Essays in Mechanism Design

Inauguraldissertation zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften der Universität Mannheim

Daniil Larionov

Abteilungssprecher: Prof. Klaus Adam, Ph.D.

Referent: Prof. Dr. Thomas Tröger

Korreferent: Prof. Dr. Martin Peitz

Tag der mündlichen Prüfung: 06.10.2022

Acknowledgements

This work would not have been possible without the help, guidance, and support of Thomas Tröger, Takuro Yamashita, and Martin Peitz. I thank Thomas Tröger for taking me on as a student, for his extremely helpful advice, and for his patience and unwavering support over the past six years. I thank Takuro Yamashita for the opportunity to visit Toulouse School of Economics, where parts of this dissertation have been written, for his continued guidance, and for making an invaluable contribution to our joint work. I thank Martin Peitz for his encouragement, support, and sharp questions and insightful comments that greatly improved this dissertation.

I also thank my coauthors Hien Pham and Shuguang Zhu for the productive collaboration and their extremely valuable contribution to Chapter 2 of this dissertation.

Throughout my PhD studies I have received financial support from the German Research Foundation through CRC TR 224 (Project B01), which I gratefully acknowledge. I thank the Graduate School of Economic and Social Sciences of the University of Mannheim for the opportunities it gave me, and its financial and administrative support.

Most importantly, I would like to express my deepest gratitude to my family. My mother Irina and my grandmother Valentina have always given me their unequivocal support, even at difficult moments. I am forever thankful to Maria for her love and her encouragement throughout this journey.

Preface

This dissertation consists of three self-contained chapters which explore how designers of economic mechanisms can make them more robust to *collusion* and *information acquisition* by the participants. Both collusion and information acquisition pose considerable challenges as they may allow the participants to coordinate their behavior and achieve outcomes unintended by the mechanism designer.

Chapter 1 (*Full Surplus Extraction from Colluding Bidders*) addresses the challenges posed by bidder collusion in repeated auction settings. Collusion is an important concern in many auctions that take place regularly and involve the same pool of bidders, especially if the number of bidders is small. Such auctions are ubiquitous in practice: procurement auctions often involve the same pool of suppliers, and auctions for natural resources, like oil and timber, often attract the same pool of potential buyers. Regularly run auctions may allow buyers to reduce competition and leave sellers with lower revenue. I study how a seller, who is facing colluding buyers, can use dynamic reserve prices to mitigate the impact of their collusion. To model the seller's concern for collusion, I introduce a new equilibrium concept: *collusive public perfect equilibrium.* For every dynamic strategy of the seller I define the corresponding "*buyer-game*" in which the seller is replaced by nature who chooses the reserve prices for the buyers in accordance with the seller's strategy. A public perfect equilibrium is collusive if the buyers cannot achieve a higher equilibrium payoff in the corresponding buyer-game. In a setting with symmetric buyers with private binary *iid* valuations and publicly revealed bids, I construct a collusive public perfect equilibrium that allows the seller to extract the entire surplus from the buyers in the limit as the discount factor goes to 1.

My construction offers new insights into how collusion can be addressed in practice. Notice that in any repeated auction setting there is a fundamental conflict between fighting collusion and revenue maximization. Revenue maximization is likely to require separation among the different valuation types of the buyers, but separation creates scope for collusion: the higher types would want to bid lower if they could. It turns out to be a good strategy for the seller to penalize abstentions and force the low-valuation buyers to overbid in exchange for the continuation of favorable terms of trade into the future where the same buyers might draw better valuations. Penalization of abstentions with high future reserve prices ensures that the buyers cannot improve their payoff by making the low-valuation types abstain in any given period, and thus forces the high-valuation types to bid higher in order to avoid the efficiency loss associated with pooling with the low-valuation types.

Chapter 2 (*First Best Implementation with Costly Information Acquisition*), which is a joint work with Hien Pham, Takuro Yamashita, and Shuguang Zhu, studies a general mechanism design setting, in which the participating agents can acquire costly and possibly correlated information. Many practical applications of mechanism design involve information acquisition by the agents. For example, bidders in an oil tract auction conduct test drills, and bidders in a spectrum auction conduct market research to better estimate the value of the license being sold. In those and in many other cases, both the information acquisition strategy and information acquisition outcomes can be hidden from the mechanism designer by the agents.

Motivated by these observations, we construct a model in which the mechanism designer proposes a mechanism to the agents, who then may acquire information about a payoff-relevant state of the world (e.g. the value of an oil tract or a spectrum license). At the outset, the mechanism designer and the agents share a common prior on the set of possible states of the world. The designer can only acquire information through the agents' reports. The agents can generate costly signals from a large signal space. We obtain a striking result: with four or more agents the designer can implement any social choice rule at zero information acquisition cost for the agents. Our solution involves a direct revelation mechanism and a set of signals for the agents constructed in such a way that any individual agent learns nothing about the state of the world from his own signal, but the designer can fully learn the state of the world from observing any pair of truthful reports. With enough pairs of agents, the identity of a unilateral deviator can always be established by the mechanism designer, who will then simply ignore the deviator's report and rely on the others' reports to determine the outcome. This feature of the mechanism ensures its incentive compatibility. If the designer can use transfers, then he can implement any social choice rule even with three agents. Transfers allow the designer to punish any profile of inconsistent reports without having to identify the agent who sent an untruthful report which caused the inconsistency.

In Chapter 3 (*Bilateral Trade with Costly Information Acquisition*), which is a joint work with Takuro Yamashita, we continue our study of mechanism design with information acquisition. Motivated in part by the question of whether first best implementation can be achieved with just two agents, we study a bilateral trade model with the same information acquisition environment. There is a buyer and a seller, who can trade a good of *ex ante* unknown quality, and an intermediary, interested in her own revenue, who designs a mechanism to facilitate their trade. In the beginning, the buyer, the seller, and the intermediary share a common prior over a set of possible qualities of the good. The intermediary proposes a mechanism to the buyer and the seller, who can then acquire information about the good's quality. We assume that the cost of information acquisition is proportional to the expected reduction in entropy. Under this assumption, we characterize the set of implementable allocations and show that the first best outcome, i.e. full surplus extraction, cannot be achieved.

Contents

Chapter 1

Full Surplus Extraction from Colluding Bidders

1.1 Introduction

Auctions rarely involve a one-shot interaction, often buyers and sellers face each other repeatedly. Procurement decisions for road construction and maintenance, to take one example, have to be made regularly and public authorities often have to deal with the same pool of potential suppliers. Auctions for electromagnetic spectrum, although less regular, often involve the same pool of potential buyers.

I model a seller who is concerned about colluding buyers and her own lack of commitment power. I assume that the seller offers an infinite sequence of first-price auctions with adjustable reserve prices and has to satisfy stringent public disclosure requirements: both the reserve prices and the buyers' bids are publicly disclosed after each round of trading. The seller can commit to her chosen reserve prices within every period, but does not have enough commitment power to fix the whole dynamic sequence of reserve prices. With respect to collusion, the seller takes a rather pessimistic stance: she expects the buyers to take her chosen strategy as given and try to collectively maximize their own payoff. To model the seller's concern for collusion, I introduce a subclass of public perfect equilibria, which I call *collusive public perfect equilibria.* For every public strategy of the seller I define the corresponding dynamic game among the buyers ("buyer-game") in which the reserve prices are chosen by Nature in accordance with the seller's strategy; I select only those public perfect equilibria of the repeated first-price auction game, in which the buyers' payoff is no smaller than the payoff they could achieve in the maximal strongly symmetric public perfect equilibrium of the corresponding buyer-game. I call the selected public perfect equilibria *collusive*. My main goal is to determine the highest payoff that the seller can obtain in a *collusive public perfect equilibrium* of the repeated auction game.

I consider buyers whose valuations are binary, independent and identically distributed across them and over time. The buyers in my model employ strongly symmetric strategies in any public perfect equilibrium of any buyer-game. In essence, the buyers are prohibited from using more complex asymmetric collusive schemes which might involve communication and/or bidding strategies dependent on each buyer's identity. While it is possible that the seller has less power against a more sophisticated cartel, it should be noted that asymmetric strategies (due to their complexity) might require explicit coordination among the buyers, and explicit coordination could be more easily detected and prevented via the traditional instruments of anti-trust policy. This paper finds a seller's strategy that is robust to collusive schemes that are simpler and more tacit, and thus harder to detect and prove to a court.

I study equilibrium outcomes as the discount factor goes to 1 and show that collusion in repeated auctions can be dealt with rather effectively: I establish that there is a *collusive public perfect equilibrium* that achieves full surplus extraction in the limit as the discount factor goes to 1, even though the seller can only set reserve prices, and stringent public disclosure requirements force her to publicly reveal bids in the end of each period. This *full-surplus-extracting collusive public perfect equilibrium* is stationary along the equilibrium path, features higher reserve prices than the static outcome and forces the buyers to bid even if their valuation is below the offered reserve price in the current period. Note that, since I am studying a restricted class of public perfect equilibria, my full surplus extraction results do not rely on any of the existing folk theorems. Since these theorems refer to the full set of public equilibrium payoffs, even the mere possibility of full surplus extraction by any *collusive public perfect equilibrium* (let alone by a cPPE of any particular structure) is not implied by the existing folk theorems.

In the full-surplus-extracting equilibrium the seller forces the buyer types to separate and punishes any off-equilibrium path deviations she can detect. In the corresponding buyer-game the buyers take the seller's threat as given and might try to deviate to a lower bidding profile. The key to the construction of the optimal equilibrium is in identifying the optimal symmetric joint deviation for the buyers and making sure that the original construction renders this joint deviation unprofitable. Since the full-surplus-extracting cPPE forces any low-type buyer to bid even when his valuation is below the reserve price, the optimal joint deviation will involve the low-type buyers abstaining from participating and receiving the punishment of zero continuation payoffs, and the high-type buyers bidding at the reserve price. There are three cases corresponding to different parameter values. In all three cases the seller extracts full surplus from the buyers. In Cases 1 and 2, the buyers' payoff in the full-surplus-extracting cPPE is exactly equal to the payoff of the optimal joint deviation. In Case 3 the proportion of the low-type buyers is so high that the optimal joint deviation would provide the buyers with a strictly lower payoff than the one they obtain along the equilibrium path.

Beyond addressing purely theoretical concerns, my results shed light on how collusion can be dealt with in practice. Note that dealing with collusion in repeated first-price auctions is especially challenging because of a fundamental conflict between revenue maximization and fighting collusion. A seller, who wants to maximize her revenue, must force the different valuation types of the buyers to separate, making the higher types bid relatively high. But separation of the different valuation types creates scope for collusion since, absent punishments, the buyers would try to coordinate on a lower bidding profile. Higher patience will only make this coordination process easier for them. What my results suggest, however, is that higher patience also allows the seller to come up with very effective punishments for colluding buyers. To effectively fight collusion, a revenue-maximizing seller should force the buyers to pay "upfront" for the continuation of favorable terms of trade, which is achieved by making the relatively low-valuation types participate even when they have to bid above their current valuations. Penalization of non-participation makes sure that the buyers cannot improve their payoff by making the lower types abstain from the auction altogether and making the higher types take their place in bidding low. Since the higher valuation types also want to avoid (inefficiently) pooling with the lower valuation types, they can do nothing but bid high.

1.1.1 Related literature

The dynamic nature of the interaction presents formidable challenges for an auction designer. Some of those challenges (e.g. intertemporal dependence of agents' private information) have been addressed by the dynamic mechanism design literature (see e.g. Pavan et al. (2014), and Bergemann and Välimäki (2019) for a review). Other important issues however remain. It is well-known that dynamic games often exhibit a multiplicity of equilibria, which makes the classical mechanism design assumption of favorable equilibrium selection harder to justify. For example, in repeated auction settings, collusive outcomes with lower revenue can be supported in equilibrium (see e.g. Skrzypacz and Hopenhayn (2004), who analyze equilibria of repeated first-price auctions and conclude that a bid rotation scheme, which leaves the seller with less revenue than optimal, can be supported even under limited observability of bids and auction outcomes). Moreover, collusive equilibria seem to be practically relevant as collusive bidding patterns are observed in many different repeated auction settings around the world (see e.g Chassang et al. (2021)).

Repeated auctions are special cases of general repeated games. Equilibria of repeated games were studied by Abreu et al. (1990), who provide a recursive characterization of equilibrium payoffs for repeated games with imperfect monitoring, and Fudenberg et al. (1994) who prove a folk theorem for these games. Athey et al. (2004) introduce (*iid*) private information into a repeated Bertrand game with imperfect monitoring and apply the recursive characterization of Abreu et al. (1990) to their game. They show that patient players can sustain high rigid prices in the optimal equilibrium, thus extracting a lot of surplus from the consumers. Their model can be translated to an auction setting with a passive seller who chooses a reserve price once and for all in the beginning of the game. In the buyer-optimal equilibrium with patient buyers such a seller would be forced to sell the good at her chosen reserve price in every period.

Even though the literature on collusion in repeated auctions and oligopolies with private information is very extensive (see Correia-da Silva (2017) for a review), very few papers are concerned with the study of how the seller's or auction designer's behavior might affect the buyers' collusion. Abdulkadiroglu and Chung (2004) consider a stage game design problem in which a committed seller proposes a mechanism that will become the stage game played repeatedly by a set of tacitly colluding buyers. The seller in their model is concerned with buyers coordinating on the buyer-optimal sequential equilibrium and designs the stage game accordingly. Similarly to my paper, Abdulkadiroglu and Chung (2004) find that there is a mechanism which extracts the entire surplus from the buyers. In the optimal mechanism all the buyers pay the same participation fee to the seller and then the partnership dissolution mechanism of Cramton et al. (1987) is run. Abdulkadiroglu and Chung (2004) however note that a non-committed seller will fall far short of full surplus extraction: in the buyer-optimal sequential equilibrium of the repeated game in which the seller moves first and proposes a mechanism, the seller's revenue will be zero. In this paper I propose a less pessimistic (from the seller's point of view) model of equilibrium coordination. While the seller in my model lacks long-term commitment, she is able to control her own strategy and does not have to coordinate on the worst equilibrium for herself. She cannot however guarantee that the buyers will coordinate on her preferred equilibrium either. The buyers could take her strategy as given and tacitly coordinate on a lower bidding profile using their continuation values to enforce collusive behavior, hence her equilibrium strategy must make such coordination unprofitable for the buyers. Although the seller has a more active role in equilibrium coordination in my model, she is more constrained in terms of feasible mechanisms: she must offer a first-price auction in every period and can only adjust reserve prices over time. The first-price auctions are widely used in practice, but give rise to severe challenges when it comes to collusive behavior under private information. A seller who wants to obtain a higher revenue should try to force the buyer types to separate, but that very separation creates a scope for collusion. I show that this conflict is resolved in favor of the seller.

