Essays on Delegation and Mechanisms

by

Volodymyr Baranovskyi

Department of Economics
Duke University

Date: ___________________________
Approved:

____
Attila Ambrus, Supervisor

____
Curtis R. Taylor

____
Xiao Yu Wang

____
Huseyin Yildirim

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2018

Abstract

Essays on Delegation and Mechanisms

by

Volodymyr Baranovskyi

Department of Economics
Duke University

Date: __________________________

Approved:

Attila Ambrus, Supervisor

Curtis R. Taylor

Xiao Yu Wang

Huseyin Yildirim

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2018

Abstract

This dissertation consists of three theoretical essays on delegation and mechanism design.

Chapter 2 is co-authored with Attila Ambrus and Aaron Kolb. We investigate competition in a delegation framework. An uninformed principal is unable to perform a task herself and must solicit proposals from two biased and imperfectly informed experts. In the focal equilibrium, the principal seeks to offset the bias of the experts, but when the experts are motivated more by ideology than career concerns, they increase the bias of their proposals in anticipation. Despite this ideological winner's curse, we show that having a second expert can benefit the principal, even if the two experts have the same biases or if the first expert is known to be unbiased. In contrast with other models of expertise, in our setting the principal prefers experts with equal rather than opposite biases. The principal may also benefit from commitment to an "element of surprise," making an ex post suboptimal choice with positive probability.

Chapter 3 is co-authored with Sergii Golovko. We study an auction environment in which after the sale, the seller has the opportunity to verify the winner's ex-post value and impose a limited punishment for "underbidding." Investigating how the seller should approach this opportunity, we show that even small penalties allow the seller to significantly increase her revenue. In our environment, the first-price auction with an optimally chosen penalty rule is optimal among all winner-pay auctions. Before the auction begins, the seller recommends a bidding strategy to the bidders.

If the auction winner bids at least as much as the seller has suggested, the winner is not punished; if, on the other hand, the winner does not bid as much as has been recommended, he is punished, with the penalty increasing as the buyer deviates more and more from the recommendation. Our results indicate several qualitative differences from standard (without ex-post punishments) auctions. In equilibrium, buyers bid more aggressively; the optimal reserve price is lower; and the revenue-equivalence principle does not hold—we state conditions under which a first-price auction is superior to a second-price auction. Our results also lead us to suggest the following recommendation for policymakers: A government may increase its revenue when auctioning publicly owned assets by providing tax concessions to buyers who submit sufficiently high bids.

Chapter 4 is co-authored with Attila Ambrus. As in Chapter 2, principal is trying to solicit information from multiple, incompletely informed experts. However, here we allow for action choice (policy) and monetary transfers to be conditional on reports. Additionally, we investigate an environment in which monetary transfers may be conditioned on realized state (ex post state verification). We show that under certain assumptions the principal can solicit all information from experts, while extracting all surplus.

Contents

List of Figures

Acknowledgements

I am deeply indebted to my advisor Attila Ambrus for his continuous guidance, encouragement and readiness to help. I am also thankful to him for engaging me as his research assistant and giving me an opportunity to learn. I would like to thank Curtis Taylor for his help and support during the whole PhD program. I want to thank Huseyin Yildirim for his valuable advice. Also I am thankful to Xiao Yu Wang for her comments and help.

I would like to thank Aaron Kolb for coauthoring Chapter 2, and Sergii Golovko for coauthoring Chapter 3.

1

Introduction

This dissertation consists of three self-contained chapters. The introductions of the first, second and third chapters are in section 2.1, 3.1 and 4.1, correspondingly.

2

A Delegation-Based Theory of Expertise (with Attila Amrus and Aaron Kolb)

2.1 Introduction

There are many situations in which a principal lacks the knowledge and expertise to perform a certain task, and therefore has to delegate the job to a qualified expert. Examples include a candidate running for office who has to hire an expert to work out her economic agenda, or the CEO of a pharmaceutical company who must delegate building a research and development division to a scientist. Further complicating the principal's situation is that experts tend to have systemic biases, preferring suboptimal actions from the principal's perspective.

In this paper we investigate a model in which a principal has to delegate a task to one of two experts. The need to delegate differentiates our model from models of expertise in which experts send cheap talk recommendations to the principal, such as Krishna and Morgan (2001b). In particular, we consider the following game. First, experts receive noisy and conditionally independent signals of a single dimensional state variable. The principal's ideal action is equal to the state, but each expert has a constant bias (either positive or negative) and a resulting ideal point different from the sender's. Next, the experts simultaneously propose actions. A proposal is assumed to bind the expert to perform the given action whenever the principal delegates the task to him.[1] The principal then chooses one of the two offers, and the corresponding action is taken by the given expert. We are motivated by situations in which the principal originally has much less knowledge about the state than the experts, and correspondingly we assume that the principal's prior is improper uniform (diffuse) over a state space represented by the real line.

The particular game form we investigate is motivated by various applications. In general, our model best applies to situations in which the principal lacks the knowledge to implement or initiate changes in the proposed actions, so all she can do is solicit different proposals and choose one of them. The assumption that proposals commit the experts corresponds to common law, according to which an offer is a statement of terms on which the offeror is willing to be bound, and it shall become binding as soon as it is accepted by the person to whom it is addressed.[2] A different application for our model is political competition: starting from Downs (1957), most papers on political competition in a Hotelling (1929) framework assume that candidates are committed to the policies they announce in the campaign, and the electorate can only choose between the policies announced by the candidates.[3] In this context the bonus corresponds to the rents from being in office. Yet another application for our model is a setting where a legislative body (floor) seeks legislative proposals for the same bill from multiple committees, using a modified rule (see Gilligan and Krehbiel (1989), Krishna and Morgan (2001a)), meaning that the floor

[1] Even if the principal might not have the knowledge to verify whether the expert indeed chose the action that he proposed, outside experts might be able to verify if that was the case and hence penalties can be imposed on experts deviating from their proposals.

[2] See Treitel (1999), p8.

[3] For theoretical motivations for this assumption, and empirical relevance in the political competition context, see Pétry and Collette (2009), Kartik et al. (2015), and papers cited therein.

cannot amend the proposals and can only accept one of the proposed bills without modification, conforming to the basic assumptions of our model.[4]

Experts have preferences over both the policy outcome and whether they are selected. In particular, to model the latter, we allow for a bonus to the chosen expert, either as a monetary payment or as a non-monetary benefit, such as increased prestige in his profession. We investigate two cases, with the bonus amount given exogenously in one case and optimally chosen by the principal in the other.

The above game is very complex in general, due to the size of the strategy space. In this paper we restrict attention to equilibria in which the experts' strategies are stationary with respect to signals, meaning that each expert's proposal is equal to his signal plus a constant. We consider focusing on such strategies, that treat states symmetrically, natural in a game with diffuse prior and preferences that are relatively stationary in the state, since in such an environment all states are perfectly symmetric (nothing distinguishes them besides the labeling, which can be arbitrarily rescaled). A further motivation comes from Ambrus and Kolb (2018), in which they examine the possibility of extending the concept of ex ante expected payoffs to a larger class of games with diffuse prior (and hence bringing them into the realm of traditional game theory, in which payoffs have to be well-defined for any strategy profile). Their paper shows that in our game expert strategies need to be restricted to be stationary in order for well-defined ex ante expected payoffs to exist. In particular, given some weak conditions on the principal's set of strategies, essentially stationary strategies are the only strategies for which well-defined limit expected payoffs exist for any strategy profile when taking a sequence of proper priors diffusing (converging in a formal sense to the diffuse prior), with the limit not depending on the particular

[4] Gilligan and Krehbiel (1989) analyze this situation with an additional option to the floor, in the form of not accepting either of the proposals and opting for a status quo outcome. As opposed to our model, Gilligan and Krehbiel (1989) assume perfectly informed experts (committees), which fundamentally changes the strategic interaction.

choice of sequence. This result shows that in order to obtain well-defined ex ante expected payoffs corresponding to all strategy profiles, one would need to restrict experts' strategies to stationary ones. In the current paper, stationary strategies are either characterized by a constant *markup* — the expert adds a fixed markup to whatever his signal is — or mixtures thereof. Here, we do not restrict experts' strategies, but simply focus on equilibria in which they play stationary strategies, and similarly to existing game theoretical models of improper prior (Friedman (1991), Klemperer (1999), Morris and Shin (2002, 2003), Myatt and Wallace (2014)), we only evaluate payoffs in the interim stage (after signal realizations).

Our first result shows that if experts play stationary strategies then we can restrict attention to the following simple strategies for the principal: always choosing expert 1's offer (effectively delegating the action choice to expert 1), always choosing expert 2's offer, always choosing the minimum of the two offers, and always choosing the maximum of the two offers. In particular, whenever the sum of markups (differences between proposed action and signal) by the experts is positive, the unique best response of the principal is always choosing the smaller of the two offers, while if the sum of the two markups is negative, the unique best response of the principal is always choosing the larger of the two offers.[5]

It is easy to show that a (Bayesian Nash) equilibrium with stationary strategies always exists in our model, in the form of delegating the task to one of the experts. Formally, one expert always proposing his ideal action conditional on the signal he observes (equal to the signal plus his bias), the other expert proposing his signal minus the first expert's bias, and the principal always delegating the task to the

[5] These strategies are also feasible for a principal who can only process information in a coarse way, being only able to make binary comparisons between two offers and lacking the ability to measure the difference between them, as consumers in Kamenica (2008). Therefore such information processing constraints would not hurt the principal in the equilibria we investigate.

first expert constitutes an equilibrium.[6] The question is whether there exist other equilibria of the game, in which the principal either always chooses the minimum or always chooses the maximum of the two proposals - hence her choice depends nontrivially on the proposals. We assume without loss of generality that the sum of the biases of the two experts is nonnegative.

When the bonus is small relative to the noise in the experts' signals, there exists an equilibrium in which the principal chooses the minimum of the two offers and the experts apply markups above their biases; if both biases are positive, this means that both experts exaggerate their biases. This result is contrary to a naive intuition that experts in competition should move toward the center. We call this equilibrium "min," as the principal chooses the minimum of the two offers, and since it is optimal in several respects (discussed below), it is our focal equilibrium. To illustrate, suppose that experts have the same biases and the bonus is small. Then in this min equilibrium both experts propose actions strictly above their ideal actions based purely on their private signals. This is because, similarly to the winner's curse phenomenon in common value auctions, being selected by the principal contains information on the other expert's signal (namely that his signal is higher), changing the optimal action of the expert. In equilibrium, proposals have to be optimal conditional on the event that the other expert's action proposal is higher.

We also show that if the experts' signals are noisy enough then, for a subset of the range of bonuses for which a min equilibrium exists, there also exists a "max" equilibrium in which the principal selects the maximum of the two proposals, and experts propose actions on average below their signals. The strategic forces are similar to those in the min equilibrium: the fact that the maximum of the two offers is selected pushes proposals downwards, and for noisy signals on average experts modify

[6] There are other Bayesian Nash equilibria on mixed strategies with the same outcome, in which one expert always proposes his ideal point, the other expert "babbles" (randomizes over possible messages he can send), and the principal always delegates the task to the first expert.

their proposals downward relative to their signals. This type of equilibrium does not exist when the signals are very precise, because then the information conveyed from being selected does not shift the optimal proposals of the experts enough to make markups negative on average. We show that even when the max equilibrium exists, the principal prefers the min equilibrium to it, and thus for further welfare comparisons we need only consider min equilibrium and simple delegation to the less-biased expert.

The feature of the above equilibria that similarly biased experts exaggerate in a particular direction (and hence the principal should choose the proposal least in the direction of the exaggeration) is in line with empirical evidence. For example, Zitzewitz (2001), Bernhardt et al. (2006) and Chen and Jiang (2006) find that financial analysts systematically exaggerate their forecasts relative to unbiased forecasts based on the analysts' information sets, while Iezzoni et al. (2012) report that 55% of doctors in a survey said that in the previous year they had been more positive about patients' prognoses than their medical histories warranted.

We compare the principal's welfare between min equilibrium and simple delegation in order to find the principal-optimal equilibrium. In general, the comparison is complicated and can go either way, but for several focal cases of interest, the optimum is the min equilibrium. The principal is always better off in the min equilibrium when the experts have the same biases (as in settings where all available experts have similar agendas), or when the experts have exactly opposite biases.[7] This result holds even when the bonus is zero and hence there is no competition among experts for being selected. The principal also prefers min equilibrium to simple delegation when one expert's bias is positive and the other's is zero; this is despite the fact that under simple delegation the unbiased expert's incentives are perfectly aligned with the

[7] The principal's payoffs are continuous in the parameters of the model for a particular type of equilibrium, hence the above comparisons are the same when the absolute values of the biases are close to each other but not exactly equal.

principal's. The intuition for the result is that in the min equilibrium the principal extracts some additional information the second expert, which reduces the variance of the chosen action, and this benefit always outweighs any cost associated with higher markups. Applied to a political setting, the latter comparison between min equilibrium and simple delegation helps explain what voters may otherwise perceive as corruption – a politician may want to seek advice from a biased expert, even if it is common knowledge that she already has access to an unbiased expert.

We also compare the principal's payoffs when experts have equal versus opposite biases, and our results here are in contrast with some of the existing literature. In our model, assuming the min equilibrium is played, having two experts with identical biases yields a higher payoff than having two antagonist experts with opposite biases. In general, the expected bias of the implemented action is smaller with antagonist experts than with experts having the same bias, but this benefit is outweighed by a higher variance of the implemented action that arises because the expert with the lower bias is selected most of the time, and so the information from the other expert's signal is only utilized to a limited extent. This result contrasts models of competition in persuasion (Milgrom and Roberts (1986), Gentzkow and Kamenica (2017)), in which antagonist experts benefit the principal by pressing each other to reveal more information,[8] and with the multi-sender cheap talk model of Krishna and Morgan (2001b), in which having a second sender with the same bias does not benefit the receiver.[9] See also Shin (1998) and Dewatripont and Tirole (1999) for

[8] Experts with identical agendas can be better for the principal than experts with opposing agendas in the persuasion model of Bhattacharya and Mukherjee (2013). The mechanism is rather different than in our paper, though: with similar experts an undesirable default action can provide strong incentives for both experts to reveal information.

[9] As opposed to cheap talk models with multiple senders and one receiver, where it tends to be better for information revelation if the senders are oppositely biased from the point of view of the receiver, in committee settings, where committee members can reveal information to each other, it helps information revelation if members have more similar preferences - see for example Li and Suen (2009).

different types of models making the case for adversarial procedures.

The principal's expected payoff depends in a complicated way on the noise in the experts' signals, and on the amount of the bonus. Hence, for these comparative statics we focus on the case of equally biased experts. Even in this case, the effect of the variance of the experts' signals is ambiguous. An increased precision of experts' signals reduces the variance of the implemented action conditional on the state. For small bonuses, this unambiguously increases the principal's expected payoff. However, for larger bonuses, it might benefit the principal in the min equilibrium if the experts increase their markups,[10] which can result from increasing the variance of the signals. We provide an exact characterization (for equally biased experts) for when a decrease in the variance of experts' signals benefits the principal.

Increasing the bonus reduces the absolute values of the experts' markups, hence bringing their proposals closer to truthful reporting, both in the min and max equilibria. Intuitively, a higher bonus increases competition among experts, leading them to decrease their proposals in the min equilibrium and increase their proposals in the max equilibrium. In the min equilibrium, this initially improves the principal's expected payoff by decreasing the expected bias of the implemented action. There is a threshold level of bonus though at which the expected bias of the implemented action becomes zero, and increasing the bonus above this threshold decreases the principal's payoff. When the bonus comes from exogenous sources, the optimal bonus from the principal's perspective is always strictly positive, and is on the interior of the interval of bonuses for which the min equilibrium exists. When the bonus is paid by the principal, the optimal bonus amount is always strictly smaller than in the previous case, and depending on the parameters it can be either strictly positive or zero.

[10] This is related to the chunkiness of the principal's possible choices in our model: for certain parameter values sticking with choosing the minimum of two proposed actions is still optimal for the principal, even though it leads to the implemented action being negatively biased. This can happen if choosing the maximum offer would lead to an even larger positive bias. These are the cases when an increase in the expectation of the minimum offer benefits the principal.

In the political competition application of the model the result implies that a small amount of office-seeking motivation can be beneficial for voters, but at higher levels a further increase in office-seek motivation can adversely affect voters' welfare.

We consider two extensions of our model, for equally biased experts. In the first one we allow the principal to commit ex ante to any mixture of simple strategies, and show that for bonuses that are not too large, such commitment leads to the same outcome as in the min equilibrium of the original game, hence the ability to commit does not improve the principal's welfare. On the other hand, we show how committing to choosing an inferior offer with some small probability - introducing an element of surprise - can improve the principal's welfare in the case of opposite biases of large magnitude. The intuition is that in the min equilibrium, the expert with the positive bias faces a very large winner's curse, and applies a markup well above his bias. The other expert faces an almost negligible winner's curse and applies a markup just slightly above his bias. Now by threatening to choose the higher offer some of the time, the principal induces the first expert to reduce his markup drastically and the second expert to raise his markup, so both markups move closer to zero. The cost of this deviation to the principal lies in mistakenly choosing the higher offer, but when the magnitude of the biases is sufficiently large, the benefit outweighs the cost. The finding that the principal can benefit from committing to a mixed strategy in certain situations is consistent with an observed pattern of regulatory uncertainty. Ederer et al. (2018) show similarly that commitment to an opaque reward scheme reduces temptation to game the system in a principal-agent environment.

In the second extension we drop the dependence of the unselected expert's payoff on the implemented action, and instead assume that the expert gets a fixed outside option payoff. This variant of the model is more realistic in market transaction situations, such as when experts are car mechanics or doctors. A car mechanic might be biased towards larger repairs than necessary, but typically he does not care about what type of repair is chosen in case a different mechanic is selected to do the job. The analysis of this version of the model is more involved, but we show that under some parameter restrictions similar min and max equilibria exist as in the baseline model. The fact that the unselected expert gets a fixed outside payment increases the experts' proposals in the min equilibrium, and decreases them in the max equilibrium.

2.2 Related Literature

The literature on delegation so far mainly focused on either the question of delegating the action choice versus retaining the right to take the action (Dessein (2002), Li and Suen (2004)) or on optimally constraining the action choices of a particular expert (Holmström (1977), Melumad and Shibano (1991), Alonso and Matouschek (2008)). Krishna and Morgan (2008) investigate how monetary incentives can be used optimally in delegation to a single agent. More related to our investigation are papers introducing policy-relevant private information on the part of candidates into the context of the classic Downs (1957) model of political competition: Heidhues and Lagerlöf (2003), Laslier and Van der Straeten (2004), Loertscher et al. (2012), Gratton (2014), and closest to our model Kartik et al. (2015), as the latter focuses on cases when voters are relatively uninformed. Similarly to our setting, in the above papers politicians receive independent private signals about the state of the world and hence the optimal policy from the electorate's point of view. The main difference relative to our model is that the politicians in the above models do not have policy preferences, and they are purely office-motivated. For this reason neither the own private information nor the rival's private information directly affects their expected payoffs, and the candidates play a zero-sum game. In contrast, in our model the experts' signals are directly payoff-relevant for them, and their interests are partially aligned, as in higher states they would both like to induce higher actions. This leads to different equilibrium dynamics than in Kartik et al. (2015), and to distinct conclusions: in particular, in their model the electorate can never strictly benefit from the presence of a second candidate (relative to just a single one).[11]

There is also a line of literature extending the Downs (1957) model framework to politicians having mixed motivation (having both policy preferences and wanting to win), as in our model, starting with Wittman (1983) and Calvert (1985). Schultz (1996), Martinelli (2001) and Martinelli and Matsui (2002) introduce asymmetric information in this context, but as opposed to our paper, with perfectly informed politicians. This leads to different conclusions, including that full revelation of information is possible in equilibrium when policy preferences are not too extreme. Callander (2011) considers a model of sequential elections with imperfectly informed politicians having mixed motivations, but the issues investigated are different from ours and inherently dynamic: searching for a good policy in a complex environment, by trial and error.

Outside the delegation literature, Prendergast (1993) considers a context in which both a worker and a manager observes the state with noise, and the worker also receives a noisy signal about the manager's observation. In this model the worker has an incentive to cater to the manager and bias her report toward what she thinks the manager's observation is. Gerardi et al. (2009) investigates aggregation of expert opinions through a particular mechanism that approximates the first best outcome if signals are very accurate. Pesendorfer and Wolinsky (2003) investigate the effects

[11] Our model approximates the model in Kartik et al. (2015) when the bonus payment is very large and so the agents mainly care about being selected. We find that for very large bonuses the only equilibria in our model involve delegating the action choice to a single agent, which is in line with the result on maximum informativeness of political competition in Kartik et al. (2015). Correspondingly, Kartik et al. (2015) discuss an extension of their model in which they show that allowing a small amount of ideological motivation for the candidates, and assuming that they are close to unbiased from the electorate's point of view implies that in equilibrium one candidate must be winning ex ante with probability close to 1. These results suggest that there is no discontinuity between no policy preference versus a small amount of policy preference for the agents. Our paper mainly focuses on cases in which agents' policy preferences are relatively important.

of being able to solicit a second opinion from a different expert, in a dynamic model in which experts are not biased but it is costly for them to gather information.

Another line of literature investigates multi-sender extensions of the cheap talk model of Crawford and Sobel (1982), and finds that under certain conditions there can be equilibria in which the receiver can extract full or almost full information from the senders (Gilligan and Krehbiel (1989), Austen-Smith (1993), Wolinsky (2002), Battaglini (2002, 2004), Ambrus and Takahashi (2008), Ambrus and Lu (2014)). As opposed to the above papers, we investigate settings in which the principal cannot solicit information from experts and then take the action choice herself. Lastly, Ottaviani and Sørensen (2006) consider a model with multiple experts with reputational concerns reporting sequentially on privately observed signals. The issues they focus on (potential herding behavior of experts) are very different than in the current paper.

2.3 Base Model

We consider the following multi-stage game with incomplete information. There are three players: a principal and two experts. The set of states of the world is $\mathbb{R}$, and we assume that the common prior distribution of states is diffuse (improper uniform).

In stage 0 state $\theta \in \mathbb{R}$ realizes. In stage 1 each expert $i = 1, 2$ receives a noisy private signal about the state of the world $s_i = \theta + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma^2)$, and ϵ_1 and ϵ_2 are independent. In stage 2 each expert i proposes an action $a_i \in \mathbb{R}$ to the principal. In stage 3 the principal chooses one of the two experts, who then implements the action he proposed in stage 2.

Let real-valued functions $a_1(s_1)$ and $a_2(s_2)$ denote the strategies of expert 1 and expert 2 respectively, while $C(a_1, a_2) \in \{1, 2\}$ is the principal's choice strategy. If action $a = a_{C(a_1,a_2)}$ is taken then the principal's payoff is $V(a, \theta) = -(a - \theta)^2$, and the payoff of expert $j = 1, 2$ is $U_j(a, \theta, C(a_1, a_2)) = \mathbb{1}\{j = C(a_1, a_2)\}B - (a - \theta -$

$b_j)^2$, where the indicator function $\mathbb{1}\{j = C(a_1, a_2)\}$ equals 1 if $j = C(a_1, a_2)$ and 0 otherwise. We call b_i the *bias* of expert i. Note that each expert experiences a quadratic loss which depends on the action chosen but not the identity of the chosen expert; in Section 2.6, we analyze the case where the quadratic loss applies only to the chosen expert. Without loss of generality, we assume that $b_1 + b_2 \geqslant 0$ and $b_1 \geqslant b_2$. We further assume that all parameters of the game are common knowledge.

In the analysis below, we focus on perfect Bayesian equilibria in which experts' strategies are stationary in the following sense: $a_i(s_i) = s_i + k_i$, where $k_1, k_2 \in \mathbb{R}$. In words, each expert applies a constant *markup* to his signal when forming a proposed action. With slight abuse of notation, we use simply $(k_1, k_2, C(a_1, a_2))$ to denote a strategy profile with constant markup strategies.

2.3.1 Best Responses to Stationary Sender Strategies

Here we analyze the principal's best response to constant markup strategies, where each expert's offer is his signal translated by a constant, and show that the best response only depends on the sum of the markups. As the principal has a quadratic loss function, her expected payoff can be decomposed into losses from the uncertainty about the true state (which is independent of her action) and the losses from the expected difference between the chosen action and the true state. Therefore the principal prefers the offer which is closer to her posterior expectation of the true state. After observing the offers, the principal's expectation about the true state is lower (higher) than the average of the experts' offers if and only if the sum of the markups is positive (negative). Figure 2.1 illustrates a case where the markups have positive sum.

Let $\arg\min\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 < a_2$, $\{2\}$ if $a_1 > a_2$, and $\{1, 2\}$ if $a_1 = a_2$. Similarly, let $\arg\max\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 > a_2$, $\{2\}$ if $a_1 < a_2$, and $\{1, 2\}$ if $a_1 = a_2$.

Theorem 1. *If experts follow constant markup strategies $a_i(s_i) = s_i + k_i$, then*

- *if $k_1 + k_2 > 0$, the principal strictly prefers the lower offer, and $C(a_1, a_2) \in \arg\min\{a_1, a_2\}$;*
- *if $k_1 + k_2 < 0$, the principal strictly prefers the higher offer, and $C(a_1, a_2) \in \arg\max\{a_1, a_2\}$;*
- *if $k_1 + k_2 = 0$, the principal is indifferent between the offers.*

Proof. After observing both offers, the principal updates her belief:

$$\theta | a_1, a_2 \sim N(\frac{1}{2}(a_1 + a_2) - \frac{1}{2}(k_1 + k_2), \frac{\sigma^2}{2}).$$

Therefore the principal's expected utility from choosing offer a_1 or a_2 is:

$$V(a_1) = \mathbb{E}[-(\theta - a_1)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_1)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_2 - a_1 - k_1 - k_2)\right]^2$$

$$V(a_2) = \mathbb{E}[-(\theta - a_2)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_2)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_1 - a_2 - k_1 - k_2)\right]^2.$$

Hence, $V(a_1) - V(a_2) = (a_2 - a_1)(k_1 + k_2)$, which immediately implies the statements in the theorem. □

Theorem 1 illustrates a contrast between our paper and the standard communication literature. In the latter, where proposals are mere communication and the principal has the ability to choose any action, the principal would back out the signals $s_i = a_i - k_i$ and would optimally take action $\frac{s_1+s_2}{2}$. Here, the principal knows that this would be the best action but is unable to take an action other than the two offers provided.

An equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be an *min equilibrium*, if the minimum proposal is accepted and on average experts adjust their signals upwards: $k_1 + k_2 \geqslant 0$

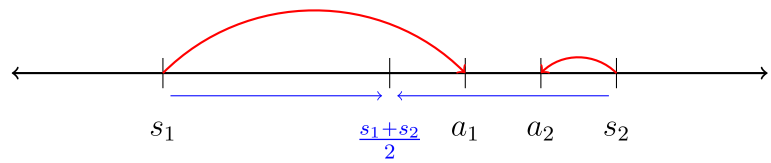

FIGURE 2.1: The principal updates beliefs and chooses the lower offer (for $k_1 > 0, k_2 < 0, k_1 + k_2 > 0$).

and $C(a_1, a_2) \in \arg\min(a_1, a_2)$. In the min equilibrium, the principal's updated expectation of the state of the world is lower than the average of the two offers. Her best response is to choose the lower offer, which is closer to her expectation, as demonstrated in Figure 1. Likewise, an equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be a *max equilibrium* if $k_1 + k_2 \leqslant 0$ and $C(a_1, a_2) \in \arg\max(a_1, a_2)$.

When $k_1 + k_2 = 0$ then the principal's posterior expectation of θ is exactly the average of the two offers, and she is indifferent between the two.

2.3.2 Simple delegation

In our game there always exist simple pure strategy equilibria in which the principal always chooses the same expert, independently of two offers, in effect delegating the decision to her. In particular, Theorem 1 implies that if expert i chooses constant markup $k_i = b_i$ and the other expert j chooses constant markup $k_j = -b_i$ then the principal is always indifferent between the two offers, and she might as well always choose expert i. Given this strategy of the principal, expert i's best response is choosing exactly markup b_i, which in expectation implements his ideal action. The other expert has no profitable deviation since his proposal is never accepted. While such an equilibrium exists for each of the two experts, it is more natural to consider the one in which the principal always chooses the expert with the smaller absolute bias, who is expert 2 by convention. These observations are summarized in the next proposition.

Proposition 2. *For $i \in \{1,2\}$, an equilibrium exists in which the principal always chooses expert i and markups are $k_i = b_i$ and $k_j = -b_i$ for $j \neq i$. The principal's expected payoff in this equilibrium is $\mathbb{E}[-(s+b_i)^2] = -\sigma^2 - b_i^2$, where $s \sim N(0,\sigma^2)$.*

2.4 Symmetric Biases

In this section we characterize stationary equilibria in pure strategies of the delegation game for the case of symmetrically biased experts: $b_1 = b_2 = b > 0$. Although we analyze general biases in Section 2.5, we start with the symmetric case for two reasons. First, in many contexts, experts' biases are fairly similar, reflecting similar background and financial interests. Second, we can derive closed form expressions for equilibrium strategies and comparative statics in this case, which helps to provide intuition for the main qualitative results.

2.4.1 Min Equilibrium

We begin by investigating equilibria, in which the principal always chooses the minimum of the two offers: $\{(k_1, k_2, a \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geqslant 0\}$. We call these *min equilibria*. Here and throughout the rest of the paper, let f and F denote the PDF and the CDF of the distribution $N(0, 2\sigma^2)$ and let $z = k_1 - k_2$.

Using the fact that the expected state given signal realizations (s_1, s_2) is $\frac{s_1+s_2}{2}$, the Agent 1's expected payoff is

$$U_1(k_1,k_2,L) = \Pr(s_1+k_1 \leqslant s_2+k_2)\mathbb{E}\left[B - \left(s_1+k_1-\frac{s_1+s_2}{2}\right)^2 | s_1+k_1 \leqslant s_2+k_2\right]$$
$$+\Pr(s_1+k_1 > s_2+k_2)\mathbb{E}\left[-\left(s_2+k_2-\frac{s_1+s_2}{2}\right)^2 | s_1+k_1 > s_2+k_2\right],$$

which he must maximize over choices of k_1. Now consider the effects of a marginal increase in k_1. For cases where expert 1 has a strictly higher offer, there is no change in expert 1's utility since the implemented action of expert 2 remains unchanged.

For cases where expert 1 has a strictly lower offer, expert 1 is still selected, but the change in his action influences his quadratic loss. For the marginal case where $s_1 + k_1 \approx s_2 + k_2$, expert 1 goes from being selected to unselected, and thus loses the associated bonus B. However, the quadratic loss term is the same regardless of which expert is chosen, and so causing expert 2 to be chosen at the margin does not influence expert 1's loss term. In equilibrium, expert 1 is best responding, and thus the amount of foregone bonus must be offset by a helpful reduction in quadratic loss. Specifically, expert 1's first order condition is

$$k_1 = b + \left(\sigma^2 - \frac{B}{2}\right) \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)}. \tag{2.1}$$

Combining this with the first order condition for expert 2, we obtain a unique equilibrium candidate in which both experts choose a markup $k_m = b + \frac{\sigma^2 - \frac{B}{2}}{\sigma\sqrt{\pi}}$.[12] This expression combines two effects: a *winner's curse* effect and a *bonus-stealing* effect. To isolate the former, suppose for a moment that $B = 0$. Clearly $k_m > b$, that is, each expert's offer is above his optimal offer were he guaranteed to be chosen, as in the case of simple delegation. The reason is that, conditional on being selected, an expert learns that his rival had a sufficiently high signal and the expert would prefer to revise his offer upward. The expert's equilibrium markup must account for this. The bonus-stealing effect acts in the opposite direction and is proportional to the size of the bonus.

As the bonus increases, experts compete more aggressively by reducing their markups. For these markups to remain part of an equilibrium, the principal must be best-responding by choosing the lower offer, which requires that each markup is positive. Equivalently, the bonus must not be too large: $B \leqslant 2\sigma^2 + 2\sqrt{\pi}\sigma b$. In what

[12] Here and throughout we use subscript or superscript m to denote min equilibrium; we will use M for max equilibrium, introduced in the next section.

follows, it is useful to define $\rho := \sigma^2 - \frac{B}{2}$.

Proposition 3. *Suppose $b_1 = b_2 = b > 0$. Then a min equilibrium exists if and only if $B \leqslant 2\sigma^2 + 2\sqrt{\pi}\sigma b$. When it exists, it is unique, and the markups are $k_1^m = k_2^m = k_m = b + \frac{\rho}{\sigma\sqrt{\pi}}$.*

It is worth noting that if $B < 2\sigma^2$, the winner's curse effect outweighs the bonus-stealing effect and $k_m > b$, or in other words, experts make offers that exceed the sum of observed signal and bias (which would be the optimal proposal solely based on own information). In the opposite case the bonus-stealing effect dominates, with $k_m \leqslant b$.

2.4.2 Max Equilibrium

If the sum of markups is negative, then the principal's best response is to choose the maximum of the two offers. We now consider *max equilibria*. The forces of the previous section are reversed. Here, an expert's winner's curse is that he would have preferred a lower offer, since being selected (with the higher offer) indicates that the rival expert had a sufficiently low signal. On the other hand, experts compete for the bonus by raising their offers. In a max equilibrium, the experts' offers are $k_M = b - \frac{\sigma^2 - \frac{B}{2}}{\sigma\sqrt{\pi}}$. For the principal to best-respond by choosing the maximum offer, it must be that each markup is nonpositive, which again places an upper bound on the bonus: $B \leqslant B_M := 2\sigma^2 - 2\sqrt{\pi}\sigma b$.

Proposition 4. *Suppose $b_1 = b_2 = b > 0$. Then a max equilibrium exists if and only if $B \leqslant 2\sigma^2 - 2\sqrt{\pi}\sigma b$. When it exists, it is unique, and the markups are $k_1^M = k_2^M = k_M = b - \frac{\rho}{\sigma\sqrt{\pi}}$.*

By inspection, the upper bound on the bonus for a max equilibrium is strictly lower than that for a min equilibrium, and in fact, it can be negative. In this case, no max equilibrium exists. The reason is that the assumption $b_1 = b_2 = b > 0$ anchors markups at a positive value, and this works in the same direction as bonus-stealing in the case of a max equilibrium. Since markups must be negative, the sum of these forces must be outweighed by the winner's curse. If b is sufficiently large, then markups would be positive even with zero bonus, which is inconsistent with max equilibrium.

