UNIVERSITY OF SOUTHAMPTON

Faculty of Mathematical Studies

DISCRETE GAMES OF INFILTRATION

by

John Michael Auger

Thesis submitted for the degree of Doctor of Philosophy

October 1991

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

Faculty of Mathematical Studies

Doctor of Philosophy

DISCRETE GAMES OF INFILTRATION

by John Michael Auger

Gal has suggested the following scenario for a two-person game. An Infiltrator starts at the first of an ordered set of p points. At discrete intervals of time $t = 1, 2, \ldots$ he chooses to move to one of the adjacent points or to stay where he is. A Guard starts from any point and at each of the same intervals of time moves to a point up to u points away. He then searches for the Infiltrator, detecting him with probability μ if the players are at the same point, and with probability zero otherwise. Neither player is aware of his opponent's moves unless detection occurs. At the last of the p points the Infiltrator is safe from detection.

We look at a number of zero-sum games which are based on this scenario. These include both infinite move games and games in which the number of moves is restricted by a time limit of n. The objective of the Infiltrator is either to reach the last point undetected, or just to evade the Guard. We show that the infinite move games have mixed strategy solutions which can be constructed from solutions to finite move games.

In addition, we study a further set of games in which the Infiltrator is also safe from detection at the first point. His objective then is to reach the last point undetected. In this context we extend the work of Lalley to a more general set of points.

Some examples of optimal strategies are presented. Finally we discuss some possible generalisations to other discrete infiltration games.

Acknowledgements

I wish to express my sincere thanks to Fred Bostock and Vic Baston my joint-supervisors. Your guidance throughout the past three years has been invaluable, and your patience and hard work limitless. You make a superb team!

My deepest appreciation is due to Sunita my wife. Your loving support and encouragement has been unfailing. Thank you for believing in me.

I am also most grateful to my parents. Thank you Mum and Dad for always backing me but never pushing me. My head for maths was a gift from you both.

Finally, I want to thank my Heavenly Father. You are the source of all life and the fount of all knowledge. I am nothing compared to you, and yet you love me.

"The fear of the Lord is the beginning of knowledge" (Proverbs I vii).

Contents

Chapter 1

INTRODUCTION

1.1 Search for an Infiltrator

On page 96 of the book *Search games* [31], Shmuel Gal poses the following problem:

> "Assume that at time $t = 0$ an infiltrator enters the set Q through a known point O on the boundary of Q, and for all $t > 0$ moves inside Q. Suppose that the searcher has to defend a 'sensitive zone' $B \subset Q$, so that he wishes to maximize the probability of capturing the infiltrator before he reaches the boundary of B (the infiltrator has the opposite goal). What are the optimal strategies of both players in this game?"

This is the opening paragraph of a section which is entitled "Search for an Infiltrator", in which Gal presents "some interesting unsolved problems which have some resemblance to the princess and monster game". We shall leave the princess and the monster until later. For the moment we concentrate on the infiltrator. As it stands the above problem is very loosely formulated, and surely this is the author's intention. The mechanics of both the players' movement and of capture have not been formulated. They are left to the imagination of the interested reader. Nevertheless, Gal does suggest one particular 'way-in' to the problem.

> "A simpler, discrete problem with a similar flavor is the following. The search set is an array of n ordered cells. At time $t = 0$, both the searcher and the hider are located in cell number 1. At the end of each time unit, the searcher can move to any cell with distance

> not exceeding a certain integer $k \geq 1$, while the hider can only move a distance of 1. Both players then stay at the chosen cell for the next time unit. The probability of capture is p for each time unit in which both players occupy the same cell (independently of previous history) and zero otherwise. The hider wins in the case that he reaches cell number n (in finite time) without getting captured, and loses otherwise." (*Search Games*, pages 96–97.)

This 'simpler' problem is the subject of this work. We consider different ways in which to approach this discrete one-dimensional problem. Later we discuss how this can then be applied more widely to other discrete infiltration games.

For the rest of this chapter we attempt to provide some perspective by tracing the developments that both allow and motivate us to tackle this problem. However, before leaving the particular source of the problem, we record one further comment made by Gal in the general context of the problem of infiltration.

> "It can be easily seen that... it is not a good policy for the hider to move in a straight line using his maximal velocity. A policy which does seem to be good for the hider is to move randomly for a certain period of time and only then to use his maximal velocity." (op. cit. page 98.)

We shall review this comment in the light of our investigations.

1.2 Search Theory

The beginnings of what has become known as Search Theory are to be found during the second world war. The US Navy wanted to know how to go about looking for enemy submarines. It was with this particular objective that their Anti-Submarine Warfare Operations Group was formed. They successfully began developing some practical ideas. In 1946, after the war had ended, one member of the group, Bernard Koopman, produced a report on their conclusions which he entitled 'Search and Screening' [43].

For a decade or so Search Theory seemed to attract little interest elsewhere. Although Koopman did extend parts of the earlier work in three papers published in 1956-57 [44, 45, 46], the bibliographies on Search Theory contain little material earlier than the mid-1960's. Exceptions to this include extensions of Koopman's work by Charnes and Cooper [18], and de Guenin [23]. Slightly different is the 1959 work by Blachman [14] and Blachman and Proschan [15] on discrete problems involving multiple targets with unknown arrival times. From 1963 Search Theory experienced a boom and by 1968 the selective bibliography of Dobbie [24] listed 52 sources in connection with search. By 1980, an internal report by Strumpfer [63] included over 400 books and articles related to Search Theory. The slow start illustrated above does have an explanation. Koopman's original report was classified until 1958!

Search Theory still has the same simple objectives with which it began. The introduction to a recent special issue of the journal *Naval Research Logistics* on Search Theory clearly states this.

> "After several decades of subsequent development, search problems are still largely of the same form as in 1942: a single target is lost, and the problem is to find it effectively with fixed resources." [62]

The field of Search Theory is commonly split into four areas. Problems are classified as either stationary or moving target, and either one or two-sided. The first of these divisions is easily understood. The second is really a question of the intelligence of the target. If it is assumed that the target has a known distribution, or that it moves with some kind of random movement, then this is known as one-sided search. If on the other hand, the target is assumed to have some intelligence which it uses to make detection as hard as possible, this is known as a two-sided search problem. It is of course the two-sided problems which have been considered using the techniques of Games Theory. Thus problems of two-sided search are known as Search Games. The target is then more correctly known as the evader, and is designated no longer by 'it' but by 'he'. It is against the background of the quarter of Search Theory known as mobile evader Search Games that our study of infiltration games must begin. We attempt to sketch this background in the following section.

Standard texts on Search Theory include those of Stone [61] (first published in 1975, but with a 2nd edition in 1989), Washburn [66] and Haley and Stone [35]. An extensive literature survey can also be found in Benkoski, Monticino and Weisinger [11].

1.3 Search Games

The two-sided game theoretical approach to search is outlined below. Some detail is inevitable yet the principles are really simple. In fact, it has been suggested that in Search Games,

> "(the) 'hide and seek' games which we used to play in our childhood are formulated as mathematical problems" [20, page 33]

The theory of two-person zero-sum games is described in more detail in Chapter 2. In Search Games the evader and the searcher are considered to be two players who choose a strategy for evasion and search respectively. These strategies are then evaluated in some way and the effectiveness of the resulting search attempt measured. It is assumed that both evader and searcher use the same measure of effectiveness. Thus a 'good' result for one of them is always equally 'bad' for the other. Therefore, given a strategy for each player, this measure of effectiveness is some numerical quantity, called the payoff, which we may consider one player gains and the other loses (or equivalently one pays to the other). We choose to assume that whenever the payoff is positive then it is paid by the searcher to the evader, although this choice is arbitrary. Thus, the evader's objective is to maximise the payoff. He is then known as the maximiser. Similarly the searcher is known as the minimiser.

There are several different possibilities for the payoff. It may be the time required before the search is effective (which may, of course, be infinite if the searcher cannot eventually ensure success). Similarly it may be the amount of resources expended by the searcher before the search is effective. If there is doubt about the certainty of eventual success, it may be the probability that the search fails. Other payoffs which involve discounting the searcher's reward can also be considered.

The strategies themselves may be many different things. Both time and space may be either discrete or continuous. The search space can be one-dimensional (a ship on a narrow channel), two-dimensional (a ship on the ocean), three-dimensional (a submarine in the ocean), and so on. It may even be a graph. When the problem is largely continuous it is often considered to be a Differential Game (see the classic book by Isaacs [40] for these), so called because of the differential equations used to define the players' motion. On the other hand, when the problem is largely discrete, more combinatorics and probability theory are usually required (see *Geometric Games* by Ruckle [59]).

Finally, the detection process itself must be described. In Search Theory generally, it has often been assumed that there is an exponential detection function which the searcher applies at different points in the search space. Thus, the probability of detection decreases the further the evader is from the point of application. In discrete space games it has often been assumed that the probability of detection is one if the players simultaneously occupy the same state and zero otherwise. However, more general frameworks are easy to imagine.

It seems that a large number of Search Games cluster around a relatively small number of particular 'real-life' military problems. More examples will be noted below, but one of the earliest and probably the most celebrated, is the Princess and Monster Game. This is the game that Gal explicitly mentions (see section 1) in connection with games of infiltration. The game was first posed by Rufus Isaacs as follows.

> "The monster P searches for the princess E, the time required being the payoff. They are both in a totally dark room Q (of any shape), but they are each cognizant of its boundary (possibly through small light admitting perforations high in the walls). Capture means the distance $PE \leq r$, a quantity small in comparison with the dimension of Q. The monster, supposed highly intelligent, moves with simple motion at a known speed. We permit the princess full freedom of locomotion." [40, example 12.4.1, page 349]

Note that by simple motion here is meant motion with no other restriction except that the speed is constant. A simpler version of this problem was also suggested (op. cit. page 350) in which the players move on the boundary of a circle. This problem was solved in the early 1970's by Alpern [1], Foreman [29] and Zelikin [68]. A discrete version of the problem was also solved by Wilson [67]. Eventually, in 1979, Gal [30] solved the Princess and Monster Problem for any convex multidimensional area.

In this work we are going to follow Gal's suggestion and adapt a discrete time and discrete space approach to games of infiltration. We shall discuss discrete Search Games a little below. It is Gal himself who has probably contributed the most to the non-differential approach to continuous Search Games. His book *Search games* [31] was the first complete work dedicated to the subject. This includes (see Appendix 1, pages 181–188) an important result concerning the existence of solutions to these problems (later extended by Alpern and Gal [2]).

In 1989, this earlier work was updated with the publication of a survey [32] of more recent developments in the field of continuous Search Games.

1.4 Discrete Games

It is obvious that we have been gradually narrowing down our field of vision. This is the final stage. We consider the important contributions that have been made in the area of discrete Search Games. An early problem to be considered here was the Bomber-Battleship Duel, which like the Princess and Monster Game was first formulated by Rufus Isaacs. The Bomber-Battleship Duel is the group name of a collection of problems all of which involve a m move time-lag, where m is some positive integer. Lee and Lee [51] describe the basic problem as follows.

> "The minimizing player, called the ship, is constrained to move on the integer lattice of the real line. In each time unit, it must move either one unit distance to the left or one unit distance to the right. The aim of the ship is to manoeuvre so as to minimize the probability of being hit by a bomb. The maximising player, called the bomber, is loaded with one bomb and it flies overhead trying to drop the bomb on the ship. The bomber can observe the movement of the ship for as long as he likes before dropping the bomb on any position. The main source of difficulty is introduced by assuming the bomb takes m units of time to reach the ship from the bomber. Thus the problem facing the bomber is to anticipate some future position of the ship to drop the bomb after observing all its past moves. The ship does not know when or where the bomb is dropped until the bomb hits. The payoff to the bomber is one if the bomb hits and zero otherwise." (page 867)

For $m = 1$ the game is trivial. The value is $1/2$ and the following strategies are optimal. The ship always moves left or right with equal probability. At any time the bomber knows that the ship is at one of two locations. Thus, at any time he drops the bomb on these locations with equal probability. For $m = 2$ the problem is non-trivial. It has been solved, using various different methods by Dubins [26], Karlin [41], Isaacs and Karlin [38], Isaacs [39] and Fergusson [28]. The value is $(3 - \sqrt{5})/2 \simeq 0.38197$. For $m = 3$ the problem is still open. The game has been shown to have a value v and bounds on v have been extensively investigated by Matula [52], Bram [17] and Lee and Lee [51, 50]. Currently, the sharpest bounds are $0.28648 \leq v \leq 0.2883686$.

Other time lag games can also be included under the title of the Bomber-Battleship Duel. Lee [48, 49], Sakaguchi [60] and Garnaev [33] have all introduced the concept of a safe region to which the battleship should eventually move. This work is considered in relation to games of infiltration in the following section. Bernhard, Colomb and Papavassilopoulos [12] have considered a time-lag game using stochastic formulations. Finally Baston and Bostock have considerd the problem as a recursive game on both a finite [6, 8] and an infinite number of states [10].

More recently Olsder and Papavassilopoulos [55] have initiated the study of what are known as Searchlight Games (these are further developed in [56, 57]). We decribe Searchlight Games by quoting from a recent article on the subject by Baston and Bostock [9].

> "In a Searchlight Game at least one of the two players has a searchlight which can be flashed a certain number of times within a given time period. A flash of the searchlight illuminates a region of known shape. The objective of a player with a searchlight is to catch his opponent in the region at the time of the flash."

Both partnerships have concentrated upon Searchlight Games played on the discrete boundary of a circle.

We leave both the Bomber-Battleship Duel and the Searchlight Games to consider the literature that surrounds a problem most recently known as the Flaming Datum problem. This may be described as follows. A helicopter (say) attempts to detect or destroy a submarine that has recently revealed its position by torpedoing a ship (the flaming datum). The helicopter uses a dipping sonar or an explosive device. The submarine's only defence is its ability to dive out of sight of the helicopter.

The Flaming Datum problem has been appoached in a number of different ways. The work by Meinardi [53] in 1964 is perhaps the earliest to consider. This is a discrete one-dimensional version of the problem. Meinardi assumes that the evader hides in a row of boxes. At each discrete time interval the searcher picks a box and examines it. The probability of finding the evader when the correct box is searched is equal to μ, where $\mu \leq 1$. If the evader is not found he may either move to a neighbouring box or remain in the same box, and the next stage is played. Importantly, both the searcher and the evader are aware of which box the searcher examines. A limit is placed on the number of stages and thus the game is finite. The payoff is the probability that the evader is found.

Another early formulation is that of Danskin [22]. Under the title 'A helicopter versus submarine search game', he considers a two-dimensional version of the Flaming Datum problem. In contrast to Meinardi he assumes that the evader gains no useful information about what the searcher is doing. Thus he assumes that the evader will not change his speed or his velocity once the game has commenced.

In the book "Geometric games and applications" [59], Bill Ruckle considers a range of different two-person zero-sum games. A considerable proportion of these games are related to problems of search and also ambush. The Lattice Ambush Game (pp. 34–37) and the Lattice Search Game (pp. 53–54) are particularly relevant to us. The discrete one-dimensional formulation of Meinardi [53] can be transferred to a lattice by considering the players to be the mobile occupants of a space-time lattice. Ruckle's games do not cover probabilities of detection that are less than one. If the searcher and the evader occupy the same point on the lattice then the evader is found. We note, however, that even these simplified problems are only partially solved by Ruckle. We are not aware that any further progress has been made with them.

Baston and Bostock [7] also consider a one-dimensional game which is closely related to the Flaming Datum problem. Their problem involves continuous motion and limits the number of searches (in fact these are anti-submarine bombs) available to the searcher. Cheong [19] approximates the two-dimensional problem by limiting to a finite number the mixed strategies available to the players. He also provides some useful discussion on the Flaming Datum problem in general.

The most recent work explicitly addressed to the Flaming Datum problem is that of Thomas and Washburn [65]. They model the two-dimensional problem with a finite grid on which the players are assumed to be able to move without restriction. Again it is assumed that both players are aware of the location and time of all past searches. This work is presented as only a rough approximation to the Flaming Datum problem due to the unrealistic motion of the players. Thus it is considered to be a simplification of the restricted motion case.

It has become clear that the games above use different assumptions in modelling the Flaming Datum problem. They approach it from different directions. If we consider just the finite versions of the problem we can identify four particular areas in which assumptions must be made. The first concerns the *players' motion.* In order to be a realistic model, some speed restrictions seem to be desirable for the players. However, as documented in Ruckle ([59] see especially the pages referred to above) and admitted in Thomas and Washburn [65], solutions are easier to come by when the motion is unrestricted. The second set of assumptions concerns the *detection probability.* This could be a function not only of current position and time, but also of the previous history of the game. However, in order to get some results, the detection probability has been greatly simplified. At times, as in Ruckle [59] it has even been taken to be uniformly equal to one. Thirdly, there is the question of the *initial position.* The players may start at known positions. Alternatively, the starting positions may be given by two probability distributions which are known to both the players. In fact it seems that the Flaming Datum problem implies that the evader must initially give away his location (eg. by the sinking of a ship), and therefore a fixed starting position might seem most appropriate. However, as some significant time period may elapse before the searcher comes into effective range, this may be questioned.

The fourth, and perhaps the most significant, set of assumptions that are made concern the players' *'knowledgeability'.* This is a term which is used to describe the information conditions which are present in a game. If, during the course of a game, a player gains useful information about his opponent's position (other than if detection occurs) then he is said to be knowledgeable in some way. This information may be incomplete. For instance, the evader might realise the searcher's position only when they are in the same state. The information may also be delayed, as in the case of the Bomber-Battleship Duel. Most of the work mentioned above assumes that at least one of the players is considerably knowledgeable. Exceptions to this are Ruckle [59], Danskin [22], and Baston and Bostock [7] (although in this last case, while it is initially assumed that neither player is knowledgeable, it is also noted that the solution would be unchanged if the evader were aware of his opponent's moves). Throughout this work we shall assume that neither player is knowledgeable. Thinking of a situation with a submarine and a helicopter, this does not seem unreasonable.

Finally we make note of the 'Cumulative Search Evasion Games' (CSEGs) recently considered by Eagle and Washburn [27]. These games differ from all of the above in the form of payoff which is assumed. They involve a cumulative score which builds up over the (finite) period of the game. This may be, for example, the number n, of meetings between the players. The conditions of movement are general enough to include all of the discrete games above. Note that, if we relate this to the games in which there is a constant probability μ, $\mu \leq 1$, of detection then the total probability of detection is given by $1-(1-\mu)^n$. Therefore the framework of CSEGs could be adapted to the Flaming Datum problem as well. However, the methods developed are admitted to be tailored specifically to the cumulative form of payoff. They do remark that, if $\mu = (1-e)^{-1}$, then $1-(1-\mu)^n = 1-e^{-1}$ which is approximately equal to n when n is small. "So CSEGs can be regarded as first order approximations to detection games" (page 496).

1.5 Infiltration and the Flaming Datum Problem

We conclude this chapter as we began it, with Gal's discrete infiltration problem. We can now observe that although infiltration seems to be a relatively new area of application, the underlying problems can be considered as generalisations of the Flaming Datum problem. An extra factor is introduced. The evader has a target which he can head for. If he gets there he is safe from detection. Note that there are two different possibilities for the payoff. If the payoff remains as the probability that the evader is not detected, then this is a true generalisation of the Flaming Datum problem. Indeed, if the safe area is far away it may have no effect on the optimal strategies for the game. On the other hand, if the payoff is taken to be the probability that safety is reached, then this is a distinct problem. When the time available is very large (or infinite) these distinct games may have the same solutions. This question is taken up again for some specific games in later chapters.

Discrete games of infiltration have received little attention. Lee [48, 49] and Sakaguchi [60] have investigated games in which the evader moves on a line and has a safe point or bunker which he must eventually get to. Garnaev [33] has dealt with the similar situation where the evader moves on the infinite two-dimensional integer grid with a general set of safe points. All of this work is based on the assumption that the probability of detection is one (although, in compensation, we note that they also limit the number of detections which the searcher is allowed to make). Attempts have been made to consider an infiltration game involving a probability of detection which may be less than one. Lalley [47] has futher simplified Gal's discrete problem by introducing another safe point where the evader starts the game. Thus, the infiltrator can choose the timing of his infiltration and remain undetected until he acts. This approach is discussed and generalised in Chapter 5. For completeness we note also the author's own work [4, 5] on infiltration games. This is also extended in Chapters 4 and 5. Finally Alpern [3] has recently solved the infinite time 'safe base' game on any discrete graph. This too is mentioned again in Chapter 5.

Chapter 2

FOUNDATIONS

2.1 Introduction

In this chapter we begin the formal analysis. We start by reviewing some elementary concepts of two-person zero-sum game theory, particularly those of solution, value and optimal strategy. In section 3 we turn to the particular games of infiltration which we are going to study. Our first aim is to consider whether solutions may be found for these games. To this end we try to use certain existence theorems, or minimax theorems as they are known.

The definitions contained here are the foundations of all that follows. Therefore in section 8 we discuss some variations on these definitions. This illustrates the relative fragility of some of the properties of these games and also tests the completeness of the set of currently known minimax theorems.

2.2 Two-Person Zero-Sum Games

We shall be very selective here. A good general introduction to Games Theory can be found, for instance, in [64]. We shall restrict ourselves to two-person zero-sum games and even then give only the bare essentials for our requirements. Further eleboration will be provided when needed.

Definition 2.2.1 *A* **two-person zero-sum game** Γ *is an ordered triple*

$$\Gamma = (X, Y, f),$$

where X and Y are sets, and f is a bounded real-valued function on $X \times Y$ the Cartesian product of X and Y.

Henceforth we shall refer to any two-person zero-sum game as simply a **game**. The two players involved we label **player I** and **player II**. The sets X and Y are known as the **strategy sets** of player I and player II respectively. The function f is known as the **payoff function**. If $x \in X$ and $y \in Y$ then $f(x,y)$ is the amount, called the **payoff** or the **payoff to player I**, of some appropriate units of utility, payed to player I (by player II). Thus, if we consider player I (resp. player II) to be able to choose a strategy from his strategy set, he does so with the objective of maximising (minimising) f. Thus player I and player II are sometimes called the **maximiser** and the **minimiser** respectively.

Recall that the payoff function f is bounded. Then fixing $\bar{x} \in X$, for all $y \in Y$, $f(\bar{x},y) \geq \inf_{y\in Y} f(\bar{x},y)$. We call $\inf_{y\in Y} f(\bar{x},y)$ the minimum payoff guaranteed by $\bar{x}$. Similarly, for fixed $\bar{y} \in Y$, $\sup_{x\in X} f(x,\bar{y})$ is the maximum payoff 'guaranteed' by $\bar{y}$. We denote by $\underline{v}(\Gamma)$ and $\overline{v}(\Gamma)$ the **security levels** for player I and player II respectively which are given by $\underline{v}(\Gamma) = \sup_{x\in X}\inf_{y\in Y} f(x,y)$ and $\overline{v}(\Gamma) = \inf_{y\in Y}\sup_{x\in X} f(x,y)$.

Lemma 2.2.2 *For all* $\Gamma = (X,Y,f)$, $\underline{v}(\Gamma) \leq \overline{v}(\Gamma)$.

Proof Clearly, for all $x \in X$ and $y \in Y$, $f(x,y) \leq \sup_{t\in X} f(t,y)$. Thus $\inf_{s\in Y} f(x,s) \leq \sup_{t\in X} f(t,y)$. As $\inf_{s\in Y} f(x,s)$ is independent of y it is $\leq \inf_{s\in Y}\sup_{t\in X} f(t,s) = \overline{v}(\Gamma)$. Finally, $\overline{v}(\Gamma)$ is independent of x, and so $\underline{v}(\Gamma) = \sup_{t\in X}\inf_{s\in Y} f(t,s) \leq \overline{v}(\Gamma)$. This completes the proof.

Definition 2.2.3 *Let* $\Gamma = (X,Y,f)$ *be a game.* Γ *has a* **solution** *if* $\underline{v}(\Gamma) = \overline{v}(\Gamma)$. *This common value is then called* **the value of the game**, *or simply* **the value**, *and is denoted by* $v(\Gamma)$.

Lemma 2.2.4 *Let* $\Gamma = (X,Y,f)$ *be a game with a solution. Then, for all* $\epsilon > 0$, *there exists a pair* $(x_\epsilon, y_\epsilon) \in X \times Y$ *such that, for all* $(x,y) \in X \times Y$,

$$f(x,y_\epsilon) - \epsilon \leq v(\Gamma) \leq f(x_\epsilon,y) + \epsilon.$$

Proof Let $\epsilon > 0$. Then, as $\inf_{y\in Y}\sup_{x\in X} f(x,y) = v(\Gamma)$, there exists $y_\epsilon \in Y$ such that $\sup_{x\in X} f(x,y_\epsilon) \leq v(\Gamma) + \epsilon$. Thus, for all $x \in X$,

$$f(x,y_\epsilon) \leq \sup_{x\in X} f(x,y_\epsilon) \leq v(\Gamma) + \epsilon.$$

Similarly, there exists $x_\epsilon \in X$ such that, for all $y \in Y$,

$$f(x_\epsilon,y) \geq \inf_{y\in Y} f(x_\epsilon,y) \geq v(\Gamma) - \epsilon.$$

The strategies x_ϵ and y_ϵ are known as **ϵ-optimal strategies** for player I and player II respectively. For any game Γ which has a solution, $v(\Gamma)$ is unique but x_ϵ and y_ϵ need not be.

Definition 2.2.5 *Let $\Gamma = (X, Y, f)$ be a game with a solution. A strategy for either player is* **optimal** *if, for all $\epsilon > 0$, it is ϵ-optimal.*

Even when a game has a solution, optimal strategies need not exist. However, if $x^* \in X$ and $y^* \in Y$ are optimal, then for all $(x, y) \in X \times Y$, $f(x, y^*) \leq f(x^*, y^*) \leq f(x^*, y)$. Although x^* and y^* are not necessarily unique, the payoff $f(x^*, y^*)$ is, and satisfies $f(x^*, y^*) = v(\Gamma)$. If X and Y are finite then it can be proved that whenever the game $\Gamma = (X, Y, f)$ has a solution there exist optimal strategies for both players.

When $\Gamma = (X, Y, f)$ does not have a solution it is common to consider the extended game Γ^*. Suitable σ-algebras are taken on X and Y. The extended strategies are probability measures on the resulting measurable spaces. The extended payoff functions are integrals. We shall see this in practice later on in the chapter. If X or Y is finite, then the extended game always has a solution.

Finally if there exist distinct $x, x' \in X$ such that, for all $y \in Y$, $f(x', y) \geq f(x, y)$ then x is said to be **dominated** by x'. Similarly for $y, y' \in Y$, if, for all $x \in X$, $f(x, y') \leq f(x, y)$, then y **dominates** y'. It is clear that the removal of a dominated maximiser strategy or a dominating minimiser strategy does not affect the existence of a solution. Moreover, if the game does have a solution, this does not change the value.

2.3 Strategies I

We now move on to the particular games of infiltration in which we are interested. In this section and the next we introduce the strategy sets and the payoff functions respectively. We define two different games which correspond to the different payoffs mentioned in Chapter 1. In section 5 we begin to extend these games.

Both the games take place on a set of p states, and it is assumed that $p \geq 2$. Let $P = \{1, 2, \ldots, p\}$ and P^∞ the set of all infinite sequences (s_r), $r = 1, 2, \ldots$, in P. We rename the players. In future, player I, will be known as **the Infiltrator** and player II as **the Guard**. Thus the Infiltrator is to be the maximiser when the payoff functions are defined. Both players are assumed to move in discrete time among the states. So we consider their strategies to be elements of P^∞. At each time, while the Infiltrator can only move between adjacent states (eg. from state two to state one, two or three), the Guard can move between states up to u states apart. Thus u is referred to as the **speed of the Guard**. We assume that $u \geq 1$, for otherwise the problems are trivial.

Definition 2.3.1 *The* **strategy sets**, *I_∞ and G_∞, for the Infiltrator and Guard respectively are the subsets of P^∞ given as follows: If $i \in I_\infty$ then*

I(a) $i_1 = 1$,

I(b) $|i_{r+1} - i_r| \leq 1$ *for all* $r \geq 1$, *and*

I(c) *if, for some* $t \geq 1$, $i_t = p$ *then, for all* $r > t$, $i_r = p$.

If $g \in G_\infty$ *then*

G(a) $g_r < p$ *for all* $r \geq 1$,

G(b) $|g_{r+1} - g_r| \leq u$, *for all* $r \geq 1$.

Let us discuss this definition. A strategy for the Infiltrator is a sequence $i_1, i_2, \ldots$ where, for each $n \in \mathbb{N}$, i_n is considered to be the state at which the Infiltrator is located at the nth stage of the game. Similarly for a strategy $g_1, g_2, \ldots$ for the Guard. As detection cannot occur at state p, the target, we ban the Guard from this state. The other conditions on the sets I_∞ and G_∞ derive from the basic properties disussed in the opening chapter. At time one, the Infiltrator is assumed to occupy the first state, and thus, for all $i \in I_\infty$, $i_1 = 1$. The conditions on adjacent terms of any element of I_∞ or G_∞ correspond to the players' speed restrictions. Condition I(c) does *not* ensure that the Infiltrator visits the target. It just guarantees that, if he ever does, he stays there.

In the following section we define two different payoff functions, Π_∞ and $\tilde{\Pi}_\infty$, thus completing the definition of two different games, $\Gamma_\infty = (I_\infty, G_\infty, \Pi_\infty)$ and $\tilde{\Gamma}_\infty = (I_\infty, G_\infty, \tilde{\Pi}_\infty)$. At the end of the chapter there is a section in which we examine some alternatives to these games which are obtained by defining different strategy sets.

We now define a topology on P^∞ and hence on I_∞ and G_∞ and consider the resulting topological spaces. In fact, we define a metric on P^∞. Let $\mathbb{N}^*$ denote the set $\mathbb{N} \cup \{\infty\}$, and consider the two functions n and d defined on $P^\infty \times P^\infty$ as follows: If $x = x_1, x_2, \ldots$ and $y = y_1, y_2, \ldots$ are both elements of P^∞ then

$$n(x,y) = \begin{cases} \min\{r | x_r \neq y_r\} & \text{if } x \neq y \\ \infty & \text{if } x = y. \end{cases}$$

For all $x, y \in P^\infty$, $d(x,y) = n(x,y)^{-1}$ where, by convention, $1/\infty = 0$. Clearly, for all x and y, $0 \leq d(x,y) \leq 1$.

Lemma 2.3.2 *d is a metric on P^∞.*

Proof Let x, y, z denote three arbitrary elements of P^∞. If $x = y$ then $d(x,y) = d(y,x) = 0$, whereas, if $x \neq y$ then $d(x,y) = d(y,x) \geq 0$.

It remains to show that $d(x,y) + d(y,z) \geq d(x,z)$. If x and y are equal then this follows immediately. Therefore let us assume that they are distinct. As x and y first differ at the $(n(x,y))$th term and y and z at the $((n(y,z))$th term, that x and z certainly cannot differ before the first of these terms. Therefore $n(x,z) \geq \min\{n(x,y), n(y,z)\}$, and so,

$$\begin{aligned} d(x,z) &= n(x,z)^{-1} \\ &\leq (\min\{n(x,y), n(y,z)\})^{-1} \\ &= \max\left\{n(x,y)^{-1}, n(y,z)^{-1}\right\} \\ &= \max\{d(x,y), d(y,z)\} \\ &\leq d(x,y) + d(y,z). \end{aligned}$$

Hence we have shown that d is reflexive, symmetric and transitive. Therefore it is a metric.

It can be seen from the above proof that the metric d is also what is known as an ultrametric. A metric u on a set X in an **ultrametric** if, for all $x, y, z \in X$, $u(x,z) \leq \max\{u(x,y), u(y,z)\}$. An ultrametric space has certain interesting properties beyond those of other metric spaces. In particular, in an ultrametric space (X, u) any ϵ-ball, $B_\epsilon(x) = \{x' \in X | u(x', x) < \epsilon\}$ is both open and closed. Moreover, for all $\epsilon > 0$, $x \in X$ and any $x' \in B_\epsilon(x)$, $B_\epsilon(x') = B_\epsilon(x)$. Some additional properties of ultrametric spaces are to be found in Dieudonné [25] (page 38, problem 4) and others in Bourbaki [16] (ex. 2.4 pages 227–228). It is sufficient for us that d possesses the general properties common to all metrics.

As a metric, d induces a topology on every subset of P^∞. In particular, we denote by $\mathcal{U}$ and $\mathcal{V}$ the collections of all open subsets of I_∞ and G_∞ respectively. The topological spaces $I = (I_\infty, \mathcal{U})$ and $G = (G_\infty, \mathcal{V})$ are then referred to as the players' **strategy spaces**.

Lemma 2.3.3 *(P^∞, d) is compact.*

Proof As (P^∞, d) is a metric space, it is compact if and only if every infinite set $Q \subseteq P^\infty$ has a limit point $\sigma \in P^\infty$. We construct such a limit point below. First, for any $Q \subseteq P^\infty$ and any finite sequence $s_1, s_2, \ldots, s_n$ in P, let

$$Q_{s_1,s_2,\ldots,s_n} = \{q \in Q \mid q_r = s_r \text{ for all } 1 \leq r \leq n\}.$$

Let Q be an infinite subset of P^∞. Observe that $Q = \bigcup_{s_1 \in P} Q_{s_1}$. As this union is finite there exists at least one value of s_1 such that Q_{s_1} is infinite. Let σ_1 denote the least such value.

Now suppose that $\sigma_1, \sigma_2, \ldots, \sigma_k$ is a finite sequence in P and that $Q_{\sigma_1,\sigma_2,\ldots,\sigma_k}$ is infinite. As above $Q_{\sigma_1,\sigma_2,\ldots,\sigma_k} = \bigcup_{s_{k+1} \in P} Q_{\sigma_1,\sigma_2,\ldots,\sigma_k,s_{k+1}}$ and there exists at least one value of s_{k+1} such that $Q_{\sigma_1,\sigma_2,\ldots,\sigma_k,s_{k+1}}$ is infinite. Let σ_{k+1} denote the least such value.

Finally let $\sigma = \sigma_1, \sigma_2, \ldots,$ be the sequence which this procedure generates. As each $\sigma_k \in P$, $\sigma \in P^\infty$. Let $\epsilon > 0$ and $r = \lceil 1/\epsilon \rceil$. As $B_{1/r}(\sigma) \cap Q = Q_{\sigma_1,\sigma_2,\ldots,\sigma_r}$, there are an infinite number of points $q \in Q$ such that $q \in B_{1/r}(\sigma) \subseteq B_\epsilon(\sigma)$. Therefore σ is a limit point of Q. This completes the proof.

Corollary 2.3.4 *(I_∞, d) and (G_∞, d) are both compact.*

Proof A closed subspace of a compact space is also compact. We show that both (I_∞, d) and (G_∞, d) are closed in (P^∞, d).

