

The Tourist Game

Marcus Pendergrass

June 9, 2007

1 The Tourist Game

There are two players, the Tourist and the Guide. Both are located at the same node in a strongly connected directed graph. The Guide is going to choose one of the available out-edges, and take the Tourist to the next node. But before he does so, the Tourist places a bet on the edge she thinks the Guide will select. If the Tourist's guess is correct, she wins, and receives a payoff proportional to her bet. The proportionality constants, or *stakes*, are chosen in a natural and fair way (see below). If she is incorrect, she loses the amount of her bet. The Guide then takes the Tourist to the next node (via the edge he chose), and the game repeats from there. The Tourist's goal is to maximize her expected fortune over the long haul. The Guide's goal is to minimize the Tourist's expected fortune.

Regarding the stakes of play in this game: suppose the players are at a vertex i whose out-degree is $n \geq 3$. If the game is played at even stakes, then the Tourist has no incentive to wager *anything* on the Guide's next move, because her chance of guessing the move correctly is less than one half (unless the Guide is playing the game very badly). To give the Tourist incentive to play, the game should be designed so that the Tourist's payoff is inversely proportional to her odds of winning, in the average case. In the average case here (i.e. averaging over all possible strategies), each of the n available destinations is equally likely to be chosen by the Guide. Therefore at this node the Tourist's odds of winning are 1 to $n - 1$ in the average case. Since she loses $n - 1$ times for every time she wins at this node, her payoff when she wins should be $n - 1$ times her bet, while when she loses she should forfeit just the amount of her bet. Note that playing with these stakes makes the Tourist Game *fair* in the average case (i.e. when each destination node is equally likely to be chosen by the Guide), in the following sense: after this step of the game, the Tourist's fortune on average is unchanged, no matter what she wagered.

2 The One Step Game

First look at the problem faced by the players at a single stage of the game. They are at some node of a directed graph, call it 0. From this node, there are n possible nodes j which they can reach by traversing a single edge. (The

case $n = 1$ is trivial, so we will focus on $n \geq 2$.) Let us assume that somehow these n possible destination nodes j have positive *values* α_j attached to them, with the following meaning: the value α_j of node j is the expected *payoff factor* associated with j . That is, if the Tourist enters node j with a fortune of F , then her expected fortune at the end of the game (however that is defined) is $\alpha_j F$. Assuming for the moment that such values α_j are defined - a subject we shall return to in Section 3 - we have the following

Definition 2.1 (The One Step Game). The Tourist and the Guide are at the top node in the diagram in Figure 2 below, and the Tourist has one dollar in her pocket. The Guide is going to choose one of the $n \geq 2$ available destinations. The Tourist bets on the branch that the Guide is going to select. If her guess is correct, her payoff is $n - 1$ times her bet, while if she is incorrect, she forfeits her bet. (See the discussion in Section 1 above.) Win or lose, her entire fortune is then multiplied by the payoff factor α_j corresponding to the branch chosen by the Guide. The Tourist's goal is to maximize her expected fortune. The Guide's goal is to minimize the Tourist's expected fortune.




Figure 1: The One Step Game

In this section we will analyze the one step game.

Strategies for the one step game. Consider first the pure (i.e. deterministic) strategies available to the players. For the Guide, there are n possible pure strategies, corresponding to the n available branches. Let \mathcal{G}_j denote the Guide's strategy of picking the branch leading to node j :

$$\mathcal{G}_j : \text{Guide chooses the branch leading to node } j, \quad (1)$$

where $j = 1, 2, \dots, n$. The Tourist's pure strategies are

$$\mathcal{T}_j : \text{wager } w_j \text{ on the Guide choosing the branch leading to node } j, \quad (2)$$

where again $j = 1, 2, \dots, n$. Notice that we are considering the wager as part of the Tourist's pure strategies, and allowing for a different wager for each pure

strategy. So technically, the Tourist has uncountably many pure strategies. The mixed (i.e. random) strategies for the Guide are probability vectors $p = (p_j, j = 1, 2, \dots, n)$, where

$$p_j = \text{probability that the Guide chooses pure strategy } \mathcal{G}_j. \quad (3)$$

The mixed strategies for the Tourist are specified by a *pair* of n -vectors (q, w) , where

$$q_j = \text{probability that the Tourist chooses pure strategy } \mathcal{T}_j, \quad (4)$$

and

$$w_j = \text{wager that the Tourist makes on the Guide choosing branch } j. \quad (5)$$

Let F be the fortune of the Tourist at the end of the one step game (i.e. after the payoff factors α_j have been applied). The Tourist's initial fortune is one dollar, so

$$E[F | \mathcal{T}_k, \mathcal{G}_j] = \begin{cases} (1 + (n-1)w_j)\alpha_j & \text{if } k = j; \\ (1 - w_k)\alpha_j & \text{if } k \neq j; \end{cases} \quad (6)$$

Thus, if the Tourist plays pure strategy \mathcal{T}_k , her expected fortune is

$$E[F | \mathcal{T}_k] = (1 + (n-1)w_k)\alpha_k p_k + (1 - w_k) \sum_{j:j \neq k} \alpha_j p_j \quad (7)$$

$$= \alpha \cdot p + w_k (n\alpha_k p_k - \alpha \cdot p), \quad (8)$$

where $\alpha = (\alpha_k : k = 1, 2, \dots, n)$, and $\alpha \cdot p$ is the dot product of α and p . Similarly, if the Guide plays his pure strategy \mathcal{G}_j , then the Tourist's expected fortune is

$$E[F | \mathcal{G}_j] = (1 + (n-1)w_j)\alpha_j q_j + \alpha_j \sum_{k:k \neq j} (1 - w_k) q_k \quad (9)$$

$$= \alpha_j (1 + nw_j q_j - w \cdot q), \quad (10)$$

where $w = (w_j : j = 1, 2, \dots, n)$. The unconditional expectation of the Tourist's fortune can be computed from either (6), (8), or (10):

$$\begin{aligned} E[F] &= \sum_k \sum_j E[F | \mathcal{T}_k, \mathcal{G}_j] q_k p_j \\ &= \sum_k E[F | \mathcal{T}_k] q_k \end{aligned} \quad (11)$$

$$= \sum_j E[F | \mathcal{G}_j] p_j \quad (12)$$

Optimal play in the one step game. The Guide wants to minimize $E[F]$, the Tourist's expected fortune. Thus he should visit high payoff nodes rarely. Consider the strategy

$$p_j = \frac{\alpha_j^{-1}}{\alpha_1^{-1} + \alpha_2^{-1} + \dots + \alpha_n^{-1}}. \quad (13)$$

In other words, the Guide chooses node j with a probability that is inversely proportional to the value of node j . Then $\alpha_j p_j = 1/\sum_k \alpha_k^{-1}$ for all j , and so $\alpha \cdot p = H$, where H is the *harmonic mean* of the values α_k :

$$H \equiv \frac{n}{\alpha_1^{-1} + \alpha_2^{-1} + \dots + \alpha_n^{-1}} \quad (14)$$

Thus, by (8) when the Guide uses this strategy we have

$$E[F | \mathcal{T}_k] = H + w_k \left(n \frac{H}{n} - H \right) = H \quad (15)$$

no matter which pure strategy \mathcal{T}_k the Tourist picks (and no matter what her wager is). Therefore by (11), the Guide can always *limit* the Tourist's expected fortune to H .

Moreover, I claim that for any Guide strategy other than the one in (13), there exists a Tourist strategy (q, w) such that $E[F] > H$. To see this, let $\hat{p} = (\hat{p}_j)$ be any strategy for the Guide other than the one given by (13). Then there is at least one index k with

$$\hat{p}_k > \frac{\alpha_k^{-1}}{\alpha_1^{-1} + \alpha_2^{-1} + \dots + \alpha_n^{-1}}.$$

For such a k , the Tourist can use her pure strategy \mathcal{T}_k , with a wager of $w_k = 1$, and from (7) her expected fortune will be

$$E[F | \mathcal{T}_k] = n\alpha_k \hat{p}_k > H,$$

proving the claim.

