Games with Finite Resources

Thomas S. Ferguson[1], **Costis Melolidakis**[2]

[1]Department of Mathematics, University of California at Los Angeles, Los Angeles, CA 90095, USA (email: `tom@math.ucla.edu`)
[2]Department of Mathematics, University of Athens, Panepistemiopolis, GR 157 84, Athens Greece

Abstract. Games with Finite Resources as defined by Gale (1957) are two-person zero-sum N-stage games in which each player has N resources and may use each resource once and only once in the N stages. Gale's theorem on these games is generalized in several directions. First the payoff is allowed to be any symmetric function of the stage payoffs. Second, the players are allowed some latitude in choosing which game is being played. Applications are given to some open questions in the area of Inspection Games. Finally the payoff is allowed to be random, thus incorporating a result of Ross (1972) on Goofspiel. Application is made to a game-theoretic version of the Generalized House Selling Problem.

Keywords: Goofspiel, Inspection Games, Generalized House Selling, Sequential Assignment Problem.

---

1. A Theorem of Gale

Games with Finite Resources are two-person zero-sum multistage games defined by David Gale (1957) to have the following structure.

Player I's resource set is $A = \{1, 2, \ldots, N\}$. Player II's resource set is $B = \{1, 2, \ldots, N\}$. Associated with these resources is an $N \times N$ payoff matrix $\boldsymbol{M} = (M(i, j))$. The game is played in N stages and each player is allowed to use each resource once and only once during these N stages.

At stage 1, the players simultaneously choose $a_1 \in A$ and $b_1 \in B$ and there is an immediate payoff of $M(a_1, b_1)$.

At stage k, the players simultaneously choose $a_k \in A - \{a_1, \ldots, a_{k-1}\}$ and $b_k \in B - \{b_1, \ldots, b_{k-1}\}$ and there is a payoff of $M(a_k, b_k)$. It is assumed that the players know which resources have been used up to stage k.

At stage N, the players use their remaining resources a_N and b_N and the payoff is $M(a_N, b_N)$. The total payoff is thus $\sum_1^N M(a_i, b_i)$.

Example 1. *Baby Goofspiel.* Player I is given the 13 hearts and Player II the 13 diamonds of an ordinary deck of cards. They simultaneously play a card from their hands and the higher card (Ace is low) wins the value of the lower card (Ace counts 1, Jack counts 11, Queen 12 and King 13). Play continues until all cards have been played (N=13 rounds). The matrix is

$$
\boldsymbol{M} = \begin{array}{c} \\ A \\ 2 \\ 3 \\ \vdots \\ Q \\ K \end{array}
\begin{array}{c}
\begin{array}{cccccc} A & 2 & 3 & & Q & K \end{array} \\
\begin{pmatrix}
0 & -1 & -1 & \dots & -1 & -1 \\
1 & 0 & -2 & \dots & -2 & -2 \\
1 & 2 & 0 & & -3 & -3 \\
\vdots & & & \ddots & & \vdots \\
1 & 2 & 3 & \dots & 0 & -12 \\
1 & 2 & 3 & \dots & 12 & 0
\end{pmatrix}
\end{array}
$$

This game of finite resources is symmetric, so the value is 0. After the first move, the game may no longer be symmetric and it seems as if the optimal strategies may be quite complex. However, Gale's remarkable result for arbitrary Games of Finite Resources is the following.

Theorem 1. *(Gale (1957)) The value of the Game With Finite Resources is*

$$V = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{N} M(i,j).$$

An optimal strategy for Player I is to choose a permutation $(a_1,\dots,a_N)$ of $\{1,2,\dots,N\}$ at random with probability $1/N!$ each, and to use a_k at stage k. Similarly for Player II.

A remarkable feature of this result is that the players' optimal strategies do not use any of the information acquired along the way. Indeed, we could alter the rules of the game so that Player II is not told the pure strategy choices of Player I as the game progresses. The lack of information does not hurt him. In addition, this strategy does not depend on the payoff matrix, $\boldsymbol{M}$. It could be that the entries to the matrix are chosen at random according to some distribution known to the players, but that only Player I is informed of the matrix that was chosen. The seeming advantage Player I gets from this information is of no use to him.

We shall refer to the optimal strategies found in Theorem 1 as *Gale strategies.*

Gale's result is somewhat more general than this. The elements of A and B are allowed to be elements of linear spaces and the payoff to be a bilinear functional of the chosen a and b. Also both players are allowed to select from a collection of N-element sets, provided that the sum of the elements in each set is the same constant. Gale's general result is extended in Theorem 4 below.

In this paper, Gale's Theorem is generalized in three directions. In Section 2, we note that Gale's result remains true if the final payoff is allowed to be any symmetric function of the stage payoffs. This allows treatment of problems where the player who wins the larger total amount is declared the winner of the game. In Section 3, we allow the players during pregame play to jointly choose the matrix to be used in the game. We also give some leeway in selecting the matrix as play progresses. This will allow us to make some far-reaching generalizations of inspection games. In Section 4, we allow the payoffs to be influenced by a sequence of random variables. This generalization contains a result of Sheldon Ross (1971) on hidden card goofspiel. Finally we examine the implication of these ideas on the sequential assignment problem of Derman, Lieberman and Ross (1972), sometimes called the generalized house selling problem.

We conclude this section with another example showing that this simple form of Gale's Theorem can already be used to extend certain results on the inspection game. This possibility was suggested in Ferguson and Melolidakis (1998).

