Cailin O’Connor

Logic and Philosophy of Science

UC Irvine

Dynamics of Inequity

What is a philosopher doing here?

Philosophy of science – studying science itself and/or doing work that extends theoretical science

My topics – misinformation, scientific communities, cultural evolution, inequity

My methods – game theory, evolutionary game theory, agent-based modeling, network modeling, experimental economics/social psych

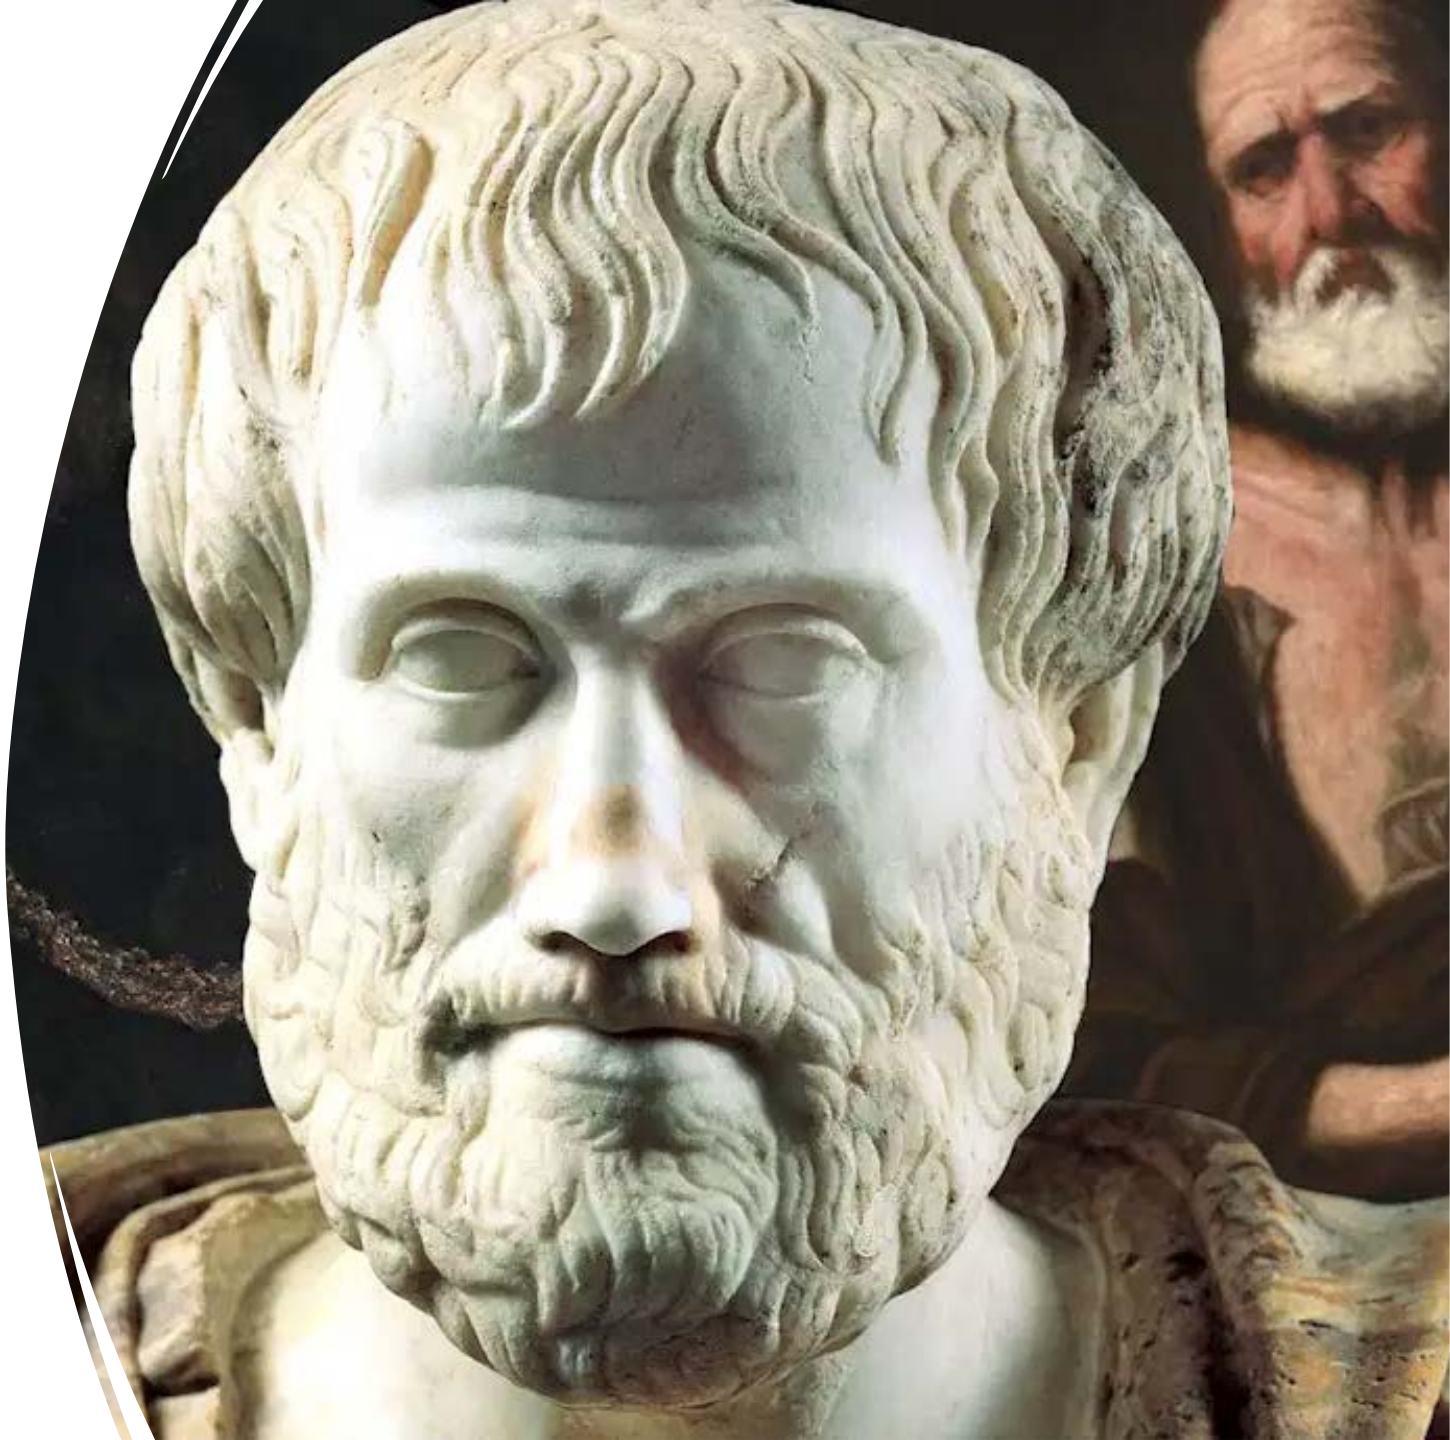

Modeling Inequity

Motivation

It is no secret that **some sorts of people tend to get more and others less**.

Many of these inequitable or unfair divisions are supported by social **conventions or norms**.

This is especially true **across social categories** (gender, race, caste etc.)

Why? **What explains the ubiquity of inequity?**

Methodology

The project uses **bargaining models from evolutionary game theory** to address this question.

Bargaining is a paradigmatic strategic interaction, and **conventions of bargaining emerge via cultural evolution.**

Game Theory and Philosophy

As we'll see, the way game theory is used in philosophy often varies from typical uses in economics, and some other fields. Main features:

1) KISS methodology
2) Focus on role of models within an argument
3) More explicit work to justify the use of models

Roadmap

1) The Nash demand game
2) The evolution of fairness
3) The evolution of unfairness
4) Power, categories, and the pull of unfairness

The Nash Demand Game

Game Theory

Players: player 1 and player 2

Strategies: any real number in [0,1]

Payoffs: Equal to demands, unless $D_1+D_2>1$, then equal to the disagreement point for each player

Information: Knowledge of game structure and opponent rationality

Infinite Equilibria

Nash equilibria – sets of strategies where no individual can deviate and improve their payoff

The Nash demand game has an **infinity of Nash equilibria** – every pair of strategies that perfectly divides the resource.

Refinements

There are many equilibrium refinements in the bargaining literature (the Nash solution, the Kalai-Smordinksy solution, etc.) – but this isn't where we're going.

Instead **we will follow a thread of literature based on empirical observations**.

The Fair Demand

In experimental settings people almost exclusively arrive at the 50-50 split.

And importantly, most cultures have **stated norms of fairness** supporting this division.

What is special about the fair split?

Roadmap

1) The Nash demand game
2) The evolution of fairness
3) The evolution of unfairness
4) Power, categories, and the pull of unfairness

Fairness Evolves

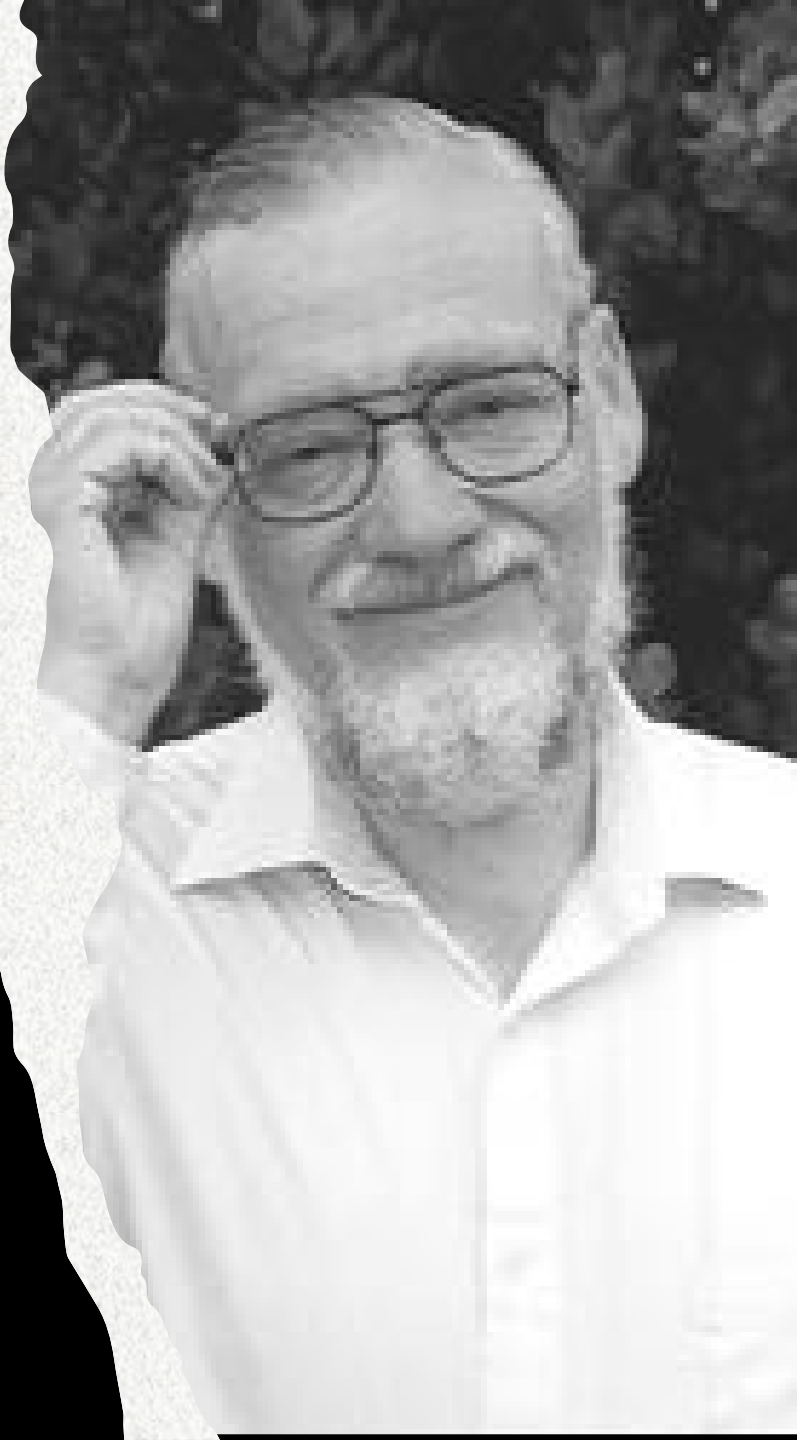

From GT to EGT

Game theoretic models use rationality assumptions to predict individual behavior in strategic scenarios.

Evolutionary game theoretic models assume a group or population repeatedly playing a game, then **evolving over time via learning or natural selection**.

Simplified NDG

		Player 2		
		Low	**Med**	**High**
Player 1	**Low**	4,4	4,5	4,6
	Med	5,4	5,5	0,0
	High	6,4	0,0	0,0

Population Structure

With a mini-game in hand, evolutionary modelers next have to decide **how a population will interact**. I.e., who do the agents in a group play this game with? How often? Under what circumstances?

Freely mixing group – agents play against every other agent with equal probability

Networks – agents are embedded in a social network which stipulates who meets whom

Positive assortment – agents meet their own strategy types with higher probability

Dynamics

Next an evolutionary model must stipulate a dynamics – **rules for how agents will change strategy over time**, based on their interactive successes and failures.

Replicator dynamics – strategies expand and contract proportional to their success (could represent natural selection, imitation, or basic learning)

Agent based learning dynamics – individual agents learn based on their payoffs (i.e., repeat successes, imitate successes of others, predict a best response)