The max equilibrium, if exists, also involves a winner's curse effect and a bonus-stealing effect, but with opposite signs than in the case of the min equilibrium: the former induces experts to lower their offers while the latter to increase their offers.

2.4.3 Principal-Optimal Equilibrium

For the case of symmetric biases, the min equilibrium is optimal among the pure-strategy, stationary equilibria discussed so far. We make use of the quadratic loss utility form, noting that the principal's expected utility $\mathbb{E}(a-\theta)^2$ (net of the bonus paid, and conditional on θ) is equal to $-\bar{b}^2 - Var$ where $\bar{b} := \mathbb{E}(a-\theta)$ is the expected bias of the chosen offer and $Var := \text{Var}(a-\theta)$ is the variance of the bias. For the equilibria discussed, we have

- simple delegation: $\bar{b} = b$, $Var = \sigma^2$
- min equilibrium: $\bar{b} = b - \frac{B}{2\sqrt{\pi}\sigma}$, $Var = (1-\frac{1}{\pi})\sigma^2$
- max equilibrium: $\bar{b} = b + \frac{B}{2\sqrt{\pi}\sigma}$, $Var = (1-\frac{1}{\pi})\sigma^2$

Comparing min and max equilibria, we see that the variance of the bias is the same due to symmetry, while $\bar{b}$ has smaller magnitude for the min equilibrium because $b > 0$. Intuitively, the principal's ability to select the lower of the two offers offsets the inherent bias of the agents, whereas selecting the higher offer does not. Comparing simple delegation with the min equilibrium, we see that Var is lower for the latter.

The reason is that the inclusion of the second expert primarily is binding when he has the lower offer, which is more likely when the first expert's signal is high. This truncates the principal's loss in such situations and reduces variance. On the other hand, expected bias can have larger magnitude in the min equilibrium than under simple delegation, as can be seen easily in the case of $b = 0$, but the condition $B \leqslant B_m$ for the existence of a min equilibrium ensures any such disadvantage is out weighed by the variance reduction. These comparisons show that the min equilibrium is principal-optimal among pure-strategy, stationary equilibria.

Proposition 5. *Suppose* $b_1 = b_2 = b > 0$*. Then whenever the min equilibrium exists, the principal prefers it to simple delegation. If, in addition, the max equilibrium also exists, the principal prefers the min equilibrium to the max equilibrium.*

In Section 2.6.1, we show further that in the case of equal biases, the min equilibrium is optimal if the principal is given commitment power and has the ability to choose a mixed strategy, provided that the bonus is not too large. In particular, commitment by the principal results in the same outcome as in the min equilibrium of the game without commitment.

2.4.4 Comparative Statics

Given that the min equilibrium is optimal among pure strategy stationary equilibria, and optimal under commitment for a large range of bonus values, from here on we investigate comparative statics of the principal's expected payoff in the min equilibrium. We highlight three nonmonotonicities with respect to input parameters: b, σ and B.

First, the principal's payoff is nonmonotonic in the common bias b of the experts; in particular, for any $B \in [0, B_m]$ such that the min equilibrium exists, the principal-optimal level of expert bias is $b^* = \frac{B}{2\sqrt{\pi}\sigma}$, which is strictly positive if $B > 0$ and is strictly increasing in B. The intuition for why the principal can benefit from an increase in the experts' bias is that this can reduce the magnitude of the expected bias of the ultimately selected action (i.e., the difference between the implemented offer and the true state). When b is small, since the principal is choosing the lower of the two offers, this expected bias is negative, despite the fact that markups are positive;[13] by increasing b, this expected bias moves closer to 0, decreasing in magnitude. The larger is B, the more aggressively experts compete by lowering their markups, and the more negative the expected bias becomes; hence there is more room for the principal to benefit by increases in b.

Second, the principal's payoff is nonmonotonic in the noise level σ in the experts' signals. Increasing σ has two effects: it unambiguously increases the variance $(1 - \frac{1}{\pi})\sigma^2$ of the chosen offer, which hurts the principal, but it can help the principal by reducing the magnitude of the expected bias. Recall that for small b, the expected bias is negative. When σ increases the winner's curse becomes more severe, and moreover, the bonus-stealing incentive decreases since a marginal reduction in an expert's offer is less likely to be pivotal. Both of these forces increase the common equilibrium markup of the experts, so the (negative) expected bias moves closer to 0. On net, the reduction in magnitude of the expected bias works opposite the increase in variance, and for certain parameter values, the former can outweigh the latter. Lemma 55 in Appendix A provides the formal comparative statics of the principal's utility in the precision of the experts' signals.

Third, the principal's payoff is nonmonotonic in the bonus B, even when this bonus is provided externally (not from the principal's pocket). Formally, suppose that the game is preceded by a stage in which the principal publicly commits to a bonus level B, and that the principal's payoff is still $-(a - \theta)^2$. While the variance

[13] This is possible because of the limited choice set the principal has: it can be that the expectation of the minimum of the two offers has a negative bias and still preferred by the principal to the maximum, provided that the latter has a positive expected bias with a larger magnitude.

of the action $Var(k_m, k_m, L) = \sigma^2 - \frac{\sigma^2}{\pi}$ does not depend on B, the expected bias $b_m = b - \frac{B}{2\sqrt{\pi}\sigma}$ is decreasing in B. The principal prefers the expected bias to be as close to 0 as possible, so her expected payoff in the min equilibrium is maximized at $B = 2\sqrt{\pi}\sigma b$, where it is equal to $-\sigma^2[1 - \frac{1}{\pi}]$. At $B = 0$, $\bar{b} = b_m := b - \frac{B}{2\sqrt{\pi}\sigma} = b \geqslant 0$ and a small increase in bonus decreases the experts' markup and benefits the principal. However, at $B = B_m$ $b_m = b - \frac{B}{2\sqrt{\pi}\sigma} = -\frac{\sigma}{\sqrt{\pi}} < 0$ and principal prefers to increase markups and correspondingly lower bonus. As a consequence, an intermediate point $B = 2\sqrt{\pi}\sigma b$ is optimal.

The principal's payoff is also non-monotonic in the bonus B when it is paid by the principal. While increasing the bonus has a direct cost to the principal, it increases the bonus-stealing incentive which reduces markups. This reduces the magnitude of the expected bias when the expected bias is positive. A simple condition determines when a positive B is optimal. Recall that the principal's payoff is $-B - b_m^2 - Var$. Starting from $B = 0$, an increase in B reduces b_m^2 at a rate $\frac{b}{\sqrt{\pi}\sigma}$, which is increasing in b due to quadratic loss and which is decreasing in σ since the bonus-stealing incentive is weaker for larger σ. The net change in the principal's payoff is positive if and only if $\frac{b}{\sqrt{\pi}\sigma} > 1$.

Further look at the relationship between the principal's payoff and the bonus. The principal chooses a strictly positive bonus if the decrease in expected bias exceeds the marginal disutility from increasing bonus: $\frac{b}{\sqrt{\pi}\sigma} > 1$. This holds for $b = 2$, $\sigma = 0.5$, but not for $b = 1$, $\sigma = 1$.

Proposition 6 summarizes the various nonmonotonicities discussed above.

Proposition 6. *Suppose $b_1 = b_2 = b > 0$. The principal's expected payoff in the min equilibrium is nonmonotonic in b and σ. It is also nonmonotonic in B, both when it is paid from exogenous sources and when it is paid by the principal.*

2.5 General Biases

Many of the economic insights obtained in the case of like biases extend to general biases. Below we show what results continue to hold, but focus on new insights.

2.5.1 Characterizing Min and Max Equilibria For General Biases

First we generalize the characterizations of the min and max equilibria. Recall the definition $\rho := \sigma^2 - \frac{B}{2}$. The first order conditions for the players are

$$k_1 = b_1 + \rho \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} \tag{2.2}$$

$$k_2 = b_2 + \rho \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)}. \tag{2.3}$$

Subtracting these yields

$$z - \rho\left[v(z) - w(z)\right] = b_1 - b_2. \tag{2.4}$$

In the proof of Theorem 7, there is a unique solution to (2.4) which we denote z^*.

Theorem 7. *There exists a threshold $B_m > 0$ such that a min equilibrium exists if and only if $B \leqslant B_m$. When it exists, it is unique and characterized by markups $k_1^m = b_1 + \rho v(z^*)$ and $k_2^m = b_2 + \rho w(z^*)$ with $k_1^m - k_2^m = z^* \geqslant 0$. Moreover, B_m lies in the interval $\left[2\sigma^2 + 2\sqrt{\pi}\sigma \max(0, b_2), 2\sigma^2 + \sqrt{\pi}\sigma(b_1 + b_2)\right]$. For $B \leqslant B_m$,*

- $z^* \geqslant b_1 - b_2 \iff B \leqslant 2\sigma^2 \iff \rho \geqslant 0$;
- $\bar{b}(k_1^m, k_2^m, L) = b_m := b_1(1 - F(z^*)) + b_2 F(z^*) - B f(z^*)$;
- $Var(k_1^m, k_2^m, L) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*))$.

The next theorem characterizes max equilibrium. Note that if $B_M < 0$, no max equilibrium exists. Recall the definition of z^* from (2.4).

Theorem 8. *There exists a threshold B_M such that a max equilibrium exists if and only if $B \leqslant B_M$. When it exists, it is unique and characterized by $k_1^M = b_1 - \rho w(z^*)$ and $k_2^M = b_2 - \rho v(z^*)$, with $k_1^M - k_2^M = z^* \geqslant b_1 - b_2$ and $B_M \in [2\sigma^2 - \sqrt{\pi}\sigma(b_1 + b_2), 2\sigma^2 - 2\sqrt{\pi}\sigma \max(0, b_2)]$. For $B \leqslant B_M$,*

- $\bar{b}(k_1^M, k_2^M, H) = b_M = b_1 F(z^*) + b_2(1 - F(z^*)) + Bf(z^*)$;
- $Var(k_1^M, k_2^M, H) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*))$.

The following corollaries are immediate from Theorems 7 and 8.

Corollary 9. *A max equilibrium exists only if a min equilibrium exists; that is, $B_M \leqslant B_m$.*

Corollary 10. *For $B \leqslant B_M$, $k_1^m - b_1 = b_2 - k_2^M$ and $k_2^m - b_2 = b_1 - k_1^M$.*

Lemma 11. *If both experts follow constant markup strategies $a_j(s_j) = s_j + k_j$ and the principal always chooses the lower offer, then*

$$\bar{b}(k_1, k_2, L) = -2\sigma^2 f(z) + k_1(1 - F(z)) + k_2 F(z);$$

$$Var(k_1, k_2, L) = \sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z));$$

$$V(k_1, k_2, L) = -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(z) - k_1^2(1 - F(z)) - k_2^2 F(z);$$

$$U_i(k_1, k_2, L) = -\sigma^2 + 2\sigma^2(k_i + k_j - 2b_i)f(z) - (k_j - b_i)^2 F(k_i - k_j) - (k_i - b_i)^2 F(k_j - k_i).$$

Notice that the equilibrium markup difference z^* as well as $\mathrm{Var}(a - \theta)$ depend on the biases of the experts only through $b_1 - b_2$. In Appendix A (Corollary A.1.1), we give the full expansion of the players' utilities in the min equilibrium.

Lemma 12. *If both experts follow constant markup strategies* $a_j(s_j) = s_j + k_j$ *and the principal always chooses the higher offer, then*

$$\bar{b}(k_1, k_2, H) = k_1 + k_2 - \bar{b}(k_1, k_2, L);$$

$$Var(k_1, k_2, H) = Var(k_1, k_2, L);$$

$$V(k_1, k_2, H) = V(-k_1, -k_2, L);$$

$$U_i(k_1, k_2, H) = -\sigma^2 - 2\sigma^2(k_i + k_j - 2b_i)f(z) - (k_i - b_i)^2 F(k_i - k_j) - (k_j - b_i)^2 F(k_j - k_i).$$

In the max equilibrium, the principal's strategy of choosing the higher offer implies that, conditional on being chosen, an expert must revise his belief and his markup downward. This downward force must be sufficiently large to ensure that the sum of markups is negative, so that the principal's choice of the higher offer is a best response. Hence, noise must be sufficiently large for the max equilibrium to exist.

2.5.2 Principal-Optimal Equilibrium

In this section we compare the principal's expected utility in the min and max equilibria, and in the case of simple delegation. We also investigate how the principal's expected payoff in equilibrium depends on the biases of the experts. The main finding is that the min equilibrium remains always preferred by the principal to the max equilibrium, and for a large range of parameter values (including the focal cases of equal biases, opposite biases, or one bias equal to zero) the min equilibrium outperforms simple delegation as well.

Note that the equilibrium markup difference z^* and $\mathrm{Var}(a - \theta)$ stay the same as in the min equilibrium. In the max equilibrium expert 1's offer is chosen with probability $F(z^*)$, which is higher than $1 - F(z^*)$ in the min equilibrium, and this shifts the expected bias higher. As a result, the expected bias is higher in the max equilibrium, making it inferior to min equilibrium for the principal. In Corollary A.1.2 we provide expressions for the players' utilities.

First we establish that whenever both min and max equilibria exist, the principal always prefers the former.

Proposition 13. *For any fixed $B \geqslant 0$, the principal prefers min equilibrium to max equilibrium whenever both exist.*

The intuition for the above result can be summarized as follows. As we pointed out earlier, in any state θ, the variances of the expected offer $\mathrm{Var}(a-\theta)$ in the min and max equilibria coincide, but the expected bias $\mathbb{E}(a-\theta)$ is higher in the max equilibrium. Furthermore, in the max equilibrium the bonus motivates experts to increase their markups, as opposed to min equilibrium in which the bonus motivates experts to decrease their markups: $b_M = b_1F(z^*) + b_2(1-F(z^*)) + Bf(z^*) \geqslant b_1(1-F(z^*)) + b_2F(z^*) - Bf(z^*) = b_m$. Hence, to conclude that the principal is better off in the min equilibrium it is enough to show that $|b_M| \geqslant |b_m|$ or, taking into account the above, $b_m + b_M \geqslant 0$. Markup differences in the min and max equilibria coincide, but the choice rule is opposite, therefore for both experts the probabilities of winning in the min and max equilibria are complementary. The direct effects of the bonus on b_m and b_M are opposite and equal in absolute value. Therefore $b_m + b_M = b_1 + b_2 \geqslant 0$, implying that the principal is better off in the min equilibrium.

We can further argue that in case the two experts are not equally biased, the expert with the lower bias also prefers min equilibrium to max equilibrium, while the expert with the higher bias has the opposite preferences. This is both because the expected action is closer to expert 2's ideal point in the min equilibrium, and closer to expert 1's ideal point in max equilibrium, and because expert 2 is chosen (and hence receives the bonus) with higher probability in the min equilibrium, and expert 1 is chosen with higher probability in the max equilibrium.

Given the above result, we now compare the principal's utility in the min equilibrium and simple delegation to the expert with the smaller absolute bias. In general this comparison is complicated, but in the next proposition we show that, for any positive difference between the experts' biases, the principal prefers simple delegation to min equilibrium if the mean of the two biases is sufficiently high. On the other hand, when one of the experts is unbiased, the principal prefers the min equilibrium to simple delegation.

Proposition 14. *Assume $B = 0$, and parameterize biases as $b_1 = b+x$ and $b_2 = b-x$ for some $x, b \geqslant 0$. For all $x > 0$, there exists a threshold $b^* > 0$ such that the principal prefers simple delegation to the min equilibrium if and only if $b > b^*$. In addition, the principal always prefers min equilibrium when $b_2 = 0$.*

In particular, Proposition 14 says that for opposite biases (parametrized by $x > 0$, $b = 0$) the principal prefers the min equilibrium. Note that the case $x = 0$ (equal biases) was treated in Proposition 5, where we showed that for any common bias b, and any bonus B such that the min equilibrium exists, the principal prefers the min equilibrium to simple delegation.

For intuition behind the above result, recall that the variance of the chosen action in the min equilibrium is always lower than that in simple delegation, and both are independent of b. As b increases, the expected bias $b - (2F(z^*) - 1)x$ in the min equilibrium and $b - x$ in the max equilibrium increase at the same rate. Since the former is higher than the latter, and losses are quadratic, this increase hurts more in the min equilibrium than in simple delegation. Once b is high enough, this disadvantage outweighs the initial advantage of lower variance.

Figure 2.2 illustrates the comparison between min equilibrium and simple delegation for the case $B = 0$, $\sigma = 1$. In the larger, unshaded region between the dashed lines, the principal prefers the min equilibrium.

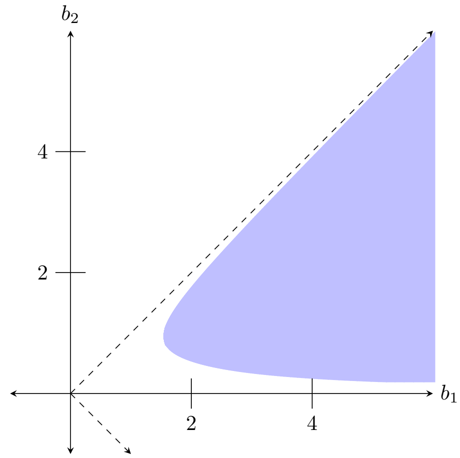

FIGURE 2.2: The principal-optimal equilibrium depending on the biases (b_1, b_2)

2.5.3 *Equal vs. Opposite Biases*

Here we compare the expected payoff the principal can achieve with two equally biased experts, $b_1 = b_2 = b$, to the expected payoff she can achieve with two oppositely biased experts, $b_1 = -b_2 = b$. Theorem 7 implies that the expected bias of the action in the case of equally biased experts is equal to b. In the case of oppositely biased experts the expected bias of the action is $b(1 - 2F(z^*)) - Bf(z^*)$, simplifying to $b(1 - 2F(z^*))$ when $B = 0$. Hence, with $B = 0$ the absolute value of the expected bias in the case of oppositely biased experts is lower than in the case of equally biased experts. However, the next proposition shows that the variance of the action is lower in the symmetric case, and in fact this effect dominates, resulting in the principal preferring to have two equally biased experts. Let $V_{symm}(b)$ be the principal's expected payoff in the min equilibrium when $b_1 = b_2 = b$, let $V_{opp}(b)$ be the principal's expected payoff in the min equilibrium when $b_1 = -b_2 = b$, and let $V_{sim}(b) = -\sigma^2 - b^2$ be the principal's expected payoff in the case of simple delegation to an expert with absolute bias b.

Proposition 15. *For any $b > 0$, $V_{symm}(b) \geqslant V_{opp}(b) \geqslant V_{sim}(b)$.*

Proof. Min equilibrium for opposite biases exists only if $B \leqslant B_m = 2\sigma^2$. From

Theorem 7

$$V_{symm}(b) = -\sigma^2 - b^2 + 2Bbf(0) + (4\sigma^4 - B^2)f^2(0)$$

$$V_{opp}(b) = -\sigma^2 - b^2 + \left(\sigma^4 - \frac{B^2}{4}\right)\frac{f^2(z^*_{opp})}{F(z^*_{opp})(1 - F(z^*_{opp}))} \geqslant -\sigma^2 - b^2 = V_{sim}(b),$$

where z^*_{opp} is the min equilibrium markup difference in the case of oppositely biased experts. As $2Bbf(0) \geqslant 0$ and $\frac{f^2(z)}{F(z)(1-F(z))}$ reaches its maximum at $z = 0$, we get $V_{symm}(b) \geqslant V_{opp}(b)$ □

2.6 Extensions

In this section, we consider two extensions of the model.

2.6.1 Commitment

We now extend the model to allow the principal to publicly commit to a mixed strategy. In particular, suppose that the principal can credibly commit to choose the lower offer with any probability $p \in [0, 1]$ and the higher offer with probability $1 - p$.

When experts have equal biases and the bonus is not too large, commitment power does not help the principal; the principal would choose $p = 1$, which is already a min equilibrium without commitment. On the other hand, in some instances, the principal benefits from introducing an "element of surprise" and optimally commits to a mixed strategy. To illustrate this point, it suffices to consider oppositely biased experts and a bonus of $B = 0$. As the common magnitude b of the biases increases, expert 1's winner's curse in the min equilibrium becomes more severe, and his markup very large. By introducing a small probability of choosing the higher offer, expert 1 is incentivized to reduce his markup, benefiting the principal. The cost to the principal of doing so is in choosing the wrong offer. When b is large, the markup reduction is large and outweighs the cost, and thus the principal can profitably deviate from $p = 1$ to an interior p.

If the principal commits to choosing the lower offer with probability p, then by Lemma 57 in Appendix A, whenever $B \leqslant 2\sigma^2$, there exists a unique equilibrium of the game between the two experts, for any pair of expert biases. In the case of $b_1 = b_2 = b$, the unique equilibrium given the pre-committed strategy of the principal involves $k_1 = k_2 = k = b + (2p-1)\frac{2\sigma^2 - B}{2\sqrt{\pi}\sigma}$. The principal's utility is thus

$$
\begin{aligned}
V &= p\left[-\sigma^2 - k^2 + \frac{2\sigma}{\sqrt{\pi}}k\right] + (1-p)\left[-\sigma^2 - k^2 - \frac{2\sigma}{\sqrt{\pi}}k\right] \\
&= \frac{\rho(2\sigma^2 + B)}{2\pi\sigma^2}(2p-1)^2 + \frac{bB}{\sqrt{\pi}\sigma}(2p-1) - \sigma^2 - b^2.
\end{aligned}
$$

The first part of Proposition 16 follows from inspection of the above. The proof of the second part of the proposition is more detailed and is deferred to Appendix A.

Proposition 16. *For all $B \leqslant 2\sigma^2$ and $b_1 = b_2 = b > 0$, $p = 1$ is the optimal strategy of the principal under commitment. For $B = 0$, $b_1 = b > 0$ and $b_2 = -b$, if b is sufficiently large, the optimal p under commitment satisfies $p \in (0, 1)$.*

2.6.2 Unselected Expert Indifferent over Actions

In the baseline model we assumed that an expert whose offer is not selected is still affected by principal's action. While this is a reasonable assumption in some contexts, in other situations it is more realistic to assume that the expert not selected by the principal receives an outside payoff that is independent of the state and the implemented action. For instance, a car mechanic is unlikely to care what kind of maintenance is done if he is not the one selected for the job. In this extension we assume that if expert i is chosen then expert j's realized payoff is normalized to be 0. We restrict attention the case of equally biased experts: $b_1 = b_2 = b > 0$.

In this version of the model we assume that B is large enough that an expert's expected payoff under simple delegation to that expert is nonnegative; that is, he prefers simple delegation to not being selected at all. Under this condition, the same simple delegation equilibria exist in this version of the model as in the baseline model. Below we show that under some conditions there also exist symmetric pure strategy equilibria that are similar to the ones characterized in the baseline model. For this to be the case, the bonus payment must be neither too low nor too large.

First we examine the conditions for the existence of the min equilibrium. Using the same notation as before, we investigate strategy profiles $\{(k_1, k_2, C(a_1, a_2) \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geqslant 0\}$. While for any such profile the principal's payoff does not change, the experts' expected payoffs should be recalculated:

$$U_i(k_i, k_j, L) = \int_{k_i - k_j}^{\infty} \left[B - \left(k_i - b - \frac{t}{2} \right)^2 - \frac{\sigma^2}{2} \right] f(t)\, dt$$

$$= \left[B - \sigma^2 - (k_i - b)^2\right] (1 - F(k_i - k_j)) + \left[2\sigma^2 (k_i - b) - \frac{1}{2}\sigma^2 (k_i - k_j) \right] f(k_i - k_j).$$

Notice that for a fixed constant markup strategy of the other expert, an expert can choose arbitrarily high constant markup and guarantee an expected payoff arbitrarily close to 0.[14] Hence, 0 is a lower bound for experts' equilibrium payoffs.

In what follows, define $\beta := \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$, $B_1 := \left(\frac{5}{2} - \frac{3\pi(8\pi - 11)}{16(\pi-1)^2} \right) \sigma^2$ and $B_2 := \frac{5}{2}\sigma^2 + 2\sqrt{\pi} b \sigma + b^2$.

Proposition 17. *A symmetric min equilibrium $k_1^m = k_2^m = k_m$ exists if and only if $B \in [B_1, B_2]$. When it exists, it is characterized by:*

- $k_1^m = k_2^m = k_m = b + \left(\sqrt{\pi} - \beta\right) \sigma$;
- $\bar{b}(k_m, k_m, L) = b + \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}} \right) \sigma$;
- $Var(k_m, k_m, L) = \left(1 - \frac{1}{\pi}\right) \sigma^2$;

[14] For this reason, we do not introduce an explicit participation constraint in this version of the model, even though such a constraint would be natural in many applications.

- $V(k_m, k_m, L) = -\left[b + \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}}\right)\sigma\right]^2 - \sigma^2 + \frac{\sigma^2}{\pi}$;

- $U_i(k_m, k_m, L) = \left[\frac{\pi-1}{\sqrt{\pi}}\beta + \frac{7}{4} - \pi\right]\sigma^2$ *for* $i = 1, 2$.

In Appendix A we show that $k_m = b + (\sqrt{\pi} - \beta)\sigma > k_m^{bas.} = b + (1 - \frac{B}{2\sigma^2})\frac{\sigma}{\sqrt{\pi}}$, hence in this version of the model experts select higher markups in the min equilibrium than in the baseline model (for parameter values for which min equilibrium exists in both model versions). The intuition behind this result is that in this alternative version of the model, the relative gain from being selected is reduced by the policy loss (that is not imposed on the expert if not selected). Since we consider B sufficiently large that expected payoffs are nonnegative, the resulting "net bonus" is still nonnegative; being selected is still preferable, conditional on having made the lower offer. It follows that an expert's equilibrium offer in either version is lower than what is ex-post optimal for that expert – that is, optimal after conditioning on both the expert's signal and having the lower offer. The smaller net bonus in the alternative version reduces the expert's incentive to marginally lower his offer in order to more frequently earn the net bonus. This reduction must be met by an offsetting reduction in his incentive to raise his offer, which is enforced through his bidding higher and thus closer to his ex-post optimum; due to quadratic losses, marginal movements toward the ex-post optimum have decreasing marginal benefits.

We note that the qualitative comparison between this extension and the baseline model is dependent upon the modeling of preferences over policy outcomes through losses. Such a model is appropriate for applications where an expert would prefer not to be associated with the project if his action would be sufficiently far from the true state; for example, this would be the case if the expert has a reputation at stake. Alternatively, one could model preferences through gains, using some single-peaked, nonnegative utility function of the distance between the action and the true state.

In that model, the comparison above would be reversed, as being selected enhances the bonus and thus experts compete more aggressively by lowering their offers.

Next we turn attention to characterizing the conditions under which a max equilibrium exists in which the principal always chooses the higher offer. Experts' expected payoffs can be calculated as:

$$U_i(k_i, k_j, H) = \int_{-\infty}^{k_i - k_j} \left[B - \left(k_i - b - \frac{t}{2} \right)^2 - \frac{\sigma^2}{2} \right] f(t)\, dt$$

$$= \left[B - \sigma^2 - (k_i - b)^2 \right] F(k_i - k_j) - \left[2\sigma^2 (k_i - b) - \frac{1}{2}\sigma^2 (k_i - k_j) \right] f(k_i - k_j)$$

As in the min equilibrium, 0 is a lower bound for experts' equilibrium payoffs.

Proposition 18. *Consider $b_1 = b_2 = b > 0$. If $\frac{b}{\sigma} > \frac{3\sqrt{\pi}}{4(\pi - 1)}$, then no symmetric max equilibrium exists. If $\frac{b}{\sigma} \leqslant \frac{3\sqrt{\pi}}{4(\pi - 1)}$, then a symmetric max equilibrium $k_1^M = k_2^M = k_M$ exists if and only if $B \in [B_1, B_2]$. When it exists, it is characterized by:*

- $k_1^M = k_2^M = k_M = b - (\sqrt{\pi} - \beta)\,\sigma$;
- $\bar{b}(k_M, k_M, H) = b - \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}} \right) \sigma$;
- $Var(k_M, k_M, H) = (1 - \frac{1}{\pi})\sigma^2$;
- $V(k_M, k_M, H) = -\left[b - \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}} \right) \sigma \right]^2 - \sigma^2 + \frac{\sigma^2}{\pi}$;
- $U_i(k_M, k_M, H) = \left[\frac{\pi - 1}{\sqrt{\pi}} \beta + \frac{7}{4} - \pi \right] \sigma^2$ *for* $i = 1, 2$.

For the max equilibrium, the difference relative to the baseline model is the mirror image of the difference described earlier for the min equilibrium. Again, the bonus is reduced by quadratic losses, but in the max equilibrium this causes markups to decrease, as experts compete less aggressively to make the higher offer.

2.7 Conclusion

We proposed a model in which a principal can choose between two imperfectly informed experts, introducing the possibility of competition in a delegation framework. We showed that a principal with limited knowledge of the decision environment can benefit from the presence of two experts, relative to a simple unconstrained delegation to one of them, even if the experts have exactly the same bias. The main reason is that in equilibria in which the selection of the expert depends nontrivially on the experts' proposals, information is utilized from both experts' private signals. The option of offering a bonus payment to the selected expert can improve the principal's payoff, by inducing the experts to report more truthfully, but only to a certain point. Lastly, committing with a small probability to choose the (in expectation) inferior proposal can benefit the principal.

As this is the first step in investigating the benefits of multiple choices of experts in a delegation problem, there are many avenues of future research. One is examining multi-dimensional environments, in which different experts differ in their dimensions of specialization. Another direction would be investigating the problem of choosing an expert to delegate a task to with a more general mechanism design approach.

3

Optimal Auctions with Ex-Post Verification and Limited Punishments (with Sergii Golovko)

3.1 Introduction

During the last few decades, auctions have been extensively used by governments for privatization. Examples include the sale of large enterprises in eastern and western Europe, the licensing of the electromagnetic spectrum in the United States, and the sale of construction contracts. In each of these auctions, the relationship between the seller and the buyer extends beyond the auction setting, as a privatized firm pays taxes, is monitored by the government, and is affected by ongoing industry regulations. As a result, the government may gain access to information regarding the buyer's true valuation. If the government realizes ex-post that the price paid in the auction was too low, it can use its regulatory power to penalize the buyer.

For example, in 2004, the Ukrainian government sold a major steel producer, Kryvorizhstal, to a consortium led by the president's son-in-law for $800 million. This price was considered to be too low, and the incident was used as a stunning example of corruption during the following presidential election. The new government deemed the instance of privatization illegal and canceled the sale, and Kryvorizhstal was eventually resold for $4.81 billion. The New York Times called the deal "Ukraine's biggest and most profitable privatization auction."[1]

This paper addresses situations in which the seller's ability to punish bidders is limited. For instance, in repeated procurement auctions, if the government receives paperwork from a bidder with a history of previous negative interactions, it may disqualify the bidder. Our results indicate that even if the probability of this occurring is quite small, its effect on the seller's expected revenue is significant. Therefore, knowing how to construct an optimal mechanism in environments with ex-post verification and limited punishments may be quite important in practice.

In our model, a profit-seeking seller uses an auction to allocate an indivisible object to one of several ex-ante identical buyers. Prior to the auction, the seller commits to both an auction format and a penalty rule as a function of the winning bid and the winner's valuation. We assume that after the auction, the winner's value can be precisely verified through, for example, observing the revenue stream. Once the winner's value is established, the seller is able to impose a punishment and the winner suffers a corresponding loss. Only a part of this punishment contributes to the seller's revenue, however. In this way, we incorporate both monetary and non-monetary punishment possibilities, as well as the costliness of imposing a punishment.

In this paper, when considering winner-pay auctions, we find the optimal selling mechanism, which includes both an auction format and a penalty rule, and study the efficiency of an optimal auction. The latter is important in practice. If the reserve price is set too high, the government may fail to sell the object. If the price is too low, however, the power can face critics from the public once the true value of object is verified. As the Economist notes in discussing the privatization of the Royal Mail though the IPO: "The Royal Mail sale was a reminder of the political risks: price

[1] The New York Times, March 15th, 2006.

an asset too high and the deal might flop; price it too low and the taxpayer feels cheated." [2]

We start our analysis by studying a first-price auction. Our first observation, which extends to other auction formats as well, is that the problem of finding an optimal penalty rule can be reduced to the design of an optimal type-dependent recommendation. Before an auction starts, the seller announces how much he expects each type of buyer to bid. The recommendation is anonymous, as it depends on the buyer's valuation rather than his identity. If the winner follows the seller's recommendation, he is not penalized (or penalized a lesser amount); otherwise, he can be penalized the maximum amount, provided he is the winner. Our second observation is that unlike in standard auctions, equilibrium recommendations, or the ones from which bidders do not find it profitable to deviate, can be non-monotonic: that is, a bidder with a high valuation may be prescribed to bid less than a bidder with a lower valuation.

Nevertheless, we show that the seller's optimal recommendation is monotonic. The penalty is used purely as a threat—the buyer is never punished, provided he follows the recommendation. The standard methods that prescribe using differential equations to derive the symmetric equilibrium are not applied in our environment. Instead, we use the following recursive procedure to find the optimal recommendation for a first-price auction. We divide the type space into finitely many intervals. If the buyer's type falls in the very left interval, the buyer is prescribed to bid his value. If the buyer's type is in the next interval, then the buyer's bid is determined in order to make him indifferent between bidding according to the recommendation without being punished and mimicking one of the buyer's types in the previous interval while being punished at the maximum amount. Similarly, for each consequent interval, the buyer is indifferent between following the recommendation and mimicking one

[2] The Economist, January 9th 2014.

of the buyer's types from all previous intervals at the cost of being punished. Using the recursive structure of equilibrium, we show that buyers bid more in this context than in a standard (without ex-post penalties) first-price auction.