Let $s \in P^\infty \backslash I_\infty$, and then s must fail to satisfy at least one of the conditions I(a), I(b) and I(c) of Definition 2.3.1. If $s_1 > 1$ then, for all $s' \in B_1(s)$, $s'_1 = s_1 > 1$ and so $s' \in P^\infty \backslash I_\infty$. If, for some r, $|s_{r+1} - s_r| > 1$, then for all $s' \in B_{\frac{1}{r+1}}(s)$, $\left|s'_{r+1} - s'_r\right| > 1$, and again $s' \in P^\infty \backslash I_\infty$. Finally, if there exist $q < t$ such that $s_q = p$, but $s_t < p$, then for all $s' \in B_{\frac{1}{t}}(s)$, $s' \notin I_\infty$. Therefore $P^\infty \backslash I_\infty$ is open and so I_∞ is closed. A similar argument shows that G_∞ is also closed.

Thus, by the previous lemma, I_∞ and G_∞ are both closed subsets of a compact space. This completes the proof.

Finally, let d^* be the metric on the cartesian product $I_\infty \times G_\infty$ given as follows: If $i, i' \in I_\infty$ and $g, g' \in G_\infty$ then

$$d^*\left[(i,g),(i',g')\right] = \max\left\{d(i,i'), d(g,g')\right\}.$$

It is well known that this, the product metric, gives rise to the standard product topology on $I_\infty \times G_\infty$. The resulting topological space, denoted by $I \times G$ is called the topological product of I and G. Note particularly that if, for all $\epsilon > 0$ and $(i,g) \in I_\infty \times G_\infty$, $B^*_\epsilon[(i,g)]$ denotes the ϵ-ball around (i,g), then

$$B^*_\epsilon[(i,g)] = B_\epsilon(i) \times B_\epsilon(g).$$

The topological product is considered in greater detail in section 5.

2.4 Payoff Functions I

We define two payoff functions on the set $I_\infty \times G_\infty$. Let $i \in I_\infty$ and $g \in G_\infty$. Both functions are based on the number of terms at which the sequences i and g agree. They differ only in the ultimate objectives of the players. Whenever the first payoff is used the Infiltrator is just trying to avoid detection. Whenever the second is used he is trying to reach the target. If we assume that if the Infiltrator reaches the target he does not subsequently leave it, then the achievement of the second of these objectives also guarantees the achievement of the first. However the converse does not hold.

We assume that if the players simultaneously occupy the same state, then the Guard detects the Infiltrator with probability μ, $0 < \mu \leq 1$. As it is easier to work with the probability of a miss, let $\lambda = 1 - \mu$. Finally, let S be a set. If S is finite then $|S|$ denotes the number of elements in S. If S is infinite then $|S| = \infty$. Also, for all $0 \leq \lambda < 1$, by convention $\lambda^\infty = 0$.

Definition 2.4.1 *Let I_∞ and G_∞ be the strategy sets given in Definition 2.3.1.*

(a) For all $(i,g) \in I_\infty \times G_\infty$ the **coincidence number**, *$\omega(i,g)$, is given by*

$$\omega(i,g) = |\{ r \in \mathbb{N} \mid i_r = g_r \}|.$$

(b) The **detection function**, *$\Pi_\infty : I_\infty \times G_\infty \to \Re$ is given as follows: If $i \in I_\infty$ and $g \in G_\infty$, then*

$$\Pi_\infty(i,g) = \lambda^{\omega(i,g)}.$$

(c) The **target function**, *$\tilde{\Pi}_\infty : I_\infty \times G_\infty \to \Re$ is given as follows: If $i \in I_\infty$ and $g \in G_\infty$, then*

$$\tilde{\Pi}_\infty(i,g) = \begin{cases} \Pi_\infty(i,g) & \text{if, for some } r \in \mathbb{N},\ i_r = p, \\ 0 & \text{otherwise.} \end{cases}$$

Both Π_∞ and $\tilde{\Pi}_\infty$ are bounded. Moreover, for all $(i,g) \in I_\infty \times G_\infty$, $\tilde{\Pi}_\infty(i,g) \leq \Pi_\infty(i,g)$. We conclude this section by considering whether these functions are continuous. We prove that, although neither is continuous, they are both semicontinuous. First note that, for any function $f : S \to T$ and $T' \subseteq T$, $f^{-1}(T') = \{s \mid f(s) \in T'\}$.

Definition 2.4.2 (Semicontinuity) *Let T be a topological space. Let f be a real-valued function on T.*

(a) f is **upper semicontinuous** *on T if, for all $c \in \Re$, $f^{-1}((-\infty, c))$ is open.*

(b) f is **lower semicontinuous** *on T if, for all $c \in \Re$, $f^{-1}((c, +\infty))$ is open.*

If T is a metric space there are equivalent definitions of semicontinuity in terms of convergent sequences. Note that, for any sequence (s_r),

$$\underline{\lim}\, s_r = \lim_{t \to \infty} \left(\inf\{s_r \mid r \geq t\}\right).$$

(a') *f is upper semicontinuous on T, if for all convergent sequences $t_i \to t$ in T, $\underline{\lim}\, f(t_i) \leq f(t)$.*

(b') *f is lower semicontinuous on T, if for all convergent sequences $t_i \to t$ in T, $\underline{\lim}\, f(t_i) \geq f(t)$.*

A function is continuous if and only if it is both upper and lower semicontinuous. Recalling that $I \times G$ denotes the topological product of I and G, we consider the semicontinuity of the functions Π_∞ and $\tilde{\Pi}_\infty$ on $I \times G$.

Lemma 2.4.3 *Π_∞ is upper semicontinuous on $I \times G$.*

Proof Let $(i^r, g^r) \to (i,g)$ be a convergent sequence in $I_\infty \times G_\infty$. If $\omega(i,g) = \infty$ then, as $r \to \infty$, $\omega(i^r, g^r) \to \infty$ and so $\Pi_\infty(i^r, g^r) \to 0$. Therefore $\underline{\lim}\, \Pi_\infty(i^r, g^r) = \lim \Pi_\infty(i^r, g^r) = 0 = \Pi_\infty(i,g)$. If $\omega(i,g) < \infty$ then there exists t such that for all $s \geq t$, $i_s \neq g_s$. Thus, for all r sufficiently large that $d^*[(i^r, g^r), (i,g)] \leq 1/t$, $\omega(i^r, g^r) \geq \omega(i,g)$ and so $\Pi_\infty(i^r, g^r) \leq \Pi_\infty(i,g)$. Therefore $\underline{\lim}\, \Pi_\infty(i^r, g^r) \leq \Pi_\infty(i,g)$. This completes the proof.

Corollary 2.4.4 *There exist $i' \in I_\infty$ and $g' \in G_\infty$ such that $\Pi_\infty(i', g)$ is not lower semicontinuous on G_∞, and $\Pi_\infty(i, g')$ is not lower semicontinuous on I_∞.*

Proof Let $i' = 1, 1, \ldots$, and let $g^r \rightarrow g$ be the convergent sequence in G_∞ such that, for all r,

$$g^r = 1, \overbrace{2, 2, \ldots, 2}^{r \text{ terms}}, 1, 1, \ldots$$

and so $g = 1, 2, 2, \ldots$. Considering the terms at which these different elements of G_∞ meet i', we find that, for all r, $\omega(i', g^r) = \infty$, but that $\omega(i', g) = 1$. Therefore $\underline{\lim}\, \Pi_\infty(i', g^r) = \lim \Pi_\infty(i', g^r) = \lambda^\infty < \lambda^1 = \Pi_\infty(i', g)$. Hence, $\Pi_\infty(i', g)$ is not lower semicontinuous on G_∞

Now let $g' = 2, 2, \ldots$, and let $i^r \rightarrow i$ be the convergent sequence in I_∞ such that, for all r,

$$i^r = \overbrace{1, 1, \ldots, 1}^{r \text{ terms}}, 2, 3, \ldots, p, p, \ldots$$

and so $i = 1, 1, \ldots$. In this case we find that, for all r, $\omega(i^r, g') = 1$, but $\omega(i, g') = 0$. Therefore $\underline{\lim}\, \Pi_\infty(i^r, g') = \lim \Pi(i, g') = \lambda^1 < \lambda^0 = \Pi_\infty(i, g')$. Hence, $\Pi_\infty(i, g')$ is not lower semicontinuous on I_∞.

We have established this corollary for future reference. Either one of the examples in the proof above is sufficient to show that Π_∞ is not continuous.

Lemma 2.4.5 *$\tilde{\Pi}_\infty$ is lower semicontinuous on $I_\infty \times G_\infty$*

Proof Let $(i^r, g^r) \rightarrow (i, g)$ be a convergent sequence in $I_\infty \times G_\infty$. If, for all s, $i_s < p$, then, by definition of $\tilde{\Pi}_\infty$, $\tilde{\Pi}_\infty(i, g) = 0$. Therefore, it follows immediately that $\underline{\lim}\, \tilde{\Pi}_\infty(i^r, g^r) \geq \tilde{\Pi}_\infty(i, g)$. If there exists t such that $i_t = p$, then, by Definition 2.3.1, for all $s \geq t$, $i_s = p$. Hence, for all r sufficiently large that $d^*[(i^r, g^r), (i, g)] \leq 1/t$, $i^r = i$ and $\omega(i^r, g^r) = \omega(i, g) < \infty$. Therefore, $\underline{\lim}\, \tilde{\Pi}_\infty(i^r, g^r) = \lim \tilde{\Pi}_\infty(i^r, g^r) = \tilde{\Pi}_\infty(i, g)$. This completes the proof.

Corollary 2.4.6 *For all $i' \in I_\infty$, $\tilde{\Pi}_\infty(i', g)$ is upper semicontinuous on G_∞, but there exists $g' \in G_\infty$ such that $\tilde{\Pi}_\infty(i, g')$ is not upper semicontinuous on I_∞.*

Proof Fix $i' \in I_\infty$ and let $g^r \rightarrow g$ be a convergent sequence in G_∞. If, for all s, $i'_s < p$, then by definition of $\tilde{\Pi}_\infty$, for all r, $\tilde{\Pi}_\infty(i', g^r) = 0 = \tilde{\Pi}_\infty(i', g)$. Therefore $\underline{\lim}\, \tilde{\Pi}_\infty(i', g^r) = \tilde{\Pi}_\infty(i', g)$. If there exists t such that $i'_t = p$ then, for all $s \geq t$, $i'_s = p$. Hence, for all r sufficiently large that $d(g^r, g) \leq 1/t$, $\omega(i', g^r) = \omega(i', g)$. Therefore $\underline{\lim}\, \tilde{\Pi}_\infty(i', g^r) = \lim \tilde{\Pi}_\infty(i', g^r) = \tilde{\Pi}_\infty(i', g)$. Hence for all $i' \in I_\infty$, $\tilde{\Pi}_\infty(i', g)$ is continuous on G_∞ and so also upper semicontinuous.

To see that this does not hold the other way around, consider the first example contained in the proof of Corollary 2.4.5. We deduce that, for all r, $\tilde{\Pi}_\infty(i^r, g') = \lambda$ but that as for all r, $i_r < p$, $\tilde{\Pi}_\infty(i, g') = 0$. Therefore $\underline{\lim}\, \tilde{\Pi}_\infty(i^r, g') = \lim \tilde{\Pi}_\infty(i^r, g') = \lambda > 0 = \tilde{\Pi}_\infty(i, g')$. Hence, $\tilde{\Pi}_\infty(i, g')$ is not upper semicontinuous on I_∞.

Definition 2.4.7 *The* **detection game**, Γ_∞, *and the* **target game**, $\tilde{\Gamma}_\infty$, *are the game triples* $\Gamma_\infty = (I_\infty, G_\infty, \Pi_\infty)$ *and* $\tilde{\Gamma}_\infty = \left(I_\infty, G_\infty, \tilde{\Pi}_\infty\right)$.

The games Γ_∞ and $\tilde{\Gamma}_\infty$ do not, in general, have solutions. If $\lambda = 0$ then the Infiltrator can be immediately detected at state one, and so $v(\Gamma_\infty) = v(\tilde{\Gamma}_\infty) = 0$. Now suppose $\lambda > 0$. For $p = 3$ the strategies $i = 1, 2, 3, 3, \ldots$ and $g = 1, 2, 2, \ldots$ are optimal in $\tilde{\Gamma}_\infty$ and $v(\tilde{\Gamma}_\infty) = \lambda^2$. For $p = 4$ it can be seen, by considering the security levels, that while $\underline{v}(\Gamma_\infty) = \underline{v}(\tilde{\Gamma}_\infty) = \lambda^3$, $\overline{v}(\Gamma_\infty) = \lambda^2$ and $\overline{v}(\tilde{\Gamma}_\infty) = \lambda$. Thus neither game has a solution. This is clearly the case for all $0 < \lambda < 1$ and $p \geq 4$.

2.5 Strategies II

Let us pause for a moment to recap and outline what we are going to do in the next couple of sections. In section 3 we defined the strategy sets I_∞ and G_∞. The games Γ_∞ and $\tilde{\Gamma}_\infty$ were then given as the ordered triples $(I_\infty, G_\infty, \Pi_\infty)$ and $(I_\infty, G_\infty, \tilde{\Pi}_\infty)$ respectively, where Π_∞ and $\tilde{\Pi}_\infty$ were specified real-valued functions on the set $I_\infty \times G_\infty$. In this section we start the construction of two further games, Γ^*_∞ and $\tilde{\Gamma}^*_\infty$, to be known as the *extended games*. Intuitively, these are the games derived from Γ_∞ and $\tilde{\Gamma}_\infty$ when the players are permitted to *randomize* their choice of strategy. This section defines what it means to randomize, and denotes by $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ the sets of all such strategies for the Infiltrator and the Guard respectively. In section 6 the real-valued functions π_∞ and $\tilde{\pi}_\infty$ on $\mathcal{I}_\infty \times \mathcal{G}_\infty$ are defined as the expectations of Π_∞ and $\tilde{\Pi}_\infty$ respectively when the players use strategies from the sets $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$. Finally, the extended games are given by the ordered triples $\Gamma^*_\infty = (\mathcal{I}_\infty, \mathcal{G}_\infty, \pi_\infty)$ and $\tilde{\Gamma}^*_\infty = (\mathcal{I}_\infty, \mathcal{G}_\infty, \tilde{\pi}_\infty)$.

We now define what it means for the players to randomize. This involves us in considering probability measures on suitable σ-algebras on the sets I_∞ and G_∞. Let us first note some definitions.

Definition 2.5.1 *Let X be a set. A collection $\mathcal{A}$ of subsets of X is a* σ-**algebra on** X *if*

(a) $X \in \mathcal{A}$,

(b) $A \in \mathcal{A} \Rightarrow A^c \in \mathcal{A}$,

(c) $A_1, A_2, \ldots \in \mathcal{A} \Rightarrow \bigcup A_j \in \mathcal{A}$,

(d) $A_1, A_2, \ldots \in \mathcal{A} \Rightarrow \bigcap A_j \in \mathcal{A}$.

As, for any infinite sequence (A_j) of subsets of X, $\bigcap A_j = (\bigcup {A_j}^c)^c$, (b) and (c) together imply (d). Similarly (b) and (d) together imply (c). If X is a set and $\mathcal{C}$ is a collection of subsets of X, it can be shown that there is a σ-algebra $\sigma(\mathcal{C})$ which is the smallest σ-algebra containing $\mathcal{C}$. (The power set of X demonstrates that there exists at least one σ-algebra containing $\mathcal{C}$. The intersection of any collection of σ-algebras can be shown itself to be a σ-algebra . Hence let $\sigma(\mathcal{C})$ be the intersection of all σ-algebras containing $\mathcal{C}$.) $\sigma(\mathcal{C})$ is called **the σ-algebra generated by $\mathcal{C}$**. If X is a set and $\mathcal{A}$ is a σ-algebra on X then the pair $(X, \mathcal{A})$ is called **a measurable space**.

Definition 2.5.2 *Let $(X, \mathcal{A})$ be a measurable space.*

(i) A **measure on** $\mathcal{A}$ *is a function $\mu : \mathcal{A} \to [0, +\infty]$ that satisfies*

(a) $\mu(\emptyset) = 0$

(b) $\mu(\bigcup A_j) = \sum \mu(A_j)$ *for every sequence (A_j) of disjoint elements of $\mathcal{A}$.*

(ii) A **probability measure on** $\mathcal{A}$ *is a measure μ on $\mathcal{A}$ such that $\mu(X) = 1$.*

If $(X, \mathcal{A})$ is a measurable space and μ is a measure on $\mathcal{A}$ then the triple $(X, \mathcal{A}, \mu)$ is called **a measure space**. When μ is also a probability measure on $\mathcal{A}$, $(X, \mathcal{A}, \mu)$ is known as a **probability space**.

In a moment we shall apply this to the sets I_∞ and G_∞. First consider another general idea which will be required later.

Definition 2.5.3 *Let $(X_1, \mathcal{A}_1)$ and $(X_2, \mathcal{A}_2)$ be measurable spaces. The σ-algebra $\mathcal{A}_1 \otimes \mathcal{A}_2$ on $X_1 \times X_2$, known as* **the product of the σ-algebras $\mathcal{A}_1$ and $\mathcal{A}_2$**, *is given by*

$$\mathcal{A}_1 \otimes \mathcal{A}_2 = \sigma(\mathcal{A}_1 \circ \mathcal{A}_2),$$

where $\mathcal{A}_1 \circ \mathcal{A}_2 = \{A_1 \times A_2 \mid A_i \in \mathcal{A}_i,\ i = 1, 2\}$.

It is well known (see Cohn [21]) that, if $(X_1, \mathcal{A}_1, \mu_1)$ and $(X_2, \mathcal{A}_2, \mu_2)$ are probability spaces, then there is a unique measure $\mu_1 \odot \mu_2$ on $\mathcal{A}_1 \otimes \mathcal{A}_2$ such that, for all $A_i \in \mathcal{A}_i$, $i = 1, 2$

$$(\mu_1 \odot \mu_2)(A_1 \times A_2) = \mu_1(A_1)\mu_2(A_2) \tag{2.1}$$

The measure $\mu_1 \odot \mu_2$ is called the **product of μ_1 and μ_2.**

For our purposes this is all that we require to know about $\mu_1 \odot \mu_2$. The full theory proves the existence of the product measure not only for probability spaces but for all measure spaces $(X_1, \mathcal{A}_1, \mu_1)$ and $(X_2, \mathcal{A}_2, \mu_2)$ which are σ-finite (Note: $(X, \mathcal{A}, \mu)$ is σ-finite if X is the union of a sequence $\{A_j\}$ of elements of $\mathcal{A}$ such that, for all j, $\mu(A_j) < +\infty$). Moreover, it shows that the measure, under $\mu_1 \odot \mu_2$, of an arbitrary element E of $\mathcal{A}_1 \otimes \mathcal{A}_2$ is given by

$$(\mu_1 \odot \mu_2)(E) = \int_{X_1} \mu_2(E_{x_1})\mu_1(dx_1) = \int_{X_2} \mu_1(E^{x_2})\mu_2(dx_2),$$

where, for all $(x_1, x_2) \in X_1 \times X_2$, $E_{x_1} = \{x_2 \in X_2 | (x_1, x_2) \in E\}$ and $E^{x_2} = \{x_1 \in X_1 | (x_1, x_2) \in E\}$. Further details about product measures can be found in any work on Measure Theory. See, for example, Cohn [21], Billingsley [13], who concentrates on probability measures, or the classic treatment by Halmos [36]. The theory which we have developed here is not yet able to handle the integrals given above. In the next section, given a measure space $(X, \mathcal{A}, \mu)$ and a $[0, +\infty]$-valued $\mathcal{A}$-measurable function f on X, we construct the integral $\int f d\mu$. For the moment it is sufficient for us that the measure $\mu_1 \odot \mu_2$ is uniquely defined and satisfies equation (2.1).

We finally return to the strategy sets I_∞ and G_∞. Recall that back in section 3 we observed that the metric d on P^∞ induces the topologies $\mathcal{U}$ and $\mathcal{V}$ on I_∞ and G_∞ respectively. The σ-algebras that we use are those generated by the open subsets of the topolgical spaces $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$. These are the σ-algebras $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$. If $T = (X, \mathcal{U})$ is any Hausdorff topological space, the σ-algebra $\sigma(\mathcal{U})$ is sometimes called the **Borel σ-algebra of T.** In this case the elements of $\sigma(\mathcal{U})$ are known as the **Borel subsets of T.**

This section ends with the definitions of the mixed strategy sets $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$. Before this, however, consider the topological product of $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$. By definition, a subset W of $I_\infty \times G_\infty$ is open under the product topology if, for all $w \in W$, there exist $U \in \mathcal{U}$ and $V \in \mathcal{V}$ such that $w \in U \times V \subseteq W$. If we denote the collection of all such open sets by $\mathcal{W}$, then the topological product of $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$ is denoted by $(I_\infty \times G_\infty, \mathcal{W})$. Thus, in addition to the

Borel σ-algebras $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$ mentioned above, we can also consider the Borel σ-algebra $\sigma(\mathcal{W})$ of $(I_\infty \times G_\infty, \mathcal{W})$.

Recalling Definition 2.5.3, we also deduce the existence of the σ-algebra $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$. Therefore $\sigma(\mathcal{W})$ and $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$ are both σ-algebras on $(I_\infty \times G_\infty, \mathcal{W})$.

Lemma 2.5.4 *If $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{W}$ denote the open subsets of the sets I_∞, G_∞ and $I_\infty \times G_\infty$ respectively, then*

$$\sigma(\mathcal{W}) \subseteq \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V}).$$

Proof The σ-algebra $\sigma(\mathcal{W})$ on $(I_\infty \times G_\infty, \mathcal{W})$ is generated by the open subsets of $(I_\infty \times G_\infty, \mathcal{W})$. By Definition 2.5.1(b), it is also generated by the closed subsets. The Lemma is thus proved if we show that for every closed subset K in $(I_\infty \times G_\infty, \mathcal{W})$, $K \in \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$.

We first deduce that as $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$ are compact spaces then, by Tychonoff's Theorem, their topological product is also compact. Moreover, as every closed subspace of a compact space is compact, K is compact too. We now exploit the fact that every open covering of a compact space has a finite subcovering.

For all $n \in \mathbb{N}$ the collection of sets $\{B^*_{1/n}[k] \mid k \in K\}$ is clearly an open covering of K. Therefore, for all n, there exists a finite subset of K, $\{k_j^{(n)} \mid j = 1, 2, \ldots, m_n\}$, such that the union

$$K_n = \bigcup_{j=1}^{j=m_n} B^*_{1/n}\left[k_j^{(n)}\right]$$

contains K.

Recall that, for all $\epsilon > 0$ and for all $(i, g) \in I_\infty \times G_\infty$, $B^*_\epsilon\left[(i, g)\right] = B_\epsilon(i) \times B_\epsilon(g)$, and therefore $B^*_\epsilon\left[(i, g)\right] \in \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$. Hence, as for all n, K_n is the union of a finite collection of such sets, $K_n \in \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$. For all $n \in \mathbb{N}$ neither $\{k_j^{(n)} \mid j = 1, 2, \ldots, m_n\}$ nor K_n is necessarily unique. It remains to show, for all possible choices of $K_1, K_2, \ldots$, that $K = \cap K_n$. For then, by Definition 2.5.1(d), $K \in \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$ as required.

By definition, for all $n \in \mathbb{N}$, $K \subseteq K_n$. Hence $K \subseteq \cap K_n$. To prove the reverse inclusion, suppose that there exist $K_1, K_2, \ldots$ and $x \in \cap K_n$ such that $x \notin K$. As for all $n \in \mathbb{N}$, $x \in K_n$, it follows that for all $n \in \mathbb{N}$ there exists $k_{j(x)}^{(n)} \in K$ such that $x \in B^*_{1/n}\left(k_{j(x)}^{(n)}\right)$. Equivalently, for all $n \in \mathbb{N}$, there exists $k_{j(x)}^{(n)} \in K$

such that $k_{j(x)}^{(n)} \in B_{1/n}^{*}(x)$. Therefore, x is a limit point of K. However, K is closed and contains all of its limit points. Thus $x \notin K$ leads to a contradiction. Hence, for all $x \in \cap K_n$, $x \in K$ and so $\cap K_n \subseteq K$. This then confirms that $K = \cap K_n$ and thus completes the proof.

Definition 2.5.5 *Let $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$ be the Borel σ-algebras of the topological spaces $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$ respectively. We define the* **extended strategy sets,** *$\mathcal{I}_\infty$ and $\mathcal{G}_\infty$, for the Infiltrator and the Guard respectively by*

(a) $\mathcal{I}_\infty = \{\iota \mid (I_\infty, \sigma(\mathcal{U}), \iota)$ *is a probability space* $\}$.

(b) $\mathcal{G}_\infty = \{\gamma \mid (G_\infty, \sigma(\mathcal{V}), \gamma)$ *is a probability space* $\}$.

We conclude by observing one particularly important property of the Borel σ-algebras $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$.

Lemma 2.5.6 *The σ-algebras $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$ contain all the singleton subsets of I_∞ and G_∞ respectively.*

Proof The proof is identical for both $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$. Let us do it for $\sigma(\mathcal{U})$.

Let $i \in I_\infty$. For all $\epsilon > 0$ the ϵ-ball $B_\epsilon(i)$ around i is open, and therefore an element of $\sigma(\mathcal{U})$. Thus, from Definition 2.5.1 and as $\{i\} = \bigcap_{n\in\mathbb{N}} B_{1/n}(i)$, it follows that the singleton subset $\{i\}$ is also an element of $\sigma(\mathcal{U})$.

We now see the intuition behind calling the sets $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ the *extensions* of I_∞ and G_∞. For suppose $i \in I_\infty$ is any pure strategy for the Infiltrator. From the preceeding lemma, $\{i\} \in \sigma(\mathcal{U})$ and hence, we can denote by $\mu_i \in \mathcal{I}_\infty$ the unique probability measure on $\sigma(\mathcal{U})$ which satisfies $\mu_i(\{i\}) = 1$. Intuitively, μ_i is the extended strategy that corresponds to the strategy i. In the same way, for any $g \in G_\infty$ we can define $\mu_g \in \mathcal{G}_\infty$ such that $\mu_g(\{g\}) = 1$. Hence we think of $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ as containing all the elements of I_∞ and G_∞ and a whole lot more.

2.6 Payoff Functions II

Our aim in this section is to define the extended payoff functions π_∞ and $\tilde{\pi}_\infty$. This requires a brief consideration of measurable functions in general and the measurability of the functions Π_∞ and $\tilde{\Pi}_\infty$ in particular.

Definition 2.6.1 *Let $(X, \mathcal{A})$ be a measurable space. A function $f : X \to \Re$ is* **$\mathcal{A}$-measurable** *if, for every $c \in \Re$, $f^{-1}((-\infty, c]) \in \mathcal{A}$.*

Lemma 2.6.2 *Π_∞ and $\tilde{\Pi}_\infty$ are both $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$-measurable on $I_\infty \times G_\infty$.*

Proof This follows from the semicontinuity established by Lemmas 2.4.3 and 2.4.5. Let $c \in \Re$. First, as $\tilde{\Pi}_\infty$ is lower semicontinuous on $I_\infty \times G_\infty$, then by definition, $\tilde{\Pi}_\infty^{-1}((c, +\infty))$ is open. This set is the complement of $\tilde{\Pi}_\infty^{-1}((-\infty, c])$. As $\sigma(\mathcal{W})$ is generated by the open sets, $\tilde{\Pi}_\infty^{-1}((-\infty, c]) \in \sigma(\mathcal{W})$.

Similarly, Π_∞ is upper semicontinuous on $I_\infty \times G_\infty$. Therefore for all $d \in \Re$, $\Pi_\infty^{-1}((-\infty, d)) \in \mathcal{W} \subseteq \sigma(\mathcal{W})$. As

$$\Pi_\infty^{-1}\{(-\infty, c]\} = \bigcap_n \Pi_\infty^{-1}\{(-\infty, c + 1/n)\}$$

it follows, from Definition 2.5.1(d) that $\Pi_\infty^{-1}((-\infty, c]) \in \sigma(\mathcal{W})$.

Finally by Lemma 2.5.4, $\sigma(\mathcal{W}) \subseteq \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$. This completes the proof.

We have already seen that, for all $\iota \in \mathcal{I}_\infty$ and all $\gamma \in \mathcal{G}_\infty$, there exists a unique probability measure $\iota \odot \gamma$, defined on $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$, which satisfies $(\iota \odot \gamma)(A \times B) = \iota(A)\gamma(B)$ for all $A \in \sigma(\mathcal{U})$ and all $B \in \sigma(\mathcal{V})$. Now that we have shown that both Π_∞ and $\tilde{\Pi}_\infty$ are $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$-measurable let us find an expression for the integrals $\int \Pi_\infty \, d(\iota \odot \gamma)$ and $\int \tilde{\Pi}_\infty \, d(\iota \odot \gamma)$.

To avoid needless repitition in considering both Π_∞ and $\tilde{\Pi}_\infty$, let $\Pi : I_\infty \times G_\infty \to \Re$ be any $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$-measurable function which takes values only from the set $\{\lambda^j \mid j = 0, 1, \ldots, \infty\}$. Recalling that, by convention, $\lambda^\infty = 0$, clearly both Π_∞ and $\tilde{\Pi}_\infty$ satisfy these conditions.

Let us define, for each $j = 0, 1, \ldots, \infty$, the subset A_j of $I_\infty \times G_\infty$ given by

$$A_j = \Pi^{-1}\left(\{\lambda^j\}\right).$$

As Π has been assumed to take only the countable set of values given above, we deduce that

$$A_j^c = \Pi^{-1}\left((-\infty, \lambda^{j+1}]\right) \cup \left[\Pi^{-1}\left((-\infty, \lambda^j]\right)\right]^c.$$

It then follows that, as Π is $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$-measurable, then $A_j \in \sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$.

Therefore, for $n = 0, 1, \ldots, \infty$ we may define the non-negative simple function $\Pi^{[n]} : I_\infty \times G_\infty \to \{0, \lambda^j \mid j = 0, 1, \ldots, n\}$ as follows. For all $(i, g) \in I_\infty \times G_\infty$,

$$\Pi^{[n]}(i, g) = \sum_{j=0}^{n} \lambda^j \, \chi_{A_j}(i, g),$$

where, for all j, χ_{A_j} is the characteristic function of the set A_j. Note that for all $(i,g) \in I_\infty \times G_\infty$ $\Pi^{[0]}(i,g) \leq \Pi^{[1]}(i,g) \leq \ldots$ and that $\lim \Pi^{[n]}(i,g) = \Pi(i,g)$. For all n, the integral $\int \Pi^{[n]}\, d(\iota \odot \gamma)$ of $\Pi^{[n]}$ with respect to $\iota \odot \gamma$ is defined to be

$$\int_{I_\infty \times G_\infty} \Pi^{[n]}\, d(\iota \odot \gamma) = \sum_{j=0}^{n} \lambda^j\ (\iota \odot \gamma)(A_j).$$

If $(X, \mathcal{A}, \mu)$ is a measure space and $f : X \to [0, +\infty]$ is a $\mathcal{A}$-measurable function, then the integral $\int f\ d\mu$ is defined in terms of integrals of simple functions such as those considered above. If we let $\mathcal{S}_+$ denote the set of all non-negative, $\mathcal{A}$-measurable, simple functions defined on X, then we define

$$\int_X f\ d\mu = \sup \left\{ \int_X g\ d\mu \mid g \in \mathcal{S}_+ \text{ and } g \leq f \right\}.$$

The monotone convergence theorem (see Cohn [21] section 2.4.1) implies that if there exists a sequence $f_0, f_1, \ldots$ of $[0, +\infty]$-valued $\mathcal{A}$-measurable simple functions on X which satisfy

$$f_0(x) \leq f_1(x) \leq \ldots \tag{2.2}$$

and

$$f(x) = \lim_{n \to \infty} f_n(x) \tag{2.3}$$

for all $x \in X$, then $\int f\ d\mu = \lim \int f_n\ d\mu$.

[In fact, the monotone convergence theorem requires only that the two displayed equations above hold *at almost every* x in X. In other words, if we consider $\bar{X} = \{x \in X |\ (2.2) \text{ or } (2.3) \text{ is false }\}$ then there exists $A \in \mathcal{A}$ such that $\bar{X} \subseteq A$ and $\mu(A) = 0$.]

Thus returning to the function $\Pi : I_\infty \times G_\infty \to [0, +\infty]$, we know that Π is, by assumption, $\sigma(\mathcal{U}) \otimes \sigma(\mathcal{V})$-measurable and that the simple functions $\Pi^{[0]}, \Pi^{[1]}, \ldots$ satisfy both (2.2) and (2.3). As, in addition,

$$\lim_{n \to \infty} \int \Pi^{[n]}\ d(\iota \odot \gamma) = \sum_{j=0}^{\infty} \lambda^j\ (\iota \odot \gamma)(A_j),$$

we deduce that

$$\int_{I_\infty \times G_\infty} \Pi\, d(\iota \odot \gamma) = \sum_{j=0}^{\infty} \lambda^j\, (\iota \odot \gamma)(A_j),$$

where, for $j = 0, 1, \ldots$, $A_j = \Pi^{-1}(\lambda^j)$.

Definition 2.6.3 *Let $\Pi_\infty : I_\infty \times G_\infty \to \Re$ and $\tilde{\Pi}_\infty : I_\infty \times G_\infty \to \Re$ be the functions given in Definition 2.4.1. Let $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ be the extended strategy sets, given by Definition 2.5.5.*

(a) The **extended detection function** $\pi_\infty : \mathcal{I}_\infty \times \mathcal{G}_\infty \to \Re$ *is given by*

$$\pi_\infty(\iota, \gamma) = \int_{I_\infty \times G_\infty} \Pi_\infty \, d(\iota \odot \gamma) \tag{2.4}$$

for all $\iota \in \mathcal{I}_\infty$ and all $\gamma \in \mathcal{G}_\infty$.

(b) The **extended target function** $\tilde{\pi}_\infty : \mathcal{I}_\infty \times \mathcal{G}_\infty \to \Re$ *is given by*

$$\tilde{\pi}_\infty(\iota, \gamma) = \int_{I_\infty \times G_\infty} \tilde{\Pi}_\infty \, d(\iota \odot \gamma) \tag{2.5}$$

for all $\iota \in \mathcal{I}_\infty$ and all $\gamma \in \mathcal{G}_\infty$.

Fubini's Theorem allows us to rewrite the integrals above as repeated integrals. Thus we can replace (2.4) by $\pi_\infty(\iota, \gamma) = \int\int \Pi_\infty \, d\iota \, d\gamma$ and (2.5) by $\tilde{\pi}_\infty(\iota, \gamma) = \int\int \tilde{\Pi}_\infty \, d\iota \, d\gamma$.

