



Decision Support

Static search games played over graphs and general metric spaces[☆]Thomas P. Oléron Evans^{*}, Steven R. Bishop*Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK**Centre for Advanced Spatial Analysis, UCL Bartlett Faculty of the Built Environment, 90 Tottenham Court Road, London W1T 4TJ, UK*

ARTICLE INFO

Article history:

Received 28 January 2013

Accepted 17 June 2013

Available online 27 June 2013

*Keywords:*Game theory
Search games
Networks
Graph theory
Metric spaces

ABSTRACT

We define a general game which forms a basis for modelling situations of static search and concealment over regions with spatial structure. The game involves two players, the searching player and the concealing player, and is played over a metric space. Each player simultaneously chooses to deploy at a point in the space; the searching player receiving a payoff of 1 if his opponent lies within a predetermined radius r of his position, the concealing player receiving a payoff of 1 otherwise. The concepts of dominance and equivalence of strategies are examined in the context of this game, before focusing on the more specific case of the game played over a graph. Methods are presented to simplify the analysis of such games, both by means of the iterated elimination of dominated strategies and through consideration of automorphisms of the graph. Lower and upper bounds on the value of the game are presented and optimal mixed strategies are calculated for games played over a particular family of graphs.

© 2013 The Authors. Published by Elsevier B.V. All rights reserved.

1. Introduction

In this paper, we define a general search and concealment game that takes full account of the spatial structure of the set over which it is played. The game is static in the sense that players do not move, but deploy simultaneously at particular spatial points and receive payoffs based on their relative positions. In this way, the static spatial search game (SSSG) provides a theoretical foundation for the study of the relative strategic value of different positions in a geography. Using the theory of metric spaces, we model situations in which the searching player may simultaneously search multiple locations based on concepts of distance or adjacency relative to the point at which they are deployed.

While the SSSG does build upon previous work, particularly that of Ruckle (1983) and White (1994), its simplicity and generality together with its explicit consideration of spatial structure set it apart from much of the literature (see Section 3 for a detailed review of related work) and lend it the versatility to describe games over a huge variety of different spaces. The primary contributions of this article are therefore to both propose a highly general model of spatial search and concealment situations, which unites several other games presented in the literature (see Section 4.2), and to present new propositions and approaches for the strategic analysis of such scenarios.

[☆] This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

^{*} Corresponding author at: Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK. Tel.: +44 (0)20 3108 3877; fax: +44 (0)20 3108 3258.

E-mail addresses: thomas.evans.11@ucl.ac.uk (T.P. Oléron Evans), s.bishop@ucl.ac.uk (S.R. Bishop).

While this paper is theoretical in nature, the SSSG provides a framework for the analysis of a diverse range of operational research questions. Aside from explicit search and concealment scenarios, the game may be used to model situations in which some structure or region must be protected against ‘attacks’ that could arise at any spatial point; for example, the deployment of security personnel to protect cities against terrorist attacks or outbreaks of rioting, security software scanning computer networks to eliminate threats, the defence of shipping lanes against piracy, the protection of a rail network against cable theft or the deployment of stewards at public events to respond to emergency situations.

We provide a brief overview of all necessary game theoretic concepts in Section 2 and a review of the literature on games of search and security in Section 3, before formally defining the SSSG, examining its relationship to other games in the literature and presenting some initial propositions in Section 4. In Section 5, we examine the SSSG on a graph and identify upper and lower bounds on the value of such games before presenting an algorithm in Section 6 which simplifies graph games by means of the iterated elimination of dominated strategies, focusing particularly on the application of the algorithm to games played on trees. Section 7 contains further results, including a way to simplify graph games through consideration of graph automorphisms and an examination of a particular type of strategy for such games, which we describe as an “equal oddments strategy”. In Section 8, we use the concept of an equal oddments strategy to find analytic solutions for a particular family of graph games, while Section 9 forms a conclusion to the paper, containing a summary of our key results and suggestions of potential avenues for further research. Two proofs, which were too complicated to include in the main text, are presented as appendices.

2. Game theoretic concepts

The definitions and notation relating to game theory used in this section are adapted from Blackwell and Girshick (1954) and Morris (1994).

When discussing two-player games, we assume the following definition:

Definition 2.1. Two-player games in normal form

A game in normal form between Players A and B , consists of:

- **strategy sets** Σ_A, Σ_B
- **payoff functions** p_A, p_B , with:

$$p_A : \Sigma_A \times \Sigma_B \rightarrow \mathbb{R}$$

$$p_B : \Sigma_A \times \Sigma_B \rightarrow \mathbb{R}$$

If the payoffs are such that for some constant c :

$$p_A(x, y) + p_B(x, y) = c, \quad \forall x \in \Sigma_A, \quad \forall y \in \Sigma_B$$

then the game is described as a constant-sum game.

The game is played by Players A and B simultaneously choosing **strategies** (described as **pure strategies** in cases where there may be any ambiguity) from their respective strategy sets $x \in \Sigma_A, y \in \Sigma_B$ and receiving payoffs $p_A(x, y), p_B(x, y)$. The objective of each player is to maximise their payoff.

In certain circumstances, we may allow players to adopt **mixed strategies**, whereby they choose their pure strategy according to a specified probability distribution. If Σ_A and Σ_B are finite, with:

$$\Sigma_A = \{x_1, \dots, x_{\kappa_A}\}$$

$$\Sigma_B = \{y_1, \dots, y_{\kappa_B}\}$$

for some positive integers κ_A, κ_B , then the mixed strategies σ_A, σ_B can simultaneously be regarded as vectors:

$$\sigma_A = (\sigma_A[x_1], \dots, \sigma_A[x_{\kappa_A}]) \in [0, 1]^{\kappa_A}$$

$$\sigma_B = (\sigma_B[y_1], \dots, \sigma_B[y_{\kappa_B}]) \in [0, 1]^{\kappa_B}$$

and as functions, which allocate probabilities to pure strategies:

$$\sigma_A : \Sigma_A \rightarrow [0, 1]$$

$$x \mapsto \sigma_A[x]$$

$$\sigma_B : \Sigma_B \rightarrow [0, 1]$$

$$y \mapsto \sigma_B[y]$$

$$\sum_{x \in \Sigma_A} \sigma_A[x] = \sum_{y \in \Sigma_B} \sigma_B[y] = 1$$

The following definitions relate to the maximum expected payoff that players can guarantee themselves through careful choice of their mixed strategies:

Definition 2.2. Values of the game

Given a two-player game, the values of the game u_A, u_B to Players A and B respectively, are defined as:

- $u_A = \max_{\tau_A} \min_{\tau_B} E[p_A(\tau_A, \tau_B)]$
- $u_B = \max_{\tau_B} \min_{\tau_A} E[p_B(\tau_A, \tau_B)]$

where τ_A and τ_B range across all possible mixed strategies for Players A and B respectively and

$$E[p_A(\tau_A, \tau_B)]$$

$$E[p_B(\tau_A, \tau_B)]$$

represent the expected payoffs to each player, given that they respectively adopt mixed (or pure) strategies τ_A and τ_B .

Definition 2.3. Optimal mixed strategies

Given a two-player constant-sum game, where the payoffs sum to $c \in \mathbb{R}$, mixed strategies σ_A, σ_B for Players A and B are described as **optimal** if and only if:

- $\min_{\tau_B} E[p_A(\sigma_A, \tau_B)] = u_A$
- $\min_{\tau_A} E[p_B(\tau_A, \sigma_B)] = u_B$

where τ_A and τ_B range across all possible mixed strategies for Players A and B respectively.

For a constant-sum game, where the payoffs sum to $c \in \mathbb{R}$, we have:

$$u_A + u_B = c \quad (1)$$

Also, provided that Σ_A and Σ_B are finite, optimal mixed strategies are guaranteed to exist for both players.

Both of these facts are consequences of the Minimax Theorem (see Morris, 1994, p. 102).

Given a constant-sum two-player game with finite strategy sets, a **solution** of the game comprises optimal mixed strategies Σ_A, Σ_B and values u_A, u_B for each Player.

The following definition allows for a crude comparison of the efficacy of different strategies.

Definition 2.4. Strategic dominance and equivalence

Consider a two-player game with strategy sets Σ_A, Σ_B and payoff functions p_A, p_B . Given particular pure strategies $x_1, x_2 \in \Sigma_A$ for Player A , we have:

- x_2 **very weakly dominates** x_1 if and only if:

$$p_A(x_2, y) \geq p_A(x_1, y), \quad \forall y \in \Sigma_B$$

- x_2 **weakly dominates** x_1 if and only if:

$$p_A(x_2, y) \geq p_A(x_1, y), \quad \forall y \in \Sigma_B$$

and $\exists y^* \in \Sigma_B$ such that:

$$p_A(x_2, y^*) > p_A(x_1, y^*)$$

- x_2 **strictly dominates** x_1 if and only if:

$$p_A(x_2, y) > p_A(x_1, y), \quad \forall y \in \Sigma_B$$

- x_2 is **equivalent** to x_1 if and only if:

$$p_A(x_2, y) = p_A(x_1, y), \quad \forall y \in \Sigma_B$$

Since the designation of the players as A and B is arbitrary, obtaining corresponding definitions of strategic dominance and equivalence for Player B is simply a matter of relabelling.

Note that weak dominance, strict dominance and equivalence are all special cases of very weak dominance. Also, strict dominance is a special case of weak dominance.

In this paper, weak dominance is of most relevance. Therefore, for reasons of clarity, the terms “dominance” and “dominated strategies” will be used to refer to weak dominance unless otherwise stated.

Since a player aims to maximise his or her payoff, we would intuitively expect that they should not play any dominated strategies.

For a general definition of dominance in game theory, see Leyton-Brown and Shoham (2008, pp. 20–23), from which the above definition was adapted.

3. Games of search and security: a review

Games of search and concealment, in which one player attempts to hide themselves or to conceal some substance in a

specified space while another player attempts to locate or capture the player or substance, have been widely studied.

3.1. Simple search games

One of the simplest search games is the well-known high-low number guessing game in which one player chooses an integer in a given range, while the other player makes a sequence of guesses to identify it, each time being informed whether the guess was too high or too low (Gal, 1974). Continuous versions of the game have also been studied (Gal, 1978; Baston et al., 1985; Alpern, 1985a).

Another simple search game involves one player attempting to locate an object that the opposing player has hidden at a location chosen from a finite or countably infinite set with no spatial structure (except a possible ordering). Variants of these games include examples where the searching player has some chance of overlooking the object despite searching the correct location (Neuts, 1963; Subelman, 1981) or where the searcher must also avoid the location of a second hidden object (Ruckle, 1990).

3.2. Search games with immobile targets

A more complicated class of search games is that in which the searching player is mobile and their target is immobile, with payoffs to each player typically (though not universally) being dependent on the amount of time that elapses before the target is located. Such games have been examined over many different types of graph (Anderson et al., 1990; Alpern, 2008, 2010; Buyang, 1995; Reijnierse & Potters, 1993; Kikuta et al., 1994; Kikuta, 2004; Pavlović, 1995), though in general the space may be a continuous region (Gal, 1979). While the starting position of the searching player is often fixed, games in which the searching player can choose their position have also been studied (Alpern, Baston, & Gal, 2008a; Dagan et al., 2008), as have games with multiple searchers (Alpern & Howard, 2000; Peshkin, 1994).

3.3. Accumulation games

Accumulation games are an extension of this concept in which there may be many hidden objects (Kikuta et al., 1997) or in which hidden objects are replaced with some continuous material that the hiding player can distribute across a set of discrete locations (Kikuta et al., 2002; Zoroa, Fernández-Sáez, & Zoroa, 2004) or across a continuous space (Ruckle et al., 2000). The payoffs in these games are typically dependent on the number of objects or the quantity of material that the searching player is able to locate.

3.4. Search games with mobile targets

Adding a further layer of complication, there is the class of search game in which both the searching player and the hiding player are mobile, including so-called “princess and monster” games. Again, the payoffs in such games are typically dependent on the amount of time that elapses before the hiding player is captured and players are typically ‘invisible’ to each other, only becoming aware of the location of their opponent at the moment of capture.

Such games have been considered over continuous one-dimensional spaces such as the circle (Alpern, 1974) and the unit interval (Alpern, Fokkink, Lindelauf, & Olsder, 2008b), over continuous graphs or networks (Alpern et al., 1986; Anderson et al., 1992; Alpern, 1985b) and over continuous two-dimensional spaces (Foreman, 1977; Garnaeu, 1991; Chkhartishvili et al., 1995). In the latter case, it is necessary to introduce the concept of a *detection radius*, with a capture occurring if the distance between the players drops below this value. In some cases, the probability of capture is allowed to vary based on the distance between the players (Garnaeu, 1992).

Analyses of search games over discrete spaces in which both searcher and hider are mobile have tended to consider spatial structure in only a very limited way. While this structure may determine the freedom of movement of the players, very little work has been done to introduce an analogous concept to the detection radius to such games. Generally, players move sequentially and may only move to locations that are sufficiently close to their current position (e.g. Eagle et al., 1991), though variants have been considered in which either the searching player (Zoroa, Fernández-Sáez, & Zoroa, 2012) or the hiding player (Thomas et al., 1991) has the freedom to move to any location regardless of adjacency or distance.

Further variations on the search game with mobile searcher and hider include games in which the searching player follows a pre-determined path and must decide how thoroughly to search each location visited (Hohzaki et al., 2000), games in which the searching player must intercept an opponent attempting to move from a given start point to a given end point (Alpern, 1992) and games with a variegated environment and the possibility that the hiding player will be betrayed by ‘citizens’ of the space (Owen et al., 2008). Such games have also been used to model predator–prey interactions (Alpern, Fokkink, Timmer, & Casas, 2011a).

3.5. Allocation games

Allocation games are a related concept, in which the searching player does not move around the space individually, but rather distributes ‘search resources’ to locate the mobile hiding player. Such games may include false information (Hohzaki, 2007) and may incorporate spatial structure by allowing the influence of resources to spread across space (“reachability”), an area which has seen “little research” (Hohzaki, 2008).

Variations on this idea include situations in which searching resources are deployed sequentially (Dendris, Kirousis, & Thilikos, 1997) or in which both players distribute resources to respectively locate or protect a hidden object (Baston et al., 2000). Cooperative allocation games, in which multiple players combine their searching resources to locate a moving target, have also been considered (Hohzaki, 2009).

3.6. Rendez-vous games

Rendez-vous games are a parallel concept to games with mobile searching and hiding players, the difference being that these games are cooperative, with both players wishing to locate the other as soon as possible (see Alpern, 2002, for an overview). Typically, in a rendez-vous game, the structure of the space is known to all, with consideration given to the amount of information available to players regarding their relative positions, and their ability to distinguish between symmetries of the space (whether they have a common understanding of “North”, for example).

Rendez-vous games have been studied over various continuous one-dimensional spaces, such as the line (Alpern et al., 1995; Lim et al., 1996; Alpern et al., 2000) and the circle (Alpern, 2000), over continuous two-dimensional spaces, such as the plane (Kikuta et al., 2010) or a general compact metric space (Alpern, 1995) and over discrete spaces, such as lattices (Alpern et al., 2005; Ruckle, 2007) and other graphs (Alpern, Baston, & Essegai, 1999). Costs may also be introduced for movement and examination of particular locations (Kikuta et al., 2007).

Work has also been done on ‘hybrid’ games of search and rendez-vous, where, for example, two agents attempt to meet without being located by a third (Alpern et al., 1998) or where the searching player does not know whether the other player is attempting to rendez-vous or to evade capture (Alpern et al., 2002).

3.7. Security games

Security games are used to model situations in which some public resource (e.g. airports, transport infrastructure, power facilities) must be protected from attack with limited defensive resources. A good introduction to the topic is provided by Tambe (2012).

Such situations tend to be modelled as Stackelberg games, where it is assumed that the defensive player first commits to some strategy to protect the vulnerable sites and that this strategy is observed by the attacking player, who then chooses an optimal response (Tambe, 2012, pp. 4–8). Stackelberg-type security games related to the mobile-searcher-immobile-hider games of Section 3.2 have also been proposed to examine optimal patrolling strategies (Alpern, Morton, & Papadaki, 2011b; Basilico, Gatti, & Amigoni, 2012).

A related concept is that of the much studied Colonel Blotto game, in which two players must simultaneously distribute a fixed quantity of discrete or continuous resources across a number of sites, each site being ‘won’ by the player who distributed the greater quantity of resources to it, with payoffs determined by the number of sites that each player wins (see Roberson, 2006). The many extensions of the Blotto game have included asymmetric versions (Tofias, Merolla, & Munger, 2007; Hortala-Vallve & Llorente-Saguer, 2012), examples in which resources are allocated to battlefields sequentially rather than simultaneously (Powell, 2009) and examples in which defensive resources are heterogenous (Cohen, 1966).

Though the deployment sites in such models are often assumed to be wholly separate, with events at one location having no effect on events at other locations, certain security games and Blotto games with strategically interdependent sites have been considered. For example, Shubik and Weber (1981) introduce the concept of a “characteristic function” for such games, which allocates values to subsets of the sites, thus allowing interdependencies between them to be captured. Other approaches to modelling such interdependence include an extension of the Colonel Blotto game in which a successful attack on a “radar site” ensures the success of attacks on other sites (Grometstein & Shoham, 1989), while Hausken (2010, 2011) discusses a classification of the underlying infrastructures of security games based on the interdependence of their sites (e.g. series, parallel, complex, ...) and Powell (2007) analyses the relative value of defending borders over protecting strategic targets directly.

Though analyses of interdependence in security games and Blotto games may be quite general (that of Shubik & Weber, 1981, for example), interdependence that arises explicitly from the spatial structure of the deployment sites has not been considered in a general setting.

