

An Introduction to Infinite Hat Problems

CHRISTOPHER S. HARDIN AND ALAN D. TAYLOR

Hat-coloring puzzles (or *hat problems*) have been around at least since 1961 (Gardner 1961), and probably longer. They gained wider public attention with a question posed and answered by Todd Ebert in his 1998 Ph.D. dissertation (Ebert 1998). The problem was presented by Sara Robinson in the April 10, 2001, Science section of *The New York Times* as follows:

Three players enter a room and a red or blue hat is placed on each person's head. The color of each hat is determined by a coin toss, with the outcome of one coin toss having no effect on the others. Each person can see the other players' hats but not his own.

No communication of any sort is allowed except for an initial strategy session before the game begins. Once they have had a chance to look at the other hats, the players must simultaneously guess the color of their own hats or pass. The group shares a hypothetical \$3 million prize if at least one player guesses correctly and no players guess incorrectly.

The same game can be played with any number of players. The general problem is to find a strategy for the group that maximizes its chances of winning the prize.

If one player guesses randomly and the others pass, the probability of a win is $1/2$. But Ebert's three-player solu-

tion is better: pass if the two visible hats are different colors, and guess the missing color if they are the same. This strategy yields a win, on average, $3/4$ of the time: of the eight possible hat assignments, it fails only on the two in which all three hats are the same color. Elwyn Berlekamp generalized this to $n = 2^k - 1$ players, using Hamming codes to show the existence of a strategy that yields a win with probability $n/(n + 1)$. Joe Buhler gives an account of this, and further variations, in Buhler (2002).

In Spring 2004, Yuval Gabay and Michael O'Connor, then graduate students at Cornell University, produced a number of hat problems involving infinitely many players, one of which was (an equivalent of) what we will call the Gabay-O'Connor hat problem:

Infinitely many players enter a room and a red or blue hat is placed on each player's head as before. Each player can see the other players' hats but not his own. Again, no communication of any sort is allowed except for an initial strategy session before the game begins. But this time, passing is not allowed and each player receives \$1 million if all but finitely many players guess correctly.

There are simple strategies ensuring that infinitely many players will guess correctly. For example, let a player guess red if he sees infinitely many red hats, and guess blue otherwise. If there are infinitely many red hats, everyone will guess red, and the players with red hats will be correct; if

there are finitely many red hats, everyone will guess blue, and the infinitely many players with blue hats will be correct.

The problem, however, seeks a strategy ensuring that all but finitely many—not just infinitely many—are correct, and this is what Gabay and O'Connor obtained using the axiom of choice. The special case in which the set of players is countable follows from a 1964 result of Fred Galvin (1965); see also Thorp (1967). Although Galvin's argument and the Gabay-O'Connor argument are similar, they are different enough that neither subsumes the other; a comparison will appear elsewhere.

As the title suggests, this paper is meant to be only an introduction to infinite hat problems, and as such proceeds in a somewhat expository manner. We have made no attempt here to say anything of the relevance of hat problems to other areas of mathematics, but the reader wishing to see some of this can begin with Galvin and Prikry (1976), George (2007), and Hardin and Taylor (2008).

The rest of the paper is organized as follows. In "The Formalism and the Finite," we set up a general framework for hat problems of the Gabay-O'Connor type, and present a few results in the finite case. "Theorems of Gabay-O'Connor and Lenstra" concerns the infinite case, which is our primary interest; we present the Gabay-O'Connor Theorem, and a theorem of Lenstra involving strategies that either make every player correct or every player incorrect. In "The Necessity of the Axiom of Choice," we discuss the necessity of the axiom of choice in the Gabay-O'Connor Theorem and Lenstra's Theorem; this section requires some basic facts about the property of Baire, so a short appendix on the property of Baire appears afterward.

Our set-theoretic notation and terminology are standard. If A is a set, then $|A|$ is the cardinality of A and A^c is the complement of A . If f is a function, then $f|A$ is the restriction of f to A , and ${}^P C$ is the set of functions mapping the set P into the set C . If x is a real number, then $\lfloor x \rfloor$ is the greatest integer that is less than or equal to x . We let $\mathbb{N} = \{0, 1, 2, \dots\}$.

