Strategic Signaling Under Higher-order Inference

by

Shih-Tang Su

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical and Computer Engineering)
in The University of Michigan
2021

Doctoral Committee:

Associate Professor Vijay Subramanian
Professor Tilman Borgers
Professor Mingyan Liu
Associate Professor David Miller
Associate Professor Grant Schoenebeck

Shih-Tang Su

shihtang@umich.edu

ORCID iD: 0000-0002-4531-3647

ACKNOWLEDGEMENTS

I would like to express my gratitude to my advisor Professor Vijay Subramanian for his guidance and support during my PhD journey. I am extremely fortunate to have an advisor who gave me the freedom to explore areas of my interests and continuously provided excellent guidance through the years of my Ph.D. research. I cannot express my thank for the incredible amount of time he spent on helping me improve my presentations and writings. I also want to thank all my committee members, Tilman Borgers, Mingyan Liu, David Miller, and Grant Schoenebeck, for their comments that improved the quality of this thesis. I also want to thank Grant Schoenebeck and David Kempe for their guidance and advice during our collaborations. Without their contribution and support, Chapter II, III, V in this thesis wouldn't have been possible. I would like to express my gratitude to Shaowei Ke and Demos Teneketzis for their excellent courses, study groups, and great mentorship.

It is my great pleasure to work with many friends at Michigan, especially Daniel Vial, Dengwang Tang, Hsu Kao, and Mehrdad Moharrami, all in Vijay's group. I learned a lot from each of them in the reading groups, overlapped courses, and a lot of creative discussions in our office.

Finally, I want to thank my parents for their years of endless love and support. Besides, I also want to thank Pusheen the cat, my favorite and the most adorable stuffed animal in the world, for its mental support during the past six years.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

Inferring the information structure of other agents is necessary to derive optimal mechanisms/signaling strategies in games with communications. However, in many real-world problems such as fake news detection, drug development, and lobbying, the (Bayesian) belief updating procedure is not one-shot, and agents' optimal signaling strategies rely on the responses of sequentially updated beliefs. This demands approaches to systematically analyze higher-order inference of agents and derive each agent's optimal manipulation of information revelation to serve her own objective. However, since the information space's size grows exponentially with the order of inference, approaches that serve both purposes can be complicated and hard to analyze. In order to advance the theoretical understanding of higher-order inference in games with communications, this thesis studies several models in social learning and information design and solves particular questions. First, we present an elegant approach to accumulate information under higher-order inference among agents in social learning models, and then explicitly construct an algorithm to achieve asymptotic learning. Second, we study two Bayesian persuasion problems, one with sequentially conducted experiments and the other with the receiver's action constrained exogenously. We propose a dynamic programming algorithm for the optimal commitments in the former problem and prove that the receiver benefits from constraints under binary-state models in the latter problem. These results highlight the fragility of optimal signaling strategies under higher-order inference. Last, we study information design problems where each sender only obtains partial information, and our results highlight the significance of the order of commitment under higher-order inference.

CHAPTER I

Introduction

1.1 Motivation

In strategic decision-making, before taking an action, agents gather information, analyze information, and infer the interactions and outcomes via surmising and exploiting other agents' information. Specifically, in games when agents have private information and communication among agents is possible, an agent not only needs to exploit a direct message received from their communications but also infers other agents' strategies to know the possible private information, i.e., information structures, behind the transmitted message; the message can be direct (e.g., via actions), or implicit (e.g., via observations). Hereafter, we call an agent who only sends messages[1] a sender and an agent who only receives messages a receiver. The study of information structures led to the establishment of two significant fields, namely, *Mechanism Design* exemplified by *Myerson* (1981); *Hurwicz and Reiter* (2006); *Börgers* (2015); and *Information Design* that was developed in *Kamenica and Gentzkow* (2011); *Rayo and Segal* (2010). In mechanism design problems, the designer aims to incentivize information revelation of senders and indirectly designs the information structure by regulating senders' action spaces. When two cornerstone properties, namely incentive compatibility and individual rationality, are satisfied, then receivers(designers) can

[1]The messages need not be these agents' private information.

infer senders' information structures via received signals and update their beliefs. In information design problems, each sender can design an information structure directly by committing to a (randomized) signaling strategy called a signal structure. Given a signal structure, each realized signal under the structure will manipulate the posterior belief of states and hence influence receivers' actions.

When the beliefs are updated only once as in static mechanism design or information design problems, e.g., principle-agent problems, sealed-bid auctions, or classical Bayesian persuasion problems, the process of inference is usually straightforward. Since senders' information is granted by nature but not other agents' signals, the receiver only needs to infer the information structure of senders, but not how the senders' information was generated. However, when the information transmission process is not one-shot, and the beliefs need to be updated iteratively, sequentially inferring the information structure, called **higher-order inference**[2] hereafter, is needed and this can be a challenging task. Agents' decision-making can be dramatically different from static models, even when strategic behaviors in signaling have not yet come to play. One such phenomenon is an information Braess's paradox[3]: in our context it is best seen from an information provider's perspective wherein providing additional source of information (with common suggested action) overturns receiver's decision. The following counter-intuitive example clearly illustrates the lack of monotonicity on agents' decisions even without any strategic behaviors from other participants.

Example 1. *Alice, Bob, and Carol play a game where Alice and Carol have private information about past participants' actions that Bob cannot access. In this game, Bob takes action, either X or Y, trying to maximize his utility. Alice and Carol can*

[2]Whereas such signalling also occurs in teams, it is a lot more complex in games.

[3]To avoid confusion with similar terminologies presented in the literature, we clarify that the information Braess's paradox in our context is different from "Informational Braess' paradox" presented in *Acemoglu et al.* (2018) or "information paradox" presented in *Yao et al.* (2019), but more closely related to the nonmonotonicity of observation errors in cascades presented in *Le et al.* (2014a).

present their private information truthfully to Bob before Bob's decision. The setting is such that following two statements hold:

- *When Alice suggests that Bob take action X and Carol stays silent, Bob's best response is to follow Alice's suggestion and take action X.*
- *When Carol suggests Bob take action X and Alice stays silent, Bob's best response is to follow Carol's suggestion and take action X.*

Given the above statements, when both Alice and Carol suggest that Bob takes action X, will Bob always follow?

*If every past participant's action is solely based on her private signal, the answer is **Yes**. However, if the higher-order inference is required, e.g., some past participants can observe their predecessors' actions and take actions based on their inferences, the answer can be **No**. The reason for this is that Bob could ascertain some of past participants' private information from both Alice's and Carol's messages and overturn his initial inference solely based on either Alice's or Carol's message. Note that this does not demand an overlap between Alice's and Carol's private information. Figure 1.1 depicts a topology of agents' observations that could make Bob takes action Y as both Alice and Carol suggest action X. The formal setup and the reasoning of the example are provided in Appendix A.1.*

Motivated by Example 1, two pivotal questions arise when higher-order inference comes to play:

- Given a topology of information propagation, how is information accumulated among (strategic) agents?
- How do (strategic) agents manipulate the information propagation to serve their own objectives?

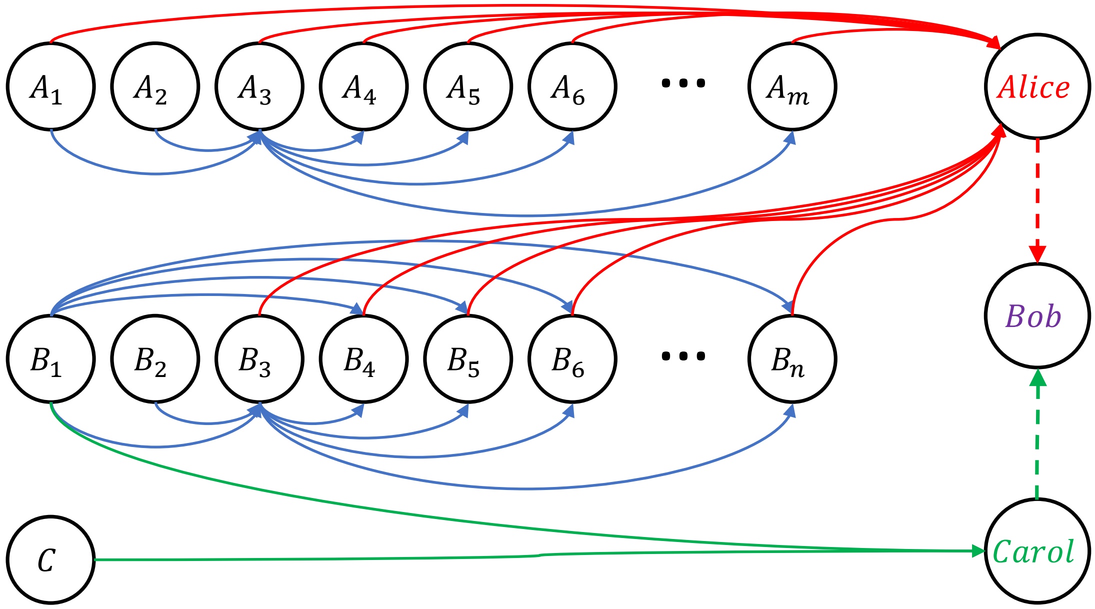

Figure 1.1: Counter-intuitive Example in Higher-order Inference

This thesis aims to advance the theoretical understanding of these two questions, in particular, studying the impact of higher-order inference. In real-world problems, answering the first question demands a systematic approach to analyzing the information accumulated among agents under a given network topology. Otherwise, a correct belief-updating procedure cannot be assured. In other words, higher-order inference has to be sequentially analyzed. To advance the understanding of the first question, we start with a social learning model where (strategic) agents collectively attempt to learn an unknown underlying state by attempting to aggregate their private signals. In such problems, we allow explicit signaling between agents that is constrained by the underlying network topology. In Chapter II, we formulate a systematic approach to analyze higher-order inference and apply the approach to analyze the problem of interest in such problems: whether asymptotic learning is possible in the model. In order to study how agents manipulate their information in an environment under sequential signal revelation, Chapter III considers a Bayesian persuasion problem where experiments are conducted sequentially to shed some light on the second question. To further understand the second question, Chapter V studies how exogenous constraints on the receiver's actions affect the senders' signals and ultimately the receiver's utility. To consider problems where both questions apply, Chapter IV studies a Bayesian persuasion problem where two senders send signals sequentially with a determined order, but where the order of commitment may vary. This studies how the inference (on top of another senders' inference) impacts the sender's signaling strategy, namely commitments. The chapter ordering reflects the flow wherein Chapter IV is a natural follow-up to Chapter III - a sequential revelation problem with two competing senders as opposed to a single sender choosing a single but complex signal structure.

This thesis can also be viewed as an exploration of the consequences of real-world constraints on strategic information disclosure. Different chapters in this thesis discuss different constraints in games with communications. First, chapter II studies constraints on "real" communications, including the network topology that facilitates the communication and limitations on (channel) capacities. Second, chapter III and Chapter V study constraints on the sender and the receiver respectively. The former studies constraints that restrict the sender's signal space, and the latter studies constraints that constrain the receiver's action distribution. Lastly, chapter IV studies endogenously enforced constraints that arise due to the presence of other (strategic) agents. We provide a summary of each chapter and its contribution below.

1.2 Summary of Thesis

1.2.1 Chapter II: Social learning with questions

Chapter II discusses our research efforts in unveiling how information can be accumulated in social learning models *Bikhchandani et al.* (1992); *Smith and Sørensen* (2000) by self-utility maximizing agents with the high-level goal avoiding an information cascade or herding behaviors and enabling (asymptotic) learning. This simpler problem setting can yield conceptual guidance on how information is accumulated without being obscured by either complex strategic inference or numerical calculations. In social learning models without communication among agents, the well-known result of *Bikhchandani et al.* (1992) in combination with *Smith and Sørensen* (2000) states the following: *Without additional source of information in social learning models, keeping the three key assumptions, namely Bayes-rationality of agents, fully observable history, and bounded signal strength, makes asymptotic learning impossible.* To study how the information can be accumulated and how higher-order inference can be systematically derived, we extend the classic BHW model in *Bikhchandani et al.* (1992) and endow agents with the ability to query their predecessors. We sketch the framework of our model below.

1. Binary state space $\theta = \{\theta_1, \theta_2\}$ with prior $p = 0.5$ (on state θ_1).

2. A countable number of Bayes-rational agents, each is taking a once-in-a-lifetime action sequentially from a binary action space with the goal of matching the true state of the world.

3. Each agent observes an informative but not revealing (binary) private signal. Besides, all private signals are conditionally independent given the state and identically distributed being generated as the output of a binary symmetric channel with the state as input.

4. Each agent can observe all the predecessors' actions before taking her action.

5. Each agent can query their immediate predecessor a question up to a predetermined capacity constraint $\bar{K}$.

The first four are assumptions in the BHW model, and the last one is motivated by real-world applications such as information dispersal in social networks such as

Facebook and Twitter. In such networks, an agent can communicate with her immediate predecessor because one must be friends with or follow others before seeing their posts/tweets. Taking agents' communications beyond classical social learning models into account, we determine whether asymptotic learning can be achieved via agents' questions. In other words, we ascertain how information is accumulated via questions and determine how to analyze higher-order inference systematically. To highlight the novelty of this model, we allow each agent to ask questions conditional on her private signal, observed history, and inference of their predecessors' possible information structures. In other words, an agent can ask questions on top of her predecessors' questions if doing this helps her maximize the probability of matching the true state. Therefore, a systematical analysis of higher-order inferences, i.e., the question on top of questions, is mandatory to correctly derive the evolution in beliefs. Otherwise, agents cannot decipher how information is accumulated by their predecessors.

In Chapter II, we present a novel perspective to view questions as means of partitioning an agent's information space. This information partitioning perspective then allows us to model the higher-order inference as transitions of a Markov process, assuming every question can be answered by the corresponding predecessor (based on the information available). Since the information-partitioning perspective and the Markov process formulation enables systematic mathematical analysis of higher-order inferences by transforming the evolution of beliefs into technically simpler matrix multiplication-based analysis, one can then answer the information accumulation among agents via any given question sequence.

After deciphering how the information is accumulated among agents, we then determine whether the queried questions under higher-order inference can help agents learn the true state asymptotically. In essence, if a fictitious social planner exists, can she tailor information structures among agents to preclude herding behavior towards the incorrect state in every realized question sequence? We answer this question by explicitly constructing a class of question sequences that achieves asymptotic learning. Based on our construction, allowing each agent to ask her immediate predecessor one yes-no question is enough to avoid herding behavior in the setting of the BHW model. This (surprising) result indicates how powerful information aggregation is when higher-order inference can be exploited explicitly.

In Section 2.6 and Section 2.7, we generalize our model to a finite signal space and a finite state space, and derive conditions where asymptotic learning can be achieved. Besides, inference under unbounded questions (no capacity constraint) in social learning models and several subtle assumptions of common knowledge and inference are discussed in Section 2.8.1 and Section 2.8.5 respectively.

1.2.2 Chapter III: Bayesian persuasion in sequential trials

Chapter III studies Bayesian persuasion models where signals are sent sequentially instead of in a one-shot manner. This setting of sequentially released signals can be found in many real-world scenarios. For example, in drug development, the drug must pass several trials from phase I to phase III to get the FDA approval. In the ideal case where the sender can design experiments in every phase, this perfectly matches the classical Bayesian persuasion model. However, when some experiments conducted in later phases are conditional on the outcome in early phases and also predetermined[4], the sender is constrained in the choice of signals available for information design; and her optimal signaling strategy can change significantly. Specifically, we are interested in problems that are naturally modeled via multi-phase trials where the interim outcomes determine the subsequent experiments. Further, we insist that some of the experiments are given in an exogenous manner. This framework aims to advance our knowledge on how the higher-order inference and the sequentially updated beliefs impact the sender's optimal signaling strategy. Since the signaling

[4]For example, several mandatory experiments in the drug development cycle may be regulated and determined by the FDA.

scheme in each phase is either pre-determined or designed by the (unique) sender, this framework avoids the challenging task of analyzing the interaction of strategies among agents under higher-order inference and provides insights on the second pivotal question mentioned earlier.

In Chapter III, the sender cannot directly design the posterior belief distribution because she cannot control all the experiments. She can only manipulate some interim beliefs to design the posterior belief distribution indirectly. In other words, designing a (Bayes) plausible distribution of the final experiments becomes the sender's objective. In Section 3.3, we develop a novel measure to evaluate different sets of possible strategies in the final round, which is then used to derive the optimal signaling strategy using a dynamic programming algorithm to design each outcome's beliefs using backward iteration; this is detailed in Section 3.4. Moreover, we (partially) address the robustness of the optimal signaling strategies in a sequential-trial model, which is a small perturbation of multi-phase models that are equivalent (in terms of expected utilities) to the classical 1-phase Bayesian persuasion in Section 3.4.4.1. The result highlights the fragility of the optimal signaling strategies under classical 1-phase Bayesian persuasion and the importance of using a tailored signaling strategy in sequential trials. Some preliminary examples deriving the optimal pre-determined experiments from the receiver's viewpoint are provided in Section 3.5.2, but a comprehensive discussion on this topic is left for future work.

1.2.3 Chapter IV: Importance of the commitment order

Chapter IV studies both the two pivotal questions raised after Example 1 jointly. In Chapter IV, we consider sequential information design problems with two senders, where each sender only observes partial information about the state of the world. For example, a company seeks advice from a lab and a factory respectively for its new product manufacturing, where the lab knows whether applying a patent improves the quality (of the product) and the factory knows which manufacturing process is the most efficient one. To study this problem in a precise manner, we specify that the signals are revealed sequentially in a Bayesian persuasion model under a predetermined order, i.e., sender S_1 always sends signals before sender S_2. Unlike most of the Bayesian persuasion literature with multiple senders, we do not assume that the commitments are made simultaneously or aligned with the order of signaling. Therefore, sender S_2 may commit before sender S_1, although her signals will always be sent after sender S_1's signal, and the problem to study is whether S_2 committing before S_1 leads to a different outcome than if S_1 commits before S_2. Once a difference is identified, a natural follow-up is to study whether there is an order of commitment that is preferable, say to the receiver.

In Chapter IV, the commitment order matters when the high-order inference changes the senders' strategic signaling. In other words, if the commitment order does not matter, the optimal commitments of both senders can be individually inferred by each sender. In such a scenario senders would have reached a consensus on each others' optimal commitments and formed an equilibrium no matter the commitment order, and the purpose of their inference is to verify the equilibrium. To study when the commitment order matters, we start by defining a core assumption on commitments, called permutation-free commitments, in Section 4.2.3 and also explain why it is a proper assumption. Using the setting of permutation-free commitments, in Section 4.3.2 we present two sets of sufficient conditions and a set of necessary conditions to identify when the commitment order matters. The results are based on whether sender S_1 and S_2 have a conflict of interest on some information sets, and whether S_2 can create a credible threat to stop S_1's aggressive signaling strategy when S_2 commits first. Interesting examples illustrating the change of optimal commitments and the gap between two different sets of sufficient conditions are discussed in Section 4.4.

1.2.4 Chapter V: Bayesian persuasion with constrained actions

Chapter V studies a Bayesian persuasion setting in which the receiver is trying to match the (binary) state of the world, but her action distribution is exogenously constrained. We consider a model where the sender's utility is partially aligned with the receiver's utility in a manner such that conditioned on the receiver's action, the sender derives higher utility when the state of the world matches the action. Using the prosecutor-judge problem in *Kamenica and Gentzkow* (2011) as an example, we consider a model where the conviction rate of the judge is constrained, e.g., within 20% range of the prior belief (historical data), and the prosecutor prefers a guilty defendant to be convicted over an innocent defendant, and an innocent defendant to be acquitted over guilty defendant. Under this setting, the sender must consider the receiver's constraints to design her optimal signaling strategy. More specifically, when constraints bind in some actions, the sender must infer the receiver's second-best actions under the signal (she finally commits to). However, based on her inference, the sender may find that some other constraints now bind according to the receiver's second-best action in this signal. Hence, the sender has to consider the signal structure jointly according to the receiver's constraints. In Chapter V, our goal is to find out when the receiver will benefit from the constraints enforced on her action distribution.

Section 5.3 analyzes scenarios when the state of the world is binary and a strong result monotone to the constraints is obtained, i.e., a harsher constraint makes the receiver better. However, Section 5.4.1 provides an example illustrating that the receiver is not always better off when the state size grows beyond binary. Although some additional assumptions are discussed in Section 5.5, conditions that guarantee the receiver benefits from a constrained action distribution are still not well understood in general (finite) state-space models. The exploration of the critical properties/assumptions in finite state-space models is left for future work.

1.3 Summary of Contributions

In summary, the major contributions of this thesis, ordered in terms of chapters, are as follows:

- When agents can communicate with their immediate predecessors, the true state can be asymptotically learned when agents accumulate the information (of compressed private signals) using a sequence of questions. Learning is achieved by viewing the communications among agents as partitions of information spaces. This viewpoint transforms the analysis of higher-order inference to the computation of transitions of a Markov process. It can be exploited in other network topologies besides the one where the immediate predecessors are queried.

- Given the sequence of a set of experiments where some of the experiments are pre-determined in a Bayesian persuasion model, we propose a dynamic programming algorithm that solves the sender's optimal signaling strategy. Moreover, when the experiments have only two phases, we provide an explicit characterization of the optimal signaling strategy using the structure of experiments.

- We present conditions when the commitment order affects the senders' signaling strategies in a two-sender, one-receiver Bayesian persuasion model where signals (from senders) are revealed sequentially. The results are primary based on whether the latter sender can credibly threaten the former sender's aggressive signaling strategies when she commits first.

- We prove that when the receiver's action distribution is constrained in a binary state-space model, the receiver experiences a higher utility when the constraints become harsher. We also provide an example where the receiver is worse off in general state-space models, indicating the difficulty of generalizing the results.

Note: Notation varies across chapters, but each is self-contained, i.e., any notation used in a chapter is defined in that chapter. Appendices use the same notation as for their corresponding chapter, e.g., Appendix B uses the same notation as Chapter II, Appendix C uses the same notation as Chapter III.

CHAPTER II

Social Learning with Questions

2.1 Introduction

The fundamental question in general social learning models *Bikhchandani et al.* (1992); *Welch* (1992); *Banerjee* (1992) is whether the agents can aggregate enough information to eventually ascertain the state of the world, i.e., learn asymptotically. Key results in *Bikhchandani et al.* (1992); *Welch* (1992); *Banerjee* (1992) state that even in the classic social learning model (also called the Bayesian observational learning model), where countably infinite Bayes-rational agents make decisions sequentially to match a binary unknown state of the world, an outcome called *Information Cascade* occurs almost surely under fully observable history and bounded likelihood ratio of signals. An information cascade is a phenomenon where it is individually optimal for agents to ignore their private signals for decision making, after observing the realized history. Since agents' private signals in the cascade will never affect their actions, information aggregation via the realized history halts once the cascade commences. Although continuing the information cascade is individually optimal for every agent (under the realized history), the termination of information aggregation could lead to a socially sub-optimal outcome where all agents, after an information cascade starts, choose a wrong action, referred to as a "wrong cascade;" decisions are made using only a finite set of stochastic private signals, which can easily indicate the wrong state.

Therefore, without an additional source of information, in social learning models, asymptotic learning is impossible while keeping only the three key assumptions:

1. Bayes rationality of agents,

2. Fully observable history, and

3. Bounded signal strength.

To achieve asymptotic learning, models that relax each of the key assumptions discussed above have been studied in the literature *Smith and Sørensen* (2000); *Sørensen* (1996); *Acemoglu et al.* (2011); *Peres et al.* (2017); *Hellman and Cover* (1970); *Drakopoulos et al.* (2013). By relaxing the assumption of either bounded signal strength *Smith and Sørensen* (2000); *Sørensen* (1996) or fully observable history *Acemoglu et al.* (2011), agents can individually learn the state of the world asymptotically since either the private signal or the revealed history grants them unbounded likelihoods using their posterior beliefs. On the other hand, by relaxing the assumption of Bayes rationality of agents *Peres et al.* (2017); *Drakopoulos et al.* (2013); *Hellman and Cover* (1970), the agents can learn collaboratively by assuming a common aspiration of learning of their successors and by implementing action strategies that balance information propagation and matching of the perceived state.

In real social networks, like Facebook or Twitter, agents usually have an additional source of information besides their private signals and history of observations; these two sources of information have already been analyzed in the classical social learning literature. Most posts shown on an agent's screens are shared from their friends whom they can also contact privately. This communication capability, given by a exogenous network topology[1], allows agents to query their friends to guide their own decisions. Since the agent can privately get some information while being trapped in an information cascade, social learning models with this additional communication

[1] A stranger usually cannot share a post on your news feed.

capability could potentially result in asymptotic learning and without weakening the three cornerstone assumptions in classical social learning models.

The idea of exploring additional communication capability in social learning models can be traced back to the following statements in *Gul and Lundholm* (1995): "*... a cascade implies that, in equilibrium, the typical agent i knows less than under the natural assumption. ... the reason is that the discrete choice sets provide an insufficient vocabulary to sustain a fully separating equilibrium, given the incentives of the players. ... Informational cascades can be eliminated by enriching the setting in a way that allows prior agents' information to be transmitted.*" The natural assumption in *Gul and Lundholm* (1995) refers to the setting where the agents can send the entire private-signal history to their successors. Since this history grows with time, this is feasible with additional inter-agent communications only if their communication channels either have unlimited capacity or have their capacity increase without bound with the agent index at an appropriate rate. It is clear that in such a scenario asymptotic learning can be easily achieved because agents can explicitly update their beliefs on the states of nature correctly using the complete private-signal history; from well known results in probability and statistics, this will eventually point to the correct state, and hence, enable learning. However, communication channels in real-world scenarios usually have capacity constraints. For example, when people suspect the credibility of a post and then query their friends that are more knowledgeable on the topic, she and/or her friends may not be willing to spend hours or days to discuss it. This then motivates the study of scenarios where the entire private-signal history cannot be communicated, especially when the agent index is large. A natural question here is whether asymptotic learning can take place with only this partial information transmission capability in place; the technical problem would be to characterize the minimal channel capacity required to achieve asymptotic learning. Another way to state the problem is as follows in the language of *Gul and*

Lundholm (1995): *Can informational cascades be avoided even when the typical agent i knows less than under the natural assumption where the knowledge of an agent is constructed using the classical information structure augmented with bounded capacity communication channels?*

With such communication constraints in place, it is natural that only the most valuable piece of available information gets transmitted. Since the value is assessed by the receiver based on her (expected) utility calculations and using her private signals, it is logical to let the receiver pose the question. In our work, the agents are myopic and Bayes-rational, so they will assess the value of the information by computing the posterior belief on the unknown state. Note that the agents will determine the domain of the knowledge of past agents using the fact that they're also myopic and Bayes-rational and that they also would have sought out specific information from their predecessors. Hence, in this work we allow agents to seek information by querying a set of predecessors determined by a given interaction topology, such that different questions (with fixed finite capacity for each response) can be asked conditional on their realized private signals in order to maximize their expected utility.

Starting from the seminal *Bikhchandani et al.* (1992); *Welch* (1992); *Banerjee* (1992) models (specifically the model of *Bikhchandani et al.* (1992) which we label as the BHW model) without modifying any of the three key assumptions, where the Maximum Aposteriori Probability (MAP) rule is still individually optimal and will be used by each agent for her decision, this work allows myopic Bayes-rational agents to query a (small and not growing with agent index, i.e., bounded) finite set of their predecessors with private questions after inferring predecessors' questions and (fully) observed history. Then we seek answers to the following three questions:

1. *What information can be retrieved by a specific agent under a given network topology? In other words, what is the information space for an agent?*

2. *If it is possible to achieve asymptotic learning with bounded capacity communi-*

cation channels, what is the minimum number of predecessors and the minimum size of questions (in terms of bits) required?

3. *Can such bounded capacity inter-agent communication capability achieve asymptotic learning beyond the BHW model and generalize to non-binary state or signal spaces?*

Main Contributions: Three main contributions are made in this chapter:

1. To the best of our knowledge, we are the first to study the endowed communication capability in social learning problems without enforcing constraints on transmittable information in networks. The presented framework unveils what the set of feasible questions[2] is and how an unlimited order of inference can be modeled and analyzed in social learning problems. The perspective of viewing questions and inference in the framework, via taking away exogenous assumptions on types of information granted in networks, can be applied to a variety of models in social learning problems.

2. To the best of our knowledge, we are the first to highlight the ability to sequentially tailor information structures among agents to achieve learning in social learning problems. The approach used in this work, namely partitioning information sets, is closely related to Bayesian persuasion *Kamenica and Gentzkow* (2011); *Rayo and Segal* (2010). Designs achieving asymptotic learning in this chapter can be viewed as a "relayed Bayesian persuasion" by persuading agents in "possibly wrong cascades" to avoid the information cascades eventually. Previewing the results, one set of agents work to assimilate, that is, relay, the information so as to allow another set of agents the possibility of stopping an ongoing cascade.

[2] Information that is possible to be retrieved and applied to update beliefs using Bayesian updates.

3. With an explicit construction of questions corresponding to the agent's possible information set and index value in the network and where agents are only allowed to query their immediate predecessors, we show that learning is achievable in 1-bit communication capacity under the BHW model. Furthermore, we also show that this scheme can be used to achieve asymptotic learning under finite-bit questions with the same network topology for finite state-space, finite signal-space social learning models but restricting to a uniform prior on underlying states of the world. Formally, in the bulk of the paper we establish that our learning outcome emerges as Bayes Nash equilibrium of the underlying dynamic game (i.e., equilibrium behavior), but in the discussion we show how learning can also be understood as a Perfect Bayesian Equilibrium (PBE) *Mas-Colell et al.* (1995); *Fudenberg and Tirole* (1991); *Osborne and Rubinstein* (1994) with sequential rationality with beliefs appropriately defined in off-equilibrium paths.

Note that in our approach, the system designer commits to a specific information structure (full public history of previous agents' actions plus private communications) without any reductions or distortions, and also provides the agents with a question guidebook, whose performance each agent can verify independently. The minimal nature of our learning achieving question guidebook also reveals the fragility of information cascades (Sec.16 in *Easley and Kleinberg* (2012)), as a small amount of strategically delivered information leads to learning; in our case though, the relaying across multiple agents is critical. The information revelation in our question guidebook is strategic in contrast to reviews in *Le et al.* (2016) that are generated via an exogenous process (and revealed only for some specific actions), and furthermore, lead to learning *Acemoglu et al.* (2017) only if the signals are sufficiently informative. A subtle but rather interesting point of our approach is the "relayed persuasion" aspect wherein we only aim to persuade particular agents chosen by our design instead of agents chosen by nature as is commonly seen in Bayesian persuasion *Hedlund* (2017);

Ely and Szydlowski (2017).

Before we proceed, we remark that generalizing results/learning in social learning models from binary state space to finite state space can be traced back to Sørensen's thesis *Sørensen* (1996). In *Sørensen* (1996), Sørensen stated the martingale properties in models *Smith and Sørensen* (2000); *Sørensen* (1996) can be generalized to arbitrary but finite states of the world and a finite set of actions, although the definition of learning may need to be revised by "adequate learning" presented in *Aghion et al.* (1991) when optimal actions on some states are not unique. Because generalizing the martingale analysis from the partition of $[0, 1]$ to a unit simplex in $\mathbb{R}^{N-1}$ makes the optimal decision rules notationally cumbersome without developing results substantially different from the binary social learning models where the scheme of information accumulation is not the main theme, Sørensen's generalization on state space did not draw much attention in the literature. However, when communication among agents are allowed, learning/cascading results can be significantly different under different information accumulation schemes because of the subtlety of simplex partitioning; and this highlights contributions in *Sørensen* (1996).

This chapter is organized as follows. We start with a discussion of related work in Section 2.2. In Section 2.3 we formally define the problem and describe the objective and solution concept. Here we also discuss the endowed communication graph and some basic assumptions on the communications. Then, in Section 2.4 we describe how to associate questions asked of neighbors with information set partitioning and also the feasibility of the questions asked. We also discuss the implications of Bayes-rationality of the agents in terms of sequential rationality. Section 2.5 describes the 1-bit questions based scheme where asymptotic learning obtains as a Bayes Nash equilibrium for the classical BHW model that considers binary state of the world and a binary signal space. Thereafter, in Section 2.6 we generalize our results to finite signal space setting but still with binary state of the world. We present the generalization to asymptotic learning for finite non-binary states of the world in Section 2.7. In Section 2.8 we discuss several subtleties of our result including how asymptotic learning can also be understood as a weak PBE of the underlying game.

2.2 Related Work

Social learning, or so called Bayesian observation learning, studies whether and how consensus (unanimous actions) can be reached among sequential Bayes-rational decision makers under incomplete information. The key result shown in *Bikhchandani et al.* (1992); *Welch* (1992); *Banerjee* (1992) says that with homogeneous and Bayes-rational agents receiving binary private signals, an ***Information cascade*** happens almost surely, with a positive probability for all but a few of the first agents to cascade to the less profitable action. Once an information cascade occurs, no future private signals are revealed. Smith and Sørensen significantly generalized the model to allow for richer signals characterized by the likelihood ratio of the two states of the world deduced from the private signals, and heterogeneous agent types[3] in *Smith and Sørensen* (2000). Using both martingale techniques and Markovian analysis, they proved that any cascade can be stopped with unbounded likelihood ratios in the private signals, hence learning toward the correct action can be achieved; Within this setting, works about the speed of learning are characterized in *Hann-Caruthers et al.* (2018); *Acemoglu et al.* (2017).

These initial works lead to considerable follow-up work towards understanding information cascades better, and towards achieving (asymptotic) learning. The vast majority of work here studies how modifying the *information structure* of the problem impacts cascades or allows for learning. As these are closest to our work, we will discuss these in detail to highlight our contributions and the differences. Aside

[3]Here a new phenomenon called "confounded learning" is demonstrated where in the long run agents of different types will herd on the same action and from the actions it will be impossible to detect the types.

from those works in the above field, there is a vast literature, e.g., *Cover* (1969); *Hellman and Cover* (1970); *Zhang et al.* (2013); *Wang and Djurić* (2015); *Le et al.* (2017b), that consider non-Bayesian agents including bounded rational players, irrational players, and algorithmic agents, and alternate history update rules as a means of achieving learning. A majority of this set of literature studies the (optimal) decision rules in decentralized (binary) hypothesis testing problem on a variety of network models *Tay et al.* (2009); *Drakopoulos et al.* (2016); *Tay et al.* (2008). Since the Bayes-rationality constraint differentiates this chapter from the work on decentralized hypothesis testing problems, we will only discuss several seminal works and highlight one work *Drakopoulos et al.* (2013), in particular, that achieves asymptotic learning with a specific set of four-state Markov chains. The approach in *Drakopoulos et al.* (2013), from the perspective of information design, is similar in spirit to partitioning information state to information sets, but for non-Bayesian agents. Last, there is also literature that allows for heterogeneous types of agents or changes the *actions and the payoff structure*, and studies the impact of these on social learning. In the literature with heterogeneous agent types *Wu* (2015); *Vaccari et al.* (2016); *Le et al.* (2017b), disagreement between agents typically leads to an information cascade but the presence of poorly informed agents surprisingly reduces the probability of a wrong cascade. Papers *Lee* (1993); *Gul and Lundholm* (1995); *Vives* (1997); *Huck and Oechssler* (1998) that modify the actions, typically consider continuum action spaces that result in learning. We will not discuss the last class of literature in detail but include the references for completeness.

As every agent is endowed with a (conditionally) independent and informative private signal about the state of the world, learning would result if these signals are communicated and collected frequently enough and in an accurate fashion. The emergence of information cascades (and herding) shows that the information assimilation part is faulty. Changing the underlying information structure either by revealing only part of the information in the database, by modifying the information in the common database or by adding new channels of information have then been the approaches that have been taken. In *Acemoglu et al.* (2018), only a (random) subset of the past agents' actions are revealed to the current agent. This feature allows some agents to take actions solely based[4] on their private signal, and the paper characterizes the properties of these subsets (called networks) such that learning results. A key result is that even with private signals having a bounded likelihood ratio, there exist networks such that learning occurs. The authors in *Mossel et al.* (2015) show that a special class of simple networks also leads to learning, where agents only see the actions of at most d past agents, and where any agent's action is only observed by agents at most L indices ahead. Recognizing that displaying no past action history would also result in Bayes-rational agents revealing their private signals, *Peres et al.* (2017) determined the minimum sequence of revelation agents that are needed for learning: each agent n reveals independently with probability p_n where we need $p_n \propto 1/n$. In *Banerjee and Fudenberg* (2004); *Smith and Sørensen* (2008), the ordering information of the subset of agents whose actions are revealed, is omitted but nevertheless learning results. For non-myopic (Bayes-rational) agents who can postpone their decisions and tolerate discounting of rewards, *Bistritz et al.* (2019) presents an equilibrium where learning can be achieved as the discount factor goes to one (infinitely patient agents). Examples of imperfect observation of history that do not lead to complete learning often feature assumptions such as deterministic sub-sampling of past agents' actions (sometimes only by a finite number of agents), unknown observation order and aggregating observations, such as in *Sgroi* (2002); *Celen and Kariv* (2004); *Callander and Hörner* (2009); *Ho et al.* (2014); *Song* (2015); *Bohren and Hauser* (2017); *Chamley and Gale* (1994); *Zhang* (2009); *Sgroi* (2003).

In *Le et al.* (2014b,a); *Monzón* (2017), imperfections are added when the informa-

[4]By Bayes-rationality, this is equivalent to the revelation of their private signal.

tion is stored in the database. Whereas this does not result in learning, discontinuous and non-monotonic behavior in the amount of imperfection added is shown; this adds to the literature on information Braess's paradoxes observed with equilibrium behavior. Along the same lines, by allowing for stochastic arrivals, in *Le et al.* (2018), the uncertainty in whether an agent arrived and didn't purchase or no agent arrived results in discontinuous and non-monotonic behavior of the wrong cascade probability, even though learning doesn't result. Finally, the impact of additional information via reviews of the item obtained when purchases are made, is studied in *Le et al.* (2016, 2017a); *Acemoglu et al.* (2017), with the main conclusions that learning requires unbounded likelihood ratios of the private signals (but not the reviews) and information Braess's paradoxes result in extremely non-intuitive behaviors.

By relaxing the assumption that agents are Bayes-rational, learning, also called optimal decision rule, has been studied in distributed hypothesis testing problems with long history. A seminal work *Cover* (1969) studied distributed binary hypothesis testing problem and achieve learning (limit probability of error zero in their language) with agents using a four-valued statistics-based algorithm. A follow-up work in *Hellman and Cover* (1970) showed that asymptotic learning is not achievable under bounded likelihood ratio of signals. General results for distributed hypothesis testing, without a specified class of network topology, were summarized/presented in *Tsitsiklis* (1989). Learning in tree networks with bounded depth/degree with algorithmic agents were respectively studied in *Tay et al.* (2009); *Drakopoulos et al.* (2016). Restricting attentions to a line-structure network, named the tandem network, although there is no learning when the signal distributions are unknown a priori *Ho et al.* (2014), learning can be achieved at the sub-exponential rate for unbounded likelihood ratio in *Tay et al.* (2008). In order to study what additional information is required to achieve learning under bounded likelihood ratio, *Drakopoulos et al.* (2013) allowed agents to see $K \geq 2$ immediate predecessors actions, instead of just one. It is then shown in *Drakopoulos et al.* (2013) that asymptotic learning can be achieved using a specific set of four-state Markov chains, for the case of $K = 2$. From the perspective of information design, this approach of designing Markov chains for learning is similar in spirit to a partitioning of information sets, but for non-Bayesian agents.

Although designing Markov chains for learning, a prototype of information set partition, is provided in *Drakopoulos et al.* (2013), conditions to achieve learning in models with Bayesian agents are usually harsher than with non-Bayesian agents. A well-known example is that even though neither Bayesian nor non-Bayesian learning procedures can achieve asymptotic learning in bounded-likelihood ratio signal with binary actions in *Cover* (1969), learning under bounded-likelihood ratio signal can be achieved in ternary actions *Koplowitz* (1975) among non-Bayesian agents, but not among Bayesian agents *Dia* (2009). The Bayesian-agent assumption restricts some particular (realized by nature) agents' action space and may hence blocks the aggregation of information.

To study how to design information structure shown to Bayes-rational agents, there is an emerging field called "Information design" following the seminal paper "Bayesian Persuasion" in *Kamenica and Gentzkow* (2011). With commitment power on the information structure (without changing the prior beliefs of agents), designed randomized signals can partition the information space and change agents' posterior beliefs to change agents' preferred actions with some probability. However, the nature of our problem, specifically the finite communication capacity, restricts the strength (richness) of the additional signal available, and as consequence it is not possible that every agent's action can stop a cascade with positive probability. To overcome this and to achieve learning, only a particular set of agents, called active agents in Definition 14, have a positive probability to stop the cascade. The purpose of the rest of agents, called silent agents in Definition 14, is to relay information to make active agents be "persuadable" (have a positive probability to stop the cascade after updating their posterior beliefs). From this perspective, the approach presented in this chapter can be viewed as a “relayed Bayesian persuasion”. Essentially, a series of silent agents convince the next active agent that she is in a wrong cascade by making the likelihood ratio effectively unbounded; this is then similar to the overturning principle that holds with an unbounded likelihood ratio *Smith and Sørensen* (2000).

2.3 Problem Formulation

2.3.1 State space, signal space, and information cascade

We start with a model that considers binary states of the world $\Theta = \{\theta_1, \theta_2\}$ where both states $\theta_i \in \Theta$ are equally likely to occur[5]. We then assume that there are a countable number of (myopic) Bayes-rational agents, each taking a single action sequentially and indexed by $t \in \{1, 2, ...\}$. At each time slot t, agent t shows up and chooses an action a_t from the action space $A = \{\bar{a}_1, \bar{a}_2\}$ with the goal of matching the true state of the world.

Formally, for every agent t, the utility function $U_t(\theta, a)$ is defined as the following:

$$U_t(\theta_i, a_t) = \begin{cases} 1, & i = t, \\ 0, & \text{otherwise.} \end{cases} \tag{2.1}$$

Hence, the objective of agent t with information I_t gathered before taking her decision, is to solve the following problem:

$$\max_{a_t \in A} \sum_{i \in \{1,2\}} \mathbb{P}(\theta_i | I_t) \mathbf{1}_{\{a_t = \bar{a}_i\}}. \tag{2.2}$$

Therefore, in the typical social learning setting, before an agent takes her action, she uses Bayes rule to assimilate two different information sources, her **private signal**

[5]The uniform prior assumption is just for simplicity of the analysis in models with binary states as the result holds more generally.

and her observed **history**. The private signal s_t of agent t is informative but not revealing and originates from a finite signal space $\mathcal{S}_t$[6]. The distribution of every signal $\Delta(\mathcal{S}_t)$ is common knowledge to every agent with index $w \geq t$. We assume all signals are conditionally independent and identically distributed, and for simplicity of analysis, $\mathcal{S}_t \equiv \mathcal{S}$ for $t \in \{1, 2, \dots\}$ but the core methodology of partitioning information sets we propose will not rely on this assumption. Besides the private signal, the full history of actions taken by agent t's predecessors $H_t \in A^{t-1}$ is perfectly revealed to agent t. With the information coming from both the private signal and the action history, the agent then computes the interim belief of the true states of the world and generates the questions that will maximize her expected payoff under the assumptions and constraints specified in Section 2.3.2. When this agent receives responses to her questions, she computes the resulting posterior belief and takes the action corresponding to the most likely state. As in *Monzón* (2017); *Le et al.* (2017a), if indifferent between the two actions, we assume that the agent follows her private signal, instead of randomizing, following the majority, or following any other such alternative action. Note that this assumption is only for the simplicity of analysis, and the results of this chapter will hold when other tie-breaking rules are enforced.

As mentioned in the previous section, a phenomenon called (permanent) information cascade may occur in social learning models. An information cascade occurs when an agent ignores her private signal when taking her decision, and a permanent information cascade is defined below:

Definition 1. *In a model of Bayesian observational learning, a* ***permanent information cascade*** *occurs if there exists a finite index T such that every agent t with $t \geq T$ will ignore her private signal while taking her action a_t. In other words, the action a_t will be the same for every choice of private signal $s_t \in \mathcal{S}$.*

[6]When all signals are generated conditionally independently given the state of the world and from an identical distribution with $\mathcal{S}_t \equiv \mathcal{S}$ for all $t \in \{1, 2, \dots\}$, and $|\mathcal{S}| = 2$, the model reduces to the seminal BHW model *Bikhchandani et al.* (1992).

When a permanent information cascade sets in, all agents imitate their immediate predecessor's action. If a permanent information cascade occurs in a social learning model, agents after some finite index T may fall in a "wrong" cascade such that every agent guesses on the wrong state of the world, unless all permanent information cascades result in the choice of the correct action. However, learning is essentially a "good" cascade as it necessarily results in agents ignoring their private signals and taking the correct action, but this behavior could only occur asymptotically. To make these notions precise, we start by defining asymptotic learning in Definition 2 below.

Definition 2. *In a model of Bayesian observational learning,* ***asymptotic learning*** *(in probability) is achieved if* $\lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_i | \theta = \theta_i) = 1$ *for all* $\theta_i \in \Theta$*. If, instead, we have* $\lim_{t\to\infty} U_t(\theta, a_t) = 1$ a.s.*, then the learning occurs almost surely.*

As mentioned earlier, in the bulk of this chapter the solution concept will be that of a Bayes Nash equilibrium, but at the end of this chapter we will also consider refinements by presenting a weak PBE for the underlying dynamic game.

Before we proceed to Section 2.4 and introduce questions/responses of agents that update agents' interim beliefs to posterior beliefs, we briefly discuss the subtlety of permanent information cascades and asymptotic learning in Bayesian observational learning models. Given a specific model, if there are permanent information cascades wherein a wrong action is taken with positive probability, then asymptotic learning is not achieved. However, a model in which permanent information cascades occurs with probability 0 may not result in asymptotic learning either. One straightforward example is a model where there is no observable history and no private communication is allowed for every agent. In such a model, every agent's action is solely based on her private signal so a permanent information cascade never occurs, but no learning occurs in this model either. Hence, given a specific model, asymptotic learning requires that both the following two conditions hold: (1) every (possible) permanent information cascade results in the correct choice (of action); and (2) agents are able to learn the true state asymptotically when permanent information cascades do not occur.

2.3.2 Deterministic network topology with finite channel capacity

In order to achieve asymptotic learning, we assume agents are endowed with communication capabilities. This assumption is motivated by people consulting with their friends before making decisions in social networks, e.g., Facebook, Twitter or Instagram. The general network topology is modeled below:

1. Each agent t has the ability to query a finite[7] set of her predecessors $\mathcal{B}_t$. where $|\mathcal{B}_t| \leq \bar{B}$ and $\bar{B} < \infty$.

2. To each predecessor n that agent t is going to query, a pre-determined $K_t \leq \overline{K}$ capacity in bits is set, where $\overline{K}$ is finite[8].

We assume that the information of $|\mathcal{B}_t|$ and K_t is common knowledge, that is, agent t can exploit the information $\{(|\mathcal{B}_n|, K_n)|n < t\}$ to decide whom she is going to query and what questions are best for her.

Once we know the query set $\mathcal{B}_t = \{n|(n,t) \in E\} \subseteq [t-1]$ for every agent t, the communication capability in this framework can be modeled by a directed acyclic[9] graph (DAG) $G(\mathbb{N}, E)$. We first assume these channels are perfectly reliable and will discuss noisy channels and the corresponding probabilistic-question models in Section 2.8.2. Given a DAG $\mathcal{G}$, agent t in the model can ask questions individually and privately to predecessors in the set $\mathcal{B}_t$ in line with the network topology before making her decision. We remark that agent t asks the set of questions **after** receiving her private signal so she can use her private information to choose the appropriate

[7]The set size can change with the index but is bounded.

[8]For readers curious about achieving asymptotic learning with fewer "total questions in expectation" by relaxing the finite-capacity assumption as agent's index goes to infinity, which enables agents to use statistical aggregation and concentration based results, a scheme is discussed in Section 2.8.1.

[9]While agent n is going to take her action, agent $t > n$ is not yet present so it is not possible for agent n to query agent t to get information benefiting her decision. Thus, the topology should be acyclic.

questions that maximize her expected payoff; issues of off-equilibrium behavior here will be discussed later on.

To highlight the power of the communication channel and to avoid cumbersome analysis of agents' "guess" on the network topology, we make the following two assumptions throughout this chapter.

Assumption 1. *The adopted DAG is common knowledge.*

The common knowledge assumption on DAG allows agents to completely infer their predecessors' questions (generally, they still cannot infer the answers) and then build questions on top of those questions. This assumption also makes the DAG deterministic for every agent. The necessity of studying the deterministic network topology is because of the lack of **information monotonicity**. Information monotonicity *Acemoglu et al.* (2011) is the core property that facilitates almost all asymptotic learning results in randomized topologies: when communication capability is allowed, then the information monotonicty seen in randomized network topologies makes it analytically simpler to demonstrate asymptotic learning. This is not sufficient for the present work owing to the difficulty of identifying or isolating the reasons for learning. In particular, with randomized network topologies, it is hard to distinguish whether the learning resulted from the proposed communication capability or the randomness in the network topology. When we transition to deterministic network topologies though, information monotonicity is not guaranteed. As a matter of fact, in Section 2.8.3.1 we present via an explicit deterministic network topology, an example where information monotonicty does not hold while staying within the framework of *Acemoglu et al.* (2011). Hence, a deterministic network topology allows us to investigate the power of learning solely from the endowed communication capability. Therefore, Assumption 1 is made to isolate the importance of communication capability in social learning problems.

Assumption 2. *There is no cost in asking questions up to the capacity constraint (K_n bits) in the network, and these questions will be truthfully answered.*

This assumption can be justified by viewing the the communication capability up to the capacity constraint as an endowment of the socially-networked agents. The (set of) predecessors, which each agent can query, are his/her friends. In addition, each agent's utility solely depends on his/her action, i.e., depends neither on their successors' actions nor the response to their successors' questions. In other words, there is no incentive for an agent to answer strategically against his/her successors (friends). This assumption essentially implies that all agents are benign, and that they believe that their friends are benign as well.

2.3.3 Line topology and perspectives on generalized DAGs

To simplify the analysis and to illustration clearly how information can be accumulated via analyzing predecessors' questions (and their possible information sets), we will focus on a line network topology in Sections 2.5-2.7. Before we proceed to the details on how information can be conveyed and accumulated via questions, we briefly discuss two different perspectives on the "general" class of DAGs to be considered. However, investigating these directions is for future work.

The first perspective is mathematical where DAGs are considered as exogenous parameters corresponding to the underlying network structures that the agents are endowed with. In this context, either the class of all DAGs should be considered or a subclass that better fits the (real world) scenario. The latter may also arise owing to platforms like Facebook, Twitter or LinkedIn designing, picking or suggesting the topology to the agents, which can be interpreted as algorithmic censorship. The second perspective is economic where we endow each agent n with the parameters $(|\mathcal{B}_n|, K_n)$, and then agent chooses the set of agents (with lower index values) to query such that the final choice is common knowledge to all agents. This perspective imposes incentive compatibility on DAGs and will (inevitably) significantly narrow the class of DAGs. Further, to our best knowledge, there is no systematic approach beyond forward verification[10] on searching and validating the set of incentive-compatible network topologies with given $(|\mathcal{B}_n|, K_n)$ for all n.

2.4 Information Conveyed through Questions

In this section, a pivotal concept will be introduced to help analyze how information can be accumulated through a sequence of questions — viewing questions as partitions of the information space. To detail this concept, we will first illustrate what class of questions are feasible to be asked (and properly answered). Then, we will introduce the process of precluding impossible history realizations (of private signal sequences) by inferring predecessors' questions, which makes information accumulation helpful to agents.

To avoid ambiguity, we carefully define the "feasibility of questions" in our work, starting with a high-level intuition. For each agent, the primary goal of asking questions is to maximize the probability of matching the state of the world. Thus, if a question is responded to with answers under some unknown distribution where the posterior beliefs cannot be explicitly calculated using the Bayes rule, an agent cannot attain information about the state of the world (via the response of this question). Consequently, asking this type of question will be useless for an agent. With this intuition in mind, this work focuses on questions that can help agents update their beliefs, called feasible questions hereafter. We start by providing a simple definition of feasible questions below.

Definition 3. *A question is feasible if it enables a well-defined Bayes update of the querying agent's belief on the state of the world.*

[10] This holds even though the set of incentive-compatible DAGs can still be recursively defined for every $(|\mathcal{B}_n|, K_n)$ given $(\mathcal{B}_i, K_i)$ $\forall i < n$, i.e. it can be defined via first-order logic.

Feasible questions defined above demand that the conditional distribution of the response be known to the agent asking these questions. However, since agents can infer their predecessors' questions and some of their predecessors' private signals according to observed actions, it is not immediate whether a question is feasible for an agent. This section will describe systematically an agent's inference process by viewing questions as tools to partition information sets. To help understand the feasibility of questions more easily, we will start with a special set of feasible questions, called deterministic questions. The response of deterministic questions has no randomness conditioned on the private signal of queried agents, the information inferred by the queried agents, and the realized history observed by the querying agent. After that, we will expand the concept to probabilistic questions to completely describe the class of feasible questions.

When we can operationally determine the set of feasible questions, the rest of this section will introduce sequences of designed questions, named question guidebooks. Question guidebooks allow the agents to accumulate information through their predecessors' questions (already asked), even though an agent does not observe the answer to their predecessors' questions. A design of questions/question guidebooks that achieves asymptotic learning in the BHW model will be presented in the next section.

As mentioned earlier, we commence our discussion with deterministic questions. With such a restriction, questions can be defined as a function of an agent's private signal and the information she can infer from the history (of actions), network topology, and her guesswork on predecessors' questions. In order to understand the information an agent can potentially exploit to ask her questions, we need to ascertain the information space and information sets for an agent.

Before revealing the history of agent t, H_t, network topology G and the set of previous questions asked by agent n's predecessors $\mathcal{Q}_t$, the pre-information space contains all possible private-signal sequences, i.e., $(s_i)_{i=1}^t, s_i \in \mathcal{S}$ forms the pre-information space of agent t. After each source of information mentioned above is revealed or inferred, a subset of private-signal sequences will be removed from the pre-information space according to the Bayesian rationality of agent t's predecessors. By considering the rationality of predecessors, the remaining set of private-signal sequences forms the information space to the agent before asking her questions, denoted by $I_t(H_t, G, \mathcal{Q}_t)$. The following example shows that the information space can be much smaller than the pre-information space.

Example 2. *Consider the classic BHW model, i.e.,* $\Theta = \{\theta_1, \theta_2\}$ *and signals* $\mathcal{S} = \{s^1, s^2\}$ *generated via a binary symmetric channel(BSC) with crossover probability* $1-p$*. The pre-information space of agent #4 contains any possible length-3 private signal sequences. When the revealed history at agent #4 is* $H_4 = \bar{a}_1\bar{a}_1\bar{a}_1$ *and no communication channel is endowed for any agent, i.e.,* $G(\mathbb{N}, E = \emptyset)$*, some private signal sequences should not belongs to the information space of agent #4. Specifically, information space* $I_4(\bar{a}_1\bar{a}_1\bar{a}_1, G(\mathbb{N}, \emptyset), \emptyset)$ *contains only two private signal sequences,* $\{s^1s^1s^1, s^1s^1s^2\}$*. Including any other private sequences in the information space of #4 will contradict the history realization assuming Bayes-rationality of the agents.*

Given the information space, a natural approach to define information sets is to let every private signal sequence in the space be an information set. However, given the constraint on the capacity or the network topology, some private signal sequences will not be distinguishable when agent t arrives. The following example extended from Example 2 will demonstrate this, and thus suggests that an amendment to the definition of elements in the information space is mandated.

Example 3. *Consider a model extended from Example 2 which only considered the first 4 agents. Assume that the revealed history up to agent #6 is* $H_6 = \{\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\}$*. Now we have* $\mathcal{B}_5 = \{3, 4\}$*,* $\mathcal{B}_6 = \{5\}$*, and* $\mathcal{B}_7 = \{6\}$ *with capacity constraint* $K = 1$ *for each channel.*

First, since agent #5 can query private signals of agent #3 and #4, the private signal sequence $\{s^1 s^1 s^2 s^2 s^2\}$ *will not be in the information space of agent #6. Otherwise, agent #5 would violate Bayes-rationality given the realized history* H_6*.*

From agent #6's perspective, she wishes to ascertain whether there are 2 s^2*s received from agent #3-#5 since in such a scenario, when she receives* s^2*, she can take an action that follows her private signal. Due to the capacity constraint and the choice of agent #6's question, agent #6 cannot distinguish between the following private signal sequences* $\{s^1s^1s^1s^2s^2, s^1s^1s^2s^1s^2, s^1s^1s^2s^2s^1\}$ *when she gets the response from agent #5 telling her there are 2* s^2*s received from agent #3-#5. Thereafter, assume that agent #7 arrives with the history being* $H_7 = \{\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\bar{\mathbf{a}}_\mathbf{2}\}$*, then the information space of agent #7 is* $\{s^1s^1s^1s^2s^2s^2, s^1s^1s^2s^1s^2s^2, s^1s^1s^2s^2s^1s^2\}$*. At this time, there is no question*[11] *that agent #7 can ask of #6 to distinguish the exact signal sequence within the set* $\{s^1s^1s^1s^2s^2s^2, s^1s^1s^2s^1s^2s^2, s^1s^1s^2s^2s^1s^2\}$*. This example suggests that we need to carefully define the (distinguishable) elements in the information space.*

Readers may wonder whether the above example is just a special case where we are unlucky and how often we will have private signal sequences that are not distinguishable given a finite capacity constraint. The following lemma states that non-distinguishable private signal sequences are inevitable when the agent index grows.

Lemma 1. *Given a finite signal space and a bounded capacity constraint on communication capability, let* A_n *be the indicator function that a non-distinguishable private signal sequence occurs up to agent* n*, then* $\lim_{n\to\infty} \mathbb{P}(A_n) = 1$*.*

Intuitively, we should define the element of information space, the information sets, by viewing each set of non-distinguishable private signal sequences to be an element of the information space, and each private signal sequence should only belong to one

[11] This is true even by removing the capacity constraint K for agent #7.

information set. Hence, we use the following definition of pooling information sets for deterministic questions.

Definition 4. *A pooling information set I_p is a (non-empty) set of private signal sequences satisfying the following two conditions:*

- *No information can be delivered (via responses) to partition I_p under the current topology.*
- *For any set of private-signal sequences $\hat{I}$ in an information space disjoint from I_p, there exists at least one set of questions respecting the capacity constraints that can partition $I_p \cup \hat{I}$ to $\{\{I_p\}, \{\hat{I}\}\}$.*

The following claim indicates that the refined definition of information sets suffices and does not miss any valid private signal sequence.

Claim 1. *The set of pooling information sets partitions the information space.*

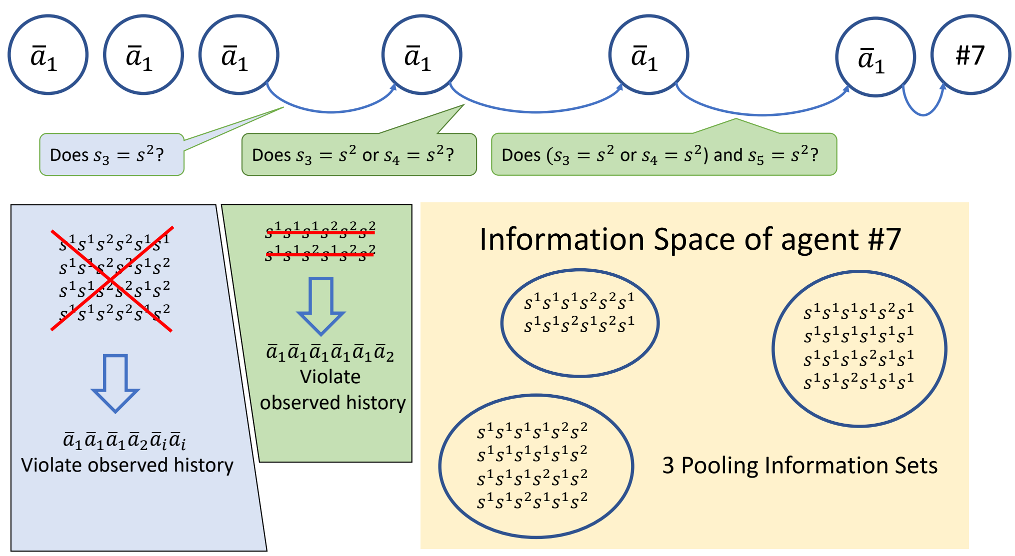

Figure 2.1: Information Space and Pooling Info Sets

Before proceeding further we note that even though for each pooling information set, there exist some questions (up to the capacity constraint) that give the current agent a positive probability for the specific pooling information set, there may not exist a set of questions (up to the capacity constraint) which can partition the information space into the set of pooling information sets. For instance, in Fig. 2.1, a 1-bit question for agent #7 has three pooling information sets as depicted in the figure. However, there is no 1-bit question that can partition the information space into three sets[12].

2.4.1 Feasibility of questions

Based on the previous discussions we know that there do not exist questions that can distinguish private signal sequences within the same pooling information set, and we can always distinguish different pooling information sets. We use these properties to define a subclass of feasible questions where each of the questions represents a partition of the information space without splitting any pooling information sets.

Definition 5. *A feasible question is deterministic if it partitions the information space without splitting any pooling information sets.*

With this definition, deterministic questions can be represented as a function.

Definition 6. *A deterministic (feasible) question q for agent t is a function q : $\mathcal{S}_t \times I_t(H_t, G, \mathcal{Q}_t) \to 2^K$, where $\mathcal{S}_t$ is the signal space of agent t, $I_t(H_t, G, \mathcal{Q}_t)$ be the information space determined by pooling information sets.*

The space of feasible questions, however, is bigger than just the deterministic questions. Conceptually it should be feasible to ask questions that include some randomness in the response such that Bayes rule can still be applied to update agents' beliefs, i.e., the responding agent is allowed to conduct an experiment (with parameters provided in the question) while responding to questions. The following example provides a non-deterministic question that should also be feasible for an agent to use.

[12] This is the main reason why we do not directly call them information sets.

Example 4. *Consider the model extended from Example 2 which only considered the first four agents, and let the crossover probability be* $1 - p = 0.4$. *Agent #5 has* $\mathcal{B}_5 = \{3, 4\}$ *and observes the history* $H_5 = \{\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\}$. *When her private signal is* s^2, *can she ask agents #3 and #4 the following question:*

> *"Please toss two biased coins first, the first one with probability of heads* $\frac{3}{4}$ *and the second one with probability of heads* $\frac{1}{5}$. *Then answer yes if getting* s^1 *with the first coin facing heads or* s^2 *with the second coin facing heads, and answer no otherwise."*

Further, can this question be answered by agents #3 and #4? Is it then possible for agent #5 to take action $\bar{a}_2$*?*

In Example 4, agent #3 and #4 know their private signal so they can follow the request of agent #5 and answer the question appropriately conditioning on the coin-tossing results. Moreover, when agent #5 gets negative responses from both agent #3 and #4, Bayes-rule will suggest her to take action $\bar{a}_2$[13]. However, this set of questions is not included in the set of deterministic feasible questions, and we need to generalize to *probabilistic questions*. Figure 2.2 depicts the difference between deterministic and probabilistic feasible questions. In probabilistic questions (right side of Figure 2.2), one pooling information set can have positive probability mass for multiple answers. Besides, when feasible probabilistic questions are given, the distribution of answers is well-defined in every pooling information set.

Definition 7. *A feasible question* Q_t *for agent* t *is probabilistic if it asks agents to draw a (different) lottery on every pooling information set with each lottery size up to the capacity constraint.*

[13]Even though the likelihood ratio of θ_2 over θ_1 is only ≈ 1.015, which is lower than the likelihood ratio based on agent #5's posterior belief in Example 3, which is 1.5.

