

Hierarchical Bayesian Persuasion*

Zhonghong Kuang[†]

Jaimie W. Lien[‡]

Jie Zheng[§]

Updated December 4, 2019; First Version May 14, 2018

ABSTRACT: We study a hierarchical Bayesian persuasion game with a sender, a receiver and several potential intermediaries, generalizing the framework of Kamenica and Gentzkow (2011, AER). The sender must be persuasive through a hierarchy of intermediaries in order to reach the final receiver, whose action affects all players' payoffs. The intermediaries care not only about the true state of world and the receiver's action, but also about their reputations, measured by whether the receiver's action is consistent with their recommendation. We characterize the subgame perfect equilibria for the optimal persuasion strategy, and show that the persuasion game has multiple equilibria but a unique payoff outcome. Among the equilibria, two natural persuasion strategies on the hierarchy arise: persuading the intermediary who is immediately above one's own position, and persuading the least persuadable individual in the hierarchy. Furthermore, we show that the order of persuasion matters, and that adding intermediaries to the hierarchy can make the sender better off. As major extensions of the model, we analyze scenarios in which intermediaries have private information, endogenized reputation, and an outside option, respectively. We also discuss as minor extensions, the endogenous choice of persuasion path, parallel persuasion, and costly persuasion. The results provide insights for settings where persuasion is prominent in a hierarchical structure, such as corporate management, higher education, and government bureaucracies.

KEYWORDS: Bayesian Persuasion, Hierarchy, Reputation, Sequential Persuasion

JEL CLASSIFICATION: C72, C73, D82, D83

*We are especially grateful to Jian Chen and Matthew Gentzkow for detailed comments which helped to improve the manuscript. For helpful comments and insights, we thank Masaki Aoyagi, Jimmy Chan, Bo Chen, Zhuoqiong Chen, Kim-Sau Chung, Vincent Crawford, Fuhai Hong, Matthew Jackson, Rongzhu Ke, Chiu Yu Ko, Ping Lin, Melody Lo, Mihai Manea, Charles Noussair, Joel Sobel, Jean Tirole, Ruqu Wang, Adam Wong, Fan Wu, Qinggong Wu, Boyu Zhang, Mu Zhang, Hangcheng Zhao, Mofei Zhao, Weijie Zhong, conference participants at the 2018 Beijing International Workshop on Microeconomics Empirics, Experiments and Theory, 6th China Conference on Organizational Economics, 2018 Annual Conference of China Information Economics Society, 2019 Asian Meeting of Econometric Society, 2019 China Meeting of Econometric Society, 2019 SWUFE Economic Theory Workshop and 30th International Conference on Game Theory, and seminar participants at Tsinghua University, Lingnan University, Chinese University of Hong Kong, and Hong Kong Baptist University. We are grateful for the funding support from National Natural Science Foundation of China (Projects No. 61661136002 and No. 71873074), Hong Kong Research Grants Council, General Research Fund (Projects No. 14500516 and 14501919), Tsinghua University Tutor Research Fund, and Chinese University of Hong Kong Direct Grants. All errors are our own.

[†]Zhonghong Kuang: Department of Management Science and Engineering, School of Economics and Management, Tsinghua University, Beijing, China (email: kuangzh.15@sem.tsinghua.edu.cn).

[‡]Jaimie W. Lien: Department of Decision Sciences and Managerial Economics, CUHK Business School, The Chinese University of Hong Kong, Hong Kong SAR, China (email: jaimie.academic@gmail.com).

[§]Jie Zheng (Corresponding Author): Department of Economics, School of Economics and Management, Tsinghua University, Beijing, China, 100084 (e-mail: jie.academic@gmail.com).

1 Introduction

An individual hoping to persuade a decision-making authority in the direction of his or her preferred outcome often does not have direct access to that person-in-charge. An entry level worker at a company may have a profitable idea, but he can neither implement it himself nor persuade the CEO directly. Instead, he must persuade his manager, who if persuaded, persuades his own manager, and so on through the corporate hierarchy until a decision-making authority is convinced. In some government bureaucracies, persuasion of officials regarding the method of implementing public policies also occurs in a hierarchical manner.

Similar persuasion structures are commonplace whenever the incentive to persuade originates from the bottom of a hierarchy, upward. When applying for graduate school, students must convince their professors that they are suitable for future study, who if well-convinced, write recommendation letters to convince admissions committees to admit the student. Similarly, job promotion processes in academic settings often require an endorsement at the department level, followed by the faculty or division level, and so on, before being approved at the university level.

We study the persuasion strategies and equilibrium characteristics in a hierarchical version of Kamenica and Gentzkow's (2011) Bayesian persuasion model. In their framework, which is increasingly applied to a variety of settings in the literature, a sender with a strictly preferred outcome over a receiver's actions, commits to a conditional signaling strategy before the realization of a state variable which affects both of their payoffs. Their framework shows that although the receiver is fully aware of the sender's strategy, he can be persuaded towards the sender's desired action. The intuition is that by committing ex-ante to a randomized strategy which conditions on the state of the world, the sender can in expectation, convince a Bayesian receiver that his personally desired action is more appropriate, compared to without such a persuasive strategy.

The Bayesian persuasion framework is well-suited to analyze and understand institutionally supported signaling strategies by the sender, in other words, settings in which the sender holds a commitment against revising his strategy after the realization of the true state of the world. As such, the framework may be specifically appropriate for studying hierarchical situations, in which some degree of bureaucracy in the system commits senders along the hierarchy to adhering to a strategy. An example is in industrial settings, in which a middle-manager may make clear his conditional endorsement policy to upper management ex-ante, before the test results of a product by his employee are known. Our study is the first to our knowledge to study this problem in a Bayesian persuasion framework.

We analyze a benchmark model and several extensions. In our benchmark model, a sender attempts to persuade a receiver through a sequence of intermediaries. Each intermediary thus serves as not only a receiver, but a sender to the next intermediary in the sequence. To capture the strategic consideration in such a sequential setting, we introduce an additional reputation concern. Each intermediary cares not only about the final decision of the receiver, but also about

whether that decision matches his own recommendation. We formalize this reputation concern by having each intermediary provide feedback to his immediately preceding intermediary about the message he intends to deliver to the next intermediary, and receives a reputation premium if the final action taken by the receiver matches his feedback.

Under this framework, we solve for the subgame perfect equilibria. Multiple persuasion equilibria exist in the game. The persuasion solution for each intermediary takes into account three components: the message directly received, the intermediary’s own threshold for being convinced, and the behavior of most difficult to convince player along the hierarchy. Among the multiple solutions are some intuitive strategies for persuasion. For example, a sender may target his persuasion strategy towards the intermediary he communicates with directly. Another approach is to target the persuasion strategy towards the most difficult person to convince in the chain. Despite the multiplicity of equilibria, our analysis shows that all of them are outcome (i.e. payoff) equivalent.

We also analyze the impact of the ordering of intermediaries on sender’s utility in the benchmark model, which yields several insights. Although direct persuasion is at least weakly preferred to indirect persuasion, perhaps surprisingly, the sender can be made better-off if an extra intermediary is introduced into a non-trivial hierarchy. In particular, the sender stands to benefit from an additional intermediary whose incentives which are most aligned with his own among all the existing members of the hierarchy. Furthermore, when the sender has his choice of orderings among the set of intermediaries, we show that the optimal situation is for this most sender-aligned intermediary to persuade the receiver directly.

Our work belongs to the expanding literature on Bayesian persuasion, following the seminal work of Kamenica and Gentzkow (2011)^[1]. Bergemann and Morris (2019)^[2] and Kamenica (2019)^[3] provide literature reviews on the general topic of information design. Since in the hierarchical structure we consider, more than one player sends a signal and more than one player receives a signal, our work is related to studies on Bayesian persuasion with multiple senders (Gentzkow and Kamenica, 2017a^[4]; Gentzkow and Kamenica, 2017b^[5]; Li and Norman, 2018a^[6]; Li and Norman, 2018b^[7]) and those on Bayesian persuasion with multiple receivers (Alonso and Camara, 2016a^[8]; Alonso and Camara, 2016b^[9]; Arieli and Babichenko, 2016^[10]; Bardhi and Guo, 2018^[11]; Chan, Gupta, Li and Wang, 2019^[12]; Hoshino, 2017^[13]; Laclau and Renou, 2016^[14]; Michaeli, 2017^[15]; Shimoji, 2016^[16]; Song and Zhao, 2019^[17]; Wang, 2015^[18]). A few studies examine settings with multiple senders and multiple receivers (Board and Lu, 2018^[19]; Koessler, Laclau and Tomala, 2018^[20]; Wu and Zheng, 2019^[21]). The key distinction between our work and those mentioned above is that in our setup we use the hierarchical structure of persuasion - the intermediaries receiving a message are also using different persuasion strategies along the chain of players.

There are also studies focusing on Bayesian persuasion with a mediator or moderator. Kosenko (2018)^[22] studies a Bayesian persuasion with mediators who can garble the signals generated from sender. Qian (2019)^[23] studies Bayesian persuasion with one moderator who can verify and choose whether to faithfully deliver the realized message. Our model differs from these papers in three

aspects. Firstly, in our framework, intermediaries can commit to information disclosure policies conditioning on the state, which is similar to the role of the sender in the original model of Kamenica and Gentzkow (2011). Secondly, instead of distorting signals, intermediaries must truthfully pass along the signal generated by preceding players. Thirdly, we introduce a reputation concern into the model, which in such a setting, enables each intermediary to be strategic along the chain.

The literature on dynamic information design also examines the release of information sequentially, but focuses on the interaction between one sender and one receiver, rather than a hierarchy (Doval and Ely, 2016^[24]; Ely, 2017^[25]; Ely, Frankel and Kamenica, 2015^[26]; Ely and Szydlowski, 2019^[27]; Felgenhauer and Loerke, 2017^[28]; Henry and Ottaviani, 2019^[29]; Horner and Skrzypacz, 2016^[30]; Orlov, Skrzypacz and Zryumov, 2017^[31]; Renault, Solan and Vieille, 2017^[32]).

Another line of research that is closely related to our work is that of hierarchical cheap talk, following the framework of Crawford and Sobel (1984).^[33] The main difference between the cheap talk and persuasion frameworks, is that the cheap talk framework asks to what degree communication without any ex-ante commitment can be informative under conflict of interest between sender and receiver, while the persuasion framework asks whether under similar conflict of interest, a communication strategy with full ex-ante commitment can persuade a Bayesian receiver. Ivanov (2010)^[34] studies the information transmission from sender to receiver through an intermediary. Ambrus, Azevedo, and Kamada (2013)^[35] introduces a chain of intermediaries in the cheap talk framework.

In the first main extension of our benchmark model of hierarchical persuasion, we study the possibility of private information in the hierarchical chain (see Kolotilin, Mylovannov, Sapechelnnyuk and Li, 2017^[36]; Kolotilin, 2018^[37], for non-hierarchical incomplete information persuasion models). In this setting, the key information that is unknown to the players in the game is the persuasion threshold of intermediaries and the final receiver, which can be calculated backwards through the hierarchy based on unknown utilities and reputation terms, in other words, how difficult they are to convince. Although a multiplicity of equilibria exist, the intuitive strategy of focusing on persuading the immediately subsequent intermediary remains, while the strategy of persuading the most difficult to persuade player may no longer be an equilibrium. To gain insights, it is useful to consider some special cases of the private information model. In the special case that only the final receiver has private information, both of these intuitive persuasion approaches are equilibria. Furthermore, for any case in which the private information of receiver is uniformly distributed, persuasion is ineffective.

Our second major extension of the benchmark model is to endogenize the aforementioned reputation concern of intermediaries using an infinitely repeated game setting. In the repeated persuasion interaction, intermediaries that are successful in implementing the desired outcome are more likely to have access to subsequent rounds of the persuasion game in which their reputation gain can be reaped. In other words, successful intermediaries have greater opportunities to engage in the persuasion interaction again compared to less successful intermediaries. We show that the

reputation premium utilized in our benchmark model can arise naturally as a result of a repeated interaction setting. Best and Quigley (2017)^[38] creates a link between future credibility and the current credibility for a long-lived sender. In our model, we consider a long-lived intermediary and show that future gains affect choices in the present.

Our last major extension considers the possibility that intermediaries have an outside option in their message set, which is to decline relaying any message altogether. An intermediary has a potential incentive to avoid relaying any message due to his reputation concern, should the undesired action be eventually taken by the receiver. However, the failure to pass a message to the next intermediary in a strictly sequential hierarchy breaks the persuasion chain. If preceding intermediaries believe that a subsequent intermediary is likely to opt out of persuasion, they are also hesitant to give their recommendations, out of reputation concern. Additionally, if an intermediary in the game with outside option observes that all preceding intermediaries have provided a message, his own concern about a subsequent intermediary opting out is lessened. In this sense, the hierarchical persuasion game with an outside messaging option bears some resemblance to a rational herding framework.

We also discuss several minor extensions of the model. We allow for an endogenized choice of the persuasion path among many possible paths of intermediaries and provide an algorithm which solves this problem. In this setting, senders will choose the path of least resistance, in the sense that the most-difficult-to-persuade person along the chosen path is in fact the easiest to persuade among all the most-difficult-to-persuade individuals on all paths. We also generalize the benchmark framework to allow for parallel persuasion paths as well as sequential paths, and show that the main results hold. Finally, we consider the case where persuasion is costly (see Gentzkow and Kamenica, 2014^[39]; Hodler, Loertscher, and Rohner, 2014^[40] on costly persuasion), and show the main results are robust under the framework of Gentzkow and Kamenica (2014), but in a hierarchical setting.

The remainder of the paper is organized as follows: Section 2 describes the main model and benchmark analyses; Sections 3 through 5 present our main extensions of the main model: Section 3 extends the model to the case of private information; Section 4 endogenizes the reputation term in an infinitely repeated game; Section 5 considers the scenario that players have outside options; Sections 6 through 8 contain further discussions of the model: Section 6 endogenizes the path of persuasion; Section 7 extends the sequential structure to include parallel persuasion paths; Section 8 extends the model to the case of costly persuasion; Section 9 concludes and describes future work. All proofs can be found in the Appendix.

2 Model and Analysis

We use the admissions process for academic programs as our referring example throughout much of the paper. Although clearly, in the real-world this example may not match our model in

terms of all aspects, it provides a useful tangible context for the forces that the model seeks to represent and analyze.

In such a context, the sender is the student wishing to apply for an academic program, while the receiver is the decision-maker in charge of admissions at the university. The intermediaries include the professor of the student, who writes a recommendation letter to the admissions committee, as well as other possible intermediaries in the chain, such as in some scenarios, the potential Ph.D. supervisor of the student, the department level admissions chair, faculty level admissions chair and so on.

In this example of hierarchy, the student's goal is to be admitted to the academic program. The student persuades the faculty letter writer of his suitability for the program. Here, note that as in Kamenica and Gentzkow (2011), persuasion can be interpreted as not only verbal communication, but presenting a collection of evidence, which in the student's case could include academic performance, research papers and other measures of suitability for further study. The professor will ex-ante design his own tests for the student, which may depend on the evidence provided by the student. The professor will first respond to the student based on the evidence, then conduct the relevant test, and eventually send the letter, which will depend on the realized outcome of the professor's test. By writing the letter in support of the student, the letter writer conveys a message to the faculty member on the department admissions committee, who after examining the contents of the letter, may also conduct their own test on the student, and convey their own recommendation to the next level of deliberation.

The process continues until reaching the receiver, or the final decision-maker in the admissions decision. Suppose that the final decision-maker has preferences over outcomes which differ from the student's. For example, they may only want to admit the student if the student is a good scholar.

2.1 Benchmark Model Setup

The hierarchical persuasion framework consists of one initiating sender and one final receiver, through n intermediaries one-by-one, subscripted by $j = 1, 2, \dots, n$. Players' payoffs depend on the state of the world $t \in T = \{\alpha, \beta\}$. We assume that the common prior probability distribution over the states T are given by: $\mathbb{P}(\alpha) = 1 - p_0, \mathbb{P}(\beta) = p_0$, where $p_0 \in [0, 1]$. The decision d is chosen by the final receiver, with potential choices denoted by $D = \{A, B\}$.¹ We refer to decision A as the *default* action and B as the *proposed* action.

The utility functions of the sender, receiver and intermediaries are state-dependent. Let $u^S(d, t)$, $u^R(d, t)$ and $u_j(d, t)$ denote the utility of sender, receiver and intermediary j respectively, that is derived from the decision d in state t . Without loss of generality, we normalize the

¹In our academic admissions example, those possible decisions are $\{reject, admit\}$.

utility associated with action A to zero.

$$u^S(A, \alpha) = u^S(A, \beta) = u^R(A, \alpha) = u^R(A, \beta) = u_j(A, \alpha) = u_j(A, \beta) = 0$$

The sender always prefers action B regardless of the state and hence tries to persuade the receiver to take the *proposed* action,

$$u^S(B, t) > 0, \forall t \in T$$

However, the receiver and intermediaries, on the other hand, prefer the *proposed* action if and only if the state is β ,²

$$u^R(B, \alpha) < 0 < u^R(B, \beta), u_j(B, \alpha) < 0 < u_j(B, \beta), \forall j = 1, \dots, n$$

We allow for heterogeneous preferences of different players, who might derive different levels of utility from each of the implemented alternatives, and assign different levels of utility loss to undesirable decisions.³ All the players' payoff structures are common knowledge.

We utilize the concept and notation of a belief threshold, which is also adopted in Wang (2015), Bardhi and Guo (2018), and Chan, Gupta, Li and Wang (2019)^[18;11;12]. Assume that the belief of the receiver about the state is $\mathbb{P}(\beta) = p$. Then his payoff will be 0 by implementing A and $pu^R(B, \beta) + (1 - p)u^R(B, \alpha)$ by implementing B . The receiver prefers action B if and only if $p > \frac{-u^R(B, \alpha)}{u^R(B, \beta) - u^R(B, \alpha)} \in (0, 1)$. Let \tilde{p}_R denote this threshold value. Then the receiver prefers B if and only if belief $\mathbb{P}(\beta)$ exceeds \tilde{p}_R . Likewise, we denote $\tilde{p}_1, \dots, \tilde{p}_n$ as the threshold belief of n intermediaries, respectively. Without the reputation term, one player is simply more difficult to convince if he has a higher threshold belief.

The reputation concern of the intermediaries is a central concept for both the mechanics and interpretation of our model in the hierarchical persuasion context. When intermediary j is trying to persuade intermediary $j + 1$, it can be reasonable to assume that before turning to intermediary $j + 2$, intermediary $j + 1$ replies to intermediary j with his preferred action $\kappa_{j+1} \in \{A, B\}$.⁴ If the action taken by the receiver is indeed this action, intermediary $j + 1$ will earn a reputation gain of R_{j+1} , which can be interpreted as his trustworthiness or status in the eyes of intermediary j . Otherwise, intermediary $j + 1$ will have no reputation gain. Reputation loss is also possible, however, the normalization is without loss of generality.⁵

²In our example, the professors, committee chairs, and final admissions decision-maker prefer to admit the student if and only if he is a good scholar.

³For example, either the faculty letter writer or the admissions chair may incur a higher disutility if a good student is rejected, or if a bad student is admitted to the program.

⁴For example, a professor may tell a student that he is happy to write the student a good letter of recommendation. A faculty level admissions chair may tell the department level admissions chair that he will forward the department level admission suggestion to the next level of evaluation.

⁵Without this reputation concern, the model will reduce to one in which all intermediaries are truth-telling in equilibrium, in the sense of simply passing forward the information received along the hierarchy.

Timeline of the Game Each intermediary participates in two stages, a *response* stage in which intermediaries reply to the preceding intermediary to report their intended message, and a *persuasion* stage in which that intermediary persuades the succeeding intermediary. Note that as the initial and terminal nodes of the hierarchy, the sender only has a *persuasion* stage while the receiver only has a *response* stage. In the *response* stage, an intermediary chooses a reply from $D = \{A, B\}$ after the preceding intermediary implements his persuasion strategy. **Figure 1** illustrates the timing of response and persuasion stages. A black dashed rectangle represents a specific meeting in which player j attempts to persuade player $j + 1$ and player $j + 1$ responds to player j .

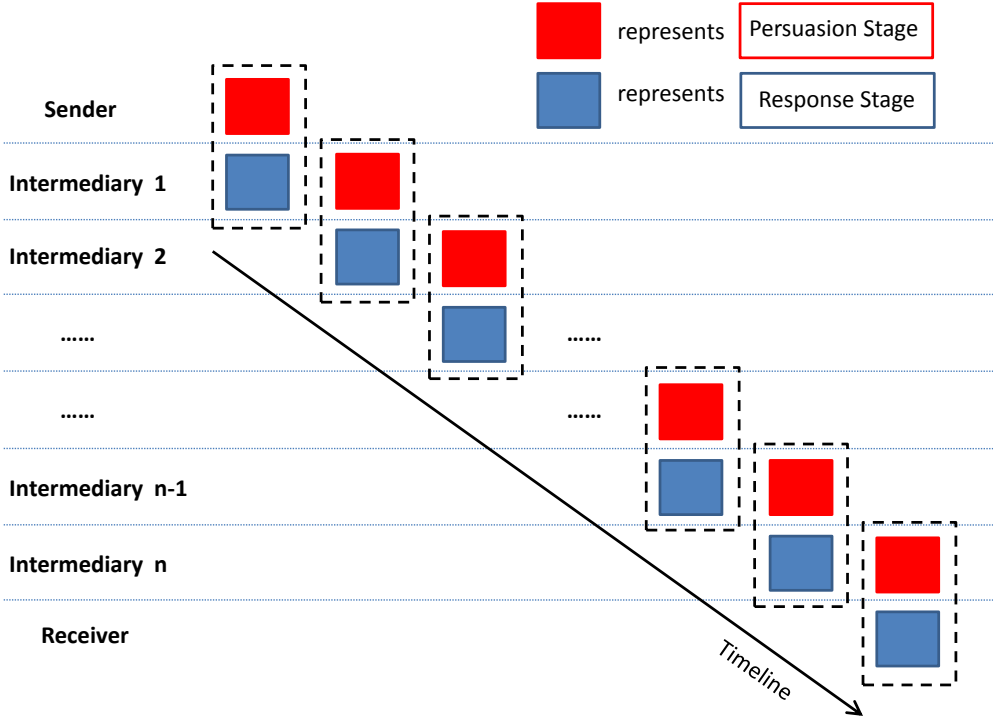


Figure 1: Stages of the Game

1. Commitment Process

We follow in the commitment feature of the original model by Kamenica and Gentzkow (2011). In their framework, the sender’s persuasion strategy is a commitment to a distribution of signals, conditional on the true state of the world. In the hierarchical setting, the sender has such a persuasion strategy, while the intermediaries commit to a distribution of signals conditional on both the state and their signal received from the sender or previous intermediaries.

We can think of this commitment to a conditional distribution as a type of formal policy which is made known to the other players in the game. For example, in the academic

recommendation case, a professor may adopt a policy to only provide a strong to very strong recommendation letter to a student who has met certain qualifications, such as specific coursework and research experience. A department admissions committee chair may have a policy to highly recommend a student for admission at the faculty level if the student has obtained certain grade point average and test score levels. In a hierarchical setting which has some elements of a bureaucracy, it may be reasonable in particular, for intermediaries to make such policies or commitment strategies.

- The sender publicly sets up a signal-generating mechanism, which consists of a family of conditional distributions $\{\pi(\cdot|t)\}_{t \in T}$ over a space of signal realizations S and hence divides the prior belief into posterior portfolios that satisfy a Bayes plausible condition.⁶
- Intermediary 1 publicly sets a response rule⁷ as well as a signal-generating mechanism which consists of a family of conditional distributions $\{\pi(\cdot|s, t)\}_{s \in S, t \in T}$ over a space of signal realizations I_1 , depending on state $t \in T$ and history $\mathcal{H}_1 \in \mathbb{H}_1$, where \mathcal{H}_1 consists of public information received by intermediary 1 in the realization process (namely signal s) and \mathbb{H}_1 denotes the space of \mathcal{H}_1 .⁸ The intermediary hence divides his incoming belief conditioning on \mathcal{H}_1 into posterior portfolios that satisfy the Bayes plausible condition as well.
- Then, sequentially and publicly, intermediaries $j = 2, \dots, n$, announce their response rules and hence divide their incoming beliefs conditioning on $\mathcal{H}_j \in \mathbb{H}_j$ into posterior portfolios, where \mathcal{H}_j consists of public information revealed by all preceding players in realization process (namely signal s, i_1, \dots, i_{j-1} and responses $\kappa_1, \dots, \kappa_{j-1}$) and \mathbb{H}_j denotes the space of \mathcal{H}_j .⁹

2. Realization Process

With these strategies in place by each player in the game, the state is realized and the course of events is as follows.

- Nature determines the true state t .
- According to the commitment strategies established, the signals are then generated one by one, each of which is observed by all successive intermediaries and the receiver.
- According to the commitment strategies established, the replies of each player for their response stage are implemented, and observed by all successive intermediaries and the receiver.
- After observing the entire history \mathcal{H}_{n+1} , the receiver chooses an action d from $\{A, B\}$.

⁶The sender publicly designs a mapping $T \rightarrow \Delta S$.

⁷Deciding when to respond A and when to respond B

⁸Intermediary 1 publicly designs a mapping $T \times \mathbb{H}_1 \rightarrow \Delta I_1$

⁹Intermediary j publicly designs a mapping $T \times \mathbb{H}_j \rightarrow \Delta I_j$.

In such a realization process, experiments are conducted sequentially and their results are revealed to all successive players, similar to the setting in Li and Norman (2018b) and Felgenhauer and Loerke (2017). For example, in the academic recommendation case, the professor writing the recommendation letter must truthfully transmit the information about coursework and research experience of the student forward through the hierarchy although there may be negative performances in some dimensions. The department admissions committee chair must also truthfully transmit all of the information from the professor, along with his own recommendation to the final decision-maker.

Note that in contrast to non-hierarchical persuasion-game models, the sender in our framework does not have direct control over what the intermediaries and receiver might observe; instead, she tries to influence the receivers' decision by setting up a signal-generating mechanism. This can also be interpreted in our academic setting as a revelation of academic performance measures, research progress, and so on, after the student establishes a study plan. The ex-ante mechanism being used for obtaining the signal, and the realized evidence that emerges are then communicated to the next decision-maker in the chain without noise. A similar process is true for each intermediary in the chain.¹⁰

The game is sequential in nature, in that each player's commitment strategy for the persuasion and response stages is set in sequence, after observing the commitment strategy of the previous player in the hierarchy. Once these strategies are set, the true state is realized, and the realization of the strategies are implemented automatically (with the assistance of a randomization device to implement stochastic elements of strategies).

2.2 Equilibrium Analysis

With the setup of the game established, we now analyze the equilibria of the game. The solution concept is essentially subgame perfect Nash equilibrium, as the game is sequential while no player has private information at the point in the game when strategies are designed.¹¹

We refer to an intermediary as *A-preferred* if he chooses *A* in the *response* stage. We refer to an intermediary as *B-preferred* if he chooses *B* in the *response* stage. We first characterize a sufficient condition that players' optimal strategy in *persuasion* stage is indeed maximizing the likelihood of his preferred action being chosen by the receiver.

¹⁰In the case of an intermediary, such as the faculty letter writer, an ex-ante stochastic policy is set and made known to all players. Based on the signal received from the student and the true state of the world, the letter writer's preset stochastic policy is implemented. In a real world setting, such a stochastic policy could involve a strong to very strong, or mediocre to moderately strong letter. However in our simplified setting, the only actions available to the faculty member are not recommend (*A*) or recommend (*B*), which means that in the current setting, a non-degenerate stochastic policy will have each of *not recommend* or *recommend* being conveyed to the next intermediary with positive probability.

¹¹This follows Kamenica and Gentzkow (2011), where note that once nature determines the true state, the state is private information to the sender. However at this point in the game, the strategies are already set by all players.

Assumption 2.1. *An intermediary j 's reputation gain has a lower bound,*

$$R_j > \max \left(-u_j(B, \alpha), u_j(B, \beta) \right) \quad (1)$$

We assume that any intermediary's reputation gain is greater than their own utility gain (or loss) in either true state, should the eventual decision be B . It is reasonable to think that intermediaries care more about their reputation compared with the sender and the receiver who in our framework only care about the final decision. Please refer to Appendix B for discussion of the role of the reputation term.