3.8. Geometric games

One of the most general and theoretical analyses of search and concealment type situations is *Geometric games and their applications* (Ruckle, 1983). In this book, the author defines a geometric game as a two-player zero-sum game (with players called “RED” and “BLUE”) played over a given set S , where the strategy sets for each player $\Sigma_{RED}, \Sigma_{BLUE}$ are subsets of the power set $\mathcal{P}(S)$ (the set of all subsets of S). Pure strategies for each player are therefore subsets $R, B \subseteq S$. The payoff to each player is a function of R and B , typically depending directly on the intersection $R \cap B$.

This concept of a geometric game allows Ruckle to model a wide variety of situations of search, ambush and pursuit, as well as a range of abstract games, taking full consideration of the structure of the space S over which the games are played.

Most published work based on Ruckle’s ideas (e.g. Baston, Bostock, & Ferguson, 1989; Zoroa et al., 1993; Zoroa, Fernández-

Sáez, & Zoroa, 1999; Zoroa, Zoroa, & Fernández-Sáez, 1999, 2001, 2003; Alpern, Baston, & Gal, 2009, 2010) has focused on specific examples of geometric games, rather than on general results.

3.9. Conclusion: motivations for the SSSG

Much of the literature on search games and related concepts has focussed on analysing specific games, rather than attempting to present general frameworks for such situations and identifying more broadly applicable results. While spatial structure may be considered for games in which players are mobile, the geography of the space over which games are played is often given little or limited consideration, particularly in the literature on security games. The concept of “reachability”, as described by Hohzaki (2008), in which a searcher or searching resource deployed at a point has influence over a neighbourhood of that point, has received very little attention. Games which concentrate purely on the strategic value of a player’s chosen position in a space, rather than on strategies for moving through the space for the purposes of search or rendez-vous, have also seen little research, at least since the work of Ruckle (1983).

4. The static spatial search game (SSSG)

The definitions and notation relating to metric spaces used in this section are from Sutherland (1975, pp. 19–44).

4.1. Definition of the SSSG

The static spatial search game (SSSG) is a two-player game played over a metric space $M = (\Omega, d)$, where Ω is a set of points x and $d : \Omega \times \Omega \rightarrow [0, \infty)$ is the metric or distance, which has the standard properties:

- (M1) $d(x, y) \geq 0$; $d(x, y) = 0 \iff x = y$
- (M2) $d(x, y) = d(y, x)$, $\forall x, y \in \Omega$
- (M3) $d(x, y) + d(y, z) \geq d(x, z)$, $\forall x, y, z \in \Omega$

The metric d reflects the spatial structure of Ω . In \mathbb{R}^n , d may be the Euclidean distance, while in a graph d may be the length of the shortest path connecting two points. However, depending on the interpretation of the game, d could also represent an abstract distance, indicating dissimilarity, difficulty of communication or perceived costs.

In specific cases, it may be sensible to relax some of these conditions. For example, in a graph that is not connected, we could allow infinite distances between vertices that are not connected by a path (an extended metric). Alternatively, to represent a directed graph, we may wish to ignore the symmetry condition (M2) (a quasi-metric; see Steen & Seebach (1970)). However, for the sake of simplicity, we do not consider such cases at this time.

We define a non-negative real number r called the detection radius and use the notation $B_r[x]$ to designate the closed ball centred on x :

$$B_r[x] = \{y \in \Omega : d(x, y) \leq r\}$$

The strategies for Player A (the searching player) and Player B (the concealing player) are specific points of Ω at which they may choose to deploy. In a single play of the game, each player simultaneously picks a point x_A, x_B from their own strategy set, $\Sigma_A, \Sigma_B \subseteq \Omega$.

For the sake of clarity, we use masculine pronouns to refer to Player A and feminine pronouns to refer to Player B throughout this paper.

We define the payoff functions for Player A and Player B respectively as:

$$p_A(x_A, x_B) = \begin{cases} 1, & x_B \in B_r[x_A] \\ 0, & \text{otherwise} \end{cases}$$

$$p_B(x_A, x_B) = 1 - p_A(x_A, x_B)$$

This is a constant-sum game and can be analysed accordingly.

In interpreting the game, we imagine that Player *B* chooses to hide somewhere in Ω , while Player *A* attempts to locate his opponent. To do this, Player *A* selects a point of Ω and searches a neighbourhood of this point. If Player *B*'s hiding place falls within the detection radius of Player *A*'s chosen point, the attempt to hide is unsuccessful and Player *B* is located. Otherwise, Player *B* remains undetected.

The game is illustrated in Fig. 1.

4.2. The SSSG and other games

The SSSG is not strictly a geometric game by Ruckle's definition (1983, p. 2), since it is not zero-sum. However, it could be transformed into a zero-sum game without altering the subsequent analysis, simply by subtracting $\frac{1}{2}$ from all payoffs. The decision that all payoffs should be 0 or 1 has been taken to ensure the clarity of the payoff matrices considered later in this paper.

Given this proviso, certain of Ruckle (1983)'s geometric games can be formulated as particular cases of the SSSG. For example, if transformed to a zero-sum game as described, game AAGV (Ruckle, 1983, p. 86; adapted from Arnold (1962)) is an example of the SSSG, with $\Omega = [0, 1] \subset \mathbb{R}$, $\Sigma_A = \Sigma_B = \Omega$ and $d(x, y) = |x - y|$.

Similarly, White (1994)'s *Games of strategy on trees* are examples of the SSSG, where Ω is the set of vertices of a tree, $\Sigma_A = \Sigma_B = \Omega$, $r = 1$ and d is the length of the shortest path between two vertices.

A game that demonstrates the potential complexity that can arise from apparently simple cases of the SSSG is the "Cookie-Cutter" game (or the "Hiding in a Disc" game), in which Player *A* chooses a point in a disc of unit radius and Player *B* simultaneously places a circular 'cookie-cutter' centred at any point of the disc, winning the game if Player *A*'s point lies within the 'cookie-cutter'. Given appropriate payoffs, this game is an example of the SSSG, where Ω is the closed unit disc, $\Sigma_A = \Sigma_B = \Omega$ and $d(\mathbf{x}, \mathbf{y}) = |\mathbf{x} - \mathbf{y}|$.

The particular case of this game with $r = 1/2$ was originally proposed by Gale et al. (1974), for which optimal mixed strategies were presented by Evans (1975). The game was extended to all $r > 0$ by Bordelon (1975), who proposed optimal mixed strategies for all $r > 1/2$, but these results were disputed by Ruckle (1983, p. 108). Ruckle's disproof was disputed in turn by Danskin

(1990), who showed that Bordelon (1975)'s results were correct for some values of $r > 1/2$, though false in general. Despite the apparent simplicity of the problem, Danskin (1990) was only able to find optimal mixed strategies for a small range of values of r around $r = 1/2$ and for all $r \geq \sqrt{2}/2$, thus illustrating the hidden complexity of many games of this form.

A particularly simple example of a game that can be represented as an SSSG is "Matching Pennies" (see Blackwell & Girshick, 1954, p. 13), in which Players *A* and *B* simultaneously call "Heads" or "Tails", with Player *A* receiving a payoff of 1 if the calls are the same and -1 otherwise, and Player *B* receiving a payoff of -1 if the calls are the same and 1 otherwise. Taking $\Sigma_A = \Sigma_B = \Omega = \{\text{Heads}, \text{Tails}\}$, $r = 0$, with d as any valid metric, this is an SSSG, again with the proviso that the payoffs must be transformed appropriately.

A more complicated example of a game that can be represented as an SSSG is the graph security game of Mavronicolas, Papadopoulos, Philippou, and Spirakis (2008) if the number of attackers is restricted to one. This game is played over an undirected graph $G = (V(G), E(G))$ with one defender and (in general) multiple attackers. Simultaneously, the defender chooses an edge and the attackers each choose a vertex. The defender receives a payoff equal to the number of attackers who choose vertices incident to his chosen edge. Each attacker receives a payoff equal to 0 if their chosen vertex is incident to the defender's edge and 1 otherwise.

Consider the graph G' obtained by inserting a new vertex at the midpoint of each of the edges of G . Let the set of new vertices created in this way be denoted $V(G')^*$ while the complete vertex set $V(G')$ includes both the new vertices and the original vertices. With the defender as Player *A*, a single attacker as Player *B*, $\Omega = V(G')$, $\Sigma_A = V(G')^*$, $\Sigma_B = V(G)$, $r = 1$ and with d as the length of the shortest path between two vertices in G' , this game is also an example of the SSSG.

The SSSG provides a framework that unites all of these games and allows for a general consideration of the relative strategic values of the different points of a space. It implicitly encompasses the concepts of reachability, interdependence based on spatial structure and the detection radius, as discussed in Section 3.

4.3. The SSSG with finite strategy sets

Consider an example of the SSSG in which the strategy sets Σ_A, Σ_B are finite. One of the simplest possible mixed strategies available to each player in such a case is the mixed strategy that allocates equal probabilities to all points in a player's strategy set. We denote these mixed strategies by ρ_A and ρ_B for Players *A* and *B* respectively.

The following proposition establishes a sufficient condition for ρ_A, ρ_B to be optimal mixed strategies:

Proposition 4.1. Consider the SSSG played over a metric space Ω , with finite strategy sets Σ_A, Σ_B and distance d . If there exists a positive integer ζ such that:

$$\begin{aligned} |B_r[x_A] \cap \Sigma_B| &= \zeta, \quad \forall x_A \in \Sigma_A \\ |B_r[x_B] \cap \Sigma_A| &= \zeta, \quad \forall x_B \in \Sigma_B \end{aligned} \tag{2}$$

then ρ_A and ρ_B are optimal mixed strategies for Players *A* and *B* respectively and $\zeta|\Sigma_A|^{-1} = \zeta|\Sigma_B|^{-1} = u$ is the value of the game to Player *A*.

Proof. Suppose that Player *A* employs the mixed strategy ρ_A which allocates a uniform probability of $|\Sigma_A|^{-1}$ to all points $x_A \in \Sigma_A$. In a particular play of the game, suppose that Player *B* deploys at point $x_B \in \Sigma_B$.




Fig. 1. The general SSSG. Players deploy simultaneously at points in some metric space Ω . Suppose Player *A* deploys at x and can search all points within a radius of r . Player *B* loses the game if deployed at y , but wins if deployed at z . A win results in a payoff of 1; a loss results in a payoff of 0.

In this situation, since $|B_r[x_B] \cap \Sigma_A| = \zeta$, the expected payoff to Player A is $\zeta|\Sigma_A|^{-1}$, the probability that Player A's point x_A lies in $B_r[x_B]$. Therefore, for any mixed strategy σ_B for Player B:

$$E[p_A(\rho_A, \sigma_B)] = \zeta|\Sigma_A|^{-1}$$

and thus:

$$u = \max_{\sigma_A} \min_{\sigma_B} E[p_A(\sigma_A, \sigma_B)] \geq \zeta|\Sigma_A|^{-1} \tag{3}$$

Now suppose that Player B employs the mixed strategy ρ_B which allocates a uniform probability of $|\Sigma_B|^{-1}$ to all points $x_B \in \Sigma_B$. In a particular play of the game, suppose that Player A deploys at point $x_A \in \Sigma_A$.

In this situation, since $|B_r[x_A] \cap \Sigma_B| = \zeta$, the expected payoff to Player A is $\zeta|\Sigma_B|^{-1}$, the probability that Player B's point x_B lies in $B_r[x_A]$. Therefore, for any mixed strategy σ_A for Player A:

$$E[p_A(\sigma_A, \rho_B)] = \zeta|\Sigma_B|^{-1}$$

and thus:

$$u = \max_{\sigma_A} \min_{\sigma_B} E[p_A(\sigma_A, \sigma_B)] \leq \zeta|\Sigma_B|^{-1} \tag{4}$$

Now, (2) together with the symmetric property of the distance (M2) imply that $\zeta|\Sigma_A| = \zeta|\Sigma_B|$ and thus that $|\Sigma_A| = |\Sigma_B|$. By (3) and (4), we therefore have:

$$|\Sigma_A|^{-1} \leq u \leq |\Sigma_B|^{-1} \Rightarrow u = |\Sigma_A|^{-1} = |\Sigma_B|^{-1}$$

and ρ_A, ρ_B are optimal mixed strategies, by Definition 2.3. \square

4.4. Dominance and equivalence in the SSSG

We can now examine strategic dominance and equivalence (see Definition 2.4) in the context of the SSSG using the notation established in Section 4.

Proposition 4.2. Consider the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d . For strategies $x_1, x_2 \in \Sigma_A, x_1 \neq x_2$, for Player A:

- x_2 very weakly dominates x_1 if and only if:

$$(B_r[x_1] \cap \Sigma_B) \subseteq (B_r[x_2] \cap \Sigma_B)$$

- x_2 weakly dominates x_1 if and only if:

$$(B_r[x_1] \cap \Sigma_B) \subset (B_r[x_2] \cap \Sigma_B)$$

- x_2 strictly dominates x_1 if and only if:

$$(B_r[x_1] \cap \Sigma_B) = \emptyset$$

$$(B_r[x_2] \cap \Sigma_B) = \Sigma_B$$

- x_2 is equivalent to x_1 if and only if:

$$(B_r[x_1] \cap \Sigma_B) = (B_r[x_2] \cap \Sigma_B)$$

This proposition states that for Player A:

- x_2 very weakly dominates x_1 if and only if, when deployed at x_2 , Player A can search every potential location of Player B that could be searched from x_1 .
- This dominance is weak if there exist potential locations of Player B that can be searched from x_2 but that cannot be searched from x_1 (inclusion is strict).
- Strict dominance only occurs in the trivial case in which no potential locations of Player B can be searched from x_1 while every potential location of Player B can be searched from x_2 .
- x_2 and x_1 are equivalent if and only if precisely the same set of potential locations of Player B can be searched from both points.

Proof. We consider each of the four parts of Definition 2.4 and show that, in the context of the SSSG, they are equivalent to the corresponding statements of Proposition 4.2. Recall that p_A takes values in $\{0, 1\}$.

- Very weak dominance

$$p_A(x_2, y) \geq p_A(x_1, y), \quad \forall y \in \Sigma_B \tag{*}$$

$$\iff [p_A(x_1, y) = 1 \Rightarrow p_A(x_2, y) = 1], \quad \forall y \in \Sigma_B$$

$$\iff [y \in B_r[x_1] \Rightarrow y \in B_r[x_2]], \quad \forall y \in \Sigma_B$$

$$\iff (B_r[x_1] \cap \Sigma_B) \subseteq (B_r[x_2] \cap \Sigma_B) \tag{**}$$

- Weak dominance

Since $(*) \iff (**)$, it suffices to observe that, if $(**)$ is assumed to be true:

$$\exists y^* \in \Sigma_B : p_A(x_2, y^*) > p_A(x_1, y^*)$$

$$\iff \exists y^* \in \Sigma_B : y^* \in B_r[x_2] \text{ and } y^* \notin B_r[x_1]$$

$$\iff (B_r[x_1] \cap \Sigma_B) \subset (B_r[x_2] \cap \Sigma_B) \quad \text{[by (**)]}$$

- Strict dominance

$$p_A(x_2, y) > p_A(x_1, y), \quad \forall y \in \Sigma_B$$

$$\iff y \in B_r[x_2] \text{ and } y \notin B_r[x_1], \quad \forall y \in \Sigma_B$$

$$\iff \begin{cases} (B_r[x_1] \cap \Sigma_B) = \emptyset, \\ (B_r[x_2] \cap \Sigma_B) = \Sigma_B \end{cases}$$

- Equivalence

$$p_A(x_2, y) = p_A(x_1, y), \quad \forall y \in \Sigma_B$$

$$\iff [y \in B_r[x_2] \iff y \in B_r[x_1]], \quad \forall y \in \Sigma_B$$

$$\iff (B_r[x_2] \cap \Sigma_B) = (B_r[x_1] \cap \Sigma_B) \quad \square$$

We now consider dominance and equivalence for Player B:

Proposition 4.3. Consider the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d . For strategies $x_1, x_2 \in \Sigma_B, x_1 \neq x_2$, for Player B:

- x_2 very weakly dominates x_1 if and only if:

$$[x_2 \in B_r[y] \Rightarrow x_1 \in B_r[y]], \quad \forall y \in \Sigma_A$$

- x_2 weakly dominates x_1 if and only if:

$$[x_2 \in B_r[y] \Rightarrow x_1 \in B_r[y]], \quad \forall y \in \Sigma_A$$

and $\exists y^* \in \Sigma_A$ such that:

$$x_1 \in B_r[y^*] \text{ and } x_2 \notin B_r[y^*]$$

- x_2 strictly dominates x_1 if and only if:

$$x_1 \in B_r[y] \text{ and } x_2 \notin B_r[y], \quad \forall y \in \Sigma_A$$

- x_2 is equivalent to x_1 if and only if:

$$[x_2 \in B_r[y] \iff x_1 \in B_r[y]], \quad \forall y \in \Sigma_A$$

This proposition states that for Player B:

- x_2 very weakly dominates x_1 if and only if, wherever Player A deploys, if he can search x_2 then he can also search x_1 .
- This dominance is weak if there exists a position for Player A from which he can search x_1 but cannot search x_2 .
- Strict dominance only occurs in the trivial case in which, wherever Player A deploys, he can search x_1 but cannot search x_2 .
- x_2 and x_1 are equivalent if and only if, wherever Player A deploys, he can search x_2 if and only if he can search x_1 .

Proof. We consider each of the four parts of Definition 2.4 and show that, in the context of the SSSG, they are equivalent to the

corresponding statements of Proposition 4.3. Recall that p_B also takes values in $\{0, 1\}$.

