The London School of Economics and Political Science

Essays on Microeconomic Theory and Behavioral Economics

Zhuoqiong Chen

A thesis submitted to the Department of Management of the London School of Economics for the degree of Doctor of Philosophy, London, September 2016.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 52,409 words (including footnotes but excluding bibliography and appendices).

Statement of Conjoint Work

I confirm that Chapter 3 was jointly co-authored with David Ong (Peking University HSBC Business School) and Ella Segev (Ben-Gurion University of the Negev). I contributed 50% of this work. I confirm that Chapter 4 was jointly co-authored with Tobias Gesche (University of Zurich). I contributed 50% of this work.

London, September 30, 2016

Zhuoqiong Chen

Abstract

The dissertation consists of four chapters. The first two chapters are devoted to exploring information acquisition and disclosure in contests. The third chapter is devoted to exploring how risk attitude affects bidding behavior in all-pay auctions. The last chapter is devoted to exploring behavioral biases in advice-giving.

In Chapter 1, I study player's incentive to spy on opponents' private information in contests. I show that each player's equilibrium effort is non-decreasing (non-increasing) in the posterior probability that the opponent has the same (a different) valuation. Accounting for the cost of spying, players are strictly better off than not spying on each other at all.

In Chapter 2, I focus on how a contest organizer should disclose information in order to achieve certain objectives. In particular, I compare private signals with public signals. I show that there is no general ranking of the two signals in terms of the performance of maximizing players' expected payoff, but public signals outperforms private signals in maximizing expected effort.

In Chapter 3 (co-authored with David Ong and Ella Segev), we extend previous theoretical work on n-players complete information all-pay auction to incorporate heterogeneous risk and loss averse utility functions. We provide sufficient and necessary conditions for the existence of equilibria with a given set of active players with any strictly increasing utility functions and characterize the players' equilibrium mixed strategies.

Finally, in Chapter 4 (co-authored with Tobias Gesche), we show experimental evidence that a one-off incentive to bias advice has a persistent effect on advisers' own actions and their future recommendations.

Acknowledgments

I would like to express my sincere gratitude to my supervisor, Kristof Madarasz, for his invaluable suggestions and continuous support of my Ph.D study, and to my advisor, Ricardo Alonso, for his stimulating discussions, endless support, patience and kindness. Their guidance helped me in all the time of research and writing of this thesis.

In addition, I have greatly benefited from many discussions with David de Meza and my former supervisor, Luis Rayo, and the support and advice from Catherine Thomas. I would also like to thank Balazs Szentes, Ronny Razin, Andrew Ellis, and Francesco Nava, for helpful comments on my papers in Chapter 1 and 2, and Erik Eyster for discussions and comments on my experimental paper in Chapter 4.

I would like to thank the fellow Ph.D students in the Management and Economics departments, for their help and support. I am indebted to my parents, Boning Chen and Xinhong Feng, who have been encouraging my academic pursuits. Finally, I am indebted to my wife, Lin Chen, who sacrifices many career opportunities to support me throughout the years.

Contents

List of Tables

List of Figures

Chapter 1

Spying in Contests

Abstract Two players compete for a prize and their valuations are private information. Before the contest, each player can covertly acquire a costly, noisy and private signal regarding the opponent's valuation. In equilibrium, each player's effort is non-decreasing in the posterior probability that the opponent has the same valuation. Accounting for the cost of spying, players are strictly better off spying when the spying technology is partially but not perfectly informative. Suppose instead that each player can, at no cost, ex ante commit to disclose a signal about her valuation to the opponent, but cannot observe realizations of the signal. Then every equilibrium involves non-disclosure by at least one player, even though some disclosure by each player would benefit both.

1.1 Introduction

Winner-take-all contests, like rent-seeking contests for monopoly rights, patent races, lobbying, political campaigns and competitions for promotion, burden participants with the prospect that their investments may yield no reward. The efforts, time and resources invested in competing for the prizes are unrecoverable, and typically, only the participant with the highest investment reaps the rewards of the contest. Thus, anticipating the rivals' intentions becomes particularly valuable; learning that rivals will invest little can save on the investment to win the prize, and, conversely, learning of an excessive investment outlay by rivals would lead a firm to avoid investing in a lost cause. This paper studies players' incentives of acquiring information about the opponents prior to winner-take-all contests.

In competing for a procurement contract, for example, suppliers spend enormous time, resources and efforts to prepare proposals for a buyer to evaluate.[1] This process is even costlier when it also involves bribing the procurement agent (Celentani and Ganuza, 2002; Burguet and Che, 2004). Since each supplier may value the contract differently, their willingness to commit resources to win the contract or to bribe the procurement agent may differ. Gathering intelligence on the opponent's valuation can prove particularly valuable. To obtain these intelligence, suppliers may hire hackers to steal information from rival supplier's computer, investigators to search through office trash or detectives to steal files from office safe, etc.[2]

The existing literature suggests that players will overall not benefit from such spying. Kovenock et al. (2015) show that the payoffs to players are the same when valuations are commonly known and when they are private information. However, perfect information about the opponent is extremely hard – if not impossible – to acquire in reality.

In this paper, we model the action of information acquisition as a continuous variable which allows players to acquire partial information. In this novel set up, we seek to answer the following questions: What's the impact of partial intelligence about opponents on a player's competitive behavior in contests? What's the implica-

[1] Airbus and Boeing spent 10 years in competing for the U.S. Air Force tanker contract and their proposals included several thousand pages (Kovenock et al., 2015).

[2] On the one hand, acquisition of intelligence can be illegal. Detectives hired by Larry Ellison, the head of Oracle, bribed the cleaning staff at Microsoft's office to gather sensitive information from the office trash until the year when the scandal was exposed by the media. In 2001, staffs of Procter&Gamble were found searching the garbage of Unilever – its competitor in the hair-care market – for "the Organics and Sunsilk brands of shampoo" which contains commercially sensible information. On the other hand, intelligence may also be acquired through legal and organized methods. Large companies usually hire competitive intelligence agencies to study their competitors; some major multinational firms like General Motors, Kodak and BP even set up their own separate competitive intelligence units (Billand et al., 2016).

tion of such spying behavior on social welfare and allocative efficiency? In addition, a large set of contests are welfare destroying in the sense that resources invested by players are wasted (Tullock, 1967; Posner, 1974), as these resources only determine the winner but do not contribute to value creation.[3] For instance, the estimated social cost of rent-seeking for the US is 22.6 percent of GNP in 1985 (Laband and Sophocleus, 1988); and it is been long argued (since (Wright, 1983)) that patent races generate wasteful duplication of effort. So how do spying activities affect total efforts in such wasteful contests?[4]

Section 1.2 presents a model of a contest with one indivisible prize and two players who have independent private valuations (IPV) for the prize. In particular, each player's valuation can be either *high* or *low*. Before participating in the contest, each player covertly acquires a costly, noisy and private *spying* signal about her opponent's valuation. The *spying* signal can be drawn from an arbitrarily large set of distribution functions. In acquiring the *spying* signal, she chooses a level of accuracy for it ranging from completely uninformative to perfectly informative.[5] She then observes both her valuation and the *spying* signal and exerts effort in the contest. The player who exerts more effort wins. This payoff structure is the same as in a first price all-pay auction. For example, in competitive procurement each supplier may have a high or low valuation of the contract, according to its estimation of profit which depends on its own production cost. The supplier can obtain intelligence regarding the opponent's valuation from its pre-existing competitive intelligence unit, and then decides how much effort to invest in the competition.

In Section 1.3.1 we consider a simplified setting where the accuracy of the *spying* signal is fixed and the signal costless. When the signals are partially informative, effort by each player in the unique symmetric equilibrium of the contest is stochastically increasing in her valuation. Furthermore, spying has a motivation/demotivation effect on a player (Proposition 1): when the player's valuation is *high*, her equilibrium effort is non-decreasing in the posterior likelihood that the opponent's valuation is also *high* ("motivation effect" of spying); instead, when her valuation is *low* then her effort is non-increasing in such a likelihood ("demotivation effect" of spying). As a result, players increase their efforts when they are perceived to be evenly matched. In competitive procurement process, a supplier only devotes more resources, efforts and time after learning the opponent might have the similar production cost, and thus a similar valuation for the procurement contract. A firm would increase investments in R&D after learning that the opponent in the patent

[3] See Congleton et al. (2008) for the overview of rent-seeking contests.

[4] The results are also applicable to other contests where efforts are productive.

[5] Accuracy is defined by rotation order (Johnson and Myatt, 2006) for tractability.

race owns a research team with the similarly background as the firm does.

Section 1.3.2 considers endogenous acquisition of *spying* signals. In the symmetric equilibrium, each player acquires a partially informative *spying* signal. Accounting for the cost of spying, players are strictly better off than not spying on each other at all (Proposition 3). Meanwhile, the total expected effort in the contest is strictly lower. Therefore, spying in wasteful contests actually contributes to social welfare by reducing duplication of efforts or other resources wasted in the competition. In procurement contests, money spent on bribery and efforts exerted in the process are lower if suppliers hack into each others' computers or steal files from rivals' offices for information. This result is particularly important because public procurement is a hotbed for bribery among OECD countries (Ehlermann-Cache and Others, 2007) which creates social inefficiencies and is hard to detect. When suppliers spy on each other, such costs are reduced without actually detecting the bribing behavior.[6] Interestingly, the same welfare outcome is not achievable if the cost of spying is zero, as then players would acquire a perfect signal about the opponent. As mentioned earlier, players' payoffs in this situation are the same as when they do not spy at all.

Spying maybe prohibitively costly (or illegal). The previous result suggests, however, that players would benefit if they were to disclose to each other a noisy signal of their valuation. This raises the question: would such disclosure be supported in equilibrium? To address the question, Section 1.4 considers a twist of the main model in which each player commits to disclose a signal about her own valuation to the opponent before the contest.[7] In doing so, she chooses an accuracy for the signal which the opponent will receive. Neither disclosing nor receiving the signal incurs any cost to any player. This corresponds to each supplier choosing its security level of office buildings or firewall.[8] Should suppliers loosen their security measures to make it easier to steal sensible files or downgrade its firewall so it is easier to hack into the computers?

Section 1.4.2 considers the case when players set up an agreement to disclose

[6] There is a similar implication to rent-seeking and lobbying contests: bribes to politicians are social costs and spying in these contests improves welfare. In the patent race example, if firms actively acquire intelligence about the opponents', say, research budgets, capabilities, breakthroughs as well as data of the new products, then there will be less duplicated investments in R&D.

[7] Most of information disclosure in contest literature often take a centralized view and analyze how a contest organizer should disclose information to players in order to maximize total effort (Lu et al., 2016; Zhang and Zhou, 2016; Chen, 2016; Serena, 2015; Denter et al., 2014). The current paper, however, takes a decentralized view and considers players disclose information to each other.

[8] Alternatively, in the patent race example, information disclosure corresponds to providing the opponent, for instances, a prototype of the new product, or samples of a new drug. Given such a piece of hard evidence (i.e. a given accuracy), the opponent firm can test the product and test results are unavailable to the firm who discloses the information.

signals to each other. If players agree to disclose partially informative signals to each other, then both players are strictly better off (Proposition 6). This is in contrast to Kovenock et al. (2015) which restricts attention to binary disclosure, i.e. non-disclosure or full disclosure, and shows that full disclosure has no impact on players' expected payoffs.[9] However, Section 1.4.3 shows that disclosing a partially informative signal is weakly dominated by disclosing an uninformative signal for each player (Lemma 7). Even though partial disclosure by both players lowers ex ante expected efforts and increases payoffs, each player can do better by adding noise to the signal disclosed to the opponent, as then the evenly matched opponent is more likely to be demotivated. Therefore, there does not exist any equilibrium in which both players disclose any partially informative signals (Proposition 7). This provides one reason for suppliers seldom sharing private information to each other in real procurement competitions or any other winner-take-all contests.

These results suggest that social welfare can be improved through mandatory, albeit imperfect, disclosure between players. In a patent race, the government can require all participating firms to disclose information about their research teams, or their research budget, or any other related information which affects their R&D investments. In lobbying competitions, lobbyists should be required to disclose their estimation of profit when their preferable policy is implemented. The results also provide a nice interpretation of why we observe spying/espionage but information disclosure/sharing in reality.

The contributions of this paper are threefold. First, the clandestine nature of spying is captured in the model as both *spying* signal realizations and accuracy of the signal are private information to each player. Microsoft was unaware of whether or how the detectives hired by Oracle acquired any information and was also unaware of what kind of information Oracle had obtained. Interestingly, the results suggest spying on each other may improve both players welfare as they receive additional information about each others' strategy and thus, are able to coordinate. Second, Section 1.3.1 of the current paper considers the all-pay auction with an information structure between incomplete information (Lu and Parreiras, 2014; Konrad, 2004; Amann and Leininger, 1996) and complete information setting (Baye et al., 1996; Ellingsen, 1991; Hillman and Riley, 1989). By varying the accuracy of the *spying* signal from completely uninformative to perfectly informative, the current paper characterized the equilibrium in an arbitrarily large set of information structure between incomplete and complete information.[10] Third, this paper provides the

[9]Yet, the allocative efficiency is compromised and the total expected effort in the contest is reduced.

[10]There has been some recent progresses in all-pay auction with common valuation and affiliated

first analytical framework of all-pay contests with endogenous information structure. The model is applied to study endogenous information acquisition (spying) and endogenous information sharing (disclosure) in the current paper, and is applicable to other endogenous information settings, including overt information acquisition, centralized information disclosure and discriminatory information acquisition, etc.

There are few studies on spying/information acquisition in contests. The paper most related to the current paper in the economics literature is Baik and Shogren (1995) who compare covert information acquisition with overt acquisition in contests. However, the paper does not model contests as a game of incomplete information and is subject to the criticism for "the negligence of strategic interdependency between the two players" (Bolle, 1996).[11] Alternatively, Zhang (2015) considered one-sided private information setting in which the public player spies on the private player in both all-pay auction and Tullock contest (Tullock, 1967). The author shows that espionage is more common in a more discriminatory contest and can be discouraged by the increasing probability of spying detection. The current paper does not consider detection of spying or double agent. Another closely related strand of literature is the information acquisition in winner-pay auction literature (Miettinen, 2013; Shi, 2012; Persico, 2000; Matthews, 1984). In all-pay auctions, Morath and Münster (2013) considers players' incentive to acquire information about their own valuation, and the decision of information acquisition is binary. Alternatively, Fang and Morris (2006) from which the research idea of the current paper was originated shows a numerical example of acquiring information about the opponent in a first price winner-pay auction. Built on Fang and Morris (2006), Tian and Xiao (2007) studies endogenous information acquisition on the opponent's valuation in the first price auction. Finally, in the IO literature, spying/espionage are also considered in entrant deterrence (Barrachina et al., 2014; Solan and Yariv, 2004), in price and quantity competition in duopoly (Kozlovskaya, 2016; Wang, 2016; Whitney and Gaisford, 1999), and in multi-market competitions (Billand et al., 2016).

The remainder of the paper are organized as follows. Section 1.2 presents the preliminaries of the main model of spying in contests. Section 1.3 shows the results on spying – endogenous information acquisition – in contests, and Section 1.4 on information disclosure. Section 1.5 concludes.

---

signals (Rentschler and Turocy, 2016; Chi et al., 2015), as well as with interdependent valuations (Siegel, 2013; Krishna and Morgan, 1997). These studies assume players can learn some information regarding the opponent from the signal about their own valuations, because of common valuation or interdependent value assumption. The current paper is different from this line of research in the independence between the information regarding the bidder's own valuation and the information regarding the opponent.

[11] See discussions in Baik and Shogren (1995), Bolle (1996) and Baik and Shogren (1996)

1.2 The model

There are two risk neutral players, indexed by $i \in \{1, 2\}$, who compete in a contest with one indivisible prize for which they have independent private valuations (IPV). Player i ($i = 1, 2$) may value the prize at θ_h with probability $p_h \in (0, 1)$ or at θ_l with probability $p_l = 1 - p_h$, where $\theta_h > \theta_l > 0$. Players know only their own valuations, and the distributions of opponents' valuations. We refer to a generic player i as "she" and her opponent, player j, as "he".

Information acquisition (spying): Player i can acquire additional information regarding the opponent by receiving a private *spying* signal (hereafter "signal") about the opponent's valuation. The possible signal realization π_i is drawn from a compact set $[\underline{\pi}, \overline{\pi}]$. Player i acquires information about θ_j by choosing from a family of joint distributions over $[\underline{\pi}, \overline{\pi}] \times \{\theta_h, \theta_l\}$

$$\{F(\pi_i, \theta_j | \alpha_i)\}_{\alpha_i \in [\underline{\alpha}, \overline{\alpha}]}$$

indexed by $\alpha_i \in [\underline{\alpha}, \overline{\alpha}]$. We refer to $F(\pi_i, \theta_j | \alpha_i)$ the signal, α_i the accuracy (defined shortly) of the signal and π_i the realization of the signal. Since the conditional distribution of π_i depends only on θ_j, and the prior distribution of θ_i is independent of the distribution of θ_j, π_i is thus independent of π_j.

Let $F(\cdot, \alpha_i)$ denote the marginal distribution of π_i with corresponding density $f(\cdot, \alpha_i)$, given any α_i. Furthermore, denoted by $F_h(\cdot, \alpha_i)$ ($F_l(\cdot, \alpha_i)$) the conditional cumulative distribution of π_i given $\theta_j = \theta_h$ ($\theta_j = \theta_l$). Let $f_h(\cdot, \alpha_i)$ and $f_l(\cdot, \alpha_i)$ be the corresponding probability density functions, and assume both are differentiable on both arguments. We assume w.l.o.g. the marginal distribution of π_i, $F(\pi_i, \alpha_i)$ is uniform on $[0, 1]$ for every given α_i,[12] i.e.,

$$p_h F_h(\pi_i, \alpha_i) + p_l F_l(\pi_i, \alpha_i) = \pi_i \tag{1.1}$$

$$p_h f_h(\pi_i, \alpha_i) + p_l f_l(\pi_i, \alpha_i) = 1 \tag{1.2}$$

given any $\alpha_i \in [\underline{\alpha}, \overline{\alpha}]$. Thus, $\underline{\pi} = 0, \overline{\pi} = 1$.

Spying cost: Player i's cost of acquiring the signal is captured by $C_i(\alpha_i)$, and $C_i(\cdot)$

[12]For any alternative signals, the marginal cumulative distribution as a random variable is always uniformly distributed. For example, if player i's, $i = 1, 2$, signal realization s_i is drawn from $F(s|\theta_h, \alpha)$ ($F(s|\theta_l, \alpha)$) conditional on player j's ($j = 2, 1$, respectively) valuation being θ_h (θ_l), then we can always define an alternative signal with realizations: $\pi_i = p_h F(s_i|\theta_h, \alpha) + p_l F(s_i|\theta_l, \alpha)$, being the *probability integral transformation* of the original signal. Thus, π_i is the percentile function of s_i, which is always distributed uniformly in $[0, 1]$.

is assumed to be convex and increasing in α_i for $i = 1, 2$, with $C_i(\underline{\alpha}) = 0$. Let $MC_i(\alpha) = \frac{\partial C_i(\alpha)}{\partial \alpha} > 0$ be the marginal cost of acquiring the signal with accuracy α.

Posterior belief: Observing π_i leads player i to update her belief on θ_j according to Bayes' rule. Denote player i's posterior belief that player j has valuation θ_h upon receiving π_i by $\mu(\pi_i, \alpha_i)$, thus

$$\mu(\pi_i, \alpha_i) = \frac{p_h f_h(\pi_i, \alpha_i)}{p_h f_h(\pi_i, \alpha_i) + p_l f_l(\pi_i, \alpha_i)}$$

According to (1.2), $\mu(\pi_i, \alpha_i) = p_h f_h(\pi_i, \alpha_i)$ and $1 - \mu(\pi_i, \alpha_i) = p_l f_l(\pi_i, \alpha_i)$. For the rest of the paper, we assume:

Assumption 1. *Monotonic likelihood ratio property (MLRP): Given any* $\alpha \in [\underline{\alpha}, \overline{\alpha}]$, $\frac{f_h(\pi,\alpha)}{f_l(\pi,\alpha)}$ *is non-decreasing in* $\pi \in [\underline{\pi}, \overline{\pi}]$.

Assumption 1 implies $\mu(\pi_i, \alpha_i)$ is non-decreasing in π_i fixing α_i.

Information order: To rank signals by accuracy, we adopt the rotation order which was first introduced by Johnson and Myatt (2006), and was applied to auction settings by Shi (2012).[13]

Definition 1 (Rotation order). *The family of distributions* $F(\pi_i, \alpha_i)$ *is rotation-ordered if there exists a point* $\pi^+ \in [0, 1]$ *such that:* $\frac{\partial f_h(\pi_i,\alpha_i)}{\partial \alpha_i} \gtreqless 0$ *if* $\pi_i \gtreqless \pi^+$, *for all* $\alpha_i \in [\underline{\alpha}, \overline{\alpha}]$.

When α_i increases, $f_h(\pi_i, \alpha_i)$ rotates *counter clockwise* around π^+, which implies the updated belief $\mu(\pi_i, \alpha_i)$ becomes steeper. Meanwhile, $f_l(\pi_i, \alpha_i)$ rotates *clockwise* around π^+, which implies the updated belief $1 - \mu(\pi_i, \alpha_i)$ also becomes steeper.

Let us consider an example with two information acquisition choices, $\alpha' < \alpha''$. By definition, it must be true that $f_h(\pi, \alpha') > f_h(\pi, \alpha'')$ when $\pi < \pi^+$ and $f_h(\pi, \alpha') < f_h(\pi, \alpha'')$ when $\pi > \pi^+$; and that $f_l(\pi, \alpha') < f_l(\pi, \alpha'')$ when $\pi < \pi^+$ and $f_l(\pi, \alpha') > f_l(\pi, \alpha'')$ when $\pi > \pi^+$. This example is shown in the Figure 1.1 and 1.2.

When player i chooses $\alpha_i = \underline{\alpha}$, i.e. to acquire a completely uninformative signal about θ_j, then any realization of the signal does not convey any information about the opponent. In this case, $f_h(\pi_i, \underline{\alpha}) = f_l(\pi_i, \underline{\alpha}) = 1$ for all $\pi_i \in [\underline{\pi}, \overline{\pi}]$, see Figure 1.9. When player i chooses $\alpha_i = \overline{\alpha}$, i.e. to acquire a perfectly informative signal about θ_j, then each realization conveys perfect information about the opponent. In this case,

[13] See Ganuza and Penalva (2010) for thorough discussion on signal ordering by precision which is defined by dispersion of the distribution of posterior estimation.

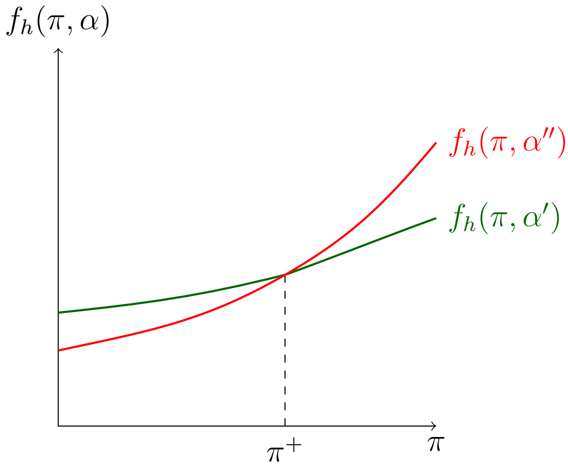

Figure 1.1: Increasing α means counter clockwise rotation of $f_h(\pi, \alpha)$

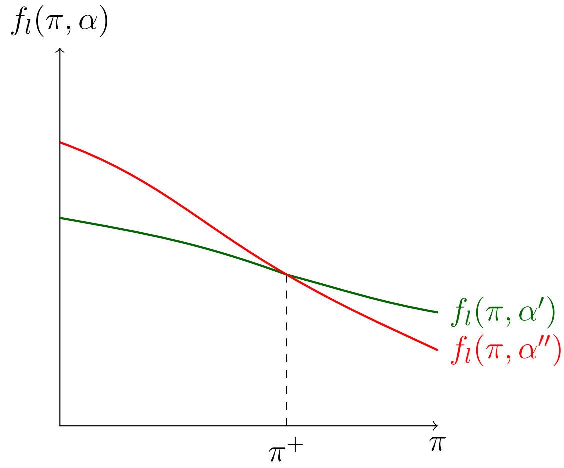

Figure 1.2: Increasing α means clockwise rotation of $f_l(\pi, \alpha)$

$f_h(\pi_i, \overline{\alpha}) = 0$ if $\pi_i \leqslant p_l$ and $f_h(\pi_i, \overline{\alpha}) = \frac{1}{p_h}$ if $\pi_i > p_l$, see Figure 1.4. Throughout the paper, we impose the following assumption:

Assumption 2. *All signals are rotation ordered around $\pi^+ = p_l$.*

$\pi^+ = p_l$ is the only rotation point such that all signals indexed from $\underline{\alpha}$ to $\overline{\alpha}$ are ordered.

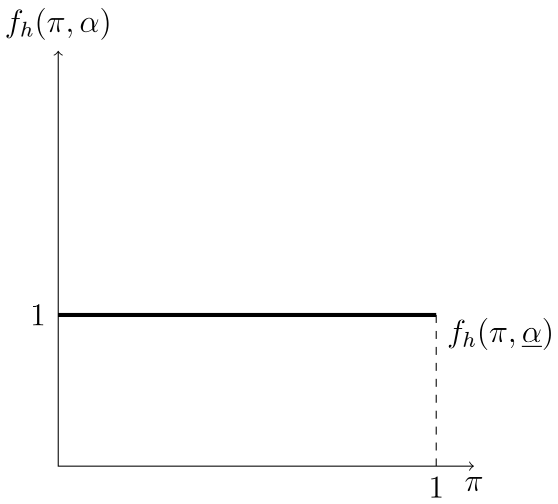

Figure 1.3: Completely uninformative signal $\underline{\alpha}$

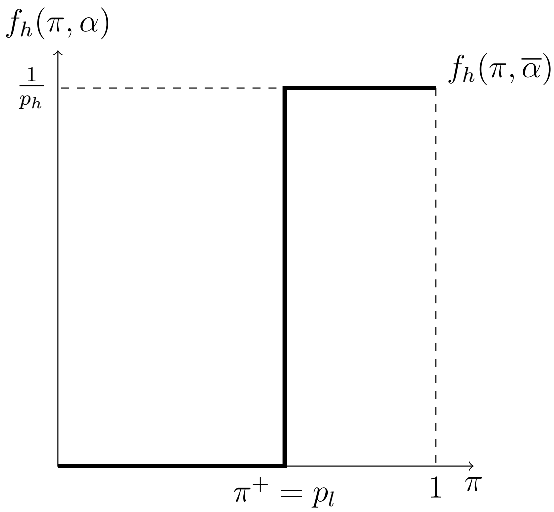

Figure 1.4: Perfectly informative signal $\overline{\alpha}$

Effort in the contest: Player i decides her effort after observing θ_i and π_i. Thus, the contest stage of the game is a Bayesian game with two-dimensional types, and the effort of player i is a two-to-one mapping:[14] $b : \{\theta_h, \theta_l\} \times [0, 1] \rightarrow R^+$.

[14]The two-to-one mapping strategy in the contest creates complications in analysis, especially in

Payoffs in the contest: Players choose their efforts in the contest simultaneously. The player who exerts higher effort wins the prize, whereas the losing player's effort is unrecoverable. Ties are broken with equal probabilities. Thus, player i with valuation θ_i exerting effort b_i earns a payoff:

$$U(b_i, b_j, \theta_i) = \begin{cases} -b_i, \text{ if } b_i < b_j \\ \theta_i - b_i, \text{ if } b_i > b_j \\ \frac{1}{2}\theta_i - b_i, \text{ if } b_i = b_j \end{cases}$$

A contest with the above payoff function, $U(b_i, b_j, \theta_i)$, is also known as a first price all-pay auction.[15]

Timing: The timing of spying in the contest game is shown in Figure 1.5. Firstly, player i chooses the accuracy α_i for the signal to be acquired on the opponent. Secondly, Nature determines the valuation profile according to the prior distribution and player i observes θ_i. Thirdly, according to θ_j and α_i, Nature determines a signal realization π_i observed by player i. Finally, player i chooses her effort b_i according to her private information (θ_i, π_i).

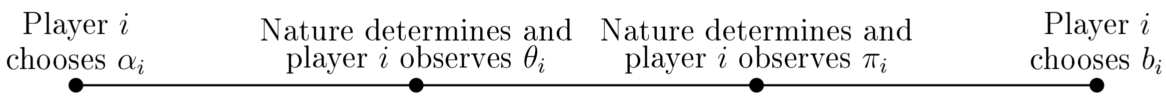

Figure 1.5: Timing of spying in the contest $(i = 1, 2)$

Social welfare and allocative efficiency: In the contest environment considered in this paper, efforts are wasted and thus, are not accounted as a part of social

characterizing the equilibrium strategy. In the previous studies involving two-to-one mapping in auctions, either the model was set up in a way that the auction has an equilibrium bidding strategy monotonically increasing with both dimensions (Tan, 2016), or the two-dimensional signal can be translated into a summary statistic which is positively correlated with the value with an assumption on the distributions of signals (Goeree and Offerman, 2003). Auctions with bi-dimensional types were considered in Fang and Morris (2006) where the authors studied a model of independent private value first-price auction in which players' valuations are drawn from a binary distribution, and each player receives a noisy two-valued signal about the opponent's valuation (with exogenous accuracy). The authors also characterized the equilibrium bidding strategy which is a two-to-one mapping. However, it is well known in the literature that the equilibrium strategy in auctions with multidimensional types are difficult to characterize, due to the fact that "monotonicity is not naturally defined" (Tan, 2016) or even non-existence of equilibrium (Jackson, 2009).

[15] Another type of model used in the contest literature is "Tullock contest" (Tullock, 1967) where player i's winning probability increases in her effort continuously instead of discontinuously as in $U(b_i, b_j, \theta_i)$ above. This type of model is sometimes called the imperfectly discriminating contest, whereas the all-pay auction is a perfectly discriminating contest in the sense that player i wins with probability one as long as her effort is larger than the opponent. See Konrad (2009) for survey of the literature.

welfare. Formally, social welfare is defined as the following:

Definition 2. *Social welfare is the total expected payoff of players:* $\Sigma_{i=1}^{2}\left[V_i(\alpha_i,\alpha_j)-C_i(\alpha_i)\right]$, *where* $V_i(\alpha_i,\alpha_j)$ *is player* i*'s equilibrium expected payoff in the contest where the profile of information acquisition choice is* (α_i,α_j).

In the context of competitive procurement, suppliers' total expected profit (net of spying cost) is the social welfare, as the effort spent in competing for the contract is unproductive. The allocative efficiency is formally defined as the following:

Definition 3. *An equilibrium of the contest is allocative efficient if the type* (θ_h, s) *player's effort is higher than type* (θ_l, t) *player's effort with probability one for any* $s,t\in[0,1]$.

Definition 3 follows from standard definition of allocative efficiency in auction literature. A competition for procurement contract is allocative efficient if suppliers with lower production cost always gets the contract.

1.3 Spying in contests

1.3.1 The contest with exogenous accuracy

In this section we first study a simplified model where both players exogenously receive a free, noisy and private *spying* signal about the opponent's valuation with the same accuracy, i.e. $\alpha_1=\alpha_2=\alpha\in[\underline{\alpha},\overline{\alpha}]$. The accuracy α is common knowledge.

Equilibrium effort in the contest

Denoted by $b_l(\pi,\alpha)$ and $b_h(\pi,\alpha)$ the effort strategy of types (θ_l,π) and (θ_h,π), respectively.

Lemma 1. *Given any* $\alpha\in[\underline{\alpha},\overline{\alpha}]$, *then in any symmetric, allocative efficient, pure strategy equilibrium of the contest, the following must be true:*

1. *Monotonicity: the type* (θ_h,π) *of player* i*'s effort is non-decreasing in* π *and the type* (θ_l,π) *of player* i*'s effort is non-increasing in* π;

2. *Continuity: both players' strategies are continuous without any atom;*

3. *Initial conditions:* $b_l(1,\alpha)=0$ *and* $b_l(0,\alpha)=b_h(0,\alpha)$.

Lemma 1 states that the efforts are monotonic in π fixing valuations in any allocative efficient equilibrium of the contest. Lemma 2 below provides the necessary condition for existence of any symmetric allocative efficient, pure strategy equilibrium.[16]

Lemma 2 (Efficiency)**.** *Given that $\alpha \in [\underline{\alpha}, \overline{\alpha}]$, there exists a symmetric, allocative efficient, pure strategy equilibrium in the contest only if $\frac{f_h(\pi,\alpha)}{f_l(\pi,\alpha)} \geqslant \frac{\theta_l}{\theta_h}$ for all $\pi \in [0,1]$.*

Note that the necessary condition given in Lemma 2 imposes a lower bound on the likelihood ratio. Thus, this is a restriction on the accuracy of the signal.

Definition 4. *Denoted by $\widehat{\alpha}$ the highest possible accuracy of a signal which satisfies $\frac{f_h(\pi,\alpha)}{f_l(\pi,\alpha)} \geqslant \frac{\theta_l}{\theta_h}$ for all $\pi \in [0,1]$.*

Thus, Lemma 2 indicates that any symmetric, pure strategy equilibrium must have $\alpha \in [\underline{\alpha}, \widehat{\alpha}]$.[17]. To understand the necessary condition in Lemma 2, recall the monotonicity property given by Lemma 1. When the type (θ_h, π_i) of player i chooses her equilibrium effort $b_h(\pi_i, \alpha)$, then she *wins* against a high valuation type opponent with probability $\int_0^{\pi_i} f_h(t,\alpha)dt$. If she instead chooses type (θ_l, π_i)'s equilibrium effort, $b_l(\pi_i, \alpha)$, then she *loses* to a low valuation opponent with probability $\int_0^{\pi_i} f_h(t,\alpha)dt$. Therefore, an increase of effort from $b_l(\pi_i, \alpha)$ to $b_h(\pi_i, \alpha)$ earns type (θ_h, π_i) of player i a gain of

$$\theta_h \left[\mu(\pi_i, \alpha) \int_0^{\pi_i} f_h(t,\alpha)dt + [1 - \mu(\pi_i, \alpha)] \int_0^{\pi_i} f_h(t,\alpha)dt \right] = \theta_h \int_0^{\pi_i} f_h(t,\alpha)dt$$

Similarly, the gain for type (θ_l, π_i) from the same increase of effort is $\theta_l \int_0^{\pi_i} f_l(t,\alpha)dt$. Allocative efficiency requires the gain of increasing the effort for type (θ_h, π_i) outweighs the cost, whereas the cost of increasing the effort outweighs the gain for type (θ_l, π_i). Since the cost are the same across the two types, this necessary condition is equivalent of $\theta_h \int_0^{\pi_i} f_h(t,\alpha)dt \geqslant \theta_l \int_0^{\pi_i} f_l(t,\alpha)dt$ for any $\pi_i \in [0,1]$, which implies $\theta_h f_h(\pi,\alpha) \geqslant \theta_l f_l(\pi,\alpha)$. Thus, we have the condition given in Lemma 2.

In light of Lemma 1 and 2, we now derive the equilibrium strategy for the symmetric, pure strategy equilibrium with efficient allocation, assuming the condition given in Lemma 2 is satisfied. The expected payoffs of types (θ_l, π) and (θ_h, π) of player i when choosing b, given that player j plays the strategy $b_h(\pi, \alpha)$ and $b_l(\pi, \alpha)$,

[16] In fact, the proof of Proposition 1 in the appendix shows that this necessary condition is also sufficient for efficiency.

[17] By MLRP and rotation order, for all $\alpha \in [\underline{\alpha}, \widehat{\alpha}]$ the condition in Lemma 2 is satisfied for all $\pi \in [0,1]$, and for all $\alpha \in (\widehat{\alpha}, \overline{\alpha}]$ the condition is not satisfied for at least $\pi = 0$.

is given by:

$$U(b|\theta_l,\pi) = \theta_l\left[1-\mu(\pi,\alpha)\right]\int_{b_l^{-1}(b,\alpha)}^{1} f_l(\Pi,\alpha)d\Pi - b \tag{1.3}$$

$$U(b|\theta_h,\pi) = \theta_h\left[\left[1-\mu(\pi,\alpha)\right]+\mu(\pi,\alpha)\int_{0}^{b_h^{-1}(b,\alpha)} f_h(\Pi,\alpha)d\Pi\right] - b \tag{1.4}$$

where $b_l^{-1}(b,\alpha)$ and $b_h^{-1}(b,\alpha)$ are the inverse of effort strategies by player j. Take the first order derivative w.r.t. b on both (1.3) and (1.4):

$$\begin{aligned}\frac{\partial b_l(\pi,\alpha)}{\partial\pi} &= -\left[1-\mu(\pi,\alpha)\right]f_l(\pi,\alpha)\theta_l\\ \frac{\partial b_h(\pi,\alpha)}{\partial\pi} &= \mu(\pi,\alpha)f_h(\pi,\alpha)\theta_h\end{aligned}$$

By the initial conditions given in Lemma 1, $b_l(1,\alpha)=0$ and $b_h(0,\alpha)=b_l(0,\alpha)$, the pure strategy equilibrium with efficient allocation is derived and given in Proposition 1. Furthermore, Proposition 1 also provides the equilibrium strategy when the allocation is not efficient, i.e. when the condition given in Lemma 2 is not satisfied.

Proposition 1. *Suppose Assumption 1 is satisfied.*

- *If $\alpha\in[\underline{\alpha},\widehat{\alpha}]$, then there exists a unique pure strategy, symmetric, allocative efficient equilibrium:*

$$\begin{aligned}b_l(\pi,\alpha) &= \theta_l\int_{\pi}^{1}\left[1-\mu\left(\Pi,\alpha\right)\right]dF_l(\Pi,\alpha)\\ b_h(\pi,\alpha) &= \theta_h\int_{0}^{\pi}\mu\left(\Pi,\alpha\right)dF_h(\Pi,\alpha)+\theta_l\int_{0}^{1}\left[1-\mu\left(\Pi,\alpha\right)\right]dF_l(\Pi,\alpha)\end{aligned}$$

- *If $\alpha\in[\widehat{\alpha},\overline{\alpha}]$, then there exists a unique, symmetric equilibrium in which the type (θ_l,π) and (θ_h,π) with $\pi>\pi^*$ play pure strategy:*

$$\begin{aligned}b_l(\pi,\alpha) &= \theta_l\int_{\pi}^{1}\left[1-\mu\left(\Pi,\alpha\right)\right]dF_l(\Pi,\alpha)\\ b_h(\pi,\alpha) &= \theta_h\int_{\pi^*}^{\pi}\mu\left(\Pi,\alpha\right)dF_h(\Pi,\alpha)\\ &+ \theta_l\int_{\pi^*}^{1}\left[1-\mu\left(\Pi,\alpha\right)\right]dF_l(\Pi,\alpha)+\frac{\theta_h\theta_l}{p_h\theta_l+p_l\theta_h}\pi^*;\end{aligned}$$

  and type (θ_l,π) and (θ_h,π) with $\pi\leqslant\pi^*$ mix over $[b_l(\pi^*,\alpha),b_h(\pi^*,\alpha)]$ according *to CDF* $\sigma_l(b|\pi,\alpha)$ *and* $\sigma_h(b|\pi,\alpha)$ *respectively:*

$$\begin{aligned}\sigma_l(b|\pi,\alpha) &= \sigma_h(b|\pi,\alpha)\\ &= \frac{p_h\theta_l + p_l\theta_h}{\theta_h\theta_l\pi^*}\left(b - \theta_l\int_{\pi^*}^{1}\left[1-\mu\left(\Pi,\alpha\right)\right]dF_l(\Pi,\alpha)\right)\end{aligned}$$

where π^* *is given by*

$$\theta_l\int_0^{\pi^*} f_l(\Pi,\alpha)d\Pi = \theta_h\int_0^{\pi^*} f_h(\Pi,\alpha)d\Pi \tag{1.5}$$

See Figure 1.6 for the allocative efficient equilibrium and Figure 1.7 for the allocative inefficient equilibrium. The intuition of the efficient equilibrium can be understood in competitive procurement process: a supplier would only spend additional effort in preparing the proposal or bribing the procurement agent after learning that the opponent is likely to be equally competitive, e.g., with similar production cost.

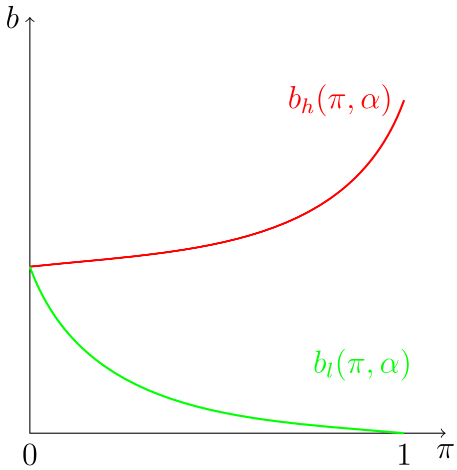

Figure 1.6: Allocative efficient equilibrium

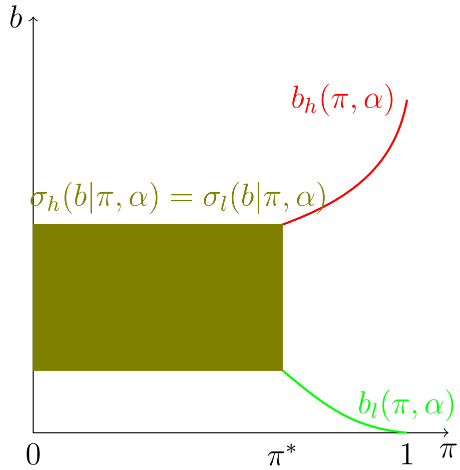

Figure 1.7: Allocative inefficient equilibrium

To gain some intuition on the equilibrium effort strategy, we rewrite type (θ_l,π) of player i's strategy as:

$$\theta_l\int_{\pi_i}^{1}\left[1-\mu(\Pi,\alpha)\right]dF_l(\Pi,\alpha) = \theta_l E\left[\left(1-\mu(\pi_j,\alpha)\right)|\pi_i \leqslant \pi_j, \theta_i = \theta_l\right] \tag{1.6}$$

Equation (1.6) suggests the type (θ_l,π_i)'s effort is her valuation times her expectation of the opponent's posterior belief that she has low valuation, conditional on she indeed has low valuation and she wins (i.e. when $\pi_i \leqslant \pi_j$).

Similarly, the part of a type (θ_h, π_i)'s effort on top of $b_l(0)$ can be rewritten in the same manner:

$$\theta_h \int_0^{\pi_i} \mu(\Pi, \alpha) dF_h(\Pi, \alpha) = \theta_h E\left[\mu(\pi_j, \alpha) | \pi_i \geqslant \pi_j, \theta_i = \theta_h\right]$$

This part of the player i's effort is θ_h times her expectation of the opponent's posterior belief that she has high valuation, conditional on she indeed has high valuation and she wins (i.e. when $\pi_i \geqslant \pi_j$).

Two observations worth mentioning.

Observation 1. *Player i's effort is first order stochastically increasing in her valuation.*

In the allocative efficient equilibrium, the high valuation player's effort is strictly higher than the low valuation player's effort. In the allocative inefficient equilibrium, the former is higher than the latter in the sense of first order stochastic dominance, as there is an interval over which players with both valuations randomize. In the inefficient equilibrium, both players' signals are sufficiently informative. The relatively lower signal realizations ($\pi < \pi^*$) credibly reveal that the opponent is a low valuation type, thus both (θ_h, π) and (θ_l, π) with $\pi < \pi^*$ are confident that the opponent has θ_l. Furthermore, player i with valuation θ_l realizes that the opponent is very likely to receive those lower signal realizations. Therefore, there cannot be an equilibrium in which type (θ_l, π) with $\pi < \pi^*$ plays pure strategy, as then she loses almost for sure and is better off to deviate to choose zero effort. This indicates both types (θ_h, π) and (θ_l, π) of the opponent with $\pi < \pi^*$ plays mixed strategy as well.[18]

In fact, type (π^*, α) is indifferent between the effort levels $b_h(\pi^*, \alpha)$ and $b_l(\pi^*, \alpha)$, by the definition of π^*. This is consistent with the idea behind Lemma 2: only type (θ_l, π) (type (θ_h, π)) with $\pi > \pi^*$ finds the cost (gain) of increasing the effort from $b_l(\pi, \alpha)$ to $b_h(\pi, \alpha)$ outweighs the gain (cost), which is why they play pure strategy with (θ_h, π)'s effort higher than (θ_l, π). For types with $\pi < \pi^*$ the condition for efficiency is no long satisfied and thus, (θ_l, π) and (θ_h, π) play mixed strategy in a common interval, as discussed above.

[18]The all-pay auction with complete information has only mixed strategy equilibrium (Baye et al., 1996) for the same reason. To see why, suppose player i with θ_h plays pure strategy b in equilibrium against the opponent with θ_l. Then the opponent would either bid slightly above b, e.g., $b + \epsilon$ – if b is lower than θ_l – or bid zero – if b is no less than θ_l. In the former, player i would find it profitable to bid slightly higher instead, say $b + 2\epsilon$, suggesting b is suboptimal. In the latter, player i would find it profitable to bid ϵ which, again, suggests b is not optimal. Therefore, the two players with different valuations play mixed strategies in a common interval suggesting a positive probability that the low valuation player wins. Hence the allocation in equilibrium is inefficient.

Observation 2. *High valuation type of player i's effort is non-decreasing in π; low valuation type of player i's effort is non-increasing in π.*

The belief that the opponent is a high valuation player – induced by higher realizations of the signal – encourages the high valuation type of player i to compete aggressively to increase the odds of winning, and discourages the low valuation type of player i to compete conservatively to save the cost of competition. In other words, players compete aggressively when there is higher posterior probability that the opponent has the same valuation, i.e. "motivation effect", and compete conservatively when there is lower posterior probability that the opponent has the different valuation, i.e. "demotivation effect". This observation is also true for the allocative inefficient equilibrium with the same intuition.

Proposition 2 below suggests players are strictly better off in the equilibrium given in Proposition 1 with accuracy $\alpha \in (\underline{\alpha}, \overline{\alpha})$. Denoted by $V_i(\alpha, \alpha)$ player i's expected payoff when both players endogenously receive a free signal about the opponent with accuracy α.

Proposition 2. *In the equilibrium given by Proposition 1:*

- *Player i's expected payoff is higher when $\alpha \in (\underline{\alpha}, \overline{\alpha})$ than when $\alpha = \underline{\alpha}$: $V_i(\alpha, \alpha) > V_i(\underline{\alpha}, \underline{\alpha})$.*
- *Furthermore, the total expected effort is strictly lower when $\alpha \in (\underline{\alpha}, \overline{\alpha})$ than when $\alpha = \underline{\alpha}$.*

The proof of the proposition is similar to that in Proposition 3 in the next section and thus, is omitted. The intuition is that the "motivation" and "demotivation" effects increase each player's marginal return of effort. In particular, a player only increases her effort when it is worthwhile to do so – when the opponent is more likely to have the same valuation as she does. In other words, given the same amount of additional effort, the additional gain of probability to win is higher when players receive the signal about the opponent.

Some features of the equilibrium

Corollary 1 suggests it is never a good news that the opponent is more likely to have high value, no matter what the player's valuation is. Denoted by $V(\theta_l, \pi, \alpha)$ and $V(\theta_h, \pi, \alpha)$ the equilibrium expected payoff for types (θ_l, π) and (θ_h, π), respectively.

Corollary 1. *In the equilibrium given by Proposition 1, the following must be true:*

(i) For all $\alpha \in [\underline{\alpha}, \overline{\alpha}]$, both the high and the low valuation types' expected payoffs are non-increasing in π:

$$\begin{aligned}
\frac{\partial V(\theta_l, \pi, \alpha)}{\partial \pi} &= p_l \theta_l \frac{\partial f_l(\pi, \alpha)}{\partial \pi} \int_{\pi}^{1} f_l(\Pi, \alpha) d\Pi \leqslant 0 \\
\frac{\partial V(\theta_h, \pi, \alpha)}{\partial \pi} &= -p_h \theta_h \frac{\partial f_h(\pi, \alpha)}{\partial \pi} \int_{\pi}^{1} f_h(\Pi, \alpha) d\Pi \leqslant 0
\end{aligned}$$

where the equalities are only satisfied when $\pi = 1$.

(ii) The pure strategies are weakly convex in the signal π:

$$\begin{aligned}
\frac{\partial^2 b_l(\pi, \alpha)}{\partial \pi^2} &= -2 p_l \theta_l f_l(\pi, \alpha) \frac{\partial f_l(\pi, \alpha)}{\partial \pi} \geqslant 0 \\
\frac{\partial^2 b_h(\pi, \alpha)}{\partial \pi^2} &= 2 p_h \theta_h f_h(\pi, \alpha) \frac{\partial f_h(\pi, \alpha)}{\partial \pi} \geqslant 0
\end{aligned}$$

and the mixed strategies are independent of π.

When both players' accuracies of signal is increased, i.e. α is increased, the equilibrium effort strategies are more sensitive to a marginal change of π, see Corollary 2.

Corollary 2 (Sensitivity). *When the signal becomes more informative in rotation order, then the slopes of $b_h(\pi, \alpha)$ and $b_l(\pi, \alpha)$ are decreased for $\pi < \pi^+$, and are increased for $\pi > \pi^+$. Furthermore, $\frac{\partial \pi^*}{\partial \alpha} > 0$.*

See Figure 1.8 for this result. For realizations lower (higher) than the rotation point, $b_h(\pi, \alpha)$ becomes flatter (steeper) while $b_l(\pi, \alpha)$ becomes steeper (flatter). Intuitively, when the signal becomes more informative, the high valuation player would not increase her effort as much as before in response to a marginal increase of π in the interval $[0, \pi^+)$, as this interval more credibly indicates that the opponent has low valuation. However, she would increase her effort more than before in response to a marginal increase of π in the interval $(\pi^+, 1]$, as this interval more credibly indicates that the opponent has high valuation. The same intuition can be applied to explain the change in sensitivity of $b_l(\pi, \alpha)$. When there is a marginal decrease of π in the interval $[0, \pi^+)$ ($(\pi^+, 1]$), the low valuation player would increase her effort more than before, as this interval more credibly indicates that the opponent has low (high) valuation.

Corollary 3 shows that the equilibrium given in Proposition 1 replicates the mixed strategy equilibrium of all-pay auction with independent private value when $\alpha = \underline{\alpha}$.

Corollary 3. *When $\alpha = \underline{\alpha}$ in the equilibrium given in Proposition 1:*

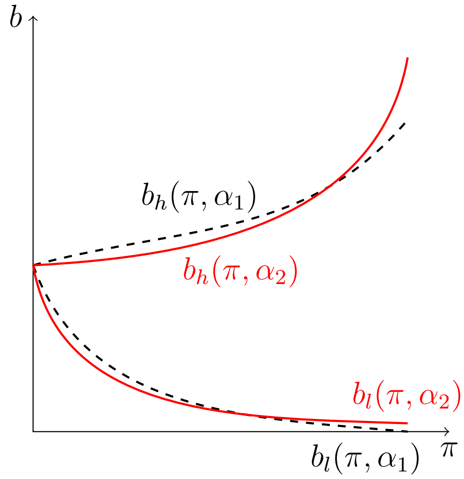

Figure 1.8: Rotation and sensitivity: rotation from α_1 to α_2 decreases the slope of effort strategies for $\pi < \pi^+$ and increases the slope for $\pi > \pi^+$

- *The ex ante distribution of equilibrium effort of player i with valuation θ_h is uniform in the interval* $[p_l\theta_l, p_l\theta_l + p_h\theta_h]$*;*
- *The ex ante distribution of equilibrium effort of player i with valuation θ_l is uniform in the interval* $[0, p_l\theta_l]$.

See Figure 1.9 for the equivalence of equilibrium in all-pay auction with IPV and the equilibrium in Proposition 1 when $\alpha = \underline{\alpha}$. The equilibrium with $\alpha = \underline{\alpha}$ simply purifies the mixed strategy equilibrium in the IPV setting. In particular, $b_h(\pi, \underline{\alpha})$ purifies $G_h(b)$ and $b_l(\pi, \underline{\alpha})$ purifies $G_l(b)$.

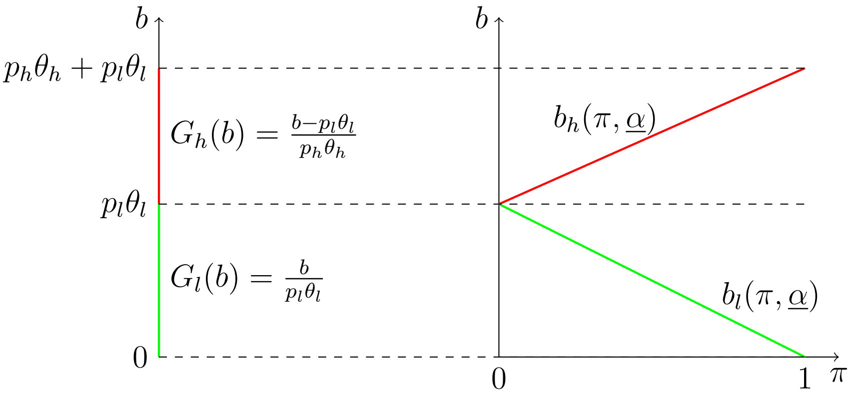

Figure 1.9: $G_h(b)$ and $G_l(b)$ are CDFs of the mixed strategy by each player with the high and the low valuation respectively in the equilibrium of all-pay auction with IPV.

