Title
Information and Political Economy: Two Illustrations

Author
Titova, Mariia

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Information and Political Economy: Two Illustrations

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Mariia Titova

Committee in charge:

Professor Renee Bowen, Co-Chair
Professor Joel Sobel, Co-Chair
Professor Simone Galperti
Professor Sebastian Saiegh
Professor Joel Watson

2021

The Dissertation of Mariia Titova is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

DEDICATION

To my mother, my friends, and my family.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I cannot begin to express my gratitude to my advisors Renee Bowen and Joel Sobel, but I will try. I thank Renee for her unconditional support and unshakable belief in me. She has always been an empathetic and compassionate mentor and made me a better researcher, a better presenter, and a better person. I thank Joel for his seemingly infinite patience and the numerous hours he spent answering my questions and reading my drafts. He has always encouraged me to be an expert in my field and taught me that I could achieve that goal through persistence and perseverance.

I am grateful to my family for never doubting my ability to excel academically and always encouraging me to aim for the stars.

I have made amazing friends during my studies. Olya, Sasha, Alex, Connor, Wes, these six years would have been bleak without you.

I am grateful to my committee members Simone Galperti, Sebastian Saiegh, and Joel Watson for their encouragement, feedback, and pushing me to get better.

It was challenging to stay optimistic during the worst academic job market to date. I thank Misha Freer, Teevrat Garg, and Denis Shishkin, who kept me sane and helped me prepare my job applications.

I express my gratitude to the UCSD Department of Economics for the funding and hospitality during my graduate studies. I am particularly grateful to Marjorie Flavin, who made me an offer of admission despite my unusual application package, and Prashant Bharadwaj, who helped me navigate the administrative side of my transition from coursework to research.

Finally, I would like to thank my teachers and mentors at the Higher School of Economics. I am particularly grateful to Fuad Aleskerov, who initially spiked my interest in economic theory, Konstantin Sonin, who inspired me to study political economy, Alexei Parakhonyak, who coauthored my first paper, and Alexei Nikitin, who mentored me through my Master's studies.

Chapters 1 and 2 of this dissertation are currently being prepared for submission for publication of the material. The dissertation author, Mariia Titova, is the sole author of this material.

VITA

2009–2013	Bachelor, Higher School of Economics
2013–2015	Master, Higher School of Economics
2016–2021	Teaching Assistant, Department of Economics University of California San Diego
2021	Doctor of Philosophy, University of California San Diego

PUBLICATIONS

"Shopping Malls, Platforms and Consumer Search" with Alexei Parakhonyak, *International Journal of Industrial Organization*, vol. 58, 2018, pp. 183–213, 2018

ABSTRACT OF THE DISSERTATION

Information and Political Economy: Two Illustrations

by

Mariia Titova

Doctor of Philosophy in Economics

University of California San Diego, 2021

Professor Renee Bowen, Co-Chair
Professor Joel Sobel, Co-Chair

This dissertation examines strategic settings in which agents have imperfect information. In the first chapter, an informed agent decides how to influence an uninformed decision-maker. In the second chapter, a group of agents decides how to learn. Both chapters discuss how these models can be applied in a political economy setting to study how politicians persuade voters and how policymakers identify the best available policies.

Chapter 1 studies persuasion with verifiable information. An informed sender with state-independent preferences sends private verifiable messages to multiple re-

ceivers attempting to convince them to approve a proposal. I find that every equilibrium outcome is characterized by each receiver's set of approved states that satisfies this receiver's obedience and the sender's incentive-compatibility constraints. That allows me to characterize the full equilibrium set. The sender-worst equilibrium outcome is one in which information unravels, and receivers act as if under complete information. The sender-preferred equilibrium outcome is the commitment outcome of the Bayesian persuasion game. In the leading application, I study targeted advertising in elections and show that by communicating with voters privately, a challenger may win elections that are unwinnable with public disclosure. As the electorate becomes more polarized, the challenger can swing unwinnable elections by targeted advertising with a higher probability.

Chapter 2 studies a model of costly sequential search among risky alternatives performed by a group of agents. The learning process stops, and the best uncovered option is implemented when the agents unanimously agree to stop or when all the projects have been researched. Both the implemented project and all the information gathered during the search process are public goods. I show that the equilibrium path implements the same project based on the same information gathered in the same order as the social planner. At the same time, due to free riding, search in teams leads to a delay at each stage of the learning process, which grows with search costs.

Chapter 1

Persuasion with Verifiable Information

1.1 Introduction

Suppose the sender attempts to convince a group of receivers to take his favorite action. The only tool available to him is hard evidence. What he can do is choose how much of it to reveal. On average, what is the best outcome that the sender can hope for?

Persuasion with verifiable information plays an essential role in electoral campaigns, product advertising, financial disclosure, and job market signaling, among many other economic situations. In politics, a challenger convinces voters to elect him over the status quo by sending fact-checked ads about his policy position on some relevant socio-economic issues, saying nothing about other issues. In business, a firm convinces consumers to adopt its product by advertising some product characteristics, not mentioning others. In finance, a CEO convinces the board of directors to approve managerial compensation by presenting some financial indicators and statements, omitting others. In labor markets, a job candidate convinces committee members to offer him a job by attaching to his application selected evidence of his qualifications.

I consider the following formal model of persuasion with verifiable information. There is an underlying continuous space of possible states of the world, which is a unit interval. The sender is fully informed about the state of the world, but his preferences do not depend on it. Receivers are uninformed about the state of the world, which to them is payoff-relevant. The sender sends a verifiable message to each receiver. Verifiability means that the message contains the truth (hard evidence is presented), but it could be vague (not all the evidence is presented). Each receiver independently chooses between two options: to approve the proposal or to reject it. There are no information spillovers between the receivers: each receiver only hears her own private message.

How does the sender convince one receiver with verifiable information? Rather than looking at the sender's messages and the receiver's beliefs, I focus on what the receiver does in every state of the world. Since she chooses between two options, we can partition the state space into two subsets: the set of approved states and the set of rejected states. My first result states that a subset of the unit interval is an equilibrium set of approved states if and only if it satisfies two constraints. Firstly, the *sender's incentive-compatibility constraint* (IC) ensures that the sender does not wish to deviate toward a fully informative strategy that induces the receiver to act as if fully informed. Secondly, the *receiver's obedience constraint* ensures that the receiver approves the proposal whenever her expected net payoff of approval is non-negative.

In the sender's least preferred equilibrium, his ex-ante odds of approval are minimized across all equilibria. The receiver learns whether the state of the world is within her complete information approval set and makes a fully informed choice.

In the sender's most preferred equilibrium, his odds of approval are maximized subject to the receiver's obedience constraint. In the sender-preferred equilibrium, the receiver approves the proposal whenever her net payoff of approval is sufficiently high, but possibly negative. That is, the sender improves his odds of approval upon full disclosure by convincing the receiver to approve when she prefers not to.

In his most preferred equilibrium, the sender pulls the "good" states that the receiver prefers to approve and the "bad" states that the receiver prefers to reject. The solution is characterized by a cutoff value: the receiver approves every state that is not too "bad". When the receiver approves, her obedience constraint binds, and she is indifferent between approval and rejection. The sender improves his ex-ante payoff over full disclosure because the receiver approves some of the "bad" states. In fact, in his most preferred equilibrium, the sender reaches the commitment payoff. This observation bridges the gap between the verifiable information literature and the Bayesian persuasion Kamenica and Gentzkow, 2011. The sender need not benefit from having ex-ante commitment power and can persuade the receiver with verifiable messages.

With many receivers, I get similar results. Every receiver makes a fully informed choice in the sender-worst equilibrium, and the sender-preferred equilibrium outcome is a commitment outcome.

Swinging Elections

Targeted advertising played an important role in the recent US Presidential Elections. In 2016, the Trump campaign used voter data from Cambridge Analytica to target voters via Facebook and Twitter. In 2008, the Obama campaign pioneered the use of social media to communicate with the electorate. Even before social media, in 2000, The Bush campaign targeted voters via direct mail. Given that the winning candidate had access to better technology or better voter data in all these cases, one may wonder whether targeted advertising was why these candidates won.[1] In other words, can targeted advertising swing electoral outcomes?

To answer that question, I apply my model to study elections. The state space is now a one-dimensional policy space with positions ranging from ultra-left (0) to ultra-right (1). The voters choose between the challenger, whose policy is unknown, and the status quo policy, which is fixed and known. Each voter prefers to vote in favor of the policy that is closest to her bliss point. In his electoral campaign, the challenger

[1]For comparison of advertising strategies between the candidates, see Kim et al. (2018) and Wylie (2019) for the 2016 election, Harfoush (2009) and Katz, Barris, and Jain (2013) for 2008, and Hillygus and Shields (2014) for 2008.

sends verifiable messages to the voters to inform them about his policy and convince them to elect him.

Suppose that winning an election requires convincing two voters, L and R, whose bliss points are located to the left and the right of the status quo policy, respectively. Observe that unless the challenger can privately advertise to each of these voters, he always loses this election. As long as these voters hold a common belief, which they do under full disclosure, no disclosure, or public disclosure by the challenger, only one of these voters expects the challenger's policy to be closer to her bliss point than the status quo. I call this election unwinnable for the challenger. Whether an election is unwinnable depends on the institution (the social choice function) and the ideology of the electorate (bliss point of the voters). For example, under the majority rule, I show that an election is unwinnable if and only if the status quo is the median voter's bliss point.

When the challenger has access to targeted advertising, he can tell different things to different voters. Recall that in his most preferred equilibrium, the sender improves his odds of approval upon full disclosure. In particular, the challenger manages to convince voter L (R) even when his policy is slightly to the right (left) of the status quo. Consequently, he can convince both voters at the same time and win unwinnable elections with positive probability. That said, the challenger only benefits from private communication if his policy is sufficiently close to the status quo: the further to the right (left) his policy is, the harder it becomes to convince voter L (R).

When a voter's bliss point moves away from the status quo, she becomes less satisfied with the status quo, and that makes her more persuadable. Consequently, when the electorate becomes more polarized, which happens when one of the voters' positions becomes more extreme, the challenger has higher odds of swinging an unwinnable election. As voter R's position moves further to the right, she becomes more persuadable also by policies further to the left of the status quo. Consequently, when voter R's bliss point shifts to the right, the challenger-preferred set of approved policies shifts to the left, toward the policies preferred by the less extreme voter L.

Related Literature

I assume that the sender uses hard evidence to communicate with the receivers. This verifiable information communication protocol was introduced by Milgrom (1981) and Grossman (1981). Other communication protocols include cheap talk by Crawford and Sobel (1982) and Bayesian persuasion by Kamenica and Gentzkow (2011). Relative to these other models of communication, Bayesian persuasion makes the sender better off because it endows him with ex-ante commitment power. Lipnowski and Ravid (2020) find that the sender's maximal equilibrium payoff from cheap talk is generally strictly lower than his payoff under commitment. Consequently, a cheap-talk sender values commitment.[2] In contrast to their result, I show that the sender does not necessarily benefit from commitment if he possesses the hard evidence to verify his messages.

There is extensive literature on applications of Bayesian persuasion models. It includes settings in which schools persuade employers to hire their graduates (Ostrovsky and Schwarz, 2010; Boleslavsky and Cotton, 2015); pharmaceutical companies persuade the FDA to approve their drug (Kolotilin, 2015); matching platforms persuade sellers to match with buyers (Romanyuk and Smolin, 2019); politicians persuade voters (Alonso and Câmara, 2016; Bardhi and Guo, 2018); governments persuade citizens through media (Gehlbach and Sonin, 2014; Egorov and Sonin, 2019). My contribution states that in all these applications, one can replace the assumption that the sender has commitment power with the assumption that the sender has hard evidence.

The leading application contributes to the growing literature on voter persuasion.

[2] Lipnowski (2020) also notes that the sender reaches the commitment outcome with cheap talk if his value function is continuous in the receiver's posterior belief. That assumption is very restrictive: when receivers choose between two options and the sender's preferences are state-independent, the sender's value function must be constant, meaning that no communication takes place under cheap talk, verifiable information, and Bayesian persuasion. I thank Elliot Lipnowski for this insight.

My results are in line with the recent findings in the information design literature on the private persuasion of strategic voters. In particular, Chan et al. (2019) confirm that the politician does better when private disclosure is allowed, and Heese and Lauermann (2019) confirm that the politician needs very little commitment power to achieve the desired outcome. In the verifiable information literature, electoral competition usually results in the full unraveling of information (Board, 2009; Janssen and Teteryatnikova, 2017; Schipper and Woo, 2019) because the candidates play a zero-sum game, and that pushes them to disclose all information voluntarily. In contrast to these papers, I consider a non-symmetric model in which one candidate has a significant advantage over his opponent in that he is the only one who can communicate with the voters. Unraveling does not necessarily occur, and the challenger can improve his odds of winning over full disclosure.

The leading application sheds more light on how political advertising, especially targeted advertising, affects electoral outcomes and why it has become widespread. DellaVigna and Gentzkow (2010) and Prat and Strömberg (2013) provide excellent surveys of the evidence of voter persuasion. First, candidates target their ads based on voters' positions on the political spectrum (George and Waldfogel, 2006; DellaVigna and Kaplan, 2007). Second, one can make a case that an increase in the availability of information catered toward certain electoral groups also counts as targeted advertising because these are the messages intended for and heard by these groups (Oberholzer-Gee and Waldfogel, 2009; Enikolopov, Petrova, and Zhuravskaya, 2011). I show that targeted political advertising may be so widespread because it allows politicians to win elections that are unwinnable otherwise.

I also contribute to the growing literature on polarization and targeted political advertising through media. As the number of media outlets increases, they become more specialized and target voters with more extreme preferences, which leads to social disagreement (Perego and Yuksel, 2018). If the electorate is polarized to begin with, so are the candidates' chosen policy platforms (Hu, Li, and Segal, 2019; Prummer, 2020). Abstracting away from candidates choosing their policies, I find that as the electorate becomes more polarized, more challengers can swing elections that are unwinnable otherwise.

This chapter is organized as follows. Section 1.2 introduces the model. Section 1.3 describes equilibrium outcomes in the game with one receiver. Section 1.4 generalizes the model to many receivers. Section 1.5 studies targeted advertising in elections. Section 1.6 is a conclusion.

1.2 Model

There is a state space $\Omega := [0,1]$ and a finite set of receivers $I := \{1,\ldots,n\}$. The game begins with the sender (him) observing the realization of the random state $\omega \in \Omega$, which is drawn from an atomless common prior distribution $p > 0$ over Ω.[3] Having observed the state, the sender sends a verifiable message $m_i \subseteq \Omega$, such that $\omega \in m_i$, to each receiver (her) $i \in I$.[4]

The sender's payoff $u_s : 2^n \to \mathbb{R}$ depends only on the subset of receivers who approve his proposal. I assume that if all receivers reject the proposal, then the sender gets the lowest payoff, which is normalized to 0. If every receiver approves the proposal, then the sender gets the highest payoff, which is normalized to 1. Also, I assume that u_s weakly increases in every receiver's action.

Assmption 1.1. The sender's payoff u_s satisfies

[3]For a compact metrizable space S, ΔS denotes the set of all Borel probability measures over S. For any $q \in \Delta\Omega$ and any measurable subset of the state space $W \subseteq \Omega$, $Q(W) = \int_W q(\omega)d\omega$ is the probability measure and $q(\cdot \mid \cdot)$ is the conditional probability distribution: $q(\omega \mid W) = 1$ if $W = \{\omega\}$ and $q(\omega \mid W) = \frac{q(\omega)}{Q(W)}$ if $Q(W) > 0$.

[4]I borrow the definition of a verifiable message as a subset of the state space that includes the true realization from Milgrom and Roberts (1986). This method satisfies normality of evidence (Bull and Watson, 2004), which means that it is consistent with both major ways of modeling hard evidence in the literature.

1. $u_s(\varnothing) = 0$ and $u_s(I) = 1$;

2. given two sets of receivers $I_1, I_2 \subseteq I$, $u_s(I_1) \leq u_s(I_2)$ if $I_1 \subseteq I_2$.

Receiver $i \in I$ chooses between approval (action 1) and rejection (action 0). Receiver i's preferences are described by a utility function $u_i : \{0,1\} \times \Omega \to \mathbb{R}$. Receiver i approves (the proposal in) state ω if her net payoff of approval $\delta(\omega) := u_i(1,\omega) - u_i(0,\omega)$ is non-negative.[5] Define receiver i's approval set as

$$\mathcal{A}_i := \{\omega \in \Omega \mid \delta_i(\omega) \geq 0\}.$$

Example 1.1 (Receiver with Spatial Preferences)**.** This example introduces the receivers with spatial preferences à la Downs (1957). Receiver i has a bliss point $v_i \in \Omega$ and compares the sender's position ω to the status quo $\omega_0 \in (0,1)$. Her net payoff of approval is $\delta_i(\omega) = -|v_i - \omega| + |v_i - \omega_0|$ and her approval set is $\mathcal{A}_i = \{\omega \in \Omega \text{ s.t. } |v_i - \omega| \leq |v_i - \omega_0|\}$. That is, she approves ω if and only if it is closer to her bliss point than the status quo.

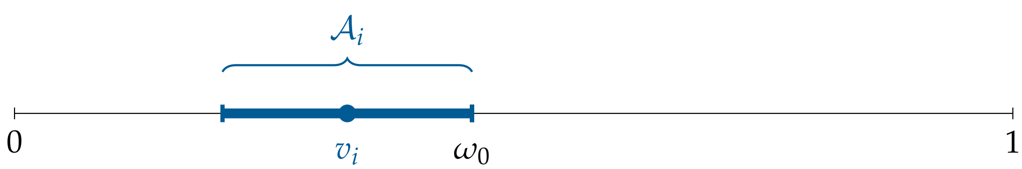

Figure 1.1. Receiver i with spatial preferences: her approval set $\mathcal{A}_i$ (solid blue) consists of points on the unit interval that are closer to her bliss point v_i than the status quo ω_0.

Under incomplete information, define receiver i's set of approval beliefs as

$$\mathcal{B}_i := \{q \in \Delta\Omega \mid \mathbb{E}_q[\delta_i(\omega)] \geq 0\}.$$

I assume that every receiver rejects the proposal under prior belief.

[5] I assume that the receiver breaks ties in favor of approval when she is indifferent, i.e. when $\delta(\omega) = 0$.

Assmption 1.2. For every receiver $i \in I$, $p \notin \mathcal{B}_i$.

Assumptions 1.1 and 1.2 ensure that without any additional information, all receivers reject the proposal and the sender gets the lowest possible payoff. The rest of the chapter studies how the sender persuades the receivers with verifiable information.

Equilibrium Outcomes

I consider Perfect Bayesian Equilibria (henceforth just *equilibria*) of this game. The sender's strategy is a probability distribution $\sigma(\cdot \mid \omega)$ over message collections $\{m_i\}_{i \in I}$, where $m_i \subseteq \Omega$ for each $i \in I$. Receiver i's approval strategy $a_i(m)$ specifies which action she takes depending on message m she receives. Receiver i's posterior belief over Ω after message m is $q_i(\cdot \mid m)$. Profiles of receivers' actions and posterior beliefs are $a := \{a_i\}_{i \in I}$ and $q := \{q_i\}_{i \in I}$, respectively.