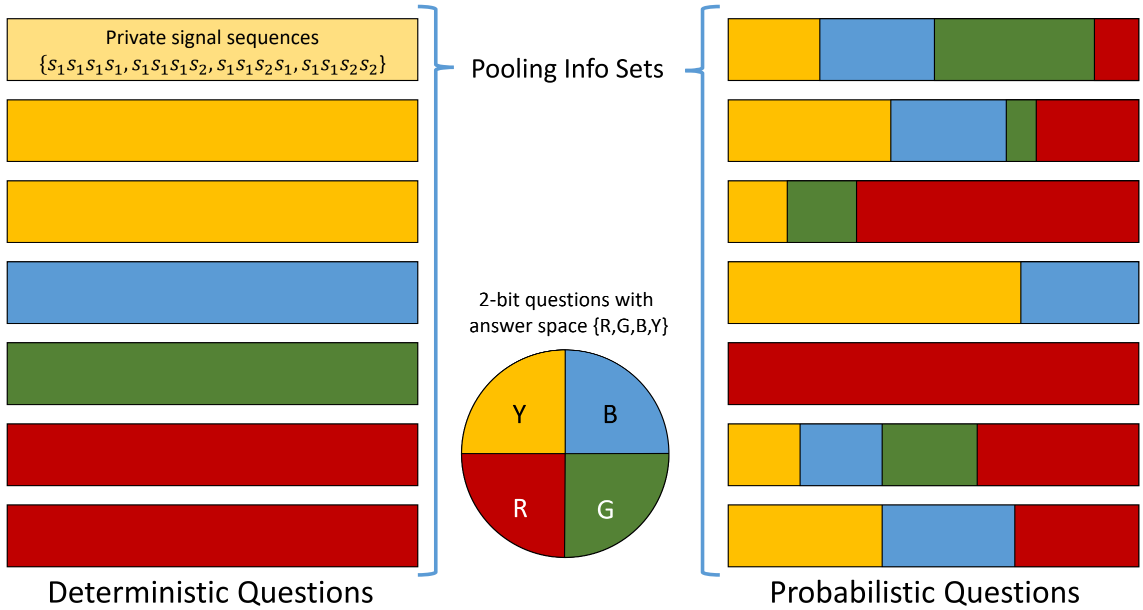

Figure 2.2: Information-set Partitioning in Deterministic and Probabilistic Questions

Allowing for probabilistic questions[14] does not complete the space of feasible questions. The following example shows other possibilities.

Example 5. *Consider the model the same as in Example 3. Agent #5 has* $\mathcal{B}_5 = \{3, 4\}$ *and agent #6 has* $\mathcal{B}_6 = \{5\}$*. Since agent #5 can query and know the private signals of agent #3 and #4 when the observed history is* $H_6 = \{\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\bar{a}_1\}$*, agent #6 can still update her posterior beliefs via the following question:*
"If the private signal sequence from agent #3-#5 is $s^1s^2s^2$*, respond A; if the private signal sequence from agent #3-#5 is* $s^1s^1s^2$*,* $s^1s^2s^1$ *or* $s^2s^1s^1$*, answer B with probability* p_1 *and C with probability* p_2*; if the private signal sequence from agent #3-#5 is* $s^2s^1s^2$ *or* $s^2s^1s^1$*, answer B or C with an arbitrary probability."*

The above question can be answered by agent #5, and agent #6 can update her beliefs on each response even though she does not know the actual distribution of B and C conditional on the state of the world. This question also can benefit agent #6 since she knows that she can take action $\bar{a}_2$ *when she gets response A and her private*

[14]We can view deterministic questions as a degeneration of probabilistic questions.

signal is s^2*.*

The tricky example in Example 5 shows that there are feasible questions outside the set of probabilistic questions: they allow for a Bayes update of beliefs and can still help agents. Fortunately, this type of questions can only arise when there are multiple responses sharing the same posterior beliefs on the state of the world. The following lemma states that we can treat these questions the same as probabilistic questions without changing the capacity constraint.

Lemma 2. *Every feasible question is either a probabilistic question or can be represented as a probabilistic question under the same capacity constraint which has the same support of posterior beliefs (updated according to possible responses).*

With the above lemma, there is no loss of generality in using the space of probabilistic questions to replace the space of feasible questions in future analyses.

2.4.2 Question guidebooks

Since the network topology is pre-determined, the set of agents that an agent can query for information is exogenous. Given the topology (assuming that we already know it), we allow the information designer to supply the agents with questions that they can ask the set of contactable predecessors $\mathcal{B}_t,\ t \in \mathbb{N}$. In general, the order of querying the agents in $\mathcal{B}_t$ matters, but in this work this issue does not arise as we only discuss the scenario $|\mathcal{B}_t| \leq 1$; this because of the difficulty of verifying incentive-compatibility of the network topology. A collection of a set of questions, supplied by the information designer, is called a *question guidebook (QGB)*. With this intuitive description in mind, we formally define a QGB.

Definition 8. *A question guidebook* Q *is a function* $Q(n, H_t, s_t, m)$ *that gives agent* t *a set of predetermined questions conditioned on agent* t*'s private signal, the observed history, and the predecessor agent* m *who agent* t *queries.*

First we recall the Assumption 2 in Section 2.3.2: agents are truthful in their responses since their (payoff relevant) action has already been taken before being queried. Next we argue that QGBs that agents are willing to follow should have two important properties: feasibility and incentive compatibility (under the current topology). We will start by defining feasible QGBs and then discuss the incentive compatibility of QGBs.

Roughly speaking, a QGB is feasible if agent t only asks feasible questions to agent m conditional on a feasible observed history. This avoids the ambiguity of what would happen if agent m does not have the information to answer her question. A critical part here is deciding whether an observed history is feasible or not. Since past agents query earlier agents, the feasibility of an observed history also depends on the QGB being used. Hence, the feasibility of QGB and observed history are coupled, and we have to define it carefully (e.g., to avoid tautological issues). This we do by the two recursive definitions below that are applied sequentially for all $t \in \{1, 2, \dots\}$ starting with $t = 1$.

Definition 9. *An observed history H_t is **feasible** under a question guidebook Q if there is a positive probability of obtaining this history given that all the agents follow the truncated question guidebook Q_{t-1}*[15]*, $\mathbb{P}(H_t|Q_{t-1}) > 0$ with $Q_0 = \emptyset$, and such truncated question guidebook Q_{t-1} is feasible.*

Definition 10. *Let Q_t denotes a truncated question guidebook Q up to agent t. Then Q_t is **feasible** if for any feasible observed history $H_t \in \mathcal{H}_t$ and private signal s_t, all questions provided by Q_t are feasible.*

In short, the definition of feasibility of history guarantees the existence of a private signal sequence $\{s_i\}_{i=1}^{t-1} \in \mathcal{S}^{t-1}$ that rationalizes each predecessor's action observed in

[15]Since the history is fully observable in our model, $H_i \subset H_j$ for all $i < j$, it is sufficient to check $\mathbb{P}(H_t|Q_{t-1}) > 0$ for the feasibility of history H_t. However, the above definition still works when the history is only partially observable by replacing Q_{t-1} by Q_w, where w is the last observable agent (if any) by agent t.

H_t. The rationality can be verified with by recursively applying backward induction via the above definitions.

With the recursive definitions in Definition 9 and Definition 10 we are ready to define the feasibility of QGBs.

Definition 11. *A question guidebook Q is* ***feasible*** *if for every agent t, under every feasible observed history H_t and s_t, all questions provided by Q are feasible questions.*

Without changing the network topology, fixing the index of agent m being queried by agent t, a QGB is *incentive-compatible* if agent t always asks a question from the set of questions that maximizes her expected payoff.

Definition 12. *A question guidebook Q is* ***incentive-compatible*** *if for every agent $t > 1$, under every feasible observed history H_t and s_t, the set of questions provided to her maximizes her expected utility among all feasible questions she can ask agents in $\mathcal{B}_t$.*

To help agents infer the predecessors' information space and then ask questions on top of their predecessors' questions to accumulate information, we make the following assumption.

Assumption 3. *The implemented QGB is common knowledge.*

The primary purpose of introducing and designing QGBs is to achieve asymptotic learning, i.e., to terminate temporary wrong information cascades. Hence, for the ease of discussion and analysis, henceforth, we concentrate on the class of QGBs only **in operation while agents are in information cascades**; this narrowing of consideration is without loss of generality in terms of the asymptotic learning objective. While an agent is not in an information cascade, her action reveals some information about her private signals[16]. This (possibly imperfectly) revealed information can be

[16]For example, in binary signal space setting, i.e., the BHW model, private signals are fully revealed via agents' actions when a cascade is not ongoing.

exploited to update successors' beliefs; ergo, questions implemented here will only shift the public belief at the moment an information cascade starts but will never overturn the outcome[17] of Bayesian observable learning problems. Note, however, that queries made when an information cascade is not occurring, could lead to faster learning, and this is not the goal of our study.

Before proceeding we discuss some exogenous feedback models via communication channels in social learning problems. These are:

- **Private-signal feedback model**: Assume that agents can get their predecessors' private signals through the channel. While $|\mathcal{S}| < 2^K$, the QGBs are feasible but not necessarily incentive-compatible, e.g., *Peres et al.* (2017); *Vial and Subramanian* (2019).

- **Summary-statistic feedback models**: Assume that agents can get some summary statistics. Usually in these models the agents are not certain about their index or the reports are inexact. The QGBs typically not feasible as $n \to \infty$, but may be incentive-compatible depending on further assumptions of agents' behavior, e.g., *Bahar et al.* (2019).

- **Decision-rule feedback**: Assume that agents follow a pre-determined decision rule to forward information. These models have a similar flavor as stochastic control problems. The QGBs are usually feasible because the Markov decision process (MDP) used has only finite states, but they need not be incentive-compatible, e.g., *Drakopoulos et al.* (2013).

2.4.3 Constructing the information space via predecessors' questions

Prior to asking questions, the current agent has to ascertain the information space of predecessors she's going to query; otherwise, she cannot decide whether she should

[17] This is in terms of whether asymptotic learning is achieved or not.

ask the questions offered by the QGB or not. Given the past questions asked by her predecessors and the observed history, each agent now can individually reconstruct the information space of every predecessor from agent #1 to her immediate predecessor in forward progression by recursively deleting infeasible private signal sequences and also inferring the actual partition of each predecessor's information space by applying questions provided by the QGB. These reconstructions, recognizing the information space of their predecessors but not fully ascertaining the information state of each her predecessors, does not require the knowledge of the answer of predecessors' questions.

Essentially, each pooling information set of the first agent being queried can be represented as a state in a state transition diagram where agent is a layer in the diagram. The capacity constraint limits the number of states in each layer. Given a QGB an agent can construct all her predecessors' information spaces and then devise questions or verify questions offered by the information designer. Since incentive-compatible questions may not be unique, this gives the information designer room to design questions to accumulate information without hurting the agents' expected payoff.

2.4.4 Temporary information cascade, indifferent questions, and helpful silent agents

Prior to designing QGBs to implement a strategy in a given topology, we know that the questions designed for agents who have a chance to stop the cascade conditional on the observed history are restricted to the set of questions that maximize their expected utility; this follows by incentive compatibility. However, for agents (if there exist any) indifferent to any question, we have the flexibility to design particular questions for them, specifically ones that will enable learning. An agent is indifferent to any question only when all the possible responses from all questions, up to the capacity constraint, do not let her take the opposite action of the current trend of action. If there is no communication capability in the model, such an event would start the (permanent) information cascade. With the aid of questions beneficial to future agents, there is the possibility that the current information cascade could be terminated because of the information accumulation via questions. Similar to results in *Smith and Sørensen* (2000) that permanent information cascade happens almost surely under full history and bounded signal strength model, the following lemma states that temporary or permanent information cascades always occur when the communication capacity is bounded.

Lemma 3. *Given a bounded signal space and assuming the capacity of each channel* $(m, n) \in E$ *is bounded, information cascades (either temporary or permanent) occur almost surely.*

Lemma 3 guarantees that QGBs will be in operation (after a cascade starts) in the model. However, different QGBs will result in different distributions of when the same[18] information cascade is stopped owing to the different choice of questions. With our goal being asymptotic learning, our aim is to not design QGBs that lead to a permanent information cascades under some history realization. As a part of this we start by pointing out a class of QGBs with a fatal flaw that we thereafter avoid in our QGB designs.

Definition 13. *A QGB is bounded if there exists a history realization up to a finite-indexed agent* t *such that actions of an infinite number of agents are determined given such a history realization under the proposed QGB.*

Since a history realization up to a finite index t has a strictly positive probability of arising, the following claim is straightforward.

Claim 2. *For every bounded QGB, almost surely we cannot achieve asymptotic learning.*

[18]While no QGBs were exercised before the start of the current cascade or all considered set of QGBs suggest the same sequence of questions before the start of a cascade.

Even though we know that we should avoid designing bounded QGBs, it is still unclear how we can design a feasible and incentive-compatible QGB that leads to asymptotic learning. In other words, it is not clear whether we can suggest a question in the set of feasible questions best for the information aggregation without violating incentive compatibility, and whether we actually have no room to design because the agent has a unique question that maximizes her expected utility. To understand the subset of feasible questions that can be used in a feasible and incentive-compatible QGB, we first classify agents into two classes; this classification emerges endogenously.

Definition 14. *Given an observed history in a feasible and incentive-compatible QGB, an agent is* ***active*** *if she has a positive probability of stopping the cascade; otherwise, she is called* ***silent****.*

To further clarify the above definition, conditional on the history, agents who may benefit from the answer to questions are active; otherwise, they are silent. A silent agent may decide not to ask questions since asking questions never benefits her. In such a scenario, her behavior is exactly the same as an agent without any communication capability. Clearly, if a silent agent will help with some probability, this probability has to be common knowledge; otherwise, the information designer and her successors have no tools to analyze it. For silent agents to help with some known probability, we have Claim 3 below.

Claim 3. *Suppose an agent is helping with a known positive probability. Then, it can be represented by another network where this agent is always helping, but every agent who queries her uses a noisy channel.*

With the above claim, the network topology is now "deterministic" since we can remove the channels of all silent agents who are not going to help successors and replace agents helping with a positive probability to agents always help by modifying the channels of agents who will query them. Thus, for the simplicity of the discussion, we make the following two assumptions.

Assumption 4. *We assume that there is no cost in asking questions.*

Assumption 5. *We assume silent agents are always willing to help their successors. That is to say, if no feasible question changes the expected payoff of an agent under the given question guidebook, an information designer can demand that this agent asks any questions*[19] *within the capacity limit.*

Assumption 4 is reasonable when only a limited number of questions are asked at no cost. Assumption 5 states that every agent is willing to help her successors when queried, in essence bringing in some level of cooperation. This is inspired by behavior in social networks, where it is common to see that people are willing to help their friend/neighbor nodes. This assumption holds (in practice) because even though no questions can benefit the current agent, her information could be beneficial to future agents and help accumulate/aggregate information. Henceforth, we will assume such behavior.

2.4.5 Equivalent statement of asymptotic learning

While restricting our attention to the topology specified in Corollary 3, we will replace the definition of asymptotic learning with the following equivalent statement.

Lemma 4. *Consider the topology in Corollary 3. Let* $e^j_{Q,t} : \mathcal{H} \times \Theta \to [0,1]$ *be a (probability) function (of the observed history (up to agent* k*)* $H_k \in \mathcal{H}$ *and the true state* $\theta_j \in \Theta$*) representing the probability of the cascade of* θ *continued in agent* k *being stopped at agent* t *conditional on the cascade continuing until agent* $t-1$*, and the question guidebook* Q *to be executed during cascades. The following two statements are equivalent:*

[19] To be provided in the question guidebook.

- $\lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_i | \theta = \theta_i) = 1$ *and* $\forall j \neq i, \lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_i | \theta = \theta_j) = 0$*;*
- $\sum_{t=k+1}^{\infty} - \ln \left(1 - e^j_{Q,t}(H_k, \theta_i)\right) = \infty$, $\sum_{t=k+1}^{\infty} - \ln \left(1 - e^j_{Q,t}(H_k, \theta_j)\right) < \infty$ *when* $j \neq i$ *and every agent* $t \geq k$ *is in a cascade.*

The first statement is the standard definition of asymptotic learning. The second statement is built on the following logic. First, agents will keep falling into cascades in this topology, either the right cascade or the wrong cascade. If every wrong cascade is stopped almost surely, then we know that we will never stay in a single wrong cascade forever, and this is stated in the first half of the second statement. Then, the second half of the second statement says that once an agent is in the right cascade, the current cascade has a positive probability of lasting forever. Hence, the second statement, combining both parts, states that as the index goes to infinity, agents will be in the right cascade almost surely, which is equivalent to the standard definition of asymptotic learning.

2.5 Asymptotic Learning via 1-bit Question in BHW model

This section will present our main analytical result in the classical BHW model, i.e., a binary signal space model with both signals generated through a conditionally independent and identically distributed binary symmetric channel. To avoid cumbersome discussions on the silent agents helping with positive probability and comparing different incentive-compatible topologies, we here assume all silent agents are willing to help and the communication capability is assumed to be $|\mathcal{B}_n| = 1$ for every agent. We also assume the network topology is a line topology which we call the **telephone-game network**[20] since like in the well-known telephone game played worldwide, each agent can only communicate with her immediate predecessor.

[20]The telephone-game network has the same topology as a tandem network, but the literature on learning in tandem networks uses one-way communication made before the next agent sees her private signal. In contrast, our questions are conditional on the received private signals. Thus, to avoid any confusion, we eschew the "tandem network" terminology.

In the first half of this section, specifically Sections 2.5.1 and 2.5.2, we will briefly review the classical BHW model, describe a high-level strategy to provide an intuition for designing QGBs, and then discuss QGBs with good properties. The second half of this section will construct a specific QGB following the strategy, argue its feasibility and incentive compatibility, and show that it satisfies the good properties introduced in Section 2.5.2. Then, the designed QGB will be analyzed to prove that asymptotic learning is achievable by using it.

2.5.1 BHW model with a telephone-game communication network

We consider the simple BHW model with an endowed telephone-game network for agents' private communications. In this model, the state of the world is binary, $\Theta = \{\theta_1, \theta_2\}$. Agents arrives sequentially, one in each different stage/time step. Each agent, once she arrives, has to take an irrevocable action $a \in \{\bar{a}_1, \bar{a}_2\}$ to guess the state of the world. Before taking her action, each agent[21] has the following sources of information to form and update her belief:

(a) For every agent t, all actions taken by her predecessors, $\{a_v | v < t\}$, are observable to her.

(b) For every agent t, the (private) signal space is binary, $S_t = \{s^1, s^2\}$, and signals are generated via a conditionally independent and identically distributed binary symmetric channel with fixed crossover probability $1 - p$, where $p > 0.5$.

(c) The responses from the previous agents whom she can query privately. Since the network topology is a line graph, every agent can only query her immediate predecessor.

Based on the total information I_t for agent t via the above sources, each agent's objective is to maximize $\mathbb{P}(a = \bar{a}|I)$ where $\bar{a}$ is the action corresponding to the state

[21] Agent 1 only has access to channel (b) from this list.

of the world θ. To avoid trivialities, we assume every agent can only ask a finite-bit (up to K bits) question for the information via channel (c). Furthermore, Assumptions 4 and 5 are also in operation.

Recall the seminal result in *Bikhchandani et al.* (1992) and *Smith and Sørensen* (2000): in the BHW model with no questions asked (i.e., the information from the source (c) is not exploited), a permanent information cascade occurs almost surely and in finite-time with a wrong cascade occurring with positive probability. Hence, a natural question that arises is: *"Is there a (general and systematic) guide for agents to query such that every wrong information cascade will be stopped?"*

Strictly speaking, the framework formulated in Section 2.4 allows us to design and analyze each agent's question manually. However, to further reduce the complexity of both the design and the analysis of questions, we only consider QGBs with some specific properties. Thus, before designing the questions, we discuss some special QGBs for which the analysis is significantly simplified.

2.5.2 Regular question guidebook and its connection to Markov chains

From the discussion in Section 2.4.3 we know that agents will need to update the posterior beliefs sequentially, and hence, to analyze them we will use dynamic systems and their transition diagrams. For the ease of design and analysis, although not restricted in the framework, we search and design QGBs that can be captured by a (small) finite number of transition matrices, which we call **regular** below:

Definition 15. *A QGB is* ***regular*** *if the following two conditions hold:*

- *Every silent agent exploits the same transition matrix to partition the pooling information sets when she gets the same private signal.*
- *When active agents know the cascade will continue after the history and their private signals are revealed, they use the same transition matrix to partition the*

pooling information sets.

The advantage of considering regular QGBs is that we can formalize the analysis of the information set partitioning via questions problem using a time-invariant Markov chain with a state size at the most 2^K, where K is the number of bits allowed for each question. By doing this, when an agent is in a specific state of the Markov chain, the only parameter that affects her action is the likelihood ratio of the observed minority versus majority, i.e., θ_j versus θ_i if we are in a cascade corresponding to θ_i, where $j \neq i$. Hereafter, such a likelihood ratio at state G_n of the Markov chain at agent t is denoted by $l^t_{G_n}$. Besides, we use θ_j represents the state different from θ_i afterward. With a relabeling of the state index, we can further restrict our attention to **aligned** regular QGBs.

Definition 16. *A regular question guidebook is* ***aligned*** *if there is a labeling of information sets such that for every silent agent n, $l^t_{G_n} \geq l^t_{G_n}$ for all $n < k$.*

An aligned regular QGB has good properties that help in analyzing the optimal actions of agents: if an active agent is going to terminate a cascade in G_k, it is also optimal for such agent to terminate a cascade in G_n for all $n < k$. Given the high-level strategy introduced in Section 2.5.3, we implement the strategy to be described in the next paragraph and show that the implemented QGB is regular and aligned.

Before proceeding we clarify why we do not restrict all active agents to use the same transition matrix, i.e., have them ask the same questions in the definition of regularity. When an active agent receives a private signal that suggests stopping the current cascade, the questions and their corresponding answers affect her utility. Therefore, restricting active agents to use the same transition matrix will make the QGB not incentive-compatible to some active agents[22]. Since we aim to design

[22]This violation of incentive compatibility can be found in all cases of signals, from a binary signal space to general finite signal spaces.

feasible and incentive-compatible question guidebooks, we cannot have such a strong restriction being enforced.

2.5.3 Threshold-based strategy

Now, we are ready to systematically design questions that can stop every wrong information cascade. Suppose we are in a temporary information cascade of $\bar{a}_i$. If the true state is θ_i, we are happy with the current cascade and willing to stay in this cascade; of course, this knowledge is known. However, if the true state is θ_j, we have to stop it eventually (hopefully within a finite horizon), otherwise, asymptotic learning cannot be achieved. Although agents never know which state is the true state, they can leverage *a priori* knowledge of the private-signal distribution. When the true state is θ_i, agents are more likely to get consecutive private signal $\{s_i ... s_i\}$ and less likely to get multiple private signals $\{s_j ... s_j\}$ in a row. Building on this rough idea, we propose the following threshold-based strategy:

Threshold-based Strategy: Iteratively monitor the event of L-consecutive observed majorities and act as follows:-

- **Step 1:** Given a predetermined threshold $L > 2$, agents in an $\bar{a}_i$ cascade ask questions to know if the event "there is a subsequence of private signals in the recent past with the number of s^i being greater than the number of s^j by at least L" occurs.

- **Step 2:** Using the information of whether the event occurs or not, an agent updates her likelihood ratio of state θ_j versus θ_i. Starting from an agent m [23] with a fixed likelihood ratio of θ_j versus θ_i (to be specified soon), there will be the first agent $n > m$ with a positive probability such that her updated likelihood ratio θ_j versus θ_i will cross (or equal to) 1. Importantly, the identity of this agent can be determined from the QGB.

[23] The "recent past" counts time from agent m onwards.

- **Step 3:** If agent n has a likelihood ratio θ_j versus $\theta_i \geq 1$, her best strategy is to stop the cascade. If she does not stop the cascade, then we use this agent n as a new starting agent (agent m in the previous step), go back to step 1, and restart a new round checking if the event occurs in the following agents.

We reiterate that the threshold-based strategy detects whether the following event occurs in every given window (i.e., between two consecutive active agents): there is a (private-signal) subsequence such that observed minorities appear at least L times more than the observed majorities. The state diagram in Figure 2.3 helps visualize this strategy. When an agent is in the first L states and her signal strengthens her belief on the current cascade, and she moves 1 step to the left in the state diagram or stays at state G_1. When an agent is in the first L states and her signal weakens her belief on the current cascade, she moves 1 step to the right. When the state is in the last state, G_{L+1} in Figure 2.3, she stays in that state till the end of current window (i.e., the appearance of the next active agent).

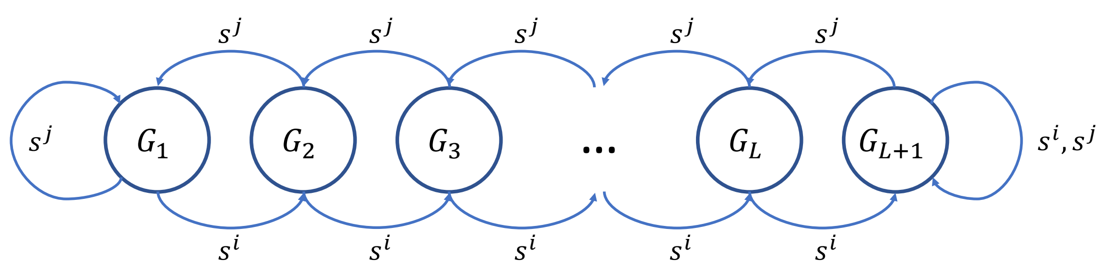

Figure 2.3: Example Diagram of a Threshold-Based Strategy

2.5.4 Implementing a threshold-based strategy

Here we implement a QGB using the threshold-based strategy introduced in Section 2.5.3. To make the design as simple as possible and to try to learn asymptotically with the minimum channel capacity, we first choose $L = 2$. When $L = 2$, questions can be answered through channels with a 1-bit capacity constraint through the proposed design detailed later in this section. In other words, in the QGB we are going to construct, we know that once two consecutive silent agents get a private signal of the same type as the current cascade, there is no chance to stop the cascade at the immediate next active agent. To make it even simpler, we will only use deterministic questions in the design of this QGB. We will demonstrate the power of probabilistic questions in Section 2.8.

Without loss of generality, we assume that the true state of the world is θ_j. In other words, an $\bar{a}_i$ cascade is a wrong cascade and an $\bar{a}_j$ cascade is a correct one. Moreover, to avoid trivial questions in the QGB, **the question guidebook operates only when a cascade is initiated**. Thus, an agent not in a cascade uses her private signal and does not need the guidebook.

Next, we illustrate the QGB by providing the corresponding Markov chains first and then detailing the questions being asked in each possible state. The Markov chains are depicted in Figure 2.4 assuming that a $\bar{a}_i$ cascade is underway[24], and they prescribe how the information space gets partitioned based on the type of the agent (active or silent), the private signal of the agent and the response from the immediate predecessor to the question (if it is asked).

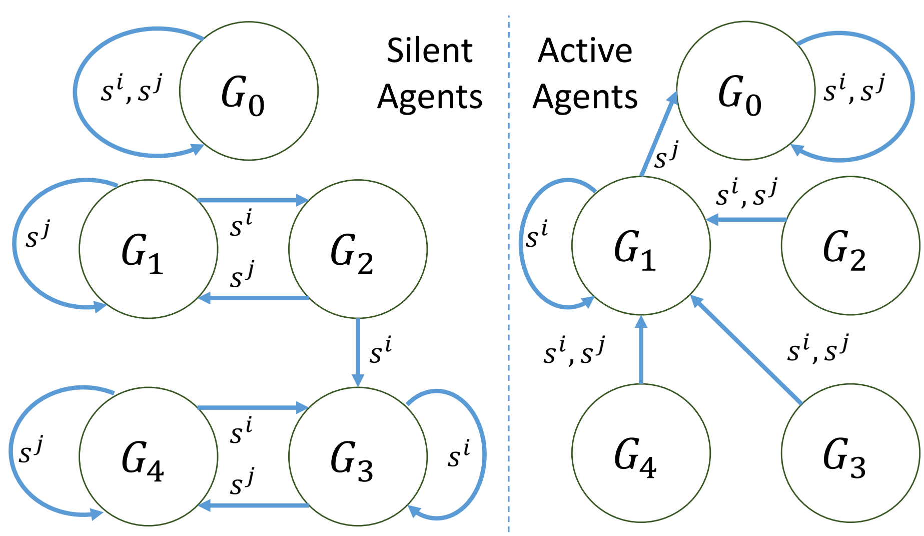

Figure 2.4: Markov chains of proposed threshold-based question guidebook

[24]In a $\bar{a}_j$ cascade, the same guidebook applies but with the θ_is and θ_js swapped.

The corresponding Markov chains in the designed QGB, as depicted in the left of Figure 2.4, endows every silent agent with the same transition matrix. This transition matrix is given by two different determined questions conditional on the silent agent's private signal[25]; and the questions and corresponding information sets are as follows.

Table 2.1: Table of questions for silent agents

	Receives s_j	Receives s_i
Question asked	Are you in $\{G_1, G_2\}$?	Are you at G_1?
Action under positive answer	Go to G_1	Go to G_2
Action under negative answer	Go to G_4	Go to G_3

As depicted in Table 2.1, while a silent agent receives private signal θ_j, she asks the question: "**Is your information set in either G_1 or G_2?**" to her immediate predecessor. Her partition is then G_1 if the answer is yes, and G_3 for a negative answer.

Au contraire, a silent agent receiving θ_i asks the question: "**Is your information set in G_1?**" to her immediate predecessor. Her partition is then G_2 for a positive answer, and G_4 otherwise.

Every active agent only cares if her immediate predecessor is in G_1 when she receives a private signal θ_i so she can stop the cascade. Hence, questions are only needed in that case. The corresponding table is attached in Table 2.2.

Table 2.2: Table of questions for active agents

<table>
<tr><td></td><td>Receives s_j</td><td>Receives s_i</td></tr>
<tr><td>Question asked</td><td>Are you at G_1?</td><td>No questions asked</td></tr>
<tr><td>Action under positive answer</td><td>Go to G_0 and stop cascade</td><td rowspan="2">Go to G_1</td></tr>
<tr><td>Action under negative answer</td><td>Go to G_1</td></tr>
</table>

Before showing that the QGB is feasible and incentive-compatible, we make the following statements to avoid any misunderstanding. We know that there exists a large set of QGBs can also achieve asymptotic learning (even with the 1-bit con-

[25]As described earlier, a silent agent's partition can never be G_0.

straint). The proposed design (and the QGB used) is elegant in simplifying the analysis and proofs, which relieves us from solving a complex recursive system of equations with four variables[26].

2.5.5 Properties of question guidebooks: feasibility and incentive compatibility

With the proposed QGB in hand, our goal is to verify its feasibility and incentive compatibility. The feasibility of this QGB is straightforward because every agent, no matter whether she is active or silent, knows her current state. Therefore, she can answer the yes-no question about her current state to pass the feasibility check. Showing incentive compatibility is equivalent to showing that G_1 is the only state that can have a likelihood ratio of θ_j over θ_i crossing 1 for any active agent. Here, we will establish a result that applies more generally than the designed QGB. We will prove that every *threshold-based* QGB has a positive probability to stop the cascade only when the immediate predecessor of an active agent is at G_1 (threshold event not holding for the current majority, i.e., $\bar{a}_i$). First, let us formally define the threshold-based QGB.

Definition 17. *A QGB is **L-threshold-based** if it satisfies the following three conditions:*

1. *Every silent agent uses the same transition matrix.*

2. *Consider every silent agent whose neighbor is in a transient state (i.e., G_i, $i < L$) of the Markov chain. She goes to state G_{i+1} upon receiving a private signal in an observed majority or goes to $G_{\max\{1,i-1\}}$ upon receiving a private signal in an observed minority.*

[26]In general, even in the 1-bit capacity model, we need to track the likelihood ratio of θ_i versus θ_j in every state. The update of likelihood ratio will make the analysis very tedious. Threshold-based question guidebooks, with the result of Lemma 6, avoid this hassle and allow us to track the likelihood ratio at only one state without loss of optimality.

3. *Consider every silent agent whose neighbor is in a recurrence state,* $G_i,\ i \geq L$. *She stays in the set of recurrence states. Furthermore, active agents continuing the cascade bring all feasible sequences of private signal to* G_1.

Given Definition 17, the following lemma claims that threshold-based QGBs possess the good properties described in 2.5.2 that aid in simplifying the analysis.

Lemma 5. *A threshold-based QGB is regular and aligned.*

Besides the properties of alignment and regularity, threshold-based QGBs have the additional technical property described below, which greatly simplifies the analysis by bypassing the update of the likelihood ratio in all states except state G_1.

Lemma 6. *In threshold-based question guidebooks, active agents can only stop the cascade at* G_1.

To show Lemma 6, we first need to guarantee there exists at least one silent agent between any pair of active agents, i.e., active agents cannot arrive consecutively. The idea of the proof is that once an active agent fails to stop a cascade, she either receives an observed-majority private signal or the response of her questions suggesting a continuation of current the cascade whatever private signal she gets. Then simple algebra rules out the possibility of consecutive active agents: see Claim 9 in Appendix B.1.6. Now, given the current active agent u_A^k, if the next active agent indexed u_A^{k+1} can stop cascade at state G_i for some $i > 1$, then agent $u_A^{k+1} - i + 1$ also has the ability to stop the cascade at G_1, which contradicts the fact that u_A^{k+1} is the next active agent. The detailed proof is in Appendix B.1.6.

Since the QGB we constructed is a threshold-based QGB and active agents stop a cascade only in G_1, Lemma 6 guarantees incentive compatibility.

2.5.6 Necessary and sufficient conditions for asymptotic learning in a threshold-based question guidebook

Before proving that the proposed QGB can achieve asymptotic learning, we first provide a necessary and sufficient condition to achieve asymptotic learning for every threshold-based QGB.

Definition 18. *Let $h^+(m)$ be the probability that the wrong cascade will be stopped when $m^k = m$, where m^k is the number of silent agents between the $(k-1)^{\text{th}}$ and k^{th} active agents. Similarly, let $h^-(m)$ represent the probability that the right cascade will be stopped when $m^k = m$.*

Recall the equivalent statement of asymptotic learning in Lemma 4. Use it we obtain the following corollary in the binary state, binary signal space case.

Corollary 1. *Given a threshold-based question guidebook Q that is operational in a cascade. The following two conditions are necessary and sufficient for the question guidebook to achieve asymptotic learning:*

1 $\lim_{k\to\infty} m^k = \infty$*;*

2 The growth rate of m^k satisfies $\Pi_{k=1}^{\infty}(1 - h^-(m^k)) > \Pi_{k=1}^{\infty}(1 - h^+(m^k)) = 0$*;*

The first condition makes the frequency of active agents go to zero as the agent index goes to infinity. Otherwise, if we have a non-zero proportion of agents that take actions according to their private signals, hence the probability of the correct action is upper bounded away from 1, and asymptotic learning cannot be achieved. Besides, we want every wrong cascade to terminate within a finite horizon, but we also need the right cascade to have a positive probability of lasting forever. This objective constrains the maximum and minimum growth rate of m^k, which is mathematically formulated in the second condition.

2.5.7 Asymptotic learning is achieved

Theorem 1. *In the telephone-game network, asymptotic learning can be achieved by a question guidebook that only requires 1-bit capacity questions.*

To show that the designed QGB satisfies the conditions in Corollary 1 we need to analyze the growth rate of m^k thoroughly. The following paragraphs will first characterize the form of m^k, then study the upper and lower bounds of its growth rate. With the bounds in hand, calculations easily show that the wrong cascade will be stopped almost surely, and the right state will last forever with positive probability (which can be lower-bounded by some constant).

Form of the number of silent agents: To study the growth rate of m^k, we need to know the exact number of agents between each pair of consecutive active agents.

Prior to this, we need to specify the functions $h^+(m)$ and $h^-(m)$ first. With threshold $L = 2$, two consecutive majority signals in a cascade will continue the cascade at the next active agent. Simple combinatorics[27] then yields $h^+(m)$ and $h^-(m)$ as follows:

$$h^+(m) = \textstyle\sum_{i=0}^{m/2} \binom{m-i-1}{i} p^{m+1-i}(1-p)^i; \tag{2.3}$$

$$h^-(m) = \textstyle\sum_{i=0}^{m/2} \binom{m-i-1}{i} p^i(1-p)^{m+1-i}. \tag{2.4}$$

Furthermore, with Lemma 6, we know that the likelihood ratio of θ_j versus θ_i at state G_1 is the only parameter that the next active agent needs to compute. Suppose we know the likelihood ratio of the first silent agent right after an active agent u_A^k, the next active agent u_A^{k+1} is the first agent who could have likelihood ratio at G_1 crossing 1. Note that $h^+(m)$ ($h^-(m)$) is the probability that an agent and her m^{th} successor

[27] Stopping a cascade needs no consecutive majority signals being observed. Hence, there must exists at least one minority signal between every two majority signals. The combinatorial equations in (2.3) and (2.4) calculate the probability of this event.

are both at G_1 conditional on the right (wrong) cascade not yet being stopped. Then the ratio of $h^+(m)$ over $h^-(m)$ is the likelihood ratio of θ_j versus θ_i conditioned on the event that the cascade continues after m silent agents. Hence, by definition of a silent agent, if the agent with index $u_A^{k-1}+m$ is silent, she must have likelihood ratio at G_1, $l_{G_1}^{H_{u_A^{k-1}+1}} \frac{h^+(m)}{h^-(m)}$, be less than 1. Given the fact that $\frac{h^+(m)}{h^-(m)}$ is a strictly increasing function of m for all $p > 0.5$, the number of silent agents m^k between active agents u_A^k to u_A^{k+1} can be mathematically defined as:

$$m^k \equiv \min \left\{ m | l_{G_1}^{H_{u_A^k+1}} \frac{h^+(m)}{h^-(m)} \geq 1 \right\}. \tag{2.5}$$

Then, since every active agent failing to stop the cascade has only one information set corresponding to G_1, there is a simple recursive form of likelihood ratio at state G_1 between first silent agents after k^{th} and $(k+1)^{\text{th}}$ active agents by using the ratio of the probability that the cascade continuous. The recursive form of $l_{G_1}^{H_{u_A^{k+1}+1}}$, a strictly decreasing[28] function of k, is given by:

$$l_{G_1}^{H_{u_A^{k+1}+1}} = l_{G_1}^{H_{u_A^k+1}} \frac{1-h^+(m^k)}{1-h^-(m^k)}. \tag{2.6}$$

With functions $h^+(m)$ and $h^-(m)$ specified, and the likelihood ration update in (2.6), m^k is non-decreasing and can be computed iteratively.

Upper bound on the growth rate of m^k: If asymptotic learning can be achieved under this QGB, every wrong cascade must be stopped, i.e., $\Pi_{k=1}^{\infty}(1-h^+(m^k)) = 0$. Thus, if we can find a sequence $w^k \geq m^k$, $\forall k$, and $\Pi_{k=1}^{\infty}(1-h^+(w^k)) = 0$, then $h^+(m^k) \geq h^+(w^k)$ guarantees $\Pi_{k=1}^{\infty}(1-h^+(m^k)) = 0$. Finding the upper bound on the growth rate of m^k is equivalent to finding the lower bound of sequence $\underline{w}_s$ such that $m^{k-1} = s-1$ and $m^{k+i-1} = s$ for all $i \leq \underline{w}_s$. Now, suppose $m^k = s$, $m^{k-1} = s-1$,

[28]See Appendix B.2.12 for the proof.

we can calculate $\underline{w}_s$ [29] to get

$$\left|\{m^t = s | l_{G_1}^{H_{u_A^{t-1}+1}} \frac{h^+(s)}{h^-(s)} \geq 1\}\right| \geq \underline{w}_s > \ln\Big(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\Big) \geq \frac{\ln(c_1(p))}{p^s + p^{2s}}, \tag{2.7}$$

where $c_1(p)$ is only a function of p.

A wrong cascade will be stopped almost surely: Taking the inequality (2.7) into the condition 2 in Corollary 1, the probability that a wrong cascade will be stopped is[30]

$$\mathbb{P}(\text{A wrong cascade can be stopped}) \tag{2.8}$$

$$= 1 - \Pi_{k=1}^{\infty}(1 - h^+(m^k)) \geq 1 - \Pi_{k=1}^{\infty}(1 - p^{m^k})$$

$$\geq 1 - \Pi_{k=1}^{\min\{j|m^{j+1}=K\}}(1 - p^{m^k}) \exp\Big(- p\ln(c_1(p)) \textstyle\sum_{n=0}^{\infty} \frac{1}{1+p^n}\Big) = 1 \tag{2.9}$$

Lower bound on the growth rate of m^k and the probability of stopping a right cascade: Similarly, to show that the probability of stopping a right cascade is strictly less than 1, we can show that a lower bound of $\Pi_{k=1}^{\infty}(1-h^-(m^k))$ is positive. Thus, we now need a lower bound on the growth rate of m^k. Using a similar technique as for (2.7), we derive an opposite inequality for $\bar{w}_s$ in Appendix B.2.11, and then the probability that a right cascade will be stopped is

$$\mathbb{P}(\text{A right cascade will be stopped}) \leq 1 - B(p)\exp\Big(\sum_{n=J}^{\infty} \frac{1}{c_4(p)c_1(p)^n - 1}\Big), \tag{2.10}$$

where $B(p) = c_2(p)e^{c_3(p)}\Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1 - h^-(k))$. Since $\sum_{n=K}^{\infty} \frac{1}{c_4(p)c_1(p)^n-1}$ converges by the ratio test, the RHS of (2.10) is less than 1, so that a right cascade can last forever with positive probability.

[29] See detailed calculations in Appendix B.2.10.

[30] See Appendix B.2.14 for detailed calculation

Now, every wrong cascade will be stopped, but there is a positive probability that a right cascade will continue forever, so the second condition is satisfied. Furthermore, since the lower bound and upper bound of growth rate $\bar{w}_s$ suggest a finite interval of $m^k = s$, m^k goes to infinity as k goes to infinity, so the first condition also holds. Thus, asymptotic learning can be achieved under this QGB. Once we are out of a cascade, agents use their private signals, which will initiate another cascade with a bias towards the right cascade. Nevertheless, every wrong cascade will be stopped in a finite horizon, and an unstoppable right cascade happens after finitely many stopped right cascades. Thus, the learning happens *almost surely* instead of just *in probability*, as in Definition 2, so that within a (random) finite horizon, learning occurs.

2.6 From Binary to Finite Signal Space

When signals are generated via a binary symmetric channel (BSC), the QGB proposed in Section 2.5 can achieve asymptotic learning. A natural question that arises with the result is whether asymptotic learning is also achieved with larger signal spaces. This section keeps all assumptions and set-ups the same as the model in Section 2.5 except for the private signals received by agents.

Here we assume that every agent receives a private signal from a finite signal space S, $|S| > 2$. Further, every signal $s_j \in S$ is informative but not revealing, i.e., $\forall i \in \{1, 2\}, s_j \in S,\ 0 < \mathbb{P}(\theta_i|s_j) < 1$. To avoid redundancy, we assume that $\mathbb{P}(\theta_i|s_j) \neq \mathbb{P}(\theta_i|s_k)$ for all $j \neq k$. Furthermore, the distribution of the signals $\Delta(S)$ is time (agent) invariant and common knowledge to every agent.

Here we will present a variant of the 1-bit threshold-based question guidebook, which we call a 1-bit lazy-active QGB, that achieves asymptotic learning with a 1-bit capacity constraint. This QGB works by passively letting a cascade emerge so that silent agents occur and then uses a similar method as our previous QGB to assimilate information in a manner to help an active agent stop a cascade, again in the same manner as before. Whereas this strategy is inefficient in terms of speed of learning (as it passively lets a cascade emerge), it nevertheless demonstrates the power of information accumulation among silent agents via small shifts of the posterior beliefs (owing to the 1-bit communication capacity constraint). Unlike in the previous section, the QGB to be detailed soon will utilize probabilistic questions (justifying our earlier careful discussion of them).

2.6.1 Lazy-active question guidebooks

In the basic BHW model with a telephone-game network, silent agents allow us to design questions we need to achieve asymptotic learning, essentially by collecting and relaying information to active agents. In the current setting, too, we again seek to leverage silent agents to help the learning process. However, with more informative signals, $|S| > 2$, at the outset, it is not clear whether silent agents will exist. Hence, we first discover a class of QGBs that guarantees the existence of a large number of silent agents and then characterize the parameters that lead to asymptotic learning.

Definition 19. *A question guidebook is called lazy-active if no active agent asks questions when she gets a weak private signal, that is, a signal where no questions will change her expected utility.*

Lemma 7. *In a telephone-game network, given any 1-bit lazy-active question guidebook, for every $\epsilon > 0$, there exists a finite integer N such that there exists a silent agent n with index $n \leq N$ with probability $\geq 1 - \epsilon$.*

An intuitive explanation of why this lemma holds is that the public belief (the likelihood ratio of the current minority (suppose state θ_2) over the majority (state θ_1)) reduces after every active agent. For each active agent taking the same action as her predecessors, she keeps ruling out feasible sequences of private signals in favor of the state θ_2. Owing to the lack of questions from the former active agents when former active agents get weak private signals in favor of θ_2 or signals in favor of θ_1, the lazy-active QGB forces the fusion of information sets. The fusion of information sets bounds the minimum rate of reduction on the public belief by some constant. Since the signals have bounded strength (otherwise, they are revealing), we will end up at an agent where the strongest signal in favor of θ_2 cannot overturn the predecessors' action ($\bar{a}_1$), and so a silent agent appears.

2.6.2 Asymptotic learning via a 1-bit lazy-active question guidebooks

Since we know that every lazy-active QGB will produce silent agents, we want to exploit these silent agents to formulate a question guidebook similar to the binary signal case discussed in Section 2.5. A simple approach is to just categorize signals into two groups C_{θ_i} and C_{θ_j}, with one in favor of action $\bar{a}_i$ and another in favor of action $\bar{a}_j$. The question a silent agent asks will only depend on whether her signal is in C_{θ_i} or C_{θ_j}. Intuitively, we then want to design the Markov chains analogous to the Markov chains in Section 2.5.

By letting silent agents ask questions only based on their private signal's group, either C_{θ_i} or C_{θ_j}, we can leverage the 1-bit threshold QGB designed in Section 4. Hence, suppose we are currently in an $\bar{a}_i$ cascade, silent agents can ask questions as in Table 2.3, similar to Table 2.1 in Section 2.5.4.

However, unlike in the binary signal space model, an active agent in this more informative signal space setting cannot stop a cascade for every signal $s \in C_{\theta_j}$. The actual subset of signals $\hat{C}_{\theta_j} \subseteq C_{\theta_j}$ that allows the current active agent k to stop a cascade will vary according to her index. This dependence arises not only because every silent agent will shift the likelihood ratio of G_1 by a constant parameter, which is the same as the model in Section 2.5. However, active agents now may appear earlier than the binary-signal model since a subset of stronger signals in favor of θ_j, $\hat{C}_{\theta_j}$, has the property $\mathbb{P}(\theta_j|\hat{C}_{\theta_j}) > \mathbb{P}(\theta_j|C_{\theta_j})$. This property may allow active agents

Table 2.3: Table of silent agents' question in the lazy-active QGB

	Private signal $s \in C_{\theta_j}$	Private signal $s \in C_{\theta_i}$
Question asked	Are you in $\{G_1, G_2\}$?	Are you at G_1?
Action under positive answer	Go to G_1	Go to G_2
Action under negative answer	Go to G_4	Go to G_3

Table 2.4: Table of active agents' question in the lazy-active QGB

<table>
<tr><th></th><th>Private signal $s \in S^k_{\theta_j}$</th><th>Private signal $s \in S \setminus S^k_{\theta_j}$</th></tr>
<tr><td>Question asked</td><td>Are you at G_1?</td><td>No questions asked</td></tr>
<tr><td>Action under positive answer</td><td>Go to G_0 and stop cascade</td><td rowspan="2">Go to G_1</td></tr>
<tr><td>Action under negative answer</td><td>Go to G_1</td></tr>
</table>

to appear earlier than in our basic model, and it also varies the update of likelihood ratio at G_1 right after every active agent. Therefore, Table 2.4 is slightly different from Table 2.2 in Section 2.5.5.

It remains to show that this design can achieve asymptotic learning. Similar to analysis in Section 2.5.7, we know that the key to this analysis is to bound the growth rate of silent agents, i.e., the variable m^k. The following lemma considerably reduces the effort needed for this.

Lemma 8. *Given the 1-bit lazy-active QGB, let m^k be the sequence of the number of silent agents between the k^{th} to $(k+1)^{th}$ active agent. Consider a binary signal space with $\mathbb{P}(\cdot|s^j) = \mathbb{P}(\cdot|s \in C_{\theta_j})$ and let the sequence of the number of silent agents in the corresponding 1-bit threshold-based QGB be $\bar{m}^k$. Then, there exist constants t_1, t_2 such that $\bar{m}^{k-t_1} \le m^k \le t_2 \bar{m}^k$ for all k.*

This lemma allows us to immediately reuse the analysis in Theorem 1 to show that we can achieve asymptotic learning. The intuition for why this lemma is true is that between every consecutive active agents, when compared to the 1-bit threshold QGB only two factors can cause the number of silent agents to be different. One factor is that unlike in the previous binary case where we stop the cascade for all signals

$s \in C_{\theta_j}$ (since there is only one), now we only stop the cascade for $s \in \hat{C}_{\theta_i} \subset C_{\theta_j}$. This will upper bound the growth rate of m^k, at most by a factor of t_2, because all the private signals are informative but not revealing. The other factor is that the strength of every group of private signals $\hat{C}_{\theta_j}$ will be stronger than C_{θ_j}, which will increase the growth rate of m^k, with the increase lower bounded by a function given by the constant t_1. The actual parameters t_1, t_2 are derived in Appendix B.1.9. With this lemma in place, we have the following proposition.

Proposition 1. *In the telephone-game network with finite signal space, there exists a QGB using 1-bit capacity questions which achieves asymptotic learning.*

2.7 Learning in a Diverse World

In Section 2.6, we studied how to achieve asymptotic learning in a binary state of the world model with finite signal spaces using a construction from the binary signal state. However, the binary state of the world assumption is a strong one. Models with non-binary states of the world, also known as diverse world, are more complex when compared to models with binary states of the world. For example, in a diverse world, the beliefs evolve in a simplex with dimension $|\Theta| - 1$ instead of in $[0, 1]$ in binary state models, and private signals partition the simplex into closed convex polytopes.

In this section, we consider a model which relaxes the binary-state assumption in Section 2.6. For simplicity of the analysis, we consider a model with a finite state space Θ with a uniform prior, a finite signal space $\mathcal{S}$, and with all other settings following the model in Section 2.5, including the telephone-game network topology[31].

[31]Note that the model in Section 2.5 only assumes bounded communication capacity on channels, the 1-bit capacity requirement for asymptotic learning is the result of the QGB design, not an imposed constraint on the model.

2.7.1 Distinguishable signals on state space

To begin with, we study the problem under a (private) signal space where asymptotic learning is possible. In other words, we focus on the models where private signals can help agents to "distinguish" different states. At first glance, the intuitive and straightforward thought is to look at models satisfying the following two conditions:

1. The signal space has cardinality equal or larger than the state space; and

2. For each state of the world θ_i, there exists a signal s_k indicating that state is the most likely state of the world, i.e., $\forall j \neq i \ \mathbb{P}(\theta_i|s_k) > \mathbb{P}(\theta_j|s_k)$.

However, these assumptions are too strong for distinguishable signals, and agents may still be able distinguish the states under a (private) signal structure which violates the above assumptions. An example of such a signal structure is provided below.

Example 6. *The state space* $\Theta = \{\theta_1, \theta_2, \theta_3\}$ *is ternary with uniform prior* $p = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. *The signal space* $S = \{s_1, s_2\}$ *is binary and satisfying the following inequality:*

$$\mathbb{P}(\theta_1|s_1) = \mathbb{P}(\theta_2|s_2) > \mathbb{P}(\theta_3|s_1) = \mathbb{P}(\theta_3|s_2) > \mathbb{P}(\theta_1|s_2) = \mathbb{P}(\theta_2|s_1)$$

To simplify the example, let's assume $\mathbb{P}(\theta_1|s_1) = \mathbb{P}(\theta_2|s_2) = \frac{1}{2}$, $\mathbb{P}(\theta_3|s_1) = \mathbb{P}(\theta_3|s_2) = \frac{1}{3}$, *and* $\mathbb{P}(\theta_1|s_2) = \mathbb{P}(\theta_2|s_1) = \frac{1}{6}$. *Given the signaling structure, states* θ_1 *and* θ_2 *can be distinguished by* s_1 *and* s_2 *respectively. Moreover, state* θ_3 *can be deciphered to be the most likely state in the sequence of private signals* $s_1 s_2$ *or* $s_2 s_1$ *by the inequality of arithmetic and geometric means.*

With the above example in mind, we have to take every possible private signal sequence into consideration to determine the distinguishability of states. The following definition defines a distinguishable signal space.

Definition 20. *A signal space $\mathcal{S}$ is distinguishable on a state space Θ if for every $\theta_i \in \Theta$, there exists a sequence of signals of finite length n, $s \in \mathcal{S}^n$, such that $\mathbb{P}(\theta_i|s) > \mathbb{P}(\theta_j|s)\ \forall j \neq i$.*

Under the assumption that the prior belief of the state of the world is uniform, that is, agents are objective Bayesians, the distinguishable signal space can be simplified as in the following corollary.

Corollary 2. *Given the uniform prior on state space Θ, a signal space $\mathcal{S}$ is distinguishable guarantees that for every pair of states $\theta_i, \theta_j \in \Theta$, there exists a signal $s \in \mathcal{S}$ such that $\mathbb{P}(\theta_i|s) \neq \mathbb{P}(\theta_j|s)$.*

Given the definition of distinguishable signal spaces, the following intuitive lemma can be proved by using an example of two states with identical probability for each signal realization.

Lemma 9. *Given the pair of state space and signal space Θ and $\mathcal{S}$, if a signal space is not distinguishable, then asymptotic learning is not achievable.*

Hereafter, we only consider state-signal pairs $(\Theta, \mathcal{S})$ that are distinguishable. According to Definition 20, we know that the least common multiple (LCM) of the shortest sequences that distinguish states, denoted by μ, is still finite under a finite state space assumption. The set of signal sequences with size μ can be partitioned into $|\Theta|$ subsets, and the set of signal sequences with index i, denoted by S^{μ}, has the property that $\mathbb{P}(\theta_i|s) > \mathbb{P}(\theta_j|s)$ for all $s \in S^{\mu}$. Based on this, we will denote the set of signal sequences with size μ, which distinguishes states, a **μ-chunk signal set** and a signal sequence within this set a **μ-chunk signal**.

2.7.2 Existence of a feasible and incentive-compatible question guidebook

To start thinking about how to construct QGBs to achieve asymptotic learning in the general model, let us put the incentive compatibility aside for the moment. By viewing each sequence of private signals with size equal to the LCM of the shortest sequences that distinguish states as a μ-chunk signal, the following lemma can be validated by periodically running a 2-state sub-QGB[32] presented in Section 2.6.2 to test whether the current cascade should be stopped.

Lemma 10. *Given a uniform prior on a state space Θ and a signal space $\mathcal{S}$ that is distinguishable, there exists a feasible QGB requiring a bounded capacity that achieves asymptotic learning.*

The above QGB violates incentive compatibility in two ways. First, agents may not be silent agents within the chunk window (because it is actually a signal sequence instead of a signal). There might be questions up to the capacity constraints that give agents a positive probability of stopping the cascade by asking narrower questions. Second, pairwise testing on states has incentive compatibility issues because continuing the current cascade (the failure of previous tests on other states) will change agents' beliefs. This change of agent's belief may cause the designed questions to violate incentive compatibility and make stopping the cascade (following the suggested action of the QGB) be an irrational action. Therefore, we have to design QGBs to overcome these two issues.

The first issue is similar to the issue we faced in the finite signal space but binary state model. The technique of bounding the shift of active agents used in Lemma 8 might again be applied here to resolve our problem once we know the set of signal chunks we are going to adopt. However, the second issue is more troublesome. In the binary state setting, there's a clear alternative to check against and stopping a current cascade is essentially the same as preferring the alternative. Such a simple association does not hold beyond binary states, and the testing a subset of states within the current cascade can unevenly discriminate against states not checked owing

[32] To be more clear, suppose we are currently in an θ_i cascade, we run a QGB presented in Section 2.6.2 for every $\theta_j \neq \theta_i$ to determine whether we should stop the cascade of θ_i because of the likelihood ratio of beliefs on θ_j versus θ_i.

to the generality of signal space. Our approach here would be to construct a set of chunk signals such that even with this uneven discrimination the drifts on the beliefs are the same. For this we will rely heavily on the distiguishability property.

To describe our approach, we start by representing each pure μ-chunk signal by posterior distribution on states induced by it. Note that we have exactly $|\Theta|$ μ-chunk signals - one for each underlying state. Using the representation based on the posterior distributions on the states, we can ascertain whether the centroid (or barycenter) of the $|\Theta|$-simplex of probability vectors (i.e., the uniform distribution), lies in the convex hull of the set of $|\Theta|$ μ-chunk signals. Distinguishability of the signal space as given in Definition 20, guarantees that this is true in a strong sense as given below.

Claim 4. *Given a distinguishable signal space $\mathcal{S}$ on a state space Θ, there exists a μ-chunk signal set $\boldsymbol{\mu}_{\mathcal{S}}$ such that the centroid of the $|\Theta|$-simplex is in the relative interior of the convex hull of $\boldsymbol{\mu}_{\mathcal{S}}$.*

With Claim 4 we know that the convex hull of $\boldsymbol{\mu}_{\mathcal{S}}$ contains contains a convex polytope similar to the $|\Theta|$-simplex with both sharing the same centroid. Let $\mathbf{1}$ be the all-ones vector of dimension $|\Theta|$, then the centroid of the $|\Theta|$-simplex is $\frac{1}{|\Theta|}\mathbf{1}$. Also define the unit vector in coordinate $i \in \Theta$ to be e_i. Then the implication from Claim 4 is that there is a small enough $\epsilon > 0$ such that convex hull of $\{\epsilon e_i + \frac{1}{|\Theta|}\mathbf{1} : i \in \Theta\}$ is strictly contained in the convex hull of $\boldsymbol{\mu}_{\mathcal{S}}$; the symmetric form of the extreme points is a consequence of the uniform prior. Each (extreme) point, namely $\epsilon e_i + \frac{1}{|\Theta|}\mathbf{1}$ for $i \in \Theta$, being in the convex hull of $\boldsymbol{\mu}_{\mathcal{S}}$ implies that each can be represented by a probabilistic feasible question sequence of length μ. Hence, we will call these extreme points probabilistic μ-chunks, with the entire collection denoted by $\boldsymbol{\mu}_{\mathcal{S}}^{P}$, and the convex hull the probabilistic μ-chunk simplex (even though the ϵ is not unique). The discussion above is summarized in the lemma below.

Lemma 11. *Given a uniform prior on state space Θ, and a signal space $\mathcal{S}$ that is distinguishable, then there exists a probabilistic μ-chunk simplex.*

Using the constructed probabilistic μ-chunk simplex, we bound the indices of active agents, and hence the number of silent agents between consecutive active agents, using the same technique as in the proof of Proposition 1. These bounds can then be used to prove that asymptotic learning is once again achievable.

Proposition 2. *Given a uniform prior on state space Θ and a distinguishable signal space $\mathcal{S}$, there exists a feasible and incentive compatible QGB using bounded communication capacity that achieves asymptotic learning.*

2.8 Discussion

Next we discuss several threads that either discuss modifications of our model or help to justify it.

2.8.1 Capacity growing with agent index

For our QGBs we assumed that each agent could only contact a bounded number of predecessors and also that the communication capacity is bounded. We discuss how relaxing these assumptions makes the learning problem a lot easier as less relaying of information needs to be carried out.

We start with the most unrestricted network: all prior agents are contactable, so that every agent is free to ask questions to any agents with a lower index than her (if any). Moreover, there are no capacity constraints on these communications. Every agent can get perfect information by asking enough questions to prior agents to achieve asymptotic learning; this would occur say by each agent using an appropriately chosen likelihood-ratio test. In the most naive case, every agent asks all past agents their private signals. Hence, this naive strategy will require $O(n)$ bits of information per agent.

Since agents are homogeneous, learning can be achieved with much reduced communication through the "backward level tracking" scheme proposed below.

We start by reasserting that if an agent is not in a cascade, she will take actions according to her private signal. For an agent in a cascade, suppose a $\bar{a}_j$ cascade, there are three different cases detailed below:

1. If she gets s_i, she plays $\bar{a}_j$ without asking any questions.

2. If she gets s_j, she asks her immediate predecessor her private signal. If her predecessor's private signal is s_i, she can safely play $\bar{a}_i$. The reason is that she knows that there must be an agent in the history who has the same value of $\#s_i - \#s_j$ and that person would've played $\bar{a}_j$. Hereafter, we set $\#s_i - \#s_j$ in the history up to agent n as the *level* of agent n.

3. If she gets s_j and her immediate predecessor also got s_j, she can safely play $\bar{a}_j$ only if she knows the current level is over 2. She learn the level by asking agents $n-1, n-2, \ldots$ about their private signals. Assuming that the agents played according to the algorithm described above, she would then know the level of each agent, and from this could compute her level too.

The step (3) may require questioning many agents. In the worst case, if the true state of the world is θ_i, for large n, agent n will likely need to query all the way back to the beginning, i.e., the first agent. However, agent n may be able to stop the querying much earlier when she knows her level relative to specific agents. Assuming that all previous agents play correctly, agent n only needs to ask until she finds either a s_j agent at her level that takes action $\bar{a}_i$ or any agent at one level lower that takes an $\bar{a}_i$ action. We call such an agent an **index agent**. The logic for this is the following: If an s_j agent at her level takes action $\bar{a}_i$, then agent n should take the same action as well. If an agent exactly one level down takes action $\bar{a}_i$, that agent must have had a non-negative level. This implies that agent n is at a strictly positive level, and so she should take action $\bar{a}_j$. Notice that while agent n does not know what level she is on, she can compute her level relative to the agents she is querying.

The remaining question is how far back an agent will need to look to find either the first agent, an agent that plays a different action, or an index agent. In the long run, if the correct state of the world is θ_j, all agents will eventually play $\bar{a}_j$. Thus s_j agents will play $\bar{a}_j$ (because the prior agent did). An s_i agent will need to look back until he finds an index agent. However, this is essentially a (downward) biased random walk, and so the expected number of steps until locating an index agent is constant. However, it is only constant in the average sense. If the graph of the level goes down for $\log(n)$ steps (as is likely at some point), then agent n will have to ask at least $\log n$ prior agents to find an index agent. This scheme will not work for agents who are allowed to ask a fixed number of queries. Our surprising result in Section 2.5 is that allowing for strategically chosen questions and relaying of information by silent agents, agents can ask just the immediate predecessor a fixed number of queries, with our achievable scheme using just one query.

2.8.2 Noisy communication channels

A noisy communication channel between agents can be represented by a $K_n \times K_n$ row stochastic matrix M, where and M_{ij} denotes the probability of the response j is received conditional on the true response is i. If the channel is noiseless, the matrix is an identity matrix with rank K_n. Given the prior belief of states $v \in \Delta(\Theta)$ where v_i represents the probability of the state is θ_i. Then a question is a $K_n \times K_n$ row stochastic matrix Q and the posterior belief of state θ_i given the response j is $\frac{v_i Q_{ij}}{\sum_{k=1}^{|\Theta|} v_k Q_{kj}}$. When a noisy communication channel gets applied to this question, the posterior belief of state i given the response j becomes

$$\frac{v_i \sum_{l=1}^{K_n} Q_{il} M_{lj}}{\sum_{k=1}^{|\Theta|} \sum_{l=1}^{K_n} v_k Q_{kl} M_{lj}},$$

where we implicitly assume that $\sum_{k=1}^{|\Theta|} \sum_{l=1}^{K_n} v_k Q_{kl} M_{lj} > 0$.

Note that this is equivalent to a question Q' where $Q'_{ij} = \sum_{l=1}^{K_n} Q_{il} M_{lj}$. The fact that Q' is a valid question holds because M is a row stochastic matrix. This guarantees that $\sum_{j=1}^{K_n} Q'_{ij} = \sum_{j=1}^{K_n} \sum_{l=1}^{K_n} Q_{il} M_{lj} = 1$. Hence, applying a noisy communication channel M on a question Q is equivalent to using a question Q', and our methodology can then be applied for the analysis of the resulting QGB. For example, when the noise communication channel parameters are common to all agents, we can carry out this transformation for all possible questions Q, and then check if there exists a Q' that satisfies our conditions for asymptotic learning. However, determining the constraints this imposes on feasible channels M for which asymptotic learning is possible, is complicated and is left for future work.