Conversely, I will now show that the Tourist can always *achieve* an expected fortune of H , no matter which strategy the Guide adopts. By (10), a strategy (q, w) satisfies $E[F | \mathcal{G}_j] = H$ for all j if and only if

$$\alpha_j (1 + nw_j q_j - w \cdot q) = H \quad (16)$$

for all $j = 1, 2, \dots, n$. Equations (16) can be written as the linear system

$$\begin{pmatrix} n-1 & -1 & -1 & \dots & -1 \\ -1 & n-1 & -1 & \dots & -1 \\ -1 & -1 & n-1 & \dots & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & -1 & \dots & n-1 \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \\ z_3 \\ \vdots \\ z_n \end{pmatrix} = \begin{pmatrix} \alpha_1^{-1} H - 1 \\ \alpha_2^{-1} H - 1 \\ \alpha_3^{-1} H - 1 \\ \vdots \\ \alpha_n^{-1} H - 1 \end{pmatrix}, \quad (17)$$

where $z_j = w_j q_j$. The column sums of the $n \times n$ matrix in (17) are all zero, implying that the matrix is singular, and that its range is the set of all n -vectors whose components sum to zero. Fortunately the right side of (17) is just such a vector. Therefore there will be a one-parameter solution family of (17). Solving (17) gives

$$z_j = t + n^{-1}H(\alpha_j^{-1} - \alpha_n^{-1}) \quad (18)$$

Assume (without loss of generality) that $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$. While t is an arbitrary parameter in (19), the fact that $z_j = w_j q_j$, where $0 \leq w_j \leq 1$ and (q_j) is a probability vector, will place limitations on t . In order for $z_j = w_j q_j$ for all j , it is necessary and sufficient that $z_j \geq 0$ for all j and $0 \leq \sum_j z_j \leq 1$. These conditions are satisfied precisely when $0 \leq t \leq \alpha_n^{-1}H/n$. Writing $t = \rho \alpha_n^{-1}H/n$, where ρ is arbitrary in $[0, 1]$, and remembering that $n^{-1}H\alpha_j^{-1} = p_j$, equation (18) becomes

$$z_j = n^{-1}H(\alpha_j^{-1} - (1 - \rho)\alpha_n^{-1}) \quad (19)$$

$$= p_j - (1 - \rho)p_n \quad (20)$$

This general solution of (17) generates optimal strategies (q, w) for the Tourist as follows: pick any probability vector q such that $q_j \geq z_j$ for all j , and set $w_j = z_j/q_j$ if $q_j > 0$. (If $q_j = 0$, then w_j can be arbitrary in $[0, 1]$.)

Let \mathbb{T} be the set of all Tourist strategies (q, w) satisfying (19) (where $z_j = q_j w_j$). We have just seen that if $(q, w) \in \mathbb{T}$, then $E[F] = H$ no matter what strategy p the Guide adopts. We now claim that for any Tourist strategy (q, w) that is *not* in \mathbb{T} , there exists a Guide strategy p such that $E[F] < H$. To see this, note that equation (10) implies that for *any* Tourist strategy (q, w) we have

$$\sum_{j=1}^n \alpha_j^{-1} E[F | \mathcal{G}_j] = n. \quad (21)$$

Each of the expected values in (21) is equal to H if and only if $(q, w) \in \mathbb{T}$. Therefore, if $(q, w) \notin \mathbb{T}$, then some of the expected values in (21) must be greater than H , and some must be less than H . For such (q, w) , the Guide can choose a strategy p concentrated on the j satisfying $E[F | \mathcal{G}_j] < H$, and thereby force $E[F] < H$. We have proven the following

Theorem 2.1 (Optimal Play In the One Step Game). For the one step game of Definition 2.1, with $n \geq 2$ branches and values $\alpha_n \geq \alpha_{n-1} \geq \dots \geq \alpha_1 > 0$, the optimal Guide strategy p is given by

$$p_j = \frac{\alpha_j^{-1}}{\sum_{\ell} \alpha_{\ell}^{-1}} = \frac{1}{n} \alpha_j^{-1} H, \quad (22)$$

where H is the harmonic mean of the values $(\alpha_j : 1 \leq j \leq n)$. There are multiple optimal Tourist strategies (q, w) , which are characterized by

$$q_j w_j = p_j - (1 - \rho)p_n \quad (23)$$

where ρ is arbitrary in $[0, 1]$. Under the optimal strategies, the Tourist's expected fortune at the end of the game is $E[F] = H$.

Corollary 2.1 (Rescaling Properties). In the one step game:

1. If the values $\alpha_j, j = 1, 2, \dots, n$ are all multiplied by a positive constant r , then the optimal strategies remain unchanged. The Tourist's expected fortune is multiplied by r .
2. If the initial fortune is scaled by a factor s , then the optimal strategies remain unchanged, provided that the wagers w_k are interpreted as the fraction of the Tourist's fortune that is wagered. The Tourist's expected fortune is multiplied by s .

Proof of Corollary 2.1. Equations (22) and (23) show that the optimal strategies for both the Tourist and the Guide are functions of the ratios of the α -values, from which the first statement follows. The second statement is trivial. \square

Corollary 2.2 (Average Wager). If the Tourist is playing optimally, then her average wager must be at least $1 - \alpha_n^{-1}H$.

Proof of Corollary 2.2. The average wager is $\sum_{j=1}^n q_j w_j$, which by (23) equals $1 - (1 - \rho) \alpha_n^{-1}H \geq 1 - \alpha_n^{-1}H$, since $\rho \in [0, 1]$. \square

In light of Corollary 2.2, I will refer to $1 - \alpha_n^{-1}H$ as the *critical wager*.

Remark 2.1. By Theorem 2.1, if the Tourist starts the one step game with a fortune of F , then her fortune on average at the end of the game will be equal to HF , where H is the harmonic mean of the values α_j . But this means that the value of node 0 is H , by our definition of value. So value “propagates” in the one step game from the leaf nodes back to the root node, via the harmonic mean. It is more convenient to keep track of the reciprocals of the values, since these propagate from leaves to root via the arithmetic mean, which is a linear operation. See Figure 2.

Example 2.1 (A simple minimum risk optimal strategy for the Tourist). Suppose the Tourist wants to “keep it simple”, minimize her risk, and yet still play optimally. She can start by choosing to wager the same amount regardless of which pure strategy she bets on: $w_j = w$ for all j . This forces $\sum_j z_j = w$. In order to risk as little of her fortune as possible, she will minimize the sum of the z_j 's by choosing $\rho = 0$ in (23), which makes her wager $w = 1 - \alpha_n^{-1}H$. (Note that this is always strictly less than 1, and is positive unless all the α -values are the same.) The associated probability vector can now be calculated via (23):

$$q_j = \frac{p_j - p_n}{1 - \alpha_n^{-1}H} \tag{24}$$

for $j = 1, 2, \dots, n$, where the p_j 's are the Guide's optimal transition probabilities.




Figure 2: Reciprocal value propagates from leaves back to the root via the arithmetic mean.

Example 2.2 (The maximum risk optimal strategy for the Tourist). If the Tourist chooses $w_j = 1$ for all j - in other words, she risks everything on her choice with probability 1 - then by the $\rho = 1$ case of (23) she can still achieve the optimal expected fortune of H by choosing $q_j = p_j$ for all j , where $p = (p_j)$ is the optimal Guide strategy. However, she stands a good chance of losing her entire fortune with strategy as well.