We use backward induction to characterize the optimal Bayesian persuasion. By Kamenica and Gentzkow (2011)^[4], the following two lemmas hold with regard to the final intermediary n . In all subsequent sections, we use \hat{p} to denote the incoming belief resulting from Bayesian updating upon history \mathcal{H} .

Lemma 2.2. *For B-preferred intermediary n with incoming belief \hat{p} , assume that the threshold belief of the receiver is \tilde{p}_R . Then the optimal strategy in persuasion stage is described as: (1) no disclosure when $\hat{p} \geq \tilde{p}_R$; (2) otherwise partial disclosure that induces posterior \tilde{p}_R with probability $\frac{\hat{p}}{\tilde{p}_R}$ and 0 with probability $1 - \frac{\hat{p}}{\tilde{p}_R}$.*

Lemma 2.3. *For A-preferred intermediary n with incoming belief \hat{p} , assume that the threshold belief of the receiver is \tilde{p}_R . Then the optimal strategy in persuasion stage is described as: (1) no disclosure when $\hat{p} \leq \tilde{p}_R$; (2) otherwise partial disclosure that induces posterior \tilde{p}_R with probability $\frac{1-\hat{p}}{1-\tilde{p}_R}$ and 1 with probability $\frac{\hat{p}-\tilde{p}_R}{1-\tilde{p}_R}$.*

Next, we characterize intermediary n 's choice of A or B in the *response* stage when the incoming belief is \hat{p} . The expected utility for intermediary n of responding with A and B are calculated as follows,

$$U_n(p|A) = \begin{cases} R_n & p \in [0, \tilde{p}_R] \\ -\frac{p(R_n - u_n(B, \beta))}{1 - \tilde{p}_R} + \frac{R_n - \tilde{p}_R u_n(B, \beta)}{1 - \tilde{p}_R} & p \in [\tilde{p}_R, 1] \end{cases}$$

$$U_n(p|B) = \begin{cases} \frac{p}{\tilde{p}_R}(R_n + u_n(B, \alpha)) + p(u_n(B, \beta) - u_n(B, \alpha)) & p \in [0, \tilde{p}_R] \\ R_n + u_n(B, \alpha) + p(u_n(B, \beta) - u_n(B, \alpha)) & p \in [\tilde{p}_R, 1] \end{cases}$$

Modified Threshold Belief Notice that $U_n(p|A)$ is decreasing in p while $U_n(p|B)$ is strictly increasing in p , and

$$U_n(0|A) > U_n(0|B)$$

$$U_n(1|A) < U_n(1|B)$$

Hence, there exists a threshold value \tilde{p}_{In} , called a **modified** threshold belief, which takes into account the reputation terms based on the strategies in Lemma 2.2 and Lemma 2.3 (the original threshold belief is \tilde{p}_n), such that intermediary n is indifferent between choosing A and B in *response* stage, $U_n(\tilde{p}_{In}|A) = U_n(\tilde{p}_{In}|B)$. Intermediary is considered to be **tougher** (from the perspective of sender) when that intermediary's modified threshold belief is higher, since he requires a higher belief to approve action B that sender strictly prefers. A comparison of modified threshold beliefs and (regular) threshold beliefs is provided in Appendix C.

Lemma 2.4. *Intermediary n responds A when $\hat{p} < \tilde{p}_{In}$, responds B when $\hat{p} > \tilde{p}_{In}$ and is indifferent between A and B when $\hat{p} = \tilde{p}_{In}$.*

Likewise, we assume there exists a modified threshold belief \tilde{p}_{Ij} for intermediary j such that intermediary j responds A when $\hat{p} < \tilde{p}_{Ij}$, responds B when $\hat{p} > \tilde{p}_{Ij}$ and is indifferent between A and B when $\hat{p} = \tilde{p}_{Ij}$. The existence and uniqueness of this strategy are shown later in the Theorem. Taking these three lemmas as benchmark cases in the backward induction process, we show that the following Bayesian persuasions reach their optimality respectively. For ease of notation, we denote the sender as intermediary 0. The following Theorem characterizes one profile of optimal signal structures and the associated sub-game perfect equilibrium, as an induction case. Before that, for any intermediary j we define the maximum modified threshold belief of all successors (intermediaries $j + 1, \dots, n$ and the receiver) as $\tilde{p}_j^{\max} = \max(\{\tilde{p}_{Ik}\}_{k=j+1}^n, \tilde{p}_R)$; the minimum modified threshold belief of all successors (intermediaries $j + 1, \dots, n$ and the receiver) is defined as $\tilde{p}_j^{\min} = \min(\{\tilde{p}_{Ik}\}_{k=j+1}^n, \tilde{p}_R)$. Notice here that the subscript of \tilde{p}_j^{\min} and \tilde{p}_j^{\max} is j rather than Ij for convenience.

In the following equilibrium, only the sender alters the information structure, which we refer to as a *one-step* equilibrium. The induction process is illustrated in [Figure 2](#) and the proof is relegated to Appendix A.

Theorem 2.5 (One-step Equilibrium).

1. *In the persuasion stage,*

- *For B-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \geq \tilde{p}_j^{\max}$; (2) otherwise partial disclosure that induces posterior \tilde{p}_j^{\max} with probability $\frac{\hat{p}}{\tilde{p}_j^{\max}}$ and 0 with probability $1 - \frac{\hat{p}}{\tilde{p}_j^{\max}}$.*
- *For A-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \leq \tilde{p}_j^{\min}$; (2) otherwise partial disclosure that induces posterior \tilde{p}_j^{\min} with probability $\frac{1-\hat{p}}{1-\tilde{p}_j^{\min}}$ and 1 with probability $\frac{\hat{p}-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}}$.*

2. *In response stage,*

- *Intermediary j responds A when $\hat{p} < \tilde{p}_{Ij}$, responds B when $\hat{p} > \tilde{p}_{Ij}$ and is indifferent*

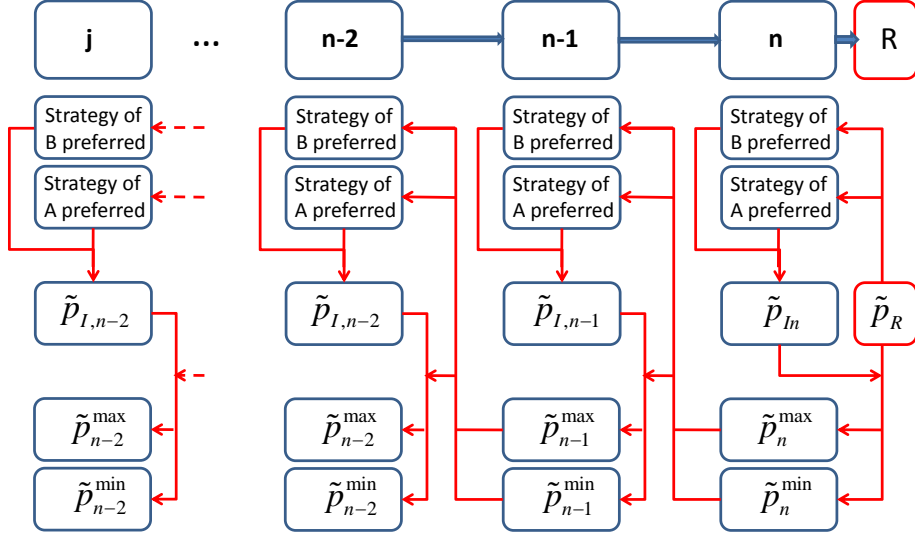


Figure 2: Induction Process of Equilibrium

between A and B when $\hat{p} = \tilde{p}_{I_j}$, where \tilde{p}_{I_j} is defined by

$$U_j(\tilde{p}_{I_j}|A) = U_j(\tilde{p}_{I_j}|B) \quad (2)$$

where

$$U_j(p|A) = \begin{cases} R_j & p \in [0, \tilde{p}_j^{\min}] \\ \frac{1-p}{1-\tilde{p}_j^{\min}} R_j + \frac{p-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}} u_j(B, \beta) & p \in [\tilde{p}_j^{\min}, 1] \end{cases}$$

$$U_j(p|B) = \begin{cases} \frac{p}{\tilde{p}_j^{\max}} (R_j + u_j(B, \alpha)) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [0, \tilde{p}_j^{\max}] \\ R_j + u_j(B, \alpha) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

The above optimal Bayesian persuasion gives us the following equilibrium strategy of the sender, (1) no disclosure when $p_0 \geq \tilde{p}_0^{\max}$; (2) otherwise partial disclosure that induces posterior \tilde{p}_0^{\max} with probability $\frac{p_0}{\tilde{p}_0^{\max}}$ and 0 with probability $1 - \frac{p_0}{\tilde{p}_0^{\max}}$. For convenience we denote the sender as intermediary 0, and \tilde{p}_0^{\max} is defined as the maximum of all modified threshold beliefs of intermediaries and receiver. No intermediary alters the information structure released by the sender. In addition, instead of using $u_j(B, \alpha), u_j(B, \beta), R_j$, the modified threshold beliefs of the players are sufficient for characterization of equilibrium. We call the intermediary with highest modified threshold the **toughest intermediary**. If modified threshold belief of toughest intermediary is higher than \tilde{p}_R ,

we can this intermediary **the toughest player**; otherwise, we call receiver the toughest player. The following corollary summarizes the features of the one-step persuasion equilibrium.

Corollary 2.6. *In the one-step Bayesian persuasion equilibrium,*

1. *Bayesian persuasion is only determined by \tilde{p}_0^{\max} .*
 - *Hierarchical Bayesian persuasion is outcome equivalent to persuading the toughest player (among intermediaries and receiver) **directly**¹²;*
 - *When belief \tilde{p}_0^{\max} is induced, all players prefer action B, and B is chosen;*
 - *When belief 0 is induced, intermediaries and the receiver prefer action A, and A is chosen.*
2. *Each intermediary provides no additional information and merely transmits the message that previous intermediaries do.*

Clearly, the analysis has not characterized all equivalent possible optimal Bayesian persuasions. The literature is mainly concerned about the payoff or welfare benefits of Bayesian persuasion compared with degenerate strategies (such as no disclosure and full disclosure) rather than focusing on fully characterizing all possible equilibria.

In the prosecutor-judge case illustrated in Kamenica and Gentzkow (2011), when the prior probability of guilt is 0.7, it is also optimal to induce posteriors 0.6 and 0.8 with equal probability. In our setting, when the incoming belief \hat{p} is lower than $\tilde{p}_{I,k+1}$, one may choose partial disclosure that induces any posterior $p \in [\tilde{p}_{I,k+1}, \tilde{p}_{k+1}^{\max}]$ with probability $\frac{\hat{p}}{p}$ and 0 with probability $1 - \frac{\hat{p}}{p}$ when $\hat{p} < \tilde{p}_{k+1}^{\max}$.

Nonetheless, there exists another intuitive optimal Bayesian persuasion which we call *myopic* equilibrium. In the one-step equilibrium, only the sender manipulates the information structure and intermediaries provide no additional information, merely transmitting what the previous intermediary does. Compared with this structure, in the *myopic* equilibrium, each player only aims to persuade the immediately subsequent player. In this case, some intermediaries provide a different signal structure, which depends on how difficult it is to persuade the subsequent intermediary.

Theorem 2.7 (Myopic Equilibrium).

1. *In the persuasion stage,*
 - *For B-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \geq \tilde{p}_{I,j+1}$; (2) otherwise partial disclosure that induces posterior $\tilde{p}_{I,j+1}$ with probability $\frac{\hat{p}}{\tilde{p}_{I,j+1}}$ and 0 with probability $1 - \frac{\hat{p}}{\tilde{p}_{I,j+1}}$.*

¹²Then all succeeding players will merely pass the information.

- For A-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \leq \tilde{p}_{I,j+1}$; (2) otherwise partial disclosure that induces posterior $\tilde{p}_{I,j+1}$ with probability $\frac{1-\hat{p}}{1-\tilde{p}_{I,j+1}}$ and 1 with probability $\frac{\hat{p}-\tilde{p}_{I,j+1}}{1-\tilde{p}_{I,j+1}}$.

2. In the response stage,

- Intermediary j responds A when $\hat{p} < \tilde{p}_{Ij}$, responds B when $\hat{p} > \tilde{p}_{Ij}$ and is indifferent between A and B when $\hat{p} = \tilde{p}_{Ij}$, where \tilde{p}_{Ij} is defined by

$$U_j(\tilde{p}_{Ij}|A) = U_j(\tilde{p}_{Ij}|B) \quad (3)$$

The difference in the myopic equilibrium compared to the one-step equilibrium is that succeeding players may provide additional information. However, the two possible posteriors received by receiver are still 0 and \tilde{p}_0^{\max} , identical to the case of the one-step equilibrium. When $\hat{p} > \tilde{p}_{I,k+1}$, the probability that action B is taken is 1 if $\hat{p} \geq \tilde{p}_{k+1}^{\max}$, and $\frac{\hat{p}}{\tilde{p}_{k+1}^{\max}}$ otherwise. The outcome and associated payoffs of the game are unchanged.

The proof is analogous to that for the *one-step* equilibrium. However, in this equilibrium, the information transmission process is quite different. We say that intermediary y is the next *node* after intermediary x , written as $x \rightarrow y$, if she is the nearest subsequent intermediary that has a higher modified threshold belief $\tilde{p}_{Iy} > \tilde{p}_{Ix}$. Mathematically, $x \rightarrow y$ if and only if,

$$y > x, \tilde{p}_{Iy} > \tilde{p}_{Ix} \quad \text{and} \quad \forall x < z < y, \tilde{p}_{Iz} \leq \tilde{p}_{Ix}$$

We say intermediaries c_0, c_1, \dots, c_m form an *increasing chain* if and only if $c_i \rightarrow c_{i+1}, \forall j = 0, \dots, m-1$. Hence, for the n intermediaries and the receiver, we can find the unique *increasing chain* starting from intermediary 1 (as the first *node*). We assume that this unique *increasing chain* is $1 \rightarrow i_1 \rightarrow \dots \rightarrow i_m$.¹³ The above optimal Bayesian persuasion gives us the following equilibrium strategy of players: no disclosure (in other words, simply pass along the message received) if the next intermediary is not in *increasing chain*; otherwise, (1) when incoming belief is $\tilde{p}_{I,i_{j-1}}$, partial disclosure that induces posterior \tilde{p}_{I,i_j} with probability $\frac{\tilde{p}_{I,i_{j-1}}}{\tilde{p}_{I,i_j}}$ and 0 with probability $\frac{\tilde{p}_{I,i_{j-1}}}{\tilde{p}_{I,i_j}}$; (2) when incoming belief is 0, no disclosure. **Figure 3** provides us a graphical illustration.

In a hierarchical setting, the forward-looking requirement needed for the senders in the game is arguably high. The sender, as well as receivers at the bottom of the hierarchy must anticipate the beliefs and actions of players who move much later in the chain of communication. A natural question is whether similar results as in the benchmark case can hold under weaker assumptions on the backward induction reasoning of players. To address this concern, which may be of greater concern in a hierarchical setting than in other persuasion contexts, we also consider a variation of the model in which players are relatively short-sighted. Instead of earning a reputation premium

¹³Intermediary i_m may be the receiver.

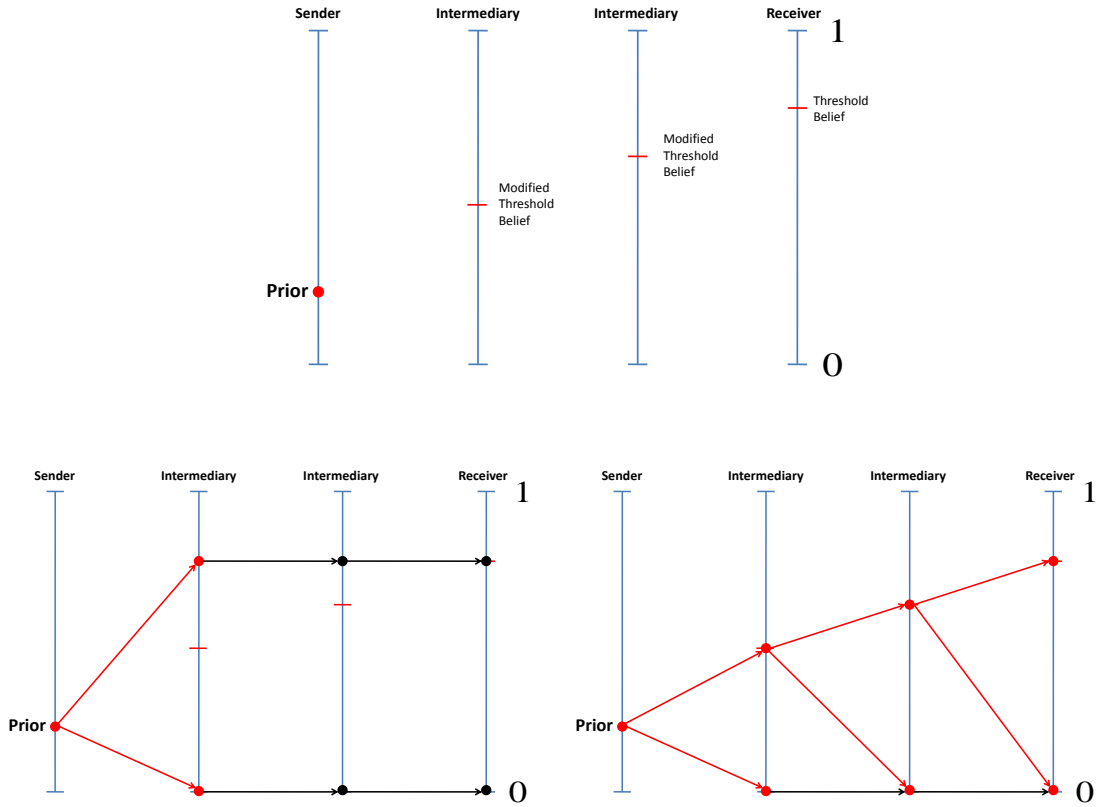


Figure 3: One-step Equilibrium (Left) and Myopic Equilibrium (Right)

and fully using backward induction reasoning in their persuasion strategy, they simply develop a threshold persuasion strategy which maximizes the probability of their preferred outcome. The equilibrium results using such an assumption are qualitatively similar to those in the benchmark case. Appendix F contains further details for interested readers.

2.3 Ordering of Intermediaries

While in the previous sections, we characterized the hierarchical persuasion equilibria, a natural question is what insights we can gain from the model in terms of the ordering of intermediaries. For example, in a situation where a faculty member seeking promotion can construct a desired chain of persuasion among colleagues of higher authority, how should such a persuasion chain be ideally constructed?

To address this question, we first consider the marginal effect of changing \tilde{p}_j^{\min} and \tilde{p}_j^{\max} for some intermediary j . This helps us establish the effect of adding an additional intermediary between intermediary j and intermediary $j + 1$. With a benchmark understanding of the impact of a single intermediary established, we then analyze the optimal order of intermediaries from the perspective of the sender.

Recall that intermediary j 's modified threshold belief is the solution to $U_j(p|A) = U_j(p|B)$ where

$$U_j(p|A) = \begin{cases} R_j & p \in [0, \tilde{p}_j^{\min}] \\ \frac{1-p}{1-\tilde{p}_j^{\min}} R_j + \frac{p-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}} u_j(B, \beta) & p \in [\tilde{p}_j^{\min}, 1] \end{cases}$$


$$U_j(p|B) = \begin{cases} \frac{p}{\tilde{p}_j^{\max}} (R_j + u_j(B, \alpha)) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [0, \tilde{p}_j^{\max}] \\ R_j + u_j(B, \alpha) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

Therefore, \tilde{p}_{Ij} moves in the same direction when \tilde{p}_j^{\min} (or \tilde{p}_j^{\max}) is changing:


- The effect of \tilde{p}_j^{\min} on \tilde{p}_{Ij} : When \tilde{p}_j^{\min} strictly decreases (increases), $U_j(p|A)$ weakly decreases (increases) everywhere and $U_j(p|B)$ remains the same. Therefore, when \tilde{p}_j^{\min} strictly decreases (increases), \tilde{p}_{Ij} weakly decreases (increases).
- The effect of \tilde{p}_j^{\max} on \tilde{p}_{Ij} : When \tilde{p}_j^{\max} strictly increases (decreases), $U_j(p|B)$ weakly decreases (increases) everywhere and $U_j(p|A)$ remains the same. Therefore, when \tilde{p}_j^{\max} strictly increases (decreases), \tilde{p}_{Ij} weakly increases (decreases).

The following four pairs of graphs illustrate the effects discussed above.


(1) \tilde{p}_j^{\min} decreases, then \tilde{p}_{Ij} decreases.




(2) \tilde{p}_j^{\min} decreases, then \tilde{p}_{I_j} remains the same.



(3) \tilde{p}_j^{\max} decreases, then \tilde{p}_{I_j} decreases.



(4) \tilde{p}_j^{\max} decreases, then \tilde{p}_{I_j} remains the same.



2.3.1 Adding an intermediary j' between j and $j + 1$

When we add another intermediary (labeled j') after an intermediary j and in front of intermediary $j + 1$, then we can solve for the modified threshold belief of this new intermediary.

If this new modified threshold belief lies within the range $[\tilde{p}_j^{\min}, \tilde{p}_j^{\max}]$, then nothing will change. If this new modified threshold belief is less than \tilde{p}_j^{\min} , then the modified threshold belief of intermediary j weakly decreases. This decrease may lead \tilde{p}_{j-1}^{\min} to decrease, \tilde{p}_{j-1}^{\max} to decrease, both, or none. Under each of those four circumstances, $\tilde{p}_{I,j-1}$ weakly decreases. Recursively, the modified threshold beliefs of all intermediaries before j weakly decrease.

If this new modified threshold belief is larger than \tilde{p}_j^{\max} , then the modified threshold belief of intermediary j weakly increases. This decrease may have the effect that either \tilde{p}_{j-1}^{\min} increases, \tilde{p}_{j-1}^{\max} increases, both, or none. Under each of those four circumstances, $\tilde{p}_{I,j-1}$ weakly increases. Recursively, the modified threshold beliefs of all intermediaries before j weakly increase.

In the benchmark model, direct communication with the receiver is weakly better than indirect communication with the receiver. However, the generalization of this statement to endogenously chosen intermediaries is not true. More people involved in the persuasion process will not necessarily make the sender worse off. In particular, adding one "nice" intermediary (who is easier to convince) can decrease the \tilde{p}_0^{\max} for the sender, and hence increase the probability of desirable action B .

Corollary 2.8. *Adding an intermediary can decrease \tilde{p}_0^{\max} in equilibrium.*

Example 2.9. *Prior probability $p_0 = \frac{1}{3}$.*

- Receiver has threshold belief $\tilde{p}_R = 0.5$.
- Player J has parameter $u_J(B, \alpha) = -3, u_J(B, \beta) = 1$ and reputation $R_J = 4$.
- Player K has parameter $u_K(B, \alpha) = -1, u_K(B, \beta) = 3$ and reputation $R_K = 4$.

If player J is the unique intermediary, then the modified threshold belief of player J is $\frac{2}{3}$. Hence, the probability that action B is taken is $\frac{1}{2}$ in the optimal Bayesian persuasion.

If player K is added between player J and the receiver, then the modified threshold belief of player K is $\frac{2}{5}$ and the modified threshold belief of player J is reduced to $\frac{5}{9}$. Hence, the probability that action B is taken is $\frac{3}{5}$ in the optimal Bayesian persuasion.

The receiver makes his decision solely based on type-dependent utility. However, the intermediaries care about their reputations. The influence of the reputation term depends on each intermediary's anticipation of the subsequent players. Adding one easy-to-convince intermediary

after the toughest intermediary may decrease \tilde{p}_j^{\min} , and hence decrease utility of the toughest intermediary when responding to the previous player with undesirable action A . By our analysis on the effect of \tilde{p}_j^{\min} on \tilde{p}_{Ij} , the modified threshold belief of the toughest player may decrease.

2.3.2 The Sender's optimal order

With the previous results established, we now consider the scenario where the parameters of the intermediaries are given, and the sender can pre-arrange the order of intermediaries. Which permutation of intermediaries is ideal from the sender's perspective? From the analysis in the benchmark model, we know that the sender seeks to minimize \tilde{p}_0^{\max} , where for ease of notation we again denote the sender as intermediary 0.

We assume that there are n intermediaries, with parameters $\{R_J, u_J(B, \beta), u_J(B, \alpha)\}_{J=1,2,\dots,n}$. We use the capital letter subscript J to represent the labeling of different players, distinguished from small letter subscript j which represents an intermediary's position in the hierarchy.

Then an order is a permutation σ defining a one-to-one mapping from $\{1, 2, \dots, n\}$ to itself. $\sigma(J) = j$ means that intermediary J is located at position j .

For further analysis, for player J , we define the **degree of sender-alignment** using the following formula,

$$\mathbb{K}_J = \left(\frac{R_J + u_J(B, \alpha)}{\tilde{p}_R} + u_J(B, \beta) - u_J(B, \alpha) \right) R_J^{-1} > 1 \quad (4)$$

The higher the value, the more aligned the interests of the intermediary are with the sender.¹⁴ The intermediary with highest degree of sender-alignment is called the **most sender-aligned intermediary**. The intermediary favors action B more if $u_J(B, \alpha)$ or $u_J(B, \beta)$ increases. When $u_J(B, \beta)$ or $u_J(B, \alpha)$ (or both) increases, \mathbb{K}_J will increase and hence he becomes a more sender-aligned player. For further analysis, we define the **inverse degree of sender-alignment** as,

$$\underline{p}_J = \mathbb{K}_J^{-1} = R_J \left(\frac{R_J + u_J(B, \alpha)}{\tilde{p}_R} + u_J(B, \beta) - u_J(B, \alpha) \right)^{-1} \quad (5)$$

which is a probability measure such that $\underline{p}_J \in (0, 1)$. If this probability measure of most sender-aligned intermediary is lower than threshold belief of sender, we call this intermediary the **most sender-aligned player**. Otherwise, we call receiver the most sender-aligned player.

Observation 2.10. • $\underline{p}_J = \tilde{p}_R$ if and only if $\tilde{p}_J = \tilde{p}_R$.

• $\underline{p}_J > \tilde{p}_R$ if and only if $\tilde{p}_J > \tilde{p}_R$

• $\underline{p}_J < \tilde{p}_R$ if and only if $\tilde{p}_J < \tilde{p}_R$

where \tilde{p}_J denotes the threshold belief of player J .

¹⁴ $\mathbb{K}_J > 1$ because $\frac{R_J + u_J(B, \alpha)}{\tilde{p}_R} + u_J(B, \beta) - u_J(B, \alpha) \geq R_J + u_J(B, \alpha) + u_J(B, \beta) - u_J(B, \alpha) = R_J + u_J(B, \beta) > R_J$

Theorem 2.11. *If $\min_J \underline{p}_J \geq \tilde{p}_R$, then all orders are equivalent and optimal for the sender; Otherwise, assume $K = \arg \min_J \underline{p}_J$, then all orders such that player K is in position n are equivalent and optimal for the sender.*

In general, the permutation of intermediaries (in terms of state-dependent utilities and reputation concerns) will influence the outcome. However, if the most sender-aligned player is either the last intermediary or the receiver, then the permutation of the preceding players (besides the sender) will not influence \tilde{p}_0^{\max} . We call player L in proof of [Theorem 2.11](#) the toughest intermediary, and we further call him toughest player if his modified threshold belief exceeds \tilde{p}_R , following our definition of **toughest** in the benchmark model.

Corollary 2.12. *\tilde{p}_0^{\max} weakly increases when removing all intermediaries besides the toughest intermediary.*

[Corollary 2.12](#) delivers a similar message as [Corollary 2.8](#). The sender may be worse off when removing intermediaries and better off when adding intermediaries. The existence of a sender-aligned intermediary in proximity to receiver may soften the toughness of other intermediaries in proximity to the sender. For example, returning to [Example 2.9](#), assume that the original persuasion hierarchy is sender, player J , player K , then the receiver. The toughest intermediary is player J . If player K is removed, \tilde{p}_0^{\max} will increase from $\frac{5}{9}$ to $\frac{2}{3}$.

2.3.3 Comparative Statics

In this section we conduct comparative statics on the ordering of the intermediaries. To do so, we first consider the case that all intermediaries share the same threshold belief, $\frac{-u(B,\alpha)}{u(B,\beta)-u(B,\alpha)}$, then consider the cases that intermediaries differ in their threshold beliefs.

