

Random Perfect Information Games

Citation for published version (APA):

Flesch, J., Predtetchinski, A., & Suomala, V. (2023). Random Perfect Information Games. *Mathematics of Operations Research*, 48(2), 708-727. <https://doi.org/10.1287/moor.2022.1277>

Document status and date:

Published: 01/05/2023

DOI:

[10.1287/moor.2022.1277](https://doi.org/10.1287/moor.2022.1277)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Random Perfect Information Games

János Flesch, Arkadi Predtetchinski, Ville Suomala

To cite this article:

János Flesch, Arkadi Predtetchinski, Ville Suomala (2022) Random Perfect Information Games. Mathematics of Operations Research

Published online in Articles in Advance 29 Aug 2022

. <https://doi.org/10.1287/moor.2022.1277>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Random Perfect Information Games

János Fleisch,^a Arkadi Predtetchinski,^{b,*} Ville Suomala^c

^aDepartment of Quantitative Economics, Maastricht University, 6200 MD Maastricht, Netherlands; ^bDepartment of Economics, Maastricht University, 6200 MD Maastricht, Netherlands; ^cDepartment of Mathematical sciences, University of Oulu, FI-90014 Oulu, Finland

*Corresponding author

Contact: j.flesch@maastrichtuniversity.nl (JF); a.predtetchinski@maastrichtuniversity.nl,  https://orcid.org/0000-0002-6658-303X (AP); ville.suomala@oulu.fi (VS)

Received: May 27, 2021

Revised: March 3, 2022

Accepted: April 5, 2022

Published Online in Articles in Advance:
August 29, 2022

MSC2020 Subject Classification: Primary:
91A70, 91A25, 91A18; secondary: 60J85

https://doi.org/10.1287/moor.2022.1277

Copyright: © 2022 INFORMS

Abstract. The paper proposes a natural measure space of zero-sum perfect information games with upper semicontinuous payoffs. Each game is specified by the game tree and by the assignment of the active player and the capacity to each node of the tree. The payoff in a game is defined as the infimum of the capacity over the nodes that have been visited during the play. The active player, the number of children, and the capacity are drawn from a given joint distribution independently across the nodes. We characterize the cumulative distribution function of the value v using the fixed points of the so-called value-generating function. The characterization leads to a necessary and sufficient condition for the event $v \geq k$ to occur with positive probability. We also study probabilistic properties of the set of player I 's k -optimal strategies and the corresponding plays.

Keywords: zero-sum game • perfect information • value • Galton–Watson measure • branching process

1. Introduction

Perfect information games are probably one of the most thoroughly researched classes of games. Descriptive set theory (Martin [20], Moschovakis [23]), computer science (Apt and Grädel [2]), logic (Van Benthem [32]), and economics (e.g., Harris [15]) all employ their own distinct methodologies to study perfect information games. Among these, two-player, zero-sum games with a semicontinuous payoff function are arguably the simplest kind, intimately related to games with closed winning sets introduced by Gale and Stewart [11]. It is this, relatively simple class of infinite perfect information games to which our study is devoted.

In this paper, we take a probabilistic point of view on perfect information games. In a nutshell, this view amounts to the following: we consider a particular space of games and a natural probability measure over this space. Nature selects a game randomly according to the given measure, the two players observe the entire realization of the game, and the play commences. It is important to stress at the outset that, once the game has been chosen by nature, the players face no further randomness: they observe the realization of the game before the play begins and may adjust their strategies depending on the game at hand. Under this approach, the value of the game is a random variable. The distribution of the value is the principal object of this study.

The probabilistic point of view described above is certainly no novelty in game theory. There is a rich and mature literature on random, one-shot (i.e., normal form) games. This literature examines, among other questions, the expected number of pure Nash equilibria (e.g., Daskalakis et al. [6], Rinott and Scarsini [29], Stanford [30], Takahashi [31]), the expected number of mixed Nash equilibria (e.g., McLennan [21], McLennan and Berg [22]), the distribution of the number of point-rationalizable strategies (Pei and Takahashi [26]), convergence properties of reinforcement learning (e.g., Galla and Farmer [12]), and the probability for the best-response dynamics to reach a Nash equilibrium (e.g., Amiet et al. [1]).

On the other hand, randomly generated extensive-form games receive less attention. One notable exception is Arieli and Babichenko [3], who consider extensive-form perfect information games of given length in which the payoffs at the end nodes are randomly generated. Their focus is on the asymptotic distribution of subgame perfect equilibrium payoffs as the length of the games increases.

Recently, percolation theory produced works on random perfect information games as researchers consider adversarial versions of the classical percolation problem.¹ Basu et al. [5] examine zero-sum perfect information games played on a random graph obtained from \mathbb{Z}^d by removing each node with a certain probability. The players take turns to move a token along the edges of the graph. Each node of the graph is only allowed to be visited by the token once. A player who has no legal moves is declared a loser. Holroyd et al. [17] consider a game on

the square lattice \mathbb{Z}^2 , whose every node is randomly classified as either a “trap,” a “target,” or an “open” node. A player can move the token to the node either to the right or above its current position. If a player moves a token into a target (respectively, a trap), she receives a payoff of +1 (respectively, -1), and if no traps or targets are ever visited, both players receive payoff 0. The study most closely related to ours is Holroyd and Martin [16]. There, three classic combinatorial games (normal, misère, and escape games) are played on a tree drawn randomly from a Galton–Watson measure.

1.1. The Model

Two players, I (Alice) and II (Bob), play against each other in a perfect information game that is chosen randomly by nature. Nature chooses a finitely branching game tree, selects the active player (either I or II) for each node of the tree, and assigns to every node a nonnegative number called a capacity. The number of children (i.e., the successors) of a node in the tree, the active player, and the capacity are drawn by nature from a joint distribution independently across the nodes. For a given node, however, the active player, the capacity, and the number of children might well be interdependent. Formally, a measure space of games is a space of marked trees (as has been introduced by Neveu [25]) and is a natural generalization of the Galton–Watson measure space of trees.

Player I’s payoff in a game is defined as the infimum of the capacity encountered in the course of play. Before play commences, both players observe the entire game, including the tree and the assignment of the active players and the capacity to its nodes. They are free to choose a strategy depending on the realization of the game at hand.

Note that, in any realization of the game, the game tree is finitely branching, and player I’s objective function is bounded and upper semicontinuous. In particular, each realization of the game has a value.

The model shares several features with and complements those in Arieli and Babichenko [3], Holroyd et al. [17], and Holroyd and Martin [16].

As in Holroyd and Martin [16], the game tree in our model is drawn from a Galton–Watson measure. In a departure from Holroyd and Martin [16], we employ a random assignment of players to the nodes of the game tree, an idea we borrow from Arieli and Babichenko [3]. The probability for player I to be assigned to a node of the game tree (called I’s activation probability) plays a key role in the analysis. The nonadversarial case when player I is assigned as an active player with probability one serves as a natural benchmark. The main advantage of the random assignment of an active player is that this modeling choice leads to a relatively tractable characterization of the distribution of the value.

A novel feature of the model is the nature of the payoff function. In Holroyd et al. [17] (as well as in the two antecedent works Basu et al. [5] and Holroyd and Martin [16]), all infinite branches are assigned the same payoff. This is but one of the specifications for which our model allows. The possibility of choosing a joint distribution of the active player, the capacity, and the number of children affords our model a lot of flexibility and allows it to accommodate many examples of nontrivial payoff functions. Some of the examples are these.

If a childless node has zero capacity and a node with at least one child has a capacity one, we obtain a payoff function assigning payoff zero to all the end nodes and payoff one to all the infinite branches. If the capacity equals the number of children, we obtain a payoff function that equals the least number of children over the nodes visited in the course of play. One other specification of interest is obtained if all nodes have two children and the capacity is uniformly distributed on $[0, 1]$. This yields a game on a complete binary tree equipped with a nontrivial payoff function over the Cantor space of plays.

The model also accommodates certain classes of finite games. Indeed, if the mean of the offspring distribution is smaller than one, then the game tree is finite almost surely. If we specify the capacity to be one whenever a node has at least one child and to be uniformly distributed on $[0, 1]$ if a node has no children, we obtain a finite game with the payoffs independently assigned to the end nodes of the tree.

Two examples are developed in detail in the main body of the paper. We also attend closely to a special case of the model called an escape model. In the escape model, the end nodes, that is, nodes without children, have capacity equal to zero. Under this specification, player I’s primary objective is to avoid the end nodes of the game tree (and, thus, “escape to infinity”) for reaching such a node results in the lowest possible payoff: zero.

1.2. The Results

We split the results into two groups: (I) the main results and (II) the corollaries.

I. The Main Result. The main result of the paper is a fixed point characterization of the distribution of the value: given $k > 0$, the probability that the value v is strictly less than k is shown to be a fixed point of the so-called value-generating function (*v*gf for short). The result leads to a necessary and sufficient condition for the event $v \geq k$ to

have a positive probability. This criterion generalizes the classical condition for (non)-extinction of a branching process.

To obtain the main result, we rely on a familiar technique of truncating a game at some period t and use the fact that the value of the truncated games converges to the value of the original game as the truncation horizon increases. The role of the value-generating function in our analysis is somewhat similar to the role of the probability-generating function in the study of branching processes. Intuitively, the vgf maps the distribution of the value at the next period to that at the current period. In particular, the t th iterate of the vgf describes the distribution of the value in the t -truncated game.

II. Corollaries. We point out several features of the distribution of the value.

Critical Activation Probability. In the escape model, the probability of the event $v \geq k$ is only positive if player I's activation probability is above the so-called k -critical level. This shows that, if player I controls too few nodes, the player has no chance of obtaining a payoff greater than k . We derive an explicit expression for the k -critical activation probability.

Distribution of the Value Conditional on the Active Player. One statistic of interest is the distribution of the value conditional on the active player or, more precisely, conditional on the event that the root of the game tree is assigned to player I (player II). We find (under mild and natural conditions) that the distribution of the value at player I's nodes first order stochastically dominates that at player II's nodes. We also explore the relationship between the conditional probability of the event $v \geq k$ at player I's (player II's) nodes as player I's activation probability approaches its k -critical value.

Asymptotic Result for Games on Complete n -ary Trees. Suppose that the game tree is a complete n -ary tree, that is, each node of the tree has exactly n children. How does the probability of the event $v < k$ behave as n becomes large? We show that it is eventually monotone and compute its limit as n goes to infinity.

Player I's k -optimal Strategies and k -optimal Subtree. Player I's strategy is said to be k -optimal if it yields a payoff of at least k against any strategy of player II. A closely related notion is that of a k -optimal subtree. It is the subtree of the game tree consisting of the nodes having a value of at least k . Player I's strategy is k -optimal if and only if it never leads outside the k -optimal subtree. In this way, the k -optimal subtree characterizes the set of k -optimal strategies. We discuss the distribution of the k -optimal subtree conditional on the event that the value of the game is at least k .

Avoidance Game. In the avoidance game, player I is prohibited from visiting the nodes controlled by player II with the exception of the nodes having no children or a single child. We think of this scenario as a proxy for the situation in which player I is reluctant to concede a turn to her opponent because of, for example, security concerns. Under this scenario, player II is being deprived of any real choice in the game and is merely a dummy. We study the distribution of the value of and the set of player I's k -optimal strategies in the avoidance game.

The paper is organized as follows. Section 2 introduces the model. Section 3 illustrates the model and gives a flavor of our results by means of two examples. We keep going back to these two examples throughout the paper. Section 4 develops the key result of the paper: the characterization of the distribution of the value. Section 5 presents the corollaries. The final section contains further examples, some discussion, and open questions.

2. Random Perfect Information Games

Two players, I (Alice) and II (Bob), play against each other in a game that is chosen randomly by nature. The game is specified by a game tree and by an assignment to its every node of the active player (either I or II) and of the capacity (a nonnegative real number). The tree, the active players, and the capacities are all generated by nature. Below, we provide the technical details on the space of games and the measure from which a game is drawn. Informally, one could think of the number of children of a node, the active player, and the capacity being drawn by nature from a joint distribution independently across the nodes.