• *Very weak dominance*

$$\begin{aligned} p_B(y, x_2) &\geq p_B(y, x_1), & \forall y \in \Sigma_A & \quad (*) \\ \iff [p_B(y, x_1) = 1 \Rightarrow p_B(y, x_2) = 1], & \forall y \in \Sigma_A \\ \iff [x_1 \notin B_r[y] \Rightarrow x_2 \notin B_r[y]], & \forall y \in \Sigma_A \\ \iff [x_2 \in B_r[y] \Rightarrow x_1 \in B_r[y]], & \forall y \in \Sigma_A & \quad (**) \end{aligned}$$

• *Weak dominance*

Since $(*) \iff (**)$, it suffices to observe that:

$$\begin{aligned} \exists y^* \in \Sigma_A : p_B(y^*, x_2) &> p_B(y^*, x_1) \\ \iff \exists y^* \in \Sigma_A : x_1 \in B_r[y^*] &\text{ and } x_2 \notin B_r[y^*] \end{aligned}$$

• *Strict dominance*

$$\begin{aligned} p_B(y, x_2) &> p_B(y, x_1), & \forall y \in \Sigma_A \\ \iff x_1 \in B_r[y] &\text{ and } x_2 \notin B_r[y], & \forall y \in \Sigma_A \end{aligned}$$

• *Equivalence*

$$\begin{aligned} p_B(y, x_2) &= p_B(y, x_1), & \forall y \in \Sigma_A \\ \iff [x_2 \in B_r[y] \iff x_1 \in B_r[y]], & \forall y \in \Sigma_A \quad \square \end{aligned}$$

The necessary and sufficient conditions for dominance and equivalence for Player B established in Proposition 4.3 can be shown to be equivalent to a simpler set of conditions, clearly analogous to those relating to dominance and equivalence for Player A seen in Proposition 4.2:

Proposition 4.4. Consider the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d . For strategies $x_1, x_2 \in \Sigma_B, x_1 \neq x_2$, for Player B:

• x_2 very weakly dominates x_1 if and only if:

$$(B_r[x_2] \cap \Sigma_A) \subseteq (B_r[x_1] \cap \Sigma_A)$$

• x_2 weakly dominates x_1 if and only if:

$$(B_r[x_2] \cap \Sigma_A) \subset (B_r[x_1] \cap \Sigma_A)$$

• x_2 strictly dominates x_1 if and only if:

$$\begin{aligned} (B_r[x_1] \cap \Sigma_A) &= \Sigma_A \\ (B_r[x_2] \cap \Sigma_A) &= \emptyset \end{aligned}$$

• x_2 is equivalent to x_1 if and only if:

$$(B_r[x_2] \cap \Sigma_A) = (B_r[x_1] \cap \Sigma_A)$$

Proof. We consider each of the four statements of Proposition 4.3 (which has already been proven) and show that they are equivalent to the corresponding statements of Proposition 4.4.

• *Very weak dominance*

$$\begin{aligned} [x_2 \in B_r[y] \Rightarrow x_1 \in B_r[y]], & \forall y \in \Sigma_A & \quad (*) \\ \iff [y \in B_r[x_2] \Rightarrow y \in B_r[x_1]], & \forall y \in \Sigma_A & \quad \text{[by (M2)]} \\ \iff (B_r[x_2] \cap \Sigma_A) \subseteq (B_r[x_1] \cap \Sigma_A) & \quad (**) \end{aligned}$$

• *Weak dominance*

Since $(*) \iff (**)$, it suffices to observe that, if $(**)$ is assumed to be true:

$$\begin{aligned} \exists z \in \Sigma_A : x_1 \in B_r[z] &\text{ and } x_2 \notin B_r[z] \\ \iff \exists z \in \Sigma_A : z \in B_r[x_1] &\text{ and } z \notin B_r[x_2] & \quad \text{[by (M2)]} \\ \iff (B_r[x_2] \cap \Sigma_A) \subset (B_r[x_1] \cap \Sigma_A) & \quad \text{[by (**)]} \end{aligned}$$

• *Strict dominance*

$$\begin{aligned} x_1 \in B_r[y] &\text{ and } x_2 \notin B_r[y], & \forall y \in \Sigma_A \\ \iff y \in B_r[x_1] &\text{ and } y \notin B_r[x_2], & \forall y \in \Sigma_A & \quad \text{[by (M2)]} \\ \iff \begin{cases} (B_r[x_1] \cap \Sigma_A) = \Sigma_A, \\ (B_r[x_2] \cap \Sigma_A) = \emptyset \end{cases} \end{aligned}$$

• *Equivalence*

$$\begin{aligned} [x_2 \in B_r[y] \iff x_1 \in B_r[y]], & \forall y \in \Sigma_A \\ \iff [y \in B_r[x_2] \iff y \in B_r[x_1]], & \forall y \in \Sigma_A & \quad \text{[by (M2)]} \\ \iff (B_r[x_2] \cap \Sigma_A) = (B_r[x_1] \cap \Sigma_A) & \quad \square \end{aligned}$$

While Proposition 4.4 is apparently simpler than Proposition 4.3, note that every part of its proof depends on the symmetric property of the distance (M2). If this condition were to be relaxed, as discussed in Section 4.1, Proposition 4.4 would not be valid and dominance and equivalence for Player B would have to be analysed on the basis of Proposition 4.3.

Definition 4.1. Pairwise Equivalence

Consider the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d . For Player A or Player B, a subset of their strategy set $\hat{\Sigma} \subseteq \Sigma_A$ or $\hat{\Sigma} \subseteq \Sigma_B$ exhibits pairwise equivalence if and only if x is equivalent to $y, \forall x, y \in \hat{\Sigma}$.

We conclude that a subset $\hat{\Sigma} \subseteq \Sigma_A$ or $\hat{\Sigma} \subseteq \Sigma_B$ exhibiting pairwise equivalence can be reduced to any singleton $\{\hat{x}\} \subseteq \hat{\Sigma}$ without altering the analysis of the game. Since all points in $\hat{\Sigma}$ are equivalent, a player would neither gain nor lose by playing another point in the set over \hat{x} .

The following proposition states that if x_2 very weakly dominates x_1 for Player A (and x_1 is adjacent to at least one potential location for Player B), then the distance between the two points must be no greater than $2r$.

Proposition 4.5. For the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d , if x_2 very weakly dominates x_1 for Player A and $B_r[x_1] \cap \Sigma_B \neq \emptyset$, then $x_2 \in B_{2r}[x_1] \cap \Sigma_A$.

Proof. If x_2 very weakly dominates x_1 for Player A and $B_r[x_1] \cap \Sigma_B \neq \emptyset$, then:

$$\begin{aligned} \emptyset \neq (B_r[x_1] \cap \Sigma_B) &\subseteq (B_r[x_2] \cap \Sigma_B) & \quad \text{[by 4.2]} \\ \Rightarrow \exists y \in B_r[x_1] \cap B_r[x_2] \\ \Rightarrow d(x_1, y) \leq r &\text{ and } d(x_2, y) \leq r & \quad \square \\ \Rightarrow d(x_1, x_2) \leq 2r & & \quad \text{[by (M2), (M3)]} \\ \Rightarrow x_2 \in B_{2r}[x_1] \cap \Sigma_A \end{aligned}$$

The condition that $B_r[x_1] \cap \Sigma_B \neq \emptyset$ simply removes trivial strategies that are very weakly dominated by every other strategy.

An analogous result holds for Player B. The proof is similar to that of 4.5 and is therefore omitted:

Proposition 4.6. For the SSSG played over a metric space Ω , with strategy sets Σ_A, Σ_B and distance d , if x_2 very weakly dominates x_1 for Player B and $B_r[x_2] \cap \Sigma_A \neq \emptyset$, then $x_2 \in B_{2r}[x_1] \cap \Sigma_B$.

In this case, the condition that $B_r[x_2] \cap \Sigma_A \neq \emptyset$ removes strategies that very weakly dominate every other strategy.

Note that both of these propositions depend on the symmetric property of the distance (M2) and the triangle inequality (M3).

4.5. Iterated elimination of dominated strategies

The concepts of dominance and equivalence provide us with a method for reducing the SSSG through an iterative process of

removing dominated strategies from Σ_A and Σ_B , reducing pairwise equivalent subsets to singletons and reassessing dominance in the new strategy sets. This is known as the iterated elimination of dominated strategies (IEDS) (see, for example, Berwanger, 2007; Börgers, 1992; Dufwenberg et al., 2002). Given any game, the aim of IEDS is to identify a simplified game, whose solutions are also solutions of the complete game. These solutions can then be identified using standard techniques (see, for example, Morris (1994, pp. 99–114)). The application of this method to games played over graphs is discussed in Section 6.

It should be noted that because we are considering weak rather than strict dominance, IEDS may not be suitable for identifying all the solutions of a particular game. The results of this form of IEDS are dependent on the order in which dominated strategies are removed (Leyton-Brown & Shoham, 2008, pp. 20–23) and some solutions may be lost. It is also necessary to observe that, while IEDS has been shown to be valid for games with finitely many possible strategies, for games with infinitely many possible strategies the process may fail (Berwanger, 2007). Indeed, such infinite games may not have solutions (Ruckle, 1983, p. 10).

However, though IEDS is not guaranteed to produce optimal mixed strategies for the SSSG in such cases, given a pair of mixed strategies σ_A, σ_B obtained by this method, it is straightforward to check whether or not they are optimal by verifying that, for some $u \in [0, 1]$, we have:

$$\begin{aligned} \inf_{x \in \Sigma_B} E[p_A(\sigma_A, x)] &= u \\ \inf_{y \in \Sigma_A} E[p_B(y, \sigma_B)] &= 1 - u \end{aligned} \tag{5}$$

where u is the value of the game to Player A.

This method is described by Blackwell and Girshick (1954, p. 60) and used extensively by Ruckle (1983) to verify proposed optimal mixed strategies for geometric games.

5. The SSSG on a graph

The definitions and notation relating to graph theory used in this section are adapted from Bondy and Murty (1976).

5.1. Definition of the SSSG on a graph

Consider a simple undirected graph G , characterised by the symmetric adjacency matrix $M = (a_{ij})$, with a set of κ vertices:

$$V(G) = \{v_1, \dots, v_\kappa\}$$

and a set of edges:

$$E(G) = \{\{v_i, v_j\} : v_i, v_j \in V(G) \text{ and } a_{ij} = 1\}$$

We suppose that G is connected, to ensure that the metric space axioms are fulfilled, but this assumption could be relaxed if we allowed for infinite distances between vertices. We also suppose that all edges of G have unit weight.

The graph game $\mathbb{G} = (G, \Sigma_A, \Sigma_B, r)$ over a finite graph G is defined to be an example of the SSSG, with $\Omega = V(G), \emptyset \neq \Sigma_A, \Sigma_B \subseteq V(G), r$ a positive real number, and the distance function $d_G(v, w)$ for $v, w \in V(G), v \neq w$, defined to be the length of the shortest path from v to w in G . We also define $d_G(v, v) = 0, \forall v \in V(G)$. The assumption that G is undirected ensures that the symmetry condition $d_G(v, w) = d_G(w, v)$ holds $\forall v, w \in V(G)$.

Setting $r = 1$, we have that $B_r[v]$, the zone that can be searched by Player A when deployed at v , is the closed neighbourhood $N[v]$ of v , the set of all vertices adjacent to v united with $\{v\}$ itself.

The game proceeds by Player A choosing a vertex $v_A \in \Sigma_A$ and Player B simultaneously choosing a vertex $v_B \in \Sigma_B$. The payoff functions are:

$$p_A(v_A, v_B) = \begin{cases} 1, & \text{if } v_A \text{ and } v_B \text{ are equal or adjacent;} \\ 0, & \text{otherwise.} \end{cases}$$

$$p_B(v_A, v_B) = 1 - p_A(v_A, v_B)$$

In this game, the pure strategies for each player are particular vertices. In the following analysis, we use the words “strategy” and “vertex” interchangeably, depending on the context.

The graph game is illustrated in Fig. 2.

5.2. Graph games with $r \neq 1$

The restriction to $r = 1$ is not a significant constraint, since any graph game can be reduced to this case by means of a minor alteration.

For situations with $r \neq 1$, we can define:

$$G' = (V(G), E(G'))$$

such that:

$$E(G') = \{\{v_i, v_j\} : v_i, v_j \in V(G) \text{ and } 1 \leq d_G(v_i, v_j) \leq r\}$$

and apply our analysis to G' with $r = 1$.

Equivalently, if $r \in \mathbb{N}$, we can replace the adjacency matrix $M = (a_{ij})$ of G with $M' = (a'_{ij})$, where:

$$a'_{ij} = \begin{cases} 0, & \text{if } b_{ij} = 0 \text{ or } i = j \\ 1, & \text{otherwise.} \end{cases}$$

with $(b_{ij}) = \sum_{q=1}^r M^q$, the matrix which shows the number of paths in G of length no greater than r connecting v_i to v_j .

It therefore suffices to exclusively study graph games \mathbb{G} with $r = 1$, since these methods for redefining G ensure that such analysis will be applicable to games for any $r \in \mathbb{N}$.

5.3. Preliminary observations

Analysis of the graph game requires the statement of some preliminary results and definitions. For these results, we use the following notation:

- $N_A[v] = N[v] \cap \Sigma_A$ Set of all potential positions for Player A in $N[v]$.
- $N_B[v] = N[v] \cap \Sigma_B$ Set of all potential positions for Player B in $N[v]$.
- $\alpha[v] = |N_A[v]|$ Number of potential positions for Player A in $N[v]$.
- $\beta[v] = |N_B[v]|$ Number of potential positions for Player B in $N[v]$.




Fig. 2. The SSSG on a graph. Players deploy simultaneously at vertices of the graph. Suppose that Player A deploys at the vertex marked w. If $r = 1$, Player B loses the game if she deploys at any of the vertices marked w or x, but wins if she deploys at any of the vertices marked y or z. If $r = 2$, Player B loses the game if she deploys at any of the vertices marked w, x or y, but wins if she deploys at either of the vertices marked z. A win results in a payoff of 1; a loss results in a payoff of 0.

- $\Delta(G) = \max_{w \in V(G)} [\text{deg}(w)]$ Maximum degree of the vertices of G .
- $\delta(G) = \min_{w \in V(G)} [\text{deg}(w)]$ Minimum degree of the vertices of G .

The following proposition states that if Player A can search a globally maximal number of potential positions for Player B from a vertex v , then v is not dominated by any other vertex and is only equivalent to those vertices which have the same closed neighbourhood as v . In a graph in which all vertices have a distinct closed neighbourhood, such as a rectangular grid graph, such a vertex v cannot be very weakly dominated by any other vertex.

Proposition 5.1. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, consider a vertex $v \in \Sigma_A$ and the subset $\Sigma_A^{(v)} = \{z \in \Sigma_A : N[z] = N[v]\}$. We have that if:

$$\beta[v] = \Delta(G) + 1 \tag{6}$$

then:

- (i) $\Sigma_A^{(v)}$ exhibits pairwise equivalence for Player A ;
- (ii) v is not very weakly dominated for Player A by any strategy in $\Sigma_A \setminus \Sigma_A^{(v)}$.

Proof. To prove (i), observe that $\forall w_1, w_2 \in \Sigma_A^{(v)}, N[w_1] = N[w_2]$. Hence, $N_B[w_1] = N_B[w_2]$ and therefore w_1 is equivalent to w_2 for Player A , so $\Sigma_A^{(v)}$ exhibits pairwise equivalence for Player A .

To prove (ii), suppose for a contradiction that v satisfies (6) and is very weakly dominated by $\hat{v} \in \Sigma_A \setminus \Sigma_A^{(v)}$.

Observe also that:

$$\text{If } \beta[w] = \Delta(G) + 1, \text{ for some } w \in V(G), \tag{7}$$

$$\text{then } N_B[w] = N[w].$$

and note that $\Delta(G) + 1$ is an upper bound for $\beta[w]$.

Now, we have that:

$$\begin{aligned} N_B[v] &\subseteq N_B[\hat{v}] \quad [\text{by very weak dominance}] \\ \Rightarrow \Delta(G) + 1 &\leq \beta[\hat{v}] \leq |N[\hat{v}]| \quad [\text{by (6)}] \\ \Rightarrow \Delta(G) + 1 &\leq \beta[\hat{v}] \leq \Delta(G) + 1 \\ \Rightarrow \Delta(G) + 1 &= \beta[v] = \beta[\hat{v}] \\ \Rightarrow N_B[v] &= N_B[\hat{v}] \quad [\text{since } N_B[v] \subseteq N_B[\hat{v}]] \\ \Rightarrow N[v] &= N[\hat{v}] \quad [\text{by (7)}] \\ \Rightarrow \hat{v} &\in \Sigma_A^{(v)} \end{aligned}$$

This is a contradiction, since we supposed that $\hat{v} \in \Sigma_A \setminus \Sigma_A^{(v)}$. \square

The next proposition is a stronger result. It states that, for Player A , given a vertex v that is known not to be very weakly dominated by any vertices outside of a certain subset S , if Player A can search strictly more potential hiding places for Player B when deployed at v than could be searched from any other vertex in S , then v is not very weakly dominated by any other vertex.

Proposition 5.2. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, consider a vertex $v \in \Sigma_A$. Let $S \subseteq \Sigma_A$ be such that $v \in S$ and such that $\Sigma_A \setminus S$ contains no vertices that very weakly dominate v for Player A . We have that if:

$$\beta[v] > \max_{w \in S \setminus \{v\}} (\beta[w]) \tag{8}$$

then v is not very weakly dominated for Player A by any other strategy in Σ_A .

Proof. For a contradiction, suppose that v satisfies (8) and is very weakly dominated by $\hat{v} \in \Sigma_A$ for Player A . We must have that $\hat{v} \in S$, since vertices outside S cannot very weakly dominate v .

From the very weak dominance, we have:

$$\begin{aligned} N_B[v] &\subseteq N_B[\hat{v}] \\ \Rightarrow \max_{w \in S \setminus \{v\}} (\beta[w]) &< \beta[v] \leq \beta[\hat{v}] \quad [\text{by (8)}] \end{aligned}$$

This is a contradiction, since $\hat{v} \in S \setminus \{v\}$. \square

The following proposition is a restatement of Proposition 4.5, reformulated in the context of the graph game. It states that any vertex that very weakly dominates v for Player A can be no more than 2 steps away from v on the graph. Its proof is identical to that of Proposition 4.5 and is thus omitted.

