Math. Scientist **18**, 73–84 (1993)
Printed in England

LET'S MAKE A DEAL

EDUARDO ENGEL,* *Harvard University and Universidad de Chile*
ACHILLES VENETOULIAS,** *Bank of America*

Abstract

The game played on the American TV show 'Let's Make a Deal' gives rise to a popular problem in elementary probability which is equivalent to the three prisoners' dilemma. This paper generalizes the problem into a sequential game with two contestants, and analyzes the decision-making process for the two players. Optimal strategies are derived for the two players, under various assumptions for the behavior of the players and the structure of the game.

THREE CURTAINS PUZZLE; THREE PRISONERS' DILEMMA; DECISION THEORY; ZERO-SUM GAME; OPTIMAL STRATEGY

AMS 1991 SUBJECT CLASSIFICATION: PRIMARY 90D15
SECONDARY 90B40, 90D26, 93E05

1. Introduction

The game played on the TV show 'Let's Make a Deal' (with Monty Hall) gives rise to the following interesting variation. There are three curtains, a (male) announcer and a (female) player. One of the curtains conceals a valuable prize and the other two are empty; only the announcer knows where the prize is. The player wins the prize if she guesses correctly which curtain conceals the prize. Initially the player selects a curtain. Then the announcer opens an empty curtain that the player did not choose. Having eliminated that curtain from contention, the announcer offers the player the option to select the remaining third curtain as her final choice. At this point the player must decide whether it is in her advantage to switch curtains or to insist on her original guess [5]–[7].

The decision the player faces is a tricky problem in elementary probability which has become a popular puzzle known as 'the three curtains puzzle'. At first it appears as if the player should be indifferent to which curtain she

---

Received 17 March 1992; revision received 27 October 1992.

*Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA and Departamento de Ingenieriá Industrial, Casilla 2777, Santiago, Chile.

**Capital Markets Division, Bank of America, 555 California Street, San Francisco, CA 94104, USA.

Research carried out while this author was teaching in the Sloan School of Management at the Massachusetts Institute of Technology and partially supported by National Science Foundation grant IST-8420614, by U. S. Army grant DAAG-29-84-K-0207 and by the Center for Computational Research in Economics and Management Science at MIT.

selects because both of the unopened curtains have the same probability of concealing the prize. However, more careful examination of the conditional probabilities reveals that this is not the case. The player's first guess conceals the prize with probability $\frac{1}{3}$, whereas the other curtain conceals the prize with probability $\frac{2}{3}$. Thus the player is not only better off if she switches to the third curtain but, by doing so, she doubles her chances of winning the prize.

This surprising conclusion is what has made the problem a favorite mathematical puzzle. The incorrect position of indifference appeals to intuition in an apparently compelling manner. As a result, most people are not eager to accept that the player should actually switch curtains because they find it counterintuitive. For this reason, some of the most prestigious business schools in the United States use this problem to educate their MBA students in decision-making: it is an example of how intuition can be misleading in making decisions under uncertainty [4]. Recently, the puzzle has attracted public attention because it appeared in nationally syndicated newspaper columns across the United States and generated long arguments as to whether the indifference position is optimal or not [1].

There are many arguments for verifying that switching curtains offers an advantage. The shortest, though not the most intuitive, argument is the following. Initially all three curtains are equally likely to conceal the prize and the player has a $\frac{1}{3}$ chance of winning the prize on her first guess. Suppose that the player decides to stay with her first guess, after the announcer opens another curtain. Clearly, nothing changes about the curtain the player has already selected. In particular, the probability of getting the prize with that curtain is the same as before, i.e. $\frac{1}{3}$. But now there are only two candidate curtains and one of them conceals the prize with probability $\frac{1}{3}$. Therefore, the other curtain conceals the prize with probability $1-\frac{1}{3} = \frac{2}{3}$. Deciding never to switch curtains is equivalent to not having the option to switch; hence the odds of winning by not switching are one in three.