A few other papers study similar settings, but none of them (to the best of my knowledge) simultaneously deals with the lack of seller's commitment and equilibrium coordination in a satisfactory way. Thomas (2005) notices that a seller could make collusion harder for the buyers by raising reserve prices, but assumes that the seller moves only once, in the beginning of time, and chooses one reserve price for the entirety of the repeated game between the buyers. Zhang (2021) studies a class of collusive agreements between bidders in a model of repeated first-price auctions, and, as a side note to his main results, shows how a revenue-maximizing seller should respond to collusion. His seller, much like the seller in Thomas (2005), moves only once and commits to a single reserve price. As the discount factor goes to 1, the seller is forced to tolerate "full collusion", in which all bids are suppressed down to the reserve price, and thus essentially makes an optimal take-it-or-leave-it offer to the colluding bidders. In contrast to the results in my paper, the revenue of a patient seller, who is restricted to choose only one reserve price once and for all, is lower than the revenue achieved under the infinite repetition of the competitive static outcome, and is therefore of course far below full surplus.

Ortner et al. (2020) are concerned with mitigating the effects of collusion in repeated procurement auctions. They propose a model with a regulator who observes the whole (infinite) bidding history and can punish colluding bidders. They construct tests for detecting collusive patterns of behavior which only allow for false negatives – therefore competitive bidders pass them with probability one. The regulator can then use the outcomes of the tests to punish the colluding bidders. My seller only has access to finite histories of bids and can only use reserve prices to punish colluding bidders.

Bergemann and Hörner (2018) also study a binary type model of first-price auctions similar to mine. The seller in their model is however passive and does not set a reserve price at all, and the buyers' valuations are perfectly persistent. They are concerned with disclosure regimes regarding the bid and winning history. In contrast to the findings in my paper, they show that the maximal disclosure regime leads to inefficient equilibria with low revenues. I show that an active seller who can adjust reserve prices over time can extract full surplus even when the full history of bids and identities of the winning buyers is publicly disclosed.

My paper is also related to the literature on collusion in static auctions. This literature was started by McAfee and McMillan (1992), who study outcomes of explicit before-auction communication in a first-price auction setting. They solve for optimal collusive schemes with ("strong collusion") and without transfers ("weak collusion"). In the optimal weak collusion scheme, the bidders bid at the reserve price as long as their valuation exceeds it and abstain otherwise. In the optimal strong collusion scheme, the colluding buyers can obtain a higher expected payoff by running a "knock-out" auction among themselves. The winner of the knock-out auction bids at the reserve price (as long as it exceeds his valuation) in the legitimate auction, and the losers are compensated for abstaining from the legitimate auction. It is however known now, that in the static setting the seller can avoid the dramatic losses from collusion via more sophisticated auction design. Che and Kim (2009) show that the second-best auction can be made collusion-proof, even when the bidders can use transfers to collude.

Finally, this paper speaks to the large literature on robustness in mechanism design (see Carroll (2019) for a comprehensive review). In my paper the seller aims to be robust to collusive behavior of the buyers.

1.1.2 Roadmap

The rest of the paper is organized as follows: Section 1.2 introduces the model of a repeated first-price auction game. In Section 1.3, I introduce the definitions of a *buyer-game* and a *collusive public perfect equilibrium*. In Section 1.4, I show how supporting collusive public perfect equilibria can be constructed to punish the seller and the buyers for deviations from the equilibrium path of *full-surplus-extracting cPPE* constructed in Sections 1.5 and 1.6. Section 1.7 briefly discusses the optimal reserve prices of the seller. Finally, Section 1.8 concludes.

1.2 Model

1.2.1 Setup

There is a seller (player 0) and $n \geq 2$ buyers (players $1, \dots, n$) who interact over infinitely many periods. The seller sells one unit of a private good in every period via a first-price auction with a reserve price. Each buyer is privately informed about his valuation type, which is drawn from a binary set $\Theta = \{\underline{\theta}, \overline{\theta}\}$, with $0 \leq \underline{\theta} < \overline{\theta}$, *iid* across periods and buyers. The probability of the low type $\underline{\theta}$ is $q \in (0, 1)$. The players share a common discount factor $\delta \in [0, 1)$.

The players play a repeated extensive form game with imperfect public monitoring. The timing of each period is as follows:

1. Seller announces a reserve price r.
2. Buyers privately learn their valuations for the good in the current period.
3. Buyers bid or abstain ($\emptyset$) in the first-price auction with the reserve price r.

4. The winner (if any) is determined, the buyers' choices are publicly disclosed.

The action set of the seller is $A_0 = \mathbb{R}_+$, the action set of each buyer is $A = \{\emptyset\} \cup \mathbb{R}_+$.

Buyer i's payoff is equal to his valuation θ_i net of his bid b_i if he wins the auction, and zero otherwise. Ties are broken by a fair coin toss. Formally,

$$u_i(r, b, \theta_i) = \begin{cases} \frac{1}{\#(\text{win})}(\theta_i - b_i), & \text{if } b_i \geq r \ \& \ \big(b_i = \max\{b_1, ..., b_n\} \text{ or } b_{-i} = \emptyset\big) \\ 0, & \text{otherwise} \end{cases},$$

where $\#(\text{win})$ stands for the number of winners in the auction, i.e. the number of buyers who placed the highest bid.

The seller's revenue is equal to the highest bid if there is a buyer who bids above the reserve price, and zero otherwise:

$$\mathcal{R}(r, b) = \begin{cases} b_i, & \text{if } b_i \geq r \ \& \ \big(b_i = \max\{b_1, ..., b_n\} \text{ or } b_{-i} = \emptyset\big) \\ 0, & \text{otherwise} \end{cases}.$$

1.2.2 One-shot auctions

Before we turn our attention to the repeated auction problem, we have to consider subgame perfect equilibria of the stage game. The intuition here is rather straightforward. If there are relatively few low types in the population (the probability q of having a low valuation is small), then the seller will prefer to trade with high types only, and will therefore set the reserve price equal to the high valuation $\overline{\theta}$. The low-type buyers will abstain while the high-type buyers will bid their valuation $\overline{\theta}$. If there are relatively many low types in the population, then the seller will prefer to trade with both types, and will therefore set the reserve price to the low valuation $\underline{\theta}$. The low-type buyers will bid their valuation while the high-type buyers will play a mixed strategy whose support lies above $\underline{\theta}$. The following proposition applies:

Proposition 1.1. One-shot auction equilibria

- *If the parameters of the model fall into the* ***High-reserve-price region*** $\big(q < \frac{n(\overline{\theta}-\underline{\theta})}{\underline{\theta}+n(\overline{\theta}-\underline{\theta})}\big)$, *then the seller sets* $r^*_{os} = \overline{\theta}$ *and generates revenue* $\mathcal{R}^*_{os} = (1 - q^n)\overline{\theta}$; *the buyers get the ex ante payoff* $v^*_{os} = 0$.

- *If the parameters of the model fall into the* ***Low-reserve-price region*** $\left(q \geq \frac{n(\bar{\theta}-\underline{\theta})}{\underline{\theta}+n(\bar{\theta}-\underline{\theta})}\right)$, *then the seller sets* $r^*_{os} = \underline{\theta}$ *and generates revenue* $\mathcal{R}^*_{os} = (1-q^n)\bar{\theta}+q^n\underline{\theta}-n(1-q)q^{n-1}(\bar{\theta}-\underline{\theta})$*; the buyers get the ex ante payoff* $v^*_{os} = (1-q)q^{n-1}(\bar{\theta}-\underline{\theta})$.

Its proof along with other details of equilibrium characterization is provided in Appendix A.1.

1.3 Collusive Public Perfect Equilibrium

1.3.1 Motivation

Let us consider the Low-reserve-price region and the infinite repetition of the associated one-shot equilibrium. Clearly, it is an equilibrium of the infinitely repeated auction game, but there is no reason to believe that the players will actually coordinate on it. In fact, buyers' collusion is a good reason to believe otherwise. Suppose that the seller sets the reserve price equal to the low valuation $\underline{\theta}$, but the buyers, instead of coordinating on their one-shot equilibrium strategies, use a different bidding profile, in which high-type buyers bid $\bar{b} = \underline{\theta}$ and the low-type buyers abstain $\underline{b} = \emptyset$ in every period. This bidding profile gives a lower revenue of $(1-q^n)\underline{\theta}$ to the seller and a higher payoff of $\frac{1}{n}(1-q^n)(\bar{\theta}-\underline{\theta})$ to the buyers. The buyers can support their new bidding profile using a "grim-trigger" strategy, which punishes deviations by moving back to the one-shot equilibrium strategies of the Low-reserve-price region; the buyers only have to make sure that the high types do not want to deviate to $\underline{\theta}+\epsilon$, i.e. whenever

$$\underbrace{(1-\delta)\frac{1-q^n}{n(1-q)}(\bar{\theta}-\underline{\theta})+\delta\frac{1}{n}(1-q^n)(\bar{\theta}-\underline{\theta})}_{\text{Payoff from } \bar{b}=\underline{\theta},\ \underline{b}=\emptyset} \geq \underbrace{(1-\delta)(\bar{\theta}-\underline{\theta})}_{\text{Today's deviation payoff}} + \underbrace{\delta(1-q)q^{n-1}(\bar{\theta}-\underline{\theta})}_{\text{Grim punishment payoff}},$$

which can be satisfied for high enough values of δ.

As we can see, the infinite repetition of the one-shot equilibrium in the Low-reserve-price region is not *"collusive"* because the buyers do not exploit their ability to collude to the fullest extent possible. A seller who has concerns about buyers' collusion should not hope to end up in such an equilibrium and needs to consider more sophisticated strategies. The seller's equilibrium strategy should however always guarantee that the buyers cannot improve their payoff similarly to how they did it in the above example. I formalize this requirement by introducing the concept of *collusive public perfect equilibrium*.

1.3.2 Definition

A *collusive public perfect equilibrium* is a strongly symmetric public perfect equilibrium that satisfies two novel requirements:

1. **Collusiveness on path.** The buyers must collude given the seller's on-path play of her equilibrium strategy. Central to this requirement is the notion of a *buyer-game* I introduce below. A buyer-game is a stochastic first-price auction game between the buyers, in which the reserve prices are determined according to the seller's strategy. *Collusiveness on path* requires that the buyers be unable to improve their payoff by moving to a different strongly symmetric public perfect equilibrium in the buyer-game induced by the seller's equilibrium strategy. In the above example of the infinite repetition of the one-shot equilibrium of the Low-reserve-price region *collusiveness on path* was violated since the buyers could improve their payoff by moving to a different equilibrium between themselves.
2. **Collusiveness off path.** The continuation play must be collusive on path in the above sense even following a seller's deviation as long as the buyers stick to their equilibrium strategies. This requirement formalizes the idea that buyers' collusive agreements cannot be broken by seller's actions. It does however allow non-collusive equilibria to be played following buyers' deviations and thus imposes no restriction on the buyers' ability to collude.

Strongly symmetric public perfect equilibrium is a public perfect equilibrium, in which buyers take symmetric actions on and off the equilibrium path. Public perfect equilibrium is an equilibrium in *public strategies*, i.e. strategies which map *public histories* into players' actions. A *public history* in the beginning of period $t+1$ is a sequence that includes all the actions taken by each player up to that period: $\big(\emptyset, (r_0, b_{10}, \ldots, b_{n0}), \ldots, (r_t, b_{1t}, \ldots, b_{nt})\big)$, where $\emptyset$ denotes the initial history. The set of those histories is given by $\mathcal{H}_0 \equiv \cup_{t=0}^{\infty}\big(A_0 \times A^n\big)^t$, with a typical period-$t$ history denoted h_0^t. Since buyers additionally observe the action taken by the seller in every period, the set of public histories at which they get to make a move is given by $\mathcal{H} \equiv \cup_{t=0}^{\infty}\big[\big(A_0 \times A^n\big)^t \times A_0\big]$ with a typical period-t history denoted h^t. A pure public strategy for the seller is a mapping $\sigma_0 : \mathcal{H}_0 \to A_0$, for the buyers it is $\sigma_i : \mathcal{H} \times \Theta \to A$.

The expected payoff of the seller in the repeated auction game is given by:

$$U_0(\sigma) = (1-\delta)\,\mathbb{E}\sum_{t=0}^{\infty}\delta^t \mathcal{R}\big(\sigma_0(h_0^t), \sigma_i(h^t, \theta_{it}), \sigma_{-i}(h^t, \theta_{-it})\big).$$

The expected payoff of the buyers $i = 1, 2, ..., n$ in the repeated auction game is given by:

$$U_i(\sigma) = (1-\delta)\,\mathbb{E}\sum_{t=0}^{\infty}\delta^t u_i\big(\sigma_0(h_0^t), \sigma_i(h^t, \theta_{it}), \sigma_{-i}(h^t, \theta_{-it}), \theta_{it}\big).$$

The above definitions extend naturally to behavioral strategies. We can now state the following definition:

Definition 1.1. Strongly symmetric public perfect equilibrium

A strategy profile $(\sigma_0^*, \sigma_1^*, ..., \sigma_n^*)$ is a strongly symmetric public perfect equilibrium if

1. *it induces a Nash equilibrium after every public history $h_0 \in \mathcal{H}_0$ and $h \in \mathcal{H}$;*
2. *$\sigma_i^*(h, \theta) = \sigma_j^*(h, \theta)$ after any public history $h \in \mathcal{H}$ for any two buyers i, j and any θ.*

The first condition of Definition 1.1 rules out non-credible threats at every public history much like subgame perfect equilibrium rules out non-credible threats in every subgame. The second condition makes sure that the buyers use symmetric bidding actions on and off the equilibrium path. Note that strongly symmetric public perfect equilibria have recursive structure: the continuation play after any public history is itself a strongly symmetric public perfect equilibrium.

All strongly symmetric public perfect equilibria I construct below, except the infinite repetition of the one-shot equilibrium in the Low-reserve-price region, satisfy the following additional assumption:

Assumption 1.1(a). Pure bidding actions along the equilibrium path

Buyers use pure bidding actions along the equilibrium path, i.e. after any public history $h \in \mathcal{H}$ consistent with the on-path play of $(\sigma_0^*, \sigma^*, \ldots, \sigma^*)$, the action $\sigma^*(h, \theta)$ is pure for both types $\theta \in \{\underline{\theta}, \overline{\theta}\}$.