Proposition 5.3. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, if w very weakly dominates v for Player A and $N_B[v] \neq \emptyset$, then:

$$w \in B_2[v] \cap \Sigma_A$$

Corollary 1. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, consider a vertex $v \in \Sigma_A$. We have that if:

$$\beta[v] > \max_{w \in S'} (\beta[w])$$

where:

$$S' = B_2[v] \cap \Sigma_A \setminus \{v\}$$

then v is not very weakly dominated for Player A by any other strategy.

Proof. This follows directly from Propositions 5.2 and 5.3, where the set S from Proposition 5.3 is defined as:

$$S = B_2[v] \cap \Sigma_A \quad \square$$

The corollary states that any vertex from which Player A can search strictly more potential hiding places for Player B than could be searched from any other valid vertex lying no more than two steps away, cannot be very weakly dominated by any other vertex. These results allow us to considerably narrow down our search for dominated and equivalent vertices.

Analogous results to Propositions 5.2 and 5.3 and Corollary 1 hold for very weak dominance for Player B . The proofs of these results are similar to those presented above and are thus omitted.

Proposition 5.4. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, consider a vertex $v \in \Sigma_B$. Let $S \subseteq \Sigma_B$ be such that $v \in S$ and such that $\Sigma_B \setminus S$ contains no vertices that very weakly dominate v for Player B . We have that if:

$$\alpha[v] < \min_{w \in S \setminus \{v\}} (\alpha[w])$$

then v is not very weakly dominated for Player B by any other strategy in Σ_B .

Proposition 5.5. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, if w very weakly dominates v for Player B and $N_A[w] \neq \emptyset$, then:

$$w \in B_2[v] \cap \Sigma_B$$

Corollary 2. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, consider a vertex $v \in \Sigma_B, \alpha[v] \neq 0$. We have that if:

$$\alpha[v] < \min_{w \in S'} (\alpha[w])$$

where:

$$S' = B_2[v] \cap \Sigma_B \setminus \{v\}$$

then v is not very weakly dominated for Player B by any other strategy.

5.4. Bounds on the value of a graph game

A first step in the analysis of a particular graph game is to determine lower and upper bounds on the values of the game (see Definition 2.2).

Recall that for a two-player constant-sum game, it makes sense to restrict discussion to the value to Player A, since this also determines the value to Player B through the condition that the sum of the two values is a fixed constant (see (1)).

Proposition 5.6. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$, let u represent the value to Player A. u is bounded as follows:

$$\min_{v \in \Sigma_B} \frac{\alpha[v]}{|\Sigma_A|} \leq u \leq \max_{v \in \Sigma_A} \frac{\beta[v]}{|\Sigma_B|}$$

Note that in the case where $\Sigma_A = \Sigma_B = V(G)$, these inequalities become:

$$\frac{\delta(G) + 1}{|V(G)|} \leq u \leq \frac{\Delta(G) + 1}{|V(G)|} \tag{9}$$

The proposition derives from a consideration of ρ_A and ρ_B , defined in Section 4.3 as the mixed strategies that allocate equal probabilities to all vertices in a player’s strategy set. If Player A employs mixed strategy ρ_A , then Player B can do no better than to deploy at the vertex whose closed neighbourhood contains the fewest possible vertices in Σ_A . The value of the game to Player A cannot therefore be less than the sum of the probabilities that ρ_A assigns to these vertices. This reasoning produces the left hand inequality. The right hand inequality follows in a similar fashion from an analysis of ρ_B as a strategy for Player B. A formal proof follows:

Proof. Suppose that Player A employs the mixed strategy ρ_A which allocates a uniform probability of $|\Sigma_A|^{-1}$ to all vertices $w \in \Sigma_A$. In a particular play of the game, suppose that Player B deploys at vertex $v \in \Sigma_B$.

In this situation, the expected payoff to Player A is $\alpha[v]|\Sigma_A|^{-1}$. Therefore, for any mixed strategy σ_B for Player B:

$$E[p_A(\rho_A, \sigma_B)] \geq \min_{v \in \Sigma_B} \frac{\alpha[v]}{|\Sigma_A|}$$

and thus:

$$u = \max_{\sigma_A} \min_{\sigma_B} E[p_A(\sigma_A, \sigma_B)] \geq \min_{v \in \Sigma_B} \frac{\alpha[v]}{|\Sigma_A|}$$

The proof of the right hand inequality is similar. \square

The following corollary is a consequence of (9):

Corollary 3. Consider the graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

where $G = (V(G), E(G))$ is a regular graph of degree D , $\Sigma_A = \Sigma_B = V(G)$, $r = 1$ and let u be the value of \mathfrak{G} to Player A. Then:

- $u = (D + 1)/|V(G)|$
- The strategy ρ that allocates a uniform probability of $|V(G)|^{-1}$ to all vertices, is an optimal mixed strategy for both players.

u can also be bounded in a different way:

Proposition 5.7. For the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$ with $r = 1$, let u represent the value to Player A and:

$$W = \left\{ W' \subseteq \Sigma_A : \bigcup_{w \in W'} N[w] \supseteq \Sigma_B \right\}$$

$$Z = \{ Z' \subseteq \Sigma_B : d_G(z_1, z_2) > 2, \forall z_1, z_2 \in Z' \}$$

Then u is bounded as follows:

$$\left[\min_{W' \in W} |W'| \right]^{-1} \leq u \leq \left[\max_{Z' \in Z} |Z'| \right]^{-1}$$

The left hand inequality is derived from consideration of the mixed strategy τ_A for Player A that allocates uniform probabilities to a minimal subset of vertices whose closed neighbourhoods cover Σ_B .

The right hand inequality is derived from consideration of the mixed strategy τ_B for Player B that allocates uniform probabilities to a maximal subset of vertices with the property that no two vertices are connected by a path of length less than 3.

Proof. First consider the left hand inequality.

Consider a subset of vertices $W' \in W$ of minimum cardinality. Suppose that Player A employs the mixed strategy τ_A that allocates uniform probability $|W'|^{-1}$ to vertices $w \in W'$ and zero probability to all other vertices.

In a particular play of the game, suppose that Player B deploys at vertex $v \in \Sigma_B$. Since $W' \in W$, we have:

$$\bigcup_{w \in W'} N[w] \supseteq \Sigma_B$$

Therefore $\exists w' \in W'$ such that $v \in N[w']$. Since τ_A allocates a probability of $|W'|^{-1}$ to w' , the expected payoff to Player A is greater than or equal to $|W'|^{-1}$. Therefore, for any mixed strategy σ_B for Player B:

$$E[p_A(\tau_A, \sigma_B)] \geq \left[\min_{W' \in W} |W'| \right]^{-1}$$

and thus:

$$u = \max_{\sigma_A} \min_{\sigma_B} E[p_A(\sigma_A, \sigma_B)] \geq \left[\min_{W' \in W} |W'| \right]^{-1}$$

Now consider the right hand inequality.

Consider a subset of vertices $Z' \in Z$ of maximum cardinality. Suppose that Player B employs the mixed strategy τ_B that allocates uniform probability $|Z'|^{-1}$ to vertices $v \in Z'$ and zero probability to all other vertices.

In a particular play of the game, suppose that Player A deploys at vertex $w \in \Sigma_A$. Since $d_G(z_1, z_2) > 2, \forall z_1, z_2 \in Z'$, we clearly have that:

$$|N[w] \cap Z'| \leq 1$$

So, in this situation, the expected payoff to Player A is less than or equal to $|Z'|^{-1}$. Therefore, for any mixed strategy σ_B for Player A:

$$E[p_A(\sigma_A, \tau_B)] \leq \left[\max_{Z' \in Z} |Z'| \right]^{-1}$$

and thus:

$$u = \max_{\sigma_A} \min_{\sigma_B} E[p_A(\sigma_A, \sigma_B)] \leq \left[\max_{Z' \in Z} |Z'| \right]^{-1} \quad \square$$

Following these results, we label the bounds on u as follows:

- $LB_1 = \min_{v \in \Sigma_B} \frac{\alpha[v]}{|\Sigma_A|}$
- $LB_2 = \left[\min_{W' \in W} |W'| \right]^{-1}$
- $UB_1 = \max_{v \in \Sigma_A} \frac{\beta[v]}{|\Sigma_B|}$
- $UB_2 = \left[\max_{Z' \in Z} |Z'| \right]^{-1}$

For a particular graph game, each of these bounds may or may not be attained. For example, consider the four graphs shown in Fig. 3. In each case, consider the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$ with




Fig. 3. Four simple graphs. Solutions of games played over these graphs and the values of the four bounds LB_1, LB_2, UB_1, UB_2 are summarised in Table 1.

$\Sigma_A = \Sigma_B = V(G)$ and $r = 1$. For such small graphs, optimal mixed strategies are easy to calculate (for example, using the method described by Morris (1994, pp. 99–114)). Table 1 summarises the optimal mixed strategies, the true values of u (the value of \mathfrak{G} to Player A) and the values of the bounds for each of the four games.

It should be noted that while LB_1 and UB_1 will generally be easy to calculate, LB_2 and UB_2 may not be, since the minimal and maximal cardinalities of $W' \in W$ and $Z' \in Z$ respectively may be difficult to determine.

6. Simplifying graph games algorithmically

6.1. An IEDS algorithm for graph games

The results of Section 4.4 and Section 5.3 allow for the creation of an IEDS algorithm (see Section 4.5) for the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$. Since the game has finitely many strategies, this approach is always a valid method for finding a solution of the game (though it may not identify all optimal mixed strategies), as discussed in Section 4.5.

The algorithm identifies vertices that may be very weakly dominated for Player A or Player B and checks for dominance and equivalence over a small subset of their surrounding vertices. Very weakly dominated vertices are eliminated and the strategy sets for the players are iteratively reduced, forming sequences $(\Sigma_{A,K})_{K \in \mathbb{N}}$ and $(\Sigma_{B,K})_{K \in \mathbb{N}}$ of subsets of Σ_A and Σ_B respectively, until there is no dominance or equivalence in the remaining vertices. The aim is to simplify the game as far as possible, such that optimal mixed strategies can be more easily identified.

Table 1

Summary of the values of bounds LB_1, LB_2, UB_1, UB_2 , the true values u to Player A and examples of optimal mixed strategies for each player for graph games played over the four graphs shown in Fig. 3, with $\Sigma_A = \Sigma_B = V(G)$ and $r = 1$. Bold cells indicate that the relevant bounds are attained. All figures are rounded to two decimal places.

| Graph | u | Bounds | | | | Optimal strategies | | | | | |
|-------|------|-------------|-------------|-------------|-------------|--|------|------|----------|------|------|
| | | LB_1 | LB_2 | UB_1 | UB_2 | Player A | | | Player B | | |
| | | | | | | Probability allocated to each vertex marked. | | | | | |
| | | | | | | X | Y | Z | X | Y | Z |
| a | 0.75 | 0.75 | 0.50 | 0.75 | 1.00 | 0.25 | N/A | N/A | 0.25 | N/A | N/A |
| b | 0.50 | 0.50 | 0.50 | 0.75 | 0.50 | 0.00 | 0.50 | N/A | 0.50 | 0.00 | N/A |
| c | 1.00 | 0.50 | 1.00 | 1.00 | 1.00 | 0.00 | 1.00 | N/A | 0.33 | 0.00 | N/A |
| d | 0.40 | 0.33 | 0.33 | 0.56 | 0.50 | 0.00 | 0.20 | 0.20 | 0.15 | 0.10 | 0.00 |

The explicit identification of vertices that cannot be dominated and the subsequent restriction of the set of vertices that should be examined when searching for dominance of a given vertex are intended to facilitate the creation of an efficient computer programme to apply IEDS in the graph game. An application of the algorithm in a very simple case is presented in Section 6.11.

For the purpose of iteration, our existing notation is extended as follows:

$$N_{A,K}[v] = N[v] \cap \Sigma_{A,K}$$

$$N_{B,K}[v] = N[v] \cap \Sigma_{B,K}$$

$$\alpha_K[v] = |N_{A,K}[v]|$$

$$\beta_K[v] = |N_{B,K}[v]|$$

6.2. Step one: transformation of G to G'

If $r \neq 1$, we transform G to G' using the method outlined in Section 5.2.

We then set the iteration variable $K = 0$ and define:

$$\Sigma_{A,0} = \Sigma_A$$

$$\Sigma_{B,0} = \Sigma_B$$

6.3. Step two: identify vertices that cannot be very weakly dominated for Player A

By Propositions 5.1 and 5.3 and Corollary 1, in a graph in which all vertices have a distinct closed neighbourhood, a vertex v for which:

$$\beta_K[v] = \Delta(G) + 1$$

or $N_{B,K}[v] \neq \emptyset$ and $B_2[v] \cap \Sigma_{A,K} = \{v\}$

or $\beta_K[v] > \max_{w \in S'(v)} (\beta_K[w])$

where:

$$S'(v) = B_2[v] \cap \Sigma_{A,K} \setminus \{v\}$$

cannot be very weakly dominated for Player A by any other vertex. We need not therefore consider such vertices when looking for dominated or equivalent strategies.

Formally, we define a reduced set of strategies for Player A:

$$\Sigma_{A,K}^- = \Sigma_{A,K} \setminus [C_{A,K}^{(0)} \cup C_{A,K}^{(1)} \cup C_{A,K}^{(2)}]$$

with:

$$C_{A,K}^{(0)} = \{v \in \Sigma_{A,K} : \beta_K[v] = \Delta(G) + 1\}$$

$$C_{A,K}^{(1)} = \{v \in \Sigma_{A,K} : N_{B,K}[v] \neq \emptyset; B_2[v] \cap \Sigma_{A,K} = \{v\}\}$$

$$C_{A,K}^{(2)} = \left\{ v \in \Sigma_{A,K} : \beta_K[v] > \max_{w \in S'(v)} (\beta_K[w]) \right\}$$

For a graph in which two vertices may have the same closed neighbourhood, Proposition 5.1 is not used, and we instead define:

$$\Sigma_{A,K}^- = \Sigma_{A,K} \setminus \left[C_{A,K}^{(1)} \cup C_{A,K}^{(2)} \right]$$

6.4. Step three: eliminate very weakly dominated vertices for Player A

Now we must look for very weakly dominated vertices for Player A in $\Sigma_{A,K}^-$.

Vertices $v \in \Sigma_{A,K}$ for which $N_{B,K}[v] = \emptyset$ are dominated automatically by any vertex $w \in \Sigma_{A,K}$ for which $N_{B,K}[w] \neq \emptyset$, while Proposition 5.3 states that for each $v \in \Sigma_{A,K}^-$ with $N_{B,K}[v] \neq \emptyset$ we need only look for vertices that very weakly dominate v in $B_2[v] \cap \Sigma_{A,K}$.

Vertices are checked sequentially according to some pre-determined ordering and very weakly dominated vertices are removed and are not considered when searching for dominance and equivalence of any remaining vertices. Sequential checking and the immediate removal of identified vertices ensure that pairwise equivalent subsets are not eliminated from $\Sigma_{A,K}$ in their entirety, but rather reduced to singletons as required.

Let the set of very weakly dominated vertices identified in this way be denoted as $\Psi_{A,K}$.

Formally, we define a new strategy set for Player A:

$$\Sigma_{A,K+1} = \Sigma_{A,K} \setminus \Psi_{A,K}$$

Note that in the trivial case where $N_{B,K}[v] = \emptyset, \forall v \in \Sigma_{A,K}$, all vertices under consideration are equivalent for Player A, so we choose any vertex $\bar{v} \in \Sigma_{A,K}$, define $\Sigma_{A,K+1} = \{\bar{v}\}$, $\Psi_{A,K} = \Sigma_{A,K} \setminus \{\bar{v}\}$ and continue to Step Four. This ensures that $\Sigma_{A,K+1}$ is not empty.

6.5. Step four: identify vertices that cannot be very weakly dominated for Player B

In a similar fashion to Step Two, we use Proposition 5.5 and Corollary 2 to define a reduced set of strategies for Player B, removing those strategies that cannot be very weakly dominated:

$$\Sigma_{B,K}^- = \Sigma_{B,K} \setminus \left[C_{B,K}^{(1)} \cup C_{B,K}^{(2)} \right]$$

with:

$$C_{B,K}^{(1)} = \{ v \in \Sigma_{B,K} : N_{A,K+1}[v] \neq \emptyset; B_2[v] \cap \Sigma_{B,K} = \{v\} \}$$

$$C_{B,K}^{(2)} = \left\{ v \in \Sigma_{B,K} : \alpha_{K+1}[v] < \min_{w \in S^v(v)} (\alpha_{K+1}[w]) \right\}$$

where:

$$S^v(v) = B_2[v] \cap \Sigma_{B,K} \setminus \{v\}$$

However, in the trivial case where $\exists v^* \in \Sigma_{B,K}$ with $N_{A,K+1}[v^*] = \emptyset$, observe that v^* (which need not be unique) very weakly dominates every other vertex for Player B, so we define $\Sigma_{B,K}^- = \Sigma_{B,K}$ and continue to Step Five.

Note the use of $N_{A,K+1}[v]$ and $\alpha_{K+1}[v]$ rather than $N_{A,K}[v]$ and $\alpha_K[v]$ at this step.

6.6. Step five: eliminate very weakly dominated vertices for Player B

Now we must look for very weakly dominated vertices for Player B in $\Sigma_{B,K}^-$.

If $\exists v^* \in \Sigma_{B,K}^-$ with $N_{A,K+1}[v^*] = \emptyset$ then v^* very weakly dominates every other vertex for Player B, as discussed, so we define $\Sigma_{B,K+1} = \{v^*\}$, $\Psi_{B,K} = \Sigma_{B,K} \setminus \{v^*\}$ and continue to Step Six.

Otherwise, Proposition 5.5 states that for each $v \in \Sigma_{B,K}^-$ we need only look for vertices that very weakly dominate v in $B_2[v] \cap \Sigma_{B,K}$.

Vertices are checked sequentially and very weakly dominated vertices are removed and are not considered when searching for dominance and equivalence of any remaining vertices.

