indicated in red and right equivalence classes in blue.

and that these elements are exactly the members of the right equivalence class of $\top$. This was shown in Lemma 5.14 for arbitrary games, but the same proof also applies to passable games. We further note the set of game values in these figures forms a lattice under $\leqslant$; this was shown in Proposition 4.13 and Lemma 6.4.

From the above, one may wonder whether the collection of passable game values over a finite atom poset is always finite. This is not the case. In fact, the following proposition, due to Eric Demer, shows that the collection of game values is always infinite when the poset A is not linearly ordered.

Proposition 10.2 (Demer)**.** *Let A be a poset with two incomparable elements a and b. Define a sequence of games by $G_0 = a$ and $G_n = \{a, b \mid G_{n-1}\}$ for all $n \geqslant 1$. Then $G_0, G_1, G_2, \ldots$ is an infinite, strictly increasing sequence of passable games over A. In particular, there are infinitely many non-equivalent passable games over A.*

Proof. We first show, by induction on n, that $G_n \leqslant G_{n+1}$ for all $n \geqslant 0$. For $n = 0$, this is clear because $G_0 = a \simeq \{a \mid a\} \leqslant \{a, b \mid a\} = G_1$. For the induction step, assume $G_n \leqslant G_{n+1}$. Then $G_{n+1} = \{a, b \mid G_n\} \leqslant \{a, b \mid G_{n+1}\} = G_{n+2}$ by Lemma 4.8. Thus, $G_0, G_1, G_2, \ldots$ is an increasing sequence of games. Next, we show that all G_n are passable. It suffices to show that $G_n \lhd G_n$ for all $n \geqslant 0$. For $n = 0$, this is trivial since $G_0 = a$ is atomic. For $n \geqslant 1$, we have $G_n \lhd G_{n-1}$ since G_{n-1} is a right option of G_n. Since we have already proved $G_{n-1} \leqslant G_n$, this implies $G_n \lhd G_n$ by Lemma 4.6. Finally, we must show that the sequence $(G_n)_{n\in\mathbb{N}}$ is strictly increasing, i.e., we must show that $G_{n+1} \nleqslant G_n$, for all $n \geqslant 0$. We show this by

Depth 0:
$G_0 = \bot$
$G_1 = a$
$G_2 = b$
$G_3 = \top$

Depth 1:
$G_4 = \{a \mid \bot\}$
$G_5 = \{b \mid \bot\}$
$G_6 = \{\top \mid \bot\}$
$G_7 = \{b \mid a\}$
$G_8 = \{\top \mid a\}$
$G_9 = \{\top \mid b\}$

Depth 2:
$G_{10} = \{b, G_8 \mid \bot\}$
$G_{11} = \{G_7 \mid \bot\}$
$G_{12} = \{G_8 \mid \bot\}$
$G_{13} = \{G_9 \mid \bot\}$
$G_{14} = \{G_9 \mid a\}$
$G_{15} = \{\top \mid a, G_5\}$
$G_{16} = \{b \mid G_4\}$
$G_{17} = \{\top \mid G_4\}$
$G_{18} = \{G_9 \mid G_4\}$
$G_{19} = \{\top \mid G_5\}$
$G_{20} = \{\top \mid G_7\}$

Depth 3:
$G_{21} = \{G_{14} \mid \bot\}$
$G_{22} = \{G_{15} \mid \bot\}$
$G_{23} = \{G_{20} \mid \bot\}$
$G_{24} = \{G_{20} \mid G_4\}$
$G_{25} = \{\top \mid G_{16}\}$
$G_{26} = \{\top \mid G_{10}\}$
$G_{27} = \{\top \mid G_{11}\}$
$G_{28} = \{G_9 \mid G_{11}\}$

Depth 4:
$G_{29} = \{G_{25} \mid \bot\}$
$G_{30} = \{\top \mid G_{21}\}$

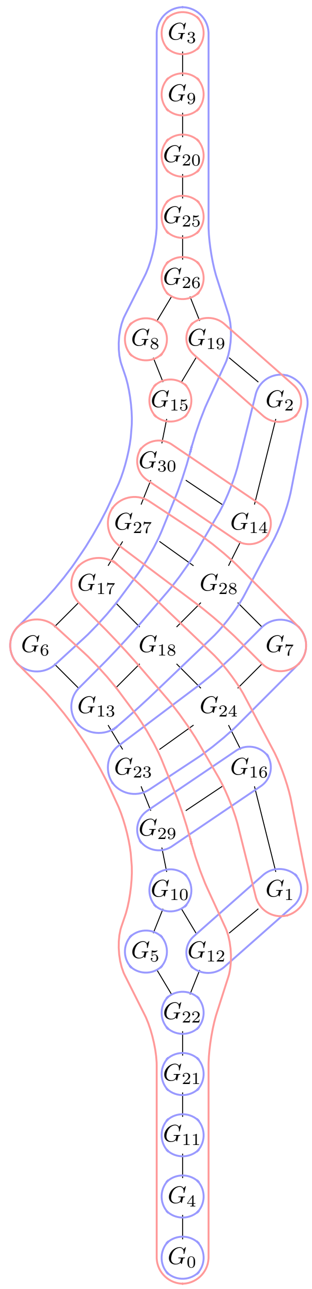

Figure 9: The passable game values over $L_4 = \{\bot, a, b, \top\}$ and their partial order. Left equivalence classes are indicated in red and right equivalence classes in blue.

induction. The base case $n = 0$ holds because $\{a, b \mid a\} \not\leqslant a$. For the induction step, assume $G_{n+1} \not\leqslant G_n$. We must show $G_{n+2} \not\leqslant G_{n+1}$, i.e., $\{a, b \mid G_{n+1}\} \not\leqslant \{a, b \mid G_n\}$. Assume, on the contrary, that $\{a, b \mid G_{n+1}\} \leqslant \{a, b \mid G_n\}$. Then by definition of $\leqslant$, we have $\{a, b \mid G_{n+1}\} \lhd G_n$. By definition of $\lhd$, there are two cases. Case 1: $\{a, b \mid G_{n+1}\} \leqslant G_n^L$ for some left option G_n^L of G_n. But this is plainly not the case, since the only left options of G_n are a and b, and neither is greater than $\{a, b \mid G_{n+1}\}$. Case 2: $G_{n+1} \leqslant G_n$, but this contradicts our induction hypothesis. Therefore $G_{n+2} \not\leqslant G_{n+1}$ as claimed. This completes the proof. $\square$

Next, one may wonder whether it is at least the case that the collection of passable game values is finite whenever the atom poset is *linearly* ordered and finite. This is also false: as shown in [6], there exist infinitely many non-equivalent passable games over L_5. It follows that the atom sets $L_1, \ldots, L_4$ (and trivially, the empty atom set) are the only ones over which the collection of passable game values is finite.