Example 2.3 (The $n = 2$ case). When $n = 2$ the stakes are even: $n - 1 = 1$. The optimal strategy for the Guide is

$$p_1 = \frac{\alpha_1^{-1}}{\alpha_1^{-1} + \alpha_2^{-1}} = \frac{\alpha_2}{\alpha_1 + \alpha_2},$$

$$p_2 = \frac{\alpha_2^{-1}}{\alpha_1^{-1} + \alpha_2^{-1}} = \frac{\alpha_1}{\alpha_1 + \alpha_2}.$$

To examine optimal strategies for the Tourist, note that the z -vector of (18) is given by

$$z_1 = t + \frac{\alpha_1^{-1} - \alpha_2^{-1}}{\alpha_1^{-1} + \alpha_2^{-1}} = t + \frac{\alpha_2 - \alpha_1}{\alpha_1 + \alpha_2}$$

$$z_2 = t,$$

where $0 \leq t \leq \alpha_2^{-1} / (\alpha_1^{-1} + \alpha_2^{-1}) = \alpha_1 / (\alpha_1 + \alpha_2)$. We can construct optimal strategies for the Tourist by taking any probability vector q with $q_k \geq z_k$ for $k = 1, 2$, and then setting $w_k = z_k / q_k$. A few such strategies are:

1. Take $t = 0$, $q_1 = 1$, and $q_2 = 0$. Then $w_1 = (\alpha_2 - \alpha_1) / (\alpha_1 + \alpha_2)$, and w_2 can be arbitrary in $[0, 1]$. This is the simple minimum risk strategy of Example 2.1, and also the so-called “critical wager” in the Lying Oracle game (see [3]).

2. Take $t = \alpha_1 / (\alpha_1 + \alpha_2)$. Then $z_1 = \alpha_2 / (\alpha_1 + \alpha_2)$, $z_2 = \alpha_1 / (\alpha_1 + \alpha_2)$. Let $q_1 = z_1 = \alpha_2 / (\alpha_1 + \alpha_2)$ and $q_2 = z_2 = \alpha_1 / (\alpha_1 + \alpha_2)$. Then $w_1 = w_2 = 1$. This is the maximum risk optimal Tourist strategy of Example 2.2: there is a good chance the Tourist will lose her entire fortune by playing this strategy, even though it is optimal in the sense of expected value.
3. Take $t = 0$, $q_1 = z_1 = (\alpha_2 - \alpha_1) / (\alpha_1 + \alpha_2)$, and $q_2 = 2\alpha_1 / (\alpha_1 + \alpha_2)$. Then $w_1 = 1$ and $w_2 = 0$. This is an example of an optimal strategy for the Lying Oracle game in which the player wagers an amount that is less than the so-called “critical wager”, at least part of the time. Note, however, that *on average*, the player is wagering the critical amount. (See Corollary 2.2.)

3 Playing the Tourist Game for a Finite Number of Steps

In this section I will illustrate how to use Theorem 2.1 to solve the *finite duration* Tourist Game. The solution technique is general, but I will use the specific graph G in Figure 3 to illustrate it. I will focus on finding the appropriate α -values, since the optimal strategies are all functions of these.




Figure 3: The Graph G

Suppose we want to play the game for just one step. Depending on whether we are initially at vertex 1, 2, or 3, we simply have a one step game on the corresponding *path tree* of length 1.




Figure 4: The Path Trees of Length 1

The leaves of each path tree have their α -values all equal to 1. The α -values for the roots of the path trees can now be computed. If the out-degree of the root is at least two, then by Theorem 2.1 the α -value of the root is the harmonic mean of the α -values of the leaves. If the out-degree of the root is one, then the α -value of the root is two. So if we play the Tourist Game on G for $m = 1$ step, the α -values for the vertices of G are

$$\alpha_{1,1} = 2, \quad \alpha_{2,1} = 1, \quad \alpha_{3,1} = 1 \tag{25}$$

(The second subscript indicates the duration of the game.) So if we play for just one step, by Theorem 2.1 the optimal strategy for the Guide is to transition from vertex i to vertex j with probability

$$p_{i,j} = \begin{cases} \alpha_{j,1}^{-1} / \sum_{k:i \rightarrow k} \alpha_{k,1}^{-1} & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \tag{26}$$

Here $i \rightarrow j$ means that there is a directed edge from node i to j in G . This gives us the Guide's *one step optimal transition matrix* $P_1 = (p_{i,j} : i, j \in V)$

$$P_1 = \begin{pmatrix} 0 & 1 & 0 \\ 1/3 & 0 & 2/3 \\ 1/5 & 2/5 & 2/5 \end{pmatrix} \quad (27)$$

Similarly, the Tourist's optimal strategies can be computed from the α -values in (25).

As I remarked earlier, it is more convenient to keep track of the reciprocals of the values, since these propagate from leaves to root via the arithmetic mean, which is a linear operation. The reciprocal α -values can be calculated all at once via a matrix multiply:

$$\begin{pmatrix} 1/2 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 & 1/2 & 0 \\ 1/2 & 0 & 1/2 \\ 1/3 & 1/3 & 1/3 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \quad (28)$$

I will refer to the matrix in this equation as the *propagation matrix*, and will denote it by $M = (m_{i,j} : i, j \in V)$. For a general digraph G , it is defined by

$$m_{i,j} = \begin{cases} 1/d(i) & \text{if } d(i) \geq 2 \text{ and } i \rightarrow j \\ 1/2 & \text{if } d(i) = 1 \text{ and } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

Here, $d(i)$ is the out-degree of node i .

Now consider playing the game for $m = 2$ steps. We have the following path trees of length 2:




Figure 5: The Path Trees of Length 2

But the circled trees just below the roots are all path trees of length 1, whose α -values we have already calculated. Therefore the reciprocal α -values for the first step of the two step game are

$$\begin{pmatrix} 0 & 1/2 & 0 \\ 1/2 & 0 & 1/2 \\ 1/3 & 1/3 & 1/3 \end{pmatrix} \begin{pmatrix} 1/2 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1/2 \\ 3/4 \\ 5/6 \end{pmatrix}, \quad (30)$$

or from (28)

$$\begin{pmatrix} 0 & 1/2 & 0 \\ 1/2 & 0 & 1/2 \\ 1/3 & 1/3 & 1/3 \end{pmatrix}^2 \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1/2 \\ 3/4 \\ 5/6 \end{pmatrix}. \quad (31)$$

So the optimal strategy for the Guide on the first step of the two step game is given by the *two step transition matrix*

$$P_2 = \begin{pmatrix} 0 & 1 & 0 \\ \frac{1/2}{1/2+5/6} & 0 & \frac{5/6}{1/2+5/6} \\ \frac{1/2}{1/2+3/4+5/6} & \frac{3/4}{1/2+3/4+5/6} & \frac{5/6}{1/2+3/4+5/6} \end{pmatrix} \quad (32)$$

$$= \begin{pmatrix} 0 & 1 & 0 \\ 3/8 & 0 & 5/8 \\ 6/25 & 9/25 & 10/25 \end{pmatrix} \quad (33)$$

Inductively, reciprocal α -values for the first step of the m -step game can be computed in the same way, via powers of the propagation matrix M :

$$\alpha_m^{-1} = M^m \mathbf{1}, \quad (34)$$

where $\alpha_m^{-1} = (\alpha_{i,m}^{-1} : i \in V)$ is the vector of reciprocal α -values for the first step of the m step game, M is the propagation matrix defined by (29), and $\mathbf{1}$ is the vector of all ones.