We further investigate how the efficiency of a first-price auction compares to that of a standard first-price auction. In our setting, the optimal reserve price is lower than it is in a standard environment. In a standard first-price auction, lowering the reserve price has two effects on the seller's revenue. First, a larger set of bidder types submit bids above the reserve price, which increases the probability of selling the object. Second, buyers adjust their behaviour by bidding less aggressively, which lowers the price, conditional on a sale. At the optimal reserve price, these two effects offset each other. In our model, the same effects are in play. However, the second effect is smaller, and it is completely absent for buyer types that are just above the reserve price (for these types, the seller effectively imposes penalties in order to extract the full surplus). Hence, the total seller's benefits from reducing the reserve price are greater than they are in the standard environment, and the two effects equalize at a lower reserve price.

By focusing our attention on monotonic recommendations, we are able to show that the first-price auction is the optimal auction format among all winner-pay auctions. For any other auction format, the seller is able to raise the same revenue by using a first-price auction instead and recommending the buyer to bid the amount a buyer of the same type would pay in expectation, conditional on winning in that auction format. The reverse of this—that there exists a monotonic bidding function of the other auction format that guarantees the same expected payment as in a first-price auction—is not always true. Perhaps surprisingly, it does not hold in a second-price auction.

While we derive the general conditions necessary for revenue equivalence to hold in the main text, a simple example of the discrete version of our model illustrates its failure. A single object is sold to one of two buyers. The values are drawn independently from the uniform distribution $V = \{2, 4, 4.9\}$. The tie is resolved by a flip of a fair coin. The seller can punish up to 60% of the value,[3] the penalty is non-monetary, and the penalty does not benefit the seller. If the object is sold through a first-price auction, the seller can extract the full surplus by recommending that each buyer bids his value. Provided his opponent follows this recommendation, a buyer does not have an incentive to deviate: If he wins the object by bidding differently from his value, then he is punished by 60% of his value that when added to 2 (the minimum bid required to be a winner), is larger than his value. Suppose now that the object is sold through a second-price auction. If there is a monotonic recommendation that allows the seller to extract the full surplus, then the buyer with the lowest valuation 2 should bid 2, and the buyer with valuation 4 should be prescribed to bid 8 (with a probability of 1/3 that she wins against an opponent with a valuation of 2 and with a probability of $1/2 \cdot 1/3 = 1/6$ that she wins against an opponent with a valuation of 4, making the expected payment, conditional on winning, equal to $(1/3 \cdot 2 + 1/6 \cdot 8)/(1/3 + 1/6) = 4$). However, the bidder with a valuation of 4.9 never finds it optimal to bid more than 8, since in that case his expected payment, conditional on winning, is greater than $(1/3 \cdot 2 + 1/3 \cdot 8)/(1/3 + 1/3) = 5$. We thus conclude that there is no symmetric monotonic equilibrium that allows the seller to extract the full surplus.

Finally, we apply the theory developed in this paper to analyze a problem of optimal taxation within a newly privatized firm. We limit taxation policies to the ones in which the future firm's profits are taxed at a flat rate that does not exceed a certain bound. We view this bound as one that is determined by law and cannot be increased. Since taxation plays the role of a limited punishment in our model, our results indicate that it should be used only as a threat to induce buyers to bid more.

[3] All our arguments go through if the seller can punish by a greater percentage than 60.

As a consequence, the government gains by credibly promising to soften taxation in the case of "fair" bidding.

RELATED LITERATURE.—In a special case of our model, when the punishments are non-monetary and are not part of the seller's revenue, our structure of verification and punishment is the same as in Mylovanov and Zapechelnyuk (2017). They study an allocation problem in an environment without monetary transfers and in which punishments are the only instruments that allow a principal to elicit truthful information from agents. As a result, the optimal allocation rule is stochastic—the agents are shortlisted with probabilities that depend on each agent type, and the object is randomly allocated to a person from the shortlist. They present the surprising finding that the principal benefits from restricting participation. In contrast, in our model, monetary transfers are allowed, and the optimal mechanism is deterministic—an object is allocated to the buyer who assigns the greatest value to it. Therefore, likewise in standard auctions, the seller benefits from more buyers participating in the auction.

In another special case of our model, when the penalty is a pure monetary transfer to the seller, our auctions are closely related to those in contingent payments literature. McAfee and McMillan (1986), in the context of procurement auctions, study linear contracts in the submitted bid and ex-post realized cost of a project, where this cost is subject to the moral hazard problem. Despite the latter, they find that an optimal linear contact is contingent on observed costs and brings in more revenue than a standard first-price auction.

In the same spirit, Hansen (1985) shows that if a seller uses auctions in which bids are made in stock or profit shares, she receives a higher expected payoff than in cash auctions. In our model, the punishment is contingent on the winner's value, and hence, as in the aforementioned two papers, the seller's revenue is larger than it is in a standard auction. In his comment to Hansen (1985), Crémer (1987) goes further and argues that if a seller is allowed to fully compensate a winner in cash, he can achieve full control of the merging firm and extract the winner's full valuation. A similar result holds in our model when the seller can apply a sufficiently high punishment. DeMarzo et al. (2005) [4] points out that if the seller is cash-constrained or initial investment is not fully contractible then Crémer's argument is not valid. They show that the seller's expected revenue is increasing in steepness of securities and maximized for call options. Our penalty design looks similar to a call, but with its exercise value increasing in winner's bid (the winner is penalized if his observed value is higher than expected based on his bid). As with classical calls in DeMarzo et al. (2005), in our setting, the first-price auction format is superior; however, here in equilibrium, it is used purely as a threat.

Our paper also relates to literature on optimal auction design. The central finding in this literature is the revenue-equivalence principle, which states that under certain conditions, two auction formats generate the same revenue. This was first observed by Vickrey (1962) in the case of first- and second-price auctions, and later generalized by Riley and Samuelson (1981) and Myerson (1981) to incorporate other auction formats as well.

The literature has identified multiple reasons for the failure of revenue equivalence by relaxing one of the assumptions in Riley and Samuelson (1981) and Myerson (1981). Thus, Milgrom and Weber (1982) study first- and second-price auctions when buyers have only partial information about their valuations. Under the assumption of interdependent values and affiliated signals, they show that a second-price auction is revenue-superior to a first-price auction. In contrast, a first-price auction is the preferable auction format when buyers are risk-averse (see, for instance, Holt Jr

[4] Skrzypacz (2013) use DeMarzo et al. (2005) as a baseline model and survey recent related literature.

(1980), Matthews (1979), and Maskin and Riley (1984)). Che and Gale (1998) derive the same result under the assumption that buyers are financially constrained. In our environment, buyers bid more than their values in the second-price auction and therefore, the same ranking holds when buyers' budgets are equal to their valuations. We suspect that in our model, this result extends to include cases in which buyers have budget constraints that may exceed their valuations and, as in the work of Che and Gale (1998), budgets are buyers' private information.

Marshall and Marx (2007), by studying buyers' behaviour in cartels that are not all-inclusive, show that a second-price auction is more susceptible to collusion than a first-price auction. Their results are in sharp contrast with previous findings[5] that all-inclusive cartels can suppress all buyers' competition, and therefore, the revenue equivalence of first- and second-price auctions continues to hold. If we allow for the possibility of collusion in our model, all-inclusive cartels may successfully escape a punishment and buy the object at the lowest price in a second-price auction. The latter is not true for a first-price auction, and therefore, similarly to Marshall and Marx (2007), we find that a first-price auction is revenue-superior.

This paper also presents its own distinct reason for the failure of revenue equivalence. Unlike in a standard second-price auction, in a second-price auction studied in this paper, bidding one's own value is not the dominant strategy. Moreover, the existence of the monotonic equilibrium that provides the same revenue to the seller as in the optimal first-price auction is not guaranteed. This causes a first-price auction format to be superior whenever such monotonic equilibrium does not exist in a second-price auction.

[5] Bidders' behaviour in all-inclusive cartels for first-price auctions has been studied by McAfee and McMillan (1992), while a similar analysis for second-price auctions has been done by Graham and Marshall (1987) and Mailath and Zemsky (1991).

3.2 The Model

A seller owns a single indivisible object and values it at zero. She uses an auction to allocate this object to one of n potential buyers. Each of them knows his exact value for the object, which is private information. It is common knowledge that bidders' values v_i are independently and identically distributed on $V = [v_L, v_H]$ with cumulative distribution function F, whose density f is continuous and positive on V. All parties, including the seller, are expected profit-maximizers. We assume that bidders are not subject to any liquidity or budget constraints. (A violation of the latter assumption is briefly discussed in Section 3.4.)

We focus our attention on the winner-pay auction, or the type of auction in which the highest bidder (if two bidders are tied, the tie is resolved by the flip of a fair coin) wins the object if his offer meets a reserve price r and he is the only one who makes a payment to the seller.[6] Each winner-pay auction is completely defined by its winner's payment function $M : \mathbb{R}_+^n \to \mathbb{R}_+$. We assume that M satisfies the following assumptions:

Assumption 1 (Anonymity)**.** $M(\mathbf{b})$ *is invariant to permutations of* $\mathbf{b}$.

Assumption 2 (Continuity)**.** $M(\mathbf{b})$ *is continuous.*

Assumption 3 (Monotonicity)**.** *For all* $\mathbf{b}, \mathbf{b}' \in \mathbb{R}_+^n$, $M(\mathbf{b}) \geqslant M(\mathbf{b}')$ *if* $\mathbf{b} \geqslant \mathbf{b}'$ *and* $M(\mathbf{b}) > M(\mathbf{b}')$ *if* $\mathbf{b} > \mathbf{b}'$. [7]

Assumption 4 (Reserve Price)**.** $M(\mathbf{b}) = r$ *if* $r = \max_i \mathbf{b}_i$.

Note that two major auction formats, first- and second-price sealed-bid auctions with reserve price r, satisfy Assumptions 1-4. More generally, each auction format

[6] The all-pay auction is the most natural example that does not belong to this class.

[7] Here, we adopt the following means of comparing two vectors $\mathbf{b}$ and $\mathbf{b}'$: $\mathbf{b} \geqslant \mathbf{b}'$ if and only if $\mathbf{b}_i \geqslant \mathbf{b}'_i$ for all i; $\mathbf{b} > \mathbf{b}'$ if and only if $\mathbf{b}_i > \mathbf{b}'_i$

with a payment function that is a linear combination of all submitted bids satisfies Assumptions 1-4.

Unlike in standard auctions, in our setting, a seller can impose a penalty on the winner. This penalty depends on the winner's bid and the ex-post verified value. We assume that the seller can commit to a penalty rule $\gamma(b, v) : \mathbb{R}_+ \times V \to \mathbb{R}_+$, where the first argument is a bid and the second argument is the ex-post verified winner's value. Similarly to Mylovanov and Zapechelnyuk (2017), we assume limited punishments: The winner with value v cannot be penalized more than $c(v)$. That is, for any $b \in \mathbb{R}_+ : \gamma(b, v) \leqslant c(v)$. We assume that this maximum penalty function $c(v)$ satisfies the following assumptions:

Assumption 5 (Continuity)**.** *$c(v)$ is continuous.*

Assumption 6 (Monotonicity of Penalty)**.** *$c(v)$ is positive and increasing.*

Assumption 7 (Monotonicity of Surplus)**.** *$v - c(v)$ is positive and increasing.*

Example 1 (Linear Penalty)**.** *The function $c(v) = av$, where $a \in (0, 1)$, satisfies Assumptions 5-7.*

Let $s \in [0, 1]$ be a useful share of the penalty that goes directly to the seller.[8] In one extreme case $s = 0$, the penalty is purely non-monetary (for instance, it may be a reputational loss), and the seller does not receive any direct benefits from it.[9] Mylovanov and Zapechelnyuk (2017) study this situation when monetary transfers are not allowed. In another extreme case $s = 1$, the penalty is purely monetary and may be interpreted as a payment, contingent on the ex-post verified winner's value v. DeMarzo et al. (2005) addresses this situation in the context of unlimited

[8] The majority of our results, in particular Theorem 20, continue to hold if we allow s to be negative. Our interpretation is that it may be costly for the seller to impose a non-monetary penalty on the winner.

[9] As our results indicate, there could be indirect benefits from such penalties, as buyers bid more aggressively in order to not be penalized.

punishments and bidders who are imperfectly informed prior to the bidding stage. Intermediate values of s correspond to a combination of monetary and non-monetary penalties or underline costliness of punishment. All parameters of the model, except for the bidders' valuations, are common knowledge. To summarize, we specify the timeline of the model below.

Timeline. Prior to the auction, the seller announces both an auction format and a penalty rule γ. Next, bidders submit their bids. Given the profile of submitted bids $\mathbf{b}$, the object is allocated to the bidder with the highest bid, provided the bid meets the reserve price r. The losing bidders make no payments to the seller, and therefore, their payoffs are equal to zero. The winning bidder, say i, obtains the object, pays $M(\mathbf{b})$ to the seller, and suffers an additional loss of $\gamma(\mathbf{b}_i, \mathbf{v}_i)$, implying that his payoff equals $\mathbf{v}_i - M(\mathbf{b}) - \gamma(\mathbf{b}_i, \mathbf{v}_i)$. This means that the seller's payoff equals the sum of the winner's transfer $M(\mathbf{b})$ and is discounted by a factor $s \in [0, 1]$ penalty $\gamma(\mathbf{b}_i, \mathbf{v}_i)$: $M(\mathbf{b}) + s\gamma(\mathbf{b}_i, \mathbf{v}_i)$.

Equilibrium. Each auction format and penalty rule γ determine a game among the bidders. A strategy for a bidder i is a Lebesgue measurable function $\beta_i : V \to \mathbb{R}_+$, which describes his bid for every possible value. As is typically done in symmetric environments (see, for instance, Krishna (2009)), we restrict our attention to symmetric equilibria—in which the value of the bid depends only on the buyer's valuation and not on his identity—and refer to β simply as a strategy. The solution concept is a Bayesian-Nash Equilibrium: Given that other bidders' values are distributed according to the distribution function F and other bidders follow strategy β, the bidder with value $v \in V$ maximizes his expected payoff by submitting a bid $\beta(v)$. Let $G(b)$ be the probability that bid b wins the auction, and $m(b)$ be the expected payment of the bidder who wins with bid b. Then, function β is an equilibrium if for all $v \in V$:

$$G(\beta(v))(v - \gamma(\beta(v), v) - m(\beta(v))) \geqslant \sup_{b \in \mathbb{R}_+} G(b)(v - \gamma(b, v) - m(b)) \tag{3.1}$$

Note that both G and m depend on the equilibrium function β. Without loss of generality, we assume that the seller imposes the greatest penalty $c(v)$ if the bidder deviates from the prescribed strategy β:

$$G(\beta(v))(v - \gamma(v) - m(\beta(v))) \geqslant \sup_{b \in \mathbb{R}_+} G(b)(v - c(v) - m(b)) \tag{3.2}$$

Here, with some abuse of notation, $\gamma(v)$ denotes $\gamma(v, \beta(v))$. Therefore, we can interpret β together with γ as either a bidding recommendation or simply a recommendation from the seller. Provided that the auction winner does not follow the recommendation, the maximum penalty is imposed. Otherwise, he is punished by $\gamma(v)$. Defined in this way, γ depends on the bidding function β. Therefore, it is convenient to define γ as a part of the equilibrium:

Definition 1. *A recommendation (γ, β) is an equilibrium of the auction with ex-post verification and limited penalties if and only if* (3.2) *holds for all $v \in V$.*

Fixing an auction format, our goal is to find an equilibrium recommendation that maximizes the seller's revenue:

Definition 2. *Given an auction format, an equilibrium recommendation (γ, β) is an optimal recommendation if it maximizes the ex-ante seller's payoff among all equilibrium recommendations in this auction format.*

The remainder of the paper is organized as follows. In the next section, we study first-price auctions. We determine the seller's optimal recommendation and prove that it is essentially uniquely defined. We complete Section 3.3 by comparing the optimal reserve price in our environment with the one in standard auctions.

In Section 3.4, we show that the optimal recommendation derived in Section 3.3 provides an upper bound on sellers' revenue in all winner-pay auctions. We also show that a second-price auction may provide a lower payoff than a first-price auction. The implications of our theory when considering the optimal taxation of a newly privatized firm are presented in Section 3.5. Our final remarks are in Section 5. All omitted proofs are collected in the Appendix B.

3.3 First-Price Auctions

In this section we derive the optimal recommendation (γ^*, β^*) in a first-price auction. We begin this section by assuming that the reserve price r is fixed. We discuss the implications of allowing the seller to choose the reserve price in Subsection 3.3.3.

In a first-price auction, the expected payment, conditional on winning the object, is simply the winning bid. Therefore, the equilibrium condition can be written as:

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant \sup_{b \in \mathbb{R}_+} G(b)(v - c(v) - b) \tag{3.3}$$

Finding the optimal recommendation is challenging because unlike in a standard auction, in our setting, the equilibrium bidding function β needs not be monotonic. Another complication is that there are equilibria in which a positive mass of types bid the same amount. The following example demonstrates both of these issues.

Example 2. *Values are uniformly distributed on* $[2, 10]$*; there are* $n = 2$ *bidders. The maximum penalty a seller can impose is* $c(v) = v/2$*. For concreteness, all penalties go directly to the seller,* $s = 1$*. There is no reserve price,* $r = 2$*. Given that* $\gamma(v) = 0$ *for all* v*, the following constitutes a symmetric equilibrium:*

$$\beta(v) = \begin{cases} 2, & \text{if } v \in [2, 4] \cup [8, 10] \\ 4, & \text{if } v \in [4, 8] \end{cases}$$

Suppose that bidder 2 follows the above strategy. We argue that it is optimal for bidder 1 to follow β also. Given the distribution of types and the prescribed strategy, bidder 1 faces his opponent's bids 2 and 4 with equal probabilities. Assume now that bidder 1 assigns value $v \in [2,4] \cup [8,10]$ to the object. If he follows the prescribed strategy and submits bid 2, then he wins the object with a probability of 1/4 and his expected utility equals $1/4(v-2)$. Bidding below 2 is not optimal, as in this case he loses the auction. If he submits a bid $b \in [2,4)$, he wins the auction with a probability of 1/2 but pays $b + v/2$ to the seller. Then his expected utility equals $1/2(v - v/2 - b) < 1/4(v-2)$. If he submits bid $b > 4$, he wins the object with certainty and pays $b + v/2$ to the seller. In this case, his expected utility equals $v/2 - b < 1/4(v-2)$. Note that we skip the case of $b = 4$. By bidding $b_\epsilon = 4 + \epsilon$ for a small but positive ϵ, the bidder wins the object with a much higher probability (1 versus 1/2) while paying just a little bit more. This results in a larger payoff for the bidder (formally a supremum over all b_ϵ in the left-hand side of (3.3) is greater than the payoff from deviating to $b = 4$). The case in which the bidder 1 assigns value $v \in [4,8]$ to the object could be similarly considered and hence, omitted.

One useful observation from the above example is that it is sufficient to check only deviations to bids b that are in close proximity to the support of bids corresponding to the prescribed strategy when verifying whether this strategy constitutes an equilibrium. The following lemma states this observation formally.

Lemma 19. *(γ, β) is an equilibrium recommendation in a first-price sealed-bid auction if and only if $\beta(v) + \gamma(v) \leqslant v$ and*

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant \sup_{v' \in V} H(v')(v - c(v) - \beta(v')), \tag{3.4}$$

holds for all $v \in V$, where $H(v')$ is the probability that bid $\beta(v')$ is among the winning bids.

Lemma 19 allows us to treat the problem of finding the optimal recommendation as one of mechanism design. $\beta(v) + \gamma(v) \leqslant v$ is an individual rationality constraint, while (3.4) describes a set of incentive compatibility constraints. Similar to Mylovanov and Zapechelnyuk (2017), the incentive constraints in our problem global rather than local: Unless $\gamma(v) = v$, a buyer never finds it optimal to deviate by mimicking the behaviour of close-by bidders' types. Therefore, the standard methods cannot be applied. The trick is in figuring out which constraints are binding.

The main result of this section is as follows:

Theorem 20. *Under the optimal seller's recommendation (γ^*, β^*_{FPA}), the bidder is never punished provided he follows the recommendation ($\gamma^*(v) = 0$), and every type of bidder is ex-ante worse off compared to in any other equilibrium.*

Moreover, β^*_{FPA} is strictly increasing on $[r, v_H]$ with $\beta^*_{FPA}(r) = r$ and is uniquely defined by the recursive formula:

$$H^*(v)(v - \beta^*_{FPA}(v)) = \sup_{v' \leqslant v - c(v)} H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \tag{3.5}$$

where $H^*(v) = F^{n-1}(v)$ if $v \geqslant r$ and 0 otherwise.

Theorem 20 states that the bidder is indifferent as to whether to follow the recommendation and avoid punishment or to choose the best deviation provided he will be punished by a maximum amount $c(v)$. Here, the best bidder's deviation is to mimic one of types $v' < v - c(v)$. To gain better intuition, consider the case of a single buyer:

Example 3. *A single buyer participates in an auction, $n = 1$. Then $\beta^*_{FPA}(v) = \min\{r + c(v), v\}$ for all $v \geqslant r$. There are two types of deviation worth considering: (1) The buyer bids below the reserve price and consequently loses the auction, and (2) The buyer bids $\beta^*_{FPA}(r) = r$ and wins the object at the lowest price, but is penalized* *by $c(v)$. For all values v such that $v - c(v) < r$, the first type of deviation is better; therefore, to be indifferent as to whether to win or to loose the auction, the bidder should bid according to $\beta^*_{FPA}(v) = v$. If $v - c(v) \geqslant r$, mimicking the type r of bidder is the best deviation. Hence, to make the buyer indifferent between winning the auction at price $\beta^*_{FPA}(v)$ and winning the auction at price r while being penalized by $c(v)$, the buyer should be prescribed to bid $\beta^*_{FPA}(v) = r + c(v)$.*

To see that (3.5) specifies a unique β^*_{FPA} (up to values at $v \leqslant r$), observe that the supremum in the left hand-side is taken over all values of v' that are bounded down away from v. Therefore, we may use previously found values of β^*_{FPA} (at all $v' \leqslant v - c(v)$) to evaluate β^*_{FPA} at v. Given that $c(v)$ is bounded away from zero by $c(v_L)$, the following algorithm defines β^*_{FPA} in finitely many steps:

The Algorithm for Finding β^*_{FPA}

1. Set $I_0 = [0, v_0]$, where $v_0 = r$, $\beta^*_{FPA}(v) = v$, and $k = 1$.

2. Set $I_k = (v_{k-1}, v_k]$, where $v_k = v_H$ if $v_H - c(v_H) < v_{k-1}$; and $v_k = w$, where w is a solution to $w - c(w) = v_{k-1}$ otherwise. For all $v \in I_k$, β^*_{FPA} is already known for all $v' < v - c(v) \leqslant v_{k-1}$, and therefore $\beta^*_{FPA}(v)$ is uniquely defined by (3.5).

3. If $v_k = v_H$, terminate the algorithm; otherwise, set $k = k + 1$ and proceed to step 2.

Note that in the second step of the algorithm, w is uniquely defined, provided $v_H - c(v_H) \geqslant v_{k-1}$, as $v - c(v)$ is a continuous and strictly increasing function. The immediate application of the above algorithm is that larger punishments result in a higher revenue for the seller.

Corollary 21. *For the two maximum penalty functions c_1 and c_2, with $c_1(v) < c_2(v)$ for all $v \in V$, the seller receives a weakly higher revenue (and strictly higher if*

$v_H - c_1(v_H) > r$) in a first-price auction with the maximum penalty function c_2 than he would in one with the maximum penalty function c_1.

We use the above algorithm to compute the optimal recommendation in Example 2:

Example 2.(Continued). *Values are uniformly distributed on $[2;10]$, there are $n = 2$ bidders, $c(v) = v/2$, and $r = 2$. Then β^*_{FPA} is defined as:*

$$\beta^*_{FPA}(v) = \begin{cases} v, & \text{if } v \in [2,4] \\ \frac{15v^2-24v-16}{16(v-2)}, & \text{if } v \in [4,10] \end{cases}$$

For comparison, the equilibrium bidding function in a standard first-price auction is $\beta_{FPA}(v) = (v+2)/2$. It is easy to verify that $\beta^*_{FPA}(v) > \beta_{FPA}(v)$ for all v. The seller's revenue equals $R^* \approx 7.19$, which is greater than $R^S \approx 4.67$—the revenue in a standard first-price auction.

The observation that buyers bid more aggressively than in a standard first-price auction can be easily generalized as such:

Corollary 22. *When following the optimal recommendation, bidders bid more than they do in a standard first-price auction. This results in a higher expected revenue for the seller.*

In example 2, the seller's ability to punish the bidder by half of his private value has a substantial effect on the seller's revenue. The seller extracts almost the full surplus—the revenue the seller receives if he sells the object with full information about each bidder's private valuation—which equals 7.33 in the example. We further show that even small punishments have a significant impact on the seller's revenue, which helps the seller to substantially reduce information loss, or the difference between the full surplus and the seller's revenue R^S in a standard first-price auction.

Example 4. *Values are uniformly distributed on* $[2; 10]$*, there are* $n = 2$ *bidders,* $c(v) = av$*, and* $r = 2$*. Figure 1 depicts the seller's revenue* $R^*(a)$ *as a function of parameter* a*. For instance, if* $a = 0.01$*, then the seller's revenue equals approximately* 5.14*—compensating for almost* 18% *of the information loss in a standard first-price auction. If the seller is able to punish the buyer by only one-thousandth of his value, she still regains* 5% *of the information loss.*

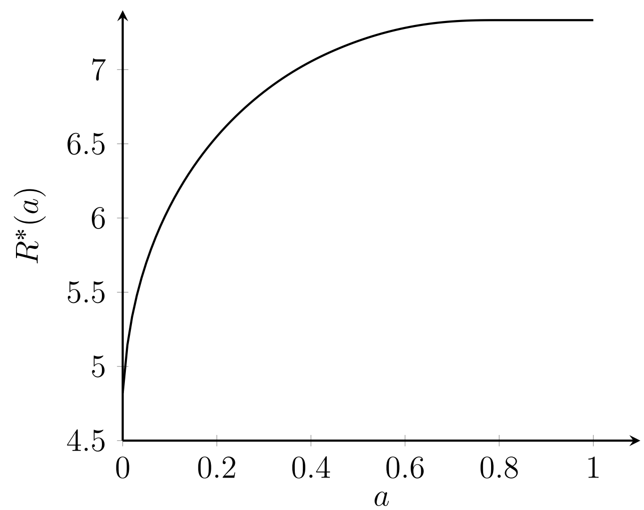

FIGURE 3.1: The Seller's Revenue in Example 4

We generalize the observation that small punishments significantly increase the seller's revenue in Proposition 23.

Proposition 23. *For a number of bidders* $n \geqslant 2$*, small linear penalties* $c(v) = av$ *have a large effect on the seller's revenue—that is:*

$$\lim_{a \to 0+} \frac{R^*(a) - R^S}{a} = \infty$$

In the proof we construct an equilibrium recommendation that prescribes a buyer to bid the weighted average of his value and his bid in a standard first-price auction.

This recommendation is not optimal, but nevertheless, small punishments have a large impact on the seller's revenue as defined in Proposition 23. Since the optimal recommendation leads to greater revenue for the seller than a constructed recommendation would, the conclusion of Proposition 23 follows.

3.3.1 Proof of Theorem 20.

We divide this proof into two steps. In the first step (Proposition 24), it is proven that $(\gamma^*, \beta^*_{FPA})$ is an equilibrium recommendation in a first-price auction. To prove this, we first assume that β^*_{FPA} defined by (3.5) is a monotonic function. Then the probability of a buyer winning the auction with value $v \geqslant r$ equals $H^*(v)$. Therefore, (3.5) implies that all incentive constraints for which $v' \leqslant v - c(v)$ are satisfied. We show that when one of these constraints is binding, it is not optimal for a type v buyer to mimic a type $v' > v - c(v)$. We complete the proof of Proposition 24 by showing that β^*_{FPA} is, in fact, a monotonic function.

In the second part (the actual proof of Theorem 20), we show that $(\gamma^*, \beta^*_{FPA})$ is an optimal recommendation. For any equilibrium recommendation (γ, β), if the object is allocated to the bidder with value v, the total surplus—the sum of the seller's and the auction winner's payoff—equals $v + (s-1)\gamma(\beta(v))$. Therefore, $(\gamma^*, \beta^*_{FPA})$ maximizes the total surplus for all possible profiles of bidders' valuations: The object is allocated to the bidder with the highest value, and he is never punished. Given that the seller's expected profit and the total surplus of the expected buyer equal the expected total surplus, it is sufficient to show that every type of buyer receives a lower expected surplus than under any other recommendation (γ, β). We show this by mathematical induction on a number of steps k of the algorithm for finding β^*_{FPA}.

Proposition 24. $(\gamma^*, \beta^*_{FPA})$ *is an equilibrium recommendation in a first-price auction.*

Proof. Assume that β^*_{FPA} is strictly increasing on $[r, v_H]$, and let $u^*(v)$ be an expected buyer's payoff given that he follows β^*_{FPA}. Then $u^*(v) = H^*(v)(v - \beta^*_{FPA})$ and equilibrium condition (3.4) can be rewritten as:

$$u^*(v) \geqslant u^*(v') + H^*(v')(v - c(v) - v') \tag{3.6}$$

for all $v, v' \in V$. Fix $v \in [r, v_H]$. By definition of β^*_{FPA} and u^*, (3.6) holds for all $v' \leqslant v - c(v)$. Next, observe that (3.5) implies that $u^*(\cdot)$ is a weakly increasing function in its argument. Then (3.6) also holds for all $v' \in (v - c(v), v]$. Hence, to show that β^*_{FPA} is an equilibrium, we have only left to prove that (3.6) holds for all $v' > v$. Suppose, by way of contradiction, that this is not true, meaning that for some $v' > v$:

$$u^*(v) < u^*(v') + H^*(v')(v - c(v) - v')$$

Let S_v be the set of all such v'. Then define $v_{inf} = \inf\{v' : v' \in S_v\}$. Take an arbitrary $v'' \in [v_{inf}, v_{inf} + c(v_L)] \cap S_v$. Since c is an increasing function, $v'' - c(v'') < v_{inf}$. Given that inequality (3.6) holds for all $v' < v_{inf}$ and using the definition of u^*, we have:

$$\begin{aligned} u^*(v'') - u^*(v) = \sup_{v' < v'' - c(v'')} H^*(v')(v'' - c(v'') - \beta^*_{FPA}(v')) - \sup_{v' < v_{inf}} H^*(v')(v - c(v) \\ - \beta^*_{FPA}(v')) \leqslant \sup_{v' < v_{inf}} H^*(v')(v'' - c(v'') - \beta^*_{FPA}(v')) - \sup_{v' < v_{inf}} H^*(v')(v - c(v) \\ - \beta^*_{FPA}(v')) \leqslant \sup_{v' < v_{inf}} H^*(v')(v'' - c(v'') - (v - c(v))) \\ \leqslant H^*(v'')(v'' - c(v'') - (v - c(v))) \leqslant H^*(v'')(v'' - (v - c(v))) \end{aligned}$$

where the second inequality follows from the triangle inequality for supremum and the third one follows from $v - c(v)$ be an increasing function.

Hence,

$$u^*(v) \geqslant u^*(v'') + H^*(v'')(v - c(v) - v'')$$

contradicting $v'' \in S_v$.

Finally, we prove that β^*_{FPA} is a strictly increasing function on $[r, v_H]$. Since $c(v) > 0$ for all $v \in V$, it is sufficient to show that $\beta^*_{FPA}(v) > \beta^*_{FPA}(w)$ for all $v > w > v - c(v)$. Using the triangle inequality for supremum once again, on the one hand,

$$\begin{aligned} u^*(v) - u^*(w) &= \sup_{v' \leqslant v - c(v)} H^*(v')(v - c(v) - \beta^*(v')) - \sup_{v' \leqslant w} H^*(v')(w - c(w) - \beta^*(v')) \\ &\leqslant \sup_{v' \leqslant w} H^*(v')(v - c(v) - \beta^*(v')) - \sup_{v' \leqslant w} H^*(v')(w - c(w) - \beta^*(v')) \\ &\leqslant \sup_{v' \leqslant w} H^*(v')(v - c(v) - (w - c(w))) \leqslant H^*(w)(v - w + c(w) - c(v)) \end{aligned}$$

On the other hand,

$$\begin{aligned} u^*(v) - u^*(w) &= H^*(v)(v - \beta^*(v)) - H^*(w)(w - \beta^*(w)) \\ &\geqslant H^*(w)(v - w + \beta^*(w) - \beta^*(v)) \end{aligned}$$

Therefore, unless $H^*(w) = 0$,

$$\beta^*(w) \leqslant \beta^*(v) + c(w) - c(v) < \beta^*(v)$$

If $H^*(w) = 0$, then $u^*(v) = u^*(w) = 0$, implying that $\beta^*(v) = v > \beta^*(w)$. □

Lemma 25. *For all $w < v$,*

$$u^*(v) - u^*(w) \leqslant H^*(v)(v - w)$$

Proof. Given that $(\gamma^*, \beta^*_{FPA})$ is an equilibrium recommendation and using the definition of β^*_{FPA} yields:

$$u^*(v) = \sup_{v' \leqslant a(v)} H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \tag{3.7}$$

for all $a(v) \geqslant v - c(v)$. Putting $a(v) = a(w) = v$, we have:

$$\begin{aligned} u^*(v) - u^*(w) &\leqslant \sup_{v' \leqslant v} H^*(v')(v - c(v) - \beta^*_{FPA}(v')) - \sup_{v' \leqslant v} H^*(v')(w - c(w) - \beta^*_{FPA}(v')) \\ &\leqslant \sup_{v' \leqslant v} H^*(v')(v - c(v) - (w - c(w))) \leqslant H^*(v)(v - w) \end{aligned}$$

□

Proof of Theorem 20. Let γ be an arbitrary penalty rule and β_{FPA} be a corresponding equilibrium in a first-price auction. Note that the total welfare under recommendation β^*_{FPA} is at least as large as it is under β. Therefore, it is sufficient to show that $u(v) \geqslant u^*(v)$ for all v, where $u(v) = G(\beta_{FPA}(v))(v - \gamma(v) - \beta_{FPA}(v))$. We prove this fact by mathematical induction on a number of steps k of the algorithm of finding β^*_{FPA}. The claim is:

Claim: For an arbitrary $k \geqslant 0$ and $v \leqslant v_k$: $u(v) \geqslant u^*(v)$.