The authors thank James Guilford, John Guilford, Hendrik Lenstra, Michael O'Connor, and Stan Wagon for al-

lowing us to include unpublished proofs that are in whole or in part due to them. Their specific contributions will be noted at the appropriate places. We also thank Andreas Blass for bringing Galvin's work to our attention, and thank the referee for many helpful suggestions.

The Formalism and the Finite

The problems we consider will resemble the Gabay-O'Connor hat problem, but we allow more generality: the set of players can be any set, there can be any number of hat colors, players do not necessarily see all other hats, and the criterion for winning is not necessarily that all but finitely many players guess correctly. So, a particular hat problem will be described by (i) the set of players, (ii) the set of possible hat colors, (iii) which hats each player can see, and (iv) a rule that indicates, given the set of players who guess correctly, whether or not they win the game. We formally define a *hat problem* to be a tuple (P, C, V, \mathcal{W}) with the following properties.

- (i) The *set of players* P is any set.
- (ii) The *set of colors* C is any set.
- (iii) The *visibility graph* V is a directed graph with P as the set of vertices. When there is an edge from a to b (which we denote by aVb or $b \in V_a$), we interpret this as meaning that a can see (the hat worn by) b . In particular, V_a is the set of players visible to a . We are only interested in cases where players cannot see their own hats, so we require that V has no edges from vertices to themselves.
- (iv) The *winning family* \mathcal{W} is a family of subsets of P . The players win iff the set of players who guess their own hat color correctly is in \mathcal{W} .

A function $g \in {}^P C$ assigns a hat color to each player; we call g a *coloring*. Given a hat problem (P, C, V, \mathcal{W}) , a *strategy* is a function $S : (P \times {}^P C) \rightarrow C$ such that for any $a \in P$ and colorings $g, b \in {}^P C$,

$$g|_{V_a} = b|_{V_a} \Rightarrow S(a, g) = S(a, b). \quad (1)$$

We think of $S(a, g)$ as the color guessed by player a under coloring g . Condition (1) ensures that this guess only depends



CHRISTOPHER S. HARDIN, B.A. Amherst 1998 and Ph.D. Cornell 2005, is currently a visiting professor at Wabash College. His publications are in the area of mathematical logic, and include (with Taylor) "A peculiar connection between the axiom of choice and predicting the future," *Amer. Math. Monthly* (2008). He enjoys climbing and music. He does not own blue jeans.

Department of Mathematics and
Computer Science
Wabash College
Crawfordsville, IN 47933
USA
e-mail: hardinc@wabash.edu



ALAN D. TAYLOR received his Ph.D. from Dartmouth in 1975, and has been at Union College ever since. He has published six books on set theory, combinatorics, fair division, and the theory of voting, including *Social Choice and the Mathematics of Manipulation* (Cambridge 2005). He still keeps trying to run track (400 meters).

Department of Mathematics
Union College
Schenectady, NY 12308
USA
e-mail: taylor@union.edu

on the hats that a can see, since $g|_{V_a} = b|_{V_a}$ means that the colorings g and b are indistinguishable to player a . We will frequently consider strategies player by player; for $a \in P$ and a strategy S , we define $S_a : {}^P C \rightarrow C$ by $S_a(g) = S(a, g)$, and call S_a a *strategy for player a* . We say that player a *guesses correctly* if $S_a(g) = g(a)$.

We call S a *winning strategy* if it ensures that the set of players who guess correctly is in the winning family, regardless of the coloring; that is, $\{a \in P : S_a(g) = g(a)\} \in \mathcal{W}$ for any coloring g .

To illustrate the kinds of questions and answers that arise within this framework, we present two results in the context of finitely many players. For the first, say that a hat problem is a *minimal hat problem* if it asks for a strategy ensuring that at least one player guesses correctly, and call such a strategy a *minimal solution*. Our first result (the second half of which is due, in part, to James Guilford and John Guilford) answers the following question.

With k players and 2 colors, how much visibility is needed to guarantee the existence of a minimal solution? What if there are k players and k colors?

