

J o u r n a l
o f
P r o b a b i l i t y
E l e c t r o n i c

Vol. 13 (2008), Paper no. 9, pages 198–212.

Journal URL
<http://www.math.washington.edu/~ejpecp/>

Two-player Knock 'em Down

James Allen Fill*
Department of Applied Mathematics and Statistics
The Johns Hopkins University
Baltimore, Maryland 21218-2682
U.S.A.
<http://www.ams.jhu.edu/~fill/>

David B. Wilson
Microsoft Research
One Microsoft Way
Redmond, Washington 98052
U.S.A.
<http://dbwilson.com>

Abstract

We analyze the two-player game of Knock 'em Down, asymptotically as the number of tokens to be knocked down becomes large. Optimal play requires mixed strategies with deviations of order \sqrt{n} from the naïve law-of-large numbers allocation. Upon rescaling by \sqrt{n} and sending $n \rightarrow \infty$, we show that optimal play's random deviations always have bounded support and have marginal distributions that are absolutely continuous with respect to Lebesgue measure.

Key words: Knock 'em Down, game theory, Nash equilibrium.

AMS 2000 Subject Classification: Primary 91A60; Secondary: 91A05.

Submitted to EJP on December 14, 2006, final version accepted January 31, 2008.

*The research of J.A.F. was supported by NSF grants DMS-0104167 and DMS-0406104 and by The Johns Hopkins University's Acheson J. Duncan Fund for the Advancement of Research in Statistics. Research carried out in part while a visiting researcher at Microsoft Research.

1 Introduction

1.1 Knock 'em Down

In the game of Knock 'em Down, a player is given n tokens which (s)he arranges into k piles, or bins. After that, a k -sided die is thrown; if the outcome is side i (which occurs with probability $p_i > 0$), then a token is knocked down from the i^{th} pile. In the event that there are no tokens in the i^{th} bin, then no tokens get knocked down. The die is thrown repeatedly until all the tokens have been knocked down. [In the original version of this game as described in [Hun98] [BF99], $n = 12$ and two fair six-sided dice are thrown, with the bin chosen being given by (1 less, let us say, than) the sum of the two numbers showing. In that case, $k = 11$ and $p_i \equiv (6 - |i - 6|)/36$. The reader may wish to note that we have altered the spelling from “Knock 'm Down” to “Knock 'em Down”.]

We consider two versions of Knock 'em Down. In solitaire Knock 'em Down, which we analyze in a separate article, there is one player, and his goal is to minimize the expected number of iterations until all the tokens have been knocked down. Solitaire Knock 'em Down is also equivalent to a two-player zero-sum game, where the payoff to the winner is the expected number of extra die throws that the other player requires to knock down all his tokens, and the goal is to maximize the expected payoff. In competitive Knock 'em Down, which we analyze in this article, it is enough merely to win, and the amount by which a player wins is irrelevant. There are two players, who each arrange their tokens into bins without seeing what the other player is doing, and then the same die is used to knock down tokens from each of the players' bins. The winner is the player whose tokens get all knocked down first (the outcome may be a tie, which we resolve with a coin flip). With competitive Knock 'em Down, as we shall see, there is an interesting Nash equilibrium.

Competitive Knock 'em Down is quite easy to analyze when $k = 2$. The result is Theorem 1 in [BF00]; the authors show that the best strategy is to use the allocation $(m, n - m)$, where m is a median of the Binomial(n, p_1) distribution. It is instructive to consider competitive Knock 'em Down in the next simplest case, in which a fair three-sided die is used. The first player may guess that the best strategy is to place $n/3$ tokens into each of the three bins (assume for convenience that n is divisible by 3). But if the first player uses that strategy, then the second player can “undercut” by placing $(n/3) - 1$ tokens into each of the first two bins and $(n/3) + 2$ tokens into the third bin. Then, for large n , the probability that the second player's third bin empties out last is only slightly larger than $1/3$, while the probability that the first player's first or second bin finishes last is only slightly smaller than $2/3$, so that the second player wins about two-thirds of the time. It turns out that an optimal strategy is to allocate approximately $n/3$ tokens to each bin, but with certain random perturbations to the bin allocations. (See Figure 1.) In game-theoretic terminology, optimal play employs a mixed (non-pure) strategy.

To simplify the analyses of both games for general k and $\vec{p} := (p_1, \dots, p_k)$, we suppose that the games are run in continuous time, with the die being thrown at instants governed by a Poisson point process with rate 1. Throwing the die at random times rather than at deterministic times has absolutely no impact on the outcome of a game of competitive Knock 'em Down; likewise, in the case of solitaire Knock 'em Down the expected (total) clearance time (i.e., time to knock down all tokens) remains unchanged. Suppose in either game that a player places ξ_i tokens into bin i , and let T_i denote the time that it takes for bin i to be cleared. Since the game is



Figure 1: An optimal strategy for (an approximation to) the continuous limiting version of two-player Knock 'em Down described in Section 1.3 when $k = 3$ and $\vec{p} = (1/3, 1/3, 1/3)$. The strategy chooses a hexagon with probability proportional to its darkness, and then allocates the chips according to the hexagon's coordinates. Hexagons near the lower-left, lower-right, or top of the triangle allocate more chips to the first, second, or third piles respectively. The top corner has coordinates $(-0.28, -0.28, +0.56)/\sqrt{3}$.

run in continuous time, the T_i 's are mutually independent, and the variable $p_i T_i$ is the sum of ξ_i independent exponential random variables with unit mean, i.e., $p_i T_i$ is a Gamma random variable with shape parameter ξ_i . The clearance time is $T = \max_i T_i$.

1.2 Asymptotic scale: root- n deviations

Partial results concerning optimal play have been derived by Art Benjamin, Matt Fluet, and Mark Huber [BF99] [BF00] [BFH01] [Flu99] for both versions of Knock 'em Down, but in neither case has optimal play been characterized analytically. For both versions, we focus here on the asymptotics of optimal play when the die remains fixed [i.e., the vector \vec{p} is held constant] and the number of tokens n becomes large. While optimal play in competitive Knock 'em Down requires a mixed strategy, it is useful to consider first what happens when a player deterministically places ξ_i tokens into the i^{th} bin ($1 \leq i \leq k$). If ξ_i is large, then by the central limit theorem, T_i is approximately normal with mean ξ_i/p_i and variance ξ_i/p_i^2 . In particular, bin i is emptied by time $T_i = (\xi_i/p_i) + O_p(\sqrt{\xi_i/p_i})$. (We have used here the usual notation $X_n = O_p(y_n)$ to mean that $\inf_n \Pr[|X_n| \leq cy_n]$ tends to 1 as $c \rightarrow \infty$; equivalently, one says that the family of distributions $\mathcal{L}(X_n/y_n)$ is *tight*.)