Corollary 4 states that the equilibrium strategy with $\alpha = \overline{\alpha}$ in Proposition 1 replicates the equilibrium of the complete information all-pay auction .

Corollary 4. *When $\alpha = \overline{\alpha}$ in the equilibrium given in Proposition 1:*

- *The type (θ_l, π) of player i with $\pi > \pi^*$ chooses zero with certainty;*
- *The type (θ_l, π) and (θ_h, π) of player i with $\pi < \pi^*$ mixing over $[0, \theta_l]$ uniformly;*
- *The type (θ_h, π) of player i with $\pi > \pi^*$ plays pure strategy $b_h(\pi, \overline{\alpha})$ which is linear in π and increases from θ_l to θ_h.*

See Figure 1.10 and 1.11 for the equivalence of the equilibrium in all-pay auction with IPV and the equilibrium in Proposition 1 with $\alpha = \overline{\alpha}$. In particular, Figure 1.10 corresponds to the case when player i has θ_l and Figure 1.11 to θ_h. For example, when player i with θ_l competes against an opponent with the same valuation in the complete information setting (which occurs with probability p_l), her equilibrium mixed strategy is given by $G_{ll}(b)$ in Figure 1.10. This part of the equilibrium is replicated by $\sigma_l(b|\pi, \overline{\alpha})$ when $\pi < \pi^+$ (which occurs with probability p_l as well).

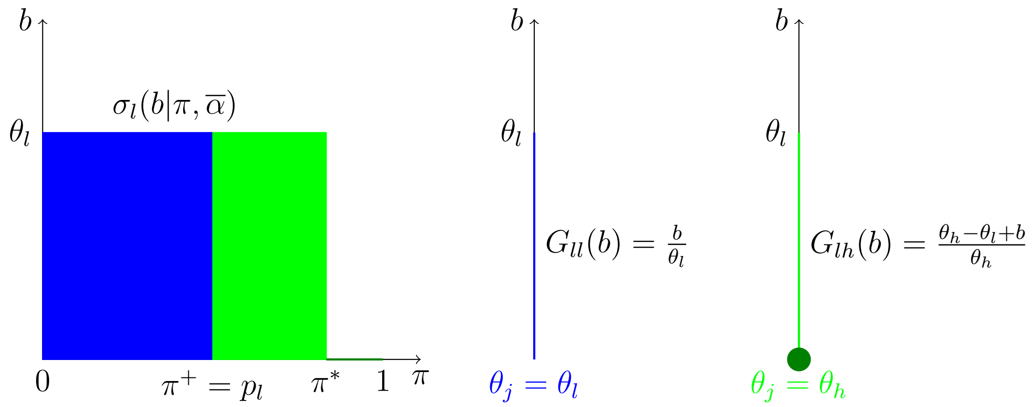

Figure 1.10: $G_{ll}(b)$ and $G_{lh}(b)$ are the CDFs of mixed strategy by a player with θ_l who encounters an opponent with θ_l and θ_h respectively in complete information all-pay auction.

1.3.2 Information acquisition – spying

In this section we characterize the symmetric equilibrium of spying, i.e. the choices of α_1 and α_2. First, define the marginal expected payoff as the following:

Definition 5. *Let $AMR(\eta, \alpha)$ be player i's marginal expected payoff from the contest when she chooses $\alpha_i = \eta$ while player j chooses $\alpha_j = \alpha$ and (wrongly) believes that player i has chosen the same.*

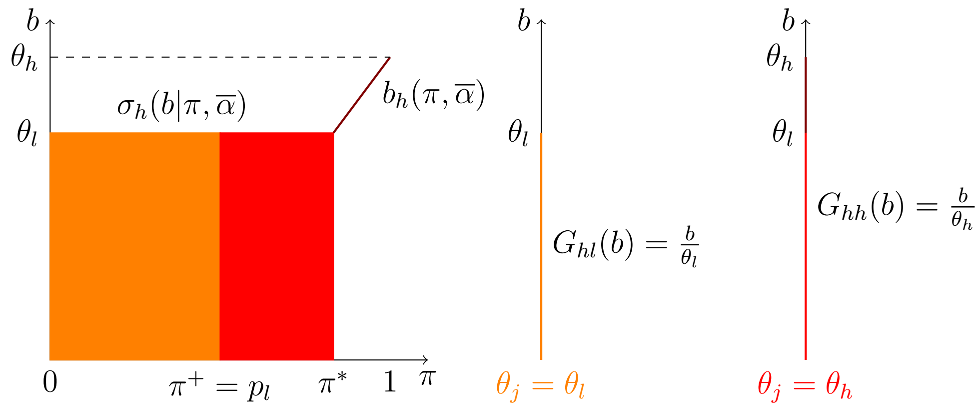

Figure 1.11: $G_{hl}(b)$ and $G_{hh}(b)$ are the CDFs of mixed strategy by a player with θ_h who encounters an opponent with θ_l and θ_h respectively in complete information all-pay auction.

Suppose player j chooses α and believes that player i has also chosen α, then the distribution of his effort will be exactly the same as in the symmetric equilibrium given by Proposition 1. Thus, for any η that player i may choose, her choice of effort is a decision problem instead of a strategic one. Increasing η would improve player i's estimation of player j's effort distribution and revise her decision accordingly, and $AMR(\eta, \alpha)$ is the marginal increase of her expected payoff due to a marginal increase of η.

Lemma 3. *$AMR(\eta, \alpha) > 0$ for all $\eta \leqslant \alpha$ and $\eta, \alpha \in [\underline{\alpha}, \overline{\alpha})$.*

Lemma 3 states that player i can always increase expected payoff in the contest by increasing η, as long as it is no larger than α.[19] It then follows that increasing η always gives player i a better estimation of player j's valuation and thus, a clearer idea of his effort distribution. However, suppose player j chooses $\overline{\alpha}$ and expects player i to do the same, then player i has no incentive to increase η.

Lemma 4. *$AMR(\eta, \overline{\alpha}) = 0$ for all $\eta \in [\underline{\alpha}, \overline{\alpha}]$.*

Let $MR(\alpha) = AMR(\alpha, \alpha)$ be the marginal expected payoff of player i through increasing η when both players have chosen α. When the marginal cost function $MC(\cdot)$ crosses with $MR(\cdot)$ from below, then there must exist an interior solution at $\eta = \alpha$ to the problem $max_\eta V(\eta, \alpha) - C(\eta)$, where $V(\eta, \alpha)$ is player i's maximum expected payoff from the contest when player j plays the symmetric equilibrium strategy believing that both players have chosen α.

[19] When $\eta > \alpha$, $AMR(\eta, \alpha)$ may still be positive, but it is irrelevant to the later results.

Proposition 3 (Information acquisition)**.** *Given any convex spying cost function $C(\alpha)$ with $C(\underline{\alpha}) = 0$, if there exists a symmetric equilibrium (α^*, α^*) where $\alpha^* \in (\underline{\alpha}, \overline{\alpha})$, i.e.*

$$MR(\alpha^*) = MC(\alpha^*),$$

then each player is strictly better off than not spying at all:

$$V(\alpha^*, \alpha^*) - C(\alpha^*) \geqslant V(\underline{\alpha}, \underline{\alpha})$$

Players are better off by spying on each other, even taking the cost of spying into account. To see the intuition, recall that in any Bayesian Nash equilibrium all players correctly anticipate the equilibrium strategy played by each type of other players. Thus, when players acquire information about the opponent's type, they in effect acquire information about the strategy that the opponent is likely to play. This means the behavior of information acquisition has a similar impact as collusion: Players anticipate their opponents' strategies on expectation and aware that their own strategies are anticipated by the opponents. The additional rent players extract through such coordination of strategies is high enough that the cost of spying is compensated for.

Proposition 3 also shows that the symmetric equilibrium choice of information acquisition α^* is within $(\underline{\alpha}, \overline{\alpha})$, i.e. between the incomplete information and complete information setting, as long as the cost function is convex enough. In fact, such information acquisition increases the social welfare to the level higher than complete information setting.

Corollary 5. *In any symmetric equilibrium of the information acquisition game as specified in Proposition 3, it must be true that $V(\alpha^*, \alpha^*) - C(\alpha^*) \geqslant V(\overline{\alpha}, \overline{\alpha})$, where the equality is satisfied when either $\alpha^* = \overline{\alpha}$ or $\alpha^* = \underline{\alpha}$.*

Corollary 5 follows directly from Proposition 3, as it had been confirmed in the literature that $V(\overline{\alpha}, \overline{\alpha}) = V(\underline{\alpha}, \underline{\alpha})$ (Kovenock et al., 2015). In the complete information setting, the loss of efficiency cancels the gain from coordination. However, when each player acquires a noisy signal about the opponent, the gain from coordination outweighs the loss of efficiency. An interesting implication of Corollary 5 is that social welfare can only be improved by spying when it is costly to do so.

Corollary 6. *When both players acquire a partially informative signal, the total expected effort they exert in the contest is strictly lower than when they do not spy on each other.*

Corollary 6 follows directly from Proposition 3. The total surplus in the contest is the sum of effort and players' expected payoffs. And such a surplus is equals to $p_l^2\theta_l+(1-p_l^2)\theta_h$ in the efficient equilibrium, and is lower in the inefficient equilibrium. The fact that both players are better off after spying on each other implies they cut a larger share of the surplus than when they don't spy at all.

1.4 Information disclosure in contests

Given the previous results that spying activities improves players' welfare, a question naturally arises: do players have incentives to disclose their private information to each other? According to Proposition 3, players are strictly better off when they both acquire a partially informative and costly signal about the opponents. Suppose players simply disclose such partially informative signals to each other, then players' welfare should be even higher than when they spy on each other because they can now avoid the cost of spying.

The model in Section 1.2 can be modified to study information disclosure in contests where each player commits to disclose a noisy signal about her own valuation to the opponent. Contrary to the spying situation where player i chooses α_i, i.e. the accuracy of the signal regarding the opponent's valuation θ_j, in the information disclosure situation player i chooses the accuracy α_j of the signal regarding her own valuation θ_i. Players can observe the accuracies of both signals. However, player i does not observe any realization of the signal she discloses – only her opponent does. A supplier competing for a procurement contract may allow the opponent to conduct some independent research/investigation on its production process, materials used in the proposal, or other information relevant to the bidding process. The supplier i controls the materials to be shared with the opponent, i.e. α_j, but the results of the investigation is not observable to supplier i, i.e. π_j.[20]

The timing of the information disclosure game is given in the Figure 1.12. In describing the timing, let $i = 1, 2$ and $j = 2, 1$. Firstly, player j chooses the accuracy α_i for the signal to be received by his opponent, player i. Secondly, Nature determines valuation profile according to the prior distribution and both players observes their own valuations. Thirdly, Nature determines signal realization π_i according to θ_j and α_i and player i observes it. Finally, player i chooses effort b_i based on her private

[20]In the patent race example, companies may disclose information by allowing each other to investigate internal research materials, chemicals, run some experiments, or provide a prototype of their product. The company who discloses the information decides what kind of investigation the opponent can conduct, yet the result of investigation is not observable to it. In the lobbying example, lobbying firms representing different companies may disclose the CVs or backgrounds of their lobbyists to each other, yet they don't know how the opponent will interpret these information.

information (θ_i, π_i). Note that the difference between information disclosure and spying is only in that the accuracy of the signal regarding the opponent is chosen by the opponent instead of the player.

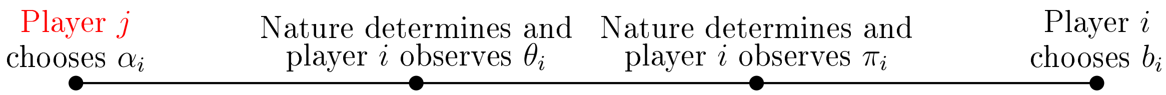

Figure 1.12: Timing of information disclosure in the contest $(i = 1, 2; j = 2, 1)$

Since the accuracy profile (α_i, α_j) is common knowledge, the first step of solving the information disclosure game is to find the equilibrium effort when players exogenously receive signals with different accuracies for free. In other words, an "asymmetric version" of Proposition 1.

1.4.1 The contest with exogenous asymmetric accuracy

Suppose each player receives a signal about the opponent, and the exogenously given accuracy of player 1's signal is α_1 whereas that of player 2's signal is α_2, where $\alpha_1 \neq \alpha_2$, i.e. the accuracy is asymmetric. Denoted by $b_{il}(\pi, \alpha_i, \alpha_j)$ and $b_{ih}(\pi, \alpha_i, \alpha_j)$ the efforts of type (θ_l, π) and (θ_h, π) respectively, when player i's accuracy is α_i and the opponent's accuracy is α_j. Similar to the symmetric α case, the monotonicity of player i's effort in equilibrium depends on her valuation, see Lemma 5.

Lemma 5. *Suppose $\alpha_1 \neq \alpha_2$, then in any allocative efficient, pure strategy equilibrium of the contest, the following must be true for both players:*

1. *Monotonicity: the type (θ_h, π) of player i's effort is non-decreasing in π and the type (θ_l, π) of player i's effort is non-increasing in π;*

2. *Continuity: both players' strategies are continuous without any atom;*

3. *Common support: $b_{1l}(1, \alpha_1, \alpha_2) = b_{2l}(1, \alpha_2, \alpha_1)$ and $b_{1h}(1, \alpha_1, \alpha_2) = b_{2h}(1, \alpha_2, \alpha_1)$;*

4. *Initial conditions: $b_{1l}(1, \alpha_1, \alpha_2) = b_{2l}(1, \alpha_2, \alpha_1) = 0$ and $b_{1l}(0, \alpha_1, \alpha_2) = b_{1h}(0, \alpha_1, \alpha_2) = b_{2l}(0, \alpha_2, \alpha_1) = b_{2h}(0, \alpha_2, \alpha_1)$.*

Part 1 of Lemma 5 implies type $(\theta_l, 1)$ of both players chooses the lowest effort, 0, and type $(\theta_h, 1)$ of both players choose the highest effort, in the allocative efficient, pure strategy equilibrium. Part 3 of Lemma 5 indicates the upper bound of each valuation of player must be the same. Part 4 is useful later in solving the equilibrium effort and in proving the uniqueness of such an equilibrium.

Lemma 6 below provides the necessary conditions for any symmetric, pure strategy and allocative efficient equilibrium to exist.

Lemma 6 (Efficiency)**.** *Suppose $\alpha_1 \neq \alpha_2$, then there exists a symmetric, pure strategy, allocative efficient equilibrium in the contest only if $\frac{f_h(\pi,\alpha_i)}{f_l(\pi,\alpha_i)} \geqslant \frac{\theta_l}{\theta_h}$ for all $\pi \in [0,1]$ and $i = 1, 2$.*

In fact, this condition is also the sufficient condition for the existence of such an equilibrium, as shown in the proof of Proposition 4. This contest is ex ante asymmetric in the sense that players' accuracy of signals about the opponent are different. However, this does not mean that the equilibrium effort strategy must be asymmetric.

For simplicity, we restrict attention to the symmetric, pure strategy and allocative efficient equilibrium, thus we make the following assumption:

Assumption 3.

$$\frac{f_h(\pi,\alpha_i)}{f_l(\pi,\alpha_i)} \geqslant \frac{\theta_l}{\theta_h}, \text{ for all } \pi \in [0,1] \text{ and } i = 1, 2.$$

Assumption 3 restricts the accuracy of both players' signal for the similar reason as the symmetric case, i.e. $\alpha_1 = \alpha_2$.

We now derive the equilibrium of the contest given $\alpha_1 \neq \alpha_2$. Denote $b_{il}^{-1}(b, \alpha_i, \alpha_j)$ and $b_{ih}^{-1}(b, \alpha_i, \alpha_j)$ the inverse effort strategy of the low and the high valuation types of player i. According to Lemma 5, the expected payoff for type (θ_l, π) and (θ_h, π) of player j when choosing an effort b can be written as $U_j(b|\theta_l, \pi)$ and $U_j(b|\theta_h, \pi)$ respectively:

$$\begin{aligned} U_j(b|\theta_l,\pi) &= \theta_l[1-\mu(\pi,\alpha_j)] \int_{b_{il}^{-1}(b,\alpha_i,\alpha_j)}^{1} f_l(\Pi,\alpha_i)d\Pi - b \\ U_j(b|\theta_h,\pi) &= \theta_h\left[(1-\mu(\pi,\alpha_j)) + \mu(\pi,\alpha_j)\int_0^{b_{ih}^{-1}(b,\alpha_i,\alpha_j)} f_h(\Pi,\alpha_i)d\Pi\right] - b \end{aligned}$$

By the first order approach and the initial conditions provided in part 4 of Lemma 5, the equilibrium strategy in the contest can be derived, as shown in Proposition 4 following.

Proposition 4. *If Assumption 3 is satisfied, then the unique pure strategy, allocative efficient equilibrium of the contest is given by:*

$$\begin{aligned} b_{il}(\pi,\alpha_i,\alpha_j) &= \theta_l \int_\pi^1 \left[1-\mu\left(\Pi,\alpha_i\right)\right] dF_l(\Pi,\alpha_j) \\ b_{ih}(\pi,\alpha_i,\alpha_j) &= \theta_h \int_0^\pi \mu\left(\Pi,\alpha_i\right) dF_h(\Pi,\alpha_j) + \theta_l \int_0^1 \left[1-\mu\left(\Pi,\alpha_i\right)\right] dF_l(\Pi,\alpha_j) \end{aligned}$$

where $i = 1, 2$ *and* $j = 2, 1$.

Recall that $1 - \mu(\cdot, \alpha_i) = p_l f_l(\cdot, \alpha_i)$ and $\mu(\cdot, \alpha_i) = p_h f_h(\cdot, \alpha_i)$, we can then make the following observation:

Observation 3. *The equilibrium effort given in Proposition 4 is symmetric, in particular:*

$$b_{il}(\pi, \alpha_i, \alpha_j) = b_{jl}(\pi, \alpha_j, \alpha_i) \text{ and } b_{ih}(\pi, \alpha_i, \alpha_j) = b_{jh}(\pi, \alpha_j, \alpha_i)$$

This also suggests the probability that type (θ_h, π_i) wins against the high valuation opponent is equal to the probability of $\pi_i > \pi_j$, and the probability that type (θ_l, π_i) wins against the low valuation opponent is equal to the probability of $\pi_i < \pi_j$.

To gain some intuition, suppose type (θ_l, π) of player i increases her effort from $b_{il}(\pi, \alpha_i, \alpha_j)$ to $b_{il}(z, \alpha_i, \alpha_j)$ which costs her an additional effort of $\theta_l \int_z^\pi [1 - \mu(\Pi, \alpha_i)]\, dF_l(\Pi, \alpha_j)$, according to Proposition 4. This increases her probability of winning by $[1 - \mu(\pi, \alpha_i)] \int_z^\pi dF_l(\Pi, \alpha_j)$. At the optimum the cost must be equal to the gain, i.e.

$$\theta_l \int_z^\pi [1 - \mu(\Pi, \alpha_i)]\, dF_l(\Pi, \alpha_j) = \theta_l [1 - \mu(\pi, \alpha_i)] \int_z^\pi dF_l(\Pi, \alpha_j) \tag{1.7}$$

which is only true when $z = \pi$, implies optimality of $b_{il}(\pi, \alpha_i, \alpha_j)$. Similarly, if type (θ_h, π) of player i increases her effort from $b_{ih}(\pi, \alpha_i, \alpha_j)$ to $b_{ih}(z, \alpha_i, \alpha_j)$ which costs her an additional effort of $\theta_h \int_\pi^z \mu(\Pi, \alpha_i)\, dF_h(\Pi, \alpha_j)$. This increases her probability of winning by $\mu(\pi, \alpha_i) \int_\pi^z dF_h(\Pi, \alpha_j)$. Again, at the optimum the cost of the additional effort must match the gain:

$$\theta_h \int_\pi^z \mu(\Pi, \alpha_i)\, dF_h(\Pi, \alpha_j) = \theta_h \mu(\pi, \alpha_i) \int_\pi^z dF_h(\Pi, \alpha_j) \tag{1.8}$$

which is only true when $\pi = z$, implies optimality of $b_{ih}(\pi, \alpha_i, \alpha_j)$.

Based on Proposition 4, the sensitivity of player i's effort to π depends on both her own and the opponent's accuracies, α_i and α_j respectively, see Corollary 7.

Corollary 7 (Sensitivity)**.** *When either* α_i *or* α_j *increases, the slope of* $b_{ih}(\pi, \alpha_i, \alpha_j)$ *and* $b_{il}(\pi, \alpha_i, \alpha_j)$ *are decreased for* $\pi < \pi^+$*, and are increased for* $\pi > \pi^+$.

When α_i increases, player i's effort choice becomes more sensitive to changes of π, similar as in Corollary 2. However, when α_j increases, that is, player j's signal becomes more informative, player i's effort choice also becomes more sensitive to π. This is because player i also knows α_j and thus, anticipates the effect on player j's effort, so she has to adjust her own effort in response.

To see the intuition, consider the gain from increasing the effort by type (θ_l, π) of player i, i.e. $\theta_l[1-\mu(\pi,\alpha_i)]\int_z^\pi dF_l(\Pi,\alpha_j)$. When α_j increases, i.e. the opponent's signal becomes more informative, this gain is increased for $\pi, z \in (0,\pi^+)$ and is decreased for $\pi, z \in (\pi^+, 1)$, as the probability that the opponent's signal realization, π_j lies in the former interval is larger but the probability that π_j lies in the latter interval is smaller. Since in equilibrium the gain must be equal to the cost, as shown in (1.7), the cost of additional effort by player i, $b_{il}(z,\alpha_i,\alpha_j) - b_{il}(\pi,\alpha_i,\alpha_j)$, must also be increased for $\pi, z \in (0,\pi^+)$ and decreased for $\pi, z \in (\pi^+, 1)$. Fixing $\pi - z$, the slope of $b_{il}(\pi,\alpha_i,\alpha_j)$ must be steeper in $\pi, z \in (0,\pi^+)$ and flatter in $\pi, z \in (\pi^+, 1)$. The intuition for the sensitivity of $b_{ih}(\pi,\alpha_i,\alpha_j)$ is similar and can be derived by referring back to (1.8).

As one might expect, each player's ex ante expected payoff in this contest with asymmetric accuracy is again higher than when none of the players receive any signal about the opponent.

Proposition 5. *When $\alpha_1 \neq \alpha_2 \in (\underline{\alpha}, \overline{\alpha})$ and suppose Assumption 3 is satisfied, then*

$$V_i(\alpha_i, \alpha_j) > V_i(\underline{\alpha}, \underline{\alpha})$$

The intuition is similar as in the symmetric exogenous accuracy setting. Next we turn to the information disclosure of players. We first show that if players set up an information disclosure agreement, then they are both better off.

1.4.2 Disclosure agreement

The players in an information disclosure agreement commit to simultaneously disclose a signal to the opponent with pre-specified accuracy. Here we focus on the symmetric agreement where the accuracies of the signals are the same. We refer to the disclosure agreement where players commit to disclose signals with accuracy α as the "information disclosure agreement α". The following result suggests that an agreement to disclose partially informative signals is beneficial to both players.

Proposition 6. *Any information disclosure agreement $\alpha \in (\underline{\alpha}, \overline{\alpha})$ makes player i ($i = 1, 2$) strictly better off than no disclosure agreement or full disclosure agreement, i.e. $V_i(\alpha, \alpha) > V_i(\underline{\alpha}, \underline{\alpha}) = V_i(\overline{\alpha}, \overline{\alpha})$.*

The proof follows directly from that of Proposition 3, thus omitted. When both players obey the agreement and discloses a partially informative signal to the opponent, they can then coordinate the efforts and thus earn higher expected payoffs. Note that the disclosure agreement is only profitable if the signals are partially informative. This suggests that there is some loss of generality to only consider full disclosure and full concealment.

Would players obey the disclosure agreement if there is no external enforcement? In other words, if players choose disclosure strategies to maximize their expected payoffs, do we have the same outcome as in the disclosure agreement equilibrium?

1.4.3 Endogenous information disclosure

To solve the equilibrium disclosure decision, (α_j^*, α_i^*), note that players' equilibrium efforts are the ones shown in Proposition 4. Thus, player i chooses α_j to maximize her equilibrium expected payoff in the contest, denoted by $V_i(\alpha_i, \alpha_j)$, given α_i, i.e. α_j is chosen to best response to α_i. The best response function of player i is derived by the first order condition of her equilibrium expected payoff in the contest, $V_i(\alpha_i, \alpha_j)$, w.r.t. α_j:

$$\frac{\partial V_i(\alpha_i, \alpha_j)}{\partial \alpha_j} = \int_0^1 \left[\frac{\theta_h + \frac{\theta_l}{p_l}}{\theta_h - \theta_l} - f_h(\Pi, \alpha_i) \left(\frac{\theta_h + \frac{p_h \theta_l}{p_l}}{\theta_h - \theta_l} - \Pi \right) - F_h(\Pi, \alpha_i) \right] \frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi = 0 \tag{1.9}$$

It then follows that no player would obey the information disclosure agreement $\alpha \in (\underline{\alpha}, \widehat{\alpha}]$ as it is strictly dominant to choose $\underline{\alpha}$ when the opponent chooses $\alpha > \underline{\alpha}$. See Lemma 7 below.

Lemma 7. *If player j chooses $\alpha_i \in (\underline{\alpha}, \widehat{\alpha}]$[21], then player i strictly prefers to choose $\alpha_j = \underline{\alpha}$.*

Player i wants to avoid the motivation effect and take advantage of the demotivation effect on the opponent. Specifically, type (θ_h, π) of player i finds it profitable to lower the accuracy of the signal she discloses, as then the high valuation opponent is relatively more likely to receive low realizations. Similarly, type (θ_l, π) of player i also finds it profitable to lower the accuracy of the signal she discloses, as then the low valuation opponent is relatively more likely to receive high realizations. Therefore, player i earns higher expected payoff ex ante by decreasing the accuracy of the signal she discloses to the opponent.

This interpretation can be demonstrated formally. Denoted by $V_i(\theta_h, \pi, \alpha_i, \alpha_j)$ and $V_i(\theta_l, \pi, \alpha_i, \alpha_j)$ player i's equilibrium expected payoff when she is type (θ_h, π) and (θ_l, π), respectively. The intuition of the above result can be shown by rearranging

[21] Recall that $\widehat{\alpha}$ is the highest possible accuracy given that the equilibrium is allocative efficient.

$V_i(\theta_h, \pi, \alpha_i, \alpha_j)$:

$$V_i(\theta_h, \pi, \alpha_i, \alpha_j) = p_h\theta_h \int_0^{\pi} \left[f_h(\pi, \alpha_i) - f_h(\Pi, \alpha_i)\right] dF_h(\Pi, \alpha_j)$$
$$-p_l\theta_l \int_0^1 f_l(\Pi, \alpha_i) dF_l(\Pi, \alpha_j) + p_l\theta_h f_l(\pi, \alpha_i) \qquad (1.10)$$

By the definition of rotation order, the larger α_i is, the larger gap between $f_h(\pi, \alpha_i)$ and $f_h(\Pi, \alpha_i)$ in (1.10). Fixing α_i, this gap is also larger for higher realizations Π, when $\Pi \in [0, \pi]$, see Figure 1.13. It is thus profitable for player i to lower the distribution of the opponent's signal, $F_h(\Pi, \alpha_j)$, stochastically. In other words, player i has an incentive to lower the accuracy of the signal she discloses to the opponent.[22] Therefore, it is then optimal for player i to choose $\alpha_j = \underline{\alpha}$, i.e. to disclose an uninformative signal to the opponent.

The similar interpretation can be applied to the type (θ_l, π) of player i's disclosure decision. Her equilibrium expected payoff is rewritten as in (1.11):

$$V_i(\theta_l, \pi, \alpha_i, \alpha_j) = p_l\theta_l \int_{\pi}^{1} \left[f_l(\pi, \alpha_i) - f_l(\Pi, \alpha_i)\right] f_l(\Pi, \alpha_j) d\Pi \qquad (1.11)$$

The larger α_i is, the larger gap between $f_l(\pi, \alpha_i)$ and $f_l(\Pi, \alpha_i)$ is. Note that this gap is larger for higher realizations Π, when $\Pi \in [\pi, 1]$. See Figure 1.14. Thus, fixing $\alpha_i > \underline{\alpha}$, it is profitable for player i to increase the distribution of the opponent's signal, $f_l(\Pi, \alpha_j)$, stochastically. Thus, it is optimal for player i to choose $\alpha_j = \underline{\alpha}$, i.e. to disclose an uninformative signal to the opponent. When $\alpha_i = \underline{\alpha}$, i.e. the

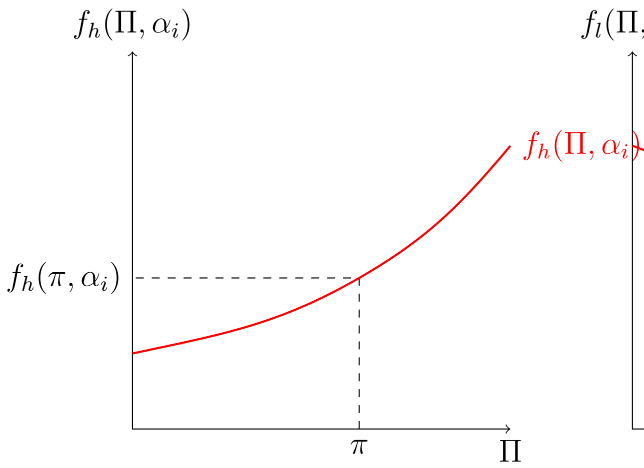

Figure 1.13

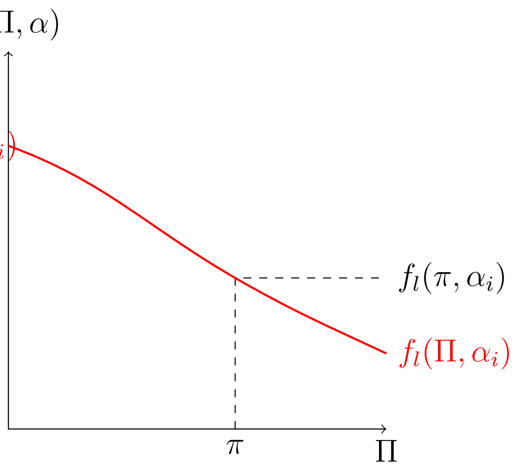

Figure 1.14

[22]The second term in (1.10), i.e. $-p_l\theta_l \int_0^1 f_l(\Pi, \alpha_i) dF_l(\Pi, \alpha_j)$, also becomes larger when α_j decreases as $f_l(\Pi, \alpha_i)$ is decreasing in Π when $\alpha_i > \underline{\alpha}$.

opponent discloses to player i an uninformative signal, then by (1.10) and (1.11), $V_i(\theta_h, \pi, \alpha_i, \alpha_j)$ and $V_i(\theta_l, \pi, \alpha_i, \alpha_j)$ becomes constants, and player i is indifferent about the accuracy of the signal she discloses to the opponent. This suggests when the opponent discloses an uninformative signal, then player i is indifferent about the accuracy of the signal she discloses. This result is shown in Lemma 8 below.

Lemma 8. *When the player j chooses $\alpha_i = \underline{\alpha}$, i.e. discloses an uninformative signal, then the disclosure decision of player i is irrelevant to his own expected payoff, as $V_i(\alpha_i = \underline{\alpha}, \alpha_j) = p_h p_l (\theta_h - \theta_l)$ for all $\alpha_j \in [\underline{\alpha}, \widehat{\alpha}]$.*

Both player's strategy are ex ante uniform. In other words, when player i observes an uninformative signal, she would expect the distribution of the opponent's effort is uniformly distributed, the same as in the IPV setting. This implies the best response of player i is to choose effort as if she is in the IPV setting. This then explains why her expected payoff is the same as in the IPV setting.

Lemma 7 and 8 jointly imply the following result.

Proposition 7. *There does not exist any equilibrium in which both player disclose an informative signal.*

Even though disclosing private information can improve total welfare, the industry cannot rely on decentralized information disclosure by players. Nevertheless, the regulator may be able to set up a minimum information disclosure requirement which specifies the accuracies of signals that players should disclose.

1.5 Conclusion

When players spy on each other, the additional information about the opponent allows them to coordinate, i.e. only exert higher effort when it is more likely that the opponent is evenly matched with the player. Such a coordination improves players' welfare even taking the cost of spying into account. This, however, is only true when spying is costly so that players acquire partially informative signals. An information disclosure agreement in which players commit to disclose a partially informative signal to each other can achieve an even better outcome, as the cost on spying is saved. However, players would unilaterally deviate by disclosing an uninformative signal if there is no external power to enforce the agreement. This is due to the incentive of players to avoid the motivation effect and to induce the demotivation effect on the opponent.

This paper yields differential policy implications dependent on nature of contests. For contests with wasteful efforts, e.g., rent-seeking, patent race and lobbying, it is advisable for regulators to impose a minimum disclosure requirement which specifies the minimum accuracy of signals players disclose to each other. For contests with productive efforts, e.g., sports tournaments, promotion contests and sales competitions, banning spying and disclosure maximizes total effort.

The model can potentially be extended in different directions. Firstly, the distribution of valuations may be generalized to a continuous distribution. Secondly, the model can be twisted a little to study overt information acquisition in which players are aware of the accuracy of the opponent's *spying* signal. Finally, the model can also be extended to contests with n players, and each player receives $n-1$ signals regarding the opponents' valuations. In order to facilitate the analysis, more restrictions may be imposed on the signal distribution.

1.6 Appendix: proofs

Proof of Lemma 1

Proof. Since now we consider α as exogenously given, we simplify the notations of effort strategy $b_h(\pi,\alpha)$ and $b_l(\pi,\alpha)$ into $b_h(\pi)$ and $b_l(\pi)$ respectively, and the posterior belief $\mu(\pi,\alpha)$ into $\mu(\pi)$. We start by proving part 1 of the lemma. Suppose in a symmetric, pure strategy equilibrium with efficient allocation, we have $b_h(\pi_1) < b_h(\pi_2)$ for $\pi_1 > \pi_2$. Then it must be true that the type (θ_h, π_1) finds the cost of increasing her effort from $b_h(\pi_1)$ to $b_h(\pi_2)$ dominates the gain from such an increase of effort, formally:

$$b_h(\pi_2) - b_h(\pi_1) \geqslant \mu(\pi_1)\Pr\{b_h(\pi_1) \leqslant b_j < b_h(\pi_2)\}\theta_h.$$

where b_j is the opponent, player j's effort. The cost must outweigh the gain to prevent type (θ_h, π_1) from deviating to $b_h(\pi_2)$. However, type (θ_h, π_2)'s gain must outweigh her cost of such an increase of effort:

$$b_h(\pi_2) - b_h(\pi_1) \leqslant \mu(\pi_2)\Pr\{b_h(\pi_1) \leqslant b_j < b_h(\pi_2)\}\theta_h.$$

where b_j is the opponent's effort. Combining the two conditions, we have $\mu(\pi_2) \geqslant \mu(\pi_1)$ which contradicts the fact that $\pi_1 > \pi_2$, due to Assumption 1. A similar argument can prove that $b_l(\pi_1) \leqslant b_l(\pi_2)$ for any $\pi_1 > \pi_2$.

To prove continuity of the strategies, i.e. the part 2 of the lemma, suppose there exists a discontinuous point on $b_h(\pi)$, say $\hat{\pi} \in (0,1)$, such that $b_h(\hat{\pi}) < b_h(\hat{\pi}+\epsilon)$ for an arbitrarily small ϵ. Then type $(\theta_h, \hat{\pi}+\epsilon)$ will find it profitable to deviate to some

$\hat{b} \in (b_h(\hat{\pi}), b_h(\hat{\pi}+\epsilon))$. Similarly, suppose there exists a discontinuous point on $b_l(\pi)$, $\tilde{\pi} \in (0,1)$, such that $b_l(\tilde{\pi}) > b_l(\tilde{\pi}+\epsilon)$ for arbitrarily small ϵ. Then type $(\theta_l, \tilde{\pi})$ will find it profitable to deviate to some $\tilde{b} \in (b_l(\tilde{\pi}+\epsilon), b_l(\tilde{\pi}))$.

To prove that there is no atom on any player's effort, suppose there exists p and q such that $1 > q > p > 0$ and that $b_h(x) = b$ where $x \in [p,q]$ and b is a constant. Then type $(\theta_h, p-\epsilon)$ will find it profitable to deviate to choosing $b+\epsilon$, as the gain of such deviation will be $\mu(p-\epsilon)\int_p^q f_h(\Pi)d\Pi$ and the cost is negligible when ϵ is arbitrarily small. A similar argument can show that there is no atom on $b_l(\pi)$.

Finally, for part 3, given part 1 is true, type $(\theta_h, 0)$ chooses the lowest effort among all types with valuation θ_h, whereas type $(\theta_l, 0)$ chooses the highest among all types with valuation θ_l. If $b_h(0) > b_l(0)$ then type $(\theta_h, 0)$ will be strictly better off by lowering her effort by a small amount ϵ satisfying $b_h(0) - \epsilon \geqslant b_l(0)$, thus $b_h(0) = b_l(0)$. Again, by part 1, the lowest effort is made by type $(\theta_l, 1)$ among all types, thus any positive effort is strictly dominated by choosing zero for $(\theta_l, 1)$. $\square$

Proof of Lemma 2

Proof. Note that allocative efficiency implies $b_h(\pi,\alpha) \geqslant b_l(\pi,\alpha)$ for all $\pi \in [0,1]$ fixing α, and the probability of tie between a high valuation and a low valuation player is zero. Then the type (θ_h, π_i)'s incentive compatibility condition such that she has no incentive to deviate to $b_l(\pi_i, \alpha)$, implies that:

$$
\begin{aligned}
& b_h(\pi_i,\alpha) - b_l(\pi_i,\alpha) \\
\leqslant\ & \left[\mu(\pi_i,\alpha)Pr\{b_j < b_h(\pi_i,\alpha)|(\theta_h,\theta_h)\} + \left[1-\mu(\pi_i,\alpha)\right]Pr\{b_j \geqslant b_l(\pi_i,\alpha)|(\theta_h,\theta_l)\}\right]\theta_h \\
=\ & \mu(\pi_i,\alpha)\int_0^{\pi_i} f_h(\Pi,\alpha)\theta_h d\Pi + (1-\mu(\pi_i,\alpha))\int_0^{\pi_i} f_h(\Pi,\alpha)\theta_h d\Pi \\
=\ & \int_0^{\pi_i} f_h(\Pi,\alpha)\theta_h d\Pi
\end{aligned}
\tag{1.12}
$$

In other words, the cost saved from choosing the lower effort (LHS of (1.12)) must be less than the gain forgone (RHS of (1.12)). This ensures that type (θ_h, π_i) does not want to deviate to choosing $b_h(\pi,\alpha)$. However, type (θ_l, π_i) should find her cost saved by choosing the lower effort outweighs her gain forgone:

$$\begin{aligned}
& b_h(\pi_i,\alpha) - b_l(\pi_i,\alpha) \\
\geqslant\ & [\mu(\pi_i,\alpha)Pr\{b_j < b_h(\pi_i,\alpha)|(\theta_l,\theta_h)\} + (1-\mu(\pi_i,\alpha))Pr\{b_j \geqslant b_l(\pi_i,\alpha)|(\theta_l,\theta_l)\}\theta_l \\
=\ & \mu(\pi_i,\alpha)\int_0^{\pi_i} f_l(\Pi,\alpha)\theta_l d\Pi + (1-\mu(\pi_i,\alpha))\int_0^{\pi_i} f_l(\Pi,\alpha)\theta_l d\Pi \\
=\ & \int_0^{\pi_i} f_l(\Pi,\alpha)\theta_l d\Pi
\end{aligned}$$

Combining the two conditions:

$$\int_0^{\pi_i} f_h(\Pi,\alpha)\theta_h d\Pi \geqslant \int_0^{\pi_i} f_l(\Pi,\alpha)\theta_l d\Pi$$

we then have $\frac{f_h(\pi,\alpha)}{f_l(\pi,\alpha)} \geqslant \frac{\theta_l}{\theta_h}$ for all $\pi_i \in [0,1]$. □

Proof of Proposition 1

Proof. The proposition is proved by checking whether type (θ_h,π) $((\theta_l,\pi))$ would deviate to any effort outside of their equilibrium support. Here we only show the idea of the proof by proving that the type (θ_h,π) of player i does not find it profitable to deviate to any effort that the type (θ_l,π) might choose in the allocative efficient equilibrium. The rest of the proofs are available in the online appendix.

We will refer to the interval $[b_l(1,\alpha), b_l(0,\alpha)]$, i.e. the equilibrium support of low valuation types, as the "low pure support", and refer to the interval $[b_h(0,\alpha), b_h(1,\alpha)]$, i.e. the equilibrium support of high valuation types, as the "high pure support". Let's start by checking whether the type (θ_h,π) of player i finds it profitable to deviate to any effort in the low pure support. This requires a comparison of the type (θ_h,π) of player i's expected payoff in the allocative efficient equilibrium:

$$\begin{aligned}
V(\theta_h,\pi,\alpha) &= \theta_h\int_0^{\pi}[\mu(\pi,\alpha)-\mu(\Pi,\alpha)]\,f_h(\Pi,\alpha)d\Pi \\
&\quad + \int_0^1 \{[1-\mu(\pi,\alpha)]\,\theta_h - [1-\mu(\Pi,\alpha)]\,\theta_l\}\,f_l(\Pi,\alpha)d\Pi,
\end{aligned}$$

to the maximum expected payoff from deviation. When deviating to $\beta \in [b_l(1,\alpha), b_l(0,\alpha)]$, the expected payoff of the type (θ_h,π) of player i given that the player j plays the allocative efficient equilibrium $b_l(\pi,\alpha)$ as given in the proposition is:

$$\widetilde{U}^l(\beta|\theta_h,\pi,\alpha) = \theta_h\,[1-\mu(\pi,\alpha)]\int_{b_l^{-1}(\beta,\alpha)}^1 f_h(\Pi,\alpha)d\Pi - \beta. \tag{1.13}$$

Obviously, among all $\beta \in [b_l(1,\alpha), b_l(0,\alpha)]$, player i would prefer to deviate to the optimal effort: $\beta^* = \arg\max_\beta \widetilde{U}(\beta|\theta_h,\pi)$. The optimal deviation effort, β^*, can be found by the first order condition with respect to β. Let the type (θ_l, t) be the one who chooses β^* in equilibrium, i.e. $b_l(t,\alpha) = \beta^*$. We can find t by the FOC of $\widetilde{U}^l(\beta|\theta_h,\pi)$ w.r.t β, and rearrange:

$$1 - \mu(\pi,\alpha) = \frac{\theta_l}{\theta_h}\frac{f_l(t,\alpha)}{f_h(t,\alpha)}(1-\mu(t,\alpha)) \tag{1.14}$$

Note that both sides are decreasing functions of their arguments, π and t, respectively. Since $\frac{\theta_l}{\theta_h}\frac{f_l(t,\alpha)}{f_h(t,\alpha)} \leqslant 1$, thus $\pi \geqslant t$. Then there must exists $\widehat{s} \in [0,1]$ such that

$$1 - \mu(\widehat{s},\alpha) \equiv \frac{\theta_l}{\theta_h}\frac{f_l(0,\alpha)}{f_h(0,\alpha)}\left[1-\mu(0,\alpha)\right] \tag{1.15}$$

For $\pi < \widehat{s}$, the LHS of the equation (1.14) is always strictly larger than the RHS, for all $t \in [0,1]$. This implies the first order derivative is positive and thus type (θ_h, π) does not want to deviate to the low pure support, whenever $\pi < \widehat{s}$. On the other hand, if $\pi \geqslant \widehat{s}$, there is always a unique interior solution of $t \in [0,1]$ satisfying equation (1.14) given π. In this case, we need to directly compare the equilibrium payoff with the payoff of choosing β^*. The maximum deviation expected payoff can be calculated by plugging β^* into (1.13). The first order derivative of the difference between the equilibrium expected payoff and the maximum deviation payoff, i.e. $V(\theta_h,\pi,\alpha) - \widetilde{U}^l(\beta^*|\theta_h,\pi,\alpha)$, w.r.t π, is, in fact, non-negative:

$$\frac{\partial\left(V(\theta_h,\pi,\alpha) - \widetilde{U}(\beta^*|\theta_h,\pi,\alpha)\right)}{\partial\pi} = \theta_h\mu'(\pi,\alpha)\left(\int_0^\pi f_h(\Pi,\alpha)d\Pi - \int_0^t f_h(\Pi,\alpha)d\Pi\right) \geqslant 0$$

This suggests this difference is non-decreasing in π. By (1.15) it can be proved that $V(\theta_h,\widehat{s}) - \widetilde{U}(b_l(0,\alpha)|\theta_h,\pi) = \theta_h\mu'(\widehat{s},\alpha)\int_0^{\widehat{s}} f_h(\Pi,\alpha)d\Pi > 0$. Thus, $V(\theta_h,\pi) - \widetilde{U}(\beta^*|\theta_h,\pi) \geqslant 0$ for all $\pi \in [\widehat{s},1]$. Therefore, type (θ_h,π) does not find it profitable to deviate from equilibrium strategy.

In the inefficient equilibrium, the proof involves checking whether the type who plays pure strategy finds it profitable to deviate to the mixed strategy support and vice versa. The uniqueness of both the allocative efficient and inefficient equilibrium are due to the initial conditions given in the Lemma 1. □

Proof of Corollary 2

Proof. The first part of the corollary is obvious after taking the first order derivative of $b_h(\pi,\alpha)$ and $b_l(\pi,\alpha)$ w.r.t α, thus is omitted. Here we show the calculation of

$\frac{\partial \pi^*}{\partial \alpha}$. Rewrite equation (1.5) in Proposition 1 to

$$\int_0^{\pi^*} [f_l(\Pi, \alpha)\theta_l - f_h(\Pi, \alpha)\theta_h] \, d\Pi = 0$$

and take first order derivative w.r.t α, we have

$$\frac{\partial \pi^*}{\partial \alpha} = \frac{\int_0^{\pi^*} \left[\frac{\partial f_h(\Pi,\alpha)}{\partial \alpha}\theta_h - \frac{\partial f_l(\Pi,\alpha)}{\partial \alpha}\theta_l\right] d\Pi}{f_l(\pi^*, \alpha)\theta_l - f_h(\pi^*, \alpha)\theta_h}$$

Since $\int_0^{\pi^*} \frac{\partial f_h(\Pi,\alpha)}{\partial \alpha} d\Pi < 0$ and $\int_0^{\pi^*} \frac{\partial f_l(\Pi,\alpha)}{\partial \alpha} d\Pi > 0$, and by (1.6) we have $f_l(\pi^*, \alpha)\theta_l - f_h(\pi^*, \alpha)\theta_h < 0$ (as $\pi^* > \pi^+$ according to (1.5)), it must be true that $\frac{\partial \pi^*}{\partial \alpha} > 0$. $\square$

Proof of Lemma 3

Proof. Suppose player j chooses $\alpha_j = \alpha$ and (wrongly) believes that player i has chosen the same, yet player i instead chooses $\alpha_i = \eta$. Thus, player j plays the symmetric equilibrium given in Propositions 1. To prove that the marginal expected payoff from increasing η by player i in this situation is positive, i.e. $AMR(\eta, \alpha) > 0$, we first find the maximum expected payoff of each type of player i. We will refer to the interval $[b_l(1, \alpha), b_l(0, \alpha)]$, i.e. the equilibrium support of low valuation types, as the "low pure support"; and refer to the interval $[b_h(0, \alpha), b_h(1, \alpha)]$, i.e. the equilibrium support of high valuation types, as the "high pure support". Denoted by $U^k(b|\theta_i, \pi_i, \eta, \alpha)$ the expected payoff of the type (θ_i, π_i) of player i chooses b in the equilibrium support of types with θ_k ($k \in \{h, l\}$) when player j chooses α and believes that player i has chosen the same, whereas player i, in fact, chooses η. Denoted by $V^k(\theta_i, \pi_i, \eta, \alpha)$ where $k \in \{h, l\}$, the maximum of $U^k(b|\theta_i, \pi_i, \eta, \alpha)$.

First, we focus on the case when $\alpha \leqslant \widehat{\alpha}$, i.e. when player j will play the symmetric, pure strategy equilibrium with efficient allocation. Note that both the type (θ_h, π) and the type (θ_l, π) of player i may choose an optimal effort in either the high or the low pure support, dependent on which interval would provide them higher expected payoff.