Replicator Dynamics

Demand 4

Avg Payoff = 4

Demand 6

Avg Payoff = .2*6 = 1.2

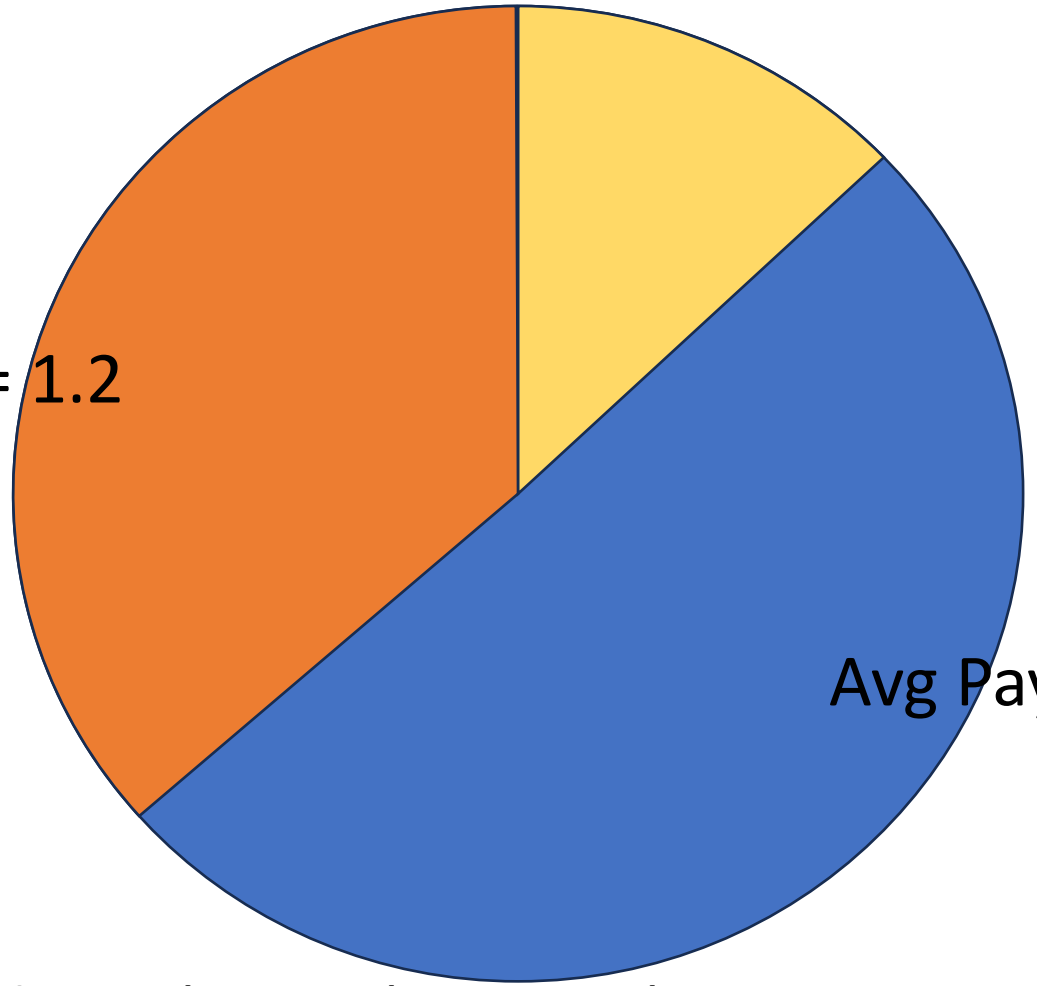

Demand 5

Avg Payoff = .5*5 + .2*5 = 3.5

Avg Population Payoff = .2*4 + .5*3.5 + .3*1.2 = 2.91

Replicator Dynamics

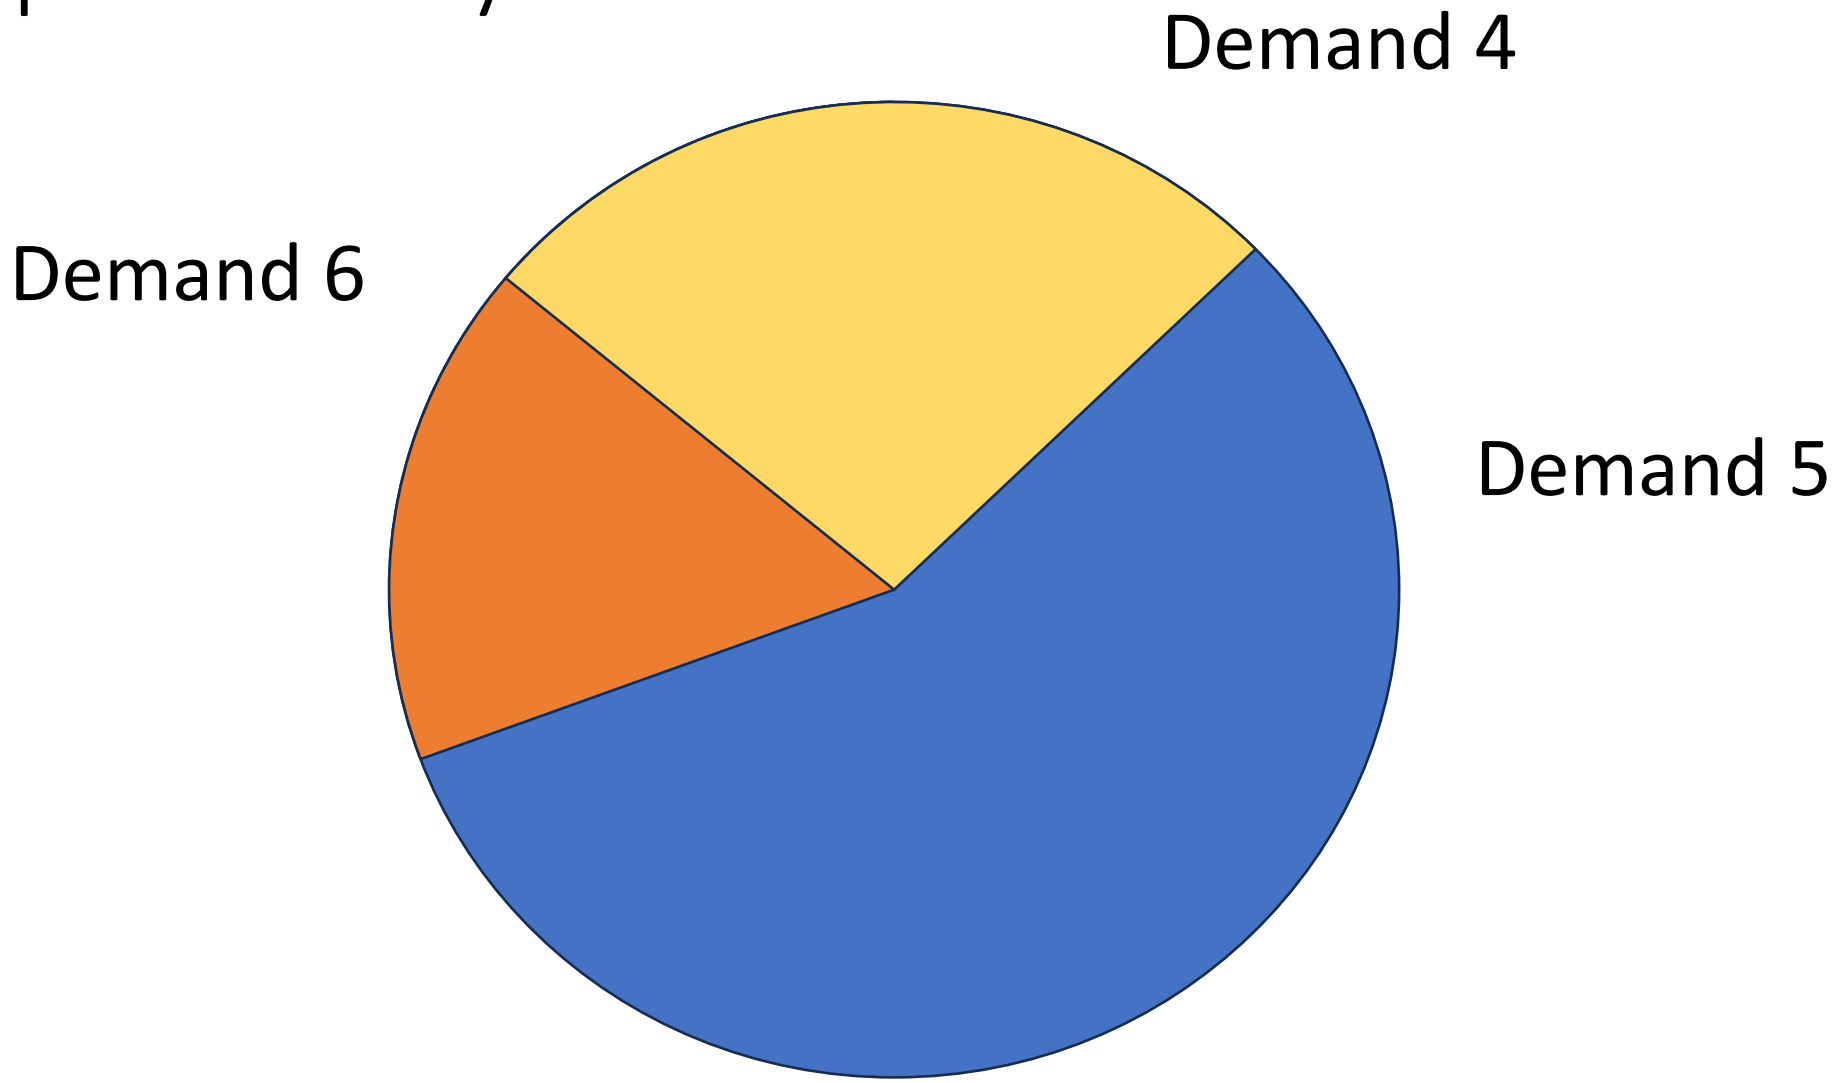

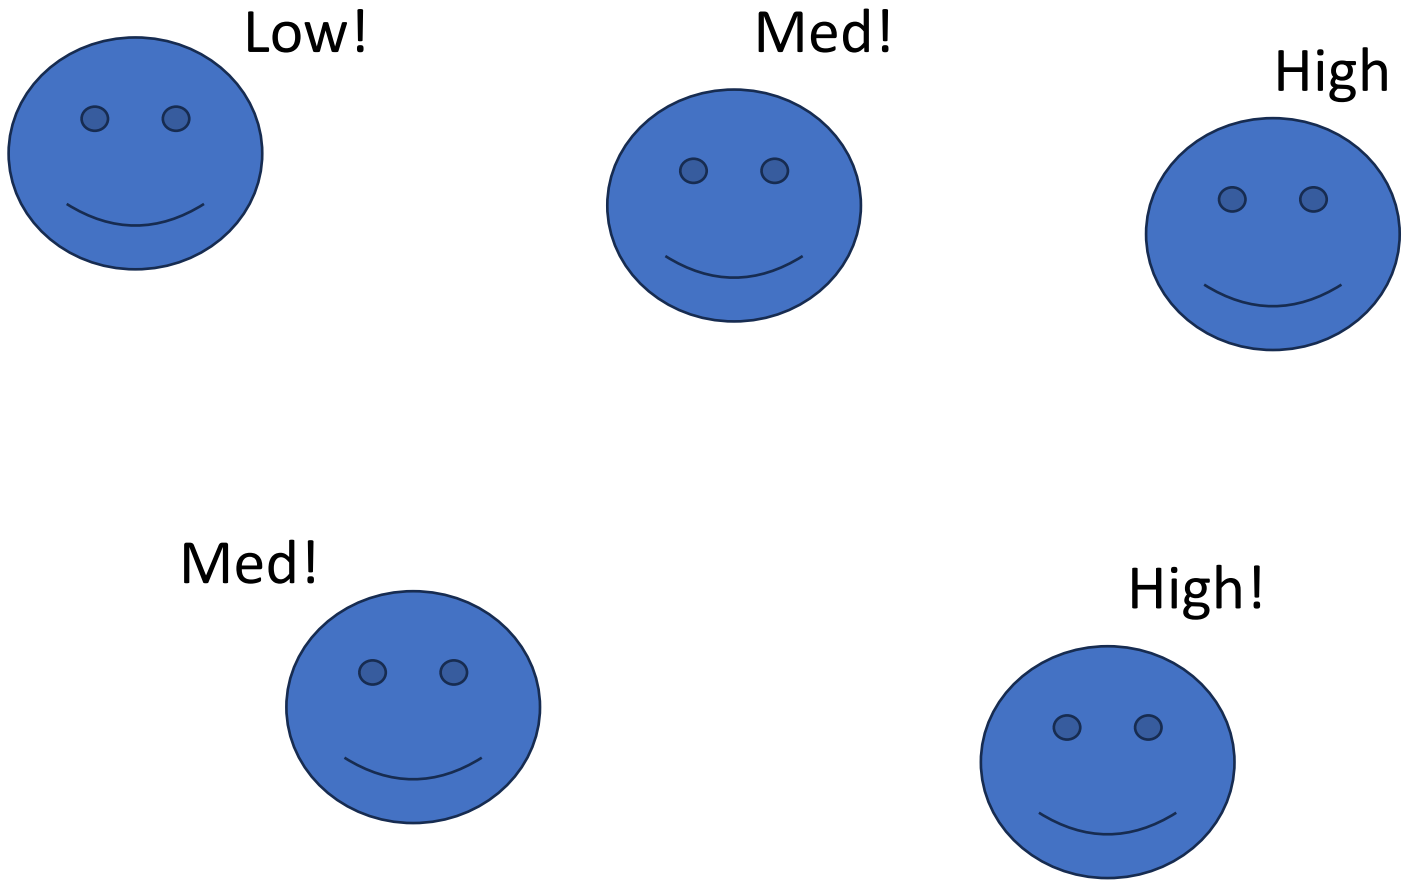
Low!
Med!
High
Med!
High!

Low!

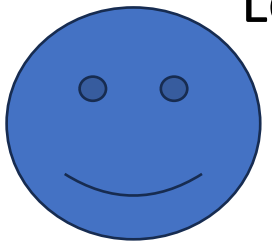

Med!

High

Med!

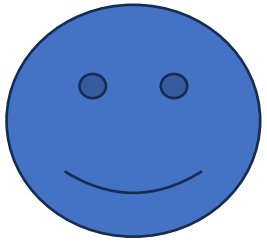

High!

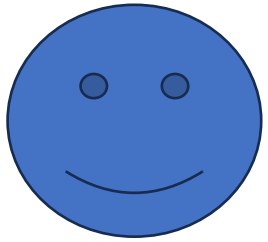

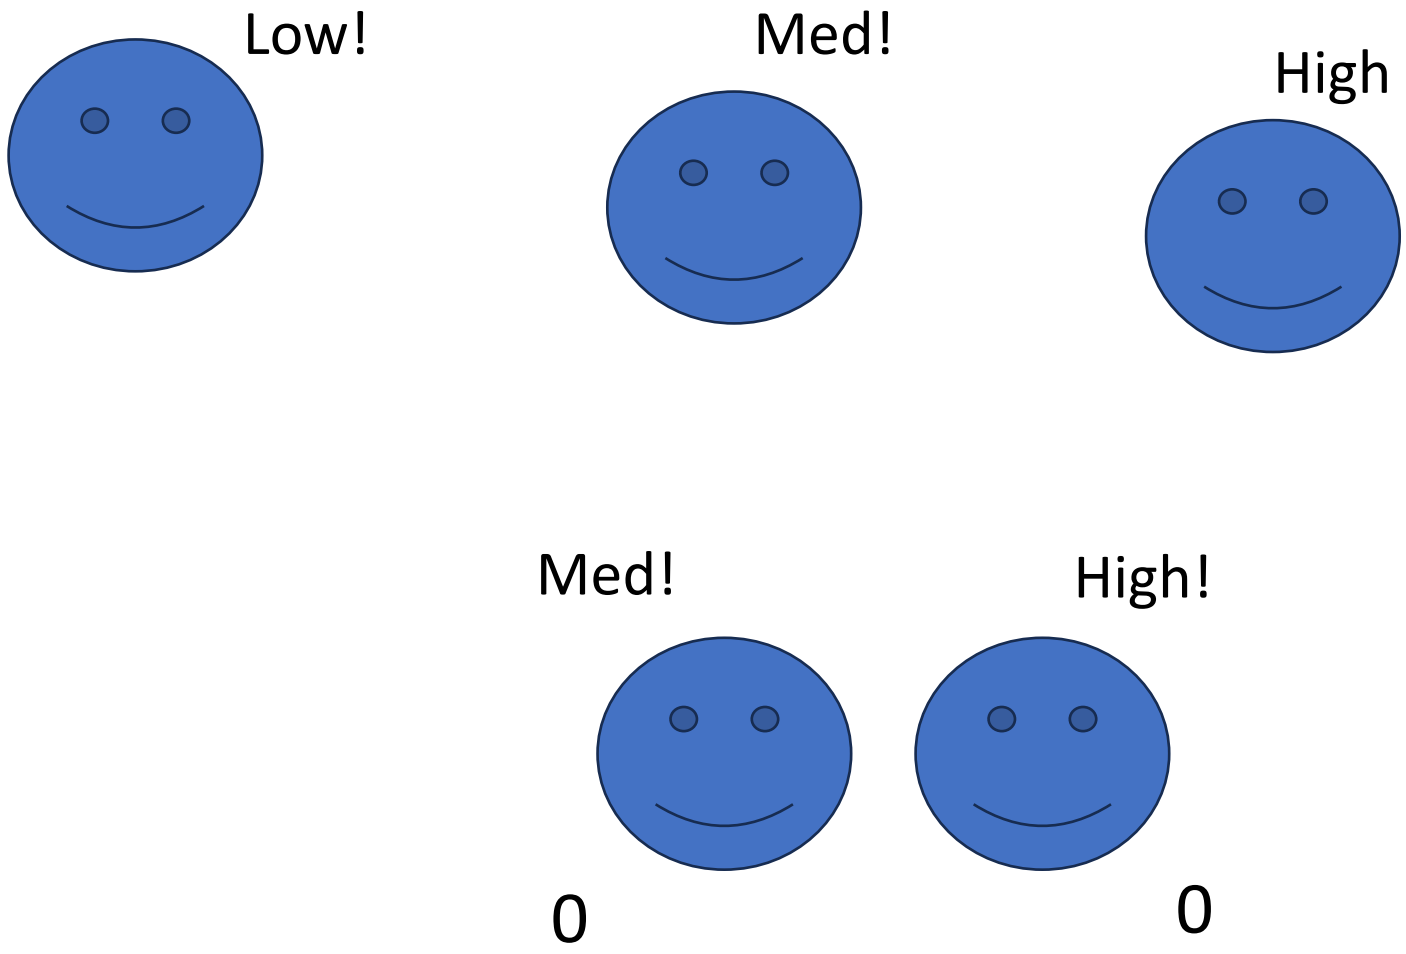
Low!
Med!
High
Med!
High!
0
0

Analysis

How do we analyze these sorts of models?

1) Find the population level equilibria (the stable end states of evolution)
2) Calculate how often the different equilibria evolve

Fairness Evolves

To summarize the typical models in this literature:

1) Fairness (a population of all medium-demands) tends to evolve with high likelihood
2) Less likely are stable “fractious” outcomes with some high and some low demanders

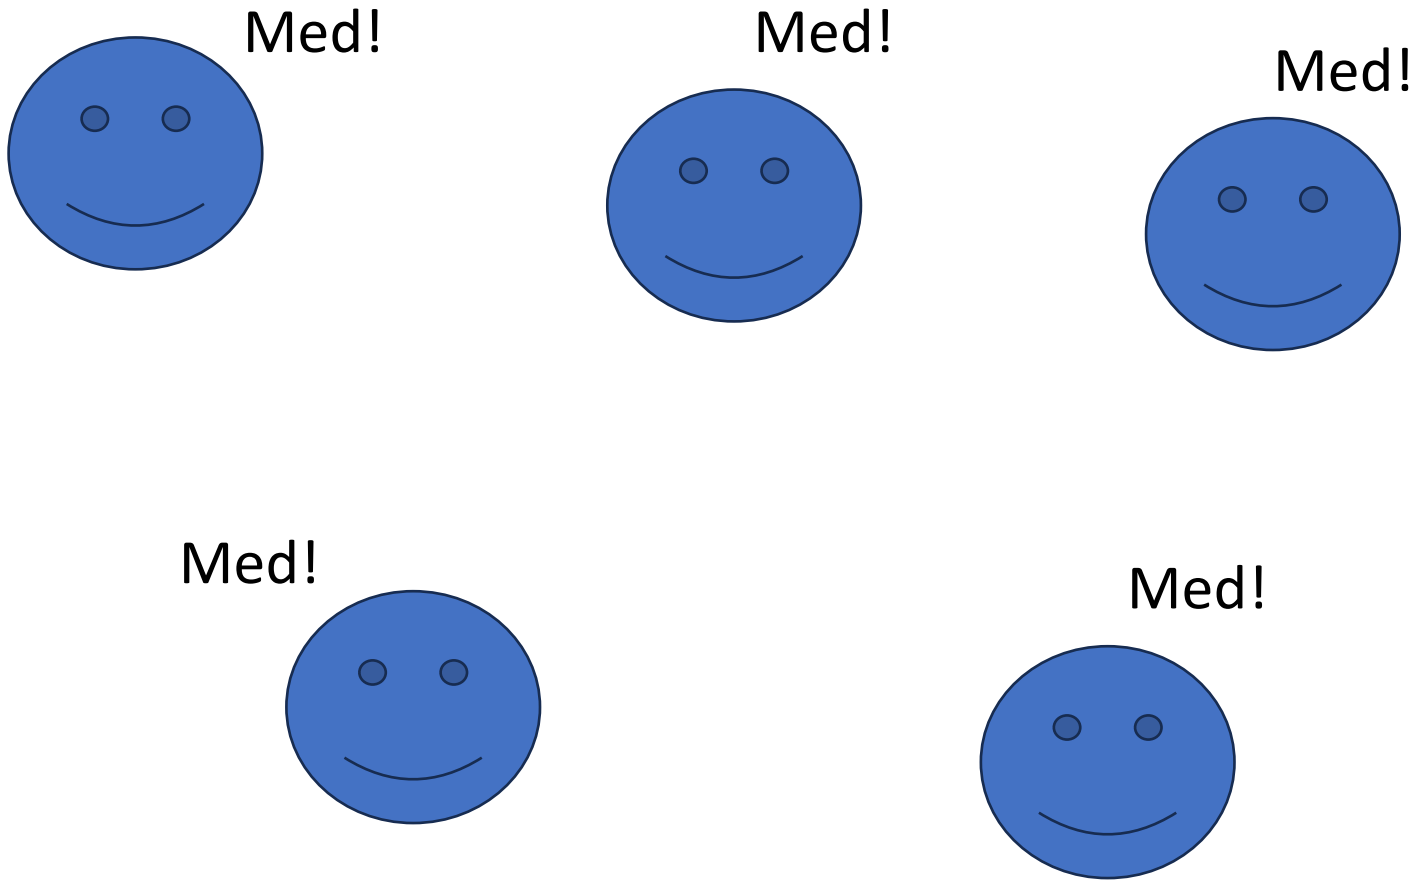
Med!
Med!
Med!
Med!
Med!

