Obervable Diversity: Cooperation, Conflict, and Stereotyping

April 4, 2002

1 Introduction

In this lecture, we focus on some of the implications of observable diversity. By observable diversity, I mean differences that cannot be hidden. For people, differences in race, gender, height, and in some cases ethnicity are observable, but differences in preferences, ideology, or religion need not be. Their revelation is a matter of choice. The difference is important. People can often avoid discrimination against revealed characteristics. If vegans are being stoned to death by carnivores, a vegan can claim to be a carnivore. In fact, she can even become a carnivore in order to prove that she is one.

In this lecture, we want to analyze how observable diversity leads to stereotyping, discrimination, differential rates of cooperation, and even conflict. Much of what we will say is speculative. There is a growing literature on observable diversity, but it is fair to say that scholars have not fleshed out all of the implications.

In all of the models that follow, agents - be they cells, ants, people, or firms - will wear hats. These hats will vary in their color, captions, or style. Initially, we assume that the agents cannot remove the hats or change their characteristics. We begin by exploring how people can use these hats to think statistically. This model helps us grasp the hat imagery. Once we've learned to hang onto our hats, we move onto bigger issues. In doing so, we'll poke fun at some scholarship by showing that diversity of hats can make almost any theory correct for some of the people. We will open a discussion about the dimensions of observability and why some become salient.

We then turn to a model of how observable diversity can implode if similarity is used to promote coordination. This model is a variant of research by Riolo, Cohen, and Axelrod on the minimal requirements for coordination. This and other models will be simplified and placed in a common framework so that we can compare across models. We conclude with a a model borrowed from Fearon and Laitin, which shows how these hats both create and quell conflict.

2 Hats

We begin by constructing some examples in which the agents all wear different color hats. We will use these to address issues about discrimination and stereotypes.

2.1 Statistical Discrimination

To begin with, let's suppose that there are three colors: red, blue, and green. Now let's suppose that there are three activities that people can do in their spare time: study, play soccer, or smoke cigarettes. Suppose that in our society there are 100 people of wearing each color hat and that they break down across activities as follows:

Color	*Study*	*Smoke*	*Soccer*
Red	40	40	20
Blue	50	10	40
Green	10	50	40

If you wanted to arrange a game of soccer, it would be easier to find people if you concentrated on those wearing blue or green hats, but in doing so you would discriminate against people wearing red hats who like to play soccer. We might stereotype the reds as being not into soccer. What we see from this extremely simple example is that discrimination need not stem from not liking people of a certain hat color. It can result from minimizing effort. Not discriminating often means extra costs.

2.2 Great Books

Our next example pokes fun at the world of ideas. Let's suppose that there are only two colors of hats: red and blue. And that people can either be socially or individually oriented. Again, we'll assume that there are 100 people wearing each color hat.

Color	*Individual*	*Social*
Red	50	50
Blue	50	50

Now, suppose that I write a book entitle *Social Reds and Individual Blues* which purports to reveal that reds are more social than blues. If reds and blues socially organize into pairs of one of each color, then 25% of all pairs will consist of a socially oriented red and an individually oriented blue. My book will strike a chord with these people and might well become a best seller. However, it's also that case that 75% of all pairs will not fit my theory. So what, I just say that it doesn't apply to everyone. This is why academic research relies on hard data, but there are examples of famous books for which there is scant aggregate evidence in support of the claims even though the book resonated with lots of people.

2.3 Sorting

Our third example looks at biases that result from sorting effects. This will be an example of how people can have different interpretations of causality based upon biased samples that result from sorting. The idea is due to Peter Streufert. Again, all hats are either red or blue. Suppose that income depends on hat color. If you have a blue hat, you have a 40% chance of being rich, a 25% chance of being middle class and a 10% chance of being poor. If you have a red hat, you have only a 10% chance of being rich, a 25% chance of being middle class and a 40% chance of being poor. If there are one hundred people each with red and blue hats and if people sort into Rich, Middle Class and Poor neighborhoods, we get the following neighborhood compositions.

Color	*Rich*	*Middle Class*	*Poor*
Red	40	25	10
Blue	10	25	40

Now, what do the people in these town think? Middle class people tend to think that there is no difference between having a red hat or a blue hat. Both are equally represented in their town. But, people in the Rich neighborhood and in the Poor neighborhood see a huge advantage to having a blue hat. To them, having a blue hat seems to be much more beneficial than it really is.

2.4 Sorting and Discrimination

We now want to combine our first example with our third example. Suppose that we put the 300 people from our first example into two towns based on some other hat attribute.

