Sequential Group Persuasion
(Extended abstract)

Arjada Bardhi[*] Yingni Guo[†]

Abstract

We study the problem of a biased sender who aims to influence a collective decision taken by a group of decisionmakers through a sequential process. The group members have heterogeneous but correlated types and differ in their thresholds of doubt. The sender, who is perfectly committed, faces each member one at a time in a pre-determined order and designs a sequence of individual-specific information devices that generate action recommendations. Each group member benefits from both private learning from her own device and observational learning from past decisionmakers. We characterize the sender-optimal policy for any level of consensus required by the collective process. In the polar cases of hierarchy (unanimous rule) and polyarchy (dictatorial rule), the optimal policy is order-independent. The hierarchy optimal policy designates a subgroup of rubberstampers, another of perfectly informed members, and a third one of partially manipulated members. We also explore how the optimal policy varies with the group cohesion, defined both as degree of correlation and distribution of thresholds of doubt. Furthermore, we remark on the optimal order of persuasion if the sender can choose whom to persuade first.

Keywords: Information control, collective decision making, sequential evaluation, social learning, hierarchy, polyarchy.

JEL Classification: D71, D72, D83.

1 Motivation and related work

A tremendous share of decision making in the economic and political realm is made within collective schemes rather than by isolated individuals. When deciding as part of a group, members understand the inherent informational and payoff interdependencies among their decisions. Such interdependencies become particularly stark when their decision making is sequential. This paper is concerned with the susceptibility of various structures of sequential collective decision making to information manipulation by a biased sender.

Consider the problem of a project promoter (or sender) facing a sequence of regulators (or receivers). Suppose that the promoter is as uninformed as the regulators about the eventual

[*]Department of Economics, Northwestern University, abardhi@u.northwestern.edu.
[†]Department of Economics, Northwestern University, yingi.guo@northwestern.edu.

success of the project, but nonetheless she prefers an approval to a rejection of the project. Regulators evaluate different but correlated aspects of the project and they also differ on their evaluative standards (i.e. some employs more demanding standards than others). The promoter sequentially provides evidence specific to those aspects to each regulator. Depending on the nature of the regulatory review, it may be that the sender needs the approval of all regulators, of only a single regulator, or of a qualified majority of them. Yet, she has to approach the regulators in an institutionally designated order and present them with evidence. What is the sender's optimal information disclosure policy? How does the probability of approval change with the number of votes needed? How do the regulators' payoffs change with the number of votes needed and the cohesion of the group? Moreover, if the sender were able to choose the order in which she informs the regulators, what would be the optimal order?

Real world instances of a biased sender attempting to persuade a group that decides sequentially are abundant. The Speaker of the House assigns a bill to pass sequentially through various committees in what is known as "sequential referral." The Speaker has considerable leverage in choosing the order of approval, the manner of presentation (i.e. whether to split the bill for consideration by various committees), and deadlines for committee decisions. U.S. presidential candidates face sequential state primaries and caucuses before receiving the official nomination of their party. An inventor interested in geographically extending the protection of his invention might apply for patent protection in multiple territories over time. A CEO who is enthusiastic about a particular project but uncertain of its eventual success has to convince first the board of directors and later a regulator to implement the project. External advocacy groups routinely approach shareholders of a company sequentially in hope of gathering support in changing a particular policy followed by the company. The flow of new ideas within organizations is usually channeled through institutionalized structures of designated approvers. In this sense, some organizations are more hierarchical: the proposal of new ideas has to battle through a sequential order of approvers before being implemented. Other organizations are more polyarchical: it is sufficient for a new idea to gain the support of only one unit of the organization. We can also imagine a mixed version of such polar structures where a new idea needs to be supported by a critical number of units of the organization before implementation.

This paper is closely related to the recent literature on optimal persuasion of multiple receivers. The nature of the information disclosure in our problem extends the framework of Kamenica and Gentzkow (2011) to a sequential setting with many receivers where both informational and payoff externalities are present. The strategic features due to collective decision-making relate our work to that of Alonso and Câmara (2015), Wang (2015) and Chan, Li, and Wang (2015) on persuasion of voters. Yet in contrast to these papers, our model allows the sender to commit to a sequence of voter-specific persuasion mechanisms in a setting with *sequential voting* and *correlated types*. Beyond this persuasion literature, related works in multi-receiver sequential information disclosure are Kremer, Mansour, and Perry (2014) and Che and Hörner (2015). A key difference between our model and theirs is that we assume a biased sender instead of a benevolent social planner. Our focus is on optimal information manipulation serving the interests of the sender rather than optimal social learning.

