Three Essays on the Economics of Information

Samuel F. Pichette

A Thesis
in
The Department
of
Economics

Presented in Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy (Economics) at
Concordia University
Montréal, Québec, Canada

April 2024

CONCORDIA UNIVERSITY
School of Graduate Studies

This is to certify that the thesis prepared

By: **Samuel F. Pichette**
Entitled: **Three Essays on the Economics of Information**

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Economics)

complies with the regulations of this University and meets the accepted standards with respect to originality and quality.
Signed by the Final Examining Committee:

______________________________ Chair
Dr. Kate de Medeiros

______________________________ External Examiner
Dr. Arianna Degan

______________________________ Examiner
Dr. Axel Watanabe

______________________________ Examiner
Dr. Ming Li

______________________________ Examiner
Dr. Szilvia Pápai

______________________________ Supervisor
Dr. Dipjyoti Majumdar

Approved by ______________________________
Dr. Jorgen Hansen, Chair
Department of Economics

__________ March 15, 2024 ______________________________
Dr. Pascale Sicotte, Dean
Faculty of Arts and Science

Abstract

Three Essays on the Economics of Information

Samuel F. Pichette, Ph.D.
Concordia University, 2024

This thesis consists of three essays on the Economics of Information, focusing on strategic information acquisition and on the design of disclosure rules.

In Chapter 2, a decision-maker relies on the information reported by a panel of experts to take an action. These experts may observe relevant information and have preferences over the decision-maker's actions. Each expert possesses three qualities: (i) the probability that he acquires information, (ii) the probability that this information is inaccurate in favor of his agenda, and (iii) the level of this error. Conscious of these qualities, the decision-maker forms her optimal panel to make the most informed decision. We show that it is generally better for the decision-maker to have experts with identical agendas if these experts cannot misreport their information, even if it is inaccurate.

In Chapter 3, we consider a variation of the model presented in Chapter 2 in which experts' information is always accurate. Moreover, each expert can now incur a cost to increase his probability of obtaining information. The objective of the decision-maker is to form a panel to make the most informed decision, while experts have their own preferences. We show that there exist levels of cost of effort such that experts in a homogeneous panel will choose to exert no effort, while experts in a diverse panel will exert some. In addition, we find that for a sufficiently high return on effort, a diverse panel is optimal.

In Chapter 4, one agent – the sender – benefits from two other agents – the receivers – taking some actions. Each receiver possesses private information referred to as his type, and can disclose it truthfully or not. The sender has a private type which she does not observe at first but can condition a communication device, observed by every agent, on its different realizations and on the reported types. In turn, this provides the receivers with action recommendations. We show that depending on the alignment of the preferences, the optimal rule ranges from fully revealing the sender's type to no revelation at all, and that most of the times only partial information is revealed.

Acknowledgement

First, I would like to express my deepest gratitude to my supervisor, Dr. Dipjyoti Majumdar, for both his academic and personal support. The road was certainly longer and more sinuous than expected, and reaching the end would not have been possible without your invaluable guidance.

I would also like express my appreciation to my thesis defence committee, namely the defence chair, Dr. Kate de Medeiros, the arms-length examiner Dr. Axel Watanabe, the external examiner Dr. Arianna Degan, and the two examiners, Dr. Ming Li and Dr. Szilvia Pápai, for their precious time.

Moreover, I would like to thank all of my classmates, professors, faculty and staff members that I had the privilege to meet during this program, and who have made this endeavor a pleasant one. In the same spirit, I am appreciative for all the support and encouragement that I received from all of my friends.

This thesis would not have been possible without the unwavering support from my family. Most crucially, I am deeply indebted to my parents, Louise and Denis, and my brother, Jacob, for their help and for always believing in me. Special thoughts for my mother who is sadly no longer with us, for your precious time and advice. You have shown me what perseverance and confidence really mean. Special thanks also to my cat, Leia Catwalker, who was there for me when I was, sometimes unknowingly, in need of a change of mind or a smile.

Finally, there is simply no words to describe how grateful I am for the support from my soulmate, Karolan, who left this world unjustifiably too soon. Let alone this thesis, I would not have even contemplated pursuing graduate studies without you. You were my rock, present in every up and down of this journey until the very end.

Dedication

In memory of Karolan Jeffrey.
I will forever cherish the time that I had with you. Always.

"All we have to decide is what to do with the time that is given to us."

— J.R.R. Tolkien, The Fellowship of the Ring

Contribution of Authors

Chapter 2 is not co-authored.

Chapter 3 is a joint work with my supervisor, Dr. Dipjyoti Majumdar.

Chapter 4 is not co-authored.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Optimal Panel Composition of Biased Experts

In the first essay (Chapter 2), we solve for the optimal panel composition of experts for a decision-maker. In our model, the decision-maker relies on the information reported by experts to choose an action. Her goal is to make the most informed decision. Each expert has his own agenda – an action by the decision-maker that he prefers the most. With some probability, an expert can be informed. That is, he observes a signal about the state of the world. However, this signal may be biased towards their own agenda, but the expert is unaware of this possibility. One can interpret this as an expert experiencing a selection bias when gathering information for his report. Once an expert observes this signal, he then decides whether to reveal it or to conceal it to the decision-maker. He cannot misreport it. Moreover, an uninformed expert must admit ignorance. The decision-maker, having received these reports (if any), then takes an action.

We first characterize the best responses for the experts. We use the *Perfect Bayesian Equilibrium* (PBE) solution concept to show that for sufficiently high levels and probabilities of error, a diverse panel – composed of experts with opposing agendas – can be preferable to a homogeneous panel – composed of experts with the same opinion. This result may be viewed as bridging the gap between two strains of the literature, in which one argues that competing interests yield higher levels of information while the other is promoting the benefits of similar agendas in providing accrued validation.

1.2 Strategic Acquisition of Costly Information from Opinionated Experts

In the second essay (Chapter 3), we show that in a setting where experts are accurate – as opposed to Chapter 2 – but instead can incur a cost to increase their probability of being informed – i.e. getting a signal about the state of the world –, then it is possible for the diverse panel to be optimal. This depends on how effective it is for an expert to convert effort into an increased informativeness probability, and how costly this process is. Indeed, we find that if the cost of effort is sufficiently high, and that exerting effort translates into an adequately high probability of being informed, then experts in a diverse panel will agree on incurring it. However, under such conditions, we find a free-riding issue present in the homogeneous panel composition, resulting in no expert putting in more work. Consequently, the diverse panel may be more appealing for the decision-maker.

This result also makes use of the PBE concept. Moreover, we characterize it for different setups. We first begin by considering the decision-maker to be risk-averse and effort to be binary. Then, we generalize our results for a risk-neutral decision-maker. We conclude by validating our results in the case of continuous levels of effort, where we find that the diverse panel can still be optimal under conditions similar to the binary scenario.

1.3 Information Design with Multiple Privately-Informed Receivers

In the third essay (Chapter 4), we consider an environment with three agents: one who is uninformed (sender), and two who are privately informed (receivers). The sender can design a disclosure mechanism that will reveal information about an unknown state of the world to each receiver individually. Both the sender and the receivers know nothing more than the distribution of this state prior to the reports from the disclosure mechanism. Each receiver has, however, payoff-relevant private information – referred to as his type – and must report it, truthfully or not, to the sender. The sender designs her disclosure mechanism in advance, before getting reports from each receiver. She can condition the information revealed – as private messages to each receiver – on the possible realizations of the state and these aforementioned reports. After observing their messages from the mechanism, the receivers chooses their actions. They cannot at any moment communicate their information to other receivers. Moreover, while the receivers' payoffs depend on the state, their types, and their collective profile of actions, the sender's payoff only depend on the receivers' actions and the state of the world.

In a binary environment, we characterize the optimal disclosure rule using the *Incentive-Compatible Bayes Correlated Equilibrium* (ICBCE) concept. The solution illustrates that depending on the payoff structures, the information revealed varies from no disclosure at all – after getting his message, a receiver knows no more information than his prior about the state – to full disclosure – when a receiver knows what the true state is with certainty. We also show that in most cases, under our settings, receivers get partial information about the state.

Chapter 2

Optimal Panel Composition of Biased Experts

2.1 Introduction

The present chapter considers aspects of strategic information transmission when experts compete for influence. Consider a situation where a decision-maker relies on advice from experts. In the classical version of the model, the decision-maker does not have information about the state of the world, but the experts do. Experts however, are interested in the decision and might behave strategically by withholding relevant information. In order to mitigate this issue, decision-makers often form panels of experts with competing interests. The idea is that competition among experts incentivizes them to share more information. There are numerous real life examples where decision-makers solicit information from a panel of experts. One such example consists of a judge who might invite testimony from both the defendant and the plaintiff. Another is depicted by television debates in a presidential primary where voters can listen to the policy positions of various candidates.

Starting with the seminal paper by Milgrom and Roberts (1986), there has been an extensive study on the issue of eliciting private information with competing experts (see for example Shin (1994), Shin (1998), Kamenica and Gentzkow (2011), Gul and Pesendorfer (2012)). Most of the existing literature, however, does not necessarily address the issue of conflict among experts and the quality of decision making. In an important series of papers, Bhattacharya and Mukherjee (2013) and Bhattacharya et al. (2018) tried to bridge this gap. In particular, Bhattacharya and Mukherjee (2013) demonstrate that it is always optimal for the decision-maker to opt for a panel of experts with extreme agendas/preferences, and that choosing similar such agendas is better – i.e. forming a homogeneous panel – than opposing ones – i.e. forming a diverse panel. In a subsequent paper, Bhattacharya et al. (2018) establish that when correlation between experts' types is low, a homogeneous panel is again optimal.

In the present chapter, we take the analysis further. We consider an environment where a decision-maker cannot commit to his actions, and where there is uncertainty over whether or not an expert possesses relevant information. Our objective is to explore the following key question: "Given that experts have extreme agendas, when is it optimal for a decision-maker to opt for a panel with diverse policy preferences instead of homogeneous ones?" In most of the literature, an expert either learns the state or he does not. Conditional on an expert learning the state, there is no ambiguity in the expert's learning. We depart from this regime by allowing experts to make mistakes in learning the state. More importantly, errors that the experts make are "biased" towards their respective agendas. When experts have the possibility of making such "agenda-specific errors", we find that the answer to the question of optimal design of a panel is nuanced. Indeed, our results suggest that for a given probability and "level of error", a diverse panel is better, the higher the quality of the experts – i.e. the probability that each experts observes a signal about the state. Moreover, for a given quality of experts, and for a given level of error, a higher probability of error – beyond a certain threshold – makes a diverse panel optimal for the decision-maker.

We follow closely the model introduced in Bhattacharya and Mukherjee (2013) and consider a persuasion game having the following features. A decision-maker wants to take an action in the unit interval $[0, 1]$. From the decision-maker's point of view, her optimal decision is to match the state of the world $\theta \in [0,1]$. Her objective is to minimize the loss which is an increasing function of the distance between the state and the action. We restrict our attention to the absolute loss function for the decision-maker – also referred to as the judge. The judge bases her action on verifiable reports from two experts. The experts have state-independent and monotonic preferences over the decision-maker's action. More importantly, an expert's preference is identified by his "*agenda* " – his most preferred action – defined as $x_i \in [0,1]$. We restrict our attention to extreme experts with either $x_i = 0$ or $x_i = 1$. Each experts $i \in \{1,2\}$ is privately informed of a signal $s_i \in [0,1]$ regarding the state θ with a probability $\alpha \in [0,1]$ – representing an expert's quality. Conditional on an expert i being informed, the signal can take two possible values: For an expert with agenda at 0 – i.e. for $x_i = 0$ –, with a probability $(1-\epsilon)$, $s_i = \theta$, and with probability $\epsilon > 0$, $s_i = \theta - \delta$, where ϵ and δ are relatively small numbers. This is a departure from the literature, which focused mainly on situations where informed experts learned the true state with certainty. The interpretation is that an informed expert with agenda $x_i = 0$ will, with a high probability, correctly observe the state. However, with a small probability ϵ, the expert makes an error and observes a signal that is removed by a distance δ from the state towards his agenda. We call ϵ the "agenda-specific error probability", and δ the "agenda-specific level of error". Likewise, for an informed expert whose agenda is at 1 – i.e. $x_i = 1$ –, with probability ϵ, the signal is $s_i = \theta + \delta$. We assume that the information is "hard" in the sense that the experts cannot misreport. In our context, even the error is verifiable. The interpretation is that an expert collects information to formulate their report. While this information may be inaccurate, the expert must disclose it truthfully – i.e. without altering it. An uninformed expert must admit ignorance, while an informed expert i must either report the signal s_i he observes or pretend to be uninformed. The decision-maker takes the action that maximizes her expected payoff, given the posterior belief about the state based on the reports by the experts.

In this framework our question is as follows: If the decision-maker could choose the experts (at the beginning of the game) based on their agendas, when is a diverse panel – composed of experts with agendas 0 and 1 – better than a homogeneous panel – composed of experts with identical agendas?

As is common in the literature, we focus on the *Perfect Bayesian Equilibrium* (PBE) of the persuasion game. However, there are a number of issues here. Unlike Bhattacharya and Mukherjee (2013), it is no longer immediate that when an expert reveals his signal s_i, the judge takes the action that matches s_i. Indeed, we show that when the error probability ϵ is small ($\epsilon < \frac{1}{2}$), the decision-maker's best response is to match the signal s_i for a homogeneous panel (most of the time). For a diverse panel, the issues are more subtle. We are able to show that the best response for the decision-maker is to either match the signal s_i or if both experts report different signals, choose the average of the two signals when ϵ is small. Given these results, it turns out that the equilibrium of this game is characterized by the decision-maker's default action y^*, taken when all experts fail to report a signal. Consequently, an expert reports a signal s_i if doing so is more favorable to him than y^*. As in Bhattacharya and Mukherjee (2013), an informed (possibly with error) expert's disclosure strategy is given by a revelation set – i.e. the set of states where he would report his signal truthfully to the decision-maker. As such, for an expert with agenda 0, his revelation set is $\Theta_0^* = [0, y^* + \delta]$. For an expert with agenda 1, it is $\Theta_1^* = [y^* - \delta, 1]$. Thus, each expert's revelation set is a set of favorable states close to his ideal action, and the judge's default action y^* is a best response to such disclosure strategies. Therefore, the joint revelation sets in a homogeneous panel covers the sets of states smaller (larger) that y^* if $x_i = 0$ ($x_i = 1$) for all i, while they cover the entire state space in a diverse panel. This is similar to the characterization in Bhattacharya and Mukherjee (2013) and Bhattacharya et al. (2018). The only difference is that the revelation sets are adjusted by the error level δ. This characterization leads to a tradeoff between diverse and homogeneous panels. Indeed, compared to a diverse panel, with a homogeneous panel the decision-maker learns the state with a higher probability if the state lies in a larger subset of the state space, while higher states – that is states outside the subset $[0, y^* + \delta]$ – are never reported. In a diverse panel, each state is reported by exactly one of the two experts (most of the time).

In our main result, we show that for given values of ϵ and δ, there exists a threshold level of the expert quality α^* such that for $\alpha > \alpha^*$, the diverse panel becomes better for the decision-maker. Likewise, for a given value of experts' quality α, and a given level of error δ, there exists a threshold level of error probability $\epsilon^* < \frac{1}{2}$ such that, for $\epsilon \in (\epsilon^*, \frac{1}{2})$, the diverse panel is optimal for the decision-maker. The intuition behind the first result is the following: as the quality of the experts increases, in the presence of a small amount of agenda-specific error (in both level and probability of error), the benefits from a homogeneous panel are offset by the benefits from the entire state space being covered by the diverse panel. Given the error level and the error probability, the value of having both experts reporting over the same set remains unchanged, but the value of reporting over different sets of states goes up. A similar intuition works for the second result. Our findings can be seen as an attempt to verify the robustness of the results in Bhattacharya and Mukherjee (2013). In addition, a significant body of the literature has compared the efficacy of a diverse panel with a setting where only one expert is responsible for gathering and revealing all information. For example, Shin (1998) compares a panel of two experts with opposing interests to one unbiased expert and demonstrates that a diverse panel reveals more information. Dewatripont and Tirole (1999) make a similar point. On the other hand, Bhattacharya and Mukherjee (2013), and Bhattacharya et al. (2018) show the optimality of the extreme homogeneous panel for a general class of problems. Our result can also be seen as an attempt to bridge the gap between these two strands of the literature.

2.1.1 Related Literature

As mentioned previously, our work is an extension of Bhattacharya and Mukherjee (2013). Indeed, they consider a model with a judge and a panel of two experts. Each of these experts has an ideal action that he would like the judge to take. The judge's goal is to pick an action as close as possible to the state. With no bias, the experts will reveal the state of the world to the judge if they are informed and if doing so improves their payoff upon the default action – i.e. if the state lies within their revelation sets. Bhattacharya and Mukherjee (2013) shows that a panel of extreme experts – homogeneous or diverse – is always optimal. Moreover, if the judge has a quadratic loss function, then the homogeneous panel of extreme experts is preferred by the judge, regardless of the probability of the experts being informed. Bhattacharya et al. (2018) refers to an expert being informed or not as his type. As opposed to Bhattacharya and Mukherjee (2013) and our work, they allow for correlation between types – i.e. for a panel of two experts, if one expert is informed there is a higher probability that the other one is as well. Informed experts observe the state, but are biased in their reporting depending on their ideal action. They show that with high correlation between types, the diverse panel is a better choice for the judge than the homogeneous panel. Moreover, if types are independent, then the homogeneous panel is generally the optimal choice.

Also related to this chapter, Milgrom (1981), Grossman (1981) and Milgrom and Roberts (1986) show that a skeptical decision-maker can elicit all information from a diverse panel of experts, assuming these latter are always informed – i.e. every expert observes the state with certainty –, an assumption that differs from our work[1]. As noted in, for example, Dye (1985), Jung and Kwon (1988), Okuno-Fujiwara et al. (1990), Shavell (1994) and Shin (1994), if there is some positive probability that an expert does not observe the state, then an informed expert may select a non-disclosure strategy when revealing his information makes him worse off. Shin (1998) shows that a panel of two experts with opposing interests reveals more information than a panel of one unbiased expert. Similarly, Dewatripont and Tirole (1999) show the benefits of advocates (opposing experts) in policy-making, arguing that in many situations it leads to higher levels of information collection by the decision-maker over a single unbiased expert. Kartik et al. (2017) demonstrate that when information acquisition is costly and endogenous for the experts – we however assume it to be costless and exogenous –, then the experts' effort level decisions are strategic substitute and increasing the number of experts may decrease these efforts. The literature on disclosure games is also relevant to our work, and one can look at Wolinsky (2002) for such example. Next, Gentzkow and Kamenica (2017) show that competition weakly increases information revelation if the information environment is *Blackwell-connected*, and that it otherwise might not be the case. Finally,

---

[1] See section 2.5 for further details on how this relates to our results.

Li and Norman (2021) argue that when having experts disclose their respective signal sequentially, adding an expert moving first weakly increases information transmission in equilibrium, and that having experts speak only once except for the first one does not change the equilibrium outcome. Moreover, having experts move sequentially makes for a weakly less informative equilibrium as opposed to the simultaneous counterpart.

The rest of the chapter is organized as follows: in Section 2.2 we introduce the model; in Section 2.3 we characterize the equilibrium generally and for the different panel compositions; we present in Section 2.4 our main results on the optimal panel composition, along with some comparative statics. In Section 2.5, we discuss the implications of decision-maker having a quadratic loss function. In addition, we briefly discuss the implications of uniform error, as against the error being bias specific. Section 2.6 concludes. The proofs of the main propositions are relegated to the appendix.

2.2 Model

We consider a model of a judge selecting a panel of experts as in Bhattacharya and Mukherjee (2013). There are two experts $i \in \{1, 2\}$ who may each observe a signal $s_i \in \Theta$ about the state of the world $\theta \in \Theta = [0, 1]$, and one judge who has no information about θ beside the common prior. The judge must choose an action $y \in \Theta$ based on the reports from the two experts $m_i(s_i) \in \Theta \cup \{\varnothing\}$. There are two qualities of expert: either an expert is *informed* ($t_i = 1$ with probability α_i), in which case he observes a signal $s_i \in \Theta$, or *uninformed* ($t_i = 0$ with probability $1-\alpha_i$), in which case he does not observe such signal ($s_i = \varnothing$). For now, we assume $\alpha_0 = \alpha_1 = \alpha$. Conditional on being informed, each expert decides if he reports the signal that he has observed ($m_i(s_i) = s_i$), or not ($m_i(s_i) = \varnothing$). However, lying is never an option; an informed expert can only report either s_i or $\varnothing$, and an uninformed expert must always report $\varnothing$. In addition, the signal observed by an expert may be *biased*. A signal is said to be *unbiased* if $s_i = \theta$, and *biased* if $s \neq \theta$. The probability that a signal is biased is ϵ. The judge wants her action y to be as close to θ

as possible. As such, the utility function of the judge is represented as

$$u^J(y, \theta) = -|y - \theta|$$

Each expert i has an ideal action $x_i \in \Theta$, and as such his utility function is defined as

$$u_i(y, x_i) = -|y - x_i|$$

We focus on *extreme* experts who have their ideal action as $x_i \in \{0, 1\}$. As mentioned, conditional on being informed an expert i can observe a biased or unbiased signal. That is, conditional on expert i being informed and $x_i = 0$:

$$s_i(\theta) = \begin{cases} \theta \text{ with probability } 1 - \epsilon \\ \max(0, \theta - \delta) \text{ with probability } \epsilon \end{cases}$$

Conditional on expert i being informed and $x_i = 1$:

$$s_i(\theta) = \begin{cases} \theta \text{ with probability } 1 - \epsilon \\ \min(\theta + \delta, 1) \text{ with probability } \epsilon \end{cases}$$

In the above signal representations, δ refers to the level of the bias. Upon observing his information about the state, an expert's pure strategy is $m_i(s_i) = \{s_i, \varnothing\}$ if informed, or $m_i = \varnothing$ if uninformed. We denote by $m = (m_1, m_2)$ the profile of reports submitted to the judge.

The judge's strategy $y(m)$ can be formulated as two distinct parts: (i) the best response $y(s_1, s_2)$ upon receiving m and $m_i \neq \varnothing$ for at least one expert; and (ii) a *default action* y^* if $m_1 = m_2 = \varnothing$. We assume that the state θ follows a uniform distribution on the interval $[0, 1]$, and that this distribution (denoted F) is common knowledge. We also assume that the judge knows about the bias of an expert given his ideal action (± δ), but cannot tell if a report $m_i(s_i) \neq \varnothing$

represents the true state of the world.

We use the Perfect Bayesian Equilibrium (PBE) solution concept. We denote $\mu(\theta|m)$ the posterior belief of the judge about the state upon observing the report profile m from the experts. A strategy profile $\langle m^*, y(m^*)\rangle$ and belief μ^* is a PBE[2] of the considered game if:

(i) Experts send their reports simultaneously. For all $i \in I$ and $\theta \in \Theta$, if an expert i is informed, then he reports his signal $m_i(s_i) = s_i$ if and only if:

$$\mathbb{E}[u_i(y(s_i, m_{-i}(s_{-i})), x_i)] \geq \mathbb{E}[u_i(y(\varnothing, m_{-i}(s_{-i})), x_i)]$$

(ii) For all report profiles m, the judge's action maximizes her expected utility:

$$y(m) = \arg\max_{y\in\Theta} \mathbb{E}_{\mu^*}[u^J(y, \theta)]$$

(iii) $\mu^*(\theta|m)$ is obtained using Bayes' rule from the strategy of the experts m^* and the common priors. Moreover, any action off-the-equilibrium path that reveals θ must lead to degenerate beliefs on θ for the judge.

In addition, in our model, due to the possibility of error on the part of the experts, it may be the case that an off-the-equilibrium path action profile reveals a signal s that differs from θ by an amount δ. We need to careful specify the off-the-equilibrium path beliefs in such cases:

- Suppose that the off-the-equilibrium path action that reveals s is taken by the expert with agenda 0. Then the judge's belief is the following: $\mu(\theta = s + \delta|s) = \epsilon$ and $\mu(\theta = s|s) = 1 - \epsilon$.
- On the other hand, if the off-the-equilibrium path action that reveals s is taken by the expert with agenda 1, then the judge's belief is the following: $\mu(\theta = s - \delta|s) = \epsilon$ and

---

[2] See Bhattacharya and Mukherjee (2013) for a discussion on the adaptation of the PBE formulation to the current problem.

$$\mu(\theta = s|s) = 1 - \epsilon.$$

2.3 Equilibrium Characterization

In this section, we present the equilibrium characterization first generally, then for the different panel compositions. These include the panels of one expert, of two identical experts, and of two opposed experts. For each we present the best-response function for the judge and the optimal default action.

2.3.1 General Characterization

As opposed to the previous models presented in Bhattacharya and Mukherjee (2013) and in Bhattacharya et al. (2018), we must also establish what is the judge's best course of action upon receiving $m \neq (\varnothing, \varnothing) \equiv \widetilde{m}$. Then we must derive the optimal default action y^* when $m = (\varnothing, \varnothing)$.

Definition 1 *An equilibrium for this game is characterized by the judge's best-response function* $y(m)$ *defined as*

$$y(m) = \begin{cases} y(\widetilde{m}) \textit{ if } m = \widetilde{m} \\ \\ y^* \textit{ if } m = (\varnothing, \varnothing) \end{cases}$$

where

$$y(\widetilde{m}) = \arg\max_{y \in \Theta} \mathbb{E}_{\mu|\widetilde{m}}[u^J(y, \theta)]$$

is the best response of the judge given the new information provided by at least one of the experts and the posterior belief $\mu(\theta|\widetilde{m}) \equiv \mu|\widetilde{m}$*, and where*

$$y^* = \arg\max_{y \in \Theta} \mathbb{E}_{\mu_0}[u^J(y, \theta)]$$

is the optimal default action from observing $m = (\varnothing, \varnothing)$ *inducing the posterior beliefs* μ_0.

It might not be in the judge's best interest to blindly choose $y(\widetilde{m}) = s' \in \widetilde{m}$ if she gets an informative report. As shown in the specific cases below, depending on the report and the bias of the expert presenting the information, the judge may be able to deduce an expert's type given his report. Moreover, in the case of the diverse panel, it is possible that the two experts submit different information as shown in the relevant section below. If, however, the judge does not get any informative signal, then depending on the panel composition it is either the case that all experts are uninformed and/or prefer the null report.

Following the above characterization, y^* serves in the definition of the upper (lower) bound of the revelation set for the expert with ideal action at 0 (1). First consider an unbiased expert at 0 and a fixed y^* (the case of the expert at 1 is analogous). If $s = \theta \leq y^*$ then the expert prefers to reveal this information to the judge than concealing it. The reverse holds if $\theta > y^*$. As such, following the notation in Bhattacharya and Mukherjee (2013), the revelation set of this expert is $\Theta_0^* = [0, y^*]$. Now if this expert were to instead be biased, then $s = \max(0, \theta - \delta)$. The expert does not know that his signal is biased. As such, $\widetilde{\Theta}_0^* = [0, y^* + \delta]$. We next proceed with characterizing the equilibrium for specific panel compositions of pertinence for this chapter.

2.3.2 Panel of One Expert

We first look at a panel consisting of only one expert at $x = 0$. The case of the expert at $x = 1$ is analogous and as such omitted. We consider the judge's best response upon observing a signal $s \neq \{\varnothing\}$. With probability ϵ, the expert is bias, and these priors are common. The judge has an absolute-value loss function – i.e. $u^J(y - \theta) = -|y - \theta|$. The expert in this case, if informed, will observe $s = \theta$ if he is unbiased, or $s = \max(0, \theta - \delta)$ if he is biased.

Proposition 1 *Consider a panel consisting of one expert with ideal action $x = 0$. The judge's absolute payoff function is $u^J(y, \theta) = -|y - \theta|$. Then the judge's best-response function upon* *observing a signal s from the expert is given as*

$$y^*(s) = \begin{cases} \frac{\delta}{2} \text{ if } s = 0 \\ s \text{ if } s \in (0, 1-\delta] \text{ and } \epsilon < \frac{1}{2} \\ s + \delta \text{ if } s \in (0, 1-\delta] \text{ and } \epsilon > \frac{1}{2} \\ s \text{ if } s \in (1-\delta, 1] \end{cases} \tag{1}$$

All proofs are relegated to the appendix. Essentially, Proposition 1 says that if the judge observes a signal $s = 0$, then she knows almost certainly that the expert is biased. Indeed, either s originates from an unbiased expert – the probability of the single event $\theta = s = 0$ over a continuous (uniform) distribution is zero – or a biased expert – where $\theta \in [0, \delta]$ leads to $s = 0$. For the latter, the true state of the world θ can be any value in $[0, \delta]$, and $y^*(s = 0) = \frac{\delta}{2}$ maximizes the judges expected payoff in that case. If the judge instead gets $s \in (1-\delta, 1]$, she knows for sure that the expert is unbiased as only an unbiased expert can give such report (the maximum signal from the biased expert being $1-\delta$ if $\theta = 1$). As such, $y^*(s) = s$ in this case. Finally, if $s \in (0, 1-\delta]$, then the judge cannot deduce if the expert is biased or not. Indeed, upon observing such signal s, the true state can either be $\theta = s$ with probability $1-\epsilon$ or $\theta = s + \delta$ with probability ϵ. Having observed s, the judge's interim payoff is maximized by choosing $y^*(s) = s$ if $\epsilon < \frac{1}{2}$ since it is more likely that the expert is unbiased. Otherwise $y^*(s) = s + \delta$ if $\epsilon > \frac{1}{2}$. For the purpose of this chapter, we look at experts that are more likely to be accurate. As such, from now on we will assume that $\epsilon < \frac{1}{2}$. One could interpret this assumption as imposing a (weak) minimum on the reliability of the experts. Finally, if the judge does not get any report from the expert ($s = \varnothing$), then she chooses the default action y^* defined below.

Given the best response function $y^*(s)$ and $\epsilon < \frac{1}{2}$, the judge must decide on an optimal default action y^*. This plays a crucial role in determining the revelation set of the expert – i.e. the set for which it is more beneficial for the expert to reveal the signal s instead of concealing it.

Proposition 2 *If $\epsilon < \frac{1}{2}$ and the panel consists of one expert at $x' = 0$, then the judge's optimal* *default action solves the maximization problem*

$$\max_{y^*} EU^J(y^*)$$

and is given by

$$y^* = \frac{1 - \alpha\epsilon\delta}{2 - \alpha} \tag{2}$$

This solution is unique. Moreover, y^* is decreasing in the bias of the expert (ϵ and δ) and increasing in his informativeness (α).

The proof of Proposition 2 is relegated to the Appendix. Observe that y^* is lower than if there were no bias introduced in this model. Indeed, in such case this corresponds to $y^* = \frac{1}{2-\alpha}$[3]. As such, the upper bound of the revelation set for this expert decreases in ϵ and δ. In other words, the judge will rely less on the report of a more biased extreme expert at zero by reducing her default action y^*.

2.3.3 Homogeneous Panel

Next we derive the judge's best response y^* when the panel is composed of two identical experts whose ideal action is $x_i = 0$ for all $i \in \{1, 2\}$. The case of both of them at 1 is analogous. Here, the two experts are essentially identical. As such, the judge's best response upon observing $\widetilde{m}$ is the same as in section 2.3.2. Keeping the assumption that $\epsilon < 1/2$, this leads to the following proposition.

Proposition 3 *Consider a panel consisting of two experts with the same ideal action $x_1 = x_2 = x'$. The judge's absolute payoff function is $u^J(y, \theta) = -|y - \theta|$. Then the judge's best-response function upon observing a report profile $\widetilde{m} \neq (\varnothing, \varnothing)$ is given as*

---

[3] See Bhattacharya and Mukherjee (2013).