Specifically, if type (θ_h, π) of player i chooses an effort in the high pure support, then her expected payoff is:

$$U^h(b|\theta_h, \pi, \eta, \alpha) = \theta_h \left[1 - \mu(\pi, \eta) + \mu(\pi, \eta) \int_0^{b_h^{-1}(b,\alpha)} f_h(\Pi, \alpha) d\Pi\right] - b. \qquad (1.16)$$

where $b_h^{-1}(b, \alpha)$ is the inverse of the equilibrium pure strategy that player j plays. By the first order condition w.r.t b we know that the type (θ_h, π) of player i must find it optimal to choose $b_h(s,\alpha)$ where s satisfies $f_h(\pi,\eta) = f_h(s,\alpha)$, which also implies $f_l(\pi,\eta) = f_l(s,\alpha)$. In other words,

$$s = f_h^{-1}(f_h(\pi,\eta),\alpha) = f_l^{-1}(f_l(\pi,\eta),\alpha) \tag{1.17}$$

When $\eta < \alpha$, by definition of rotation order, there always exists an s satisfying (1.17) for all $\pi \in [0,1]$. Therefore, the maximum expected payoff for type (θ_h,π) when she chooses the optimal effort in the high pure support is given by:

$$\begin{aligned} V^h(\theta_h,\pi,\eta,\alpha) &= \theta_h \int_0^s \left[\mu(\pi,\eta) - \mu\left(\Pi,\alpha\right)\right] f_h(\Pi,\alpha) d\Pi + \left[1-\mu(\pi,\eta)\right]\theta_h \\ &\quad -\theta_l \int_0^1 \left[1-\mu\left(\Pi,\alpha\right)\right] f_l(\Pi,\alpha) d\Pi \end{aligned}$$

Suppose instead that type (θ_h,π) of player i chooses an effort in the low pure support, then her expected payoff is:

$$U^l(b|\theta_h,\pi,\eta,\alpha) = \theta_h \left[1-\mu(\pi,\eta)\right] \int_{b_l^{-1}(b,\alpha)}^1 f_h(\Pi,\alpha) d\Pi - b \tag{1.18}$$

where $b_l^{-1}(b,\alpha)$ is the inverse of the equilibrium pure strategy that player j will play. The first order condition w.r.t b requires:

$$f_l(\pi,\eta) = \frac{f_l(\widehat{s},\alpha)}{f_h(\widehat{s},\alpha)} \frac{\theta_l}{\theta_h} f_l(\widehat{s},\alpha)$$

meaning that player i would find it optimal to exert an effort $b_l(\widehat{s},\alpha)$ when she chooses η. Note that $f_l(s,\alpha) = f_l(\pi,\eta) < f_l(\widehat{s},\alpha)$, thus, $s > \widehat{s}$. The maximum expected payoff for the type (θ_h,π) when choosing an effort level in the low pure support is thus,

$$V^l(\theta_h,\pi,\eta,\alpha) = \theta_h \left[1-\mu(\pi,\eta)\right] \int_{\widehat{s}}^1 f_h(\Pi,\alpha) d\Pi - \theta_l \int_{\widehat{s}}^1 \left[1-\mu\left(\Pi,\alpha\right)\right] f_l(\Pi,\alpha) d\Pi$$

The difference between the two maximum expected payoffs is

$$
\begin{aligned}
& V^h(\theta_h, \pi, \eta, \alpha) - V^l(\theta_h, \pi, \eta, \alpha) \\
= \ & \theta_h \int_0^s \left[\underbrace{\mu(s, \alpha) - \mu(\Pi, \alpha)}_{>0} \right] f_h(\Pi, \alpha) d\Pi \\
& + \theta_h \int_0^{\widehat{s}} \left[[1 - \mu(\pi, \eta)] - [1 - \mu(\Pi, \alpha)] \frac{f_l(\Pi, \alpha)}{f_h(\Pi, \alpha)} \frac{\theta_l}{\theta_h} \right] f_h(\Pi, \alpha) d\Pi \\
> \ & \theta_h \int_0^s [\mu(\pi, \eta) - \mu(\Pi, \alpha)] f_h(\Pi, \alpha) d\Pi + \theta_h \int_0^{\widehat{s}} [\mu(\Pi, \alpha) - \mu(\pi, \eta)] f_h(\Pi, \alpha) d\Pi \\
> \ & \theta_h \int_0^{\widehat{s}} [\mu(\pi, \eta) - \mu(\Pi, \alpha)] f_h(\Pi, \alpha) d\Pi + \theta_h \int_0^{\widehat{s}} [\mu(\Pi, \alpha) - \mu(\pi, \eta)] f_h(\Pi, \alpha) d\Pi \\
= \ & 0
\end{aligned}
$$

Therefore, type (θ_h, π) of player i's expected payoff when choosing η is $V^h(\theta_h, \pi, \eta, \alpha)$, and thus, the marginal expected payoff from increasing η is given by

$$\frac{\partial}{\partial \eta} V^h(\theta_h, \pi, \eta, \alpha) = -p_h \theta_h \frac{\partial f_h(\pi, \eta)}{\partial \eta} \int_s^1 f_h(\Pi, \alpha) d\Pi$$

Now we turn to the types with the low valuation. If type (θ_l, π) of player i chooses an effort in the low pure support, i.e. $[b_l(1, \alpha), b_l(0, \alpha)]$, then her expected payoff is:

$$U^l(b|\theta_l, \pi, \eta, \alpha) = \theta_l [1 - \mu(\pi, \eta)] \int_{b_l^{-1}(b,, \alpha)}^1 f_l(\Pi, \alpha) d\Pi - b \tag{1.19}$$

where $b_l^{-1}(b,, \alpha)$ is the inverse of the equilibrium pure strategy that player j will play. By the first order condition w.r.t b we know that the type (θ_l, π) of player i should optimally choose $b_l(t, \alpha)$ where t satisfies $f_l(\pi, \eta) = f_l(t, \alpha)$, which also implies $f_h(\pi, \eta) = f_h(t, \alpha)$. In other words,

$$t = f_h^{-1}(f_h(\pi, \eta), \alpha) = f_l^{-1}(f_l(\pi, \eta), \alpha). \tag{1.20}$$

Note that (1.17) and (1.20) suggests $t = s$. When $\eta < \alpha$, by definition of rotation order, there always exists t and s satisfy (1.20) and (1.17) for all $\pi \in [0, 1]$, respectively. Therefore, the expected payoff for the type (θ_l, π) is given by:

$$V^l(\theta_l, \pi, \eta, \alpha) = \theta_l \int_t^1 [\mu(\Pi, \alpha) - \mu(\pi, \eta)] f_l(\Pi, \alpha) d\Pi$$

Suppose instead that type (θ_l, π) of player i chooses an effort in the high pure support, i.e. $[b_h(0,\alpha), b_h(1,\alpha)]$, then her maximum expected payoff is:

$$\begin{aligned}V^h(\theta_l,\pi,\eta,\alpha) &= p_h\theta_l\int_0^{\hat{t}}\left[f_h(\pi,\eta)-f_h(\Pi,\alpha)\frac{f_h(\Pi,\alpha)}{f_l(\Pi,\alpha)}\frac{\theta_h}{\theta_l}\right]f_l(\Pi,\alpha)d\Pi \\ &\quad +p_h\theta_l\int_0^1\left[f_h\left(\Pi,\alpha\right)-f_h(\pi,\eta)\right]f_l(\Pi,\alpha)d\Pi\end{aligned}$$

Again, we need to compare the two maximum expected payoff to determine whether type (θ_l,π) of player i should choose an effort in $[b_l(1,\alpha), b_l(0,\alpha)]$ or $[b_h(0,\alpha), b_h(1,\alpha)]$. It turns out that the former earns the type (θ_l,π) higher expected payoff:

$$\begin{aligned}& V^h(\theta_l,\pi,\eta,\alpha) \\ = \;& p_h\theta_l\int_0^{\hat{t}}\left[f_h(\pi,\eta)-f_h(\Pi,\alpha)\frac{f_h(\Pi,\alpha)}{f_l(\Pi,\alpha)}\frac{\theta_h}{\theta_l}\right]f_l(\Pi,\alpha)d\Pi \\ & +p_h\theta_l\int_0^{t}\left[f_h\left(\Pi,\alpha\right)-f_h(\pi,\eta)\right]f_l(\Pi,\alpha)d\Pi+p_h\theta_l\int_t^1\left[f_h\left(\Pi,\alpha\right)-f_h(\pi,\eta)\right]f_l(\Pi,\alpha)d\Pi \\ < \;& p_h\theta_l\int_0^{t}\left[1-\frac{f_h(\Pi,\alpha)}{f_l(\Pi,\alpha)}\frac{\theta_h}{\theta_l}\right]f_h\left(\Pi,\alpha\right)f_l(\Pi,\alpha)d\Pi+p_h\theta_l\int_t^1\left[f_h\left(\Pi,\alpha\right)-f_h(\pi,\eta)\right]f_l(\Pi,\alpha)d\Pi \\ < \;& p_h\theta_l\int_t^1\left(f_h\left(\Pi,\alpha\right)-f_h(\pi,\eta)\right)f_l(\Pi,\alpha)d\Pi \\ = \;& V^l(\theta_l,\pi,\eta,\alpha)\end{aligned}$$

Therefore, type (θ_l,π) of player i's expected payoff is $V^l(\theta_l,\pi,\eta,\alpha)$, and thus, the marginal expected payoff from increasing η is given by

$$\frac{\partial}{\partial\eta}V^l(\theta_l,\pi,\eta,\alpha)=-p_h\theta_l\frac{\partial f_h(\pi,\eta)}{\partial\eta}\int_t^1 f_l(\Pi,\alpha)d\Pi$$

Now, it can be proved that the ex ante marginal expected payoff from increasing η is positive:

$$\begin{aligned}AMR(\eta,\alpha) &= \int_0^1\left[p_l\frac{\partial}{\partial\eta}V^l(\theta_l,\pi,\eta,\alpha)+p_h\frac{\partial}{\partial\eta}V^h(\theta_h,\pi,\eta,\alpha)\right]d\pi \\ &= -p_h\int_0^1\left[\underbrace{p_l\theta_l\int_t^1 f_l(\Pi,\alpha)d\Pi+p_h\theta_h\int_s^1 f_h(\Pi,\alpha)d\Pi}_{A}\right]\frac{\partial f_h(\pi,\eta)}{\partial\eta}d\pi \\ &= p_h\left[p_l\theta_l\int_\nu^1 f_l(\Pi,\alpha)d\Pi+p_h\theta_h\int_\nu^1 f_h(\Pi,\alpha)d\Pi\right]\int_{\pi^+}^1\frac{\partial f_h(\pi,\eta)}{\partial\eta}d\pi \\ &\quad -p_h\left[p_l\theta_l\int_\xi^1 f_l(\Pi,\alpha)d\Pi+p_h\theta_h\int_\xi^1 f_h(\Pi,\alpha)d\Pi\right]\int_{\pi^+}^1\frac{\partial f_h(\pi,\eta)}{\partial\eta}d\pi \\ &> 0\end{aligned}$$

where $\nu \in [0, \pi^+]$ and $\xi \in [\pi^+, 1]$. Recall that $s = t$ which is why the third equality is true. Note that by definition of rotation order, $-\int_0^{\pi^+} \frac{\partial f_h(\pi,\eta)}{\partial \eta} d\pi = \int_{\pi^+}^1 \frac{\partial f_h(\pi,\eta)}{\partial \eta} d\pi > 0$, and also that the term A is decreasing with s and t. Thus, by applying the intermediate value theorem for integrals, we have the inequality given at the end.

We have proved that $AMR(\eta, \alpha) > 0$ given that $\alpha \leqslant \widehat{\alpha}$. The proof for the part when $\alpha \leqslant \widehat{\alpha}$ follows the same idea, i.e. by checking for profitable deviation. This part of the proof is omitted and is available in the online appendix.

□

Proof of Lemma 4

Proof. Recall from Section 1.2, when the signals players receive are perfectly informative, we have $f_h(\pi_i, \overline{\alpha}) = 0$ if $\pi_i \leqslant p_l$ and $f_h(\pi_i, \overline{\alpha}) = \frac{1}{p_h}$ if $\pi_i > p_l$; and correspondingly, that $f_l(\pi_i, \overline{\alpha}) = \frac{1}{p_l}$ if $\pi_i \leqslant p_l$ and $f_l(\pi_i, \overline{\alpha}) = 0$ if $\pi_i > p_l$. The notations of this proof follows from Lemma 3.

When type (θ_l, π) of player i chooses an effort in the low pure support, the expected payoff given that player j playing the symmetric equilibrium in the contest with exogenously given $\overline{\alpha}$ is:

$$\begin{aligned} U^l(b|\theta_l, \pi, \eta, \overline{\alpha}) &= \theta_l \left[1 - \mu(\pi, \eta)\right] \int_{b_l^{-1}(b,\overline{\alpha})}^1 f_l(\Pi, \overline{\alpha}) d\Pi - b \\ &= 0 \end{aligned}$$

When type (θ_l, π) of player i chooses an effort in the mixed support, the expected payoff is

$$\begin{aligned} U^m(\theta_l, \pi, \eta, \overline{\alpha}) &= \theta_l \int_{\pi^*}^1 \left(\mu(\pi, \eta) - \mu\left(\Pi, \overline{\alpha}\right)\right) f_l(\Pi, \overline{\alpha}) d\Pi \\ &= 0 \end{aligned}$$

When the type (θ_l, π) of player i chooses an effort in the high pure support, the expected payoff is:

$$U^h(b|\theta_l, \pi, \eta, \overline{\alpha}) = \theta_l \left[1 - \mu(\pi, \eta) + \mu(\pi, \eta) \int_0^{b_h^{-1}(b,\overline{\alpha})} f_l(\Pi, \overline{\alpha}) d\Pi\right] - b$$

The first order derivative of $U^h(b|\theta_l, \pi, \eta, \overline{\alpha})$ w.r.t. b is:

$$\frac{f_h(\pi, \eta) f_l(\widehat{s}, \overline{\alpha})\theta_l}{f_h(\widehat{s}, \overline{\alpha}) f_h(\widehat{s}, \overline{\alpha})\theta_h} - 1 < 0 \tag{1.21}$$

Thus, the type (θ_l, π) of player i does not want to choose any effort in the high pure support. She finds it optimal to choose an effort in the mixed support and thus, her maximum expected payoff is zero.

Now we turn to type (θ_h, π) of player i's optimal effort given she chooses η. When she chooses an effort in the high pure support, the expected payoff is:

$$U^h(b|\theta_h, \pi, \eta, \overline{\alpha}) = \theta_h \left((1 - \mu(\pi, \eta)) + \mu(\pi, \eta) \int_0^{b_h^{-1}(b,\overline{\alpha})} f_h(\Pi, \overline{\alpha}) d\Pi \right) - b$$

and by taking the first order derivative w.r.t. b we have $\frac{f_h(t,\overline{\alpha})}{f_h(\pi,\eta)} - 1 = -1 \leqslant 0$. Thus, player i does not want to choose any effort in the high pure support and again, she finds it optimal to choose an effort in the mixed effort, which earns her zero expected payoff.

When the type (θ_h, π) of player i chooses an effort in the low pure support, the expected payoff is:

$$U^l(b|\theta_h, \pi, \eta, \overline{\alpha}) = \theta_h(1 - \mu(\pi, \eta)) \int_{b_l^{-1}(b,\overline{\alpha})}^1 f_h(\Pi, \overline{\alpha}) d\Pi - b$$

The first order derivative w.r.t b gives

$$\frac{f_l(\pi, \eta) f_h(\widehat{t}, \overline{\alpha})}{f_l(\widehat{t}, \overline{\alpha}) f_l(\widehat{t}, \overline{\alpha})} \frac{\theta_h}{\theta_l} - 1 = +\infty$$

Thus, the type (θ_h, π) of player i does not want to choose any effort in the low pure support. In other words, the optimal effort is always in the mixed support.

When type (θ_h, π) of player i chooses an effort in the mixed support, the expected payoff is:

$$U^m(b|\theta_h, \pi, \eta, \overline{\alpha}) = V^m(\theta_h, \pi, \eta, \overline{\alpha}) \quad = \quad p_l \theta_h f_l(\pi, \eta) \int_{\pi^*}^1 f_h(\Pi, \overline{\alpha}) d\Pi$$

The first order derivative of $V^m(\theta_h, \pi, \eta, \alpha)$ w.r.t η is:

$$\frac{\partial V^m(\theta_h, \pi, \eta, \overline{\alpha})}{\partial \eta} = -p_h \theta_h \frac{\partial f_h(\pi, \eta)}{\partial \eta} \int_{\pi^*}^1 f_h(\Pi, \overline{\alpha}) d\Pi$$

We are now able to calculate the ex ante marginal expected payoff of player i

w.r.t η:

$$\begin{aligned} AMR(\eta,\overline{\alpha}) &= p_h \int_0^1 \frac{\partial U^m(b|\theta_h,\pi,\eta,\overline{\alpha})}{\partial \eta} d\pi \\ &= -p_h^2\theta_h \int_0^1 \frac{\partial f_h(\pi,\eta)}{\partial \eta} \int_{\pi^*}^1 f_h(\Pi,\overline{\alpha}) d\Pi \\ &= 0 \end{aligned}$$

□

Proof of Proposition 3

Proof. Given that $AMR(\eta,\alpha) > 0$ for any $\eta < \alpha \in [\underline{\alpha},\overline{\alpha}]$, it must be true that $MR(\alpha) > 0$. Suppose there exists some convex cost function with marginal cost function $MC(\cdot)$ which crosses $MR(\cdot)$ only once from below, and there is an interior solution $\alpha^* \in (\underline{\alpha},\overline{\alpha})$ of $MR(\alpha) = MC(\alpha)$. Then there exists an equilibrium of the information acquisition game in which both players chooses α^* and $MR(\alpha^*) = MC(\alpha^*)$.

We now turn to prove the second part of the proposition. By the first order conditions derived from (1.19) and (1.16) we know that type (θ_l,π) of player i who chose $\underline{\alpha}$ should choose $b_l(t,\alpha)$ when t satisfies $1 = f_l(\pi,\underline{\alpha}) = f_l(t,\alpha)$, i.e. $t = \widehat{\pi}$. Type (θ_h,π) of player i should choose $b_h(s,\alpha)$ when s satisfies $1 = f_h(\pi,\underline{\alpha}) = f_h(s,\alpha)$ thus $s = \widehat{\pi}$. The maximum expected payoff for the type (θ_l,π) and $(\theta_h,\pi,)$ are given by:

$$\begin{aligned} V(\theta_l,\pi,\underline{\alpha},\alpha) &= p_h\theta_l \int_{\widehat{\pi}}^1 \left[f_h(\Pi,\alpha) - f_h(\pi,\underline{\alpha})\right] dF_l(\Pi,\alpha) > 0 \\ V(\theta_h,\pi,\underline{\alpha},\alpha) &= p_h \underbrace{\int_0^{\widehat{\pi}} \left[1 - f_h(\Pi,\alpha)\right]\left[f_h(\Pi,\alpha)\theta_h - f_l(\Pi,\alpha)\theta_l\right] d\Pi}_{>0} \\ &\quad + p_l\theta_l \underbrace{\int_{\widehat{\pi}}^1 \left[1 - f_l(\Pi,\alpha)\right] f_l(\Pi,\alpha) d\Pi}_{>0} + p_l\left(\theta_h - \theta_l\right) \end{aligned}$$

This suggests player i's ex ante expected payoff when she chose $\underline{\alpha}$ given player j plays symmetric equilibrium in the contest believing both players have chosen $\alpha > \underline{\alpha}$, i.e. $V(\underline{\alpha},\alpha) = \int_0^1 p_l V(\theta_l,\pi,\underline{\alpha},\alpha) + p_h V(\theta_h,\pi,\underline{\alpha},\alpha) d\pi$, is larger than when both players chose $\underline{\alpha}$, i.e. $p_h p_l\left(\theta_h - \theta_l\right)$. Note that $V(\underline{\alpha},\alpha)$ may not be the maximum ex ante expected payoff of player i, but the fact that the above is larger than $p_h p_l\left(\theta_h - \theta_l\right)$ suggests the optimal must be always larger.

Since Lemma 3 has shown that $AMR(\eta, \alpha) > 0$, suggesting $V(\eta, \alpha)$ is increasing in η, it must be true that

$$V(\alpha^*, \alpha^*) - C(\alpha^*) \geqslant V(\underline{\alpha}, \alpha) - C(\underline{\alpha}) > V(\underline{\alpha}, \underline{\alpha}) = V(\overline{\alpha}, \overline{\alpha})$$

The inequality equality is due to the optimality of α^*. The inequality is due to $V(\underline{\alpha}, \alpha) \geqslant p_h p_l(\theta_h - \theta_l) = V(\underline{\alpha}, \underline{\alpha})$ as shown above, and the fact that $C(\underline{\alpha}) = 0$.

Therefore, player i's expected payoff in the entire game is strictly higher than when both players choosing $\underline{\alpha}$, i.e. not spying on each other. Furthermore, by $V(\underline{\alpha}, \underline{\alpha}) = V(\overline{\alpha}, \overline{\alpha})$, player i's expected payoff in the game is also higher than when both players receive a perfect signal about the opponent for free. □

Proof of Lemma 5

Proof. For ease of notation, in this proof we simplify the notations of effort choice functions $b_{ih}(\pi, \alpha_i, \alpha_j)$ and $b_{il}(\pi, \alpha_i, \alpha_j)$ into $b_{ih}(\pi)$ and $b_{il}(\pi)$, respectively. We start by proving part 1. Suppose in a pure strategy equilibrium with efficient allocation, we have $b_{ih}(\pi_1) < b_{ih}(\pi_2)$ for some $\pi_1 > \pi_2$. Then type (θ_h, π_1) of player i must find the cost of increasing her effort from $b_{ih}(\pi_1)$ to $b_{ih}(\pi_2)$ dominates the gain from such increase, formally:

$$b_{ih}(\pi_2) - b_{ih}(\pi_1) \geqslant \mu(\pi_1) \Pr\{b_{ih}(\pi_1) \leqslant b_{jh} < b_{ih}(\pi_2)\}\theta_h.$$

where b_{jh} is player j's effort. The LHS of the above is the cost from increasing effort from $b_{ih}(\pi_1)$ to $b_{ih}(\pi_2)$, and the RHS is the gain from doing so. The cost must outweigh the gain to prevent type (θ_h, π_1) from deviating to $b_{ih}(\pi_2)$. However, type (θ_h, π_2) would require her gain outweighs her cost of such increase of effort:

$$b_{ih}(\pi_2) - b_{ih}(\pi_1) \leqslant \mu(\pi_2) \Pr\{b_{ih}(\pi_1) \leqslant b_{jh} < b_{ih}(\pi_2)\}\theta_h.$$

Combining the two condition, we have $\mu(\pi_2) \geqslant \mu(\pi_1)$ which contradicts to $\pi_1 > \pi_2$, due to Assumption 1. Similar arguments can prove that $b_{il}(\pi_1) \leqslant b_{il}(\pi_2)$ for any $\pi_1 > \pi_2$.

To prove continuity of effortding strategies in part 2, suppose there exists a discontinuous point on player i's effort strategy $b_{ih}(\pi)$, $\hat{\pi} \in (0, 1)$, such that $b_{ih}(\hat{\pi}) < b_{ih}(\hat{\pi} + \epsilon)$ for an arbitrarily small ϵ. Then type $(\theta_h, b_{jh}^{-1}(b_{ih}(\hat{\pi} + \epsilon)))$ of player j will find it profitable to deviate to some $\hat{b} \in (b_{ih}(\hat{\pi}), b_{ih}(\hat{\pi} + \epsilon))$, where $b_{jh}^{-1}(\cdot)$ is the inverse of player j's effort strategy. Similarly, suppose there exists a discontinuous point on $b_{il}(\pi)$, $\tilde{\pi} \in (0, 1)$, such that $b_{il}(\tilde{\pi}) > b_{il}(\tilde{\pi} + \epsilon)$ for arbitrarily small ϵ.

Then type $(\theta_l, b_{jl}^{-1}(b_{il}(\tilde{\pi})))$ of player j will find it profitable to deviate to some $\tilde{b} \in (b_{il}(\tilde{\pi}+\epsilon), b_{il}(\tilde{\pi}))$, where $b_{jl}^{-1}(\cdot)$ is the inverse of player j's effort strategy.

To prove that there is no atom on any player's effort, suppose there exists p and q such that $1 > q > p > 0$ and that $b_{ih}(x) = b$ where $x \in [p, q]$ and b is a constant. Then by continuity there must be a type $(\theta_h, b_{jh}^{-1}(b-\epsilon))$ of player j who chooses $b - \epsilon$, and he will find it profitable to deviate to choosing $b + \epsilon$, as the gain of such deviation will be $\mu(b_{jh}^{-1}(b-\epsilon), \alpha_j) \int_p^q f_h(\Pi, \alpha_i) d\Pi > 0$ and the cost is negligible when ϵ is arbitrarily small. A similar argument can show that there is no atom on $b_{il}(\pi)$.

For part 3, given part 1 is true, type $(\theta_h, 1)$ of player i chooses the highest effort among all types in an allocative efficient equilibrium, whereas type $(\theta_l, 1)$ chooses the lowest effort among all types in an allocative efficient equilibrium. Thus, it must be true that $b_{il}(1) = b_{jl}(1) = 0$ as these are the lower bound of equilibrium support. They must be the same and cannot be positive. It must also be true that $b_{ih}(1) = b_{jh}(1)$ as these are the highest effort exerted by players, and in any equilibrium the equality is satisfied.

Finally, by part 1, type $(\theta_h, 0)$ of player i chooses the lowest effort among all types with valuation θ_h, whereas type $(\theta_l, 0)$ chooses the highest among all types with valuation θ_l. Suppose $b_{ih}(0) > b_{il}(0)$, then it implies that there is a gap in the equilibrium support of player i's effort. This cannot be part of any equilibrium as then player j would not choose any effort in $[b_{il}(0), b_{ih}(0)]$, which contradicts the optimality of $b_{ih}(0)$, as player i would want to deviate to any $b \in (b_{il}(0), b_{ih}(0))$. Suppose $b_{ih}(0) < b_{il}(0)$, then in any equilibrium with efficient allocation, it must be true that $b_{jl}(0) \leqslant b_{ih}(0) < b_{il}(0) \leqslant b_{jh}(0)$. But then this implies there is a gap in the equilibrium support of player j's effort, which we showed above to be impossible in any equilibrium with efficient allocation. Therefore, in any equilibrium with efficient allocation, it must be true that $b_{il}(0) = b_{ih}(0)$ for $i = 1, 2$. Now we prove $b_{il}(0) = b_{ih}(0) = b_{jl}(0) = b_{jh}(0)$. Without loss of generality, suppose $b_{il}(0) = b_{ih}(0) > b_{jl}(0) = b_{jh}(0)$, but then this contradicts efficient allocation as $b_{il}(0) > b_{jh}(0)$. Thus, in any pure strategy equilibrium with efficient allocation, it must be true that $b_{il}(0) = b_{ih}(0) = b_{jl}(0) = b_{jh}(0)$. □

Proof of Proposition 4

Proof. There are two steps to take to prove the proposition. First, we show that the equilibrium strategies of each valuation type given in the proposition are indeed the optimal strategy in their equilibrium support. Second, we show that each type do not want to deviate to any effort level outside of their equilibrium support.

Following the notation in the previous proofs, we refer to the equilibrium support of low valuation types as the "low pure support", and refer to the equilibrium support of high valuation types as the "high pure support".

Given that player j chooses his strategy according to the proposition. Suppose type (θ_l, π) of player i chooses an alternative effort level $b = b_{il}(s, \alpha_i, \alpha_j)$, then her expected payoff is

$$U_i(b|\theta_l, \pi, \alpha_i, \alpha_j) = \theta_l \int_s^1 [\mu(\Pi, \alpha_i) - \mu(\pi, \alpha_i)] dF_l(\Pi, \alpha_j)$$

Thus,

$$V_i(\theta_l, \pi, \alpha_i, \alpha_j) - U_i(b, \theta_l, \pi, \alpha_i, \alpha_j) = \theta_l \int_\pi^s [\mu(\Pi, \alpha_i) - \mu(\pi, \alpha_i)] dF_l(\Pi, \alpha_j) \geqslant 0$$

regardless of whether $\pi \geqslant s$ or $\pi < s$.

Suppose type (θ_h, π) of player i chooses an alternative effort level $b = b_{ih}(t, \alpha_i, \alpha_j)$, then her expected payoff is

$$\begin{aligned} U_i(b|\theta_h, \pi, \alpha_i, \alpha_j) &= \theta_h[(1 - \mu(\pi, \alpha_i)) + \mu(\pi, \alpha_i) \int_0^t f_h(\Pi, \alpha_j) d\Pi] \\ &\quad -\theta_h \int_0^t \mu(\Pi, \alpha_i) dF_h(\Pi, \alpha_j) - \theta_l \int_0^1 [1 - \mu(\Pi, \alpha_i)] dF_l(\Pi, \alpha_j) \end{aligned}$$

Again, compare this payoff to the equilibrium payoff:

$$V_i(\theta_h, \pi, \alpha_i, \alpha_j) - U_j(b, \theta_h, \pi, \alpha_i, \alpha_j) = \theta_h \int_t^\pi [\mu(\pi, \alpha_i) - \mu(\Pi, \alpha_i)] dF_h(\Pi, \alpha_j) \geqslant 0$$

regardless of $\pi \geqslant t$ or $\pi \leqslant t$. Thus, the strategy given in the proposition is indeed optimal for players if they choose efforts in the equilibrium support.

Now we turn to the case when each valuation type deviates by choosing an effort in the other valuation type's support, e.g. the high valuation type chooses an effort in the support of the low valuation type's support. When the type (θ_h, π) of player i deviates to an effort level β in the low pure support of player j, that is $\beta \in [0, b_{jl}(0, \alpha_j, \alpha_i)]$, then the expected payoff given the opponent playing equilibrium strategy $b_{jl}(\pi, \alpha_j, \alpha_i)$ is:

$$\widetilde{U}_i(\beta|\theta_h, \pi, \alpha_i, \alpha_j) = \theta_h[1 - \mu(\pi, \alpha_i)] \int_{b_{jl}^{-1}(\beta, \alpha_j, \alpha_i)}^1 f_h(\Pi, \alpha_j) d\Pi - \beta.$$

Among all the possible deviating efforts player i would prefer to deviate to the effort level that maximizes the deviation expected payoff, i.e. $\beta^* = \arg\max_\beta \widetilde{U}_i(\beta|\theta_h, \pi, \alpha_i, \alpha_j)$.

β^* can be found by the first order condition with respect to β:

$$f_l(\pi, \alpha_i) = \frac{\theta_l}{\theta_h} \frac{f_l(t, \alpha_j)}{f_h(t, \alpha_j)} f_l(t, \alpha_i) \tag{1.22}$$

where t is given by $b_{jl}(t, \alpha_j, \alpha_i) = \beta^*$, i.e. type (θ_l, t) of player j bids β^* in equilibrium. It is easy to check that both sides of (1.22) are decreasing functions of their arguments, π and t, respectively. Furthermore, Assumption 3 implies $f_l(\pi, \alpha_i) \leqslant f_l(t, \alpha_i)$ and thus, $\pi \geqslant t$. Then, there must exists some $\widehat{\pi}$ satisfying

$$f_l(\widehat{\pi}, \alpha_i) \equiv \frac{\theta_l}{\theta_h} \frac{f_l(0, \alpha_j)}{f_h(0, \alpha_j)} f_l(0, \alpha_i)$$

If the equality in Assumption 3 is satisfied at $\pi = 0$, then $\widehat{\pi} = 0$. For $\pi < \widehat{\pi}$, we always have the LHS of the equation (1.22) strictly larger than the RHS, for all $t \in [0, 1]$. This implies the first order derivative is positive and thus type (θ_h, π) of player i doesn't want to deviate.

For $\pi \geqslant \widehat{\pi}$, there always exists a unique solution of equation (1.22) given π. In this case, we need to directly compare the equilibrium payoff with the payoff of choosing β^*. The difference between the equilibrium expected payoff and the optimal deviation payoff:

$$\begin{aligned} & V_i(\theta_h, \pi, \alpha_i, \alpha_j) - \widetilde{U}_i(\beta^* | \theta_h, \pi, \alpha_i, \alpha_j) \\ = \; & p_h \theta_h \int_0^{\pi} \left[f_h(\pi, \alpha_i) - f_h(\Pi, \alpha_i) \right] f_h(\Pi, \alpha_j) d\Pi \\ + \; & p_l \theta_h \int_0^{t} \left(f_l(\pi, \alpha_i) - \frac{f_l(\Pi, \alpha_i)\theta_l}{f_h(\Pi, \alpha_j)\theta_h} f_l(\Pi, \alpha_j) \right) d\Pi \end{aligned}$$

is increasing with π, as its first order derivative w.r.t π is positive (as $\pi \geqslant t$):

$$\begin{aligned} & \frac{\partial \left(V_i(\theta_h, \pi, \alpha_i, \alpha_j) - \widetilde{U}_i(\beta^* | \theta_h, \pi, \alpha_i, \alpha_j) \right)}{\partial \pi} \\ = \; & p_h \frac{\partial f_h(\pi, \alpha_i)}{\partial \pi} \theta_h \left(\int_0^{\pi} f_h(\Pi, \alpha_j) d\Pi - \int_0^{t} f_h(\Pi, \alpha_j) d\Pi \right) \\ \geqslant \; & 0 \end{aligned}$$

Note that in the above derivation we applied equation (1.22). Since we also know that

$$V_i(\theta_h, \widehat{\pi}, \alpha_i, \alpha_j) - \widetilde{U}_i(\beta^* | \theta_h, \widehat{\pi}, \alpha_i, \alpha_j) = p_h \theta_h \int_0^{\widehat{\pi}} \left[f_h(\widehat{\pi}, \alpha_i) - f_h(\Pi, \alpha_i) \right] f_h(\Pi, \alpha_j) d\Pi > 0$$

the difference is thus, positive. Therefore, type (θ_h, π) of player i does not find it profitable to deviate to any effort in $[0, b_{jl}(0, \alpha_j, \alpha_i)]$.

When a type (θ_l, π) of player i deviates to an effort level β in the high pure support, that is $\beta \in [b_{jl}(0, \alpha_j, \alpha_i), b_{jh}(1, \alpha_j, \alpha_i)]$, the expected payoff given the opponent playing equilibrium strategy $b_{jh}(\pi, \alpha_j, \alpha_i)$ is:

$$\widetilde{U}_i(\beta|\theta_l, \pi, \alpha_i, \alpha_j) = \theta_l \left[\mu(\pi, \alpha_i) \int_0^{b_{jh}^{-1}(\beta, \alpha_j, \alpha_i)} f_l(\Pi, \alpha_j) d\Pi + (1 - \mu(\pi, \alpha_i)) \right] - \beta.$$

Again, we find the optimal deviation effort $\beta^* = \arg\max_\beta \widetilde{U}_i(\beta|\theta_l, \pi, \alpha_i, \alpha_j)$ by the first order condition with respect to β:

$$f_h(\pi, \alpha_i) = \frac{\theta_h}{\theta_l} \frac{f_h(s, \alpha_j)}{f_l(s, \alpha_j)} f_h(s, \alpha_i) \tag{1.23}$$

where s is given by $b_{jh}(s, \alpha_j, \alpha_i) = \beta^*$. It is easy to check that both sides are increasing functions of their arguments, π and s, respectively. Furthermore, Assumption 3 implies $f_h(\pi, \alpha_i) \geqslant f_h(s, \alpha_i)$ and thus, $\pi \geqslant s$. Then, there must be some $\widehat{\widehat{\pi}}$ satisfies

$$f_h(\widehat{\widehat{\pi}}, \alpha_i) = \frac{\theta_h}{\theta_l} \frac{f_h(0, \alpha_j)}{f_l(0, \alpha_j)} f_h(0, \alpha_i)$$

If the equality in condition (3) is satisfied at $\pi = 0$, then we must have $\widehat{\widehat{\pi}} = 0$. For $\pi < \widehat{\widehat{\pi}}$, we always have the LHS of the equation (1.23) strictly less than the RHS, for all $s \in [0, 1]$. This implies the first order derivative is negative and thus type (θ_l, π) doesn't want to deviate.

For $\pi \geqslant \widehat{\widehat{\pi}}$, there always exists a unique valuation of s satisfying equation (1.23) given π. In this case, we need to compare the equilibrium payoff with the payoff of choosing β^*. The difference between the two

$$\begin{aligned}
& V_i(\theta_l, \pi, \alpha_i, \alpha_j) - \widetilde{U}_i(\beta^*|\theta_l, \pi, \alpha_i, \alpha_j) \\
= \; & p_l \theta_l \int_0^\pi \left[f_l(\Pi, \alpha_i) - f_l(\pi, \alpha_i) \right] f_l(\Pi, \alpha_j) d\Pi \\
& + p_h \int_0^s \left[\frac{f_h(\Pi, \alpha_j)\theta_h}{f_l(\Pi, \alpha_j)\theta_l} f_h(\Pi, \alpha_i) - f_h(\pi, \alpha_i) \right] f_l(\Pi, \alpha_j) \theta_l d\Pi
\end{aligned}$$

is positive because its first order derivative w.r.t π is positive:

$$\begin{aligned}&\frac{\partial\Big(V_i(\theta_l,\pi,\alpha_i,\alpha_j)-\widetilde{U}_i(\beta^*|\theta_l,\pi,\alpha_i,\alpha_j)\Big)}{\partial\pi}\\=\;&p_h\theta_l f_h'(\pi,\alpha_i)\left(\int_0^{\pi} f_l(\Pi,\alpha_j)d\Pi-\int_0^{s} f_l(\Pi,\alpha_j)d\Pi\right)\geqslant 0\end{aligned}$$

and $V_i(\theta_l,\widehat{\widehat{\pi}},\alpha_i,\alpha_j)-\widetilde{U}_i(\beta^*|\theta_l,\widehat{\widehat{\pi}},\alpha_i,\alpha_j)=p_l\theta_l\int_0^{\widehat{\widehat{\pi}}}\Big[f_l(\widehat{\widehat{\pi}},\alpha_i)-f_l(\pi,\alpha_i)\Big]f_l(\Pi,\alpha_j)d\Pi>0$. Thus, there is no profitable deviation for any type of player i. □

Proof of Proposition 5

Proof.

Definition 6. *Let $AMR(\eta,\alpha_i,\alpha_j)$ be player i's marginal expected payoff from the contest by choosing η when player j chooses α_j and (wrongly) believes that player i has chosen α_i.*

The above definition of marginal expected payoff from the contest is different from the symmetric case only in the player j's belief of α_i. In the current asymmetric setting, player j chooses α_j but believes that player i chooses α_i. Thus, j plays according to the asymmetric equilibrium given in Proposition 4. By the same logic as in the symmetric case, player i would find it profitable to increase η as it provides a more accurate estimation of j's effort distribution.

First we prove the following lemma:

Lemma 9. *$AMR(\eta,\alpha_i,\alpha_j)>0$ for all $\eta\leqslant\alpha_i$ and $\eta,\alpha_i,\alpha_j\in[\underline{\alpha},\widehat{\alpha})$.*

Given that $AMR(\eta,\alpha_i,\alpha_j)>0$ for any $\eta<\alpha_i\in[\underline{\alpha},\overline{\alpha}]$, it must be true that $AMR(\alpha_i,\alpha_i,\alpha_j)>0$.

If player i chooses $\eta=\underline{\alpha}$, and suppose player j still believes that player i chose α_i and plays the symmetric equilibrium strategy in the contest as given in the Proposition 4, then type (θ_l,π) and (θ_h,π) of player i's maximum expected payoff are given by:

$$
\begin{aligned}
U_i^l(\theta_l,\pi,\underline{\alpha},\alpha_i,\alpha_j) &= \theta_l(1-\mu(\pi,\underline{\alpha}))+\theta_l\int_0^{\widehat{\widehat{\pi}}}\left(\underbrace{\mu(\pi,\underline{\alpha})-\mu(\Pi,\alpha_i)\frac{\theta_h}{\theta_l}\frac{f_h(\Pi,\alpha_j)}{f_l(\Pi,\alpha_j)}}_{>0}\right)f_l(\Pi,\alpha_j)d\Pi \\
U_i^h(\theta_h,\pi,\underline{\alpha},\alpha_i,\alpha_j) &= p_h\underbrace{\int_0^{\widehat{\pi}}\left[1-f_h(\Pi,\alpha_i)\right]\left[f_h(\Pi,\alpha_j)\theta_h-f_l(\Pi,\alpha_j)\theta_l\right]d\Pi}_{>0} \\
&\quad +p_l\theta_l\underbrace{\int_{\widehat{\pi}}^1\left[1-f_l(\Pi,\alpha_i)\right]f_l(\Pi,\alpha_j)d\Pi}_{>0}+p_l\left(\theta_h-\theta_l\right)
\end{aligned}
$$

where $\widehat{\pi}$ satisfies $f_h(\pi,\underline{\alpha})=1=f_h(\widehat{\pi},\alpha_i)$ and $\widehat{\widehat{\pi}}$ is the solution for $\widehat{t}$ in:

$$
f_l(\pi,\underline{\alpha})=1=\frac{f_l(\widehat{t},\alpha_j)}{f_h(\widehat{t},\alpha_j)}\frac{\theta_l}{\theta_h}f_l\left(\widehat{t},\alpha_i\right)
$$

Given that $AMR(\eta,\alpha_i,\alpha_j)>0$, $V_i(\eta,\alpha_i,\alpha_j)$ must be increasing in η. This means

$$
V_i(\alpha_i,\alpha_j)\geqslant U_i(\underline{\alpha},\alpha_i,\alpha_j)>V_i(\underline{\alpha},\underline{\alpha})
$$

where $U_i(\underline{\alpha},\alpha_i,\alpha_j)=\int_0^1 p_lU_i^l(\theta_l,\pi,\underline{\alpha},\alpha_i,\alpha_j)+p_hU_i^h(\theta_h,\pi,\underline{\alpha},\alpha_i,\alpha_j)d\pi$. The first inequality is due to the optimality of α_i. The second inequality is due to $U_i(\underline{\alpha},\alpha_i,\alpha_j)>p_hp_l(\theta_h-\theta_l)$ shown above.

□

Proof of Lemma 7

Proof. The marginal ex ante expected payoff when player i increases α_j is:

$$
\begin{aligned}
&\frac{\partial V_i(\alpha_i,\alpha_j)}{\partial\alpha_j} \\
&= p_h^2\left(\theta_h-\theta_l\right)\int_0^1\left[\underbrace{\frac{\theta_h+\frac{\theta_l}{p_l}}{\theta_h-\theta_l}-f_h(\Pi,\alpha_i)\left(\frac{\theta_h+\frac{p_h\theta_l}{p_l}}{\theta_h-\theta_l}-\Pi\right)-F_h(\Pi,\alpha_i)}_{L(\Pi)}\right]\frac{\partial f_h(\Pi,\alpha_j)}{\partial\alpha_j}d\Pi
\end{aligned}
$$

Note that the terms inside the bracket, $L(\Pi)$, is monotonically decreasing with Π as

$$\frac{\partial L(\Pi)}{\partial \Pi} = -\frac{\partial f_h(\Pi, \alpha_j)}{\partial \Pi}\left(\frac{\theta_h + \frac{p_h \theta_l}{p_l}}{\theta_h - \theta_l} - \Pi\right) < 0$$

Thus, applying the intermediate value theorem for integrals can prove that $\frac{\partial V_i(\alpha_i, \alpha_j)}{\partial \alpha_j} < 0$. Rewrite $\frac{\partial V_i(\alpha_i, \alpha_j)}{\partial \alpha_j}$ for player i into:

$$\begin{aligned} & \int_0^{\pi^+} L(\Pi)\frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi + \int_{\pi^+}^1 L(\Pi)\frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi \\ = \; & L(\zeta)\int_0^{\pi^+} \frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi + L(\iota)\int_{\pi^+}^1 \frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi < 0 \end{aligned}$$

where $\zeta \in [0, \pi^+]$ and $\iota \in [\pi^+, 1]$. Recall that $\int_0^{\pi^+} \frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi + \int_{\pi^+}^1 \frac{\partial f_h(\Pi, \alpha_j)}{\partial \alpha_j} d\Pi = 0$, and since $L(\zeta) > L(\iota)$, we thus have $\frac{\partial V_i(\alpha_i, \alpha_j)}{\partial \alpha_j} < 0$. □

Proof of Lemma 8

Proof. The interim expected payoff for each type of player i in equilibrium are given by (1.10) and (1.11) in the main text. Then, the ex ante interim expected payoff for θ_h and θ_l can be found:

$$\begin{aligned} V_i(\theta_h, \alpha_i, \alpha_j) &= p_h\theta_h \int_0^1 \int_0^{\pi} \left[f_h(\pi, \alpha_i) - f_h(\Pi, \alpha_i)\right] f_h(\Pi, \alpha_j) d\Pi d\pi \\ &\quad + p_l \int_0^1 \int_0^1 \left[f_l(\pi, \alpha_i)\theta_h - f_l(\Pi, \alpha_i)\theta_l\right] f_l(\Pi, \alpha_j) d\Pi d\pi \\ V_i(\theta_l, \alpha_i, \alpha_j) &= p_l\theta_l \int_0^1 \int_{\pi}^1 f_l(\pi, \alpha_j) f_l(\Pi, \alpha_i) d\Pi d\pi - p_l\theta_l \int_0^1 \int_{\pi}^1 f_l(\Pi, \alpha_j) f_l(\Pi, \alpha_i) d\Pi d\pi \end{aligned}$$

And thus, the ex ante expected payoff for player i can be calculated by

$$V_i(\alpha_i, \alpha_j) = p_h V_i(\theta_l, \alpha_i, \alpha_j) + p_l V_i(\theta_h, \alpha_i, \alpha_j)$$

Let $\alpha_i = \underline{\alpha}$, i.e. when the opponent shares no information to player i, then player i's ex ante expected payoff is a constant: $V_i(\alpha_i = \underline{\alpha}, \alpha_j) = p_h p_l (\theta_h - \theta_l)$ □

Proof of Proposition 7

Proof. To prove the proposition, we need to show that the first order conditions for both player 1 and 2 are both satisfied given the strategy profiles specified in the proposition. That is,

$$
\begin{aligned}
&\frac{\partial V_1(\alpha_1,\alpha_2)}{\partial \alpha_2} \\
= \;& p_h^2 (\theta_h - \theta_l) \int_0^1 \left[\frac{\theta_h + \frac{\theta_l}{p_l}}{\theta_h - \theta_l} - f_h(\Pi,\alpha_1) \left(\frac{\theta_h + \frac{p_h \theta_l}{p_l}}{\theta_h - \theta_l} - \Pi \right) - F_h(\Pi,\alpha_1) \right] \frac{\partial f_h(\Pi,\alpha_2)}{\partial \alpha_2} d\Pi = 0 \\
&\frac{\partial V_2(\alpha_2,\alpha_1)}{\partial \alpha_1} \\
= \;& p_h^2 (\theta_h - \theta_l) \int_0^1 \left[\frac{\theta_h + \frac{\theta_l}{p_l}}{\theta_h - \theta_l} - f_h(\Pi,\alpha_2) \left(\frac{\theta_h + \frac{p_h \theta_l}{p_l}}{\theta_h - \theta_l} - \Pi \right) - F_h(\Pi,\alpha_2) \right] \frac{\partial f_h(\Pi,\alpha_1)}{\partial \alpha_1} d\Pi = 0
\end{aligned}
$$

For the symmetric equilibrium, Lemma 8 has shown that given the opponent choosing $\underline{\alpha}$, player i is indifferent between any $\alpha \in [\underline{\alpha}, \widehat{\alpha})$ as she always receives the same expected payoff as in the IPV setting without spying. This suggests $(\underline{\alpha}, \underline{\alpha})$ is a symmetric equilibrium. Lemma 8 also implies the above FOCs are satisfied with $(\underline{\alpha}, \underline{\alpha})$. Now, we prove that this equilibrium is the unique symmetric equilibrium. Suppose both players choose some $\alpha_1 = \alpha_2 = \alpha > \underline{\alpha}$, then Lemma 7 implies, say, player 1 finds it profitable to deviate to $\alpha_2 = \underline{\alpha}$. Thus, there is no symmetric equilibrium in which $\alpha_1 = \alpha_2 = \alpha > \underline{\alpha}$.

Now we turn to asymmetric equilibria. From Lemma 8 and Lemma 7, it can be seen that any strategy profile with one player chooses $\underline{\alpha}$ and the other player chooses $\alpha > \underline{\alpha}$ is an equilibrium. Without loss of generality, assume that player i chooses $\alpha_j = \underline{\alpha}$ and her opponent chooses $\alpha_i > \underline{\alpha}$. By Lemma 7, player i strictly prefers to choose $\alpha_j = \underline{\alpha}$ when $\alpha_i > \underline{\alpha}$. By Lemma 8, player j is indifferent with choosing any $\alpha_i \in [\underline{\alpha}, \widehat{\alpha})$, thus he has no profitable deviation. Therefore, any strategy profile (α_1, α_2) with either $\alpha_1 = \underline{\alpha}$ and $\alpha_2 > \underline{\alpha}$, or $\alpha_1 > \underline{\alpha}$ and $\alpha_2 = \underline{\alpha}$, is an equilibrium of the information disclosure game. $\square$

Chapter 2

Information Disclosure in Contests: Private vs. Public Signals

Abstract Two players with independent private valuations compete in the first-price all-pay auction. Apart from each player's own valuation, she also observes a noisy signal regarding the opponent's valuation. We characterize the unique symmetric equilibrium of the contest when the signal is (1) conditionally independent private, or (2) public. In the former, each player's expected payoff (expected effort) is always higher (lower) than when they do not receive any signals regarding the opponent. In the latter, the expected payoff (expected effort) can be either higher or lower than when they do not receive any signals, and the upper bound of expected payoffs is characterized. Numerical examples show that some private signals may induce higher expected payoffs than the maximum payoff induced by any public signals, and that some public signals can always induce higher total expected efforts than the maximum expected efforts induced by any private signals.

2.1 Introduction

In contests, e.g., rent-seeking, lobbying, promotion, and patent race, the amount of effort to be invested often hinges on the player's belief about the opponents' competitiveness. The contest organizer can thus manipulate the belief via disclosing information about the opponents to manipulate competitive behavior indirectly. The objective of the organizer may be to increase total effort. For example, the employer who organizes a promotion contest wish to boost employees' efforts. Alternatively, the objective may also be to increase players' payoffs. For instances, the objective of a trade association is to protect the interests of member firms who compete in the same industry.

In this paper, we study how the organizer should disclose information before contests to fulfill the above objectives. The contest is modeled by the first-price all-pay auction with two players and one prize. Players' private valuations for the prize are independently drawn from a binary distribution, $\{v_h, v_l\}$ with $v_h > v_l$. In addition to their own valuations, they can also observe a binary distributed signal ($\{h, l\}$) disclosed by the contest organizer and which contains information about the opponent's valuation.We consider two classes of disclosure policies: disclosing via conditionally independent private (hereafter private) signals or via public signals. In the former, each player receives a signal conditional on the opponent's valuation and the signal realization is private information to the player. In particular, with probability q the signal is "correct", i.e. $\Pr(s_i = h|v_{-i} = v_h) = \Pr(s_i = l|v_{-i} = v_l) = q$. In the latter, both players observe a public signal conditional on the valuation profile. In particular, when both players have v_h (v_l), they observe the signal h with probability k_h (k_l); when players have different valuations they observe the signal h with probability r. For each class of disclosure policies, the contest organizer chooses these parameters before the contest to fulfill her objectives.

We characterized unique symmetric equilibrium in the contest under both classes of disclosure policies. The equilibrium strategy may be non-monotonic in the sense that the high valuation player's effort may be lower than that of the low valuation player. With partially informative private signals, on one hand, each player's expected payoff is always higher than when they do not receive any signals about the opponent, and the total expected effort is always lower. The public signals, on the other hand, can be designed to maximize players' expected payoff and its upper bound is also strictly higher than when players do not receive any signals about the opponents. Numerical examples suggest that there is no general ranking between the two classes of signals in maximizing players' expected payoffs. The organizer can also design the public signals in a way that the total expected effort is strictly higher than when players do not receive any signals – thus it is also higher than the total effort in the contest with any private signals. Therefore, some public signals dominate all private signals in terms of increasing total expected effort.

The prior literature has shown that the total effort in the contest can be boosted by concealing all players' private information in all-pay auctions (Fu et al., 2014; Lu et al., 2016), partially revealing such information to the opponents in Tullock contest (Serena, 2015), providing reviews (Gershkov and Perry, 2009) of previous performance or publicly announcing (Aoyagi, 2010) it in multi-stage contests, disclosing opponent's previous performance (Sheremeta, 2010) or rent-seeking expenditure (Fallucchi et al., 2013) in rent-seeking contests. The study closely related to the current paper is Lu et al. (2016) who extend the study on the partial information disclosure in Tullock contests (Serena, 2015) to the first-price all-pay auction. According to Serena (2015), the partial disclosure policy is a mapping from the anonymous[1] valuation profile, $\{v_h, v_h\}$, $\{v_h, v_l\}$ and $\{v_l, v_l\}$, to a binary decision between *Concealing* or *Disclosing* the profile to both players, i.e. C or D. For example, $\{C, C, D\}$ corresponds to the disclosure policy which conceals the valuation profile $\{v_h, v_h\}$, i.e. both players have high valuation, and $\{v_h, v_l\}$, i.e. when the players have different valuation, and discloses the profile only when it is $\{v_l, v_l\}$, i.e. both players have low valuation. Such partial disclosure policies are special cases of public signals considered in the Section 2.4 of the current paper. Lu et al. (2016) shows that the disclosure policy $\{C, C, D\}$ maximizes each player's expected payoff, and the maximum is $\min\{p_h(v_h - v_l), p_h p_l v_h\}$. Interestingly, this coincides with the maximum of expected payoffs characterized in the current paper. However, the current paper shows that there exists a broad set of public signals which can induce this maximum. When the objective is to maximize total effort, Lu et al. (2016) shows that fully concealing the valuation profile, i.e. $\{C, C, C\}$, maximizes the total expected effort, whereas Section 2.4 of the current paper shows that some public signals can induce strictly higher efforts.

Fang and Morris (2006) was the first to consider the setting in which players receive conditional independent private signals about opponents in winner-pay auctions. In the first-price auction with binary distribution of valuations, the low valuation type of player always bid his valuation and the high valuation type of player's bid increases first order stochastically in her signal regarding the opponent. The authors show that the revenue in the first-price auction is lower than in the second-price auction in which players still bid their own valuations in equilibrium

[1] The disclosure policy is anonymous in the sense that the policy depends on the type profile that does not differentiate the identities of players. See Serena (2015) for more details.

even when they receive a signal about the opponent. Azacis and Vida (2015) then generalized the model to any correlated signals and characterized the optimal signal which maximizes player's expected payoff. This is closely related to the second part of the current paper which considers public signal. Also built on Fang and Morris (2006), Tian and Xiao (2007) studied a model in which players endogenously acquire costly information about their opponent in the first-price auction.

In winner-pay auctions, auctioneers may find it profitable to selectively disclose information to bidders (Li and Shi, 2013; Eso and Szentes, 2007; Ganuza and Penalva, 2010; Bergemann and Pesendorfer, 2007). In the current paper, we show how a contest organizer can influence the competitive behavior of players via disclosing a signal regarding the opponents' private information. This type of information disclosure in contests has been studied in one-sided private information setting (Zhang and Zhou, 2016; Denter et al., 2014), or two-sided private information with disclosure policies (Lu et al., 2016; Serena, 2015) which are special cases of the public signals considered in the current paper.

Finally, the Bayesian persuasion literature (Kamenica and Gentzkow, 2011) has studied how a sender should design signals which reveal information to voters (Alonso and Câmara, 2016, 2015; Wang, 2013), or to consumers (Anderson and Renault, 2006; Johnson and Myatt, 2006; Rayo and Segal, 2010), about their payoff relevant states in order to influence their behavior.

2.2 Preliminaries

The Contest: Two players compete for an indivisible prize in a contest. Player i's ($i \in \{1,2\}$) private valuation is independently drawn from the binary distribution: $v_i = v_h$ with probability $p_h \in (0,1)$, and $v_i = v_l$ with probability $p_l \in (0,1)$, where $v_h > v_l > 0$ and $p_h + p_l = 1$. In addition to her own valuation, player i also observes an additional signal, $s_i \in \{h, l\}$, regarding her opponent's valuation v_{-i}. The distribution of the signal will be discussed in detail shortly.

Players choose their efforts, (b_i, b_{-i}), simultaneously. The player who chooses higher effort wins and both players incur the costs of their own efforts. Ties are broken with equal probability. Thus, the contest is equivalent to the first-price all-pay auction. The player with the valuation v_i chooses effort b_i earns the following expected payoff:

$$U_i(b_i, b_{-i}, v_i) = \begin{cases} -b_i, \text{ if } b_i < b_{-i} \\ v_i - b_i, \text{ if } b_i > b_{-i} \\ \frac{1}{2}v_i - b_i, \text{ if } b_i = b_{-i} \end{cases}$$

Equilibrium: In the contest as a Bayesian game, each player's type has two dimensions, i.e. valuation and signal: (v_i, s_i). Denote by $G_{(v_i,s_i)}(b)$ the cumulative distribution function of effort in the equilibrium mixed strategy of type (v_i, s_i) of player i, and denote by $G_i(b)$ player i's ex ante cumulative distribution of equilibrium effort. Formally, a Bayesian Nash Equilibrium is defined as the following.

Definition 7. *A Bayesian Nash Equilibrium (BNE) of the contest is a vector of strategies $\boldsymbol{G} = (G_1, G_2)$ such that for all $b_i \in supp[G_{(v_i,s_i)}]$, we have*

$$b_i \in \arg\max_b U_i(b, v_i, s_i; G_{-i})$$

In Section 2.3, we consider private signals where the distribution of s_i is contingent on realizations of the opponent's valuation v_{-i}. In Section 2.4, we consider public signals where the distribution of s_i is contingent on realizations of valuation profile (v_i, v_{-i}). Finally, we compare the two signals in terms of players' expected payoff and total expected effort in Section 2.5.

