

14.773 Political Economy of Institutions and Development.

Lecture 10 and 11. Information, Beliefs, Norms and Politics

Daron Acemoglu

MIT

February-March 2026

Information-Behavior Feedbacks

- An important dimension of politics is about beliefs. For example, voters may be uncertain about how distortionary redistributive policies will be or about the strengths of some rival groups. They may also be uncertain about others' beliefs or expected behavior—thus about “norms” as well.
- Most interestingly, there may be information-behavior feedbacks: certain types of information leads to different behavior patterns, which then support different types of information.
- One extreme case of this is “self-confirming equilibria”: certain beliefs support behavior that then do not create an opportunity for information that contradicts these beliefs to emerge.
- One application of these ideas is in Piketty (1995): when voters think that these are very distortionary, then they may choose low redistribution. But then the society may not learn about true consequence of redistributive policies.
- Start with classic political economy models, and then consider norms.

Redistribution and Mobility: Model

- An individual i of generation t has utility

$$U_{it} = \hat{y}_{it} - \frac{1}{2\alpha} e_{it}^2,$$

where $\hat{y}_{it} = (1 - \tau)y_{it} + T$ is after-tax income, is $y_{it} \in \{0, 1\}$ is earned income (and can be thought of as success or failure), τ is the tax rate, T is a lump-sum transfer, and e_{it} is the effort level.

- Suppose that success depends on effort and also on

$$\mathbb{P}(y_{it} = 1 \mid e_{it} = e \text{ and } y_{it-1} = 0) = \pi_0 + \theta e,$$

and

$$\mathbb{P}(y_{it} = 1 \mid e_{it} = e \text{ and } y_{it-1} = 1) = \pi_1 + \theta e,$$

where $\pi_1 \geq \pi_0$.

- The gap between these two parameters is the importance of “inheritance” in success, whereas θ is the importance of “hard work”.
- The vector of parameters (θ, π_0, π_1) is unknown.

Model (continued)

- At any given point in time, individuals will have a posterior over this policy vector μ_{it} , shaped by their dynasty's prior experiences as well as other characteristics in the society that they may have observed.
- The only policy tool is a tax rate on output, which is then redistributed lump sum.
- Let total output under tax rate τ be $Y(\tau)$.
- This implies that given an expectation of a tax rate τ , an individual with a successful or unsuccessful parent denoted by $z = 1$ or $z = 0$ will choose

$$e_z(\tau, \mu) \in \arg \max_e \mathbb{E}_\mu [(1 - \pi_z - \theta e) \tau Y(\tau) + (\pi_z + \theta e) ((1 - \tau) + \tau Y(\tau))] - \frac{1}{2\alpha} e^2,$$

where the expectation is over the parameters.

Effort and Voting

- It can be easily verified that

$$e_z(\tau, \mu) = e(\tau, \mathbb{E}_\mu \theta) = \alpha(1 - \tau) \mathbb{E}_\mu \theta.$$

- Therefore, all that matters for effort is the expectation about the parameter θ .
- Now given this expectation, individuals will also choose the tax rate by voting.
- Individuals vote for the tax rates that maximizes “expected social welfare” $\mathbb{E}_{\mu_{it}} V_t$ (why is this conditional on μ_{it} ?).
- Given the quadratic utility function, it can be verified that individuals have single peaked preferences, with bliss point given by

$$\tau(\mu_{it}) \in \arg \max \mathbb{E}_{\mu_{it}} V_t.$$

- An application of the median voter theorem then gives the equilibrium tax rate is the median of these bliss points.

Evolution of Beliefs

- How will an individual update their beliefs? Straightforward application of Bayes rule gives the evolution of beliefs.
- For example, for an individual $i \in \mathcal{I}$ with a successful or unsuccessful parent denoted by $z = 1$ or $z = 0$, starting with beliefs μ_{it} , with support $S[\mu_{it}]$, we have that for any $(\theta, \pi_0, \pi_1) \in S[\mu_{it}]$, we have

$$\mu_{it+1}(\theta, \pi_0, \pi_1) = \mu_{it}(\theta, \pi_0, \pi_1) \frac{\pi_z + \theta e(\tau_t, \mathbb{E}_{\mu_{it}} \theta)}{\int [\pi'_z + \theta' e(\tau_t, \mathbb{E}_{\mu_{it}} \theta)] d\mu_{it}}.$$

- Note that individuals here are not learning from the realized tax rate, simply from their own experience. This is because individuals are supposed to have “heterogeneous priors”. They thus recognize that others have beliefs driven by their initial priors, which are different from theirs and there is no learning from initial priors.
- Is this just a consequence of heterogeneous priors?

Evolution of Beliefs (continued)

- Standard results about Bayesian updating, in particular from the *martingale convergence theorem*, imply the following:

Proposition

The beliefs of individual $i \in \mathcal{I}$, μ_{it} , starting with any initial beliefs μ_{i0} almost surely converges to a stationary belief $\mu_{i\infty}$.

- But if beliefs converge for each dynasty, then the median also converges, and thus equilibrium tax rates also converge.

Proposition

Starting with any distribution of beliefs in the society, the equilibrium tax rate τ_t almost surely converges to a stationary tax rate τ_∞ .

Limits of Learning

- The issue, however, is that this limiting tax rate need not be unique, because the limiting stationary beliefs are not necessarily equal to the distribution that puts probability 1 on truth.
- The intuition for this is the same as “self confirming” equilibria, and can be best seen by considering an extreme set of beliefs in the society that lead to $\tau = 1$ (because effort doesn’t matter at all).
- If $\tau = 1$, then nobody exerts any effort and there is no possibility that anybody can learn that effort actually matters.

Limits of Learning (continued)

- The characterization of the set of possible limiting beliefs is straightforward.
- Define $M^*(\tau)$ be the set of beliefs that are “self consistent” at the tax rate τ in the following sense:
- For any $\tau \in [0, 1]$, we have

$$\begin{aligned} M^*(\tau) = & \{ \mu : \text{for all } (\theta, \pi_0, \pi_1) \in S[\mu], \\ & \pi_z + \theta e(\tau, \mathbb{E}_\mu \theta) = \pi_z^* + \theta^* e(\tau, \mathbb{E}_\mu \theta) \\ & \text{for } z = 0, 1 \text{ and } (\theta^*, \pi_0^*, \pi_1^*) \in S[\mu] \}. \end{aligned}$$

- Intuitively, these are the set of beliefs that generate the correct empirical frequencies in terms of upward and downward mobility (success and failure) given the effort level that they imply.
- Clearly, if the tax rate is in fact τ and $M^*(\tau)$ is not a singleton, a Bayesian cannot distinguish between the elements of $M^*(\tau)$: they all have the same observable implications.