The critical element behind the solution of the puzzle is that the announcer knows *both* which curtain conceals the prize *and* which curtain the player has chosen *before* he opens a curtain. He is deprived of the opportunity to eliminate the player's guess even if that curtain is empty. Therefore, the announcer's choice does not contribute anything to the likelihood that the player's curtain conceals the prize. On the contrary, the fact that, *having had the opportunity to do so*, the announcer did not open the curtain which the player did not choose, does increase the probability of finding the prize behind that curtain. This observation provides a heuristic explanation for why the curtain which the player did not choose is more likely to conceal the prize.

The object of this paper is to analyze the decision-making process for the player in a more general version of the game, in which there are n available curtains and the announcer successively eliminates k curtains. Each time the announcer eliminates a curtain he offers the player an additional guess. The player wins if she guesses correctly which curtain conceals the prize on her last guess. In this case the player faces a series of k decisions and can choose from an overwhelming number of possible strategies. This paper determines whether, and under what conditions, there exists an optimal strategy for the player. Various assumptions for the behavior of the announcer and the structure of the game are studied. The main results of the paper are presented in an informal manner in the remainder of this section. Section 2 gives a precise definition of the game, and Sections 3 and 4 contain the mathematical results.

In the special case of a memoryless player (i.e. a player who does not remember the course of the game) not only does there exist an optimal strategy but a complete ordering of all possible strategies can be obtained. This topic is beyond the scope of the present paper and is presented elsewhere [2].

The intuition for many of the results in this paper is evident in the second explanation given above for the original puzzle. When the announcer opens a curtain, he is prevented from discarding the player's current curtain and is, therefore, not giving out any information on the likelihood of finding the prize behind that curtain. At the same time, the announcer's choice increases the likelihood of finding the prize behind one of the curtains he could, but did not, choose. Thus it seems plausible, and turns out to be the case, that the last time the player has the option to switch curtains, the curtains that are more likely to conceal the prize are those that she has never visited before. In other words, any strategy that switches to an unvisited curtain on the last round of the game is optimal. Furthermore, the probability of finding the prize behind one of the unvisited curtains does not depend on the strategy that the player followed during the earlier part of the game. The previous rounds only prevent the announcer from discarding certain curtains, and this is taken into account by all the *move-to-a-new-curtain-on-the-last-round* strategies. Choosing an unvisited curtain on the last round gives the contestant the largest probability of winning the prize regardless of what she did on all the previous rounds (Theorem 3.1).

Section 4 extends the original game in two directions. First, the number of rounds becomes random and unknown but independent from the actions of the player and the announcer. Among all the optimal strategies for the game with a known number of rounds, only one remains optimal. This is the strategy where the player moves to an unvisited curtain on every move (Theorem 4.1). Second, the number of rounds remains known but the game becomes a zero-sum game where the announcer attempts to minimize the player's probability of winning. Again, among all the *move-to-a-new-curtain-on-the-last-round* strategies only one remains optimal. This is the strategy where the player stays with her initial guess until the last move and then chooses a new curtain (Theorem 4.2); obviously this is the only strategy which is not affected by the announcer's behavior. Thus either a random number of rounds or a malevolent announcer reduce the optimal strategies from the infinitely many *move-to-a-new-curtain-on-the-last-round* to a single strategy, one for each case. Unfortunately, these two strategies do not coincide, and this approach does not yield a unique optimal strategy for the zero-sum game with a random, unknown number of rounds. For this most general game, finding the optimal strategies for both the player and the announcer remains an open problem. A conjecture for the form of these strategies is presented in Section 4.3.

2. Definition of the game

The definition of the general version of the game is as follows. There is one announcer, one player and n nodes (curtains). The nodes are labeled $1,2,\ldots,n$. One of these nodes, the *winning* node, contains a prize; only the announcer knows where the prize is. There is no loss of generality in assuming that all nodes are equally likely to contain the prize. The player wins the prize by guessing correctly, at the end of the game, which the winning node is.

The game begins with the player making an initial guess (i.e. selecting a node). Afterwards, the announcer and the player take turns in making a move, with the announcer moving first. A move is the selection of, or the *visit* to, a node. The announcer makes a move by choosing a node which does not contain the prize and eliminating it from further consideration, by revealing its contents; however, the announcer cannot choose the player's most recent guess. The player makes a move by choosing a node—not already eliminated by the announcer—which represents her current guess for the winning node. One move by the announcer and one move by the player comprise one round (stage, complete move).