2.8.3 Deterministic and randomized topology

Unlike a deterministic network where the topology is the common knowledge of every agent, another common assumption in the literature is to use a randomized network topology where the statistics of the network topology is common knowledge. In more detail, the randomized network model assumes that the distributions of edges are common knowledge, but the realized set of predecessors $\mathcal{B}_n = \{m | (m, n) \in E\}$ that agent n can communicate is private information of agent n. Under this assumption, if an information cascade occurs and the information that is allowed to get from predecessors with communication channels is restricted to their private signals, it is analogous to the model in *Acemoglu et al.* (2011) with neighbor set $N_n = \mathcal{B}_n \cup \{i\} \cup \{i+1\}$, where agent $i, i+1$ are the agents to start the cascade. Then, Theorem

4 in *Acemoglu et al.* (2011) states the condition that asymptotic learning can be achieved under randomized networks; the result relies on an information monotonicity property arising from the randomized network topology assumption. In contrast, things are much tricky in deterministic network models because an informational Braess's paradox might occur. In the context of our model an informational Braess's paradox occurs when two information sources that both individually suggest agent n to take action $\bar{a}_j$ (and ignore her private signal) may collectively make agent n take action $\bar{a}_j$ instead. An example is constructed below to illustrate this phenomenon that distinguishes deterministic and randomized network problems.

2.8.3.1 Informational Braess's paradox in deterministic networks

We present a counter-intuitive example where informational monotonicity fails in a deterministic network. In a nutshell, two (or more) sources of information both suggest an agent ignoring her private signal and taking action $\bar{a}_j$ may eventually make the agent taking action $\bar{a}_j$.

Consider the observable history is represented by network $\mathcal{G}$ contains two subnetwork $G_1(U, E_U)$ and $G_2(V, E_V)$ and two nodes W_1, W_2. Let $e^U_{i,j}$ represents the edge between U_i and U_j, and similarly for $e^V_{i,j}$.

Now, define the topology of the first subnetwork $G_1(U, E_U)$ such that $E_U = \{e^U_{1,3}, e^U_{2,3}, e^U_{3,j} \forall j \in \{[J] \setminus \{1,2,3\}\}\}$, where J is a fixed integer. Then define the topology of the second subnetwork $G_2(V, E_V)$ such that $E_V = \{e^V_{1,3}, e^V_{2,3}, e^V_{1,k}, e^V_{3,k} \forall k \in \{[K] \setminus \{1\}, \{2\}, \{3\}\}\}$, where K is another fixed integer.

Then, the whole network $\mathcal{G}$ is defined as follows: $\mathcal{G} = (U \cup V \cup \{W_1, W_2\}, E_U \cup E_V \cup E_{W_2})$, where $E_{W_2} = \{(i, W_2) | i \in U \cup V \cup W_1 \setminus \{U_2, V_1, V_2\}\}$. In short, W_2 can observe all agents' actions except for agents U_2, V_1, and V_2. We assume that agents take actions sequentially. Specifically, agent X_i takes action before agent X_j for all $i < j$, $X \in \{U, V, W\}$ and agent X_i takes action before agent W_2 for all $X \in \{U, V\}$

and for i in $[J]$ or $[K]$ respectively. The specific order of taking actions for agents across sets U or V is immaterial. The directed graph is depicted in Figure 2.5. Given

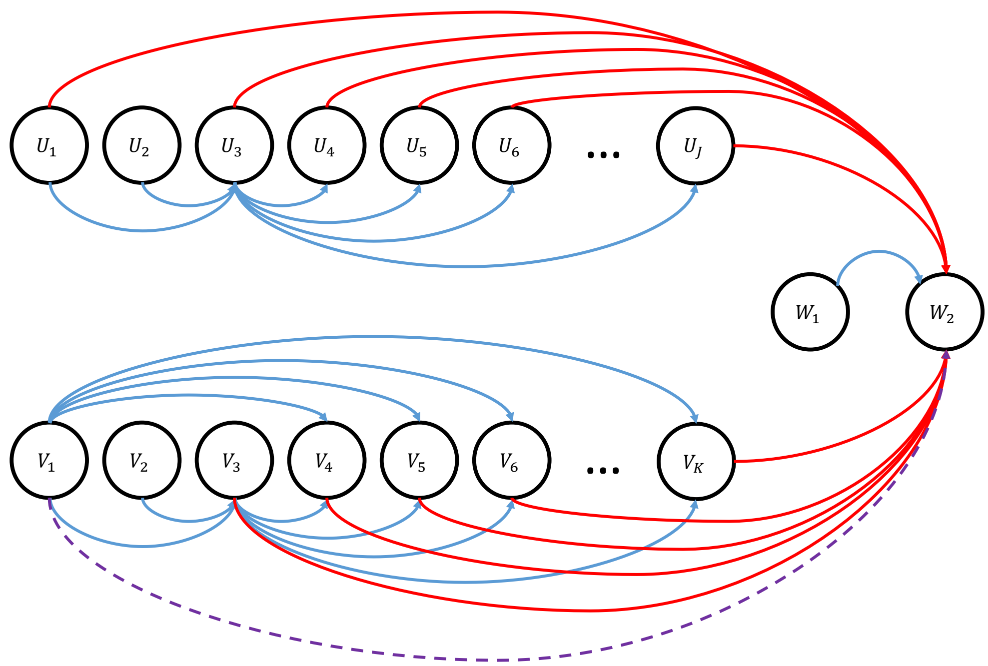

Figure 2.5: Deterministic network example to illustrate an informational Braess's paradox in social learning

this topology, consider a realized observation of W_2 below:

- Every agent U_j with $j \geq 3$ takes action $a_{U_j} = \bar{a}_i$,and $a_{U_1} = \bar{a}_i$.
- Every agent V_k with $k \geq 3$ takes action $a_{V_k} = \bar{a}_j$.
- Agent W_1 takes action $a_{W_1} = \bar{a}_i$.

By calculating the likelihood ratio at W_2, we know that W_2 should ignore her private signal and take action $\bar{a}_i$. In other words, we know that when agent W_2 observes the realized history above, there is an information cascade on $\bar{a}_i$.

Now, if on top of these observations, agent W_2 now can observe the action taken by agent V_1 (the purple dashed line in the figure) and $a_{V_1} = \bar{a}_j$, which when taken on its own will suggest that agent W_2 take action $\bar{a}_i$. Surprisingly, if we calculate the likelihood ratio of W_2 using the observation of both sources of information (original red lines and new purple line), agent W_2, instead of being in an information cascade and taking action $\bar{a}_i$, will continue to ignore her private signal and remain in an information cascade but will take action $\bar{a}_j$ when $K \geq 5$. In this case, although both sources of information suggest an information cascade, by considering the feasible sequence of the private signals using both sources of information, the agent can choose to initiate an information cascade opposite to the one without the new information. Hence, the new information still leaves agent W_2 confounded.

Instead of presenting the detailed calculations of the likelihood ratio, we point out the high-level idea of why such an informational Braess's paradox can happen in deterministic networks. In deterministic networks, agents know the whole network topology. Therefore, they can build high-order beliefs conditional on the past agent's topology to get some information about some predecessors' actions that they cannot observe directly. In the example shown above, before observing the action taken by V_1, agent W_2 doesn't know if the action $a_{V_k} = \bar{a}_j$ she observes is taken by agent V_k using her private signal or if V_k is in an information cascade on action $\bar{a}_j$. Therefore, she needs to consider both scenarios to update her posterior belief if a_{V_1} (the purple line in the figure) is not observable. However, once she knows that $a_{V_1} = \bar{a}_i$, she knows that every $a_{V_k},\ k \geq 3$ is taken by her predecessor's private signal instead of arising from within a cascade. When K is large enough, this will prompt agent W_2 to drop her private signal, and hence to initiate an information cascade on action $\bar{a}_j$.

2.8.4 Extension from Bayesian Nash equilibrium to perfect Bayesian equilibrium

Based on our analysis and proposed algorithms, this chapter establishes asymptotic learning in social learning models under the Bayesian Nash equilibrium (BNE) concept. However, perfect Bayesian equilibrium (PBE) can also be achieved with the following belief system: every off-equilibrium path, i.e., an information set of zero probability, leads to an act solely based on the corresponding agent's private signal when Bayes rule cannot be applied. We will detail the belief system below.

To get a PBE refinement of our BNE, we must establish a belief system that can address both on-equilibrium and off-equilibrium behaviors. Since the definitions and discussions on feasibility and incentive compatibility of QGBs are built on the equilibrium path, we highlight that a QGB must contain a feasibility check of questions; otherwise, the truthful-report assumption is not well-defined when an agent deviates. To formulate a game tree and extend our equilibrium concept from a BNE to a PBE, we extend a feasible and incentive-compatible QGB with an additional feature for queried agents to truthfully report that the questions are not feasible.

To construct the game tree, we first emphasize that an agent's information set contains the following information:

- A private signal of the agent,
- The state of the world,
- All the private signals of previous agents,
- Responses received by previous agents (including any realizations of any random variables used for the responses),
- The response the agent receives,
- The (public) history of actions.

Given her information set (a node in the game tree) containing the information mentioned above, an agent makes two strategic decisions: (1) asking questions, (2) taking action. Theoretically, the node corresponding to this agent in the game tree has the descendants' number equal to the product of the action space and the cardinality of the set of feasible questions (according to her information set). However, this agent's successors cannot distinguish many branches because the responses to questions are private. Thus, when a QGB is given, an equivalent game tree contains 6 descendants in every node, detailed below:

- The first two descending nodes represent that the agent asks questions suggested by the QGB, and takes action $\bar{a}_1$, $\bar{a}_2$, respectively.

- The next two descending nodes represent the agent who does not ask questions suggested by the QGB. Moreover, her questions make the successor's questions infeasible. Since the agent has two different choices on actions, these two nodes represent that the agent takes action $\bar{a}_1$, $\bar{a}_2$ respectively.

- The last two descending nodes capture a set of corner cases. In this set of corner cases, the agent does not ask questions suggested by the QGB. However, her successor asks questions not in the QGB either. However, her deviation makes her successor's questions become feasible. This type of corner case will occur when a sequence of agents deviate, but they (accidentally) deviate to the same (another) QGB. Note that an agent cannot know the exact deviation (on questions) of her predecessors before asking her questions. Hence, this type of cases cannot occur based on agents' collaboration on deviation. In short the successor who deviates will be surprised when her questions are appropriately answered. Since the agent has two different choices on actions, these two nodes represent the cases that the agent takes action $\bar{a}_1$, $\bar{a}_2$ respectively.

First, we explain why the actions are sequentially rational, and the belief system is consistent in the first two nodes. When the agent follows the QGB, her successor's questions (based on the QGB) can be properly answered. Thus, it is either in the on-equilibrium path or off-equilibrium path where the agent takes (Bayes) irrational action. Bayes rule can be applied when it is in the on-equilibrium path, and the belief system based on the BNE analysis is consistent. Hence, to make a consistent belief system for these two nodes, we only need to look at the case where the agent takes (Bayes) irrational action. In our model, a QGB only operates when a cascade occurs. Therefore, the agent who takes (Bayes) irrational action is either active or silent. When she is a silent agent, her successors' belief update does not depend on her actions. Hence, her deviation on actions does not affect the belief-updating procedure. That is to say, if the QGB forms a BNE, it also guarantees that using the MAP rule strategy is sequential rational for her successor. Thus, the belief system is consistent when a silent agent deviates. Now, suppose the agent who deviates is active. Then her deviation either mistakenly continues a cascade or mistakenly stops a cascade. In both cases, Bayes rule can be applied for her successor for belief updating, and the MAP rule strategy is sequentially rational for her successor. (Actually, her successor cannot know whether this active agent deviates from her (sequentially-rational) action, so Bayes rule is applied as in the no-deviation case, which makes the belief consistent here.)

Second, in the second two nodes where agents deviate from the QGB, and her deviation makes her successor's question infeasible, we will use the following strategy to make the belief system consistent and the action (sequentially) rational.

- The immediate successor forms her belief based on her private signal.
- The immediate successor takes an MAP-rule based action using her private signal.
- All other successors view the immediate successor as the first agent and apply the same QGB when a cascade starts.

Since the Bayes rule cannot be applied for the immediate successor, PBE allows us to formulate an arbitrary belief for the immediate successor. After the immediate successor, the following agents can use Bayes rule by applying the original QGB and believe that the original QGB is used (after the immediate successor) with probability one. Thus, the belief is consistent for these two nodes, and applying the MAP rule is again sequentially rational.

Last, in corner cases where an agent and her successor deviate to another QGB, the questions can be appropriately answered till another deviation on questions or a probability-zero history occur. Hence, Bayes rule can be applied in these off-equilibrium paths whenever possible.

Given the belief system we formulate based on the QGB, actions in the on-equilibrium path are sequentially rational, and Bayes rule can be applied not only in the on-equilibrium path but also in off-equilibrium paths whenever possible. Thus, we can extend the equilibrium concept in this chapter from BNE to PBE.

2.8.5 A subtly weaker assumption on common knowledge of the question guidebook

At the end of Section 2.3, we assumed that the whole question guidebook is common knowledge. However, it is no loss to use a slightly weaker assumption that only ***the feasibility and the compatibility of the proposed question guidebook*** is common knowledge. To see the difference between these two assumptions, we have to consider the set of question guidebooks containing agent index dependent questions.

Consider the scenario that we have two different question guidebooks, Q_1 and Q_2, both achieving asymptotic learning. Besides, these two QGBs will become operational only when a cascade is initiated. Consider a joint question guidebook Q^* which uses questions in Q_1 when an odd-index agent initiates a cascade and uses questions in Q_2 when an even-indexed agent initiates the cascade. (This question guidebook can still achieve asymptotic learning, but that is not our main point here.) Applying the original assumption, we need to disclose the whole question guidebook to all agents. However, it is straightforward that after an agent observes the history and knows that the cascade was initiated by an odd-indexed agent or an even-indexed one, half of the new QGB is useless for her for updating her posterior beliefs. By restricting to the assumption which commits to revealing to all the feasibility and incentive compatibility of the QGB, the best response of every agent is to ask the questions suggested by the QGB when a part of the QGB is common knowledge in this scheme.

2.8.6 Importance of the existence of silent agents

In order to justify why we think the existence of silent agents is critical for achieving asymptotic learning, we provide a QGB example where every agent is active, but asymptotic learning cannot be achieved.

Consider the signal space $S = \{s_A, s_B, s_a, s_b\}$, where $\mathbb{P}(\theta_i|s_A) > \mathbb{P}(\theta_i|s_a) > \mathbb{P}(\theta_i|s_b) > \mathbb{P}(\theta_i|s_B)$. We say that s_A, s_B are strong signals in favor of θ_i, θ_j and s_a, s_b are weak signals in favor of θ_i, θ_j. For the simplicity of analysis, we assume that $\mathbb{P}(\theta_i|s_a) + \mathbb{P}(\theta_i|s_b) = 1$ and $\mathbb{P}(\theta_i|s_A) + \mathbb{P}(\theta_i|s_B) = 1$, and we also assume that $\mathbb{P}(s_a) = \mathbb{P}(s_b) = \mathbb{P}(s_A) = \mathbb{P}(s_B) = \frac{1}{4}$. Suppose also that a strong signal is enough to overturn the action indicated by two weaker signals, i.e., $\frac{\mathbb{P}(\theta_j|s_B)}{\mathbb{P}(\theta_i|s_B)}\left(\frac{\mathbb{P}(\theta_j|s_a)}{\mathbb{P}(\theta_i|s_a)}\right)^2 > 1$ and $\frac{\mathbb{P}(\theta_i|s_A)}{\mathbb{P}(\theta_j|s_A)}\left(\frac{\mathbb{P}(\theta_i|s_b)}{\mathbb{P}(\theta_j|s_b)}\right)^2 > 1$, the following haste-waste question guidebooks will make every agent be an active agent, but asymptotic learning cannot be achieved.

Definition 21. *A haste-waste question guidebook is a QGB suggesting the following questions, where assume the first agent takes action $\bar{a}_i$ without loss of generality.*

The second agent will take action $\bar{a}_i$ when she gets signal s_B. When she gets signal s_b, she should ask the question whether the first agent got signal s_a. When she gets s_a, no questions can improve her expected utility, but we let the second agent ask the question whether the first agent got signal s_a. When she gets s_A, no question (or trivial question) is suggested for her.

The third agent will take action $\bar{a}_i$ when she gets signal s_a, s_b or s_A. When she gets signal s_B, she should ask the second agent the question whether the first two agents' private signal is one of the following two cases: $s_1 s_2 \in \{s_a s_a, s_A s_b\}$. When she gets *one of the signals s_a or s_b, we also have her enquire if the first two agents' private signals are both s_a.*

If the fourth agent gets signal s_B, then she should ask the third agent if she either got s_B or got s_b and the positive response from the second agent. She will take $\bar{a}_j$ if she gets a positive response. And take action $\bar{a}_i$ otherwise. Now, when she gets signal s_a, s_b, we want her asks if the third agent got either s_a or s_b and whether she received a positive response from the second agent.

For every agent with index ≥ 5, when she gets private signal s_a, s_b or s_B, she asks her immediate predecessor the following question: "Did you get signal s_a or s_b and receive a positive response from your question?" Then, she will take action $\bar{a}_j$ only if she gets private signal s_B and positive response from her immediate predecessor.

Now the following lemma can be easily verified.

Lemma 12. *The haste-waste QGB is feasible and incentive compatible.*

However, even though the haste-waste QGB is feasible and incentive-compatible, it is straightforward that this QGB cannot achieve asymptotic learning. The reason is that once an active agent with index ≥ 3 gets the private signal s_A, all the future responses will be negative in this question guidebook. The failure of this QGB indicates that without the existence of silent agents, it might be hard for us to design a QGB to take care of all the possible private signals to still achieve asymptotic learning.

2.8.7 Incentive compatibility on network topology

Having discussed QGBs at length, we now revisit the issue of incentive compatibility on topology mentioned in Section 2.3.2. Given a fixed topology, agents can easily verify the feasibility and incentive compatibility of QGBs by iteratively constructing the information space using the predecessors' questions. In other words, we already know the set of questions that can maximize agent n's expected payoff when she is asked to query agent m. Now, to study the determination of the network topology in an incentive-compatible manner, we seek to determine the set of agents who should be queried for the current agent.

When $|\mathcal{B}_n| \geq 2$, the order of questions matters. This holds because the response that agent n gets from the first predecessor she chooses to query would update her beliefs and change the questions to the next predecessors she would like to query. Hence, to verify the incentive-compatibility of a given network topology, when $|\mathcal{B}_n| \geq 2$, requires that one checks all combinations of the order in every $\mathcal{B}_n$. Unfortunately, we have not found out a systematic approach to searching and verifying to incentive compatibility of a class of network topologies give the information of $|\mathcal{B}_n|$ for all n when $|\mathcal{B}_n| \geq 2$. Hence, we leave this for future work.

However, when $|\mathcal{B}_n| \leq 1$ for every n, the incentive-compatibility check can be carried out systematically. On social networks such as Twitter or Facebook, the $|\mathcal{B}_n| \leq 1$ assumption insists that an agent can only contact one of their predecessors before making decisions. In these environments, agents take actions (e.g., pressing the "like" button, sharing, retweeting) relatively quickly, and querying a large number of their predecessors becomes impractical and even not possible (because an agent may not know all her friends' friends). Considering this, we restrict attention to the setting that only allows agents to query (at most) one of their predecessors. This restriction avoids the issue of the order of questions, and the property in the following proposition helps us analyze the means to aggregate information through questions.

Proposition 3. *Given a set of network topologies with the property $|\mathcal{B}_t| \leq 1$ $\forall t$ and K-bit communication capacity for each channel ($K \geq 1$), there exists an incentive-compatible network topology that every agent n with $|\mathcal{B}_n| = 1$ chooses $\mathcal{B}_n = \{m\}$ satisfying one of the following conditions:*

- $|\mathcal{B}_m| = 1$ *and* $\forall m < i < n$, $|\mathcal{B}_i| = 0$;

- *m is the first agent in the current cascade.*

Proposition 3 tells that when every agent is allowed to query at most one of their predecessors, querying the agent that asked questions most recently is one of the best responses to every agent who is endowed with the communication power. The proof idea is to show that asking the one who asked questions right before the current agent can benefit the agent more than asking questions to any predecessor without communication capacity or any other predecessors asking questions before the most recent one. In the proof process, we find out that although the topology is incentive-compatible, it may not be the unique incentive-compatible network topology. The non-uniqueness of the topology arises when the current history is heavily pointing to one state and no questions, whoever the current Bayesian agent queries, will let her take the opposite action irrespective of the private signal she receives. In other words, this Bayesian agent now is indifferent among all questions to all predecessors up to her capacity constraint. In a more restricted case, we can show the uniqueness of a telephone game network topology.

Corollary 3. *Given $|\mathcal{B}_n| = 1\ \forall n$ and insisting that an agent has to commit to the identity of the agent that she will query before the observing the history, there exists a unique incentive-compatible network topology that every agent n with $|\mathcal{B}_n| = 1$ chooses $\mathcal{B}_n = \{n-1\}$ such for all $n > 1$.*

2.9 Conclusion

By removing all exogenous assumptions on types of information transmission in networks, we study the power of communication capability on Bayes-rational agents and present a framework on how to utilize communication channels via querying. Questions help to partition information spaces, and enable a systematic analysis on unlimited order of inference among Bayes-rational agents in such social learning problems. When we restrict our attention to a basic social network model where every agent can only query the one sharing the post with her (her immediate predecessor), we show that agents can avoid wrong cascades and learn asymptotically by asking just a 1-bit binary question to the preceding agent in the seminal sequential social learning model *Bikhchandani et al.* (1992). To show this, we developed a question guidebook that agents can use to ask questions such that each queried agent can always answer the question, and agents always ask a question that best serves their interest. We then extended our result from binary signal space to discrete (finite) signal spaces by designing a question guidebook to seek information about a set of signal sequences instead of a particular signal sequence. Finally, we extended our result from binary state space to finite signal space by constructing a signal-chunk simplex within the convex hull of the signal-chunk set and design sequences of questions to distinguish the state of the world. Our results demonstrate the power of strategic communications while also alluding to the complexity of designing and analyzing them.

CHAPTER III

Persuasion in Sequential Trials

3.1 Introduction

Information design studies how informed agents (senders) in a game can persuade uninformed agents (receivers) to take specific actions by influencing the uninformed agents' beliefs via information disclosure. The canonical Kamenica-Gentzkow model *Kamenica and Gentzkow* (2011) is one where the sender can commit to an information disclosure policy (signaling strategy) before learning the true state. Once the state is realized, a corresponding (randomized) signal is sent to the receiver. The receiver then takes an action, which results in payoffs for both the sender and the receiver. Senders in information design problems only need to manipulate the receivers' beliefs with properly chosen signals. The manipulated beliefs will create the right incentives for them to spontaneously take specific actions that benefit the sender (in expectation); the signals must account for the receiver's incentive compatibility constraints.

Motivated by many real-world problems listed in Section 3.1.1, this chapter studies persuasion schemes where the sender is constrained in the choice of signals available for information design. Specifically, we are interested in problems that are naturally modeled via multi-phase trials where the interim outcomes determine the subsequent experiments. Further, we insist that some of the experiments are given in an exogenous manner. This feature imposes restrictions on the sender's signaling space, and without it, we would have a classical Bayesian persuasion problem with an enlarged signal space. Our goal is to study the impact of such constraints on the optimal signaling scheme, and in particular, to contrast it with optimal signaling schemes in classical Bayesian persuasion. The following motivating example describes a possible real-world scenario for a sender.

Example 7. *A start-up is seeking funds from a venture capital firm. The process for this will typically involve multiple rounds of negotiation and evaluation: some of these will be demonstrations (or a pitch) of the start-up's core business idea, and the others will be assessments by the venture capital firm following their own screening procedures. The start-up will have to follow the venture capital firm's screening procedures but choose its product demonstrations. Based on these stipulations, the start-up needs to design its demonstrations to maximize its chance of getting funded.*

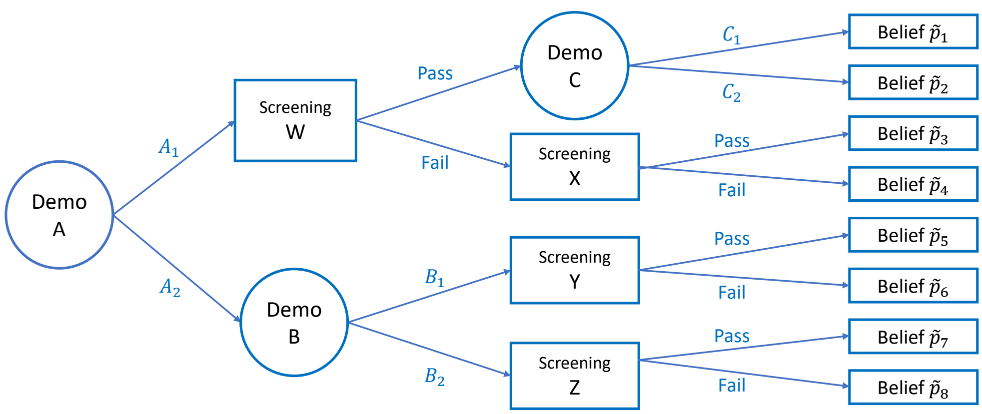

Figure 3.1: Example of a negotiation process between a start-up company and a venture capital firm.

In the example above, the start-up (sender) has to generate an optimal information disclosure scheme to get the desired funds from the venture capital firm (receiver). Then the screening procedures set by the venture capital firm are analogous to our determined experiments, and the demonstrations carried out by the start-up are the (sender) designed experiments. For example, in Figure 3.1 we present one plausible interaction where the start-up company designs demonstrations A, B, and C (circles in the figure) and the venture capital firm designs screening examinations W, X, Y, and Z (rectangles in the figure). Whereas we have illustrated this example via a balanced tree, if we have an unbalanced tree owing to the receiver deciding in the middle, we can modify it to a balanced tree by adding the required number of dummy stages.

The sender's reduced flexibility in choosing signaling strategies in a sequential trial with some predetermined experiments distinguishes our work from the growing literature on dynamic information design. Our model considers a problem with the following features: a static state space, a sequential information disclosure environment, and a signaling space restricted by some exogenous constraints. Models with a static state space, an unrestricted signal space but a variety of sequential information disclosure environment have been studied to capture features in different real-world problems: e.g., with multiple senders *Li and Norman* (2018a); *Forges and Koessler* (2008), with costly communication *Honryo* (2018); *Nguyen and Tan* (2021), allowing for sequential decision making *Ely and Szydlowski* (2020), or with partial commitment *Au* (2015); *Nguyen and Tan* (2021). The work in *Ely* (2017); *Farhadi and Teneketzis* (2020); *Meigs et al.* (2020) considers sequential information disclosure problems in dynamic state models where the sender has information about the underling dynamically changing state with the setting being either state change detection *Ely* (2017); *Farhadi and Teneketzis* (2020) or routing games *Meigs et al.* (2020). Although several works *Gradwohl et al.* (2021); *Dughmi et al.* (2016); *Le Treust and Tomala* (2019) also consider constrained signaling schemes, these works either consider the signal space to be smaller than the action space *Gradwohl et al.* (2021); *Dughmi et al.* (2016) or consider a noisy signaling environment *Le Treust and Tomala* (2019). The sequential information disclosure in our model, which actually enlarges the signal space, makes our work different from *Gradwohl et al.* (2021); *Dughmi et al.* (2016); *Le Treust and*

Tomala (2019). Besides, Doval and Ely consider the combination of mechanism design and information design to study the interaction of both approaches and the implications of different equilibrium concepts from dynamic games in a sequential information design environment in *Doval and Ely* (2020). The approach and results in *Doval and Ely* (2020) are not relevant to our work's core contributions. However, there are relations to our model's extension considering the receiver's optimal determined experiments; those are discussed in Section 3.5.2.

Aside from the literature of dynamic information design, the topic of a constrained sender has received significant attention in the literature with the constraints arising in a varied manner. The work in (*Perez-Richet and Prady*, 2012; *Perez-Richet*, 2014) proves that pooling equilibria result if the receiver either performs a certification or validation process; the results allow for the signals of the sender to be exogenously constrained. There is a large body of work that considers scenarios where the constraints arise either due to communication costs (*Gentzkow and Kamenica*, 2014; *Hedlund*, 2015; *Carroni et al.*, 2020; *Nguyen and Tan*, 2021), capacity limitations (*Le Treust and Tomala*, 2019), the sender's signal playing multiple roles (*Boleslavsky and Kim*, 2018) (e.g. incentivizing a third party to take a payoff-relevant action), or costs to the receiver for acquiring additional information (*Matyskova*, 2018). These works characterize either the applicability of the concavization approach (*Kamenica and Gentzkow*, 2011), the optimal signaling structure, or the conditions for the optimality of certain signaling structures. There is also a burgeoning literature on constraints on the sender arising from a privately informed receiver (*Kolotilin et al.*, 2017; *Guo and Shmaya*, 2019; *Doval and Skreta*, 2018; *Candogan*, 2020; *Candogan and Strack*, 2021). The main contributions here are to characterize the optimal signaling structure with a key aspect being the fact that the sender constructs a different signal for each receiver type. Finally, the single experiment design work (*Kolotilin*, 2015) studies constraints on the sender that arise from the receiver having access to some publicly available information correlated with the state world; the latter is similar in construction to our determined experiments. Within this context, the paper then studies comparative statics on the sender's utility based on the quality of the sender's information or the public information; at a high level, these conclusions are related to our results in Section 3.5.2 where we allow the receiver to design some or all of the determined experiments.

In many real-world applications, the number of phases in the sequential trial procedure is often large: for example, the number of experiments required in developing and releasing new drugs or advertisements aiming to overturn a brand image. When the sender needs to design a large number of experiments jointly, an algorithmic approach exploiting the structure of models is desired for solving the optimal signaling strategy instead of solving using an LP formulation where the number of constraints grows exponentially; the constraints arise from the receiver's incentive compatibility constraints. In this aspect our problem shares some similarity with the work in on algorithmic information design, such as *Dughmi and Xu* (2019), *Dughmi* (2017), *Babichenko and Barman* (2017), and *Celli et al.* (2019).

According to the motivating example illustrated in Figure 3.1, the persuasion problem considers a sequence of experiments where experiments further along in the tree depend on the outcomes of previous phases. The experiment to be run in each phase is either exogenously determined or chosen by the sender. In the game, the sender chooses designed experiments with the knowledge of prior, the determined experiments, and the receiver's utility function, but before the state of the world is realized. After the sender commits to the experiments (i.e., the signaling strategy), the state of the world is realized, and a specific sequence of experiments is conducted based on the realization of the underlying random variables. The receiver then takes an action depending on the entire sequence of outcomes. The prior, the utility functions of the sender and the receiver, the determined experiments, and the designed experiments (after the sender finalizes them) are assumed to be common knowledge. We study this problem for binary states of the world, first for a two-phase binary outcome trial, and then generalize to multi-phase binary outcome trials. We then discuss how to generalize to non-binary experiments (still with an underlying binary state space). Finally, we consider games with an additional stage where the receiver moves before the sender to decide some or all of the determined experiments, perhaps with some constraints.

Contributions: The main contributions of this work are three folds:

1. To the best of our knowledge, we are the first to study how a sender should design signaling schemes in a multi-phase trial where some experiments are exogenously determined (before the game starts).

2. When the experiments have binary outcomes, we explicitly solve the optimal signaling scheme in a two-phase trial. Using structural insights gained from the two-phase trial, we present a dynamic programming algorithm to solve the sender's optimal signaling in a general multi-phase trial via backward iteration.

3. Using structural insights of the optimal signaling scheme for two-phase trials, we study the impact of constraints on the sender via the determined experiments by contrasting the performance with the classical Bayesian persuasion setting and when using classical Bayesian persuasion optimal signaling schemes when the sender is constrained. As a part of this, we also provide a sufficient condition for when a sequential trial is equivalent to classical Bayesian persuasion with a potentially enlarged signal space.

3.1.1 Real-world problems with multiple phases

In the introduction, we present a motivating interview example. Here, we provide more real-world Bayesian persuasion problems that contain multi-phase trials.

3.1.1.1 Graduate application

A junior undergraduate who wants to apply for graduate schools must submit her transcript and three recommendation letters from her professors. She wants to get offers no matter whether she is mature enough or not. The admission committee, who always wants to recruit mature students, will first do a GPA screening and then read the recommendation letters. The student can choose to take any combination of courses as she wishes, but the recommendation letters must come from the professors of taken courses. There are some easy courses with a high average GPA, but those instructors' recommendation letters provide minimal information of whether she is mature enough. On the contrary, there are lab courses that the instructor can tell whether she is mature enough but these have a higher chance of getting a low average GPA. The student has to decide the best combination of courses that she will take based on the above distinction.

3.1.1.2 Multi-round interview

A new graduate wants to get an offer from a company. The department head of the company wants to hire a person who is competent for the position. To get an offer, the new graduate has to pass the interview. The interview has multiple rounds, and she can choose different strategies in each round to sell herself. According to her strength and personality traits that she shows in previous rounds, the department head will ask different questions in the final interview to evaluate whether the new graduate is competent for the position and decide whether the offer should be granted.

3.1.1.3 Product Launch

A product development team invents a new product and wants to launch it. The company has to decide whether it should launch such a new product or not. The product development team can decide the number of simplified trial products that will be sent out. Based on those trial products' customer feedback, the company will decide to offer it in different selected branches. Given the feedback and the sales report of branches, the company will decide whether to launch the product or not.

3.1.1.4 Advertisement of New products

A company wants to sell a new product to a targeted customer, e.g., a cleaning robot, and the customer wants to buy the product if it is worth the money. The company can choose different advertisement methods to reach out to the customer, e.g., mailed ads, in-store posters, online ads, etc. After seeing the ads, the customer may choose different ways to evaluate whether it is worth buying the new product, conditional on how she encountered the ads. Hence, a company must take the customer's possible evaluation technique into account while determining the best way to reach out to the customer.

3.1.1.5 Buy-out

A start-up wants to be acquired by a tech giant. The tech giant always wants to decide whether the acquisition will benefit its growth. To showcase its business model's potential, the start-up can first choose a consulting firm to evaluate some of its business. Based on the consulting firm's evaluation report, an accounting firm will audit the value of the business (or parts of it), and the tech giant will decide its acquisition strategy based on both reports.

3.2 Problem Formulation

There are two agents, a sender (Alice) and a receiver (Bob), participating in the game. We assume binary states of the world, $\Theta = \{\theta_1, \theta_2\}$, with prior belief $p \coloneqq \mathbb{P}(\theta_1)$ known to both agents. The receiver has to take an action $a_1, a_2 \in A$ which can be thought of as a prediction of the true state. We assume that the receiver's utility[1] is given by $U_R(a_i, \theta_j) = 1_{\{i=j\}}$ for all $i, j \in \{1, 2\}$. To preclude discussions on trivial cases[2] and simplify the analysis, we assume that the sender always prefers the action a_1, and her utility is assumed to be $U_S(a_1, \theta_i) = 1$ and $U_S(a_2, \theta_i) = 0$ for all $i \in \{1, 2\}$.

Before the receiver takes his action, a trial consisting of multiple phases will be run, and the outcome in each phase will be revealed to him. In each phase, the sender will conduct an experiment, which is chosen according to the outcomes in earlier phases. Hence, the experiment outcomes in earlier phases affect the interim belief and influence the possible (sequence of) experiments that will be conducted afterward.

In the most sender-friendly setup where the sender can choose any experiment in each phase without any constraints, the problem is equivalent to the classical Bayesian persuasion problem with an enlarged signal space. However, when some experiments are pre-determined conditional on a set of outcomes, the sender must take these constraints into account to design her optimal signaling structure.

To present our results on the influence of multiple phases on the sender's signaling strategy[3], we start with a model of two-phase trials with binary-outcome experiments in the rest of this section. We then analyze the optimal signaling strategy of this model in Section 3.3. After that, we will introduce the general model of multiple-phase trials with binary-outcome experiments and propose a systematic approach to analyze the optimal signaling structure in Section 3.4.

---

[1]Setting the receiver's utility in a binary-state model to be $U_R(a_i, \theta_j) = 1_{\{i=j\}}$ for all $i, j \in \{1, 2\}$ is without loss of generality. To make $U_R(a_i, \theta_j) = 0 \ \forall i \neq j$, we can normalize the utility of the lower utility action (under each given state) to zero. To symmetrize the receiver's utility difference in both states, we can keep the sender's utility unchanged and map the original problem to a new one with a modified (scaled) prior.

[2]In the binary state model, with state-dependent (and action-dependent) preferences, the sender's preferences on actions have two possibilities vis-a-vis the receiver's preference. The first case is the preferences are perfectly aligned, which leads to a trivial strategy of a truthful revelation. The second case is the preferences are polar opposites, where no information will be revealed. Thus, we only consider cases where the sender's preference is state-independent.

[3]Assuming that the sender does not have full freedom to design the experiments in each phase.

3.2.1 Model of two-phase trials with binary-outcome experiment

There are two phases in the trial: phase I and phase II. Unlike in the classical Bayesian persuasion problem, our goal is for the sender to not have the ability to choose the experiments to be conducted in both phases of the trial. We will start by assuming that the sender can choose any binary-outcome experiment in phase I, but both the phase-II experiments (corresponding to the possible outcomes in phase I) are determined. Formally, in phase I, there is a binary-outcome experiment with two possible outcomes $\omega_1 \in \Theta_1 = \{\omega_A, \omega_B\}$ and each outcome corresponds to a determined binary-outcome experiment, E_A or E_B, which will be conducted in phase II, respectively. The sender can design the experiment in phase I via choosing a probability pair $(p_1, p_2) \in [0,1]^2$, where $p_i = \mathbb{P}(\omega_A|\theta_i)$. Once the probability pair (p_1, p_2) is chosen, the interim belief of the true state $\mathbb{P}(\theta_1|\omega_1)$ can be calculated while respecting the prior as follows:

$$\mathbb{P}(\theta_1|\omega_A) = \frac{pp_1}{pp_1 + (1-p)p_2}, \qquad \mathbb{P}(\theta_1|\omega_B) = \frac{p(1-p_1)}{p(1-p_1) + (1-p)(1-p_2)}. \tag{3.1}$$

On the other hand, the phase-II experiments are given in an exogenous manner beyond the sender's control. In phase II, one of the binary-outcome experiments, $E \in \{E_A, E_B\}$ will be conducted according to the outcome, ω_A or ω_B of the phase-I experiment. If ω_A is realized, then experiment E_A will be conducted in phase II; if ω_B is realized, the experiment E_B will be conducted in phase II. Similarly, we can denote the possible outcomes $\omega_2 \in \Theta_2^X = \{\omega_{XP}, \omega_{XF}\}$ when the experiment E_X is conducted, where notation P, F can be interpreted as passing or failing the experiment. Likewise, the phase-II experiments can be represented by two probability pairs $E_1 = (q_{A1}, q_{A2})$ and $E_2 = (q_{B1}, q_{B2})$, where q_{Xi} denotes the probability that the outcome ω_{XP} is realized conditional on the experiment E_X and the state θ_i, i.e., $q_{Xi} = \mathbb{P}(\omega_{XP}|\theta_i, E_X)$.

In real-world problems, regulations, physical constraints, and natural limits are usually known to both the sender and the receiver before the game starts. Hence, we assume that the possible experiments E_1, E_2 that will be conducted in phase II are common knowledge in this chapter. Given the pairs $(q_{A1}, q_{A2}), (q_{B1}, q_{B2})$, the sender's objective is to maximize her expected utility by manipulating the posterior belief (of state θ_1) in each possible outcome of phase II. However, since the phase II experiments are predetermined, the sender can only indirectly manipulate the posterior belief by designing the probability pair (p_1, p_2) of the phase-I experiment. As the sender prefers the action a_1 irrespective of the true state, her objective is to select an optimal probability pair (p_1, p_2) to maximize the total probability that the receiver is willing to take action a_1. Recalling the receiver's utility function discussed above, the receiver's objective is to maximize the probability of the scenarios where the action index matches the state index. Thus, the receiver will take action a_1 if the posterior belief $\mathbb{P}(\theta_i|\omega) \geq \frac{1}{2}$ and take action a_2 otherwise. After taking the receiver's objective into the account, the sender's optimization problem can be formulated as below:

$$
\begin{aligned}
&\max_{p_1,p_2} \sum_{\omega_2 \in \{\omega_{AP}, \omega_{AF}, \omega_{BP}, \omega_{BF}\}} \mathbb{P}(a_1, \omega_2) \qquad (3.2)\\
\text{subject to}\quad & \Big(\mathbb{P}\big(\theta_1|\omega_{AY}, q_{A1}, q_{A2}, p_1, p_2\big) - \tfrac{1}{2}\Big)\Big(\mathbb{P}(a_1, \omega_{AY}) - \tfrac{1}{2}\Big) \geq 0 \;\forall\; Y \in \{P, F\},\\
& \Big(\mathbb{P}\big(\theta_1|\omega_{BY}, q_{B1}, q_{B2}, p_1, p_2\big) - \tfrac{1}{2}\Big)\Big(\mathbb{P}(a_1, \omega_{BY}) - \tfrac{1}{2}\Big) \geq 0 \;\forall\; Y \in \{P, F\},\\
& \mathbb{P}(\omega_{AP}) = pp_1q_{A1} + (1-p)p_2q_{A2},\\
& \mathbb{P}(\omega_{BP}) = p(1-p_1)q_{B1} + (1-p)(1-p_2)q_{B2},\\
& \mathbb{P}(\omega_{AP}) + \mathbb{P}(\omega_{AF}) = pp_1 + (1-p)p_2,\\
& \mathbb{P}(\omega_{BP}) + \mathbb{P}(\omega_{BF}) = p(1-p_1) + (1-p)(1-p_2),\\
& \mathbb{P}(a_1, \omega_2) \in [0,1] \;\forall \omega_2 \in \{\omega_{AP}, \omega_{AF}, \omega_{BP}, \omega_{BF}\}, \qquad p_1, p_2 \in [0,1].
\end{aligned}
$$

In the sender's optimization problem (3.2), the first two inequalities are the Incentive-Compatibility (IC) constraints that preclude the receiver's deviation. The IC constraints can be satisfied when both terms in the brackets are positive or negative. That is to say; the sender can only persuade the receiver to take action a_1/a_2 when the posterior belief (of θ_1) is above/below 0.5. While we have written the IC constraints in a nonlinear form for compact presentation, in reality they're linear constraints. The next four equations are constraints that make the sender's commitment (signaling strategy) Bayes plausible[4]. Hence, there are 4 IC constraints and 4 Bayes-plausible constraints in the optimization problem for a two-phase trial. However, in an N-phase trial, both the number of IC constraints and the number of Bayes-plausible constraints will expand to 2^N each. Although the linear programming (LP) approach can solve this optimization problem, solving this LP problem in large Bayesian persuasion problems can be hard *Dughmi and Xu* (2019) unless we accept a weaker ϵ-equilibrium concept. Hence, instead of solving this optimization problem via an LP, we aim to leverage structural insights discovered in the problem to derive the sender's optimal signaling structure.

We end this section by emphasizing that when there is at least one determined experiment and at least one designed experiment, this model is the only non-trivial two-phase trial configuration. In other configurations such that the sender can design at least one of the phase-II experiments in a two-phase trial, the model will be equivalent to a single-phase trial (with a different prior if the experiment in phase-I is a determined experiment).

3.3 Binary-outcome Experiments in Two-phase Trials

In this section, the sender's optimization problem presented in Section 3.2.1 is solved starting with the simplest non-trivial case. There are only two phases in the trial studied here, and from this, we will develop more insight into how different types

[4]A commitment is Bayes-plausible *Kamenica and Gentzkow* (2011) if the expected posterior probability of each state equals its prior probability, i.e., $\sum_{\omega\in\Theta}\mathbb{P}(\omega)\mathbb{P}(\theta_i|\omega) = \mathbb{P}(\theta_i)$.

of experiments (determined versus sender-designed) influence the optimal signaling strategy of the sender. To be more specific, we will analyze how two determined experiments (in phase II) and one sender-designed experiment (in phase I) will impact the sender's optimal signaling strategy. Before presenting the general case, we will discuss a subset class of two-phase trials similar to single-phase trials. In this class of two-phase trials, in one of the phase-II experiments, called a *trivial* experiment, the outcome distribution is independent of the true state. Trivial experiments *Basu* (1975), also called non-informative or Blackwell non-informative experiments in some literature, are frequently used as benchmarks to compare the agents' expected utility change under different signaling schemes/mechanisms, e.g., *Nguyen and Tan* (2021); *Li and Zhou* (2016); *Meigs et al.* (2020). This two-phase model with a trivial experiment tries to capture real-world problems with one actual (and costly) experiment, e.g., clinical trials, venture capital investments, or space missions. Since the experiment is costly, a screening procedure is provided to decide whether it is worth conducting the experiment. We will then analyze the optimal signaling strategy in the general scenario, where both experiments in phase II are *non-trivial*.

3.3.1 Experiments with screenings

As mentioned above, we start by analyzing the sender's optimal strategy (signaling structure) in a simple scenario where there is one *non-trivial* experiment conducted in phase II. This non-trivial experiment is conducted only if the sample passes/fails the screening. The sender's authority on choosing the probability pair (p_1, p_2) controls the screening process. To avoid any ambiguity, we first define a *trivial* experiment.

Definition 22. *An experiment E is trivial if the distribution of its outcomes $\triangle_E$ is independent of the state of the world: $\triangle_E = \triangle_{E|\theta_i}$ for all $\theta_i \in \Theta$.*

In other words, if a trivial experiment (in phase II) is conducted, the posterior belief of the state will stay the same as the interim belief derived in (3.1). When there exists a trivial experiment in the two phase-II trial options, then Lemma 13 states that the sender and the receiver's expected utility under the optimal signaling strategy is the same as in the classical Bayesian persuasion problem with a single-phase trial.

Lemma 13. *In a binary-state Bayesian persuasion problem, we are given two different persuasion schemes: one is a single-phase trial, and another one is a two-phase trial with a sender-designed phase-I experiment and a trivial experiment in phase II, e.g.,* $q_{A1} = q_{A2}$*. Under each scheme's optimal signaling strategy, both sender and receiver's expected utilities are the same in these two schemes.*

In the single-trial classical Bayesian persuasion setting, the optimal signaling strategy only mixes the two possible states in one outcome (e.g., when the prosecutor claims the suspect is guilty). On the other outcome, the sender reveals the true state with probability one (e.g., when the prosecutor says the suspect is innocent). When there is a trivial experiment in phase II, the other experiment (supposing that it will be conducted at outcome ω_B) will be rendered defunct by the sender's choice of experiments in phase I[5]. This phenomenon occurs because the sender can always choose to reveal the true state when the non-trivial experiment is to be conducted, i.e., by setting $\mathbb{P}(\theta_1|E_B) = 1$ or $\mathbb{P}(\theta_2|E_B) = 1$. With this choice, the classical Bayesian persuasion solution can be replicated. In essence, having a trivial experiment as one of the Phase-II trials does not constrain the sender.

3.3.2 Assumptions and partial strategies

Next, we detail the optimal signaling strategy in our two-phase trial setting with general binary-outcome experiments. To aid in the presentation and to avoid repetition, we make two assumptions without loss of generality and introduce several explanatory concepts before the analysis.

Claim 5. *We can make the following assumptions without loss of generality:*

[5] As the interim belief is either 0 or 1.

- *The probability of passing a phase-II experiment when the true state is* θ_1 *is greater than or equal to the probability of passing a phase-II experiment when the true state is* θ_2*, i.e.,* $q_{A1} \geq q_{A2}$ *and* $q_{B1} \geq q_{B2}$.

- *When the true state is* θ_1*, the experiment conducted when outcome* ω_A *occurs is more informative*[6] *than the experiment conducted when outcome* ω_B *occurs, i.e.,* $q_{A1} \geq q_{B1}$.

The sender's strategy consists of the following: choice of phase-I experiment parameters (p_1, p_2) and the persuasion strategies in phase-II for each outcome of the phase-I experiment. To understand better the choices available to the sender and her reasoning in determining her best strategy, we will study the possible persuasion strategies in phase-II; these will be called partial strategies to distinguish them from the entire strategy. Given the assumptions above on phase-II experiments, it will turn out we can directly rule out one class of partial strategies from the sender's consideration. The other set of partial strategies will need careful assessment that we present next.

Claim 6. *When the inequalities of the assumption* $q_{A1} \geq q_{A2}$ *and* $q_{B1} \geq q_{B2}$ *in Claim 5 are strict, for any phase-II experiment* $E_X \in \{E_A, E_B\}$*), taking action* a_1 *when* E_X *fails but taking action* a_2 *when* E_X *passes is not an incentive compatible strategy for the receiver for any interim belief* $\mathbb{P}(\theta_1|E_X) \in [0, 1]$.

The above claim can be verified by comparing the posterior belief $\mathbb{P}(\theta_1|\omega)$ of each possible outcome $\omega \in \{\omega_{XP}, \omega_{XF}\}$ and its corresponding receiver's best response using the IC constraints in (3.2). Therefore, upon the outcome of a phase-I experiment being revealed (to be either ω_A or ω_B), the sender only has three different classes of "partial strategies" by which to persuade the receiver:

[6] In terms of the Blackwell informativeness from *Blackwell* (1953).

(α_X) Suggest action a_1 only when the phase-II experiment outcome is a pass, i.e., ω_{XP} occurs;

(β_X) Suggest action a_1 no matter the result of phase-II experiment; and

(γ_X) Suggest action a_2 irrespective of the result of phase-II experiment, which is equivalent to not persuading the receiver to take the sender-preferred action.

Given these three classes of partial strategies and the freedom to choose different partial strategies based on the phase-I experiment's outcome, the sender can use any combination of these 3^2 choices to form a set of strategies $\mathcal{S}$. To simplify the representation, we use (c_A, d_B), $c, d \in \{\alpha, \beta, \gamma\}$ to represent a strategy of the sender. Note that to specify a strategy S within the set of strategies, i.e., $S \in \mathcal{S}$, the probability pair (p_1, p_2) has to be determined first. Before we analyze the different strategies, we discuss the relationship between the given phase-II experiments, partial strategies, and the incentive-compatibility requirements[7] from the sender's side. To avoid ambiguity, hereafter, when we mention incentive compatibility/incentive compatible (IC) requirements/IC strategies, we mean the condition/requirements/strategies of a sender's commitment satisfying the following statement[8]: for every possible realized signal under this commitment, the receiver taking the sender-suggested action is incentive-compatible.

[7]This is to avoid any profitable deviation from the receiver side.

[8]To avoid confusion, the term "incentive-compatible commitment" is analogous to the concept of incentive-compatible mechanism from mechanism design. A commitment in the information design literature is a signaling mechanism and its corresponding suggested actions under every possible signal realization. While viewing a commitment as a mechanism, a commitment is incentive-compatible when the receiver's best response is to take sender-suggested actions for every possible signal realization under this commitment. In short, it is incentive compatible for the receiver to follow suggested actions in this signaling mechanism. Therefore, an incentive-compatible commitment only guarantees the best response from the receiver side (according to this commitment) but not the sender side.

3.3.3 Constraints given by phase-II experiments

By her choice of the experiment in phase-I, the sender decides how to split the prior into the interim beliefs for the two experiments available in phase-II. The resulting interim-beliefs lead to certain partial strategies at stage-II being applicable, i.e., incentive compatible (for the receiver). In other words, the probability pair (p_1, p_2) must make each partial strategy (of this strategy) yield the maximum utility for the receiver. These requirements constrain the sender's choice of (p_1, p_2), and she needs to account for the (reduced) choice while deciding the split of the prior. Table 3.1 summarizes the impact in terms of the parameters of the phase-I experiment via primary[9] requirements on (p_1, p_2) driven by the incentive compatibility while using each class of partial strategies. Hereafter, when we use IC requirements without specifically mentioned, we mean primary IC requirements. From the entries in the table, it is clear that the phase-II experiments (indirectly) limit the sender's strategy selection where this limitation arises due to the receiver's IC requirements for each partial strategy (when that partial strategy is used).

Partial strategy	Primary IC requirement	Partial strategy	Primary IC requirement
α_A	$p_1 \geq \frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2$	α_B	$1-p_1 \geq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}(1-p_2)$
β_A	$p_1 \geq \frac{1-p}{p}\frac{1-q_{A2}}{1-q_{A1}}p_2$	β_B	$1-p_1 \geq \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}}(1-p_2)$
γ_A	$p_1 \leq \frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2$	γ_B	$1-p_1 \leq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}(1-p_2)$

Table 3.1: IC requirements on the sender's signalling strategy based on the persuasion strategy in phase-II experiments

With this in mind, the sender's experiment design in phase I, essentially, is to select between different combinations of these partial strategies such that each partial strategy satisfies the constraint listed in Table 3.1. Hence, we next seek to understand

[9]Strictly speaking, for α_A and α_B partial strategies, another (secondary) IC constraint needs to be discussed - $p_1 \leq \frac{1-p}{p}\frac{1-q_{A2}}{1-q_{A1}}p_2$, $1-p_1 \leq \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}}(1-p_2)$ should be satisfied, respectively. This is to preclude the receiver's deviation to a β_X strategy. However, a deviation from a α_X to a β_X strategy benefits the sender. Hence, when the interim belief $\mathbb{P}(\theta_1|E_X)$ gives a profitable deviation from α_X to β_X to the receiver, the sender will directly use a β_X partial strategy instead of using α_X and anticipating the receiver's deviation.

how these IC constraints collectively determine the sender's strategy selection. To answer this, we first discuss the relationship between partial strategies, IC requirements, and the sender's expected utility.

From the sender's perspective, each partial strategy and its corresponding signals provide a path to persuade (or dissuade) the receiver to take action a_1. Since the sender's objective is to maximize the probability that action a_1 is taken, she would like to use the "most efficient"[10] partial strategy pair to persuade the receiver when the prior p falls in the region such that the optimal signaling strategy is non-trivial[11]. To better understand the "efficiency" of using partial strategies, we need to evaluate each partial strategy under a given phase-II experiment E_X.

- α_X strategy: To persuade receiver to take action a_1 via this partial strategy, the sender needs satisfy the IC requirement, i.e., to ensure that $\mathbb{P}(\theta_1|\omega_{XP}) \geq \frac{1}{2}$. Hence, the interim belief $\mathbb{P}(\theta_1|E_X)$ must satisfy $\mathbb{P}(\theta_1|E_X) \in [\frac{q_{X2}}{q_{X1}+q_{X2}}, \frac{1-q_{X2}}{2-q_{X1}-q_{X2}}]$, otherwise a commitment using α_X partial strategy will never be incentive-compatible. From the sender's perspective, the most efficient strategy to persuade the receiver using α_X partial strategy is to design the phase-I experiment such that $\mathbb{P}(\theta_1|E_X) = \frac{q_{X2}}{q_{X1}+q_{X2}}$. At this interim belief, the sender experiences a relative expected utility[12] $2q_{X1}$. When $\mathbb{P}(\theta_1|E_X) \in (\frac{q_{X2}}{q_{X1}+q_{X2}}, \frac{1-q_{X2}}{2-q_{X1}-q_{X2}})$, the sender's marginal expected utility when the interim belief increases is $q_{X1} - q_{X2}$.

- β_X strategy: To persuade receiver to take action a_1 with this partial strategy the sender needs to ensure that both inequalities $\mathbb{P}(\theta_1|\omega_{XP}) \geq \frac{1}{2}$ and $\mathbb{P}(\theta_1|\omega_{XF}) \geq \frac{1}{2}$

[10]Mathematically, the efficiency of a strategy is defined as the ratio

$$\frac{\mathbb{P}(a_1|\text{interim belief, the partial strategy used })}{\mathbb{P}(\theta_1|\text{interim belief, the partial strategy used})}.$$

[11]It is well known that Bayesian persuasion problems and the study of optimal signaling strategies are more interesting/critical in a prior region where the sender cannot persuade the receiver to take sender's most preferred action (a_1) with probability one. For example, in the seminal prosecutor-judge example *Kamenica and Gentzkow* (2011) this holds when $p < \frac{1}{2}$.

[12]The expected utility relative to the prior.

hold. Given the assumption in Claim 5, namely $q_{X1} \geq q_{X2}$, the only constraint that can be tight is $\mathbb{P}(\theta_1|\omega_{XF}) \geq \frac{1}{2}$. Hence, IC commitments using β_X partial strategy exist only when the interim belief $\mathbb{P}(\theta_1|E_X) \geq \frac{1-q_{X2}}{2-q_{X1}-q_{X2}}$. From the sender's perspective, the most efficient strategy to persuade the receiver using β_X partial strategy is to design the phase-I experiment such that $\mathbb{P}(\theta_1|E_X) = \frac{1-q_{X2}}{2-q_{X1}-q_{X2}}$ with the resulting relative expected utility $1+\frac{1-q_{X1}}{1-q_{X2}}$. Unlike an α_X partial strategy where the sender still gets a positive utility gain when the interim belief increases, for a β_X partial strategy the sender's marginal expected utility gain when the interim belief increases is 0 when $\mathbb{P}(\theta_1|E_X) > \frac{1-q_{X2}}{2-q_{X1}-q_{X2}}$.

- γ_X strategy: Given that the sender suggests the receiver to take action a_2 in this strategy, the sender's expected utility is 0 when using this partial strategy.

According to the discussion above, it is clear that the sender will not use the set of strategies corresponding to (γ_A, γ_B) unless the prior $p = 0$[13]. Additionally, we also know that different partial strategies provide different relative expected utility to the sender. When partial strategies are used in the most efficient manner, the relative expected utility under α_X partial strategy is at most $2q_{X1}$, and the average expected utility under β_X partial strategy is at most $1+\frac{1-q_{X1}}{1-q_{X2}}$.

Since these two values capture the best scenario that the sender can achieve by tailoring the interim belief under the given experiment, we define this pair of ratios, $(2q_{X1}, 1+\frac{1-q_{X1}}{1-q_{X2}})$ as a function of (the given) experiment, denoted by $PerP(E_X)$; henceforth, this pair is called the *persuasion potential.*

Definition 23. *Given an experiment* $E_X = (q_{X1}, q_{X2})$*, the persuasion potential of this experiment,* $PerP(E_X)$*, is the pair* $\left(2q_{X1}, 1+\frac{1-q_{X1}}{1-q_{X2}}\right)$.

To provide some insights on the importance and use of the persuasion potential we preview Corollary 4. Corollary 4 states that the sender only uses partial strategies in

[13]When prior $p = 0$, no signaling strategy can increase the sender's expected utility since the receiver knows the true state is always θ_2.

the most efficient manner, i.e., $\mathbb{P}(\theta_1|E_X) = \frac{q_{X2}}{q_{X1}+q_{X2}}$ when a partial strategy α_X is used and $\mathbb{P}(\theta_1|E_X) = \frac{1-q_{X2}}{2-q_{X1}-q_{X2}}$ when a partial strategy β_X is used. Thus, the persuasion potential can simplify the sender's search for the optimal signaling strategy. When a particular partial strategy is used in the most efficient manner described in the above parameter, the interim belief is now determined. Therefore, the sender does not need to search for the optimal signaling strategy from the whole set of IC strategies but only needs to search from a small number of strategies that generate the particular interim beliefs.

3.3.4 Persuasion ratio and the optimal signaling structure

Since the sender wants to maximize the total probability of action a_1, she needs to compare different sets of strategies formed by different pairs of partial strategies. To compare[14] each set of strategies, we introduce the persuasion ratio of a set of strategies for a given value of the prior.

Definition 24. *Given a set of incentive-compatible strategies $\mathcal{S}$, e.g., $\mathcal{S} = (c_A, d_B)$ with $c, d \in \{\alpha, \beta, \gamma\}$ which satisfying IC requirements, the persuasion ratio (PR) of the set of strategies $\mathcal{S}$ is the maximum total probability of action a_1 is taken (under a strategy within the set) divided by the prior p, $PR(\mathcal{S}, p) = \max_{S \in \mathcal{S}} \frac{\mathbb{P}(a_1|S)}{p}$.*

Careful readers may notice that if we multiply the persuasion ratio with the prior, the value will be the (maximum) expected utility the sender can achieve from the given set of strategies. Since the sender's expected utility will be monotone increasing in the prior p in this framework regardless of which set of strategies the sender adopts, the persuasion ratio for a given prior can be viewed as the relative utility gain this set of strategies can offer to the sender under this prior. Therefore, given a specific prior, if a set of strategies has a higher persuasion ratio with respect to another set of

[14] Instead of a brute-force search to identify the optimal signaling strategy, we develop structural properties which will be used for our generalization.

strategies, the sender should use the strategies in the former set instead of the latter set.

According to this discussion, we can draw a persuasion ratio curve for each set of strategies as the prior is varied in $[0, 1]$. Abusing notation, we represent the persuasion ratio curve by $PR(\mathcal{S})$. Optimization may need to carry out for each value of the prior. However, structural insights presented in the following two lemmas considerably simplify the analysis. Properties presented in Lemma 14 narrow down the space where the sender needs to search for the optimal signaling strategy. This allows us to depict persuasion ratio curves $PR(S)$ for some basic strategies. On top of that, Lemma 15 provides a systematic approach to derive persuasion ratio curves for all types of strategies.

Lemma 14. *Given a type of strategy $\mathcal{S}$ and a prior p, there exists an optimal strategy $S \in \mathcal{S}$, i.e., the sender's expected utility with S equals the product of the persuasion ratio of $\mathcal{S}$ and the prior p, which satisfies one of the following two conditions:*

- *At least one IC requirements of the constituent partial strategies is tight under S; or*
- *There is a signal that will be sent with probability 1 under S.*

Before discussing the simplifications that Lemma 14 yields in terms of the key properties for solving the problem, we give an intuitive outline of the proof of Lemma 14. When the IC requirements of the two partial strategies are not tight and both signals are sent with non-zero probability in a strategy S, the sender can increase her expected utility by slightly raising the probability of the signal with a higher persuasion ratio (and adjust the probability of the other signal to respect the prior) to form a strategy S_+. The sender can keep doing this 'slight' modification of her strategies until either one of the IC conditions is satisfied or the signal is sent w.p. 1.

Given this lemma, the persuasion ratio curve of the following types of strategies: (α_A, γ_B), (β_A, γ_B), (γ_A, α_B), (γ_A, β_B) can be determined immediately since the IC requirement can never be tight for the γ class partial strategy. For the remaining four types of strategies: (c_A, d_B), $c, d \in \{\alpha, \beta\}$, the following lemma aids in solving for the strategy meeting the persuasion ratio without a point-wise calculation. As a preview of the result of Lemma 16, once we derive the persuasion ratio curve for each type of strategy via Lemma 15, we can immediately identify the optimal signaling strategy by overlaying those curves in one figure.

Lemma 15. *Given the persuasion ratio curves of the types of strategies (c_A, γ_B) and (γ_A, d_B), denoted by $PR((c_A, \gamma_B), p)$ and $PR((\gamma_A, d_B), p)$, respectively, the persuasion ratio curve of the set of incentive compatible strategies (c_A, d_B), denoted by $PR((c_A, d_B), p)$, is the generalized concave hull* [15] *of functions $PR((c_A, \gamma_B), p)$ and $PR((\gamma_A, d_B), p)$:*

$$PR((c_A, d_B), p) = \max_{\substack{x,u,v \in [0,1], \\ xu+(1-x)v=p}} xPR((c_A, \gamma_B), u) + (1-x)PR((\gamma_A, d_B), v)), \tag{3.3}$$

when $\mathcal{S}_{(c_A, d_B)}(p) \neq \emptyset$, where $\mathcal{S}_{(c_A, d_B)}(p)$ is the set of incentive-compatible (c_A, d_B) strategies at p.

The proof of Lemma 15 uses the structure of the sender's expected utility when no γ partial strategy is used in the types of strategies employed. The sender's expected utility function under (c_A, d_B) at prior p can be represented as a linear combination of her utility function under (c_A, γ_B) at prior u and her utility function under (γ_A, d_B) at prior v, then we prove that the optimization problem for solving for the optimal phase-I experiment parameters (p_1, p_2) under a prior p can be transformed to the maximization problem in the statement of Lemma 15. As we have discussed the

[15] If $PR((c_A, \gamma_B), \cdot)$ and $PR((\gamma_A, d_B), \cdot)$ were the same function $f(\cdot)$, then this would be its convex hull.

means to determine each type of strategy's persuasion ratio curve, the persuasion ratio curve of the optimal signaling strategy can be determined using the following Lemma 16.