Definition 2.6.4 *The* **extended detection game**, Γ^*_∞, *and the* **extended target game**, $\tilde{\Gamma}^*_\infty$, *are the game triples $\Gamma^*_\infty = (\mathcal{I}_\infty, \mathcal{G}_\infty, \pi_\infty)$ and $\tilde{\Gamma}^*_\infty = (\mathcal{I}_\infty, \mathcal{G}_\infty, \tilde{\pi}_\infty)$.*

2.7 The Existence of Solutions

We observed in the first half of the chapter that the basic games Γ_∞ and $\tilde{\Gamma}_\infty$ do not necessarily have solutions. The motivation behind the second half of the chapter has been the hope that the extended games Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ do have solutions. We see below that our hopes are fulfilled. The existence of solutions to Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ follows from one of the minimax theorems mentioned earlier. However, before this theorem can be stated, there is one more property to define.

Definition 2.7.1 *Let $T = (X, \mathcal{U})$ be a Hausdorff topological space, and $\sigma(\mathcal{U})$ the Borel σ-algebra on T. A measure μ on $\sigma(\mathcal{U})$ is* **regular** *if*

(a) *for every compact subspace K of T, $\mu(K) < \infty$,*

(b) *for every $S \in \sigma(\mathcal{U})$, $\mu(S) = \inf \{\mu(U) \mid S \subseteq U$ and U open$\}$,*

(c) *for every $U \in \mathcal{U}$ $\mu(U) = \sup \{\mu(K) \mid K \subseteq U$ and K compact$\}$.*

Theorem 2.7.2 (Glicksberg, 1950) *Let X and Y be compact metric spaces, and $\mathcal{B}(X)$ and $\mathcal{B}(Y)$ the sets of all regular probability measures on the Borel algebras of X and Y respectively. Let $f : X \times Y \to \Re$.*

(i) If f is upper semicontinuous then

$$\max_{\mu\in\mathcal{B}(X)} \inf_{\nu\in\mathcal{B}(Y)} \iint f(x,y)\, d\mu(x)\, d\nu(y) = \inf_{\nu\in\mathcal{B}(Y)} \sup_{\mu\in\mathcal{B}(X)} \iint f(x,y)\, d\mu(x)\, d\nu(y).$$

(ii) If f is lower semicontinuous then

$$\min_{\mu\in\mathcal{B}(X)} \sup_{\nu\in\mathcal{B}(Y)} \iint f(x,y)\, d\mu(x)\, d\nu(y) = \sup_{\nu\in\mathcal{B}(Y)} \inf_{\mu\in\mathcal{B}(X)} \iint f(x,y)\, d\mu(x)\, d\nu(y).$$

The two parts (i) and (ii) are different versions of the same result. One is obtained from the other by considering the function $f' = -f$. For the proof see Glicksberg [34].

Provided that we can show that the extended strategy sets contain only regular measures, we can apply Glicksberg's Theorem to the games Γ^*_∞ and $\tilde{\Gamma}^*_\infty$. In fact, it is known that when a topological space is also a metric space all probability measures on the Borel σ-algebra are regular. This result can be found in Parthasarathy [58] (Chapter II, section 1).

Theorem 2.7.3 *The extended games Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ both have solutions. Moreover there exists $\iota^* \in \mathcal{I}_\infty$ and $\gamma^* \in \mathcal{G}_\infty$ such that ι^* is optimal in Γ^*_∞ and γ^* is optimal in $\tilde{\Gamma}^*_\infty$.*

Proof Consider Γ^*_∞ first. By Lemma 2.3.4 both $(I_\infty, \mathcal{U})$ and $(G_\infty, \mathcal{V})$ are compact metric spaces. By Lemma 2.4.3, Π_∞ is upper semicontinuous on $I_\infty \times G_\infty$. Finally $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ are precisely the sets of regular probability measures on the Borel σ-algebras $\sigma(\mathcal{U})$ and $\sigma(\mathcal{V})$ respectively. Therefore, by version (i) of Glicksberg's Theorem, $\max_{\iota\in\mathcal{I}_\infty} \inf_{\gamma\in\mathcal{G}_\infty} \pi_\infty(\iota,\gamma) = \inf_{\gamma\in\mathcal{G}_\infty} \sup_{\iota\in\mathcal{I}_\infty} \pi_\infty(\iota,\gamma)$ and so Γ^*_∞ has a solution. Moreover if $\iota^* \in \mathcal{I}_\infty$ is any extended strategy which satisfies $\inf_{\gamma\in\mathcal{G}_\infty} \pi_\infty(\iota^*,\gamma) = \max_{\iota\in\mathcal{I}_\infty} \inf_{\gamma\in\mathcal{G}_\infty} \pi_\infty(\iota,\gamma)$, then ι^* is an optimal strategy for the Infiltrator.

The solution of $\tilde{\Gamma}^*_\infty$ and the existence of an optimal Guard strategy $\gamma^* \in \mathcal{G}_\infty$ is almost identical. In this case we use the lower semicontinuity of $\tilde{\pi}_\infty$ (Lemma 2.4.5) and version (ii) of Glicksberg's Theorem.

Let us denote the value of these games by v_∞ and $\tilde{v}_\infty$ respectively. As for all $i \in I_\infty$ and $g \in G_\infty$, $\tilde{\Pi}_\infty(i,g) \leq \Pi_\infty(i,g)$, we immediately deduce that $\tilde{v}_\infty \leq v_\infty$.

2.8 Alternative Approaches

There are some questions that naturally arise from what we have done. Possibly the most immediate is this: Are not the games Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ really the same? As the game goes on is not the Infiltrator bound, eventually, to go to the target, whether he has to or not? This question is put off until the next chapter. It is really outside the scope of the current chapter. Our purpose here has been to find whatever results we could concerning the *existence* of solutions and optimal strategies. In this context there are still a couple of questions that remain.

One question concerns the extent of Theorem 2.7.3. We have proved that both Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ have solutions and that in each game one of the players has an optimal strategy. Can we go on and prove that both players have optimal strategies?

A second question relates to the choices we have made in the way we have defined the games. By the strategy sets and the payoff functions we have used we have affected the mathematical structure of the games. There were other possibilities available to us. We decided to reject them. How do these decisions affect the existence of solutions and the values of the games when solutions exist?

We tackle these questions together. As a first step let us question the particular minimax theorem we have used. In the literature there are many minimax theorems. There are results which have supplemented and generalised Glicksberg's Theorem of 1950. Could we do better with another theorem? One of the more recent results is that of Alpern and Gal [2]. They used the earlier theorem of Kneser [42] to prove the following result.

Theorem 2.8.1 (Alpern/Gal, 1988) *Let $(X, \mathcal{A})$ be a measurable space and let Y be a compact Hausdorff space. Let $f : X \times Y \to \Re$ be a measurable function which is bounded below and lower semicontinuous on Y for fixed $x \in X$. Let M be any convex set of probability measures on $(X, \mathcal{A})$. Then*

$$\min_{\mu\in\mathcal{B}(Y)} \sup_{\nu\in M} \iint f(x,y)\; d\mu(x)\, d\nu(y) = \sup_{\nu\in M} \min_{\mu\in\mathcal{B}(Y)} \iint f(x,y)\; d\mu(x)\, d\nu(y),$$

where $\mathcal{B}(Y)$ is the set of regular Borel probability measures on Y.

Letting X be a topological space with $\mathcal{A}$ its Borel algebra, and letting M be equal to the set of regular Borel probability measures on $(X, \mathcal{A})$, they obtain the following generalisation of Glicksberg's Theorem. As before we give two versions of the result.

Corollary 2.8.2 (Alpern/Gal, 1988) *Let X be a topological space and Y a compact Hausdorff space, and let $\mathcal{B}(X)$ and $\mathcal{B}(Y)$ denote the sets of regular probability measures on the Borel σ-algebras on X and Y respectively. Let $f : X \times Y \to \Re$ be a measurable function.*

(a) If f is bounded below and lower semicontinuous on Y for fixed $x \in X$, then

$$\min_{\mu \in \mathcal{B}(Y)} \sup_{\nu \in \mathcal{B}(X)} \iint f(x,y)\, d\mu(x)\, d\nu(y) = \sup_{\nu \in \mathcal{B}(X)} \min_{\mu \in \mathcal{B}(Y)} \iint f(x,y)\, d\mu(x)\, d\nu(y).$$

(b) If f is bounded above and upper semicontinuous on Y for fixed $x \in X$, then

$$\max_{\mu \in \mathcal{B}(Y)} \inf_{\nu \in \mathcal{B}(X)} \iint f(x,y)\, d\mu(x)\, d\nu(y) = \inf_{\nu \in \mathcal{B}(X)} \max_{\mu \in \mathcal{B}(Y)} \iint f(x,y)\, d\mu(x)\, d\nu(y).$$

The reader is referred to the original [2] for proof of both the theorem and the corollary. The corollary clearly follows immediately from the theorem once it has been established that $\mathcal{B}(X)$ is a convex set.

As Corollary 2.8.2 is a generalisation of Glicksberg's Theorem we could use it to prove Theorem 2.7.3. However it gives us nothing more. If, for fixed $i \in I_\infty$, Π_∞ was lower semicontinuous on G_∞, we could use Corollary 2.8.2(a) and deduce the existence of an optimal Guard strategy for Γ^*_∞. Similarly, if for fixed $g \in G_\infty$, $\tilde{\Pi}_\infty$ was upper semicontinuous on I_∞ we could use Corollary 2.8.2(b) to deduce the existence of an optimal Infiltrator strategy for $\tilde{\Gamma}^*_\infty$. Unfortunately we have shown (see Corollaries 2.4.4 and 2.4.6) that neither of these properties hold.

These optimal strategies may still exist. It is just that neither Glicksberg's Theorem nor the Alpern/Gal result guarantees it. For the moment we leave this question and look at the second problem mentioned. We have considered an alternative minimax theorem. Let us now consider an alternative game.

Consider first our original target game $\tilde{\Gamma}_\infty = \left(I_\infty, G_\infty, \tilde{\Pi}_\infty\right)$. By definition of $\tilde{\Pi}_\infty$ it is no use for the Infiltrator to permanently stay away from the target. If he does so the payoff is bound to be zero. Therefore let

$$I'_\infty = \{i \in I_\infty |\text{ for some } t \in \mathbb{N},\ i_t = p\}.$$

If $i \in I'_\infty$ then there exists t such that, $i_s = p$ if and only if $s \geq t$. Note that as $I' \subseteq P^\infty$, (I'_∞, d) is a metric space. Now consider the new game $\left(I'_\infty, G_\infty, \tilde{\Pi}_\infty\right)$.

Lemma 2.8.3 $\tilde{\Pi}_\infty$ *is continuous on* $I'_\infty \times G_\infty$.

Proof Let $(i^r, g^r) \to (i, g)$ be a convergent sequence in $I'_\infty \times G_\infty$. As $i \in I'_\infty$ there exists t such that, for all $s \geq t$, $i_s = p$. Thus for all r sufficiently large that $d^*[(i^r, g^r), (i, g)] \leq \frac{1}{t}$, $i^r = i$ and $\omega(i^r, g^r) = \omega(i, g)$. Therefore $\lim \tilde{\Pi}_\infty(i^r, g^r) = \tilde{\Pi}_\infty(i, g)$.

As $\tilde{\Pi}_\infty$ is continuous on $I'_\infty \times G_\infty$ it can also be shown to be measurable. Therefore we can consider the resulting extended game $(\mathcal{I}'_\infty, \mathcal{G}_\infty, \tilde{\pi}_\infty)$ where $\mathcal{I}'_\infty$ denotes the Borel probability measures on I'_∞. Our only difficulty is with the metric space (I'_∞, d).

Lemma 2.8.4 *(I'_∞, d) is not compact.*

Proof Recall that a metric space is compact if and only if every infinite subset has a limit point. We shall give a subset of I'_∞ which has no limit point in I'_∞. Let (i^r), $r \geq 1$ be the sequence in I'_∞ such that, for $r \geq 1$,

$$i^r = \overbrace{1, 1, \ldots, 1}^{r \text{ terms}}, 2, \ldots, p, p, \ldots$$

In the larger set I_∞, this is a convergent sequence and the infinite set $\{i^r | r \geq 1\}$ has a unique limit point $i = 1, 1, \ldots$. As $i \notin I'_\infty$, $\{i^r | r \geq 1\}$ has no limit point in I'_∞. Therefore I'_∞ is not compact.

This prevents us from using Glicksberg's Theorem on $(\mathcal{I}'_\infty, \mathcal{G}_\infty, \tilde{\pi}_\infty)$. However it does illustrate the strength of the Alpern/Gal result. From Corollary 2.8.2(a) we deduce that $(\mathcal{I}'_\infty, \mathcal{G}_\infty, \tilde{\pi}_\infty)$ has a solution and that there is an optimal strategy for the Guard. The value of this game is clearly the same as $\tilde{v}_\infty$ and an optimal strategy or ϵ-optimal strategy here is also optimal and ϵ-optimal in $\tilde{\Gamma}^*_\infty$. This confirms our earlier conclusions concerning $\tilde{\Gamma}^*_\infty$. Unfortunately, as I'_∞ is not compact, it still does not guarantee the existence of an optimal strategy for the Infiltrator.

There are other games that we can define which have the same solutions as Γ^*_∞ or $\tilde{\Gamma}^*_\infty$. There are alternatives for both the strategy sets and the payoff functions. However it is not clear that any of these give us more details concerning the solutions of Γ^*_∞ or $\tilde{\Gamma}^*_\infty$.

In summary, we have partially answered the questions posed at the start of this section. We are unable, even using a recent more general minimax theorem, to establish the existence of optimal strategies for both players. Neither do alternative strategy sets or payoff functions seem to help in this. An example in a later chapter will partly justify these conclusions.

Chapter 3

FINITE GAMES

3.1 Introduction

Let us change our approach. We are still interested in Γ^*_∞ and $\tilde{\Gamma}^*_\infty$ and their solutions. The previous chapter has shown that there are solutions to these games, but not much else. We know that the values v_∞ and $\tilde{v}_\infty$ exist, but what are they?

It is this kind of question that leads us into the current chapter. Here we study 'finite games' — that is, games on which we impose a finite time limit. By deducing the properties of these games, especially as the time limit becomes larger and larger, insights emerge into games without time limits. We arrange for these to be none other than the infinite games Γ^*_∞ and $\tilde{\Gamma}^*_\infty$.

3.2 Games with Time Limits

We define two sequences of games: the finite detection games $\Gamma_1, \Gamma_2, \ldots$, and the finite target games $\tilde{\Gamma}_1, \tilde{\Gamma}_2, \ldots$. For $n \in \mathbb{N}$, the games Γ_n and $\tilde{\Gamma}_n$ are equivalent to the infinite games Γ_∞ and $\tilde{\Gamma}_\infty$ except that there is also a time limit n.

Definition 3.2.1 *Let $n \in \mathbb{N}$.*

(i) The finite strategy sets, I_n and G_n, for the Infiltrator and the Guard respectively, are the subsets of P^n given as follows. If $i \in I_n$ then

i(a) $i_1 = 1$,

i(b) $|i_{r+1} - i_r| \leq 1$ *for all* $1 \leq r < n$, *and* i(c) *if, for some* $t \geq 1$, $i_t = p$ *then for all* $r > t$, $i_r = p$.

If $g \in G_n$ *then*

g(a) $g_r < p$ *for all* $r \geq 1$,

g(b) $|g_{r+1} - g_r| \leq u$, *for all* $1 \leq r < n$.

(ii) The payoff functions, $\Pi_n, \tilde{\Pi}_n : I_n \times G_n \to \Re$, *are given as follows. For all* $(i, g) \in I_n \times G_n$,

$$\Pi_n(i, g) = \lambda^{\omega_n(i,g)},$$

where $\omega_n(i, g) = |\{r \leq n | i_r = g_r\}|$, *and*

$$\tilde{\Pi}_n(i, g) = \begin{cases} \Pi_n(i, g) & \text{if } i_r = p \text{ for some } r \leq n, \\ 0 & \text{otherwise.} \end{cases}$$

(iii) The **finite detection game** *and the* **finite target game** *are the game triples* $\Gamma_n = (I_n, G_n, \Pi_n)$ *and* $\tilde{\Gamma}_n = \left(I_n, G_n, \tilde{\Pi}_n\right)$ *respectively.*

Before continuing, let us make a few observations which follow on immediately from the definitions:

(i) For all $n \in \mathbb{N}$ and all $(i, g) \in I_n \times G_n$, $\tilde{\Pi}_n(i, g) \leq \Pi_n(i, g)$.

(ii) For all $n \in \mathbb{N}$ and $i \in I_n$, then for $1 \leq r \leq n$, $i_r \leq r$. Hence, for all r, $i_r \leq n$. This implies (iii).

(iii) If $n < p$, then for all $(i, g) \in I_n \times G_n$, $\tilde{\Pi}_n(i, g) = 0$.

(iv) For all $n \geq p$, any strategy $i \in I_n$ such that $i_r < p$ for all $r \leq n$ is dominated in the game $\tilde{\Gamma}_n$ by all strategies not of this form.

For all $n \in \mathbb{N}$, $I_n = K_n \cup L_n$, where

$$\begin{aligned} K_n &= \{i \in I_n | \text{there exists } 1 \leq t \leq n \text{ such that } i_r = p \Leftrightarrow r \geq t\} \\ L_n &= \{i \in I_n | \text{for all } 1 \leq r \leq n,\ i_r < p\}, \end{aligned}$$

K_n is the set of Infiltrator strategies which visit the target at some time and then settle there. L_n is the set of strategies which never visit the target. These sets are disjoint. Clearly, if $n < p$ then K_n is empty and $I_n = L_n$. If $n \geq p$, every element of L_n is dominated in the target game $\tilde{\Gamma}_n$.

What then of solutions to these games? If we recall, from section 2.1, our earlier definition of solutions and optimal strategies, we may observe that, for any $n \in \mathbb{N}$, as I_n and G_n are both finite then Γ_n, say, has a solution if and only if there exists a pair of strategies $(i^*, g^*) \in I_n \times G_n$ such that, for all $(i, g) \in I_n \times G_n$,

$$\Pi_n(i, g^*) \leq \Pi_n(i^*, g^*) \leq \Pi_n(i^*, g).$$

In this case, $v(\Gamma_n) = \Pi_n(i^*, g^*)$.

We saw above that, for all $n, p \in \mathbb{N}$ such that $n < p$, the function $\tilde{\Pi}_n$ is uniformly zero. Hence, for $n < p$, $v(\tilde{\Gamma}_n) = 0$ and all strategies are optimal.

We now show also that, if $p = 2$, then both Γ_n and $\tilde{\Gamma}_n$ have solutions for all $n \in \mathbb{N}$. For suppose that, for all $n \in \mathbb{N}$, the strategy pair $(i^n, g^n) \in I_n \times G_n$ is given as follows:

$$(i_r^n, g_r^n) = \begin{cases} (1,1) & \text{for } r = 1, \\ (2,1) & \text{for } 1 < r \leq n. \end{cases}$$

Then, for all $n \in \mathbb{N}$, $\Pi_n(i^n, g^n) = \lambda$, and, for all $(i, g) \in I_n \times G_n$, $\Pi_n(i, g^n) \leq \lambda \leq \Pi_n(i^n, g)$. Thus, for all $n \in \mathbb{N}$, Γ_n has a solution, $v(\Gamma_n) = \lambda$, and (i^n, g^n) is optimal. Moreover, for all $n \in \mathbb{N}$, and all $(i, g) \in I_n \times G_n$, $\tilde{\Pi}_n(i, g^n) \leq \lambda \leq \tilde{\Pi}_n(i^n, g)$, and so $\tilde{\Gamma}_n$ also has a solution, $v(\tilde{\Gamma}_n) = \tilde{\Pi}_n(i^n, g^n) = \lambda$ and (i^n, g^n) is again optimal.

So much for when $p = 2$. We see in Chapter 4 that when $p = 3$ there is again a simple solution for the target game (although not for the detection game). However, for $p \geq 4$, no such general optimal strategies exist. It has long been known, though, that in games with finite strategy sets solutions can be found by introducing probability vectors. This is what we shall consider next.

Let $n \in \mathbb{N}$. Then denote the set of all **probability vectors** on I_n [resp. G_n] by $\mathcal{I}_n$ [$\mathcal{G}_n$]. If $\iota \in \mathcal{I}_n$ [$\gamma \in \mathcal{G}_n$], and $i \in I_n$ [$g \in G_n$], then let $\iota(i)$ [$\gamma(g)$] denote the probability, under ι [γ], that strategy i [g] is chosen. If for some $\iota \in \mathcal{I}_n$ and $i \in I_n$, $\iota(i) > 0$ then we say that i is a **non-zero component** of ι.

Definition 3.2.2 *Let $n \in \mathbb{N}$. The nth extended finite detection game, is the game triple*

$$\Gamma_n^* = (\mathcal{I}_n, \mathcal{G}_n, \pi_n),$$

where (a) the strategy sets $\mathcal{I}_n$ and $\mathcal{G}_n$ are as given above, and (b) the payoff function $\pi_n : \mathcal{I}_n \times \mathcal{G}_n \to \Re$ is given as follows: For all $(\iota, \gamma) \in \mathcal{I}_n \times \mathcal{G}_n$,

$$\pi_n(\iota, \gamma) = \sum_{i \in I_n} \sum_{g \in G_n} \Pi_n(i, g)\ \gamma(g)\iota(i).$$

Definition 3.2.3 *Let $n \in \mathbb{N}$. The nth extended finite target game, is the game triple*

$$\tilde{\Gamma}_n^* = (\mathcal{I}_n, \mathcal{G}_n, \tilde{\pi}_n),$$

where (a) the strategy sets $\mathcal{I}_n$ and $\mathcal{G}_n$ are as given above, and (b) the payoff function $\tilde{\pi}_n : \mathcal{I}_n \times \mathcal{G}_n \to \Re$ is given as follows: For all $(\iota, \gamma) \in \mathcal{I}_n \times \mathcal{G}_n$,

$$\tilde{\pi}_n(\iota, \gamma) = \sum_{i \in I_n} \sum_{g \in G_n} \tilde{\Pi}_n(i, g)\ \gamma(g)\iota(i).$$

For all $n \in \mathbb{N}$, Γ_n^* and $\tilde{\Gamma}_n^*$ are finite matrix games. Therefore by the Von Neumann minimax theorem [54], all of these games have a solution. Let v_n and $\tilde{v}_n$ denote the values of Γ_n^* and $\tilde{\Gamma}_n^*$ respectively. As I_n and G_n are finite, both games have optimal strategies (not just ϵ-optimal strategies) for each player.

Recall, from above, that if $n < p$ then $v(\tilde{\Gamma}_n) = 0$. This implies that, for all $n < p$, $\tilde{v}_n = 0$ as well. On the other hand, if we define the Infiltrator strategy $\bar{\imath} \in I_n$ such that $\bar{\imath}_r = r$ for all $1 \le r \le p$ and p for all $p < r \le n$, then $\bar{\imath}$ ensures a payoff of at least λ^{p-1} in the game Γ_n^*. So, for all n, $v_n \ge \lambda^{p-1}$. Moreover, as for all n and all $(i, g) \in I_n \times G_n$, $\tilde{\Pi}_n(i, g) \le \Pi_n(i, g)$, it follows that $v_n \ge \tilde{v}_n$. In fact, if for some $n \ge p$ there exists an optimal Infiltrator strategy for Γ_n^* whose non-zero components all visit the target, then $\tilde{v}_n = v_n$. For suppose $\iota^* \in \mathcal{I}_n$ satisfies $\iota^*(i) > 0$ only if i is of the form $i_r = p$ for some $r \le n$. Then, for all $g \in G_n$ and all $i \in I_n$ such that $\iota^*(i) > 0$, $\tilde{\Pi}_n(i, g) = \Pi_n(i, g)$. Therefore, for all $\gamma \in \mathcal{G}_n$, $\tilde{\pi}_n(\iota^*, \gamma) = \pi_n(\iota^*, \gamma) \ge v_n$ as ι is optimal in Γ_n^*. This implies $\tilde{v}_n \ge v_n$ and hence they are equal.

In the next chapter we give some examples that show, for relatively small n, no such optimal strategy exists in $\tilde{\Gamma}_n^*$, and that, for all $p \ge 3$, $\tilde{v}_n < v_n$. However, for sufficiently large n we find that v_n and $\tilde{v}_n$ do get arbitrarily close. Before we can prove this however, there are some other results that we require.

Lemma 3.2.4 *(v_n) is a decreasing sequence.*

Proof Let $\gamma \in \mathcal{G}_n$ be an optimal strategy in Γ_n^*, $n \ge 1$. Form $\gamma^* \in \mathcal{G}_{n+1}$ by adding an arbitrary $(n+1)$th term to the end of each non-zero component of γ. Now, for all $i \in I_{n+1}$, denote by i' the element of I_n obtained by removing the term i_{n+1}. The longer the game goes on, the better are the Guard's chances of detecting his opponent. Thus, for all $i \in I_{n+1}$, $\pi_{n+1}(i, \gamma^*) \le \pi_n(i', \gamma) \le v_n$ as γ is optimal in Γ_n^*. Therefore $v_{n+1} \le v_n$.

Lemma 3.2.5 $(\tilde{v}_n)$ *is an increasing sequence.*

Proof We have already seen that, for $n < p$, $\tilde{v}_n = 0$. Moreover, of course $\tilde{v}_p \geq 0 = \tilde{v}_{p-1}$. Now let us assume that $n \geq p$ and deduce that $\tilde{v}_{n+1} \geq \tilde{v}_n$.

If $\iota \in \mathcal{I}_n$ is optimal in $\tilde{\Gamma}^*_n$ then we may assume that every non-zero component of ι is not strictly dominated, and thus belongs to the set K_n. Thus every non-zero component of ι has p as its last term. Now form $\iota^* \in \mathcal{I}_{n+1}$ by adding to each non-zero component of ι an $(n+1)$th term, also equal to p. Under ι^*, if the Infiltrator is not detected he is certain to reach the target by time n. So, if for all $g \in G_{n+1}$ we denote by g' the element of G_n obtained when the term g_{n+1} is removed, then $\tilde{\pi}_{n+1}(\iota^*, g) = \tilde{\pi}_n(\iota, g') \geq \tilde{v}_n$, as ι is optimal. Hence $\tilde{v}_{n+1} \geq \tilde{v}_n$.

For all $n \in \mathbb{N}$ the functions Π_n and $\tilde{\Pi}_n$ take values only in the range $[0, 1]$. Thus the sequences (v_n) and $(\tilde{v}_n)$ are bounded by 0 and 1 respectively. Therefore, as a consequence of the previous Lemmas, they are both convergent and we denote their limits by v_{lim} and $\tilde{v}_{\text{lim}}$ respectively. Observe that, as for all $n \geq p$, $v_n \geq \lambda^{p-1} > 0$, we can deduce that $v_{\text{lim}} \geq \lambda^{p-1} > 0$.

At a later point, we shall prove that $v_{\text{lim}} = v_\infty$ and $\tilde{v}_{\text{lim}} = \tilde{v}_\infty$. Before we can tackle the first of these proofs, however, there is one key result which we must verify. This is stated below as Lemma A. It will be proved shortly.

3.3 Results I

If $\iota \in \mathcal{I}_n$, $n < \infty$, and S is a subset of I_n, then $\iota(S) = \sum_{s \in S} \iota(s)$.

Lemma A *For all $\epsilon > 0$ there exists $n[\epsilon] \geq p$ and $\iota[\epsilon] \in \mathcal{I}_{n[\epsilon]}$ such that $\iota[\epsilon]$ is optimal in $\Gamma^*_{n[\epsilon]}$ and $\iota[\epsilon]\left(L_{n[\epsilon]}\right) \leq \epsilon$.*

Let us discuss Lemma A a little. Intuitively, it says that provided you take a suitably large time limit n, the game Γ^*_n has an optimal Infiltrator strategy which is almost certain to try and take him to the target. He virtually disregards the possibility of avoiding detection by remaining 'at large' rather than by seeking refuge. That means that the Infiltrator is almost playing as if he were in the game $\tilde{\Gamma}^*_n$ instead of Γ^*_n. In fact, as we shall see, this Lemma will enable us to deduce that as $n \to \infty$ the values of Γ^*_n and $\tilde{\Gamma}^*_n$ lend to a common limit.

We prove Lemma A below. The basis of the argument is as follows: Suppose the Infiltrator always insisted on using, with positive probability, some strategies which never take him to the target. In response the Guard could play as follows. He chooses a game Γ_n^* whose value v_n is close to $v_{\lim}$ and finds one of his optimal strategies γ in this game. He then chooses another game Γ_N^*, $N \gg n$, and takes as his strategy in this game probability vector $\bar{\gamma}$ which is constructed by using γ until time n and thereafter visiting the states $1, 2, \ldots, p-1$ in a random way. This is an effective strategy. Whatever the Infiltrator does the Guard has ensured he is detected with probability $1 - v_n$ within the first n time units. Thus, during the whole duration of the game Γ_N^*, the Infiltrator can avoid detection with no more than probability v_n. As n is sufficiently large, both v_n and v_N are arbitrarily close to $v_{\lim}$ and hence to each other. In particular, if during the whole of the game and especially in the time between n and N, the Infiltrator remains at large, then, if N is also sufficiently larger than n, then the Guard's random movements ensure that the Infiltrator is almost certain to be detected.

In order to put this more rigorously we must start by considering a means by which the Guard can move in a random way between the states. If $u \geq p-2$, where u is the speed of the Guard, then at each time the next state can truly be chosen at random. However, for $u < p-2$, the Guard's movement is restricted. We give a method which the Guard can use for all $u \geq 1$.

Notation *Let $P' = P \backslash \{p\}$. Let $a, b \in (P')^{2p-3}$ denote the sequences where, for $1 \leq r \leq 2p-3$, a_r and b_r are given by*

$$a_r = \begin{cases} 1 & \text{if } r = 1, \\ r-1 & \text{if } 2 \leq r \leq p, \\ (2p-1) - r & \text{if } p < r \leq 2p-3, \end{cases}$$

and

$$b_r = \begin{cases} r & \text{if } 1 \leq r < p, \\ (2p-2) - r & \text{if } p \leq r \leq 2p-3. \end{cases}$$

These sequences are set out adjacently, term by term, in the table below. Observe that, thus defined, a and b are elements of both I_{2p-3} and G_{2p-3}. We are only concerned with their use as Guard strategies.

r	1	2	3	4	$\cdots$	p-2	p-1	p	p+1	$\cdots$	2p-5	2p-4	2p-3
a_r	1	1	2	3	$\cdots$	$\cdots$	p-2	p-1	p-2	$\cdots$	$\cdots$	3	2
b_r	1	2	3	$\cdots$	$\cdots$	p-2	p-1	p-2	$\cdots$	$\cdots$	3	2	1

Lemma 3.3.1 *Let $a, b \in G_{2p-3}$ be given as above. Let $s = s_1, \ldots, s_{2p-3} \in (P')^{2p-3}$ satisfy $|s_{r+1} - s_r| \leq 1$ for all $r \geq 1$. Then, there exists t, $1 \leq t \leq 2p-3$ such that $s_t = a_t$ or $s_t = b_t$.*

Proof Assume that, for all $t \in [1, 2p-3]$, $s_t \neq a_t$ and $s_t \neq b_t$.

As $a_1 = b_1 = 1$ it follows that $s_1 \geq 2$. Let us suppose that, for some k, $1 \leq k \leq p-2$, we know that $s_k \geq k+1$. From above, $|s_{k+1} - s_k| \leq 1$, hence $s_{k+1} \geq s_k - 1 \geq k$. However, as by assumption $s_{k+1} \neq a_{k+1}$ or b_{k+1}, and since for $k+1 \leq p-1$, $a_{k+1} = k$ and $b_{k+1} = k+1$, it follows that $s_{k+1} \geq k+2 = (k+1)+1$. Therefore, by induction, for all k, $1 \leq k \leq p-1$, $s_k \geq k+1$.

However, if we take $k = p-1$, then this implies $s_{p-1} \geq p$. As $P' = \{1, \ldots, p-1\}$, this contradicts $s \in (P')^{2p-3}$. Thus our assumption was false. There does exist some t, $1 \leq t \leq 2p-3$, such that $s_t = a_t$ or $s_t = b_t$.

Now, using the previous Lemma, we shall show how for any Guard strategy $g \in G_n$, $n < \infty$, a random movement can be added to the end of g. First another piece of notation. If $s' \in (P')^{k'}$ and $s'' \in (P')^{k''}$ are two sequences, then let $s' \circ s'' \in (P')^{k'+k''}$ be the sequence $s'_1, \ldots, s'_{k'}, s''_1, \ldots, s''_{k''}$. We refer to $s' \circ s''$ as the concatenation of s' and s''.

Definition 3.3.2 (θ-operators) *(a) For all $g \in G_n$, $n < \infty$, the strategy $\hat{g} \in G_{n+p}$ is given by*

$$\hat{g}_r = \begin{cases} g_r & 1 \leq r \leq n, \\ \max\{1, g_n - r + n\} & n < r \leq n+p. \end{cases}$$

Observe that, for all $g \in G_n$, $\hat{g}_{n+p} = 1$.

(b) For all $k \in \mathbb{N}$, let $C(k) = \left\{ s^1 \circ \cdots \circ s^k \mid \text{for all } 1 \leq j \leq k,\ s^j = a \text{ or } b \right\}$.

Clearly $C(k) \subset G_{k(2p-3)}$ and $|C(k)| = 2^k$.

(c) For all $g \in G_n$, $n < \infty$, and $k \in \mathbb{N}$, let $C(g;k) = \{\hat{g} \circ c \mid c \in C(k)\}$.

Clearly $C(g;k) \subset G_{m(n;k)}$, where $m(n;k) = (n+p) + (2p-3)k$, and again $|C(g;k)| = |C(k)| = 2^k$ which is independent of g and n.

(d) For all $k \in \mathbb{N}$, the θ-operator, θ_k is a function on $\bigcup_{n \in \mathbb{N}} G_n$, defined as follows: If $g \in G_n$, $n < \infty$, then $\theta_k(g)$ is a vector on $G_{m(n;k)}$ such that, for all $h \in G_{m(n;k)}$,

$$(\theta_k(g))(h) = \begin{cases} 1/2^k & \text{if } h \in C(g;k), \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $\sum_{h \in G_{m(n;k)}} (\theta_k(g))(h) = \sum_{h \in C(g;k)} (\theta_k(g))(h) = \sum_{h \in C(g;k)} (1/2)^k = 1$, because $|C(g;k)| = 2^k$. Therefore $\theta_k(g)$ is a probability vector on $G_{m(n;k)}$ and thus an element of $\mathcal{G}_{m(n;k)}$.

Recall that, for all $n \in \mathbb{N}$, $L_n \subseteq I_n$ is the set of Infiltrator strategies which never visit the target.