Upon observing the realization of the game, the two players decide on their strategies. The game is zero-sum, player I being the maximizer. Player I's payoff is the smallest capacity encountered during play. A special case is the so-called escape model, in which the capacity is set equal to zero whenever the node has no children. In the escape model, player I's primary objective is to avoid childless nodes (i.e., "escape to infinity") because arriving at a childless node leads to the smallest possible payoff, zero.

Table 1. Summary of notation (with references to subsections).

p	The primitive distribution (a probability measure on $S = \{I, II\} \times [0, \infty) \times \mathbb{N}$)	2.1
ι, γ, ξ	The three coordinate functions on S (the active player, the capacity, the number of children)	2.1
q_i	The probability $p(\iota = i)$ (player i 's activation probability)	2.1
$p_i(k, n)$	The probability $p(\{\iota = i\} \cap \{\gamma \geq k\} \cap \{\xi = n\})$	2.1
$G_i(k, x)$	Generating function for the sequence $\{p_i(k, n)\}_{n \in \mathbb{N}}$	2.1
ω	A game, the triple $(T_\omega, \iota_\omega, \gamma_\omega)$	2.3
$\Omega, \mathbb{P} = \mathbb{P}_p$	the space of games and a probability measure on Ω	2.4
$E(i, k, n)$	the event $\iota(\omega) = i, \gamma(\omega) \geq k, \xi(\omega) = n$	2.3
σ_i	A strategy for player $i \in \{I, II\}$	2.6
$v = v_\omega$	The value of the game ω	2.7
q_c	Critical activation probability	5.1
f	The value generating function (vgf)	4.1
$d(k)$	A key parameter: $\mathbb{E}_p(\mathbf{1}_{\{\iota=I\} \cap \{\gamma \geq k\}} \xi) + p(\{\iota = II\} \cap \{\gamma \geq k\} \cap \{\xi = 1\})$	4.1
ω_i	The truncated game	4.2
$v_i = v_{\omega_i}$	The value of the truncated game	4.2
$\alpha(k)$	Probability of the event $v < k$	4.3
$\beta(k)$	Probability of the event $v \geq k$	4.3

Each realization of the game is a zero-sum game with bounded upper semicontinuous payoffs. As is well-known, each such game has a value. The distribution of the value is at the focus of the paper.

Table 1 summarizes the notation.

2.1. The Primitive Distribution

Let $\mathbb{N} = \{0, 1, \dots\}$ and $\mathbb{N}_+ = \{1, 2, \dots\}$.

The *primitive distribution*, denoted by p , is a joint probability distribution of the triple of random variables (ι, γ, ξ) with ι taking values in the set $\{I, II\}$, γ in $[0, \infty)$, and ξ in \mathbb{N} . As will be clear shortly, the three random variables represent the active player, the capacity, and the number of children at a particular node of the game tree. Formally, we treat p as a probability measure on the Borel subsets of $S = \{I, II\} \times [0, \infty) \times \mathbb{N}$ (where, of course, $\{I, II\}$ and \mathbb{N} are equipped with the discrete topology) and ι, γ , and ξ as the three coordinate functions on S . For $(i, k, n) \in S$, let

$$q_i =: p(\{\iota = i\})$$

and

$$p_i(k, n) =: p(\{\iota = i\} \cap \{\gamma \geq k\} \cap \{\xi = n\}). \quad (2.1)$$

The number q_i is called *player i 's activation probability*. Occasionally, we also write simply q for q_I (and $1 - q$ for q_{II}). We call the marginal of p on ξ *offspring distribution*.

Define, for $x \in [0, 1]$,

$$G_i(k, x) =: \mathbb{E}_p(\mathbf{1}_{\{\iota=i\} \cap \{\gamma \geq k\}} x^\xi) = \sum_{n \in \mathbb{N}} p_i(k, n) x^n, \quad (2.2)$$

where \mathbb{E}_p is the expectation with respect to the measure p . In particular, $G_i(k, 0) = p_i(k, 0)$. The functions G_I and G_{II} completely characterize the probability measure p . Note that $G_i(k, x)$ is the generating function for the sequence $\{p_i(k, n)\}_{n \in \mathbb{N}}$.

At this point, we allow for any joint distribution of ι, γ , and ξ . In the sequel, however, we attend closely to two classes of models.

The first special case of importance is the escape model. We say that p is an *escape model* if $p(\{\gamma > 0\} \cap \{\xi = 0\}) = 0$ and $p(\{\gamma < k\}) > 0$ for each $k > 0$. The first condition imposes that all childless nodes of the game tree have zero capacity. By virtue of this condition, player I 's primary objective in the escape model is making sure that the game never stops for reaching an end node results in the lowest possible payoff: zero. The second condition (which could be equivalently stated as saying that the essential infimum of γ is zero) is, in a sense, without loss of generality. If $p(\{\xi = 0\}) > 0$, it is implied by the first condition. And, if $p(\{\xi = 0\}) = 0$, one could redefine the capacity by subtracting its essential infimum, obtaining a strategically equivalent specification.

The second important class of models are activation-independent models. We say that p is *activation-independent* if the random variable (γ, ξ) is independent of ι . Note that an activation-independent model still permits any joint distribution of γ and ξ . Activation-independent models lend themselves to the study of comparative statics

(of the distribution of the value) with respect to the activation probability of player I. To study comparative statistics, we fix a marginal of p on (γ, ξ) and think of p as being parameterized by its marginal on ι , that is, by q .

2.2. Trees

Let $H = \cup_{n \in \mathbb{N}} \mathbb{N}_+^n$ denote the set of finite sequences of positive natural numbers, including the empty sequence \emptyset . The length of a sequence $h \in H$ is the number $n \in \mathbb{N}$ such that $h \in \mathbb{N}_+^n$; in particular, \emptyset has length zero. For a sequence $h = (j_0, \dots, j_n)$ of length $n + 1$, the sequences $\emptyset, (j_0), (j_0, j_1), \dots, (j_0, \dots, j_n)$ are said to be the prefixes of h .

A tree is a subset T of H containing \emptyset such that, whenever a nonempty sequence is an element of T , all its prefixes are elements of T as well. A tree T is ordered if, for each $h \in T$, there is a natural number $\xi(h) \in \mathbb{N}$ such that, for $j \in \mathbb{N}_+$, the sequence (h, j) is an element of T if and only if $1 \leq j \leq \xi(h)$. Thus, $\xi(h)$ is the number of children of the node h in T ; if $\xi(h) = 0$, then h is an end node of T . Note that an ordered tree is finitely branching.

An infinite branch of T is an infinite sequence $(j_0, j_1, \dots) \in \mathbb{N}_+^{\mathbb{N}}$ such that $(j_0, \dots, j_t) \in T$ for every $t \in \mathbb{N}_+$.

2.3. Games

We formally define a game as a particular type of a marked ordered tree, in which the markings on a tree represent the active player and the capacity. The formalism of marked trees (*arbre marqué*) is borrowed from Neveu [25].

A game ω is a triple $(T_\omega, \iota_\omega, \gamma_\omega)$, where T_ω is an ordered tree, and ι_ω and γ_ω are functions on T_ω with values in $\{I, II\}$ and in $[0, \infty)$, respectively. The three elements of the game are, respectively, the *game tree*, the assignment of the *active player* to the nodes of the game tree, and the assignment of the *capacity* to the nodes of the game tree. For $h \in T_\omega$, we let $\xi_\omega(h)$ denote the number of children of the node h in the tree T_ω . Let Ω be the set of games. Given a game $\omega \in \Omega$ and a node $h \in T_\omega$, one defines a subgame $\omega(h)$ of ω starting at the node h . Formally, $\omega(h) =: (s^{-1}[T_\omega], \iota_\omega \circ s, \gamma_\omega \circ s)$, where $s: H \rightarrow H$ is a shift operator given by $s(h') = (h, h')$.

We endow Ω with a topology generated by the subbase consisting of sets of the form $\{\omega \in \Omega : h \notin T_\omega\}$ and $\{\omega \in \Omega : h \in T_\omega, \iota_\omega(h) = i, k_0 < \gamma_\omega(h) < k_1\}$, where $h \in H, i \in \{I, II\}$, and k_0 and k_1 are rational numbers. The space Ω is Polish.

Note that the map $(\iota(\emptyset), \gamma(\emptyset), \xi(\emptyset)) : \Omega \rightarrow S$ given by $\omega \mapsto (\iota_\omega(\emptyset), \gamma_\omega(\emptyset), \xi_\omega(\emptyset))$ is a continuous function. The map $\omega \rightarrow \omega(1)$ is a continuous map from $\{\omega \in \Omega : \xi_\omega(\emptyset) \geq 1\}$, a clopen subset of Ω , into Ω . For $(i, k, n) \in S$, let

$$E(i, k, n) = \{\omega \in \Omega : \iota_\omega(\emptyset) = i, \gamma_\omega(\emptyset) \geq k, \xi_\omega(\emptyset) = n\}. \quad (2.3)$$

As is usual, we drop the subscript ω from our notation whenever this does not lead to a confusion.

2.4. A Measure on the Space of Games

The following result is essentially due to Neveu [25]: there exists a unique measure $\mathbb{P} = \mathbb{P}_p$ on the Borel subsets of Ω satisfying the following two conditions:

1. The random variable $(\iota(\emptyset), \gamma(\emptyset), \xi(\emptyset))$ is distributed according to p .
2. For each $(i, k, n) \in S$, where $n \in \mathbb{N}_+$, if $p_i(k, n) > 0$, then the random variables $\omega(1), \dots, \omega(n)$ are independent under the conditional measure $\mathbb{P}(\cdot | E(i, k, n))$, and each is distributed according to \mathbb{P} .

The measure \mathbb{P} on the space of games is a natural generalization of the Galton–Watson measure on the space of ordered trees. In fact, the marginal of \mathbb{P} on T_ω is a Galton–Watson measure.

2.5. How a Game Is Played

Consider a game $\omega \in \Omega$. The play in ω starts at the root \emptyset . Suppose that, at some stage of the game, a node h of T_ω is reached. If h is an end node, the game ends. Otherwise, the active player $\iota_\omega(h)$ chooses one of the children $j \in \{1, \dots, \xi_\omega(h)\}$ of h , and the node (h, j) of T_ω is reached at the next stage. Thus, a play of the game either leads to an end node of T_ω or induces an infinite branch of T_ω . Let $h_0 = \emptyset, h_1, h_2, \dots$ be a sequence, finite or infinite, of nodes of T_ω successively visited in the course of the game. The payoff to player I is defined as $\inf \{\gamma_\omega(h_0), \gamma_\omega(h_1), \dots\}$. Player I's goal is to maximize the payoff, and player II's goal is to minimize it.

One can view the payoff function in the game ω as defined on the Baire space $\mathbb{N}_+^{\mathbb{N}}$, the space of infinite sequences of positive natural numbers, equipped with the product topology. Under this interpretation, the payoff is a bounded upper semicontinuous function.

2.6. Strategies

A strategy for player I in the game ω is a function σ_I that assigns a number $\sigma_I(h) \in \{1, \dots, \xi_\omega(h)\}$ to each node $h \in T_\omega$ with $\iota_\omega(h) = I$ and $\xi_\omega(h) \geq 1$. The interpretation is that, at node h , the strategy σ_I recommends player I to move to the child $(h, \sigma_I(h))$ of h . A strategy σ_{II} for player 2 is defined in a similar way. A pair of strategies (σ_I, σ_{II}) either

leads to an end node of T_ω eventually or it induces an infinite branch of T_ω . Let $u_\omega(\sigma_I, \sigma_{II})$ denote the corresponding payoff.

2.7. The Value and k -optimal Strategies

As a consequence of Martin's ([20]) determinacy theorem, the game ω has a value

$$v_\omega =: \sup_{\sigma_I} \inf_{\sigma_{II}} u_\omega(\sigma_I, \sigma_{II}) = \inf_{\sigma_{II}} \sup_{\sigma_I} u_\omega(\sigma_I, \sigma_{II}).$$

The value is the highest payoff that player I can guarantee to receive; at the same time, it is the lowest payoff that player II can force upon player I.