For the case where all intermediaries share the same threshold belief, while the reputation terms may be different across intermediaries, the results are summarized in the following corollary.

Corollary 2.13. *When all intermediaries share the same threshold belief \tilde{p}_I , if*

- $\tilde{p}_I = \tilde{p}_R$, then $\tilde{p}_0^{\min} = \tilde{p}_0^{\max} = \tilde{p}_R$, irrespective of the ordering.
- $\tilde{p}_I > \tilde{p}_R$, then intermediary with the largest reputation term is the most sender-aligned intermediary and the receiver is the most sender-aligned player.
- $\tilde{p}_I < \tilde{p}_R$, then intermediary with the smallest reputation term is the most sender-aligned intermediary and the most sender-aligned player.

If intermediaries and the receiver all share the same threshold belief, then degree of sender-alignment of all players J , \mathbb{K}_J , equals $1/\tilde{p}_R$, which is unrelated to the reputation term. If $\tilde{p}_I > \tilde{p}_R$, then degree of sender-alignment of all players J can be represented as $(R_J/\tilde{p}_R -$

positive number) R_J^{-1} , which increases as R_J increasing. If $\tilde{p}_I < \tilde{p}_R$, then degree of sender-alignment of all player J can be represented as $(R_J/\tilde{p}_R + \text{positive number})R_J^{-1}$, which increases as R_J increasing.

For the case that intermediaries differ in their threshold beliefs, without loss of generality, we assume that the reputation terms for all players are identical. This is reasonable for the purpose of addressing our question, because scaling will not change the behavior of intermediaries.

Corollary 2.14. *When all intermediaries share the same reputation term R_J , if*

- *All intermediaries share the same $u_J(B, \beta)$, then intermediary with the largest $u_J(B, \alpha)$ is the most sender-aligned intermediary.*
- *All intermediaries share the same $u_J(B, \alpha)$, then intermediary with the largest $u_J(B, \beta)$ is the most sender-aligned intermediary.*

We identify the most sender-aligned intermediary in the above two corollaries. The optimal ordering for the sender is to simply put the most sender-aligned intermediary in front of the receiver (as long as the receiver is not the most sender-aligned player).

Recall that from the benchmark model results, from the perspective of the sender, direct persuasion of the receiver is at least weakly better than indirect persuasion. However, by considering the ordering of intermediaries in this section, we learn that if persuasion must be indirect, perhaps surprisingly, the sender can become better off by adding another intermediary to the hierarchy. Specifically, if a potential intermediary has naturally aligned interests with the sender, the sender finds it beneficial to add him to the persuasion chain. Furthermore, to the sender, the ideal position of such an intermediary in the hierarchy is for him to engage in direct persuasion with the receiver.

2.3.4 Numerical Examples

The previous analysis provides us with the answer regarding how small of \tilde{p}_0^{\min} and \tilde{p}_0^{\max} the sender can obtain when having manipulation power on ordering of intermediaries. The following numerical examples will show us that the modified threshold beliefs of players are related to the ordering, although \tilde{p}_0^{\max} is unrelated to the ordering under these two examples.

Example 2.15 (Modified Threshold Beliefs Change I). *The hierarchy is comprised of three intermediaries and one receiver. The threshold belief for the receiver is 0.5. All three intermediaries share the same state-dependent utility $u_I(B, \beta) = 2$ and $u_I(B, \alpha) = -1$. The reputation term for the three players **A**, **B**, **C** are 4, 5, 6 respectively. The following table summarizes the modified threshold beliefs of the three intermediaries under each of the six different orderings.*

- *Player **A** is the most sender-aligned player and his modified threshold belief will not change.*
- *The receiver is the toughest player.*

- *Modified threshold belief of both **B** and **C** change.*

<i>Modified Threshold Belief</i>	A	B	C
<i>S, A, B, C, R</i>	0.4444	0.4545	0.4615
<i>S, A, C, B, R</i>	0.4444	0.4545	0.4590
<i>S, B, A, C, R</i>	0.4444	0.4512	0.4615
<i>S, B, C, A, R</i>	0.4444	0.4512	0.4554
<i>S, C, A, B, R</i>	0.4444	0.4545	0.4554
<i>S, C, B, A, R</i>	0.4444	0.4512	0.4554

The previous example shows that the modified threshold belief of the most sender-aligned player is invariant under all orderings. The following example with different utility values for the intermediaries shows the modified threshold belief of toughest player is invariant under all orderings.

Example 2.16 (Modified Threshold Beliefs Change II). *The hierarchy is comprised of three intermediaries and one receiver. The threshold belief for receiver is 0.5. All three intermediaries share the same state-dependent utility $u_I(B, \beta) = 1$ and $u_I(B, \alpha) = -2$. The reputation term for the three players **A**, **B**, **C** are 4, 5, 6 respectively. The following table summarizes the modified threshold beliefs of the three intermediaries under each of the six different orderings. We can observe that*

- *The receiver is the most sender-aligned player.*
- ***A** is the toughest player and his modified threshold belief will not change.*
- *Modified threshold belief of both **B** and **C** change.*

<i>Modified Threshold Belief</i>	A	B	C
<i>S, A, B, C, R</i>	0.5556	0.5455	0.5385
<i>S, A, C, B, R</i>	0.5556	0.5455	0.5410
<i>S, B, A, C, R</i>	0.5556	0.5488	0.5385
<i>S, B, C, A, R</i>	0.5556	0.5488	0.5446
<i>S, C, A, B, R</i>	0.5556	0.5455	0.5446
<i>S, C, B, A, R</i>	0.5556	0.5488	0.5446

3 Private Information

We now consider the scenario that each player in the chain of persuasion has private information, represented in the model by their types. Different types of players (intermediaries and receiver) may have different values of $u(B, \alpha)$, $u(B, \beta)$, and different types of intermediaries may also have different values of the reputation gain term. For simplicity the private information represented by a single random variable, can be considered as being directly over each intermediary's

modified threshold belief and the receiver's threshold belief. Later in this section, we show this simplification is reasonable.

We use Θ_R to denote the type space for the receiver and Θ_j to denote the type space for intermediary j . We assume that $|\Theta_R| \leq \infty$ and $|\Theta_j| \leq \infty$ for all j . The cumulative distribution function of prior distribution of the receiver's type is common knowledge and is denoted as $F_R : \Theta_R \rightarrow [0, 1]$. The cumulative distribution function of prior distribution of intermediary j 's type is common knowledge and is denoted as $F_j : \Theta_j \rightarrow [0, 1]$. These cumulative distribution functions are able to directly represent the modified threshold beliefs in place of the triple $(R_j, u_j(B, \alpha), u_j(B, \beta))$.

Each player's strategy is affected by the incoming belief from the previous intermediary as well as the type distributions of subsequent players, but not the realization of the types of players preceding him in the path of persuasion. Let θ_R be a representative element of Θ_R and let θ_j be a representative element of Θ_j . Compared to the benchmark model, the state-dependent utilities can be represented by conditional utility expressions

- The utility function of the receiver depends on private information type θ_R , $u^R(d, t|\theta_R)$. Then threshold belief $\tilde{p}_R(\theta_R) = \frac{-u^R(B, \alpha|\theta_R)}{u^R(B, \beta|\theta_R) - u^R(B, \alpha|\theta_R)}$ is a sufficient statistic to characterize the behavior of the receiver.
- The utility function of intermediary j depends on private information type θ_j , $u_j(d, t|\theta_j)$.
- Furthermore, the reputation term of intermediary j is now a random parameter depends on private information type θ_j , $R_j(\theta_j)$.

Therefore, our original assumptions are adjusted to accommodate the private information setting as follows

Assumption 3.1. *The minimum reputation gain has a lower bound, for all $\theta_j \in \Theta_j$*

$$R_j(\theta_j) > \max \left(-u_j(B, \alpha|\theta_j), u_j(B, \beta|\theta_j) \right) \quad (6)$$

Nonetheless, the randomness of $R_j(\theta_j)$ can be incorporated into $\tilde{p}_{I_j}(\theta_j)$ because $R_j(\theta_j)$ is only used to determine $\tilde{p}_{I_j}(\theta_j)$ when analyzing behavior in the response stage of intermediary j . When analyzing the behavior for the persuasion stage of intermediary $j - 1$, the random component is already captured by the modified threshold beliefs $\{\tilde{p}_{I_j}(\theta_j)\}_{\theta_j \in \Theta_j}$.

We denote a distribution of posteriors by $\tau \in \Delta(\Delta(T))$. Therefore, the optimal strategy for a player is a function $\tau(p) : \Delta(T) \rightarrow \Delta(\Delta(T))$. The optimal strategy is a function which generates a distribution of posteriors based on one's own belief about the state, ie. a function from a posterior $\Delta(T)$ to distribution of posteriors $\Delta(\Delta(T))$. Keeping with previous terminology, in the private information case, no disclosure refers to the situation that the message of the intermediary in question is fully uninformative. When $|T| = 2$, the degree of freedom (number of parameters

of the system that may vary independently) of the posterior distribution is 1. Here for a binary distribution, when probability of one type is given as p , then the probability of the other type is $1 - p$. In other words, we can use one parameter to denote the distribution. Therefore, the distribution of posteriors is denoted by $\tau \in \Delta([0, 1])$ and $\tau(p) : [0, 1] \rightarrow \Delta([0, 1])$. Hence the concavification process works on a one dimensional function $V(p)$. The following proposition shows that, for such a one dimensional function, the result of concavification must take a special form, which guarantees the computational tractability of further analysis on private information.

Proposition 3.2. *In the case of private information, the optimal strategy is always equivalent to a partition over $[0, 1]$ by a series of cutoff values $0 = \gamma_1 < \dots < \gamma_l = 1$, each interval is associated with a strategy taken from $\{\mathbb{C}, \mathbb{D}\}$, where \mathbb{C} represents full concealment, and \mathbb{D} represents partial disclosure that generates two posteriors. When prior p is located in the range $[\gamma_k, \gamma_{k+1}]$ with an associated strategy \mathbb{C} , the optimal strategy is no disclosure. When prior p is located in the range $[\gamma_k, \gamma_{k+1}]$ with an associated strategy \mathbb{D} , the optimal strategy induces two posteriors, γ_k, γ_{k+1} .*

The proposition states that under private information, the result of optimal Bayesian persuasion must take this special form. Cutoff values are determined by concavification. In the proof provided in the Appendix, we will provide an algorithm to generate all cutoff values. The full disclosure policy is equivalent to $\gamma_1 = 0, \gamma_2 = 1$ and \mathbb{D} when $p \in [0, 1]$. The full concealment policy is equivalent to $\gamma_1 = 0, \gamma_2 = 1$ and \mathbb{C} when $p \in [0, 1]$. The optimal policy in the prosecution example in Kamenica and Gentzkow (2011) is equivalent to the following,

$$\begin{cases} \mathbb{D} & p \in [0, 0.5] \\ \mathbb{C} & p \in [0.5, 1] \end{cases}$$

Note that this proposition hold not only in hierarchical Bayesian persuasion case, but also in other situations as long as the state is binary.

3.1 Persuading the Receiver

In the previous section, we showed that for all players in the hierarchical Bayesian persuasion game, their optimal strategies take the specific form described in [Proposition 3.2](#). Thus, we can continue to solve the equilibrium by backward induction without solving for the closed form expression of value function V as well as its concave closure \hat{V} for specific players. To examine the persuasion of the receiver, we first consider a simple model without intermediaries as illustrated in Kamenica and Gentzkow (2011). The state-dependent payoff of the receiver is random and may take on several values. For each realization of payoffs, we can compute the threshold belief of such realizations. Those cutoff beliefs are then sufficient for the sender. The sender takes a similar strategy as in Kamenica and Gentzkow (2011), except that he needs to decide what posterior belief he wants to generate, which is 0.5 in the Kamenica and Gentzkow (2011) prosecutor example.

We demonstrate that such a posterior belief must be one of the possible values of the threshold belief of the receiver. Mathematically, all γ values (except the two endpoints 0,1) introduced in [Proposition 3.2](#) must take the threshold belief for some type of receiver. However, not all values of threshold beliefs are taken as γ values. [Example 3.3](#) below provides insights. Compared with the benchmark model, under private information, for some posteriors, some types of receivers may reject the proposal.

Formally, if a B-preferred sender induces belief p in the receiver, then the probability that the receiver finally chooses B can be expressed as the probability that the receiver has lower threshold belief than p ,

$$V_{BR}(p) = \Pr(d = B|p) = \mathbb{E}_{\theta_R}[I(\tilde{p}_R(\theta_R) \geq p)|\theta_R] = F_R(p) \quad (7)$$

where $F_R(p)$ denotes the cumulative distribution function of the threshold belief of the receiver. $V_{BR}(p)$ has the following properties, (1) weakly increasing, (2) right-continuous, (3) piecewise constant (if $|\Theta_R| < \infty$).

For a better understanding, we use the following cumulative distribution function as a discrete example to illustrate the construction of \hat{V}_{BR} and \hat{V}_{AR} as well as the optimal strategy represented by [Proposition 3.2](#).

Example 3.3. ¹⁵ *Assume there are four different types of receivers,*

1. *20 percent of the receivers have threshold belief 0.24;*
2. *40 percent of the receivers have threshold belief 0.3;*
3. *24 percent of the receivers have threshold belief 0.6;*
4. *16 percent of the receivers have threshold belief 0.96.*

We can derive the $V_{BR}(p)$ and $V_{AR}(p)$ for the receiver from the cumulative distribution function,

$$V_{BR}(p) = F_R(p) = \begin{cases} 0 & p \in [0, 0.24) \\ 0.2 & p \in [0.24, 0.3) \\ 0.6 & p \in [0.3, 0.6) \\ 0.84 & p \in [0.6, 0.96) \\ 1 & p \in [0.96, 1] \end{cases} \quad V_{AR}(p) = 1 - F_R(p) = \begin{cases} 1 & p \in [0, 0.24) \\ 0.8 & p \in [0.24, 0.3) \\ 0.4 & p \in [0.3, 0.6) \\ 0.16 & p \in [0.6, 0.96) \\ 0 & p \in [0.96, 1] \end{cases}$$

¹⁵In this example, the concave closure is a continuous function while the original function is right-continuous only.

The concave closures of $V_{BR}(p)$ and $V_{AR}(p)$ are solved as follow

$$\hat{V}_{BR}(p) = \begin{cases} 2p & p \in [0, 0.3] \\ \frac{4p+1.8}{5} & p \in [0.3, 0.6] \\ \frac{4p+5.16}{9} & p \in [0.6, 0.96] \\ 1 & p \in [0.96, 1] \end{cases} \quad \hat{V}_{AR}(p) = \begin{cases} 1 & p \in [0, 0.24] \\ \frac{7.68-7p}{6} & p \in [0.24, 0.96] \\ 4-4p & p \in [0.96, 1] \end{cases}$$

We can compute the concave closure $\hat{V}_{BR}(p)$ for B-preferred player and $\hat{V}_{AR}(p)$ for A-preferred player, which are shown in the left and the right brackets respectively.




Figure 4: B-preferred (Left) and A-preferred (Right)

In the left picture, for a B-preferred sender, the $V_{BR}(p)$ is represented by the orange solid line while $\hat{V}_{BR}(p)$ is represented by the red dashed line. The optimal strategy written in form of Proposition 3.2 is

$$\begin{cases} \mathbb{D} & p \in [0, 0.3] \\ \mathbb{D} & p \in [0.3, 0.6] \\ \mathbb{D} & p \in [0.6, 0.96] \\ \mathbb{C} & p \in [0.96, 1] \end{cases}$$

cutoff values are 0, 0.3, 0.6, 0.96 and 1. Indeed these D 's represent different disclosure strategies, inducing different posteriors.

In the right picture, for an A-preferred sender, the $V_{AR}(p)$ is represented by the orange solid line while $\hat{V}_{AR}(p)$ is represented by the red dashed line. The optimal strategy written in form of Proposition 3.2 is

$$\begin{cases} \mathbb{C} & p \in [0, 0.24] \\ \mathbb{D} & p \in [0.24, 0.96] \\ \mathbb{D} & p \in [0.96, 1] \end{cases}$$

cutoff values are 0, 0.24, 0.96 and 1.

For the further analysis, we characterize some properties for $\hat{V}_{BR}(p)$ and $\hat{V}_{AR}(p)$,

1. Both $\hat{V}_{BR}(p)$ and $\hat{V}_{AR}(p)$ are concave
2. $\hat{V}_{BR}(1) = \hat{V}_{AR}(0) = 1$
3. $\hat{V}_{BR}(0) = \hat{V}_{AR}(1) = 0$
4. \hat{V}_{BR} is strictly increasing when $\hat{V}_{BR}(p) < 1$ while \hat{V}_{AR} is strictly decreasing when $\hat{V}_{BR}(p) < 1$.

Lemma 3.4. $\hat{V}_{BR}(p) + \hat{V}_{AR}(p) > 1$ for $p \in (0, 1)$ unless $F_R(p) = p$.

This lemma will be useful in section 3.4.1 during discussion of the special case of the receiver with uniformly distributed private information.

3.2 Persuading the Last Intermediary

Just as in the benchmark model, the strategy of an intermediary consists of two parts, the response rule in the response stage, and Bayesian persuasion signals in the persuasion stage. For the last intermediary in the sequence, his or her behavior in the persuasion stage is the same as that of a sender in the model without intermediaries, which is analyzed in the previous subsection. Now, we examine the response stage. We must characterize whether intermediary n responds to intermediary $n - 1$ with A or B , when the induced belief is p , given his type θ_n . The most important question is whether the modified threshold belief defined in the benchmark model still exists. If so, is it the unique one?

We first define the expected utility for intermediary n with the following features

1. Responding A/B to intermediary $n - 1$
2. Intermediary n generates posterior p to the receiver
3. The type of the intermediary n is θ_n
4. The type of the receiver is θ_R

by the expressions

$$\begin{aligned}
U_n(p|A, \theta_n, \theta_R) &= \Pr(d = A)R_n + \Pr(d = B) \left(pu_n(B, \beta|\theta_n) + (1 - p)u_n(B, \alpha|\theta_n) \right) \\
&= (1 - \mathcal{I}(p|\theta_R))R_n + \mathcal{I}(p|\theta_R) \left(pu_n(B, \beta|\theta_n) + (1 - p)u_n(B, \alpha|\theta_n) \right) \\
U_n(p|B, \theta_n, \theta_R) &= \Pr(d = A) \cdot 0 + \Pr(d = B) \left(R_n + pu_n(B, \beta|\theta_n) + (1 - p)u_n(B, \alpha|\theta_n) \right) \\
&= \mathcal{I}(p|\theta_R) \left(R_n + pu_n(B, \beta|\theta_n) + (1 - p)u_n(B, \alpha|\theta_n) \right)
\end{aligned}$$

respectively, where boolean variable $\mathcal{I}(p|\theta_R) = \mathbb{I}(\tilde{p}_R(\theta_R) \leq p) \in \{0, 1\}$ represents whether the receiver will choose action B or not when the intermediary passes posterior p to the receiver. Firstly, it is easy to see that $U_n(p|B, \theta_n, \theta_R)$ is increasing in p because when p increases, both $\mathcal{I}(p|\theta_R)$ and $R_n + pu_n(B, \beta) + (1-p)u_n(B, \alpha)$ are non-negative and increasing. Secondly, we have the following boundary values, which help us prove the existence and uniqueness of the modified threshold belief.

$$\begin{aligned} U_n(0|A, \theta_n, \theta_R) &= R_n \\ U_n(1|A, \theta_n, \theta_R) &= u_n(B, \beta) \\ U_n(0|B, \theta_n, \theta_R) &= 0 \\ U_n(1|B, \theta_n, \theta_R) &= R_n + u_n(B, \beta) \end{aligned}$$

In addition, we have $U_n(p|A, \theta_n, \theta_R) \leq R_n - \mathcal{I}(p|\theta_R)(R_n - u_n(B, \beta)) \leq R_n$ for all p . Therefore, if we take the expectation over θ_R , we have the expected utility for A-preferred and B-preferred intermediaries n given his type is θ_n under two extreme posteriors (0 and 1),

$$\begin{aligned} \mathbb{E}_{\theta_R}[U_n(0|A, \theta_n, \theta_R)] &= R_n \\ \mathbb{E}_{\theta_R}[U_n(1|A, \theta_n, \theta_R)] &= u_n(B, \beta) \\ \mathbb{E}_{\theta_R}[U_n(0|B, \theta_n, \theta_R)] &= 0 \\ \mathbb{E}_{\theta_R}[U_n(1|B, \theta_n, \theta_R)] &= R_n + u_n(B, \beta) \end{aligned}$$

Then by Bayesian persuasion, for an A-preferred player, the best he can achieve is the concave closure of $\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)]$, $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$; for a B-preferred player, the best he can achieve is the concave closure of $\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)]$, $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$. We then need to compare whether response A or B is better.

We claim that $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$ is decreasing in p while $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$ is increasing in p . The latter statement holds because each $U_n(p|B, \theta_n, \theta_R)$ is increasing. As for the former statement, we have shown that $U_n(p|A, \theta_n, \theta_R) \leq R_n$ and hence $\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)] \leq R_n$. Since $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$ is a concave function, if the derivative is positive for some value p , then $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$ is strictly increasing for all probabilities $[0, p]$. However, $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$ cannot exceed $R_n = \mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(0|A, \theta_n, \theta_R)])$. This provides a contradiction.

The following lemma illustrates the *single crossing property*.

Lemma 3.5. $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)]) = \mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$ has a unique solution and we define it as $p = \tilde{p}_{In}(\theta_n)$.

Hence, there exists a threshold value $\tilde{p}_{In}(\theta_n)$ for type θ_n , called the modified threshold belief such that intermediary n is indifferent between choosing A and B in response stage. Here, we can

see that the modified threshold belief $\tilde{p}_{In}(\theta_n)$ is sufficient to represent the behavior of intermediary n with type θ_n , which justifies our definition about the cumulative distribution in Section 3 (paragraph 2).

Now we move to the persuasion stage of intermediary $n - 1$. Take the B-preferred intermediary $n - 1$ as an example, we first compute the probability that the receiver chooses B given that intermediary n has type θ_n when intermediary $n - 1$ passes posterior p to intermediary n . If p does not exceed $\tilde{p}_{In}(\theta_n)$, then intermediary n with type θ_n becomes A-preferred and tries to persuade receiver to maximize the probability of choosing A (which is $\hat{V}_{AR}(p)$), hence the probability of choosing B is $1 - \hat{V}_{AR}(p)$. If p exceeds $\tilde{p}_{In}(\theta_n)$, then intermediary n with type θ_n becomes B-preferred and tries to persuade the receiver to maximize the probability of choosing B (which is $\hat{V}_{BR}(p)$). Therefore, the probability that the receiver chooses B is defined as

$$V_{Bn}(p|\theta_n) = \begin{cases} 1 - \hat{V}_{AR}(p) & p < \tilde{p}_{In}(\theta_n) \\ \hat{V}_{BR}(p) & p \geq \tilde{p}_{In}(\theta_n) \end{cases} \quad (8)$$

where both $\hat{V}_{AR}(p)$ and $\hat{V}_{BR}(p)$ are derived from the previous subsection and are irrelevant to θ_n . Therefore, on behalf of intermediary $n - 1$, when he passes p to intermediary n , the probability that B is chosen is the expectation of $V_{Bn}(p|\theta_n)$ taken over θ_n .

$$V_{Bn}(p) = \int_{\theta_n} V_{Bn}(p|\theta_n) dF_n(\theta_n) \quad (9)$$

Likewise, if intermediary $n - 1$ is A-preferred, $V_{An}(p)$ has the following expression

$$V_{An}(p) = \int_{\theta_n} V_{An}(p|\theta_n) dF_n(\theta_n) \quad (10)$$

where

$$V_{An}(p|\theta_n) = \begin{cases} \hat{V}_{AR}(p) & p \leq \tilde{p}_{In}(\theta_n) \\ 1 - \hat{V}_{BR}(p) & p > \tilde{p}_{In}(\theta_n) \end{cases} \quad (11)$$

Hence, we can calculate $\hat{V}_{Bn}(p) = \mathbf{Co}(V_{Bn}(p))$ and $\hat{V}_{An}(p) = \mathbf{Co}(V_{An}(p))$. **Proposition 3.2** tells us that these calculation processes are straightforward. Hence, we can use \hat{V}_{An} and \hat{V}_{Bn} directly.

Although the calculation process of concave closure is generally difficult, **Proposition 3.2** shows that with a binary state, the optimal strategy of intermediary $n - 1$ is always equivalent to a partition associated with strategy \mathbb{C}, \mathbb{D} and can be computed in polynomial time efficiently, even with an infinite number of types. Then, we can recover the strategies from \hat{V}_{An} and \hat{V}_{Bn} .

3.3 Persuading Intermediary j

We now analyze how to persuade intermediary j for a general j . Mathematically, this means solving \hat{V}_{Bj} and \hat{V}_{Aj} from $\hat{V}_{B,j+1}$ and $\hat{V}_{A,j+1}$ (In section 3.2, we solve \hat{V}_{Bn} and \hat{V}_{An} from \hat{V}_{BR} and

\hat{V}_{AR}). The methodology for proving the single crossing property is exactly the same as previously. For any intermediary j , the optimal disclosure strategy in the persuasion stage is unrelated to his own type. However, the modified threshold belief of intermediary j depends on his own type. Please refer to the Appendix for details.

3.4 Some Special Cases

While the previous subsection derived results for the general case, it is helpful to examine some special cases in order to gain specific insights which cannot be easily observed from the general results. We first discuss the special case that the receiver has a uniformly distributed type. We then separately consider the special case that only the receiver has private information.

3.4.1 Uniformly Distributed Receiver

The previous analysis have focused on understanding when the sender benefits from persuasion and deriving the optimal signal. However, when is persuasion ineffective? The case of a uniformly type-distributed receiver provides insights on this question. The following proposition shows that for a uniformly type-distributed receiver, Bayesian persuasion cannot do any better than degenerate strategies, including full concealment and full disclosure.

Proposition 3.6. *Bayesian persuasion cannot generate higher revenue for all $p \in [0, 1]$ if $F_R(p) = p$.*

This proposition applies [Lemma 3.4](#). This is not the only situation that persuasion is ineffective. As long as $V_{Bj}(p) = p$ and $V_{Aj}(p) = 1 - p$ for some intermediary j , then for all preceding intermediaries, persuasion is ineffective.

3.4.2 Deterministic Intermediaries

Another special case is that only the receiver has private information while all intermediaries' information are public. This case is somewhere between benchmark model and general version of private information model. Hence, when we derive $V_{Bn}(p)$ and $V_{An}(p)$, there will be no uncertainty and the results will be

$$V_{Bn}(p) = \begin{cases} 1 - \hat{V}_{AR}(p) & p < \tilde{p}_{In} \\ \hat{V}_{BR}(p) & p \geq \tilde{p}_{In} \end{cases}$$

$$V_{An}(p) = \begin{cases} \hat{V}_{AR}(p) & p \leq \tilde{p}_{In} \\ 1 - \hat{V}_{BR}(p) & p > \tilde{p}_{In} \end{cases}$$

and the concave closures are

$$\hat{V}_{Bn}(p) = \begin{cases} \frac{p\hat{V}_{BR}(\tilde{p}_{In})}{\tilde{p}_{In}} & p < \tilde{p}_{In} \\ \hat{V}_{BR}(p) & p \geq \tilde{p}_{In} \end{cases}$$

$$\hat{V}_{An}(p) = \begin{cases} \hat{V}_{AR}(p) & p \leq \tilde{p}_{In} \\ \frac{(1-p)\hat{V}_{AR}(\tilde{p}_{In})}{1-\tilde{p}_{In}} & p > \tilde{p}_{In} \end{cases}$$

Applying the same induction process as [Theorem 2.5](#), we can derive the general formula for $\hat{V}_{B,j+1}(p)$ and $\hat{V}_{A,j+1}(p)$. From the previous analysis, we can conclude that the modified threshold belief of each intermediary exists uniquely.