Lemma 3.3.3 *For all $g \in G_n$, $n < \infty$, all $k \in \mathbb{N}$, and all $i \in L_{m(n;k)}$,*

$$\pi_{m(n;k)}(i, \theta_k(g)) \leq \left(\frac{1+\lambda}{2}\right)^k .$$

Proof Let $i \in L_{m(n;k)}$. From the definition above,

$$\pi_{m(n;k)}(i, \theta_k(g)) = \frac{1}{2^k} \sum_{h \in C(g;k)} \Pi_{m(n;k)}(i, h). \tag{3.1}$$

We shall consider the elements of this sum individually.

First, however, recall that, as $i \in L_{m(n;k)}$, for all $1 \leq r \leq m(n;k)$, $i_r < p$. For all $j = 1, \ldots, k$, let $t_j = (n+p) + (2p-3)(j-1)$. Therefore, applying Lemma 3.3.1 to the period $[t_j + 1, \ldots, t_{j+1}]$, we deduce there exists $c^j = a$ or b, and not necessarily unique, such that the part of i which corresponds to this period meets c^j at least once (In other words, for each $j = 1, \ldots, k$, if the sequences $i_{t_j+1}, \ldots, i_{t_{j+1}}$ and ${c^j}_1, \ldots, {c^j}_{2p-3}$ are compared term by term, then they coincide at least once) . The sequence $\hat{g} \circ c^1 \circ \cdots \circ c^k$ is clearly an element of $C(g;k)$. Now we partition the elements of the set $C(g;k)$ according to their relationship to this particular element.

For $z = 0, 1, \ldots, k$, there exist(s) $\binom{k}{z}$ elements of $C(g;k)$ of the form $h = \hat{g} \circ s^1 \circ \cdots \circ s^k$ where $s^j = c^j$ for precisely z values of j. Moreover, for each such h, $\omega_{m(n;k)}(i, h)$, the number of times i and h meet, is equal to at least z. Therefore $\Pi_{m(n;k)}(i, h) \leq \lambda^z$. As z runs from 0 through k every element of $C(g;k)$ is accounted for. Therefore

$$\pi_{m(n;k)}(i, \theta_k(g)) \leq \frac{1}{2^k} \sum_{r=0}^{k} \lambda^r \binom{k}{r} = \frac{1}{2^k}(1+\lambda)^k.$$

This completes the proof.

One step remains before we can give a proof of Lemma A. We shall show how, for any probability vector $\gamma \in \mathcal{G}_n$, $n \in \mathbb{N}$, random motion is added to γ. The following definition closely follows Definition 3.3.2. In particular, observe from that, that for all $k, n \in \mathbb{N}$ if $g, g' \in G_n$ are distinct, then $C(g;k)$ and $C(g';k)$ are disjoint subsets of $G_{m(n;k)}$.

Definition 3.3.4 (Θ-operators) *For all $k \in \mathbb{N}$, the Θ-operator, Θ_k, is a function on $\cup_{n\in\mathbb{N}}\mathcal{G}_n$, defined as follows: If $\gamma \in \mathcal{G}_n$, $n < \infty$, then $\Theta_k(\gamma)$ is a vector on $G_{m(n;k)}$ such that, for all $h \in G_{m(n;k)}$ then*

$$(\Theta_k(\gamma))(h) = \begin{cases} (1/2)^k \gamma(g) & \textit{if } h \in C(g;k), \textit{ for some } g \in G_n, \\ 0 & \textit{otherwise.} \end{cases}$$

Recall that for all $k \in \mathbb{N}$ and all $g \in G_n$, $n < \infty$, $|C(g;k)| = 2^k$. Also, from above, if $g \neq g'$, then $C(g;k) \cap C(g';k) = \emptyset$. Therefore, if $\gamma \in \mathcal{G}_n$,

$$\sum_{h \in G_{m(n;k)}} (\Theta_k(\gamma))(h) = \sum_{g \in G_n} \left(\sum_{h \in C(g;k)} \frac{1}{2^k}\gamma(g) \right) = \sum_{g \in G_n} \gamma(g) = 1$$

and so $\Theta_k(\gamma)$ is a probability vector. Thus $\Theta_k(\gamma) \in \mathcal{G}_{m(n;k)}$.

Lemma 3.3.5 *For all $\gamma \in \mathcal{G}_n$, $n < \infty$, all $k \in \mathbb{N}$, and all $i \in L_{m(n;k)}$,*

$$\pi_{m(n;k)}(i, \Theta_k(\gamma)) \leq \left(\frac{1+\lambda}{2} \right)^k.$$

Proof This is just a corollary to the previous Lemma. By the above definition,

$$\begin{aligned}
\pi_{m(n;k)}(i, \Theta_k(\gamma)) &= \sum_{g \in G_n} \left(\frac{1}{2^k} \sum_{h \in C(g;k)} \Pi_{m(n;k)}(i,h) \right) \gamma(g) \\
&= \sum_{g \in G_n} \pi_{m(n;k)}(i, \theta_k(g)) \;\; \gamma(g) \text{ from equation [3.1]}, \\
&\leq \sum_{g \in G_n} \left(\frac{1+\lambda}{2} \right)^k \gamma(g) \text{ by Lemma 3.3.3} \\
&= \left(\frac{1+\lambda}{2} \right)^k \text{ as required.}
\end{aligned}$$

Finally, let us restate and prove Lemma A.

Lemma 3.3.6 (Lemma A) *For all $0 < \epsilon < 1$, there exists $n[\epsilon] \geq p$ and $\iota[\epsilon] \in \mathcal{I}_{n[\epsilon]}$ such that $\iota[\epsilon]$ is optimal in $\Gamma^*_{n[\epsilon]}$ and $\iota[\epsilon](L_{n[\epsilon]}) \leq \epsilon$.*

Proof We assume the result to be false and show that this leads to a contradiction. Hence assume that there exists $0 < \epsilon < 1$ such that, for all $n \geq p$, every $\iota \in \mathcal{I}_n$ which is optimal in Γ_n^* satisfies $\iota(L_n) > \epsilon$.

Take $\gamma \in \mathcal{G}_n$ to be an optimal Guard strategy in Γ_n^*, for some $n \in \mathbb{N}$. Then, for some $k \in \mathbb{N}$, let $\bar{\gamma} = \Theta_k(\gamma)$. By definition of the Θ-operator, $\bar{\gamma} \in \mathcal{G}_m$, where $m = m(n; k) \geq p$. Now take any $\iota \in \mathcal{I}_m$ which is optimal for the Infiltrator in Γ_m^*. As $m \geq p$ then, by assumption, ι satisfies $\iota(L_m) > \epsilon$. Let us now consider the payoff $\pi_m(\iota, \bar{\gamma})$.

We observed earlier that, for all $n \geq p$, the Infiltrator strategies for Γ_n which were not strictly dominated belonged to either K_n or L_n. Hence, we deduce here that $\iota(K_m) + \iota(L_m) = 1$, and

$$\pi_m(\iota, \bar{\gamma}) = \sum_{i \in K_m} \pi_m(i, \bar{\gamma})\iota(i) + \sum_{i \in L_m} \pi_m(i, \bar{\gamma})\iota(i). \tag{3.2}$$

Suppose that $i \in K_m$. Let i' denote that element of I_n obtained by removing the terms $i_{n+1}, i_{n+2}, \ldots, i_m$ from i. The longer the game goes on, the worse are the Infiltrator's chances of staying undetected. Thus, as γ is optimal in Γ_n^*, we deduce that

$$\pi_m(i, \bar{\gamma}) \leq \pi_n(i', \gamma) \leq v_n.$$

Now suppose that $i \in L_m$. Recalling that $\bar{\gamma} = \Theta_k(\gamma)$ we can immediately apply Lemma 3.3.5. Hence,

$$\pi_m(i, \bar{\gamma}) \leq \left(\frac{1+\lambda}{2}\right)^k.$$

Finally, by assumption we know that $\iota(K_m) = 1 - \iota(L_m) < 1 - \epsilon$ and certainly that $\iota(L_m) \leq 1$. Therefore, returning to equation (3.2), we now have that

$$\begin{aligned}
\pi_m(\iota, \bar{\gamma}) &\leq v_n \sum_{i \in K_m} \iota(i) + \left(\frac{1+\lambda}{2}\right)^k \sum_{i \in L_m} \iota(i) \\
&< v_n(1-\epsilon) + \left(\frac{1+\lambda}{2}\right)^k \\
&= v_m - \left[\epsilon v_n - (v_n - v_m) - \left(\frac{1+\lambda}{2}\right)^k\right].
\end{aligned}$$

For all $n \in \mathbb{N}$, $v_n \geq v_{\lim} > 0$; moreover it was initially supposed that $\epsilon > 0$. Thus $\epsilon v_n \geq \delta$, where δ is itself > 0. However, by taking suitably large values of n and k, the expressions $(v_n - v_m)$ (which is non-negative since $m > n$, and (v_n) is decreasing) and $\left(\frac{1+\lambda}{2}\right)^k$ can both be made arbitrarily small. Therefore, there exist n and k such that $\pi_m(\iota, \bar{\gamma}) < v_m$. This contradicts the fact that ι is optimal in Γ_m^*. Hence our original assumption was false. This proves the Lemma.

Lemma A enables us to prove two other important results, the first concerning the infinite game Γ_∞^*. But, before approaching again the value of Γ_∞^*, we recall the strategy sets $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$. By Definition 2.5.5 these are the sets of probability measures on the Borel-algebras on I_∞ and G_∞ respectively. We now consider an important subset of each of these sets.

A measure μ on a measurable space $(X, \mathcal{A})$ is said to be **discrete** if there are denumerably many points $x_i \in X$ and scalars $m_i \in [0, \infty]$ such that, for all $A \in \mathcal{A}$, $\mu(A) = \sum_{x_i \in A} m_i$ (see Billingsley [13], page 134). It is clear that a discrete probability measure belonging to $\mathcal{I}_\infty$ or $\mathcal{G}_\infty$ is a probability vector on some denumerable subset of I_∞ or G_∞ respectively. These are the kind of probability measures we shall now make use of.

There will be several occasions in the future on which these discrete measures will be particularly important. In particular we encounter them when we extend strategies from the finite games and use them in Γ_∞^* and $\tilde{\Gamma}_\infty^*$. It is these extensions that must now be defined.

Definition 3.3.7 *Let $n < \infty$.*

(a) Let $i \in I_n$. The extended strategy $\bar{i} \in I_\infty$ is given by

$$\bar{i}_r = \begin{cases} i_r, & r \leq n, \\ i_n & r > n. \end{cases}$$

Note that, for all $i', i'' \in I_n$, $\bar{i}' \neq \bar{i}'' \Leftrightarrow i' \neq i''$.

(b) Let $\iota \in \mathcal{I}_n$, and let $X_\iota = \{i \in I_n \mid \iota(i) > 0\}$. The function $\bar{\iota}: I_\infty \to \Re$ is then given as follows. For all $i \in I_\infty$,

$$\bar{\iota}(i) = \begin{cases} \iota(j) & \text{if there exists } j \in X_\iota \text{ such that } \bar{j} = i, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\bar{\iota}$ is a probability vector on I_∞. Thus, by the discussion above, $\bar{\iota} \in \mathcal{I}_\infty$.

(c) Let $g \in G_n$ and $\gamma \in \mathcal{G}_n$. The extended strategies $\overline{g} \in G_\infty$ and $\overline{\gamma} \in \mathcal{G}_\infty$ are defined as in (a) and (b) above.

Observe that we need not have defined $\bar{i}$ and $\overline{g}$ so precisely. It would be sufficient for $\bar{i}$ and $\bar{g}$ to be any elements of I_∞ and G_∞ such that, for all $1 \leq r \leq n$, $\bar{i}_r = i_r$ and $\overline{g}_r = g_r$.

Theorem 3.3.8 $v_\infty = v_{\lim}$.

Proof Let $\epsilon \in (0,1)$. By Lemma A, there exists $n[\epsilon] \geq p$ and $\iota[\epsilon] \in \mathcal{I}_{n[\epsilon]}$ such that $\iota[\epsilon]$ is optimal in $\Gamma^*_{n[\epsilon]}$ and $\iota[\epsilon](L_{n[\epsilon]}) \leq \epsilon$. Recall that the non-zero components of $\iota[\epsilon]$ cannot be strictly dominated. Therefore they are elements of either $K_{n[\epsilon]}$ or $L_{n[\epsilon]}$. Consequently, we deduce that, for all $h \in G_{n[\epsilon]}$,

$$\pi_{n[\epsilon]}(\iota[\epsilon], h) = \sum_{i \in K_{n[\epsilon]}} \Pi_{n[\epsilon]}(i,h)\ \iota[\epsilon](i) + \sum_{i \in L_{n[\epsilon]}} \Pi_{n[\epsilon]}(i,h)\ \iota[\epsilon](i) \geq v_{n[\epsilon]},$$

and hence, as $\sum_{i \in L_{n[\epsilon]}} \Pi_{n[\epsilon]}(i,h)\ \iota[\epsilon](i) \leq \sum_{i \in L_{n[\epsilon]}} \iota[\epsilon](i) \leq \epsilon$, that

$$\sum_{i \in K_{n[\epsilon]}} \Pi_{n[\epsilon]}(i,h)\ \iota[\epsilon](i) \geq v_{n[\epsilon]} - \epsilon. \tag{3.3}$$

Now let us consider the strategy $\overline{\iota[\epsilon]} \in \mathcal{I}_\infty$, formed according to Definition 3.3.7(b). Also, for all $g \in G_\infty$, let g' denote the element of $G_{n[\epsilon]}$ obtained when all of the terms except $g_1, \ldots, g_{n[\epsilon]}$ are removed from g. We observe that, for all $k \in K_{n[\epsilon]}$ which are non-zero components of $\iota[\epsilon]$, $\overline{k} \in K_\infty$, and also $\omega(\overline{k}, g) = \omega_{n[\epsilon]}(k, g')$, and so $\Pi_\infty(\overline{k}, g) = \Pi_{n[\epsilon]}(k, g')$. Therefore, for all $g \in G_\infty$,

$$\begin{aligned}
\pi_\infty(\overline{\iota[\epsilon]}, g) &= \sum_{i \in I_\infty} \Pi_\infty(i,g)\ \overline{\iota[\epsilon]}(i) \\
&\geq \sum_{i \in K_\infty} \Pi_\infty(i,g)\ \overline{\iota[\epsilon]}(i) \\
&= \sum_{k \in K_{n[\epsilon]}} \Pi_\infty(\overline{k},g)\ \iota[\epsilon](k) \qquad \text{by Definition 3.3.7(b),} \\
&= \sum_{k \in K_{n[\epsilon]}} \Pi_{n[\epsilon]}(k,g')\ \iota[\epsilon](k) \\
&\geq v_{n[\epsilon]} - \epsilon \qquad \text{from equation (3.3),} \\
&\geq v_{\lim} - \epsilon \qquad \text{as } (v_n) \text{ decreases.}
\end{aligned}$$

We now complete the proof by deducing the existence of a strategy $\overline{\gamma} \in \mathcal{G}_\infty$ such that, for all $i \in I_\infty$, $\pi_\infty(i, \overline{\gamma}) \leq v_{\lim} + \epsilon$. First we take $\gamma \in \mathcal{G}_n$ to be optimal in Γ_n^*, where n is sufficiently large to satisfy $v_n - v_{\lim} \leq \epsilon$. Then define $\overline{\gamma} \in \mathcal{G}_\infty$ according to Definition 3.3.7(c). Finally, for all $i \in I_\infty$, let $i' = i_1, \ldots, i_n$. Clearly, for all $i \in I_\infty$, $\pi_\infty(i, \overline{\gamma}) \leq \pi_n(i', \gamma)$. Therefore, as γ is optimal in Γ_n^*,

$$\pi_\infty(i, \overline{\gamma}) \leq v_n \leq v_{\lim} + \epsilon.$$

This completes the proof.

Let us immediately move on to the next important result which stems from Lemma A.

Theorem 3.3.9 $v_{\lim} = \tilde{v}_{\lim}$.

Proof Let $\epsilon \in (0,1)$. By Lemma A, there exists $n[\epsilon] \geq p$ and $\iota[\epsilon] \in \mathcal{I}_{n[\epsilon]}$ such that $\iota[\epsilon]$ is optimal in $\Gamma^*_{n[\epsilon]}$ and $\iota[\epsilon](L_{n[\epsilon]}) \leq \epsilon$. As in the previous proof (see equation (3.3)), we deduce that, for all $g \in G_{n[\epsilon]}$,

$$\sum_{i \in K_{n[\epsilon]}} \Pi_{n[\epsilon]}(i, g)\ \iota[\epsilon](i) \geq v_{n[\epsilon]} - \epsilon.$$

Moreover, recall that for all $n \geq p$, if $i \in K_n$ and $g \in G_n$, then the payoffs $\Pi_n(i, g)$ and $\tilde{\Pi}_n(i, g)$ are equal. Therefore, for all $g \in G_{n[\epsilon]}$,

$$\begin{aligned}
\tilde{\pi}_{n[\epsilon]}(\iota[\epsilon], g) &= \sum_{i \in I_{n[\epsilon]}} \tilde{\Pi}_{n[\epsilon]}(i, g)\ \iota[\epsilon](i) \\
&\geq \sum_{i \in K_{n[\epsilon]}} \tilde{\Pi}_{n[\epsilon]}(i, g)\ \iota[\epsilon](i) \\
&= \sum_{i \in K_{n[\epsilon]}} \Pi_{n[\epsilon]}(i, g)\ \iota[\epsilon](i) \quad \text{changing the payoff function,} \\
&\geq v_{n[\epsilon]} - \epsilon \quad \text{from above.}
\end{aligned}$$

Hence $\tilde{v}_{n[\epsilon]} \geq v_{n[\epsilon]} - \epsilon$.

As (v_n) and $(\tilde{v}_n)$ are decreasing and increasing sequences respectively, we deduce that, for all $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that $v_n - \tilde{v}_n \leq \epsilon$. It follows from this that $v_{\lim} - \tilde{v}_{\lim} \leq 0$. We already know that, for all $n \in \mathbb{N}$, $\tilde{v}_n \leq v_n$, and hence that $\tilde{v}_{\lim} \leq v_{\lim}$. This completes the proof.

It will be seen that, after considering Theorems 3.3.8 and 3.3.9 above, we are just one link away from proving the chain of equalities

$$v_\infty = v_{\lim} = \tilde{v}_{\lim} = \tilde{v}_\infty.$$

3.4 Results II

We lack a proof that $\tilde{v}_{\text{lim}} = \tilde{v}_\infty$. In order to prove this, there is again one particularly important preliminary result. We shall distinguish this by giving it the title Lemma B. The proof of Lemma B requires a construction by which optimal Guard strategies from a sequence of finite games are combined to form a strategy for the infinite game $\tilde{\Gamma}^*_\infty$.

The following definition makes use of the notation associated with the θ- and Θ-operators. The reader is referred to Definitions 3.3.2 and 3.3.4 for a reminder of this.

Definition 3.4.1 (ψ- and Ψ-strategies) *Let $k \in \mathbb{N}$.*

(a) The sequence (t_n), $n \geq 0$, is given as follows: $t_0 = 1$ and for $n \geq 1$, $t_n = m(t_{n-1}; k)$.

Observe that (t_n) is monotonic increasing.

(b) Let $n \in \mathbb{N}$. The ψ-strategy, $\psi_k^{[n]}$, is defined as follows: First take $\gamma \in \mathcal{G}_{t_{n-1}}$ to be any optimal strategy in $\tilde{\Gamma}^*_{t_{n-1}}$, and form the strategy $\Theta_k(\gamma) \in \mathcal{G}_T$, where $T = m(t_{n-1}; k) = t_n$. Then, recalling Definition 3.3.7(b), let $\psi_k^{[n]} \in \mathcal{G}_\infty$ be given by $\psi_k^{[n]} = \overline{\Theta_k(\gamma)} \in \mathcal{G}_\infty$.

Thus, the strategy $\psi_k^{[n]} \in \mathcal{G}_\infty$ behaves like γ until time t_{n-1}, then has a time of random motion until time t_n, and from then on is defined arbitrarily. Let us denote the set of non-zero components of all the strategies $\psi_k^{[1]}, \psi_k^{[2]}, \ldots$ by X. Observe that X is a countable subset of G_∞.

(c) Let $z \in \mathbb{N}$. The Ψ-strategy, $\Psi_{k,z}$, is defined as follows: $\Psi_{k,z}$ is a vector on X such that, if $g \in X$, then

$$\Psi_{k,z}(g) = \frac{1}{z}\sum_{n=1}^{z} \psi_k^{[n]}(g).$$

Clearly $\Psi_{k,z}$ is a probability vector X. As X is a countable set, $\Psi_{k,z}$ is a discrete measure and hence an element of $\mathcal{G}_\infty$.

Lemma 3.4.2 (Lemma B) *For all $\epsilon > 0$ there exists $\gamma[\epsilon] \in \mathcal{G}_\infty$ such that, for all $i \in I_\infty$,*

$$\tilde{\pi}_\infty(i, \gamma[\epsilon]) \leq \tilde{v}_{\text{lim}} + \epsilon.$$

Proof We show first that, for all $i \in I_\infty$, and all $k, n \in \mathbb{N}$, if $\Psi_{k,z}$ is defined as above, then

$$\tilde{\pi}_\infty(i, \Psi_{k,z}) \leq \tilde{v}_{\lim} + \left[\left(\frac{1+\lambda}{2} \right)^k + \frac{1}{z} \right].$$

Let $i \in I_\infty$. As we are looking at a target game rather than a detection game, we may assume that at some time $\tau \in \mathbb{N}$ the Infiltrator reaches the target and settles there. Thus $i_r = p$ if and only if $r \geq \tau$. One immediate consequence of this is that, by Definition 2.4.1, $\tilde{\pi}_\infty(i, \gamma) = \pi_\infty(i, \gamma)$ for all $\gamma \in \mathcal{G}_\infty$. In order to pinpoint the value of τ relative to the sequence (t_n), let $s = \min\{r | t_r \geq \tau\}$. By Definition 2.3.1, $\tau \geq p$ and hence $s \geq 1$. Thus we have that

$$t_0 < \cdots < t_{s-1} < \tau \leq t_s < t_{s+1} < \cdots$$

To evaluate the required payoff, observe that, by Definition 3.4.1(c),

$$\tilde{\pi}_\infty \left(i, \Psi_{k,z} \right) = \frac{1}{z} \sum_{n=1}^{z} \tilde{\pi}_\infty \left(i, \psi_k^{[n]} \right).$$

We now consider the individual payoffs $\tilde{\pi}_\infty \left(i, \psi_k^{[n]} \right)$ in three cases, according to whether $n <$, $=$ or $> s$.

Suppose $n < s$, for then we deduce that $t_n < \tau$. Thus, if we compare i with the Guard strategy $\psi_k^{[n]}$, we see that the Infiltrator stays at large until after the Guard's period of random motion. In other words, if we let $i' \in I_{t_n}$ denote the sequence $i_1, \ldots, i_{t_n}$, then $\tau > t_n$ implies that

$$\tilde{\pi}_\infty \left(i, \psi_k^{[n]} \right) = \pi_\infty \left(i, \psi_k^{[n]} \right) \leq \pi_{t_n} \left(i', \Theta_k(\gamma) \right).$$

As $\tau > t_n$, then $i' \in L_{t_n}$. Recall that, by Lemma 3.3.5, for all $j \in L_{t_n}$, $\pi_{t_n} \left(j, \Theta_k(\gamma) \right) \leq \left(\frac{1+\lambda}{2} \right)^k$. Therefore, for all $n < s$,

$$\tilde{\pi}_\infty \left(i, \psi_k^{[n]} \right) \leq \left(\frac{1+\lambda}{2} \right)^k.$$

This completes the first case.

In fact, if $z \leq s - 1$, then this is the only case we need. It follows immediately that

$$\tilde{\pi}_\infty \left(i, \Psi_{k,z} \right) \leq \frac{1}{z} \sum_{n=1}^{z} \left(\frac{1+\lambda}{2} \right)^k = \left(\frac{1+\lambda}{2} \right)^k$$

and we are finished. So now assume that $z \geq s$. This forces us to consider the other two cases mentioned above.

First let $n = s$. In this case we have that $t_{n-1} < \tau \leq t_n$, and so, if we consider the Guard strategy $\psi_k^{[s]}$, then the Infiltrator reaches the target during the Guard's period of random movement. Thus the Guard cannot guarantee any probability of detecting the Infiltrator. The only upper bound for the payoff that we deduce is

$$\tilde{\pi}_\infty\left(i, \psi_k^{[s]}\right) \leq 1.$$

Now for the third case suppose that $n > s$. Thus $\tau \leq t_{n-1} < t_n$. Here, the Infiltrator reaches that target by time t_{n-1} and we can thus exploit the fact that $\psi_k^{[n]}$ is formed from a strategy γ which is optimal in $\tilde{\Gamma}^*_{t_{n-1}}$. Hence, if we let i'' denote the sequence $i_1, \ldots, i_{t_{n-1}}$, we have

$$\tilde{\pi}_\infty\left(i, \psi_k^{[n]}\right) = \tilde{\pi}_{t_{n-1}}\left(i'', \gamma\right) \leq \tilde{v}_{t_{n-1}} \leq \tilde{v}_{\text{lim}}.$$

This completes the third case.

Therefore, reconstructing our original sum, we deduce that

$$\begin{aligned}
\tilde{\pi}_\infty\left(i, \Psi_{k,z}\right) &\leq \frac{1}{z}\left[\sum_{n=1}^{s-1}\left(\frac{1+\lambda}{2}\right)^k + 1 + \sum_{n=s+1}^{z} \tilde{v}_{\text{lim}}\right] \\
&\leq \frac{1}{z}\left[z\left(\frac{1+\lambda}{2}\right)^k + 1 + z\tilde{v}_{\text{lim}}\right] \\
&= \tilde{v}_{\text{lim}} + \left[\left(\frac{1+\lambda}{2}\right)^k + \frac{1}{z}\right]
\end{aligned}$$

as required.

Finally, it is clear that for all $\epsilon > 0$, then, by taking k and z sufficiently large and letting $\gamma[\epsilon] = \Psi_{k,z}$ then, for all $i \in I_\infty$,

$$\tilde{\pi}_\infty\left(i, \gamma[\epsilon]\right) \leq \tilde{v}_{\text{lim}} + \epsilon.$$

Theorem 3.4.3 $\tilde{v}_\infty = \tilde{v}_{\text{lim}}$.

Proof After Lemma B, we need only show that, for all $\epsilon > 0$ there exists $\iota[\epsilon] \in \mathcal{I}_\infty$ such that, for all $g \in G_\infty$,

$$\tilde{\pi}_\infty\left(\iota[\epsilon], g\right) \geq \tilde{v}_{\text{lim}} - \epsilon.$$

First let $\iota \in \mathcal{I}_n$ be any optimal strategy in $\tilde{\Gamma}_n^*$, where $n \geq p$ and is sufficiently large to ensure that $\tilde{v}_{\text{lim}} - \tilde{v}_n \leq \epsilon$. Now, recalling Definition 3.3.7, let $\iota[\epsilon] = \overline{\iota}$. As ι is optimal we may assume that any non-zero component j is an element of K_n, and thus $\overline{j} \in K_\infty$. In fact, if for all $g \in G_\infty$ we let $g' \in G_n$ denote the sequence $g_1, \ldots, g_n$, then clearly $\omega(\overline{j}, g) = \omega_n(j, g')$. Therefore, for all $g \in G_\infty$,

$$\tilde{\pi}_\infty\left(\iota[\epsilon], g\right) = \tilde{\pi}_n(\iota, g') \geq \tilde{v}_n,$$

as ι is optimal. As $\tilde{v}_n \geq \tilde{v}_{\text{lim}} - \epsilon$, the proof is complete.

Theorem 3.4.4 $v_\infty = \tilde{v}_\infty$.

Proof This is immediate from the last three Theorems.

Chapter 4

EXAMPLES

4.1 Introduction

In this chapter we take a break from the theory to look at some examples. When p, the number of states, is small we can directly find both values and optimal or ϵ-optimal strategies. From these few examples certain important features emerge. For example, we meet for the first time a type of Infiltrator strategy known as a wait-and-run strategy.

We consider in turn the different games for $p = 2$, 3 and 4. In each section we look at both the infinite games Γ_∞^* and $\tilde{\Gamma}_\infty^*$, and the finite games Γ_n^* and $\tilde{\Gamma}_n^*$. The solutions vary with the value of u, the Guard's speed. In particular we describe games with $u = 1$ as slow Guard games and those with $u \geq p - 2$ as fast Guard games. In the earlier chapters the general theory has played down the role of the parameters p, n, u and λ. Here we find that the form of the solution depends on their relative values. In general we have considered that $p \geq 2$, $u \geq 1$, $0 \leq \lambda < 1$ and $0 \leq n \leq \infty$. Now let us be more specific. As a start note that if $\lambda = 0$ then μ, the probability of detection, is one and so the Guard can detect the Infiltrator immediately. Thus all the games are trivial. Throughout this chapter let us assume that $\lambda > 0$.

In considering examples we encounter finite matrix games. These are games in which the strategy sets are finite and the extended strategy sets just probability vectors. It is conventional to denote these games by matrices in which the rows and columns correspond to the strategies of the maximiser and minimiser respectively. Let us illustrate this by looking at the detection game Γ_n^* when $p = 3$, $n = 3$ and $u \geq 1$. The set I_3 contains five elements which correspond to the rows of the matrix below. As $p = 3$ the Guard can only move between states one and two and so $|G_3| = 2^3 = 8$. However, as any element of G_3 of the form $g = 2, g_2, g_3$ dominates the element $g' = 1, g_2, g_3$, we need only consider the four strategies which start at state one. These correspond to the columns in the matrix below.

	1 1 1	1 1 2	1 2 1	1 2 2
1, 2, 3	λ	λ	λ^2	λ^2
1, 2, 2	λ	λ^2	λ^2	λ^3
1, 2, 1	λ^2	λ	λ^3	λ^2
1, 1, 2	λ^2	λ^3	λ	λ^2
1, 1, 1	λ^3	λ^2	λ^2	λ

Moreover, the Infiltrator strategies $1, 2, 2$ and $1, 2, 1$ are both dominated by $1, 2, 3$. Intuitively this is because, if the Infiltrator moves to the state adjacent to the target, he should move to the target on his next go. Thus we further reduce the matrix to

	1 1 1	1 1 2	1 2 1	1 2 2
1, 2, 3	λ	λ	λ^2	λ^2
1, 1, 2	λ^2	λ^3	λ	λ^2
1, 1, 1	λ^3	λ^2	λ^2	λ

The value can now be calculated to be $2\lambda(1 + \lambda)^2/L$, where $L = 2(3 + \lambda)$. The probability distributions $(2 + 2\lambda, 2, 2)/L$ and $(1, 1, 2 + \lambda, 2 + \lambda)/L$ are optimal for the Infiltrator and Guard respectively.

When a game is expressed as a matrix, it is common to refer to the rows and columns as the pure strategies for the players. For $n < \infty$, the sets of pure strategies are I_n and G_n. In general, probability distributions on the pure strategies are known as mixed strategies. Thus, for $n < \infty$, the sets of mixed strategies are $\mathcal{I}_n$ and $\mathcal{G}_n$. In fact, for the infinite games we shall also to refer to I_∞ and G_∞ as the sets of pure strategies and $\mathcal{I}_\infty$ and $\mathcal{G}_\infty$ as the sets of mixed strategies.

This chapter is intended to give insight into the kind of strategies that are often optimal. It also demonstrates some of the problems encountered in looking for general solutions to these games.

4.2 Two State Games

Lemma 4.2.1 *If $p = 2$, then for all $u \geq 1$, $0 < \lambda < 1$, and $n \geq 2$,*

$$v_\infty = \tilde{v}_\infty = v_n = \tilde{v}_n = \lambda.$$

Proof When $p = 2$ the Infiltrator starts in the state adjacent to the target. Clearly it is optimal for him to move immediately to safety. The Guard can only stay at state one. He has one chance to detect the Infiltrator. Hence the values of the games are all λ, the probability that the Infiltrator avoids detection on this one occasion.

4.3 Three State Games

Once $p > 2$, the games Γ^*_∞, $\tilde{\Gamma}^*_\infty$, Γ^*_n and $\tilde{\Gamma}^*_n$ begin to develop different optimal strategies and different values. However for $p = 3$ we can still deal with most of them together.

Lemma 4.3.1 *If $p = 3$, then for all $u \geq 1$, $0 < \lambda < 1$, and $n \geq 3$,*

$$v_\infty = \tilde{v}_\infty = \tilde{v}_n = \lambda^2$$

Proof Consider first the target games. The Infiltrator's optimal strategy is again to move directly to the target. This leaves him open to detection on just two occasions and hence guarantees a payoff of at least λ^2. The Guard plays optimally by following the sequence $1, 2, 2, \ldots$ for as long as the game lasts (till the time limit if there is one, or forever otherwise). For, unless the Infiltrator never goes to the target, in which case the payoff is zero, he must always meet the sequence above at least twice. Thus the Guard ensures that the payoff is no greater than λ^2. Hence, for $n \geq 3$, $\tilde{v}_n = \tilde{v}_\infty = \lambda^2$. Finally, from Lemma 3.4.4, $v_\infty = \tilde{v}_\infty$. This completes the proof.

We have just shown that v_∞, the value of the game Γ^*_∞, is λ^2. But now let us consider the strategies in this game. It is again optimal for the Infiltrator to head straight for the target, as this guarantees a payoff of at least λ^2. However, as the Infiltrator need never move from state one, the Guard strategy $1, 2, 2, \ldots$ is not optimal here. In fact, as we now show, there is no optimal strategy for the Guard.

Let us suppose that there were an optimal strategy for the Guard. Of course it need not be pure. However, to keep the payoff down to λ^2 against the Infiltrator strategy $1,2,3,3,\ldots$, the Guard must start by visiting state one and then state two. To also keep the payoff to λ^2 against $1,1,2,3,3,\ldots$, he must then stay at state two at time 3. By considering all Infiltrator strategies of the form $\overbrace{1,1,\ldots,1}^{\theta \text{ terms}},2,3,3\ldots$, where $\theta \geq 1$, we deduce that the Guard's only optimal strategy must be the pure strategy $1,2,2,\ldots$. However, as we saw above, this fails to restrict the payoff to λ^2 against the Infiltrator pure strategy $1,1,\ldots$. Therefore we have a contradiction. There is no optimal strategy for the Guard.

Let $\epsilon > 0$ and we construct an ϵ-optimal strategy as follows. At time 1 the Guard moves to state one. At each time thereafter, independently of his previous moves, he moves to state one with probability δ and state two with probability $1-\delta$, where $0 < \delta \ll 1$. If the Infiltrator stays at state one forever then he is detected with probability one, since, as $n \to \infty$, $\lim \lambda(\delta\lambda + 1 - \delta)^n = 0$. Moreover, if he moves from state one, it is clearly best to do so immediately and then make straight for the target. Thus, the Infiltrator reaches the target with probability at most $\lambda\,[\lambda(1-\delta)+\delta]$. By taking δ sufficiently small, this probability can be made arbitrarily close to λ^2. In fact, if $\delta \leq \frac{\epsilon}{\lambda(1-\lambda)}$, then this is an ϵ-optimal strategy for the Guard.

