MS&E 336/CS 366: Computational Social Choice. Winter 2019-20
Course URL: http://www.stanford.edu/~ashishg/msande336.html.
Instructor: Ashish Goel, Stanford University.

Lecture 6, 1/29/2020. Scribed by Kaleigh Mentzer.

Schelling Segregation

Schelling segregation is a surprising result where even mild homophily leads to segregation. We begin with a grid, where agents of 2 different types reside on the grid points. The initial placement of the agents is random, and some grid points may be unoccupied. Of 4 possible neighbors (N,S,E,W), each agent prefers to have at least two neighbors be the same type as itself. If this condition is not satisfied, the agent will move, switching to an unoccupied spot selected uniformly at random from spots that do satisfy this property, if such a point exists. This process is iterated over many time steps, and over time, the two types become quite separated. A simulation can be found at:

http://nifty.stanford.edu/2014/mccown-schelling-model-segregation/

The surprising aspect of this model is that the realized segregation is much stronger than the preferences that created it. Even though agents only want at least half their neighbors to be like them, this drives behavior that creates a much higher percentage of neighbors to be the same type. As one might presume, if the preferences are increased such that agents want at least 3 neighbors to be the same, the resulting segregation is even more pronounced.

This very simple model shows how residential segregation might arise in spite of only mild preferences of the residents. Despite the simplicity of the model, which neglects heterogeneity amongst houses, preferences, etc., it describes a broader emergent pattern that is seen in practice. Schelling shared the Nobel Prize in economics in 2005, for this and other similar work. We will not discuss this model quantitiatively in this class, but this is definitely an important qualititative result to know about.

Polarization

As Lord, et. al. found in 1979, people receive opinion-confirming new information differently than opinion-conflicting information. This phenomena is well known in social psychology, and is frequently referred to as biased assimilation. It is also sometimes referred to as attitude polarization. In conjunction with homophily, biased assimilation can be used to explain polarization.

Urn Dynamics

We use urn dynamics as a simplified, stylized model of how other people's opinion impacts a voter. Consider two voters, i and j. Each voter has their own urn, containing some mixture of red and blue balls. Denote the fraction of red balls for voter i at time t to be $x_i(t)$. This fraction represents voter i's preference for red.

Consider the following two scenarios:

Scenario 1: Voter i draws a ball from voter j's urn. Vote i then adds a ball of that color to their own urn, and discards a randomly selected ball already in their urn to keep the total number of balls constant. In this scenario, voter i places a high value on voter j's opinion, regardless of if it matches their own opinion.

Scenario 2: Voter i draws 2 balls, one ball from voter j's urn and one ball from their own urn. If the color of the two drawn balls matches, voter i puts another ball of that color into their urn, and discards a random preexisting ball to keep the total number constant. If the color of the two drawn balls does not match, the drawn balls are returned to their respective urns and no further changes are made. This is an example of biased assimilation, only letting another person's opinion influence your own opinion if it matches your preexisting beliefs.

Formal Definition of Biased Assimilation

We begin by defining a notion of other people's opinions. $S_i(t)$ is the weighted average of the opinions of i's neighbors, or more formally,

$$S_i(t) = \frac{\sum_{j \neq i} w_{ij} x_j(t)}{\sum_{j \neq i} w_{ij}}$$

In the previous class, we considered the following assumptions on weights:

1. $w_{ij} \geq 0$
2. $w_{ii} > 0$
3. $\sum_j w_{ij} = 1$

In this class, we change to only having the assumption:

1. $w_{ii} > 0$

By and large, we can translate between these assumptions by normalization.

With these definitions, we can define the notion of biased assimilation.

$$x_i(t+1) = \frac{w_{ii} x_i(t) + (x(t))^b S_i(t)}{w_{ii} + (x_i(t))^b S_i(t) + (1 - x_i(t))^b (1 - S_i(t))}$$

This definition includes a bias factor, b, which allows additional control of how strongly the voter weights other's opinions relative to their own. If $b = 0$, we get DeGroot. If $b = 1$, we get the second scenario in the urn dynamics.

Two Island Network

Consider a graph with two very well connected components. Edges within these components have the weight p_{same}. These two components are sparsely connected by a smaller number of edge, each with weight p_{diff}. Assuming $p_{same} + p_{diff} = 1$, then biased assimilation in this context can be expressed as

$$x_L(t+1) = p_{same}x_L(t) + x_L(t)(p_{same}x_L(t) + p_{diff}x_R(t)$$

Further, we notice by symmetry that $x_R(t) = 1 - x_L(t)$, so we only need to keep track of a single fraction. We will address this problem in more detail in the homework, illustrating that this model leads to polarization.