If $\xi_i \equiv nq_i$, where the q_i 's are nonnegative and sum to unity, then the clearance time T satisfies $T = n \max_i (q_i/p_i) + O_p(\sqrt{n})$. Since $\max_i (q_i/p_i) \geq 1$, this suggests that the optimal allocation is $\xi_i \approx np_i$, and indeed for the solitaire game this was first shown in [BFH01]. We show in our companion paper on solitaire Knock 'em Down that the optimal choice is $\xi_i = np_i + O(\sqrt{n})$ for a judicious choice of the $O(\sqrt{n})$ perturbations. For competitive Knock 'em Down, our Theorem 1.2 below shows similarly that an optimal mixed strategy will never choose ξ_i differing from np_i by more than order \sqrt{n} .

Define the overplay of an allocation $\vec{\xi}$ relative to $n\vec{p}$ to be

$$\max_i \left(\frac{\xi_i}{p_i} - n \right).$$

In the following lemma, $\{\vec{\xi} \text{ beats } \vec{\eta}\}$ is the event that the corresponding clearance times satisfy the strict inequality $T_{\vec{\xi}} < T_{\vec{\eta}}$, or else $T_{\vec{\xi}} = T_{\vec{\eta}}$ and $\vec{\xi}$ wins the coin toss. The notation extends naturally to $\{\alpha \text{ beats } \beta\}$ for mixed strategies α, β ; see the start of Section 2 for careful analogous definitions in the continuous-game setting that we introduce in Section 1.3.

Lemma 1.1. *If the overplay of allocation $\vec{\eta}$ exceeds the overplay of allocation $\vec{\xi}$ by at least $w\sqrt{kn}/\min_i p_i$, then $\Pr[\vec{\eta} \text{ beats } \vec{\xi}] < 1/w^2$.*

Proof. Let j be the bin maximizing η_j/p_j . In order for allocation $\vec{\eta}$ to win, there must be some bin $\ell \neq j$ for which $\vec{\eta}$'s bin j clears out before $\vec{\xi}$'s bin ℓ . The clearance time T_0 of $\vec{\eta}$'s bin j has mean η_j/p_j and variance $\eta_j/p_j^2 \leq n/p_j^2$, while the clearance time T_ℓ of $\vec{\xi}$'s bin $\ell \neq j$ has mean $\xi_\ell/p_\ell \leq (\eta_j/p_j) - (w\sqrt{kn}/\min_i p_i)$ and variance $\leq \xi_\ell/p_\ell^2$. Since T_0 and T_ℓ are independent, by Chebyshev's inequality

$$\Pr[0 < T_\ell - T_0] \leq \frac{(n/p_j^2) + (\xi_\ell/p_\ell^2)}{(w\sqrt{kn}/\min_i p_i)^2} \leq \frac{1 + (\xi_\ell/n)}{w^2 k},$$

and since the support of $T_\ell - T_0$ is \mathbb{R} , the first above inequality is strict. Upon summing over $\ell \neq j$, we find that the probability that there is some bin $\ell \neq j$ for which $T_0 < T_\ell$ is less than $1/w^2$. \square

Theorem 1.2. *No optimal competitive Knock 'em Down (mixed) strategy ever overplays by more than $\left(3\sqrt{3 + \sqrt{5}}\sqrt{kn} + 1\right) / \min_i p_i$.*

Proof. Let α be any optimal strategy, and let $a(x)$ be the probability that strategy α picks an allocation with an overlay of x or more. We show successively, in each case using Lemma 1.1, that

$$a\left(\frac{w\sqrt{kn} + 1}{\min_i p_i}\right) < p_1, \quad a\left(\frac{2w\sqrt{kn} + 1}{\min_i p_i}\right) < p_2, \quad a\left(\frac{3w\sqrt{kn} + 1}{\min_i p_i}\right) = 0,$$

for $w = \sqrt{3 + \sqrt{5}}$, $p_1 = \frac{-1 + \sqrt{5}}{2}$, and $p_2 = \frac{3 - \sqrt{5}}{2}$. Let $p = 1 - (1/w^2) = (1 + \sqrt{5})/4$.

Suppose first that $a((w\sqrt{kn} + 1)/\min_i p_i) \geq p_1$. Consider pure (i.e., deterministic) strategy β which always plays $n\vec{p}$ (rounded to an integer vector summing to n). Because of the rounding, β may overplay $n\vec{p}$ slightly, but never more than $1/\min_i p_i$. By Lemma 1.1, strategy β beats α with probability $> p_1 \times p = 1/2$, contradicting the optimality of α .

Suppose next that $a((2w\sqrt{kn} + 1)/\min_i p_i) \geq p_2$. With positive probability ($> 1 - p_1$), strategy α overplays by no more than $(w\sqrt{kn} + 1)/\min_i p_i$, so we can define a strategy β which plays according to strategy α conditioned to overplay by no more than $(w\sqrt{kn} + 1)/\min_i p_i$. When β plays against α , by Lemma 1.1, β wins with probability $> (1 - p_1) \times \frac{1}{2} + p_2 \times p = 1/2$, which is again a contradiction to the optimality of α .