When $p = 3$, the only games we have not found solutions for are the finite detection games. As these are finite they have both values and optimal mixed strategies. For example, consider the game Γ_3^*. We saw in section 1 that this game can be reduced to the matrix

	1,1,1	1,1,2	1,2,1	1,2,2
1,2,3	λ	λ	λ^2	λ^2
1,1,2	λ^2	λ^3	λ	λ^2
1,1,1	λ^3	λ^2	λ^2	λ

Optimal strategies for the Infiltrator and Guard are given by the probability distributions $\iota = (2+2\lambda, 2, 2)/L$ and $\gamma = (1, 1, 2+\lambda, 2+\lambda)/L$, where $L = 2(3+\lambda)$. The value $v_3 = \pi_3(\iota, \gamma) = 2\lambda(1+\lambda)^2/L$.

For all $0 < \lambda < 1$, $2\lambda(1+\lambda)^2/L > \lambda^2$. In fact we can see that, for all $n \in \mathbb{N}$, $v_n > v_\infty = \lambda^2$. For suppose there exists t such that $v_t \leq \lambda^2$. Any optimal strategy for the Guard in Γ_t^* then ensures that the Infiltrator is detected with at least probability $1 - v_t \geq 1 - v_\infty$. By arbitrarily extending one of these optimal strategies we can construct a mixed strategy for Γ_∞^* which also ensures detection with at least probability $1 - v_\infty$. However this would then be an optimal strategy for the Guard. This is a contradiction for, as we have just proved, the Guard has no optimal strategy in Γ^*_∞. Therefore, for all $n \in \mathbb{N}$, $v_n > v_\infty$. As, from the previous chapter, for all n, $\tilde{v}_n \leq \tilde{v}_\infty = v_\infty$, we also deduce that $\tilde{v}_n < v_n$.

Let us return to the matrix above. The pure strategies $1,2,3$, $1,1,2$ and $1,1,1$ are examples of what are sometimes known as **wait-and-run strategies** (we think it is Lalley [47] who is responsible for the first use of this title). In our detection games, the set W of wait-and-run strategies is given as follows. For all n and p

$$W = \{\overbrace{1,1,\ldots,1}^{\theta \text{ terms}},2,3,\ldots,p-1,p,p,\ldots \mid \text{where } 1 \leq \theta \leq n\ \}.$$

Under any wait-and-run strategy the Infiltrator waits for some time at state one and then moves straight towards the target. He never retreats and never loiters, except at state one or the target. If $\theta > n - p + 1$, the wait-and-run strategy ends up short of the target. Such a strategy is of no use in a target game but can feature (as above) in an optimal strategy for a detection game. If we denote by $\tilde{W}$ the wait-and-run strategies for the target games, then if $n \geq p$

$$\tilde{W} = \{\overbrace{1,1,\ldots,1}^{\theta \text{ terms}},2,3,\ldots,p-1,p,p,\ldots \mid \text{where } 1 \leq \theta \leq n - p + 1\ \},$$

and if $n < p$ then $\tilde{W}$ is empty.

When $p = 3$ and $n = 3$ we have seen that we need only consider wait-and-run strategies. If $p = 3$, this is the case for any $n \leq \infty$. In general this is not so, although, as we shall see, wait-and-run strategies occur frequently in the optimal strategies of many games.

4.4 Four State Games

It is worth briefly considering what goes on when $p = 4$. We will look at just a couple of the games. This is sufficient to illustrate some of the difficulties that begin to arise. The first result is valid only when the Guard is at least twice as fast as the Infiltrator. For the first time the Guard's speed affects the solution. This is not suprising since, when $p = 2$ or 3 his speed (assuming he is no slower than his opponent) is clearly unimportant. However, once $p \geq 4$, if $u \geq 2$ the Guard may be able to exploit his ability to move between states one and three without having to go via state two.

		state			
		one	two	three	four
time	1	w_2w_1			
	2	w_2	w_1		
	3		w_2	w_1	
	4			w_2	w_1
	≥ 5				w_2w_1

		state		
		one	two	three
time	1	g_2g_1		
	2	g_2	g_1	
	3			g_2g_1
	4			g_2g_1
	≥ 5			g_2g_1

The tables above give four pure strategies for the game $\tilde{\Gamma}_n^*$, where $5 \leq n \leq \infty$, under the assumption that $u \geq 2$. We consider the pure strategies w_1 and w_2 for the Infiltrator, and g_1 and g_2 for the Guard. The notation emphasises that both of the Infiltrator strategies are wait-and-run strategies. Observe that g_2 requires a Guard speed of at least 2, since it involves a jump from state one to state three. Let $w^* \in \mathcal{I}_n$ and $g^* \in \mathcal{G}_n$ be the mixed strategies according to which each player chooses one of his two pure strategies at random.

Lemma 4.4.1 *If $p = 4$, then for all $u \geq 2$, $0 < \lambda < 1$ and $n \geq 5$,*

$$\tilde{v}_\infty = \tilde{v}_n = \frac{\lambda^2}{2}(1+\lambda)$$

and w^* and g^* are optimal strategies.

Proof First consider the Guard's best replies if he knows that the Infiltrator is using strategy w^*. If there is a best reply there must be at least one which is pure. It is easiest to think as w_1 and w_2 as two separate Infiltrators. Then let $\bar{g} \in G_n$ denote the Guard strategy $1, 2, 3, 3, \ldots$. Under $\bar{g}$, the Guard may detect w_1 or w_2 at time 1, w_1 at time 2 and at time 3, and w_2 at time 4. Thereafter both w_1 and w_2 are safe at the target. As the Guard has precisely three chances at detecting w_1, and two at w_2, $\tilde{\Pi}_n(w_1, \bar{g}) = \lambda^3$ and $\tilde{\Pi}_n(w_2, \bar{g}) = \lambda^2$. Thus $\tilde{\pi}_n(w^*, \bar{g}) = \frac{1}{2}(\lambda^2 + \lambda^3)$.

In fact, $\bar{g}$ is a best reply to w^*. For suppose that the Guard uses a pure strategy g which meets the path of one of w_1 and w_2 on n_1 occasions, and the other on n_2 occasions. Thus $n_1, n_2 > 0$ and, as, except at time 1 the Guard can meet at most one of w_1 and w_2 at a time, $n_1 + n_2 \leq 5$. Without loss of generality we may assume that $n_1 \leq 2$. Therefore, for all $0 < \lambda < 1$,

$$\tilde{\pi}_n(w^*, g) - \tilde{\pi}_n(w^*, \bar{g}) = \frac{1}{2}\left(\lambda^{n_1} + \lambda^{n_2} - \lambda^2 - \lambda^3\right)$$

$$
\begin{aligned}
&\geq \frac{1}{2}\left(\lambda^{n_1} - \lambda^2 - \lambda^3 + \lambda^{5-n_1}\right) \\
&= \frac{1}{2}\left(\lambda^3 - \lambda^{5-n_1}\right)\left(\lambda^{n_1-3} - 1\right) \\
&\geq 0 \text{ as } n_1 \leq 2 \text{ and } \lambda < 1.
\end{aligned}
$$

Now suppose that the Guard is known to be playing strategy g^*. Note that of course both elements of g^* are best replies to w^*. It is clear that whatever route to the target the Infiltrator takes, he must cross the path of one of g_1 and g_2 at least three times, and the other at least twice. Hence, for any $i \in I_n$, $\tilde{\pi}_n(i, g^*) \leq \frac{1}{2}(\lambda^2 + \lambda^3)$. This completes the proof.

Note that the pure strategy w_1 represents an immediate dash to the target by the Infiltrator. Taken alone, this ensures the Infiltrator a payoff of only λ^3 and so is not optimal. The Infiltrator increases his chances of success by varying the timing of his dash to safety.

Finally, let us keep $p = 4$ and $u \geq 2$, and look at the detection game Γ^*_∞. By Theorem 3.4.4 we know that $v_\infty = \tilde{v}_\infty = \frac{1}{2}(\lambda^2 + \lambda^3)$. Moreover it is clear that w^*, as given above, is also optimal in the detection game Γ^*_∞. However, as the Infiltrator is not bound ever to move to the target, g^* is not necessarily optimal for the Guard. However, using the same technique as in the case $p = 3$, we can adapt g^* and obtain an ϵ-optimal strategy for Γ^*_∞. Let $\epsilon > 0$ and let γ^* be the Guard strategy which is identical to g^* for the first two time units and at each later time locates him, independently of his previous locations, at state three with probability $1 - 2\delta$ and each of states one and two with probability δ, where $0 < \delta \ll 1/2$. What is the Infiltrator's best reply to γ^*? During the first two time units, whatever he does, the Infiltrator remains undetected with probability $\frac{1}{2}(\lambda + \lambda^2)$. But assuming he is still undetected, what should he do next? As this is not a detection game he does not have to go to the target, but if he never does then he is detected with probability one. If he does go, he should do so immediately and then the final probability of avoiding detection is given by $\frac{1}{2}(\lambda + \lambda^2)\left[(1 - 2\delta)\lambda + 2\delta\right]$. By taking δ sufficiently small, this probability can be made arbitrarily close to $\frac{1}{2}\left(\lambda^2 + \lambda^3\right)$. In fact, if $\delta \leq \frac{\epsilon}{\lambda^2(1-\lambda^2)}$, then γ^* is ϵ-optimal.

We conclude this section on $p = 4$ with a comment on the game $\tilde{\Gamma}^*_\infty$ when $u < 2$. The Guard can no longer move directly between states one and three, and so has no effective speed advantage over the Infiltrator. We have no solution for this apparently simple problem. However we suggest that the Infiltrator can exploit his opponent's limitations and that, if $u < 2$, $\tilde{v}_\infty > \frac{1}{2}(\lambda^2 + \lambda^3)$.

The significance of the Guard's speed is clear. For fixed n, p and λ, both v_n and $\tilde{v}_n$ are decreasing functions of u. The faster the Guard can move the less chance the Infiltrator has of success. There are two values of u which we consider to be particularly important. When $u = 1$ the Guard and the Infiltrator have the same maximum speed, and we describe the Guard as a **slow Guard**. When $u \geq p - 2$ the Guard can move between any of the $p - 1$ states in P', and he is a **fast Guard**.

In these simple examples when $p = 4$ we have found that it is easier to find solutions when there is a fast Guard. That this is not true in general is suggested by the partial solution for finite detection games obtained in Chapter 6. The solution to be presented there is only valid when the Guard is slow.

Before leaving this discussion, let us return to the comment of Gal concerning games of infiltration, which we quoted in the introduction.

> " It can be easily seen that ... it is not a good policy for the hider to move in a straight line using his maximal velocity. A policy which does seem to be good for the hider is to move randomly for a certain period of time and only then to use his maximal velocity." ([31], page 98)

For the linear problems which we are looking at we agree that it is not a good policy for the Infiltrator simply to dash straight for the target. We suggest that, especially when $u = 1$, a policy which does seem good is for the Infiltrator to stay put for a variable period of time and only then to use his maximal velocity. In section 6.4 there is an example of a game in which $u = 1$ and yet the unique optimal strategy for the Infiltrator plays, with positive probability, a pure strategy which moves out from state one only to return there later. We suggest, however, that it is unusual to find that this method of 'losing oneself' is optimal.

Chapter 5

SAFE BASE GAMES

5.1 Introduction

We shall consider an adaptation of our original problem. First recall the extended target game $\tilde{\Gamma}_n^*$, $n < \infty$. This ought to be familiar to us by now. The game involves p states, which we call states one, two, etc. up to p. The Infiltrator starts at state one and, moving between adjacent states, has to make his way to the target state p within a time limit of n. The Guard, who can move between states which are up to $u \geq 1$ states apart, tries to prevent this by detecting the Infiltrator before he can reach the target. If the two players occupy the same state detection occurs with probability $1 - \lambda$, where $0 \leq \lambda < 1$, and if it does not then neither of them is aware that it might have occurred. The full details of the game $\tilde{\Gamma}_n^*$, $n < \infty$, are found in the previous chapters.

Now suppose that we add an extra state to the game. This is not simply any other state, but a special base state, or state zero. In the adapted game, which we shall denote by Λ_n^*, the Infiltrator starts at state zero where he is safe from detection. He chooses when to move from here and after he does so is located at state one. The game then continues as before. The total time allowed for him to reach the target, including whatever time he stays at the base, is still n. The Guard is a aware that his opponent is starting from this safe base but he receives no information when the Infiltrator moves to state one.

It was S.P.Lalley who first suggested the game Λ_n^*. Therefore we shall refer to it as the Lalley Game. In [47] he addresses the problem when the Guard's speed $u = 1$. He finds the value of the game and gives an optimal strategy for each player. In the following section we shall consider his work and suggest how it is possible to extend his conclusions.

For the rest of the chapter we give a generalisation of this safe base game. We extend the state space. Instead of a line of states joining the base to the target we consider there to be $k \geq 1$ such lines of different lengths. These lines intersect only at the base and the target, therefore we think of them as k arcs joining these states. The time limit n and the detection probability $1 - \lambda$ are as in the earlier games. We consider the Guard's mobility to be such that, at each step, he can move between any two states (except the base and the target). We refer to this as the K-Arc Game. The value of the game is found and a pair of optimal strategies given. By looking at our results when k, the number of arcs, is equal to 1 we have the solution to the Lalley Game Λ_n^* when u, the speed of the Guard is $\geq p - 1$. This can be compared with Lalley's solution for $u = 1$ given earlier in the chapter.

5.2 The Lalley Game

Lalley has proposed a game which is different from that originally posed by Gal. The ammendment he makes is to introduce a safe base state, or state zero where the Infiltrator is to be found at time 0. The Infiltrator may leave the base at any time after this and, when he does so, he moves to state one. The game then continues as before. The objective of the Infiltrator is to reach the target within the time limit n, without being detected by the Guard. Thus this is a target game rather than a detection game. The Guard cannot move to the base to detect the Infiltrator, neither is he told when he leaves the base. The other parameters λ, p and u are defined as in the original target game. A pure strategy for the Infiltrator is a function $I : \{0, \ldots, n\} \to P \cup \{0\}$ such that $I(0) = 0$ and, for all $t \geq 1$, $|I(t) - I(t-1)| \leq 1$. A pure strategy for the Guard is a function $G : \{1, \ldots, n\} \to P \backslash \{p\}$ such that, for all $t \geq 2$, $|G(t) - G(t-1)| \leq 1$. As the there are only a finite number of states, and as the time limit is finite, the number of pure strategies is finite. Each player chooses a probability vector over his set of pure strategies. These probability vectors are known as mixed strategies. Lalley himself calls this a 'One Dimensional Infiltration Game', we call it the Lalley Game and denote it by Λ_n^*.

In [47] Lalley considers only the case where the Guard's speed $u = 1$. Letting

q and r be the unique integers which satisfy

$$p - 2 = q(n - p + 1) + r, \quad \text{where } q \geq 0 \text{ and } 0 \leq r < n - p + 1, \tag{5.1}$$

he proves that $v(\Lambda_n^*) = l(n, p, \lambda)$, where

$$l(n, p, \lambda) = \lambda^{q+2}\left(\frac{r}{n - p + 1}\right) + \lambda^{q+1}\left(1 - \frac{r}{n - p + 1}\right). \tag{5.2}$$

He gives a pair of optimal mixed strategies for Λ_n^*. We shall discuss these strategies but not give a complete proof.

The Infiltrator must reach the target to win. Starting from the base he has p states to travel and hence he can stay at the base until no later than time $n - p$. The Infiltrator strategy X is defined as follows: A starting time s is chosen at random from the set $\{1, 2, \ldots, n - p + 1\}$. The Infiltrator waits at his base until time $s - 1$, and then proceeds, full speed ahead, towards the target. The path of X, $X(t)$, $0 \leq t \leq n$, is given by

$$X(t) = \begin{cases} 0 & \text{if } 0 \leq t < s, \\ t - s + 1 & \text{if } s \leq t \leq s + p - 2, \\ p & \text{if } s + p - 1 \leq t \leq n. \end{cases}$$

For each choice of s, X can clearly be called a wait-and-run strategy. Note that here, however, that the Infiltrator waits at state zero and not at state one. Note also that no two of these wait-and-run paths intersect except at the base or the target. The mixed strategy X chooses each of these paths with probability $1/(n - p + 1)$.

We shall prove that X ensures the Infiltrator a payoff of at least $l(n, p, \lambda)$. Under X, if undetected, the Infiltrator reaches the target within the time limit. Thus the Guard has $n - 1$ chances to detect the Infiltrator who is equally likely to be on any one of $n - p + 1$ non-intersecting paths. The following Lemma shows that the Guard's best reply to X is to cover the wait-and-run paths as evenly as possible. Note first that, by adding $n - p + 1$ to both sides, equation (5.1) can be rewritten as

$$n - 1 = (q + 1)(n - p + 1) + r \quad \text{where } q \geq 0 \text{ and } 0 \leq r < n - p + 1. \tag{5.3}$$

Lemma 5.2.1 *Let S be the set of all $(n - p + 1)$-tuples $(s_1, s_2, \ldots, s_{n-p+1})$ of non-negative integers such that $\sum s_i = n - 1$. Then*

$$\min_{(s_1, \ldots, s_{n-p+1}) \in S} \left(\sum_{i=1}^{n-p+1} \lambda^{s_i} \right) = r\lambda^{q+2} + (n - p + 1 - r)\lambda^{q+1}.$$

Proof If $\lambda = 0$ the sum $\sum \lambda^{s_i}$ is always zero and the result follows immediately. Now let us suppose $0 < \lambda < 1$.

Let $S' \subseteq S$ be the set of $(n-p+1)$-tuples $(s'_1, \ldots, s'_{n-p+1})$ such that there are r values of i for which $s'_i = q+2$, while for the remainder $s'_i = q+1$. For all $(s_1, \ldots, s_{n-p+1}) \in S'$, $\sum \lambda^{s_i} = r\lambda^{q+2} + (n-p+1-r)\lambda^{q+1}$. We show that S' is the set on which the minimum of $\sum \lambda^{s_i}$ is achieved.

Suppose that $(s_1, \ldots, s_{n-p+1}) \in S$ but $\notin S'$. As $\sum s_i = n-1$ there exist j, k such that $s_j - s_k \geq 2$. Now let $(s'_1, \ldots, s'_{n-p+1}) \in S$ be given by $s'_i = s_i - 1$ if $i = j$, $s_i + 1$ if $i = k$, and s_i otherwise. Comparing the power series we see that $\sum \lambda^{s_i} - \sum \lambda^{s'_i} = (1-\lambda)(\lambda^{s_k} - \lambda^{s'_j}) > 0$, since $s_k \leq s_j - 2 < s'_j$. Hence $(s_1, \ldots, s_{n-p+1})$ does not achieve the minimum of $\sum \lambda^{s_i}$. This is so for all elements of $S \backslash S'$. Hence the result follows.

Therefore, against X, a best reply for the Guard meets r of the wait-and-run paths $q+2$ times and the remainder $q+1$ times. Thus X ensures a payoff greater than or equal to

$$\frac{1}{n-p+1}\left[\lambda^{q+2}r + \lambda^{q+1}(n-p+1-r)\right],$$

which, by (5.2) is equal to $l(n,p,\lambda)$.

Recall that, for the Guard, play really starts at time 1. By this time of course the first wait-and-run strategy has already left the base. The Guard's optimal strategy is made up of best replies to X. These particular best replies are all of a type Lalley describes as 'orderly fallback' strategies. That means that the Guard falls back from state one to state $p-1$ in the following orderly manner: He uses the path of the first wait-and-run strategy for either $q+1$ or $q+2$ time units, then moves onto the second wait-and-run path, again for either $q+1$ or $q+2$ time units, and so on. Whether he spends $q+1$ or $q+2$ time units on a particular path is not just chosen at random. In total he must spend $q+2$ time units on r of the wait-and-run paths and $q+1$ on the remainder. Thus, the total duration of this strategy is $r(q+2)+(n-p+1-r)(q+1) = (q+1)(n-p+1)+r = n-1$, as required.

More rigorously, the mixed strategy Y is defined as follows: Let the random variables $\xi_0, \xi_1, \ldots, \xi_{n-p}$ be obtained by sampling without replacement from an urn containing r balls marked $q+1$ and $(n-p+1)-r$ balls marked q. The Guard must occupy site $x=1$ at time $t=1$, retreat at full speed for ξ_0 time units, wait one time unit, then retreat at full speed for ξ_1 time units, wait one time unit etc., finally retreating at full speed for ξ_{n-p} time units. In total this manoeuvre lasts precisely $1+\xi_0+1+\xi_1+\ldots+1+\xi_{n-p} = n-1$ time units. The trajectory of Y is somewhat complicated to write. For all $1 \leq t \leq n-1$, $Y(t) = t-m$, where m is the unique integer $0 \leq m \leq n-p$ such that $m+1+\sum_{i=0}^{m-1}\xi_i \leq t \leq m+1+\sum_{i=0}^{m}\xi_i$.

Lalley shows that when Y is played against any Infiltrator strategy, the payoff is at most $l(n,p,\lambda)$. We shall omit the proof of this; it may be found in [47]. Thus $l(n,w,\lambda)$ is the value, and X and Y are optimal strategies.

We note that Lalley's solutions can be extended. He states that the speed of the Guard $u=1$. However, nowhere is this required. The Infiltrator strategy X is just as effective against a Guard of any speed $u>1$. We shall return to this point at the end of the chapter.

5.3 The K-Arc Game

For the remainder of this chapter we consider in detail the K-Arc Game. This is a generalisation of the Lalley Game that we have just discussed. The K-Arc Game retains the notion of a safe base state where the Infiltrator is initially located, but the arrangement of the states is more complex. It is simplest to consider the states as vertices of a graph Γ, in which every vertex is joined by a single edge to all the other vertices to which, from there, the Infiltrator can move next. Γ is made up of two distinguished vertices which are joined by k non-intersecting paths of varying vertex length. These distinguished vertices are the base and the target and are denoted by v_A and v_B respectively. The objective of the Infiltrator is to reach the target within the time limit n, undetected by the Guard. The speed of the Guard, u, is considered to be great enough for him to move freely between the vertices. He has no chance of detecting his opponent at either the base or the target. Therefore v_A and v_B are excluded from his playing space. As in the other games we have considered, if the players simultaneously occupy the same vertex then detection occurs with probability $1-\lambda$, $0 \leq \lambda < 1$. If detection does not occur then neither player realises that it could have done. We give a rigorous analysis of this game. It is a finite game and so is guaranteed to have a solution in terms of probability vectors. The value of the game is found and compared to that of the Lalley Game. Optimal strategies are also found.

5.4 The Model

For the rest of this chapter let Γ denote a graph made up of $k \geq 1$ non-intersecting vertex paths joining the base and the target. We shall abuse the notation by also using Γ to refer to the target game which is played on this graph.

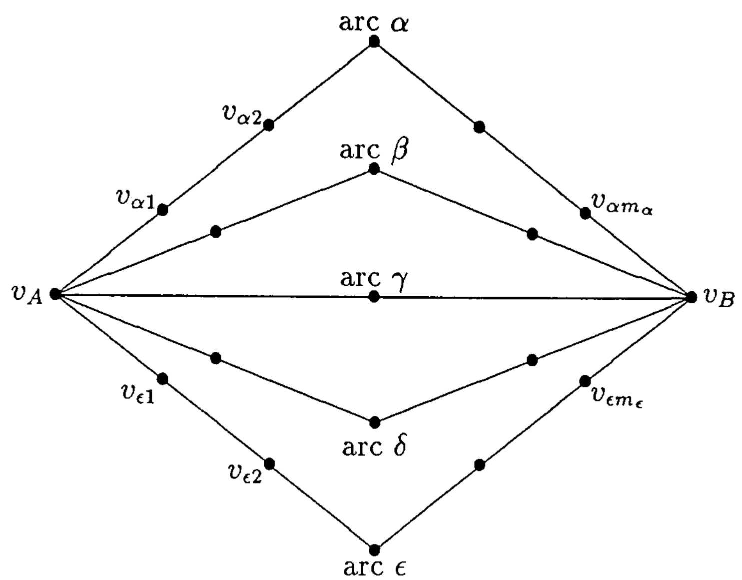

We need to introduce some more notation. Let us define

V, the set of **vertices** of Γ;

$V' = V \backslash \{v_A, v_B\}$, the set of **interior vertices;**

A, the set of discrete disjoint paths between v_A and v_B, the **arcs** of Γ, which we consider to have a fixed arbitrary order;

m_α, where $\alpha \in A$, the length of arc α, ie the number of interior vertices on arc α;

$m = \sum_{\alpha \in A} m_\alpha = |V'|$, the total number of interior vertices;

$v_{\alpha j}$, where $\alpha \in A$ and $j \in \{1, \ldots, m_\alpha\}$, the jth interior vertex on arc α, ordered from base to target;

$T_0 = \{0, 1, \ldots, n\}$, the **universal time set**, which we consider to be chronologically ordered;

$T = \{1, \ldots, n-1\}$, the **real time set.**

A pure strategy for the Infiltrator is a function I, $I : T_0 \rightarrow V$, such that $I(0) = v_A$ and, for $t = 1, 2, \ldots, n$, $I(t)$ is a vertex adjacent to or equal to $I(t-1)$. A pure strategy for the Guard is a function J, $J : T \rightarrow V'$. Note that here there is no restriction on the speed of the Guard. He can move between any interior vertices. We have restricted the codomain of J to exclude v_A and v_B since they are safe for the Infiltrator. Similarly we have restricted the domain to exclude $t = 0$ and $t = n$ since initially the Infiltrator is safe, and at the end of the game, if the Infiltrator is still at large in V', then he has failed in his objective and so, in either case, the Guard's position is immaterial. Let I and J be pure strategies for the Infiltrator and the Guard respectively. We say that I and J **meet** if there exists a $t \in T$ such that $I(t) = J(t)$. The total number of meetings is sometimes denoted by $\omega(I, J)$. The payoff $E(I, J)$ when these strategies are played is then given by $E(I, J) = \lambda^{\omega(I,J)}$.

The sets of pure strategies are finite. Each player chooses a probability distribution on his pure strategies. These probability distributions are known as mixed strategies. If I^* and J^* are mixed strategies $E(I^*, J^*)$, the payoff when I^* and J^* are played, is the expected value of $\lambda^{\omega(I,J)}$ under the joint distribution of I^* and J^*.

To avoid trivial situations we assume that every arc has at least one interior point. Further, if Γ were to contain any arcs of length greater than $n - 1$ then the Infiltrator would be unable to use these arcs to reach the target within the time limit. Therefore we make the following assumption.

A(0) For all $\alpha \in A$, $1 \leq m_\alpha \leq n - 1$.

5.5 Infiltrator Strategies

We consider only those Infiltrator pure strategies which are not dominated. Thus we assume that

A(i) The Infiltrator will never return to the base once he has left it.

A(ii) The Infiltrator will never leave the target once he has reached it.

A(iii) The Infiltrator will never move to a vertex from which he would have no chance of reaching the target within the time limit.

One result of A(i) and A(ii) is that we may restrict our attention to those pure Infiltrator strategies which use precisely one arc to reach the target. Furthermore, for any I, we can define $s(I)$ and $t(I)$, to be the **last base time** and **first target time**, respectively so that, for $0 \leq t \leq n$,

$$I(t) = \begin{cases} v_A & \text{if and only if } t \leq s(I), \\ v_B & \text{if and only if } t \geq t(I). \end{cases}$$

Between $s(I)$ and $t(I)$ the path of the Infiltrator is certain to be along a single arc α. However, just how he travels along α is uncertain. He may move straight from the base to the target, he may loiter along the way, he may even retreat at times. We shall give an optimal strategy which is made up of only pure strategies in which he does the first of these things. As in the Lalley Game we shall see that the Infiltrator waits and then dashes. In the K-Arc Game however he can vary not only the timing but also the route of his dash.

Let us introduce the set of pure strategies which we will use to construct an optimal strategy. For $\alpha \in A$, and $\tau \in \{1, \ldots, n - m_\alpha\}$, the **wait-and-run strategy**, $I_{\alpha\tau}$, is given by

$$I_{\alpha\tau}(t) = \begin{cases} v_A & 0 \leq t < \tau, \\ v_{\alpha(t-\tau+1)} & \tau \leq t < \tau + m_\alpha, \\ v_B & \tau + m_\alpha \leq t \leq n. \end{cases}$$

Under $I_{\alpha\tau}$ the Infiltrator waits at base until $t = \tau - 1$, and then proceeds full speed ahead, along arc α, towards the target, arriving at $t = \tau + m_\alpha$. Thus $s(I_{\alpha\tau}) = \tau - 1$ and $t(I_{\alpha\tau}) = \tau + m_\alpha$.

Let W denote the set of all wait-and-run strategies, and w the number of elements in W. Then $w = \sum_{\alpha \in A} n - m_\alpha = kn - m$. Note the similarity between these strategies and those given for the Lalley Game in section 2.

Later it will be useful to have an ordering $<$ on W, and so we introduce this here. Let $W = \{W_1, \ldots, W_w\}$ be the ordered set of wait-and-run strategies, satisfying the following: $W' < W''$ if and only if either

(i) $t(W') < t(W'')$, or

(ii) $t(W') = t(W'')$, and $s(W') < s(W'')$, or

(iii) $t(W') = t(W'')$, $s(W') = s(W'')$, and the arc corresponding to W' precedes the arc corresponding to W'' in the ordering of A.

For all $1 \leq j \leq w$, there exists $\alpha \in A$ and $1 \leq i \leq n - m_\alpha$ such that $W_j = I_{\alpha i}$. Suppose that $W_{j'}$ is another wait-and-run strategy on arc α, but that $W_{j'}$ leaves the base before W_j. Thus $W_{j'} = I_{\alpha i'}$, where $1 \leq i' \leq i - 1$. We deduce from this that if $W_j = I_{\alpha i}$ then $i \leq j$.

It is immediate that, for all $1 \leq j < w$, $t(W_j) \leq t(W_{j+1})$. However, if $t(W_j) < t(W_{j+1})$ then, we can also deduce that $t(W_{j+1}) = t(W_j)+1$. For otherwise there would be another wait-and-run strategy along the same arc as W_j which would come between W_j and W_{j+1} in the ordering. Thus, for all $1 \leq j < w$, $t(W_j) \leq t(W_{j+1}) \leq t(W_j) + 1$, and hence $t(W_j) \geq t(W_w) - (w - j) = n - w + j$ since the last wait-and-run path must reach the target right at the time limit n.

Finally, we denote by I^* the mixed strategy for the Infiltrator which chooses pure strategy I with probability $1/w$ if $I \in W$, and with probability zero otherwise. Let J be any pure strategy for the guard. The payoff when the Infiltrator and the Guard use strategies I^* and J respectively is given by

$$E(I^*, J) = \frac{1}{w}\sum_{j=1}^{w} \lambda^{\omega(W_j, J)}.$$

where $\omega(W_j, J)$ is the number of meetings between W_j and J.

Before giving our first result let $z = \lfloor (n-1)/w \rfloor$, where $\lfloor\ \rfloor$ denotes the floor function, and

$$f(n, w, \lambda) = \lambda^{z+1}\left(\frac{n-1}{w} - z\right) + \lambda^z\left(1 + z - \frac{n-1}{w}\right). \tag{5.4}$$

Note here that, with z as defined above,

$$n - 1 = wz + (n - 1 - wz), \quad \text{where } z \geq 0 \text{ and } 0 \leq n - 1 - wz < w. \tag{5.5}$$

Lemma 5.5.1 *For all pure Guard strategies J, $E(I^*, J) \geq f(n, w, \lambda)$.*

Proof Let $1 \leq j \leq w$. If the Infiltrator travels along path W_j, and if he is not detected along the way, then he reaches the target by time n. Thus, if the Infiltrator uses strategy I^*, the Guard has at most $n-1$ opportunities to detect him. Moreover, since the paths of any two of the wait-and-run strategies meet only at the base and the target, at each of his opportunities the Guard can meet at most one element of W.

Clearly, if J is a best reply to I^* then J will meet an element of W at each of the times $1, \ldots, n-1$. Thus, if we let $\omega_j = \omega(W_j, J)$, $\sum_{j=1}^{w} \omega_j = n - 1$ and

$E(I^*, J) = (1/w)\sum_{i=1}^{w} \lambda^{\omega_i}$. Comparing Lemma 5.2.1 we see that this expression is minimized by taking the ω_i's to be as equal as possible. Thus, as $n - 1 = wz + (n - 1 - wz)$ and $0 \leq n - 1 - wz < w$, $E(I^*, J)$ is minimized when $n - 1 - wz$ of the ω_i's are equal to $z + 1$, and the remainder equal to z. Therefore, for all J,

$$E(I^*, J) \geq \frac{1}{w}\left[(n - 1 - wz)\lambda^{z+1} + (w - n + 1 + wz)\lambda^{z}\right] = f(n, w, \lambda).$$

This $f(n, w, \lambda)$ turns out to be the value of the game. However finding an optimal strategy for the Guard is much more complicated. It is not clear from the proof above whether there exist best replies of the form mentioned. It is not immediately obvious that the Guard will be able to meet $n - 1 - wz$ of the wait-and-run paths $z + 1$ times and the remainder z times. All we have done is shown that he can do no better than this.

In the three following sections we shall show that such best replies do exist. Depending on the properties of graph Γ, we find different examples of best replies to I^* which are pure. These examples are then combined to construct optimal mixed strategies for the Guard.

5.6 Guard Strategies I

We shall now split the problem down into three cases. First we compare w, the number of wait-and-run strategies, with n, the time limit. The first case arises when $w < n$ and we shall consider this shortly. The other cases arise when $w \geq n$, and we will look at these in the following sections.

CASE A : Γ satisfies $w < n$.

Note from the definition of z that $w < n$ if and only if $z \geq 1$. Therefore, for all graphs Γ which belong to Case A, $z \geq 1$.

In both this case and the next we use the following general method: First consider the Infiltrator to be restricted to only pure strategies which are wait-and-run strategies. Find a Guard's optimal strategy in this case. Finally show that this strategy is also optimal if the Infiltrator is no longer restricted to wait-and-run strategies.

The first result here does not involve the Guard at all. However it is fundamental to the strategies he uses. It concerns the times at which the different wait-and-run strategies are to be found 'at large' on the interior vertices of Γ.

Lemma 5.6.1 *Let Γ belong to Case A and let (t_j) be a sequence of w elements of T such that $t_1 < t_2 < \cdots < t_w$. Then, for all $1 \le j \le w$, $s(W_j) < t_j < t(W_j)$.*

Proof First observe that, as (t_j) is a strictly increasing sequence of integers between 1 and $n-1$, then $1 \le t_1 \le t_2 - 1 \le t_3 - 2 \le \ldots \le t_w - (w-1) \le n-w$. Thus, for all $1 \le j \le w$, $j \le t_j \le n - w + (j-1)$.

