EXTENDED SPRAGUE-GRUNDY THEORY FOR LOCALLY FINITE GAMES, AND APPLICATIONS TO RANDOM GAME-TREES

James B. Martin
Department of Statistics, University of Oxford, United Kingdom
martin@stats.ox.ac.uk

Received: 2/2/21, Revised: 6/5/21, Accepted: 8/28/21, Published: 12/20/21

Abstract

The Sprague-Grundy theory for finite games without cycles was extended to general finite games by Cedric Smith and by Aviezri Fraenkel and coauthors. We observe that the same framework used to classify finite games also covers the case of locally finite games (that is, games where any position has only finitely many options). In particular, any locally finite game is equivalent to some finite game. We then study cases where the directed graph of a game is chosen randomly, and is given by the tree of a Galton-Watson branching process. Natural families of offspring distributions display a surprisingly wide range of behavior. The setting shows a nice interplay between ideas from combinatorial game theory and ideas from probability.

1. Introduction

Among the plethora of beautiful and intriguing examples to be found in Elwyn Berlekamp, John Conway, and Richard Guy's *Winning Ways* is the game of Fair Shares and Varied Pairs ([1, Chapter 12]). The game is played with some number of almonds, which are arranged into heaps. A move of the game consists of either

- dividing any heap into two or more equal-sized heaps (hence "fair shares"); or
- uniting any two heaps of different sizes (hence "varied pairs").

The only position from which no move is possible is the one where all the almonds are completely separated into heaps of size 1. When that position is reached, the player who has just moved is the winner.

Fair Shares and Varied Pairs is a *loopy* game: the directed graph of game positions has cycles, so the game can return to a previously visited position. The way in which the loopiness manifests itself depends on the number of almonds:

- With 3 or fewer almonds, there are no cycles. The game is *non-loopy.*
- With 4 to 9 almonds, the graph has loops, but all positions are equivalent to finite nim heaps. Hence in any position, either the first player has a winning strategy, or the second player has a winning strategy; furthermore, the same is true for the (disjunctive) sum of any two positions, or for the sum of a position with a nim heap. Berlekamp, Conway and Guy call such behavior *latently loopy.* "This kind of loopiness is really illusory; unless the winner wants to take you on a trip, you won't notice it."
- With 10 almonds, still any position has either a forced win for the first player or a forced win for the second player. However, now there exist some *patently loopy* positions which are not equivalent to finite nim heaps. If one takes the sum of two such positions, or the sum of such a position with a nim heap, one can obtain a game where neither player has a winning strategy – the game is drawn with best play.
- With 11 or more almonds, there exist *blatantly loopy* positions where the game is drawn with best play.

In this article we explore similar themes, but we concentrate particularly on cases where the possibility of draws comes not necessarily from cycles in the game-graph, but instead from infinite paths. (Although the game-graph may be infinite, from any given position there will be be only finitely many possible moves.)

We also focus on situations where the directed graph of the game is chosen at random. The randomness is only in the choice of the graph – that is, of the "rules of the game". All the games themselves will be combinatorial games in the usual sense, with full information and with no randomness.

Here is an example. We will consider a population where each individual reproduces with some given probability $p \in (0, 1)$. If an individual reproduces, it has 4 children. We start with a single individual (the "root"). With probability $1 - p$, the root has no children, and with probability p, the root has 4 children, forming generation 1. If the root does have children, then in turn each of those children itself has no children with probability $1 - p$, and has 4 children with probability p. The collection of those families forms generation 2, whose members again go on to reproduce in the same way, and so on. All the decisions are made independently. From the family tree of this process, we form a directed graph by taking the individuals as vertices, and adding an arc from each vertex to each of its children. This is an example of a *Galton-Watson tree* (or *Bienaymé tree*). In the game played on this tree, from every position there are either 0 or 4 possible moves. Again we consider normal play – if there are no moves possible from a position, the next player to move loses.

Note for example that the tree could be trivial: with probability $1-p$ it consists of just a single vertex. Or it could be larger but finite (its size can be any integer which is congruent to 1 mod 4), as shown in the example in Figure 1.1. But the tree can also be infinite.

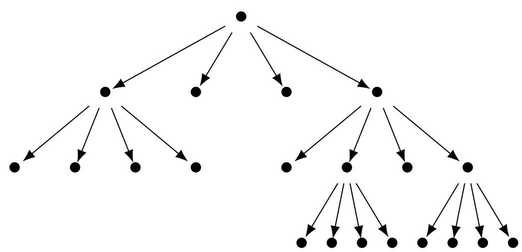

Figure 1.1: An example of a finite directed graph that could arise from the Galton-Watson tree model considered in the introduction with out-degrees 0 and 4.

The game played with such a tree as its game-graph displays very interesting parallels with that of FAIR SHARES AND VARIED PAIRS described above. The behavior depends on the value of the parameter p. We will find that there are thresholds $a_0 = 1/4$, $a_1 \approx 0.52198$, $a_2 = 5^{3/4}/4 \approx 0.83593$ such that the following hold.

- For $p \leq a_0$, the tree is finite with probability 1.
- For $a_0 < p \leq a_1$, there is positive probability that the tree is infinite. However, with probability 1, all its positions are equivalent to finite nim heaps, and so in particular every position has a winning strategy for one or the other of the players.
- For $a_1 < p \leq a_2$, still with probability 1 every position has a winning strategy for one player or the other. However, there is now positive probability that the tree has positions which are not equivalent to finite nim heaps. The sum of two such games, or the sum of such a game with a nim heap, may be drawn with best play.
- For $p > a_2$, with positive probability the tree has positions which are drawn with best play.

1.1. Background and Outline of Results

The equivalence of any finite loop-free impartial game to a nim heap was shown independently by Roland Sprague and by Patrick Grundy in the 1930s. Richard Guy was a key figure in developing and broadening the scope of the Sprague-Grundy theory in the next couple of decades, notably for example in his 1956 paper with Cedric Smith [7].

An extension of the Sprague-Grundy theory to finite games which may contain cycles was first described by Smith [13] and was developed extensively in a series of works by Aviezri Fraenkel and coauthors (for example [4, 2, 5]). As well as *finite-rank* games that are equivalent to nim heaps, one now additionally has *infinite-rank* games which are not equivalent to nim heaps. The "extended Sprague-Grundy value" (or "loopy nim value") of such a game is written in the form $\infty(\mathcal{A})$, where $\mathcal{A} \subset \mathbb{N}$ is the set of nim values of the game's finite-rank options. These infinite-rank games may either be first-player wins (if $0 \in \mathcal{A}$) or draws (if $0 \notin \mathcal{A}$). Again, we have equivalence between two games if and only if they have the same (extended) Sprague-Grundy value.

In [13] Smith already envisages extensions of the theory to infinite games, involving ordinal-valued Sprague-Grundy functions. An extension of a different sort to infinite graphs was done by Fraenkel and Rahat [3], who extend the finite non-loopy Sprague-Grundy theory to infinite games which are *locally path-bounded*, in the sense that for any vertex of the game-graph, the set of paths starting at that vertex has finite maximum length.

In this paper we observe that the extended Sprague-Grundy values which classify finite games are also enough to classify the class of *locally finite games*, in which every position has finitely many options. As a result, any such locally finite (perhaps cyclic) game is equivalent to a finite (perhaps cyclic) game.

We then focus in particular on applying the theory to games whose directed graph is given by a *Galton-Watson tree*, of which the 0-or-4 tree described in the previous section is an example. Galton-Watson trees provide an extremely natural model of a random game-tree. They have a self-similarity which can be described as follows: the root individual has a random number of children (distributed according to the *offspring distribution*), and then conditional on that number of children, the sub-trees of descendants of each of those children are independent and have the same distribution as the original tree.

Games on Galton-Watson trees (including normal play, misère play, and other variants) are studied by Alexander Holroyd and the current author in [9]. There, a particular focus was on determining which offspring distributions give positive probability of a draw, and on describing the type of phase transition that occurs between the sets of distributions with and without draws. In this paper we concentrate on normal play; but, armed with the extended Sprague-Grundy theory, we can investigate, for example, whether infinite-rank positions occur in games without draws (the case analogous to Berlekamp, Conway and Guy's "patently loopy" behavior described above). This setting shows a very nice interplay between ideas from combinatorial game theory and ideas from probability.