Limits of Learning (continued)

- Now the following result is immediate.

Proposition

Starting with any initial distribution of beliefs in society $\{\mu_{i0}\}_{i \in \mathcal{I}}$, we have that

- For all $i \in \mathcal{I}$, $\mu_{i\infty}$ exists and is in $M^*(\tau_\infty)$, and
- τ_∞ is the median of $\{\tau(\mu_{i\infty})\}_{i \in \mathcal{I}}$.

- This proposition of course does not rule out the possibility that there will be convergence to beliefs corresponding to the true parameter values regardless of initial conditions. But it is straightforward from the above observations establish the next result:

Limits of Learning (continued)

Proposition

Suppose \mathcal{I} is arbitrarily large. Then for any $\{\mu_{i\infty}\}_{i\in\mathcal{I}} \in M^*(\tau_\infty)$ such that τ_∞ is the median of $\tau(\mu_{i\infty})$, there exists a set of initial conditions such that there will be convergence to beliefs $\{\mu_{i\infty}\}_{i\in\mathcal{I}}$ and tax rate τ_∞ with probability one.

- This proposition implies that a society may converge and remain in equilibria with very different sets of beliefs and these beliefs will support different amounts of redistribution.
- Different amounts of redistribution will then lead to different tax rates, which “self confirm” these beliefs because behavior endogenously adjusts to tax rates.

Interpretation

- Therefore, according to this model, one could have the United States society converge to a distribution of beliefs in which most people believe that θ is high and thus vote for low taxes, and this in turn generates high social mobility, confirming the beliefs that θ is high.
- Many more Europeans believe that θ is low (and correspondingly $\pi_1 - \pi_0$ is high) and this generates more redistribution and lower social mobility.
- Neither Americans nor Europeans are being “irrational”.

Discussion

- How to interpret these results?
- Perhaps a good approximation to the formation of policemen individuals are not “hyper rational”.
- But why don't different societies learn from each other?
- How likely is this process to lead to multiple stable points?

Voting and Experimentation

- Information is in general acquired dynamically, as a result of past political choices.
- Example: Economic or social reforms
 - Reforms make winners and losers, whose identities are unknown *ex ante*.
 - Fernandez and Rodrik (1991): resistance to trade liberalization because of losers' fear that they will not be compensated.
- But in a dynamic context, there are new effects that make political actors even more averse to the information and experimentation.
- Strulovici (2010): two novel reasons for this:
 - Loser trap (can't return to status quo).
 - Winner frustration (can't exploit new alternative).

Illustration

- Ann, Bob and Chris go to the restaurant every week-end.
- They always choose their restaurant by majority rule.
- A new restaurant has opened.
- If any one of them could choose *alone* future restaurants, he or she would try the new one now.
- However, it is possible that all three will vote against trying this restaurant.

Illustration (continued)

- Experimentation with new alternatives is less attractive when one has to share power.
- Sharing control induces two opposite control loss effects, which have different implications.
 - *Loser trap.* If Ann and Bob like the new restaurant, they will impose it to Chris in the future, even if he does not like it.
 - *Winner frustration.* If only Ann likes the new restaurant, she will be blocked by Bob and Chris. So the “risk” of trying a new restaurant need not be rewarded even for those who do turn out to like it.
- Majority-based experimentation is also shorter than the socially efficient outcome.
- New winners induce more experimentation from remaining voters.

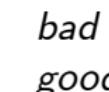
Model: Single Agent Problem

- Safe (S) and risky (R) actions.
- R can be good or bad. Agent type initially unknown.
- Continuous time with fixed discount rate, infinite horizon.
- At each instant, one action (S or R) is chosen.

Model: Single Agent Problem (continued)

- Payoffs:

$$S \rightarrow s > 0$$

R 

 bad : 0

good: positive reward $g > 0$ starting at some Poisson arrival time

- bad (loser) $<$ safe $<$ good (winner).
- Bayesian updating of beliefs:

$$\frac{dp_t}{dt} = -\lambda p_t (1 - p_t)$$

where λ arrival rate of good outcome from the risky action and p_t belief at time t that risk action is good (or the agent is of good type).

Model: Single Agent Problem (continued)

- Equilibrium: Experiment up to some level of belief $p^{SD} < p^{\text{myopic}}$
- This is because of the option value of experimentation.

Model: Collective Decision-Making

- N (odd) agents.
- Publicly observed payoffs.
- Types are iid. Initially, $\text{Prob}[good] = p_0$ for all.
- Arrival times also independent across agents.
- State variables (k, p) where k is number of *sure winners*, and $p = \text{Prop}[good]$ for *unsure voters*.
- Equilibrium concept: *Markov Voting Equilibrium*
- At any time, chose the action preferred by majority (given that the same rule holds in the future).
- Equilibrium can be solved by backward induction on number of sure winners.

Markov Voting Equilibrium

- A *Majority Voting Equilibrium (MVE)* is a mapping $C : (k, p) \rightarrow \{S, R\}$ such that $C = R$ if $k > k_N = (N - 1)/2$ and $C = S$ if $k \leq k_N$ and

$$\begin{aligned} ru(k, p) &= \lambda pg + \lambda p[w(k + 1, p) - u(k, p)] + & (1) \\ &\quad \lambda p(n - k - 1)[u(k + 1, p) - u(k, p)] - \lambda p(1 - p) \frac{\partial u}{\partial p} > s \end{aligned}$$

where u and w are the value of functions of unsure voters and sure winners when voting rule C determines future votes.

Collective Decision-Making: Structure of Equilibrium

- Now threshold belief $p^G(k)$ for stopping when there has been k people revealed to be of good type until now.
- Monotonicity: $p^G(k)$ is decreasing in k .
- Intuition: Good news for any one prompts remaining unsure voters to experiment more.
 - Why? Suppose to the contrary that experimentation stops when a new winner is observed.
 - Then, risky action pays lower expected payoffs and has no option value.
 - Therefore, experimentation was not optimal when the news arrived: contradiction.