Assumption 1.1(a) itself is not restrictive since we can find a full-surplus-extracting strongly symmetric public perfect equilibrium that belongs to the class of equilibria allowed by Assumption 1.1(a). However, I make a similar assumption in the next subsection (Assumption 1.1(b)), which forces the buyers to play the same class of equilibria in any buyer-game, restricting the set of collusive schemes they could use. It remains an open question whether Assumptions 1.1(a) and 1.1(b) could be dispensed with.

Collusiveness on path

To define *collusiveness on path* formally, we have to introduce the notion of a *buyer-game* induced by a seller's strategy. To define the states in the buyer-game, we need to define the *path automaton* of a seller's strategy[1]. In order to do that, fix a particular pure public strategy[2] of the seller σ_0. Let $\tilde{\mathcal{H}}_0(\sigma_0)$ be the set of histories consistent with the seller's play of σ_0 and any profile of buyers' strategies[3]. Two histories h_0 and h_0' from $\tilde{\mathcal{H}}_0(\sigma_0)$ are called σ_0-equivalent if they prescribe the same continuation play for the seller according to σ_0, i.e. $\sigma_0|_{h_0} = \sigma_0|_{h_0'}$. Let Ω be the resulting set of equivalence classes with ω^0 being the equivalence class of the initial history $\varnothing$. The path automaton representation of σ_0 is defined as follows:

Definition 1.2. Path automaton of a seller's strategy

The path automaton of σ_0 is the tuple $\left(\Omega, \omega^0, r, \tau\right)$, where

- *$r : \Omega \to A_0$ is the decision rule satisfying $r(\omega) = \sigma_0(h_0)$ for any $h_0 \in \omega$.*
- *$\tau : \Omega \times A^n \to \Omega$ is the transition function satisfying $\tau(\omega, b) = w'$ iff for any history $h_0 \in w$ the concatenated history $(h_0, r(\omega), b) \in w'$.*

We can now introduce the definition of the buyer-game induced by σ_0:

Definition 1.3. Buyer-game

Let $\left(\Omega, \omega^0, r, \tau\right)$ be the path automaton of σ_0. The buyer-game induced by σ_0 is a stochastic game between the buyers where:

[1]Unlike an automaton representation, the path automaton of a seller's strategy assumes that the seller never deviates from σ_0, and therefore represents only part of her repeated game strategy. See also Kandori and Obara (2006) who employ a similar definition of a path automaton in the context of repeated games with private monitoring.

[2]It is without loss of generality to restrict attention to pure strategies of the seller, since our goal is to construct a full-surplus-extracting collusive public perfect equilibrium, which can be achieved under this restriction.

[3]A typical element of $\tilde{\mathcal{H}}_0(\sigma_0)$ can be written as $h_0^t = \left(\varnothing,\ \left(\sigma_0(\varnothing),\ b_0\right),\ \left(\sigma_0(h_0^1),\ b_1\right),\ \ldots,\ \left(\sigma_0(h_0^{t-1}),\ b_{t-1}\right)\right)$; where $h_0^1 = \left(\sigma_0(\varnothing),\ b_0\right)$, $h_0^2 = \left(\left(\sigma_0(\varnothing),\ b_0\right),\ \left(\sigma_0(h_0^0),\ b_1\right)\right)$, etc.

- *The set of states is Ω, with the initial state ω^0. State transitions occur according to τ.*
- *The set of actions for each buyer is A, i.e. is as defined in the repeated auction game.*
- *The set of valuations for each buyer is Θ, i.e. is as defined in the repeated auction game.*
- *The utility of buyer i with type θ_i bidding b_i in state ω is*

$$\tilde{u}_i(\omega, b, \theta_i) = \begin{cases} \frac{1}{\#(\text{win})}(\theta_i - b_i), & \text{if } b_i \geq r(\omega) \ \& \ \big(b_i = \max\{b_1, ..., b_n\} \text{ or } b_{-i} = \emptyset\big) \\ 0, & \text{otherwise} \end{cases},$$

where $\#(\text{win})$ stands for the number of winners in the auction.

Let us look at the strongly symmetric public perfect equilibria of the buyer-game induced by σ_0. A public history at period $t+1$ in the buyer-game includes all states and bids up to period $t+1$: $\big(\omega_0, (b_{10}, \dots, b_{n0}), \dots, \omega_t, (b_{1t}, \dots, b_{nt}), \omega_{t+1}\big)$. Let $\mathbf{H}(\sigma_0)$ be the set of these public histories. A public strategy in the buyer game is a function $\rho_i : \mathbf{H}(\sigma_0) \times \Theta \to A$. This definition of public strategy extends naturally to behavior strategies. A strongly symmetric public perfect equilibrium in the buyer-game induced by a seller's strategy σ_0 is defined as follows:

Definition 1.4. Strongly symmetric public perfect equilibrium in the buyer-game

A strategy profile $(\rho_1^*, \dots, \rho_n^*)$ is a strongly symmetric public perfect equilibrium equilibrium of the buyer-game induced by σ_0 if

1. *It induces a Nash equilibrium after any public history $\mathbf{h} \in \mathbf{H}(\sigma_0)$.*
2. *$\rho_i^*(\mathbf{h}, \theta) = \rho_j^*(\mathbf{h}, \theta)$ after any public history $\mathbf{h} \in \mathbf{H}(\sigma_0)$ for any two buyers i, j and any θ.*

Recall that by Assumption 1.1(a) the buyers use pure bidding actions along the equilibrium path of any strongly symmetric public perfect equilibrium of the repeated auction game. The following Assumption 1.1(b) restrict the buyers to play equilibria from the same class in the buyer game.

Assumption 1.1(b). Pure bidding actions along the equilibrium path

Buyers use pure bidding actions along the equilibrium path in the buyer-game induced by σ_0, i.e. after any public history $\mathbf{h} \in \mathbf{H}(\sigma_0)$ consistent with the on-path play of $(\rho^*, \dots, \rho^*)$, *the action* $\rho^*(\mathbf{h}, \theta)$ *is pure for both types* $\theta \in \{\underline{\theta}, \overline{\theta}\}$.

Assumption 1.1(b) does not allow the buyers to collude by moving to a strongly symmetric public perfect equilibrium of the buyer game that exhibits mixed actions along the equilibrium path. It is in principle possible that the buyers could collectively benefit from using mixed actions along the equilibrium of the buyer-game induced by the full-surplus-extracting collusive equilibrium constructed below. It can be shown that the simplest collusive schemes with mixed actions do not help the buyers to improve their payoff[4]. The larger question of whether Assumption 1.1(b) could be dispensed with remains open.

We can now use the above definitions to formally introduce the notion of *collusiveness on path*.

Definition 1.5. Collusiveness on path

A strongly symmetric public perfect equilibrium $(\sigma_0^*, \sigma^*, ..., \sigma^*)$ *of the repeated auction game is collusive on path if there is no strongly symmetric public perfect equilibrium with pure actions along the equilibrium path (i.e. satisfying Assumption 1.1(b)) in the buyer-game induced by* σ_0^**, whose equilibrium payoff exceeds the buyer payoff from* $(\sigma_0^*, \sigma^*, ..., \sigma^*)$ *in the repeated auction game.*

Collusiveness off path

Recall that the requirement of *collusiveness off path* formalizes the idea that buyers' collusive agreements cannot be broken by seller's actions. More specifically, if the buyers have played their equilibrium actions up to the current period, then they must collude on path from the next period on no matter what the seller has played. The formal definition is as follows:

Definition 1.6. Collusiveness off path

Suppose $(\sigma_0^*, \sigma^*, \ldots, \sigma^*)$ *is a strongly symmetric public perfect equilibrium of the repeated auction game. Consider an alternative seller's strategy* σ_0' *and let* $h_0^t \in \mathcal{H}_0$ *be a period-*t *history consistent with the on-path play of* $(\sigma_0', \sigma^*, \ldots, \sigma^*)$*. If the continuation equilibrium* $(\sigma_0^*|_{h_0^t}, \sigma^*|_{h_0^t}, \ldots, \sigma^*|_{h_0^t})$ *is collusive on path for any such* h_0^t *and* σ_0'*, then* $(\sigma_0^*, \sigma^*, \ldots, \sigma^*)$ *is collusive off path.*

[4]For example, some stationary schemes, in which the high types mix over two bidding actions on path, do not improve the buyers' payoff because of their efficiency loss vis-à-vis fully separating behavior

We can now state the main definition:

Definition 1.7. Collusive public perfect equilibrium

A strongly symmetric public perfect equilibrium of the repeated auction game is a collusive public perfect equilibrium if it is collusive on and off path.

Remark 1.1. *Observe that the infinite repetition of the one-shot equilibrium in the High-reserve-price region is a collusive public perfect equilibrium in the sense of Definition 1.7. First of all it is clearly a strongly symmetric public perfect equilibrium. To show collusiveness on path, observe that the buyers get zero payoff along the equilibrium path, and it is not possible for them to improve their payoff once the seller's on path play is fixed: bidding below* $\overline{\theta}$ *leads to a zero payoff as well, bidding above* $\overline{\theta}$ *can only lead to losses. Since after a deviation by any player, the players return to the same equilibrium in the next period, it is also collusive off path.*

1.4 Supporting collusive equilibria

A seller who intends to actively fight collusion has to come up with punishment strategies for the buyers who are suspected of coordinating their bidding behavior. Since our ultimate goal is to construct a *collusive public perfect equilibrium*, in which the seller extracts the entire surplus from the buyers in the limit as the discount factor goes to 1, the punishment has to be as severe as possible. The most severe punishment that the seller could construct in principle involves leaving zero payoff to the buyers. Our goal in this section is to establish that a threat of such a severe punishment can be made credible if the discount factor is sufficiently high.

1.4.1 Repetition of the one-shot equilibrium in the High-reserve-price region

It is easy to see that the threat of severe punishment is immediately available to the seller if the parameters belong to High-reserve-price region. Since the one-shot equilibrium payoff of the buyers is already equal to zero, the seller can always reduce the equilibrium payoff of the buyers to zero, no matter what the value of δ is by switching to the infinite repetition of the one-shot equilibrium. Moreover since the equilibrium reserve price is extremely high, there is no room for collusion in this equilibrium:

Lemma 1.1. *Suppose that the parameters of the model belong to High-reserve-price region, then the infinite repetition of the equilibrium of Proposition 1.1 (with $r^* = \overline{\theta}$ in every period) is a collusive public perfect equilibrium in the sense of Definition 1.7.*

Proof. The buyers get zero payoff along the equilibrium path. It is not possible for them to improve their payoff once the seller's strategy is fixed: bidding below $\overline{\theta}$ is impossible, bidding above $\overline{\theta}$ can only lead to losses. □

1.4.2 Low-revenue collusive equilibria in the Low-reserve-price region

Suppose now that the parameters of the model belong to the Low-reserve-price region. Unlike in the High-reserve-price region, it might be harder for the seller to reduce the buyers' payoff to zero when she prefers trading with both types in the one-shot auction game. It nevertheless turns out to be possible when the seller is patient enough. To provide the appropriate punishments to the seller, I first construct *collusive public perfect equilibria* which leave the seller with little revenue. I will then use these equilibria to support a high-reserve-price equilibrium, in which the seller sets $r = \overline{\theta}$ along the equilibrium path and the buyers get zero equilibrium payoffs. This high reserve price equilibrium equilibrium will then be used to support the full-surplus-extracting equilibrium in Section 1.6.

Low-revenue separating equilibrium

I will now construct a separating equilibrium with low (but non-zero) revenue that can be supported for high enough discount factors. Since our aim is to find a low-revenue equilibrium, it is reasonable to try to force the seller to set $r = 0$ along the equilibrium path and have the low type of each buyer bid zero in every period. I denote the high type's bid by $\overline{b}$.

First, we have to make sure that the on-schedule incentive compatibility conditions are satisfied, i.e. that the low type does not want to emulate the behavior of the high type and vice versa. A low type $\underline{\theta}$ obtains in every period: $\frac{q^{n-1}}{n}\underline{\theta}$ and a high type's payoff in each period is: $\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b})$[5]. If a low type buyer attempts to mimic a high type buyer's behavior, his payoff is going to be: $\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b})$, thus the low type incentive compatibility is given by:

$$\frac{q^{n-1}}{n}\underline{\theta} \geq \frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b}),$$

which is equivalent to $\overline{b} \geq \frac{1-q^{n-1}}{1-q^n}\underline{\theta}$. Since we are attempting to minimize the seller's revenue, it is reasonable to select the minimal possible bid for a high type buyer:

$$\overline{b}^* = \frac{1-q^{n-1}}{1-q^n}\underline{\theta}.$$

The *ex ante* equilibrium payoff of each buyer:

$$v^*_{\text{lrs}} = \frac{1}{n}\left[(1-q^n)\left(\overline{\theta} - \frac{1-q^{n-1}}{1-q^n}\underline{\theta}\right) + q^n\underline{\theta}\right] = \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \tag{1.1}$$

The resulting revenue of the seller:

$$\mathcal{R}^*_{\text{lrs}} = (1-q^n)\frac{1-q^{n-1}}{1-q^n}\underline{\theta} + q^n 0 = (1-q^{n-1})\underline{\theta}.$$

Recall that in the one-shot equilibrium of the Low-reserve-price region in Proposition 1.1, the *ex ante* equilibrium payoff for each bidder is given by $v^*_{\text{os}} = (1-q)q^{n-1}(\overline{\theta}-\underline{\theta})$. Comparing the static equilibrium payoff in Proposition 1.1 and the payoff in (1.1), we obtain:

$$\begin{aligned}
v^*_{\text{lrs}} - v^*_{\text{os}} &= \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] - (1-q)q^{n-1}(\overline{\theta}-\underline{\theta}) \\
&= \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta} - n(1-q)q^{n-1}(\overline{\theta}-\underline{\theta})\right] \\
&= \frac{1}{n}\left[(1-q^n - n(1-q)q^{n-1})(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \\
&= \frac{1}{n}\left[\left((1-q)\sum_{k=0}^{n-1} q^k - n(1-q)q^{n-1}\right)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \\
&= \frac{1}{n}\left[(1-q)\left(\sum_{k=0}^{n-1} q^k - nq^{n-1}\right)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \\
&> \frac{1}{n}\left[(1-q)\left(nq^{n-1} - nq^{n-1}\right)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] = \frac{1}{n}q^{n-1}\underline{\theta} > 0,
\end{aligned}$$

[5]Interested readers will find the calculation of separating equilibrium payoffs in Appendix A.2.