Now consider W_j, where $1 \le j \le w$. Recalling the ordering of W we saw earlier that $W_j = I_{\alpha i}$ where $\alpha \in A$ and $1 \le i \le j$. We also observed that $t(W_j) \ge n - w + j > n - w + (j-1)$. Therefore, we have that

$$s(W_j) = s(I_{\alpha i}) = i - 1 < i \le j \le t_j \le n - w + (j-1) < t(W_j).$$

This completes the proof.

Corollary 5.6.2 *Let Γ belong to Case A and let (t_j) be a sequence of w elements of T such that $t_1 < t_2 < \cdots < t_w$. Then there exists a pure Guard strategy J such that, for all $j = 1, \ldots, w$, J meets W_j at time t_j.*

Proof This follows from the Lemma above. Construct J as follows: Let $t \in T$. If there exists $1 \le j \le w$ such that $t = t_j$ then let $J(t) = W_j(t)$. Otherwise choose $J(t)$ arbitrarily on V'.

Corollary 5.6.3 *Let Γ belong to Case A and let $\left(t_j^{(1)}\right)$, $\left(t_j^{(2)}\right)$ and $\left(t_j^{(z)}\right)$, $j = 1, \ldots, w$, be z sequences, each of w elements of T, such that $t_1^{(1)} < \cdots < t_1^{(z)} < \cdots < t_w^{(1)} < \cdots < t_w^{(z)}$. Then there exists a pure Guard strategy J which, for all $j = 1, \ldots, w$, meets W_j at times $t_j^{(1)}, t_j^{(2)}, \ldots, t_j^{(z)}$.*

Proof Construct J as follows: Let $t \in T$. If there exist $1 \le j \le w$ and $1 \le r \le z$ such that $t = t_j^{(r)}$ then let $J(t) = W_j(t)$. Otherwise choose $J(t)$ arbitrarily on V'.

We can now begin the construction of special sets of pure strategies for the Guard. These sets, analogous to the set of wait-and-run strategies for the Infiltrator, will form the basis of optimal mixed strategies for the Guard when Γ belongs to Case A. In general, we will use Corollary 5.6.3 to construct particular

Guard strategies each of which meets each W_j, $j = 1, \ldots, w$, z times as above, and so has $0 \leq (n-1) - wz < w$ opportunities for additional meetings. We refer to these as basic meetings and spare meetings respectively. Thus we construct these pure Guard from two parts, a 'basic' part and a 'spare' part. The basic part of one of these strategies meets each W_j precisely z times, while the spare part meets some of them an additional time.

For any integer q let $a(q)$ be the unique integer such that $a(q)$ is equal to q modulo w, and $0 < a(q) \leq w$. Finally, recall that, as $z = \lfloor (n-1)/w \rfloor$, then $0 \leq n - 1 - wz < w$.

Definition 5.6.4 *Let $1 \leq i \leq w$. The* **spare time set**, *$A_i \subseteq T$, is given by*

$$A_i = \{a(i), a(i+1), \ldots, a(i+n-2-wz)\}.$$

If $z = (n-1)/w$, A_i is empty by convention.

If $q_1, \ldots, q_w$ are consecutive integers then $a(q_1), \ldots, a(q_w)$ are all distinct. therefore, as $n - 1 - wz < w$, for all $1 \leq i \leq w$, $|A_i| = n - 1 - wz$. Thus A_i', the complement of A_i in T, contains precisely wz elements, which, when they are put in order, we denote by

$$a_{i1}^{(1)} < \cdots < a_{i1}^{(z)} < \cdots < a_{iw}^{(1)} < \cdots < a_{iw}^{(z)}.$$

We now define a total of w pure strategies for the Guard. The paths of these pure strategies are given according to which wait-and-run strategies they meet along the way. For all $1 \leq i \leq w$ the basic part of the strategy is defined on A_i' and the spare part on A_i. Note that if $t \in A_i'$ then $t = a_{ij}^{(k)}$ where $1 \leq j \leq w$ and $1 \leq k \leq z$. Similarly if $t \in A_i$ then $1 \leq t \leq w$. Recall also that, for $1 \leq t \leq w$, $s(W_t) < t < t(W_t)$, and so, at time t, the wait-and-run strategy W_t is somewhere 'at large' among the interior vertices. Therefore from Corollary 5.6.3 we deduce the existence of the following pure strategies.

Definition 5.6.5 *Let $1 \leq i \leq w$. The pure Guard strategy J_i is defined as follows:*

$$J_i(t) = \begin{cases} W_j(t) & \text{if } t = a_{ij}^{(k)} \in A_i', \\ W_t(t) & \text{otherwise.} \end{cases}$$

Let $1 \leq i \leq w$. We refer to $J_i|A_i'$, the restriction of J_i to domain A_i', as the **basic part** of J_i and to $J_i|A_i$ as the **spare part**. Finally, recall our earlier definition of the mixed strategy I^*. Now let J^* denote the mixed strategy for the Guard which chooses the pure strategy J with probability $1/w$ if $J = J_1, \ldots, J_w$, and with probability zero otherwise.

Suppose that I is any pure strategy for the Infiltrator. Recalling that $\omega(I, J)$ denotes the number of meetings between I and J, the payoff when the Infiltrator and the Guard use strategies I and J^* respectively is $E(I, J^*) = \frac{1}{w}\sum_{i=1}^{w} \lambda^{\omega(I,J_i)}$.

We now consider the payoff when one of these Guard strategies $J_1, \ldots, J_w$ is played against one of the wait-and-run strategies $W_1, \ldots, W_w$. For all $1 \leq i, j \leq w$ we calculate $\omega(W_j, J_i)$ by considering separately how many times W_j meets the two different parts of J_i.

Lemma 5.6.6 *Let Γ belong to Case A and let $1 \leq j \leq w$. The wait-and-run strategy W_j meets (i) the basic part of each of the J_i's exactly z times, and (ii) the spare part of exactly $n - 1 - wz$ of the J_i's exactly once, and the spare part of the rest not at all.*

Proof (i) Let $1 \leq i, j \leq w$. The basic part of J_i meets W_j if and only if, for some $t \in A_i'$, $J_i(t) = W_j(t)$. By Definition 5.6.5, this occurs at the z times $a_{ij}^{(1)}, a_{ij}^{(2)}, \ldots, a_{ij}^{(z)}$.

(ii) Let $1 \leq i, j \leq w$. Then, by Definition 5.6.5, W_j meets the spare part of J_i if and only if j is in the spare time set of J_i, that is if

$$j \in A_i = \{a(i), \ldots, a(i + [n - 2 - wz])\}.$$

To deduce for which i this is so, first note that if $n - 1 - wz = 0$ then each A_i is empty and so trivially j belongs to none of them. If, on the other hand, $n - 1 - wz > 0$, then by definition of A_i, we see that $j \in A_i$ if and only if

$$i \in \{a(j - [n - 2 - wz]), \ldots, a(j - 1), a(j)\}.$$

Since $j - [n - 2 - wz], \ldots, j$ are $n - 1 - wz$ consecutive integers and as $n - 1 - wz < w$, it follows that $a(j - [n - 2 - wz]), \ldots, a(j)$ are all distinct. Therefore j appears in A_i precisely once for exactly $n - 1 - wz$ values of i, and not at all for the remaining values of i.

Corollary 5.6.7 *Let $1 \leq j \leq w$, and consider the wait-and-run strategy W_j.*

$$|\{i \mid \omega(W_j, J_i) = \omega\}| = \begin{cases} w - n + 1 + wz & \text{if } \omega = z, \\ n - 1 - wz & \text{if } \omega = z + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Proof Follows immediately from the Lemma above.

Therefore, for all $1 \leq j \leq w$,

$$E(W_j, J^*) = \frac{1}{w}\left[(n-1-wz)\lambda^{z+1} + (w-n+1+wz)\lambda^z\right] = f(n,w,\lambda).$$

We have already proved that, for all pure strategies J for the Guard, $E(I^*, J) \geq f(n,w,\lambda)$. Hence, if the Infiltrator is restricted to only wait-and-run strategies, J^* is an optimal strategy for the Guard and the value of this restricted game is $f(n,w,\lambda)$.

We now need to broaden our sights again to consider the behaviour of the J_i's when they are played against any pure Infiltrator strategy I, not necessarily of the wait-and-run type. In order to do this we look again at the general form of I. We can describe the path of I according to which wait-and-run paths it meets.

First however recall the assumptions A(i), A(ii) and A(iii) which we assume to apply to any pure strategy I. We deduce that, for all $0 \leq t \leq n$ there exists $g(t)$, $1 \leq g(t) \leq w$, such that $I(t) = W_{g(t)}(t)$. Moreover, as distinct wait-and-run strategies coincide only at the base and the target, then for all $s(I) < t < t(I)$, $g(t)$ is unique. We now extend g by determining that, for all $0 \leq t \leq s(I)$, $g(t) = g(s(I)+1)$, while for all $t(I) \leq t \leq n$, $g(t) = g(t(I)-1)$. Thus, for all $t \in T_0$, at time t, I is on the path of the wait-and-run strategy $W_{g(t)}$. This function g provides the link from the set of all pure Infiltrator strategies back into the smaller set of wait-and-run strategies.

It is clear that, while he is strictly between the base and the target, the Infiltrator can move from the path of one wait-and-run strategy to another. However there is an important restriction on this due to the limitation on the Infiltrator's speed. While he is strictly between the base and the target, he is unable, except by moving to the target, to move onto a wait-and-run path that left the base before the one he is currently on. This idea is developed below in terms of the function g and the ordering $<$ on W.

Lemma 5.6.8 *Let I be a pure strategy for the Infiltrator. Then the function g is monotonic increasing.*

Proof This holds trivially for t in the domains $[1, s(I)+1]$ and $[t(I)-1, n]$, as g is constant in these intervals.

Consider now the position of the Infiltrator at some time $t \in [s(I)+1, t(I)-1)$, and the three choices of move he may have. Firstly he may advance to the next vertex and, in so doing, stay on the wait-and-run path he is already on. In this case $g(t+1) = g(t)$. Alternatively, providing that to do so is consistent with assumptions A(i) and A(iii), he may either loiter where he is, or retreat, and in both of these cases, by failing to advance, he will move himself off the wait-and-run path $W_{g(t)}$ that he is currently on, and onto a new path. Thus $g(t+1) \neq g(t)$. As this new path reaches the target strictly later than $W_{g(t)}$, it follows from the ordering $<$, that $g(t+1) > g(t)$. Thus, for all $1 \leq t < n$, $g(t+1) \geq g(t)$.

We can now extend Lemma 5.6.6 by showing the effectiveness of the strategies $J_1, \ldots, J_w$ against any pure Infiltrator strategy.

Lemma 5.6.9 *Let Γ belong to Case A. Let I be any pure Infiltrator strategy. I meets (i) the basic part of each J_i at least z times, and (ii) the spare parts of at least $n-1-wz$ of the J_i's at least once.*

Proof (i) Let $1 \leq i \leq w$. By Lemma 5.6.6(i), the basic part of J_i meets every W_j precisely z times. For $1 \leq j \leq w$, let s_j and t_j denote respectively the times of the first and last meeting beteen the basic part of J_i and the wait-and-run strategy W_j. By Definition 5.6.5, $s_j = a_{ij}^{(1)}$ and $t_j = a_{ij}^{(z)}$, and $s_1 \leq t_1 < s_2 \leq t_2 < \cdots < s_w \leq t_w$. Note that, for all j, $t_j \geq jz \geq j$.

Now let I be any pure Infiltrator strategy. At time $t(I)-1$, the path of I has yet to reach the target but does so on the next move. Therefore, for some $\alpha \in A$,

$$I\left(t(I)-1\right) = v_{\alpha m_\alpha} = W_{g(t(I)-1)}\left(t(I)-1\right).$$

As the basic part of J_i and $W_{g(t(I)-1)}$ meet z times up to and including the time that $W_{g(t(I)-1)}$ reaches $v_{\alpha m_\alpha}$, it follows that $t_{g(t(I)-1)} \leq t(I)-1$. Thus, if we let $\theta = \min\left\{x \in T_0 | t_{g(x)} \leq x\right\}$, then $0 \leq \theta \leq t(I)-1$.

Thus, at time θ, I is following the wait-and-run strategy $W_{g(\theta)}$ and, by this time, the basic part of J_i and $W_{g(\theta)}$ have already met z times. Now let ϕ, $0 \leq \phi \leq \theta$, denote the time at which I first joins the path of $W_{g(\theta)}$. There are two possibilities. If $\phi > 0$, then, by Lemma 5.6.8, $g(\phi-1) < g(\phi)$, and so, as $\phi - 1 < \theta$ it follows that $\phi - 1 < t_{g(\phi-1)} < s_{g(\phi)} = s_{g(\theta)}$. If $\phi = 0$ then trivially $\phi \leq s_g(\theta)$. So, in either case $\phi \leq s_{g(\theta)}$, and we deduce that I joins the path of $W_{g(\theta)}$ before the basic part of J_i and $W_{g(\theta)}$ meet for the first time. Hence it is certain that, throughout the interval $[\phi, \theta]$, I is on the path of $W_{g(\theta)}$ and hence meets the basic part of J_i exactly z times.

(ii) Let $1 \leq j \leq w$. By Lemma 5.6.6(ii), the wait-and-run strategy W_j is met exactly once by the spare parts of exactly $n-1-wz$ of the J_i's. Moreover, by Definition 5.6.5, these meetings all take place at time j.

Now consider the pure Infiltrator strategy I, and note that as $g(n-1) \leq w \leq n-1$, we can define ψ to be the least $t \in T_0$ such that $g(t) \leq t$. Thus, at time ψ, I is on the path of $W_{g(\psi)}$ and since $g(\psi) \leq \psi$, by this time $W_{g(\psi)}$ has already met the spare parts of $n-1-wz$ of the J_i's. Let ξ, $0 \leq \xi \leq \psi$, denote the time at which I first joined $W_{g(\psi)}$. Again there are two possibilities. If $\xi > 0$, then $g(\xi - 1) < g(\xi)$ and since $\xi - 1 < \psi$, $\xi - 1 < g(\xi - 1) < g(\xi) = g(\psi)$. If $\xi = 0$, then trivially $\xi \leq g(\psi)$. So, in either case, as $\xi \leq g(\psi)$, we deduce that I joined the path of $W_{g(\psi)}$ before $W_{g(\psi)}$ had met the spare parts of the $n-1-wz$ of the J_i's. Hence in the interval $[\xi, \psi]$ I must meet the spare parts of all $n-1-wz$ of the J_i's which $W_{g(\psi)}$ meets. The result then follows.

Corollary 5.6.10 *Let I be any pure Infiltrator strategy.*

$$|\{i \,|\, \omega(I, J_i) \geq \omega\}| \geq \begin{cases} w & if\, \omega = z, \\ n-1-wz & if\, \omega = z+1. \end{cases}$$

Proof Follows immediately from the Lemma above when the basic parts and the spare parts of $J_1, \ldots, J_w$ are reassembled.

Therefore, for all pure Infiltrator strategies I,

$$E(I, J^*) \leq \frac{1}{w}\left[(n-1-wz)\lambda^{z+1} + (w-n+1+wz)\lambda^z\right] = f(n, w, \lambda).$$

This, together with Lemma 5.5.1 allows us to state without further proof the conclusion to Case A.

Lemma 5.6.11 *Let Γ belong to Case A. The strategies I^* and J^* are optimal for the Infiltrator and the Guard and $v(\Gamma) = f(n, w, \lambda)$.*

5.7 Guard Strategies II

We turn our attention to what happens when, w, the number of wait-and-run strategies is greater than or equal to the time limit n. In a moment we shall proceed in a manner similar to that of Case A and construct another set of w pure strategies for the Guard analogous to $J_1, \ldots, J_w$. This approach depends on an application of Lemma 5.6.1 to a subgraph of Γ and hence puts another condition on Γ. This situation is covered in the current section under the title Case B. The final possibility is considered in the following section and this is Case C.

When k, the number of arcs, is equal to one, then $w = n - m_\alpha < n$, and hence Γ belongs to Case A. Therefore, if Γ does not belong to Case A we may assume that $k \geq 2$. In both of the following cases we shall obtain the value of Γ by deleting one or more arcs from Γ and studying the resulting subgraphs.

Suppose $k \geq 2$. If $\alpha \in A$ we denote by Γ_α the subgraph obtained by removing all the interior vertices $v_{\alpha 1}, \ldots, v_{\alpha m_\alpha}$ from graph Γ. Recalling our earlier convention, let Γ_α also denote the game which is played on this subgraph when the parameters λ and n are left unchanged. Observe also that, if Γ satisfies assumption A(0), then so too will Γ_α. Moreover, any pure strategy in Γ_α will also be a pure strategy in Γ.

Let $\alpha \in A$. We shall partition W, the set of wait-and-run strategies, according to whether its elements use arc α or not. Denote by $\Omega = \{\Omega_1, \Omega_2, \ldots, \Omega_{n-m_\alpha}\}$ the set of wait-and-run strategies which do use arc α to reach the target. We assume that Ω is ordered by $<$, and hence, for all $1 \leq i \leq n - m_\alpha$, $\Omega_i = I_{\alpha i}$. Similarly, if $w_\alpha = w - (n - m_\alpha)$ is the number of wait-and-run strategies which remain, and so do not use arc α, denote this set by $\Theta = \{\Theta_1, \Theta_2, \ldots, \Theta_{w_\alpha}\}$. Again we assume that Θ is ordered by $<$. Θ is the set of wait-and-run strategies for the game Γ_α.

CASE B :

(i) Γ **satisfies** $w \geq n$.

(ii) There exists $\alpha \in A$ **such that** $w_\alpha < n$.

For all $\alpha, \beta \in A$, if $m_\alpha \leq m_\beta$ then $w_\alpha \leq w_\beta$. Thus, without loss of generality we may assume that α is an arc of minimum vertex length.

As in Case A we start by supposing that the Infiltrator is restricted to only wait-and-run strategies. We are concerned with the number of times that the Guard can meet the different elements of W. Since, in this case, $w \geq n$ he may be unable to meet them all. We can no longer apply Lemma 5.6.1 or its corollaries directly to the graph Γ. However, thanks to condition (ii), there is a subgraph Γ_α to which Lemma 5.6.1 may be applied.

Hence suppose that $S' \in T$ contains the w_α elements, $s'_1, s'_2, \ldots, s'_{w_\alpha}$. Then there exists a pure Guard strategy J, for the game Γ_α, such that, for all $1 \leq j \leq w_\alpha$, $J(s'_j) = \Theta_j(s'_j)$. Moreover, of course, J is also a pure strategy for the full game Γ. In this way we proceed to construct pure Guard strategies in the full game which are composed of basic parts which meet each element of Θ, and spare parts which meet some of the elements of Ω as well.

For any integer q let $b(q)$ be the unique integer such that $b(q)$ is equal to q mod $n - m_\alpha$ and $1 \leq b(q) \leq n - m_\alpha$.

Definition 5.7.1 *For $1 \leq i \leq n - m_\alpha$, we define the* **spare time set**, *$B_i \subseteq T$, as*

$$B_i = \left\{b(i), b(i+1), \ldots, b(i+n-2-w_\alpha)\right\}.$$

If $w_\alpha = n - 1$, then by convention, for all i, B_i is empty.

If $q_1, \ldots, q_{n-m_\alpha}$ are consecutive integers, then $b(q_1), \ldots, b(q_{n-m_\alpha})$ are all distinct. Therefore, as $n - 1 - w_\alpha < w - w_\alpha = n - m_\alpha$, for all $1 \leq i \leq n - m_\alpha$, $|B_i| = n - 1 - w_\alpha$. Thus B'_i, the complement of B_i in T, has w_α elements, which we denote by

$$0 < b_{i1} < \cdots < b_{iw_\alpha} < n.$$

Let us now give a defintion which is analogous to that of the pure strategies $J_1, \ldots, J_w$ in the previous section.

Definition 5.7.2 *Let $1 \leq i \leq n - m_\alpha$. The pure Guard strategy K_i is defined as follows:*

$$K_i(t) = \begin{cases} \Theta_j(t) & \text{if } t = b_{ij} \in B'_i, \\ \Omega_t(t) & \text{otherwise.} \end{cases}$$

Let $1 \leq i \leq n - m_\alpha$. We refer to $K_i|B'_i$ and $K_i|B_i$ as the basic and spare parts of K_i respectively.

We now consider the payoff when one of the Infiltrator wait-and-run strategies $W_1, \ldots, W_w$ is played against one of these Guard strategies $K_1, \ldots, K_{n-m_\alpha}$. As in Case A we calculate $\omega(W_j, K_i)$ by considering separately the basic part and the spare part of K_i.

Lemma 5.7.3 *Let $1 \leq j \leq w$, and consider the wait-and-run strategy W_j.*

(i) *If $W_j \in \Theta$, then W_j meets the basic part of each of the K_i's exactly once, and the spare part of none of them.*

(ii) *If $W_j \in \Omega$, then W_j meets the basic part of none of the K_i's, and the spare part of exactly $n-1-w_\alpha$ of them precisely once.*

Proof (i) Let $W_j = \Theta_x$, $1 \leq x \leq w_\alpha$. Θ_x does not meet the spare part of any of the K_i's simply because, for $t \in B_i$, by Definition 5.7.2, $K_i(t)$ is defined only on the elements of $\Omega = W \backslash \Theta$. For fixed $1 \leq i \leq n - m_\alpha$, now consider the basic part of K_i. The basic part of K_i meets Θ_x if and only if there exists $t \in B_i'$ such that $K_i(t) = \Theta_x(t)$. By Definition 5.7.2, this happens at precisely $t = b_{ix}$, and at no other value of t.

(ii) Let $W_j = \Omega_y$, $1 \leq y \leq n - m_\alpha$. This time Ω_y does not meet the *basic* part of any of the K_i's because, by Definition 5.7.2, for $t \in B_i'$, $K_i(t)$ is defined only on the elements of $\Theta = W \backslash \Omega$. Now consider for which $i \in \{1, \ldots, n - m_\alpha\}$, Ω_y meets the spare part of K_i. By Definition 5.7.2, the necessary and sufficient condition for this is that $y \in B_i$, that is

$$y \in \{b(i), b(i+1), \ldots, b(i+n-2-w_\alpha)\}.$$

To deduce for which i this holds, first note that if $n-1-w_\alpha = 0$, then each B_i is empty, and so trivially j belongs to none of them. If, on the other hand, $n-1-w_\alpha > 0$, then by definition of the B_i's, it can be seen that $j \in B_i$ if and only if

$$i \in \{b(j-[n-2-w_\alpha]), \ldots, b(j-1), b(j)\}.$$

Since $j-[n-2-w_\alpha], \ldots, j$ are $n-1-w_\alpha$ consecutive integers, and $n-1-w_\alpha = n-1-w+(n-m_\alpha) \leq n-m_\alpha$, it follows that $b(j-[n-2-w_\alpha]), \ldots, b(j)$ are all distinct. Therefore j appears in B_i precisely once for exactly $n-1-w_\alpha$ values of i, and not at all for the remaining $n-m_\alpha-(n-1-w_\alpha) = w-n+1$ values of i.

Having introduced these $n-m_\alpha$ K_i's, we now produce another w_α pure strategies for the Guard which brings the total to w, and allows us to construct the mixed strategy M^*, which plays them all with equal probability. This time we reverse the order of things. These Guard strategies are guaranteed, in their basic parts to meet all of the elements of Ω, and in their spare parts to meet some of the elements of Θ.

For any integer q let $c(q)$ denote the unique integer such that $c(q)$ is equal to $q \bmod w_\alpha$ and $1 \leq c(q) \leq w_\alpha$.

Definition 5.7.4 *For $1 \leq i \leq w_\alpha$, we define the* **spare time set**, *$C_i \subseteq T$, as*

$$C_i = \{c(i), c(i+1), \ldots, c(i+m_\alpha-2)\},$$

If $m_\alpha = 1$, then by convention, for all i, C_i is empty.

If $q_1, \ldots, q_{w_\alpha}$ are consecutive integers then $c(q_1), \ldots, c(q_{w_\alpha})$ are all distinct. Therefore, as $m_\alpha - 1 < w_\alpha$, $|C_i| = m_\alpha - 1$. Thus C_i', the complement of C_i in T, contains $n - m_\alpha$ elements, which we denote by

$$0 < c_{i1} < \ldots < c_{i(n-m_\alpha)} < n.$$

Definition 5.7.5 *Let $1 \leq i \leq w_\alpha$. We define the pure strategy for the Guard, L_i, as follows:*

$$L_i(t) = \begin{cases} \Omega_j(t) & \textit{if } t = c_{ij} \in C_i', \\ \Theta_t(t) & \textit{otherwise.} \end{cases}$$

Let $1 \leq i \leq w_\alpha$. We refer to $L_i|C_i'$ and $L_i|C_i$ as the basic part and spare part of L_i respectively.

Lemma 5.7.6 *Let $1 \leq j \leq w$, and consider the wait-and-run strategy W_j.*

(i) *If $W_j \in \Theta$, then W_j meets the basic part of none of the L_i's, and the spare part of exactly $m_\alpha - 1$ of them precisely once.*

(ii) *If $W_j \in \Omega$, then W_j meets the basic part of each of the L_i's, and the spare part of none.*

Proof This proof is almost identical to that of Lemma 5.7.3, and hence we omit it.

From Definitions 5.7.2 and 5.7.5 we have a total of w different pure Guard strategies, namely $K_1, \ldots, K_{n-m_\alpha}$ and $L_1, \ldots L_{w_\alpha}$. Let us arbitrarily rename them $M_1, \ldots, M_w$. Then M^* is the mixed strategy for the Guard which chooses pure strategy J with probability $1/w$ if $J = M_1, \ldots, M_w$, and with probability zero otherwise. Combining Lemma 5.7.3 and Lemma 5.7.6, we have the following result.

Lemma 5.7.7 *Let $1 \leq j \leq w$, then consider the wait-and-run strategy W_j.*

$$|\{i \,|\, \omega(W_j, M_i) = \omega\}| = \begin{cases} w - n + 1 & \textit{if } \omega = 0, \\ n - 1 & \textit{if } \omega = 1, \\ 0 & \textit{otherwise.} \end{cases}$$

Proof Determine first whether W_j belongs to Θ or to Ω. If $W_j \in \Theta$, then by Lemma 5.7.3(i) and Lemma 5.7.6(i), W_j meets the basic parts of every one of the $n - m_\alpha$ K_i's precisely once, and the spare parts of exactly $m_\alpha - 1$ of the L_i's precisely once, but nothing else. Since $(n - m_\alpha) + (m_\alpha - 1) = n - 1$, the result follows.

Similarly, if $W_j \in \Omega$, then by Lemma 5.7.3(ii) and Lemma 5.7.6(ii), W_j meets the spare parts of exactly $n - 1 - w_\alpha$ of the K_i's precisely once, and the basic parts of every one of the w_α L_i's precisely once, but nothing else. Again, as $(n - 1 - w_\alpha) + w_\alpha = n - 1$, this completes the proof.

Suppose Γ belongs to Case B. It follows from above that, for all $1 \leq j \leq w$,

$$E(W_j, M^*) = \frac{1}{w}\left[(n-1)\lambda + (w - n + 1)\right].$$

Recall from Lemma 5.5.1 that, for all J, $E(I^*, J) \geq f(n, w, \lambda)$. However, from equation (5.4), as $w \geq n$ then $z = 0$, and so $f(n, w, \lambda) = \lambda\left(\frac{n-1}{w}\right) + \left(1 - \frac{n-1}{w}\right)$. Hence, if the Infiltrator is restricted to wait-and-run strategies, I^* and M^* are optimal and the value of the game is $f(n, w, \lambda)$.

Finally it remains to see what happens when we take away this restriction on the Infiltrator.

Lemma 5.7.8 *Let I be any pure strategy for the Infiltrator. Then*

$$|\{i \,|\, \omega(I, M_i) \geq 1\}| \geq n - 1.$$

Proof Relate I to the set W of wait-and-run strategies using the function g which we defined in the previous section. Consider whether the path of I uses arc α or a different arc. In either case the result follows from Lemma 5.7.7 in the same way that, in the previous section, Lemma 5.6.9 and Corollary 5.6.10 follow from Lemma 5.6.6 and Corollary 5.6.7. To avoid twice repeating, almost identically, the proof of Lemma 5.6.9. we omit the details.

Therefore, for any I, $E(I, M^*) \leq (1/w)\left[(n-1)\lambda + (w - n + 1)\right]$. This brings us to the conclusion of Case B.

Lemma 5.7.9 *Let Γ belong to Case B. Then I^* and M^* are optimal strategies for the Infiltrator and the Guard, and $v(\Gamma) = f(n, w, \lambda)$.*

Proof If Γ belongs to Case B, then $z = 0$ and so, from equation (5.4), $f(n, w, \lambda) = (1/w)\,[(n-1)\lambda + (w-n+1)]$. Now let (I, J) be any pair of pure strategies. From Lemma 5.5.1 and above, it follows that

$$E(I, M^*) \leq (1/w)\,[(n-1)\lambda + (w-n+1)] \leq E(I^*, J).$$

This completes the proof.

5.8 Guard Strategies III

In the previous section we started to look at those games in which w, the number of wait-and-run strategies, is greater than or equal to n, the time limit. In Case B we examined what happens when, for at least one of the subgraphs Γ_α, $w_\alpha < n$, where $w_\alpha = w - (n - m_\alpha)$. We now consider the final possibility.

CASE C :

(i) Γ satisfies $w \geq n$.

(ii) For all $\alpha \in A$, $w_\alpha \geq n$

The consequence of (i) and (ii) is that we are unable to use Lemma 5.5.1 to directly study the behaviour of either Γ itself, or any of the games Γ_α. Hence we must adopt another approach, but one which will allow us to build upon what we already know from the previous sections. We first recall that if k, the number of arcs in Γ, is 1, then $w(= n - m) < n$, and so Γ belongs to Case A. Similarly, if $k = 2$, and Γ consists of the two arcs α and β, either again $w < n$, or $w \geq n$ and $w_\alpha(= n - m_\beta) < n$, and so Γ belongs to either Case A or Case B. Hence, if Γ belongs to Case C we may assume that $k \geq 3$.

It transpires that the value of a graph belonging to Case C is the same as for one from Case A or B. We shall use induction on the number of arcs to prove this, and to determine the form of an optimal strategy N^* for the Guard. I^*, the uniform mixed strategy on W, the set of Wait-and- Run strategies, again turns out to be optimal for the Infiltrator. Unlike J^* and M^* from Cases A and B, which are both composed of precisely w components and which are readily found in any particular game, N^* will not be given explicitly. The form of N^* is constructed through the induction as the graph Γ is built up arc by arc. The reader may, by repeating the procedure given in the inductive argument, construct N^* for any game belonging to Case C. However the intricacy of this procedure sharply contrasts with the ease in which J^* or M^* may be constructed for a graph belonging to either of the other two cases. It is conjectured that an optimal strategy for the Guard of a more direct form, analogous to J^* and M^*, does exist. If this is so, it has eluded our attempts at detection.

Let us now consider the inductive hypothesis.

Hypothesis $\mathbf{P}_k$, $k \in \mathbb{N}$: Let Γ be a graph on k arcs which belongs to Case C. Then there exist $\mu(\Gamma) \in \mathbb{N}$ such that $\mu(\Gamma)|(k-1)!$, and $w\mu(\Gamma)$ pure Guard strategies $N_1, N_2 \ldots, N_{w\mu(\Gamma)}$, which are not neccessarily all distinct, such that, for any Infiltrator pure strategy I,

$$|\{i \mid \omega(I, N_i) \geq 1\}| \geq \mu(\Gamma)(n-1).$$

In other words, inclucing multiplicities if they are not all distinct, the Infiltrator must meet at least $(n-1)\mu(\Gamma)$ of these $w\mu(\Gamma)$ Guard strategies. If, for all $k \in \mathbb{N}$, P_k holds then, letting N^* denote the mixed strategy which chooses at random between the N_i's, we can show that, for all I, $E(I, N^*) \leq f(n, w, \lambda)$. We shall return to this once we have verified the hypothesis. The fact that $\mu(\Gamma)$ divides $(k-1)!$ is not, in fact, essential to the induction that follows. It is included simply to give the reader some idea of an upper bound on the number of pure Guard strategies required.

We consider first a graph Γ which belongs to Case C, and for which $k = 3$. Label the three arcs of Γ as α, β and γ. We delete arc α and consider the game Γ_α. Note that Γ_α must satisfy the conditions for Case B in section 7. Hence there are w_α pure Guard strategies, say $M_1^\alpha, \ldots, M_{w_\alpha}^\alpha$, which make up an optimal strategy in Γ_α. By applying Lemma 5.7.8 to Γ_α, any pure Infiltrator strategy in Γ_α meets at least $n-1$ of these M_i^α's.

We repeat this procedure with the games Γ_β and Γ_γ. Thus we deduce the existence of the pure Guard strategies $M_1^\beta, \ldots, M_{w_\beta}^\beta$ and $M_1^\gamma, \ldots, M_{w_\gamma}^\gamma$, which make up his optimal strategies in the games Γ_β and Γ_γ respectively. Therefore there exist a total of $w_\alpha + w_\beta + w_\gamma = 2w$ pure Guard strategies, say $N_1, \ldots, N_{2w}$, not necessarily distinct, which may be played in Γ. We show that these strategies satisfy the conditions of the hypothesis P_k when $k = 3$.

First note that we have $\mu(\Gamma) = 2$, and so, since $k = 3$, $\mu(\Gamma)|(k-1)!$. Now we must check that the other condition of the inductive hypothesis is satisfied.

If we let I be any pure Infiltrator strategy, then, relabelling the arcs if necessary, we may assume that I uses arc α. Then, by the same argument as used above, we deduce that I meets at least $n-1$ of the M_i^β's and at least $n-1$ of the M_i^γ's, giving a total of at least $2(n-1)$ of the N_i's, thus ensuring the inductive hypothesis is satisfied. Therefore, for $k=3$, P_k does hold and $\mu(\Gamma)=2$.

Now assume that, for some $K \geq 3$, P_K holds. Then consider the proposition P_{K+1}. Therefore we we take a game Γ, belonging to Case C, and which is played on $K+1$ arcs. Let α be an arc of Γ, and look at the game Γ_α. This is a game which is played on K arcs. Since $w_\alpha \geq n$, Γ_α belongs to either Case B or to Case C.

If Γ_α belongs to Case B, then there exist the w_α pure Guard strategies

$$M_1^\alpha, \ldots, M_{w_\alpha}^\alpha,$$

given by applying Definition 5.7.2 and Definition 5.7.5 to the game Γ_α, such that, by Lemma 5.7.8, every pure Infiltrator strategy in Γ_α meets at least $n-1$ of the M_i^α's. If, on the other hand, Γ_α belongs to Case C, then since Γ_α has K arcs and as we are assuming that P_k holds for $k=K$, we may immediately apply the inductive hypothesis to Γ_α.