What about the Tourist's fortune at the end of the m step game? This depends on the vertex at which the players start the game. Let

$$X_m = \text{position of the players at step } m, \quad (35)$$

where $m = 0, 1, 2, \dots$, and

$$F_m = \text{fortune of the Tourist at step } m. \quad (36)$$

(We always assume $F_0 = 1$.) If both players are playing optimally, the m step game is equivalent to a one step game with α -values given by (34). Therefore by Theorem 2.1

$$E[F_m | X_0 = i] = \alpha_{i,m} \quad (37)$$

Theorem 3.1 (Optimal Play in the m Step Game). The optimal strategies for the first step of the m -step Tourist Game are the same as the optimal strategies for the one step game with reciprocal α -values given by

$$\alpha_m^{-1} = M^m \mathbf{1}, \quad (38)$$

where M is the propagation matrix defined by (29), and $\mathbf{1}$ is the vector of all ones. The optimal Guide and Tourist strategies are calculated from the α -values in (38) using Theorem 2.1. If both players are playing optimally, the Tourist's expected fortune at the end of the m step game is given by

$$E[F_m | X_0 = i] = \alpha_{i,m} \quad (39)$$

4 Limiting Strategies and the Discounted Game

Theorem 3.1 shows that the optimal strategies on the first move of an m step game depend on m . What happens in the limit as m goes to infinity? The answer turns out to depend on some special properties of the propagation matrix M of equation (29). Obviously the propagation matrix is nonnegative. Since the graph G is strongly connected, it is also *irreducible* (see [1], Definition 2.1 page 5, and Theorem 3.2 page 78). Such matrices have many beautiful properties. For our purposes, the relevant ones are collected here in

Theorem 4.1 (Maximal Eigenvalue of a Irreducible Nonnegative Matrix). Let A be an irreducible nonnegative square matrix. Then

- A has a positive *maximal eigenvalue* r , with the property that any other eigenvalue λ of A satisfies $|\lambda| < r$.
- r is at least as big as the minimum row (or column) sum of A , and no bigger than the maximum row (or column) sum of A .
- There is a positive eigenvector x associated with r .
- The eigenspace associated with r (and containing x) has dimension 1.
- No other eigenvector of A is positive.

For proofs of these statements, see Chapter 1 of [1].

Theorem 4.1 applies to the matrix M of (29). The row sums of M are all either 1 or $1/2$, so the maximal eigenvalue r of M lies between $1/2$ and 1. I will first show that $\lim_{m \rightarrow \infty} r^{-m} M^m$ exists and is non-zero. Since the eigenspace of M associated with r is of dimension 1, the Jordan canonical form (see [2]) of M can be written in the form of the following partitioned matrix:

$$J = \left[\begin{array}{c|c} r & 0 \\ \hline 0 & \tilde{J} \end{array} \right] \quad (40)$$

Here J is $n \times n$, where n is the number of vertices in G , and \tilde{J} is the $(n-1) \times (n-1)$ block diagonal matrix consisting of the Jordan blocks associated with the other eigenvalues of M . Thus each block of \tilde{J} is upper triangular, with an eigenvalue along the diagonal, ones along the super-diagonal, and zeros elsewhere. So $M = PJP^{-1}$, where P is a nonsingular change of basis matrix. It follows that

$$r^{-m} M^m = P \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & r^{-m} \tilde{J}^m \end{array} \right] P^{-1} \quad (41)$$

Now $r^{-1} \tilde{J}$ is an upper triangular matrix whose diagonal entries are all less than 1 in absolute value. The m^{th} power of any such matrix approaches the zero matrix as m approaches infinity. (This isn't too hard to see.) So

$$\lim_{m \rightarrow \infty} r^{-m} M^m = P \left[\begin{array}{c|c} 1 & 0 \\ \hline 0 & 0 \end{array} \right] P^{-1} = xy^T \quad (42)$$

where x is the first column of P and y^T is the first row of P^{-1} . To see that the limit is non-zero, note that (40) along with $MP = PJ$ implies that the first column of P is an eigenvector of M corresponding to r , while $P^{-1}M = JP^{-1}$ implies that the first row of P^{-1} is an eigenvector of M^T corresponding to r . (Equivalently, y is a *left* eigenvector of M corresponding to r .) Thus Theorem 4.1 guarantees that the entries of x and y are all strictly positive. Therefore the limiting matrix is not only non-zero, but actually positive.

For future reference, I want to point out that the limiting matrix in (42) can be calculated without having to find the full Jordan basis matrix P . Indeed, *any* pair of right and left eigenvectors x and y of M will do the job. All we need to do is normalize x and y so that they satisfy $x^T y = 1$, which insures that they can act as the first column and row of P and P^{-1} . Thus we can write

$$\lim_{m \rightarrow \infty} r^{-m} M^m = \frac{xy^T}{x^T y} \quad (43)$$

where x and y are *any* right and left eigenvectors of M corresponding to r .

What does all this have to do with limiting strategies? Consider the *discounted game* on G , played as follows: after the payoff at each step of the Tourist Game, we discount the Tourist's fortune by some *discount rate* $\hat{r} \in (0, 1]$. Thus, if the Tourist has fortune F_m after the payoff from stage m of the game, we "tax" her at rate $1 - \hat{r}$, and she begins stage $m + 1$ of the game with a fortune of $\hat{r}F_m$. For the one step game of Section 2, this is obviously just a rescaling of the α -values by the factor \hat{r} , which by Corollary 2.1 leaves the optimal strategies for both players unchanged. Similarly, the α -values for the first step of discounted m step game are given by $\hat{r}^m \alpha_m$, where α_m is given by (38), so that the optimal strategies for the finite duration Tourist Game are unaffected by discounting as well. If we set the discount rate equal to the maximal eigenvalue of the propagation matrix M ,

$$\hat{r} = r,$$

then the vector α of limiting α -values for the discounted game is given by

$$\alpha^{-1} \equiv \lim_{m \rightarrow \infty} r^{-m} M^m \mathbf{1} \quad (44)$$

$$= \frac{xy^T}{x^T y} \mathbf{1}. \quad (45)$$

Moreover, the limiting vector α^{-1} is itself an eigenvector of M corresponding to the eigenvalue r :

$$M\alpha^{-1} = M \lim_{m \rightarrow \infty} r^{-m} M^m \mathbf{1} \quad (46)$$

$$= \lim_{m \rightarrow \infty} r^{-m} M^{m+1} \mathbf{1} \quad (47)$$

$$= r \lim_{m \rightarrow \infty} r^{-(m+1)} M^{m+1} \mathbf{1} \quad (48)$$

$$= r\alpha^{-1} \quad (49)$$

The limiting strategies for the finite duration Tourist game - either discounted or not - can be calculated from the vector α using Theorem 2.1. As far as the Guide is concerned, his *limiting transition matrix* is given by

$$p_{i,j} = \begin{cases} \alpha_j^{-1} / \sum_{k:i \rightarrow k} \alpha_k^{-1} & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (50)$$

It is perhaps almost obvious that this limiting strategy is in fact optimal for the infinite duration Tourist Game.

As for the Tourist, by Theorem 3.1 her limiting strategies at nodes of out-degree two or more are characterized by

$$q_{i,j} w_{i,j} = p_{i,j} - (1 - \rho) p_{i,\min}, \quad (51)$$

where $p_{i,\min} = \min\{p_{i,k} : k \in V, \text{ and } i \rightarrow k \text{ in } G\}$, ρ is arbitrary in $[0, 1]$, and the $p_{i,j}$ are given by (50). (At nodes of out-degree one, the Tourist obviously bets everything on the destination node.) Again, it is fairly clear that these limiting strategies are in fact optimal for the Tourist in the infinite duration Tourist Game, provided that $\rho < 1$. (If the Tourist chooses $\rho = 1$, then with probability one she will eventually go bankrupt in the infinite duration game.)

