A countable hat game and the importance of measurability

Jalex Stark

December 29, 2016

Abstract

We study a classic guessing game played by an infinite sequence of players wearing randomly colored hats. We use the axiom of choice to give a strategy that always wins, and then we show the main reults: any measurable strategy wins with probability strictly less than 1. The proof of the latter is an application of the Kolmogorov 0-1 law. The main result has likely been considered before, but the author has been unable to find a reference. Several further questions are posed at the end of section 2.

1 The game

We start with a classic puzzle.

Problem 1.1. n people stand in a line. The game master puts either a black hat or a white hat on each person's head, independently with equal probability. Each person can see the hats of all of the people in front of them, but not of those people behind them. Starting with the person in the back, the game master asks each player to announce a guess for their own hat color. Each player hears all of the guesses of the people who go before them. They win if all but the first person guesses correctly. The n players collude beforehand to come up with a fixed joint strategy. What is the highest probability with which they can win?

To be more precise, let's identify white hats with the symbol 0 and black hats with the symbol 1. Let's call the color of the hat on the n^{th} player's head X_n. Let's call the guess announced by the n^{th} player Y_n. The strategy of the first player can be described by some function $f_1^{(n)} : \{0,1\}^{n-1} \to \{0,1\}$ with the formula $Y_1 = f_1^{(n)}(X_2, \dots, X_n)$. The i^{th} player's strategy is also a function $f_i^{(n)} : \{0,1\}^{n-1} \to \{0,1\}$, but now remember that the i^{th} player hears the guesses of the players that come before them. We can write $Y_i = f_i(Y_1, \dots, Y_{i-1}, X_{i+1}, \dots, X_n)$. The players win the game if for all $i > 1$, we have $Y_i = X_i$.

Solution to Problem 1.1. They can win with probability 1. That is, we'll find a solution that works for all assignments to the X_i. In fact, all of the $f_i^{(n)}$ will be the same function: XOR. This is the function that returns 0 if an even number of arguments are 1 and returns 1 if an odd number of arguments are 1. We'll write the XOR of two bits a and b as $a \oplus b$ and the XOR of n bits $a_1, \ldots, a_n$ as $\bigoplus_{i=1}^n a_i$ or as $a_1 \oplus a_2 \oplus \cdot \oplus a_n$. $\oplus$ is in fact addition in the field with two elements. In particular, it is commutative and associative, has 0 as the identity, and satisfies $a \oplus a = 0$ for all a. We'll define the strategy functions as:

$$\begin{aligned} f_i^{(n)}(Y_1, \ldots, Y_{i-1}, X_{i+1}, \ldots, X_n) &= Y_1 \oplus \cdots \oplus Y_{i-1} \oplus X_{i+1} \oplus \cdots \oplus X_n \\ &= \bigoplus_{j=0}^{i-1} Y_j \oplus \bigoplus_{j=i+1}^{n-1} X_j \end{aligned}$$

To see that the strategy wins with probability 1, we show by induction on k that $X_k = Y_k$ for all $k > 1$. For $k = 2$, we see that

$$\begin{aligned} Y_2 = Y_1 \oplus \bigoplus_{j=3}^{n} X_j &= \left(X_2 \oplus \bigoplus_{j=3}^{n} X_j \right) \oplus \bigoplus_{j=3}^{n} X_j \\ &= X_2 \oplus \left(\bigoplus_{j=3}^{n} X_j \oplus \bigoplus_{j=3}^{n} X_j \right) = X_2. \end{aligned}$$

Now assume that $X_i = Y_i$ for all $i < k$. Then we have

$$Y_k = Y_1 \oplus \left(\bigoplus_{j=2}^{k-1} Y_j \oplus \bigoplus_{k+1}^{n} X_j \right) = \bigoplus_{1 \le j \le n} X_j \oplus \bigoplus_{\substack{1 \le j \le n \\ j \ne k}} X_j = X_k. \tag{1}$$

□

Problem 1.2. What if we have a countably infinite sequence of players?