One tool on which we rely heavily is the study of the behavior of the game-graph when the set $\mathcal{P}$ of its second-player-winning positions is removed. This reduction behaves especially nicely in the Galton-Watson setting. For example, if we take a Galton-Watson tree for which draws have probability 0, condition the root to be a first-player win, and remove the set $\mathcal{P}$, then the remaining component connected to the root is again a Galton-Watson tree, with a new offspring distribution. Combining iterations of this procedure with recursions involving the probability generating function of the offspring distribution yields a lot of information about the infinite-rank positions that can occur in the tree.

We finish by presenting three particular examples of families of offspring distribution: the Poisson case, the geometric case, and the 0-or-4 case described above. In these examples alone we see a surprisingly wide variety of different types of behavior.

We now briefly describe the organisation of the paper.

In Section 2 we describe the extended Sprague-Grundy theory for locally finite games. Although the setting is new, the results can be written in a form which is almost identical to that of the finite case. We proceed in a way that closely parallels the presentation of Siegel from Section IV.4 of [12] (with some variations of notation). The proofs given in [12] also carry over to the current setting essentially unchanged, and for that reason we do not reproduce them here. A reader who is not already familiar with the extended Sprague-Grundy theory for finite games may like to start with that section of [12] before reading on further here.

In Section 3 we discuss the operation of removing $\mathcal{P}$-positions from a locally finite game, and examine its effect on the Sprague-Grundy values of the positions which remain. For the particular case of trees, we give an interpretation involving *mex labellings* (labellings of the vertices of the tree by natural numbers which obey mex recursions at each vertex).

In Section 4, we introduce games on Galton-Watson trees, and develop the analysis via graph reductions and generating function recursions.

Finally, examples of particular offspring distributions are studied in Section 5.

2. Extended Sprague-Grundy Theory for Games with Infinite Paths

In this section we introduce basic notation and definitions, and then describe the extended Sprague-Grundy theory for locally finite games. The results look identical to those that have previously been written for the case of finite games. Proofs of these results, written for the case of finite games but equally applicable here, can be found in Section IV.4 of [12]. However, note that formally speaking, the content of the results is different; this is not just because the scope of the statements is broader, but also because the definition of equivalence is different (see the discussion in Section 2.3).

2.1. Directed Graphs and Games

We will represent impartial games by directed graphs. If V is a directed graph, we call the vertices of V *positions*. If there is an arc from x to y in V, we write $y \in \Gamma(x)$ (or $y \in \Gamma_V(x)$ if we want to specify the graph V) – here $\Gamma(x)$ is the set of *options* (i.e., out-neighbors) of x. We say that the graph V is *locally finite* if all its vertices have finite out-degree; that is, $\Gamma(x)$ is a finite set for each vertex x. We may be deliberately loose in using the same symbol V to refer both to the graph and to the set of vertices of the graph.

Informally, we consider two-player games with alternating turns; each turn consists of moving from a position x to a position y, where $y \in \Gamma(x)$. We consider normal play: if we reach a *terminal* position, meaning a vertex with outdegree 0, then the next player to move loses. Since the graphs we consider may have cycles or infinite paths, it may be that play continues forever without either player winning.

Formally, a *locally finite game* is a pair $G = (V, x)$ where V is a locally finite directed graph (which is allowed to contain cycles) and x is a vertex of V. We will often write just x instead of (V, x) when the graph V is understood. For example, for the outcome function $\mathcal{O}$, the Sprague-Grundy function $\mathcal{G}$, and the rank function (all defined below), we will often write $\mathcal{O}(x)$, $\mathcal{G}(x)$, and $\text{rank}(x)$, rather than $\mathcal{O}((V, x))$, $\mathcal{G}((V, x))$, and $\text{rank}((V, x))$. We use the fuller notation when we need to consider more than one graph simultaneously (for example when considering disjunctive sums of games, or when considering operations which reduce a graph by removing some of its vertices).

Let V be a directed graph and o a vertex of V. If o has in-degree 0, and if for every $x \in V$, there exists a unique directed walk from o to x, then we say that V is a *tree* with *root* o. If x and y are vertices of a tree V with $y \in \Gamma_V(x)$, we may say that y is a *child* of x in V. We write $\text{height}(x)$ for the *height* of x, which is the number of arcs in the path from o to x.

2.2. Outcome Classes

For a graph V, each position $x \in V$ falls into one of three outcome classes.

- If the first player has a winning strategy from x, then we write $x \in \mathcal{N}$, or $\mathcal{O}(x) = \mathcal{N}$, and say that x is an $\mathcal{N}$-position.
- If the second player has a winning strategy from x, then we write $x \in \mathcal{P}$, or $\mathcal{O}(x) = \mathcal{P}$, and say that x is an $\mathcal{P}$-position.
- If neither player has a winning strategy from x, so that with optimal play the game continues forever without reaching a terminal position, we write $x \in \mathcal{D}$, or $\mathcal{O}(x) = \mathcal{D}$, and say that x is an $\mathcal{D}$-position.

Theorem 2.1. *Let V be a locally finite graph, and $x \in V$.*

- *x is a $\mathcal{P}$-position if and only if every $y \in \Gamma(x)$ is an $\mathcal{N}$-position.*
- *x is an $\mathcal{N}$-position if and only if some $y \in \Gamma(x)$ is a $\mathcal{P}$-position.*
- *x is a $\mathcal{D}$-position if and only if no $y \in \Gamma(x)$ is a $\mathcal{P}$-position, but some $y \in \Gamma(x)$ is a $\mathcal{D}$-position.*

2.3. Disjunctive Sums and Equivalence Between Games

Let V and W be directed graphs. We define $V \times W$ to be the directed graph whose vertices are $\{(x, y), x \in V, y \in W\}$, and which has an arc from (u, v) to (x, y) if and only if either $u = x$ and $y \in \Gamma_W(v)$, or $x \in \Gamma_V(u)$ and $v = y$. If V and W are both locally finite, then so is $V \times W$.

If $G = (V, x)$ and $H = (W, y)$ are locally finite games, we define their (disjunctive) *sum* $G + H$ to be the locally finite game $(V \times W, (x, y))$.

We have the following interpretation. A position of $V \times W$ is an ordered pair of a position of V and a position of W. To make a move in the sum of games, from position (x, y) of $V \times W$, one must either move from x to one of its options in V, or from y to one of its options in W (and not both). The position (x, y) is terminal for $V \times W$ if and only if x is terminal for V and y is terminal for W.

Now we define equivalence between two locally finite games G and H. The games G and H are said to be *equivalent*, denoted by $G = H$, if $\mathcal{O}(G + X) = \mathcal{O}(H + X)$ for every locally finite game X.

Note here that we have defined equivalence within the class of locally finite games: we required the equality to hold for every locally finite game X. The definition (and the meaning of the results below) would be different if X ranged over a different set. However, it will follow from the extended Sprague-Grundy theory below that this equivalence extends both the equivalence within the class of finite loopfree graphs, and that within the class of finite graphs. That is, two finite games are equivalent within the class of finite games if and only if they are equivalent with the class of locally finite games; also two finite loopfree games are equivalent within the class of finite loopfree games if and only if they are equivalent within the class of finite games.

2.4. The Rank Function and the Sprague-Grundy Function

Let V be a locally finite directed graph. We recursively define $\mathcal{G}_n(x)$ for $x \in V$ and $n \geq 0$ as follows. First, let

$$\mathcal{G}_0(x) = \begin{cases} 0, & \text{if } x \text{ is terminal;} \\ \infty, & \text{otherwise.} \end{cases}$$

Then for $n \geq 1$ and given x, write $m = \text{mex}\{\mathcal{G}_{n-1}(y), y \in \Gamma(x)\}$, and let

$$\mathcal{G}_n(x) = \begin{cases} m, & \text{if for each } y \in \Gamma(x), \text{ either } \mathcal{G}_{n-1}(y) \leq m, \text{ or there is } z \in \\ & \Gamma(y) \text{ with } \mathcal{G}_{n-1}(z) = m; \\ \infty, & \text{otherwise.} \end{cases}$$

Proposition 2.2. *Let $x \in V$. Then either:*

- *$\mathcal{G}_n(x) = \infty$ for all n; or*
- *there exist m and n_0 such that*

$$\mathcal{G}_n(x) = \begin{cases} \infty, & \text{if } n < n_0; \\ m, & \text{if } n \geq n_0. \end{cases}$$

In the light of Proposition 2.2, we can now define the *extended Sprague-Grundy function* $\mathcal{G}$ in the case of a locally finite graph V. Let $x \in V$. If the second case of Proposition 2.2 holds, and $\mathcal{G}_n(x) = m$ for all sufficiently large n, then $\mathcal{G}(x) = m$. Otherwise, we write

$$\mathcal{G}_n(x) = \infty(\mathcal{A}),$$

where $\mathcal{A}$ is the finite set defined by

$$\mathcal{A} = \{a \in \mathbb{N} : \mathcal{G}(y) = a \text{ for some } y \in \Gamma(x)\}.$$

We then define the *rank* of x, written $\text{rank}(x)$, to be the least n such that $\mathcal{G}_n(x)$ is finite, or ∞ if no such n exists. (Hence the finite-rank vertices are those x with $\mathcal{G}(x) = m \in \mathbb{N}$, while the infinite-rank vertices are those x with $\mathcal{G}(x) = \infty(\mathcal{A})$ for some $\mathcal{A} \subset \mathbb{N}$.)