Over atom sets where the collection of passable game values is infinite, we can still enumerate passable games of small depth.

- Over the set L_5, there are exactly 5 canonical-form passable games of depth 0, 10 such games of depth 1, 40 games of depth 2, 178 games of depth 3, and 2962 games of depth 4, for a total of 3195 canonical-form passable games of depth up to 4.
- Over the poset $A = \{\bot, a, b, \top\}$, with a and b incomparable (see Figure 6(b)), there are exactly 4 canonical-form passable games of depth 0, 11 such games of depth 1, and 291 games of depth 2, for a total of 306 games of depth up to 2. We enumerated more than 43000 distinct passable games of depth 3 before running out of memory.
- Over the poset $A = \{\bot, a, b, c, \top\}$, with a, b, and c incomparable (see Figure 2(b)), there are exactly 5 canonical-form passable games of depth 0 and 33 such games of depth 1. We enumerated more than 1.8 million distinct passable games of depth 2 before running out of memory.

11. Realizability as Hex Positions

As we saw, the value of a Hex position is always a passable game. However, it is not a priori clear whether *every* passable game arises as the value of some Hex position — and indeed, we will show in Proposition 11.2 that this is not in general the case.

11.1. Realizable Game Values

We start by reporting some positive results on Hex-realizable game values over small atom posets.

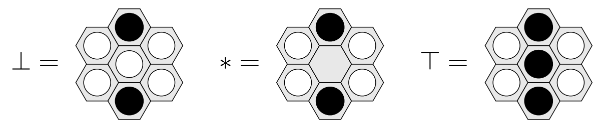

Figure 10: 2-terminal Hex positions realizing all passable game values over $L_2 = \mathbb{B}$.

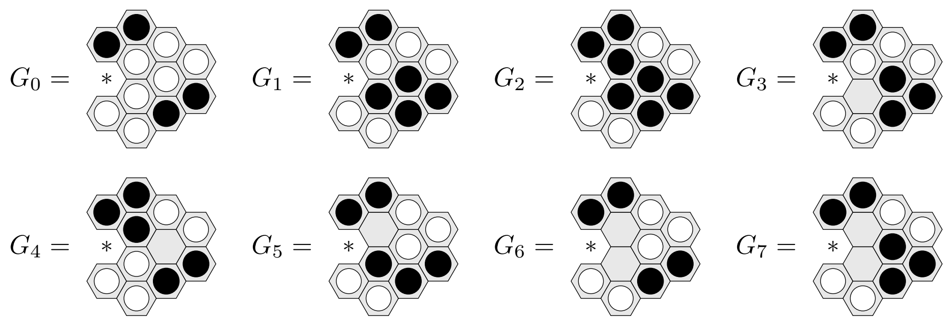

Figure 11: 2-terminal Hex positions with gap, realizing all passable game values over L_3.

- For 2-terminal regions in Hex, the outcome poset is $\mathbb{B} = \{\bot, \top\}$, and there are three passable game values: $\bot$, $*$, and $\top$. All of these passable game values are realizable as Hex positions, as shown in Figure 10.
- For 2-terminal Hex regions with gap (see Figure 5), the outcome poset is $L_3 = \{\bot, a, \top\}$. Recall from Figure 8 that there are 8 passable game values over this poset. All of them are realizable as Hex positions, as shown in Figure 11.
- For one-sided forks (see Figure 7), the outcome poset is also $L_3 = \{\bot, a, \top\}$. In this setting too, all 8 passable game values are realizable as Hex positions, as shown in Figure 12.
- Figure 13 shows a one-sided fork with an additional gap marked "$*$" between the two non-edge terminals. We consider the gap to be part of the region. The outcome poset for this type of region is the 4-element linearly ordered set $L_4 = \{\bot, a, b, \top\}$, with $\bot < a < b < \top$. The four outcomes are:
  - $\bot$: White's board edges are connected.
  - a: Neither player's board edges are connected, and White occupies the gap.
  - b: Neither player's board edges are connected, and Black occupies the gap.

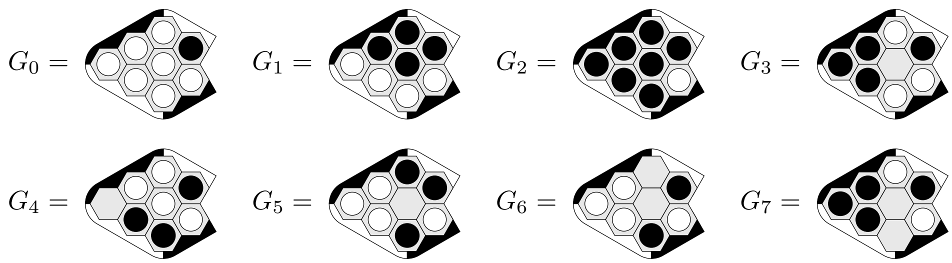

Figure 12: One-sided forks, realizing all passable game values over L_3.

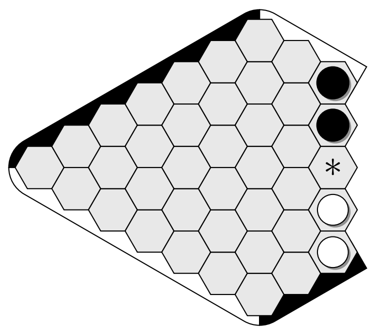

Figure 13: A 3-terminal region with four board edges and a gap, or "one-sided fork with gap".

- $\top$: Black's board edges are connected.

Recall from Figure 9 that there are 31 passable game values over L_4. All of them are realizable in Hex by one-sided forks with gap, as shown in Figure 14.