Town A

Color	*Study*	*Smoke*	*Soccer*
Red	20	20	20
Blue	10	10	10
Green	10	10	10

Town B

Color	*Study*	*Smoke*	*Soccer*
Red	20	20	0
Blue	40	0	30
Green	0	40	30

Town A is perfectly balanced in terms of hat colors and activities. We would expect almost no stereotyping. However, there is a cost of separating off this idyllic distribution because now we have no red soccer player, no blue smokers, and no green studiers in Town B. Thus, it is unlikely that people in town B will not stereotype reds as non athletes and blues as goody goodies, and greens as jocks or partiers. What we see in this example is how sorting could exacerbate stereotyping.

3 The Cats in the Hats

I call this the "Cats in the Hats" model because the agents, who are sufficiently unsophisticated that we can think of them as felines, identify with other cats who have similar hats. Each hat is a blue baseball cap with two numbers on it separated by a comma. The first number is the agent's *type*. The second number is her *tolerance*. I assume that both numbers are integers from 1 to 100.

These cats reside in a complex ecosystem filled with exotic vegetation. As a result, these cats wind up covered with burrs, dirt, etc .... At the end of each day, the cats return to their cat house to assigned rooms, some of which are empty. The rooms are arranged in a giant circle so that each cat has exactly two neighboring rooms. Each evening, each cat has the opportunity to help clean his neighbors. If he helps to clean a neighboring cat, it lowers his immunity by one unit. But, if another cat helps to clean him, then the cat's own immunity increases by three units. Thus, a cat that cleans two cats and is cleaned by two cats increases his immunity by three units a day. The decision to clean neighboring cats depends upon those numbers of the cats' hats. If the neighbor's number and the cat's own number differ by an amount less than or equal to the cat's tolerance, then the cat cleans his neighbor. Otherwise, he does not.

To complete the model, we need to link these immunity levels with reproduction. To accomplish this, we assume that each cat begins with an immunity of ten and that immunity decreases by one each period in addition to whatever immunity change results from cleaning. If a cat's immunity falls to zero, it dies. Cat's also dies randomly with a very small probability. Immunity values have a maximum of twenty, and if cat's immunity reaches twenty, it can produce an offspring (the offspring has a number and a tolerance within two of the parent) so long as there is room in the building. The offspring lands in the nearest open room to the right in the circular building.

Runs of this model depend on the initial conditions. Here is just one example in a 100 room building.

Example 1: 50 cats each with numbers between 1 and 100 and tolerances between 1 and 10: Initially tolerances are rather high (an average of about five). After about ten iterations, many of the cats are dead, only a lucky two or three survive and they have been good enough to reproduce. These two survivors quickly spread offspring through the population because they create neighbors just like them. Meanwhile tolerances creep up and after about 60 periods, the average tolerance is around six and the entire building is full. It is full of only cats who all have numbers near the two founders. For example, if one survivor had the number 6 and the other had the number 74, all of the cats will have numbers between 4 and 8 or between 72 and 76. By about 100 iterations, only one of these two survivor's offspring remain. One of the two types gets the upper hand and takes over the building. The tags have wiped out the diversity. The building stays full but tolerance begins to fall because being less tolerant is better. We can think of this an implicit enforcement of "pure" type. At period 500 average tolerance has fallen to about one, and the building starts to have some vacancies. By period 600, the building is only half full. After this purification of the types, the population creeps back up and hovers near full capacity. However, tolerance continues to fall. After 1000 periods, average tolerance is only 0.74 and almost all of the cats have hats with the exact same number.

What this model can tell us is that these hats enable cooperation but the cooperation may be local and therefore diversity can get wiped out. The norm of cooperating with someone of your own hat color may be easier to get started than a norm of cooperating with anyone. We take up this idea next.

4 Red Hat Green Hat

Our final model considers a two color hat model in which the hats also have numbers. The caveat is that only people with red hats can read the numbers on the red hats and only people with green hats can read the numbers on the green hats. Suppose that people play a standard prisoners' dilemma game.

Town A

	Cooper	*Defect*
Cooper	2,2	0,5
Defect	5,0	4,4

Suppose that each person plays this game against a random set of people each period. Cooperation could be sustained through a strategy of punishing people who deviate. However, if a blue deviates on a red, the reds do not know which blue deviated. Therefore, as Fearon and Laitin point out in their APSR paper, there are two options. One, the reds could punish all of the blues if a blue deviates, or two, the blues could punish a blue who deviates against a red. This all will work out fine unless people are more likely to play against their own color than the other. In this case, even if all of the reds punish all of the blues, it still may be in the interests of a blue to deviate when the probability of running into a red is low.

The model makes four interesting points. First, cooperation can be sustained despite the informational asymmetry. Two, punishing people wearing your own color hat is far more effective - you only punish one person. Three, in group policing is more robust to noise. If someone makes a mistake in an act of passion or drunkenness you don't get the whole system falling into a punishment regime. And fourth, in group policing is not without faults. The reds could pretend to punish someone who deviated against a blue and instead split the spoils.

5 Homework

Give an example of statistical discriminating due to sorting.