Some examples of extended Sprague-Grundy values can be found in Figure 3.1.

Theorem 2.3.

(a) $\mathcal{G}(x) = 0$ if and only if $\mathcal{O}(x) = \mathcal{P}$.

(b) If $\mathcal{G}(x)$ is a positive integer, then $\mathcal{O}(x) = \mathcal{N}$.

(c) If $\mathcal{G}(x) = \infty(\mathcal{A})$ for a set $\mathcal{A}$ with $0 \in \mathcal{A}$, then $\mathcal{O}(x) = \mathcal{N}$.

(d) $\mathcal{G}(x) = \infty(\mathcal{A})$ for some $\mathcal{A}$ with $0 \notin \mathcal{A}$ if and only if $\mathcal{O}(x) = \mathcal{D}$.

Theorem 2.3 tells us that the Sprague-Grundy value of a position determines its outcome class. In fact, much more is true: the Sprague-Grundy values of two games determines the Sprague-Grundy value, and hence the outcome class, of their sum. The algebra of the Sprague-Grundy values is the same as in the case of finite loopy graphs, and full details can be found at the end of Section IV.4 of [12]. Again the proofs carry over unchanged to the locally finite setting. We note a few particular consequences.

Theorem 2.4. *Let G and H be locally finite games.*

(a) $G+H$ *has infinite rank if and only if at least one of G and H have infinite rank.*

(b) If both G and H have infinite rank than $\mathcal{G}(G+H)=\infty(\emptyset)$, and in particular $\mathcal{O}(G+H)=\mathcal{D}$.

(c) If $\mathcal{G}(G)=m\in\mathbb{N}$, then G is equivalent to $*m$, a nim heap of size m.

(d) G and H are equivalent if and only if $\mathcal{G}(G)=\mathcal{G}(H)$.

Corollary 2.5. *Every locally finite game is equivalent to some finite game.*

We finish the section by recording the following consequence of the construction of the extended Sprague-Grundy function, in a form which will be useful for later reference.

Proposition 2.6. *Let V be a locally finite graph, and $x\in V$. Then the following are equivalent:*

(a) $\operatorname{rank}(x)\leq n$ *and* $\mathcal{G}(x)=m$*;*

(b) the following two properties hold:

(i) for each i with $0\leq i\leq m-1$, there exists $y_i\in\Gamma(x)$ such that $\operatorname{rank}(y)<n$ *and* $\mathcal{G}(y_i)=i$*;*

(ii) for all $y\in\Gamma(x)$, either $\operatorname{rank}(y)<n$ *and $\mathcal{G}(y)<m$, or there is $z\in\Gamma(y)$ with* $\operatorname{rank}(z)<n$ *and $\mathcal{G}(z)=m$.*

3. Reduced Graphs

Let $k\geq 0$. We will say that a locally finite directed graph V is *k-stable* if whenever $x\in V$ has infinite rank – that is, whenever $\mathcal{G}((V,x))=\infty(\mathcal{A})$ for some $\mathcal{A}$, then $\{0,1,\ldots,k\}\subseteq\mathcal{A}$.

Note that by Theorem 2.3(d), being 0-stable is equivalent to being *draw-free*: every position of V has a winning strategy either for the first player or for the second player.

Let $\mathcal{P}_V$ be the set of $\mathcal{P}$-positions of the graph V, in other words those $x\in V$ with $\mathcal{G}((V,x))=0$. Consider the graph $R(V):=V\setminus\mathcal{P}_V$ which results from removing the $\mathcal{P}$-positions from V (and retaining all arcs between remaining vertices). More generally, for $k\geq 1$ let $R^k(V)$ be the graph resulting from removing all vertices x with $\mathcal{G}((V,x))<k$.

Theorem 3.1. *Let V be a locally finite directed graph, and let $x\in R(V)$.*

(a) *If x has finite rank in V, then also x has finite rank in $R(V)$; specifically,*

$$\mathcal{G}((R(V),x)) = \mathcal{G}((V,x)) - 1.$$

(b) *Suppose additionally that V is draw-free. If x has infinite rank in V, then also x has infinite rank in $R(V)$; specifically, if $\mathcal{G}((V,x)) = \infty(\mathcal{A})$ for some $\mathcal{A}$ (in which case necessarily $0 \in \mathcal{A}$), then*

$$\mathcal{G}((R(V),x)) = \infty(\mathcal{A} - 1),$$

where $\mathcal{A} - 1$ denotes the set $\{a \geq 0 : a + 1 \in \mathcal{A}\}$.

If V is not draw-free, then the conclusion of part (b) may fail; removing the $\mathcal{P}$-positions may convert infinite-rank vertices to finite-rank vertices (either $\mathcal{P}$-positions or finite-rank $\mathcal{N}$-positions). See Figure 3.1 for an example.

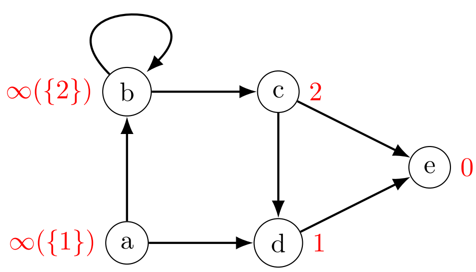

Figure 3.1: The conclusion of Theorem 3.1(b) may fail when the graph is not draw-free. Here, removing the unique $\mathcal{P}$-position e from the graph on the left, to give the graph on the right, converts the position a from infinite rank to finite rank. The extended Sprague-Grundy values are shown by the nodes in red.

Corollary 3.2. *Let $k \geq 1$.*

(a) *Suppose that $V, R(V), \ldots, R^{(k)}(V)$ are all draw-free. Then $R^{(k+1)}(V) = R(R^{(k)}(V))$.*

(b) *V is k-stable if and only if $V, R(V), \ldots, R^{(k)}(V)$ are all draw-free.*

Proof of Theorem 3.1. (a) For the first part, we use induction on the rank of x in V. We claim that if $x \in R(V)$ has $\text{rank}((V,x)) = n$ and $\mathcal{G}((V,x)) = m > 0$, then $\text{rank}((R(V),x) \leq n$ and $\mathcal{G}((R(V),x)) = m - 1$.

Any x with rank 0 in V is in $\mathcal{P}_V$ and hence is not a vertex of $R(V)$, so the claim holds vacuously for x with $\text{rank}((V,x)) = 0$.

Now for $n > 0$, suppose the claim holds for all x with $\text{rank}((V,x)) < n$, and consider $x \in R(V)$ with $\text{rank}((V,x)) = n$ and $\mathcal{G}((V,x)) = m$.

From Proposition 2.6 we have the following properties:

(i) for each $i = 0, \ldots, m-1$, there exists $y_i \in \Gamma_V(x)$ such that $\text{rank}((V, y_i)) < n$ and $\mathcal{G}((V, y_i)) = i$;

(ii) for all $y \in \Gamma_V(x)$, either $\text{rank}((V, y)) < n$ and $\mathcal{G}((V, y)) < m$, or there is $z \in \Gamma_V(y)$ with $\text{rank}((V, z)) < n$ and $\mathcal{G}((V, z)) = m$.

Applying the induction hypothesis we get:

(i) for each $i = 1, \ldots, m-1$, there exists $y_i \in \Gamma_{R(V)}(x)$ such that $\text{rank}((R(V), y_i)) < n$ and $\mathcal{G}((R(V), y_i)) = i - 1$;

(ii) for all $y \in \Gamma_{R(V)}(x)$, either $\text{rank}(R(V), y) < n$ and $\mathcal{G}((R(V), y)) < m - 1$, or there is $z \in \Gamma_{R(V)}(y)$ with $\text{rank}((R(V), z)) < n$ and $\mathcal{G}((R(V), z)) = m - 1$.