2.3 Private signals

In the private signals setting, player i's signal is generated as the following:

$$\begin{aligned} Pr(s_i = l|v_{-i} = v_l) = Pr(s_i = h|v_{-i} = v_h) &= q \in [\frac{1}{2}, 1] \\ Pr(s_i = h|v_{-i} = v_l) = Pr(s_i = l|v_{-i} = v_h) &= 1 - q. \end{aligned}$$

That is, q is the probability that player i receives a "correct" signal. The signal s_i is player i's private information. Thus, the type space becomes two dimensional with four types in total: $(v_i, s_i) \in \{v_h, v_l\} \times \{h, l\}$. Denote by $\Pr(v_{-i}|s_i)$ the probability that the opponent's valuation is v_{-i} conditional on player i's signal s_i. Upon receiving a signal s_i, player i updates her belief according to Baye's rule:

$$\begin{aligned} \Pr(v_h|h) &= \frac{p_h q}{p_h q + p_l (1-q)} \\ \Pr(v_l|h) &= \frac{p_l (1-q)}{p_h q + p_l (1-q)} \\ \Pr(v_h|l) &= \frac{p_h (1-q)}{p_h (1-q) + p_l q} \\ \Pr(v_l|l) &= \frac{p_l q}{p_h (1-q) + p_l q} \end{aligned}$$

To facilitate later analyzes, we define the following condition:

Condition 1. $\Pr(v_l|h)v_h \geqslant \Pr(v_l|l)v_l$*;*

Alternatively, we refer to the condition with the opposite inequality, i.e., $\Pr(v_l|h)v_h \leqslant \Pr(v_l|l)v_l$, as the Condition $\neg 1$.[2]

In the above conditions, Condition 1 is equivalent of $(1-q)v_h \geqslant qv_l$, i.e., $q \leqslant q^* \equiv \frac{v_h}{v_h+v_l}$, and Condition $\neg 1$ is equivalent of $(1-q)v_h \leqslant qv_l$, i.e., $q \geqslant q^* \equiv \frac{v_h}{v_h+v_l}$. Condition 1, in fact, is the sufficient condition for the existence of monotonic strategy equilibrium (MSE). In all-pay auctions with one dimensional affiliated signals, the existence of MSE depends on whether a "monotonicity condition" is satisfied, which states that the product of the conditional probability of the opponent's signal and the player's valuation increases in the player's signal (Rentschler and Turocy, 2016; Chi et al., 2015; Krishna and Morgan, 1997). In such single dimensional settings, players obtain information regarding both their own valuation and the opponent's valuation from one signal. In the two dimensional signal setting considered in the current paper, however, players obtain the two sorts of information from separated channels. Condition 1 therefore, is the two dimensional version of "monotonicity condition".

2.3.1 Equilibrium

The following result shows that the structure of equilibrium depends on whether Condition 1 or $\neg 1$ is satisfied.

Proposition 8. *If $q \in [\frac{1}{2}, 1]$, then there exists a unique symmetric equilibrium in which all types randomize over connected supports. When the Condition 1 is satisfied, then*

- *type (v_l, h) mixes over $[0, \bar{b}_{(v_l,h)}]$ uniformly according to*

$$G_{(v_l,h)}(b) = \frac{p_l(1-q)+p_hq}{p_l(1-q)^2v_l}b$$

- *type (v_l, l) mixes over $[\bar{b}_{(v_l,h)}, \bar{b}_{(v_l,l)}]$ uniformly according to*

$$G_{(v_l,l)}(b) = \frac{p_h(1-q)+p_lq}{p_lq^2v_l}b$$

- *type (v_h, l) mixes over $[\bar{b}_{(v_l,l)}, \bar{b}_{(v_h,l)}]$ uniformly according to*

$$G_{(v_h,l)}(b) = \frac{p_h(1-q)+p_lq}{p_h(1-q)^2v_h}b$$

[2]Note that Condition 1 and $\neg 1$ are equivalent when the equalities are satisfied simultaneously.

- *type* (v_h, h) *mixes over* $[\bar{b}_{(v_h,l)}, \bar{b}_{(v_h,h)}]$ *uniformly according to*

$$G_{(v_h,h)}(b) = \frac{p_l(1-q) + p_h q}{p_h q^2 v_h} b$$

where

$$\begin{aligned}
\bar{b}_{(v_l,h)} &= \frac{p_l(1-q)^2 v_l}{p_l(1-q) + p_h q} \\
\bar{b}_{(v_l,l)} &= \bar{b}_{(v_l,h)} + \frac{p_l q^2 v_l}{p_h(1-q) + p_l q} \\
\bar{b}_{(v_h,l)} &= \bar{b}_{(v_l,l)} + \frac{p_h(1-q)^2 v_h}{p_h(1-q) + p_l q} \\
\bar{b}_{(v_h,h)} &= \bar{b}_{(v_h,l)} + \frac{p_h q^2 v_h}{p_l(1-q) + p_h q}
\end{aligned}$$

When the Condition ¬1 is satisfied, then

- *type* (v_h, h) *mixes over* $[\bar{b}, \bar{b}_{(v_h,h)}]$ *uniformly according to*

$$G_{(v_h,h)}(b) = \frac{p_h q + p_l(1-q)}{p_h q^2 v_h} b - \frac{p_h q v_h + p_l v_l}{p_h q^2 v_h} (1-q),$$

  and mixes over $[\underline{b}, \bar{b}]$ *according to*

$$G_{(v_h,h)}(b) = \frac{1}{2q-1}\left(\frac{q}{v_h} - \frac{1-q}{v_l}\right) b - \frac{1-q}{2q-1} \frac{p_l(1-q)}{p_h q + p_l(1-q)} \frac{q v_l - (1-q) v_h}{q v_h - (1-q) v_l} \frac{v_h - v_l}{v_h}$$

- *type* (v_h, l) *and* (v_l, l) *mix over* $[\underline{b}, \bar{b}]$ *uniformly according to*

$$G_{(v_h,l)}(b) = G_{(v_l,l)}(b) = \frac{1}{2q-1}\left(\frac{q}{v_l} - \frac{1-q}{v_h}\right) b - \frac{1-q}{2q-1} \frac{p_l(1-q)}{p_h q + p_l(1-q)} \frac{v_h - v_l}{v_h}$$

- *type* (v_l, h) *mixes over* $[\underline{b}, \bar{b}]$ *uniformly according to*

$$\begin{aligned}
G_{(v_l,h)}(b) &= \frac{1}{2q-1}\left(\frac{q}{v_h} - \frac{1-q}{v_l}\right) b \\
&\quad - \frac{1-q}{2q-1} \frac{p_l(1-q)}{p_h q + p_l(1-q)} \frac{q v_l - (1-q) v_h}{q v_h - (1-q) v_l} \frac{v_h - v_l}{v_h} - \frac{v_h - v_l}{(1-q) v_l - q v_h}
\end{aligned}$$

  and mixes over $[0, \underline{b}]$ *according to*

$$G_{(v_l,h)}(b) = \frac{p_h q + p_l(1-q)}{p_l (1-q)^2 v_l} b$$

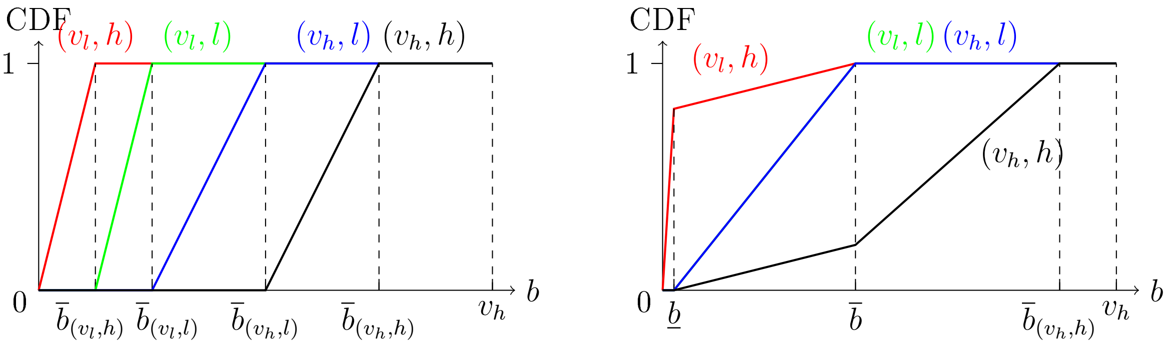

Figure 2.1: Example: $p_h = \frac{1}{2}, v_h = 2, v_l = 1$ thus $q^* = \frac{2}{3}$. The non-overlapping equilibrium is given in the left panel, where $q = \frac{1}{2}$; the overlapping equilibrium is given in the right panel, where $q = \frac{3}{4}$.

where

$$
\begin{aligned}
\bar{b}_{(v_h,h)} &= \frac{p_h q v_h + p_l(1-q)v_l}{p_h q + p_l(1-q)} \\
\bar{b} &= \frac{q\left[p_h q + (p_l - p_h)(1-q)\right] v_h - (1-q)^2 p_l v_l}{\left[p_h q + p_l(1-q)\right]\left[q v_h - (1-q)v_l\right]} v_l \\
\underline{b} &= \frac{p_l (1-q)}{p_h q + p_l(1-q)} \frac{(1-q)\, v_l}{q v_h - (1-q)v_l} (v_h - v_l)
\end{aligned}
$$

Specifically, if Condition 1 is true, the model has a unique non-overlapping equilibrium in the sense that types play mixed strategies distributed on contagious non-overlapping intervals, similar to the independent private valuation setting (Konrad, 2004). If, however, Condition ¬1 is true, the model has an overlapping equilibrium in the sense that the supports of all types intersect at a common interval which becomes a singleton when Condition ¬1 satisfies with equality. This is consistent with complete information setting (Baye et al., 1996) which corresponding to the case when $q = 1$. See Example 1 for the graphical structure of the equilibrium.

Example 1. *Suppose $p_h = \frac{1}{2}, v_h = 2, v_l = 1$ thus $q^* = \frac{2}{3}$. Thus, when $q = \frac{1}{2} < q^*$ the equilibrium supports of all types are non-overlapping, and the upper bounds are given by: $\bar{b}_{(v_l,h)} = 0.25$, $\bar{b}_{(v_l,l)} = 0.5$, $\bar{b}_{(v_h,l)} = 1$ and $\bar{b}_{(v_h,h)} = 1.5$. This equilibrium is shown in the left panel of Figure 2.1. When $q = \frac{3}{4} > q^*$, the equilibrium supports of all types are overlapping, and the bounds of supports are given by: $\underline{b} = 0.05$, $\bar{b} = 0.85$ and $\bar{b}_{(v_h,h)} = 1.75$. The overlapping equilibrium is shown in the right panel of Figure 2.1.*

2.3.2 Expected payoff and total expected effort

Proposition 9 indicates that both players are better off with any partially informative private signals.

Proposition 9. *Player i ($i = 1, 2$) earns strictly higher expected payoff when $q \in (\frac{1}{2}, 1)$ than when $q = \frac{1}{2}, 1$; the total expected effort is strictly lower when $q \in (\frac{1}{2}, 1)$ than when $q = \frac{1}{2}, 1$.*

Figures 2.2 and 2.3 are two examples of player i's expected payoff and the total expected effort. As can be seen from the figures, the difference between overlapping and non-overlapping equilibrium has significant impact on the expected effort. The total expected effort experiences a sudden drop at $q = q^* \equiv \frac{v_h}{v_h + v_l}$ during the transition from the non-overlapping to the overlapping equilibrium. The expected payoff is, however, continuous at the cutoff value of q.

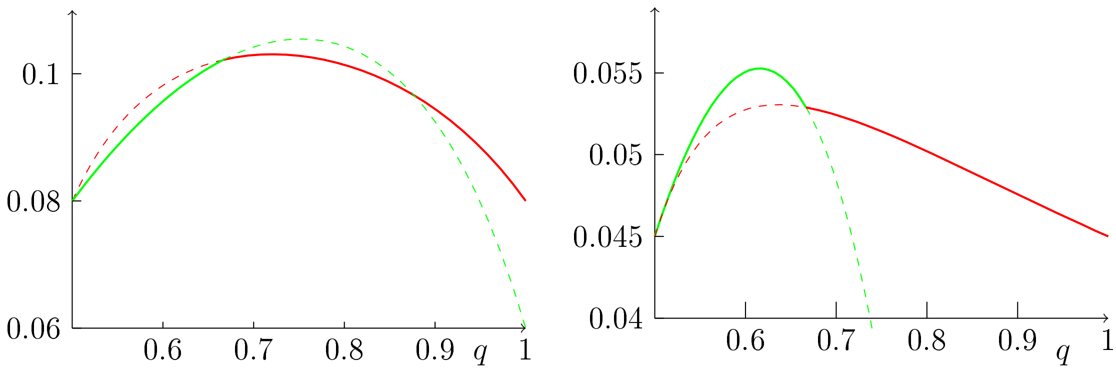

Figure 2.2: Expected payoff when $v_h = 1$ and $v_l = 0.5$: left $p_h = 0.2$; right $p_h = 0.9$

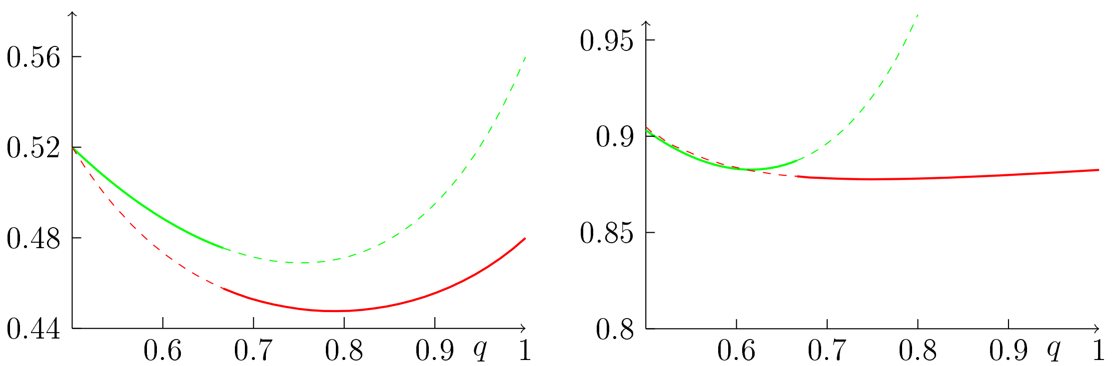

Figure 2.3: Revenue/Total expenditure when $v_h = 1$ and $v_l = 0.5$: left $p_h = 0.2$; right $p_h = 0.9$

The above example is consistent with Morath and Münster (2008) in the sense that the total expected effort is lower when $q = 1$ than when $q = \frac{1}{2}$. The fact that players earn higher expected payoff with $q \in (\frac{1}{2}, 1)$ suggests the total expected effort when $q \in (\frac{1}{2}, 1)$ must be lower than when $q = \frac{1}{2}$, i.e., when players receive uninformative signals about opponents.

2.4 Public signals

Now we turn to public signals. We focus only on the symmetric distribution of public signals, which are generated as the following:

$$\begin{aligned} k_h &= \Pr(s_1 = s_2 = h | v_1 = v_2 = v_h) \\ k_l &= \Pr(s_1 = s_2 = l | v_1 = v_2 = v_l) \\ r &= \Pr(s_1 = s_2 = h | v_i \neq v_{-i}) \end{aligned}$$

where $k_h, k_l, r \in [0, 1]$. k_h (k_l) is the probability that a high (low) valuation player receives the signal realization h (l) conditional on both players having high (low) valuation. r is the probability that a player receives the signal realization h when players have different valuations. This is summarized in Table 2.1. We refer to the vector (k_h, k_l, r) as the "public signal (k_h, k_l, r)".

	(v_h, h)	(v_h, l)	(v_l, h)	(v_l, l)
(v_h, h)	$p_h^2 k_h$	0	$p_h p_l r$	0
(v_h, l)	0	$p_h^2(1 - k_h)$	0	$p_h p_l(1 - r)$
(v_l, h)	$p_h p_l r$	0	$p_l^2(1 - k_l)$	0
(v_l, l)	0	$p_h p_l(1 - r)$	0	$p_l^2 k_l$

Table 2.1: Public signal

Denote by $\Pr(v_{-i}|v_i, s_i)$ the probability that player $-i$ has value v_{-i} conditional on player i has value v_i and receives signal s_i. Note that $\Pr(v_{-i}|v_i, s_i) = \Pr(v_{-i}, s_{-i}|v_i, s_i)$, e.g. $\Pr(v_h, h|v_l, h) = \Pr(v_h|v_l, h)$, since the signal realization is common knowledge. Thus, the conditional probabilities can be written down as:

$$\begin{aligned}
\Pr(v_h|v_l,h) &= \frac{p_h r}{p_h r + p_l\left(1-k_l\right)} \\
\Pr(v_h|v_h,h) &= \frac{p_h k_h}{p_h k_h + p_l r} \\
\Pr(v_h|v_l,l) &= \frac{p_h\left(1-r\right)}{p_h\left(1-r\right) + p_l k_l} \\
\Pr(v_h|v_h,l) &= \frac{p_h\left(1-k_h\right)}{p_h\left(1-k_h\right) + p_l\left(1-r\right)}
\end{aligned}$$

In order to determine the equilibrium strategies for types of players who receive signal h, we define the Conditions 2 and 3:

Condition 2. $\Pr(v_l|v_h,h)v_h \geqslant \Pr(v_l|v_l,h)v_l$;

Condition 3. $\Pr(v_h|v_h,h)v_h \geqslant \Pr(v_h|v_l,h)v_l$;

To determine the equilibrium strategies for types (v_h,l) and (v_l,l), we define the Conditions 4 and 5:

Condition 4. $\Pr(v_l|v_h,l)v_h \geqslant \Pr(v_l|v_l,l)v_l$;

Condition 5. $\Pr(v_h|v_h,l)v_h \geqslant \Pr(v_h|v_l,l)v_l$

All the above conditions can again be understood in analogous to the "monotonicity condition" in the previous literature. Indeed, as shown in Proposition 10, Conditions 2-5 ensures the existence of monotonic strategy equilibrium in which the player with high valuation randomizes in the support higher than and non-overlapping with the support of the low valuation player. In addition, we also characterized the unique equilibrium when Conditions ¬2-¬5 are satisfied, i.e., when the directions of inequalities in Conditions 2-5 are reversed.[3]

Before going into the next section, note that Conditions ¬2 and ¬3 cannot be both satisfied, as Condition ¬2 implies

$$\Pr(v_h|v_h,h)v_h = v_h - \Pr(v_l|v_h,h)v_h \geqslant v_h - \Pr(v_l|v_l,h)v_l > \Pr(v_h|v_l,h)v_l$$

which is in contradiction to Condition ¬3. Similarly, Conditions ¬4 and ¬5 cannot be both satisfied, as Condition ¬4 implies

$$\Pr(v_h|v_h,l)v_h = v_h - \Pr(v_l|v_h,l)v_h \geqslant v_h - \Pr(v_l|v_l,l)v_l > \Pr(v_h|v_l,l)v_l$$

[3]For instances, Condition ¬2 is $\Pr(v_l|v_h,h)v_h \leqslant \Pr(v_l|v_l,h)v_l$, and Condition ¬5 is $\Pr(v_h|v_h,l)v_h \leqslant \Pr(v_h|v_l,l)v_l$.

which contradicts Condition ¬5.

Therefore, in the equilibrium analysis we need to consider the cases when the following conditions are satisfied: Conditions 2 and 3 (4 and 5), Conditions 2 and ¬3 (4 and ¬5), Conditions ¬2 and 3 (¬4 and 5).

2.4.1 Equilibrium

Proposition 10. *When players receive the public signal (k_h, k_l, r), the unique equilibrium is symmetric, and all types randomizes over connected supports.*
Specifically, for type (v_h, h) and (v_l, h):

- *If Conditions 2 and 3 are satisfied, then type (v_l, h) mixes over $[0, \bar{b}_{(v_l,h)}]$ and (v_h, h) mixes over $[\bar{b}_{(v_l,h)}, \bar{b}_{(v_h,h)}]$ according to CDF $G_{(v_l,h)}(b)$ and $G_{(v_h,h)}(b)$, respectively:*

$$\begin{aligned} G_{(v_l,h)}(b) &= \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l} b \\ G_{(v_h,v_h)}(b) &= \frac{p_h k_h + p_l r}{p_h k_h v_h} b - \frac{v_l}{v_h} \frac{p_l(1-k_l)}{p_h k_h} \frac{p_h k_h + p_l r}{p_h r + p_l(1-k_l)} \end{aligned}$$

where $\bar{b}_{(v_l,h)} = \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} v_l$ and $\bar{b}_{(v_h,h)} = \frac{p_h k_h}{p_h k_h + p_l r} v_h + \bar{b}_{(v_l,h)}$.

- *If Conditions 2 and ¬3 are satisfied, then type (v_h, h) mixes over $[\underline{b}_{(v_h,h)}, v_l]$ according to CDF $G_{(v_h,h)}(b)$:*

$$\begin{aligned} G_{(v_h,h)}(b) &= \frac{r(p_h r + p_l(1-k_l))v_h - (1-k_l)(p_h k_h + p_l r)v_l}{p_h(r^2 - k_h(1-k_l))v_h v_l} b \\ &\quad - (1-k_l)\frac{p_l r + p_h k_h}{p_h(r^2 - k_h(1-k_l))v_h}(v_h - v_l) \end{aligned}$$

while type (v_l, h) mixes over $[\underline{b}_{(v_h,h)}, v_l]$ according to CDF $G_{(v_l,h)}(b)$:

$$\begin{aligned} G_{(v_l,h)}(b) &= \frac{-k_h(p_h r + p_l(1-k_l))v_h + r(p_l r + p_h k_h)v_l}{p_l(r^2 - k_h(1-k_l))v_h v_l} b \\ &\quad + r\frac{p_l r + p_h k_h}{p_l(r^2 - k_h(1-k_l))v_h}(v_h - v_l) \end{aligned}$$

and mixes over $[0, \underline{b}_{(v_h,h)}]$ according to CDF $G_{(v_l,h)}(b)$:

$$G_{(v_l,h)}(b) = \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l} b$$

where $\underline{b}_{(v_h,h)} = \frac{(1-k_l)(p_l r + p_h k_h)}{r(p_h r + p_l(1-k_l))v_h - (1-k_l)(p_l r + p_h k_h)} v_l (v_h - v_l) v_l$.

- *If Conditions ¬2 and 3 are satisfied, then type* (v_h, h) *mixes over* $[0, \overline{b}_{(v_l,h)}]$ *according to CDF* $G_{(v_h,h)}(b)$*:*

$$G_{(v_h,h)}(b) = \frac{(1-k_l)(p_h k_h + p_l r) v_l - r(p_h r + p_l(1-k_l)) v_h}{p_h(k_h(1-k_l) - r^2) v_l v_h} b$$

and mixes over $[\overline{b}_{(v_l,h)}, v_h]$ *according to CDF* $G_{(v_h,h)}(b)$*:*

$$G_{(v_h,h)}(b) = \frac{p_h k_h + p_l r}{p_h k_h v_h} b - \frac{p_l r}{p_h k_h}$$

while type (v_l, h) *mixes over* $[0, \overline{b}_{(v_l,h)}]$ *according to* $G_{(v_l,h)}(b)$*:*

$$G_{(v_l,h)}(b) = \frac{k_h(p_l(1-k_l) + p_h r) v_h - r(p_l r + p_h k_h) v_l}{p_l(k_h(1-k_l) - r^2) v_h v_l} b$$

where $\overline{b}_{(v_l,h)} = \frac{p_l\left(k_h(1-k_l)-r^2\right)v_h v_l}{k_h(p_l(1-k_l)+p_h r)v_h - r(p_l r + p_h k_h)v_l}$

For type (v_h, l) *and* (v_l, l)*:*

- *If Conditions 4 and 5 are satisfied, then type* (v_l, l) *mixes over* $[0, \overline{b}_{(v_l,l)}]$ *and type* (v_h, l) *mixes over* $[\overline{b}_{(v_l,l)}, \overline{b}_{(v_h,l)}]$ *according to CDF* $G_{(v_l,l)}(b)$ *and* $G_{(v_h,l)}(b)$*, respectively:*

$$\begin{aligned} G_{(v_l,l)}(b) &= \frac{p_h(1-r) + p_l k_l}{p_l k_l v_l} b \\ G_{(v_h,l)}(b) &= \frac{p_h(1-k_h) + p_l(1-r)}{p_h(1-k_h) v_h} b - \frac{v_l}{v_h} \frac{p_l k_l}{p_h(1-k_h)} \frac{p_h(1-k_h) + p_l(1-r)}{p_h(1-r) + p_l k_l} \end{aligned}$$

where $\overline{b}_{(v_l,l)} = \frac{p_l k_l}{p_h(1-r)+p_l k_l} v_l$ *and* $\overline{b}_{(v_h,l)} = \frac{p_h(1-k_h)}{p_h(1-k_h)+p_l(1-r)} v_h + \overline{b}_{(v_l,l)}$.

- *If Conditions 4 and ¬5 are satisfied, then type* (v_h, l) *mixes over* $[\underline{b}_{(v_h,l)}, v_l]$ *according to CDF* $G_{(v_h,l)}(b)$*:*

$$\begin{aligned} G_{(v_h,l)}(b) &= \frac{(1-r)(p_h(1-r) + p_l k_l) v_h - k_l(p_h(1-k_h) + p_l(1-r)) v_l}{p_h((1-r)^2 - (1-k_h) k_l) v_h v_l} b \\ &\quad - k_l \frac{p_l(1-r) + p_h(1-k_h)}{p_h((1-r)^2 - (1-k_h) k_l) v_h} (v_h - v_l) \end{aligned}$$

while type (v_l, l) *mixes over* $[\underline{b}_{(v_h,l)}, v_l]$ *according to CDF* $G_{(v_l,l)}(b)$

$$\begin{aligned} G_{(v_l,l)}(b) &= \frac{-(1-k_h)(p_h(1-r) + p_l k_l) v_h + (1-r)(p_l(1-r) + p_h(1-k_h)) v_l}{p_l((1-r)^2 - (1-k_h) k_l) v_h v_l} b \\ &\quad + (1-r) \frac{p_l(1-r) + p_h(1-k_h)}{p_l((1-r)^2 - (1-k_h) k_l) v_h} (v_h - v_l) \end{aligned}$$

and mixes over $[0, \underline{b}_{(v_h,l)}]$ *according to CDF* $G_{(v_l,l)}(b)$

$$G_{(v_l,l)}(b) = \frac{p_h(1-r)+p_l k_l}{p_l k_l v_l} b,$$

where $\underline{b}_{(v_h,l)} = \frac{k_l(p_h(1-k_h)+p_l(1-r))}{(1-r)(p_l k_l + p_h(1-r))v_h - k_l(p_h(1-k_h)+p_l(1-r))v_l}\left(v_h - v_l\right) v_l.$

- *If Conditions* $\neg 4$ *and 5 are satisfied, then type* (v_h, l) *mixes over* $[0, \overline{b}_{(v_l,l)}]$ *according to CDF* $G_{(v_h,l)}(b)$*:*

$$G_{(v_h,l)}(b) = \frac{k_l\left(p_h(1-k_h)+p_l(1-r)\right)v_l - (1-r)\left(p_h(1-r)+p_l k_l\right)v_h}{p_h\left((1-k_h)k_l - (1-r)^2\right)v_l v_h} b$$

and mixes over $[\overline{b}_{(v_l,l)}, v_h]$ *according to CDF* $G_{(v_h,l)}(b)$*:*

$$G_{(v_h,l)}(b) = \frac{p_h(1-k_h)+p_l(1-r)}{p_h(1-k_h)v_h} b - \frac{p_l(1-r)}{p_h(1-k_h)}$$

while type (v_l, l) *mixes over* $[0, \overline{b}_{(v_l,l)}]$ *according to CDF* $G_{(v_l,h)}(b)$*:*

$$G_{(v_l,h)}(b) = \frac{k_h\left(p_h(1-r)+p_l k_l\right)v_h - r\left(p_h(1-k_h)+p_l(1-r)\right)v_l}{p_l\left((1-k_h)k_l - (1-r)^2\right)v_h v_l} b$$

where $\overline{b}_{(v_l,l)} = \frac{p_l((1-k_h)k_l-(1-r)^2)v_h v_l}{(p_l k_l + p_h(1-r))(1-k_h)v_h - (1-r)(p_h(1-k_h)+p_l(1-r))v_l}.$

See the following for a numerical example for the structure of the equilibrium when players receive "h". The structure of the equilibrium when players receive "l" is qualitatively the same.

Example 2. *Suppose* $p_h = \frac{1}{2}, v_h = 2, v_l = 1$. *When* $(k_h, k_l, r) = (\frac{2}{3}, \frac{2}{3}, \frac{1}{3})$*, then Conditions 2 and 3 are satisfied. Thus, the corresponding equilibrium mixed strategy is given in the left panel of the Figure 2.4. When* $(k_h, k_l, r) = (\frac{1}{10}, \frac{2}{3}, \frac{1}{3})$*, then Conditions 2 and* $\neg 3$ *are satisfied. Thus, the corresponding equilibrium mixed strategy is given in the middle panel of the Figure 2.4. When* $(k_h, k_l, r) = (\frac{2}{3}, \frac{2}{3}, \frac{1}{10})$*, then Conditions* $\neg 2$ *and 3 are satisfied. Thus, the corresponding equilibrium mixed strategy is given in the right panel of the Figure 2.4.*

2.4.2 Optimal public signal

Given the equilibrium strategy we can identify the optimal information structure which maximizes each player's expected payoff.

Proposition 11. *Player i's maximum expected payoff is* $\min\{p_h(v_h - v_l), p_h p_l v_h\}$.

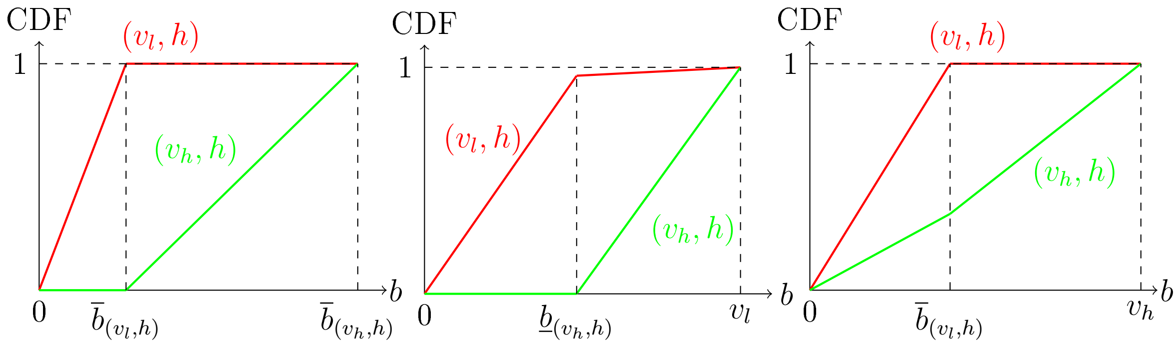

Figure 2.4: Equilibrium mixed strategies when $v_h = 2$, $v_l = 1$ and $p_h = 0.5$. (k_h, k_l, r): left: $(\frac{2}{3}, \frac{2}{3}, \frac{1}{3})$; middle $(\frac{1}{10}, \frac{2}{3}, \frac{1}{3})$; right: $(\frac{2}{3}, \frac{2}{3}, \frac{1}{10})$.

When $p_h v_h \geq v_l$ the maximum expected payoff is $p_h p_l v_h$ and when $p_h v_h \leq v_l$ the maximum expected payoff is $p_h(v_h - v_l)$. See Table 2.2 and Table 2.3 for two examples of the public signal which maximizes expected payoff of each player.

	(v_h, h)	(v_h, l)	(v_l, h)	(v_l, l)
(v_h, h)	p_h^2	0	$p_h p_l$	0
(v_h, l)	0	0	0	0
(v_l, h)	$p_h p_l$	0	0	0
(v_l, l)	0	0	0	p_l^2

Table 2.2: The optimal public signal under Conditions 2, 3, 4 and 5

	(v_h, h)	(v_h, l)	(v_l, h)	(v_l, l)
(v_h, h)	0	0	0	0
(v_h, l)	0	p_h^2	0	$p_h p_l$
(v_l, h)	0	0	p_l^2	0
(v_l, l)	0	$p_h p_l$	0	0

Table 2.3: The optimal public signal under Conditions ¬2, 3, 4 and ¬5

There are some common features of the two examples: 1) the high valuation player only receives one of the signal realizations (in Table 2.2 the high valuation type only receives h whereas in Table 2.3 the high valuation type only receives l), and when she does so she cannot differentiate the valuation of the opponent; 2) the low valuation type of player either only competes against a high valuation type or only competes against a low valuation type, i.e., she can differentiate her opponent's valuation. This drives the low valuation types to choose zero effort with probability 1 when competing against a high valuation type, and the high valuation player randomizes in an interval with the lower bound equals zero.

The examples given in Tables 2.2 and 2.3 are consistent with the studies concerning the partial disclosure policy $\{C, C, D\}$ in the all-pay auction (Lu et al., 2016) and in Tullock contest (Serena, 2015). In particular, such a disclosure policy requires concealment of valuation profile when it is $\{v_h, v_h\}$ or $\{v_l, v_l\}$, and fully disclose the profile to both players when it is $\{v_h, v_l\}$.[4] Proposition 12 below indicates that there exists a much broader set of public signals than $\{C, C, D\}$ or the signals given in Table 2.2 and 2.3, which maximize each player's expected payoff.

Proposition 12. *There exists an open set of public signal* (k_h, k_l, r) *which maximizes player* i*'s* $(i = 1, 2)$ *expected payoff.*

See an example to understand Proposition 12. Suppose Conditions 2, ¬3, 4 and ¬5 are satisfied, then the expected payoff of the player is always $p_h(v_h - v_l)$. In the equilibrium under this combination of conditions, type (v_h, h) of player i mixes over $[\underline{b}_{(v_h,h)}, v_l]$ and type (v_h, l) of player i mixes over $[\underline{b}_{(v_h,l)}, v_l]$, thus a player with high valuation earns $v_h - v_l$. But a player with low valuation mixes over $[0, v_l]$ and thus, earns zero. Therefore, whenever $p_h v_h \leq v_l$, we can use any public signal satisfying Conditions 2, ¬3, 4 and ¬5 to maximize players' expected payoff. Since the values (k_h, k_l, r) satisfying these conditions are not unique, the proposition follows.

2.5 Comparing private and public signals

In this section, we compare private and public signals in terms of increasing players' expected payoffs or the total expected effort. We start by showing some numerical examples which suggest that there is no general ranking between the two signals in terms of increasing players' expected payoffs. On the one hand, the following example suggests expected payoff can be larger with public than private signals.

Example 3. *Suppose* $p_h = \frac{1}{2}$, $v_h = 2$ *and* $v_l = 1$. *According to Proposition 11, the maximum expected payoff for player* i *with public signals is* $p_h p_l v_h = p_h(v_h - v_l) = \frac{1}{2}$. *The expected payoff with private signals in non-overlapping equilibrium is* $-\frac{5}{2}\left(q - \frac{13}{20}\right)^2 + \frac{49}{160}$ *which takes the maximum at* $\frac{49}{160} < \frac{1}{2}$ *when* $q = \frac{13}{20}$, *and with overlapping equilibrium is* $\frac{23}{36} - \frac{1}{9(3q-1)} - \frac{1}{3}q$ *which takes the maximum at* $\frac{11}{36} < \frac{1}{2}$ *when* $q = \frac{2}{3}$.

On the other hand, the following example suggests the expected payoff with private signals can be larger than $p_h(v_h - v_l)$, i.e., the maximum expected payoff of each player with public signals when $p_h v_h \leqslant v_l$.

Example 4. *Suppose* $p_h = \frac{2}{10}$, $v_h = 2$, $v_l = 1$ *and* $q = 0.7$. *Firstly, note that the cutoff value of* q *is* $\frac{2}{3}$, *thus with private signals the set of parameters entails an overlapping equilibrium, thus the expected payoff is* 0.2297. *Secondly, note also*

[4] Since the disclosure policy they consider is *anonymous*, $\{v_h, v_l\}$ represents any valuation profile when the two players' valuations are different.

that $p_h v_h - v_l = -0.6 < 0$, *thus the maximum expected payoff with public signals is* $p_h(v_h - v_l) = 0.2 < 0.2297$.

Similarly, the following example suggests the expected payoff with private signals can be larger than $p_h p_l v_h$, i.e., the maximum expected payoff of each player with public signals when $p_h v_h \geqslant v_l$.

Example 5. *Suppose* $p_h = \frac{1}{2}$, $v_h = 100$, $v_l = 1$ *and* $q = 0.7$. *Firstly, note that the cutoff value* q^* *is* $\frac{100}{101}$, *thus with private signals the set of parameters entail a non-overlapping equilibrium, thus the expected payoff is* 27.844. *Secondly, note also that* $p_h v_h - v_l = 49 > 0$, *thus the maximum expected payoff with public signals is* $p_h p_l v_h = 25 < 27.844$.

Let us turn to the comparison of signals in terms of increasing total expected effort. Recall that Proposition 9 indicates that the total expected effort with private signals when $q \in (\frac{1}{2}, 1)$ is always lower than when $q = \frac{1}{2}$. Thus, if there exists an example in which the total expected effort with public signals is higher than that in the IPV setting, i.e., $q = \frac{1}{2}$, then it must be true that public signals outperform private signals.

Proposition 13. *There exists an open set of public signals with which the total expected effort in the contest is higher than that with any private signal.*

Again, we provide an example showing that the total expected effort can be higher with public signal than in IPV setting.

Example 6. *Suppose* $(k_h, k_l, r) = (\frac{1}{10}, \frac{2}{3}, \frac{1}{3})$ *and* $p_h = \frac{1}{2}$, $v_h = 2$, $v_l = 1$. *The total expected effort in this case is 1.2553, which is larger than the expected effort in the IPV setting,* $p_h^2 v_h + (1 - p_h^2) v_l = 1.25$.

Recall that private signals always lower total expected effort, thus the example indicates that the public signal dominates private signal in maximizing total expected effort.

2.6 Conclusion

When players receive additional information regarding the opponent's valuation, they are always better off if the information is disclosed through conditional independent signals. They may be worse off if the information is disclosed through public signals. To maximize total expected effort, it is advised to disclose information through some public signals.

There are multiple directions to generalize the current paper. Firstly, the signals can be partially correlated and/or take more than two values. Since none of the two information structure considered in the current paper dominants the other, it can then be expected that a partially correlated signal might perform better. Secondly, there is an emerging literature on auctions with general information structure (Bergemann et al., 2015). It is also interesting to consider the lower or upper bound of players' expected payoff or total expected effort when there is no restrictions on information structure.

2.7 Appendix

Proof of Proposition 8

Proof. The non-overlapping part is proven by showing that no profitable deviation exists. Here, we show the process of checking type (v_h, h)'s profitable deviation. The checking of other types' deviation can be done in the same fashion and thus is omitted.

The expected payoff of type (v_h, h) when choosing an effort within her own equilibrium support, $(\bar{b}_{(v_h,l)}, \bar{b}_{(v_h,h)})$:

$$\frac{p_l(1-q)+p_hq(1-q)}{p_l(1-q)+p_hq}+\frac{p_hq^2}{p_l(1-q)+p_hq}G_{(v_h,h)}(b)v_h-b$$

Plug in (v_h, h)'s mixed strategy $G_{(v_h,h)}(b)$, the expected payoff is $v_h - \bar{b}_{(v_h,h)}$, which is exactly her equilibrium payoff. Now we check whether type (v_h, h) want to deviate to the supports of other players.

If type (v_h, h) deviate to (v_h, l) 's support, the expected payoff becomes

$$\left\{\frac{p_l(1-q)}{p_l(1-q)+p_hq}+\frac{p_hq(1-q)}{p_l(1-q)+p_hq}G_{(v_h,l)}(b)\right\}v_h-b$$

plug in the equilibrium mixed strategy of (v_h, l), $G_{(v_h,l)}$, and rearrange, the parameter of effort, b, becomes

$$\frac{p_lq+p_h(1-q)}{p_l(1-q)+p_hq}\frac{q}{1-q}-1 \tag{2.1}$$

which is also the first order derivative of the above expected payoff function w.r.t b. If $p_h \leqslant p_l$, that is, expression (2.1) is positive, then type (v_h, h) can increase her payoff by increasing b, until it reaches the upper bound of (v_h, l)'s support, $\bar{b}_{(v_h,l)}$, which is also the lower bound of (v_h, h)'s support. This suggests deviating to (v_h, l)'s support is not profitable. If, however, $p_h > p_l$ and thus (2.1) is negative, type (v_h, h) should choose the lower bound of (v_h, l)'s support, $\overline{b}_{(v_l,l)}$, instead of any effort higher. Thus we need to check whether the expected payoff of choosing $\overline{b}_{(v_l,l)}$ is higher than (v_h, h)'s equilibrium expected payoff when choosing an effort within her own support.

Let (v_h, h)'s equilibrium expected payoff be $\pi^*_{(v_h,h)}$ and her payoff from choosing $\overline{b}_{(v_l,l)}$ be $\pi_{(v_h,h)}(\overline{b}_{(v_l,l)})$, then the difference between the two:

$$\pi^*_{(v_h,h)} - \pi_{(v_h,h)}(\overline{b}_{(v_l,l)}) = p_h(1-q)\left[\frac{q}{p_h q + p_l(1-q)} - \frac{1-q}{p_h(1-q) + p_l q}\right] v_h > 0$$

Thus we have shown that type (v_h, h) do not want to deviate to (v_h, l)'s support. An important observation is that the expected payoff from a type deviate to another type's support is always a linear function of b, due to the all-pay rule. This fact ensures no strictly profit maximizing effort exists between the boundaries of any types' support. This means a simpler way of checking the equilibrium is to compare the equilibrium payoffs of each type with the payoffs from choosing each types' upper bounds of their equilibrium supports.

Now the only thing left to check is the profitability of choosing $\overline{b}_{(v_l,h)}$ and zero. When (v_h, h) chooses $\overline{b}_{(v_l,h)}$, the expected payoff is

$$\frac{(1-p)(1-q)q}{pq + (1-p)(1-q)} v_h - \frac{(1-p)(1-q)^2}{pq + (1-p)(1-q)} v_l$$

The gap between her equilibrium expected payoff and the above is:

$$\frac{(1-p)(1-q)}{pq + (1-p)(1-q)} [(1-q)v_h - qv_l]$$

which is positive when $(1-q)v_h > qv_l$. It is trivial to show that choosing zero cannot be more profitable. Thus it is not profitable for (v_h, h) to choose outside of her equilibrium support.

When $qv_l \geqslant (1-q)v_h$, the proof, again, consists of showing the indifference when choosing an effort inside the equilibrium support, and no profitable deviation exists. It is easy to check that all types are indifferent when choosing an effort in $[\overline{b}, \underline{b}]$, thus it is omitted. Here, we show that type (v_h, h) doesn't find it profitable to deviate on $[0, \underline{b}]$ and that type (v_h, l) doesn't want to deviate to $[\overline{b}, \overline{b}_{(v_h,h)}]$.

If type (v_h, h) deviate to $(0, \underline{b})$, then the expected payoff is

$$\frac{p_l(1-q)q}{p_h q + p_l(1-q)} G_{(v_l,h)}(b) v_h - b = \frac{qv_h - (1-q)v_l}{(1-q)\, v_l} b$$

is increasing with b since $qv_h > (1-q)v_l$.

If type (v_h, l) deviate to $(0, \underline{b})$

$$\frac{p_l q^2}{p_h(1-q)+p_h q} G_{(v_l,h)}(b) v_h - b = \frac{q^2((1-q)p_l + p_h q)v_h - (1-q)^2 \left(p_h(1-q)+p_l q\right) v_l}{v_l \left(q-1\right)^2 \left(p_h(1-q)+p_l q\right)} b$$

It's increasing because

$$\begin{aligned} & q^2(p_h q + (1-q)p_l)v_h - (1-q)^2 \left(p_h(1-q)+p_l q\right) v_l \\ > \; & q^2(p_h q + (1-q)p_l)v_l - (1-q)^2 \left(p_h(1-q)+qp_l\right) v_l \\ = \; & (2q-1)\left(p_h(1-q)+p_h q^2 + qp_l(1-q)\right) v_l > 0 \end{aligned}$$

If type (v_l, l) deviate to $(0, \underline{b})$, the expected payoff:

$$\frac{p_h(2q-1)}{(1-q)\left(p_h(1-q)+p_l q\right)} b$$

increasing with b.

If type (v_h, l) deviate to $(\overline{b}, \overline{b}_{(v_h,h)})$, the parameter of b in the expected payoff

$$\frac{-p_l(2q-1)}{q\left(p_h(1-q)+p_l q\right)} b$$

decreasing with b.

If type (v_l, l) deviate to $(\overline{b}, \overline{b}_{(v_h,h)})$, the parameter of b in the expected payoff:

$$\begin{aligned} & \frac{(1-q)^2\left(p_h q + p_l(1-q)\right) v_l - q^2 \left(p_h(1-q)+p_l q\right) v_h}{q^2 v_h \left(p_h(1-q)+p_l q\right)} \\ < \; & \frac{(1-q)^2\left(p_h q + p_l(1-q)\right) - q^2 \left(p_h(1-q)+p_l q\right)}{q^2 v_h \left(p_h(1-q)+p_l q\right)} v_l \\ = \; & -\frac{(2q-1)\left(q(p_h\left(1-q\right)+p_l q) + (1-q)p_l\right)}{q^2 v_h \left(p_h(1-q)+p_l q\right)} v_l < 0 \end{aligned}$$

If type (v_l, h) deviate to $(\overline{b}, \overline{b}_{(v_h,h)})$, the parameter of b in the expected payoff:

$$\frac{(1-q)v_l - qv_h)}{qv_h} = -\frac{qv_h - (1-q)v_l}{qv_h} < 0$$

Thus no type has profitable deviation. The uniqueness of the symmetric equilibrium is proven in the first chapter of this thesis. $\square$

Proof of Proposition 9

Proof. Note first that when $q = \frac{1}{2}$, the model is equivalent to the IPV setting and thus it is well known that the expected payoff of a player is $p_h p_l(v_h - v_l)$. Note also that when $q = 1$, the model is equivalent to the complete information setting, thus the expected payoff of a generic player is also $p_h p_l(v_h - v_l)$.

Under non-overlapping equilibrium, that is, when $(1-q)v_h \geq qv_l$, the expected payoff of each type are the following:
(v_h, h)'s payoff:

$$\pi^{ove}_{(v_h,h)} = v_h - \frac{p_h q^2 v_h + p_l(1-q)^2 v_l}{p_h q + p_l(1-q)} - \frac{p_l q^2 v_l + p_h(1-q)^2 v_h}{p_h(1-q) + p_l q}$$

(v_h, l)'s payoff:

$$\pi^{ove}_{(v_h,l)} = \frac{p_l q v_h - p_l q^2 v_l}{p_h(1-q) + p_l q} - \frac{p_l(1-q)^2 v_l}{p_h q + p_l(1-q)}$$

(v_l, l)'s payoff:

$$\pi^{ove}_{(v_l,l)} = \frac{p_l q(1-q) v_l}{p_h(1-q) + p_l q} - \frac{p_l(1-q)^2 v_l}{p_h q + p_l(1-q)}$$

(v_l, h)'s payoff: $\pi^{ove}_{(v_l,h)} = 0$

Thus, the ex ante expected payoff of a player is

$$\begin{aligned}
\pi &= p((p_h q + p_l(1-q))\pi^{ove}_{(v_h,h)} + (p_h(1-q) + p_l q)\pi^{ove}_{(v_h,l)}) + p_l(p_h(1-q) + p_l q)\pi^{ove}_{(v_l,l)} \\
&= p_h p_l(v_h - v_l) + p_h p_l(2q-1)[\frac{p_h(1-q)(p_h q + p_l(1-q))v_h}{(p_h(1-q) + p_l q)(p_h q + p_l(1-q))} \\
&- \frac{(-p_h - q - 3p_h q^2 - p_h^2 q + 2p_h^2 q^2 + 4p_h q + q^2)v_l}{(p_h(1-q) + p_l q)(p_h q + p_l(1-q))}] \\
&\geqslant p_h p_l(v_h - v_l) + p_h(2q-1)\frac{p_l(1-q)v_l}{p_h q + p_l(1-q)} \\
&> p_h p_l(v_h - v_l)
\end{aligned}$$

The greater equality used the condition $(1-q)v_h \geq qv_l$. Recall that both SPA and the non-overlapping equilibrium in the model ensures efficient allocation, thus the social surplus are the same across the two auctions. The total expected effort is thus higher in SPA than in APA.

If $(1-q)v_h \leqslant qv_l$, i.e., under overlapping equilibrium, the expected payoff are the following:

(v_h, h)'s payoff:

$$\begin{aligned}\pi^{non}_{(v_h,h)} &= v_h - \bar{b}_{(v_h,h)} \\ &= \frac{p_l(1-q)}{p_hq + p_l(1-q)}(v_h - v_l)\end{aligned}$$

(v_h, l)'s payoff:

$$\begin{aligned}\pi^{non}_{(v_h,l)} &= \frac{p_lq^2}{p_h(1-q)+p_lq}G_{(v_l,h)}(\underline{b})v_h - \underline{b} \\ &= \frac{(v_h - v_l)}{qv_h - (1-q)v_l}(\frac{p_lq^2}{p_h(1-q)+p_lq}v_h - \frac{p_l(1-q)^2)}{p_hq+p_l(1-q)}v_l)\end{aligned}$$

(v_l, l)'s payoff:

$$\begin{aligned}\pi^{non}_{(v_l,l)} &= \frac{p_lq(1-q)}{p_h(1-q)+p_lq}G_{(v_l,h)}(\underline{b})v_l - \underline{b} \\ &= \frac{(1-q)v_l}{qv_h - (1-q)v_l}(\frac{q}{p_h(1-q)+p_lq} - \frac{(1-q)}{p_hq+p_l(1-q)})p_l(v_h - v_l)\end{aligned}$$

(v_l, h)'s payoff: $\pi^{non}_{(v_l,h)} = 0$.
Thus player's surplus

$$\begin{aligned}\pi &= p_hp_l(v_h - v_l)\frac{q(p_hq + p_l(1-q))v_h - (1-q)(-p_h - 3q + 2p_hq + 2)v_l}{(qv_h - (1-q)v_l)(p_hq + p_l(1-q))} \\ &= p_hp_l(v_h - v_l)(1 + \frac{(1-q)(2q-1)v_l}{(qv_h - (1-q)v_l)(p_hq + p_l(1-q))}) \\ &> p_hp_l(v_h - v_l)\end{aligned}$$

This completes the proof of the first part of the equilibrium, now we turn to the second part. The social surplus, SS, is $p_l^2v_l + (1 - p_l^2)v_h$ in the non-overlapping equilibrium as the allocation is efficient, and is less than that in the overlapping equilibrium. Proposition 9 shows each player's expected payoff when $q \in (\frac{1}{2}, 1)$ is higher than when $q = \frac{1}{2}, 1$, i.e., $\pi(q) > p_hp_l(v_h - v_l)$. Thus the total expected effort equals the social surplus minus the joint expected payoff of players which is

$$\begin{aligned}R &= SS - 2\pi(q) \\ &\leqslant p_l^2v_l + (1 - p_l^2)v_h - 2\pi \\ &< p_l^2v_l + (1 - p_l^2)v_h - 2p_hp_l(v_h - v_l) \\ &= p_h^2v_h + (1 - p_h^2)v_l\end{aligned}$$

This completes the proof. $\square$

Proof of Proposition 10

Lemmas 10, 11, 12 proves that the strategy profile given under the Conditions 2 and 3 is indeed the unique equilibrium.

Lemma 10. *When Conditions 2 and 3 are satisfied, then types (v_h, h) and (v_l, h) randomize in non-overlapping supports. Furthermore, the support of type (v_h, h) is higher than the support of (v_l, h).*

Proof. For the first part of the lemma, suppose both the two types randomize in an interval, then in this interval it must be true that

$$\begin{aligned}\left(\frac{p_h k_h}{p_h k_h + p_l r} G_{(v_h,h)} + \frac{p_l r}{p_h k_h + p_l r} G_{(v_l,h)}\right) v_h - b &= K_{(v_h,h)} \\ \left(\frac{p_h r}{p_h r + p_l(1-k_l)} G_{(v_h,h)} + \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{(v_l,h)}\right) v_l - b &= K_{(v_l,h)}\end{aligned}$$

Thus,

$$\begin{aligned}G_{(v_h,h)}(b) &= \frac{r\left(p_h r + p_l\left(1-k_l\right)\right) v_h - \left(1-k_l\right)\left(p_h k_h + p_l r\right) v_l}{p_h\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l} b \\ &\quad - \left(1-k_l\right) \frac{p_l r + p_h k_h}{p_h\left(r^2 - k_h\left(1-k_l\right)\right) v_h}\left(v_h - v_l\right) \\ G_{(v_l,h)}(b) &= \begin{cases} \frac{-k_h(p_h r + p_l(1-k_l)) v_h + r(p_l r + p_h k_h) v_l}{p_l(r^2 - k_h(1-k_l)) v_h v_l} b + \frac{r(p_l r + p_h k_h)(v_h - v_l)}{p_l(r^2 - k_h(1-k_l)) v_h}, & \text{for } b \in [\underline{b}_{(v_h,h)}, v_l] \\ \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l) v_l} b, & \text{for } b \in [0, \underline{b}_{(v_h,h)}] \end{cases}\end{aligned}$$

In this case, the slop of $G_{(v_h,h)}(b)$ is

$$\begin{aligned}&\frac{r\left(p_h r + p_l\left(1-k_l\right)\right) v_h - \left(1-k_l\right)\left(p_h k_h + p_l r\right) v_l}{p_h\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l} \\ =\; &\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_h k_h + p_l r\right) \frac{\frac{p_l r}{(p_h k_h + p_l r)} v_h - \frac{p_l(1-k_l)}{(p_h r + p_l(1-k_l))} v_l}{p_l p_h\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l}\end{aligned}$$

In this case, the slop of $G_{(v_l,h)}(b)$ is

$$\begin{aligned}&\frac{-k_h\left(p_h r + p_l\left(1-k_l\right)\right) v_h + r\left(p_l r + p_h k_h\right) v_l}{p_l\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l} \\ =\; &\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_l r + p_h k_h\right) \frac{\frac{p_h r}{(p_h r + p_l(1-k_l))} v_l - \frac{p_h k_h}{(p_l r + p_h k_h)} v_h}{p_h p_l\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l}\end{aligned}$$

For the slop of $G_{(v_h,h)}(b)$ to be positive and the Condition 2 to be satisfied, it must be true that $(r^2 - k_h(1-k_l)) > 0$, for the slop of $G_{(v_l,h)}(b)$ to be positive and the Condition 3 to be satisfied, it must be true that $(r^2 - k_h(1-k_l)) < 0$. Thus, when Conditions 2 and 3 both satisfied, type (v_h, h) and (v_l, h)'s support cannot be overlapping.