Collective Decision-Making: Comparison

- We have that $p^G(k)$ is always greater than what social planner maximizing utilitarian welfare would choose.
- This is because of loser trap and winner frustration.

Comparative Statics

- - Experimentation decreases if N increases (enough): $p(k, N)$ almost increases in N .
 - Agents behave myopically as $N \rightarrow \infty$
- For N above some threshold, agents prefer safe action even if trying risky action would immediately reveals types.

Alternative Rules

- Suppose R requires unanimous approval.
- This gets rid of the loser trap.
- However, this increases winner frustration, since R is less likely to be played in the long run.
- Which rule performs better depends on the relative strengths of the two effects.

Cycles of Conflict

- Conflict (between ethnic groups, religious groups, countries, ideologies, social classes, rival individuals) is endemic.
- Why? Part of it may be related to incorrect information ("misperceptions") and relatedly to **fear** of actions, intentions and behavior of the other party as Thucydides emphasized long ago.
- Often continuing cycles of conflict between different groups. Partly related to information:

Group A's actions look aggressive

- ⇒ Group B thinks Group A is aggressive
- ⇒ Group B acts aggressively
- ⇒ Group A thinks Group B is aggressive
- ⇒ Group A acts aggressively ...

Examples

- Spirals in the World Horowitz (2000) on ethnic conflict:
“The fear of ethnic domination and suppression is a motivating force for the acquisition of power as an end . . . The imminence of independence in Uganda aroused ‘fears of future ill-treatment’ along ethnic lines. In Kenya, it was ‘Kikuyu domination’ that was feared; in Zambia, ‘Bemba domination’; and in Mauritius, [‘Hindu domination’] . . . ”
- Serbo-Croatian War (DellaVigna et al, 2011).
- Protestant-Catholic Conflict in Northern Ireland.
- Trade (Guiso, Sapienza, and Zingales, 2009, Bottazzi, Da Rin, and Hellmann, 2011).
- Political polarization (Sunstein, 2006).

Ebbs and Flows of Conflict

- But not ever-lasting continuous conflict.
- Ethnic conflict in Africa way down in last 20 years.
- France and Germany not on brink of war, and trade a lot.
- Conflict and distrust in Balkans greatly diminished.
- Political polarization in U.S. was probably as bad or worse in first third of 20th century.

Idea for Cycles

- Once Groups A and B are both acting aggressively, aggression becomes uninformative of their true types.
- Once this happens, one group will experiment with cooperation, which causes trust to restart.
- Conflict spirals cannot last forever, because if they did the informational content of conflict would eventually dissipate.

Model

- Timing and Actions 2 groups, A and B. Time $t = 0, 1, 2, \dots$
- Overlapping generations.
- At time t , one active player: player t .
- Player t takes pair of actions $(x_t, y_t) \in \{0, 1\} \times \{0, 1\}$.
- t even \implies player t from Group A.
- t odd \implies player t from Group B.

Model: Information

- Before player t takes actions, observes noisy signal $\tilde{y}_{t-1} \in \{0, 1\}$.

$$Pr(\tilde{y}_{t-1} = 1 | y_{t-1} = 1) = 1 - \pi$$

$$Pr(\tilde{y}_{t-1} = 1 | y_{t-1} = 0) = 0.$$

- Each group is either normal or bad.
- If normal, all representatives are normal types.
- If bad, all representatives are bad types.

$$Pr(\text{bad}) = \mu_0 \in (0, \mu^*).$$

Model

- For bad player t , playing $(x_t = 0, y_t = 0)$ is dominant strategy.
- For normal player t , utility function is

$$u(x_t, \tilde{y}_{t-1}) + u(\tilde{y}_t, x_{t+1}).$$

- Assume “subgame” between neighboring players is coordination game, and $(1, 1)$ is Pareto-dominant equilibrium: $u(1, 1) > u(0, 1)$,
 $u(0, 0) > u(1, 0)$,
 $u(1, 1) > u(0, 0)$.

Equilibrium

- What happens in (sequential) equilibrium?
- Normal player t plays $x_t = 1$ if and only if $\tilde{y}_{t-1} = 1$.
- So normal player 0 plays $y_0 = 1$ if and only if μ_0 is below some threshold μ^* :

$$\mu^* \equiv \left(\frac{u(1, 1) - u(0, 0)}{u(1, 1) - u(1, 0)} \right).$$

- If normal player 1 sees $\tilde{y}_0 = 1$, learns other group is good, and plays $y_1 = 1$.
- If normal player 1 sees $\tilde{y}_0 = 0$, posterior belief that other group is bad rises to

$$\mu_1 = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\pi} > \mu_0.$$

- Plays $y_1 = 0$ if and only if $\mu_1 > \mu^*$. Holds if π small.

Equilibrium (continued)

- Equilibrium Suppose up to time t normal players play $y_t = 0$ when $\tilde{y}_{t-1} = 0$.
- Then normal player t 's posterior when $\tilde{y}_{t-1} = 0$ is

$$\mu_t = \frac{\mu_0}{\mu_0 + (1 - \mu_0)(1 - (1 - \pi)^t)}.$$

- Observe that μ_t is decreasing in t , $\mu_t \rightarrow \mu_0$ as $t \rightarrow \infty$, and $\mu_0 < \mu^*$.
- But this implies that there is first time t at which $\mu_t \leq \mu^*$. Call it T .
- Normal player T plays $y_T = 1$ **even if he sees a bad signal**.
- But now normal player $T + 1$ faces same problem as player 1.
- This implies a **cycle of conflict**.

Equilibrium (continued)

Proposition

Assume $\mu_0 < \mu^*$ and $\mu_t \neq \mu^*$ for all t .

Then the baseline model has a unique sequential equilibrium.

It has the following properties:

- At every time $t \neq 0$: mod : T , normal player t plays good actions $(x_t = 1, y_t = 1)$ if she gets the good signal and plays bad actions $(x_t = 0, y_t = 0)$ if she gets the bad signal.
- At every time $t = 0$: mod : T , normal player t plays the good action $x_t = 1$ toward player $t - 1$ if and only if she gets the good signal, but plays the good action $y_t = 1$ toward player $t + 1$ regardless of her signal.
- Bad players always play bad actions $(x_t = 0, y_t = 0)$.