Example 2. *An Inspection Game.* The inspector has $k \le N$ agents of differing abilities. The smuggler has $l \le N$ shipments of differing values that he must ship within the next N days. The inspector may use each agent only once and at most one agent per day. The inspector decides sequentially which agent to use on which day. The smuggler decides sequentially which contraband to smuggle on which day. If agent i is inspecting on the day that contraband j is being sent, the payoff to the inspector is m_{ij}, $i = 1, \ldots, k$, $j = 1, \ldots, l$. If the smuggler sends contraband j on a day the inspector is not inspecting, the payoff to the smuggler is s_j. On days the smuggler does nothing, the payoff is zero. The inspector learns after the fact of any successful smuggling and of the type of contraband smuggled. However, the smuggler only learns of those inspectors used on the days he tries to smuggle.

We may set this problem up as a game of finite resources as follows. Let the first k elements of $A = \{1, 2, \ldots, N\}$ denote the agents and the last $N - k$ elements denote do-nothings. Similarly, let the first l elements of $B = \{1, 2, \ldots, N\}$ denote the different contrabands and the last $N - l$ elements denote do-nothings. The associated matrix $\boldsymbol{M}$ is

$$M(i,j) = \begin{cases} m_{ij} & \text{if } 1 \le i \le k \text{ and } 1 \le j \le l \\ -s_j & \text{if } k < i \le N \text{ and } 1 \le j \le l \\ 0 & \text{if } l < j \le N \end{cases}$$

The value of the game is $V = (1/N)[\sum_{i=1}^{k} \sum_{j=1}^{l} m_{ij} - (N-k) \sum_{j=1}^{l} s_j]$ and the Gale strategies are optimal. This result for $l = 1$ and all m_{ij} equal is one of the problems solved by Dresher (1962) (see Section 8). Later, Thomas and Nisgav (1976) independently found the same result. The case of arbitrary l and all m_{ij} equal and all s_j equal is in Sakaguchi (1977). The case of arbitrary l and $m_{ij} = 1$ for $i \le h$, $m_{ij} = -1$ for $h < i \le k$, and $s_j = 0$ is in Nakai (1980).

2. Allowing a more general payoff

As a first generalization of games with finite resources, we allow the terminal payoff to be a more general function of the stage payoffs rather than just the sum. Suppose the final payoff of the game is some function, $g(M(a_1, b_1), \ldots, M(a_N, b_N))$, where $M(a_1, b_1), \ldots, M(a_N, b_N)$ are the stage payoffs. The following theorem states that Gale's

Theorem remains true if the function $g(x_1, \ldots, x_N)$ is symmetric in its arguments, i.e. $g(x_1, \ldots, x_N) = g(x_{\pi(1)}, \ldots, x_{\pi(N)})$ for all $\pi \in \Pi_N$, the set of all permutations $(\pi(1), \ldots, \pi(N))$ of $(1, \ldots, N)$. An example given in Gale's paper shows that this result is not necessarily true for arbitrary functions, g. The proof of Theorem 2 follows Gale's proof closely with some simplification.

Theorem 2. *The optimal strategies given in Theorem 1 are still optimal if the final payoff is a symmetric function, g, of the stage payoffs, but the value is now*

$$V = \frac{1}{N!} \sum_{\pi \in \Pi_N} g(M(1, \pi(1)), \ldots, M(N, \pi(N))).$$

Proof. We show that if Player II, say, uses a Gale strategy, then all choices of the first move by Player I lead to the same expected payoff. We show this by induction on N. It is trivially true for $N = 1$. Suppose it is true for games of $N-1$ stages. If Player I uses a_1 and Player II uses b_1 on the first stage of the N-stage game, the induction hypothesis implies that the average payoff is

$$V(a_1, b_1) = \frac{1}{(N-1)!} \sum_{\pi \in \Pi_N(a_1, b_1)} g(M(1, \pi(1)), \ldots, M(a_1, b_1), \ldots, M(N, \pi(N))),$$

where $\Pi_N(a_1, b_1) = \{\pi \in \Pi_N : \pi(a_1) = b_1\}$. This uses the assumed symmetry of g. If Player II chooses b_1 from $\{1, \ldots, N\}$ at random with probability $1/N$ each, then the average payoff to Player I is

$$\begin{aligned} V(a_1) &= \frac{1}{N} \sum_{j=1}^{N} V(a_1, j) \\ &= \frac{1}{N} \sum_{j=1}^{N} \frac{1}{(N-1)!} \sum_{\pi \in \Pi_N(a_1, j)} g(M(1, \pi(1)), \ldots, M(a_1, j), \ldots, M(N, \pi(N))) \\ &= \frac{1}{N!} \sum_{\pi \in \Pi_N} g(M(1, \pi(1)), \ldots, M(N, \pi(N))) = V. \end{aligned}$$

This completes the induction. ∎

In Example 1 above, we might say the player wins the game if he has the most points at the end of the game. The final payoff is the sign of the sum of the stage payoffs, $g(x_1, \ldots, x_N) = \mathrm{sgn}(x_1 + \cdots + x_N)$, and so is a symmetric function of the stage payoffs. Or in Example 2, the inspector might win if he catches the spy at least once. More interesting would be to let the payoff be the standard deviation of the stage payoffs, with Player I trying to achieve a wide spread of the payoffs, and Player II trying to achieve a level payoff.

Example 3. *Divide and Conquer*, by Claude Soucie. This game is described by Sid Sackson (1994) p.135-136. Player I is given five cards drawn at random from a deck of ten cards consisting of the $2, 3, \ldots, 10$, and Q. The remaining five cards are given to player II.

The game is played in two rounds of five stages each. At each stage of the first round, the players simultaneously choose a card from their hand and the winner for that stage is determined by the rules: The higher card wins, unless the lower card is one unit below the higher card, or unless the value of the lower card divides the value of the higher card (Q is considered 12), in which case the lower card wins. The cards played at each stage are set aside and when the five stages are finished, the two players exchange their initial hands and play five more stages with the same rules. The winner is the player who has won at least six stages in the two rounds combined.