Now we prove that the support of (v_h, h) must be higher than the support of (v_l, h). Suppose instead that the type (v_l, h) mixes over the interval $[\widehat{b}, \widetilde{b}]$, but the type (v_h, h) randomizes in the interval $[0, \widehat{b}]$, as the lowest possible effort for each player must be 0. However, this then implies type (v_h, h) must earns an expected payoff of 0, which cannot be true in any equilibrium as she can also deviate by choosing v_l to earn positive payoff. □

Lemma 11. *When Conditions 2 and 3 are satisfied, the mixed strategies given in the proposition form a symmetric equilibrium.*

Proof. When Conditions 2 and 3 are satisfied, we first show that a player with type (v_l, h) is indifferent in the equilibrium support. By plugging in the mixed strategy $G_{(v_l,h)}(b)$ in equilibrium, the expected payoff indeed equals zero:

$$G_{(v_l,h)}(b)\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)}v_l - b = 0$$

For type (v_h, h), we plug in $G_{(v_h,h)}(b)$ and the expected payoff is also constant and equals the expected payoff given in the proposition:

$$\left(\frac{p_l r}{p_h k_h + p_l r} + G_{(v_h,h)}(b)\frac{p_h k_h}{p_h k_h + p_l r}\right)v_h - b = \frac{p_l r}{p_h k_h + p_l r}v_h - \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)}v_l$$

Note that Condition 2 guarantees the expected payoff of type (v_h, h) is non-negative. Now we check for profitable deviations when each type deviates to choose outside of her equilibrium support. When type (v_h, h) deviate to the support of (v_l, h), the expected payoff becomes

$$\begin{aligned}
& G_{(v_l,h)}(b)\frac{p_l r}{p_h k_h + p_l r}v_h - b \\
= \; & \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l}b\frac{p_l r}{p_h k_h + p_l r}v_h - b \\
= \; & \left(\frac{r}{(1-k_l)}\frac{p_h r + p_l(1-k_l)}{p_h k_h + p_l r}\frac{v_h}{v_l} - 1\right)b
\end{aligned}$$

This expected payoff is increasing with b given the Condition 2 is satisfied. Thus, type (v_h, h) does not want to deviate to the support of (v_l, h). When type (v_l, h)

deviates to the support of (v_h, h), the expected payoff is:

$$
\begin{aligned}
& \left(\frac{p_l\,(1-k_l)}{p_h r + p_l\,(1-k_l)} + G_{(v_h,h)}(b)\frac{p_h r}{p_h r + p_l\,(1-k_l)}\right) v_l - b \\
= \;& \frac{p_l\,(1-k_l)}{p_h r + p_l\,(1-k_l)} v_l \\
+ \;& \left(\frac{p_h k_h + p_l r}{p_h k_h v_h} b - \frac{v_l}{v_h}\frac{p_l\,(1-k_l)}{p_h k_h}\frac{p_h k_h + p_l r}{p_h r + p_l\,(1-k_l)}\right)\frac{p_h r}{p_h r + p_l\,(1-k_l)} v_l - b \\
= \;& \left(\frac{v_l}{v_h}\frac{r}{k_h}\frac{p_h k_h + p_l r}{p_h r + p_l\,(1-k_l)} - 1\right) b \\
+ \;& v_l \left(\frac{p_l\,(1-k_l)}{p_h r + p_l\,(1-k_l)} - \frac{v_l}{v_h}\frac{(1-k_l)}{k_h} p_l r \frac{r p_l + k_h p_h}{\left(r p_h + (1-k_l)\,p_l\right)^2}\right)
\end{aligned}
$$

This expected payoff is decreasing in b given the condition if Condition 3 is satisfied. Thus, this is not a profitable deviation. □

Lemma 12. *When Conditions 2 and 3 are satisfied, then there is no asymmetric equilibrium.*

Proof. Denote by $\overline{b}_{1(v_h,h)}$ and $\overline{b}_{2(v_h,h)}$ the upper bound of equilibrium support of player 1 and 2, respectively. From a similar argument from the proof of Lemma above, for both players the type (v_h, h) must choose higher than (v_l, h). Thus, it must be true in any equilibrium that $\overline{b}_{1(v_h,h)} = \overline{b}_{2(v_h,h)}$. If there exists an asymmetric equilibrium, then it must be true that $\overline{b}_{1(v_l,h)} \neq \overline{b}_{2(v_l,h)}$. Suppose without loss that $\overline{b}_{1(v_l,h)} > \overline{b}_{2(v_l,h)}$. Since type (v_l, h) of player 1 is indifferent between any effort in $[0, \overline{b}_{2(v_l,h)}]$, thus her expected payoff

$$
\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{2(v_l,h)}(b) v_l - b = 0
$$

which gives

$$
G_{2(v_l,h)}(b) = \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l} b
$$

and $\overline{b}_{2(v_l,h)} = \frac{p_l(1-k_l)v_l}{p_h r + p_l(1-k_l)} = \overline{b}_{(v_l,h)}$ as given in the proposition. Since type (v_l, h) of player 2 is also indifferent between any effort in $[0, \overline{b}_{2(v_l,h)}]$, thus her expected payoff

$$
\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{1(v_l,h)}(b) v_l - b = 0
$$

which then gives

$$
G_{1(v_l,h)}(b) = \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l} b
$$

and it can be shown that $G_{1(v_l,h)}(\overline{b}_{(v_l,h)}) = 1$. Thus, it must be true that $\overline{b}_{1(v_l,h)} =$

$\bar{b}_{2(v_l,h)}$. □

Lemmas 13, 14, 15 proves that the strategy profile given under the Conditions 2 and ¬3 is indeed the unique equilibrium.

Lemma 13. *When Conditions 2 and ¬3 are satisfied, then in any equilibrium types (v_h, h) and (v_l, h) randomize in overlapping supports. Furthermore, the upper bound of supports $\bar{b}_{(v_h,h)} = \bar{b}_{(v_l,h)}$. Finally, the expected payoff of (v_l, h) must be 0 and the expected payoff of (v_h, h) must be $v_h - v_l$.*

Proof. Suppose types (v_h, h) and (v_l, h) randomize in non-overlapping equilibrium. Then the support of (v_h, h) must be higher than (v_l, h), thus type (v_l, h)'s expected payoff must be zero. However, by choosing $\bar{b}_{(v_h,h)}$ the type (v_l, h)'s expected payoff must be $v_l - \bar{b}_{(v_h,h)} = \frac{p_h r}{p_h r + p_l(1-k_l)} v_l - \frac{p_h k_h}{p_h k_h + p_l r} v_h > 0$, as the Condition ¬3 is strictly satisfied. Thus, in any symmetric equilibrium it cannot be true that the supports are overlapping.

In any symmetric equilibrium, it cannot be true that both types earn positive payoff, as one of the types must have the lower bound of support equals 0. Suppose both types have lower bound equals 0, then both earn a payoff of 0. In that case, the indifference conditions in the overlapping interval of their supports are

$$\begin{aligned}
\left(\frac{p_h k_h}{p_h k_h + p_l r} G_{(v_h,h)}(b) + \frac{p_l r}{p_h k_h + p_l r} G_{(v_l,h)}(b)\right) v_h - b &= 0 \\
\left(\frac{p_h r}{p_h r + p_l(1-k_l)} G_{(v_h,h)}(b) + \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{(v_l,h)}(b)\right) v_l - b &= 0
\end{aligned}$$

and thus

$$\begin{aligned}
G_{(v_h,h)}(b) &= \frac{r\left(p_h r + p_l\left(1-k_l\right)\right) v_h - \left(1-k_l\right)\left(p_h k_h + p_l r\right) v_l}{p_h\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l} b \\
G_{(v_l,h)}(b) &= \frac{-k_h\left(p_h r + p_l\left(1-k_l\right)\right) v_h + r\left(p_h k_h + p_l r\right) v_l}{p_l\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l} b
\end{aligned}$$

By letting $G_{(v_h,h)}(b) = G_{(v_l,h)}(b) = 1$, we have

$$\bar{b}_{(v_h,h)} = \frac{p_h p_l\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l}{\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_h k_h + p_l r\right)} \frac{1}{\frac{p_l r}{\left(p_h k_h + p_l r\right)} v_h - \frac{p_l(1-k_l)}{\left(p_h r + p_l(1-k_l)\right)} v_l}$$

and

$$\bar{b}_{(v_l,h)} = \frac{p_h p_l\left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l}{\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_l r + p_h k_h\right)} \frac{1}{-\frac{p_h k_h}{\left(p_l r + p_h k_h\right)} v_h + \frac{p_h r}{\left(p_h r + p_l(1-k_l)\right)} v_l}$$

thus

$$\begin{aligned}
& \bar{b}_{(v_h,h)} - \bar{b}_{(v_l,h)} \\
= \ & \frac{p_h p_l \left(r^2 - k_h\left(1-k_l\right)\right) v_h v_l}{\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_h k_h + p_l r\right)} \left(\frac{1}{\frac{p_l r v_h}{(p_h k_h + p_l r)} - \frac{p_l(1-k_l) v_l}{(p_h r + p_l(1-k_l))}} - \frac{1}{\frac{p_h r v_l}{(p_h r + p_l(1-k_l))} - \frac{p_h k_h v_h}{(p_l r + p_h k_h)}} \right) \\
< \ & \frac{p_h p_l \left(r^2 - k_h\left(1-k_l\right)\right) v_h}{\left(p_h r + p_l\left(1-k_l\right)\right)\left(p_h k_h + p_l r\right)} \left(\frac{1}{\frac{p_l r}{(p_h k_h + p_l r)} - \frac{p_l(1-k_l)}{(p_h r + p_l(1-k_l))}} - \frac{1}{\frac{p_l r}{(p_l r + p_h k_h)} - \frac{p_l(1-k_l)}{(p_h r + p_l(1-k_l))}} \right) \\
= \ & 0
\end{aligned}$$

This means $\bar{b}_{(v_h,h)} < \bar{b}_{(v_l,h)}$. Then it must be true that $\bar{b}_{(v_l,h)} \leqslant v_l$ as any effort above v_l is strictly dominated to type (v_l, h). But then type (v_h, h) has an incentive to choose $\bar{b}_{(v_l,h)}$ to earn $v_h - \bar{b}_{(v_l,h)} \geqslant v_h - v_l > 0$. Thus, it cannot be true that both types have lower bound of support equals 0.

So the only case left is type (v_h, h) earns positive expected payoff whereas type (v_l, h) earns 0. Thus, the lower bound of (v_h, h)'s support must be positive. It cannot be true that $\bar{b}_{(v_h,h)} < \bar{b}_{(v_l,h)} < v_l$, as then the expected payoff of (v_l, h) would be positive. It cannot be true that $\bar{b}_{(v_h,h)} < \bar{b}_{(v_l,h)} = v_l$, as type (v_h, h) will increase the effort until $\bar{b}_{(v_l,h)}$ in the interval $[\bar{b}_{(v_h,h)}, \bar{b}_{(v_l,h)}]$. In particular, we have for type (v_l, h) in the interval $[\bar{b}_{(v_h,h)}, \bar{b}_{(v_l,h)}]$ that:

$$\left(\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{(v_l,h)}(b) + \frac{p_h r}{p_h r + p_l(1-k_l)} \right) v_l - b = 0$$

thus

$$G_{(v_l,h)}(b) = \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)} \frac{b}{v_l} - \frac{p_h r}{p_l(1-k_l)}$$

Now if type (v_h, h) increase the effort from $\bar{b}_{(v_h,h)}$ to $b \in (\bar{b}_{(v_h,h)}, \bar{b}_{(v_l,h)})$ the expected payoff increases by

$$\begin{aligned}
& \frac{p_l r}{p_h k_h + p_l r} \left[G_{(v_l,h)}(b) - G_{(v_l,h)}(\bar{b}_{(v_h,h)}) \right] v_h - \left[b - \bar{b}_{(v_h,h)} \right] \\
= \ & \left[\frac{p_l r}{p_h k_h + p_l r} \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)} \frac{v_h}{v_l} - 1 \right] \left[b - \bar{b}_{(v_h,h)} \right] > 0
\end{aligned}$$

According to Condition 2 this is positive. Thus we must have $\bar{b}_{(v_h,h)} \geqslant \bar{b}_{(v_l,h)}$. Suppose $\bar{b}_{(v_h,h)} > \bar{b}_{(v_l,h)}$, then in the interval $[\bar{b}_{(v_l,h)}, \bar{b}_{(v_h,h)}]$, type (v_h, h) must be indifferent

$$\left(\frac{p_h k_h}{p_h k_h + p_l r} G_{(v_h,h)}(b) + \frac{p_l r}{p_h k_h + p_l r} \right) v_h - b = \widehat{\pi}_{(v_h,h)}$$

nd thus the mixed strategy is

$$G_{(v_h,h)}(b) = \frac{p_h k_h + p_l r}{p_h k_h}\frac{b + \widehat{\pi}_{(v_h,h)}}{v_h} - \frac{p_l r}{p_h k_h}$$

If type (v_l, l) increases her effort from $\overline{b}_{(v_l,h)}$ to $b \in (\overline{b}_{(v_l,h)}, \overline{b}_{(v_h,h)})$ then the expected payoff increases by

$$\begin{aligned} & \frac{p_h r}{p_h r + p_l(1-k_l)}\left[G_{(v_h,h)}(b) - G_{(v_h,h)}(\overline{b}_{(v_l,h)})\right] v_l - \left[b - \overline{b}_{(v_l,h)}\right] \\ = & \left[\frac{p_h r}{p_h r + p_l(1-k_l)}\frac{p_h k_h + p_l r}{p_h k_h}\frac{v_l}{v_h} - 1\right]\left[b - \overline{b}_{(v_l,h)}\right] > 0 \end{aligned}$$

According to Condition ¬3 this is positive. Thus, it must be true that $\overline{b}_{(v_h,h)} = \overline{b}_{(v_l,h)}$. Consider the interval $[\underline{b}_{(v_h,h)}, \overline{b}_{(v_h,h)}]$, where type (v_h, h) and (v_l, h)'s indifference conditions must be

$$\begin{aligned} \left(\frac{p_h k_h}{p_h k_h + p_l r}G_{(v_h,h)}(b) + \frac{p_l r}{p_h k_h + p_l r}G_{(v_l,h)}(b)\right) v_h - b &= \widehat{\pi}_{(v_h,h)} \\ \left(\frac{p_h r}{p_h r + p_l(1-k_l)}G_{(v_h,h)}(b) + \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)}G_{(v_l,h)}(b)\right) v_l - b &= 0 \end{aligned}$$

and thus

$$\begin{aligned} G_{(v_h,h)}(b) &= \frac{r\left(p_h r + p_l\left(1-k_l\right)\right)v_h - \left(1-k_l\right)\left(p_h k_h + p_l r\right)v_l}{p_h\left(r^2 - k_h\left(1-k_l\right)\right)v_h v_l}b \\ & \quad - \left(1-k_l\right)\frac{p_l r + p_h k_h}{p_h\left(r^2 - k_h\left(1-k_l\right)\right)v_h}\widehat{\pi}_{(v_h,h)} \\ G_{(v_l,h)}(b) &= \frac{-k_h\left(p_h r + p_l\left(1-k_l\right)\right)v_h + r\left(p_l r + p_h k_h\right)v_l}{p_l\left(r^2 - k_h\left(1-k_l\right)\right)v_h v_l}b + r\frac{p_l r + p_h k_h}{p_l\left(r^2 - k_h\left(1-k_l\right)\right)v_h}\widehat{\pi}_{(v_h,h)} \end{aligned}$$

Given that $\overline{b}_{(v_h,h)} = \overline{b}_{(v_l,h)} = \beta$, we have $G_{(v_h,h)}(\beta) = G_{(v_l,h)}(\beta) = 1$, this implies $\widehat{\pi}_{(v_h,h)} = v_h - v_l$ and $\beta = v_l$. □

Lemma 14. *When Conditions 2 and ¬3 are satisfied, the mixed strategies given in the proposition form a symmetric equilibrium.*

Proof. When Conditions 2 and ¬3 are satisfied, we show that type (v_l, h) and (v_h, h) are indifferent in their equilibrium support. In the interval $[\underline{b}_{(v_h,h)}, v_l]$, for type (v_l, h), the expected payoff is calculated by plugging in the expression of $G_{(v_h,h)}(b)$ and $G_{(v_l,h)}(b)$, the expected payoff is zero:

$$\left(G_{(v_h,h)}(b)\frac{p_h r}{p_h r + p_l\left(1-k_l\right)} + G_{(v_l,h)}(b)\frac{p_l\left(1-k_l\right)}{p_h r + p_l\left(1-k_l\right)}\right)v_l - b = 0$$

Similarly, type (v_h, h)'s expected payoff can be shown to be a constant $v_h - v_l$:

$$\left(G_{(v_h,h)}(b)\frac{p_h k_h}{p_h k_h + p_l r} + G_{(v_l,h)}(b)\frac{p_l r}{p_h k_h + p_l r}\right) v_h - b = v_h - v_l$$

Since only type (v_l, h) is choosing an effort in the interval $[0, \underline{b}_{(v_h,h)}]$, by plug the $G_{(v_l,h)}(b)$ in, her expected payoff in this interval is

$$G_{(v_l,h)}(b)\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} v_l - b = 0$$

Now we prove that both types do not want to deviate to any effort outside of their equilibrium support. When type (v_h, h) deviates to $[0, \underline{b}_{(v_h,h)}]$, then her expected payoff would be

$$\begin{aligned} & G_{(v_l,h)}(b)\frac{p_l r}{p_h k_h + p_l r} v_h - b \\ = \quad & \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)v_l} b \frac{p_l r}{p_h k_h + p_l r} v_h - b \\ = \quad & b\frac{r(p_h r + p_l(1-k_l))v_h - (1-k_l)(p_h k_h + p_l r)v_l}{v_l(1-k_l)(rp_l + k_h p_h)} \end{aligned}$$

This is increasing in b. □

Lemma 15. *When Conditions 2 and ¬3 are satisfied, then there is no asymmetric equilibrium.*

Proof. If player 1 has $\bar{b}_{1(v_h,h)} = \bar{b}_{1(v_l,h)} = v_l$, then player 2 must have $\bar{b}_{2(v_h,h)} = \bar{b}_{2(v_l,h)} = v_l$. To see why, suppose $\bar{b}_{2(v_h,h)} < \bar{b}_{2(v_l,h)} = v_l$, then by the same argument in the previous lemma, type (v_h, h) of player 1 is strictly better off by reallocating probability mass from the interval $(\bar{b}_{2(v_h,h)}, \bar{b}_{2(v_l,h)})$ to v_l. Similarly, if $\bar{b}_{2(v_h,h)} = v_l > \bar{b}_{2(v_l,h)}$, then the previous lemma indicates that type (v_l, h) is strictly better off by reallocating probability mass from the interval $(\bar{b}_{2(v_l,h)}, \bar{b}_{2(v_h,h)})$ to v_l. Thus, it must be true that $\bar{b}_{1(v_h,h)} = \bar{b}_{1(v_l,h)} = \bar{b}_{2(v_h,h)} = \bar{b}_{2(v_l,h)} = v_l$, which means the expected payoff are the same as in the unique symmetric equilibrium.

By $G_{1(v_h.h)}(\underline{b}_{1(v_h,h)}) = 0$, it can be checked that $\underline{b}_{1(v_h,h)} = \underline{b}_{(v_h,h)}$ as in the symmetric equilibrium. Similarly, it can be found that $\underline{b}_{2(v_h,h)} = \underline{b}_{(v_h,h)}$ by $G_{2(v_h.h)}(\underline{b}_{2(v_h,h)}) = 0$. Therefore, $\underline{b}_{1(v_h,h)} = \underline{b}_{2(v_h,h)} = \underline{b}_{(v_h,h)}$ in any equilibrium. □

Lemmas 16, 17, 18 proves that the strategy profile given under the Conditions ¬2 and 3 is indeed the unique equilibrium.

Lemma 16. *When Conditions ¬2 and 3 are satisfied, then in any equilibrium types (v_h, h) and (v_l, h) randomize in overlapping supports. Furthermore, the expected payoff for both types are zero.*

Proof. Suppose (v_h, h) and (v_l, h) randomize in non-overlapping supports in a symmetric equilibrium, then again it must be true that (v_h, h)'s support is higher than (v_l, h). This implies type (v_h, h)'s expected payoff must be $v_h - \overline{b}_{(v_h,h)} = \frac{p_l r}{p_h k_h + p_l r} v_h - \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} v_l \leqslant 0$, as Condition ¬2 is satisfied. Therefore, the two types must have overlapping supports.

Again it cannot be true that both types earn positive payoff. Now, suppose type (v_h, h) earns positive payoff and type (v_l, h) earns zero. Thus $\underline{b}_{(v_h,h)} > \underline{b}_{(v_l,h)} = 0$. In the interval $[0, \underline{b}_{(v_h,h)}]$, type (v_l, h)'s indifference condition is

$$\frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} G_{(v_l,h)}(b) v_l - b = 0$$

thus

$$G_{(v_l,h)}(b) = \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)} \frac{b}{v_l}$$

Now if type (v_h, h) decreases her effort from $\underline{b}_{(v_h,h)}$ to $b \in \left(0, \underline{b}_{(v_h,h)}\right)$, her expected payoff increses by

$$\begin{aligned} & \frac{p_l r}{p_h k_h + p_l r} \left[G_{(v_l,h)}(b) - G_{(v_l,h)}(\underline{b}_{(v_h,h)})\right] v_h - \left[b - \underline{b}_{(v_h,h)}\right] \\ = & \left[\frac{p_l r}{p_h k_h + p_l r} \frac{p_h r + p_l(1-k_l)}{p_l(1-k_l)} \frac{v_h}{v_l} - 1\right] \left[b - \underline{b}_{(v_h,h)}\right] > 0 \end{aligned}$$

According to Condition ¬2, the above is positive. Thus, it is profitable for type (v_h, h) to decrease the effort until 0. This then implies the expected payoff of type (v_h, h) must also be 0. This then implies $\overline{b}_{(v_h,h)} = v_h$, as any $v_h > \overline{b}_{(v_h,h)}$ suggests type (v_h, h) earns positive expected payoff. □

Lemma 17. *When Conditions ¬2 and 3 are satisfied, the mixed strategies given in the proposition form a symmetric equilibrium.*

Proof. When Conditions ¬2 and 3 are satisfied, both type (v_h, h) and (v_l, h) are indifferent in the equilibrium support $[0, \overline{b}_{(v_l,h)}]$, and they both get zero:

$$\begin{aligned} \left(G_{(v_h,h)}(b)\frac{p_h r}{p_h r + p_l\left(1-k_l\right)} + G_{(v_l,h)}(b)\frac{p_l\left(1-k_l\right)}{p_h r + p_l\left(1-k_l\right)}\right) v_l - b &= 0 \\ \left(G_{(v_h,h)}(b)\frac{p_h k_h}{p_h k_h + p_l r} + G_{(v_l,h)}(b)\frac{p_l r}{p_h k_h + p_l r}\right) v_h - b &= 0 \end{aligned}$$

Type (v_h, h) also get zero when choosing an effort in $[\bar{b}_{(v_l,h)}, v_h]$ as

$$\left(\frac{p_l r}{p_h k_h + p_l r} + G_{(v_h,h)}(b)\frac{p_h k_h}{p_h k_h + p_l r}\right) v_h - b = 0$$

Type (v_l, h) does not want to deviate to $[\bar{b}_{(v_l,h)}, v_h]$ as

$$\begin{aligned} & \left(G_{(v_h,h)}(b)\frac{p_h r}{p_h r + p_l (1-k_l)} + \frac{p_l (1-k_l)}{p_h r + p_l (1-k_l)}\right) v_l - b \\ = & \left(\frac{r (p_l r + p_h k_h)}{k_h (p_l (1-k_l) + p_h r)}\frac{v_l}{v_h} - 1\right) b + v_l \left(\frac{p_l (1-k_l)}{p_l (1-k_l) + p_h r} + \frac{r}{k_h}\frac{p_l r}{p_l (1-k_l) + p_h r}\right) \end{aligned}$$

which is decreasing in b. Thus none of the two types want to deviate. □

Lemma 18. *When Conditions ¬2 and 3 are satisfied, then there is no asymmetric equilibrium.*

Proof. Only thing to check is that $\bar{b}_{1(v_l,h)} = \bar{b}_{2(v_l,h)}$. Suppose $\bar{b}_{1(v_l,h)} > \bar{b}_{2(v_l,h)}$. Type (v_h, h) and (v_l, h) of player 2's expected payoff is

$$\begin{aligned} \left(\frac{p_h k_h}{p_h k_h + p_l r}G_{1(v_h,h)}(b) + \frac{p_l r}{p_h k_h + p_l r}G_{1(v_l,h)}(b)\right) v_h - b &= 0 \\ \left(\frac{p_h r}{p_h r + p_l(1-k_l)}G_{1(v_h,h)}(b) + \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)}G_{1(v_l,h)}(b)\right) v_l - b &= 0 \end{aligned}$$

when means

$$\begin{aligned} G_{1(v_h,h)}(b) &= \frac{r (p_h r + p_l (1-k_l)) v_h - (1-k_l) (p_h k_h + p_l r) v_l}{p_h (r^2 - k_h (1-k_l)) v_h v_l} b \\ G_{1(v_l,h)}(b) &= \frac{-k_h (p_h r + p_l (1-k_l)) v_h + r (p_l r + p_h k_h) v_l}{p_l (r^2 - k_h (1-k_l)) v_h v_l} b \end{aligned}$$

Then according to $G_{1(v_l,h)}(\bar{b}_{1(v_l,h)}) = 1$, we have $\bar{b}_{1(v_l,h)} = \bar{b}_{(v_l,h)}$. Similarly we can find $G_{2(v_l,h)}(b)$, and according to $G_{2(v_l,h)}(\bar{b}_{2(v_l,h)}) = 1$, we have $\bar{b}_{2(v_l,h)} = \bar{b}_{(v_l,h)}$. □

Finally, the part of the proof for types (v_h, l) and (v_l, l) can be obtained by exchanging k_h with $1-k_h$, k_l with $1-k_l$ and r with $1-r$ in the above proofs.

Proof of Proposition 11

Proof. Note first that there are 3×3 possible cases as it is not possible to have ¬2 and ¬3 satisfied simultaneously or ¬4 and ¬5 satisfied simultaneously.

Case 1: When Conditions 2, ¬3, 4 and ¬5 are satisfied, then by the equilibrium strategy given in the Proposition 3, the expected payoffs of each type are: Type

(v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = v_h - v_l$; type (v_l, h)'s expected payoff: $V_{(v_l,h)}(k_h, k_l, r) = 0$; and type (v_h, l)'s expected payoff: $V_{(v_h,l)}(k_h, k_l, r) = v_h - v_l$; type (v_l, l)'s expected payoff: $V_{(v_l,l)}(k_h, k_l, r) = 0$.

Thus, player i's expected payoff is $p_h (v_h - v_l)$ for all values of (k_h, k_l, r) satisfying Conditions 2, ¬3, 4 and ¬5. Suppose $p_h v_h \geqslant v_l$ then according to Condition ¬3

$$\begin{aligned} & \frac{p_h k_h}{p_h k_h + p_l r} v_h - \frac{p_h r}{p_h r + p_l (1 - k_l)} v_l \\ \geqslant & \left(\frac{k_h}{p_h k_h + p_l r} - \frac{p_h r}{p_h r + p_l (1 - k_l)} \right) v_l \end{aligned}$$

And according to Condition ¬5

$$\begin{aligned} & \frac{p_h (1 - k_h)}{p_h (1 - k_h) + p_l (1 - r)} v_h - \frac{p_h (1 - r)}{p_h (1 - r) + p_l k_l} v_l \\ \geqslant & \left(\frac{(1 - k_h)}{p_h (1 - k_h) + p_l (1 - r)} - \frac{p_h (1 - r)}{p_h (1 - r) + p_l k_l} \right) v_l \end{aligned}$$

After rearrange, we have $(1 - k_l) k_h \leqslant p_h r (r - k_h)$ and $k_l (1 - k_h) \leqslant p_h (1 - r) (k_h - r)$ which then implies $k_h = r$. But this is inconsistent with ¬3 and ¬5 since we also have $p_h v_h - \frac{p_h r}{p_h r + p_l (1-k_l)} v_l \geqslant \frac{p_l (1-k_l)}{p_h r + p_l (1-k_l)} v_l \geqslant 0$ which contradicts ¬3; and since we also have $p_h v_h - \frac{p_h (1-r)}{p_h (1-r) + p_l k_l} v_l \geqslant \frac{p_l k_l}{p_h (1-r) + p_l k_l} v_l \geqslant 0$ which contradicts ¬5. Therefore, it must be true that $p_h v_h \leq v_l$, and thus, $p_h (v_h - v_l) \leqslant p_h p_l v_h$.

Case 2: When Conditions 2, 3, 4 and 5 are satisfied, then the expected payoffs of each type are: Type (v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = \frac{p_l r}{p_h k_h + p_l r} v_h - \frac{p_l (1-k_l)}{p_h r + p_l (1-k_l)} v_l$; type (v_l, h)'s expected: $V_{(v_l,h)}(k_h, k_l, r) = 0$; type (v_h, l)'s expected payoff: $V_{(v_h,l)}(k_h, k_l, r) = \frac{p_l (1-r)}{p_h (1-k_h) + p_l (1-r)} v_h - \frac{p_l k_l}{p_h (1-r) + p_l k_l} v_l$; type (v_l, l)'s expected: $V_{(v_l,l)}(k_h, k_l, r) = 0$

Thus, player i's expected payoff is

$$\begin{aligned} V(k_h, k_l, r) &= p_h (p_h k_h + p_l r) V_{(v_h,h)} + p_h (p_h (1 - k_h) + p_l (1 - r)) V_{(v_h,l)} \\ &= p_h p_l v_h - \left((1 - k_l) \frac{p_h k_h + p_l r}{p_h r + p_l (1 - k_l)} + k_l \frac{p_h (1 - k_h) + p_l (1 - r)}{p_h (1 - r) + p_l k_l} \right) p_h p_l v_l \end{aligned}$$

Note that the expected payoff is maximized if the second term is zero, which is when $k_h = k_l = r = 1$, and the Conditions 2 and 3 now becomes $p_l v_h \geqslant 0$ and $p_h v_h - v_l \geqslant 0$ The Conditions 4 and 5 are irrelevant in this case as the probability of receiving a signal l is zero for players with v_h. Thus, the maximum is $p_h p_l v_h$. Note that Condition 3 implies $p_h (v_h - v_l) \geqslant p_h p_l v_h = V(k_h, k_l, r)$.

Case 3: When Conditions ¬2, 3, 4 and 5 are satisfied, then the expected payoffs of each type are: Type (v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = 0$; type (v_l, h)'s expected payoff: $V_{(v_l,h)}(k_h,k_l,r)=0$; type (v_h,l)'s expected payoff: $V_{(v_h,l)}(k_h,k_l,r)=\frac{p_l(1-r)}{p_h(1-k_h)+p_l(1-r)}v_h-\frac{p_lk_l}{p_h(1-r)+p_lk_l}v_l$; type (v_l,l)'s expected payoff: $V_{(v_l,l)}(k_h,k_l,r)=0$

Thus, player i's expected payoff is

$$\begin{aligned} V(k_h,k_l,r) &= p_h\left(p_h\left(1-k_h\right)+p_l\left(1-r\right)\right)\left(\frac{p_l(1-r)}{p_h(1-k_h)+p_l(1-r)}v_h-\frac{p_lk_l}{p_h(1-r)+p_lk_l}v_l\right) \\ &= p_hp_l(1-r)v_h-\frac{p_h\left(1-k_h\right)+p_l\left(1-r\right)}{p_h(1-r)+p_lk_l}p_hp_lk_lv_l \end{aligned}$$

and when $k_l=r=0$, $V(k_h,k_l,r)$ reaches its maximum $p_hp_lv_h$. Check the conditions when $k_l=r=0$ in the order of Conditions ¬2, 3, 4 and 5:

$$\begin{aligned} -v_l &\leqslant 0 \text{ and } v_h\geqslant 0 \\ \frac{p_l}{p_h(1-k_h)+p_l}v_h-v_l &\geqslant 0 \text{ and } \frac{p_h\left(1-k_h\right)}{p_h\left(1-k_h\right)+p_l}v_h-v_l\geqslant 0 \end{aligned}$$

Thus Conditions ¬2 and 3 are satisfied, but Conditions 4 and 5 impose some restrictions on k_h: $k_h\in[1-\frac{p_l}{p_h}\frac{v_h-v_l}{v_l},1-\frac{p_l}{p_h}\frac{v_l}{v_h-v_l}]$. Since k_h is restricted to be between zero and one, we need $1-\frac{p_l}{p_h}\frac{v_l}{v_h-v_l}\geqslant 0$ which implies $p_hv_h\geqslant v_l$, thus, $p_h(v_h-v_l)\geqslant p_hp_lv_h=V(k_h,k_l,r)$.

Case 4: When Conditions ¬2, 3, ¬4 and 5 are satisfied, then the expected payoffs of each type are: Type (v_h,h)'s expected payoff: $V_{(v_h,h)}(k_h,k_l,r)=0$; type (v_l,h)'s expected payoff: $V_{(v_l,h)}(k_h,k_l,r)=0$; type (v_h,l)'s expected payoff: $V_{(v_h,l)}(k_h,k_l,r)=0$; type (v_l,l)'s expected payoff: $V_{(v_l,l)}(k_h,k_l,r)=0$. Thus, player i's ex ante expected payoff is 0.

Case 5: When Conditions 2, 3, ¬4 and 5 are satisfied, then the expected payoffs of each type are: type (v_h,h)'s expected payoff: $V_{(v_h,h)}(k_h,k_l,r)=\frac{p_lr}{p_hk_h+p_lr}v_h-\frac{p_l(1-k_l)}{p_hr+p_l(1-k_l)}v_l$; type (v_l,h)'s expected payoff: $V_{(v_l,h)}(k_h,k_l,r)=0$; type (v_h,l)'s expected payoff: $V_{(v_h,l)}(k_h,k_l,r)=0$; type (v_l,l)'s expected: $V_{(v_l,l)}(k_h,k_l,r)=0$

Thus, player i's ex ante expected payoff is:

$$V(k_h,k_l,r)=p_hp_lrv_h-\frac{p_hk_h+p_lr}{p_hr+p_l\left(1-k_l\right)}p_hp_l\left(1-k_l\right)v_l$$

Let $k_l=r=1$, then it is maximized at $p_hp_lv_h$. See below that Conditions ¬4 and 5 are satisfied, whereas Conditions 2 and 3 impose some restrictions on k_h: $k_h\geqslant\frac{p_l}{p_h}\frac{v_l}{v_h-v_l}$:

$$\begin{aligned} \frac{p_l}{p_hk_h+p_l}v_h &\geqslant 0 \text{ and } \frac{p_hk_h}{p_hk_h+p_l}v_h-v_l\geqslant 0 \\ -v_l &\leqslant 0 \text{ and } v_h-v_l\geqslant 0 \end{aligned}$$

Since k_h has to be between zero and one, we need $\frac{p_l}{p_h}\frac{v_l}{v_h-v_l} \leqslant 1$ which implies $v_l \leqslant p_h v_h$, and thus $p_h(v_h - v_l) \geqslant p_h p_l v_h$.

Case 6: When Conditions $\neg 2$, 3, 4 and $\neg 5$ are satisfied, then the expected payoffs of each type are: Type (v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = 0$; type (v_l, h)'s expected payoff: $V_{(v_l,h)}(k_h, k_l, r) = 0$; type (v_h, l)'s expected payoff: $V_{(v_h,l)}(k_h, k_l, r) = v_h - v_l$; type (v_l, l)'s expected payoff: $V_{(v_l,l)}(k_h, k_l, r) = 0$.

Thus, player i's ex ante expected payoff is:

$$V(k_h, k_l, r) = p_h \left(p_h \left(1-k_h\right) + p_l \left(1-r\right)\right) \left(v_h - v_l\right)$$

which is maximized at $p_h \left(v_h - v_l\right)$ when $k_h = r = 0$. In this case $\neg 5$ implies $p_h v_h \leqslant v_l$, then $p_h \left(v_h - v_l\right) \leqslant p_h p_l v_h$

Case 7: When Conditions 2, 3, 4 and $\neg 5$ are satisfied, then the expected payoffs of each type are: Type (v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = \frac{p_l r}{p_h k_h + p_l r} v_h - \frac{p_l(1-k_l)}{p_h r + p_l(1-k_l)} v_l$; type (v_l, h)'s expected payoff: $V_{(v_l,h)}(k_h, k_l, r) = 0$; type (v_h, l)'s expected payoff: $V_{(v_h,l)}(k_h, k_l, r) = v_h - v_l$; type (v_l, l)'s expected payoff: $V_{(v_l,l)}(k_h, k_l, r) = 0$.

Thus, player i's ex ante expected payoff is:

$$\begin{aligned} V(k_h, k_l, r) &= p_h \left(p_h k_h + p_l r\right) V_{(v_h,h)} + p_h \left(p_h \left(1-k_h\right) + p_l \left(1-r\right)\right) V_{(v_h,l)} \\ &= p_h \left(p_l + p_h \left(1-k_h\right)\right) v_h - p_h \left(1 - \frac{p_h r}{p_h r + p_l \left(1-k_l\right)} \left(p_h k_h + p_l r\right)\right) v_l \end{aligned}$$

Since the Condition 3 is equivalent of $p_h k_h v_h - \frac{p_h r}{p_h r + p_l(1-k_l)} \left(p_h k_h + p_l r\right) v_l \geqslant 0$, we thus have $V(k_h, k_l, r) \leqslant p_h \left(v_h - v_l\right)$. In other words, when the Condition 3 is binding, the expected payoff reaches its maximum of $p_h \left(v_h - v_l\right)$. Thus, any public signal (k_h, k_l, r) satisfies the Condition 3 when it is binding and satisfies Condition 2, 4, $\neg 5$ maximizes the expected payoff.

Suppose $p_h v_h \geqslant v_l$, then $\neg 5$ implies $0 \leqslant \frac{p_h(1-k_h)}{p_h(1-k_h)+p_l(1-r)} v_h - \frac{p_h(1-r)}{p_h(1-r)+p_l k_l} v_l \geqslant \left(\frac{(1-k_h)}{p_h(1-k_h)+p_l(1-r)} - \frac{p_h(1-r)}{p_h(1-r)+p_l k_l}\right) v_l$ and thus, $\frac{(1-k_h)}{p_h(1-k_h)+p_l(1-r)} \leqslant \frac{p_h(1-r)}{p_h(1-r)+p_l k_l}$ thus $k_l \left(1-k_h\right) \leqslant p_h \left(1-r\right) \left(k_h - r\right)$. Similarly, 3 implies $\frac{k_h}{p_h k_h + p_l r} \leqslant \frac{p_h r}{p_h r + p_l(1-k_l)}$ thus $\left(1-k_l\right) k_h \leqslant p_h r \left(r - k_h\right)$. Thus, it must be true that $k_h = r$. Thus, $p_h v_h - \frac{p_h(1-r)}{p_h(1-r)+p_l k_l} v_l \geqslant \frac{p_l k_l}{p_h(1-r)+p_l k_l} v_l$ and $\neg 5$ implies $k_l = 0$. But this is inconsistent with 3 binding, as $p_h v_h - \frac{p_h r}{p_h r + p_l} v_l \geqslant \frac{p_l}{p_h r + p_l} v_l > 0$. Therefore, we must have $p_h v_h \leqslant v_l$ and thus, $p_h \left(v_h - v_l\right) \leqslant p_h p_l v_h = V(k_h, k_l, r)$.

Case 8: When Conditions 2, $\neg 3$, 4 and 5 are satisfied, then the expected payoffs of each type are: Type (v_h, h)'s expected payoff: $V_{(v_h,h)}(k_h, k_l, r) = v_h - v_l$; type (v_l, h)'s expected payoff: $V_{(v_l,h)}(k_h, k_l, r) = 0$; type (v_h, l)'s expected payoff:

$V_{(v_h,l)}(k_h,k_l,r) = \frac{p_l(1-r)}{p_h(1-k_h)+p_l(1-r)}v_h - \frac{p_lk_l}{p_h(1-r)+p_lk_l}v_l$; type (v_l,l)'s expected payoff: $V_{(v_l,l)}(k_h,k_l,r) = 0$.

Thus, player i's ex ante expected payoff is:

$$\begin{aligned} V(k_h,k_l,r) &= p_h\left(p_hk_h+p_lr\right)V_{(v_h,h)}(k_h,k_l,r) + p_h\left(p_h\left(1-k_h\right)+p_l\left(1-r\right)\right)V_{(v_h,l)}(k_h,k_l,r) \\ &= p_h\left(p_hk_h+p_l\right)v_h - p_h\left(p_hk_h+p_lr+\frac{p_h\left(1-k_h\right)+p_l\left(1-r\right)}{p_h(1-r)+p_lk_l}p_lk_l\right)v_l \end{aligned}$$

It is decreasing with k_l, thus let $k_l=0$, and we have

$$V(k_h,k_l,r) = p_h\left(p_hk_h\left(v_h-v_l\right)-p_lrv_l\right)+p_hp_lv_h$$

which is increasing with k_h and decreasing with r. Now given that $k_l=0$, the conditions become the following (in the order of Conditions 2, $\neg$3, 4 and 5)

$$\begin{aligned} \frac{p_lr}{p_hk_h+p_lr}v_h - \frac{p_l}{p_hr+p_l}v_l &\geqslant 0 \text{ and } \frac{p_hk_h}{p_hk_h+p_lr}v_h - \frac{p_hr}{p_hr+p_l}v_l \leqslant 0 \\ \frac{p_l\left(1-r\right)}{p_h(1-k_h)+p_l(1-r)}v_h &\geqslant 0 \text{ and } \frac{p_h\left(1-k_h\right)}{p_h\left(1-k_h\right)+p_l\left(1-r\right)}v_h - v_l \geqslant 0 \end{aligned}$$

The Condition 5 then implies: $p_h\left(v_h-v_l\right)-p_lv_l \geqslant p_hk_h\left(v_h-v_l\right)-p_lrv_l$ and thus, $V(k_h,k_l,r) \leqslant p_h\left(v_h-v_l\right)$. To reach the maximum, Condition 5 must be binding, i.e., $\frac{p_h(1-k_h)}{p_h(1-k_h)+p_l(1-r)}v_h = v_l$. Suppose $p_hv_h \geqslant v_l$, then we have $v_l = \frac{p_h(1-k_h)}{p_h(1-k_h)+p_l(1-r)}v_h \geqslant \frac{(1-k_h)}{p_h(1-k_h)+p_l(1-r)}v_l$ thus $k_h \geqslant r$. This then violates Condition $\neg$3 as $\frac{p_hk_h}{p_hk_h+p_lr}v_h - \frac{p_hr}{p_hr+p_l}v_l \geqslant \left(\frac{k_h}{p_hk_h+p_lr} - \frac{p_hr}{p_hr+p_l}\right)v_l \geqslant \left(1-\frac{p_hr}{p_hr+p_l}\right)v_l > 0$. Thus, it must be true that $p_h\left(v_h-v_l\right) \leqslant p_hp_lv_h$.

Case 9: When Conditions 2, $\neg$3, $\neg$4 and 5 are satisfied, then the expected payoffs of each type are: Type (v_h,h)'s expected payoff: $V_{(v_h,h)}(k_h,k_l,r) = v_h-v_l$; type (v_l,h)'s expected payoff: $V_{(v_l,h)}(k_h,k_l,r) = 0$; type (v_h,l)'s expected payoff: $V_{(v_h,l)}(k_h,k_l,r) = 0$; type (v_l,l)'s expected payoff: $V_{(v_l,l)}(k_h,k_l,r) = 0$.

Thus, player i's ex ante expected payoff is:

$$V(k_h,k_l,r) = p_h\left(p_hk_h+p_lr\right)\left(v_h-v_l\right)$$

which is maximized at $p_h\left(v_h-v_l\right)$ when $k_h=r=1$. In this case $\neg$3 implies $p_h(v_h-v_l) \leqslant p_hp_lv_h$.

Therefore, the maximum expected payoff is $\min\{p_hp_lv_h, p_h(v_h-v_l)\}$. □

Proof of Proposition 13

Proof. First, we show that with all possible set of parameters of the all-pay auction, i.e., any $p_h \in (0,1)$ and $v_h > v_l$, there exists a public signal which raises at least the same total expected effort as the IPV setting, which is the highest total expected effort the all-pay auction can raise with private signal. Suppose $k_h = r = 1 - k_l = s$, then it can be easily checked that Conditions 2, 3, 4 and 5 are satisfied. Thus, all types randomize in non-overlapping intervals. Furthermore, player i's expected effort equals:

$$p_l(p_h r + p_l(1-k_l))\frac{1}{2}\frac{p_l\,(1-k_l)}{p_h r + p_l\,(1-k_l)}v_l \tag{2.2}$$

$$+\quad p_h(p_h k_h + p_l r)\left(\frac{p_l\,(1-k_l)}{p_h r + p_l\,(1-k_l)}v_l + \frac{1}{2}\frac{p_h k_h}{p_h k_h + p_l r}v_h\right) \tag{2.3}$$

$$+\quad p_l(p_h(1-r) + p_l k_l)\frac{1}{2}\frac{p_l k_l}{p_h(1-r) + p_l k_l}v_l \tag{2.4}$$

$$+\quad p_h(p_h(1-k_h) + p_l(1-r))\left(\frac{p_l k_l}{p_h(1-r) + p_l k_l}v_l + \frac{1}{2}\frac{p_h(1-k_h)}{p_h(1-k_h) + p_l(1-r)}v_h\right) \tag{2.5}$$

where (2.2) is $p_l(p_h r + p_l(1-k_l))$ times the expected effort of type (v_l, h), (2.3) is $p_h(p_h k_h + p_l r)$ times the expected effort of type (v_h, h), (2.4) is $p_l(p_h(1-r) + p_l k_l)$ times the expected effort of type (v_l, l), and (2.5) is $p_h(p_h(1-k_h) + p_l(1-r))$ times the expected effort of type (v_h, l). Let

$$k_h = r = 1 - k_l = s \in [0,1], \tag{2.6}$$

then the above becomes $p_h^2 v_h + (1 - p_h^2)v_l$ which is equivalent of the total expected effort in the IPV setting, i.e., the maximum total expected effort with the private signal. Thus, for any value of p_h and $v_h > v_l$, we can always let the public signal satisfies (2.6). This means using public signal can at least raise a total expected effort no less than using private signal.

Next, we provide an example in which public signal raises higher total expected effort.

Example 7. *Suppose $(k_h, k_l, r) = (\frac{1}{10}, \frac{2}{3}, \frac{1}{3})$ and $p_h = \frac{1}{2}$. The total expected effort in this case is 1.2553, which is larger than the total expected effort in the IPV setting, i.e., 1.25.*

Therefore, using public signal, the all-pay auction can always raise a total expected effort equals to that in the IPV setting, and with some set of parameters (e.g. the open ball centered at the point $(k_h, k_l, r) = (\frac{1}{10}, \frac{2}{3}, \frac{1}{3})$), public signal induces higher expected effort. $\square$

Chapter 3

Heterogeneous Risk/Loss Aversion in Complete Information All-Pay Auctions

Zhuoqiong Chen, David Ong and Ella Segev

Abstract We extend previous theoretical work on n-players complete information all-pay auction to incorporate heterogeneous risk and loss averse utility functions. We provide sufficient and necessary conditions for the existence of equilibria with a given set of active players with any strictly increasing utility functions and characterize the players' equilibrium mixed strategies. Assuming that players can be ordered by their risk aversion (player a is more risk averse than player b if whenever player b prefers a certain payment over a given lottery so will player a), we find that, in equilibrium, the more risk averse players either bid higher (in terms of first order stochastic dominance of their mixed strategy cumulative distribution) than the less risk averse players and win with higher ex-ante probability – or they drop out. Furthermore, while each player's expected bid decreases with the other players' risk aversion, her expected bid increases with her own risk aversion. Thus, increasing a player's risk aversion creates two opposing effects on total expected bid. A sufficient condition for the total expected bid to decrease with a player's risk aversion is that this player is relatively more risk averse compared to the rest of the players. Our findings have important implications for the literature on gender differences in competitiveness and for gender diversity in firms that use personnel contests for promotions.

3.1 Introduction

Sunk cost contests, where effort is unrecoverable, are pervasive. Especially important are those where the winners need only perform slightly better to take all (Frank and Cook, 2010; Rosen, 1981). These are in effect all-pay auctions in their incentive structure. Indeed, all-pay auctions theory has been used to study many types of contest and tournaments, e.g., rent seeking contest and lobbying (Baye et al., 1993; Ellingsen, 1991; Hillman and Riley, 1989), election campaigns (Che and Gale, 1998), R&D races (Dasgupta et al., 1982), college admission (Andreoni and Brownback, 2014; Hickman, 2014), and job promotion (Rosen, 1986). In these contests, the risk of lost effort, opportunities, or resources to individuals can be significant. Furthermore, even contests between organizations, like firms, can involve significant loss to individuals to the extent that decisions are made by individual CEOs and managers who care about the consequences of those decisions on their own welfare, through such mechanisms, for example, as options in compensation packages (Bertrand, 2009), and of course, in promotions and in dismissals based upon relative performance. However, despite the importance of risk in such contests which can be modeled as all-pay auctions, the modeling of all-pay auction incentives has generally been restricted to risk neutral players or to specific utility functional forms (Parreiras and Rubinchik, 2010; Klose and Schweinzer, 2014) or to local approximations (Fibich et al., 2006)[1]. Moreover, in the case of gender, the difference in risk aversion is observable. Observability is important because there is accumulating evidence of a gender difference in risk aversion, where women are found to be more risk averse than men (Charness and Gneezy, 2012; Croson and Gneezy, 2009). These observable gender differences in risk attitudes and their interactions with all-pay auction incentives in the business world could contribute to an explanation of the paucity of women among top executives (Bertrand, 2009), particularly in entrepreneurial settings (Coates et al., 2009).

In order to fill this gap in the theory of all-pay auctions, we extend Baye, Kovenock, and De Vries (1996)'s n-player, complete information all-pay auction model to incorporate heterogeneous risk averse players. We provide sufficient and necessary conditions for any equilibrium to exist and more importantly, closed-form solutions to the equilibrium strategies for any strictly increasing utility functions, focusing on weakly concave utility functions as well as loss averse utility function. After characterizing equilibrium strategies, we derive novel comparative statistics

[1]Siegel (2009) gives a general framework of finding equilibria with heterogeneous players, but does not explicitly characterize the equilibria or provide comparative statics of heterogeneous risk averse players as we do here.

for equilibria in which active players randomize continuously from 0 to the common value of the prize, given that players can be ordered by their risk aversion (player a is more risk averse than player b if whenever player b prefers a certain payment over a given lottery so will player a).

We find that, in equilibrium, the more risk averse players either bid higher than the less risk averse players (in terms of first order stochastic dominance of their mixed strategy cumulative distribution) and win with higher ex-ante probability – or they drop out. When players are homogeneous in their risk aversion, the total expected bid decreases with their risk aversion. We find, surprisingly, in the heterogeneous risk aversion case, that while each player's expected bid decreases with the other players' risk aversion, her expected bid increases with her own risk aversion. Thus, increasing a player's risk aversion creates two opposing effects on total expected bid. A sufficient condition for the total expected bid to decrease with a player's risk aversion is that this player is relatively more risk averse compared to the rest of the players.

With only two risk aversion types of players, we show that the total expected bid decreases monotonically with the share of the more risk averse players, when the difference between the two types is not too large. Our findings have important implications for the literatures on gender differences in competitiveness and for gender diversity in firms that use personnel contests for promotions. We discuss these implications after the main results.