The game begins and ends with the player making a move. Therefore, the above definition entails the convention that the player's initial move provides a starting point and does not constitute part of any of the rounds of the game. The number of nodes, n, is fixed. The number of rounds of the game is determined by a random variable K, whose distribution is known to both the player and the announcer but whose actual value is unknown to them. The game where K is constant and equal to k is denoted by $G(k)$, with the understanding that $k = 0$ corresponds to the game where the player simply chooses a node and the game ends. This emphasizes the fact that in $G(k)$ the player actually makes $k+1$ moves. It should also be noted that the game $G(k)$ does not make sense unless there are at least two nodes left (not eliminated) after the announcer has discarded k nodes. Therefore, the integers n and k are subject to the constraint $n \geqslant k+2$ or, more generally, the support of K is contained in $\{1,2,\ldots,n-2\}$.

At any round of the game, the nodes available to the player are all the nodes except for the ones which have been eliminated by the announcer in the previous rounds. The nodes available to the announcer are all the nodes except for the ones already eliminated, the one containing the prize and the one which represents the player's most recent move.

The probability of winning is a function of the strategies used by the player and the announcer. The player's strategy is defined as the collection of conditional distributions of her kth move ($k = 1,2,\ldots$), given all feasible histories of the game at that point. The announcer's strategy is defined analogously. The *history*, h_k, for any round k captures the information contained in the course of the previous $k-1$ rounds. Let v_k and d_k denote the nodes which are respectively visited and discarded on the kth round (v_0 denoting the player's initial selection). The history for the kth round is

$$h_k = (v_0, d_1, v_1, d_2, v_2, \ldots, d_{k-1}, v_{k-1}, d_k) \quad (k \geqslant 1).$$

The d's are all distinct and $v_k \notin \{d_1, d_2, \ldots, d_k\}$, for all k. The $2k$-tuplet h_k represents the knowledge on which the player bases her action in the kth round. The $(2k+1)$-tuplet (h_{k-1}, v_{k-1}) represents the knowledge on which the announcer bases his selection for the kth discard. For notational purposes it is convenient to extend the definition of h_k to the case $k = 0$ by defining $h_0 = (v_0)$. The tuplets h_k $(k = 0, 1, \ldots)$ form the space $\mathscr{H}$ of all possible histories for the game.

The announcer is either fair or malevolent (the announcer can also be helpful [2], but that case is not discussed here). When the announcer is fair, his conditional distribution on any round is uniform over all the nodes he is allowed to discard at that point. In this case, the problem is to determine the strategy that maximizes the player's probability of winning for a given distribution of $\boldsymbol{K}$. When the announcer is malevolent, his conditional distribution on any round is the one that minimizes the player's overall probability of winning the game. In this case the game is a zero-sum game, and the problem is to determine the value of the game as well as the strategies (for the announcer and the player) which attain that value.

For the remainder of this paper, let $\boldsymbol{R}$ denote the random variable that determines the winning node, and assume (without loss of generality) that $\boldsymbol{R}$ has a uniform distribution on $\{1, 2, \ldots, n\}$; let $\boldsymbol{V}_k$ denote the random variable that determines the node which the player visits on the kth round; let $\boldsymbol{D}_k$ denote the random variable that determines which node the announcer discards on the kth round and let $\mathscr{D}_k = \{d_1, \ldots, d_k\}$ denote the set of discarded nodes on the first k rounds. Let also

$$\boldsymbol{H}_k = (\boldsymbol{V}_0, \boldsymbol{D}_1, \boldsymbol{V}_1, \boldsymbol{D}_2, \boldsymbol{V}_2, \ldots, \boldsymbol{D}_{k-1}, \boldsymbol{V}_{k-1}, \boldsymbol{D}_k)$$

denote the random variable which describes the history of a game up to the kth round, h_k denote a history that has positive probability of occurrence under the strategies being considered and the distribution for the number of rounds, h_j (for $0 \leqslant j \leqslant k-1$) denote the sub-history of the first j rounds that is determined by h_k, and $\boldsymbol{H}_j$ denote the random variable that describes h_j.