Lemma 16. *The persuasion ratio curve of the optimal signaling strategy $PR^*(p)$ is the upper envelope of the different types of strategies' persuasion curves. Further, the optimal signaling strategy (under a given prior) is the strategy that reaches the frontier of the persuasion ratio curve (at that prior).*

Since a higher persuasion ratio indicates a higher (sender's) expected utility for every given prior, it is straightforward that the sender will choose the upper envelope of the persuasion curves of the different types of strategies. Because the the set $\mathcal{S}_{(c_A,d_B)}(p)$ can be empty for some prior values with a corresponding persuasion ratio $PR(c_A, d_B), p) = 0$, the main effort in proving Lemma 16 is to show the existence of an incentive-compatible[16] commitment on the frontier of the persuasion ratio curve at every possible prior. Finally, once the persuasion ratio curve of the optimal signaling strategy is determined, we can immediately infer an optimal signaling strategy S^* under a specific prior: this is a strategy that has the ratio of total probability of action a_1 to the prior equalling the optimal persuasion ratio at that prior.

We emphasize that Lemmas 14, 15, and 16 provide a structural approach to solve for optimal signaling strategies in two-phase trials. In Section 3.4, we will leverage this structural approach to solve the more complex multi-phase trial problem with a dynamic programming algorithm. For a two-phase trial or to solve the last two phases of a trial with more than two phases, the following corollary that uses the persuasion potential, can further simplify the sender's optimization procedure.

Corollary 4. *Let $\Gamma^*(p)$ represents the optimal signaling strategy at prior p. If $\mathbb{P}(a_1|\Gamma^*(p)) < 1$, then the following two statements are true:*

[16] Again, from the receiver's perspective.

- *When a partial strategy α_X is used in $\Gamma^*(p)$, the interim belief is* $\mathbb{P}_{\Gamma^*(p)}(\theta_1|E_X) = \frac{q_{X2}}{q_{x1}+q_{X2}}$.

- *When a partial strategy β_X is used in the optimal signaling strategy $\Gamma^*(p)$, the interim belief is* $\mathbb{P}_{\Gamma^*(p)}(\theta_1|E_X) = \frac{1-q_{X2}}{2-q_{x1}-q_{X2}}$.

Given the result of Corollary 4, the Equation (3.3) in Lemma 15 reduces to a linear equation. Hence, the comparison in Lemma 16 and the computation in Lemma 15 can be reduced to a comparison of the (unique) corresponding IC strategies (if one exists) under the interim belief listed in Corollary4 for different types of strategies.

3.3.5 Comparison with the classical Bayesian persuasion

Given the optimal signaling strategy derived in Lemma 16, one natural follow-up question is the quantification of the sender's utility improvement obtained by adopting the optimal signaling strategy in comparison to using strategies structurally similar to the optimal strategies in binary state of the world. Using the concavification approach presented in *Kamenica and Gentzkow* (2011), the optimal strategy classical Bayesian persuasion in binary-state problems always mixes two possible states in one signal and reveals the true state on the other signal. Hence, we define a class of strategies structurally similar to the Bayesian persuasion optimal strategy of binary-state problems below.

Definition 25. *With binary states of the world, a (binary-state) Bayesian persuasion (BBP) strategy is a strategy that "mixes two possible states in one signal and reveals the true state on the other signal".*

Given the model defined in Section 3.2.1, a BBP strategy is forced to use at least one γ_X partial strategy[17]. Given a fixed type of strategy, e.g., (α_A, γ_B), an

[17] This is true because using either α_X, β_X strategies requires a mixture of two possible states in one signal.

optimal BBP strategy using this type of strategy can be solved by the concavification approach after calculating the sender's expected utility curve of interim beliefs. After solving the optimal BBP strategy of a given strategy type via the concavification approach respectively, the optimal BBP strategy is the strategy in the set of $\{(\alpha_A, \gamma_B), (\beta_A, \gamma_B), (\gamma_A, \alpha_B), (\gamma_A, \beta_B)\}$ which yields the highest expected utility for the sender. Figure 3.2 plots the sender's expected utility for the optimal signaling strategy and the optimal BBP strategy under a given pair of phase-II experiments: $(q_{A1}, q_{A2}) = (0.8, 0.2), (q_{B1}, q_{B2}) = (0.7, 0.3)$. The blue line in Figure 2 is the benchmark of the sender's maximum expected utility in a single-phase scenario, which would be the optimal performance if one of the phase-II trials were a trivial experiment. As we can see, the sender's expected utility is lowered owing to the determined phase-II experiments. For low-priors, the optimal signaling strategy derived in Section 3.3.4 and the optimal BBP strategy give the sender the same expected utility[18]. However, as the prior increases, a utility gap between the optimal signaling strategy derived in Section 3.3.4 and the optimal BBP strategy appears and then increases until the receiver will take a_1 with probability one. The utility gap starts when the optimal signaling strategy uses strategies (α_A, β_B) or (α_B, β_A) which are not considered in BBP strategies.

3.4 Binary-outcome Experiments in Multi-phase Trials

In this section, we generalize the approach we derived in Section 3.3 to study problems with multi-phase trials. In section 3.2.1, we built a model for binary-outcome experiments in two-phase trials. Here, we generalize the model to multi-phase trials and then propose a dynamic programming algorithm to solve for the optimal signaling strategy. The state for the dynamic program will be the interim belief on the state of

[18]This holds because the optimal signaling strategy in the low-prior region for this example is (α_A, γ_B).

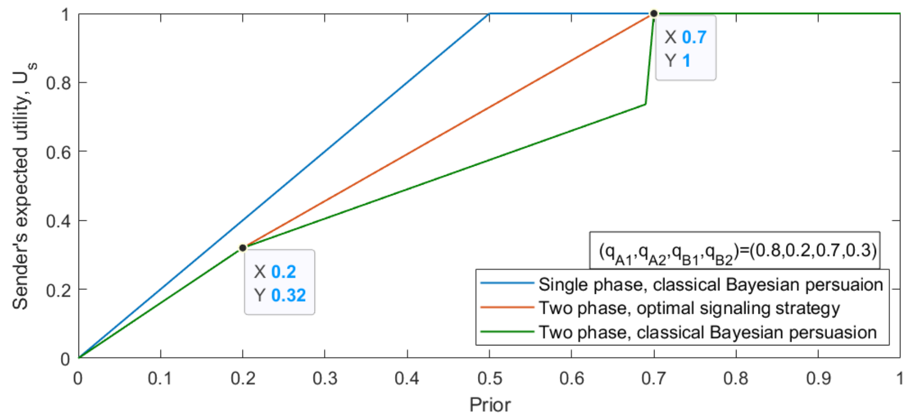

Figure 3.2: Sender's utility under different problem settings and strategies

the world that results at any node in the extensive-form delineation of the problem. As the belief at each level is determined based on the actions in earlier stages (if any), in the backward iteration procedure, we will determine the optimal choice of experiments by the sender (if there is a choice) for any possible interim belief. In this dynamic programming, there is only a terminal reward that arises from the receiver's action based on the outcome of the final trial and based on the receiver's resulting posterior belief on the state of the world.

3.4.1 Model of binary-outcome experiments in multi-phase trials

There are N phases in a trial where one binary-outcome experiment will be conducted in each phase. However, as in the two-phase-trial settings, the specific experiment conducted in each phase is determined by the earlier phases' outcomes. Therefore, for an N-phase trial, we can draw a height-N binary tree where each leaf node represents an outcome revealed to the receiver, and each non-leaf node represents an experiment. With this binary tree, experiment $E_{i,j}$ represents the j^{th} experiment to be conducted at level i. When j is odd, an experiment $E_{i,j}$ will be conducted only if the experiment $E_{i-1,(j+1)/2}$ is conducted and passed. Similarly, when j is even, an experiment $E_{i,j}$ will be conducted only if the experiment $E_{i-1,j/2}$ is conducted but it

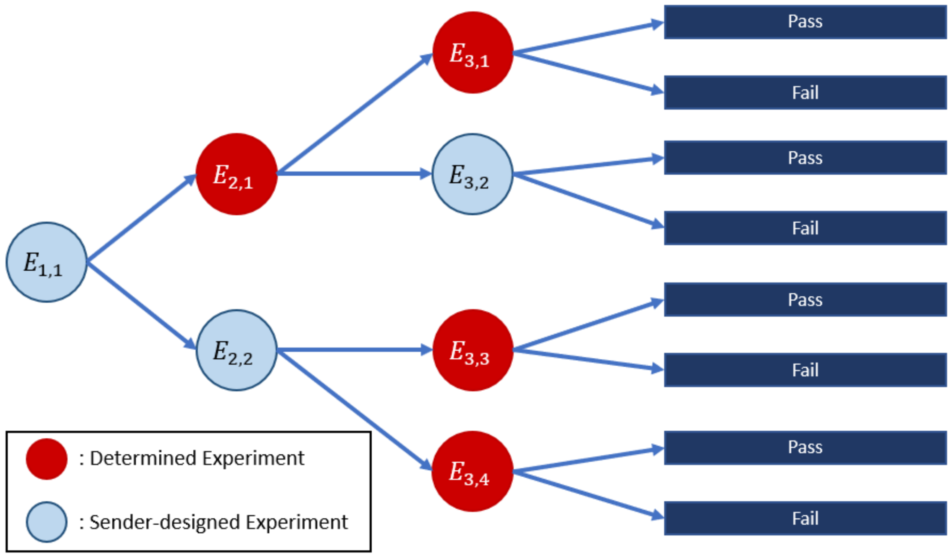

Figure 3.3: A 3-phase trial example

fails. In this binary tree, some experiments, e.g, $E_{i,j}$, are determined with parameters (e.g., (q_{ij1}, q_{ij2})) known to both the sender and the receiver. However, some experiments, e.g., $E_{k,l}$, can be designed by the sender; all the parameters are chosen before any experiment is conducted and are common knowledge. In such experiments, the sender can choose a probability pair, e.g., $(p_{kl1}, p_{kl2}) \in [0, 1]^2$. In contrast to the model defined in Section 3.2.1, here determined experiments and sender-designed experiments can be at any level of the binary tree with the placement arbitrary but carried out before the sender chooses her signals. In other words, unlike the model we considered in Section 3.2.1, a sender may be able to design an experiment at phase N (final level before the receiver takes a decision) owing to a determined experiment outcome at phase $N-1$. Fig. 3.3 provides an example of a three-phase trial.

In this model, both the sender and the receiver are assumed to know the prior, the experiments that are determined (their location on the tree and their parameters), and the sender-designed experiments (their location on tree). The sender has to design all the experiments she has the flexibility to choose simultaneously and before the game starts (when the state is realized); the designed experiments' parameters are then revealed to the receiver (again before the game starts). Given the experiments designed by the sender and the realized outcome of a sequence of experiments, the receiver will take an action to guess the true state of the world. For simplicity of analysis, we keep the sender and the receiver's utilities the same as in Section 3.2, i.e., $U_R(a_i, \theta_i) = 1$ and $U_R(a_i, \theta_j) = 0$ for all $i, j \in \{1, 2\}, i \neq j$; and $U_S(a_1, \theta_i) = 1$ and $U_S(a_2, \theta_i) = 0$ for all $i \in \{1, 2\}$. Then, the sender's objective is to jointly design the set of experiments that she has the flexibility to choose to maximize her expected utility, which is nothing but the probability of the receiver taking action a_1. Before proceeding, we point out that this model can be easily generalized to unbalanced binary trees straightforwardly by adding dummy nodes with determined trivial experiments defined in Section 3.3.1 to construct an equivalent balanced binary tree.

3.4.2 Determined versus sender-designed experiments

Given the model, the sender can manipulate the phase-K interim belief only when designing an experiment at phase $K-1$. If an experiment at phase $K-1$ is determined, then the phase-K interim belief given the experiment's outcome in phase $K-1$ is just a function of phase-$K-1$ interim belief. Therefore, figuring out how these two types of experiments, determined and sender-designed experiments, will influence every given phase's interim belief is the key to solving for the optimal signaling strategy. We start by noting that if the posterior belief at a leaf node is given, then the receiver's action is determined - he will take the action with the highest posterior probability unless there is a tie, in which case he is indifferent and will follow the sender's recommendation. Therefore, we can use backward iteration and the principle of optimality to determine the optimal signaling. We start by considering the last[19] phase's experiments when the sender can design them.

[19] They are phase-N experiments at level $N-1$ of the binary tree.

3.4.2.1 Experiments at phase N

Recall the result we have discussed in Section 3.3.2, a determined experiment in the last phase (phase-II in Section 3.3.2) limits the sender's strategy choice to one of three partial strategies. Besides, the best scenario that the sender can achieve via using these partial strategies (without violating the IC requirement) is captured by the persuasion potential of the determined experiment. However, when there is a sender-designed experiment at phase N and the interim belief[20] $\tilde{p} \leq \frac{1}{2}$, the sender can always design the experiment to make two states equally likely when this experiment passes and reveal the sender-disliked state[21] when this experiment fails. If we cast this sender-designed experiment in terms of a determined experiment, the sender-designed experiment will have a persuasion potential $(2, 2)$[22]. Thus, no matter the type of experiment at phase N, we can capture the sender's optimal set of partial strategies via a persuasion potential.

3.4.2.2 Experiments in phase $N-1$

In the second-last phase, results in Section 3.3.4 describe a sender-designed experiment's role in the optimal signaling strategy: pick the strategy on the frontier of all persuasion-ratio curves. However, if the experiment is determined in the second-last phase, an additional constraint on the interim belief between the second-last phase and the last phase is enforced. That is to say, the set of (feasible) strategies that we can use to search for the persuasion ratio shrinks. For example, if we want to look at the persuasion ratio of a set of strategies (c_A, d_B) in a two phase trial with a deter-

[20] If the interim belief $\tilde{p} > \frac{1}{2}$, then the sender can design a trivial experiment $E_X = (q_{X1}, q_{X2})$ where $q_{X1} = q_{X2}$, which provides no further information to the receiver and results in the posterior belief same as the interim belief of the experiment. Here, the sender can easily convince the receiver to take the sender-preferred action with probability 1.

[21] The sender and the receiver's preferred actions are opposite to each other at this state.

[22] When the sender gets to design an experiment $E_X = (q_{X1}, q_{X2})$, she can choose $q_{X1} = 1, q_{X2} > 0$ to make the first element of the persuasion potential to be 2. Similarly, she can choose $q_{X1} = q_{X2} \in [0, 1]$ to make the second element of the persuasion potential to be 2. (When $q_{X1} = q_{X2}$, then $\frac{1-q_{X1}}{1-q_{X2}} = 1$ if $q_{X1} = 1$ by L'Hôpital's rule.)

mined phase-I experiment $E_X = (q_{X1}, q_{X2})$, the persuasion ratio of a given strategy becomes a function of parameters in (E_X, E_A, E_B) (when that strategy satisfies the IC requirements):

$$PR((c_A, d_B), p) = \frac{\mathbb{P}_{E_A,E_B}\Big(a_1\Big|(c_A, d_B), \big\{\mathbb{P}(\theta_1|\omega),\ \omega \in \{\omega_{AP}, \omega_{AF}, \omega_{BP}, \omega_{BF}\}\big\}\Big)}{p},$$

where $\mathbb{P}(\theta_1|\omega_{AP}) = \frac{pq_{X1}q_{A1}}{pq_{X1}q_{A1}+(1-p)q_{X2}q_{A2}}$, $\mathbb{P}(\theta_1|\omega_{BP}) = \frac{p(1-q_{X1})q_{B1}}{p(1-q_{X1})q_{B1}+(1-p)(1-q_{X2})q_{B2}}$,

$\mathbb{P}(\theta_1|\omega_{AF}) = \frac{pq_{X1}(1-q_{A1})}{pq_{X1}(1-q_{A1})+(1-p)q_{X2}(1-q_{A2})}$,

$\mathbb{P}(\theta_1|\omega_{BF}) = \frac{p(1-q_{X1})(1-q_{B1})}{p(1-q_{X1})(1-q_{B1})+(1-p)(1-q_{X2})(1-q_{B2})}$.

Fortunately, after enforcing the constraints, the process of searching for the optimal signaling strategy under a determined experiment is the same as the sender-designed experiment, i.e., pick the strategy in the frontier of all persuasion ratio curves. Therefore, at each possible branch of phase $N-1$, we can plot an optimal persuasion ratio curve capturing the sender's optimal signaling strategy at phase $N-1$ and phase N.

3.4.2.3 Experiments in earlier phases

Now we consider experiments in earlier phases. Let's start with a determined experiment. Given a determined experiment in phase-K, e.g., $E_{k,i} = (q_{Ki1}, q_{Ki2})$, if we have already solved the optimal persuasion ratio curves of its succeeding phase-$(K+1)$ experiments $E_{K+1,2i-1}$ and $E_{K+1,2i}$, i.e., $PR^*_{K+1,2i-1}(p)$ and $PR^*_{K+1,2i}(p)$ are known, then the optimal persuasion ratio curve at this determined phase-K experiment $E_{k,i}$ is just a linear combination of $PR^*_{K+1,2i-1}(p)$ and $PR^*_{K+1,2i}(p)$ can be written as follows:

$$\begin{aligned} PR^*_{K,i}(p) &= (pq_{Ki1} + (1-p)q_{Ki2})PR^*_{K+1,2i-1}\Big(\frac{pq_{Ki1}}{pq_{Ki1} + (1-p)q_{Ki2}}\Big) + (p(1-q_{Ki1}) \\ &+ (1-p)(1-q_{Ki2}))PR^*_{K+1,2i}\Big(\frac{p(1-q_{Ki1})}{p(1-q_{Ki1}) + (1-p)(1-q_{Ki2})}\Big) \end{aligned} \tag{3.4}$$

For a sender-designed experiment at phase K, e.g., $E_{K,j}$, if we have already solved the optimal persuasion ratio curves of its succeeding phase-$(K + 1)$ experiments $E_{K+1,2j-1}$ and $E_{K+1,2j}$, the sender's best design at $E_{K,j}$ is to find a linear combination of $PR^*_{K+1,2j-1}(p)$ and $PR^*_{K+1,2j}(p)$ which yield the highest persuasion ratio for every phase-K interim belief p. Since the persuasion ratio curve is monotone decreasing in the belief, the optimal persuasion ratio curve can be constructed similar to Lemma 15 as shown in the following claim.

Claim 7. *Given two persuasion ratio curves at phase $K + 1$, $PR^*_{K+1,2j-1}(p)$ and $PR^*_{K+1,2j}(p)$, the optimized persuasion ratio curve $PR^*_{K,j}(p)$ at phase K is the maximum convex combination of $PR^*_{K+1,2j-1}(p)$ and $PR^*_{K+1,2j}(p)$, i.e.,*

$$PR^*_{K,j}(p) = \max_{x,u,v\in[0,1],xu+(1-x)v=p} xPR^*_{K+1,2j-1}(u) + (1-x)PR^*_{K+1,2j}(v) \tag{3.5}$$

3.4.3 Dynamic programming approach for the optimal signaling strategy

Given the analysis above, we can derive the optimal signaling strategy using a dynamic programming approach outlined in Algorithm 1.

ALGORITHM 1: Dynamic programming approach for multi-phase trials

Input: The set of determined experiments $\mathbf{E_D}$, the binary tree structure
Output: The optimal persuasion ratio curve
1. For each experiment at phase N, $E_{N,i}$ $i \in \{1, ..., 2^N\}$, solve its persuasion potential $Prep(E_{N,i})$
2. For each experiment at phase $N-1$, $E_{N-1,i}$ $i \in \{1, ..., 2^{N-1}\}$, find the optimal persuasion ratio curve using $(Prep(E_{N,2i-1}), Prep(E_{N,2i}))$.
3. K=N-2
4. while $K > 0$ **do**
    For each experiment at phase K, $E_{K,i}$ $i \in \{1, ..., 2^K\}$, find the optimal persuasion ratio curve using equation (3.4) or Claim 7
    K=K-1
end
5. Return the optimal persuasion ratio curve at phase 1

3.4.4 Multi-phase model and classical Bayesian persuasion

At the end of this section, we mention a class of special multi-phase trials where the sender's expected utility under the optimal signaling strategy is equivalent to utility obtained from a single-phase Bayesian persuasion model. Inspired by the two-phase example with a trivial experiment in 13, the sender can implement a signaling strategy similar to single-phase Bayesian persuasion when there exists a trivial experiment in the last phase and she can design experiments in earlier phases. When the sender can design all earlier phases, she can voluntarily reduce the signal space in effect via designing the experiment $E_{i,j}$ to be $E_{i,j} = (1, 1)$ or $E_{i,j} = (0, 0)$, i.e., a revealing experiment. By doing this, the sender can reduce the multi-phase trial to an equivalent two-phase trial model and then a straightforward extension of Lemma 13 will hold when there exists a trivial experiment in the last phase. Figure 3.4 depicts this extension in a 5-phase example. A lemma that further generalizes the class of multi-phase models where the sender has the same expected utility as single-phase Bayesian persuasion problem is presented below after the definition of a necessary pruning process.

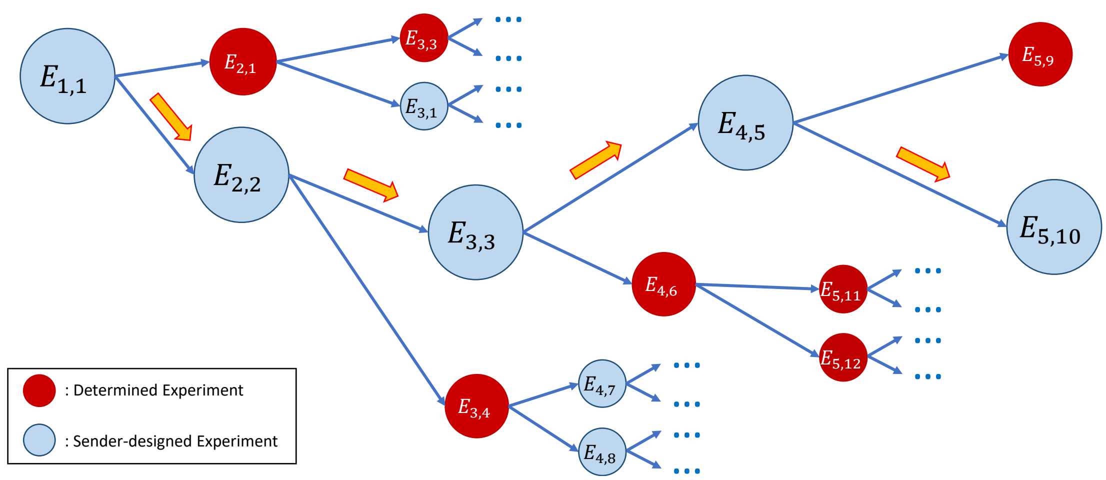

Figure 3.4: A 5-phase example with sender's expected utility equivalent to classical Bayesian persuasion

Definition 26. *Given an N-phase trial model M, a pruned N-phase trial model $Prun(M)$ is a subset of M which starting from the leaf nodes recursively replaces every subtree satisfying the following condition by a revealing experiment $E_\theta = (1, 0)$:*

- *The node has a trivial (determined) experiment E_X with at least one of its succeeding experiments a non-trivial (determined) experiment.*

Note that the pruned tree will potentially be unbalanced.

Lemma 17. *Given an N-phase trial M with binary-outcome experiments, if there exists a pruned N-phase trial model $Prun(M)$ such that the following two conditions hold, then the sender's expected utility is given by an equivalent single-phase Bayesian persuasion model.*

1. *For every non-trivial determined experiment, its sibling is either a trivial or a sender-designed experiment.*

2. *There exists a least one sender-designed experiment in each experiment sequence*[23] *of $Prun(M)$.*

We provide some intuition for Lemma 17. The first condition guarantees that the sender's signaling space is (effectively) not restricted in this pruned model. The second condition states that the sender can design posterior beliefs in every possible outcome in this pruned model. Generalizing Lemma 17 to a complete characterization when a multi-phase trial is equivalent to a single-phase Bayesian persuasion model is for future work.

3.4.4.1 Robustness of signaling strategies under small perturbation

Given the statement of Lemma 17, we are curious about the robustness of sender's optimal signaling strategy under small perturbations. Here, we want to discuss a 4-phase example which initially satisfies the conditions of Lemma 17 and see how small

[23] From root to leaf experiments.

perturbations on trivial experiments will affect the sender's expected utility and the optimal signaling strategy.

To begin with, we consider a 4-phase example depicted in Figure 3.5, where S^i denotes the i^{th} sender-designed experiment, N^i denotes the i^{th} non-trivial experiment, and T^i denotes the i^{th} trivial experiment. The phase-I experiment is always S^1, then depends on the outcome of S^1, the phase-II experiment can be either T^1 or N^1 as depicted in the figure. To avoid ambiguity, for each experiment E^i, $q_1^{E^i}, q_2^{E^i}$ represent the probability of going to the left node when the true state is θ_1, θ_2, respectively. In this example, the prior belief of θ_1 is set to be $p = 0.4$, and all parameters of pre-determined experiments (trivial or non trivial) are provided in Table 3.2.

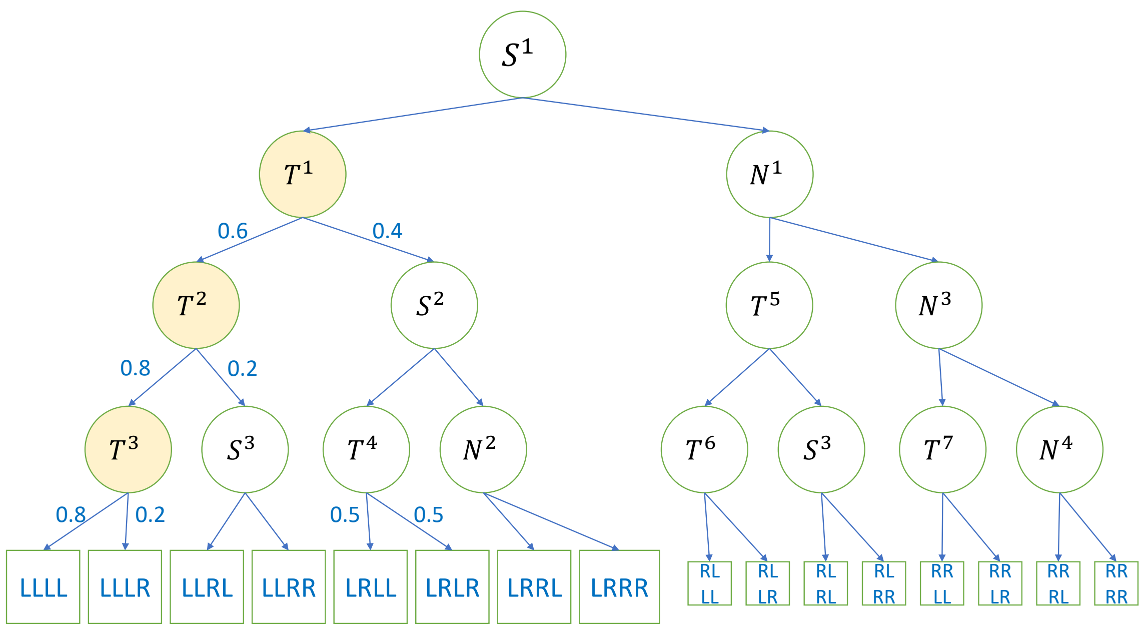

Figure 3.5: A 4-Phase Trial Satisfying Lemma 17

Table 3.2: Table of experiment parameters

Experiment	T^1	T^2	T^3	T^4	T^5	T^6	T^7	N^1	N^2	N^3	N^4
q_1	0.6	0.8	0.8	0.5	0.7	0.7	0.7	0.8	0.6	1	0.5
q_2	0.4	0.2	0.2	0.5	0.7	0.7	0.7	0.5	0.4	0.5	0.5

When there is no perturbation, the sender's optimal signaling strategy will design

S^1 to be $(q_1^{S^1}, q_2^{S^2}) = (1, \frac{2}{3})$ and the sender's expected utility is 0.8. Now, we want to investigate small perturbations in T^2 and T^3. Since every trivial experiment T^i has $q_1^{T^i} = q_2^{T^i}$, we assume that the perturbation always occurs in $q_2^{T^i}$ to avoid redundancy.

Example 1: Perturbation in T_3 When there is a small perturbation of $q_2^{T^3}$, the posterior of the outcome $LLLL$ and $LLLR$ will now be different. In the original signaling strategy, the sender will make $\mathbb{P}(\theta_1|LLLL) = 0.5$ and $\mathbb{P}(\theta_1|LLLR) = 0.5$ to persuade the receiver taking action a_1. Under the perturbation, one of the outcome will have posterior belief less than 0.5 and the receiver's best response will switch to a_2. Thus, the sender will have to decide whether she will accommodate this perturbation and design a new experiment at the very beginning (S^1) or using the same signaling strategy and reluctantly accept the loss. This utility change towards the perturbation is plotted in Figure 3.6. We want to note that other examples with perturbation in the phase phase or in a trivial experiment where none of its succeeding experiments is a sender-designed experiment will have the figures with the shape similar to this one.

Example 2: Perturbation in T_2 Instead of a perturbation in the last phase, we want to look at a perturbation occurs in a trivial experiment where one[24] of its succeeding experiment is a sender-designed experiment. In Figure 3.7, different from Figure 3.6, keeping the same strategy will never be an optimal when the perturbation increases the total probability on the side where its succeeding node is a sender-designed experiment. The reason is that when there is a sender-designed experiment in its succeeding, the sender can decide whether to accommodate the perturbation in earlier phases S^1 or in later phases S^3.

[24]If both of its succeeding experiments are sender-designed experiments, then the perturbation in this trivial experiment will not affect the sender's expected utility if it originally satisfies the conditions in Lemma 17. The reason is the sender can design these two sender-designed experiments and adjust the sender-designed experiment in an earlier phase than this perturbation to cancel out the effect of a small perturbation.

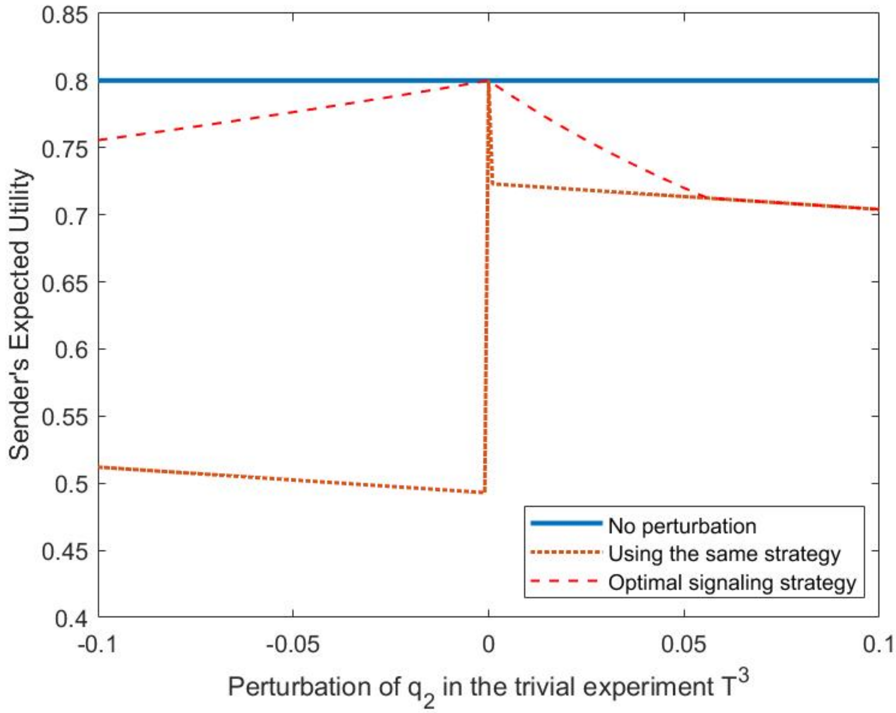

Figure 3.6: Perturbation at T^3

With above two small examples, we can see there's a significant utility drop even with a small perturbation. This suggests that even the difference between q_1^E and q_2^E is small, we should still view it as a non-trivial experiment instead of treating it as a trivial experiment.

3.5 Discussion

3.5.1 Non-binary-outcome experiments

In Section 3.4, we proposed a dynamic programming algorithm to solve the optimal signaling strategy in binary-outcome experiments. When the experiments have more than two possible outcomes, the dynamic programming approach still works if complexity is not an issue. Again, supposing that we can solve the last two phases in non-binary-outcome experiments directly via the LP, then the earlier phases can be solved by a dynamic programming approach similar to Section 3.4.

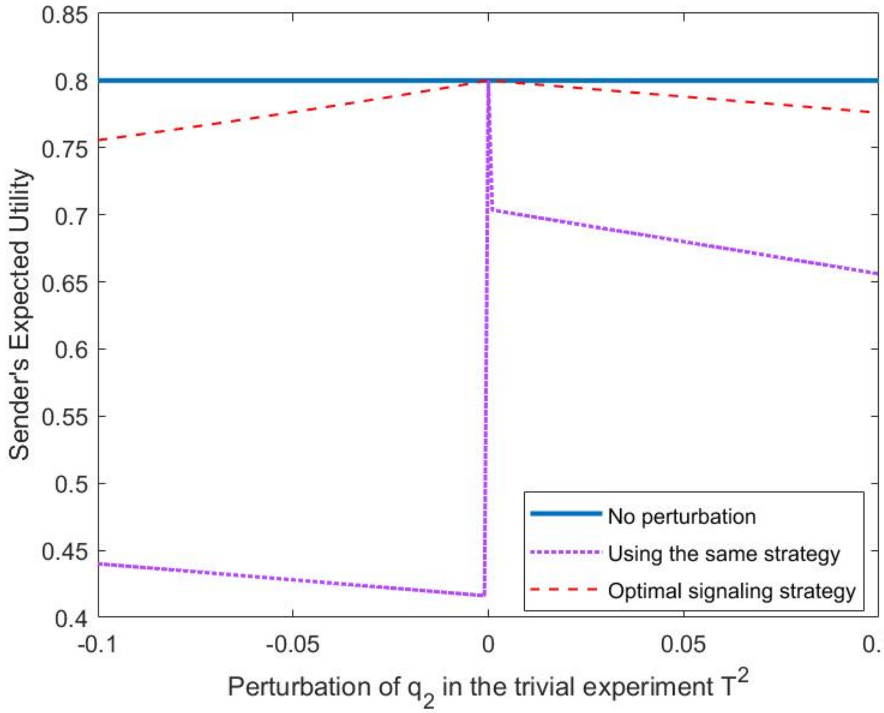

Figure 3.7: Perturbation at T^2

To detail the dynamic programming operation in each experiment not in the last two phases, we again treat determined experiments and sender-designed experiments differently. Let's start with a determined experiment. Given a determined experiment (assuming its index is i) in K^{th} phase with n possible outcomes, denoted as $E_{K,i} = (q_{Ki11}, ..., q_{Kin1}, q_{Ki12}, ..., q_{Kin2})$, if we have already solved the optimal persuasion ratio curves in its succeeding phase, i.e., phase $(K+1)$ with experiments $E^i_{K+1,1}$, $E^i_{K+1,2}$, ..., $E^i_{K+1,n}$, and the curves are represented by $PR^*_{K+1,(i,1)}(p)$, $PR^*_{K+1,(i,2)}(p)$, ..., $PR^*_{K+1,(i,n)}(p)$, then the optimal persuasion ratio curve at this determined phase-K experiment $E_{K,i}$ is a linear combination of the persuasion curves $PR^*_{K+1,(i,1)}(p)$, $PR^*_{K+1,(i,2)}(p)$, ..., $PR^*_{K+1,(i,n)}(p)$. It can be represented in the following form:

$$PR^*_{K,i}(p) = \sum_{j=1}^{n}(pq_{Kij1} + (1-p)q_{Kij2})PR^*_{K+1,(i,j)}\Big(\frac{pq_{Kij1}}{pq_{Kij1} + (1-p)q_{Kij2}}\Big). \qquad (3.6)$$

When we have an n-outcome experiment $E_{K,m}$ at phase K that the sender can design, the sender's optimal persuasion ratio curve is the generalized concave hull of all of its succeeding phase-$(K+1)$ experiments' persuasion ratio curves $PR^*_{K+1,(m,1)}(p)$, $PR^*_{K+1,(m,2)}(p)$, ..., $PR^*_{K+1,(m,n)}(p)$. This convex optimization function can be written in the following form:

$$\begin{aligned} PR^*_{K,m}(p) &= \max_{u_1,\ldots u_n,\ v_1,\ldots v_n} \sum_{j=1}^{n} u_k PR^*_{K+1,(m,j)}(v_j) \\ \text{subject to} &\qquad \sum_{j=1}^{n} u_j = 1,\ \sum_{j=1}^{n} u_j v_j = p \end{aligned} \tag{3.7}$$

Given the convex optimization problem in (3.7), we know that for each sender-designed experiment, we have to solve a convex optimization problem that has $2n-2$ parameters. At this point, we do not know whether there is a simplification to solve this convex optimization directly in an efficient manner. As an alternative, we will propose a transformation to transform any n-outcome experiments into a sequence of binary-outcome experiments.

Lemma 18. *Given a n-outcome experiment* $E = [(q_{11}, q_{21}), (q_{12}, q_{22}), ..., (q_{1n}, q_{2n})]$ *satisfying* $\sum_{i=1}^{n} q_{1i} = 1$ *and* $\sum_{i=1}^{n} q_{2i} = 1$*, we can replace this experiment* E *by at most* $\lceil log_2 n \rceil$ *phases of binary-outcome experiments.*

Lemma 18 can be proved by constructing a level-$\lceil log_2 n \rceil$ tree with binary-outcome experiments such that the state θ_1 has probability q_{1i} to be in the i^{th} leaf node and the state θ_2 has probability q_{2i} to be in the same node. Let's use a ternary-outcome experiment as an example. With the tuple (Input,Outcomes) representing experiments, consider an experiment $E = (I_1, \{O_1, O_2, O_3\})$ with a given $[(q_{11}, q_{21}), (q_{12}, q_{22}), (q_{13}, q_{23})]$. This experiment E is equivalent to a sequence of 3 binary-outcome experiments $E_1 = (I_1, \{I_2, I_3\})$, $E_2 = (I_2, \{O_1, O_2\})$, and $E_3 = (I_3, \{O_3, O_3\})$. Every belief evolution $p \to (p_1, p_2, p_3)$ generated via E can be reproduced using E_1, E_2, and E_3. Hence, we can set $E_1 = (q_{11} + q_{12}, q_{21} + q_{22})$, $E_2 = (\frac{q_{11}}{q_{11}+q_{12}}, \frac{q_{21}}{q_{21}+q_{22}})$, and $E_3 = (1, 1)$ replace E by a set of 2-phases binary-outcome experiments. The explicit construction of the general formulation can be found in the Appendix C.2.6.

With Lemma 18, for every given non-binary experiment, predetermined or sender-designed, we can produce a set of equivalent binary-outcome experiments. Thus, our DP algorithm (Algorithm 1) can also be applied to solve multi-phase trials with non-binary (outcome) experiments with an increased number of phases.

3.5.2 Receiver's optimal experiments

In many real-world applications, the receiver can design/determine some of the experiments ahead of time for all future senders. For example, the FDA can determine regulations (the experiments) ahead of time for future trials. In such scenarios, the receiver aims to design the experiments to maximize her expected utility by circumscribing the sender's ability to tailor the interim beliefs. However, while determining the experiments, we assume that the receiver has no information about the prior in each future case or has minimal information about the prior in future cases. In such scenarios, by viewing the future prior (distribution) as a choice made by nature, the receiver now faces uncertainty on nature's strategy. This problem falls into the problem category of picking the decision rule under subjective/non-additive expected utilities with this interpretation. Assuming the axiom of uncertainty aversion, maximin strategies that were axiomatized by *Gilboa and Schmeidler* (1989), are favored in when neither the degree of the uncertainty nor agents' uncertainty aversion level is negligible. In the rest of this discussion, we assume that the receiver uses the maximin strategy to design/choose the optimal experiments and discuss two examples in binary-outcome two-phase trials. The first example considers the simplest case where the receiver can design both phase-II experiments but has no knowledge about the prior distribution. The second example considers a more complex scenario where the receiver can only design one of the two phase-II experiments. However, he knows the possible priors in future cases[25]. Before presenting the examples, we pause to add that at the moment, we are unable to extend our analysis to multi-phase trials where the receiver can devise experiments in the middle phases. We believe that the results of *Doval and Ely* (2020) could be useful in this context but again leave this exploration to future work.

3.5.2.1 When receiver controls phase-II experiments in a two-phase trial

Let's use the same model analyzed in Section 3.3 as an example to shed some light on the receiver's optimal experiments. First, we consider the case where the receiver can control both experiments in phase II. In other words, the receiver can design parameters $q_{A1}, q_{A2}, q_{B1}, q_{B2}$. To avoid redundant discussions and to simplify the analysis, we still assume these parameters satisfy the assumptions made in Claim 5. When there are no real-world constraints on how precise the experiments can be, i.e., $q_{A1}, q_{A2}, q_{B1}, q_{B2} \in [0, 1]$, and the receiver uses a maximin strategy among all possible priors, his optimal experiments are stated in the following lemma:

Lemma 19. *If the receiver can design phase-II experiments with no constraints, then her optimal experiments* E_A, E_B *is* $E_A = (q_{A1}, q_{A2}) = (1, 0)$ *and* $E_B = (q_{B1}, q_{B2}) = (1, 0)$. *In other words, the experiment outcome will reveal the true states regardless of the prior and interim beliefs.*

When we allow the receiver to design experiments without any constraints, the receiver will pick two revealing experiments in phase-II to guarantee that he will take the preferred action, whatever the state is. This outcome occurs due to the extreme flexibility available to the receiver.

[25]Specifically, the receiver knows the support of the prior but the distribution of the prior is still unknown.

3.5.2.2 Example: Receiver can only pick one of the phase-II experiment from a finite set

This example discusses the receiver's optimal experiment selection if one of the phase-II experiments is already fixed (by other external means) and she is only allowed to pick the second experiment from a given finite set of experiments. To simplify the discussion, hereafter, we assume that the fixed experiment is E_A where the receiver knows (q_{A1}, q_{A2}). In this example, the receiver can pick the experiment E_B from a finite set $\mathbf{E_B}$, and she also knows that the future prior will be in the range $[a, b]$.

To make the result more general, we rule out the second assumption in Claim 5, i.e., we do not assume that experiment E_A is more (Blackwell) informative than the experiment E_B when the true state is θ_1. However, we still assume that the probability of passing a phase-II trial when the true state is θ_1 is greater than or equal to the probability of passing a phase-II trial when the true state is θ_2, i.e., $q_{A1} \geq q_{A2}$ and $q_{B1} \geq q_{B2}$.

When the receiver uses the maximin strategy in knowing the prior region $[a, b]$ under the utility function $U_R((,a_)i, \theta_i) = \mathbf{1}_{\{i=j\}}$, the receiver's objective is to find the experiment $E_B \in \mathbf{E_B}$ that minimizes her maximum probability of taking the wrong action in the prior region $[a, b]$. For example, given $E_A = (1, 0)$, $\mathbb{E_B} = \{E_B, E_{\hat{B}}\}$ with $q_{B1} = q_{\hat{B}1}$ and $q_{B2} \geq q_{\hat{B}2}$ and $[a, b] = [0, \frac{q_{B2}}{q_{B1}+q_{B2}}]$, the receiver will choose use experiments $(E_A, E_{\hat{B}})$ instead of (E_A, E_B). If the receiver chooses (E_A, E_B), he will guess incorrectly more often for every prior $p \in [0, \frac{q_{B2}}{q_{B1}+q_{B2}}]$ than when using experiment $(E_A, E_{\hat{B}})$. When the set of available experiment $\mathbf{E_B}$ is small, comparing the receiver's maximin expected utility for each pair of (E_A, E_B^i), $E_B^i \in \mathbf{E_B}$ in a given prior region is a viable approach. However, when $\mathbf{E_B}$ is large, an approach that efficiently searches through the available experiments (E_A, E_B^i), $E_B^i \in \mathbf{E_B}$ is required. A preliminary result here provides a partial order of experiments in terms of the receiver's utility. To help in formulating the partial order of experiments, we define the strong notion of inferiority of experiments.

Definition 27. *An experiment E_X is inferior to an experiment E_Y if the receiver's expected utility of choosing (E_A, E_Y) is higher than the expected utility of choosing (E_A, E_X) for all priors in the specified region and for any given experiment E_A.*

Given the above definition of inferiority, a partial ordered set of $\mathbf{E_B}$ can be built using Lemma 20:

Lemma 20. *Given two experiments E_X, E_Y, E_X is inferior to E_Y if one of the following two conditions is satisfied:*

$$(1)\ q_{Y1} \leq \frac{2-q_{Y2}}{3-2q_{Y2}} \ \textit{and} \ \max\{2q_{X1}-1, \frac{1-q_{X1}}{1-q_{X2}}\} < \frac{1-q_{Y1}}{1-q_{Y2}},$$
$$(2)\ q_{Y1} > \frac{2-q_{Y2}}{3-2q_{Y2}},\ q_{X_1} < q_{Y1},\ \textit{and}\ \frac{1-q_{X1}}{1-q_{X2}} < \frac{1-q_{Y_1}}{1-q_{Y2}}.$$

Lemma 20 can help the receiver to reduce the search of optimal experiments from a large set of experiment candidates, especially when the available experiments is not a countable set. The following example demonstrates how Lemma 20 reduces the set of searching the experiments.

Example: Searching for the optimal experiments from two sets Consider a given experiment $E_A = (0.7, 0.5)$. The receiver now can choose an experiment E_B from two different sets. The first set has low q_{B1} and low q_{B2}. It has to satisfying the following requirements:

1. The experiment has a bounded pass rate, $0.1 < q_{B2} \leq q_{B1} \leq 0.7$.

2. The experiment has a limit of informativeness $q_{B1} - q_{B2} \leq 0.25$ and $\frac{q_{B1}}{q_{B2}} \leq 1.6$.

The second set has high pass rate q_{B1}, q_{B2} but it q_{B1}, q_{B2} are functions satisfying the following requirements:

1. The experiment has a bounded pass rate, $0.7 \leq q_{B1} \leq 0.9$.

2. q_{B2} is a function of q_{B1}, $q_{B2} = 1.05 - 2.5(1 - q_{B1})$.

The receiver wants to choose the experiment from one of these two sets to maximize her expected utility.

Given Lemma 20, every other experiment in the first set is inferior to the experiment $(\frac{2}{3}, \frac{5}{12})$. (Every experiment in the first set satisfying $q_{B1} \leq \frac{2-q_{B2}}{3-q_{B2}}$, hence, applying the first condition of Lemma 20 can help us build the partial order of experiments in the first set.) Hence, if the receiver will pick E_B from the first set, she should choose the experiment $(\frac{2}{3}, \frac{5}{12})$.

In the second set, every experiment satisfies $q_{B1} > \frac{2-q_{B2}}{3-q_{B2}}$. Hence, we can use the second condition of Lemma 20 to build the partial order of this set of experiments. In this set, every other experiment is inferior to the experiment $(0.9, 0.8)$. Thus, the receiver only needs to consider the experiment $(0.9, 0.8)$.

Now, we want to look at if we can order $(\frac{2}{3}, \frac{5}{12})$ and $(0.9, 0.8)$ directly using the inferiority defined in Definition 27. Unluckily, between this pair neither is inferior to the other. Hence, we needs to look at whether the pair of experiments $\{(0.7, 0.5), (\frac{2}{3}, \frac{5}{12})\}$ gives the sender a higher expected utility than the pair $\{(0.7, 0.5), (0.9, 0.8)\}$. After the calculation, the pair $\{(0.7, 0.5), (\frac{2}{3}, \frac{5}{12})\}$ gives the sender higher (or equal) expected utility for every prior. Hence, we can claim the receiver's optimal experiment in this example is $(\frac{2}{3}, \frac{5}{12})$.

3.6 Summary of the Chapter

Chapter III studies Bayesian persuasion problems in multi-phase trials with determined experiments, and highlights the significant difference between single-phase and multi-phase trials: conducted experiments can be correlated in multi-phase trials, and determined experiments can constrain the sender's signalling strategy. When the experiments have binary outcomes, we explicitly analyze the two-phase trial and propose a dynamic programming approach to solve multi-phase trials using backward iteration. The proposed dynamic programming approach helps the sender systematically update her interim beliefs, and makes a design of strategic (and complex) signaling scheme possible. We also study non-binary-outcome experiments, discuss scenarios where the determined experiments can be designed by the receiver, and analyze the robustness of signaling strategies under small perturbation (from classical Bayesian persuasion). Results in these discussions demonstrates the fragility of optimal signaling strategies in environments where signals are revealed sequentially, and highlight the significant of tailored signaling strategies in information design problems under higher-order inference.

CHAPTER IV

Importance of the Commitment Order

4.1 Motivation – Inferring Other Senders' Commitments

In Chap III, we studied how a sender systematically analyzes and designs her optimal commitment in environments where signals/experiment results are generated sequentially. When none of the experiments are determined, the model in Chap III can also be interpreted as a model with multiple senders with aligned utilities, where each sender only acts once. This interpretation is closely related to the work *Li and Norman* (2018b). In the Li-Norman model, senders commit sequentially to post experiments, and every experiment is conditional independent of each other given the state. With this conditional independence, Li and Norman analyze optimal signaling strategies for heterogeneous[1] senders with same (and full[2]) information about the state. Their primary contribution is to prove that this sequential persuasion cannot yield a more informative equilibrium in comparison to simultaneous persuasion *Gentzkow and Kamenica* (2016, 2017) and is strictly less informative in binary state models. This result stands on the conditional independence of the experiments and relies on the assumption that the order of (senders') commitments perfectly matches the order of signals/experiments. In some real-world problems, this alignment may

[1]Senders with different utility functions.

[2]The sender can access all possible experiments under the "reduced" state space merging each state sets not distinguishable by the sender to a state.

not hold, and the participants may have the choice about who should commit first. The following mock example elucidates the possibility of such scenarios in real-world problems.

Example 8. *A company X searches for the most efficient way to produce its new product. For this purpose, it needs to decide between the following combinations: whether to apply a particular patent or not, and which manufacturing process (advanced or mature) it should adopt. The company sponsors a university lab Y to do some simulations for its decision on patent application. Based on the simulation results, the company then ask its partner factory Z to do some trial runs to determine the manufacturing process. However, lab Y always prefers the patent to be applied, and the factory Z wants the advanced process to be adopted. Thus, Y may choose simulation environments in favor of the patent being applied, and Z may choose the trial run parameters in favor of the advanced process. However, all the simulation and trial-run setups have to be submitted to the company beforehand. This brings out two questions: First, does the order of submission matter? Second, if it does, from the company's perspective, which party should submit their setups first, Y or Z?*

In the above example, Y and Z play the role of senders with commitments, and X is the receiver. In problems similar to the above example, senders have partial (maybe correlated) information of the state of the world. Further, the sender who sends the signal later, Z in the example, may be able to commit to a signaling scheme that depends on the state of the world and the realized signal of the other sender, Y in the example. We postpone a detailed illustration of the motivating example after discussing the model and the assumptions for our results.

4.2 Problem Formulation

There are two senders S_1, S_2 and one receiver R in the game. The state space Θ is finite, and the receiver's action space A is also finite. For simplicity of analysis, we assume that the state and action space have the same size, i.e., $|\Theta| = |A|$, and action a_θ is the receiver's unique best response of state $\theta \in \Theta$. To further simplify the analysis, we assume that the receiver only obtains (positive) utility upon matching the state correctly, i.e.,

$$U_R(\theta_i, a_j) \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \tag{4.1}$$

Before the true state is realized (by nature), the senders and the receiver have a common prior belief of the distribution of the state $p \in \Delta(\Theta)$. However, both senders have some private information about the true state when the true state is realized. To avoid some trivial cases and also to avoid overlap with the existing literature, we make the following four assumptions on senders' private information:

1. Each sender only has partial information of the true state. To avoid redundancy, we assume each sender S_k's information space $\mathcal{I}_k$ has the size $|\mathcal{I}_k| < |\Theta|$, $k \in \{1, 2\}$.

2. Senders S_1 and S_2 have different private information, i.e., $\mathcal{I}_1 \neq \mathcal{I}_2$.

3. Neither sender is more knowledgeable[3] than the other, and the private information of each sender is generated via an onto function $F_k : \Theta \to \mathcal{I}_k$. In other words, F_k is a partition of Θ, and neither F_1 is more Blackwell informative than F_2 nor F_2 is more Blackwell informative than F_1.

4. The true state can be revealed under both senders' truth-telling strategies. In

[3] In terms of Blackwell informativeness in *Blackwell* (1953).

other words, for every $\theta_i \in \Theta$, there exists an $I_1^x \in \mathcal{I}_1$ and an $I_2^y \in \mathcal{I}_2$ such that $\mathbb{P}(\theta_i | I_1^x, I_2^y) = 1$.

After both senders get their private information, S_1 sends a signal $\omega_1 \in \Omega_1$ to influence the receiver's belief. After observing the signal ω_1, S_2 sends a signal $\omega_2 \in \Omega_2$ to influence the receiver's belief. After receiving signals ω_1 and ω_2, the receiver has to take action $a \in A$. Extending the sender-preferred tie-breaking rules in *Kamenica and Gentzkow* (2011), we assume that the receiver breaks a tie in favor of S_1 first. When the action set maximizing both S_1 and receiver's utilities is not a singleton, the receiver takes an (arbitrary) action maximizing S_2's utility from the set.

In this model, both senders commit to their respective signaling strategies before receiving their private information. Moreover, S_1 and S_2 make their commitment sequentially under a pre-determined order. That is to say, the sender who commits later can exploit another sender's commitment to design her signaling strategy. We will first detail of the game flow (when the order of commitment in steps 2 and 3 is determined), and then discuss how the sender's commitment can depend on the other sender's signaling strategy based on the commitment order.

4.2.1 Procedure of the game

The game in our model evolves as follows:

1. Prior belief of the state $p \in \Delta(\Theta)$ is disclosed to all participants.

2. S_k commits her signaling strategy, where the index $k \in \{1, 2\}$ is determined beforehand and becomes common knowledge.

3. The other sender, S_l, commits to her signaling strategy, where $l \neq k$.

4. Nature chooses a realized state $\theta \in \Theta$.

5. S_1 gets her private information and sends a signal ω_1 according to her committed signaling strategy.

6. S_2 gets her private information and observes S_1's realized signal ω_1. Then, she sends a signal ω_2 according to her committed signaling strategy.

7. The receiver observes the realized signals ω_1, ω_2 and takes action $a \in A$ to maximize the probability of matching the state.

4.2.2 Receiver's best response

Given the procedure of the game, the receiver takes action after observing the realization of both senders' signals, ω_1, ω_2. Given that her objective is to match the state correctly, the receiver's best response is taking action $a^* \in A$ such that

$$a^* \in \arg\max_{a \in A} \mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_R(\theta, a)|\omega_1, \omega_2].$$

However, the tie-breaking rule, although indifferent to the receiver, plays a significant role on senders' signaling strategy, especially when multiple senders are considered in the model. As mentioned above, we assume the receiver breaks a tie in favor of S_1 first and then S_2. The formal expression of the tie-breaking rule is presented below:

Let $A_R \coloneqq \arg\max_{a \in A} \mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_R(\theta, a)|\omega_1, \omega_2]$, $A_{S_1} \coloneqq \arg\max_{a' \in A_R} \mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_{S_1}(\theta, a')|\omega_1, \omega_2]$,

$$\mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_{S_1}(\theta, a^*)|\omega_1, \omega_2] = \max_{a' \in A_R} \mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_{S_1}(\theta, a')|\omega_1, \omega_2], \tag{4.2}$$

$$\mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_{S_2}(\theta, a^*)|\omega_1, \omega_2] = \max_{a'' \in A_{S_1}} \mathbb{E}_{p,\Gamma_1,\Gamma_2}[U_{S_2}(\theta, a')|\omega_1, \omega_2]. \tag{4.3}$$

To avoid ambiguity, $\mathbb{E}_{p,\Gamma_1,\Gamma_2}[U(\theta, a)|\omega_1, \omega_2]$ represents the expected utility conditional on the realized signals ω_1 and ω_2 under the prior p, S_1's commitment Γ_1, and S_2's commitment Γ_2, given action a is taken. In short, the expectation is taken only on the prior. Equation (4.2) states that the receiver chooses an (arbitrary) action in the set which maximizes S_1's expected utility while she is indifferent. Equation (4.3) states if there is still a tie after maximizing both S_1 and receiver's expected utility, the receiver chooses an (arbitrary) action in the set which maximizes S_2's expected utility. For simplicity of representation, we abuse the notation and let $a^*(\omega_1, \omega_2)$ represents the receiver's best response under ω_1 and S_2's signal ω_2 in the rest of the chapter.

4.2.3 Assumption on commitments

As mentioned earlier, we detail why S_l may be able to utilize another sender S_k's commitment when she can commit later. The key is the pre-determined order of the sequential signaling. Since S_1 always sends signal before S_2, the signaling schemes of the senders are asymmetric. In other words, S_1 cannot commit to a signaling scheme depending on the realization of ω_2, but S_2 can commit to a signaling scheme based on the realization of ω_1. Next we clarify the information that S_2 can exploit while making her commitment. To avoid ambiguity, we discuss both commitment orders and will introduce our cornerstone assumption, namely **permutation-free** commitments. Before the discussion, we note that the superscript f, s of the commitment notation Γ^f, Γ^s denote the commitment order, and the subscript denote the sender and another sender's commitment (if available), e.g., Γ^s_{2,Γ^f_1} denotes S_2's commitment when she commits after S_1 under S_1's commitment Γ^f_1.

Case 1: S_1 commits to her signaling strategy first

- Since sender S_1 has no observation of sender S_2's signal while sending her signal and has to commit first, her commitment Γ^f_1 can be represented as a function $\Gamma^f_1 : \mathcal{I}_1 \to \Delta(\Omega_1)$, where $|\Omega_1| \leq |\mathcal{I}_1|$.

- Sender S_2 can utilize sender S_1's commitment and realized signal to make her commitment. Hence, sender S_2's commitment Γ^s_{2,Γ^f_1} can be represented as a function $\Gamma^s_{2,\Gamma^f_1} : \mathcal{I}_2 \times \Omega_1 \to \Delta(\Omega_2)$.

Case 2: S_2 commits to her signaling strategy first

- Even though S_1 knows sender S_2's commitment Γ_2^f, she still has no observation of sender S_2's signal while sending her signal. Hence, her commitment Γ^s_{1,Γ_2^f} can be represented as a function $\Gamma^s_{1,\Gamma_2^f} : \mathcal{I}_1 \to \Delta(\Omega_1)$, where $|\Omega_1| \leq |\mathcal{I}_1|$. In other words, she can adjust her commitment based on S_2's commitment, but cannot tailor her commitments by utilizing S_2's realized signals.

- S_2 wants to utilize sender S_1's realized signal to make her (optimal) commitment. However, S_2 has to commit before sender S_1 in this case. Since we don't want each signal token to possess an implied meaning, thereby allowing S_1 to act against S_2 by just reordering the signal tokens, we assume that the S_2 commits to a **permutation-free commitment**.

Definition 28. *A commitment $\Gamma_2^f : \mathcal{I}_2 \times \Omega_1 \to \Delta(\Omega_2)$ is **permutation-free** if for every given S_1's commitment Γ_1^s, there is no permutation matrix M such that*

$$\mathbb{E}_{p,\Gamma_1^s}[U_{S_2}|\Gamma_2^f(I_2, M(\omega_1))] > \mathbb{E}_{p,\Gamma_1^s}[U_{S_2}|\Gamma_2^f(I_2, \omega_1)].$$

Next, we justify why we believe that the permutation-free assumption on commitments is an appropriate assumption when S_2 has to commit first. When we exploit Bayesian persuasion in real-world problems, commitments are usually implemented via experiments (which can be operated/supervised by a third party). Hence, when S_2 has to commit first and wants to exploit S_1's realized signal, she can propose different experiments depending on S_1's signal realization. Unless we assume signals have implied meaning or that S_2 has some prior knowledge about S_1's signaling strategy, S_2 cannot know the mapping from signal ω_1 to the interim distribution of the states. However, when the experiment of S_2 has to be executed, S_1's signal realization and her commitments are both common knowledge. Hence, S_2 can commit to experiments with her private signal and the mock signal realizations of S_1, and ask the operator of her experiments to use the permutation/reordering of S_1's signal tokens that works best for her. These experiments can be appropriately executed, and this permutation-free commitment can be fulfilled. Moreover, violating the permutation-free property implies that S_1 and S_2 have an ex-ante consensus on the meaning of S_1's signal tokens, but S_1 can violate their ex-ante consensus to act against S_2. Besides, when S_2 may suffer loss from sender S_1 permuting her signal tokens, S_2's experiments will reveal her (partial) private information only if S_1 cannot exploit this revelation against her via permutations. This will significantly reduce the utility of being able to choose the commitment order. To further illustrate this, we demonstrate via a numerical example that S_2's optimal commitment ends up being independent of S_1's signal realization when permutation-free commitments are not enforced; the rest of the chapter will analyze the case of permutation-free commitments.

Example 9. *There are four possible states of the world $\Theta = \{TL, TR, BL, BR\}$ (top-bottom and left-right) with prior distribution depicted in Table 4.3. Sender S_1 knows whether the state of the world is top or bottom $\mathcal{I}_1 = \{T, B\}$ and sender S_2 knows whether the state of the world is left or right $\mathcal{I}_2 = \{L, R\}$. S_1 can send signals $\omega_1 \in \{\omega_{1,1}, \omega_{1,2}\}$ and S_2 can send signals $\omega_2 \in \{\omega_{2,1}, \omega_{2,2}\}$. We assume that both senders' utilities are as given in Table 4.2, and this example requires sender S_2 to commit before S_1. All participants' utility functions and the (prior) state distribution are common knowledge.*

Table 4.1: Distribution of the states in Example 9

	L	R
T	0.1	0.2
B	0.5	0.2

Table 4.2: Utilities of senders in Example 9

	a_{TL}	a_{TR}	a_{BL}	a_{BR}
(U_{S_1}, U_{S_2})	2,0	4,2	3,0	0,3

Given the above prior and the utility functions, sender S_1 and S_2's objective functions are the following:

$$\Gamma_1^{s*} = \arg\max_{\Gamma_1^s \in \mathbf{\Gamma_1}} \mathbb{E}_p[U_{S_1}|\Gamma_2^s]$$

$$s.t.\ \mathbb{E}_p[U_{S_2}|\Gamma_1^s(\Gamma_2^{s*})] \geq \mathbb{E}_p[U_{S_2}|\Gamma_1^{s'}(\Gamma_2^{s^*})],$$

$$\textit{where } \Gamma_1^{s'} \textit{ is a commitment that swaps the signals of } \Gamma_1^s, i.e.,$$

$$\mathbb{P}_{\Gamma_1^s}(\omega_{1,1}|T) = \mathbb{P}_{\Gamma_1^{s'}}(\omega_{1,2}|T) \textit{ and } \mathbb{P}_{\Gamma_1^s}(\omega_{1,1}|B) = \mathbb{P}_{\Gamma_1^{s'}}(\omega_{1,2}|B).$$

$$\Gamma_2^{f*} = \arg\max_{\Gamma_2^f \in \mathbf{\Gamma_2}} \mathbb{E}_p[U_{S_2}|\Gamma_1^{s*}(\Gamma_2^s)]$$

Since S_1 cannot threaten S_2 via signal reordering, S_2 can commit to signaling strategies depending on the realization of S_1's signal. For simplicity of the representation, let's assume $\omega_{1,1}$ has higher posterior belief on S_1's private information T than $\omega_{1,2}$. S_2's optimal commitment is the following:

- *While observing $\omega_{1,1}$, send $\omega_{2,1}$ with probability 1 no matter what private information she knows.*
- *While observing $\omega_{1,2}$ and given her private information is R, send $\omega_{2,1}$ with probability 1.*
- *While observing $\omega_{1,2}$ and given her private information is L, send signal $\omega_{2,1}$ with probability $0.4 - \epsilon$ and $\omega_{2,2}$ with probability $0.6 + \epsilon$, where ϵ is infinitesimal, i.e., $|\epsilon| \ll 1$.*

Given S_2's commitment above, the optimal commitment of S_1 is the following:

- *Given her private information T, send $\omega_{1,1}$.*
- *Given her private information B, send $\omega_{1,1}$ with probability 0.4 and send $\omega_{1,2}$ with probability 0.6.*

Given this pair of commitments, the expected utilities of all the agents are given by $(U_{S_1}, U_{S_2}, U_R) = (2.86, 1.88, 0.5)$.