(i) *for* $x' = 0$*:*

$$y^*(\widetilde{m}) = \begin{cases} \frac{\delta}{2} \text{ if } s = 0 \\ \\ \overline{s} \text{ if } \overline{s} \in (0,1] \end{cases} \tag{3}$$

where

$$\overline{s} = \begin{cases} \max\{s_1, s_2\} \text{ if } (s_1, s_2) \neq (\varnothing, \varnothing) \\ \\ s_k \text{ if } s_k \neq \varnothing \text{ and } s_{-k} = \varnothing, k \in \{1,2\} \end{cases}$$

(ii) *for* $x' = 1$*:*

$$y^*(\widetilde{m}) = \begin{cases} s \text{ if } s \in [0,1) \\ \\ 1 - \frac{\delta}{2} \text{ if } s = 1 \end{cases} \tag{4}$$

where

$$\overline{s} = \begin{cases} \min\{s_1, s_2\} \text{ if } (s_1, s_2) \neq (\varnothing, \varnothing) \\ \\ s_k \text{ if } s_k \neq \varnothing \text{ and } s_{-k} = \varnothing, k \in \{1,2\} \end{cases}$$

In the first case of Proposition 3, we state that if the signal profile consists of two informative reports, then only the highest of the two is relevant. Indeed, because the judge is aware of the possible bias the experts may have towards their ideal action, upon receiving two reports she knows that the highest one corresponds to the true state – and is submitted by an unbiased expert. A similar assessment can be made for the other case. Consequently, the proof is similar to Proposition 1 and is therefore omitted.

Using this best response function, we now derive the ex ante expected utility for the judge EU^J and use it to find the optimal default action y^*. This, along with the aforementioned $y(\widetilde{m})$, completes the best-response function of the judge. We only present the case of the homogeneous panel at 0 as the case of $x_1 = x_2 = 1$ is analogous.

Proposition 4 *If* $\epsilon < \frac{1}{2}$ *and the panel is homogeneous at* $x' = 0$*, then the judge's optimal default* *action solves the maximization problem*

$$\max_{y^*} EU^J(y^*)$$

and is given by

$$y^* = \frac{1 - \delta((\alpha\epsilon)^2 + 2\alpha\epsilon(1-\alpha))}{1 + (1-\alpha)^2} \tag{5}$$

This solution is unique. Moreover, y^* is decreasing in the bias of the expert (ϵ and δ) and increasing in his informativeness (α).

Proposition 4 shows that higher bias lowers the default action. One can see that the terms within parenthesis on the numerator is strictly positive and multiplied by $-\delta$ which is strictly negative by definition. It presents a similar intuition to Proposition 2 regarding the reliance of the judge on the experts depending on their bias. The proof is relegated to the Appendix.

2.3.4 Diverse Panel

We now consider a panel of two different extreme experts: one whose ideal action is $x_1 = 0$ and the other $x_2 = 1$. We first compute the judge's best response when she observes at least one signal. Recall that for the expert at 0 his biased signal is $s_1 = \max\{0, \theta - \delta\}$, and for the expert at 1 it is $s_2 = \min\{\theta + \delta, 1\}$. Before proceeding with the characterization of the judge's best-response function, observe that for the diverse panel it will never be the case that the experts with submit identical reports (excluding when $\theta = y^*$ and both experts are informed and unbiased or when both are uninformed). Indeed, for this panel design, either (i) $\Theta_0^* \cap \Theta_1^* = \{y^*\}$, (ii) $\Theta_{0b}^* \cap \Theta_1^* = [y^*, y^* + \delta]$, (iii) $\Theta_0^* \cap \Theta_{1b}^* = [y^* - \delta, y^*]$ or (iv) $\Theta_{0b}^* \cap \Theta_{1b}^* = [y^* - \delta, y^* + \delta]$, where we index by b the revelation set of a biased expert. Because of the truthful report requirement, if both experts are informed and at least one expert is biased then $s_1 \neq s_2$. If both are informed and unbiased, then only if $\theta = y^*$ can both submit $s_1 = s_2 = \theta$. If only one expert is informed, then the signals obviously differ. If both are uninformed, then $s_1 = s_2 = \varnothing$. Proposition 5 details the judge's best-response function.

Proposition 5 *Consider a panel consisting of two experts with opposing ideal actions* $x_1 = 0$ *and* $x_2 = 1$*. The judge's absolute payoff function is* $u^J(y, \theta) = -|y - \theta|$*. Then the judge's best-response function upon observing a report profile* $\widetilde{m} \neq (\varnothing, \varnothing)$ *is given as*

$$y(s_1, s_2) = \begin{cases} \frac{\delta}{2} \textit{ if } s_1 = 0 \textit{ and } s_2 = \varnothing \\ s_1 \textit{ if } s_1 \in (0, 1] \textit{ and } s_2 = \varnothing \\ s_2 \textit{ if } s_2 \in [0, 1) \textit{ and } s_1 = \varnothing \\ 1 - \frac{\delta}{2} \textit{ if } s_2 = 1 \textit{ and } s_1 = \varnothing \\ \frac{s_1 + s_2}{2} \textit{ if } s_2 - s_1 = 2\delta \textit{ or } s_2 - s_1 = \delta \end{cases} \tag{6}$$

If the judge receives only one informative report, then she acts as if the panel only consisted of the expert who had submitted it. The main novel part of Proposition 5 comes from the last case of $y(s_1, s_2)$. When the two experts submit different reports, the judge can either deduce the true state if $s_2 - s_1 = 2\delta$ by choosing $y(s_1, s_2) = \frac{s_1 + s_2}{2} = \theta$, or decide on a middle action if $s_2 - s_1 = \delta$. Note that for this latter case, due to the symmetric nature of our parametrization, the judge can in fact decide on any $y(s_1, s_2) \in [s_1, s_2]$. For simplification and without loss of generality, we decided on the midpoint strategy. This option is also intuitively more meaningful, stating that whenever the judge gets two different reports she chooses the compromise between the two. We next compute the optimal default action y^* for the judge when the panel is diverse.

Proposition 6 *If* $\epsilon < \frac{1}{2}$ *and the panel is diverse with* $x_1 = 0$ *and* $x_2 = 1$*, then the judge's optimal default action solves the maximization problem*

$$\max_{y^*} EU^J(y^*)$$

and is given by

$$y^* = \frac{1}{2} \tag{7}$$

With $\alpha_1 = \alpha_2 = \alpha$*,* $\delta_1 = \delta_2 = \delta$*, and* $\epsilon_1 = \epsilon_2 = \epsilon$*, this solution is invariant to the parameters.*

For the purpose of this chapter, it is most useful to focus on the symmetric case and is without loss of generality. Indeed, our objective is to evaluate the robustness of the optimality of the homogeneous panel composition. If we were to consider any combination of $\delta_1 > \delta_2$ and $\epsilon_1 > \epsilon_2$, then one expert would be more biased than the other[4]. As such, a homogeneous panel experts with the lowest bias would be preferred by the judge[5]. Now due to this symmetric nature of the problem, this is the same result as if there were no bias. These assumptions are without loss of generality and only serves to simplify the calculations. In the following section, we establish the conditions for the diverse panel to be chosen by the judge over the homogeneous one.

2.4 Optimal Panel Composition

In this section, we explore the different panel compositions – diverse and homogeneous – and the parametric conditions that makes one the better choice from the judge's perspective. In particular, we show the requirements on α, δ, and ϵ for the diverse panel to be optimal.

Let y^D and y^H be the optimal default actions for the diverse and homogeneous panels respectively. Similarly, let $EU^D(y^D)$ and $EU^H(y^H)$ be the corresponding expected utility functions for the judge given the default action y^k, $k \in \{D, H\}$.

We define the function

$$g(\epsilon; \alpha, \delta) = EU^D(y^D) - EU^H(y^H) \tag{8}$$

The diverse panel is preferred if $g(\epsilon; \alpha, \delta) > 0$. By definition, the function g is continuous and differentiable everywhere on the interval $\epsilon \in [0, 1]$.

Theorem 1 *We claim the following:*

(i) $g(\epsilon; \alpha, \delta)$ *is strictly convex in* ϵ *for any* $\alpha, \delta \in (0, 1)$*;*

---

[4] See Bhattacharya and Mukherjee (2013) for an explanation on why assuming $\alpha_1 = \alpha_2 = \alpha$ is without loss of generality.

[5] The asymmetrical case only increases the revelation set of the expert who has the highest likelihood of revealing the true state. This effectively works in favour of the homogeneous panel which is not helpful for the purpose of this essay.

(ii) $g(0; \alpha, \delta) < 0$ *for any* $\alpha, \delta \in (0, 1)$*;*

(iii) $g\left(\frac{1}{2}; \alpha, \delta\right) > 0$ *for some* $\alpha, \delta \in (0, 1)$*;*

(iv) *there exists values of* α *and* δ *such that* (iii) *holds, and given* (i) *and* (ii)*, there exists* $\epsilon^* < \frac{1}{2}$ *such that* $g(\epsilon^*; \alpha, \delta) = 0$*. Moreover,* $g(\cdot)$ *is increasing around* ϵ^**. As such, for* $\bar{\epsilon} \in (\epsilon^*, \frac{1}{2})$ *it must be that* $g(\bar{\epsilon}) > 0$*.*

Theorem 1 establishes that if conditions (i), (ii), and (iii) are satisfied for some $\alpha, \delta \in (0, 1)$, then if ϵ is above the threshold ϵ^* the diverse panel yields a higher ex ante expected utility than the homogeneous panel. However, as shown in Figure 2.1 below, this requires a relatively high level of bias (δ).

2.4.1 Comparative Statics

Fix the level of informativeness (e.g. $\alpha = \frac{1}{2}$). First, observe in Figure 2.1 that for low levels of bias δ the homogeneous panel is always preferred. This is shown by the red curve. Second, the blue curve illustrates that slightly increasing δ does not make the diverse panel optimal. Indeed, the function $g(\cdot)$ crosses the horizontal axis at a point $\epsilon^* \approx 0.7599 > \frac{1}{2}$ which contradicts the assumption that $\epsilon < \frac{1}{2}$. Third, we can find values of α and δ that does allow for a feasible solution in our context. Indeed, the green curve in Figure 2.1 crosses the horizontal axis at $\epsilon^* \approx 0.3713$ and, as such, for any $\epsilon > \epsilon^*$ the diverse panel is preferred.

It is also worth pointing out that all three lines share the same origin ($g(0; \alpha, \delta) = -0.025$). One can see from the definition of g that δ that if $\epsilon = 0$, then the level of bias is immaterial to this problem. Also, setting $\epsilon = 0$ makes our model identical to Bhattacharya and Mukherjee (2013). Regarding the parameter α, it does play a role in determining the curvature of g as depicted in Figure 2.2. While both the purple and the orange lines are convex, they do not share the same origin anymore. Moreover, the latter is decreasing for small values of ϵ. Notice that this also illustrates the role of informativeness in determining the panel composition. For a fixed level of bias (e.g. $\delta = 0.3$), if $\alpha = 0.5$ the diverse panel is never optimal given our restriction on admissible

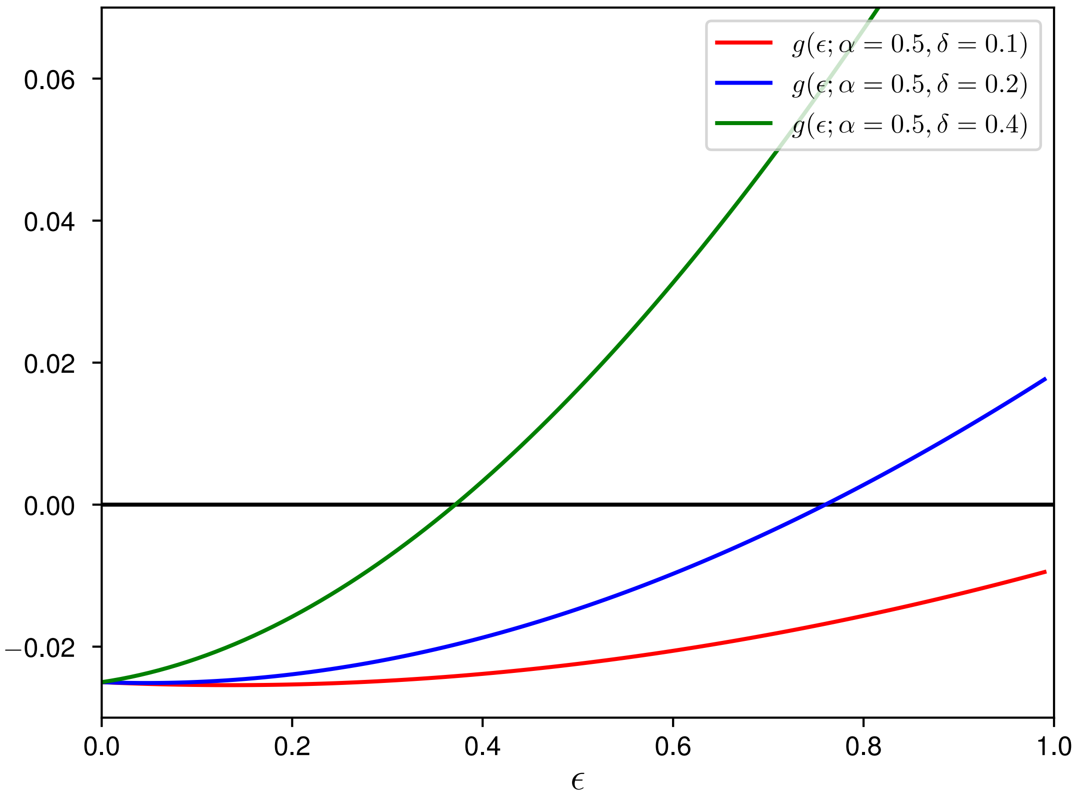

Figure 2.1: Function $g(\epsilon)$ given $\alpha = 0.5$ and $\delta \in \{0.1, 0.2, 0.4\}$

values of ϵ ($\epsilon^* \approx 0.5205$). However, increasing it to $\alpha = 0.9$ allows for a panel of experts with opposing agendas to be the better choice if $\epsilon > \epsilon^* \approx 0.4319$.

Diving deeper into the role of α, we rewrite g as a function of α and fix the values for ϵ and δ. Figure 2.3 fixes δ and consider different values for ϵ, while 2.4 does the opposite using the same values. Both of them depict a similar story: there is a non-monotonic relationship between g and α. Indeed, lower levels of α work in favor of the homogeneous panel. However, depending on the bias (in terms of level δ or probability ϵ), there is a point α^* above which larger values of α makes the diverse panel relatively better. For higher bias, $\alpha^* \leq 0$ and as such the diverse panel is always optimal. At the opposite, for lower bias this value is such that the homogeneous panel is the better option for any $\alpha \in [0, 1]$. At the middle ground of these two scenarios, there is a

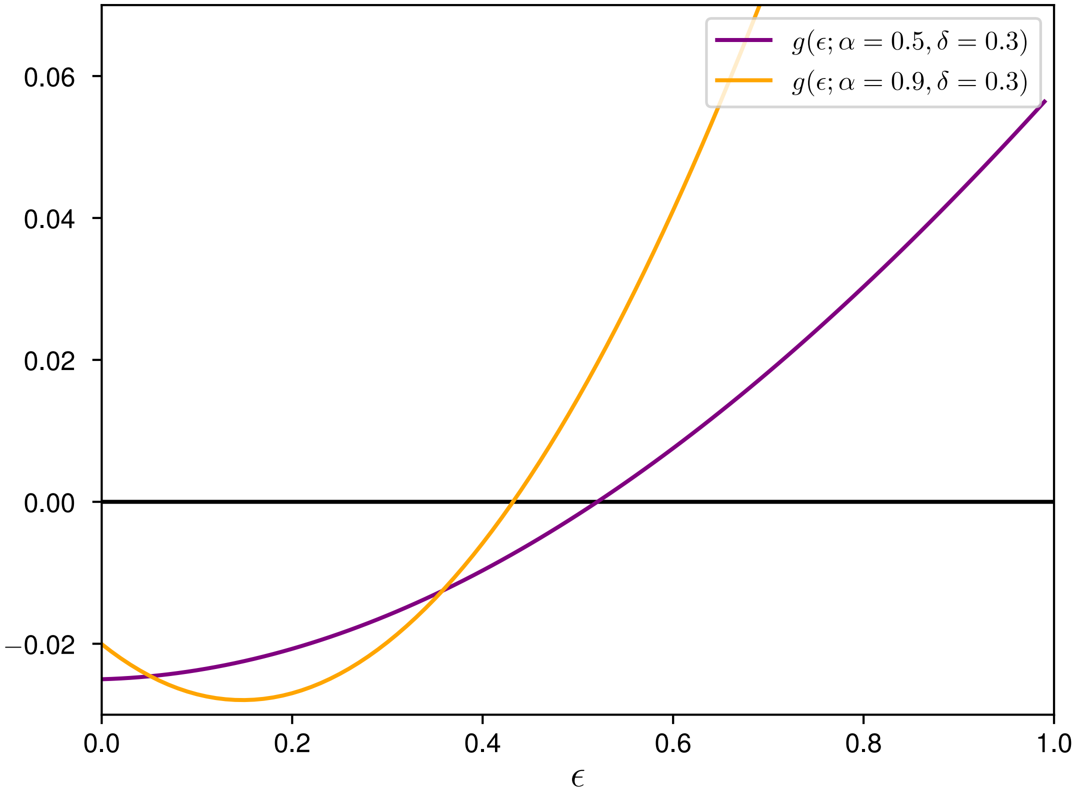

Figure 2.2: Function $g(\epsilon)$ given $\alpha \in \{0.5, 0.9\}$ and $\delta = 0.3$

$\widetilde{\alpha} > \alpha^*$ that marks the cutoff at $g(\alpha|\epsilon', \delta') = 0$. Therefore, for $\alpha > \widetilde{\alpha}$, the diverse panel is the preferred panel by the decision-maker. This shows the non-monotonicity of the returns from the informativeness α of the experts, which is not specific to our model. Indeed, Figure 2.5 shows that in the absence of bias – returning to the model presented in Bhattacharya and Mukherjee (2013) – we still observe this relationship. This also shows the indifference points between the two panel compositions at $\alpha \in \{0, 1\}$. Adding bias essentially moves one of the extreme points towards the center, depending on the agenda profile of the homogeneous panel. In the current case, we have focused on $x = (0, 0)$ which changes the position of the right point. By doing so, we allow for values of $\alpha < 1$ that makes the diverse panel optimal. Intuitively, very informed experts makes the homogeneous panel less profitable for the judge. Recall from Proposition 4 that for any positive bias, the default action for the homogeneous panel is strictly less than 1 even if $\alpha = 1$. Therefore, it is always the case that $\Theta^*_{00} \subset \Theta = [0, 1]$. This is not the case for the diverse panel, for which at $\alpha = 1$ the revelation set covers the entire set of states Θ. This illustrates the arbitrage between a lower occurrence of more reliable information – from the homogeneous panel, at $\alpha = 1$ the probability of receiving a biased report is ϵ^2 over the revelation set $\Theta^*_{00} = [0, y^*_h]$, with $y^*_h < 1$ – and a higher occurrence of poorer information – from the diverse panel, at $\alpha = 1$ the probability of receiving a biased report is ϵ over the revelation set $\Theta^*_{00} = [0, 1] = \Theta$. Next we summarize our main results with the help of an example. We show how the inclusion of bias affects the choice of experts in the composition of a panel.

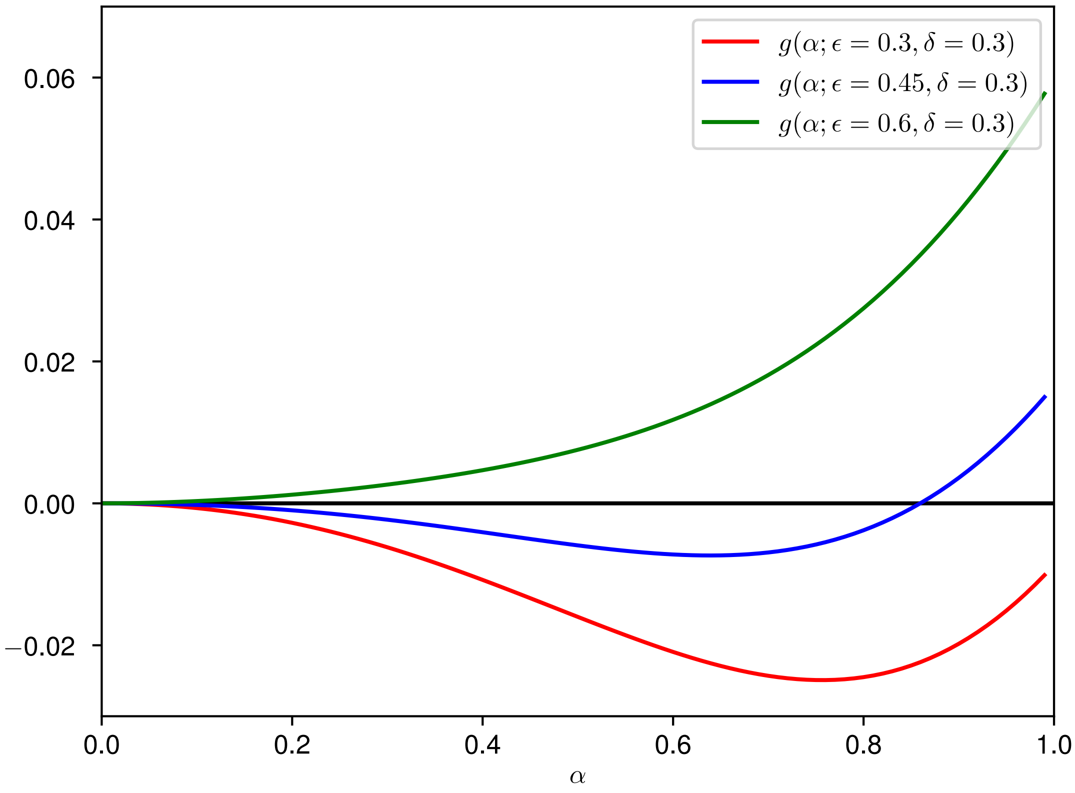

Figure 2.3: Function $g(\alpha)$ given $\epsilon \in \{0.3, 0.45, 0.6\}$ and $\delta = 0.3$

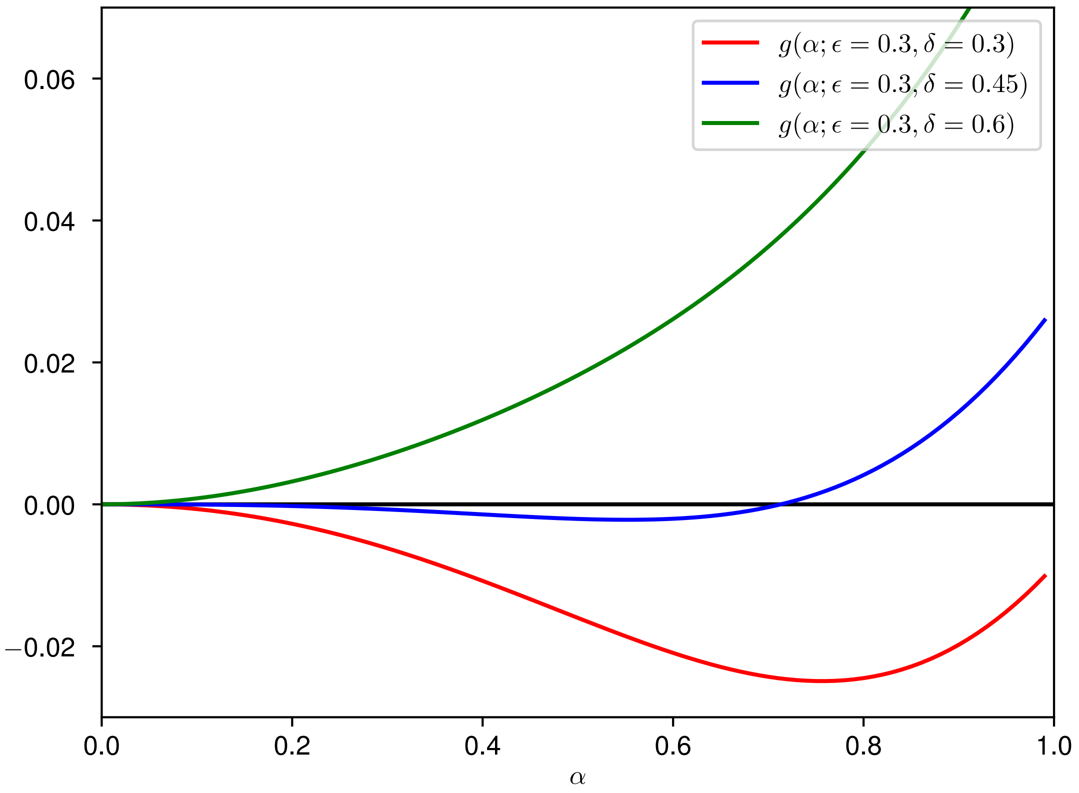

Figure 2.4: Function $g(\alpha)$ given $\epsilon = 0.3$ and $\delta \in \{0.3, 0.45, 0.6\}$

2.4.2 Example: Public Policy

In this section, we illustrate our results through a simple example. We depict the situation in which a government decides on the composition of a panel of experts to provide information. The inputs from these experts are then used by the government to decide on an appropriate public policy.

Concretely, a government wants to establish the share of the firms in the country's total pollution. By doing so, the government is looking for an appropriate level of taxation to finance environmental countermeasures. As such, it asks for available experts from both the industry – to represent the firms – and from environmental groups. From this pool of experts, the government

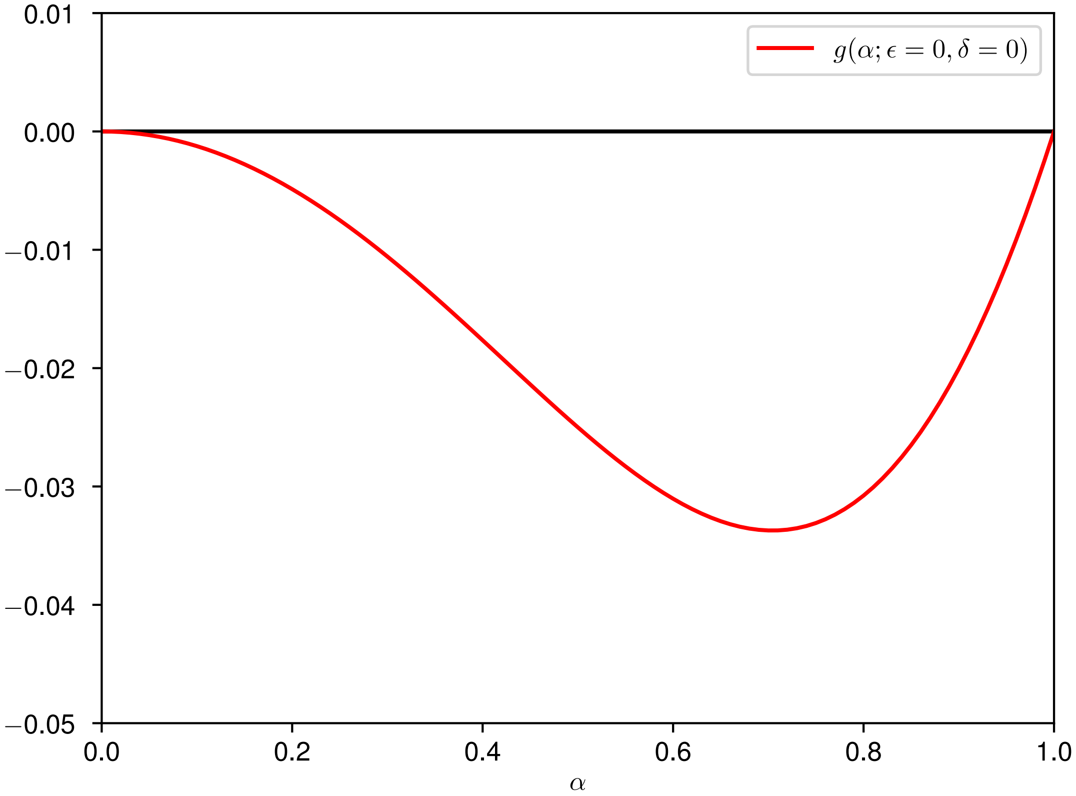

Figure 2.5: Function $g(\alpha)$ given $\epsilon = \delta = 0$

then forms a panel to provide the most information about the actual share of responsibility of the firms. Naturally, because a higher share will result in higher penalties, the firms would like the government to agree on a null level of responsibility. At the opposite, environmental groups would rather have the authority decides on a share of 100%.

Using the model presented previously, the government acts as the judge while the firms and environmental groups as the experts. We also proceed by referring to the latter as singular experts from each position. One can think of this as having a unique representative expert for each opinion who aggregates all the recommendations formulated by its members into a single (average) report. Each side can contain misinformed members that tends to submit reports closer to their preferences. This is described as their agenda-specific bias.

Formally, let the environmental groups (E) and the firms (F) have their ideal actions defined as $x_E = 1$ and $x_F = 0$. For $i \in E, F$, the payoff functions of the experts are identical:

$$u_i(x_i, y) = -|x - y|$$

Similarly, the government's (G) payoff function is

$$u_G(y, \theta) = -|y - \theta|$$

Let us first consider the case where there is no agenda-specific bias. That is $\epsilon = 0$ or $\delta = 0$ (or both). We also assume that $\alpha_E = \alpha_F = \alpha = 0.8$. The default action for the homogeneous panel composition is given as

$$y_h^* = 1 - \frac{1}{1 + (1 - \alpha)^2} = 1 - \frac{25}{26} \approx 0.038$$

Because of the symmetry between the precision levels α of the experts, the default action for the diverse panel is

$$y_d^* = \frac{1}{2}$$

Using these results, we can establish the optimality of the homogeneous panel. Indeed, we find that

$$EU^J(x = (0, 0)) \approx -0.019 > -0.05 = EU^J(x = (0, 1))$$

As in Bhattacharya and Mukherjee (2013), we obtain that the homogeneous panel is optimal[6]. This is to be expected in the absence of bias. Let us now change this and consider the probability and level of bias to be $\delta = \epsilon = 0.4$. Again, we assume that both experts share these traits. Since this preserves the symmetry, the optimal default action for the diverse panel is left unchanged.

---

[6] See Proposition 6 in Bhattacharya and Mukherjee (2013).

Regarding the homogeneous panel, we obtain that

$$y_h^* = 1 - \frac{1 - \delta((\alpha\epsilon)^2 + 2\alpha\epsilon(1-\alpha))}{1 + (1-\alpha)^2} \approx 0.127$$

Immediately observe that the government, aware of the possibility of getting biased reports, adjusts its default decision accordingly. Under such conditions, the diverse panel is now optimal:

$$EU^J(x = (0,0)) \approx -0.119 < -0.103 = EU^J(x = (0,1))$$

In the next section, we further explain the results obtained in this section, and discuss some limitations of our model. We also explore another type of bias, and describe how our work relates to the literature.