Definition 1.1. *A triple (σ, a, q) is an equilibrium if*

(i) $\forall \omega \in \Omega$, $\sigma(\cdot \mid \omega)$ *is supported on* $\arg\max_{m_1,\dots,m_n} u_s\big(\{i \in I \mid a_i(m_i) = 1\}\big)$, *s.t.* $\omega \in m_i$, $\forall i \in I$.

The following conditions hold for every receiver $i \in I$:

(ii) $\forall m \subseteq \Omega$, $a_i(m) = \mathbb{1}(q_i(\cdot \mid m) \in \mathcal{B}_i)$;

(iii) $\forall m \subseteq \Omega$ *such that* $\int_{\Omega} \sigma_i(m \mid \omega) d\omega > 0$, $q_i(\omega \mid m) = \frac{\sigma_i(m \mid \omega) \cdot p(\omega)}{\int_{\Omega} \sigma_i(m \mid \omega') \cdot p(\omega') d\omega'}$, *where σ_i is the marginal distribution of messages heard on the equilibrium path by receiver i;*

(iv) $\forall m \subseteq \Omega$, $\operatorname{supp} q_i(\cdot \mid m) \subseteq m$.

In words, *(i)* states that the sender sends a collection of messages with positive probability only if it maximizes his payoff; *(ii)* states that each receiver approves the proposal whenever her expected net payoff of approval is non-negative under her posterior belief; *(iii)* states that receivers' posterior beliefs are Bayes-rational on the equilibrium path; *(iv)* states that the receivers' posterior beliefs on and off the path are concentrated on the states in which the message is available to the sender.

An outcome of the game specifies what action receivers take in every state of the world.

Definition 1.2.

- *An outcome* $\alpha = \{\alpha_i\}_{i \in I}$ *specifies* $\forall i \in I$ *and* $\forall \omega \in \Omega$ *the probability* $\alpha_i(\omega) \in [0,1]$ *that receiver* i *approves the sender's proposal in state* ω.

- *An outcome is an equilibrium outcome if it corresponds to some equilibrium.*[6]

Some outcomes are deterministic, meaning that in every state ω each receiver either approves or rejects the proposal with certainty.[7] Consequently, for each receiver, we can partition Ω into states of approval and states of rejection.

Definition 1.3.

- *An outcome* α *is deterministic if* $\alpha_i(\omega) \in \{0,1\}$ *for every* $i \in I$ *and* $\omega \in \Omega$.

- *The set of approved states* W_i *of receiver* $i \in I$ *in deterministic outcome* α *is*

$$W_i := \{\omega \in \Omega \mid \alpha_i(\omega) = 1\}.$$

1.3 One Receiver

Let us first focus on the case with one receiver, i.e. $I = \{1\}$. For ease of exposition, I drop all receiver-relevant subscripts i. By Assumption 1.1, the sender gets 1 if the receiver approves and 0 otherwise. By Assumption 1.2, the receiver rejects the proposal under the prior belief.

[6]Specifically, if there exists equilibrium (σ, a, q) such that $\forall i \in I$ and $\forall \omega \in \Omega$, $\alpha_i(\omega) = \int_{\mathcal{M}_i} \sigma_i(m \mid \omega) dm$, where $\mathcal{M}_i := \{m \subseteq \Omega \mid a_i(m) = 1\}$ is the set of messages that convince receiver i to approve.

[7]Although each receiver breaks ties in favor of approval, the sender may be playing a mixed strategy in state ω, and then in that state the receiver may be approving the proposal with a probability between 0 and 1.

1.3.1 Direct Implementation

Consider a deterministic equilibrium outcome with a set of approved states W. Suppose that the sender learns that $\omega \in \mathcal{A}$. One message that is available to the sender in this state (and unavailable in every other state) is $\{\omega\}$. Since that message is verifiable, upon receiving it, the receiver learns with certainty that the state is ω. Since ω is in the receiver's approval set, she approves the proposal after hearing that message. Then, for every $\omega \in \mathcal{A}$, the receiver should be approving every $\omega \in \mathcal{A}$ in every deterministic equilibrium, or else the sender has a profitable deviation towards full disclosure. That gives rise to the sender's incentive-compatibility constraint

$$\mathcal{A} \subseteq W. \tag{IC}$$

Next, if the receiver approves every state in W, then she expects that on average, her net payoff of approval is non-negative. Thus, we obtain the receiver's obedience constraint

$$p(\cdot \mid W) \in \mathcal{B}. \tag{obedience}$$

The first result of this chapter allows us to restrict attention to sets of approved states $W \subseteq \Omega$ that satisfy these two constraints.

Theorem 1.1. *Suppose $n = 1$. Then, every equilibrium outcome is deterministic. Furthermore, $W \subseteq \Omega$ is an equilibrium set of approved states if and only if it satisfies the sender's* (IC) *and the receiver's* (obedience) *constraints.*

The proofs of Theorem 1.1 and other results are in the appendix. Here I describe the intuition behind this result. First, in every equilibrium outcome, the receiver either approves or rejects the proposal in every state of the world. Suppose, on the contrary, that in some state, the receiver approves and rejects with positive probability. Since the receiver approves sometimes, the sender has access to at least one message that convinces the receiver to approve. Then, the sender can deviate and send that message with certainty so that the receiver approves with probability one. Hence, all equilibrium outcomes are deterministic.

Next, if W is an equilibrium set of approved states, it satisfies the sender's (IC) constraint, or else the sender can deviate to full disclosure. To see why W also satisfies the receiver's (obedience) constraint, implement this set of approved states directly. Specifically, let the sender send message W from $\omega \in W$ and message $\Omega \setminus W$ from $\omega \notin W$. Intuitively, the (obedience) constraint states that the receiver interprets message W as a recommendation to approve. It holds because if the sender induces approval in every state in W in the original equilibrium, he also induces approval with the pooling message W.

Finally, suppose that $W \subseteq \Omega$ satisfies (IC) and (obedience). Then, we can construct an equilibrium that directly implements the set of approved states W. Let the sender send message W from every state within W and message $\Omega \setminus W$ from every state outside of W. Then, the receiver interprets message W as a recommendation to approve by the (obedience) constraint. Off the equilibrium path, let the receiver be "skeptical" and assume that any unexpected message comes from the worst possible state. Then, the sender does not have profitable deviations: if $\omega \in W$, he is getting the highest possible payoff; if the state is not in W, the sender cannot replicate message W because $\omega \notin W$, and the receiver rejects after every other message.

Note that Theorem 1.1 is a version of the communication revelation principle for games with verifiable information. According to Myerson (1986) and Forges (1986), any equilibrium outcome of a mediated sender-receiver game may be implemented truthfully and obediently. In the present context, it translates into *(i)* the sender truthfully revealing the state of the world to the mediator, *(ii)* the mediator translating this report into an action recommendation for the receiver, and *(iii)* the receiver obediently following her recommendation. Which equilibrium outcome is implemented is decided by the mediator at step *(ii)*. Conveniently, Theorem 1.1 also provides the necessary and sufficient conditions for a set of approved states to be implementable in equilibrium.

1.3.2 Equilibrium Range and Value of Commitment

For the purposes of characterizing equilibrium outcomes, Theorem 1.1 allows us to restrict attention to sets $W \subseteq \Omega$ satisfying (IC) and (obedience). I rank equilibria in terms of the sender's ex-ante utility, which is the same as his ex-ante odds of approval and equals $P(W)$, the prior measure of the set of approved states.

In the sender-worst equilibrium, the set of approved states $\underline{W}$ minimizes the sender's ex-ante utility across all equilibria. Thus, the (IC) constraint binds and $\underline{W} = \mathcal{A}$. In this equilibrium, the receiver approves the proposal if and only if she approves it under complete information. Hence, the sender-worst equilibrium is outcome-equivalent to *full disclosure* (also known as *full unraveling*), salient in the verifiable information literature.[8]

In the sender-preferred equilibrium, the set of approved states $\overline{W}$ maximizes the sender's ex-ante utility across all equilibria. Mathematically,

$$\overline{W} = \arg\max_{W \subseteq \Omega} P(W), \quad \text{subject to} \quad \begin{array}{c} \mathcal{A} \subseteq W, \\ p(\cdot \mid W) \in \mathcal{B}. \end{array} \tag{1.1}$$

To find the sender-preferred equilibrium, we would increase the ex-ante measure of the set of approved states W so long as the receiver, when approving, expects that her net payoff of approval is non-negative, on average. Because the state space is continuous, $\overline{W}$ makes the receiver exactly indifferent between approval and rejection, binding her (obedience) constraint.

[8]See, e.g., Milgrom (1981), Grossman (1981), Milgrom and Roberts (1986) and review by Milgrom (2008).

Theorem 1.2. *When $n = 1$, the sender-preferred set of approved states $\overline{W}$ is characterized by a cutoff value $c^* > 0$ such that*

- *the receiver almost surely approves the proposal if $\delta(\omega) > -c^*$ and rejects it if $\delta(\omega) < -c^*$;*[9]
- *whenever the receiver approves the proposal, her expected net payoff of approval is zero: $\mathbb{E}_p\left[\delta(\omega) \mid \overline{W}\right] = 0$.*

Furthermore, the sender-preferred equilibrium outcome is a commitment outcome.

First, notice that the receiver's (obedience) constraint binds, or else we could increase the value of the objective while still satisfying that constraint. I prove the first part of Theorem 1.2 by contradiction. Suppose that the sender-preferred set of approved states $\overline{W}$ in not characterized by a cutoff value of the receiver's net payoff of approval. Then, there exist two sets $X, Y \subseteq \Omega$ of positive and equal measure, such that $\overline{W}$ includes X, $\overline{W}$ does not include Y, yet the receiver has a higher net payoff of approving any state in Y over any state in X. Consider an alternative set of approved states W^* that replaces X with Y, i.e. $W^* = (\overline{W} \setminus X) \cup Y$. The sender has the same ex-ante payoff at W^* and $\overline{W}$ because sets X and Y have the same measure. Yet, the (obedience) constraint for W^* is loose, while for $\overline{W}$ it is binding. That happens because every state in Y is "cheaper" in terms of the constraint than each state in X. Thus, we can improve upon both $\overline{W}$ and W^*, which is a contradiction.

Next, let us compare the problems of *(i)* finding the sender-preferred equilibrium outcome and *(ii)* finding the commitment outcome. In *(i)*, we maximize the ex-ante measure of the set of approved states subject to (IC) and (obedience) constraints. In *(ii)*, the sender maximizes his ex-ante utility subject to an obedience-like constraint of the receiver. Crucially, under commitment, the sender does not face an incentive-compatibility constraint. Also, a commitment outcome may not be deterministic.

[9] Almost surely with respect to the prior distribution p of the state of the world ω.

A commitment outcome is characterized by a cutoff value of the receiver's net payoff of approval for the same reason $\overline{W}$ is.[10] That is, the receiver certainly approves (rejects) the states with a net payoff of approval above (below) some threshold. Furthermore, that threshold is negative, and the receiver certainly approves every state in her approval set. Hence, any commitment outcome satisfies the sender's incentive-compatibility constraint.

In a non-deterministic commitment outcome, the sender induces both actions of the receiver with positive probabilities on some set $\mathcal{D} \subseteq \Omega$. Since any commitment outcome is characterized by a cutoff value, the receiver's net payoff of approval must be the same for every state in $\mathcal{D}$. Rather than making a mixed recommendation, partition the set of these states in two and let the sender recommend one action on each subset with certainty. Due to the continuity of the state space, such partitioning does not affect the objective function or the obedience constraint of the receiver. As a result, there exists deterministic commitment outcome. Since this commitment outcome satisfies the sender's incentive-compatibility constraint, it is an equilibrium outcome.

Example 1.2 (Receiver with Spatial Preferences: Equilibrium Range)**.** Suppose that the receiver has spatial preferences described in Example 1.1. In the sender-worst equilibrium, the set of approved states is $\underline{W} = \mathcal{A}$, and the receiver approves the proposal if and only if the sender's position is closer to her bliss point than the status quo.

To find the sender-preferred set of approved states $\overline{W}$, we maximize the measure of set $W \subseteq \Omega$ subject to the receiver's (obedience) constraint. According to Theorem 1.2, in the sender-preferred equilibrium, the receiver approves some states outside of her approval set. However, there is a cutoff for how far the sender's position could be to be approved. When approving, the receiver expects that the sender's position and the status quo are equidistant from her bliss point. Furthermore, the sender-preferred

[10] Alonso and Câmara (2016) prove that if the state space is finite, then the solution under commitment features a cutoff state.

equilibrium outcome is a commitment outcome, meaning that the sender need not benefit from having ex-ante commitment power. Figure 1.2 illustrates the equilibrium range.

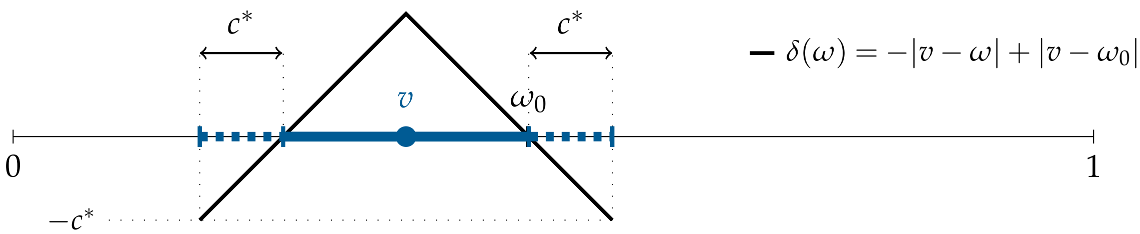

Figure 1.2. The sender-worst set of approved states $\underline{W} = \{\omega \in \Omega \text{ s.t. } |v-\omega| \leq |v-\omega_0|\}$ (solid blue) and the sender-preferred set of approved states $\overline{W} = \{\omega \in \Omega \text{ s.t. } |v-\omega| \leq |v-\omega_0| + c^*\}$ (solid plus dotted blue), where c^* solves $\mathbb{E}_p[|v-\omega| \mid \overline{W}] = |v-\omega_0|$.

1.4 Many Receivers

Having assumed that the receivers solve independent problems, I get similar results in the many-receiver case.[11]

Theorem 1.3. *The following statements about the sender's ex-ante payoff $\overline{u}_s$ are equivalent:*

1. *$\overline{u}_s$ is reached in equilibrium;*

2. *$\overline{u}_s$ is given by*

$$\overline{u}_s = \int_{\Omega} u_s(\{i \in I \mid \omega \in W_i\}) \cdot p(\omega) d\omega,$$

where for every receiver $i \in I$, $W_i \subseteq \Omega$ is her set of approved states, which satisfies

- *sender's (IC) constraint $\mathcal{A}_i \subseteq W_i$,*
- *receiver's obedience constraint $p(\cdot \mid W_i) \in \mathcal{B}_i$.*

[11] That is, receiver i's utility does not depend on other receivers' actions, and receiver i's message is private and observed by her only.

The proof of the theorem follows the same steps as the proof of Theorem 1.1. The only substantial difference is that Theorem 1.3 characterizes the sender's equilibrium ex-ante utility, while Theorem 1.1 characterizes the equilibrium sets of approved states. The reason is that with many receivers, some equilibrium outcomes are not deterministic. That happens because the sender may not try his hardest to convince the receivers whose approval does not strictly increase his payoff.

According to Theorem 1.3, when characterizing the sender's equilibrium ex-ante utility, we can restrict attention to collections of sets of approved states $(W_1, \ldots, W_n)$, each of which satisfies the IC and obedience constraints for each receiver. Moreover, the sender's ex-ante utility only depends on $(W_1, \ldots, W_n)$ and the prior distribution.

Once again, in the sender-worst equilibrium, in which the sender's ex-ante utility is minimized across all equilibria, the sender does as well as under full disclosure. The set of approved states of receiver $i \in I$ is $\underline{W}_i = \mathcal{A}_i$, and each receiver makes her decision as if under complete information.

The sender-preferred equilibrium outcome is characterized by the collection of sets of approved states that maximizes the sender's ex-ante utility across all equilibria, i.e. subject to every receiver's obedience constraint and every incentive-compatibility constraint of the sender. When there are many receivers, the sender need not benefit from having commitment power, either.

Theorem 1.4. *The sender's ex-ante payoff in the sender-preferred equilibrium is the commitment payoff.*

The proof of Theorem 1.4 follows the same steps as the proof of Theorem 1.2. That is, I show that if we take an arbitrary commitment outcome, we can find a deterministic commitment outcome with the same payoff of the sender. That deterministic commitment outcome satisfies every (IC) constraint of the sender, meaning that it is also an equilibrium outcome.

In general, the problem of finding the sender-preferred equilibrium outcome is computationally hard.[12] In the following section, I make additional assumptions on the sender's payoff and study elections.

1.5 Targeted Advertising in Elections

In this section, I show that targeted advertising helps politicians swing elections. I compare communication via targeted advertising to public disclosure. In the first case, the politician sends a private message to each voter, for example, through social media. Targeted advertising is an application of the main model. In the second case, the politician sends a public message to all voters. Public disclosure is not an application of the main model. However, analysis of that case is simple because the voters share a common prior belief, and if they receive the same message, they will also share a common posterior belief.

In this application, Ω is the policy space, with positions ranging from far-left (0) to far-right (1). The sender is a politician who challenges the status quo. The challenger is privately informed about his policy $\omega \in \Omega$, while the receivers hold a prior belief p. The challenger receives 1 if he wins the election and 0 otherwise. The outcome of the election is decided by the social choice function u_s that satisfies Assumption 1.1. For example, the election may be decided by a simple majority: the challenger wins the election if and only if the majority of receivers approve his policy, i.e. $u_s(X) = 1 \iff |X| > n/2$.

The set of receivers I is now the electorate, and the receivers are sincere voters with spatial preferences. Firstly, each voter chooses expressively, and not strategically,

[12] Babichenko and Barman (2016) show that the problem of finding the commitment outcome is NP-hard when the sender's utility is submodular; Arieli and Babichenko (2019) find the commitment outcome for the case of supermodular utility; Kamenica (2019) note that if sender's utility is separable in receiver's actions, then the sender determines the optimal signal receiver by receiver and faces a set of independent problems of a single-receiver variety.

between the challenger and the status quo.[13] Secondly, I assume that the approval set of voter $i \in I$ is $\mathcal{A}_i = \{\omega \in \Omega \text{ s.t. } |v_i - \omega| \leq |v_i - \omega_0| - \varepsilon\}$, where $\varepsilon > 0$.[14] That is, voter i approves policies that are closer than the status quo to her bliss point by at least ε.

Observe that the preferences of the electorate can be summarized by the preferences of at most two voters whose bliss points are located closest to the status quo.

Definition 1.4. *Voter* $L = \arg \max_{i \in I,\, v_i < \omega_0} v_i$ *is the* *left representative voter* *and voter* $R = \arg \min_{j \in I,\, v_j > \omega_0} v_j$ *is the* *right representative voter.*

First, notice that as a voter's bliss point moves away from the status quo, her approval set expands to include more policies of the challenger. Put differently, the further a voter's bliss point is from the status quo, the easier it is for the challenger to convince her. As a result, if the challenger convinces the left (right) representative voter, he also convinces all voters with bliss points further to the left (right). Second, voters with bliss points on opposite sides of the status quo have incompatible preferences. Intuitively, left voters prefer to approve the left policies of the challenger, while right voters prefer to approve policies on the right. These observations are summarized in Corollary 1.1 and illustrated in Figure 1.3.