Equilibrium (continued)

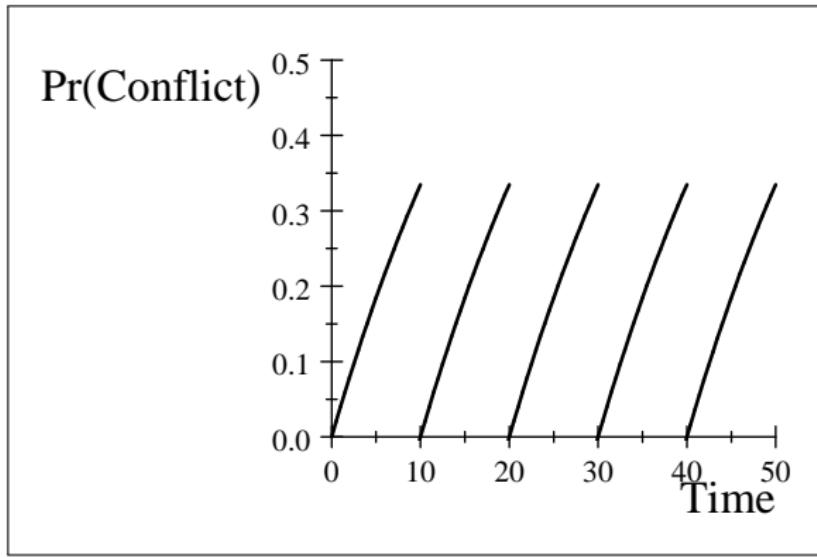


Figure: A Cycle of Conflict

Equilibrium (continued)

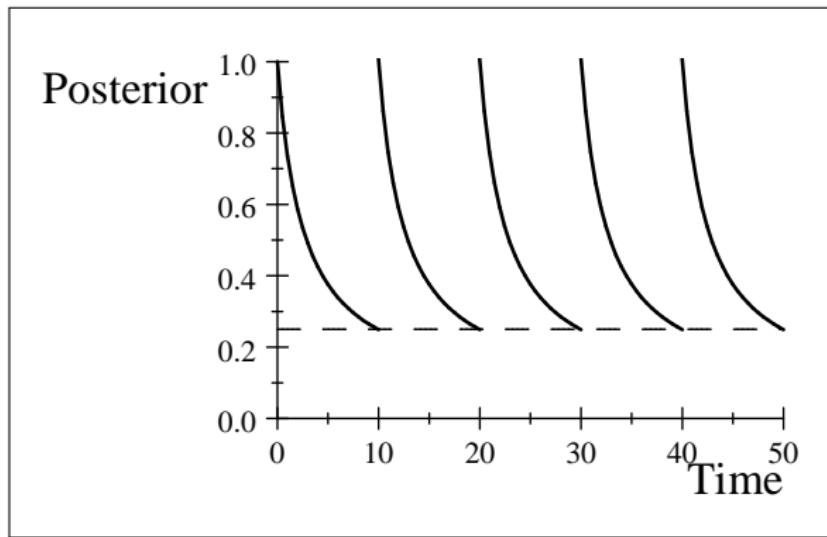


Figure: The Corresponding Cycle of Beliefs

Comparative Statics

Proposition

The cycle period T has the following properties:

- *It is increasing in $u(0, 0)$, decreasing in $u(1, 0)$, and decreasing in $u(1, 1)$.*
- *It is increasing in the prior probability of the bad type μ_0 .*
- *It is decreasing in the error probability π .*

Comparative Statics (continued)

Proposition

Welfare If player t 's payoff is u_t , define social welfare to be

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=0}^N u_t.$$

Suppose both groups are normal. Then:

- *The limit of social welfare as $\pi \rightarrow 0$ is less than the efficient level $2u(1, 1)$.*
- *For any sequence $(\pi_n, \mu_{0,n})$ converging to $(0, 0)$ as $n \rightarrow \infty$, the limit of social welfare as $n \rightarrow \infty$ equals the efficient level $2u(1, 1)$.*
- *The limit of no misperception is not the same as the perfect information game because any conflict lasts so much longer in that limit.*

Conclusion

- Important feedbacks between beliefs and political/public actions.
- Important high-level questions are:
 - Does the presence of political economy lead to biased or less accurate learning/belief formation?
 - Does imperfect information exacerbate political economy conflicts? Does it lead to new types of inefficiencies?
 - Are there feedback cycles leading from bad politics to bad information to bad politics?
 - How can these issues be empirically operationalized?

Endogenous Norms of Cooperation

- Different focus: how does “cooperation” (or “solution to collective action problem”) emerge, and why does “history” affect the outcome of such cooperation games? How and why do norms of corporation change?
 - Why does a history of distrust leads to distrust? How do we understand “social norms” and why do they persist?
 - Why does a society sometime break out of a history of distrust and change social norms?
 - Why does “collective action” differ across societies and why does it seem to change abruptly from time to time?
 - What is the role of leadership and “prominence”?
- Simple model based on Acemoglu and Jackson (2011).

Model

- Consider an overlapping-generations model where agents live for two periods. We suppose for simplicity that there is a single agent in each period (generation), and each agent's payoffs are determined by his interaction with agents from the two neighboring generations (older and younger agents).
- The action played by the agent born in period t is denoted $A_t \in \{H, L\}$, corresponding to "High" and "Low" actions (also can be interpreted as "honest" and "dishonest" actions).
- An agent chooses an action only once, in the first period of his or her life and that is played in both periods. This can be thought of as a proxy for a case where there is discretion, but also a high cost of changing behavior later in life.