This is slightly more complicated than the other examples because Theorem 2 must be applied twice. First, consider the second round. Let x denote the number of wins Player I has achieved upon entering the second round. Then conditional on x, the total payoff, being a function of the sum of the remaining stage payoffs, is a symmetric function of those payoffs. So Theorem 2 applies to the second round, and the optimal strategies are Gale strategies, guaranteeing some $v(x)$ for each x. Therefore, in the first round, it is as if the two players play a finite resource game, the evaluation function being $v(x)$, which is again a function of the sum of the first round payoffs and hence symmetric. This shows that in the first round the two players will also play Gale strategies. Since their optimal strategies are independent of the particular $\mathcal{M}$ initially chosen, this chance move is of no consequence.

3. Allowing the payoff matrix to be chosen jointly by the players

Suppose that instead of a single matrix $\boldsymbol{M}$, a set of $N \times N$ matrices, $\{\boldsymbol{M}_{s,t}\}$ for $s = 1, \ldots, S$ and $t = 1, \ldots, T$, is given. Let Γ be the game in which Player I chooses s, and Player II chooses t simultaneously, and then the finite resource game with matrix $\boldsymbol{M}_{st}$ is played à la Gale. We have the following rather obvious theorem.

Theorem 3. *The game* Γ *has value* $V = \mathrm{Val}(\boldsymbol{M}^*)$, *where* $\boldsymbol{M}^*$ *is the* $S \times T$ *matrix*

$$\boldsymbol{M}^* = \left(\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} M_{st}(i,j) \right)$$

and Val *denotes the value operator applied to matrices. An optimal strategy for Player I is to choose* s *according to any optimal strategy for the matrix game* $\boldsymbol{M}^*$ *and then choose a permutation of* $(1, 2, \ldots, N)$ *with probability* $1/N!$ *each. Similarly for player II.*

Once s and t have been chosen, the players' optimal strategies are independent of the information received and of the matrix selected and of its values. This is also the solution if one of the players is not informed of which game is being played and his opponent is informed. However, it is assumed that once s and t have been chosen, they may not be changed. Certain applications to inspection games require that we allow the players to select the matrix as the game proceeds. For this we need a different approach.

Let A and B be nonempty (possibly infinite) sets representing the possible actions of the players. Let S and T denote (possibly infinite) sets of permissible sequences of actions for the N-stage game. We take S and T to be arbitrary nonempty sets of N-tuples of elements of A and B respectively, satisfying the sole condition that they are closed under permutations. That is, we assume

$$\begin{aligned} \boldsymbol{s} = (a_1, \ldots, a_N) \in S \quad &\text{implies} \quad (a_{\pi_1}, \ldots, a_{\pi_N}) \in S \quad \text{for all } \pi \in \Pi_N \\ \boldsymbol{t} = (b_1, \ldots, b_N) \in T \quad &\text{implies} \quad (b_{\pi_1}, \ldots, b_{\pi_N}) \in T \quad \text{for all } \pi \in \Pi_N. \end{aligned}$$

The N-stage game is played as follows. Given are A, B, S, T and a real valued function u on $A \times B$. At the first stage, Player I chooses some $a_1 \in A$ that is the first element of some $\boldsymbol{s} \in S$. Simultaneously, Player II chooses some $b_1 \in B$ that is the first member of some $\boldsymbol{t} \in T$. Then a_1 and b_1 are announced to the players, and Player I receives an amount $u(a_1, b_1)$ from Player II, where u is some known real function defined on $A \times B$. Then play proceeds to the next stage where Players I and II simultaneously choose $a_2 \in A$ and $b_2 \in B$ such that (a_1, a_2) and (b_1, b_2) are the first two elements of some $\boldsymbol{s} \in S$ and $\boldsymbol{t} \in T$. This continues for all N stages, the final payoff being the sum of the stage payoffs.

For a given element $\boldsymbol{s} \in S$, we let $\boldsymbol{s}^*$ denote the Gale strategy that selects an equiprobable random permutation of the elements of $\boldsymbol{s}$. Similarly for $\boldsymbol{t}^*$. The following theorem states that if two given strategies satisfy a strong saddle-point property, they will be optimal.

Theorem 4. *If there exists* $\hat{\boldsymbol{s}} = (\hat{a}_1, \ldots, \hat{a}_N) \in S$ *and* $\hat{\boldsymbol{t}} = (\hat{b}_1, \ldots, \hat{b}_N) \in T$ *such that*

$$\max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} u(a_i, \hat{b}_j) = \sum_{i=1}^{N} u(\hat{a}_i, \hat{b}_j) \quad \textit{for all } j = 1, \ldots, N \tag{1}$$

and

$$\min_{\boldsymbol{t} \in T} \sum_{j=1}^{N} u(\hat{a}_i, b_j) = \sum_{j=1}^{N} u(\hat{a}_i, \hat{b}_j) \quad \textit{for all } i = 1, \ldots, N, \tag{2}$$

then the strategies $\hat{\boldsymbol{s}}^*$ *and* $\hat{\boldsymbol{t}}^*$ *are optimal, and the value is*

$$V = \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{N} u(\hat{a}_i, \hat{b}_j). \tag{3}$$

Proof. We will show that if Player II uses $\hat{\boldsymbol{t}}^*$, then Player I can achieve no more than V. The theorem will then follow since by symmetry Player I would be able to achieve at least V using $\hat{\boldsymbol{s}}^*$.