Let $k \geq 0$. We say that player I's strategy σ_I is k -optimal in ω if $u_\omega(\sigma_I, \sigma_{II}) \geq k$ for each player II's strategy σ_{II} in ω . In each game $\omega \in \Omega$, player I has a v_ω -optimal strategy.² Likewise, player II's strategy σ_{II} is said to be k -optimal in ω if $u_\omega(\sigma_I, \sigma_{II}) \leq k$ for each player I's strategy σ_I . Player II might not have a v_ω -optimal strategy, but he does have a $(v_\omega + \epsilon)$ -optimal strategy for each $\epsilon > 0$. If the capacity is supported on a finite set, then player II also has a v_ω -optimal strategy (see, e.g., Laraki et al. [18]).

For $h \in T_\omega$, we write $v_\omega(h)$ to denote $v_{\omega(h)}$, the value of the subgame $\omega(h)$ of ω .

3. Two Examples

This section introduces two examples that serve as the illustration throughout the paper. The data reported here is based on the main results derived in the following section (see Section 4.3). We hope to convey the flavor of our main findings and help the reader anticipate the developments in the rest of the paper.

Note that both examples are instances of an activation-independent escape model.

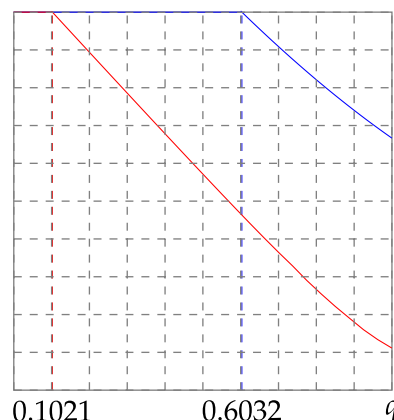
Example 3.1. Suppose that the number of children ξ has the geometric distribution $p(\xi = n) = (1-l)^n$, where $0 < l < 1$. Recall that $E_p(\xi) = l/(1-l)$.

To complete the description of the example, assume that ι and ξ are independent and $\gamma = 1$ whenever $\xi > 0$ and $\gamma = 0$ if $\xi = 0$. This specification of the capacity is a particularly important special case of the model. Under this specification, the sole goal of player I is to avoid end nodes of the tree. The value of the game is either zero or one, and it is zero precisely when player II can force a play to reach an end node. We say that player I wins the game if $v = 1$, and that player II wins the game if $v = 0$. We let α denote the probability of the event $\{v = 0\}$ and by β the probability of $\{v = 1\}$.

Consider first the nonadversarial scenario, the scenario under which all nodes are assigned to player I (i.e., $q = 1$). Of course, in the nonadversarial case, $v = 0$ precisely when the game tree T_ω has no infinite branches, that is, in the event of "extinction" of the game tree. Because, under the measure \mathbb{P} , the distribution of the game tree is a Galton–Watson measure, classical results on the branching processes apply (e.g., Athreya and Ney [4]): the probability of extinction of the game tree is $\alpha = 1$ if $l/(1-l) \leq 1$ and is $\alpha = (1-l)/l$ otherwise.

Turning to the general case, Figure 1 depicts α as a function of q for $l = 0.6$ (blue) and $l = 0.9$ (red). Predictably, α is decreasing in q with $q = 1$ corresponding to the nonadversarial scenario. The figure reveals that, if player I controls too few nodes, the player has no chance of winning the game. The critical activation probability q_c for player I is 0.6032 for $l = 0.6$ and 0.1021 for $l = 0.9$.

Figure 1. (Color online) The probability $\alpha = \mathbb{P}(v = 0)$ in Example 3.1 as a function of q for $l = 0.6$ (blue) and $l = 0.9$ (red).



We find that $\alpha < 1$ (or, equivalently, that $\beta > 0$) if and only if $q > q_c$, where the critical activation probability is given by³

$$q_c = \frac{(1-l)(1-l+l^2)}{l^2(2-l)}.$$

Whenever $q > q_c$, the probability of $\{v = 0\}$ is given by

$$\alpha = \frac{1}{2}(2-l)(1-q) + \frac{1}{2}\sqrt{4\frac{(1-l)^2}{l^2} + (2-l)^2(1-q)^2}.$$

Example 3.2. Let $n \geq 2$ and suppose that, under the measure p , $\xi = n$ almost surely (so that the game tree T_ω is a complete n -ary tree), that γ and ι are independent, and that γ is uniformly distributed on $[0, 1]$. Clearly, the value v_ω of any game ω is an element of $[0, 1]$.

First, consider the nonadversarial scenario ($q = 1$). In the nonadversarial scenario, the probability that $\{v_\omega \geq k\}$ is positive if and only if $(1-k)n > 1$. The rationale behind this fact is as follows. If player I controls all nodes, then $v_\omega \geq k$ if and only if the tree T_ω has an infinite branch that passes only through the nodes with a capacity no smaller than k . Thus, we may remove all the children of any node with a capacity smaller than k and let T_ω^k be an infinite component of the root of the remaining subgraph. Then, $\{v_\omega \geq k\}$ is exactly the event of nonextinction of the tree T_ω^k . One can see that the tree T_ω^k is distributed according to a Galton–Watson measure generated by an offspring distribution with mean $(1-k)n$. We may, thus, conclude that T_ω^k has an infinite branch with positive probability if and only if $(1-k)n > 1$.

Turning to the general case, we find that, for $k > 0$, the probability that $\{v \geq k\}$ is positive if and only if $(1-k)nq > 1$. In particular, if $q \leq 1/n$, then $v = 0$ almost surely: if player I controls too few nodes, then player II is able to force the play to go through the nodes with a vanishingly small capacity, thus ensuring the payoff of zero.

Let us define the k -critical activation probability for player I to be $q_c(k) = ((1-k)n)^{-1}$ if $(1-k)n > 1$ and to be one otherwise. Then, $\mathbb{P}(v \geq k) > 0$ if and only if $q > q_c(k)$. For the binary and ternary trees, the probability of the event $\{v < k\}$ allows for a simple closed-form solution.⁴ For the binary tree, we have

$$\mathbb{P}(v < k) = \begin{cases} 1 & \text{if } 2(1-k)q \leq 1 \\ \frac{k}{(1-k)} \frac{1}{(2q-1)} & \text{if } 1 < 2(1-k)q. \end{cases} \quad (3.1)$$

For the ternary tree,

$$\mathbb{P}(v < k) = \begin{cases} 1 & \text{if } 3(1-k)q \leq 1 \\ \frac{1}{2}(2-3q) + \frac{1}{2}\sqrt{(2-3q)^2 + 4\frac{k}{(1-k)}} & \text{if } 1 < 3(1-k)q. \end{cases} \quad (3.2)$$


The left panel of Figure 2 displays $\mathbb{P}(v < 0.05)$ for the binary tree (in blue) and the ternary tree (in red) as a function of q . Note, in particular, the 0.05-critical probability: $q_c(0.05) = 0.5263$ in the case of the binary tree and $q_c(0.05) = 0.3509$ for the ternary. It emerges from the figure that enlarging the tree does not necessarily benefit player I: if $q = 0.7$, then the probability that $v < 0.05$ goes up from 0.1316 for the binary tree to 0.1848 for the ternary.

The right panel of Figure 2 depicts $\mathbb{P}(v < k)$ as a function of k (that is, the cumulative distribution function of the value) for the binary tree (blue) and the ternary tree (red), assuming that $q = 0.7$.

4. The Distribution of the Value

In this section, we derive the main result of the paper: a fixed-point characterization of the distribution of the value. In the first two subsections, we introduce our tools: a family of the so-called value-generating functions and the truncated games. The last subsection states the main results.

The value-generating functions reflect the recursive nature of the value. Effectively, they represent the Shapley operator. Intuitively, they map the distribution of the value in the next period to that in the current period. The t th iterate of the value-generating functions at zero determines the distribution of the value in the t -truncated game (Lemma 4.4). This and the fact that the value of the truncated games converge (from above) to the value of the infinite game eventually yield the fixed-point characterization of the distribution of the value (Theorem 4.1).

Figure 2. (Color online) The probability $\mathbb{P}(v < k)$ in Example 3.2.

Notes. (a) The probability of $\{v < 0.05\}$ for the binary (blue) and ternary (red) trees as a function of q . (b) The cumulative distribution of v for the binary (blue) and the ternary (red) trees, $q = 0.7$.

4.1. The Value-Generating Function

Define the *value generating function* (vgf) $f : [0, \infty) \times [0, 1] \rightarrow [0, 1]$ associated with the primitive distribution p by

$$f(k, x) = 1 - G_I(k, 0) - G_I(k, 1) + G_I(k, x) - G_{II}(k, 1 - x). \quad (4.1)$$

We write $f_k : [0, 1] \rightarrow [0, 1]$ to denote the function $x \mapsto f(k, x)$.

Define

$$d(k) = \mathbb{E}_p(\mathbf{1}_{\{\iota=I\} \cap \{\gamma \geq k\}} \xi) + p(\{\iota = II\} \cap \{\gamma \geq k\} \cap \{\xi = 1\}). \quad (4.2)$$

This quantity turns out to be one of the key parameters of our model.

The following lemma summarizes the relevant properties of the vgf.

Lemma 4.1. *Let $k \geq 0$.*

- i. *The function f_k is continuous and nondecreasing on $[0, 1]$. It is differentiable any number of times at any point of $(0, 1)$.*
- ii. *It holds that*

$$\begin{aligned} f_k(0) &= p(\{\gamma < k\}), \\ f_k(1) &= p(\{\gamma < k\} \cup \{\xi \geq 1\}), \\ \lim_{x \uparrow 1} \frac{\partial f_k}{\partial x}(x) &= d(k). \end{aligned}$$

- iii. *There exists a point $c \in [0, 1]$ such that f_k is concave on $[0, c]$ and convex on $[c, 1]$.*

iv. *Suppose that f_k is not the identity map. If $s \in (0, 1)$ is a fixed point of f_k , then $x < f_k(x)$ for each $x \in (0, s)$ and $f_k(x) < x$ for each $x \in (s, 1)$. In particular, f_k has at most one fixed point in $(0, 1)$.*

v. *The function f_k has no fixed point in $[0, 1)$ if and only if all of the following three conditions are satisfied: $f_k(0) > 0$, $f_k(1) = 1$, and $d(k) \leq 1$.*

Proof. We write $f = f_k$, and $G_i(x)$ for $G_i(k, x)$.

Claim (i). The function G_i is a generating function for the sequence $\{p_i(k, n)\}_{n \in \mathbb{N}}$. It follows that (Grimmett and Stirzaker [14, section 5.1]) both these functions are continuous on $[0, 1]$, and they may be differentiated term by term any number of times at any point $x \in (0, 1)$. It holds that $f'(x) = G'_I(x) + G'_{II}(1 - x) \geq 0$ for each $x \in (0, 1)$.

Claim (ii). This follows by a direct computation.

Claim (iii). First, we argue that either f is linear or it has at most one inflection point in $(0, 1)$, that is, only one point $x \in (0, 1)$ such that $f''(x) = 0$.

For each $x \in (0, 1)$, it holds that $f'''(x) = G_I'''(x) + G_{II}'''(1 - x)$, where $G_i'''(x) \geq 0$ for both $i \in \{I, II\}$. In particular, f'' is nondecreasing on $(0, 1)$. Hence, either f has at most one inflection point in $(0, 1)$ or there exists a nondegenerate interval, say (c_0, c_1) , on which f'' and, hence, also f''' vanish. In the latter case, $G_I'''(x) = 0$ and $G_{II}'''(1 - x) = 0$ for each $x \in (c_0, c_1)$. It then follows that $p(\{\gamma \geq k\} \cap \{\xi \geq 3\}) = 0$. But, in this case, f is either linear or quadratic, and it, therefore, has no inflection points, proving the assertion.

If f is a linear function, we can set $c = 0$. Now, suppose that f has at most one inflection point in $(0, 1)$. Define c to be zero if $f''(x) > 0$ for all $x \in (0, 1)$, to be one if $f''(x) < 0$ for all $x \in (0, 1)$, and otherwise to be the unique point of $(0, 1)$ such that $f''(c) = 0$. Recalling that f'' is nondecreasing, we conclude that f is concave on $[0, c]$ and is convex on $[c, 1]$.

Claim (iv). The claim is trivial if f is a linear function. So suppose that f is not linear. By the earlier conclusion, f has at most one inflection point in $(0, 1)$. We distinguish two cases.