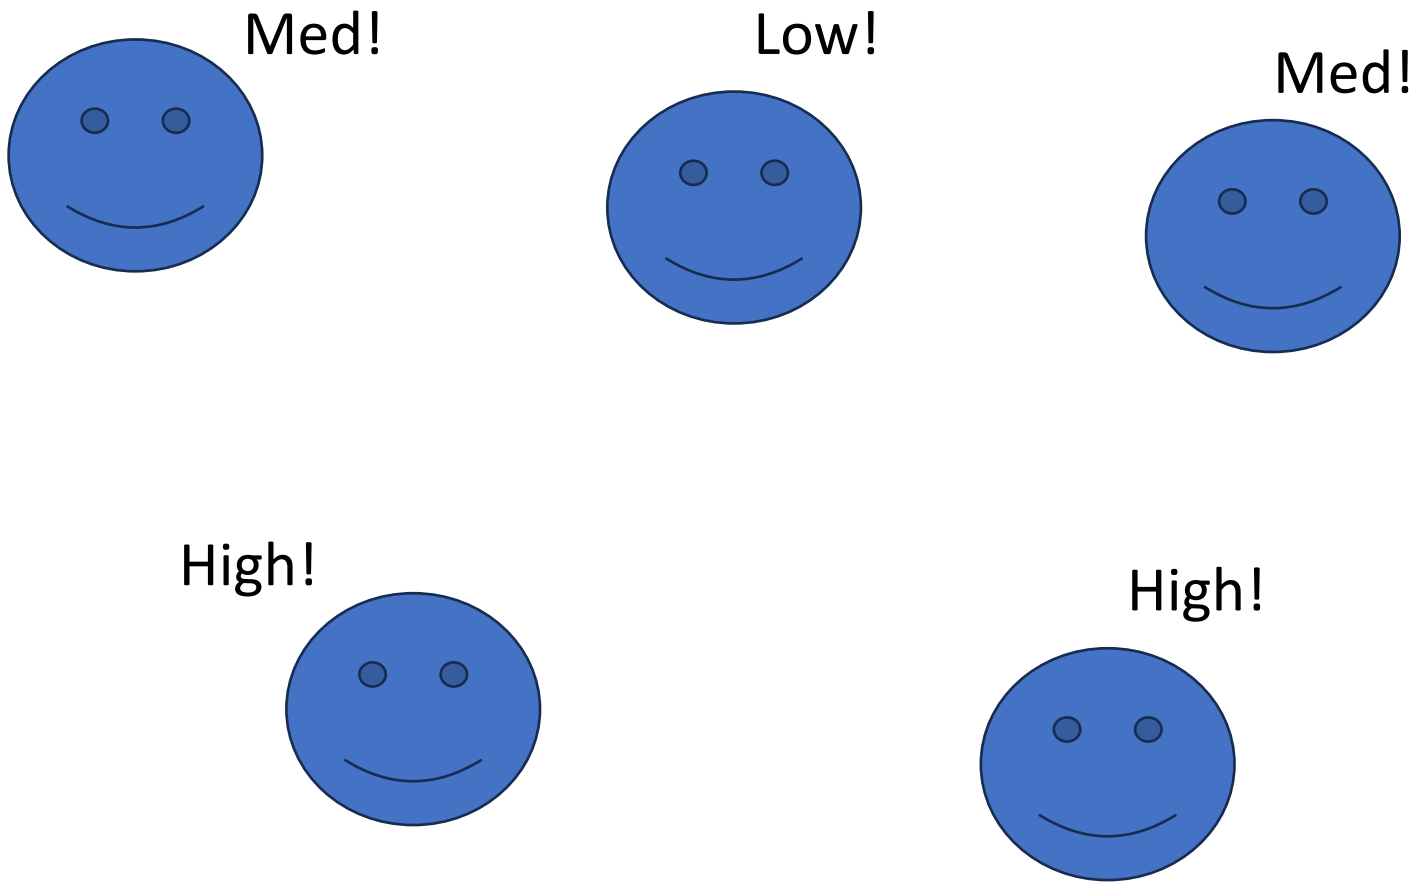
Med!
Low!
Med!
High!
High!

Explaining Fairness

The general project takes these models to help **explain the prevalence of fair bargaining norms**, and to support a picture where justice emerges naturally.

Further work outlines conditions in which fairness is not just likely, but certain, to emerge (i.e., situations where individuals are on networks or positively assort) and arguing for the relevance of these conditions to human groups.

Roadmap

1) The Nash demand game
2) The evolution of fairness
3) The evolution of unfairness
4) Power, categories, and the pull of unfairness

Bad news: things aren't always fair

I could list thousands of examples:

1) Salary/hiring biases
2) Landowner/sharecropper divisions
3) Attentional biases by teachers (and others)
4) Division of household labor
5) Sales of cars, bus tickets, etc.

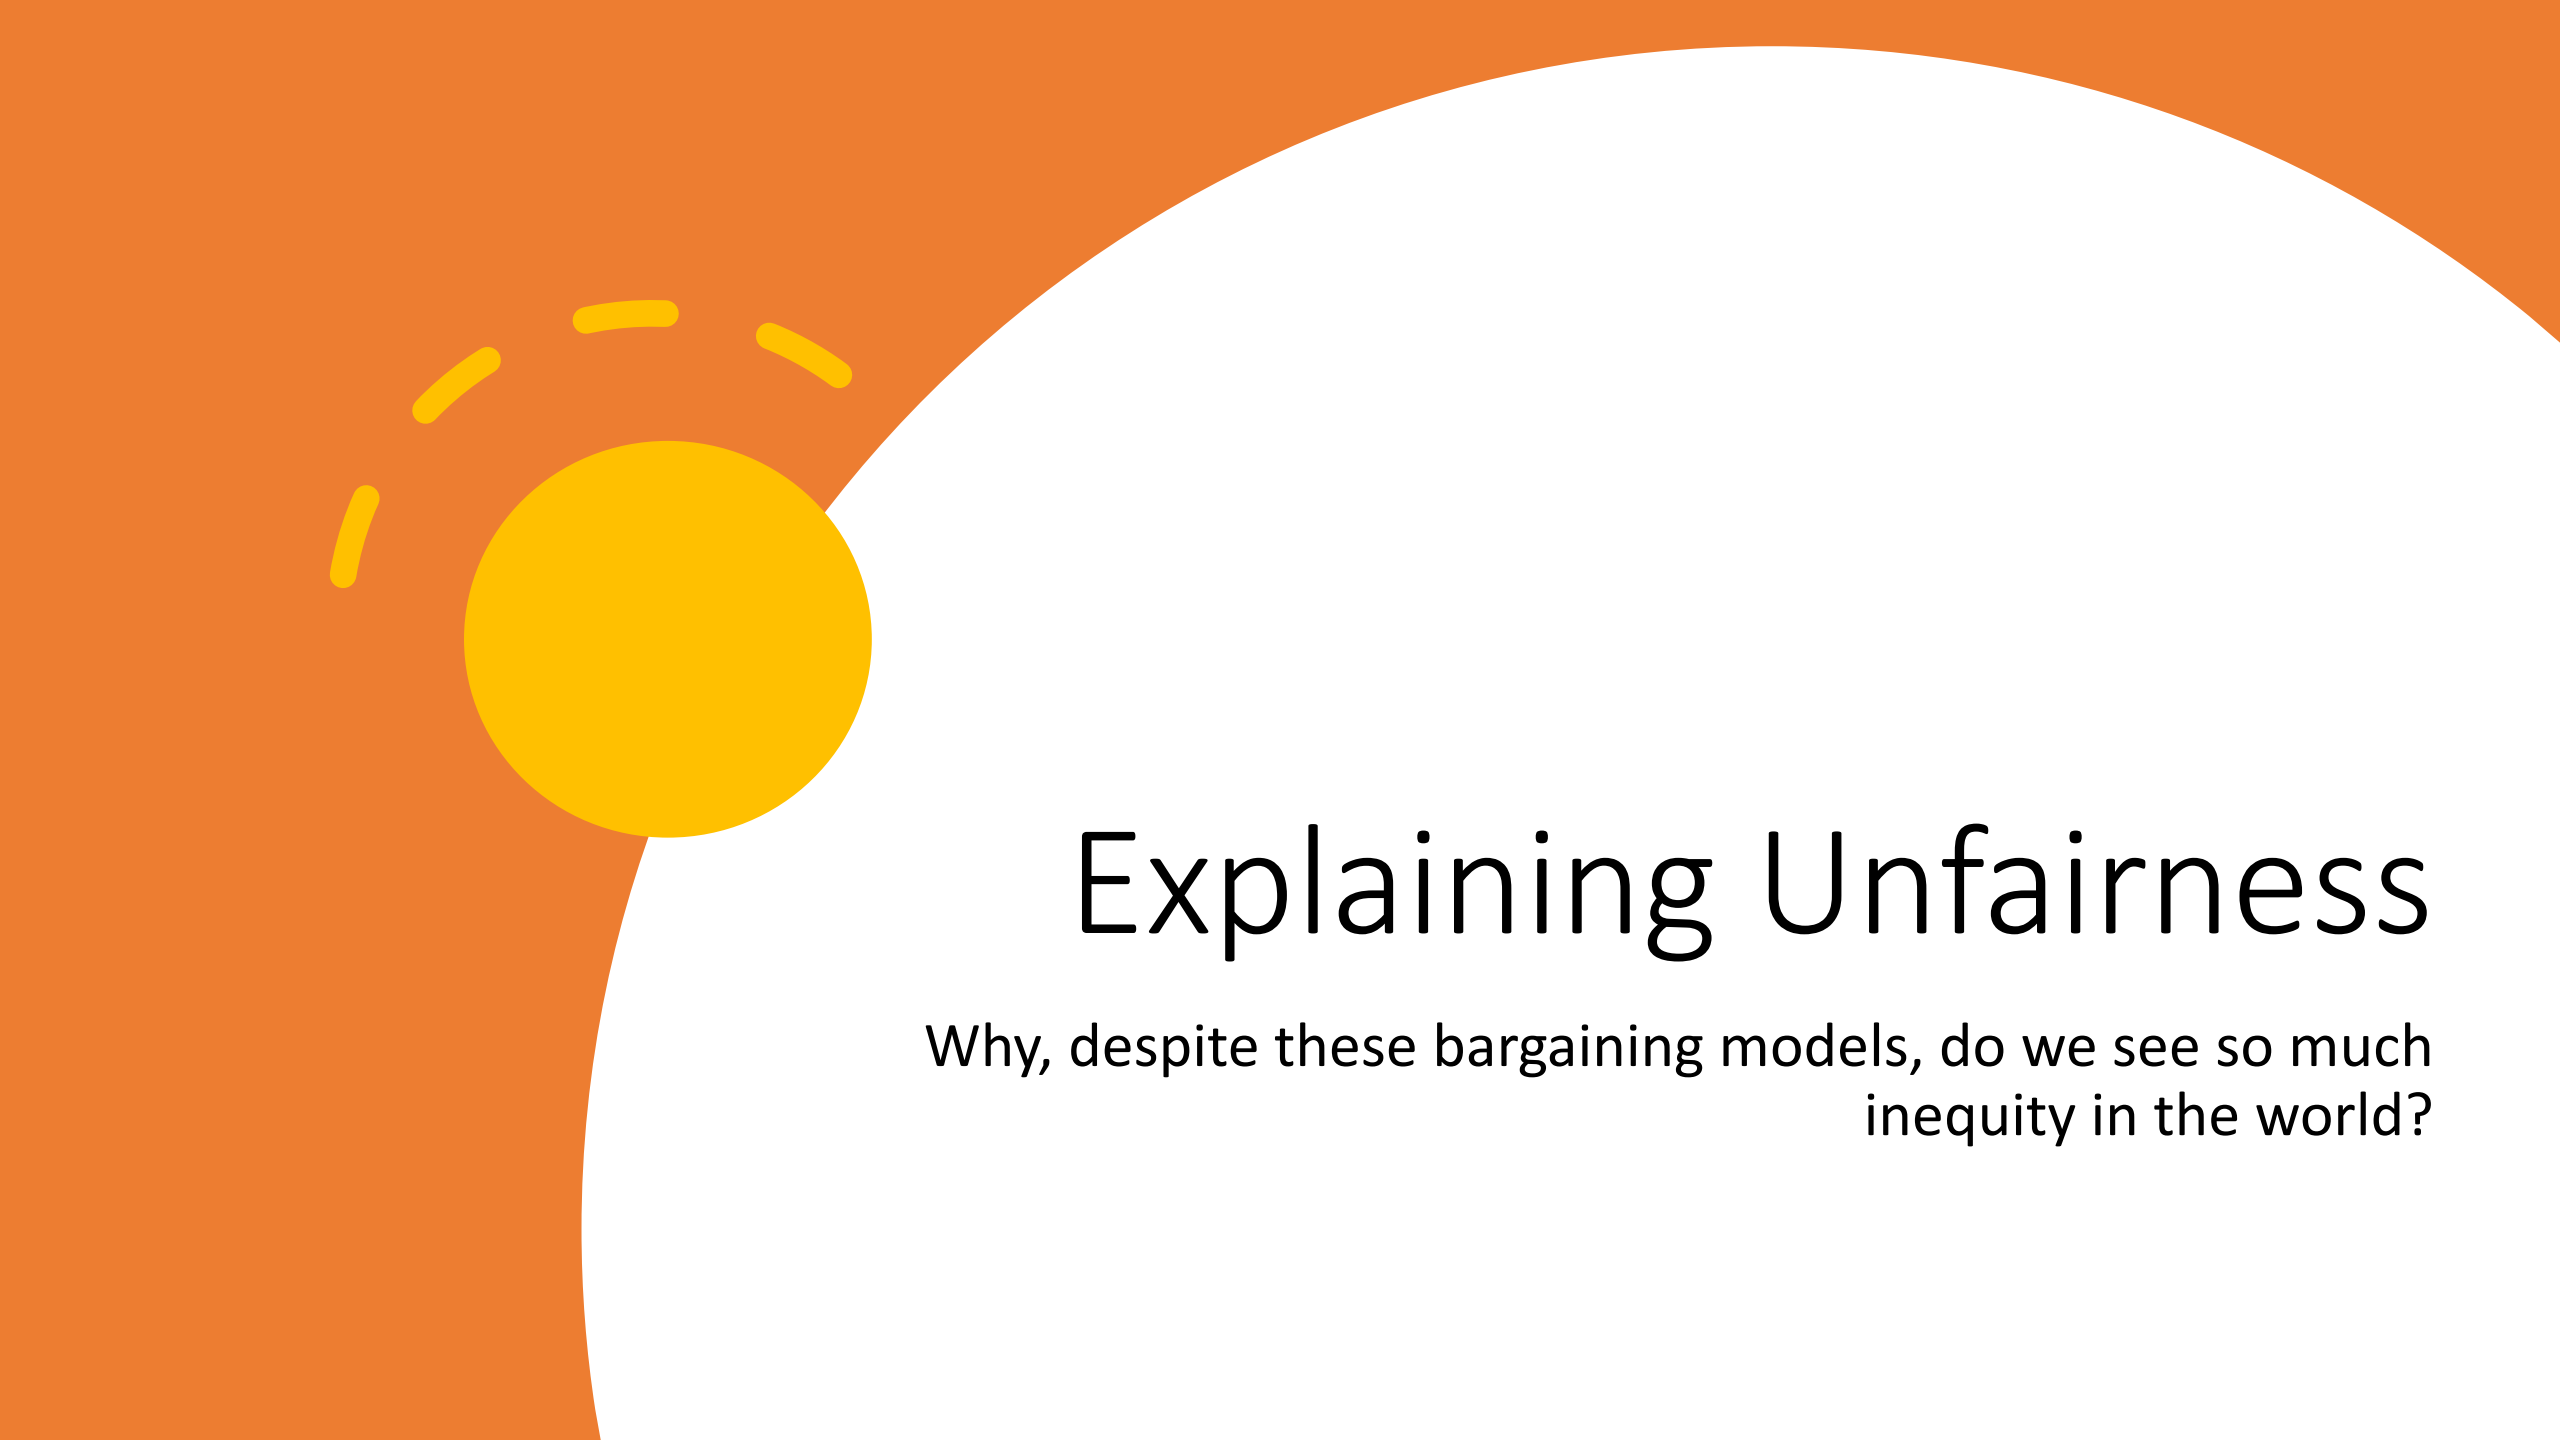
Explaining Unfairness
Why, despite these bargaining models, do we see so much inequity in the world?

Social Categories and Unfairness

Notably, many cases of inequity relate to **social categories**.

