How to build a probability-free casino

Adam Chalcraft
CCR–La Jolla
dachalc@ccrwest.org

Randall Dougherty
CCR–La Jolla
rdough@ccrwest.org

Chris Freiling
Cal State San Bernardino
cfreilin@csusb.edu

Jason Teutsch*
Ruprecht-Karls-Universität Heidelberg
teutsch@math.uni-heidelberg.de

January 13, 2012

Abstract

Casinos operate by generating sequences of outcomes which appear unpredictable, or *random*, to effective gamblers. We investigate relative notions of randomness for gamblers whose wagers are restricted to a finite set. Some sequences which appear unpredictable to gamblers using wager amounts in one set permit unbounded profits for gamblers using different wager values. In particular, we show that for non-empty finite sets A and B, every A-valued random is B-valued random if and only if there exists a $k \geq 0$ such that $B \subseteq A \cdot k$.

You enter a casino and there is an "idiot" betting on a sequence of coin flips with wagers of the following values: \$2, \$3, and \$7. He always bets on heads and cycles through these values:

$$\$2, \$3, \$7, \$2, \$3, \$7, \ldots$$

You get a "normal" stack of values: \$1, \$5, and \$10 with which to bet. The idiot becomes rich and moves to Beverly Hills, and you move into a cardboard box. What happened?

*Research supported by the Deutsche Forschungsgemeinschaft grant ME 1806/3-1.

Most casinos extract money from gamblers by exploiting the *law of large numbers* from probability theory which says that almost surely the number of heads approaches one half the total coin flips in the limit. The title of "probability-free" casino indicates that we examine an alternate principle under which a casino might operate profitably. In our paradigm, a naïve "idiot," as in the example above, seduces the gambler into believing he can duplicate the idiot's sure win. Our main result identifies a sequence in which betting with wager values from one set one can gain arbitrary wealth whereas gambling using another set of wager values results in guaranteed bankruptcy.

A *sequence* is a denumerable list of coin flip outcomes, with "h" for heads and "t" for tails. In the following discussion, $\{\mathtt{h},\mathtt{t}\}^*$ denotes all finite strings of coin flips, and $\{\mathtt{h},\mathtt{t}\}^\omega$ is the class of infinite sequences of coin flips. A computable function M from $\{\mathtt{h},\mathtt{t}\}^*$ to non-negative reals which satisfies the *fairness condition*

$$M(\sigma) = \frac{M(\sigma\mathtt{h}) + M(\sigma\mathtt{t})}{2}$$

for all $\sigma \in \{\mathtt{h},\mathtt{t}\}^*$ is called a *martingale* [2]. Informally, a martingale is a gambling strategy, and the M represents the gambler's capital at each position.

For any sequence $X \in \{\mathtt{h},\mathtt{t}\}^\omega$, $X \upharpoonright n$ denotes the first n coin flips of X, and $|\cdot|$ denotes the length of a string, size of a set, or absolute value of a real number. A martingale M *succeeds* on $X \in \{\mathtt{h},\mathtt{t}\}^\omega$ if M achieves arbitrarily large sums of money over X, that is, $\limsup_n M(X \upharpoonright n) = \infty$[1]. Otherwise X *defeats* M. For a martingale M and $\sigma \in \{\mathtt{h},\mathtt{t}\}^*$, $M(\sigma\mathtt{h}) - M(\sigma)$ is called the *wager at* σ. If the wager at σ is positive, then M *predicts* "heads" at σ; if the wager at σ is negative, then M *predicts* "tails" at σ.

Given a set V of reals, we say that a martingale is *V-valued* if for all σ the wager of M at σ belongs to V, unless M does not have enough capital in which case the wager at σ is 0.

Definition (Bienvenu, Stephan, Teutsch [1])**.** A sequence $X \in \{\mathtt{h},\mathtt{t}\}^\omega$ is *V-valued random* if no V-valued martingale succeeds on X.

[1] Equivalently $\lim_n M(X \upharpoonright n) = \infty$ since we only consider martingales whose nonzero wagers are bounded by some $\epsilon > 0$. If some martingale succeeded lim sup but not lim, then there is some least interval $[k\epsilon, (k+1)\epsilon)$, k a nonnegative integer, which the martingale's capital visits infinitely often, and after all but finitely many such visits the next positive wager of the martingale must be correct. Thus some other martingale can succeed lim by agreeing with the successful lim sup martingale at these positions of minimal capital.