Base, $k = 0$**.** For all $v \leqslant v_0 = r$: $u^*(v) = 0$ and hence, $u(v) \geqslant u^*(v)$.

Induction Step, $k = l + 1$**.** Assume that $u(v) \geqslant u^*(v)$ for all $v \in [0, v_l]$. Fix an arbitrary $v \in (v_l, v_{l+1}]$. Recall that by the definition of β^*_{FPA}:

$$u^*(v) = \sup_{v' \leqslant v - c(v)} H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \tag{3.8}$$

Also, since β_{FPA} is an equilibrium recommendation, it should satisfy:

$$u(v) \geqslant \sup_b G(b)(v - c(v) - b) \tag{3.9}$$

Hence, it is sufficient to show that for an arbitrary $v' \leqslant v - c(v)$, there exists $b \in \mathbb{R}_+$ such that:

$$H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \leqslant G(b)(v - c(v) - b) \tag{3.10}$$

Fix an arbitrary $v' \leqslant v - c(v)$. There are two possibilities: (1) $\beta^*_{FPA}(v') > \beta_{FPA}(w)$ for all $w \leqslant v'$ and (2) There exists $w \leqslant v'$ such that $\beta^*_{FPA}(v') \leqslant \beta_{FPA}(w)$.

Case 1. $\beta^*_{FPA}(v') > \beta_{FPA}(w)$ **for all** $w \leqslant v'$**.** Then define $b = \beta^*_{FPA}(v')$. This implies that by bidding b in the auction with the seller's recommendation β_{FPA}, and imposing that others bid according to a prescribed strategy β_{FPA}, a bidder wins the object in each situation in which his opponents' values do not exceed v'. Hence, $G(b) \geqslant H^*(v')$, which combined with $b = \beta^*_{FPA}(v')$ gives us (3.10).

Case 2. There exists $w \leqslant v'$ such that $\beta^*_{FPA}(v') \leqslant \beta_{FPA}(w)$. Define $b = \beta_{FPA}(w)$. Since $w \leqslant v_l$, by the induction hypothesis $u(w) \geqslant u^*(w)$, which implies

$$G(b)(w - b) \geqslant H^*(w)(w - \beta^*_{FPA}(w)) \tag{3.11}$$

Also, by Lemma 25, we have

$$H^*(w)(w - \beta^*_{FPA}(w)) \geqslant H^*(v')(v' - \beta^*_{FPA}(v')) - H^*(v')(v' - w) = H^*(v')(w - \beta^*_{FPA}(v')) \tag{3.12}$$

Combining (3.11) and (3.12) yields:

$$G(b)(w - b) \geqslant H^*(v')(w - \beta^*_{FPA}(v')) \tag{3.13}$$

From (3.13), it follows that $G(b) \geqslant H^*(v')$ as $\beta^*_{FPA}(v') \leqslant b$. Therefore,

$$\begin{aligned} G(b)(v - c(v) - b) &= G(b)(v - c(v) - w) + G(b)(w - b) \\ &\geqslant H^*(v')(v - c(v) - w) + H^*(v')(w - \beta^*_{FPA}(v')) \\ &= H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \end{aligned}$$

where the inequality follows from $G(b) \geqslant H^*(v')$, $w \leqslant v - c(v)$, and (3.13). Hence, (3.10) holds. □

3.3.2 Uniqueness of Optimal Recommendation

Is optimal recommendation unique? According to Theorem 20, every equilibrium recommendation (γ, β_{FPA}) that provides the same ex-ante seller's payoff as the optimal recommendation $(\gamma^*, \beta^*_{FPA})$ should satisfy $u(v) = u^*(v)$ for all but the zero measure of bidder's types v. Here, as before, $u(v)$ is type v bidder's expected payoff, which corresponds to recommendation (γ, β_{FPA}), with $u^*(v)$ defined in a similar manner. In addition (γ, β_{FPA}) needs to maximize the expected total surplus: The highest-value bidder wins an object, and the sum of the penalty transfers equals zero. For $s < 1$, the latter is possible only if no penalties are imposed on the equilibrium path. Therefore, $\gamma = \gamma^*$. But then $u(v) = u^*(v)$ implies that $\beta_{FPA}(v) = \beta^*_{FPA}(v)$ for all but positive measure of types v:

Corollary 26. *For every useful share of penalty $s \in [0, 1)$ there exists a unique optimal recommendation $(\gamma^*, \beta^*_{FPA})$.*

Before discussing the case $s = 1$, recall that we derive β^*_{FPA} based on the assumption that out of equilibrium, the behaviour of type v bidder is punished by a maximum amount $c(v)$, given that he is the winner. However, such a penalty is not necessary. In particular, if there is a single buyer, he has no incentive to bid more than prescribed, even if there is no penalty imposed. In the same vein, local deviations can be punished more deliberately: In a single-buyer case (see example 3), to prevent type v buyer, where v is such that $v - c(v) < r$, from submitting bid $v' \in [r, v]$, it is sufficient to punish such deviation by $v - v'$. This result holds more generally:

Corollary 27. *Let penalty rule $\tilde{\gamma}$ be defined as*

$$\tilde{\gamma}(b, v) = \begin{cases} 0, \;\; \textit{if } b \geqslant \beta^*_{FPA}(v) \\ v - v', \;\; \textit{if } b = \beta^*_{FPA}(v') \textit{ and } v - c(v) < v' < v \\ c(v), \textit{otherwise} \end{cases} \tag{3.14}$$

Given penalty rule $\tilde{\gamma}$, β^*_{FPA} is a symmetric equilibrium in a first-price auction.

Coming back to our discussion about the uniqueness of the optimal recommendation, if $s = 1$, we may specify a different equilibrium recommendation (γ, β_{FPA}) such that the total payment to the seller remains unchanged. To do this, define a bidding function β_{FPA} as

$$\beta_{FPA}(v) = \begin{cases} \beta^*_{FPA}(v), \;\text{if } v < \tilde{v} \\ (1 - \alpha)\beta^*_{FPA}(v) + \alpha\beta^*_{FPA}(\tilde{v}), \;\text{otherwise} \end{cases} \tag{3.15}$$

where $\tilde{v}$ is an arbitrary value in the interval $(v_H - c(v_H), v_H]$ and $\alpha \in (0, 1)$.

The penalty rule γ is defined, as outlined above, in order to keep the total payments to the seller unchanged: $\gamma(v) = \beta^*_{FPA}(v) - \beta_{FPA}(v)$. All deviations are punished by the maximum amount $c(v)$. To sustain (γ, β_{FPA}) as an equilibrium recommendation, it is sufficient to show that no type v bidder would aim to deviate by submitting bid $\beta_{FPA}(w)$ for some $w \geqslant \tilde{v}$. We show that for a sufficiently small α, the net transfers to the seller under any such deviation are greater than the net transfers to the seller under equilibrium β^*_{FPA}, provided the bidders are punished according to $\tilde{\gamma}$ defined by (3.15). This makes such deviations unprofitable.

Lemma 28. *For a sufficiently small but positive α recommendation, (γ, β_{FPA}), where β_{FPA} is defined according to* (3.15) *and $\gamma(v) = \beta^*_{FPA}(v) - \beta_{FPA}(v)$, constitutes an equilibrium in a first-price auction.*

Also, if there is a single buyer, the equilibrium bidding function β does not need to be monotonic. An alternative to this is presented in example 3, where the equilibrium recommendations $\beta_{FPA}(v) = r$ and $\gamma(v) = \min\{v - r, c(v)\}$ provide the same payoff $\min\{v, r + c(v)\}$ to the seller. Therefore, the uniqueness of the optimal recommendation is established up to the net transfers received by the seller, and the monotonicity of the bidding function holds if there are at least two bidders in the auction.

Corollary 29. *For a useful share of the penalty $s = 1$ and any optimal recommendation (γ, β_{FPA}), the net transfers from the winner with a value v to the seller are equal to $\beta^*_{FPA}(v)$ $(\gamma(v) + \beta_{FPA}(v) = \beta^*_{FPA}(v))$. Moreover, if the number of bidders $n \geqslant 2$, β_{FPA} is monotonic on $[r, v_H]$.*

3.3.3 Optimal Reserve Price

Suppose now that the seller sets a reserve price r in such a way as to maximize her expected revenue. Is the optimal reserve price higher or lower than the optimal reserve price in a standard first-price auction? In other words, do access to verification technology and the ability to impose limited punishments necessarily lead to improvements in auction efficiency?

To answer this question, we follow Myerson (1981) by assuming that virtual valuation $\psi(\cdot)$ that is defined as

$$\psi(v) = v - \frac{1 - F(v)}{f(v)}$$

is an increasing function (regular case). In a standard auction, the seller finds it optimal to set a reserve price equal to v_L if $\psi(v_L) \geqslant 0$, and $\psi^{-1}(0)$ otherwise. In the latter case, the seller excludes some bidders from the auction by setting the reserve price above their valuations.

In our setting, the following result holds:

Theorem 30. *In a regular case, the optimal reserve price is weakly lower than in a standard auction. Moreover, if $\psi(v_L) < 0$, the optimal reserve price is strictly lower than in a standard auction.*

Let r_S be the optimal reserve price in a standard first-price auction. We find that for every reserve price $r > r_S$, the seller gains at least as much as in a standard auction when reducing the reserve price by a small amount. In a regular case, such a reduction of the reserve price is beneficial for the seller in a standard first-price auction, and hence, beneficial in our environment as well. The complete proof of Theorem 30 can be found in the Appendix B.

Unlike in standard auctions, deriving the explicit general formula[10] for the reserve price is possible only in very special cases.

Corollary 31. *If $c(v_H)$ is sufficiently high to make $v_H - c(v_H) \leqslant v_L$, it is optimal to set the reserve price equal to v_L.*

Proof. Provided $v_H - c(v_H) \leqslant v_L$, $\beta^*_{FPA}(v) = v$, and therefore, it is optimal to set $r = v_L$ and extract the full surplus. □

For intuition behind Theorem 30, consider a case of a single buyer—an example in which the reserve price can be found explicitly and is not necessarily equal to v_L. In this case, the problem of choosing the optimal reserve price parallels[11] the problem that the monopolist faces when selling homogeneous goods to a continuum of buyers, assuming that a measure of buyers whose valuations do not exceed v is equal to $F(v)$. Then a standard auction is essentially equivalent to a non-discrimination case: The same price p is set for all buyers. Therefore, the demand equals $q = 1 - F(p)$, and the total revenue is $TR = qp = p(1 - F(p))$. Then the marginal revenue equals

$$MR_S = \frac{dTR}{dp} \cdot \frac{dp}{dq} = p - \frac{1 - F(p)}{f(p)} = \psi(p)$$

and given that, at the optimal price $MR_S = 0$, $r_S = \psi^{-1}(0)$, as outlined above.

According to Example 3, in our environment, the monopolist has a partial ability to discriminate among buyers. To maximize her profit, she charges buyer of type v $\min\{c(v), v - p\}$ in addition to the regular price p. The total revenue then equals $TR = \int\limits_{p}^{v_H} \min\{c(v) + p, v\} dF(v)$, and the marginal revenue equals

$$MR_V = \frac{dTR}{dp} \cdot \frac{dp}{dq} = p - \frac{1 - F(v_p)}{f(p)},$$

[10] Given the specific distribution function, the derivation of the optimal reserve price is straightforward.

[11] See Bulow and Roberts (1989) for a discussion of how this observation generalizes to more than one bidder in standard auctions.

where $v_p = v_H$ if $v_H - c(v_H) < p$ and is otherwise determined as a solution to $v_p - c(v_p) = p$. Since $v_p > v$, the marginal revenue MR_V is greater than the marginal revenue MR_S in a non-discrimination case. Consequently, the optimal reserve price r_V is lower than r_S.

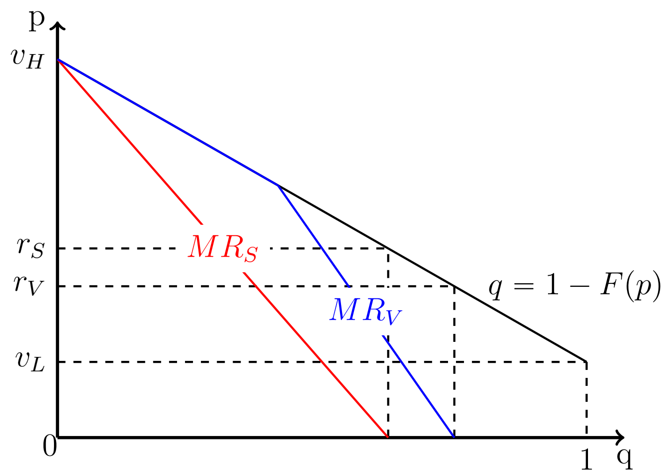

FIGURE 3.2: Comparison of Optimal Reserve Prices for a Single-Buyer Case

Why is the marginal revenue MR_V greater than MR_S? To gain further insight into this, consider a small reduction in price p (that corresponds to a small increase in quantity). The marginal benefit from this due to more buyers entering the market in both discriminatory and non-discriminatory cases equals $pf(p)$. However, the marginal losses are different in these cases. In non-discriminatory cases, all buyers above the reserve price now pay less, and the marginal loss equals $1 - F(p)$. In discriminatory cases, the buyers with valuations just above the reserve price pay their values and hence are not affected by this change. Only the buyers with valuations above v_p benefit from a reduction in prices. Therefore, the marginal loss that equals $1 - F(v_p)$ is lower than it is in non-discriminatory cases, which implies that $MR_V > MR_S$.

Define a modified virtual valuation $\phi(p)$ as

$$\phi(p) = p - \frac{1 - F(v_p)}{f(p)}$$

Provided ϕ is a monotonic function, we obtain an explicit formula for the reserve price

Proposition 32. *In a single-buyer case, if the modified virtual valuation $\phi(\cdot)$ is monotonic, then the optimal reserve price equals v_L if $\phi(v_L) \geqslant 0$ and equals $\phi^{-1}(0)$ otherwise.*

In the proof of Theorem 30, we show that the marginal gain from decreasing the reserve price r is at least as large as $nF^{n-1}(r)f(r)\phi(r)$. Therefore,

Corollary 33. *If the modified virtual valuation $\phi(\cdot)$ is monotonic, then the optimal reserve price equals v_L if $\phi(v_L) \geqslant 0$ and does not exceed $\phi^{-1}(0)$ otherwise.*

Note, that monotonicity of the modified virtual valuation holds for all convex distribution functions. In particular, if the valuations are distributed uniformly on $[v_L, v_H]$ and the penalty rule is linear $c(v) = av$, then the optimal reserve price in a single-buyer case equals $\max\{\frac{1-a}{2-a} v_H, v_L\}$. In Example 2 ($v_L = 2$, $v_H = 10$ and $a = 0.5$), the optimal reserve price for a single-buyer case equals 10/3. For more than one buyer, it can be verified that the optimal reserve price equals 2, and hence, no buyers are excluded from the auction. For comparison, the reserve price in a standard first-price auction is 5.

Remark. *The optimal reserve price may depend on the number of bidders. An optimal auction may be efficient for number of bidders $n \geqslant 2$ even if it is not efficient in a single-bidder case.*

3.4 Other Auction Formats

In this section, we look beyond a first-price auction format. First we consider an arbitrary winner-pay auction. Let (γ, β) be an equilibrium recommendation. In this section, we restrict our attention to monotonic recommendations:

Definition 3. *Recommendation (γ, β) is a monotonic recommendation if β is a strictly increasing function on the interval $[r, v_H]$.*

The main result of this section is

Theorem 34. *For arbitrary monotonic recommendation (γ, β) of a winner-pay auction, there exists a recommendation $(\gamma_{FPA}, \beta_{FPA})$ of a first-price auction that provides the same ex-ante payoffs for all parties (including the seller) as (γ, β).*

We prove Theorem 34 by showing that, for a given monotonic recommendation (γ, β) of a winner-pay auction, the seller can achieve the same revenue by using the first-price auction format with seller's recommendation (γ, β_{FPA}), where $\beta_{FPA}(v) = m(\beta(v))$. That is, the seller recommends that a buyer bid the same amount he will bid in expectation, conditional on winning in the other auction format. Therefore, according to Theorem 20, the first-price auction with an optimally chosen recommendation $(\gamma^*, \beta^*_{FPA})$ provides a (weakly) greater revenue to the seller than any other winner-pay auction.

To show that other auction format may fail to provide the same amount of revenue to the seller as the optimal first-price auction format, we investigate a buyer's behaviour in a second-price auction. In a second-price auction, the winner pays the value of the second-highest bid to the seller. Let (γ, β_{SPA}) be an equilibrium recommendation in a second-price auction. Recall that without loss of generality, any violation of the equilibrium strategy β_{SPA} by a buyer with value v is punished by the maximum amount $c(v)$. In the same vein that bidding one's value is the dominant strategy in the second-price auction, the only possible deviation worth considering[12] is bidding $v - c(v)$.

Lemma 35. (γ, β_{SPA}) *is an equilibrium recommendation of the second-price auction if and only*

$$G(\beta_{SPA}(v))(v - \gamma(v) - m(\beta_{SPA}(v))) \geqslant G(b)(v - c(v) - m(b))$$

for $b = v - c(v)$.

Let $u(v)$ be the expected utility of type v buyer that corresponds to a symmetric equilibrium β_{SPA} in the second-price auction. As before, let $u^*(v)$ be the expected utility of type v buyer that corresponds to a symmetric equilibrium β^*_{FPA} in the optimal first-price auction. Based on Theorem 34 and Theorem 20, we may conclude that $u(v) \geqslant u^*(v)$ for all $v \in V$. Hence, the second-price auction is revenue-equivalent to the first-price auction with recommendation $(\gamma^*, \beta^*_{FPA})$ if $u(v) = u^*(v)$ for all but the zero measure of types v. To simplify arguments, we assume that a fraction of the penalty that goes directly to the seller $s < 1$. Therefore, $\gamma(v) = 0$ for all but the zero measure of types v. To conclude:

$$F^{n-1}(v)(v - \beta^*_{FPA}(v)) = F^{n-1}(v)(v - m(\beta_{SPA}(v)))$$

for all but the zero measure of types $v > r$. Therefore, $\beta^*_{FPA}(v) = m(\beta_{SPA}(v))$ for all but the zero measure of types $v > r$.

If there is a single buyer $n = 1$, $m(\beta_{SPA}(v)) = r$ for all $v > r$. However, according to Example 3, $\beta^*_{FPA}(v) = \min\{v, r + c(v)\} > r$ for all $v > r$. Therefore, we obtain

Proposition 36. *If there is a single buyer and* $s < 1$*, a second-price auction provides the seller with a strictly lower payoff than the optimal first-price auction.*

[12] If a buyer with value v decides to bid differently from $\beta_{SPA}(v)$, then the buyer's net value of penalty payments is $v - c(v)$. Hence, given that he deviates, the dominant strategy is to bid $v - c(v)$.

Assume now that $n \geqslant 2$. Then given that the expected unconditional payment of type v buyer equals $rF^{n-1}(r) + \int_r^v \beta_{SPA}(t)dF^{n-1}(t)$, the equilibrium condition (3.5) for the first-price auction can be rewritten as

$$F^{n-1}(v)v - rF^{n-1}(r) - \int_r^v \beta_{SPA}(t)dF^{n-1}(t) = \sup_{v' \in V} H^*(v')(v - c(v) - m(\beta_{SPA}(v')))$$

Here as before $H^*(v') = F^{n-1}(v')$ if $v' \geqslant r$ and $H^*(v') = 0$ otherwise. Note that according to Lemma 35, the value of v' that maximizes the right-hand side of the above equality satisfies[13] $\beta_{SPA}(v') = v - c(v)$. Differentiating both sides of above equation with respect to v and applying the Envelope Theorem yields:

$$\beta_{SPA}(v) = v + \frac{F(v)}{(n-1)f(v)} - \frac{H^*(\beta_{SPA}^{-1}(v - c(v)))}{(n-1)F^{n-2}(v)f(v)} \left(1 - \frac{dc(v)}{dv}\right) \tag{3.16}$$

where we assume that $\beta_{SPA}(v) = v$ for all $v \leqslant r$ and $\beta_{SPA}^{-1}(w) = \sup\{v : \beta_{SPA}(v) \leqslant w\}$ for all w for which the inverse of β_{SPA} fails to exist. Therefore, we have proven:

Proposition 37 (Revenue Equivalence)**.** *If $s < 1$, the expected revenue in a second-price auction is the same as it is in the first-price auction (with an optimally chosen recommendation) if and only if there exists a strictly increasing bidding function β_{SPA} that satisfies* (3.16) *for all $v > r$.*

Analogous to the first-price auction, let β^*_{SPA} denote the monotonic solution of (3.16), provided it exists. One important feature of β^*_{SPA} is that the bid exceeds the value[14]. We use this fact to introduce the recursive algorithm for finding β^*_{SPA}

[13] To be more rigorous, it is the maximum value of v' such that $\beta_{SPA}(v') \leqslant v - c(v)$.

[14] Recall that in an optimal first-price auction buyers bid more than in a standard first-price auction. Consequently, the expected payment by type v bidder is also larger than it is in a standard auction. Provided that β^*_{SPA} and β^*_{FPA} lead to the same expected payment by type v bidder, buyers need to bid more than in a standard second-price auction (that is more than v).

similarly to how we find that of the first-price auction. Note that this algorithm also provides a constructive way of checking the conditions of Proposition 37

The Algorithm for Finding β^*_{SPA}

1. Set $I_0 = [0, v_0]$, where $v_0 = r$; $\beta^*_{SPA}(v) = v$ and $k = 1$.

2. Set $I_k = (v_{k-1}, v_k]$, where $v_k = v_H$ if $v_H - c(v_H) < v_{k-1}$; and $v_k = w$, w is a solution to $w - c(w) = v_{k-1}$ otherwise. For all $v \in I_k$, $v - c(v) < v_{k-1}$ and therefore $(\beta^*)^{-1}_{SPA}(v - c(v))$ is well defined and hence, $\beta^*_{SPA}(v)$ is evaluated by (3.16).

3. If β^*_{SPA} is non-monotonic on $[0, v_k]$, then terminate the algorithm: Revenue equivalence does not hold.

4. If $v_k = v_H$, terminate the algorithm; otherwise set $k = k + 1$ and proceed to step 2.

When does revenue equivalence fail for $n \geqslant 2$? This question can now be rephrased as: Under what conditions does the above algorithm terminate at step 3? Assume it terminates for $k = 1$. For all $v \in I_1$, equation (3.16) is especially simple

$$\beta_{SPA}(v) = v + \frac{F(v)}{(n-1)f(v)} \tag{3.17}$$

For β_{SPA} to be non-monotonic on I_1, the density function f needs to increase sufficiently "fast" on some sub-interval of I_1 to compensate for the increase in the distribution function F and in v on the same sub-interval. We illustrate this possibility using a numerical example similar to the one presented in the Introduction for a discrete version of our model.

Example 5. *Values are distributed on* $[2, 14]$ *according to the absolutely continuous distribution function with corresponding density function*

$$f(v) = \begin{cases} \frac{1}{50}, & \text{if } v \in [2, 11) \\ \frac{19}{100}(v-13) + \frac{2}{5}, & \text{if } v \in [11, 13] \\ \frac{2}{5}, & \text{if } v \in (13, 14] \end{cases}$$

There are two bidders $n = 2$*, the reserve price* $r = 2$*, and the maximum penalty function* $c(v) = 0.9v$*.*

Since $v_H - c(v_H) = 14 - 0.9 \cdot 14 < 2 = v_L$*,* $I_1 = (2, 14]$*, it is possible to extract the full surplus in the first-price auction. Applying formula* (3.17)*, we can compute that: (1) For all* $v < 11$*,* $\beta_{SPA}(v) = 2v - 2$*, and (2) For all* $v > 13$*,* $\beta_{SPA}(v) = v + (3/5 + 2/5(v-13))/(2/5) = 2v - 11.5$*. Therefore, for all* $v > 13$ *and* $w \in [10, 11]$*,* $\beta_{SPA}(v) < \beta_{SPA}(w)$*, implying that there is no monotonic solution for* (3.16)*.*

The above example can be generalized even for those cases in which the maximum penalty that the seller can impose on a buyer is sufficiently small.

Proposition 38. *For the arbitrary maximum penalty function* $c(v)$ *and* $s < 1$*, there exists an absolutely continuous distribution function* F *on* $[v_L, v_H]$ *such that the seller's revenue in a second-price auction is strictly lower than in an optimal first-price auction.*

Therefore, the failure of revenue equivalence is due to the possible non-existence of monotonic equilibrium that generates the desirable buyers expected (conditional on winning) payments.

In a first-price auction, a buyer's payment is his bid and consequently, there is always a monotonic equilibrium, provided the desirable expected payment is a monotonic function in a bidder's value.

We outline two other reasons why the first-price auction format may be the preferable auction format in our environment. First, assume that buyers are financially constrained and cannot afford to pay more than their valuations. Then in a second-price auction, each buyer's bid cannot exceed his value, implying that the seller's expected revenue does not exceed that of a standard auctions. Therefore, according to Corollary 22, the first-price auction provides a strictly larger revenue to the seller.

Second, consider the possibility of collusion. For simplicity, assume that all bidders form a cartel. In a second-price auction, cartel bidders may successfully escape penalties and buy the object at a reserve price: The bidder with the highest value v bids $\beta^*_{SPA}(v)$, and the rest of the cartel members submit bids that are not serious (at or below r). In contrast, in a first-price auction, it is impossible for a cartel to buy the object at a reserve price: Even if the winner with value v wins the auction at price r, he will be punished by $c(v)$ unless $v = r$ and therefore, the total revenue equals $r + c(v)$.

3.5 Application: Selling a Firm by a Government.

Consider a government that wants to sell a firm via an auction. To simplify the exposition, we assume that there is no reserve price. The firm generates a revenue $R(e) = 2\sqrt{\alpha e}$, if its owner exerts an effort e and his productivity level equals α. The cost of effort is $C(e) = e$. There are n potential buyers, who are heterogeneous in their productivity levels α. Assume that productivity levels are distributed independently and identically across bidders according to the distribution function F on $[\alpha_L, \alpha_U]$. All information besides private ability levels is common knowledge. In addition to payments received during the auction stage, the government receives a tax on profit $t\pi$, where π is the profit generated by the winning bidder and $t \in (0,1)$ is the tax rate.

Given his productivity level α, the winner maximizes the net profit of taxes: $\pi(e) = (1-t)(2\sqrt{\alpha e} - e)$. Let v_α be the maximum value of profit the entrepreneur with productivity level α can generate. For our specific functional form, $v_\alpha = \alpha$.

If we interpret $c(v) = tv$ as the maximum penalty that the government can impose on the buyers, then the current specification fits our model. There is one exception, though: The winner could generate any value $v \in (-\infty, v_\alpha]$, while in our model, we assume that the ex-post value is not under the control of the winner.

In particular, a standard auction corresponds to a situation in which penalty rule $\gamma(b, v) = c(v)$ for all values of bids b and ex-post profits v. We call this situation a flat taxation—the government imposes the same level of taxation independent of the winning bid and profit generated by the winner. Ignoring for a moment that the winner is able to generate a lower value than v_α, we know from Section 3.3 that flat taxation is sub-optimal. The government can do strictly better by promising to not impose any taxation if the buyer submits a bid according to the prescribed strategy. We derive a lower bound on gains from following the optimal taxation policy:

Proposition 39. *Under the optimal taxation policy, a government can gain at least*

$$n(1-t)\int_{\alpha_L}^{\alpha_U}(F^{n-1}(\alpha) - F^{n-1}((1-t)\alpha))(1-F(\alpha))d\alpha$$

compared to flat taxation.

Coming back to the possibility of a buyer generating profit below v_α, we find that it is never profitable for a buyer to plan at the bidding stage to use this opportunity after the auction is over[15]. In fact, the only possibility we need to consider is when a buyer who pretends to be of a lower type α' submits a bid $\beta(\alpha')$ and then produces α'. This results in the ex-ante utility of a type α', $u(\alpha')$ that is lower than the ex-ante utility of a type α, $u(\alpha)$. In all other cases, the deviation is detectable and therefore the maximum taxation t is imposed, making it unprofitable.

[15] Note that ex-post the winner may in fact find it optimal to produce below v_α

3.6 Conclusions

In this paper, we study an optimal auction design problem under the assumption that a buyer's private information is verified ex-post and that a seller could impose limited punishments. We have derived an optimal penalty rule for a first-price auction and have shown that small penalties have a significant impact on the seller's revenue. Before an auction begins, the seller makes recommendations regarding how much each type of buyer should bid in the auction. If the auction winner bids at least as much as the seller has prescribed, he is never punished; if he does not bid that much, the penalty increases as the buyer deviates more and more from the recommendation. The optimal recommendation is monotonic—the buyer with the highest value wins the auction, provided his bid meets a reserve price—and a buyer bids more than in a standard first-price auction. When the seller sets a reserve price, she sets it lower than in standard auctions.

We have shown that when a seller is restricted to using monotonic recommendations, a first-price auction is the optimal auction format among all winner-pay auctions. We have further proven that under certain conditions, a first-price auction is revenue superior to a second-price auction. The revenue comparison comes from the fact that a monotonic equilibrium strategy, which leads to a desirable buyers expected payments, may fail to exist in a second-price auction. This finding contrasts sharply with existing auction literature, as under a wide range of assumptions, the dominant strategy for a buyer in a standard second-price auction is to bid his own value. In our environment, a buyer bids more than his value and the existence of monotonic equilibrium is not guaranteed. We conjecture that this revenue comparison holds more generally to incorporate the situations, when the seller is not restricted to using monotonic recommendations.

Our results offer, perhaps, a surprising suggestion for policymakers: A government can increase the competition when auctioning publicly owned assets by providing privileges to buyers who submit sufficiently high bids. This advice, however, should be implemented cautiously, as the exact privileges depend on the nature of information available to the government after the auction. Therefore, such policies should be used only in cases in which the primitives of the environment have been studied in-depth and are well-understood. We have illustrated one such possible application in Section 3.5, where we investigate an optimal taxation policy of a newly privatized firm.

We consider our paper as a first step towards understanding how monetary transfers and limited punishments work together in environments in which a seller has access to additional information regarding buyers' valuations ex-post. The natural extensions of our work include considering a model in which buyers are imperfectly informed prior to an auction or buyers' values are correlated or drawn from the different distributions. We leave these and other possible extensions for future research.

4

Soliciting Information from Biased Experts (with Attila Ambrus)

4.1 Introduction

In Chapter 2, we analyzed a model that introduces competition among experts into a delegation framework. In particular, we considered a setting with principal who has to delegate a task to one of two experts. First, the experts receive noisy and conditionally independent signals of a single dimensional state variable. The principal's ideal action is equal to the state, but each expert has a constant bias and a resulting ideal point different from the sender's. Next, the experts simultaneously propose actions. A proposal is assumed to bind the expert to perform the given action whenever the principal delegates the task to him. The principal then chooses one of the two offers, and the corresponding action is taken by the given expert.

While the game form we adopted in the Chapter 2 fits various applications, in other settings a principal has the means to solicit information from experts more efficiently. For this reason, here we pursue a general mechanism design in similar situations. We assume there is a finite number of experts, who each obtain noisy observations of a state variable. These signals can be conditionally independent on the state, or correlated. Before the experts receive their signals, the principal can commit to a mechanism that specifies the action choice conditional on the recommendation profile, and also a monetary transfer to each participating expert, that is again can be conditioned on the recommendations of all participating experts. We investigate two versions of the model. In the version with ex post state verification, we also allow the monetary transfers (but not the action, which has to be chosen before state verification) to depend on the realized state. In the version with no ex post state verification monetary transfers are only allowed to be conditioned on the vector of recommendations. In both versions of the model we assume that there is a lower bound on the amount of transfer from the expert to the principal (limited liability or liquidity constraint). Each expert decides whether to participate in the mechanism or receive an outside option payoff not influenced by what happens afterwards. Then the state realizes and the experts receive their private signals. Next, each participating expert sends a recommendation to the principal. Action choice and monetary payments are then implemented by the mechanism.

If the state space is discrete, we show that under a full rank information condition (from an arbitrary expert's probability distribution over the state space/other experts' signals, we can uniquely identify a corresponding probability distribution over received signals), that holds generically, the principal can solicit all information from experts if the liquidity constraint does not bind. The optimal mechanism we provide implements the socially optimal action, which maximizes the sum of all participating agents' utilities, while keeping experts at their outside options. In such a way, the principal extracts all surplus, and further we refer to such equilibria as full extraction.

If the state space is a continuum, again provided that the minimum transfer from the principal to the expert is not too high, we show that the full extraction is possible if experts' signals are independent, and payments can be conditioned on the realized state. In the model without ex post state verification, we look for the full extraction mechanisms which make experts indifferent among all possible recommendations. The existence of such mechanisms is equivalent to the existence of a solution of a Fredholm equation of first kind, and the corresponding conditions for the latter are provided. Also, in spirit of Chapter 2, we consider a quadratic-diffuse-normal example and show that in both versions of the model, full extraction is possible. In addition, if we consider a limited liability case - we require all payments to be not less than the specified threshold, and if this threshold is low enough, full extraction is still possible.

We also show that if the minimum transfer from the principal to the agent is high enough to bind then it is no longer true that the optimal action always implements the socially optimal action. In some states the action specified by the mechanism gets distorted relative to it, so as to reduce the monetary transfers from the expert in other states. Nevertheless, contingent monetary payments above the minimum transfer can still be required in some states in the optimal mechanism.

Our results are similar to the classical full extraction results by Cremer and McLean (1985), Cremer and McLean (1988) and McAfee and Reny (1992), although the contexts are different. In our setting besides the question of how much money the principal can extract from (or has to pay to) experts, the question of what actions get implemented by the optimal mechanism are of primary interest as well. Moreover, as in our setting there is a partial common interest between the principal and the experts (all of their optimal actions are increasing in the state), transfers from the principal to the agent can be bounded from below in a mechanism extracting the full surplus, and we explicitly characterize the lower bound in some cases when the state space is the whole real line.

