

Chapter 2

The Finite Setting

2.1 Background

Although our primary interest in this monograph is with the infinite, we begin with a discussion of hat problems in which the set A of agents is finite and visibility is given by a directed graph V on A (the visibility graph). The set K of colors is typically a natural number k and the set C of colorings is the entire set ${}^A K$. Thus if f and g are two colorings in ${}^A K$ and a is an agent, then $f \equiv_a g$ iff $f(b) = g(b)$ for every $b \in V(a)$.

As indicated in the introduction, we think of each agent a as trying to guess his own hat color via a guessing strategy $G_a : C \rightarrow K$ satisfying $G_a(f) = G_a(g)$ whenever $f \equiv_a g$. The predictor P arising from these guessing strategies is given by $P(f)(a) = G_a(f)$, and agent a guesses correctly for f if $G_a(f) = f(a)$; equivalently, if $P(f)(a) = f(a)$.

Most of what is known in the finite case (where agents cannot pass) can be found in a single paper entitled *Hat Guessing Games*, by Steven Butler, Mohammed Hajiaghayi, Robert Kleinberg, and F. Thomson Leighton [BHKL08]. Some of these results were obtained independently and appeared in [HT08a].

The rest of this chapter is organized as follows. In Sect. 2.2 we illustrate the notion of a minimal predictor, including a result from [BHKL08] on bipartite visibility graphs that we will need in Chap. 4. In Sect. 2.3, we consider optimal strategies, and in Sect. 2.4 we present another result from [BHKL08] that uses the Tutte-Berge formula to completely solve the prediction problem for finite symmetric graphs. In Sect. 2.5 we consider the situation in which the agents have different color-sets for their hats, and in Sect. 2.6 we look at some variants of the standard hat problem, one of which will arise again in Chap. 3. We conclude in Sect. 2.7 with a discussion of some open questions.

2.2 Minimal Predictors

In the present context, a minimal predictor will be one that guarantees at least one correct guess. Thus, with n agents, we can ask how much visibility is needed for the existence of a minimal predictor. Stated differently, for a fixed number of colors, we seek a characterization of those visibility graphs that yield a minimal predictor. Our first theorem answers this for the case of two colors and the case of n colors; the result appears as Theorem 1 in [HT08a], although most of it can be derived from results in [BHL08]. But first we need a lemma that confirms an intuition about how many agents guess correctly on average.

Lemma 2.2.1. *In an n -agent, k -color hat problem, for any particular predictor, the average number of agents who guess correctly is n/k . (The average is taken over all colorings.)*

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

Theorem 2.2.2. *For an n -agent, 2-color hat problem, there is a predictor ensuring at least one correct guess iff the visibility graph has a cycle. For an n -agent, n -color hat problem, there is a predictor ensuring at least one correct guess iff the visibility graph is complete.*

Proof. Suppose first that there are two colors. For the right-to-left direction, assume the visibility graph has a cycle. Fix an agent on the cycle and let his strategy be to guess assuming his hat and the one on the agent just ahead of him on the cycle are the same color. The other agents on the cycle guess according to the opposite assumption; they assume that their hat and that of the agent just ahead of them on the cycle are different colors. To see that this works, assume that the first agent on the cycle has a red hat and that everyone on the cycle guesses incorrectly using this strategy. Then the second agent on the cycle has a green hat, the third agent on the cycle has a green hat, and so on until we're forced to conclude that the first agent on the cycle also has a green hat, contrary to what we assumed.

For the other direction, we can appeal to Corollary 1.3.4 which ensures that if there is no cycle in the visibility graph V , then for every predictor there is a coloring for which everyone guesses incorrectly.