3.2 The model

There are m players who have a common valuation, $v_1 = \cdots = v_m = v$ for the prize[2]. Players compete in an all-pay auction for one prize by submitting a bid (exerting an effort): x_i. The vector of bids is denoted $(x_1, x_2, \ldots, x_m)$. The payoff function in an all-pay auction is given by:

$$\pi_i(x_1, x_2, \ldots, x_m) = \begin{cases} -x_i & \text{if} \quad \exists j, \quad x_j > x_i \\ v_i - x_i & \text{if} \quad x_j < x_i \quad \text{for all} \quad j \end{cases}.$$

Moreover, there exists some tie breaking rule to determine the winner in case there is more than one bidder with the highest bid. Any tie breaking rule is applicable in our model. We assume that players are risk/loss averse with strictly

[2]The model we will present can be trivially extended to the case in which one player has higher valuation, while all other players have the same lower valuation. However, when there are finite many possible valuations, the interaction between valuation and risk attitude significantly complicates the model. We leave this for future work.

increasing utility functions which we denote by $U_1(x), U_2(x), \ldots, U_m(x)$. These utilities are common knowledge and potentially different from each other. We discuss two cases separately: 1) risk averse and 2) loss averse. For case 1), we assume only continuity and concavity of the utility functions. For case 2), we assume that the utility functions take the following form:

$$U_i(x) = \begin{cases} g_i(x) & \text{if} \quad x > 0 \\ 0 & \text{if} \quad x = 0 \\ l_i(x) & \text{if} \quad x < 0 \end{cases}, \tag{3.1}$$

where the utility from gains, $g_i(x)$ is a concave function while the utility from losses, $l_i(x)$ is a convex function, and both are continuous in their domains.

In this paper, we focus on mixed strategy equilibria. In any such equilibria, any active player (a player who bids a positive amount with positive probability) i is indifferent between all the bids in her equilibrium support. Formally, that means,

$$\rho U_i(v-x) + (1-\rho)U_i(-x) = U_i(E\pi_i\,(x, b_{-i})) \tag{3.2}$$

where x is in the support of the player's equilibrium strategy, ρ denotes the probability that bidder i wins when she bids x, and $E\pi_i\,(x, b_{-i})$ is the certainty equivalent of bidding x given the other players bid b_{-i}. We can rewrite equation (3.2) as:

$$\Pr(i \text{ wins}|x, b_{-i}) = \rho = \frac{U_i(E\pi_i) - U_i(-x)}{U_i(v-x) - U_i(-x)}$$

We define $K_{U_i}(x)$ to facilitate the analysis of the mixed strategy equilibria.

$$K_{U_i}(x) = \frac{U_i(0) - U_i(-x)}{U_i(v-x) - U_i(-x)},$$

In our analysis of equilibria, the equilibrium probability of winning that makes player i indifferent between bidding $x > 0$ and bidding zero (since bidding zero yields a zero payoff for sure) will be equal to $K_{U_i}(x)$. We sometimes abuse notation and write $K_{U_i}(x)$ as $K_i(x)$.

In what follows, we will exploit the following important property of $K_{U_i}(x)$. Its magnitude only depends on player i's risk attitude and not on any other players' risk attitude or bids. This is already evident in the definition. In fact, we will show that $K_{U_i}(x)$ is monotonic in player i's risk aversion in the lemma below. All proofs are in the Appendix. We first define increasing risk aversion and increasing loss aversion.

Definition 8. *A concave utility function $U(\cdot)$ represents a more risk averse player*

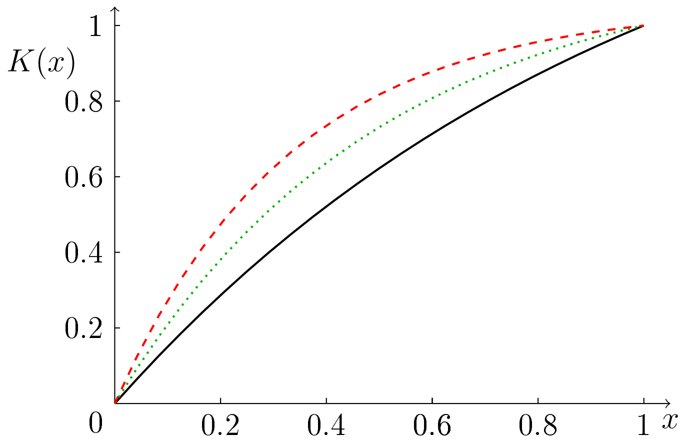

Figure 3.1: $K(x)$ increases with risk aversion.

than the concave utility function $\tilde{U}(\cdot)$, if for any lottery l over a set of prizes Z, the lottery's certainty equivalent is smaller under U than under $\tilde{U}$. In that case, we say that the risk aversion of the player increases from $\tilde{U}$ to U.

Definition 9. *For a player with a utility function $U(\cdot)$ of the form (3.1), a convex loss averse function $l(\cdot)$ represents a more loss averse player than the convex function $\tilde{l}(\cdot)$, if $l(x) < \tilde{l}(x)$ for all $x < 0$.*

Lemma 19. *If $U_i(x)$ is concave, then for any $x \in (0, v)$, $K_i(x)$ increases with player i's risk aversion, i.e., if U_i represents a more risk averse player than $\tilde{U}_i$ then $K_{U_i}(x) > K_{\tilde{U}_i}(x)$ for any $x \in (0, v)$. When players are both risk and loss averse as described above by the utility function of the form (3.1), then for any x, $K_i(x)$ increases with player i's loss aversion.*

Note that Lemma 19 above also suggests that the function $K(x)$ of different players will never cross if the players can be ordered by their risk or loss aversions. The following is an example of the function $K(x)$ when the player has CARA utility function:

Example 8. *If a player has CARA utility function: $U_i(c) = 1 - e^{-\beta_i c}$ and $v = 1$, then $K_i(x) = \frac{1-e^{-\beta_i x}}{1-e^{-\beta_i}}$. In this case, player i is more risk averse than player j if $\beta_i > \beta_j$. In figure 3.1 we plot $K_i(x)$ for $\beta = 1$ (black solid), 2 (green dotted), and 3 (red dashed).*

In fact, the sufficient and necessary conditions for the existence of equilibrium that we find below, and the closed-form expressions for the mixed strategies we provide in the next section, rely only on the assumption that the utility functions are strictly increasing, i.e., utility functions do not have to be rankable by their certainty equivalent. The rest of the results apply to any utility function that is rankable by their certainty equivalent, irrespective of whether the utility function is risk averse, loss averse, or even risk seeking. All the results with risk loving players can be derived analogously, as long as the more risk loving utilities have higher certainty equivalent than the less risk loving utilities for every lottery. We focus only on risk and loss averse utilities due to their ubiquity in the literature.

3.3 Equilibrium

In this section, we characterize the sufficient and necessary conditions for the existence of any possible equilibrium, and then, we characterize the mixed strategies in all of these equilibria. We also highlight some interesting features of the equilibria. In discussing these features, for simplicity, we focus only on the equilibria in which all active players randomize on the entire interval $[0, v]$.

3.3.1 Existence and closed-form solution

Our first proposition, Proposition 14 provides the necessary and sufficient conditions for the existence of an equilibrium in which a given subset of players is active. We start by defining an active player.

Definition 10. *A player is active when she bids zero with a probability strictly less than* 1. *A player is inactive when she bids zero with probability* 1.

Let $B \subseteq \{1, \ldots, m\}$ be a set of players. For convenience and without loss of generality, we assign $i = 1, 2, ..., |B|$ as the index for the active players.

Proposition 14. *An equilibrium in which a set $B \subseteq \{1, \ldots, m\}$ of players is active, and*

1. *players $i = 1, 2, \ldots, h$, where $2 \leqslant h \leqslant |B|$, randomize continuously over $[0, v]$, and*

2. *players $i = h+1, h+2, \ldots, |B|$ randomize continuously over $[b_i, v]$ and have an atom at zero, with $0 = b_h < b_{h+1} \leqslant b_{h+2} \leqslant \ldots \leqslant b_{|B|} \leqslant b_{|B|+1} = v$, and*

3. *players $j = |B| + 1, \ldots, m$ are inactive,*

   exists if and only if the following conditions hold for all $0 \leqslant t \leqslant |B| - h$:

   (I) Incentive Constraints: $\prod_{l \leqslant h+t} K_l(x) \leqslant K_j^{h+t-1}(x)$, for all $x \in [b_{h+t}, b_{h+t+1}]$ and $j > h + t$;

(II) Feasibility Constraints: $\prod_{l\leqslant h+t} K_l(x) \leqslant K_i^{h+t-1}(x)$, *for all* $x \in [b_{h+t}, b_{h+t+1}]$ *and* $i \leqslant h+t$.

We now prove the proposition and characterize the equilibrium bidding strategies. We first restrict our attention to the case where all the active players in the equilibrium randomize continuously on $[0,v]$, i.e., $b_{h+1} = b_{h+2} = \ldots = b_{|B|} = 0$. In this case $h = |B|$, and we can simplify the constraints for the existence of this equilibrium to:

(i) Incentive Constraints: $\prod_{l\in B} K_l(x) \leqslant K_j^{|B|-1}(x)$, for all $x \in [0,v]$ and for all $j > |B|$;

(ii) Feasibility Constraints: $\prod_{l\in B} K_l(x) \leqslant K_i^{|B|-1}(x)$, for all $x \in [0,v]$ and for all $i \leqslant |B|$.

Let player i's mixed strategy cumulative distribution function denoted by $G_i(x)$. For player $i \leqslant |B|$ and $x \in [0,v]$, we have in equilibrium:

$$\prod_{l\neq i, l\in B} G_l(x)U_i(v-x) + (1 - \prod_{l\neq i, l\in B} G_l(x))U_i(-x) = U_i(0) \tag{3.3}$$

where $G_l(x)$ is the probability player $l \neq i$ bid lower than x, so $\prod_{l\neq i, l\in B} G_l(x)$ is player i's probability of winning when bidding x. Note that by bidding zero the player's payoff is zero with certainty, and therefore, she must be indifferent between bidding any $x \in (0,v]$ or getting zero. We thus have for all $i \leqslant |B|$ and $x \in [0,v]$

$$\prod_{l\neq i, l\in B} G_l(x) = K_i(x)$$

We solve this system of $|B|$ equations and get the equilibrium strategy of player $i \in B$

$$G_i(x) = \left(\prod_{l\in B, l\neq i} K_l(x)\right)^{\frac{1}{|B|-1}} K_i(x)^{-\frac{|B|-2}{|B|-1}} \tag{3.4}$$

We are now able to calculate the probability of winning of an inactive player who deviates to some positive bid x:

$$\prod_{i\in B} G_i(x) = \prod_{i\in B}\left(\left(\prod_{l\in B, l\neq i} K_l(x)\right)^{\frac{1}{|B|-1}} K_i(x)^{-\frac{|B|-2}{|B|-1}}\right) = \prod_{l\in B} K_l(x)^{\frac{1}{|B|-1}} \tag{3.5}$$

Note that if the probability of winning given in equation (3.5) is less than $K_j(x)$, where $j \notin B$, as indicated by the incentive constraint for player j, then player j will earn an expected payoff less than zero should he bid any positive amount x, which makes him worse off than staying inactive. Therefore, the incentive constraints ensure that inactive players do not want to deviate to positive bids.

The feasibility constraints guarantee that the mixed strategies played by the active players are well defined (between zero and one). This can be obtained from equation (3.4). By restricting $G_i(x) \leqslant 1$, we have

$$\prod_{l\in B} K_l\left(x\right) \leqslant K_i^{|B|-1}\left(x\right)$$

which is the feasibility constraint for player $i \in B$. Together we have shown that if both the incentive and the feasibility constraints on the K functions hold, then the strategy profile defined by (3.4) constitutes an equilibrium.

The above derivation can be extended to the case with $0 < b_{h+1} \leqslant b_{h+2} \leqslant \ldots \leqslant b_{|B|} \leqslant v$. Specifically, by definition, there are exactly $h+t$ players (players $1, ..., h+t$) who place bids in the interval $[b_{h+t}, b_{h+t+1}]$, and thus, a system of $h+t$ equations for any bid $x \in [b_{h+t}, b_{h+t+1}]$. The equilibrium strategy, incentive and feasibility constraints can then be derived through the same procedure shown above. The incentive constraints now guarantee not only that an inactive player will not want to deviate and bid a positive bid but also that an active player will not want to deviate and bid outside her support. Moreover, none of the players would want to deviate to any bid above v, as they will earn negative payoff for sure.

We now characterize the equilibrium bidding strategies. Assume an equilibrium strategy profile as described in Proposition 14, then the equilibrium strategies for the active players must make each active player indifferent between any point on her support and a payoff of zero. (Recall that active players have zero in their support which yields a zero payoff.) Therefore, from (3.3) we must have for $\forall x \in [b_{|B|}, v]$

$$G_i\left(x\right) = \left(\prod_{l\in B} K_l(x)\right)^{\frac{1}{|B|-1}} K_i(x)^{-1}, \tag{3.6}$$

where $i = 1, \ldots, |B|$. For $t = h+1, h+2, \ldots, |B|-1$; we have for $\forall x \in [b_t, b_{t+1}]$

$$G_i\left(x\right) = \left(\frac{\prod_{l\leqslant t} K_l(x)}{\prod_{l>t} G_l(b_l)}\right)^{\frac{1}{t-1}} K_i(x)^{-1} \tag{3.7}$$

where $i = 1, 2, \ldots, t$; and

$$G_k(x) = G_k(b_k) \tag{3.8}$$

where $k = t+1, \ldots, |B|$. Finally, for $\forall x \in [0, b_{h+1}]$

$$G_i(x) = \left(\frac{\prod_{l \leqslant h} K_l(x)}{\prod_{l>h} G_l(b_l)} \right)^{\frac{1}{h-1}} K_i(x)^{-1} \tag{3.9}$$

where $i = 1, 2, \ldots, h$; and

$$G_k(x) = G_k(b_k) \tag{3.10}$$

where $k = h+1, \ldots, m$.

This completes the proof of the proposition.

Example 9. *Assume there are three players with CARA utility functions $U_i(x) = 1 - e^{-\beta_i x}$ and a valuation $v = 1$ for $i = 1, 2, 3$. Assume that $\beta_3 = 1$, $\beta_2 = 2$, $\beta_1 = 10$. Then, there exists no equilibrium in which all three players are active and all randomize continuously on $[0, 1]$ since the feasibility constraint is violated on $[0.13035, 1]$. Specifically, we have $G_3(x) = (K_1(x)K_2(x))^{\frac{1}{2}} K_3(x)^{-\frac{1}{2}} = \left(\frac{1-e^{-2x}}{1-e^{-2}} \frac{1-e^{-10x}}{1-e^{-10}}\right)^{\frac{1}{2}} \left(\frac{1-e^{-x}}{1-e^{-1}}\right)^{-\frac{1}{2}}$, which is larger than 1 for $x \geqslant 0.13035$. However there exists an equilibrium in which only players 2 and 3 are active and they randomize continuously on the interval $[0, 1]$ according to the following strategies: $G_2(x) = K_3(x) = \frac{1-e^{-x}}{1-e^{-1}}$ and $G_3(x) = K_2(x) = \frac{1-e^{-2x}}{1-e^{-2}}$ since then all the conditions hold.*

The incentive constraints determine who participates and who does not. The inactive players require better odds of winning (higher $K(x)$) for each positive bid (x) than what the active players in equilibrium can provide. The feasibility constraints impose a restriction on active players: they cannot be too different in terms of risk attitudes. This condition restricts the level of heterogeneity of active players. According to Proposition 1, our model entails multiple equilibria in which different numbers of players are active in equilibrium. However, this fact does not restrict the power of our theory in making predictions either for empirical or experimental data, since in reality we can generally observe the number of active players, especially if players play over multiple rounds.

3.3.2 Some features of equilibria

In real life competitions, it is not uncommon for participants to differ in observable characteristics like gender, ethnicity, culture...etc. It is then important to examine whether risk attitudes associated with these characteristics help to explain the difference in competitive behaviour. For example, women are under-represented in the elites of many competitive industries (Bertrand, 2009), yet women are also more likely to achieve academic success (Angrist et al., 2009; DiPrete and Buchmann, 2013; Fortin et al., 2015)[3]. Importantly, our results below are consistent with this empirical evidence.

Corollary 8. *If there exists an equilibrium where the set of all active players, B, randomize continuously on the interval $[0, v]$, and these players can be ranked by their risk aversions: $K_{B_1}(x) \geqslant K_{B_2}(x) \geqslant \ldots \geqslant K_{B_{|B|}}(x)$ for all $x \in [0, v]$, then the cumulative distribution function of player s's strategy first order stochastically dominates that of player t for every $t > s$ and the players' expected bids have the same ranking as their levels of risk aversion.*

Corollary 8 suggests that more risk averse players bid higher in expectation than less risk averse players among all active players. Given that the more risk averse a player is, the higher she bids conditional on her being active in equilibrium, one may expect that her probability of winning is also higher. Corollary 9 indicates that this conjecture is generally true but is not always the case.

Corollary 9. *Assume an equilibrium where the set of all active players, B, randomize continuously on the interval $[0, v]$. For any two active players, $s, t \in B$, if player s is more risk averse than player t, i.e., $K_s(x) \geqslant K_t(x)$ for all $x \in [0, v]$, then player s's probability of winning is higher or equal to that of player t if $K_t(x)$ dominates $K_s(x)$ in terms of the reverse hazard rate, i.e., $\frac{K_t'(x)}{K_t(x)} \geqslant \frac{K_s'(x)}{K_s(x)}$ for all $x \in [0, v]$.*

Note that $K_s(x)$, $K_t(x)$ are also the joint cumulative distributions of opponents' bids that players s and t are competing against (e.g., $\Pi_{l \in B, l \neq s} G_l(x) = K_s(x)$), respectively. Corollary 9 suggests that the more risk averse player s is more likely to win the contest compared to player t, if in player t's view (as measured by $K_t(x)$) the contest is sufficiently more competitive (i.e., dominates in terms of reverse hazard rate) than in player s's view (as measured by $K_s(x)$).

Interestingly, the more risk averse players not only bid higher and win with higher probability, they are also more likely to dropout in the following sense.

Corollary 10. *Assume an equilibrium where the set of all active players, B, randomize continuously on the interval $[0, v]$. If for some $i \in B$ and $j \notin B$, we have $K_i(x) \geqslant K_j(x)$ for all $x \in [0, v]$, then the existence of the equilibrium with the set B of active players implies the existence of another equilibrium with the set $\widetilde{B}$ of active players who randomize continuously on the interval $[0, v]$, where $\widetilde{B} = (B \cup \{j\}) \setminus \{i\}$.*

Corollary 10 suggests that the conditions for the existence of the equilibrium in which a relatively more risk averse player bids actively is sufficient for the existence of the equilibrium in which a less risk averse player bids actively, holding

[3]See detailed discussion and literature review in section 3.5.

all other active and inactive players constant, but the opposite is not necessarily true. This implies that mere differences in risk attitudes can result in different non-entry/dropouts decisions, without having heterogeneity in valuations or incomplete information. The player with the higher risk aversion may not participate in the competition because she finds that the potential returns from bidding any positive amount do not sufficiently compensate her for the risk.

One implication of this finding is that the well established gender difference in risk aversion (Croson and Gneezy, 2009) alone may be sufficient to explain differences in participation rates found in gender differences in competitiveness experiments (Niederle and Vesterlund, 2007)[4], without the need to hypothesize gender differences in competitiveness, confidence, or other characteristics.

A question naturally follows: are the dropouts always of the players who are more risk averse than the active ones? The answer is not necessarily. Example 10 suggests there might exist equilibria in which the intermediary risk aversion players drop out.

Example 10. *Assume there are three players with CARA utility functions* $U_i(x) = 1 - e^{-\beta_i x}$ *and a valuation* $v = 1$ *for* $i = 1, 2, 3$. *Assume also that* $\beta_1 = 2, \beta_2 = 1, \beta_3 = \frac{1}{2}$. *Then there exists an equilibrium in which only players* 1 *and* 3 *are active, while player* 2 *is inactive. See (3.11) below for the incentive constraint for player* 2 *and (3.12) for the feasibility constraints for player* 1 *and* 3.

$$\text{Incentive constraint:} \qquad K_1(x)K_3(x) = \frac{1-e^{-2x}}{1-e^{-2}}\frac{1-e^{-\frac{1}{2}x}}{1-e^{-\frac{1}{2}}} \leqslant \frac{1-e^{-x}}{1-e^{-1}} = K_2(x) \tag{3.11}$$

$$\text{Feasibility constraints:} \qquad K_1(x) = \frac{1-e^{-2x}}{1-e^{-2}} \leqslant 1 \text{ and } K_3(x) = \frac{1-e^{-\frac{1}{2}x}}{1-e^{-\frac{1}{2}}} \leqslant 1 \tag{3.12}$$

Thus, there exists an equilibrium in which the most and the least risk averse players are active while the player with the intermediary risk aversion is inactive.

3.4 Comparative statics

We now discuss the effect of increasing players' risk aversion on their expected bids. Our results in this section are derived for the equilibrium in which all active players randomize continuously in $[0, v]$. We first show in subsection 3.4.1 that if players are homogeneous in their risk attitude, then increasing all players' risk aversion decreases the total expected bid. Then, we show in subsection 3.4.2 that in contrast

[4]These papers examine entry into what are in effect all-pay auctions to measure gender differences in competitiveness. See Niederle (2014) for a recent survey.

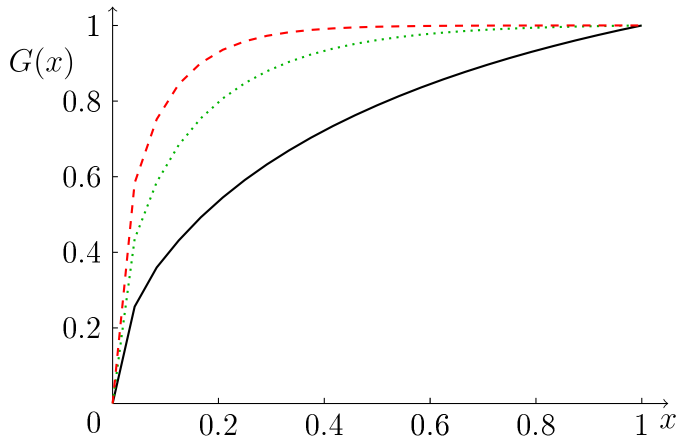

Figure 3.2: Homogeneously increasing all players' risk aversion decreases $G(x)$.

in the case with heterogeneous risk aversion, each player's expected bid increases with her own risk aversion, though it still decreases with other active players' risk aversions.

3.4.1 Homogeneous risk aversion

The equilibrium strategy with homogeneous risk aversion is a special case of the equilibrium strategy with heterogeneous risk aversion derived above.

Lemma 20. *Assume all the players are homogeneous, and there exists an equilibrium where a set B of players randomize continuously on the interval $[0, v]$ and all other players are inactive. If all the players' risk aversion increases homogeneously, then their bids and the total expected bid in the equilibrium in which the same set B of players randomize continuously on the interval $[0, v]$ are decreased in terms of first order stochastic dominance.*

We illustrate Lemma 20 with the following example.

Example 11. *Assume there are three players, each with the CARA utility function: $U_i(x) = 1 - e^{-\beta x}$ and valuation $v = 1$. Figure 3.2 shows the unique symmetric equilibrium strategy when $\beta = 1$ (black solid), 5 (green dotted), and 10 (red dashed). It is clear that as all players become more risk averse, the distribution function of their bids decreases in terms of first order stochastic dominance, i.e., the probability that they bid below x for any $x \in [0, v]$ is higher when they are more risk averse. The total expected bid decreases from 0.812 to 0.357 and then to 0.184.*

In the complete information all-pay auction with homogenous risk averse players, the total expected bid decreases in the players' risk aversion. As they become more risk averse, all players require better odds of winning in order to be compensated for the same risk. To maintain each others' indifference conditions as required by equilibrium, all players bid lower in the sense of first order stochastic dominance to compensate each other for the disutility of risk.

3.4.2 Heterogeneous risk aversion

In this section, we first show in Lemma 21 that each player's expected bid increases with her own risk aversion, but decreases with other active players' risk aversions. With the insights from this result, we characterize the sufficient condition for the total expected bid to decrease when the more risk averse players become even more risk averse in Proposition 15. We again assume that players can be ordered according to their risk aversion. Without loss of generality, let $K_1(x) \geqslant K_2(x) \geqslant \ldots \geqslant K_m(x)$ for all $x \in [0, v]$, so that player 1 is the most risk averse player.

Lemma 21. *Assume there exists an equilibrium where a set B of players randomize continuously on the interval $[0, v]$ and all other players are inactive. Assume, furthermore, that the level of risk aversion for some player $i \in B$ has increased, i.e., $K_i(x)$ changes to $\tilde{K}_i(x) \geqslant K_i(x)$ for every $x \in [0, v]$. Assume that after this change, there still exists an equilibrium where the set B of players randomize continuously on the interval $[0, v]$, and all other players are inactive. Then, the expected bid of player i increases with her level of risk aversion, while the expected bid of player k decreases with player i's level of risk aversion, for $k \in B, k \neq i$.*

The following example illustrates this result.

Example 12. *Assume there are three players with CARA utility functions $U_i = 1 - e^{-\beta_i c}$ with $\beta_1 = 1, \beta_2 = 0.5, \beta_3 = 0.1$ and valuation $v = 1$ for $i = 1, 2, 3$. Then, player 1 is the most risk averse and $K_1(x) \geqslant K_2(x) \geqslant K_3(x)$ for all $x \in [0, 1]$. In the equilibrium in which all three players are active and randomize continuously on the interval $[0, 1]$, the equilibrium strategies of the players (the CDFs of their mixed strategies) are given in the left part of figure 3.3: $G_1(x)$ (black) $\leqslant G_2(x)$ (green) $\leqslant G_3(x)$ (red). Assume now that β_1 changes to $\tilde{\beta}_1 = 1.2$, then the players strategies change to the dashed lines as in the right part of figure 3.3. It can be seen that player 1's mixed strategy (black) increases to $\widetilde{G}_1(x) \leqslant G_1(x)$ while players 2's (green) and 3's (red) mixed strategy decreases to $\widetilde{G}_2(x) \geqslant G_2(x)$ and $\widetilde{G}_3(x) \geqslant G_3(x)$ in the sense of first order stochastic dominance, respectively.*

Next we discuss the effect of a change in the risk attitude of an active player on the total expected bid in equilibrium. We first interpret the intuition behind Lemma

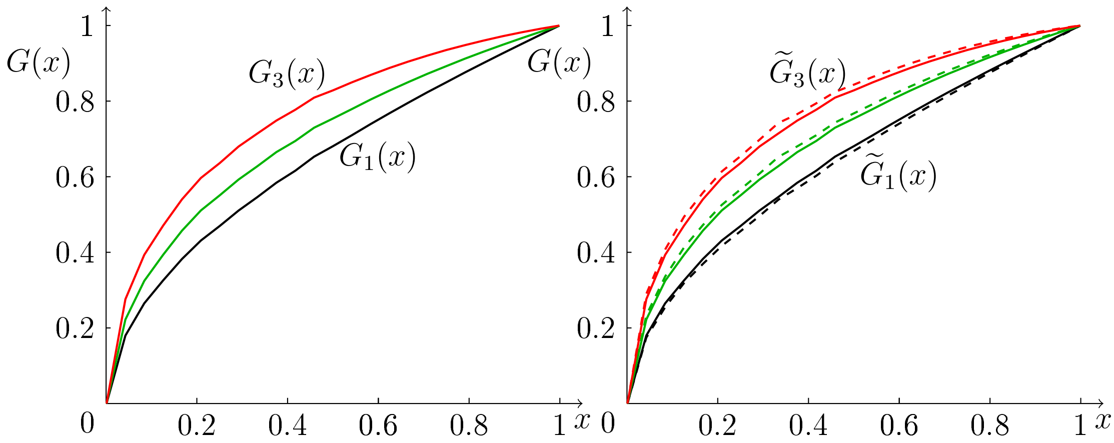

Figure 3.3: A player's bid increases with her own risk aversion and decreases with other players' risk aversions.

21. In a mixed strategy equilibrium, any active player t is made indifferent between any of his bids by the strategies of the other players. When player t becomes even slightly more risk averse, the other players will have to lower their bids to ensure player t stays indifferent (in order for an equilibrium with the same set of active bidders to continue to exist). Thus, by equilibrium strategy (3.4), increasing t's risk aversion has two effects on total expected bid, fixing the same set of active players:

1. Player t bids higher, since the CDF of her new equilibrium strategy first order stochastically dominates the CDF of her equilibrium strategy before she became more risk averse;

2. The rest of the players bid lower, since their CDF decrease in the sense of first order stochastic dominance when $K_t(x)$ increases.

The net effect on total expected bids is not obvious. We provide in Proposition 15 a sufficient condition for the total expected bid to decrease when one player's risk aversion increases, assuming the equilibrium with the same set of active players still exists after the increase of the player's risk aversion.

Proposition 15. *Assume an equilibrium with a set B of active players who randomize continuously on the interval $[0,1]$. For an active player i, if*

$$K_i(x) \geqslant \frac{|B|-2}{\sum_{l\in B, l\neq i} K_l(x)^{-1}}, \quad \textit{for all } x \in [0,v] \tag{3.13}$$

then, the total expected bid decreases in i's risk aversion.

Note that the r.h.s. of (3.13) can be rewritten into the harmonic mean of the $K(x)$ functions of the rest of the active players multiplied by a constant:

$$\frac{|B|-2}{\sum_{l\in B,l\neq i} K_l(x)^{-1}} = \frac{|B|-1}{\sum_{l\in B,l\neq i} K_l(x)^{-1}} \frac{|B|-2}{|B|-1}$$

Thus, condition (3.13) requires that player i be sufficiently risk averse compared to the rest of the active players to guarantee that an increase in her risk aversion decreases the total expected bid. See example 13 for an illustration of Proposition 15.

Example 13. *Assume there are three players $B = \{1, 2, 3\}$ who have CARA utility functions with $\beta_1 = 2$, $\beta_2 = 1$, $\beta_3 = \frac{1}{2}$, and valuation $v = 1$. Then, $K_1(x) = \frac{1-e^{-2x}}{1-e^{-2}}$, $K_2(x) = \frac{1-e^{-x}}{1-e^{-1}}$, $K_3(x) = \frac{1-e^{-\frac{1}{2}x}}{1-e^{-\frac{1}{2}}}$. Note that the condition (3.13) for $i = 1$ is satisfied:*

$$K_1(x) = \frac{1-e^{-2x}}{1-e^{-2}} \geqslant \left(\frac{1-e^{-\frac{1}{2}}}{1-e^{-\frac{1}{2}x}} + \frac{1-e^{-1}}{1-e^{-x}}\right)^{-1} = \left(K_2^{-1}(x) + K_3^{-1}(x)\right)^{-1}$$

and the total expected bid in the equilibrium in which all three players are active and randomize continuously on $[0, 1]$ is:

$$3 - \int_0^1 \left(K_2(x)K_3(x)\right)^{\frac{1}{2}} \left(K_1(x)\right)^{-\frac{1}{2}} dx - \int_0^1 \left(K_1(x)K_3(x)\right)^{\frac{1}{2}} \left(K_2(x)\right)^{-\frac{1}{2}} dx$$
$$- \int_0^1 \left(K_1(x)K_2(x)\right)^{\frac{1}{2}} \left(K_3(x)\right)^{-\frac{1}{2}} dx = 0.779$$

Assume now that we increase player 1's risk aversion to $\beta_1 = 3$, then the total expected bid decreases to 0.723.

Many real-life competitions are composed of participants with evidently different risk attitudes, e.g., mixed gender contests. We now analyze how the composition of contests of two different risk types affects participation. Formally, assume there are two sets of contestants in the competition: type 1 players with risk attitude defined by $K_1(x)$; and type 2 players with attitude defined by $K_2(x)$. There are m players in total. Let μ be the percentage share of type 1 players. Thus, the total number of type 1 players is μm, and similarly, the total number of type 2 players is $(1-\mu)m$. Note that $\mu \in \{0, \frac{1}{m}, \frac{2}{m}, ..., 1\}$ so that μm and $(1-\mu)m$ are always integers. All players have the same valuation v for the prize and $K_1(x) \geqslant K_2(x)$ for all $x \in [0, v]$. Type 1 players are thus more risk averse than type 2 players.

Based on Proposition 14, the feasibility constraints are the only conditions required for the existence of the equilibrium in which all m players are active and randomize continuously on $[0, v]$. The following result suggests that the feasibility constraints are sensitive to the number of players and the difference between the two types' risk aversion.

Corollary 11. *There exists an equilibrium in which all m players are active and randomize continuously on $[0, v]$ if and only if the number of type 1 players satisfies*

$$\mu m \leqslant \frac{\ln K_2(x)}{\ln K_2(x) - \ln K_1(x)}, \text{ for all } x \in [0, v]. \tag{3.14}$$

In other words, the number of the more risk averse players must be bounded from above given $K_1(x)$ and $K_2(x)$, to ensure that all players are active. The bound given by the r.h.s. of (3.14) depends on how different the two types of players are in terms of their risk preferences. The bound is lower when the difference between the two risk preferences is larger.

Corollary 12. *For any $\mu \in [0, 1]$, there exists an equilibrium in which all players randomize continuously on the interval $[0, v]$ if and only if*

$$K_1^m(x) \leqslant K_2^{m-1}(x), \text{ for all } x \in [0, v]. \tag{3.15}$$

Corollary 12 follows directly from Corollary 11. There always exists an equilibrium with all players active for any $\mu \in [0, 1]$, when (3.14) is satisfied for $\mu = 1$, which boils down to the inequality (3.15).

Corollary 13. *When there are two risk aversion types of players with $K_1(x) > K_2(x)$ for $x \in (0, v)$, assume condition (3.15) in Corollary 12 is satisfied. Then the total expected bid is monotonically decreasing with the share of the $K_1(x)$ players, μ, if*

$$K_2(x) \geqslant \frac{m-2}{m-1} K_1(x) \tag{3.16}$$

Corollary 13 explicates the transition in terms of total revenue from the case where all players are homogeneously less risk averse to the case where all players are homogeneously more risk averse. As the share of the more risk averse players increases, the total expected bid monotonically decreases. According to (3.16), this is true if the two types of players are not too different, as

$$K_1(x) \geqslant K_2(x) \geqslant \frac{m-2}{m-1} K_1(x)$$

has to hold.

3.5 Discussion

Our findings suggest the possibility that the higher risk aversion of women can simultaneously lead them to avoid participating in all-pay auction type incentives, while bidding higher and having a higher probability of winning than men, when they do participate. Heterogeneous risk aversion, therefore, could be an important factor in explaining most of the stylized facts about gender differences in competitiveness, including women's greater reluctance to enter contests with all-pay auction incentives, like elections, unless they have a good chance of winning (Fulton et al., 2006), women's greater willingness to exert effort in preparation (Duckworth and Seligman, 2006) and their higher odds of success in academic contests (Angrist et al., 2009; DiPrete and Buchmann, 2013; Fortin et al., 2015), and women's greater reluctance to enter laboratory contest (Niederle and Vesterlund, 2007), where either effort does not affect performance or for which they cannot prepare.

Moreover, our findings imply that if women are more risk averse than men, they will simultaneously work harder than men and decrease everyone's effort in the firm in personnel contests that have an all-pay auction structure. In these contests, if men and women are not too different in their levels of risk aversions, then a higher share of women may lead to increased odds of women dropping out. This result itself suggests an alternative and possibly more parsimonious explanation for the paucity of women in the upper management of firms in highly competitive industries.

However, our finding that the more risk averse player bids higher, and therefore, has a higher probability of winning, while at the same time depressing the bids of others suggests a further possible reason. While women may be more likely to win internal personnel contests, firms that promote women according to their individual competitiveness may suffer a general decrease in its competitiveness from the competition diminishing externality that women's greater risk aversion imposes on other personnel. Thus, firms which discriminate against women could do better against firms that do not. This finding suggests a potentially important exception to the intuition that competitive markets should eliminate taste-based discrimination (Becker, 2010). Moreover, our findings suggest that equilibrium discrimination against women in such industries should be stronger in countries in which the genders are more similar, i.e., developed rather than developing countries. Thus, the prospects for greater representation by women in competitive industries are not reassuring if women are indeed also less competitive than men (Niederle and Vesterlund, 2007). However, recent evidence suggests that when risk aversion is fully controlled for, women may actually be more competitive than men (Chen et al., 2015).

3.6 Appendix

Proof of Lemma 19

If player i becomes more risk averse (from $\tilde{U}_i$ to U_i), then the certainty equivalent of winning $v-x$ with probability $K_{\tilde{U}_i}(x)$ and $-x$ with probability $\left(1-K_{\tilde{U}_i}(x)\right)$ is less than zero, which is the certainty equivalent of this gamble with a utility function $\tilde{U}_i$. Therefore, the player can be restored to indifference between winning zero for sure and the gamble above only if the probability of winning the larger prize $v-x$ increases. Thus $K_{U_i}(x) > K_{\tilde{U}_i}(x)$. For loss averse players, we rewrite their $K_{U_i}(x)$ function as

$$K_{U_i}(x) = \frac{-l_i(-x)}{g_i(v-x) - l_i(-x)} = 1 - \frac{g_i(v-x)}{g_i(v-x) - l_i(-x)}$$

Thus, when player i gets more loss averse, $l_i(-x)$ gets smaller and $K_{U_i}(x)$ increases.

Proof of Corollary 8

If player s is more risk averse than player t, where $s, t \in B$, then we have $K_s(x) \geqslant K_t(x)$ for $x \in [0, v]$. Based on the equilibrium strategy given in (3.4), the difference in the distributions of their mixed strategies is:

$$\begin{aligned} G_s(x) - G_t(x) &= \left(\prod_{l \in B, l \neq s} K_l(x)\right)^{\frac{1}{|B|-1}} K_s(x)^{-\frac{|B|-2}{|B|-1}} - \left(\prod_{l \in B, l \neq t} K_l(x)\right)^{\frac{1}{|B|-1}} K_t(x)^{-\frac{|B|-2}{|B|-1}} \\ &= \left(\prod_{l \in B} K_l(x)\right)^{\frac{1}{|B|-1}} \left(K_s(x)^{-1} - K_t(x)^{-1}\right) \leqslant 0 \end{aligned}$$

Thus, player s's expected bid is higher than player t's and the cumulative distribution function of player s first order stochastically dominates the cumulative distribution function of player t. Therefore, the ranking of expected bids is the same as the ranking of risk aversion.

Proof of Corollary 9

The expected probability of winning for player s is given by (note that $K_s(v) = K_t(v) = 1$):

$$\int_0^v K_s(x)dG_s(x) = 1 - \int_0^v G_s(x)dK_s(x)$$

For player t:

$$\int_0^v K_t(x)dG_t(x) = 1 - \int_0^v G_t(x)dK_t(x)$$

Thus, the difference between the probabilities of winning is:

$$\begin{aligned}
& \int_0^v K_s(x)dG_s(x) - \int_0^v K_t(x)dG_t(x) \\
= & \int_0^v \left[G_t(x)dK_t(x) - G_s(x)dK_s(x)\right] \\
= & \int_0^v \left(\prod_{l\in B} K_l(x)\right)^{\frac{1}{|B|-1}} \left[\frac{dK_t(x)}{K_t(x)} - \frac{dK_s(x)}{K_s(x)}\right]
\end{aligned} \tag{3.17}$$

Therefore, the difference is non-negative if

$$\frac{dK_t(x)}{K_t(x)} - \frac{dK_s(x)}{K_s(x)} = \frac{K_t'(x)}{K_t(x)} - \frac{K_s'(x)}{K_s(x)} \geqslant 0$$

for all $x \in [0, v]$.[5]

Proof of Corollary 10

To prove the corollary, we need to show that after the replacement of player i with player j, the incentive constraint of player i and the feasibility constraint for player j are both satisfied. Note that after the replacement of j and i, the l.h.s of the incentive constraints and the feasibility constraints are weakened from $\prod_{l\in B} K_l(x)$ to $K_j(x)\prod_{l\in B, l\neq i} K_l(x) \leqslant \prod_{l\in B} K_l(x)$. Thus, we only need to show that the incentive constraint for player i:

$$K_j(x)\prod_{l\in B, l\neq i} K_l(x) \leqslant K_i^{|B|-1}(x)$$

[5] The reverse hazard rate dominance is not implied by the fact that $K_t(x)$ first order stochastically dominates $K_s(x)$. In fact, the reverse hazard rate dominance implies first order stochastic dominance. However, it is easy to show that the reverse hazard rate dominance condition is equivalent to first order stochastic dominance for CARA and CRRA utility functions. Readers can refer to Appendix B in Krishna (2009) which provides a useful introduction of stochastic dominance.

and the feasibility constraint for player j:

$$K_j(x) \prod_{l \in B, l \neq i} K_l(x) \leqslant K_j^{|B|-1}(x)$$

are satisfied after the replacement. Furthermore, since j is inactive in B, her incentive constraint must hold:

$$\prod_{l \in B} K_l(x) \leqslant K_j^{|B|-1}(x)$$

Therefore, player i's incentive constraint is satisfied since

$$K_j(x) \prod_{l \in B, l \neq i} K_l(x) \leqslant \frac{1}{K_i(x)} K_j^{|B|}(x) \leqslant K_i^{|B|-1}(x)$$

Player j's feasibility constraint in the new equilibrium is also satisfied since:

$$K_j(x) \prod_{l \in B, l \neq i} K_l(x) \leqslant \prod_{l \in B} K_l(x) \leqslant K_j^{|B|-1}(x)$$

Proof of Lemma 20

When players are homogeneous, we have $K_1(x) = K_2(x) = ... = K_m(x)$. By the equilibrium strategy given in (3.4), the strategy under homogeneous risk aversion is given by

$$G_i(x) = K(x)^{\frac{1}{|B|-1}}, \quad \text{where} \quad i \in B$$

It is then obvious that any active player i's bid is decreased in the sense of first order stochastic dominance when all players become more risk averse. The total expected bid can be calculated as

$$R = \sum_{i=1}^{|B|} R_i = \sum_{i=1}^{|B|} \int_0^v x dG_i(x)$$

where

$$R_i = \int_0^v x dG_i(x) = v - \int_0^v G_i(x) dx$$

is the expected bid of any player i. The second equality follows from integration by parts. Since $G_i(x)$ for $i \in B$ is increased, R_i is decreased and thus the total expected bid R is decreased when $K(x)$ increases for $x \in (0, v)$.

Proof of Lemma 21

An active player i's expected bid is given by

$$R_i = \int_0^v x dG_i(x)$$

where in equilibrium we have (from (3.4))

$$G_i(x) = \left(\prod_{l \in B, l \neq i} K_l(x)\right)^{\frac{1}{|B|-1}} K_i(x)^{-\frac{|B|-2}{|B|-1}}$$

Therefore, when $K_i(x)$ changes to $\tilde{K}_i(x) \geqslant K_i(x)$, then $G_i(x)$ decreases for every $x \in (0, v)$, and therefore, R_i increases. Moreover, for any other active players $k \in B$, $k \neq i$ we have

$$G_k(x) = \left(\prod_{l \in B, l \neq i,k} K_l(x)\right)^{\frac{1}{|B|-1}} K_k(x)^{-\frac{|B|-2}{|B|-1}} K_i(x)^{\frac{1}{|B|-1}}$$

Therefore, when $K_i(x)$ changes to $\tilde{K}_i(x) \geqslant K_i(x)$, then $G_k(x)$ increases for every $x \in (0, v)$, and therefore, R_k decreases.

Proof of Proposition 15

As each player j's expected bid can be written as $R_j = \int_0^v x dG_j(x)$, we can write the total expected bid R as:

$$R = \sum_{j \in B} R_j = |B| v - \int_0^v \sum_{j \in B} G_j(x) dx.$$

Rewrite the second term:

$$\int_0^v \sum_{j \in B} G_j(x) dx = \int_0^v (\prod_{l \in B} K_l(x))^{\frac{1}{|B|-1}} \sum_{j \in B} K_j(x)^{-1} dx. \tag{3.18}$$

Thus, the marginal effect of an increase of $K_i(x)$ for every given $x \in (0, v)$ on (3.18) can be written as:

$$\begin{aligned}
&\int_0^v \frac{d\left(\sum_{j\in B} G_j(x)\right)}{dK_i(x)} dx \\
= \;&\int_0^v \frac{1}{|B|-1}\left(\prod_{l\in B} K_l(x)\right)^{\frac{1}{|B|-1}} K_i(x)^{-1}\left(\sum_{l\in B, l\neq i} K_l(x)^{-1} - (|B|-2)K_i(x)^{-1}\right) dx,
\end{aligned}$$

This expression is positive if $\sum_{l\in B,l\neq i} K_l(x)^{-1} - (|B|-2)K_i(x)^{-1} \geqslant 0$ for all $x \in [0, v]$, which is condition (3.13). Therefore, the marginal effect on R is negative if the condition (3.13) is satisfied.

Proof of Corollary 11

Based on Proposition 1, the feasibility constraints in the current context are:

$$K_1(x)^{\mu m} K_2(x)^{(1-\mu)m} \leqslant K_1^{m-1}(x)\,, \text{ for all } x \in [0, v] \tag{3.19}$$

$$K_1(x)^{\mu m} K_2(x)^{(1-\mu)m} \leqslant K_2^{m-1}(x)\,, \text{ for all } x \in [0, v] \tag{3.20}$$

Rewrite the equation (3.19):

$$\left(\frac{K_1(x)}{K_2(x)}\right)^{\mu m} \leqslant \frac{K_1^{m-1}(x)}{K_2(x)^m}$$

i.e.,

$$\mu \leqslant 1 + \frac{\ln K_1^{-1}(x)}{m\left(\ln K_1(x) - \ln K_2(x)\right)} \tag{3.21}$$

Since the r.h.s. of inequality (3.21) is always larger than one, the feasibility constraint (3.19) always holds. Rewrite the equation (3.20):

$$\left(\frac{K_1(x)}{K_2(x)}\right)^{\mu m} \leqslant K_2^{-1}(x)$$

i.e.,

$$\mu \leqslant \frac{\ln K_2^{-1}(x)}{m\left(\ln K_1(x) - \ln K_2(x)\right)} \tag{3.22}$$

which is the condition in the corollary.

Proof of Corollary 12

Let the r.h.s. of inequality (3.22) be no less than one:

$$\frac{1}{m}\frac{\ln K_2^{-1}(x)}{\ln K_1(x) - \ln K_2(x)} \geqslant 1$$

After rearrange, we have

$$K_1^m(x) \leqslant K_2^{m-1}(x)$$

Therefore, whenever $K_1^m(x) \leqslant K_2^{m-1}(x)$ for all x, we have that for all μ there is an equilibrium in which all players active.

Proof of Corollary 13

Let $\mu \in \{0, \frac{1}{m}, \frac{2}{m}, ..., \frac{m-1}{m}\}$ be the current share of $K_1(x)$ players. Substitute a $K_2(x)$ player with a $K_1(x)$ player. Then, by Proposition 15 the total expected bid decreases if

$$K_2(x) \geqslant \frac{m-2}{\frac{\mu m}{K_1(x)} + \frac{m-1-\mu m}{K_2(x)}} \tag{3.23}$$

It can be verified that the r.h.s. of the above inequality is increasing with μ, and thus, is less than $\frac{m-2}{m-1}K_1(x)$. Therefore, condition (3.16) is sufficient for condition (3.23), and we have proved that increasing μ decreases total expected bid.

Chapter 4

Persistent Bias in Advice-Giving

Zhuoqiong Chen and Tobias Gesche

Abstract We show that a one-off incentive to bias advice has a persistent effect on advisers' own actions and their future recommendations. In an experiment, advisers obtained information about a set of three differently risky investment options to advise less informed clients. The riskiest option was designed such that it is only preferred by risk-seeking individuals. When advisers are offered a bonus for recommending this option, half of them recommend it. In contrast, in a control group without the bonus only four percent recommend it. After the bonus was removed, its effect remained: In a second recommendation for the same options but without a bonus, those advisers who had previously faced it are almost six times more likely to recommend the riskiest option compared to the control group. A similar increase is found when advisers make the same choice for themselves. To explain our results we provide a theory based on advisers trying to uphold a positive self-image of being incorruptible. Maintaining a positive self-image then forces them to be consistent in the advice they give, even if it is biased.

4.1 Introduction

When making risky decisions, we often seek advice. Doctors, investment advisers, scientists, and other experts have specific skills and knowledge to assess the potential consequences of important choices. Their job is to use their specialized information and skills to provide recommendations which are supposed to be in the best interest of patients, investors, politicians, and other clients. However, advisers may face a conflict of interest. Often, third parties pay commissions or create situations such that advisers owe them and then bias advice in their interest.[1] Advisers who give in to such third-party incentives can morally accommodate this behaviour by convincing themselves that they would have given the same advice, even if there had not been such a conflict of interest. For example, when a financial adviser recommends an investment fund as opposed to a less risky asset because of a sales commissions, this can later be justified by believing that it would have been the appropriate advice anyway. However, to uphold such a justification, the adviser has to act consistently. That is, an adviser has to issue the same biased advice even when the conflict of interest does not exist anymore.

This paper presents evidence for such persistent effects from advisers' conflicts of interest. In an experiment, we offer advisers a bonus which pays if they recommend less informed clients an investment option that is preferred only by risk-seeking individuals. Among advisers in a control group without such a bonus, almost no-one recommends this risky option. In contrast, almost half of the advisers to whom the bonus was offered do recommend it. Afterwards, advisers have to choose for themselves among the same options and then make a second recommendation for another client. For these tasks, it was explicitly stated that there would not be any bonus. Our results show that advisers who were previously exposed to the bonus were six times more likely to recommend the risky option than those who were not. We also find a similar increase in the probability that advisers choose the risky option for themselves. In consequence, being exposed to a conflict of interest in advice-giving in one single instance creates an externality on the advice which another client receives and the adviser's own choices.

We present a behavioural mechanism which can explain such persistent effects

[1] For example, US financial advisers administered more than $38 trillion for more than 14 million clients in 2011 (SEC, 2011). Despite laws like the Dodd-Frank Act which require them to "[...]*to act in the best interest of the customer*" (United States Congress, 2010, Sec. 913g), they receive sales commissions and bias their advice accordingly (Mullainathan et al., 2012; Malmendier and Shanthikumar, 2014). Other experts face such conflicts of interest too: Although supposed to be impartial, doctors reciprocate gifts from pharmaceutical companies (Dana and Loewenstein, 2003; Cain and Detsky, 2008) and scientists are dependent on industries sponsoring their research (Hilgartner, 2000; Taylor and Giles, 2005).

on repeated advice and advisers' own choices. It is based on the human tendency to interpret own actions to infer one's own morality (Mazar et al., 2008; Benabou and Tirole, 2011). To avoid a negative and immoral self-image, biased advisers can perceive their recommendations as those which they actually should have recommended, had they actually been impartial. However, when advisers morally accommodate their corrupted behaviour in such a manner, they have to stick with their advice. The reason is that changing it, in particular when the conflict of interest disappears, would signal to themselves that their initial advice was corrupted, and therefore, that they acted immorally.

Our results also show more exactly what advisers take as a reference for giving impartial advice and thus, how they try to keep a positive self-image: In principle, an adviser can internally disguise the fact that his advice was biased by forming a motivated belief (Kunda, 1990) about the clients' preferences, for example that a client is sufficiently risk-seeking.[2] In the adviser's view it is then in the client's best interest and therefore moral to recommend the risky option, even though the actual motive is the conflict of interest. This would not put the adviser under any pressure to act accordingly for himself, since his motivated belief is only about the client's risk preferences, not his own. However, prior research has shown that when forming beliefs about others' preferences, in particular risk preferences, we do so by starting from our own (Mullen et al., 1985; Faro and Rottenstreich, 2006). The question "What would I choose if I were in the client's situation?" then also determines what an adviser should recommend. Under such a rule, advisers who want to perceive themselves as incorruptible should then also choose for themselves what they have recommended to others. Our data indicate that this is the case: Having been exposed to a bonus leads advisers to choose the risk-seeking option more often. This is in line with the recent findings of Linnainmaa et al. (2016). In a large sample, they show that financial advisers hold the same expensive, under-performing portfolios as their clients, even after having left the industry.

Related literature: Our work combines findings from self-signaling, motivated beliefs, and self-deception to obtain new insights about their implications in the context of advice-giving. It captures the fact that people assign informational value to their actions to infer about their personal traits (Bodner and Prelec, 2003; Benabou and Tirole, 2004) and in particular their moral values (Benabou and Tirole, 2011). Self-signaling then means that actions are also influenced by the consequences they subsequently have on peoples' self-image. For example, Mazar et al. (2008) argue

[2]Without referring to any actual gender roles we will call advisers and clients "he" and "she", respectively.

that we often do not lie as much as we, in principle, could because strong, outright lies would damage our self-perception of being honest and moral persons. Gneezy et al. (2012) present the seemingly paradoxical finding that sales under a pay-as-you-want scheme are lower than under a low, fixed price. They explain the consumers' reluctance to set a sufficiently low pay-as-you-want-price with consumers' desire to not perceive themselves as greedy. Related to this, Fallis et al. (2015) report that the demand for goods which a share of the sales price is donated is increasing in this price. They also present evidence that this is due to the decrease in social image utility which consumer derive from purchasing such good-donation-bundles.

Prior research has also shown that when it comes to morally-ladden situations, people form self-serving assessments about what norms should apply and about others' preferences when it helps them to obtain a positive, moral self-image. Loewenstein et al. (1993) give subjects information about legal cases. These subjects then differ strongly in what they consider as appropriate, fair settlement values for these cases after they argued in fictitious roles of being the plaintiff as opposed to the defendant. Di Tella and Pérez-Truglia (2015) show evidence that people form beliefs about others behaving anti-socially, i.e. that others steal from a common pot, in order to justify their own anti-social behaviour of not splitting the pot equally. People also employ uncertainty and ambiguity in a related manner to form self-serving beliefs and probability assessments which allow them to obfuscate their own immoral behaviour (Haisley and Weber, 2010).