Formally, we now ask about a countable sequence of independent identically distributed (iid) hat colors $(X_i)_{i \in \mathbb{N}}$ and a countable sequence $(f_i)_{i \in \mathbb{N}}$ of strategy functions. We define $Y_i = f_i(Y_1, \ldots Y_{i-1}, X_{i+1}, \ldots)$ and we say the players win if $X_i = Y_i$ for all $i \ge 1$.

In the finite case, essentially the same strategy worked for every n, and the correctness of this strategy follows from an induction proof. Naïvely, we should be able to use this to give a limiting strategy in the infinite case.

"Solution". Set $Y_k^{(n)} = \bigoplus_{i=1}^{k-1} Y_i \oplus \bigoplus_{i=k+1}^{n} X_i$. This is what the k^{th} person would announce if they intended to play the game with only the first n players in line. Set Y_k equal to $\lim_n Y_k^{(n)}$, ignoring for now the question of whether this

limit exists. We want to argue that $X_k = Y_k$ for all $k > 1$. We proceed by induction.

$$\begin{aligned}
Y_k &= \lim_m Y_1 \oplus \bigoplus_{i=2}^{k-1} Y_i \oplus \bigoplus_{i=k+1}^{m} X_i \\
Y_k &= \lim_m Y_1 \oplus \bigoplus_{i=2}^{k-1} X_i \oplus \bigoplus_{i=k+1}^{m} X_i \\
Y_k &= \lim_m Y_1 \oplus \bigoplus_{i=2}^{m} X_i \oplus X_k \\
Y_k &= Y_1 \oplus X_k \oplus \lim_m \bigoplus_{i=2}^{m} X_i \\
Y_k &= Y_1 \oplus X_k \oplus Y_1 \\
Y_k &= X_k.
\end{aligned}$$

In the fourth line, we use the fact that adding a constant is a continuous function, and that applying continuous functions commutes with taking limits. □

There is of course one gaping flaw in this argument: with probability 1, the limits defining the Y_i do not exist. Then we're not even close to defining a strategy yet: we've only solved the case where all but finitely many of the hats are 0. At the same time, our "proof" really was using properties of the limit to say things about the performance of our strategy. What limit properties did we need to make it work?

Proposition 1.1. *There exists a function* $\lim^*$ *from sequences of reals to reals satisfying the following properties:*

1. *If* (x_n) *converges, then* $\lim^*_n(x_n)$ *agrees with* $\lim_{n\to\infty} x_n$.
2. $\lim^*_n x_n$ *is always defined.*
3. *If* f *is a continuous function, then* $f(\lim^*_n x_n) = \lim^*_n f(x_n)$.

Theorem 1.2. *There is a strategy which wins the countable hat game on every input.*

Proof. From our earlier "solution", replace lim with $\lim^*$ from proposition 1.1. □

A function $\lim^*$ satisfying the criteria of the proposition (and a few other properties) is called an *ultralimit.* Their existence requires the Axiom of Choice, which means that the resulting strategy is not easily described. In particular, there is no good answer to the question, what should the first person say if all hats are colored 1? Indeed, there are different ultralimits the give different answers to this question. We use the Axiom of Choice to nonconstructively choose a way to answer all such questions simultaneously and consistently. The rest of this article will argue that this use of the axiom of choice is essential: if we try to do this in an explicit, constructive way, then we will fail. First, we'll need to introduce some concepts from measure theory.

2 Nice strategies don't work

2.1 A touch of probability theory

We'll quickly review just enough measure theory to state and prove the main theorem. For more detail, I recommend Omer Tamuz's wonderful lecture notes. [2]

Definition 2.1 (Probability space)**.** A *probability space* is a triple $(\Omega, \Sigma, \Pr)$. Ω is a set known as the *sample space.* $\Sigma \subseteq \mathcal{P}(\Omega)$ is a σ-algebra which we'll call the *event space.* $\Pr : \Sigma \to [0, 1]$ is the *probability measure*, which must satisfy the axioms of a countably additive measure.