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

For the other direction, assume that there are n agents and n colors, and assume the visibility graph is not complete. Let P be any predictor. We must show that there is a coloring in which every agent guesses incorrectly. Suppose agent a does not see agent b 's hat (with $a \neq b$), and pick a coloring in which agent a guesses correctly. If we change the color of agent b 's hat to match agent b 's guess, agent a will not change his guess, and we will have a coloring in which a and b guess correctly. By Lemma 2.2.1, the average number of agents who guess correctly is $n/n = 1$; because we have a coloring with at least two agents guessing correctly, there must be another coloring in which less than one (namely, zero) agents guess correctly. \square

We conclude this section with one other result about minimal predictors. It was first established in [BHKLO8] and later rediscovered by Daniel J. Velleman [Vel11] in his solution to a question left open in [HT10]. Velleman's result is in Chap. 4.

Theorem 2.2.3. *For every k there is a finite bipartite graph V such that if visibility is given by V with agent a seeing agent b only if a is adjacent to b , then there is a predictor for the k -color hat problem ensuring at least one correct guess.*

Proof. Thinking of the problem as involving two teams that see each other, the trick is to get team 1 to guess in such a way that if they're all wrong, there are relatively few possibilities for how team 2 is colored.

Let team 2 have $k - 1$ agents. There are only finitely many possible individual strategies for agents on team 1, so let team 1 have one agent for each possible individual strategy. We claim that given a coloring f of team 1 and k distinct colorings g_1, \dots, g_k of team 2, someone on team 1 is correct for at least one of these colorings; for, at least one of the agents on team 1 will guess a different color for each of g_1, \dots, g_k , so one of these guesses must agree with f . In light of that claim, team 2 does the following: assuming every guess on team 1 is wrong leaves at most $k - 1$ different possible colorings of team 2, so with $k - 1$ agents, they can let agent i among them guess according to the i th such possibility. \square

Note that with k colors, both teams need at least $k - 1$ agents, as the following argument shows. Suppose team 2 has fewer than $k - 1$ agents, and that team 1 has any set (possibly infinite) of agents. Fix any predictor P . We will commit to coloring team 2 with a constant coloring among $0, \dots, k - 2$. Let f_i be the coloring guessed by team 1 when team 2 is all colored i , for $i = 0, \dots, k - 2$. There are only $k - 1$ such f_i s, so for each agent a on team 1, there is a color that differs from $f_0(a), \dots, f_{k-2}(a)$, so we can fix a coloring h of team 1 that makes them all wrong when team 2 is all colored i , for $0 \leq i \leq k - 2$. Now, look at what team 2 guesses when they see h . Since they have only $k - 2$ agents, there is some color among $0, \dots, k - 2$ that none of them guesses; color everyone on team 2 this color.

2.3 Optimal Predictors

As we said in the preface, an optimal predictor achieves a degree of correctness that is maximal in some sense. For the case of a visibility graph that is complete, there is a very satisfying result that we present below. It occurs as Theorem 2 in [BHKL08] and as Theorem 3 in [HT08a], although it was first proved for two colors by Peter Winkler [Win01] and later generalized to k colors by Uriel Feige [Fei04].

By way of motivation, recall that Lemma 2.2.1 showed that, regardless of strategy, if there are n agents and k colors, the number who guess correctly will *on average* be n/k . 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 n/k is essentially the correct answer.

Theorem 2.3.1. *Consider the hat problem with $|A| = n$, $|K| = k$, and a complete visibility graph V . Then there exists a predictor ensuring that $\lfloor n/k \rfloor$ agents guess correctly, but there is no predictor ensuring that $\lfloor n/k \rfloor + 1$ agents guess correctly.*

Proof. The strategy ensuring that $\lfloor n/k \rfloor$ agents guess correctly is obtained as follows. Choose $k \times \lfloor n/k \rfloor$ of the agents (ignoring the rest) and divide them into $\lfloor n/k \rfloor$ pairwise disjoint groups of size k . Regarding each of the groups as a k -agent, k -color hat problem, we can apply Theorem 2.2.2 to get a strategy for each group ensuring that (precisely) one in each group guesses correctly. This yields $\lfloor n/k \rfloor$ correct guesses altogether, as desired.