In this paper, we connect these findings to obtain insights about their lasting implications in the context of advice-giving. Closely related to our results is Gneezy et al. (2016): In several experiments, the authors show that advisers bias their recommendations relatively strongly when they learn about their conflict of interest before they receive the information about a client's decision situation. When they first learn about the situation, then consider what to recommend, and then about the conflict of interest, their advice is less biased. Following Trivers (2011), they label this behaviour self-deception. Our theory and results describe behaviour which is in line with such self-deception, i.e. that advisers effectively bias their own choices. We make the point that the reason for this behaviour and the consistency in advisers' biased recommendation is that advisers try to avoid a negative self-inference.[3] This also relates to Konow (2000), who examines a dictator game where the pie to be split

[3] Falk and Zimmermann (2016) show that agents also act consistently to signal their skills to a principal. In Falk and Zimmermann (2015), they provide evidence that people act consistently without any external observers. The general idea which underlies the mechanism we propose also applies in these settings: Acting inconsistently shows that one's first action was somehow flawed, acting consistently therefore avoids such a inference to oneself and/or outside observers.

is dependent on the dictator's and the recipient's prior joint effort. He finds that dictators who allocate themselves larger shares of the pie interpret their personal contribution in establishing the common pot more favorably than outside observers. Documenting the persistence of such a self-serving bias, these dictators apply their persistent biased judgment about others' effort when they act as outside observers themselves.

Recent findings on actual advisers' behaviour by Linnainmaa et al. (2016) relate to ours. Using matched data on about 5900 Canadian financial advisers and their more than 580,000 clients, these studies show that the most important determinants of advice to these clients are not the clients' personal characteristics, but rather the identity of their advisers. Even more important in our context, they show that these recommendations to clients are also reflected by the choices which advisers make for their own portfolios. For example, advisers prefer the return-chasing and actively managed funds they sell to clients also for themselves. This is puzzling since these investments do not perform better than the market. When fees are subtracted, clients' and their advisers' investments even significantly under-perform relative to the market. Our results and the theory we propose resonate with these findings. In addition, the experimental setup we use allows to abstract from concerns of advisers self-selecting into suitable environments which may drive such findings (e.g. risk-seeking advisers who choose to sell risky investments with sales commissions).

We identify a strong, causal, and lasting effect of bonuses in advice-giving. Our findings therefore contribute to the recent literature on the adverse effects of bonus payments (Agarwal and Itzhak, 2014; Bénabou and Tirole, 2016). We also point out the role of self-signaling in such a setting which connects directly to the recent research on the work culture and self-perception of those working in the financial industry (Cohn et al., 2014; Zingales, 2015). However, our findings apply also outside this specific financial context to advice on risky decisions more generally.

In the remainder of this paper, we present our findings in more detail. The next section describes a mechanism of how moral and self-image concerns can lead to persistent bias after advisers have faced a conflict of interest. Section 3 explains the design and procedures of the experiment in which we investigate this mechanism. Section 4 derives predictions and section 5 presents our results. Section 6 concludes by reviewing these results with respect to their implications for the economics of motivated beliefs, advice giving and its regulation. An appendix contains a formal model in which the predictions are derived; it also contains further data analysis and the experimental instructions.

4.2 Mechanism

In this section, we describe a behavioural mechanism in which advisers' concerns to appear impartial and moral can lead to the opposite behaviour – a persistent bias in their advice. The framework presented here also provides the assumptions that underlie a formal model which can be found in the appendix. To analyze an adviser's behaviour, we assume his overall utility to depend on three parts: 1) consumption utility derived from monetary payoffs, 2) the moral cost of not giving impartial advice, and 3) diagnostic (dis-)utility of learning from actions which reveal that one's previous advice was biased.

While the first element of an adviser's overall utility is standard, the second reflects the fact that advisers might feel compelled, and often are, to act solely in a client's best interest. Not doing so then creates a moral cost. To determine when such a cost occurs, the question then arises what constitutes a "client's best interest", i.e. what constitutes impartial advice. We assume that an adviser can form a belief about his clients' preferences and therefore about the utility that clients experience when they follow his advice. Giving advice which does not maximize this assumed utility of the client would then be a violation of giving impartial advice and creates the moral cost. However, predicting others' preferences is inherently difficult. This applies in particular for risk preferences (Hsee and Weber, 1997; Eckel and Grossman, 2008; Harrison et al., 2013), even when the inference is conducted by trained financial advisers and there is no conflict of interest (Roth and Voskort, 2014). In the presence of external incentives which creates such a conflict, the uncertainty in estimating others' risk preferences can be instrumentalized in a self-serving manner: Advisers may form a belief about their clients' preferences such that their, potentially biased, advice is compatible with it.

However, there are limits to such self-serving beliefs. It is a robust psychological fact that people base their inferences about others' preferences on their own (Marks and Miller, 1987), in particular for risk preferences (Faro and Rottenstreich, 2006).[4] In consequence, advisers' own preferences also play a role in determining what is impartial advice. We capture this by assuming that advisers incur a moral cost when they recommend an option which they would not choose for themselves if they were in the client's position.

[4]Though initially coined by Ross et al. (1977) as a "false consensus effect", the falsity of estimates of others' preferences based on one's own is not evident. Works by Hoch (1987) and Dawes (1990) demonstrate that such projection is not just statistically correct; they also show that people can often improve their accuracy in predicting others' preferences by relying more strongly on their own. Engelmann and Strobel (2000) show that subjects do so when they are incentivized to make accurate predictions.

The third factor which matters for advisers is the diagnostic (dis-)utility they derive when they learn to have given biased advice, based on a model of self-signaling (Bodner and Prelec, 2003; Benabou and Tirole, 2011). In contrast to the moral cost of acting immorally, this dis-utility only occurs to an adviser *after* he has biased his advice, at the point when his later actions indicate exactly this fact to him. This can be captured by a dual-self model in which the "diagnostic self" of an adviser learns ex post about the other self's motive for giving advice, e.g. whether prior advice was issued impartially and therefore was morally sound or whether it was corrupted. The important implication of such an inference is that advisers can only uphold a positive and self-serving belief of their prior motives for giving advice as long as they do not take actions which are incompatible with this.[5] Dual-self models have been used previously to explore how people infer about themselves, in particular their moral behaviour (Benabou and Tirole, 2004; Grossman, 2015). Here, we use it as a crucial device to describe the trade-off between keeping self-serving beliefs about one's own motives and taking contradictory actions.[6]

These three components together then have implications for how and, most importantly for how long, a conflict of interest affects advisers' choices and their recommendations. To see this, consider an adviser who issued a biased advice, thus an adviser whose pecuniary payoff for biasing advice outweighed his moral cost of doing so. If he is also concerned about the self-image, he then needs to continue to give the same biased advice again, especially when the conflict of interest has disappeared. The reason is that in order to later entertain the (counterfactual) idea that his initial advice was unbiased, it should be unaffected by the presence of an external incentive. Changing advice when the the conflict of interest disappears would then signal just the opposite. When an adviser's own preference stipulates what he should recommend to a client, this mechanism has even further consequences. This is because such a rule implies that in order to perceive oneself as unbiased, an adviser has to act according to his biased advice for himself.

In consequence, a behavioural trait which generally seems to be desirable, the preference to perceive oneself as a moral person, can lead to persistent biases in the context of advice giving. In addition, it can have a lasting effect on advisers'

[5] In essence, this reflects the desire to avoid cognitive dissonance (Festinger and Carlsmith, 1959) – a discrepancy between one's actions and one's beliefs about what is the norm one should follow (for economic models of cognitive dissonance, see Akerlof and Dickens (1982) and Rabin (1994)). For a discussion about how cognitive dissonance and motivated (self-)perception relate see also Kunda (1992).

[6] Apart from enabling us to capture this cognition, it also captures the fact that the inferring self "forgets" about the other self's motives. This is in line with research showing that people cannot perfectly recall their past decision motives nor foresee their future ones (Kahneman et al., 1997; Loewenstein and Schkade, 1999)

own choices to the degree that they assign diagnostic value to them. With this behavioural mechanism in mind, we set up the following experimental design to explore it in more detail.

4.3 Experimental design and procedures

At the beginning of the experiment, subjects were allocated to computer terminals in cubicles where instructions were shown to them on screen. Subjects acting as advisers were then informed that they would get GBP 5.00 as a show-up fee for participating in the experiment and that there would be further possibilities to earn money. They were also informed that they would act as advisers for clients who would be drawn from the same pool of subjects for a future experimental session and that clients would also receive the same show-up fee.

It was then explained to advisers that they would have to recommend which out of three investments, referred to as option A, B, and C, their clients should take. They were told that clients would only know that option A's payoff would depend more on luck than option C's while option B is intermediate in this regard. They were also told that clients would not know the options' payoffs or the associated probabilities. Advisers were informed that they, as advisers, would soon learn these exact parameters of the investment options before they had to make a recommendation.

The advisers' superior information was then given to them on a paper sheet which explained the three investment options in detail (for a copy of this sheet and the experimental interface see the appendix). The text on the sheet explained the following procedure of how an option's payoff was determined: After an option was chosen, a six-sided die would be rolled. Depending on the chosen option, this would then yield either a safe payoff or a lottery. This lottery was described as a (fair) coin toss with heads yielding GBP 20 and tails nothing. The following table which was also on that sheet summarizes how the die's result maps into these possible outcomes, depending on the chosen investment option:

Die equal to:	**Option A**	**Option B**	**Option C**
1 or 2	lottery: GBP 20 or 0	safe payment: GBP 12	safe payment: GBP 12
3 or 4	lottery: GBP 20 or 0	lottery: GBP 20 or 0	safe payment: GBP 8
5 or 6	lottery: GBP 20 or 0	lottery: GBP 20 or 0	lottery: GBP 20 or 0

Table 4.1: Description of the investment options as shown to advisers, "lottery" is a coin toss.

The text explained this procedure in detail and also contained several examples.

Note that a choice among the three compound lotteries which these three options represent, allows to categorize the underlying risk preferences.[7] Comparing the differences between option A and B, only those who are willing to give up a safe payoff of GBP 12 to play a lottery with an expected payoff of GBP 10 instead, i.e. risk-seeking individuals, choose option A. Conversely, option C is preferred to option B only by those who want to sacrifice an expected payoff of GBP 10 for a safe payoff of GBP 8. Thus, only risk-averse individuals choose option C. Accordingly, Option B is chosen by individuals who are neither sufficiently risk-averse nor sufficiently risk-seeking. Reflecting this ordering based on risk-preferences we will henceforth, with slight abuse of the precise meaning, refer to option A/B/C as the "risk-seeking/neutral/averse option".

Step 1 – First recommendation R1:

After having studied the instructions and choice situations, advisers were asked to make a recommendation to clients. For this, they had to write the sentence "I recommend you to choose option A/B/C", depending on what they wanted to advise, on a piece of paper which had their cubicle number on it. They were instructed to put this recommendation into an envelope, close it, and then click on a button on their screen. The envelope was then collected by an experimenter and put into a box. Before they made their recommendations, they were told that at the end of the experiment, one of the envelopes would be randomly drawn from the box to be presented to a client and that the corresponding cubicle number would be read aloud. An adviser thus knew that he would eventually know whether his recommendation was chosen to a be shown to a client.

Step 2 – Own choice O:

After all advisers had written down their recommendation R1 and all envelopes were collected, they were informed that they would now have to choose an investment option for themselves. Advisers were previously not informed about this step. The procedure was the same as for issuing advice: Subjects had to write on a sheet "I choose option A/B/C." and then put it in an envelope. An experimenter came by and collected the envelopes and put it in a separate box. Again, they were informed that at the end of the experiment, one of the envelopes would be chosen randomly, its number would be announced aloud, and that the respective adviser would be asked later to roll the die to determine his chosen option's payoff. Ex-ante, the choice situation and its implementation probability was thus the same as the one on which they had previously advised a client on.

[7] This choice between possible sub-lotteries within a compound lottery is essentially a stripped-down version of a similar task used previously by Hsee and Weber (1997).

Step 3 – Second recommendation R2

After advisers made their own choice O, they were asked to make a second recommendation. The procedure was exactly the same as for R1, including the collection of envelopes in a separate box, sampling one from it and announcing its number. Again, advisers did not know in advance about this step. Advisers were also informed that this second advice, if it was sampled, would be shown to a different client in the same future session with clients.

Step 4 - Questionnaire and implementation:

After all recommendations were collected, subjects filled out a short questionnaire which elicited personal characteristics. The experimenter then sampled one envelope from each of the boxes which contained the envelopes for R1, O, and R2 and announced the respective cubicle numbers. Subjects were then paid out in private based on whether they were offered a bonus and their recommendations; the subject in each session whose own choice O was sampled also rolled the die and received the corresponding payoff.

NO BONUS versus BONUS treatment: The above describes the experimental procedure in our baseline condition to which we will refer as NO BONUS. Our experimental manipulation was to offer some advisers a bonus for recommending the risk-seeking option A in R1. We will refer to this treatment as BONUS. After having been informed about the advice they had to give and how to do so, but before seeing the sheet with the detailed information about the investment options, every second adviser (in total 48) in a given session was randomly determined to be in that treatment. These advisers were informed that they would get a bonus of GBP 3 if they recommended option A. This bonus was only paid for subject's first recommendation R1. For those advisers who were offered the bonus, there were explicit notifications on the screens which explained the O and R2 tasks which clearly stated that there would not be any bonus for these tasks.[8] This within-session, across-subjects intervention with regard to the bonus is the only difference between our NO BONUS and BONUS.

Verifiability: In order to ensure that advisers believed that a recommendation, if randomly chosen to be shown to a client, would be actually seen by the client we allowed advisers to sign their recommendations and to address the envelopes to themselves. Advisers were explained that if their recommendation was chosen to be shown a client, the sheet would be signed by the respective client. In case

[8]Since advisers' payoff in BONUS do not depend on the clients' decisions, they were not explicitly informed about whether clients would learn about the bonus. Also none of the advisers asked for this information. In the session with clients, they were informed of the bonus when they received a recommendation R1 from an adviser who had been in the BONUS treatment.

that the corresponding adviser had provided us with his or her address, this subject would then get a copy of the signed recommendation by post. In addition, they were informed that this mailing would also contain information on how they could see the original, signed receipts which were deposited with the lab's official record depository. Subjects were informed of this before making their first recommendation. Since an adviser knew that he would know whether his envelope was sampled, this procedure pre-committed us to actually show the sampled advice letters to actual clients.

General procedures: Throughout the experiment, we enforced a strict no communication policy. We conducted eight sessions, each with 11 to 14, in total 99, subjects acting as advisers. Advisers earned on average GBP 6.68 ($9.51 at the time of the experiment) while no session lasted longer than 45 minutes. All subjects were students across several degrees and fields of studies. Table 12 in the appendix shows descriptive statistics. The experimental sessions were conducted in late January 2016 at the London School of Economics's Behavioural Research Lab with subjects from its pool. The experimental interface was implemented using zTree (Fischbacher, 2007). A week after the eight adviser sessions, we invited 16 additional subjects from the same pool for an additional session. In this session, they acted as clients and received the sampled recommendations from the previous adviser sessions, made their choices, and were paid their resulting payoffs. In this paper, we only focus on advisers and their recommendations.[9]

4.4 Predictions

In this section, we derive predictions for our experiment. They are based on the assumptions which we described in section 2, thus on advisers maximizing their overall utility from pecuniary payoffs, the moral cost of giving in-appropriate advice, and the self-image concern. Given our treatment intervention, we make the predictions with regards to how often the risk-seeking option A is recommended and chosen. All predictions derived and presented in this section are also derived in the formal, mathematical model which can be found in the appendix.

Predictions for R1: In NO BONUS, there is no pecuniary gain of issuing any specific recommendation. Since this is the first choice which an adviser makes it does

[9] With only 16 client observations which are not balanced over treatments (only three are eventually with recommendations from BONUS; recall that the probability of a recommendation being chosen is independent of the treatment), any analysis of client would have limited statistical power. However in the two experiments of Gneezy et al. (2016) which are in a related setting but have much more client observations, clients followed advisers in 74% and 85% of all cases, respectively.

not have signaling value with regards to past behaviour. Absent pecuniary motives, only the moral cost of issuing inappropriate advice therefore remains. Beliefs about client's preference can be formed in a self-serving way, i.e. such that they suit an adviser's recommendation, up to the point that they contradict his own preference. To minimize the cost from recommending something that one would not choose for oneself, advisers thus recommend option A only if they prefer it. Thus, only risk-seeking adviser recommend option A.

In the BONUS treatment this is different: Advisers are now paid for recommending option A and derive pecuniary utility from the bonus when they do so. Clearly, those who would have recommended it anyhow, i.e. risk-seeking advisers, also recommend it in this treatment and in addition, get the bonus. However, those who would not have recommended it in the NO BONUS because they do not prefer it themselves now face a trade-off: When the moral cost of recommending something they would not choose for themselves are smaller than the pecuniary value of the bonus, they recommend option A. Otherwise, they recommend their preferred option. In both cases, they hold self-serving beliefs about the client's preference which is compatible with their issued advice. Assuming that some advisers have sufficiently low moral cost and follow the offered bonus, we get the following prediction:

Prediction 1. *There are more advisers in BONUS than in NO BONUS who recommend option A for the first recommendation R1.*

Predictions for O: In contrast to the first recommendation, advisers now make choices for themselves. The moral cost of giving inappropriate advice are therefore absent. Since the NO BONUS did not feature a bonus, there was no incentive to act immorally and to give biased advice. In consequence, there is no concern about drawing any (negative) inference from the own choice about one's preceding advice. The only relevant decision criterion is thus one's own risk preference and only risk-seeking advisers should choose option A for themselves in NO BONUS.

The own choice situation in BONUS and the NO BONUS is identical. Differences in behaviour must occur because advisers in BONUS have previously been exposed to the bonus and, potentially, have given in to it. To the degree that they assign diagnostic value to their choices, advisers' own choices can then reveal to themselves that they were corrupted by the bonus: Advisers who recommended option A in R1 should, in order to appear as having given appropriate advice, also prefer it for themselves. In order to uphold the self-image that they were not corruptible, advisers who recommended option A just for the bonus must then mimic the incorruptible ones by choosing option A for themselves.[10] However, these advisers lose expected pecuniary utility because they choose the option which they do not actually prefer. In consequence, only those corruptible advisers who have sufficiently high image concerns, relative to their loss in expected pecuniary utility, choose option A for themselves, in addition to the incorruptible, risk-seeking ones. Note, however, that this only applies if own choices have sufficient diagnostic value, i.e. if advisers acknowledge the reverse implication of "I should recommend to my client what I would choose in her situation". Under the assumption that advisers assign such diagnostic value to their own choices we predict the following:

Prediction 2. *There are more advisers in BONUS than in NO BONUS who recommend option A for the the own choice O.*

Second recommendation R2: The predictions for the second recommendation combine insights from above. In NO BONUS, an adviser's pecuniary utility is unaffected by his second recommendation. Also, absent any previous bonus to give inappropriate advice, self-signaling concern do not play any role either. Accordingly, only the moral cost for giving inappropriate advice matters, as in R1. A previously formed self-serving belief coincides with the previous recommendation. For this recommendation, an adviser's own preference was the determining factor so that again, only risk-seeking advisers recommend option A (again).

In the BONUS treatment, the second recommendation does not entail any bonus either. However, the bonus which was offered to advisers in R1 opens the possibility that this recommendation was biased and therefore, the concern for signaling one's own corruptibility matters. Advisers who truly prefer option A can then minimize the moral cost of giving inappropriate advice and the self-signaling concern by recommending option A again in R2. As outlined above, advisers who do not prefer option A but recommended it in R1 for the bonus may mimic the incorruptible ones by choosing option A in O to prevent dis-utility from learning that they gave biased advice. Following the same logic, they can then mimic the incorruptible ones by re-recommending option A in R2. Note that the situations in R2 and R1 are identical, except for the bonus. Therefore, an inconsistency is more directly attributable

[10] In terms of a signaling model, this is an equilibrium where corruptible advisers pool with those who truly prefer option A. In principle, there could be other equilibria where corruptible advisers and those who truly prefer option A pool on choosing non-A options, together with incorruptible advisers who actually prefer these options. However, in terms of self-signaling, these are rather unrealistic equilibria. This is so because in such equilibria, those who behaved morally obfuscate their behaviour while those who behaved immorally do not. We therefore exclude them. We discuss this in more detail in the formal model in the appendix. There, we also show that these excluded equilibria do not even need to exist. In contrast, the former one where corruptible advisers mimic incorruptible ones by choosing option A does always exist.

to one's corruptibility; the second recommendation should have higher diagnostic value than the own choice. We thus get the following prediction:

Prediction 3. *There are more advisers in BONUS than in NO BONUS who recommend option A for the second recommendation R2.*

Conditional on a scenario in which at least some advisers are corrupted by the bonus, thus that prediction 1 is true, our design enables us to investigate two main questions. First, by testing prediction 3 we can find evidence for self-image concerns which cause repeated bias in advice-giving. If advisers are only steered by pecuniary incentives and not by the diagnostic value of their actions, we would not expect differences between BONUS and NO BONUS. In addition, comparing the own choice O across treatments allows to test whether they also have diagnostic value. If they do not, advisers should just implement their preferred choices which, due to random treatment assignment, should not differ between BONUS and NO BONUS. However, if prediction 2 is also confirmed, this indicates that advisers make choices which are, from a purely pecuniary point of view, sub-optimal just to appear incorruptible. It would therefore indicate that they assign diagnostic value to their own actions.

With this in mind, we will next examine the actual advisers' behaviour in our experiment. Before doing so, it is noteworthy that the proposed mechanism is, in principle, also capable of explaining the findings by Gneezy et al. (2016). They report on an experiment in which they expose advisers to a bonus and to a decision situation similar to ours. They then examine the effect of when this exposure to the bonus happens. They find that recommendations are less affected by the bonus when advisers learn about it after they have first considered what to recommend. In contrast, when they know about the bonus before such a consideration, their following advice is more biased. If the act of actively considering what to recommend also has diagnostic value, then changing one's actual recommendation afterward, once one has learned about the bonus, would also signal one's corruptibility. If in contrast an adviser knows from the beginning about the bonus, this can already be taken into account when initially considering what to recommend. He can then form a self-serving belief which supports his biased consideration and therefore also the actual recommendation. This would prevent a negative self-inference.

4.5 Results

Results for R1: This is where our treatment manipulation occurred. In the BONUS treatment, advisers were paid a bonus to recommend option A. Accordingly, we expect some to give in to this incentive and recommend it. This is also what we observe: In the NO BONUS only 3.9% of advisers recommend option A in their first recommendation. In contrast, about half of all advisers (54.2%) in BONUS recommend this option – an increase by 50.3 percentage points which is highly significant (Fisher exact test: $p = 0.000$).[11] Figure 4.1 shows the overall distribution of choices across these treatments:

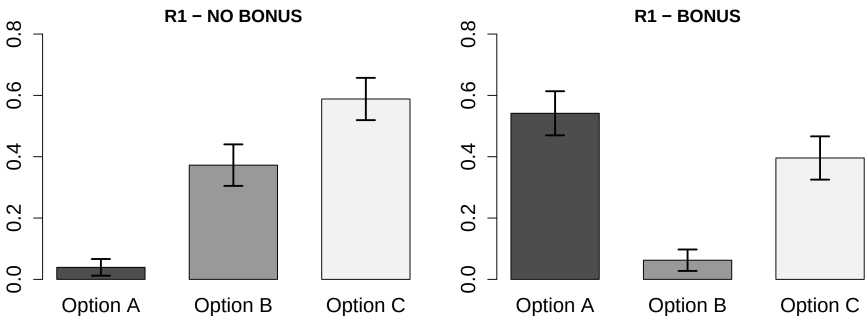

Figure 4.1: Frequency for each option being recommended in R1, bars depict standard errors.

We also employed a parametric approach via the following linear probability model which allows us to control for the effect of remaining heterogeneity across treatments or sessions:

$$\text{Prob}[r_{1,i} = A] = \alpha + \beta \cdot BONUS_i + \boldsymbol{\delta} \cdot \mathbf{c}_i + \boldsymbol{\gamma} \cdot \mathbf{s}_i + \epsilon_i \tag{4.1}$$

In the above, $r_{1,i}$ is subject i's first recommendation out of the set of possible recommendations $\{A, B, C\}$ and $BONUS_i$ is a dummy indicating whether this subject was randomly assigned to the treatment BONUS. The vector $\mathbf{c}_i$ collects control variables which indicate a subject's age, gender, monthly budget, dummies for regions of origin, the highest degree a subject holds or pursues and his or her fields of studies. Control dummies for each session are collected in $\mathbf{s}_i$. The error term ϵ_i captures idiosyncratic noise in the decision for an adviser's recommendation. Table 4.2 presents the results when controls are successively added. It shows that the increase of about 50 percentage points in the probability of recommending option A is almost unaffected by the addition of these controls and remains highly significant. We also repeat the same estimation procedure by probit and do not find any qualitative differences (see table 8 in the appendix). We therefore note that our treatment manipulation worked and that prediction 1 is confirmed.

[11]Although we have directed hypothesis, the reported p-value here and in the the following always refer to more conservative two-sided hypotheses.

	(1)	(2)	(3)	(4)
BONUS	0.502***	0.497***	0.489***	0.481***
	(0.078)	(0.076)	(0.093)	(0.092)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	99	99
Adjusted R^2	0.304	0.323	0.280	0.310

Table 4.2: OLS estimates of the probability to recommend option A in R1 robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
personal controls: age, gender, monthly budget, subject's region of origin and field of studies

It is also noteworthy that our results indicate that when offered a bonus, almost half of our subjects do *not* recommend option A. If subjects were confused or indecisive we would expect them all to take the money. However, there is something which stops a significant share, 45.8% (t-test: $p = 0.000$), of all advisers in BONUS from recommending this option, even for money. The notion of advisers refusing to recommend it because they consider it inappropriate or immoral advice is consistent with this observation.

Own choice O: For their own choice, no bonus is paid to advisers in both conditions. Figure 4.2 displays their choices. In the baseline NO BONUS we observe that 9.8%

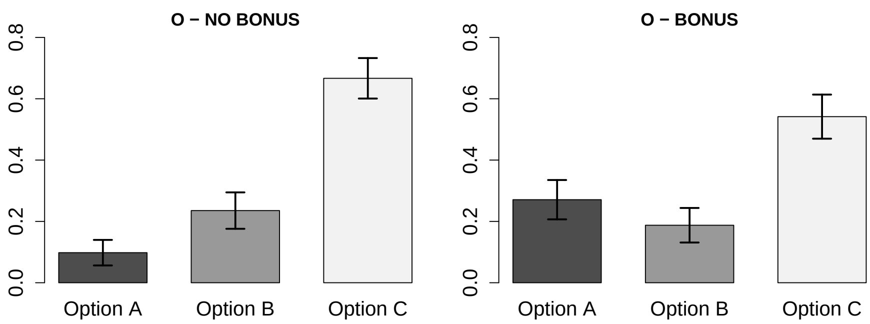

Figure 4.2: Frequency for each option being chosen in O, bars depict standard errors.

choose option A for themselves. In BONUS however, when advisers were *previously* offered the bonus for their first recommendation, 27.1% of all advisers, almost three times as much as in NO BONUS, choose the risk-seeking option A for themselves. This increase by 17.3 percentage point is significant (Fisher exact test: $p = 0.036$).

This finding is also confirmed when we re-estimate model (4.1) with a dummy indicating that an adviser chooses option O for himself as the dependent variable. Table 4.3 reports the corresponding results when the same control variables as in the preceding analysis are successively added. The effect of being in BONUS even increases and this pattern is again similar when the model is estimated by probit (see table 9 in the appendix). Therefore, we regard prediction 2 as confirmed.

	(1)	(2)	(3)	(4)
BONUS	0.173**	0.178**	0.219**	0.218**
	(0.077)	(0.081)	(0.095)	(0.087)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	99	99
Adjusted R^2	0.040	0.010	0.065	0.088

Table 4.3: OLS estimates of the probability to choose option A for oneself in O
robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
personal controls: age, gender, monthly budget, subject's region of origin and field of studies

Given these findings, it is helpful to recall the mechanism which underlies our prediction since we can examine this causal channel more closely. The mechanism argues that if advisers assign diagnostic value to their own choice, they have to act according to their (biased) advice in order not to self-signal that they were corrupted. Our findings for R1 indicate that the bonus corrupted about half of all advisers; it leads to an increase of recommending option A by 50.3 percentage points for BONUS relative to NO BONUS. The findings on advisers' own choice O just presented, show that there is an increase of 17.3 percentage points for those who were potentially corrupted, i.e. those who were exposed to the bonus. These estimated probabilities then imply the share of advisers who choose option A for themselves because they have previously given in to the bonus but do not want to self-signal their corruptibility is given by 34.4% ($\triangleq 0.173/0.503$).[12] This estimate shows that

[12] This follows from re-arranging the following: The observed increase between NO BONUS and BONUS in own choices for A (17.3%) has, according to the described mechanism, to equal advisers' propensity of feeling compelled to choose option A for themselves due to their previous recommendation for it, multiplied with the increase in the probability of them recommending option A as caused by the bonus (50.3%). To capture this effect in our regression framework, we implemented the following two-stage procedure: In the first stage, we took our regression results for (4.1) to obtain an estimate of how strongly the bonus lead advisers to recommend option A. To see how this causal channel affected their own choice, denoted by c_i, we then estimated in a second step the model $\text{Prob}[c_i = A] = \alpha + \beta \cdot \widehat{r_{1,i} = A} + \delta \cdot \mathbf{c}_i + \gamma \cdot \mathbf{s}_i + \epsilon_i$ where $\widehat{r_{1,i} = A}$ is the predicted probability of adviser i recommending option A because i is exposed to the bonus, thus we take $Bonus_i$ and our first-stage results to instrument $r_{1,i}$. The estimate for β in the second stage then reflects the causal effect of the bonus on the probability of choosing option A for oneself. The point estimates range from 0.344 to 0.452, depending on the specification, and are significant

more than a third of those advisers who were put on the spot by biasing their recommendations and then having to choose for themselves behaved consistently by choosing option A for themselves.

We can also take our choice rate for option A in NO BONUS, which is 9.8%, as an estimate of how many people actually prefer it independent of possible image concerns due to the bonus. Adding this to the above estimate, we would expect that a total of 44.2%(=34.4%+9.8%) of the advisers in BONUS who initially recommended option A in R1 behaved consistently and also chose it in O. What we empirically observe is that 42.3% of the advisers in BONUS who initially recommended option A exhibit such a behaviour, a percentage which is not different from the expected one (t-test: p = 0.850). Furthermore, this observed frequency also means that a significant share of advisers in treatment BONUS who have initially recommended option A, 57.7% (t-test: p = 0.000), do *not* choose it for themselves. Again, if advisers were just confused and took the bonus as an indication of what they should recommend, we would expect them all to also act accordingly for themselves.

Second recommendation R2: For their second recommendation, the decision situation for advisers in NO BONUS is the same as for their first. Accordingly, we expect a similar pattern of recommendations. The left panel of figure 4.3 shows the recommendation frequencies for each option. Comparing it to the left panel of

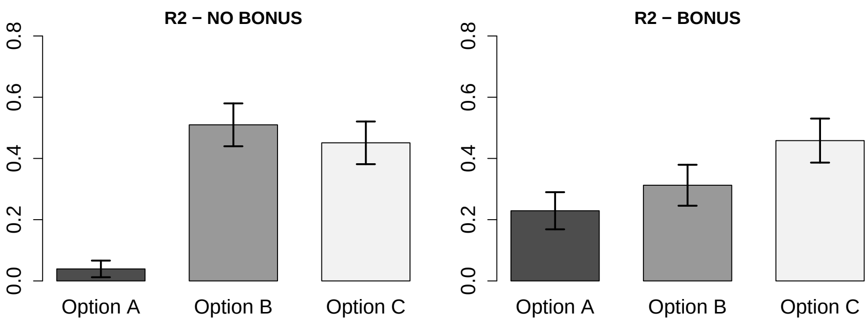

Figure 4.3: Frequency for each option being recommended in R1, bars depict standard errors.

figure 4.1 shows that this is largely the case: 82.4% of the advisers in NO BONUS recommend again exactly the same option they recommend initially. In particular, exactly the small minority of 3.9% of the advisers who recommended option A

($p < 0.05$). Strictly speaking, the results of this two-stage procedure may however be biased since the exclusion restriction for the instrument $Bonus_i$ could be violated (being in the BONUS treatment could influence the own choice via channels other than the first recommendation). Given the fit to our above estimates and observations, we however consider the results of this procedure noteworthy.

recommends it again.

This picture is very different when we compare this to the recommendations in BONUS. Although there is no bonus for recommending option A in R2 either, the rate of recommendation for option A is almost five times as high as in the NO BONUS: 22.9% of those advisers who had previously been exposed to the bonus recommend option A, a significant increase by 19.0 percentage points relative to the NO BONUS (Fisher exact test: $p = 0.007$). This is also confirmed by a regression analysis which re-estimates model (4.1) when a dummy which indicates whether option A is recommended in the second recommendation is the dependent variable. Table 4.4 presents the results and shows that this point estimate even increases. Again,

	(1)	(2)	(3)	(4)
BONUS	0.190***	0.203***	0.211**	0.213**
	(0.067)	(0.067)	(0.092)	(0.087)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	99	99
Adjusted R^2	0.070	0.073	0.038	0.064

Table 4.4: OLS estimates of the probability to recommend option A in R2
robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
personal controls: age, gender, monthly budget, subject's region of origin and field of studies

this pattern is also observed for probit estimates (see table 10 in the appendix). We therefore treat prediction 3 as confirmed.

As above for advisers' own choices O, we can estimate the causal effect of having given in to the bonus on the repeated recommendation for the risk-seeking option. The initial effect of an increase in the probability of recommending A in R1 due to the bonus was estimated by 50.3 percentage points. The observed increase of 19.0 percentage point in R2 then implies that, in expectation, 37.8% ($\triangleq 0.190/0.503$) of advisers recommend option A again just because they have previously given in to the bonus.[13] To estimate the frequency of advisers in BONUS who recommend option A twice we add the 3.9% who do so in the NO BONUS treatment as an estimate for the proportion of those who recommend it for reasons unrelated to the bonus. The implied point estimate from this decomposition is 41.7%(=37.8%+3.9%). This

[13] We also repeated the two-step instrumental-variable-procedure as explained in footnote 12. That is, we estimate the probability of recommending option A again in R2 when one's first recommendation R1 has been biased the bonus. With the same caveat as described there applying here, the resulting IV-estimates of this causal channel range from 0.372 to 0.442 percentage point, depending on the specification, and are significant ($p < 0.01$).

estimate in the region of the actually observed frequency of advisers in BONUS who re-recommend option A in R2: It is given by 34.6% which is not statistically differed from the above estimate (t-test: $p = 0.417$).

Further results: There are some further findings which support our theory and its underlying assumptions. Given our previous results, we expect high consistency between advisers' own choices and their first recommendation when there is no conflict of interest. Our results are largely in line with this: Table 4.5a) shows the frequencies of advisers choosing for themselves, conditional on their first recommendation in NO BONUS. Only the off-diagonal entries are not in line with this prediction. They amount to a total of 17.7% of the observation in this treatment; 82.3% of our observations in NO BONUS are therefore in line with the predicted consistency. In

a) NO BONUS

		O = A	O = B	O = C
R1 =	A	3.9%	0.0%	0.0%
	B	2.0%	23.5%	11.8%
	C	3.9%	0.0%	54.9

b) BONUS

		O = A	O = B	O = C
R1 =	A	22.9%	8.3%	22.9%
	B	0.0%	6.3%	0.0%
	C	4.2%	4.2%	31.3

Table 4.5: Frequencies of advisers' own choices O conditional on their first recommendation R1.

BONUS, our theory predicts that some of those who have previously recommended option A stick to it in order to avoid a negative self-image. Other advisers who have recommended it but who do not have sufficiently strong image concerns choose their preferred option instead. Accordingly, we can explain the diagonal entries in table 4.5b) plus the off-diagonal ones in the first row. Again, this leaves only a small fraction of 8.4% of our observations unexplained.

We find similar results with regards to the consistency between advisers' first and second recommendations. Table 4.6a) and b) show the respective conditional frequencies across our experimental conditions. In NO BONUS, noise is somewhat

a) NO BONUS

		R2 = A	R2 = B	R2 = C
R1 =	A	3.9%	0.0%	0.0%
	B	0.0%	35.3%	2.0%
	C	3.9%	15.7%	43.1%

b) BONUS

		R2 = A	R2 = B	R2 = C
R1 =	A	18.8%	16.7%	18.8%
	B	0.0%	6.3%	0.0%
	C	4.2%	8.3%	27.1%

Table 4.6: Frequencies of advisers' second recommendations R2 conditional on their first R1.

higher than for the previous comparison. We observe a total of 21.6% to be inconsistent, i.e. to be outside table 4.6a)'s diagonal. However, one should note firstly, that these inconsistencies are primarily due to switches from having initially recommended option C and then option B, thus between neighboring, non-risk-seeking options. Secondly, almost eighty percent of recommendations are consistent and thus in line with our theory. With regard to variations in the BONUS treatment the results are even stronger. In total, 87.5% of our observations fall into an explainable pattern, thus are either on the diagonal or the first row. Overall, the consistency predicted by our theory can be observed in at least four fifth of the relevant cases and often, in even higher proportions.

Further evidence comes from our exit questionnaire. It contained a question on advisers' general risk attitudes. More precisely, it asked subjects to indicate on an 11-point Likert-scale "How willing are you to take risk, in general?". Although this question was not incentivized, answers to it has previously been shown to correlate with peoples' actual choices under risk. While in NO BONUS, the average response was 5.0 points, it increased by almost one point or alternatively, 39.8% of its standard deviation, to 5.9 points in the treatment BONUS. This increase is marginally statistically significant (Wilcoxon ranksum-test: $p = 0.059$).[14] This result becomes even stronger, both numerically and statistically, in an OLS regression analysis when additional control variables are included. Table 4.7 represents the results from estimating model (4.1) when the dependent variable is this self-assessed risk-measure and controls are successively added. The results are also robust to estimation via ordered probit (see table 11 in the appendix). This increase in an adviser's self-stated risk measure is consistent with our theory: Advisers who have previously given in to the bonus can self signal that this advice was appropriate from their point of view when they consider themselves as more risk-seeking.[15] Once again, this is also consistent with advisers who are not just confused about their choices and

[14] Due to a data-glitch in the first two sessions, we had to collect the risk-measure along with the other post-experimental questionnaire separately. When we exclude these sessions, the increase is 1.1 points, 46% of the measure's standard deviation, and is similarly significant (Wilcoxon ranksum-test: $p = 0.062.$). The same pattern (higher point estimates and slightly lower but still significant p-values) holds when we exclude these observations from the regressions reported in table 4.7. Note that our primary data on the recommendations R1/R2 and own choices O were not affected by this data glitch since they were collected by advisers writing them on paper.

[15] We also repeated the two step instrumental variable procedure laid out along with its caveats in footnote 12. This allows us to estimate the effect on the risk measure through having recommended option A by instrumenting this choice via an advisers' random exposure to the bonus. The estimated coefficient ranges from a 1.8 to 2.2, depending on the specification and are significant ($p < 0.05$). Given the first stage increase in the probability of recommending option A due to the bonus of 50.3 percentage points, the implied causal increase of 0.9 to 1.1 (1.8×0.503 to 2.2×0.503) is consistent with these estimates.

	(1)	(2)	(3)	(4)
BONUS	0.914**	1.030**	1.244**	1.306**
	(0.453)	(0.436)	(0.576)	(0.590)
Constant	4.961***	3.534***	5.185**	5.819*
	(0.335)	(0.634)	(2.284)	(3.090)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	99	99
Adjusted R^2	0.040	0.199	0.374	0.415

Table 4.7: OLS estimates on the self-assessed preference for risk (Likert scale, 0 to 10)
robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
personal controls: age, gender, monthly budget, subject's region of origin and field of studies

recommendations but who, on the contrary, do even understand the more general behavioural implications of their recommendations outside the given set of options.

4.6 Conclusion

In this paper, we provide experimental evidence that incentives to bias advice have a lasting and causal effect on both, advisers' future recommendations for risky decisions and their own choices. When advisers are paid a bonus to recommend an investment option which is only preferred by risk-seeking individuals, about half of them recommend it. Without such a bonus only four percent do so. Prior exposure to a bonus leads a significant share of advisers to re-recommend this option when there is no bonus anymore and even to choose it for themselves. We provide a psychological mechanism which is capable of explaining these findings. It is based on advisers' desire to not self-signal their corruptibility. This forces them to be consistent in their recommendations and own choices, even when this means to bias further advice and even their own choices. With this theory we can consistently decompose the recommendation and choice pattern of advisers in our experiment. We estimate that around 35 to 40 percent of those advisers whose advice has been corrupted by the bonus engage in such continuing deception of advisers and also of themselves in order to preserve a positive self-image.

A straightforward policy implication of our findings is therefore that removing advisers' conflicts of interest does not necessarily eliminate their effect on advice giving. For example, the Retail Distribution Review (RDR) in the UK whose step-wise implementation started in 2013 bans commission-based financial advice. Our results indicate that, while it may improve such advice in the long run, its full effects may be considerably delayed. Experienced advisers, who have spent their hitherto professional life in an environment which featured such incentives, will likely exhibit persistent biases in their recommendations.

Our proposed mechanism also has profound consequences on how accountable advisers feel. It implies that it is the desire to see oneself as a moral, impartial adviser which can lead to exactly the opposite behaviour. Those who stop giving biased advice after bonuses are removed identify themselves as having previously been corrupted. In contrast, those who continue to give biased advice do so just to avoid this inference and therfore, do not feel corrupted. In consequence, the awareness of acting in a corrupted manner and actually giving biased advice do not coincide, in fact they are asymmetric. This provides challenges for the remedy of the biases resulting from conflicts of interest as those who do the damage might not even feel culpable. Given the demand for advice in many situations, we think that exploring these mental processes by advisers and the adverse consequences it has on their job is a fruitful avenue for further research.

Appendix A – A simple model of self-signaling and corrupted advice-giving

In the following, we formally derive three predictions I to III which are analogous to their respective counterparts, predictions 1 to 3 in then main text. These derivations are based on a formal model presented below with assumptions capturing those described in section 2.

First recommendation R1: We consider an adviser who recommends a client which action out of a discrete set $\mathcal{S}$ to take. In our experiment, these are three investment options A, B, and C, thus $\mathcal{S} = \{A, B, C\}$. We denote an adviser's (first) recommendation by $r_1 \in \mathcal{S}$. In addition, there is a bonus $b(r_1)$ which depends on the issued recommendation. In our experiment, an adviser gets a bonus b if he recommends option A, otherwise he does not get any bonus. We thus have $b(r_1) = b \cdot \mathbb{1}[r_1 = A]$. We denote the utility which advisers get from a pecuniary payoff x by the strictly increasing vNM-utility function $u(x)$.

In addition, an adviser i suffers dis-utility $k_i > 0$ to the extend that he recommends an option which is not in the client's best interest. What constitutes a client's best interest is based on two factors: First, it is the choice c^* which the adviser would make if he had to make the client's decision for himself, thus $c^* = \arg\max_{c \in \mathcal{S}} \mathrm{E}[u(c)]$. Second, we allow the adviser to hold a (motivated) belief about the client's preferences. This is captured by the vNM-utility function $\tilde{u}$ which denotes the adviser's belief about the client's preference. We can then denote the implied optimal choice, based on this first-order belief, by $\tilde{c}^* = \arg\max_{c \in \mathcal{S}} \mathrm{E}[\tilde{u}(c)]$. We let $\gamma \in [0, 1]$ denote the weight which advisers assign to their own preference in determining what it is the client's best interest as opposed to optimal recommendations based on their first-order beliefs about the client's preferences. An adviser's overall utility of recommending r_1 is then given by the following expression:

$$\tilde{V}(r_1) = u(b \cdot \mathbb{1}[r_1 = A]) - k_i\big(\, \gamma \cdot \mathbb{1}[r_1 \neq c^*] + (1 - \gamma) \cdot \mathbb{1}[r_1 \neq \tilde{c}^*]\big) \quad (2)$$

This allows several interpretations: When $\gamma = 1$, the question of what constitutes appropriate, morally sound advice is the same as "What would I choose if I were in the client's position?". Conversely, $\gamma = 0$ means that only what an adviser beliefs about others' preferences, not his personal consideration, is relevant for issuing appropriate advice. Values of γ within the unit interval can represent situations in between or when an adviser believes that a client has utility represented by u with probability γ and otherwise represented by $\tilde{u}$. The magnitude of k_i then scales concerns about issuing unsuited advice relative to pecuniary payoffs.

Advisers can form a belief about the client's preferences in a self-serving manner. That is, whenever they issue a recommendation r_1 they can maximize their overall utility by self-servingly believe that the clients' preferences $\tilde{u}$ are such that $\tilde{c}^* = r_1$. In this regard, γ can also be interpreted as how far such a self-serving belief can be formed, independently of and adviser's own preferences. Therefore, the recommendation r_1 which maximizes (2) is the maximizer of the following, more simple, expression:

$$v(r_1) = u(b \cdot \mathbb{1}[r_1 = a)]) - \gamma k_i \cdot \mathbb{1}[r_1 \neq c^*] \tag{3}$$

We let K_{c^*} denote the cdf of the distribution of an adviser's moral cost k_i, conditional on this adviser preferring option c^*, e.g. $K_A(x) = \Pr[k_i \leq x | c^* = A]$. For simplicity, we assume that each of these conditionals cdf's has pdf which is strictly positive over its support.[16] We also let $\alpha_{c^*} > 0$ denote the share in the population of advisers who have preferred action $c^* \in \mathcal{S}$.[17] For easier notification, we let $\alpha = \alpha_A$, i.e. in our experiment α is the share of advisers who are sufficiently risk-seeking to choose option A. We assume the above distributions and parameters to be common knowledge.[18]

R1 – NO BONUS: Since there is no incentive to bias advice, only the second part of (3) matters. This is maximized by $r_1^* = c^*$. In consequence, the share of advisers

[16]Results do not change when the cdfs are allowed to be partially non-increasing, as long as at least one of the pdfs has some mass on sufficiently low values, i.e. that $K_{c^*}(\min\{u(b) - u(0), \mathrm{E}[u(c^*) - u(A)]\}) > 0$ for at least one $c^* \in \mathcal{S} \setminus \{A\}$.

[17]In consequence, the unconditional cdf $\Pr[k_i \leq x]$ is given by $\sum_{c^* \in \mathcal{S}} \alpha_{c^*} K_{c^*}(x)$.

[18]Note that when the signaling concern refers to a dual-self model where advisers ex-post infer their own type from actions, this common prior only refers to these selves. A common prior between individuals is not required.

who recommend option A equals α.

R1 – BONUS: For those who have $c^* = A$, it follows from (3) that they should also recommend it. For those with $c^* \neq A$, they can either recommend option A nevertheless to earn the bonus or they recommend their preferred option $c^* \neq A$ and obtain a utility of $u(0)$. Advisers who do not prefer option A then recommend it if and only if $\gamma k_i < u(b) - u(0)$. By using the convention that $K_{c^*}\left(\frac{u(b)-u(0)}{\gamma}\right)|_{\gamma=0} = \lim_{x\to+\infty} K_{c^*}(x) = 1$ we can then define the following

$$\beta \equiv \sum_{c^*\in\mathcal{S}\setminus\{A\}} \alpha_{c^*} K_{c^*}\left(\frac{u(b)-u(0)}{\gamma}\right) = \alpha_B K_B\left(\frac{u(b)-u(0)}{\gamma}\right) + \alpha_c K_c\left(\frac{u(b)-u(0)}{\gamma}\right) > 0$$

Thus with a bonus, a share β of advisers is corrupted by the bonus and recommends option A, in addition to the share α who would have recommended this option anyhow.

Given the same expected population of advisers across BONUS and NO BONUS, as achieved by random treatment assignment, we can then state the following:

Prediction I. $\Pr[r_1 = A \mid \text{bonus}\,] = \alpha + \beta > \Pr[r_1 = A \mid \text{no bonus}\,] = \alpha$

It will be helpful to categorize advisers along three behavioural types $\theta \in \{1, 2, 3\}$. These types reflect the motives underlying their recommendation r_1 as follows:

Type 1 ($\theta = 1$): Advisers who have $c^* = A$ and recommend $r_1 = c^*$, share α.

Type 2 ($\theta = 2$): Advisers who have $c^* \neq A$ but recommend $r_1 \neq c^*$, share β.

Type 3 ($\theta = 3$): Advisers who have $c^* \neq A$ and recommend $r_1 = c^*$, share $1 - \alpha - \beta$.

Type 1 and 3 advisers give the same advice they would have given had the bonus been absent. Type-2-advisers are corrupted: They recommend option A not because they prefer it but because they were paid to do so. Note that this above categorization of types also applies in the NO BONUS-treatment, the respective shares however differ: Share α also recommends option A without a bonus. Type-2-advisers do not exist in this treatment thus we can treat β as if it were equal to zero and the share of type-3-advisers is given by $1 - \alpha$.

Own choice O: The extent to which advisers take their own choice as a "diagnosis" of the moral type in R1 is given by $\lambda \geq 0$. A value $\lambda \in (0,1)$ would reflect that choosing for one-self is not exactly the same as recommending to others but also that is not unrelated; $\lambda = \gamma$ is then a natural case. In general, we assume that λ is some increasing function Λ of γ with $\lambda = \Lambda(\gamma) = 0$ if and only if $\gamma = 0$. This means that own choices only have diagnostic value to assess an adviser's previous recommendation when his own preference is, at least partly, relevant for issuing appropriate advice.

When λ is positive, an adviser's own choice $c \in S$ signals his underlying motives for his previous recommendation in R1. In particular, an adviser can potentially infer that he was a type-2-adviser according to the above classification. The cost of inferring that one is such a type, thus that one-seld is corruptible yield image dis-utility $l_i > 0$. By denoting the expected utility from choosing a lottery $c \in \mathcal{S}$ by $\mathrm{E}[u(c)]$, the overall utility of advisers is then given by

$$V(c|r_1) = \mathrm{E}[\mathrm{u(c)}] - \lambda l_i \cdot \Pr[\theta = 2|r_1, c] \tag{4}$$

As before, we assume that l_i can be described by a family of commonly known conditional cdfs $(L_{c^*})_{c^* \in \mathcal{S}}$, e.g. $L_A(x) = \Pr[l_i \leq x | c^* = A]$.[19]

O – NO BONUS: When there was no prior bonus, there are no type-2 advisers. In consequence, $\Pr[\theta = 2|r_1, c] \leq \Pr[\theta = 2] = 0$ holds and $c = c^*$ maximizes (4) via $\mathrm{E}[u(c)]$. The share of advisers choosing option A for themselves is thus given by α.

O – BONUS: We start with the case that $\lambda > 0$. First note that type-3-advisers who have previously recommended $r_1 \neq A$ cannot infer to be type-2-advisers, i.e. $\Pr[\theta = 2|r_1 \neq A, c] = 0$. All type-3-advisers therefore choose $c = r_1 = c^* \neq A$ to maximize (4). Type-1 and type-2 advisers can however both infer to be type-2 and would then suffer dis-utility l_i because they have the same initial recommendation $r_1 = A$. Denote the likelihood that a type-1-adviser chooses $c = A$ with $\tau_c = \Pr[c = A|\theta = 1]$

[19] This effectively constitutes a intrapersonal signaling game where an adviser of type (k_i, l_i) sends a message $(c|r_1)$ and then gets dis-utility when he infers from this that his type is such the he behaves according to the behavioural type $\theta = 2$.

and that a type-2-adviser makes the same choice with $\pi_c = \Pr[c = A|\theta = 2]$. One then gets the following for the corresponding posteriors:

$$\Pr[\theta = 2|c = A, r_1 = A] = \frac{\pi_c \cdot \beta}{\tau_c \cdot \alpha + \pi_c \cdot \beta} \tag{5}$$

$$\Pr[\theta = 2|c \neq A, r_1 = A] = \frac{(1 - \pi_c) \cdot \beta}{(1 - \tau_c) \cdot \alpha + (1 - \pi_c) \cdot \beta} \tag{6}$$

It is easily verified that $\Pr[\theta = 2|c \neq A, r_1 = A] \geq \Pr[\theta = 2|c = A, r_1 = A]$ whenever $\tau_c \geq \pi_c$. If this condition holds, type-1-advisers who choose $c \neq A$ suffer for two reasons: First, they loose expected pecuniary utility by choosing a suboptimal choice $c = A \neq c^*$. Second, they expect dis-utility from damage to self-image which is at least as big as when they had chosen their preferred option. In consequence, there is only one equilibrium with $\tau_c \geq \pi_c$ in which $\tau_c = 1$ and all type-1-advisers are consistent by choosing $r_1 = c = A$. While other equilibria with $\tau_c < \pi_c$ cannot be excluded but also do not need to exist, the one with $\tau_c = 1$ is a natural candidate: In it, type-1-advisers who are not corrupted by the bonus do also not deviate from their preferred choice just because of the fear of perceiving themselves as corruptible type-2-advisers while type-2-adviser, who want to uphold a positive self-image, might do so. Also, while there is always the equilibrium with $\tau_c = 1$, those with $\tau_c < \pi_c$ may not even exist.[20]

Type 2-advisers then face a trade-off: They would not like to choose option A for themselves, since for them $c^* \neq A$holds. However, if they switch from their first recommendation to their preferred option, they then generate a perfect signal of being type-2 since all other types are consistent by choosing $c = r_1$ and therefore, $\Pr[\theta = 2|c \neq A, r_1 = A] = 1$ holds. Using the posterior (5), a type-2-adviser therefore chooses his preferred option $c^* \neq A$ if and only if

$$\mathrm{E}[u(c^*)] - \lambda l_i > \mathrm{E}[u(A)] - \lambda l_i \cdot \frac{\pi_c \cdot \beta}{\alpha + \pi_c \cdot \beta}$$

That is, an adviser reveals himself when his image concern is sufficiently low, i.e.