Let the set of very weakly dominated vertices identified in this way be denoted as $\Psi_{B,K}$.

Formally, we define a new strategy set for Player B:

$$\Sigma_{B,K+1} = \Sigma_{B,K} \setminus \Psi_{B,K}$$

6.7. Step six: condition for continued iteration

If no very weakly dominated vertices were identified at Steps Three and Five ($\Psi_{A,K} = \Psi_{B,K} = \emptyset$), then no further pure strategy dominance or equivalence exists. In this case, proceed to Step Seven. Otherwise, increase the iteration variable K by one and return to Step Two.

6.8. Step seven: find optimal mixed strategies for the simplified game

Construct the payoff matrix for the simplified game defined by the strategy sets $\Sigma_{A,K}, \Sigma_{B,K}$. If optimal mixed strategies for this game can be found (for example, using the method described by Morris (1994), pp. 99–114), they are optimal mixed strategies for the complete graph game (see Morris, 1994, pp. 48–49; Berwanger, 2007, p. 2).

6.9. Termination and efficiency of the algorithm

Note that this algorithm must terminate for some $K \in \mathbb{N}$. The condition for continued iteration (Step Six) is only fulfilled if new pure strategy dominance or equivalence is identified and this necessarily results in the elimination of strategies from one player's strategy set. Also, graph games have been defined on graphs with a finite set of κ vertices, which implies that $\Sigma_{A,0}$ and $\Sigma_{B,0}$ are finite sets. Therefore the integer sequence:

$$(|\Sigma_{A,K}| + |\Sigma_{B,K}|)_{K \in \mathbb{N}} \tag{10}$$

is strictly decreasing and positive, and so must terminate for some $K = K' \in \mathbb{N}$.

Though a thorough investigation of the computational complexity of the algorithm is beyond the scope of this paper (and would depend to some extent on its precise implementation), a crude measure of the algorithm's efficiency can be determined by considering the number of pairs of vertices that must be tested for very weak dominance before the algorithm terminates. Given a graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$ with $|V(G)| = \kappa$, an (extremely conservative) upper bound for this number can be determined as follows.

At each of Steps Three and Five, a particular vertex can be tested against no more than $\kappa - 1$ other vertices. Therefore, the number of vertex pairs tested for very weak dominance at a single iteration certainly does not exceed $2\kappa^2$. Also, since (10) is a strictly decreasing positive sequence and $|\Sigma_{A,0}|, |\Sigma_{B,0}| \leq \kappa$, the number of iterations clearly cannot exceed 2κ . Therefore, for any graph game, the total number of vertex pairs that must be tested for very weak dominance before the algorithm terminates will not exceed a cubic function of the number of vertices κ .

Alternatively, since the algorithm only tests vertex pairs for very weak dominance if the distance between them is less than 3, at each of Steps Three and Five, a particular vertex will be tested against no more than $\Delta(G)^2$ other vertices (where $\Delta(G)$ is the maximum degree of the vertices of the graph after the completion of Step One). Therefore, the number of vertex pairs tested in a single iteration certainly does not exceed $2\Delta(G)^2\kappa$. Again using the fact that the number of iterations cannot exceed 2κ , we may conclude that, for any graph game, the total number of vertex pairs tested before termination will not exceed a quadratic function of $\Delta(G)\kappa$.

Which of these bounds is more useful will depend on the structure of the graphs under consideration.

6.10. The algorithm applied to games on trees

In this section and the next, we demonstrate that the algorithm of Section 6.1 offers a distinct advantage over the method proposed by White (1994) for solving games played on trees (with $r = 1$). Proposition 6.2 establishes that the algorithm always succeeds in reducing games of this type to cases that are trivially simple to solve, while Section 6.11 presents a simple example which shows that the algorithm can also be applied to games on graphs that are not trees. Also note that White (1994) exclusively considers situations where $\Sigma_A = \Sigma_B = V(G)$, while the algorithm presented here is not restricted to such cases.

Recall that a tree is a graph in which there is a unique path between any pair of vertices. In particular, if G is a tree, we may designate a vertex $t \in V(G)$ as the “root” of G , such that any vertex $v \neq t$ has precisely one neighbour v_p (the “parent” of v) for which:

$$d_G(t, v_p) = d_G(t, v) - 1$$

For any other neighbour v_c of v (a “child” of v) we have:

$$d_G(t, v_c) = d_G(t, v) + 1$$

As in Section 4.3, let ρ_A and ρ_B represent the mixed strategies for each player that allocate equal probabilities to all vertices in their strategy set. The following proposition is a restatement of Proposition 4.1 in the context of graph games, with $\zeta = 1$.

Proposition 6.1. Consider the graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

with $r = 1$. If we have:

$$\alpha[v] = 1, \quad \forall v \in \Sigma_B$$

$$\beta[w] = 1, \quad \forall w \in \Sigma_A$$

then ρ_A and ρ_B are optimal mixed strategies for Players A and B respectively and $|\Sigma_A|^{-1} = |\Sigma_B|^{-1} = u$ is the value of the game to Player A.

The following proposition states that, for graph games on trees with $r = 1$, either the value of the game to Player A is zero and analysis of the game is trivial or the algorithm of Section 6.1 reduces the game to the form given in Proposition 6.1, for which ρ_A and ρ_B have been shown to be optimal mixed strategies.

Proposition 6.2. Consider the graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

where G is a tree and $r = 1$. When applied to \mathfrak{G} , the algorithm of Section 6.1 terminates for some $K = K' \in \mathbb{N}$, such that:

• Either:

$$\Sigma_{A,K'} = \{w\}, \quad \beta_{K'}[w] = 0$$

$$\Sigma_{B,K'} = \{v\}, \quad \alpha_{K'}[v] = 0$$

and u , the value of the game to Player A, is 0;

• Or:

$$\alpha_{K'}[v] = 1, \quad \forall v \in \Sigma_{B,K'}$$

$$\beta_{K'}[w] = 1, \quad \forall w \in \Sigma_{A,K'}$$

and thus, by Proposition 6.1, the mixed strategies $\rho_{A,K'}$ and $\rho_{B,K'}$, which allocate equal probabilities to all vertices in the players' respective strategy sets $\Sigma_{A,K'}$ and $\Sigma_{B,K'}$, are optimal mixed strategies and the value of the game to Player A is:

$$u = |\Sigma_{A,K'}|^{-1} = |\Sigma_{B,K'}|^{-1}$$

A proof of the proposition is presented in Appendix B.

6.11. An application of the algorithm

A computer programme was written in Python (Python Software Foundation, 2012) using NumPy (Numpy Developers, 2012) to implement the algorithm of Section 6.1.

The single example presented here is clearly very simple and is included purely to demonstrate that the algorithm can be applied to games on graphs other than trees, though the extent to which it is able to simplify such game varies greatly. Note particularly that this example has $r \neq 1$ and $\Sigma_A \neq \Sigma_B$.

Consider the graph game $\mathfrak{G}^* = (G^*, \Sigma_A, \Sigma_B, r)$, where $r = 2$ and G^* is the graph over 14 vertices v_0 to v_{13} defined by the adjacency matrix:

$$M = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

with:

$$\Sigma_A = V(G^*)$$

$$\Sigma_B = \{v_1, v_2, v_5, v_6, v_8, v_{10}, v_{11}, v_{12}, v_{13}\}$$

The graph is shown in Fig. 4.

The algorithm reduces the game to the following case:

- Remaining strategies for Player A : v_1, v_7, v_8, v_{13}
- Remaining strategies for Player B : v_0, v_2, v_6, v_{12}
- Payoff matrix (for Player A):

| | | | | | |
|----------|----------|----------|-------|-------|----------|
| | | Player B | | | |
| | | v_0 | v_2 | v_6 | v_{12} |
| Player A | v_1 | 0 | 1 | 1 | 1 |
| | v_7 | 1 | 0 | 1 | 1 |
| | v_8 | 1 | 1 | 0 | 1 |
| | v_{13} | 1 | 1 | 1 | 0 |

Calculating the optimal mixed strategies from this matrix is simple, for example, using the method described by Morris (1994, pp. 99–114). Alternatively, observe that we may apply Proposition 4.1 to this reduced game, setting $\zeta = 3$.

An optimal mixed strategy for Player A is to play vertices v_1, v_7, v_8, v_{13} , each with probability 0.25; an optimal mixed strategy for Player B is to play vertices v_0, v_2, v_6, v_{12} , each with probability 0.25. The value of the game to Player A is 0.75.

7. Further methods for analysing graph games

7.1. Exploiting the automorphisms of a graph game

The definitions and notation relating to group theory used in this section are from Neumann, Stoy, and Thompson (1994).

An automorphism of a graph $G = (V(G), E(G))$ is a mapping of the graph to itself, which preserves adjacencies. The following definition is adapted from Bondy and Murty (1976, pp. 5–7):

Definition 7.1. Consider a simple undirected graph $G = (V(G), E(G))$.

- An **automorphism** ϕ of G is a permutation of the vertices:

$$\begin{aligned} \phi : V(G) &\rightarrow V(G) \\ v &\mapsto \phi(v) \end{aligned}$$

such that:

$$\{v, w\} \in E(G) \iff \{\phi(v), \phi(w)\} \in E(G)$$

- The **automorphism group** $\Gamma(G)$ is the permutation group formed by the set of all automorphisms of G , with the operation of composition.
- The **orbit** $\mathcal{O}_H[v]$ of a vertex $v \in V(G)$ under a subgroup $H \leq \Gamma(G)$ is the set:

$$\mathcal{O}_H[v] = \{w \in V(G) : \exists \phi \in H \text{ with } \phi(v) = w\}$$

The orbits of the vertices of G under H form a partition of $V(G)$.

Note that this definition is valid for simple graphs. For more general graphs, which may include directed edges, loops, or multiple edges connecting the same vertices, a permutation of the edges also needs to be specified.

We extend the concept of a graph automorphism to that of a graph game automorphism, an automorphism of the graph that also preserves the strategic status of the vertices:

Definition 7.2. Consider a graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

with $r = 1$. A **graph game automorphism** ϕ of \mathfrak{G} is an automorphism of G , such that:

- $v \in \Sigma_A \iff \phi(v) \in \Sigma_A, \quad \forall v \in V(G)$
- $v \in \Sigma_B \iff \phi(v) \in \Sigma_B, \quad \forall v \in V(G)$

We can also define automorphism groups and vertex orbits in terms of graph game automorphisms:

Proposition 7.1. Consider a graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

with $r = 1$.

- The set of all graph game automorphisms of \mathfrak{G} , with the operation of composition, forms a subgroup $\Gamma(\mathfrak{G})$ (the **graph game automorphism group**) of $\Gamma(G)$.




Fig. 4. The graph G . Square vertices lie in Σ_A only, while circular vertices lie in $\Sigma_A \cap \Sigma_B$.

- The orbits of the vertices of G under any subgroup $H \leq \Gamma(\mathfrak{G})$ form a partition of $V(G)$.

For $\Gamma(\mathfrak{G})$ to be a subgroup of $\Gamma(G)$, we require that $\Gamma(\mathfrak{G})$ is closed under composition, contains the identity and contains an inverse for every element. Since the only restriction on $\Gamma(\mathfrak{G})$ is that Σ_A, Σ_B are invariant under its elements, these conditions are clearly satisfied. The second part of the proposition is true of any permutation group.

The following proposition describes a relationship between certain optimal mixed strategies of \mathfrak{G} and the graph game automorphism group $\Gamma(\mathfrak{G})$.

Proposition 7.2. Consider a graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

with $V(G) = \{v_1, \dots, v_\kappa\}$ and $r = 1$. Let $\mathcal{O}[v]$ denote the orbit of v under some subset H of the graph game automorphism group $\Gamma(\mathfrak{G})$.

There exists a pair of optimal mixed strategies:

$$\sigma_A = (\sigma_A[v_1], \dots, \sigma_A[v_\kappa])$$

$$\sigma_B = (\sigma_B[v_1], \dots, \sigma_B[v_\kappa])$$

where σ_A and σ_B respectively allocate probabilities $\sigma_A[v_i]$ and $\sigma_B[v_i]$ to vertex v_i , such that $\forall i, j \in \{1, \dots, \kappa\}$:

$$\mathcal{O}[v_i] = \mathcal{O}[v_j] \Rightarrow \begin{cases} \sigma_A[v_i] = \sigma_A[v_j] \\ \sigma_B[v_i] = \sigma_B[v_j] \end{cases} \quad (11)$$

Note that in particular the proposition is true for $H = \Gamma(\mathfrak{G})$. A proof of the proposition is presented in Appendix B.

Proposition 7.2 implies that, when looking for optimal mixed strategies for a graph game \mathfrak{G} , it suffices to consider those strategies that allocate equal probability to all vertices lying in the same orbit under graph game automorphisms. For graphs with high numbers of symmetries, this can significantly simplify the analysis of the game.

Furthermore, the fact that the proposition is valid for any subgroup $H \leq \Gamma(\mathfrak{G})$ ensures that even if not all graph game automorphisms of \mathfrak{G} are known, to find a pair of optimal mixed strategies, it is sufficient to restrict consideration to those strategies that allocate equal probability to all vertices lying in the same orbit under the subgroup generated by those graph game automorphisms that can be identified.

Note that, though stated and proved in a different setting, this proposition is strongly related to Theorem 2.1 of Zoroa et al. (1993, pp. 526–528), which relates to games played over sets that admit certain transformations. However, Zoroa et al. require that the transformations considered be commutative, which is not necessarily true of the automorphisms of a graph.

7.2. Example: Using automorphisms to find solutions of \mathfrak{G}

Consider the graph game $\mathfrak{G} = (G_{5,5}, \Sigma_A, \Sigma_B, r)$, where $G_{5,5}$ is the 5×5 rectangular grid graph depicted in Fig. 5, $\Sigma_A = \Sigma_B = V(G_{5,5})$ and $r = 1$. Each player has 25 possible pure strategies: $V(G_{5,5}) = \{v_1, \dots, v_{25}\}$.

Note that since both Players may deploy at any vertex, $\Gamma(\mathfrak{G}) = \Gamma(G_{5,5})$, so we consider all automorphisms of $G_{5,5}$. These automorphisms are rotations through multiples of $\pi/2$, reflection about a vertical axis and compositions of these transformations. The vertices can be partitioned into six orbits, which we denote as $\mathcal{O}^{(i)}$, $i \in \{1, \dots, 6\}$. The vertices belonging to each orbit are indicated in Fig. 5.

By Proposition 7.2, there exist optimal mixed strategies σ_A, σ_B for Players A and B, which allocate the same probability to all vertices lying in the same orbit. Let $\omega_A^{(i)}$ and $\omega_B^{(i)}$ be the respective

probabilities that σ_A and σ_B allocate to individual vertices in orbit $\mathcal{O}^{(i)}$, and define:

$$A_i = |\mathcal{O}^{(i)}|\omega_A^{(i)}, \quad B_i = |\mathcal{O}^{(i)}|\omega_B^{(i)}, \quad \forall i \in \{1, \dots, 6\} \tag{12}$$

Observe that:

$$\sum_{i=1}^6 A_i = \sum_{i=1}^6 B_i = 1$$

A_i and B_i represent the total probability that σ_A and σ_B respectively allocate to all vertices in orbit $\mathcal{O}^{(i)}$.

Let $\sigma^{(1)}, \dots, \sigma^{(6)}$ be the mixed strategies (for either player), which allocate uniform probability to all vertices in the corresponding orbit and zero probability to other vertices:

$$\sigma^{(i)} = (\sigma^{(i)}[\nu_1], \dots, \sigma^{(i)}[\nu_{25}])$$

with:

$$\sigma^{(i)}[\nu_j] = \begin{cases} |\mathcal{O}^{(i)}|^{-1} & , \nu_j \in \mathcal{O}^{(i)} \\ 0 & , \text{otherwise.} \end{cases}$$

We see that σ_A and σ_B are linear combinations of the $\sigma^{(i)}$, weighted by the orbit probabilities A_i and B_i :

$$\sigma_A = \sum_{i=1}^6 A_i \sigma^{(i)}, \quad \sigma_B = \sum_{i=1}^6 B_i \sigma^{(i)}$$

This observation allows us to take a different perspective on the problem. Suppose that we treat the $\sigma^{(i)}$ as if they were pure strategies. The following strategies would be optimal strategies of such a game:

$$\tau_A = (\tau_A[\sigma^{(1)}], \dots, \tau_A[\sigma^{(6)}])$$

$$\tau_B = (\tau_B[\sigma^{(1)}], \dots, \tau_B[\sigma^{(6)}])$$

with:

$$\tau_A[\sigma^{(i)}] = A_i, \quad \forall i \in \{1, \dots, 6\}$$

$$\tau_B[\sigma^{(i)}] = B_i, \quad \forall i \in \{1, \dots, 6\}$$

To find these optimal mixed strategies, we analyse the matrix of expected payoffs to Player A for each of these six strategies against one another:

| | | Player B | | | | | |
|----------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|
| | | $\sigma^{(1)}$ | $\sigma^{(2)}$ | $\sigma^{(3)}$ | $\sigma^{(4)}$ | $\sigma^{(5)}$ | $\sigma^{(6)}$ |
| Player A | $\sigma^{(1)}$ | 1/4 | 1/4 | 0 | 0 | 0 | 0 |
| | $\sigma^{(2)}$ | 1/4 | 1/8 | 1/4 | 1/4 | 0 | 0 |
| | $\sigma^{(3)}$ | 0 | 1/4 | 1/4 | 0 | 1/4 | 0 |
| | $\sigma^{(4)}$ | 0 | 1/4 | 0 | 1/4 | 1/2 | 0 |
| | $\sigma^{(5)}$ | 0 | 0 | 1/4 | 1/2 | 1/4 | 1 |
| | $\sigma^{(6)}$ | 0 | 0 | 0 | 0 | 1 | 1 |




Fig. 5. The 5 × 5 rectangular grid graph $G_{5,5}$. Vertices marked i belong to orbit $\mathcal{O}^{(i)}$ under $\Gamma(G_{5,5})$.

