Persuading Crowds

Caio Lorecchio
Universitat de Barcelona
School of Economics

EEA - ESEM Congress 2022

Motivation

People follow the wisdom of crowds.

- Consumers buy popular brands because popularity is seen as an indicator of quality.
- High trading volume stocks attract new investors.
- Bank runs arise from a small number of people withdrawing money.

This study examines crowd manipulation through dynamic information disclosure.

Motivation

> *When, however, it is proposed to imbue in the mind of a crowd with ideas and beliefs [...] leaders have recourse to different expedients. The principal of them are three in number and clearly defined - affirmation, repetition and contagion. Their action is somewhat slow, but its effects, once produced, are very lasting.*
>
> — Gustave le Bon, The Crowd: A Study of the Popular Mind

Model

A long-lived principal and an infinite sequence of short-lived agents.

Agents want to match their choices with an unknown (fixed) state.

Principal wants to maximize the (expected, discounted) no. of agents choosing some action.

- Examples: seller/buyers, advisor/investors, central bank/account holders ...

Model

Each agent has three sources of information about the state.

Private: Distributions are conditionally i.i.d, but each one observes his signal realization only. (1)

Public:
- Past actions: $\Rightarrow$ may be informative about past private signals. (2)
- Messages: $\Rightarrow$ principal commits to an information policy. (3)

Principal cannot observe private signals and cannot censor info about past actions.

Main result

Principal has informational advantage over agents... but when it is best to use it?

how to use it?

Main result

Principal has informational advantage over agents... but when it is best to use it?

how to use it?

For a class of private information structure, a characterization:

Social learning is optimal $\Leftrightarrow$ private info unfavorable to principal's preferred action is rare enough.

- Less likely that someone breaks a herd on principal's preferred choice.

Related literature

Many papers discuss how a benevolent planner affects social learning:

- Glazer, Kremer and Perry (2015);Che and Hörner (2017); Smith, Sørensen and Tian (2021).

Few study dynamics when principal and agents have conflicting interests:

- censorship - Sgroi (2002); Nikiforov (2015).
- persuasion in a static environment - Inostroza and Pavan (2022).

I investigate information provision using a dynamic persuasion approach.

- Renault, Solan and Vieille (2014); Ely (2017).

Model

Payoffs

Nature draws a state $\theta := \{H, L\}$ with equal prior probability.

Short-lived agent $t \in \mathbb{N}$ arrives and chooses $a \in \{h, \ell\}$.

He wants to match his action with the unknown state:

$$u(h, H) = u(\ell, L) = 1 > 0 = u(h, L) = u(\ell, H).$$

A long-lived principal gains 1 every time an agent chooses h; 0 otherwise.

Private information structure

Conditionally i.i.d. private signals generate private beliefs q about $\theta = L$.

- Private signal distributions $\rightarrow$ private belief distributions.
- Private beliefs will be conditionally $i.i.d.$ as well.

Assume that the unconditional distribution of private beliefs is absolutely continuous.

- It ensures density g and common support $[\underline{q}, \bar{q}]$.

Assume in addition that private signals are informative: $\underline{q} < 1/2 < \bar{q}$.

The relevant private information structure is $(\underline{q}, \bar{q}, g)$.

- Private beliefs are bounded if $[\underline{q}, \bar{q}] \subset [0, 1]$; unbounded if $[\underline{q}, \bar{q}] = [0, 1]$.

Belief update and action probabilities

Assume away for a moment any intervention from the principal.

Agent t combines private belief q with public belief p_t about $\theta = H$.

He chooses h whenever it yields the highest expected utility, or $q \leq p_t$.

- There will be conditional probabilities of t choosing h under p_t.
- Remark: they are functions of the private information structure $(\underline{q}, \bar{q}, g)$.

Belief update and action probabilities

Given p_t, next period's public belief is a Bayesian inference from the past action:

$$p_{t+1} = \begin{cases} \varphi_\ell(p_t) \leq p_t, \\ \varphi_h(p_t) \geq p_t. \end{cases}$$

Let $\alpha(p_t)$ denote the expected probability of agent t choosing h given p_t.

Example: $q \sim U(\underline{q}, \bar{q})$

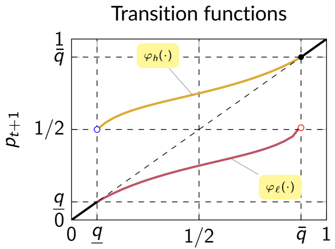

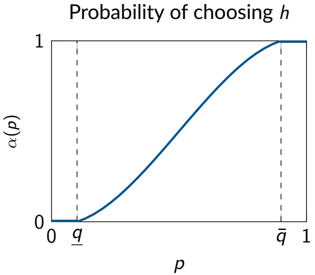

Information policy

Principal *commits* to a public information policy, consisting of

1. a message space;
2. conditional probabilities over this space, for every public history.

Every public history at t leads to a public belief p_t.

Every message realization at t leads to an induced belief ρ_t.

Bayes plausibility: the expected value of ρ_t given public history equals p_t.