Using Proposition 2.6 again we conclude that $\text{rank}((R(V), x)) \leq n$ and $\mathcal{G}((R(V), x)) = m - 1$, completing the induction step.

(b) Now we suppose that in addition V is draw-free. We first want to show that if x has finite rank in $R(V)$, then it also has finite rank in V. In this case we work by induction on the rank of x in $R(V)$.

If x has rank 0 in $R(V)$, (i.e., if x is terminal in $R(V)$), then all options of x in V are in $\mathcal{P}_V$ (i.e., $\mathcal{G}((V, y)) = 0$), which gives $\mathcal{G}((V, x)) = 1$.

Now let $n > 1$. Assume that any vertex with rank less than n in $R(V)$ has finite rank in V, and consider any vertex x with rank n in $R(V)$, say $\mathcal{G}_n((R(V), x)) = m$.

Then using Proposition 2.6 again,

(i) There are $y_0, y_1, \ldots, y_{m-1} \in \Gamma_{R(V)}(x)$ such that for each i, $\text{rank}((R(V), y_i)) < n$ and $\mathcal{G}((R(V), y_i)) = i$. Then by the induction hypothesis, $\text{rank}((V, y_i)) < \infty$, and part (a) gives $\mathcal{G}((V, y_i)) = i + 1$.

(ii) For all $y \in \Gamma_{R(V)}(x)$, either $\text{rank}((R(V), y)) < n$ and $\mathcal{G}((R(V), y)) < m$, or there is $z \in \Gamma_{R(V)}(y)$ such that $\text{rank}((R(V), z)) < n$ and $\mathcal{G}((R(V), z)) = m$. By the induction hypothesis and part (a) again, then either $\mathcal{G}(V, y) < m + 1$ or there is such a z with $\mathcal{G}(V, z) = m + 1$.

Now consider two possibilities. Either there is $y \in \Gamma_V(x)$ with $\mathcal{G}((V, y)) = 0$. Then for some large enough n' we get $\mathcal{G}_{n'}((V, x)) = m+1$, and indeed x has finite rank in V. Alternatively, there is no such y. Then if x had infinite rank in V, we would have $\mathcal{G}((V, x)) = \infty(\mathcal{A})$ for some $\mathcal{A}$ with $0 \notin \mathcal{A}$. This would contradict the assumption that V is draw-free. Hence again x must have finite rank in V, as required. □

3.1. Mex Labellings, and Interpretation of *k*-stability in the Case of Trees

The material in this section is not used in the later analysis, but it aims to give helpful intuition about the notion of k-stability in the case of trees, showing that it can be interpreted in terms of consistency of the set of vertices labelled $0, 1, \ldots, k$ across all labellings which locally respect the mex recursions.

Let V be a locally finite directed graph. We call a function $f : V \to \mathbb{N}$ a *mex labelling* of V if for all $x \in V$, $f(x) = \text{mex}\{f(y), y \in \Gamma_V(x)\}$.

Of course, if V is finite and loop-free, then there is a unique mex labelling f of V given by $f(x) = \mathcal{G}((V, x))$ for $x \in V$.

Notice also that any locally finite tree has at least one mex labelling. To see this we can consider the sequence of finite graphs $(V_n, n \in \mathbb{N})$, where V_n is the induced subgraph of V containing all vertices x such that $\text{height}(x) \leq n$. Each such V_n is finite and loop-free, and so has a mex labelling f_n. In any mex labelling, the vertex x has value no greater than the out-degree of x (which is finite by assumption). Then a compactness/diagonalisation argument shows that there exists a labelling $f : V \to \mathbb{N}$ which, on any finite subset $W \subset V$, agrees with infinitely many of the f_n. In particular, for any vertex x, f agrees with one of the f_n on $\{x\} \cup \Gamma_V(x)$. Then f obeys the mex recursion at every such vertex x, so f is indeed a mex labelling of V.

Proposition 3.3. *Let V be a locally finite tree, and $k \in \mathbb{N}$.*

(a) Suppose V is k-stable. Then the set $\{x \in V : f(x) = k\}$ is the same for all mex labellings f of V, and is equal to $\{x \in V : \mathcal{G}((V, x)) = k\}$.

(b) Suppose V is not k-stable, but is $(k-1)$-stable. (Ignore the vacuous condition of $(k-1)$-stability for $k = 0$.) Let $x \in V$ with $\mathcal{G}(x) = \infty(\mathcal{A})$ for some $\mathcal{A}$ not containing k. Then there are mex labellings f and f' of V with $f(x) = k$, $f'(x) \neq k$.

Note that the conclusion of part (b) can fail even for graphs which are acyclic in the sense of having no directed cycles. See Figure 3.2 for an example. (The method of proof below makes clear that the result does extend to bipartite graphs with no directed cycles.)

Proof. We start by proving that if $x \in V$ has finite rank with $\mathcal{G}(x) = m$, then $f(x) = m$ for all mex labellings f of V. (This holds for any locally finite directed graph V.)

We proceed by induction on $\text{rank}(x)$. Let f be any mex labelling of V.

If $\text{rank}(x) = 0$, then x has no options. Then $\mathcal{G}(x) = 0$, and so $f(x) = \text{mex}(\emptyset) = 0$.

Now suppose $\text{rank}(x) = n > 0$ and $\mathcal{G}(x) = m$, and that the statement holds for all vertices of rank less than n.

From Proposition 2.6, for each i with $0 \leq i \leq m - 1$, there exists $y_i \in \Gamma(x)$ with $\mathcal{G}(y_i) = i$ and $\text{rank}(y_i) < n$. Hence $f(y_i) = i$.

Also for every $y \in \Gamma(x)$ with $\mathcal{G}(y) \geq m$, there is $z \in \Gamma(y)$ with $\text{rank}(z) < n$ and $\mathcal{G}(z) = m$. Then $f(z) = m$, and hence $f(y) \neq m$.

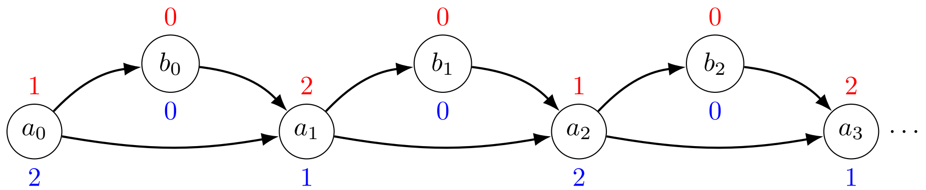

Figure 3.2: An example showing the conclusion of Proposition 3.3(b) can fail even for "loop-free" graphs (i.e., graphs with no directed cycle). The directed graph with vertex set $\{a_i, i \in \mathbb{N}\} \cup \{b_i, i \in \mathbb{N}\}$, and arcs from a_i to b_i, from a_i to a_{i+1}, and from b_i to b_{i+1} for each i. There are two mex labellings, one shown in red above the vertices and the other shown in blue below the vertices. Every position has Sprague-Grundy value $\infty(\emptyset)$, and the graph is not 0-stable. However, the positions b_i have value 0 in both mex labellings, while the positions a_i have non-zero values in both mex labellings.

Thus x has options on which f takes value $0, 1, \ldots, m-1$, but no option on which f takes value m. This gives $f(x) = m$ as required.

To complete the proof of part (a), suppose that V is k-stable, and let f be any mex labelling of V. Then any vertex x with infinite rank has $\mathcal{G}(x) = \infty(\mathcal{A})$ for some $\mathcal{A}$ with $k \in \mathcal{A}$. Hence there exists $y \in \Gamma(x)$ with $\mathcal{G}(y) = k$, giving $f(y) = k$. Then $f(x) \neq k$. So indeed, the set of vertices x with $f(x) = k$ is exactly the set of x with $\mathcal{G}(x) = k$.

We turn to part (b), starting with the case $k = 0$. Suppose that V is a locally finite tree which is not 0-stable. Let x be any vertex with $\mathcal{G}(x) = \infty(\mathcal{A})$ for some $\mathcal{A}$ not containing 0 (that is, $x \in \mathcal{D}$).

Take any $n \geq \text{height}(x)$. Since the game from position x is drawn, if we consider the game on the truncated graph V_n described just before the statement of the proposition, so that all vertices at height n become terminal, then position x becomes a first-player win if $n - \text{height}(x)$ is odd, and a second-player win if $n - \text{height}(x)$ is even.