The optimal mixed strategies τ_A and τ_B can be computed from this payoff matrix using standard methods (see, for example, Morris (1994, pp. 99–114)), giving:

$$(A_1, \dots, A_6) = 69^{-1}(16, 28, 10, 4, 6, 5)$$

$$(B_1, \dots, B_6) = 69^{-1}(16, 28, 10, 4, 6, 5)$$

Using (12) we then find:

$$(\omega_A^{(1)}, \dots, \omega_A^{(6)}) = 138^{-1}(8, 7, 5, 2, 3, 10)$$

$$(\omega_B^{(1)}, \dots, \omega_B^{(6)}) = 138^{-1}(8, 7, 5, 2, 3, 10)$$

This fully determines the optimal mixed strategies σ_A and σ_B for \mathfrak{G} . These optimal mixed strategies are illustrated in Fig. 6.

Note that by identifying the six orbits of $V(G)$ under graph game automorphisms of \mathfrak{G} , rather than solving a game with a 25×25 payoff matrix, we needed only to solve a game with a 6×6 payoff matrix. A similar method could be applied to any graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$ with $r = 1$, for which non-trivial graph game automorphisms can be found.

7.3. Equal oddments strategies

The example of the graph game \mathfrak{G} played over the 5×5 rectangular grid $G_{5,5}$ exhibits two curious and related properties.

Firstly, and most obviously, the optimal mixed strategies that were calculated for each player are identical $\sigma_A = \sigma_B = \sigma$. This means that the optimal mixed strategy for Player A allocates identical probabilities to the vertices of $G_{5,5}$ as does the optimal mixed strategy for Player B (though these optimal mixed strategies need not be unique). Given the interpretation of the game, this is a surprising result. Player B is attempting to hide from Player A, so we might have expected that her best mixed strategy would involve avoiding vertices at which Player A was more likely to deploy.

Note that the method of Section 7.1 does not produce identical optimal mixed strategies for all graphs, nor for all grid graphs. For example, when applied to the 6×6 grid graph $G_{6,6}$, the method produces distinct optimal mixed strategies σ_A, σ_B (see Fig. 7) that appear to be more consistent with our intuition, in that vertices to which σ_A allocates fairly high probability seem to be allocated fairly low probability by σ_B and vice versa.

Secondly, in Fig. 6, observe that for all vertices $v \in V(G_{5,5})$, the sum of the probability oddments allocated by σ to the vertices in the closed neighbourhood $N[v]$ is equal to 22. Formally, we make the following definition:




Fig. 6. Probability oddments (integer values proportional to the probabilities) allocated to each vertex of $G_{5,5}$ by the (identical) optimal mixed strategies σ_A and σ_B for the graph game $\mathfrak{G} = (G_{5,5}, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G_{5,5})$. Darker shading indicates lower probability oddments.




Fig. 7. Probability oddments allocated to each vertex of $G_{6,6}$ by the optimal mixed strategies σ_A and σ_B for the graph game $\mathfrak{G} = (G_{6,6}, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G_{6,6})$. Darker shading indicates lower probability oddments.

Definition 7.3. Given a graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G) = \{v_1, \dots, v_\kappa\}$, an **equal oddments strategy** σ is a vector:

$$\sigma = (\sigma[v_1], \dots, \sigma[v_\kappa])$$

with:

$$0 \leq \sigma[v_i] \leq 1, \quad \forall i \in \{1, \dots, \kappa\}$$

$$\sum_{i=1}^{\kappa} \sigma[v_i] = 1$$

such that:

$$\sum_{w \in N[v_j]} \sigma[w] = u, \quad \forall j \in \{1, \dots, \kappa\}$$

For some $u \in \mathbb{R}$. We call u the **neighbourhood sum** of σ .

The following proposition links those situations in which Players A and B have identical optimal mixed strategies with the existence of equal oddments strategies.

Proposition 7.3. Given a graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G)$, then the following statements are equivalent:

- σ is an equal oddments strategy of \mathfrak{G} with neighbourhood sum $u \in \mathbb{R}$.
- σ is an optimal mixed strategy of \mathfrak{G} for both players and $u \in \mathbb{R}$ is the value of the game to Player A.

Proof. Suppose that σ is an equal oddments strategy of \mathfrak{G} with neighbourhood sum $u \in \mathbb{R}$. Observe that for all pure strategies $v \in V(G)$ for either player, we have:

$$E[p_A(v, \sigma)] = u$$

$$E[p_B(\sigma, v)] = 1 - u$$

Therefore, by (5), σ is an optimal mixed strategy of \mathfrak{G} for both players and u is the value of the game to Player A.

To prove the opposite implication, suppose that σ is an optimal mixed strategy of \mathfrak{G} for both players and that u is the value of the game to Player A. Let:

$$V(G) = \{v_1, \dots, v_\kappa\}$$

$$\sigma = (\sigma[v_1], \dots, \sigma[v_\kappa])$$

Also let:

$$m_- = \min_{v \in V(G)} \left[\sum_{w \in N[v]} \sigma[w] \right]$$

$$m_+ = \max_{v \in V(G)} \left[\sum_{w \in N[v]} \sigma[w] \right]$$

and let v_-, v_+ be vertices for which this minimum and maximum are respectively attained.

Assume (to derive a contradiction) that σ is not an equal oddments strategy. Then:

$$m_- < m_+ \tag{13}$$

Therefore:

$$m_+ = E[p_A(v_+, \sigma)] \geq E[p_A(\sigma, \sigma)] = u$$

$$1 - m_- = E[p_B(\sigma, v_-)] \geq E[p_B(\sigma, \sigma)] = 1 - u$$

Since σ is an optimal mixed strategy, the above inequalities must be equalities. Specifically:

$$m_- = m_+ = u$$

This contradicts (13). Therefore σ is an equal oddments strategy. \square

Proposition 7.3 means that, given a graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G)$, if we can find a distribution of positive real numbers across the vertices such that the sum of these numbers in any closed neighbourhood is equal to a constant, and we scale these numbers to produce a valid probability distribution across the vertices, this distribution defines a mixed strategy that is optimal for both players. This offers an alternative approach to proving **Corollary 3** on optimal mixed strategies for games played over regular graphs, since the mixed strategy ρ , as defined in the corollary, is clearly an equal oddments strategy.

8. Games on poly-level graphs

8.1. Poly-level graphs

The concepts discussed in Section 7 suggest a potential approach for finding general expressions for the optimal mixed strategies of games played over certain families of graphs. To demonstrate this approach, we consider games played over a specific family of graphs, which we describe as **poly-level graphs**. The vertices in such graphs are arranged in levels and each vertex exhibits a local structural similarity in the way that it is connected to other vertices in its own level and to those in the levels above and below.

Definition 8.1. Poly-level graph

A graph $G = (V(G), E(G))$ is called a **poly-level graph** if and only if the vertices can be partitioned into h subsets $L_1, L_2, \dots, L_h \subseteq V(G)$ called **levels**, with $|L_i| = c_i$, such that each vertex in L_i is adjacent to precisely:

- j other vertices in L_i (the **intradegree**);
- k vertices in L_{i+1} (the **superdegree**), for $i \neq h$;
- l vertices in L_{i-1} (the **subdegree**), for $i \neq 1$;
- 0 vertices in any other level.

where j is a non-negative integer and h, k, l, c_1, \dots, c_h are positive integers.

Fig. 8 provides a visual representation of a general poly-level graph, while Figs. 9–13 show specific examples of such graphs. In each example, vertices in L_1 are coloured black, with higher levels being indicated by progressively lighter shades.

The appearance of a poly-level graph may be highly symmetric (Figs. 9 and 10) or quite irregular (Fig. 11). Also, two poly-level graphs with the same parameters may be topologically quite different (Fig. 12), while the parameters used to describe a particular poly-level graph are generally not unique (Fig. 13).

Note that it suffices to specify the number of vertices c_1 in L_1 to determine the number of vertices c_i in any level L_i , as established in the following simple proposition:

Proposition 8.1. Given a poly-level graph G , with levels L_1, \dots, L_h , superdegree k , subdegree l and $|L_i| = c_i$, we have:

$$c_i = (kl^{-1})^{i-1} c_1, \quad \forall i \in \{1, \dots, h\}$$

Proof. The number of edges linking vertices in L_i to vertices in L_{i+1} can be expressed in two forms, which must be equal:

$$lc_i = kc_{i-1}, \quad \forall i \in \{2, 3, \dots, h\}$$

Thus, the c_i form a geometric progression with common ratio kl^{-1} and the proposition follows immediately. \square

The following proposition establishes two intuitively obvious constraints on the parameters of a poly-level graph.




Fig. 8. Visual representation of a poly-level graph. c_i represents the number of vertices in level L_i ; j , k and l represent the number of edges connecting a single vertex in the level at the tail of the arrow to vertices in the level at the head of the arrow.




Fig. 9. Poly-level graph with $h = 4, j = 2, k = 1, l = 1, c_1 = 4$.




Fig. 10. Poly-level graph with $h = 3, j = 2, k = 2, l = 1, c_1 = 3$.

Proposition 8.2. Given a poly-level graph G , with levels L_1, \dots, L_h , intradegree j , superdegree k , subdegree l and $|L_i| = c_i$, we have that $\forall i \in \{1, \dots, h\}$:

- (a) $(kl^{-1})^{i-1} c_1 \in \mathbb{Z}$
- (b) $(kl^{-1})^{i-1} c_1 \geq j + 1$

These constraints arise immediately from the fact that the number of vertices c_i in any level L_i (replaced in Proposition 8.2 by the expression from Proposition 8.1) must (a) be an integer and (b) exceed the intradegree j .

8.2. Equal oddments solutions on poly-level graphs

Given a poly-level graph G , from Section 7.3, we know that if there exists a probability distribution across the vertices of the graph such that the sum of the probabilities in any closed neighbourhood is equal to some constant $u \in (0, 1]$, then this probability distribution defines an optimal mixed strategy for both players for the game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1, \Sigma_A = \Sigma_B = V(G)$, and the value of the game to Player A is u .

Suppose that such a distribution exists, and suppose that this distribution allocates an equal probability ω_i to each vertex in level L_i . Through consideration of the structure of the graph, we can write the following equations, which must hold for all $i \in \{2, 3, \dots, h-1\}$:




Fig. 11. Poly-level graph with $h = 2, j = 3, k = 1, l = 2, c_1 = 16$.




Fig. 12. Two topologically different poly-level graphs with identical parameters $h = 2, j = 1, k = 2, l = 1, c_1 = 2$.




Fig. 13. A poly-level graph that admits two different sets of parameters: (a) $h = 3, j = 1, k = 1, l = 1, c_1 = 2$ and (b) $h = 2, j = 1, k = 2, l = 1, c_1 = 2$.

$$\left. \begin{aligned} (j+1)\omega_1 + k\omega_2 &= u \\ l\omega_{h-1} + (j+1)\omega_h &= u \\ l\omega_{i-1} + (j+1)\omega_i + k\omega_{i+1} &= u \end{aligned} \right\} \quad (14)$$

We also have the following constraints, to ensure that the ω_i define a valid probability distribution:

$$\sum_{i=1}^h \omega_i c_i = 1 \quad (15)$$

$$\omega_i \geq 0, \quad \forall i \in \{1, 2, \dots, h\} \quad (16)$$

Note that it is not necessary to explicitly include the constraint $\omega_i \leq 1$, since this is implied by (15) and (16).

We now perform the change of variables:

$$w_i = u^{-1}(j+k+l+1)\omega_i, \quad \forall i \in \{1, \dots, h\} \quad (17)$$

and introduce additional unknowns $w_i, \forall i \in \mathbb{Z}$.

The question of finding a function ω_i of i that satisfies (14)–(16) can be reformulated as follows:

- Given a non-negative integer j and positive integers h, k, l, c_1 satisfying the conditions of Proposition 8.2, find a function w_i of i , such that, $\forall i \in \mathbb{Z}$:

$$lw_{i-1} + (j+1)w_i + kw_{i+1} = j+k+l+1 \quad (18)$$

- Subject to the boundary conditions:

$$w_0 = w_{h+1} = 0 \quad (19)$$

- With the constraint:

$$w_i \geq 0, \quad \forall i \in \{1, 2, \dots, h\} \quad (20)$$

If suitable w_i can be found, an equal oddments solution is then given by:

$$u = c_1^{-1}(j+k+l+1) \left[\sum_{i=1}^h (kl^{-1})^{i-1} w_i \right]^{-1} \quad (21)$$

$$\omega_i = w_i u (j+k+l+1)^{-1}, \quad \forall i \in \{1, \dots, h\} \quad (22)$$

Applying the change of variables (17): (18) and (19) are derived from (14), (20) is derived from (16), (21) is derived from (15) and Proposition 8.1, while (22) is immediate.

8.3. Exact solutions

To find a solution to the problem, let:

$$w_i = w_i^{GS} + w_i^{PS}$$

where w_i^{PS} is any particular solution of (18), and w_i^{GS} is the most general solution of the homogeneous difference equation:

$$lw_{i-1} + (j+1)w_i + kw_{i+1} = 0, \quad \forall i \in \mathbb{Z} \quad (23)$$

An obvious candidate for the particular solution is:

$$w_i^{PS} = 1$$

For w_i^{GS} , there are three cases to consider:

8.4. Case 1: $(j+1)^2 > 4kl$

Look for a solution of the form:

$$w_i^{GS} = A_+ \lambda_+^i + A_- \lambda_-^i$$

where A_+, A_- are arbitrary constants and λ_+, λ_- are unknown constants to be determined. Substituting into (23), we find:

$$\lambda_{\pm} = \frac{-(j+1) \pm \sqrt{(j+1)^2 - 4kl}}{2k}$$

So:

$$w_i = A_+ \lambda_+^i + A_- \lambda_-^i + 1$$

Applying the boundary conditions gives:

$$A_+ + A_- = -1$$

$$A_+ \lambda_+^{h+1} + A_- \lambda_-^{h+1} = -1$$

which can be solved to give:

$$A_+ = \frac{1 - \lambda_-^{h+1}}{\lambda_-^{h+1} - \lambda_+^{h+1}}$$

$$A_- = \frac{1 - \lambda_+^{h+1}}{\lambda_+^{h+1} - \lambda_-^{h+1}}$$

Note that in this case, $\lambda_+ \neq \lambda_-$, and therefore A_+ and A_- both exist.

8.5. Case 2: $(j+1)^2 = 4kl$

Look for a solution of the form:

$$w_i^{GS} = (A + iB)\lambda^i$$

where A, B are arbitrary constants and λ is an unknown constant to be determined. Substituting into (23), we find:

$$\lambda = -\frac{j+1}{2k}$$

So:

$$w_i = (A + iB)\lambda^i + 1$$

Applying the boundary conditions gives:

$$A = -1$$

$$B = \frac{1 - \lambda^{-[h+1]}}{h+1}$$

Again, A and B exist for all values of the parameters.

8.6. Case 3: $(j+1)^2 < 4kl$

Look for a solution of the form:

$$w_i^{GS} = \lambda^i [P \cos(\theta i) + Q \sin(\theta i)]$$

where P, Q are arbitrary constants and λ, θ are unknown constants to be determined. Substituting into (23), we find:

$$\lambda = \sqrt{l/k} \quad (24)$$

$$\theta = \arg \left[-(j+1) + i\sqrt{4kl - (j+1)^2} \right] \quad (25)$$

So:

$$w_i = \lambda^i [P \cos(\theta i) + Q \sin(\theta i)] + 1$$

Provided that $\theta[h+1]$ is not an integer multiple of π , applying the boundary conditions gives:

$$P = -1$$

$$Q = \frac{\cos(\theta[h+1]) - \lambda^{-[h+1]}}{\sin(\theta[h+1])}$$

If $\theta[h+1]$ is an integer multiple of π then no values of P and Q can be found to satisfy the boundary conditions (19) and no solution w_i exists, except in the particular case where $\theta[h+1]$ is an even multiple of π and $k=l$. In the latter instance, the boundary conditions are

satisfied for all values of Q , with $P = -1$, identifying an infinite family of possible solutions.

8.7. Summary

To summarise, given any poly-level graph G with suitable parameters h, j, k, l, c_1 , we have found w_i satisfying (18) and (19) in all cases except where $(j + 1)^2 < 4kl$ and $\theta[h + 1]$ is an integer multiple of π (barring the special case where $\theta[h + 1]$ is an even multiple of π and $k = l$), where θ is defined as in (25).

Note that the non-existence of such a w_i in the specific cases mentioned does not necessarily imply that no equal oddments strategy exists for the corresponding graph game, simply that any such strategy cannot be expressed as the solution to a difference equation of the form discussed.

Where such a w_i does exist, it remains to check whether constraint (20) is satisfied. If so, then there exists an equal oddments solution to the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and $\Sigma_A = \Sigma_B = V(G)$, defined by (21) and (22), which allocates a probability of ω_i to each vertex in level L_i . By Proposition 7.3, this equal oddments solution is an optimal mixed strategy of \mathfrak{G} for both players.

Table 2 summarises the equal oddments strategies for the graphs shown in Figs. 9–11.

In this section, we have demonstrated that the concepts outlined in Section 7 may be used to find general expressions for optimal mixed strategies of games played over a particular family of graphs. It may therefore be possible to use or adapt this approach to seek optimal mixed strategies for games played over other families of graphs, thus suggesting a potentially valuable focus for future work.

9. Conclusions and further work

In this paper, we have defined a search and concealment game, the SSSG, which differs from the games considered in the bulk of the literature both in its generality and in its explicit consideration of the strategic interdependence of positions based on spatial structure. We believe that formulating the game in terms of metric spaces may allow new tools and results from this area of mathematics to be applied to search and concealment problems, facilitating the identification of optimal mixed strategies in a variety of different cases.