Social categories are categories that are widely recognized within a society, and which members of that society use to guide their behaviors towards others.

Primary social categories are those that are most important for shaping interpersonal behavior.

Across societies, **gender** and **age** always constitute primary social categories, and commonly **race**, **class**, **caste**, and **religious affiliation** do as well.

Modeling Social Categories

As noted, in most evolutionary models, you have a single, unitary group of actors who interact and evolve.

In order to capture the evolution of behavior with social categories, one needs a model with **multiple types or tags**.

Norms for Un-Fairness Evolve

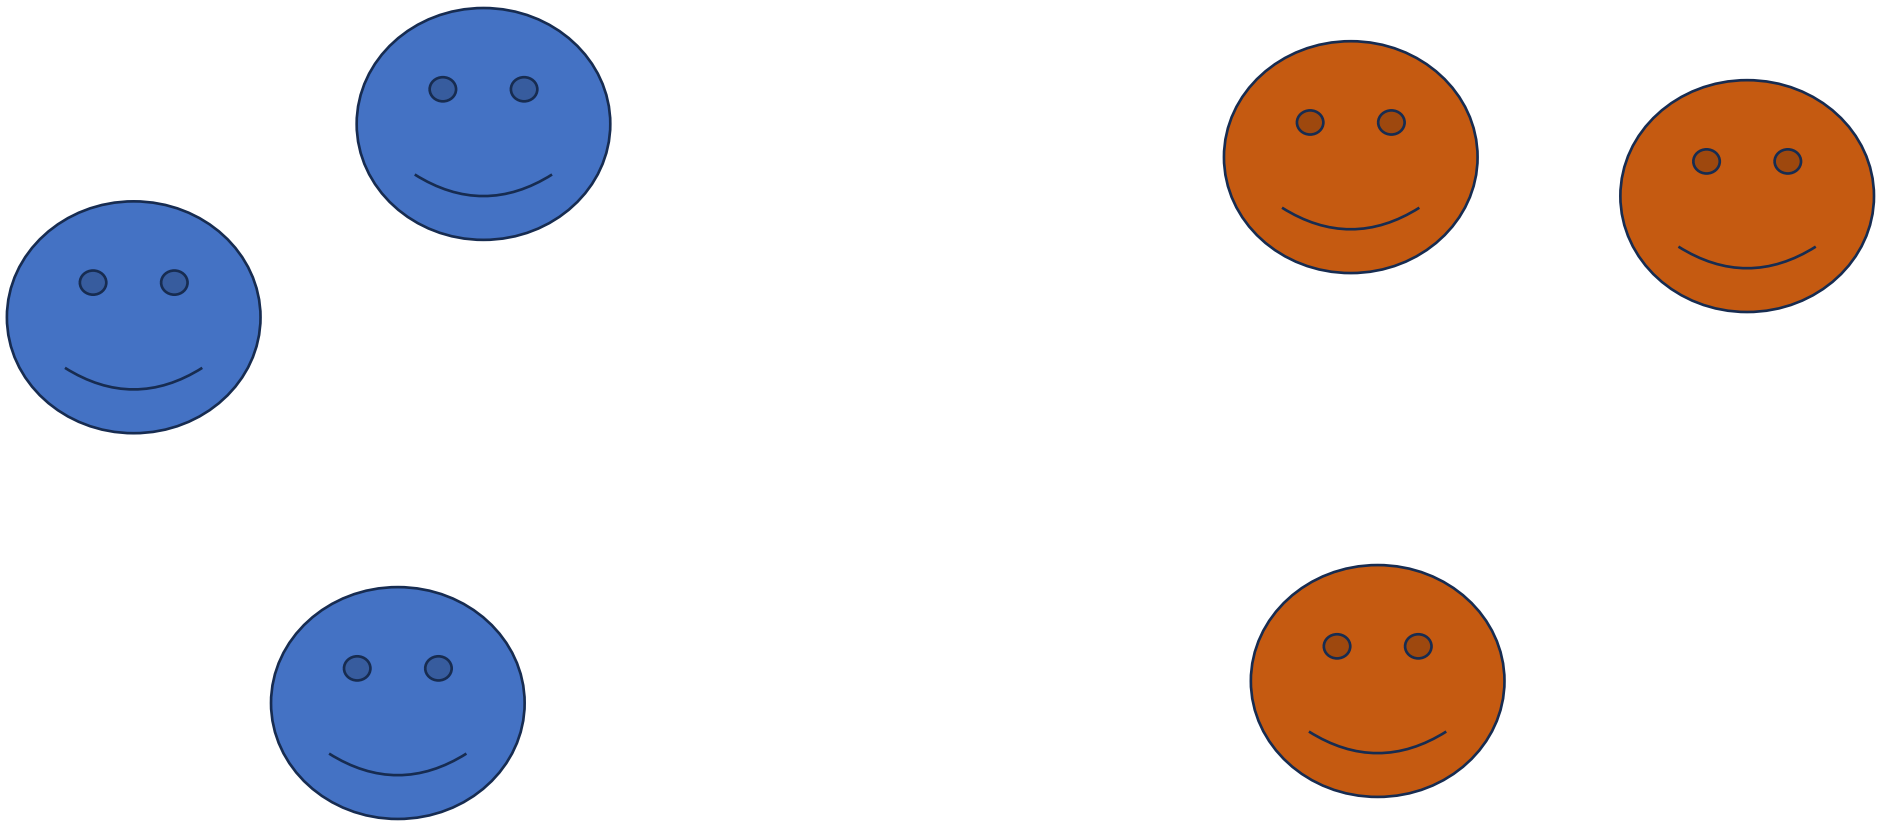

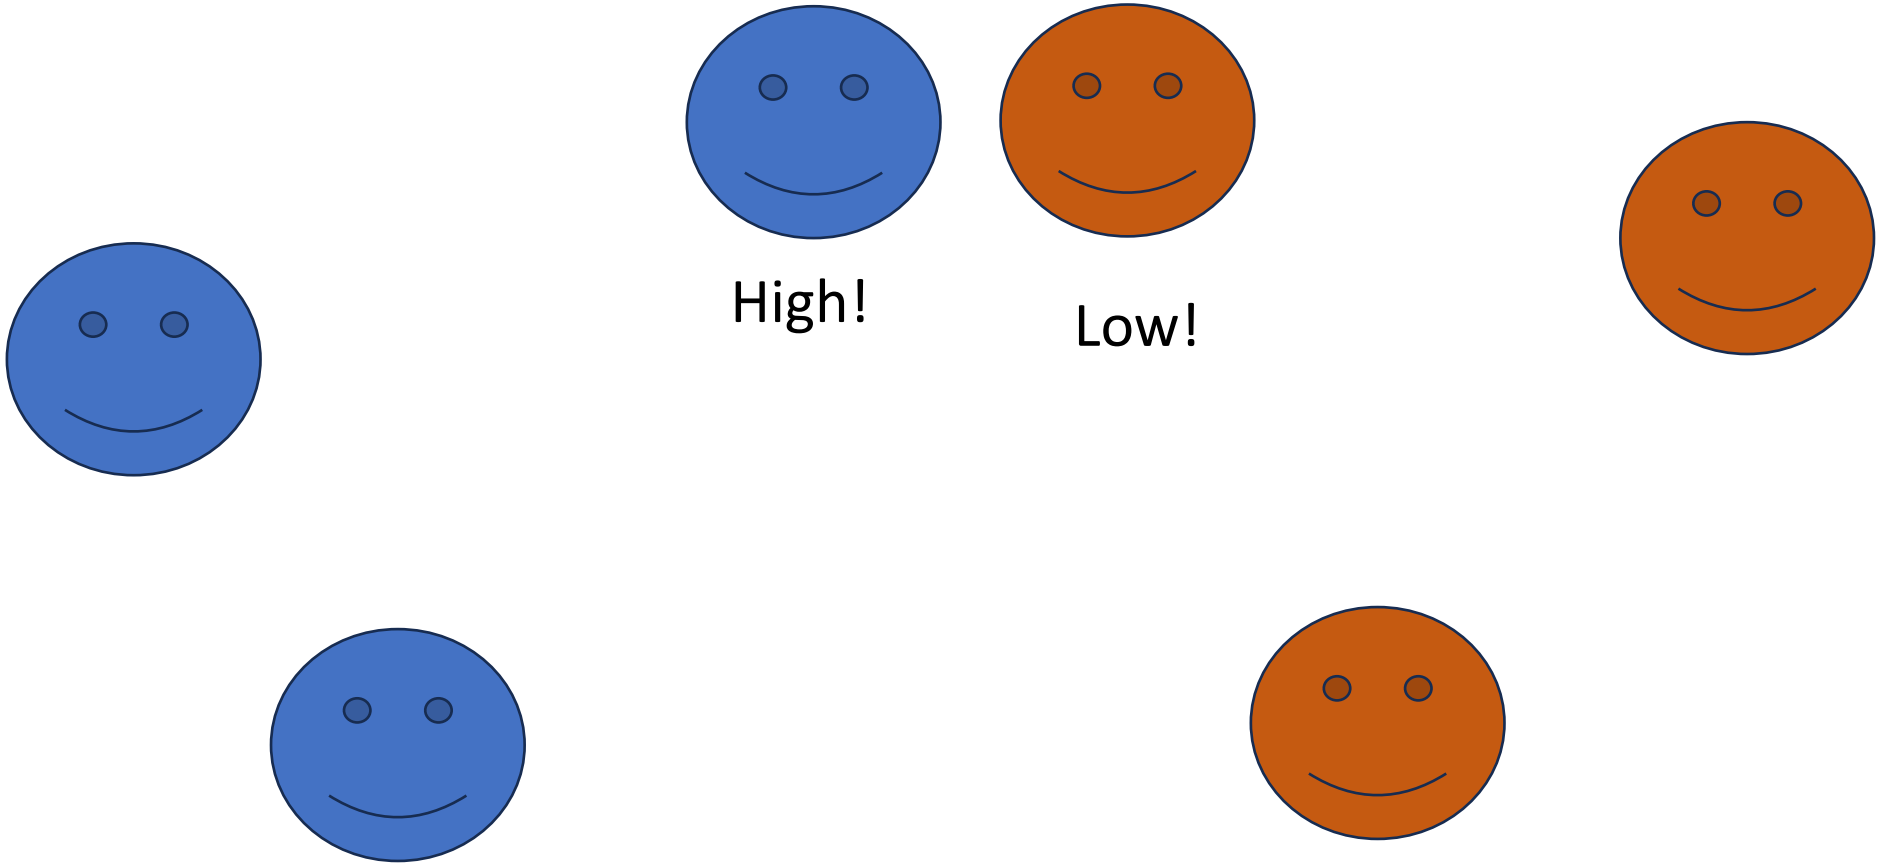
High!
Low!