Case 1. $c \leq s$. We know that $f(s) = s$ and $f(1) \leq 1$ and that f is convex on $[s, 1]$. This implies that $f(x) \leq x$ for each $x \in [s, 1]$. Moreover, if there existed a point $x \in (s, 1)$ such that $f(x) = x$, the function f would be the identity map on $[s, 1]$, implying a continuum of inflection points, a contradiction.

We also have $x < f(x)$ for each $x \in [c, s)$. For, otherwise, convexity of f on $[c, 1]$ would imply that f is the identity map on $[s, 1]$, contradicting the fact that it has at most one inflection point in $(0, 1)$. We also know that $0 \leq f(0)$ and $c < f(c)$. Concavity of f on $[0, c]$ now implies that $x < f(x)$ for each $x \in (0, c]$.

Case 2. $s \leq c$. One applies a similar reasoning to the intervals $[0, s]$, $[s, c]$, and $[c, 1]$.

Claim (v). (Necessity). Suppose that f has no fixed point in $[0, 1)$. Obviously, then, $0 < f(0)$. If there existed a point $x \in [0, 1]$ such that $f(x) < x$, the intermediate value theorem would imply that f has a fixed point in $(f(0), f(x))$, contradicting the supposition. Thus, $x \leq f(x)$ for each $x \in [0, 1]$. In particular, $f(1) = 1$. If $c = 1$, then f is concave on $[0, 1]$. Then, $xf'(x) \leq f(x) - f(0)$ for each $x \in (0, 1)$. Letting $x \uparrow 1$, we obtain $d(k) \leq 1$. If $c < 1$, then f is convex on $[c, 1]$. In this case, $f'(x)(1 - x) \leq f(1) - f(x) \leq 1 - x$ for each $x \in (c, 1)$, and thus, $d(k) \leq 1$.

Claim (vi). (Sufficiency). Suppose that $0 < f(0)$, $f(1) = 1$, and $d(k) \leq 1$. Toward a contradiction, let $s \in (0, 1)$ be a fixed point of f . Clearly, this means that f cannot be concave on $[0, 1]$; hence, $c < 1$, and so f is convex on $[c, 1]$. On the other hand, by Claim (iv), $f(x) < x$ for each $x \in (s, 1)$. But, then, there is a point $x_0 \in (c, 1)$ such that $1 < f'(x_0)$, and because $f'(x_0) \leq f'(x)$ for all $x \in (x_0, 1)$, we obtain $1 < f'(x_0) \leq d(k)$, a contradiction. \square

Example 4.1. In the setup of Example 3.1, we have

$$G_i(1, x) = q_i \frac{(1-l)x}{1-lx}, \tag{4.3}$$

from which we derive the vgf:

$$f_1(x) = 1 - l(1-x) \left(q_I \frac{1}{1-lx} + q_{II} \frac{1-l}{1-l+lx} \right).$$

The function is displayed in the left panel of Figure 3 for $l = 0.9$ and $q = 0.5$ (blue) and for $l = 0.6$ and $q = 0.5$ (red). Especially, the first of these (the blue) displays a feature characteristic of all vgf's: it is concave on an interval $[0, c]$ for some $0 < c < 1$ and convex on $[c, 1]$.

In the setup of Example 3.2, the vgf is given by

$$G_i(k, x) = (1-k)q_i x^n, \tag{4.4}$$

and hence,

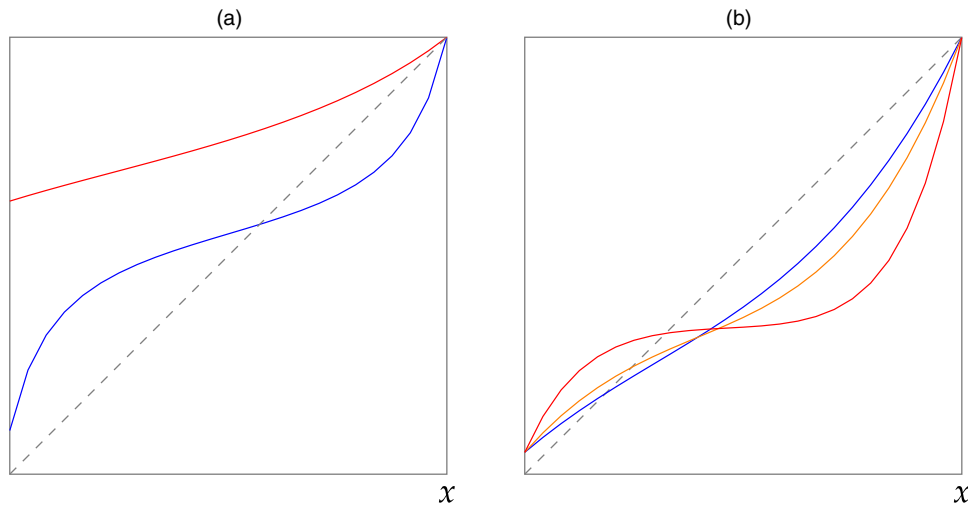
$$f_k(x) = 1 - (1-k)q_I + (1-k)q_I x^n - (1-k)q_{II}(1-x)^n.$$

It is pictured in the right panel of Figure 3 with $k = 0.05$ and $q = 0.7$ for $n = 3$ (blue), $n = 4$ (orange), and $n = 8$ (red).

4.2. The Truncated Game

Given a game ω and time $t \in \mathbb{N}$, we define a version of the game ω played on a game tree truncated at time t . The idea is straightforward: the game ω_t lasts for no more than t periods, and player I's payoff is the smallest capacity along the nodes visited prior to the deadline t . Given a game $\omega = (T_\omega, \iota_\omega, \gamma_\omega)$, define the *truncated game* ω_t to be the triple $(T_{\omega_t}, \iota_{\omega_t}, \gamma_{\omega_t})$, where T_{ω_t} is the subset of nodes of T_ω having the length of at most t , and ι_{ω_t} and γ_{ω_t} are the restrictions of ι_ω and γ_ω , respectively, to T_{ω_t} . In particular, ω_0 has a tree consisting of a single node, namely, the empty sequence \emptyset .

Figure 3. (Color online) vgf.



Notes. (a) The vgf in Example 3.1 for $l = 0.9$ and $q = 0.5$ (blue) and for $l = 0.6$ and $q = 0.5$ (red). (b) The vgf in Example 3.2 with $k = 0.05$ and $q = 0.7$ for $n = 3$ (blue), $n = 4$ (orange), and $n = 8$ (red).

We first establish measurability of the value of a game ω . We remark that the measurability of the value cannot be taken for granted. It is known that the value of a Borel-parameterized infinite perfect information game need not be Borel-measurable (Moschovakis [23]); counterexamples to measurability of the value are also given in other contexts, for instance, for simultaneous move games (Prikry and Sudderth [27]). The following positive result can be linked to the fact that player I’s payoff function is, in each game ω , upper semicontinuous.

Lemma 4.2. Let $t \in \mathbb{N}_+$.

- i. The map $\Omega \rightarrow \Omega$ given by $\omega \mapsto \omega_t$ is continuous.
- ii. The map $v_t : \Omega \rightarrow \mathbb{R}$ defined by $\omega \mapsto v_{\omega_t}$ is continuous.
- iii. For each $\omega \in \Omega$, $v_{\omega_0} \geq v_{\omega_1} \geq \dots$ is a nonincreasing sequence converging to v_ω . Consequently, the map $v : \Omega \rightarrow \mathbb{R}$, $\omega \mapsto v_\omega$ is upper semicontinuous and, hence, measurable.

Proof. Item (i) can be checked directly using the definition of the subbase of the topology on Ω . Item (ii) can be shown by an induction on t . We turn to item (iii).

Let $v_{\omega_t} = k$, and let σ_{II} be player II’s k -optimal strategy in the game ω_t . Note that, ω_t being essentially a finite game, both players have an optimal strategy. The strategy σ_{II} guarantees that, within t periods of time, the play visits a node with a capacity of at most k . Clearly, then, the same strategy is k -optimal in ω_{t+1} . Thus, $v_{\omega_t} \geq v_{\omega_{t+1}}$.

Now, let $v_\omega = k$. Also, let $\epsilon > 0$.

First, we have $v_{\omega_t} \geq k$ for each $t \in \mathbb{N}$ because each k -optimal strategy of player I in ω is also a k -optimal strategy in ω_t .

To see that the sequence v_{ω_t} converges to $v_\omega = k$, fix some $(k + \epsilon)$ -optimal strategy for player II in ω , say σ_{II} . Let $W \subseteq T$ denote a tree consisting of the nodes $h \in T_\omega$ that (i) could be reached when player II is using strategy σ_{II} and (ii) have the property that $\gamma(h') > k + \epsilon$ for each prefix h' of h . Then, the tree W is well-founded, that is, it has no infinite branches. Because it is a locally finite tree, it is actually finite. Let t be the height of W . We conclude that any play of the game ω consistent with σ_{II} reaches a node with a capacity of $k + \epsilon$ or less by period t . But this means that σ_{II} is a $(k + \epsilon)$ -optimal strategy in ω_t . Thus, $v_{\omega_t} \leq k + \epsilon$, and hence, $\inf_{t \in \mathbb{N}} v_{\omega_t} \leq k + \epsilon$. Because $\epsilon > 0$ is arbitrary, we have shown that $\inf_{t \in \mathbb{N}} v_{\omega_t} \leq k$. \square

We turn to the probabilistic properties of the value. Recall that $\omega(h)$ denotes the subgame of the game ω starting at a node $h \in T_\omega$. Recall also the definition of the events $E(i, k, n)$ in (2.3).

Lemma 4.3. Let $t \in \mathbb{N}_+$.

- i. It holds that $(\omega_t)(\ell) = (\omega(\ell))_{t-1}$ for $\omega \in \Omega$ and each $\ell \in \{1, \dots, \xi_\omega(\emptyset)\}$.
- ii. Take $(i, k, n) \in S$ such that $p_i(k, n) > 0$. For $\ell \in \{1, \dots, n\}$ let $v_t(\ell) : E(i, k, n) \rightarrow \mathbb{R}$ denote the random variable $\omega \mapsto v_{(\omega_t)(\ell)}$. Under the conditional measure $\mathbb{P}(\cdot | E(i, k, n))$, the random variables $v_t(1), \dots, v_t(n)$ are independent, and each is distributed like the random variable v_{t-1} .

Proof. Item (i) follows immediately from the definition of the truncated game. Item (ii) follows from item (i) and the fact that, under the conditional measure $\mathbb{P}(\cdot | E(i, k, n))$, the random variables $\omega(1), \dots, \omega(n)$ are independent and each is distributed like the random variable ω . \square

Lemma 4.4. *Take $k > 0$. Let $\alpha_t = \mathbb{P}(v_t < k)$ and $\beta_t = \mathbb{P}(v_t \geq k)$. Then, the sequence $\alpha_0 \leq \alpha_1 \leq \dots$ converges to $\mathbb{P}(v < k)$. For each $t \in \mathbb{N}$, we have $\alpha_t = f_k^{t+1}(0)$, where f_k^{t+1} denotes the $(t+1)$ -fold iterate of f_k .*

Proof. By Lemma 4.2, $\{v_0 < k\} \subseteq \{v_1 < k\} \subseteq \dots$ is a nondecreasing sequence of events converging to $\{v < k\}$. The first claim follows.

Note that $v_t = v_{\omega_t} \leq \gamma_{\omega_t}(\emptyset) = \gamma_{\omega}(\emptyset) < k$ everywhere on the complement of the event $\cup \{E(i, k, n) : i \in \{I, II\}, n \in \mathbb{N}\}$. Consider the event $E(I, k, 0)$. Because $v_t = \gamma_{\omega}(\emptyset) \geq k$ holds everywhere on $E(I, k, 0)$, we have

$$\mathbb{P}(v_t \geq k | E(I, k, 0)) = 1;$$

whenever $E(I, k, 0)$ has a positive probability.