Suppose finally that $a((3w\sqrt{kn} + 1)/\min_i p_i) = \delta > 0$. Let β be the strategy which plays according to α conditioned to overplay by no more than $(3w\sqrt{kn} + 1)/\min_i p_i$. When β plays against α , by Lemma 1.1, strategy β wins with probability $> [(1 - \delta) \times \frac{1}{2}] + [\delta \times (1 - p_2) \times p] = \frac{1}{2}$, which is again a contradiction. \square

1.3 A continuous game

From Theorem 1.2 we see that any optimal strategy for competitive Knock 'em Down plays allocations deviating by $O_p(\sqrt{n})$ from the naïve law-of-large-numbers allocation $n\vec{p}$. For given n and allocation $\vec{\xi}$, it is thus natural to define numbers x_i so that $\xi_i = p_i n + x_i \sqrt{n}$. Then

$$\begin{aligned} T_i &\sim \text{Gamma}(\xi_i, 1)/p_i \\ &\sim \text{Normal}\left(n + \frac{x_i}{p_i}\sqrt{n}, \frac{n}{p_i} + \frac{x_i^2}{p_i^2}\sqrt{n}\right) \\ &\sim n + \text{Normal}\left(\frac{x_i}{p_i}, \frac{1}{p_i}\right)\sqrt{n}. \end{aligned}$$

(By \sim we mean that the total variation distance between the two distributions tends to 0 as $n \rightarrow \infty$ with p_i fixed and x_i bounded.) The player chooses the x_i 's so that

$$\sum_{i=1}^k x_i = 0$$

and $p_i n + x_i \sqrt{n}$ is an integer, and the player's tokens are all exhausted at time approximately

$$n + \sqrt{n} \max_i \left(\frac{x_i}{p_i} + \frac{Z_i}{\sqrt{p_i}} \right),$$

where the Z_i 's are independent standard normal random variables.

Thus the large- n asymptotics of either version of Knock 'em Down is effectively a continuous game (called *solitaire/competitive continuous Knock 'em Down*), where the first player chooses real numbers x_i satisfying $\sum_i x_i = 0$, and his clearance time is

$$T_{\vec{x}} = \max_i \left(\frac{x_i}{p_i} + \frac{Z_i}{\sqrt{p_i}} \right).$$

The second player similarly chooses numbers y_i summing to 0, with clearance time $T_{\vec{y}}$ defined using the same Z_i 's. The sequel provides various rigorous connections between n -token Knock 'em Down and our continuous game.

We present here two indications that even qualitative analysis of the continuous game is not entirely trivial. First, the naïve pure strategy $\vec{x} = \vec{0}$ has no optimal response. Indeed, the responding player can undercut by playing $(-\varepsilon, -\varepsilon, \dots, +(k-1)\varepsilon)$, and letting $\varepsilon \downarrow 0$ provides strategies which give the respondent asymptotically the optimal probability $(k-1)/k$ of winning, but no strategy achieves this probability. Second, it is not immediately clear that our two-player continuous game has a value (in the game-theoretic sense). There are standard tools for proving that a continuous game has a value, such as results of Ky Fan [Fan53], but our payoff function is rather severely discontinuous at certain points and the tools require a semicontinuous payoff function. Continuous games without values do exist [SW57], but it turns out that our game does indeed have a value. One can prove this by a suitable comparison of our game with a “ties go to player 1” modification of the game having upper semicontinuous payoff function, but we will give a somewhat more direct proof whose basic idea is simply to pass to the limit from n -token optimal strategies.

1.4 Guide to later sections

With *solitaire continuous Knock 'em Down*, the optimal strategy is deterministic, and we are able to characterize it for general p_i 's [FW06]. With *competitive continuous Knock 'em Down*, optimal play is random (i.e., mixed) with a rather complicated distribution (see Figure 1). Even in the simplest nontrivial instance, where $k = 3$ and $\vec{p} = (1/3, 1/3, 1/3)$, we are unable to calculate an optimal strategy. However, for general $k \geq 3$ and \vec{p} , we are able to derive some basic results about optimal play. In Section 2 we show, for example, that any optimal strategy for the continuous game has absolutely continuous marginals. In Section 3 we prove that the continuous game has an optimal strategy. In Section 4 we show that a good strategy for the continuous game can be converted to good strategies for the n -token games by rounding; in particular, an optimal continuous strategy can be converted to asymptotically optimal n -token strategies. Finally, in Section 5 we list some open problems arising from our work.

2 Properties of optimal play

The main result of this section (Theorem 2.5) asserts that the marginals of any optimal strategy for the continuous game described in Section 1.3 are absolutely continuous. This will follow from an estimate which holds (see Lemma 2.4) equally well for the discrete game.

We begin by establishing some terminology for the continuous game; analogous terminology will be employed for the discrete game. Recall that

$$T_{\vec{x}} = \max_i \left(\frac{x_i}{p_i} + \frac{Z_i}{\sqrt{p_i}} \right)$$

is the clearance time for a player using allocation $\vec{x} = (x_1, \dots, x_k)$ with $x_1 + \dots + x_k = 0$. Here Z_1, \dots, Z_k are independent standard normal random variables. Recall that if player 1 (say) uses allocation \vec{x} and player 2 uses allocation \vec{y} , then player 1 wins if and only if $T_{\vec{x}} < T_{\vec{y}}$; for short, we simply say “ \vec{x} beats \vec{y} ”, in which case player 2 pays 1 unit (utile) to player 1.

The usual game-theoretic analysis takes the payoff $K(\vec{x}, \vec{y})$ to be the expected amount won by player 1, namely $\Pr[\vec{x} \text{ beats } \vec{y}] - \Pr[\vec{y} \text{ beats } \vec{x}]$; for convenience we instead define $K(\vec{x}, \vec{y})$ to be half this difference.

For any two mixed strategies α and β , $\Pr[\alpha \text{ beats } \beta] = \iint \Pr[\vec{x} \text{ beats } \vec{y}] \alpha(d\vec{x}) \beta(d\vec{y})$ and we extend the definition of K to

$$\begin{aligned} K(\alpha, \beta) &:= \iint K(\vec{x}, \vec{y}) \alpha(d\vec{x}) \beta(d\vec{y}) = \frac{1}{2} (\Pr[\alpha \text{ beats } \beta] - \Pr[\beta \text{ beats } \alpha]) \\ &= \frac{1}{2} \iint (\Pr[\vec{x} \text{ beats } \vec{y}] - \Pr[\vec{y} \text{ beats } \vec{x}]) \alpha(d\vec{x}) \beta(d\vec{y}). \end{aligned} \quad (2.1)$$

We say that α *dominates* β if $K(\alpha, \beta) > 0$. A strategy α is *optimal* if no mixed strategy (equivalently, no pure strategy) dominates it.

Notice that payoffs are unaffected if we flip a fair coin to decide ties. That is, in the above definitions we may (and do) without effect redefine $\Pr[\vec{x} \text{ beats } \vec{y}]$ as $\Pr[T_{\vec{x}} \leq T_{\vec{y}}]$. Here we have introduced the convenient notation $\Pr[A \leq B]$ to mean $\Pr[A < B] + \frac{1}{2} \Pr[A = B]$; similarly, $\Pr[A \leq B \leq C]$ is shorthand for $\Pr[A < B < C] + \frac{1}{2} \Pr[A = B < C] + \frac{1}{2} \Pr[A < B = C] + \frac{1}{6} \Pr[A = B = C]$. (In our use of “ $\Pr[A \leq B \leq C]$ ” below, the third term always vanishes because $\Pr[A < C] = 1$.)