Theorem 3.7. *For B-preferred intermediary j with incoming belief \hat{p} , the probability that B is chosen by the receiver is $\hat{V}_{B,j+1}(\hat{p})$ in equilibrium, and for A-preferred intermediary j with incoming belief \hat{p} , the probability that A is chosen by the receiver is $\hat{V}_{A,j+1}(\hat{p})$ in equilibrium, where*

$$\hat{V}_{B,j+1}(p) = \begin{cases} \frac{p\hat{V}_{BR}(\tilde{p}_j^{\max})}{\tilde{p}_j^{\max}} & p < \tilde{p}_j^{\max} \\ \hat{V}_{BR}(p) & p \geq \tilde{p}_j^{\max} \end{cases}$$

$$\hat{V}_{A,j+1}(p) = \begin{cases} \hat{V}_{AR}(p) & p \leq \tilde{p}_j^{\min} \\ \frac{(1-p)\hat{V}_{AR}(\tilde{p}_j^{\min})}{1-\tilde{p}_j^{\min}} & p > \tilde{p}_j^{\min} \end{cases}$$

where $\hat{V}_{BR} = \mathbf{Co}(F_R)$ and $\hat{V}_{AR} = \mathbf{Co}(1-F_R)$. \tilde{p}_j^{\max} and \tilde{p}_j^{\min} are defined as $\tilde{p}_j^{\max} = \max\left(\{\tilde{p}_{Ik}\}_{k=j+1}^n, 0\right)$ and $\tilde{p}_j^{\min} = \min\left(\{\tilde{p}_{Ik}\}_{k=j+1}^n, 1\right)$.

\tilde{p}_j^{\max} and \tilde{p}_j^{\min} do not include \tilde{p}_R anymore, since the information of \tilde{p}_R is already captured by $\hat{V}_{AR}(p)$ and $\hat{V}_{BR}(p)$. When $j = n$, $\tilde{p}_j^{\max} = 0$ and $\tilde{p}_j^{\min} = 1$.

The equilibrium strategy is same as in the benchmark model. Both the one-step equilibrium and myopic equilibrium, as well as other equilibria situated between these two extreme cases remain as equilibria. However, the differences are (1) when $p > \tilde{p}_j^{\max}$, for B-preferred intermediary j , the probability of decision B is no longer 1; (2) when $p < \tilde{p}_j^{\min}$, for A-preferred intermediary j , the probability of decision A is no longer 1.

4 Endogenous Reputation

The exogenous reputation term in the model up to now can be seen as the discounted future benefit in current terms. However, the nature of this benefit may be based on the relationship implied by successful recommendation. For example, assume that a professor writes a recommendation letter for a student to the admission committee and the student is accepted to the program. Firstly, this long-run relationship may generate benefits for the professor compared to a

situation where the student is rejected from the program. Secondly, when the rejection is observed by other students, fewer students may approach the professor in the future. Similarly, when a professor only commits to a mediocre recommendation and the student is subsequently admitted, the long-run credibility of the professor’s recommendation may also be harmed in the eyes of other students. These effects make the long-run benefit potentially more important for the professor than the admission decision itself. However, up to this point we have treated the reputation gain from effective persuasion as exogenously given.

In this section, we endogenize future benefits through an infinitely repeated game instead of using the previously considered exogenous reputation term. In an infinitely repeated game setting, the discounted future benefit is influenced by subsequent strategies, while the discounted future benefit is used to find the equilibrium strategy as in previous section. The relationship between the persuasion policy and the associated revenue function is shown in **Figure 5**.

The benchmark case merely considers the impact of future payoffs on the optimal policy but not the impact of the optimal policy on future payoffs. Therefore, we apply dynamic programming techniques to handle this endogenous influence. We show that the equilibrium essentially follows the same pattern as before. We first consider the simpler case of only one intermediary in the hierarchy, then analyze the case of multiple intermediaries.




Figure 5: Iteration Process with One Intermediary

4.1 One Intermediary

Consider a repeated game with an infinite horizon. In the first stage, sender, receiver and an intermediary play a hierarchical Bayesian persuasion game drawn from the stage game set X . If receiver took the same action that intermediary responded to the sender in time k , then another hierarchical Bayesian persuasion game is played in time $k + 1$, drawn from same stage game set X .

The sender and receiver may not be the same individuals as in previous stages. If the receiver took the other action, intermediary is not considered credible anymore and receives η in all following stages. In order to make the game Markovian and hence solvable under dynamic programming, we assume that the stage game in $k + 1$ only depends on the stage game in k , but we do not impose a requirement for independence of the stage games. The probability that stage game in $k + 1$ is x_b given that the stage game in k is x_a is defined as π_{ab} , which is irrelevant to k . Then the path of stage games is a discrete time Markov chain.

Hence, in each stage, the state of the intermediary can be characterized by $X \cup \{N\}$ where $x \in X$ represents the stage game that an intermediary faces if involved, while N represents non-involvement due to a previously failed recommendation. We assume that once a recommendation fails, that intermediary will not be invited to participate in further persuasion processes anymore,

$$\Pr(x_{k+1} = N | x_k = N) = 1 \quad (12)$$

Note that this assumption is for simplicity only, and that the results hold for any situation where the probability of being able to participate in future persuasion stage games is less than 1, conditional on a failed recommendation.

The model is thus a dynamic control problem (or a Markovian decision problem) with stationary discrete system $x_{k+1} = f(x_k, u_k, w_k)$, where x_k represents the state variable, u_k represents the Bayesian persuasion strategy and w_k represents a disturbance that depends on x_k, u_k only. In this section, we apply the notation from control theory and dynamic programming. No participation, N , is an absorbing state. The Bayesian persuasion strategy consists of two parts, one in each stage (response and persuasion stages). The disturbance w_k exists because the outcome of Bayesian persuasion is probabilistic.

The utility $g(x, u, w)$, for the intermediary in stage game x with policy u and random variable w , is discounted by factor $\delta \in (0, 1)$ and bounded by M , $\forall x, u, w, |g(x, u, w)| \leq M$.¹⁶ function given that the state is now x under policy u . We use $J^*(x)$ to represent the optimal utility function. Hence, $J^*(N) = \frac{\eta}{1-\delta}$. In the following assumption, $\delta \mathbb{E}_x J^*(x | x \neq N)$ represents the future discounted utility under the optimal policy. We use $u_I(d, t|x)$ to represent the stage utility of the intermediary derived from decision d in state t .

Assumption 4.1. For all $x \in X$, $\max \left(-u_I(B, \alpha|x), u_I(B, \beta|x) \right) < \delta \mathbb{E}_x J^*(x | x \neq N) - \frac{\eta\delta}{1-\delta}$

$\delta \mathbb{E}_x J^*(x | x \neq N) - \frac{\eta\delta}{1-\delta}$ is the incremental revenue from a successful recommendation. Under this assumption, the intermediary maximizes the probability of his preferred action in the persuasion stage. Now, we move to the general policy. For revenue function J , we define mapping T as

$$(TJ)(x) = \max_u \mathbb{E}_w \{g(x, u, w) + \delta J(f(x, u, w))\} \quad (13)$$

¹⁶This notation is applied in control theory and followed from *Dynamic Programming and Optimal Control, Volume II* (Dimitri P. Bertsekas, 2012), <http://athenasc.com/dpbook.html>. We use $J_u(x)$ to represent the utility

which denotes the optimal revenue function for the one-stage problem that has stage revenue g and terminal cost δJ .

Lemma 4.2 (Convergence Theorem and Bellman Equation Theorem). *The optimal revenue function satisfies*

$$J^*(x) = \lim_{N \rightarrow \infty} (T^N J)(x) \quad (14)$$

and

$$J^*(x) = (TJ^*)(x) \quad (15)$$

Furthermore, J^* is the unique solution of this equation.

Proof. The above two theorems can be proved by a contraction mapping fixed-point theorem^[41].¹⁷ □

Although the optimal revenue is pinned down uniquely, the optimal policy is not necessarily unique. We illustrate the optimal control policy u^* under optimal revenue J^* . That is,

$$J^*(x) = \mathbb{E} \left\{ g(x, u^*, w) + \delta J^*(f(x, u^*, w)) \right\} \quad (16)$$

We claim that the optimal policy is a threshold policy with a unique threshold. Given stationary policy u is employed by intermediary after time k , it is required that u maximizes the expected utility given that u is applied in all subsequent stages.

$$u \in \arg \max_v \left(g(x, v, w) + \delta J_u(f(x, v, w)) \right) \quad (17)$$

The negative proposition tells us that if u is not a maximizer of $g(x, u, w) + \delta J_u(f(x, u, w))$, u cannot be an optimal policy. Given any non-threshold policy u , we can compute the expectation of future benefits $\delta J_u(x) - \frac{\eta \delta}{1-\delta}$ and plug it into the reputation term. Then in time k , the optimal policy should be a threshold policy via the same reasoning as the benchmark model. Hence, we conclude the following lemma.

Lemma 4.3. *The optimal policy in response stage is a threshold policy. Mathematically, in equilibrium, there exists a threshold value \tilde{p} depending on state x only such that (1) responds B when incoming belief is larger than \tilde{p} ; (2) responds A when incoming belief is smaller than \tilde{p} .*

Last but not least, we prove the uniqueness of the modified threshold belief for the intermediary. The policy u can be represented by a series of parameters $\tilde{p}_I(u|x)$, which is known as the modified threshold belief under game x . Given J^* and the stage game x , the reputation term can be determined uniquely as in the benchmark case, and hence we can solve the optimal modified threshold belief. We can also prove uniqueness by regarding $J(x)$ as a continuous function of

¹⁷For details, please refer to *Dynamic Programming and Optimal Control, Volume II* (Dimitri P. Bertsekas, 2012), <http://athenasc.com/dpbook.html>

$\{\tilde{p}_I(u|x)\}_{x \in X}$ and using the continuity property. $J(x)$ as well as the optimal threshold belief are uniquely determined, but without a closed form solution in general, since both the game set and policy set are infinite. However, with finite size of X , the value of J^* can be solved iteratively by numerical linear algebra.

Example 4.4 (Endogenized Reputation with $|X| = 1$).

Setup. There is only one possible stage game x . In this stage game, the prior probability is $p_0 = \Pr(\beta) = 0.2$. The discount factor is $\delta = 0.9$. $\eta = -10$ and hence $J^*(N) = -100$. The receiver has threshold belief $\tilde{p}_R = 0.5$, state-dependent utilities for the intermediary are given by $u_I(B, \alpha) = -1$, $u_I(B, \beta) = 10$.

Analysis. Assume that intermediary has threshold \tilde{p}_I in the stage game. We need to calculate the present value $J(\tilde{p}_I)$ if applying this threshold policy in all subsequent games. The sender will generate two possible posteriors, $\max(\tilde{p}_I, \tilde{p}_R)$ with probability $\frac{p_0}{\max(\tilde{p}_I, \tilde{p}_R)}$ and 0 with probability $1 - \frac{p_0}{\max(\tilde{p}_I, \tilde{p}_R)}$. The revenue in the stage game is hence

$$\begin{aligned} \mathbb{E}g(\tilde{p}_I) &= \begin{cases} \frac{p_0}{\tilde{p}_R}(\tilde{p}_R u_I(B, \beta) + (1 - \tilde{p}_R)u_I(B, \alpha)) & \tilde{p}_I < \tilde{p}_R \\ \frac{p_0}{\tilde{p}_I}(\tilde{p}_I u_I(B, \beta) + (1 - \tilde{p}_I)u_I(B, \alpha)) & \tilde{p}_I > \tilde{p}_R \end{cases} \\ &= \begin{cases} 1.8 & \tilde{p}_I < 0.5 \\ 2.2 - \frac{0.2}{\tilde{p}_I} & \tilde{p}_I > 0.5 \end{cases} \end{aligned}$$

The present value of future benefit is hence $J(\tilde{p}_I) = \frac{\mathbb{E}g(\tilde{p}_I)}{1-\delta}$ and $\delta J(\tilde{p}_I) = \frac{\delta \mathbb{E}g(\tilde{p}_I)}{1-\delta}$ for the time horizon starting from the next period. Given this revenue function, we iteratively solve for the optimal threshold policy. The payoff of the intermediary with response A and B are calculated separately,

$$\begin{aligned} U_I(p|A) &= \begin{cases} \delta J(\tilde{p}_I) & p \in [0, \tilde{p}_R] \\ \frac{1-p}{1-\tilde{p}_R} \delta J(\tilde{p}_I) + \frac{p-\tilde{p}_R}{1-\tilde{p}_R} (u_I(B, \beta) + \delta J^*(N)) & p \in [\tilde{p}_R, 1] \end{cases} \\ U_I(p|B) &= \begin{cases} \frac{p}{\tilde{p}_R} (\delta J(\tilde{p}_I) + \tilde{p}_R u_I(B, \beta) + (1 - \tilde{p}_R)u_I(B, \alpha)) + (1 - \frac{p}{\tilde{p}_R}) (\delta J^*(N)) & p \in [0, \tilde{p}_R] \\ \delta J(\tilde{p}_I) + u_I(B, \alpha) + p(u_I(B, \beta) - u_I(B, \alpha)) & p \in [\tilde{p}_R, 1] \end{cases} \end{aligned}$$

Fixed Point. Since $U_I(\tilde{p}_R|A) < U_I(\tilde{p}_R|B)$ because $\tilde{p}_R > \frac{-u_I(B, \alpha)}{u_I(B, \beta) - u_I(B, \alpha)}$. We can take $J(\tilde{p}_I) =$

$\frac{\mathbb{E}g(\tilde{p}_I)}{1-\delta} = \frac{1.8}{1-0.1} = 18$ into those formula, and obtain

$$U_I(p|A) = \begin{cases} 16.2 & p \in [0, 0.5] \\ 112.4 - 192.4p & p \in [0.5, 1] \end{cases}$$

$$U_I(p|B) = \begin{cases} 221.4p - 90 & p \in [0, 0.5] \\ 15.2 + 11p & p \in [0.5, 1] \end{cases}$$

We can then solve for the modified threshold belief $\tilde{p}_I = 0.4797$.

Iteration. If we begin with any threshold $\tilde{p}_I^{(0)} < 0.5$ as the initial value of iterations, we can solve for the above modified threshold belief in one iteration, that is $\tilde{p}_I^{(1)} = \tilde{p}_I$. If we begin with any threshold $\tilde{p}_I^{(0)} > 0.5$ as the initial value, then

$$U_I(p|A) = \begin{cases} 19.8 - \frac{1.8}{\tilde{p}_I^{(0)}} & p \in [0, 0.5] \\ (119.6 - \frac{3.6}{\tilde{p}_I^{(0)}}) - (199.6 - \frac{3.6}{\tilde{p}_I^{(0)}})p & p \in [0.5, 1] \end{cases}$$

$$U_I(p|B) = \begin{cases} (228.6 - \frac{3.6}{\tilde{p}_I^{(0)}})p - 180 & p \in [0, 0.5] \\ 18.8 - \frac{1.8}{\tilde{p}_I^{(0)}} + 11p & p \in [0.5, 1] \end{cases}$$

which means that $\tilde{p}_I^{(1)} < 0.5$ and hence $\tilde{p}_I^{(2)} = \tilde{p}_I$.

Penalty. If the size of the penalty is reduced from 100 to 10, then we have

$$U_I(p|A) = \begin{cases} 16.2 & p \in [0, 0.5] \\ 31.4 - 30.4p & p \in [0.5, 1] \end{cases}$$

$$U_I(p|B) = \begin{cases} 59.4p - 9 & p \in [0, 0.5] \\ 15.2 + 11p & p \in [0.5, 1] \end{cases}$$

The modified threshold belief decreases to $\tilde{p}_I = 0.4242$.

4.2 Multiple Intermediaries

We now incorporate multiple intermediaries into the repeated game setting. Following the same basic setup, in the first stage, sender, receiver, as well as all intermediaries play a hierarchical Bayesian persuasion game drawn from the stage game set X . If receiver took the same action that intermediary j responded to intermediary $j - 1$ with at time k , then the player located at j th position will be included in another hierarchical Bayesian persuasion game at time $k + 1$, and his or her location in that game will be position j as well. This is merely for simplification, and we

later show that having the same position in the hierarchy as previously is not required. If the receiver took the other action, an intermediary is no longer considered credible and receives η in all subsequent stages. Similarly to the one intermediary case, we assume that stage game set does not change in each period. The evolution of the stage game is Markovian, without requiring independence.

When considering multiple intermediaries, we introduce a new parameter, the hierarchical position j . We use $\mathbf{u} = \{\mathbf{u}_j\}_{j=1,\dots,n}$ that consists of strategies for all n intermediaries, to represent the strategy profile in time k , where $\mathbf{u}_j = \{u_j(x)\}_{x \in X}$ represents the strategy of the intermediary with position j in different stage games. We use $J_u(x|j)$ to represent the revenue function for intermediary j given that

1. The stage game is now x ;
2. In all coming stage games, players applied policy \mathbf{u} , in other words, intermediary i applies $\mathbf{u}_i(x')$ if the stage game changes to x' ;
3. In this stage game, strategies of intermediaries $j + 1, \dots, n$ are solved backward under point 2.

Let us use $J^*(x|j)$ to represent the optimal revenue function for intermediary j under optimal policy profiles \mathbf{u}^* . The optimality of \mathbf{u}^* says the following: if in all upcoming stage games, players applied policy \mathbf{u}^* , then in this stage game, strategies of all intermediaries $1, \dots, n$ are \mathbf{u}^* by backward induction. For absorbing state N , $J^*(N|j) = \frac{c}{1-\delta}$ for all j . Using the convergence theorem and Bellman equation theorem, we obtain the following two insights:

(1) The optimal revenue function has a fixed point under the mapping T , $T\mathbf{J}^*(x) = \mathbf{J}^*(x)$. This method is similar to the previous section but we can now decompose $T = T_1 T_2 \dots T_n$ where T_j denotes the mapping that solves for the optimal strategy of intermediary j while keeping strategies of others unchanged, such that (a) the discounted future benefit is fixed at δJ , (b) all succeeding intermediaries apply subgame perfect equilibrium, given the discounted future benefit δJ .

(2) The optimal revenues can be attained by iteration if the discount factor is strictly less than 1 and the stage outcome is bounded. Let u_j^m denote the optimal strategy of intermediary j in iteration m . For some specific stage game x , the iteration process of solving $\mathbf{u}^{m+1} = \{u_1^{m+1}, \dots, u_n^{m+1}\}$ are described as follows:

$$\begin{aligned}
T_n : u_n^{m+1}(x) &\in \arg \max_{u_n} \mathbb{E}_w \{g(x, \mathbf{u}(\mathbf{n}, \mathbf{m}), w|n) + \delta J_{\mathbf{u}^m}(f(x, \mathbf{u}(\mathbf{n}, \mathbf{m}), w)|n)\} \\
T_{n-1} : u_{n-1}^{m+1}(x) &\in \arg \max_{u_{n-1}} \mathbb{E}_w \{g(x, \mathbf{u}(\mathbf{n}-\mathbf{1}, \mathbf{m}), w|n-1) + \delta J_{\mathbf{u}^m}(f(x, \mathbf{u}(\mathbf{n}-\mathbf{1}, \mathbf{m}), w)|n-1)\} \\
&\dots \\
T_2 : u_2^{m+1}(x) &\in \arg \max_{u_2} \mathbb{E}_w \{g(x, \mathbf{u}(\mathbf{2}, \mathbf{m}), w|2) + \delta J_{\mathbf{u}^m}(f(x, \mathbf{u}(\mathbf{2}, \mathbf{m}), w)|2)\} \\
T_1 : u_1^{m+1}(x) &\in \arg \max_{u_1} \mathbb{E}_w \{g(x, \mathbf{u}(\mathbf{1}, \mathbf{m}), w|1) + \delta J_{\mathbf{u}^m}(f(x, \mathbf{u}(\mathbf{1}, \mathbf{m}), w)|1)\}
\end{aligned}$$

where $\mathbf{u}(j, m)$ denotes the strategy that intermediaries with a label larger than j take in iteration $m + 1$ while intermediaries with a label less than or equal to j take the following strategy in iteration m ,

$$\begin{aligned}
\mathbf{u}(n, m) &= (\mathbf{u}_1^m, \mathbf{u}_2^m, \dots, \mathbf{u}_{n-1}^m, \mathbf{u}_n^m) \\
\mathbf{u}(n-1, m) &= (\mathbf{u}_1^m, \mathbf{u}_2^m, \dots, \mathbf{u}_{n-1}^m, \mathbf{u}_n^{m+1}) \\
&\dots \\
\mathbf{u}(2, m) &= (\mathbf{u}_1^m, \mathbf{u}_2^m, \dots, \mathbf{u}_{n-1}^{m+1}, \mathbf{u}_n^{m+1}) \\
\mathbf{u}(1, m) &= (\mathbf{u}_1^m, \mathbf{u}_2^{m+1}, \dots, \mathbf{u}_{n-1}^{m+1}, \mathbf{u}_n^{m+1})
\end{aligned}$$

and

$$\mathbf{u}_j^m = \{u_j^m(x)\}_{x \in X}$$

Let us further examine the subscript $\mathbf{u}(n, m)$ at the term δJ that represents discounted future revenue. When computing the iteration $m + 1$, we use the optimal revenue computed at iteration m for all n backward induction processes in T_1 to T_n . However, at iteration $m + 1$, each time we solve for the new strategy of intermediary j , we assume that all subsequent players apply the new strategy in the current stage game in the spirit of backward induction. We show the iteration process graphically in [Figure 6](#).

In the one intermediary model, we showed that under some sufficient conditions, the optimal strategy in the persuasion stage maximizes the likelihood of the preferred action. Additionally, the optimal policy in the response stage is a threshold policy. When incorporating multiple intermediaries, we need to modify the assumptions to make it optimal to maximize the preferred action in the persuasion stage. Previously, we showed the multi-intermediary version of the convergence theorem and Bellman equation theorem. Therefore, with multiple intermediaries, the equilibrium strategies follow the same pattern as the benchmark case and the case of endogenous reputation with one intermediary.

When the same position condition is violated, the agent in the j th position who recommends successfully joins the next game and changes to a random position, we use $\delta \mathbb{E}_j J_{\mathbf{u}(n, m)}(f(x, \mathbf{u}(\cdot, m), w)|j)$ to replace $\delta J_{\mathbf{u}(n, m)}(f(x, \mathbf{u}(\cdot, m), w)|j)$. The previously mentioned results hold. Consider again the example of a professor writing a recommendation letter: When a professor with position j successfully recommends a student, when writing a letter for a different student, he may not have the same position in the admissions hierarchy as in the application process of the previous student.

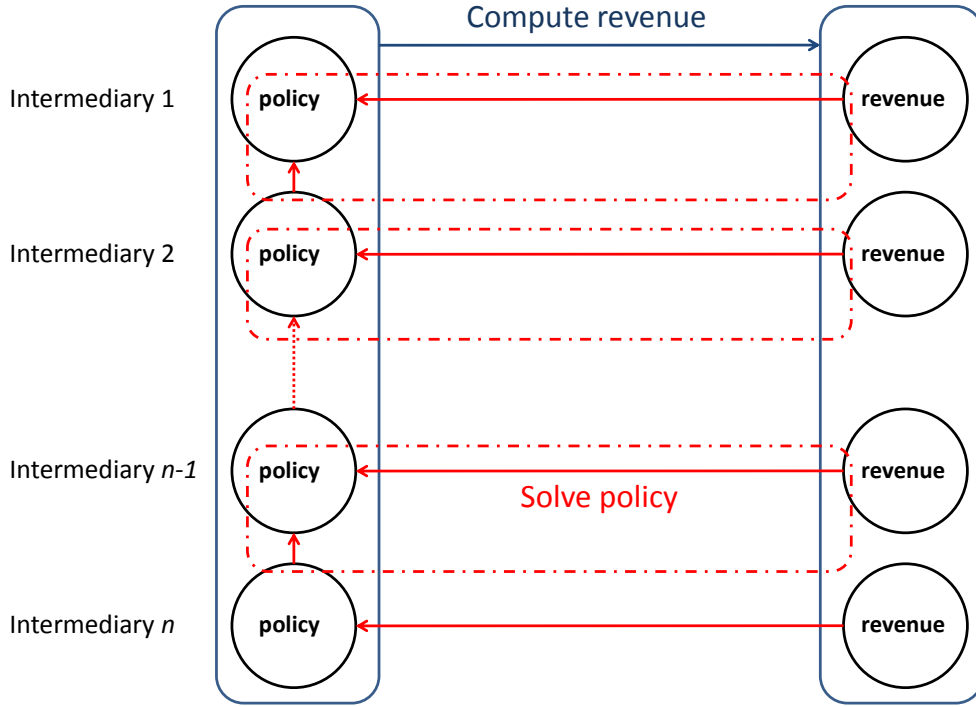


Figure 6: Iteration Process with Multiple Intermediaries

5 Outside Option

5.1 Background

In many real world persuasion settings, the strict threshold we analyzed before may not exist. In particular, an intermediary may lack the incentive to pass on any information, or may refuse to respond to previous player. This phenomenon can be addressed by incorporating an outside option into the benchmark model.

We provide the intermediary with a third choice in response stage, called **Refuse to Answer**. In a hierarchical structure, if someone chooses this outside option, the persuasion chain is broken and the final decision by the receiver will not be made. All previous intermediaries receive a large enough penalty such that intermediaries try to avoid this scenario because their recommendation has already been made, and hence their reputation is at stake. The introduction of an outside option is natural when persuasion process has a fixed cost.

The utility of the outside option can be interpreted as the incremental utility saved from the persuasion process. In the benchmark model, the minimum possible utility of an intermediary is attained at the modified threshold belief. If the fixed cost C_j does not exceed such a minimum, the equilibrium will not change. Otherwise, it is the case that intermediary j (1) **Refuses to Answer** when the incoming belief falls into some interval $[\tilde{p}_j^A, \tilde{p}_j^B]$; (2) responds A when the incoming belief falls into some interval $[0, \tilde{p}_j^A]$; and (3) responds B when the incoming belief falls into some interval

$[\tilde{p}_j^B, 1]$. It is possible that the expected gain from the outside option is so high such that some intermediary always chooses **Refuse to Answer**. This double threshold policy reflects many situations in reality where an individual has a belief in between 'yes' and 'no', and chooses to remain silent.

Based on the benchmark model of Section 2, we set up a model incorporating the outside option of intermediaries. The hierarchical structure includes one sender, one receiver and n intermediaries just as before. The state space and action space are both binary. Besides receiving state-dependent utility and reputation gains, each intermediary has a third option at the response stage, **Refuse to Answer**, which provides utility C_j with certainty. For simplicity, if the outside option is exercised, all preceding players receive a large enough punishment that can for simplicity be regarded as $-\infty$.

5.2 Outside Option in the Hierarchical Structure

Following backward induction, we first analyze the behavior of the last intermediary in the chain. Under Assumption 2.1, the *persuasion* stage of the last intermediary is the same as previously, maximizing the preferred action. In the *response* stage, with different levels of C , intermediary n may have different levels of $\tilde{p}_n^A, \tilde{p}_n^B$. This is illustrated in the three figures below.




Figure 7: Different Values of $\tilde{p}_n^A, \tilde{p}_n^B$ with Different Levels of the Outside Option

- The figure on the left shows the case when C is comparable, $0 \leq \tilde{p}_n^A < \tilde{p}_n^B \leq 1$.
- The figure in the middle shows the case when C is small, $0 < \tilde{p}_n^A = \tilde{p}_n^B < 1$.
- The figure on the right shows the case when C is large, in which **Refuse to Answer** is the dominant strategy.

To avoid the trivial case that some intermediary always chooses **Refuse to Answer**, we further assume that the utility of outside option C_j for intermediary j will not exceed his reputation gain R_j .

Assumption 5.1. $C_j \leq R_j$.

This assumption guarantees that when the state is deterministic, an intermediary will not **Refuse to Answer**.

We then move to the persuasion stage of intermediary $n - 1$. We first define

- $\tilde{p}_j^{\max} = \max\left(\{\tilde{p}_k^B\}_{k=j+1}^n, \tilde{p}_R\right)$
- $\tilde{p}_j^{\min} = \min\left(\{\tilde{p}_k^A\}_{k=j+1}^n, \tilde{p}_R\right)$

instead. This definition is consistent with benchmark model because $\tilde{p}_{Ij} = \tilde{p}_j^A = \tilde{p}_j^B$ when there is no outside option. We want to show that an equilibrium strategy in the *persuasion* stage is exactly the same as it is described in Theorem 2.5. That is, for a B-preferred intermediary, no disclosure when the incoming belief is higher than some threshold, and partial disclosure to 0 and such threshold when incoming belief is lower. For A-preferred, no disclosure when incoming belief is lower than some threshold, and partial disclosure to 1 and such threshold when incoming belief is higher.