For $A \subseteq \mathbb{R}$, let $A \cdot k$ denote the set $\{x \cdot k : x \in A\}$. The purpose of this article is to prove the following theorem which answers a question posed by Bienvenu, Stephan, and Teutsch [1].

Theorem. *Let A and B be non-empty finite sets of computable real numbers. Then every A-valued random is B-valued random if and only if there exists a $k \geq 0$ such that $B \subseteq A \cdot k$.*

Proof. Assume that $B \subseteq A \cdot k$ for some $k \geq 0$, and suppose some B-valued martingale S succeeds on a sequence X. Let M be a martingale which bets k times whatever S bets at each position (and starts with k times S's capital). Then M's capital equals k times S's capital everywhere, so M succeeds on X.

Conversely, suppose that there is no k such that $B \subseteq A \cdot k$. We shall exhibit a B-valued martingale S and a sequence X such that S succeeds on X but no A-valued martingale succeeds on X. The martingale S represents the capital of a *stooge* who works in collusion with the *casino* X. $M_0, M_1, M_2, \ldots$ is a list of the A-valued martingales, or *customers*, who bet on the casino's coin flips alongside the stooge. We shall build our casino so as to destroy all customers while allowing the stooge to win arbitrarily large amounts of money. Our casino sequence will be noncomputable since its construction depends on access to a list of all A-valued martingales.

The overall heuristic is that each customer must try to copy the bets of the casino's stooge, otherwise the casino can just outright destroy the customer while simultaneously helping the stooge to win. But by the hypothesis that $B \not\subseteq A \cdot k$, no customer can exactly copy the stooge's bets over an interval of positions containing all possible wagers for the stooge. In particular, any customer must overbet or underbet somewhere along this interval, and then the casino can exploit the customer's misstep.

Assume $B = \{b_1, b_2, \ldots, b_n\}$. We fix S's strategy now so as to be independent of the choice of X:

$$S(\sigma\mathtt{h}) = S(\sigma) + b_i \qquad \text{where } |\sigma| \equiv i \bmod n,$$

unless $S(\sigma) < |b_i|$ in which case $S(\sigma\mathtt{h}) = S(\sigma)$. In other words, S simply cycles through all possible wager values in B. The initial capital of S is chosen to be at least $\max_i |b_i|$.

It suffices to build X so as to meet for all e the following requirements:

R_e**:** Ruin M_e (or reach a stage where M_e stops betting).

S_e**:** The stooge's capital reaches $\$e$ at some position.

If all requirements are satisfied, then all M_e's are obliterated and S succeeds. Our method for satisfying these requirements is based on the following heuristics:

- M_e always tries to "copy" S's bets. Otherwise X can hurt M_e without harming S's capital.
- Once S is sufficiently richer than M_e, X just hurts both of them until M_e dies.

The order of priority for satisfying the requirements above is:

$$S_0 > R_0 > S_1 > R_1 > S_2 > R_2 > \cdots$$

Higher priority requirements may injure lower priority R_e requirements that have already been satisfied, but any particular requirement will only be injured finitely often.

Basic module for R_e. Our method to defeat a single customer is as follows. We build finite extensions whose limit is X. We shall assume that at the current position the customer and the stooge always both bet on heads (or both bet on tails), since the strategy for the casino is otherwise straightforward: help the stooge while hurting the customer.

Every $|B|$ positions, we compute the ratio k of the stooge's capital to the customer's capital. At those positions where k times the magnitude of the customer's wager exceeds the magnitude of the stooge's wager, we choose a value for X which hurts both the customer and stooge, and otherwise we help them both.

Intuitively, the casino hurts the customer if he bets riskier than the stooge; otherwise the casino helps the stooge get ahead. By assumption the customer cannot match any multiple of the stooge's $|B|$ wager values, hence $|B|$ consecutive positions will suffice to commit a slight increase in the ratio of stooge-to-customer capital.

We shall show that the sequence of k's is strictly increasing and tends to infinity, at least until the customer runs out of money. The latter condition must eventually occur because, due to the assumption that wagers are restricted to a finite set of values, for sufficiently large k the casino always hurts the customer.

We now give the full construction. Assume that X has been defined at all positions less than p. We describe how to make the next proper extension of X. Let e be the least index such that either:

1. S's capital has not reached \$$e$ at any position before p, or

2. M_e makes a non-zero wager at position p.

If e satisfies the first case, then we set X to agree with S's wager at position p, and we say that requirement S_e *acts* at this position. Otherwise M_e made a non-zero wager at position p, and we say that e is the index which *receives attention* for the next $|B|$ positions (unless some higher priority requirement interrupts). In this case, we try to satisfy R_e.