Timing

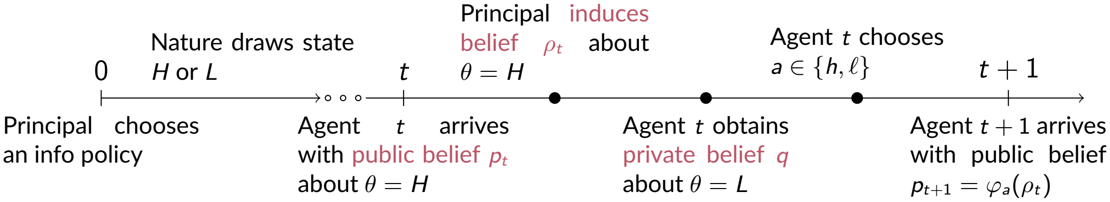

0
Nature draws state H or L
Principal chooses an info policy
t
Agent t arrives with public belief p_t about $\theta = H$
Principal induces belief ρ_t about $\theta = H$
Agent t obtains private belief q about $\theta = L$
Agent t chooses $a \in \{h, \ell\}$
$t + 1$
Agent $t+1$ arrives with public belief $p_{t+1} = \varphi_a(\rho_t)$

Reformulation - direct policies

Any information policy generates stochastic processes $\{\rho_t\}_{t\in\mathbb{N}}$ and $\{p_t\}_{t\in\mathbb{N}}$ satisfying

1. Bayes plausibility: $\mathbb{E}[\rho_t|p_t] = p_t$ for every realization p_t;
2. Laws of motion: $p_{t+1} = \varphi_h(\rho_t)$ with prob. $\alpha(\rho_t)$ or $p_{t+1} = \varphi_\ell(\rho_t)$ with prob. $1 - \alpha(\rho_t)$ for every realization ρ_t.

The converse also holds.

Lemma (Ely, 2017)

Any processes satisfying (1) and (2) can be generated by a direct policy in which (i) the message space is the belief space; (ii) the conditional dist. depend only on the current public beliefs.

Reformulation - principal's problem

For every p_t, the optimal information policy solves

$$V(p_t) = \underset{\tau\,:\,\mathbb{E}_\tau[\rho]=p_t}{Max}\ \mathbb{E}_\tau\left[(1-\delta)\alpha(\rho) + \delta\Big(\alpha(\rho)V(\varphi_h(\rho)) + (1-\alpha(\rho))V(\varphi_\ell(\rho))\Big)\right].$$

A dynamic concavification algorithm, with some particularities:

1. There are multiple laws of motion.
2. Even if the state is fixed, private information generates dynamics.

Main result

Informational cascades

Suppose that $p_t \in [\bar{q}, 1]$.

- Without intervention, agent t chooses h no matter private beliefs.
- Principal has no reason to disclose additional information.

What if $p \leq \underline{q}$?

Proposition

With bounded private beliefs, for any positive public belief $p \leq \underline{q}$, it is optimal to induce beliefs $\rho^- = 0$ and $\rho^+ > \underline{q}$.

Belief convergence

The cascade sets now are $\{0\}$ and $[\bar{q}, 1]$.

The public belief process is a martingale $\Rightarrow$ It converges a.s. to a random variable p_∞.

Proposition

Almost surely $p_\infty \in \{0\} \cup [\bar{q}, 1]$.

Social implications:

- There are social benefits from getting additional information from a conflicted source.
- An info cascade on principal's worst action emerges only under complete learning.

Single disclosure

Here is a simple information policy:

1. for every $p \in [\bar{q}, 1]$, do not disclose any additional information;
2. for every $p \notin \{0\} \cup [\bar{q}, 1]$, induce beliefs 0 and $\bar{q}$.

With it, no agent learns from past actions (no social learning).

This is called the single disclosure policy.

Are there private information structures leading to single disclosure being optimal?

Valuable social learning - bounded private beliefs

Assumption: the unconditional private belief density g is log-concave on $(\underline{q}, \bar{q})$.

Theorem

With bounded private beliefs, single disclosure is optimal if and only

$$\lim_{q \uparrow \bar{q}} g(\bar{q}) \geq \frac{1}{4(1-\bar{q})\bar{q}^2}.$$

Social learning is optimal $\Leftrightarrow$ private info unfavorable to principal's preferred action is rare enough.

Valuable social learning - bounded private beliefs

The theorem follows from two auxiliary results:

Proposition

Single disclosure is optimal if and only if $\alpha(p) \leq V^{sd}(p) \;\; \forall p \in (\underline{q}, \bar{q})$.

- It ensures that I can derive conditions for optimal social learning depending only on α.

Proposition

If g *is log-concave, then the probability* α *of choosing* h *is convex-concave in* $(\underline{q}, \bar{q})$.

The last step to prove the theorem is to use the properties of α being concave near $\bar{q}$.

Valuable social learning - unbounded private beliefs

Theorem

With unbounded private beliefs, single disclosure is never optimal.

- Principal always encourages some social learning in this case.
- When private beliefs are unbounded, principal has less informational advantage.

Conclusion

I study persuasion with social learning in a stylized principal-agents interaction.

- A long-lived principal commits to a public information policy,
- but has to deal with private information and social learning.

Under what circumstances should she encourage social learning?

- For log-concave private belief densities, a tail characterization:
- $\Leftrightarrow$ private info unfavorable to her preferred action is rare enough.

What are the implications for learning in society?

- An info cascade on principal's worst action emerges only under complete learning.