Formally we have the following Theorem:

Theorem 5.2 (Equilibrium with Outside Option).

1. In the persuasion stage,

- For B-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \geq \tilde{p}_j^{\max}$; (2) partial disclosure that induces posterior \tilde{p}_j^{\max} with probability $\frac{\hat{p}}{\tilde{p}_j^{\max}}$ and 0 with probability $1 - \frac{\hat{p}}{\tilde{p}_j^{\max}}$.
- For A-preferred intermediary j with incoming belief \hat{p} , the following Bayesian persuasion process is optimal: (1) no disclosure when $\hat{p} \leq \tilde{p}_j^{\min}$; (2) partial disclosure that induces posterior \tilde{p}_j^{\min} with probability $\frac{1-\hat{p}}{1-\tilde{p}_j^{\min}}$ and 1 with probability $\frac{\hat{p}-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}}$.

2. In the response stage,

- Intermediary j responds A when $\hat{p} \leq \tilde{p}_j^A$, responds B when $\hat{p} \geq \tilde{p}_j^B$ and **Refuse to Answer** between A and B when $\hat{p} \in [\tilde{p}_j^A, \tilde{p}_j^B]$, where \tilde{p}_j^A and \tilde{p}_j^B are defined as the following. Let \tilde{p}_{Ij} solves $U_j(p|A) = U_j(p|B)$,

(a) If $C_j \leq U_j(\tilde{p}_{Ij}|A)$, then $\tilde{p}_j^A = \tilde{p}_j^B = \tilde{p}_{Ij}$

(b) If $C_j > U_j(\tilde{p}_{Ij}|A)$, then \tilde{p}_j^A solves $U_j(p|A) = C_j$ and \tilde{p}_j^B solves $U_j(p|B) = C_j$.

Functions $U_j(p|A)$ and $U_j(p|B)$ are defined as,

$$U_j(p|A) = \begin{cases} R_j & p \in [0, \tilde{p}_j^{\min}] \\ \frac{1-p}{1-\tilde{p}_j^{\min}} R_j + \frac{p-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}} u_j(B, \beta) & p \in [\tilde{p}_j^{\min}, 1] \end{cases}$$

$$U_j(p|B) = \begin{cases} \frac{p}{\tilde{p}_j^{\max}} (R_j + u_j(B, \alpha)) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [0, \tilde{p}_j^{\max}] \\ R_j + u_j(B, \alpha) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

The strategy within critical values \tilde{p}_j^A and \tilde{p}_j^B will not change the concave closure, hence will not change the equilibrium strategy.

The introduction of the **Refuse to Answer** option has implications for the ease of persuasion through the hierarchical chain. The possibility of refusal to answer by an intermediary later in the chain implies a risk of reputation loss for those intermediaries earlier in the chain who have already chosen A or B as their reply. Thus, intermediaries earlier in the chain of persuasion will tend to be more hesitant in providing a definitive response (A or B) about their recommendation. On the other hand, for intermediaries later or higher in the chain, conditional that the persuasion process calls for them to make a response, it can be inferred that all previous intermediaries have not chosen **Refuse to Answer**. In equilibrium, the later intermediaries are thus less likely to **Refuse to Answer**. The structure of the information transmission bears some similarity to that of an information cascade (Bikhchandani, Hirshleifer and Welch, 1992)^[42].

The following example demonstrates that the presence of an outside option reduces the likelihood of the sender's preferred action being taken, and reduces his welfare.

Example 5.3 (Outside Option with One Intermediary). *The hierarchy is comprised of one intermediary and one receiver. The prior probability is $\frac{1}{3}$. The threshold belief for the receiver is 0.5. The intermediary has state-dependent utility $u_I(B, \beta) = 2$ and $u_I(B, \alpha) = -1$. The reputation term for the intermediary is 6. If there is no outside option for the intermediary, the modified threshold belief is $\tilde{p}_I = \frac{5}{9}$, and hence action B is taken with probability $\frac{3}{5} = 0.6$. If the intermediary has outside option 5.7, then $\tilde{p}_I^A = 0.55$ and $\tilde{p}_I^B = \frac{17}{30}$. Then, action B is taken with probability $\frac{10}{17} \approx 0.588$. The presence of an outside option makes sender worse off.*

6 Discussion I: Choosing a Persuasion Path

Our analysis thus far has been confined to a single path of persuasion in the hierarchy. In this section, we allow for multiple persuasion paths between sender and receiver, allowing the sender to choose his optimal hierarchical path.

6.1 Background

In many real world settings, the persuasion path is not necessarily unique. For example, when applying for graduate school, students may convince any available professor to write a recommendation letter hoping that they have credibility with someone in the admission committee. Another route for the student is to convince his supervisor, who if persuaded, helps the student to request a recommendation letter from other senior professors. Modeling the links between professors as a social network, the objective of the student is to find the best persuasion path that maximizes his or her probability of a successful application.

6.2 Settings

Setup We follow the basic framework (**State** and **Action**), described in the benchmark model. Each intermediary j is uniquely characterized by $u_j(B, \alpha), u_j(B, \beta)$ and reputation term R_j .

Directed Graph In graph theory, a directed graph (or digraph) is a graph that is a set of vertices connected by edges, where the edges have a direction associated with them. In formal terms, a directed graph is an ordered pair $G = (V, E)$ where

1. V is a set whose elements are called vertices, nodes, or points with representative element denoted v_i ;
2. E is a set of ordered pairs of vertices, called arrows, directed edges, directed arcs, or directed lines with representative entry denoted (v_i, v_j) .

The aforementioned definition does not allow a directed graph to have multiple arrows with same source and target nodes, which coincides with the setting we aim to analyze here.

Path We call a sequence of vertices $\langle v_1, v_2, \dots, v_m \rangle$ a **path** if and only if

$$\forall i = 1, \dots, m - 1, (v_i, v_{i+1}) \in E$$

Social Network We assume that the sender (denoted as vertex s), the receiver (denoted as vertex r) as well as n intermediaries (denoted as vertex $1, \dots, n$) are involved in a directed graph. The size of the graph is polynomial at the number of vertices.¹⁸ To avoid the trivial case, we assume that there exists at least one path from sender to receiver. For path $\langle s, v_1, v_2, \dots, v_m, r \rangle$, we can calculate the modified threshold belief for each vertex v , \tilde{p}_v , and we call $\max(\{\tilde{p}_{v_i}\}_i, \tilde{p}_R)$ the **path threshold**.

Sender's Problem The sender chooses a persuasion path applying hierarchical Bayesian persuasion, in order to maximize the probability that the receiver chooses action B . Off-path vertices are ignored after the choice of persuasion path. An alternative setup where the sender chooses who to persuade next, and intermediary j chooses who will be intermediary $j + 1$ is not equivalent to the sender's problem we examine here, because under such scenario an intermediary may choose a different path with a different preferred action.

Timeline The game consists of two stages,

1. (Choice of Path) Sender chooses a persuasion path starting from vertex s to vertex r .

¹⁸ $|V| = n + 2$ and $|E| \leq (n + 2)(n + 1)$. When $|E| = (n + 2)(n + 1)$, the graph is complete since there is a direct connection between all $(n + 2)(n + 1)$ ordered pairs.

2. (Persuasion) Sender, receiver and intermediaries located on the chosen path play the hierarchical Bayesian persuasion game.

Based on the optimal design analysis in the benchmark model, we can conclude that direct communication, if available, cannot be worse than indirect communication.

Proposition 6.1. *If sender can communicate with the receiver directly, then he will do so.*

For the general case of indirect communication analysis, we need to find a sequence of vertices $\langle s, v_1, v_2, \dots, v_m, r \rangle$ such that the maximum threshold among the intermediaries and receiver is minimized. To build up to the result, we begin with the shortsighted setting in which each player makes their decision based on the threshold belief (not the modified threshold belief) that is independent of subsequent players in the persuasion path. Please refer to the Appendix for details.

6.3 Fully Forward-looking Case

Since the modified threshold belief is determined by \tilde{p}_j^{\min} and \tilde{p}_j^{\max} , for the fully forward-looking case we need to begin at the receiver node r and find the path backwards through the hierarchy until finding the sender node s . In order to avoid interaction effects, we require the graph to be a directed acyclic graph.

A directed acyclic graph (DAG) is a finite directed graph with no directed cycles. That is, it consists of finitely many vertices and edges, with each edge directed from one vertex to another, such that there is no way to start at any vertex v and follow a consistently-directed sequence of edges that eventually loops back to v again. Equivalently, a DAG is a directed graph that has a topological ordering, a sequence of the vertices such that every edge is directed from earlier in the sequence to later in the sequence.

In computer science, a topological sort or topological ordering of a directed graph is a linear ordering of its vertices such that for every directed edge uv from vertex u to vertex v , u comes before v in the ordering. For instance, the vertices of the graph may represent tasks to be performed, and the edges may represent constraints that one task must be performed before another; in this application, a topological ordering is just a valid sequence for the tasks. A topological ordering is possible if and only if the graph has no directed cycles, that is, if it is a directed acyclic graph (DAG). Any DAG has at least one topological ordering, and algorithms are known for constructing a topological ordering of any DAG in linear time.

We can compute the modified threshold beliefs of intermediaries in reverse direction of the topological order. In a directed acyclic graph, when calculating the modified threshold belief for some intermediary, modified threshold beliefs of intermediaries that are located in all possible subsequent paths have been calculated already. Please see the following example.

Example 6.2 (Fully Forward-looking). The graph structure as well as three parameters for each intermediary are shown in *Figure 8*. The operational process of the algorithm is shown in *Figure 14* in the Appendix. The resulting optimal persuasion path is shown in *Figure 9*.

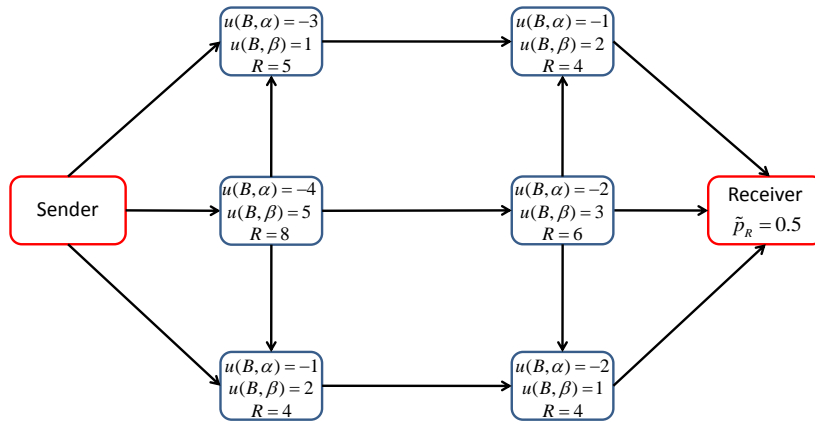


Figure 8: Graph Structure




Figure 9: Optimal Persuasion Path

The sender solves a shortest path problem when choosing a persuasion path. There are two

different ways to solve the shortest path problem, forward induction and backward induction. With fully forward-looking intermediaries, the modified threshold belief depends on the subsequent path, so we must use the backward induction approach.

7 Discussion II: Parallel Bayesian Persuasion

7.1 Background

Considering a network of individuals in a persuasion hierarchy, one possibility is that more than one individual needs to be persuaded at a time, in other words, some intermediaries need to be persuaded in parallel. For example, when applying for graduate school, admission committees typically require more than one recommendation letter. Similarly, in job promotions in academia, the Dean may consider the opinions of several professors in making the final decision. In such circumstances, the sender must persuade multiple intermediaries simultaneously, and each intermediary gives their advice to the receiver separately. We begin with the case of one intermediary per path in a parallel structure, and then generalize to the case of multiple intermediaries in the parallel structure in Section 7.4.

7.2 Setting

Setup We analyze a parallel Bayesian persuasion setup with one sender and one receiver through n parallel intermediaries, denoted as $j = 1, 2, \dots, n$. Binary state, binary action and modified threshold belief \tilde{p}_{Ij} are defined as in previous sections.¹⁹ We further assume that the sender applies public persuasion. In other words, for each signal realization, the incoming beliefs of all parallel intermediaries are the same under the common prior assumption.

Timeline The parallel persuasion game consists of two stages.

Stage 1: The sender publicly sets up a signal-generating mechanism, which consists of a family of conditional distributions $\{\pi(\cdot|t)\}_{t \in T}$ over a space of signal realizations S , and hence divides the prior belief into posterior portfolios that satisfy the Bayes' plausible condition. We denote this signal as π .

Stage 2: all n intermediaries choose a partition in the lattice structure $T \times [0, 1]$ simultaneously, denoted as π_1, \dots, π_n .

Our setup is closely related to Gentzkow and Kamenica (2017)^[4], which introduces a lattice structure to persuasion with multiple senders. The receiver observes all signals generated by the intermediaries, and hence has partition $\pi \vee (\bigvee_{j=1}^n \pi_j)$. The receiver has posterior p with distribution $\langle \pi \vee (\bigvee_{j=1}^n \pi_j) \rangle$ and then chooses an action from $\{A, B\}$. Their model corresponds nearly identically

¹⁹There exist slight difference on definition of \tilde{p}_{Ij} . Beliefs of intermediaries are correlated using lattice structure we introduced later. However, the mathematical formula would not change.

to Stage 2.

7.3 Result

Similarly to the benchmark model, we can define $\tilde{p}_P^{\max} = \max(\{\tilde{p}_{I_j}\}_j, \tilde{p}_R)$ as the minimum modified threshold belief and $\tilde{p}_P^{\min} = \min(\{\tilde{p}_{I_j}\}_j, \tilde{p}_R)$ as the maximum modified threshold belief. Again, \tilde{p}_P^{\max} quantifies how difficult it is for a (B-preferred) sender to persuade the intermediaries. Here, subscript P denotes parallel. By backward induction, we first obtain the following lemma about stage 2.

Lemma 7.1. *If some signal realization s induced posterior p such that at least one intermediary is B-preferred and at least one intermediary is A-preferred, then given signal realization s , the receiver is eventually fully informed after persuasions of the intermediaries.*

We now analyze stage 1. For the sender, if he induces some posterior p that is inside the range $(\tilde{p}_P^{\min}, \tilde{p}_P^{\max})$, then the receiver is eventually fully informed. Otherwise, all intermediaries will reach a consensus and make no further manipulations to the information structure. The following theorem characterizes the optimal disclosure policy applied by sender, which follows the same pattern as in the benchmark model.

Theorem 7.2. *The sender's optimal Bayesian persuasion is (1) no disclosure when $p_0 \geq \tilde{p}_P^{\max}$; (2) partial disclosure that induces posterior \tilde{p}_P^{\max} with probability $\frac{p_0}{\tilde{p}_P^{\max}}$ and 0 with probability $1 - \frac{p_0}{\tilde{p}_P^{\max}}$.*

The intuition behind the theorem is quite similar to the benchmark hierarchical model. For both hierarchical and parallel persuasion, intermediaries preferring different actions is always an undesirable situation for the sender. Under parallel persuasion, this conflict will finally result in full disclosure, which is unfavorable. The situation is even worse in hierarchical persuasion. An A-preferred intermediary will further confound a small probability on β in α , which further lowers the probability of $\Pr(d = B)$. Therefore, following Kamenica and Gentzkow (2011), the sender tries to generate exactly two posteriors such that one is lowest posterior such that all succeeding players prefers B at the same time while the other is 0.

7.4 Combining Hierarchical and Parallel Structures

Using the results above, we can extend our result to a more complex general structure: (1) a parallel hierarchical (PH) structure and (2) a hierarchical parallel (HP) structure. Both of them are special cases of directed acyclic graphs that avoid interactions. The following two figures highlight the network structure, where the blue boxes represent intermediaries.

Parallel Hierarchical Bayesian Persuasion

In this scenario, the sender persuades the receiver through n parallel paths, denoted as $j = 1, 2, \dots, n$. In each path j , there exist m_j intermediaries, denoted as $k_j = 1, 2, \dots, m_j$. Binary




Figure 10: Parallel Hierarchical Structure

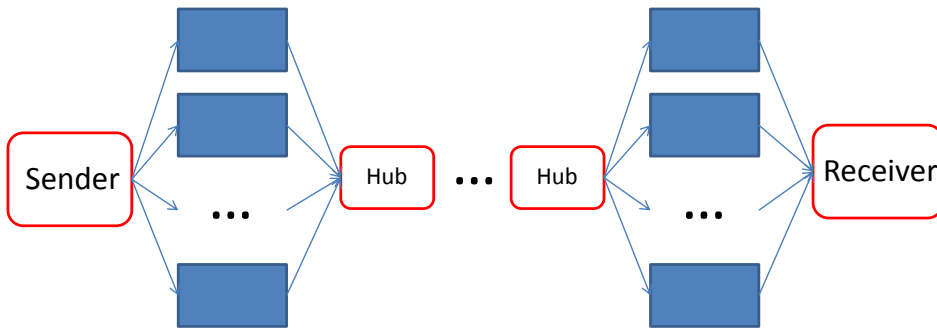


Figure 11: Hierarchical Parallel Structure

state, binary action and modified threshold belief are defined identically to in previous sections. We let $\tilde{p}_{PH}^{\max(j)}$ denote the maximum modified threshold belief of path j , and $\tilde{p}_{PH}^{\max} = \max_j \tilde{p}_{PH}^{\max(j)}$.

Consider the typical academic admissions process. When applying for admission, a student is often asked for supporting materials such as recommendation letters regarding different categories, such as academics, sports and arts. For each recommendation letter, a student may potentially persuade a reputable individual in that area indirectly. The admission committee makes a decision based on the letters across the different categories.

Hierarchical Parallel Bayesian Persuasion

The sender persuades the receiver through n hubs, denoted as $j = 1, 2, \dots, n$. Hub j influences hub $j + 1$ through m_j parallel intermediaries, denoted as $k_j = 1, 2, \dots, m_j$. Binary state, binary action and modified threshold belief are defined the same as in previous sections. We let $\tilde{p}_{HP}^{\max(j)}$ denote the maximum threshold belief connecting hub j and $j + 1$, and $\tilde{p}_{HP}^{\max} = \max(\max_j \tilde{p}_{HP}^{\max(j)}, \max_j \tilde{p}_{Ij}, \tilde{p}_R)$.

In such a network structure, the hubs are more important than intermediaries. Consider the example of an entry level worker having a profitable idea. To persuade the manager, the worker may ask several engineers to persuade the manager, who if convinced, persuades his director through some external experts on the topic, and the director may finally persuade the CEO through some members of board. The manager and director here are the hubs while the engineers and external experts serve as intermediaries.

Corollary 7.3. *In both settings described above, the sender's optimal Bayesian persuasion is (1) no disclosure when $p_0 \geq \tilde{p}_L^{\max}$; (2) partial disclosure that induces posterior \tilde{p}_L^{\max} with probability $\frac{p_0}{\tilde{p}_L^{\max}}$ and 0 with probability $1 - \frac{p_0}{\tilde{p}_L^{\max}}$, where $L = PH$ for the parallel hierarchical structure and $L = HP$ for the hierarchical parallel structure.*

Once again, the results follow the same pattern as the benchmark model. Here we have analyzed two special cases of directed acyclic graphs. We leave the generalized analysis of directed acyclic graphs as our future work.

8 Discussion III: Costly Persuasion

8.1 Background

In the Bayesian persuasion literature, we often regard the signal generation process as an investigation. However, not all signals are costless. The previous approach we used is not generally feasible if signals are costly. In that case, the sender's payoff is not fully determined by the posterior; given the posterior, the payoff also depends on the signal (due to its cost). Since one cannot express the sender's payoff as a value function over beliefs, the concavification approach does not work generally. We need to add some restrictions on form of the cost function to make it solvable under concavification. By introducing into the hierarchical persuasion model *posterior separable* cost functions (Gentzkow and Kamenica, 2014^[39]; Matyskova, 2018), for which the entropy based cost and residual variance is a typical example, the hierarchical problem is solvable under concavification.

8.2 Setting

We follow the binary state, binary action and reputation term setup as our benchmark model. Gentzkow and Kamenica (2014)^[39] introduce a family of cost functions that is compatible with the concavification approach to deriving the optimal signal. Matyskova (2018) uses the same family of cost functions to analyze a model where the receiver has additional costly information acquisition, and calls such cost functions *posterior separable*.

The cost of signals is defined by a *measurement of uncertainty* $H(\mu)$, as a mapping from a distribution to a real number. Two examples are *entropy* as proposed by Shannon (1948) and *residual variance* (Gentzkow and Kamenica, 2014^[39]). The cost of signals is proportional to the expected reduction in uncertainty. In our binary setting, without loss of generality, we assume $H(\mu) > 0$ and $H(0) = H(1) = 0$. When the probability of state β is p , entropy gives us

$$H_{\text{ENT}}(p) = -p \ln p - (1 - p) \ln(1 - p) \quad (18)$$

and residual variance gives us

$$H_{\text{RES}}(p) = p(1 - p) \quad (19)$$

Take the prosecution example in Kamenica and Gentzkow (2011) as an example here, at prior distribution $p = 0.3$, $H_{\text{ENT}}(p) = 0.611$ and $H_{\text{RES}}(p) = 0.210$. If the prosecutor applies a no disclosure policy, then the uncertainty reduction is 0. If prosecutor applies a full disclosure policy, then the uncertainty measurement for two posteriors are 0 for both measurements and both posteriors. The uncertainty reduction is 0.611 under entropy and 0.210 under residual variance. Now we consider the optimal signal, that is the posteriors are 0 with 40 percent and 0.5 with 60 percent. The uncertainty reductions are given by

$$\begin{aligned} H_{\text{ENT}}(0.3) - \left(0.4H_{\text{ENT}}(0) + 0.6H_{\text{ENT}}(0.5)\right) &= 0.195 \\ H_{\text{RES}}(0.3) - \left(0.4H_{\text{RES}}(0) + 0.6H_{\text{RES}}(0.5)\right) &= 0.060 \end{aligned}$$

Then in the persuasion stage, players need to compute the concave closure of the weighted summation of the raw utility and uncertainty measurement. In the previous section, any intermediary j tries to maximize the probability of their preferred action being taken by the receiver. However, when introducing uncertainty, the relative weight of uncertainty which we denote by $\lambda > 0$ and the winning probability may create variation in the net utility of the preferred action. Especially for an intermediary that is indifferent between two actions, she always applies a no disclosure policy. Hence, the previous two optimal Bayesian persuasion strategies do not hold generally under costly persuasion.

8.3 Persuading the Last Intermediary

To obtain a general result, the first difficulty we encounter is that in the persuasion stage, an intermediary is no longer purely a preferred action maximizer, based on our previous definition. Since persuasion is costly, each intermediary must take this into consideration. Additionally, we cannot put the uncertainty measurements directly into the winning probability, we must incorporate them into the utility terms $u_j(B, \alpha)$ and $u_j(B, \beta)$. It is intuitive that when the cost of a signal is relatively low, the optimal behavior of an intermediary in the persuasion stage is still preferred action maximization. Here we characterize such a condition.

We first consider the behavior of the last intermediary. For a B-preferred intermediary n , the utility when inducing belief p to the receiver is

$$\begin{cases} 0 & \text{if } p < \tilde{p}_R \\ R_n + pu_n(B, \beta) + (1 - p)u_n(B, \alpha) & \text{if } p \geq \tilde{p}_R \end{cases}$$

Under Assumption 2.1, when there is no cost, then the optimal policy is partial disclosure that generates posterior 0 and \tilde{p}_R when $p < \tilde{p}_R$, and no disclosure otherwise. If we incorporate the uncertainty measure, then concavification is applied on

$$V_{Bn}(p) = \begin{cases} \lambda H(p) & \text{if } p < \tilde{p}_R \\ R_n + pu_n(B, \beta) + (1 - p)u_n(B, \alpha) + \lambda H(p) & \text{if } p \geq \tilde{p}_R \end{cases}$$

The conditions

$$\begin{aligned} V'_{Bn}(0) &\leq \frac{V_{Bn}(\tilde{p}_R) - V_n(0)}{\tilde{p}_R} \\ V'_{Bn}(\tilde{p}_R) &\leq \frac{V_{Bn}(\tilde{p}_R) - V_n(0)}{\tilde{p}_R} \end{aligned}$$

are required to keep the optimal policy unchanged: (1) When $p < \tilde{p}_R$ is still optimal to generate posteriors 0 and \tilde{p}_R , and (2) when $p \geq \tilde{p}_R$, no disclosure is still optimal. By re-arranging, we obtain the following assumption

Assumption 8.1.

$$\begin{aligned} R_n &\geq \lambda \tilde{p}_R H'(0) - \tilde{p}_R u_n(B, \beta) - (1 - \tilde{p}_R)u_n(B, \alpha) - \lambda H(\tilde{p}_R) \\ R_n &\geq -u_n(B, \alpha) + \lambda \left(\tilde{p}_R H'(\tilde{p}_R) - H(\tilde{p}_R) \right) \end{aligned}$$

for reputation term R_n .

Clearly, entropy based uncertainty measurement never satisfies this assumption because $H'(0) = +\infty$. The assumption provides us with the lower bound on reputation term conditioning on c and $H(\cdot)$ such that intermediary n is a preferred action maximizer.

For a A-preferred intermediary n , the utility when inducing belief p to the receiver is

$$\begin{cases} R_n & \text{if } p < \tilde{p}_R \\ pu_n(B, \beta) + (1-p)u_n(B, \alpha) & \text{if } p \geq \tilde{p}_R \end{cases}$$

Under Assumption 2.1, when there are no costs, then the optimal policy is disclosure when $p < \tilde{p}_R$ and partial disclosure associated with \tilde{p}_R and 1 otherwise. If we incorporate the uncertainty measure, then concavification is applied on

$$V_{An}(p) = \begin{cases} R_n + \lambda H(p) & \text{if } p \leq \tilde{p}_R \\ pu_n(B, \beta) + (1-p)u_n(B, \alpha) + \lambda H(p) & \text{if } p > \tilde{p}_R \end{cases}$$

We require

$$\begin{aligned} V'_{An}(1) &\geq \frac{V_{An}(1) - V_{An}(\tilde{p}_R)}{1 - \tilde{p}_R} \\ V'_{An}(\tilde{p}_R) &\geq \frac{V_{An}(1) - V_{An}(\tilde{p}_R)}{1 - \tilde{p}_R} \end{aligned}$$

to keep the optimal policy unchanged. By re-arranging, we obtain the following

Assumption 8.2.

$$\begin{aligned} R_n &\geq u_n(B, \beta) - (1 - \tilde{p}_R)(u_n(B, \beta) - u_n(B, \alpha)) - (1 - \tilde{p}_R)H'(1)\lambda - \lambda H(\tilde{p}_R) \\ R_n &\geq u_n(B, \beta) - (1 - \tilde{p}_R)\lambda H'(\tilde{p}_R) - \lambda H(\tilde{p}_R) \end{aligned}$$

for reputation term R_n .

Next, we characterize intermediary n 's choice of A or B in the *response* stage when the incoming belief is p . The expected utility for intermediary n of responding with A and B are calculated as follows,

$$\begin{aligned} U_n(p|A) &= \begin{cases} R_n & p \in [0, \tilde{p}_R] \\ \frac{(1-p)(R_n + cH(\tilde{p}_R))}{1 - \tilde{p}_R} + \frac{(p - \tilde{p}_R)u_n(B, \beta)}{1 - \tilde{p}_R} - cH(p) & p \in [\tilde{p}_R, 1] \end{cases} \\ U_n(p|B) &= \begin{cases} \frac{p}{\tilde{p}_R}(R_n + \tilde{p}_R u_n(B, \beta) + (1 - \tilde{p}_R)u_n(B, \alpha) + cH(\tilde{p}_R)) - cH(p) & p \in [0, \tilde{p}_R] \\ R_n + u_n(B, \alpha) + p(u_n(B, \beta) - u_n(B, \alpha)) & p \in [\tilde{p}_R, 1] \end{cases} \end{aligned}$$

We need to prove that the modified threshold belief exists uniquely. The boundary conditions are

$$\begin{aligned}
U_n(0|A) &= R_n \\
U_n(1|A) &= u_n(B, \beta) \\
U_n(0|B) &= 0 \\
U_n(1|B) &= R_n + u_n(B, \beta)
\end{aligned}$$

For $p \in [0, \tilde{p}_R]$,

$$\frac{\partial^2 U_n(p|B)}{\partial p^2} = -\lambda H''(p) > 0$$

For $p \in [\tilde{p}_R, 1]$,

$$\frac{\partial^2 U_n(p|A)}{\partial p^2} = -\lambda H''(p) > 0$$

If $U_n(\tilde{p}_R|A) > U_n(\tilde{p}_R|B)$, then the unique solution exists in $[\tilde{p}_R, 1]$. If $U_n(\tilde{p}_R|A) < U_n(\tilde{p}_R|B)$, then the unique solution exists in $[0, \tilde{p}_R]$. Therefore, the modified threshold belief of intermediary n exists uniquely.