A point in Ω can be thought of as a fixed outcome of an underlying random process. The event space Σ tells you what kinds of things you can and can't talk about. The probability measure Pr tells you how likely each event is.

Example 2.2. Let $\Omega = [0, 1] \subseteq \mathbb{R}$ and say that a set $X \subseteq \Omega$ is an event (that is, $X \in \Sigma$) iff X can be obtained by starting from intervals and applying countable unions, countable intersections, and complements. (This is known as the *Borel σ-algebra.*) Finally, let $\Pr[(a, b)] = b - a$, and extend the probability measure by countable additivity and complements.

We can think of this probability measure as picking a uniform random point on $[0, 1]$.

Example 2.3. Let $\Omega' = \{0, 1\}^{\mathbb{N}}$ be the space of countably infinite bit-strings and say that a set is an event if it can be obtained from "prefix-sets" of the form $A_w = \{\omega \in \Omega' : w \subseteq \omega\}$ by applying countable unions, countable intersections, and complements. (In the previous, w is a finite bit string and we say that $w \subseteq \omega$ if w is a prefix of ω, i.e. w has length $|w|$ and the first $|w|$ bits of ω are the bits of w.) Finally, let $\Pr[A_w] = 2^{-|w|}$ and extend the probability measure by countable additivity and complements. We can think of this probability measure as picking a string by flipping an infinite sequence of independent fair coins and recording the results.

These two probability spaces are in fact isomorphic by the map that sends a real number to its binary representation.

Definition 2.4. Let $f : \Omega \to \Omega'$ be a function between measure spaces. We say that f is *measurable* if, for every event $A \in \Sigma'$, the set $[f \in A] := \{\omega \in \Omega : f(\omega) \in A\}$ is an event in Σ. Note that the notin of measurability depends on the measures Σ and Σ'; we usually supress this from the notation, since the choice of measure is usually clear.

Intuitively, we should think of measurable functions as "the nice functions" between probability spaces. For example, the Banach–Tarski paradox requires the existence of non-measurable functions. Finally, we can state the main theorem.

Theorem 2.5. *If $(f_i)_{i\in\mathbb{N}}$ is a strategy for the countably infinite hat game which wins with probability 1, then f_1 is not a measurable function.*

The proof will be an application of Kolmogorov's 0-1 law, which applies to sequences of independent events.

Definition 2.6. We say two events A, B are *independent* if one of the following equivalent conditions holds:

- $\Pr[A \wedge B] = \Pr[A]\Pr[B]$, or
- $\Pr[A \mid B] = \Pr[A]$, or
- $\Pr[B \mid A] = \Pr[B]$.

Where $\Pr[A \mid B]$ denotes the conditional probability of A given B. We say that A is independent of the finite collection $\{B_i\}$ if A is independent from every event formed by intersections and complements of the B_i. We say that a collection of events $\mathcal{C}$ is independent if each event $A \in \mathcal{C}$ is independent from the subcollection $\mathcal{C} \setminus A$.

Theorem 2.7 (Kolmogorov)**.** *Let $(A_1, A_2, \ldots)$ be a a countable sequence of independent events. Suppose that $A = f(A_1, A_2, \ldots)$ for some measurable function f and that A is independent from every finite subcollection of A_i. (We say that A is a* tail event.*) Then* $\Pr[A] \in \{0, 1\}$.