Remark 11.1. While the values of all 31 Hex positions shown in Figure 14 are distinct when viewed as abstract combinatorial games, Demer pointed out that several of them are actually equivalent *as Hex positions*, due to the global decisiveness of $\top$ and $\bot$ in these positions (see Section 9.2). Specifically, we have

$$G_{21} \simeq^g G_{22}, \quad G_{10} \simeq^g G_{29}, \quad G_{30} \simeq^g G_{15}, \quad G_{26} \simeq^g G_{25}.$$

For example, the values of G_{10} and G_{29} are $\{b, \{\top \mid a\} \mid \bot\}$ and $\{\{\top \mid \{b \mid \{a \mid \bot\}\}\} \mid \bot\}$, respectively. These were shown to be equivalent under global decisiveness in Example 9.8 and Remark 9.9. The games G_{15} and G_{30} are their duals. The remaining equivalences can be shown by a similar argument. Perhaps surprisingly, apart from the four equivalences shown above, there are no additional

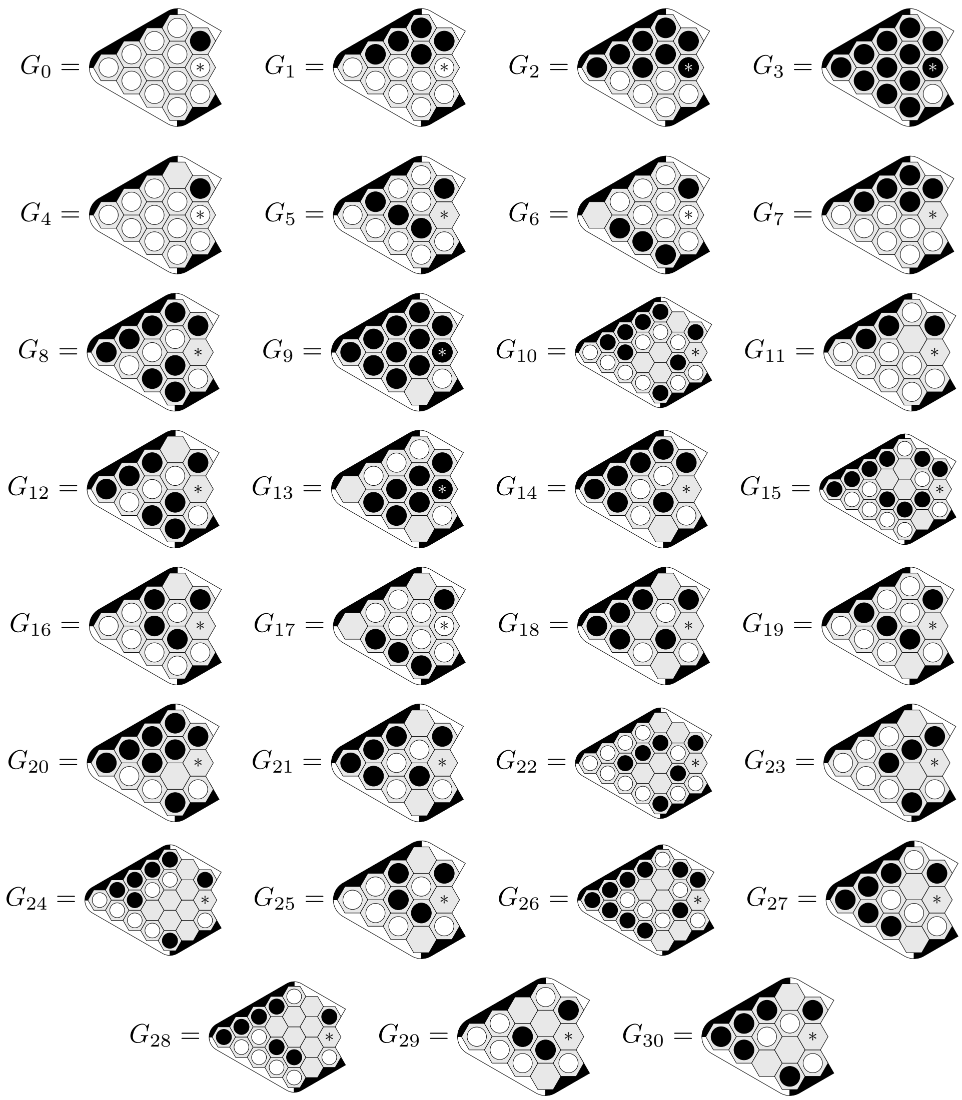

Figure 14: One-sided forks with gap, realizing all passable game values over L_4.

collapses of the order under global decisiveness. In other words, among the 31 games of Figure 9 or Figure 14, there are no pairs such that $G_i \leqslant^g G_j$ but $G_i \nleqslant G_j$, except for $G_{22} \leqslant^g G_{21}$, $G_{29} \leqslant^g G_{10}$, $G_{15} \leqslant^g G_{30}$, and $G_{25} \leqslant^g G_{26}$.

We also note that if two Hex positions realize combinatorial game values that are inequivalent, or even inequivalent up to global decisiveness, it does not necessarily follow that they are inequivalent as Hex positions. Since equivalence of Hex positions in a region is defined relative to all possible ways of embedding the region in a larger *Hex* position, it is possible that equivalence of Hex positions is a strictly coarser relation than equivalence of abstract game values. However, apart from the phenomenon of global decisiveness, which is not specific to Hex, no examples of this are currently known.

11.2. Unrealizable Game Values

So far, we have seen several settings in which all abstract passable games over a given outcome poset were realizable as Hex positions. The following proposition shows that this is not true in general.

Proposition 11.2. *Consider 4-terminal Hex positions. As usual, let $\top$ denote the outcome where Black connects all terminals, and let $\bot$ denote the outcome where Black connects no terminals. The abstract game value $* = \{\top \mid \bot\}$, which is passable, cannot be realized by any 4-terminal Hex position.*

The proof relies on the following lemma.

Lemma 11.3. *Consider a Hex region that is completely filled with black and white stones. Changing the color of a single stone from black to white can increase the number of Black's connected components by at most 2.*

Proof. Let x be the cell whose color is being changed. Before the color change, x belongs to one black connected component. Since x is a hexagon, it has at most 6 neighbors; therefore, after the color change, at most 3 black connected components can be adjacent to x. Since all black connected components that are not adjacent to x are unaffected by the color change, this shows that the total number of black connected components increases by at most 2. $\square$

Proof of Proposition 11.2. Suppose, for the sake of contradiction, that there is some 4-terminal position with value $\{\top \mid \bot\}$. Then Black has a first-player strategy that allows Black to connect all 4 terminals. Now suppose White goes first and plays in some cell x. Black can simply ignore White's move and follow Black's original strategy. (If the strategy ever calls for Black to play at x, Black can simply pass, or arbitrarily play elsewhere.) This will result in an outcome where all 4 terminals would be connected if x was occupied by a black stone. By Lemma 11.3, changing a single black stone to white can break this connected component into at most 3 parts, which implies that at least 2 of Black's terminals are still connected. Therefore, White cannot achieve outcome $\bot$, contradicting our initial assumption. □

Using essentially the same proof, we can easily derive generalizations of Proposition 11.2, such as the following.