In the m -step game, the Tourist's expected fortune under optimal play satisfies

$$E[F_m | X_0 = i] = \alpha_{i,m}, \quad (52)$$

where the $\alpha_{i,m}$ are given by (34). But the vector $r^{-m} \alpha_m^{-1}$ converges to the limiting vector α^{-1} of (44) as m goes to infinity. Therefore

$$\lim_{m \rightarrow \infty} r^m E[F_m | X_0 = i] = \lim_{m \rightarrow \infty} r^m \alpha_{i,m} \quad (53)$$

$$= \alpha_i \in (0, \infty), \quad (54)$$

where the α_i are the reciprocals of the components of the limiting vector α^{-1} of (44). But $r^m F_m$ is the Tourist's discounted fortune at the end of the m -step game. Hence, α_i represents the limiting discounted fortune at the end of the m -step game. It follows that in the undiscounted game, the expected value of the Tourist's fortune F_m grows exponentially, as r^{-m} . (Recall that $1/2 \leq r \leq 1$.) Moreover, in the discounted game, the Tourist's expected fortune remains bounded and nonzero for all time precisely when the discount rate is $\hat{r} = r$. Lower values of \hat{r} force the Tourist's expected fortune to 0, while higher values of \hat{r} force the expected fortune to infinity.

Example 4.1 (Limiting Strategies). Go back to the graph of Figure 3 with propagation matrix

$$M = \begin{pmatrix} 0 & 1/2 & 0 \\ 1/2 & 0 & 1/2 \\ 1/3 & 1/3 & 1/3 \end{pmatrix}.$$

All the column sums of M are $5/6$, implying that its maximal eigenvalue is $r = 5/6$. Calculating right and left eigenvectors corresponding to r yields

$$x = (9, 15, 16)^T \quad \text{and} \quad y = (1, 1, 1)^T, \quad (55)$$

so we get

$$\lim_{m \rightarrow \infty} r^{-m} M^m = \frac{xy^\top}{x^\top y} = \begin{pmatrix} 9/40 & 9/40 & 9/40 \\ 3/8 & 3/8 & 3/8 \\ 2/5 & 2/5 & 2/5 \end{pmatrix}. \quad (56)$$

The limiting α -vector is given by

$$\alpha^{-1} = \begin{pmatrix} 9/40 & 9/40 & 9/40 \\ 3/8 & 3/8 & 3/8 \\ 2/5 & 2/5 & 2/5 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \frac{1}{40} \begin{pmatrix} 27 \\ 45 \\ 48 \end{pmatrix}, \quad (57)$$

or

$$\alpha = (40/27, 8/9, 5/6)^\top \approx (1.481, 0.889, 0.833)^\top. \quad (58)$$

Now we can calculate the limiting strategy for the Guide in the discounted game; by (50) it is given by the transition matrix

$$P = \begin{pmatrix} 0 & 1 & 0 \\ 9/25 & 0 & 16/25 \\ 9/40 & 15/40 & 16/40 \end{pmatrix} \quad (59)$$

Since optimal strategies are the same for the discounted and undiscounted games, this is also the limiting Guide strategy in the undiscounted game.

For the Tourist, the limiting strategies are characterized by $q_{i,j} w_{i,j} = p_{i,j} - (1 - \rho) p_{i,\min}$. For simplicity, let's assume that the Tourist's wager at node i is the same for every possible destination: $w_{i,j} \equiv w_i$ for all $j \in V$ such that $i \rightarrow j$. Then the Tourist's optimal guessing probabilities are given by

$$Q = \begin{pmatrix} 0 & 1 & 0 \\ \frac{9\rho_2}{(7+18\rho_2)} & 0 & \frac{(9\rho_2+7)}{(7+18\rho_2)} \\ \frac{9\rho_3}{(13+27\rho_3)} & \frac{(9\rho_3+6)}{(13+27\rho_3)} & \frac{(9\rho_3+7)}{(13+27\rho_3)} \end{pmatrix} \quad (60)$$

where ρ_i are arbitrary in $[0, 1)$, $i = 2, 3$. The corresponding optimal wagers for the Tourist are

$$w = (1, 7/25 + 18\rho_2/25, 13/40 + 27\rho_3/40)^\top \quad (61)$$

The Tourist's minimum risk optimal strategy is obtained by setting all the ρ_i 's equal to zero, yielding

$$Q = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 6/13 & 7/13 \end{pmatrix} \quad (62)$$

and

$$w = (1, 7/25, 13/40)^\top \quad (63)$$

Summarizing the discussion in this section, we have

Theorem 4.2 (Limiting Strategies in the m step Tourist Game). The optimal strategies for the first step of the m step Tourist Game on a strongly connected graph G tend towards limits as m goes to infinity. The limiting α -values for the m step Tourist Game are given by

$$\alpha^{-1} = \frac{xy^T}{x^T y} \mathbf{1}, \quad (64)$$

where x and y are any right and left eigenvectors respectively corresponding to the maximal eigenvalue $r \in [1/2, 1]$ of the propagation matrix M for G . Moreover, α^{-1} is itself a positive right eigenvector of M corresponding to r . The Guide's limiting strategy is given by the transition matrix $P = (p_{i,j}, i, j \in V)$, where

$$p_{i,j} = \begin{cases} \alpha_j^{-1} / \sum_{k:i \rightarrow k} \alpha_k^{-1} & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (65)$$

The Tourist's limiting strategies are characterized by

$$q_{i,j} w_{i,j} = \begin{cases} p_{i,j} - (1 - \rho_i) p_{i,\min} & \text{if } d(i) > 1 \text{ and } i \rightarrow j \\ 1 & \text{if } d(i) = 1 \text{ and } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (66)$$

where $p_{i,\min} = \min\{p_{i,k} : k \in V, \text{ and } i \rightarrow k \text{ in } G\}$, and the ρ_i are arbitrary in $[0, 1)$. Under optimal play, the Tourist's expected fortune F_m satisfies

$$\lim_{m \rightarrow \infty} r^m E[F_m | X_0 = i] = \alpha_i \in (0, \infty), \quad (67)$$

where the α_i are the reciprocals of the components of the vector α^{-1} of (64).

5 Limiting Strategies are Optimal for the Infinite Duration Tourist Game

Theorem 5.1 (Optimal Strategies for the Infinite Duration Tourist Game). The limiting strategies of Theorem 4.2 are optimal for the infinite duration Tourist Game.

Proof. Suppose the players are initially at node i . Consider the following class of strategies: on the first move of the game, the players can choose any strategy they desire, but on all subsequent moves, the players must play according to their limiting strategies. It suffices to prove that no such strategy can beat consistently playing the limiting strategies.

First, look at things from the Guide's point of view.

$$E[F_1 | X_0 = i, X_1 = j, G_i = k] = \begin{cases} 1 + (n-1) w_{i,k} & \text{if } j = k \\ 1 - w_{i,k} & \text{otherwise} \end{cases} \quad (68)$$

So

$$E[D_m | X_0 = i, X_1 = j, G_1 = k] = E[D_1 | X_0 = i, X_1 = j, G_1 = k] r^{m-1} \alpha_{j, m-1} \quad (69)$$

and

$$\begin{aligned} E[D_m | X_0 = i, G_1 = k] &= \sum_j E[D_m | X_0 = i, X_1 = j, G_1 = k] p_{i,j} \quad (70) \\ &= \sum_j E[D_1 | X_0 = i, X_1 = j, G_1 = k] r^{m-1} \alpha_{j, m-1} p_{i,j} \quad (71) \end{aligned}$$

$$\begin{aligned} &= r(1 + (n-1)w_{i,k}) r^{m-1} \alpha_{k, m-1} p_{i,k} \\ &\quad + r(1 - w_{i,k}) \sum_{j:j \neq k} r^{m-1} \alpha_{j, m-1} p_{i,j} \quad (72) \end{aligned}$$

So

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i, G_1 = k] = r(1 + (n-1)w_{i,k}) \alpha_k p_{i,k} \quad (73)$$

$$\begin{aligned} &\quad + r(1 - w_{i,k}) \sum_{j:j \neq k} \alpha_j p_{i,j} \\ &= r(\alpha \cdot p_{(i)} + w_{i,k}(n\alpha_k p_{i,k} - \alpha \cdot p_{(i)})) \quad (74) \end{aligned}$$

From the results of Section 2, the expression in the outer parentheses here is equal to H_i for all k if and only if the Guide is playing his limiting strategy at node i on the first move; furthermore, if the Guide plays a different strategy at node i , then the Tourist can choose her wagers such that for at least one k

$$\alpha \cdot p_{(i)} + w_{i,k}(n\alpha_k p_{i,k} - \alpha \cdot p_{(i)}) > H_i. \quad (75)$$

Therefore, if the Guide is playing his limiting strategy at node i , we have

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i, G_1 = k] = rH_i = \alpha_i, \quad (76)$$

for all k , so that

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i] = \alpha_i, \quad (77)$$

no matter what strategy the Tourist employs on the first move, while if the Guide is not playing his limiting strategy at node i , then on the first move the Tourist can concentrate her strategy on those k such that (75) holds, and thereby force

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i] < \alpha_i. \quad (78)$$

This proves that the limiting strategy for the Guide is in fact optimal for the infinite duration game.