We have examined the way in which the game theoretic concepts of dominance and equivalence of strategies are manifested in the context of the SSSG and have presented various methods for analysing the SSSG played over a graph, including:

- The reduction of graph games with detection radius $r \neq 1$ to games with $r = 1$.
- The formulation of lower and upper bounds on the value of a graph game.
- An algorithm that applies the concept of IEDS in the explicit context of graph games and which has been demonstrated to reduce games played on trees to cases in which optimal mixed strategies can be immediately determined.

- A method for simplifying the the analysis of graph games using automorphisms of the graph.
- The introduction of the concept of an “equal oddments strategy”, and the demonstration that such mixed strategies are optimal for both players.
- The presentation of explicit optimal mixed strategies for a particular family of graph games; those played over “poly-level graphs”.

In terms of future work, of immediate interest would be the extension of the results and methods of Sections 5 and 7, which have been formulated explicitly for the SSSG played over a graph, to the SSSG over a general metric space. It is our belief that analogues of many of these results could probably be formulated through the careful application of topology and measure theory. This would be an important first step in increasing the applicability and relevance of the SSSG.

With regard to the algorithm of Section 6, it would be useful to have explicit results on its computational efficiency and on how the structure of a graph affects the extent to which a corresponding graph game can be simplified using IEDS. It may be possible, for example, to identify other families of graph game (aside from those played on trees) for which the the method is particularly effective.

It would also be valuable to find analytic expressions for the optimal mixed strategies of a wider range of graph games, perhaps extending the ideas used to analyse games played over poly-level graphs to games played over rectangular lattices. For highly irregular graphs, algorithmic methods to find equal oddments strategies could be sought. Where equal oddments strategies do not exist, but where equal oddments ‘distributions’ that fail to satisfy (16) can nonetheless be found, an investigation into the connection between the true optimal mixed strategies and these invalid equal oddments distributions may be instructive.

Thinking more broadly, there are many ways in which the game could be extended. For example, by increasing the number of searching and concealing players, by allowing players to deploy at multiple points, by investigating situations in which payoffs are dependent on location, by relaxing the metric space axioms to consider directed graphs and other spaces, by investigating the game over weighted graphs or by allowing the players to move around the space.

For the SSSG as described, a particular area of interest is the analysis of games in which Ω is a region of \mathbb{R}^2 , where the searching player can search a disc of radius r . An understanding of such cases would make the game more applicable in real operational research scenarios, such as the deployment of patrol ships to locate pirates over areas of the ocean or searching archaeological sites for features of historical interest. If optimal mixed strategies cannot be obtained directly, one approach to such continuous situations would be to examine a discretised version of the game, superimposing a rectangular grid over the region in question and analysing a corresponding graph game over this grid.

Finally, it would be valuable to investigate the initial ways in which the SSSG could be profitably applied to real world scenarios. For example, following the riots in England in the summer of 2011, the game could be used to examine possible deployment strategies

Table 2

Table detailing the equal oddments solutions to the poly-level graphs shown in Figs. 9–11. ω_i is the probability that the equal oddments strategy (an optimal mixed strategy for both players) allocates to each vertex in level L_i .

| Figure | Poly-level graph parameters | | | | | Equal oddments solution | | | | Game values | |
|--------|-----------------------------|-----|-----|-----|-------|-------------------------|------------|------------|------------|-------------|-------|
| | h | j | k | l | c_1 | ω_1 | ω_2 | ω_3 | ω_4 | u_A | u_B |
| 9 | 4 | 2 | 1 | 1 | 4 | 3/40 | 1/20 | 1/20 | 3/40 | 11/40 | 29/40 |
| 10 | 3 | 2 | 2 | 1 | 3 | 1/15 | 0 | 1/15 | N/A | 1/5 | 4/5 |
| 11 | 2 | 3 | 1 | 2 | 16 | 3/64 | 1/32 | N/A | N/A | 7/32 | 25/32 |

for the police. Manchester city centre, where a number of incidents occurred over a relatively compact area (Rogers, Sedghi, & Evans, 2011), could be represented as a graph to which the algorithm of Section 6 would be applied to determine the optimal location for a police unit (the searching player) to protect retail centres from a band of rioters (the concealing player). After the deployment of a first police unit, vertices within the detection radius could be removed and new optimal mixed strategies could be calculated on the rest of the graph to determine the placement of a second unit. It may be of interest to test the strategies obtained from such a process by simulation of an outbreak of rioting and to compare the results with the actual events of summer 2011.

Acknowledgements

The authors acknowledge the financial support of the UK Engineering and Physical Sciences Research Council (EPSRC) under the grant ENFOLD-ing – Explaining, Modelling, and Forecasting Global Dynamics, reference EP/H02185X/1.

The authors thank Professor Sir Alan Wilson of University College London for his helpful comments and Dr. Lynda White of Imperial College London for the benefit of her advice and insight. Thanks are also due to the anonymous reviewers, whose constructive comments contributed significantly to the final version of this article.

Appendix A. Proof of Proposition 6.2

We reproduce the proposition here for ease of reference:

Proposition 6.2. Consider the graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

where G is a tree and $r = 1$. When applied to \mathfrak{G} , the algorithm of Section 6.1 terminates for some $K = K' \in \mathbb{N}$, such that:

- Either:

$$\begin{aligned} \Sigma_{A,K'} &= \{w\}, & \beta_{K'}[w] &= 0 \\ \Sigma_{B,K'} &= \{v\}, & \alpha_{K'}[v] &= 0 \end{aligned}$$

and u , the value of the game to Player A, is 0;

- Or:

$$\begin{aligned} \alpha_{K'}[v] &= 1, & \forall v \in \Sigma_{B,K'} \\ \beta_{K'}[w] &= 1, & \forall w \in \Sigma_{A,K'} \end{aligned}$$

and thus, by Proposition 6.1, the mixed strategies $\rho_{A,K'}$ and $\rho_{B,K'}$, which allocate equal probabilities to all vertices in the players' respective strategy sets $\Sigma_{A,K'}$ and $\Sigma_{B,K'}$, are optimal mixed strategies and the value of the game to Player A is:

$$u = |\Sigma_{A,K'}|^{-1} = |\Sigma_{B,K'}|^{-1}$$

Proof. To derive a contradiction, assume that, when applied to \mathfrak{G} , the algorithm of Section 6.1 terminates for some $K = K' \in \mathbb{N}$ for which at least one of the following two conditions holds:

- $\exists v' \in \Sigma_{B,K'}$ such that $\alpha_{K'}[v'] > 1$
- $\exists w' \in \Sigma_{A,K'}$ such that $\beta_{K'}[w'] > 1$

We will demonstrate that this assumption is false by identifying two vertices in $\Sigma_{A,K'}$ or $\Sigma_{B,K'}$, one of which very weakly dominates the other for the appropriate player. This will be a contradiction, since the algorithm only terminates when no dominance or equivalence can be found in the reduced strategy sets.

Define the following subsets of $V(G)$:

$$\begin{aligned} S_1 &= \{s \in \Sigma_{B,K'} : \alpha_{K'}[s] > 1\} \\ S_2 &= \{s \in \Sigma_{A,K'} : \beta_{K'}[s] > 1\} \\ S &= S_1 \cup S_2 \end{aligned}$$

Now choose a vertex $t \in V(G)$ to designate as the root of G and let $s' \in S$ be such that:

$$d_G(t, s') = \max_{s \in S} [d_G(t, s)] \quad (26)$$

Since $S = S_1 \cup S_2$ there are two possible (non-mutually exclusive) cases: $s' \in S_1$ or $s' \in S_2$.

If $s' \in S_1$ then $s' \in \Sigma_{B,K'}$ and at least two vertices in $N_{K'}[s']$ lie in $\Sigma_{A,K'}$. We consider two possible scenarios.

- First, suppose that at least one of the children of s' lies in $\Sigma_{A,K'}$. So, since $s' \in S_1$, there exist distinct $s_1, s_2 \in N_{A,K'}[s']$ such that:

$$d_G(t, s_1) \leq d_G(t, s_2) = d_G(t, s') + 1 \quad (27)$$

Observe that neither s_2 nor any of its children lies in $\Sigma_{B,K'}$, since this would imply that $s_2 \in S_2 \subseteq S$, which is a contradiction by (26) and (27). Therefore:

$$N_{B,K'}[s_2] = \{s'\} \subseteq N_{B,K'}[s_1]$$

which means that s_1 very weakly dominates s_2 for Player A, by Proposition 4.2.

- Now suppose that none of the children of s' lies in $\Sigma_{A,K'}$. Since $s' \in S_1$, this means that s' and its parent both lie in $\Sigma_{A,K'}$. i.e. there exist $s_1, s_2 \in N_{A,K'}[s']$ (where $s_2 = s'$) such that:

$$d_G(t, s_1) < d_G(t, s_2) = d_G(t, s')$$

If none of the children of s_2 lies in $\Sigma_{B,K'}$, then:

$$N_{B,K'}[s_2] \subseteq N_{B,K'}[s_1]$$

which means that s_1 very weakly dominates s_2 for Player A, by Proposition 4.2.

Otherwise, at least one child of s_2 lies in $\Sigma_{B,K'}$. Call this child s'' and note that:

$$d_G(t, s'') = d_G(t, s') + 1 \quad (28)$$

Observe that neither s'' nor any of its children lies in $\Sigma_{A,K'}$, since this would imply that $s'' \in S_1 \subseteq S$, which is a contradiction by (26) and (28). Therefore:

$$N_{A,K'}[s''] = \{s'\} \subseteq N_{A,K'}[s']$$

which means that s'' very weakly dominates s' for Player B, by Proposition 4.4.

So if $s' \in S_1$, in both possible scenarios, we have identified very weak dominance in one of the strategy sets $\Sigma_{A,K'}$ or $\Sigma_{B,K'}$. In the case where $s' \in S_2$, very weak dominance can also be identified by following very similar logic (the complete proof is omitted). As previously discussed, this is a contradiction and thus our initial assumption was false. Therefore:

$$\begin{aligned} \alpha_{K'}[v] &\leq 1, & \forall v \in \Sigma_{B,K'} \\ \beta_{K'}[w] &\leq 1, & \forall w \in \Sigma_{A,K'} \end{aligned} \quad (29)$$

To complete the proof, note that:

- Any vertex $w \in \Sigma_{A,K'}$ for Player A such that $\beta_{K'}[w] = 0$ is clearly very weakly dominated by any other vertex for Player A.
- Any vertex $v \in \Sigma_{B,K'}$ for Player B such that $\alpha_{K'}[v] = 0$ clearly very weakly dominates any other vertex for Player B.

- Statements (*) and (**) are equivalent:

$$\begin{aligned} [\alpha_{K'}[v] = 0; \quad \forall v \in \Sigma_{B,K'}] & \quad (*) \\ [\beta_{K'}[w] = 0; \quad \forall w \in \Sigma_{A,K'}] & \quad (**) \end{aligned}$$

- Since the algorithm has terminated, no very weak dominance exists in $\Sigma_{A,K'}$ or $\Sigma_{B,K'}$.

These facts, combined with (29) imply that:

- Either:

$$\begin{aligned} \Sigma_{A,K'} &= \{w\}, \quad \beta_{K'}[w] = 0 \\ \Sigma_{B,K'} &= \{v\}, \quad \alpha_{K'}[v] = 0 \end{aligned}$$

and u , the value of the game to Player A, is clearly 0;

- Or:

$$\begin{aligned} \alpha_{K'}[v] &= 1, \quad \forall v \in \Sigma_{B,K'} \\ \beta_{K'}[w] &= 1, \quad \forall w \in \Sigma_{A,K'} \end{aligned}$$

and, by Proposition 6.1, $\rho_{A,K'}$ and $\rho_{B,K'}$ are optimal mixed strategies for each player and the value of the game to Player A is:

$$u = |\Sigma_{A,K'}|^{-1} = |\Sigma_{B,K'}|^{-1}$$

This proves the proposition. \square

Appendix B. Proof of Proposition 7.2

The definitions and notation relating to group theory used in this section are from Neumann et al. (1994). The proof requires the orbit-stabilizer theorem (adapted from Neumann et al., 1994, p. 62):

Theorem 9.1. Orbit-Stabilizer Theorem

Given a group Γ , which acts on a finite set X :

$$|\text{Stab}_\Gamma(x)| = \frac{|\Gamma|}{|\mathcal{O}_\Gamma[x]|}, \quad \forall x \in X$$

where $\mathcal{O}_\Gamma[x]$ is the orbit of x under the action of Γ , and $\text{Stab}_\Gamma(x)$ is the stabilizer of x under Γ :

$$\text{Stab}_\Gamma(x) = \{g \in \Gamma : gx = x\}$$

To prove Proposition 7.2, we will use the following corollary of the theorem. It states that if two vertices v , w of a graph game lie in the same orbit with respect to graph game automorphisms, then the number of graph game automorphisms that map v to some other vertex z is the same as the number that map w to z .

Corollary 4. Consider a subgroup H of the graph game automorphism group $\Gamma(\mathfrak{G})$ for the graph game $\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$, with $r = 1$ and let $\mathcal{O}[v]$ be the orbit of $v \in V(G)$ under H .

Then, for all $v, w, z \in V(G)$, we have:

$$\begin{aligned} \mathcal{O}[v] &= \mathcal{O}[w] \\ \Rightarrow |H[v, z]| &= |H[w, z]| = \begin{cases} \frac{|H|}{|\mathcal{O}[v]|}, & z \in \mathcal{O}[v] \\ 0, & z \notin \mathcal{O}[v] \end{cases} \end{aligned}$$

where we define:

$$H[v, z] = \{\phi \in H : \phi(v) = z\}$$

Proof. Clearly, if $\mathcal{O}[v] = \mathcal{O}[w]$ and $z \notin \mathcal{O}[v]$ then by the definition of an orbit under H , we have that:

$$|H[v, z]| = |H[w, z]| = 0$$

We claim that in the case where $z \in \mathcal{O}[v]$, we have:

$$|H[v, z]| = |\psi_{v,z} \text{Stab}_H(v)| \tag{30}$$

where $\psi_{v,z} \in H$ is any graph game automorphism in H such that $\psi_{v,z}(v) = z$ and $\psi_{v,z} \text{Stab}_H(v)$ is the left coset of $\text{Stab}_H(v)$ in H containing $\psi_{v,z}$:

$$\psi_{v,z} \text{Stab}_H(v) = \{\psi_{v,z} \circ \phi \in H : \phi \in \text{Stab}_H(v)\}$$

where \circ represents the operation of composition.

To justify the claim, we first observe that:

$$\psi \in \psi_{v,z} \text{Stab}_H(v) \Rightarrow \psi \in H[v, z]$$

because:

$$\begin{aligned} \psi &\in \psi_{v,z} \text{Stab}_H(v) \\ \Rightarrow \psi(v) &= \psi_{v,z}(\phi(v)), \quad \phi(v) \in \text{Stab}_H(v) \\ &= \psi_{v,z}(v) \\ &= z \\ \Rightarrow \psi &\in H[v, z] \end{aligned}$$

We also observe that:

$$\psi \in H[v, z] \Rightarrow \psi \in \psi_{v,z} \text{Stab}_H(v)$$

because:

$$\begin{aligned} \psi &\in H[v, z] \\ \Rightarrow \psi &= \psi_{v,z} \circ (\psi_{v,z}^{-1} \circ \psi) \end{aligned}$$

and:

$$\psi_{v,z}^{-1}(\psi(v)) = \psi_{v,z}^{-1}(z) = v$$

so:

$$\psi_{v,z}^{-1} \circ \psi \in \text{Stab}_H(v)$$

and therefore:

$$\psi \in \psi_{v,z} \text{Stab}_H(v)$$

This demonstrates the claim (30).

Now, combining this result with a standard result about cosets (Neumann et al., 1994, p. 3) gives us that:

$$|H[v, z]| = |\psi_{v,z} \text{Stab}_H(v)| = |\text{Stab}_H(v)|$$

So by the orbit-stabilizer theorem:

$$|H[v, z]| = \frac{|H|}{|\mathcal{O}[v]|}$$

Clearly, bearing in mind that $|\mathcal{O}[v]| = |\mathcal{O}[w]|$, identical logic could be employed to demonstrate that:

$$|H[w, z]| = \frac{|H|}{|\mathcal{O}[v]|}$$

This proves the corollary. \square

We now prove Proposition 7.2, reproduced here for ease of reference:

Proposition 7.2. Consider a graph game:

$$\mathfrak{G} = (G, \Sigma_A, \Sigma_B, r)$$

with $V(G) = \{v_1, \dots, v_\kappa\}$ and $r = 1$. Let $\mathcal{O}[v]$ denote the orbit of v under some subset H of the graph game automorphism group $\Gamma(\mathfrak{G})$.

There exists a pair of optimal mixed strategies:

$$\sigma_A = (\sigma_A[v_1], \dots, \sigma_A[v_\kappa])$$

$$\sigma_B = (\sigma_B[v_1], \dots, \sigma_B[v_\kappa])$$

where σ_A and σ_B respectively allocate probabilities $\sigma_A[v_i]$ and $\sigma_B[v_i]$ to vertex v_i , such that $\forall i, j \in \{1, \dots, \kappa\}$:

$$\mathcal{O}[v_i] = \mathcal{O}[v_j] \Rightarrow \begin{cases} \sigma_A[v_i] = \sigma_A[v_j] \\ \sigma_B[v_i] = \sigma_B[v_j] \end{cases} \quad (11)$$

Proof. Let τ_A, τ_B be any optimal mixed strategies for Players A and B:

$$\begin{aligned} \tau_A &= (\tau_A[v_1], \dots, \tau_A[v_\kappa]) \\ \tau_B &= (\tau_B[v_1], \dots, \tau_B[v_\kappa]) \end{aligned}$$

Given $\phi \in H$, let:

$$\begin{aligned} \tau_{A,\phi} &= (\tau_A[\phi(v_1)], \dots, \tau_A[\phi(v_\kappa)]) \\ \tau_{B,\phi} &= (\tau_B[\phi(v_1)], \dots, \tau_B[\phi(v_\kappa)]) \end{aligned}$$

So $\tau_{A,\phi}, \tau_{B,\phi}$ are the mixed strategies that respectively allocate probabilities $\tau_A[\phi(v_i)], \tau_B[\phi(v_i)]$ to vertex v_i , for all $i \in \{1, \dots, \kappa\}$.