However, when S_2 *cannot commit to permutation-free commitments,* S_1*'s objective function becomes the following*

$$\Gamma_1^{s*} = \arg\max_{\Gamma_1^s \in \mathbf{\Gamma_1}} \mathbb{E}_p[U_{S_1}|\Gamma_2^s]$$

If S_2 *commits to the same signaling strategy,* S_1 *can simply swap* $\omega_{1,1}$ *with* $\omega_{1,2}$ *to get a higher expected utility. Under the threat of* S_1*'s reordering of signals,* S_2*'s commitment reduces to the following form. No matter the signal sent by* S_1*,* S_2 *commits to the following signaling scheme:*

- *When* S_2*'s private information is* R*, send* $\omega_{2,1}$ *with probability 1.*
- *When* S_2*'s private information is* L*, send signal* $\omega_{2,1}$ *with probability* $0.4-\epsilon$ *and* $\omega_{2,2}$ *with probability* $0.6+\epsilon$*, where* ϵ *is infinitesimal, i.e.,* $|\epsilon| \ll 1$.

Under S_2's commitment above, S_1's optimal commitment stays the same. However, the expected utility of agents changes to $(U_{S_1}, U_{S_2}, U_R) = (2.74, 1.76, 0.56)$.

In this example, when permutation-free commitments are not enforced, S_2 is under the threat of S_1's reordering on signals, so her optimal strategy is to make her commitment independent of S_1's signal realizations. Here, both senders suffer (utility losses) from the lack of permutation-free commitments.

4.2.4 Objectives of senders

Before we proceed, we list the objective functions of senders for both commitment orders. For simplicity of representation, we abuse notations and let Γ_k^{f*} denotes the optimal commitment when S_k commits first, and Γ_k^{s*} denotes the optimal commitment when S_k commits last.

When S_1 commits first, the objective functions of S_1 and S_2 are the following:

$$\begin{aligned}
\Gamma_1^{f*} &= \arg\max_{\Gamma_1^f \in \mathbf{\Gamma_1}} \mathbb{E}_{p,\Gamma_1^f,\Gamma_2^s(\Gamma_1^f)}[U_{S_1}] \\
&\quad \text{s.t. } \mathbb{E}_{p,\Gamma_1^f,\Gamma_2^s(\Gamma_1^f)}[U_{S_2}] = \max_{\Gamma_2^{s\prime} \in \mathbf{\Gamma_2}} \mathbb{E}_{p,\Gamma_1^f,\Gamma_2^{s\prime}}[U_{S_2}], \qquad (4.4) \\
\Gamma_2^{s*} &= \arg\max_{\Gamma_2^s \in \mathbf{\Gamma_2}} \mathbb{E}_{p,\Gamma_1^{f*},\Gamma_2^s}[U_{S_2}]. \qquad (4.5)
\end{aligned}$$

When S_2 commits first with permutation-free commitments, the objective functions of S_1 and S_2 are the following:

$$\begin{aligned}
\Gamma_1^{s*} &= \arg\max_{\Gamma_1^s \in \mathbf{\Gamma_1}} \mathbb{E}_{p,\Gamma_1^s,\Gamma_2^{f*}}[U_{S_1}] \text{ s.t. } \mathbb{E}_{p,\Gamma_1^s,\Gamma_2^{f*}}[U_{S_2}] = \max_{\Gamma_1^{s\prime} \in \mathbf{\Gamma_1^M}(\Gamma_1^s)} \mathbb{E}_{p,\Gamma_1^{s\prime},\Gamma_2^{f*}}[U_{S_2}], \\
&\quad \text{where } \mathbf{\Gamma_1^M}(\Gamma_1^s) \text{ is a set of commitments where the signals are permuted} \\
&\quad \text{when compared to } \Gamma_1^s. \qquad (4.6) \\
\Gamma_2^{f*} &= \arg\max_{\Gamma_2^f \in \mathbf{\Gamma_2}} \mathbb{E}_{p,\Gamma_1^s,\Gamma_2^f}[U_{S_2}] \text{ s.t. } \mathbb{E}_{p,\Gamma_1^s,\Gamma_2^f}[U_{S_1}] = \max_{\Gamma_1^{s\prime} \in \mathbf{\Gamma_1}} \mathbb{E}_{p,\Gamma_1^{s\prime},\Gamma_2^f}[U_{S_1}]. \qquad (4.7)
\end{aligned}$$

4.3 When the Commitment Order Matters

In previous subsections, we discussed the rationale for why the permutation-free property on commitments is an appropriate assumption in sequential commitment problems. In this section, we discuss conditions when the commitment order matters, i.e., when different expected utilities are obtained for different commitment orders so that senders or the receiver would prefer between them. We start by presenting a numerical example that shows that the commitment order can result in a credible threat even though the order of signaling stays unchanged.

4.3.1 Commitment order and credible threats

Let's consider a game with two senders and one receiver. There are 4 possible states of the world $\Omega = \{TL, TR, BL, BR\}$ (top-bottom and left-right) similar to the one discussed earlier. Again, S_1 knows whether the state of the world is top or bottom, i.e., $\mathcal{I}_1 = \{T, B\}$, and S_2 knows whether the state of the world is left or right, i.e., $\mathcal{I}_2 = \{L, R\}$. The distribution of the state of the world (prior) is common knowledge, listed in Table 4.3. The utilities of senders are in Table 4.4. We remind the reader that the receiver's goal is to guess the state of the world accurately: she receives payoff 1 when she guesses correctly and 0 otherwise.

Table 4.3: Distribution of the states

	L	R
T	0.1	0.2
B	0.4	0.3

Table 4.4: Utilities of senders

	a_{TL}	a_{TR}	a_{BL}	a_{BR}
(U_{S_1}, U_{S_2})	1,0	2,2	0,0	0,3

In this game, S_1 prefers top over bottom and S_2 prefers right over left. Now, we study the optimal commitments of S_1 and S_2 under different commitment orders. First, let's start with the scenario where S_1 make her commitment first.

4.3.1.1 S_1 commits first

If S_2 didn't appear in the game, S_1 would act as if in a classical 1-sender Bayesian persuasion model: sending a signal which suggests action a_{TL} or a_{TR} when she knows the true state is in the top half and perform a mixed signaling strategy when the true state is in the bottom half. However, in the presence of S_2, S_1 knows that as long as she can guarantee that the posterior probability of BR is smaller than or equal to TR, S_2, for the sake of maximizing her expected payoff, will commit to a signaling strategy which helps S_1 since BL is inferior to TR for S_2. Bearing this in mind, S_1 will make the following commitment by anticipating S_2 commitment.

S_1's commitment: Send a suggestion of T with probability 1 when her private information is T and probability $\frac{2}{3}$ when her private information is B.

Now, since S_1 has made her commitment, S_2 makes her commitments on top of it. S_2's optimal commitment is as follows:

S_2's commitment

- While observing S_1's signal suggesting T, send a suggestion R with probability 1 when her private information is R and probability $\frac{3}{4} - \epsilon$ when her private information is L. And send a suggestion of L with probability $\frac{1}{4} + \epsilon$ while observing L, where ϵ is infinitesimal, i.e., $|\epsilon| \ll 1$.

- While observing S_1's signal suggesting B, send a suggestion R with probability 1 when her private signal is R and probability $\frac{3}{4} - \epsilon$ when her private signal L. And send a suggestion of L with probability $\frac{1}{4} + \epsilon$ while observing L, where ϵ is infinitesimal, i.e., $|\epsilon| \ll 1$.

We notice that S_2's commitment is independent of S_1's realized signal. This is merely a happenstance in this example; this makes the analysis simpler too. Given the commitments above, the expected utility of S_1,S_2 and Receiver can be calculated below:

$$\mathbb{E}_p[U_{S_1}] = (0.2 + 0.3 \times \tfrac{2}{3} + 0.4 \times \tfrac{2}{3} \times \tfrac{3}{4}) \times 2 + 0 = 1.2 \quad (4.8)$$

$$\begin{aligned} \mathbb{E}_p[U_{S_2}] &= (0.2 + 0.3 \times \tfrac{2}{3} + 0.4 \times \tfrac{2}{3} \times \tfrac{3}{4}) \times 2 + \tfrac{1}{3} \times (0.3 + 0.4 \times \tfrac{3}{4}) \times 3 \\ &= 1.2 + 0.6 = 1.8 \end{aligned} \quad (4.9)$$

$$\mathbb{E}_p[U_R] = 0.2 + 0.3 \times \tfrac{1}{3} + 0.4 \times \tfrac{1}{3} \times \tfrac{1}{4} = \tfrac{1}{3} \quad (4.10)$$

4.3.1.2 S_2 commits first

Now, consider the game where S_2 commits first. Although S_2 does not know S_1's signaling strategy while making her commitment, her commitment can still use S_1's signal realization. This is because, when her experiment is conducted, S_1's signal realization is common knowledge. Therefore, she can commit to the following signal and use this to threaten S_1 credibly.

S_2's commitment

- While observing S_1's suggestion on T, send a suggestion R with probability 1, whatever her private information is.

- While observing S_1's suggesting on B, send a suggestion R with probability 1 when her private information is R and probability $\frac{3}{4}$ when her private information is L. And send a suggestion L otherwise.

Given this commitment of S_2, S_1, now, cannot hope that S_2 will help her when S_1 signals her private information is T. Hence, S_1 has to make a more conservative commitment.

S_1's commitment

- Send a suggestion T with probability 1 when her private information is T and probability $\frac{1}{2}$ when S_1's private information is B.

- Send a suggestion B otherwise.

Now, we have to check whether S_2 will deviate from her original commitment to verify the optimality from S_2's perspective. Since the receiver breaks a tie in favor of S_1 and then S_2, using any signal L with probability > 0 when S_1 sends a suggestion T will strictly lower her expected utility. When S_1 suggests B, the true state is either BL or BR, and her commitment already maximizes her expected utility by making the receiver weakly prefer a_{BR} (given the receiver's tie-breaking rule and the indifference of S_1's utility between a_{BL} and a_{BR}). Hence, her original commitment is the best response to the S_1's commitment.

Calculating the expected utility of agents, we obtain:

$$\mathbb{E}_p[U_{S_1}] = (0.2 + 0.3 \times \tfrac{1}{2} + 0.4 \times \tfrac{1}{2}) \times 2 + 0 = 1.1 \tag{4.11}$$

$$\begin{aligned}\mathbb{E}_p[U_{S_2}] &= (0.2 + 0.3 \times \tfrac{1}{2} + 0.4 \times \tfrac{1}{2}) \times 2 + \frac{1}{2} \times (0.3 + 0.4 \times \tfrac{3}{4}) \times 3 \\ &= 1.1 + 0.9 = 2\end{aligned} \tag{4.12}$$

$$\mathbb{E}_p[U_R] = 0.2 + 0.3 \times \tfrac{1}{2} + 0.4 \times \tfrac{1}{2} \times \tfrac{1}{4} = 0.4 \tag{4.13}$$

4.3.1.3 Comparison of utilities

Next, we compare the utilities for senders and the receiver between the two commitment orders and assess their value for each agent. Both senders want to commit first, and the receiver has a preference on letting S_2 commit first. Moreover, if senders can trade their commitment order and S_1 is initially endowed the first commitment slot, every transfer $x \in (0.1, 0.2)$ from S_2 to S_1 makes both senders better off by trading the first commitment slot from S_1 to S_2. The example shows an interesting phenomenon: when senders only have partial information and have to send signals sequentially, they can collaborate on the signals they commit and collude on the order of commitment if there is a market for senders to trade their commitment position. For readers curious about whether there is a scenario that both senders and the receiver are better off after a transfer of the commitment slot, an example is provided in Section 4.4.3.

4.3.2 Main results

Motivated by the above example, we first present an easily rationalized proposition capturing a set of scenarios where the commitment order matters and prove a generalized theorem afterward. The proposition captures scenarios similar to the example just discussed, wherein S_1's and S_2's interests are aligned if S_1 reveals her private information truthfully. In these scenarios, both senders want to persuade the receiver to take the same action. This implies that senders can collaborate on their signaling strategies in some states. We call this type of states **collaborative states** and define them formally below.

Definition 29. *A state $\hat{\theta} \in \Theta$ is a collaborative state if the following two conditions hold:*

1. $\hat{\theta} \in \arg\max_{\theta \in \hat{I}_1} U_{S_1}(\theta, a_\theta)$,

2. $\hat{\theta} \in \arg\max_{\theta \in \hat{I}_1} U_{S_2}(\theta, a_\theta)$,

where $\hat{I}_1$ is the information set containing state $\hat{\theta}$, i.e., $\hat{I}_1 = F_1(\hat{\theta})$.

This definition states that when the true state is $\hat{\theta}$ and S_1 reveals her private information truthfully, S_2 is willing to collaborate with S_1 to persuade the receiver towards taking their most preferred action. Note the asymmetry of the definition vis-a-vis the sender identities. This occurs because S_1 signals before S_2.

Given this definition, we present our first result about conditions when the commitment order matters.

Proposition 4. *Given a prior p and both senders' utility functions U_{S_1}, U_{S_2}, the commitment order matters if there exists a private information set $I_1 \in \mathcal{I}_1$ satisfying the following conditions:*

1. *There exists a collaborative state $\hat{\theta} \notin I_1$ satisfying the following two inequalities:*

$$U_{S_2}(\hat{\theta}, a_{\hat{\theta}}) < \mathbb{E}_{p,\Gamma_2^*(I_1)}[U_{S_2}(\theta, a^*(I_1, \omega_2))|I_1] \tag{4.14}$$

$$U_{S_1}(\hat{\theta}, a_{\hat{\theta}}) > \mathbb{E}_{p,\Gamma_2^*(I_1)}[U_{S_1}(\theta, a^*(I_1, \omega_2))|I_1], \tag{4.15}$$

   where $\Gamma_2^*(I_1)$ denotes S_2's optimal signaling strategy while S_1 truthfully reveals her private signal I_1, and $a^*(I_1, \omega_2)$ denotes the receiver's best response under I_1 and S_2's signal ω_2.

2. Let $\hat{\mathbf{\Gamma}}_1$ be the set of commitments with a signal $\hat{\omega}_1 \in \Omega_1$ satisfying $\mathbb{P}(I_1|\hat{\omega}_1) = 1$. For any commitment order, the optimal commitment(s) of S_1 must belong to this set, i.e., $\Gamma_1^{f*}, \Gamma_1^{s*} \subseteq \hat{\mathbf{\Gamma}}_1$.

3. The preference orders of S_1 and S_2 on $\Phi(I_1)$ are polar opposites, where $\Phi(I_1)$ is the receiver's possible optimal action set given S_1's private information I_1, i.e., $\Phi(I_1) = \{a_\theta | \theta \in I_1\}$.

We provide a detailed proof in Appendix D.1.1. The core of the proof is to construct a credible threat using a collaborative state. To make the construction possible, we need the existence of another signal which gives S_1 a lower expected utility (based on the receiver's best response). This other signal gives S_1 an incentive to commit to a mixed signaling strategy, which mixes this signal with the signal corresponding to the collaborative state. Otherwise, S_1's private information can be elicited by S_2 when S_1's private information contains the collaborative state, and credible threats are no longer needed since S_1's and S_2's interests are aligned. The existence of this other signal is implied by condition 2 in non-trivial scenarios, where the details are stated in Claim 8. Moreover, conditions 1 and 3 guarantee the conflict of interest between S_1 and S_2's on mixing other states with the collaborative state. Hence, if S_1 gets to commit first, she can propose a more aggressive mixed strategy because the best commitment of S_2, after S_1's commitment, is to collaborate with S_1's mixed strategy. However, if S_2 can commit first, she is willing to collaborate, but she does not want S_1 to mix other states with this collaborative state. Hence, she can create a credible threat to reduce (but not eliminate) S_1's strategy of mixing other states with the collaborative one. Hence, the commitment order matters in this scenario.

Claim 8. *When the inequality (4.15) holds, condition 2 in Proposition 4 implies that one of the following two conditions is true:*

1. *The signal $\hat{\omega}_1$ of condition 2 in Proposition 4 is not the signal realization which gives S_1 the highest expected utility in S_1's optimal commitments.*

2. *Every signal realization gives S_1 the same expected utility.*

Note that while Proposition 4 provides a set of sufficient conditions when the commitment order matters in a constructive manner, each condition has its deficiency. First, the requirement of a collaborative state is a strong condition. S_2 may be willing to partially collaborate with S_1 after she has maximized the probability of her most-preferred action under a given I_1. In other words, a collaboration may happen in a state where its corresponding action is S_2's second-best action (under the given information set of S_1). Second, the verification of condition 2 demands some knowledge of the structure of the commitment space[4]. When the state space grows, this can be a tedious task without additional constraints. Third, condition 3 guarantees the uniqueness of S_2's optimal signaling strategy while S_1 truthfully reveals her private signal I_1 and a conflict of interest between S_1 and S_2 on I_1. However, a polar opposite preference ordering is only one type where a conflict of interest between S_1 and S_2 happens. As long as S_1 and S_2 have an opposite preference on a mixture of two information sets of S_1, S_2 may be able to construct a credible threat using this opposite preference. We provide an example where the commitment order matters but none of the conditions in Proposition 4 hold in Section 4.4.3. Hence, to remove the need for the knowledge of the commitment space, the requirement of collaborative states, and the assumption of polar opposite preference ordering, we will present a set of sufficient conditions generalized from Proposition 4. We start by defining the best response signaling strategy of S_2 under an observed signal of S_1.

Definition 30. *Given a prior p and S_1's commitment Γ_1, a signal realization ω_1 indicates a conditional distribution of S_1's information sets, $\Delta(\mathcal{I}_1|\omega_1)$, denoted as*

[4]When the utilities can be written in a matrix form, it is usually straightforward to see whether there is an information set of S_1 gives her a lower utility comparing to the collaborative state.

q^{ω_1}. *Moreover, since the best response of* S_2 *and the best response of receiver will stay unchanged respectively under two signals* ω_1, ω_1' *with the same information set distribution, i.e.,* $q^{\omega_1} = q^{\omega_1'}$, *we abuse notation to avoid duplication and use* $a^*(q^{\omega_1}, \omega_2)$ *to represent* $a^*(\omega_1, \omega_2)$.

Definition 31. *Given a prior* p *and a distribution of* S_1*'s information set* q^{ω_1}, $G(p, q^{\omega_1})$ *denotes the set of best response signaling strategies while observing* S_1*'s information set distribution* q^{ω_1}. *Mathematically, an element* $g \in G(p, q^{\omega_1})$ *is a function that maps* $I_2 \in \mathcal{I}_2$ *to a signal distribution* $r \in \Delta(\Omega_2)$ *such that* $\mathbb{E}_{p,g}[U_{S_2}|q^{\omega_1}, r] \geq \mathbb{E}_{p,g'}[U_{S_2}|q^{\omega_1}, r]$ *for all* $g' \notin G(p, q^{\omega_1})$. *For simplicity of the representation,* $G(p, q^{\omega_1} + \alpha q^{\bar{\omega}_1})$ *denotes* S_2*'s best response under* S_1*'s mixed signal* $\hat{\omega}_1$ *such that* $q^{\hat{\omega}_1} = \frac{1}{1+\alpha} q^{\omega_1} + \frac{\alpha}{1+\alpha} q^{\bar{\omega}_1}$. *Since* S_2*'s expected utility is maximized in every best response strategy* $g \in G(p, q^{\omega_1})$, *we use* $\mathbb{E}_p[U_{S_2}|\omega_1, G(p, q^{\omega_1})]$ *to represent* S_2*'s expected utility under* S_1*'s information set distribution* q^{ω_1}, *prior* p, *and* S_1*'s signal* ω_1.

Given the above definitions, we introduce a sufficient condition of when the commitment order matters.

Theorem 2. *Given a prior* p, *the commitment order matters if there exists a pair of* S_1*'s information set* I_1^x, I_1^y *satisfying the following conditions:*

1. *There exist two parameters* $\alpha > \beta > 0$ *and a signaling strategy* $\hat{\Gamma}_2$ *satisfying the following conditions:*

   (a) $G(p, I_1^x + \alpha I_1^y) = G(p, I_1^x + \beta I_1^y)$,

   (b) *Let* $\omega_1^\alpha, \omega_1^\beta$ *be two mock signals of* S_1 *such that* $\mathbb{P}(I_1^x|\omega_1^\alpha) = \frac{1}{1+\alpha}$, $\mathbb{P}(I_1^y|\omega_1^\alpha) = \frac{\alpha}{1+\alpha}$, $\mathbb{P}(I_1^x|\omega_1^\beta) = \frac{1}{1+\beta}$, *and* $\mathbb{P}(I_1^y|\omega_1^\beta) = \frac{\beta}{1+\beta}$, *then*

   $$\hat{\Gamma}_2(p, \omega_1^\beta) \in G(p, I_1^x + \beta I_1^y) \text{ and } \hat{\Gamma}_2(p, \omega_1^\alpha) \notin G(p, I_1^x + \alpha I_1^y),$$

   (c) $\mathbb{E}_p[U_{S_1}|\omega_1^\alpha, \hat{\Gamma}_2] < \frac{\mathbb{P}(I_1^x)+\beta\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)} \mathbb{E}_p[U_{S_1}|\omega_1^\beta, \hat{\Gamma}_2] + \frac{(\alpha-\beta)\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)} \mathbb{E}_p[U_{S_1}|I_1^y, \hat{\Gamma}_2]$,

(d) $\mathbb{E}_p[U_{S_2}|\omega_1^{\beta}, G(p, I_1^x + \beta I_1^y)] < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]$.

2. *Sender S_1 has a higher expected utility under I_1^x than I_1^y, regardless of S_2's tie-breaking rule. In other words,*

$$\min_{g \in G(p, I_1^x)} \mathbb{E}_{p,g}[U_{S_1}|I_1^x] > \max_{g' \in G(p, I_1^y)} \mathbb{E}_{p,g'}[U_{S_1}|I_1^y].$$

3. *There is no commitment pair (Γ_1, Γ_2) such that*

$$\sum_{\theta_i \in I_1^y} \mathbb{P}_{p,\Gamma_1,\Gamma_2}(a_i) = 0 \ \textit{or} \ \sum_{\theta_i \in I_1^x} \mathbb{P}_{p,\Gamma_1,\Gamma_2}(a_i) = 0$$

Although the conditions in Theorem 2 look technical, there are intuitions behind each condition. The first condition in Theorem 2 states that the sender S_2 prefers a separation of sender S_1's information sets I_1^x and I_1^y. Moreover, if S_2 can commit first, she can make a credible threat on her commitment to constrain the mixing of I_1^x and I_1^y in sender S_1's signaling strategies. This generalizes condition 3 in Proposition 4. Besides, the combination of condition 1.(a) and 1.(d) guarantees that S_2 can collaborate with S_1, and this generalizes the collaborative state requirement in Proposition 4. The second condition in Theorem 2 states that the sender S_1 prefers a mixture of I_1^x and I_1^y in her signaling strategy. Combining the intuition behind the first two conditions, the two senders have opposite preferences on the mixture of I_1^x and I_1^y. Based on the first two conditions, senders S_1, S_2 are already in a scenario where the sender S_2 is capable of constraining S_1's mixture of I_1^x and I_1^y, and a collaboration between S_1 and S_2 is possible under I_1^x. Hence, the proof of Theorem 2 shares the same general principle with the proof of Proposition 4, i.e., constructing a credible threat using an information set where both sender collaborate on persuasion. The third condition is a regularity condition which guarantees that the prior p has put enough probability mass in states $\theta \in I_1^x$ and $\theta \in I_1^y$. This is to avoid corner cases where the senders S_1 and S_2 have opposite preferences on the mixture of I_1^x and I_1^y, but their opposite preferences do not matter because the sender S_1 can commit to some signaling strategies which never suggests an action $a_i \in \{a_j | \theta_j \in I_1^x \cup I_1^y\}$. The statement of condition 3 may demand a search of the whole space of possible commitments, which could make this theorem difficult to apply in models with a large state space. However, in most practical problems, verifying condition 3 can be straightforward by inspecting the prior distribution. Hence, we believe condition 3 doesn't undermine the contribution of Theorem 2.

The first two conditions in Theorem 2 capture the scenario when a credible threat by S_2 can be issued, hence constraining S_1's signaling strategies. This leads to the question of whether there are simpler and more basic requirements that make credible threats possible without analyzing the best response of commitments. Moreover, results in Theorem 2 and Proposition 4 depend on prior. We also aim for more basic conditions that are prior independent. Last, we observe that different tie-breaking rules of the receiver may play a role in determining whether the commitment order matters. However, it is hard to determine the criticality of the receiver's tie-breaking rule under sufficient conditions because each tie-breaking rule only plays a role under a small set of priors. Hence, we present the following theorem to state a set of necessary conditions of scenarios where the commitment order could matter.

Theorem 3. *Given a prior p and both senders' utility functions U_{S_1} and U_{S_2}, if both conditions below are violated, then the commitment order does not matter:*

1. *There exists a set of states $(\theta_\alpha, \theta_\beta, \theta_\gamma)$ satisfying the following properties:*

   (a) *θ_α and θ_β are in the same information set of S_1 but θ_γ is not, i.e.,$\exists! I_1^x \ni \theta_\alpha, \theta_\beta$ and $\theta_\gamma \in I_1^y \neq I_1^x$. Besides, θ_α and θ_β are not in the same information set of S_2, i.e., $\exists\ I_2^x \neq I_2^y$ s.t. $\theta_\alpha \in I_2^x$ and $\theta_\beta \in I_2^y$. But θ_γ is in the same information set as either θ_α or θ_β, i.e., $\theta_\gamma \in I_2^x$ or $\theta_\gamma \in I_2^y$. An* *example relationship is depicted in Figure 4.1.*

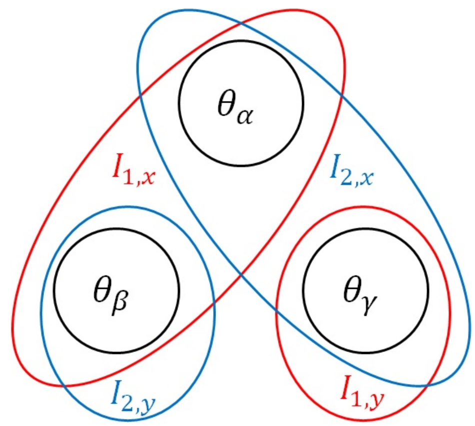

Figure 4.1: An example information structure of a state tuple in condition (a)

(b) The utilities of S_1 and S_2 satisfy

$$U_{S_1}(a_\alpha) > \max\{U_{S_1}(a_\beta), U_{S_1}(a_\gamma)\}, \tag{4.16}$$

$$U_{S_2}(a_\gamma) > U_{S_2}(a_\alpha) > U_{S_2}(a_\beta). \tag{4.17}$$

2. The receiver's tie-breaking rule is belief dependent.

The first condition in Theorem 3 captures a requirement of an alignment between S_1's and S_2's utilities. The requirement is considerably weaker than the definition of collaborative states and condition 1 of Theorem 2. The commitment order may matter when both sender wants to collaborate in a pair of states, but they have a conflict of interest on mixing a third state[5] with this pair. When both senders have no conflict of interests or never collaborate, then the commitment order will not matter. To avoid confusion, we note here that the violation of Theorem 3 does not demand fully-aligned preference orderings between S_1 and S_2, or polar opposite preference orderings between S_1 and S_2. Senders can still collaborate in a set of states and have a conflict of interests in another set of states, as long as these two

[5]This third state is in a different information set of S_1.

sets do not intersect. The second condition states that if the receiver's tie-breaking rule depends on her beliefs, then a sender may tailor her signaling strategies to make the tie-breaking rule in favor of her instead of the other sender. When this occurs, the sender committing first suffers from the tie-breaking rule. The sender committing later can always tailor the receiver's beliefs on top of the earlier sender's commitment and make the tie-breaking rule favor her instead.

4.4 Discussion

We present three examples in this discussion section to demonstrate interesting phenomena. The first example demonstrates an interesting scenario where the commitment order matters. In this scenario, the only reason that drives S_1 to get the first commitment slot is because she can stay silent while committing first. The second example compares sequential commitments with simultaneous commitments. It gives an example where a sequential commitment setting gives the receiver higher expected utility and more informative signals than simultaneous commitments. The third example demonstrates a gap between Proposition 4 and Theorem 2, where the conditions in Proposition 4 are violated but the commitment order still matters. The third example also demonstrates a scenario where transferring the commitment slot increases the utilities of both senders and also the receiver; hence, the social welfare increases. Beside these examples, we remark on the difficulty of developing general algorithms to derive the optimal commitments under incomplete information of senders at the end of this chapter.

4.4.1 Silence is golden: commit to send a non-informative signal

We present an interesting example where commitment order matters, but senders commit completely opposite signaling strategies while committing first. When S_1 commits first, her optimal signaling strategy is either to stay silent or to send non-informative signals. However, when S_2 commits first, her optimal commitment is a truth-telling strategy, and her truth-telling strategy forces S_1's optimal commitment to become truth-telling as well. Hence, in the example detailed below, if the receiver can choose the commitment order, she prefers S_2 committing first.

Example 10. *The state space is ternary* $\Theta = \{\theta_1, \theta_2, \theta_3\}$ *with prior distribution* $(\mathbb{P}(\theta_1), \mathbb{P}(\theta_2), \mathbb{P}(\theta_3)) = \big(0.2, 0.5, 0.3\big)$*.* S_1 *knows the state is* θ_3 *or not, but cannot distinguish* θ_1 *and* θ_2*,* $\mathcal{I}_1 = \{\{\theta_1, \theta_2\},\ \{\theta_3\}\}$*.* S_2 *knows the state is* θ_1 *or not, but cannot distinguish* θ_2 *and* θ_3*,* $\mathcal{I}_2 = \{\{\theta_1\}, \{\theta_2, \theta_3\}\}$*.* S_1 *and* S_2*'s utilities only depend on the receiver's action listed in Table 4.5.*

Table 4.5: Utilities of senders

	a_1	a_2	a_3
(U_{S_1}, U_{S_2})	3,1	0,0	1,3

The receiver wants to match the state and the action's index, and her utility is 1 *if correctly matched and* 0 *otherwise.*

When S_1 commits first, she knows that S_2 cannot distinguish θ_2 and θ_3. Hence, as long as the receiver will guess on a_2 as opposed to a_3 given the posterior belief of the states, S_2 is willing to collaborate with S_1 to suggest a_1. Hence, to maximize S_1's expected utility, her optimal commitment is to stay silent, i.e., commit to a no-informative signal because $\mathbb{P}(\theta_2) > \mathbb{P}(\theta_3)$ in the prior. Given S_1's non-informative commitment, S_2's optimal commitment is to suggest a_1 while she knows the state is θ_1, and suggest a_1 with probability 0.4 when the state is either θ_2 or θ_3 (since S_2 cannot distinguish between θ_2 or θ_3).

When S_2 commits first, she knows S_1 is willing to help her to persuade the receiver to take a_2 when the information set contains only θ_2 and θ_3. Hence, her optimal commitment is, to tell the truth, sending a signal to reveal θ_1 when the true state is θ_1, and sending a signal to tell she knows the state is θ_2 or θ_3 otherwise. Given

S_2's truth-telling commitment and underlying prior, S_1 knows if she doesn't reveal θ_2 when the true state is θ_2, the receiver will take a_3 when the state is θ_2. Hence, her optimal signaling strategy becomes a truth-telling strategy as well.

4.4.2 Comparison to simultaneous commitments

When senders only have partial information (of the state of the world), sequential commitments may bring more informative signals jointly for the receiver. This is different from Li-Norman's result, where senders have complete information. Hence, we want to compare our example in Section 4.3.1 with simultaneous commitments (and simultaneous signaling) from the receiver's perspective. In order to make a fair comparison, the prior and sender's utilities are the same as the example in Section 4.3.1.

Table 4.6: Distribution of the states

	L	R
T	0.1	0.2
B	0.4	0.3

Table 4.7: Utilities of senders

	a_{TL}	a_{TR}	a_{BL}	a_{BR}
(U_{S_1}, U_{S_2})	1,0	2,2	0,0	0,3

From S_2's perspective, since TR is better than TL and BR is better than BL based on her utility function, her commitment will try to persuade the receiver in the best possible manner for taking action TR or BR, regardless what S_1 sends. Hence, S_2's optimal commitment is the following:

- When her private information is R, send a signal to suggest R w.p. 1.
- When her private information is L, send a signal to suggest R with probability $\frac{3}{4}$ and send a suggestion L with probability $\frac{1}{4}$.

From S_1's perspective, even though she cannot observe S_2's commitment and signals, she can anticipate S_2's optimal commitment via S_2's utility function and the prior. Hence, although signals are sent and commitments are made simultaneously, S_1 can use her inference on S_2's optimal commitment to make her commitment. Thus, S_1's commitment is the following:

- When her private information is T, send a signal to suggest T w.p. 1.
- When her private information is B, send a signal to suggest T with probability $\frac{2}{3}$ and send a suggestion B with probability $\frac{1}{3}$.

Given the above commitments, the receiver's utility is calculated below:

$$\mathbb{E}_p[U_R] = 0.1 \times 0 + 0.2 \times 1 + 0.4 \times \tfrac{1}{3} \times \tfrac{1}{4} + 0.3 \times \tfrac{1}{3} = \frac{1}{3}.$$

In Section 4.3.1, the receiver's utility is $\frac{1}{3}$ while S_1 commits first and 0.4 while S_2 commits first. In this example, simultaneous commitment (and signaling) grants the receiver the same utility as the case where S_1 commits first. However, it is still a lower utility compared to the case where S_2 commits first. We note that even though the commitments are made simultaneously, and signals are sent simultaneously, the equilibrium commitment pairs may not be unique when senders only have partial information about the state of the world. The non-uniqueness of equilibrium commitment pairs makes it difficult to make a fair comparison as we now need to include equilibrium selection as well.

4.4.3 The gap between Proposition 4 and Theorem 2

We provide an example where the commitment order matters, but conditions in Proposition 4 are violated. To highlight that a collaborative state is not necessary for S_1 to mix her private information sets. Besides, a conflict of interest between S_1 and

S_2 on a mixture of a pair of S_1's information sets does not demand (polar) opposite preference ordering (on corresponding actions) in any of S_1's private information sets.

The example we provide requires 7 states. However, to simplify the description of S_1's and S_2's private information structures, we add two dummy states with prior probability 0 and expand the example to a 9-state model. To avoid confusion and to simplify the discussion on sender's preference ordering, we skip the definition of S_1's and S_2's utilities on dummy states.

Example 11. *There are $3 \times 3 = 9$ states of the world $\theta_{XY} \in \Theta$, where $X \in \{T, M, B\}$ and $Y \in \{L, C, R\}$. S_1 knows the information of X and S_2 knows the information of Y, i.e., $\mathcal{I}_1 = \{T, M, B\}$ and $\mathcal{I}_2 = \{L, C, R\}$. The prior on the state is provided in Table 4.8, and the utilities of S_1 and S_2 are provided in Table 4.9. To make notations less cumbersome, we abuse the notation and use a_{XY} to represent $a_{\theta_{XY}}$.*

Table 4.8: Distribution of the states

	L	C	R
T	0.05	0	0.05
M	0.12	0.2	0.02
B	0.35	0.21	0

Table 4.9: Utilities of senders

	a_{TL}	a_{TR}	a_{ML}	a_{MC}	a_{MR}	a_{BL}	a_{BC}
(U_{S_1}, U_{S_2})	0,0	5,5	1,0	2,2	0,3	0,1	1,4

In this example, all the conditions in Proposition 4 do not hold. Although there exists a (unique) collaborative state θ_{TR} satisfying condition 1 and we can find a S_1's private information $I_1 = B$ satisfying condition 2, S_1's and S_2's preference order on states $\{\theta_{BL}, \theta_{BC}\}$ are fully aligned, instead of a polar opposite preference required by condition 3. Actually, persuading receiver to take action a_{TR} is best not only for S_1 but also for S_2. Hence, S_2 will not threaten S_1 on her signals suggesting a_{TR}. In short, the signaling strategy on persuading a_{TR} will be independent of the commitment order.

However, the commitment order matters in this example. S_1 prefers a mixture of her private information M, B to persuade the receiver taking action a_{MC}, but S_2 prefers a separation of S_1's private information M and B. Note that θ_{MC} is not a collaborative state because S_2 will first persuade the receiver to take a_{MR}, and then help S_1 persuade the receiver to take a_{MC} when she cannot increase the total probability $\mathbb{P}(a_{MR})$. To verify that the commitment order matters in this example, we calculate the utilities of the senders S_1, S_2 and the receiver under both commitment orders. For simplicity of representation, ω_1^X denotes S_1's signal suggesting the action set $\{a_{XL}, a_{XC}, a_{XR}\}$, $X \in \{T, M, B\}$ and ω_2^Y denotes S_2's signal suggesting the action set $\{a_{TY}, a_{MY}, a_{BY}\}$, $Y \in \{L, C, R\}$. In other words, ω_1^X implies $\mathbb{P}(\theta_{XY}) = \max\{\mathbb{P}(\theta_{TY}), \mathbb{P}(\theta_{MY}), \mathbb{P}(\theta_{BY})\}$ for every $Y \in \{L, C, R\}$ and ω_2^Y implies $\mathbb{P}(\theta_{XY}) = \max\{\mathbb{P}(\theta_{XL}), \mathbb{P}(\theta_{XC}), \mathbb{P}(\theta_{XR})\}$ for every $X \in \{T, M, B\}$. We will use ω_1^X and ω_2^Y to state the commitments in this example.

4.4.3.1 S_1 commits first

First, both S_1 and S_2 wants to maximize the probability of a_{TR} according to their utility functions. Hence, S_1 will first maximize the probability of signal ω_1^T.

- Send ω_1^T with probability 1 when her private information $I_1 = T$.
- Send ω_1^T with probability 0.25 when her private information $I_1 = M$.
- Send ω_1^T with probability $\frac{1}{7}$ when her private information $I_1 = B$.

Now, conditional on S_1 not sending signal ω_1^T, the distribution of states reduces to the following table, where $w = 0.735$ normalizes the total probability to be 1.

According to Table 4.10, S_2 knows the true state is θ_{MR} when her private information $I_2 = R$, regardless what signaling strategy S_1 uses to mix M and B. Since S_2 prefers θ_{MR} than θ_{MC} than θ_{ML}, S_2 will first persuade the receiver taking action a_{MR} when $\omega_1 \neq \omega_1^T$. When S_2 has maximized the total probability of a_{MR} taken by

Table 4.10: Distribution of the states when $\omega_1 \neq \omega_1^T$

	L	C	R
T	0	0	0
M	$0.09/w$	$0.15/w$	$0.015/w$
B	$0.3/w$	$0.18/w$	0

receiver conditional on her signaling strategies, S_2 is willing to collaborate with S_1 on persuading the receiver to take action a_{MC} instead of a_{ML}. Therefore, S_1's optimal mixture of information sets M and B is to maximize the conditional probability of B in her signal, while ensuring that the posterior belief of θ_{MC} is no less than θ_{MB}. Based on the concavification approach in *Kamenica* (2019), the posterior beliefs of states conditional on S_1 sending signal ω_1^M is presented in Table 4.11, where $w = 0.655$ normalizes the total probability to be one.

Table 4.11: Distribution of the states when S_1 sends ω_1^M

	L	C	R
T	0	0	0
M	$0.09/w$	$0.15/w$	$0.015/w$
B	$0.25/w$	$0.15/w$	0

Given S_1's signal which mixes information sets T, M, and B and the optimal mixture of M and B derived above, S_1 still has a positive probability on information set B for which she has no choice but to tell the truth. By calculating the conditional probabilities of each state conditional on S_1's private information and her optimal signaling strategy presented above, we summarize a complete commitment of S_1 below and depict the diagram in Figure 4.2.

- When S_1's private information is T, send ω_1^T with probability 1.
- When S_1's private information is M, send ω_1^T with probability 0.25 and send ω_1^M with probability 0.75.
- When S_1's private information is B, send ω_1^T with probability $\frac{1}{7}$, send ω_1^M with

probability $\frac{5}{7}$ and send ω_1^B with probability $\frac{1}{7}$.

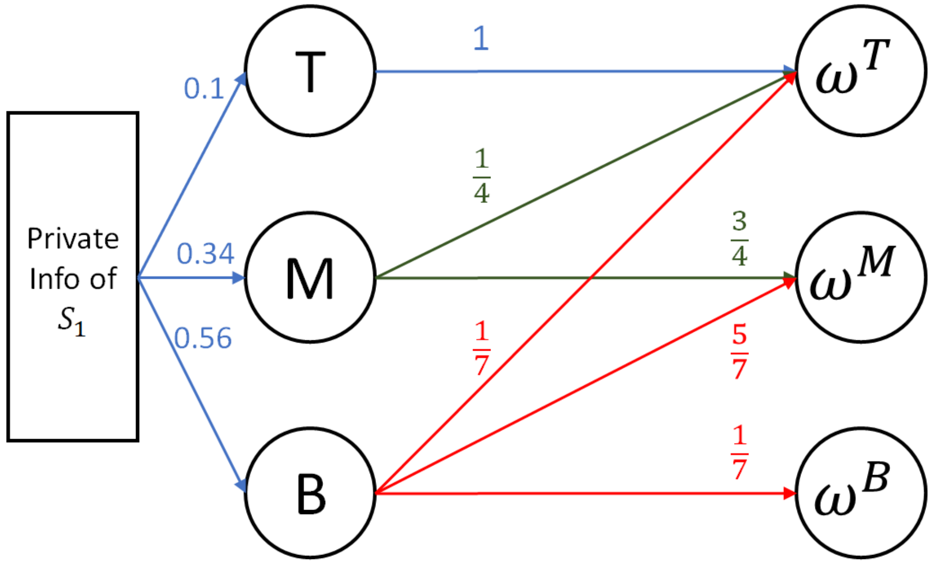

Figure 4.2: S_1's optimal commitment when she commits first

Given S_1's optimal commitment, deriving S_2's optimal commitment is equivalent to solving a classical single-sender Bayesian persuasion problem for each S_1's realized signal when S_2 commits after S_1. Thus, S_2's optimal signaling strategy can be calculated in a straightforward manner via the concavification approach that solves classical Bayesian persuasion problems. Since the result of concavification approach can be easily verified via calculating S_2's utility change under small deviations, we omit the numerical computations of the concavification procedure and present S_2's optimal commitment below:

- While observing ω_1^T, send ω_2^R with probability 1, whatever her private information is.

- While observing ω_1^M, send signals using the following strategy:

  – Send ω_2^R with probability 1 when her private information is R, with probability 0.1 when her private information is C, and with probability 0.06 when her private information is L.

  - Send ω_2^C with probability 0.9 when her private information is C and with probability 0.54 when her private information is L.
  - Send ω_2^L with probability 0.4 when her private information is L.

- While observing ω_1^B, send signals using the following strategy:
  - Send ω_2^C with probability 1 when her private information is C and with probability 0.6 when her private information is L.
  - Send ω_2^L with probability 0.4 when her private information is L.

We want to highlight that S_1 plays an aggressive strategy which relies on S_2's collaboration to persuade the receiver towards taking S_1's preferred action. When S_1 sends ω_1^M, the interim belief of the state distribution is Table 4.11 presented earlier:

We notice that state θ_{BL} has the highest probability if S_2 stays silent. Without the signal of S_2, the receiver will take action a_{BL} and S_1 experiences the lowest utility. However, since a_{MC} gives S_2 a higher utility than a_{ML}, S_2 will collaborate with S_1 after she maximizes the total probability of a_{MR}. This allows S_1 to commit to the above signaling strategy.

Calculating the expected utility of agents, we obtain:

$$\begin{aligned}
\mathbb{E}_p[U_{S_1}] &= (0.05 + 0.05 + 0.03 + 0.05 + 0.005 + 0.05 + 0.03) \times 5 \\
&+ (0.09 \times 0.54 + 0.15 \times 0.9 + 0.25 \times 0.54 + 0.15 \times 0.9) \times 2 \\
&+ (0.25 \times 0.4 + 0.09 \times 0.46) + 0.03 \times 2 = 2.4336 \qquad (4.18) \\
\mathbb{E}_p[U_{S_2}] &= (0.05 + 0.05 + 0.03 + 0.05 + 0.005 + 0.05 + 0.03) \times 5 \\
&+ (0.09 \times 0.54 + 0.15 \times 0.9 + 0.25 \times 0.54 + 0.15 \times 0.9) \times 2 \\
&+ (0.015 \times 4 + 0.09 \times 0.06) \times 3 + 0.03 \times 2 \times 4 + 0.02 \\
&= 2.6884 \qquad (4.19) \\
\mathbb{E}_p[U_R] &= 0.05 + 0.02 \times 0.75 + 0.2 \times 0.75 \times 0.9 + 0.12 \times 0.75 \times 0.36 \\
&+ 0.35 \times \tfrac{1}{7} \times 0.4 + 0.21 \times \tfrac{1}{7} \approx 0.286. \qquad (4.20)
\end{aligned}$$

To summarize, the expected utilities are

$$[\mathbb{E}_p[U_{S_1}], \mathbb{E}_p[U_{S_2}], \mathbb{E}_p[U_R]] = [2.4336, 2.6884, 0.286].$$

4.4.3.2 S_2 commits first

Similarly, both S_1 and S_2 aim to maximize the probability of a_{TR} according to their utility functions. Hence, when S_1 maximizes the probability of signal ω_1^T, S_2 has no incentive to restrict her signal ω_1^T. Therefore, before S_2 makes her commitment, she knows S_1's signaling strategy on ω_1^T below.

- Send ω_1^T with probability 1 when her private information $I_1 = T$.
- Send ω_1^T with probability 0.25 when her private information $I_1 = M$.
- Send ω_1^T with probability $\frac{1}{7}$ when her private information $I_1 = B$.

Now, conditional on S_1 not sending signal ω_1^T, the distribution of states reduces to

Table 4.10 presented earlier. S_2's problem then reduces to designing an optimal commitment under Table 4.10. Opposite to S_1, S_2 prefers a separation of S_1's information sets M and B. Therefore S_2 will restrict S_1's mixture of M and B in her signal ω_1^M as much as she can. Since S_1 will never use a signal ω_1 which mixes M and B such that $\mathbb{P}(MC|\omega_1) \leq \mathbb{P}(BC|\omega_1)$[6], S_2 can only restrict S_1's mixture by tailoring her signal on her private signal L. Ideally, S_2 has two choices: (1) increase the probability of signal ω_2^R while receiving ω_1^M with private signal L, (2) increase the probability of signal ω_2^C while receiving ω_1^M with private signal L. However, since S_1 has utility 0 under a_{MR}, increasing the probability of signal ω_2^R while receiving ω_1^M with private signal L cannot threaten S_1. This is because S_1 is better off when S_2 fails to persuade the receiver to take action a_{MR}[7]. Thus, S_2 can only tailor the probability of signal ω_2^C while receiving ω_1^M with private signal L to restrict S_1. Given the inequality $\mathbb{P}_p(\theta_{BL}) > \mathbb{P}_p(\theta_{BC})$ from the prior, S_2 can reduce the probability $\mathbb{P}(\omega_2^L|\omega_1^M)$ to 0 to restrict S_1's mixture of M and B. In other words, S_1 acknowledges that the receiver will take either a_{MR} or a_{MC} when ω_1^M is sent. Given the interim distribution of Table 4.10, S_2's optimal signaling scheme (after using the concavification approach to derive ω_2^R) is to send signal ω_2^R with probability 0.1 and ω_2^C with probability 0.9 when she observes ω_1^M and her private information is L. Last, when S_2 observes ω_1^B, she immediately knows S_1's private signal is B (because mixing T or M to suggest either a_{BL} or a_{BM} does not benefit S_1). Thus, a concavification approach analogous to the one-sender, binary-state Bayesian persuasion problem solves S_2's optimal signaling strategy in this case. To sum up, we list S_2's optimal signaling strategy below:

- While observing ω_1^T, send ω_2^R with probability 1, whatever her private information is.

- While observing ω_1^M, send signals using the following strategy:

[6] Otherwise, the receiver will never take action a_{MC} suggested by S_1.

[7] In this case, the receiver will take action a_{ML} instead of a_{MR}, and S_1 obtains a higher utility.

- Send ω_2^R with probability 1 when her private information is R, with probability 0.1 when her private information is C, and with probability 0.1 when her private information is L.
- Send ω_2^C with probability 0.9 when her private information is C or L.

- While observing ω_1^B, send signals using the following strategy:
  - Send ω_2^C with probability 1 when her private information is C and with probability 0.6 when her private information is L.
  - Send ω_2^L with probability 0.4 when her private information is L.

Based on S_2's commitment, S_1 knows she cannot commit to an aggressive signaling strategy and wish S_2 to help her persuading the receiver to take action a_{MC}. Moreover, her maximum mixture of M and B is restricted by S_2 with the satisfaction of the inequality $\mathbb{P}(\theta_{BL}|\omega_1^M) \leq \mathbb{P}(\theta_{MC}|\omega_1^M)$. Thus, given the distribution presented in Table 4.10, S_1 can only send signal ω_1^M with probability $\frac{1}{2}$ when her private information is B and signal ω_1^T is not sent. In short, S_1's optimal commitment is to first maximize ω_1^T, and then maximize ω_1^M under the restriction of S_2's signaling strategy. After calculating the probability of ω_1^T and the conditional probability of ω_1^M, we summarize S_1's optimal commitment below and a diagram in Figure 4.3.

- When S_1's private information is T, send ω_1^M with probability 1.
- When S_1's private information is M, send ω_1^T with probability 0.25 and send ω_1^M with probability 0.75.
- When S_1's private information is B, send ω_1^T with probability $\frac{1}{7}$, send ω_1^M with probability $\frac{3}{7}$ and send ω_1^B with probability $\frac{3}{7}$.

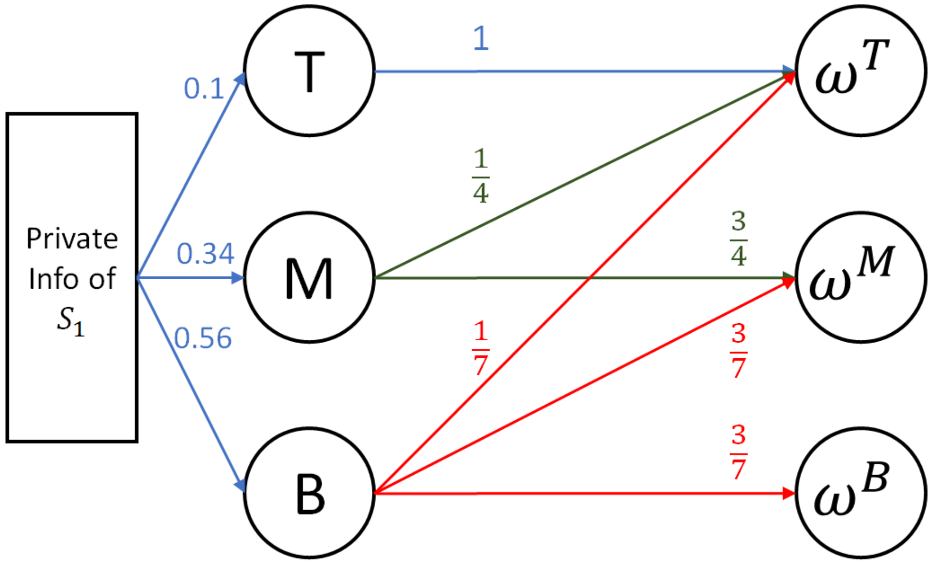

Figure 4.3: S_1's optimal commitment when S_2 commits first

Calculating the expected utility of agents, we obtain:

$$\begin{aligned}
\mathbb{E}_p[U_{S_1}] &= (0.05 + 0.05 + 0.03 + 0.05 + 0.005 + 0.05 + 0.03) \times 5 \\
&+(0.09 \times 0.9 + 0.15 \times 0.9 + 0.15 \times 0.9 + 0.09 \times 0.9) \times 2 \\
&+(0.15 \times 0.6 + 0.09) = 2.207 \qquad (4.21)\\
\mathbb{E}_p[U_{S_2}] &= (0.05 + 0.05 + 0.03 + 0.05 + 0.005 + 0.05 + 0.03) \times 5 \\
&+(0.09 \times 0.9 + 0.15 \times 0.9 + 0.15 \times 0.9 + 0.09 \times 0.9) \times 2 \\
&+(0.015 \times 3 + 0.009 \times 2) \times 3 + 0.09 \times 2 \times 4 + 0.06 = 3.158 \qquad (4.22)\\
\mathbb{E}_p[U_R] &= 0.05 + 0.02 \times 0.75 + 0.2 \times 0.75 \times 0.9 \\
&+0.35 \times \tfrac{3}{7} \times 0.4 + 0.21 \times \tfrac{3}{7} = 0.35. \qquad (4.23)
\end{aligned}$$

To sum up, the expected utilities are $[\mathbb{E}_p[U_{S_1}], \mathbb{E}_p[U_{S_2}], \mathbb{E}_p[U_R]] = [2.207, 3.158, 0.35]$.

4.4.3.3 Comparison of utilities

Similar to the example in Section 4.3.1, both senders want to commit first, and the receiver has a preference on letting S_2 commit first in this example. However, the unique collaborative state in this example violates conditions in Proposition 4 and the signaling strategies (the probability of suggesting a_{TR}) are the same in both commitment orders. The commit order matters because of the collaboration in persuading a_{MC}, even though a_{MC} is not S_2's most preferred action under $I_1 = M$. Moreover, if senders can trade their commitment order and S_1 is initially endowed the first commitment slot, every transfer $x \in (0.2266, 0.4616)$ from S_2 to S_1 makes both senders and the receiver better off by trading the first commitment slot from S_1 to S_2. If the transfer from the receiver to S_1 is also allowed (even though the receiver doesn't participate in this trade), the maximum compensation S_1 could receive from giving out the first commitment slot grows to 0.5276. This example demonstrates a case where both senders' and the receiver's utilities can increase when the first commitment slot is transferred from S_1 to S_2. Besides, Theorem 2 generalizes sufficient conditions when the commitment order matters that not captured in Proposition 4.

4.4.4 The challenge of developing a general algorithm to derive the optimal commitments

Finally, we conclude by remarking that a general algorithm which solves the optimal commitments in information design problems with multiple senders where each sender obtains incomplete private information is challenging to develop. The reason is because the (required) level of inference depends on the information structures given by the problem. Even when the commitment order aligns with the signaling order, the backward-iteration algorithm presented in *Li and Norman* (2018b) cannot solve the optimal commitments when senders do not possess full information (of the state of the world). In general, deriving the optimal commitments demands an iterative reasoning of the other sender's optimal signaling strategy under different interim beliefs of states. The iterative reasoning is straightforward when senders information spaces are binary, but becomes cumbersome when senders' information spaces grows, even from binary to ternary in Example 11. In this Chapter, we present examples where the optimal commitments can be derived and verified via straightforward reasoning to demonstrate our results. We believe deriving a general algorithm which iteratively reason senders' optimal signaling strategies under incomplete information will lead to further significant contributions to the field of information design, but this is beyond the scope of this thesis.

CHAPTER V

Inference Under Exogenous Constraints

5.1 Introduction

In this chapter, we study situations akin to the following stylized dialogue, which will likely be familiar to anyone who has ever served on hiring committees:

ALICE: I see that you wrote strong recommendation letters for your Ph.D. graduates Carol and Dan. Can you compare them for us?

BOB: They are both great! Carol made groundbreaking contributions to ...; Dan made groundbreaking contributions to ....

ALICE: Which of the two would you say is stronger?

BOB: They are hard to compare. You really need to interview both of them!

ALICE: We can only invite one of them for an interview.

BOB: I guess Carol is a bit stronger.

What happened in this example? Alice and Bob were involved in a signaling setting, in which Bob had an informational advantage. Bob's goal was to get as many of his students interviews as possible, while Alice's goal was to only invite the strong students. While Bob knew which of his students were strong (or how strong), Alice had to rely on the information she could obtain from Bob. As is standard in signaling settings, Bob could use this fact to improve his own utility. In this sense, the example initially was virtually identical to the standard "judge/prosecutor" example in the seminal paper of *Kamenica and Gentzkow* (2011).

However, a change happened along the way. When Alice revealed that she was constrained in her actions (to one interview at most), this changed the utility that Bob could obtain from his previous strategy. For example, if he had insisted on not ranking the students, Alice might have flipped a coin. Implicitly, while Bob wanted *both* of his students to obtain interviews, when forced to choose, he knew he would obtain higher utility from the stronger of his students being interviewed. In this sense, his utility function was "partially aligned" with Alice's; this partial alignment, coupled with Alice's constraint, resulted in Alice obtaining more information, and thus higher utility.

The main goal of this chapter is to investigate to what extent the behavior illustrated informally in the dialogue above arises in a standard model of Bayesian persuasion. Specifically, if the utilities of the sender and receiver are "partially aligned," will it always benefit a receiver to be more constrained in how she can choose her actions?

5.1.1 The model: an overview

Our model — described in full detail in Section 5.3 — is based on the standard Bayesian persuasion model of *Kamenica and Gentzkow* (2011). For our main result, we assume that the state space is binary: $\Theta = \{\theta_1, \theta_2\}$. These states could correspond to a student being bad/good in our introductory example, a defendant being innocent/guilty in the example of *Kamenica and Gentzkow* (2011), or a stock being about to go up or down. The sender and receiver share a common prior p for the distribution of the state θ. In addition, the sender will observe the actual state θ, but only after committing to a signaling scheme (also called information structure).

A signaling scheme is a (typically randomized) mapping $\Gamma : \Theta \to \Omega$. The receiver observes the (typically random) signal $\omega = \Gamma(\theta)$; based on this observation, she takes an action $a \in A$. Here, we assume that — like the state space — the action space is binary, i.e., $A = \{a_1, a_2\}$. Based on the true state of the world and the action taken by the receiver, both the sender and receiver derive utilities $U_S(\theta, a)$, $U_R(\theta, a)$. The receiver will choose her action (upon observing ω) to maximize her own expected utility; the sender, knowing that the receiver is rational, will commit to a signaling scheme to maximize his expected utility under rational receiver behavior.

Motivated by many practical applications, we assume that the receiver prefers to match the state of the world, in the sense that $U_R(\theta_1, a_1) \geq U_R(\theta_1, a_2)$ and $U_R(\theta_2, a_2) \geq U_R(\theta_2, a_1)$. For instance, in our introductory example, Alice prefers to interview strong candidates and to not interview weak ones; in the judge-prosecutor example, the judge prefers convicting exactly the guilty defendants; and an investor prefers to buy stocks that will go up and sell stocks that will go down. Our assumption about the "partial alignment" of the sender and receiver utilities is formalized as an *action-matching* preference of the sender, stated as follows: $U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2)$. That is, if a candidate is being interviewed, Bob prefers it to be a strong candidate over a weak one (but may still prefer a weak candidate being interviewed over a strong/weak candidate *not* being interviewed); similarly, if a prosecutor sees a defendant convicted, he would prefer the defendant to be guilty (but may still prefer an innocent candidate being convicted over an innocent defendant going free); similarly, an investment platform may want to entice a client to buy stock, but conditioned on the client buying stock, the platform may prefer for the stock to go up.

In addition to the assumption of partial alignment, our main addition to the standard Bayesian persuasion model is to consider constraints on the receiver's actions. Specifically, we assume that there are (lower and upper) bounds $\underline{b}$, $\overline{b}$ on the probability with which the receiver is allowed to take action[1] a_1. Such a constraint corresponds to a department only being willing to interview at most 10% of their applicants, a judge having a quota for how many defendants (at most) to convict, or a conference having an upper bound on its number/fraction of accepted papers. Such a constraint creates dependencies between the receiver's actions under different received signals, and may force her to randomize between actions, contrary to the standard Bayesian persuasion setting in which the receiver may deterministically choose any utility-maximizing action conditioned on the observed signal ω. To see this, consider a prior under which a candidate is strong with probability $\frac{1}{3}$, and the receiver obtains utility 1 from interviewing a strong candidate and -1 from interviewing a weak candidate (and 0 from not interviewing). If the sender reveals no information, the receiver would prefer to interview no candidates, but a lower-bound constraint may force her to do so, in which case she would randomize the decision to interview the smallest total number of candidates. We write $a^* : \Omega \to A$ for the receiver's (typically randomized) mapping from signals to actions. Note that the constraint applies across all sources of randomness (the state of the world, the sender's randomization, and the receiver's randomization), so it is required that $\underline{b} \leq \mathbb{P}_{p,\Gamma,a^*}(a^*(\Gamma(\theta)) = a_1) \leq \overline{b}$.

To avoid trivialities, we assume that $\mathbb{P}_p(\theta = \theta_1) \in [\underline{b}, \overline{b}]$, that is, if the sender revealed the state of the world perfectly, the receiver would be allowed to match it. We say that a receiver with constraints $(\underline{b}', \overline{b}')$ is *more constrained* than one with constraints $(\underline{b}, \overline{b})$ iff $\underline{b}' \geq \underline{b}$ and $\overline{b}' \leq \overline{b}$.

5.1.2 Our results

Our main result is that when the state of the world is binary, a receiver is always (weakly) better off when more constrained. We state this result here informally, and revisit it more formally in Section 5.4.

[1] This naturally implies constraints of $1 - \overline{b}, 1 - \underline{b}$ on the probability of taking action a_2. We discuss a more general model and its specialization to binary actions in Section 5.3.3.

Theorem 4 (stated informally)**.** *Consider a Bayesian persuasion setting in which the state and action spaces are binary, the receiver is trying to match the state of the world, and the sender is action-matching. Then, a more constrained receiver always obtains (weakly) higher utility than a less constrained one.*

Unfortunately, this insight does not extend to more fine-grained states of the world: even for a ternary state of the world, there are examples with partially aligned sender and receiver in which a more constrained receiver is strictly worse off. We discuss such an example in depth in Section 5.5. It is possible to obtain some positive results recovering versions of Theorem 4 by imposing additional constraints on the sender's and receiver's utility functions. However, many of these constraints are strong, and have only limited applicability to real-world settings. We discuss some of these approaches in Section 5.6 — whether there are less restrictive conditions recovering Theorem 4 for more states of the world is an interesting direction for future work.

5.2 Related Work

In general, information design as an area is concerned with situations in which a better-informed sender or information designer can influence the behavior of other agents via the provision of information. The literature generally studies problems in which the underlying game between the agents is given and fixed, but where the sender can influence the outcome by an appropriate choice of information to be disclosed. The core difference between Bayesian persuasion *Kamenica and Gentzkow* (2011); *Bergemann and Morris* (2013, 2016); *Kamenica* (2019); *Bergemann and Morris* (2019) and other standard paradigms that study information transmission (such as cheap talk *Crawford and Sobel* (1982), verifiable messages *Grossman* (1981); *Milgrom* (1981) or signaling games *Spence* (1978)) is the commitment power of the sender.

In Bayesian persuasion models, the sender moves first and commits to a (typically randomized) mapping from states of the world to signals. Subsequently, the sender observes the state of the world and applies the mapping. Based on the mapping and the observed signal, the rational recipients (called receivers) choose actions.

The study of Bayesian persuasion was initiated in the seminal work of *Kamenica and Gentzkow* (2011) and *Rayo and Segal* (2010). In their work, the sender can *commit to sending any distribution of messages* before (accurately) observing the state of the world; the receiver, on the other hand, only has knowledge of the prior. The full commitment setting allows for an equivalence to an alternate model where the sender publicly chooses the amount of information (regarding the state of the world) he will privately observe and then (strategically) decides how much of this information to share with the receiver via verifiable messages. Follow-up work in *Bergemann and Morris* (2013, 2016) established a useful and important equivalence between the set of outcomes achievable via information design and Bayes correlated equilibria. Since these seminal works, there has been a large body of work on Bayesian persuasion with theoretical developments as well as a multitude of applications. To keep our discussion focused, for the broader literature, we refer the reader to two recent survey articles (*Kamenica*, 2019; *Bergemann and Morris*, 2019).