Take $n \geq 1$ and consider the event $E(I, k, n)$. We have

$$E(I, k, n) \cap \{v_t < k\} = E(I, k, n) \cap \bigcap_{j=1}^n \{v_t(j) < k\}.$$

Hence, using Lemma 4.3, we compute

$$\mathbb{P}(v_t < k | E(I, k, n)) = \mathbb{P}(\bigcap_{j=1}^n \{v_t(j) < k\} | E(I, k, n)) = \prod_{\ell=1}^n \mathbb{P}(v_{t-1} < k) = \alpha_{t-1}^n,$$

whenever $E(I, k, n)$ has a positive probability. Thus,

$$\begin{aligned} \mathbb{P}(\{v_t \geq k\} \cap \{t_{\omega}(\emptyset) = I\}) &= \sum_{n \in \mathbb{N}} \mathbb{P}(\{v_t \geq k\} \cap E(I, k, n)) \\ &= \sum_{n \in \mathbb{N}} \mathbb{P}(E(I, k, n)) \mathbb{P}(\{v_t \geq k\} | E(I, k, n)) \\ &= p_I(k, 0) + \sum_{n \in \mathbb{N}} p_I(k, n)(1 - \alpha_{t-1}^n) \\ &= G_I(k, 0) + G_I(k, 1) - G_I(k, \alpha_{t-1}). \end{aligned}$$

Consider the event $E(II, k, 0)$. Because $v_t = \gamma_{\omega}(\emptyset) \geq k$ holds everywhere on $E(II, k, 0)$, we have

$$\mathbb{P}(v_t \geq k | E(II, k, 0)) = 1;$$

whenever $E(II, k, 0)$ has a positive probability.

Let $n \geq 1$ and consider the event $E(II, k, n)$. We have

$$E(II, k, n) \cap \{v_t \geq k\} = E(II, k, n) \cap \bigcap_{j=1}^n \{v_t(j) \geq k\}.$$

Hence, using Lemma 4.3, we obtain

$$\mathbb{P}(v_t \geq k | E(II, k, n)) = \mathbb{P}(\bigcap_{j=1}^n \{v_t(j) \geq k\} | E(II, k, n)) = \prod_{j=1}^n \mathbb{P}(v_{t-1} \geq k) = \beta_{t-1}^n;$$

whenever $E(II, k, n)$ has a positive probability. Thus,

$$\begin{aligned} \mathbb{P}(\{v_t \geq k\} \cap \{t_{\omega}(\emptyset) = II\}) &= \sum_{n \in \mathbb{N}} \mathbb{P}(\{v_t \geq k\} \cap E(II, k, n)) \\ &= \sum_{n \in \mathbb{N}} \mathbb{P}(E(II, k, n)) \mathbb{P}(v_t \geq k | E(II, k, n)) \\ &= \sum_{n \in \mathbb{N}} p_{II}(k, n) \beta_{t-1}^n \\ &= G_{II}(k, \beta_{t-1}). \end{aligned}$$

It now follows that

$$\beta_t = \mathbb{P}(v_t \geq k) = G_I(k, 0) + G_I(k, 1) - G_I(k, \alpha_{t-1}) + G_{II}(k, \beta_{t-1}),$$

and hence, recalling the definition of the vgf (4.1), we obtain $\alpha_t = f_k(\alpha_{t-1})$. Unraveling this recursive relation yields $\alpha_t = f_k^t(\alpha_0)$. Finally, recall that ω_0 is a trivial game, having a single node, \emptyset . Consequently, $v_0 = \gamma_{\omega}(\emptyset)$, and so $\alpha_0 = \mathbb{P}(v_0 < k) = p(\gamma < k) = f_k(0)$, where the last equality is by Lemma 4.1(ii). Thus, $\alpha_t = f_k^{t+1}(0)$. \square

4.3. The Main Results

We are in a position to derive the main result of the paper. For $k > 0$, define $\alpha(k) = \mathbb{P}(v < k)$ and $\beta(k) = \mathbb{P}(v \geq k)$.

Theorem 4.1. *Let $k > 0$. Then, $\alpha(k)$ is the smallest fixed point of the function f_k .*

Proof. The result follows from Lemma 4.4. Let s denote the smallest fixed point of f_k . Because $\alpha_{t+1} = f_k(\alpha_t)$ and because the sequence $\alpha_0, \alpha_1, \dots$ converges to $\alpha(k)$, continuity of f_k implies that $\alpha(k) = f_k(\alpha(k))$ so that $\alpha(k)$ is a fixed point of f_k . Thus, $s \leq \alpha(k)$. On the other hand, because f_k is nondecreasing, we obtain by induction that $f_k^{t+1}(0) \leq s$. Hence, by $\alpha_t \leq s$ for each $t \in \mathbb{N}$ and, therefore, $\alpha(k) \leq s$. \square

Theorem 4.2. *Let $k > 0$. Then, $\beta(k) = 0$ if and only if all three of the following conditions are satisfied:*

$$p(\{\gamma < k\}) > 0, \quad (4.5)$$

$$p(\{\gamma \geq k\} \cap \{\xi = 0\}) = 0, \quad (4.6)$$

$$d(k) \leq 1. \quad (4.7)$$

Theorem 4.2 follows directly from Theorem 4.1 and Lemma 4.1.

That Condition (4.5) is necessary for $\beta(k) = 0$ follows from the fact that, in the event that the root of the tree has no children but a capacity of at least k , then also the value of the game is at least k . And that Condition (4.6) is necessary for $\beta(k) = 0$ follows because, in a game in which the capacity is at least k almost surely, also the value is at least k almost surely.

Condition (4.7) is much more subtle. Recall that $d(k)$ as defined in (4.2) is a sum of two terms: the expectation of the random variable $1_{\{\iota=I\} \cap \{\gamma \geq k\}} \xi$ and the probability of the event $\{\iota = II\} \cap \{\gamma \geq k\} \cap \{\xi = 1\}$. It is perhaps only natural that the first term affects (4.7): intuitively, the higher the expected number of nodes assigned to player I with a capacity of at least k , the easier it is for player I to secure a payoff of k . The second term is more difficult to interpret. Let us suggest one possible explanation: a player assigned to a node with a single child has no real choice of action at that node. Therefore, it is inconsequential who is being assigned to control the nodes having a single child: if one reassigns the nodes with a single child from player II to player I, one obtains a strategically equivalent game.

We revisit Condition (4.7) in the next section. For activation-independent escape models, the condition can be reformulated in terms of player I's activation probability. The discussion of $d(k)$ is further continued in Section 5.5.

Theorem 4.2 subsumes the classical criterion for the (non-)extinction of a branching process. To see this, suppose that $\iota = I$, that $\gamma = 1$ whenever $\xi \geq 1$, and that $\gamma = 0$ if $\xi = 0$ almost surely under the primitive distribution p . Then, the value of a game ω is either one or zero, depending on whether the game tree T_ω has an infinite branch or not. Taking $k = 1$, Theorem 4.2 reads $\mathbb{P}(v = 1) > 0$ if and only if $p(\{\xi = 0\}) = 0$ or $\mathbb{E}_p(\xi) > 1$. This can be easily seen to be equivalent to the classic condition: $p(\{\xi = 1\}) = 1$ or $\mathbb{E}_p(\xi) > 1$.

5. Corollaries

We explore several features of the distribution of the value.

When discussing the examples of Section 3, we have noted that the probability of the event $\{v \geq k\}$ is positive only if player I's activation probability is larger than a certain critical level. Subsection 5.1 derives an expression for the k -critical level of player I's activation probability in the context of an activation-independent escape model.

Subsection 5.2 introduces into our study the distribution of the value conditional on the active player. In any activation-independent model, the distribution of the value at player I's nodes first order stochastically dominates that at player II's nodes, and both are nondecreasing (in the sense of first order stochastic dominance) with respect to player I's activation probability. We also take a close look at the (conditional) probability of the event $\{v \geq k\}$ as player I's activation probability approaches its k -critical value.

Subsection 5.3 discusses an asymptotic result for games defined on complete n -ary trees as n becomes large.

Subsection 5.4 introduces the notion of a k -optimal subtree and discusses its relation to the set of player I's k -optimal strategies.

In subsection 5.5, we study the so-called avoidance game, the game in which player I is prohibited from visiting the nodes controlled by player II with the exception of the nodes having no children or a single child. In the avoidance game, player II is, thus, deprived of any real choice and is merely a dummy. As the avoidance game incorporates an additional restriction on player I's moves, it is harder for player I to play. We show nonetheless that, if player I is able to secure a payoff of at least k in the original game, then she is able to do so

in the avoidance game. The avoidance game also leads to an interesting interpretation of the key parameter $d(k)$.

For the most part, we omit (rather straightforward) proofs. They can be found, along with other results, in our working paper (Flesch et al. [8]).

5.1. Critical Activation Probability

Because an escape model satisfies Conditions (4.5) and (4.6) of Theorem 4.2 for any $k > 0$, we obtain the following:

Corollary 5.1. *If p is an escape model, then $\beta(k) = 0$ if and only if $d(k) \leq 1$.*

We already noted an interesting feature of the two examples of Section 3: the probability of the event $\{v \geq k\}$ is only positive if player I's activation probability is above a certain critical level (see Figures 1 and 2). Here, we give a general expression for player I's k -critical activation probability for an activation-independent escape model.

If p satisfies activation-independence, $d(k)$ can be rewritten as

$$d(k) = q \cdot \mathbb{E}_p(1_{\{\gamma \geq k\}} \xi) + (1 - q) \cdot p(\{\gamma \geq k\} \cap \{\xi = 1\}).$$

Solving the inequality $d(k) \leq 1$ for q , we obtain $q \leq q_c(k)$, where $q_c(k)$, player I's k -critical activation probability, is defined by the following expression:

$$q_c(k) = \begin{cases} 0 & \text{if } \mathbb{E}_p(1_{\{\gamma \geq k\}} \xi) = \infty, \\ \frac{1 - p(\{\gamma \geq k\} \cap \{\xi = 1\})}{\mathbb{E}_p(1_{\{\gamma \geq k\}} \xi) - p(\{\gamma \geq k\} \cap \{\xi = 1\})} & \text{if } 1 < \mathbb{E}_p(1_{\{\gamma \geq k\}} \xi) < \infty, \\ 1 & \text{otherwise.} \end{cases} \quad (5.1)$$

We summarize the discussion as follows:

Corollary 5.2. *If p is an activation-independent escape model, then $\beta(k) = 0$ if and only if $q \leq q_c(k)$, where $q_c(k)$ is given by (5.1).*

A further interpretation of the two terms appearing in (5.1), namely, $\mathbb{E}_p(1_{\{\gamma \geq k\}} \xi)$ and $p(\{\gamma \geq k\} \cap \{\xi = 1\})$, is offered by Corollary 5.3.

5.2. Distribution of the Value Conditional on the Active Player

Apart from $\alpha(k)$ and $\beta(k)$, the probabilities of the events $\{v < k\}$ and $\{v \geq k\}$, one might also be interested in the probabilities of these events conditional on the root of the tree \emptyset being assigned to player I or player II. For $i \in \{I, II\}$, assuming that $q_i > 0$, let $\alpha_i(k) = \mathbb{P}(v < k \mid \iota(\emptyset) = i)$ and $\beta_i(k) = \mathbb{P}(v \geq k \mid \iota(\emptyset) = i)$. Of course, the conditional probabilities are related to the unconditional ones by $\alpha(k) = q_I \alpha_I(k) + q_{II} \alpha_{II}(k)$ and likewise for $\beta(k)$.

For the rest of this section, fix a $k > 0$. We suppress the dependence on k whenever convenient, writing, for example, α , β , α_i , β_i , and q_c in place of $\alpha(k)$, etc.

Recall that, in an activation-independent model, the random variables (γ, ξ) and ι are independent. For such models, a very intuitive result holds: the value tends to be higher at nodes controlled by player I than at nodes controlled by II. The precise statement is probabilistic: the conditional distribution of the value given the event $\{\iota(\emptyset) = I\}$ first order stochastically dominates that given $\{\iota(\emptyset) = II\}$. And, second, the value is a "nondecreasing" function of player I's activation probability, again in the sense of first order stochastic dominance.

A further insight can be gained by focusing on activation-independent escape models. For such models, we can compute the limits of the ratios β_i/β as player I's activation probability q approaches the k -critical level. These limits allow one to estimate how much larger the conditional probability β_I is compared with β_{II} , assuming that player I controls just enough nodes for the probability β to be positive. As it turns out, the ratios are determined by two quantities: $\mathbb{E}_p(1_{\{\gamma \geq k\}} \xi)$ and $p(\{\gamma \geq k\} \cap \{\xi = 1\})$, the same two that enter into the expression for the k -critical activation probability (5.1).