The problem we investigate is also similar to classical principal-agent problems with hidden information. In particular, our model is similar to the full commitment case of Krishna and Morgan (2008), however, they assume that the expert has perfect information, which makes their problem fundamentally different from ours (their techniques cannot be applied in our context).

Finally, the analysis of the version of the model with ex post state verification is related to the literature on scoring rules (see for example Brier (1950), Savage (1971), Winkler et al. (1996) and Kilgour and Gerchak (2004). However, this literature assumes that experts do not have policy preferences over the action choice, they do not adopt a mechanism design approach, and they typically assume that experts have immutable beliefs as opposed to making Bayesian inferences.

4.2 Model

We consider the following multi-stage game with incomplete information. There are $n+1$ players: a principal and n experts. The set of possible states of the world is Ω. The probability distribution of state of the world θ is common prior with cumulative distribution function F. The principal takes action a, and her policy preferences are represented by $U_P(a,\theta)$, while policy preferences for expert i are given by $U_i(a,\theta)$.

Stage 1: The principal commits to her future action choice as a function of all further received experts' recommendations (reports): $a = \phi(r_1, .., r_n)$, $\phi : (\Omega \cup \{\varnothing\})^n \to \mathbb{R}$, and makes a take-or-leave offer to every expert i ($i \in \{1, .., n\}$) - a payment function $m_i(r_i, r_{-i}, \theta) : (\Omega \cup \{\varnothing\})^n \times \Omega \to \mathbb{R}$, which depends on further expert i's own recommendation r_i, as well as vector of all other players' recommendations r_{-i} and realized state θ - amount that is paid by principal to expert i at the end of the game, after all recommendations are observed.

Stage 2: Every expert accepts or rejects his received offer. If expert i *does not accept* m_i, he quits, receives his outside option u_{0i} and *does not participate* in the further game. For convenience, we denote his recommendation as $r_i = \varnothing$ and denote

$m_i(\varnothing, r_{-i}) = 0$.

Stage 3: $\theta \in \Omega$ realizes.

Stage 4: Each expert i receives a noisy private signal about the state of the world s_i, where $s = (s_1, .., s_n)$ is distributed with cumulative distribution function G_θ.[1]

Stage 5: Every participating expert i proposes a recommendation (report) $r_i \in \mathbb{R}$ to the principal.

Stage 6: The principal observes the recommendations and state θ, implements pre-committed $a = \phi(r_1, .., r_n)$, and fulfills her payment obligations. The players' payoffs are quasilinear with respect to monetary payments. The principal's utility is $U_P(a, \theta) - \sum_{i=1}^{n} m_i(r_i, r_{-i}, \theta)$, while each participating expert i gets $U_i(a, \theta) + m_i(r_i, r_{-i}, \theta)$.

4.2.1 Model Without Ex Post State Verification

In many environments, the realized state is non-contractible, non-verifiable, or just verification is prohibitively costly. In such situations, payment functions m_i cannot depend on θ, that is, $m_i = m_i(r_i, r_{-i})$. Further, we refer to such a model as *model without ex post state verification*. Naturally, this model is of our special interest.

4.3 Socially Optimal Equilibrium and Sufficient Condition for Its Existence

First, note that by Revelation Principle, we can concentrate on *direct mechanisms*, that is, incentive compatible perfect Bayesian equilibria in which experts truthfully recommend their own signals: $r_i = s_i$ for every i.

We call a (Bayesian Nash) equilibrium *socially optimal equilibrium*, if each participating expert recommends his own signal $r_i = s_i$ and the principal takes a *socially optimal action* - an action that, conditional on the principal's full information (her

[1] Note that we do not require signals to be independent.

prior and participating experts' revealed signals), maximizes the sum of her and participating experts' expected utilities.

Lemma 40. *If for any set of participating experts a socially optimal equilibrium where each participating expert i gets utility equal to his outside option u_{0i} exists, then one of these equilibria is the principal-optimal equilibrium.*

Proof. Assume contrary, that for any set of participating experts, a socially optimal equilibrium where each participating expert i gets utility equal to his outside option u_{0i} exists, but none of them is the principal-optimal. This means, there exists an equilibrium profile σ, in which the principal's utility is higher than in the before mentioned equilibria.

For the set of experts participating in σ, there exists a social optimal equilibrium σ^* where each participating expert i gets utility equal to his outside option u_{0i}. But in the social optimal equilibrium σ^* the sum of the principal's and the participating experts' expected utilities is maximized, and in any equilibrium each participating expert gets at least his outside option value. Therefore, in σ^* the principal gets at least as much as in σ, and we get a contradiction. □

Hence, in the presence of the mentioned equilibria, the principal-optimal equilibrium payoff is the maximum among her payoffs in these 2^n equilibria. If the principal-optimal equilibrium is a socially optimal equilibrium where each participating expert i gets utility equal to his outside option u_{0i}, we call it a *full extraction* equilibrium and refer to this situation as full extraction.

Denote by $g_i^A(r, s)$ an expected payoff of player i when he recommends $r_i = r$ after observing signal $s_i = s$, while A is the set of participating experts, all of them recommend their own signals ($r_j = s_j$ for all $j \in A$ and $j \neq i$), and the principal chooses the socially optimal action. If the principal can choose payments such that for every A, $i \in A$ and $r_i \in \Omega$, $g_i^A(r_i, s_i) \leqslant g_i^A(s_i, s_i) = u_{0i}$, then she can achieve full extraction.

4.4 Model with Ex Post State Verification

In this section, we analyze the model with ex post state verification, that is, m_i can be conditioned on θ.

First, we show that if the principal can condition her payments to any expert i on vector of all experts' reports $r = (r_i, r_{-i})$, full extraction is possible. Moreover, the principal can fully insure all experts, that is for any realization of (θ, s, r_i), expert i gets exactly his outside option, u_{i0}.

Proposition 41. *If the principal can commit to payment function being conditional on vector of all experts' reports $r = (r_1, .., r_n)$ and the realized state θ, then the principal achieves full extraction.*

Proof. Let us fix any participation set, A. The principal commits to choose socially optimal action; she proposes $m_i(r_i, \theta) = u_{i0} - U_i(\phi(r_i, r_{-i}), \theta)$ to all experts in A and $m_i(r_i, \theta) = u_{i0} - U_i(\phi(r_i, r_{-i}), \theta) - 1$ for all other experts (not in A). Then all experts besides those in A choose not to participate and it is a best response for every expert in A to participate and report his own signal.

Finally, going through all possible subsets of $\{1, .., n\}$, principal chooses the principal-optimal equilibrium and achieves full extraction. □

In the following subsections, we analyze the model when payments to expert i are pure individual performance-based, that is, can be conditioned on his own report r_i and realized state θ only: $m_i = m_i(r_i, \theta)$.

4.4.1 Discrete Case

In this section, we consider the discrete case: $\Omega = \{1, ..., S\}$. The probability distribution of state of the world θ is common prior: $\mathbb{P}[\theta = s] = \pi_s$, $\pi_s \geqslant 0$, $\sum_{s=1}^{S} \pi_s = 1$. The vector of experts' signals $s = (s_1, .., s_n)$ is distributed with probability mass function $f_\theta : \Omega^n \to [0, 1]$: $\mathbb{P}(s = (x_1, .., x_n)|\theta) = f_\theta(x_1, .., x_n)$.

To find the expected payoff $g_i^A(r_i, s_i)$, let us denote by $P_i'(\theta|s_i)$ a probability that expert i assigns to the state being θ after observing his own signal s_i. Further, denote by q_{i,s_i} a $1 \times S$ vector of such probabilities written in increasing order:

$$q_{i,s_i} = (P_i'(1|s_i), .., P_i'(S|s_i)),$$

and by Q_i a $S \times S$ matrix, k-th row of which is $q_{i,k}$ $(1 \leqslant k \leqslant S)$.

Lemma 42. *Take any $A \subseteq \{1, .., n\}$. If for every $i \in A$ matrix Q_i is non-singular, then for any $g_i^0(r_i, s_i) \in \mathbb{R}^{|\Omega|^2}$, there exists a payment function $m_i(r_i, \theta)$ such that $g_i^A(r_i, s_i) \equiv g_i^0(r_i, s_i)$.*

Proof.

$$g_i^A(r_i, s_i) = \sum_{\theta \in \Omega} P_i'(\theta|s_i) m_i(r_i, \theta) + \mathbb{E}[U_i(a, \theta)|A, s_i, a = \phi(r_i, s_{-i})]$$

Define $g_{i,r_i}^A = (g_i^A(r_i, 1), .., g_i^A(r_i, S))'$, $Pol_{i,r_i,s_i}^A = \mathbb{E}[U_i(a, \theta)|A, s_i, a = \phi(r_i, s_{-i})]$, $Pol_{i,r_i} = (Pol_{i,r_i,1}, .., Pol_{i,r_i,S})'$. Finally, denote by m_{i,r_i} a $1 \times S$ vector of monetary payments $m_i(r_i, \theta)$ written according to increasing order in θ:

$$m_{i,r_i} = (m_i(r_i, 1), .., m_i(r_i, S)).$$

Then for any fixed r_i, $g_i^A(r_i, s_i) \equiv g_i^0(r_i, s_i)$ is equivalent to

$$Q_i \times m_{i,r_i} + Pol_{i,r_i}^A = g_{i,r_i}^0.$$

As matrix Q_i is non-singular, for any $g^0_{i,r_i} \in \mathbb{R}^{|\Omega|}$, there is a vector $m_{i,r_i} \in \mathbb{R}^{\Omega^{m-1}}$ such that the matrix equation above is satisfied:

$$m_{i,r_i} = Q_i^{-1}[g^0_{i,r_i} - Pol^A_{i,r_i}].$$

Going through all $r_i \in \Omega$, and joining these cases, we get that for any $g^0_i(r_i, s_i) \in \mathbb{R}^{|\Omega|^2}$, there exists a payment function $m_i(r_i, \theta) : \Omega^2 \to \mathbb{R}$ such that $g_i(r_i, s_i) \equiv g^0_i(r_i, s_i)$. □

Proposition 43. *If for every expert i matrix Q_i is non-singular, then the principal can achieve full extraction.*

Proof. Let us fix any participation set, A. For every expert $i \in A$ take $g^0_i(r_i, s_i) \equiv u_{i0}$. By Lemma 42, for each $i \in A$ there exists a payment function $m^*_i(r_i, \theta)$ such that $g^A_i(r_i, s_i) \equiv g^0_i(r_i, s_i) \equiv u_{i0}$. Then, for any $A \subseteq \{1, .., n\}$, there exists a socially optimal equilibrium with participating set A, where the principal commits to the socially optimal action and the payment functions $m^*_i(r_i, \theta)$, all participating experts recommend $r_i = s_i$ and get u_{i0} in expectation.

Finally, the principal chooses the principal-optimal equilibrium by comparing all subsets of $\{1, .., n\}$ and achieves full extraction. □

Q_i is non-singular if and only if the rows of Q_i are linearly independent. This means that from an arbitrary expert i's probability distribution over Ω, we can uniquely identify his corresponding probability distribution over received signals (although expert i gets only one signal).

Our results here are fully analogous to classical spanning conditions of Cremer and McLean (1985) and Cremer and McLean (1988).

4.4.2 Continuous Case: Independent Signals

In this section, we consider the continuous case: $\Omega = \mathbb{R}$, F and G_θ are continuous cumulative distribution functions. Also, assume that signals $\{s_i\}_{i=1}^n$ are conditionally independent given θ.

The following proposition shows that in the case of independent signals the full extraction is possible.

Proposition 44. *If signals are conditionally independent, then the principal can achieve full extraction.*

Proof. Take any participating set A. Without loss of generality, let $A = \{1, .., k\}$. The principal maximizes the total welfare of herself and all participating experts: based on a vector of recommendations $r = (r_1, ..., r_n)$, she chooses a socially optimal action - action $a = \phi(r)$, that maximizes

$$\mathbb{E}\left[U_P(a,\theta) + \sum_{i \in A} U_i(a,\theta) \middle| s_1 = r_1, ..., s_k = r_k\right]. \tag{4.1}$$

Denote a conditional density of θ given expert i's signal s_i by $h_i(\cdot|s_i)$.

If all other experts report truthfully, expert i's expected utility after observing $s_i = s$ and reporting $r_i = r$ is:

$$g_i^A(s,r) = \mathbb{E}[U_i(a,\theta) + m_i(r_i,\theta)|A, a = \phi(r_i, r_{-i}), s_i = s, r_i = r, r_{-i} = s_{-i}]$$

$$= \mathbb{E}[U_i(\phi(r, s_{-i}),\theta) + m_i(r,\theta)|A, s_i = s] = \int_{\mathbb{R}} \left[V_i^A(r,x) + m_i(r,x)\right] h_i(x|s)dx,$$

where $V_i^A(r,x) = \mathbb{E}[U_i(\phi(r,s_{-i}),x)|A, \theta = x] = \mathbb{E}[U_i(\phi(r,s_{-i}),x)|A, s_i = s, \theta = x]$.

If the principal proposes payment function $m_i(r,\theta) \equiv u_{0i} - V_i^A(r,\theta)$, then:

$$V_i^A(r,x) + m_i(r,x) \equiv u_{0i},$$

and, consequently,

$$g_i^A(s_i, r_i) \equiv u_{i0},$$

making the truthful reporting optimal for every expert $i \in A$. □

Note also that under payment functions from the proof above the experts are indifferent among all possible recommendations.

However, if signals are correlated, then

$$g_i^A(s,r) = \int_{\mathbb{R}} \left[V_i^A(s,r,x) + m_i(r,x)\right] h_i(x|s)dx,$$

$$V_i^A(s,r,x) \equiv \mathbb{E}[U_i(\phi(r,s_{-i}),x)|A,s_i = s,\theta = x] \neq \mathbb{E}[U_i(\phi(r,s_{-i}),x)|A,\theta = x],$$

$V_i^A(s,r,x)$ is not constant in s and the proof above cannot be directly applied.

4.4.3 Continuous Case: Correlated Signals with 2 Experts

Consider the case when number of experts $n = 2$ and their signals are correlated conditional on θ. Here, we concentrate on mechanisms that make experts indifferent among all possible recommendations.

Conditional on all other experts report truthfully, expert i's expected utility after observing $s_i = s$ and reporting $r_i = r$ is:

$$g_i^A(s,r) = \int_{\mathbb{R}} \left[V_i^A(s,r,x) + m_i(r,x)\right] h_i(x|s)dx,$$

where

$$V_i^A(s,r,x) = \mathbb{E}[U_i(\phi(r,s_{-i}),x)|A,s_i = s,\theta = x]$$

In order to get an expected payoff function $g_i^A(s,r) \equiv u_{i0}$, for every r, we need to solve a Fredholm equation of first kind,

$$\int_{\mathbb{R}} m_i(r,x)h_i(x|s)dx = u_{i0} - \int_{\mathbb{R}} V_i^A(s,r,x)h_i(x|s)dx,$$

with kernel function $h_i(x|s)$ and unknown $m_i(r,x)$.

Square Integrable Conditional Densities

If $h_i(x|s)$ **is square-integrable**, that is $\int_{\mathbb{R}}\int_{\mathbb{R}} h_i^2(x|s)dxds < +\infty$, then from Hilbert-Schmidt theory $h_i(x|s)$ can be represented as

$$h_i(x|s) = \sum_{i=1}^{\infty} \lambda_i v_i(x) w_i(s),$$

where v_i, w_i and λ_i are its left and right singular functions and singular values, correspondingly (see, for e.g., Groetsch (2007), p. 10).

Then, according to Picard's Criterion (Groetsch (2007), p. 11), if $P_i(r,s) \equiv u_{i0} - \int_{\mathbb{R}} V_i(s,r,x)h_i(x|s)dx \in L_2(-\infty,+\infty)$ a solution exists if and only and

$$\sum_{i=1}^{\infty} \frac{(P_i(r,s), w_i(s))^2}{\lambda_i^2} < +\infty,$$

where

$$(P_i(r,s), w_i(s)) = \int_{\mathbb{R}} P_i(r,s) w_i(s) ds.$$

If this condition holds, the solution can be written as

$$m_i(r,x) = \sum_{i=1}^{\infty} \frac{(P_i(r,s), w_i(s))}{\lambda_i} v_i(x).$$

4.4.4 Quadratic-diffuse-normal example

In spirit of Chapter 2, we consider the following case: $U_P(a,\theta) = -(a-\theta)^2$ and $U_i(a,\theta) = -\alpha_i(a-\theta-b_i)^2$, common prior is uniform improper (diffuse), and noises are normally distributed: $\varepsilon_i \sim N(0,\sigma^2)$ with common correlation $Corr(\varepsilon_i,\varepsilon_j) = \rho \geqslant 0$. Here $\alpha_j \geqslant 0$ represents a level of expert j's interest in the implemented decision and we normalize principal's level of interest to be 1.

The next proposition, in particular, states that in any socially optimal equilibrium with a fixed set of participating experts, the principal's action is the same function of signals.

Proposition 45. *In any socially optimal equilibrium with participating set* $S \subseteq \{1, 2, .., n\}$*:* $r_i = s_i$ *for any* $i \in S$ *and* $a = \frac{1}{|S|}\sum_{i\in S} r_i + b_S^*$*, where* $b_S^* = \frac{\sum_{i\in S}\alpha_i b_i}{1+\sum_{i\in S}\alpha_i}$.

Proof. In any socially optimal equilibrium with experts offering their signals: $r_i = s_i$, the principal maximizes the sum of her and participating experts' expected utilities, conditional on r_i, $i \in S$. In other words, she chooses a to maximize

$$\mathbb{E}[-(a-\theta)^2 - \sum_{i\in S}\alpha_i(a-\theta-b_i)^2|\{r_i\}_{i\in S}] =$$

$$-\left(1+\sum_{i\in S}\alpha_i\right)a^2 + \left[2(1+\sum_{i\in S}\alpha_i)\mathbb{E}[\theta|\{r_i\}_{i\in S}] + 2\sum_{i\in S}\alpha_i b_i\right]a$$

$$-(1+\sum_{i\in S}\alpha_i)\mathbb{E}[\theta^2|\{r_i\}_{i\in S}] - 2\sum_{i\in S}\alpha_i b_i\mathbb{E}[\theta|\{r_i\}_{i\in S}] - \sum_{i\in S}\alpha_i b_i^2$$

As $\mathbb{E}(\theta|\{r_i\}_{i\in S}) = \frac{1}{|S|}\sum_{i\in S} r_i$, maximizing this expression with respect to a gives

$$a^* = \mathbb{E}[\theta|\{r_i\}_{i\in S}] + \frac{\sum_{i\in S}\alpha_i b_i}{1+\sum_{i\in S}\alpha_i} = \frac{1}{|S|}\sum_{i\in S} r_i + \frac{\sum_{i\in S}\alpha_i b_i}{1+\sum_{i\in S}\alpha_i}$$

□

Corollary 46. *In any socially optimal equilibrium with experts getting utilities equal to their outside options, the principal's utility is*

$$-(b_S^*)^2 - \sum_{i\in S}\alpha_i(b_S^*-b_i)^2 - \left(1+\sum_{i\in S}\alpha_i\right)\frac{(|S|+(|S|-1)|S|\rho)\sigma^2}{|S|^2} - \sum_{i\in S}u_{0i}$$

As can be seen, the principal's utility in a socially optimal equilibrium is decreasing in correlation coefficient ρ, variance of noises σ^2 and experts' outside options u_{0i}.

Next, recall that in order to find a principal-optimal equilibrium, it is sufficient to show the existence of socially optimal equilibria.

Proposition 47. *For any of $2^n - 1$ non-empty sets of participating experts, a socially optimal equilibrium where participating experts get utilities equal to their outside options exists.*

Proof. Fix a set of participating experts, S. The principal commits to choose a socially optimal action, she proposes payment functions $m_i(r_i, \theta) = \frac{2}{|S|}\alpha_i(b_S^* - b_i)(r_i - \theta) + u_{0i} + \alpha_i[(b_i - b_S^*)^2 + \frac{\sigma^2}{|S|}(1 + (|S| - 1)\rho)]$ to $i \in S$ and $m_i(r_i, \theta) = u_{0i}$ to $i \notin S$;[2] all experts from S report truthfully: $r_i = s_i$, the remaining experts do not participate if the principal acts according to the described strategy, all experts do not participate otherwise.

Notice that the principal faces infinite losses from deviation as then no expert participates, therefore, to show that this profile is a socially optimal equilibrium, it is enough to check that the truthful report, $r_i = s_i$, is a best response for each expert i to other players' strategies. Since $a - \theta - b_i = \frac{1}{|S|}(r_i + (|S| - 1)\theta + \sum_{j \in S, j \neq i} \epsilon_j) + b_S^* - \theta - b_i = \frac{1}{|S|}(r_i - \theta + \sum_{j \in S, j \neq i} \epsilon_j) + b_S^* - b_i = \frac{1}{|S|}(r_i - s_i + \sum_{j \in S} \epsilon_j) + b_S^* - b_i$, then $(a - \theta - b_i | s_i) \sim N(\frac{1}{|S|}(r_i - s_i) + b_S^* - b_i, \frac{\sigma^2}{|S|}(1 + (|S| - 1)\rho))$ and participating expert i obtains

$$
\begin{aligned}
&\mathbb{E}[-\alpha_i \left(a - \theta - b_i\right)^2 + m_i(r_i, \theta) \, | s_i] = \\
&-\alpha_i \left[\left(\frac{1}{|S|} r_i - \frac{1}{|S|} s_i + b_S^* - b_i \right)^2 + \frac{\sigma^2}{|S|} \left(1 + (|S| - 1)\rho\right) \right] \\
&+ \frac{2}{|S|}\alpha_i(b_S^* - b_i)(r_i - s_i) + u_{0i} + \alpha_i \left[(b_i - b_S^*)^2 + \frac{\sigma^2}{|S|} \left(1 + (|S| - 1)\rho\right) \right] \\
&= u_{0i} - \frac{1}{|S|^2}\alpha_i(r_i - s_i)^2.
\end{aligned}
$$

Therefore, for each expert i, to participate and choose $r_i = s_i$ is a best response. As a consequence, the considered strategy profile is indeed a socially optimal equilibrium

[2] These payment functions are not unique, see limited liability case for another example.

where all experts' expected utilities are equal to their outside options. □

Notice that in these socially optimal equilibria the principal commits to her strategies from Proposition 45. To get her optimal equilibrium, the principal just chooses the best of these $2^n - 1$ equilibria. Therefore, she achieves full extraction.

4.4.5 Non-Binding Limited Liability for the Quadratic-diffuse-normal example

Let us return to the quadratic-diffuse-normal example. Up to this moment we assumed the unlimited liability on the part of the experts. Now we suppose that $m_i(r_i, r_{-i}, \theta) \geqslant l_i$, that is, for any offer r_i and realized state θ payment to the expert i is at least l_i (no matter what recommendation he gave and the realized state).

In the following proposition, we show that if l_i is not larger than u_{0i}, the socially optimal equilibria still exist.

Proposition 48. *If $l_i \leqslant u_{0i}$, then the principal can achieve full extraction.*

Proof. Fix a set of participating experts, S. The principal commits to choose $a = \frac{1}{|S|}\sum_{i\in S} r_i + b_S^*$ with $b_S^* = \frac{\sum_{i\in S}\alpha_i b_i}{1+\sum_{i\in S}\alpha_i}$ if the set of participating experts is S and to choose $a = 0$ otherwise, she proposes payment functions $m_i(r_i, \theta) = \frac{1}{|S|^2}\alpha_i(r_i - \theta)^2 + \frac{2}{|S|}\alpha_i(b_S^* - b_i)(r_i - \theta) + u_{0i} + \alpha_i[(b_i - b_S^*)^2 + \frac{(|S|-1)\sigma^2}{|S|^2}(1 + (|S|)\rho)]$ and $m_i(r_i, \theta) = u_{0i}$ for $i \notin S$; all experts from S report truthfully: $r_i = s_i$, the remaining experts do not participate if the principal acts according to the described strategy, all experts do not participate otherwise.

Notice that the principal faces infinite losses from deviation as then no expert participates, therefore, to show that this profile is a socially optimal equilibrium it is enough to check that the truthful report is a best response for each expert i to other player's strategies. Since $a - \theta - b_i = \frac{1}{|S|}(r_i + (|S|-1)\theta + \sum_{j\in S, j\neq i}\epsilon_j) + b_S^* - \theta - b_i = \frac{1}{|S|}(r_i - s_i + \sum_{j\in S}\epsilon_j) + b_S^* - b_i$, then $(a - \theta - b_i|s_i) \sim N(\frac{1}{|S|}(r_i - s_i) + b_S^* - b_i, \frac{\sigma^2}{|S|}(1 +$

$(|S|-1)\rho))$ and participating expert i obtains

$$\mathbb{E}[-\alpha_i\left(a-\theta-b_i\right)^2+m_i(r_i,\theta)\,|s_i]=$$

$$-\alpha_i\left[\left(\frac{1}{|S|}r_i-\frac{1}{|S|}s_i+b_S^*-b_i\right)^2+\frac{\sigma^2}{|S|}\left(1+(|S|-1)\rho\right)\right]$$

$$+\frac{1}{|S|^2}\alpha_i\left[(r_i-s_i)^2+\sigma^2\right]+\frac{2}{|S|}\alpha_i(b_S^*-b_i)(r_i-s_i)+u_{0i}$$

$$+\alpha_i\left[(b_i-b_S^*)^2+\frac{(|S|-1)(1+|S|\rho)\sigma^2}{|S|^2}\right]=u_{0i}$$

Therefore, any expert i's strategy yields him u_{0i} and the considered strategy profile is indeed a socially optimal equilibrium.

Finally check the limited liability requirement. For payment function m_i, denoting $r_i-\theta\equiv z$ we obtain $m_i(z)=\frac{1}{|S|^2}\alpha_i z^2+\frac{2}{|S|}\alpha_i(b_S^*-b_i)z+u_{0i}+\alpha_i[(b_i-b_S^*)^2+\frac{(|S|-1)\sigma^2}{|S|^2}(1+|S|\rho)]$, the global minimum of m_i is $u_{0i}+\frac{|S|-1}{|S|^2}(1+|S|\rho)\alpha_i\sigma^2\geqslant u_{0i}+\frac{|S|-1}{|S|^2}\alpha_i\sigma^2\geqslant l_i$. □

4.4.6 Binding Limited Liability

Consider the simplest case with one expert and 2-state space, $\Omega=\{0,1\}$. The common prior over states is $\pi_0=\pi_1=\frac{1}{2}$. Both the principal and the expert face quadratic losses: the principal utility is $U_P(a,\theta)=-(a-\theta)^2$, while the expert's preferences are represented by $U_1=-\alpha(a-\theta-b)^2$. Without loss of generality, assume that the expert is biased upwards, that is, $b\geqslant 0$.

The principal's signal is equal to the true state with probability $p>\frac{1}{2}$: $\mathbb{P}(s_1=\theta)=p$ and $\mathbb{P}(s_1=1-\theta)=1-p$. Denote by m_{ij} a payment that the expert gets from the principal if he recommends i, while the realized state is j. Limited liability here means that minimum payment the expert gets for his advise is $l\geqslant 0$, implying $m_{ij}\geqslant l$. With the limited liability on payments, we can consider the case when expert's Individual Rationality constraint is not binding (his outside option is low enough) - formally, we *drop Stage 2 of the game*. However, we allow the principal not to hire the expert: in this case, she just chooses $a = \frac{1}{2}$ and her utility is $-\frac{1}{4}$ - that is the level we need to compare with.

Recall that if the principal chooses to hire the expert, Revelation Principle implies that we can concentrate on direct mechanisms. The principal solves the following problem:

$$\max_{a_0,a_1,m_{00},m_{01},m_{10},m_{11}} W = \frac{1}{2}p(-a_0^2 - m_{00}) + \frac{1}{2}(1-p)(-a_1^2 - m_{10})$$
$$+ \frac{1}{2}(1-p)(-(a_0-1)^2 - m_{01}) + \frac{1}{2}p(-(a_1-1)^2 - m_{11})$$

s.t.

$$p(-\alpha(a_0-b)^2 + m_{00}) + (1-p)(-\alpha(a_0-1-b)^2 + m_{01})$$
$$\geqslant p(-\alpha(a_1-b)^2 + m_{10}) + (1-p)(-\alpha(a_1-1-b)^2 + m_{11})$$
$$p(-\alpha(a_1-1-b)^2 + m_{11}) + (1-p)(-\alpha(a_1-b)^2 + m_{10})$$
$$\geqslant p(-\alpha(a_0-1-b)^2 + m_{01}) + (1-p)(-\alpha(a_0-b)^2 + m_{00})$$
$$m_{00} \geqslant l$$
$$m_{01} \geqslant l$$
$$m_{10} \geqslant l$$
$$m_{11} \geqslant l$$

Recall that the principal hires the expert if $W > \frac{1}{4}$ only. Using the method of Lagrange multipliers, we arrive to the next results.

Proposition 49. *The principal hires the expert in one of the following cases:*
a) $0 \leqslant b \leqslant p - \frac{1}{2}$ *and* $l < (p-\frac{1}{2})^2$*. In this case,* $a_0 = 1-p$*,* $a_1 = p$ *and all the transfers are minimal:* $m_{00} = m_{01} = m_{10} = m_{11} = l$*.*
b) $p - \frac{1}{2} < b < \min\{2, 1+\alpha\}(p-\frac{1}{2})$ *and* $l < b(2p-1-b)$*. In this case,* $a_0 = \frac{3}{2} - 2p + b$*,*

$a_1 = \frac{1}{2} + b$ and no transfers are minimal: $m_{00} = m_{01} = m_{10} = m_{11} = l$.

c) $(1+\alpha)(p-\frac{1}{2}) \leqslant b < (1+\frac{1}{\alpha})(p-\frac{1}{2})$ and $l < \frac{\alpha^2}{1-\alpha^2}((1+\frac{1}{\alpha})(p-\frac{1}{2})-b)^2$, then $a_0 = 1-p+\frac{\alpha b}{1+\alpha}$, $a_1 = \frac{-\alpha+(1+\alpha)p-\alpha b}{1-\alpha}$, $m_{10} = m_{11} = l$, while m_{00} and m_{01} are any transfers that satisfy

$$\begin{cases} \frac{4\alpha^2}{1-\alpha^2}((1+\frac{1}{\alpha})(p-\frac{1}{2})-b) + m_{00} - m_{01} \geqslant 0 \\ pm_{00} + (1-p)m_{01} = l + \frac{4\alpha^2((1+\frac{1}{\alpha})(p-\frac{1}{2})-b)(b-(1+\alpha)(p-\frac{1}{2}))}{(1-\alpha^2)^2}. \end{cases}$$

In case a), the preferences of the agents are closely aligned: since b is small enough, the expert's preferred point after signal 0, $1-p+b$, is closer to $1-p$ than to p, implying he reports truthfully even without monetary incentives. We can call this solution *the first best* for the principal.

In case b), the principal cannot achieve her first best solution, she finds it beneficial to keep the expert indifferent between recommending 0 and 1 after signal 0. As b is not too high, she achieves this by distorting actions only while keeping minimum payments: a small increase in bias relative to a) makes the actions' distortions (which lead to losses quadratic in the bias increase) more beneficial than paying excessive monetary transfers (linear in the bias increase).

In case c), which exists for $\alpha < 1$ only, keeping the expert indifferent between recommending 0 and 1 after signal 0 by distorting actions only becomes too expensive for the principal. As a consequence, after recommendation 0 the principal chooses to pay some non-minimal transfers, while choosing the action $a_0 = 1-p+\frac{\alpha b}{1+\alpha}$ that maximizes the sum of her and the agent's utilities, which is optimal for the principal due to monetary balancing. At the same time, to minimize monetary payments (which are used to compensate the expert's utility difference between policies a_1 and a_0 after signal 0), the principal decreases a_1 as b increases - this distortion allows her to pay lower monetary transfers. Moreover, the decrease in a_1 makes the principal even better off for $b < 2p-1$ ($a_1 = 1-p$ for $b = 2p-1$), and for $b > 2p-1$ the policy losses are offset by decreased monetary transfers. As b approaches the right bound of the interval, $(1+\frac{1}{\alpha})(p-\frac{1}{2})$, both a_0 and a_1 tend to $\frac{1}{2}$, while transfers tend to 0: when bias reaches the bound, informative equilibrium no longer exists.

4.5 Model without Ex Post State Verification

In this section, we look whether the results of the previous section hold for the model without ex post state verification. Below, we see that the mentioned results do not extend only for the continuous case model with independent signals.

4.5.1 Discrete Case

In the discrete case, results of previous section hold as well for the model without ex post state verification. Here and further relevant omitted proofs are found in Appendix C.

Take any participation set A. Without loss of generality, let $A = \{1, .., m\}$. Denote by $P_i(s_{-i}|s_i)$ a probability that expert i assigns to a vector of other experts' signals being $s_{-i} = (s_1, .., s_{i-1}, s_{i+1}, .., s_m)$ after observing his own signal s_i. Further, denote by q_{i,s_i} a $1 \times S^{m-1}$ vector of such probabilities written according to increasing lexicographic order:

$$q_{i,s_i} = (P_i((1, .., 1, 1)|s_i), P_i((1, .., 1, 2)|s_i), .., P_i((S, .., S, S-1), P_i((S, .., S, S)|s_i)),$$

and by Q_i^m a $S \times S^{m-1}$ matrix, k-th row of which is $q_{i,k}$ $(1 \leqslant k \leqslant S)$.

Lemma 50. *If matrix Q_i^m has full rank (rank(Q_i^m)=S), then for any $g_i^0(r_i, s_i) \in \mathbb{R}^{|\Omega|^2}$, there exists a payment function $m_i(r_i, r_{-i})$ such that $g_i(r_i, s_i) \equiv g_i^0(r_i, s_i)$.*

Proposition 51. *If for every expert i $(1 \leqslant i \leqslant m)$ matrix Q_i^m has full rank (rank(Q_i^m) $= S$), then for a set of agents $\{1, .., m\}$ there exists a socially optimal equilibrium where each participating expert i gets utility equal to his outside option.*

Q_i is non-singular if and only if the rows of Q_i^m are linearly independent. This means that from an arbitrary expert i's probability distribution over the set of other experts' signals Ω^{m-1}, we can uniquely identify a corresponding probability distribution over the set of expert i's received signals (although expert i gets only one signal).