THEOREM 1 *A k -player, 2-color hat problem has a minimal solution iff the visibility graph has a cycle. A k -player, k -color hat problem has a minimal solution iff the visibility graph is complete.*

To prove Theorem 1, it will help to have a lemma that confirms an intuition about how many players guess correctly on average.

LEMMA 2 *In a k -player, n -color hat problem, for any particular strategy, the average number of players who guess correctly is k/n . (The average is taken over all colorings.)*

Proof. Suppose there are k players and n colors. Let S be any strategy. It suffices to show that any particular player a is correct in 1 out of n colorings. Given any assignment of hat colors to all players other than a , player a 's guess will be determined; of the n ways to extend this hat assignment to a , exactly one will agree with a 's guess. \square

Proof of Theorem 1. Suppose first that there are 2 colors. The right-to-left direction is easy; assuming that the visibility graph has a cycle, the strategy is for a designated player on the cycle to guess that his hat is the same color as that of the player immediately ahead of him on the cycle, while all the others on the cycle guess that their hat color is the opposite of the player immediately ahead of them. To see that this works, assume that the first player on the cycle has a red hat and that everyone on the cycle guesses incorrectly using this strategy. Then the second player on the cycle has a blue hat, the third player on the cycle has a blue hat, and so on until we're forced to conclude that the first player on the cycle has a blue hat, which we assumed not to be the case.

For the other direction, we show that if there is no cycle in the visibility graph V , then for every strategy there is a coloring for which everyone guesses incorrectly. To do

this, we first note that because V has no cycles, we can assign each player a rank as follows: a has rank k if there is a directed path of length k beginning at a , but none of length $k + 1$. Now, if there is a directed edge from vertex a to vertex b , then the rank of player a is strictly greater than the rank of player b . Thus, a player can only see hats of players of strictly smaller rank. Hence, given any strategy, we can assign hat colors in order of rank to make everyone wrong: once we have colored the hats of players of rank $< k$, the guesses of players of rank k are determined, and we can then color their hats to make them wrong.

Now suppose there are k colors. For the right-to-left direction, assuming the visibility graph is complete, the strategy is as follows. Number the players $0, 1, \dots, k - 1$, and the colors $0, 1, \dots, k - 1$, and for each i , let s_i be the mod k sum of the hats seen by player i . The plan is to have player i guess $i - s_i \pmod k$ as the color of his hat. If the colors of all the hats add to $i \pmod k$, then player i will be the one who guesses correctly. That is, if $c_0 + \dots + c_{k-1} = i \pmod k$ then $c_i = i - s_i \pmod k$.

For the other direction, assume that there are k players and k colors, and assume the visibility graph is not complete. Let S be any strategy. We must show that there is a coloring in which every player guesses incorrectly. Suppose player a does not see player b 's hat (with $a \neq b$), and pick a coloring in which player a guesses correctly. If we change the color of player b 's hat to match player b 's guess, player a will not change his guess, and we will have a coloring in which a and b guess correctly. By Lemma 2, the average number of players who guess correctly is $k/k = 1$; because we have a coloring with at least 2 players guessing correctly, there must be another coloring in which fewer than 1 (namely, zero) players guess correctly. \square

Our second result along these same lines is also due, in part, to James Guilford and John Guilford (the $n = 2$ case appears in Winkler 2001). It answers the following question.

With k players and n colors, how many correct guesses can a strategy guarantee, assuming the visibility graph is complete?

Lemma 2 shows us that, regardless of strategy, the number who guess correctly will *on average* be k/n . But this is very different from ensuring that a certain fraction will guess correctly regardless of luck or the particular coloring at hand. Nevertheless, the fraction k/n is essentially the correct answer.

THEOREM 3 *Consider the hat problem with $|P| = k$, $|C| = n$, and a complete visibility graph V . Then there exists a strategy ensuring that $\lfloor k/n \rfloor$ players guess correctly, but there is no strategy ensuring that $\lfloor k/n \rfloor + 1$ players guess correctly.*

Proof. The strategy ensuring that $\lfloor k/n \rfloor$ players guess correctly is obtained as follows. Choose $n \times \lfloor k/n \rfloor$ of the players (ignoring the rest) and divide them into $\lfloor k/n \rfloor$ pairwise disjoint groups of size n . Regarding each of the groups as an n -player, n -color hat problem, we can apply the previ-

ous theorem to obtain a strategy for each group ensuring that (precisely) one in each group guesses correctly. This yields $\lfloor k/n \rfloor$ correct guesses altogether, as desired.