Corollary 1.1. *If L and R are representative voters, then*

1. *if L (R) prefers to approve challenger's policy, then so does every voter with a bliss point to her left (right), i.e.*

$$\mathcal{A}_L \subset \mathcal{A}_i \text{ and } \mathcal{B}_L \subset \mathcal{B}_i, \forall i \in I \text{ such that } v_i < v_L,$$

[13] The theory of sincere voting was pioneered by Brennan and Lomasky (1993), Brennan and Hamlin (1998), and reviewed by Hamlin and Jennings (2011). There is a large body of evidence that the behavior of voters in large elections is consistent with sincere voting, e.g., in U.S. national elections (Kan and Yang, 2001; Degan and Merlo, 2007), Spanish General elections (Artabe and Gardeazabal, 2014), Israeli General elections (Felsenthal and Brichta, 1985).

[14] ε is the status quo bias; $\varepsilon > 0$ rules out situations wherein the challenger with the status quo policy always wins the election.

$$\mathcal{A}_R \subset \mathcal{A}_j \text{ and } \mathcal{B}_R \subset \mathcal{B}_j, \forall j \in I \text{ such that } v_j > v_R;$$

2. *approval sets and sets of approval beliefs of voters L and R do not intersect, i.e.*

$$\mathcal{A}_L \cap \mathcal{A}_R = \varnothing \text{ and } \mathcal{B}_L \cap \mathcal{B}_R = \varnothing.$$

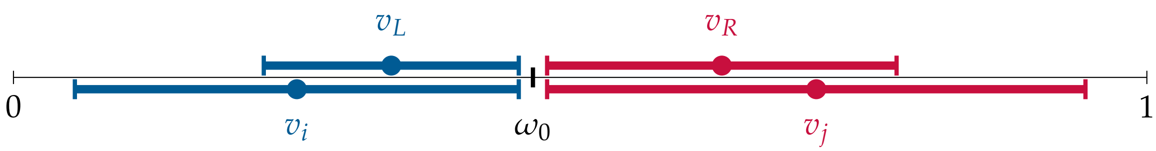

Figure 1.3. Voter i is convinced if voter L is convinced: her approval set includes L's approval set (solid blue lines). Voters L and R have incompatible preferences: their approval sets do not intersect.

1.5.1 Swinging Unwinnable Elections

Part 2 of Corollary 1.1 implies that voters L and R *never* both approve the challenger's policy when they hold the same belief. Thus, if representative voters L and R are jointly pivotal, the challenger always loses the election under common belief.

Definition 1.5. *Election with representative voters L and R is unwinnable for the challenger under common belief if for all $X \subseteq I$, $u_s(X) = 1$ if and only if $\{L, R\} \in X$.*

Whether an election is unwinnable is determined by the institution (the social choice function) and the ideology (bliss points of the voters). For example, under the simple majority rule, we arrive at a version of the median voter theorem.[15] Intuitively, for an election to be unwinnable, there may not be a majority of voters located on either side of the status quo.

[15] Black (1948) states the median voter theorem as "If Ω is a single-dimensional issue and all voters have single-peaked preferences defined over Ω, then ω_0, the median position, could not lose under majority rule."

Corollary 1.2. *Under the simple majority rule, an election is unwinnable for the challenger under common belief if and only if ω_0 is the median voter's bliss point.*

With targeted advertising, the challenger can say different things to different voters. The voters will no longer hold the same belief, which opens up a possibility of winning (with positive probability) an unwinnable election. Here I show how the challenger can convince representative voters L and R, persuading who is sufficient to win *any* unwinnable election. I focus on the best-case scenario for the challenger and thus consider the sender-preferred equilibrium.

By Theorem 1.3, we can restrict attention to a pair of sets of approved policies (W_L, W_R). In the sender-preferred equilibrium, we maximize the challenger's odds of convincing the representative voters subject to their obedience constraints:

$$\max_{W_L, W_R} P(W_L \cap W_R)$$

$$\text{subject to } p(\cdot \mid W_i) \in \mathcal{B}_i, \text{ for } i \in \{L, R\}.$$

The following theorem describes the solution to this problem.

Theorem 1.5. *In the sender-preferred equilibrium of an unwinnable election with representative voters L and R, if ε is small enough,*

- *the set of approved policies $\overline{W}_i$ of voter $i \in \{L, R\}$ is an interval $[a_i, b_i] \supset \mathcal{A}_i$;*
- *the challenger wins the election if his policy is in the interval $[a_R, b_L]$ with $a_R < \omega_0 < b_L$;*
- *the challenger's ex-ante odds of winning the election are positive. i.e. $P([a_R, b_L]) > 0$.*

To understand the intuition behind this result, recall that when voter L (R) is the only receiver, the challenger can convince her to approve his policy even when his policy is slightly to the right (left) of the status quo. I illustrated that in Figure 1.2 of

Example 1.2. One thing that the challenger can do under private communication is treat each voter as if she is the only receiver. If his policy is close enough to the status quo and ε is small enough, the challenger convinces both voters at the same time and swings an unwinnable election. However, he can do even better. To convince voter L (R), the challenger needs to make her believe that his policy is on average to the left (right) of the status quo. To induce that belief, the challenger could pull left (right) policies within this voter's approval set with some of the right (left) policies preferred by her counterpart. More precisely, voter L's (R's) message would include her approval set and as many policies to the right (left) of the status quo as this voter's obedience constraint permits. This solution is illustrated in Figure 1.4.

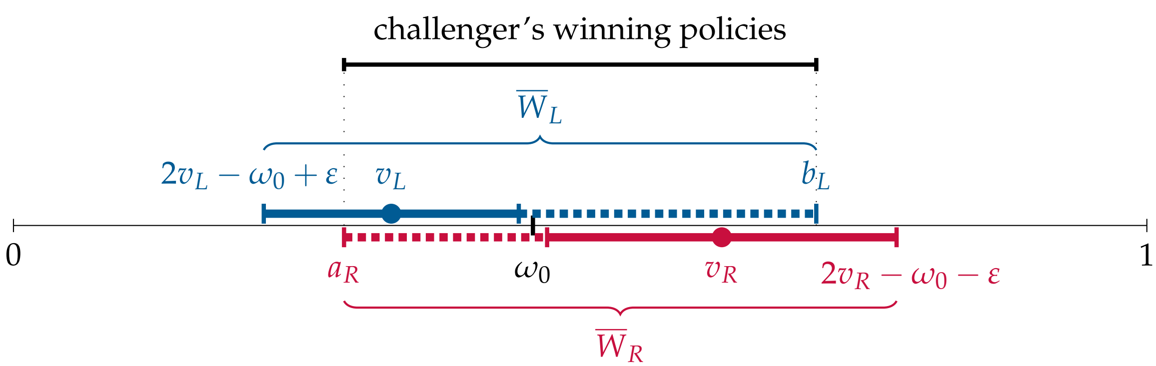

Figure 1.4. The sender-preferred sets of approved policies $\overline{W}_L$ (in blue) and $\overline{W}_R$ (in red). $\overline{W}_i$ consists of voter i's approval set (solid) and policies preferred by voter $j \neq i$ (dotted). The challenger wins the election by convincing both voters when his policy is in $\overline{W}_L \cap \overline{W}_R = [a_R, b_L]$.

Comparative Statics

Assume for the rest of this section that the prior is uniform.[16] Notice that the distance from a voter's bliss point to ω_0 measures this voter's persuadability.

Definition 1.6. *Suppose that $p \sim U[0,1]$. Then,*

- *voter i is more persuadable than voter j if $|v_i - \omega_0| > |v_j - \omega_0|$, where $i, j \in I$;*

[16]The prior is chosen to be uniform for ease of exposition. Similar results hold for any prior distribution.

- *consider electorates I and I' with representative voters $\{L,R\}$ and $\{L',R'\}$. I' is <u>more polarized</u> than I if $v'_L \leq v_L < \omega_0 < v_R \leq v'_R$.*

In words, the further from the status quo the voter's bliss point is, the less satisfied she is with the status quo policy, and that makes her more persuadable. I say voter $i \in I$ *becomes more persuadable* if $|v_i - \omega_0|$ increases. The electorate *becomes more polarized* when either representative voter becomes more persuadable. Figure 1.5 illustrates the dynamics of the numerical solution to the problem of finding the sender-preferred equilibrium as voter R becomes more persuadable (and the electorate becomes more polarized). Theorem 1.6 summarizes the comparative statics.

Theorem 1.6. *Suppose that $p \sim U[0,1]$. In the sender-preferred equilibrium of an unwinnable election with representative voters L and R,*

- *as R becomes more persuadable, the challenger's ex-ante odds of winning $P([a_R,b_L])$ increase;*
- *suppose $|v_L - \omega_0| = |v_R - \omega_0|$, meaning that neither voter is more persuadable than the other. Then, as R becomes more persuadable, the set of challenger's winning policies $[a_R,b_L]$ shifts to the left, i.e. a_R and b_L decrease.*

In words, as voter R becomes more persuadable, it becomes easier for the challenger to swing the election by targeting, in the sense that his ex-ante odds of winning increase. Furthermore, R becomes more persuadable by policies further to the left, meaning that the set of winning policies shifts to the left, also. When voter R is far enough to the right, her obedience constraint no longer binds (as in the top exhibit of Figure 1.5), and the sender-preferred set of approved policies is the same as if voter L was the only receiver.

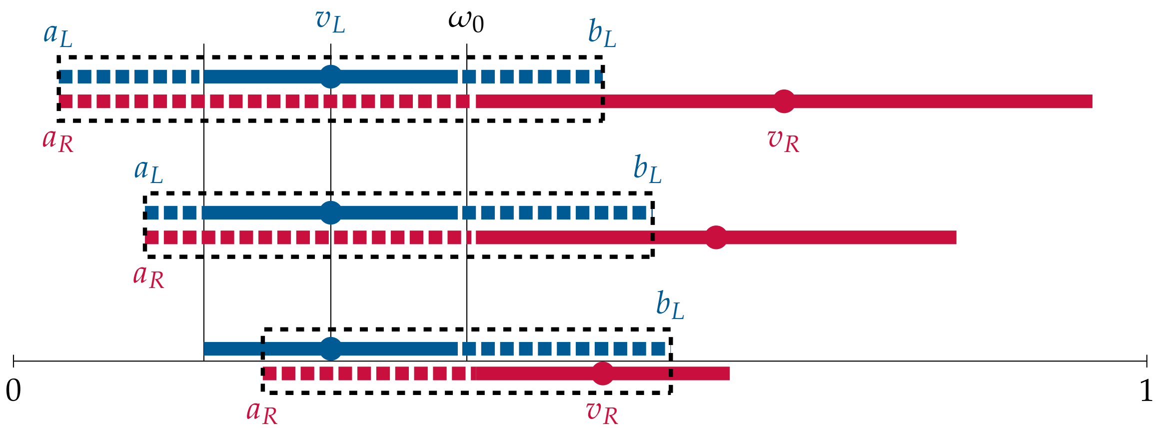

Figure 1.5. Comparative statics as voter R moves to the right (bottom to top): her approval set (solid red area) expands; she is convinced by more policies on the left (dashed red area); the set of challenger's winning policies (dashed black area) moves to the left and expands.

1.6 Conclusion

This chapter argued that the sender need not benefit from having commitment power and can persuade the receivers with verifiable information only. This result is useful in applications, especially in the context of elections, where assuming that the sender has hard evidence is more plausible than assuming that the sender has commitment power.

While illustrated in the simplified framework, the observation that targeted advertising helps challenger swing elections holds for more than one dimension and any social choice rule. Because targeting leads to election outcomes that are different from the complete-information outcomes, one can argue that targeted advertising is bad for democracy. Certain policy implications, especially concerning restricting the collection and use of personal data by the candidates in their electoral campaigns, should be considered.

Chapter 1 is currently being prepared for submission for publication of the material. The dissertation author, Mariia Titova, is the sole author of this material.

Chapter 2

Collaborative Search for a Public Good

2.1 Introduction

Suppose a team of agents is facing a problem, solving which benefits them all. Before making a collective choice, the team members must engage in a costly search to learn the possible solutions. How efficient is collaborative search?

Collaborative search for a public good takes place in many economic situations. In politics, policymakers identify the best available policies. In organizations, committee members search for the most qualified candidate. In consumer search, family members look for a house to move to. In research and development, scientists decide which idea to pursue. Broadly speaking, any situation that involves sequential social learning and that results in the final project benefiting everyone can be studied using this model. I examine inefficiencies that arise as an artifact of sequential searching in teams rather than individually.

I model the sequential search process after the seminal model of Weitzman (1979). There are two team members and a finite number of boxes. Each box contains an uncertain reward. To learn the contents of a box, one needs to open it, which comes at a cost. At each stage of the game, one agent is randomly chosen to decide between three alternatives: she could open a box of her choice, do nothing, or propose to terminate the game. The game ends if a termination offer is extended and accepted or if there are no more boxes left to open. At the end of the game, both players collect the highest reward among all the opened boxes.[1] I study *(i)* the optimal order of search among alternatives, *(ii)* incentives to free ride on colleague's search efforts, *(iii)* the efficiency of searching in teams.

My most important result is that, compared to the socially optimal protocol of an individual searcher, the team will use the same search order and stopping rule. In other words, the policymakers identify the same policy, the committee members find the same candidate, the family moves to the same house, and the scientists pursue the same research project as if these choices were made by the social planner. However, team search may be inefficient due to the free riding effect: agents procrastinate at each stage of the search process and hope their colleagues exert the search effort instead.

In the symmetric equilibrium, the chosen player acts similarly to how she would act had she been searching alone. More precisely, if she decides to open a box, she opens the "best" box according to Weitzman (1979), i.e. the box with the highest reservation value. She also wants to stop and proposes to terminate the game at the same threshold – when no boxes are "good enough" to be opened, i.e. the highest reservation value among the unopened boxes is lower than best uncovered reward so far. In that case, her opponent always accepts the termination offer. Consequently, the search order and termination protocol on the equilibrium path are those of the social planner. The only difference is that agents free ride when search costs are sufficiently high: the chosen player only opens a box sometimes, and does nothing the rest of the time. As a result, delay arises at each stage of the learning process.

Compared to searching by herself, each agent is doing better on average when searching with a companion. Intuitively, their learning protocol is the same, but each

[1]Weitzman's setup was generalized by Olszewski and Weber (2015) to have the searcher's payoff depends on all discovered prizes, and by Doval (2018) to allow the searcher to forego inspection costs and take any unopened box. I focus on Weitzman's setup because of the simple form of the solution that allows for a straightforward comparison of the team vs. social planner results.

team member only pays the search cost about half the time. However, team search is inefficient because every box is opened with a delay. Delay occurs at each phase of the learning process in the sense that it takes time to open each consecutive box. How long it takes to open a box depends on the distribution of the reward and the best uncovered option so far. As players open boxes, two effects take place. First, because the search protocol prescribes to open ex-ante better boxes first, ex-ante worse boxes remain towards the end. Hence, the next best box to be opened is less attractive than the previous box and takes longer to open. Second, when someone opens a box, the best uncovered option improves. That decreases the time it takes to open the following box because the players are eager to collect the higher reward. As a result, the time it takes to open the next box may go up or down.

Related Literature

First, this chapter contributes to the literature on collective experimentation. In their seminar paper, Bolton and Harris (1999) extend the classic two-armed bandit problem to a many-agent setting. Because information is a public good, agents have an incentive to wait and let their colleagues experiment instead, which is known as the *free rider effect*. At the same time, the prospect of others experimenting forces every agent to experiment more, which is known as the *encouragement effect*. The strength of each of these effects depends on the problem. For example, if the bandit is exponential, as in Keller, Rady, and Cripps (2005), then only the free-riding effect is present. With Poisson bandits, Keller and Rady (2010) show that the encouragement effect dominates. All these papers focus on a two-armed bandit that has one safe, one risky arm. In this chapter, I consider a multi-armed bandit, such that the outcome of an arm is revealed after just one experiment. This allows me to study the *order* as well as the *stopping rule* of the search process for the public good.

The literature on collective experimentation is closely related to the literature on delegation and approval of experimentation. When a principal delegates experimentation to agents, the optimal mechanism exhibits tolerance for early failure (Manso, 2011; Lewis, 2012), but asymmetric information leads to less experimentation, lower success rate, and more variance in success rates (Halac, Kartik, and Liu, 2016). When the agent is privately informed about the state of the world, optimal dynamic mechanisms feature cutoff approval rules (Guo, 2016; McClellan, 2019). I abstract away from the principal's incentive-compatibility problem and study individual and team incentives instead. I conclude that from the outside (i.e. the manager's) perspective, search by an individual is more efficient because it happens without any delay. At the same time, each agent would rather search with a companion because that allows her to pay the search cost less often.

This chapter contributes to the literature on the search by committees. Group experimentation that ends with a vote is usually inefficient because committee members experience loser trap and winner frustration (Strulovici, 2010), are less picky and more conservative than a single agent (Albrecht, Anderson, and Vroman, 2010), or because they communicate before the game (Compte and Jehiel, 2010). In my setting, the agents are collectively searching for the public good. Since their preferences are aligned, agents agree to end the learning process when the social planner would, which is true for any social choice function.

This chapter also contributes to the literature on collaboration in teams. When agents work on a project as a team, inefficiencies arise because team members have an incentive to procrastinate (Bonatti and Hörner, 2011) and due to lack of communication (Campbell, Ederer, and Spinnewijn, 2014). The size of the inefficiency is minimized if the manager dynamically decreases the size of the team as the project nears completion (Georgiadis, 2015). My findings confirm that agents free ride when searching in teams. At the same time, they prefer to search as part of a team because when they procrastinate, there is a chance that their partner exerts the costly search effort.

Finally, this chapter contributes to the literature on the dynamic provision of public goods. Efficiency is usually not achieved because socially optimal projects are not completed (Fershtman and Nitzan, 1991; Admati and Perry, 1991; Kessing, 2007), completed with a delay (Marx and Matthews, 2000; Compte and Jehiel, 2004), or completed at a lower scale (Bowen, Georgiadis, and Lambert, 2019). I show that when searching for the public good, team members effectively use the socially optimal search protocol, albeit with some free riding when search costs are large enough. Consequently, all socially optimal projects are searched through, and the only source of inefficiency is the delay.

The rest of this chapter is organized as follows. Section 2.2 introduces the dynamic model of sequential search among risky alternatives. Section 2.3 describes the model with one alternative and discusses comparative statics and welfare implications. Section 2.4 generalizes the model to the case of finitely many alternatives. Section 2.5 is a conclusion.

2.2 Model

Two agents sequentially search for a public good. At each stage, one agent is chosen randomly with a probability of $1/2$. The chosen player has an option to *(i)* open exactly one box of her choice, *(ii)* do nothing, or *(iii)* propose to terminate the game. In the latter case, her opponent chooses between accepting and rejecting this proposal.

Each public good project is represented by a box that contains a stochastic prize. Initially, there is a finite number of unopened boxes. Box $b_k = (c_k, F_k)$ contains an uncertain reward $x_k \sim F_k(\cdot)$ distributed independently of all other rewards. If the chosen player decides to open this box, she pays the search cost c_k and players wait one period to learn its contents. Once the contents are revealed, a new stage starts immediately, and a new player is chosen. The game ends if the chosen player proposes termination and the opponent accepts the offer, or if there are no more boxes left to open. In either case, each player collects the highest reward they uncovered during the search. The initial fallback reward is z_0.