[20] If the dis-utility of not choosing option A although one prefers it is too large, type-1 would not choose another option just to appear less as a type 2.

when $l_i < \frac{\alpha+\pi_c\beta}{\lambda\alpha}(\mathrm{E}[u(c^*)] - \mathrm{E}[u(A)])$. For this, they have to take into account that by not choosing option A, they decrease π_c. This in turn simplifies pooling and thereby raises the opportunity cost of such a choice. It follows that, in equilibrium, the share of type-2-advisers who choose option A to uphold a positive image balance this effect. This share is therefore given by the solution to the following expression:

$$\begin{aligned}1-\pi_c &= \sum_{c^*\in\mathcal{S}\setminus\{A\}} \alpha_{c^*} L_{c^*}\left(\frac{\alpha+\pi_c\beta}{\lambda\alpha}(\mathrm{E}[u(c^*)]-\mathrm{E}[u(A)])\right)\\ &= \alpha_B L_B\left(\frac{\alpha+\pi_c\beta}{\lambda\alpha}(\mathrm{E}[u(B)]-\mathrm{E}[u(A)])\right)+\alpha_C L_C\left(\frac{\alpha+\pi_c\beta}{\lambda\alpha}(\mathrm{E}[u(C)]-\mathrm{E}[u(A)])\right)\end{aligned} \tag{7}$$

Note that for all values of $\pi_c \in [0,1]$, the above RHS is strictly positive and non-decreasing in π_c. Also note that from $\alpha = \alpha_A > 0$ it holds that $\sum_{c^*\in\mathcal{S}\setminus\{A\}} \alpha_{c^*} L_{c^*}(x) < \sum_{c^*\in\mathcal{S}} \alpha_{c^*} L_{c^*}(x) \leq 1$ for every $x \in \mathbb{R}_{++}$.[21] The above RHS is therefore strictly less than one. Since the LHS is strictly decreasing in π_C and takes all values in $[0,1]$ over that interval, there is a unique solution $\pi_c^* \in (0,1)$ to (12). Also note that since the RHS of (12) is decreasing in λ, the implied consistency in own choice π_c^* is also strictly increasing in this parameter.

Now consider $\lambda = 0$: The second part in (4) does not count then and irrespective of their prior behaviour, all advisers choose c^*. This is equivalent to $\pi_c^* = 0$.

In summary, share α of type-1-advisers initially recommend and then choose for themselves option A. Type-3-advisers initially recommend and then choose their preferred non-A option. Type-2-advisers, whose total share is given by β, split in two subgroups: Advisers in the first subgroup who represent share $\pi_c^*\beta$ of all advisers choose option A to uphold a positive image. Advisers in the second subgroup with population share $(1-\pi_c^*)\beta$ put their own payoff above image concerns and choose their preferred non-A options. The first sub-group then has mass only when they advisers assign diagnostic value to their choices, thus if $\gamma > 0$. Assuming that this is true, the following predictions can then be stated:

[21] This also holds under the condition for weakly-increasing cdfs laid out in footnote 16 (it is the reason for the second expression in the *min*-term).

Prediction II. Suppose $\lambda > 0$. Then $\Pr[c = A|\text{bonus}] = \alpha + \pi_c^*\beta > \Pr[c = A|\text{no bonus}] = \alpha$

Second recommendation R2: As before, the dis-utility of inferring to be corruptible, thus to be a type-2-adviser, is given by $l_i > 0$. Since the advice in R2 is the same as in R1 we do not discount the diagnostic value by some $\lambda < 1$. The recommendation does not affect the adviser himself but the client. We thus assume, as for the first recommendation, that he suffers dis-utility from giving inappropriate advice, measured by k_i. Note that advisers initially formed a self-serving belief about $\tilde{u}$. In consequence, they have to stick to it. This means that there is additional dis-utility $k_i(1-\gamma)$, of not living up to one's prior motivated belief to $\tilde{c}^* = r_1$. An adviser's ex-ante utility from giving recommendation r_2, given his prior actions and beliefs, is then described by

$$V(r_2|r_2, r_1, c) = -k_i\left(\gamma \cdot \mathbb{1}[r_2 \neq c^*] + (1-\gamma) \cdot \mathbb{1}[r_2 \neq r_1]\right) - l_i \cdot \Pr[\theta = 2|r_2, c, r_1] \quad (8)$$

R2 – NO BONUS: Again, without a previous bonus type-2-advisers do not exist and $\Pr[\theta = 2|r_1, c] \leq \Pr[\theta = 2] = 0$. Since $r_1 = c = c^*$ was chosen initally, recommending $r_2 = c^*$ then maximizes (8). The share of advisers recommending option A (again) in NO BONUS is therefore α.

R2 – BONUS: For type-3-advisers, their previous behaviour with $c = r_1 = c^* \neq A$ prevents them from inferring to be type-2-advisers since $\Pr[\theta = 2|r_2, c = r_1 \neq A] \leq \Pr[\theta = 2|c = r_1 \neq A] = 0$. Since for them $c = r_1 = c^*$ holds, they maximize (8) by recommending $r_2 = c = r_1 = c^* \neq A$.

First, consider the case that own actions in O had diagnostic value, thus $\lambda > 0$ and therefore $\pi_c^* \in (0, 1)$. Share $1 - \pi_c^*$ of type-2-advisers has then already revealed himself as such. For them, $\Pr[\theta = 2|r_2, c \neq r_1 = A] = \Pr[\theta = 2|c \neq r_1 = A] = 1$ applies. Their second recommendation r_2 is thus unaffected by image concerns. Accordingly, $r_2 = c^* \neq A$ maximizes (8) when $\gamma > \frac{1}{2}$ and $r_2 = r_1 = A$ when it holds that $\gamma \in (0, \frac{1}{2}]$.[22]

[22] Note that when $\gamma > \frac{1}{2}$, their prior, self-serving belief leads even those advisers who have

It follows that the mass of candidates for continued pooling with type-1-advisers in R2 is given by the overall share $\pi_c^*\beta > 0$ of advisers who has not yet revealed themselves to be type-2. They, together with type-1-advisers have a history of $c = r_1 = A$. By denoting the likelihood that a type-1-adviser chooses $r_2 = A$ with $\tau_{r_2} = \Pr[r_2 = A|\theta = 1]$ and the corresponding probability for a type-2-adviser who has not revealed himself by $\pi_{r_2} = \Pr[r_2 = A|\theta = 2, c = r_1 = A]$ we get the following posteriors:

$$\Pr[\theta = 2|r_2 = c = r_1 = A] = \frac{\pi_{r_2} \cdot \pi_c^*\beta}{\tau_{r_2} \cdot \alpha + \pi_{r_2} \cdot \pi_c^*\beta} \tag{9}$$

$$\Pr[t = 2|r_2 \neq A, c = r_1 = A] = \frac{(1 - \pi_{r_2}) \cdot \pi_c^*\beta}{(1 - \tau_{r_2}) \cdot \alpha + (1 - \pi_{r_2}) \cdot \pi_c^*\beta} \tag{10}$$

Analogously to the comparison of (5) and (6), (10) is larger than (9) whenever $\tau_{r_2} \geq \pi_{r_2}$. Repeating the analogous reasoning for an equilibrium with $\tau_c = 1$ in O, there is an equilibrium where all type-1-advisers choose option A for their second recommendation, thus with $\tau_{r_2} = 1$. This is the equilibrium on which we focus (see the above discussion on this selection for O, the same arguments carry over to R2).

Type 2-advisers who have so far not revealed themselves through inconsistent actions (i.e. $c \neq r_1 = A$) face again a trade-off: On the one hand, they could recommend their preferred choice $r_2 = c^* \neq A$ to prevent the cost γk_i of giving inappropriate advice, based on their personally preferred action. However, this would then reveal them to be type-2s and get them dis-utility l_i. In addition, they would give inappropriate advice based on their self-serving belief $\tilde{c}^* = A = r_1$ they formed in R1 which would now create costs of $(1 - \gamma)k_i$ when they recommend $r_2 \neq r_1$. The alternative is to continue in recommending option A to pool with type-1-advisers and therefore uphold a positive self-image. By using (9), together with $\tau_{r_2} = 1$, a type-2-adviser then recommends $r_2 = c = r_1 = A \neq c^*$ if and only if

$$-k_i\gamma - l_i \cdot \frac{\pi_{r_2} \cdot \pi_c^*\beta}{\alpha + \pi_{r_2} \cdot \pi_c^*\beta} > -k_i(1 - \gamma) - l_i \;\Leftrightarrow\; \frac{k_i}{l_i}(2\gamma - 1) < \frac{\alpha}{\alpha + \pi_{r_2}} \cdot \pi_c^* \tag{11}$$

In consequence, a type-2-adviser who re-issues biased advice by recommending $r_2 =$

already revealed themselves to re-issue their biased advice for option A.

A has low concerns of giving inappropriate advice (k_i) relative to their image concern (l_i). To formalize this, it will be useful to denote the family of cdfs of the ratio distribution k_i/l_i, conditional on an adviser's preferred option c^*, by $(R_{c^*})_{c^* \in \mathcal{S}}$. For example, a typical member is $R_B = \Pr[k_i/l_i \leq x | c^* = B]$.[23]

First consider the case that $\gamma > \frac{1}{2}$. Again, revealing one-self by recommending a non-A option increases the opportunity cost of doing so as pooling becomes easier. In equilibrium, advisers take this into account. From (11), it then follows that the share $\pi^*_{r_2}$ of hitherto not revealed type-2-advisers who continue to pool with type-1s has to solve the following expression:

$$\begin{aligned} \pi_{r_2} &= \sum_{c^* \in \mathcal{S} \setminus \{A\}} \alpha_{c^*} R_{c^*} \left(\frac{\alpha}{(2\gamma - 1)(\alpha + \pi_{r_2} \cdot \pi^*_c \beta)} \right) \\ &= \alpha_B R_B \left(\frac{\alpha}{(2\gamma - 1)(\alpha + \pi_{r_2} \cdot \pi^*_c \beta)} \right) + \alpha_C R_C \left(\frac{\alpha}{(2\gamma - 1)(\alpha + \pi_{r_2} \cdot \pi^*_c \beta)} \right) \end{aligned} \tag{12}$$

By analogous reasoning as for the RHS of (12), the above RHS is strictly less than one. It is also non-increasing in π_{r_2}. Therefore, there has to be a unique intersection $\pi^*_{r_2} \in (0, 1)$ with the 45-degree line over the unit interval. We then get the following:

Prediction III.a) $\Pr[r_2 = A|\text{bonus}] = \alpha + \pi^*_{r_2}\pi^*_c\beta > \Pr[c = A|\text{no bonus}] = \alpha$ when $\gamma \in (\frac{1}{2}, 1]$.

Alternatively, if $\gamma \in (0, \frac{1}{2}]$ the second inequality in (11) is always fulfilled since its RHS is strictly positive while the LHS is strictly negative. It then follows that $\pi^*_{r_2} = 1$ and all of the unrevealed type-2s choose $r_2 = A$. In addition, the share $1 - \pi^*_c$ who have previously revealed themselves also choose $r_2 = A$ (see above). This prediction then follows:

Prediction III.b) $\Pr[r_2 = A|\text{bonus}] = \alpha + \beta > \Pr[c = A|\text{no bonus}] = \alpha$ when $\gamma \in (0, \frac{1}{2}]$.

Lastly, consider $\gamma = 0$. Own choices then have no diagnostic value as $\lambda = 0$. The main difference to the preceding analysis is that not choosing $c = r_1 = A$ for

[23]Since k_i and l_i are positively-valued and their distributions are commonly known, R_{c^*} is defined and also commonly known.

type-2-advisers does not necessarily reveal them to be of this type. In consequence, there is no mass $\pi_c\beta$ of candidates for *continued* pooling but *all* type-2-adviser are candidates for pooling with the moral type-1-advisers in R2 and none has previously revealed. The mass of those who potentially mimic type-1-advisers is thus given by β. The analogs to the inference posteriors (9) and (10) are equivalent to setting $\pi_c = 1$ in these expression.[24] Also, they are independent of the adviser's previous choice c since it does not have diagnostic value because $\lambda = \gamma = 0$ applies. Expression (11) then becomes

$$- l_i \cdot \frac{\pi_{r_2} \cdot \beta}{\alpha + \pi_{r_2} \cdot \beta} > -k_i - l_i \quad \Leftrightarrow \quad \frac{k_i}{l_i} > -\frac{\alpha}{\alpha + \pi_{r_2}} \tag{13}$$

and is always fulfilled, thus all type-2-advisers re-recommend $r_2 = A$:

Prediction III.c) $\Pr[r_2 = A|\text{bonus}] = \alpha + \beta > \Pr[c = A|\text{no bonus}] = \alpha$ when $\gamma = 0$.

From predictions III.a) through III.c) we get that for any weight $\gamma \in [0, 1]$, option A is more often re-recommended in BONUS than in NO BONUS, thus prediction 3 in the main text.

[24]Note that in slight contradiction to the initial definition of π_c as the share of type-2-advisers which behaves consistently in the own choice, setting this value equal to one does not mean that all behave consistently. It is however mathematically equivalent to this situation since the choice c has no diagnostic value. This is the same as if all type-2-advisers would have pooled with type-1-advisers. In both cases, the mass for (continued) pooling is the same and given by β.

Appendix B – Further data and analysis

	(1)	(2)	(3)	(4)
BONUS	0.440***	0.420***	0.441***	0.467***
	(0.047)	(0.041)	(0.059)	(0.060)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	89$^{\diamond}$	89$^{\diamond}$

Table 8: Average marginal effect of probit estimates for recommending option A in R1

	(1)	(2)	(3)	(4)
BONUS	0.170**	0.175**	0.210**	0.217***
	(0.074)	(0.071)	(0.090)	(0.071)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	99	79$^{\diamond}$	79$^{\diamond}$

Table 9: Average marginal effect of probit estimates for choosing option A for oneself in O

For the above tables:
Robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.
Personal controls: age, gender, monthly budget, subject's region of origin and field of studies.
Observations with $^{\diamond}$: some combinations of the control variables predicted outcomes perfectly which is why the respective observations are not used in the ML-estimation.

	(1)	(2)	(3)	(4)
BONUS	0.194***	0.251***	0.182**	0.284***
	(0.071)	(0.072)	(0.087)	(0.107)
Personal Controls	no	yes	no	yes
Session Controls	no	no	yes	yes
Observations	99	87$^\diamond$	81$^\diamond$	66$^\diamond$

Table 10: Average marginal effect of probit estimates for recommending option A in R2

	(1)	(2)	(3)	(4)
BONUS	0.409**	0.508**	0.584**	0.624***
	(0.206)	(0.212)	(0.230)	(0.234)
Personal controls	no	yes	no	yes
Session controls	no	no	yes	yes
Observations	99	99	99	99

Table 11: Ordered probit estimates on the self-assessed preference for risk (Likert scale, 0 to 10)

For the above tables:

Robust standard error in parentheses, significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Personal controls: age, gender, monthly budget, subject's region of origin and field of studies.

Observations with $^\diamond$: some combinations of the control variables predicted outcomes perfectly which is why the respective observations are not used in the ML-estimation.

	NO BONUS		BONUS		OVERALL		rank-sum/χ^2-test
	mean	s.d.	mean	s.d.	mean	s.d.	p-value
age	24.824	8.002	23.208	5.411	24.040	6.882	0.264
male	0.451	0.070	0.354	0.070	0.404	0.050	0.339
region of origin							0.194
UK or Ireland	0.196	0.401	0.063	0.244	0.131	0.034	-
other Europe	0.137	0.348	0.188	0.394	0.162	0.370	-
N. America/Australia/New Zealand	0.020	0.140	0.083	0.279	0.051	0.220	-
South America	0.039	0.196	0.021	0.144	0.030	0.172	-
Asia	0.608	0.493	0.645	0.483	0.626	0.486	-
other	0.000	0.000	0.000	0.000	0.000	0.000	-
degree							0.220
bachelor	0.607	0.493	0.500	0.505	0.555	0.050	-
master	0.353	0.483	0.479	0.504	0.414	0.050	-
phd	0.000	0.000	0.000	0.000	0.000	0.000	-
other postgraduate	0.000	0.000	0.020	0.144	0.101	0.100	-
none	0.039	0.196	0.000	0.000	0.020	0.014	-
subject							0.261
economics/business/finance	0.216	0.415	0.375	0.489	0.293	0.457	-
other social sciences	0.353	0.483	0.229	0.425	0.293	0.458	-
psychology	0.059	0.237	0.021	0.144	0.040	0.198	-
public administration	0.039	0.196	0.062	0.244	0.051	0.220	-
math/sciences/engineering	0.157	0.367	0.083	0.279	0.121	0.328	-
arts or humanities	0.157	0.367	0.146	0.357	0.152	0.360	-
other	0.020	0.140	0.083	0.279	0.051	0.220	-
monthly budget (in GBP)	606.275	450.719	640.00	563.775	622.626	506.328	0.964
number of observations	51		48		99		

Table 12: Summary statistics for advisers' personal characteristics and dummy variable based on categorical data. The rightmost column provides p-values for a randomization check between NO BONUS and BONUS (Wilcoxon rank-sum tests for the variables age and budget; χ^2-tests for the remaining categorical variables).

Appendix C – Experimental instructions

The following pages contain screenshots of instructions shown to subjects in ztree and on the information about the investment options printed on paper. They are presented in the order as they were seen by the subjects in the experiment.

- Screen 1: Welcome stage and general instructions
- Screens 2a and 2b: Explanation for R1. Two screens which explain the client's choice situation, the adviser's role, and the investment options.
- Information on the investment options shown to advisers, printed on paper
- Screen 2c: Instructions for giving the first recommendation R1
- Screen 3: Instructions for making the own choice O
- Screen 4: Instructions for giving the second recommendation R2
- Screen 5: Exit questionnaire

The screens show the information shown to advisers in treatment BONUS. The parts which are not shown to advisers in NO BONUS are put in square brackets.

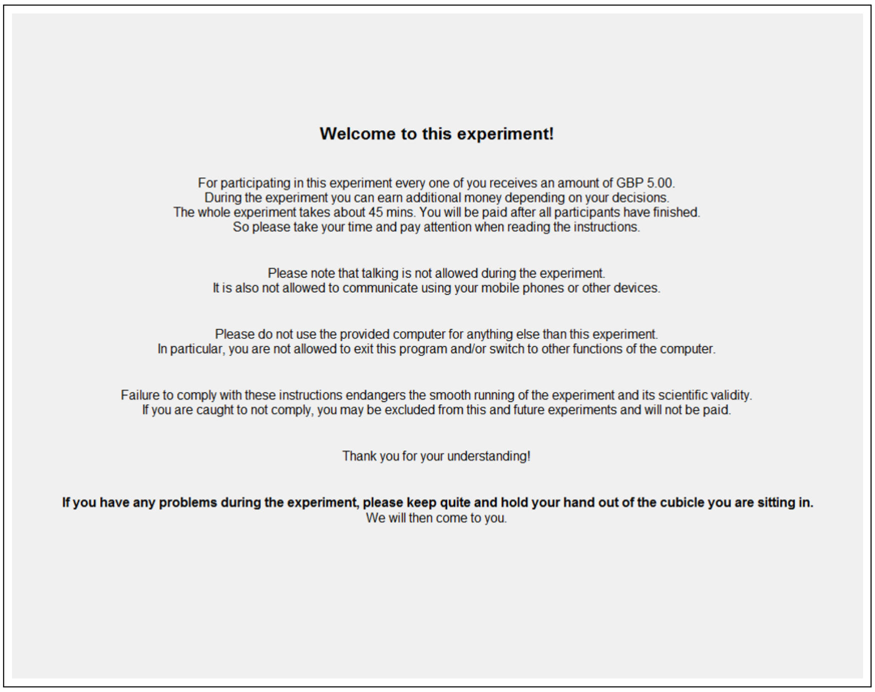

Welcome to this experiment!

For participating in this experiment every one of you receives an amount of GBP 5.00.
During the experiment you can earn additional money depending on your decisions.
The whole experiment takes about 45 mins. You will be paid after all participants have finished.
So please take your time and pay attention when reading the instructions.

Please note that talking is not allowed during the experiment.
It is also not allowed to communicate using your mobile phones or other devices.

Please do not use the provided computer for anything else than this experiment.
In particular, you are not allowed to exit this program and/or switch to other functions of the computer.

Failure to comply with these instructions endangers the smooth running of the experiment and its scientific validity.
If you are caught to not comply, you may be excluded from this and future experiments and will not be paid.

Thank you for your understanding!

If you have any problems during the experiment, please keep quite and hold your hand out of the cubicle you are sitting in.
We will then come to you.

Screen 1

General Information

Your role:

All subjects in the current experimental session are assigned the role of an **advisor**.
As an advisor, you will give a recommendation to a client.
These clients will be subjects in another experiment at the LSE's Behavioral Research Lab.

How it works

In this future experiment with clients, each of them has to choose one out of three options, A, B or C.
Here is what will be shown to the client:
"Each option will earn different monetary payoffs.
Option A presents a possibility to earn a high or a low payoff, depending on luck.
Option B adds the possibility to earn some amount between the high and low payoff, option C increases that possibility."

Clients however do NOT know more about this situation than the above text when they choose an option.
You, as an advisor, will soon learn what exactly these options are.
Afterwards, you have to recommend one option to a client.

Verification

You will have to write down your recommendation on paper and put it into an envelope. If you want, you can address the envelope to yourself.
At the end of this experiment, we will randomly choose one of the recommendations given here to be shown to a client.
If your recommendation is chosen to be shown to a client the following happens:

- We will read out loudly your cubical number (not the name) of that recommendation. You therefore know that you have been chosen.
- We will ask the client who will receive your recommendation to sign it.
- If you wrote your address on the envelope, we will mail you a copy of your recommendation signed by the client
- We will also mail you information of how you can retrive the receipt signed by that client from the lab's official record depository.
- The client will only see your written recommendation, not the envelope which potentially bears your name.

With this procedure, you can verify whether a client has actually gotten your advice, should your recommendation be drawn and you self-addressed the envelope.

[**Your bonus**

You receive a bonus of GBP 3.00 for recommending **Option A**
The bonus will be paid independently of whether your recommendation is chosen to be shown to a client.]

I understood. Please proceed.

You will now learn precisely how a chosen option affects a client's payoffs in addition to the GBP 5.00 they get (as you will) for coming here.

A risky choice

You have to choose one out of the following three options to recommend to a client.
This will determine the client's payoff as follows:

Option A

- Client rolls a six-sided die;
- For any number of the die: client flips a coin and earns GBP 20.00 when the coin shows "Heads"; or nothing when the coin shows "Tails".

Option B

- Client rolls a six-sided die;
- Die shows 1 or 2: client earns an amount of GBP 12.00;
- Die shows 3, 4, 5 or 6: client flips a coin and earns GBP 20.00 when the coin shows "Heads"; or nothing when the coin shows "Tails".

Option C

- Client rolls a six-sided die;
- Die shows 1 or 2: client earns an amount of GBP 12.00;
- Die shows 3 or 4: client earns an amount of GBP 8.00;
- Die shows 5 or 6: client flips a coin and earns GBP 20.00 when the coin shows "Heads"; or nothing when the coin shows "Tails".

[Note: Your bonus of GBP 3.00 which you get for recommending option A is independent of a client's choice.]

Please look now at the paper instructions. It contains a summary of the above and a table which lists all possible outcomes.

Please study the table and examples carefully.
You will soon have to make a recommendation to the client. As said, the client knows nothing of the above.
If you are ready click "Continue" below.

Continue.

Screens 2a (top) and 2b (bottom)

A risky choice

One of the following options must be chosen. Then the following happens:

Option A:

- Roll die: for every outcome, play the lottery.

Option B:

- Roll die: if it shows 1 or 2, one earns GBP 12.00 for sure;
- Roll die: if it shows 3, 4, 5 or 6, one has to play the lottery

Option C: receive a chance to roll the same six-sided die:

- Roll die: if it shows 1 or 2, one earns GBP 12.00 for sure;
- Roll die: if it shows 3 or 4, one earns GBP 8.00 for sure;
- Roll die: if it shows 5 or 6, one has to play the lottery

The lottery:

For the lottery one has to toss a coin. “Heads” then yields GBP 20.00, “Tails” nothing.

Each row of the table below represents a possible result of the die. The columns describe the possible consequences, depending on the chosen option.

Die equal to....	**Option A** is chosen	**Option B** is chosen	**Option C** is chosen
1 or 2	lottery: GBP 20 or 0	GBP 12	GBP 12
3 or 4	lottery: GBP 20 or 0	lottery: GBP 20 or 0	GBP 8
5 or 6	lottery: GBP 20 or 0	lottery: GBP 20 or 0	lottery: GBP 20 or 0

Example:

Suppose the die yielded 3: If option A or B was chosen before, one has to play the lottery. If option C was chosen, one would have gotten GBP 8.00 for sure instead.

Suppose the die yielded 1. If option B or C was chosen before, one gets GBP 12.00 for sure. If option A was chosen, one plays the lottery instead.

Suppose the die yielded 6. Independently of the chosen option one plays the lottery.

Information sheet shown to advisers
(It was placed face down on each adviser's table with the following print on its back: "Information – do not turn until explicitly told so".)

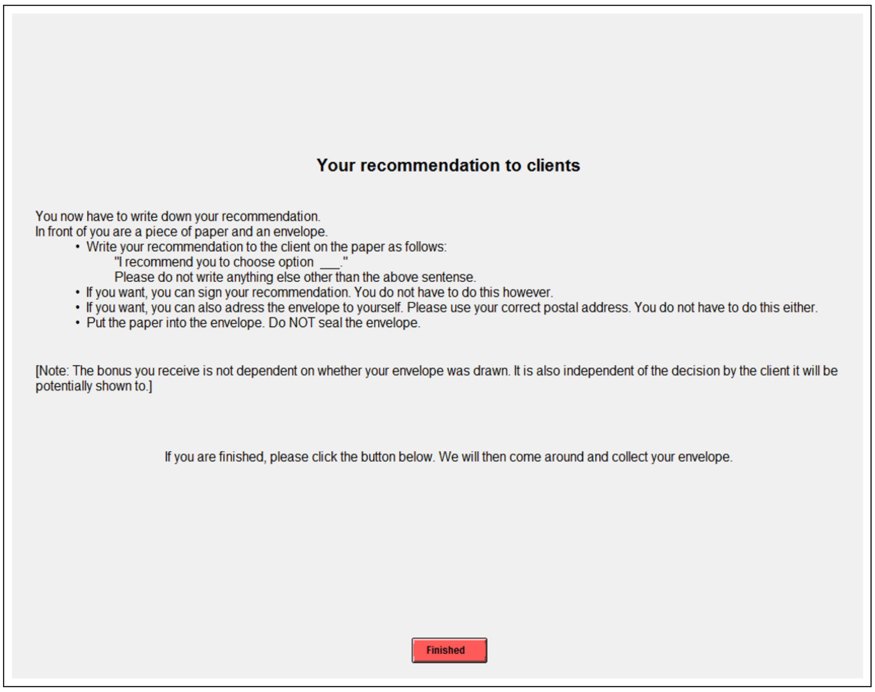

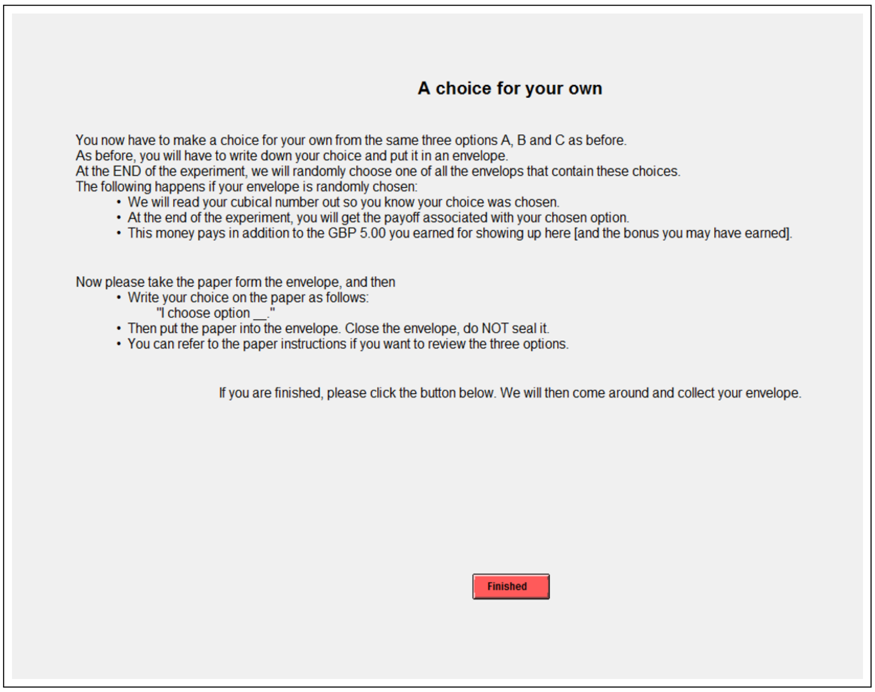

Screens 2c (top) and 3 (bottom)

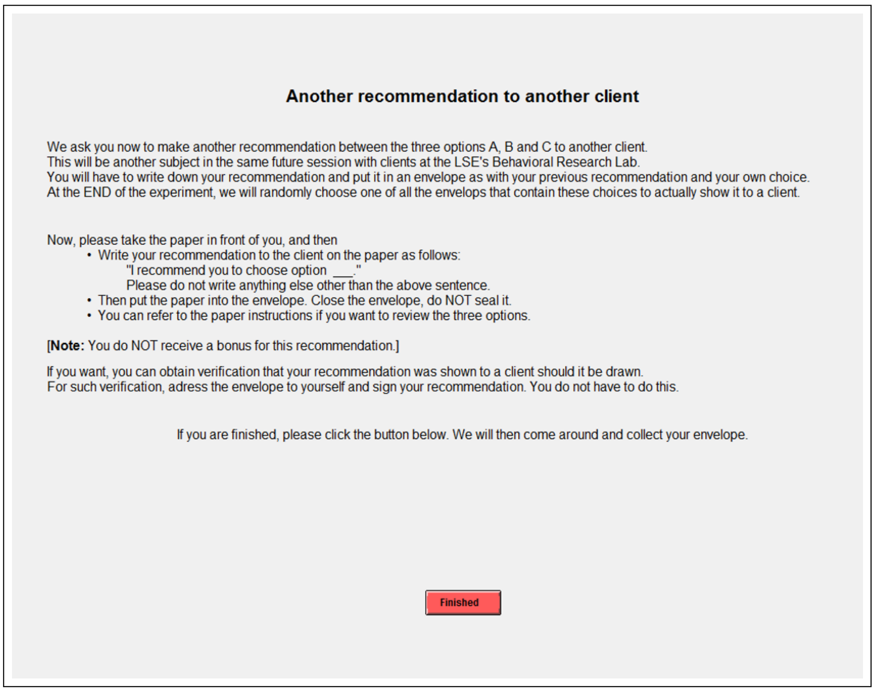

Another recommendation to another client

We ask you now to make another recommendation between the three options A, B and C to another client.
This will be another subject in the same future session with clients at the LSE's Behavioral Research Lab.
You will have to write down your recommendation and put it in an envelope as with your previous recommendation and your own choice.
At the END of the experiment, we will randomly choose one of all the envelops that contain these choices to actually show it to a client.

Now, please take the paper in front of you, and then

- Write your recommendation to the client on the paper as follows:
  "I recommend you to choose option ___."
  Please do not write anything else other than the above sentence.
- Then put the paper into the envelope. Close the envelope, do NOT seal it.
- You can refer to the paper instructions if you want to review the three options.

[**Note:** You do NOT receive a bonus for this recommendation.]

If you want, you can obtain verification that your recommendation was shown to a client should it be drawn.
For such verification, adress the envelope to yourself and sign your recommendation. You do not have to do this.

If you are finished, please click the button below. We will then come around and collect your envelope.

Finished

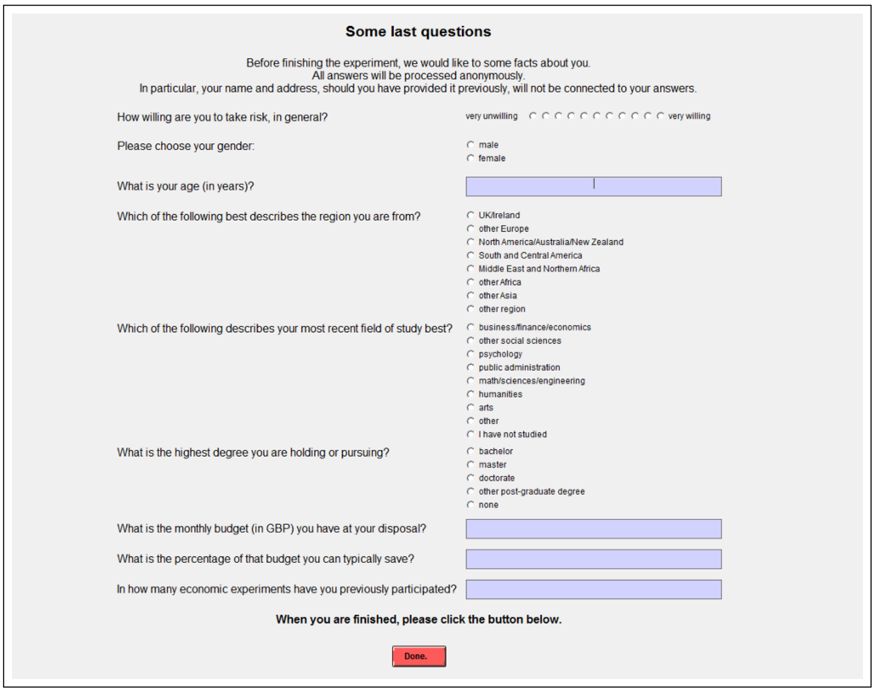

Some last questions

Before finishing the experiment, we would like to some facts about you.
All answers will be processed anonymously.
In particular, your name and address, should you have provided it previously, will not be connected to your answers.

How willing are you to take risk, in general? — very unwilling ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ very willing

Please choose your gender:
- male
- female

What is your age (in years)?

Which of the following best describes the region you are from?
- UK/Ireland
- other Europe
- North America/Australia/New Zealand
- South and Central America
- Middle East and Northern Africa
- other Africa
- other Asia
- other region

Which of the following describes your most recent field of study best?
- business/finance/economics
- other social sciences
- psychology
- public administration
- math/sciences/engineering
- humanities
- arts
- other
- I have not studied

What is the highest degree you are holding or pursuing?
- bachelor
- master
- doctorate
- other post-graduate degree
- none

What is the monthly budget (in GBP) you have at your disposal?

What is the percentage of that budget you can typically save?

In how many economic experiments have you previously participated?

When you are finished, please click the button below.

Done.

Screens 4 (top) and 5 (bottom)

Bibliography

Agarwal, S. and B. D. Itzhak (2014). Loan prospecting and the loss of soft information. *NBER Working Paper Series 19945*.

Akerlof, G. A. and W. T. Dickens (1982). The economic consequences of cognitive dissonance. *American Economic Review 71*(3), 437–447.

Alonso, R. and O. Câmara (2015). Political disagreement and information in elections. Technical report, mimeo.

Alonso, R. and O. Câmara (2016). Persuading voters. *American Economic Review* (forthcoming).

Amann, E. and W. Leininger (1996). Asymmetric all-pay auctions with incomplete information: the two-player case. *Games and Economic Behavior 14*(1), 1–18.

Anderson, S. P. and R. Renault (2006). Advertising content. *The American Economic Review 96*(1), 93–113.

Andreoni, J. and A. Brownback (2014). Grading on a curve, and other effects of group size on all-pay auctions. Technical report, National Bureau of Economic Research.

Angrist, J., D. Lang, and P. Oreopoulos (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics*, 136–163.

Aoyagi, M. (2010). Information feedback in a dynamic tournament. *Games and Economic Behavior 70*(2), 242–260.

Azacis, H. and P. Vida (2015). Collusive communication schemes in a first-price auction. *Economic Theory 58*(1), 125–160.

Baik, K. H. and J. F. Shogren (1995). Contests with spying. *European Journal of Political Economy 11*(3), 441–451.

Baik, K. H. and J. F. Shogren (1996). Contests with spying: Reply. *European Journal of Political Economy 12*(4), 735–737.

Barrachina, A., Y. Tauman, and A. Urbano (2014). Entry and espionage with noisy signals. *Games and Economic Behavior 83*, 127–146.

Baye, M. R., D. Kovenock, and C. G. De Vries (1993). Rigging the lobbying process: an application of the all-pay auction. *American Economic Review 83*(1), 289–294.

Baye, M. R., D. Kovenock, and C. G. De Vries (1996). The all-pay auction with complete information. *Economic Theory 8*(2), 291–305.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Benabou, R. and J. Tirole (2004). Willpower and Personal Rules. *Journal of Political Economy 112*(4), 848–886.

Benabou, R. and J. Tirole (2011). Identity, morals, and taboos: Beliefs as assets. *Quarterly Journal of Economics 126*(2), 805–855.

Bénabou, R. and J. Tirole (2016). Bonus Culture: Competitive Pay, Screening, and Multitasking. *Journal of Political Economy*, forthcoming.

Bergemann, D., B. A. Brooks, and S. Morris (2015). First price auctions with general information structures: Implications for bidding and revenue.

Bergemann, D. and M. Pesendorfer (2007). Information structures in optimal auctions. *Journal of Economic Theory 137*(1), 580–609.

Bertrand, M. (2009). CEOs. *Annual Review of Economics 1*(1), 121–150.

Billand, P., C. Bravard, S. Chakrabarti, and S. Sarangi (2016). Business Intelligence and Multimarket Competition. *Journal of Public Economic Theory 18*(2), 248–267.

Bodner, R. and D. Prelec (2003). Self-Signaling and Diagnostic Utility in Everyday Decision Making. In I. Brocas and J. D. Carrillo (Eds.), *The Psychology of Economic Decisions Vol.1*, Volume 1, pp. 105–126. Oxford University Press.

Bolle, F. (1996). Contests with spying: A comment. *European Journal of Political Economy 12*(4), 729–734.

Burguet, R. and Y.-K. Che (2004). Competitive procurement with corruption. *RAND Journal of Economics*, 50–68.

Cain, D. M. and A. S. Detsky (2008). Everyone's a Little Bit Biased (Even Physicians). *Journal of the American Medical Association (JAMA) 299*(24), 2893–2895.

Celentani, M. and J.-J. Ganuza (2002). Corruption and competition in procurement. *European Economic Review 46*(7), 1273–1303.

Charness, G. and U. Gneezy (2012). Strong evidence for gender differences in risk taking. *Journal of Economic Behavior & Organization 83*(1), 50–58.

Che, Y.-K. and I. L. Gale (1998). Caps on political lobbying. *American Economic Review 88*(3), 643–651.

Chen, Z. (2016). Information Disclosure in Contests: Private vs. Public Signals.

Chen, Z. C., D. Ong, and R. M. Sheremeta (2015). The gender difference in the value of winning. *Economics Letters 137*, 226–229.

Chi, C.-K., M. Pauli, and V. Juuso (2015). Common value all-pay auctions.

Coates, J. M., M. Gurnell, and A. Rustichini (2009). Second-to-fourth digit ratio predicts success among high-frequency financial traders. *Proceedings of the National Academy of Sciences 106*(2), 623–628.

Cohn, A., E. Fehr, and M. A. Maréchal (2014). Business culture and dishonesty in the banking industry. *Nature 7529*(516), 86–89.

Congleton, R. D., A. L. Hillman, and K. A. Konrad (2008). *40 Years of Research on Rent Seeking 2: Applications: Rent Seeking in Practice*, Volume 2. Springer Science & Business Media.

Croson, R. and U. Gneezy (2009). Gender differences in preferences. *Journal of Economic literature*, 448–474.

Dana, J. and G. Loewenstein (2003). A social science perspective on gifts to physicians from industry. *Journal of the American Medical Association (JAMA) 290*(2), 252–255.

Dasgupta, P., R. J. Gilbert, and J. E. Stiglitz (1982). Invention and innovation under alternative market structures: The case of natural resources. *Review of Economic Studies 49*(4), 567–582.

Dawes, R. (1990). The potential nonfalsity of the false consens effect. In R. M. Hogarth (Ed.), *Insights in Decision Making: A Tribute to Hillel J. Einhorn*, pp. 179–199. University of Chicago Press.

Denter, P., J. Morgan, and D. Sisak (2014). 'Where Ignorance is Bliss,'Tis Folly to Be Wise': Transparency in Contests. *Available at SSRN 1836905*.

Di Tella, R. D. and R. Pérez-Truglia (2015). Conveniently Upset: Avoiding Altruism by Distorting Beliefs About Others. *American Economic Review 105*(11), 3416–3442.

DiPrete, T. A. and C. Buchmann (2013). *The Rise of Women: The Growing Gender Gap in Education and What it Means for American Schools: The Growing Gender Gap in Education and What it Means for American Schools*. Russell Sage Foundation.

Duckworth, A. L. and M. E. P. Seligman (2006). Self-discipline gives girls the edge: Gender in self-discipline, grades, and achievement test scores. *Journal of Educational Psychology 98*(1), 198.

Eckel, C. C. and P. J. Grossman (2008). Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior and Organization 68*(1), 1–17.

Ehlermann-Cache, N. and Others (2007). *Bribery in public procurement: methods, actors and counter-measures.* OECD.

Ellingsen, T. (1991). Strategic buyers and the social cost of monopoly. *American Economic Review 81*(3), 648–657.

Engelmann, D. and M. Strobel (2000). The false consensus effect disappears if representative information and monetary incentives are given. *Experimental Economics 3*(3), 241–260.

Eso, P. and B. Szentes (2007). Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies 74*(3), 705–731.

Falk, A. and F. Zimmermann (2015). Information Processing and Commitment. *mimeo*.

Falk, A. and F. Zimmermann (2016). Consistency as a Signal of Skills. *Management Science*, forthcoming.

Fallis, A. G., X. Luo, and Z. Fang (2015). Self-Signaling and Prosocial Behavior: A Cause Marketing Mobile Field Experiment. *NBER Working Paper Series 21475.*

Fallucchi, F., E. Renner, and M. Sefton (2013). Information feedback and contest structure in rent-seeking games. *European Economic Review 64*, 223–240.

Fang, H. and S. Morris (2006). Multidimensional private value auctions. *Journal of Economic Theory 126*(1), 1–30.

Faro, D. and Y. Rottenstreich (2006). Affect, Empathy, and Regressive Mispredictions of Others' Preferences Under Risk. *Management Science 52*(4), 529–541.

Festinger, L. and J. M. Carlsmith (1959). Cognitive consequences of forced compliance. *Journal of Abnormal Psychology 58*(2), 203–210.

Fibich, G., A. Gavious, and A. Sela (2006). All-pay auctions with risk-averse players. *International Journal of Game Theory 34*(4), 583–599.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Fortin, N. M., P. Oreopoulos, and S. Phipps (2015). Leaving boys behind. *Journal of Human Resources 50*(3).

Frank, R. H. and P. J. Cook (2010). *The winner-take-all society: Why the few at the top get so much more than the rest of us.* Random House.

Fu, Q., Q. Jiao, and J. Lu (2014). Disclosure policy in a multi-prize all-pay auction with stochastic abilities. *Economics Letters 125*(3), 376–380.

Fulton, S. A., C. D. Maestas, L. S. Maisel, and W. J. Stone (2006). The sense of a woman: Gender, ambition, and the decision to run for congress. *Political Research Quarterly 59*(2), 235–248.

Ganuza, J.-J. and J. S. Penalva (2010). Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica 78*(3), 1007–1030.

Gershkov, A. and M. Perry (2009). Tournaments with midterm reviews. *Games and Economic Behavior 66*(1), 162–190.

Gneezy, A., U. Gneezy, G. Riener, and L. D. Nelson (2012). Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences 109*(19), 7236–7240.

Gneezy, U., S. Saccardo, M. Serra-Garcia, and R. van Veldhuizen (2016). Motivated Self-deception and Unethical Behavior. *mimeo*.

Goeree, J. K. and T. Offerman (2003). Competitive bidding in auctions with private and common values. *Economic Journal 113*(489), 598–613.

Grossman, Z. (2015). Self-Signaling and Social-Signaling in Giving. *Journal of Economic Behavior & Organization 117*, 26–39.

Haisley, E. C. and R. A. Weber (2010). Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior 68*(2), 614–625.

Harrison, G. W., M. I. Lau, E. E. Rutström, and M. Tarazona-Gómez (2013). Preferences over social risk. *Oxford Economic Papers 65*(1), 25–46.

Hickman, B. R. (2014). Pre-college human capital investment and affirmative action: a structural policy analysis of US college admissions. *University of Chicago, mimeo*.

Hilgartner, S. (2000). *Science on stage: Expert advice as public drama.* Stanford University Press.

Hillman, A. L. and J. G. Riley (1989). Politically contestable rents and transfers. *Economics & Politics 1*(1), 17–39.

Hoch, S. J. (1987). Perceived consensus and predictive accuracy: The pros and cons of projection. *Journal of Personality and Social Psychology 53*(2), 221–234.

Hsee, C. K. and E. U. Weber (1997). A fundamental prediction error: self-other discrepancies in risk preference. *Journal of Experimental Psychology: General 126*(1), 45–53.

Jackson, M. O. (2009). Non-existence of equilibrium in Vickrey, second-price, and English auctions. *Review of Economic Design 13*(1-2), 137–145.

Johnson, J. P. and D. P. Myatt (2006). On the simple economics of advertising, marketing, and product design. *American Economic Review 96*(3), 756–784.

Kahneman, D., P. P. Wakker, and R. Sarin (1997). Back to Bentham? Explorations of Experienced Utility. *The Quarterly Journal of Economics 112*(2), 375–405.

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

Klose, B. S. and P. Schweinzer (2014). Auctioning risk: The all-pay auction under mean-variance preferences. *University of Zurich, Department of Economics, Working Paper* (97).

Konow, J. (2000). Fair Shares : Accountability and Cognitive Dissonance in Allocation Decisions Fair Shares. *American Economic Review 90*(4), 1072–1091.

Konrad, K. A. (2004). Altruism and envy in contests: an evolutionarily stable symbiosis. *Social Choice and Welfare 22*(3), 479–490.

Konrad, K. A. (2009). Strategy and dynamics in contests. *OUP Catalogue*.

Kovenock, D., F. Morath, and J. Munster (2015). Information sharing in contests. *Journal of Economics & Management Strategy 24*(3), 570–596.

Kozlovskaya, M. (2016). Industrial espionage in duopoly games.

Krishna, V. (2009). *Auction theory*. Academic press.

Krishna, V. and J. Morgan (1997). An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory 72*(2), 343–362.

Kunda, Z. (1990). The Case for Motivated Reasoning. *Psychological Bulletin 108*(3), 480–498.

Kunda, Z. (1992). Can Dissonance Theory Do It All? *Psychological Inquiry 4*(3), 337–339.

Laband, D. N. and J. P. Sophocleus (1988). The social cost of rent-seeking: First estimates. *Public Choice 58*(3), 269–275.

Li, H. and X. Shi (2013). Discriminatory information disclosure.

Linnainmaa, J. T., B. T. Melzer, and A. Previtero (2016). The Misguided Beliefs of Financial Advisors. *mimeo*.

Loewenstein, G., S. Issacharoff, C. Camerer, L. Babcock, C. Camerer, and L. Babcock (1993). Self-Serving Assessments of Fairness and Pretrial Bargaining. *Journal of Legal Studies 22*(1), 135–159.

Loewenstein, G. and D. Schkade (1999). Wouldn't It Be Nice? Predicting Future Feelings. In D. Kahneman, E. Diener, and N. Schwarz (Eds.), *Well-being: The* *foundations of hedonic psychology*, Chapter 5, pp. 85–105. Russell Sage Foundation.

Lu, J., H. Ma, and Z. Wang (2016). Ranking Disclosure Policies in All-Pay Auctions with Incomplete Information.

Lu, J. and S. O. Parreiras (2014). Asymmetric all-pay auctions with incomplete information, the two player case redux.

Malmendier, U. and D. M. Shanthikumar (2014). Do security analysts speak in two tongues? *Review of Financial Studies 27*(5), 1287–1322.

Marks, G. and N. Miller (1987). Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin 102*(1), 72–90.

Matthews, S. (1984). Information acquisition in discriminatory auctions. *Bayesian models in economic theory, ed. by M. Boyer, and R. Kihlstrom 49*, 1477–1500.

Mazar, N., O. Amir, and D. Ariely (2008, December). The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *Journal of Marketing Research 45*(6), 633–644.

Miettinen, P. (2013). Information acquisition during a Dutch auction. *Journal of Economic Theory 148*(3), 1213–1225.

Morath, F. and J. Münster (2008). Private versus complete information in auctions. *Economics Letters 101*(3), 214–216.

Morath, F. and J. Münster (2013). Information acquisition in conflicts. *Economic Theory 54*(1), 99–129.

Mullainathan, S., M. Noeth, and A. Schoar (2012). The Market For Financial Advice. An Audit Study. *NBER Working Paper Series 17929.*

Mullen, B., J. L. Atkins, D. S. Champion, C. Edwards, D. Hardy, J. E. Story, and M. Vanderklok (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology 21*(3), 262–283.

Niederle, M. (2014). Gender. In *Handbook of Experimental Economics.*

Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 1067–1101.

Parreiras, S. O. and A. Rubinchik (2010). Contests with three or more heterogeneous agents. *Games and Economic Behavior 68*(2), 703–715.

Persico, N. (2000). Information acquisition in auctions. *Econometrica 68*(1), 135–148.

Posner, R. A. (1974). Social costs of monopoly and regulation.

Rabin, M. (1994). Cognitive dissonance and social change. *Journal of Economic Behavior & Organization 23*, 177–194.

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

Rentschler, L. and T. L. Turocy (2016). Two-bidder all-pay auctions with interdependent valuations, including the highly competitive case. *Journal of Economic Theory 163*, 435–466.

Rosen, S. (1981). The economics of superstars. *American Economic Review 71*(5), 845–858.

Rosen, S. (1986). Prizes and Incentives in Elimination Tournaments. *American Economic Review*, 701–715.

Ross, L., D. Greene, and P. House (1977). The "false consensus effect": An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology 13*(3), 279–301.

Roth, B. and A. Voskort (2014). Stereotypes and false consensus: How financial professionals predict risk preferences. *Journal of Economic Behavior and Organization 107*, 553–565.

SEC (2011). Study on Investment Advisers and Broker-Dealers. *Report by the staff of the U.S. Securities and Exchange Commission.*

Serena, M. (2015). Harnessing beliefs to stimulate efforts. *Available at SSRN 2686543*.

Sheremeta, R. M. (2010). Expenditures and information disclosure in two-stage political contests. *Journal of Conflict Resolution 54*(5), 771–798.

Shi, X. (2012). Optimal auctions with information acquisition. *Games and Economic Behavior 74*(2), 666–686.

Siegel, R. (2009). All-pay contests. *Econometrica 77*(1), 71–92.

Siegel, R. (2013). Asymmetric Contests with Interdependent Valuations.

Solan, E. and L. Yariv (2004). Games with espionage. *Games and Economic Behavior 47*(1), 172–199.

Tan, X. (2016). Information revelation in auctions with common and private values. *Games and Economic Behavior 97*, 147–165.

Taylor, R. and J. Giles (2005). Cash interests taint drug advice. *Nature 437*, 1070–1071.

Tian, G. and M. Xiao (2007). Endogenous Information Acquisition on Opponents' Valuations in Multidimensional First Price Auctions.

Trivers, R. (2011). *The folly of fools: The logic of deceit and self-deception in human life.* Basic Books.

Tullock, G. (1967). The welfare costs of tariffs, monopolies, and theft. *Economic Inquiry 5*(3), 224–232.

United States Congress (2010). Dodd-Frank Wall Street Reform and Consumer Protection Act.

Wang, T. (2016). Espionage and Disclosure of Cost Information in Cournot Duopoly.

Wang, Y. (2013). Bayesian persuasion with multiple receivers. *Available at SSRN 2625399*.

Whitney, M. E. and J. D. Gaisford (1999). An Inquiry Into the Rationale for Economic Espionage*. *International Economic Journal 13*(2), 103–123.

Wright, B. D. (1983). The economics of invention incentives: Patents, prizes, and research contracts. *The American Economic Review 73*(4), 691–707.

Zhang, H. (2015). Three Essays on the Incentives for Information Acquisition and Information Sharing in Competitive Environments.

Zhang, J. and J. Zhou (2016). Information Disclosure in Contests: A Bayesian Persuasion Approach. *Economic Journal*.

Zingales, L. (2015). Presidential Address: Does Finance Benefit Society? *Journal of Finance 70*(4), 1327–1363.