The literature closest to our work studies information design with a constrained sender: the constraints arise through a diverse set of assumptions. The work in *Perez-Richet and Prady* (2012); *Perez-Richet* (2014) shows that pooling equilibria result if the receiver either prefers lower complexity (for a certification process) or performs a validation of the sender's signal; this holds whether the signals of the sender are exogenously constrained or not. A growing body of work considers constraints on the sender that arise either due to communication costs for signaling *Gentzkow and Kamenica* (2014); *Hedlund* (2015); *Carroni et al.* (2020); *Nguyen and Tan* (2021), capacity limitations for signaling *Le Treust and Tomala* (2019), the sender's signal serving multiple purposes (such as convincing a third party to take a payoff-relevant action) *Boleslavsky and Kim* (2018), or costs to the receiver for acquiring additional information *Matyskova* (2018). The contributions are then to characterize either the applicability of the concavification approach of *Kamenica and Gentzkow* (2011), the optimal signaling structure, or the conditions for the optimality of certain signaling structures. In *Kolotilin* (2015), constraints on the sender arise from the receiver having access to some publicly available information. Within this context, *Kolotilin* (2015) studies comparative statics on the sender's utility based on the quality of the sender's information or the public information. There is also a burgeoning literature on constraints on the sender arising from a privately informed receiver (e.g., *Kolotilin et al.* (2017); *Guo and Shmaya* (2019); *Doval and Skreta* (2018); *Candogan* (2020); *Candogan and Strack* (2021). The main contributions in this line of research are to characterize the optimal signaling structure with a key aspect being the fact that the sender constructs a different signal for each receiver type.

Based on the discussion above, clearly there is significant literature studying a constrained sender's optimal signaling scheme and utility. However, work that studies constraints on the *receiver*, or their impact on the receiver's utility, is extremely limited. To the best of our knowledge, *Babichenko et al.* (2020) is the only work to analyze a constrained receiver problem. The authors impose *ex ante* and *ex post* constraints on the receiver's posterior beliefs, characterize the dimensionality of the optimal signaling structure and develop low-complexity approximate welfare maximizing algorithms. In our work, we have two important differences: first, we impose constraints on the receiver's *actions* as opposed to posterior beliefs; and second, we explore when these constraints result in increased utility for the receiver.

5.3 Problem Formulation

Our model is based on the standard Bayesian persuasion model of *Kamenica and Gentzkow* (2011). Two players, a sender and a receiver, interact in a signaling game. The sender can observe the state of the world, while the receiver can take an action. The sender can convey information about the state of the world to the sender. Both players receive utility as a function of both the state of the world and the action chosen by the receiver. Since their utility functions typically do not align, the sender will be strategic in the information he reveals to the receiver.

5.3.1 State of the World, Actions, and Utilities

The (random) state of the world θ is drawn from a state space Θ. For our main result, we assume that the state space is binary ($\Theta = \{\theta_1, \theta_2\}$); however, we define the model in more generality. The sender and receiver share a common-knowledge prior distribution $p \in \Delta(\Theta)$ for θ. When the state space is binary, this prior is fully characterized by $p = \mathbb{P}_p(\theta = \theta_1)$.

Only the receiver can take an action $a \in A$. Again, for our main result, we assume that the action space is binary: $A = \{a_1, a_2\}$. Both the sender's and the receiver's utilities are functions of the true state θ and the action taken; they are captured by the functions $U_S : \Theta \times A \to \mathbb{R}$ and $U_R : \Theta \times A \to \mathbb{R}$. As discussed in Section 5.1.1, we assume that the receiver tries to match the state of the world with her action.

Definition 32 (State-Matching Receiver)**.** *We say that the receiver's utility function is* state-matching *if it satisfies the following: for all* i, j, k *with* $i \leq j \leq k$ *or* $i \geq j \geq k$, *we have that*

$$U_R(\theta_i, a_j) \geq U_R(\theta_i, a_k). \tag{5.1}$$

When the state of the world is binary, the condition simplifies to:

$$U_R(\theta_1, a_1) \geq U_R(\theta_1, a_2) \text{ and } U_R(\theta_2, a_2) \geq U_R(\theta_2, a_1). \tag{5.2}$$

In words, a state-matching receiver always prefers an action closer to the true state of the world; however, the definition does not enforce any comparisons between choosing an action that is too high vs. too low compared to the true state.

The key notion for our analysis is a partial alignment of the sender's utility with the receiver's. This is captured by the fact that the sender, given any fixed action, would prefer states closer to the action, expressed in the following definition:

Definition 33 (Action-Matching Sender)**.** *We say that the sender's utility function is* action-matching *if it satisfies the following: for all i, j, k with $i \leq j \leq k$ or $i \geq j \geq k$, we have that*

$$U_S(\theta_j, a_i) \geq U_S(\theta_k, a_i). \tag{5.3}$$

When the state of the world is binary, the condition simplifies to:

$$U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1) \text{ and } U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2). \tag{5.4}$$

In words, an action-matching sender always prefers a state of the world closer to the action chosen by the receiver; again, we do not enforce any comparisons between states that are higher vs. lower than the chosen action.

Notice the difference between Inequalities (5.3) and (5.4) vs. (5.1) and (5.2): (5.1) and (5.2) compare the receiver's utilities when the state of the world is fixed and the action is changed, while (5.3) and (5.4) compare the sender's utilities when the action is fixed and the state of the world is changed. That is, given that the receiver takes a particular action, the sender derives higher utility when that action more closely matches the state of the world than when it does not. Again, a justification for this assumption is discussed in Section 5.1.1.

5.3.2 Signaling Schemes

Before the receiver takes her action, the sender can send a signal ω to reveal (partial) information about the state of the world. More precisely, prior to observing the state of the world θ, the sender commits to a signaling scheme Γ, which is a mapping $\Gamma : \Theta \to \Delta(\Omega)$. For our purposes Γ is conveniently characterized by the probability with which each signal is sent conditional on the state. We write $\mathbb{P}(a_j|\theta_i) = \mathbb{P}(\Gamma(\theta) = \omega_j|\theta = \theta_i) \in [0, 1]$ for the probability that signal ω_j is sent conditional on the state of the world being θ_i. We write $\overline{\Gamma}_j = \sum_i \mathbb{P}_p(\theta = \theta_i) \cdot \mathbb{P}(a_j|\theta_i)$ for the probability of sending the signal ω_j.

The receiver is Bayes-rational, and her objective is to maximize her expected utility after observing the signal. The expected utility derived from action a when observing ω_j can be written as

$$\mathbb{E}[U_R(\omega_j, a)] = \sum_{\theta_i \in \Theta} \mathbb{P}(\theta = \theta_i|\Gamma(\theta) = \omega_j) \cdot U_R(\theta_i, a) \;=\; \sum_{\theta_i \in \Theta} \tfrac{\mathbb{P}(\theta=\theta_i)\cdot\mathbb{P}(a_j|\theta_i)}{\overline{\Gamma}_j} \cdot U_R(\theta_i, a).$$

Thus, barring other constraints (which we will introduce below), the receiver chooses an action a in $\arg\max_a \mathbb{E}[U_R(\omega_j, a)]$. Following most of the literature in the field of information design, we assume that the receiver breaks ties in favor of an action most preferred by the sender. The following very useful alternative view has been observed in the prior literature (see, e.g., (*Bergemann and Morris*, 2016)): instead of sending abstract signals, the sender can without loss of generality send the receiver a recommended action a_j. The sender must ensure that Γ is such that the receiver will always voluntarily follow the recommendation. In other words, the recommended action a_j must always be in $\arg\max_a \mathbb{E}[U_R(\omega_j, a)]$. This constraint ensures *ex-post* *incentive compatibility (EPIC)* of the signaling scheme, and is often referred to as an *obedience constraint*.

We write $a^* : \Omega \to \Delta(A)$ for the receiver's (possibly randomized) best-response function. In the setting described so far, there is actually no need for the receiver to randomize, and she can always choose any arbitrary deterministic $a^*(\omega_j) \in \arg\max_a \mathbb{E}[U_R(\omega_j, a)]$. However, as we will see in Section 5.3.3, the situation changes when the receiver is constrained. For a receiver strategy a^*, we write $\pi_{i,j} = \mathbb{P}(a^*(\omega_j) = a_i)$ for the probability that the receiver, upon observing signal ω_j, chooses action a_i.

The sender's objective is to design a signaling strategy which maximizes his expected utility in the subgame perfect equilibrium. That is, he chooses Γ so as to maximize his expected utility (under all sources of randomness)

$$\mathbb{E}_{\theta\sim p,\omega\sim\Gamma(\theta),a\sim a^*_\Gamma(\omega)}[U_S(\theta, a)],$$

assuming a best response a^*_Γ from the receiver.

5.3.3 Constrained Receiver

Our main conceptual departure from prior work is that we consider constraints on the receiver, restricting the probability with which actions can be chosen. In a general setting, such constraints are lower and upper bounds on the probability of taking each action, i.e., $\underline{b}_a$ and $\overline{b}_a$ for each a. Formally, we require that for each action a_i, the combination of the sender's signaling scheme Γ and the receiver's response a^* satisfy

$$\underline{b}_{a_i} \leq \sum_j \overline{\Gamma}_j \cdot \pi_{i,j} \ \leq \ \overline{b}_{a_i}. \tag{5.5}$$

The constraints are common knowledge among the sender and receiver. When the state space is binary, the constraints can be simplified: they are fully characterized by the lower and upper bounds $\underline{b} = \max(\underline{b}_{a_1}, 1 - \overline{b}_{a_2}), \overline{b} = \min(\overline{b}_{a_2}, 1 - \underline{b}_{a_1})$ for the probability with which the receiver can choose action a_1.

The focus of our work is on whether being (more) constrained helps the receiver, by forcing an action-matching sender to disclose "more" information. Without any further assumptions, this is trivially false. For example, suppose that the state of the world is known to be θ_1 with probability 1, and both the sender and the receiver obtain utility 1 when the receiver chooses action a_1, and 0 otherwise. If the constraint specified that a_1 must be taken with probability 0, and a_2 with probability 1, then of course, the receiver (and the sender) would be worse off. In order to allow us to clearly articulate the question of whether a constrained receiver obtains more information, we require that perfect state matching would always be feasible for the receiver, if the true state were revealed:

Definition 34 (Implementable and Feasible Constraints)**.** *Consider a set of constraints $\langle \underline{b}_{a_i}, \overline{b}_{a_i} \rangle$ for all $a_i \in A$. We say that the constraints are* implementable *if and only if $\sum_i \underline{b}_{a_i} \leq 1 \leq \sum_i \overline{b}_{a_i}$. The constraints are* feasible *if and only if $\underline{b}_{a_i} \leq \mathbb{P}_p(\theta = \theta_i) \leq \overline{b}_{a_i}$ for all i. For the special case of a binary state space, a constraint $\langle \underline{b}, \overline{b} \rangle$ is feasible if and only if $\underline{b} \leq p \leq \overline{b}$.*

Notice that when constraints are not implementable, there is no strategy for the receiver to satisfy all constraints. When constraints are feasible, then with full information, perfect state matching can be implemented by the receiver.

We say that the constraints $\langle \underline{b}_{a_i}, \overline{b}_{a_i} \rangle$ are *more binding* (or the receiver is more constrained by them) than $\langle \underline{b}'_{a_i}, \overline{b}'_{a_i} \rangle$ if and only if $\underline{b}'_{a_i} \leq \underline{b}_{a_i}$ and $\overline{b}_{a_i} \leq \overline{b}'_{a_i}$ for all i. When the state space is binary, the condition simplifies: the constraint $\langle \underline{b}, \overline{b} \rangle$ is more binding than $\langle \underline{b}', \overline{b}' \rangle$ if and only if $\underline{b}' \leq \underline{b}$ and $\overline{b} \leq \overline{b}'$.

We note that the presence of a constraint may force the receiver to randomize between actions, even possibly actions that are not optimal. For a simple example, suppose that the state of the world is binary and determined by a fair coin flip, and the receiver obtains utility 2 from matching state θ_2, 1 from matching state θ_1, and 0 for not matching the state. If the sender reveals no information, then a receiver constrained by — say — $\underline{b} = \overline{b} = \frac{1}{2}$, would have to flip a fair coin to decide which action to choose, even though the optimal strategy would be to always choose a_2.

While the receiver's best response a^* may in general (have to) be randomized, we show that there is always an optimal signaling strategy for the sender such that the receiver will play a deterministic strategy a^*. Notice that the following proposition does not even require feasibility in the sense that the prior distribution satisfies the constraints: it merely requires that the constraints allow for the existence of *any* signaling scheme and corresponding receiver strategy.

Proposition 5. *Assume that* $|\Omega| \geq |A|$, *and let* $\langle \underline{b}_{a_i}, \overline{b}_{a_i} \rangle$ *(for all* i*) be implementable constraints on the receiver. Then, for any signaling scheme* $\hat{\Gamma}$*, there exists another signaling scheme* Γ *under which the sender has at least the same utility as under* $\hat{\Gamma}$*, and such that the receiver's best response* a^*_Γ *is deterministic. In particular, there is a sender-optimal strategy under which the receiver responds deterministically.*

The proof of Proposition 5 is by an explicit construction of such a strategy Γ where the receiver's best response is deterministic. The detailed construction is provided in Appendix E.1.1. In general, most of the literature on Bayesian persuasion assumes that the signal space is at least as large as the action space (which is enough to obtain sender-optimal strategies, and find them via an LP (*Kolotilin*, 2018) when EPIC holds). Hence, we make the same assumption that $|\Omega| \geq |A|$ in Proposition 5.

Henceforth, we will restrict attention to signaling schemes with deterministic best response functions a^* without loss of optimality. However, the sender still has to ensure that following the deterministic recommendation is incentive compatible for the receiver. Since the receiver is constrained, her space of deviations is only to best-response functions satisfying the constraints. This is captured by the following definition:

Definition 35. *Let $\Gamma : \Theta \to \Omega$ be a direct signaling scheme for the sender, i.e., making action recommendations and assuming $\Omega = A$. Let A^* be the set of all randomized mappings $a^* : \Omega \to A$ (characterized by $\pi_{i,j}$) satisfying the following inequalities for all actions a_j:*

$$\underline{b}_{a_j} \leq \sum_i \overline{\Gamma}_i \cdot \pi_{i,j} \ \leq \ \bar{b}_{a_j}.$$

Then, Γ is ex ante incentive compatible *if and only if for all feasible response functions $a^* \in A^*$,*

$$\sum_i \overline{\Gamma}_i \cdot \mathbb{E}[U_R(\omega_i, a_i)] \geq \sum_i \sum_j \overline{\Gamma}_i \cdot \pi_{i,j} \cdot \mathbb{E}[U_R(\omega_i, a_j)].$$

Note that the presence of constraints forces us to deviate from the standard EPIC requirement in the literature. Definition 35 bears similarity to definitions in (*Babichenko et al.*, 2020; *Doval and Skreta*, 2018; *Candogan and Strack*, 2021), where ex ante constraints are considered.

5.4 Our Main Result

In this section, we present the main result of this paper.

Theorem 5. *Consider a Bayesian persuasion setting in which the state and action spaces are binary. The receiver is state-matching, and the sender is action-matching. Let $\langle \underline{b}, \overline{b} \rangle$ and $\langle \underline{b}', \overline{b}' \rangle$ be two feasible constraints such that $\langle \underline{b}, \overline{b} \rangle$ is more binding than $\langle \underline{b}', \overline{b}' \rangle$, and let Γ, Γ' be the sender's optimal signaling schemes under these constraints.*

Then, the receiver is no worse off under Γ than under Γ'.

Proof. At a high level, the intuition for the proof is as follows. Based on the discussion in Section 5.3.3, the constraints on the receiver actually translate into constraints on the sender in the optimization problem. Because the sender's signaling schemes are more constrained, he has to reveal more information. However, this intuition is not complete — after all, the constraints may entice the sender to reveal *less* information. Furthermore, as we see in Section 5.5, when the state space is not binary, a more constrained receiver may be worse off.

Let Γ, Γ' be the sender-optimal signaling schemes under the two constraints, and let $\mathbb{P}(a_j|\theta_i)$, $\mathbb{P}_{'}(a_j|\theta_i)$ be their corresponding conditional probabilities of sending the signal ω_j in state θ_i. By Proposition 5, w.l.o.g., under the sender-optimal strategies Γ, Γ', the sender recommends an action to the receiver, and the receiver deterministically follows the recommendation. That is, the signal ω_i can be associated with the action a_i for $i = 1, 2$. Our proof is based on distinguishing four cases, depending on the sender's utility:

1. $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_2, a_1)$

   In this case, for every state, the sender prefers the same action as the receiver. Since the sender's and the receiver's preferences are fully aligned, the sender's optimal strategy is to fully reveal the state of the world. Since the constraints are feasible, the receiver can perfectly match the state of the world under both constraints, and hence obtains the same utility under both constraints.

2. $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \leq U_S(\theta_2, a_1)$

   In this case, the sender always prefers action a_1. Since the sender is action-matching, $U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2)$. Combining these inequalities, we obtain that the sender's utility function satisfies the following total order:

$$U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2).$$

This implies that

$$U_S(\theta_1, a_1) - U_S(\theta_1, a_2) \geq U_S(\theta_2, a_1) - U_S(\theta_2, a_2). \tag{5.6}$$

We now show that $\mathbb{P}(a_2|\theta_1) = 0$. An identical proof also shows that $\mathbb{P}'(a_2|\theta_1) = 0$. We distinguish two cases:

- If $\mathbb{P}(a_2|\theta_1) > 0$ and $\mathbb{P}(a_1|\theta_2) > 0$, then the sender could move some probability mass $\epsilon > 0$ from recommending a_2 under θ_1 to recommending a_1, and in return move the same amount from recommending a_1 under θ_2 to recommending a_2. Because the receiver is state-matching, she will still follow the sender's recommendation, and the total probability with which each action is played stays unchanged, so the strategy is still feasible. By Equation (5.6), the sender's utility (weakly) increases. By choosing ϵ as large as possible, we arrive at the claim or at the following case.
- If $\mathbb{P}(a_2|\theta_1) > 0$ and $\mathbb{P}(a_1|\theta_2) = 0$, then $\overline{\Gamma}_{a_1} = p \cdot \mathbb{P}(a_1|\theta_1) < p \leq \bar{b}$. Therefore, it is feasible for the sender to always send the signal a_1 when the state is θ_1 (i.e., decrease $\mathbb{P}(a_2|\theta_1)$ to 0 and increase $\mathbb{P}(a_1|\theta_1)$ by the same amount). Again, because the receiver is state-matching, she will still follow the sender's recommendation, and because $U_S(\theta_1, a_1) \geq U_S(\theta_1, a_2)$, the sender is weakly better off.

Because $U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2)$ and $\mathbb{P}(\Gamma(\theta_1) = a_1) = 1$ proved above, the sender will also send a_1 with as much probability as possible when the state is θ_2, subject to not violating the receiver's incentive to play a_1 and not exceeding the upper bound $\bar{b}$ (or $\bar{b}'$). In other words, the sender maximizes $\mathbb{P}(\Gamma(\theta_2) = a_1)$ subject to $\mathbb{P}(\theta_1|\overline{\Gamma}_{a_1}) \geq \mathbb{P}(\theta_2|\overline{\Gamma}_{a_1})$ and $\bar{b} \geq \overline{\Gamma}_{a_1}$ (or $\bar{b}' \geq \overline{\Gamma}_{a_1}$). Since $\mathbb{P}(\Gamma(\theta_1) = a_1) = 1$, the incentive constraint is independent of the bound $\bar{b}$ (or $\bar{b}'$), while the latter is more restricted for $\bar{b} \leq 2p$ (or $\bar{b}' \leq 2p$). Therefore, given $\bar{b} \leq \bar{b}'$, the receiver is weakly better off under the constraint $\bar{b}$ than under $\bar{b}'$.

3. $U_S(\theta_1, a_1) \leq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_2, a_1)$

4. $U_S(\theta_1, a_1) \leq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_2, a_1)$

   This case is symmetric to the previous one. Here, the roles of a_1 and a_2 (and θ_1 and θ_2) are reversed, and the important constraint becomes the lower bound $\underline{b}$ (and $\underline{b}'$) rather than the upper bound $\bar{b}$.

5. $U_S(\theta_1, a_1) \leq U_S(\theta_1, a_2)$ and $U_S(\theta_2, a_2) \leq U_S(\theta_2, a_1)$

   In this case, the fact that the sender is action-matching together with the assumed inequalities implies that

$$U_S(\theta_2, a_2) \overset{\text{AM}}{\geq} U_S(\theta_1, a_2) \geq U_S(\theta_1, a_1) \overset{\text{AM}}{\geq} U_S(\theta_2, a_1) \geq U_S(\theta_2, a_2).$$

   Thus, the sender's utility is the same, regardless of the state and action. As a result, the sender is indifferent between all signaling schemes. In particular, fully revealing the state is an optimal strategy for the sender for any constraint.

Thus, for all four cases, the receiver will be no worse off under the more binding constraint. □

5.5 Failure of the main result with larger state spaces

Unfortunately, contrary to the case of binary state and action spaces, when the state and action spaces are larger, a state-matching receiver and action-matching sender (and feasible constraints) are not enough to ensure that the receiver is always better off when more constrained. Consider the utilities given in Table 5.1. There are three states in the world, and correspondingly three actions. The prior over the states is uniform.

	θ_1	θ_2	θ_3
a_1	10	10	0
a_2	0	2	2
a_3	0	0	1

(a) Sender's Utility

	θ_1	θ_2	θ_3
a_1	4	2	0
a_2	0	3	1
a_3	0	1	3

(b) Receiver's Utility

Table 5.1: Sender's and Receiver's Utility in the example where a constrained receiver is worse off

Notice that the receiver is state-matching, and the sender is action-matching.

Unconstrained Receiver. First, consider an unconstrained receiver. The sender's optimal signaling scheme Γ is to recommend action a_1 whenever the state of the world is θ_1 or θ_2, and recommend action a_3 otherwise.

To verify that the receiver follows the recommendation, one simply compares the utility from the alternative actions: when the sender recommends a_1, following the recommendation gives the receiver expected utility $\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 2 = 3$, while action a_2 would give utility $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 3 = \frac{3}{2}$, and a_3 would give $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot = \frac{1}{2}$. For the recommendation of a_3, the receiver gets to match the state deterministically, so following the recommendation is optimal. Because the signaling scheme is even ex post incentive compatible for the receiver, it is most definitely ex ante incentive compatible.

To see that this signaling scheme is optimal for the sender, first observe that for states θ_1 and θ_2, the sender obtains the maximum possible utility of 10 over all actions. For state θ_3, the sender would prefer the receiver to play action a_2. However, the only way to get the sender to play a_2 is to mix at least one unit of probability of θ_2 per unit of probability of θ_3. While this increases the sender's utility for the unit of probability from θ_3 from 1 to 2, it decreases his utility for the unit of probability from θ_2 from 10 (since the receiver played a_1) to 2. Thus, the given signaling scheme is sender-optimal.

Under this signaling scheme, the receiver's expected utility can be calculated as $\frac{2}{3} \cdot (\frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 2) + \frac{1}{3} \cdot 3 = 3$.

Adding a non-trivial constraint. Now, consider a receiver constrained by an upper bound $\bar{b}_{a_1} = \frac{1}{2}$. Table 5.2 shows the sender-optimal signaling scheme. Here, the entries show the conditional probability $\mathbb{P}(a_j|\theta_i)$ of recommending action a_j (i.e., sending signal ω_j) when the state is θ_i.

	θ_1	θ_2	θ_3
ω_1	1	$\frac{1}{2}$	0
ω_2	0	$\frac{1}{2}$	$\frac{1}{2}$
ω_3	0	0	$\frac{1}{2}$

Table 5.2: Sender-optimal signaling scheme when $\bar{b}_{a_1} = \frac{1}{2}$

First, notice that action a_1 is recommended with probability $\frac{1}{2}$, so the constraint is satisfied. Second, the receiver will follow the sender's recommendation, as can be checked by comparing her utility from each of the three actions conditioned on any signal. (In the case of receiving ω_2, she is indifferent between a_2 and a_3 — recall that we assume tie breaking in favor of the sender.) Again, the given signaling scheme is even ex post incentive compatible, so in particular, it is also ex ante incentive compatible.

To see that the signaling scheme is optimal for the sender, first notice that he induces action a_1 (under states θ_1 or θ_2) with the maximum probability of $\frac{1}{2}$. Also, notice that using all of the probability from θ_1 to induce a_1 is optimal for the sender, because under θ_1, if any action other than a_1 is played, the sender's utility is 0. Because $\frac{1}{6}$ unit of probability from θ_2 yields a recommendation of a_1, at most $\frac{1}{6}$ can yield a recommendation of a_2, which gives the next-highest utility for the sender. And because the receiver will choose a_2 only when the conditional probability of θ_2

is at least as large as that of θ_3, action a_2 is induced with the maximum possible probability of $\frac{1}{3}$. Inducing any other actions for any of the states would yield the sender utility 0. Hence, the given signaling scheme is optimal for the sender.

Under this signaling scheme, the receiver's expected utility is $\frac{1}{2} \cdot (\frac{2}{3} \cdot 4 + \frac{1}{3} \cdot 2) + \frac{1}{3}(\frac{1}{2} \cdot 3 + \frac{1}{2} \cdot 1) + \frac{1}{6} \cdot 3 = \frac{17}{6}$.

Thus, the constrained receiver's utility of $\frac{17}{6}$ is lower than the unconstrained receiver's of 3.

5.6 Discussion

We showed that a state-matching receiver, facing an action-matching sender under a binary state space, obtains weakly higher utility when more constrained. We believe that such behavior is in fact observed in the real world: for example, recommenders tend to be more careful in whom they nominate for particularly selective awards or positions.

5.6.1 Larger state/action spaces

As we discussed in Section 5.5, our results do not carry over to larger state spaces. Indeed, even for state spaces with three states, in which the receiver tries to minimize the distance between the action and the state of the world, there are counter-examples under which a constrained receiver is worse off.

While the result does not hold in full generality with three (or more) states, by imposing additional conditions, a positive result can be recovered:

Proposition 6. *Assume that the state space has size $|\Theta| = 3$, and that the receiver is state-matching and the sender is action-matching. In addition, assume that the following two conditions are satisfied.*

1. *The sender has a monotone[2] preference over actions across all states, i.e.,* $U_S(\theta_i, a_1) \geq U_S(\theta_i, a_2) \geq U_S(\theta_i, a_3)$ *for all* i.

2. *For every state* i, *the receiver is worse off choosing an action* $j < i$ *that is too low compared to choosing an action* $k > i$ *that is too high[3]: that is,* $U_R(\theta_i, a_j) \leq U_R(\theta_i, a_k)$ *for all* $j < i < k$.

Then, a more constrained receiver is never worse off than a less constrained one.

The additional assumptions on the sender side capture a stronger version of the utility relationship of the interesting cases in the proof of Theorem 5. They are motivated in many of our cases: for instance, a letter writer may want to obtain the highest possible honor (or salary) for a student, or a prosecutor may want to maximize the sentence of a defendant.

The additional assumption on the receiver side would capture a cautious department or judge, who would prefer to err on the side of not inviting weak candidates (or giving awards to undeserving candidates), or giving the defendant a sentence that is too low rather than ever giving too high of a sentence.

While Proposition 6 shows that with enough assumption, a positive result can be recovered, we believe that the assumptions are still rather restrictive, meaning that the proposition is likely of limited interest. The proof involves a long and tedious case distinction, and we provided it in Appendix E.1.2.

For fully general state spaces (i.e., $n = |\Theta| \geq 3$), we can currently obtain a positive result only by imposing even more assumptions on the utility functions. In addition to the (generalization of) the assumptions from Proposition 6, we can make the following assumptions: (1) Whenever $j < i$, the sender's utility difference between actions $j < j'$ is larger under state θ_i than under state $\theta_{i'}$ for $i' > i$. In other words,

[2]The result holds symmetrically if the order is reversed.

[3]Notice that in the case $|\Theta| = 3$, this constraint only applies to $i = 2, j = 1, k = 3$. We phrase it more generally to set the stage for a further generalization below.

when the state of the world is smaller, the sender is more sensitive to changes in the receiver's action. (2) For any fixed state θ_i, the receiver's utility as a function of j (the action) is increasing and *convex* for $j \leq i$, and decreasing and convex for $j \geq i$. By adding these two assumptions, we can again obtain a result that a constrained receiver is always weakly better off than an unconstrained one. While it is possible to construct reasonably natural applications which satisfy these conditions, the conditions are far from covering a broad class of Bayesian persuasion settings. For this reason, we are not including a proof of this result, instead considering the discussion as a point of departure towards identifying less stringent assumptions that may enable positive results.

Whether there is a broad and natural class of Bayesian persuasion instances with more than 2 states of the world in which the insight "A more constrained receiver is better off" from Theorem 5 carries over is an interesting direction for future research.

5.6.2 Finding optimal signaling schemes

While the main focus of our work is on the receiver's utility when more constrained, our model also raises an interesting computational question, as briefly discussed in Section 5.3.3. In particular, we do not know whether there is a polynomial-time algorithm which — given the sender's and receiver's utility functions as well as the constraints on the receiver — finds a sender-optimal signaling scheme. Since probability constraints on receivers (quotas) are quite natural in many signaling settings, this constitutes an interesting direction for future work.

The main difficulty in applying standard techniques is that the constraints may force the receiver to play an ex post suboptimal action. The standard linear program for the sender's optimization problem *Dughmi and Xu* (2019) maximizes the sender's expected utility subject to the constraint that the receiver is incentivized to play the sender's recommended action. To appreciate the difference, consider a setting in which the state of the world is uniform over $\{\theta_1, \theta_2\}$, and the sender and receiver both obtain utility 1 if the receiver plays action a_1, and 0 otherwise. Without any constraints, the sender need not send any signal, and the receiver would simply play action a_1. But if the receiver is constrained to playing action a_1 with probability exactly $\frac{1}{2}$, then she must randomize, including the (always suboptimal) action a_2 with probability $\frac{1}{2}$. By Proposition 5, the randomization can be pushed to the sender instead, but when the sender recommend action a_2, it will be ex post suboptimal for the receiver to follow the recommendation. Indeed, an LP requiring deterministic ex post obedience from the sender would become infeasible for this setting. Whether the sender's optimization problem can still be cast as a different LP, or solved using other techniques, is an interesting direction.

We remark here that the preceding example does not have a state-matching receiver. If the receiver is state-matching and the constraints are feasible, then full revelation of the state is ex post incentive compatible for the receiver. This implies that the linear program for optimizing the sender's utility over ex post incentive compatible signaling schemes has a feasible solution. However, since the LP is more restricted, it is not at all clear that its optimum solution maximizes the sender's utility when the recommendation does not have to be ex post incentive compatible.

5.6.3 Alignment of the action-matching sender

Results in this Chapter stand on the assumption that the sender has a partial-aligned utility with the receiver. To emphasize the importance of the assumption of sender's partial alignment, we present an example where the receiver can be worse off under constraints without an action-matching sender.

Consider the following 2×2 example with prior probability on θ_1 is $p = 0.25$, assuming $\epsilon \ll 1$.

Given the sender's and receiver's utilities, we know the sender prefers action a_1

	θ_1	θ_2
a_1	2	3
a_2	1	0

(a) Sender's Utility

	θ_1	θ_2
a_1	1	ϵ
a_2	0	1

(b) Receiver's Utility

Table 5.3: Sender's and Receiver's Utility in the example without partial alignment

in both states, and the receiver prefers action a_i when the true state is θ_i, $i \in \{1, 2\}$. For simplicity of representation, ω_i denotes the sender's signal suggesting action a_i, $i \in \{1, 2\}$.

We will compare receiver's expected utilities in the following two settings:

1. There are (effectively) no constraints, i.e., $\left[\bar{b}_1 = 1, \underline{b}_1 = 0, \bar{b}_2 = 1, \underline{b}_2 = 0\right]$

2. A constraint profile binds the sender-preferred action to its prior probability, i.e., $\left[\bar{b}_1 = 0.25, \underline{b}_1 = 0, \bar{b}_2 = 1, \underline{b}_2 = 0\right]$

Setting 1 is a classical 1-sender Bayesian persuasion problem, where the sender's optimal signaling strategies can be solved by the concavification approach presented in *Kamenica and Gentzkow* (2011). The sender's optimal signaling strategy is the following:

- Send ω_1 with probability 1 when the true state is θ_1, with probability $\frac{1}{3}$ when the true state is θ_2.

- Send ω_2 with probability $\frac{2}{3}$ when the true state is θ_2

Given this commitment, the receiver's expected utility is $\frac{1+\epsilon}{2}$ while receiving ω_1 (because θ_1 and θ_2 are equally likely to occur), and her expected utility is 1 while receiving ω_2. Given the current prior, the receiver's expected utility is $\frac{3+\epsilon}{4}$ via straightforward computations.

In setting 2, the sender cannot send signal ω_1 as often as in the unconstrained case. When she is forced to reduce ω_1, she prefers to reduce θ_2 from ω_1 over θ_1.

This is because $U_S(\theta_2, a_1) - U_S(\theta_2, a_2) > U_S(\theta_1, a_1) - U_S(\theta_1, a_2)$. However, reducing θ_2 from ω_1 may cause the failure of persuasion when the posterior belief violates the inequality $\mathbb{P}(\theta_2|\omega_1) \leq \mathbb{P}(\theta_1|\omega_1)$. Hence, the sender's optimal signaling strategy is to maximize the total probability of ω_1 under constraints as the posterior belief satisfies $\mathbb{P}(\theta_2|\omega_1) = \mathbb{P}(\theta_1|\omega_1)$. Thus, the sender's optimal commitment is the following:

- Send ω_1 with probability $\frac{1}{2}$ when the true state is θ_1, with probability $\frac{1}{6}$ when the true state is θ_2.
- Send ω_2 with probability $\frac{1}{2}$ when the true state is θ_1, with probability $\frac{5}{6}$ when the true state is θ_2

Given the above signaling strategy, the receiver's best response under ω_1 is taking action a_1, with expected utility $\frac{1+\epsilon}{2}$. And her best response under ω_2 is taking action a_2, with expected utility $\frac{5}{6}$. Hence, the receiver's expected utility under the current prior is $\frac{6+\epsilon}{8}$ under the current constraint profile.

To sum up, the receiver's expected utility in setting 1, $\frac{3+\epsilon}{4}$, is higher than the receiver's utility in setting 2, $\frac{6+\epsilon}{8}$. This shows an example where the receiver can be worse off under constraints when the sender's utility is not partially aligned with the receiver.

5.6.4 A potential issue: sender's tie-breaking rule on signaling strategies

This chapter presents that a state-matching receiver is always better off under constraints in a binary state-space model with an action-matching sender. However, when the sender is not action-matching or the state space is non-binary, the sender's optimal signaling strategy may not be unique. Hence, the sender's choice of signaling scheme could give the receiver different payoffs. For example, consider a sender (Bob) who has a constant utility regardless of actions and states, i.e., he does not care what the receiver does. One optimal strategy is to reveal complete information, and another is to flip a fair coin. Suppose Bob decides to use different signaling strategies under different constraints, e.g., flip a coin to a constrained receiver but reveal complete information to an unconstrained one. In this case, Bob's choice of signaling strategy (tie-breaking rules) becomes the sole factor that impacts the receiver's utility. This example does not fit our model and goal not only because the sender is not action-matching but because the constraints never bind the sender's signaling strategies. However, we are aware that the sender's choice of signaling strategies may be a potential issue for generalizing non-binary state models. To be more precise, we are interested in whether the following scenario exists:

1. The sender has a set of optimal signaling schemes in the unconstrained case, call it $\mathbf{\Gamma}_o$.

2. (Binding Requirement) There is a set of feasible constraints that bind the sender's signaling scheme under the current prior, i.e., every $\Gamma \in \mathbf{\Gamma}_o$ is not an obedient signaling scheme. In other words, the sender's expected utility decreases under this binding constraint.

3. The sender has a new set of signaling schemes under this set of constraints, call it $\mathbf{\Gamma}_c$.

4. There exists a $\Gamma \in \mathbf{\Gamma}_o$ and $\hat{\Gamma} \in \mathbf{\Gamma}_c$ such that the receiver's expected utility under $\hat{\Gamma}$ is lower than her utility under Γ.

At this point, we have not found any instance under the state-matching receiver and action-matching sender satisfying the above scenario. However, proving the non-existence of the above scenario is a prerequisite for establishing results for general state-space models. Currently, this issue is still unresolved, and the exploration of it is left for future work.

CHAPTER VI

Conclusion

This thesis studied problems where higher-order inference plays a critical role in games. We first highlighted the importance of systematical approaches/algorithms to derive higher-order inference, which facilitated the belief updating procedure and then pointed out the subtlety and the difficulty of deriving optimal signaling schemes when strategic behaviors come to play. With results and contributions presented in each chapter, we first conclude and describe three unifying insights presented in this thesis and discuss future directions.

6.1 Unified Insights

6.1.1 Belief updating procedure and the partition of information spaces

In Chapter II, we presented a pivotal perspective to view the design of queried information (called questions in Chapter II) as the design of partitions on information spaces. By designing the partition of information spaces, we can systematically analyze the information accumulation among agents in social learning models and determine whether (asymptotic) learning occurs. In Chapter III, we presented a dynamic programming approach with backward iteration to help the sender make partitions on the information space on top of current partitions (determined by experiments executed in later phases). The algorithm helps the sender to iteratively design the signaling strategies based on the realized sequence of signals (generated from the experiment conducted in each phase) to persuade the receiver with a tailored partition of the information space. These results in Chapter II and Chapter III highlight the importance of systematical (and algorithmic) partitioning approaches on sequential belief updating. With a systematical approach to partition the information space, we can then analyze the higher-order inferences systematically.

6.1.2 Subtlety of strategic senders' optimal commitments

In Chapter III, the tailored optimal commitments under pre-determined experiments (presented in Section 3.4) and the fragility of optimal commitments under small perturbations (presented in Section 3.4.4.1) both indicated that the sender's optimal signaling strategy demands a tailored design on the signal sequence's structure (i.e., the distribution of posterior distributions). In Chapter IV, all the examples highlighted the same fact: the interaction between two strategic senders and their inferences on the other's signaling strategies make the optimal commitments heavily dependent on their (private) information structures. In Chapter V, exogenously constraints forced the sender to redesign her optimal signaling strategy. Our results stated that the receiver is better off in binary state-space models, and we discussed the difficulty of developing approaches to solve optimal commitments under exogenously constraints in general state-space models in Section 5.6.2. To sum up, results in Chapter III-V indicate the subtlety of optimal commitment when strategic behaviors (on signaling) come to play, even though higher-order inference can be systematically analyzed like in the models presented in Chapter III.

6.1.3 Commitment order and level of inference

In Chapter IV, we presented conditions on when the commitment order matters. From the perspective of the level of inference, the sender who commits first is the one who is forced to stop her inference (and commit her signaling strategies) before the other sender. Therefore, examples in Section 4.4 where both senders prefer committing first and the statement in Theorem 3 where senders prefer committing second showed that there is no rule of thumb to guide strategic agents on the level of inference that is optimal (for them to design their signaling strategies), especially in the setting where senders have incomplete information about the state of the world. Besides, the demand of recursive reasoning in objective functions (stated in Section 4.2.4) and the challenge of developing algorithms (discussed in Section 4.4.4) for the optimal commitments highlight the difficulties of systematical approaches to analyze the interaction of higher level inferences among agents under asymmetric information.

6.2 Future Directions

There are a variety of future directions for Chapter III, IV, and V detailed below:

6.2.1 Future directions for Chapter III

In Chapter III, we analyze the optimal signaling strategy for the sender when some of her experiments are pre-determined. A natural follow-up extension is to study the receiver's optimal experiments if she can design (some of) pre-determined experiments beforehand. The main difficulty of this direction is to make proper assumptions on the receiver's room to design experiments that fit the scenarios that occur in real-world problems. Another direction is to generalize the algorithms (as well as the structural results) beyond the binary state-space model. In binary state-space models, the structural results in Lemma 15, 14, 16, and Corollary 4 allow us to avoid solving complex LP problem via solving some linear equations. However, these structural results cannot be directly generalized to non-binary state-space models. Thus, deriving structural results and algorithms to solve the optimal signaling strategy in sequential BP problems under non-binary state-space models is a future direction to explore.

6.2.2 Future directions for Chapter IV

Given the results in Theorem 2 and Theorem 3, we have a set of necessary conditions and a set of sufficient conditions. A future direction is to derive necessary and sufficient conditions on when the commitment order matters. In this direction, we may need additional assumptions on both senders' tie-breaking rules. Besides, this direction requires a deeper understanding of the significance of the prior's influence on the commitment order.

Besides, several future directions can be explored to generalize the result by lifting or weakening the assumptions made in Chapter IV. For example, one direction is to consider problems where one sender is more (Blackwell) informative than the other. This direction can advance the understanding of the generality of the less-informative-commitment results presented in *Gentzkow and Kamenica* (2016) and *Li and Norman* (2018b). Another direction generalizes the receiver's utility function under some structures, e.g., unimodal and distance-based utility, state-matching with a preference order on states.

Another future direction is to generalize our results to more than two senders, even under a restrictive assumption where each sender's information space is binary. We suspect the analysis will be really cumbersome in this extension if the issue presented in Section 4.4.4 has not yet been resolved.

6.2.3 Future directions for Chapter V

As mentioned in Section 5.5, a future direction is to generalize the results to non-binary state-space models with mild additional assumptions on sender/receiver's utility functions. Based on the example presented in 5.5, we expect some monotonicity assumptions on sender's and receiver's utility difference aligning with the preference order might be mandatory to restrict sender's room of design. Suppose we want to generalize the results to non-binary state-space models but keep assumptions about the sender/receiver's utility function. In that case, a direction can be explored by enforcing additional assumptions on the constraint profiles. By shrinking the space of constraints to some particular constraint types, we may generalize the results to non-binary state-space models.

APPENDICES

APPENDIX A

Motivational Example in Introduction

A.1 Motivational Example in Detail

Alice, Bob, and Carol play a game where Alice and Carol have private information about past participants' actions that Bob cannot access. In this game, Bob takes action, either X or Y, trying to guess the state of the world $\Theta = \{X, Y\}$. If Bob makes a correct guess, his utility is 1, and 0 otherwise. Alice and Carol can present their private information truthfully to Bob and suggest an action before Bob's decision. I_A denotes the information sent by Alice, I_C denotes the information sent by Carol. R_A denotes the action suggested by Alice, and R_C is the action suggested by Carol. Alice and Carol's utilities U_A, U_C are 1 when Bob takes the action they suggest, and 0 otherwise.

In this example three groups of past agents A_i, B_j, C where $i \in \{1, ..., m\}$, $j \in \{1, ..., n\}$ have played the same game before and their utility functions are the same as Bob. Before her decision, each agent can observe a subset of other agents' decisions and an i.i.d private signal with an accuracy q. The network topology is shown in 1.1, where each agent's observation is drawn with blue curves.

Given the above topology, some of past agent's decisions may be made because of herding behavior. Herding behavior occurs when agents' observation overwhelmingly suggests to guess on a particular state and ignore their private signal. For example, if agent B_3 observes the same action X by B_1 and B_2, her best strategy is to take action X and ignore her private signal (since she is Bayes rational).

Now, Alice and Carol can observe a set of past agents' actions. Alice's observation is drawn with red curves in Figure 1.1, and Carol's observation is drawn with green curves in Figure 1.1.

In our example, Alice observes dramatically different actions taken in group A and group B. Every agents A_i she observes taken action X, where $i \neq 2$. However, every agents B_j she observes has taken action Y, where $j \neq 1, 2$. Based on her observation, Alice cannot know whether agents A_i's or B_j's action has used her private signal or it is the result of herding behavior when $i, j \geq 3$. Thus, Alice has to consider both possibilities to calculate the (Bayesian) interim belief. With each past agent's private signal having the same accuracy q, Alice's interim belief will be in favor of action X. Hence, Alice will tell Bob what she observes and suggest action X.

At the same time, Carol observes the action X taken by agent B_1 and C. Thus, Carol will tell Bob what he observes and suggest action X.

If Bob only gets information and suggestions from Alice or Carol, his best decision is to take action X (since Bob has no other private information to help his decision). However, suppose both Alice and Carol's observation come to Bob, Bob learns that every agent B_j's action Y with $i \geq 3$, observed by Alice, is based on their private information because Carol tells him that B_1 takes action X instead of Y. This inference overturns the interim belief calculated by Alice (because in Alice's analysis, a sequence of actions Y is based on each past agent's private information has a much smaller probability than it is based on herding behaviors. According to Bob's inference, he should take action Y even though both Alice and Carol suggest action

X, and an information Braess's paradox occurs in this motivational example.

To avoid misunderstanding, we carefully design the information sources of Alice and Carol to make their information sources has no intersections. Therefore, we can claim that Bob's overturn decision is not based on removing duplicate information sources from Alice and Carol but based on the inference on the information structures of agents observed by Alice or Carol.

APPENDIX B

Appendices of Chapter II

We put the proofs of technical claims and calculations that support the arguments mentioned in the main chapter in Appendix B.2 to avoid the troublesome of searching the primary proofs of lemmas/propositions/theorems with the mixture of those technical parts.

B.1 Major proofs of Chapter II

B.1.1 Proof of Lemma 1

Proof. First, if the cardinality of the private signal space is greater than the cardinality of the state space[1], $|\mathcal{S}| \geq |\Theta|$, then the existence of non-distinguishable private signal sequences is guaranteed because there exists a finite n such that $|\mathcal{S}|^n > |\Theta|^n \times 2^{|K_n|}$, where K_n is the capacity constraint for agent n satisfying $\lim_{n\to\infty} K_n < \infty$.

Then, while $|\mathcal{S}| \leq |\Theta|$, we will show that a non-distinguishable private signal sequence will occur in cascade continuing for a long enough period. When a cascade starts, the realized history stops providing the information of an agent's private signal to her successors. Hence, utilizing the communication capability is the only way to

[1] This holds because the action space is the same as the state space.

distinguish the private signal sequence of agents in a cascade. When an agent n is in a cascade, her observation and her signal can help her distinguish at most $|\Theta| \times 2^{|K_n|}$ different private signal sequences based on the response of her questions and the observed history. Hence, when this agent is in a cascade with length $m > \frac{\ln|\Theta|+|K_n|\ln 2}{\ln|\mathcal{S}|}$, there must exists some some private signal sequence being not distinguishable to this agent because $|\mathcal{S}|^m > |\Theta| \times 2^{|K_n|}$.

Given the assumption of bounded signal strength, there must exist a set of private signal subsequences with a finite (and fixed) length $l > m$ such that each subsequence will start a cascade and continue at least m periods. When a subsequence in this set occurs, the private information sequence is not distinguishable by every agent who arrives later than the m^{th} one in the cascade owing to the capacity constraint. Given that all private signals are i.i.d. drawn, by the infinite monkey theorem, subsequence in the above set appears almost surely as the length of private signals n goes to infinity. □

B.1.2 Proof of Lemma 2

Proof. According to Definition 3, a feasible question allows the corresponding agent to update her belief using Bayes rule. To deal with non-probabilistic but feasible questions such as questions in Example 5, we must prove that every non-probabilistic but feasible question can be mapped to a probabilistic question under the same capacity constraints.

To make non-probabilistic questions feasible, the randomness (the distribution of responses) that is not well-defined in these questions should not change the posterior belief no matter what the realizations are; otherwise, Bayes rule cannot be applied. With this feature of non-probabilistic but feasible questions, we then prove this lemma by constructing an injective function that maps from a non-probabilistic but feasible question to a probabilistic question under the same capacity constraints.

Suppose we are given a K-bit non-probabilistic question Z with posterior beliefs of responses $q_1^Z, q_2^Z, ..., q_{2^K}^Z$ updated using Bayes rule, we there exists a k-bit probabilistic question Z' with posterior beliefs of responses $q_1^{Z'}, q_2^{Z'}, ..., q_{2^K}^{Z'}$ such that there is an injective function maps from $q_1^Z, q_2^Z, ..., q_{2^K}^Z$ to $q_1^{Z'}, q_2^{Z'}, ..., q_{2^K}^{Z'}$, where each possible response in Z' represents a unique posterior belief, i.e. $q_m^{Z'} \neq q_n^{Z'}$ $\forall m \neq n$. Such mapping exists because each possible response in Z has a posterior belief derived using the Bayes rule. The number of possible posterior beliefs in Z must be less than or equal to the number of possible responses (2^K). Hence, the probabilistic question Z' can be formulated under the same capacity constraint. □

B.1.3 Proof of Lemma 3

Proof. When there is no information cascade occurring, the strongest message an agent can get from a question is still bounded owing to the bounded signal strength and the finite channel capacity. However, when there is no cascade, each agent's action reveals her preferred state based on her private signal. In other words, an agent's action can be interpreted as a signal telling her successor the set of private signals she receives. Since each agent's private signals are i.i.d. generated, agent t's interim belief (after observing the history but before receiving the response of her question) of the true state goes to 1 and the interim belief of any other state goes to 0 as t goes to infinity almost surely by the strong law of large number (sLLN). This result by sLLN indicates that the likelihood ratio of the true state over any other state is unbounded almost surely. Therefore, there must exist an agent with a finite index such that none of the responses to any feasible questions can overturn her best-response action suggested from her observed history. When this phenomenon occurs, an information cascade starts. □

B.1.4 Proof of Lemma 4

Proof. We remind the reader that $e^j_{Q,t} : \mathcal{H} \times \Theta \to [0, 1]$ is the (probability) function (of the observed history up to agent k, $H_k \in \mathcal{H}$, and the true state $\theta_j \in \Theta$) that represents the probability of the cascade of θ continued in agent k being stopped at agent $t \geq k + 1$ conditional on the cascade continues till agent $t - 1$, assuming that the question guidebook Q executes in cascades. We want to verify the following two statements are equivalent:

- $\lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_i|\theta_i) = 1$ and $\forall j \neq i, \lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_i|\theta_j) = 0$;
- $\sum_{t=k+1}^{\infty} - \ln\left(1 - e^j_{Q,t}(H_k, \theta_i)\right) = \infty$, $\sum_{t=k+1}^{\infty} - \ln\left(1 - e^j_{Q,t}(H_k, \theta_j)\right) < \infty$ when $j \neq i$, and every agent $t \geq k$ is in a cascade.

The first equation in the second statement, $\sum_{t=k+1}^{\infty} - \ln\left(1 - e^j_{Q,t}(H_k, \theta_i)\right) = \infty$, implies all the wrong cascade will be stopped in finite time. This is because the equation guarantees the product $\Pi_{t=k+1}^{\infty}(1 - e^j_{Q,t}(H_k, \theta_i)) = 0$. Besides, if $\sum_{t=k+1}^{\infty} - \ln\Big(1 - e^j_{Q,t}(H_k, \theta_j)\Big) < \infty$ when $j \neq i$ and every agent $t \geq k$ is in a cascade, then the right cascade has a positive probability to last forever because $\Pi_{t=k+1}^{\infty}(1 - e^j_{Q,t}(H_k, \theta_j)) > 0$. Therefore, $\lim_{n\to\infty} \mathbb{P}(\{a_t = \theta_i|\theta_i) = 1$ and $\lim_{n\to\infty} \mathbb{P}(\{a_t = \theta_i|\theta_j) = 0$ are satisfied because every wrong cascade will be stopped almost surely; and once an agent is in the right cascade, the current cascade has a positive probability to last forever. Hence, the second statement, combining both parts, states that as the index goes to infinity, agents will be in the right cascade almost surely, equivalent to the standard definition of asymptotic learning. □

B.1.5 Proof of Lemma 5

Proof. Before showing Lemma 6, we first present a useful property in threshold-based QGBs in Claim 9, with detailed proof provided in Appendix B.2.6.

Claim 9. *Given a threshold-based QGB which is feasible and incentive-compatible, active agents cannot arrive consecutively.*

Now, since a threshold-based QGB satisfies the definition of regular QGBs, we only need to prove that a threshold-based QGB is aligned. Given that every active agent goes to G_1 once knowing she cannot stop the cascade and Claim 9 asserts that the next agent must be a silent agent, the first silent agent arriving after an active agent must start from G_1.

Given the first silent agent starts from G_1 and the definition of threshold-based QGBs (Definition 17), an agent in state G_u has $v-u$ more signal s_j than an agent in state G_v for every $1 \leq u < v$ when we are in an $\bar{a}_i$ cascade. Hence, the likelihood ratios of θ_j over θ_i in these two state must satisfy the inequality $l^n_{G_u} \geq l^n_{G_v}$. Therefore, the threshold-based QGB is aligned based on the transitive property of inequality. □

B.1.6 Proof of Lemma 6

Proof. Given that every active agent goes to G_1 once she knows that she cannot stop the cascade, Claim 9 stated in Appendix B.1.5 asserts that the next agent must be a silent agent. With the claim, we can prove this lemma by contradiction. Starting from G_1, we assume the next active agent is a^{k+1} who can stop the cascade at G_i for some $i > 1$. Now, consider that likelihood ratio of state G_1 at agent $a^{k+1}-i+1$, $l^{a^{k+1}-i+1}_{G_1}$, $l^{a^{k+1}-i+1}_{G_1} \geq l^{a^{k+1}}_{G_i}$; this is because of the allowed transitions in the Markov chain of silent agents. Agent a^{k+1} can stop the cascade at G_i and this requires $l^{a^{k+1}}_{G_i} \geq \frac{1-p}{p}$. It implies that $l^{a^{k+1}-i+1}_{G_1} \geq \frac{1-p}{p}$. Now, agent $a^{k+1}-i+1$ must be an active agent, which contradicts that next active agent is a^{k+1}. □

B.1.7 Proof of Corollary 1

First, we show that asymptotic learning is not achievable if one of these two conditions fails.

The check for the first condition is straightforward. If it fails, then there exists an N such that $m^k < N$ for all k. Now, $h^-(m^k)$ is lower-bounded by $h^-(N) > 0$. As we all know, $\lim_{n\to\infty}(1 - h^-(N))^n = 0$, a correct cascade will always be stopped in a finite time horizon, and asymptotic learning is not possible.

For the second condition, if $\Pi_{k=1}^{\infty}(1 - h^+(m^k)) \neq 0$, then there is a positive probability that a wrong cascade lasts forever. Obviously, asymptotic learning cannot be achieved. If $\Pi_{k=1}^{\infty}(1 - h^-(m^k)) = \Pi_{k=1}^{\infty}(1 - h^+(m^k)) = 0$, then we will keep stopping every cascade whether it is a right cascade or not. Therefore, $\lim_{t\to\infty} \mathbb{P}(a_t = \bar{a}_\theta) < 1$ because we will have a positive probability of either being in a wrong cascade or not being in a cascade.

In the other direction, the first condition guarantees that the frequency of active agent goes to zero when we continue operating a threshold-based QGB. However, this does not guarantee that we will have infinite number of active agents. Combining the first condition and the right half of the second condition, i.e., $\Pi_{k=1}^{\infty}(1 - h^+(m^k)) = 0$, we can ensure that a threshold-based QGB will generates infinite number of active agent. This is because $h^+(m^k) > 0$ for every finite m^k according to the equation (2.3), but $\Pi_{k=1}^{\infty}(1-h^+(m^k)) = 0$. With infinite number of active agents, $\Pi_{k=1}^{\infty}(1-h^+(m^k)) = 0$ guarantees that all wrong cascades will be stopped, and $\Pi_{k=1}^{\infty}(1 - h^-(m^k)) > 0$ ensures that some right cascades will continue forever under the existence of infinite number of active agents. Hence, then asymptotic learning is achieved.

B.1.8 Proof of Lemma 7

Recall that a silent agent appears when an agent will not take action different from her predecessor, whatever private signal and the response of questions she gets. On the other hand, if agent t is active and gets the strongest signal against her predecessor's action, she should be able to take the opposite action in at least one of the following two circumstances.

First, agent t and her immediate predecessor $(t-1)$ both get the strongest signal against the $(t-1)^{\text{th}}$ agent's action a_{t-1}, but agent $t-1$ got a negative response from agent $t-2$.

Second, agent t gets the strongest signal against her immediate predecessor's action a_{t-1}, and agent $t-1$ gets the strongest signal belong to the set E_{t-1}, where E_{t-1} is the set of signals that agent $t-1$ will not ask questions.

Now, suppose every agent up to $t-1$ is active. Define q_t to be the probability that agent t will take action $\bar{a}_j$, and let the strongest signal in favor of state θ_j be s^{j*} and its probability is p_j^*, then the public belief (i.e., the likelihood ratio of θ_j versus θ_i) when $a_n = \bar{a}_i$ for $n \in [t-1]$ is upper-bounded by $\Pi_{n=1}^{t-1} \frac{1-\mathbb{P}(\theta_j|s^{j*})q_n}{\mathbb{P}(\theta_i|s^{j*})q_n}$. Note that this upper-bound is strictly less than 1 since we assume that both states are equally likely to happen.

In order to bound q_t, we can take the minimum probability of the two circumstances stated above to get the following inequality:

$$q_t \geq \min\left\{q_{t-1}p_j^*, \mathbb{P}\big(u|u,v \in E_{n-1}, \mathbb{P}(\theta_j|u) > \mathbb{P}(\theta_j|v)\big)\right\} \geq \min\left\{\min_{s\in S}\mathbb{P}(s), \frac{p_j^*}{1+p_j^*}\right\} > c.$$

Thus, we know that q_t is bounded above c for all i so that the public belief conditional on $a_t = \bar{a}_i$ for $t \in [n]$ will converge 0 as $n \to \infty$. Recall the two circumstance stated above, the minimum public belief making at least one of the two circumstance to happen (in other words, agent t is an active agent) requires the public belief up to $t-2$ be $\geq \left(\frac{\mathbb{P}(\theta_i|s^{j*})}{\mathbb{P}(\theta_j|s^{j*})}\right)^2$. Therefore, there must exists a finite number T such that agent T an silent agent when $a_t = \bar{a}_i$ for $t \in [T-1]$. Now, as $n \to \infty$, $\mathbb{P}(\{\exists m + T < n, a_{m+t} = \bar{a}_i \ \forall t \in [T]\}) \to 1$. Thus, the lemma is proved.

B.1.9 Proof of Lemma 8

First, we can bound the factor that the growth rate m^k may be faster than $\bar{m}^k$, which we call the early jump. Recall the ratio of $h^+(\bar{m})$ over $h^-(\bar{m})$ is the likelihood ratio of θ_j versus θ_i conditioned on the event that the cascade continues after $\bar{m}$ silent agent. Then calculation in Appendix B.2.14 bounds the lower bound of $\frac{h^+(\bar{m+1})}{h^-(\bar{m+1})} / \frac{h^+(\bar{m})}{h^-(\bar{m})}$, which we denote by γ. The early jump given the signal strengths in our model, will be bounded by at most $\frac{\mathbb{P}(\theta_j|s^{j*})\mathbb{P}(\theta_i|\hat{C}_{\theta_j})}{\mathbb{P}(\theta_i|s^{i*})\mathbb{P}(\theta_i|\hat{C}_{\theta_j})}\frac{1}{\gamma}$.

Now, we want to bound the factor which may lower the growth rate of m^k. The scenario that leads to the slowest growth rate of m^k is when an active agent only stops the cascade when she gets s^{j*}, the strongest signal which favors $\bar{a}_j$. Reviewing the analysis in Appendix B.2.11, we can get a sequence $\bar{w}_s$ such that $|\{m^k = s|l_{G_1}^{H_{a^{k-1}+1}}\frac{h^+(s)}{h^-(s)} \geq 1\}| \geq \bar{w}_s$ for all $s \in \mathbb{N}$, which will only vary by at most a constant ratio in the logarithm term. Recalling the original $\bar{w}_s = \frac{\ln(c_2(p))}{\ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right)}$, we note that the new growth rate is lowered because active agents only stop cascade in a subset of $\hat{C}_{\theta_j}$, but it can be upper-bounded by $n_s \leq \frac{\ln(c_2(p))}{\ln\left(\frac{1-h^-(s)p_B^*}{1-h^+(s)p_j^*}\right)} \leq (1+\frac{1}{p_j^*})\frac{\ln(c_2(p))}{\ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right)}$ when s is large. Hence, we can set $t_2 = 1 + \frac{1}{p_j^*}$ and the proof is complete here.

B.1.10 Argument of Proposition 1

Instead of rewriting all the algebra in Appendix B.2.9, B.2.10, B.2.11, B.2.12, B.2.14, and B.2.15 by fixing notation p, $1-p$ to the corresponding variables here, we use Lemma 7 and the result in Appendix B.2.10, B.2.12 to capture the core of this proposition.

When the growth rate of m^k goes faster than $\bar{m}^k$, all we need to check is whether the wrong cascade will still be stopped with probability 1. Recall the equation in

Appendix B.2.10,

$$\begin{aligned}
& \mathbb{P}(\text{A wrong cascade will stop}) \\
\geq\ & 1 - A(p)\Pi_{n=J}^{\infty}(1-p^n)^{\frac{\ln(c_1(p))}{p^{n-1}(1+p^{n-1})}} \\
\geq\ & 1 - A(p)\exp\Big(p\sum_{n=0}^{\infty}\frac{\ln(c_1(p))}{p^n(1+p^{n-1})}(-p^n)\Big) \\
=\ & 1 - A(p)\exp\Big(-p\ln(c_1(p))\sum_{n=0}^{\infty}\frac{1}{1+p^n}\Big) = 1,
\end{aligned}$$

where $A(p) = \Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1-p^{m^k})$. We know that lowering m^k to $m^k - t_1$ will introduce a constant parameter p^{-t_1} in the denominator, which will not change the convergence properties of the series, and the wrong cascade will still be stopped w.p. 1.

When the growth rate of m^k goes slower than $\bar{m}^k$, we need to check whether the right cascade will last forever with a positive probability. First, let us recall the equation in Appendix B.2.12.

$$\mathbb{P}(\text{A right cascade will be stopped}) \leq 1 - B(p)\exp\Big(\sum_{n=J}^{\infty}\frac{1}{c_4(p)c_1(p)^n - 1}\Big),$$

where $B(p) = c_2(p)e^{c_3(p)}\Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1-h^-(k))$.

Now, the change of rate with ratio bounded by t_2 is equivalent to multiplying $\frac{1}{t_2}$ to each term in the summation within the exponential function. This only increases the probability that a right cascade will be stopped but not the property that $\mathbb{P}(\text{A right cascade will be stopped}) < 1$, which will remain unchanged.

B.1.11 Proof of Lemma 9

If the signal space not distinguishable, then there exists a state θ_i such that the following inequality holds:

$$\mathbb{P}(\theta_i|s) \leq \mathbb{P}(\theta_j|s)\ \forall j \neq i \text{ for every finite-length sequence of signal } s. \tag{B.1}$$

When strictly inequality holds for every $j \neq i$, an agent will never take action to guess state θ_i, and the asymptotic learning can not be achieved. When equality holds for some $j \neq i$, an agent can not distinguish θ_i and θ_j in any signal sequence. In this case, asymptotic learning can not be achieved either.

B.1.12 Proof of Lemma 10

Since we only need to care about the feasibility, we can pick every other possible state different from the current cascade suggests operating a 1-bit QGB proposed in Section 2.6 respectively. Since we have a constant-finite number of states and the true state is either θ_i or θ_j, the process of checking other states θ_k will not alter whether asymptotic learning is achievable in the 1-bit QGB checking whether θ_i or θ_j is the true state. Hence, operating a 1-bit QGB for every state different from the current suggested state in the cascade, in turn, can achieve asymptotic learning.

B.1.13 Proof of Lemma 11

The uniform prior guarantees that the prior is in the center of the $|\Theta|$ simplex where every points in the simplex represents a distribution of the state space. A distinguishable signal space guarantees that there must be a finite-length sequence of signals such that the convex hull of the signal sequences has a positive measure (in the simplex). Besides, the prior must be an inner point in the convex hull of the signal sequences. Therefore, a construction of μ-chunks must exist such that the

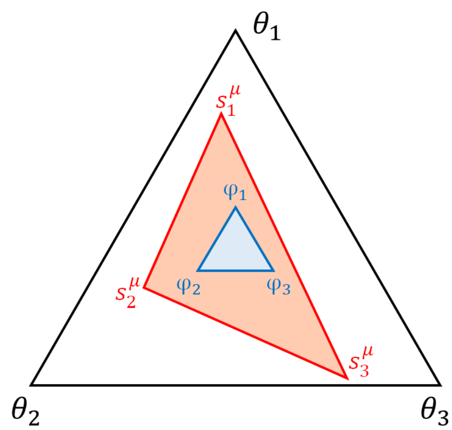

Figure B.1: A signal chunk in a simplex with 3 states

convex hull of these μ-chunks forms an isotropically scaled simplex co-centered with the state-space simplex. The following figure illustrates the existence of such μ-chuck simplex, denoted by the set of ϕ_i.