Now look at it from the Tourist's point of view.

$$E[F_1 | X_0 = i, X_1 = j] = (1 + (n-1)w_{i,j})q_{i,j} + \sum_{k:k \neq j} (1 - w_{i,k})q_{i,k} \quad (79)$$

$$= 1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)} \quad (80)$$

It follows that

$$E[D_1 | X_0 = i, X_1 = j] = r(1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)}) \quad (81)$$

and

$$E[D_m | X_0 = i, X_1 = j] = r(1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)})r^{m-1}\alpha_{j,m-1} \quad (82)$$

and so

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i, X_1 = j] = r(1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)})\alpha_j. \quad (83)$$

Now by the results of Section 2, $(1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)})\alpha_j = H_i$ for all j such that $i \rightarrow j$ if and only if the Tourist is playing one of her limiting strategies at node i on the first move; furthermore, if the Tourist plays a different strategy at node i , then for at least one j we have

$$(1 + nw_{i,j}q_{i,j} - w_{(i)} \cdot q_{(i)})\alpha_j < H_i. \quad (84)$$

Therefore, if the Tourist is playing one of her limiting strategies at node i , we have

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i] = rH_i = \alpha_i, \quad (85)$$

no matter what strategy the Guide employs, while if the Tourist is not playing a limiting strategy at node i , then on the first move the Guide can concentrate his strategy on those j such that (84) holds, and thereby force

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = i] < \alpha_i. \quad (86)$$

Therefore, the limiting strategies for the Tourist are in fact optimal for the infinite duration game. This completes the proof. \square

6 The Markov Chain Associated with Optimal Play

The optimal strategies of Theorem 5.1 give rise to a Markov chain on the graph in which the game is being played. As a prelude to studying this chain, let's summarize the results up to now, and introduce a little notation that will be helpful in the sequel. Recall that $d(i)$ is the out-degree of vertex i in G . Denote

$$e_i = \begin{cases} d(i) & \text{if } d(i) \geq 2 \\ 2 & \text{if } d(i) = 1 \end{cases} \quad (87)$$

(Since G is strongly connected, $d(i) \geq 1$ for all $i \in V$.) Then the propagation matrix $M = (m_{i,j} : i, j \in V)$ is given by

$$m_{i,j} = \begin{cases} 1/e_i & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (88)$$

By Theorem 4.2 the reciprocal value vector α^{-1} is an eigenvector of M corresponding to the maximal eigenvalue $r = r$:

$$\begin{aligned} r\alpha_i^{-1} &= \sum_j m_{i,j}\alpha_j^{-1} \\ &= \frac{1}{e_i} \sum_{j:i \rightarrow j} \alpha_j^{-1} \end{aligned} \quad (89)$$

The optimal transition probabilities for the Guide are given by

$$p_{i,j} = \begin{cases} \alpha_j^{-1} / \sum_{k:i \rightarrow k} \alpha_k^{-1} & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (90)$$

Note that by equation (89),

$$p_{i,j} = \begin{cases} \alpha_i / (e_i r \alpha_j) & \text{if } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (91)$$

Let X_m denote the vertex of G which contains the players at step m of the game. Since it is the Guide who determines the path that the players take through the graph, the transition probabilities (91) govern the Markov chain $(X_m : m \geq 0)$. One thing we would like to know is the invariant measure for this chain, which would tell us the long-term fraction of time that the players spend at each vertex of G . We will soon take up this question.

The Tourist's strategy at node i consists of a probability vector $q_{(i)} = (q_{i,j} : j \in V)$, and a wager vector $w_{(i)} = (w_{i,j} : j \in V)$. Here, $q_{i,j}$ represents the probability that the Tourist bets that the Guide chooses node j as the next destination node, given that the players are currently at node i , and $w_{i,j}$ is the wager she places on this outcome. For nodes with out-degree at least 2, equation (16) characterizes optimal play on the part of the Tourist, which in the context of the game on the digraph G reads:

$$1 + e_i w_{i,j} q_{i,j} - w_{(i)} \cdot q_{(i)} = \alpha_j^{-1} H_i, \quad (92)$$

for all j such that $i \rightarrow j$, where H_i is the harmonic mean of the values α_k such that $i \rightarrow k$. The solution of this linear system is given by Theorem 2.1 as

$$q_{i,j} w_{i,j} = p_{i,j} - (1 - \rho_i) p_{i,\min}, \quad (93)$$

where $p_{i,\min} = \min\{p_{i,k} : k \in V \text{ and } i \rightarrow k \text{ in } G\}$, and where ρ_i is an arbitrary constant in $[0, 1)$ (which can depend on i). For nodes of out-degree 1, the

optimal strategy for the Tourist is obviously $w_{i,j} = q_{i,j} = 1$ for the unique node j such that $i \rightarrow j$. So the optimal strategies for the Tourist are characterized by

$$q_{i,j}w_{i,j} = \begin{cases} p_{i,j} - (1 - \rho_i)p_{i,\min} & \text{if } d(i) > 1 \text{ and } i \rightarrow j \\ 1 & \text{if } d(i) = 1 \text{ and } i \rightarrow j \\ 0 & \text{otherwise} \end{cases} \quad (94)$$

Hence we can represent the optimal strategies for the Tourist with two matrices, $Q = (q_{i,j} : i, j \in V)$ and $W = (w_{i,j} : i, j \in V)$.

The whole point of the Tourist's play is to maximize her expected fortune. The strategy (94) accomplishes this. But the Tourist might also be interested in what her average fortune is when she is located at a specific node of G . This is another issue we will take up soon. But first, we discuss the invariant measure for the chain.

6.1 The Invariant Measure for the Chain

A probability measure μ on the vertices of G is said to be *invariant* for the Markov chain $(X_m : m \geq 0)$ if X_1 has the distribution μ whenever X_0 does. Thus, if μ is invariant for the chain, then *all* the random variables X_m have the distribution μ , and therefore the distribution of the players' location in the graph does not change with time. In terms of the transition matrix P , the measure μ is invariant if and only if $\mu^T P = \mu^T$; i.e.

$$\mu_j = \sum_{i \in V} \mu_i p_{i,j}$$

By equation (91), μ is invariant for our chain $(X_m : m \geq 0)$ if and only if

$$\mu_j = \sum_{i:i \rightarrow j} \mu_i \frac{\alpha_i}{r e_i \alpha_j},$$

or equivalently

$$r \alpha_j \mu_j = \sum_{i:i \rightarrow j} \alpha_i \mu_i \frac{1}{e_i} \quad (95)$$

$$= \sum_{i \in V} \alpha_i \mu_i m_{i,j} \quad (96)$$

This says that the vector $(\alpha_j \mu_j : j \in V)$ is a left eigenvector of the propagation matrix M corresponding to its maximal eigenvalue r . Thus to *find* the invariant measure, we need only find any left eigenvector y of M associated with r (recall that the eigenspace has dimension 1), and set

$$\mu_i = c y_i \alpha_i^{-1}, \quad (97)$$

where c is chosen to make the sum of the μ_i 's equal to 1. But recall that the vector of reciprocal values $\alpha^{-1} = (\alpha_i^{-1} : i \in V)$ is itself a right eigenvector of

M corresponding to r . So the invariant measure is simply the componentwise product of the left and right eigenvectors of the propagation matrix, suitably scaled. Interesting.