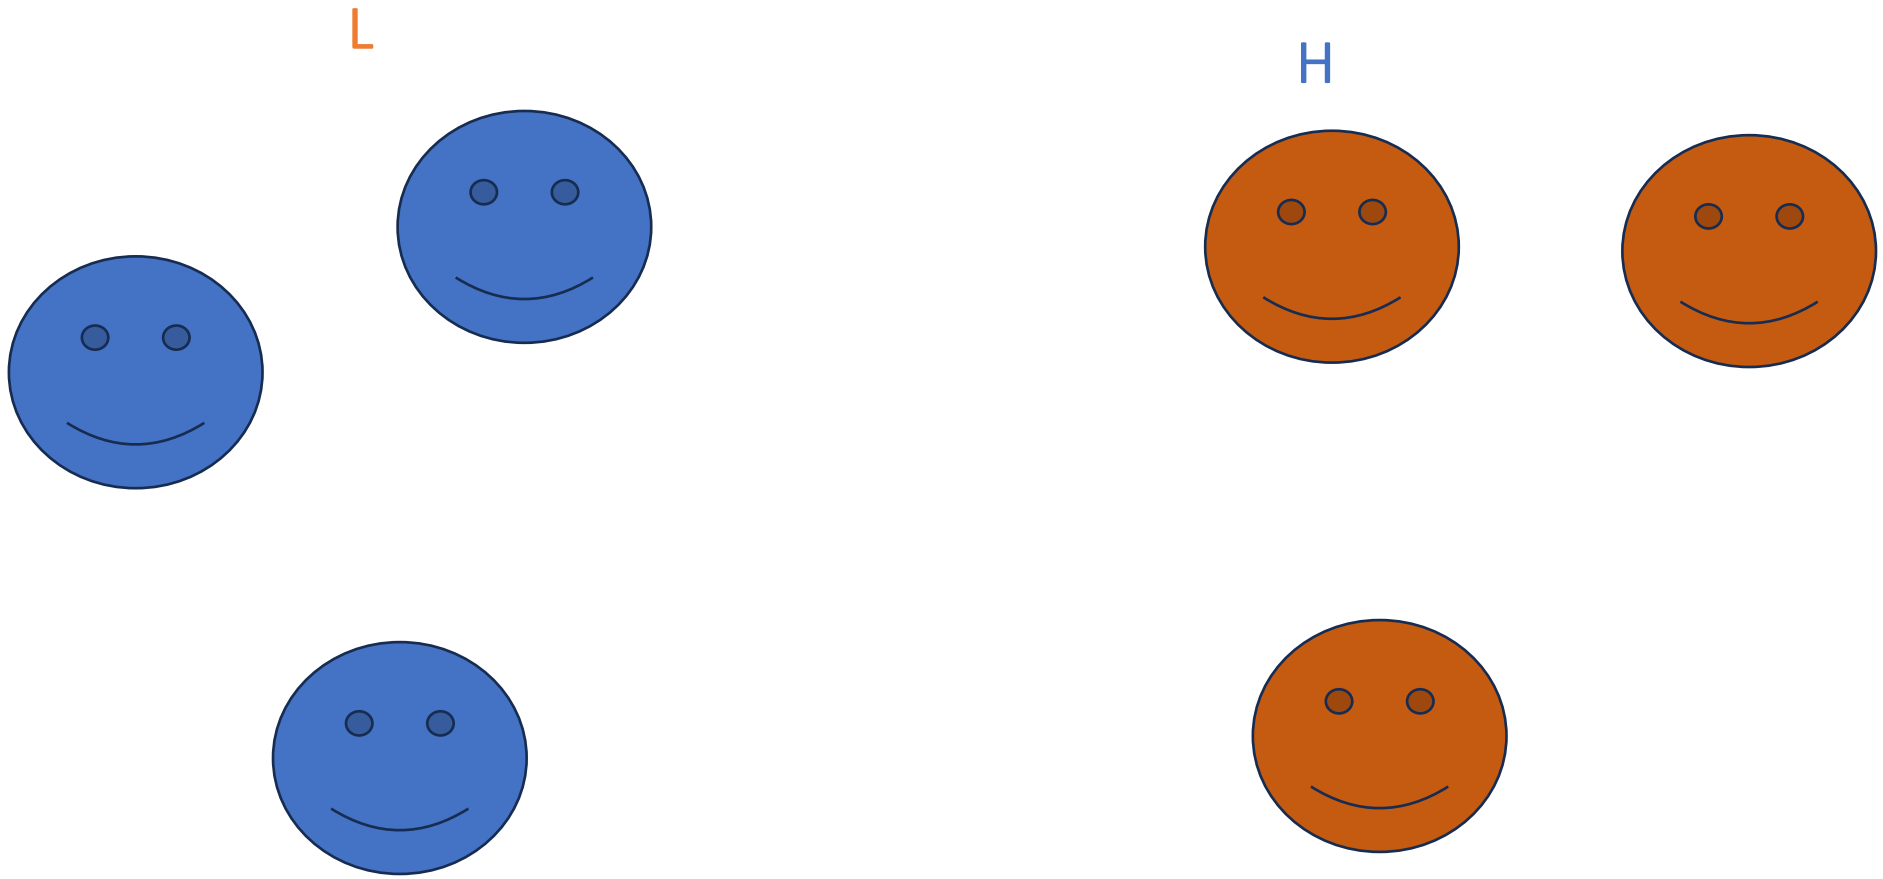
L
H

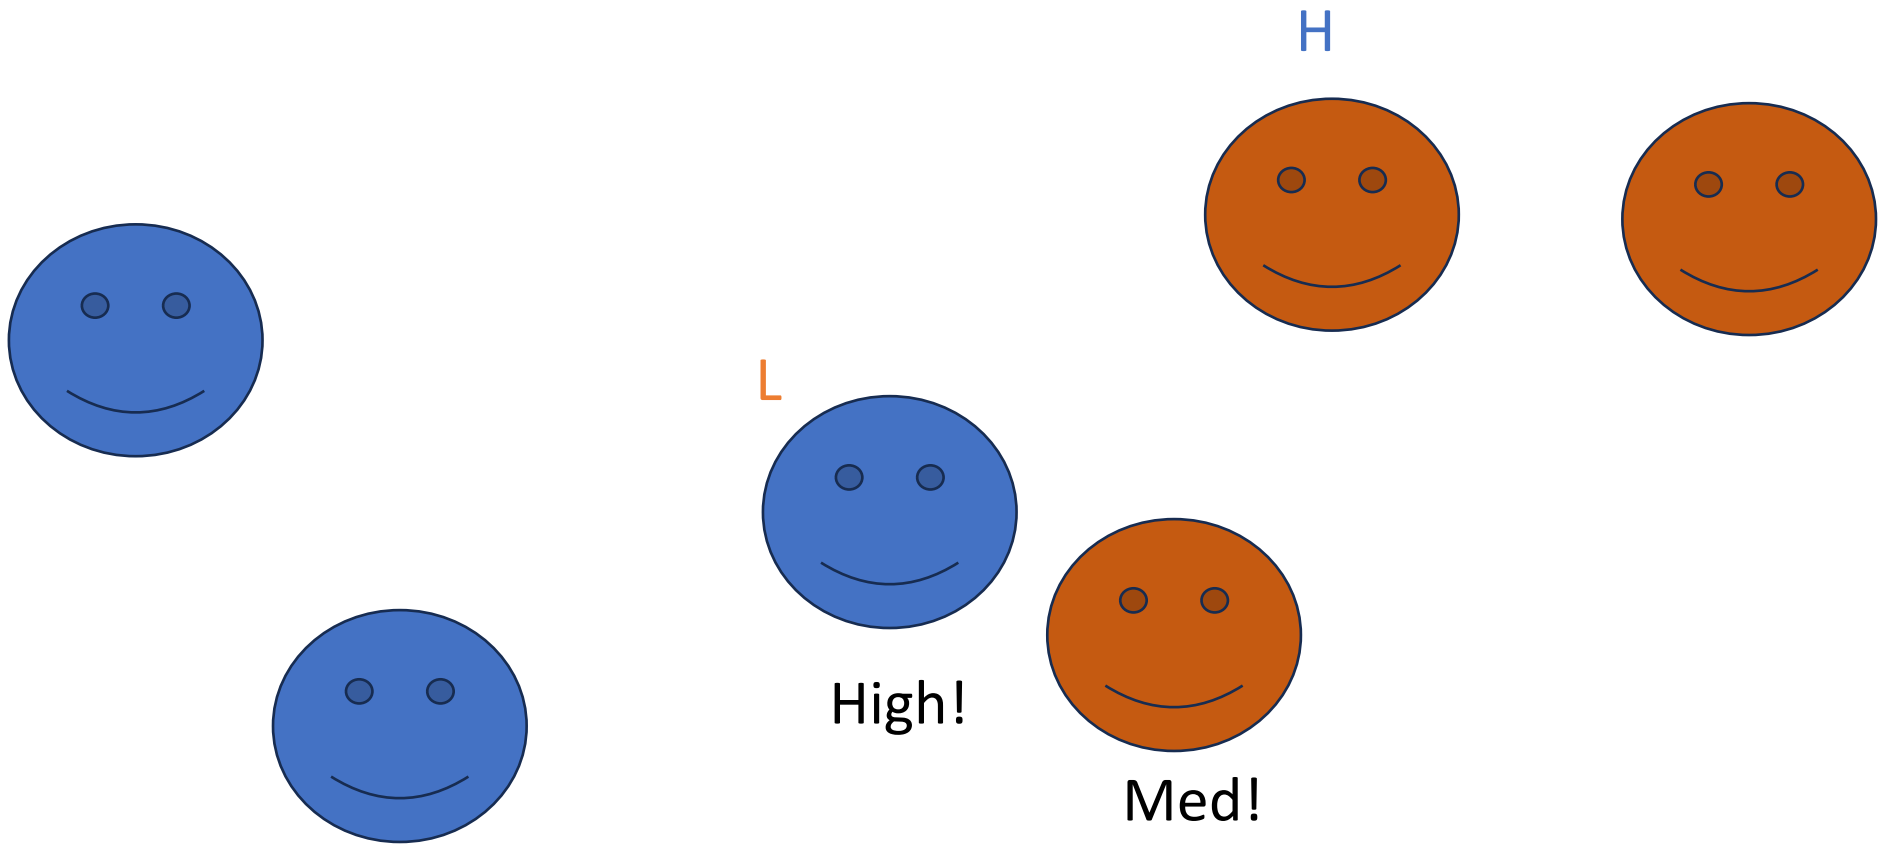
H
L
High!
Med!

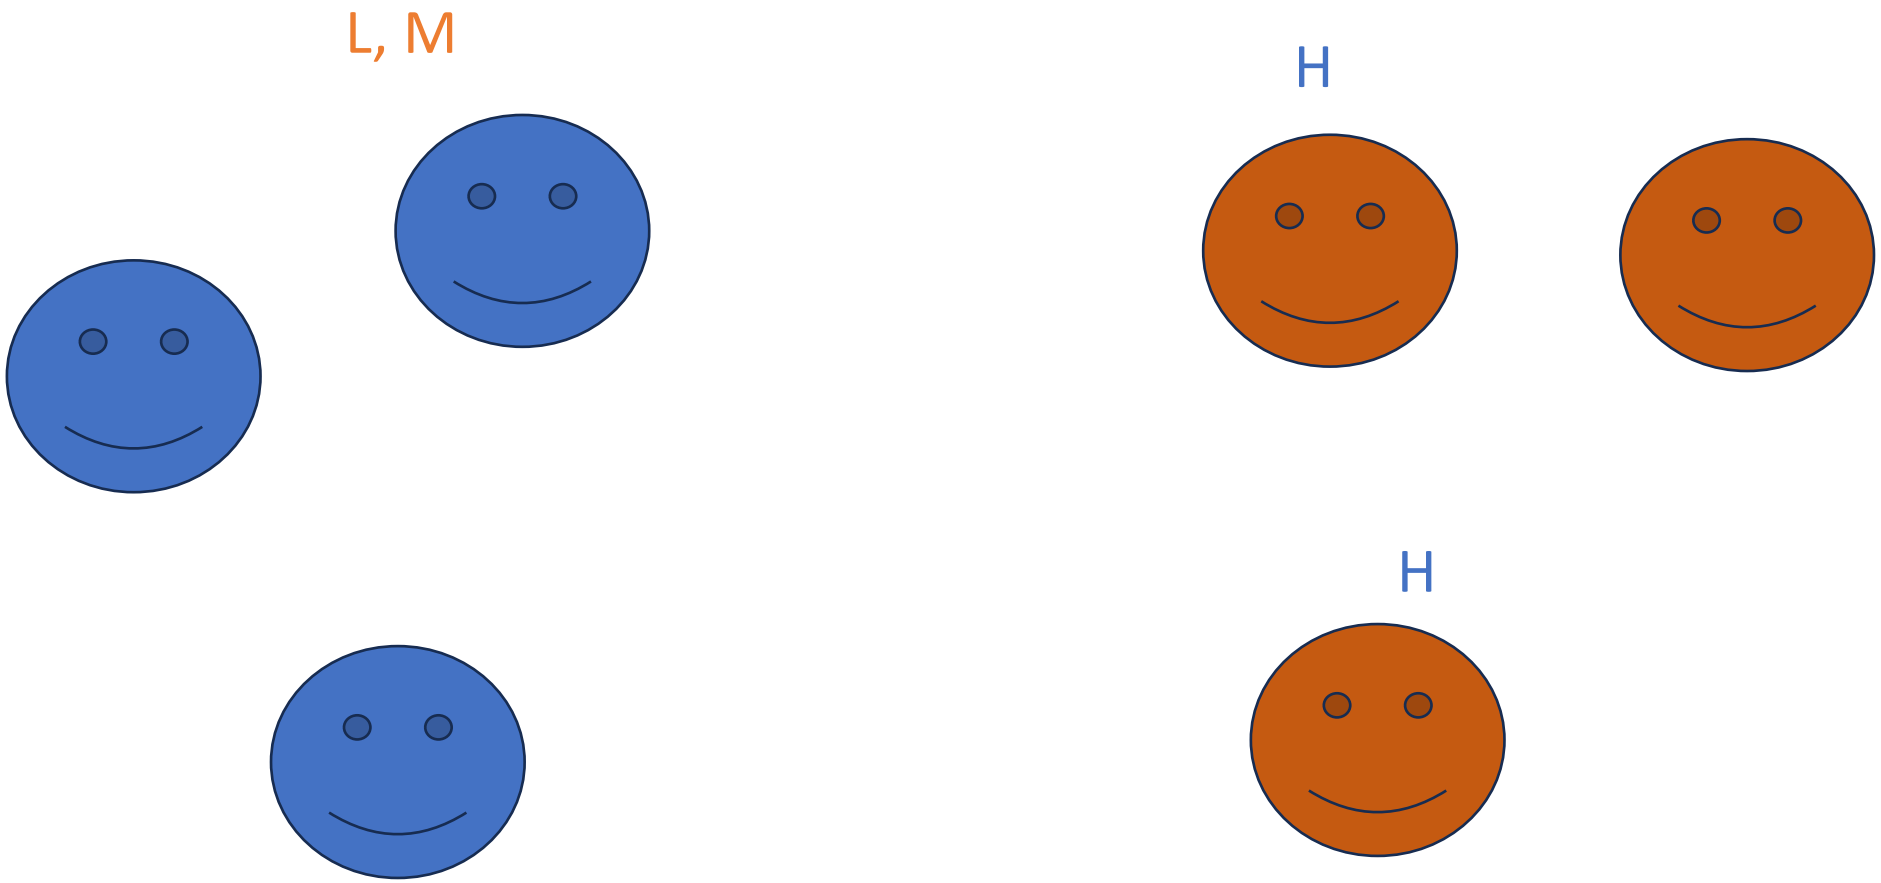
L, M
H
H

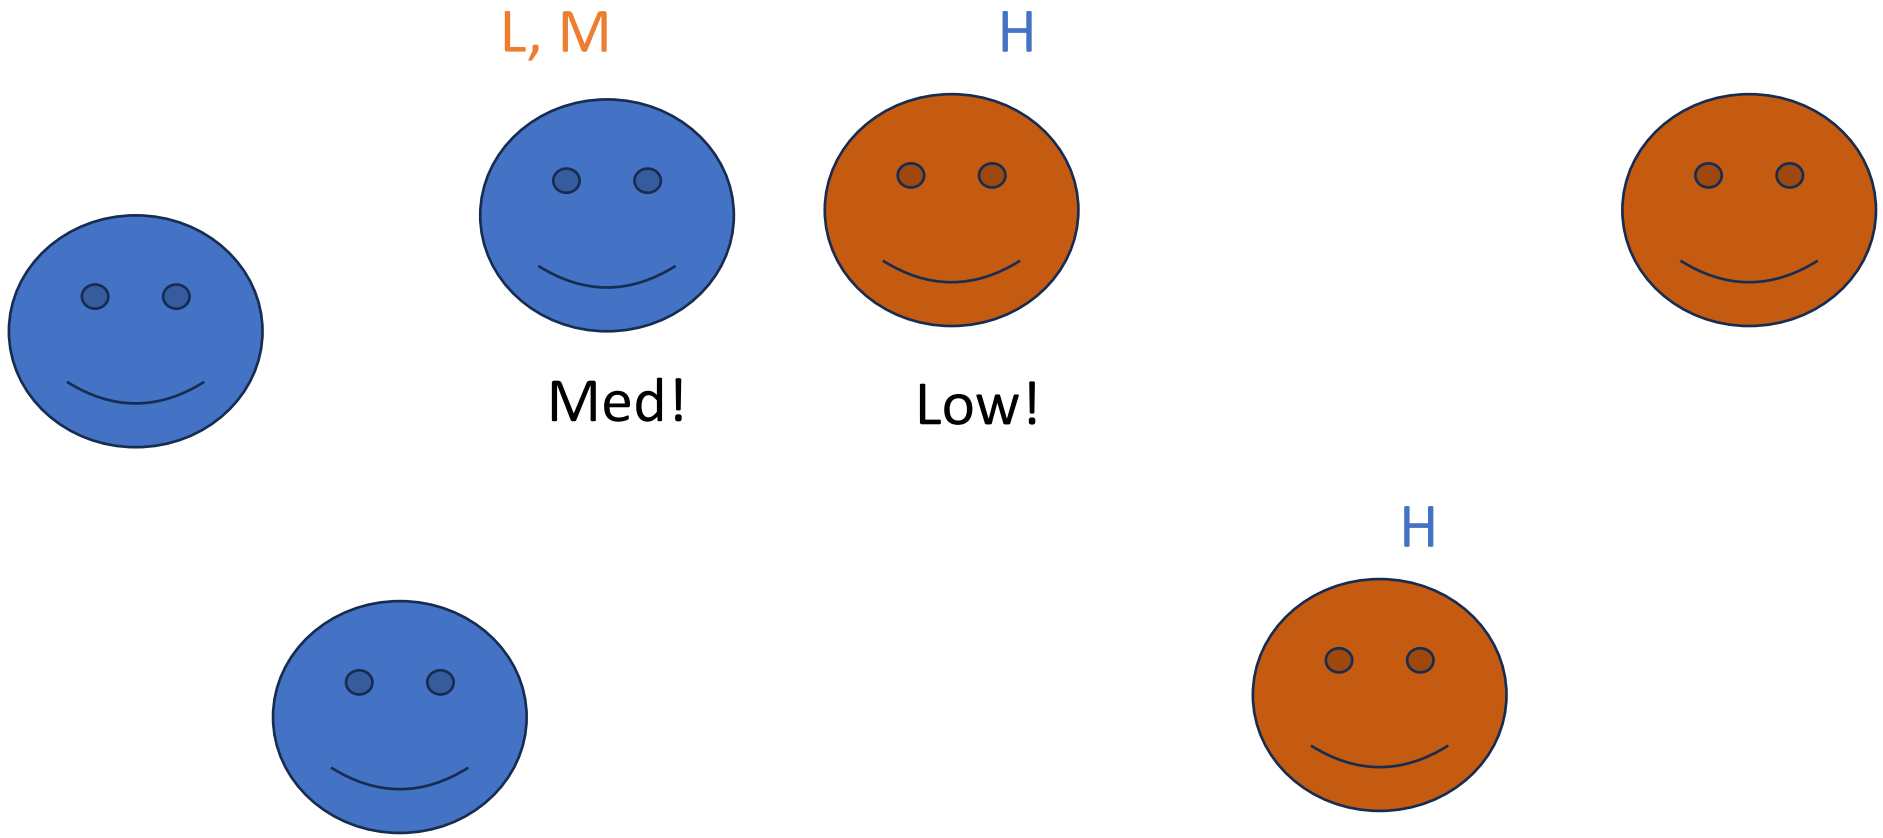
L, M
H
Med!
Low!
H

L, M, L

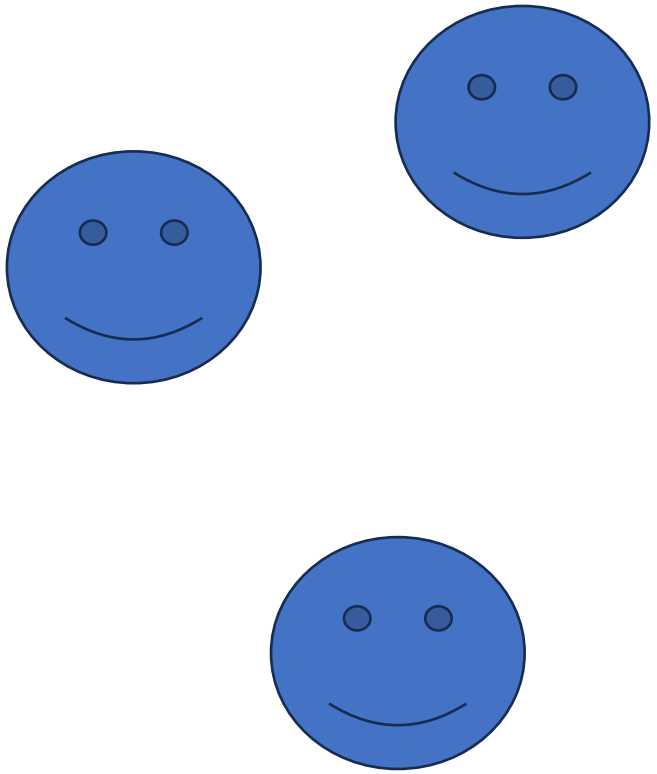

H, M

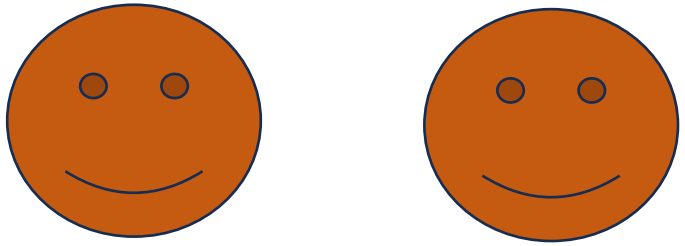

H

L, M, L

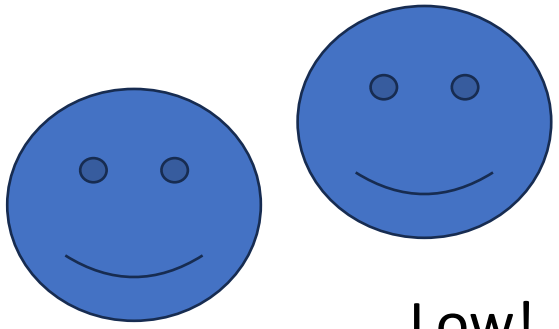

Med!

Low!