2.5 Discussion

By now, we have established that there exists a cutoff ϵ^* above which the judge will prefer to form a panel of experts with opposed agendas. Moreover, that cutoff decreases in δ and α. Essentially, a larger expected bias results in a lower cutoff ϵ^*. In turn, a lower cutoff means that the judge is more likely to prefer a diverse panel composition. This nuance complements the intuition presented in Bhattacharya and Mukherjee (2013): one can think of the homogeneous panel as a singular expert who is more informed. When deciding on the panel composition, the judge balances the benefits from being more likely to get a report but on a shorter range of states (homogeneous panel), versus having a lower probability but on the entire set of states (diverse panel). Without bias, the homogeneous panel is unambiguously preferred for all values of α[7]. However, as α approaches 1, the difference between the two shrinks to the point were, at $\alpha = 1$, the judge is indifferent between the two compositions. With bias, it depends on ϵ^* as explained before, which in turn depends on α and δ. Now the aforementioned trade-off has an extra layer: For the judge, are

---

[7] See Bhattacharya and Mukherjee (2013)

the benefits of an higher probability of getting a report worth the loss from a smaller revelation set, when this report in itself is less valuable by being possibly biased? Indeed, conditional on experts being informed, the probability of getting a biased report from a homogeneous panel is $2\epsilon - \epsilon^2$, which is greater than ϵ for the diverse panel. Theorem 1 shows that there is a value for ϵ above which the judge prefers the less probable but more reliable report from the diverse panel.

For ease of presentation, we have focused on absolute utility functions for both the judge and the experts. Our results qualitatively hold – to some extent – if we were to consider a quadratic function instead (e.g. $u^J(y, \theta) = -(y - \theta)^2$ for the judge, and $u_i(y, x_i) = -(y - x_i)^2$ for the experts). Starting with the experts, recall that they play a rather passive role in our model. Indeed, they simply observe a signal and only reveal it if it lies in their respective revelation set. Therefore, this new utility form bears no significance for our results on their side. However, such a change would affect the judge's best response function. One can follow the same procedure as in Proposition 1 to see that if $s = 0$ or $s \in (1-\delta, 1]$, the best response remains unchanged. For $s \in (0, 1-\delta]$, it is now $y(s) = s + \epsilon\delta$. Proposition 7 formalizes the judge's best response function.

Proposition 7 *Consider a panel consisting of one expert with ideal actions $x = 0$. The judge's quadratic payoff function is $u^J(y, \theta) = -(y - \theta)^2$. Then the judge's best-response function upon observing a report $\widetilde{m} \neq \varnothing$ is given as*

$$y(s) = \begin{cases} \frac{\delta}{2} \text{ if } s = 0 \\ s + \epsilon\delta \text{ if } s \in (0, 1-\delta] \\ s \text{ if } s \in (1-\delta, 1] \end{cases}$$

This proof, along with the following counterparts for the other panel compositions, are omitted as they follow closely their respective propositions from our original model. Essentially for this panel, the first and last cases are identical. For $s \in (0, 1 - \delta)$, optimizing the interim payoff with the quadratic function gives the above result. Similarly, Proposition 8 presents the best response for the homogeneous panel.

Proposition 8 *Consider a panel consisting of two experts with the same ideal action $x_1 = x_2 = 0$. The judge's quadratic payoff function is $u^J(y, \theta) = -(y - \theta)^2$. Then the judge's best-response function upon observing a report profile $\widetilde{m} \neq (\varnothing, \varnothing)$ is given as*

$$y^*(\widetilde{m}) = \begin{cases} \frac{\delta}{2} \textit{ if } s = 0 \\ \overline{s} \textit{ if } \overline{s} \in (0, 1-\delta] \\ m_1 = m_2 = \theta \textit{ if } \overline{s} \in (1-\delta, 1] \end{cases} \tag{9}$$

where

$$\overline{s} = \begin{cases} \max\{s_1, s_2\} \textit{ if } (s_1, s_2) \neq (\varnothing, \varnothing) \\ s_k + \epsilon\delta \textit{ if } s_k \neq \{\varnothing\} \textit{ and } s_{-k} = \varnothing, k \in \{1, 2\} \end{cases}$$

Proposition 9 below covers the diverse panel scenario.

Proposition 9 *Consider a panel consisting of two experts with opposing ideal actions $x_1 = 0$ and $x_2 = 1$. The judge's quadradic payoff function is $u^J(y, \theta) = -(y - \theta)^2$. Then the judge's best-response function upon observing a report profile $\widetilde{m} \neq (\varnothing, \varnothing)$ is given as*

$$y(s_1, s_2) = \begin{cases} \frac{\delta}{2} \textit{ if } s_1 = 0 \textit{ and } s_2 = \varnothing \\ s_1 + \epsilon\delta \textit{ if } s_1 \in (0, 1-\delta] \textit{ and } s_2 = \varnothing \\ s_1 \textit{ if } s_1 \in (1-\delta, 1] \textit{ and } s_2 = \varnothing \\ s_2 - \epsilon\delta \textit{ if } s_2 \in [\delta, 1) \textit{ and } s_1 = \varnothing \\ s_2 \textit{ if } s_2 \in [0, \delta) \textit{ and } s_1 = \varnothing \\ 1 - \frac{\delta}{2} \textit{ if } s_2 = 1 \textit{ and } s_1 = \varnothing \\ \frac{s_1 + s_2}{2} \textit{ if } s_2 - s_1 = 2\delta \textit{ or } s_2 - s_1 = \delta \end{cases} \tag{10}$$

With these best responses, it can be shown that there exists values of ϵ that makes the diverse panel the better option for the judge.

Theorem 2 *If the judge has a quadratic utility function, then for some values of* $\alpha \in (0, 1)$ *and* $\delta \in (0, 1)$ *there exists values of* $\epsilon^* \in (\underline{\epsilon}, \overline{\epsilon})$ *for which the diverse panel is preferred by the judge.*

Again, we omit the formal proof as it follows the steps of Theorem 1. Notice that here there is a range of values for ϵ^*. This is due to the quadratic utility function having a nonlinear relationship with the parameters of our model. Appendix B provides details about the intuition behind Theorem 2. Therefore, instead of having a clear cutoff ϵ^* as in our results using the absolute payoff function, we get, in some of the cases that allow it, an interval for ϵ where the diverse panel is optimal. Further details about the specific characterization of the equilibrium are beyond the scope of this essay.

Next, we look at a modification of our definition of bias. So far, we have assumed that an informed expert would be prone to unconsciously select information sources that were more aligned with his ideal action, thus resulting in a bias directed towards it. Now, let us simply assume that an informed expert may make a mistake in his interpretation of the information, thus effectively removing the direction of the bias. More precisely, instead of observing a possibly biased signal $s = \theta \pm \delta \in [0, 1]$ with probability ϵ, an informed expert observes with certainty a signal that is uniformly distributed around θ. That is, $s \sim U[\theta - \min\{\delta, \theta, 1-\theta\}, \theta + \min\{\delta, \theta, 1-\theta\}]$. Including $\min\{\delta, \theta, 1 - \theta\}$ serves three purposes: (i) it takes care of the corner issues, (ii) as $\mathbb{E}[\theta|s] = \theta$ the best-response function of the judge is $y(s) = s$ for all $s \in \Theta$, and (iii) it makes for an intuitive interpretation for extreme states. Indeed, the closer θ is to one extreme of the interval, the smaller is the range of the uniform distribution. As such, the different information sources should agree more about it. Moreover, on average an expert is not expected to make a mistake. For example, with 99% of the 3000 scientific peer-reviewed articles agreeing that contemporary climate change originates from human activity[8], it is extremely likely that different experts will have the same information – and as such, in terms of our model, observe very similar signals about the state. Using the same (symmetrical) settings, we find that the results are similar to those in Bhattacharya and Mukherjee (2013). In fact, the judge's strategy – which can be resumed in this context by y^* – is

---

[8] See Lynas et al. (2021).

identical, but because of the possibility of an interpretation mistake by the expert(s) $EU^J(y^*)$ is reduced.

If we remove the bias in our model, experts can only decide whether or not they disclose the state if they observe it. This reverts our model back to Bhattacharya and Mukherjee (2013) and, as such, we obtain the same result that the homogeneous panel is optimal. Now in this context, take $\alpha = 1$. Then the homogeneous panel with $y^H = 1$ is equivalent for the judge to the diverse panel with $y^D = \frac{1}{2}$. In addition to $\alpha = 1$, let us allow experts to report intervals containing θ. This is a specific model in the spirit of what is presented more generally in Milgrom and Roberts (1986). Then, consistent with their Corollary 2, we get that the diverse panel is optimal. To see this, fix y^* and consider a homogeneous panel at 0 along with a diverse panel with one expert at 0 and another at 1. The homogeneous panel will disclose $I^H = [0, \theta]$ if $\theta \leq y^*$ and $I_0^H = [0, 1]$ otherwise. The diverse panel, however, will have the signal intervals of the experts – I^0 and I^1 for experts at 0 and 1 respectively – be such that $I^0 \cap I^1 = \{y^*\}$ for all $\theta \in \Theta$. Indeed, notice that for the expert at 0 – the case of the expert at 1 is analogous –, the lower bound of I^0 is his ideal action $x_1 = 0$. As such, he only has to decide on the upper bound $\overline{\theta}$. Take $\theta > y^*$. If $\overline{\theta} = \theta$ we have the aforementioned result. If $\overline{\theta} > \theta$ then $I^0 \cap I^1 = [\theta, \overline{\theta}]$ and the judge will choose an action in that interval. Since the lower bound minimizes the expert's absolute loss, the strategy $\overline{\theta} > \theta$ is dominated by $\overline{\theta} = \theta$.

One limitation of our model is that the bias is exogenous. Indeed, popular beliefs suggest that individuals with more extreme opinions may be more likely to be biased towards it. For example, suppose that $\epsilon(x)$ (or $\delta(x)$) is an increasing function of $|\frac{1}{2} - x|$. While this is beyond the scope of this chapter, intuitively this makes the choice of extreme experts less likely for the judge as there is now a cost possibly outweighing the benefits from choosing such panel. Also not covered in this essay is the scenario of costly information acquisition for the experts. Indeed, we have assumed here that observing the state is costless for an informed expert. If, for example, experts were required to exert some effort that affects their probability of being informed α, then it may be the case that for some cost function (depending on effort) the diverse panel is optimal. Briefly, an expert in a homogeneous (extreme) panel benefits from the other expert's effort, therefore allowing him to reduce his own. The other expert reacts similarly. This is not an issue in a diverse (extreme) panel because there is no such interaction between the experts. In fact, the expert at $x = 0$ suffer a greater loss if he does not observe the state while the other expert does and it is optimal for him to reveal it. Assuming that α is increasing in effort, this leads to $\alpha^D > \alpha^H$; the probability that an expert is informed in a diverse panel is greater than in a homogeneous panel. Then we get that the judge, for some cost function depending on effort, might prefer the diverse panel. The details of such results are left for future research.

2.6 Conclusion

Common intuition suggests that a decision-maker should ask different parties with diverging opinions about information relevant to decide on a policy. By creating competition among the experts consulted in this persuasion game, numerous research in the literature (e.g. Milgrom and Roberts (1986)) have shown that a diverse panel of experts with strongly opposed interests will lead to fully informed decisions. However, this is assuming that experts always possess the relevant information. Bhattacharya and Mukherjee (2013) show that if with some probability they are uninformed, then a homogeneous panel of extreme experts is instead optimal. In this chapter, we extend this latter article by including a degree of bias from the experts. Focusing on small directional error (i.e. error towards their own agenda), we provide a sufficient condition where the diverse panel is optimal for the decision-maker.

Changing the judge's absolute payoff function to a quadratic one affects the details of our results, but the spirit of our conclusion stays unchanged. There, in some instances the diverse panel may be the better choice. Next, we also looked at a generic bias that is not directed towards one's agenda. We show that it essentially only lowers the decision-maker's expected utility, but does not alter the result in Bhattacharya and Mukherjee (2013) regarding the optimality of the homogeneous panel. Finally, a significant limitation of our model is that for an informed expert – an expert who does get to observe the state –, it is costless or effortless for him to acquire the information. This nullifies any moral-hazard concerns that could otherwise be relevant to this problem.

Appendix A

Proof of Proposition 1

Let $a(s)$ denote the action chosen by the judge upon observing s.

If $s = 0$, then either the judge received an unbiased signal and $s = \theta$, or a biased signal and $\theta \in [0, \delta]$. Then the judge's expected payoff is

$$\begin{aligned} EU^J_{s=0} &= \epsilon \int_0^\delta -|a(0) - \theta| dF(\theta) \\ &= \epsilon \int_0^a -(a - \theta) dF(\theta) + \epsilon \int_a^\delta -(\theta - a) dF(\theta) \\ &= \frac{-\epsilon}{2}(2a^2 - 2a\delta + \delta^2) \\ \frac{\partial EU^J_{s=0}}{\partial a} &= 0 \implies a = \frac{\delta}{2} \end{aligned}$$

This is optimal since $\frac{\partial^2 EU^J_{s=0}}{\partial a^2} = -2\epsilon < 0$. Therefore, $a(0) = \frac{\delta}{2}$.

If $s \in (1 - \delta, 1]$, the judge knows that the expert is unbiased. Therefore, $a(s) = s$ for such signal.

If $s \in (0, 1 - \delta]$, then the judge's interim payoff once she has observed s is

$$\begin{aligned} EU^J_s &= -\epsilon|a(s) - (\theta - \delta)| - (1 - \epsilon)|a(s) - \theta| \\ &= -\epsilon|a(s) - (s + \delta)| - (1 - \epsilon)|a(s) - s| \end{aligned}$$

Suppose that $a(s) = s + d$, and $d \in \mathbb{R}$. Then

$$\begin{aligned} EU^J_s &= -\epsilon|(s + d) - (s + \delta)| - (1 - \epsilon)|(s + d) - s| \\ &= -\epsilon|d - \delta| - (1 - \epsilon)|d| \end{aligned}$$

Now we want to maximize the interim payoff EU_s^J. As such,

$$\begin{aligned}\max_d EU_s^J &= \max_d -\epsilon|d-\delta| - (1-\epsilon)|d| \\ &= \max_d \left\{ \max_{d\in(-\infty,0]} -\epsilon|d-\delta| - (1-\epsilon)|d|, \max_{d\in[0,+\infty)} -\epsilon|d-\delta| - (1-\epsilon)|d| \right\}\end{aligned}$$

This procedure is without loss of generality. Indeed, consider a fixed level of ϵ. Then $EU_s^J(d) \in (-\infty, +\infty)$. For some $d = d'$ such that d' corresponds to $\arg\max_d EU_s^J$, then either $d' \in (-\infty, 0]$ or $d' \in [0, +\infty)$ or $d' = 0$. As such, if for example $\max_d EU_s^J = \widetilde{EU}$ at $d' \in (-\infty, 0]$, then it must be that

$$\max_d \left\{ \max_{d\in(-\infty,0]} -\epsilon|d-\delta| - (1-\epsilon)|d|, \max_{d\in[0,+\infty)} -\epsilon|d-\delta| - (1-\epsilon)|d| \right\} = \widetilde{EU}$$

at that same d'. A similar argument can be made for the other cases.

For $d \in (-\infty, 0]$, observe that EU_s^J is decreasing in d. Therefore, $d = 0$ maximizes EU_s^J and the interim payoff is $-\epsilon\delta$.

For $d \in [0, +\infty)$, the maximum depends on ϵ:

- If $\epsilon > \frac{1}{2}$, then $d = \delta$ and $EU_s^J(d) = -(1-\epsilon)\delta$
- If $\epsilon < \frac{1}{2}$, then $d = 0$ and $EU_s^J(d) = -\epsilon\delta$
- If $\epsilon = \frac{1}{2}$, then $d \in [0, \delta]$ and $EU_s^J(d) = -\frac{\delta}{2}$

As such, the judge's best response $y^*(s)$ given a signal s is

$$y^*(s) = \begin{cases} \frac{\delta}{2} \text{ if } s = 0 \\ s \text{ if } s \in (0, 1-\delta] \text{ and } \epsilon < \frac{1}{2} \\ s + \delta \text{ if } s \in (0, 1-\delta] \text{ and } \epsilon > \frac{1}{2} \\ s \text{ if } s \in (1-\delta, 1] \end{cases}$$

■

Proof of Proposition 2

The judge's ex ante expected utility can be written as

$$\begin{aligned}
EU^J &= (1-\alpha)\int_0^1 -|y^* - \theta| dF + \alpha\left\{\epsilon \int_{y^*+\delta}^1 -|y^* - \theta| dF + (1-\epsilon)\int_{y^*}^1 -|y^* - \theta| dF\right\} \\
&\quad + \alpha\left\{\epsilon \int_0^{\delta} -\left|\frac{\delta}{2} - \theta\right| dF + \epsilon \int_{\delta}^{y^*+\delta} -|y^*(s) - \theta| dF + (1-\epsilon)\int_0^{y^*} -|y^*(s) - \theta| dF\right\} \\
&= (1-\alpha)\Bigg\{\int_0^{y^*} -(y^* - \theta) dF + \int_{y^*}^1 -(\theta - y^*) dF\Bigg\} \\
&\quad + \alpha\epsilon\Bigg\{\int_0^{\frac{\delta}{2}} -\left(\frac{\delta}{2} - \theta\right) dF + \int_{\frac{\delta}{2}}^{\delta} -\left(\theta - \frac{\delta}{2}\right) dF + \int_{\delta}^{y^*+\delta} -(\delta) dF + \int_{y^*+\delta}^1 -(\theta - y^*) dF\Bigg\} \\
&\quad + \alpha(1-\epsilon)\Bigg\{\int_0^{y^*} -(s - \theta) dF + \int_{y^*}^1 -(\theta - y^*) dF\Bigg\} \\
&= (1-\alpha)\Bigg\{\int_0^{y^*} -(y^* - \theta) dF + \int_{y^*}^1 -(\theta - y^*) dF\Bigg\} \\
&\quad + \alpha\epsilon\Bigg\{\int_0^{\frac{\delta}{2}} -\left(\frac{\delta}{2} - \theta\right) dF + \int_{\frac{\delta}{2}}^{\delta} -\left(\theta - \frac{\delta}{2}\right) dF + \int_{\delta}^{y^*+\delta} -(\delta) dF + \int_{y^*+\delta}^1 -(\theta - y^*) dF\Bigg\} \\
&\quad + \alpha(1-\epsilon)\Bigg\{\int_0^{y^*} -(\theta - \theta) dF + \int_{y^*}^1 -(\theta - y^*) dF\Bigg\} \\
&= \frac{1-\alpha}{2}\left\{-y^{*2} - (1-y^*)^2\right\} + \frac{\alpha\epsilon}{2}\left\{\frac{\delta^2}{2} - (1-y^*)^2\right\} - \alpha\epsilon\delta y^* - \frac{\alpha(1-\epsilon)}{2}\left\{-(1-y^*)^2\right\}
\end{aligned}$$

The above is continuously differentiable in y^*. As such, the first-order condition implies that

$$\frac{\partial EU^J}{\partial y^*} = (1-\alpha)(-2y^*+1) + \alpha\epsilon(1-y^*) - \alpha\epsilon\delta + \alpha(1-\epsilon)(1-y^*) = 0$$
$$\implies y^* = \frac{1-\alpha\epsilon\delta}{2-\alpha}$$

Second-order condition shows that EU^J is strictly concave:

$$\frac{\partial^2 EU^J}{\partial {y^*}^2} = \alpha - 2 < 0$$

and therefore $y^* = \frac{1-\alpha\epsilon\delta}{2-\alpha}$ is the unique solution that maximizes EU^J.

■

Proof of Proposition 4

The expected utility of the judge is given by

$$
\begin{aligned}
EU^J = {} & (1-\alpha)^2 \int_0^1 -|y^* - \theta| dF \\
& + \ \alpha^2 \Bigg\{ \epsilon^2 \Bigg[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*+\delta} -|s - \theta| dF + \int_{y^*+\delta}^{1} -|y^* - \theta| dF \Bigg] \\
& + \ 2\epsilon(1-\epsilon) \Bigg[\int_0^{y^*} -|s - \theta| dF + \int_{y^*}^{y^*+\delta} -|(s+\delta) - \theta| dF + \int_{y^*+\delta}^{1} -|y - \theta| dF \Bigg] \\
& + \ (1-\epsilon)^2 \Bigg[\int_0^{y^*} -|s - \theta| dF + \int_{y^*}^{1} -|y^* - \theta| dF \Bigg] \Bigg\} \\
& + \ 2\alpha(1-\alpha) \Bigg\{ \epsilon \Bigg[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*+\delta} -|s - \theta| dF + \int_{y^*+\delta}^{1} -|y - \theta| dF \Bigg] \\
& + \ (1-\epsilon) \Bigg[\int_0^{y^*} -|s - \theta| dF + \int_{y^*}^{1} -|y - \theta| dF \Bigg] \Bigg\}
\end{aligned}
$$

$$
\begin{aligned}
= {} & \frac{-(1-\alpha)^2}{2} \left[y^{*2} + (1-y^*)^2 \right] \\
& + \ \alpha^2 \Bigg\{ \epsilon^2 \left[\frac{-\delta^2}{4} - \delta y^* - \frac{1}{2}((1-y^*)^2 - \delta^2) \right] \\
& + \ 2\epsilon(1-\epsilon) \left[-\delta^2 - \frac{1}{2}((1-y^*)^2 - \delta^2) \right] \\
& - \ \frac{(1-\epsilon)^2}{2}(1-y^*)^2 \Bigg\} \\
& + \ 2\alpha(1-\alpha) \Bigg\{ \epsilon \left[\frac{-\delta^2}{4} - \delta y^* - \frac{1}{2}((1-y^*)^2 - \delta^2) \right] + (1-\epsilon) \left[\frac{-1}{2}(1-y^*)^2 \right] \Bigg\}
\end{aligned}
$$

First order condition to find y^*:

$$\frac{\partial EU^J}{\partial y^*} = 0$$
$$\implies -(1-\alpha)^2(2y^*-1) - \delta((\alpha\epsilon)^2 + 2\alpha\epsilon(1-\alpha))y^* + (2\alpha\epsilon - (\alpha\epsilon)^2)(1-y^*)$$
$$+\ (\alpha^2(1-\epsilon)^2 + 2\alpha(1-\alpha)(1-\epsilon))(1-y^*) = 0$$
$$\implies y^* = \frac{1-\delta((\alpha\epsilon)^2 + 2\alpha\epsilon(1-\alpha))}{1+(1-\alpha)^2} \tag{11}$$

■

Proof of Proposition 5

First consider the case where both are biased. The judge will receive two reports, s_1 and s_2. The difference between these reports is $s_2 - s_1 = 2\delta$. Thus, the judge is able to deduce that both are biased, and that the true state lies exactly at the midpoint of their reports, $\frac{s_1+s_2}{2} = \theta$. Hence the midpoint is the best response.

Second, if only one of the experts is biased, then the judge will not be able to identify which one is making a mistake. This is because regardless of which expert is biased, the difference in their reports will be δ. The judge knows that one of the experts is biased, and the other is not. The probability that expert 1 (at $x_1 = 0$) is biased and that expert 2 (at $x_2 = 1$) is not is $\epsilon(1-\epsilon)$. Then the true state is $\theta = s_1 + \delta = s_2$. Similarly, the probability that expert 1 is unbiased and expert 2 is biased is $\epsilon(1-\epsilon)$. For this the true state is $\theta = s_2 - \delta = s_1$.

Let β be the judge's belief that expert 1 is the one that is biased. Let $a(s_1, s_2)$ be the action chosen by the judge. Then the interim expected utility of the judge is given by

$$EU^J(a(s_1,s_2)|s_1,s_2) = -\beta|a(s_1,s_2) - \theta(s_2)| - (1-\beta)|a(s_1,s_2) - \theta(s_1)|$$

where $\theta(s_i)$ stands for $\theta = s_i$. The above can be reformulated as below by substituting the values of $\theta(s_i)$:

$$EU^J(a(s_1, s_2)|s_1, s_2) = -\beta|a(s_1, s_2) - s_2| - (1 - \beta)|a(s_1, s_2) - s_1|$$

Here the first term corresponds to the case where expert 1 is biased while expert 2 is not, and the second term is the reverse scenario. One can follow a similar procedure as in section 2.3.2 to find the optimal value of d when considering the optimal action $a^*(s_1, s_2)$. Consequently, we get that

$$a^*(s_1, s_2) = \begin{cases} s_1 \text{ if } \beta < \frac{1}{2} \\ s_2 \text{ if } \beta > \frac{1}{2} \\ \tilde{s} \in [s_1, s_2] \text{ if } \beta = \frac{1}{2} \end{cases}$$

In the current settings, using Bayes' rule we get that

$$\beta = \mathbb{P}(\text{expert 1 is biased} \mid \text{only one expert is biased}) = \frac{\epsilon(1-\epsilon)}{2\epsilon(1-\epsilon)} = \frac{1}{2}$$

Therefore, upon receiving the reports (s_1, s_2) such that $s_2 - s_1 = \delta$, the judge has no reason to believe that one expert is more likely to be biased than the other. As such, she is indifferent between any action $\tilde{s} \in [s_1, s_2]$. For consistency with the case where $s_2 - s_1 = 2\delta$, we will assume that in this case as well the judge will choose the midpoint, that is $a(s_1, s_2) = \frac{s_1+s_2}{2}$. Combining everything gives the judge's best response function.

■

Proof of Proposition 6

The expected utility of the judge is given by

$$
\begin{aligned}
EU^J = (1-\alpha)^2 & \int_0^1 -|y^* - \theta| dF \\
+ \ \alpha(1-\alpha) & \Bigg\{ \epsilon \left[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*+\delta} -|s_1 - \theta| dF + \int_{y^*+\delta} -|y^* - \theta| dF \right] \\
& \qquad + (1-\epsilon) \left[\int_0^{y^*} -|s_1 - \theta| dF + \int_{y^*}^1 -|y^* - \theta| dF \right] \Bigg\} \\
+ \ \alpha(1-\alpha) & \Bigg\{ \epsilon \left[\int_0^{y^*-\delta} -|y^* - \theta| dF + \int_{y^*-\delta}^{1-\delta} -|s_2 - \theta| dF + \int_{1-\delta}^1 - \left| 1 - \frac{\delta}{2} - \theta \right| dF \right] \\
& \qquad + (1-\epsilon) \left[\int_0^{y^*} -|y^* - \theta| dF + \int_{y^*}^1 -|s_2 - \theta| dF \right] \Bigg\} \\
+ \ \alpha^2 & \Bigg\{ \epsilon^2 \Bigg[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*-\delta} -|s_1 - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{s_1 + s_2}{2} - \theta \right| dF \\
& \qquad + \int_{y^*+\delta}^{1-\delta} -|s_2 - \theta| dF + \int_{1-\delta}^1 - \left| 1 - \frac{\delta}{2} - \theta \right| dF \Bigg] \\
& \quad + \epsilon(1-\epsilon) \Bigg[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*-\delta} -|s_1 - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{s_1 + s_2}{2} - \theta \right| dF \\
& \qquad + \int_{y^*+\delta}^{1-\delta} -|s_2 - \theta| dF + \int_{1-\delta}^1 -|s_2 - \theta| dF \Bigg] \\
& \quad + \epsilon(1-\epsilon) \Bigg[\int_0^{\delta} -|s_1 - \theta| dF + \int_{\delta}^{y^*-\delta} -|s_1 - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{s_1 + s_2}{2} - \theta \right| dF \\
& \qquad + \int_{y^*+\delta}^{1-\delta} -|s_2 - \theta| dF + \int_{1-\delta}^1 - \left| 1 - \frac{\delta}{2} - \theta \right| dF \Bigg] \\
& \quad + (1-\epsilon)^2 \Bigg[\int_0^{y^*} -|s_1 - \theta| dF + \int_{y^*}^1 -|s_2 - \theta| dF \Bigg] \Bigg\}
\end{aligned}
$$

$$
\begin{aligned}
&= (1-\alpha)^2 \int_0^1 -|y^* - \theta| dF \\
&\quad + \ \alpha(1-\alpha)\left\{ \epsilon \left[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*+\delta} -|\theta - \delta - \theta| dF + \int_{y^*+\delta} -|y^* - \theta| dF \right] \right. \\
&\qquad\qquad\qquad \left. + (1-\epsilon) \left[\int_0^{y^*} -|\theta - \theta| dF + \int_{y^*}^{1} -|y^* - \theta| dF \right] \right\} \\
&\quad + \ \alpha(1-\alpha)\left\{ \epsilon \left[\int_0^{y^*-\delta} -|y^* - \theta| dF + \int_{y^*-\delta}^{1-\delta} -|\theta + \delta - \theta| dF + \int_{1-\delta}^{1} - \left| 1 - \frac{\delta}{2} - \theta \right| dF \right] \right. \\
&\qquad\qquad\qquad \left. + (1-\epsilon) \left[\int_0^{y^*} -|y^* - \theta| dF + \int_{y^*}^{1} -|\theta - \theta| dF \right] \right\} \\
&\quad + \ \alpha^2 \left\{ \epsilon^2 \left[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*-\delta} -|\theta - \delta - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{(\theta-\delta) + (\theta+\delta)}{2} - \theta \right| dF \right. \right. \\
&\qquad\qquad \left. + \int_{y^*+\delta}^{1-\delta} -|\theta + \delta - \theta| dF + \int_{1-\delta}^{1} - \left| 1 - \frac{\delta}{2} - \theta \right| dF \right] \\
&\qquad + \epsilon(1-\epsilon) \left[\int_0^{\delta} - \left| \frac{\delta}{2} - \theta \right| dF + \int_{\delta}^{y^*-\delta} -|\theta - \delta - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{(\theta-\delta) + \theta}{2} - \theta \right| dF \right. \\
&\qquad\qquad \left. + \int_{y^*+\delta}^{1-\delta} -|\theta - \theta| dF + \int_{1-\delta}^{1} -|\theta - \theta| dF \right] \\
&\qquad + \epsilon(1-\epsilon) \left[\int_0^{\delta} -|\theta - \theta| dF + \int_{\delta}^{y^*-\delta} -|\theta - \theta| dF + \int_{y^*-\delta}^{y^*+\delta} - \left| \frac{\theta + (\theta+\delta)}{2} - \theta \right| dF \right. \\
&\qquad\qquad \left. + \int_{y^*+\delta}^{1-\delta} -|\theta + \delta - \theta| dF + \int_{1-\delta}^{1} - \left| 1 - \frac{\delta}{2} - \theta \right| dF \right] \\
&\qquad \left. + (1-\epsilon)^2 \left[\int_0^{y^*} -|\theta - \theta| dF + \int_{y^*}^{1} -|\theta - \theta| dF \right] \right\}
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{(1-\alpha)^2}{2}(-y^{*2} - (1-y^*)^2) \\
&\quad + \ \alpha(1-\alpha)\left\{\epsilon\left[\frac{-\delta^2}{4} - \delta y^* - \frac{1}{2}((1-y^*)^2 - \delta^2)\right] - \frac{(1-\epsilon)}{2}(1-y^*)^2\right\} \\
&\quad + \ \alpha(1-\alpha)\left\{\epsilon\left[\frac{1}{2}(\delta^2 - y^{*2}) - \delta(1-y^*) + \frac{\delta^2}{4}\right] - \frac{(1-\epsilon)}{2}y^{*2}\right\} \\
&\quad + \ \alpha^2\Bigg\{\epsilon^2\left[-\delta(y-2\delta) - \delta(1-y^*-2\delta)\right] + \epsilon(1-\epsilon)\left[-\delta(y^*-\delta) - \delta^2 - \delta(1-y^*-\delta)\right]\Bigg\}
\end{aligned}
$$

$$
= 3(\alpha\delta\epsilon)^2 + \frac{\alpha\delta^2\epsilon}{2} - \alpha\delta\epsilon - (1-\alpha)y^{*2} + (1-\alpha)y^* - \frac{(1-\alpha)}{2}
$$

First order condition to find y^*:

$$
\begin{aligned}
\frac{\partial EU^J}{\partial y^*} &= 0 \\
&\implies -2(1-\alpha)y^* = -(1-\alpha) \\
&\implies y^* = \frac{1}{2} \qquad (12)
\end{aligned}
$$

■

Proof of Theorem 1

Below are the functions $EU^D(y^D)$ and $EU^H(y^H)$.

$$
EU^D(y^D) = \frac{\alpha\epsilon\delta}{2}(6\alpha\epsilon\delta + \delta - 2) - \frac{(1-\alpha)}{4}
$$

Let $P = (\alpha\epsilon)^2 + 2\alpha\epsilon(1-\alpha)$ and $Q = 1 + (1-\alpha)^2$.

$$
EU^H(y^H) = \frac{-[(Q-1+\delta P)^2 + (1-\alpha)^2(1-\delta P)^2]}{2Q^2} - \frac{\delta P(1-\delta P)}{Q} + \frac{\alpha\epsilon\delta^2}{2}\left[\frac{5\alpha\epsilon}{2} - 3\alpha + 1\right]
$$

With these we now proceed with the proof. For simplification and ease of reading, we write $g(\epsilon; \alpha, \delta)$ as $g(\epsilon)$.