Recall that by equation (64) of Theorem 4.2 the vector of reciprocal α -values satisfies $\alpha^{-1} = (x^T y)^{-1} x y^T \mathbf{1}$, where x and y are any right and left eigenvectors of the propagation matrix M , and $\mathbf{1}$ is the vector of all ones. This means that

$$\alpha_i^{-1} = (x^T y)^{-1} x_i \sum_j y_j \quad (98)$$

Combining this with (97) gives

$$\mu_i = c (x^T y)^{-1} x_i y_i \sum_j y_j. \quad (99)$$

Summing over i gives $c = \left(\sum_j y_j\right)^{-1}$, from which we get

$$\mu_i \alpha_i = \frac{y_i}{\sum_j y_j}, \quad (100)$$

and hence

$$\sum_i \mu_i \alpha_i = 1. \quad (101)$$

Thus, when the chain is in equilibrium, the average value experienced by the players in the discounted game is 1.

6.2 Steady State Fortunes

Let D_m denote the Tourist's *discounted* fortune at step m of the game. When the chain is in *equilibrium* - i.e. when the distribution of X_m is governed solely by the invariant measure μ , and no longer varies with time - then the discounted fortune will also have a stationary distribution. In particular, the conditional expectations

$$\eta_j = E[D_m | X_m = j] \quad (102)$$

will not depend on m . The η_j 's represent how much money the Tourist can expect to have in her pocket when she is currently at node j . She would like to know what these numbers are.

To solve equations (102), condition on where the players were on the previous time step. Given that $X_m = j$ and $X_{m-1} = i$, we have

$$\begin{aligned} D_m &= r D_{m-1} \left[(1 + (e_i - 1) w_{i,j}) \xi(\text{Tourist bets on destination } j) \right. \\ &+ \left. \sum \{ (1 - w_{i,k}) \xi(\text{Tourist bets on destination } k) ; k : i \rightarrow k \text{ and } k \neq j \} \right]. \end{aligned} \quad (103)$$

(This holds for nodes i of out-degree 1 as well.) Taking expectations now

$$E[D_m | X_m = j, X_{m-1} = i, D_{m-1}] = rD_{m-1}[(1 + (e_i - 1)w_{i,j})q_{i,j} + \sum\{(1 - w_{i,k})q_{i,k}\}] \quad (104)$$

which reduces to

$$E[D_m | X_m = j, X_{m-1} = i, D_{m-1}] = rD_{m-1}[1 + e_i w_{i,j} q_{i,j} - w_{(i)} \cdot q_{(i)}] \quad (105)$$

Integrating over D_{m-1} now yields

$$E[D_m | X_m = j, X_{m-1} = i] = rE[D_{m-1} | X_{m-1} = i][1 + e_i w_{i,j} q_{i,j} - w_{(i)} \cdot q_{(i)}] \quad (106)$$

If the Tourist is playing optimally, then

$$1 + e_i w_{i,j} q_{i,j} - w_{(i)} \cdot q_{(i)} = \alpha_j^{-1} H_i \quad (107)$$

(See equations (16) and (92).) But $H_i = \frac{1}{r} \alpha_i$ from equation (89), so

$$E[D_m | X_m = j, X_{m-1} = i] = rE[D_{m-1} | X_{m-1} = i] \frac{\alpha_i}{r\alpha_j} \quad (108)$$

From this it follows that

$$E[D_m | X_m = j] = \sum_{i:i \rightarrow j} E[D_{m-1} | X_{m-1} = i] \frac{\alpha_i^2 \mu_i}{r\alpha_j^2 \mu_j e_i} \quad (109)$$

If the chain is in equilibrium, then the two conditional expectations are η_j and η_i respectively, so we get

$$r\alpha_j^2 \mu_j \eta_j = \sum_{i:i \rightarrow j} \alpha_i^2 \mu_i \eta_i \frac{1}{e_i}. \quad (110)$$

Therefore the vector $(\alpha_j^2 \mu_j \eta_j : j \in V)$ is a left eigenvector of the propagation matrix M corresponding to the maximal eigenvalue r . But we have already seen that $(\alpha_j \mu_j : j \in V)$ is also a left eigenvector of M corresponding to r , and that the eigenspace has dimension 1. Therefore $\alpha_j \eta_j = C$, where C is a constant.

To determine the value of C , start the discounted game in equilibrium: $\text{Prob}(X_0 = j) = \mu_j$. Then $E[D_m | X_m = j]$ does not depend on m , and so neither does

$$E[D_m] = \sum_j E[D_m | X_m = j] \text{Prob}(X_m = j) = \sum_j \eta_j \mu_j. \quad (111)$$

But we also know from Theorem 4.2 that $\lim_{m \rightarrow \infty} E[D_m | X_0 = j] = \alpha_j$. So in equilibrium, $E[D_m | X_0 = j] = \alpha_j$ for all m . Therefore

$$E[D_m] = \sum_j E[D_m | X_0 = j] \text{Prob}(X_0 = j) = \sum_j \alpha_j \mu_j = 1 \quad (112)$$

by (101). Therefore $\sum_j \eta_j \mu_j = 1$, which fixes the constant C above. The interpretation is this: the overall average discounted fortune of the Tourist as the players move through the graph is one.

Summarizing the discussions in this section, we have

Theorem 6.1 (Markov Chain Properties). The invariant measure for the Tourist Game on the graph G is given by

$$\mu_i = \frac{y_i \alpha_i^{-1}}{\sum_j y_j}, \quad (113)$$

where y is any left eigenvector of the propagation matrix M associated with its maximal eigenvalue r . The steady state fortunes $\eta_i = E[D_m | X_m = i]$ satisfy

$$\eta_i = C \alpha_i^{-1}, \quad (114)$$

where the constant C is chosen such that $\sum_i \eta_i \mu_i = 1$.

Example 6.1 (Markov Chain Properties). Continuing with the analysis of the Tourist Game on the graph G of Figure 3, we find that the invariant measure is

$$\mu = \frac{1}{40} (9, 15, 16)^T, \quad (115)$$

and the steady state fortunes are

$$\eta = \frac{1}{40} (9, 15, 16)^T. \quad (116)$$

It is something of an accident that in this case the invariant measure and steady state fortunes have the same value. It is due to the fact that the left eigenvector y for this example is the vector of all ones.

7 Games With Terminal States

Under Construction.

8 Applications to the Lying Oracle Game

;; insert explanation of Lying Oracle game (fair case) here ;;

8.1 At Most m Lies in Any Block of Length n

The (m, n) -game is defined by the constraint that the Oracle can lie at most m times in any block of n predictions. The interesting cases are when $1 \leq m \leq n - 1$. All of these games are equivalent to playing the Tourist game on a certain directed graph.

Example 8.1 (The $(2, 3)$ -game). In this game, the Oracle can lie at most two times in any block of three predictions. To find the equivalent Tourist Game, begin by writing down all the allowed lying patterns of length $n = 3$ (these will become the vertices of the Tourist Game graph G), and then map out the allowed transitions between these nodes.

state	left child (t)	right child(ℓ)
ttt	ttt	$t\ell\ell$
$t\ell\ell$	$t\ell t$	$t\ell\ell$
$t\ell t$	ℓtt	$\ell t\ell$
$t\ell\ell$	$\ell\ell t$	
ℓtt	ttt	$t\ell\ell$
$\ell t\ell$	$t\ell t$	$t\ell\ell$
$\ell\ell t$	$\ell\ell t$	$\ell t\ell$

This transition table corresponds to the strongly connected digraph in Figure 8.1. Notice that transitions to states in the left hand column of the figure correspond to the Oracle telling the truth, while transitions to the states in the right hand column correspond to the Oracle telling a lie. Each allowed lying pattern in the $(2, 4)$ -game corresponds to at least one path in this graph, and each path in this graph corresponds to exactly one allowed lying pattern. In the context of the Lying Oracle game, play always begins with the players located at the “ ttt ” state.