4.5.2 Continuous Case: Independent Signals

In this subsection, we show that the results for our base model are not automatically extended for the model without ex post state verification.

Recall that $h_i(\cdot|s_i)$ is a conditional density of θ given s_i. Also, denote by $\rho_i(\cdot|s_i)$ a conditional density of $s_{-i} \in \mathbb{R}^{k-1}$ given s_i.

If all other experts report truthfully, expert i's utility after observing $s_i = s$ and reporting $r_i = r$ is:

$$\begin{aligned}
g_i(s,r) &= \mathbb{E}[U_i(a,\theta) + m_i(r_i, r_{-i})|a = \phi(r_i, r_{-i}), s_i = s, r_i = r, r_{-i} = s_{-i}] \\
&= \mathbb{E}[U_i(\phi(r, s_{-i}), \theta) + m_i(r, s_{-i})|s_i = s] \\
&= \int_{\mathbb{R}^{k-1}} \mathbb{E}[U_i(\phi(r, s_{-i}), \theta) + m_i(r, s_{-i})|s_i = s, s_{-i} = x]\rho_i(x|s)dx \\
&= \int_{\mathbb{R}^{k-1}} \mathbb{E}[U_i(\phi(r, x), \theta) + m_i(r, x)|s_i = s, s_{-i} = x]\rho_i(x|s)dx,
\end{aligned}$$

As $W_i(s,r,x) \equiv \mathbb{E}[U_i(\phi(r,x),\theta)|s_i = s, s_{-i} = x] \neq \mathbb{E}[U_i(\phi(r,x),\theta)|s_{-i} = x]$, we cannot apply the approach from the previous subsection.

In order to get an expected payoff function $g_i(s,r) \equiv u_{i0}$, for every r, we need to solve a Fredholm equation of first kind,

$$\int_{\mathbb{R}^{k-1}} m_i(r,x)\rho_i(x|s)dx = u_{i0} - \int_{\mathbb{R}^{k-1}} W_i(s,r,x)\rho_i(x|s)dx,$$

with kernel function $\rho_i(x|s)$ and unknown $m_i(r,x)$.

Hence, the existence of payment functions under which the experts are indifferent among all possible recommendations is not guaranteed in the absence of ex post state verification.

4.5.3 Continuous Case: Correlated Signals with 2 Experts

In order to get an expected payoff function $g_i(s,r) \equiv u_{i0}$, for every r, we need to solve a Fredholm equation of first kind,

$$\int_{\mathbb{R}} m_i(r,x)\rho_i(x|s)dx = u_{i0} - \int_{\mathbb{R}} W_i(s,r,x)\rho_i(x|s)dx,$$

with kernel function $\rho_i(x|s)$ and unknown $m_i(r,x)$.

Square Integrable Conditional Density

If $\rho_i(x|s)$ **is square-integrable**, that is $\int_{\mathbb{R}}\int_{\mathbb{R}} \rho_i^2(x|s)dxds < +\infty$, using Hilbert-Schmidt theory $\rho_i(x|s)$ can be written as

$$\rho_i(x|s) = \sum_{i=1}^{\infty} \mu_i y_i(x) z_i(s),$$

where y_i, z_i and μ_i are its left and right singular functions and singular values, correspondingly (Groetsch (2007)).

Then, a solution exists if $Q_i(r,s) \equiv u_{i0} - \int_{\mathbb{R}} W_i(s,r,x)\rho_i(x|s)dx \in L_2(-\infty,+\infty)$ and

$$\sum_{i=1}^{\infty} \frac{(Q_i(r,s), z_i(s))^2}{\mu_i^2} < +\infty,$$

where

$$(Q_i(r,s), z_i(s)) = \int_{\mathbb{R}} Q_i(r,s) z_i(s) ds$$

If this condition holds, the solution can be written as

$$m_i(r,x) = \sum_{i=1}^{\infty} \frac{(Q_i(r,s), z_i(s))}{\mu_i} y_i(x).$$

4.5.4 Quadratic-diffuse-normal example

The results for the quadratic-diffuse-normal example hold even if payments are not dependent on the realized state θ.

Proposition 52. *In the model without ex post state verification, the principal can achieve full extraction.*

4.5.5 Non-Binding Limited Liability for the Quadratic-diffuse-normal example

In the following proposition, we show that if l_i is not larger than u_{0i}, the social optimal equilibrium in which all experts participate still exists in the absence ex post state verification.

Proposition 53. *In the model without ex post verification with $l_i \leqslant u_{0i}$ for every i, the principal can achieve full extraction.*

4.6 Conclusions

In this chapter, we allow the principal to reward experts conditional on their recommendations and, in some scenarios, on the realized state. We show that in the discrete case the full surplus extraction by principal is possible under the standard full rank condition. In continuous case, we show that the full extraction is possible if experts' signals are independent and payments can be conditioned on the realized state. In addition, we prove that full extraction is possible in the environment of Chapter 2 with monetary transfers, even if the latter cannot be conditioned on the realized state and the experts' signals are correlated. This result holds even with limited liability, if the minimal payment is below a specified threshold.

5

Conclusion

This dissertation consists of three self-contained chapters. The conclusions of the first, second and third chapters are in section 2.7, 3.6 and 4.6, respectively.

Appendix A

Appendices of Chapter 2

We begin with proofs for the general results of Section 2.5, after which many of the results of Section 2.4 can be handled readily as special cases.

A.1 Proofs for Section 2.5

We first provide an auxiliary lemma.

Lemma A.1.1. *Let* $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$*, where* ϵ_1 *and* ϵ_2 *are independent, and define* $\xi(k_1, k_2) := \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$*,* $\eta(k_1, k_2) := \max(\epsilon_1 + k_1, \epsilon_2 + k_2)$*. Then*

$$\mathbb{E}\xi(k_1, k_2) = -2\sigma^2 f(k_1 - k_2) + k_1(1 - F(k_1 - k_2)) + k_2 F(k_1 - k_2);$$

$$\mathbb{E}\eta(k_1, k_2) = 2\sigma^2 f(k_1 - k_2) + k_1 F(k_1 - k_2) + k_2(1 - F(k_1 - k_2));$$

$$\mathbb{E}\xi^2(k_1, k_2) = \sigma^2 - 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2(1 - F(k_1 - k_2)) + k_2^2 F(k_1 - k_2);$$

$$\mathbb{E}\eta^2(k_1, k_2) = \sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2 F(k_1 - k_2) + k_2^2(1 - F(k_1 - k_2)).$$

Proof. Let $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$ be independent random variables. By defining $\xi(k_1, k_2) =$

$\min(\epsilon_1 + k_1, \epsilon_2 + k_2)$ and $\eta(k_1, k_2) = \max(\epsilon_1 + k_1, \epsilon_2 + k_2)$, we have

$$\mathbb{E}\xi(k_1, k_2) = -2\sigma^2 f(k_1 - k_2) + k_1(1 - F(k_1 - k_2)) + k_2 F(k_1 - k_2);$$

$$\mathbb{E}\eta(k_1, k_2) = 2\sigma^2 f(k_1 - k_2) + k_1 F(k_1 - k_2) + k_2(1 - F(k_1 - k_2));$$

$$\mathbb{E}\xi^2(k_1, k_2) = \sigma^2 - 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2(1 - F(k_1 - k_2)) + k_2^2 F(k_1 - k_2);$$

$$\mathbb{E}\eta^2(k_1, k_2) = \sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) + k_1^2 F(k_1 - k_2) + k_2^2(1 - F(k_1 - k_2)).$$

Recall that f and F denote the PDF and the CDF of $N(0, 2\sigma^2)$. Let h and H denote the PDF and the CDF of the distribution $N(0, \sigma^2)$.

We start with a calculation of auxiliary integrals. The first two are directly computed by substitution:

$$I_1(k_1, k_2) = \int_{\mathbf{R}} h(t - k_1)h(t - k_2)dt = \left| z = t - \frac{k_1 + k_2}{2} \right| = f(k_1 - k_2)$$

$$I_2(k_1, k_2) = \int_{\mathbf{R}} th(t - k_1)h(t - k_2)dt = \left| z = t - \frac{k_1 + k_2}{2} \right| = \frac{1}{2}(k_1 + k_2)f(k_1 - k_2)$$

The third one is calculated by parts:

$$I_3(k_1, k_2) = \int_{\mathbf{R}} H(t - k_1)dh(t - k_2) = H(t - k_1)h(t - k_2)|_{t=-\infty}^{t=\infty} - I_1(k_1, k_2)$$
$$= -f(k_1 - k_2)$$

The fourth one

$$I_4(k_1, k_2) = \int_{\mathbf{R}} H(t - k_1)h(t - k_2)dt = \int_{\mathbf{R}} H(t + k_2 - k_1)h(t)dt := \phi(k_2 - k_1)$$

$0 \leqslant \phi(z) \leqslant 1$ and $\phi(z)$ is well-defined. Then for any z: $\phi'(z) = I_1(z, 0) = f(z)$. Integrating back, we get $\phi(z) = F(z) + C$. Integrating by parts, we receive $\phi(z) = 1 - \phi(-z)$. These equalities imply $C = 0$ and $\phi(z) = F(z)$. As a consequence, $I_4(k_1, k_2) = F(k_2 - k_1) = 1 - F(k_1 - k_2)$.

Finally, compute two integrals, that will be used for the calculations of RV moments:

$$I_5(k_1,k_2) = \int_{\mathbf{R}} tH(t-k_1)h(t-k_2)dt$$
$$= -\sigma^2 I_3(k_1,k_2) + k_2 I_4(k_1,k_2)$$
$$= \sigma^2 f(k_1-k_2) + k_2(1-F(k_1-k_2))$$

$$I_6(k_1,k_2) = \int_{\mathbf{R}} t^2 H(t-k_1)h(t-k_2)dt$$
$$= -\sigma^2 \int_{\mathbf{R}} tH(t-k_1)dh(t-k_2) + k_2 I_5(k_1,k_2)$$
$$= \sigma^2[I_2(k_1,k_2) + I_4(k_1,k_2)] + k_2 I_5(k_1,k_2)$$
$$= \frac{1}{2}(k_1+3k_2)\sigma^2 f(k_1-k_2) + (k_2^2+\sigma^2)(1-F(k_1-k_2))$$

Now, return to the RVs. The PDF and the CDF of η are given by:

$$F_\eta(t) = \Pr(\eta < t) = \Pr(\epsilon_1 + k_1 < t, \epsilon_2 + k_2 < t) = H(t-k_1)H(t-k_2)$$
$$f_\eta(t) = H(t-k_1)h(t-k_2) + H(t-k_2)h(t-k_1).$$

Using the obtained density function, calculate the moments of η:

$$\mathbb{E}\eta = \int_{\mathbf{R}} tf_\eta(t)dt = I_5(k_1,k_2) + I_5(k_2,k_1) = 2\sigma^2 f(k_1-k_2) + k_1 F(k_1-k_2)$$
$$+k_2(1-F(k_1-k_2));$$

$$\mathbb{E}\eta^2 = \int_{\mathbf{R}} t^2 f_\eta(t)dt = I_6(k_1,k_2) + I_6(k_2,k_1)$$
$$= \sigma^2 + 2(k_1+k_2)\sigma^2 f(k_1-k_2) + k_1^2 F(k_1-k_2) + k_2^2(1-F(k_1-k_2)).$$

Finally, from the identities $\xi+\eta \equiv \epsilon_1+k_1+\epsilon_2+k_2$ and $\xi^2+\eta^2 \equiv (\epsilon_1+k_1)^2+(\epsilon_2+k_2)^2$, we obtain the moments for ξ:

$$\mathbb{E}\xi = \mathbb{E}(\epsilon_1 + k_1 + \epsilon_2 + k_2 - \eta) = k_1 + k_2 - \mathbb{E}\eta$$
$$= -2\sigma^2 f(k_1-k_2) + k_1(1-F(k_1-k_2)) + k_2 F(k_1-k_2);$$

$$\mathbb{E}\xi^2 = \mathbb{E}[(\epsilon_1+k_1)^2 + (\epsilon_2+k_2)^2 - \eta^2] = 2\sigma^2 + k_1^2 + k_2^2 - \mathbb{E}\eta^2$$
$$= \sigma^2 - 2(k_1+k_2)\sigma^2 f(k_1-k_2) + k_1^2(1-F(k_1-k_2)) + k_2^2 F(k_1-k_2).$$

$\square$

Proof of Lemma 11. After observing the signal s_i, expert i does a Bayesian update of his beliefs:

$\theta|s_i \sim N(s_i, \sigma^2)$ and $s_j|s_i \sim N(s_i, 2\sigma^2)$.

Since the principal chooses the lower offer, she accepts a_i iff $s_j > s_i + k_i - k_j$. Denote by g the PDF of $N(s_i, 2\sigma^2)$. Hence, the expected utility of expert i

$$U_i(k_1, k_2, L) = \int_{s_i+k_i-k_j}^{\infty} \mathbb{E}\left[B - (a_i - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j$$

$$+ \int_{-\infty}^{s_i+k_i-k_j} \mathbb{E}\left[-(a_j - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j$$

As $(\theta|s_i, s_j) \sim N(\frac{s_j+s_j}{2}, \frac{\sigma^2}{2})$, $a_i = s_i + k_i$, $a_j = s_j + k_j$, we obtain

$$U_i(k_1, k_2, L) = \int\limits_{s_i+k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j$$

$$+ \int\limits_{-\infty}^{s_i+k_i-k_j} \left[-\left(k_j - b_i + \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j.$$

Now make a substitution $t = s_j - s_i$ and denote by f and F the PDF and CDF of $N(0, 2\sigma^2)$.

$$U_i(k_1, k_2, L) = \int\limits_{k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt$$

$$+ \int\limits_{-\infty}^{k_i-k_j} \left[-\left(k_j - b_i + \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt.$$

Note that $U_i(k_1, k_2, L)$ does not depend on signal s_i, which is intuitive for the improper prior. As $\int\limits_a^{\infty} tf(t)\, dt = 2\sigma^2 f(a)$ and $\int\limits_{-\infty}^{\infty} t^2 f(t)\, dt = 2\sigma^2$, we get the expression for $U_i(k_i, k_j, L)$.

Now in state θ, the principal's action a is distributed as $\theta+\xi$, where $\xi = \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$; $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$, ϵ_1 and ϵ_2 are independent.

Therefore, from Lemma A.1.1 the expected bias of the accepted offer is

$$\bar{b}(k_1, k_2, L) = \mathbb{E}\xi(k_1, k_2) = -2\sigma^2 f(k_1 - k_2) + k_2 F(k_1 - k_2) + k_1(1 - F(k_1 - k_2))$$

and the expected utility of the principal is

$$\begin{aligned} V(k_1, k_2, L) &= -\mathbb{E}(a - \theta)^2 = -\mathbb{E}(\theta + \xi - \theta)^2 = -\mathbb{E}\xi^2(k_1, k_2) \\ &= -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) - k_2^2 - (k_1^2 - k_2^2)(1 - F(k_1 - k_2)). \end{aligned}$$

Finally, the variance of the chosen offer is

$$\begin{aligned} Var(k_1, k_2, L) =& - V(k_1, k_2, L) - \bar{b}^2(k_1, k_2, L) \\ =& \sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z)). \end{aligned}$$

□

The following lemma provides several useful bounds. The statements are immediate corollaries of Sampford (1953).

Lemma 54. *The following inequalities hold for all $x \in \mathbb{R}$:*

- $0 < v'(x) < \frac{1}{2\sigma^2}$;
- $0 > w'(x) > -\frac{1}{2\sigma^2}$;
- $v''(x) > 0$.

Proof of Theorem 7. We start by showing that $U_i(k_1, k_2, L)$ is a single-peaked function of k_i. Taking a derivative w.r.t. k_i yields

$$\begin{aligned} U_i'(k_i) &= -2\left[(k_i - b_i)(1 - F(k_i - k_j)) - \rho f(k_i - k_j)\right] \\ &= -2(1 - F(k_i - k_j))\left[k_i - b_i - \rho\frac{f(k_i - k_j)}{1 - F(k_i - k_j)}\right]. \end{aligned}$$

Let $g(k_i)$ denote the term in square brackets above. Lemma 54 implies that $g'(k_i) = 1 - \rho\lambda(k_i - k_j) \geqslant 1 - \sigma^2\lambda(k_i - k_j) > 0$. Additionally, we have $\lim_{x\to\pm\infty} g(x) = \pm\infty$. Combining these facts, U_i has a unique critical point, which is a global maximum.

We now look for min equilibria. The FOCs for the experts are equivalent to:

$$k_1 - b_1 - \rho\frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} = 0 \tag{A.1}$$

$$k_2 - b_2 - \rho\frac{f(k_1 - k_2)}{F(k_1 - k_2)} = 0. \tag{A.2}$$

Subtracting (A.2) from (A.1) substituting $z = k_1 - k_2$, we get

$$z - \rho\left[\frac{f(z)}{1 - F(z)} - \frac{f(z)}{F(z)}\right] = b_1 - b_2. \tag{A.3}$$

Denote $l(z) = -\rho\left[\frac{f(z)}{1-F(z)} - \frac{f(z)}{F(z)}\right] + z$. Using a) from Lemma 54, we obtain

$$l'(z) = 1 - \rho\left[v'(z) - w'(z)\right] + 1 \geqslant 1 - \sigma^2[v'(z) - w'(z)] > 0$$

Now $l(z)$ is continuous, strictly increasing on $\mathbb{R}$, and ranges from $-\infty$ to $+\infty$. Therefore (2.4) has a unique solution, z^*; we use $z(B)$ to denote explicitly the dependence on B.

Using this solution, we get (k_1^m, k_2^m) as the only critical point and check that this point satisfies both initial FOCs. As $U_i(k_i, k_j, L)$ is a single-peaked function of k_i, (k_1^m, k_2^m) is a pair of best responses.

As it was shown in Theorem 1, choosing the lower offer is the BR strategy for the principal iff $k_1 + k_2 \geqslant 0$, or equivalently

$$b_1 + b_2 + \rho\left[\frac{f(z^*)}{1 - F(z^*)} + \frac{f(z^*)}{F(z^*)}\right] \geqslant 0.$$

Also the LHS of (2.4) is equal to 0 at $z = 0$, and therefore $z^* \geqslant 0$ and $k_1^m - k_2^m \geqslant 0$.

Define a function $m(B) = b_1 + b_2 + \rho[v(z(B)) + w(z(B)]$; the min equilibrium exists if and only if $m(B) \geqslant 0$.

1) For $B \leqslant 2\sigma^2$: $m(B) \geqslant 0$, therefore the min equilibrium exists.

2) Next, we show that $m(B)$ is decreasing in B in the region $B \geqslant 2\sigma^2$.

$$m'(B) = -\frac{1}{2}[v(z(B)) + w(z(B))] + \rho[\lambda(z(B)) - \lambda(-z(B))]z'(B). \tag{A.4}$$

Differentiating equation (2.4) at point B, we get:

$$z'(B) - \rho[\lambda(z(B)) + \lambda(-z(B))]z'(B) + \frac{1}{2}[v(z(B)) - w(z(B))] = 0. \tag{A.5}$$

By substituting (A.5), the second term of (A.4) becomes

$$+\rho[\lambda(z(B)) - \lambda(-z(B))]\frac{-\frac{1}{2}[v(z(B)) - w(z(B))]}{1 - \rho\left[\lambda(z(B)) + \lambda(-z(B))\right]}$$

$$\leqslant -\rho[\lambda(z(B)) - \lambda(-z(B))]\frac{\frac{1}{2}[v(z(B)) - w(z(B))]}{-\rho\left[\lambda(z(B)) + \lambda(-z(B))\right]}$$

$$= \frac{1}{2}\frac{\lambda(z(B)) - \lambda(-z(B))}{\lambda(z(B)) + \lambda(-z(B))}[v(z(B)) - w(z(B))]$$

$$\implies m'(B) \leqslant -\frac{1}{2}[v(z(B)) + w(z(B))] + \frac{1}{2}[v(z(B)) - w(z(B))] = -w(z(B)) < 0.$$

3) From Lemma 54 the hazard rate v is convex, so for any real x, $v(x) + w(x) = v(x) + v(-x) \geqslant 2v(0) > 0$, and $m(B)$ tends to $-\infty$ as B tends to ∞.

From 1)-3) follows that there exists $B_m : m(B) \geqslant 0$ iff $B \leqslant B_m$. Also

$$\left(\frac{B_m}{2} - \sigma^2\right)[v(z(B_m)) + w(z(B_m))] = b_1 + b_2. \tag{A.6}$$

As $z(B_m)$ satisfies equation (2.4), we have:

$$\left(\frac{B_m}{2} - \sigma^2\right)[v(z(B_m)) - w(z(B_m))] + z(B_m) = b_1 - b_2. \tag{A.7}$$

From the previous discussion and (A.6) we have $B_m \geqslant 2\sigma^2$. Also, (A.6) and the inequality $v(x) + w(x) = v(x) + v(-x) \geqslant 2v(0) = \frac{2}{\sqrt{\pi}\sigma}$ give an upper bound on B_m:

$$\left(\frac{B_m}{2} - \sigma^2\right)\frac{2}{\sqrt{\pi}\sigma} \leqslant b_1 + b_2.$$

Subtracting (A.7) from (A.6), we get a lower bound on B_m:

$$2b_2 = (B_m - 2\sigma^2)w(z(B_m)) - z(B_m) \leqslant (B_m - 2\sigma^2)w(0) = (B_m - 2\sigma^2)\frac{1}{\sqrt{\pi}\sigma}.$$

Finally, we calculate the expected bias of the chosen offer, its variance, and players' utilities:

$$\begin{aligned}\bar{b}(k_1^m, k_2^m, L) &= -2\sigma^2 f(z^*) + k_2^m F(z^*) + k_1^m(1 - F(z^*)) \\ &= -2\sigma^2 f(z^*) + b_2 F(z^*) + \rho f(z^*) + b_1(1 - F(z^*)) + \rho f(z^*) \\ &= b_1(1 - F(z^*)) + b_2 F(z^*) - Bf(z^*);\end{aligned}$$

$$\begin{aligned}Var(k_1^m, k_2^m, L) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) \\ + (z^*)^2 F(z^*)(1 - F(z^*)).\end{aligned}$$

□

Corollary A.1.1. *In the min equilibrium,*

$$\begin{aligned}V(k_1^m, k_2^m, L) &= -\sigma^2 - b_1^2(1 - F(z^*)) - b_2^2 F(z^*) + B(b_1 + b_2)f(z^*) + \Delta(z^*) \\ U_1(k_1^m, k_2^m, L) &= -\sigma^2 - (b_1 - b_2)^2 F(z^*) + B(1 - F(z^*)) - B(b_1 - b_2)f(z^*) + \Delta(z^*) \\ U_2(k_1^m, k_2^m, L) &= -\sigma^2 - (b_1 - b_2)^2(1 - F(z^*)) + BF(z^*) + B(b_1 - b_2)f(z^*) + \Delta(z^*),\end{aligned}$$

where $\Delta(z^*) := \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$.

Proof. Immediate from applying Lemma 11 to the markups given by Theorem 7. □

Proof of Lemma 12. Since the principal chooses the highest offer, she chooses a_i iff $s_j < s_i + k_i - k_j$. Using arguments similar to used in Lemma 11, we find the expected utility of expert i:

$$U_i(k_1,k_2,H) = \int_{k_i-k_j}^{\infty} \mathbb{E}[-(a_j-\theta-b_i)^2|s_j]f(s_j)\,ds_j$$

$$+\int_{-\infty}^{k_i-k_j} \mathbb{E}[B-(a_i-\theta-b_i)^2|s_j]f(s_j)\,ds_j = \int_{k_i-k_j}^{\infty}\left[-\left(k_j-b_i+\frac{s_j}{2}\right)^2-\frac{\sigma^2}{2}\right]f(s_j)\,ds_j$$

$$+\int_{-\infty}^{k_i-k_j}\left[B-\left(k_i-b_i-\frac{s_j}{2}\right)^2-\frac{\sigma^2}{2}\right]f(s_j)\,ds_j = (B-(k_i-b_i)^2)F(k_i-k_j)-\sigma^2$$

$$-(k_j-b_i)^2[1-F(k_i-k_j)]-2\sigma^2(k_i+k_j-2b_i)f(k_i-k_j).$$

In state θ the principal's action a is distributed as $\theta+\eta$, where $\eta \sim \max(\epsilon_1+k_1,\epsilon_2+k_2)$; $\epsilon_1,\epsilon_2 \sim N(0,\sigma^2)$, ϵ_1 and ϵ_2 are independent.

From Lemma A.1.1 the expected bias of the accepted offer is

$$\bar{b}(k_1,k_2,H) = \mathbb{E}\eta(k_1,k_2) = 2\sigma^2 f(k_1-k_2)+k_2(1-F(k_1-k_2))+k_1F(k_1-k_2).$$

The expected utility of the principal is

$$\begin{aligned} V(k_1,k_2,H) &= -\mathbb{E}(a-\theta)^2 = -\mathbb{E}\eta^2(k_1,k_2) \\ &= -\sigma^2-2(k_1+k_2)\sigma^2 f(k_1-k_2)-k_2^2-(k_1^2-k_2^2)F(k_1-k_2). \end{aligned}$$

The variance of the chosen offer is

$$\begin{aligned} Var(k_1,k_2,H) =& -V(k_1,k_2,H)-\bar{b}^2(k_1,k_2,H) \\ =&\sigma^2-4\sigma^4f^2(z)-2\sigma^2zf(z)(2F(z)-1)+z^2F(z)(1-F(z)). \end{aligned}$$

□

Proof of Theorem 8. The proof is analogous to that of Theorem 7. The FOCs for experts are now:

$$k_1-b_1+\rho\frac{f(k_1-k_2)}{F(k_1-k_2)} = 0 \tag{A.8}$$

$$k_2-b_2+\rho\frac{f(k_1-k_2)}{1-F(k_1-k_2)} = 0. \tag{A.9}$$

Subtracting equation (A.8) from equation (A.9) yields (A.3). Principal optimality holds if and only if $k_1 + k_2 \leqslant 0$, or equivalently

$$b_1 + b_2 - \rho \left[\frac{f(z^*)}{1 - F(z^*)} + \frac{f(z^*)}{F(z^*)} \right] \leqslant 0, \tag{A.10}$$

Define a function $n(B) = b_1 + b_2 - \rho[v(z(B)) + w(z(B)]$, where $z(B)$ is given by equation (2.4). For $B > 2\sigma^2$, we have $n(B) > 0$, and thus a max equilibrium does not exist. Observe further that $n(2\sigma^2) = b_1 + b_2 \geqslant 0$. Since $m(B) + n(B) = 2(b_1 + b_2)$ and $m'(B) < 0$, we have $n'(B) > 0$. It follows that if $n(0) \leqslant 0$, then there exists $B_M \in [0, 2\sigma^2]$ such that $n(B) \leqslant 0$ iff $B \leqslant B_M$. Therefore

$$\left(\frac{B_M}{2} - \sigma^2 \right) [v(z(B_M)) + w(z(B_M))] = -(b_1 + b_2). \tag{A.11}$$

Also $z(B_M)$ satisfies equation (2.4), and therefore

$$\left(\frac{B_M}{2} - \sigma^2 \right) [v(z(B_M)) - w(z(B_M))] + z(B_M) = b_1 - b_2. \tag{A.12}$$

From the previous discussion and (A.11) we have $B_M \leqslant 2\sigma^2$. Also (A.11) and the inequality $v(x) + w(x) \geqslant 2v(0) = \frac{2}{\sqrt{\pi}\sigma}$ give the lower bound

$$\left(\frac{B_M}{2} - \sigma^2 \right) \frac{2}{\sqrt{\pi}\sigma} \geqslant -(b_1 + b_2).$$

Summing (A.12) and (A.11), we get the upper bound

$$-2b_2 = \left(\frac{B_M}{2} - \sigma^2 \right) v(z(B_M)) + z(B_M) \geqslant (B_M - 2\sigma^2) 2v(0).$$

Finally, we compute the following:

$$\begin{aligned}
\bar{b}(k_1^M, k_2^M, H) &= 2\sigma^2 f(z^*) + k_1^M F(z^*) + k_2^M (1 - F(z^*)) \\
&= 2\sigma^2 f(z^*) + b_1 F(z^*) - \rho f(z^*) + b_2(1 - F(z^*)) - \rho f(z^*) \\
&= b_1 F(z^*) + b_2(1 - F(z^*)) + B f(z^*);
\end{aligned}$$

$$Var(k_1^M, k_2^M, H) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1)$$
$$+(z^*)^2 F(z^*)(1 - F(z^*)).$$

□

Recall the definition of $\Delta(z^*)$ from Corollary A.1.1.

Corollary A.1.2. *In the max equilibrium,*

$$V(k_1^M, k_2^M, H) = -\sigma^2 - b_1^2 F(z^*) - b_2^2(1 - F(z^*)) - B(b_1 + b_2) f(z^*) + \Delta(z^*)$$
$$U_1(k_1^M, k_2^M, H) = -\sigma^2 - (b_1 - b_2)^2(1 - F(z^*)) + BF(z^*) + B(b_1 - b_2) f(z^*) + \Delta(z^*)$$
$$U_2(k_1^M, k_2^M, H) = -\sigma^2 - (b_1 - b_2)^2 F(z^*) + B(1 - F(z^*)) - B(b_1 - b_2) f(z^*) + \Delta(z^*).$$

Proof. Lemma 12 applied to Theorem 8. □

Proof of Proposition 13. From Theorems 7 and 8

$$V_{min} - V_{max} = (2F(z^*) - 1)(b_1^2 - b_2^2) + 2B(b_1 + b_2) f(z^*) \geqslant 0,$$

with equality if and only if either $b_1 + b_2 = 0$ or both $B = 0$ and $b_1 = b_2$. □

Proof of Proposition 14. Delegation to expert 2 alone yields principal utility $V^S(b, x) = -\sigma^2 - (b - x)^2$, while min equilibrium yields

$$V^m(b, x) = -\sigma^2 - b^2 - x^2 + (2F(z^*) - 1)2bx + (k_1 - b - x)(k_2 - b + x).$$

The difference between these is

$$V^m(b, x) - V^S(b, x) = -2(1 - F(z^*))bx + (k_1 - b - x)(k_2 - b + x)$$
$$= -2(1 - F(z^*))bx + \sigma^4 v(z^*) w(z^*). \quad \text{(A.13)}$$

Recall that z^* is independent of b. For fixed $x > 0$ then the existence of b^* follows from the fact that this expression is decreasing linearly in b and positive for $b = 0$.

Next, consider $b_2 = 0$ and $b_1 = b > 0$. The principal's utility in the min equilibrium is

$$V = -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(z) - k_1^2(1 - F(z)) - k_2^2 F(z),$$

which we aim to show is greater than $-\sigma^2$ as under simple delegation to expert two. Using the expressions $k_1 = b + \sigma^2 \frac{f(z)}{1-F(z)}$ and $k_2 = \sigma^2 \frac{f(z)}{F(z)}$, this is true if and only if

$$2\left(b + \sigma^2 \frac{f}{1-F} + \sigma^2 \frac{f}{F}\right)\sigma^2 f > \left(b^2 + 2b\sigma^2 \frac{f}{1-F} + \sigma^4 \left(\frac{f}{1-F}\right)^2\right)(1-F)$$
$$+\sigma^4 \left(\frac{f}{F}\right)^2 F.$$

Using $b = z - \sigma^2 \frac{f(2F-1)}{F(1-F)}$ and simplifying, this is equivalent to

$$\sigma^4(4F-1) > z(zF(1-F) - 2\sigma^2 f(2F-1)).$$

As $z > 0$, the left hand side is positive; we now show that the right hand side is negative. Let $h(z) := 2\sigma^2 f(2F-1) - zF(1-F)$, which we aim to show is positive. Then $h'(z) = 2f^2 - F(1-F)$. As shown in Sampford (1953), $k(z) := \frac{f^2}{F(1-F)}$ is decreasing for $z \geqslant 0$. It is easy to verify that $2k(0) > 1$ and that $\lim_{z\to+\infty} k(z) = 0$. It follows that there is a unique positive solution to $h'(z) = 0$. It is also easy to check that $h'(0) > 0$ and that $\lim_{z\to+\infty} h(z) = 0$. Together these facts imply that $h(z) > 0$ for all $z > 0$, as desired. □

A.2 Proofs for Section 2.4

Proof of Proposition 3. If $b_1 = b_2 = b > 0$, then upper and lower bounds on B_m from Theorem 7 coincide, so $B_m = 2\sigma^2 + 2\sqrt{\pi}\sigma b$. The experts' markups are $k_1^m = k_2^m = k_m = b + \frac{\rho\sigma}{\sigma\sqrt{\pi}}$ and $z^* = 0$. □

Proof of Proposition 4. If $b_1 = b_2 = b > 0$, then upper and lower bounds on B_M from Theorem 8 coincide, so $B_M = 2\sigma^2 - 2\sqrt{\pi}\sigma b$. The experts' markups are $k_1^M = k_2^M = k_M = b - \frac{\rho}{\sigma\sqrt{\pi}}$ and $z^* = 0$. □

Proof of Proposition 5. The comparison of the min and max equilibria follows as a special case of Proposition 13. The comparison of the min equilibrium and simple delegation is immediate by inspection of (A.13) in the proof of Proposition 14 by setting $x = 0$. $\square$

Given $B > 0$, let σ^* denote the unique positive solution to the equation $(4\pi - 4)\sigma^4 + 2\sqrt{\pi}bB\sigma = B^2$.