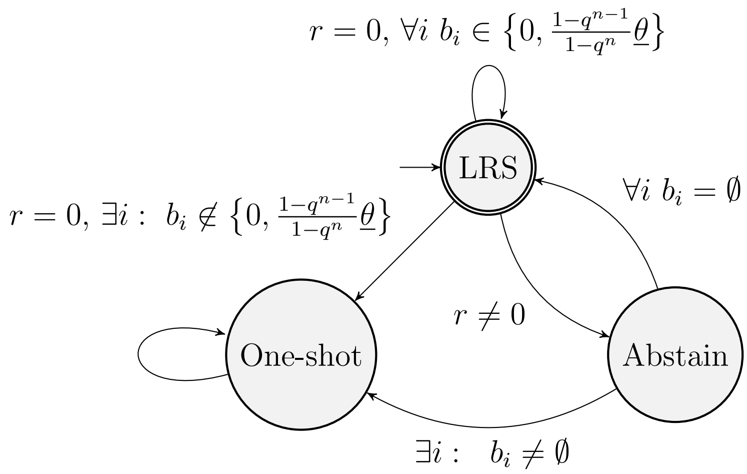

Figure 1.1: Low-revenue separating (LRS) strategy profile

which suggests that the chosen on-path behavior of the buyers can be supported by the threat of switching to the infinite repetition of the one-shot equilibrium. We can now formulate the full definition of the strategy profile:

Definition 1.8. Low-revenue separating strategy profile

(i) Along the equilibrium path:

(a) Seller sets $r^* = 0$,

(b) Any low-type buyer bids $\underline{b}^* = 0$,

(c) Any high-type buyer bids $\overline{b}^* = \frac{1-q^{n-1}}{1-q^n}\underline{\theta}$,

(ii) If at any history following $r = 0$ *in every period a bid outside of* $\{\underline{b}^*, \overline{b}^*\}$ *is made, then the game switches to the infinite repetition of the one-shot equilibrium of the Low-reserve-price region forever.*

(iii) Both buyer types abstain whenever $r > 0$.

(iv) After any history along which a positive bid has been observed following $r > 0$, *the game switches the infinite repetition of the one-shot equilibrium of the Low-reserve-price region forever.*

The low-revenue separating strategy profile is illustrated by Figure 1.1. The following proposition shows that the low-revenue separating strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game for high values of the discount factor.

Proposition 1.2. *Suppose that the parameters of the model belong to the Low-reserve-price region. There exists δ^* such that for all $\delta \in [\delta^*, 1)$ the low-revenue separating strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game. Along the equilibrium path the buyers will obtain the payoff of $v^*_{lrs} = \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right]$, and the seller will get $\mathcal{R}^*_{lrs} = (1-q^{n-1})\underline{\theta}$.*

Proof. Consider first the incentives of the seller. It is clear that the seller does not want to deviate: if she attempts a one-shot deviation to $r > 0$, her revenue will become $(1-\delta)0 + \delta(1-q^{n-1})\underline{\theta} = \delta(1-q^{n-1})\underline{\theta}$ (because all the buyers will abstain following $r > 0$), which can never exceed his equilibrium revenue of $(1-q^{n-1})\underline{\theta}$.

Now turn to the buyers. Consider first the public histories along which neither of the players has deviated. Incentive compatibility will require for a high-type buyer:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \\ \geq (1-\delta)\max\{q^{n-1}\overline{\theta}, \overline{\theta}-\overline{b}^*\} + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}),$$

and for a low type buyer:

$$(1-\delta)\frac{q^{n-1}}{n}\underline{\theta}+\delta\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \\ \geq (1-\delta)\max\{q^{n-1}\underline{\theta}, \underline{\theta}-\overline{b}^*\} + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}).$$

Before dealing with these constraints, consider a public history along which the seller has deviated to $r > 0$ in the current period. The equilibrium strategy of the buyers prescribes abstaining from participation if the reserve price is set above zero. The associated incentive compatibility condition of a high-type buyer is given by:

$$\delta\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] > (1-\delta)(\overline{\theta}-r) + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}).$$

For a low-type buyer it is given by:

$$\delta\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \geq (1-\delta)(\underline{\theta}-r) + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}).$$

The best deviation for $r > 0$ is the one for the high type and when $r \approx 0$. This deviation is unprofitable whenever:

$$\delta \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] \geq (1-\delta)\overline{\theta} + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}). \tag{1.2}$$

Notice that the incentive compatibility condition in (1.2) implies all of the above incentive compatibility conditions since the on-path payoff in each of them can only be higher and the deviation payoff can only be lower than in (1.2). The incentive compatibility condition in (1.2) is satisfied for all δ such that:

$$\delta > \frac{n\overline{\theta}}{n\overline{\theta} + \left(1-q^n-n(1-q)q^{n-1}\right)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}} \equiv \delta^*. \tag{1.3}$$

Since $1 - q^n - n(1-q)q^{n-1} > 0$, we can conclude that $\delta^* \in [0, 1)$. □

Zero-revenue pooling equilibrium

It is natural to ask the question whether the seller can be forced to give away the good in every period for free (clearly the worst possible outcome for the seller in this setup). That would require the seller to set the reserve price $r = 0$ along the equilibrium path and the buyers to bid $b^* = 0$ along the equilibrium path. The buyers' payoff would be equal to:

$$v^*_{\text{zrp}} = (1-q)\frac{1}{n}(\overline{\theta}-r) + q\frac{1}{n}(\underline{\theta}-r) = \frac{\mathbb{E}(\theta)}{n}. \tag{1.4}$$

Comparing the buyer's payoff in (1.4) to the *ex ante* payoff of the buyers in the low-revenue separating equilibrium, we get:

$$\begin{aligned}
v^*_{\text{lrs}} - v^*_{\text{zrp}} &= \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] - (1-q)\frac{1}{n}\overline{\theta} - q\frac{1}{n}\underline{\theta} \\
&= \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta} - (1-q)\overline{\theta} - q\underline{\theta}\right] \\
&= \frac{1}{n}\left[(q-q^n)\overline{\theta} + (-1+q^n+q^{n-1}-q)\underline{\theta}\right] \\
&= \frac{1}{n}\left[q(1-q^{n-1})\overline{\theta} - (1+q)(1-q^{n-1})\underline{\theta}\right] \\
&= \frac{1-q^{n-1}}{n}\left[q\overline{\theta} - (1+q)\underline{\theta}\right],
\end{aligned}$$

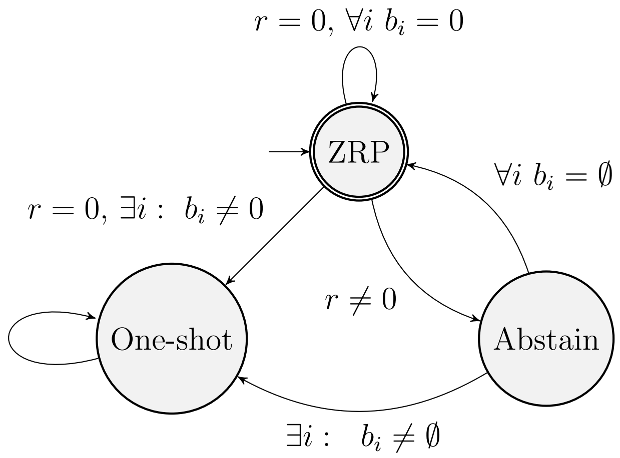

Figure 1.2: Zero-revenue pooling (ZRP) strategy profile

which means that $v^*_{\text{lrs}} < v^*_{\text{zrp}}$ whenever $q\overline{\theta} - (1+q)\underline{\theta} < 0$ or

$$q < \frac{\underline{\theta}}{\overline{\theta} - \underline{\theta}}. \tag{1.5}$$

We can now formulate the full definition of the zero-revenue pooling strategy profile:

Definition 1.9. Zero-revenue pooling strategy profile

(i) Along the equilibrium path

(a) Seller sets $r^* = 0$,

(b) Both buyer types bid 0,

(ii) If at any history following $r = 0$ in every period a bid $b \neq 0$ is placed, then the game switches to the infinite repetition of the one-shot equilibrium of the Low-reserve-price region forever.

(iii) Both buyer types abstain whenever $r > 0$.

(iv) After any history along which a bid has been observed following $r > 0$, the play of the game switches to the infinite repetition of the one-shot equilibrium of the Low-reserve-price region forever.

The zero-revenue pooling strategy profile is illustrated by Figure 1.2. The following proposition shows that the zero-revenue pooling strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game whenever the condition in (1.5) is satisfied.

Proposition 1.3. *Suppose that the parameters of the model belong to the Low-reserve-price region, and suppose further that the condition in (1.5) is satisfied, then there exists $\delta^* \in [0,1)$ such that for all $\delta > \delta^*$ the zero-revenue pooling strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game.*

Proof. Consider first the seller's incentives. The seller does not have any incentive to deviate because she would end up with zero revenue regardless of the reserve price, which makes setting $r = 0$ one of the optimal choices.

Consider now one of the buyers who is contemplating a deviation. Consider first a public history along which neither player has deviated, the best available deviation after such a history is for the high type to bid $0 + \epsilon$ for some small ϵ. This deviation will be detected by both the seller and the competing buyer. The competing buyer would then have to punish the deviator by switching to the infinite repetition of the one-shot equilibrium with $r = \underline{\theta}$ and competitive bidding, enforcing the continuation value of $(1-q)q^{n-1}(\overline{\theta} - \underline{\theta})$. The associated incentive compatibility condition for the high type is then given by:

$$(1-\delta)\frac{1}{n}\overline{\theta} + \delta\frac{\mathbb{E}(\theta)}{n} \geq (1-\delta)\overline{\theta} + \delta(1-q)q^{n-1}(\overline{\theta} - \underline{\theta}). \tag{1.6}$$

Consider now a public history along which the seller has deviated to $r > 0$ in the current period. The equilibrium strategy prescribes abstaining from participation for both buyers in the current period. The payoff from following the equilibrium strategy is thus $\delta\,\mathbb{E}(\theta)/n$. The best deviation available to the buyers is for the high type to bid r and get the good with the payoff of $\overline{\theta} - r$. Since this deviation is automatically detected by the seller and the competing buyers, the game then switches to the infinite repetition of the one-shot equilibrium, thus resulting in the incentive compatibility condition given by:

$$\delta\frac{\mathbb{E}(\theta)}{n} \geq (1-\delta)(\overline{\theta} - r) + \delta(1-q)q^{n-1}(\overline{\theta} - \underline{\theta}).$$

Clearly this deviation is most profitable when $r \approx 0$, therefore we could rule out all such deviations if we made sure that the following condition holds:

$$\delta \frac{\mathbb{E}(\theta)}{n} \geq (1-\delta)\overline{\theta} + \delta(1-q)q^{n-1}(\overline{\theta} - \underline{\theta}). \tag{1.7}$$

Recall now the no-deviation condition in (1.6). Clearly its left-hand side is strictly above the left-hand side of (1.7). As the respective right-hand sides are identical, it is obvious then that (1.7) implies (1.6). The condition in 1.7 is satisfied whenever

$$\delta \geq \frac{n\overline{\theta}}{n\overline{\theta} + q\underline{\theta} + (1-q)\overline{\theta} - n(1-q)q^{n-1}(\overline{\theta} - \underline{\theta})} \equiv \delta^*. \tag{1.8}$$

Note that the critical value of the discount factor δ^* defined in (1.8) is in $[0, 1)$ as long as $q\underline{\theta} + (1-q)\overline{\theta} - n(1-q)q^{n-1}(\overline{\theta} - \underline{\theta}) = n(v^*_{\text{zrp}} - v^*_{\text{os}})$ is strictly positive. Recall that the payoff from the low revenue separating equilibrium v^*_{lrs} always exceeds the one-shot equilibrium payoff v^*_{os}. Under the assumption that $q < \frac{\underline{\theta}}{\overline{\theta} - \underline{\theta}}$ in (1.5) we have $v^*_{\text{zrp}} > v^*_{\text{lrs}} > v^*_{\text{os}}$, which establishes the claim. □

Observe that both the low-revenue separating equilibrium and the zero-revenue pooling equilibrium lead to the same buyer-game. This buyer game is a repeated first-price auction game in which the reserve price is set to zero. In an optimal strongly symmetric public perfect equilibrium of this game the buyers either pool or separate along the equilibrium path. If they pool, then their optimal equilibrium payoff is equal to the buyers' payoff in the zero-revenue pooling equilibrium. If they separate, then their optimal equilibrium payoff is equal to the buyer's payoff in the low-revenue pooling equilibrium. Thus, depending on the parameter values, either the zero-revenue pooling equilibrium is collusive, or the low-revenue separating equilibrium is collusive. The following proposition, whose proof is relegated to Appendix A.3, establishes this claim formally.

Proposition 1.4. *If $q \geq \frac{\underline{\theta}}{\overline{\theta} - \underline{\theta}}$, then the low-revenue separating equilibrium of Proposition 1.2 is collusive in the sense of Definition 1.7, otherwise the zero-revenue pooling equilibrium of Proposition 1.3 is collusive in the sense of Definition 1.7.*

Proof. See Appendix A.3. □

1.4.3 High-reserve-price equilibrium in the Low-reserve-price region

Having constructed equilibria with low revenue in the previous sections, we can now proceed to characterize some of the high(er) revenue equilibria in which the seller actively fights collusion among the buyers. Suppose that the seller sets $r = \overline{\theta}$ along the equilibrium path. Clearly the optimal response of the buyers is to bid $\overline{\theta}$ for the high type and to abstain for the low type. This equilibrium therefore leaves zero rents to the buyers, but is inefficient and therefore does not allow the seller to extract full surplus. It does, however, allow the seller to credibly threaten the buyers with zero continuation value (as does the repetition of the one-shot equilibrium in the High-reserve-price region). In the full-surplus-extracting equilibria of Section 1.6 the buyers can therefore be incentivized to give up almost the entire surplus along the equilibrium path.