For the second part, we use Lemma 2.2.1. For any predictor, the average number of agents who guess correctly will be n/k , and $n/k < \lfloor n/k \rfloor + 1$, so no predictor can guarantee at least $\lfloor n/k \rfloor + 1$ agents guess correctly for each coloring. \square

2.4 The Role of the Tutte-Berge Formula

In this section, we consider only symmetric visibility graphs on a finite set A of agents and we present a very nice result from [BHKL08] that specifies exactly how successful a predictor for two colors can be for such a visibility graph. The obvious strategy with a symmetric visibility graph $V = (A, E)$ is to pair up agents who can see each other, and then have them use the trivial two-agent strategy that we described in Chap. 1. Remarkably, this obvious strategy turns out to be optimal.

A *matching* M for a graph V is a collection of pairwise disjoint edges, where we are thinking of an edge as a two-element set. Thus, the *size* of the matching M is literally the cardinality of M ; that is, the number of edges in the matching. A vertex is said to be *covered* by the matching M if it is in the union of M , that is, if it is an endpoint of one of the edges in M . Thus, if M is a matching of maximum size for a finite symmetric visibility graph, then there is a two-color predictor ensuring that at least $|M|$ agents guess correctly. The theorem below shows that no predictor can ensure more.

First, however, we need to discuss the so-called Tutte-Berge formula for the maximum size of a matching for a graph V . William Tutte's original contribution [Tu47] was in characterizing those graphs V for which there is a matching that covers every vertex of V . His starting point was with an obvious necessary condition for such a matching: Every set $S \subseteq A$ of vertices must have at least as many points as there are odd-sized components in $A - S$. The point is that in a component of $V - S$, vertices can appear in the matching only when paired with either another element of that component, or with a vertex in S . If the component has odd size, at least one of the vertices in the component will need to be paired with an element of S . And different components require different vertices in S . Tutte showed that this necessary condition for a matching covering all vertices of V was also sufficient.

Claude Berge's generalization [B58] of Tutte's result involves the finer analysis resulting from the "deficiency" of a set S in having enough vertices to handle each of the leftover vertices in the components of odd size in $V - S$. Notationally, let $\mathcal{O}(V - S)$ denote the set of components of odd size in $V - S$, and let $\text{def}(S)$, the *deficiency* of S , be given by $\text{def}(S) = |\mathcal{O}(V - S)| - |S|$.

It now follows from what we've said that for every set $S \subseteq V$, every matching will leave at least $\text{def}(S)$ vertices not covered. Thus, if M is a matching of maximum size, then for every set $S \subseteq V$, we can cover at most $V - \text{def}(S)$ vertices, and so we must have $|M| \leq \frac{1}{2}(|V| - \text{def}(S))$. Berge's contribution was to show that equality always holds for some $S \subseteq V$. This is the Tutte-Berge formula.

With this at hand, we can now establish the following from [BHKL08].

Theorem 2.4.1. *Let V be any finite graph and consider the corresponding two-color hat problem with symmetric visibility given by V . Let M be a matching of maximum size for V . Then there is a predictor ensuring $|M|$ correct guesses, and there is no predictor ensuring $|M| + 1$ correct guesses.*

Proof. The predictor ensuring $|M|$ correct guesses is the one described in the first paragraph of this section. What must be shown is that no predictor P can do better. So let S be a set of agents (vertices) as in the Tutte-Berge formula, wherein $|M| \leq \frac{1}{2}(|V| - \text{def}(S))$. Let W_1, \dots, W_j denote the components of odd size in $V - S$, and let $Y = (V - S) - (W_1 \cup \dots \cup W_j)$. Thus A is the disjoint union of S , Y , and the W_i s. We begin by placing hats on S arbitrarily. Any agent in W_i sees only other agents in W_i or agents in S . Because hats have been placed on agents in S , we can regard the predictor P as operating on W_1 alone, and by Theorem 2.3.1 we can place hats so as to make at most $\frac{1}{2}(|W_1| - 1)$ of the agents in W_1 guess correctly. We do this for each W_i . Finally, we place hats on the agents in Y so that at most half of them guess correctly. It now follows that the total number of agents guessing correctly for this hat coloring is at most