Lemma 55. *Consider $b_1 = b_2 = b > 0$. The principal's payoff is related to the noise level σ as follows:*

- *If $0 < B < \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$, there exists a non-empty interval $\sigma \in [\frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2}, \sigma^*)$, where min equilibrium exists and the principal's expected payoff is increasing in σ; when $\sigma > \sigma^*$, the principal's expected payoff is decreasing in σ.*

- *If $B = 0$ or $B \geqslant \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$, the principal's expected payoff is always decreasing in σ.*

Proof. From Proposition 3, the symmetric min equilibrium exists if $B \leqslant 2\sigma^2 + 2\sqrt{\pi}b\sigma$ or, equivalently, $\sigma \geqslant \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2}$. The principal's expected payoff in the min equilibrium is equal to $V = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - (1-\frac{1}{\pi})\sigma^2$. Then $V'(\sigma) = \frac{B}{\sqrt{\pi}\sigma^2}(\frac{B}{2\sqrt{\pi}\sigma} - b) - 2(1-\frac{1}{\pi})\sigma = -\frac{1}{2\pi\sigma^3}[(4\pi-4)\sigma^4 + 2\sqrt{\pi}bB\sigma - B^2] > 0$ if and only if $\sigma < \sigma^*$.

If $B = 0$, $V'(\sigma) < 0$. Otherwise, denote $\sigma_0 = \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2} > 0$.

The interval $[\sigma_0, \sigma^*)$ is non-empty if and only if $0 > (4\pi-4)\sigma_0^4 + 2\sqrt{\pi}bB\sigma_0 - B^2 = (4\pi-4)\sigma_0^4 - 2B\sigma_0^2 + B[2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 - B] = (4\pi-4)\sigma_0^4 - 2B\sigma_0^2$ or, equivalently, $0 > (2\pi-2)\sigma_0^2 - B = (2\pi-2)\sigma_0^2 - (2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0) = (2\pi-4)\sigma_0^2 - 2\sqrt{\pi}b\sigma_0$. The latter holds if and only if $\sigma_0 < \frac{\sqrt{\pi}b}{\pi-2}$ or, equivalently, $B = 2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 < \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$. $\square$

Lemma 56. *Suppose that bonus B is paid by the principal. The principal's optimal choice of B depends on the bias-to-noise ratio as follows:*

- *If $\frac{b}{\sigma} \leqslant \sqrt{\pi}$, then the principal pays no bonus: $B = 0$.*

- *If $\frac{b}{\sigma} > \sqrt{\pi}$, then the principal chooses bonus $B = 2\sqrt{\pi}\sigma^2\left[\frac{b}{\sigma} - \sqrt{\pi}\right]$. The increase in the principal's expected payoff, relative to when the bonus is restricted to be 0, is equal to $(b - \sqrt{\pi}\sigma)^2$.*

Proof. By Proposition 13, only min equilibrium should be considered. Therefore, we seek to maximize the quadratic function $V(k_m, k_m, L) - B = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \sigma^2 + \frac{\sigma^2}{\pi} - B$ on the interval $B \in [0, 2\sigma^2 + 2\sqrt{\pi}\sigma b]$.

a) First consider $\frac{b}{\sigma} \leqslant \sqrt{\pi}$. In this case the maximum is achieved at $B = 0$ and is equal to $V(k_m, k_m, L|B = 0) = -b^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$.

Hence, if $\frac{b}{\sigma} < \sqrt{\pi}$, then the principal pays no bonus.

b) Now consider $\frac{b}{\sigma} \geqslant \sqrt{\pi}$. In this case the principal achieves maximum in the min equilibrium at $B_R = 2\sqrt{\pi}\sigma\left(b - \sqrt{\pi}\sigma\right)$ (min equilibrium exists for this point) and is equal to $V(k_m, k_m, L|B = B_R) = (\pi - 1 + \frac{1}{\pi})\sigma^2 - 2\sqrt{\pi}\sigma b$.

Her gains comparatively to $B = 0$ (if she is legally restricted from paying bonuses) are equal to:

$$V(k_m, k_m, L|B = B_R) - V(k_m, k_m, L|B = 0) = \left(b - \sqrt{\pi}\sigma\right)^2.$$

□

Proof of Proposition 6. For nonmonotonicity in b, note that the principal's payoff in the min equilibrium is $-\bar{b}^2 - Var = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$, which is clearly increasing in b for $b \in \left[0, \frac{B}{2\sqrt{\pi}\sigma}\right]$. Nonmonotonicity in σ is established by Lemma 55, and nonmonotonicity in B is established by Lemma 56. □

A.3 Proofs for Section 2.6

Lemma 57. *Consider the game played between experts given a fixed principal strategy p. For $B \leqslant 2\sigma^2$, this game has a unique equilibrium in constant markup strategies,* *characterized by differentiable markup functions $k_1(p)$ and $k_2(p)$, with $k_1(p) + k_2(p)$ increasing.*

Proof. In what follows, define

$$v(z) := \frac{f(z)}{1 - F(z)}$$

$$w(z) := \frac{f(z)}{F(z)}$$

$$J_v(z; p) := \frac{1 - F(z)}{\frac{p}{2p-1} - F(z)}$$

$$J_w(z; p) := \frac{F(z)}{\frac{1-p}{2p-1} + F(z)}$$

$$H(z; p) := \frac{f(z)}{\frac{p}{2p-1} - F(z)} = v(z)J_v(z)$$

$$K(z; p) := \frac{f(z)}{\frac{1-p}{2p-1} + F(z)} = w(z)J_w(z)$$

where f and F are the PDF and CDF of $N(0, 2\sigma^2)$.

Note that for all $z \in \mathbb{R}$ and $p \in [0, 1]$, we have the identities $H(z; 1-p) \equiv -K(z; p)$ and $H(-z; p) \equiv K(z; p)$.

Given p, the experts' FOCs are

$$k_1 = b_1 + \left(\sigma^2 - \frac{B}{2}\right) \frac{f(z)(2p - 1)}{p(1 - F(z)) + (1 - p)F(z)} \tag{A.14}$$

$$= b_1 + \left(\sigma^2 - \frac{B}{2}\right) H(z; p) \tag{A.15}$$

$$k_2 = b_2 + \left(\sigma^2 - \frac{B}{2}\right) \frac{f(z)(2p - 1)}{pF(z) + (1 - p)(1 - F(z))} \tag{A.16}$$

$$= b_2 + \left(\sigma^2 - \frac{B}{2}\right) K(z; p), \tag{A.17}$$

where $z = k_1 - k_2$. Note that for $p = \frac{1}{2}$, we have simply $k_1 = b_1$ and $k_2 = b_2$, and

$H(z;p)$ and $K(z;p)$ are continuous in p for all $p \in [0,1]$. Henceforth, we consider $p \neq \frac{1}{2}$.

Combining (A.14) and (A.16) yields

$$D(z;p) := b_1 - b_2 + \left(\sigma^2 - \frac{B}{2}\right)[H(z;p) - K(z;p)] - z = 0. \tag{A.18}$$

It is easy to verify that the $D(z;p)$ is symmetric in p about $p = \frac{1}{2}$, so it suffices for the moment to consider only $p \in (\frac{1}{2}, 1)$. That a solution to (A.18) exists is evident upon examining the limiting behavior; as z approaches $+\infty$ or $-\infty$, $D(z;p)$ approaches $+\infty$ or $-\infty$, respectively. Every solution to (A.18) determines a unique pair (k_1, k_2) that solves (A.14) and (A.16). We now show that there is a unique solution z^* to (A.18); it suffices to show that the $D(z;p)$ has derivative strictly less than 0. We have $\frac{\partial H(z;p)}{\partial z} := v'(z)J_v(z) + v(z)J_v'(z)$. Now $v'(z) > 0$ by Lemma 54 and since $p \geqslant \frac{1}{2}$, $J_v(z) < 1$ and $J_v'(z) < 0$. Thus by Lemma 54, $\frac{\partial H(z;p)}{\partial z} < v'(z) < \frac{1}{2\sigma^2}$. By similar reasoning, $\frac{\partial K(z;p)}{\partial z} > w'(z) > -\frac{1}{2\sigma^2}$. If $\frac{\partial[H(z;p)-K(z;p)]}{\partial z} < 0$, the claim immediately follows, and if not, we have $\left(\sigma^2 - \frac{B}{2}\right)\frac{\partial[H(z;p)-K(z;p)]}{\partial z} < \sigma^2 \frac{2}{2\sigma^2} = 1$, as desired.

It remains to show that for the unique solution (k_1^*, k_2^*) to (A.14) and (A.16), each expert is indeed best-responding to the other. Consider two cases.

Case I: $p > \frac{1}{2}$. By the arguments above, the right side of (A.14) has derivative w.r.t. z (and hence k_1) less than 1. Since the labeling of experts is arbitrary, the second order condition is satisfied for each expert.

Case II: $p < \frac{1}{2}$. By the identity $H(z;1-p) = -K(z;p)$, (A.16) becomes $k_2 = b_2 - \left(\sigma^2 - \frac{B}{2}\right)H(z;1-p)$. By familiar arguments, the right side has derivative w.r.t. k_2 less than 1, and since experts are labeled arbitrarily, second order conditions hold.

Clearly, (A.14) and (A.16) are equations of differentiable functions of k_1, k_2, and p, and are nonconstant in the k_i. By the implicit function theorem, there exist differentiable functions $k_1(p), k_2(p)$ that solve these equations for each p. Finally, we show that $k_1(p) + k_2(p)$ is increasing. It is useful now to label players so that $b_1 > b_2$, and thus $z^*(p) > 0$. Combining (A.14) and (A.16), we obtain

$$S(k_1, k_2, z; p) := b_1 + b_2 + \left(\sigma^2 - \frac{B}{2}\right)[H(z;p) + K(z;p)] - (k_1 + k_2) = 0 \quad \text{(A.19)}$$

Note that $z^*(p)$ is decreasing in p for $p < \frac{1}{2}$ and increasing for $p > \frac{1}{2}$. To see this, consider (A.18). Since $D(z;p)$ is symmetric in p about $p = \frac{1}{2}$, it suffices to consider $p > \frac{1}{2}$. We have

$$\frac{\partial[H(z;p) - K(z;p)]}{\partial p} = \frac{f(z)}{(2p-1)^2}\left[\left(\frac{p}{2p-1} - F(z)\right)^{-2} - \left(\frac{1-p}{2p-1} + F(z)\right)^{-2}\right].$$

Now for all $z > 0$, $F(z) > \frac{1}{2}$ and thus $\frac{1-p}{2p-1} + F(z) > \frac{p}{2p-1} - F(z) > 0$. It follows that $\frac{\partial D(z;p)}{\partial p} > 0$. By earlier claims, $D(z;p)$ is decreasing in z, so $z^*(p)$ must be increasing in p for $p > \frac{1}{2}$. By symmetry, $z^*(p)$ is decreasing in p for $p < \frac{1}{2}$. Let $C = \sigma^2 - B/2 \geqslant 0$. By totally differentiating (A.18) w.r.t p, we obtain

$$z'(p) = \frac{C(H_p - K_p)}{1 - C(H_z - K_z)}. \quad \text{(A.20)}$$

From (A.19), we have

$$\begin{aligned} k_1'(p) + k_2'(p) &= C[(H_z + K_z)z'(p) + H_p + K_p] \\ &= C(H_z + K_z)\frac{C(H_p - K_p)}{1 - C(H_z - K_z)} + C(H_p + K_p). \end{aligned}$$

The above is nonnegative if and only if

$$\begin{aligned} &(H_z + K_z)C(H_p - K_p) + (H_p + K_p)(1 - C(H_z - K_z)) \geqslant 0 \\ \iff &H_p + K_p + 2CK_zH_p - 2CH_zK_p \geqslant 0 \quad \text{(A.21)} \end{aligned}$$

The LHS of (A.21) is at least $H_p + K_p + 2Cw'(z)H_p - 2Cv'(z)K_p$, which by Lemma 54 is positive, as desired.

For $p < \frac{1}{2}$, we use the substitution $p' = 1 - p$ and apply the identity $H(z;p) \equiv K(z;1-p)$; we obtain $k_1(p)+k_2(p)-b_1-b_2 = -(k_1(p')+k_2(p')-b_1-b_2)$. Differentiating both sides and noting that $\frac{dp'}{dp} = -1$, we get that $k_1'(p)+k_2'(p) = k_1'(1-p)+k_2'(1-p) > 0$. □

Proof of Proposition 16. The uniqueness of the equilibrium of the game played between experts given p comes from Lemma 57. For $b_1 = b_2 = b > 0$, the expression for the principal's payoff is the weighted average of payoffs from choosing the minimum or maximum offer, $pV(k,k,L)+(1-p)V(k,k,H)$, where $V(k_1,k_2,L)$ and $V(k_1,k_2,H)$ are defined in Lemmas 11 and 12, respectively. Using $B \leqslant 2\sigma^2$ and taking a derivative w.r.t. p, the payoff is increasing in p and maximized at $p = 1$. This establishes the first claim of the proposition.,

For the second claim, suppose $B = 0$, $b_1 = b > 0$ and $b_2 = -b$. Let $V(p)$ denote the principal's utility from commitment to p. That $b_1 + b_2 = 0$ implies $V(0) = V(1)$ is shown in the proof of Proposition 13. Therefore, it suffices to show that $V'(1) < 0$. Given p, the markups satisfy

$$k_1(p) = b + \frac{(2p-1)f(z(p))}{W(p)}$$

$$k_2(p) = -b + \frac{(2p-1)f(z(p))}{1 - W(p)}$$

$$\implies z(p) = 2b + \frac{(2p-1)f(z)(1-2W(p))}{W(p)(1-W(p))},$$

where $W(p) := p(1 - F(z(p))) + (1-p)F(z(p))$. Differentiating with respect to p and solving for $z'(1)$ yields

$$z'(1) = \frac{f(4F+1)(2F-1)}{F^2(1-F)^2 - F(1-F)(f'(2F-1)+2f^2) - f^2(2F-1)^2}.$$

The principal's utility is

$$V(p) = 2(2p-1)f(k_1+k_2) - k_1^2 W - k_2^2(1-W).$$

By differentiating with respect to p, evaluating at $p = 1$, substituting in the above expression for $z'(1)$ and simplifying, we obtain

$$V'(1) = \frac{f^2}{F^2(1-F)^2}\left(1 + \frac{g_1}{g_2}\right), \text{ where}$$
$$g_1 := (2F-1)(2F(1-F)f' + f^2(2F-1)),$$
$$g_2 := F^2(1-F)^2 - f^2(2F^2 - 2F + 1) - f'(2F-1)F(1-F).$$

We claim that $g_2 > 0$, and thus it suffices to show that $g_2 < -g_1$ for sufficiently large z. To see this, note that by Lemma 54, $f' < \frac{1-F}{2} - \frac{f^2}{1-F}$, and thus

$$g_2 > F(1-F)^2\left[\frac{1}{2} - \frac{f^2}{F(1-F)}\right].$$

It is easy to verify that $\frac{f^2}{F(1-F)} < \frac{1}{2}$ holds globally, and thus $g_2 > 0$ as desired. To see that $g_2 < -g_1$ for sufficiently large z, note that by algebra this comparison is equivalent to

$$F(1-F) + f'(2F-1) \leqslant 2f^2. \tag{A.22}$$

Using Lemma 54 again and simplifying, a sufficient condition for (A.22) is $2F - \frac{1}{2} < \frac{f^2}{(1-F)^2}$. The left hand side is bounded above by $\frac{3}{2}$, while the right hand side is increasing and unbounded above; thus for sufficiently large z, (A.22) holds. Finally, since z is increasing and unbounded above as a function of b, the second part of the proposition holds. □

Proof of Proposition 17. First, we calculate marginal utilities:

$$U_i'(k_i) = -2(k_i - b)[1 - F(k_i - k_j)] + \left[\frac{1}{2}\sigma^2 + 2\rho + \frac{1}{4}(k_i + k_j - 2b)^2\right] f(k_i - k_j).$$

Here, setting $U_i'(k) = 0$ gives two critical points:

$$k = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \tfrac{5}{2}\right)\sigma^2 + B} \text{ and } k = b + \sqrt{\pi}\sigma + \sqrt{\left(\pi - \tfrac{5}{2}\right)\sigma^2 + B}.$$

The second derivative is:

$$U_i''(k_i) = -2[1 - F(k_i - k_j)] + f(k_i - k_j)\times$$

$$\times \left[2(k_i - b) + \frac{1}{2}(k_i + k_j - 2b) + \left(\frac{B}{2\sigma^2} - \frac{5}{4}\right)(k_i - k_j) - \frac{1}{8\sigma^2}(k_i - k_j)(k_i + k_j - 2b)^2\right].$$

We get that only $k^* = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ is a local maximum of experts' utility functions.

Optimality for the principal holds if and only if $k^* \geqslant 0$ or, equivalently, $B \leqslant b^2 + 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2$.

Calculating, we get that $U_i(k^*, k^*, L) = \frac{(\pi-1)\sigma}{\sqrt{\pi}}\sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} - (\pi - \frac{7}{4})\sigma^2$. As we noted earlier, a necessary condition for equilibrium is $U_i(k^*, k^*, L) \geqslant 0$ or, equivalently, $B \geqslant (\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2})\sigma^2$. Therefore, min equilibrium may exist only if

$$B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi - 11)}{16(\pi - 1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right].$$

To finish the proof, we show that if B lies on this interval, then $k = k^*$ is a global maximum of $U_1(k, k^*, L)$.

Denote $g(k) = -2(k - b) + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] v(k - k^*)$.[1]

Then $U_1'(k) = -2(k - b)[1 - F(k - k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k - k^*) = [1 - F(k - k^*)]g(k)$ and $sign(U_1'(k)) = sign(g(k))$.

The first and second derivatives of g are

$$g'(k) = -2 + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] v'(k - k^*) + \frac{1}{2}(k + k^* - 2b)v(k - k^*)$$

$$g''(k) = \left[\frac{\sigma^2}{2} + 2\rho - B + \frac{1}{4}(k + k^* - 2b)^2\right] v''(k - k^*) + (k + k^* - 2b)v'(k - k^*)$$

$$+\frac{1}{2}v(k - k^*).$$

[1] Recall that $v(k - k^*) = \frac{f(k-k^*)}{1-F(k-k^*)})$.

Consider two cases.

1. $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2\right]$. Here $k^* \geqslant b$.

a) On the interval $k < b$, $U_1'(k) > 0$ and hence there is no point of maximum there.

b) On the interval $k \geqslant b$ we also have that $k + k^* - 2b \geqslant 0$. As all v, v' and v'' are strictly positive functions, $g''(k) > 0$. As $g'(k^*) < 0$ and $g'(+\infty) > 0$, hence there exists $k^{**} > k^*$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$, $g(k)$ is increasing. As $g(b) > 0$, $g(k^*) = 0$, $g(k^{**}) < 0$ and $g(+\infty) > 0$, then there exists $k_0 > k^{**} : g(k_0) = 0$. In summary, $g(k)$ is negative only on (k^*, k_0). Consequently, $U_1(k)$ is increasing on $[b, k^*)$, decreasing on (k^*, k_0), increasing for $k > k_0$. Hence, to show that k^* is a maximum on the interval $k \geqslant b$ it is sufficient to verify that $U_1(k^*) \geqslant U_1(+\infty) = 0$, which has already been done.

2. $B \in \left[\frac{5}{2}\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right]$. Here $k^* \leqslant b$.

a) On the interval $k < k^*$: $U_1'(k) > 2(b - k^*)[1 - F(k^* - k^*)]$

$+ \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^* + k^* - 2b)^2\right] f(k - k^*) = b - k^* - \frac{b-k^*}{f(0)} f(k - k^*)$

(as k^* is a solution of $\frac{5}{2}\sigma^2 - B + (k^* - b)^2 = \frac{k^*-b}{f(0)}$).

Therefore $U_1'(k) > b - k^* - \frac{b-k^*}{f(0)} f(k - k^*) \geqslant b - k^* - \frac{b-k^*}{f(0)} f(0) = 0$.

b) On the interval $k \in (k^*, b]$: $\frac{U_1'(k)}{f(k-k^*)} = -2(k - b)\frac{1-F(k-k^*)}{f(k-k^*)} +$

$+ \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] < 2(b - k^*)\frac{1-F(k^*-k^*)}{f(k^*-k^*)} + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^* + k^* - 2b)^2\right]$

$= \frac{b-k^*}{f(0)} - \frac{b-k^*}{f(0)} = 0$, hence $U_1'(k) < 0$.

c) On the interval $k \in \left(b, 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}\right]$ we also have

$\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2 \leqslant 0$. Then $U_1'(k) = -2(k - b)[1 - F(k - k^*)] +$

$+ \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k - k^*) < 0$.

d) On the interval $k > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$ we also have $k + k^* - 2b > 0$. Hence, on this interval $g''(k) > 0$. Also notice that $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then two cases are possible: (i) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) \geqslant 0$. Then on the whole interval $g'(k) > 0$ and $g(k)$ is increasing. As $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$ and $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Now $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$. Hence, to show that k^* is a global maximum it is enough to check that $U_1(k^*) \geqslant U_1(+\infty) = 0$, which has already been done.

(ii) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then there exists $k^{**} > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$ $g(k)$ is increasing. As $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Then $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$ and k^* is a global maximum as $U_1(k^*) \geqslant U_1(+\infty) = 0$.

We now verify that $k_m = k^* > k_m^{bas.}$:

$$k_m = b + \left(\sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}\right)\sigma > k_m^{bas.} = b + \left(1 - \frac{B}{2\sigma^2}\right)\frac{\sigma}{\sqrt{\pi}}$$

$$\iff \sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}} > \left(1 - \frac{B}{2\sigma^2}\right)\frac{1}{\sqrt{\pi}}$$

$$\iff \sqrt{\pi} + \frac{1}{\sqrt{\pi}}\left(\frac{B}{2\sigma^2} - 1\right) > \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$$

$$\iff \pi + \frac{B}{\sigma^2} - 2 + \frac{1}{\pi}\left(\frac{B}{2\sigma^2} - 1\right)^2 > \pi + \frac{B}{\sigma^2} - \frac{5}{2} \iff \frac{1}{2} + \left(\frac{B}{2\sigma^2} - 1\right)^2 > 0.$$

□

Proof of Proposition 18. Start with calculation of marginal utilities:

$$U_i'(k_i) = -2(k_i - b_i)F(k_i - k_j) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k_i + k_j - 2b)^2\right] f(k_i - k_j)$$

Consider the symmetric case: $k_1 = k_2 = k$. The FOCs give two critical points:

$k = b - \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ and $k = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$.

Next, calculate second derivatives:

$$U_i''(k_i) = -2F(k_i - k_j)f(k_i - k_j)\times$$

$$\left[2(k_i - b_i) + \frac{1}{2}(k_i + k_j - 2b_i) + \left(\frac{B}{2\sigma^2} - \frac{5}{4}\right)(k_i - k_j) - \frac{1}{8\sigma^2}(k_i - k_j)(k_i + k_j - 2b_i)^2\right]$$

We get that only $k^* = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ satisfies SOCs.

In order to satisfy principal optimality we need $k^* \leqslant 0$ or, equivalently, both $\frac{b}{\sigma} \leqslant \sqrt{\pi}$ and $B \leqslant b^2 - \sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2$.

Calculating, we get that $U_i(k^*, k^*, H) = \frac{(\pi-1)\sigma}{\sqrt{\pi}}\sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} - (\pi - \frac{7}{4})\sigma^2$ (the same as in the min equilibrium). As in the min equilibrium case, a necessary condition is $B \geqslant \left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2$.

From previous arguments max equilibrium may exist only if $\frac{b}{\sigma} \leqslant \sqrt{\pi}$ and

$$B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi - 11)}{16(\pi - 1)^2}\right)\sigma^2, b^2 - 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2\right).$$

This interval is non-empty if and only if $\frac{b}{\sigma} \leqslant \frac{3\sqrt{\pi}}{4(\pi-1)}$. Note also that $B \leqslant b^2 - 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2 \leqslant \frac{5}{2}\sigma^2$.

To finish the proof we show that if B lies on this interval, then $k = k^*$ is not only a local, but also a global maximum of $U_1(k, k^*, L)$.

Denote $r(k) = -2(k - b) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w(k - k^*)$ (remind that $w(k - k^*) = \frac{f(k-k^*)}{F(k-k^*)}$).

Then $U_1'(k) = -2(k - b)F(k - k^*) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k - k^*) = F(k - k^*)r(k)$ and $sign(U_1'(k)) = sign(r(k))$.

First and second derivatives of $r(k)$ are:

$r'(k) = -2 - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w'(k - k^*) - \frac{1}{2}(k + k^* - 2b)w(k - k^*)$;

$r''(k) = -\left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w''(k - k^*) - (k + k^* - 2b)w'(k - k^*) - \frac{1}{2}w(k - k^*)$.

Notice that as $\frac{b}{\sigma} \leqslant \sqrt{\pi}$, $B \leqslant b^2 - \sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2 \leqslant \frac{5}{2}\sigma^2$ and $k^* = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} \leqslant b$.

a) On interval $k > b$ $U'(k) < 0$, so there is no candidate for maximum there.

b) On interval $k < b$ we also have $k + k^* - 2b \leqslant 0$. As $w > 0$, $w' < 0$, $w'' > 0$, we have $r''(k) < 0$.

As $r'(-\infty) > 0$ and $r'(k^*) < 0$, there exists $k^{**} < k^* < b$: $r(k)$ is increasing for $k < k^{**}$, $r(k)$ is decreasing for $k > k^{**}$. As also $r(-\infty) < 0$, $r(k^* - 0) > 0$ and $r(k^* + 0) < 0$, there exists k_0: $r(k) > 0$ only on (k_0, k^*). Therefore, $U_1(k)$ is decreasing on $k < k_0$, increasing on (k_0, k^*), decreasing on (k^*, b). Hence, k^* is a global maximum if $U_1(k^*) \geqslant U_1(-\infty) = 0$, which has already been shown. □

Appendix B

Appendices of Chapter 3

Proof of Lemma 19. *Necessity.* Let (γ, β) be an equilibrium recommendation. Then (3.3) implies that

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant G(b)(v - c(v) - b) \tag{B.1}$$

holds for all $v \in V$ and $b \in \mathbb{R}_+$. Since the above in particular holds for $b = 0$, we have $\beta(v) + \gamma(v) \leqslant v$. Then it is sufficient to show that for arbitrary $v, v' \in V$

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant H(v')(v - c(v) - \beta(v')) \tag{B.2}$$

Fix $v, v' \in V$. If $\beta(v') \geqslant v - c(v)$ then, given that $G(\beta(v')) \leqslant H(v')$, (B.2) immediately follows from (B.1) by setting $b = \beta(v')$. Therefore, assume that $\beta(v') < v - c(v)$. For arbitrary $\epsilon > 0$ (B.1) yields

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant G(\beta(v') + \epsilon)(v - c(v) - (\beta(v') + \epsilon))$$

Observe that $G(\beta(v') + \epsilon) \geqslant H(v')$ and therefore by taking a limit in the right-hand side of the above inequality when ϵ goes to zero we obtain (B.2).

Sufficiency. Assume now that (B.2) holds for all $v, v' \in V$ and $\beta(v) + \gamma(v) \leqslant v$. We want to show that (B.1) holds for all $v \in V$ and $b \in \mathbb{R}_+$. Fix $v \in V$ and $b \in \mathbb{R}_+$. If $v - c(v) \leqslant b$ then (B.1) trivially holds as left-hand side is non-negative.

Let $v - c(v) > b$. If $b = \beta(v')$ for some $v' \in V$ then (B.1) is an immediate application of (B.2). Assume that $b \neq \beta(v')$ and let $S = \{v' : \beta(v') > b\}$. If $S = \varnothing$, then $G(b) = 1$. Next consider a sequence $v'_{(k)}$ such that $H(v'_{(k)})$ converges to 1. Without loss of generality[1], we may assume that $\beta(v'_{(k)})$ also convergent sequence and let b' be a limit point of this sequence. Note that $b' \leqslant b$. For any $k \in \mathbb{N}$ (B.2) implies:

$$G(\beta(v))(v - \gamma(v) - \beta(v)) \geqslant H(v'_k)(v - c(v) - \beta(v'_k)) \quad \text{(B.3)}$$

Then taking a limit of the right-hand side of the above inequality and taking into account that $G(b) = 1$ and $b' < b$, (B.2) follows.

Suppose now that $S \neq \varnothing$ and let $b_{min} = \inf\{\beta(v') : v' \in S\}$. Then there exists a sequence $v'_{(k)} \subseteq S$ such that $\beta(v'_k)$ converges to b_{min}. As before, without loss of generality, we may assume that sequence $H(v_{(k)})$ is also convergent. Let $\underline{H}$ be a limit of this sequence. Since $H(v_k) \geqslant G(b)$ for all k, $\underline{H} \geqslant G(b)$. Similarly, to the above, taking the limit of the right-hand side of inequality (B.3) then immediately implies (B.2).

□

Proof of Corollary 21. Let $(\beta^*_{FPA,i}, \gamma^*(v))$, $i \in \{1, 2\}$ and $\gamma^*(v) = 0$, be the optimal recommendation in a first-price auction with the maximum penalty function c_i. If $v_H - c_1(v_H) \leqslant r$ then $v_H - c_2(v_H) \leqslant r$ and hence, $\beta^*_{FPA,i}(v) = v$ for all $v \in V$ and $i \in \{1, 2\}$. Therefore, in both cases, when a maximum penalty function equals to c_1 and when it equals to c_2, the seller gets the same revenue.

Assume now that $v_H - c_1(v_H) > r$. We will do a proof by mathematical induction

[1] Otherwise, since $\beta(V) \subset [0, b]$ we may always choose a convergent subsequence.

on the number of steps k of the algorithm for finding the optimal recommendation, to show that

Claim: For arbitrary $k \geqslant 1$, $\beta^*_{FPA,2}(v) \geqslant \beta^*_{FPA,1}(v)$ for all $v \in (r, v_k]$ with strict inequality for all $k > 1$.

Base, $k = 1$. For all $v \in (r, v_1]$: $\beta^*_{FPA,2}(v) = \beta^*_{FPA,1}(v) = v$.

Induction Step, $k = l + 1$. Let $\beta^*_{FPA,2}(v) \geqslant \beta^*_{FPA,1}(v)$ for all $v \in (r, v_l]$. Then for all $v \in (v_l, v_{l+1}]$ we obtain

$$\begin{aligned} H^*(v)(v - \beta^*_{FPA,2}(v)) &= \sup_{v' < v - c_2(v)} H^*(v')(v - c_2(v) - \beta^*_{FPA,2}(v')) \\ &< \sup_{v' < v - c_1(v)} H^*(v')(v - c_1(v) - \beta^*_{FPA,1}(v')) \\ &= H^*(v)(v - \beta^*_{FPA,1}(v)) \end{aligned}$$

where the inequality follows from $c_2(v) > c_1(v)$ and the induction step. Taking into account that $H^*(v) = F^{n-1}(v) > 0$, we have $\beta^*_{FPA,2}(v) > \beta^*_{FPA,1}(v)$, which ends the proof of the above claim.

Since $\beta^*_{FPA,2}(v) \geqslant \beta^*_{FPA,1}(v)$ for all $v \in [r, v_H]$, with a strict inequality for a positive measure of types v, the seller's revenue is higher when the maximum penalty function equals to $c_2(v)$. □

Proof of Corollary 22. Let β_{FPA} be an equilibrium bidding function in a standard auction. We will do a proof by mathematical induction on the number of steps k of the algorithm of finding β^*_{FPA}. The claim is:

Claim: For arbitrary $k \geqslant 1$, $\beta^*_{FPA}(v) > \beta_{FPA}(v)$ holds for all $v \in (r, v_k]$.

Base, $k = 1$. For all $v \in (r, v_1]$: $\beta^*_{FPA}(v) = v > \beta_{FPA}(v)$.