Proposition 11.4. *Suppose Black has a strategy that allows Black to connect all n terminals in an n-terminal position. If White is given k free moves, then Black still has a strategy for connecting at least $n - 2k$ terminals.*

Results like Propositions 11.2 and 11.4 suggest that in Hex, there is some kind of limit on the amount of advantage a player can gain with a single move. In combinatorial game theory terms, while Hex is a "hot" game (players never prefer passing to making a move), it is not a "very hot" game (there is an inherent limit on how much a single move can achieve).

In Proposition 10.2, we saw that the class of passable abstract games over most outcome posets is infinite. However, in light of Proposition 11.2, we know that in general, not all abstract game values are realizable as Hex positions. This leaves open the question whether certain types of Hex regions (such as 3-terminal regions) admit finitely or infinitely many non-equivalent Hex-realizable combinatorial values. In fact, since the counterexample of Proposition 11.2 requires at least four terminals, it is not even currently known whether there exists an abstract passable 3-terminal value that is not realizable as a 3-terminal Hex position. One of the simplest passable 3-terminal values for which no Hex realization is currently known is $\{a, \{\top \mid b\} \mid \{a \mid \bot\}, b\}$.

Note that Lemma 11.3, and therefore Propositions 11.2 and 11.4, are specific to Hex, and do not apply to other monotone set coloring games, or even planar connection games. For example, on a game board that contains an octagonal cell, it is trivially possible to find a 4-terminal position with value $\{\top \mid \bot\}$, namely, a single empty octagon surrounded by alternating black and white stones.

This leaves open another question, namely, whether all passable game values can be realized by planar connection games. Failing that, we may ask whether they are at least realizable by vertex Shannon games, or failing that, by arbitrary monotone set coloring games. While the first two questions remain open, the third question was recently answered in the affirmative: all (short) passable game values are realizable as monotone set coloring games [7].

12. Application: Minimal Connecting Sets in $k \times n$ Hex

The theory developed in this paper has several potential applications to Hex, and some of these will be discussed in subsequent papers. Here, we will briefly discuss

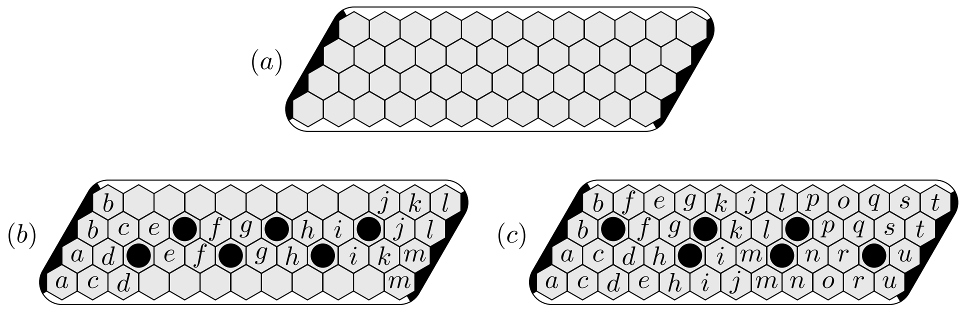

Figure 15: (a) A Hex board of size 4×12. (b) and (c): two virtual connections for Black.

one such application.

12.1. Problem Statement

Consider a Hex board of size $k \times n$, as shown for $k = 4$ and $n = 12$ in Figure 15(a). It is well-known that White has a first-player winning strategy when $n \geqslant k$ and a second-player winning strategy when $n > k$ [14, 8]. The question we would like to answer is: what is the minimal number of black stones that must be placed on the board to create a virtual connection between Black's edges? Here, by "virtual connection", we mean that Black will win the game even if White moves first, or equivalently, that the game has combinatorial value $\top$. Two examples of such virtual connections are shown in Figures 15(b) and (c). Both of these virtual connections use 6 stones. We will later prove that this is the minimal number required for a board of size 4×12.

To see that the patterns of stones in Figures 15(b) and (c) are indeed virtual connections, it suffices to consider a simple pairing strategy. Whenever White moves in a cell that is labelled with a letter, Black moves in the other cell with the same letter. It is not hard to see that this strategy guarantees a Black connection. We note that although both patterns use 6 stones, the pattern in Figure 15(b) requires less "space": the unlabeled cells are not required for Black's connection, and may as well be occupied by White. On the other hand, experienced Hex players will have no trouble seeing that if any of the labelled cells in Figures 15(b) and (c) are occupied by White, Black no longer has a virtual connection. We also note that the patterns shown in Figures 15(b) and (c) both generalize to larger board sizes; obvious continuations of these patterns work for boards of size 4×15, 4×18, 4×21, and so on.

Note that we are not claiming that *every* virtual connection admits a simple pairing strategy. We are merely saying that this is the case for the ones shown in

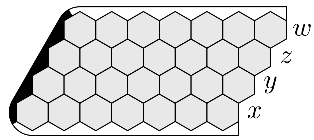

Figure 16: An open region of height 4.

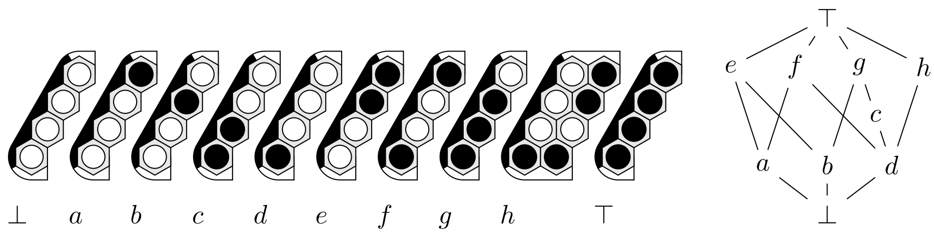

Figure 17: The atomic outcomes for open regions of height 4.

Figures 15(b) and (c).

So how would one go about proving minimality? More importantly, how would one do this for, say, all boards of size $k \times n$, for fixed k and arbitrary n? Combinatorial game theory is the perfect tool for this job. We will illustrate the method for $k = 4$, but it also works for other fixed values of k. (However, the computations get exponentially harder when k increases.)

12.2. Open Regions

We start by considering *open regions* of height 4. By this, we mean a region that includes the left, top, and bottom edges, but not the right edge, as shown in Figure 16. For an open region that is completely filled with stones, the outcome class is determined by the following information: whether White's edges are connected to each other (outcome $\perp$), and if they are not connected: which of the cells adjacent to the region (labelled x, y, z, w in Figure 16) are connected to the left edge by a black chain, and which such cells are connected to each other by a black chain. It turns out that for open regions of height 4, there are exactly 10 distinct outcomes, representatives of which are shown in Figure 17. For example, for outcome f, the cells x, z, and w are connected to the left edge (and therefore to each other), but y is not. For outcome h, only x is connected to the left edge, but y, z, and w are connected to each other. Figure 17 also shows the partial order on these outcomes.