H, M

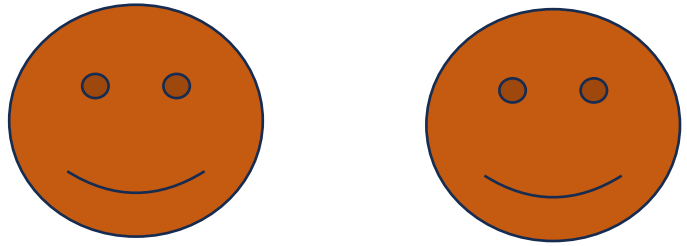

H

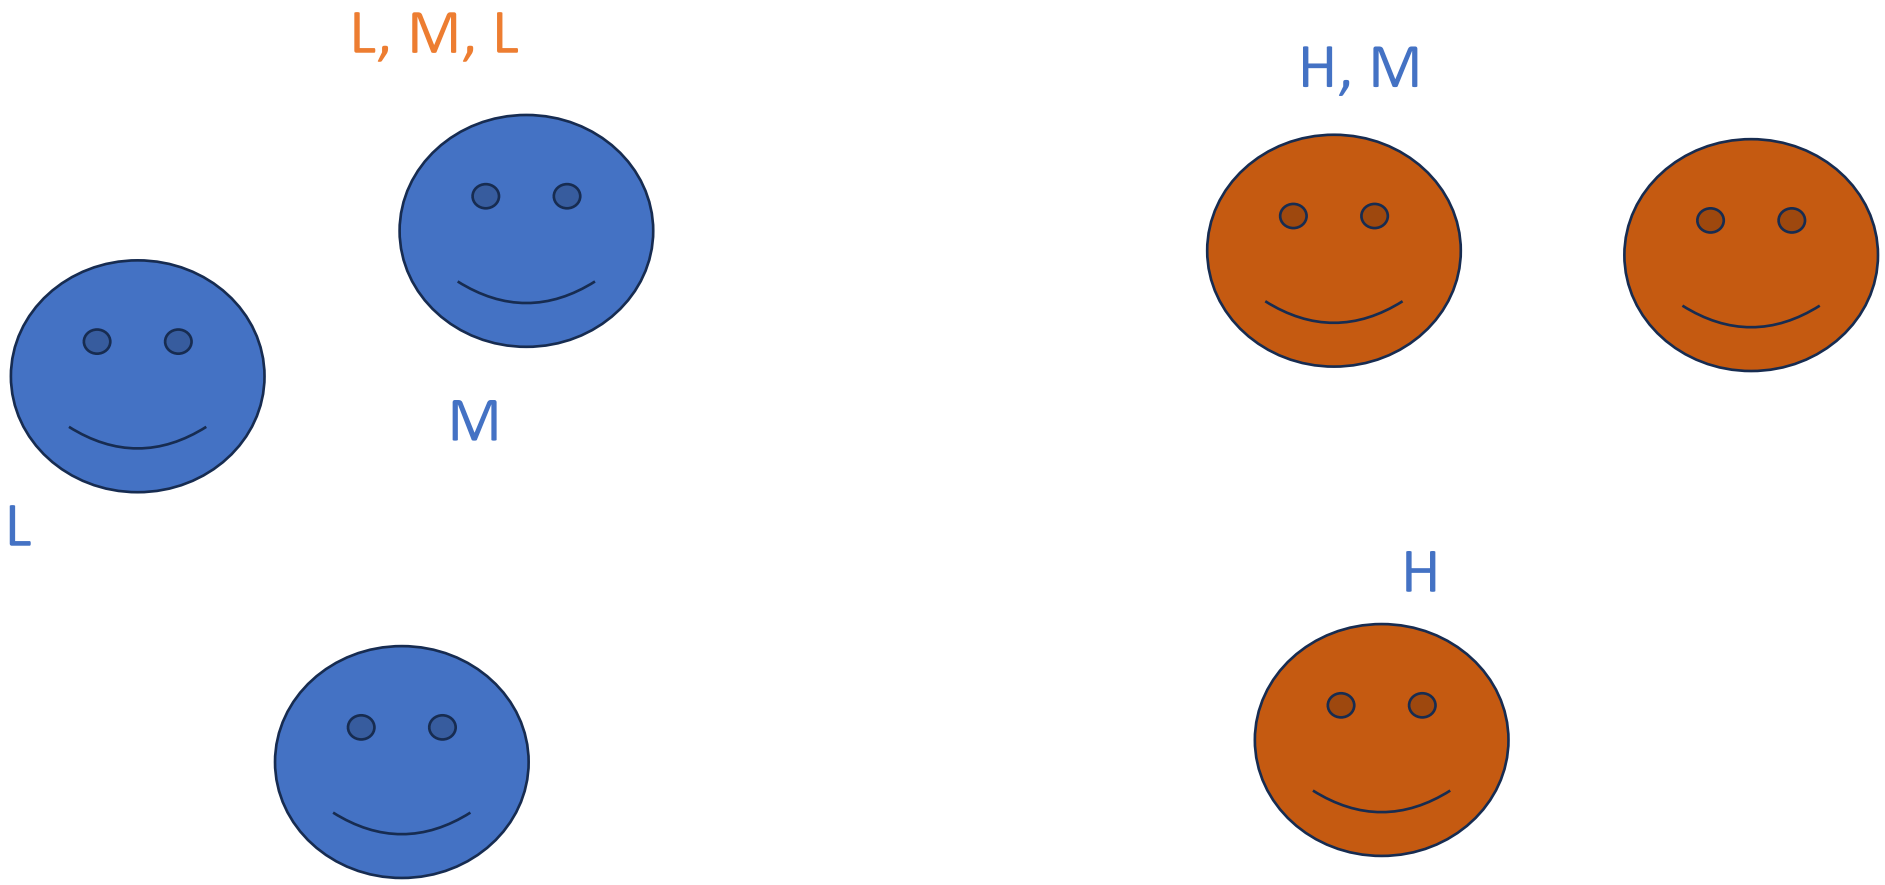
L, M, L
M
L
H, M
H

New Equilibria

		Player 2		
		Low	**Med**	**High**
Player 1	**Low**	4,4	4,5	4,6
	Med	5,4	5,5	0,0
	High	6,4	0,0	0,0

New Equilibria

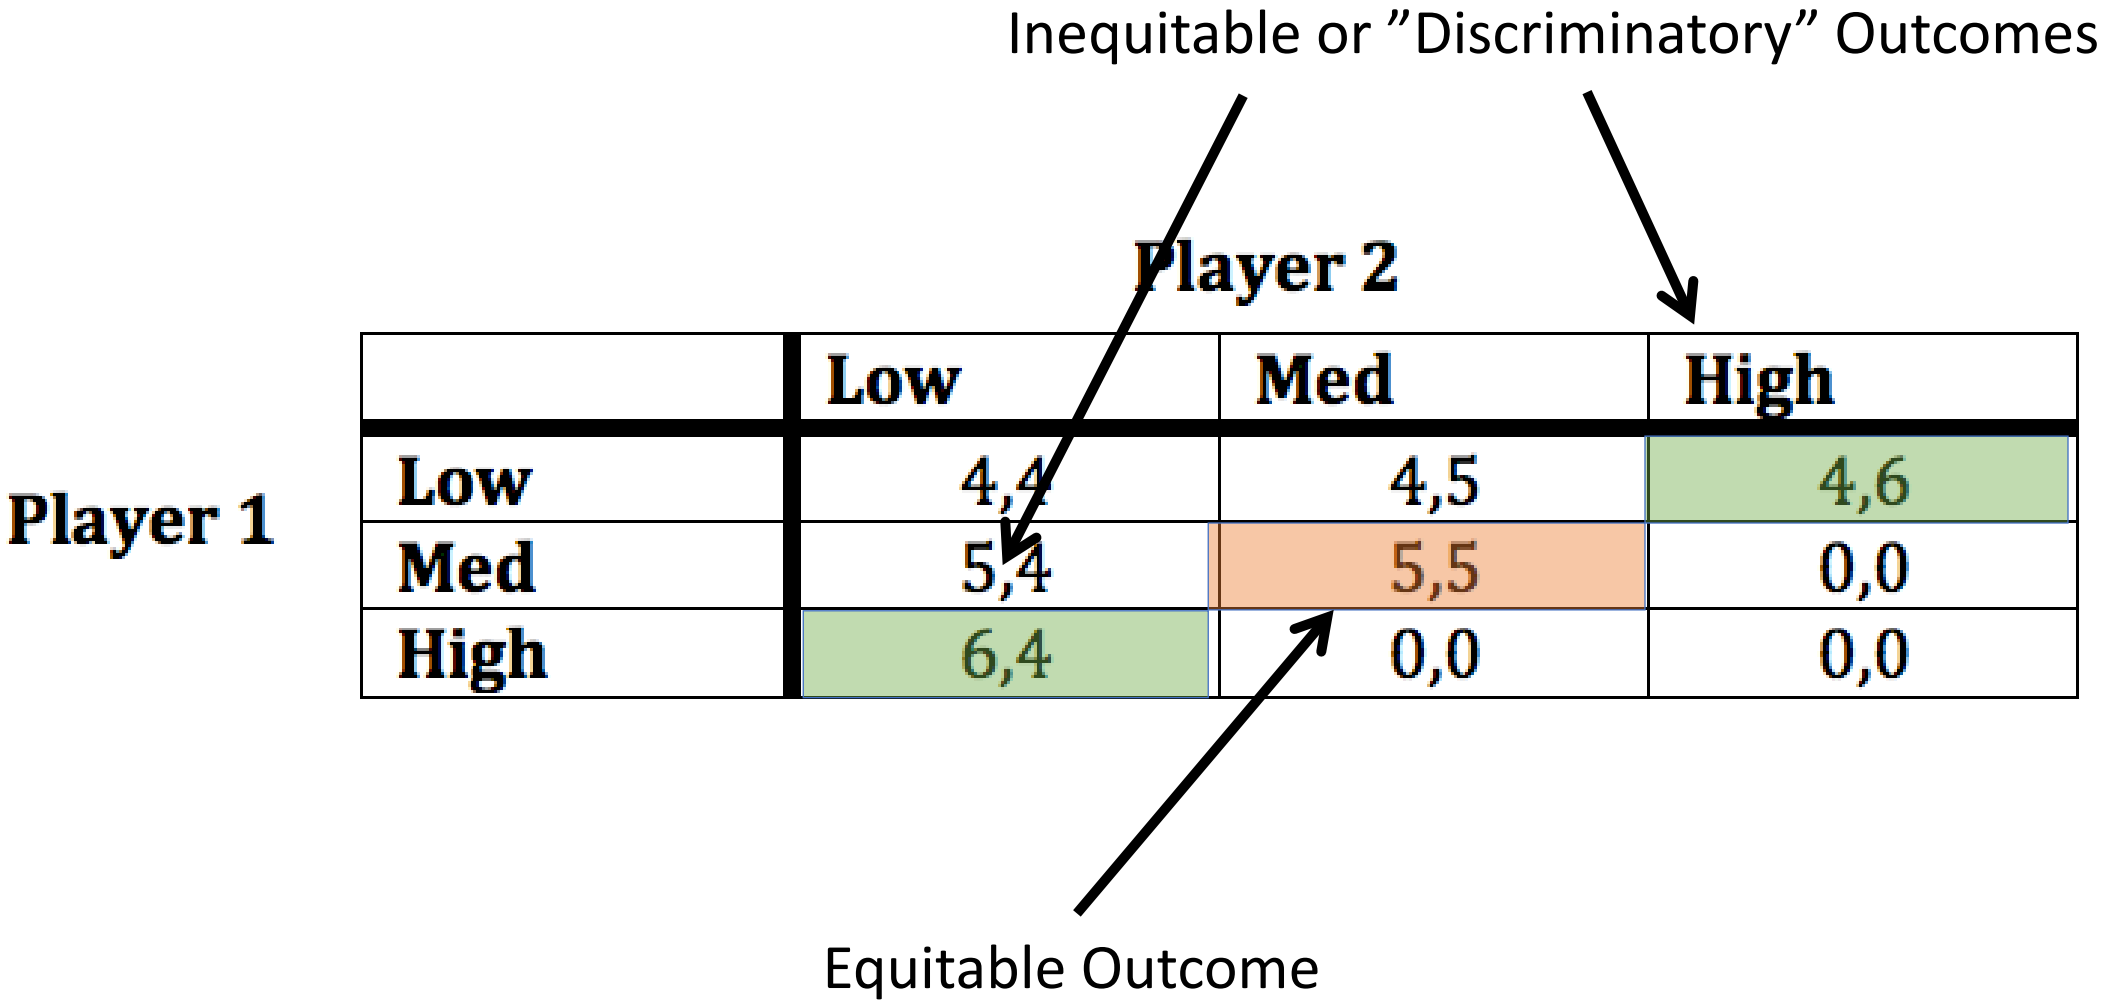

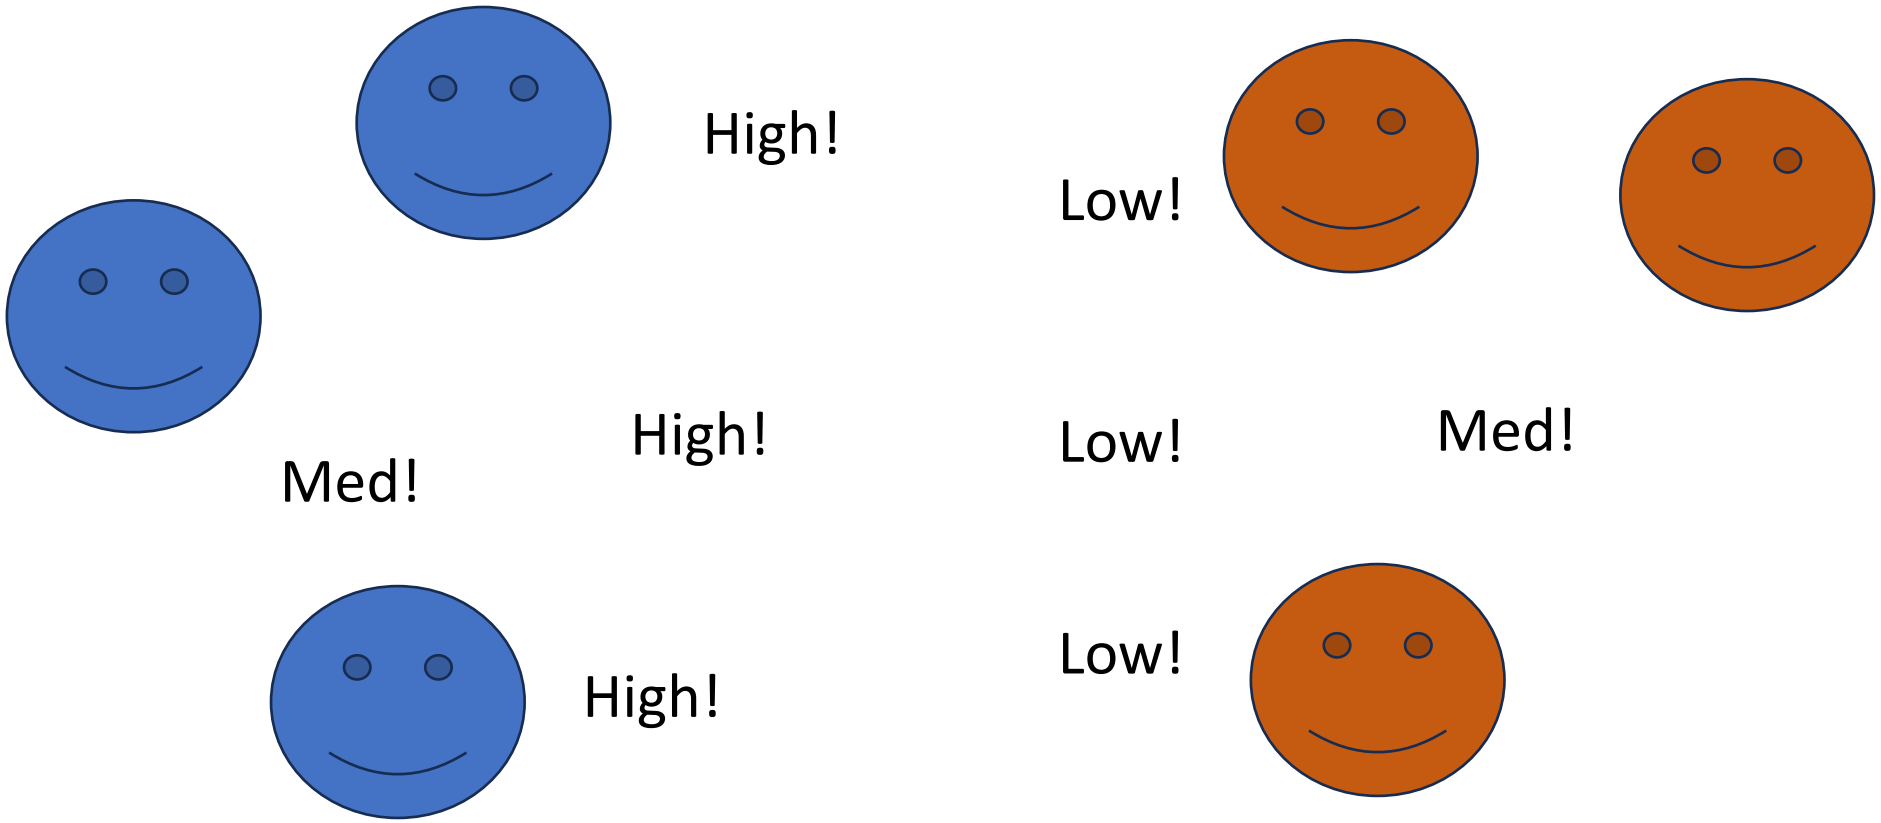
High!
Low!
High!
Low!
Med!
Med!
High!
Low!

L,L,L,L,L

L,L,L,L,L

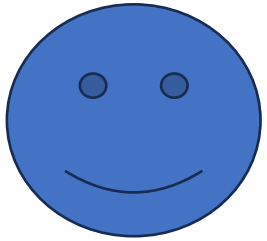

L,L,L,L,L

H,H,H,H,H

H,H,H,H,H

H,H,H,H,H

Symmetry Breaking

In models with one group, the symmetric (fair) outcome is special.

Tags create a symmetry breaker.

Now rules like “blues get more and oranges get less”, can be efficient. Actors can coordinate unfair expectations using social categories.

Emergence of Unfairness

Axtell, Epstein, and Young (2001) find that these kinds of ‘discriminatory’ conventions emerge robustly in their models.

Much subsequent work by myself and others confirms that **this is a robust, widespread result.**

Preconditions for Unfairness

Notably, for unfairness to emerge in these models, you need **fairly minimal preconditions**.

1) Social Categories
2) Conditional Behavior
3) Learning

These factors are pervasive and hard to eliminate.

Take-Away

An important message from this little history of modeling is that when you have social categories, **inequity is not hard to generate.**

Thus models do not generically predict fairness. Unfairness is also possible.

And in further work I've argued that **inequity is the more likely cultural evolutionary outcome.**

Roadmap

1) The Nash demand game
2) The evolution of fairness
3) The evolution of unfairness
4) Power, categories, and the pull of unfairness

Power and the Emergence of Bargaining

In “Power, Bargaining, and Collaboration” Justin Bruner and I show how **powerful groups tend to evolve to get more**.