It will be convenient to view the contest between allocations \vec{x} and \vec{y} as follows. Let $m_i := \max(x_i, y_i)$ for $1 \leq i \leq k$. For convenience we will refer to $\vec{m} = (m_1, \dots, m_k)$ as an allocation, even though we are now working on the \sqrt{n} -scale of deviations and the sum $m_1 + \dots + m_k$ may exceed 0. Correspondingly, on the scale of tokens, define $M_i := p_i n + m_i \sqrt{n}$. Let $I \equiv I_{\vec{m}}$ be the bin that is cleared last when the bin sizes are M_1, \dots, M_k . [In the continuous game, I is defined correspondingly as $\arg \max_i (\frac{m_i}{p_i} + \frac{Z_i}{\sqrt{p_i}})$, where the Z_i 's are independent standard normals.] If $x_I > y_I$, then \vec{y} wins; if $x_I < y_I$, then \vec{x} wins; and if $x_I = y_I$, then the game is resolved by a coin flip. (Thus, overall, $\Pr[\vec{x} \text{ wins}] = \Pr[x_I \leq y_I]$.) To compare various strategies, we will couple the $I_{\vec{m}}$'s for various values of \vec{m} , taking the viewpoint that the same random sequence of die tosses will be made for a given game regardless of the allocations \vec{x} and \vec{y} that the two players use (in the continuous game, that the same vector of Z_i 's will be used). Then increasing m_i while leaving the other m_j 's ($j \neq i$) fixed may change I from a value different from i to i , but otherwise I will not change.

Lemma 2.1. *If \vec{m} is incremented by $1/\sqrt{n}$ in position i , then*

$$\Pr[I \text{ changes}] \leq (1 + o(1))/\sqrt{2\pi p_i n} = O(1/\sqrt{n}).$$

Proof. Let $M := M_i$, so that M is incremented by 1 to produce from \vec{m} a new allocation \vec{m}' . Ignore bin i for the moment, and let $T := \max_{j \neq i} T_j$ be the time at which all remaining bins are cleared for allocation \vec{m} (or \vec{m}'). In order for $I_{\vec{m}} \neq I_{\vec{m}'}$, it must be that bin i was selected exactly M times up through time T . Conditional on T , the number of times bin i was selected is Poisson distributed with parameter $\lambda := Tp_i$, and the probability of exactly M selections is

$$\frac{e^{-\lambda} \lambda^M}{M!} \leq \frac{e^{-M} M^M}{M!} \leq \frac{1}{\sqrt{2\pi M}},$$

Thus (unconditionally) $\Pr[I \text{ changes}] \leq 1/\sqrt{2\pi M}$. We use this bound when n is large and (say) $M \geq p_i n - n^{2/3}$. When M is smaller, we instead use $\Pr[I \text{ changes}] \leq \Pr[I_{\vec{m}'} = i]$, which is $o(1/\sqrt{n})$ (indeed, is exponentially small in n). \square

Lemma 2.2. *For fixed \vec{p} , the distribution of $I_{\vec{m}}$ is “locally flat” as a function of \vec{m} : for bins h , i , and j (not necessarily distinct), with e_h and e_i denoting unit vectors,*

$$\Pr[I_{\vec{m}} = j] + \Pr[I_{\vec{m} + \frac{e_h}{\sqrt{n}} + \frac{e_i}{\sqrt{n}}} = j] = \Pr[I_{\vec{m} + \frac{e_h}{\sqrt{n}}} = j] + \Pr[I_{\vec{m} + \frac{e_i}{\sqrt{n}}} = j] + O(1/n).$$

Proof. Let $F_t(\vec{M})$ denote the probability that the game with allocations given by \vec{M} has clearance time $\leq t$. If $Y_{\ell,t}$ denotes the number of times that bin ℓ has been selected through time t , then the Y_{ℓ} 's are independent Poisson processes with respective intensity parameters p_{ℓ} . We have

$$F_t(\vec{M}) = \Pr[\bigwedge_{\ell=1}^k Y_{\ell,t} \geq M_{\ell}] = \prod_{\ell=1}^k \Pr[\text{Poisson}(p_{\ell}t) \geq M_{\ell}].$$

The probability that the clearance time falls in $(t, t + dt)$ with bin j the last bin cleared is

$$[F_t(\vec{M} - e_j) - F_t(\vec{M})] p_j dt = \Delta_j F_t(\vec{M} - e_j) p_j dt,$$

where e_j is the j^{th} unit vector and Δ_j denotes the difference operator

$$\Delta_j f(\vec{x}) = f(\vec{x}) - f(\vec{x} + e_j).$$

Letting $\vec{M}' = \vec{M} - e_j$ we therefore have

$$\Pr[I_{\vec{M}} = j] = p_j \int_0^{\infty} \Delta_j F_t(\vec{M}') dt,$$

where here we are viewing $I = I_{\vec{M}}$ as a function of $\vec{M} = \vec{p}n + \vec{m}\sqrt{n}$. In this notation, the desired bound we will establish is

$$\Delta_h \Delta_i \Pr[I_{\vec{M}} = j] = p_j \int_0^{\infty} \Delta_h \Delta_i \Delta_j F_t(\vec{M}') dt = O(1/n). \quad (2.2)$$

Since we keep p_j fixed as $n \rightarrow \infty$, we see from (2.2) that we may treat h, i , and j symmetrically. Thus there are three cases to consider: (i) h, i , and j all distinct, (ii) $h = i \neq j$, (iii) $h = i = j$. Let us consider first the case (i) of distinct h, i, j :

$$\begin{aligned} \Delta_h \Delta_i \Pr[I_{\vec{M}} = j] &= \int_0^\infty \Pr[Y_{h,t} = M_h] \Pr[Y_{i,t} = M_i] \\ &\quad \times \Pr[Y_{j,t} = M_j - 1] p_j \left(\prod_{\ell \notin \{h,i,j\}} \Pr[Y_{\ell,t} \geq M_\ell] \right) dt, \\ |\Delta_h \Delta_i \Pr[I_{\vec{M}} = j]| &\leq \int_0^\infty \frac{1}{\sqrt{2\pi M_h}} \frac{1}{\sqrt{2\pi M_i}} \Pr[Y_{j,t} = M_j - 1] p_j dt \\ &= \frac{1}{\sqrt{2\pi M_h}} \frac{1}{\sqrt{2\pi M_i}}. \end{aligned}$$