Corollary 5.3. *Suppose that p is an activation-independent model. Then,*

- i. $\alpha_I \leq \alpha \leq \alpha_{II}$ and the probabilities α_I , α , and α_{II} are nonincreasing functions of player I's activation probability q on $[0, 1]$.
 - ii. $\beta_I \geq \beta \geq \beta_{II}$ and the probabilities β_I , β , and β_{II} are nondecreasing functions of player I's activation probability q on $[0, 1]$.
- Suppose that p is an activation-independent escape model such that $\mathbb{E}_p(1_{\{\gamma \geq k\}} \xi) > 1$. Then,
- iii. The ratio β_I/β is nonincreasing, and the ratio β_{II}/β is nondecreasing as functions of player I's activation probability q on $(q_c, 1)$.

iv. As $q \downarrow q_c$, α approaches 1 and β approaches 0. Moreover,

$$\lim_{q \downarrow q_c} \frac{\beta_I}{\beta} = \mathbf{E}_p(\mathbf{1}_{\{\gamma \geq k\}} \xi), \tag{5.2}$$

$$\lim_{q \downarrow q_c} \frac{\beta_{II}}{\beta} = p(\{\gamma \geq k\} \cap \{\xi = 1\}). \tag{5.3}$$

Example 5.1. Let us start with a numerical illustration. Consider the setup of Example 3.1 and let $l = 0.9$. The expected number of children is then $\mathbf{E}_p(\xi) = 9$, whereas the probability for a node to have a single child is $p(\xi = 1) = 0.09$. We may, thus, expect that, as q approaches the 1-critical level $q_c = 0.1021$, the probability for player I to win at the own node (i.e., at the node controlled by player I) is approximately 100 times larger than that at a node controlled by player II. And, indeed, we find that, for $q = 0.11$, the ratio β_I/β_{II} is approximately 92.

We have

$$\alpha_I = \frac{1-l}{1-l\alpha}, \quad \beta_I = \frac{l\beta}{1-l+l\beta},$$

$$\alpha_{II} = \frac{(1-l)^2 + (2l-l^2)\alpha}{1-l+l\alpha}, \quad \beta_{II} = \frac{l(1-l)\beta}{1-l\beta}.$$

The probabilities α , α_I , and α_{II} as functions of q are pictured in the left panel of Figure 4. In accordance with the preceding corollary, the α is the middle line (the same as that in Figure 1), α_{II} is the top line, and α_I is the bottom line. As q approaches the 1-critical level (see Example 3.1 for an explicit expression), β_I/β converges to $l/(1-l)$, the expected number of children, whereas β_{II}/β converges to $l(1-l)$, the probability for a node to have exactly one child.

As q approaches one, both α and α_I converge to the same limit, the probability of extinction of the game tree (i.e., that the game tree is finite), in our case $1/9$. On the other hand, α_{II} converges to $3/5$. This is the probability that the root has no children or that, for at least one child of the root, the corresponding subtree is finite.

Example 5.2. Consider now Example 3.2. We have

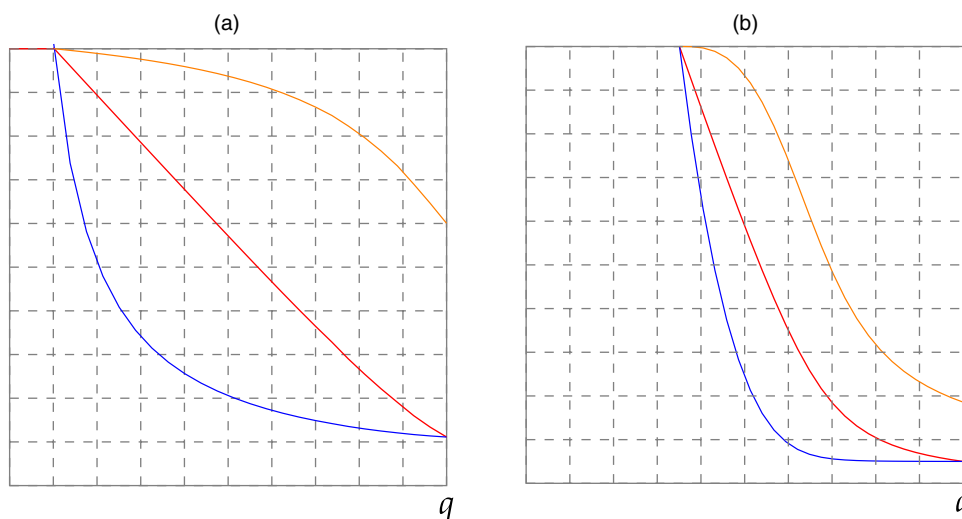
$$\alpha_I = 1 - (1-k) + (1-k)\alpha^n, \quad \beta_I = (1-k) - (1-k)(1-\beta)^n,$$

$$\alpha_{II} = 1 - (1-k)(1-\alpha)^n, \quad \beta_{II} = (1-k)\beta^n.$$

See the right panel of Figure 4 in which α , α_I , and α_{II} are pictured for the ternary tree.

As q approaches the k -critical level $q_c = 1/(1-k)n$, the ratio β_I/β approaches $(1-k)n$, whereas β_{II}/β approaches zero. We interpret the latter fact as follows: if player I’s activation probability is just above the k -critical level, she can only guarantee a payoff of at least k at her own nodes; the probability for player I to secure a payoff of at least k starting at a node controlled by player II is negligible.

Figure 4. (Color online) The probabilities α (red), α_I (blue), and α_{II} (orange) as functions of q . (a) Example 3.1 with $l = 0.9$. (b) Example 3.2 with $n = 3$ and $k = 0.05$.



5.3. An Asymptotic Result for Games on Complete n -ary Trees

We consider a special case of the model in which (as in Example 3.2) the game tree is the complete n -ary tree. Fixing a particular joint distribution of the active player and the capacity, we study the probability of the event $\{v < k\}$ as n becomes large.

Corollary 5.4. *Consider a sequence $p = p_0, p_1, \dots$ of probability measures on S such that the marginal of p_n on (ι, γ) does not depend on n and $p_n(\xi = n) = 1$. Let $k > 0$ be such that $0 < p(\{\gamma < k\})$. Denote $p(\{\gamma \geq k\} \cap \{\iota = i\})$ by ρ_i .*

- *The sequence $\mathbb{P}_{p_n}(v < k)$ is eventually monotone. It is*
 - A. *Constant if $\rho_I = 0$.*
 - B. *Eventually increasing if $\frac{1}{2} < \rho_I$ and $0 < \rho_{II}$.*
 - C. *Eventually decreasing otherwise.*
- *It converges to $1 - \rho_I$.*

Proof. The vgf corresponding to p_n is given by

$$f_n(x) = 1 - \rho_I(1 - x^n) - \rho_{II}(1 - x)^n. \quad (5.4)$$

Let us write x_n for $\mathbb{P}_{p_n}(v < k)$. Recall that x_n is the smallest fixed point of the function f_n .

A. If $\rho_I = 0$, then Theorem 4.2 implies that $\mathbb{P}_{p_n}(v \geq k) = 0$ for each $n \in \mathbb{N}$, and we are done.

For the rest of the proof, we assume that $\rho_I > 0$. Because $\rho_I > 0$, Condition (4.7) of Theorem 4.2 is violated for $n \in \mathbb{N}$ sufficiently large as the left-hand side of the inequality is $d(k) = n\rho_I$. Consequently, $x_n < 1$ for $n \in \mathbb{N}$ sufficiently large. Moreover, it holds that $0 < x_n$ for each $n \in \mathbb{N}$ since $0 < p(\{\gamma < k\}) = f_n(0) \leq f_n(x_n) = x_n$. We find that

$$0 < x_n < 1 \text{ for } n \in \mathbb{N} \text{ sufficiently large.}$$

We use this fact together with Lemma 4.1(iv) repeatedly.

Noting that $\rho_I + \rho_{II} = p(\{\gamma \geq k\}) < 1$, we can subdivide (C) into three subcases as follows: (C1) $0 < \rho_I$ and $0 = \rho_{II}$, (C2) $0 < \rho_I < \frac{1}{2}$ and $0 < \rho_{II}$, (C3) $\frac{1}{2} = \rho_I$ and $0 < \rho_{II}$.

C1. Because $f_n(x) = 1 - \rho_I(1 - x^n)$, we find that $f_{n+1}(x_n) < f_n(x_n) = x_n$, implying that $x_{n+1} < x_n$. We conclude that the sequence $\{x_n\}_{n \in \mathbb{N}}$ is eventually decreasing. Hence, $(x_n)^n \rightarrow 0$. Taking the limit of $x_n = f_n(x_n)$, we find that $x_n \rightarrow 1 - \rho_I$.

For the rest of the proof, we assume that $\rho_{II} > 0$. Define

$$y_n = 1 - \frac{1}{1 + \left(\frac{\rho_{II}}{\rho_I}\right)^{\frac{1}{n-1}}}.$$

As is easy to check, the equation $f_n(x) = f_{n+1}(x)$ admits exactly three solutions in $[0, 1]$, namely, zero, y_n , and one. Moreover,

$$f_n(x) < f_{n+1}(x) \text{ for } x \in (0, y_n), \quad (5.5)$$

$$f_{n+1}(x) < f_n(x) \text{ for } x \in (y_n, 1). \quad (5.6)$$

Indeed, this holds as, at $x = 0$, we have $f_n(0) = f_{n+1}(0)$ and $f'_n(0) < f'_{n+1}(0)$, and at $x = 1$, we have $f_n(1) = f_{n+1}(1)$ and $f'_n(1) < f'_{n+1}(1)$. It also holds that

$$\lim_{n \rightarrow \infty} y_n = \frac{1}{2} \quad \text{and} \quad \lim_{n \rightarrow \infty} f_n(y_n) = 1 - \rho_I. \quad (5.7)$$

B. By (5.7), $f_n(y_n) < y_n$ for large $n \in \mathbb{N}$. Hence, $x_n < y_n$ for large $n \in \mathbb{N}$. It follows by (5.5) that $x_n = f_n(x_n) < f_{n+1}(x_n)$; hence, $x_n < x_{n+1}$ for large $n \in \mathbb{N}$. We conclude that the sequence $\{x_n\}_{n \in \mathbb{N}}$ is eventually increasing. Because $x_n < y_n$ for large $n \in \mathbb{N}$, we have $\lim (x_n)^n \leq \lim (y_n)^n = 0$. And because $\{x_n\}_{n \in \mathbb{N}}$ is bounded away from zero, also $\lim (1 - x_n)^n = 0$. Taking the limit of $x_n = f_n(x_n)$, we find that $x_n \rightarrow 1 - \rho_I$.

C2. By (5.7), $y_n < f_n(y_n)$ for large $n \in \mathbb{N}$. Hence, $y_n < x_n$ for large $n \in \mathbb{N}$. It follows by (5.6) that $f_{n+1}(x_n) < f_n(x_n) = x_n$; hence, $x_{n+1} < x_n$ for large $n \in \mathbb{N}$. We conclude that the sequence $\{x_n\}_{n \in \mathbb{N}}$ is eventually decreasing. Hence, $\lim (x_n)^n = 0$. And because the sequence is bounded away from zero, also $\lim (1 - x_n)^n = 0$. Taking the limit of $x_n = f_n(x_n)$, we find that $x_n \rightarrow 1 - \rho_I$.

C3. Because $\rho_{\text{II}} < \frac{1}{2} = \rho_{\text{I}}$, we find that $\frac{1}{2} < f_n\left(\frac{1}{2}\right)$, and hence, $\frac{1}{2} < x_n$. On the other hand, $y_n < \frac{1}{2}$. We conclude that $y_n < x_n$ for all $n \in \mathbb{N}$. The rest of the argument is identical to that in C2. \square

Example 5.3. In the setup of Example 3.2, we have

$$\lim_{n \rightarrow \infty} \mathbb{P}_{p_n}(v < k) = 1 - q + qk$$

for each $0 < k \leq 1$. The sequence $\mathbb{P}_{p_n}(v < k)$ is (A) constant if $q = 0$ or $k = 1$, (B) eventually increasing if $1/2 < q < 1$ and $0 < k < 1 - 1/2q$, and (C) eventually decreasing otherwise.