Induction Step, $k = l + 1$. Let $\beta^*_{FPA}(v) > \beta_{FPA}(v)$ for all $v \in (r, v_l]$. Fix, $v \in (v_l, v_{l+1}]$. Since $\beta_{FPA}(v)$ is a symmetric equilibrium in a standard auction, it should be the case that

$$H^*(v)(v - \beta_{FPA}(v)) \geqslant H^*(v')(v - \beta_{FPA}(v'))$$

for all $v' \in V$. Consequently,

$$\begin{aligned} H^*(v)(v - \beta_{FPA}(v)) &> \sup_{v' < v - c(v)} H^*(v')(v - c(v) - \beta_{FPA}(v')) \\ &\geqslant \sup_{v' < v - c(v)} H^*(v')(v - c(v) - \beta^*_{FPA}(v')) \\ &= H^*(v)(v - \beta^*_{FPA}(v)) \end{aligned}$$

where the second inequality follows from induction assumption. Therefore, we proved, my method of mathematical induction, that $\beta^*_{FPA}(v) > \beta_{FPA}(v)$. □

Proof of Proposition 23. Consider a recommendation $(\beta_{FPA}(v), \gamma(v))$ such that $\gamma(v) = 0$ and $\beta_{FPA}(v)$ for each v is a convex combination of a type v buyer's bid in a standard first-price auction $\beta^S_{FPA}(v)$ and his value v:

$$\beta_{FPA}(v) = (1 - k(a)) \cdot \beta^S_{FPA}(v) + k(a)v$$

where

$$k(a) = \frac{\sqrt{M^2a^2 + 4Ma} - Ma}{2}$$

and

$$M = \min_{v \in V} \frac{(n-1)vf(v)}{F(v)} > 0$$

Note, that for all $a \in [0, 1]$, $k(a) \in [0, 1)$ with $k(0) = 0$. If all the bidders follow this recommendation, then the revenue equals to $R(a) = (1 - k(a))R^S + k(a)R_{max}$, where R^S is the seller's revenue in a standard first-price auction and R_{max} denotes a value of a full surplus. Evaluating the derivative of $R(a)$ at $a = 0$ yields

$$\left.\frac{dR(a)}{a}\right|_{a=0} = \left.\frac{dk(a)}{da}\right|_{a=0} \cdot (R_{\max} - R^S) = \infty$$

If recommendation $(\beta_{FPA}(v), \gamma(v))$ is an equilibrium, then

$$\lim_{a \to 0} \frac{R^*(a) - R^S}{a} \geqslant \lim_{a \to 0} \frac{R(a) - R(0)}{a} = \left.\frac{dR(a)}{a}\right|_{a=0} = \infty$$

Therefore, we are left to show that the proposed recommendation $(\beta_{FPA}(v), \gamma(v))$ is an equilibrium. Since $\beta_{FPA}(v)$ is monotonic and does not exceed v, it is sufficient to check that the incentive compatibility constraints

$$F^{n-1}(v)(v - \beta_{FPA}(v)) \geqslant F^{n-1}(v')((1-a)v - \beta_{FPA}(v'))$$

hold for all v and v' from the interval $[r, v_H]$. Given that

$$\beta^{S}_{FPA}(v) = v - \frac{1}{F^{n-1}(v)} \int_{r}^{v} F^{n-1}(t)dt$$

and the way we defined β_{FPA}, the above inequality can be rewritten as

$$(1-k(a)) \int_{r}^{v} F^{n-1}(t)dt \geqslant F^{n-1}(v')((1-a)v - v') + (1-k(a)) \int_{r}^{v'} F^{n-1}(t)dt \quad \text{(B.4)}$$

For fixed $v \in [r, v_H]$ and $a \in [0, 1]$ we need only to check that the above inequality holds for v' at which the right-hand side achieves its maximum on V (clearly, if inequality (B.4) holds for all $v' \in V$, it also holds for all $v' \in [r, v_H]$). Provided that v' can take values from the closed interval $[v_L, v_H]$ and that the right-hand side is continuously differentiable, the maximum is attained at either end-points—v_L and v_H—or at values of v' at which the derivative of the right-hand side equals to 0. Note that inequality (B.4) is clearly satisfied at $v' = v_L$ as the left-hand side is non-negative and the right-hand side is non-positive at $v' = v_L$. Differentiate the right-hand side of inequality (B.4) with respect to v', we obtain

$$\frac{dRHS(v')}{dv'} = (n-1)F^{n-2}(v')f(v')((1-a)v - v') - F^{n-1}(v') + (1-k(a))F^{n-1}(v')$$

For all $v' > (1-a)v$, the derivative of $RHS(v')$ is negative and therefore, the maximum is never attained at $v' = v_H$. Therefore, we left with verifying that inequality (B.4) holds for all $v' = v^*$ at which the derivative of $RHS(v')$ equals to 0:

$$v^* + \frac{k(a)F(v^*)}{(n-1)f(v^*)} = (1-a)v$$

Substituting the expression of v^* into (B.4) and rearranging terms we have

$$(1-k(a)) \int_{v^*}^{v} F^{n-1}(t)dt \geqslant \frac{k(a)F^n(v^*)}{(n-1)f(v^*)} \tag{B.5}$$

Next observe that the left-hand side of the last inequality satisfies

$$\begin{aligned}(1-k(a)) \int_{v^*}^{v} F^{n-1}(t)dt &\geqslant (1-k(a))F^{n-1}(v^*)(v-v^*) \\ &\geqslant (1-k(a))F^{n-1}(v^*)\left(av + \frac{k(a)F(v^*)}{(n-1)f(v^*)}\right)\end{aligned}$$

Therefore, to show that (B.5) holds, it is sufficient to check that

$$(1-k(a))\left(av + \frac{k(a)F(v^*)}{(n-1)f(v^*)}\right) \geqslant \frac{k(a)F^(v^*)}{(n-1)f(v^*)}$$

or what is equivalent

$$(1-k(a))av \geqslant k^2(a)\frac{F(v^*)}{(n-1)f(v^*)}$$

The latest is true as

$$\frac{k^2(a)}{1-k(a)} = aM = a \cdot \min_{v \in V} \frac{(n-1)vf(v)}{F(v)} \leqslant a \cdot \frac{(n-1)v^*f(v^*)}{F(v^*)}$$

□

Proof of Corollary 27. Since $\tilde{\gamma}(b,v) = \gamma^*(b,v)$ for all $b < \beta^*_{FPA}(v - c(v))$, we need only to insure that equilibrium condition (3.1) holds for all $b \geqslant \beta^*_{FPA}(v - c(v))$.

Following the logic of Lemma 25 we consider only b for which there exists v' such that $\beta^*_{FPA}(v') = b$. Then for all $v' \in [v - c(v), v]$ using monotonicity of u^* yields

$$\begin{aligned} u^*(v) \geqslant u^*(v') &= H^*(v')(v - (v - v') - \beta^*_{FPA}(v)) \\ &= H^*(v')(v - \tilde{\gamma}(\beta^*_{FPA}(v'), v') - \beta^*_{FPA}(v)) \end{aligned}$$

For $v' > v$, using Lemma 25 we obtain

$$u^*(v) \geqslant u^*(v') - H^*(v')(v' - v) = H^*(v')(v - \tilde{\gamma}(\beta^*_{FPA}(v'), v) - \beta^*_{FPA}(v'))$$

Therefore, β^*_{FPA} is an equilibrium induced by penalty rule $\tilde{\gamma}$.

□

Proof of Lemma 28. Let α be small enough positive number such that $\beta^*_{FPA}(v) - \beta_{FPA}(v) \leqslant \lambda = \min\{c(\tilde{v}), \tilde{v} - (v_H - c(v_H))\}$ for all $v \in [\tilde{v}, v_H]$. Given that $\beta_{FPA}(v)$ coincides with $\beta^*_{FPA}(v)$ for all $v \leqslant \tilde{v}$, we need only to check that equilibrium condition (3.6) holds for all $v \in V$ and $v' \geqslant \tilde{v}$. Provided our choice of λ for all $v \in V$ and $v' \geqslant \tilde{v}$

$$\beta^*_{FPA}(v) + \tilde{\gamma}(\beta^*_{FPA}(v'), v) \leqslant \beta_{FPA}(v) + c(v)$$

where $\tilde{\gamma}$ is defined in Corollary 27. According to Corollary 27, β^*_{FPA} is an equilibrium in the first-price auction with penalty rule $\tilde{\gamma}$. Therefore,

$$u^*(v) \geqslant H^*(v')(v - \tilde{\gamma}(\beta^*_{FPA}(v'), v) - \beta^*_{FPA}(v')) \geqslant H^*(v')(v - c(v) - \beta^*_{FPA}(v'))$$

To complete the proof recall that by construction of β_{FPA}, the expected utility of bidder with type v is the same under equlbrium recommendation (γ, β_{FPA}) and $(\gamma^*, \beta^*_{FPA})$. □

Proof of Theorem 30. With some abuse of notation, let $\beta^*_{FPA}(v, r)$ denotes the amount a buyer with value v bids under the optimal recommendation, provided the reserve price equals to r. For concreteness, assume that $\beta^*_{FPA}(v, r) = v$ if $v < r$. We first derive a lower bound for expected gains in the seller's profit when the reserve price is reduced by a small amount $\Delta r > 0$. We do this by showing that for all $v > v_L$

$$\beta^*_{FPA}(v,r) - \beta^*_{FPA}(v,r-\Delta r) \leqslant \frac{\Delta r F^{n-1}(r)}{F^{n-1}(v)} \tag{B.6}$$

(B.6) is proved by mathematical induction on the number of steps k of the algorithm of finding β^*_{FPA} when reserve price is $r - \Delta r$. The claim is:

Claim: For arbitrary $k \geqslant 1$ (B.6) holds for all $v \leqslant v_k$.

Base, $k = 1$. For all $v \leqslant v_1 = v_{r-\Delta r}$: $\beta^*_{FPA}(v, r-\Delta r) = \beta^*_{FPA}(v,r) = v$.

Induction Step, $k = l+1$. Let (B.6) holds for all $v \leqslant v_l$. For arbitrary $v \in (v_l, v_{l+1}]$ we may write the equilibrium condition as

$$F^{n-1}(v)(v - \beta^*_{FPA}(v, r-\Delta r)) = \sup_{v' < v - c(v)} F^{n-1}(v')(v - c(v) - \beta^*_{FPA}(v', r - \Delta r))$$

We derive an upper bound on the right-had side of the above equality, by splitting supremum in the right-hand side into two parts: 1) for all $v' \leqslant r$ and 2) for all $v' > r$. For the first part we obtain:

$$\begin{aligned}\sup_{v' \leqslant r} F^{n-1}(v')(v - c(v) - \beta^*_{FPA}(v', r-\Delta r)) &\leqslant F^{n-1}(r)(v - c(v) - (r - \Delta r)) \\ &= F^{n-1}(r)(v - c(v) - r) + \Delta r F^{n-1}(r) \\ &\leqslant F^{n-1}(v)(v - \beta^*_{FPA}(v,r)) + \Delta r F^{n-1}(r)\end{aligned}$$

For the second part we obtain:

$$\begin{aligned}&\sup_{v' \in (r, v-c(v)]} F^{n-1}(v')(v - c(v) - \beta^*_{FPA}(v', r-\Delta r)) \\ &\qquad \leqslant \sup_{v' \in (r, v-c(v)]} F^{n-1}(v')(v - c(v) - \beta^*_{FPA}(v', r)) + \Delta r F^{n-1}(r) \\ &\qquad \leqslant F^{n-1}(v)(v - \beta^*_{FPA}(v,r)) + \Delta r F^{n-1}(r)\end{aligned}$$

where the first inequality follows from induction assumption as $v' \leqslant v - c(v) \leqslant v_l$ and the second one is an equilibrium condition when reserve price equals to r.

Combining two parts together we have:

$$F^{n-1}(v)(v-\beta^*_{FPA}(v,r-\Delta r)) \leqslant F^{n-1}(v)(v-\beta^*_{FPA}(v,r)) + \Delta r F^{n-1}(r)$$

Dividing both sides of above inequality by $F^{n-1}(v)$ and rearranging the terms result in (B.6). Hence, by method of mathematical induction (B.6) holds for all $v > v_L$.

Let v be the highest value among the bidders. Then by reducing the reserve price by Δr, the seller gains by making a sale at price v whenever $v \in [r-\Delta r, r]$, is indifferent whenever $v \in [r, v_{r-\Delta r}]$, and suffers a loss that is bounded by expression in the right-hand side of (B.6) whenever $v > v_{r-\Delta r}$. Therefore, the expected gain from reducing the reserve price by Δr is at least as large as:

$$LB(\Delta r) = \int\limits_{r-\Delta r}^{r} v dF^n(v) - \int\limits_{v_{r-\Delta r}}^{v_H} \frac{\Delta r F^{n-1}(r)}{F^{n-1}(v)} dF^n(v)$$

which after simplification results in

$$LB(\Delta r) = \int\limits_{r-\Delta r}^{r} v dF^n(v) - n\Delta r F^{n-1}(r)(1-F(v_{r-\Delta r}))$$

Taking the limit from the last expression divided by Δr when Δr goes to zero, results in:

$$\lim_{\Delta r \to 0} \frac{LB(\Delta r)}{\Delta r} = nF^{n-1}(r)f(r)\left(r - \frac{1-F(v_r)}{f(r)}\right)$$

Note that the expression in the brackets is strictly larger than the virtual valuation $\psi(r)$. Hence, whenever the reserve price r exceeds the optimal reserve price r_S in standard auctions, the seller finds it profitable to reduce the reserve price by a small amount. If $r_S > v_L$ then at $r = r_S$, unlike the standard auctions, the gains are still positive and therefore the seller benefits from reducing the reserve price below r_S. Finally, the existence of optimal reserve price in our environment follows directly from continuity of $\beta^*(v, r)$ in the second argument, which is an immediate application of (B.6).

□

Proof of Theorem 34. Let (γ, β) be an optimal monotonic recommendation in some winner-pay auction. Then

$$H^*(v)(v - \gamma(v) - m(\beta(v))) \geqslant \sup_{v' \in v} H^*(v')(v - c(v) - m(\beta(v')))$$

Consider now a recommendation (γ, β_{FPA}), where $\beta_{FPA}(v) = m(\beta(v))$. Provided that $m(\beta(\cdot))$ is a monotonic function Lemma 25 guarantees that (γ, β_{FPA}) is an equilibrium recommendation in the first-price auction. Therefore, we complete the proof by showing that

Lemma 58. *Let $\beta(\cdot)$ be strictly increasing function on $[r, v_H]$. Then $m(\beta(\cdot))$ is also a strictly increasing function on $[r, v_H]$.*

Proof. Fix arbitrary $v, v' \in [r, v_H]$ such that $v' > v$. For any $w \in \mathbb{R}_+^{n-1}$, let $y_k(w)$ be a k-dimensional vector with coordinates equal to the first k coordinates of vector w, and $z_k(w)$ be a $n-1$-dimensional vector, first k coordinates of which coincides with first k coordinates of w and each of the last $n-k-1$ coordinates equals to v. Also for arbitrary k-dimensional vector x, let $\mathbb{F}(x) = \prod_{i=1}^{k} F(x_i)$. Then using monotonicity of M, for all natural $k \leqslant n-1$ we have

$$\begin{aligned} F^{n-1}(v)m(\beta(v)) &= \int\limits_{\{w \in V^{n-1}: w_i < v\}} M(v, w) d\mathbb{F}(w) \\ &\leqslant F^{n-1-k}(v) \int\limits_{\{w \in V^{n-1}: w_i < v\}} M(v, z_k(w)) d\mathbb{F}(y_k(w)) \end{aligned}$$

which implies that

$$F^k(w)m(\beta(v)) \leqslant \int\limits_{\{w\in V^{n-1}:w_i<v\}} M(v, z_k(w))d\mathbb{F}(y_k(w)), \tag{B.7}$$

Note, that above also holds for $k = 0$ if with some abuse of notation we assume that $M(v, z_0(w)) = \int\limits_{\{w\in V^{n-1}:w_i<v\}} M(v, z_0(w))d\mathbb{F}(y_0(w))$. Finally,

$$\begin{aligned} F^{n-1}(v')m(\beta(v')) &= \int\limits_{\{w\in V^{n-1}:w_i<v'\}} M(v, w)d\mathbb{F}(w) \\ &\geqslant \sum_{k=0}^{n-1} C_n^k [F(v') - F(v)]^{n-1-k} \int\limits_{\{w\in V^{n-1}:w_i<v\}} M(v, z_k(w))d\mathbb{F}(y_k(w)) \\ &\geqslant \sum_{k=0}^{n-1} C_n^k [F(v') - F(v)]^{n-1-k} F^k(v)m(\beta(v)) = F^{n-1}(v')m(\beta(v)), \end{aligned}$$

where the first inequality follows from monotonicity of M and second follows from (B.7). Hence, $m(\beta(v')) > m(\beta(v))$. □

□

Proof of Proposition 38. To simplify the exposition, we assume that $r = v_L$[2]. Let, as before, v_1 denotes a valuation such that $v_1 - c(v_1) = v_L$ ($v_1 = v_H$ if $v_H - c(v_H) > v_L$), and let $\tilde{v}$ be an arbitrary valuation such that $\tilde{v} \in (v_L, v_1)$. Then define a density function f as

$$f(v) = \begin{cases} a \text{ if } v \in [v_L, \tilde{v} - \epsilon] \\ a + \frac{(k-1)a}{\epsilon}(v - \tilde{v} + \epsilon) + ka \text{ if } v \in (\tilde{v} - \epsilon, \tilde{v}) \\ ka \text{ if } v \in [\tilde{v}, v_H] \end{cases}$$

where $k > 1$, $\epsilon \in (0, \tilde{v} - v_L)$ and $a = 1/(\epsilon(k-1)/2 + k(v_H - v_L) + (1-k)(\tilde{v} - v_L))$ to guarantee that the density function f integrates to 1. Given the specific form

[2] The proof is easily extended to the case of $r > v_L$.

of a, we may always choose ϵ small enough such that $n\epsilon + \frac{1}{ka} < v_H - v_L$ (which is equivalent to $\epsilon < 2(k-1)(\tilde{v} - v_L)/(2nk + k - 1)$).

Note that $F(v) = a(v - v_L)$ for all $v \in [v_L, \tilde{v} - \epsilon]$ and $F(v) = 1 + ka(v - v_H)$ for all $v \in [\tilde{v}, v_H]$. Therefore,

$$\beta_{SPA}(v) = v + \frac{F(v)}{(n-1)f(v)} = \begin{cases} \frac{1}{n-1}(nv - v_L) \text{ , if } v \in [v_L, \tilde{v} - \epsilon] \\ \frac{1}{n-1}(nv - v_H + \frac{1}{ka}) \text{ , if } v \in [\tilde{v}, v_1] \end{cases}$$

Since $n\epsilon + \frac{1}{ka} < v_H - v_L$, $\beta_{SPA}(\tilde{v} - \epsilon) > \beta_{SPA}(\tilde{v})$ and hence, $\beta_{SPA}(v)$ is non-monotonic.

□

Proof of Proposition 39. Under the optimal taxation the expected utility of type v buyer should satisfy:

$$u^*(v) = \max_{v'} F^{n-1}(v')((1-t)v - \beta^*_{FPA}(v'))$$

Taking into account that maximum is achieved at $v' \leqslant (1-t)v$ and applying the Envelope Theorem yields:

$$(u^*)'(v) = (1-t)F^{n-1}(v') \leqslant (1-t)F^{n-1}((1-t)v)$$

Under the flat taxation, the net of taxes type v buyer's valuation equals to $(1-t)v$. Then in the first-price auction with symmetric equilibrium β_{FPA} the expected utility of type v buyer equals to

$$u(v) = \max_{v'} F^{n-1}(v')((1-t)v - \beta_{FPA}(v))$$

Since maximum is achieved at $v' = v$, applying the Envelope Theorem results in:

$$u'(v) = (1-t)F^{n-1}(v)$$

Since in both cases final allocation is efficient, the difference between seller's expected payoff under optimal and flat taxation equals to:

$$
\begin{aligned}
n\int_{v_L}^{v_H} u(v) - u^*(v)dF(v) &= n\int_{v_L}^{v_H} (u^*(v) - u(v))d(1 - F(v)) \\
&= n\int_{v_L}^{v_H} (u'(v) - (u^*)'(v))(1 - F(v))dv \\
&= n(1-t)\int_{v_L}^{v_H} (F^{n-1}(v) - F^{n-1}((1-t)v))dv
\end{aligned}
$$

where the second inequality follows from integration by parts. □

Appendix C

Appendices of Chapter 4

Proof of Lemma 50.

$$g_i(r_i, s_i) = \sum_{k=1}^{S} \mathbb{P}(\theta = k|s_i) \left(\mathbb{E}[m_i(r_i, s_{-i})|\theta = k, s_i] + \mathbb{E}[U_i(a, \theta)|\theta = k, a = \phi(r_i, s_{-i})]\right)$$

$$= \mathbb{E}[m_i(r_i, s_{-i})|s_i] + \sum_{k=1}^{S} \mathbb{P}(\theta = k|s_i)\mathbb{E}[U_i(a, \theta)|\theta = k, a = \phi(r_i, s_{-i})]$$

$$= \sum_{s_{-i} \in \Omega^{m-1}} P_i(s_{-i}|s_i) m_i(r_i, s_{-i}) + \mathbb{E}[U_i(a, \theta)|s_i, a = \phi(r_i, s_{-i})]$$

If we denote $g_{i,r_i} = (g_i(r_i, 1), .., g_i(r_i, S))'$, $Pol_{i,r_i,s_i} = \mathbb{E}[U_i(a, \theta)|s_i, a = \phi(r_i, s_{-i})]$, $Pol_{i,r_i} = (Pol_{r_i,1}, .., Pol_{r_i,S})'$. Finally, denote by m_{i,r_i} a $1 \times S^{m-1}$ vector of monetary payments $m_i(r_i, s_{-i})$ written according to increasing lexicographic order in s_{-i}:

$$m_{i,r_i} = (m_i(r_i; 1, .., 1, 1), m_i(r_i; 1, .., 1, 2), .., m_i(r_i; 1.., S, S-1), m_i(r_i; S, .., S, S)).$$

Then for any fixed r_i, $g_i(r_i, s_i) \equiv g_i^0(r_i, s_i)$ is equivalent to

$$Q_i^m \times m_{i,r_i} + Pol_{i,r_i} = g_{i,r_i}^0.$$

As matrix Q_i^m has full rank (rank(Q_i^m)=S), for any $g^0_{i,r_i} \in \mathbb{R}^{|\Omega|}$, there is a vector $m_{i,r_i} \in \mathbb{R}^{\Omega^{m-1}}$ such that the matrix equation above is satisfied. Going through all $r_i \in \Omega$, and joining these cases, we get that for any $g_i^0(r_i, s_i) \in \mathbb{R}^{|\Omega|^2}$, there exists a payment function $m_i(r_i, r_{-i}) : \Omega^m \to \mathbb{R}$ such that $g_i(r_i, s_i) \equiv g_i^0(r_i, s_i)$. □

Proof of Proposition 51. For every expert i ($1 \leqslant i \leqslant m$) take $g_i^0(r_i, s_i) \equiv u_{i0}$. By Lemma 50, for each i there exists a payment function $m_i^*(r_i, r_{-i})$ such that $g_i(r_i, s_i) \equiv g_i^0(r_i, s_i) \equiv u_{i0}$. Then there exists a socially optimal equilibrium, where the principal commits to the socially optimal action and the payment functions $m_i^*(r_i, r_{-i})$, all participating experts recommend $r_i = s_i$ and get u_{i0} in expectation. □

Proof of Proposition 52. We can slightly modify the strategy profile from Proposition 47 by only changing payment functions m_i. The principal commits to choose a socially optimal action, she proposes payment functions $m_i(r_i, r_{-i}) = \frac{2}{n}\alpha_i(b^* - b_i)(r_i - \frac{1}{n-1}\sum_{j\neq i} r_j) + u_{0i} + \alpha_i[(b_i - b^*)^2 + \frac{\sigma^2}{n}(1 + (n-1)\rho)]$ to all experts; all experts report truthfully if the principal acts according to the described strategy, all experts do not participate otherwise. Then

$$\mathbb{E}[m_i(r_i, r_{-i})|s_i] = \frac{2}{n}\alpha_i(b^* - b_i)(r_i - s_i) + u_{0i} + \alpha_i\left[(b_i - b^*)^2 + \frac{\sigma^2}{n}(1 + (n-1)\rho)\right]$$

as before and the whole line of the proof in Proposition 47 stays the same here. □

Proof of Proposition 53. The principal commits to choose $a = \frac{1}{n}\sum_{i=1}^n r_i + b^*$ with $b^* = \frac{\sum_{i=1}^n \alpha_i b_i}{1+\sum_{i=1}^n \alpha_i}$ if all experts participate and to choose $a = 0$ otherwise, she proposes payment functions $m_i(r_i, r_{-i}) = \frac{1}{n^2}\alpha_i(r_i - \frac{1}{n-1}\sum_{j\neq i} r_j)^2 + \frac{2}{n}\alpha_i(b^* - b_i)(r_i - \frac{1}{n-1}\sum_{j\neq i} r_j) + u_{0i} + \alpha_i[(b_i - b^*)^2 + \frac{((n-2)+(n^2-2n+2)\rho)\sigma^2}{(n-1)n}]$ to all experts; all experts report truthfully if the principal acts according to the described strategy, all experts do not participate otherwise.

Again the principal faces infinite losses from deviation, so it is enough to check that the truthful report is a best response for each expert i to other players' strategies. As we already know, $(a-\theta-b_i|s_i) \sim N(\frac{1}{n}(r_i-s_i)+b^*-b_i, \frac{\sigma^2}{n}(1+(n-1)\rho))$. Also, $(\frac{1}{n-1}\sum_{j\neq i} s_j|s_i) \sim N(s_i, \frac{(1-\rho)n}{n-1}\sigma^2)$ and participating expert i obtains

$$\mathbb{E}[-\alpha_i \left(a-\theta-b_i\right)^2 + m_i(r_i, r_{-i}) \left|s_i\right] =$$

$$-\alpha_i \left[\left(\frac{1}{n}r_i - \frac{1}{n}s_i + b^* - b_i\right)^2 + \frac{\sigma^2}{n}(1+(n-1)\rho)\right]$$

$$+\frac{1}{n^2}\alpha_i \left[(r_i-s_i)^2 + \frac{(1-\rho)n}{n-1}\sigma^2)\right] + \frac{2}{n}\alpha_i(b^*-b_i)(r_i-s_i)$$

$$+u_{0i} + \alpha_i \left[(b_i-b^*)^2 + \frac{((n-2)+(n^2-2n+2)\rho)\sigma^2}{(n-1)n}\right] = u_{0i}$$

Any expert i's strategy yields him u_{0i} and the considered strategy profile is a socially optimal equilibrium.

Finally, let us check the limited liability requirement. For payment function m_i, denoting $r_i - \frac{1}{n-1}\sum_{j\neq i} r_j \equiv z$, we obtain $m_i(z) = \frac{1}{n^2}\alpha_i z^2 + \frac{2}{n}\alpha_i(b^*-b_i)z + u_{0i} + \alpha_i[(b_i - b^*)^2 + \frac{(n-2)(1+(n-2)\rho)\sigma^2}{(n-1)n}]$, the global minimum of $m_i(z)$ is $u_{0i} + \frac{((n-2)+(n^2-2n+2)\rho)\sigma^2}{(n-1)n} \geqslant u_{i0} + \frac{(n-2)\alpha_i\sigma^2}{(n-1)^2} \geqslant l_i$. □

Bibliography

Alonso, R. and N. Matouschek (2008). Optimal delegation. *The Review of Economic Studies 75*, 259–293.

Ambrus, A. and A. Kolb (2018). On defining ex ante payoffs in games with diffuse prior. *Unpublished working paper*.

Ambrus, A. and S. Lu (2014). Almost fully revealing cheap talk with imperfectly informed senders. *Games and Economic Behavior 88*, 174–189.

Ambrus, A. and S. Takahashi (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics 3*, 1–27.

Austen-Smith, D. (1993). Interested experts and policy advice: multiple referrals under open rule. *Games and Economic Behavior 5*, 3–43.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica 70*, 1379–1401.

Battaglini, M. (2004). Policy advice with imperfectly informed experts. *Advances in Theoretical Economics 4*, 1–32.

Bernhardt, D., M. Campello, and E. Kutsoati (2006). Who herds? *Journal of Financial Economics 80*, 657–675.

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

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review 78*, 1–3.

Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *The Journal of Political Economy 97*, 1060–1090.

Callander, S. (2011). Searching for good policies. *American Political Science Review 105*, 643–662.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science 29*, 69–95.

Che, Y.-K. and I. Gale (1998). Standard auctions with financially constrained bidders. *The Review of Economic Studies 65*, 1–21.

Chen, Q. and W. Jiang (2006). Analysts' weighting of private and public information. *Review of Financial Studies 19*, 319–355.

Crawford, V. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*, 1431–1451.

Crémer, J. (1987). Auctions with contingent payments: Comment. *The American Economic Review 77*, 746.

Cremer, J. and R. P. McLean (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica 53*, 345–361.

Cremer, J. and R. P. McLean (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica 56*, 1247–1257.

DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2005). Bidding with securities: Auctions and security design. *The American Economic Review 95*, 936–959.

Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies 69*, 811–838.

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

Downs, A. (1957). *An economic theory of democracy.* New York: Harper & Row.

Ederer, F., R. Holden, and M. A. Meyer (2018). Gaming and strategic opacity in incentive provision. *The RAND Journal of Economics, forthcoming*.

Friedman, D. (1991). A simple testable model of double auction markets. *Journal of Economic Behavior and Organization 15*, 47–70.

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

Gerardi, D., R. McLean, and A. Postlewaite (2009). Aggregation of expert opinions. *Games and Economic Behavior 65*, 339–371.

Gilligan, T. and K. Krehbiel (1989). Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science 33*, 459–490.

Graham, D. A. and R. C. Marshall (1987). Collusive bidder behavior at single-object second-price and english auctions. *The Journal of Political Economy 95*, 1217–1239.

Gratton, G. (2014). Pandering and electoral competition. *Games and Economic Behavior 84*, 163–179.

Groetsch, C. (2007). Integral equations of the first kind, inverse problems and regularization: a crash course. In *Journal of Physics: Conference series*, Volume 73, pp. 012001. IOP Publishing.

Hansen, R. G. (1985). Auctions with contingent payments. *The American Economic Review 75*, 862–865.

Heidhues, P. and J. Lagerlöf (2003). Hiding information in electoral competition. *Games and Economic Behavior 42*, 48–74.

Holmström, B. (1977). *On incentives and control in organizations*. Stanford University.

Holt Jr, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *The Journal of Political Economy 88*, 433–445.

Hotelling, H. (1929). Stability in competition. *The Economic Journal 39*, 41–57.

Iezzoni, L., S. Rao, C. DesRoches, C. Vogeli, and E. Campbell (2012). Survey shows that at least some physicians are not always open or honest with patients. *Health Affairs 31*, 383–391.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *The American Economic Review 98*, 2127–2149.

Kartik, N., F. Squintani, and K. Tinn (2015). Information revelation and pandering in elections. Technical report, mimeo, Columbia University.

Kilgour, D. M. and Y. Gerchak (2004). Elicitation of probabilities using competitive scoring rules. *Decision Analysis 1*, 108–113.

Klemperer, P. (1999). Auction theory: a guide to the literature. *Journal of Economic Surveys 13*, 227–286.

Krishna, V. (2009). *Auction theory*. Academic press.

Krishna, V. and J. Morgan (2001a). Asymmetric information and legislative rules: Some ammendments. *American Political Science Review 95*, 435–452.

Krishna, V. and J. Morgan (2001b). A model of expertise. *Quarterly Journal of Economics 116*, 747–775.

Krishna, V. and J. Morgan (2008). Contracting for information under imperfect commitment. *The RAND Journal of Economics 39*, 905–925.

Laslier, J.-F. and K. Van der Straeten (2004). Electoral competition under imperfect information. *Economic Theory 24*, 419–446.

Li, H. and W. Suen (2004). Delegating decisions to experts. *Journal of Political Economy 112*, 311–335.

Li, H. and W. Suen (2009). Viewpoint: Decision-making in committees. *Canadian Journal of Economics 42*, 359–392.

Loertscher, S. et al. (2012). Location choice and information transmission. Technical report, mimeo, University of Melbourne.

Mailath, G. J. and P. Zemsky (1991). Collusion in second price auctions with heterogeneous bidders. *Games and Economic Behavior 3*, 467–486.

Marshall, R. C. and L. M. Marx (2007). Bidder collusion. *Journal of Economic Theory 133*, 374–402.

Martinelli, C. (2001). Elections with privately informed parties and voters. *Public Choice 108*, 147–167.

Martinelli, C. and A. Matsui (2002). Policy reversals and electoral competition with privately informed parties. *Journal of Public Economic Theory 4*, 39–61.

Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica 52*, 1473–1518.

Matthews, S. (1979). *Risk aversion and the efficiency of first and second price auctions*. College of Commerce and Business Administration, University of Illinois at Urbana-Champaign.

McAfee, R. P. and J. McMillan (1986). Bidding for contracts: a principal-agent analysis. *The RAND Journal of Economics 17*, 326–338.

McAfee, R. P. and J. McMillan (1992). Bidding rings. *The American Economic Review 82*, 579–599.

McAfee, R. P. and P. J. Reny (1992). Correlated information and mechanism design. *Econometrica 60*, 395–421.

Melumad, N. and T. Shibano (1991). Communication in settings with no transfers. *The RAND Journal of Economics 22*, 173–198.

Milgrom, P. and J. Roberts (1986). Relying on the information of interested parties. *The RAND Journal of Economics 17*, 18–32.

Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica 50*, 1089–1122.

Morris, S. and H. Shin (2002). The social value of private information. *The American Economic Review 92*, 1521–1534.

Morris, S. and H. Shin (2003). Global games: Theory and applications. Econometrics Society Monographs.

Myatt, D. and C. Wallace (2014). Central bank communication design in a Lucas-Phelps economy. *Journal of Monetary Economics 63*, 64–79.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research 6*, 58–73.

Mylovanov, T. and A. Zapechelnyuk (2017). Optimal allocation with ex post verification and limited penalties. *The American Economic Review 107*, 2666–2694.

Ottaviani, M. and P. Sørensen (2006). Professional advice. *Journal of Economic Theory 126*, 120–142.

Pesendorfer, W. and A. Wolinsky (2003). Second opinions and price competition: Inefficiency in the market for expert advice. *The Review of Economic Studies 70*, 417–437.

Pétry, F. and B. Collette (2009). Measuring how political parties keep their promises: A positive perspective from political science. Volume 15 of *Studies in Public Choice*, pp. 65–80.

Prendergast, C. (1993). A theory of "yes men". *The American Economic Review 83*, 757–770.

Riley, J. G. and W. F. Samuelson (1981). Optimal auctions. *The American Economic Review 71*, 381–392.

Sampford, M. (1953). Some inequalities on Mill's ratio and related functions. *The Annals of Mathematical Statistics 24*, 130–132.

Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association 66*, 783–801.

Schultz, C. (1996). Polarization and inefficient policies. *The Review of Economic Studies 63*, 331–344.

Shin, H. (1998). Adversarial and inquisitorial procedures in arbitration. *The RAND Journal of Economics 29*, 378–405.

Skrzypacz, A. (2013). Auctions with contingent payments – an overview. *International Journal of Industrial Organization 31*, 666–675.

Treitel, G. (1999). *The law of contract* (10 ed.). Sweet & Maxwell.

Vickrey, W. (1962). Auctions and bidding games. *Recent advances in game theory 29*, 15–27.

Winkler, R. L., J. Munoz, J. L. Cervera, J. M. Bernardo, G. Blattenberger, J. B. Kadane, D. V. Lindley, A. H. Murphy, R. M. Oliver, and D. Ríos-Insua (1996). Scoring rules and the evaluation of probabilities. *Test 5*, 1–60.

Wittman, D. (1983). Candidate motivation: A synthesis of alternative theories. *American Political Science Review 77*, 142–157.

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

Zitzewitz, E. (2001). Measuring herding and exaggeration by equity analysts and other opinion sellers. Technical report, mimeo, Stanford University.

Biography

Volodymyr Baranovskyi was born in Zhytomyr, Ukraine on June 10, 1985. Volodymyr received his Bachelor's Degree in Applied Mathematics (2006) and a subsequent Master's Degree in Applied Mathematics (2008) from Taras Shevchenko National University of Kyiv. He also holds a M.A. in Economics (2011) from Kyiv School of Economics and a M.A in Economics (2017) from Duke University.