B.1.14 Argument of Proposition 2

With the result in Proposition 1, if agents already know the true state is either θ_i or θ_j in a finite state space model, then agents can learn the true state asymptotically by operating a QGB learning whether the true state is θ_i or θ_j same as the binary-state, finite-signal QGB proposed in Proposition 1. Hence, if we can construct a QGB which determines the continuation/termination of the current cascade by verifying the state corresponds to the current cascade is more likely to occur than any other state via a pairwise check, then we can learn the true state asymptotically if checking each pair of states does not affect the likelihood ratio of other states (not in this pair). In other words, the process of checking a pair of states is neutral to other states. Here, we will provide a high-level construction of the QGB that can be used to learn in a finite state space model.

First, given the result of Lemma 11, we can construct probabilistic μ-chunks which is an isotropically scaled simplex co-centered with the state-space simplex. Now, if we use the QGB developed in Proposition 1 with the probabilistic μ-chunks signals to determine the continuation or termination of the current cascade (assuming on θ_i) because of the belief of a particular state (assuming on θ_j) not in the current cascade, the likelihood ratio of other pair of states, assuming θ_k and θ_l, will stay unchanged. This is because these probability μ-chunks signals forms an isotropically scaled simplex co-centered with the state-space simplex, and every probability μ-chunks signal (assuming μ_s^j) has $\mathbb{P}(\theta_k|\mu_s^j) = \mathbb{P}(\theta_l|\mu_s^j)$ for every pair (θ_k, θ_l) with $k, l \neq i$. Thus, updating the likelihood ratio of θ_j over θ_i via these μ-chunk signals in the window to determine whether the agent should stop a cascade because of θ_j will not change the likelihood ratio of θ_k over θ_l for every $k, l \neq j$. Hence, if we only consider two states, θ_i in the current cascade and θ_j not in the cascade, using the QGB developed for Proposition 1 can asymptotically learn the state θ_j when θ_j is the true state. Now, given that this checking procedure is neutral to any other states not checked in current sequence of questions (where the responses generates probabilistic $\mu-$chunk signals), we can periodically perform different QGBs to determine the continuation or termination of the $\bar{a}_i$ cascade for every pair of state (θ_i, θ_j), $j \neq i$. Since Proposition 1 guarantees that every wrong cascades will be stopped and some correct cascade will last forever in each pair of check, when the prior is uniform and the signal space is distinguishable, asymptotic learning is achievable.

B.1.15 Proof of Lemma 12

First, questions asked by the first four agents are feasible. For every agent with index $n > 5$, she will ask questions when she gets s_a or s_b and then will forward the past positive/negative response. Therefore every agent's question when she gets s_a, s_b, or s_B can be answered because of her immediate predecessor's effort on asking the question just for forwarding the information to her (when her immediate predecessor gets s_a or s_b). This QGB is now feasible.

To show the incentive compatibility, we can immediately get that the first three agents' questions maximize their expected utility. Hence, we only need to verify this for the fourth agent and agents with index $n \geq 5$.

For agent 4, when she gets s_B, she can overturn the $\bar{A}$ action when the third agent also gets s_B, and it is contained as a part of her question. When she get s_B and agent 3 gets s_b, the positive response tells her that the first two agents are one of the two cases $s_1 s_2 \in \{s_a s_a, s_A s_b\}$, while the negative response tells her $s_1 s_2 \in \{s_A s_A, s_A s_a, s_a s_A\}$. To maximize her expected payoff, she should take action $\bar{B}$ while getting a positive response and take $\bar{A}$ while getting a negative response. When agent 3 gets s_a and agent 4 gets a positive response, she is in one of the two cases, where $s_1 s_2 s_3 s_4 \in \{s_a s_a s_a s_B, s_A s_b s_a s_B\}$. By our setting $\frac{\mathbb{P}(\theta=B|s_B)}{\mathbb{P}(\theta=A|s_B)}\left(\frac{\mathbb{P}(\theta=B|s_a)}{\mathbb{P}(\theta=A|s_a)}\right)^3 = 1$, taking action $\bar{B}$ maximizes her expected utility. However, when she gets a negative response from agent 3, she should take action $\bar{A}$. By checking all the circumstances, we verify that the action designed for agent 4 in this QGB is incentive compatible.

Last, we need to check all agents with index $n \geq 5$. When agent n gets a positive response, it implies that agent 3 to agent $n-1$ are either gets s_a or s_b. Since s_a, s_b happen with equal probability and $\mathbb{P}(A|s_a) = \mathbb{P}(B|s_b)$, every agent that gets a positive response has the same posterior belief as agent 4 getting a positive response. However, getting a negative response, at agent n, will give her a strictly lower posterior belief (that the true state of the world is B) than agent $n-1$. Therefore, taking action $\bar{B}$ only when she gets a positive response maximizes her expected utility. Finally, we can claim that the haste-waste QGB is feasible and incentive compatible.

B.1.16 Proof of Proposition 3

To prove the incentive compatibility, we first prove every agent n's strategy choosing $\mathcal{B}_n = \{k\}$ with $|\mathcal{B}_k| = 0$ is a weakly dominated strategy. Suppose agent n queries agent k with $|\mathcal{B}_k| = 0$, the (private[2]) information she can know from agent k is a particular agent's (agent k here) private signal. However, if she queries another agent m with $|\mathcal{B}_m| = 1$, she can also know a particular agent's (agent m here) private signal. Besides, when she queries agent m, she can also query the information of the response of agent m's question. Hence, querying agent m weakly dominates querying agent k.

Next, we prove querying the most recent predecessor m with $|\mathcal{B}_m| = 1$ weakly dominates querying any other earlier predecessors. Given the rationality of the predecessors, if an agent wants to ask the predecessor m' with $|\mathcal{B}_{m'}| = 1$ and $m' < m$, she can ask the successor of m', called m'', who queries agent m' to get the same information under the capacity constraint. However, querying agent m'' weakly dominates querying agent m' because agent m'' possesses her private signal unknown to agent m' and observes a longer realized history (because she arrives after m'). Using the same logic, querying m'''s successor who queries her also weakly dominates querying m'' directly. Thus, querying the most recent predecessor m with $|\mathcal{B}_m| = 1$ weakly dominates querying any other earlier predecessors.

Last, although querying the most recent predecessor m with $|\mathcal{B}_m| = 1$ is a weakly dominate strategy for agent n, it is not possible when there is no predecessor m with $|\mathcal{B}_m| = 1$. In this case, asking the private signal of an arbitrary agent who is in the current cascade is an incentive compatible strategy. Hence, asking the first agent in the current cascade is incentive compatible and Proposition 3 holds.

B.2 Technical Claims and Calculations

B.2.1 Proof of Corollary 2

We proof this Corollary by contradiction. Suppose there exists a pair of state θ_i, θ_j such that for every signal $s \in \mathcal{S}$, we have $\mathbb{P}(\theta_i|s) = \mathbb{P}(\theta_j|s)$. Then for any length-$n$

[2]The information cannot be inferred from the history realization.

sequence $(s_n) \in \mathcal{S}^n$, we will have $\mathbb{P}(\theta_i|(s_n)) = \mathbb{P}(\theta_j|(s_n))$. This violates our definition of distinguishability in Definition 20.

B.2.2 Proof of Claim 1

Proof. First, we want to show that every pair of pooling information sets are disjoint. If two pooling information sets E, F are not disjoint, if $\{E \cap F\} \subset E$, or $\{E \cap F\} \subset F$, it will violate the definition that there is no information that can be delivered from the current topology to partition E or F. Therefore, the only possibility is that $\{E \cap F\} = E = F$, which means E, F both indicate the same pooling information set.

Now, the second statement in the definition of pooling information set guarantees that every private signal sequence must belong to at least one pooling information set. Given the above proof that pooing information sets are disjoint, pooling information sets partition the information space. □

B.2.3 Proof of Claim 2

Proof. Since the given (length t) history realization has a positive probability $\geq (1-p)^t$, this QGB will lead to a wrong action with a probability no less than $(1-p)^t$. Hence, asymptotic learning is not achievable. □

B.2.4 Proof of Claim 3

Suppose an agent helps with a known probability r, her successor(s) will get an appropriate response with probability r but a non-informative response with probability $1-r$ with the distribution of response of non-helpful responses well-defined. This can be modeled as a noisy channel where an information can be correctly transmitted with probability r. Given the transformation of questions from perfect channel to noisy channel discussed in Section 2.8.2, we can transform every successor's question from a perfect channel to a noisy channel. Hence, Claim 3 stands here.

B.2.5 Proof of Claim 4

Given the signal space is distinguishable, for every state θ_i, there must exists a sequence of signals s such that $\mathbb{P}(\theta_i|s) > \mathbb{P}(\theta_j|s)\ \forall j \neq i$. Now, for each state θ_i, we can choose a sequence of signals satisfying $\mathbb{P}(\theta_i|s) > \mathbb{P}(\theta_j|s)\ \forall j \neq i$ and set it to be an element in the $\mu-$chuck signal set μ_s. Since every vertex of the $|\Theta|$-simplex (which represents the true state is a particular state w.p. 1) has a corresponding $\mu-$chuck signal in μ_s such the distance between the vertex and the $\mu-$chuck signal is less than the distance between the vertex and the centroid of the $|\Theta|$-simplex, the convex-hull of this $\mu-$chuck signal set must contains the centroid of the $|\Theta|$-simplex as an inner point.

B.2.6 Proof of Claim 9

First, we note that the occurrence of active agents means that the cascade is still ongoing and the question guidebook operational. Then assume that active agent a^k is in state $G_j \in \mathcal{G}^*$, where $\mathcal{G}^*$ is defined as the set of states at which she can stop the cascade.

We know for sure that the likelihood ratio of agent a^k+1 at state G_1, denoted by $l_{G_1}^{a^k+1}$ is less than $\frac{1-p}{p}$. The reason is that G_1 at a^k+1 contains two types of events. The first types is where a^k is in $G_k \in \mathcal{G}^*$ but receiving an observed majority. Given $l_{G_i}^{a^k} < 1$ for all $G_i \in \mathcal{G}^*$, $(l_{G_1}^{a^k+1}|a^k \text{ at } G_i \in \mathcal{G}^*) < \frac{1-p}{p}$; note that agent a^k does not stop the cascade. The other type is a^k was at $G_i \notin \mathcal{G}^*$. Here also we get $(l_{G_1}^{a^k+1}|a^k \text{ at } G_i \notin \mathcal{G}^*) < \frac{1-p}{p}$ because a^k cannot stop the cascade even after receiving observed minority in any of those states.

Therefore, we can conclude that $l_{G_1}^{a^k+1} < \frac{1-p}{p}$, which guarantees that agent a^k+1 is a silent agent, and not an active agent.

B.2.7 Claim 10 and its proof

Claim 10. *The likelihood ratio* $l_{G_1}^{H_{a^k+1}}$ *is a strictly decreasing function of* k.

Proof. First, $\frac{h^+(m)}{h^-(m)}$ is a strictly increasing function of m (please see calculations in Appendix B.2.14). Since $h^+(1) > h^-(1) > 0$ and given the form in (2.5), $l_{G_1}^{H_{a^k+1}}$ is an strictly decreasing function. $\square$

B.2.8 Claim 11 and its Proof

Claim 11. $\ln\left(\dfrac{1-h^-(s)}{1-h^+(s)}\right)$ *can be upper-bounded by* $p^s + p^{2s}$.

Proof. First, $\ln\left(\dfrac{1-h^-(s)}{1-h^+(s)}\right) \leq \ln\left(\dfrac{1}{1-h^+(s)}\right)$.

Then, using the upper bound in Appendix B.2.13 and Taylor's expansion, we know that

$$\ln\left(\frac{1}{1-h^+(s)}\right) \leq -\ln(1-p^s) \leq p^s + p^{2s}.$$

$\square$

B.2.9 Upperbound on m^k

Finding an upper bound of m^k is analogous to finding a sequence $\underline{w}_s$ such that

$$\left|\left\{m^k = s | l_{G_1}^{H_{a^{k-1}+1}} \frac{h^+(s)}{h^-(s)} \geq 1\right\}\right| \geq \underline{w}_s \ \forall\ s \in \mathbb{N}.$$

The following calculations derive the value of $\underline{w}_s$.

First, the number of $m^k = s$ is equivalent to difference of indices between the k^{th} active agent to the next active agent with index j such that $l_{G_1}^{H_{a^j+1}} \frac{h^+(s)}{h^-(s)} < 1$. Hence, we get the following equation:

$$\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| = \min\{n | l_{G_1}^{H_{a^{k+n}+1}} \frac{h^+(s)}{h^-(s)} < 1\}.$$

Now, using the recursive formula of likelihood ratio stated in (2.6), we get the following equation:

$$\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| = \min\left\{n | l_{G_1}^{H_{a^k+1}} \left(\frac{1-h^+(s)}{1-h^-(s)}\right)^n \frac{h^+(s)}{h^-(s)} < 1\right\}.$$

We know that $l_{G_1}^{H_{a^{k-1}+1}} \frac{h^+(s)}{h^-(s)} > 1$ because $m^k = s$ and $m^{k-1} = s-1$, so taking logarithms we obtain

$$\begin{aligned}
&\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| \\
&= \min\left\{n | n \ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right) > \ln\left(\frac{h^+(s)}{h^-(s)}\right) + \ln(l_{G_1}^{H_{a^k+1}})\right\} \\
&\geq \min\left\{n | n \ln\left(\tfrac{1-h^-(s)}{1-h^+(s)}\right) > \ln\left(\tfrac{h^+(s)}{h^-(s)}\right) + \ln\left(l_{G_1}^{H_{a^k+1}}\right) - \ln\left(l_{G_1}^{H_{a^{k-1}+1}}\right) - \ln\left(\tfrac{h^+(s-1)}{h^-(s-1)}\right)\right\}
\end{aligned}$$

Using the recursive form in (2.6) once again, we get

$$\begin{aligned}
&\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| \\
&\geq \min\left\{n | n \ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right) > \ln\left(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\right) + \ln\left(\frac{1-h^+(s-1)}{1-h^-(s-1)}\right)\right\} \\
&\geq \min\left\{n | (n+1) \ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right) > \ln\left(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\right)\right\}
\end{aligned}$$

From this point onwards, we need to compute a lower bound of $\ln\left(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\right)$, and an upper bound of $\ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right)$. For ease of reading we put the derivations of the calculations of the lower bound of $\ln\left(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\right)$ in Appendix B.2.14, and the upper bound of $\ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right)$ using the Claim 11 in Appendix B.2.8 (by using the fact that $\ln\left(\frac{1-h^-(s)}{1-h^+(s)}\right) < \ln\left(\frac{1}{1-h^+(s)}\right)$ first).

Combining the result in AppendixB.2.14, which shows $\ln\left(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\right) \geq c_1(p)$, where $c_1(p)$ is only a function of p, and the result in Appendix B.2.8, $\ln(\dfrac{1-h^-(s)}{1-h^+(s)}) \leq$

$p^s + p^{2s}$, we can pick $\underline{w}_s$ as

$$\underline{w}_s = \frac{\ln(c_1(p))}{p^s + p^{2s}}. \tag{B.2}$$

B.2.10 Calculation of the probability that a wrong cascade will be stopped

First, we use the lower bound of $h^+(m)$ derived in Appendix B.2.13 to get the first inequality

$$\mathbb{P}(\text{A wrong cascade will stop}) = 1 - \Pi_{k=1}^{\infty}(1 - h^+(m^k)) \geq 1 - \Pi_{k=1}^{\infty}(1 - p^{m^k}).$$

Now, we use the sequence $\underline{w}_s$ that upperbounds m^k to replace the product and apply the lower-bound result in B.2.14 for large enough J.

$$\begin{aligned}
& \mathbb{P}(\text{A wrong cascade will stop}) \\
\geq\; & 1 - A(p)\Pi_{n=J}^{\infty}(1-p^n)^{\frac{\ln(c_1(p))}{p^{n-1}(1+p^{n-1})}} \\
\geq\; & 1 - A(p)\exp\Big(p\sum_{n=0}^{\infty}\frac{\ln(c_1(p))}{p^n(1+p^{n-1})}(-p^n)\Big) \\
=\; & 1 - A(p)\exp\Big(-p\ln(c_1(p))\sum_{n=0}^{\infty}\frac{1}{1+p^n}\Big) = 1,
\end{aligned}$$

where $A(p) = \Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1 - p^{m^k})$.

B.2.11 Lowerbound on m^k

Finding a lower bound of m^k is analogous to finding a sequence of $\bar{w}_s$ such that $|\{m^k = s | l_{G_1}^{H_{a^{k-1}+1}} \frac{h^+(s)}{h^-(s)} \geq 1\}| \geq \bar{w}_s$ for all $s \in \mathbb{N}$. In contrast to what we did in computing $\underline{w}_s$, we now want to upperbound the following equation:

$$\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| = \min\left\{n | l_{G_1}^{H_{a^k+1}} \Big(\frac{1-h^+(s)}{1-h^-(s)}\Big)^n \frac{h^+(s)}{h^-(s)} < 1\right\}.$$

Given that $m^k = s$ and $m^{k-1} = s-1$, $l_{G_1}^{H_{a^k+1}} \frac{h^+(s-1)}{h^-(s-1)} < 1$ because $m^k = s$ and $m^{k-1} = s-1$, we can lower bound the equation and then take logarithms on both sides of the inequality to get

$$\begin{aligned}
&\left|\left\{m^t = s | l_{G_1}^{H_{a^t+1}} \frac{h^+(s)}{h^-(s)} \geq 1, t \geq k\right\}\right| \\
&= \min\left\{n \Big| n \ln\Big(\frac{1-h^-(s)}{1-h^+(s)}\Big) > \ln\Big(\frac{h^+(s)}{h^-(s)}\Big) + \ln(l_{G_1}^{H_{a^k+1}})\right\} \\
&\leq \min\left\{n \Big| n \ln\Big(\tfrac{1-h^-(s)}{1-h^+(s)}\Big) > \ln\Big(\tfrac{h^+(s)}{h^-(s)}\Big) + \ln\Big(l_{G_1}^{H_{a^k+1}}\Big) - \ln\Big(l_{G_1}^{H_{a^k+1}}\Big) - \ln\Big(\tfrac{h^+(s-1)}{h^-(s-1)}\Big)\right\} \\
&= \min\left\{n \Big| n \ln\Big(\frac{1-h^-(s)}{1-h^+(s)}\Big) > \ln\Big(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\Big)\right\}
\end{aligned}$$

Similarly, we can replace $\ln\Big(\frac{h^+(s)h^-(s-1)}{h^-(s)h^+(s-1)}\Big)$ by an upper bound $c_2(p)$ derived in Appendix B.2.14 for large enough s. However, unlike what we did for upper bounding m^k, now we keep the form $\ln\Big(\frac{1-h^-(s)}{1-h^+(s)}\Big)$ and bring this into the calculation of the probability that a right cascade will being stopped. In other words, $\bar{w}_s = \frac{\ln(c_2(p))}{\ln\big(\frac{1-h^-(s)}{1-h^+(s)}\big)}$.

B.2.12 Calculation of the probability that a right cascade will be stopped

Using the lower bound derived in Appendix B.2.11, we get the following inequality:

$$\begin{aligned}
&\mathbb{P}(\text{A right cascade will be stopped}) \\
&= 1 - \Pi_{k=1}^{\infty}(1 - h^-(m^k)) \\
&\leq 1 - A(p)\Pi_{n=J}^{\infty}(1 - h^-(n))^{\frac{\ln(c_2(p))}{ln\big(\frac{1-h^-(n)}{1-h^+(n)}\big)}} \\
&= 1 - A(p)\exp\Big(\sum_{n=J}^{\infty} \frac{\ln(c_2(p))}{\ln\big(\frac{1-h^-(n)}{1-h^+(n)}\big)} \ln(1-h^-(n))\Big),
\end{aligned}$$

where $A(p) = \Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1-h^{-}(k))$, and J is the same as in Appendix B.2.10.

Now, rest of the calculations are just using well-known upper and lower bounds on the logarithm. With calculations detailed in B.2.15, we get that

$$\mathbb{P}(\text{A right cascade will be stopped}) \leq 1 - B(p)\exp\Big(\sum_{n=J}^{\infty}\frac{1}{c_4(p)c_1(p)^n - 1}\Big),$$

where $B(p) = c_2(p)e^{c_3(p)}\Pi_{k=1}^{\min\{j|m^{j+1}=J\}}(1-h^{-}(k))$.

Now, using the fact that $c_1(p) > 1$ in Appendix B.2.14, for a large enough n, $|\frac{c_4(p)c_1(p)^n-1}{c_4(p)c_1(p)^{n+1}-1} - c_1(p)| \leq \epsilon$. Hence, by ratio test, we can claim that $\sum_{n=K}^{\infty}\frac{1}{c_4(p)c_1(p)^n-1}$ converges, and
$\mathbb{P}(\text{A right cascade will be stopped}) < 1$.

B.2.13 Calculation of the upper bound of $h^+(m)$

First, in this question guidebook, $h^+(m)$ has a closed form given by

$$h^+(m) = \begin{cases} \sqrt{\frac{p}{4-3p}}2^{-(m+1)}\left[\left(p+\sqrt{p(4-3p)}\right)^{m+1} - \left(p-\sqrt{p(4-3p)}\right)^{m+1}\right], \\ \text{where } m \text{ is even;} \\ \sqrt{\frac{p}{4-3p}}\left(\frac{p}{2}\right)^{\frac{m+1}{2}} \\ \quad\times\left[\left(2-p+\sqrt{p(4-3p)}\right)^{\frac{m+1}{2}} - \left(2-p-\sqrt{p(4-3p)}\right)^{\frac{m+1}{2}}\right], \\ \text{where } m \text{ is odd.} \end{cases}$$

With the closes form of $h^+(m)$, a simple bound of $\dfrac{h^+(m+1)}{h^+(m)} < p$ holds for all $m \in \mathbb{N}$.

For readers curious about the exact difference of $ph^+(m) - h^+(m+1)$, the form is $p^{m+2}\Big[{}_2F_1(-\frac{m}{2}-\frac{1}{2}, -\frac{m}{2}, -m-1, 4(1-\frac{1}{p})) - {}_2F_1(-\frac{m}{2}+\frac{1}{2}, -\frac{m}{2}, -m-1, 4(1-\frac{1}{p}))\Big]$, where ${}_2F_1(\cdot)$ is a hypergeometric function. We will not need this detail as getting the bound $\frac{h^+(m+1)}{h^+(m)} < p$ is good enough for our results.

B.2.14 Convergence and bounds of $\frac{h^+(m+1)}{h^-(m+1)} / \frac{h^+(m)}{h^-(m)}$

For simplicity of analysis, suppose m is even. Recall the closed form of $h^+(m)$ and $h^-(m)$ as follows:

$$h^+(m) = \sqrt{\tfrac{p}{4-3p}} 2^{-(m+1)}[(p+\sqrt{p(4-3p)})^{m+1} - (p-\sqrt{p(4-3p)})^{m+1}]$$

$$\begin{aligned} & h^-(m) \\ & = \sqrt{\tfrac{1-p}{1+3p}} 2^{-(m+1)} \left[(1-p+\sqrt{1+2p-3p^2})^{m+1} - (1-p-\sqrt{1+2p-3p^2})^{m+1}\right] \end{aligned}$$

Consider $\dfrac{h^+(m+2)}{h^-(m+2)} / \dfrac{h^+(m)}{h^-(m)}$.

$$\begin{aligned} & \frac{h^+(m+2)}{h^-(m+2)} \frac{h^-(m)}{h^+(m)} \\ = \; & \frac{(p+\sqrt{p(4-3p)})^{m+3} - (p-\sqrt{p(4-3p)})^{m+3}}{(1-p+\sqrt{1+2p-3p^2})^{m+3} - (1-p-\sqrt{1+2p-3p^2})^{m+3}} \\ \times \; & \frac{(1-p+\sqrt{1+2p-3p^2})^{m+1} - (1-p-\sqrt{1+2p-3p^2})^{m+1}}{(p+\sqrt{p(4-3p)})^{m+1} - (p-\sqrt{p(4-3p)})^{m+1}} \\ = \; & \frac{(p+\sqrt{p(4-3p)})^2}{(1-p+\sqrt{1+2p-3p^2})^2} \frac{1-\frac{(p-\sqrt{p(4-3p)})^{m+3}}{(p+\sqrt{p(4-3p)})^{m+3}}}{1-\frac{(p-\sqrt{p(4-3p)})^{m+1}}{(p+\sqrt{p(4-3p)})^{m+1}}} \frac{1-\frac{(1-p-\sqrt{1+2p-3p^2})^{m+1}}{(1-p+\sqrt{1+2p-3p^2})^{m+1}}}{1-\frac{(1-p-\sqrt{1+2p-3p^2})^{m+3}}{(1-p+\sqrt{1+2p-3p^2})^{m+3}}} \\ \leq \; & \frac{(p+\sqrt{p(4-3p)})^2}{(1-p+\sqrt{1+2p-3p^2})^2} \cdot \left[1 + \left(\frac{p-\sqrt{p(4-3p)}}{p+\sqrt{p(4-3p)}}\right)^{m+1} + \left(\frac{p-\sqrt{p(4-3p)}}{p+\sqrt{p(4-3p)}}\right)^{m+2} \right. \\ & \left. - \left(\frac{1-p-\sqrt{1+2p-3p^2}}{1-p+\sqrt{1+2p-3p^2}}\right)^{m+1} - \left(\frac{1-p-\sqrt{1+2p-3p^2}}{1-p+\sqrt{1+2p-3p^2}}\right)^{m+2}\right] \end{aligned}$$

Since $\frac{h^+(m+2)}{h^-(m+2)} \frac{h^-(m)}{h^+(m)}$ converges to $\frac{(p+\sqrt{p(4-3p)})^2}{(1-p+\sqrt{1+2p-3p^2})^2}$ exponentially fast, for large enough m, we have $\left|\dfrac{h^+(m+2)}{h^-(m+2)} \dfrac{h^-(m)}{h^+(m)} - \dfrac{(p+\sqrt{p(4-3p)})^2}{(1-p+\sqrt{1+2p-3p^2})^2}\right| < \epsilon$.

Given the fact that $\frac{(p+\sqrt{p(4-3p)})^2}{(1-p+\sqrt{1+2p-3p^2})^2} > 1$ for all $p > 0.5$, there exists functions $c_1(p), c_2(p) > 1$ such that $c_1(p) < \dfrac{h^+(m+2)}{h^-(m+2)} \dfrac{h^-(m)}{h^+(m)} < c_2(p)$ for all m greater than some natural number $N \in \mathbb{N}$.

B.2.15 Bound the probability of stopping a right cascade

$$
\begin{aligned}
&\mathbb{P}(\text{a right cascade will be stopped})\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=K}^{\infty}\frac{-h^-(n)}{\ln(1+\frac{h^+(n)-h^-(n)}{1-h^+(n)})}\Big)\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=K}^{\infty}\frac{-h^-(n)}{\frac{h^+(n)-h^-(n)}{1-h^+(n)}(1-\frac{h^+(n)-h^-(n)}{1-h^+(n)})}\Big)\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=K}^{\infty}\frac{-1}{(\frac{h^+(n)}{h^-(n)}-1)(1-\frac{h^+(n)-h^-(n)}{1-h^+(n)})\frac{1}{1-h^+(n)}}\Big)\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=1}^{\infty}\frac{-1}{(\frac{h^+(n)}{h^-(n)}-1)(1-\frac{p^n-0}{1-p^n})\frac{1}{1-p^n}}\Big)\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=1}^{\infty}\frac{-1}{(\frac{h^+(n)}{h^-(n)}-1)\frac{1-2p^n}{(1-p^n)^2}}\Big)\\
&\leq 1 - A(p)\exp\Big(\ln(c_2(p))\sum_{n=1}^{\infty}\frac{-1}{(\frac{h^+(n)}{h^-(n)}\frac{1-2p^n}{(1-p^n)^2}-1}\Big)\\
&\leq 1 - B(p)\exp\Big(\sum_{n=K}^{\infty}\frac{1}{c_4(p)c_1(p)^n-1})\Big),
\end{aligned}
$$

where $B(p) = c_2(p)e^{c_3(p)}\Pi_{k=1}^{\min\{j|m^{j+1}=K\}}(1-h^-(k))$

APPENDIX C

Appendices of Chapter III

C.1 Additional Discussions of Chapter III: Trade-off between efficiency and productivity

In Bayesian persuasion problems, the signal strength (posterior beliefs of states under signals) and the signal distribution represent two different notions. The signal strength of a signal represents how efficiently[1] this signal can persuade the receiver. The signal distribution represents how productive a signal strategy is towards deriving benefit for the sender under the current prior. However, in many (single-trial) Bayesian persuasion problems, e.g., binary-state space Bayesian persuasion, Bayesian persuasion with senders having lexicographic preference on actions, it is hard to evaluate the role that each notion plays in the optimal signaling strategy respectively. The reason is the insufficient freedom on experiment design. In these problems, the sender can always use the (set of) most efficient signals in optimal signaling strategies. Given the (set of) most efficient signals, the most productive distribution can be

[1]A signal ω is more efficient than ω' if the following statement is true: when we use the same total probability of the receiver' preferred state, which corresponds to the sender's suggested action of the signal (θ_1 in this work), to implement both signals, the receiver's total probability of taking suggested action (ϕ_1) while receiving ω is higher than the receiver's total probability of taking suggested action while receiving ω'.

solved via a greedy approach[2]. However, if we consider multi-phase trials, the trade-off between efficiency and productivity becomes the sender's primary issue (even in the binary-state space Bayesian persuasion problem with a sender with lexicographic preferences on actions) because of the enlarged signal space. That is to say, if we want to get a better understanding of how the signal strength and the signal distribution plays a role in Bayesian persuasion problems, it is not necessary to build complex models with stylized utility functions or an increased state space. Here we provide an example where the most efficient signal is not used in the sender's optimal signaling strategy owing to the trade-off between efficiency and productivity.

Example:

Given two pre-determined phase-II experiments A, B with $(q_{A1}, q_{A2}) = (0.8, 0.5)$ and $(q_{B1}, q_{B2}) = (0.75, 0.15)$, we can derive the persuasion potential of E_A and E_B, where $E_A = (1.6, 1.4)$ and $E_B = (1.5, \frac{9}{7})$. Under this persuasion potential, the most efficient signal will be a signal $\bar{\omega}$ with $(\mathbb{P}(\theta_1|\bar{\omega}, E_A), \mathbb{P}(\theta_2|\bar{\omega}, E_A)) = (\frac{5}{13}, \frac{8}{13})$. As we can see, when the sender uses this signal, the sender is using the partial strategy α_A in her signaling strategy, and the IC constraint for α_A is tight. Here, we provide an example where the signal $\bar{\omega}$ is not in the support of the sender's optimal signaling strategy.

When the prior belief on state θ_1, p, is $\frac{2}{3}$, we can solve for the optimal signaling strategy for the sender using the approach developed in Section 3.3. The optimal signaling strategy when $p = \frac{2}{3}$ is (β_A, α_B) with $\mathbb{P}(E_A, \theta_1) = \frac{55}{84}$ and $\mathbb{P}(E_A, \theta_2) = \frac{22}{84}$. The persuasion ratio of this optimal signaling strategy is $\frac{471}{336} \approx 1.402$ and the sender's expected utility is $\frac{157}{168} \approx 0.935$. As we can see, the most efficient signal $\bar{\omega}$ is not in the support of the optimal signaling strategy.

However, if the sender insists on using the signal $\bar{\omega}$, her best signaling strategy

[2] Given the sender's preference on actions, the sender has a preference on signals (in terms of the efficiency). In this set of Bayesian persuasion problems, the sender can greedily pick the most efficient signal where the prior allows iteratively to come up the optimal signaling strategy.

when $p = \frac{2}{3}$ is (α_A, β_B) with $\mathbb{P}(E_A, \theta_1) = \frac{5}{46}$ and $\mathbb{P}(E_A, \theta_2) = \frac{8}{46}$. The corresponding persuasion ratio is $\frac{123}{92} \approx 1.337$ and the sender's expected utility is $\frac{41}{46} \approx 0.891$, which is lower than $\frac{157}{168}$ under the optimal signaling strategy.

In the above example, we discuss a case where using the most efficient signal is non-optimal in a two-phase trial's Bayesian persuasion problem. This example suggests that the sender has to consider the trade-off between efficiency and productivity to derive the optimal signaling strategy in multi-phase trials, which may not be necessary for single-trial Bayesian persuasion problems.

C.2 Proofs of Chapter III

C.2.1 Proof of Lemma 13

We derive the optimal signaling strategy in each scheme and the participants expected utility to validate this claim. First, under different prior regions, the sender's optimal[3] strategy in a single-phase trial is to commit to the following signaling strategy:

$p \leq \frac{1}{2}$	θ_1	θ_2
ω_1	p	p
ω_2	0	1-2p

$p > \frac{1}{2}$	θ_1	θ_2
ω_1	p	1-p
ω_2	0	0

Given the above signaling strategy, the sender and receiver's expected utility can be calculated in the following table:

Table C.1: Table: Sender and Receiver's expected utility in single-phase trials

	$\mathbb{E}[U_S(a, \theta)]$	$\mathbb{E}[U_R(a, \theta)]$
$p \leq \frac{1}{2}$	2p	1-p
$p > \frac{1}{2}$	1	p

In a two-phase trial with a trivial experiment in phase II, we now solve for the sender's optimal signaling strategy. Without loss of generality, let's assume that the

[3]This is to maximize the probability that the receiver will take action a_1.

trivial experiment is the experiment E_A, $q_{A1} = q_{A2}$, and another experiment E_B has $q_{B1} > q_{B2}$. Given these two experiments, the sender's optimal signaling strategy can be derived via solving the following optimization problem:

$$\max_{p_1,p_2} \mathbb{P}(\omega_A)\mathbf{1}_{\{\mathbb{P}(\theta_1|\omega_A)\geq\frac{1}{2}\}} + \mathbb{P}(\omega_{BP})\mathbf{1}_{\{\mathbb{P}(\theta_1|\omega_{BP})\geq\frac{1}{2}\}} + \mathbb{P}(\omega_{BF})\mathbf{1}_{\{\mathbb{P}(\theta_1|\omega_{BF})\geq\frac{1}{2}\}}$$

$$\begin{aligned}
\text{subject to } & \mathbb{P}(\omega_A) = pp_1 + (1-p)p_2 \\
& \mathbb{P}(\omega_{BP}) = p(1-p_1)q_{B1} + (1-p)(1-p_2)q_{B2} \\
& \mathbb{P}(\omega_{BF}) = p(1-p_1)(1-q_{B1}) + (1-p)(1-p_2)(1-q_{B2}) \\
& \mathbb{P}(\theta_1|\omega_A) = \frac{pp_1}{pp_1 + (1-p)p_2} \\
& \mathbb{P}(\theta_1|\omega_{BP}) = \frac{p(1-p_1)q_{B1}}{p(1-p_1)q_{B1} + (1-p)(1-p_2)q_{B2}} \\
& \mathbb{P}(\theta_1|\omega_{BF}) = \frac{p(1-p_1)(1-q_{B1})}{p(1-p_1)(1-q_{B1}) + (1-p)(1-p_2)(1-q_{B2})}
\end{aligned}$$

After solving this optimization problem, we can derive the sender and receiver's expected utility in Table C.2:

Table C.2: Sender and Receiver's expected utility under a trivial experiment

	$\mathbb{E}[U_S(a,\theta)]$	$\mathbb{E}[U_R(a,\theta)]$
$p \leq \frac{1}{2}$	2p	1-p
$p > \frac{1}{2}$	1	p

Note that when $q_{B1} = 1$, the sender may have multiple optimal signaling strategies when the prior p is low. However, no matter which strategy the sender chooses, the sender and the receiver's expected utilities are the same.

After checking Table C.1 and Table C.2, we validate the fact that these two schemes will provide the sender and the receiver the same expected utility under optimal signaling strategies.

C.2.2 Proof of Lemma 14

We prove this lemma by explicitly verifying which strategies maximize the sender's expected utility. First, let's recall the IC requirement for each partial strategies listed in Section 3.3.3.

Partial strategy	IC requirement	Partial strategy	IC requirement
α_A	$p_1 \geq \frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2$	α_B	$1-p_1 \geq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}(1-p_2)$
β_A	$p_1 \geq \frac{1-p}{p}\frac{1-q_{A2}}{1-q_{A1}}p_2$	β_B	$1-p_1 \geq \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}}(1-p_2)$
γ_A	$p_1 \leq \frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2$	γ_B	$1-p_1 \leq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}(1-p_2)$

1. **When** $S \in (\alpha_A, \gamma_B)$ **or** $S \in (\gamma_A, \alpha_B)$

   Suppose S is a (α_A, γ_B) type strategy, the sender's expected utility under S is $pp_1q_{A1} + (1-p)p_2q_{A2}$. Given the definition of persuasion ratio, $pp_1q_{A1} + (1-p)p_2q_{A2}$ is maximized when $p_1 = 1, p_2 = \frac{pp_1}{1-p}\frac{q_{A1}}{q_{A2}}$ or $p_1 = 1, p_2 = 1$. Hence, in the former scenario where $p_2 = \frac{pp_1}{1-p}\frac{q_{A1}}{q_{A2}}$, IC requirement of α_A is tight in S. In the later scenario, the sender sends signal ω_A with probability one because the prior is larger enough to make the posterior belief of θ_1 conditional on ω_{AP} to be $\geq \frac{1}{2}$. Hence, the lemma holds when S is a (α_A, γ_B) type strategy.

   Similarly, if S is a (γ_A, α_B) type strategy, the sender's expected utility is $p(1-p_1)q_{B1} + (1-p)(1-p_2)q_{B2}$. When sender's expected utility is maximized at S, the parameter of (p_1, p_2) is either $p_1 = 0, p_2 = 1 - \frac{pq_{B1}}{(1-p)q_{B2}}$ or $p_1 = 0, p_2 = 0$. The IC requirement of α_B is tight in the former case and $\mathbb{P}(\omega_2) = 1$ in the later case. Thus, the lemma holds when S is a (α_B, γ_A) type strategy.

2. **When** $S \in (\beta_A, \gamma_B)$ **or** $S \in (\gamma_A, \beta_B)$

   Suppose S is a (β_A, γ_B) type strategy, the sender's expected utility is $pp_1 + (1-p)p_2$. The strategy $S \in (\gamma_A, \beta_B)$ that maximizes the sender's utility is at $p_1 = 1, p_2 = \frac{p(1-q_{A1})}{(1-p)(1-q_{A2})}$ when $p \leq \frac{1-qA2}{2-q_{A2}-q_{A1}}$ or at $p_1 = p_2 = 1$ when $p > \frac{1-qA2}{2-q_{A2}-q_{A1}}$. In the former case, the IC requirement of β_A is tight. In the later case, the sender will send signal ω_1 with probability 1. Hence, the lemma holds when S is a (β_A, γ_B) type strategy.

3. **When** $S \in (\alpha_A, \beta_B)$ **or** $S \in (\beta_A, \alpha_B)$

   Suppose S is a (α_A, β_B) type strategy, the sender's expected utility is

$$
\begin{aligned}
&pp_1 q_{A1} + p(1-p_1) + (1-p)p_2 q_{A2} + (1-p)(1-p_2)) \\
=&1 - p_1 p(1-q_{A1}) - p_2(1-p)(1-q_{A2})
\end{aligned}
$$

   when the inequality of IC requirements $\frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2 \leq p_1 \leq \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}}p_2 + (1 - \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}})$ is satisfied. Now, let's discuss the strategy S under different prior region.

   (a) When $p < \frac{q_{A2}}{q_{A1}+q_{A2}}$, the IC inequality will never be satisfied. Hence, there is no incentive compatible strategy in the set of strategy $S \in (\alpha_A, \beta_B)$ in this region.

   (b) When $\frac{q_{A2}}{q_{A1}+q_{A2}} \leq p \leq \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$, the sender's expected utility is maximized when both inequalities of IC requirements are tight, i.e.,

$$
(p_1, p_2) = \left(\frac{1 - \frac{1-p}{p}\frac{1-q_{B2}}{1-q_{B1}}}{1 - \frac{q_{A1}}{q_{A2}}\frac{1-q_{B2}}{1-q_{B1}}}, \frac{\frac{p}{1-p} - \frac{1-q_{B2}}{1-q_{B1}}}{\frac{q_{A2}}{q_{A1}} - \frac{1-q_{B2}}{1-q_{B1}}} \right).
$$

   (c) When $p > \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$, the sender's expected utility is maximized when $p_1 = p_2 = 0$. In other words, the sender will send signal ω_2 w.p.1.

   After analyzing prior regions (a), (b), and (c), we verify the lemma holds when S is a (α_A, β_B) type strategy.

   Similarly, we can use the same approach to verify when S is a (β_A, α_B) type strategy.

Suppose S is a (β_A, α_B) type strategy, the sender's expected utility is

$$\begin{aligned} & pp_1 + (1-p)p_2 + p(1-p_1)q_{B1} + (1-p)(1-p_2)q_{B2} \\ =& 1 - (1-p_1)p(1-q_{B1}) - (1-p_2)(1-p)(1-q_{B2}) \end{aligned}$$

when the inequality of IC requirements $\frac{1-p}{p}\frac{1-q_{A2}}{1-q_{A1}}p_2 \leq p_1 \leq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}p_2 + (1 - \frac{1-p}{p}\frac{q_{B2}}{q_{B1}})$ is satisfied. Now, let's discuss where the strategy S at under different prior region.

(a) When $p < \frac{q_{B2}}{q_{B1}+q_{B2}}$, the IC inequality will never be satisfied. Hence, there is no incentive compatible strategy in the set of strategy $S \in (\alpha_A, \beta_B)$ in this region.

(b) When $\frac{q_{B2}}{q_{B1}+q_{B2}} \leq p \leq \frac{1-p}{p}\frac{q_{A2}}{q_{A1}}(1-p_2)$, the sender's expected utility is maximized when both inequalities of IC requirements are tight, i.e.,

$$(p_1, p_2) = \left(\frac{1 - \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}}{1 - \frac{1-q_{A1}}{1-q_{A2}}\frac{q_{B2}}{q_{B1}}}, \frac{\frac{p}{1-p} - \frac{q_{B2}}{q_{B1}}}{\frac{1-q_{A2}}{1-q_{A1}} - \frac{q_{B2}}{q_{B1}}} \right).$$

(c) When $p > \frac{1-q_{A2}}{2-q_{A1}-q_{A2}}$, the sender's expected utility is maximized when $p_1 = p_2 = 1$. In other words, the sender will send signal ω_1 w.p.1.

After analyzing prior regions (a), (b), and (c), we verify the lemma holds when S is a (β_A, α_B) type strategy.

4. **When** $S \in (\alpha_A, \alpha_B)$

Suppose S is a (α_A, α_B) type strategy, the sender's expected utility is

$$pp_1q_{A1} + p(1-p_1)q_{B1} + (1-p)p_2q_{A2} + (1-p)(1-p_2))q_{B2}$$

when the inequality of IC requirements $\frac{1-p}{p}\frac{q_{A2}}{q_{A1}}p_2 \leq p_1 \leq \frac{1-p}{p}\frac{q_{B2}}{q_{B1}}p_2 + (1 - \frac{1-p}{p}\frac{q_{B2}}{q_{B1}})$ is satisfied. Now, let's discuss the strategy S under different prior region.

(a) When $p < \min\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$, the IC inequality will never be satisfied. Hence, there is no incentive compatible strategy in the set of strategy $S \in (\alpha_A, \alpha_B)$ in this region.

(b) When $\min\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\} \leq p \leq \max\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$, the sender's expected utility is maximized when both inequalities of IC requirements are tight, i.e.,

$$(p_1, p_2) = \left(\frac{1-\frac{1-p}{p}\frac{q_{B2}}{q_{B1}}}{1-\frac{q_{A1}}{q_{A2}}\frac{q_{B2}}{q_{B1}}}, \frac{\frac{p}{1-p}-\frac{q_{B2}}{q_{B1}}}{\frac{q_{A2}}{q_{A1}}-\frac{q_{B2}}{q_{B1}}} \right).$$

(c) When $p > \max\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$, the sender's expected utility is maximized when exactly one IC requirement α_A or α_B is tight.

(For readers curious about which requirement will be tight, the requirement α_A is tight when $p+(1-p)(2q_{A1}-1)\frac{q_{A2}}{q_{A1}}+2pq_{B1}-(1-\frac{1-p}{p}\frac{q_{B2}}{q_{B1}})p(1-2q_{B1}) \geq 0$, otherwise, α_B is tight.)

After analyzing prior regions (a), (b), and (c), we verify the lemma holds when S is a (α_A, β_B) type strategy.

5. **When** $S \in (\beta_A, \beta_B)$

Suppose S is a (β_A, β_B) type strategy, the sender's expected utility is 1. However, the IC inequalities of this strategy can be satisfied only if the prior $p \geq \max\left\{\frac{1-q_{A2}}{2-q_{A1}-q_{A2}}, \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}\right\}$.

When $p < \min\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$, the IC inequality will never be satisfied. Hence, there is no incentive compatible strategy in the set of strategy $S \in (\alpha_A, \alpha_B)$ in this region.

C.2.3 Proof of Lemma 15

First, since (γ_A, γ_B) strategy has a persuasion ratio 0 when $\mathcal{S}_{\gamma_A,\gamma_B}(p) \neq \emptyset$, the statement in the lemma for each set of strategies: (α_A, γ_B), (β_A, γ_B), (γ_A, α_B), and (γ_A, β_B) is true. Hence, we only need to verify this lemma for the following four sets of strategies: (1) (α_A, α_B), (2) (α_A, β_B), (3) (β_A, α_B), and (4) (β_A, β_B).

Let's start with the type of strategies (α_A, α_B). The IC requirements of this type of strategy can be satisfied when $p \geq \min\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$, i.e., $\mathcal{S}_{\alpha_A,\alpha_B}(p) \neq \emptyset$ when $p \geq \min\left\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\right\}$. When the IC requirements are satisfied, the sender's expected utility under a strategy $\{(\alpha_A, \alpha_B), p_1, p_2\}$ is

$$\mathbb{E}_p[U_S|(\alpha_A, \alpha_B), p_1, p_2] = pp_1q_{A1} + p(1-p_1)q_{B1} + (1-p)p_2q_{A2} + (1-p)(1-p_2))q_{B2}.$$

Besides, the sender's expected utility under the strategy types $\{(\alpha_A, \gamma_B), p_1, p_2\}$ and $\{(\gamma_A, \alpha_B), p_1, p_2\}$ are the following:

$$\mathbb{E}[U_S(p|(\alpha_A, \gamma_B), p_1, p_2)] = pp_1q_{A1} + (1-p)p_2q_{A2},$$
$$\mathbb{E}[U_S(p|(\gamma_A, \alpha_B), p_1, p_2)] = p(1-p_1)q_{B1} + (1-p)(1-p_2))q_{B2}.$$

Given the result in Lemma 14, we know that $\mathbb{E}[U_S(p|(\alpha_A, \gamma_B), p_1, p_2)]$ can be represented by a scaled optimized signaling strategy under another prior u such that $xu = pp_1$ and $\mathbb{E}[U_S(p|(\alpha_A, \gamma_B), p_1, p_2)] = xuPR(\alpha, \gamma, u)$. (Because an optimized strategy $\arg\max_{S\in(\alpha_A,\gamma_B,p'_1,p'_2)} \mathbb{E}[U_R(S, p')]$ must have $p'_1 = 1$ in the optimized S.) Similarly, we can represent $\mathbb{E}[U_S(p|(\gamma_A, \alpha_B), p_1, p_2)]$ as a scaled optimized signaling strategy under another prior v such that $(1-x)v = p(1-p_1)$ and

$\mathbb{E}[U_S(p|(\gamma_A, \alpha_B), p_1, p_2)] = (1-x)vPR(\gamma, \alpha, u)$. Hence,

$$\begin{aligned}
&\max_{p_1,p_2} \mathbb{E}[U_S(p|(\alpha_A, \alpha_B), p_1, p_2)] \\
=&(\max_{p_1,p_2} pp_1q_{A1} + p(1-p_1)q_{B1} + (1-p)p_2q_{A2} + (1-p)(1-p_2))q_{B2}) \\
=&\max_{p_1,p_2} \left(\mathbb{E}[U_S(p|(\alpha_A, \gamma_B), p_1, p_2)] + \mathbb{E}[U_S(p|(\gamma_A, \alpha_B), p_1, p_2)]\right) \\
=&\max_{x,u,v} xuPR((\alpha_A, \gamma_B), u) + (1-x)vPR((\gamma_A, \alpha_B), v) \ s.t. \ xu + (1-x)v = p
\end{aligned}$$

Now, let $y = \frac{xu}{p}$, we can represent the persuasion ratio curve of (α_A, α_B) strategy to be the following:

$$\begin{aligned}
&PR((\alpha_A, \alpha_B), p) \\
=&\max_{p_1,p_2} \mathbb{E}[U_S(p|(\alpha_A, \alpha_B), p_1, p_2)]/p \\
=&\max_{x,u,v} \frac{xu}{p} PR((\alpha_A, \gamma_B), u) + \frac{(1-x)v}{p} PR((\gamma_A, \alpha_B), v) \ s.t. \ xu + (1-x)v = p \\
=&\max_{y,u,v} \{yPR((\alpha_A, \gamma_B), u) + (1-y)PR((\gamma_A, \alpha_B), v) \big| yu + (1-y)v = p \ y \in [0,1]\}
\end{aligned}$$

Using the same approach, we can represent the maximum expected utility of the set of strategy (α_A, β_B), (β_A, α_B), (β_A, β_B) in the form of a maximum convex combination of two persuasion ratio curves by using the corresponding two of four utility functions listed below:

$$\begin{aligned}
\mathbb{E}[U_S(p|(\alpha_A, \gamma_B), p_1, p_2)] &= pp_1q_{A1} + (1-p)p_2q_{A2}, \\
\mathbb{E}[U_S(p|(\gamma_A, \alpha_B), p_1, p_2)] &= p(1-p_1)q_{B1} + (1-p)(1-p_2))q_{B2}, \\
\mathbb{E}[U_S(p|(\beta_A, \gamma_B), p_1, p_2)] &= pp_1 + (1-p)p_2, \\
\mathbb{E}[U_S(p|(\gamma_A, \beta_B), p_1, p_2)] &= p(1-p_1) + (1-p)(1-p_2).
\end{aligned}$$

C.2.4 Proof of Lemma 16

Given persuasion ratio curves of all types of strategies, the sender, in order to maximize her expected utility, will choose the type of strategies with the highest persuasion ratio. (Because the persuasion ratio indicates a "normalized" utility with respect to the prior.) When we plot the persuasion ratio the sender chooses for every prior, the optimal persuasion ratio curve is the upper envelope of the different types of strategies' persuasion ratio curves. We want to explain why the persuasion ratio that the sender chooses to maximize her expected utility must correspond to an incentive-compatible strategy. According to Definition 24, every non-zero point on the persuasion ratio curve of a given type of strategy corresponds to at least one incentive-compatible strategy. Hence, we need to show that at every prior $p \in (0, 1)$, there must exist a type of strategy with a persuasion ratio > 0.

1. When $0 < p < \frac{q_{A2}}{q_{A1}+q_{A2}}$, the type of strategy (α_A, γ_B) has an incentive compatible strategies with persuasion ratio $2q_{A1}$.

2. When $\frac{q_{A2}}{q_{A1}+q_{A2}} \leq p < \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$, the type of strategy (α_A, β_B) has an inventive compatible strategy with persuasion ratio $\geq \min\{2q_{A1}, 1 + \frac{1-q_{B1}}{1-q_{B2}}\}$.

3. When $\frac{1-q_{B2}}{2-q_{B1}-q_{B2}} \leq p < 1$, the type of strategy (γ_A, β_B) has an inventive compatible strategy with persuasion ratio $\frac{1}{p}$.

Given the fact (1), (2), and (3), the value of the optimal persuasion ratio curve is non-zero for every $p \in (0, 1)$. Hence, the optimal signaling strategy is the strategy that reaches the frontier of the persuasion ratio curve.

C.2.5 Proof of Lemma 17

We prove this lemma by constructing an optimal signaling strategy in the pruned model, which has the same expected utility as the single-phase Bayesian persuasion problem.

Let's assume there is an N-phase model M for which the pruned model $Prun(M)$ satisfies the above two conditions stated in Lemma 17 and a prior $p < \frac{1}{2}$. We construct a signaling strategy that gives the sender the expected utility is equal to $2p$ (which is the same as the single-phase Bayesian persuasion.

First, for each non-trivial experiment E_X, design its preceding sender-designed experiment such that $\mathbb{P}(\theta_1|\omega_X) = \mathbb{P}(\theta_2|\omega_X) = 0$. In other words, we assign probability 0 to the experiment E_X. After the above process, if there exists some sender-designed experiments, design these experiments to be (arbitrary chosen) trivial experiment with $p_1 = p_2$.

Now, let $\mathbb{E}_{\mathbb{F}}$ be the set of sender-designed experiments such that for every $E_f \in \mathbb{E}_{\mathbb{F}}$, there is no sender-designed experiment between this experiment E_f and the experiment at root E_o. In other words, $\mathbb{E}_{\mathbb{F}}$ is the set of first-visited sender-designed experiments starting from the root. For every experiment $E_f = (p_1, p_2) \in \mathbb{E}_{\mathbb{F}}$ with succeeding experiments E_l, E_r (on the binary tree of experiments), let $(p_1, p_2) = (1, \frac{p}{1-p})$ if the succeeding experiment E_r is a non-trivial experiment, otherwise let $(p_1, p_2) = (0, 1 - \frac{p}{1-p})$. Hence, the posterior belief $\mathbb{P}(\theta_1|\omega)$ in this pruned model $\mathbb{M}$ is either 0 or $\frac{1}{2}$. Since the whole signal distribution must respect the prior, the sender will have an expected utility $2p$ under this construction, which is the same as the classical Bayesian persuasion in single-phase model.

C.2.6 Proof of Lemma 18

We present a construction to substitute a given non-binary experiment $E = \{q_{1,1}, ..., q_{1,n}; q_{2,1}, ..., q_{2,n}\}$ by a sequence of binary outcome experiments.

First, we construct a level $\lceil \log_2 n \rceil$ of binary-outcome experiments - $\{E'_{1,1}, E'_{2,1}, E'_{1,2}, ..., E'_{\lceil \log_2 n \rceil, 2^{\lceil \log_2 n \rceil}}\}$ - that we will use for replacement. Specifically, let

$E_{i,j}$ to be the experiment below:

$$\left(\frac{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k}}{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k} + \sum_{l=2^{i-1}\cdot(j-2)+1}^{2^{i-1}\cdot(j-1)} q_{1,l}}, \frac{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{2,k}}{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k} + \sum_{l=2^{i-1}\cdot(j-2)+1}^{2^{i-1}\cdot(j-1)} q_{2,l}}\right)$$

when i is even, and

$$\left(\frac{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k}}{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k} + \sum_{l=2^{i-1}\cdot j+1}^{2^{i-1}\cdot(j+1)} q_{1,l}}, \frac{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{2,k}}{\sum_{k=2^{i-1}\cdot(j-1)+1}^{2^{i-1}\cdot j} q_{1,k} + \sum_{l=2^{i-1}\cdot j+1}^{2^{i-1}\cdot(j+1)} q_{2,l}}\right)$$

when i is odd.

Note that the denominator sums to 1 in $E_{1,1}$, and the nominator cancels the denominator in the next level. Hence, when we multiple terms from the first level to the last level, the k^{th} outcome should yield $(q_{k,1}, q_{k,2})$.

C.2.7 Proof of Lemma 19

Given the receiver's utility function defined in this work, i.e., $U_R(a_i, \theta_j) = \mathbf{1}_{\{i=j\}}$, the receiver's utility is maximized when the sender commits to a truth-telling signaling strategy. Given the phase-II experiments $E_A = (1, 0)$ and $E_B = (1, 0)$ stated in the lemma, the posterior belief in each possible outcome $\mathbb{P}(\theta_1|\omega_{AP})$, $\mathbb{P}(\theta_1|\omega_{AF})$, $\mathbb{P}(\theta_1|\omega_{BP})$, $\mathbb{P}(\theta_1|\omega_{BF}) \in \{0, 1\}$ for any $(p_1, p_2) \in [0, 1]^2$ where the sender chooses phase-I experiment. In other words, the receiver knows the true state for every possible realized outcome $\omega \in \{\omega_{AP}, \omega_{AF}, \omega_{BP}, \omega_{BF}\}$. Hence, for any prior p, the receiver will have an utility $\mathbb{E}[U_R(\omega, p)] = 1$, which is the maximum utility the receiver can have under the given utility function.

C.2.8 Proof of Lemma 20

First, let's consider the case when $q_{Y1} \leq \frac{2-q_{Y2}}{3-2q_{Y2}}$.

1. Suppose a given experiment E_A has $\max\{2q_{A1}-1, \frac{1-q_{A2}}{1-q_{A1}}\} \leq \frac{1-q_{Y2}}{1-q_{Y1}}$, the optimal signaling strategy of (E_A, E_Y) when $\mathbb{P}(a_1|\Pi) < 1$ is (γ_A, β_Y). Hence, if $\max\{2q_{X1}-1, \frac{1-q_{X1}}{1-q_{X2}}\} < \frac{1-q_{Y1}}{1-q_{Y2}}$, the persuasion ratio curve of any type of IC strategy (c_A, d_X) with $c, d \in \{\alpha, \beta, \gamma\}$ has a persuasion ratio $\leq 1 + \frac{1-q_{Y2}}{1-q_{Y1}}$, for every possible prior. Hence, the lemma holds in this case.

2. Suppose a given experiment E_A has $\max\{2q_{A1}-1, \frac{1-q_{A2}}{1-q_{A1}}\} > \frac{1-q_{Y2}}{1-q_{Y1}}$ and $2q_{A1} > 1 + \frac{1-q_{A2}}{1-q_{A1}}$, the optimal signaling strategy of $(E_A, E_B), B \in \{X, Y\}$ when $\mathbb{P}(a_1|\Pi) < 1$ is (α_A, γ_B), (α_A, β_B), or (β_A, γ_B) depending on the prior region. Given $\max\{2q_{X1}-1, \frac{1-q_{X1}}{1-q_{X2}}\} < \frac{1-q_{Y1}}{1-q_{Y2}}$, the persuasion ratio of (α_A, β_Y) is higher than (α_A, β_X) for every prior p such that (α_A, β_Y), (α_A, β_X) have IC strategies. Hence, the lemma holds in this case.

3. Suppose a given experiment E_A has $\max\{2q_{A1}-1, \frac{1-q_{A2}}{1-q_{A1}}\} > \frac{1-q_{Y2}}{1-q_{Y1}}$ and $2q_{A1} < 1 + \frac{1-q_{A2}}{1-q_{A1}}$, the optimal signaling strategy of $(E_A, E_B), B \in \{X, Y\}$ when $\mathbb{P}(a_1|\Pi) < 1$ is always a (β_A, γ_B) type strategy. Hence, the statement of lemma holds in this case.

After considering all scenarios under $q_{Y1} \leq \frac{2-q_{Y2}}{3-2q_{Y2}}$, we know that if $\max\{2q_{X1} - 1, \frac{1-q_{X1}}{1-q_{X2}}\} < \frac{1-q_{Y1}}{1-q_{Y2}}$, then E_X is inferior to E_Y.

Now, let's consider the case when $q_{Y1} > \frac{2-q_{Y2}}{3-2q_{Y2}}$. E_X is inferior to E_Y requires for every prior p, the inequality of optimal persuasion ratio $PR^*_{E_A,E_x}(p) \leq PR^*_{E_A,E_x}(p)$ holds. Since the experiment E_A can be arbitrary, E_X being inferior to E_Y requires the persuasion ratio curve of different type of strategies under (E_A, E_x) always being below the persuasion ratio curve of different type of strategies under (E_A, E_Y). Using the partial strategy analysis in Section 3.3.3, the above condition requires that every class of partial strategies c_B, $c \in \{\alpha, \beta\}$, $B \in \{X, Y\}$, using experiment E_X is always inefficient when compared to using experiment E_X. That is to say, $2q_{X_1} < 2q_{Y1}$ and $1 + \frac{1-q_{X1}}{1-q_{X2}} < 1 + \frac{1-q_{Y_1}}{1-q_{Y2}}$ according to the persuasion potential of E_X and E_Y. Hence, when $q_{Y1} > \frac{2-q_{Y2}}{3-2q_{Y2}}$, the condition derived from persuasion potential guaranteeing that E_X is inferior to E_Y satisfies the second condition stated in the lemma.

C.2.9 Proof of Corollary 4

When the prior $p \geq \min\{\frac{1-q_{A2}}{2-q_{A1}-q_{A2}}, \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}\}$, the sender has an optimal signaling strategy to persuade the receiver to take action a_1 with probability one by $(p_A, p_B) = (1, 1)$ when $\frac{1-q_{A2}}{2-q_{A1}-q_{A2}} \leq \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$ or by $(p_A, p_B) = (0, 0)$ when $\frac{1-q_{A2}}{2-q_{A1}-q_{A2}} > \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$. Hence, to prove this lemma, we only need to consider the prior region $p \in [0, \min\{\frac{1-q_{A2}}{2-q_{A1}-q_{A2}}, \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}\})$.

First, let's consider $p \in [0, \min\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}]$. When prior falls in this region, four type of strategies: (α_A, α_B), (α_A, β_B), (β_A, α_B), (β_A, β_B) cannot satisfy the IC requirements[4]. Hence, the optimal signaling strategy can only come from (α_A, γ_B), (β_A, γ_B), (γ_A, α_B), (γ_A, β_B) types of strategies. Given the result of Lemma 14, the IC requirement of α or β partial strategy used in the optimal signaling strategy must be tight. This implies that the interim belief of α_X strategy must be $\frac{q_{X2}}{q_{X1}+q_{X2}}$ and the interim belief of β_X strategy $\frac{1-q_{X2}}{2-q_{X1}-q_{X2}}$ (otherwise the IC requirement is not tight). Under these interim belief, the α_X, β_X partial strategy is used in the most efficient manner as described in Section 3.3.3 and the corollary holds here.

Second, let's consider $p \in [\min\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}, \max\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}]$. When the prior falls in this region, the optimal signaling strategy can come from one of the 6 types of strategies[5]: (α_A, γ_B), (β_A, γ_B), (γ_A, β_B), (α_A, α_B), (α_A, β_B), (β_A, α_B). Recall the analysis in the proof of Lemma 14, when a strategy in the type of (α_A, β_B), (β_A, α_B) or (α_A, α_B) maximizes the sender's expected utility in this prior region, both IC requirements are tight. Besides, when a strategy in the type of (β_A, γ_B) or (γ_A, β_B) maximizes the sender's expected utility in this prior region, the IC requirement of

[4]See Proof of Lemma 14 for details.

[5](γ_A, α_B) cannot be an optimal signaling strategy in this region because $q_{A1} \geq q_{B1}$ is assumed in our framework.

β_X is tight. Hence, the only possibility for this corollary to fail is when the optimal signaling strategy is a (α_A, γ_B) type strategy. Now, we prove that the optimal signaling strategy is never an (α_A, γ_B) type strategy in this prior region by comparing the sender's maximum expected utility between (α_A, γ_B) and (α_A, β_B) type strategies.

Under (α_A, γ_B) type strategy, the sender's expected utility is maximized when $(p_1, p_2 = 1)$. Hence, $U_S((\alpha_A, \gamma_B), p) = pq_{A1} + (1-p)q_{A2}$ in this prior region.

Under (α_A, β_B) type strategy, the sender's expected utility is maximized when $\mathbb{P}(\theta_1|E_A) = \frac{q_{A2}}{q_{A1}+q_{A2}}$ and $\mathbb{P}(\theta_1|E_B) = \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}$.

Hence, $U_S((\alpha_A, \beta_B), p) = 2\frac{\mathbb{P}(\theta_1|E_B)-p}{\mathbb{P}(\theta_1|E_B)-\mathbb{P}(\theta_1|E_A)}\mathbb{P}(\theta_1|E_A)q_{A1} + (1 - \frac{\mathbb{P}(\theta_1|E_B)-p}{\mathbb{P}(\theta_1|E_B)-\mathbb{P}(\theta_1|E_A)})$.