Suppose then that Player II uses $\hat{\boldsymbol{t}}^*$. We may assume without loss of generality that all $\hat{b}_j$ are distinct, since we could enlarge B with dummy equivalents to the duplicated $\hat{b}_j$ and extend the definition of u appropriately. We allow Player I considerably more leeway in choosing strategies. We allow him to choose a different strategy for each $\hat{b}_j$. That is, we pretend that Player I may hire N brokers, $j = 1, 2, \ldots, N$, each of which may use a different strategy, where the payoff broker j receives by choosing a is

$$u_j(a, b) = \begin{cases} u(a, b) & \text{if } b = \hat{b}_j \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Player I receives the sum of the returns from all the brokers. Clearly this cannot hurt Player I because all the brokers could use the same strategy and this would give him what he could achieve on his own. We show by induction on N, that the largest expected return broker j can achieve against the strategy $\hat{\boldsymbol{t}}^*$ of Player II is

$$V_j = \max_{\boldsymbol{s} \in S} \frac{1}{N} \sum_{i=1}^{N} u(a_i, \hat{b}_j). \tag{5}$$

The result is obviously true for $N = 1$. Suppose it is true for $N - 1$. Suppose broker j uses an arbitrary a_1 on the first stage. With probability $1/N$, Player II uses $\hat{b}_j$ and the payoff is $u(a_1, \hat{b}_j)$. With probability $(N-1)/N$, Player II uses one of the other elements of $\hat{\boldsymbol{t}}$ and by the induction hypothesis the expected payoff is at most

$$\frac{1}{N-1} \max_{\boldsymbol{s} \in S(a_1)} \sum_{i=2}^{N} u(a_i, \hat{b}_j) \tag{6}$$

where $S(a_1)$ is the set of permissible sequences of actions remaining to Player I after he uses a_1 at the first stage,

$$S(a_1) := \{\boldsymbol{s} = (a_2, \ldots, a_N) : \ (a_1, a_2, \ldots, a_N) \in S\}. \tag{7}$$

The total expected payoff to broker j, using a_1 on the first round, is at most

$$\begin{aligned} \frac{1}{N} u(a_1, \hat{b}_j) + \frac{N-1}{N} \frac{1}{N-1} \max_{\boldsymbol{s} \in S(a_1)} \sum_{i=2}^{N} u(a_i, \hat{b}_j) &= \frac{1}{N} [u(a_1, \hat{b}_j) + \max_{\boldsymbol{s} \in S(a_1)} \sum_{i=2}^{N} u(a_i, \hat{b}_j)] \\ &\leq \frac{1}{N} \max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} u(a_i, \hat{b}_j) = V_j, \end{aligned}$$

completing the induction. Therefore, the total payoff to Player I is bounded above by the sum

$$\sum_{j=1}^{N} V_j = \sum_{j=1}^{N} \frac{1}{N} \max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} u(a_i, \hat{b}_j).$$

This is equal to V from condition (1). ■

An example showing the usefulness of allowing A and S to be infinite sets is given in Gale's paper. Gale takes S and T to N-tuples of m-dimensional and n-dimensional vectors respectively, $\boldsymbol{s} = (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N)$ and $\boldsymbol{t} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_N)$, closed under permutations, and satisfying

$$\sum_{i=1}^{N} \boldsymbol{a}_i = \boldsymbol{a}, \quad \text{and} \quad \sum_{i=1}^{N} \boldsymbol{b}_i = \boldsymbol{b},$$

where $\boldsymbol{a}$ and $\boldsymbol{b}$ are fixed m- and n-dimensional vectors. The payoff at stage i when the players use $\boldsymbol{a}_i$ and $\boldsymbol{b}_i$ is $u(\boldsymbol{a}_i, \boldsymbol{b}_i) = \boldsymbol{a}_i' \boldsymbol{M} \boldsymbol{b}_i$ where $\boldsymbol{M}$ is a fixed m by n matrix. The overall payoff for given $\boldsymbol{s}$ and $\boldsymbol{t}$ is the sum of the stage payoffs,

$$\sum_{i=1}^{N} \boldsymbol{a}_i' \boldsymbol{M} \boldsymbol{b}_i.$$

The conclusion is that the value is $\boldsymbol{a}' \boldsymbol{M} \boldsymbol{b}/N$ and the Gale strategies are optimal.

Theorem 4 may be considered as a generalization of this result. One has only to check that conditions (1) and (2) are satisfied. But $\sum_{i=1}^{N} \boldsymbol{a}_i' \boldsymbol{M} \boldsymbol{b}_j = \boldsymbol{a}' \boldsymbol{M} \boldsymbol{b}_j$. so that the maximum is independent of $\boldsymbol{s} \in S$ and condition (1) is satisfied. Similarly for condition (2).

Example 4. *The Many-Agent Inspection Game.* Consider again the N-stage inspection game. Suppose Player I has Q types of agents, with q_i inspection agents of type i, $i = 1, \ldots, Q$. The agents of type i are allowed at most k_i inspections total, which may be allocated in any manner I wishes during the next N days (but with no more than q_i inspections of type i on a single day). Without loss of generality we assume that $k_i \leq N q_i$. We also assume that a total of no more than c agents may be used on a single day.

Player II has $1 \leq \ell < N$ identical contraband which must be delivered during the next N days. II is not allowed to deliver more than 1 contraband per day.

The first stage is played in the following way: Player I chooses the number and types of agents to be dispatched. Hence, I chooses $\boldsymbol{a}_1 := (a_{11}, \ldots, a_{1Q})$, with $0 \leq a_{1i} \leq \min\{q_i, k_i\}$, $i = 1, \ldots, Q$ and $\sum_{i=1}^{Q} a_{1i} \leq c$. Simultaneously, II chooses $b_1 = 0$ or $b_1 = 1$ as the number of contraband sent. The payoff to I for the first stage is given by $u(\boldsymbol{a}_1, b_1)$, where u is a function defined on the possible team and contraband allocations. Then, each player is informed of the other player's actions and the game moves to the next stage where the whole process is repeated. The total payoff is the sum of the stage payoffs.