5.4. Player I’s k -optimal Strategies and k -optimal Subtree

The k -optimal subtree is the largest subtree of the game tree containing the root and consisting of the nodes in which the value (of the corresponding subgame) is at least k . The motivation for this notion stems from its connection to the set of player I’s k -optimal strategies: player I’s strategy is k -optimal if and only if it never requires player I to make a move leading outside the k -optimal subtree. The k -optimal subtree generalizes the notion of the reduced family tree of a branching process: the latter consists of the individuals having an infinite line of descent (e.g., Athreya and Ney [4, section I.D.12], Lyons and Peres [19, section 5.7]).

The number $k > 0$ is fixed throughout this section. We write $\alpha = \alpha(k)$ and $\beta = \beta(k)$, and we assume that $\beta > 0$.

Consider a game $\omega \in \Omega$ with a value $v_\omega \geq k$. Define the k -optimal subtree, T_ω^* , to be the largest subtree of T_ω consisting of the nodes having a value of at least k and containing \emptyset . Equivalently, one can define T_ω^* recursively as follows: the empty sequence \emptyset is an element of T_ω^* . For $h \in T_\omega^*$ and $j \in \{1, \dots, \xi_\omega(h)\}$, declare (h, j) to be an element of T_ω^* if $v_\omega(h, j) \geq k$.

Intuitively, the tree T_ω^* consists of player I’s moves that allow her to “defend” a payoff of k against player II. To be able to guarantee the payoff of at least k , all player I needs to do is never choose an action leading outside T_ω^* . More formally, consider player I’s strategy σ_{I} . Then, σ_{I} is k -optimal if and only if, whenever player I’s node h is in T_ω^* , σ_{I} selects a child $\sigma_{\text{I}}(h) \in \{1, \dots, \xi_\omega(h)\}$ of h such that the successor node $(h, \sigma_{\text{I}}(h))$ is also in T_ω^* .

Define p^* to be the following probability distribution on the set \mathbb{N} of natural numbers:

$$p^*(0) = p_{\text{I}}(k, 0)\beta^{-1} + p_{\text{II}}(k, 0)\beta^{-1}, \text{ and}$$

$$p^*(n) = \sum_{m \geq n} p_{\text{I}}(k, m) \binom{m}{n} \beta^{n-1} \alpha^{m-n} + p_{\text{II}}(k, n)\beta^{n-1} \text{ for each } n \geq 1.$$

A technical remark is in order. The k -optimal subtree is, in general, not an ordered tree. However, one obtains an ordered copy of T_ω^* simply renaming the children of a node of T_ω^* using consecutive natural numbers. It is to such an ordered copy of the k -optimal subtree that Theorem 5.1 refers.

Theorem 5.1. *Suppose that $\beta > 0$. Conditional on the event $\{\omega \in \Omega : v_\omega \geq k\}$, the distribution (of the ordered copy) of the k -optimal subtree T_ω^* is the Galton–Watson measure with p^* as the offspring distribution.*

The intuition is as follows. Consider a node h of the game tree T_ω with $v_\omega(h) \geq k$. Then, h has no children with a value of at least k precisely when h has no children in T_ω and has a capacity of at least k . Conditioning on the event $\{\omega \in \Omega : v_\omega(h) \geq k\}$, the event of probability β , we obtain $p^*(0)$.

Turning to the second expression, we argue as follows: consider again a node h of the game tree T_ω with $v_\omega(h) \geq k$. The node h has exactly n children with a value of at least k in either of the following cases: (I) h is player I’s node, it has a capacity of at least k , and it has $m \geq n$ children in T_ω of which exactly n have a value of at least k ; (II) h is player II’s node, it has a capacity of at least k , and it has exactly n children in T_ω , all of which have a value of at least k . Conditioning on the event $\{\omega \in \Omega : v_\omega(h) \geq k\}$, we obtain $p^*(n)$.

We remark that the mean of the offspring distribution p^* is

$$e(k) =: \mathbb{E}_p(1_{\{\text{I}\}} \mathbb{1}_{\{\gamma \geq k\}} \xi + 1_{\{\text{II}\}} \mathbb{1}_{\{\gamma \geq k\}} \xi \beta^{\xi-1}), \tag{5.8}$$

and that it is bounded below by $d(k)$.

5.5. The Avoidance Game

Consider a scenario under which player I is prohibited from taking an action that leads to player II’s node with more than one child. We think of this scenario as a proxy for the situation in which player I is reluctant to concede a turn of the game to her opponent. Such behavior on the part of player I effectively deprives player II of

any real choice in the game, rendering him a dummy: indeed, the only nodes assigned to him that could ever be visited during play are end nodes or the nodes with a single child.

What payoff can player I still guarantee (with positive probability) under this scenario? It is clear that the avoidance game is harder for her to play. And yet, as is shown, if player I can guarantee a payoff of at least k with positive probability in the absence of any restrictions on her behavior, she can do so without ever visiting player II's nodes with more than one child.

Given a game ω , define the *avoidance game* ω' by setting the capacity to zero at any player II node with more than one child: thus, $T_{\omega'} = T_{\omega}$, $\iota_{\omega'} = \iota_{\omega}$, whereas $\gamma_{\omega'}$ is defined as follows: $\gamma_{\omega'}(h) = 0$ whenever $\iota_{\omega}(h) = \text{II}$ and $\xi_{\omega}(h) \geq 2$, and $\gamma_{\omega'}(h) = \gamma_{\omega}(h)$ otherwise. Setting capacity at a node to zero effectively prohibits player I from visiting that node. Clearly, the value $v_{\omega'}$ of the game ω' is not greater than the value of ω .

Define a primitive distribution p' by letting $p'_i(k, n)$ be zero if $i = \text{II}$, $n \geq 2$, and $k > 0$, and letting it be equal to $p_i(k, n)$ otherwise. We state the following lemmas for the sake of completeness. The proofs are easy and are omitted.

Lemma 5.1. *The map $\omega \mapsto \omega'$, $\Omega \rightarrow \Omega$ is Borel measurable.*

Lemma 5.2. *The distribution of the game ω' is given by the measure $\mathbb{P}_{p'}$: that is, $\mathbb{P}_p(\{\omega' \in B\}) = \mathbb{P}_{p'}(B)$ for each Borel set $B \subseteq \Omega$.*

The main message of this section is the following rather surprising result (it follows from Lemma 5.2, Theorem 4.2, and Corollary 5.1).

Corollary 5.5. *Let $k > 0$ and suppose that $p(\{\gamma < k\}) > 0$. Then, $\mathbb{P}_p(\{v_{\omega} \geq k\}) > 0$ if and only if $\mathbb{P}_{p'}(\{v_{\omega'} \geq k\}) > 0$. If, moreover, p is an escape model, then the two conditions are equivalent to $d(k) > 1$.*

Example 5.4. Consider the setup of Example 3.1. The vgf corresponding to the measure p' (where we take $k = 1$) is given

$$f'_1(x) = 1 - l(1-x) \left(q_I \frac{1}{1-lx} + q_{\text{II}}(1-l) \right).$$

Computing the smallest fixed point of the function, we obtain

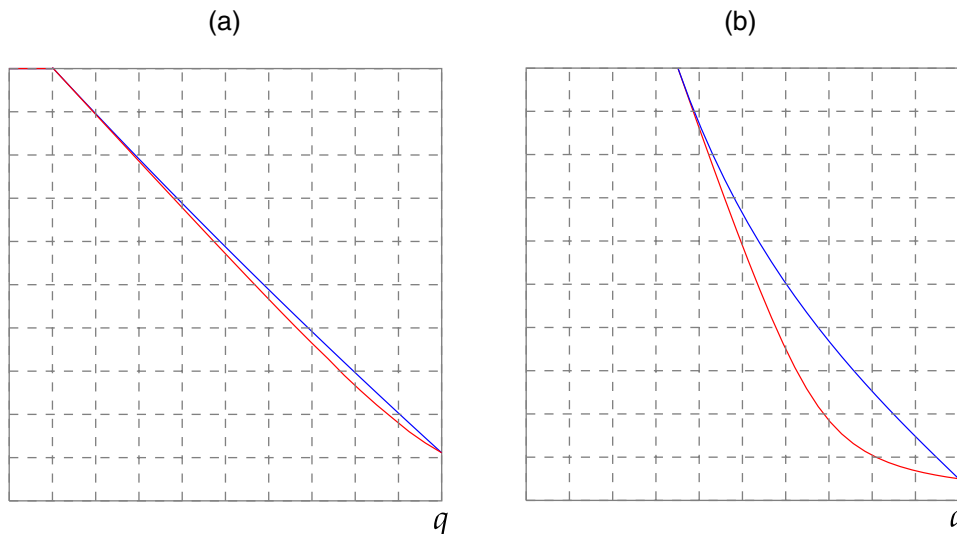
$$\mathbb{P}(v_{\omega'} = 0) = \frac{1 - l + l^2 - l^2q}{l(1 - l + l^2 + l(1-l)q)}.$$

The probabilities $\mathbb{P}(v_{\omega'} = 0)$ and $\mathbb{P}(v_{\omega} = 0)$ are pictured in Figure 5 (the latter is the same as the red line in Figure 1).

Consider Example 3.2. Note that, because all nodes have n children, in ω' , all the nodes controlled by player II have capacity zero. The vgf corresponding to p' is

$$f'_k(x) = 1 - (1-k)q + (1-k)qx^n,$$

Figure 5. (Color online) (a) The probabilities of the events $\{v_{\omega'} = 0\}$ (blue) and $\{v_{\omega} = 0\}$ (red) in Example 3.1 as functions of q for $l = 0.9$. (b) The probabilities of the events $\{v_{\omega'} < 0.05\}$ (blue) and $\{v_{\omega} < 0.05\}$ (red) in Example 3.1 with $n = 3$ for $k = 0.05$.



and its smallest fixed point for the ternary case $n = 3$ is

$$\mathbb{P}(v_{\omega'} < k) = -\frac{1}{2} + \sqrt{\frac{1}{(1-k)q} - \frac{3}{4}}.$$

This probability, along with $\mathbb{P}(v_{\omega} < k)$ is pictured in the right panel of Figure 5 for $k = 0.05$ as a function of q (the latter graph is the same as that in Figure 2).

We conclude this section by reflecting on the nature of k -optimal strategies in ω' and comparing them to those in ω .

Fix a $k > 0$. One can think of player I's k -optimal strategy in ω' as a play in the tree T_{ω} that never visits a node with a capacity less than k or a node in which player II has more than one move. Clearly, a k -optimal strategy in ω' is also a k -optimal strategy in ω , but in general, player I has k -optimal strategies in ω that are unavailable to her in ω' .

Recall that the k -optimal subtree (see Section 5.4) characterizes player I's k -optimal strategies. Thus, suppose that $v_{\omega'} \geq k$ and consider T_{ω}^* and $T_{\omega'}^*$: the former is the k -optimal subtree of ω , whereas the latter is the k -optimal subtree of the avoidance game ω' . As can be seen from the fact that $v_{\omega'}(h) \leq v_{\omega}(h)$ for each node $h \in T_{\omega'}$, the k -optimal trees of ω' and of ω are nested: $T_{\omega'}^* \subseteq T_{\omega}^*$.

Applying Theorem 5.1 and Equation (5.8) to the probability measure p' and using Lemma 5.2, we obtain the following statement.

Corollary 5.6. *Suppose that $\mathbb{P}_p(\{v_{\omega'} \geq k\}) > 0$. Conditional on the event $\{\omega \in \Omega : v_{\omega'} \geq k\}$, the distribution (of the ordered copy) of the tree $T_{\omega'}^*$ is the Galton–Watson measure generated by an offspring distribution with mean $d(k)$.*

This result leads to an interpretation of $d(k)$ as the average number of k -optimal actions player I has in the avoidance game or, equivalently, as the average number of actions that allow her to guarantee a payoff of at least k while at the same time avoiding visiting nodes at which the player's opponent has more than one move.

One can compare the two trees with the help of branching numbers (Lyons and Peres [19, section 1.8]).