Then we can apply again the compactness argument mentioned before the statement of Proposition 3.3, separately for odd n and even n. This yields two mex labellings f and f', one of which gives value 0 to x, and the other of which gives a strictly positive value to x, as required. This completes the proof of part (b) in the case $k = 0$.

Now we extend to $k > 0$. Suppose V is $(k-1)$-stable but not k-stable. As in Corollary 3.2, we can apply the reduction operator $k-1$ times, removing all the vertices $y \in V$ with $\mathcal{G}((V, y)) < k$, to arrive at the graph $R^k(V)$.

Any $v \in \mathbb{R}^k(V)$ either has $\mathcal{G}((V, x)) = m$ for some finite $m \geq k$, or $\mathcal{G}((V, x)) = \infty(\mathcal{A})$ for some $\mathcal{A}$ with $\{0, \ldots, k-1\} \subseteq \mathcal{A}$. It is then easy to check that whenever

$\hat{f} : R^k(V) \mapsto \mathbb{N}$ is a mex labelling of $R^k(V)$, we can obtain a mex labelling $f : V \mapsto \mathbb{N}$ of V by defining

$$f(x) = \begin{cases} \mathcal{G}((V,x)), & \text{if } \mathcal{G}((V,x)) < k \\ \hat{f}(x) + k, & \text{otherwise.} \end{cases} \tag{3.1}$$

Let $x \in V$ with $\mathcal{G}(V,x) = \infty(\mathcal{A})$ for some $\mathcal{A}$ containing $0, \ldots, k-1$ but not k. Then, by applying Theorem 3.1 k times, we have $x \in R^k(V)$ and $\mathcal{G}(R^k(V),x)) = \infty(\mathcal{B})$ where $\mathcal{B} = \mathcal{A} - k$. In particular, $0 \notin \mathcal{B}$ (that is, the position x in $R^k(V)$ is a draw). We wish to show that there are mex labellings f, f' of V such that $f(x) = k$ and $f'(x) \neq k$. In light of (3.1), it is enough to show that there are mex labellings $\hat{f}, \hat{f}'$ of $R^k(V)$ such that $\hat{f}(x) = 0$ and $\hat{f}'(x) > 0$.

Since x is a draw in $R^k(V)$, we would like to use the same approach as in the $k = 0$ case. The situation is more complicated since the graph $R^k(V)$ may not be connected. However, the graph $R^k(V)$ is a union of finitely or countably many disjoint trees. Any labelling which restricts to a mex labelling of each tree component is a mex labelling of the whole graph. So it suffices to find mex labellings of the tree component of $R^k(V)$ which contains x, one of which assigns value 0 to x and another of which assigns strictly positive value to x. This indeed can be done using the same compactness argument used in the $k = 0$ case.

This completes the proof of part (b). □

4. Random Game-trees

4.1. Galton-Watson Trees

A *Galton-Watson* (or *Bienaymé*) *branching process* is constructed as follows. We fix some *offspring distribution* which is a probability distribution $\mathbf{p} = (p_k, k \in \mathbb{N})$ on the non-negative integers. The process begins with a single individual, called the root. The root individual has a random number of children, distributed according to the offspring distribution, which form generation 1. Then each of the members of generation 1 has a number of children according to the offspring distribution, forming generation 2, and so on. All family sizes are independent. See for example [6] for a basic introduction, and [11] for much more depth including a rigorous construction.

We derive a directed graph from the process by regarding each individual as a vertex, and putting an arc to each child from its parent. In this way each vertex of the graph has in-degree 1, except for the root which has in-degree 0. We call the resulting graph a *Galton-Watson tree.* This tree has a natural self-similarity property: conditional on the number of the children of the root being k, the subtrees rooted at those children are independent and each one has the distribution of the original Galton-Watson tree.

We assume always that $p_0 > 0$, so that the tree can have terminal vertices.

A key role in what follows will be played by the *probability generating function* of the offspring distribution, defined by

$$\phi(s) = \sum_{k\geq 0} p_k s^k.$$

The function ϕ is strictly increasing on the interval $[0, 1]$, and maps $[0, 1]$ bijectively to the interval $[p_0, 1]$.

A fundamental result is a criterion for the tree to be infinite, in terms of the mean $\mu = \sum_{k\geq 0} kp_k = \phi'(1)$ of the offspring distribution $\mathbf{p}$. Excluding the trivial case $p_1 = 1$ (where with probability 1 the tree consists of a single path) one has that whenever $\mu \leq 1$, the tree is finite with probability 1, and whenever $\mu > 1$, there is positive probability for the tree to be infinite.

If $d = \sup\{k : p_k > 0\}$ is finite, we say the offspring distribution has *maximum out-degree* d. Otherwise we say that the offspring distribution has unbounded vertex degrees.

4.2. Galton-Watson Games

We will consider *Galton-Watson games* i.e., games whose directed graph is a Galton-Watson tree T.

We start with a very simple lemma which helps simplify the language.

Lemma 4.1. *Consider a Galton-Watson tree T, with root o. Let $\mathcal{C}$ be any set of possible Sprague-Grundy values. The following are equivalent:*

(a) $\mathbb{P}(\mathcal{G}((T, o)) \in \mathcal{C}) > 0$*;*

(b) $\mathbb{P}(\mathcal{G}((T, u)) \in \mathcal{C}$ *for some* $u \in T) > 0$*.*

For example, the tree T is draw-free with probability 1 if and only if the probability that the root is drawn is 0. So we do not need to distinguish carefully between saying that "the tree has draws with positive probability" and that "the root is drawn with positive probability". More generally, the tree T is k-stable with probability 1 if and only if the probability that $\mathcal{G}(T, o) = \infty(\mathcal{A})$ for some $\mathcal{A}$ not containing $\{0, 1, \ldots, k\}$ is 0.

Proof of Lemma 4.1. Trivially (a) implies (b). On the other hand, if (a) fails, so that that $\mathbb{P}(\mathcal{G}(T, o)) \in \mathcal{C} = 0$, then the self-similarity of the Galton-Watson tree, the fact that the tree has at most countably many vertices, and the countable additivity of probability measures, combine to give that $\mathbb{P}(\mathcal{G}((T, u)) \in \mathcal{C}$ for some $u \in T) = 0$ also, so that (b) also fails. □

The question of when a Galton-Watson game has positive probability to be a draw was considered in [9].

Let $\mathcal{P}_n$ be the set of vertices from which the second player has a winning strategy that guarantees to win within $2n$ moves (n by each player), and let P_n be the probability that $o \in \mathcal{P}_n$. Note that $o \in \mathcal{P}_n$ if and only if for every child u of o, u itself has a child in $\mathcal{P}_{n-1}$. This leads to the following recursion for the probabilities P_n in terms of the generating function.

$$P_n = 1 - \phi(1 - \phi(P_{n-1})). \tag{4.1}$$

Now let P be the probability that $o \in \mathcal{P}$. We have $P = \lim_{n\to\infty} P_n$. Taking limits in (4.1), and using the fact that the generating function ϕ is continuous and increasing on $[0, 1]$, we obtain part (a) of the following result. A similar approach involving the probability of winning strategies for the first player gives part (b). For full details, see [9].

Proposition 4.2 (Theorem 1 of [9])**.** *Define a function $h : [0, 1] \to [0, 1]$ by*

$$h(s) = 1 - \phi\big(1 - \phi(s)\big). \tag{4.2}$$

(a) $P := \mathbb{P}((T, o) \in \mathcal{P})$ is the smallest fixed point of h in $[0, 1]$.

(b) If $N := \mathbb{P}((T, o) \in \mathcal{N})$, then $1 - N$ is the largest fixed point of h in $[0, 1]$.

Corollary 4.3. *$D := \mathbb{P}((T, o) \in \mathcal{D}) = 1 - N - P$ is positive if and only if the function h defined by (4.2) has more than one fixed point in $[0, 1]$.*

Note that h defined in (4.2) is the second iteration of the function $1 - \phi$. The function $1 - \phi$ is continuous and strictly decreasing, mapping $[0, 1]$ to $[1 - p_0, 0]$. It follows that $1 - \phi$ has precisely one fixed point in $[0, 1]$, and that fixed point is also a fixed point of h. So Corollary 4.3 tells us that the game has positive probability of draws if and only if h has further fixed points which are not fixed points of $1 - \phi$.