Our motivating examples are similar to those behind an earlier effort by Caillaud and Tirole (2007) to link persuasion with collective decision making. Unlike in our model, evidence in their framework is perfectly informative, so the sender only chooses which decisionmakers to provide this evidence to. The nature of persuasion is fundamentally different in our work as we allow the sender to vary how informative each regulator's evidence is. Other related literatures include: 1) bandwagoning in sequential voting (Callander (2007), Ali and Kartik (2012)), 2) optimal information acquisition in committees (Gerardi and Yariv (2008), Gershkov and Szentes (2009)), 3) optimal stopping in collective experimentation (Chan et al. (2015), Strulovici (2010)), and 4) project evaluation in various organizational architectures (Sah and Stiglitz (1986), Stiglitz (1988), Sah (1991)). A crucial difference between our model and these papers is that we endogenize the information structure available to the group of decisionmakers.

2 The Model

We consider a communication game between a sender S and a sequence of n regulators $\{R_i\}_{i=1}^n$. These regulators must decide whether to adopt a project sponsored by the sender. Each regulator's payoff from adopting this project depends on his type, and can be either positive or negative.

In particular, before the game starts, nature chooses a state randomly from two alternatives: $\omega \in \{G, B\}$. Let $p_0 = \Pr(\omega = G)$ denote the common prior. After choosing the state, nature chooses randomly and independently the type of each regulator. Regulator R_i's type is denoted by θ_i which can be either high or low, i.e., $\theta_i \in \{H, L\}$. If $\omega = G$, the probability of being H is λ_1, while if $\omega = B$, the probability of being L is λ_0:

$$\Pr(\theta_i = H|\omega) = \begin{cases} \lambda_1 & \text{if } \omega = G, \\ \lambda_0 & \text{if } \omega = B. \end{cases}$$

We assume that $0 \leq \lambda_0 \leq \lambda_1 \leq 1$, so R_i is more likely to be of H type under state G than under state B. In the special case where $\lambda_1 = 1 = 1 - \lambda_0$, the regulators' types are perfectly correlated. At the other extreme, the regulators' types are independent if $\lambda_1 = \lambda_0$. Conditional on the state, the regulators' types are drawn independently. All types are initially unobservable to all parties, but the parameters $p_0, \lambda_1, \lambda_0$ are common knowledge. The prior belief that any regulator is of type H is $p_0\lambda_1 + (1 - p_0)\lambda_0$.

The sender meets the regulators sequentially, so the game lasts for n stages. In stage $i \in \{1, ..., n\}$, the sender meets R_i and designs an information device which generates a signal about R_i's type. We assume that the realized signal from stage i's information device is only observed by R_i. After observing this signal, R_i decides whether to approve or disapprove the project, and this decision is publicly observed by all parties. Then, the game moves to the next stage. We assume that the sender perfectly commits to a sequence of information devices at the beginning of the game.

We let $d_i \in \{1, 0\}$ represent R_i's approval decision. Here, d_i being 1 means that R_i approves and $d_i = 0$ means that he disapproves. Whether the project is eventually adopted depends on the decisions of all regulators. We let $d \in \{1, 0\}$ represent the eventual adoption decision. We consider three collective decision rules: hierarchy, polyarchy and k-majority. If the hierarchy rule is used, the project is adopted if all regulators approve. That is, $d = 1$ if $d_i = 1$ for every i. Under the polyarchy rule, the project is adopted if at least one regulator approves (i.e. $d = 1$ if there exists some i for which $d_i = 1$). Under the k-majority rule, the project is adopted if at least $k \in \{2, ..., n-1\}$ regulators approve (i.e. $d = 1$ if $|\{i : d_i = 1\}| \geq k$).