Armed with this outcome poset, we can now calculate the value of positions in

$P =$ 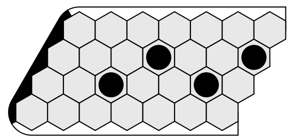

$$\mathrm{val}(P) = \{\{\top \mid g\}, \{\top \mid \{\{\top \mid h\}, \{\top \mid e\} \mid \{h \mid \bot\}, \{e \mid \bot\}\}\} \mid \{g \mid \{\{g \mid d\}, \{g \mid b\} \mid \{d \mid \bot\}, \{b \mid \bot\}\}\}, \{\{\{\top \mid h\}, \{\top \mid g\} \mid \{h \mid d\}, \{g \mid d\}\}, \{\{\top \mid h\}, \{\top \mid e\} \mid \{h \mid \bot\}, \{e \mid \bot\}\}, \{\{\top \mid g\}, \{\top \mid e\} \mid \{g \mid b\}, \{e \mid b\}\} \mid \{\{h \mid d\} \mid \bot\}, \{\{g \mid d\}, \{g \mid b\} \mid \{d \mid \bot\}, \{b \mid \bot\}\}, \{\{e \mid b\} \mid \bot\}\}\}$$

Figure 18: A position P in an open region and its value $\mathrm{val}(P)$.

open regions. For example, Figure 18 shows such a position P and its combinatorial value, which we write $\mathrm{val}(P)$. The point of such combinatorial values is not that they should be readable and understandable by humans, but that they can be calculated and compared.

12.3. Column-Wise Computation

It is important that values of open regions can be computed efficiently. Note that the most naive method is not efficient. The Hex position in Figure 18 has 24 empty cells. If we were to calculate its value by brute force, we would first obtain a game with 24 left and right options, each of which has 23 left and right options and so on, to depth 24. This game would be of astronomical size, containing roughly 10^{31} atomic positions. Although its canonical form (shown in Figure 18) is much smaller, it would not be feasible to compute it in this way.

The key to computing values of open positions efficiently is the following. Consider a *column*, which is a region of size 4×1 with top and bottom edges (but no left and right edges), as shown in Figure 19. A column has exactly 16 distinct atomic outcomes, namely, all of the ways of filling the column with black and white stones, as shown in Figure 20. Given an open position P of size $4 \times n$ and a column C, we write $P + C$ for the open position of size $4 \times (n + 1)$ obtained by adding C to the right of P. This naturally gives rise to a function on outcome classes, i.e., for $P \in \{\bot, a, \ldots, h, \top\}$ and $C \in \{k_0, \ldots, k_{15}\}$, we define $f(P, C)$ as the outcome class of $P + C$. For example, we have $f(c, k_{10}) = d$, because

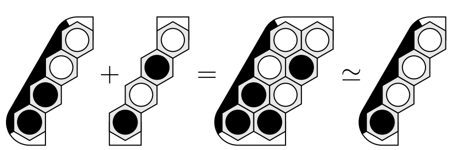

This function f on atomic games can then be extended to non-atomic games as in Section 8.3. Namely, if P and C are (not necessarily atomic) positions for an open

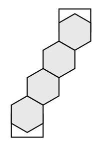

Figure 19: A column of height 4.

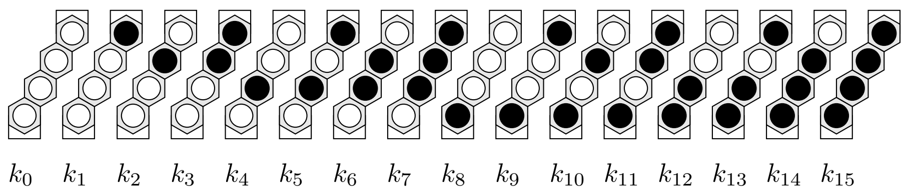

Figure 20: The atomic outcomes for columns of height 4.

region and a column, then $\mathrm{val}(P+C) = \mathrm{val}(P) +_f \mathrm{val}(C)$. From now on, by a slight abuse of notation, we just write $+$ instead of $+_f$, so that we can use the notation $P+C$ whether P and C are positions, outcomes, or values. We occasionally confuse positions with their values or vice versa.

We now have a reasonably efficient method for computing values of open positions. Namely, we can start from a position of size 4×0, then add one column, reduce to canonical form, add the next column, reduce to canonical form, and so on.

12.4. Best Patterns

Now that we have a good method for calculating the value of open $4 \times n$ positions, we need one more ingredient before we can prove results on the minimality of virtual connections on $4 \times n$-boards. We certainly cannot do this by calculating the values of all possible arrangements of black stones. The final key idea is that we do not need to calculate the values of all such arrangements. It suffices to calculate the values of the *best* arrangements. For brevity, we will use the term *pattern* for a position in an open region that uses only black stones.

Definition 12.1. A pattern P is *unacceptable* if $\mathrm{val}(P) \lhd \bot$. In this case, White has a winning move in P, so clearly P can never be part of a virtual connection for Black.

Definition 12.2. To each acceptable pattern P, we associate a triple $P^* = (G, s, n)$, where:

- $G = \mathrm{val}(P)$ is the combinatorial value of P,

- s is the number of black stones used, and
- n is the *width* of P, i.e., the number of columns in P.

For convenience, we also define $P^* = (\bot, \infty, 0)$ if P is an unacceptable pattern. For example, for the pattern P from Figure 18, we have $P^* = (G, 4, 7)$, where G is the value shown in Figure 18. We define an ordering on such triples by $(G, s, n) \leqslant (G', s', n')$ if $G \leqslant G'$ and $s' \leqslant s$ and $n \leqslant n'$. We also write $P \leqslant P'$ if $P^* \leqslant P'^*$, and we say that P' is *at least as good as* P. When $P \leqslant P'$ and $P' \leqslant P$, we say that P and P' are *equivalent*. Note that all unacceptable patterns are equivalent to each other, and are worse than all acceptable patterns.

Informally, the ordering on triples is justified as follows. If two patterns P and P' cover the same distance with the same number of black stones, but P' has a higher combinatorial value than P, then P' is clearly better, in the sense that everything that can be achieved with P can also be achieved with P'. In other words, if extending P with some additional columns results in a virtual connection, then so does adding the same columns to P'. Next, if P and P' have the same combinatorial value and cover the same distance, but P' uses fewer stones, then P' is the better of the two. Finally, if P and P' have the same combinatorial value and use the same number of stones, but P' covers a greater distance, then again P' is better.

We note that if $P = P' + C$, then P^* is completely determined by P'^* and C. In particular, if P and Q are equivalent, then so are $P + C$ and $Q + C$.