8.4 Persuading the Intermediary j

Now we consider the analysis for the general intermediary. Assume that

1. For B-preferred intermediary j , the utility when inducing belief p to intermediary $j + 1$ is

$$\begin{cases} 0 & \text{if } p < \tilde{p}_j^{\max} \\ R_j + pu_j(B, \beta) + (1-p)u_j(B, \alpha) & \text{if } p \geq \tilde{p}_j^{\max} \end{cases}$$

2. For A-preferred intermediary j , the utility when inducing belief p to intermediary $j + 1$ is

$$\begin{cases} R_j & \text{if } p < \tilde{p}_j^{\min} \\ pu_j(B, \beta) + (1-p)u_j(B, \alpha) & \text{if } p \geq \tilde{p}_j^{\min} \end{cases}$$

Similarly, we can write out the lower bound on R_j (similar to previous formulas on the reputation term). In other words, we require reputation term to be large enough. As for the specific lower bound, the derivation is similar to assumption 8.1 and 8.2) such that in the persuasion stage, all intermediaries are maximizers of the probability of their preferred action. Now, we consider the response stage of intermediary j . We characterize intermediary j 's choice of A or B in the *response* stage when the incoming belief is p . The expected utility for intermediary j of responding

with A and B are calculated as follows,

$$U_j(p|A) = \begin{cases} R_j & p \in [0, \tilde{p}_j^{\min}] \\ \frac{(1-p)(R_j + \lambda H(\tilde{p}_j^{\min}))}{1 - \tilde{p}_j^{\min}} + \frac{(p - \tilde{p}_j^{\min})u_j(B, \beta)}{1 - \tilde{p}_j^{\min}} - \lambda H(p) & p \in [\tilde{p}_j^{\min}, 1] \end{cases}$$

$$U_j(p|B) = \begin{cases} \frac{p}{\tilde{p}_j^{\max}}(R_j + \tilde{p}_j^{\max}u_j(B, \beta) + (1 - \tilde{p}_j^{\max})u_j(B, \alpha) + \lambda H(\tilde{p}_j^{\max})) - \lambda H(p) & p \in [0, \tilde{p}_j^{\max}] \\ R_j + u_j(B, \alpha) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

We need to show that an intersection point exists uniquely. In this general case, the boundary conditions and second order conditions still hold,

$$\begin{aligned} U_j(0|A) &> U_j(0|B) \\ U_j(1|A) &< U_j(0|B) \\ \frac{\partial^2 U_j(p|A)}{\partial p^2} &> 0 \quad \text{when } p \in [\tilde{p}_j^{\min}, 1] \\ \frac{\partial^2 U_j(p|B)}{\partial p^2} &> 0 \quad \text{when } p \in [0, \tilde{p}_j^{\max}] \end{aligned}$$

Depending on the signs of $U_j(\tilde{p}_j^{\min}|A) - U_j(\tilde{p}_j^{\min}|B)$ and $U_j(\tilde{p}_j^{\max}|A) - U_j(\tilde{p}_j^{\max}|B)$, we can illustrate the uniqueness of the intersection point in three different situations,

1. If $U_j(\tilde{p}_j^{\min}|A) > U_j(\tilde{p}_j^{\min}|B)$ and $U_j(\tilde{p}_j^{\max}|A) < U_j(\tilde{p}_j^{\max}|B)$, then the unique solution exists in $[\tilde{p}_j^{\min}, \tilde{p}_j^{\max}]$.
2. If $U_j(\tilde{p}_j^{\min}|A) < U_j(\tilde{p}_j^{\min}|B)$, then the unique solution exists in $[0, \tilde{p}_j^{\min}]$.
3. If $U_j(\tilde{p}_j^{\max}|A) > U_j(\tilde{p}_j^{\max}|B)$, then the unique solution exists in $[\tilde{p}_j^{\max}, 1]$.

Therefore, the modified threshold belief of intermediary j exists uniquely.

The equilibrium strategies of intermediaries in the case of costly persuasion share the same pattern with benchmark model: a threshold policy in the response stage and similar disclosure policy as described in Theorem 2.5. In order for costly persuasion to share the same pattern as benchmark model, we need to guarantee (1) the objective in the persuasion stage is preferred action maximization, and (2) the modified threshold belief exists uniquely. For the first condition, a costly signal requires a higher lower bound for the reputation term. As the cost level λ increases, the lower bound of the reputation term needed increases. However, for the second condition, there are no additional requirements for parameters except for the concavity of $H(\cdot)$.

9 Conclusion

Situations in which individuals must be convinced through a chain of command are abundant in many formal application procedures and bureaucracies. In this paper we analyze Bayesian persuasion in a hierarchical setting, in which a sender persuades a receiver through a series of intermediaries.

By introducing reputation considerations of the intermediaries in the chain, we solve for the equilibrium persuasion strategies, and identify two intuitive hierarchical persuasion approaches out of the multiplicity of equilibria: focusing on persuading the immediately subsequent intermediary, and focusing on persuading the most difficult to convince along the hierarchy. In each of our extensions, we examine whether these intuitive strategies are sustained as equilibrium persuasion approaches. It is worth noting that although the hierarchical Bayesian persuasion game has multiple equilibria, these equilibria are payoff equivalent, which pins down the payoff-based prediction power of the model. By studying the determinants of modified threshold beliefs, we learn that adding another easily convinced intermediary may be beneficial for the sender. In addition, we characterize the optimal ordering of intermediaries for the sender, which provides insights on seeking the consensus of agents through a persuasion chain.

We analyzed three main extensions of the benchmark model. First, a natural extension is to incorporate the presence of private information among the intermediaries, which we capture by an uncertainty by other players about an intermediary's persuadability. The private information setting favors the strategy of persuading the immediately subsequent intermediary over persuading the most difficult to persuade player in this setting. The intuition is straightforward. The incomplete information of intermediaries creates uncertainty, which may hurt the intermediaries due to their reputation concern. Persuading the toughest player aggregates the uncertainty from different intermediaries, which leads to greater uncertainty for each intermediary. Each intermediary will strictly prefer persuading the subsequent player.

Our second extension seeks to justify the reputation concern utilized in the benchmark model by showing that such reputation concern arises naturally, in a repeated sequential persuasion setting. By applying infinite horizon dynamic programming, we show that reputation concerns indeed arise under scenarios in which intermediaries have a higher likelihood to participate in future persuasion activities when their recommendation is successful.

The third extension we consider is a realistic one in practice, that any intermediary has the option to break the chain of persuasion by declining to send any message. The influence of this outside option by the intermediaries can be severe, as all preceding intermediaries who gave messages may incur large costs when someone takes the outside option, thus drawing similarities to the herding literature in terms of intermediaries decisions to send messages at all.

Three minor extensions or discussions of the model are also considered, mainly allowing the sender to choose among different potential persuasion paths, allowing for parallel persuasion in

the hierarchy, and incorporating the possibility of costly persuasion. In the case of the first two minor extensions, we show that the main concepts of our baseline result are generally robust to endogenizing the persuasion path, and allowing parallel persuasion activities, respectively. In the case of costly persuasion, we derive the conditions under which the main pattern of equilibria found in the benchmark model holds.

We view the Bayesian persuasion framework as particularly well-suited for studying communication in a hierarchy, in the context of bureaucratic settings. Bureaucracies have the feature that individuals who are members of the hierarchy have the incentive to behave strategically with respect to their reputations, while also being able in theory, to set the kind of commitment strategies that are characteristic of the Bayesian persuasion approach, in order to appear impartial to the message conveyed.

Our main baseline result, that the persuasion game has many equilibria which are all payoff-equivalent, implies that if members of a bureaucracy are game theoretic in their choices, there could be many ways to effectively persuade, but only a single possible outcome for members of the bureaucracy in terms of their benefits obtained. Thus, the exact method of persuasion may be inconsequential for bureaucrats' payoffs. The model also provides explanations for political behavior within bureaucracies. In particular, we demonstrate the theoretical justification behind the intuition that a bureaucrat benefits from inviting an easy-to-convince colleague to the persuasion chain, which can explain the political support for like-minded bureaucrats by those who seek to persuade policy-makers.

The model extension with private information of intermediaries shows that in situations where bureaucrats have asymmetric information, the strategy of persuading the immediately higher-up bureaucrat is preferred over targeting the most difficult to persuade. This strategy is close to the seemingly rule-of-thumb style behavior we often observe in bureaucracies. Our model shows that a driving force for such a strategy is the compounding of uncertainties under private information of members of the bureaucracy, rather than institutional barriers in accessing and persuading the bureaucrats of a higher rank.

Finally, the model extension that gives intermediaries the option to decline sending any message, informs us on the consequences of possible halted communication in bureaucracies. Interacted with the reputation concern of bureaucrats, the possibility that some other bureaucrat higher in the chain might refuse to answer leads lower level bureaucrats to be relatively hesitant in making a response. However, bureaucrats higher up in the hierarchy can in some sense 'free ride' from the information implied in lower level bureaucrats' definitive responses sent.

Our model does have several limitations which are still unsolved in the literature on Bayesian persuasion. Future work can consider persuasion under a non-binary action space, a generalization of the conflict of interest format, and persuasion in a generalized network structure. Advancements in these areas can further enhance the applicability of the Bayesian persuasion approach in modeling real world phenomena.

A Proofs

Proof of **Theorem 2.5**:

1. Persuasion Stage

When $j = n$, the results in the Theorem hold according to **Lemma 2.2**, **Lemma 2.3**, and **Lemma 2.4**. Assume that for $j = k + 1, \dots, n$, the results in the theorem hold. We then prove that for $j = k$, the results in the theorem hold.

There are three different orderings for $\tilde{p}_{I,k+1}, \tilde{p}_{k+1}^{\min}, \tilde{p}_{k+1}^{\max}$ and there are two different preferred actions. Hence, there are six different circumstances altogether. **Table 1** illustrates the probability of the preferred action when intermediary k induces belief p to intermediary $k + 1$. Among the three orderings, the probability given *B-preferred* and the probability given *A-preferred* have a summation of 1.



Table 1: Probability of preferred action when intermediary k induces belief p to intermediary $k + 1$

	$\tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$	$\tilde{p}_{k+1}^{\min} < \tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\max}$	$\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{I,k+1}$
<i>B-preferred</i>	$\begin{cases} 0 & 0 \leq p < \tilde{p}_{I,k+1} \\ \frac{p}{\tilde{p}_{k+1}^{\max}} & \tilde{p}_{I,k+1} \leq p < \tilde{p}_{k+1}^{\max} \\ 1 & \tilde{p}_{k+1}^{\max} \leq p \leq 1 \end{cases}$	$\begin{cases} 0 & 0 \leq p < \tilde{p}_{k+1}^{\min} \\ \frac{p - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} & \tilde{p}_{k+1}^{\min} \leq p < \tilde{p}_{I,k+1} \\ \frac{p}{\tilde{p}_{k+1}^{\max}} & \tilde{p}_{I,k+1} \leq p < \tilde{p}_{k+1}^{\max} \\ 1 & \tilde{p}_{k+1}^{\max} \leq p \leq 1 \end{cases}$	$\begin{cases} 0 & 0 \leq p < \tilde{p}_{k+1}^{\min} \\ \frac{p - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} & \tilde{p}_{k+1}^{\min} \leq p < \tilde{p}_{I,k+1} \\ 1 & \tilde{p}_{I,k+1} \leq p \leq 1 \end{cases}$
<i>A-preferred</i>	$\begin{cases} 1 & 0 \leq p \leq \tilde{p}_{I,k+1} \\ \frac{\tilde{p}_{k+1}^{\max} - p}{\tilde{p}_{k+1}^{\max}} & \tilde{p}_{I,k+1} < p \leq \tilde{p}_{k+1}^{\max} \\ 0 & \tilde{p}_{k+1}^{\max} < p \leq 1 \end{cases}$	$\begin{cases} 1 & 0 \leq p \leq \tilde{p}_{k+1}^{\min} \\ \frac{1-p}{1 - \tilde{p}_j^{\min}} & \tilde{p}_{k+1}^{\min} < p \leq \tilde{p}_{I,k+1} \\ \frac{\tilde{p}_{k+1}^{\max} - p}{\tilde{p}_{k+1}^{\max}} & \tilde{p}_{I,k+1} < p \leq \tilde{p}_{k+1}^{\max} \\ 0 & \tilde{p}_{k+1}^{\max} < p \leq 1 \end{cases}$	$\begin{cases} 1 & 0 \leq p \leq \tilde{p}_{k+1}^{\min} \\ \frac{1-p}{1 - \tilde{p}_j^{\min}} & \tilde{p}_{k+1}^{\min} < p \leq \tilde{p}_{I,k+1} \\ 0 & \tilde{p}_{I,k+1} < p \leq 1 \end{cases}$



For all six circumstances, we can easily find the concave closure graphically, as shown in the following figures.

Here is an example, we provide an explanation of Case (1) illustrated in the first panel. For other cases illustrated in other panels, we omit the detailed reasoning which is analogous to that of Case (1).

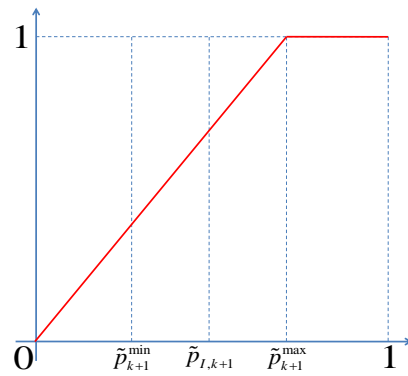
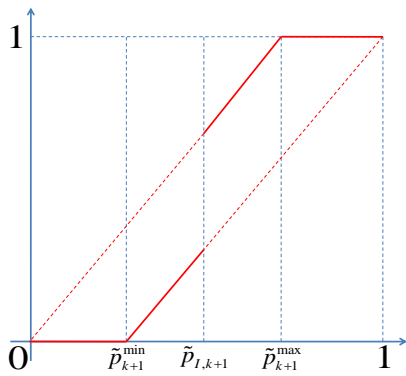
Case (1): B-preferred, $\tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$.





Case (2): A-preferred, $\tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$.





Case (3): B-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\max}$.





Case (4): A-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\max}$



Case (5): B-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{I,k+1}$



Case (6): A-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{I,k+1}$



For case (1), intermediary k is B-preferred, and $\tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$.

If intermediary $j+1$ receives belief \hat{p} , then he or she will reply A to intermediary j if $\hat{p} < \tilde{p}_{I,k+1}$ and reply B to intermediary j if $\hat{p} > \tilde{p}_{I,k+1}$.

If $\hat{p} > \tilde{p}_{I,k+1}$, then intermediary $k+1$ applies the strategy described in the theorem: (1) no disclosure when $\hat{p} \geq \tilde{p}_{k+1}^{\max}$, (2) partial disclosure that induces posterior \tilde{p}_{k+1}^{\max} with probability $\frac{\hat{p}}{\tilde{p}_{k+1}^{\max}}$, and 0 with probability $1 - \frac{\hat{p}}{\tilde{p}_{k+1}^{\max}}$. All successive players' messages are fully uninformative according to the induction hypothesis. The probability that action B is taken is then 1 if $\hat{p} \geq \tilde{p}_{k+1}^{\max}$, and $\frac{\hat{p}}{\tilde{p}_{k+1}^{\max}}$ otherwise.

If $\hat{p} < \tilde{p}_{I,k+1}$, then all successive players' messages are fully uninformative (intermediary $k+1$ provides nothing, intermediary $k+2$ is also A-preferred, and so on) according to the induction hypothesis, and the final decision will be A .

Therefore, the probability of the preferred action been taken by receiver when intermediary k induces belief p to intermediary $k+1$ is,

$$\begin{cases} 0 & 0 \leq p < \tilde{p}_{I,k+1} \\ \frac{\hat{p}}{\tilde{p}_{k+1}^{\max}} & \tilde{p}_{I,k+1} \leq p \leq \tilde{p}_{k+1}^{\max} \\ 1 & \tilde{p}_{k+1}^{\max} \leq p \leq 1 \end{cases}$$

The concave closure of this is

$$\begin{cases} \frac{\hat{p}}{\tilde{p}_{k+1}^{\max}} & 0 \leq p \leq \tilde{p}_{k+1}^{\max} \\ 1 & \tilde{p}_{k+1}^{\max} \leq p \leq 1 \end{cases}$$

because the slope of the function inside the interval $[\tilde{p}_{I,k+1}, \tilde{p}_{k+1}^{\max}]$ is exactly the same as the slope of the line connecting the origin point to $(\tilde{p}_{I,k+1}, \frac{\tilde{p}_{I,k+1}}{\tilde{p}_{k+1}^{\max}})$.

Since \tilde{p}^{\max} and \tilde{p}^{\min} have the following iterative relationships,

$$\tilde{p}_k^{\max} = \begin{cases} \tilde{p}_{k+1}^{\max} & \tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\max} \\ \tilde{p}_{I,k+1} & \tilde{p}_{I,k+1} > \tilde{p}_{k+1}^{\max} \end{cases} \quad (20)$$

$$\tilde{p}_k^{\min} = \begin{cases} \tilde{p}_{I,k+1} & \tilde{p}_{I,k+1} < \tilde{p}_{k+1}^{\min} \\ \tilde{p}_{k+1}^{\min} & \tilde{p}_{I,k+1} > \tilde{p}_{k+1}^{\min} \end{cases} \quad (21)$$

Therefore, when $j = k$, the result regarding the persuasion stage in the Theorem holds.

2. Response Stage

The expected utilities for intermediary k by responding with A and B are calculated respec-

tively as follow,

$$\begin{aligned}
U_k(p|A) &= \begin{cases} R_k & p \in [0, \tilde{p}_k^{\min}] \\ \frac{1-p}{1-\tilde{p}_k^{\min}} R_k + \frac{p-\tilde{p}_k^{\min}}{1-\tilde{p}_k^{\min}} u_k(B, \beta) & p \in [\tilde{p}_k^{\min}, 1] \end{cases} \\
U_k(p|B) &= \begin{cases} \frac{p}{\tilde{p}_k^{\max}} (R_k + u_k(B, \alpha)) + p(u_k(B, \beta) - u_k(B, \alpha)) & p \in [0, \tilde{p}_k^{\max}] \\ R_k + u_k(B, \alpha) + p(u_k(B, \beta) - u_k(B, \alpha)) & p \in [\tilde{p}_k^{\max}, 1] \end{cases}
\end{aligned}$$

Since $U_k(p|A)$ is decreasing in p while $U_k(p|B)$ is strictly increasing in p , and

$$\begin{aligned}
U_k(0|A) &> U_k(0|B) \\
U_k(1|A) &< U_k(1|B)
\end{aligned}$$

there exists a modified threshold belief \tilde{p}_{Ik} such that intermediary k is indifferent between choosing A and B in *response* stage, $U_k(\tilde{p}_{Ik}|A) = U_k(\tilde{p}_{Ik}|B)$. Therefore, our theorem holds when $j = k$, which completes our proof. \square

Proof of Theorem 2.11. We first obtain a lower bound on \tilde{p}_0^{\min} . Sometimes, $\tilde{p}_0^{\min} = \tilde{p}_R$, which means that the receiver is actually the most sender-aligned player in the hierarchy. If not, we can see that \tilde{p}_0^{\min} is bounded by the inverse degree of sender-alignment.

Lemma A.1. *If for some permutation, $\tilde{p}_0^{\min} < \tilde{p}_R$, then*

$$\tilde{p}_0^{\min} \geq \min_J \underline{p}_J = (\max_J \mathbb{K}_J)^{-1} \quad (22)$$

Proof. Assume that \tilde{p}_0^{\min} is attained for player K located at position k , $\tilde{p}_0^{\min} = \tilde{p}_{Ik}$. Then \tilde{p}_{Ik} is defined by the intersection point of $U_k(p|A)$ and $U_k(p|B)$. Since we must have $\tilde{p}_{Ik} \leq \tilde{p}_k^{\min} \leq \tilde{p}_k^{\max}$, then \tilde{p}_{Ik} solves

$$R_K = \frac{p}{\tilde{p}_k^{\max}} (R_K + u_K(B, \alpha)) + p(u_K(B, \beta) - u_K(B, \alpha)) \quad (23)$$

where the left hand side is the expression of $U_k(p|A)$ when $p \leq \tilde{p}_k^{\min}$, and the right hand side is the expression of $U_k(p|B)$ when $p \leq \tilde{p}_k^{\max}$. Then by the fact that $\tilde{p}_k^{\max} \geq \tilde{p}_R$, we have the following relationship,

$$R_K \leq \frac{p}{\tilde{p}_R} (R_K + u_K(B, \alpha)) + p(u_K(B, \beta) - u_K(B, \alpha)) \quad (24)$$

after rearranging, we obtain

$$\tilde{p}_{Ik} \geq R_K \left(\frac{R_K + u_K(B, \alpha)}{\tilde{p}_R} + u_K(B, \beta) - u_K(B, \alpha) \right) = \underline{p}_K \geq \min_J \underline{p}_J$$

\square

We then claim that such lower bound can be attained under some specific permutations, in which the most sender-aligned intermediary talks to the receiver directly.

Lemma A.2. *If it is true that (1) there exists some permutation such that $\tilde{p}_0^{\min} < \tilde{p}_R$, (2) $K = \arg \min_J \underline{p}_J$, then for all permutations σ such that $\sigma(K) = n$, we have*

$$\tilde{p}_0^{\min} = \min_J \underline{p}_J = \underline{p}_K \quad (25)$$

Proof. We can verify that $\tilde{p}_{In} = \underline{p}_K$ solves $U_n(p|A) = U_n(p|B)$ and $\tilde{p}_{In} < \tilde{p}_R = \tilde{p}_n^{\min} = \tilde{p}_n^{\max}$. Then $\tilde{p}_0^{\min} \leq \tilde{p}_{In}$ because \tilde{p}_0^{\min} is the minimum among all modified threshold beliefs. According to the previous lemma, we can conclude that $\tilde{p}_0^{\min} = \min_J \underline{p}_J$. \square

Both [Lemma A.1](#) and [Lemma A.2](#) require the condition that there exists some permutation such that $\tilde{p}_0^{\min} < \tilde{p}_R$. However, when does such a permutation exist? The following lemma shows us that such permutation exists if and only if $\min_J \underline{p}_J < \tilde{p}_R$. As we definition in main body, if this condition is met, we call the most sender-aligned intermediary the most sender-aligned player. If this relationship is not met, we call receiver the most sender-aligned player. Recall from [Observation 2.10](#), that the $\min_J \underline{p}_J < \tilde{p}_R$ condition is equivalent to the existence of at least one player with threshold belief smaller than \tilde{p}_R , which provides us with an easier determination rule on whether we can obtain $\tilde{p}_0^{\min} < \tilde{p}_R$.²⁰ To summarize,

$$\min_J \underline{p}_J < \tilde{p}_R \iff \min_J \tilde{p}_J < \tilde{p}_R$$

Lemma A.3. \bullet *If $\min_J \underline{p}_J < \tilde{p}_R$ and $K = \arg \min_J \underline{p}_J$, then for all permutations σ such that $\sigma(K) = n$, we have*

$$\tilde{p}_0^{\min} = \min_J \underline{p}_J \quad (26)$$

- \bullet *If $\min_J \underline{p}_J \geq \tilde{p}_R$, under all permutations σ , $\tilde{p}_0^{\min} = \tilde{p}_R$.*

Proof. If $\min_J \underline{p}_J < \tilde{p}_R$, and $K = \arg \min_J \underline{p}_J$, by ordering player K as intermediary n , we can get a threshold belief $\tilde{p}_{In} = \underline{p}_K < \tilde{p}_R$. Then, condition (1) of the previous lemma is satisfied, and hence $\tilde{p}_0^{\min} = \min_J \underline{p}_J$.

If $\min_J \underline{p}_J \geq \tilde{p}_R$, we need to show that under all permutations, the modified threshold belief of all intermediaries cannot be less than \tilde{p}_R , which is equivalent to proving that $U_j(\tilde{p}_R|B) \leq U_j(\tilde{p}_R|A)$ for all intermediaries j . We can prove this by induction. For intermediary n , we have the following equivalent inequalities,

²⁰Note however, that \underline{p}_J and \tilde{p}_J may not be minimized by the same J .

$$\begin{aligned}
R_n \left(\frac{R_n + u_n(B, \alpha)}{\tilde{p}_R} + u_n(B, \beta) - u_n(B, \alpha) \right)^{-1} &\geq \tilde{p}_R \\
\tilde{p}_R \left(\frac{R_n + u_n(B, \alpha)}{\tilde{p}_R} + u_n(B, \beta) - u_n(B, \alpha) \right) &\leq R_n \\
U_n(\tilde{p}_R|B) = R_n + u_n(B, \alpha) + \tilde{p}_R(u_n(B, \beta) - u_n(B, \alpha)) &\leq R_n = U_n(\tilde{p}_R|A)
\end{aligned}$$

Assume that for intermediary j , $\tilde{p}_j^{\min} = \tilde{p}_R$, we have the following equivalent inequalities,

$$\begin{aligned}
R_j \left(\frac{R_j + u_j(B, \alpha)}{\tilde{p}_R} + u_j(B, \beta) - u_j(B, \alpha) \right)^{-1} &\geq \tilde{p}_R \\
\tilde{p}_R \left(\frac{R_j + u_j(B, \alpha)}{\tilde{p}_R} + u_j(B, \beta) - u_j(B, \alpha) \right) &\leq R_j \\
R_j + u_j(B, \alpha) + \tilde{p}_R(u_j(B, \beta) - u_j(B, \alpha)) &\leq R_j
\end{aligned}$$

Multiplying the $R_j + u_j(B, \alpha)$ term by a constant $\frac{\tilde{p}_R}{\tilde{p}_j^{\max}}$ that is less than or equal to 1, the left hand side is reduced to,

$$U_j(\tilde{p}_R|B) = \frac{\tilde{p}_R}{\tilde{p}_j^{\max}}(R_j + u_j(B, \alpha)) + \tilde{p}_R(u_j(B, \beta) - u_j(B, \alpha)) \leq R_j = U_j(\tilde{p}_R|A)$$

We can then conclude that $\tilde{p}_0^{\min} = \tilde{p}_R$. □

For clarity, we summarize the main conditions on permutations of intermediaries so far as follows:

- **Statement 1.** There exists some permutation such that $\tilde{p}_0^{\min} < \tilde{p}_R$.
- **Statement 2.** There exists some permutation such that $\tilde{p}_0^{\min} = \min_J \underline{p}_J$.
- **Statement 3.** $\min_J \underline{p}_J < \tilde{p}_R$.

Lemma A.2 shows that Statement 1 \rightarrow Statement 2 and the first part of **Lemma A.3** shows that Statement 3 \rightarrow Statement 1. The second part of **Lemma A.3** shows that Statement $\neg 3 \rightarrow$ Statement $\neg 1$ through Statement 4, where

- **Statement $\neg 1$.** In all permutations, $\tilde{p}_0^{\min} \geq \tilde{p}_R$.
- **Statement $\neg 3$.** $\min_J \underline{p}_J \geq \tilde{p}_R$.
- **Statement 4.** In all permutations, $\tilde{p}_0^{\min} = \tilde{p}_R$.

Then statements 1 and 3 are equivalent, which addresses our original question. The following lemma tells us that from the perspective of sender, the most sender-aligned player (if it is not the receiver) should be the intermediary who communicates with the receiver directly.