If the project is eventually adopted, the sender's payoff is 1 which is independent of the state. The payoff of regulator R_i depends on his type θ_i. If $\theta_i = H$, the project yields a payoff of 1 to R_i; if $\theta_i = L$, the project yields a negative payoff $g_i < 0$ to R_i. Here, $g_i < 0$ captures the loss of R_i when the project is adopted when R_i's type is L. If the project is not approved, all parties' payoffs are 0. Therefore, R_i prefers the project to be adopted if and only if his type is H, whereas the goal of the sender is to maximize the probability that the project is adopted. The following assumption focuses our analysis on parameter values for which none of the regulators prefers approval under the prior belief:

Assumption 1. *For any i, g_i is sufficiently low so that R_i cannot be persuaded with probability 1 given the prior belief p_0:*

$$g_i < g^* := \frac{p_0\lambda_1 + (1-p_0)\lambda_0}{\{p_0\lambda_1 + (1-p_0)\lambda_0\} - 1}. \tag{1}$$

At the beginning of stage i, the public history includes the information devices in the first $i-1$ stages and the decisions of the first $i-1$ regulators. Let H^{i-1} stand for the set of histories at the end of stage $i-1$, with h^{i-1} being a typical element. The set of histories before stage 1 is the empty set, i.e., $H^0 = \emptyset$. The sender commits to a communication policy, which in the general setup is a sequence of functions $\{\tilde{M}_i\}_{i=1,\dots,n}$, where $\tilde{M}_i : H^{i-1} \to M^i$ is a mapping from the set of histories H^{i-1} to the set of possible information devices which can be presented to R_i. Without loss, we focus on the canonical information devices which simply make action recommendations. Any canonical information device in stage i can be characterized by two numbers $(\pi_i^H, \pi_i^L) \in [0, 1]^2$. Here, π_i^H is the probability that R_i is recommended to approve if R_i's type is H, and π_i^L is the probability that R_i is recommended to approve if R_i's type is L.

3 Preliminary results

In the following paragraphs we present some of the main results we have obtained so far for the cases of a hierarchy ($k = n$) and a polyarchy ($k = 1$). We are currently working on the case of intermediate voting rules and expect to present results on this case as well.

3.1 Hierarchy $(k = n)$

In the polar case of a perfectly correlated hierarchy, any sequence of optimal information devices is determined by the threshold of the most demanding regulator, i.e.

$$\pi_i^H = 1 \text{ for all } i, \qquad \prod_{i=1}^{n} \pi_i^L = \frac{p_0}{(p_0 - 1)\min_i g_i}.$$

In particular, the policy in which the approval by the strictest regulator is rubberstamped by all other regulators is optimal. In the other polar case of independent types, the optimal policy is the myopic policy that sets all IC constraints binding.

In a hierarchy with imperfectly correlated types, the order in which the regulators are faced by the sender is inconsequential, as it affects neither the probability of final approval nor the optimal information device offered to each regulator. In the optimal policy the sender provides more precise information to more demanding regulators. When a regulator's type is H, his interest and the sender's are aligned and he is recommended to approve with probability 1. On the other hand, the probability of such a recommendation being made in the state at which their interests are misaligned is decreasing in the threshold of doubt of the regulator. The more skeptic a regulator is, the more truthful is the information revealed to her by the information device:

$$\pi_i^H = 1 \text{ for all } i, \qquad \pi_i^L \leq \pi_j^L \text{ for } g_i < g_j.$$

The optimal policy when regulators' types are imperfectly correlated virtually divides the group of regulators into at most three subgroups: the most lenient regulators who rubberstamp, the most demanding regulators who are provided perfectly truthful recommendations, and an intermediate subgroup that is persuaded with some moderate probability. The extreme subgroups (i.e. the most lenient and the most demanding regulators) are left with a positive expected payoff by the optimal policy. This stands in sharp contrast to the case of perfectly correlated regulator types, in which the sender optimally focuses his persuasion efforts only on the most demanding regulator and asks all other more lenient regulators to rubberstamp this decision.

We use our characterization of the optimal policy to draw implications about the distribution of the rejection time for the project. A sequence of regulators that orders them monotonically from the strictest to the most lenient minimizes the expected rejection time. Moreover, when a strict regulator is moved earlier in the decisionmaking sequence by being swapped with a more lenient regulator, the original sequence first order stochastically dominates the new sequence in terms of the distribution of rejection time.