Cardinal Social Choice

To this point, we have only considered ordinal rankings; that is, we only get an ordered list of preferences, but no information about the intensity of those preferences.

Consider the following example. Three different voters, A, B, and C, have the same ordinal ranking of ice cream flavors: Vanilla, Chocolate, Strawberry, then Peanut Butter. However, they all have different intensity of preferences. The following table shows the 'utility' that each voter gets from each option.

	A	B	C
Vanilla	1	1	.25
Chocolate	.75	1	.25
Strawberry	.5	1	.25
Peanut Butter	.25	0	.25

We notice that voter B is actually indifferent between the first 3 items, and voter C is indifferent between all options. Perhaps this additional information might change how we would like to apply voting rules.

To help describe the intensity of preferences, we introduce two different notions: utility and cost. Utility is how much benefit the voter gets from an option, while the cost is a measure of the voter's disagreement or "disutility" if a particular option is selected. Specifically, voter i has utility $u_{ih} \geq 0$ for alternative j, and voter i has cost $c_{ij} \geq 0$ for alternative j.

Distortion

For both cost and utility, we can define the concept of distortion. The distortion of alternative a can be written in terms of utility as

$$D(a) = \frac{\max_{j \in A} \sum_{i \in V} u_{ij}}{\sum_{i \in V} u_{ia}}$$

and in terms of cost as

$$D(a) = \frac{\sum_{i \in V} c_{ia}}{\min_{j \in A} \sum_{i \in V} c_{ij}}$$

A distortion of 1 indicates that you have found the optimal solution.

Obtaining a small distortion seems as though it might be a difficult problem. The social choice function does not know the costs or utilities, so we might not expect to be able to effectively optimize for this criteria. However, under the right set of reasonable assumptions, distortion becomes a useful notion. The following is a simple example where we can obtain a reasonable distortion.

We begin by assuming all voters and all alternatives lie on a line. We represent the alternatives as $A_j \in [0,1]$ and the voters as $v_i \in [0,1]$. The cost of an alternative for a voter is the distance between them on the line, or $c_{ij} = |A_j - v_i|$. To simplify, we say that the voters and the alternatives are the same, i.e. $V = A$. Note that this is a strong structure imposed on this problem.

One option is to choose the median voter as the winner. This is the choice that minimizes the total cost in this model, thereby minimizing the distortion as well. Furthermore, it turns out that this solution is analogous to choosing the Condorcet winner. Therefore, any Condorcet consistent rule will obtain a distortion of 1 under these assumptions.

Impossibility Results

Cardinal preferences, in the worst case, achieves the same distortion result as ordinal rankings. Intuitively, this makes sense because cardinal preferences capture the information in ordinal rankings, and perhaps additional information as well. This provides one type of limit on the behavior we can anticipate from cardinal voting mechanisms.

There are impossibility results relating to utility and cost. For utility, no deterministic algorithm can get any finite bound on the distortion. Furthermore, no randomized algorithm can do better than 1/M. For cost, there are no deterministic or randomized algorithm that can capture a finite bound.

To see the problems that may arise from the deterministic case, consider the two voters: with

V_1	V_2
A	B
B	A

the corresponding preferences Which we can see creates an unbounded resulting distortion if we

U	A	B
V_1	0	0
V_2	0	1

decided to break ties by always choosing A.

To see the problems associated with randomized algorithms, consider the profile with the corre-

V_1	V_2	$\ldots$	V_m
c_1	c_2		c_m
$\vdots$	$\vdots$		$\vdots$

sponding preferences With arbitrary m, we can see that the resulting distortion can be arbitrarily

U	c_1	c_2	$\dots$	c_m
V_1	1	0	$\dots$	0
V_2	0	0		
$\vdots$	$\vdots$		$\ddots$	
V_m	0			0

bad.

To avoid this impossibility result, we would like to impose additional structure to make the problem tractable. For utility, the typical assumptions typically are that $u_{ij} \geq 0$ and $\sum_j u_{ij} = 1$, essentially that the utilities are normalized to be of equal magnitude. For cost, we generalize the concept of voters lying on a line. The assumption is typically that there exists a metric space, some notion of distance, on $V \cup A$ such that $c_{ij} = d(V_i, A_j)$. Costs are often preferred because this mathematical formulation is often useful.

References

[1] A. Goel. Lecture 6 Slides. https://docs.google.com/presentation/d/1Uf4SxW3CqOIVPrs PpfZunGgtD95Z1GAgc5zxJglCrR4/edit#slide=id.g7d0d0f3009_0_62

[2] Polarization Dynamics Slides. https://web.stanford.edu/ ashishg/msande336_handouts/ attitude.pdf