If both M_h and M_i are of order n (as when good strategies are used), then we obtain the desired bound of $|\Delta_h \Delta_i \Pr[I_{\vec{M}} = j]| = O(1/n)$. Otherwise, if (say) $M_i = o(n)$, we may instead use the bound

$$|\Delta_i \Pr[I_{\vec{M}} = j]| \leq \Pr[I_{\vec{m}+e_i} \neq I_{\vec{m}}] \leq \Pr[I_{\vec{M}+e_i} = i],$$

which is of much smaller order than $1/n$. Thus $|\Delta_h \Delta_i \Pr[I_{\vec{M}} = j]| = O(1/n)$ in any case.

Now suppose (ii) $h = i \neq j$. In this case

$$\begin{aligned} \Delta_i \Delta_i \Pr[I_{\vec{M}} = j] &= \int_0^\infty (\Pr[Y_{i,t} = M_i] - \Pr[Y_{i,t} = M_i + 1]) \\ &\quad \times \Pr[Y_{j,t} = M_j - 1] p_j \left(\prod_{\ell \notin \{i,j\}} \Pr[Y_{\ell,t} \geq M_\ell] \right) dt, \\ |\Delta_i \Delta_i \Pr[I_{\vec{M}} = j]| &\leq \int_0^\infty |\Pr[Y_{i,t} = M_i] - \Pr[Y_{i,t} = M_i + 1]| \\ &\quad \times \Pr[Y_{j,t} = M_j - 1] p_j dt \\ &\leq \sup_t |\Pr[Y_{i,t} = M_i] - \Pr[Y_{i,t} = M_i + 1]|. \end{aligned}$$

Let $\lambda := p_j t$ and $M := M_i + 1$. For fixed M , the expression

$$\left| e^{-\lambda} \frac{\lambda^{M-1}}{(M-1)!} - e^{-\lambda} \frac{\lambda^M}{M!} \right| = e^{-\lambda} \frac{\lambda^{M-1}}{(M-1)!} \left| 1 - \frac{\lambda}{M} \right|$$

is easily seen to be maximized when λ is one of the values $\lambda = M \mp \sqrt{M}$, and so

$$|\Delta_i \Delta_i \Pr[I_{\vec{M}} = j]| \leq \frac{1}{\sqrt{2\pi M_i}} \frac{1}{\sqrt{M_i + 1}}.$$

As in case (i), we conclude that $|\Delta_i \Delta_i \Pr[I_{\vec{M}} = j]| = O(1/n)$.

In case (iii) $h = i = j$, the desired bound follows from the equality

$$\sum_j \Delta_i \Delta_i \Pr[I_{\vec{M}} = j] = \Delta_i \Delta_i 1 = 0$$

and the bounds for case (ii) $h = i \neq j$. □

Define the δ -undercut of \vec{x} to be the mixed strategy that increments \vec{x} by $(k-1)\delta$ in a uniformly random coordinate, and decrements \vec{x} by δ in the remaining coordinates. (In the discrete setting, δ will be a multiple of $1/\sqrt{n}$.) The preceding lemma implies

Corollary 2.3. *Uniformly in allocations \vec{x} and \vec{y} such that $|x_i - y_i| > (k-1)\delta$ for each coordinate i ,*

$$|\Pr[\delta\text{-undercut of } \vec{x} \text{ beats } \vec{y}] - \Pr[\vec{x} \text{ beats } \vec{y}]| = O(\delta^2).$$

Proof. We prove this in the discrete setting; the result for the continuous setting follows by taking limits. We assume without loss of generality that the bins are numbered so that $x_i > y_i$ for $1 \leq i \leq \ell$ and $x_i < y_i$ for $\ell + 1 \leq i \leq k$. Recall that $\vec{m} = \max(\vec{x}, \vec{y})$. Let \vec{x}' denote the random δ -undercut of \vec{x} , and $\vec{m}' = \max(\vec{x}', \vec{y})$. We have

$$\Pr[\vec{x} \text{ beats } \vec{y}] = \sum_{j=\ell+1}^k \Pr[I_{\vec{m}} = j],$$

and because \vec{x} and \vec{y} differ by more than $(k-1)\delta$ in each coordinate,

$$\Pr[\vec{x}' \text{ beats } \vec{y}] = \sum_{j=\ell+1}^k \Pr[I_{\vec{m}'} = j].$$

For each coordinate i , $\mathbb{E}[m'_i] = m_i$, and $|m'_i - m_i| = O(\delta)$, so the corollary follows from the local-flatness property we proved in Lemma 2.2. \square

Lemma 2.4. *For any fixed \vec{p} with $k \geq 3$ bins, for any coordinate i , for any $\delta \geq 1/\sqrt{n}$, and for any optimal strategy α , if \vec{x} and \vec{y} are independent draws from α , then $\Pr[|x_i - y_i| \leq \delta] = O(\delta)$, where the constant implicit in $O(\delta)$ depends only upon \vec{p} .*

The same proof works for both the discrete and continuous versions of the game.

Proof. We construct a strategy β which attempts to beat strategy α by undercutting it, to wit: β picks an allocation \vec{x} from α , but, rather than playing \vec{x} , strategy β instead plays the δ -undercut of \vec{x} . By analyzing how β fares against α we will be able to bound $\Pr[|x_i - y_i| \leq \delta]$.

When β and α are pitted against each other, we will take the viewpoint that \vec{x} and \vec{y} are independently drawn from α and a fair coin is flipped; if the coin lands heads, then α plays \vec{x} and β plays the δ -undercut of \vec{y} , while if the coin lands tails, then α plays \vec{y} and β plays the δ -undercut of \vec{x} . Recall that I is the last bin to be cleared when the allocations are \vec{x} and \vec{y} . Without the undercutting, bin I is owned by β with probability $1/2$. (We say that a player “owns” the last bin to be cleared if he was the player who placed more chips in that bin, or lost the coin toss in the event of tie.) Let I' be the bin last to be cleared *with* the undercutting.