Let us denote by S and T the sets of permissible sequences of actions of Players I and II in the generalized inspection game we are discussing, i.e.

$$S := \{\boldsymbol{s} = (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N) : \boldsymbol{a}_\nu = (a_{\nu 1}, \ldots, a_{\nu Q}),\ 0 \leq a_{\nu i} \leq q_i,$$
$$\sum_{j=1}^{Q} a_{\nu j} \leq c,\ \sum_{\nu=1}^{N} a_{\nu i} = k_i,\ \nu = 1, \ldots, N,\ i = 1, \ldots, Q\}$$

and

$$T := \{\boldsymbol{t} = (b_1, \ldots, b_N) : b_i \in \{0, 1\}, i = 1, \ldots, N, \sum_{i=1}^{N} b_i = \ell\}.$$

Both S and T are closed under permutations.

Note that T consists of one element and its permutations. This implies that condition (2) of Theorem 4 is satisfied. If we make the assumption that there is no payoff in those stages where Player II does not act, i.e.

$$u(\boldsymbol{a}, 0) = 0 \quad \text{for all } \boldsymbol{a}, \tag{8}$$

then condition (1) is also satisfied with $\hat{a}$ achieving the maximum in $\sum_1^N u(a_i, 1)$. The value is

$$V = \frac{\ell}{N} \max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} u(\boldsymbol{a}_i, 1), \tag{9}$$

and there are optimal Gale strategies. A counterexample to this result when condition (8) is not satisfied is given after the discussion.

This result extends and simplifies certain results in the inspection game literature. The game with $\ell = 1$, $Q = 1$, $q_1 = 1$, $k_1 = k$ and $c = 1$ is the game, discussed in Example 2 and solved by Dresher (1962) using recursive techniques. This game was extended by Thomas and Nisgav (1976) to $q_1 = 2$, $c = 2$. The inspector may use a total of k inspections, but if the two agents inspect on the same day, there may be a better result (or a higher probability of catching the smuggler). Thomas and Nisgav did not give a complete solution to this problem but instead showed how to use a linear program to solve it. They also proposed looking at the problem with $Q = 2$, $q_1 = q_2 = 1$ and $c = 2$. The agent of type 1 can be used at most k_1 times and the agent of type 2 can be used at most k_2 times.

Baston and Bostock (1991) consider this second (more difficult) problem and derive a closed form solution using recursive techniques. Garnaev (1994) succeeds in finding closed form solutions to the problem of three types of agents, $Q = 3$, $q_1 = q_2 = q_3 = 1$. These problems are indeed difficult to solve if the multistage game approach is used to solve them.

It may be checked that (9) provides the same answer as that found in the particular cases solved in the inspection game literature. As an example, suppose $\ell = 1$, $Q = 2$, and $q_1 = q_2 = 1$, $c = 2$, and let us use as notation for the payoff, $u((1,1),1) = p$, $u((1,0),1) = p_1$, $u((0,1),1) = p_2$, and 0 elsewhere. For a particular $\boldsymbol{s} \in S$, let c_0 be the number of $(1,1)$ appearing, c_1 the number of $(1,0)$ appearing, and c_2 the number of $(0,1)$ appearing. From (9), the value will be M/N, where $M := \max\{c_0 p + c_1 p_1 + c_2 p_2 : c_0 + c_1 = k_1, c_0 + c_2 = k_2\}$. Assuming $k_1 \le k_2 \le N$ and rescaling, we get the main theorem (Theorem 2.6) of Baston and Bostock (1991).

Using Theorem 4, we may also treat problems in which Player II has several distinct contraband that must be sent. For example, there are N possibly distinct contraband that must be sent, but that at most one contraband can be sent per day, so that condition (2) is still satisfied. Also suppose that the payoff for sending contraband j when action $\boldsymbol{a}$ is being used by the inspector has the form

$$u(\boldsymbol{a}, j) = c_j \phi(\boldsymbol{a})$$

for some arbitrary function ϕ and some constants $c_j \geq 0$. The worthless contraband with $c_j = 0$ correspond to the "do-nothing" actions. Then condition (1) is also satisfied because

$$\max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} u(\boldsymbol{a}_i, j) = c_j \max_{\boldsymbol{s} \in S} \sum_{i=1}^{N} \phi(\boldsymbol{a}_i)$$

so that the maximizing $\boldsymbol{s}$ may be chosen independent of j.

Counterexample when $u(\boldsymbol{a}, 0) \not\equiv 0$. Suppose $N = 3$, $Q = 1$, $q_1 = k_1 = c = 2$, and $\ell = 1$. Player I has just two permissible sequences, $\boldsymbol{s}_1 = (0, 1, 1)$ and $\boldsymbol{s}_2 = (0, 0, 2)$ and their permutations. Suppose

$$\begin{array}{lll} u(0,0) = 0 & u(1,0) = 0 & u(2,0) = -1 \\ u(0,1) = 0 & u(1,1) = 0 & u(2,1) = 2. \end{array}$$

Then both Gale strategies, $\boldsymbol{s}_1^*$ and $\boldsymbol{s}_2^*$, give an average payoff of zero against $\boldsymbol{t}^*$. However, the strategy: Play $a_1 = 0$ first. If Player II uses 0, then continue with $a_2 = 0$ and $a_3 = 2$. If Player II uses 1, then continue with $a_2 = 1$ and $a_3 = 1$. Against $\boldsymbol{t}^*$, this gives an expected payoff of 1/3. (If a counterexample with u non-decreasing in its first argument is required, change $u(1, 0) = 1$ and $u(2, 0) = 1$.)