This is actually a slight weakening of the fact that is usually referred to as "Kolmogorov's 0-1 law". For a more general statement, see any probability theory text. [2]

Proof of theorem 2.5. Assume f_1 is measurable, so that in particular, $\Pr[Y_1 = 0]$ and $\Pr[Y_1 = 1]$ are well-defined. We'll derive a contradiction by showing that they must be equal by symmetry and that they must be in $\{0, 1\}$ by Kolmogorov. (This is absurd since they must also sum to 1.) , Let $\Omega = \{0, 1\}^{\mathbb{N}}$ be the space of possible hat assignments, with probability measure given by a countable sequence of independent fair coin flips. Let $p_k : \Omega \to \Omega$ be the function which flips the k^{th} bit of a sequence. p_k is a measure-preserving bijection. Let $W \subseteq \Omega$ be the set of strings on which $X_i = Y_i$ for all $i > 1$. By assumption, W is a probability 1 event. Define $W_k = W \cap p_k(W)$. Since p_k is measure-preserving, $\Pr[p_k(W)] = 1$ and indeed $\Pr[W_k] = 1$.

Let $\omega \in W_k$. Consider the game as played on ω and on $p_k(\omega)$, as seen from the k^{th} player's perspective for $k > 1$. We have

$$\omega_k = f_k(f_1(\omega), \omega_1, \ldots, \omega_{k-1}, \omega_{k+1}, \ldots) \tag{2}$$
$$p_k(\omega)_k = f_k(f_1(p_k(\omega)), \omega_1, \ldots, \omega_{k-1}, \omega_{k+1}, \ldots) \tag{3}$$

The bits defined in the above two lines are unequal. Since f_k is a deterministic function, we must have that its inputs on the two lines are unequal. So $f_1(\omega) \neq f_1(p_k(\omega))$. Since p_k is measure preserving and W_k is p_k-invariant, we have

$$\Pr[Y_1 = 0 \mid W_k] = \Pr[Y_1 = 1 \mid W_k]. \tag{4}$$

Since W_k is a probability 1 event, we get the same for the unconditional probabilities: $\Pr[Y_1 = 0] = \frac{1}{2} = \Pr[Y_1 = 1]$. Now let $w \in \{0,1\}^{k-1}$ and let A_w be the event that the first $k-1$ hat colors are given by w. Since A_w is a positive probability event and W_k is a probability 1 event, we have that $\Pr[Y_1 = 0 \mid A_w] = \Pr[Y_1 = 0 \mid A_w \wedge W_k]$. The argument leading to equation (4) holds for any p_k-invariant subevent of W, (W_k being the maximal such event) in particular for $W_k \cap A_w$. Therefore,

$$\begin{aligned}\Pr[Y_1 = 0 \mid A_w] &= \Pr[Y_1 = 0 \mid A_w \wedge W_k] \\ &= \frac{1}{2} = \Pr[Y_1 = 1 \mid A_w \wedge W_k] = \Pr[Y_1 = 1 \mid A_w].\end{aligned} \tag{5}$$

Then $\frac{1}{2} = \Pr[Y_1 = 0 \mid A_w] = \Pr[Y_1 = 0]$, so the event $[Y_1 = 0]$ is independent of the choices of the first $k-1$ hat colors. k was arbitrary, so by Kolmogorov, $\Pr[Y_1 = 0] \in \{0, 1\}$. Contradiction! □

2.2 Further questions

We've proven that in this game, all measurable strategies have probability less than 1. I suspect that in fact they all have win probability 0 but have been unable to prove this.

Problem 2.1. Either show that all measurable strategies win with probability 0 or exhibit a measurable strategy winning with positive probability.

How about a weaker statement, that measurable strategies have win probabilities bounded away from 1?

Problem 2.2. Either find a sequence of measurable strategies whose win probabilities approach 1 or prove that no such sequence exists.

Here's a fun fact which helps point out the kinds of obstructions that exist in finding proofs for the above facts: assuming Choice, for every $p \in [0,1]$, there is a strategy which wins with probability exactly p. (Hint: to win with probability $\frac{1}{2}$, take the probability 1 strategy and have the 0^{th} person say the wrong thing if the 1^{st} person has a black hat.)

We can also think about relaxing the win condition.