The on-path behavior in this equilibrium can be supported either by the threat of switching to the low-revenue separating equilibrium or by the threat of switching to the zero-revenue pooling equilibrium. The full definition of the strategy profile is as follows:

Definition 1.10. High-reserve-price strategy profile

(i) At any history in which the seller has always set $r^* = \overline{\theta}$

- *(a) The seller sets* $r^* = \overline{\theta}$*,*
- *(b) Any low-type buyer abstains,*
- *(c) Any high-type buyer bids* $\overline{\theta}$*.*

(ii) If $q \geq \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$ *(low-revenue separating equilibrium is collusive), then*

- *Following any observation of* $r < \overline{\theta}$ *in period* t*, the buyers abstain in period* t *and the low-revenue separating equilibrium is played from period* $t+1$ *on.*
- *Following any observation of* $r < \overline{\theta}$ *in period* t*, if any of the buyers fails to abstain in period* t*, the one-shot equilibrium of the Low-reserve-price region is infinitely repeated from period* $t+1$ *on.*

(iii) If $q < \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$ *(zero-revenue pooling equilibrium is collusive), then*

- *Following any observation of* $r < \overline{\theta}$ *in period* t*, the buyers abstain in period* t *and the zero-revenue pooling equilibrium is played from period* $t+1$ *on.*

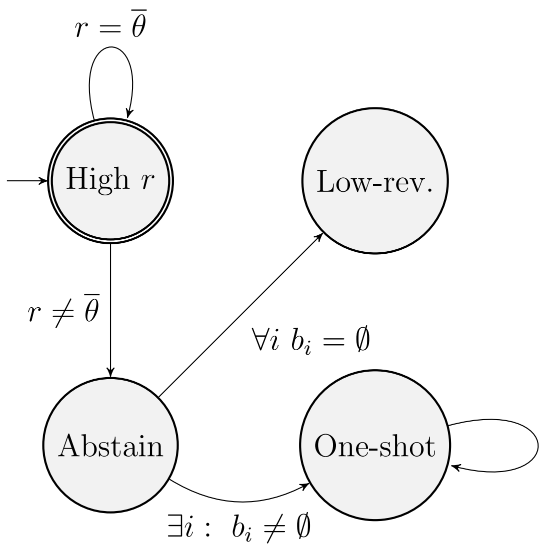

Figure 1.3: High-reserve-price (High r) strategy profile

- *Following any observation of $r < \overline{\theta}$ in period t, if any of the buyers fails to abstain in period t and places a positive bid above r, the one-shot equilibrium of the Low-reserve-price region is infinitely repeated from period $t+1$ on.*

The high-reserve-price strategy profile is illustrated by Figure 1.3. The following proposition shows that it is a strongly symmetric public perfect equilibrium of the repeated auction game for high values of the discount factor.

Proposition 1.5. *Suppose that the parameters of the model belong to the Low-reserve-price region, then there exists $\delta^* \in [0, 1)$ such that for all $\delta > \delta^*$ the high-reserve-price strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game. The buyers get the payoff $v^*_{\text{hrp}} = 0$, the seller gets the revenue of $\mathcal{R}^*_{\text{hrp}} = (1 - q^n)\overline{\theta}$.*

Proof. (ii) *Low-revenue separating equilibrium is collusive*

It is easy to see that the seller does not want to deviate in any period. Along the equilibrium path, her revenue is equal to $\mathcal{R}^*_{\text{hrp}} = (1 - q^n)\overline{\theta}$. If she deviates to any $r < \overline{\theta}$, then her revenue is $(1 - \delta)0 + \delta(1 - q^{n-1})\underline{\theta} = \delta(1 - q^{n-1})\underline{\theta} < (1 - q^n)\overline{\theta}$.

Buyers get zero payoffs along the equilibrium path. Following $r = \overline{\theta}$ neither type wants to deviate: bidding leads to a negative payoff for the low type in the current period, and abstaining does not improve the payoff of the high type in the current period. It remains to make sure that buyers do not want to deviate from the proposed strategy following an observation of a lower reserve price $r < \overline{\theta}$. It is required that both types prefer abstaining in the current period and playing the low-revenue separating equilibrium to bidding r (the lowest possible bid) and playing the one-shot equilibrium in the continuation game, i.e. for type $\theta_i \in \{\underline{\theta}, \overline{\theta}\}$

$$\delta \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] > (1-\delta)(\theta_i - r) + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}).$$

The best deviation obtains for the high type at $r = 0$:

$$\delta \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\underline{\theta}) + q^{n-1}\underline{\theta}\right] > (1-\delta)\overline{\theta} + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}),$$

which is the same condition as in (1.2) satisfied for all δ defined in (1.3).

(iii) *Zero-revenue pooling equilibrium is collusive*

Just as in the previous case, the seller's revenue is equal to $\mathcal{R}^*_{\text{hrp}} = (1-q^n)\overline{\theta}$. She does not want to deviate since deviation leads to zero revenue forever.

As before the best deviation is for a high type buyer whenever the seller deviates to a reserve price $r > 0$ near zero. The condition is:

$$\delta \frac{\mathbb{E}(\theta)}{n} \geq (1-\delta)\overline{\theta} + \delta(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}),$$

which is identical to the no deviation condition in (1.7), and therefore leads to the same threshold for the discount factors as in (1.8).

□

The following corollary is immediate:

Corollary 1.1. *The high-reserve-price equilibrium of Proposition 1.5 is collusive in the sense of Definition 1.7.*

Proof. Holding the seller's equilibrium strategy fixed, it is impossible for the two buyers to improve their payoff even if they perfectly coordinate: bidding higher leads to negative payoffs, bidding lower is impossible. □

Since this equilibrium leaves the buyers with zero payoffs, we can now use it to construct full-surplus-extracting equilibria by threatening the buyers who deviate off-schedule with zero continuation values.

1.5 High-revenue collusive equilibria

In this section I will introduce a class of collusive public perfect equilibria that allow the seller to extract full surplus in the limit as δ goes to 1. These equilibria are stationary and separating along the equilibrium path, i.e. in each of them any low-type buyer bids $\underline{b}$, and a high-type buyer bids $\bar{b}$, while the seller sets the reserve price to $r = \underline{b}$ in every period along the equilibrium path. The full description of the class of strategy profiles I am considering is given by the following definition.

Definition 1.11. High-revenue strategy profile

Fix a pair of bids $(\bar{b}, \underline{b})$. The corresponding ***high-revenue strategy profile*** *is described as follows.*

(i) Along the equilibrium path

- *Seller sets a reserve price equal to the equilibrium bid of a low type buyer $r = \underline{b}$,*
- *Any low-type buyer bids $\underline{b}$.*
- *Any high-type buyer bids $\bar{b}$.*

(ii) If the parameters of the model belong to the High-reserve-price-region, then

- *If at any history following r in every period a bid outside of $\{\underline{b}, \bar{b}\}$ is placed, the play of the game switches to the infinite repetition of the one-shot equilibrium of the High-reserve-price region forever.*
- *If in period t the seller sets $r' \neq r$, then the buyers play the one-shot equilibrium with reserve price r' in period t, and the play of the game switches to the infinite repetition of the one-shot equilibrium of the High-reserve-price region forever.*

(iii) If the parameters of the model belong to the Low-reserve-price region, then

- *If at any history following* r *in every period a bid outside of* $\{\underline{b}, \overline{b}\}$ *is placed, then the play of the game switches to the high-reserve-price equilibrium of Proposition 1.5 forever.*
- *both types abstain in period* t *if* $r' \neq r$ *is observed in period* t*, and from* $t+1$ *on the play of the game game switches to the low-revenue separating equilibrium when it is collusive, (i.e. when* $q \geq \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$*) or to the zero-revenue pooling equilibrium when it is collusive (i.e. when* $q < \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$*).*
- *After any history along which a bid has been observed following* $r' \neq r$*, the game switches to the infinite repetition of the one-shot equilibrium of he Low-reserve-price region forever.*

The high-revenue strategy profile is illustrated by Figure 1.4.

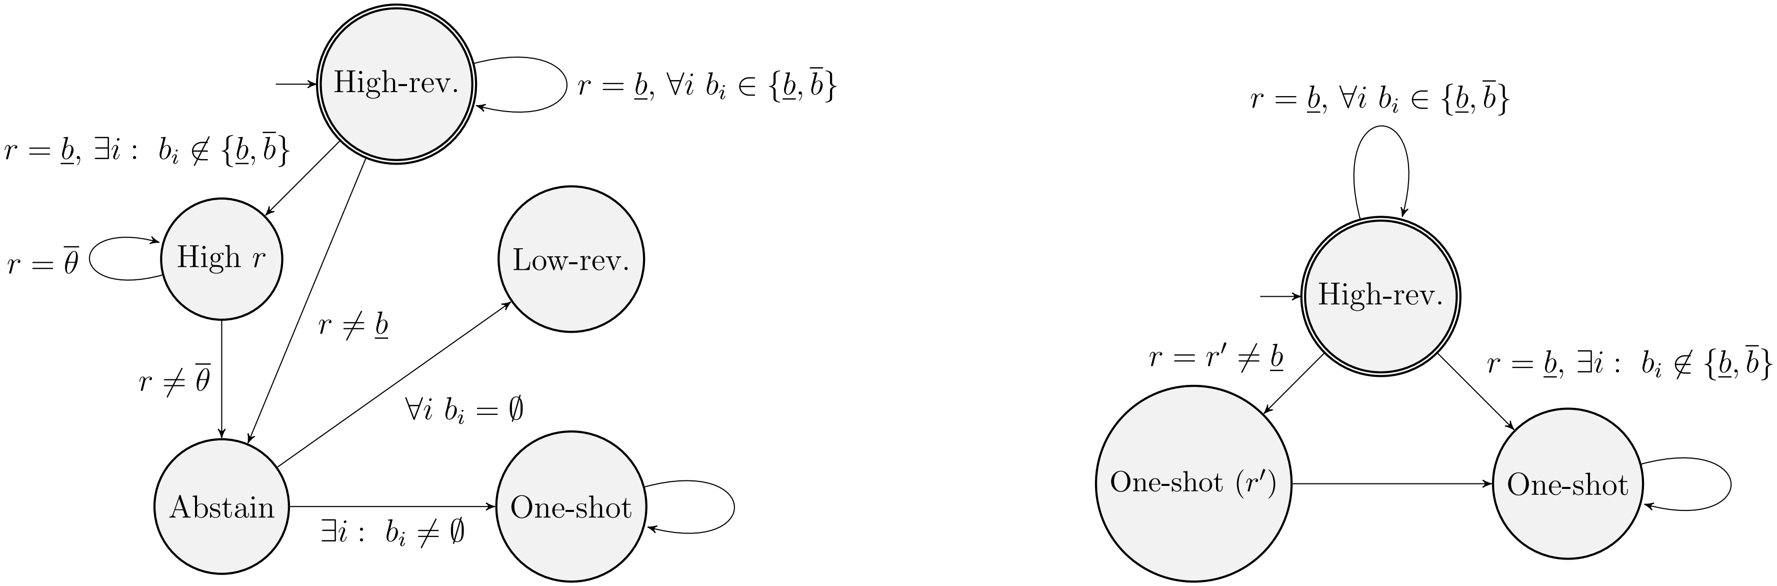

(a) Low-reserve-price region

(b) High-reserve-price region

Figure 1.4: High-revenue strategy profile

Having discussed the structure of the high-revenue strategy profiles, I can set up the following *revenue maximization problem*:

$$\mathcal{RM}: \ \mathcal{R}^*_{\text{fse}} \equiv \max_{\bar{b},\underline{b},v} (1-q^n)\bar{b} + q^n\underline{b}, \quad \text{s.t.}$$

$$\text{(Eq-payoff)} \qquad v = \frac{1}{n}\left[(1-q^n)(\bar{\theta}-\bar{b}) + q^n(\underline{\theta}-\underline{b})\right];$$

Incentive constraints:

$$\text{(LowIC)} \qquad (1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}) + \delta v \geq 0,$$

$$\text{(HighIC-up)} \qquad (1-\delta)\frac{1-q^n}{n(1-q)}(\bar{\theta}-\bar{b}) + \delta v \geq (1-\delta)(\bar{\theta}-\bar{b}),$$

$$\text{(HighIC-down)} \qquad (1-\delta)\frac{1-q^n}{n(1-q)}(\bar{\theta}-\bar{b}) + \delta v \geq (1-\delta)q^{n-1}(\bar{\theta}-\underline{b}),$$

$$\text{(HighIC-on-sch)} \qquad \frac{1-q^n}{n(1-q)}(\bar{\theta}-\bar{b}) \geq \frac{q^{n-1}}{n}(\bar{\theta}-\underline{b});$$

No-collusion constraints:

$$\text{(No-col-sep-1)} \qquad v \geq \frac{(1-\delta)(1-q^n)(\bar{\theta}-\underline{b})}{n(1-\delta(1-q)^n)},$$

$$\text{(No-col-sep-2)} \qquad v \geq \frac{(1-\delta)\left[(1-q^n)(\bar{\theta}-\underline{b}) + q^n(\underline{\theta}-\underline{b})\right]}{n(1-\delta q^n)},$$

$$\text{(No-col-pool)} \qquad v \geq \frac{1}{n}\left[(1-q)(\bar{\theta}-\underline{b}) + q(\underline{\theta}-\underline{b})\right];$$

A solution to the revenue maximization problem in $\mathcal{RM}$ is a pair of bids $(\bar{b}^*, \underline{b}^*)$ together with a buyer payoff v^*_{fse}. In the next lemma, I will show that the high-revenue strategy profile corresponding to $(\bar{b}^*, \underline{b}^*)$ is a *collusive public perfect equilibrium* for high enough values of δ as long as the solution to $\mathcal{RM}$ induces a well-defined separating equilibrium (i.e. $\bar{b}^* > \underline{b}^*$), the low-type buyers bid strictly above their valuation (i.e. $\underline{b}^* > \underline{\theta}$), and the seller achieves a higher revenue than in the high reserve price equilibrium (i.e. $\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\bar{\theta}$). In Section 1.6, I will solve $\mathcal{RM}$, verify that its solution satisfies the aforementioned conditions for sufficiently high values of δ, and show that the maximal revenue goes to full surplus as δ goes to 1.

Lemma 1.2. *Suppose $(\bar{b}^*, \underline{b}^*, v^*_{\text{fse}})$ solve the revenue maximization problem $\mathcal{RM}$. Suppose further that $\underline{\theta} < \underline{b}^* < \bar{b}^*$ and $\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\bar{\theta}$, then the high-revenue strategy profile corresponding to $(\bar{b}^*, \underline{b}^*)$ (as defined by 1.11) is a collusive public perfect equilibrium of the* *repeated auction game if*

(i) the parameters of the model belong to the High-reserve-price region; or

(ii) the parameters of the model belong to the Low-reserve-price region and δ satisfies

- *condition (1.3) if $q \geq \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$ (low-revenue separating equilibrium is collusive),*
- *condition (1.8) if $q < \frac{\underline{\theta}}{\overline{\theta}-\underline{\theta}}$ (zero-revenue pooling equilibrium is collusive).*

Proof. Let us show first that the high-revenue strategy profile is a strongly symmetric public perfect equilibrium of the repeated auction game. Strong symmetry follows from Definition 1.11, thus we only need to check the players' incentives. I start with the buyers.