4. Allowing Random Payoffs.

We may allow the stage payoffs to be influenced by a sequence of random variables. The results depend on whether the random variables are observed before or after the players make their choices at each stage.

Example 5. *Goofspiel.* As in Baby Goofspiel, Player I is given the 13 hearts and Player II the 13 diamonds from a deck of cards. In the game of (adult) Goofspiel, the 13 spades are shuffled and placed in the center of the table. In the first round, the top card of the spade pile is turned over for the players to see. Then they each choose a card from their hands and simultaneously play it on the table. The player who has played the higher card wins the value of the spade card showing. If the players have chosen the same card, there is no payoff. Then all three cards are removed from the game and play proceeds to the next round. Play ends after 13 rounds.

Goofspiel, also called GOPS (game of pure strategy), is described in Luce and Raiffa (1957) and studied by Ross (1971). It is indeed a difficult game to analyze. No one knows what the optimal strategy is for the first round. However, in a related game called *Hidden Card Goofspiel*, the card on top of the spade pile is turned over only after the players make their choices. For this game, Ross shows there is a very Gale-like solution, in which the players have optimal strategies that choose cards from the hands equally likely, ignoring information received along the way. The solution of this game does not seem to be derivable directly from Gale's Theorem. So it is of interest to discover a general result that will contain Gale's Theorem and the result of Ross on Hidden Card Goofspiel.

We present two theorems. In the first, the random variables are observed before the players make their choices. The set-up is as for Gale's Theorem. Players I and II each have N resources to be used one at a time in the N stages. There is a payoff matrix $M(i,j)$ for $i = 1,\ldots,N$, $j = 1,\ldots,N$, and there is also a sequence of random variables, $X_1, X_2, \ldots, X_N$, assumed to be a martingale (i.e. for all $n \geq 1$ $E(X_{n+1}|X_1,\ldots,X_n) = X_n$ a.s.). At stage k, if the players choose resources a_k and b_k, the stage payoff is $X_k \cdot M(a_k, b_k)$. The overall payoff is the sum of the stage payoffs. We have

Theorem 5. *If at each stage k, X_k is observed before the players make their choices, and if the $\{X_i\}_{i=1}^N$ form a martingale with finite mean, $E(X_1) = \mu$, then the value of the game is $V = \mu(1/N)\sum\sum M(i,j)$, and the Gale strategies are optimal.*

Proof. Follow Gale's proof (or the proof of Theorem 2) but in the inductive step, suppose Player II uses the Gale strategy and Player I sees X_1 and uses a_1. Then conditional on X_1, the expected payoff is

$$\begin{aligned}
&\frac{1}{N}\sum_{k=1}^{N}[X_1 M(a_1,k) + E(X_2|X_1)\frac{1}{N-1}\sum_{i\neq a_1}\sum_{j\neq k} M(i,j)] \\
&\qquad = X_1 \frac{1}{N}\sum_{k=1}^{N}[M(a_1,k) + \frac{1}{N-1}\sum_{i\neq a_1}\sum_{j\neq k} M(i,j)] \qquad (10)\\
&\qquad = X_1 \frac{1}{N}\sum_{i=1}^{N}\sum_{k=1}^{N} M(i,k).
\end{aligned}$$

This does not depend on a_1, and its expected value is V. ■

In this theorem, the players may not only ignore their opponents choices, they may ignore the X's as well. We may even let one of the players choose the distribution of the martingale as play proceeds, provided it starts with $E(X_1) = \mu$. That is Player I, say, can choose at stage $k > 1$ the conditional distribution of X_k given $X_1, \ldots, X_{k-1}$, provided it satisfies $E(X_k|X_1,\ldots,X_{k-1}) = X_{k-1}$.

If the players get to observe the X's only after they make their choices, we can weaken the condition on the distribution of the X's.

Definition. *A sequence of random variables, $X_1, X_2, \ldots$, is said to have stable mean or to be mean-stable if for all $k \geq 1$,*

$$E(X_k|X_1,\ldots,X_{k-1}) = E(X_{k+1}|X_1,\ldots,X_{k-1}) \quad \text{a.s.}$$

For $k = 1$ this means that $E(X_1) = E(X_2)$. Note that if $X_1, X_2, \ldots$ is mean-stable, then at any stage k, the conditional expectation of all future observations is the same, that is, for all k and $n > k$, $E(X_n|X_1,\ldots,X_{k-1}) = E(X_k|X_1,\ldots,X_{k-1})$. This is because

$$\begin{aligned}
E(X_{k+2}|X_1,\ldots,X_{k-1}) &= E(E(X_{k+2}|X_1,\ldots,X_k)|X_1,\ldots,X_{k-1}) \\
&= E(E(X_{k+1}|X_1,\ldots,X_k)|X_1,\ldots,X_{k-1}) \\
&= E(X_{k+1}|X_1,\ldots,X_{k-1}) = E(X_k|X_1,\ldots,X_{k-1})
\end{aligned}$$

and so on by induction.

A martingale is automatically mean-stable. But so also is any sequence, X_i, that is exchangeable, i.e. the distribution of $X_1, \ldots, X_N$ is invariant under permutations of the subscripts. For example, the X_i could arise as a sample from a normal distribution whose unknown mean has been chosen at random from a normal distribution with mean μ.