When we look at the sender's maximum expected utility under these two types of strategies, we can notice that $U_S((\alpha_A, \gamma_B), p) = U_S((\alpha_A, \beta_B), p)$ at $p = \frac{q_{A2}}{q_{A1}+q_{A2}}$, which is the minimum prior in this region. When p increases, the marginal utility gain at (α_A, γ_B) type strategy is $q_{A1} - q_{A2}$, and the marginal utility gain at (α_A, β_B) type strategy is

$$\frac{-2}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)}\mathbb{P}(\theta_1|E_A)q_{A1} + \frac{1}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} = \frac{1 - 2q_{A1}\mathbb{P}(\theta_1|E_A)}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)})$$

Thus, the utility difference $U_S((\alpha_A, \beta_B), p) - U_S((\alpha_A, \gamma_B), p)$ in this prior region is

$$\begin{aligned}
& p \times \Big(\frac{1 - 2q_{A1}\mathbb{P}(\theta_1|E_A)}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} - (q_{A1} - q_{A2})\Big) \\
=& \frac{p}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} \times \Big(1 - 2q_{A1}\mathbb{P}(\theta_1|E_A) - (\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A))(q_{A1} - q_{A2})\Big) \\
=& \frac{p}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} \times \Big(1 - \mathbb{P}(\theta_1|E_A)(q_{A1} + q_{A2}) - \mathbb{P}(\theta_1|E_B)(q_{A1} - q_{A2})\Big) \\
=& \frac{p}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} \times \Big(1 - q_{A2} - \mathbb{P}(\theta_1|E_B)(q_{A1} - q_{A2})\Big) \\
=& \frac{p}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} \times \Big(1 - \frac{(1 - q_{B2})q_{A1} + (1 - q_{B1})q_{A2}}{2 - q_{B1} - q_{B2}}\Big) \\
>& \frac{p}{\mathbb{P}(\theta_1|E_B) - \mathbb{P}(\theta_1|E_A)} \times \Big(1 - q_{A1}\Big) \geq 0.
\end{aligned}$$

Given the calculation above, we know that (α_A, γ_B) type strategy will never be an optimal signaling strategy in this prior region. Hence, the corollary holds in this prior region.

Finally, we consider $p \in [\max\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}, \min\{\frac{1-q_{A2}}{2-q_{A1}-q_{A2}}, \frac{1-q_{B2}}{2-q_{B1}-q_{B2}}\}]$. When the prior falls in this region, all strategies except the (β_A, β_B) strategies can satisfy the IC requirements. Given the tightness analysis in the proof of Lemma 14, the corollary will fail if the optimal signaling strategy comes from type (α_A, γ_B), (γ_A, α_B), or (α_A, α_B). To rule out (α_A, γ_B), (γ_A, α_B), we use exactly the same calculation as we did in the prior region $p \in [\min\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}, \max\{\frac{q_{A2}}{q_{A1}+q_{A2}}, \frac{q_{B2}}{q_{B1}+q_{B2}}\}]$. Hence, the only interesting case to study is whether (α_A, α_B) type strategy can be an optimal signaling strategy in this prior region.

Suppose (α_A, α_B) type strategy is an optimal signaling strategy in this prior region, at most one IC requirement, which is α_A, can be tight. When the IC requirement of α_A is not tight, the strategy is not an optimal signaling strategy since one the sender can increase her expected payoff by increasing $\mathbb{P}(\omega_A)$ and reduce $\mathbb{P}(\omega_A)$ respecting the prior. When the IC requirement of α_A is tight and the current (α_A, β_B)-type strategy, called $W_{\alpha,\beta}$ is not an incentive compatible strategy for the receiver, the following modification can further increase the sender's expected utility:

- Reduce $\mathbb{P}(\theta_1, \omega_B)$ by ϵ but increase $\mathbb{P}(\theta_1, \omega_A)$ by ϵ;
- Given the above ϵ, reduce $\mathbb{P}(\theta_2, \omega_B)$ by $\frac{1-q_{A1}}{1-q_{A2}}\epsilon$ but increase $\mathbb{P}(\theta_2, \omega_A)$ by $\frac{1-q_{A1}}{1-q_{A2}}\epsilon$.

When $W_{\alpha,\beta}$ is not an incentive compatible strategy for the receiver, there must exists an $\epsilon > 0$ such that the above modification constructs a new incentive compatible strategy, called $W^*_{alpha,\beta}$. When ϵ is picked such that $W^*_{\alpha,\beta}$ is an incentive compatible strategy for the receiver, the strategy type (α_A, β_B) (with the strategy $W^*_{\alpha,\beta}$ constructed above) gives a higher expected utility than the strategy type (α_A, α_B) under the same parameter of (p_1, p_2). Hence, (α_A, α_B) is never an optimal signaling strategy (in this prior region) and the corollary holds.

C.2.10 Proof of Claim 5

We prove each assumptions by a replacement of "equivalant" experiment. First, suppose we have a phase-II experiment $E_X,\ X \in \{A, B\}$ which violates our first assumption with $q_{X1} \leq q_{X2}$. We can construct an equivalent experiment E'_X to replace the experiment E_X. Given $E_X = (q_{X1}, q_{X2})$ with $q_{X1} \leq q_{X2}$, we consider another experiment $E'_X = (q_{X'1}, q_{X'2})$ such that $q_{X'1} = q_{X2}$ and $q_{X'2} = q_{X1}$. This experiment E'_X can be implemented in two steps, the first step flips the index of state, and the second step operates the experiment same as E_X. For every interim belief $\tilde{p}_X$, we claim this replacement of experiment will not affect the sender's design room of optimal signaling strategies. Given a posterior belief $\mathbb{P}(\theta_1|\tilde{p}_X, \omega_{XP})$ (or $\mathbb{P}(\theta_1|\tilde{p}_X, \omega_{XF})$) is generated when the experiment E_X passes (or fails), the same posterior belief can be generated when the experiment E'_X fails (or passes). Mathematically,

$$
\begin{aligned}
\mathbb{P}(\theta_1|\tilde{p}_X, \omega_{XP}) &= \frac{\tilde{p}_X q_{X1}}{\tilde{p}_X q_{X1} + (1-\tilde{p}_X) q_{X2}} = \frac{\tilde{p}_X q_{X'2}}{\tilde{p}_X q_{X'2} + (1-\tilde{p}_X) q_{X'1}} \\
&= 1 - \frac{(1-\tilde{p}_X) q_{X'1}}{\tilde{p}_X q_{X'2} + (1-\tilde{p}_X) q_{X'1}} = \mathbb{P}(\theta_1|\tilde{p}_X, \omega_{X'F}) \\
\mathbb{P}(\theta_1|\tilde{p}_X, \omega_{XF}) &= \frac{\tilde{p}_X (1-q_{X1})}{\tilde{p}_X (1-q_{X1}) + (1-\tilde{p}_X)(1-q_{X2})} \\
&= \frac{\tilde{p}_X (1-q_{X'2})}{\tilde{p}_X (1-q_{X'2}) + (1-\tilde{p}_X)(1-q_{X'1})} \\
&= 1 - \frac{(1-\tilde{p}_X)(1-q_{X'1})}{\tilde{p}_X (1-q_{X'2}) + (1-\tilde{p}_X)(1-q_{X'1})} = \mathbb{P}(\theta_1|\tilde{p}_X, \omega_{X'P})
\end{aligned}
$$

Hence, the first assumption in Claim 5 can be made without loss of generality.

Suppose we have phase-II experiments E_A, E_B such that $q_{A1} \leq q_{B1}$, then we can replace experiments E_A, E_B by experiments $E_{A'}, E_{B'}$ such that $q_{A'1} = q_{B1}$ and $q_{B'1} = q_{A1}$ to satisfy the assumption 2. Similarly, we can prove this operation (of replacement) will not affect the sender's design room by proving the posterior belief $\mathbb{P}(\theta_1|\omega_{AY}) =$

$\mathbb{P}(\theta_1|\omega_{B'Y}),\ Y \in \{P, F\}$ and the posterior belief $\mathbb{P}(\theta_1|\omega_{BY}) = \mathbb{P}(\theta_1|\omega_{A'Y}),\ Y \in \{P, F\}$. For the conciseness, we only present the math details for $Y = P$ to avoid redundant equations replacing $q_{A1}, q_{B1}, q_{A'1}, q_{B'1}$ with $(1-q_{A1}), (1-q_{B1}), (1-q_{A'1}), (1-q_{b'1})$ respectively.

$$
\begin{aligned}
\mathbb{P}(\theta_1|\omega_{AP}) &= \frac{pq_{A1}}{pq_{A1} + (1-p)q_{A2}} = \frac{pq_{B'1}}{pq_{B'1} + (1-p)q_{B'2}} = \mathbb{P}(\theta_1|\omega_{B'P}) \\
\mathbb{P}(\theta_1|\omega_{BP}) &= \frac{pq_{B1}}{pq_{B1} + (1-p)q_{B2}} = \frac{pq_{A'1}}{pq_{A'1} + (1-p)q_{A'2}} = \mathbb{P}(\theta_1|\omega_{A'P})
\end{aligned}
$$

Hence, the second assumption in Claim 5 can be made without loss of generality.

C.2.11 Proof of Claim 6

Given that $q_{X1} \geq q_{X2}$, let's calculate the posterior of outcome ω_{XP} and ω_{XF}. Suppose an interior beleif of $\mathbb{P}(\theta_1|\omega_X) = \tilde{p}$, the posterior belief $\mathbb{P}(\theta_1|\omega_{XP}) = \frac{\tilde{p}q_{X1}}{\tilde{p}q_{X1}+(1-\tilde{p})q_{X2}}$ and the posterior belief $\mathbb{P}(\theta_1|\omega_{XF}) = \frac{\tilde{p}(1-q_{X1})}{\tilde{p}(1-q_{X1})+(1-\tilde{p})(1-q_{X2})}$.

Given that $q_{X1} \geq q_{X2}$,

$$
\begin{aligned}
\mathbb{P}(\theta_1|\omega_{XP}) &= \frac{\tilde{p}q_{X1}}{\tilde{p}q_{X1} + (1-\tilde{p})q_{X2}} \\
&> \frac{\tilde{p}q_{X1}}{\tilde{p}q_{X1} + (1-\tilde{p})q_{X1}} = \frac{\tilde{p}}{\tilde{p} + (1-\tilde{p})} = \frac{\tilde{p}(1-q_{X1})}{\tilde{p}(1-q_{X1}) + (1-\tilde{p})(1-q_{X1})} \\
&> \frac{\tilde{p}(1-q_{X1})}{\tilde{p}(1-q_{X1}) + (1-\tilde{p})(1-q_{X2})} = \mathbb{P}(\theta_1|\omega_{XF}).
\end{aligned}
$$

Hence, taking action a_1 when E_X fails but taking action a_2 when E_X passes is not an incentive compatible strategy of the receiver for any interim belief $\mathbb{P}(\theta_1|E_X) \in [0, 1]$.

C.2.12 Proof of Claim 7

Given $PR^*_{K+1,2j-1}(p)$ and $PR^*_{K+1,2j}(p)$, We show that there exists an incentive compatible strategy $S_{K,j,+}$ starting from the experiment $E_{K,j}$ satisfies the following equation:

$$PR_{K,j}(S_{K,j,+}, p) = \max_{x,u,v \in [0,1], xu+(1-x)v=p} xPR^*_{K+1,2j-1}(u) + (1-x)PR^*_{K+1,2j}(v) \quad \text{(C.1)}$$

We prove it by mathematical induction. When $K = N - 2$, Kemma 16 tells that every point on an optimal persuasion ratio curve at $N - 1$ corresponds to at least one incentive compatible strategy. Hence, a strategy $S_{N-2,j,+}$ which uses the experiment $E_{N-2,j} = \left(\frac{xu}{p}, \frac{x(1-u)}{1-p}\right)$ at phase K given the optimal IC signaling strategies at $E_{N-1,2j-1}$, $E_{N-1,2j}$ at phase $N - 1$ satisfying equation (C.1). The reason is that the role of experiment $E_{N-2,i}$ is to determining what incentive compatible strategies will be used at phase $N-1$ by influencing the interim belief at phase $N-1$. As long as the strategies taken at phase $N - 1$ satisfies the IC requirements, $S_{N-1,j,+}$ is also an incentive compatible strategy.

Suppose the statement is true at phase $K + 1$ and the corresponding optimal signaling strategies are $S_{K+1,2j-1,+}$ and $S_{K+1,2j,+}$, the strategy $S_{K,j,+}$ with $E_{K,j} = \left(\frac{xu}{p}, \frac{x(1-u)}{1-p}\right)$ such that the parameters x, u, v correspond to a maximized the equation (C.1) is an incentive compatible strategy. This is because for every other strategy $S'_{K,j,+}$ not using both $S_{K+1,2j-1,+}$ and $S_{K+1,2j,+}$, suppose $S'_{K+1,2j-1,+}$ and $S'_{K+1,2j,+}$, the sender can replace $S'_{K+1,2j-1,+}$ by $S_{K+1,2j-1,+}$ and $S'_{K+1,2j-1,+}$ by $S_{K+1,2j,+}$ to construct another strategy no worse than the $S'_{K,j,+}$. Moreover, the strategy $S_{K,j,+}$ uses the parameters that maximize the equation (C.1) guarantees that there is no other combination of $S_{K+1,2j-1,+}$ and $S_{K+1,2j,+}$ gives the sender a higher expected utility. Hence, $S_{K,j,+}$ is an incentive compatible strategy in phase K if both $S_{K+1,2j-1,+}$ and $S_{K+1,2j,+}$ are incentive compatible strategies. Since it is true in every strategy $S_{N-2,j,+}$ at phase $K = N - 2$, the the mathematical induction guarantees that $S_{K,j,+}$ is an incentive compatible strategy for every $K \leq N - 2$ and $j \in [1, 2^K]$ and the proof is complete here.

APPENDIX D

Appendices of Chapter IV

D.1 Proofs of Chapter IV

D.1.1 Proof of Proposition 4

Proof. According to condition 2, we have a signal $\hat{\omega}_1$ which tells S_1's private information is I_1, i.e., $\mathbb{P}(I_1|\hat{\omega_1}) = 1$. First, we prove S_1 does not prefer sending the signal $\hat{\omega}_1$. When S_1 sends signal $\hat{\omega}_1$, S_2 knows S_1's private information is I_1. Let $\Gamma_2^*(I_1)$ denote S_2's optimal signaling strategy while receiving $\hat{\omega}_1$, inequality (4.15) states S_1's utility under action $a_{\hat{\theta}}$ is higher than her expected utility conditional on her signal $\hat{\omega}_1$. Thus, S_1 prefers a mixture of her information set I_1 and the information set which contains the collaborative state $\hat{\theta}$, called $\hat{I}_1$.

We notice that if S_2's optimal signaling strategy while receiving $\hat{\omega}_1$ is not unique, the conditional expectation $\mathbb{E}_{p,\Gamma_2^*(I_1)}[U_{S_1}(\theta, a^*(I_1, \omega_2))|I_1]$ may not be well-defined since different S_2's tie-breaking strategies may vary S_1's expected utility. To deal with the non-uniqueness of $\Gamma_2^*(I_1)$, we assume the polar opposite preference ordering on the receiver's possible optimal action set given S_1's private information I_1, called $\Phi(I_1)$ in condition 3. Therefore, when S_2 is indifferent between two different actions a_m and a_n

where $a_m, a_n \in \Phi(I_1)$, S_1 is indifferent between a_m and a_n too. Thus, the conditional expectation $\mathbb{E}_{p,\Gamma_2^*(I_1)}[U_{S_1}(\theta, a^*(I_1, \omega_2))|I_1]$ in inequality (4.15) is well-defined, and S_1 prefers a mixture of I_1 and $\hat{I}_1$ when the inequality (4.15) holds.

Next, we prove S_2 prefers a separation of I_1 and $\hat{I}_1$. Given assumption 4 in the problem formulation, where the true state can be revealed under both senders' truth-telling strategies, S_2 knows the true state while observing the signal realization $\hat{\omega}_1$ and her private information I_2, and her optimal signaling strategy under $\hat{\omega}_1$ can be calculated directly via the concavification approach — this is equivalent to solving classic Bayesian persuasion problem in *Kamenica and Gentzkow* (2011) with finite states. According to the inequality 4.14 stated in condition 1, S_2 enjoys a higher expected utility while observing $\hat{\omega}_1$ than observing the signal telling her S_1's private information is $\hat{I}_1$. Thus, S_2 prefers a separation of I_1 and $\hat{I}_1$.

When condition 3 in Proposition 4 holds, i.e., preference orders of senders on $\Phi(I_1)$ are polar opposite, S_2's expected-utility maximization strategy (via the concavification approach) will directly minimize S_1 expected utility when $\hat{\omega_1}$ is realized. Therefore, from S_1's objective, S_1 should minimize the probability $\mathbb{P}(\hat{\omega_1})$ to reduce the probability where S_2 observes $\hat{\omega_1}$. Hence, S_1 will increase $\mathbb{P}(\omega_1|I_1)$ for every ω_1 such that $\mathbb{E}_p[U_{S_1}|\omega_1] > \mathbb{E}_p[U_{S_1}|\hat{\omega_1}]$. This includes the signal which suggests the receiver to take the collaborative action $a_{\hat{\theta}}$. Hence, when S_1 commits first, she will maximize the probability of suggesting the collaborative state $\hat{\theta}$ via mixing I_1 with $\hat{I}_1$. Since $\hat{\theta}$ is a collaborative state, S_2 will collaborate with S_1 while observing S_1's signal which suggests $a_{\hat{\theta}}$.

However, when sender S_2 commits first, the inequality (4.14) stated in condition 1, i.e., $U_{S_2}(\hat{\theta}, a_{\hat{\theta}}) < \mathbb{E}_{p,\Gamma_2^*(I_1)}[U_{S_2}(\theta, a^*(I_1, \omega_2))|I_1]$, guarantees that S_2 prefers the realized signal $\hat{\omega_1}$ instead of the realized signal which suggests her to collaborate toward $\hat{\theta}$ via a mixture of I_1 and $\hat{I}_1$. Therefore, when sender S_2 commits first, she will commit to a signaling strategy which minimizes the probability of I_1 mixed with $\hat{I}_1$ (in order to increase the total probability of observing $\hat{\omega_1}$). Since condition 2 guarantees that $\mathbb{P}(\hat{\omega}_1) > 0$ for every optimal commitments, we will never fall into the corner case where S_1 and S_2 have an opposite preference on the mixture of I_1 and $\hat{I}_1$, but one of the sender's strategy of minimizing/maximizing the mixture of I_1 and $\hat{I}_1$ cannot be achieved owing to the low value of $\mathbb{P}(I_1)$ under the given prior. Hence, according to the opposite strategies of S_1 and S_2 on the mixture of I_1 and $\hat{I}_1$, the commitment order matters. $\square$

D.1.2 Proof of Claim 8

Proof. We prove this claim by contradiction. First, assume that both conditions are violated and the signal $\hat{\omega}_1$ revealing I_1 is the signal realization which gives S_1 the highest expected utility under an equilibrium derived by solving the objectives stated in Section 4.2.4, where each sender's commitment is an optimal one under the given commitment order. Then, let signal $\bar{\omega}_1$ represents the signal that S_1 sends when the true state is $\hat{\theta}$, the violation of condition 2 guarantees that the following inequality holds

$$\mathbb{E}_p[U_{S_1}(a^*_{\Gamma_1^*,\Gamma_2^*}(\hat{\omega}_1,\omega_2))] > \mathbb{E}_p[U_{S_1}(a^*_{\Gamma_1^*,\Gamma_2^*}(\bar{\omega}_1,\omega_2))],$$

where $a^*_{\Gamma_1^*,\Gamma_2^*}(\omega_1,\omega_2)$ is the best response of the receiver under a realized pair of signal ω_1,ω_2 and both senders' commitments Γ_1^* and Γ_2^*. Since the signal $\hat{\omega}_1$ reveals the private signal I_1, we can construct another commitment Γ_1' below, where $\epsilon \ll 1$:

1. $\mathbb{P}(\hat{\omega}_1') = \mathbb{P}(\hat{\omega}_1) + \epsilon\mathbb{P}(\bar{\omega}_1)$ and $\mathbb{P}(\theta|\hat{\omega}_1') = \frac{1}{1+\epsilon}\mathbb{P}(\theta|\hat{\omega}_1) + \frac{\epsilon}{1+\epsilon}\mathbb{P}(\theta|\bar{\omega}_1)$ for all $\theta \in \Theta$;
2. $\mathbb{P}(\bar{\omega}_1') = (1-\epsilon)\mathbb{P}(\bar{\omega}_1)$, and the interim beliefs of $\bar{\omega}_1'$ and $\bar{\omega}_1$ are the same, i.e., $\mathbb{P}_{p,\Gamma_1'}\big(\theta|\bar{\omega}_1'\big) = \mathbb{P}_{p,\Gamma_1}\big(\theta|\bar{\omega}_1\big)$ for all $\theta \in \Theta$;
3. For every ω_1 under the commitment Γ_1' such that $\omega_1 \neq \bar{\omega}_1$ and $\omega_1 \neq \hat{\omega}_1$, $\mathbb{P}_{p,\Gamma_1'}(\omega_1) = \mathbb{P}_{p,\Gamma_1^*}(\omega_1)$ and $\mathbb{P}_{p,\Gamma_1'}\big(\theta|\omega_1\big) = \mathbb{P}_{p,\Gamma_1^*}\big(\theta|\omega_1\big)$ for all $\theta \in \Theta$;

As assumed earlier, $\hat{\omega}_1$ is the signal which gives S_1 the highest expected utility with $\mathbb{P}(I_1|\hat{\omega}_1) = 1$, and $\bar{\omega}_1$ is the signal S_1 sends to collaborate with S_2 on persuading the receiver to take $a_{\hat{\theta}}$. Given the inequality (4.15), S_2's optimal signaling strategy belongs to one of the following two cases:

Case 1: S_2 cannot elicit the collaborative state $\hat{\theta}$ under the signal $\hat{\omega}_1'$. In this case, because $\epsilon \ll 1$, the optimal commitment of sender S_2 stays unchanged between Γ_1^* and Γ_1' since S_2 cannot successfully persuade the receiver to take $a_{\hat{\theta}}$ and she has already chosen her optimal partition of I_1 in her commitment. Therefore, when $\epsilon \ll 1$, changing the commitment from Γ_1^* to Γ_1' in this case will not affect S_2's optimal signaling strategy. However, the revised commitment Γ_1' has a higher total probability of $\hat{\omega}_1'$ comparing to the signal $\hat{\omega}_1$ under Γ_1^*. Given the inequality $\mathbb{E}_p[U_{S_1}(a^*_{\Gamma_1^*,\Gamma_2^*}(\hat{\omega}_1, \omega_2))] > \mathbb{E}_p[U_{S_1}(a^*_{\Gamma_1^*,\Gamma_2^*}(\bar{\omega}_1, \omega_2))]$ and the knowledge that both the optimal commitments of S_2 and the receiver's best response stay unchanged, S_1 will enjoy a higher expected utility under the the commitment Γ_1', which contradicts the assumption that Γ_1^* is the optimal commitment of S_1. Hence, one of conditions in Claim 8 must hold in this case.

Case 2: S_2 can elicit the collaborative state $\hat{\theta}$ under the signal $\hat{\omega}_1'$. This case occurs when S_2's private information set $I_2 \ni \hat{\theta}$ contains no state that belongs to S_1's information set I_1, i.e., $\theta_i \notin I_2, \forall \theta_i \in I_1, \hat{\theta} \in I_2$. Therefore, even under the signal $\hat{\omega}_1'$ mixes $\hat{\omega}_1$ with the signal containing the collaborate states, S_2 can still persuade the receiver to take $a_{\hat{\theta}}$ with some probability (according to the prior). However, after maximizing the signal she used to persuade the action $a_{\hat{\theta}}$, her optimal signaling strategy conditional on her not sending signal persuading $a_{\hat{\theta}}$ is the same as her optimal signaling strategy while receiving $\hat{\omega}_1$. Therefore, S_2's optimal signaling changes by a small amount (because $\epsilon \ll 1$). Now, since $\hat{\theta}$ is a collaborative state, S_2's adjustment on her optimal signaling strategy from $\hat{\omega}_1$ under Γ_1^* to $\hat{\omega}_1'$ under Γ_1' not only benefits her but also benefits S_1. Moreover, after maximizing the probability of the signal persuading $a_{\hat{\theta}}$, the S_2's optimal signaling strategy stays unchanged (reduced to the case 1 scenario). Hence, S_1 enjoys a higher expected utility under the the commitment Γ_1' in this case (because S_2's optimal signaling strategy gives S_1 a higher utility than the utility in case 1, and S_1 enjoys a higher expected utility in case 1).

In both cases, S_1's expected utility increases. This violates the assumption that S_1 has the highest utility in $\hat{\omega}_1$ when condition 2 is violated. Hence, one of the conditions in Claim 8 must hold.

□

D.1.3 Proof of Theorem 2

First, let us review the statements in condition 1 for the readability of the proof.

Condition 1: There exist two parameters $\alpha > \beta > 0$ and a signaling strategy $\hat{\Gamma}_2$ satisfying the following conditions:

(a) $G(p, I_1^x + \alpha I_1^y) = G(p, I_1^x + \beta I_1^y)$,

(b) Let $\omega_1^\alpha, \omega_1^\beta$ be two mock signals of S_1 such that $\mathbb{P}(I_1^x|\omega_1^\alpha) = \frac{1}{1+\alpha}$, $\mathbb{P}(I_1^y|\omega_1^\alpha) = \frac{\alpha}{1+\alpha}$, $\mathbb{P}(I_1^x|\omega_1^\beta) = \frac{1}{1+\beta}$, and $\mathbb{P}(I_1^y|\omega_1^\beta) = \frac{\beta}{1+\beta}$, then

$$\hat{\Gamma}_2(p, \omega_1^\beta) \in G(p, I_1^x + \beta I_1^y) \text{ and } \hat{\Gamma}_2(p, \omega_1^\alpha) \notin G(p, I_1^x + \alpha I_1^y),$$

(c) $\mathbb{E}_p[U_{S_1}|\omega_1^\alpha, \hat{\Gamma}_2] < \frac{\mathbb{P}(I_1^x)+\beta\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)}\mathbb{E}_p[U_{S_1}|\omega_1^\beta, \hat{\Gamma}_2] + \frac{(\alpha-\beta)\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)}\mathbb{E}_p[U_{S_1}|I_1^y, \hat{\Gamma}_2]$,

(d) $\mathbb{E}_p[U_{S_2}|\omega_1^\beta, G(p, I_1^x + \beta I_1^y)] < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]$.

Given (a), if S_1 commits first and the two mock signals $\omega_1^\alpha, \omega_1^\beta$ described in (b) are used in S_1's signaling strategy, $G(p, I_1^x + \alpha I_1^y) = G(p, I_1^x + \beta I_1^y)$ states that the set of S_2's optimal signaling strategies under ω_1^α is the same as the set of S_2's optimal signaling strategies under ω_1^β. In other words, ω_1^α and ω_1^β are two different mixing schemes using S_1's information sets I_1^x and I_1^y that elicit the same response from S_2. According to the statement of condition 2, i.e., $\min_{g \in G(p, I_1^x)} \mathbb{E}_{p,g}[U_{S_1}|I_1^x] > \max_{g' \in G(p, I_1^y)} \mathbb{E}_{p,g'}[U_{S_1}|I_1^y]$, when S_2's action and the receiver's action both stay unchanged, S_1 prefers signal ω_1^α over signal ω_1^β. The logic is the following: First, suppose the receiver takes action $a \in I_1^x$, sending signal ω_1^α mixes a higher portion of I_1^y with I_1^x than sending signal ω_1^β. This implies that the receiver takes action $a \in I_1^x$ with a higher total probability. Based on the inequality $\min_{g \in G(p, I_1^x)} \mathbb{E}_{p,g}[U_{S_1}|I_1^x] > \max_{g' \in G(p, I_1^y)} \mathbb{E}_{p,g'}[U_{S_1}|I_1^y]$, increasing the total probability of a set of actions $a \in I_1^x$ by reducing the same amount of probability on a set of actions $a' \in I_1^y$ benefits S_1; second, suppose the receiver takes action $a' \in I_1^y$ under both signal ω_1^α and signal ω_1^β, S_1 still prefers using the signal ω_1^α when condition 2 holds, since under this circumstance signal ω_1^α uses less probability mass of states in I_1^x than signal ω_1^β. In short, combining the statements in condition 1.(a) and condition 2, we know S_1 prefers sending signal ω_1^α over signal ω_1^β. Furthermore, when we consider the statements in condition 1.(c) and 2, the inequality in 1.(c) implies S_1 can persuade the receiver to take action $a \in I_1^x$ under signal ω_1^β. (Otherwise, the inequality is violated since the receiver will take the same action $a'_a \in I_1^y$ under signal ω_1^α, signal ω_1^β, and the signal which reveals I_1^y.) Therefore, combining statements in condition 1.(a), condition 1.(c), and condition 2, we know S_1 can persuade the receiver taking actions $a \in I_1^x$ under signal ω_1^α and signal ω_1^β, and S_1 prefers sending signal ω_1^α.

Now let's explore S_2's preference on signals/information sets. Given condition 1.(d), $\mathbb{E}_p[U_{S_2}|\omega_1^\beta, G(p, I_1^x + \beta I_1^y)] < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]$ states that S_2 prefers a (pure) signal eliciting S_1's information I_1^y over a mixed signal ω_1^β. Moreover, condition 1.(a) states that the set of S_2 optimal signaling strategies under ω_1^β is the same as the set of S_2 optimal signaling strategies under ω_1^α. Thus, with $\alpha > \beta$ and the inferred fact that S_1 can persuade the receiver taking $a \in I_1^x$ under signal ω_1^α and signal ω_1^β, we can derive the inequality $\mathbb{E}_p[U_{S_2}|\omega_1^\alpha, G(p, I_1^x + \alpha I_1^y)] < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]$ from $\mathbb{E}_p[U_{S_2}|\omega_1^\beta, G(p, I_1^x + \beta I_1^y)] < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]$. Combining the two inequalities derived above, we can write the following inequality:

$$\max\left\{\mathbb{E}_p[U_{S_2}|\omega_1^\alpha, G(p, I_1^x + \alpha I_1^y)], \mathbb{E}_p[U_{S_2}|\omega_1^\beta, G(p, I_1^x + \beta I_1^y)]\right\} < \mathbb{E}_p[U_{S_2}|I_1^y, G(p, I_1^y)]. \tag{D.1}$$

Given the inequality (D.1), we can conclude that S_2 prefers a separation of S_1's information sets I_1^x and I_1^y.

Now, condition 1.(b) and condition 1.(c) guarantee the existence of a signaling strategy $\hat{\Gamma}_2$ satisfying the following three conditions:

- $\hat{\Gamma}_2(p, \omega_1^\beta) \in G(p, I_1^x + \beta I_1^y)$,
- $\hat{\Gamma}_2(p, \omega_1^\alpha) \notin G(p, I_1^x + \alpha I_1^y)$,
- $\mathbb{E}_p[U_{S_1}|\omega_1^\alpha, \hat{\Gamma}_2] < \frac{\mathbb{P}(I_1^x)+\beta\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)}\mathbb{E}_p[U_{S_1}|\omega_1^\beta, \hat{\Gamma}_2] + \frac{(\alpha-\beta)\mathbb{P}(I_1^y)}{\mathbb{P}(I_1^x)+\alpha\mathbb{P}(I_1^y)}\mathbb{E}_p[U_{S_1}|I_1^y, \hat{\Gamma}_2]$.

Therefore, $\hat{\Gamma}_2$ is a signaling strategy that can threaten S_1. Given $\hat{\Gamma}_2$, S_1 is better off by committing to signaling strategies using signal ω_1^β instead of signal ω_1^α, guaranteed by the condition 1.(c). This increases the total probability of the signal which elicits S_1's information I_1^y, and then increases S_2's expected utility (because S_2 prefers separation over mixture on the information sets I_1^x and I_1^y). In summary, given conditions (a)-(d) in condition 1 and the inequality in condition 2, S_2 has a conflict of interest with S_1 on the mixture of I_1^x and I_1^y. However, S_2 is willing to (partially) collaborate with S_1 under I_1^x when S_2 has no alternatives (because $G(p, I_1^x + \alpha I_1^y) = G(p, I_1^x + \beta I_1^y)$).

At a high level, a collaboration under I_1^x benefits both S_1 and S_2 no matter who commits first. However, S_1 and S_2 hold different opinions on the scale of collaboration, where S_1 want to maximize the total probability of collaboration, but S_2 only wants to collaborate when she has no alternative (signaling strategies that are incentive-compatible to the receiver). Given the first two conditions, if the prior distribution supports it, then both senders will collaborate on I_1^x. Condition 1.(b) and 1.(c) guarantees that S_2 can construct a credible threat to reduce the total probability of collaboration when she commits first, and this makes the commitment order matter. Finally, condition 3 serves as a regularity condition which ensures that the prior distribution supports the collaboration and the (potential) credible threat. Hence, when three conditions in Theorem 2 hold, the commitment order matters.

D.1.4 Proof of Theorem 3

The idea of this proof is to prove that the commitment order may matter if one of the conditions is not violated.

First, if the receiver's tie-breaking rule is belief-dependent, the sender who commits later can tailor the posterior beliefs to make the tie-breaking rule favor her. Hence, unless all tie-breaking decisions made by the receiver are indifferent to both senders S_1 and S_2, both prefer to commit last. In order to verify whether the commitment order matters when condition 1 holds, we assume that condition 2 is violated.

Suppose the receiver's tie-breaking rule is belief independent; we then prove that satisfying condition 1 may make the commitment order matters. The idea of proving the above statement is to construct the minimum requirements such that (partial) collaboration and credible threats can both occur. We will illustrate the conditions required for a (partial) collaboration and a credible threat below, respectively.

First, we derive the minimum requirement of a (partial) collaboration. S_1 and S_2 will collaborate when their preferences (ordering) on at least a pair of states align, called θ_α and θ_β hereafter for simplicity of representation. Moreover, if collaboration occurs in θ_α and θ_β, S_1 must be unable to distinguish θ_α and θ_β using her private information, i.e.,$\exists! I_1^x \ni \theta_\alpha, \theta_\beta$. Otherwise, S_1 will separate these two states herself (if it is best for her) and no collaboration occurs. Given assumption 4 in this chapter, S_2 can learn the true state via her private information when S_1 reveals her information truthfully. This demands θ_α and θ_β must belong to different information sets of S_2.

To make a collaboration on θ_α and θ_β possible, S_1 and S_2 must have aligned preference orders on these two states. Without loss of generality, we assume $U_{S_1}(a_\alpha) > U_{S_1}(a_\beta)$ and $U_{S_2}(a_\alpha) > U_{S_2}(a_\beta)$. Collecting the points together, to construct a (partial) collaboration, we need the following conditions to hold:

- $\exists! I_1^x \ni \theta_\alpha, \theta_\beta$
- $\exists! I_2^x \ni \theta_\alpha$ and $\exists! I_2^y \ni \theta_\beta$ such that $I_2^x \neq I_2^y$.
- $U_{S_1}(a_\alpha) > U_{S_1}(a_\beta)$
- $U_{S_2}(a_\alpha) > U_{S_2}(a_\beta)$

Next, we derive the minimum requirements for a (potential) conflict of interest. Before the analysis, we pause to note that a conflict of interest is not guaranteed under our requirements, because we do not specify the utilities of S_1 and S_2 on the receiver's best-response action corresponding to other states within the same information set of S_1 (if any). Precisely, when I_1^y is the information set involved in a (potential) conflict of interest between senders, we do not specify $U_{S_1}(\theta_\tau, a_\tau)$ and $U_{S_2}(\theta_\tau, a_\tau)$ for every state $\theta_\tau \in I_1^y$ not belonging to the set of states where we construct a (potential) conflict of interest.

A (potential) conflict of interest arises when S_1 and S_2 prefer different mixed strategies on a set of S_1's information sets. Since we are interested in the minimum requirements, we search for conflict of interests on a pair of information sets, called I_1^x and I_1^y, where the information set I_1^x is the same set where collaboration may occur. (Because we demand the occurrence of collaboration and a conflict of interests. If only collaborations occur, S_1 and S_2 can reach a consensus on their (collectively) optimal commitments. If only a conflict of interests occurs, S_1 will commit to a maximin strategy, and S_2 will commit to a minimax strategy. This is because S_1 always sends signals before S_2.)

Before constructing a (potential) conflict of interest, we assume θ_γ is a state that belongs to S_1's information I_1^y. As presented earlier, a (potential) conflict of interest requires that S_1 and S_2 have difference preference on a mixture of I_1^x and I_1^y. Here we assume[1] S_1 prefers a mixture over separation and S_2 prefers a separation over mixture.

Since S_1 prefers a mixture of I_1^x and I_1^y, there must exists a state $S_\gamma \in I_1^y$ such that $U_{S_1}(a_\alpha) > U_{S_1}(a_\gamma)$, where $\theta_\alpha \in I_1^x$ as used earlier. Similarly, when S_2 prefers a separation over a mixture, there must exists a state $S_\gamma \in I_1^y$ such that $U_{S_2}(a_\alpha) < U_{S_2}(a_\gamma)$. To avoid a violation of assumption 1, i.e., each sender only has partial information about the state of the world, in a ternary state space corner case, we let θ_γ belongs to either I_2^x or I_2^y. To sum up, to construct a minimum (potential) conflict of interest op top of a partial collaboration, the following additional conditions hold:

- $U_{S_1}(a_\alpha) > U_{S_1}(a_\gamma)$
- $U_{S_2}(a_\alpha) < U_{S_2}(a_\gamma)$
- Either $\theta_\gamma \in I_2^x$ or $\theta_\gamma \in I_2^y$.

Now, we have constructed the minimum requirement of collaboration and the minimum requirement of a conflict of interest. The collection of information structure requirements is the statement of condition 1.(a) in Theorem 3, and the requirements on S_1's and S_2's utilities are summarized in the inequalities in condition 1.(b). Therefore, when condition 1 is satisfied, the commitment order may matter because of the co-occurrence of collaboration and a conflict of interest. Thus, when both conditions are violated, the commitment order does not matter.

[1] For the opposite conflict of interest, i.e., S_1 prefers a separation over mixture and Sender2 prefers a mixture over separation, the whole analysis is exactly analogous with a swap of senders' utility inequalities. The inequalities restricting S_1 will now restrict S_2, and vice versa. In short, we just swap S_1 with S_2 in condition 1.(b).

APPENDIX E

Appendices of Chapter V

E.1 Proofs of Chapter V

E.1.1 Proof of Proposition 5

Proof. We will give an explicit construction of such a strategy. Let $\hat{\Gamma}$ be any signaling scheme. Let $a^*_{\hat{\Gamma}}$ be the receiver's (randomized) best response. Recall that $\pi_{i,j}$ is the probability with which the receiver plays a_i when receiving the signal ω_j. We will first construct an intermediate signaling scheme Γ', and from it the final signaling scheme Γ.

As a first step, the signaling scheme maps to an expanded space $\Omega' = \Omega \times A$. When observing the state θ_k, the sender sends the signal (ω_j, a_i) with probability $\mathbb{P}(a_j|\theta_k) \cdot \pi_{i,j}$. In other words, the sender performs exactly the randomization that the receiver would perform, and makes the corresponding recommendation to the receiver. Conditioned on the signal ω_j, the signal's second component a_i reveals no information about the state of the world. Therefore, because the distribution of a_i is exactly the distribution that $a^*_{\hat{\Gamma}}(\omega_j)$ uses, it is a best response for the receiver (and satisfies the constraints) to deterministically[1] follow the sender's "recommendation" a_i when receiving the signal (ω_j, a_i).

Finally, following the standard approach for reducing the size of the signal space, we "compress" all signals under which the receiver chooses the same action into one signal. That is, under the final signaling scheme Γ, whenever the sender was going to send (ω_j, a_i) for any j under Γ', the sender simply sends a_i. Because it is a best response for the receiver to deterministically choose a_i for all received (ω_j, a_i), it is still a best response to follow the recommendation a_i.

Thus, we have constructed a signaling scheme Γ such that the receiver plays a deterministic best response, and the number of signals employed by the sender is at most $|A|$.

Finally, to prove the existence of a sender-optimal signaling scheme with deterministic receiver response, let $\hat{\Gamma}$ be any sender-optimal signaling scheme. The existence of a signaling scheme, and thus a sender-optimal one, follows because the constraints are implementable by assumption. Then, applying the previous argument to $\hat{\Gamma}$ gives the desired optimal signaling scheme with deterministic receiver responses. □

E.1.2 Proof of Proposition 6

When a sender is monotone, $U_S(\theta_i, a_j) \geq U_S(\theta_i, a_k)$ for all $j < k$ implies the sender will never persuade a higher-indexed state to lower-indexed action solely, i.e., the sender will not mix θ_1 and θ_2 to persuade the receiver to take a_2. Note that the monotonicity doesn't rule out the sender to persuade a middle-indexed state using a mixture of higher-index state and lower-indexed action. Thus, in a ternary state model, there are five types of persuading schemes that the sender may use:

1. Mix θ_1 and θ_2 to persuade the receiver taking a_1,

[1]Note that it is optimal for the receiver to follow the recommendation due to the overall constraints. In isolation, the receiver may be better off deviating for some signals — however, doing so would violate a constraint, or come at the expense of having to choose an even more suboptimal action under another signal.

2. Mix θ_1 and θ_3 to persuade the receiver taking a_1,

3. Mix θ_2 and θ_3 to persuade the receiver taking a_2,

4. Mix θ_1, θ_2, and θ_3 to persuade the receiver taking a_1,

5. Mix θ_1, θ_2, and θ_3 to persuade the receiver taking a_2.

Now, let us consider how a non-trivial (binding) constraint affects each persuading scheme, respectively.

Type 1: Mix θ_1 and θ_2 to persuade the receiver to take a_1: When a constraint binds the sender, the inequalities $U_S(\theta_1, a_1) \geq U_S(\theta_2, a_1)$ and $U_S(\theta_2, a_2) \geq U_S(\theta_1, a_2)$ guarantees that $U_S(\theta_1, a_1) - U_S(\theta_1, a_2) \geq U_S(\theta_2, a_1) - U_S(\theta_2, a_2)$. Thus, the sender's adjustment (to satisfy the constraint) will reduce $\mathbb{P}(a_1|\theta_2)$ instead of $\mathbb{P}(a_1|\theta_1)$. Let's call the probability mass of θ_2 the sender reduced from suggesting a_1 is $\hat{p}$. Now, the sender can mix this $\hat{p}$ amount of θ_2 with θ_3 to persuade the receiver taking a_2 when $\omega_3 \neq 0$ in the original signaling scheme (otherwise there is no θ_3 the sender can mix with θ_2 under the current prior). Therefore, the constraint benefits the receiver when the true state is θ_2 but hurts the receiver when the true state is θ_3. Therefore, we have to derive the sender's signaling strategy which is used to persuade the receiver to take a_2 to determine whether the receiver is better off or worse off. The optimal signaling strategy for the sender (if the prior supports it) is to make the receiver indifferent between a_2 and a_3. In other words, the sender can mix $\hat{p}$ amount of θ_2 with $\beta\hat{p}$ amount of θ_3 satisfying the following equality:

$$\hat{p}U_R(\theta_2, a_2) + \beta\hat{p}U_R(\theta_3, a_2) = \hat{p}U_R(\theta_2, a_3) + \beta\hat{p}U_R(\theta_3, a_3).$$

Hence, we get $\beta = \frac{U_R(\theta_2,a_2)-U_R(\theta_2,a_3)}{U_R(\theta_3,a_3)-U_R(\theta_3,a_2)}$.

Now, we can compare the receiver's gain and loss owing to the constraints. First, the receiver benefits from taking action a_1 to a_2 when the true state is θ_2 with probability mass $\hat{p}$. However, the receiver suffers from taking action a_3 to a_2 when the true state is θ_3 with probability mass $\beta\hat{p}$. Hence, her overall utility change is

$$
\begin{aligned}
&\hat{p}\big(U_R(\theta_2,a_2)-U_R(\theta_1,a_2)\big)-\beta\hat{p}\big(U_R(\theta_3,a_3)-U_R(\theta_3,a_2)\big)\\
=&\hat{p}\big(U_R(\theta_2,a_2)-U_R(\theta_1,a_2)\big)-\frac{U_R(\theta_2,a_2)-U_R(\theta_2,a_3)}{U_R(\theta_3,a_3)-U_R(\theta_3,a_2)}\hat{p}\big(U_R(\theta_3,a_3)-U_R(\theta_3,a_2)\big)\\
=&\hat{p}\big(U_R(\theta_2,a_2)-U_R(\theta_1,a_2)\big)-\hat{p}\big(U_R(\theta_2,a_2)-U_R(\theta_2,a_3)\big)\\
=&\hat{p}\big(U_R(\theta_2,a_3)-U_R(\theta_2,a_1)\big)
\end{aligned}
$$

Given the second condition $U_R(\theta_2,a_1)\leq U_R(\theta_2,a_3)$ in the proposition, the receiver is not worse-off in this type of persuading scheme when a constraint binds.

Type 2: Mix θ_1 and θ_3 to persuade the receiver to take a_1: When a constraint binds the sender, unimodality $U_S(\theta_1,a_1)\geq U_S(\theta_3,a_1)$ and $U_S(\theta_3,a_3)\geq U_S(\theta_1,a_3)$ implies $U_S(\theta_1,a_1)-U_S(\theta_1,a_3)\geq U_S(\theta_3,a_1)-U_S(\theta_3,a_3)$. Hence, the sender's adjustment (to satisfy the constraint) will reduce $\mathbb{P}(a_1|\theta_3)$ instead of $\mathbb{P}(a_1|\theta_1)$. Since a_3 is the least preferred action across all states based on the sender's utility, the only chance the receiver can be hurt is when $\overline{\Gamma}_3=0$ in the original optimal signaling scheme. Then the sender can further increase the probability mass of θ_3 in ω_2 without changing the receiver's best response $a^*(\omega_2)$. When this occurs, the receiver's utility change is

$$
\big(U_R(\theta_3,a_3)-U_R(\theta_3,a_1)\big)-\big(U_R(\theta_3,a_3)-U_R(\theta_3,a_2)\big)=U_R(\theta_3,a_2)-U_R(\theta_3,a_1).
$$

Given the receiver's unimodality on states, $U_R(\theta_3,a_1)\leq U_R(\theta_3,a_2)$, the receiver is no worse off.

Type 3: Mix θ_2 and θ_3 to persuade the receiver to take a_2: Given the unimodality of the sender's preference on the state under given actions, the sender will reduce $\mathbb{P}(a_2|\theta_3)$ when a constraint binds. However, in this scenario, the sender cannot mix the excess probability of θ_3 in ω_1 to persuade the receiver to take a_1; otherwise, the original signaling scheme is not the optimal scheme. Hence, the sender can only increase $\mathbb{P}(a_3|\theta_3)$ and make the receiver strictly better off.

Type 4: Mix θ_1, θ_2, and θ_3 to persuade the receiver to take a_1: This type occurs when the prior probability on θ_1 is high enough to allow the sender to mix not only θ_2 but also θ_3 to persuade the receiver to take a_1. When a constraint binds, unimodality tells that the sender may reduce either $\mathbb{P}(a_1|\theta_2)$, $\mathbb{P}(a_1|\theta_3)$, or both to satisfy the constraint. In each case, the receiver is better off owing to the unimodality on state-matching since we are moving the probability mass from $\mathbb{P}(a_1|\theta_2)$ to $\mathbb{P}(a_2|\theta_2)$ or moving the probability mass from $\mathbb{P}(a_1|\theta_3)$ to $\mathbb{P}(a_2|\theta_3)$ or $\mathbb{P}(a_3|\theta_3)$.

Type 5: Mix θ_1, θ_2, and θ_3 to persuade the receiver to take a_2: When the sender's utility change from $U_S(\theta_1, a_1)$ and $U_S(\theta_1, a_2)$ is small, her optimal signaling strategy may use a small probability of θ_1 to increase the total probability of persuading the receiver to take a_2. (The sender may benefit from mixing θ_1 and θ_3 to persuade the receiver to take a_2.) When a constraint binds, the sender will reduce $\mathbb{P}(a_2|\theta_1)$ and $\mathbb{P}(a_2|\theta_3)$ to satisfy the constraint requirements. When this occurs, the sender may mix θ_1 with θ_3 trying to persuade the receiver to take a_1. However, given the optimality of the original signaling scheme, the sender cannot use all the reduced probability mass $\mathbb{P}(a_2|\theta_1)$ and $\mathbb{P}(a_2|\theta_3)$ to persuade the receiver to take a_1. (Otherwise, the sender should do it without the constraints too, leading to a higher expected utility.) Hence, the sender's optimal usage of the reduced probability mass $\mathbb{P}(a_2|\theta_1)$ and $\mathbb{P}(a_2|\theta_3)$ is to split it into two sets. The first set is a mixture of θ_1 and θ_3 which makes the receiver indifferent between a_1 and a_2 or between a_1 and a_2. The second set truthfully reveals θ_3. Now, from the receiver's perspective, since she is indifferent between a_1 and a_2 or between a_1 and a_3 in the first set, the sender's adjustment procedure is equivalent to eliciting true state θ_3 from a mixture of states where her best-response action is a_2 during the whole process. Hence, the receiver is always better off when this occurs.

To conclude, based on the 5 types, the receiver is no worse off when $U_R(\theta_2, a_1) \leq U_R(\theta_2, a_3)$ and the sender has a monotone preference across all states.

BIBLIOGRAPHY

BIBLIOGRAPHY

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011), Bayesian learning in social networks, *The Review of Economic Studies*, *78*(4), 1201–1236.

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2017), Fast and slow learning from reviews, *Tech. rep.*, National Bureau of Economic Research.

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2018), Informational Braess' paradox: The effect of information on traffic congestion, *Operations Research*, *66*(4), 893–917.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991), Optimal learning by experimentation, *The review of economic studies*, *58*(4), 621–654.

Au, P. H. (2015), Dynamic information disclosure, *The RAND Journal of Economics*, *46*(4), 791–823.

Babichenko, Y., and S. Barman (2017), Algorithmic aspects of private Bayesian persuasion, in *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, vol. 67, p. 34, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

Babichenko, Y., I. Talgam-Cohen, and K. Zabarnyi (2020), Bayesian persuasion under ex ante and ex post constraints, *arXiv preprint arXiv:2012.03272*.

Bahar, G., R. Smorodinsky, and M. Tennenholtz (2019), Social learning and the innkeeper's challenge, in *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 153–170.

Banerjee, A., and D. Fudenberg (2004), Word-of-mouth learning, *Games and economic behavior*, *46*(1), 1–22.

Banerjee, A. V. (1992), A simple model of herd behavior, *The quarterly journal of economics*, *107*(3), 797–817.

Basu, D. (1975), Statistical information and likelihood [with discussion], *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 1–71.

Bergemann, D., and S. Morris (2013), Robust predictions in games with incomplete information, *Econometrica*, *81*(4), 1251–1308.

Bergemann, D., and S. Morris (2016), Bayes correlated equilibrium and the comparison of information structures in games, *Theoretical Economics*, *11*(2), 487–522.

Bergemann, D., and S. Morris (2019), Information design: A unified perspective, *Journal of Economic Literature*, *57*(1), 44–95.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992), A theory of fads, fashion, custom, and cultural change as informational cascades, *Journal of political Economy*, *100*(5), 992–1026.

Bistritz, I., N. Heydaribeni, and A. Anastasopoulos (2019), Do informational cascades happen with non-myopic agents?, *arXiv preprint arXiv:1905.01327*.

Blackwell, D. (1953), Equivalent comparisons of experiments, *The annals of mathematical statistics*, pp. 265–272.

Bohren, J. A., and D. Hauser (2017), Bounded rationality and learning: A framework and a robustness result, *Working paper, submitted to Econometrica*.

Boleslavsky, R., and K. Kim (2018), Bayesian persuasion and moral hazard, available at SSRN 2913669.

Börgers, T. (2015), *An introduction to the theory of mechanism design*, Oxford University Press, USA.

Callander, S., and J. Hörner (2009), The wisdom of the minority, *Journal of Economic theory*, *144*(4), 1421–1439.

Candogan, O. (2020), Reduced form information design: Persuading a privately informed receiver, available at SSRN 3533682.

Candogan, O., and P. Strack (2021), Optimal disclosure of information to a privately informed receiver, available at SSRN 3773326.

Carroni, E., L. Ferrari, and G. Pignataro (2020), Does costly persuasion signal quality?, available at SSRN.

Celen, B., and S. Kariv (2004), Observational learning under imperfect information, *Games and Economic Behavior*, *47*(1), 72–86.

Celli, A., S. Coniglio, and N. Gatti (2019), Computing optimal ex-ante correlated equilibria in two-player sequential games, in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 909–917.

Chamley, C., and D. Gale (1994), Information revelation and strategic delay in a model of investment, *Econometrica: Journal of the Econometric Society*, pp. 1065–1085.

Cover, T. M. (1969), Hypothesis testing with finite statistics, *The Annals of Mathematical Statistics*, *40*(3), 828–835.

Crawford, V. P., and J. Sobel (1982), Strategic information transmission, *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

Dia, M. (2009), On decision making in tandem networks, Master's thesis, Massachusetts Institute of Technology.

Doval, L., and J. C. Ely (2020), Sequential information design, *Econometrica*, *88*(6), 2575–2608.

Doval, L., and V. Skreta (2018), Constrained information design: Toolkit, arXiv preprint arXiv:1811.03588.

Drakopoulos, K., A. Ozdaglar, and J. N. Tsitsiklis (2013), On learning with finite memory, *IEEE Transactions on Information Theory*, *59*(10), 6859–6872.

Drakopoulos, K., A. Ozdaglar, and J. N. Tsitsiklis (2016), When is a network epidemic hard to eliminate?, *Mathematics of Operations Research*, *42*(1), 1–14.

Dughmi, S. (2017), Algorithmic information structure design: A survey, *ACM SIGecom Exchanges*, *15*(2), 2–24.

Dughmi, S., and H. Xu (2019), Algorithmic Bayesian persuasion, *SIAM Journal on Computing*, pp. STOC16–68.

Dughmi, S., D. Kempe, and R. Qiang (2016), Persuasion with limited communication, in *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 663–680.

Easley, D., and J. Kleinberg (2012), Networks, crowds, and markets, *Cambridge Books*.

Ely, J. C. (2017), Beeps, *American Economic Review*, *107*(1), 31–53.

Ely, J. C., and M. Szydlowski (2017), Moving the goalposts, *Tech. rep.*, Working paper.

Ely, J. C., and M. Szydlowski (2020), Moving the goalposts, *Journal of Political Economy*, *128*(2), 468–506.

Farhadi, F., and D. Teneketzis (2020), Dynamic information design: A simple problem on optimal sequential information disclosure, *Available at SSRN 3554960*.

Forges, F., and F. Koessler (2008), Long persuasion games, *Journal of Economic Theory*, *143*(1), 1–35.

Fudenberg, D., and J. Tirole (1991), *Game Theory*, vol. 1, 1 ed., The MIT Press.

Gentzkow, M., and E. Kamenica (2014), Costly persuasion, *American Economic Review*, *104*(5), 457–62.

Gentzkow, M., and E. Kamenica (2016), Competition in persuasion, *The Review of Economic Studies*, *84*(1), 300–322.

Gentzkow, M., and E. Kamenica (2017), Bayesian persuasion with multiple senders and rich signal spaces, *Games and Economic Behavior*, *104*, 411–429.

Gilboa, I., and D. Schmeidler (1989), Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics*, *18*(2), 141 – 153, doi:https://doi.org/10.1016/0304-4068(89)90018-9.

Gradwohl, R., N. Hahn, M. Hoefer, and R. Smorodinsky (2021), Algorithms for persuasion with limited communication, in *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 637–652, SIAM.

Grossman, S. J. (1981), The informational role of warranties and private disclosure about product quality, *The Journal of Law and Economics*, *24*(3), 461–483.

Gul, F., and R. Lundholm (1995), Endogenous timing and the clustering of agents' decisions, *Journal of political Economy*, *103*(5), 1039–1066.

Guo, Y., and E. Shmaya (2019), The interval structure of optimal disclosure, *Econometrica*, *87*(2), 653–675.

Hann-Caruthers, W., V. V. Martynov, and O. Tamuz (2018), The speed of sequential asymptotic learning, *Journal of Economic Theory*, *173*, 383–409.

Hedlund, J. (2015), Persuasion with communication costs, *Games and Economic Behavior*, *92*, 28–40.

Hedlund, J. (2017), Bayesian persuasion by a privately informed sender, *Journal of Economic Theory*, *167*, 229–268.

Hellman, M. E., and T. M. Cover (1970), Learning with finite memory, *The Annals of Mathematical Statistics*, pp. 765–782.

Ho, J., W. P. Tay, and T. Q. Quek (2014), Robust detection and social learning in tandem networks, in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pp. 5457–5461, IEEE.

Honryo, T. (2018), Dynamic persuasion, *Journal of Economic Theory*, *178*, 36–58.

Huck, S., and J. Oechssler (1998), Informational cascades with continuous action spaces, *Economics Letters*, *60*(2), 163–166.

Hurwicz, L., and S. Reiter (2006), *Designing economic mechanisms*, Cambridge University Press.

Kamenica, E. (2019), Bayesian persuasion and information design, *Annual Review of Economics*, *11*, 249–272.

Kamenica, E., and M. Gentzkow (2011), Bayesian persuasion, *American Economic Review*, *101*(6), 2590–2615.

Kolotilin, A. (2015), Experimental design to persuade, *Games and Economic Behavior*, *90*, 215–226.

Kolotilin, A. (2018), Optimal information disclosure: A linear programming approach, *Theoretical Economics*, *13*(2), 607–635.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017), Persuasion of a privately informed receiver, *Econometrica*, *85*(6), 1949–1964.

Koplowitz, J. (1975), Necessary and sufficient memory size for m-hypothesis testing, *IEEE Transactions on Information Theory*, *21*(1), 44–46.

Le, T., V. Subramanian, and R. Berry (2018), Bayesian learning with random arrivals, in *Information Theory (ISIT), 2018 IEEE International Symposium on*, pp. 6990–6995, IEEE.

Le, T. N., V. G. Subramanian, and R. A. Berry (2014a), The impact of observation and action errors on informational cascades, in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pp. 1917–1922, IEEE.

Le, T. N., V. G. Subramanian, and R. A. Berry (2014b), The value of noise for informational cascades, in *Information Theory (ISIT), 2014 IEEE International Symposium on*, pp. 1101–1105, IEEE.

Le, T. N., V. G. Subramanian, and R. A. Berry (2016), Quantifying the utility of imperfect reviews in stopping information cascades, in *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 6990–6995, IEEE.

Le, T. N., V. G. Subramanian, and R. A. Berry (2017a), Information cascades with noise, *IEEE Transactions on Signal and Information Processing over Networks*, *3*(2), 239–251.

Le, T. N., V. G. Subramanian, and R. A. Berry (2017b), Learning from randomly arriving agents, in *Communication, Control, and Computing (Allerton), 2017 55th Annual Allerton Conference on*, pp. 196–197, IEEE.

Le Treust, M., and T. Tomala (2019), Persuasion with limited communication capacity, *Journal of Economic Theory*, *184*, 104,940.

Lee, I. H. (1993), On the convergence of informational cascades, *Journal of Economic theory*, *61*(2), 395–411.

Li, F., and P. Norman (2018a), On Bayesian persuasion with multiple senders, *Economics Letters*, *170*, 66–70.

Li, F., and P. Norman (2018b), Sequential persuasion, *Available at SSRN 2952606.*

Li, J., and J. Zhou (2016), Blackwell's informativeness ranking with uncertainty-averse preferences, *Games and Economic Behavior*, *96*, 18–29.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995), *Microeconomic theory*, vol. 1, Oxford university press New York.

Matyskova, L. (2018), Bayesian persuasion with costly information acquisition, cERGE-EI Working Paper Series.

Meigs, E., F. Parise, A. Ozdaglar, and D. Acemoglu (2020), Optimal dynamic information provision in traffic routing, *arXiv preprint arXiv:2001.03232*.

Milgrom, P. R. (1981), Good news and bad news: Representation theorems and applications, *The Bell Journal of Economics*, pp. 380–391.

Monzón, I. (2017), Aggregate uncertainty can lead to incorrect herds, *American Economic Journal: Microeconomics*, *9*(2), 295–314.

Mossel, E., A. Sly, and O. Tamuz (2015), Strategic learning and the topology of social networks, *Econometrica*, *83*(5), 1755–1794.

Myerson, R. B. (1981), Optimal auction design, *Mathematics of operations research*, *6*(1), 58–73.

Nguyen, A., and T. Y. Tan (2021), Bayesian persuasion with costly messages, *Journal of Economic Theory*, *193*, 105,212.

Osborne, M. J., and A. Rubinstein (1994), *A course in game theory*, MIT press.

Peres, Y., M. Z. Racz, A. Sly, and I. Stuhl (2017), How fragile are information cascades?, *arXiv preprint arXiv:1711.04024*.

Perez-Richet, E. (2014), Interim bayesian persuasion: First steps, *American Economic Review*, *104*(5), 469–74.

Perez-Richet, E., and D. Prady (2012), Complicating to persuade?, unpublished.

Rayo, L., and I. Segal (2010), Optimal information disclosure, *Journal of political Economy*, *118*(5), 949–987.

Sgroi, D. (2002), Optimizing information in the herd: Guinea pigs, profits, and welfare, *Games and Economic Behavior*, *39*(1), 137–166.

Sgroi, D. (2003), Irreversible investment and the value of information gathering, *Economics Bulletin*, *21*(4).

Smith, L., and P. Sørensen (2000), Pathological outcomes of observational learning, *Econometrica*, *68*(2), 371–398.

Smith, L., and P. Sørensen (2008), Rational social learning with random sampling, *Available at SSRN 1138095*.

Song, Y. (2015), Social learning with endogenous network formation, *arXiv preprint arXiv:1504.05222*.

Sørensen, P. N. (1996), Rational social learning, Ph.D. thesis, Massachusetts Institute of Technology.

Spence, M. (1978), Job market signaling, in *Uncertainty in economics*, pp. 281–306, Elsevier.

Tay, W. P., J. N. Tsitsiklis, and M. Z. Win (2008), On the subexponential decay of detection error probabilities in long tandems, *IEEE Transactions on Information Theory*, *54*(10), 4767–4771.

Tay, W. P., J. N. Tsitsiklis, and M. Z. Win (2009), Bayesian detection in bounded height tree networks, *IEEE Trans. Signal Process.*, *57*(10), 4042 – 4051.

Tsitsiklis, J. N. (1989), Decentralized detection, in *Advances in Statistical Signal Processing, Vol. 2: Signal Detection*, edited by H. V. Poor and J. B. Thomas, JAI Press.

Vaccari, S., M. Scarsini, and C. Maglaras (2016), Social learning in a competitive market with consumer reviews, *Columbia Business School Research Paper No. 16-52*.

Vial, D., and V. Subramanian (2019), Local non-bayesian social learning with stubborn agents, in *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 551–552.

Vives, X. (1997), Learning from others: a welfare analysis, *Games and Economic Behavior*, *20*(2), 177–200.

Wang, Y., and P. M. Djurić (2015), Social learning with bayesian agents and random decision making, *IEEE Transactions on Signal Processing*, *63*(12), 3241–3250.

Welch, I. (1992), Sequential sales, learning, and cascades, *The Journal of finance*, *47*(2), 695–732.

Wu, J. (2015), Helpful laymen in informational cascades, *Journal of Economic Behavior & Organization*, *116*, 407–415.

Yao, J., Z. Cheng, J. Dai, A. Chen, and S. An (2019), Traffic assignment paradox incorporating congestion and stochastic perceived error simultaneously, *Transportmetrica A: Transport Science*, *15*(2), 307–325.

Zhang, Y. (2009), Robust information cascade with endogenous ordering, *working paper. Available at: http://www. gtcenter. org/Archive/2010/Conf/Zhang1085. pdf.*

Zhang, Z., E. K. Chong, A. Pezeshki, and W. Moran (2013), Hypothesis testing in feedforward networks with broadcast failures, *IEEE Journal of Selected Topics in Signal Processing*, *7*(5), 797–810.