We first show that if τ_A is an optimal mixed strategy for Player A, then $\tau_{A,\phi}$ is also an optimal mixed strategy for Player A, for all $\phi \in H$.

For a particular $\phi \in H$, consider the graph $G' = (V(G'), E(G'))$ and the sets Σ'_A, Σ'_B constructed by relabelling the vertices of G such that:

$$v_i \mapsto \phi^{-1}(v_i), \quad \forall i \in \{1, \dots, \kappa\}$$

$\tau_{A,\phi}$ and $\tau_{B,\phi}$ are clearly optimal mixed strategies for the resulting game $\mathfrak{G}' = (G', \Sigma'_A, \Sigma'_B, r)$.

However, since $\phi \in H \subseteq \Gamma(\mathfrak{G})$, Σ_A, Σ_B are invariant under ϕ (see Definition 7.2) and thus:

$$\Sigma'_A = \Sigma_A, \quad \Sigma'_B = \Sigma_B$$

Also, because ϕ is an automorphism, we have:

$$V(G') = V(G), \quad E(G') = E(G)$$

Therefore $\mathfrak{G}' = \mathfrak{G}$ and so $\tau_{A,\phi}$ and $\tau_{B,\phi}$ are optimal mixed strategies of \mathfrak{G} for Players A and B respectively, as required.

Now, let:

$$\begin{aligned} \sigma_A &= \frac{1}{|H|} \sum_{\phi \in H} \tau_{A,\phi} \\ \sigma_B &= \frac{1}{|H|} \sum_{\phi \in H} \tau_{B,\phi} \end{aligned}$$

These mixed strategies are optimal mixed strategies for \mathfrak{G} , because any weighted average of optimal mixed strategies is itself an optimal mixed strategy.

It remains to prove that σ_A and σ_B satisfy Property (11). Observe that, $\forall i \in \{1, \dots, \kappa\}$:

$$\begin{aligned} \sigma_A[v_i] &= \frac{1}{|H|} \sum_{j=1}^{\kappa} |H[v_i, v_j]| \tau_A[v_j] \\ \sigma_B[v_i] &= \frac{1}{|H|} \sum_{j=1}^{\kappa} |H[v_i, v_j]| \tau_B[v_j] \end{aligned} \quad (31)$$

where $H[v_i, v_j]$ is defined as in Corollary 4.

Using the corollary, (31) can be rewritten as follows, $\forall v \in V(G)$:

$$\begin{aligned} \sigma_A[v] &= \frac{1}{|\mathcal{O}[v]|} \sum_{w \in \mathcal{O}[v]} \tau_A[w] \\ \sigma_B[v] &= \frac{1}{|\mathcal{O}[v]|} \sum_{w \in \mathcal{O}[v]} \tau_B[w] \end{aligned}$$

This demonstrates that σ_A and σ_B satisfy Property (11) and thus proves Proposition 7.2. \square

References

Alpern, S. (1974). *Differential games and control theory. The search game with mobile hider on the circle*. New York: Dekker, pp. 181–200.

Alpern, S. (1985a). Search for point in interval, with high-low feedback. *Mathematical Proceedings of the Cambridge Philosophical Society*, 98(3), 569–578.

Alpern, S. (1985b). The search value of a network. *Networks*, 15(2), 229–238.

Alpern, S. (1992). Infiltration games on arbitrary graphs. *Journal of Mathematical Analysis and Applications*, 163(1), 286–288.

Alpern, S. (1995). The rendezvous search problem. *SIAM Journal on Control and Optimization*, 33(3), 673–683.

Alpern, S. (2000). Asymmetric rendezvous search on the circle. *Dynamics and Control*, 10(1), 33–45.

Alpern, S. (2002). Rendezvous search: A personal perspective. *Operations Research*, 50(5), 772–795.

Alpern, S. (2008). Hide-and-seek games on a tree to which Eulerian networks are attached. *Networks*, 52(3), 162–166.

Alpern, S. (2010). Search games on trees with asymmetric travel times. *SIAM Journal on Control and Optimization*, 48(8), 5547–5563.

Alpern, S., & Asic, M. (1986). Ambush strategies in search games on graphs. *SIAM Journal on Control and Optimization*, 24(1), 66–75.

Alpern, S., Baston, V., & Gal, S. (2008a). Network search games with immobile hider, without a designated searcher starting point. *International Journal of Game Theory*, 37(2), 281–302.

Alpern, S., & Baston, V. J. (2005). Rendezvous on a planar lattice. *Operations Research*, 53(6), 996–1006.

Alpern, S., Baston, V. J., & Essegai, S. (1999). Rendezvous search on a graph. *Journal of Applied Probability*, 36(1), 223–231.

Alpern, S., Baston, V. J., & Gal, S. (2009). Searching symmetric networks with utilitarian-postman paths. *Networks*, 53(4), 392–402.

Alpern, S., & Beck, A. (2000). Pure strategy asymmetric rendezvous on the line with an unknown initial distance. *Operations Research*, 48(3), 498–501.

Alpern, S., Fokkink, R., & Kikuta, K. (2010). On Ruckle's conjecture on accumulation games. *SIAM Journal on Control and Optimization*, 48(8), 5073–5083.

Alpern, S., Fokkink, R., Lindelauf, R., & Olsder, G.-J. (2008b). The Princess and Monster game on an interval. *SIAM Journal on Control and Optimization*, 47(3), 1178–1190.

Alpern, S., Fokkink, R., Timmer, M., & Casas, J. (2011a). Ambush frequency should increase over time during optimal predator search for prey. *Journal of the Royal Society Interface*, 8(64), 1665–1672.

Alpern, S., & Gal, S. (1995). Rendezvous search on the line with distinguishable players. *SIAM Journal on Control and Optimization*, 33(4), 1270–1276.

Alpern, S., & Gal, S. (2002). Searching for an agent who may or may not want to be found. *Operations Research*, 50(2), 311–323.

Alpern, S., & Howard, J. V. (2000). Alternating search at two locations. *Dynamics and Control*, 10(4), 319–339.

Alpern, S., & Lim, W. S. (1998). The symmetric rendezvous-evasion game. *SIAM Journal on Control and Optimization*, 36(3), 948–959.

Alpern, S., Morton, A., & Papadaki, K. (2011b). Patrolling games. *Operations Research*, 59(5), 1246–1257.

Anderson, E. J., & Aramendia, M. A. (1990). The search game on a network with immobile hider. *Networks*, 20(7), 817–844.

Anderson, E. J., & Aramendia, M. A. (1992). A linear programming approach to the search game on a network with mobile hider. *SIAM Journal on Control and Optimization*, 30(3), 675–694.

Arnold, R.D. (Jan 1962). Avoidance in one dimension: a continuous-matrix game. Interim Research Memorandum – 10 277843, Operations Evaluation Group, Washington DC, USA.

Basilico, N., Gatti, N., & Amigoni, F. (2012). Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence*, 184, 78–123.

Baston, V. J., & Bostock, F. A. (1985). A high-low search game on the unit interval. *Mathematical Proceedings of the Cambridge Philosophical Society*, 97(2), 345–348.

Baston, V. J., Bostock, F. A., & Ferguson, T. S. (1989). The number hides game. *Proceedings of the American Mathematical Society*, 107(2), 437–447.

Baston, V. J., & Garnae, A. Y. (2000). A search game with a protector. *Naval Research Logistics*, 47(2), 85–96.

Berwanger, D. (2007). Admissibility in infinite games. In W. Thomas & P. Weil (Eds.), *STACS 2007. Lecture notes in computer science* (Vol. 4393, pp. 188–199). Berlin: Springer-Verlag.

Blackwell, D., & Girshick, M. A. (1954). *Theory of games and statistical decisions*. New York: Dover Publications (1979).

Bondy, J. A., & Murty, U. S. R. (1976). *Graph theory with applications*. London: Macmillan Press.

Bordelon, D. J. (1975). Solutions of elementary problems – E 2469 – Editor's comment. *The American Mathematical Monthly*, 82(5), 522.

Börgers, T. (1992). Iterated elimination of dominated strategies in a Bertrand-Edgeworth model. *Review of Economic Studies*, 59(1), 163–176.

Buyang, C. (1995). Search-hide games on trees. *European Journal of Operational Research*, 80(1), 175–183.

Chkhartishvili, A. G., & Shikin, E. V. (1995). Simple search games on an infinite circular cylinder. *Mathematical Notes*, 58(5), 1216–1222.

Cohen, N. D. (1966). An attack-defense game with matrix strategies. *Naval Research Logistics Quarterly*, 13(4), 391–402.

- Dagan, A., & Gal, S. (2008). Network search games, with arbitrary searcher starting point. *Networks*, 52(3), 156–161.
- Danskin, J. M. (1990). On the cookie-cutter game - search and evasion on a disk. *Mathematics of Operational Research*, 15(4), 573–596.
- Dendris, N. D., Kirousis, L. M., & Thilikos, D. M. (1997). Fugitive-search games on graphs and related parameters. *Theoretical Computer Science*, 172(1), 233–254.
- Numpy Developers (2012). NumPy. <<http://numpy.scipy.org>> Accessed 06.06.12.
- Dufwenberg, M., & Stegeman, M. (2002). Existence and uniqueness of maximal reductions under iterated strict dominance. *Econometrica*, 70(5), 2007–2023.
- Eagle, J. N., & Washburn, A. R. (1991). Cumulative search-evasion games. *Naval Research Logistics*, 38(4), 495–510.
- Evans, R. (1975). Solutions of elementary problems - E 2469. *The American Mathematical Monthly*, 82(5), 521–522.
- Foreman, J. G. (1977). Differential search games with mobile hider. *SIAM Journal on Control and Optimization*, 15(5), 841–856.
- Python Software Foundation (2012). Python programming language <www.python.org>. Accessed 06.06.12.
- Gal, S. (1974). A discrete search game. *SIAM Journal on Applied Mathematics*, 27(4), 641–648.
- Gal, S. (1978). A stochastic search game. *SIAM Journal on Applied Mathematics*, 34(1), 205–210.
- Gal, S. (1979). Search games with mobile and immobile hider. *SIAM Journal on Control and Optimization*, 17(1), 99–122.
- Gale, D., & Glassey, C. R. (1974). Problems and solutions - E 2469. *The American Mathematical Monthly*, 81(4), 405.
- Garnaev, A. Y. (1991). Search game in a rectangle. *Journal of Optimization Theory and Applications*, 69(3), 531–542.
- Garnaev, A. Y. (1992). A remark on the princess and monster search game. *International Journal of Game Theory*, 20(3), 269–276.
- Grometstein, A. A., & Shoham, D. (1989). Colonel Richard's game. *The Lincoln Laboratory Journal*, 2(2), 235–246.
- Hausken, K. (2010). Defense and attack of complex and dependent systems. *Reliability Engineering & System Safety*, 95(1), 29–42.
- Hausken, K. (2011). Protecting complex infrastructures against multiple strategic attackers. *International Journal of Systems Science*, 42(1), 11–29.
- Hohzaki, R. (2007). Discrete search allocation game with false contacts. *Naval Research Logistics*, 54(1), 46–58.
- Hohzaki, R. (2008). A search game taking account of attributes of searching resources. *Naval Research Logistics*, 55(1), 76–90.
- Hohzaki, R. (2009). A cooperative game in search theory. *Naval Research Logistics*, 56(3), 264–278.
- Hohzaki, R., & Iida, K. (2000). A search game when a search path is given. *European Journal of Operational Research*, 124(1), 114–124.
- Hortala-Vallve, R., & Llorente-Saguer, A. (2012). Pure strategy Nash equilibria in non-zero sum colonel Blotto games. *International Journal of Game Theory*, 41(2), 331–343.
- Kikuta, K. (2004). A search game on a cyclic graph. *Naval Research Logistics*, 51(7), 977–993.
- Kikuta, K., & Ruckle, W. H. (1994). Initial point search on weighted trees. *Naval Research Logistics*, 41(6), 821–831.
- Kikuta, K., & Ruckle, W. H. (1997). Accumulation games, Part 1: Noisy search. *Journal of Optimization Theory and Applications*, 94(2), 395–408.
- Kikuta, K., & Ruckle, W. H. (2002). Continuous accumulation games on discrete locations. *Naval Research Logistics*, 49(1), 60–77.
- Kikuta, K., & Ruckle, W. H. (2007). Rendezvous search on a star graph with examination costs. *European Journal of Operational Research*, 181(1), 298–304.
- Kikuta, K., & Ruckle, W. H. (2010). Two point one sided rendezvous. *European Journal of Operational Research*, 207(1), 78–82.
- Leyton-Brown, K., & Shoham, Y. (2008). Essentials of game theory: A concise, multidisciplinary introduction. In *Synthesis lectures on artificial intelligence and machine learning*. San Rafael, California: Morgan & Claypool.
- Lim, W. S., & Alpern, S. (1996). Minimax rendezvous on the line. *SIAM Journal on Control and Optimization*, 34(5), 1650–1665.
- Mavronicolas, M., Papadopoulos, V., Philippou, A., & Spirakis, P. (2008). A network game with attackers and a defender. *Algorithmica*, 51(3), 315–341.
- Morris, P. (1994). *Introduction to game theory*. New York: Universitext Springer.
- Neumann, P. M., Stoy, G. A., & Thompson, E. C. (1994). *Groups and geometry*. Oxford, UK: Oxford Science Publications, Oxford University Press.
- Neuts, M. F. (1963). A multistage search game. *Journal of the Society for Industrial and Applied Mathematics*, 2(2), 502–507.
- Owen, G., & McCormick, G. H. (2008). Finding a moving fugitive. A game theoretic representation of search. *Computers & Operations Research*, 35(6), 1944–1962.
- Pavlović, L. (1995). A search game on the union of graphs with immobile hider. *Naval Research Logistics*, 42(8), 1177–1189.
- Peshkin, M. (Ed.). (Dec 1994). The rendezvous and coordinated search problems. In *IEEE - 33rd conference on decision and control, Lake Buena Vista, FL*.
- Powell, R. (2007). Defending against terrorist attacks with limited resources. *American Political Science Review*, 101(3), 527–541.
- Powell, R. (2009). Sequential, nonzero-sum Blotto: Allocating defensive resources prior to attack. *Games and Economic Behavior*, 67(2), 611–615.
- Reijnierse, J. H., & Potters, J. A. M. (1993). Search games with immobile hider. *International Journal of Game Theory*, 21(4), 385–394.
- Roberson, B. (2006). The Colonel Blotto game. *Economic Theory*, 29(1), 1–24.
- Rogers, S., Sedghi, A., & Evans, L. (2011). UK riots: Every verified incident - Interactive map. The Guardian (11 August). <www.guardian.co.uk/news/datablog> Accessed 20.04.12.
- Ruckle, W. H. (1983). *Geometric games and their applications. No. 82 in research notes in mathematics*. Boston: Pitman.
- Ruckle, W. H. (1990). The gold-mine game. *Journal of Optimization Theory and Applications*, 64(3), 641–650.
- Ruckle, W. H. (2007). Rendez-vous search on a rectangular lattice. *Naval Research Logistics*, 54(5), 492–496.
- Ruckle, W. H., & Kikuta, K. (2000). Continuous accumulation games in continuous regions. *Journal of Optimization Theory and Applications*, 106(3), 581–601.
- Shubik, M., & Weber, R. J. (1981). Systems defense games - Colonel-Blotto, command and control. *Naval Research Logistics*, 28(2), 281–287.
- Steen, L. A., & Seebach, J. A. Jr. (1970). *Counterexamples in topology*. New York: Springer-Verlag.
- Subelman, E. J. (1981). A hide-search game. *Journal of Applied Probability*, 18(3), 628–640.
- Sutherland, W. A. (1975). *Introduction to metric and topological spaces*. Oxford: Oxford University Press.
- Tambe, M. (2012). *Security and game theory: Algorithms, deployed systems, lessons learned* (1st ed.). New York: Cambridge University Press.
- Thomas, L. C., & Washburn, A. R. (1991). Dynamic search games. *Operations Research*, 39(3), 415–422.
- Tofias, M., Merolla, J., & Munger, M. (Apr 2007). Of colonels and generals: Understanding asymmetry in the Colonel Blotto game, prepared for MPSA 2007 Panel 34-6 "Computational Models". <www.uwm.edu/tofias>.
- White, L. V. (Oct 1994). Games of strategy on trees [Unpublished], Technical Report Series, Imperial College London.
- Zoroa, N., Fernández-Sáez, M. J., & Zoroa, P. (1999). A game related to the number of hides game. *Journal of Optimization Theory and Applications*, 103(2), 457–473.
- Zoroa, N., Fernández-Sáez, M. J., & Zoroa, P. (2004). Search and ambush games with capacities. *Journal of Optimization Theory and Applications*, 123(2), 431–450.
- Zoroa, N., Fernández-Sáez, M. J., & Zoroa, P. (2012). Patrolling a perimeter. *European Journal of Operational Research*, 222(3), 571–582.
- Zoroa, N., & Zoroa, P. (1993). Some games of search on a lattice. *Naval Research Logistics*, 40(4), 525–541.
- Zoroa, N., Zoroa, P., & Fernández-Sáez, M. J. (1999). A generalization of Ruckle's results for an ambush game. *European Journal of Operational Research*, 119(2), 353–364.
- Zoroa, N., Zoroa, P., & Fernández-Sáez, M. J. (2001). New results on a Ruckle problem in discrete games of ambush. *Naval Research Logistics*, 48(1), 98–106.
- Zoroa, N., Zoroa, P., & Fernández-Sáez, M. J. (2003). Raid games across a set with cyclic order. *European Journal of Operational Research*, 145(3), 684–692.