Letting E be the event that $|x_i - y_i| \leq (k-1)\delta$ for some coordinate i and E^c its complement, we may express

$$\begin{aligned} \Pr[\alpha \text{ beats } \beta] &= \Pr[\{\alpha \text{ beats } \beta\} \cap E] + \Pr[\alpha \text{ beats } \beta \mid E^c] \Pr[E^c] \\ &\leq \Pr[\{\alpha \text{ beats } \beta\} \cap E] + \left(\frac{1}{2} + O(\delta^2)\right) \Pr[E^c] \end{aligned}$$

by Corollary 2.3. The optimality of α implies

$$0 \leq \Pr[\{\alpha \text{ beats } \beta\} \cap E] - \frac{1}{2} \Pr[E] + O(\delta^2). \quad (2.3)$$

Now, conditioning on \vec{x} and \vec{y} ,

$$\begin{aligned} \Pr[\alpha \text{ beats } \beta \mid \vec{x}, \vec{y}] &= \Pr[\beta \text{ owns } I' \neq I \mid \vec{x}, \vec{y}] + \Pr[\beta \text{ owns } I' = I \mid \vec{x}, \vec{y}] \\ &\leq \Pr[I' \neq I \mid \vec{x}, \vec{y}] + \Pr[\beta \text{ owns } I \mid \vec{x}, \vec{y}]. \end{aligned}$$

By Lemma 2.1, $\Pr[I' \neq I \mid \vec{x}, \vec{y}] = O(\delta)$. To compute $\Pr[\beta \text{ owns } I \mid \vec{x}, \vec{y}]$, we condition on I . For example, conditionally given the event $\delta < |x_I - y_I| < (k-1)\delta$, the player using β owns I with probability $1/2$ if the ‘‘overcut bin’’ [the bin with allocation incremented by $(k-1)\delta$] chosen is *not* bin I and with probability 1 if it is. Thus, if $\delta < |x_I - y_I| < (k-1)\delta$, then

$$\Pr[\beta \text{ owns } I \mid \vec{x}, \vec{y}, I] = \frac{1}{2} \left(1 - \frac{1}{k}\right) + \frac{1}{k} = \frac{1}{2} + \frac{1}{2k}.$$

The other entries in the following formula are computed similarly:

$$\Pr[\beta \text{ owns } I \mid \vec{x}, \vec{y}, I] = \begin{cases} \frac{1}{2} & \text{if } |x_I - y_I| > (k-1)\delta \\ \frac{1}{2} + \frac{1}{4k} & \text{if } |x_I - y_I| = (k-1)\delta \\ \frac{1}{2} + \frac{1}{2k} & \text{if } \delta < |x_I - y_I| < (k-1)\delta \\ \frac{1}{4} + \frac{3}{4k} & \text{if } |x_I - y_I| = \delta \\ \frac{1}{k} & \text{if } |x_I - y_I| < \delta. \end{cases}$$

Thus, conditional on \vec{x} and \vec{y} but not I ,

$$\begin{aligned} \Pr[\alpha \text{ beats } \beta \mid \vec{x}, \vec{y}] - \frac{1}{2} &= O(\delta) \\ &+ \frac{1}{2k} \Pr[\delta \leq |x_I - y_I| \leq (k-1)\delta \mid \vec{x}, \vec{y}] - \frac{k-2}{2k} \Pr[|x_I - y_I| \leq \delta \mid \vec{x}, \vec{y}], \end{aligned}$$

and so unconditionally

$$\begin{aligned} \Pr[\{\alpha \text{ beats } \beta\} \cap E] - \frac{1}{2} \Pr[E] &= O(\delta) \Pr[E] \\ &+ \frac{1}{2k} \Pr[\delta \leq |x_I - y_I| \leq (k-1)\delta] - \frac{k-2}{2k} \Pr[|x_I - y_I| \leq \delta]. \end{aligned}$$

Substituting this into (2.3), and then rearranging,

$$\begin{aligned} \frac{k-2}{2k} \Pr[|x_I - y_I| \leq \delta] &\leq O(\delta) \Pr[E] + O(\delta^2) + \frac{1}{2k} \Pr[\delta \leq |x_I - y_I| \leq (k-1)\delta], \\ (k-1) \Pr[|x_I - y_I| \leq \delta] &\leq O(\delta) \Pr[E] + O(\delta^2) + \Pr[|x_I - y_I| \leq (k-1)\delta]. \end{aligned}$$

Recall from Theorem 1.2 that we have $O(\sqrt{n})$ bounds on the overlay or underplay of the optimal strategy α . It follows that there is a positive constant q (depending on \vec{p}) such that $\Pr[I = i \mid \vec{x}, \vec{y}] \geq q$ for any coordinate i and plays \vec{x} and \vec{y} that α might make. Recalling also that E is the event that $|x_i - y_i| \leq (k-1)\delta$ for some coordinate i , we see that

$$\Pr[E] \leq \frac{1}{q} \Pr[|x_I - y_I| \leq (k-1)\delta].$$

Letting $r := k - 1$, we have then that, for some c and all $\delta \geq 1/\sqrt{n}$,

$$\Pr[|x_I - y_I| \leq \delta] \leq c\delta^2 + \frac{1 + c\delta}{r} \Pr[|x_I - y_I| \leq r\delta].$$

When $r = 1$ (i.e., $k = 2$) this inequality is uninformative, but otherwise we may iterate it to show, for $j \geq 1$ and $\delta \geq 1/\sqrt{n}$,

$$\frac{\Pr[|x_I - y_I| \leq \delta]}{(1 + c\delta) \times \cdots \times (1 + cr^{j-1}\delta)} \leq c[\delta^2 + \cdots + r^{j-1}\delta^2] + \frac{1}{r^j} \Pr[|x_I - y_I| \leq r^j\delta].$$

We take $j = \lceil \log(1/\delta)/\log r \rceil$ so that $(1 + c\delta) \times \cdots \times (1 + cr^{j-1}\delta) = O(1)$ and both terms on the right-hand side are $O(\delta)$. Thus

$$\Pr[|x_I - y_I| \leq \delta] = O(\delta).$$

Recalling again that $\Pr[|x_I - y_I| \leq \delta] \asymp \Pr[\cup_i \{|x_i - y_i| \leq \delta\}]$ yields the lemma. \square

Theorem 2.5. *For competitive continuous Knock 'em Down with $k \geq 3$ bins, each marginal distribution of any optimal strategy is absolutely continuous with respect to Lebesgue measure, and has a square-integrable density function.*