Both players are risk-neutral and wish to maximize the expected present value of the best uncovered reward. The search costs are sunk because they are paid during the search process, while the reward is only realized upon the end of the game. The players discount the time at the exponential rate $\delta = e^{-r\Delta t}$, where Δt is the length of the time interval between the stages. This chapter aims to find a dynamic rule that describes the optimal search protocol. This dynamic rule should specify which box (if any) to open, when to propose to end the search process, and when to accept the termination offer.

2.2.1 Dynamic Problem

Let $\mathcal{B}$ denote the set of unopened boxes and z be the best uncovered reward. I focus on the stationary Markov perfect equilibrium. It is a subgame perfect equilibrium in which strategies depend only on the payoff-relevant history, i.e. the pair $(z, \mathcal{B})$.

A stationary Markov strategy for player i is a pair $a_i := (a_i^{ch}, a_i^{op})$ that specifies which action she takes when she is chosen and when she is the opponent, respectively. With slight abuse of notation, $a_i^{ch}(z, \mathcal{B}) \in A^{ch}(z, \mathcal{B}) := \{\varnothing, T\} \cup \mathcal{B}$, meaning that when chosen, player i decides between doing nothing, proposing termination, and opening one of the boxes in $\mathcal{B}$. When she receives a termination proposal, $a_i^{op}(z, \mathcal{B}) \in A^{op}(z, \mathcal{B}) := \{0, 1\}$, so she can reject or accept it. The mixed stationary Markov strategy of player i is denoted by $\alpha_i = (\alpha_i^{ch}, \alpha_i^{op})$, where $\alpha_i(z, \mathcal{B}) \in \Delta A^{ch}(z, \mathcal{B}) \times \Delta A^{op}(z, \mathcal{B})$.

In state $(z, \mathcal{B})$, let $\Phi_i^{ch}(z, \mathcal{B}; \boldsymbol{a}_j)$ be the highest possible payoff that player i can achieve when she is chosen at this stage, given that player $j \neq i$ plays the Markov strategy $\boldsymbol{a}_j$. Let $\Phi_i^{op}(z, \mathcal{B}; \boldsymbol{a}_j)$ be her continuation value when she is the opponent. Also, let

$$\overline{\Phi}_i(z, \mathcal{B}; \boldsymbol{a}_j) := \frac{\Phi_i^{ch}(z, \mathcal{B}; \boldsymbol{a}_j) + \Phi_i^{op}(z, \mathcal{B}; \boldsymbol{a}_j)}{2}$$

be the average value function that accounts for the fact that each period player i is chosen with probability 1/2. It then follows that

$$\Phi_i^{ch}(z,\mathcal{B};\boldsymbol{a}_j) = \max_{a_i^{ch}\in\{\varnothing,T\}\cup\mathcal{B}} \begin{cases} z, & \text{if } a_i^{ch} = T \wedge \boldsymbol{a}_j^{op} = 1, \\ \delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j), & \text{if } a_i^{ch} = T \wedge \boldsymbol{a}_j^{op} = 0, \text{ or } a_i^{ch} = \varnothing, \\ -c_k + \delta\Big[\overline{\Phi}_i(z,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)F_k(z) + \int_z^{+\infty}\overline{\Phi}_i(x,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)dF_k(x)\Big], & \text{if } a_i^{ch} = b_k \in \mathcal{B}. \end{cases} \tag{2.1}$$

In words, if she proposes termination and her offer is accepted, player i receives z immediately. If her offer of termination is rejected or if she does nothing, then the next period starts, the time between periods is discounted by a factor of δ, and roles are reset. If she opens box b_k, she pays the search cost c_k immediately. Next period, contents of the box are revealed, best uncovered reward and the set of available boxes are updated, and roles are reset.

The value function of player i when she is the opponent is

$$\Phi_i^{op}(z,\mathcal{B};\boldsymbol{a}_j) = \max_{a_i^{op}\in\{0,1\}} \begin{cases} z, & \text{if } \boldsymbol{a}_j^{ch} = T \wedge a_i^{op} = 1, \\ \delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j), & \text{if } \boldsymbol{a}_j^{ch} = T \wedge a_i^{op} = 0, \text{ or } \boldsymbol{a}_j^{ch} = \varnothing, \\ \delta\Big[\overline{\Phi}_i(z,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)F_k(z) + \int_z^{+\infty}\overline{\Phi}_i(x,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)dF_k(x)\Big], & \text{if } \boldsymbol{a}_j^{ch} = b_k \in \mathcal{B}. \end{cases} \tag{2.2}$$

When she is the opponent, player i chooses between accepting and rejecting a termination proposal. If she accepts, her payoff is z. If she rejects or if the chosen player did nothing, the next stage begins, and the state is unchanged. If the chosen player opens a box, player i observes its contents without paying the search cost.

Player i's value functions $\Phi_i^{ch}(z, \mathcal{B}; \boldsymbol{\alpha}_j)$ and $\Phi_i^{op}(z, \mathcal{B}; \boldsymbol{\alpha}_j)$ given player j's mixed Markov strategy $\boldsymbol{\alpha}_j$ are calculated by taking expectation of (2.1) and (2.2) with respect to $\boldsymbol{\alpha}_j(z, \mathcal{B})$.

When all boxes are open, the players collect the best uncovered reward, i.e.

$$\Phi_i^{ch}(z, \varnothing) = \Phi_i^{op}(z, \varnothing) = \overline{\Phi}_i(z, \varnothing) = z. \tag{2.3}$$

Definition 2.1. *Profile of strategies $(\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2)$ is a Markov perfect equilibrium if for every player $i \in \{1, 2\}$ and any possible state $(z, \mathcal{B})$, if $\boldsymbol{a}_i \in \operatorname{supp} \boldsymbol{\alpha}_i$, then $\boldsymbol{a}_i^{ch}$ maximizes $\Phi_i^{ch}(z, \mathcal{B}; \boldsymbol{\alpha}_j)$ and $\boldsymbol{a}_i^{op}$ maximizes $\Phi_i^{op}(z, \mathcal{B}; \boldsymbol{\alpha}_j)$ subject to the boundary condition* (2.3).

Since the players are symmetric, I focus on the symmetric equilibria.

2.3 One Box

Let $\mathcal{B}$ contain just one box b and z be the safe option. Note that when initially there is only one box, z equals the initial fallback reward z_0.

Suppose that there is only one unopened box, i.e. $\mathcal{B} = \{b\}$. Let

$$S(z, F) := \mathbb{E}[\max\{z, x\}] = zF(z) + \int_z^{+\infty} x dF(x)$$

be the expected value of the best uncovered reward after opening the box. That is, with probability $\text{Prob}(x \leq z) = F(z)$, the reward in the box is lower than z, in which case the best uncovered option is not updated. Otherwise, the reward x discovered in the box becomes the new safe option.

Recall that, according to Weitzman (1979), social planner opens the box if and only if the expected benefit net of the search cost exceeds the safe option:

$$-c + \delta S(z, F) \geq z. \tag{SR}$$

Weitzman shows that there exists a unique z that solves the binding *social rationality* condition (SR).

Definition 2.2. *Reservation value $\bar{z}$ of box $b = (c, F)$ solves $-c + \delta S(\bar{z}, F) = \bar{z}$.*

Weitzman also shows that (SR) holds if and only if $z \leq \bar{z}$. Notice that if (SR) holds, then for the chosen player, doing nothing weakly dominates proposing termination.

How does the opponent respond to a termination offer? The best she can do by refusing to terminate the game is wait until she is chosen and open the box herself. Due to the discounting of future payoffs and the uncertainty regarding the period in which she is chosen, the expected payoff from opening the box is multiplied by a factor of $\frac{1}{2}\delta + \left(\frac{1}{2}\delta\right)^2 + \cdots = \frac{\delta}{2-\delta}$. Thus, the opponent rejects a termination offer if and only if the following *individual rationality* condition holds:

$$\frac{\delta}{2-\delta} \cdot [-c + \delta S(z, F)] \geq z. \tag{IR}$$

While Weitzman (1979) defines $\bar{z}$ as the threshold for opening the box, I define z^R as the threshold for rejecting a termination proposal in favor of opening the box when she is chosen.

Definition 2.3. *Rejection threshold z^R of box $b = (c, F)$ solves $\frac{\delta}{2-\delta} \cdot [-c + \delta S(z^R, F)] = z^R$.*

It is straightforward to show that z^R is unique and that (IR) condition holds if and only if $z \leq z^R$.[2] Furthermore, $z^R \leq \bar{z}$. In other words, (IR) implies (SR), but not vice versa. Intuitively, if the box is good enough that the player is willing to wait to open it in the future, then the box is good enough that she is willing to open it today.

[2]The formal proof of this statement and all other results can be found in the appendix.

When $z \leq \bar{z}$, the chosen player opens the box, does nothing, or mixes between these two actions. It is easy to see that a symmetric equilibrium in *pure* stationary Markov strategies often does not exist. When a player is chosen and knows that her counterpart will open the box (do nothing) when chosen, she is better off doing nothing (opening the box). Consequently, when (SR) holds, the chosen player mixes between opening the box and doing nothing. Let π be the equilibrium probability of opening the box. The chosen player must be indifferent between *(i)* opening the box today and *(ii) someone* opening the box in the future, which translates into

$$-c + \delta S(z, F) = \frac{\pi\delta}{1 - (1-\pi)\delta} \cdot \Big[-\frac{c}{2} + \delta S(z, F) \Big]. \tag{2.4}$$

In words, her expected payoff from not opening the box today is the surplus $\delta S(z, F)$ from the box being opened eventually, less her having to pay the search cost c *half of the time* on average, infinitely discounted according to the time discount factor δ and the probability $1 - \pi$ that the box is not opened in the current period.

By solving the indifference condition above, we obtain the equilibrium probability of opening the box π as a function of the safe option z. Theorem 2.1 summarizes the search and termination protocol for the model with one box.

Theorem 2.1. *Let z be the safe option and $\mathcal{B} = \{b\}$. In the symmetric equilibrium,*

the chosen player

- *if $z \leq \bar{z}$, opens the box with probability*

$$\pi(z) = \begin{cases} \dfrac{2(1-\delta)}{\delta c} \cdot \big[-c + \delta S(z, F) \big] < 1 & \text{if } c > S(z, F) \cdot \dfrac{2\delta(1-\delta)}{2-\delta}, \\ 1 & \text{otherwise}, \end{cases}$$

and does nothing with probability $1 - \pi(z)$;

– *proposes to terminate the game if* $z > \bar{z}$.

the opponent

– *rejects the termination proposal if* $z \le z^R$;

– *accepts it otherwise.*

Comparing this to the optimal search and stopping protocol of an individual searcher, we can see that on the equilibrium path, the box is eventually opened as long as $z \le \bar{z}$, the same cutoff as in Weitzman (1979). Put differently, the box is opened if and only if it is socially optimal to do so. If the search cost if large enough, there is a chance that the chosen player does nothing instead of opening the box. Hence, it may take several periods to open the same box that the social planner opens right away. Consequently, collaborative search results in delay as a consequence of the free riding. I measure the size of the delay and discuss comparative statics in the section that follows.

In the case of one box, we can refer to the state of the problem as just z, and the value functions take a simple form.

Corollary 2.1. *In the symmetric equilibrium, if* z *is the safe option and one box remains to be opened, the value functions are*

$$\begin{aligned}
\Phi^{ch}(z) &= \max\left\{z, -c + \delta S(z,F)\right\}, \\
\Phi^{op}(z) &= \max\left\{z, \frac{2-\delta}{\delta} \cdot \left[-c + \delta S(z,F)\right]\right\}, \\
\overline{\Phi}(z) &= \max\left\{z, \frac{1}{\delta} \cdot \left[-c + \delta S(z,F)\right]\right\}.
\end{aligned}$$

Consider the case when the box is eventually opened, i.e. when $z \le -c + \delta S(z,F)$. Notice that the chosen player can guarantee herself the payoff of an individual searcher by opening the box. Since she opens the box with a positive probability, her indifference implies that her value function is exactly that of the social planner. On average, however, each searcher is doing strictly better than the social planner since $\overline{\Phi}(z) = \frac{1}{\delta}\Phi^{ch}(z) > \Phi^{ch}(z)$. The key to understanding this result is recalling that the chosen player can wait to become an opponent, in which case her value function is strictly higher because $\Phi^{op}(z) = \frac{2-\delta}{\delta}\Phi^{ch}(z) > \Phi^{ch}(z)$. Simply speaking, each player prefers to search in a team because then there is a chance that her companion pays the search cost.

Lemma 2.1. *In the symmetric equilibrium, the value functions $\Phi^{ch}(z)$ and $\overline{\Phi}(z)$ increase as*

- *the search cost c decreases,*
- *the length of the time interval between stages Δt decreases.*

The opponent's value function $\Phi^{op}(z)$ decreases in search cost; $\frac{\partial \Phi^{op}(z)}{\partial \Delta t}$ could be positive or negative.

When the search cost decreases, two effects take place. First of all, the expected net benefit of opening the box $-c + \delta S(z, F)$ increases. Secondly, the reservation value $\overline{z}$ of the box also increases, i.e. this box becomes ex-ante more attractive. As a result, the box is now opened for values of z for which it was not opened before, which drives the value functions even higher. The same arguments apply for Φ^{ch} and $\overline{\Phi}$ when the discount factor δ increases due to the shorter wait between stages of the game.

The dynamics of the opponent's value function with respect to Δt are inconclusive. On the one hand, a higher discount factor leads to a higher continuation value because the box becomes ex-ante more attractive. On the other hand, as Δt decreases, the incentive to free ride increases, and that drives the equilibrium probability of opening the box down. Consequently, the opponent's continuation value drops since it is determined by the likelihood of the chosen player opening the box. Either effect may prevail, depending on other parameters.

2.3.1 Welfare Implications and Comparative Statics

Recall that it is socially optimal to open the box immediately whenever (SR) holds. According to Theorem 2.1, the chosen player opens the box with certainty when the search cost is low enough, and with probability less than one otherwise. Consequently, when the search cost is high enough, there is welfare loss due to the delay.

Given the time interval between stages Δt and the probability π that the box is opened each round, I define the expected delay before the box is opened as

$$D(\pi,\Delta t) = 0 \cdot \pi + \Delta t \cdot \pi \cdot (1-\pi) + 2\Delta t \cdot (1-\pi)^2 \cdot \pi + \cdots = \Delta t \cdot \frac{1-\pi}{\pi}.$$

To understand the severity of the delay in a collaborative search environment, I analyze how the equilibrium probability $\pi(z)$ of opening the box varies with the search cost c and the safe option z.

Recall that by Theorem 2.1, $\pi(z)$ equals to one (meaning that there is no delay) for low enough values of the search cost, and is between zero and one (there is delay) when the search costs are sufficiently high. In particular,

- if $c > \bar{c} := S(\bar{z},F) \cdot \frac{2\delta(1-\delta)}{(2-\delta)}$, the search costs are so large that there is an interior solution $\pi \in (0,1)$ for every $z \in [0,\bar{z}]$;

- if $c < \underline{c} := S(0,F) \cdot \frac{2\delta(1-\delta)}{(2-\delta)}$, the search costs are so small that opening the box right away is strictly dominant for all $z \in [0,\bar{z}]$;

- if $c \in [\underline{c},\bar{c}]$, then there is an interior solution $\pi(z) \in (0,1)$ for $z < \tilde{z}$ that solves $c = S(\tilde{z},F) \cdot \frac{2\delta(1-\delta)}{(2-\delta)}$ and the box is opened right away for $z \geq \tilde{z}$.

The properties of $\pi(z)$ are summarized in Lemma 2.2 and illustrated in Figure 2.1.

Lemma 2.2. *Let $z \leq \bar{z}$ be the safe option and $\mathcal{B} = \{b\}$. In the symmetric equilibrium, the probability that the chosen player opens the box $\pi(z)$ has the following properties.*

1. *If* $c \geq \bar{c}$, *then* $\pi(z) \in (0,1)$ *and is strictly increasing and strictly convex;*

2. *if* $c \leq \underline{c}$, *then* $\pi(z) = 1$*;*

3. *if* $c \in (\underline{c}, \bar{c}_k)$, *then* $\pi(z)$ *exhibits the same properties as in case (1) for* $z \in [0, \tilde{z}]$, *and as in case (2) for* $z \in [\tilde{z}, \bar{z}]$.

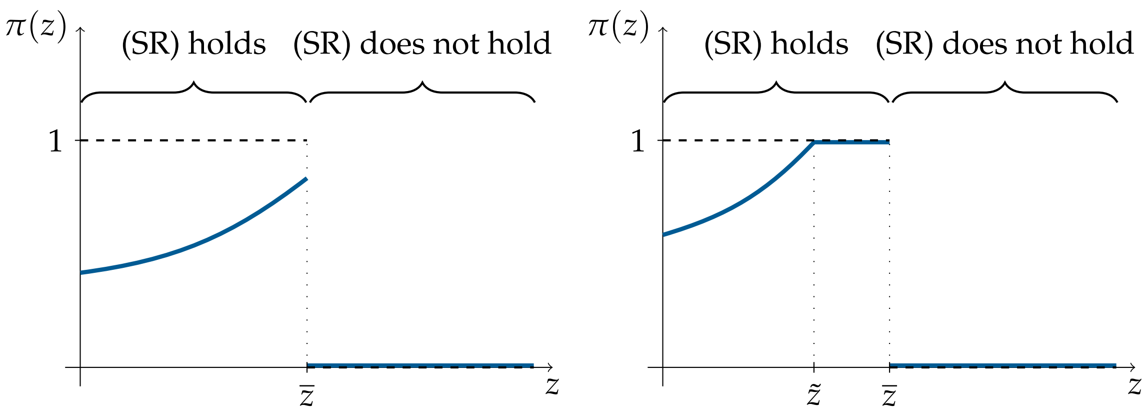

(a) $c \geq \bar{c}$. **(b)** $c \in (\underline{c}, \bar{c})$.

Figure 2.1. The probability $\pi(z)$ that the chosen player opens the box in the symmetric equilibrium. The dashed black lines represent the efficient level of $\pi(z)$.

Strikingly, the equilibrium probability opening the box is *increasing* in the value of the safe option at an increasing rate. The key to understanding this result is recalling that the safe option z is not the outside option for the chosen player because she cannot unilaterally deviate to collect it (her opponent must accept the termination offer). The effective outside option is to open the box today, and increasing z makes it more appealing to do so. Once the box is opened, the game ends, and *at least* z is collected. To remain indifferent between opening the box and not, the chosen player must rationally expect that the box is more likely to be opened in the future, which drives $\pi(z)$ up and the delay down.

Notice the lowest delay occurs at $z = \bar{z}$, when the chosen player is indifferent between *(i)* opening the box, *(ii)* not opening the box, and *(iii)* collecting the safe option

$\overline{z}$ (if she could). At that point, the expected benefit from opening the box $-c + \delta S(z, F)$ is the highest among all $z \in [0, \overline{z}]$ which, by the logic described above, makes opening the box sooner more desirable.[3] Delay is the highest when the safe option is zero when there seems to be the most to gain from opening the box. However, the expected reward from opening the box is actually the lowest compared to higher safe options. Thus, the indifference condition dictates that the box is the least likely to be opened.