For the second part, we use Lemma 2. For any strategy, the average number of players who guess correctly will be $\lfloor k/n \rfloor$, and $\lfloor k/n \rfloor < \lfloor k/n \rfloor + 1$, so no strategy can guarantee at least $\lfloor k/n \rfloor + 1$ players guess correctly for each coloring. \square

Theorem 3, and most of Theorem 1, were obtained independently by Butler, Hajiaghay, Kleinberg, and Leighton (2008; see this for a considerably more detailed investigation of the finite context).

With two colors and an even number of players, Theorem 3 says that—with collective strategizing—the on-average result of 50% guessing correctly can, in fact, be achieved with each and every coloring. But it also says that this is the best that can be done by collective strategizing. In the finite case, this latter observation does little more than provide proof for what our intuition suggests: collective strategizing notwithstanding, the on-average result of 50% cannot be improved in a context wherein guesses are simultaneous. The infinite, however, is very different, and it is to this that we next turn.

Theorems of Gabay-O'Connor and Lenstra

We begin with a statement and proof of what we will call the Gabay-O'Connor Theorem. As stated, this result is strong enough to solve the Gabay-O'Connor hat problem and to allow us to derive Lenstra's Theorem (below) from it. (One can use an arbitrary filter in place of the collection of cofinite sets, with essentially the same proof, to generalize the result.)

THEOREM 4 (GABAY-O'CONNOR) *Consider the situation in which the set P of players is arbitrary, the set C of colors is arbitrary, and every player sees all but finitely many of the other hats. Then there exists a strategy under which all but finitely many players guess correctly. Moreover, the strategy is robust in the sense that each player's guess is unchanged if the colors of finitely many hats are changed.*

Proof. For $b, g \in {}^P C$, say $b \approx g$ if $\{a \in P : b(a) \neq g(a)\}$ is finite; this is an equivalence relation on ${}^P C$. By the axiom of choice, there exists a function $\Phi : {}^P C \rightarrow {}^P C$ such that $\Phi(b) \approx b$, and if $b \approx g$, then $\Phi(b) = \Phi(g)$. Thus, Φ is choosing a representative from each equivalence class. Notice that for each coloring h , each player a knows the equivalence class $[h]$, and thus $\Phi(h)$, because the player can see all but finitely many hats. The strategy is then to have the players guess their hat colors according to the chosen representative of the equivalence class of the coloring; more formally, we are letting $S_a(b) = \Phi(b)(a)$. For any coloring b , since this representative $\Phi(b)$ only differs from b in finitely many places, all but finitely many players will guess correctly. Also, if finitely many hats change colors, the equivalence class remains the same and players keep the same guesses. \square

Theorem 4 is sharp in the sense that even with countably many players and two colors, no strategy can ensure that, for a fixed k , all but k players will guess correctly, even if everyone sees everyone else's hat. The reason is

that any such strategy would immediately yield a strategy for $2k + 1$ players in which more than 50% would guess correctly each time, contradicting Lemma 2.

The following theorem was originally obtained by Hendrik Lenstra using techniques (described below) quite different from our derivation of it here from Theorem 4.

THEOREM 5 (LENSTRA) *Consider the situation in which the set P of players is arbitrary, $|C| = 2$, and every player sees all of the other hats. Then there exists a strategy under which everyone's guess is right or everyone's guess is wrong.*

Proof. Let S be a strategy as in Theorem 4. A useful consequence of the robustness of S is that, for a given coloring b , a player a can determine $S_b(b)$ for every player b . Since we are assuming players can see all other hats, a also knows the value of $b(b)$ for every $b \neq a$. So, we may define a strategy T by letting $T_a(b) = S_a(b)$ iff $\{|b \in P : b \neq a \text{ and } S_b(b) \neq b(b)\}|$ is an even number. That is, the players take it on faith that, when playing S , an even number of players are wrong: if they see an even number of errors by others, they keep the guess given by S , and otherwise they switch.