Hence, it follows that, in either case, there are $w_\alpha \mu(\Gamma_\alpha)$ not neccessarily distinct pure Guard strategies for Γ_α,

$$N_1^\alpha, \ldots, N_{\mu(\Gamma_\alpha) w_\alpha}^\alpha,$$

where $\mu(\Gamma_\alpha) | (K-1)!$, and such that any pure Infiltrator strategy in Γ_α meets, counting multiplicities, at least $\mu(\Gamma_\alpha)(n-1)$ of the N_i^α's.

Note that these N_i^α's, originally defined as pure Guard strategies in the game Γ_α, may also be used in the full game Γ. Hence, corresponding to each arc $\alpha \in A$, we have $w_\alpha \mu(\Gamma_\alpha)$ pure Guard strategies for Γ, where, for each $\alpha \in A$, $\mu(\Gamma_\alpha)$ divides $(K-1)!$.

Now let ν denote the lowest common multiple of the $\mu(\Gamma_\alpha)$'s as α varies in A. Note that ν must also divide $(K-1)!$. If, for all $\alpha \in A$, we count all of the N_i^α's exactly $\nu / \mu(\Gamma_\alpha)$ times, then altogether we have $w\mu(\Gamma)$ not necessarily distinct pure Guard strategies for Γ, say

$$N_1, \ldots, N_{w\mu(\Gamma)},$$

where

$$w\mu(\Gamma) = \sum_{\alpha \in A} w_\alpha \left(\frac{\nu}{\mu(\Gamma_\alpha)} \right) \mu(\Gamma_\alpha) = \nu \sum_{\alpha \in A} w_\alpha = w\nu K,$$

and so $\mu(\Gamma) = \nu K$, and hence $\mu(\Gamma)$ divides $K!$.

We check that these N_i's satisfy P_{K+1}. Let I be any pure strategy for the Infiltrator. Recall that I uses precisely one arc to cross from base to target, and denote this arc by α_I. Note that I, a pure strategy in the game Γ is also a pure strategy in each of the games Γ_α, providing $\alpha \neq \alpha_I$, and so we deduce from above that, counting multiplicities, I meets at least $\sum_{\alpha \in A \setminus \{\alpha_I\}} (\nu/\mu(\Gamma_\alpha))\mu(\Gamma_\alpha)(n-1) = \nu K(n-1)$ of the N_i's.

Thus, we have shown that if, for some $K \geq 3$, P_K holds, then P_{K+1} also holds. Therefore, since we have proved that P_3 holds, it follows, by induction that P_k holds for all $k \geq 3$.

Let N^* be the mixed strategy for the Guard which choses the pure strategy J with probability $1/w\mu(\Gamma)$ if $J = N_1, \ldots, N_{w\mu(\Gamma)}$, and with probability zero otherwise. It then follows that, for all I,

$$\begin{aligned} E(I, N^*) &\leq \frac{1}{w\mu(\Gamma)}[(n-1)\mu(\Gamma)\lambda + (w-n+1)\mu(\Gamma)] \\ &= \frac{1}{w}[(n-1)\lambda + (w-n+1)] \end{aligned}$$

We can now conclude Case C.

Lemma 5.8.1 *Let Γ belong to Case C. Then I^* and N^* are optimal strategies for the Infiltrator and the Guard, and $v(\Gamma) = f(n, w, \lambda)$.*

Proof If Γ belongs to Case C then $z = 0$. Hence, from equation (5.4), $f(n, w, \lambda) = (1/w)[(n-1)\lambda + (w-n+1)]$. Now let (I, J) be any pair of pure strategies. Then, from Lemma 5.5.1 and above $E(I, N^*) \leq f(n, w, \lambda) \leq E(I^*, J)$. This completes the proof.

NOTE: The reader may have observed that this induction need not have started at $k = 3$. It has already been noted that, for $k = 1$ or 2, the set of games belonging to Case C is in fact empty, and so P_k is satisfied vacuously. However, by starting the induction from $k = 3$ and including the direct verification of P_3, insight is given into the process of constructing the N_i's which would otherwise have been lacking.

5.9 The Solution

By considering Cases A, B and C we have completed our analysis of all non-trivial games Γ. To bring our conclusions together, we make one final definition.

Definition 5.9.1 *Let Γ be a K-Arc Game. We define G^*, a mixed strategy for the Guard in Γ, as follows.*

$$G^* = \begin{cases} J^* & \text{if } \Gamma \text{ belongs to Case } A, \\ M^* & \text{if } \Gamma \text{ belongs to Case } B, \\ N^* & \text{if } \Gamma \text{ belongs to Case } C. \end{cases}$$

Theorem 5.9.2 *Let Γ be a K-Arc Game. Then G^* and I^* are optimal strategies for the Guard and the Infiltrator respectively, and the value of the game is given by $v(\Gamma) = v(n, w, \lambda)$.*

Proof Follows directly from Lemmas 5.6.11, 5.7.9 and 5.8.1.

We now draw attention to some of the consequences of this theorem. The first point to be noted is that the value of the game, $v(\Gamma)$, does not depend on the exact form of the graph Γ. Of course, Γ must consist of k disjoint arcs joining v_A and v_B, but the actual distribution of the vertices on the arcs is not of importance to the value of the game. All games played on graphs consisting of k arcs and m interior vertices, provided that there are between one and $n-1$ interior vertices on each arc, have the same value.

Consider also the situation where the number of interior vertices, $m > 1$, is fixed although k may vary. It can be shown that $v(\Gamma)$ is here minimised if $k = 1$ and maximised if $k = m$. These instances correspond to the two cases in which the Guard can exploit his superior freedom of movement to the greatest and least advantage respectively.

Now, for a fixed graph Γ, with $k \geq 1$, we consider what happens as the time limit, n, becomes very large. In all of the cases considered earlier, the value tends to $\lambda(1/k) + (1 - 1/k)$.

If the time limit is actually dispensed with, we must note that Γ is now an infinite game. In this case an optimal strategy for the Guard is simply, at each time, to choose at random an arc, and to occupy the first interior vertex on that arc. An ϵ-optimal strategy for the Infiltrator is to use the optimal strategy I^* in the finite game covered above when there is a time limit $n(\epsilon)$, where

$$n(\epsilon) \geq \frac{(m-k)(1-\lambda)}{\epsilon k^2} + \frac{m}{k}.$$

The value of the infinite game is $\lambda(1/k) + (1 - 1/k) = 1 - (1 - \lambda)/k$.

Alpern [3] has generalised this result concerning infinite move games. He shows, by invoking Menger's Theorem ([37], theorem 5.9), that for any finite connected graph with two vertices distinguished as the base and the target, the infinite move infiltration game has a solution in mixed strategies. He shows that, if k is interpreted as the number of elements in the smallest set which separates the base and the target (sometimes known as the 'min cut'), then the value is still equal to $1 - (1 - \lambda)/k$

Finally, let us move from an arbitrary graph back to the simple line. Let $k = 1$ and consider how the solution above compares to that of the Lalley Game Λ_n^*. If $k = 1$, then $w = n - m_\alpha < n$ and hence the game belongs to Case A. In fact, when $k = 1$, the optimal strategies I^* and J^* correspond precisely to the strategies X and Y given by Lalley. Unfortunately, the different approaches used obscure this correspondence. We show, however, that both solutions yield the same value.

When Γ consists of only one arc α, then as $w = n - m_\alpha$, we deduce that

$$\begin{aligned} v(\Gamma) &= \lambda^{z+1}\left(\frac{n-1}{n-m_\alpha} - z\right) + \lambda^z\left(1 + z - \frac{n-1}{n-m_\alpha}\right) \\ &= \lambda^{z+1}\left(\frac{y}{n-m_\alpha}\right) + \lambda^z\left(1 - \frac{y}{n-m_\alpha}\right), \end{aligned}$$

where $y = n - 1 - z(n - m_\alpha)$, or equivalently $m_\alpha - 1 = (z-1)(n - m_\alpha) + y$.

Now observe that, to compare the K-Arc Game when there is a single arc with the Lalley Game, we have to relate the notation by noting that $m_\alpha = p - 1$, and as $p - 2 = q(n - p + 1) + r$, then $y = r$ and $z = q + 1$. Thus, in the notation of section 2, we see that $v(\Gamma)$ is equal to

$$\lambda^{q+2}\left(\frac{r}{n+p-1}\right) + \lambda^{q+1}\left(1 - \frac{r}{n-p+1}\right),$$

which is precisely the value of the Lalley Game when $u = 1$.

This equality is somewhat surprising. The K-Arc Game has a Guard with no speed restriction. The Lalley Game with $u = 1$ has a Guard who is no faster than the Infiltrator. Yet, when $k = 1$ the values of these games are the same. When the players are playing on just the single arc, the Guard's speed advantage is of no use to him at all. In fact we deduced this earlier when we noted that in Lalley's solution there is no point at which the Infiltrator exploits his opponent's speed restriction. The alternative solution we have found here gives independent confirmation of this.

Chapter 6

SLOW GUARD DETECTION

6.1 Introduction

We now consider the game Γ_n^*, $n < \infty$. By constructing a mixed Guard strategy $\gamma^* \in \mathcal{G}_n$ we obtain an upper bound on the value of the game which, providing $p \geq 3$, is independent of p. Moreover, by constructing a mixed Infiltrator strategy $\iota^* \in \mathcal{I}_n$ we show that, when the speed of the Guard $u = 1$, then for certain values of the parameters λ and n, this upper bound is equal to the value.

Later we consider an example which shows what happens when the conditions mentioned above no longer hold and the value is strictly less than our upper bound.

Throughout this chapter it is to be assumed that $u = 1$.

6.2 An Upper Bound for the Value

We assume that $n \geq 3$, and that $3 \leq p \leq \infty$. Other than this we are not presently interested in their values. It is unimportant whether $n < p$, or $n \geq p$. We construct a mixed Guard strategy for the game Γ_n^*.

Definition 6.2.1 *Let $n \geq 3$ and $3 \leq p \lneq \infty$.*

(i) For $j = 0, \ldots, n-1$, the pure Guard strategy $h^j \in G_n$ is given by

$$h^j = \begin{cases} 1, 1, 2, 2, \ldots\ldots\ldots, 2 & \text{if } j = 0, \\ 1, 2, 2, 2, \ldots\ldots\ldots, 2 & \text{if } j = n-1, \\ 1, \underbrace{2, 2, \ldots, 2}_{j \text{ terms}}, 1, \ldots, 1 & \text{otherwise.} \end{cases}$$

(ii) The vector $q = (q_0, \ldots, q_{n-1})$ *is given by*

$$q_j = \frac{1}{A_n} \times \begin{cases} 1 & j = 0, \\ 1+\lambda & j = 1, \\ (1+\lambda)^{j-1} & j = 2, \ldots, n-2, \\ < (1+\lambda)^{j-1} - \lambda & j = n-1, \end{cases}$$

where $A_n = 1 + \sum_{r=0}^{n-2} (1+\lambda)^r$.

Since $A_n \sum_{r=0}^{n-1} q_j = 1 + (1+\lambda) + (1+\lambda)^1 + \cdots + (1+\lambda)^{n-2} - \lambda = A_n$, it follows that q is a probability vector.

(iii) The mixed strategy $\gamma^* \in \mathcal{G}_n$ *is given by*

$$\gamma^*(g) = \begin{cases} q_j & \textit{if there exists } j \textit{ such that } g = h^j, \\ 0 & \textit{otherwise.} \end{cases}$$

As we shall need to refer to this information easily we illustrate γ^* by means of the table below.

j	h^j		$A_n \times q_j$
0	$1,1,2,2,\ldots$	$\ldots,2$	1
1	$1,2,1,1,\ldots$	$\ldots,1$	$1+\lambda$
2	$1,2,2,1,1,\ldots$	$\ldots,1$	$(1+\lambda)^1$
3	$1,2,2,2,1,1,\ldots$	$\ldots,1$	$(1+\lambda)^2$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
θ	$1,\underbrace{2,2,\ldots,2}_{\theta \text{ terms}},1,1,\ldots$	$\ldots,1$	$(1+\lambda)^{\theta-1}$
$(2 \leq \theta \leq n-2)$			
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$n-2$	$1,2,2,\ldots$	$\ldots,2,1$	$(1+\lambda)^{n-3}$
$n-1$	$1,2,2,\ldots$	$\ldots,2$	$(1+\lambda)^{n-2} - \lambda$

Before continuing let us explain a piece of notation. Let s be a sequence with n terms. If $1 \leq \theta \leq n-1$, then it is conventional that the notation $s = \overbrace{s_1, \ldots, s_1}^{\theta \text{ terms}}, s_2, \ldots$ means that the first θ terms are equal to s_1, the next term is s_2 and the other terms can be anything. When $\theta = n-1$ there are no other terms.

Similarly, if $\theta = n$ the above notation simply denotes the constant sequence $s_1, \ldots, s_1$.

Remember that we are interested in the game Γ_n^*, $3 \le n < \infty$. Let us think about the Infiltrator's strategies in this game if it is assumed that the Guard plays strategy γ^*. Under γ^* the Guard is always to be found at either state one or state two. Therefore, if the Infiltrator ever moves to state two, it is best for him to immediately move on and never to return to state two. The best pure strategies to play in reply to γ^* are all of the form

$$x = \overbrace{1, 1, \ldots, 1}^{\theta \text{ terms}}, 2, 3, \ldots, \quad \text{where } 1 \le \theta \le n, \text{ and, if } t > \theta + 1,\ x_t > 2.$$

Let us suppose $X_n \subseteq I_n$ is the set of all pure strategies of this form.

Lemma 6.2.2 *For all $n \ge 3$, $3 \le p \le \infty$ and $x \in X_n$,*

$$A_n \pi_n(x, \gamma^*) = \lambda(1+\lambda)^{n-1}$$

Proof Let $x \in X_n$. By definition of γ^*, we can decompose the payoff as

$$\pi_n(x, \gamma^*) = \sum_{j=0}^{n-1} q_j \, \Pi_n\left(x, h^j\right) = \sum_{j=0}^{n-1} q_j \, \lambda^{\omega_n(x, h^j)}.$$

Now as $x = \overbrace{1, 1, \ldots, 1}^{\theta \text{ terms}}, 2, 3, \ldots$ as above, then we consider separately the cases $\theta = 1$, $2 \le \theta \le n-1$, and $\theta = n$. If $\theta = 1$, $x = 1, 2, 3, \ldots$, so $\omega_n(x, h^j)$, the number of meetings between x and h^j, is 1 if $j = 0$, and 2 otherwise. Therefore

$$\begin{aligned}
A_n \pi_n(x, \gamma^*) &= \lambda + \lambda^2 \left[(1+\lambda) + \sum_{j=2}^{n-1} (1+\lambda)^{j-1} - \lambda \right] \\
&= \lambda + \lambda^2 + \lambda^3 + \lambda^2(1+\lambda) \frac{(1+\lambda)^{n-2} - 1}{(1+\lambda) - 1} - \lambda^3 \\
&= \lambda + \lambda^2 + \lambda^3 + \lambda(1+\lambda)^{n-1} - \lambda(1+\lambda) - \lambda^3 \\
&= \lambda(1+\lambda)^{n-1}
\end{aligned}$$

If $2 \le \theta \le n-1$, then

$$\omega_n(x, h^j) = \begin{cases} 3 & j = 0, \\ \theta - j & j = 1, \ldots, \theta - 1, \\ 2 & j = \theta, \ldots, n-1, \end{cases}$$

and so

$$
\begin{aligned}
&A_n\pi_n(x,\gamma^*)\\
&= \lambda^3+\lambda^{\theta-1}(1+\lambda)+\sum_{j=2}^{\theta-1}\lambda^{\theta-j}(1+\lambda)^{j-1}+\lambda^2\left[\sum_{j=\theta}^{n-1}(1+\lambda)^{j-1}-\lambda\right]\\
&= \lambda^3+\lambda^{\theta-1}+\lambda^{\theta}+\lambda^{\theta-1}\sum_{j=2}^{\theta-1}\left(\frac{1+\lambda}{\lambda}\right)^{j-1}+\lambda^2(1+\lambda)^{\theta-1}\sum_{j=1}^{n-\theta}(1+\lambda)^{j-1}\\
&\quad -\lambda^3\\
&= \lambda^3+\lambda^{\theta-1}+\lambda^{\theta}+\lambda^{\theta-1}\left(\frac{1+\lambda}{\lambda}\right)\frac{\left(\frac{1+\lambda}{\lambda}\right)^{\theta-2}-1}{\left(\frac{1+\lambda}{\lambda}\right)-1}\\
&\quad +\lambda^2(1+\lambda)^{\theta-1}\frac{(1+\lambda)^{n-\theta}-1}{(1+\lambda)-1}-\lambda^3\\
&= \lambda^3+\lambda^{\theta-1}+\lambda^{\theta}+\lambda(1+\lambda)^{\theta-1}-\lambda^{\theta-1}(1+\lambda)+\lambda(1+\lambda)^{n-1}\\
&\quad -\lambda(1+\lambda)^{\theta-1}-\lambda^3\\
&= \lambda(1+\lambda)^{n-1}
\end{aligned}
$$

Finally, if $\theta=n$, then $\omega(x,h^j)=2$ if $j=0$, and $n-j$ otherwise. Therefore

$$
\begin{aligned}
&A_n\pi_n(x,\gamma^*)\\
&= \lambda^2+\lambda^{n-1}(1+\lambda)+\sum_{j=2}^{n-2}\lambda^{n-j}(1+\lambda)^{j-1}+\lambda\left[(1+\lambda)^{n-2}-\lambda\right]\\
&= \lambda^2+\lambda^{n-1}+\lambda^{n}+\lambda^{n-1}\sum_{j=2}^{n-2}\left(\frac{1+\lambda}{\lambda}\right)^{j-1}+\lambda(1+\lambda)^{n-2}-\lambda^2\\
&= \lambda^2+\lambda^{n-1}+\lambda^{n}+\lambda^{n-1}\left(\frac{1+\lambda}{\lambda}\right)\frac{\left(\frac{1+\lambda}{\lambda}\right)^{n-3}-1}{\left(\frac{1+\lambda}{\lambda}\right)-1}+\lambda(1+\lambda)^{n-2}-\lambda^2\\
&= \lambda^2+\lambda^{n-1}+\lambda^{n}+\lambda^2(1+\lambda)^{n-2}-\lambda^{n-1}(1+\lambda)+\lambda(1+\lambda)^{n-2}-\lambda^2\\
&= (\lambda+\lambda^2)(1+\lambda)^{n-2}\\
&= \lambda(1+\lambda)^{n-1}.
\end{aligned}
$$

Corollary 6.2.3 *For all* $3\le p\le\infty$ *and* $n\ge 3$, $v_n\le\dfrac{\lambda(1+\lambda)^{n-1}}{A_n}$.

Proof We noted above that any best reply to γ^* which is pure must be contained in X_n. Thus, from the previous Lemma, for all $i\in I_n$, $A_n\pi_n(i,\gamma^*)\le\lambda(1+\lambda)^{n-1}$ and so $v_n\le\lambda(1+\lambda)^{n-1}/A_n$.

6.3 An Equalising Strategy

We shall now construct a mixed strategy ι^* which equalises against $h^0, \ldots, h^{n-1}$. That is, for all $0 \le j \le n-1$, $A_n \pi_n(\iota^*, h^j) = \lambda(1+\lambda)^{n-1}$. In addition ι^* gives the same payoff against any Guard strategy that never moves beyond state two. In the particular case that $\lambda(1+\lambda)^{n-2} < 1$, such strategies are the best replies against ι^*, and ι^* and γ^* are then optimal.

Definition 6.3.1 *Let $n < \infty$ and $2 \le p \le \infty$.*

(i) For $j = 1, \ldots, n$ the jth **wait-and-run strategy**, *$w^j \in I_n$, is the sequence*
$$\underbrace{1, 1, \ldots, 1}_{j \text{ terms}}, 2, 3, \ldots, p-1, p, p, \ldots, p$$

(ii) The vector $p = (p_1, \ldots, p_n)$ is given by
$$p_j = \frac{1}{A_n} \times \begin{cases} (1+\lambda)^{n-j-1} & 1 \le j \le n-1, \\ 1 & j = n, \end{cases}$$
where $A_n = 1 + \sum_{r=0}^{n-2} (1+\lambda)^r$.

Observe that $A_n \sum_{j=1}^{n} p_j = 1 + 1 + (1+\lambda)^1 + \cdots + (1+\lambda)^{n-2} = A_n$, and so p is a probability vector.

(iii) The mixed strategy $\iota^* \in \mathcal{I}_n$ is defined by
$$\iota^*(i) = \begin{cases} p_j & \textit{if there exists } j \textit{ such that } i = w^j, \\ 0 & \textit{otherwise.} \end{cases}$$

Note that, for $n - p + 1 < j \le n$, w^j does not make it to the target. These strategies would be no good in the target game but are useful here because the Infiltrator does not have to reach the target. As $n \to \infty$, $\sum_{j=n-p+2}^{n} p_j \to 0$ and so these strategies become increasingly less important. We illustrate ι^* by means of the table on the following page.

Let us now consider the Guard's best replies to ι^*. There must be a best reply which is pure so we shall concentrate on how different members of G_n perform against ι^*. We observe immediately that, if $g \in G_n$ is a best reply then, as the

Infiltrator is known to start at state one, $g_1 = 1$. Now consider the situation if we take a Guard strategy which only ever visits states one and two.

j	w^j		$A_n \times p_j$
1	$1, 2, 3, \ldots, p, p, \ldots$	$\ldots, p$	$(1+\lambda)^{n-2}$
2	$1, 1, 2, 3, \ldots, p, p, \ldots$	$\ldots, p$	$(1+\lambda)^{n-3}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
θ	$\underbrace{1, 1, \ldots, 1}_{\theta \text{ terms}}, 2, 3, \ldots$	$\ldots, p$	$(1+\lambda)^{n-\theta-1}$
$(1 \leq \theta \leq n-1)$			
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$n-1$	$1, 1, \ldots\ldots$	$\ldots, 1, 2$	1
n	$1, 1, \ldots\ldots$	$\ldots, 1$	1

Lemma 6.3.2 *If $g \in G_n$ satisfies $g_1 = 1$ and, for all $2 \leq t \leq n$, $g_t = 1$ or 2, then $A_n \pi_n(\iota^*, g) = \lambda(1+\lambda)^{n-1}$.*

Proof We shall use two approaches here according to whether g ends at state one or state two. If $g_n = 1$, then by a sequence of identical steps we find that, against ι^*, g gives the same payoff as the strategy $1, 1, \ldots, 1$. Similarly, if $g_n = 2$, g gives the same payoff as the strategy $1, 2, 2, \ldots, 2$. Our first step is to show that these two strategies both give a payoff of $\lambda(1+\lambda)^{n-1}/A_n$ against ι^*.

Let $h^{n-1} = 1, 2, 2, \ldots, 2$ (note this is consistent with how h^{n-1} was given earlier in Definition 6.2.1). The payoff $\pi_n(\iota^*, h^{n-1})$ depends upon how many times h^{n-1} meets each of the wait-and-run strategies $w^1, \ldots, w^n$. Clearly h^{n-1} meets all of them once at state one and then each of $w^1, \ldots, w^{n-1}$ a second time at state two. Therefore

$$
\begin{aligned}
A_n \pi_n(\iota^*, h^{n-1}) &= \sum_{j=1}^{n} A_n \, p_j \, \lambda^{\omega_n(w^j, h^{n-1})} \\
&= \lambda^2 \left[(1+\lambda)^{n-2} + (1+\lambda)^{n-3} + \cdots + 1 \right] + \lambda \\
&= \lambda^2 \left[\frac{(1+\lambda)^{n-1} - 1}{(1+\lambda) - 1} \right] + \lambda \\
&= \lambda(1+\lambda)^{n-1}.
\end{aligned}
$$

Similarly if $h^n = 1, 1, \ldots, 1$, then for $1 \leq j \leq n$, $\omega_n(w^j, h^n) = j$. Therefore

$$
\begin{aligned}
A_n \pi_n(\iota^*, h_n) &= \sum_{j=1}^{n-1} \lambda^j (1+\lambda)^{n-j-1} + \lambda^n \\
&= \lambda(1+\lambda)^{n-2} \sum_{j=1}^{n-1} \left(\frac{\lambda}{1+\lambda}\right)^{j-1} + \lambda^n \\
&= \lambda(1+\lambda)^{n-2} \left[\frac{1 - \left(\frac{\lambda}{1+\lambda}\right)^{n-1}}{1 - \left(\frac{\lambda}{1+\lambda}\right)} \right] + \lambda^n \\
&= \lambda(1+\lambda)^{n-1}.
\end{aligned}
$$

Now suppose that $g \in G_n$ is any Guard strategy which satisfies $g_1 = 1$, $g_n = 2$ and, for all $1 < t < n$, $g_t \leq 2$. There must exist a last time r at which g is at state one. Note that $1 \leq r \leq n-1$. If $r = 1$, then $g = h^{n-1}$, so suppose that $r \geq 2$. Representing the path of g during the period $[r, n]$ by bullets • we illustrate this in the diagram below. Note that the diagonal arrows represent the paths of some of the wait-and-run strategies.

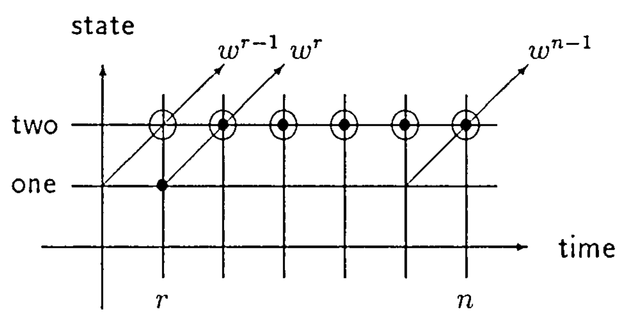

Now compare to this the strategy $h \in G_n$ which is the same as g except that $h_r = 2$. The path of h is represented by circles ∘ in the diagram above. Comparing how many times g and h meet the different wait-and-run strategies observe that, although h meets w^{r-1} once more than does g, because of spending one time unit less at state one, it meets each of $w^r, w^{r+1}, \ldots, w^n$ once less. Therefore

$$
\omega_n(w^j, h) = \begin{cases} \omega_n(w^j, g) & 1 \leq j \leq r-2, \\ \omega_n(w^j, g) + 1 & j = r-1, \\ \omega_n(w^j, g) - 1 & r \leq j \leq n. \end{cases}
$$

Moreover if up to time $r-1$, g visits state one on ϕ occasions, then

$$
\omega_n(w^j, g) = \begin{cases} \phi & j = r-1, \\ \phi + 2 & r \leq j \leq n-1, \\ \phi + 1 & j = n. \end{cases}
$$

Therefore

$$
\begin{aligned}
& A_n\left[\pi_n(\iota^*, g) - \pi_n(\iota^*, h)\right] \\
&= A_n \sum_{j=1}^{n} p_j \left[\lambda^{\omega_n(w^j,g)} - \lambda^{\omega_n(w^j,h)}\right] \\
&= (1-\lambda)\lambda^{\phi}(1+\lambda)^{n-r} + \left(1-\lambda^{-1}\right)\lambda^{\phi+2}\sum_{j=r}^{n-1}(1+\lambda)^{n-j-1} \\
&\quad + \left(1-\lambda^{-1}\right)\lambda^{\phi+1} \\
&= (1-\lambda)\lambda^{\phi}\left[(1+\lambda)^{n-r} - \lambda\sum_{j=0}^{n-r-1}(1+\lambda)^j - 1\right] \\
&= (1-\lambda)\lambda^{\phi}\left[(1+\lambda)^{n-r} - \lambda\frac{(1+\lambda)^{n-r}-1}{(1+\lambda)-1} - 1\right] \\
&= 0.
\end{aligned}
$$

Thus by repeated application of this step we deduce that for any g which starts at state one and ends at state two, $A_n\pi_n(\iota^*, g) = A_n\pi_n(\iota^*, h^{n-1}) = \lambda(1+\lambda)^{n-1}$, where h^{n-1} is the Guard strategy which moves immediately to state two and then stays there.

Now suppose that $g \in G_n$ is any Guard strategy which satisfies $g_1 = 1$, $g_n = 1$ and, for all $1 < t < n$, $g_t \leq 2$. If g never visits state two then $g = h^n$. Suppose that it does and hence there is a last time r such that $g_r = 2$. Observe that $2 \leq r \leq n-1$. The path of g during the period $[r, n]$ is represented by bullets on the diagram below.

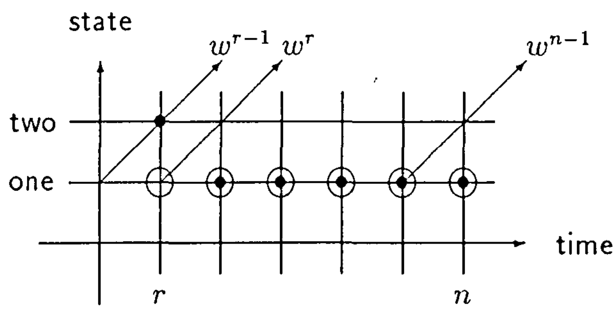

Compare with this the strategy $h \in G_n$ which is the same as g except that $h_r = 1$. The path of h is represented by circles. Observe that here

$$
\omega_n(w^j, h) = \begin{cases} \omega_n(w^j, g) & 1 \leq j \leq r-2, \\ \omega_n(w^j, g) - 1 & j = r-1, \\ \omega_n(w^j, g) + 1 & r \leq j \leq n. \end{cases}
$$

Moreover, if, up to time $r-1$, g visits state one on ϕ occasions then

$$\omega_n(w^j,g) = \begin{cases} \phi+1 & j=r-1, \\ \phi+(j-r) & r \le j \le n. \end{cases}$$

Therefore

$$\begin{aligned}
&A_n\left[\pi_n(\iota^*,g)-\pi_n(\iota^*,h)\right] \\
&= A_n\sum_{j=1}^{n} p_j\left[\lambda^{\omega_n(w^j,g)}-\lambda^{\omega_n(w^j,h)}\right] \\
&= \left(1-\lambda^{-1}\right)\lambda^{\phi+1}(1+\lambda)^{n-r}+(1-\lambda)\sum_{j=r}^{n-1}\lambda^{\phi+j-r}(1+\lambda)^{n-j-1} \\
&\quad +(1-\lambda)\lambda^{\phi+n-r} \\
&= (1-\lambda)\lambda^{\phi}\left[-(1+\lambda)^{n-r}+(1+\lambda)^{n-r-1}\sum_{j=r}^{n-1}\left(\frac{\lambda}{1+\lambda}\right)^{j-r}+\lambda^{n-r}\right] \\
&= (1-\lambda)\lambda^{\phi}\left[-(1+\lambda)^{n-r}+(1+\lambda)^{n-r-1}\frac{1-\left(\frac{\lambda}{1+\lambda}\right)^{n-r}}{1-\left(\frac{\lambda}{1+\lambda}\right)}+\lambda^{n-r}\right] \\
&= 0.
\end{aligned}$$

Thus, by repeated application of this step we deduce that for any g which starts and ends at state one, $A_n\pi_n(\iota^*,g)=A_n\pi_n(\iota^*,h^n)=\lambda(1+\lambda)^{n-1}$, where h^n is the pure Guard strategy which always stays at state one. This completes the proof.

We have shown that all Guard strategies which start at state one and then move only between states one and two yield the same payoff against ι^*. We shall now see that if λ satisfies $\lambda(1+\lambda)^{n-2}<1$, these are the only pure strategies which are best replies to ι^*. Recall again that for any $g \in G_n$ to be a best reply it must at least satisfy $g_1=1$.

Lemma 6.3.3 *Let λ satisfy $\lambda(1+\lambda)^{n-2}<1$. Let $g \in G_n$, $g_1=1$ and suppose there exists $3 \le t \le n$ such that $g_t \ge 3$. Then there exists $h \in G_n$ such that $h_1=1$, for all t, $h_t=1$ or 2, and $\pi_n(\iota^*,g)-\pi_n(\iota^*,h)>0$.*

Proof Suppose $g \in G_n$ is a strategy which moves beyond state two. There are two steps which we use to gradually improve g and eventually arrive at the strategy h required. Considering the path of g drawn with bullets on the diagram below, the first step flattens the bumps in the path and the second reduces the slope of the final ascent. The improved path h is represented by the circles.

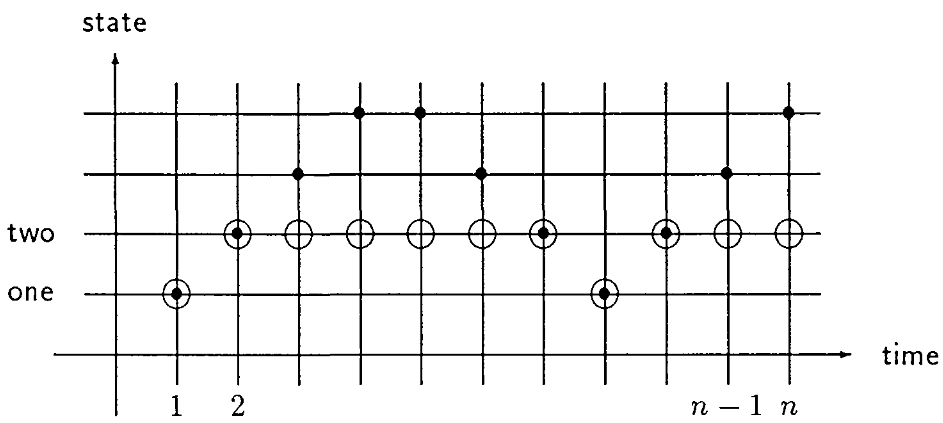

First suppose the Guard strategy g moves beyond state two towards the target, but at some time later takes a backward step.

Thus, there exist $2 \leq \theta < p-1$ and $\theta+1 \leq r < s \leq n$ such that in the period $[r-1, s]$, the path of g is given by

$$g_t = \begin{cases} \theta & t = r-1, s, \\ \theta+1 & \text{otherwise.} \end{cases}$$

Now compare this with the strategy h which is the same as g except that, for $r \leq t \leq s-1$, $h_t = \theta$. The paths of g and h, in the interval $[r-1, s]$, are represented on the diagram below by the bullets and circles respectively.