We can now systematically enumerate the *best* patterns of each width n, i.e., patterns that are maximal with respect to the preorder $\leqslant$. When several patterns are equivalent, we only keep one of them, selected arbitrarily. The enumeration proceeds as follows: For $n = 0$, there is only one pattern, and it is automatically best. Now suppose we have the list $P_1, \ldots, P_m$ of the best patterns of width n. We generate a list of patterns of width $n + 1$ by adding each of the 16 possible column patterns (using blank cells and black stones) to each of $P_1, \ldots, P_m$. Of these, we keep only the best ones, throwing out the rest.

Perhaps surprisingly, the number of best patterns does not grow exponentially as the width n increases; it remains bounded. In fact, in an appropriate sense, the sequence eventually repeats. Figure 21 shows all best patterns (up to equivalence) of width 4, 5, 6, and 7. Each pattern P is shown with its triple P^*. The game values $G_1, \ldots, G_{28}$ have not been written out as they can be quite long; notice, however, that several values appear more than once. As is evident from Figure 21, for each best pattern with triple $(G, s, 4)$, there is a best pattern with triple $(G, s + 2, 7)$, and vice versa. As remarked above, the best triples of width $n + 1$ are completely determined by the best triples of width n. It follows that the entire sequence of best triples repeats in the following sense: for all $n \geqslant 4$, (G, s, n) is a best triple if and only if $(G, s + 2, n + 3)$ is a best triple. In other words, whenever the width increases by 3, exactly 2 additional black stones are required to achieve the same

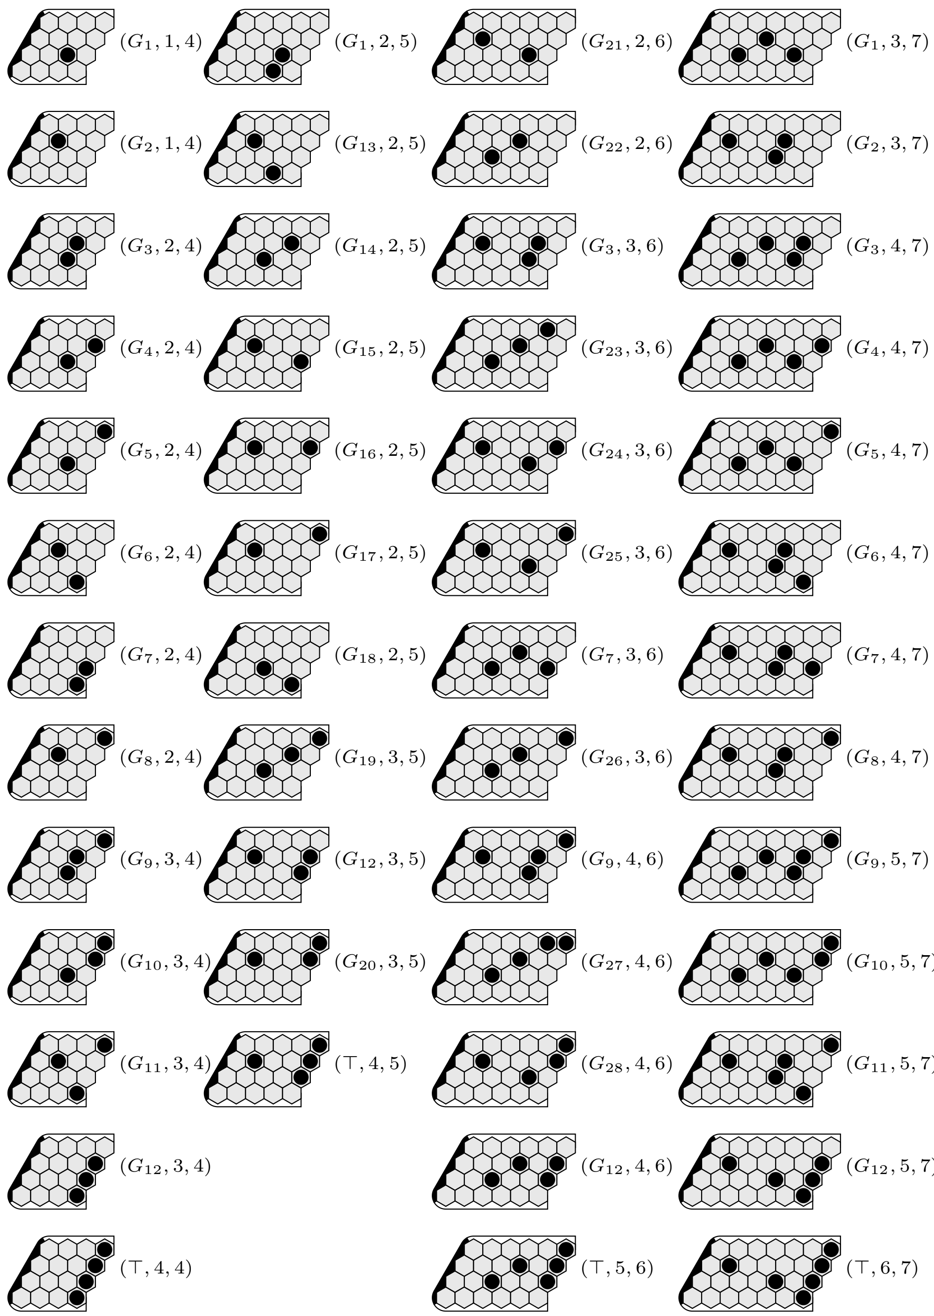

Figure 21: All best patterns of height $k = 4$ and widths $n = 4, 5, 6, 7$.

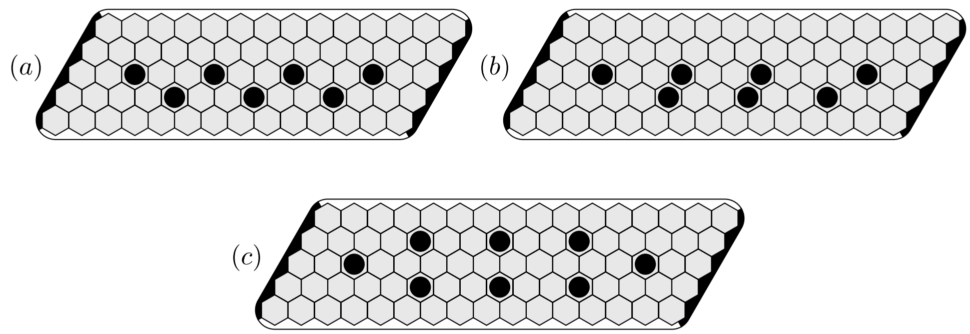

Figure 22: Optimal virtual connections on boards of size 5×14, 5×15, and 5×16.

combinatorial value. Moreover, Figure 21 shows that the minimum number of stones in an acceptable pattern of width 4, 5, and 6 is 1, 2, and 2, respectively. These facts imply the following theorem.