Next, I discuss the comparative statics of the expected delay with respect to various model parameters. In general, as the value of opening the box today increases, so does the continuation value of not opening the box today. To remain indifferent, the chosen player must rationally expect a higher probability of the box being opened in the future, as prescribed by equation (2.4). Hence, *less delay* is associated with *higher reservation values* due to

- lower search cost c: the root of free riding lies in the unwillingness to pay the search cost. Reducing the search cost reduces the incentive to do nothing when chosen;
- "better" rewards: changing $F(\cdot)$ to $G(\cdot)$ such that $G(x) \leq F(x)\ \forall x$ leads to a higher expected reward $S(z, G)$ and higher π. Performing a mean-preserving spread on $F(\cdot)$ (making the box riskier) has the same effect.

The dynamic of the equilibrium probability of opening the box π with respect to the time between stages Δt is inconclusive. Recall that decreasing Δt increases the discount factor δ. On the one hand, increasing δ increases the value of opening the box today, and by the argument discussed above, decreases the delay. On the other hand, higher δ increases the players' willingness to wait for their opponent to perform the search, which drives the delay up. Either effect may prevail.

[3] For every $z < \overline{z}$ it is true that $-c + \delta S(z, F) < -c + \delta S(\overline{z}, F) = \overline{z}$.

2.4 Many Boxes

Let $\boldsymbol{b}$ be the box with the highest reservation value $\overline{\mathbf{z}}$ among the unopened boxes in $\mathcal{B}$, i.e.

$$\boldsymbol{b} := \arg\max_{b_k \in \mathcal{B}} \overline{z}_k \text{ and } \overline{\mathbf{z}} := \max_{b_k \in \mathcal{B}} \overline{z}_k.$$

In state $(z, \mathcal{B})$, if $z > \overline{\mathbf{z}}$, then opening *any* leads to the highest payoff for the chosen player. As such, any termination proposal in this state is accepted. Next, suppose $z \leq \overline{\mathbf{z}}$. When does player i reject a termination proposal?

To reject the termination proposal, player i must expect that the value of continuing the game and opening some boxes exceeds z. When player i rejects the proposal, she is chosen next period with probability $1/2$, and with probability $1/2$ she faces another termination proposal.[4] Because player i faces a termination proposal whenever she is not chosen, her problem is effectively the problem of an individual searcher who discounts her payoff with a factor of $\delta/(2-\delta)$, instead of δ. According to Weitzman (1979), player i's optimal policy is to open the boxes in the order of decreasing reservation value. When the highest reservation value becomes less than the maximum observed reward, she proposes termination, and her opponent agrees. Theorem 2.2 summarizes the best response of the opponent to a termination proposal.

Theorem 2.2. *Player i's best response to player j's termination proposal in state* $(z, \mathcal{B})$ *is to*

- *reject it if and only if* $z \leq \delta \overline{\Phi}_i(z, \mathcal{B})$,
- *accept it otherwise,*

where for any state $(\tilde{z}, \tilde{\mathcal{B}})$ *such that* $\tilde{z} \geq z$ *and* $\tilde{\mathcal{B}} \subseteq \mathcal{B}$, *player i's discounted average value*

[4] When termination offer is rejected, next stage begins with the same state of the problem. Since we are considering stationary Markov strategies, if player j proposes termination in this period, she also proposed termination in the next period.

function $\overline{\Phi}_i$ is recursively defined by

$$\delta\overline{\Phi}_i(\tilde{z},\tilde{\mathcal{B}}) = \frac{\delta}{2-\delta}\max\begin{cases}\tilde{z}, \\ \max\limits_{b_k\in\tilde{\mathcal{B}}}\Big(-c_k+\delta\big[\overline{\Phi}_i(\tilde{z},\tilde{\mathcal{B}}\setminus b_k)F_k(\tilde{z})+\int\limits_{z}^{+\infty}\overline{\Phi}_i(x,\tilde{\mathcal{B}}\setminus b_k)dF_k(x)\big]\Big),\end{cases}$$

$$\overline{\Phi}_i(\tilde{z},\varnothing)=\tilde{z}.$$

Notice that there is no explicit solution for when to accept the termination proposal, unlike in the one-box case. To make her decision, player i needs to iterate the search process forward until no more boxes are left. That said, when calculating the expected value of rejecting the proposal, she uses the socially optimal search protocol. The reason is that her discounted average value function $\delta\overline{\Phi}_i$ satisfies the Bellman equation of an individual searcher with discount factor $\delta/(2-\delta)$. Thus, when player j proposes termination whenever she is chosen, player i opens boxes in the order of decreasing reservation values, and proposes termination when the social planner would. That offer is accepted because it maximizes player j's payoff.

Two simple arguments provide sufficient conditions for when player i accepts and rejects the termination proposal.

Lemma 2.3. *In state $(z,\mathcal{B})$, let $\boldsymbol{b}$ be the box with the highest reservation value $\overline{\mathbf{z}}$ among the unopened boxes in $\mathcal{B}$ and let $\mathbf{z}^R$ be the rejection cutoff of that box. Then,*

- *$z \leq \mathbf{z}^R$ is sufficient to reject the termination offer,*
- *$z > \overline{\mathbf{z}}$ is sufficient to accept the termination offer.*

Recall that $\boldsymbol{b}$ is the "best" unopened box. Intuitively, if z is low enough that player i wants to wait and open just box $\boldsymbol{b}$, then she rejects the termination proposal. Conversely, if z is high enough that player i does not want to open every box $\boldsymbol{b}$ when she is eventually chosen, then she rejects the termination proposal.

Next, I consider the problem of the chosen player. If $z > \overline{\boldsymbol{z}}$, proposing termination is the weakly dominant strategy. In this case, her termination proposal is accepted by Lemma 2.3, the game ends, and both players receive z.

On the other hand, if $z \leq \overline{\boldsymbol{z}}$, then proposing termination is weakly dominated by doing nothing. Similarly to the one-box case, let us look for a symmetric mixed-strategy equilibrium. Given player j's mixed strategy, player i should be indifferent between every action she plays with positive probability. In particular, to be indifferent between opening two or more boxes, player i should expect the same average continuation value after opening each of them. However, according to Weitzman (1979), the value of opening the box with the highest continuation value exceeds the value of opening any other box.

Consequently, player i can only be indifferent between opening box $\boldsymbol{b}$ and doing nothing. When player i opens box $\boldsymbol{b}$, she pays the search cost $\boldsymbol{c}$, and moves on to the next stage of the problem with reset roles, fewer boxes, and a potentially higher uncovered reward. When she does nothing, next stage begins, roles are reset, but the state of the problem remains the same. As a result, the probability of opening box $\boldsymbol{b}$ solves player i's indifference condition between these two options, given that player j plays the same mixed strategy. The chosen player's search protocol in the mixed-strategy symmetric equilibrium is described in Theorem 2.3.

Theorem 2.3. *In state $(z, \mathcal{B})$, let $\boldsymbol{b} = (\boldsymbol{c}, \boldsymbol{F})$ be the box with the highest reservation value $\overline{\boldsymbol{z}}$ among the unopened boxes in $\mathcal{B}$. In the symmetric equilibrium, the chosen player*

- *if $z \leq \overline{\boldsymbol{z}}$,*
  - *opens box $\boldsymbol{b}$ with probability $\boldsymbol{\pi}(z, \mathcal{B}) = \min\left\{\frac{2(1-\delta)}{\delta \boldsymbol{c}} \cdot \left[-\boldsymbol{c} + \delta \overline{\Phi}^{\boldsymbol{b}}(z, \mathcal{B})\right], 1\right\}$,*
  - *does nothing otherwise;*

- *proposes termination if* $z > \bar{\mathbf{z}}$.

Here, in state $(\tilde{z}, \tilde{\mathcal{B}})$ *such that* $\tilde{z} \geq z$ *and* $\tilde{\mathcal{B}} \subseteq \mathcal{B}$*, the average value* $\overline{\Phi}$ *satisfies*

$$\delta\overline{\Phi}(\tilde{z}, \tilde{\mathcal{B}}) = \max\Big\{\tilde{z}, \max_{b_k \in \tilde{\mathcal{B}}}\big(-c_k + \delta\overline{\Phi}^{b_k}(\tilde{z}, \tilde{\mathcal{B}})\big)\Big\},$$

$$\overline{\Phi}(\tilde{z}, \varnothing) = \tilde{z},$$

and the average value $\overline{\Phi}^{b_k}$ *after opening box* $b_k = (c_k, F_k) \in \tilde{\mathcal{B}}$ *is*

$$\overline{\Phi}^{b_k}(\tilde{z}, \tilde{\mathcal{B}}) := \overline{\Phi}(\tilde{z}, \tilde{\mathcal{B}} \setminus b_k) F_k(\tilde{z}) + \int_{z}^{+\infty} \overline{\Phi}_i(x, \tilde{\mathcal{B}} \setminus b_k) dF_k(x).$$

Recall that in Weitzman (1979), the individual searcher makes decisions myopically: at each stage of the search process, she compares the reservation value of the best unopened box to the highest reward uncovered so far. With two searchers, this is no longer true. Both the chosen player and the opponent take the future into account when making decisions: the chosen player needs the future value function to calculate $\boldsymbol{\pi}$, while the opponent needs it to calculate the acceptance cutoff rule. At the same time, the *order* and the *stopping rule* on the equilibrium path are identical to that of Weitzman (1979): the box with the highest reservation value is opened next, and the game is terminated when the best uncovered reward exceeds the reservation value of all unopened boxes. The only difference is that, whenever $\boldsymbol{\pi} < 1$, box $\boldsymbol{b}$ is opened with a delay. Delay occurs because each player hopes that her opponent ends up opening the box and paying the search cost. Depending on the size of the search costs, *every* box may be opened with a delay.

While Theorem 2.3 implicitly defines $\boldsymbol{\pi}(z, \mathcal{B})$, the equilibrium probability that the chosen player opens box $\boldsymbol{b}$, the one-box case provides a convenient lower bound. In particular, it is straightforward to show that $\overline{\Phi}^{\boldsymbol{b}}(z, \mathcal{B}) \geq \overline{\Phi}^{\boldsymbol{b}}(z, \{\boldsymbol{b}\}) = S(z, F)$, meaning that the average value after opening box $\boldsymbol{b}$ is higher if $\boldsymbol{b}$ is not the only box left to open. Consequently, $\boldsymbol{\pi}(z,\mathcal{B}) \geq \boldsymbol{\pi}(z,\{\boldsymbol{b}\})$, where $\boldsymbol{\pi}(z,\{\boldsymbol{b}\})$ is the equilibrium probability of opening box $\boldsymbol{b}$ when $\boldsymbol{b}$ is the only box to be opened. Theorem 2.1 describes the one-box case and provides an explicit expression for $\boldsymbol{\pi}(z,\{\boldsymbol{b}\}) = \pi(z)$. This lower bound is useful for determining whether the box is opened with a delay or not. In particular, if $\boldsymbol{c} \leq S(z,\boldsymbol{F}) \cdot \frac{2\delta(1-\delta)}{2-\delta}$, Theorem 2.1 concludes that box $\boldsymbol{b}$ is opened without a delay. In that case, the chosen player prefers to open the box immediately, rather than wait for her counterpart to exert the low enough search cost.

Corollary 2.2. *In state $(z,\mathcal{B})$, let $\boldsymbol{b} = (\boldsymbol{c},\boldsymbol{F})$ be the box with the highest reservation value $\overline{\mathbf{z}}$ among the unopened boxes in $\mathcal{B}$. In the symmetric equilibrium, the value functions are*

$$\begin{aligned}
\Phi^{ch}(z,\mathcal{B}) &= \max\left\{z, -\boldsymbol{c} + \delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right\}, \\
\Phi^{op}(z,\mathcal{B}) &= \max\left\{z, \frac{2-\delta}{\delta} \cdot \left[-\boldsymbol{c} + \delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right]\right\}, \\
\overline{\Phi}(z,\mathcal{B}) &= \max\left\{z, \frac{1}{\delta} \cdot \left[-\boldsymbol{c} + \delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right]\right\},
\end{aligned}$$

where for any state $(\tilde{z},\tilde{\mathcal{B}})$ such that $\tilde{z} \geq z$ and $\tilde{\mathcal{B}} \subseteq \mathcal{B}$, the average value $\overline{\Phi}^{b_k}$ after opening box $b_k = (c_k, F_k) \in \tilde{\mathcal{B}}$ is

$$\overline{\Phi}^{b_k}(\tilde{z},\tilde{\mathcal{B}}) = \overline{\Phi}(\tilde{z},\tilde{\mathcal{B}} \setminus b_k)F_k(\tilde{z}) + \int_{z}^{+\infty} \overline{\Phi}_i(x,\tilde{\mathcal{B}} \setminus b_k)dF_k(x).$$

When chosen, each player does as well as an individual searcher because she bears the search cost whenever she decides to open a box. When not chosen, each player does better than an individual searcher because if a box is opened, she does not pay the search cost. As a result, on average, each player prefers to search with a colleague rather than on her own.

Recall from Section 2.3 that boxes with higher reservation values are opened faster. In the multi-box setting, boxes with the higher reservation value are opened earlier rather than later, suggesting that the delay grows over the search process. On the other hand, the delay also decreases with the value of the outside option z, which grows as more rewards are uncovered. In the end, it is unclear whether the players shirk more or less as the search process goes on.

2.5 Conclusion

This chapter examined a model of sequential search for a public good performed by a team of agents. I showed that group search results in the socially optimal search and stopping rule. However, delay occurs at every stage of the learning process because agents free ride. Overall, the team manager prefers to delegate research to individual agents, but each agent prefers to search with a teammate.

Chapter 2 is currently being prepared for submission for publication of the material. The dissertation author, Mariia Titova, is the sole author of this material.

Appendix A

Supplemental Material

A.1 Omitted Proofs for Chapter 1

Definition A.1.1. *In equilibrium* (σ, a, q), *for every receiver* $i \in I$, *let*

- $\mathcal{M}_i := \{m \subseteq \Omega \mid a_i(m) = 1\}$ *be the set of messages that convince receiver i to approve;*
- $\mathcal{W}_i := \{\omega \in \Omega \mid \exists m \in \mathcal{M}_i \text{ s.t. } \omega \in m\}$ *be the set of states in which the sender has access to at least one message that convinces receiver i to approve;*
- $\overline{\mathcal{W}}_i := \{\omega \in \Omega \mid \alpha_i(\omega) = 1\} \subseteq \mathcal{W}_i$ *be the set of states in which this receiver approves the proposal with probability* 1.

Note that $\mathcal{A}_i \subseteq \mathcal{W}_i$: if $\omega \in \mathcal{A}_i$, then $\{\omega\} \in \mathcal{M}_i$ since $q_i(\cdot \mid \{\omega\}) = p(\cdot \mid \{\omega\}) \in \mathcal{B}_i$. Also,

$$\forall \omega \in \overline{\mathcal{W}}_i, \quad \int_{\mathcal{M}_i} \sigma_i(m \mid \omega) dm = 1,$$

i.e. to convince the receiver in state ω with certainty, the sender must be sending her convincing messages, and convincing messages only.

Lemma A.1.1. *In equilibrium* (σ, a, q), *for every receiver* $i \in I$, *the set* $\overline{\mathcal{W}}_i \cup \mathcal{A}_i$ *satisfies receiver i's* (obedience) *constraint* $p(\cdot \mid X_i \cup \mathcal{A}_i) \in \mathcal{B}_i$.

Proof. Every message $m \in \mathcal{M}_i$ convinces the receiver to approve the proposal:

$$\int_{\text{supp } q_i(\cdot \mid m)} \delta_i(\omega) \cdot q_i(\omega \mid m) d\omega \geq 0.$$

Notice that supp $q_i(\cdot \mid m) \subseteq m$ because messages are verifiable. Furthermore, $m \subseteq \mathcal{W}_i$ because if $\omega \in m$ such that $m \in \mathcal{M}_i$, then $\omega \in \mathcal{W}_i$. On the equilibrium path, the inequality above becomes

$$\int_{\mathcal{W}_i} \delta_i(\omega) \cdot \frac{\sigma_i(m \mid \omega) \cdot p(\omega)}{\int_{\mathcal{W}_i} \sigma_i(m \mid \omega') \cdot p(\omega') d\omega'} d\omega \geq 0 \iff \int_{\mathcal{W}_i} \delta_i(\omega) \cdot \sigma_i(m \mid \omega) \cdot p(\omega) d\omega \geq 0.$$

Integrate the above inequality over all $m \in \mathcal{M}_i$:

$$\int_{\mathcal{M}_i} \int_{\mathcal{W}_i} \delta_i(\omega) \cdot \sigma_i(m \mid \omega) \cdot p(\omega) d\omega dm \geq 0.$$

Next, partition $\mathcal{W}_i$ into $\overline{\mathcal{W}}_i$, $\mathcal{A}_i \setminus \overline{\mathcal{W}}_i$, and $\mathcal{W}_i \setminus (\overline{\mathcal{W}}_i \cup \mathcal{A}_i)$ and observe that

$$\int_{\mathcal{M}_i} \int_{\overline{\mathcal{W}}_i} \delta_i(\omega) \cdot \sigma_i(m \mid \omega) p(\omega) d\omega dm$$

$$= \int_{\overline{\mathcal{W}}_i} \delta_i(\omega) p(\omega) \underbrace{\int_{\mathcal{M}_i} \sigma_i(m \mid \omega) dm}_{=1, \forall \omega \in \overline{\mathcal{W}}_i} d\omega = \int_{\overline{\mathcal{W}}_i} \delta_i(\omega) p(\omega) p\omega;$$

$$\int_{\mathcal{M}_i} \int_{\mathcal{A}_i} \delta_i(\omega) \sigma_i(m \mid \omega) p(\omega) d\omega dm$$

$$= \int_{\mathcal{A}_i} \underbrace{\delta_i(\omega)}_{\geq 0 \forall \omega \in \mathcal{A}_i} p(\omega) \underbrace{\int_{\mathcal{M}_i} \sigma_i(m \mid \omega) dm}_{\leq 1} d\omega \leq \int_{\mathcal{A}_i} \delta_i(\omega) p(\omega) p\omega;$$

$$\int_{\mathcal{M}_i} \int_{\mathcal{W}_i \setminus (\overline{\mathcal{W}}_i \cup \mathcal{A}_i)} \underbrace{\delta_i(\omega)}_{\leq 0 \forall \omega \notin \mathcal{A}_i} \sigma_i(m \mid \omega) p(\omega) d\omega dm \leq 0.$$

As a result,

$$\int_{\overline{\mathcal{W}}_i \cup \mathcal{A}_i} \delta_i(\omega) p(\omega) p\omega \geq \int_{\mathcal{M}_i} \int_{\mathcal{W}_i} \delta_i(\omega) \cdot \sigma_i(m \mid \omega) \cdot p(\omega) d\omega dm \geq 0 \Longrightarrow p(\cdot \mid \overline{\mathcal{W}}_i \cup \mathcal{A}_i) \in \mathcal{B}_i.$$

□

Proof of Theorem 1.1

PartI, every equilibrium outcome is deterministic. Suppose, on the contrary, that there exists an equilibrium outcome α such that $\alpha(\omega) \in (0,1)$ for some $\omega \in \Omega$. Then, $\alpha(\omega) > 0$ implies that $\sigma(m_\omega \mid \omega) > 0$ and $q(\cdot \mid m_\omega) \in \mathcal{B}$ for some $m_\omega \subseteq \Omega$. Then, the sender has a profitable deviation to $\widetilde{\sigma}(m_\omega \mid \omega) = 1$. His payoff in state ω increases from $\alpha(\omega) < 1$ to 1.