Model

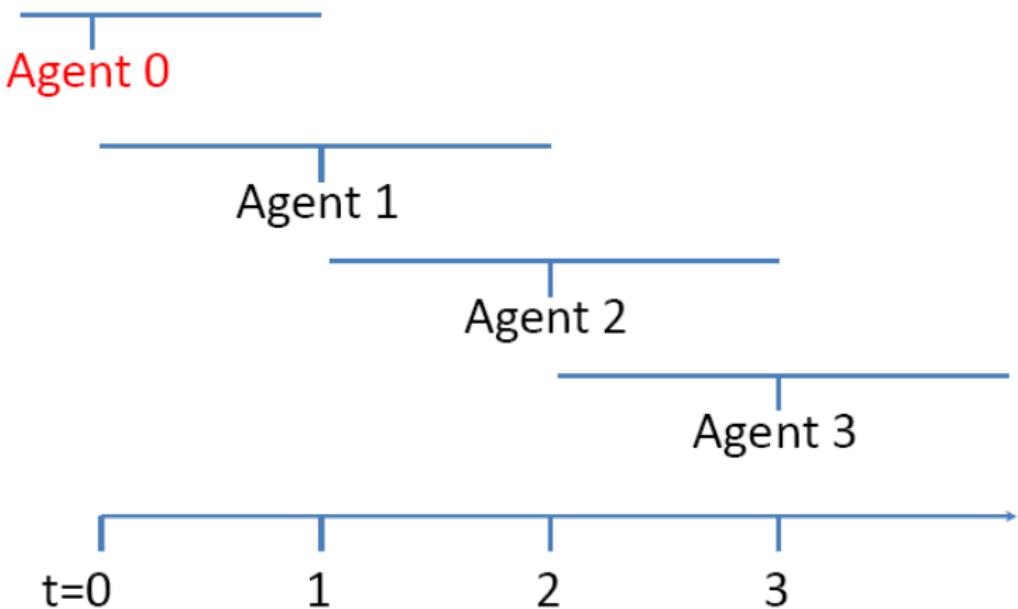
- The stage payoff to an agent playing A when another agent plays A' is denoted $u(A, A')$. The total payoff to the agent born at time t is

$$(1 - \lambda) u(A_t, A_{t-1}) + \lambda u(A_t, A_{t+1}), \quad (2)$$

where A_{t-1} designates the action of the agent in the previous generation and A_{t+1} is the action of the agent in the next generation.

- Implicit assumption: choose single “pattern of behavior” A_t against both generations
 - $\lambda \in [0, 1]$ is a measure of how much an agent weighs the play with the next generation compared to the previous generation.
 - When $\lambda = 1$ an agent cares only about the next generation’s behavior, while when $\lambda = 0$ an agent cares only about the previous generation’s actions. We do not explicitly include a discount factor, since it is subsumed by λ .

Demographics



Model (continued)

- The stage payoff function $u(A, A')$ is given by the following matrix:

	H	L
H	β, β	$-\alpha, 0$
L	$0, -\alpha$	$0, 0$

where β and α are both positive.

- This payoff matrix captures the notion that, from the static point of view, both honesty and dishonesty could arise as social norms—i.e., both (H, H) and (L, L) are static equilibria given this payoff matrix. (H, H) is clearly the Pareto optimal equilibrium, and depending on the values of β and α , it may be the risk dominant equilibrium as well.

Endogenous and Exogenous Agents

- There are four types of agents in this society.
- First, agents are distinguished by whether they choose an action to maximize the utility function given in (2). We refer to those who do so as “endogenous” agents.
- In addition to these endogenous agents, who will choose their behavior given their information and expectations, there will also be some committed or “exogenous” agents, who will choose an exogenously given action.
 - This might be due to some irrationality, or because some agents have a different utility function.

Endogenous and Exogenous Agents (continued)

- Any given agent is an “exogenous type” with probability 2π (independently of all past events), exogenously committed to playing each of the two actions, H and D , with probability $\pi \in (0, \frac{1}{2})$, and think of π as small.
- With the complementary probability, $1 - 2\pi > 0$, the agent is “endogenous” and chooses whether to play H or D , when young and is stuck with the same decision when old.
- Any given agent is also “prominent” with probability q (again independent). Information about prominent agents will be different. Thus:

	non-prominent	prominent
endogenous	$(1 - 2\pi)(1 - q)$	$(1 - 2\pi)q$
exogenous	$2\pi(1 - q)$	$2\pi q$

- Let us refer to endogenous non-prominent agents as *regular*.

Signals, Information and Prominent Agents

- A noisy signal of an action taken by a non-prominent agent of generation t is observed by the agent in generation $t + 1$.
- No other agent receives any information about this action.
- In contrast, the actions taken by prominent agents are perfectly observed by all future generations.

Information Structure

- Let h^{t-1} denotes the public history at time t , which includes a list of past prominent agents and their actions up to and including time $t-1$. We denote the set of h^{t-1} histories by \mathcal{H}^{t-1} .
- We write $h_t = (T, a)$ if at time t the agent is both prominence type $T \in \{P, N\}$ and has taken action $a \in \{H, L\}$ if $T = P$ (if $T = N$, his action is not part of the public history).
- In addition to observing $h^{t-1} \in \mathcal{H}^{t-1}$, an agent of generation t , when born, receives a signal $s_t \in [0, 1]$ about the behavior of the agent of the previous generation (where the restriction to $[0, 1]$ is without loss of any generality). This signal has a continuous and distribution described by a density function $f_H(s)$ if $A_{t-1} = H$ and $f_L(s)$ if $A_{t-1} = L$.

Information Structure (continued)

- Without loss of generality, we order signals such that higher s has a higher likelihood ratio for H ; i.e., so that

$$g(s) \equiv \frac{f_L(s)}{f_H(s)}$$

is nonincreasing in s .

- Suppose also that it is strictly decreasing, so that we have *strict Monotone Likelihood Ratio Property (MLRP)* everywhere.
- Suppose, without loss of any generality, that $s \in [0, 1]$, so that 0 is the worst signal for past H and 1 best signal for past H .
- Let $\Phi(x, s)$ denote the posterior probability that $A_{t-1} = H$ given $s_t = s$ under the belief that an endogenous agent of generation $t - 1$ plays H with probability x . This is:

$$\Phi(x, s) \equiv \frac{f_H(s)x}{f_H(s)x + f_L(s)(1-x)} = \frac{1}{1 + g(s) \frac{1-x}{x}}. \quad (3)$$

The game begins with a prominent agent at time $t = 0$ playing

Strategies

- Let us use N to denote regular agents and P to denote prominent agents.
- With this notation, we can write the strategy of an endogenous agent of generation t (who may or may not be regular) as:

$$\sigma_t : \mathcal{H}^{t-1} \times [0, 1] \times \{P, N\} \rightarrow [0, 1],$$

written as $\sigma_t(h^{t-1}, s, T)$ where $h^{t-1} \in \mathcal{H}^{t-1}$ is the public history of play, $s \in [0, 1]$ is the signal of the previous generation's action, and $T \in \{P, N\}$ denotes whether or not the current agent is prominent.

- The number $\sigma_t(s, h^t, T)$ corresponds to the probability that the agent of generation t plays H .