Theorem 6. *If at each stage k, X_k is observed after the players make their choices, and if the $\{X_i\}_{i=1}^N$ is mean stable with finite mean $E(X_1) = \mu$, then the value of the game is $V = \mu(1/N)\sum\sum M(i,j)$, and the Gale strategies are optimal.*

Proof. In the proof of Theorem 5, equation (10) is replaced by

$$\frac{1}{N}\sum_{k=1}^{N}[E(X_1)M(a_1,k) + E(E(X_2|X_1))\frac{1}{N-1}\sum_{i\neq a_1}\sum_{j\neq k}M(i,j)]$$

$$= E(X_1)\frac{1}{N}\sum_{k=1}^{N}[M(a_1,k) + \frac{1}{N-1}\sum_{i\neq a_1}\sum_{j\neq k}M(i,j)]. \blacksquare$$

In Hidden Card Goofspiel, $N = 13$ and the distribution of $X_1, \ldots, X_{13}$ is uniform over all permutations of $\{1, 2, \ldots, 13\}$ and so is exchangeable. Therefore, Theorem 6 contains the result of Ross on Hidden Card Goofspiel. In Goofspiel, the card is turned up before play, so neither Theorem 5 nor Theorem 6 applies. In fact, Ross has a very interesting result about Goofspiel. Namely, if one of the players chooses his cards completely at random (i.e. uses a Gale strategy), then the optimal response of the other player is to match the card turned up from the spade pile. This shows that Theorem 5 would be false under the hypothesis that the $X_1, \ldots, X_N$ is an exchangeable sequence. Sakaguchi (1977) and Nakai (1980) have studied the problem of Theorem 5 under the assumption that the X_i are independent and identically distributed.

5. Application to a game-theoretic version of the generalized house selling problem.

The generalized house selling (GHS) problem was introduced by Derman, Lieberman and Ross (1972) as a dynamic programming model and was subsequently generalized in various directions (Albright (1974, 1977), Righter(1987, 1988, 1989)). A game-theoretic version of the problem, extending work of Brams and Davis (1978) and Sakaguchi (1980) on peremptory jury challenges, has been treated by Nakai (1982).

In the original version (Derman et al.), a company has on hand N workers, whose values are known and denoted by $y_1, y_2, \ldots, y_N$, assumed to be arranged in nondecreasing order. A sequence of i.i.d. random variables, $X_1, X_2, \ldots, X_N$, representing jobs, arrive sequentially, and each job as it arrives must be assigned to one of the available workers. Thus, at each stage k, $k = 1, \ldots, N$, the controller observes X_k and chooses a y_{n_k} from the $N - k + 1$ remaining workers. Then y_{n_k} is removed from the set of available workers, the payoff for this stage being the product, $X_k y_{n_k}$, and the process proceeds to the next stage.

The total payoff is the sum of the payoffs over all of the stages. Derman et al. showed that when there are ν stages to go, there exist constants, $-\infty = a_{0\nu} \leq a_{1\nu} \leq a_{2\nu} \leq \ldots \leq a_{\nu\nu} = +\infty$, which may depend on the distribution of the X's but are independent of the y's, so that the policy that maximizes the expected total payoff is to pick from the remaining y's the t-th in increasing order if the observed $X_{N-\nu+1}$ falls in the interval $(a_{t-1,\nu}, a_{t,\nu}]$.

Here, we are interested in the following two-person zero-sum game-theoretic generalization of the GHS problem, which we will call the GHS game. Players I and II, with resource sets $Y = \{y_1, \ldots, y_N\}$ and $Z = \{z_1, \ldots, z_N\}$ respectively, play the following game. At each stage k, $1 \leq k \leq N$, they make simultaneous choices y_{m_k}, z_{n_k} and player I gets from II a payoff of $X_k y_{m_k} z_{n_k}$, where X_k, $k = 1, \ldots, N$, is some sequence of random variables. Whether the two players observe the X's before or after they make their choices will be discussed subsequently. The y_i's and the z_j's can be positive, negative or zero. The total payoff is the sum of the stage payoffs. Both players learn the stage payoff and their opponent's move before the next stage begins. Then, Theorems 5 and 6 lead to the following

Corollary 2. *(a) If the sequence X_k, $k = 1, \ldots, N$, is mean-stable with mean μ and is observed by the two players after they make their moves at each stage, then, the value of the game is $v = (\mu/N)(\sum y_i)(\sum z_j)$. (b) If the sequence X_k, $k = 1, \ldots, N$, is a martingale having mean μ and is observed by the two players before they make their moves at each stage, then, the value of the game is $v = (\mu/N)(\sum y_i)(\sum z_j)$. In both cases (a) and (b), the Gale strategy is an optimal strategy for both players.*

Notes. 1. In case (b), the distributions of the X sequence may be chosen along the way by either player without affecting the result. For example at stage k, the distribution of X_k may be chosen arbitrarily by Player I provided $E(X_k | X_1, \ldots, X_{k-1}) = X_{k-1}$. The player who makes this choice could even be decided by a random mechanism.

2. Instead of $X_k y_{m_k} z_{n_k}$, the payoff may be $X_k f(y_{m_k}, z_{n_k})$ at each stage k, $k = 1, \ldots, N$, for some bilinear function f, with the corresponding change in the value formula.

As is the case with Goofspiel, the above result is not necessarily valid if the X's are exchangeable and observed before the players make their moves. In the following theorem, we mimic the result of Ross on Goofspiel by finding Player I's optimal response in this case if the Gale strategy is used by Player II. This will then show that if a player plays completely at random, then the other player has a better strategy than playing completely at random. To this end we will need the following result.