PartII: consider equilibrium (σ, a, q) with the set of approved states W. W must satisfy the sender's (IC) constraint, or else the sender can deviate to full disclosure. Next, using Definition A.1.1, $\overline{\mathcal{W}} = W$, and by Lemma A.1.1, the (obedience) constraint holds.

PartIII: suppose that $W \subseteq \Omega$ satisfies (IC) and (obedience). Let $\sigma(W \mid \omega) = \mathbb{1}(\omega \in W)$ and $\sigma(\Omega \setminus W \mid \omega) = \mathbb{1}(\omega \in \Omega \setminus W)$ be the sender's strategy. On the path, receiver only hears two messages, W and $\Omega \setminus W$, and her posterior belief is $q(\cdot \mid W) = p(\cdot \mid W) \in \mathcal{B}$ by (obedience) and $q(\cdot \mid \Omega \setminus W) = p(\cdot \mid \Omega \setminus W) \notin \mathcal{B}$. In words, the sender sends two messages and the receiver interprets them as a recommendation to approve or reject. Off-the-path, i.e. following any message $m \neq W, \Omega \setminus W$, let the receiver have "skeptical beliefs"

$$\forall m \subseteq \mathcal{A}, \ \text{supp}\, q(\cdot \mid m) \subseteq m, \text{so that } q(\cdot \mid m) \in \mathcal{B},$$

$$\forall m \not\subseteq \mathcal{A}, m \neq W, \ \text{supp}\, q(\cdot \mid m) \subseteq m \setminus \mathcal{A}, \text{so that } q(\cdot \mid m) \notin \mathcal{B}$$

that assign positive probability to states within the approval set if and only if the message comprises of these states only. Then, the sender does not have profitable deviations.

Proof of Theorem 1.2

Suppose that $\overline{W}$ solves a relaxed problem

$$\max_{W \subseteq \Omega} \int_W p(\omega) d\omega, \quad \text{subject to} \quad \int_W \delta(\omega) p(\omega) d\omega \geq 0. \tag{A.1}$$

Observe that since $\delta(\omega) \geq 0$ for every $\omega \in \mathcal{A}$, we have $\mathcal{A} \subseteq \overline{W}$. Hence, $\overline{W}$ also solves (1.1). Furthermore, the obedience constraint binds, i.e. $\int_W \delta(\omega) p(\omega) d\omega = 0$. If it does not, increase the value of the objective function while still satisfying the constraint.

Next, suppose that $\overline{W}$ is not characterized by a cutoff value of the receiver's net payoff of approval $\delta(\cdot)$. Then, there must exist $X, Y \subseteq \Omega$ such that

- $P(X) = P(Y) > 0$;
- $\forall \omega \in X, \forall \omega' \in Y, \delta(\omega) < \delta(\omega')$;
- $X \subseteq \overline{W}$ and $Y \subseteq \Omega \setminus \overline{W}$.

In words, the sender-preferred set of approved states includes a positive-measure set X, does not include a positive-measure set Y, yet the receiver has a higher net payoff of approving any state in Y over any state in X.

Now, let $W^* := \left(\overline{W} \setminus X\right) \cup Y$. Observe that the value of the objective function is the same for $\overline{W}$ and W^*:

$$P(\overline{W}) = P(\overline{W} \setminus X) + P(X) = P(\overline{W} \setminus X) + P(Y) = P(W^*).$$

The obedience constraint for $\overline{W}$ is

$$\int_{\overline{W}\setminus X} \delta(\omega)p(\omega)d\omega + \int_{X} \delta(\omega)p(\omega)d\omega = 0.$$

The obedience constraint for W^* is

$$\int_{W^*\setminus Y} \delta(\omega)p(\omega)d\omega + \int_{Y} \delta(\omega)p(\omega)d\omega > 0,$$

where the last inequality follows from

1. $W^* \setminus Y = \overline{W} \setminus X$, so the first term in both constraints is the same;

2. $\int_X \delta(\omega)p(\omega)d\omega < \int_Y \delta(\omega)p(\omega)d\omega$, so the second term in the second constraint is strictly larger.

We have found that W^* retains the sender's ex-ante utility at the same level as $\overline{W}$. At the same time, the obedience constraint for $\overline{W}$ is binding, whereas for W^* it is loose. Since the obedience constraint is binding at the optimum, W^*, and thus $\overline{W}$, do not maximize the objective function, which brings us to a contradiction. Hence, $\overline{W}$ is characterized by a cutoff value of the receiver's net payoff of approval $\delta(\cdot)$.

Next, I show that the sender-preferred equilibrium outcome $\overline{\alpha}(\omega) := \mathbb{1}(\omega \in \overline{W})$ is a commitment outcome. Consider the problem of finding the optimal commitment protocol $(\sigma^{BP}, a^{BP}, q^{BP})$. According to Kamenica and Gentzkow (2011), that problem may be simplified to finding an optimal *straightforward* experiment σ^{BP} that is supported on set $\{s^+, s^-\}$, where s^+ induces posterior $q^+ \in \mathcal{B}$ and recommends that the receiver approves the sender's proposal and s^- induces posterior $q^- \notin \mathcal{B}$ and recommends rejection. The outcome takes form of $\alpha(\omega) = \text{Prob}(s^+ \mid \omega)$, and the sender's problem under commitment becomes

$$\max_{\alpha} \int_{\Omega} \alpha(\omega)p(\omega)d\omega, \quad \text{subject to} \quad \begin{array}{c} \forall \omega \in \Omega, 0 \leq \alpha(\omega) \leq 1, \\ \int_{\Omega} \delta(\omega) \cdot \alpha(\omega)p(\omega)d\omega \geq 0. \end{array} \tag{A.2}$$

Observe that any commitment outcome α^{BP} is characterized by a cutoff value $c^{BP} > 0$, meaning that

$$\begin{aligned} &\alpha^{BP}(\omega) = 1 && \text{if } \delta(\omega) > -c^{BP}, \\ &\alpha^{BP}(\omega) \in [0,1], && \text{if } \delta(\omega) = -c^{BP}, \\ &\alpha^{BP}(\omega) = 0, && \text{if } \delta(\omega) < -c^{BP}. \end{aligned}$$

α^{BP} is characterized by a cutoff value for the same reason why $\overline{W}$ is. If it was not, then there exist $X, Y \subseteq \Omega$ such that

- $\int_X \alpha^{BP}(\omega)p(\omega)d\omega = \int_Y (1 - \alpha^{BP}(\omega))p(\omega)d\omega$;
- $\forall \omega \in X, \forall \omega' \in Y, \delta(\omega) < \delta(\omega')$;
- $\forall \omega \in X, \alpha^{BP}(\omega) > 0$ and $\forall \omega \in Y, \alpha^{BP}(\omega) < 1$.

Then, letting $\alpha^*(\omega) = \alpha^{BP}(\omega)$ for all $\omega \notin X \cup Y$, $\alpha^*(\omega) = 1$ if $\omega \in Y$, $\alpha^*(\omega) = 0$ if $\omega \in X$ leads to the same level of the objective function and a looser constraint.

Notice that the problem of finding the sender-preferred equilibrium set of approved states (A.1) is the sender's problem under commitment (A.2) with an additional constraint $\alpha(\omega) \in \{0,1\}$ for every $\omega \in \Omega$. Hence, if there exists a deterministic commitment outcome $\widetilde{\alpha}(\omega) := \mathbb{1}(\omega \in \widetilde{W})$, then $\widetilde{W} = \overline{W}$, meaning that the sender-preferred equilibrium outcome is a commitment outcome.

Next, taking an arbitrary commitment outcome α^{BP}, let $\mathcal{D} := \{\omega \in \Omega \mid 0 < \alpha^{BP}(\omega) < 1\}$ be the set of states the receiver approves and rejects with a positive probability. Since α^{BP} is characterized by the cutoff value c^{BP}, for every $\omega \in \mathcal{D}$, $\delta(\omega) = -c^{BP}$.

Next, let $\widetilde{\alpha}(\omega) = \alpha^{BP}(\omega)$ for all $\omega \notin \mathcal{D}$ and $\widetilde{\alpha}(\omega) = \mathbb{1}(\omega \in X)$ for all $\omega \in \mathcal{D}$, where $X \subseteq \mathcal{D}$ solves

$$\int_{\mathcal{D}} \alpha^{BP}(\omega) \cdot p(\omega) d\omega = \int_{\mathcal{D}} \widetilde{\alpha}(\omega) \cdot p(\omega) d\omega = P(X).$$

Now compare the commitment outcome α^{BP} and the candidate outcome $\widetilde{\alpha}$, keeping in mind that they only differ on $\mathcal{D}$. The value of the sender's objective function is the same:

$$\int_{\mathcal{D}} \alpha(\omega) p(\omega) d\omega = \int_{\mathcal{D}} \widetilde{\alpha}(\omega) p(\omega) d\omega = P(X);$$

the constraint is also the same:

$$\int_{\mathcal{D}} \underbrace{\delta(\omega)}_{=-c^{BP}, \forall \omega \in \mathcal{D}} \cdot \alpha^{BP}(\omega) p(\omega) d\omega = -c^{BP} \cdot \int_{\mathcal{D}} \widetilde{\alpha}(\omega) p(\omega) d\omega = -c^{BP} \cdot P(X).$$

Consequently, $\widetilde{\alpha}(\omega) = \mathbb{1}(\omega \in \mathcal{D}_1 \cup X)$ is a *deterministic* commitment outcome. As a result, the sender-preferred equilibrium outcome $\overline{\alpha}(\omega) = \mathbb{1}(\omega \in \mathcal{D}_1 \cup X)$ is a commitment outcome.

Proof of Theorem 1.3

$\Rightarrow$: consider equilibrium outcome α with the ex-ante utility of the sender $\overline{u}_s$. Let $X_i := \{\omega \in \Omega \mid \alpha_i(\omega) = 1\}$ be the set of states in which the sender convinces receiver $i \in I$ to approve the proposal with certainty. For very $i \in I$, set $W_i = X_i \cup \mathcal{A}_i$ satisfies the sender's (IC) constraint, and by Lemma A.1.1, W_i also satisfies receiver i's (obedience) constraint.

If $(W_1, \ldots, W_n)$ is the collection of the receivers' sets of approved states, then the sender's ex-ante utility equals

$$\int_{\Omega} u_s(\{i \in I \mid \omega \in W_i\}) \cdot p(\omega) d\omega,$$

because receiver i approves the proposal if and only if $\omega \in W_i$. What remains to show is that this expression equals $\overline{u}_s$, the ex-ante utility of the sender in the original equilibrium. That is true because if in state $\omega \in \Omega$ receiver $i \in I$ is convinced

- with certainty, then $\omega \in W_i$;
- with probability less than 1 and $\omega \in \mathcal{A}_i$, then her action is inconsequential to the sender's utility; adding ω to W_i does not change the sender's utility in state ω;
- with probability less than 1 and $\omega \notin \mathcal{A}_i$, then her action is inconsequential to the sender's utility; removing ω to W_i does not change the sender's utility in state ω.

As a result, $\overline{u}_s$ equals the expression above.

$\Leftarrow$: consider collection $(W_1, \ldots, W_n)$ of receivers' sets of approved states, each of which satisfies the sender's (IC) and receiver's (obedience) constraints. Then, let the sender's strategy satisfy $\sigma_i(W_i \mid \omega) = \mathbb{1}(\omega \in W_i)$ and $\sigma_i(\Omega \setminus W_i \mid \omega) = \mathbb{1}(\omega \in \Omega \setminus W_i)$, for every receiver $i \in I$. Then, given the same skeptical off-the-path beliefs of the receivers as in Theorem 1.1, none of the players have profitable deviations and the direct implementation constitutes an equilibrium.

Proof of Theorem 1.4

Consider the problem of finding the optimal commitment protocol $(\sigma^{BP}, a^{BP}, q^{BP})$. According to Kamenica and Gentzkow (2011), the problem may be simplified to finding an optimal *straightforward* experiment σ^{BP} that is supported on set $(S_1, \ldots, S_n)$, where $S_i = \{s_i^+, s_i^-\}$ is the private set of *straightforward* signal realizations of receiver $i \in I$. Signal realization s_i^+ induces posterior $q_i^+ \in \mathcal{B}_i$ and recommends that receiver i approves the proposal and s_i^- induces posterior $q_i^- \notin \mathcal{B}_i$ and recommends rejection. The commitment outcome is

$$\alpha^{BP} = \arg \max_{\alpha_i, \forall i \in I} \int_{\Omega} \sum_{T \subseteq 2^I} \alpha(T,\omega) \cdot u_s(T) \cdot p(\omega) d\omega, \text{ subject to } \forall i \in I$$

- $\forall \omega \in \Omega$, $0 \leq \alpha_i(\omega) \leq 1$;
- receiver i's obedience constraint $q_i^+ \in \mathcal{B}_i$, which is $\int_{\Omega} \delta_i(\omega) \cdot \alpha_i(\omega) \cdot p(\omega) d\omega \geq 0$,

where $\alpha(T,\omega) := \prod_{i \in T} \alpha_i(\omega) \cdot \prod_{j \in I \setminus T} (1 - \alpha_i(\omega))$ is the probability that receivers in $T \subseteq I$ approve the proposal and the receivers in $I \setminus T$ reject it. Notice that if $\alpha_i(\omega) = \mathbb{1}(\omega \in W_i^j)$ for all $i \in I$, then $\alpha(T,\omega) = \mathbb{1}(T = \{i \in I \mid \omega \in W_i\})$, and the sender's problem becomes

$$\max_{W_i \subseteq \Omega, \forall i \in I} \int_{\Omega} u_s(\{i \in I \mid \omega \in W_i\}) \cdot p(\omega) d\omega,$$

subject to receiver i's obedience constraint $p(\cdot \mid W_i) \in \mathcal{B}_i$, for all $i \in I$. What remains to show is that *(i)* there exists a deterministic commitment outcome, and *(ii)* every set of approved states W_i induced by that outcome satisfies the sender's (IC) constraint. I construct a deterministic commitment outcome $\widetilde{\alpha}$ in a sequence of steps.

<u>*Step0*</u>: start with $\widetilde{\alpha} = \alpha^{BP}$;

<u>*Step1*</u>: if, for some $i \in I$ and $\omega \in \mathcal{A}_i$, $\alpha_i^{BP}(\omega) < 1$, then let $\widetilde{\alpha}_i(\omega) = 1$. This weakly increases the objective, loosens receiver i's obedience constraint, and does not alter other receivers' obedience constraints. Note that this case only arises when the sender's payoff in state ω does not strictly increase in receiver i's action;

<u>*Step2*</u>: if, for some $i \in I$, this receiver's obedience constraint does not bind, then let $\widetilde{\alpha}_i(\omega) = 0$ for every ω such that $\alpha_i^{BP}(\omega) < 1$. In those states, the sender could have increased $\alpha_i^{BP}(\omega)$ by tightening receiver i's obedience constraint, but did not do so because convincing this receiver in this state did not increase his payoff;

<u>*Step3*</u>: if, for some receiver $i \in I$ and set $\mathcal{D} \subseteq \Omega$, $\alpha_i^{BP}(\omega) \in (0,1)$ for every $\omega \in \mathcal{D}$, and

this receiver's obedience constraint binds, then we follow the steps on the proof of Theorem 1.2. Rewrite receiver i's obedience constraint as

$$\int_{\mathcal{D}} \delta_i(\omega) \cdot \alpha_i^{BP}(\omega) p(\omega) d\omega = - \int_{\Omega \setminus \mathcal{D}} \delta_i(\omega) \cdot \alpha_i^{BP}(\omega) p(\omega) d\omega := \mathcal{I}_i.$$

Since $\alpha_i(\omega) \in (0,1)$ on $\mathcal{D}_i$, then $\delta_i(\omega)$ is constant on $\mathcal{D}_i$. Next, let $\widetilde{\alpha}_i(\omega) = \mathbb{1}(\omega \in X)$ for all $\omega \in \mathcal{D}$, where $X_i \subseteq \mathcal{D}_i$ solves

$$\int_{\mathcal{D}_i} \alpha_i(\omega) \cdot p(\omega) d\omega = \int_{X_i} p(\omega) d\omega = P(X_i).$$

Step4: if for $i \in I$ and $\omega \in \Omega$, $\alpha_i(\omega) \in \{0,1\}$, then let $\widetilde{\alpha}_i(\omega) = \alpha_i(\omega)$.

At this point, $\widetilde{\alpha}_i$, $\forall i \in I$, is a deterministic commitment outcome that satisfies all of the sender's (IC) constraints. Consequently, it is also the sender-preferred equilibrium outcome.

Proof of Corollary 1.1

1. By contradiction, suppose $\exists q \in \mathcal{B}_L$ such that $q \notin \mathcal{B}_i$. Notice that because $v_i < v_L < \omega_0$, we have $|v_i - \omega| = |v_i - v_L| + |v_L - \omega_0|$, that is, v_L is located between v_i and ω_0. Then,

$$q \notin \mathcal{B}_i \iff \mathbb{E}_q[|v_i - \omega|] > |v_i - \omega_0| - \varepsilon$$

$$= |v_i - v_L| + |v_L - \omega_0| - \varepsilon \overset{q \in \mathcal{B}_L}{\geq} |v_i - v_L| + \mathbb{E}_q[|v_L - \omega|].$$

   We have arrived at a violation of the triangle inequality, which for every $\omega \in \Omega$ states that $|v_i - \omega| \leq |v_i - v_L| + |v_L - \omega|$. $\mathcal{B}_L \subseteq \mathcal{B}_i$ implies that $\mathcal{A}_L \subseteq \mathcal{A}_i$, because $\omega \in \mathcal{A}$ if and only if belief that puts probability 1 on state ω belongs to $\mathcal{B}$. The proof for voter R is analogous.

2. By the definition of the set of approval beliefs, for every $i \in \{L, R\}$

$$q \in \mathcal{B}_i \iff \int_{\Omega} |v_i - \omega| \cdot q(\omega) d\omega \leq |v_i - \omega_0| - \varepsilon.$$

Adding up the right-hand sides for $i \in \{L, R\}$,

$$q \in \mathcal{B}_L \cap \mathcal{B}_R \Longrightarrow \int_{\Omega} [|v_L - \omega| + |\omega - v_R|] \cdot q(\omega) d\omega$$

$$\leq |v_L - \omega_0| + |\omega_0 - v_R| - 2\varepsilon < |v_L - v_R|.$$

The right hand side violates the triangle inequality, according to which $|v_L - \omega| + |\omega - v_R| \geq |v_L - v_R|$ for every $\omega \in \Omega$. This proves that $\mathcal{B}_L \cap \mathcal{B}_R = \emptyset$. Since $\mathcal{B}_i$ includes beliefs that put probability 1 on $\omega \in \mathcal{A}_i$ for $i \in \{L, R\}$, $\mathcal{A}_L \cap \mathcal{A}_R = \emptyset$.