Problem 2.3. Play the same game, but now say that the players win if all but finitely many guess their hat color correctly. Give a measurable strategy winning with probability 1 or prove that none exists.

A nonexistence result here would imply our main theorem as a corollary. As a first step towards this problem, we might consider the win condition where all but the first k people guess correctly. Examining the structure of our proof, we show that the conditional min-entropy of $(Y_1, \ldots, Y_k)$ is at least 1, conditioned on any finite substring of the hats. In the case $k = 1$, this proves that Y_1 is uniform and tail, so we can apply Kolmogorov. Perhaps this bound on the conditional min-entropy is enough to apply some stronger 0-1law, or maybe we need a different technique altogether. This also tells us that our proof is not already enough to handle the case of m hat colors (wherein we need to show $\log m$ conditional min-entropy in order to prove that Y_1 is uniform and tail.)

Problem 2.4. Modify the game so that all players guess simultaneously, i.e. Y_k is a function only of $(X_{k+1}, X_{k+2}, \ldots)$. Again require that all but finitely many people guess correctly. It is well-known that with the axiom of choice, there is a strategy that always wins. Give a measurable strategy winning with probability 1 or show that none exists.

3 Appendix: Ultralimits

Let's recall the usual definition of a limit of a sequence over $\mathbb{R}$, and then play with it until we can get an idea for how to define $\lim^*$.

Definition 3.1 (Limit of a sequence, I)**.**

$$\lim_{n\to\infty} x_n = L \Leftrightarrow \forall\varepsilon\exists N[|x_n - L| < \varepsilon \text{ if } n > N] \tag{6}$$

In words, we say that (x_n) converges to L if once we fix $\varepsilon > 0$, there is some largest N so that x_N is ε-far away from L. Reframing this, we could say that the set of n so that x_n is ε-far from L is finite.

$$\lim_{n\to\infty} x_n = L \Leftrightarrow \forall\varepsilon[\{n : |x_n - L| > \varepsilon\} \text{ is finite}] \tag{7}$$

Instead of saying that the set of bad indices is finite, we might say that the set of good indices is *cofinite.* To make this notationally cleaner, let's introduce the collection of all cofinite sets $\mathcal{F}_0 = \{A \subseteq \mathbb{N} : |\mathbb{N} \setminus A| < \infty\}$. This collection is sometimes referred to as the *Fréchet filter.*

Definition 3.2 (Limit of a sequence, II)**.**

$$\lim_{n\to\infty} x_n = L \Leftrightarrow \forall\varepsilon[\{n : |x_n - L| \le \varepsilon\} \in \mathcal{F}_0] \tag{8}$$

Take a moment to convince yourself that this definition coincides with Definition 3.1. We want to define a notion of limit that converges *more often*, so we want to make our definition *easier to satisfy.* In order to do this, let's enlarge $\mathcal{F}_0$.

Definition 3.3 ($\mathcal{F}$-limit of a sequence)**.** Let $\mathcal{F} \supseteq \mathcal{F}_0$.

$$\lim_{n\to\mathcal{F}} x_n = L \Leftrightarrow \forall\varepsilon \, \{n : |x_n - L| \le \varepsilon\} \in \mathcal{F} \tag{9}$$

Our goal is to find $\mathcal{F}$ so that $\lim_{n\to\mathcal{F}}$ satisfies the properties of Proposition 1.1 and is also a nice notion of limit. For a moment, let's restrict ourselves to 0-1 sequences. Let's suppose that $x_n \in \{0,1\}$ for all n and $\lim_{n\to\mathcal{F}} x_n = 1$. Taking any $\varepsilon \in (0,1)$ in definition 3.3 shows that $\{i : x_i = 1\} \in \mathcal{F}$.

If (x_n) is any sequence from $\{0,1\}^{\mathbb{N}}$, then its $\mathcal{F}$-limit must be either 0 or 1. Then of the two sets $X_1 = \{i : x_i = 0\}$ and $X_1 = \{i : x_i = 1\}$, at least one must be in $\mathcal{F}$.