In order to ensure that the ratio of capitals between S and M_e looks no less favorable to S than the last time M_e made a non-zero wager, we introduce the following function[2]:

$$e\text{-}\mathtt{savings}(X \restriction p) = \max_{u<p}\big\{S(X \restriction u+1) : (\exists j < e)\, [u \text{ was the latest position} < p \text{ for which } M_j \text{ made a non-zero wager on } X \restriction u]\big\} \cup \{0\}.$$

For each index j with priority higher than e, we look at the position of the most recent bet of M_j and define e-`savings` to be the maximum capital of the stooge over all such positions. The casino will prevent the stooge's capital from ever dropping below e-`savings` so long as R_j does not force this to happen for some $j < e$. The idea is that no requirement with priority less than R_e can harm X's progress in defeating M_e. Now define the ratio:

$$k(p) = \frac{S(X \restriction p) - e\text{-}\mathtt{savings}(X \restriction p)}{M_e(X \restriction p)}.$$

For each u with $p \leq u < p + |B|$ where M_e makes a nonzero wager, define the predicate

$$P(u) \quad = \quad \left[k(p) > \frac{|S\text{'s wager at } X \restriction u|}{|M_e\text{'s wager at } X \restriction u|}\right].$$

$P(u)$ holds when M_e commits a greater fraction of capital (relative to position p) than S does.

We now enter a loop which lasts for up to $|B|$ positions. We define p according to the loop, and then successively $p+1, p+2, \ldots, p+|B|-1$ in the same way unless for some $j < e$ the martingale M_j makes a non-zero wager at one of these positions. In this case, the loop breaks and instead j

[2]For the first reading of this proof, consider the simplified problem where we diagonalize against only a *single* strategy M_0. In this case we can just take e-`savings` to be the constant zero function.

receives attention: we define a new function k and predicate P based on j and the current position and initiate another length $|B|$ loop analogously. Eventually one of these loops will complete $|B|$ repetitions, and at this point we will have defined our next extension of X. For the remainder of the algorithm description, we shall assume that the loop below lasts for the full $|B|$ repetitions.

Assume that X has been defined for all positions less than u. Now define $X(u)$ according to the first of the following cases which is satisfied:

$$X(u) = \begin{cases} S\text{'s prediction at } u & \text{if } M_e \text{ bets opposite from } S \text{ at } u \\ & \quad \text{or } M_e \text{ bets } 0, \\ \neg M_e\text{'s prediction at } u & \text{if } P(u) \text{ holds}, \\ M_e\text{'s prediction at } u & \text{if } P(u) \text{ fails}, \end{cases} \tag{1}$$

so that M_e gets hurt iff she wagers a greater percentage of capital than S (and matches S's prediction) or makes the opposite prediction of S. This completes the construction of X.

We now verify that X defeats every A-valued martingale. At every position u at which M_e receives attention and

$$k(p) > \frac{\max\{|b| : b \in B\}}{\min\{|a| : a \in A\}}, \tag{2}$$

M_e must lose money by (1). Thus it suffices to show that $k(p)$ exceeds this bound (2) for all sufficiently large p.

Let us examine what happens to k when M_e receives attention at position p and the loop lasts for the full $n = |B|$ repetitions. That is, we assume no higher priority requirement receives attention during these n steps. Let $s_1, s_2, \ldots, s_n$ be the dollars won (possibly negative) by S at positions $p, p+1, \ldots, p+n-1$, and let $m_1, m_2, \ldots, m_n$ be the dollars won by M_e at these positions. Let q be the next position at which M_e receives attention. For clarity in the calculation (3), let $s = s_1 + s_2 + \ldots + s_n$ and $m = m_1 + m_2 + \ldots + m_n$. If we further assume that p is sufficiently large that after position p no index less than e receives attention and no S_j with $j < e$ ever acts, then e-`savings` does not change between p and q and thus

$$\begin{aligned} k(q) \cdot [M_e(X \upharpoonright p) + m] &= [S(X \upharpoonright p) - e\text{-}\mathtt{savings}(X \upharpoonright p)] + s \\ &= k(p) \cdot M_e(X \upharpoonright p) + s \\ &= [k(p) \cdot M_e(X \upharpoonright p) + k(p) \cdot m] + [-k(p) \cdot m + s], \end{aligned}$$

and so

$$k(q) = k(p) + \frac{-k(p) \cdot m + s}{M_e(X \restriction p) + m}. \tag{3}$$

Let $c = |B| \cdot \max\{|a| : a \in A\}$. Then

$$M_e(X \restriction p) + m \leq M_e(X \restriction p) + c. \tag{4}$$

In particular, the denominator in (3) increases by at most a constant each time M_e receives attention.