To see that T works, let b be a given coloring. When $\{|b \in P : S_b(b) \neq b(b)\}|$ is even, every guess given by T will be correct: the players who were already correct under S will see an even number of errors (under S), and keep their guess; the players who were wrong under S will see an odd number of errors and will switch. When $\{|b \in P : S_b(b) \neq b(b)\}|$ is odd, the opposite occurs, and every guess given by T will be incorrect: the players who would be correct under S will see an odd number of errors and will switch (to the incorrect guess); the players who would be wrong under S will see an even number of errors and will stay (with the incorrect guess). \square

The assumption that everyone can see everyone else's hat in Theorem 5 is necessary. That is, if player a could not see player b 's hat, then changing player b 's hat would change neither his nor player a 's guess, but player b would go from wrong to right or vice-versa, and player a would not.

Lenstra's Theorem can be generalized from two colors to the case in which the set of colors is an arbitrary (even infinite) Abelian group. The conclusion is then that, for a given coloring, everyone's guess will differ from his true hat color by the same element of the group. Intuitively, the strategy is for everyone to take it on faith that the (finite) group sum of the differences between the true coloring and the guesses provided by the Gabay-O'Connor Theorem is the identity of the group. (Variants of this observation were made independently by a number of people.)

Lenstra's original proof is certainly not without its charms, and goes as follows. If we identify the color red with the number zero and the color blue with the number one, then we can regard the collection of all colorings as a vector space over the two-element field. The collection W of all colorings with only finitely many red hats is a subspace, and the function that takes each such coloring to zero if the number of red hats is even, and one otherwise, is a linear functional defined on W . The axiom of choice guarantees that this linear functional can be extended to the

whole vector space. Moreover, a coloring is in the kernel iff the changing of one hat yields a coloring that is not in the kernel. Hence, the strategy is for each player to guess his hat color assuming that the coloring is in the kernel. If the coloring is, indeed, in the kernel, then everyone guesses correctly. If not, then everyone guesses incorrectly.

Another proof of Lenstra's Theorem, at least for the case where the set of players is countably infinite, was found by Stan Wagon. It uses the existence (ensured by AC) of a so-called *non-principal ultrafilter on P* —that is, a collection \mathcal{U} of subsets of P that contains no finite sets, that is closed under finite intersections, and that contains exactly one of X and X^c for every $X \subseteq P$. Wagon's proof goes as follows. Label the players by natural numbers and call an integer a "red-even" if the number of red hats among players $0, 1, \dots, a$ is even. Player a 's hat color affects which integers $b > a$ are red-even in the sense that changing player a 's hat color causes the set of red-even numbers greater than a to be complemented. The strategy is for player a to make his choice so that, if this choice is correct, then the set of red-even numbers is in the ultrafilter \mathcal{U} . The strategy works because either the set of red-even numbers is in \mathcal{U} (in which case everyone is right) or the set of red-even numbers is not in \mathcal{U} (in which case everyone is wrong).

The Necessity of the Axiom of Choice

Some nontrivial version of the axiom of choice is needed to prove Lenstra's Theorem or the Gaby-O'Connor Theorem. Specifically, if we take the standard axioms of set theory (ZFC) and replace the axiom of choice with a weaker principle known as *dependent choice*, the resulting system ZF + DC is not strong enough to prove Lenstra's Theorem or the Gaby-O'Connor Theorem, even when restricted to the case of two colors and countably many players. Historically, the precursor to our results here is a slightly weaker observation (in a different but related context) of Roy O. Davies that was announced in Silverman (1966). The reader does not need any familiarity with ZF + DC; all that must be understood is that, as an axiom system, ZF + DC is weaker than ZFC, and somewhat stronger than ZF (set theory with the axiom of choice removed altogether).

To follow our argument, some basic facts about the property of Baire are needed; to this end, the appendix gives a short introduction to the property of Baire. As an aid to intuition, having the property of Baire is somewhat analogous to being measurable, whereas being meager (see appendix) is somewhat analogous to having measure 0. (The two notions should not be conflated too much, however: the real numbers can be written as the disjoint union of a measure 0 set and a meager set.)