Suppose $\lim_{n\to\mathcal{F}} x_n = 1$ and $\lim_{n\to\mathcal{F}} y_n = 1$. Then we'd like $\lim_{n\to\mathcal{F}} x_n y_n = 1$ as well. This holds iff $\{i : x_i = 1 \wedge y_i = 1\} \in \mathcal{F}$. This set is exactly $X_1 \cap Y_1$, where X_1 is as defined above and Y_1 is similar. We conclude that $\mathcal{F}$ should be closed under intersection.

Suppose $\lim_{n\to\mathcal{F}} x_n = 1$ and $\lim_{n\to\mathcal{F}} y_n = 0$. Then we hope $\lim_{n\to\mathcal{F}}(x_n + y_n) = 1$ also. In the special case that $Y_1 \subseteq X_1$, this amounts to requiring that the limit stays 1 when we flip some of the bits from 0 to 1. In other words, we ask that $\mathcal{F}$ is closed upwards.

Summarizing the above discussion, we want $\mathcal{F}$ to satisfy the following definition.

Definition 3.4. We say that $\mathcal{F} \subseteq \mathcal{P}(\mathbb{N})$ is a *nonprincipal ultrafilter* if it is...

(i) closed upwards, i.e. $B \supseteq A \in \mathcal{F}$ implies $B \in \mathcal{F}$

(ii) closed under intersection, i.e. $A, B \in \mathcal{F}$ implies $A \cap B \in \mathcal{F}$

(iii) nonprincipal, i.e. contains the Fréchet filter

(iv) nontrivial, i.e. does not contain $\emptyset$

(v) "ultra" or maximal, i.e. for every set A, either $A \in \mathcal{F}$ or $\mathbb{N} \setminus A \in \mathcal{F}$.

Axiom 3.6 is a form of the Axiom of Choice from which we can easily derive the existence of nonprincipal ultrafilters.

Definition 3.5. Let $(P, \leq)$ be a partially ordered set. We say that a subset $C \subseteq P$ is a *chain* if it is totally ordered, i.e. for every $a, b \in C$ we have $a \leq b$ or $b \leq a$. We say that c is an *upper bound for C* if for all $a \in C$, we have $a \leq c$.

Axiom 3.6 (Zorn's Lemma)**.** *Suppose $(P, \leq)$ is a poset such that that every chain in P has an upper bound. Then P has a maximal element, i.e. an element m such that $m \leq m'$ only if $m = m'$.*

For a direct statement of the Axiom of Choice and a proof that Zorn's Lemma is equivalent, see any introductory set theory text.

Proposition 3.7. *Nonprincipal ultrafilters exist.*

Proof. Consider the poset whose elements are set families satisfying the first four properties of definition 3.4 of $\mathcal{F}_0$, ordered by set inclusion (such families are called *filters*). It is easy to check that the union of a chain of filters is again a filter. Therefore chains in this poset have upper bounds, and we conclude that there are maximal elements, i.e. nonprincipal ultrafilters. □

A nice introduction to ultrafilters extending our treatment can be found in [1].

4 Acknolwedgements

Thanks to Professor Omer Tamuz for posing this problem as part of his probability class. Thanks to William Hoza and Tynan Ochse for providing feedback on early drafts and posing some of the questions in §2.2. Thanks to Aaron Anderson for helping work out the proof of the main result. Thanks to Sarah Deretic and Kevin Shu for useful conversations during the write-up.

References

[1] David Galvin. "Ultrafilters, with applications to analysis, social choice and combinatorics". In: *unpublished notes* (2009). URL: `https://www3.nd.edu/~dgalvin1/pdf/ultrafilters.pdf`.

[2] Omer Tamuz. "Lecture notes on Probability". In: *unpublished notes* (2016). URL: `http://people.hss.caltech.edu/~tamuz/teaching/ma144a/lectures.pdf`.