In “Power by Association”, Travis LaCroix and I extend this to situations **with heterogenous groups.**

Operationalizing Power

To do this, we need to **add power to these models.**

We consider different ways of capturing power. Today I will just discuss one natural one, which goes back to John Nash.

Player 2

Player 1	**Low**	**Med**	**High**
Low	4,4	4,5	4,6
Med	5,4	5,5	D,d
High	6,4	D,d	D,d

Replicator dynamics

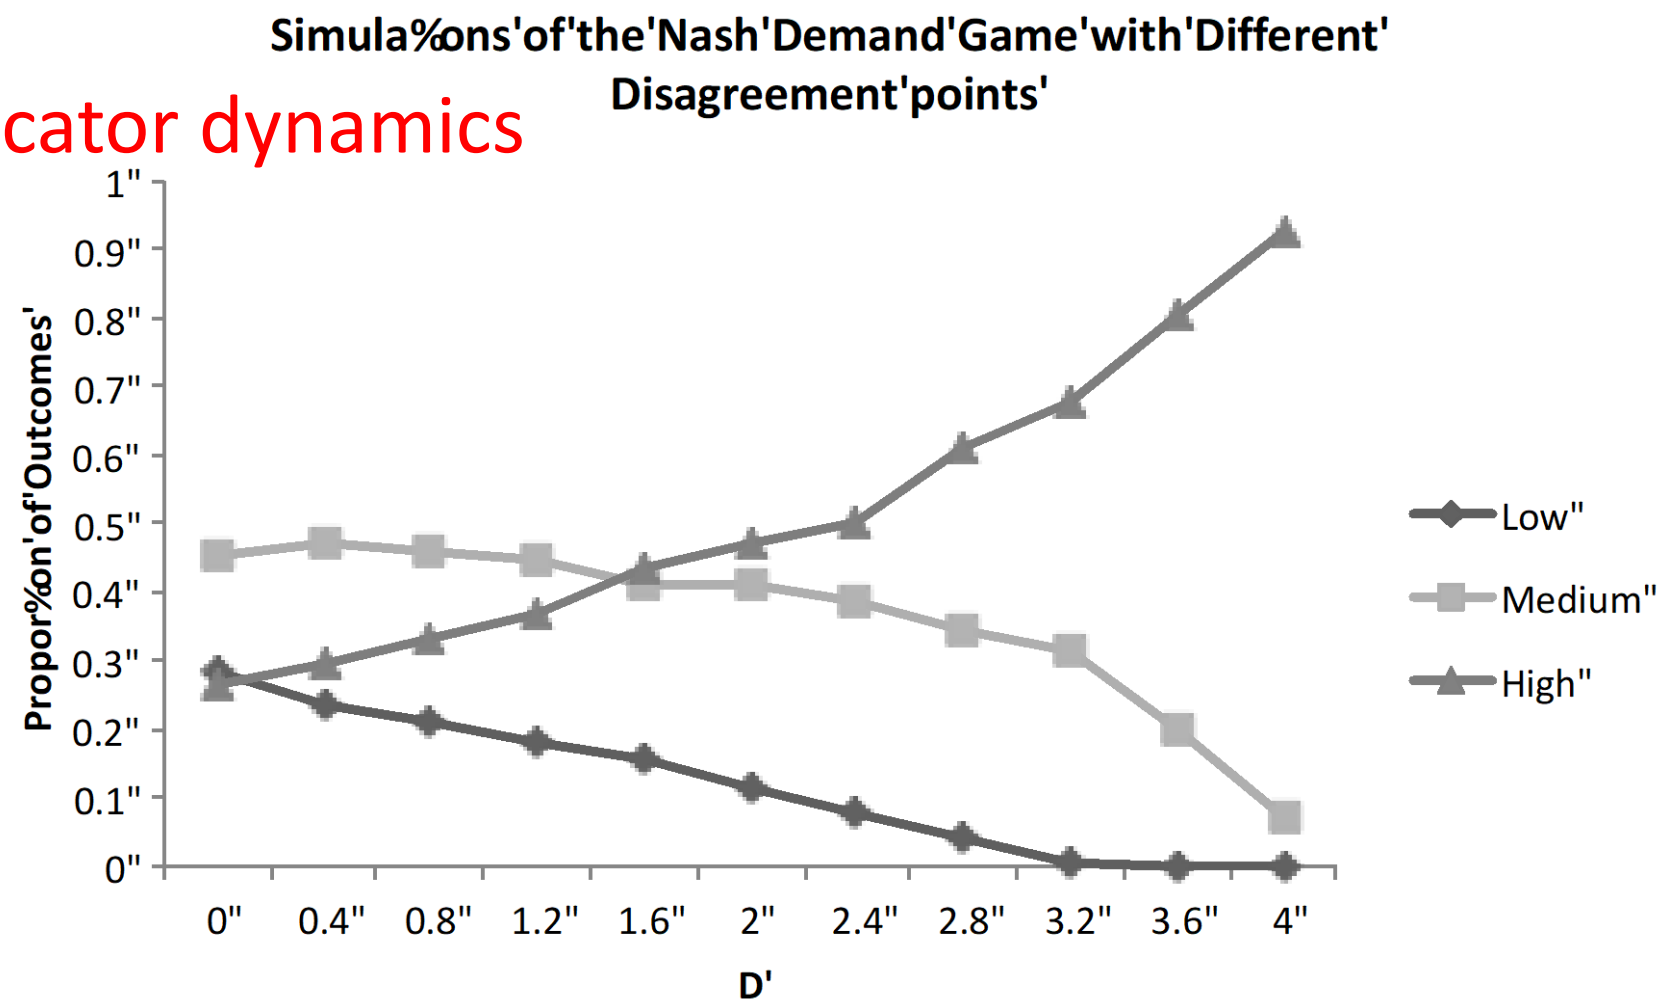

Powerful Individuals

Notably, this sort of pattern can arise even from adding just a few powerful individuals to one group.

So if, say, women are, on average, less empowered than men, the entire group may be disadvantaged by bargaining norms. This extends to women who are, themselves, empowered.

Agent based model

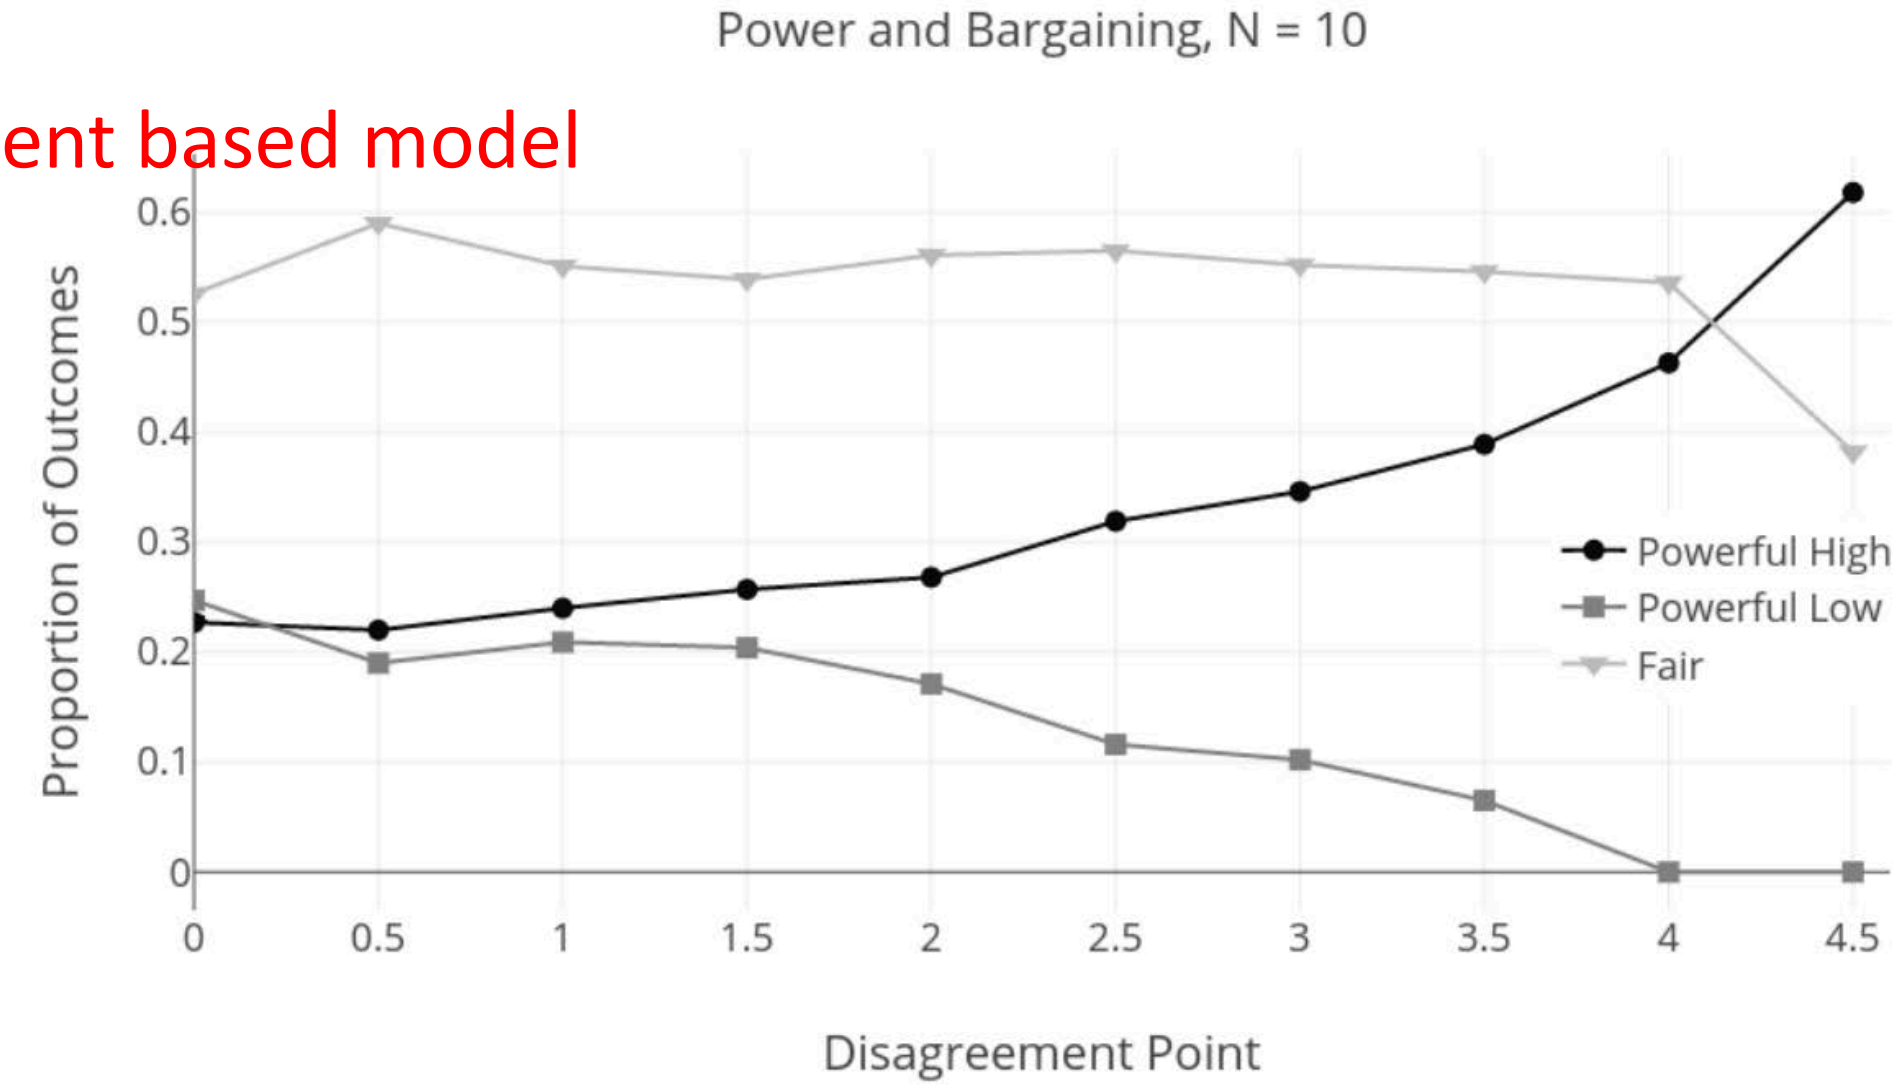

Why?

This happens because the “powerful” group has relatively little incentive to learn low demands. They do not get a significant advantage for doing so, compared to receiving their disagreement point.

This shifts their learning patterns to push them towards preferable outcomes.