Incentive compatibility of the buyers. Consider histories in which every player has stayed on the equilibrium path up to period t. Suppose first that the parameters of the model fall into High-reserve-price region (i.e. $r = \overline{\theta}$ is optimal in the one-shot game). If the seller deviates in period t, the play from $t+1$ is a public perfect equilibrium by construction. Since the buyers receive zero continuation values from $t+1$ on, they will play the one-shot equilibrium in period t for a given reserve price as if the game ends tomorrow, hence the buyers do not want to deviate in period t. Suppose now that the parameters of the model fall into the Low-reserve-price region (i.e. $r = \underline{\theta}$ is optimal in the one-shot game). If the seller deviates in period t, then the equilibrium strategy dictates that the buyers abstain in period t. Since a buyer's deviation triggers the switch to the infinite repetition of the one-shot equilibrium of Low-reserve-price region, it is not profitable for the buyers as long as δ satisfies conditions (1.3) or (1.8) by the argument employed in the construction of the low-revenue separating or zero-revenue pooling equilbria respectively.

Suppose now that the seller does not deviate in period t, and consider the buyers' incentives. Let us start with on-schedule deviations, i.e. attempts to mimic the behavior of the other type. The on-schedule deviation is unprofitable of a low-type buyer as long as:

$$\underbrace{\frac{q^{n-1}}{n}(\underline{\theta} - \underline{b}^*)}_{\text{Equilibrium reward}} \geq \underbrace{\frac{1-q^n}{n(1-q)}(\underline{\theta} - \overline{b}^*)}_{\text{Mimic the high type}}.$$

This incentive compatibility condition is satisfied since $\underline{\theta} < \underline{b}^* < \overline{b}^*$ by assumption: if a low-type buyer deviates to $\overline{b}^*$, then he receives a lower payoff with a higher probability, which cannot be profitable. The on-schedule deviation is unprofitable for a high-type buyer as long as:

$$\underbrace{\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)}_{\text{Equilibrium reward}} \geq \underbrace{\frac{q^{n-1}}{n}(\overline{\theta}-\underline{b}^*)}_{\text{Mimic the low type}},$$

which is the incentive constraint (HighIC-on-sch) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*)$, and is therefore satisfied.

Consider now off-schedule deviations. First of all, we must make sure that a low-type buyer is actually willing to participate in the auction as opposed to abstaining and getting the forever punishment of high reserve price, i.e. that the following condition is satisfied:

$$\underbrace{(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}^*)+\delta v^*_{\text{fse}}}_{\text{Equilibrium payoff of a low-type buyer}} \geq (1-\delta)\underbrace{0}_{\text{Abstain today}} + \delta\underbrace{0}_{\text{Switch to } r=\overline{\theta} \text{ forever}} = 0,$$

which is the incentive constraint (LowIC) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$. If a low-type buyer deviates to a higher off-schedule bid, then he receives a negative expected reward in the period of the attempted deviation (since $\underline{\theta} < \underline{b}^*$) and zero continuation value, which cannot be profitable for someone who receives a positive payoff along the equilibrium path. We can therefore conclude that the remaining off-schedule incentive constraints of a low-type buyer are satisfied at $\left(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}}\right)$.

Consider now high-type buyers who contemplate off-schedule deviations. A high-type buyer could deviate upwards which would guarantee him winning the auction with probability 1. The best upward deviation is to $\overline{b}^* + \epsilon$ which gives the deviating high-type buyer a payoff almost equal to $\overline{\theta} - \overline{b}^*$. For this deviation to be unprofitable, we must have:

$$\underbrace{(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}}}_{\text{Equilibrium payoff of a high-type buyer}} \geq (1-\delta)\underbrace{(\overline{\theta}-\overline{b}^*)}_{\text{Deviate to } \overline{b}^*+\epsilon} + \delta\underbrace{0}_{\text{Switch to } r=\overline{\theta} \text{ forever}} = (1-\delta)(\overline{\theta}-\overline{b}^*),$$

which is the incentive constraint (HighIC-up) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$.

A high-type buyer could also deviate downwards and win the auction only in the case when all his competitors are low-type buyers, that is with probability q^{n-1}. In this case the best deviation is to $\underline{b}^* + \epsilon$ with a payoff almost equal to $\overline{\theta} - \underline{b}^*$. For this deviation to be unprofitable, we must have:

$$\underbrace{(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*) + \delta v^*_{\text{fse}}}_{\text{Equilibrium payoff of a high type buyer}}$$
$$\geq (1-\delta)\underbrace{q^{n-1}(\overline{\theta}-\underline{b}^*)}_{\text{Deviate to }\underline{b}^*+\epsilon} + \delta \underbrace{0}_{\text{Switch to } r=\overline{\theta} \text{ forever}} = (1-\delta)q^{n-1}(\overline{\theta}-\overline{b}^*),$$

which is the incentive constraint (HighIC-down) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$.

Incentive compatibility of the seller. Consider now the seller's incentives. Recall that we have $\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\overline{\theta}$ by assumption. If the parameters of the model belong to the High-reserve-price region, a deviating seller would receive the payoff of $(1-\delta)\,\mathcal{R}^*_r + \delta(1-q^n)\overline{\theta}$ where $\mathcal{R}^*_r$ is the revenue achieved by the seller in the one-shot auction game with the reserve price equal to r. In the High-reserve-price region the optimal reserve price for the seller is $r = \overline{\theta}$ with the associated revenue of $(1-q^n)\overline{\theta}$. Thus a deviating seller would not be able to get more than $(1-\delta)(1-q^n)\overline{\theta} + \delta(1-q^n)\overline{\theta} = (1-q^n)\overline{\theta}$ which cannot exceed $\mathcal{R}^*_{\text{fse}}$. If the parameters of the model belong to Low-reserve-price region, a deviating seller would receive either 0 (if the zero-revenue pooling equilibrium is collusive), or $\delta(1-q^{n-1})\underline{\theta}$ (if the low revenue separating equilibrium is collusive), neither of which can exceed $\mathcal{R}^*_{\text{fse}}$.

Other histories. Neither the seller nor the buyers want to deviate after any of the other histories by construction of continuation equilibria, hence the high-revenue strategy profile corresponding to the bids $(\overline{b}^*, \underline{b}^*)$ is a strongly symmetric public perfect equilibrium.

Buyer-game. We must make sure that the public perfect equilibrium we have constructed is indeed collusive in the sense of Definition 1.7. To do that, we shall consider the buyer-game induced by the seller's equilibrium strategy. This buyer game is a stochastic game with two states. The game starts in the low reserve price state ω^l, in which the reserve price is equal to $r(\omega^l) = \underline{b}^*$, and remains in that state unless a bid outside of $\{\underline{b}^*, \overline{b}^*\}$ is placed by at least one buyer, in which the game transitions to the high reserve price state ω^h, in which the reserve price is $r(\omega^h) = \overline{\theta}$. The high reserve price state is absorbing, i.e. once the high reserve price state is achieved, the game remains in that state forever. The full definition of this high-revenue buyer-game is as follows:

Definition 1.12. ***High-revenue buyer-game***

- *The set of states is* $\Omega = \{\omega^l, \omega^h\}$*, the initial state is* $\omega^0 = \omega^l$*.*
- *The set of actions for each buyer is* A*, i.e. as defined in the repeated auction game.*
- *The transitions between states occur according to* τ*:*

$$\tau(\omega^l, b) = \begin{cases} \omega^l, & \text{if } b \in \{\underline{b}^*, \overline{b}^*\}^n \\ \omega^h, & \text{otherwise} \end{cases},$$
$$\tau(\omega^h, b) = \omega^h, \ \forall b.$$

- *The set of valuations for each buyer is* Θ*, i.e. is as defined in the repeated auction game.*
- *The utility of buyer* i *with type* θ_i *bidding* b_i *in state* ω *is*

$$\tilde{u}_i(\omega, b, \theta_i) = \begin{cases} \frac{1}{\#(\text{win})}(\theta_i - b_i), & \text{if } b_i \geq r(\omega) \ \& \ \big(b_i = \max\{b_1, ..., b_n\} \text{ or } b_{-i} = \emptyset\big) \\ 0, & \text{otherwise} \end{cases},$$

where $\#(\text{win})$ *stands for the number of winners in the auction.*

The definition of collusive public perfect equilibria (Definition 1.7) requires that the buyers be unable to play a strongly symmetric public perfect equilibrium of the high-revenue buyer-game in Definition 1.12 that improves their payoff. I first show that the buyers' strategy in any strongly symmetric public perfect equilibrium of the buyer-game must be monotonic:

Lemma 1.3. Monotonicity lemma

Consider the high-revenue buyer-game in Definition 1.12. Any strongly symmetric public perfect equilibrium of this buyer-game satisfies monotonicity: pick any history of play that leads to state ω^l*, if* $\overline{b}$ *is the equilibrium bidding action of a high-type buyer and* $\underline{b}$ *is the equilibrium bidding action of a low-type buyer after that history, then* $\overline{b} \geq \underline{b}$*.*

Proof. See Appendix A.4. □

The Monotonicity lemma shows any high-type buyer must always place a higher bid than any low type buyer in any symmetric public perfect equilibrium of the buyer-game whenever the current state is ω^l. Recall that when the current state is ω^h, the reserve price is equal to $\overline{\theta}$, and thus the buyers cannot get more than zero in any continuation equilibrium is that state. Since they cannot get a negative payoff in any continuation equilibrium either, they must be getting zero once the game is stuck in state ω^h. As I restrict attention to pure strategies along the equilibrium path, the resulting *ex ante* payoff from bidding $(\overline{b}, \underline{b})$ in state ω^l is given by:

$$\hat{u}_{\omega^l}(\overline{b}, \underline{b}) \equiv \begin{cases} \frac{1}{n}\left[(1-q^n)(\overline{\theta}-\overline{b}) + q^n(\underline{\theta}-\underline{b})\right] & \text{if } \overline{b} > \underline{b} \\ \frac{1}{n}\left[(1-q)(\overline{\theta}-\underline{b}) + q(\underline{\theta}-\underline{b})\right] & \text{if } \overline{b} = \underline{b} \end{cases}$$

where whenever $b < r(\omega^l)$, the convention is to set $\theta - b = 0$ for the respective type $\theta \in \{\overline{\theta}, \underline{\theta}\}$

Consider now the optimal collusion problem in the high-revenue buyer-game and ignore all the aspects of incentive compatibility except monotonicity. Since all the remaining incentive compatibility constraints are ignored, the following maximization problem provides an upper bound on symmetric equilibrium payoffs in the buyer-game:

$$\max_{\{\overline{b}_t, \underline{b}_t\}_{t=0}^{+\infty}} (1-\delta) \sum_{t=0}^{+\infty} \delta^t \hat{u}_\omega(\overline{b}_t, \underline{b}_t) \quad \text{s.t.} \tag{1.9}$$

(i) $\overline{b}_t \geq \underline{b}_t$,

(ii) Transition function τ.

where $\hat{u}_{\omega^l}(\overline{b}, \underline{b})$ is defined above, and $\hat{u}_{\omega^h}(\overline{b}, \underline{b})$ is assumed to be equal to zero without loss of generality. The optimization problem in (1.9) is a Markov decision problem. It follows from Blackwell (1965) that, if this problem has a solution, it must also have a stationary solution. I therefore consider two kinds of stationary monotonic bidding profiles: separating and pooling.

Separating profiles. Suppose first that the buyers coordinate on a separating bidding profile in the high-revenue buyer-game under consideration. If both types bid on schedule, then clearly there is only one option: $\underline{b} = \underline{b}^*$ and $\overline{b} = \overline{b}^*$ with the payoff equal to v^*_{fse}. *If all buyers of type $\overline{\theta}$ bid on schedule and all buyers of type $\underline{\theta}$ bid off schedule*, then the off-schedule action of any low-type buyer will be immediately detected by the seller and punished with zero continuation values. Since the punishment will not occur if and only if all buyers have high types (i.e. with probability $(1-q)^n$), the resulting payoff will be:

$$v = (1-\delta)\frac{1}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}) + q^n(\underline{\theta}-\underline{b})\big] + \delta(1-q)^n v.$$

Recall that we assume $\underline{b}^* > \underline{\theta}$, hence by incentive compatibility we must have $\overline{b}^* < \overline{\theta}$. Then the optimal solution here is to coordinate on the bidding profile in which any high-type buyer bids the low equilibrium bid $\underline{b}$ and any low-type buyer abstains, i.e. choose $\underline{b}^* = \emptyset$ and $\overline{b} = \underline{b}^*$, which results in the payoff:

$$v(\underline{b}^*, \emptyset) = \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\underline{b}^*)\big]}{n(1-\delta(1-q)^n)}.$$

The no-collusion constraint (No-col-sep-1) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ gives us $v^*_{\text{fse}} \geq v(\underline{b}^*, \emptyset)$.

If all buyers of type $\overline{\theta}$ bid off schedule and all buyers of type $\underline{\theta}$ bid on schedule, then the off-schedule action of any high-type buyer will be immediately detected by the seller and punished with zero continuation values. Since the punishment will not occur if and only if all buyers have low types (i.e. with probability q^n), the resulting payoff will be:

$$v = (1-\delta)\frac{1}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}) + q^n(\underline{\theta}-\underline{b})\big] + \delta q^n v,$$

which can be solved for v' to get:

$$v = \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\overline{b}) + q^n(\underline{\theta}-\underline{b})\big]}{n(1-\delta q^n)}.$$

The optimal solution here is for the low types to choose $\underline{b} = \underline{b}^*$ and for the high types to choose $\overline{b} = \underline{b}^* + \epsilon$, with the resulting payoff of:

$$v(\underline{b}^* + \epsilon, \underline{b}^*) = \frac{(1-\delta)\left[(1-q^n)(\overline{\theta} - \underline{b}^*) + q^n(\underline{\theta} - \underline{b}^*)\right]}{n(1-\delta q^n)}.$$

The no-collusion constraint (No-col-sep-2) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ gives us $v^*_{\text{fse}} \geq v(\underline{b}^* + \epsilon, \underline{b}^*)$.