Now let us inspect the numerator in (3). Define the non-negative constant

$$\epsilon = \inf_{k>0} \left[\max_i \min_j |k \cdot a_j - b_i| \right]$$

where the a_j's are the elements of A. Recall that $B \not\subseteq A \cdot k$ by assumption, so by continuity ϵ is nonzero. We claim:

- $k(p) \cdot m_i \leq s_i$ for all $i \leq n$, and
- $k(p) \cdot m_i \leq s_i - \epsilon$ for at least one such i.

We obtain from this claim the following bound for the numerator:

$$k(p) \cdot m = k(p) \cdot \sum_{i=1}^{n} m_i \leq -\epsilon + \sum_{i=1}^{n} s_i = s - \epsilon. \tag{5}$$

Combining our observations on the numerator and denominator (4),(5) with (3), we see that

$$k(q) \geq k(p) + \frac{\epsilon}{M_e(X \restriction p) + c}. \tag{6}$$

The first part of the claim follows by inspecting each of the cases in (1). If M_e and S predict opposite outcomes or M_e bets 0, the result is immediate because X hurts M_e while helping S at the same time. If $P(p+i)$ fails, then X helps both S and M_e at this position and in particular $k(p) \cdot m_i \leq s_i$. If on the other hand $P(p+i)$ is true, then X hurts both S and M_e at this position, and in particular $k(p) \cdot m_i \leq s_i$ because m_i and s_i are both negative.

It remains to establish the second part of the claim. Appealing to our original hypothesis, we find that B is not a subset of $A \cdot k(p)$. Since S cycles through all possible bets for B over the n positions $p, p+1, \ldots, p+n-1$, there exists some $r \leq n$ such that $k(p) \cdot m_r \leq s_r - \epsilon$.

Let $p_0, p_1, p_2, \ldots$ be the positions where M_e receives attention, with p_0 large enough that no lesser index than e ever requires attention again and

S_j has been satisfied for all $j < e$. These positions exist by induction on the following argument. As noted in (4), $M_e(X \upharpoonright p_{i+1}) \leq M_e(X \upharpoonright p_i) + c$. So by induction on (6) and divergence of the harmonic series we have:

$$k(p_r) \geq k(p_0) + \sum_{i=1}^{r} \frac{\epsilon}{M_e(X \upharpoonright p_0) + c \cdot i} \geq \frac{\epsilon}{\max\{M_e(X \upharpoonright p_0), c\}} \cdot \sum_{i=2}^{r+1} \frac{1}{i} \to \infty.$$

It follows that either:

- For all sufficiently large r, $k(p_r)$ satisfies the desired inequality (2),
- or else M_e stops betting at some stage.

Either way, M_e does not succeed on X.

Since each M_j ($j \leq e$) eventually stops betting, there is some stage at at which point S_e becomes satisfied (if it was not already) because the algorithm chooses an index to focus on according to case 1. Hence S succeeds. □

The fact that there is no effective enumeration of $M_0, M_1, M_2, \ldots$ prevents the sequence X from being computable. In a real-world "online" environment, however, a casino may force gamblers to place their wagers within a fixed number of seconds. In this online case, the analogous diagonal sequence X' depends only on a list of time-bounded martingales, which can be effectively enumerated, and therefore X' is computable.

Question. *Does the theorem above generalize to the case where A and B are infinite?*

As a particular instance, does betting with multiples of some minimum value, i.e. the set of integers, have as much power as gambling with nonzero wagers merely bounded away from zero, i.e. $\{x \in \mathbb{Q} : |x| \geq 1\}$ [1]?

References

[1] Laurent Bienvenu, Frank Stephan, and Jason Teutsch. How powerful are integer-valued martingales? In Fernando Ferreira, Benedikt Löwe, Elvira Mayordomo, and Luís Mendes Gomes, editors, *Programs, Proofs, Processes (CiE 2010)*, volume 6158 of *Lecture Notes in Computer Science*, pages 59–68. Springer-Verlag, Berlin, Heidelberg, 2010.

[2] Rodney G. Downey and Denis R. Hirschfeldt. *Algorithmic randomness and complexity.* Theory and Applications of Computability. Springer, New York, 2010.