Two particular families of offspring distributions had been considered earlier. The Binomial$(2, p)$ case was studied by Holroyd in [8]. The case of the Poisson offspring family is closely related to the analysis of the *Karp-Sipser algorithm* used to find large matchings or independent sets of a graph, which was introduced by Karp and Sipser in [10]; the link to games is not described explicitly in that paper, but the choice of notation and terminology makes clear that the authors were aware of it.

One particular focus of [9] was on the nature of the phase transitions between the set of offspring distributions without draws, and the set of offspring distributions with positive probability of draws. This transition can be either continuous or discontinuous. Without going into precise details, we illustrate with a couple of examples.

Example 4.4 (Poisson distribution – continuous phase transition)**.** The Poisson(λ) offspring family was considered in Proposition 3.2 of [9]. The game has probability 0 of a draw if $\lambda \le e$, and positive probability of a draw if $\lambda > e$. The phase transition is illustrated in Figure 4.1. For $\lambda \le e$, the function h has only one fixed point, while for $\lambda > e$, h has three fixed points. The additional fixed points emerge continuously from the original fixed point as λ goes above e. Note that the probability of a draw at the critical point itself is 0; more strongly, we have the draw probability $\mathbb{P}(o \in \mathcal{D})$ is a continuous function of λ.

Example 4.5 (A discontinuous phase transition)**.** Consider a family of offspring distributions with $p_0 = 1 - a$, $p_2 = a/2$, $p_{10} = a/2$, where $a \in [0, 1]$. This family is used in the proof of Proposition 4(i) of [9]. Again there is some critical point $a_c \approx 0.979$ such that there is positive probability of a draw for $a > a_c$ and not for $a < a_c$. However, unlike in the Poisson case above, at the critical point itself, the function h already has three fixed points, and the probability $\mathbb{P}(o \in \mathcal{D})$ jumps discontinuously from 0 for $a < a_c$ to approximately 0.61 at $a = a_c$ itself. The difference in the nature of the emergence of the additional fixed points of h can be seen by comparing Figure 4.1 and Figure 4.2.

4.3. Existence of Infinite Rank Vertices in Galton-Watson Games

Now we go beyond the question of whether draws have positive probability, to ask more generally about the extended Sprague-Grundy values that occur in a Galton-Watson game. A specific question will be whether, when draws are absent, there are still some infinite rank positions. As suggested by Corollary 3.2, we can investigate the k-stability of the tree T by looking at whether draws occur for the reduced trees $R^k(T)$. The reduction operator behaves particularly nicely in the setting of a Galton-Watson tree.

Theorem 4.6. *Consider a Galton-Watson tree T whose offspring distribution $(p_n, n \ge 0)$ has probability generating function ϕ.*

Suppose the tree is draw-free with probability 1. Let P be the probability that the root o is a $\mathcal{P}$-position.

Condition on the event $\mathcal{O}(o) = \mathcal{N}$, and consider the graph obtained by removing all the $\mathcal{P}$-positions. Let $T^{(1)}$ denote the component connected to the root o in this graph. Then $T^{(1)}$ is again a Galton-Watson tree rooted at o, whose offspring distribution has probability generating function given by

$$\phi^{(1)}(s) = \frac{1}{1-P}\left[\phi\big(P + s(1-P)\big) - \phi\big(s(1-P)\big)\right]. \tag{4.3}$$

Proof. Since we assume that T has no draws, each vertex of T is either a $\mathcal{P}$-position or an $\mathcal{N}$-position. The type of a vertex is determined by the subtree rooted at that

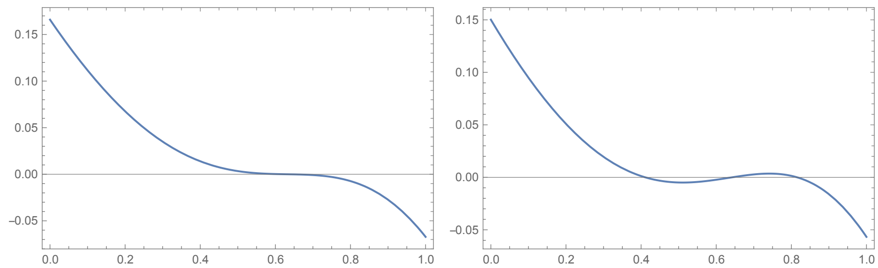

Figure 4.1: An illustration of the phase transition from the non-draw to the draw region, for Poisson(λ) offspring distributions (see Example 4.4). The two plots show the function $h(s) - s$ for $s \in [0, 1]$, where h is defined by (4.2). The fixed points of h are those s where $h(s) - s$ crossing the horizontal axis. On the left, $\lambda = 2.7$, just below the critical point $\lambda = e$; the function h has just one fixed point. On the right, $\lambda = 2.8$, just above the critical point; now h has three fixed points.

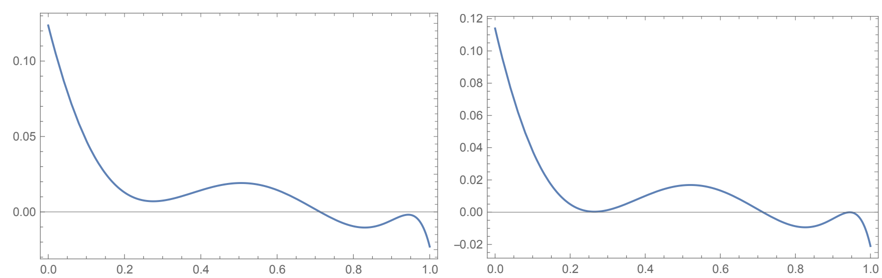

Figure 4.2: The phase transition for the family of offspring distributions given in Example 4.5, with $p_0 = 1 - a, p_2 = a/2, p_{10} = a/2$. Again the function $h(s) - s$ is shown for $s \in [0, 1]$. On the left $a = 0.977$, and on the right $a = 0.979 \approx a_c$. Unlike in Figure 4.1, at the critical point there are already multiple fixed points of h; at a_c, the draw probability jumps from 0 to a positive value around 0.681, which is the distance between the minimum and maximum fixed points of h.

vertex. Conditional on the number of children of the root, the subtrees rooted at each child are independent and each have the same distribution as the original tree. In particular, each child is independently a $\mathcal{P}$-position with probability P, and an $\mathcal{N}$-position with probability $1-P$.

This gives us a two-type Galton-Watson process. We have the familiar recursion that a vertex is an $\mathcal{N}$-position if and only if at least one of its children is a $\mathcal{P}$-position.

We condition on the root being of type $\mathcal{N}$, and retain only its $\mathcal{N}$-type children, and the $\mathcal{N}$-type children of those children, and so on. This gives a one-type Galton-Watson process, and its offspring distribution is the distribution of the number of $\mathcal{N}$-type children of the root in the original process, conditional on the root having type $\mathcal{N}$.

The probability that the root has k $\mathcal{P}$-children and m $\mathcal{N}$-children is

$$p_{m+k}\binom{m+k}{k}P^k(1-P)^m.$$

We can sum over $k \geq 1$ to obtain the probability that the root is of type $\mathcal{N}$ and has m $\mathcal{N}$-children. Finally, we can condition on the event that the root has type $\mathcal{N}$ (which has probability $1-P$), to obtain that the conditional probability that the root has m $\mathcal{N}$-children given that it has type $\mathcal{N}$ is

$$p_m^{(1)} := \frac{1}{1-P}\sum_{k=1}^{\infty} p_{m+k}\binom{m+k}{k}P^k(1-P)^m.$$

Finally we want to calculate the probability generating function $\phi^{(1)}(s) := \sum_{m\geq 0} s^m p_m^{(1)}$ of this distribution. This can easily be done using the binomial theorem to arrive at the form given in (4.3). □

Combining Corollary 3.2 and Theorem 4.6 is the key to studying the infinite-rank vertices of our Galton-Watson tree T; see the strategy described at the beginning of Section 5.

We finish this section with a result about the possible infinite Sprague-Grundy values that can occur in a Galton-Watson game. Essentially, the value $\infty(\mathcal{A})$ has positive probability to appear for every finite $\mathcal{A}$ which is not ruled out either by k-stability or by finite maximum vertex degree. Most notably, part (a)(i) says that for a tree which has draws and for which the offspring distribution has infinite support, *all* finite $\mathcal{A}$ have positive probability.

Proposition 4.7. *Consider the game on a Galton-Watson tree.*

(a) Suppose there is positive probability of a draw.