Lemma A.4. *If $\min_J \underline{p}_J < \tilde{p}_R$ and $K = \arg \min_J \underline{p}_J$. We move player K to position n and keep the relative positions of all other players unchanged, then \tilde{p}_0^{\max} will weakly decrease.*

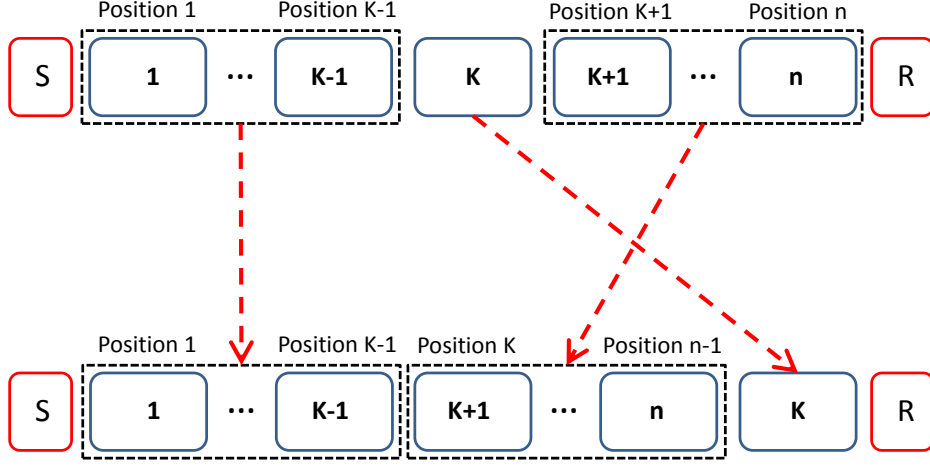


Figure 12: Changing Position

Figure 12 shows us such a procedure where S in the red box denotes the sender, R in the red box denotes the receiver, and blue boxes with player labels represent intermediaries.

Proof. Without loss of generality, we assume that the permutation is an identity mapping before changing the position, i.e., $\sigma(j) = j$. Then by letting K move to the position n , player $j = 1, \dots, K-1$ remain at their relative positions while player $j = K+1, \dots, n$ now moves to position $j-1$. After changing position, player n (now in position $n-1$) faces a new $U(p|A)$ and $U(p|B)$ where \tilde{p}^{\min} decreases from \tilde{p}_R to \underline{p}_K while \tilde{p}^{\max} remains at \tilde{p}_R . Then his modified threshold belief will weakly decrease. That makes both \tilde{p}^{\min} and \tilde{p}^{\max} faced by player $n-1$ weakly decrease. The reasoning process is shown in Figure 13, where \downarrow represents the associated value for some specific individual player (such individual may no longer be in same position after moving player K) which weakly decreases compared with the original permutation. Therefore, \tilde{p}_0^{\max} weakly decreases. \square

The intuition behind this lemma is that we are weakly better off if the intermediary who communicates with the decision-maker is the most sender-aligned player. Directly from the previous lemma, when searching for the optimal order (with lowest \tilde{p}_0^{\max}), it is outcome equivalent (having the same \tilde{p}_0^{\max}) to search within some specific order where player K is intermediary n . The following lemma will tell us that the order of other players is irrelevant when $\min_J \underline{p}_J < \tilde{p}_R$.

Lemma A.5. • *If $\min_J \underline{p}_J < \tilde{p}_R$ and $K = \arg \min_J \underline{p}_J$, as long as player K is intermediary n , \tilde{p}_0^{\max} are the same, irrespective of the ordering.*

• *If $\min_J \underline{p}_J \geq \tilde{p}_R$, then \tilde{p}_0^{\max} are the same, irrespective of the ordering.*

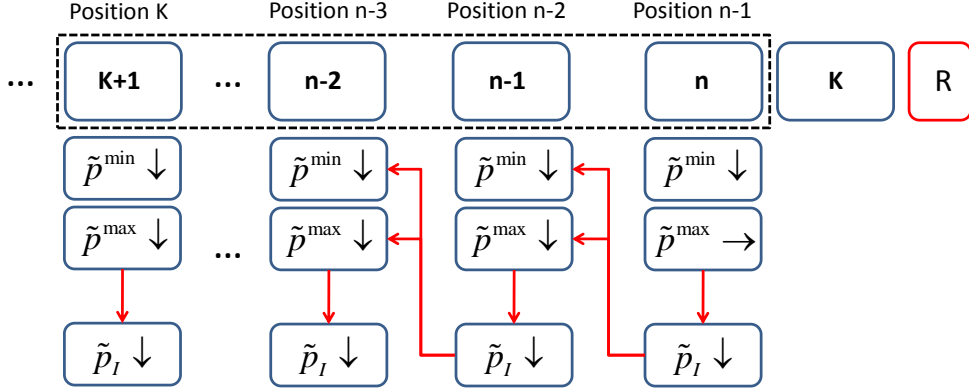


Figure 13: How Modified threshold Beliefs Change

Proof. Sender tries to minimize \tilde{p}_0^{\max} .

Case (1) If $\min_J \underline{p}_J < \tilde{p}_R$ and $K = \arg \min_J \underline{p}_J$, from previous the analysis, for all intermediaries except for intermediary n , $\tilde{p}_j^{\min} = \underline{p}_K$. Then the modified threshold belief for player $J \neq K$ is solved by letting

$$U_J(p|A) = U_J(p|B)$$

since the following relationship directly follows from the fact that K minimizes \underline{p}_J ,

$$U_J(\underline{p}_K|A) \geq U_J(\underline{p}_K|B)$$

we can conclude that modified threshold belief of player J must be no less than \underline{p}_K , then at the intersection point,

$$U_J(p|A) = \frac{1-p}{1-\underline{p}_K} R_J + \frac{p-\underline{p}_K}{1-\underline{p}_K} u_J(B, \beta)$$

and we have

$$U_J(p|B) \leq R_J + u_J(B, \alpha) + p(u_J(B, \beta) - u_J(B, \alpha))$$

Then the solution of the following equation gives us the lower bound for the modified threshold belief of player J ,

$$\frac{1-p}{1-\underline{p}_K} R_J + \frac{p-\underline{p}_K}{1-\underline{p}_K} u_J(B, \beta) = R_J + u_J(B, \alpha) + p(u_J(B, \beta) - u_J(B, \alpha))$$

and denoted as \bar{p}_J ,

$$\bar{p}_J = \frac{\underline{p}_K R_J - u_J(B, \beta)}{R_J - u_J(B, \beta) + (1 - \tilde{p}_K)(u_J(B, \beta) - u_J(B, \alpha))} > \frac{\underline{p}_K R_J - \underline{p}_K u_J(B, \beta)}{R_J - u_J(B, \beta)} = \underline{p}_K$$

Then $\max_J \bar{p}_J$ gives us the lower bound for \tilde{p}_0^{\max} .

We now prove that for any ordering, we can achieve this lower bound. Let $L = \arg \max_J \bar{p}_J$ (later we will call him toughest player because he has the maximum modified threshold belief), and assume l is L 's position.

Step 1 Prove that $\tilde{p}_l^{\max} \leq \bar{p}_L$. By contrast, we assume that $\tilde{p}_l^{\max} > \bar{p}_L$ and this modified threshold belief is attained by player M at position m . Then \tilde{p}_{Im} is the intersection point of

$$\frac{1 - p}{1 - \underline{p}_K} R_m + \frac{p - \underline{p}_K}{1 - \underline{p}_K} u_m(B, \beta) = R_m + u_m(B, \alpha) + p(U_m(B, \beta) - u_m(B, \alpha))$$

which means that $\tilde{p}_{Im} = \bar{p}_M > \bar{p}_L$. This provides a contradiction because \bar{p}_L reaches the maximum.

Step 2 Prove that $\tilde{p}_{ll} = \bar{p}_L$. From the definition of \bar{p}_L , we have

$$U_l(\bar{p}_L|A) = \frac{1 - \bar{p}_L}{1 - \underline{p}_K} R_L + \frac{\bar{p}_L - \underline{p}_K}{1 - \underline{p}_K} u_L(B, \beta) = R_L + u_L(B, \alpha) + \bar{p}_L(u_L(B, \beta) - u_L(B, \alpha)) = U_l(\bar{p}_L|B)$$

The last equation holds because $\bar{p}_L \geq \tilde{p}_l^{\max}$. \tilde{p}_{ll} solves $U_l(p|A) = U_l(p|B)$, hence $\tilde{p}_{ll} = \bar{p}_L$.

Step 3 Prove that for all preceding players W at positions $w \leq l$, $\tilde{p}_{Iw} \leq \bar{p}_L$. We prove this by backward induction. The case $w = l$ holds trivially. Then for the induction process, assume that for some w , we have $\tilde{p}_w^{\max} = \bar{p}_L$ and $\tilde{p}_w^{\min} = \underline{p}_K$. We need to prove that $\tilde{p}_{I,w-1} \leq \bar{p}_L$, which is equivalent to $U_w(\bar{p}_L|A) \leq U_w(\bar{p}_L|B)$, directly derived from $\bar{p}_L \geq \bar{p}_W$ for all W :

$$\begin{aligned} U_w(\bar{p}_L|A) &\leq U_w(\bar{p}_L|B) \\ \frac{1 - \bar{p}_L}{1 - \underline{p}_K} R_W + \frac{\bar{p}_L - \underline{p}_K}{1 - \underline{p}_K} u_W(B, \beta) &\leq R_W + u_W(B, \alpha) + \bar{p}_L(u_W(B, \beta) - u_W(B, \alpha)) \\ -\bar{p}_L R_W + \bar{p}_L u_w(B, \beta) &\leq -\underline{p}_K R_w + u_w(B, \alpha) + (1 - \underline{p}_K) \bar{p}_L(u_w(B, \beta) - u_w(B, \alpha)) \\ \underline{p}_K R_W - u_W(B, \alpha) &\leq \bar{p}_L R_W - \bar{p}_L u_W(B, \beta) + (1 - \underline{p}_K) \bar{p}_L(u_W(B, \beta) - u_W(B, \alpha)) \\ \bar{p}_W &\leq \bar{p}_L \end{aligned}$$

All inequalities are equivalent and the last inequality holds trivially.

Case (2) If $\min_J \underline{p}_J < \tilde{p}_R$, then $\tilde{p}_0^{\min} = \tilde{p}_R$ for all possible permutations according to **Lemma A.3**. The remaining analysis is similar to the previous case except we use \tilde{p}_R instead of \underline{p}_K . \square

\square

Proof of Proposition 3.2. Assume that utility function $V(\cdot)$ is defined over probability space $[0, 1]$. The proof of the proposition follows from several observations.

We first establish the following Lemma, which illustrates that two posteriors are sufficient for implementing optimal Bayesian persuasion signal.

Lemma A.6. *If \mathbb{C} is not the optimal strategy for belief p , then all possible optimal disclosure policies are outcome-equivalent to a policy that consists of two posteriors that are minimum and maximum of the original posteriors profile.*

Proof. If the posteriors profile consists of more than 2 posteriors, consider the arbitrary three posteriors in profile, $p_1 < p_2 < p_3$. We claim that

$$V(p_2) = \frac{p_2 - p_1}{p_3 - p_1} V(p_3) + \frac{p_3 - p_2}{p_3 - p_1} V(p_1) \quad (27)$$

in other words, p_2 and

$$\begin{array}{ll} p_1 & \text{with probability } \frac{p_3 - p_2}{p_3 - p_1} \\ p_3 & \text{with probability } \frac{p_2 - p_1}{p_3 - p_1} \end{array}$$

are mutually replaceable. Otherwise, there exists a better disclosure policy, and the original one cannot be optimal. If the weight of p_2 in the posterior is q , then we increase the weight of p_1 by $\frac{p_3 - p_2}{p_3 - p_1} q$ and p_3 by $\frac{p_2 - p_1}{p_3 - p_1} q$, then we obtain another optimal policy that contains fewer posteriors. This procedure can be repeated until there are only two posteriors included. Furthermore, these two posteriors are the minimum and maximum of the original posteriors, respectively. \square

From now on, without loss of generality, we assume that if no disclosure is not optimal for belief p , then optimal Bayesian persuasion must consist of two posteriors. The following lemma builds the bridge from analysis of a specific prior to analysis of all priors in some interval.

Lemma A.7. *If it is optimal to induce posteriors $p_l < p < p_h$ for belief p , then it is optimal to induce posteriors $p_l < p' < p_h$ for all beliefs $p' \in (p_l, p_h)$.*

Proof. The following disclosure policy is optimal for belief p ,

$$\begin{array}{ll} p_h & \text{with probability } \frac{p - p_l}{p_h - p_l} \\ p_l & \text{with probability } \frac{p_h - p}{p_h - p_l} \end{array}$$

Assume (for a contradiction) that there exists a $p' \in (p_l, p_h)$ such that inducing beliefs $p'_l < p' < p'_h$

is better than inducing beliefs $p_l < p < p_h$, which indicates

$$\frac{p' - p'_l}{p'_h - p'_l} V(p'_h) + \frac{p'_h - p'}{p'_h - p'_l} V(p'_l) > \frac{p' - p_l}{p_h - p_l} V(p_h) + \frac{p_h - p'}{p_h - p_l} V(p_l) \quad (28)$$

Then we can find a better disclosure policy for belief p ,

$$\begin{aligned} p_h & \text{ with probability } \frac{p - p_l}{p_h - p_l} - \varepsilon \frac{p' - p_l}{p_h - p_l} \\ p_l & \text{ with probability } \frac{p_h - p}{p_h - p_l} - \varepsilon \frac{p_h - p'}{p_h - p_l} \\ p'_h & \text{ with probability } \varepsilon \frac{p' - p'_l}{p'_h - p'_l} \\ p'_l & \text{ with probability } \varepsilon \frac{p'_h - p'}{p'_h - p'_l} \end{aligned}$$

where $0 < \varepsilon < \min(\frac{p-p_l}{p'-p_l}, \frac{p_h-p}{p_h-p'})$. (Contradiction) □

The previous lemma illustrates algebraically that if for some probability $p \in [\gamma', \gamma'']$, it is optimal to generate two posteriors γ' and γ'' , then these two posteriors are optimal for all prior probabilities between γ' and γ'' . We can prove this lemma geometrically as following. From the optimality for p to generate γ' and γ'' , function V must lie below the line connecting $(\gamma', V(\gamma'))$ and $(\gamma'', V(\gamma''))$ for all $p \in [0, 1]$, which gives us an upper bound of concave closure \hat{V} . Meanwhile, it is possible for all priors between γ' and γ'' to reach this upper bound, which proves its optimality. The next lemma demonstrates how we can merge many intervals for the same prior if we have already found several pairs of γ', γ'' .

Lemma A.8. *If there exist multiple optimal disclosure policies for belief p , associated posteriors are denoted as $\{p_l^i, p_h^i\}_{i=1, \dots, n}$, then a disclosure policy generating posteriors as $\min_i p_l^i$ and $\max_i p_h^i$ is also optimal.*

Proof. It is also optimal for the following policy,

$$\begin{array}{lll}
p_h^1 & \text{with probability} & \frac{1}{m} \frac{p - p_l^1}{p_h^1 - p_l^1} \\
p_l^1 & \text{with probability} & \frac{1}{m} \frac{p_h^1 - p}{p_h^1 - p_l^1} \\
\dots & \dots & \dots \\
\dots & \dots & \dots \\
p_h^i & \text{with probability} & \frac{1}{m} \frac{p - p_l^i}{p_h^i - p_l^i} \\
p_l^i & \text{with probability} & \frac{1}{m} \frac{p_h^i - p}{p_h^i - p_l^i} \\
\dots & \dots & \dots \\
\dots & \dots & \dots \\
p_h^m & \text{with probability} & \frac{1}{m} \frac{p - p_l^m}{p_h^m - p_l^m} \\
p_l^m & \text{with probability} & \frac{1}{m} \frac{p_h^m - p}{p_h^m - p_l^m}
\end{array}$$

By [Lemma A.6](#), this is outcome equivalent to generating $\min_i p_l^i$ and $\max_i p_h^i$. \square

Finally, we show that for all prior distributions for which no disclosure is not the optimal strategy, we can always find the longest interval.

Lemma A.9. *If \mathbb{C} is not optimal for belief p , there exists $p_l < p$ and $p_h > p$ such that it is optimal to induce posteriors p_l and p_h . There exists no other optimal policy with p'_l and p'_h such that $0 \leq p'_l < p_l$ or $p_h < p'_h \leq 1$.*

Proof. Since no disclosure is not optimal for belief p , by Lemma 1, we assume it is optimal to induce p_L and p_H . Hence the linear function goes through $(p_L, V(p_L))$ and $(p_H, V(p_H))$ is no less than the original function V everywhere in the interval $[0, 1]$,

$$V(p) \leq \frac{p - p_L}{p_H - p_L} V(p_H) + \frac{p_H - p}{p_H - p_L} V(p_L) \quad (29)$$

If for some p , $V(p) < \frac{p - p_L}{p_H - p_L} V(p_H) + \frac{p_H - p}{p_H - p_L} V(p_L)$, then it cannot be included in the optimal posteriors profile. In other words, only solution to the following equation can be introduced in optimal posteriors profile,

$$V(p) = \frac{p - p_L}{p_H - p_L} V(p_H) + \frac{p_H - p}{p_H - p_L} V(p_L) \quad (30)$$

Assume this equation has multiple solutions $p_1, \dots, p_m \in [0, 1]$. Let p_l and p_h be the minimum and maximum. If $p_l = 0$ or $p_h = 1$, it is trivial we cannot find such p'_l and p'_h . If $p_l > 0$, then

$V(p'_l) < \frac{p-p_L}{p_H-p_L}V(p_H) + \frac{p_H-p}{p_H-p_L}V(p_L)$ for all $p'_l \in [0, p_l)$ because p_l is the minimum solution. If $p_h < 1$, then $V(p'_h) < \frac{p-p_L}{p_H-p_L}V(p_H) + \frac{p_H-p}{p_H-p_L}V(p_L)$ for all $p'_h \in (p_h, 1]$ because p_l is the maximum solution. \square

Consider the following divide and conquer algorithm for $p \in [L, H]$. Any posteriors that are not belonging to $[L, H]$ cannot be introduced in optimal posteriors profile for prior $p \in [L, H]$. We begin with an arbitrary prior $p \in [L, H]$.

Case (1). If \mathbb{C} is optimal, then we can find the largest connected interval $[p_l, p_h]$ including p that \mathbb{C} is optimal by trials, and run the algorithm in sub-interval $[L, p_l]$ and $[p_h, H]$.

Case (2). If \mathbb{C} is not optimal, we can always find p_l and p_h satisfying the condition of Lemma 4. It is optimal to induce p_l and p_h when $p \in [p_l, p_h]$ and never optimal for $p \in [L, p_l) \cup (p_h, H]$ to induce any posterior in $[p_l, p_h]$.

Then, any posteriors that are not belongs to $[L, p_l)$ cannot be introduced in optimal posteriors profile for prior in $[L, p_l]$; any posteriors that are not belongs to $[p_h, H]$ cannot be introduced in optimal posteriors profile for prior in $(p_h, H]$. Then, we run the algorithm for both sub-interval, $[L, p_l)$ and $(p_h, H]$. The following pseudo-code illustrates the procedure,

Algorithm 1 Find Cutoff Values Inside $[L, H]$

```

1: For arbitrary  $p$ 
2: Find largest interval  $[p_l, p_h]$  that has same strategy with  $p$ .
3: if  $p_l = L$  and  $p_h = H$  then
4:   return
5: end if
6: if  $p_l > L$  then
7:   Report  $p_l$  as cutoff values.
8:   Find Cutoff Values Inside  $[L, p_l]$ .
9: end if
10: if  $p_h < H$  then
11:   Report  $p_h$  as cutoff values.
12:   Find Cutoff Values Inside  $[p_h, H]$ .
13: end if

```

\square

Proof of Lemma 3.4. The concave closure $\hat{V}_{BR}(p)$ (or $\hat{V}_{AR}(p)$) is the smallest concave function that is everywhere weakly greater than $F_R(p)$ (or $1 - F_R(p)$). Therefore, $\hat{V}_{BR}(p) \geq p$ for $p \in (0, 1)$ while $\hat{V}_{AR}(p) \geq 1 - p$ for $p \in (0, 1)$. Assume there exists a $p' \in (0, 1)$ such that $F_R(p') \neq p'$. If $F_R(p') > p'$, then $\hat{V}_{BR}(p) > p$ for all $p \in (0, 1)$. Since $\hat{V}_{AR}(p) \geq 1 - p$ for $p \in (0, 1)$, $\hat{V}_{BR}(p) + \hat{V}_{AR}(p) > 1$ for $p \in (0, 1)$. If $F_R(p') < p'$, then $\hat{V}_{AR}(p) > 1 - p$ for all $p \in (0, 1)$. Since $\hat{V}_{BR}(p) \geq p$ for $p \in (0, 1)$, $\hat{V}_{BR}(p) + \hat{V}_{AR}(p) > 1$ for $p \in (0, 1)$. \square

Proof of Lemma 3.5. Given that $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$ is concave increasing, $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$ is strictly increasing when $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)]) \in [0, R_n + u_n(B, \beta))$. Therefore, $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|B, \theta_n, \theta_R)])$

is strictly increasing when its value is inside the interval $[u_n(B, \beta), R_n]$. Since $\mathbf{Co}(\mathbb{E}_{\theta_R}[U_n(p|A, \theta_n, \theta_R)])$ is concave decreasing, the solution exists uniquely. \square

Proof of Proposition 3.6. When $F_R(p) = p$, we have $V_{BR}(p) = \hat{V}_{BR}(p) = p$ and $V_{AR}(p) = \hat{V}_{AR}(p) = 1 - p$. By induction, we can prove that $V_{Bj} = \hat{V}_{Bj} = p$ and $V_{Aj}(p) = \hat{V}_{Aj}(p) = 1 - p$ for all j . \square

Proof of Lemma 7.1. By contrast we assume the receiver is not fully informed when observing the signal generating by intermediaries. Then either the B-preferred intermediary or the A-preferred intermediary will benefit from full disclosure because of the following reason. If for some signal realizations, some posterior $p \in (0, 1)$ is induced, and under this posterior the receiver chooses one action, then the intermediary that prefers the opposite action will benefit from a full disclosure based on those signal realizations. \square

Proof of Theorem 7.2. Assume the posterior generated by realization s is p , then the utility for the sender is,

$$u(p) = \begin{cases} 0 & 0 \leq p < \tilde{p}_P^{\min} \\ p & \tilde{p}_P^{\min} \leq p < \tilde{p}_P^{\max} \\ 1 & \tilde{p}_P^{\max} \leq p \leq 1 \end{cases} \quad (31)$$

The concave closure of $u(p)$ is

$$\mathbf{Co}u(p) = \begin{cases} \frac{p}{\tilde{p}_P^{\max}} & 0 \leq p \leq \tilde{p}_P^{\max} \\ 1 & \tilde{p}_P^{\max} \leq p \leq 1 \end{cases} \quad (32)$$

\square

B Discussion of the Role of the Reputation Term R_j

The utility for intermediary j when the response is $r \in \{A, B\}$, the final decision is $d \in \{A, B\}$ and the state is $T \in \{\alpha, \beta\}$, is defined as

$$R_j \mathbb{I}(r = d) + u_j(d, t) \quad (33)$$

Intermediary j maximizes the probability of his or her preferred action if and only if intermediary j prefers his or her preferred action for all possible posteriors $p \in [0, 1]$,

$$\forall r, \forall p, R_j + pu_j(r, \beta) + (1 - p)u_j(r, \alpha) \geq pu_j(-r, \beta) + (1 - p)u_j(-r, \alpha) \quad (34)$$

where $-r = \begin{cases} A & r = B \\ B & r = A \end{cases}$. This equation is further equivalent to:

$$\forall r, \forall t, R_j + u_j(r, t) \geq u_j(-r, t) \quad (35)$$

Otherwise, intermediary j prefers A at state α and prefers B at state β , which makes full disclosure the best strategy.

If preferred action of intermediary j is A , then, it is required that

$$\begin{aligned} R_j + u_j(A, \alpha) &\geq u_j(B, \alpha) \\ R_j + u_j(A, \beta) &\geq u_j(B, \beta) \end{aligned}$$

The first equation holds because $u_j(B, \alpha) < u_j(A, \alpha)$. The second equation implies $R_j > u_j(B, \beta)$.

If preferred action of intermediary j is B , then, it is required that

$$\begin{aligned} R_j + u_j(B, \alpha) &\geq u_j(A, \alpha) \\ R_j + u_j(B, \beta) &\geq u_j(A, \beta) \end{aligned}$$

The second equation holds because $u_j(B, \beta) > u_j(A, \beta)$. The first equation implies $R_j > -u_j(B, \alpha)$.

Therefore, we need Assumption 2.1.

C Comparison Between \tilde{p}_j and \tilde{p}_{Ij}

Since \tilde{p}_j is defined by $u_j(B, \alpha)$ and $u_j(B, \beta)$:

$$\tilde{p}_j = \frac{-u_j(B, \alpha)}{u_j(B, \beta) - u_j(B, \alpha)} \quad (36)$$

At the same time, the modified threshold belief is solved by

$$U_j(p|A) = U_j(p|B) \quad (37)$$

where

$$U_j(p|A) = \begin{cases} R_j & p \in [0, \tilde{p}_j^{\min}] \\ \frac{1-p}{1-\tilde{p}_j^{\min}} R_j + \frac{p-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}} u_j(B, \beta) & p \in [\tilde{p}_j^{\min}, 1] \end{cases}$$

$$U_j(p|B) = \begin{cases} \frac{p}{\tilde{p}_j^{\max}} (R_j + u_j(B, \alpha)) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [0, \tilde{p}_j^{\max}] \\ R_j + u_j(B, \alpha) + p(u_j(B, \beta) - u_j(B, \alpha)) & p \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

If we plug \tilde{p}_j into above two equations, we will have

$$U_j(\tilde{p}_j|A) = \begin{cases} R_j & \tilde{p}_j \in [0, \tilde{p}_j^{\min}] \\ \frac{1-\tilde{p}_j}{1-\tilde{p}_j^{\min}} R_j + \frac{\tilde{p}_j-\tilde{p}_j^{\min}}{1-\tilde{p}_j^{\min}} u_j(B, \beta) & \tilde{p}_j \in [\tilde{p}_j^{\min}, 1] \end{cases}$$

$$U_j(\tilde{p}_j|B) = \begin{cases} \frac{\tilde{p}_j}{\tilde{p}_j^{\max}} (R_j + u_j(B, \alpha)) - u_j(B, \alpha) & \tilde{p}_j \in [0, \tilde{p}_j^{\max}] \\ R_j & \tilde{p}_j \in [\tilde{p}_j^{\max}, 1] \end{cases}$$

$\tilde{p}_{Ij} < \tilde{p}_j$ is equivalent to

$$U_j(\tilde{p}_{Ij}|B) > U_j(\tilde{p}_{Ij}|A) \quad (38)$$

because \tilde{p}_{Ij} solves $U_j(p|B) = U_j(p|A)$ and $U_j(p|B)$ ($U_j(p|A)$) is increasing (decreasing).

From now on, we set $\mathbb{A} = -u_j(B, \alpha)$, $\mathbb{B} = u_j(B, \beta)$.

Case (1) When $\tilde{p}_j \leq \tilde{p}_j^{\min}$, $\tilde{p}_{Ij} > \tilde{p}_j$:

$$\begin{aligned} & U_j(\tilde{p}_{Ij}|B) - U_j(\tilde{p}_{Ij}|A) \\ &= \frac{\mathbb{A}/(\mathbb{A} + \mathbb{B})}{\tilde{p}_j^{\max}} (R_j - \mathbb{A}) + \mathbb{A} - R_j \\ &= \frac{\mathbb{A}/(\mathbb{A} + \mathbb{B}) - \tilde{p}_j^{\max}}{\tilde{p}_j^{\max}} (R_j - \mathbb{A}) \\ &= \frac{\tilde{p}_j - \tilde{p}_j^{\max}}{\tilde{p}_j^{\max}} (R_j - \mathbb{A}) \\ &< 0 \end{aligned} \quad (39)$$

Case (2) When $\tilde{p}_j \geq \tilde{p}_j^{\max}$, $\tilde{p}_{Ij} < \tilde{p}_j$:

$$\begin{aligned}
& U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) \\
&= R_j - \frac{\mathbb{B}/(\mathbb{A} + \mathbb{B})}{1 - \tilde{p}_j^{\min}} R_j - \frac{\mathbb{A}/(\mathbb{A} + \mathbb{B}) - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} \mathbb{B} \\
&= \frac{\mathbb{A}/(\mathbb{A} + \mathbb{B}) - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} (R_j - \mathbb{B}) \\
&= \frac{\tilde{p}_j - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} (R_j - \mathbb{B}) \\
&> 0
\end{aligned} \tag{40}$$

Case (3) When $\tilde{p}_j \in (\tilde{p}_j^{\min}, \tilde{p}_j^{\max})$

$$U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) = \frac{\tilde{p}_j - \tilde{p}_j^{\max}}{\tilde{p}_j^{\max}} (R_j - \mathbb{A}) + \frac{\tilde{p}_j - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} (R_j - \mathbb{B}) \tag{41}$$

Then,

$$\begin{aligned}
& U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0 \\
\iff & \frac{\tilde{p}_j - \tilde{p}_j^{\min}}{1 - \tilde{p}_j^{\min}} (R_j - \mathbb{B}) > \frac{\tilde{p}_j^{\max} - \tilde{p}_j}{\tilde{p}_j^{\max}} (R_j - \mathbb{A}) \\
\iff & \frac{R_j - \mathbb{B}}{R_j - \mathbb{A}} > \frac{(\tilde{p}_j^{\max} - \tilde{p}_j)(1 - \tilde{p}_j^{\min})}{\tilde{p}_j^{\max}(\tilde{p}_j - \tilde{p}_j^{\min})}
\end{aligned}$$

Since $\frac{(\tilde{p}_j^{\max} - \tilde{p}_j)(1 - \tilde{p}_j^{\min})}{\tilde{p}_j^{\max}(\tilde{p}_j - \tilde{p}_j^{\min})}$ is monotonic decreasing when $\tilde{p}_j \in (\tilde{p}_j^{\min}, \tilde{p}_j^{\max})$ because numerator is positive and decreasing while denominator is positive and increasing.