Let BP be the assertion that every set of reals has the property of Baire. It is known (assuming ZF is consistent) that ZF + DC cannot disprove BP (Judah and Shelah 1993). (This was established earlier, assuming the existence of a large cardinal, in [Solovay 1970].) It follows that ZF + DC cannot prove any theorem that contradicts BP, as any such proof could be turned into a proof that BP fails. We will show that Lenstra's Theorem and the Gaby-O'Connor Theorem contradict BP, and thus ZF + DC cannot prove Lenstra's Theorem or the Gaby-O'Connor Theorem. Al-

though BP is useful for establishing results such as these, one should note that BP is false in ZFC (for instance, ZFC can prove Lenstra's Theorem, which contradicts BP).

Throughout this section, we take the set P of players to be the set \mathbb{N} of natural numbers, and we take the two colors to be 0 and 1. The topology and measure on $\mathbb{N}\{0, 1\}$ are the usual ones. That is, if s is a finite sequence of 0s and 1s, then the set $[s]$ of all infinite sequences of 0s and 1s that extend s is a basic open set whose measure is 2^{-n} , where n is the length of s . Identifying $\mathbb{N}\{0, 1\}$ with the binary expansions of reals in $[0, 1]$, this is Lebesgue measure. The topology is that of the Cantor set.

Let T_k be the measure-preserving homeomorphism from $\mathbb{N}\{0, 1\}$ to itself that toggles the k th bit in a sequence of 0s and 1s. Call a set $D \subseteq \mathbb{N}\{0, 1\}$ a *toggle set* if there are infinitely many values of k for which $T_k(D) \cap D = \emptyset$.

The next lemma is key to the results in this section; its proof makes use of the following observation. If a set D has the property of Baire but is not meager, then there exists a nonempty open set V such that the symmetric difference of D and V is meager. Hence, if we take any basic open set $[s] \subseteq V$, it then follows that $[s] - D$ is meager.

LEMMA 6 *Every toggle set with the property of Baire is meager.*

Proof. Assume for contradiction that D is a nonmeager toggle set with the property of Baire, and choose a basic open set $[s]$ such that $[s] - D$ is meager. Because D is a toggle set, we can choose k greater than the length of s such that $T_k(D) \cap D = \emptyset$. It now follows that $[s] \cap D \subseteq [s] - T_k(D)$. But $T_k([s]) = [s]$, because k is greater than the length of s . Hence, $[s] \cap D \subseteq [s] - T_k(D) = T_k([s]) - T_k(D) = T_k([s] - D)$. Thus, $[s] \cap D$ is meager, as was $[s] - D$. This means that $[s]$ itself is meager, a contradiction. \square

With these preliminaries, the following theorem (of ZF + DC) shows that Lenstra's Theorem contradicts BP, and hence it cannot be proven without some nontrivial version of the axiom of choice.

THEOREM 7 *Consider the situation in which the set P of players is countably infinite, there are two colors, and each player sees all of the other hats. Assume BP. Then for every strategy there exists a coloring under which someone guesses correctly and someone guesses incorrectly.*

Proof. Assume that S is a strategy and let D denote the set of colorings for which S yields all correct guesses, and let I denote the set of colorings for which S yields all incorrect guesses. Notice that both D and I are toggle sets, since changing the hat on one player causes his (unchanged) guess to switch from right to wrong or vice versa. If D and I both have the property of Baire, then both are meager. Choose $b \in \mathbb{N}\{0, 1\} - (D \cup I)$. Under b , someone guesses correctly and someone guesses incorrectly. \square

In ZFC, nonmeager toggle sets do exist: as seen in the previous proof, if all toggle sets are meager, then Lenstra's Theorem fails, but Lenstra's Theorem is valid in ZFC.

We derived Lenstra's Theorem from the Gabay-O'Connor Theorem, so Theorem 7 also shows that some non-trivial version of the axiom of choice is needed to prove the Gabay-O'Connor Theorem. However, the Gabay-O'Connor Theorem, even in the case of two colors and countably many players, is stronger than the assertion that the Gabay-O'Connor hat problem has a solution: the theorem does not require that players can see *all* other hats, and it produces not just a strategy, but a robust strategy. The following theorem (of ZF + DC) shows that any solution to the Gabay-O'Connor hat problem, even in the countable case, contradicts BP and hence requires some non-trivial version of the axiom of choice.