Proof of Theorem 1.5

Recall that $\delta_i(\omega) = |v_i - \omega_0| - |v_i - \omega| - \varepsilon$ is voter i's net payoff of approval. Her (obedience) constraint is:

$$p(\cdot \mid W_i) \in \mathcal{B}_i \iff \int_{W_i} \delta_i(\omega) p(\omega) d\omega \geq 0$$

$$\iff \int_{W_i \setminus \mathcal{A}_i} - \underbrace{\delta_i(\omega)}_{<0, \forall \omega \notin \mathcal{A}_i} \cdot p(\omega) d\omega \leq \int_{\mathcal{A}_i} \underbrace{\delta_i(\omega)}_{>0, \forall \omega \in \mathcal{A}_i} p(\omega) d\omega := \mathcal{I}_i.$$

Notice that when $\omega \notin \mathcal{A}_i$, $-\delta_i(\omega)$ reflects the distance from point ω to the approval set of voter i. The voter's obedience constraint states that the expected distance from the challenger to the voter's approval set must not exceed a known quantity $\mathcal{I}_i$, which reflects how persuadable this voter is. For example, Figure A.1 – part *(a)* illustrates how under uniform prior, voter R's obedience constraint states that the area under the function $\delta_R(\omega)$ over the approval set (it equals $\mathcal{I}_R$) must exceed the area over the same function outside of the approval set. Adding point x to $W_L \cap W_R$ increases the objective function by $p(x)$ and costs $-\delta_i(x)p(x) \cdot \mathbb{1}(x \notin \mathcal{A}_i)$ to each voter $i \in \{L, R\}$. Consequently, $x \notin \mathcal{A}_i$ is "cheaper" in terms of i's obedience constraint than $y \notin \mathcal{A}_i$ if $\delta_i(x) \geq \delta_i(y)$. Points in the approval set of the voter are "free" in terms of the obedience constraint of that voter.

Relying on these observations, the following arguments, illustrated in Figure A.1, part *(b)*, prove that $W_L = [a_L, b_L]$ with $a_L \leq a$ and $b_L > \omega_0 - \varepsilon$. Letting $a = 2v_L - \omega_0 + \varepsilon$ and $b = 2v_R - \omega_0 - \varepsilon$ be the left boundary of L's approval set and right boundary of R's approval set, respectively, we get

- $[a, \omega_0 - \varepsilon] \subseteq W_L$ because it is the approval set of voter L;
- if $x_1 \in [0, a)$ and $x \in W_L$, then $\forall y_1 \in [\omega_0 - \varepsilon, b]$ such that $|a - x_1| \geq |y_1 - \omega_0 + \varepsilon|$, $y_1 \in W_L$;
- if $x_1 \in [0, a)$ and $x \in W_L$, then $\forall x \in (x_1, a]$, $x \in W_L$;
- if $y_1 \in (\omega_0 - \varepsilon, b]$ and $y_1 \in W_L$, then $\forall y \in [\omega_0 - \varepsilon, y_1)$, $y \in W_L$;
- if $y_2 \in (b, 1]$ and $y_2 \in W_L$, then $\forall y \in [\omega_0 - \varepsilon, y_2)$, $y \in W_L$.

Finally, $b_L > \omega_0 - \varepsilon$ because $\mathcal{I}_L > 0$, and for small enough ε, $b_L > \omega_0$.

Proof of Theorem 1.6

Given convincing message $[a_R, b_R] \supseteq [\omega_0 + \varepsilon, 2v_R - \omega - \varepsilon] = \mathcal{A}_R$, voter R's obedience constraint becomes

$$\int_{a_R}^{\omega_0+\varepsilon} (\omega_0 - \omega)p(\omega)d\omega + \int_{2v_R-\omega_0-\varepsilon}^{b_R} (\omega_0 - \omega - 2v_R)p(\omega)d\omega$$

$$\leq \int_{\omega_0+\varepsilon}^{v_R} (\omega - \omega_0)p(\omega)d\omega + \int_{v_R}^{2v_R-\omega_0-\varepsilon} (2v_R + \omega - \omega_0)p(\omega)d\omega.$$

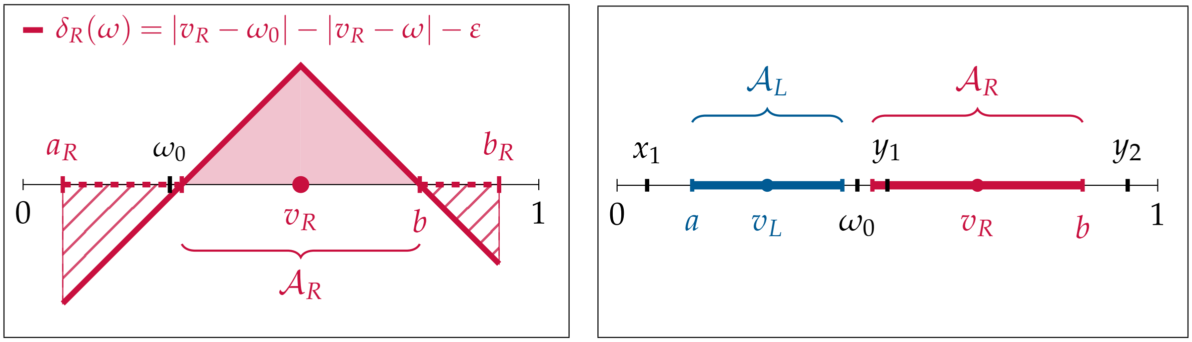

(a) Voter R's net payoff of approval. Under uniform prior, her obedience constraint states that the solid area exceeds the dashed area.

(b) Approval sets of the voters and points x_1, y_1, y_2.

Figure A.1. Why sender-preferred convincing messages are intervals.

The derivative of the left-hand side of this inequality with respect to v_R is negative and equals $-2P([2v_R - \omega_0 - \varepsilon, b_R])$, while the derivative of the right-hand side with respect to v_R is positive and equals $2P([v_R, 2v_R - \omega_0 - \varepsilon])$. Consequently, as v_R increases, voter R's obedience constraint loosens, and that is true for any prior distribution. Hence, the solution, specifically, the challenger's ex-ante odds of winning, can only improve.

Now suppose $|v_L - \omega_0| = |v_R - \omega_0|$ and let $a = 2v_L - \omega_0 + \varepsilon$ be the left boundary of L's approval set, and let $b = 2v_R - \omega_0 - \varepsilon$ be the right boundary of R's approval set. Voters' (obedience) constraints are symmetric about ω_0, implying that the solution $\overline{W}_L \cap \overline{W}_R$ is symmetric, as well, i.e. $|b_L - \omega_0| = |\omega_0 - a_R|$. Here, a_R solves voter R's obedience constraint $-\int_{a_R}^{\omega_0+\varepsilon} \delta_R(\omega)d\omega = \int_{\omega_0+\varepsilon}^{b} \delta_R(\omega)d\omega > 0$. For small enough ε, $a_R < \omega_0 - \varepsilon$ (from obedience, $a_R < \omega_0 + \varepsilon$) and $b_R > \omega_0 + \varepsilon$, implying that $a_R < \omega_0 < b_L$ and the challenger swings the election with a positive probability.

As v_R increases, voter R's obedience constraint loosens, while voter L's obedience constraint remains the same. An increase in the value of the objective function is thus obtained by decreasing both a_R and b_L because

- b_L cannot increase because it is obtained from the binding obedience constraint of voter L that was not affected by the change in v_R;

- for high enough v_R, $\int_{\omega_0+\varepsilon}^{b} \delta_R(\omega)d\omega > -\int_{a_R}^{\omega_0+\varepsilon} \delta_R(\omega)d\omega$, meaning that the optimal message that convinces voter L has to be optimally shifted to the left and becomes $[a_L, b_L']$, with $a_L < a$ and $b_L' < b_L$;

- voter L's obedience constraint: $\int_{a}^{\omega_0-\varepsilon} \delta_L(\omega)d\omega \geq -\int_{a_L}^{a} \delta_L(\omega)d\omega - \int_{\omega_0-\varepsilon}^{b_L} \delta_L(\omega)d\omega$. Because b_L is further from v_L than a is, removing $b_L - d$ from the message that convinces voter L and replacing it with $a - d$ (for some $d > 0$) loosens voter L's obedience constraint and keeps the value of the objective the same. That means that as b_L decreases, a_L decreases even more;

- the above argument stops working when $b_L - \omega_0 = a - a_L$. At that point, voter R is so persuadable that only voter L's constraint binds. The problem boils down to persuading just voter L, is characterized in Theorem 1.2, and no further changes in a_R and b_L are observed.

A.2 Omitted Proofs for Chapter 2

Lemma A.2.1. *Let* $H(z,b) := \frac{\delta}{2-\delta} \cdot [-c + S(z,F)]$ *and let* z^R *solve* $H(z^R, b) = z^R$. *Then,* (IR) *holds if and only if* $z \leq z^R$.

Proof. Since $\frac{\partial S(z,F)}{\partial z} = F(z) \in [0,1]$, $H(z,b)$ is increasing in z at the rate less than one. Since $z^R = H(z^R, b)$, $z \leq H(z,b)$ if and only if $z \leq z^R$. □

Proof of Theorem 2.1 and Corollary 2.1

With one box, for each z we have two states of the world, the box being closed and the box being open. With a slight abuse of notation, below I call z the state of the world when the box is closed. Once the box is open, boundary condition (2.3) states that both value functions equal z.

When player i faces a termination proposal, that is, when player j plays $\boldsymbol{a}_j$ such that $\boldsymbol{a}_j^{ch} = T$, her Bellman equation in state z is

$$\Phi_i^{op}(z;\boldsymbol{a}_j) = \max_{a_i^{op} \in \{0,1\}} \begin{cases} z, & \text{if } a_i^{op} = 1, \\ \delta\overline{\Phi}_i(z;\boldsymbol{a}_j), & \text{if } a_i^{op} = 0. \end{cases}$$

To calculate her average value function $\overline{\Phi}_i$, we first calculate her value function when she is chosen as

$$\Phi_i^{ch}(z;\boldsymbol{a}_j) = \max_{a_i^{ch} \in \{\varnothing,T,b\}} \begin{cases} z, & \text{if } a_i^{ch} = T \wedge \boldsymbol{a}_j^{op} = 1, \\ \delta\overline{\Phi}_i(z;\boldsymbol{a}_j), & \text{if } a_i^{ch} = T \wedge \boldsymbol{a}_j^{op} = 0, \text{ or } a_i^{ch} = \varnothing, \\ -c + \delta S(z,F), & \text{if } a_i^{ch} = b. \end{cases}$$

Player i rejects a termination proposal if and only if $\delta\overline{\Phi}_i(z;\boldsymbol{a}_j) \geq z$, that is, if she expects a higher payoff from continuing the game than from terminating it. Since player j is not helping her open the box (she plays $\boldsymbol{a}_j^{ch} = T$), player i would only reject the termination proposal if she expects to open the box herself when chosen, i.e. $\Phi_i^{ch}(z;\boldsymbol{a}_j) = -c + \delta S$. Then, using the facts that *(i)* it may take her time to get chosen, $\Phi_i^{op}(z;\boldsymbol{a}_j) = \delta\overline{\Phi}_i(z;\boldsymbol{a}_j)$, and *(ii)* her average value equals $\overline{\Phi}_i = \frac{\Phi^{ch}+\Phi^{op}}{2}$, we get that player i rejects the termination offer if and only if

$$\frac{\delta}{2-\delta} \cdot [-c + \delta S(z,F)] \geq z.$$

Note that the inequality above is the (IR) condition and it holds if and only if $z \leq z^R$ according to Lemma A.2.1.

Next consider the decision of player i when she is chosen. First suppose that the social rational condition (SR) does not hold, i.e. $z > \overline{z}$. In this case, player i does not want to open the box, she would rather propose termination, since she knows that the offer will be accepted. As a result, $\Phi_i^{ch}(z) = \Phi_i^{op}(z) = z$.

Next suppose that (SR) holds, i.e. $z \leq -c + \delta S(z, F)$. Player i's value function when she is chosen becomes

$$\Phi_i^{ch}(z; \boldsymbol{a}_j) = \max_{a_i^{ch} \in \{\varnothing, b\}} \begin{cases} \delta \overline{\Phi}_i(z; \boldsymbol{a}_j), & \text{if } a_i^{ch} = \varnothing, \\ -c + \delta S(z, F), & \text{if } a_i^{ch} = b, \end{cases}$$

since her payoff from proposing termination is weakly worse than her payoff of opening the box. When she is the opponent, her value function is

$$\Phi_i^{op}(z; \boldsymbol{a}_j) = \begin{cases} \delta \overline{\Phi}_i(z; \boldsymbol{a}_j), & \text{if } \boldsymbol{a}_j^{ch} = \varnothing, \\ \delta S(z, F), & \text{if } \boldsymbol{a}_j^{ch} = b. \end{cases}$$

Let π_j be player j's mixed strategy when she is chosen. $\pi_j(z)$ denotes the probability that player j opens the box (i.e. plays action $\boldsymbol{a}_j^{ch} = b$) in state z. Then, player i's value function when she is the opponent is

$$\Phi_i^{op}(z; \pi_j) = \pi_j \cdot \delta S(z, F) + (1 - \pi_j) \cdot \delta \overline{\Phi}_i(z; \pi_j).$$

For player i to play a mixed strategy in the symmetric equilibrium, she must be indifferent between opening the box and doing nothing, i.e.

$$\Phi_i^{ch}(z; \pi_j) = \delta \overline{\Phi}_i(z; \pi_j) = -c + \delta S(z, F).$$

Solving these equations given that $\overline{\Phi}_i = \frac{\Phi^{ch} + \Phi^{op}}{2}$, we get that in the symmetric equilibrium

$$\Phi^{ch}(z) = \delta \overline{\Phi}(z) = -c + \delta S(z, F), \quad \Phi^{op}(z) = \frac{2 - \delta}{\delta} \cdot [-c + \delta S(z, F)],$$

$$\pi(z) = \frac{2(1-\delta)}{\delta c} \cdot [-c + \delta S(z,F)].$$

Note that there is no interior solution $\pi \in (0,1)$ when opening the box strictly dominates doing nothing, which happens if and only if

$$-c + \delta S(z,F) > \delta \overline{\Phi}_i(z;\pi_j) \text{ and } \Phi_i^{op}(z;\pi_j) = \delta S(z,F) \iff c < S(z,F) \cdot \frac{2\delta(1-\delta)}{2-\delta}.$$

Summarizing our findings for the cases when (SR) does and does not hold, the value functions in the symmetric equilibrium are

$$\begin{aligned}\Phi^{ch}(z) &= \max\{z, -c + \delta S(z,F)\},\\ \Phi^{op}(z) &= \max\left\{z, \frac{2-\delta}{\delta} \cdot [-c + \delta S(z,F)]\right\},\\ \overline{\Phi}(z) &= \max\left\{z, \frac{1}{\delta} \cdot [-c + \delta S(z,F)]\right\}.\end{aligned}$$

Proof of Lemma 2.1

Firstly, let us find how the reservation value $\overline{z}$ changes with c and δ (Δt enters expressions via $\delta = e^{-r\Delta t}$ only). Recall that $\overline{z}$ is implicitly defined by

$$\overline{z} = -c + \delta S(z,F).$$

Differentiating this equation with respect to the variables of interest, we get

$$\frac{\partial \overline{z}}{\partial c} = -1 + \delta F(\overline{z}) \cdot \frac{\partial \overline{z}}{\partial c} \Rightarrow \frac{\partial \overline{z}}{\partial c} = \frac{-1}{1-\delta F(\overline{z})} < 0,$$

$$\frac{\partial \overline{z}}{\partial \delta} = S(z,F) + \delta F(\overline{z}) \cdot \frac{\partial \overline{z}}{\partial \delta} \Rightarrow \frac{\partial \overline{z}}{\partial \delta} = \frac{S(z,F)}{1-\delta F(\overline{z})} > 0,$$

using the fact that $\frac{\partial S(z)}{\partial z} = F(z)$.

Since the reservation value increases as c decreases and δ increases, the condition for opening the box $z \leq \bar{z}$ is met for a wider range of z. Since $\frac{\partial \Phi^{ch}(z)}{\partial \delta} = S(z,F) > 0$ and $\frac{\partial \overline{\Phi}(z)}{\partial \delta} = \frac{c}{\delta^2} > 0$ when $z \leq \bar{z}$, it follows directly that $\Phi^{ch}(z)$ and $\overline{\Phi}(z)$ increase. The same argument can be applied for dynamics of $\Phi^{op}(z)$ with respect to c. The sign of $\frac{\partial \Phi^{op}(z)}{\partial \delta}$ is inconclusive.

Proof of Lemma 2.2

When there is an interior solution,

$$\pi(z) = \frac{2(1-\delta)}{\delta c} \cdot [-c + \delta S(z,F)].$$

Then,

$$\frac{\partial \pi(z)}{\partial z} = \frac{2(1-\delta)}{c} \cdot F(z) > 0, \text{ and } \frac{\partial^2 \pi(z)}{\partial z^2} = \frac{2(1-\delta)}{c} \cdot f(z) > 0.$$

Proof of Theorem 2.2

When player i faces a termination proposal, that is, when player j plays $\boldsymbol{a}_j$ such that $\boldsymbol{a}_j^{ch} = T$ and $\boldsymbol{a}_j^{op} = 1$ whenever $z > \bar{\mathbf{z}}$, her Bellman equation in state $(z,\mathcal{B})$ is

$$\Phi_i^{op}(z,\mathcal{B};\boldsymbol{a}_j) = \max_{a_i^{op} \in \{0,1\}} \begin{cases} z, & \text{if } a_i^{op} = 1, \\ \delta \overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j), & \text{if } a_i^{op} = 0. \end{cases}$$

Consequently, player i rejects the proposal if and only if $z \leq \delta \overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j)$. To calculate her average value function $\overline{\Phi}_i$, we first calculate her value function when she is chosen:

$$\Phi_i^{ch}(z,\mathcal{B};\boldsymbol{a}_j) =$$

$$\max_{a_i^{ch}\in\{\varnothing,T\}\cup\mathcal{B}} \begin{cases} z, & \text{if } a_i^{ch}=T\wedge z>\overline{\boldsymbol{z}}, \\ \delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j), & \text{if } a_i^{ch}=T\wedge z\le\overline{\boldsymbol{z}}, \text{ or } a_i^{ch}=\varnothing, \\ -c_k+\delta\big[\overline{\Phi}_i(z,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)F_k(z)+\int_z^{+\infty}\overline{\Phi}_i(x,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)dF_k(x)\big], & \text{if } a_i^{ch}=b_k. \end{cases}$$

Using the definition of player i's average value function, $\overline{\Phi}_i = \frac{\Phi^{ch}+\Phi^{op}}{2}$, we get that

$$\delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j) = \frac{\delta}{2-\delta}\max\begin{cases} z, \\ \max_{b_k\in\mathcal{B}}\Big(-c_k+\delta\overline{\Phi}_i^{b_k}(z,\mathcal{B};\boldsymbol{a}_j)\Big), \end{cases}$$

where

$$\overline{\Phi}_i^{b_k}(z,\mathcal{B};\boldsymbol{a}_j) := \overline{\Phi}_i(z,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)F_k(z)+\int_z^{+\infty}\overline{\Phi}_i(x,\mathcal{B}\setminus b_k;\boldsymbol{a}_j)dF_k(x)$$

This equation with boundary condition (2.3) allow us to calculate $\delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j)$ implicitly. Player i's best response to a termination proposal in state $(z,\mathcal{B})$ is to reject it if and only if $z\le\delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j)$.