If Φ denotes a strategy for the player and Ψ denotes a strategy for the announcer (given a certain distribution for the number of rounds, $\boldsymbol{K}$), the probability of winning is denoted by $P_{\Phi,\Psi}\{\text{win}\}$. When the announcer is fair and the number of rounds is fixed in advance, the probability of winning is simply denoted by $P_{\Phi}\{\text{win}\}$. For the remainder of the article, it is understood that probabilities conditional on events of zero probability are equal to 0.

3. The main optimality result

This section considers the case where the number of rounds is fixed in advance (and equal to k) and the announcer is fair. In this game the player's problem is to choose the k-strategy Φ which maximizes $P_{\Phi}\{\boldsymbol{V}_k = \boldsymbol{R}\}$ over all possible strategies of length k.

Let C_N abbreviate a *change-to-a-new-node* move where the player chooses one of the previously unvisited nodes at random. Theorem 3.1 shows that the class of optimal strategies is equal to the class of strategies where the player moves to an unvisited node on the last round, that is, the class of strategies where the kth move is a C_N. The next two lemmas are needed to prove Theorem 3.1.

Lemma 3.1. For any node $r \notin \mathcal{D}_k$:

$$P_\Phi\{H_k = h_k \mid R = r\} = \frac{P_\Phi\{V_0 = v_0\}}{\prod_{j=0}^{k-1} (n-j-2+y_j)} \prod_{j=1}^{k-1} P_\Phi\{V_j = v_j \mid H_j = h_j\},$$

where

$$y_j = \begin{cases} 1, & \text{if } r = v_j, \\ 0, & \text{if } r \neq v_j. \end{cases}$$

Proof. Since

$$\begin{aligned} P_\Phi\{H_k = h_k \mid R = r\} &= P_\Phi\{H_{k-1} = h_{k-1}, V_{k-1} = v_{k-1}, D_k = d_k \mid R = r\} \\ &= P\{D_k = d_k \mid H_{k-1} = h_{k-1}, V_{k-1} = v_{k-1}, R = r\} \\ &\quad \times P_\Phi\{V_{k-1} = v_{k-1} \mid H_{k-1} = h_{k-1}, R = r\} \\ &\quad \times P_\Phi\{H_{k-1} = h_{k-1} \mid R = r\}, \end{aligned}$$

and on the kth round the announcer has either $n-k$ or $n-k-1$ available nodes (depending on whether $v_{k-1} = r$ or not) it follows that

$$P_\Phi\{H_k = h_k \mid R = r\} = \frac{P_\Phi\{V_{k-1} = v_{k-1} \mid H_{k-1} = h_{k-1}, R = r\} P_\Phi\{H_{k-1} = h_{k-1} \mid R = r\}}{n-k-1+y_{k-1}}.$$

Successive applications of this formula and the observation that

$$P_\Phi\{V_{j-1} \mid H_{j-1}, R\} = P_\Phi\{V_{j-1} \mid H_{j-1}\}$$

yield the result.

Lemma 3.2. For $k \geqslant 2$,

$$\sum_{\text{all } h_k} P_\Phi\{V_0 = v_0\} \prod_{j=1}^{k-1} P_\Phi\{V_j = v_j \mid H_j = h_j\} = (n-1)(n-2)\cdots(n-k),$$

where the sum is over all possible histories $h_k = (v_0, d_1, v_1, \ldots, v_{k-1}, d_k)$, i.e. over all v_i's and d_i's in $\{1, 2, \ldots, n\}$ such that $v_j \notin \{d_1, \ldots, d_{j-1}\}$ and $d_j \notin \{d_1, d_2, \ldots, d_{j-1}, v_{j-1}\}$.