What's more,

$$\begin{aligned}
\lim_{\tilde{p}_j \rightarrow \tilde{p}_j^{\min} +} \frac{(\tilde{p}_j^{\max} - \tilde{p}_j)(1 - \tilde{p}_j^{\min})}{\tilde{p}_j^{\max}(\tilde{p}_j - \tilde{p}_j^{\min})} &= +\infty \\
\lim_{\tilde{p}_j \rightarrow \tilde{p}_j^{\max} -} \frac{(\tilde{p}_j^{\max} - \tilde{p}_j)(1 - \tilde{p}_j^{\min})}{\tilde{p}_j^{\max}(\tilde{p}_j - \tilde{p}_j^{\min})} &= 0
\end{aligned}$$

hence, we can define $q = \frac{(\tilde{p}_j^{\max} - \tilde{p}_j)(1 - \tilde{p}_j^{\min})}{\tilde{p}_j^{\max}(\tilde{p}_j - \tilde{p}_j^{\min})} \in (0, \infty)$. We now consider as R_j changes, whether $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ or not.

If $\mathbb{A} = \mathbb{B}$, then left hand side $\frac{R_j - \mathbb{B}}{R_j - \mathbb{A}}$ is always 1. $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ if and only if $q < 1$.

If $\mathbb{A} > \mathbb{B}$, then left hand side $\frac{R_j - \mathbb{B}}{R_j - \mathbb{A}} \in (1, \infty)$. If $q \leq 1$, then $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ always holds. If $q > 1$, then $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ holds if and only if $R_j < \frac{\mathbb{A} - \mathbb{B}}{q - 1}$.

If $\mathbb{A} < \mathbb{B}$, then left hand side $\frac{R_j - \mathbb{B}}{R_j - \mathbb{A}} \in (0, 1)$. If $q \geq 1$, then $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ never holds. If $q < 1$, then $U_j(\tilde{p}_j|B) - U_j(\tilde{p}_j|A) > 0$ holds if and only if $R_j > \frac{\mathbb{B} - \mathbb{A}}{1 - q}$.

We summarize the results in the following three tables,

R_j	$\mathbb{A} > B$	$\mathbb{A} = B$	$\mathbb{A} < B$
$q < 1$	Always	Always	$R_j > \frac{B - \mathbb{A}}{1 - q}$
$q = 1$	Always	Never	Never
$q > 1$	$R_j < \frac{\mathbb{A} - B}{q - 1}$	Never	Never

Table 2: When $\tilde{p}_{I_j} < \tilde{p}_j$ Holds

R_j	$\mathbb{A} > B$	$\mathbb{A} = B$	$\mathbb{A} < B$
$q < 1$	Never	Never	$R_j = \frac{B - \mathbb{A}}{1 - q}$
$q = 1$	Never	Always	Never
$q > 1$	$R_j = \frac{\mathbb{A} - B}{q - 1}$	Never	Never

Table 3: When $\tilde{p}_{I_j} = \tilde{p}_j$ Holds

R_j	$\mathbb{A} > B$	$\mathbb{A} = B$	$\mathbb{A} < B$
$q < 1$	Never	Never	$R_j < \frac{B - \mathbb{A}}{1 - q}$
$q = 1$	Never	Never	Always
$q > 1$	$R_j > \frac{\mathbb{A} - B}{q - 1}$	Always	Always

Table 4: When $\tilde{p}_{I_j} = \tilde{p}_j$ Holds

D Details about Persuading the Intermediary j

We first define the expected utility for intermediary j with the following features

1. Responding A/B to intermediary $j - 1$
2. Intermediary j generates posterior p to the Intermediary $j + 1$.
3. Type of intermediary j is θ_j
4. Type of intermediary $j + 1$ is θ_{j+1}

by the expressions

$$\begin{aligned} U_j(p|A, \theta_j, \theta_{j+1}) &= (1 - \mathcal{I}_{j+1}(p|\theta_{j+1}))R_j + \mathcal{I}_{j+1}(p|\theta_{j+1})(pu_j(B, \beta|\theta_j) + (1 - p)u_j(B, \alpha|\theta_j)) \\ U_j(p|B, \theta_j, \theta_{j+1}) &= \mathcal{I}_{j+1}(p|\theta_{j+1})(R_j + pu_j(B, \beta|\theta_j) + (1 - p)u_j(B, \alpha|\theta_j)) \end{aligned}$$

respectively, where boolean variable $\mathcal{I}_{j+1}(p|\theta_{j+1}) = I(\tilde{p}_{I,j+1}(\theta_{j+1}) \leq p) \in \{0, 1\}$ represents whether receiver will choose action B or not when intermediary j passes posterior p to intermediary $j + 1$. Firstly, it is easy to see that $U_j(p|B, \theta_j, \theta_{j+1})$ is increasing in p because when p increases, both $\mathcal{I}_{j+1}(p|\theta_{j+1})$ and $R_j + pu_j(B, \beta|\theta_j) + (1 - p)u_j(B, \alpha|\theta_j)$ are nonnegative and increasing. Secondly, we have the following boundary values, which help us prove the existence and uniqueness of the modified threshold belief.

$$\begin{aligned} U_j(0|A, \theta_j, \theta_{j+1}) &= R_j \\ U_j(1|A, \theta_j, \theta_{j+1}) &= u_j(B, \beta) \\ U_j(0|B, \theta_j, \theta_{j+1}) &= 0 \\ U_j(1|B, \theta_j, \theta_{j+1}) &= R_j + u_j(B, \beta) \end{aligned}$$

In addition, we have $U_j(p|A, \theta_j, \theta_{j+1}) \leq R_j - \mathcal{I}_{j+1}(p|\theta_{j+1})(R_j - u_j(B, \beta)) \leq R_j$. Therefore, if we take the expectation over θ_{j+1} , we have the expected utility for A-preferred and B-preferred intermediaries j given his type is θ_j under two extreme posteriors (0 and 1),

$$\begin{aligned} \mathbb{E}_{\theta_{j+1}}[U_j(0|A, \theta_j, \theta_{j+1})] &= R_j \\ \mathbb{E}_{\theta_{j+1}}[U_j(1|A, \theta_j, \theta_{j+1})] &= u_j(B, \beta) \\ \mathbb{E}_{\theta_{j+1}}[U_j(0|B, \theta_j, \theta_{j+1})] &= 0 \\ \mathbb{E}_{\theta_{j+1}}[U_j(1|B, \theta_j, \theta_{j+1})] &= R_j + u_j(B, \beta) \end{aligned}$$

Then by Bayesian persuasion, for an A-preferred player, the best he can achieve is the concave closure of $\mathbb{E}_{\theta_{j+1}}[U_j(0|A, \theta_j, \theta_{j+1})]$, $\mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|A, \theta_j, \theta_{j+1})])$; for B-preferred player, the best he

can achieve is the concave closure of $\mathbb{E}_{\theta_{j+1}}[U_j(0|B, \theta_j, \theta_{j+1})]$, $\mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|B, \theta_j, \theta_{j+1})])$. We then need to compare whether response A or B is better.

By similar demonstration as persuading last intermediary, $\mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|A, \theta_j, \theta_{j+1})])$ is decreasing in p while $\mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|B, \theta_j, \theta_{j+1})])$ is increasing in p . Therefore, *single crossing property* still holds.

Lemma D.1. $\mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|A, \theta_j, \theta_{j+1})]) = \mathbf{Co}(\mathbb{E}_{\theta_{j+1}}[U_j(p|B, \theta_j, \theta_{j+1})])$ has unique solution $p = \tilde{p}_{I_j}(\theta_j)$.

Hence, there exists a threshold value $\tilde{p}_{I_j}(\theta_j)$ for type θ_j , called modified threshold belief such that intermediary j is indifferent between choosing A and B in response stage.

Proof of Theorem 5.2. Since the relationship among $\tilde{p}_{k+1}^A, \tilde{p}_{k+1}^B, \tilde{p}_{k+1}^{\min}, \tilde{p}_{k+1}^{\max}$ has six possible outcomes,

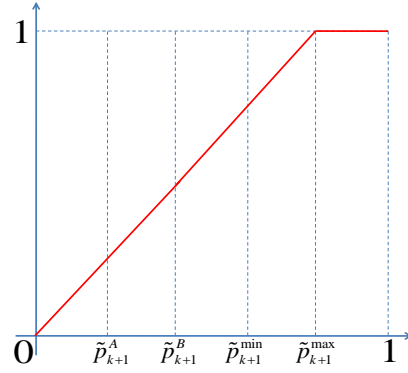

$$\begin{array}{ll} \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} & \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max} \\ \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B & \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max} \\ \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B & \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B \end{array}$$

In the induction process, we need to verify that these are indeed equilibrium strategies in 12 different cases which considering the possible preferred actions for intermediary j . The 12 different cases are



- **B-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$; **A-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$
- **B-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$; **A-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$
- **B-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$; **A-preferred**, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$
- **B-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$; **A-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$
- **B-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$; **A-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$
- **B-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B$; **A-preferred**, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B$

Due to the nature of concavification, the boundary cases are included. In the induction processes, all 12 cases can be verified. For concavification of those different cases, please see the following graphical illustrations. Intermediary k avoids to induce any posterior between \tilde{p}_{k+1}^A and \tilde{p}_{k+1}^B .



Case (1): B-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$.





Case (2): A-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max}$





Case (3): B-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$





Case (4): A-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$





Case (5): B-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$





Case (6): A-preferred, $\tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$





Case (7): B-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$.





Case (8): A-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max}$.





Case (9): B-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$.



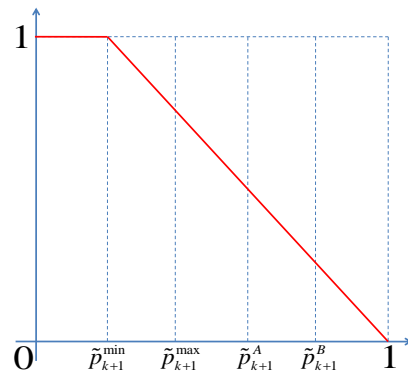

Case (10): A-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^B$



Case (11): B-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B$



Case (12): A-preferred, $\tilde{p}_{k+1}^{\min} < \tilde{p}_{k+1}^{\max} < \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^B$



Finally, the following relationships hold between intermediaries,

$$\tilde{p}_k^{\max} = \begin{cases} \tilde{p}_{k+1}^{\max} & \tilde{p}_{k+1}^B < \tilde{p}_{k+1}^{\max} \\ \tilde{p}_{k+1}^B & \tilde{p}_{k+1}^B > \tilde{p}_{k+1}^{\max} \end{cases} \quad (42)$$

$$\tilde{p}_k^{\min} = \begin{cases} \tilde{p}_{k+1}^A & \tilde{p}_{k+1}^A < \tilde{p}_{k+1}^{\min} \\ \tilde{p}_{k+1}^{\min} & \tilde{p}_{k+1}^A > \tilde{p}_{k+1}^{\min} \end{cases} \quad (43)$$

□

E Operational Process of Example 6.2

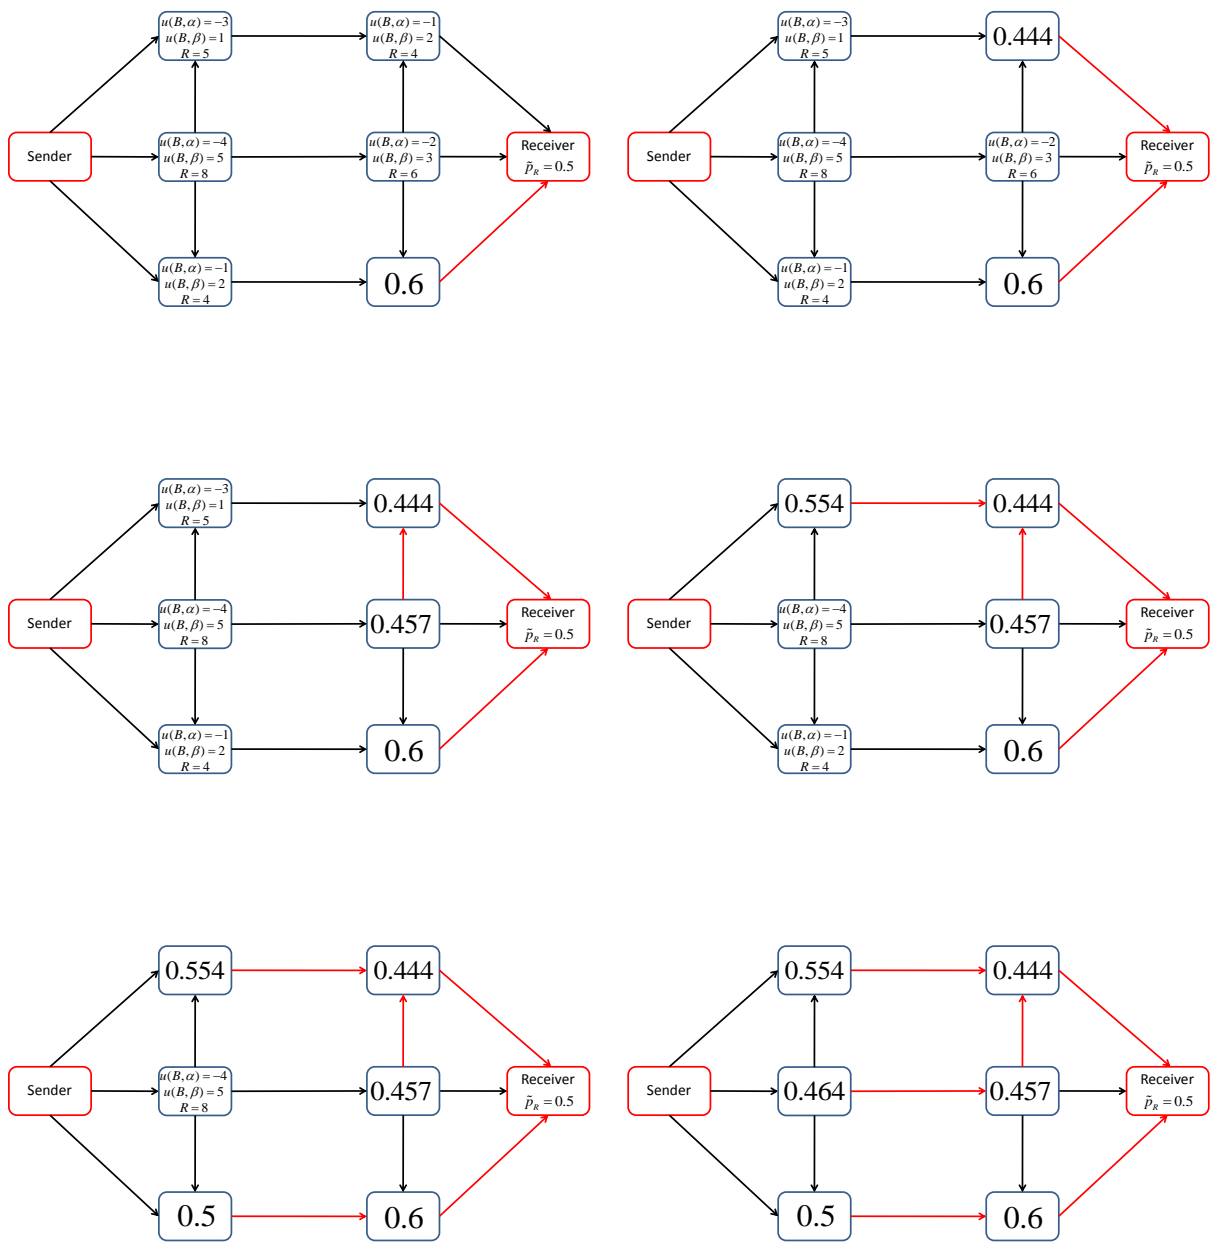


Figure 14: Running Algorithm

F Contents about Shortsighted Player

F.1 Benchmark Model

The results of the previous analysis only depend on the modified threshold beliefs of players and the equilibrium analysis is based on the following standard assumptions,

- The game theoretic rationality of all players are common knowledge.
- The parameters of all players are common knowledge.
- The computational power of all players is unrestricted.

Each player can then calculate their own modified threshold beliefs based on the subsequent players' modified threshold beliefs accordingly. Under alternative conditions, such as limited computational ability, this may not be the case. Indeed, the forward-looking and backward induction requirement on players in the benchmark model is arguably high.

When a player is uncertain or ambiguous about the subsequent players, it is natural to respond heuristically, such as according to a standard threshold belief. Previous results derived in the benchmark model are qualitatively similar if standard threshold beliefs are used to substitute for modified threshold beliefs. By short-sighted, it means that player maximized the likelihood of his preferred action being implemented by the receiver.

While the exact persuasion strategies differ from the benchmark model, the structure of persuasion based on the magnitudes of own and subsequent intermediaries' threshold beliefs, maintains the same threshold characteristics described in the one-step equilibrium.

F.2 Choosing Persuasion Path

A modified version of Dijkstra's algorithm is applied to find the optimal persuasion **Path**, as well as resulting path threshold, $\phi(r)$, as it is shown in the algorithm **FindPath-S**(G, s, r). Here, $\phi(x)$ denotes the maximum threshold belief when the persuasion destination is x .

The interpretation of the algorithm is as follows,

- **Line 1** initializes the values of $\phi(\cdot)$.
- Each iteration in **Line 2-11** adds some node into Q :
 - **Line 3** generates all vertices that not included in Q but only one step from Q in set N .
 - Each iteration in **Line 4-7** computes $\phi(n)$ for all $n \in N$ and records its preceding vertices.
 - **Line 8-10** adds all vertices with minimax threshold belief in N into Q
- **Line 12** initializes the recovering of the optimal path.

Algorithm 2 FindPath-S(G, s, r)

```
1:  $Q = \{s\}, \phi(s) = 0$  and  $\forall v \neq s, \phi(v) = +\infty$ .
2: while  $r \notin Q$  do
3:    $N = \{v \notin Q \mid \exists q \in Q, (q, v) \in E\}$ .
4:   for  $n \in N$  do
5:      $c = \min(\phi(q) \mid q \in Q, (q, n) \in E)$ .
6:      $\mathbf{Pre}(n) = \{q \in Q \mid \phi(q) = c\}$ 
7:      $\phi(n) = \max(\tilde{p}_n, c)$ .
8:   end for
9:   for  $u = \arg \min_n \phi(n)$  do
10:     $Q = Q \cup \{u\}$ 
11:   end for
12: end while
13:  $x = r$ , add  $x$  in Path.
14: while  $x \neq s$  do
15:   Find any element  $y \in \mathbf{Pre}(x)$ .
16:    $y = x$ .
17:   Add  $x$  at the head of Path.
18: end while
19: return Path and  $\phi(r)$ .
```

- Each iteration in **Line 13-17** backtracks one step.
- **Line 18** outputs the result.

Theorem F.1. *Algorithm **FindPath-S**(G, s, r) returns an optimal persuasion path. All optimal persuasion paths can be generated from the algorithm. The algorithm completes in polynomial time.*

The validity and efficiency can be proved by the same method as Dijkstra’s algorithm. The reader is referred to the last subsection for the detailed proof.

The main intuition behind our algorithm is that if direct communication is unavailable, then the optimal persuasion may not choose the shortest path. A sequence of friendly intermediaries can outperform a path that includes a stubborn intermediary.

F.3 Details about Dijkstra’s Algorithm

Introduction Dijkstra’s algorithm is an algorithm for finding the shortest paths between nodes in a graph. This algorithm was proposed by computer scientist Edsger W. Dijkstra in 1956 and published three years later. For a given source node in the graph, the algorithm iteratively finds the shortest path between that node and every other. Hence, by stopping the algorithm once the shortest path to the destination node has been reached, the algorithm finds the shortest path between s (sender node) and r (receiver node).

Distance However, the definition of distance is different from the standard framework in shortest path problem. For a directed graph $G = (V, E)$, we define the distance from s to v through one path as the maximum modified threshold believes of nodes that involved in such path. The proof of correctness are inductive.

Hypothesis For each node $v \in Q$, $\phi(v)$ is considered the minimax (modified) threshold belief from s to v ; and for each unvisited node u , $\phi(u)$ is assumed the minimax (modified) threshold belief from s to u when traveling via visited nodes only. This assumption is only considered if a path exists, otherwise the distance is set to infinity.

Base Case When there is just one visited node, namely the initial node s , in which case the hypothesis is trivial.

Induction Assume the hypothesis for $n - 1$ visited nodes. In which case, we choose an edge (v, u) where u has the least $\phi(u)$ (**Line 8-10**) of any unvisited nodes and the edge (v, u) is such that $\phi(u) = \max(\phi(v), \tilde{p}(u))$ (**Line 6**). $\phi(u)$ is considered to be the minimax (modified) threshold belief from s to u because if there were a path with shorter path, and if w was the first unvisited node on that path then by the original hypothesis $\phi(w) > \phi(u)$ which creates a contradiction. Similarly if there was a shorter path to u without using unvisited nodes, and if the last but one node on that path were w , then we would have had $\phi(u) = \max(\phi(w), \tilde{p}(u))$, also a contradiction. After processing u it will still be true that for each unvisited nodes w , $\phi(w)$ will be the shortest distance from s to w using visited nodes only, because if there were a shorter path that doesn't go by u we would have found it previously, and if there were a shorter path using u we would have updated it when processing u .

References

- [1] E. Kamenica and M. Gentzkow, “Bayesian persuasion,” *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.
- [2] D. Bergemann and S. Morris, “Information design: A unified perspective,” *Journal of Economic Literature*, vol. 57, no. 1, pp. 44–95, 2019.
- [3] E. Kamenica, “Bayesian persuasion and information design,” *Annual Review of Economics*, vol. forthcoming, 2019.
- [4] M. Gentzkow and E. Kamenica, “Bayesian persuasion with multiple senders and rich signal spaces,” *Games & Economic Behavior*, vol. 104, 2017.
- [5] M. Gentzkow and E. Kamenica, “Competition in persuasion,” *Review of Economic Studies*, vol. 84, pp. 300–322, 2017.
- [6] L. Fei and N. Peter, “On bayesian persuasion with multiple senders,” *Economics Letters*, vol. 170, pp. 66–70, 2018.
- [7] F. Li and P. Norman, “Sequential persuasion,” *Working Paper*, 2018.
- [8] R. Alonso and O. Camara, “Bayesian persuasion with heterogeneous priors,” *Journal of Economic Theory*, vol. 165, pp. 672–706, 2016.
- [9] R. Alonso and O. Camara, “Persuading voters,” *American Economic Review*, vol. 106, no. 11, pp. 3590–3605, 2016.
- [10] I. Arieli and Y. Babichenko, “Private bayesian persuasion,” *Working Paper*, 2016.
- [11] A. Bardhi and Y. Guo, “Modes of persuasion toward unanimous consent,” *Theoretical Economics*, vol. 13, no. 3, pp. 1111–1150, 2018.
- [12] J. Chan, S. Gupta, F. Li, and Y. Wang, “Pivotal persuasion,” *Journal of Economic Theory*, vol. 180, pp. 178–202, 2019.
- [13] T. Hoshino, “Using strategic uncertainty to persuade multiple agents,” *Working Paper*, 2017.
- [14] L. Marie and R. Ludovic, “Public persuasion,” *Working Paper*, 2016.
- [15] B. Michaeli, “Divide and inform: Rationing information to facilitate persuasion,” *The Accounting Review*, vol. 92, no. 5, 2017.
- [16] M. Shimoji, “Rationalizable persuasion,” *Working Paper*, 2016.
- [17] S. Yangbo and Z. Mofei, “Designing noise in persuasion,” *Working Paper*, 2019.
- [18] Y. Wang, “Bayesian persuasion with multiple receivers,” *Working Paper*, 2015.

- [19] S. Board and J. Lu, “Competitive information disclosure in search markets,” *Journal of Political Economy*, vol. 126, no. 5, pp. 1965–2010, 2018.
- [20] F. Koessler, M. Laclau, and T. Tomala, “Interactive information design,” *Working Paper*, 2018.
- [21] F. Wu and J. Zheng, “Double directional persuasion,” *Working Paper*, 2019.
- [22] K. Andrew, “Mediated persuasion,” *Working Paper*, 2018.
- [23] Y. Qian, “Persuasion through a strategic moderator,” *Working Paper*, 2019.
- [24] L. Doval and J. Ely, “Sequential information design,” *Working Paper*, 2019.
- [25] J. C. Ely, “Beeps,” *American Economic Review*, vol. 107, no. 1, pp. 31–53, 2017.
- [26] J. Ely, A. Frankel, and E. Kamenica, “Suspense and surprise,” *Journal of Political Economy*, vol. 123, no. 1, pp. 215–260, 2015.
- [27] J. Ely and M. Szydlowski, “Moving the goalposts,” *Journal of Political Economy*, vol. forthcoming, 2019.
- [28] M. Felgenhauer and P. Loerke, “Bayesian persuasion with private experimentation,” *International Economic Review*, vol. 58, no. 3, pp. 829–856, 2017.
- [29] E. Henry and M. Ottaviani, “Research and the approval process: the organization of persuasion,” *American Economic Review*, vol. 109, no. 3, pp. 911–955, 2019.
- [30] J. Horner and A. Skrzypacz, “Selling information,” *Journal of Political Economy*, vol. 124, no. 6, pp. 1515–1562, 2016.
- [31] D. Orlov, A. Skrzypacz, and P. Zryumov, “Persuading the principal to wait,” *Working Paper*, 2018.
- [32] J. Renault, E. Solan, and N. Vieille, “Optimal dynamic information provision,” *Games & Economic Behavior*, vol. 104, pp. 329–349, 2017.
- [33] V. P. Crawford and J. Sobel, “Strategic information transmission,” *Econometrica*, vol. 50, no. 6, pp. 1431–1451, 1982.
- [34] M. Ivanov, “Communication via a strategic mediator,” *Journal of Economic Theory*, vol. 145, no. 2, pp. 869–884, 2010.
- [35] A. Ambrus, E. M. Azevedo, and Y. Kamada, “Hierarchical cheap talk,” *Theoretical Economics*, vol. 8, no. 1, pp. 233–261, 2013.
- [36] A. Kolotilin, T. Mylovanov, A. Zapechelnuk, and L. Ming, “Persuasion of a privately informed receiver,” *Econometrica*, vol. 85, no. 6, pp. 1949–1964, 2017.

- [37] A. Kolotilin, “Optimal information disclosure: A linear programming approach,” *Theoretical Economics*, vol. 13, no. 2, pp. 607–635, 2018.
- [38] B. James and Q. Daniel, “Persuasion for the long run,” *Working Paper*, 2017.
- [39] M. Gentzkow and E. Kamenica, “Costly persuasion,” *American Economic Review*, vol. 104, no. 5, pp. 457–462, 2014.
- [40] R. Hodler, S. Loertscher, and D. Rohner, “Persuasion, binary choice, and the costs of dishonesty,” *Economics Letters*, vol. 124, no. 2, pp. 195–198, 2014.
- [41] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Volume II*. Athena Scientific, 2012.
- [42] S. Bikhchandani, D. Hirshleifer, and I. Welch, “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.