(i) The function $g(\epsilon)$ is continuously differentiable on the interval $[0, 1]$. As such,

$$\begin{aligned}
\frac{\partial^2 g(\epsilon)}{\partial \epsilon^2} &= \frac{\alpha^2\delta(-12\alpha^2\delta\epsilon^2 + 24\alpha^2\delta\epsilon - \alpha^2\delta - 24\alpha\delta\epsilon + 2\alpha\delta - 6\delta + 4)}{2(1 + (1-\alpha)^2)} > 0 \\
&\implies -12\alpha^2\delta\epsilon^2 + 24\alpha^2\delta\epsilon - \alpha^2\delta - 24\alpha\delta\epsilon + 2\alpha\delta + 6\delta + 4 > 0 \\
&\implies 4 + 2\alpha\delta + 24\alpha^2\delta\epsilon + 6\delta > 24\alpha\delta\epsilon + \alpha^2\delta + 12\alpha^2\delta\epsilon^2
\end{aligned}$$

which holds for any $\alpha, \delta \in [0, 1]$ and $\epsilon \leq \frac{1}{2}$.

(ii) Observe that $P(\epsilon = 0) = 0$. Then,

$$g(0) = \frac{-\alpha^2(1-\alpha)}{4[1 + (1-\alpha)^2]} < 0$$

which is strictly smaller than 0 for all $\alpha \in (0, 1)$.

(iii) Substituting $\epsilon = \frac{1}{2}$ into the function g, we get that

$$g\left(\tfrac{1}{2}\right) = \frac{\alpha^2(29\alpha^2\delta^2 - 52\alpha\delta^2 - 16\alpha\delta + 8\alpha + 60\delta^2 + 8\delta - 8)}{32(1 + (1-\alpha)^2)}$$

Solving for δ, we get that $g\left(\frac{1}{2}\right) > 0$ if

$$\delta > \frac{-4(1-2\alpha) + \sqrt{-232\alpha^3 + 712\alpha^2 - 960\alpha + 496}}{29\alpha^2 - 52\alpha + 60}$$

which is strictly positive for all $\alpha \in (0, 1)$.

(iv) Given (i), (ii) and (iii), we conclude that one of the two following cases must hold:

(a) $g\left(\frac{1}{2}\right) \leq 0$;

(b) $g\left(\frac{1}{2}\right) > 0$;

For (a), given some values of α and δ, $g\left(\frac{1}{2}\right) \leq 0$ for all $\epsilon \in (0, \frac{1}{2})$. Therefore we conclude that in this case the homogeneous panel is preferred (with indifference at $g\left(\frac{1}{2}\right) = 0$). For (b), (i), (ii) and (iii) implies that $g(\cdot)$ must intercept the horizontal axis only once on the interval $\epsilon \in \left[0, \frac{1}{2}\right]$. Let ϵ^* be that value of ϵ such that $g(\epsilon^*) = 0$. Thus, by continuity, for $\bar{\epsilon} \in [0, \epsilon^*)$, then $g(\bar{\epsilon}) < 0$ and we conclude that the homogeneous panel is preferred by the judge. However, if $\bar{\epsilon} \in \left(\epsilon^*, \frac{1}{2}\right)$, then $g(\bar{\epsilon}) > 0$ and the diverse panel is the better choice. At $\bar{\epsilon} = \epsilon^*$, the judge is indifferent between the two panel compositions (since $g(\bar{\epsilon} = \epsilon^*) = 0$).

■

Appendix B

We proceed to show details about our results in Theorem 2. The judge's payoff function is now quadratic. To illustrate the more complex interactions between $g(\cdot)$ and the parameters α, δ and ϵ, Figure 2.6 depicts how $g(\epsilon; \alpha, \delta)$ reacts to different values of δ, while keeping α fixed. Figure 2.7 sets α as the variable for $g(\cdot)$, and considers again different values of δ[9]. Interestingly, as shown in Figures 2.6, the function $g(\epsilon)$ is no longer convex.

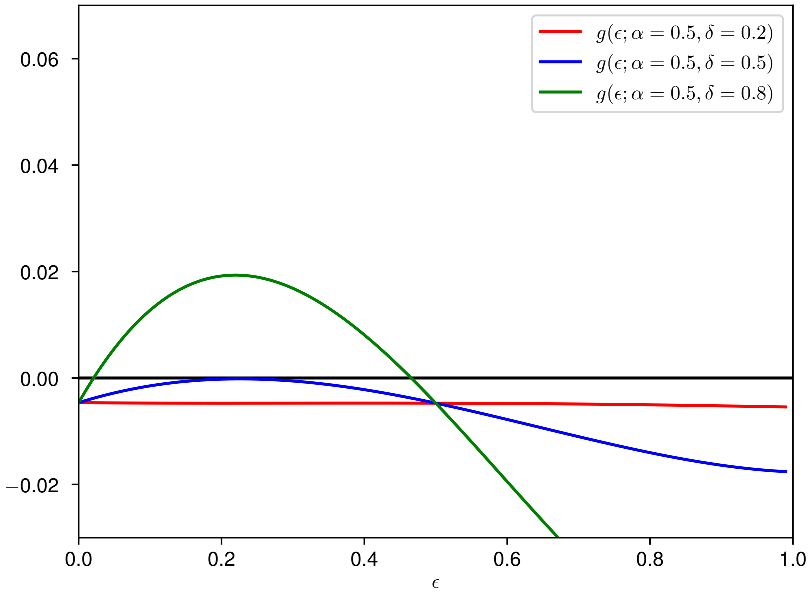

Figure 2.6: Function $g(\epsilon)$ given $\alpha = 0.5$ and $\delta \in \{0.2, 0.5, 0.8\}$

Importantly, Figures 2.6 and 2.7 demonstrate that depending on the parametric conditions, there may exists values of ϵ for which $g(\cdot)$ is positive. As such, for those values the diverse panel will lead to an higher expected utility for the judge than the homogeneous panel. Observe that as mentioned in Theorem 2, we now have a range of values for ϵ that makes the diverse panel the preferred composition for the judge. Intuitively, under the quadratic loss function the judge now adjusts her action according to the expected bias in the report. As such, for low bias the loss from

---

[9] Considering different levels of ϵ instead does not affect the graph drastically.

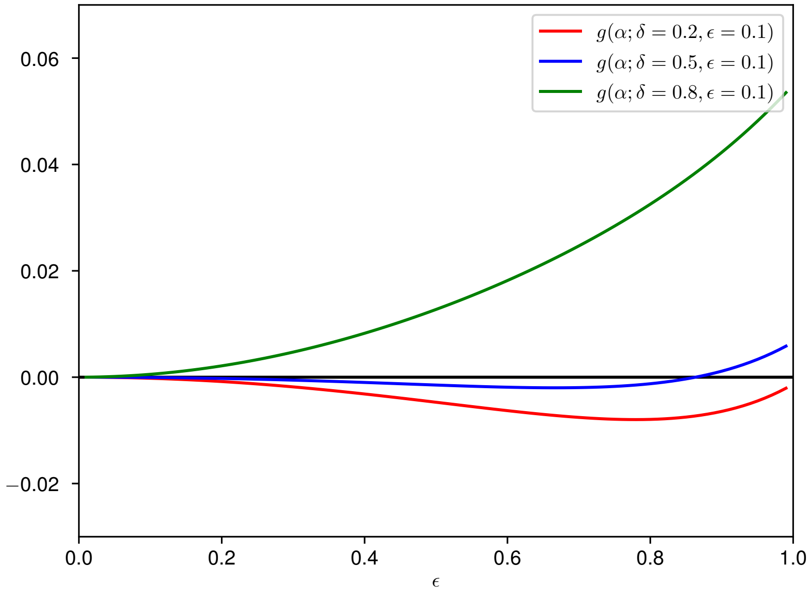

Figure 2.7: Function $g(\epsilon)$ given $\alpha = 0.5$ and $\delta \in \{0.2, 0.5, 0.8\}$

this adjustment – detrimental if the report is in fact the true state – is minimal and justified for risk aversion purposes. The homogeneous panel composition is then optimal. However, as the bias increases, this is no longer justifiable as the loss from deviating from the true state increases exponentially. At this point, the diverse panel is preferred as the deviation from the report to prevent the possible bias is too costly. Indeed, this refers to the aforementioned arbitrage between the lower occurrence of obtain more reliable information from the homogeneous panel, and the opposite from the diverse panel. As ϵ increases, the revelation set of the homogeneous panel shrinks and therefore becomes less attractive for the judge.

Chapter 3

Strategic Acquisition of Costly Information from Opinionated Experts

3.1 Introduction

Decision-makers often rely on experts to obtain information. Common intuition suggests that the decision-maker would benefit by consulting multiple experts, especially if these experts have opposing interests. This view is behind the design of panels of experts in many decision-making situations. For examples, a judge listens to experts from both the plaintiffs and the defendants, or a government, seeking to implement a new environmental policy, collects opinions from both the proponents and the opponents of environmental protection.

Starting with the seminal paper by Milgrom and Roberts (1986), the literature on persuasion has demonstrated that in many situations, the decision-maker benefits by considering a panel of experts with opposing interests (see also for example, Shin (1994), Shin (1998), Kamenica and Gentzkow (2011), Gul and Pesendorfer (2012)). On the other hand, in two important papers, Bhattacharya and Mukherjee (2013) and Bhattacharya et al. (2018) have shown that there may arise situations where it is optimal for the decision-maker to opt for like-minded experts with extreme agendas rather than experts whose preferences are opposed. Much of the literature has however focused on revelation of costless information[1]. There are, however, many scenarios where it is reasonable to assume that experts can acquire information at a cost. To illustrate this, consider a government-level environmental agency listening to different advocacy groups. Each of these groups – referred to as experts – can decide on how much evidence to gather, how much time to spend in assessing it, etc. All these processes are costly, and affect the quality of the information that an expert may possess.

The objective of the present chapter is to assess the impact of costly information acquisition in a multiple expert disclosure game. In particular, we ask the following questions: (i) "In a model where experts have extreme agendas, when is it the case that experts will invest in acquiring information?", and (ii) "Should a panel consist of experts with similar or opposed opinions?"

We find that the answers to the above questions are nuanced. We compare a panel that is composed of experts with opposed preferences – referred to as a *diverse* panel as in Chapter 2 –, with one where the experts have identical preferences – a *homogeneous* panel. The investment to acquire information is costly and affects the "quality" of the experts. That is, the probability for an expert to acquire relevant information. In this setting, we uncover situations where experts who are diametrically opposed undertake costly investment to improve their quality, while those who share a similar agenda do not. Moreover, we identify cases for which a diverse panel is better for the decision-maker than a homogeneous one.

As in Chapter 2, we follow closely the model introduced in Bhattacharya and Mukherjee (2013). A decision-maker has to choose an action in the interval $[0, 1]$. From the decision-maker's point of view, her optimal action is to match the state of the world $\theta \in [0, 1]$. Throughout this essay, we assume that θ is uniformly distributed. The objective of the decision-maker is to minimize the loss which is an increasing function of the distance between the state of the world and the action she takes. In our baseline model, the focus is on the quadratic loss function for the judge. The

---

[1] Some notable exceptions are Kartik et al. (2017), and Dewatripont and Tirole (1999). We discuss the related literature in details below.

decision-maker bases her action on verifiable reports from two experts. The experts have state-independent and monotonic preferences over the decision-maker's action. More importantly, an expert's preference is identified by his "*agenda* " – i.e. his most preferred action. We restrict our attention to the case where the most preferred action for the experts are either 0 or 1. We qualify experts with such preferences to be *extreme*. Each expert $i \in \{1, 2\}$ is privately informed of the state with a probability α_i. As in the previous chapter, we interpret α_i as the quality of expert i. So far the model is exactly the same as in Bhattacharya and Mukherjee (2013). However, unlike previous papers, in our present setup, quality is not exogenous. For each expert i, the quality parameter α_i can take two possible values $\alpha_i \in \{\alpha_L, \alpha_H\}$ with $0 < \alpha_L < \alpha_H < 1$. With an effort at a cost $c > 0$, an expert can choose quality level α_H. The cost for the low level of quality α_L is normalized to zero.

We focus on the *Perfect Bayesian Equilibrium* (PBE) of the resulting game. As in Bhattacharya and Mukherjee (2013), the equilibrium is characterized by the decision-maker's "default action" – the action that the decision-maker would choose in the event where no information is revealed by the experts. An uninformed expert – who did not observe any relevant information about the state – must admit ignorance. On the other hand, an informed expert – who did observe the state – with agenda 0 reveals states that are smaller than the default action. Likewise, an informed expert with agenda 1 will only report states that are higher than that.

We compare the diverse panel with a *left*-homogeneous panel – where both experts have agenda 0. We show that there exists cut off levels $\underline{c}$ and $\bar{c}$ that depend on α_L and α_H such that, whenever $c \in [\underline{c}, \bar{c}]$, experts in a homogeneous panel choose the lower level of quality α_L while both experts in a diverse panel choose the higher and more costly level of quality α_H. Moreover, there exists a cutoff level α^* such that, if $\alpha_H > \alpha^*$, then it is optimal for the decision-maker to choose a diverse panel.

It is immediate that the results depend critically on the risk attitudes of the decision-maker, the preferences of the experts, as well as the distribution of the state. In our baseline model, the preference of the decision-maker is characterized by quadratic loss function, while the preferences of the experts are characterized by absolute loss function. In addition, the state θ is uniformly distributed. So far we have been able to extend our result to cases where the decision-maker's preference is characterized by absolute loss function. However, what happens when we change the distribution of the state and /or the preferences of experts, remain open questions. We hope to pursue these questions in future research.

3.1.1 Related Literature

Our work is an extension of Bhattacharya and Mukherjee (2013) in which they consider a model with a judge and a panel of two experts. Each of these experts have an ideal action that he would like the judge to take. The judge's goal is to pick an action as close as possible to the state. The experts reveal the state of the world to the judge if they are informed and if doing so improves their payoff upon the default action – i.e. if the state lies within their revelation set. Bhattacharya and Mukherjee (2013) show that a panel of extreme experts – homogeneous or diverse – is always optimal. Moreover, if the judge has a quadratic loss function, then the homogeneous panel of extreme experts is preferred by the judge, regardless of the probability of the experts being informed. Bhattacharya et al. (2018) refer to an expert being informed or not as their type. As opposed to Bhattacharya and Mukherjee (2013) and our work, they allow for correlation between types – i.e. for a panel of two experts, if one expert is informed there is a higher probability of the other one is as well. Informed experts observe the state, but are biased in their reporting depending on their ideal actions. They show that with high correlation between types, the diverse panel is a better choice for the judge than the homogeneous panel. Moreover, if types are independent, then the homogeneous panel is generally the optimal choice.

The early works on persuasion games (See for example Milgrom (1981), Grossman (1981) and Milgrom and Roberts (1986)) show that when messages are verifiable, a skeptical decision-maker can elicit all information from a diverse panel of experts. However, the underlying assumption in those models is that every expert observes the state with certainty – an assumption that is absent in our model. In fact, the vehicle for costly information acquisition in our model is the choice of higher quality to observe the state (at a cost). As noted in, for example, Okuno-Fujiwara et al. (1990), Shavell (1994) and Shin (1994), if there is some positive probability that an expert does not observe the state, then an informed expert can conceal information by selective non-disclosure. Shin (1998) shows that a panel of two experts with opposing interests reveals more information than a panel of one unbiased expert. Similarly, Dewatripont and Tirole (1999) show the benefits of advocates (opposing experts) in policy-making, arguing that in many situations it leads to higher levels of information collection by the decision-maker over a single unbiased expert. Related to our work, Kartik et al. (2017) demonstrate that when information acquisition is costly and endogenous for the experts, then the experts' effort level decisions are strategic substitutes and increasing the number of experts may decrease these efforts. However, they focus on a discrete environment – with discrete states of the world and discrete signal realizations for the experts – in which the experts have preferences that are linear in the decision-maker's beliefs. Moreover, effort linearly increases the probability that an expert observes a signal about the state. Recently, in a general model, Gentzkow and Kamenica (2017) show that competition weakly increases information revelation if the information environment is *Blackwell-connected*, and that it otherwise may not be the case. Finally, Li and Norman (2021) argue that when having experts disclose their respective signal sequentially, adding an expert moving first weakly increases information transmission in equilibrium, and that having experts speak only once except for the first one does not change the equilibrium outcome. Moreover, having experts move sequentially makes for a weakly less informative equilibrium as opposed to the scenario where experts move simultaneously. All these papers are related to our work, and we have extensively made use of the previous results.

The rest of the chapter is organized as follows: in Section 3.2 we define the model; in Section 3.3 we characterize the equilibrium generally and for the different panel compositions; we present in Section 3.4 our main results on the optimality of a panel composed of opposed experts; in Section 3.5 we address the robustness of our results to changes in our model; we wrap up this essay with some closing thoughts in Section 3.6.

3.2 Model

We consider a model in which a decision-maker – referred to as the *judge* – must decide on the composition of a panel of experts, similar to Bhattacharya and Mukherjee (2013). There are two experts and an unknown state of the world $\theta \in \Theta = [0, 1]$. Moreover, we assume that θ is uniformly distributed (and denote the distribution as F). Each expert $i \in I = \{1, 2\}$ may observe θ with probability $\alpha_i \in (0, 1)$. We refer to this as the quality of an expert: an *informed* expert is one who observed the state, while an *uninformed* expert did not. The judge must decide on an action $y \in \Theta$ based on the report(s) from the experts. An informed expert may choose to report the state ($m_i = \theta$) or not ($m_i = \varnothing$), whereas an uninformed expert must admit ignorance ($m_i = \varnothing$). That is, misreporting θ is not permitted in our model. The goal of the judge is to select an action closest to the true state θ. To that end, we represent her utility function as

$$u^J(y, \theta) = -(y - \theta)^2$$

Each expert possesses an ideal action – also referred to as his agenda. We focus on *extreme* experts whose agenda $x_i \in \{0, 1\}$. Consequently, they want the judge to take an action closest to their preference:

$$u_i(x_i, y) = -|x_i - y|$$

For ease of presentation, we refer to the loss functions $v^J(y, \theta) = -u^J(y, \theta)$ and $v_i(x_i, y) = -u_i(x_i, y)$. We assume the experts to be risk-neutral. A panel of two experts can be *homogeneous* if the profile of agendas – defined as $x = (x_1, x_2)$ – is $x = (0, 0)$ or $x = (1, 1)$, or it can be *diverse* if $x = (0, 1)$.

An expert i may influence his probability of being informed by privately exerting some level of effort denoted e_i. That is, $\alpha_i(e_i)$ is increasing in e_i. The judge does not observe these effort levels. The experts privately know their own effort level, but not that of other experts. On the other hand, all experts know the panel composition. Moreover, we assume the same function α for both experts. To this effort is associated a cost $c(e_i)$, also increasing in e_i and common to the experts. We assume two levels of effort – i.e. $e_i \in \{0, e\}$ and $e > 0$ – such that $\alpha(0) = \alpha_L$, $\alpha(e) = \alpha_H > \alpha_L$, $c(0) = 0$ and $c(e) = c > 0$. Moreover, α_L and α_H lie strictly between 0 and 1.

We denote by $m = (m_1, m_2)$ the profile of messages submitted by the two experts. As we will establish, the judge's strategy for a given panel composition and for a given profile of qualities for the two experts simplifies to the choice of a default action y^*. This default action for the judge corresponds to her response to a non-informative report profile $m^0 = (\varnothing, \varnothing)$. Indeed, if one of the experts submits $m_i \neq \varnothing$, then the judge learns the true state and chooses $y(m) = m_i = \theta$ – and gets her maximum payoff $u^J(y(m), \theta) = 0$.

We use the Perfect Bayesian Equilibrium (PBE) solution concept. We denote $\mu(\theta|m)$ the posterior belief of the judge about the state upon observing the report profile m from the experts. We use $-i$ to refer to the expert other than expert i. A strategy profile $\langle m^*, y(m^*)\rangle$ and belief μ^* is a PBE[2] of the considered game if:

(i) Every expert chooses their level of effort simultaneously. For all $i \in I$, an expert i's level of effort e_i minimizes their expected loss given the levels chosen by the other experts:

$$e_i(e_{-i}) = \arg \min_{e_i \in \mathbb{R}_+} \mathbb{E}[v_i(y(m), x_i)]$$

(ii) For all $i \in I$, and for all $\theta \in \Theta$, if an expert i is informed then they report their signal $m_i = \theta$ if and only if:

$$\mathbb{E}[v_i(y(\theta, m_{-i})), x_i)] \leq \mathbb{E}[v_i(y(\varnothing, m_{-i})), x_i)]$$

(iii) For all report profile m, the judge's action minimizes their expected loss:

$$y(m) = \arg \min_{y \in \Theta} \mathbb{E}_{\mu^*}[v^J(y, \theta)]$$

---

[2] See Bhattacharya and Mukherjee (2013) for a discussion on the adaptation of the PBE formulation to the current problem.

(iv) $\mu^*(\theta|m)$ is obtained using Bayes' rule from the strategy of the experts m^* and the common priors. Moreover, any action off-the-equilibrium path that reveals θ must lead to degenerate beliefs on θ for the judge.

An informed expert i's strategy is characterized by his "revelation set" Θ_i^* that is the set of states over which he reports truthfully. Point (ii) leads to the characterization of expert i's *revelation set*. Given the single-peaked preferences for the experts, the revelation sets turn out to be intervals. Indeed, an expert with $x_i = 0$ will prefer to reveal the state to the judge if $\theta \leq y^*$ (recall that y^* is the default action for the judge) and choose to remain silent otherwise. Similarly, an expert with $x_i = 1$ will reveal θ if $\theta \geq y^*$. We provide further details about the equilibrium in the next section.

3.3 Equilibrium Characterization

In this section, we adapt the model presented in Bhattacharya and Mukherjee (2013) to account for effort levels exerted by the experts. We closely follow their presentation as it can easily be extended to our model.

Given an agenda profile $x = (x_1, x_2)$ and $x_i \in [0, 1]$ for all $i \in \{1, 2\}$, observe that if one of the experts reveals the state – i.e. $m_i = \theta$ for one expert $i \in \{1, 2\}$ –, then the judge optimally chooses it as her action – i.e. $y(m) = m_i = \theta$ – as it minimizes her ex post loss. In contrast, if the report profile $m = (m_1, m_2) = (\varnothing, \varnothing)$, then she chooses the default action $y(m) = y^*$. As such, given an agenda profile x for the panel, the judge's strategy is completely characterized by the default action, as noted also in Bhattacharya and Mukherjee (2013) and in Bhattacharya et al. (2018).

Indeed, fix the quality for each expert i, $\alpha_i \in \{\alpha_L, \alpha_H\}$. Note that given a profile of qualities (α_1, α_2), the best response for the judge, as well as the optimal actions by the experts, are given as in Bhattacharya and Mukherjee (2013)[3]. We restate their results for completion.

Result 1 *(Bhattacharya and Mukherjee (2013), Proposition 1) There always exists a PBE of this*

---

[3] See Proposition 1 in Bhattacharya and Mukherjee (2013)

game. Moreover, in any PBE of this game an informed expert's strategy is:

$$m_i^*(\theta) = \begin{cases} \theta \ \ \textit{if} \ \ \theta \in \Theta_i^* = \{\theta \in \Theta \mid \|\theta - x_i\| \leq \|y^* - x_i\|\} \\ \\ \varnothing \ \ \textit{otherwise} \end{cases}$$

and the judge's strategy is:

$$y(m) = \begin{cases} \theta \ \ \textit{if} \ \ m_i(\theta) = \theta \textit{ for some } i \\ \\ y^* \ \ \textit{otherwise} \end{cases}$$

where

$$y^* = \arg\max_{y' \in \Theta} \int u^J(y';\theta)dF(\theta|m = \varnothing)$$

and y^* is the unique solution to the above maximization problem.

As for the experts, an informed expert i's report m_i is only relevant if $m_{-i} = \varnothing$. Then, his decision will decide what will be the action of the judge between $y(m) = \theta$ or $y(m) = y^*$. As noted in the previous section, this leads to the definitions of the revelation sets Θ_i^* given an expert i's agenda x_i. Formally, if $x_i = 0$ then $\Theta_i^*(x_i, y^*) = [0, y^*]$. On the other hand, if $x_i = 1$ then $\Theta_i^*(x_i, y^*) = [y^*, 1]$. Given an agenda profile $x = (x_1, x_2)$, we can define a panel's revelation set Θ_{x_1,x_2}^* as the set of states for which at least one expert reveals θ. Therefore, for the left-homogeneous panel – for which $x = (0, 0)$ – the revelation set corresponds to $\Theta_{0,0}^*(y^*) = [0, y^*]$. Similarly, for the right-homogeneous panel – for which $x = (1, 1)$ – it is $\Theta_{1,1}^*(y^*) = [y^*, 1]$. For a diverse panel – for which $x = (0, 1)$ – this set is $\Theta_{0,1}^*(y^*) = [0, 1]$. These also illustrates the intuitive trade-off between a diverse panel composition ($x_1 \neq x_2$) and a homogeneous one ($x_1 = x_2$). Indeed, as pointed out in Bhattacharya and Mukherjee (2013), the question as to which panel is the better choice for the judge boils down to either (i) getting to observe the true state across the whole set of states but with a lower probability (diverse panel), or (ii) getting a higher probability of observing θ but on a smaller interval (homogeneous panel). It turns out that in their original model, the homogeneous composition is optimal[4]. Throughout this essay, we denote by h a (left) homogeneous panel and by d a diverse panel.

Our contribution to this model is the inclusion of costly effort on the experts' side. Now, since experts privately decide on how much effort they want to exert, the choice of panel by the judge may influence these levels of effort. To that end, we define below the optimal default actions y_h^* and y_d^* for the homogeneous panel with $x = (0, 0)$ and the diverse panel with $x = (0, 1)$ respectively[5]. Given a panel with agenda profile $x = (0, 0)$, the judge's ex ante expected loss given the expected effort profile $\hat{e} = (\hat{e}_1, \hat{e}_2)$ is given by

$$EL_h^J(y_h^*) = (1 - \alpha(\hat{e}_1))(1 - \alpha(\hat{e}_2)) \int_0^{y_h^*} v^J(y_h^*, \theta)dF(\theta) + \int_{y_h^*}^1 v^J(y_h^*, \theta)dF(\theta) \tag{1}$$

and y_h^* corresponds to the optimal default action y_h solving

$$y_h^* = \arg \min_{y_h \in [0,1]} EL_h^J(y_h) \tag{2}$$

Indeed, an expert i with $x_i = 0$ will report θ if he observes $\theta \in [0, y_h^*]$. At the opposite, such expert will never reveal a $\theta \in [y_h^*, 1]$. Observe that the judge's expected loss comes from the situations where she does not get any report. As such, the first term in (1) corresponds to the state lying in the interval where it would be revealed if it were observed by the experts, but both experts are not informed. The second term depicts the scenario in which the state would never be revealed.

Similarly, given a panel with agenda profile $x = (0, 1)$, the judge's ex ante expected loss given the expected effort profile $\hat{e} = (\hat{e}_1, \hat{e}_2)$ is given by

$$EL_d^J(y_d^*) = (1 - \alpha(\hat{e}_1)) \int_0^{y_d^*} v^J(y_d^*, \theta)dF(\theta) + (1 - \alpha(\hat{e}_2)) \int_{y_d^*}^1 v^J(y_d^*, \theta)dF(\theta) \tag{3}$$

---

[4] See Proposition 6 in Bhattacharya and Mukherjee (2013).

[5] The case for the homogeneous panel with $x = (1, 1)$ is analogous.

and y_d^* corresponds to the optimal default action y_d solving

$$y_d^* = \arg \min_{y_d \in [0,1]} EL_d^J(y_d) \tag{4}$$

The first term corresponds to θ belonging to the revelation set of the uninformed expert at 0, while the second term shows the case where it belongs to the revelation set of the uninformed expert at 1. Because we are changing only α_i from the model presented in Bhattacharya and Mukherjee (2013), we obtain the same functional forms for y_d^* and y_h^*. Indeed, the difference here is that α is a function of effort[6].

In the following subsections, we cover the different panel compositions individually. Using the specifications of our model, we derive the optimal default action for each case. We conclude this section by comparing the homogeneous and diverse panels, and identify the conditions for the latter being the better choice for the judge.

3.3.1 Homogeneous Panel

We begin by characterizing the equilibrium for the homogeneous panel, focusing on the agenda profile $x = (0, 0)$. The complementary case of $x = (1, 1)$ is analogous and therefore omitted. Throughout this section, we will refer to α_H and α_L for $\alpha(e_i > 0)$ and $\alpha(e_i = 0)$ respectively to simplify the presentation. Recall that we assumed $\alpha(e_i)$ to be the same function for all experts. Also, since experts are identical but for their agendas in the diverse panel, it is without loss of generality to declare that the experts effectively decide between $e_i = 0$ or $e_i = e' > 0$. That is, they share the same set of choices for their levels of effort.

We now proceed with writing down the expected loss of expert i in the context of our model,

---

[6] See Proposition 6 in Bhattacharya and Mukherjee (2013).

using α_i and α_j, $i \neq j$, to define each expert's probability of being informed:

$$\begin{aligned}
EL_i^E &= (\alpha_i + \alpha_j - \alpha_i\alpha_j) \int_0^{y_h^*} \theta dF(\theta) \\
&\quad + (1-\alpha_i)(1-\alpha_j) \int_0^{y_h^*} y_h^* dF(\theta) + \int_{y_h^*}^1 y_h^* dF(\theta) + c(e_i) \\
&= (\alpha_i + \alpha_j - \alpha_i\alpha_j)\frac{y_h^{*2}}{2} + (1-\alpha_i)(1-\alpha_j)y_h^{*2} + (1-y_h^*)y_h^{*2} + c(e_i) \qquad (5)
\end{aligned}$$

Observe that the probability that at least one expert observes the state is $\alpha_i\alpha_j + \alpha_i(1-\alpha_j) + \alpha_j(1-\alpha_i) = \alpha_i + \alpha_j - \alpha_i\alpha_j$. Then, with one of them reporting it to the judge if $\theta \in [0, y_h^*]$, both experts get a loss of $|x - \theta| = \theta$. Otherwise, they get $|x - y_h^*| = y_h^*$.

We know that if both experts in the homogeneous panel choose α_H, then this panel will always be preferred by the judge[7]. As such, we are interested in uncovering the conditions under which both experts will exert $e_i = e_j = 0$. This can also be described as a free-riding issue where in a panel with aligned agenda, no expert has an incentive to exert any effort and all of them rely on others to exert some.