Proof. The proof is by induction on the number of rounds k. For $k = 2$,

$$\begin{aligned} &\sum_{\text{all } h_2} P_\Phi\{V_0 = v_0\} P_\Phi\{V_1 = v_1 \mid H_1 = h_1\} \\ &\quad = \sum_{v_0, d_1 \neq v_0, v_1 \neq d_1, d_2 \notin \{d_1, v_1\}} P_\Phi\{V_0 = v_0\} P_\Phi\{V_1 = v_1 \mid H_1 = h_1\} \end{aligned}$$

$$= \sum_{\text{all } h_1} P_\Phi\{V_0 = v_0\} \sum_{v_1 : v_1 \neq d_1} P_\Phi\{V_1 = v_1 \mid \boldsymbol{H}_1 = h_1\} \sum_{d_2 : d_2 \notin \{v_1, d_1\}} 1$$

$$= (n-2) \sum_{\text{all } h_1} P_\Phi\{V_0 = v_0\}$$

$$= (n-2) \sum_{v_0} \sum_{d_1 : d_1 \neq v_0} P_\Phi\{V_0 = v_0\}$$

$$= (n-2)(n-1).$$

For $k = p+1$, the quantity of interest is

$$\sum_{\text{all } h_{p+1}} P_\Phi\{V_0 = v_0\} \prod_{i=1}^{p} P_\Phi\{V_i = v_i \mid \boldsymbol{H}_i = h_i\}$$

$$= \sum_{\text{all } h_p} P_\Phi\{V_0 = v_0\} \prod_{i=1}^{p-1} P_\Phi\{V_i = v_i \mid \boldsymbol{H}_i = h_i\}$$

$$\times \sum_{v_p : v_p \notin \mathcal{D}_p} \sum_{d_{p+1} : d_{p+1} \notin \{v_p\} \cup \mathcal{D}_p} P_\Phi\{V_p = v_p \mid \boldsymbol{H}_p = h_p\},$$

and a calculation similar to the case $k = 2$ shows that

$$\sum_{v_p : v_p \notin \mathcal{D}_p} \sum_{d_{p+1} : d_{p+1} \notin \{v_p\} \cup \mathcal{D}_p} P_\Phi\{V_p = v_p \mid \boldsymbol{H}_p = h_p\} = n-p-1.$$

Thus,

$$\sum_{\text{all } h_{p+1}} P_\Phi\{V_0 = v_0\} \prod_{i=1}^{p} P_\Phi\{V_i = v_i \mid \boldsymbol{H}_i = h_i\}$$

$$= (n-p-1) \sum_{\text{all } h_p} P_\Phi\{V_0 = v_0\} \prod_{i=1}^{p-1} P_\Phi\{V_i = v_i \mid \boldsymbol{H}_i = h_i\}$$

and, assuming the assertion holds for all strategies of length p, the result follows from the induction hypothesis.

Theorem 3.1. The collection of optimal strategies is the (non-empty) set of strategies for which the final move is a C_N. That is, a strategy is optimal if and only if (with it) the player always can, and does, move to an unvisited node on the last round. In the game with n nodes and k rounds, the winning probability of any such strategy is

$$\frac{n-1}{n(n-k-1)}.$$

Proof. Successive conditioning and Lemma 3.1, together with the uniform distribution of $\boldsymbol{R}$ and the fact that $P_\Phi\{V_k \mid \boldsymbol{H}_k, \boldsymbol{R}\} = P_\Phi\{V_k \mid \boldsymbol{H}_k\}$, give

$$P_\Phi\{\text{win}\} = \frac{1}{n} \sum_{r=1}^{n} P_\Phi\{V_k = r \mid \boldsymbol{R} = r\}$$

$$= \frac{1}{n} \sum_{r=1}^{n} \sum_{\text{all } h_k} P_\Phi\{H_k = h_k, V_k = r \mid R = r\}$$

$$= \frac{1}{n} \sum_{r=1}^{n} \sum_{\text{all } h_k} P_\Phi\{V_k = r \mid H_k = h_k\} P_\Phi\{H_k = h_k \mid R = r\}$$

$$= \frac{1}{n} \sum_{r=1}^{n} \sum_{\text{all } h_k} \left(P_\Phi\{V_k = r \mid H_k = h_k\} P_\Phi\{V_0 = v_0\} \times \prod_{j=1}^{k-1} P_\Phi\{V_j = v_j \mid H_j = h_j\} \Big/ \prod_{j=0}^{k-1} (n-j-2+y_j) \right),$$

with the y_j's as in Lemma 3.1 ($y_j = 1$, if $v_j = r$; $y_j = 0$, otherwise). An upper bound for $P_\Phi\{\text{win}\}$ can be obtained from the last expression by taking all the y_j's equal to zero. This substitution and Lemma 3.2 lead to:

$$P_\Phi\{\text{win}\} \leqslant \frac{1}{n \prod_{j=0}^{k-1} (n-j-2)} \sum_{r=1}^{n} \sum_{\text{all } h_k} \left(P_\Phi\{V_k = r \mid H_k = h_k\} P_\Phi\{V_0 = v_0\} \times \prod_{j=1}^{k-1} P_\Phi\{V_j = v_j \mid H_j = h_j\} \right)$$

$$= \frac{1}{n \prod_{j=0}^{k-1} (n-j-2)} \sum_{\text{all } h_k} P_\Phi\{V_0 = v_0\} \prod_{j=1}^{k-1} P_\Phi\{V_j = v_j \mid H_j = h_j\}$$

$$= \frac{n-1}{n(n-k-1)}.$$

This inequality becomes an equality if and only if, for every r and every history h_k with positive probability under Φ (i.e. $P_\Phi\{H_k = h_k\} > 0$), either $y_0 = y_1 = \cdots = y_{k-1} = 0$ or $P_\Phi\{V_k = r \mid H_k = h_k\} = 0$ (or both). In the first case ($y_0 = y_1 = \cdots = y_{k-1} = 0$), equality holds because the corresponding term in the sum is equal to the same term in the upper bound. In the second case ($P_\Phi\{V_k = r \mid H_k = h_k\} = 0$), equality holds because the corresponding term in the sum (and in the upper bound) is equal to zero. Thus, whenever $P_\Phi\{V_k = r \mid H_k = h_k\} > 0$, equality holds if and only if *all* the y_j's are equal to 0 or $r \notin \{v_0, v_1, \ldots, v_{k-1}\}$. Therefore, a necessary and sufficient condition for equality is that all the values that V_k can take with (strictly) positive probability correspond to nodes which were not visited in the first $k-1$ rounds. It follows that a strategy is optimal if and only if its final move is a C_N.

The last theorem completely characterizes the class of optimal strategies for the general game with k rounds and a fair announcer. A C_N is at any point at least as good as any other move (this can be seen by applying the theorem to the games with $1,2,\ldots,k-1$ rounds) and a final C_N guarantees the maximum probability of winning. Thus the strategy $C_N\cdots C_N$ is both greedy and optimal, and all the strategies that end in a C_N are equivalent.

It should be emphasized that Theorem 3.1 is a general result which encompasses all the possible strategies for the player, non-randomized and randomized alike. Also, the theorem assumes that in the player's strategy the final C_N move is *always* feasible. This assumption may require a constraint between the number of nodes and the number of rounds which is more restrictive than $n \geqslant k+2$ (see Section 2). For example, if the player chooses the greedy strategy $C_N\cdots C_N$, then necessarily $n \geqslant 2k+1$; otherwise, on the last round the player will not be able to visit a new node, if all the nodes discarded by the announcer have never been visited by the player.

4. Extensions

4.1. *The announcer's behavior.* This section concerns the case of a zero-sum game, where the announcer wishes to minimize the player's probability of winning. The game has a value because of the well-known minimax theorem [3]. Theorem 4.1 gives the value of the game and shows that the player attains this value if and only if she adapts the strategy of staying with her original choice until the last round, and then moves to another curtain. This strategy is denoted by $S\cdots SC$, with S denoting a *stay-on-the-same-node* move where the player's conditional distribution has all its mass on the current node. The optimal strategy for the announcer (denoted by Δ) is to randomize among all nodes available to him on every round.

The proof of Theorem 4.1 requires Lemma 4.1, which is of interest in its own right. Let θ denote the follow-the-player (or, eliminate-old-nodes) strategy for the announcer. Lemma 4.1 shows that whenever the player chooses a strategy that ends in a C_N move (and is different from $S\cdots SC_N$), a malevolent announcer is better off by following θ than by being fair.

Lemma 4.1. For any strategy of the form ΦC_N,

$$P_{\Phi C_N,\theta}\{\text{win}\} \leqslant P_{\Phi C_N,\Delta}\{\text{win}\},$$

with a strict inequality unless $\Phi = S\cdots S$.