Theorem 12.3. *For all $n \geqslant 4$, the minimal number of stones required for a virtual connection between Black's edges on a board of size $4 \times n$ is exactly $\lceil \frac{2}{3}n - 2 \rceil$. In table form:*

Width of board	4	5	6	7	8	9	10	11	12	13	14	15	16
Minimum stones	1	2	2	3	4	4	5	6	6	7	8	8	9

This confirms — at least in the case of boards of height 4 — what Hex players have long suspected: namely that the asymptotic cost of a "long distance" virtual connection is 2 stones per 3 columns of distance travelled. The optimal connections claimed in Theorem 12.3 can be achieved by patterns similar to the ones shown in Figure 15.

The size of minimal connecting sets is also known for boards of size $1 \times n$, $2 \times n$, $3 \times n$, and $5 \times n$, and is n, $\lceil \frac{2}{3}(n-1) \rceil$, $\lceil \frac{2}{3}n - 1 \rceil$, and $\lceil \frac{2}{3}n - 3 \rceil$, respectively, with the exception of the 5×6 board, where 2 black stones are required. This can be proved by the same method, although for the cases $k = 1, 2, 3$, there exist simpler proofs that do not use combinatorial game theory. Some typical optimal connections for $5 \times n$ are shown in Figure 22; note that in addition to the two long-distance patterns we already saw in Figure 15, an additional new pattern emerges in Figure 22(c). It uses more space, but still requires 2 stones per 3 columns of distance.

12.5. An Inductive Proof of Theorem 12.3

We have derived Theorem 12.3 from a computation, rather than giving a proof in the traditional sense. Of course, it is possible to convert this computation into a proof by induction. In this section, we show how to do this. Given a subset S of any preordered set, let us write $\Downarrow S$ for the down-closure of S, i.e., $\Downarrow S = \{x \mid \text{there exists } y \in S \text{ such that } x \leqslant y\}$.

Definition 12.4. We define the *cost* of a pattern P as

$$\mathrm{cost}(P) = 6 + 3s - 2n,$$

where s is the number of black stones and n is the number of columns in P. We define the *benefit* of P as

$$\mathrm{ben}(P) = \begin{cases} -\infty & \text{if } P \text{ is unacceptable} \\ 0 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_{21}, G_{22}\}, \\ 1 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_1, G_2\}, \\ 2 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_{13}, G_{14}, G_{15}, G_{16}, G_{17}, G_{18}\}, \\ 3 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_3, G_7, G_{23}, G_{24}, G_{25}, G_{26}\}, \\ 4 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_4, G_5, G_6, G_8\}, \\ 5 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_{12}, G_{19}, G_{20}\}, \\ 6 & \text{otherwise, if } \mathrm{val}(P) \in \Downarrow\{G_9, G_{27}, G_{28}\}, \\ 7 & \text{otherwise. Note that in this case, } \mathrm{val}(P) \in \Downarrow\{\top\}. \end{cases}$$

Here, $G_1, \ldots, G_{28}$ are the same values that are shown in Figure 21. Note that ben is a monotone function. We also define the cost of a column to be $\mathrm{cost}(C) = 3s - 2$, where s is the number of black stones in it. Note that $\mathrm{cost}(P + C) = \mathrm{cost}(P) + \mathrm{cost}(C)$.

The point of this definition is that we will prove that a position's benefit is never greater than its cost. The key to this result is the following lemma.

Lemma 12.5. *Whenever P is a pattern and C is a column containing only blank cells and black stones,* $\mathrm{ben}(P + C) \leqslant \mathrm{ben}(P) + \mathrm{cost}(C)$.

Proof. We first consider the special case where $\mathrm{val}(P) \in \{G_1, \ldots, G_{28}, \top\}$. Since there are only 29 patterns and 16 columns to consider, this special case of the lemma can be proved by exhaustively checking all 464 cases.

Now consider some arbitrary pattern P and column C. If $\mathrm{ben}(P) = -\infty$, then so is $\mathrm{ben}(P+C)$ and there is nothing to show. Otherwise, $\mathrm{ben}(P)$ is some number. To illustrate the proof, we consider the case $\mathrm{ben}(P) = 4$; all other cases are similar. By the definition of ben, we know that $\mathrm{val}(P) \leqslant G_4$, $\mathrm{val}(P) \leqslant G_5$, $\mathrm{val}(P) \leqslant G_6$, or $\mathrm{val}(P) \leqslant G_8$. We consider the case $\mathrm{val}(P) \leqslant G_4$; again, the other cases are similar. Now we have

$$\begin{array}{rcll} \mathrm{ben}(P + C) & \leqslant & \mathrm{ben}(G_4 + C) & \text{by monotonicity of ben,} \\ & \leqslant & \mathrm{ben}(G_4) + \mathrm{cost}(C) & \text{by the above special case,} \\ & \leqslant & 4 + \mathrm{cost}(C) & \text{since } \mathrm{ben}(G_4) \leqslant 4 \text{ by definition of ben,} \\ & = & \mathrm{ben}(P) + \mathrm{cost}(C) & \text{by assumption, since } \mathrm{ben}(P) = 4. \end{array}$$

This proves the lemma. □

Lemma 12.6. *For all patterns P of width $n \geqslant 4$,* $\mathrm{ben}(P) \leqslant \mathrm{cost}(P)$.

Proof. We prove this by induction on n. For the base case $n = 4$, there are finitely many patterns to check. Moreover, it suffices to check the 13 patterns shown in column 1 of Figure 21, because they are best, i.e., they maximize the benefit for a given cost, or equivalently, minimize the cost for a given benefit.