Proof of Lemma 2.3

Letting $\boldsymbol{b}=(\boldsymbol{c},\boldsymbol{F})$, we have

- $z>\overline{\boldsymbol{z}} \iff z>-\boldsymbol{c}+\delta S(z,\boldsymbol{F})$, meaning that $\delta\overline{\Phi}_i(z,\mathcal{B})=\frac{\delta}{2-\delta}\cdot z$ and $\Phi_i^{op}=z$, and the termination offer is accepted;
- $z\le\boldsymbol{z}^R \iff z\le\frac{\delta}{2-\delta}\cdot[-\boldsymbol{c}+\delta S(z,\boldsymbol{F})]$, meaning that

$$\frac{\delta}{2-\delta}\cdot\Big\{-\boldsymbol{c}+\delta\big[\overline{\Phi}_i(z,\mathcal{B}\setminus\boldsymbol{b})\boldsymbol{F}(z)+\int_z^{+\infty}\overline{\Phi}_i(x,\mathcal{B}\setminus\boldsymbol{b})d\boldsymbol{F}(x)\big]\Big\}$$

$$\geq \frac{\delta}{2-\delta} \cdot [-\boldsymbol{c} + \delta S(z, \boldsymbol{F})] \geq z,$$

and the termination offer is rejected in favor of opening at least one box.

Proof of Theorem 2.3 and Corollary 2.2

Consider the problem of the chosen player. If $z > \bar{\boldsymbol{z}}$, proposing termination is the weakly dominant action. In this case, her termination proposal is accepted by Lemma 2.3, game ends, and both players receive z.

Next suppose that $z \leq \bar{\boldsymbol{z}}$. Player i's value function when she is chosen is

$$\Phi_i^{ch}(z, \mathcal{B}; \boldsymbol{a}_j) = \max_{a_i^{ch} \in \varnothing \cup \mathcal{B}} \begin{cases} \delta \overline{\Phi}_i(z, \mathcal{B}; \boldsymbol{a}_j), & \text{if } a_i^{ch} = \varnothing, \\ -c_k + \delta \overline{\Phi}_i^{b_k}(z, \mathcal{B}; \boldsymbol{a}_j), & \text{if } a_i^{ch} = b_k \in \mathcal{B}, \end{cases}$$

where

$$\overline{\Phi}_i^{b_k}(z, \mathcal{B}; \boldsymbol{a}_j) := \overline{\Phi}_i(z, \mathcal{B} \setminus b_k; \boldsymbol{a}_j) F_k(z) + \int_z^{+\infty} \overline{\Phi}_i(x, \mathcal{B} \setminus b_k; \boldsymbol{a}_j) dF_k(x)$$

denotes the average value function after opening box $b_k \in \mathcal{B}$.

When player i is the opponent, her value function is

$$\Phi_i^{op}(z, \mathcal{B}; \boldsymbol{a}_j) = \begin{cases} \delta \overline{\Phi}_i(z, \mathcal{B}; \boldsymbol{a}_j), & \text{if } \boldsymbol{a}_j^{ch} = \varnothing, \\ \delta \overline{\Phi}_i^{b_k}(z, \mathcal{B}; \boldsymbol{a}_j), & \text{if } \boldsymbol{a}_j^{ch} = b_k \in \mathcal{B}. \end{cases}$$

Let π_j be player j's mixed strategy when she is chosen. $\pi_j^k(z, \mathcal{B})$ denotes the probability that player j opens box $b_k \in \mathcal{B}$ and $\pi_j^{\varnothing}(z, \mathcal{B})$ denotes the probability that she does nothing. Then, player i's value function when she is the opponent is

$$\Phi_i^{op}(z, \mathcal{B}; \pi_j) = \pi_j^{\varnothing}(z, \mathcal{B}) \cdot \delta \overline{\Phi}_i(z, \mathcal{B}; \boldsymbol{a}_j) + \sum_{b_k \in \mathcal{B}} \pi_j^k(z, \mathcal{B}) \cdot \delta \overline{\Phi}_i^{b_k}(z, \mathcal{B}; \boldsymbol{a}_j).$$

To play a mixed strategy π_i, player i must be indifferent between all the actions that she plays with positive probability. i.e.

$$\delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{a}_j) = -c_k + \delta\overline{\Phi}_i^{b_k}(z,\mathcal{B};\boldsymbol{a}_j), \text{ for all } b_k \in \mathcal{B} \text{ such that } \pi_i^k(z,\mathcal{B}) > 0.$$

Recall that according to Weitzman (1979), when facing two boxes with different reservation values, strictly prefers to open the box with the higher reservation value first. Consequently, in the symmetric equilibrium, player i cannot be indifferent between opening all the boxes. She can only be indifferent between doing nothing and opening the box with the highest reservation value. Let $\boldsymbol{\pi}_j > 0$ be the probability that player j opens box $\boldsymbol{b}$ and $1 - \boldsymbol{\pi}_i$ be the probability that she does nothing. Player i's value functions become

$$\Phi_i^{op}(z,\mathcal{B};\boldsymbol{\pi}_j) = (1-\boldsymbol{\pi}_j)\cdot\delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{\pi}_j) + \boldsymbol{\pi}_j\cdot\delta\overline{\Phi}_i^{\boldsymbol{b}}(z,\mathcal{B};\boldsymbol{\pi}_j),$$

$$\Phi_i^{ch}(z,\mathcal{B};\boldsymbol{\pi}_j) = \delta\overline{\Phi}_i(z,\mathcal{B};\boldsymbol{\pi}_j) = -\boldsymbol{c} + \delta\overline{\Phi}_i^{\boldsymbol{b}}(z,\mathcal{B};\boldsymbol{\pi}_j).$$

Solving these equations, we get that in the symmetric equilibrium each player opens box $\boldsymbol{b}$ with probability

$$\boldsymbol{\pi}(z,\mathcal{B}) = \min\left\{\frac{2(1-\delta)}{\delta\boldsymbol{c}}\cdot\left[-\boldsymbol{c}+\delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right],1\right\}.$$

Summarizing our findings for the cases when (SR) does and does not hold for box $\boldsymbol{b}$, the value functions in the symmetric equilibrium are

$$\begin{aligned}
\Phi^{ch}(z,\mathcal{B}) &= \max\left\{z, -\boldsymbol{c}+\delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right\},\\
\Phi^{op}(z,\mathcal{B}) &= \max\left\{z, \frac{2-\delta}{\delta}\cdot\left[-\boldsymbol{c}+\delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right]\right\},\\
\overline{\Phi}(z,\mathcal{B}) &= \max\left\{z, \frac{1}{\delta}\cdot\left[-\boldsymbol{c}+\delta\overline{\Phi}^{\boldsymbol{b}}(z,\mathcal{B})\right]\right\}.
\end{aligned}$$

Bibliography

ADMATI, ANAT R. and MOTTY PERRY (1991), "Joint Projects without Commitment", *The Review of Economic Studies*, 58, 2 (Apr. 1991), p. 259.

ALBRECHT, JAMES, AXEL ANDERSON, and SUSAN VROMAN (2010), "Search by Committee", *Journal of Economic Theory*, 145, 4 (July 2010), pp. 1386-1407.

ALONSO, RICARDO and ODILON CÂMARA (2016), "Persuading Voters", *American Economic Review*, 106, 11 (Nov. 2016), pp. 3590-3605.

ARIELI, ITAI and YAKOV BABICHENKO (2019), "Private Bayesian Persuasion", *Journal of Economic Theory*, 182 (July 2019), pp. 185-217.

ARTABE, ALAITZ and JAVIER GARDEAZABAL (2014), "Strategic Votes and Sincere Counterfactuals", *Political Analysis*, 22, 2 (Jan. 2014), pp. 243-257.

BABICHENKO, YAKOV and SIDDHARTH BARMAN (2016), "Computational Aspects of Private Bayesian Persuasion", pp. 1-16.

BARDHI, ARJADA and YINGNI GUO (2018), "Modes of Persuasion toward Unanimous Consent", *Theoretical Economics*, 13, 3, pp. 1111-1149.

BLACK, DUNCAN (1948), "On the Rationale of Group Decision-making", *Journal of Political Economy*, 56, 1, pp. 23-34.

BOARD, OLIVER (2009), "Competition and Disclosure", *The Journal of Industrial Economics*, 57, 1 (Feb. 2009), pp. 197-213.

BOLESLAVSKY, RAPHAEL and CHRISTOPHER COTTON (2015), "Grading Standards and Education Quality", *American Economic Journal: Microeconomics*, 7, 2 (May 2015), pp. 248-279.

BOLTON, PATRICK and CHRISTOPHER HARRIS (1999), "Strategic Experimentation", *Econometrica*, 67, 2 (Mar. 1999), pp. 349-374.

BONATTI, ALESSANDRO and JOHANNES HÖRNER (2011), "Collaborating", *American Economic Review*, 101, 2 (Apr. 2011), pp. 632-663.

BOWEN, T. RENEE, GEORGE GEORGIADIS, and NICOLAS S. LAMBERT (2019), "Collective Choice in Dynamic Public Good Provision", *American Economic Journal: Microeconomics*, 11, 1 (Feb. 2019), pp. 243-298.

BRENNAN, GEOFFREY and ALAN HAMLIN (1998), "Expressive Voting and Electoral Equilibrium", *Public Choice*, 95, 1-2, pp. 149-175.

BRENNAN, GEOFFREY and LOREN LOMASKY (1993), *Democracy and Decision: The Pure Theory of Electoral Preference*, Cambridge University Press.

BULL, JESSE and JOEL WATSON (2004), "Evidence Disclosure and Verifiability", *Journal of Economic Theory*, 118, 1 (Sept. 2004), pp. 1-31.

CAMPBELL, ARTHUR, FLORIAN EDERER, and JOHANNES SPINNEWIJN (2014), "Delay and Deadlines: Freeriding and Information Revelation in Partnerships", *American Economic Journal: Microeconomics*, 6, 2 (May 2014), pp. 163-204.

CHAN, JIMMY, SEHER GUPTA, FEI LI, and YUN WANG (2019), "Pivotal Persuasion", *Journal of Economic Theory*, 180 (Mar. 2019), pp. 178-202.

COMPTE, OLIVIER and PHILIPPE JEHIEL (2004), "Gradualism in Bargaining and Contribution Games", *Review of Economic Studies*, 71, 4 (Dec. 2004), pp. 975-1000.

— (2010), "Bargaining and Majority Rules: a Collective Search Perspective", *Journal of Political Economy*, 118, 2 (Apr. 2010), pp. 189-221.

CRAWFORD, VINCENT P. and JOEL SOBEL (1982), "Strategic Information Transmission", *Econometrica*, 50, 6 (Nov. 1982), p. 1431.

DEGAN, ARIANNA and ANTONIO MERLO (2007), "Do Voters Vote Sincerely?", *NBER Working Paper 12922*.

DELLAVIGNA, STEFANO and MATTHEW GENTZKOW (2010), "Persuasion: Empirical Evidence", *Annual Review of Economics*, 2, 1 (Sept. 2010), pp. 643-669.

DELLAVIGNA, STEFANO and ETHAN KAPLAN (2007), "The Fox News Effect: Media Bias and Voting", *The Quarterly Journal of Economics*, 122, 3 (Aug. 2007), pp. 1187-1234.

DOVAL, LAURA (2018), "Whether or Not to Open Pandora's Box", *Journal of Economic Theory*, 175 (May 2018), pp. 127-158.

DOWNS, ANTHONY (1957), "An Economic Theory of Political Action in a Democracy", *Journal of Political Economy*, 65, 2 (Apr. 1957), pp. 135-150.

EGOROV, GEORGY and KONSTANTIN SONIN (2019), "Persuasion on Networks", *SSRN Electronic Journal*.

ENIKOLOPOV, RUBEN, MARIA PETROVA, and EKATERINA ZHURAVSKAYA (2011), "Media and Political Persuasion: Evidence from Russia", *American Economic Review*, 101, 7 (Dec. 2011), pp. 3253-3285.

FELSENTHAL, DAN S. and AVRAHAM BRICHTA (1985), "Sincere and Strategic Voters: An Israeli Study", *Political Behavior*, 7, 4, pp. 311-324.

FERSHTMAN, CHAIM and SHMUEL NITZAN (1991), "Dynamic Voluntary Provision of Public Goods", *European Economic Review*, 35, 5 (July 1991), pp. 1057-1067.

FORGES, FRANÇOISE (1986), "An Approach to Communication Equilibria", *Econometrica*, 54, 6 (Nov. 1986), p. 1375.

GEHLBACH, SCOTT and KONSTANTIN SONIN (2014), "Government Control of the Media", *Journal of Public Economics*, 118 (Oct. 2014), pp. 163-171.

GEORGE, LISA M and JOEL WALDFOGEL (2006), "The New York Times and the Market for Local Newspapers", *American Economic Review*, 96, 1 (Feb. 2006), pp. 435-447.

GEORGIADIS, GEORGE (2015), "Projects and Team Dynamics", *The Review of Economic Studies*, 82, 1 (Jan. 2015), pp. 187-218.

GROSSMAN, SANFORD J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality", *The Journal of Law and Economics*, 24, 3 (Dec. 1981), pp. 461-483.

GUO, YINGNI (2016), "Dynamic Delegation of Experimentation", *American Economic Review*, 106, 8 (Aug. 2016), pp. 1969-2008.

HALAC, MARINA, NAVIN KARTIK, and QINGMIN LIU (2016), "Optimal Contracts for Experimentation", *The Review of Economic Studies*, 83, 3 (July 2016), pp. 1040-1091.

HAMLIN, ALAN and COLIN JENNINGS (2011), "Expressive Political Behaviour: Foundations, Scope and Implications", *British Journal of Political Science*, 41, 3 (July 2011), pp. 645-670.

HARFOUSH, RAHAF (2009), *Yes We Did! An Inside Look at How Social Media Built the Obama Brand*, New Riders, p. 199.

HEESE, CARL and STEPHAN LAUERMANN (2019), "Persuasion and Information Aggregation in Elections", *Working Paper*.

HILLYGUS, D. SUNSHINE and TODD G. SHIELDS (2014), *The Persuadable Voter: Wedge Issues in Presidential Campaigns*, Princeton University Press, p. 267.

HU, LIN, ANQI LI, and ILYA SEGAL (2019), "The Politics of Personalized News Aggregation", *Working Paper* (Oct. 2019).

JANSSEN, MAARTEN C. W. and MARIYA TETERYATNIKOVA (2017), "Mystifying but not Misleading: when does Political Ambiguity not Confuse Voters?", *Public Choice*, 172, 3-4 (Sept. 2017), pp. 501-524.

KAMENICA, EMIR (2019), "Bayesian Persuasion and Information Design", *Annual Review of Economics*, 11, pp. 249-272.

KAMENICA, EMIR and MATTHEW GENTZKOW (2011), "Bayesian Persuasion", *American Economic Review*, 101, 6 (Oct. 2011), pp. 2590-2615.

KAN, KAMHON and C. C. YANG (2001), "On Expressive Voting: Evidence from the 1988 U.S. Presidential Election", *Public Choice*, 108, 3-4, pp. 295-312.

KATZ, JAMES EVERETT., M. BARRIS, and A. JAIN (2013), *The Social Media President : Barack Obama and the Politics of Digital Engagement*, Palgrave Macmillan, p. 215.

KELLER, GODFREY and SVEN RADY (2010), "Strategic Experimentation with Poisson Bandits", *Theoretical Economics*, 5, 2, pp. 275-311.

KELLER, GODFREY, SVEN RADY, and MARTIN CRIPPS (2005), "Strategic Experimentation with Exponential Bandits", *Econometrica*, 73, 1 (Jan. 2005), pp. 39-68.

KESSING, SEBASTIAN G. (2007), "Strategic Complementarity in the Dynamic Private Provision of a Discrete Public Good", *Journal of Public Economic Theory*, 9, 4 (Aug. 2007), pp. 699-710.

KIM, YOUNG MIE, JORDAN HSU, DAVID NEIMAN, COLIN KOU, LEVI BANKSTON, SOO YUN KIM, RICHARD HEINRICH, ROBYN BARAGWANATH, and GARVESH RASKUTTI (2018), "The Stealth Media? Groups and Targets behind Divisive Issue Campaigns on Facebook", *Political Communication*, pp. 1-29.

Kolotilin, Anton (2015), "Experimental Design to Persuade", *Games and Economic Behavior*, 90 (Mar. 2015), pp. 215-226.

Lewis, Tracy R. (2012), "A Theory of Delegated Search for the Best Alternative", *The RAND Journal of Economics*, 43, 3 (Sept. 2012), pp. 391-416.

Lipnowski, Elliot (2020), "Equivalence of Cheap Talk and Bayesian Persuasion in a Finite Continuous Model", *note*.

Lipnowski, Elliot and Doron Ravid (2020), "Cheap Talk With Transparent Motives", *Econometrica*, 88, 4, pp. 1631-1660.

Manso, Gustavo (2011), "Motivating Innovation", *The Journal of Finance*, 66, 5 (Oct. 2011), pp. 1823-1860.

Marx, Leslie M. and Steven Matthews (2000), "Dynamic Voluntary Contribution to a Public Project", *Review of Economic Studies*, 67, 2 (Apr. 2000), pp. 327-358.

McClellan, Andrew (2019), "Experimentation and Approval Mechanisms", *Mimeo*, pp. 1-77.

Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications", *The Bell Journal of Economics*, 12, 2, p. 380.

— (2008), "What the Seller Won't Tell You: Persuasion and Disclosure in Markets", *Journal of Economic Perspectives*, 22, 2 (Mar. 2008), pp. 115-131.

Milgrom, Paul R. and John Roberts (1986), "Relying on the Information of Interested Parties", *The RAND Journal of Economics*, 17, 1, p. 18.

Myerson, Roger B. (1986), "Multistage Games with Communication", *Econometrica*, 54, 2 (Mar. 1986), p. 323.

Oberholzer-Gee, Felix and Joel Waldfogel (2009), "Media Markets and Localism: Does Local News en Español Boost Hispanic Voter Turnout?", *American Economic Review*, 99, 5 (Dec. 2009), pp. 2120-2128.

Olszewski, Wojciech and Richard Weber (2015), "A More General Pandora Rule?", *Journal of Economic Theory*, 160 (Dec. 2015), pp. 429-437.

Ostrovsky, Michael and Michael Schwarz (2010), "Information Disclosure and Unraveling in Matching Markets", *American Economic Journal: Microeconomics*, 2, 2 (May 2010), pp. 34-63.

Perego, Jacopo and Sevgi Yuksel (2018), "Media Competition and Social Disagreement", *Working Paper*.

Prat, Andrea and David Strömberg (2013), "The Political Economy of Mass Media", *Advances in Economics and Econometrics*, Cambridge University Press, pp. 135-187.

Prummer, Anja (2020), "Micro-Targeting and Polarization", *Journal of Public Economics*, 188 (Aug. 2020), p. 104210.

Romanyuk, Gleb and Alex Smolin (2019), "Cream Skimming and Information Design in Matching Markets", *American Economic Journal: Microeconomics*, 11, 2 (May 2019), pp. 250-276.

Schipper, Burkhard C. and Hee Yeul Woo (2019), "Political Awareness and Microtargeting of Voters in Electoral Competition", *Quarterly Journal of Political Science*, 14, 1, pp. 41-88.

STRULOVICI, BRUNO (2010), “Learning While Voting: Determinants of Collective Experimentation”, *Econometrica*, 78, 3, pp. 933-971.

WEITZMAN, MARTIN L. (1979), “Optimal Search for the Best Alternative”, *Econometrica*, 47, 3 (May 1979), p. 641.

WYLIE, CHRISTOPHER (2019), *Mindf*ck: Cambridge Analytica and the Plot to Break America*, Random House, p. 270.