We denote the strategy profile of all agents by the sequence

$$\sigma = (\sigma_1(h^0, s, T), \sigma_2(h^1, s, T), \dots, \sigma_t(h^t, s, T), \dots).$$

Semi-Markovian Strategies

- For the focus here, the most relevant equilibria involve agents ignoring histories that come before the last prominent agent (in particular, it will be apparent that these histories are not payoff-relevant provided others are following similar strategies).
- Let us refer to these as *semi-Markovian* strategies.
- Semi-Markovian strategies are specified for endogenous agents as functions $\sigma_{\tau}^{SM} : \{H, D\} \times [0, 1] \times \{P, N\} \rightarrow [0, 1]$, written as $\sigma_{\tau}^{SM}(a, s, T)$ where $\tau \in \{1, 2, \dots\}$ is the number of periods since the last prominent agent, $a \in \{H, D\}$ is the action of the last prominent agent, $s \in [0, 1]$ is the signal of the previous generation's action, and again $T \in \{P, N\}$ is whether or not the current agent is prominent.
- Let us denote a semi-Markovian by the sequence $\sigma^{SM} = (\sigma_1^{SM}(a, s, T), \sigma_2^{SM}(a, s, T), \dots, \sigma_t^{SM}(a, s, T), \dots)$.
- With some abuse of notation, write $\sigma_t = H$ or D to denote a strategy (or a semi-Markovian strategy) that corresponds to playing honest (dishonest) with probability one.

Equilibrium Definition

- Perfect Bayesian Equilibrium or Sequential Equilibrium.
- Only need to be careful when $q = 0$.
- Define *greatest* and *least* equilibria, and focus on greatest equilibria.

Towards Equilibrium Behavior

- Let ϕ_{t-1}^t be the probability that the agent of generation t assigns to the agent from generation $t-1$ choosing $A = H$
- Let ϕ_{t+1}^t be the probability that the agent of generation t assigns, conditional on herself playing $A = H$, to the agent from generation $t+1$ choosing $A = H$.
- Payoff from L : 0
- Payoff from H :

$$(1 - \lambda) [\phi_{t-1}^t \beta - (1 - \phi_{t-1}^t) \alpha] + \lambda [\phi_{t+1}^t \beta - (1 - \phi_{t+1}^t) \alpha].$$
- Then an endogenous agent of generation t will prefer to play $A = H$ only if

$$(1 - \lambda) \phi_{t-1}^t + \lambda \phi_{t+1}^t \geq \frac{\alpha}{\beta + \alpha} \equiv \gamma. \quad (4)$$

- Parameter γ a convenient way of summarizing relative payoffs (and also “basin of attraction” of L ; so the greater is γ , the more attractive it is $A = L$).

Cutoff Strategies

- We say that a strategy σ is a *cutoff strategy* if for each t , h^{t-1} such that $h_{t-1} = N$ and $T_t \in \{P, N\}$, there exists s_t^* such that $\sigma_t(h^t, s, T_t) = 1$ if $s > s_t^*$ and $\sigma_t(h^t, s, T_t) = 0$ if $s < s_t^*$.
 - Clearly, setting $\sigma_t(h^t, s, T) = 1$ (or 0) for all s is a special case of a cutoff strategy.
- Cutoff strategy profile can be represented by the sequence of cutoffs

$$c = \left(c_1^N(h_0), c_1^P(h_0), \dots, c_t^N(h_{t-1}), c_t^P(h_{t-1}), \dots \right).$$

- Given strict MLRP, all equilibria will be in cutoff strategies.
- Define greatest equilibria using the Euclidean distance on cutoffs.

Equilibrium Characterization

Proposition

- ① *All equilibria are in cutoff strategies.*
- ② *There exists an equilibrium in semi-Makovian cutoff strategies.*
- ③ *The set of equilibria and the set of semi-Markovian equilibria form complete lattices, and the greatest (and least) equilibria of the two lattices coincide.*

Understanding History-Driven Behavior

- Look for a unique equilibrium given by history:
 - When following prominent H , will all endogenous agents play H ?
 - When following prominent L , will all endogenous agents play L ?
- In such an equilibrium, social norms of High and Low emerge and persist, but not forever, since there might be switches because of exogenous prominent agents.
- Related question: when is this the greatest equilibrium?

Understanding History-Driven Behavior (continued)

- Recall that an endogenous agent of generation t will prefer to play $A = H$ only if

$$(1 - \lambda) \phi_{t-1}^t + \lambda \phi_{t+1}^t \geq \frac{\alpha}{\beta + \alpha} \equiv \gamma. \quad (5)$$

- H is a *unique* best response for all if

$$(1 - \lambda) \phi_{t-1}^t + \lambda \phi_{t+1}^t \geq \gamma$$

$$\gamma_H^* \equiv (1 - \lambda) \Phi(1 - \pi, 0) + \lambda \pi \geq \gamma.$$

- L is a *unique* best response for all if

$$(1 - \lambda) \phi_{t-1}^t + \lambda \phi_{t+1}^t < \gamma$$

$$\gamma_L^* \equiv (1 - \lambda) \Phi(\pi, 1) + \lambda(1 - \pi) < \gamma.$$

Understanding History-Driven Behavior (continued)

Proposition

- ① If $\gamma < \gamma_H^*$, then following $a = H$ at date $t = 0$, the unique continuation equilibrium involves all (prominent and non-prominent) endogenous agents playing H .
- ② If $\gamma > \gamma_L^*$, then following $a = L$ at date $t = 0$, the unique continuation equilibrium involves all endogenous agents playing L .
- ③ If $\gamma_L^* < \gamma < \gamma_H^*$, then there is a unique equilibrium driven by the starting condition: all endogenous agents take the same action as the action of the prominent agent at date $t = 0$.

- Interpretation: persistent, but not everlasting, social norms.

Understanding History-Driven Behavior (continued)

- The condition that $\gamma_L^* < \gamma < \gamma_H^*$ boils down to

$$\lambda(1 - 2\pi) < (1 - \lambda) [\Phi(1 - \pi, 0) - \Phi(\pi, 1)]. \quad (6)$$

- It requires that λ be sufficiently small, so that sufficient weight is placed on the past. Without this, behavior would coordinate with future play, which naturally leads to a multiplicity.
- It also requires that signals are not too strong (so $\Phi(1 - \pi, 0) - \Phi(\pi, 1) > 0$), as otherwise players would react to information about the most recent past generation and could change to High behavior if they had a strong enough signal regarding the past play and would also expect the next generation to have good information.