Proof. If X and Y are i.i.d. draws from the marginal distribution, then, by Lemma 2.4, $\Pr[|X - Y| \leq \delta] = O(\delta)$. From this it is straightforward (see, e.g., [Mat95, Theorem 2.12]) to check that the distribution is absolutely continuous with respect to Lebesgue measure, and that the density function is square-integrable. \square

Note that the support of *any* measure with nonzero absolutely continuous part has positive Lebesgue measure, and that *any* subset of the line with positive Lebesgue measure has positive 1-dimensional Hausdorff measure. Thus we have

Corollary 2.6. *For $k \geq 3$ bins, the support of any optimal strategy for continuous Knock 'em Down has Hausdorff dimension at least 1.*

It is natural to guess that the true Hausdorff dimension is $k - 1$, and indeed that optimal strategies are absolutely continuous with respect to $(k - 1)$ -dimensional Lebesgue measure.

3 Existence of an optimal continuous strategy

While every game in which each player has a finite number of options will have a value (which of course is 0 for n -token competitive Knock 'em Down), there are continuous games without a value [SW57]. Thus Theorem 3.1 below has nontrivial content.

Recall the definition of K at (2.1). For the n -token game, we regard a mixed strategy α_n as a probability measure on k -tuples $\vec{x} = (x_1, \dots, x_k)$ with vanishing sum, as described in Section 1, and we define the payoff function K_n on this \vec{x} -scale. We say that a sequence (α_n) of strategies for the n -token competitive Knock 'em Down games is *asymptotically optimal* if

$$\min_{\vec{y} \in A_n} K_n(\alpha_n, \vec{y}) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

Here the min is taken over the finite (but growing, as $n \rightarrow \infty$) set A_n of (normalized) actions (allocations) available for n -token Knock 'em Down.

Theorem 3.1. *The continuous game has value 0, and there is at least one optimal strategy. Indeed, any subsequential weak limit of any asymptotically optimal sequence (α_n) of strategies for n -token competitive Knock 'em Down is an optimal strategy for the continuous game.*

Key ingredients to the proof of this theorem are Theorem 1.2 and Lemma 2.4. The converse to this theorem is proved in Corollary 4.3.

Proof of Theorem 3.1. By Theorem 1.2 and Lemma 1.1, any sequence (α_n) of asymptotically optimal strategies (probability measures) is tight. Therefore there is a subsequential weak limit α which is a probability measure (see e.g. [Str93, Theorem 3.1.9]). It is easy to check that α is concentrated on k -tuples with vanishing sum, so that α is a mixed strategy for the continuous game, and (using Lemma 2.4) that α has atomless marginals. We claim that α is optimal, and then we see immediately that the continuous game has value 0.

The proof that α is optimal is rather routine. To avoid double subscripts, we henceforth innocuously assume that the full sequence (α_n) converges weakly to α . Fix any pure strategy \vec{y} for the continuous game. Let \vec{y}_n denote a rounding of \vec{y} to a pure strategy for K_n ; the details of the rounding procedure are irrelevant for our purposes here, as long as $\vec{y}_n \rightarrow \vec{y}$. Using the facts that α has atomless marginals, $\alpha_n \xrightarrow{w} \alpha$, K is continuous away from pairs (\vec{x}, \vec{y}) for which $x_i = y_i$ for some i , and (for any $\varepsilon > 0$) $|K_n(\vec{x}, \vec{y}_n) - K(\vec{x}, \vec{y})| \rightarrow 0$ uniformly for \vec{x} satisfying $|x_i - y_i| > \varepsilon$ for $i = 1, \dots, k$, it follows (we omit the details) that $K_n(\alpha_n, \vec{y}_n) \rightarrow K(\alpha, \vec{y})$. But by the optimality of α_n , $K_n(\alpha_n, \vec{y}_n) \geq 0$ for every n , so $K(\alpha, \vec{y}) \geq 0$ and α is optimal. \square

Remark 3.2. (a) If it happens to be true that there exists a *unique* optimal strategy α_0 for K , then Theorem 3.1 implies that

$$\alpha_n \xrightarrow{w} \alpha_0 \text{ as } n \rightarrow \infty. \quad (3.1)$$

(b) If $p_1 = \dots = p_k$, it is possible to choose the strategies α_n to be symmetric. If so, and if it happens that there exists a unique symmetric optimal strategy α_0 for K , then (3.1) holds.

4 Asymptotically optimal play of Knock 'em Down

The main result of this section (see Corollary 4.3) is that when an optimal strategy α_0 for the continuous game K is “rounded” to produce allocations for n -token Knock 'em Down, the result is an asymptotically optimal strategy. The drawback here is that we do not know how to construct such an α_0 , but we might at least hope to find a not unreasonably suboptimal $\tilde{\alpha}$ for K , such as the one discussed in Remark 4.2. We are thus motivated to show more generally (see Theorem 4.1) that “rounding” of any strategy $\tilde{\alpha}$ with atomless marginals and bounded support gives a strategy for Knock 'em Down whose worst-case payoff (i.e., payoff against an opponent playing the best possible response) is asymptotically at least as large as the worst-case payoff in game K from use of $\tilde{\alpha}$.

Theorem 4.1. *If $\alpha_n \xrightarrow{w} \tilde{\alpha}$, where $\tilde{\alpha}$ has atomless marginals and bounded support, then*

$$\liminf_{n \rightarrow \infty} \min_{\vec{y} \in A_n} K_n(\alpha_n, \vec{y}) \geq \inf_{\vec{y}} K(\tilde{\alpha}, \vec{y}). \quad (4.1)$$

Proof. Let us define

$$\kappa_n := \min_{\vec{y} \in A_n} K_n(\alpha_n, \vec{y}), \quad (4.2)$$

and let $\vec{y}_n \in A_n$ be a strategy achieving this minimum. Since the sequence (α_n) converges weakly, it is tight, so by Lemma 1.1, \vec{y}_n remains bounded. Let $n_\ell \uparrow \infty$ be any sequence for which $\lim_{\ell \rightarrow \infty} \kappa_{n_\ell} = \liminf_{n \rightarrow \infty} \kappa_n$. By compactness, we know that there is a subsequence $\tilde{n}_\ell \uparrow \infty$ and a continuous-game allocation \vec{y} such that $\vec{y}_{\tilde{n}_\ell} \rightarrow \vec{y}$.