Proposition 1 *In a homogeneous panel with $x = (0, 0)$, an expert will always prefer to exert no effort if the following condition is satisfied:*

$$\frac{1}{2}(1-\alpha_L)(\alpha_H - \alpha_L)y_h^{*2} \leq c \qquad (6)$$

Moreover, this expert will do so regardless of the level of effort exerted by the other expert.

Proposition 1 shows that for a sufficiently high cost of effort, every expert in a homogeneous panel will exert no effort independently of what the other expert chooses to do. This impacts the judge's choice of panel since some of the benefit from selecting a homogeneous panel – namely, the relatively higher probability of observing the state on a smaller interval $\Theta_{0,0}^* \subset [0, 1]$ – will be diminished. The proof is relegated to the Appendix.

---

[7] See Proposition 6 in Bhattacharya and Mukherjee (2013).

Corollary 1 *If both experts in an homogeneous panel with $x = (0, 0)$ exert the same low level of effort α_L, then the optimal default action for the judge is*

$$y_h^* = \frac{1}{2 - \alpha_L} \tag{7}$$

The proof of Corollary 1 is a straightforward application of the results in Proposition 6 from Bhattacharya and Mukherjee (2013) and is relegated to the Appendix. Next, we proceed with a similar analysis for the diverse panel.

3.3.2 Diverse Panel

We now characterize the equilibrium for the diverse panel consisting of the agenda profile $x = (0, 1)$. The expected loss of expert i with $x_i = 0$ is provided below:

$$\begin{aligned} EL_i^E &= \alpha_i \int_0^{y_d^*} \theta dF(\theta) + (1 - \alpha_i) \int_0^{y_d^*} y_d^* dF(\theta) \\ &\quad + \alpha_j \int_{y_d^*}^1 \theta dF(\theta) + (1 - \alpha_j) \int_{y_d^*}^1 y_d^* dF(\theta) + c(e_i) \\ &= \alpha_i \frac{y_d^{*2}}{2} + (1 - \alpha_i) y_d^{*2} + \alpha_j \frac{(1 - y_d^{*2})}{2} + (1 - \alpha_j)(1 - y_d^*) y_d^* + c(e_i) \end{aligned} \tag{8}$$

Each expert would prefer if the other does not exert any effort, as it only increases their expected loss by some constant. Indeed, from the perspective of expert i, one can see that increasing α_j would increase EL_i^E.

Now, we want to show that for some cost c, it is beneficial for both experts to exert the highest level of effort. Since the two share the same loss function, if one has an incentive to pick α_H, so does the other expert.

Proposition 2 *In a diverse panel with $x = (0, 1)$, both expert will prefer to exert some effort if the* *following condition is satisfied:*

$$\frac{1}{2}(\alpha_H - \alpha_L)y_d^{*2} \geq c \tag{9}$$

Proposition 2 states that if the cost of effort is sufficiently low, both experts will prefer to select a high level of effort. Intuitively, if one expert chooses to exert some effort and the other does not, then the judge will adjust her default action accordingly. This will result in a larger revelation set for the former expert. The proof is relegated to the Appendix.

Corollary 2 *If both experts in a diverse panel with $x = (0, 1)$ choose the same high level of effort, then the optimal default action for the judge is*

$$y_d^* = \frac{1}{2} \tag{10}$$

The proof of Corollary 2 is also relegated to the Appendix. With propositions 1 and 2, we have a range of admissible values for c corresponding to

$$\frac{1}{2}(1 - \alpha_L)(\alpha_H - \alpha_L)y_h^{*2} \leq c \leq \frac{1}{2}(\alpha_H - \alpha_L)y_d^{*2} \tag{11}$$

Proposition 3 *With an absolute loss function for the experts, there always exists a value of $c \in [\underline{c}, \overline{c}]$, with $\underline{c}$ and $\overline{c}$ corresponding to the L.H.S. and R.H.S. of (11) respectively, since $\underline{c} < \overline{c}$ for all α_L such that $0 < \alpha_L < \alpha_H < 1$.*

If c lies within the range prescribed in (11), then we see the homogeneous panel preferring to exert no effort – with $e = 0$ leading to $\alpha(0) = \alpha_L$ – and the diverse panel choosing the higher level of effort – with $e > 0$ leading to $\alpha(e) = \alpha_H$. The proof is relegated to the Appendix. With these two different levels of effort, we now demonstrate when the diverse panel composition is optimal for the judge.

3.4 Optimality of the Diverse Panel

We look at the conditions on the precision levels of the diverse and homogeneous panels, α^d and α^h respectively, that make the judge opt for the the former composition. Clearly, the judge perfers the diverse panel composition if it provides her with a lower expected loss than the homogeneous panel, given the levels of effort of the experts.

Proposition 4 *Given the effort profiles* $e^d = (e_1^d, e_2^d)$ *and* $e^h = (e_1^h, e_2^h)$ *leading to the quality profiles* $\alpha^h = (\alpha_L, \alpha_L)$ *and* $\alpha^d = (\alpha_H, \alpha_H)$ *for the diverse and homogeneous panels respectively, the judge prefers the diverse panel composition to the homogeneous one, that is*

$$EL_J^d(y_d^*, e^d) \leq EL_J^h(y_h^*, e^h) \tag{12}$$

if the precision levels of each panel, $\alpha^d = \alpha_H$ *and* $\alpha^h = \alpha_L$*, satisfy*

$$\alpha_H \geq 1 - 4\left(1 - \frac{1}{2 - \alpha_L}\right)^2 \tag{13}$$

and c *satisfies* (11).

The proof of this proposition is relegated to the Appendix. What Proposition 4 establishes is the following. Let $\alpha_H > \alpha_L$ and let c satisfy condition (11). Then for the diverse panel to be better for the judge than a homogeneous one, the difference between α_H and α_L has to be sufficiently high. Indeed, as a function of α_L, one can plot the cutoff level α^* such that for $\alpha_H > \alpha^*$, the diverse panel is optimal for the judge. This we demonstrate below in the following corollary.

Corollary 3 *With a quadratic loss function for the judge and an absolute loss function for the experts, there always exists a value of* $c \in \mathbb{R}_+$ *that satisfies* (11)*. Moreover, the judge prefers the diverse panel composition to the homogeneous one if the precision levels of each panel,* $\alpha^d = \alpha_H$ *and* $\alpha^h = \alpha_L$*, satisfy* (13) *and* c *satisfies* (11).

Observe that (13) is strictly increasing in α_L for $\alpha_L \in [0, 1)$ and that it always lie above the 45° line. Let $g(\alpha_L)$ to be the difference between the LHS of (13) and α_L as defined below:

$$g(\alpha_L) = 1 - 4\left(1 - \frac{1}{2 - \alpha_L}\right)^2 - \alpha_L \tag{14}$$

Notice that g is increasing in α_L on the interval $\alpha \in [0, \alpha_L^*]$, until it reaches its maximum at $\alpha_L^* = 3 - \sqrt{5} \approx 0.764$. Thereafter it is strictly decreasing on $\alpha \in [\alpha_L^*, 1]$. Intuitively, if experts have a default quality $\alpha_L = \alpha_L^*$, then for them to exert a positive level of effort it must yield the highest increase in precision $\alpha_H \geq \frac{1}{2}(3\sqrt{5} - 5) \approx 0.854$. Then, as $\alpha_L \geq \alpha_L^*$ gets closer to 1, this difference decreases until $\alpha_L = \alpha_H = 1$ where the choice of panel composition has no significance – the judge will always get a report from the experts. See Figure 3.1 for a graphical representation of g.

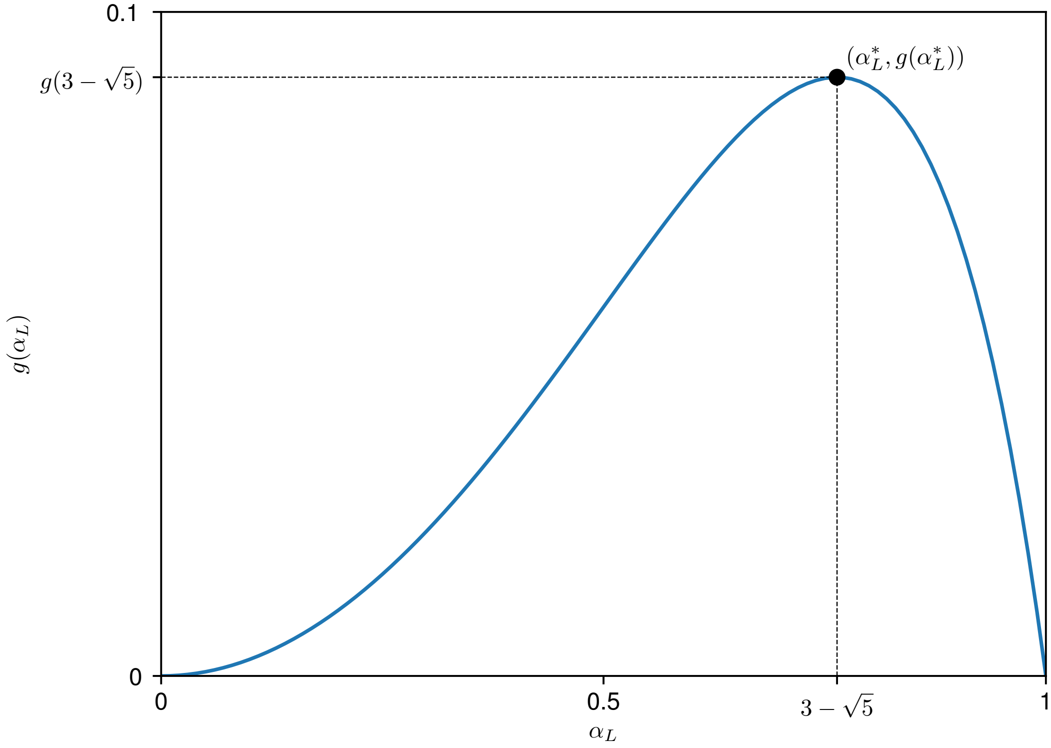

Figure 3.1: The required difference between α_H and α_L for an optimal diverse panel

Proposition 4 essentially depicts the following situation: if the return on effort is sufficiently large and at an appropriate cost such that this investment is beneficial when experts are in competition, then experts on a diverse panel will choose the higher level of effort. However, when experts share the same agenda, the inherent lack of competition generate a free-riding issue resulting in none of the experts exerting any effort. As such, the diverse panel can appear ex ante more informative to the judge, therefore making it the better choice. In the next section, we discuss some extensions and limits of our model.

3.5 Discussion

In this section, we discuss the robustness of our results to modifications of the judge's loss function assumed in our model. We also discuss the case of continuous effort levels.

3.5.1 Modifications to the Judge's Loss Function

We have shown that if conditions (11) and (13) are satisfied, then the diverse panel composition is optimal for the judge. These results critically depend on the risk attitudes of the judge, the preferences of the experts, as well as the distribution of the state. Below we provide a partial answer to these questions by considering the judge's preference to be captured by absolute loss function instead of the quadratic one. Observe that this change makes her risk neutral in such scenario. Suppose that the judge's preference is captured by the following function:

$$v_i^J(y,\theta) = -u^J(y,\theta) = |y-\theta|$$

Then, by repeating the steps in Proposition 4 for (13), we obtain that the diverse panel is preferred if

$$\alpha_H \geq 1 - 8\left(1 - \frac{1}{2-\alpha_L}\right)^2 \tag{15}$$

Note that condition (11) remains the same since the preferences of the experts are unchanged. Clearly, this new condition (15) is less stringent than its original counterpart (13). For example, in this case for any $\alpha_L \leq 0.7$, whenever $\alpha_H > \alpha_L$ and c satisfies (11), the judge would find a diverse panel better than a homogeneous one. In other words, the constraint on α_H given by (15) binds over a smaller region. This is not surprising given that Bhattacharya et al. (2018) reaches similar conclusions with regards to the judge's attitude towards risk.

3.5.2 Continuous Levels of Effort

Let us now consider the case where effort is continuous such that $e_i \in [0, 1]$ for all $i \in I$. Let $\alpha(e_i)$ and $c(e_i)$ be twice continuously differentiable functions on that same interval, common to all experts. Moreover, let us assume α to be concave and c to be convex in the effort levels. Finally, we assume the loss functions for both the judge and the experts to be described by quadratic loss functions. This is a departure from our baseline model in that there we have assumed that the experts' preferences are characterized by absolute loss function. This departure allows us to deal with differentiability issues in a tractable manner. The rest of our original model is kept intact. Observe that this does not change meaningfully the formulations for the expected utilities – one only has to consider α as a function of e here. Define e_i^h as the level of effort exerted by an expert i in a homogeneous panel.

Proposition 5 *In a homogeneous panel with $x = (0, 0)$, the optimal level of effort by expert i, e_i^h, is obtained by solving*

$$\frac{2}{3}(1 - \alpha(e_j^h))\alpha'(e_i^h)y_h^{*3} = c'(e_i^h) \tag{16}$$

Proposition 5 is the direct counterpart of Proposition 1, and its proof is therefore omitted. Equation (16) is obtained by deriving the first-order condition with respect to e_i^h from (5). The LHS of (16) represents the marginal benefit of effort for expert i, while the RHS gives its marginal cost. Therefore, if the former is smaller than the latter, then $e_i^{h*} = 0$. Moreover, notice that expert

i's level of effort depends on the action of expert j.

Corollary 4 *Fix $e_j^h = e \in [0,1]$. If for $e_i^h = 0$ we have*

$$\frac{2}{3}(1-\alpha(e_j^h))\alpha'(0)y_h^{*3} \leq c'(0) \tag{17}$$

then $e_i^{h*} = 0$ for all $i \in I$.

The proof of Corollary 4 is also omitted. From the concavity of α and the convexity of c, if at $e_i^h = 0$ we have that (17) holds then increasing e_i^h will still have (17) satisfied. Similarly, let e^d be the level of effort in a diverse panel.

Proposition 6 *In a diverse panel with $x = (0,1)$, the optimal level of effort by expert i, e_i^d is obtained by solving*

$$\frac{2}{3}\alpha'(e_i^d)y_d^{*3} = c'(e_i^d) \tag{18}$$

Again, the proof of Proposition 6 is straightforward and is therefore omitted. As opposed to the homogeneous panel, observe that expert i's effort is independent of e_j^d.

Corollary 5 *Suppose that the following holds:*

$$\frac{2}{3}\alpha'(0)y_d^{*3} > c'(0) \tag{19}$$

Then for a diverse panel with $x = (0,1)$, for each expert i there exists an optimal level of effort $e_i^d = e^{d*} > 0$ for all $i \in I$, where e^{d*} is given by (18).

Corollary 6 *Suppose that the following holds:*

$$\frac{2}{3}(1-\alpha(0))\alpha'(0)y_h^{*3} < c'(0) < \frac{2}{3}\alpha'(0)y_d^{*3}$$

Let $\alpha_H = \alpha(e^{d^*})$ *and* $\alpha_L = \alpha(0)$ *where* e^{d^*} *is given by* (18)*. Then experts in a left-homogeneous panel will choose* α_L *while experts in a diverse panel will choose* α_H*. Moreover, if*

$$\alpha_H \geq 1 - 4\left(1 - \frac{1}{2 - \alpha_L}\right)^2$$

then it is optimal for the judge to choose a diverse panel.

3.6 Conclusion

We consider an environment in which a decision-maker relies on information provided by experts to take an action. These experts may not always possess information, and have preferences regarding the action taken. Similar works on such models, notably Bhattacharya and Mukherjee (2013), show that a homogeneous panel of identical, extreme experts – whose opinions lie at one extremity of the interval of possible actions – is optimal for the decision maker. This is contrary to common intuition suggesting that competition among the agendas of the experts leads to more information being revealed – a claim also supported by Milgrom (1981). We show that if experts can exert effort – which comes at a cost – to increase their probability of observing relevant information, then it is possible to find levels of productivity and cost of effort such that a diverse, extreme panel is optimal – with experts having opinions at both extremities of the interval of actions. Under such conditions, the homogeneous panel exhibits a free-riding issue, where none of the experts will exert effort. All the while, experts in the diverse panel will choose a positive level of effort that if sufficiently large, may make such panel the preferred option for the decision-maker. We show the robustness of these results in environments where the decision-maker is risk-neutral or risk-averse and effort is binary. We also cover the case of continuous effort with both the judge and the experts being risk-averse, and obtain similar conclusions.

Appendix

Proof of Proposition 1

To show this, we proceed by deriving the best response of an expert i while fixing the level of effort of expert j. For ease of presentation, we will define $i = 1$ and $j = 2$.

Case 1: $\alpha_2 = \alpha_L$

We start by assuming that expert 2 does not exert any effort. Then, if expert 1 chooses $\alpha_1 = \alpha_H$, he gets

$$EL_1^E = (\alpha_H + \alpha_L - \alpha_H\alpha_L)\frac{y_h^{*2}}{2} + (1-\alpha_H)(1-\alpha_L)y_h^{*2} + (1-y_h^*)y_h^{*2} + c \tag{20}$$

where $c(e_i > 0) = c > 0$. If he instead chooses α_L, he gets

$$EL_1^E = (2\alpha_L - \alpha_L^2)\frac{y_h^{*2}}{2} + (1-\alpha_L)^2 y_h^{*2} + (1-y_h^*)y_h^{*2} \tag{21}$$

Now, we say that we observe free-riding if (20) $\leq$ (21). That is

$$(20) \leq (21) \implies \frac{1}{2}(1-\alpha_L)(\alpha_H - \alpha_L)y_h^{*2} \leq c \tag{22}$$

If condition (22) is satisfied, then $\alpha_1 = \alpha_L$ is a best response for expert 1 if $\alpha_2 = \alpha_L$.

Case 2: $\alpha_2 = \alpha_H$

Next we assume that expert 2 does exert a positive level of effort. Then, if expert 1 chooses $\alpha_1 = \alpha_H$, he gets

$$EL_1^E = (2\alpha_H - \alpha_H^2)\frac{y_h^{*2}}{2} + (1-\alpha_H)^2 y_h^{*2} + (1-y_h^*)y_h^{*2} + c \tag{23}$$

where $c(e_i > 0) = c > 0$. If he instead chooses α_L, he gets

$$EL_1^E = (\alpha_L + \alpha_H - \alpha_L\alpha_H)\frac{{y_h^*}^2}{2} + (1-\alpha_L)(1-\alpha_H){y_h^*}^2 + (1-y_h^*){y_h^*}^2 \tag{24}$$

Now, we say that we observe free-riding if (24) $\leq$ (23). That is

$$(24) \leq (23) \implies \frac{1}{2}(1-\alpha_H)(\alpha_H-\alpha_L){y_h^*}^2 \leq c \tag{25}$$

If condition (25) is satisfied, then $\alpha_1 = \alpha_L$ is a best response for expert 1 if $\alpha_2 = \alpha_H$. Observe that since $\alpha_L < \alpha_H$, (22) $\implies$ (25) as (22) is stricter.

■

Proof of Corollary 1

With $\alpha_i = \alpha_j = \alpha_L$, (1) can be written as

$$\frac{(1-\alpha_L)^2{y_h^*}^3}{3} + \frac{(1-y_h^*)^3}{3} \tag{26}$$

Deriving w.r.t. y_h^* and solving the FOC, we get

$$EL_h^J(y_h^*) = (1-\alpha_L)^2{y_h^*}^2 - (1-y_h^*)^2 = 0$$
$$\implies y_h^* = \frac{1}{2-\alpha_L}$$

■

Proof of Proposition 2

Keeping the same notation as in the proof of Proposition 1, observe that fixing α_2 results in its associated terms to be treated as constants. So we have that

$$\begin{aligned} & EL_1^E(\alpha_1 = \alpha_H, \alpha_2) \geq EL_1^E(\alpha_1 = \alpha_L, \alpha_2) \\ \implies & -(\alpha_H - \alpha_L)\frac{y_d^{*2}}{2} + (\alpha_H - \alpha_L)y_d^{*2} \geq c \\ \implies & (\alpha_H - \alpha_L)\frac{y_d^{*2}}{2} \geq c \end{aligned} \tag{27}$$

■

Proof of Corollary 2

With $\alpha_i = \alpha_j = \alpha_H$, (3) can be written as

$$EL_d^J(y_d^*) = \frac{(1-\alpha_H)y_d^{*3}}{3} + \frac{(1-\alpha_H)(1-y_d^*)^3}{3} \tag{28}$$

Deriving w.r.t. y_h^* and solving the FOC, we get

$$\begin{aligned} & (1-\alpha_H)y_h^{*2} - (1-\alpha_H)(1-y_h^*)^2 = 0 \\ \implies & y_d^* = \frac{1}{2} \end{aligned}$$

■

Proof of Proposition 3

To demonstrate this, we show that

$$\frac{1}{2}(1-\alpha_L)(\alpha_H - \alpha_L)y_h^{*2} \leq \frac{1}{2}(\alpha_H - \alpha_L)y_d^{*2} \tag{29}$$

Using the values of y_h^* and y_d^* in Corollaries (1) and (2), we have

$$\frac{1-\alpha_L}{(2-\alpha_L)^2} \leq \frac{1}{4}$$
$$\implies 4(1-\alpha_L) \leq (2-\alpha_L)^2$$

which holds for any value of $\alpha_L \in \mathbb{R}$.

■

Proof of Proposition 4

From (12), we get

$$\frac{1}{3}\left((1-\alpha_1^d)y_d^{*3} + (1-\alpha_2^d)(1-y_d^*)^3\right) \leq \frac{1}{3}\left((1-\alpha_1^h)(1-\alpha_2^h)y_h^{*3} + (1-y_h^*)^3\right)$$

Proposition 1 and Proposition 2 lead to $\alpha_1^h = \alpha_2^h = \alpha^h = \alpha_L$ and $\alpha_1^d = \alpha_2^d = \alpha^d = \alpha_H$. Therefore, substituting $y_d^* = \frac{1}{2}$ and $y_h^* = \frac{1}{2-\alpha^h}$, we have

$$\frac{1-\alpha^d}{4} \leq \frac{(1-\alpha^h)^2}{(2-\alpha^h)^3} + \frac{(1-\alpha^h)^3}{(2-\alpha^h)^3}$$
$$\implies \alpha^d \geq 1 - 4\left(1 - \frac{1}{2-\alpha^h}\right)^2 \tag{30}$$

If c satisfies (11), then the above implies

$$\implies \alpha_H \geq 1 - 4\left(1 - \frac{1}{2-\alpha_L}\right)^2$$

■

Chapter 4

Information Design with Multiple Privately-Informed Receivers

4.1 Introduction

We study the problem of one agent – the information designer – having to design a communication rule that will recommend some other agents – the receivers – to each take a specific action. These agents have private information – referred to as their types – which they can report, truthfully or not, to the information designer. The latter also has her own private information – also referred to as the state of the world. What makes this setup different than the standard *cheap talk* problems is that the information designer must create this communication device prior to receiving the reports and learning the state, and cannot alter in any way the resulting action recommendations – this is described as committing to the device in both the information design and the persuasion literature.

One practical example of such situation is that of a political or a company leader who wants to convince as many agents as possible to vote in favor of a new project, while waiting for the approval of an internal committee – i.e. the state of the world. It is equally likely that the committee approves it or not. Suppose that both the leader and the agents start with the same information. The leader is asked to prepare their argument in advance and run it by the legal department – making it impossible for the leader to alter her message to the agents. The leader prefers the project to be implemented regardless of what the agents think of it, while each agent may approve it or not – which corresponds to their type and is also assumed to be uniformly distributed. She then gathers everyone's opinion about the project, along with the internal committee's approval and concerns, and delivers the final voting recommendation to each agent individually.

In this example, the issue from the leader's perspective is how to tailor the recommendation to each agent and each type such that it maximizes the probability that the project is accepted. To accomplish that, the leader has to design a message that will provide the agents with an incentive to report their type truthfully and to follow the recommendation – this is referred to as the *incentive compatibility* constraint in the literature. This is therefore closely related to the level of alignment between the leader's preferences and the agents' ones. On the one hand, if these are perfectly aligned – everybody prefers the project to be accepted in every realization of the agents' individual types and the state – then there is no reason for the leader to conceal any information. Therefore, full revelation would be optimal here. On the other hand, if the preferences are completely misaligned – the agents always prefer the project to be discarded – then there is nothing that the leader can do to convince them otherwise and revealing information serves him no purpose. However, between these two extreme cases lies the more interesting aspect of this problem. If preferences are *partially* (mis)aligned – for example, some agents belong to the extreme cases mentioned previously while others may have preferences that differ depending on their own type and the internal committee's decision – then it may be the case that partial revelation – i.e. revealing information without making obvious what are the state and other agents' true types – improves the leader's expected utility. In this scenario, the latter could choose to fully disclose the level of approval from the committee if it is positive – by an unequivocal action recommendation to vote in favour –, and send a mixed signal – with some probability recommend to vote in favour and with some other probability to vote against – if the committee's verdict is neutral or negative. Again, the main issues here are that the communication rule – the conditional probability distributions over the action profiles for each agent's opinion and each possible report from the internal committee – must be such that every agent has no incentive to lie about their opinion regarding the project, while also preferring to obey the voting recommendation from the leader.

The main objective of this chapter is to characterize the solution of the information design problem with multiple agents. We do so by considering a binary environment: one designer, two receivers, two states of the world, with each receiver having two possible types and two possible actions. Both the state and the types are independently and uniformly distributed. Moreover, we focus on private messages. This essay also aims to provide comparative statics regarding the impact of the private information held by the different agents on the design of an incentive compatible communication rule. Regarding the information held by the receivers, assuming a uniform distribution over their types corresponds to the highest amount of private information. Indeed, such distribution makes it hardest for the designer to accurately predict the receivers' types. At the opposite end of this assumption, we have the case of a degenerate distribution which falls back to the problem of information design without private information.

In our main result, we characterize the optimal communication rule under different payoff functions for the sender and the receivers. We show that the level of alignment between their preferences affects the level of information revealed. To that end, we regroup our solutions into three categories: full revelation, no revelation, and partial revelation. Full revelation occurs when every receiver, after observing their recommended action from the disclosure mechanism, knows what the true state is with certainty. No revelation relates to the opposite situation where the communication rule leaves every receiver's prior belief unchanged – i.e. it does not lead to any updating. In between these two cases lies the case of partial revelation. The primary objective of this chapter is to present the conditions that lead the optimal disclosure rule to fall into one of these three categories.

4.1.1 Related Literature

There is an extensive literature on communication with multiple receivers, which includes Farrell and Gibbons (1989), Caillaud and Tirole (2007) and Goltsman and Pavlov (2011) to name a few. Recent papers, such as Inostroza and Pavan (2023), Zeng (2023) and Ziegler (2023) for example, discuss the problem of persuading multiple audiences. Candogan and Strack (2023) and Guo (2019), among others, look at information design and persuasion models, respectively, with privately-informed receivers.

Related to our work, Glazer and Rubinstein (2004) explore the case of a privately-informed speaker (sender) that wishes to influence the actions of an uninformed listener (receiver). The speaker does so by designing a mechanism that sends messages given her type with the objective to maximize her expected utility, but achieves this by minimizing the probability of the listener committing a mistake (choosing an undesired action). One aspect that contrast with the literature is that the commitment to the disclosure mechanism comes from the listener instead of the speaker. Commitment to the mechanism is crucial to distinguish the current problem with the *cheap talk* literature, as it prevents the information designer from deferring to an alternative message once she observes the state of the world – sometimes referred to as the sender's type. Rayo and Segal (2010) notably introduces commitment as they study the problem of a sender who is a priori uninformed and must commit to a disclosure mechanism to influence the action of an uninformed receiver. To do so, the sender must design a persuasion mechanism that conditions on the set of states to generate some messages – denoted as signals – intended to the receiver. The goal of these is to influence the receiver's equilibrium choice of action in a favourable manner from the sender's perspective, by the mean of modifying their beliefs regarding the true state of the world. They characterize the optimal disclosure rule – obtained by solving the linear program that maximizes the sender's expected utility – under the assumption that the receiver's private reservation value was drawn from a uniform distribution. Continuing on the subject of persuasion games, Kamenica and Gentzkow (2011) use a "concavification" method to find the optimal signal structure when there is one sender and one uninformed receiver. This method requires the definition of the sender's indirect utility from the receiver's posterior beliefs under the disclosure mechanism, and then defines the optimal mechanism as the smallest concave closure over this indirect utility function. More recently, Kolotilin (2018) tackles a similar problem but with a privately-informed receiver instead. In his paper, both the state and the receiver's private information – i.e. type – are continuous intervals, and their utility is continuous and single-crossing with respect to the messages. The optimal mechanism is found by the mean of a linear program maximizing the sender's expected utility subject to a feasibility constraint – that is, the rule creates beliefs that are consistent with the priors – and an obedience constraint – with two actions considered here, optimality requires the agent to be indifferent between the two of them. Using this approach, he considers the problem of a grade-disclosure policy to which schools should adhere for employment statistics purposes. In a similar setup, Kolotilin et al. (2017) characterize the optimal solution, but with the addition that the receiver has to report their type – and possibly misreport it – to the sender who then takes it into account when designing his disclosure rule. In their model, they describe the necessary and sufficient conditions for an optimum and characterize the solution as an experiment – which does not depend on the report. Moving closer to the mechanism design field, Bergemann and Morris (2016) look at information structures and their relative sufficiency to compare them. More important to the current project is the introduction of the *Bayes Correlated Equilibrium* (BCE) concept, which relies on an *obedience* constraint to define the set of possible equilibria. This expansion of the notion of correlated equilibrium in incomplete information games to account for incentive constraints plays a crucial role in our analysis. Using this solution concept, Bergemann and Morris (2017) examine multiple cases from the simple one sender and one receiver, to the more complex one sender and multiple privately-informed receivers, and describe the sets of possible information structures for each of them. Finally, Taneva (2019) solves the one sender and two (uninformed) receivers problem by setting up the extensive form of the game. She also characterizes the optimal solution under different payoff combinations, and discusses the impact of the payoff structures on the solution. Our work follows closely her framework and aims at extending the results to privately-informed receivers.

The present chapter proceeds as follow: Section 4.2 covers the generic model of this essay; Section 4.3 presents numerical examples and their solutions; Section 4.4 contains our main results in a binary environment.

4.2 Model

There is one sender – denoted by 0 – and 2 receivers – denoted by $i \in I = \{1, 2\}$. Denote the set of players as $N = I \cup \{0\}$. Each receiver has a finite set of actions $A_i = \{a_i^0, ..., a_i^k\}$. Generic actions are represented by a_i and $a = \{a_i\}_{i \in I}$ refers to a generic action profile. We denote by A the set of all action profiles. There is a finite set of states of the world Θ with θ denoting a generic element of that set. The sender and the receivers share a common prior $\phi \in \text{int}(\Delta(\theta))$ that is commonly known. Each receiver i has payoff relevant private information, summarized by his *type* t_i drawn from a finite set T_i. We denote by $t = (t_1, t_2)$ a type profile and by T the set of all type profiles. For each possible state of the world, the types of the receivers are drawn according to the distribution $\pi(t|\theta)$. Therefore $\pi : \Theta \to \Delta(T)$, which is also commonly known.

The sender (also referred to as the designer) has a utility function $v : \Theta \times A \to \mathbb{R}$ that depends on the actions taken by the receivers and the realized state of the world. Notice that the payoff for the sender does not depend directly on the types of the receivers. Her expected payoff, however, does depend on them through the disclosure mechanism which influences the actions chosen by the receivers for various realizations of their respective payoff relevant types.

The payoff of each receiver, on the other hand, depends on the realized state of the world, the payoff relevant type profile, and the action profile chosen by all receivers. As such, the payoff of each receiver i is a function: $u_i : \Theta \times T \times A \to \mathbb{R}$. Note that it depends only on i's realized type and not on the entire type profile. In that sense, this is a *private values* model.