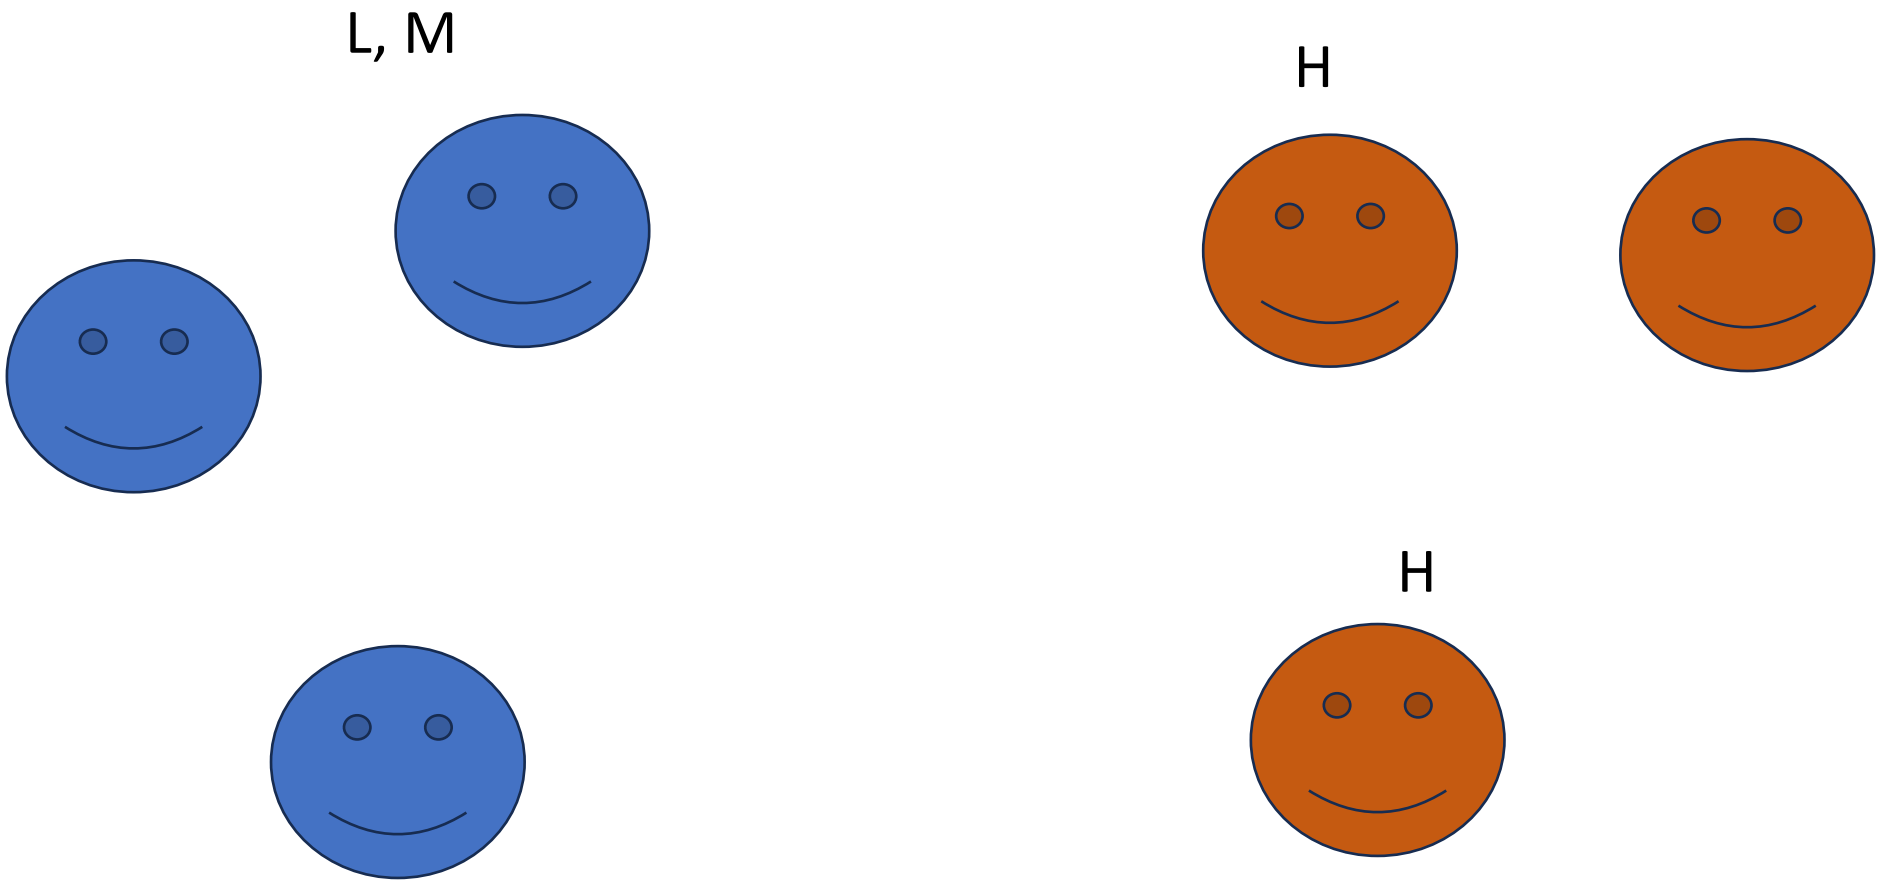
L, M
H
H

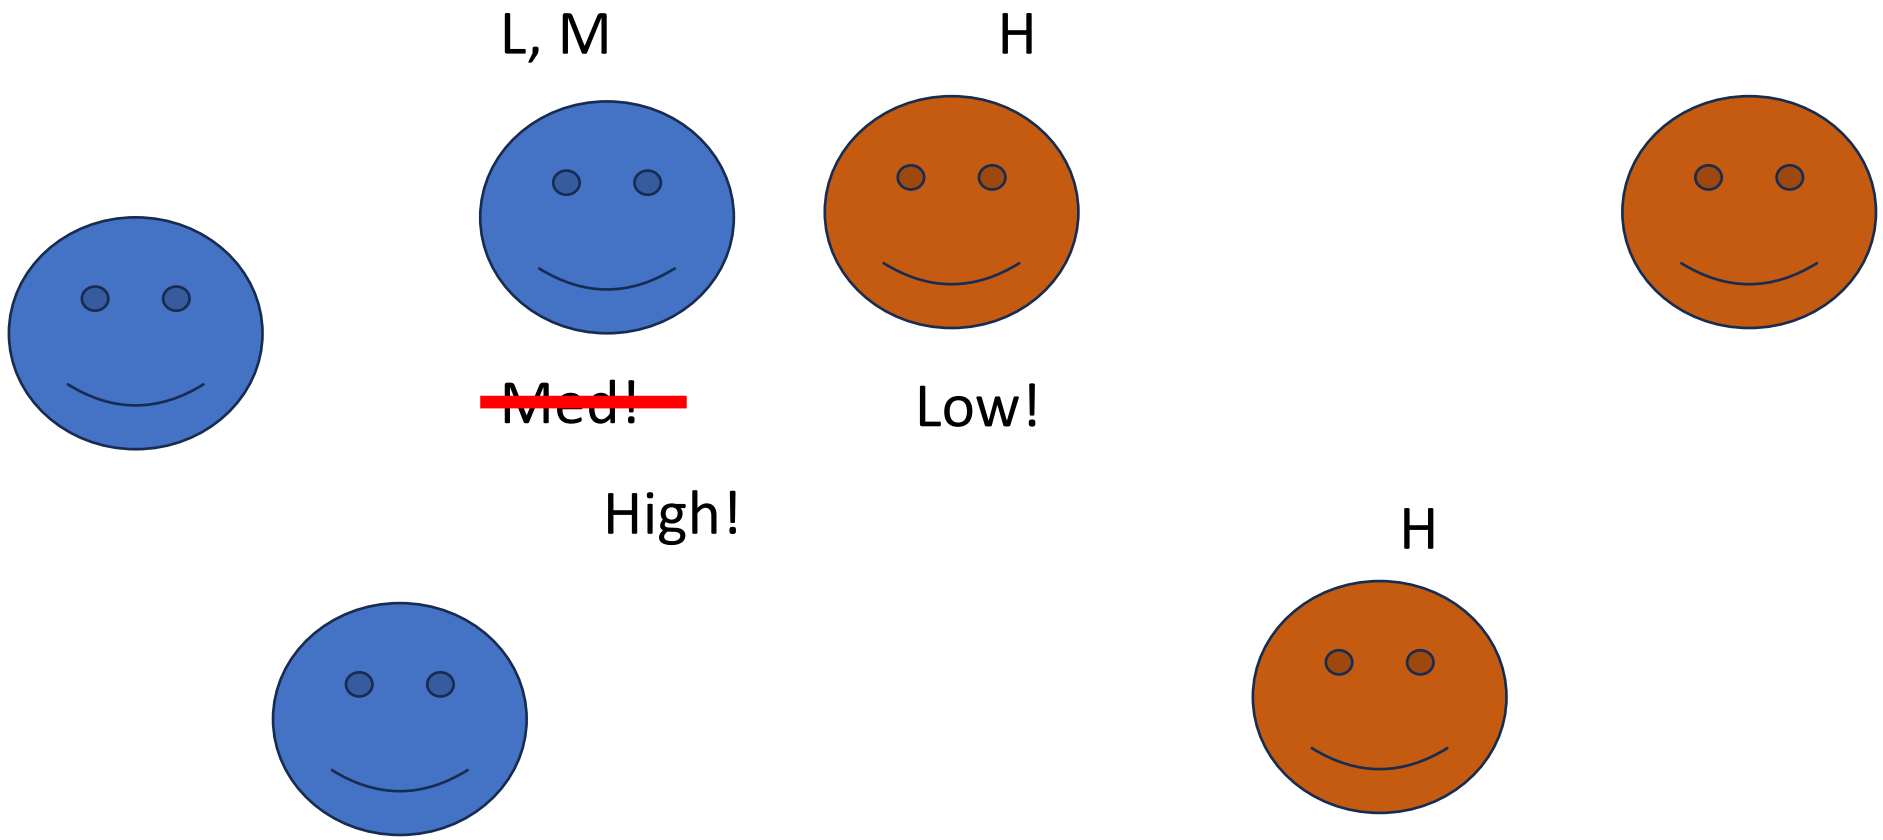
L, M
H
~~Med!~~
Low!
High!
H

Take-Away

One general take-away is that **power asymmetries between groups tend to promote inequity**.

Given that power asymmetries **are the rule in the real world**, this supports the picture where unfair outcomes are expected.

Extending the Paradigm

A number of papers have used this paradigm to think about the emergence of categorical unfairness, generally asymmetries promote unfairness, and it is robust

1) Power imbalances
2) Minority disadvantage
3) Information asymmetries
4) Intersectional disadvantage
5) Homophily
6) Dividing products of joint labor
7) The emergence of categories

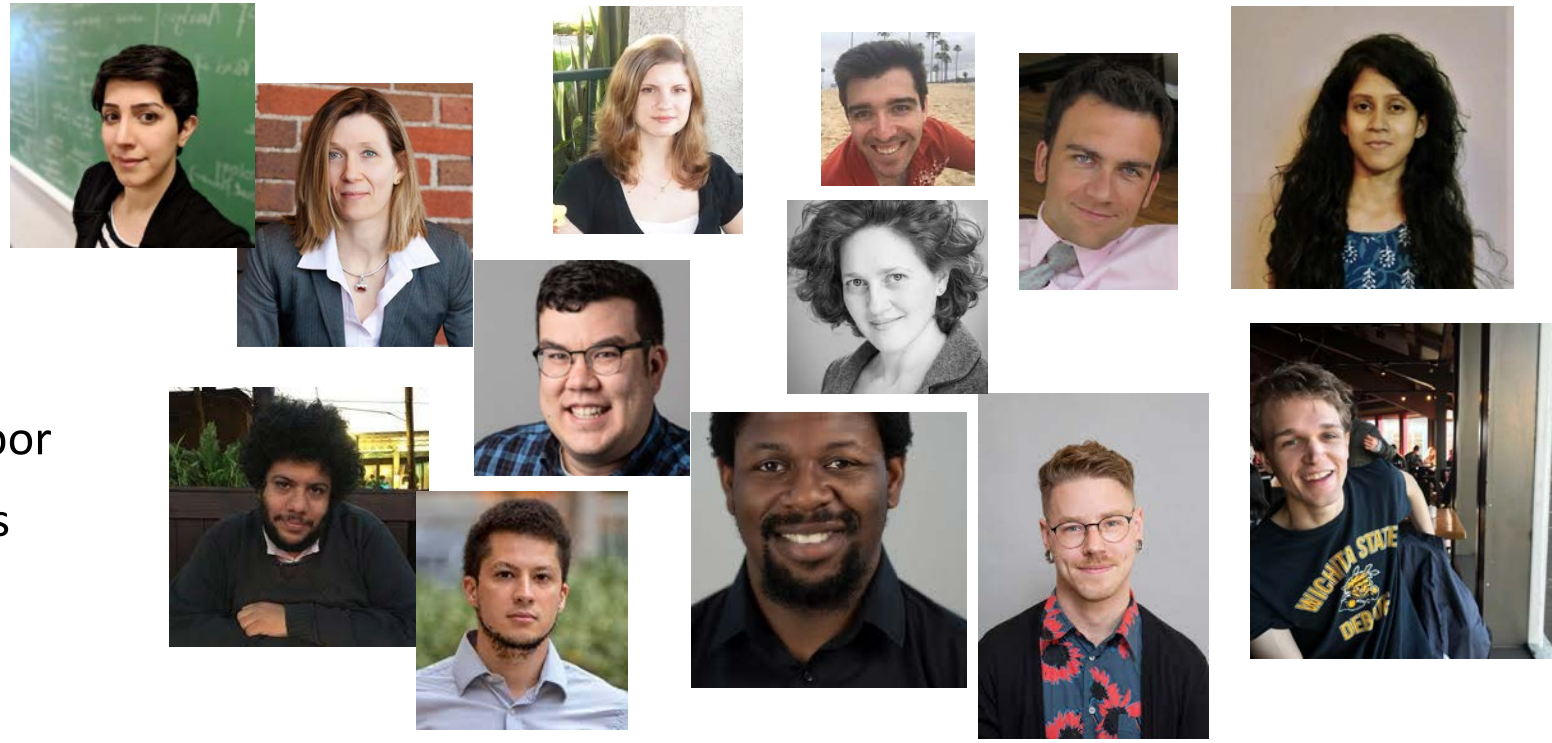

Revolution in Permanence

Altogether, this contributes to the picture where **unfairness is easy to get**.

In thinking about countering unfairness, then, we should not expect to do it and be done. **Unfairness emerges endogenously**.

Equity may not be an achievable state, but rather a process or activity.

Wrapping Up

1) Because of the presence of social categories, we shouldn't expect emergent bargaining patterns to be equitable or "fair".
2) This helps square earlier evolutionary game theoretic work with observations of categorical inequity
3) Power asymmetries between groups tend to increase the probability of unfairness
4) It may be useful, in promoting equity, to think about the cultural evolutionary pressures at work

Thank you!

Many thanks to all the collaborators here – Justin, Liam, Nathan, Olufemi, Travis– and others not listed – Hannah Rubin, Aydin Mohseni, Calvin Cochran, and Mike Schneider.

Thanks to Daniel Saunders, Michaela Popa, Sahar Heydari-Fard, and others for their insightful work on these topics.

Thanks to NSF and so many folks for comments.

Modeling Social Change

1) To the extent possible, and where appropriate, weaken the informational strength of categories
2) Actual behavioral change is necessary to disrupt “equilibria” – like those that constitute discriminatory bargaining conventions
3) When unfairness causes power asymmetries, it may be harder to disrupt unfair systems
4) Redistributing wealth/power may help
5) The conditions necessary for inequity to emerge are very minimal. For this reason, we should treat it as a problem we are always solving

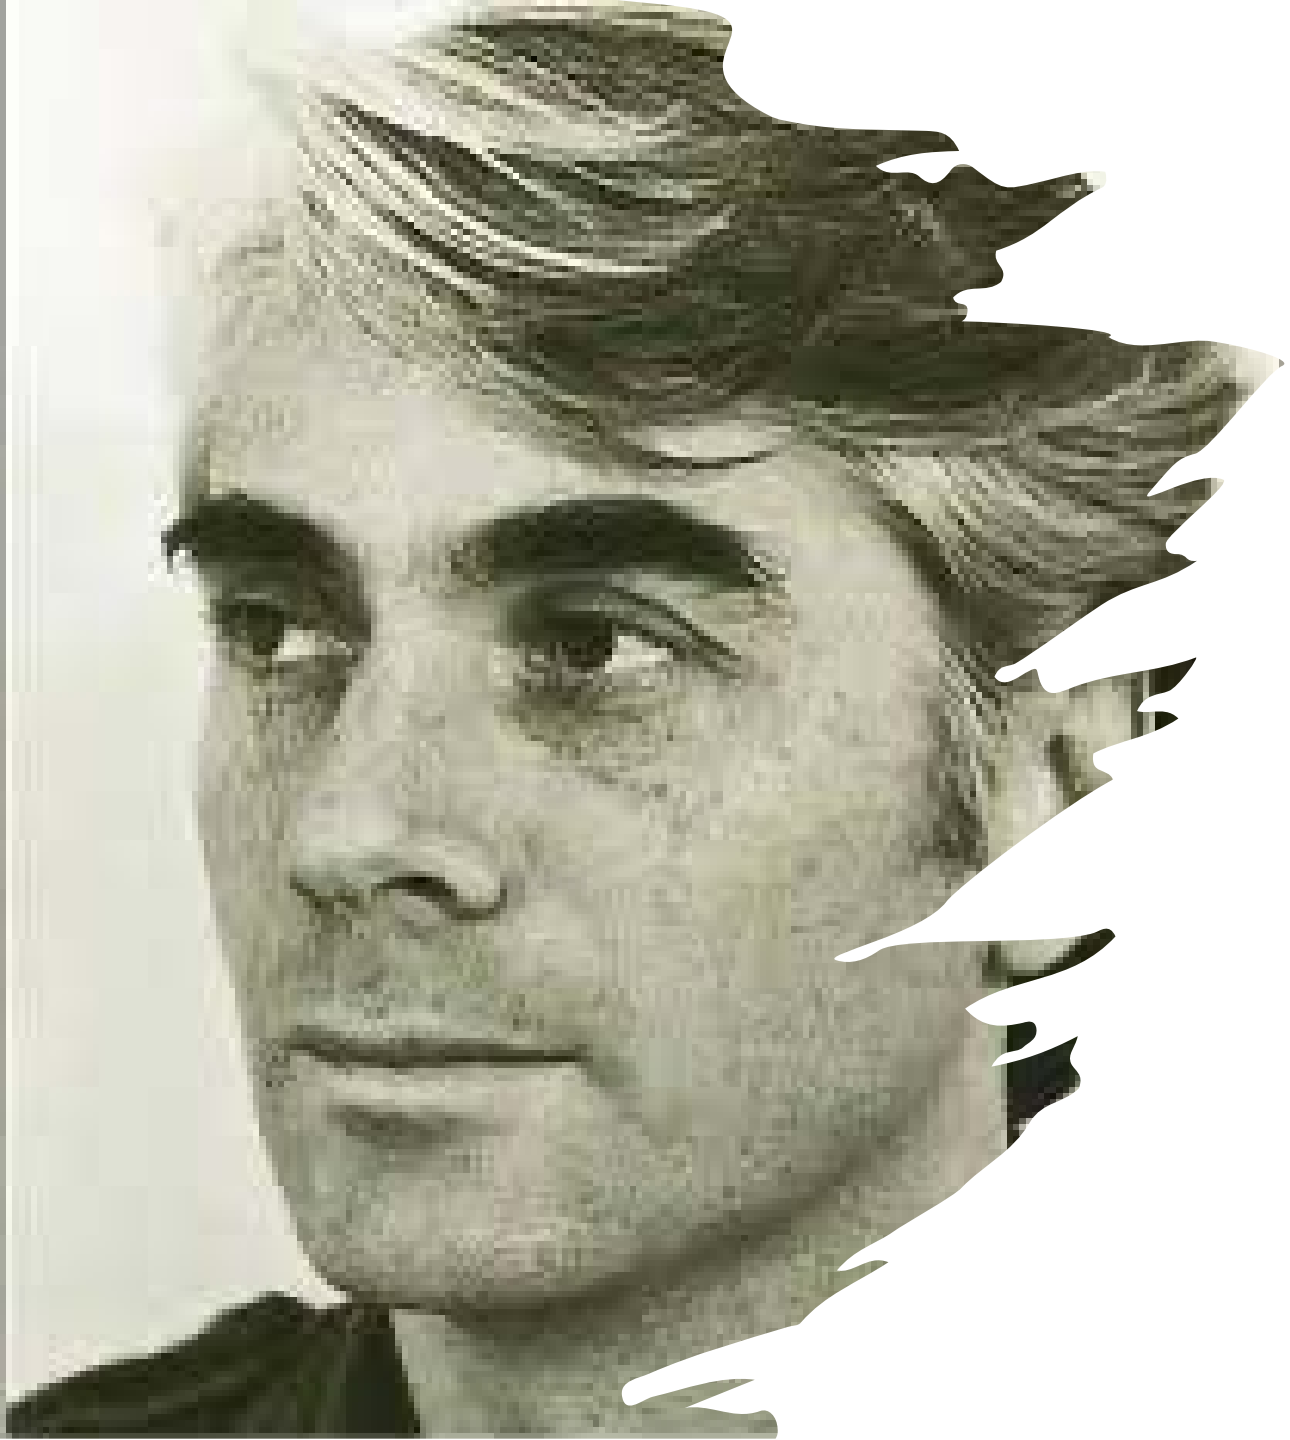

One Last Stab

Nozick's account of historical justice holds that distributions resulting from just initial acquisitions and just transfers are themselves just.

Our models arguably meet these conditions – but categorically inequitable distributions emerge endogenously.

A general take-away might be that just theorizing about complex historical processes can miss important possibilities.

Minority (Dis)advantage

Minority Disadvantage

Justin Bruner showed how in some circumstances a minority group can end up with a disadvantage in the emergence of bargaining conventions.

This happens because a small group learns more quickly about their out-group, and is thus more reactive.

Disadvantage vs. Advantage

I subsequently showed how minority *dis*advantage is more likely when:

1) Actors display out-group biases

2) Minority groups are economically disadvantaged (and thus risk-averse)

Experimental Settings

Aydin Mohseni, Hannah Rubin, and I provided experimental support for this effect.

Springer Link

Published: 25 October 2019

On the emergence of minority disadvantage: testing the cultural Red King hypothesis

Aydin Mohseni, Cailin O'Connor & Hannah Rubin

Synthese **198**, 5599–5621 (2021) | Cite this article