Understanding History-Driven Behavior (continued)

- Focusing on the greatest equilibrium:
- Let

$$\bar{\gamma}_H \equiv (1 - \lambda) \Phi(1 - \pi, 0) + \lambda (1 - \pi). \quad (7)$$

- Thus relative to γ_H^* , more “optimistic” expectations about the future.

Proposition

The greatest equilibrium is such that:

- (i) *following a prominent play of L, there is a low social norm and all endogenous agents play L if and only if $\bar{\gamma}_L < \gamma$; and*
- (ii) *following a prominent play of H, there is a high social norm and all endogenous agents play H if and only if $\gamma \leq \bar{\gamma}_H$.*

Thus, endogenous players always follow the play of the most recent prominent player in the greatest equilibrium if and only if $\bar{\gamma}_L < \gamma \leq \bar{\gamma}_H$.

General Characterization of Greatest Equilibrium

- Let

$$\hat{\gamma}_H \equiv (1 - \lambda) \Phi(1 - \pi, 1) + \lambda (1 - \pi).$$

- This is the expectation of $(1 - \lambda)\phi_{t-1}^t + \lambda\phi_{t+1}^t$ for an agent who believes that any regular agent preceding him or her played H and sees the most optimistic signal, and believes that all subsequent endogenous agents will play H .
- Above, this threshold, no regular agent would ever play H .

General Characterization (continued)

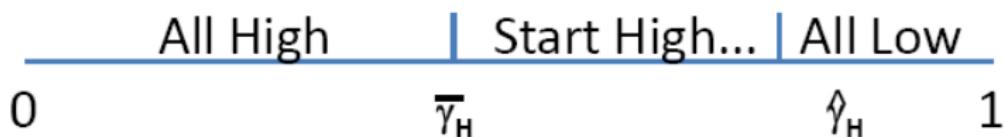
Proposition

In the greatest equilibrium:

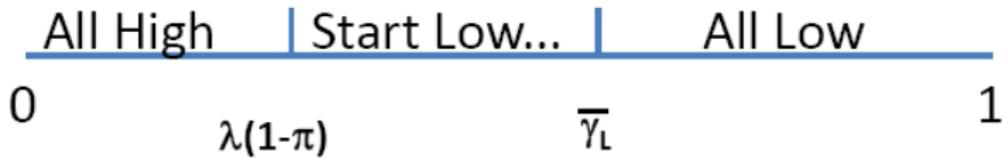
- ① If $\gamma \leq \lambda(1 - \pi)$, then all endogenous agents play H in all circumstances, and thus society has a stable high behavioral pattern.
- ② If $\lambda(1 - \pi) < \gamma \leq \bar{\gamma}_H$, then following a prominent play of H (but not following the prominent play of L) all endogenous agents play H .
- ③ If $\bar{\gamma}_L < \gamma \leq \bar{\gamma}_H$, then following a prominent play of L , all endogenous agents play L , and so all endogenous players follow the play of the most recent exogenous prominent player.
- ④ If $\bar{\gamma}_H < \gamma$, then endogenous agents play L for at least some signals, periods, and types even following a prominent play of H .
- ⑤ If $\hat{\gamma}_H < \gamma$, then all endogenous agents who do not immediately follow a prominent H play L regardless of signals or types.

General Characterization of Greatest Equilibrium (continued)

Last prominent was High



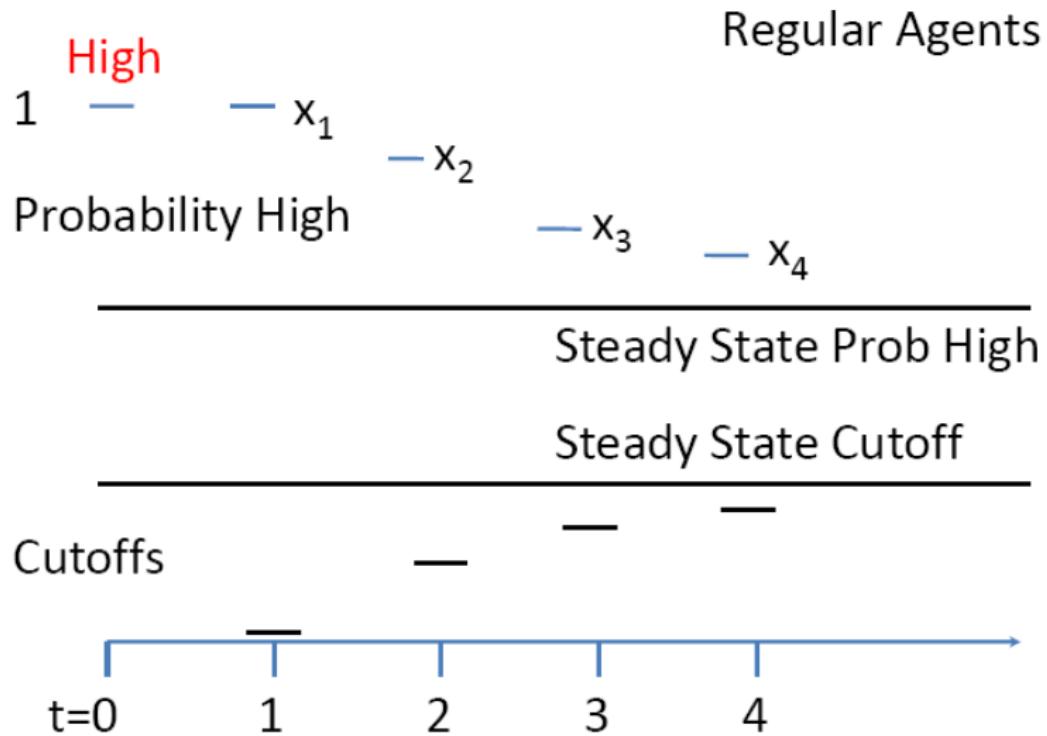
Last prominent was Low



Reversion of Play

- What happens when all High and all Low are not stable social norms?
- Answer: play reverts from an *extreme* (started by a prominent agent) to a steady-state distribution.
 - Start with *H*
 - Next player knows previous is *H* with probability 1
 - Next player knows previous endogenous played *H*, but this has probability $1 - \pi$, so *action made depend on signal*
 - In fact, even stronger, because she knows that her signals will be interpreted as not necessarily coming from *H*.
 - Next player knows previous play was *H* with probability $< 1 - \pi$.