However, there are simpler graphs that also represent the $(2, 3)$ -game. The basic insight needed to find them is this: if two states have exactly the same children, then the paths starting from these states generate exactly the same lying patterns. Hence, the states can be “contracted”, by merging the two nodes, and eliminating any multiple edges in the resulting graph. The new graph will have one less vertex, but it will generate exactly the same lying patterns as the old graph. By performing this operation repeatedly, we arrive at a graph with drastically fewer nodes than the one we started with. In the case of the $(2, 3)$ -game, we can reduce the graph of Figure 8.1 to the graph of Figure 8.1.

Again, transitions to the state in the left column represent the Oracle telling the truth, while transitions to the states in the right column represent the Oracle telling a lie. Play starts from state 0. Applying the Tourist Game theory to this graph yields discount rate of

$$r = 0.9196 \tag{117}$$

and an α -vector of

$$\alpha = (0.8792, 1.0475, 1.6170) . \tag{118}$$

The optimal strategies and associated invariant measure for the chain are shown in Figure 8.1. The graph on the left shows the Oracle’s transition probabilities




Figure 6: A graph for the (2,3)-game




Figure 7: The simplest graph for the (2,3)-game

(the edge labels) and the invariant measure (the vertex labels). The graph on the right shows the Player's minimum risk optimal strategy, with the guessing probabilities (edge labels) and the wagers (vertex labels).

The invariant measure for the chain tells us something important about the long-term fraction of time in which the Oracle tells the truth. Recall that transitions to node 0 represent the Oracle telling the truth, while transitions to nodes 1 or 2 represent the Oracle telling a lie. The invariant measure gives the fraction of time that the players are located at each of these nodes. Hence, we can conclude that the long-term fraction of time that the Oracle is telling the truth is

$$\mu_0 = 0.6184, \tag{119}$$

while the long-term fraction of time that the Oracle is lying is

$$\mu_1 + \mu_2 = 0.3816. \tag{120}$$

Thus, while the Oracle *could* lie up to 67% of the time in the (2,3)-game, in order to play optimally he should lie just over 38% of the time.




Figure 8: Optimal strategies for the (2,3)-game. The graph on the left shows the Oracle's transition probabilities (edge labels) and the invariant measure for the chain (vertex labels). The graph on the right shows the Player's minimum risk guessing probabilities (edge labels) and wagers (vertex labels).

The α -values also say something important about the game. Play always begins at node 0, which has an α -value of 0.8792. Recall that by equation (67) of Theorem 4.2, this α -value represents the limiting expected fortune of the Player in the discounted game:

$$\lim_{m \rightarrow \infty} E[D_m | X_0 = 0] = \alpha_0 = 0.8792 \quad (121)$$

Thus, if the Player plays the discounted game starting with \$1, then as play continues she will have on average about 88¢ in her pocket. Of course, $D_m = r^m F_m$, so the Player's fortune in the undiscounted game satisfies

$$\lim_{m \rightarrow \infty} \frac{E[F_m | X_0 = 0]}{r^{-m} \alpha_0} = 1 \quad (122)$$

In other words, on average the undiscounted fortune grows as $0.8792 \cdot 1.0874^m$.

Example 8.2 (The (2,4)-game). In the (2,4)-game, the Oracle can lie at most two times in any block of four predictions. Using the same technique as in Example 8.1 we arrive at the following equivalent Tourist Game graph:

As before, transitions to states in the left column of this figure represent times when the Oracle tells the truth, while transitions to states in the right column represent times when the Oracle tells a lie. Play always starts from state 0.

Applying the Tourist Game theory to this graph results in a discount rate of

$$r = 0.8571 \quad (123)$$

and a value vector of

$$\alpha = (0.7689, 0.8695, 1.3180, 1.0764, 1.4906, 2.2595). \quad (124)$$




Figure 9: The graph for the $(2, 4)$ -game.

The resulting optimal strategies and invariant measure are shown in the next figure. Overall, the Oracle is lying 30.43% of the time, while the Player's expected fortune in the undiscounted game grows as $0.7689 \cdot 1.1667^m$.




Figure 10: Optimal strategies for the $(2, 4)$ -game. The graph on the left shows the Oracle's transition probabilities (edge labels) and the invariant measure for the chain (vertex labels). The graph on the right shows the Player's minimum risk guessing probabilities (edge labels) and wagers (vertex labels).

Example 8.3 (The exponential growth rate as a function of m and n). It is interesting to consider the exponential growth rate of the Player's fortune $\beta = 1/r$ as a function of m and n . Figure 8.3 plots the exponential rate of growth of the Player's average fortune, $\beta(km, kn)$ as a function of k , for several different values of m/n . This plot was generated from data produced by a collection of Java classes that I have written.

Two open questions:

1. For fixed m and n , what is $\lim_{k \rightarrow \infty} \beta(km, kn)$?
2. In the same spirit, let $\ell(m, n)$ be the overall probability that the Oracle lies in the (m, n) -game. What is $\lim_{k \rightarrow \infty} \ell(km, kn)$?

Example 8.4 (The $(1, n)$ -game). Under construction.

Example 8.5 (The $(n - 1, n)$ -game). Under construction.

9 The Tourist Game Without Odds-Weighting

Under construction.

References

- [1] Henryk Minc, *Nonnegative Matrices*, John Wiley & Sons, New York, 1988

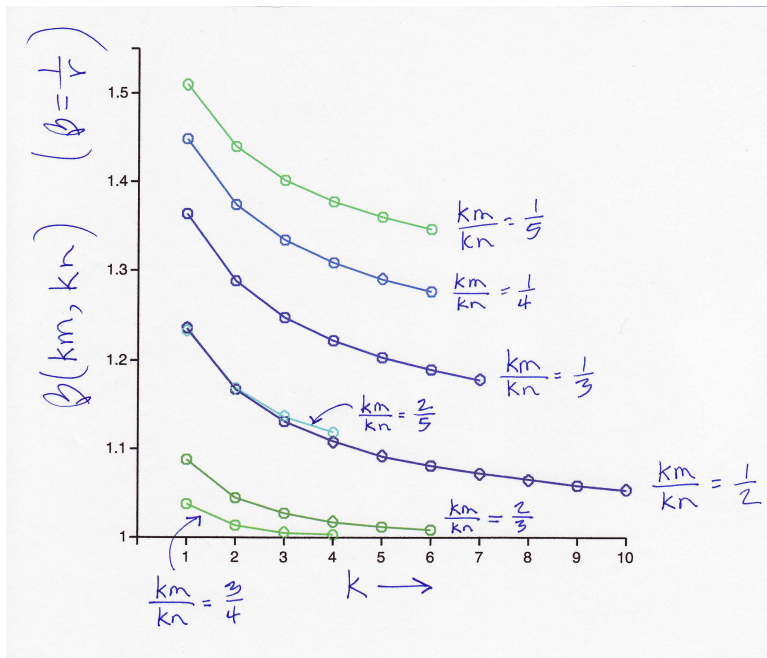


Figure 11: $\beta(km, kn)$ is the exponential growth rate of the Player's expected fortune in the (km, kn) -game. This plot shows $\beta(km, kn)$ as a function of k for several values of m and n . Plot generated in Maple from data produced by a collection of Java classes written by the author.

- [2] Peter Lancaster and Miron Tismenetsky, *The Theory of Matrices*, second edition, Academic Press, New York, 1985
- [3] Robb Koether and John Osoinach, "Outwitting the Lying Oracle", *Mathematics Magazine*, 78 (2005), 98-109.