For the induction step, suppose the claim is true for patterns of width n, and consider a pattern $P + C$ of width $n + 1$. Then $\mathrm{ben}(P + C) \leqslant \mathrm{ben}(P) + \mathrm{cost}(C) \leqslant \mathrm{cost}(P) + \mathrm{cost}(C) = \mathrm{cost}(P + C)$, where the first inequality holds by Lemma 12.5 and the second inequality holds by the induction hypothesis. $\square$

Proof of Theorem 12.3. We have already found virtual connections with the claimed number of stones; what must be proved is the minimality. Let P be a virtual connection of width n, using s stones. Since P is an acceptable pattern, we have $0 \leqslant \mathrm{ben}(P)$, hence by Lemma 12.6, $0 \leqslant \mathrm{cost}(P)$, i.e., $0 \leqslant 6+3s-2n$, i.e., $s \geqslant \frac{2}{3}n-2$. Since s is an integer, we have $s \geqslant \lceil \frac{2}{3}n - 2\rceil$, as claimed. $\square$

13. Conclusion and Future Work

We described a theory of combinatorial games over partially ordered atom sets that is appropriate for Hex and other monotone set coloring games. The fundamental theorem about these games is that monotone and passable games are the same thing up to equivalence; however, passable games are the more robust class, as they are closed under taking canonical forms. Passable games support many of the usual operations of combinatorial game theory, including sums and opposite games. We briefly discussed the notion of global decisiveness, under which certain games that are otherwise non-equivalent can become equivalent. We enumerated all passable games up to equivalence for certain small atom posets, and proved that the class of passable games is infinite in all other cases. We showed that many game values over small atom posets are realizable as Hex positions, but also proved that there are some values that are not realizable in Hex. Finally, we illustrated the usefulness of this theory by proving a theorem about the minimal size of connecting sets in $k \times n$ Hex.

Many problems are left open. It is not currently known whether every passable game value can be realized by a planar connection game or a vertex Shannon game. It is also not known whether every passable abstract game value for a 3-terminal region is realizable as a Hex position, though we proved this to be false for 4-terminal regions. Indeed, it is not even known whether the number of combinatorial values that are realizable as 3-terminal Hex positions is finite or infinite. The theory of global decisiveness has not been worked out in any detail; in particular, no simple recursive characterization of the globally decisive order has been given, and no canonical form is known for games up to global decisiveness. Finally, we found the size of minimal connecting sets in $k \times n$ Hex for $k \leqslant 5$. While the method generalizes to larger k, the computations get exponentially harder as k increases. The answer to this question is not yet known for $k \geqslant 6$.

There are many potential applications of this theory that can be explored. The theory of passable games can be applied to the analysis of Hex positions, and might potentially lead to useful optimizations in computer Hex. Specifically, for positions that partition the board into multiple disjoint regions, current computer Hex implementations typically have to search many sequences of moves that alternate between the regions. It may be possible to speed up the analysis by analyzing each region separately. One situation where this is likely to be useful is when a region is already *settled*, meaning that its combinatorial value is atomic. In that case, no further moves in that region need to be explored at all. Combinatorial game theory may also lend itself to a more systematic analysis of ladders and ladder escapes in Hex. It can also be used to reason about, and potentially simplify the theory of, capture and domination, which helps players find locally optimal moves and avoid inferior moves. Another possible application of this theory is to help resolve the inferiority or non-inferiority of template intrusions, such as Conjecture 1 of [11]. Some of these applications are already being worked on and will appear in forthcoming papers.

Acknowledgements. I owe an enormous debt of gratitude to Eric Demer, a master Hex player who taught me most of what I know about the game, and who has spent countless late-night hours discussing the combinatorial game theory of Hex with me. Eric has been an invaluable source of insight, examples, and counterexamples, and he has generously permitted me to include many of his contributions in this paper. Among other things, he suggested the region of Figure 13 as an example of a Hex region with 4 linearly ordered outcomes, helped find several of the Hex positions in Figure 14, and introduced me to the concept of global decisiveness and explained its importance. He also provided valuable feedback on an earlier draft of this paper. I would also like to thank the anonymous referee for valuable comments. Any errors are of course mine. This research was supported by NSERC.

References

[1] E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning Ways for Your Mathematical Plays, Volumes 1 and 2*, A. K. Peters, 2nd edn., 2001.

[2] Y. Björnsson, R. Hayward, M. Johanson, and J. van Rijswijck, Dead cell analysis in Hex and the Shannon game, in A. Bondy, J. Fonlupt, J.-L. Fouquet, J.-C. Fournier, and J. L. Ramírez Alfonsín, editors, *Graph Theory in Paris: Proceedings of a Conference in Memory of Claude Berge*, Trends Math., Birkhäuser, 2006, pp. 45–59.

[3] C. Browne, *Hex Strategy: Making the Right Connections*, A. K. Peters, 2000.

[4] D. J. Bush and M. Rice, The game of Pex, https://boardgamegeek.com/boardgame/39288/pex.

[5] J. H. Conway, *On Numbers and Games*, A. K. Peters, 2nd edn., 2001.

[6] E. Demer and P. Selinger, There are infinitely many monotone games over L_5, 2022, available from arXiv:2203.13971.

[7] E. Demer, P. Selinger, and K. Wang, All passable games are realizable as monotone set coloring games, 2021, available from arXiv:2111.10351.

[8] M. Gardner, Mathematical games: Concerning the game of Hex, which may be played on the tiles of the bathroom floor, *Scientific American* **197** (1957)(1), 145–151.

[9] R. B. Hayward and B. Toft, *Hex: The Full Story*, CRC Press, 2019.

[10] P. Hein, Vil de laere Polygon?, *Politiken* (December 26, 1942), translated in [9].

[11] P. Henderson and R. B. Hayward, Probing the 4-3-2 edge template in Hex, in H. van den Herik, X. Xu, Z. Ma, and M. Winands, editors, *Computers and Games (CG 2008)*, vol. 5131 of *Lecture Notes in Comput. Sci.*, Springer, 2008, pp. 229–240.

[12] P. Henderson and R. B. Hayward, Captured-reversible moves and star decomposition domination in Hex, *Integers* **13** (2013), #G1.

[13] R. Milley, *Restricted Universes of Partizan Misère Games*, Ph.D. thesis, Department of Mathematics and Statistics, Dalhousie University, 2013.

[14] J. Nash, Some games and machines for playing them, Tech. Rep. D-1164, Rand Corporation, 1952.

[15] OEIS Foundation Inc., The On-Line Encyclopedia of Integer Sequences, sequence A000108, https://oeis.org/A000108.

[16] S. Reisch, Hex ist PSPACE-vollständig, *Acta Inform.* **15** (1981), 167–191.

[17] P. Selinger, On the combinatorial value of Hex positions, 2021, version 1, available from arXiv:2101.06694v1.

[18] M. Seymour, *Hex: A Strategy Guide*, Online book, available at http://www.mseymour.ca/hex_book/, 2019.

[19] M. Seymour, Hex puzzle collection, http://www.mseymour.ca/hex_puzzle/, 2020.

[20] J. van Rijswijck, *Set Colouring Games*, Ph.D. thesis, Department of Computing Science, University of Alberta, 2006.

[21] Y. Yamasaki, Theory of division games, *Publ. Res. Inst. Math. Sci.* **14** (1978), 337–358.