As is standard in the literature (for example, see Gossner (2000)), it is possible to separate the problem into two parts. The first is a "basic game" – denoted G – which consists of the set of players I, their respective action sets A_i and utility functions u_i, along with a full support common prior ψ over the states of the world Θ shared by all players. In other words, the basic game is $G = \big((A_i, u_i)_{i=1}^{I}, \Theta, \psi\big)$. In addition to the basic game, as in Bergemann and Morris (2017), there is an *information structure* S that consists of the set of type profiles T as the space of signals and, for each possible realization of the state of the world θ, the conditional prior distribution over type profiles – i.e. $S = (T, \{\pi(\cdot \mid \theta)\}_{\theta\in\Theta})$. Together the basic game G and the information structure S define a game of incomplete information (G, S).

We are interested in the problem of the sender who has the obligation to commit to provide players with *additional* information, if any, in order to induce them to make specific choices. Viewed as an extensive form game between the sender and the receivers, and following Bergemann and Morris (2017), the timing of the game is as follows:

1. Sender picks and commits to a correlated communication device;
2. Nature reveals θ to Sender, and t_i to Receiver i;
3. Each Receiver i sends a message about t_i to the information designer;
4. Each Receiver observes an action recommendation *privately* according the communication device/rule;
5. Each Receiver chooses an action a_i based on his true type, and the recommendation;
6. Payoffs are realized.

At this point, there is no restriction on the messages transmitted between the two parties. However, following Bergemann and Morris (2016), we will employ a *Revelation Principle* argument and restrict our attention to direct messages where the receivers only report types, and the sender only submit action recommendations. Since we focus on finding the *best* equilibrium, in our settings this restriction is without loss of generality[1]. However, in a general sense, this may not always

---

[1] See Proposition 2 in Taneva (2019)

be the case[2]. Given this restriction to direct messages, the sender's problem is to choose among communication devices which we call *decision rules*.

Definition 1 *A communication rule is a function* $f : \Theta \times T \to M$*, where* $M = A$ *is a finite set of (direct) messages. That is, for every type profile* t *and for every state of the world* θ*,* $f(a|t,\theta)$ *is the probability that the sender recommends the action profile* $a = (a_1, a_2)$*, such that for all* (t, θ)*,*

$$\sum_{m \in M} f(a|t,\theta) = 1.$$

Moreover,

$$0 \leq f(x|t,\theta) \leq 1 \quad \forall x \in A, \quad ,\forall t \in T, \quad \forall \theta \in \Theta.$$

Definition 2 *A communication rule* f *is direct if* $M = \Delta(A)$*, with* $m = (m_1, m_2) = (a_1, a_2) = a$*. It then sends action recommendations conditional on* t *and* θ*.*

Under these settings, every receiver i must report a type (any) and choose an action after observing the sender's recommendation, which is communicated privately to them. Consequently, any decision rule must provide the receivers with incentives to report their types *truthfully* in equilibrium. In addition, the decision rule must provide each receiver i with incentives to follow his private recommendation *obediently* under the assumption that every other receiver is obedient as well. The notion of equilibrium we consider is the Bayesian Nash Equilibrium. More precisely, following Bergemann and Morris (2017), we restrict our attention to decision rules that satisfy the notion of incentive compatible Bayes' Correlated Equilibrium. The idea behind this notion of incentive compatibility originates in Myerson[3]. We explain it below.

Consider a receiver i of type t_i. If every receiver reports his type honestly to the sender and

---

[2] In two seminal papers, Caroll (2016) and Mathevet et al. (2020) show that when there are multiple receivers, in the presence of private information on the part of the receivers, the revelation principle may break down when looking at *worst* equilibria

[3] See for example Myerson (1991)

obeys the recommendations of the communication rule, then the expected utility of receiver i is

$$U_i(f|t_i) = \sum_{t_j \in T_j} \sum_{a \in A} \sum_{\theta \in \Theta} \phi(\theta)\pi(t_i, t_j|\theta) f\Big((a_i, a_j)|(t_i, t_j), \theta\Big) u_i\Big((a_i, a_j), (t_i, t_j), \theta\Big).$$

Suppose now that receiver i reports $s_i \in T_i$ and for each action recommendation a_i chooses an action $\delta_i(a_i)$, with $\delta : A_i \to A_i$. If all the other receivers are honest and obedient, then the expected payoff of receiver i is

$$U_i(f, \delta_i, s_i|t_i) = \sum_{\theta \in \Theta} \sum_{t_j \in T_j} \sum_{a \in A} \phi(\theta)\pi_i(t_i, t_j|\theta) f(a|s_i, t_j, \theta) u_i\Big((a_j, \delta_i(a_i)), (t_i, t_j), \theta\Big)$$

Definition 3 *A decision rule $f : T \times \Theta \to \Delta(A)$ is incentive compatible if for every $i \in I$, for every $t_i, s_i \in T_i$, and for every $\delta_i : A_i \to A_i$,*

$$U_i(f|t_i) \geq U_i(f, \delta_i, s_i|t_i). \tag{1}$$

Following Myerson (1982), Myerson (1991) and Bergemann and Morris (2017), we say that a rule f is incentive compatible if it leads to an *Incentive-Compatible Bayes Correlated Equilibrium* (ICBCE) of (G, S). The inequality (1) is referred to in Myerson (1991) as the *generalized incentive constraint*. It ensures two conditions. First, that receiver i of type t_i has incentives to report his type truthfully. In addition, after observing and updating his beliefs on the information contained in the action recommendation, he finds it optimal to obey the recommendation. Thus there are two separate constraints embedded in (1): *truth-telling* and *obedience*. In addition, inequality (1) also accounts for "double-deviation" – where a receiver may find it optimal to misreport and following a recommendation, choose an action different from the recommended one. For purposes of exposition, below we describe in details the concepts of truth telling and obedience. Let $h_i : A_i \to A_i$ be the identity map – i.e. for all $a_i \in A_i$, $h_i(a_i) = a_i$.

Definition 4 *A decision rule $f : T \times \Theta \to \Delta(A)$ satisfies truth-telling if for each $i \in I$, and for* *each $t_i \in T_i$,*

$$U_i(f|t_i) \geq U_i(f, h_i, s_i|t_i). \tag{2}$$

Definition 5 *A decision rule $f : T \times \Theta \to \Delta(A)$ satisfies obedience if for each $i \in I$, for all $t_i \in T_i$, and for all $\delta_i : A_i \to A_i$,*

$$U_i(f|t_i) \geq U_i(f, \delta_i, t_i|t_i). \tag{3}$$

4.2.1 Sender's Problem

As mentioned previously, the sender must commit to a decision rule before she gets to observe the state and before the receivers submit their reported types. Consequently, the sender's ex-ante expected utility is

$$V_i(f) = \sum_{\theta\in\Theta} \sum_{t\in T} \sum_{a=(a_i,a_j)\in A} \psi(\theta)\pi(t|\theta)f(a|t,\theta)v(a_i, a_j). \tag{4}$$

The objective of the sender is the following:

$$\begin{aligned} \text{maximize}_f \quad & V_i(f) \\ \text{subject to} \quad & (1). \end{aligned} \tag{5}$$

Given that the sets of actions, the sets of types and the set of states of the world are all finite, and given that the incentive constraints are linear in the probabilities, for any (G, S), the set of decision rules f that are ICBCE of (G, S) constitute a convex polytope. This result is standard, but we include for completeness.

Proposition 1 *For the game (G, S), the set of decision rules f that are ICBCE of (G, S) constitute a convex polytope.*

The proof of Proposition 1 is relegated to the Appendix. In the next section we consider some examples to illustrate the problem of interest.

4.3 Examples: Applications in Binary Environments

4.3.1 No Private Information

Let us first look at the same setup as in Taneva (2019). Again, there are one sender and two receivers $i \in \{1, 2\}$, but the latter have no private information – i.e. no types are assigned to the receivers. There are two states of the world $\{\theta^0, \theta^1\} = \Theta$ and the receivers can choose among two possible actions $\{a^0, a^1\} = A$ which are the same for the both of them $(A_1 = A_2 = A)$. We assume the states θ to be uniformly distributed, as in Taneva (2019).

We define the sender's utility function $v : A_1 \times A_2 \to \mathbb{R}$ to be independent of the state of the world. It ultimately only depend on the pair of actions (a_1, a_2) selected by the receivers. Formally, the sender gets a payoff of 1 if the two receivers choose different actions, and 0 otherwise – as illustrated below.

$$v(a^1, a^2) = \begin{cases} 1 & \text{if } a^1 \neq a^2 \\ \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

We define receiver i's utility function $u_i : A_1 \times A_2 \times \Theta \to \mathbb{R}$ such that it depends on the pair of actions (a_1, a_2) and the state of the world θ. Table 4.1 depicts the payoffs of the two receivers. For now, we set $c = 2$ and $d = 1$, and make the payoffs symmetric with respect to the state of the world.

Receiver 2

	$\theta = \theta^0$	a_0	a_1
Receiver 1	a_0	(c, c)	$(d, 0)$
	a_1	$(0, d)$	$(0, 0)$

Receiver 2

	$\theta = \theta^1$	a_0	a_1
Receiver 1	a_0	$(0, 0)$	$(0, d)$
	a_1	$(d, 0)$	(c, c)

Table 4.1: Receivers' payoffs – Without private information

As mentioned previously, the communication device $f : \Theta \to \Delta(A \times A)$ maps the state of the world into a probability distribution over the combinations of actions taken by the receivers. We represent the solution to the optimization problem f^* by the array shown in Table 4.2 below. Each table corresponds to one state of the world (θ^0 or θ^1), and the rows (columns) represent the action recommendations sent to Receiver 1 (Receiver 2). All-in-all, each table depicts a probability distribution over the set of action combinations conditional on the state of the world.

$\theta = \theta^0$

	a_2^0	a_2^1
a_1^0	1/5	2/5
a_1^1	2/5	0

$\theta = \theta^1$

	a_2^0	a_2^1
a_1^0	0	2/5
a_1^1	2/5	1/5

Table 4.2: Optimal disclosure rule for the binary case with no type

Evidently, the results are identical to the those in Taneva (2019). The sender and the receivers each get an ex ante expected utility of 0.8. Note that with no communication device f, the sender's expected utility would be 0.5 while the receivers would get 0.75.

4.3.2 Payoff-Relevant Private Information

Let us now look at an example of the main problem of this chapter. Here we examine the case with two common types and two states of the world, both uniformly distributed and formally defined as $T = \{t^0, t^1\}$ and $\Theta = \{\theta^0, \theta^1\}$. There are two actions available to the receivers, such that $A = \{a^0, a^1\}$. Table 4.3 shows the payoff function for the two receivers, while Table 4.4 shows the sender's payoffs. We assume $c = 2$, $d = 1$, $l = 1$ and $m = n = 0$.

As mentioned previously, we assume a uniform distribution over the set of types for both receivers. Table 4.5 shows the joint probability distribution over types. Under these settings, the optimal communication rule is as shown in Table 4.6.

In this case, the sender's expected payoff is 0.714 while for both receivers it is 0.857. These

$\theta = \theta^0$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	(c, c)	$(d, 0)$
a_1^1	$(0, d)$	$(0, 0)$

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	$(c, 0)$	(d, d)
a_1^1	$(0, 0)$	$(0, c)$

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	$(0, c)$	$(0, 0)$
a_1^1	(d, d)	$(c, 0)$

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	$(0, 0)$	$(0, d)$
a_1^1	$(d, 0)$	(c, c)

$\theta = \theta^1$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	$(0, 0)$	$(0, d)$
a_1^1	$(d, 0)$	(c, c)

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	$(0, c)$	$(0, 0)$
a_1^1	(d, d)	$(c, 0)$

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	$(c, 0)$	(d, d)
a_1^1	$(0, 0)$	$(0, c)$

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	(c, c)	$(d, 0)$
a_1^1	$(0, d)$	$(0, 0)$

Table 4.3: Receivers' payoffs – Generic case

$\theta = \theta^0$

	a_2^0	a_2^1
a_1^0	m	l
a_1^1	l	n

$\theta = \theta^1$

	a_2^0	a_2^1
a_1^0	n	l
a_1^1	l	m

Table 4.4: Sender's payoffs – Generic case

represent an improvement over the case without communication, with expected utilities of 0.5 and 0.75 respectively.

In the next section, we detail our main results. We show how different payoff structures, for both sender and the two receivers, can affect the optimal communication rule.

Receiver 2

Receiver 1	$\theta = \theta^0$	t_2^0	t_2^1
	t_1^0	$1/4$	$1/4$
	t_1^1	$1/4$	$1/4$

Receiver 2

Receiver 1	$\theta = \theta^1$	t_2^0	t_2^1
	t_1^0	$1/4$	$1/4$
	t_1^1	$1/4$	$1/4$

Table 4.5: Conditional probability distributions over the types

$\theta = \theta^0$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	0.572	0.214
a_1^1	0.214	0

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	0	0.357
a_1^1	0.643	0

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	0	0.643
a_1^1	0.357	0

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	0	0.214
a_1^1	0.214	0.572

$\theta = \theta^1$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	0	0.214
a_1^1	0.214	0.572

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	0	0.643
a_1^1	0.357	0

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	0	0.357
a_1^1	0.643	0

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	0.572	0.214
a_1^1	0.214	0

Table 4.6: Optimal information structure for the case with uniform priors over types

4.4 Main Result

We now discuss how the parameters of the problem – namely c, d, l, m, n – affect the optimal communication rule. Throughout the different cases below, we keep the distributions over the states and types uniform, and use the same payoff functions as in our last example above. That is, the receivers' payoffs are as shown in Table 4.3, and the sender's payoffs are as shown in Table 4.4.

The layout of the communication rule is depicted in Table 4.7. Variables r, q, p and w represent the probability distribution over the set of action profiles for the two receivers when they report the same type. On the other hand, x, y, z and v play the same role but when their reports are different. As these two sets of variables represents probability distributions, it is required that $r+q+p+w=1$ and $x+y+z+v=1$. This layout is considered to be symmetric between type profiles, as both (t^0, t^0) and (t^1, t^1) share the same variables. As shown in Theorem 1, the optimal rule may not always be symmetric.

$\theta = \theta^0$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	r	q
a_1^1	p	w

$t_1 = t^0$, $t_2 = t^1$

	a_2^0	a_2^1
a_1^0	x	y
a_1^1	z	v

$t_1 = t^1$, $t_2 = t^0$

	a_2^0	a_2^1
a_1^0	x	z
a_1^1	y	v

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	w	p
a_1^1	q	r

$\theta = \theta^1$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	w	p
a_1^1	q	r

$t_1 = t^0$, $t_2 = t^1$

	a_2^0	a_2^1
a_1^0	v	z
a_1^1	y	x

$t_1 = t^1$, $t_2 = t^0$

	a_2^0	a_2^1
a_1^0	v	y
a_1^1	z	x

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	r	q
a_1^1	p	w

Table 4.7: Communication rule layout – Symmetric w.r.t. types

Next, the sender's expected utility is defined as

$$
\begin{aligned}
EV = \pi(t^0, t^0)[\psi(\theta^0) + \psi(\theta^1)][rm + (q+p)l + wn] \\
+ \ \pi(t^1, t^1)[\psi(\theta^0) + \psi(\theta^1)][rn + (q+p)l + wm] \\
+ \ \pi(t^0, t^1)[\psi(\theta^0) + \psi(\theta^1)][xm + (y+z)l + vn] \\
+ \ \pi(t^1, t^0)[\psi(\theta^0) + \psi(\theta^1)][xm + (y+z)l + vn]
\end{aligned}
$$

We refer to the obedience, truth-telling and incentive compatibility equations as they are detailed in the Appendix. Next, we formally define the condition for a communication rule to be symmetric.

Definition 6 *Let C be an array representing a communication rule, $C_{t,\theta}$ be the sub-array for the type profile $t \in T$ and state $\theta \in \Theta$, and $C^{\alpha,\beta}$ be the element at the α, β coordinate. Then C is symmetric if, given a receiver i's set of actions A_i, for every state $\theta \in \Theta$, type profile $t \in T$, and receiver $i \in I$,*

$$C_{t,\theta}^{\alpha,\beta} = C_{t,\theta'}^{(|A_i|+1-\alpha,\, |A_j|+1-\beta)}$$

$\forall i, j \in I$, $|A_k| > 1$, $k \in I$ and $|I| = 2$.

Proposition 2 below describes the conditions under which a symmetric rule is optimal. Its proof is relegated to the Appendix.

Proposition 2 *Let V be sender's payoff array such that $V_{t\theta}$ is the sub-array of type profile $t \in T$ and state $\theta \in \Theta$, and and $V^{\alpha,\beta}$ be the element at the α, β coordinate. A symmetrical communication rule is optimal for Sender if*

(i) $V_{\theta}^{\alpha,\beta} = V_{\theta'}^{(|A_i|+1-\alpha,\, |A_j|+1-\beta)}$

(ii) $V(\theta) = V^{\top}(\theta)$

(iii) $\mathbb{E}V((a_i, a_j)) = \mathbb{E}V((a_j, a_i)), \quad \forall a_k \in A_k,\ k \in I$

$\forall i, j \in I$, $|A_k| > 1$, $k \in I$ and $|I| = 2$.

We now present the main results of this chapter. For ease of presentation, we split it into two parts: Proposition 3 encapsulates the fundamental findings of Theorem 1. Theorem 1 elaborates on every optimal communication rule allowed by the different parametric conditions of our model. Proposition 3 categorizes some of them into three types of information disclosure. We say that a communication rule is fully revealing if every receiver i can correctly deduce the state from his private recommendation. Next, we say that it induces no revelation if for each receiver i, the recommendation does not lead to any updating of his private belief regarding the state of the world. Lastly, we refer to partial revelation to describe situations that lie in between the aforementioned extreme cases. Proposition 3 is based on Theorem 1. The proof of the latter is relegated to the Appendix.

Proposition 3 *Consider different values of l, m and n for the sender's payoff, and c, d for the receivers' payoffs.*

(i) The optimal communication device fully reveals the state if:

- $m > l > n$ *and* $c = 2d$
- $m > l > n$, $m + n > 2l$, $m > \frac{l(c+d)-n(2d-c)}{2c-d}$, $c < d < 2c$

(ii) The optimal communication device reveals no information about the state if:

- $l > m \lesseqgtr n$ *and* $c < d$
- $l < m = n$ *and* $c > d$ *for* $r = \frac{1}{2}$

(iii) The optimal communication device partially reveals the state if, for example:

- $l > m \lesseqgtr n$ *and* $c > d$ *(Symmetric w.r.t. types)*
- $m \lesseqgtr n > l$, $c > 2d$ *(Non-symmetric w.r.t. types)*

Theorem 1 *For each of the following cases regarding the relationship between l, m and n and between c and d, the optimal communication rule is given below.*

1. *If* $l > m \geq n$ *or* $l > n \geq m$ *&* $c > d$*, then* $r = 1 - 2q$, $q = p = \frac{c+d}{2(3c+d)}$, $x = v = 0$, $y = \frac{3c-d}{2(3c+d)}$, $z = 1 - y$

2. *If* $l > m > n$ *or* $l > n > m$ *&* $c < d$*, then* $r = 0$, $q = p = \frac{1}{2}$, $x = v = 0$, $y = z = \frac{1}{2}$

3. *If* $l < m = n$ *&* $c > d$*, then* $r \in [\frac{1}{2}, 1]$, $p = q = 0$, $x = v = \frac{1}{2}$, $y = z = 0$

4. *If* $l < m = n$ *&* $c < d$*, then* $r = \frac{3d-c}{2(c+3d)}$, $q = p = 0$, $x = v = \frac{c+d}{2(c+3d)}$, $z = 0$

5. *If* $m > n > l$ *or* $n > m > l$ *&* $c > d$*, then*
   (a) *for* $c \leq 2d$, $r = 1$, $x = \frac{2c-d}{c+d}$, $y = z = 0$
   (b) *for* $c > 2d$*, the communication rule is no more symmetric with respect to the type profiles* (t^0, t^0) *and* (t^1, t^1)*. In that case, we define* r, q, p *and* $\hat{r}, \hat{q}, \hat{p}$ *associated with these two type profiles respectively. At optimum,* $r = 1$, $\hat{r} = \frac{2d}{c}$, $x = 1$

6. *If* $m > n > l$ *or* $n > m > l$ *&* $c < d$*, then*
   (a) *for* $\frac{d(c+d)}{(c+3d)(d-c)} \geq \frac{m-n}{m+n-2l}$*, at optimum* $r = \hat{r} = \frac{3d-c}{2(c+3d)}$, $q = \hat{q} = p = \hat{p} = 0$, $x = v = \frac{c+d}{2(c+3d)}$, $z = 0$
   (b) *otherwise,* $r = \frac{d-c}{d}$, $\hat{r} = 0$, $q = \hat{q} = p = \hat{p} = 0$, $y = 1$

7. *If* $m > l > n$ *&* $c > d$*, then*
   (a) *for* $c = 2d$*, at optimum* $r = x = 1$
   (b) *for* $c < 2d$*, at optimum* $r = 1$, $x = \frac{2c-d}{c+d}$, $y = z = 0$
   (c) *for* $c \leq 2d$ *and* $\frac{4(2c+d)(c+d)}{13c^2-d^2} > \frac{m-n}{l-n}$*, at optimum* $r = 1 - 2q$, $q = \frac{c+d}{2(3c+d)}$, $x = v = 0$, $y = \frac{3c-d}{2(3c+d)}$
   (d) *for* $d(2 + \sqrt{3}) > c > 2d$*, at optimum* $r = 1$, $\hat{r} = \frac{d(2c-d)}{c^2}$, $\hat{q} = \frac{(c-d)^2}{2c^2}$, $x = \frac{(c-d)(3c-d)}{2c^2}$, $z = 1 - x$
   (e) *for* $c > 2d$*, at optimum* $r = x = 1$, $q = \hat{q} = 0$, $\hat{r} = \frac{2d}{c}$

    (f) *for* $c > 2d$ *and* $\frac{d(2d(c-d)^2-2c(c-2d))}{c(d(c-d)^2-c(c-2d))} > \frac{4(m-l)}{m-n}$, *at optimum* $r = x = 1$, $\hat{r} = \frac{2(c-2d)}{(c-d)^2}$, $\hat{q} = 2 - \frac{\hat{r}c}{d}$

8. *If* $m > l > n$ *&* $c < d$*, then*

    (a) *for* $d < 2c$, $m + n \geq 2l$ *and* $m > \frac{l(c+d)-n(2d-c)}{2c-d}$, *at optimum* $r = 1$ *and* $y = 1$

    (b) *for* $d < 2c$, $m + n \leq 2l$ *and* $m > \frac{l(c+d)-n(2d-c)}{2c-d}$, *at optimum* $q = \frac{1}{2}$ *and* $x = \frac{3(d-c)}{2(2d-c)}$, $y = 1 - x$

    (c) *for* $d \geq 2c$ *or* $m+n \geq 2l$ *and* $m < \frac{l(c+d)-n(2d-c)}{2c-d}$, *at optimum* $r = p = q = \frac{1}{3}$, $x = \frac{d-c}{2d-c}$ *and* $y = 1 - x$

    (d) *for* $m \leq m^*_{(b)}$, *or for* $m \leq m^*_{(c)}$, *at optimum* $r = \frac{(c-d)(y-1)}{c+d} + x$, $q = p = 0$, $x = v = \frac{y(2d^2-3cd-c^2)+3cd-d^2}{2d(c+d)}$, $y = \frac{2d(2cd-c^2-d^2)}{5cd^2-cd-c^3-3d^3}$, $z = 0$. *These grant the sender a higher expected utility compared to the above cases (b) and (c) respectively, where* $m^*_{(b)}$ *and* $m^*_{(c)}$ *are thresholds for* m *for (b) and (c) presented in the Appendix.*

Theorem 1 shows that the extent to which the state is revealed depends on the level of alignment between the payoffs of the sender and the receivers. Observe that for the type t^0, the payoff c is associated with his preference for "matching the state" – that is, playing a^k when $\theta = \theta^k$ – when the other do so as well, while d is obtained by this type when he is the only receiver matching it. At the opposite, t^1 gets c for not matching the state. Moreover, t^0 receives c when both receivers choose $a^k \neq \theta^k$, and d for being the only one. As for the sender, she gets m when both receivers are matching the state, n when none does so, and l when they play different actions. Regarding the communication rule, r and w are associated with receivers selecting the same action when they report identical types, and x and v when their reports are different. Similarly, q and p are the probabilities of recommending $a_1 \neq a_2$ when $t_1 = t_2$, and y and z when $t_1 \neq t_2$. Refer to Tables 4.3, 4.4 and 4.7 for a complete presentation of these variables.

Now recall that both the type profiles and the states are assumed to be uniformly distributed. As shown in Proposition 3, observe that full revelation occurs when the sender has a clear preference for a specific action profile, while receivers' preferences are either symmetrically opposed between types ($c = 2d$ in (i)) or moderately misaligned with the sender's ($c < d < 2c$ in (i)). At the opposite, the optimal communication rule reveals no information when both parties prefer miscoordination ($c < d$ in (ii)). Observe also that it might be that while coordination is preferred by every player ($l < m = n$ and $c > d$ as in (ii)), no revelation can be optimal (if $r = \frac{1}{2}$) because of the indifference of the sender between the two matching action profiles (a^0, a^0) and (a^1, a^1). In most cases, however, partial revelation yields the highest expected utility for the sender. Indeed, in the previous case with preferences towards coordination, if $r = 1$ then the state is perfectly revealed if the reported types are identical, while no information is transmitted otherwise. This is also the case if preferences are conflicting. Moreover, even under some degree of alignment, if receivers strongly favor some action profiles ($m \gtreqless n > l$ and $c > 2d$), this can lead to non-symmetric communication rules with respect to type profiles where (t^0, t^0) and (t^1, t^1) do not exhibit a symmetric distribution over the set of action profiles in a given state. Finally, notice that full and no revelation are rarely optimal in our model. Because we have modelled the receivers' preferences to be in opposition with regards to their types, it is to be expected that cases of full revelation almost never occur. Indeed, if we were to change both parties' payoffs such that their preferences are (i) aligned and (ii) differentiable between states, then full revelation would be optimal. Moreover, as stated in Proposition 3, no revelation is preferred when preferences are aligned and ex ante indifference between action profiles is optimal for every agent. For all other cases, the chosen communication device reveals partial information about the state. We further distinguish between two types of partial revelation: Symmetric and non-symmetric communication rules. Given a state θ, the former exhibits symmetrical distributions over action profiles between the types profiles (t^0, t^0) and (t^1, t^1) while the latter does not. Because the two receivers are, in our settings, identical in terms of preferences given a type t^k and $k \in \{0, 1\}$, the rule will always be symmetrical between (t^0, t^1) and (t^1, t^0). Indeed, it is easy to see that changing the positions of the receivers – i.e. making player 2 to be player 1 and vice-versa – does not affect our results. However, there are cases where it is beneficial for the sender to submit different recommendations to the profiles (t^0, t^0) and (t^1, t^1). For example, when the two sides have strong, opposite preferences towards specific action profiles.

4.5 Conclusion

We have extended the information design problem in Taneva (2019) by incorporating private information held by the receivers. Using a revelation principle argument, we have focused on action recommendations and have identified the communication rules that satisfy *incentive compatibility*[4]. In a binary environment, our main results (Proposition 3 and Theorem 1) characterize the optimal communication rule depending on the payoff structures of a designer and two privately-informed receivers. We show that while in some cases full disclosure of the true state can be optimal, in (most) other cases upon privately observing their recommendation, receivers cannot deduce perfectly the state of the world. Moreover, while the designer does not observe the receivers' types but only their reports, it is sometimes beneficial for her to construct a communication rule that is not symmetrical across type profiles. That is, even if the payoff structures between two type profiles are symmetric, tailoring the recommendation rule differently to each of them can result in a higher expected payoff for the designer.