(i) If the vertex degrees are unbounded, then for any finite $\mathcal{A} \subset \mathbb{N}$, there is positive probability that $\mathcal{G}(o) = \infty(\mathcal{A})$.

(ii) If the maximum out-degree is d, then there is positive probability that $\mathcal{G}(o) = \infty(\mathcal{A})$ if and only if $\mathcal{A} \subset \{0, 1, \ldots, d\}$ with $|\mathcal{A}| \leq d - 1$.

(b) For $k \geq 1$, suppose that the tree is $(k-1)$-stable with probability 1, but has positive probability not to be k-stable.

(i) If the vertex degrees are unbounded, then for any finite $\mathcal{A} \subset \mathbb{N}$ which contains $\{0, \ldots, k-1\}$, there is positive probability that $\mathcal{G}(o) = \infty(\mathcal{A})$.

(ii) If the maximum out-degree is d, then there is positive probability that $\mathcal{G}(o) = \infty(\mathcal{A})$ if and only if $\{0, 1, \ldots, k-1\} \subseteq \mathcal{A} \subset \{0, 1, \ldots, d\}$ with $|\mathcal{A}| \leq d - 1$.

Proof. First we note that all finite Sprague-Grundy values have positive probability, up to the maximum out-degree d if there is one. This is easy by induction. We know that value 0 is possible since any terminal position has value 0. If values $0, 1, \ldots, k-1$ are possible, and it is possible for the root to have degree k or larger, then there is positive probability that the set of values of the children of the root is precisely $\{0, 1, \ldots, k-1\}$, giving value k to the root as required.

Now for part (a), since draws are possible, the value $\infty(\mathcal{B})$ has positive probability for some $\mathcal{B}$ not containing 0. In that case, there is positive probability for all the children of the root to have value $\infty(\mathcal{B})$, and then the root has value $\infty(\emptyset)$.

So the value $\infty(\emptyset)$ has positive probability. Now if $\mathcal{A}$ is any finite set such that the number of children of the root can be as large as $|\mathcal{A}| + 1$, then there is positive probability that the set of values of the children of the root is precisely $\mathcal{A} \cup \{\infty(\emptyset)\}$, and in that case the value of the root is $\infty(\mathcal{A})$ as required.

Finally, if $|\mathcal{A}|$ is greater than or equal to the maximum degree, then the value $\infty(\mathcal{A})$ is impossible, since any vertex with such a value must have at least one child with value m for each $m \in \mathcal{A}$, and additionally at least one child with infinite rank.

We can derive the result for part (b) by applying part (a) to the Galton-Watson tree $T^{(k)}$ obtained by conditioning the root to have Sprague-Grundy value not in $\{0, 1, \ldots, k-1\}$, and removing all the vertices with values $\{0, 1, \ldots, k-1\}$ from the graph, as described above. Theorem 3.1 tells us that if the resulting tree has positive probability to have a node with value $\infty(\mathcal{A})$, then the original tree has positive probability to have a node with value $\infty(\mathcal{B})$ where $\mathcal{B} = \{b \geq k : b - k \in \mathcal{A}\} \cup \{0, 1, \ldots, k-1\}$, and the desired results follow. □

Remark 4.8. Suppose we have a Galton-Watson tree T with positive probability to be infinite, and a set $\mathcal{C}$ of Sprague-Grundy values with $\mathbb{P}(\mathcal{G}(o) \in \mathcal{C}) > 0$. A straightforward extension of Lemma 4.1 says that conditional on T being infinite, with probability 1 there exists $u \in T$ with $\mathcal{G}(u) \in \mathcal{C}$.

Combining with Proposition 4.7, we get the following appealing property. If T has unbounded vertex degrees, and positive probability of draws, then conditional on

T being infinite, with probability 1, vertices with every possible extended Sprague-Grundy value are found in the tree.

5. Examples

First we lay out how to use the results of the previous sections to address the question of which infinite-rank Sprague-Grundy values have positive probability for a given Galton-Watson tree T.

Let ϕ be the probability generating function of the offspring distribution of T. To examine whether T can have draws, we apply the criterion given Corollary 4.3: T is draw-free with probability 1 if and only if the function $h(s) = 1 - \phi(1 - \phi(s))$ has a unique fixed point.

If so, we use the procedure in Theorem 4.6. We condition the root to be a $\mathcal{N}$-position, we remove all the $\mathcal{P}$-positions, and we retain the connected component of the root, to obtain a new Galton-Watson tree $T^{(1)}$ with an offspring distribution whose probability generating function is $\phi^{(1)}$. We then examine whether or not this new generating function gives a draw-free tree.

If it does, we can repeat the procedure again, producing a new generating function which we call $\phi^{(2)}$, corresponding to removing the positions with Sprague-Grundy values 0 and 1 from the original tree.

If we perform k reductions and still have a draw-free tree at every step, this tells us that our original tree was k-stable with probability 1.

If the iteration of this procedure never produces a tree with positive probability of a draw, then the original tree had probability 0 of having infinite-rank vertices. (Note that for example if at any step we arrive at a tree which is sub-critical, i.e., whose offspring distribution has mean less than or equal to 1 and which therefore has probability 1 to be of finite size, then we know that every further reduction must give rise to a draw-free tree.)

We now apply this strategy to a few different examples of families of offspring distributions. We see a surprising range of types of behavior.

Example 5.1 (Poisson case, continued)**.** Galton-Watson trees with Poisson offspring distribution behave particularly nicely under the graph reduction operation. This allows us to give a complete analysis of the Poisson case without any need for calcuations or numerical approximation.

The tree has positive probability to be infinite precisely when $\lambda > 1$. We already saw in Example 4.4 that there is positive probability of a draw precisely when $\lambda > e$.

Suppose we are in the $\lambda \leq e$ case without draws. So each node is a $\mathcal{P}$-node (with probability P) or a $\mathcal{N}$-node (with probability $1 - P$).

By basic properties of the Poisson distribution, the number of $\mathcal{P}$-children of the root is Poisson(λP)-distributed, and the number of $\mathcal{N}$-children of the root is

Poisson($\lambda(1-P)$)-distributed, and the two are independent.

If we condition the root to have at least one $\mathcal{P}$-child, and then remove all its $\mathcal{P}$-children, then because of the independence of the number of $\mathcal{P}$-children and the number of $\mathcal{N}$-children, we are simply left with a Poisson($\lambda(1-P)$) number of children.

So we again have a Poisson Galton-Watson tree, but now with a new parameter $\lambda^{(1)} < \lambda$. Since $\lambda^{(1)} < e$, the new tree is still draw-free with probability 1.

Hence, to adapt the terminology of the introduction, in the Poisson case we may see a "blatantly infinite" game once $\lambda > e$, but for $\lambda \leq e$ we are at worst "latently infinite". There is no λ which gives "patently infinite" behavior whereby draws are absent but infinite rank vertices have positive probability.

Example 5.2 (Degrees 0 and 4)**.** We return to the example in the introduction, where all outdegrees are 0 or 4. We have $p_4 = p$ and $p_0 = 1 - p$ for some $p \in (0, 1)$.

If $p \leq a_0 := 1/4$ then the mean offspring size is less than or equal to 1, and the tree is finite with probability 1.

One can show algebraically that there is positive probability of a draw if and only if $p > a_2 := 5^{3/4} \approx 0.83593$. Namely one can obtain that the function h defined in (4.2) has derivative less than 1 on $[0, 1]$ for all $p \leq a_2$ (except for a single point in the case $p = a_2$), and so r has just one fixed point for such p. Meanwhile for $p > a_2$ there is a fixed point s^* of the function $1 - \phi$ for which $h'(s^*) > 1$, and this can be used to show that r has at least two further fixed points. Corollary 4.3 then gives the result.

Between a_0 and a_2 there exist no draws, but the tree is infinite with positive probability, so we may ask whether there can exist positions with infinite rank.

Numerically, we observe a phase transition around the point $a_1 \approx 0.52198$. For $p \leq a_1$, we know that the tree T has zero probability of a draw, and we observe that the same is also true for the trees $T^{(1)}$ and $T^{(2)}$ (their maximum out-degrees are 3 and 2 respectively, so their generating functions $\phi^{(1)}$ and $\phi^{(2)}$ are cubic and quadratic respectively. The tree $T^{(3)}$ has vertices of out-degrees only 0 and 1, and will also be finite with probability 1, so we do not need to examine $T^{(k)}$ for any higher k.)