$$\begin{aligned} &|S| + \frac{1}{2}(|W_1| - 1) + \dots \\ &+ \frac{1}{2}(|W_j| - 1) + \frac{1}{2}|Y| = |S| + \frac{1}{2}(|W_1| + \dots + |W_j| - j + |Y|) \end{aligned}$$

$$\begin{aligned}
&= |S| + \frac{1}{2}(|V - S| - j) \\
&= \frac{1}{2}|S| + \frac{1}{2}|V - S| + \frac{1}{2}|S| - \frac{1}{2}j \\
&= \frac{1}{2}|V| + \frac{1}{2}|S| - \frac{1}{2}j \\
&= \frac{1}{2}(|V| - \text{def}(S)) \\
&= |M|. \quad \square
\end{aligned}$$

2.5 A Variable Number of Hat Colors

Suppose we have finitely many agents, each of whom can see all of the others. When does there exist a minimal predictor? If the same set of colors is used for each agent, we already know: there is a minimal predictor iff there are at least as many agents as colors. But what if different agents have different sizes of sets from which their hat colors are drawn?

Suppose then that $p = \{0, \dots, p-1\}$ is our set of agents and that $c_0, \dots, c_{p-1} \in \omega - \{0\}$. We will assume that agent i 's hat color will be in the set $c_i = \{0, \dots, c_i - 1\}$, and we will let the tuple $(c_0, c_1, \dots, c_{p-1})$ encode this problem. Our first observation is that if $\sum_i 1/c_i < 1$, then there is no minimal predictor because, as in Lemma 2.2.1, the average number of correct guesses will be less than one. A natural question here is whether or not the converse holds. That is, if $\sum_i 1/c_i \geq 1$, must there be a minimal predictor? The following theorem shows that the answer is yes.

Theorem 2.5.1. *Let $p \in \omega$ and $c_0, \dots, c_{p-1} \in \omega - \{0\}$, and consider the hat problem in which p is the set of agents and the set of colorings is $\{f \in {}^p\omega : (\forall i \in p)(f(i) \in c_i)\}$; that is, agent i has c_i possible hat colors. Let $r = \sum_i 1/c_i$. Provided no agent sees himself, the average number of correct guesses will be r , regardless of the predictor. If every agent sees every other agent, then there is a predictor under which the number of correct guesses is always $\lfloor r \rfloor$ or $\lceil r \rceil$. In particular, a predictor ensuring at least one correct guess exists iff $r \geq 1$.*

Proof. The fact that the average number of correct guesses will be r is again by the same argument as in Lemma 2.2.1.

Suppose now that every agent sees every other agent. We define the predictor P as follows. Colorings can be seen as elements of the group $C = \mathbf{Z}_{c_0} \oplus \dots \oplus \mathbf{Z}_{c_{p-1}}$ in the obvious fashion. Define $\pi : \mathbf{R} \rightarrow \mathbf{R}/\mathbf{Z}$ (as a group homomorphism) by $\pi(x) = x + \mathbf{Z}$ (that is, we are projecting modulo \mathbf{Z}). Define $\varphi : C \rightarrow \mathbf{R}/\mathbf{Z}$ by

$$\varphi(f) = \pi\left(\sum_{k \in p} \frac{f(k)}{c_k}\right).$$

Define $d_k \in \mathbf{R}$ by $\sum_{j < k} 1/c_j$, let $\hat{I}_k = [d_k, d_k + 1/c_k) \subseteq \mathbf{R}$, and let $I_k = \pi[\hat{I}_k]$. The intervals \hat{I}_k lie end-to-end, and have total length r , so when we project to \mathbf{R}/\mathbf{Z} , each point in \mathbf{R}/\mathbf{Z} occurs in $\lfloor r \rfloor$ or $\lceil r \rceil$ of the I_k .