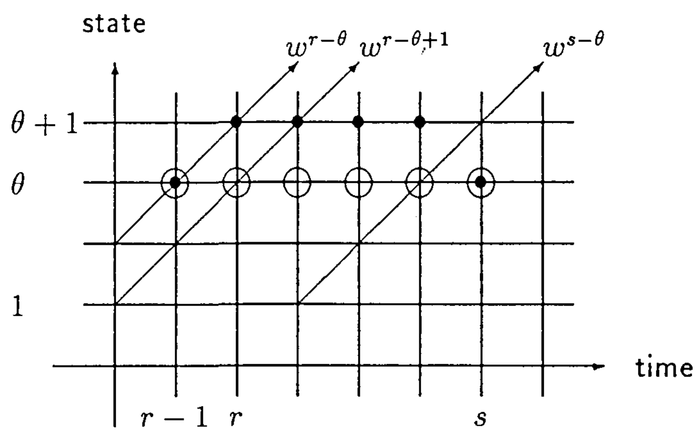

It is obvious that, for j outside the interval $[r-\theta, s-\theta]$, g and h meet the path of strategy w^j at precisely the same times. In fact, we deduce that

$$\omega_n(w^j,h) = \begin{cases} \omega_n(w^j,g)-1 & j=r-\theta, \\ \omega_n(w^j,g)+1 & j=s-\theta, \\ \omega_n(w^j,g) & \text{otherwise.} \end{cases}$$

Finally we note that if, before time r, the strategy g spends ϕ units of time at state one, then $\omega_n(w^{r-\theta},g) \geq \phi+2$ whereas $\omega_n(w^{s-\theta},g) = \phi$.

So, comparing how g and h perform against ι^* we see that

$$\begin{aligned} & A_n\left[\pi_n(\iota^*,g) - \pi_n(\iota^*,h)\right] \\ = \;& A_n \sum_{j=1}^{n} p_j \left[\lambda^{\omega_n(w^j,g)} - \lambda^{\omega_n(w^j,h)}\right] \\ \geq \;& \left(1-\lambda^{-1}\right)\lambda^{\phi+2}(1+\lambda)^{n-r+\theta-1} + (1-\lambda)\lambda^{\phi}(1+\lambda)^{n-s+\theta-1} \\ = \;& (1-\lambda)\lambda^{\phi}(1+\lambda)^{n-s+\theta-1}\left[-\lambda(1+\lambda)^{s-r}+1\right] \\ \geq \;& (1-\lambda)\lambda^{\phi}(1+\lambda)^{n-s+\theta-1}\left[-\lambda(1+\lambda)^{n-3}+1\right] \quad \text{as } r \geq 3 \text{ and } s \leq n, \\ > \;& 0 \quad \text{since } \lambda(1+\lambda)^{n-3} < \lambda(1+\lambda)^{n-2} < 1. \end{aligned}$$

Secondly, suppose that the Guard strategy g makes a move beyond state two towards the target but this time never subsequently takes a backward step. Thus $g_n = \theta$, where $\theta \geq 3$. Moreover there exists a time r, $\theta \leq r \leq n$ at which g first arrives at state θ.

Thus, in the period $[r-1, n]$,

$$g_t = \begin{cases} \theta-1 & t=r-1, \\ \theta & \text{otherwise.} \end{cases}$$

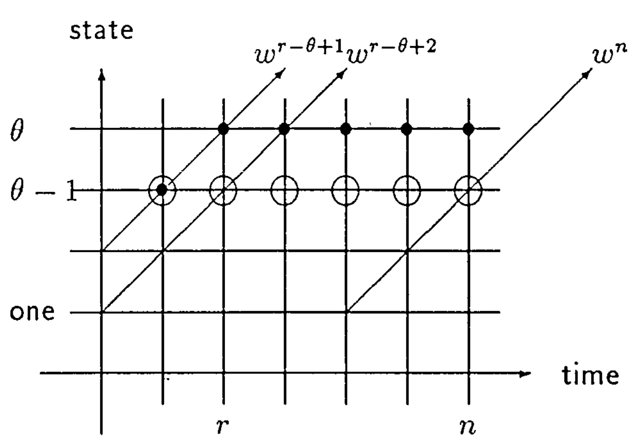

In the diagram above the path of g is represented by the bullets. Now compare g with the strategy h which is the same as g except that, for $r \leq t \leq n$, $h_t = \theta - 1$. The path of h is represented by circles. Considering the times at which g and h meet the wait-and-run strategies we deduce that

$$\omega_n(w^j, h) = \begin{cases} \omega_n(w^j, g) - 1 & j = r - \theta + 1, \\ \omega_n(w^j, g) + 1 & j = n - \theta + 2, \\ \omega_n(w^j, g) & \text{otherwise.} \end{cases}$$

Moreover, if the strategy g spends precisely ϕ units of time at state one, then $\omega_n(w^{r-\theta+1}, g) \geq \phi + 2$ whereas $\omega_n(w^{n-\theta+2}, g) = \phi$. Therefore

$$\begin{aligned} & A_n \left[\pi_n(\iota^*, g) - \pi_n(\iota^*, h)\right] \\ = \;& A_n \sum_{j=1}^{n} p_j \left[\lambda^{\omega_n(w^j,g)} - \lambda^{\omega_n(w^j,h)}\right] \\ \geq \;& \left(1 - \lambda^{-1}\right) \lambda^{\phi+2}(1+\lambda)^{n-r+\theta-2} + (1-\lambda)\lambda^{\phi}(1+\lambda)^{\theta-3} \\ = \;& (1-\lambda)\lambda^{\phi}(1+\lambda)^{\theta-3}\left[-\lambda(1+\lambda)^{n-r+1} + 1\right] \\ \geq \;& (1-\lambda)\lambda^{\phi}(1+\lambda)^{\theta-3}\left[-\lambda(1+\lambda)^{n-2} + 1\right] \quad \text{as } r \geq \theta \geq 3, \\ > \;& 0 \quad \text{as } \lambda(1+\lambda)^{n-2} < 1. \end{aligned}$$

Finally suppose that g is any pure Guard strategy which starts at state one but subsequently moves beyond state two. Repeated application of the two steps above produces a strategy $h \in G_n$ such that $h_1 = 1$, for all t, $h_t = 1$ or 2, and $\pi_n(\iota^*, g) - \pi_n(\iota^*, h) > 0$.

Theorem 6.3.4 *Let $n \geq 3$, $3 \leq p \leq \infty$ and $\lambda(1+\lambda)^{n-2} < 1$. Then ι^* and γ^* are optimal in Γ_n^* and $v_n = \lambda(1+\lambda)^{n-1}/A_n$.*

Proof Follows immediately from Lemmas 6.2.2, 6.3.2 and 6.3.3.

6.4 When Wait-and-Run is Not Good Enough

If $n = 3$, $3 \leq p \leq \infty$ and $u = 1$, the game Γ_n^* is represented by the matrix below, with the rows and columns representing the pure strategies for the Infiltrator and the Guard respectively (only non-dominating Guard strategies are included). Note that entry $1, 2, 3^*$ is available to the Guard if and only if $p > 3$.

	1,2,3*	1,2,2	1,2,1	1,1,2	1,1,1
1,2,3	λ^3	λ^2	λ^2	λ	λ
1,2,2	λ^2	λ^3	λ^2	λ^2	λ
1,2,1	λ^2	λ^2	λ^3	λ	λ^2
1,1,2	λ	λ^2	λ	λ^3	λ^2
1,1,1	λ	λ	λ^2	λ^2	λ^3

From above, if $\lambda(1+\lambda) < 1$, then the probability distributions $\iota^* = (1+\lambda, 0, 0, 1, 1)/A_3$ and $\gamma^* = (0, 1, 1+\lambda, 1, 0)/A_3$ are optimal for the Infiltrator and Guard respectively, and $v_3 = \lambda(1+\lambda)^2/A_3$. This could be checked on the matrix above. Now let us consider what happens when $\lambda(1+\lambda) \geq 1$.

If $p = 3$ and $\lambda(1+\lambda) \geq 1$, then ι^* and γ^* are still optimal. Again this can be checked manually above. It is, in fact, a consequence of Lemma 6.3.2, since the Guard can only move between states one and two.

However, if $p > 3$ and $\lambda(1+\lambda) \geq 1$ the solution is different. There is a strategy for both players which is an equalising strategy. Let

$$\iota = \gamma = (1 - 2\alpha - 2\beta, \alpha, \alpha, \beta, \beta)$$

where $\alpha = (\lambda^2 + \lambda - 1)/(5\lambda^2 + 8\lambda - 1)$ and $\beta = (\lambda^2 + 2\lambda)/(5\lambda^2 + 8\lambda - 1)$. Then it can be verified above that, for all $(i, g) \in I_3 \times G_3$, $\pi_3(i, \gamma) = \pi_3(\iota, \gamma) = \pi_3(\iota, g)$, and thus

$$v_3 = \pi_3(\iota, \gamma) = \frac{(\lambda^5 + 4\lambda^4 + 5\lambda^3 + 2\lambda^2)}{(5\lambda^2 + 8\lambda - 1)}.$$

For all $\lambda > \left(\sqrt{5} - 1\right)/2 \approx 0.618$, $\lambda(1+\lambda) > 1$ and it can be shown that, as given above, $v_3 < \lambda(1+\lambda)^2/A_3$. Hence γ^*, as defined in section 2, is no longer optimal for the Guard. Also the equalizing strategy given above involves the Infiltrator playing more than just wait-and-run strategies. In fact, if he is restricted to only probability distributions over the set of wait-and-run strategies, it is the strategy $(2+\lambda, 0, 0, 1+\lambda, 1+\lambda)/(4+3\lambda)$ which is optimal for both players. The value $\bar{v}_3$ of this restricted game is $(2\lambda + 2\lambda^2 + 2\lambda^3 + \lambda^4)/(4 + 3\lambda)$. For all λ such that $\lambda(1+\lambda) > 1$, $\bar{v}_3 < v_3$. Therefore in the original, unrestricted game it is not optimal for the Infiltrator to use only probability distributions over the wait-and-run strategies.

When the Infiltrator can restrict his attention to wait-and-run strategies he is dealing with a subset of his strategies which can be easily ordered. It seems unsurprising that in some of these cases optimal strategies can be found which follow some sort of pattern. One example of this has been given in this chapter.

6.5 One More Detection Game

The previous sections can be used in one more way. Observe that if $\lambda = 0$ then the condition $\lambda(1+\lambda)^{n-2} < 1$ is satisfied for all $n \geq 2$. This corresponds to the game in which the probability of a miss is zero. In this case, as the Guard may immediately detect the Infiltrator at state one, the value of all target and detection games is zero.

But what if the Guard is unable to make this initial detection? We consider the game which is identical to the slow Guard detection game except that the Infiltrator cannot be detected until time 2. We shall show that, provided λ satisfies $\lambda(1+\lambda)^{n-2} < 1$, the mixed strategies ι^* and γ^* are still optimal and the value of this game is $(1+\lambda)^{n-1}/A_n$.

Repeating the order in which we approached the original problem, consider the Guard first. Define the mixed strategy γ^* according to Definition 6.2.1. Note that again the best replies to γ^* are all contained in X_n and that the equivalent to Lemma 6.2.2 is that against γ^*, for all $x \in X_n$, the payoff is $(1+\lambda)^{n-1}/A_n$ (Note the absence of the multiplier λ which corresponded to the probability of a miss at time 1).

Now consider again the Infiltrator strategy ι^* as given by Definition 6.3.1. We need no longer assume that a best reply to ι^* must start at state one. However, as at time 2 the Guard should be at either state one or two, we can follow through the proof of Lemma 6.3.2. We deduce that for all $g \in G_n$ which safisfies $g_t = 1$ or 2 for $2 \leq t \leq n$, the payoff against ι^* is $(1+\lambda)^{n-1}/A_n$. Finally we must consider how ι^* performs against any $g \in G_n$. From above we can assume that $g_2 = 1$ or 2. Following the proof of Lemma 6.3.3 we deduce that providing $\lambda(1+\lambda)^{n-2} < 1$, the best replies to ι^* which are pure are all of the type described above. This completes the proof that in this amended game when $\lambda(1+\lambda)^{n-2} < 1$, ι^* and γ^* are still optimal and the value is $(1+\lambda)^{n-1}/A_n$.

In particular this gives the solution to the game when $\lambda = 0$. This is non-trivial and the value is $1/n$. In this case the optimal strategies ι^* and γ^* involve the players choosing randomly between their respective sets of n pure strategies.

Finally, this adapted game with $\lambda = 0$ can also be studied when there is a fast Guard. It seems likely that, providing n and p are not too small, the Guard can ensure a payoff lower then $1/n$. However the solutions to these games are more complex. For example, consider the adapted game when $\lambda = 0$, $n = 4$ and $p > 4$.

Note that the Guard might as well start from state one. Moreover if at time 2 he searches state two there is no point in him searching state three next time. Omitting these and other dominating pure strategies for the Guard, the game matrix is given below. The value of this game is $\frac{4}{23}$. An optimal strategy for the Infiltrator is $(2,0,1,2,1,0,1,2,2,2,2,2,4,4)/23$ and for the Guard $(0,1,2,0,0,1,0,2,2,1,0,2,2,2,4,4)/23$. Of course $\frac{4}{23} < \frac{1}{4}$.

	1 1 1 1	1 1 1 2	1 1 1 3	1 1 1 4	1 1 2 1	1 1 2 2	1 1 2 3	1 1 2 4	1 1 3 1	1 1 3 2	1 1 3 3	1 2 1 1	1 2 1 2	1 2 1 3	1 2 2 1	1 2 2 2
1,2,3,4	1	1	1	0	1	1	1	0	0	0	0	0	0	0	0	0
1,2,3,3	1	1	0	1	1	1	0	1	0	0	0	0	0	0	0	0
1,2,3,2	1	0	1	1	1	0	1	1	0	0	0	0	0	0	0	0
1,2,2,3	1	1	0	1	0	0	0	0	1	1	0	0	0	0	0	0
1,2,2,2	1	0	1	1	0	0	0	0	1	0	1	0	0	0	0	0
1,2,2,1	0	1	1	1	0	0	0	0	0	1	1	0	0	0	0	0
1,2,1,2	0	0	0	0	1	0	1	1	1	0	1	0	0	0	0	0
1,2,1,1	0	0	0	0	0	1	1	1	0	1	1	0	0	0	0	0
1,1,2,3	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0
1,1,2,2	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0
1,1,2,1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0
1,1,1,2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
1,1,1,1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

Chapter 7

CONCLUSIONS

7.1 Recapitulation

In this final chapter we discuss possible extensions of this work on discrete problems of infiltration. We consider some further conclusions which can be drawn from the present work, and some areas for further research. First, however, we recap on the results we have found so far.

In Chapters 2, 3, 4 and 6 we have concentrated on particular discrete infiltration games which are played on the line. In Chapter 2 we defined infinite games in which there was no restriction on the number of moves made by the players. We considered two payoffs, the probability of reaching the target and the probability of remaining undetected. Thus, we started with two basic games — a target game and a detection game. Theorem 2.7.3 showed that both of these games have solutions in terms of mixed strategies. Moreover, it was also shown that in each game, at least one of the players has an optimal strategy.

In Chapter 3 we considered an alternative approach to the games defined the chapter before. We defined two further sets of games in which there is a limit on the number of moves. As these games are finite they were known to possess mixed strategy solutions. By considering the properties of these games as the time limit tends to infinity, we deduced information about the original infinite games. Hence, in Theorem 3.4.4, we deduced that the infinite target game and the infinite detection game have the same value. We also showed that ϵ-optimal strategies for these games could be obtained from optimal strategies for the finite games.

In Chapters 4 and 6 we attempted to illustrate some of the features of solutions to certain finite and infinite games. In particular, in Chapter 4, we found that, in some infinite detection games, there does not exist an optimal mixed strategy for the Guard. Also, providing the number of states is at least three and the probability of detection is less than one, the value of a finite target game is strictly less than that of the finite detection game with the same time limit. It was also demonstrated how important are the speed of the Guard, u, and the probability of a miss, λ. The form of the solutions could change considerably at certain threshold values of these parameters. In this chapter wait-and-run strategies were also encountered for the first time. We found that, while these pure strategies play an important role in many optimal strategies for the Infiltrator, sometimes other pure strategies are also required. In Chapter 6 we gave a partial solution for finite detection games.

In Chapter 5 we considered an adaptation of the discrete infiltration problems. A safe point was introduced from which the Infiltrator then started. In this case it only made sense to consider the probability of reaching the target. In a detection game the Infiltrator could just stay put forever. Discussing Lalley's work in this area, we extended his conclusions and solved a generalisation of the problem on $k \geq 1$ discrete arcs. It was illustrated that the addition of a safe base considerably simplified the problem on the line. For example, the Guard's speed no longer affects the value. It is of no use to him to have any speed advantage over the Infiltrator. Moreover, in these safe base games, we also demonstrated the sufficiency for the Infiltrator of mixed strategies which are composed only of wait-and-run strategies.

7.2 Generalisations

There are clearly many directions in which generalisations could be made. We shall concentrate on some of those in which both time and space remain discrete, and in which the players gain no useful information about one another's moves. As mentioned in the introduction, to change these assumptions would lead us into a completely different set of problems, some of which were discussed there. We are going to briefly discuss some extensions that can be made in the following areas of the problem: the playing space, the detection probability, and the initial conditions.

We have concentrated upon the discrete line. Although in Chapter 5 we have looked at a safe base game on several arcs, this is the only time on which we have moved away from the linear problems. There are several reasons for this. Perhaps the most influential has been the richness of this apparently simple set of problems. We have had a desire to find the values for some of these games and get at least some idea of what optimal or ϵ-optimal strategies would look like. Hence the solutions given in Chapter 4 and Chapter 6 are tailored particularly to these linear problems. Another reason for looking so closely at the line has been to keep things as simple as possible. At times our notation has been cumbersome enough. It was felt that to present the basic ideas of Chapter 2 and Chapter 3 any more generally would only have made matters worse.

However, it does seem likely that these basic ideas can be appropriated for a wider range of discrete search problems. Consider the following infinite move situation. Let P be any *finite* set of states. For every state $v \in P$, let P_v^I be the set of states to which the Infiltrator can move from state v, and P_v^G the analogous set for the Guard. A pure strategy for the Infiltrator could be defined as an infinite sequence $(i_r) \in P^\infty$ which satisfies, for all $r \geq 1$, $i_{r+1} \in P_{i_r}^I$, and some other conditions relating to starting position and rules about reaching the target. Likewise for the Guard. The Infiltrator could start at any distinguished state, and have a general set S of target states. The details are not in themselves important, although, as we have seen, it takes very little to affect the compactness of a strategy space. Suppose that it were possible to define a pair of compact strategy spaces and an appropriate semicontinuous payoff function. It would then follow from Glicksberg [34] or Alpern and Gal [2] that this game had a solution in mixed strategies. More importantly, we could introduce a time limit and consider the 'finite version' of this game. It seems probable that, as was found for the linear games in Chapter 3, optimal or ϵ-optimal strategies for the infinite game could then be constructed out of solutions to the finite game. If this were so, Theorem 3.3.8 ($v_\infty = v_{\lim}$) would also have its analogue in a more general discrete game.

What happens when the set of states is infinite? In particular, let us consider this in the light of the games we introduced in Chapters 2 and 3. When $p = \infty$ it clearly makes no sense to talk about a target game, so we need only consider the detection games Γ_∞ and Γ_n, $n < \infty$. When $p = \infty$, it can be shown that the space (I_∞, d) is still compact, and, if we assume that for all $g \in G_\infty$ and all r, $g_r \leq r$, then so too is (G_∞, d). As Π_∞ is upper semicontinuous on $I \times G$ (the topological product of (I_∞, d) and (G_∞, d)), we deduce from Glicksberg's Theorem that the extended game Γ_∞^* has a solution and there is an optimal strategy for the Infiltrator.

If the speed of the Guard $u = 1$ then we do not know the value of this infinite state game. However, if $u \geq 2$, then the value is zero. The Guard strategy below ensures that the payoff is zero no matter what the Infiltrator does. The crucial factor here is that wherever the Infiltrator is, the Guard can always be sure of 'overtaking' him within a finite amount of time. At time 1 the Guard knows that his opponent must be at state one and so he also moves there. There is a probability μ that detection occurs immediately. At time 2 the Infiltrator may be either at state one or state two, and the Guard chooses one of these states at random. Thus, after time 2 the probability that the Infiltrator has not been detected is precisely $(1 - \mu)(1 - \mu/2)$. Whichever state the Guard moved to at time 2, at time 3 he can still move to state three, which he does. Let $t_1 = 3$. From here on the Guard chooses his moves according to the following rule.

Suppose that at time t_j, $j \geq 1$, the Guard is at state t_j. Therefore he can be sure that the Infiltrator is not any further from state one than he is. He then retraces his steps to state one, choosing at random one of the paths $t_j, t_j - 1, t_j - 2, \ldots, 2, 1$ and $t_j - 1, t_j - 2, \ldots, 2, 1, 1$ each of which takes t_j moves. Assuming that the Infiltrator was not detected before time t_j then, during the period $[t_j, 2t_j]$ he must meet at least one of the above paths. Thus the probability of detection during this period is at least $\frac{1}{2}\mu$. The Guard then moves out to state t_{j+1} where $t_{j+1} = 4t_j - 1$, moving two states forward on each move. Hence he arrives at state t_{j+1} at precisely time $2t_j + \frac{(4t_j-1)-1}{2} = t_{j+1}$. He then repeats the manouevre, retracing his steps again, choosing at random between the paths $t_{j+1}, t_{j+1} - 1, \ldots, 2, 1$ and $t_{j+1} - 1, \ldots, 2, 1, 1$. Assuming that the Infiltrator was not detected before time t_{j+1}, there is again a probability of at least $\frac{1}{2}\mu$ that he is detected during the period $[t_{j+1}, 2t_{j+1}]$. If the Guard continues in this way, by time t_n, $n \in \mathbb{N}$, the probability that the Infiltrator is still undetected is at most $(1 - \mu)\left(1 - \frac{\mu}{2}\right)^n$. As $n \to \infty$ this tends to zero for all $0 < \mu \leq 1$. Therefore detection is guaranteed with proability one. The value of the game is zero.

For other problems with an infinite set of states, it is unclear whether similar detection games are sure to have solutions. The infinite state problem is certainly an interesting one. In the context of evasion it can be related to the problem in which an evader is attempting to escape from some region. He may either try to leave the region by using known escape routes, or he can try and lose himself in the interior. If the region is a large one, it may be appropriate to consider there to be an infinite number of states among which he can move.

The comments we have made concerning the playing space could be repeated in the context of both the detection probability and the initial conditions. The technique of studying the limiting behaviour of a sequence of finite games seems quite robust. We make a few comments about the detection probability before discussing more thoroughly the initial conditions.

We have assumed throughout that there is a constant probability of detection μ, $0 < \mu \leq 1$. It is also possible to make μ depend on the state, the time, and even the past history of the game (so that, for instance, the second time that the Guard looks in a particular state he does so more carefully). Another idea is that the Infiltrator could be detected with positive probability even when the Guard is in another state. In this way, the continuous notion of an exponential detection function could be discretized.

Finally we come to the question of the initial conditions. By this we mean the position of the two players at the start of play. We have assumed that the Infiltrator starts at some known state, whereas the Guard can choose his starting state. However, as we have also assumed that, except in the safe base games, the Guard can start in the same state as the Infiltrator, and so the Guard can always play in such a way that there is a probability of μ that the Infiltrator is detected immediately. Both of these assumptions may be questioned.

In particular, it seems interesting to suggest the following generalisation. Suppose that there is a time period T, $T \geq 0$, at the beginning of the game, during which the Infiltrator is undetectable. Thus, if $T > 0$, he has some time in which to 'get himself lost'. Of course this problem is covered by the general detection function that we suggested above. But let us suppose that, after time T the detection probability is again constant, so the problem is only a relatively slight generalisation of that which we have studied. If $T = 0$, this is the original game.

We can deduce the solution for $T = 1$. This corresponds to the situation in which the Guard can no longer have an attempt at detecting his opponent at state one. It is clear that optimal and ϵ-optimal mixed strategies for the original game are also optimal and ϵ-optimal here. However, the values of all the games, finite or infinite, detection or target, are increased by a factor of $1/\lambda$. This is because the probability of the Infiltrator surviving beyond the first move of the game is now one instead of λ, and otherwise all the probabilities are the same. If $T \geq 2$ then the solution has no such simple relationship to the original problem. For from the second move onwards, the Infiltrator can start to lose himself among the states.

Alternatively, the Guard could be assumed to start from a given state away from the start of the Infiltrator. This would again give the Infiltrator some chance to 'get lost' before the Guard is within range.

Our final point concerns what we describe as the 'computability' of the solutions. One of our main aims has been to demonstrate the use of the finite approximations to Search Games that are obtained by imposing a time limit on play. The strength of this technique is that optimal mixed strategies are known to exist for all finite games. However, for this to be practically useful, it must be feasible to calculate the solutions to these finite games.

In theory, a finite game can be solved using a computer. In practice however, if the number of pure strategies is too large, we are told this is still infeasible. We are unqualified to give any opinions on this question. The point we can make is that the pure strategy sets can sometimes be substantially reduced without changing the value of the game. For instance, we have discussed the importance to the Infiltrator of the pure strategies known as wait-and-run strategies. We have seen that, in the solution given for finite detection games in Chapter 6, the Infiltrator could here restrict himself to mixed strategies over only the wait-and-run strategies, rather than over the whole of I_n. A general question to ask is this. In what circumstances can a restriction like this (to wait-and-run strategies, or any other subset of the pure strategy sets) be made? Note also that, in the infinite games, by considering only wait-and-run strategies, the set of pure strategies is brought down to a countable number. If the pure Guard strategies could be similarly restricted, then these problems could be studied as infinite matrix games.

7.3 In Conclusion

In conclusion then we present this work as a small contribution to the theory of Search Games. We have given a complete solution (the value and two optimal or ϵ-optimal strategies) to a few games of infiltration, most significantly the safe base games of Chapter 5. But there is also included here a collection of more general theorems and observations. Although these are based around the discrete infiltration problem originally suggested by Gal, it is hoped that some of these ideas and conclusions may prove fruitful in stimulating and directing further research in other kinds of discrete search problems.

Bibliography

[1] Alpern S., 1974: The search game with mobile hider on the circle, *Differential games and control theory*, E.O.Roxin, P.T.Liu and R.L.Sternberg (eds.), Dekker, New York, 181–200.

[2] Alpern S. and S.Gal, 1988: A mixed strategy minimax theorem, *SIAM Journal on Control and Optimization*, **26**, 1–5.

[3] Alpern S., 1991: Infiltration games on arbitrary graphs, to appear.

[4] Auger J.M., 1991: On discrete games on infiltration, *Differential Games — Developments in Modelling and Computation*, R.P.Hämäläinen and H.K.Ehtamo (eds.), Springer-Verlag, Berlin, 144–150.

[5] Auger J.M., 1991: An infiltration game on k arcs, *Naval Research Logistics*, **38**, 511–529.

[6] Baston V.J.D. and F.A.Bostock, 1988: An evasion game with barriers, *SIAM Journal on Control and Optimization*, **26**, 1099–1105.

[7] Baston V.J.D. and F.A.Bostock, 1989: A one-dimensional helicopter-submarine game, *Naval Research Logistics*, **36**, 479–490.

[8] Baston V.J.D. and F.A.Bostock, 1990: An evasion game on a finite tree, *SIAM Journal on Control and Optimization*, **28**, 671–677.

[9] Baston V.J.D. and F.A.Bostock, 1991: A basic searchlight game, *Journal of Optimization Theory and Applications*, to appear.

[10] Baston V.J.D. and F.A.Bostock, 1991: An evasion game with an infinite number of states, *SIAM Journal on Control and Optimization*, submitted.

[11] Benkoski S.J., M.G.Monticino and J.R.Weisinger, 1991: A survey of the search theory literature, *Naval Research Logistics*, **38**, 469–494.

[12] Bernhard P., A.-L.Colomb and G.P.Papavassilopoulos, 1987: Rabbit and hunter game: two discrete stochastic formulations, *Computers and Mathematics with Applications*, **13**, 205–225.

[13] Billingsley P., 1979: *Probability and measure*, Wiley, New York.

[14] Blachman N.M., 1959: Prolegomena to optimum discrete search procedures, *Naval Research Logistics Quaterly*, **6**, 273–281.

[15] Blachman N.M and Proschan F., 1959: Optimum search for objects having unknown arrival times, *Operations Research*, **7**, 625–638.

[16] Bourbaki N., 1966: *Elements of mathematics, General topology part 2*, Addison-Wesley, Reading, Massachusetts.

[17] Bram J., 1973: The discrete evasion game, Center for Naval Analysis Research Contributions No. 244.

[18] Charnes A. and Cooper W.W., 1958: The theory of search: optimum distribution of search effort, *Management Science*, **5**, 44–50.

[19] Cheong E., 1988: A helicopter submarine search game, MSc thesis, Naval Postgraduate School, Monterey, California.

[20] Chudnovsky, D.V. and G.V.Chudnovsky, 1989: *Search theory, some recent developments*, Dekker, New York.

[21] Cohn D.L., 1980: *Measure theory*, Birkhauser, Boston.

[22] Danskin J., 1968: A helicopter versus submarine search game, *Operations Research*, **16**, 509–517.

[23] de Guenin J., 1961: Optimum distribution of effort: an extension of the Koopman basic theory, *Operations Research*, **9**, 1–7.

[24] Dobbie J.M., 1968: A survey of search theory, *Operations Research*, **16**, 525–537.

[25] Dieudonné J.A., 1960: *Foundations of modern analysis*, Academic Press, New York.

[26] Dubins L.E., 1957: A discrete evasion game, Contributions to the Theory of Games, Volume III, *Annals of Mathematical Studies*, **39**, 231–255.

[27] Eagle J.N. and A.R.Washburn, 1991: Cumulative search-evasion games, *Naval Research Logistics*, **38**, 495–510.

[28] Ferguson T., 1967: On discrete evasion games with a two-move information lag, *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, **1**, 453–462.

[29] Foreman J.G., 1974: The princess and monster on the circle, *Differential games and control theory*, E.O.Roxin, P.T.Liu and R.L.Sternberg (eds.), Dekker, New York, 231–240.

[30] Gal S., 1979: Search games with mobile and immobile hider, *SIAM Journal on Control and Optimization*, **17**, 99–122.

[31] Gal S., 1980: *Search Games*, Academic Press, New York.

[32] Gal S., 1989: Continuous search games, *Search theory, some recent developments*, D.V. and G.V. Chudnovsky (eds.), Dekker, New York, 33–53.

[33] Garnaev A., 1989: On a discrete game with time-lag in a plane, *Automation and Remote Control*, **50**, 1480–1487.

[34] Glicksberg I.L., 1950: Minimax theorem for upper and lower semi-continuous payoffs, RAND Corporation, Research Memo RM-478, 4pp.

[35] Haley K.B. and L.D.Stone, 1980: *Search theory and applications*, Plenum, New York.

[36] Halmos P.R., 1950: *Measure theory*, Van Nostrand, New York.

[37] Harary F., 1969: *Graph theory*, Addison-Wesley, Reading, Massachusetts.

[38] Isaacs R. and S.Karlin, 1954: A game of aiming and evasion, RAND Memorandum No. RM-1316.

[39] Isaacs R., 1955: The problem of aiming and evasion, *Naval Research Logistics Quarterly*, **2**, 47-67.

[40] Isaacs R., 1965: *Differential Games*, Wiley, New York.

[41] Karlin S., 1957: An infinite move game with a lag, Contributions to the Theory of Games, Volume III, *Annals of Mathematical Studies*, **39**, 257–272.

[42] Kneser H., 1952: Sur un théorème fondamental de la théorie des jeux, *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences Serie A*, **234**, 2418–2420.

[43] Koopman B.O., 1946: Search and screening, OEG report no. 56, The Summary Reports Group of the Columbia University Division of War Research.

[44] Koopman B.O., 1956: The theory of search, part I. Kinematic bases, *Operations Research*, **4**, 324–346.

[45] Koopman B.O., 1956: The theory of search, part II. Target detection, *Operations Research*, **4**, 503–531.

[46] Koopman B.O., 1957: The theory of search, part III. The optimum distribution of searching effort, *Operations Research*, **5**, 613–626.

[47] Lalley S.P., 1988: A one-dimensional infiltration game, *Naval Research Logistics*, **35**, 441-44

[48] Lee K.T., 1983: An evasion game with a destination, *Journal of Optimization Theory and Applications*, **41**, 359–372.

[49] Lee K.T., 1983: A firing game with a time lag, *Journal of Optimization Theory and Applications*, **41**, 547–558.

[50] Lee K.T. and L.Y.Lee, 1991: On a discrete game of Isaacs, *Journal of Optimization Theory and Applications*, to appear.

[51] Lee L.Y. and K.T.Lee, 1990: The discrete evasion game with three-move lag, *Naval Research Logistics*, **37**, 867–874.

[52] Matula D.W., 1966: Games of sequence prediction, Report No. ORC 66-3, Operations Research Center, University of California, Berkeley.

[53] Meinardi J., 1964: *Theory of games*, A. Mensch (ed.), Elsevier, New York, 285–299.

[54] Neumann J. von and O.Morgenstern, 1947 (first edition 1944): *Theory of games and economic behavior*, Princeton University Press, Princeton.

[55] Olsder G.J. and Papavassilopoulos G.P., 1986: On a finite state space pursuit evasion game with dynamic information, *Proceedings of the 25th International Conference on decision and control.*

[56] Olsder G.J. and Papavassilopoulos G.P., 1988: about when to use a searchlight, *Journal of Mathematical Analysis and Applications*, **136**, 466–478.

[57] Olsder G.J. and Papavassilopoulos G.P., 1988: A markov chain with dynamic information, *Journal of Optimization Theory and Applications*, **59**, 467–486.

[58] Parthasarathy K.R., 1967: *Probability measures on metric spaces*, Academic Press, New York.

[59] Ruckle, W.H., 1983: *Geometric games and their applications*, Pitman, Boston.

[60] Sakaguchi M., 1986: A pursuit game on a line with a bunker for the evader, *Mathematica Japonica*, **31**, 449–461.

[61] Stone L.D., 1989: *Theory of optimal search*, 2nd edition, Operations Research Society of America, Arlington, VA.

[62] Stone L.D. and A.R.Washburn, 1991: Introduction, special issue on Search Theory, *Naval Research Logistics*, **38**, 465–468.

[63] Strumpfer J., 1980: Search theory index, Technical note no. TN-017-80, Institute for Maritime Technology, Simonstown, South Africa.

[64] Thomas L.C., 1984: *Games, theory and applications*, Wiley, New York.

[65] Thomas L.C. and A.R.Washburn, 1991: Dynamic search games, *Operations Research*, **39**, 415–422.

[66] Washburn A.R., 1981: *Search and detection*, Operations Research Society of America, Rosslyn, VA.

[67] Wilson D.J., 1972: Isaac's princess and monster game on the circle, *Journal of Optimization Theory and Applications*, **9**, 265–288.

[68] Zelikin M.I., 1972: On a differential game with incomplete information, *Soviet Mathematics Doklady*, **13**, 228–231.