The supremum of $|K_n(\vec{x}, \vec{y}) - K(\vec{x}, \vec{y})|$ over any bounded set of (\vec{x}, \vec{y}) tends to 0 as $n \rightarrow \infty$; in particular, ties do not cause a problem here. Since the sequence (α_n) is tight and the sequence \vec{y}_n is bounded, we have that κ_n differs by $o(1)$ from

$$\hat{\kappa}_n := K(\alpha_n, \vec{y}_n).$$

For $\varepsilon > 0$ define

$$F_\varepsilon := \{\vec{x} : |x_i - y_i| \leq \varepsilon \text{ for some } i\},$$

and let F_ε^c denote the complement of this set. We have that

$$K(\vec{x}, \vec{y}_{\tilde{n}_\ell}) \rightarrow K(\vec{x}, \vec{y}) \quad \text{as } \ell \rightarrow \infty$$

uniformly for $\vec{x} \in F_\varepsilon^c$. Thus, fixing $\varepsilon > 0$,

$$\hat{\kappa}_{\tilde{n}_\ell} \geq \int_{F_\varepsilon^c} K(\vec{x}, \vec{y}) \alpha_{\tilde{n}_\ell}(d\vec{x}) - \frac{1}{2} \alpha_{\tilde{n}_\ell}(F_\varepsilon) - o(1).$$

Now, by an argument (omitted here) very much like one used in the detailed proof of Theorem 3.1 one finds

$$\lim_{\ell \rightarrow \infty} \hat{\kappa}_{\tilde{n}_\ell} \geq K(\tilde{\alpha}, \vec{y}) - 2\tilde{\alpha}(F_\varepsilon).$$

Letting $\varepsilon \downarrow 0$ and using the continuity of the marginals of $\tilde{\alpha}$, we obtain

$$\liminf_{n \rightarrow \infty} \kappa_n = \liminf_{n \rightarrow \infty} \hat{\kappa}_n = \lim_{\ell \rightarrow \infty} \hat{\kappa}_{\tilde{n}_\ell} \geq K(\tilde{\alpha}, \vec{y}) \geq \inf_{\vec{y}} K(\tilde{\alpha}, \vec{y}). \quad \square$$

Remark 4.2. Even very simple-minded mixed strategies can offer substantial improvement over pure strategies. To illustrate this, we consider $k = 3$ and $\vec{p} = (1/3, 1/3, 1/3)$. Numerical experimentation suggested that the uniform distribution $\tilde{\alpha}$ over the simplex of all triples $\vec{x} = (x_1, x_2, x_3)$ summing to zero and satisfying $x_i \geq -1/6$ for all i might be a good strategy for the continuous game. Indeed, numerical explorations indicate that the best response to $\tilde{\alpha}$ is the pure strategy $(0, 0, 0)$ and thence that

$$\text{RHS}(4.1) \doteq -0.0101219$$

[a huge improvement on the worst-payoff value $-1/6 \doteq -0.1666667$ resulting from use of the naïve pure strategy $(0, 0, 0)$ in place of $\tilde{\alpha}$]. The strategy $\tilde{\alpha}$ was “rounded” to a strategy α_{180} for 180-chip Knock ’em Down by taking α_{180} to be uniform over allocations placing at least 58 chips in each bin. [Note $60 - (1/6)\sqrt{180} \doteq 57.8$.] Then, according to *Mathematica*, the value of κ_{180} at (4.2) is

$$\kappa_{180} \doteq -0.0165257,$$

improving on the worst possible payoff of $\doteq -0.0920653$ for the naïve allocation $(60, 60, 60)$.

Corollary 4.3. *If $\alpha_n \xrightarrow{w} \alpha_0$, where α_0 is optimal for the continuous game K , then α_n is asymptotically optimal for the n -token Knock ’em Down game K_n , in the sense that κ_n defined at (4.2) vanishes in the limit.*

5 Open problems

We have proved the existence of an optimal strategy for two-player continuous Knock 'em Down, but we do not have an explicit description of optimal play even when $k = 3$ and $p_1 = p_2 = p_3 = 1/3$. We know that the marginal distributions of optimal play are absolutely continuous with respect to Lebesgue measure. Consequently the set of pure strategies supporting optimal play will have dimension at least 1; perhaps the dimension is $k - 1$.

Acknowledgments

We are grateful to Yuval Peres and Elchanan Mossel for useful discussions, and the referee for helpful comments. We thank Ed Scheinerman, whose son Jonah played Knock 'em Down in his third-grade math class taught by Bonnie Nagel, for bringing the game to our attention.

References

- [BF99] Arthur T. Benjamin and Matthew T. Fluet. The best way to Knock 'm Down. *The UMAP Journal*, 20(1):11–20, 1999.
- [BF00] Arthur T. Benjamin and Matthew T. Fluet. What's best? *Amer. Math. Monthly*, 107(6):560–562, 2000. MR1767066
- [BFH01] Arthur T. Benjamin, Matthew T. Fluet, and Mark L. Huber. Optimal token allocations in Solitaire Knock 'm Down. *Electron. J. Combin.*, 8(2):Research Paper 2, 8 pp., 2001. In honor of Aviezri Fraenkel on the occasion of his 70th birthday. MR1853253
- [Fan53] Ky Fan. Minimax theorems. *Proc. Nat. Acad. Sci. U.S.A.*, 39:42–47, 1953. MR0055678
- [Flu99] Matthew T. Fluet. Searching for optimal strategies in Knock 'm Down, 1999. Senior thesis, Harvey Mudd College, Claremont, CA.
- [FW06] James A. Fill and David B. Wilson. Solitaire Knock 'em Down, 2006. Manuscript.
- [Hun98] Gordon Hunt. Knock 'm down. *Teaching Stat.*, 20(2):59–62, 1998.
- [Mat95] Pertti Mattila. *Geometry of Sets and Measures in Euclidean Spaces, fractals and rectifiability*, volume 44 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995. MR1333890
- [Str93] Daniel W. Stroock. *Probability Theory, an analytic view*. Cambridge University Press, Cambridge, 1993. MR1267569
- [SW57] Maurice Sion and Philip Wolfe. On a game without a value. In *Contributions to the theory of games, vol. 3*, Annals of Mathematics Studies, no. 39, pages 299–306. Princeton University Press, Princeton, N. J., 1957. MR0093742