---

[4] See Bergemann and Morris (2017)

Appendix A

Simplified Constraints

Obedience

Player 1

- $(a^0, t^0) \to (a^1, t^0)$

$$\begin{aligned}&\pi(t^0,t^0)\psi(\theta^0)(rc+qd)+\pi(t^0,t^1)\psi(\theta^0)(xc+yd)\\ \geq\ &(\pi(t^0,t^0)\psi(\theta^1)(pc+wd)+\pi(t^0,t^1)\psi(\theta^1)(zc+vd))\end{aligned}$$

- $(a^1, t^0) \to (a^0, t^0)$

$$\begin{aligned}&\pi(t^0,t^0)\psi(\theta^1)(rc+qd)+\pi(t^0,t^1)\psi(\theta^1)(xc+yd)\\ \geq\ &(\pi(t^0,t^0)\psi(\theta^0)(pc+wd)+\pi(t^0,t^1)\psi(\theta^0)(zc+vd))\end{aligned}$$

- $(a^0, t^1) \to (a^1, t^1)$

$$\begin{aligned}&\pi(t^1,t^1)\psi(\theta^1)(rc+qd)+\pi(t^1,t^0)\psi(\theta^1)(vc+yd)\\ \geq\ &(\pi(t^1,t^1)\psi(\theta^0)(pc+wd)+\pi(t^1,t^0)\psi(\theta^0)(zc+xd))\end{aligned}$$

- $(a^1, t^1) \to (a^0, t^1)$

$$\begin{aligned}&\pi(t^1,t^1)\psi(\theta^0)(rc+qd)+\pi(t^1,t^0)\psi(\theta^0)(vc+yd)\\ \geq\ &(\pi(t^1,t^1)\psi(\theta^1)(pc+wd)+\pi(t^1,t^0)\psi(\theta^1)(zc+xd))\end{aligned}$$

Player 2

- $(a^0, t^0) \rightarrow (a^1, t^0)$

$$\begin{aligned} &\pi(t^0, t^0)\psi(\theta^0)(rc + pd) + \pi(t^1, t^0)\psi(\theta^0)(xc + yd) \\ \geq &\ (\pi(t^0, t^0)\psi(\theta^1)(qc + wd) + \pi(t^1, t^0)\psi(\theta^1)(zc + vd)) \end{aligned}$$

- $(a^1, t^0) \rightarrow (a^0, t^0)$

$$\begin{aligned} &\pi(t^0, t^0)\psi(\theta^1)(rc + pd) + \pi(t^1, t^0)\psi(\theta^1)(xc + yd) \\ \geq &\ (\pi(t^0, t^0)\psi(\theta^0)(qc + wd) + \pi(t^1, t^0)\psi(\theta^0)(zc + vd)) \end{aligned}$$

- $(a^0, t^1) \rightarrow (a^1, t^1)$

$$\begin{aligned} &\pi(t^1, t^1)\psi(\theta^1)(rc + pd) + \pi(t^0, t^1)\psi(\theta^1)(vc + yd) \\ \geq &\ (\pi(t^1, t^1)\psi(\theta^0)(qc + wd) + \pi(t^0, t^1)\psi(\theta^0)(zc + xd)) \end{aligned}$$

- $(a^1, t^1) \rightarrow (a^0, t^1)$

$$\begin{aligned} &\pi(t^1, t^1)\psi(\theta^0)(rc + pd) + \pi(t^0, t^1)\psi(\theta^0)(vc + yd) \\ \geq &\ (\pi(t^1, t^1)\psi(\theta^1)(qc + wd) + \pi(t^0, t^1)\psi(\theta^1)(zc + xd)) \end{aligned}$$

Truth-telling

Player 1

- $t^0 \to t^1$

$$\begin{aligned}
&\pi(t^0,t^0)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^0,t^1)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd))\\
\geq\ &(\pi(t^0,t^0)(\psi(\theta^0)(xc+zd)+\psi(\theta^1)(xc+zd))+\pi(t^0,t^1)(\psi(\theta^0)(wc+pd)+\psi(\theta^1)(wc+pd)))
\end{aligned}$$

- $t^1 \to t^0$

$$\begin{aligned}
&\pi(t^1,t^1)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^1,t^0)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))\\
\geq\ &(\pi(t^1,t^1)(\psi(\theta^0)(vc+zd)+\psi(\theta^1)(vc+zd))+\pi(t^1,t^0)(\psi(\theta^0)(wc+pd)+\psi(\theta^1)(wc+pd)))
\end{aligned}$$

Player 2

- $t^0 \to t^1$

$$\begin{aligned}
&\pi(t^0,t^0)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^1,t^0)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd))\\
\geq\ &(\pi(t^0,t^0)(\psi(\theta^0)(xc+zd)+\psi(\theta^1)(xc+zd))+\pi(t^1,t^0)(\psi(\theta^0)(wc+qd)+\psi(\theta^1)(wc+qd)))
\end{aligned}$$

- $t^1 \to t^0$

$$\begin{aligned}
&\pi(t^1,t^1)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^0,t^1)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))\\
\geq\ &(\pi(t^1,t^1)(\psi(\theta^0)(vc+zd)+\psi(\theta^1)(vc+zd))+\pi(t^0,t^1)(\psi(\theta^0)(wc+qd)+\psi(\theta^1)(wc+qd)))
\end{aligned}$$

Incentive compatibility

Example: Player 1

- $\delta(a^0) = \delta(a^1) = a^0 \,; t^0 \to t^1$

$$\pi(t^0,t^0)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^0,t^1)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd))$$
$$\geq (\pi(t^0,t^0)\psi(\theta^0)(xc+zd+yc+vd)+\pi(t^0,t^1)\psi(\theta^0)(wc+pd+qc+rd))$$

- $\delta(a^0) = \delta(a^1) = a^0 \,; t^1 \to t^0$

$$\pi(t^1,t^1)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^1,t^0)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))$$
$$\geq (\pi(t^1,t^1)\psi(\theta^1)(vc+zd+yc+xd)+\pi(t^1,t^0)\psi(\theta^1)(wc+pd+qc+rd))$$

- $\delta(a^0) = \delta(a^1) = a^1 \,; t^0 \to t^1$

$$\pi(t^0,t^0)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^0,t^1)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd))$$
$$\geq (\pi(t^0,t^0)\psi(\theta^1)(xc+zd+yc+vd)+\pi(t^0,t^1)\psi(\theta^1)(wc+pd+qc+rd))$$

- $\delta(a^0) = \delta(a^1) = a^1 \,; t^1 \to t^0$

$$\pi(t^1,t^1)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^1,t^0)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))$$
$$\geq (\pi(t^1,t^1)\psi(\theta^0)(vc+zd+yc+xd)+\pi(t^1,t^0)\psi(\theta^0)(wc+pd+qc+rd))$$

- $\delta(a^0) = a^1$ & $\delta(a^1) = a^0$; $t^0 \to t^1$

$$\begin{aligned} &\pi(t^0,t^0)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^0,t^1)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd)) \\ &\geq (\pi(t^0,t^0)(\psi(\theta^0)(yc+vd)+\psi(\theta^1)(yc+vd))+\pi(t^0,t^1)(\psi(\theta^0)(qc+rd)+\psi(\theta^1)(qc+rd))) \end{aligned}$$

- $\delta(a^0) = a^1$ & $\delta(a^1) = a^0$; $t^1 \to t^0$

$$\begin{aligned} &\pi(t^1,t^1)(\psi(\theta^0)(rc+qd)+\psi(\theta^1)(rc+qd))+\pi(t^1,t^0)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd)) \\ &\geq (\pi(t^1,t^1)(\psi(\theta^0)(yc+xd)+\psi(\theta^1)(yc+xd))+\pi(t^1,t^0)(\psi(\theta^0)(qc+rd)+\psi(\theta^1)(qc+rd))) \end{aligned}$$

Example: Player 2

- $\delta(a^0) = \delta(a^1) = a^0$; $t^0 \to t^1$

$$\begin{aligned} &\pi(t^0,t^0)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^1,t^0)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd)) \\ &\geq (\pi(t^0,t^0)\psi(\theta^0)(xc+zd+yc+vd)+\pi(t^1,t^0)\psi(\theta^0)(wc+qd+pc+rd)) \end{aligned}$$

- $\delta(a^0) = \delta(a^1) = a^0$; $t^1 \to t^0$

$$\begin{aligned} &\pi(t^1,t^1)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^0,t^1)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd)) \\ &\geq (\pi(t^1,t^1)\psi(\theta^1)(vc+zd+yc+xd)+\pi(t^0,t^1)\psi(\theta^1)(wc+qd+pc+rd)) \end{aligned}$$

- $\delta(a^0) = \delta(a^1) = a^1$; $t^0 \to t^1$

$$\begin{aligned} &\pi(t^0,t^0)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^1,t^0)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd)) \\ &\geq (\pi(t^0,t^0)\psi(\theta^1)(xc+zd+yc+vd)+\pi(t^1,t^0)\psi(\theta^1)(wc+qd+pc+rd)) \end{aligned}$$

- $\delta(a^0) = \delta(a^1) = a^1$; $t^1 \to t^0$

$$\pi(t^1,t^1)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^0,t^1)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))$$
$$\geq (\pi(t^1,t^1)\psi(\theta^0)(vc+zd+yc+xd)+\pi(t^0,t^1)\psi(\theta^0)(wc+qd+pc+rd))$$

- $\delta(a^0) = a^1$ & $\delta(a^1) = a^0$; $t^0 \to t^1$

$$\pi(t^0,t^0)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^1,t^0)(\psi(\theta^0)(xc+yd)+\psi(\theta^1)(xc+yd))$$
$$\geq (\pi(t^0,t^0)(\psi(\theta^0)(yc+vd)+\psi(\theta^1)(yc+vd))+\pi(t^1,t^0)(\psi(\theta^0)(pc+rd)+\psi(\theta^1)(pc+rd)))$$

- $\delta(a^0) = a^1$ & $\delta(a^1) = a^0$; $t^1 \to t^0$

$$\pi(t^1,t^1)(\psi(\theta^0)(rc+pd)+\psi(\theta^1)(rc+pd))+\pi(t^0,t^1)(\psi(\theta^0)(vc+yd)+\psi(\theta^1)(vc+yd))$$
$$\geq (\pi(t^1,t^1)(\psi(\theta^0)(yc+xd)+\psi(\theta^1)(yc+xd))+\pi(t^0,t^1)(\psi(\theta^0)(pc+rd)+\psi(\theta^1)(pc+rd)))$$

Appendix B

Proofs

Proof of Proposition 1

Let f^1 and f^2 be feasible solutions yielding $V_i(f^1) = V_i(f^2) = v^*$ for the designer. Define $\bar{f} = \alpha f^1 + (1-\alpha)f^2$ with $\alpha \in (0,1)$. The goal is to show that $\bar{f}$ cannot be a solution if the set of feasible solutions is convex. Suppose the contrary. By definition, f^1, f^2 and $\bar{f}$ are all incentive compatible. Now, regarding $\bar{f}$ this means that

$$U_i(\bar{f}, h_i|t_i, t_i) \geq U_i(f, \delta_i|t_i, \hat{t_i})$$

given the identity map $h_i : A_i \to A_i$, i's true and reported types t_i and s_i. The above must also hold for

$$\begin{aligned} & U_i(\bar{f}, h_i|t_i, t_i) \geq U_i(f^1, \delta_i|t_i, s_i) \\ \Longrightarrow\ & \alpha U_i(f^1, h_i|t_i, t_i) + (1-\alpha)U_i(f^2, h_i|t_i, t_i) \geq U_i(f^1, \delta_i|t_i, s_i) \\ \Longrightarrow\ & U_i(f^2, h_i|t_i, t_i) \geq U_i(f^1, h_i|t_i, t_i) \end{aligned}$$

Then by definition the last inequality must hold with equality, proving the claim.

■

Proof of Proposition 2

Let $\mathbb{E}V$ and $\widetilde{\mathbb{E}V}$ be sender's expected payoffs under the symmetric and non-symmetric communication rules C and $\widetilde{C}$ respectively. By definition,

$$\mathbb{E}V = \sum_{a\in A,\, t\in T,\, \theta\in\Theta} \psi(\theta)\pi(t|\theta)C(a|t,\theta)V(a,\theta)$$

C is optimal if $\mathbb{E}V \geq \widetilde{\mathbb{E}V}$. Suppose that, on the contrary, C is not optimal given assumptions (i) and (ii). This implies that

$$\sum_{a\in A,\, t\in T,\, \theta\in\Theta} \psi(\theta)\pi(t|\theta)\widetilde{C}(a|t,\theta)V(a,\theta) \geq \sum_{a\in A,\, t\in T,\, \theta\in\Theta} \psi(\theta)\pi(t|\theta)C(a|t,\theta)V(a,\theta) \tag{1}$$

Without loss of generality, suppose that $\widetilde{C}$ is identical to C but for a specific type profile $\widetilde{t} \in T$, for which two action profiles are getting different recommendation probabilities than under C. Then this can only hold if V is not symmetric and therefore induces a contradiction.

■

Proof of Theorem 1

We will focus on the constraints that are binding.

Case 1: $l > m = n$ & $c > d$

Given that $m = n$, the optimal rule is symmetric, thus $p = q$ and $v = x$. This in turn implies that $w = 1 - 2q - r$ and $z = 1 - 2x - y$. The sender's expected utility can therefore be written as

$$EV = \frac{1}{2}(x - q)(m + n - 2l)$$

and thus the sender wants to maximize q and minimize x. From obedience,

$$r \geq \frac{q(c - 3d) - x(3c - d)}{c + d} - y + 1$$

Because of the symmetry in the communication rule, every case of obedience resolves into the above inequality. The same applies for truth-telling and incentive compatibility. Truth-telling is given by

$$r \geq \frac{c + d - 2d(x + y)}{2c} - q$$

As for incentive compatibility, the inequality becomes

$$r \geq q + y - x$$

Non-negativity implies that

$$w \geq 0 \implies 1 - 2q - r \geq 0 \tag{C1}$$

$$z \geq 0 \implies 1 - 2x - y \geq 0 \tag{C2}$$

Consider $x = 0$ and (C1), then obedience implies that

$$y \geq q\left(\frac{3c - d}{c + d}\right)$$

and incentive compatibility requires

$$y \leq 1 - 3q$$

A graphical representation is given below in Figure 4.1. The black circle represents the optimal values for q and y.

Case 2: $l > m = n$ & $c < d$

Observe that only the incentive compatibility constraint changes to now be

$$r \leq q + y - x$$

since $c - d < 0$ by definition. This effectively removes the upper bound on q in the previous case, making (C1) bind. Moreover, obedience stays the same as in Case 1. A graphical representation is given below in Figure 4.2.

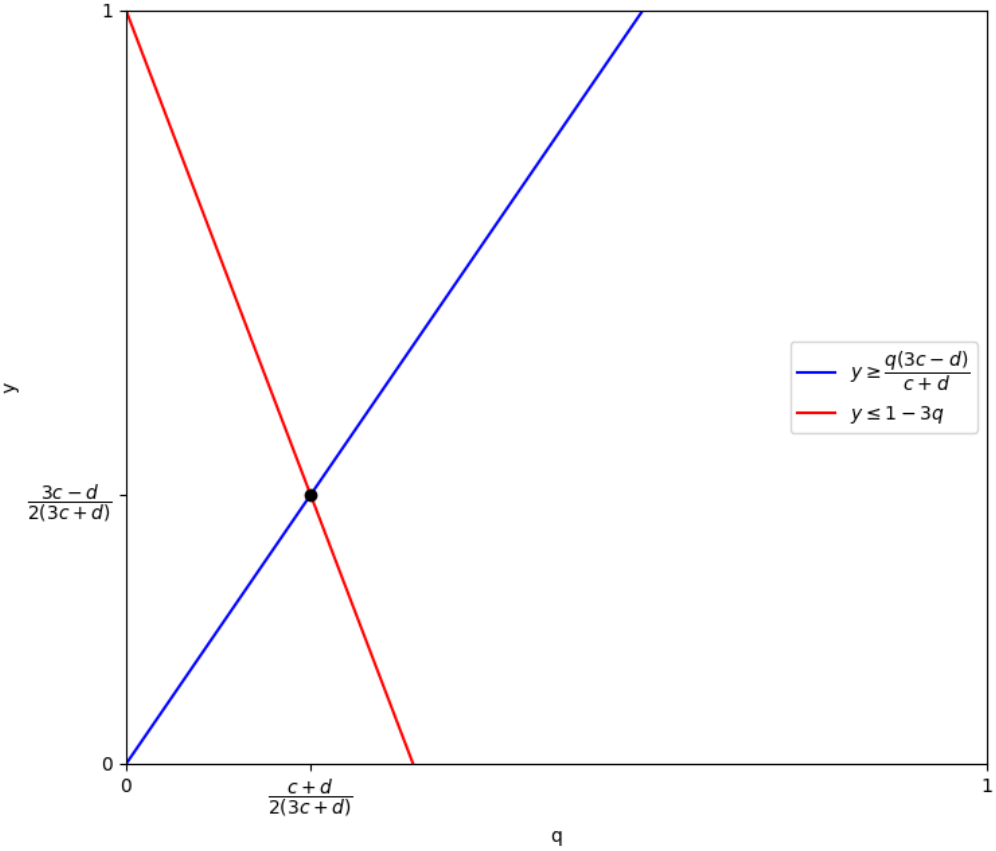

Figure 4.1: Optimal values for q and y in Case 1

Case 3: $l < m = n$ & $c > d$

Here, while the actual form of sender's expected utility does not change, the objective does. Indeed, with $l < m = n$, sender now wants to maximize x instead. Consider $q = y = 0$, then the truth-telling constraint is

$$r \geq \frac{c+d-2xd}{2c}$$

Below is the graphical representation (Figure 4.3). Here, the optimal values of r lie on the highlighted line.

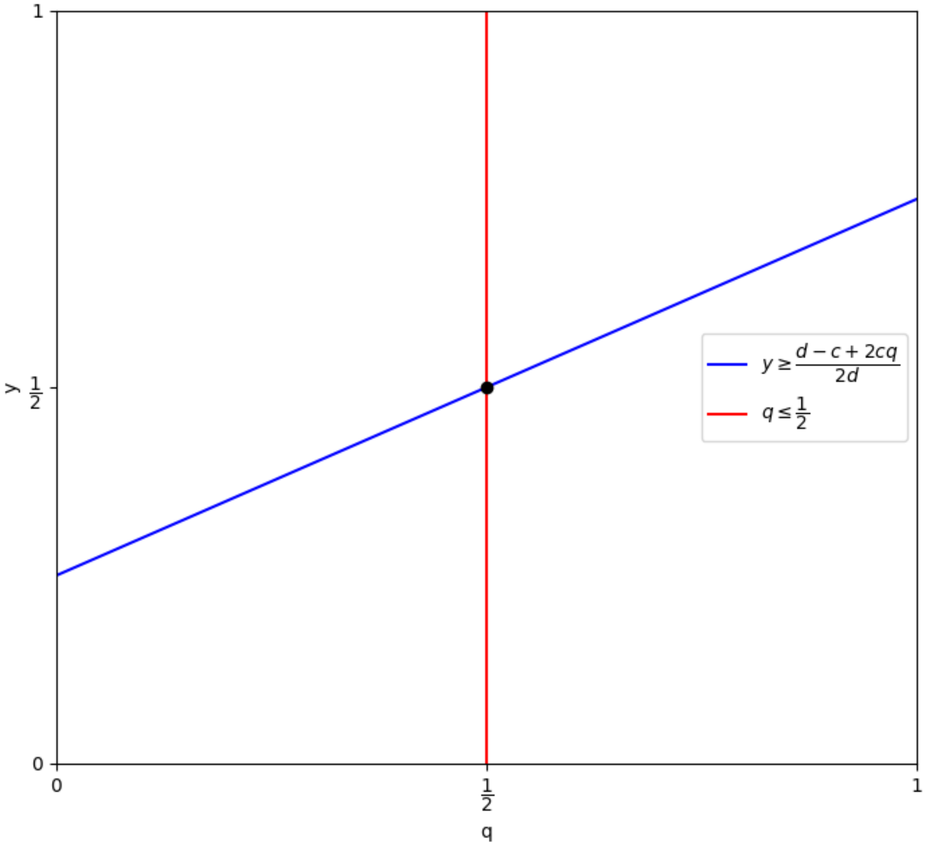

Figure 4.2: Optimal values for q and y in Case 2

Case 4: $l < m = n$ & $c < d$

As before, the incentive compatibility constraint

$$r \leq q + y - x$$

imposes an upper bound on r, and therefore binds. Using this in the obedience constraint, along with $q = 0$, we get

$$r \geq \frac{1}{2} - \frac{2xc}{c+d}$$

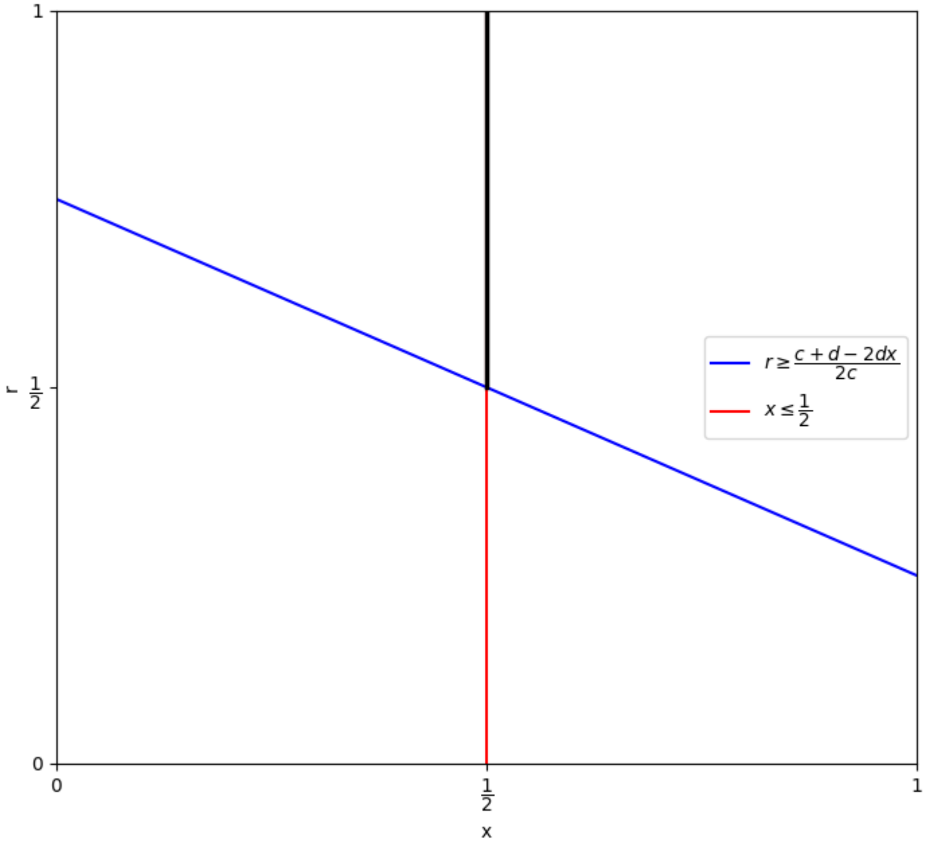

Figure 4.3: Optimal values for r and x in Case 3

From maximizing x, constraint (C2) binds, thus

$$r \leq 1 - 3x$$

Figure 4.4 shows the feasible values for r and x, and the optimum.

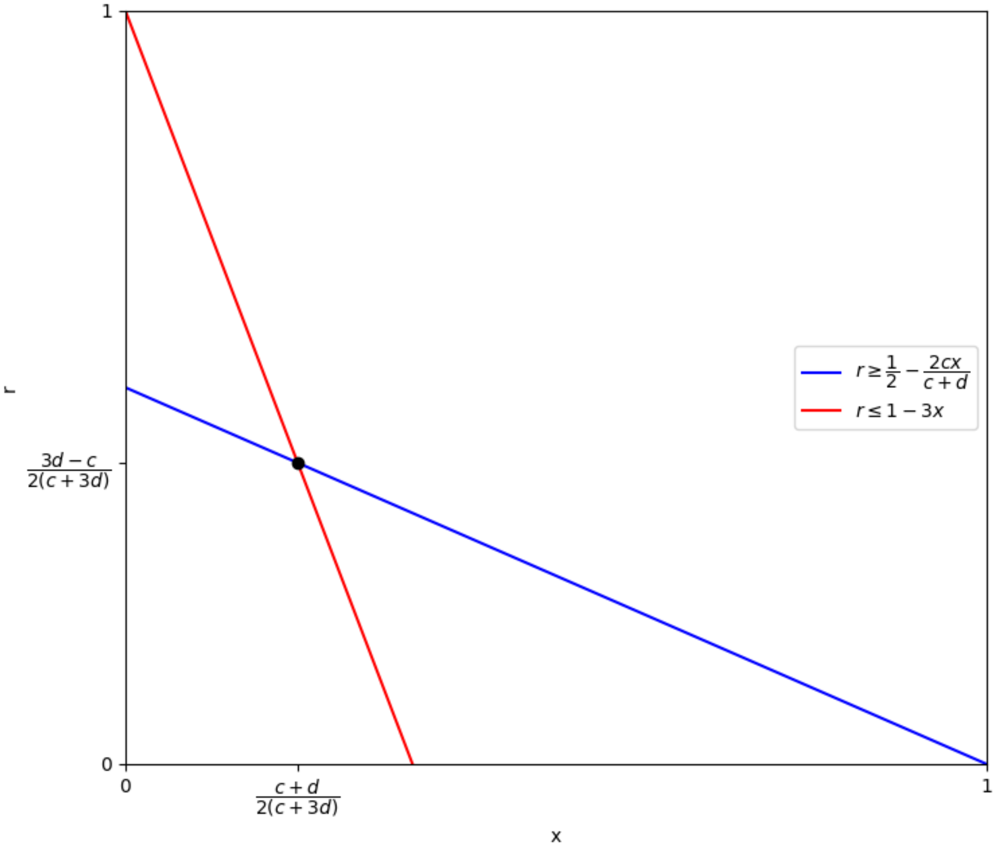

Figure 4.4: Optimal values for x and y in Case 4

Case 5: $l > m > n$ & $c > d$

Here, observe that since $l > m > n$, the goal for the sender stays the same – that is, to maximize q and minimize x. As such, sender's expected utility is defined as

$$EV = \frac{1}{2}\{x(m-n) + (y+z)(l-n) - q(m+n-2l)\}$$

Therefore, the optimal rule is effectively the same as in Case 1.

Case 6: $l > m > n$ & $c < d$

Similarly to Case 5, the results are the same as in Case 2.

Case 5' & Case 6': $l > n > m$

Changing the ordering of m and n in sender's preferences in these settings does not affect the results. The intuition behind this is that variables other than q and x only play a parametric role in the solution. They serve to meet the restrictions imposed by the constraints.

Case 7: $m > n > l$ **&** $c > d$

From this section and onward, the environment allows, at times, for both symmetric and non-symmetric solutions. At this point, let us rewrite the relevant constraint used throughout the remaining cases. There will now be two sets of constraints: one that keeps the symmetry between the two types of the receivers, and another one that does not. For the former, the layout of the communication rule does not change, and therefore – assuming that $c > d$ for now – the constraints are

Obedience for type t^0

$$r \geq \frac{q(c-3d) - 2yd + z(c-d) + 2d}{c+d} - x \tag{Ob – t^0}$$

Obedience for type t^1

$$r \geq \frac{q(c-3d) - y(c-d) + 2zc - c + d}{c+d} + x \tag{Ob – t^1}$$

Truth-telling for type t^0 *&* t^1

$$r \geq \frac{1}{2} - q - \frac{d(y-z)}{2c} \tag{TT}$$

Incentive compatibility for type t^0

$$r \geq q - \frac{x(c+d) - y(c-2d) + zd - d}{c-d} \tag{IC – t^0}$$

Incentive compatibility for type t^1

$$r \geq q + \frac{x(c+d) + y(2c-d) + zc - c}{c-d} \tag{IC – t^1}$$

Observe that if $c < d$, then only the incentive compatibility constraints related to r being smaller or equal to their respective right-hand sides bind. For the second scenario, in which the symmetry between the two types is broken, then the layout of the communication rule becomes

$\theta = \theta^0$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	r	q
a_1^1	p	w

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	x	y
a_1^1	z	v

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	x	z
a_1^1	y	v

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	$\hat{w}$	$\hat{p}$
a_1^1	$\hat{q}$	$\hat{r}$

$\theta = \theta^1$

$t_1 = t_2 = t^0$

	a_2^0	a_2^1
a_1^0	w	p
a_1^1	q	r

$t_1 = t^0, t_2 = t^1$

	a_2^0	a_2^1
a_1^0	v	z
a_1^1	y	x

$t_1 = t^1, t_2 = t^0$

	a_2^0	a_2^1
a_1^0	v	y
a_1^1	z	x

$t_1 = t_2 = t^1$

	a_2^0	a_2^1
a_1^0	$\hat{r}$	$\hat{q}$
a_1^1	$\hat{p}$	$\hat{w}$

Table 4.8: Communication rule layout – Non-symmetric w.r.t. types

Under this new layout, we can rewrite the relevant constraints as

Obedience for type t^1

$$\hat{r} \geq \frac{q(c-3d) - y(c-d) + 2zc - c + d}{c+d} + x \qquad \text{(Ob' – } t^1\text{)}$$

Incentive compatibility for type t^1

$$\hat{r}c - rd \geq qc - \hat{q}d + 2yc - yd + x(c+d) + zc - c \qquad \text{(IC' – } t^1\text{)}$$

Again, these are the constraint that will be required to solve the remaining cases. Regarding the current case, sender's objective function is

$$EV = \frac{1}{2}\{x(m-n) - (y+z)(n-l) - q(m+n-2l)\}$$

Thus, the priority is to maximize x, and to minimize q, y and z. With this in mind. let us begin by considering $q = y = z = 0$. The incentive constraints for t^1 is now

$$r \geq \frac{x(c+d) - c}{c-d}$$

Consider $x = 1$, this means that

$$r \geq \frac{d}{d-c}$$

which is smaller than 1 if $c \geq 2d$. Therefore, if $d < c \leq 2d$, $x \leq 1$, and therefore if $r = 1$ it must be that

$$x \leq \frac{2c-d}{c+d}$$

Observe that under these assumptions, sender's expected utility is

$$EV = \frac{1}{2}x(m-n)$$

Let us now take a look at what happens if we drop the symmetry between types. Then, sender's expected utility becomes (dropping the constant terms as before to ease reading)

$$EV = \frac{1}{4}\{(r-\hat{r})(m-n)+2q(l-n)+2\hat{q}(l-m)\}+\frac{1}{2}\{x(m-n)+(y+z)(l-n)\}$$

Assuming $q=\hat{q}=y=z=0$, we obtain

$$\hat{EV} = \frac{1}{4}\{(r-\hat{r})(m-n)\}+\frac{1}{2}\{x(m-n)\}$$

and $\hat{EV} > EV$ if $r > \hat{r}$. Thus, it might be beneficial for sender to go that road, as long as this condition is met. From $(IC' - t^1)$, we get that

$$\hat{r} \geq \frac{d(r+x)}{c}$$

It follows that this will hold with equality, since EV decreases in $\hat{r}$. Also, $\hat{r} = \frac{2d}{c}$ if $r = x = 1$, and as such $\hat{r} \leq 1$ if $c \geq 2d$. This means that if $c \geq 2d$, it is indeed beneficial for the sender to drop the symmetry between types. This can be seen graphically in Figure 4.5 if $c \leq 2d$,

On the other hand, if $c \geq 2d$, the graph and optimal values are as shown in Figure 4.6.

Case 7’: $n > m > l$ & $c > d$

Similar to Case 7.

Case 8: $m > n > l$ & $c < d$

Let us first verify the type-symmetric communication rule. Observe that here, the receivers prefer their actions to not be the same. Thus, as opposed to the previous case where type t^0 preferred matching actions, both types have their ideal outcomes as mismatching actions. One can compare the obedience constraints of both types to realize that because of the uniform distributions of types and states, and due to the symmetric payoffs, this requires $x = v$ if $r = \hat{r}$ and $q = \hat{q}$. Under these

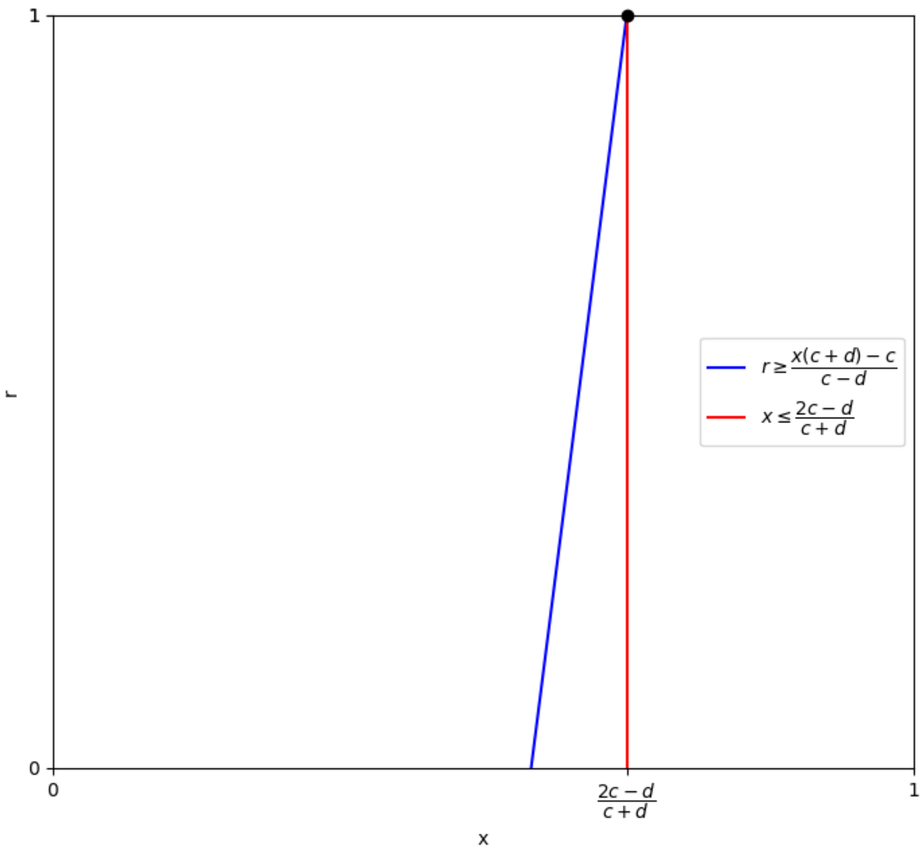

Figure 4.5: Optimal values for r and x in Case 7 if $c \leq 2d$

settings, sender's expected utility falls back to

$$EV = \frac{1}{2}(x-q)(m+n-2l)$$

Let $q = 0$ and (C2) hold with equality, then obedience requires that

$$r \geq \frac{-x(3c-d)}{c+d} - y - 1$$

and incentive compatibility imposes that

$$r \leq y - x$$

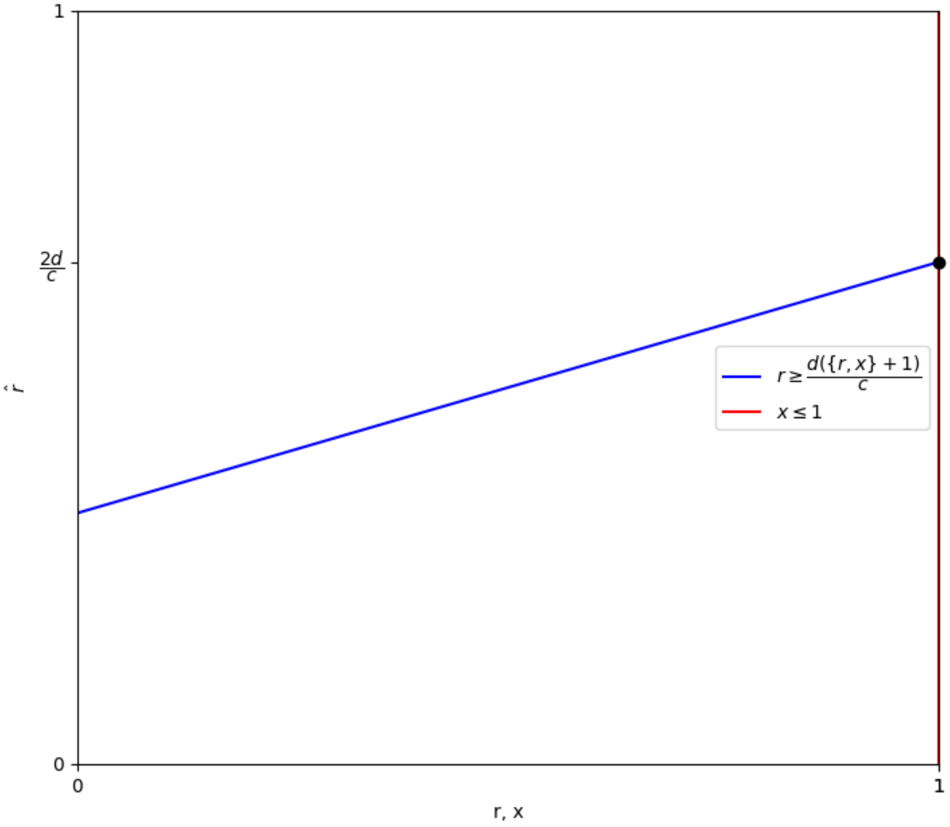

Figure 4.6: Optimal values for r, $\hat{r}$ and x in Case 7 if $c \geq 2d$

which holds with equality as the goal is to maximize x and minimize q. Using the latter to solve the obedience inequality, we get that

$$r \geq \frac{1}{2} - \frac{2xc}{c+d}$$

Moreover, (C2) and incentive compatibility together require that

$$r \leq 1 - 3x$$

Figure 4.7 shows these equations and the optimal values.