Lemma 1. *Let $v(x, \{y_t\}_{t=1}^N)$ be the value function of the maximization GHS problem conditional on $X_1 = x$ and suppose that the sequence of the X's is mean-stable. Then for $c \geq 0$, $cv(x, \{y_t\}_{t=1}^N) = v(x, \{cy_t\}_{t=1}^N)$.*

Proof: Multiplication of the y's with $c \geq 0$ keeps the ordering unchanged and hence the optimal policies in all problems are identical. Writing down the optimality equations and using induction completes the proof. ■

Assume now that in the GHS problem the stage payoff depends on two mean-stable sequences of random variables with finite means, independent one from the other, X_k, $k = 1, \ldots, N$, and Z_k, $k = 1, \ldots, N$, with $Z_k > 0$ a.s. At each stage the controller observes the corresponding X before and the corresponding Z after he makes his move (i.e. choice of the corresponding y). Let us denote the joint distribution of N of the Z variables by F_N and let us call this the two sequence GHS problem. We then have,

Theorem 7. *The conditional value $V(x, \{y_t\}_{t=1}^N, F_N)$ given $X_1 = x$ for the two sequence GHS problem is given by $v(x, \{E[Z]y_t\}_{t=1}^N)$, where $E[Z]$ is the expectation of any Z variable. The optimal policy at each stage k, $k = 1, \ldots, N$, is to play optimally in the one sequence (X) corresponding GHS problem.*

Proof: The proof is by induction on N. It is easily shown for $N = 2$. Assume it is true for $N - 1$. Then, the optimality equation and the induction hypothesis give

$$\begin{aligned} V(x, \{y_t\}_{t=1}^N, F_N) &= \max_{1 \le i \le N} \left\{ xy_i E[Z_1] + E\left[V(X_2, \{y_t\}_{\substack{t=1 \\ t \ne i}}^N, F_{N-1}) \middle| X_1 = x \right] \right\} \\ &= \max_{1 \le i \le N} \left\{ xy_i E[Z_1] + E\left[v(X_2, \{E[Z_2|Z_1]y_t\}_{\substack{t=1 \\ t \ne i}}^N) \middle| X_1 = x \right] \right\}. \end{aligned}$$

Using $E[Z_2|Z_1] \ge 0$ and the independence of the X and Z sequences, Lemma 1 implies that

$$\begin{aligned} V(x, \{y_t\}_{t=1}^N, F_N) &= \max_{1 \le i \le N} \left\{ xy_i E[Z_1] + E\left[E[Z_2|Z_1] v(X_2, \{y_t\}_{\substack{t=1 \\ t \ne i}}^N) \middle| X_1 = x \right] \right\} \\ &= \max_{1 \le i \le n} \left\{ xy_i E[Z_1] + E[Z_2] E\left(v(X_2, \{y_t\}_{\substack{t=1 \\ t \ne i}}^N) \middle| X_1 = x \right) \right\} \end{aligned}.$$

Factoring out $E[Z] = E[Z_1] = E[Z_2]$, and using the optimality equation for the one sequence problem and finally Lemma 1, we conclude that

$$\begin{aligned} V(x, \{y_t\}_{t=1}^N, F_N) &= E[Z] \max_{1 \le i \le N} \left\{ xy_i + E\left(v(X_2, \{y_t\}_{\substack{t=1 \\ t \ne i}}^N) \middle| X_1 = x \right) \right\} \\ &= E[Z] v(x, \{y_t\}_{t=1}^N) \\ &= v(x, \{E[Z]y_t\}_{t=1}^N). \end{aligned}$$

■

The optimal policy of Theorem 7 may be significantly better than playing randomly. This may be already be seen in the game of (adult) Goofspiel with arbitrary N. The value is zero if both players play randomly or if both players play optimally. Suppose Player II plays randomly and Player I uses the optimal policy as given by Ross or Theorem 7 (match the card that is turned up). When card i is turned up, Player I wins i with probability $(i-1)/N$ and loses i with probability $(N-i)/N$. Player I's expected return is then $\sum_{i=1}^N i((i-1) - (N-i))/N = (N+1)(N-1)/6$.

References

Albright SC (1974) Optimal sequential assignments with random arrival times. Management Science 21:60-67.

Albright SC (1977) A Bayesian approach to a generalized house selling problem. Management Science 24:432-440.

Baston VJ, Bostock FA (1991) A generalized inspection game. Naval Research Logistics 38:171-182.

Brams SJ, Davis MD (1978) Optimal jury selection: A game-theoretic model for the exercise of peremptory challenges. Operations Research 26:966-991.

Derman C, Lieberman G, Ross SM (1972) A sequential assignment problem. Management Science 18:349-355.

Dresher M (1962) A sampling inspection problem in arms control agreements: a game-theoretic analysis. Memorandum RM-2972-ARPA, The RAND Corporation, Santa Monica, California.

Ferguson TS, Melolidakis C (1998) On the Inspection Game. Naval Research Logistics 45:327-334.

Gale D (1957) Information in games with finite resources. Annals of Mathematics Studies 39:141-145.

Garnaev AY (1994) A remark on the customs and smuggler game. Naval Research Logistics 41:287-293.

Luce RD, Raiffa H (1957) Games and Decisions, Reprinted (1989) by Dover Publications, Inc., New York.

Nakai T (1980) Generalized goofspiel. J. Operations Research Society Japan 23:156-170.

Nakai T (1982) A time sequential game related to the sequential assignment problem. J. Operations Research Society of Japan 25:129-139.

Righter R (1987) The stochastic sequential assignment problem with random deadlines. Prob. Engineering Information Sciences 1:189-202.

Righter R (1988) Job scheduling to minimize expected weighted flowtime on uniform processors. Systems Control Letters 10:211-216.

Righter R (1989) A resource allocation problem in a random environment. Operations Research 37:211-216.

Ross SM (1971) Goofspiel—The game of pure strategy. Journal of Applied Probability 8:621-625.

Sackson S (1994) Card Games Around the World, Dover Publ. Inc., New York.

Sakaguchi M (1977) A sequential allocation game for targets with varying values. J. Operations Research Society of Japan 20:182-193.

Sakaguchi M (1980) A time-sequential game for sums of random variables. J. Operations Research Society of Japan 23:81-90.

Thomas MU, Nisgav Y (1976) An infiltration game with time dependent payoff. Naval Research Logistics Quarterly 23:297-302.