Corollary 5.7. *Suppose that p is an escape model. Let $k > 0$ and suppose that $d(k) > 1$. Then, \mathbb{P}_p -almost surely*

- i. *On the event $\{\omega \in \Omega : v_{\omega} \geq k\}$, the branching number of T_{ω}^* equals $e(k)$.*
- ii. *On the event $\{\omega \in \Omega : v_{\omega'} \geq k\}$, the branching number of $T_{\omega'}^*$ equals $d(k)$.*

Proof. The corollary rests on the following well-known fact (see, e.g., Lyons and Peres [19, section 1.8 and corollary 5.10]): consider a Galton–Watson measure on trees generated by the offspring distribution with a mean greater than one. Almost surely on the event of nonextinction of the tree, its branching number equals the mean of the offspring distribution.

Using Theorem 5.1 and the remarks thereafter, we argue thus: conditional on the event $\{v_{\omega} \geq k\}$, the distribution of T_{ω}^* is the Galton–Watson measure with the mean number of offspring given by (5.8). This number, $e(k)$, is bounded below by $d(k)$ and is, therefore, greater than one. Moreover, because p is an escape model, T_{ω}^* has no end nodes; in particular, it is extinct nowhere on the event $\{v_{\omega} \geq k\}$. Thus, we conclude that the branching number of T_{ω}^* is $e(k)$ almost everywhere on $\{v_{\omega} \geq k\}$.

Using Corollary 5.6, we argue that, conditional on the event $\{v_{\omega'} \geq k\}$, the distribution of $T_{\omega'}^*$ is the Galton–Watson measure with mean $d(k) > 1$. Because p is an escape model, so is p' , implying that $T_{\omega'}^*$ has no end nodes; in particular, it is extinct nowhere on $\{v_{\omega'} \geq k\}$. Thus, we conclude that the branching number of $T_{\omega'}^*$ is $d(k)$ almost everywhere on $\{v_{\omega'} \geq k\}$. \square

We already note the fact that $d(k) \leq e(k)$. Whenever $d(k) < e(k)$, one may conclude that, almost surely on the event $\{\omega \in \Omega : v_{\omega'} \geq k\}$, the tree $T_{\omega'}^*$ is a proper subset of T_{ω}^* . This indicates that, typically, player I has k -optimal strategies in ω that are not available to her in ω' .

6. Discussion and Open Questions

6.1. On Player II's k -optimal Strategies

We reflect on the difference in the nature of k -optimal strategies of player I and those of player II.

Consider the following player I strategy for a game with value k : always choose the youngest (i.e., the one with the lowest index) child of the current node having a value of at least k . This strategy (let us call it the *one-step* strategy) is k -optimal. The counterpart of the one-step strategy for player II—always choosing the youngest child with a value no greater than k —need not be k -optimal. For instance, take a game in which player II is the only player, each node has two children, the younger child (the one with the index one) has capacity one,

whereas the older child has capacity zero. The one-step strategy is not 0-optimal. Any 0-optimal strategy for player II must choose the older child at some point.

A question arises whether examples such as this one are exceptional in the measure-theoretic sense. Would not player II's one-step strategy be k -optimal in all games with a value of at most k apart from a set of measure zero? In general, the answer is no.

Consider Example 3.1 with $l = 0.9$ and $q = 0.1$, which is slightly below the critical level (of 0.1021) so that $\alpha = 1$. Thus, player II does have a 0-optimal strategy with probability one. Consider what happens if player II uses the one-step strategy instead. Because all nodes of the game tree have value zero with probability one, the one-step strategy boils down to choosing the first child.

Let $S_\omega \subseteq T_\omega$ denote the tree consisting of the nodes that are consistent with player II's one-step strategy. The tree S_ω is distributed according to a Galton–Watson measure with the mean $0.1 \cdot 9 + 0.9 \cdot 0.9 \cdot 1 = 1.71$ (because a node controlled by player I has nine children on average, whereas a node controlled by player II has no children or exactly one child, depending on whether in T_ω the node had any children). Thus, S_ω has an infinite branch with positive probability. And, hence, player I has a positive probability of winning against player II's one-step strategy. We conclude that, with positive probability, player II's one-step strategy is not 0-optimal.

This illustrates that describing player II's k -optimal strategies requires a different approach than that we have employed in Section 5.4. Finding a suitable approach is an interesting direction for future work.

6.2. On Finite Games

Our model encompasses certain classes of finite games. To illustrate, suppose that the mean of the offspring distribution $E_p(\xi)$ is smaller than one, that the capacity equals one whenever $\xi > 0$, and that it is uniformly distributed on $[0, 1]$ if $\xi = 0$. Under these assumptions, the game tree T_ω is finite almost surely, and the payoff equals the capacity at the end node. Thus, we have a model of finite games with payoffs randomly and independently assigned to the end nodes of the game tree.

The model outlined above resembles that in Arieli and Babichenko [3] with the difference that our game tree is random; in particular, there is a positive probability for the root of the game tree to be its only node. Thus, the two models are distinct. Developing a framework to encompass both these models as special cases might be a fruitful avenue for future research.

6.3. Subgame Perfect Equilibrium in Multiplayer Perfect Information Games

Multiplayer perfect information games with semicontinuous payoffs are the subject of much work (Flesch and Predtetchinski [7], Flesch et al. [9, 10], Purves and Sudderth [28]). Nevertheless, we believe that the modeling technique of this paper, suitably adapted, could offer a new perspective on the topic. Here, we suggest one particular question.

Consider perfect information games played by an infinite sequence of players $0, 1, 2, \dots$, player $t \in \mathbb{N}$ moving only once, in period t . Suppose that each has a lower semicontinuous payoff. It is known that not all games of this class admit a subgame perfect ϵ -equilibrium. But how “large” is the set of games that do have one? One possible approach to this question is probabilistic: suppose that each player's payoff function is generated randomly by independently assigning the capacities to the nodes as is done in this work. What is the measure of games that have a subgame perfect ϵ -equilibrium?

6.4. On the Nonadversarial Case

The nonadversarial case is the special case of the model with $q_I = 1$ so that all the nodes of the game tree are assigned to player I. We have mentioned the nonadversarial case when discussing the examples in Section 3; in both examples, the nonadversarial case boils down to the question of (non-)extinction of a certain branching process. This observation can be generalized to escape models.

Consider an escape model with $q_I = 1$, and let T_ω^k be the subtree of T_ω consisting of the nodes with a capacity of at least k . Then, the value of ω is at least k precisely when T_ω^k has an infinite branch. The tree T_ω^k is governed by a Galton–Watson measure, and the corresponding offspring distribution is that of the random variable $1_{\{\gamma \geq k\}} \xi$. In particular, $\{v_\omega \geq k\}$ has a positive probability if and only if $E_p(1_{\{\gamma \geq k\}} \xi) > 1$ or $p(\{\gamma \geq k\} \cap \{\xi = 1\}) = 1$.

In general, the nonadversarial case is not entirely trivial. For example, if the capacity is independent of the number of children, player I might strive to finish the game by reaching an end node and do so sooner rather than later lest the player encounter a node with a low capacity. Consider, for example, the nonadversarial case of the model described in Section 6.2. The value is then the maximum of a random number η of independent uniformly distributed random variables, in which the random variable η is the number of end nodes in the family tree of a subcritical branching process. See Nariyuki [24] for the properties of its distribution.

Acknowledgments

The authors are grateful to Andrés Perea, Christian Seel, and Xavier Venel for their questions and suggestions.

Endnotes

- ¹ After our manuscript was accepted for publication, we became aware of the paper by Garnier and Ziliotto [13].
- ² Note a slight departure from the standard terminology: what we call a v_ω -optimal strategy would usually be called an optimal strategy.
- ³ The general expression for the critical activation probability q_c is (5.1). The expression for α follows by Theorem 4.1 using Example 4.1.
- ⁴ These expressions can be derived easily using Example 4.1 and Theorem 4.1.

References

- [1] Amiet B, Collevocchio A, Scarsini M, Zhong Z (2019) Pure Nash equilibria and best-response dynamics in random games. *Math. Oper. Res.* 46(4):1552–1572.
- [2] Apt KR, Grädel E, eds. (2011) *Lectures in Game Theory for Computer Scientists* (Cambridge University Press, Cambridge, UK).
- [3] Arieli I, Babichenko Y (2016) Random extensive form games. *J. Econom. Theory* 166:517–535.
- [4] Athreya KB, Ney PE (1972) *Branching Processes* (Springer Berlin, Heidelberg, Germany).
- [5] Basu R, Holroyd AE, Martin JB, Wästlund J (2016) Trapping games on random boards. *Ann. Appl. Probab.* 26(6):3727–3753.
- [6] Daskalakis C, Dimakis AG, Mossel E (2011) Connectivity and equilibrium in random games. *Ann. Appl. Probab.* 21(3):987–1016.
- [7] Flesch J, Predtetchinski A (2017) A characterization of subgame-perfect equilibrium plays in Borel games of perfect information. *Math. Oper. Res.* 42(4):1162–1179.
- [8] Flesch J, Predtetchinski A, Suomala V (2021) Random perfect information games. Preprint, submitted April 21, <https://arxiv.org/abs/2104.10528>.
- [9] Flesch J, Herings PJJ, Maes J, Predtetchinski A (2019) Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information. *Econom. Theory* 29:1–25.
- [10] Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E, Vrieze K (2010) Perfect-information games with lower-semicontinuous payoffs. *Math. Oper. Res.* 35(4):742–755.
- [11] Gale D, Stewart FM (1953) Infinite games with perfect information. Kuhn HW, Tucker AW, eds. *Contributions to the Theory of Games* (Princeton University Press, Princeton, NJ), vol. 2, 245–266.
- [12] Galla T, Farmer JD (2013) Complex dynamics in learning complicated games. *Proc. Natl. Acad. Sci. USA* 110(4):1232–1236.
- [13] Garnier G, Ziliotto B (2021) Percolation games. Preprint, submitted October 23, <https://arxiv.org/abs/2110.12227>.
- [14] Grimmett G, Stirzaker D (2001) *Probability and Random Processes* (Oxford University Press, UK).
- [15] Harris C (1985) Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica* 53:613–628.
- [16] Holroyd AE, Martin JB (2019) Galton–Watson games. Preprint, submitted April 8, <https://arxiv.org/abs/1904.04150>.
- [17] Holroyd AE, Marcovici I, Martin JB (2019) Percolation games, probabilistic cellular automata, and the hard–core model. *Probab. Theory Related Fields* 174(3–4):1187–1217.
- [18] Laraki R, Maitra AP, Sudderth WD (2013) Two-person zero-sum stochastic games with semicontinuous payoff. *Dynamic Games Appl.* 3:162–171.
- [19] Lyons R, Peres Y (2017) *Probability on Trees and Networks* (Cambridge University Press, Cambridge, UK).
- [20] Martin DA (1975) Borel determinacy. *Ann. Math.* 102(2):363–371.
- [21] McLennan A (2005) The expected number of Nash equilibria of a normal form game. *Econometrica* 73(1):141–174.
- [22] McLennan A, Berg J (2005) Asymptotic expected number of Nash equilibria of two-player normal form games. *Games Econom. Behav.* 51(2):264–295.
- [23] Moschovakis YN (2009) *Descriptive Set Theory* (American Mathematical Society, Providence, RI).
- [24] Nariyuki N (2005) On the number of vertices with a given degree in a Galton–Watson tree. *Adv. Appl. Probab.* 37(1):229–264.
- [25] Neveu J (1986) Arbres et processus de Galton–Watson. *Annales de l’I.H.P. Probabilités et statistiques* 22(2):199–207.
- [26] Pei T, Takahashi S (2019) Rationalizable strategies in random games. *Games Econom. Behav.* 118:110–125.
- [27] Prikry K, Sudderth WD (2016) Measurability of the value of a parametrized game. *Internat. J. Game Theory* 45:675–683.
- [28] Purves RA, Sudderth WD (2011) Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.* 36(3):468–473.
- [29] Rinott Y, Scarsini M (2000) On the number of pure strategy Nash equilibria in random games. *Games Econom. Behav.* 33(2):274–293.
- [30] Stanford W (1996) The limit distribution of pure strategy Nash equilibria in symmetric bimatrix games. *Math. Oper. Res.* 21(3):726–733.
- [31] Takahashi S (2008) The number of pure Nash equilibria in a random game with nondecreasing best responses. *Games Econom. Behav.* 63(1):328–340.
- [32] Van Benthem J (2014) *Logic in Games* (MIT Press, Cambridge, MA).