If buyers of both types bid off schedule, then the seller will punish them in the first period with probability 1, and the resulting payoff will be:

$$v = (1-\delta)\frac{1}{n}\left[(1-q^n)(\overline{\theta} - \overline{b}) + q^n(\underline{\theta} - \underline{b})\right] + \delta 0.$$

Since it must be that $\overline{b}^* < \overline{\theta}$, the best bidding profile here is for the high types to choose $\overline{b} = \underline{b}^* + \epsilon$ and for the low types to choose $\underline{b} = \emptyset$ with the payoff of:

$$v(\underline{b}^* + \epsilon, \emptyset) = (1-\delta)\frac{1}{n}(1-q^n)(\overline{\theta} - \underline{b}^*),$$

which is clearly below $v(\underline{b}^*, \emptyset)$ and therefore below v^*_{fse}.

Pooling profiles. The buyers might find it optimal to pool instead of separating. *If the buyers pool on schedule*, then their collusive scheme is never detected by the seller. Clearly the optimal pooling on schedule is achieved at $\underline{b}^*$ with the resulting payoff of:

$$v(\underline{b}^*, \underline{b}^*) = \frac{1}{n}\left[(1-q)(\overline{\theta} - \underline{b}^*) + q(\underline{\theta} - \underline{b}^*)\right] \tag{1.10}$$

The no-collusion constraint (No-col-pool) of the revenue maximization problem $\mathcal{RM}$ evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ gives us $v^*_{\text{fse}} \geq v(\underline{b}^*, \underline{b}^*)$.

Note that the payoff from *pooling off-schedule* cannot exceed v^*_{fse}. If the buyers coordinate on any off-schedule bid above $\underline{b}^*$ they will get a fraction of the payoff in 1.10 since they will be punished by the seller with probability 1. Abstaining from the auction altogether cannot be optimal as long as $v^*_{\text{fse}} \geq 0$, which it is by incentive compatibility.

We therefore conclude that no strongly symmetric public perfect equilibrium payoff in the high-revenue buyer game corresponding to $(\underline{b}^*, \overline{b}^*)$ can exceed v^*_{fse}, and therefore the high-revenue strategy profile corresponding to $(\underline{b}^*, \overline{b}^*)$ is a *collusive public perfect equilibrium* of the repeated auction game in the sense of Definition 1.7. □

1.6 Full surplus extraction

Let us now solve the revenue maximization problem $\mathcal{RM}$. There are three cases depending on which constraints are binding; the parameter values corresponding to each case are illustrated by Figure 1.5. In Case 1, (No-col-sep-1) and (LowIC) constraints are binding with both types being indifferent between their payoff in the full-surplus-extracting cPPE and the payoff they could have obtained by coordinating on the bidding profile $(\underline{b}^*, \emptyset)$. Case 1 does not always apply because its solution candidate does not always satisfy the (HighIC-up) incentive compatibility constraint: if n is high enough, the winning probability of a high-type buyer is so low that such a buyer would prefer to win with probability 1 by placing a slightly higher bid and suffer the punishment of zero continuation values. We therefore have to consider Case 2, in which (HighIC-up) and (No-col-sep-1) are binding and the remaining constraints are slack. Case 2 equilibrium candidate in turn does not apply for high values of q: in this case the (HighIC-down) incentive compatibility constraint will be violated. Intuitively, if the mass of low types is sufficiently large, then a high type buyer will have a fairly high chance of winning by bidding just above the low type equilibrium bid even though placing such a bid is severely punished. In Case 3, only (HighIC-up) and (HighIC-down) are binding, and the remaining constraints are slack, which implies that the buyers do not have a strict incentive to collude.

The remaining constraints in the revenue maximization problem are never binding. Consider first the on-schedule incentive compatibility constraint of a high-type buyer (HighIC-on-sch). This constraint essentially puts an upper bound on the high-type equilibrium bid (if a high-type buyer is asked to bid a lot more than a low-type buyer, he might find it profitable to deviate to the low-type bid and get a much higher reward with a smaller winning probability), but we have already included a constraint that does the same, the no-collusion constraint (No-col-sep-1). Indeed, if a high-type buyer is asked to place a very high bid in every period, then the buyers might find it profitable to collude on a lower bidding profile, and such a collusion scheme is prevented by (No-col-sep-1). The restriction on equilibrium

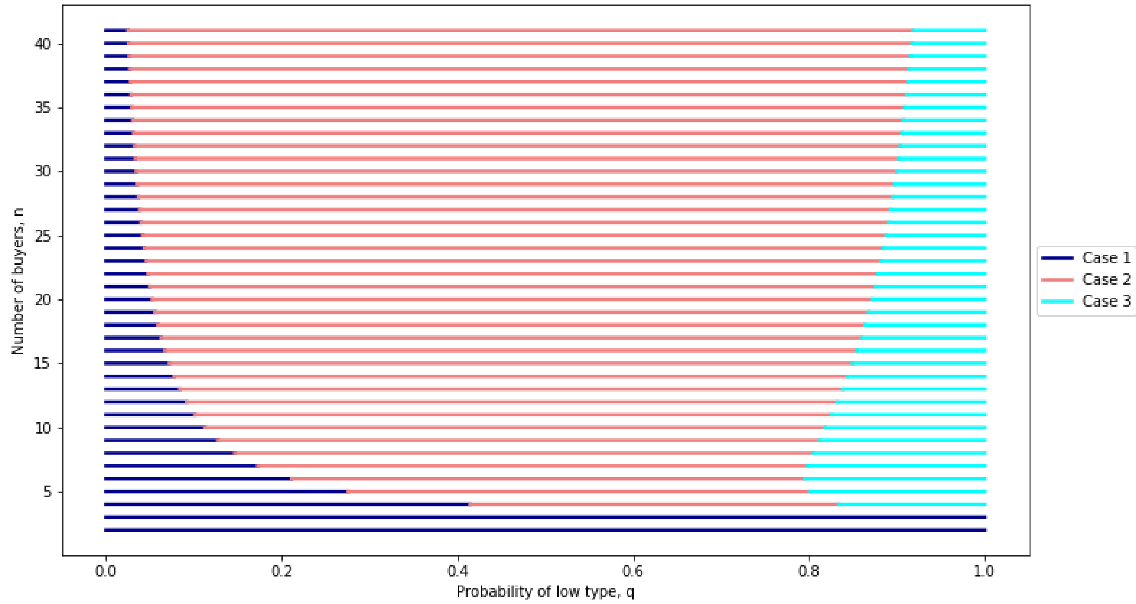

Figure 1.5: *Parameters corresponding to Cases 1, 2, and 3. For each number of buyers n, the respective line shows which values of q belong to Cases 1, 2, and 3.*

bids imposed by (No-col-sep-1) is more severe than the one imposed by the on-schedule incentive compatibility of a high type buyer. Clearly, if the more severe restriction were the one imposed by incentive compatibility, we would be unlikely to consider collusion an important problem in an auction setting with adverse selection.

The two remaining no-collusion constraints, (No-col-sep-2) and (No-col-pool), are also non-binding in all three cases, which means that the optimal optimal collusion scheme for the buyers always involves bidding $\underline{b}^*$ for the high types and abstaining for the low types. Collusion by pooling on schedule turns out to be particularly inefficient as it leads to negative payoffs for the buyers for δ close to 1, while the buyers' payoff in the full-surplus-extracting cPPE is non-negative by construction. Collusion by leaving the low types on schedule and moving the high types off schedule does not outperform the optimal collusion scheme because it leads to punishments for the high types, who, as opposed to the low types, get a positive payoff in every period. The gain from bidding lower made by the high types in this collusion scheme is completely offset by the severity of the seller's punishment.

In the following subsections I will construct the solutions to the revenue maximization problem $\mathcal{RM}$ in each of the three cases. I will show that the revenue-maximizing bidding profiles can indeed be supported in the collusive public perfect equilibrium with the corresponding high-revenue strategy profiles (as defined by 1.11), and derive the conditions on the parameters of the model for each of the three cases. In all three cases the seller will be able to extract full surplus from the buyers in the limit as the discount factor δ goes to 1.

Case 1: High expected valuation/Small number of buyers

Recall that in Case 1, the no-collusion constraint (No-col-sep-1) and the low-type incentive compatibility constraint (LowIC) bind at the optimum of the revenue maximization problem $\mathcal{RM}$.

Full surplus extraction cPPE, Case 1.

- *Equilibrium conditions:*

$$\begin{aligned} &\text{(No-col-sep-1)} && v^*_{\text{fse}} = \frac{(1-\delta)(1-q^n)(\overline{\theta}-\underline{b}^*)}{n(1-\delta(1-q)^n)}, \\ &\text{(LowIC)} && (1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}^*)+\delta v^*_{\text{fse}} = 0, \\ &\text{(Eq-payoff)} && v^*_{\text{fse}} = \frac{1}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big]. \end{aligned}$$

- *Parameter restriction:*

$$q < \frac{1-q^n}{n(1-q)}.$$

The solution to this system of equilibrium conditions is provided in Appendix A.5.1. I will derive the condition on the parameters in the course of proving Proposition 1.6 below. The resulting equilibrium payoff for a low-type buyer conditional upon winning with $\underline{b}^*$ is:

$$\underline{\theta}-\underline{b}^* = \frac{-\delta q\big(1-q^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}; \tag{1.11}$$

for a high-type buyer winning with $\overline{b}^*$ we have:

$$\overline{\theta}-\overline{b}^* = \frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}; \tag{1.12}$$

and for a high-type buyer winning with $\underline{b}^*$

$$\overline{\theta} - \underline{b}^* = \frac{q^n\big(1-\delta(1-q)^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}. \tag{1.13}$$

The *ex ante* equilibrium payoff is:

$$v^*_{\text{fse}} = \frac{1}{n}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}. \tag{1.14}$$

The equilibrium bids can be immediately computed from the payoffs in 1.11 and 1.12:

$$\underline{b}^* = \underline{\theta} + \frac{\delta q\big(1-q^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}, \tag{1.15}$$

$$\overline{b}^* = \overline{\theta} - \frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}. \tag{1.16}$$

I first show that the equilibrium bids in (1.15) and (1.16) satisfy the condition of Lemma 1.2.

Lemma 1.4. $\underline{\theta} < \underline{b}^* < \overline{b}^*$.

Proof. (i) $\underline{\theta} < \underline{b}^*$ is equivalent to $\underline{\theta} - \underline{b}^* < 0$, which is true since $-\delta q(1-q^n)(\overline{\theta}-\underline{\theta}) < 0$.

(ii) $\underline{b}^* < \overline{b}^*$ is equivalent to $\overline{\theta} - \underline{b}^* > \overline{\theta} - \overline{b}^*$, which is true since $1-\delta(1-q)^n > 1-\delta(1-q)$ because $(1-q)^n < (1-q)$ for any $q \in (0,1)$ and $n \geq 2$.

□

I now show that the bids in (1.15) and (1.16) can in fact be supported in a *collusive public perfect equilibrium* for a high values of δ:

Proposition 1.6. *Suppose that $q < \frac{1-q^n}{n(1-q)}$. Suppose further that $\underline{b}^*$ and $\overline{b}^*$ are as defined in (1.15) and (1.16) respectively, then there exists a critical discount factor δ^*, such that for all $\delta \in [\delta^*, 1)$ the high-revenue strategy profile corresponding to $(\overline{b}^*, \underline{b}^*)$ (as defined by 1.11) is a collusive public perfect equilibrium of the repeated auction game in the sense of Definition 1.7. Moreover, the seller achieves full surplus extraction in the limit as δ goes to* 1.

Proof sketch. The complete proof is provided in Appendix A.7.1. Here I briefly sketch the main arguments. Recall that by Lemma 1.2 and Lemma 1.4, it is enough to check that

$\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\overline{\theta}$ and that the remaining constraints in the revenue maximization problem $\mathcal{RM}$ are satisfied at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ for high enough δ. I start with the seller's revenue.

Seller's revenue. The seller's revenue is equal to the full surplus net of the equilibrium payoff of the buyers:

$$\mathcal{R}^*_{\text{fse}} = (1-q^n)\overline{\theta} + q^n\underline{\theta} - nv^*_{\text{fse}}.$$

Recall that nv^*_{fse} is given by:

$$nv^*_{\text{fse}} = \frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} \xrightarrow[\delta\to 1]{} 0.$$

and therefore the seller extracts full surplus in the limit as δ goes to 1 and $\mathcal{R}^*_{\text{fse}} \approx (1-q^n)\overline{\theta} + q^n\underline{\theta}$ for δ close enough to 1, which clearly exceeds $(1-q^n)\overline{\theta}$.

Incentive constraints. All of the remaining incentive constraints in in the revenue maximization problem $\mathcal{RM}$ are non-binding at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ for all δ high enough and all values of q and n, except the incentive constraint (HighIC-up). There is a region of q and n, where this constraint cannot be satisfied even for δ close to 1. To see why, observe that (HighIC-up) can be rewritten as:

$$\delta v^*_{\text{fse}} \geq (1-\delta)\left(1 - \frac{1-q^n}{n(1-q)}\right)(\overline{\theta} - \overline{b}^*)$$

Plugging the respective payoffs from (1.12) and (1.14) in, we obtain:

$$\frac{\delta}{n}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} \geq (1-\delta)\left(1 - \frac{1-q^n}{n(1-q)}\right)\frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)},$$

which simplifies to:

$$\delta \geq \frac{1}{1-q} - \frac{1-q^n}{n(1-q)^2}. \tag{1.17}$$

The condition on δ identified in (1.17) can only be satisfied if the right-hand side of this inequality is strictly below 1, which is only true whenever:

$$q < \frac{1-q^n}{n(1-q)},$$

which gives is satisfied in Case 1 by assumption.

No-collusion constraints. We check that the no-collusion constraints (No-col-sep-2) and (No-col-pool) are satisfied, or, in other words, that in the corresponding buyer-game pooling at $\underline{b}^*$ and bidding $(\underline{b}^* + \epsilon, \underline{b}^*)$ does not improve the buyers' payoff. If the buyers decide to bid $(\underline{b}^* + \epsilon, \underline{b}^*)$ in the buyer-game, their payoff will be:

$$\begin{aligned} v(\underline{b}^* + \epsilon, \underline{b}^*) &= \frac{(1-\delta)\big[(1-q^n)(\overline{\theta} - \underline{b}^*) + q^n(\underline{\theta} - \underline{b}^*)\big]}{n(1-\delta q^n)} \\ &= \frac{(1-\delta) q^n (1-q^n)\big(1 - \delta(1-q)^n - \delta q\big)(\overline{\theta} - \underline{\theta})}{n(1-\delta q^n)\big(\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)\big)}. \end{aligned}$$

We must make sure that that $v^*_{\text{fse}} \geq v(\underline{b}^* + \epsilon, \underline{b}^*)$, which is equivalent to:

$$1 \geq \frac{1 - \delta(1-q)^n - \delta q}{1 - \delta q^n} \Leftrightarrow (1-q)^n \geq -q + q^n,$$

which is true since the right-hand side of $(1-q)^n \geq -q + q^n$ is strictly negative, and the left-hand side is strictly positive.