Proof. Let $|A|$ denote the cardinality of the set A and suppose that the game has k rounds. Let $\mathcal{V}_j$ denote the set $\{V_0, V_1,\ldots,V_j\}$ of the random variables which specify the nodes that the player visits in the first j rounds $(0 \leqslant j \leqslant k)$. Using this notation,

$$\begin{aligned} P_{\Phi C_N,\theta}\{\text{win}\} &= P_{\Phi C_N,\theta}\{V_k = R\} \\ &= \sum_{r=1}^{n} P_{\Phi C_N,\theta}\{V_k = r \mid R = r\}P\{R = r\} \\ &= \sum_{r=1}^{n} P_{\Phi C_N,\theta}\{V_k = r, r \notin \mathcal{V}_{k-1} \mid R = r\}P\{R = r\} \end{aligned}$$

$$= \frac{1}{n} \sum_{r=1}^{n} P_{\Phi C_N, \theta}\{V_k = r, r \notin \mathcal{V}_{k-1} \mid R = r\}$$

$$= \frac{1}{n} \sum_{r=1}^{n} (P_{\Phi C_N, \theta}\{V_k = r \mid r \notin \mathcal{V}_{k-1}, R = r\}$$

$$\times P_{\Phi C_N, \theta}\{r \notin \mathcal{V}_{k-1} \mid R = r\}).$$

The probability $P_{\Phi C_N, \Psi}\{V_k = r \mid r \notin \mathcal{V}_{k-1}, R = r\}$ is

$$P_{\Phi C_N, \Psi}\{V_k = r \mid r \notin \mathcal{V}_{k-1}, R = r\} = \frac{1}{n - |\mathcal{D}_k \cup \mathcal{V}_{k-1}|},$$

for any strategy Ψ of the announcer. In particular, when $\Psi = \theta$ this probability is just

$$\frac{1}{n-k-1}.$$

Thus,

$$\begin{aligned} P_{\Phi C_N, \theta}\{\text{win}\} &= \frac{1}{n(n-k-1)} \sum_{r=1}^{n} P_{\Phi C_N, \theta}\{r \notin \mathcal{V}_{k-1} \mid R = r\} \\ &= \frac{1}{n(n-k-1)} \sum_{r=1}^{n} P_{\Phi C_N, \theta}\{V_0 \neq r, V_1 \neq r, \ldots, V_{k-1} \neq r \mid R = r\} \\ &\leqslant \frac{1}{n(n-k-1)} \sum_{r=1}^{n} P_{\Phi C_N, \theta}\{V_0 \neq r \mid R = r\} \\ &= \frac{1}{n(n-k-1)} \sum_{r=1}^{n} \frac{n-1}{n} \\ &= \frac{n-1}{n(n-k-1)} \\ &= P_{\Phi C_N, \Delta}\{\text{win}\} \quad \text{(by Theorem 3.1)}, \end{aligned}$$

and it is easy to see that an equality obtains only if $\Phi = S \cdots S$.

Theorem 4.1. In the $G(k)$ zero-sum game, $S \cdots SC$ is the maxmin strategy for the player, Δ is the minimax strategy for the announcer and $(n-1)/n(n-k-1)$ is the value of the game.

Proof. From Theorem 3.1 and the fact that $P_{S \cdots SC, \Psi}\{\text{win}\}$ does not depend on Ψ,

$$\begin{aligned} \min_{\Psi} \max_{\Phi} P_{\Phi, \Psi}\{\text{win}\} &\leqslant \max_{\Phi} P_{\Phi, \Delta}\{\text{win}\} \\ &= \frac{n-1}{n(n-k-1)} \end{aligned}$$

$$= \min_{\Psi} P_{S\cdots SC,\Psi}\{\text{win}\}$$

$$\leqslant \max_{\Phi} \min_{\Psi} P_{\Phi,\Psi}\{\text{win}\}.$$

Since $\max_{\Phi} \min_{\Psi} P_{\Phi,\Psi}\{\text{win}\} \leqslant \min_{\Psi} \max_{\Phi} P_{\Phi,\Psi}\{\text{win}\}$ is always true, the value of the game is $(n-1)/n(n-k-1)$ and $S\cdots SC$ is a maxmin strategy. Lemma 4.1 shows that any other strategy ending in C_N is not maxmin and Theorem 3.1 shows that any strategy not ending in C_N cannot be maxmin.