For our predictor, agent k assumes that the coloring f satisfies $\varphi(f) \in I_k$ and guesses accordingly. (This is a well-defined predictor: agent k knows the value of $\varphi(f)$ up to a multiple of $1/c_k$, and exactly one of these multiples would put $\varphi(f)$ in I_k , because I_k is left-closed right-open with length $1/c_k$.) Now, agent k will guess correctly iff $\varphi(f) \in I_k$, and as observed above, this occurs for $\lfloor r \rfloor$ or $\lceil r \rceil$ values of k .

When $r \geq 1$, of course, the number of correct guesses under this predictor is at least one, so we have a minimal predictor; when $r < 1$, the average number of correct guesses is less than one, so there can be no minimal predictor. \square

2.6 Variations on the Standard Hat Problem

In addition to the kind of hat problem we are considering, several interesting variants have arisen over the years. We'll mention two of these here.

The first variant is the following. Ten prisoners are lined up facing forward, and red and green hats are placed on their heads. Each prisoner will be asked to make a verbal guess as to the color of his hat, and, before guessing, each will be able to hear the guesses of those prisoners behind him as well as seeing the hats of those prisoners in front of him. If at most one guesses incorrectly, all will go free.

No strategy can ensure that the first prisoner will guess correctly, but if the first player uses his guess to announce red iff he sees an even number of red hats, then all of the others can use this signal (and the knowledge that the others are also using it) to correctly guess their hat color.

The extension of this problem and its solution to countably many agents will be given in Chap. 3. We also show there that these signaling problems are equivalent (in ZF + DC) to the kind of non-signaling problems that we are considering.

The second variant goes as follows. There are again ten prisoners, this time wearing shirts that are numbered one through ten. Each prisoner has a hat with a number on it matching the number on his shirt. The warden confiscates the hats and places them randomly in boxes numbered one through ten. One-by-one, the prisoners are called to the room and allowed to open nine of the ten boxes. If all ten prisoners find their own hats, then all go free. If any one of the ten fails, they all remain in prison. Find a strategy yielding a 90% chance that all will go free.

The solution is for prisoner i to begin by looking in box i . If he sees hat j , then he next looks in box j . And so on. The only way for any prisoner to lose using this strategy is for the placement of the hats to correspond to one of the $9!$ ways to place 10 numbers around a circle. But there are $10!$ permutations, so the chance of failure is only $9!/10! = 1/10$.

2.7 Open Questions

There is no shortage of questions that could be stated here, but we'll introduce a bit of notation in order to state the most obvious (and perhaps the most difficult). For positive integers k , n , and m , let $P_k(n, m)$ denote the collection of all directed graphs on n vertices (“ n -graphs”) for which there is a predictor ensuring that at least m agents guess correctly when there are k hat colors. The results in this section show that:

- $P_2(n, 1)$ is the collection of n -graphs with a cycle.
- $P_n(n, 1)$ is the collection of n -graphs that are complete.
- $P_k(n, \lfloor n/k \rfloor + 1)$ is the empty collection.

Question 2.7.1. Can one characterize the graphs in $P_k(n, m)$ for other values of k , n , and m ?

There are also two questions related to the theorem on bipartite graphs; the first is from [BHKL08], and the second is due to Velleman.

Question 2.7.2. Is there a bipartite graph ensuring the existence of a minimal predictor for k colors whose size is polynomial in k ?

Question 2.7.3. Given cardinals (possibly finite, treated as sets) c, m, k , with $k \leq c$, what is the smallest size of a family F of functions from m to c such that, for every subset A of m of size k , $f|_A$ is one-to-one for some $f \in F$?

There is an old saying, variously attributed to everyone from the French Minister Charles Alexandre de Calonne (1734–1802) to the singer Billie Holiday (1915–1959), that goes roughly as follows: “The difficult is done at once; the impossible takes a little longer.” More to the point, Stanislaw Ulam (1909–1984) provided the adaptation that says, “The infinite we shall do right away; the finite may take a little longer.” With this in mind, we leave the finite.

The Mathematics of Coordinated Inference

A Study of Generalized Hat Problems

Hardin, C.S.; Taylor, A.D.

2013, XI, 109 p., Hardcover

ISBN: 978-3-319-01332-9