Reversion of Play (continued)



Reversion of Play (continued)

- Let us denote the cutoffs used by prominent and non-prominent agents τ periods after the last prominent agent by c_τ^P and c_τ^N respectively.
- We say that high play *decreases* over time if $(c_\tau^P, c_\tau^N) \leq (c_{\tau+1}^P, c_{\tau+1}^N)$ for each τ .
- We say that high play *strictly decreases* over time, if in addition, we have that when $(c_\tau^P, c_\tau^N) \neq (0, 0)$, $(c_\tau^P, c_\tau^N) \neq (c_{\tau+1}^P, c_{\tau+1}^N)$.
- The concepts of high play increasing and strictly increasing are defined analogously.

Reversion of Play (continued)

Proposition

- ① *In the greatest and least equilibria, cutoff sequences (c_τ^P, c_τ^N) are monotone. Thus, following a prominent agent choosing H , (c_τ^P, c_τ^N) are nondecreasing and following a prominent agent choosing L , they are non-increasing.*
- ② *If $\gamma > \bar{\gamma}_H$, then in the greatest equilibrium, high play strictly decreases over time following high play by a prominent agent.*
- ③ *If $\gamma < \bar{\gamma}_L$, then in the greatest equilibrium, high play strictly increases over time following low play by a prominent agent.*

- But important asymmetry from switching from L to H vs from H to L
 - As we will see next, endogenous prominent agents would not like to the latter, but would prefer to do the former, so the first type of switches are driven by both exogenous and endogenous prominent agents, while the second only by exogenous prominent agents.

Reversion of Play (continued)

- The following is an immediate corollary of Proposition 12.

Corollary

As the distance from the last prominent agent grows ($\tau \rightarrow \infty$), cutoffs in the greatest equilibrium converge and the corresponding distributions of play converge to a stationary distribution. Following a choice of H by the last prominent agent, this limiting distribution involves only H by all endogenous agents if and only if $\gamma \leq \bar{\gamma}_H$. Similarly, following a choice of L by the last prominent agent, this limiting distribution involves L by all endogenous agents if and only if $\gamma \geq \bar{\gamma}_L$.

Leadership: Breaking the Low Social Norm

- Can promise breaker social norm of L ?
 - Regular agents may be stuck in L for reasons analyzed so far.
 - But prominence, greater visibility in the future, can enable “leadership”
- Idea:
 - endogenous prominent agents can always break the social norm
 - when “all L ” is not the unique equilibrium after prominent L , endogenous prominent agents will like to break the social norm of L and start a switch to H

Leadership: Breaking the Low Social Norm (continued)

Proposition

① Suppose that the last prominent agent played L and

$$\tilde{\gamma}_L \leq \gamma < \tilde{\gamma}_H \equiv (1 - \lambda)\Phi(\pi, 0) + \lambda(1 - \pi). \quad (8)$$

Then there exists a fixed cutoff below 1 (after at least one period) such that an endogenous prominent agent chooses *High* and breaks the *Low social norm* if the signal is above the cutoff.

② Suppose that $\gamma < \tilde{\gamma}_L$ and $\gamma < \tilde{\gamma}_H$. Suppose that the last prominent agent has played L . Endogenous prominent agents have cutoffs below 1 that decrease with time such that if the signal is above the cutoff then in a greatest equilibrium the endogenous prominent agent will choose *H* and break a low social norm.

③ Moreover, in either case if $\gamma < \gamma_H^*$, the above are the unique continuation equilibrium.

Role of Prominence

- Prominence provides greater visibility and thus coordinates future actions.
- Crucially: common knowledge of visibility.
 - Without this, prominence is less effective.

Role of Prominence (continued)

- Suppose there is a starting non-prominent agent at time 0 who plays *High* with probability $x_0 \in (0, 1)$, where x_0 is known to all agents who follow, and generates a signal for the first agent in the usual way.
- All agents after time 1 are not prominent.
- In every case all agents (including time 1 agents) are endogenous with probability $(1 - 2\pi)$.

Scenario 1: The agent at time 1 is not prominent and his or her action is observed with the usual signal structure.

Scenario 2: The agent at time 1's action is observed perfectly by the period 2 agent, but not by future agents.

Scenario 2': The agent at time 1 is only observed by the next agent according to a signal, but then is subsequently perfectly observed by all agents who follow from time 3 onwards.

Scenario 3: The agent at time 1 is prominent, and all later agents are viewed with the usual signal structure.

Role of Prominence (continued)

- Clearly, as we move from Scenario 1 to Scenario 2 (or 2') to Scenario 3, we are moving from a non-prominent agent to a prominent one
- Let us focus again on the greatest equilibrium and let $c^k(\lambda, \gamma, f_H, f_L, q, \pi)$ denote the cutoff signal above which the first agent (if endogenous) plays *High* under scenario k as a function of the underlying setting.

Proposition

The cutoffs satisfy $c^2(\cdot) \geq c^3(\cdot)$ and $c^1(\cdot) \geq c^2(\cdot) \geq c^3(\cdot)$, and there are settings $(\lambda, \gamma, f_H, f_L, q, \pi)$ for which all of the inequalities are strict.

Multiple Agents

- Now suppose n agents within each generation, and random matching; unless there is a prominent agent, in which case all those from previous and next generations match with the prominent agent.
- If no prominent agent, then observe a signal generated by the action of a randomly generated agent from the previous generation.
- Results generalize, except but now we can do comparative statics with respect to n .

Multiple Agents (continued)

Proposition

In the model with n agents within each generation, there exist greatest and least equilibria. In the greatest equilibrium:

- ① *following a prominent play of Low, there is a Low social norm and all endogenous agents play Low (i.e., $\bar{\sigma}_\tau^{SM}(a = Low, s, T) = Low$ for all s, T and all $\tau > 0$) if and only if $\bar{\gamma}_L^n < \gamma$; and*
- ② *following a prominent play of High, there is a High social norm and all endogenous agents play High (i.e., $\bar{\sigma}_\tau^{SM}(a = High, s, T) = High$ for all s, T and all $\tau > 0$) if and only if $\gamma \leq \bar{\gamma}_H^n$.*

The threshold $\bar{\gamma}_H^n$ is increasing in n and the threshold $\bar{\gamma}_L^n$ is decreasing in n , so that both High and Low social norms following, respectively, High and Low prominent play, emerge for a larger set of parameter values.

- Intuition: signals less informative, thus history matters more.