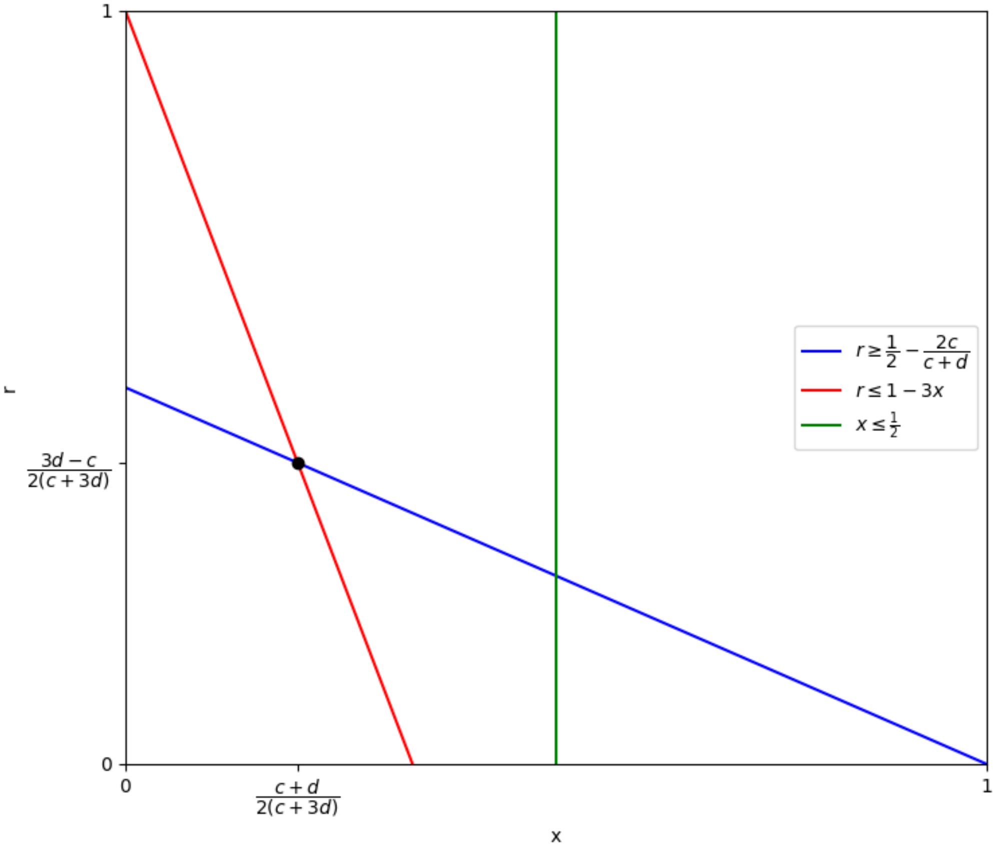

Figure 4.7: Optimal values for r and x in Case 8

Next, let us now see if sender can benefit from allowing the communication rule to be non-symmetric with respect to types. Let us again assume that $q = \hat{q} = z = 0$, and as for the type-symmetric rule of this case $v = x$ and $z = 1 - 2x - y$. Then obedience requires that

$$\hat{r} \geq 1 - y + \frac{x(d-3c)}{c+d}$$

Observe that if $d > 3c$, then $y > 0$. Next, looking at incentive compatibility, we have

$$r \leq \frac{\hat{r}c - (x-y)(d-c)}{d}$$

Since the goal for sender is to minimize $\hat{r}$, (C2) binds such that $y = 1-2x$. Using this in obedience,

$$\hat{r} \geq \frac{x(3d-c)}{c+d}$$

This will hold with equality, such that the incentive constraint results in

$$r \leq \frac{x(2c^2+3cd-3d^2)}{d(c+d)} + \frac{d-c}{d}$$

and as such, at best $r = \frac{d-c}{d}$. Notice also that the above has a negative slope if $d \geq c\left(\frac{3+\sqrt{33}}{6}\right)$. Recall that

$$\hat{EV} = \frac{1}{4}(r-\hat{r})(m-n) + \frac{1}{2}x(m-n+2l)$$

and that if we had a type-symmetric communication rule,

$$EV = \frac{1}{2}(x-q)(m+n-2l)$$

Thus, we can affirm that going the non-symmetric route is beneficial for sender only if $\hat{EV} > EV$. Using the maximum for r in the non-symmetric variation, and the optimal value of $x = \frac{c+d}{2(c+3d)}$ from its symmetric counterpart, we get that $\hat{EV} > EV$ if

$$\frac{d(c+d)}{(c+3d)(d-c)} < \frac{m-n}{m+n-2l}$$

Graphically, this is depicted in Figure 4.8.

Case 8': $n > m > l$ & $c < d$

Similar to Case 8.

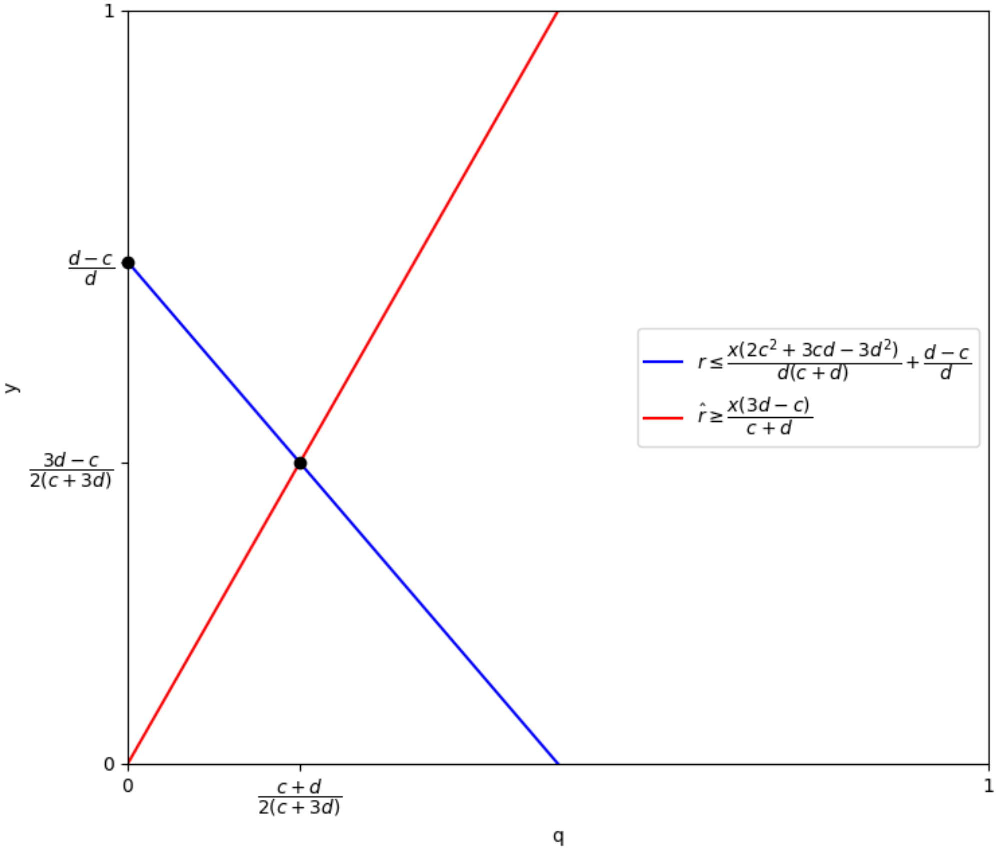

Figure 4.8: Optimal values for q and y in Case 8 in the non-symmetric setting

Case 9: $m > l > n$ & $c > d$

Recall sender's expected utility

$$EV = \frac{1}{2}\{x(m-n) + (y+z)(l-n) - q(m+n-2l)\}$$

Here, maximizing x is a priority over y and z, but q contributes positively to sender's expected payoff if $m + n < 2l$. As before, let us start with the type-symmetric rule, and let $y = z = 0$. Then, the incentive constraint for t^1 is given by

$$r \geq q + \frac{x(c+d) - c}{c-d}$$

If $m + n > 2l$, then $q = 0$ and the above becomes

$$r \geq \frac{x(c+d) - c}{c - d}$$

which is less or equal than 1 if $c \geq 2d$. If, on the contrary, $c < 2d$, then for $r = 1$ the incentive constraint requires that

$$x \leq \frac{2c - d}{c + d}$$

Figures 4.9 and 4.10 show the optimal values of r and x if $c = 2d$ and $c < 2d$. As we will cover shortly, $c \geq 2d$ is a matter of the non type-symmetric rule.

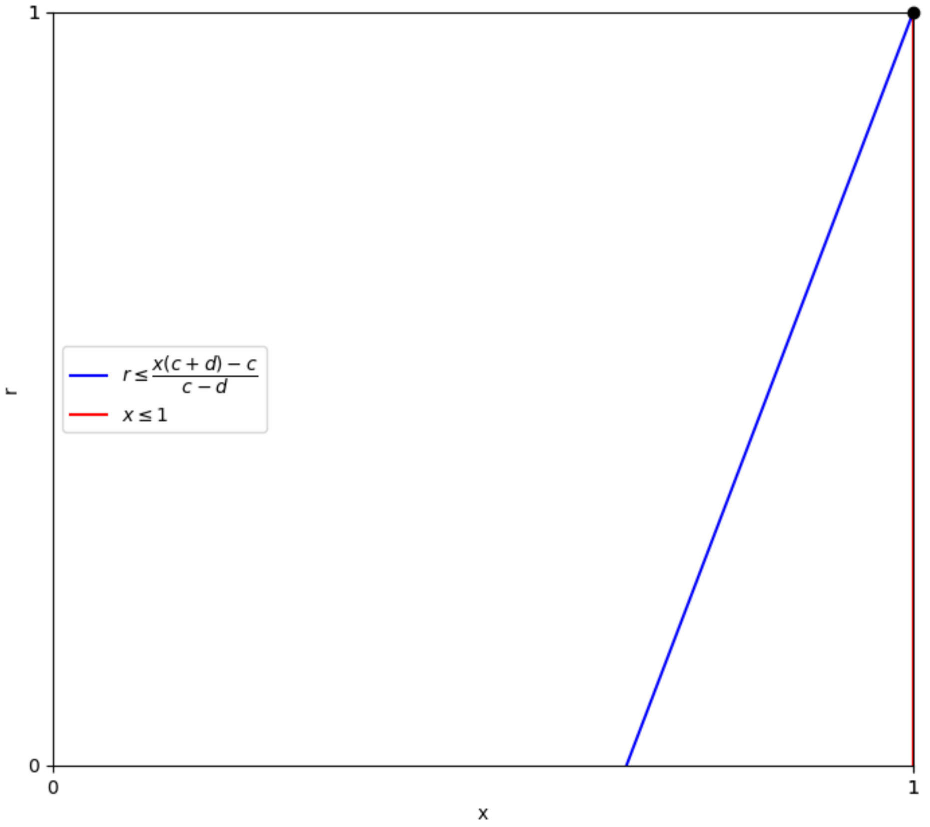

Figure 4.9: Optimal values for r and x in Case 9 if $c = 2d$

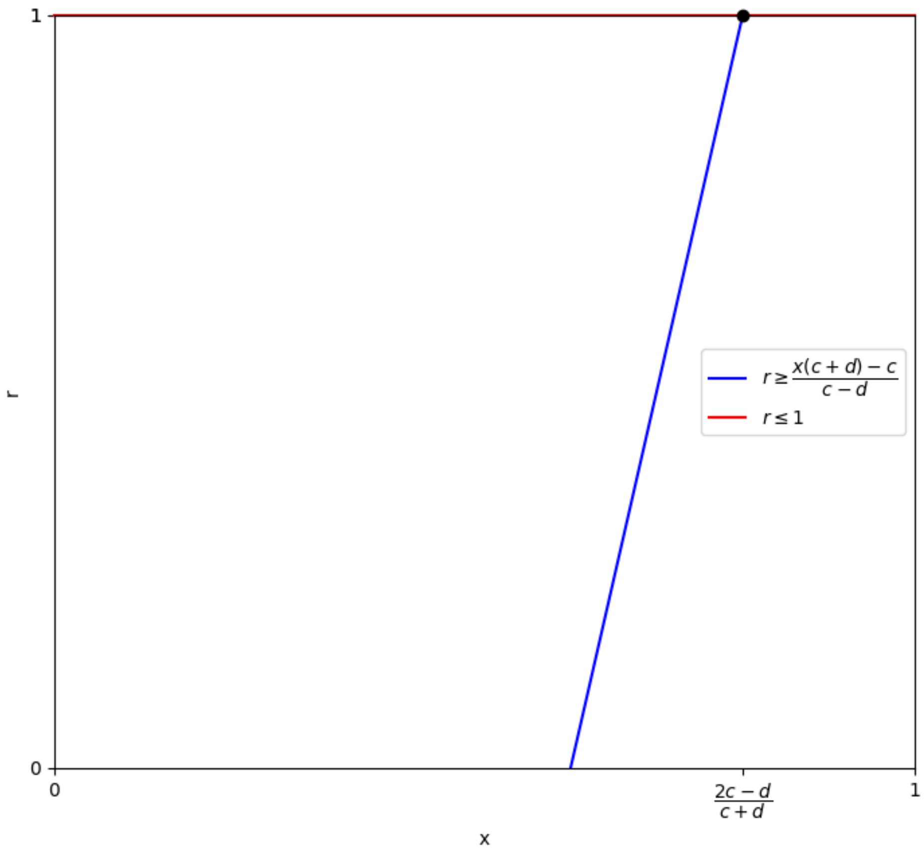

Figure 4.10: Optimal values for r and x in Case 9 if $c < 2d$

Before proceeding to the non-symmetric rule, observe that if x is small enough, it might be beneficial for sender to have $y, z > 0$ and $v = x = 0$. Under these assumptions, we again have the incentive constraint as

$$r \geq q + y$$

For this part, let us assume that $m + n < 2l$, such that EV increases in q. Using the above to solve the obedience constraint, we have

$$q \leq \frac{(1-2y)(c+d)}{4d}$$

The non-negativity constraint thus implies that $y \leq \frac{1}{2}$. Since the goal is to maximize q, (C1) binds and the incentive constraint reduces to

$$y \leq 1 - 3q$$

Figure 4.11 depicts the values of q as a function of y.

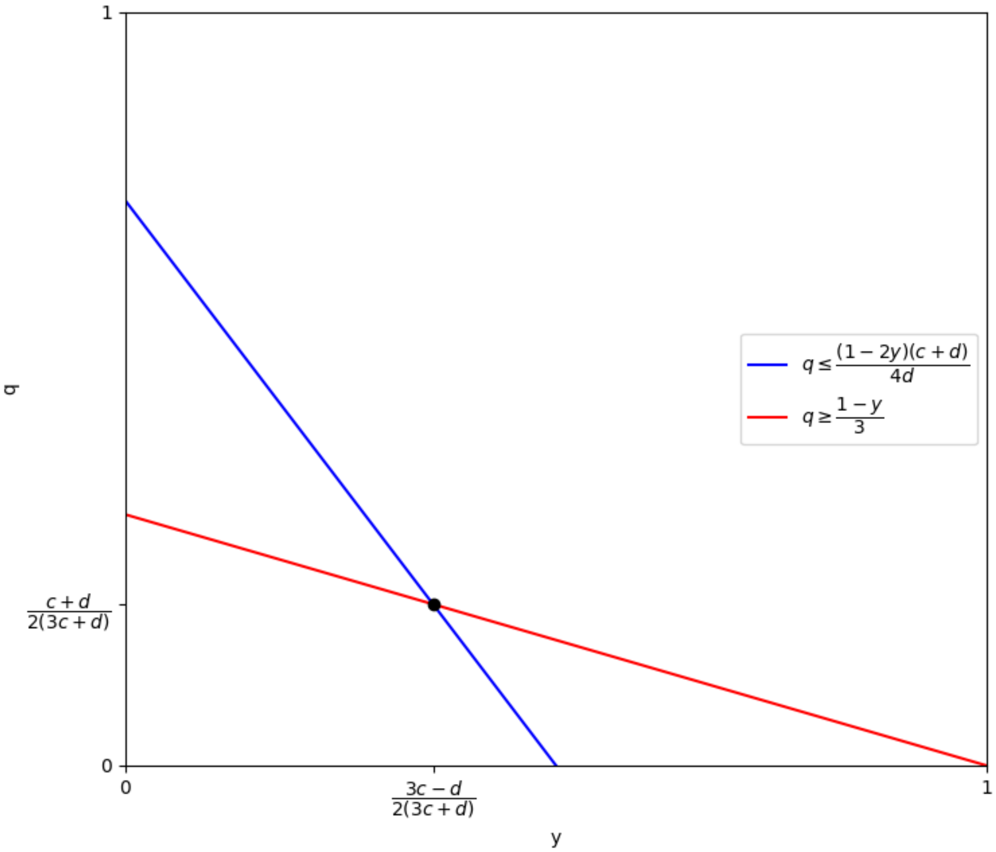

Figure 4.11: Optimal values for q and y in Case 9 in the non-symmetric setting

Now, defining the optimal communication rule from these values benefits sender if it effectively increases her expected utility $\widetilde{EV} > \overline{EV}$, when $c \leq 2d$. There are two possible instances of such scenario:

(a) $\widetilde{EV} > EV$: Improvement upon the symmetric case. This holds if

$$\frac{4(2c+d)(c+d)}{13c^2-d^2} > \frac{m-n}{l-n}$$

(b) $\widetilde{EV} > EV^3$: Improvement upon Subcase 3 (covered ulteriorly). This requires that the following inequality to be satisfied

$$2(11c^3-7c^2d-cd^2+2d^3)(l-n)+2(3c+d)(c-d)^2(m-l) > 2(13c^3-13c^2d+2d^3)(m-n)$$

For the non type-symmetric rule, because of the increase in complexity, we will split it into different subcases. Each of these subcases implies simplifying assumptions.

Subcase 1: $r = x = 1$**,** $q = \hat{q} = 0$

From the obedience constraint for t^1,

$$\hat{r} \geq \frac{2d}{c+d}$$

and from incentive compatibility,

$$\hat{r} \geq \frac{2d}{c}$$

The latter binds as it is the highest out of the two. This is an improvement from the type-symmetric rule for the sender if the expected utility she gets in this case, EV^1, is higher than EV – the expected utility from the type symmetric rule. Since $r = x = 1$, this is the case if $r - \hat{r} > 0$, which is the case if $c > 2d$. Observe that if $c < 2d$, then the incentive constraints creates a contradiction.

Subcase 2: $r = x = 1$**,** $\hat{q} > 0$

From obedience for t^1,

$$\hat{r} \geq \frac{\hat{q}(c-3d)+2d}{c+d}$$

and from incentive compatibility,

$$\hat{q} \geq 2 - \frac{\hat{r}c}{d}$$

Combining the two gives

$$\hat{r} \geq \frac{2(c-2d)}{(c-d)^2}$$

Now, we have to consider two possible scenarios: (a) $EV^2 > EV^1$ if $c > 2d$ and (b) $EV^2 > EV$ if $c \leq 2d$

(a) The requirement for this to be the case is

$$\frac{d(2d(c-d)^2 - 2c(c-2d))}{c(d(c-d)^2 - c(c-2d))} > \frac{4(m-l)}{m-n}$$

(b) The requirement for this part is that

$$\frac{c(c-2d)}{d(c-d)^2} > \frac{1}{2}$$

which has no solution for $c < 2d$. Thus, the type symmetric rule is better in this case.

Subcase 3: $r = 1$**,** $x < 1$**,** $\hat{q} > 0$

Let $y = 1 - x - y - z = 0$. Then, incentive compatibility for t^1 gives

$$\hat{r} \geq \frac{1 - \hat{q} + x}{c}$$

Observe from the above inequality that sender wants to maximize x, and to minimize $\hat{r}$ and as such, the type-asymmetric equivalent of (C1) ($1 - 2\hat{q} - \hat{r} \geq 0$) will hold with equality. This allows us to simplify the above and get

$$\hat{q} \leq \frac{c - d(1+x)}{2c - d}$$

We can do the same with obedience, such that

$$\hat{q} \geq \frac{x(c-d)}{3c-d}$$

With these inequalities, we can then identify the optimal values of x, $\hat{q}$ and $\hat{r}$:

$$x = \frac{(c-d)(3c-d)}{2c^2} \qquad\qquad \hat{q} = \frac{(c-d)^2}{2c^2}$$

Now it must be that $z = 1 - x \geq 0$, which requires that $c < d(2+\sqrt{3})$. Finally, we can deduce that

$$\hat{r} = 1 - 2\hat{q} = \frac{d(2c-d)}{c^2}$$

As before, this can be beneficial to sender under the following scenarios:

(a) $EV^3 > EV^1$ if $c > 2d$. This holds if

$$m > \frac{4c^2l - 2cdn + d^2l}{4c^2 - 2cd + d^2}$$

(b) $EV^3 > EV$ if $c \leq 2d$. This is true as long as the following condition is met

$$(4cd - c^2 - d^2)(l-n) - (c-d)^2(m-l) > \frac{2d^2(2c-d)(m-n)}{c+d}$$

Case 10: $m > l > n$ & $c < d$

Let us again begin this case by looking at the type-symmetric communication rule. Observe that here, $r = x = 1$ cannot be a solution. Indeed, from the incentive constraint of t^1, this would mean

$$1 \leq \frac{d}{c-d}$$

which does not hold for $c < d$. Suppose instead that $x = 1$, then the same constraint requires

$$r \leq q + \frac{d}{d-c}$$

which from non-negativity imposes that

$$q \geq \frac{d}{d-c}$$

which is greater than 1 and therefore generates a contradiction. Hence, it must be that $x < 1$. Let us look at the incentive constraint.

$$(r-q)(c-d) \geq xd + y(c-d)$$

Therefore, depending on if $m + n - 2l$ is positive or not, sender may want to minimize q. If it is negative, then let $q = \frac{1}{2}$, its maximum value, and let $z = 0$ and (C2) bind. Then the above gives

$$x \leq \frac{3(d-c)}{2(2d-c)}$$

Then, sender's expected utility is

$$EV = \frac{1}{2}\left(\frac{(d-2c)m + 3(c-2d)n + l(5d-c)}{2(2d-c)}\right)$$

Instead, if $m + n - 2l \geq 0$, then let $q = 0$ and $r = 1$ – for the goal is to maximize x – in the above inequality, such that

$$c - d \geq xd + y(c-d)$$

This implies that $x = 0$ and $y = 1$, and $EV = \frac{1}{2}(l-n)$. If we relax the $r = 1$ assumption, then from the same procedure we get that

$$x \leq \frac{d-c}{2d-c}$$

Then, sender's expected payoff is

$$EV = \frac{1}{2}\left(\frac{(d-2c)m + 4(c-2d)n + (7d-2c)l}{3(2d-c)}\right)$$

Notice that relaxing $r = 1$ always benefits sender if $d \geq 2c$, and otherwise might depend on the values of m, n and l. Below is the condition for it to be beneficial.

$$m < \frac{l(c+d) - n(2d-c)}{2c-d}$$

Now, sender could also opt for a symmetric rule such that $v = x$. Then, let $q = 0$ and $z = 1 - 2x - y = 0$ to match sender's objective. The obedience constraint becomes

$$r \geq \frac{y(c-d) - c + d}{c+d} + x$$

which will bind if x is to be maximized. Using this, the incentive compatibility for t^1 reduces to

$$x \leq \frac{y(2d^2 - 3cd - c^2) + 3cd - d^2}{2d(c+d)}$$

The above will hold with equality for the same reason mentioned before, and we can substitute x into the obedience for t^0 to get the value of y

$$y \geq \frac{2d(2cd - c^2 - d^2)}{5cd^2 - cd - c^3 - 3d^3}$$

These last two inequalities hold with equality as required by the maximization problem, and represent the optimal values for x and y. These values of r, x and y yields a higher expected utility to the sender compared to $q = \frac{1}{2}$, $x = \frac{3(d-c)}{2(2d-c)}$ and $y = 1 - x$ if

$$\begin{aligned} m &\leq m^*_{(b)} \\ &= \frac{(2l-n)(c^4 + c^2d + 8cd^3) - 3c^3d(2l-3n) - c^2d^2(2l+11n) - cd^2(4l-5n) - d^4(8l-7n)}{c^3(c+3d) - c^2d(13d-1) + cd^2(8d+1) - d^4} \end{aligned}$$

Moreover, they are also preferred to $r = p = q = \frac{1}{3}$, $x = \frac{d-c}{2d-c}$ and $y = 1 - x$ if

$$m \leq m^*_{(c)}$$
$$= \frac{c^2 l(c^2+d) - 2c^3 d(2l-3n) + c^2 d^2(3l-16n) + cd^3(3l+10n) - 2cd^2(l-n) - 2d^4(l+n)}{c^3(c+2d) - c^2 d(13d-1) + d^3(13c-4d)}$$

For clarity, we categorize every case as follows:

- Cases 1, 5 and 5' are regrouped in point 1
- Cases 2, 6 and 6' are regrouped in point 2
- Case 3 corresponds to point 3
- Case 4 corresponds to point 4
- Case 7 and 7' are regrouped in point 5 and its subcases
- Case 8 and 8' are regrouped in point 6 and its subcases
- Case 9 corresponds to point 7, where
    - (a) (b) and (c) are integrated in the body of Case 9
    - (d) is represented by Subcase 3
    - (e) is represented by Subcase 1
    - (f) is represented by Subcase 2
- Case 10 corresponds to point 8

■

Bibliography

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bergemann, D. and Morris, S. (2017). Information design: A unified perspective. *SSRN Electronic Journal*.

Bhattacharya, S., Goltsman, M., and Mukherjee, A. (2018). On the optimality of diverse expert panels in persuasion games. *Games and Economic Behavior*, 107:345–363.

Bhattacharya, S. and Mukherjee, A. (2013). Strategic information revelation when experts compete to influence. *The RAND Journal of Economics*, 44(3):522–544.

Caillaud, B. and Tirole, J. (2007). Consensus building: How to persuade a group. *American Economic Review*, 97(5):1877–1900.

Candogan, O. and Strack, P. (2023). Optimal disclosure of information to privately informed agents. *Theoretical Economics*, 18(3):1225–1269.

Caroll, G. (2016). Informationally robust trade and limits to contagion. *Journal of Economic Theory*, (166):334–361.

Dewatripont, M. and Tirole, J. (1999). Advocates. *Journal of Political Economy*, 107(1):1–39.

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of accounting research*, pages 123–145.

Farrell, J. and Gibbons, R. (1989). Cheap talk with two audiences. *The American Economic Review*, 79(5):1214–1223.

Gentzkow, M. and Kamenica, E. (2017). Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322.

Glazer, J. and Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736.

Goltsman, M. and Pavlov, G. (2011). How to talk to multiple audiences. *Games and Economic Behavior*, 72(1):100–122.

Gossner, O. (2000). Comparison of information structures. *Games and Economic Behavior*, 30:44–63.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483.

Gul, F. and Pesendorfer, W. (2012). The war of information. *The Review of Economic Studies*, 79(2):707–734.

Guo, Y. (2019). Information transmission and voting. *Economic Theory*, 72(3):835–868.

Inostroza, N. and Pavan, A. (2023). Adversarial coordination and public information design. *Available at SSRN 4531654*.

Jung, W.-O. and Kwon, Y. K. (1988). Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting research*, pages 146–153.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kartik, N., Lee, F. X., and Suen, W. (2017). Investment in concealable information by biased experts. *The RAND Journal of Economics*, 48(1):24–43.

Kolotilin, A. (2018). Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017). Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964.

Li, F. and Norman, P. (2021). Sequential persuasion. *Theoretical economics*, 16:639–675.

Lynas, M., Houlton, B. Z., and Perry, S. (2021). Greater than 99% consensus on human caused climate change in the peer-reviewed scientific literature. *Environmental Research Letters*, 16(11):114005.

Mathevet, L., Perego, J., and Taneva, I. (2020). On information design in games.

Milgrom, P. and Roberts, J. (1986). Relying on the information of interested parties. *The RAND Journal of Economics*, 17(1):18.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380.

Myerson, R. (1982). Optimal coordination mechanisms in generalized principal-agent problems.

Myerson, R. B. (1991). *Game theory: analysis of conflict.* Harvard university press.

Okuno-Fujiwara, M., Postlewaite, A., and Suzumura, K. (1990). Strategic information revelation. *The Review of Economic Studies*, 57(1):25.

Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987.

Shavell, S. (1994). Acquisition and disclosure of information prior to sale. *The RAND Journal of Economics*, 25(1):20.

Shin, H. S. (1994). The burden of proof in a game of persuasion. *Journal of Economic Theory*, 64(1):253–264.

Shin, H. S. (1998). Adversarial and inquisitorial procedures in arbitration. *The RAND Journal of Economics*, 29(2):378.

Taneva, I. (2019). Information design. *American Economic Journal: Microeconomics*, 11(4):151–185.

Wolinsky, A. (2002). Eliciting information from multiple experts. *Games and Economic Behavior*, 41(1):141–160.

Zeng, Y. (2023). Derandomization of persuasion mechanisms. *Journal of Economic Theory*, 212:105690.

Ziegler, A. (2023). Persuading an audience: Testing information design in the laboratory. Technical report, Tinbergen Institute Discussion Paper.