Hence for $p \in (a_0, a_1]$, we have the "latently infinite" phase where all Sprague-Grundy values are finite with probability 1.

However, for $p \in (a_1, a_2]$ we observe that the function $h^{(1)}(s) := 1 - \phi^{(1)}(1 - \phi^{(1)}(s))$ has more than one fixed point. Consequently, there is positive probability of a draw in the tree $T^{(1)}$. The tree T has positive probability not to be 1-stable, and so to have positions of infinite rank.

The behavior of h, $h^{(1)}$ and $h^{(2)}$ around the phase transition point $p = a_1$ is shown in Figure 5.1. Although the precise nature and location of this phase transition is only found numerically, it is not hard to show rigorously that for p just above a_0, the functions $h^{(1)}$ and $h^{(2)}$ have only one fixed point, while for p just below a_2, the

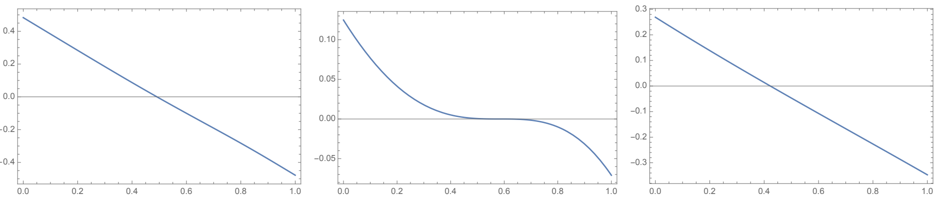

Figure 5.1: The case of the 0-or-4 distribution from Example 5.2, with $p = 0.52198 \approx a_1$. From left to right the three graphs show the functions $h(s) - s$, $h^{(1)}(s) - s$, and $h^{(2)}(s) - s$. As p moves through the critical point a_1, the function $h^{(1)}$ acquires multiple fixed points. For $p \leq a_1$, the tree has only finite-rank vertices. For $p > a_1$, the tree no longer has probability 1 to be 1-stable, and for example the Sprague-Grundy value $\infty(0)$ has positive probability.

function $h^{(1)}$ has more than one fixed point, so that the family of distributions does display all four of the "finite", "latently infinite", "patently infinite" and "blatantly infinite" types of behavior.

Example 5.3 (Geometric case)**.** We now consider the family of geometric offspring distributions, with $p_k = q^k(1-q)$ for $k = 0, 1, 2, \dots$, for some $q \in (0, 1)$.

Rather surprisingly, there is no q for which draws have positive probability! See for example Proposition 3(iii) of [9]. (This shows for example that the property of having positive probability of draws is not monotone in the offspring distribution. If we take any $\lambda > e$, then as discussed above, the Poisson(λ) distribution has positive probability of draws, but for q sufficiently large, this distribution is stochastically dominated by a Geometric(q) distribution, which does not have draws.)

However, other interesting phase transitions for the geometric family do occur. Numerically, we observe that there are critical values $q_0 = 1/2, q_1 \approx 0.88578, q_2 \approx 0.88956, q_3 \approx 0.923077$ such that the following hold.

- For $q \leq 0.5$, the tree is finite with probability 1.
- For $q \in (0.5, q_1]$, there are infinite paths with positive probability, but the tree is 3-stable with probability 1. In fact for q sufficiently close to 0.5, the tree $T^{(1)}$ is finite with probability 1, and so in fact the tree is k-stable for all k, i.e., all positions have finite rank (the latently infinite phase). It seems plausible that in fact the latently infinite phase continues all the way to q_1, but we do not know how to demonstrate that.
- For $q \in (q_1, q_2]$, with positive probability the tree is not 3-stable; however it continues to be 2-stable.
- For $q \in (q_2, q_3]$, with positive probability the tree is not 2-stable; however it continues to be 1-stable.

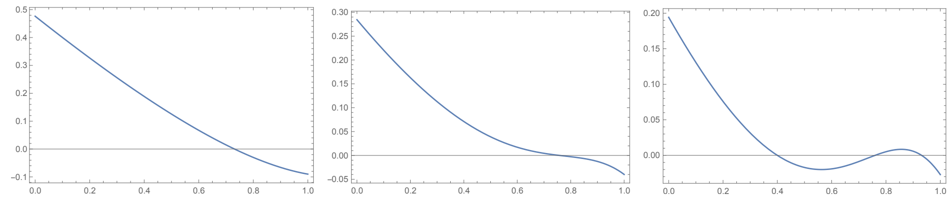

Figure 5.2: The geometric case of Example 5.3 with $q = 0.91 \in [q_2, q_3)$. As in Figure 5.1, we plot the functions $h(s)-s$, $h^{(1)}(s)-s$, and $h^{(2)}(s)-s$. The functions h and $h^{(1)}$ have unique fixed points, but the function $h^{(2)}$ has multiple fixed points; so the tree has probability 1 to be 1-stable, but has probability less than 1 of being 2-stable.

- For $q \geq q_3$, with positive probability the tree is not 1-stable (but as we know, it continues to be 0-stable, in other words draw-free, for all q).

Except for the transition at q_0, the precise nature and location of all the phase transitions above are only found numerically. However, with a sufficiently precise analysis one could rigorously establish in each case a smaller interval on which the claimed behavior holds (for example we could find some sub-interval of the claimed interval (q_2, q_3) on which to show that $h^{(1)}$ has only one fixed point while $h^{(2)}$ has more than one fixed point).

In summary, the three families in Examples 5.1-5.3 show a wide variety of behaviors. In the Poisson case, one has existence of draws whenever one has existence of positions with infinite rank. In the 0-or-4 case, there is additionally a phase wth infinite rank vertices but no draws. In the geometric case, it is the phase with draws which is missing; however, one sees additional phase transitions, losing 3-stability, 2-stability, and 1-stability step by step as the parameter increases.

We end with a question.

Question 5.4. Does there exist for every $k \in \mathbb{N}$ an offspring distribution for which the Galton-Watson tree is k-stable with probability 1, but nonetheless infinite rank positions exist with positive probability? Numerical explorations have so far only produced examples up to $k = 2$ (for example, the Geometric(q) case with $q \in (q_1, q_2]$ described above).

Acknowledgments. Many thanks to Alexander Holroyd and Omer Angel for conversations relating to this work, particularly during the Montreal summer workshop on probability and mathematical physics at the CRM in July 2018. I am grateful to an anonymous referee, whose comments have considerably improved the paper.

References

[1] E. R. Berlekamp, J. H. Conway, and R. K. Guy, *Winning Ways for Your Mathematical Plays, Volume 2*, 2nd edition, CRC Press, 2003.

[2] A. S. Fraenkel and Y. Perl, Constructions in combinatorial games with cycles, in *Infinite and Finite Sets (to Paul Erdős on his 60th birthday), Volume 2*, ed. A. Hajnal, R. Rado, and V. Sós, pp. 667-699, North-Holland, 1975.

[3] A. S. Fraenkel and O. Rahat, Infinite cyclic impartial games, *Theor. Comp. Sci.* **252** (2001), 13-22.

[4] A. S. Fraenkel and U. Tassa, Strategy for a class of games with dynamic ties, *Comput. Math. Appl.* **1** (1975), 237-254.

[5] A. S. Fraenkel and Y. Yesha, The generalized Sprague-Grundy function and its invariance under certain mappings, *J. Comb. Thy. A* **43** (1986), 165-177.

[6] G. Grimmett and D. Welsh, *Probability: an Introduction*, 2nd edition, Oxford University Press, Oxford, 2014.

[7] R. K. Guy and C. A. B. Smith, The G-values of various games, *Math. Proc. Camb. Phil. Soc.* **52** (1956), 514–526.

[8] A. E. Holroyd, *Percolation Beyond Connectivity*, PhD thesis, Univ. Cambridge, 2000.

[9] A. E. Holroyd and J. B. Martin, Galton-Watson games, *Random Structures & Algorithms*, to appear, DOI 10.1002/rsa.21008.

[10] R. M. Karp and M. Sipser, Maximum matching in sparse random graphs, in *22nd Annual Symposium on Foundations of Computer Science*, pp. 364-375, IEEE, 1981.

[11] J.-F. Le Gall, Random trees and applications, *Probab. Surveys* **2** (2005), 245-311.

[12] A. N. Siegel, *Combinatorial Game Theory*, Amer. Math. Soc., Providence, Rhode Island, 2013.

[13] C. A. Smith, Graphs and composite games, *J. Comb. Thy.* **1** (1966), 51-81.