THEOREM 8 *Consider the case of the Gabay-O'Connor hat problem in which the set of players is countably infinite. Assume BP. Then for every strategy there exists a coloring under which the number of players guessing incorrectly is infinite.*

Proof. Assume that S is a strategy and, for each k , let D_k denote the set of colorings for which S yields all correct guesses from players numbered k and higher. Notice that each D_k is a toggle set, since changing the hat on a player higher than k causes his (unchanged) guess to switch from right to wrong. If all the D_k s have the property of Baire, then all are meager. Let D be the union of the D_k s, and choose $b \in {}^{\mathbb{N}}\{0, 1\} - D$. Under b , the number of people guessing incorrectly is infinite. \square

Theorems 7 and 8 can be recast in the context of Lebesgue measurability to show that both Lenstra's Theorem and the Gabay-O'Connor Theorem imply the existence of nonmeasurable sets of reals. However, to show that ZF + DC cannot prove the existence of nonmeasurable sets of reals, one must assume the consistency of ZFC plus the existence of a large cardinal (Solovay 1970, Shelah 1984). Although this is not a particularly onerous assumption, it is why we favored the presentation in terms of the property of Baire.

It turns out that with infinitely many colors, some non-trivial version of the axiom of choice is needed to obtain a strategy ensuring even one correct guess; this will appear elsewhere.

Appendix: The Property of Baire

DEFINITION 9 A subset N of a topological space is *nowhere dense* if the interior of its closure is empty. A set is *meager* if it is the union of countably many nowhere dense sets.

A set B has the *property of Baire* if it differs from an open set by a meager set; that is, there is an open set V and a meager set M such that $B \Delta V = M$ (equivalently, $B = V \Delta M$), where Δ denotes symmetric difference.

A topological space is a *Baire space* if its nonempty open sets are nonmeager.

THEOREM 10 (BAIRE CATEGORY THEOREM) *Every nonempty complete metric space is a Baire space.*

We do not show the proof here, but it can be carried out in ZF + DC. For the special cases of the reals and Cantor space, the proof can be carried out in ZF.

REFERENCES

- Steven Butler, Mohammad T. Hajiaghayi, Robert D. Kleinberg, and Tom Leighton. Hat guessing games. *SIAM Journal of Discrete Mathematics* 22:592–605, 2008.
- Joe P. Buhler. Hat tricks. *Mathematical Intelligencer* 24:44–49, 2002.
- Todd Ebert. *Applications of recursive operators to randomness and complexity*. PhD thesis, University of California at Santa Barbara, 1998.
- Fred Galvin. Problem 5348. *American Mathematical Monthly* 72:1136, 1965.
- Martin Gardner. *The 2nd Scientific American Book of Mathematical Puzzles & Diversions*. Simon and Schuster, New York, 1961.
- Alexander George. A proof of induction? *Philosophers' Imprint* 7:1–5, March 2007.
- Fred Galvin and Karel Prikry. Infinitary Jonsson algebras and partition relations. *Algebra Universalis* 6:367–376, 1976.
- Christopher S. Hardin and Alan D. Taylor. A peculiar connection between the axiom of choice and predicting the future. *American Mathematical Monthly* 115:91–96, February 2008.
- Haim Judah and Saharon Shelah. Baire property and axiom of choice. *Israel Journal of Mathematics* 84:435–450, 1993.
- Saharon Shelah. Can you take Solovay's inaccessible away? *Israel Journal of Mathematics* 48:1–47, 1984.
- D. L. Silverman. Solution of problem 5348. *American Mathematical Monthly* 73:1131–1132, 1966.
- Robert Solovay. A model of set theory in which every set of reals is Lebesgue measurable. *Annals of Mathematics* 92:1–56, 1970.
- B. L. D. Thorp. Solution of problem 5348. *American Mathematical Monthly* 74:730–731, 1967.
- Peter Winkler. Games people don't play. In David Wolfe and Tom Rodgers, editors, *Puzzlers' Tribute*, pages 301–313. A. K. Peters, Ltd., 2001.