3.2 Polyarchy $(k = 1)$

Interestingly enough, the immateriality of the order of the regulators does not extend to the case of a polyarchy. In order to disentangle the role of the informational and payoff externalities simultaneously present in our problem, we study separately the cases of a polyarchy with and without payoff externalities.[1] In the latter case, each regulator receives an outside option normalized at zero upon rejecting the project, while in the former, a regulator who rejects the project receives the payoff from the continuation game involving latter regulators. We find that when payoff externalities are absent, a sender who faces perfectly correlated regulators chooses to persuade only the most lenient of them and ignores the rest of the group– this insight is independent of the order. But when types are imperfectly correlated, the sender does not necessarily find it optimal to maximize the probability that the most lenient regulator approves the project and to provide $\pi_i^H = 1$ for such a regulator. We provide examples in which the sender rationally sets $\pi_i^H < 1$ for earlier lenient regulators as caution against the negative information externality on later regulators' decisions generated by their rejections. The probability of a collective approval by the group is highest when the regulators are ordered from the most lenient to the most demanding.

The introduction of payoff externalities to the polyarchy necessitates taking a stance on off-equilibrium conjectures held by regulators. We employ a notion of trembling-hand perfection, in which a regulator slightly trembles when recommended approval (or rejection) with certainty regardless of his type. This refinement essentially makes every regulator pivotal with strictly positive probability, building a bridge with our results for a polyarchy without payoff externalities. It continues to be the case that when the regulators are perfectly correlated, the sender finds it optimal to persuade only the most lenient regulator and ignore all other regulators, despite the order in which she faces them. When the types are imperfectly correlated, the approval IC constraints for all regulators bind, leaving each of them with an expected payoff of zero.

4 Bibliography

Ali, S. Nageeb and Navin Kartik. "Herding with collective preferences." *Economic Theory* 51 (2012): 601-626.

Alonso, Ricardo and Odilon Camara. "Persuading voters." Working paper. November, 2015.

Caillaud, Bernard and Jean Tirole. "Consensus building: How to persuade a group." *The American Economic Review* 97.5 (2007): 1877-1900.

Callander, Steven. "Bandwagons and Momentum in Sequential Voting." *The Review of Economic Studies* 74.3 (2007): 653-684.

Chan, Jimmy, Fei Li, and Yun Wang. "Discriminatory persuasion: How to convince a group." Working Paper. November 2015.

Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leeat Yariv. "Deliberating collective deci-

[1] Notice that the case with payoff externalities is the relevant case we build upon when moving to k-majority.

sions." Working Paper. December 2015.

Che, Yeon-Koo and Johannes Hörner. "Optimal design for social learning." Working Paper. April 2015.

Gerardi, Dino and Leeat Yariv. "Information acquisition in committees." *Games and Economic Behavior* 62.2 (2008): 436-459.

Gershkov, Alex and Balazs Szentes. "Optimal voting schemes with costly information acquisition" *Journal of Economic Theory* 144 (2009): 36–68.

Kamenica, Emir and Matthew Gentzkow. "Bayesian persuasion." *The American Economic Review* 101.6 (2011): 2590-2615.

Kremer, Ilan, Yishay Mansour, and Motty Perry. "Implementing the 'Wisdom of the Crowd'." *Journal of Political Economy* 122.5 (2014): 988-1012.

Sah, Raaj Kumar. "Fallibility in Human Organizations and Political Systems." *Journal of Economic Perspectives* 5.2 (1991): 67-88.

Sah, Raaj Kumar and Joseph E. Stiglitz. "The architecture of economic systems: Hierarchies and polyarchies." *The American Economic Review* 76.4 (1986): 716-727.

Sah, Raaj Kumar and Joseph E. Stiglitz. "Committees, hierarchies, and polyarchies." *The Economic Journal* 98 (1988): 451-470.

Strulovici, Bruno. "Learning while voting: Determinants of collective experimentation." *Econometrica* 78.3 (2010): 933-971.

Wang, Yun. "Bayesian persuasion with multiple receivers." Working Paper. February 2015.