In the light of this result, the player does not have to worry much about encountering an adverse announcer. The player can ensure herself, regardless of the announcer's behavior, of a probability of winning which is equal to the largest probability of winning with a fair announcer. The catch, of course, is that Theorem 4.1 requires that the number of rounds k be known; otherwise, $S\cdots SC_N$ cannot be realized.

4.2. *Unknown number of rounds.* This section considers the case where the number of rounds k is unknown but determined by a random variable K with a known distribution. In this game, the player has a strategy which is optimal for any value of K and, therefore, optimal regardless of k. This strategy is the greedy strategy $C_N\cdots C_N$ (always move to a new node) which is obviously not affected by the value of K. The next theorem shows that $C_N\cdots C_N$ is the unique optimal strategy when k is not known but determined by a probabilistic mechanism.

Theorem 4.2. Consider a game with a fair announcer in which the number of rounds is determined by some probabilistic mechanism. Let K denote the random variable that describes the length of the game and assume that the support of K is equal to $\{1,2,\ldots,l\}$, with $l < \frac{1}{2}n$. Then, the strategy $C_N\cdots C_N$ (l C_N's) is the unique strategy which maximizes the player's probability of winning.

Proof. From

$$P_{\Phi}\{\text{win}\} = \sum_{k=1}^{l} P_{\Phi}\{\text{win} \mid K = k\} P\{K = k\}$$

and the assumption that all the $P_{\Phi}\{K = k\}$ terms in the sum are positive, it follows that if there exists a strategy Φ which maximizes $P_{\Phi}\{\text{win} \mid K = k\}$ for all k between 1 and l it must be optimal. The strategy $C_N\cdots C_N$ is well-defined (since $l < \frac{1}{2}n$), and Theorem 3.1 implies that it is the only strategy which satisfies this property.

4.3. *The zero-sum game with an unknown number of rounds.* The minimax theorem shows that the game has a value even when the announcer is malevolent and the number of rounds is random. At this point, however, no general result is available for the optimal strategies of the two contestants.

On the basis of Lemma 4.1, it is tempting to conjecture that the follow-the-player strategy (θ) is minimax for the announcer. Yet this strategy suffers from the drawback that it may force the announcer to effectively reveal the identity of the winning node (if the player accidentally stumbles upon that node during the course of the game). Based on this observation, we conjecture that the announcer's optimal strategy consists of randomizing on every round between the follow-the-player and the fair strategy. We also conjecture that the player's optimal strategy randomizes between the following moves: (a) move to an unvisited node, (b) stay on her current node and (c) return (with some positive probability) to the last node she visited, when the announcer did not follow her. The probabilities with which the announcer and player randomize on every round depend on the history of the game and the distribution of K.

Acknowledgements

The authors thank Libby Black for motivating this work, Marco Scarsini for his valuable suggestions, Mark Knowles, Jason Papastavrou and Evangelos Tabakis for their careful reading of the first draft, and the referee for a clarification on the formulation of the main result.

References

[1] ENGEL, E. AND VENETOULIAS, A. (1991) Monty Hall's probability puzzle. *Chance* **4** (2), 6–10.
[2] ENGEL, E. AND VENETOULIAS, A. (1989) Let's make a deal. Technical Report #TR-66, Center for Computational Research in Economics and Management Science, MIT, Cambridge, Massachusetts.
[3] FERGUSON, T. S. (1967) *Mathematical Statistics: A Decision Theoretic Approach.* Academic Press, New York.
[4] KAHNEMAN, D., SLOVIC, P. AND TVERSKY, A. (1982) *Judgement under Uncertainty: Heuristics and Biases.* Cambridge University Press, New York.
[5] NALEBUFF, B. (1987) Puzzles: choose a curtain, duel-ity, two point conversions, and more. *Journal of Economic Perspectives* **1**, 157–163.
[6] SELVIN, S. (1975) A problem in probability. In 'Letters to the Editor', *The American Statistician* **29**, 67.
[7] SELVIN, S. (1975) On the Monty Hall problem. In 'Letters to the Editor', *The American Statistician*, **29**, 134.