If the buyers coordinate on pooling at $\underline{b}^*$ in the buyer-game, they will obtain:

$$\begin{aligned} v(\underline{b}^*, \underline{b}^*) =& \frac{1}{n}\big[(1-q)(\overline{\theta} - \underline{b}^*) + q(\underline{\theta} - \underline{b}^*)\big] = \\ &= \frac{\big((1-q) q^n \big(1 - \delta(1-q)^n\big) - \delta q^2 (1-q^n)\big)(\overline{\theta} - \underline{\theta})}{n\big(\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)\big)}. \end{aligned} \tag{1.18}$$

Consider the numerator of (1.18) in the limit as δ goes to 1:

$$\begin{aligned} &(1-q) q^n \big(1 - (1-q)^n\big) - q^2(1-q^n) \\ &= (1-q)\left[q^n\big(1-(1-q)^n\big) - q^2 \sum_{k=0}^{n-1} q^k\right] \\ &= (1-q)\left[- q^n (1-q)^n - q^2 \sum_{k=0}^{n-3} q^k - q^{n+1}\right] < 0 \end{aligned}$$

Hence the payoff from pooling at $\underline{b}^*$ in (1.18) is strictly negative for all δ sufficiently close to 1, and therefore cannot exceed v^*_{fse} for δ around 1. □

Case 1: the restriction on the parameters

The full surplus extraction equilibrium of Proposition 1.6 can only be sustained if $q < \frac{1-q^n}{n(1-q)}$. It is easy to check that this condition can be satisfied for any q as long as $n = 2$ or $n = 3$, but only for some q if $n \geq 4$. Indeed consider $n = 2$ first:

(I) $n = 2$. In this case the condition becomes:

$$2q < \frac{1-q^2}{1-q} \Leftrightarrow 2q < 1 + q \Leftrightarrow q < 1,$$

which is obviously true.

(II) $n = 3$. In this case the condition becomes:

$$3q < \frac{1-q^3}{1-q} \Leftrightarrow 3q < 1 + q + q^2 \Leftrightarrow 0 < 1 - 2q + q^2 \Leftrightarrow 0 < (1-q)^2,$$

which is also obviously true for any $q \in (0, 1)$.

(III) $n = 4$ In this case the condition becomes:

$$4q < \frac{1-q^4}{1-q} \Leftrightarrow 4q < 1 + q + q^2 + q^3 \Leftrightarrow 0 < 1 - 3q + q^2 + q^3$$
$$\Leftrightarrow 0 < (1-q)(-q^2 - 2q + 1) \Leftrightarrow 0 < -q^2 - 2q + 1,$$

which is only true for $q \in (0, -1 + \sqrt{2})$.

It it however possible to establish that for any number of players n there will be some values of q falling into Case 1:

Proposition 1.7. *The equation $1 - q^n = nq(1-q)$ has a unique solution q^* on $(0, 1)$ for any $n \geq 4$. Moreover for all $q < q^*$ it is true that $q < \frac{1-q^n}{n(1-q)}$ and vice versa.*

Proof. See Appendix A.8.1. □

The above proposition essentially shows that for every $n \geq 4$ the restriction divides the interval $(0, 1)$ into two parts. In the left part of the segment one will find the values of q that fall into Case 1, and in the right part of the segment one will find the values of q that fall into Cases 2 and 3. Figure 1.5 provides an illustration and also suggests that, as n goes to infinity, lower and lower values of q fall into Case 1 until there are none left in the limit. Indeed, it is easy to see that

$$\lim_{n\to\infty} nq(1-q)-(1-q^n)=+\infty,$$

implying that, for any fixed value of q, the parameter restriction does not hold for all sufficiently high n.

Case 2: Medium expected valuation

Recall that in Case 2, the no-collusion constraint (No-col-sep-1) and the upward incentive compatibility constraint of a high-type buyer (HighIC-up) bind at the optimum of the revenue maximization problem $\mathcal{RM}$.

Full surplus extraction cPPE, Case 2.

- *Equilibrium conditions:*

$$\begin{aligned}
&\text{(No-col-sep-1)} && v^*_{\text{fse}}=\frac{(1-\delta)(1-q^n)(\overline{\theta}-\underline{b}^*)}{n(1-\delta(1-q)^n)},\\
&\text{(HighIC-up)} && (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}}=(1-\delta)(\overline{\theta}-\overline{b}^*),\\
&\text{(Eq-payoff)} && v^*_{\text{fse}}=\frac{1}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big].
\end{aligned}$$

- *Parameter restrictions:*

$$\begin{aligned}
&q\geq\frac{1-q^n}{n(1-q)},\\
&(1-q^n)(1-q)>q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big].
\end{aligned}$$

The complete solution to the system of equilibrium conditions is provided in Appendix A.5.2. I will derive the restrictions on the parameters in the course of the proof of Proposition 1.8. First, define $D(\delta)$ as:

$$D(\delta)=q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]+(1-q^n)\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big].$$

The payoff of a low-type buyer who wins by bidding $\underline{b}^*$ is given by:

$$\underline{\theta} - \underline{b}^* = -\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta} - \underline{\theta}). \tag{1.19}$$

The payoff of a high type buyer who wins by bidding $\overline{b}^*$ is given by:

$$\overline{\theta} - \overline{b}^* = \frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta} - \underline{\theta}), \tag{1.20}$$

and the payoff of a high type buyer who wins by bidding $\underline{b}^*$ is given by:

$$\overline{\theta} - \underline{b}^* = \frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q) - (1-q^n)\big](\overline{\theta} - \underline{\theta}). \tag{1.21}$$

The resulting *ex ante* equilibrium payoff of the buyers is:

$$v^*_{\text{fse}} = \frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q) - (1-q^n)\big](\overline{\theta} - \underline{\theta}). \tag{1.22}$$

Note that as δ goes to 1, $D(\delta)$ goes to:

$$D(1) = q^n\big(1-(1-q)^n\big)\big[n(1-q) - (1-q^n)\big] + (1-q^n)(1-q^n)(1-q) > 0,$$

hence we can conclude that $D(\delta)$ is strictly positive for all δ sufficiently close to 1[6].

The equilibrium bids of each type can be computed from the payoffs in (1.19) and (1.20):

$$\underline{b}^* = \underline{\theta} + \frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta} - \underline{\theta}), \tag{1.23}$$

$$\overline{b}^* = \overline{\theta} - \frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta} - \underline{\theta}). \tag{1.24}$$

I first show that the equilibrium bids satisfy the condition of Lemma 1.2.

Lemma 1.5. *Suppose $q \geq \frac{1-q^n}{n(1-q)}$, and δ is sufficiently close to 1, then $\underline{\theta} < \underline{b}^* < \overline{b}^*$.*

[6] More precisely, for all δ satisfying

$$\delta > \frac{(1-2q^n)\big(n(1-q) - (1-q^n)\big)}{(1-q^n)^2\big(n(1-q) - q(1-q^{n-1})\big)}$$

Proof. (i) To see that $\underline{\theta} < \underline{b}^*$ for sufficiently high δ, observe that

$$\underline{\theta} - \underline{b}^* \xrightarrow[\delta \to 1]{} -\frac{1}{D(1)}(1-q^n)(1-q)(1-q^n)(\overline{\theta} - \underline{\theta}) < 0.$$

(ii) The proof of $\underline{b}^* < \overline{b}^*$ is provided in Appendix A.6.

□

I now proceed to establish that the bidding profile $(\overline{b}^*, \underline{b}^*)$ can indeed be played along the equilibrium path of a *collusive public perfect equilibrium* of the repeated auction game:

Proposition 1.8. *Suppose that* $q \geq \frac{1-q^n}{n(1-q)}$ *and* $(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]$. *Suppose further that* $\overline{b}^*$ *and* $\underline{b}^*$ *are defined by (1.23) and (1.24) respectively, then there exists a critical discount factor* δ^**, such that for all* $\delta \in [\delta^*, 1)$ *the high-revenue strategy profile corresponding to* $(\overline{b}^*, \underline{b}^*)$ *(as defined by 1.11) is a collusive public perfect equilibrium of the repeated auction game in the sense of Definition 1.7. Moreover, the seller achieves full surplus extraction in the limit as* δ *goes to* 1*.*

Proof sketch. The complete proof is provided in Appendix A.7.2. As in the previous case, I only provide a sketch of the main argument in the main text. By Lemma 1.2 and Lemma 1.5, it is enough to check that $\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\overline{\theta}$ and that the remaining constraints in the revenue maximization problem $\mathcal{RM}$ are satisfied at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ for high enough δ. I start with the seller's revenue.

Seller's revenue. The seller's revenue is equal to the full surplus net of the equilibrium payoff of the buyers::

$$\mathcal{R}^*_{\text{fse}} = (1-q^n)\overline{\theta} + q^n\underline{\theta} - nv^*_{\text{fse}}.$$

nv^*_{fse} is equal to:

$$nv^*_{\text{fse}} = \frac{1-\delta}{D(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}),$$

which goes to zero in the limit as δ goes to 1 (recall that $D(\delta)$ converges to a strictly positive number). Thus $\mathcal{R}^*_{\text{fse}} \approx (1-q^n)\overline{\theta} + q^n\underline{\theta}$ for δ close enough to 1, which clearly exceeds $(1-q^n)\overline{\theta}$.

Incentive constraints. Since we have relaxed the low type's incentive compatibility constraint (LowIC), we must now make sure that this constraint is satisfied in the relevant parameter region. Recall that a low-type buyer must be willing to participate in the bidding with the bid $\underline{b}^*$ as opposed to abstaining and getting a zero payoff:

$$(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}^*)+\delta v^*_{\text{fse}}\geq 0.$$

Plugging the payoffs defined in (1.19) and (1.22) into the above constraint, I obtain:

$$-(1-\delta)\frac{q^{n-1}}{n}\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta})$$
$$+\delta\frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})\geq 0,$$

which simplifies to:

$$\delta\leq\frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2},$$

which is true since $\frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2}\geq 1$ by assumption that $q\geq\frac{1-q^n}{n(1-q)}$.

The remaining incentive constraints in $\mathcal{RM}$ are all non-binding at $(\overline{b}^*,\underline{b}^*,v^*_{\text{fse}})$ for high values of δ and all values of q and n, except for the constraint associated with a downward deviation of a high-type buyer (HighIC-down). Recall that a high-type buyer could deviate to $\underline{b}^*+\epsilon$ and win whenever all of his competitors are low types. For this deviation to be unprofitable, his payoff must satisfy:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}}\geq(1-\delta)q^{n-1}(\overline{\theta}-\underline{b}^*).$$

Plugging the payoffs defined in (1.20), (1.21), and (1.22) into the above inequality, I obtain:

$$(1-\delta)\frac{1-q^n}{n(1-q)}\frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta}-\underline{\theta})$$
$$+\delta\frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})$$
$$\geq(1-\delta)q^{n-1}\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}),$$

which simplifies to:

$$\delta(1-q^n)(1-q)\geq q^{n-1}\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big],$$

which can only be satisfied when for δ high enough:

$$(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big],$$

which is the second parameter restriction of Case 2.

No collusion constraints. I check that the no-collusion constraints (No-col-pool) and (No-col-sep-2) are satisfied, or, equivalently, that pooling at $\underline{b}^*$ or bidding $(\underline{b}^*+\epsilon,\underline{b}^*)$ cannot help the buyers to improve their payoff in the buyer-game induced by the seller's equilibrium strategy. Suppose first that the buyers attempt to bid according to $(\underline{b}^*+\epsilon,\underline{b}^*)$, then their payoff will be equal to:

$$\begin{aligned}
v(\underline{b}^*+\epsilon,\underline{b}^*) &= \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\underline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big]}{n(1-\delta q^n)}\\
&= \frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{n(1-\delta q^n)D(\delta)}\times\\
&\times\bigg(\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]-\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\bigg),
\end{aligned}$$

which is exceeded by v^*_{fse} for δ sufficiently close to 1 as long as q and n satisfy the following inequality:

$$(1-q^n)(1-q) > \big(q^n-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big],$$

which is implied by the the second parameter restriction of Case 2 (also given below in (1.25)) since $q^n-(1-q)^n < q^{n-1}\big(1-(1-q)^n\big)$.

If the buyers try to coordinate on pooling at $\underline{b}^*$, their payoff will be:

$$\begin{aligned}
v(\underline{b}^*,\underline{b}^*) =& \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big]\\
&= \frac{1}{n}\bigg[(1-q)\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})\\
&-q\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta})\bigg],
\end{aligned}$$

$v(\underline{b}^*+\epsilon,\underline{b}^*)$ converges to a strictly negative number as δ goes to 1. Indeed, in the limit

$v(\underline{b}^* + \epsilon, \underline{b}^*)$ is given by:

$$\frac{(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\left[q^n\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]-q(1-q^n)(1-q^n)\right],$$

which is strictly negative since $(1-q^n)(1-q^n) > q^{n-1}\big[n(1-q)-(1-q^n)\big]$ (see Appendix A.7.2 for the proof of this claim).

□

Case 2: the restrictions on the parameters

Consider the second parameter restriction of Case 2:

$$(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]. \tag{1.25}$$

The pairs of q and n satisfying this restriction (together with the restriction $q \geq \frac{1-q^n}{n(1-q)}$) are illustrated by Figure 1.5. In the following proposition I establish that the set of q satisfying (1.25) is non-empty for any $n \geq 4$ and that there are values q that do not satisfy (1.25) for every $n \geq 4$.

Proposition 1.9. *The equation*

$$(1-q^n)(1-q) = q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]$$

has a solution on $q \in (0,1)$ for every $n \geq 4$.

Proof. See Appendix A.8.2. □

Observe that the range of q expands as n increases. In the next proposition I establish that any $q \in (0,1)$ will satisfy condition (1.25) for all sufficiently high values of n:

Proposition 1.10. *For all $q \in (0,1)$*

$$\lim_{n\to\infty}\Big((1-q^n)(1-q)-q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]\Big) = 1-q > 0.$$

Proof. See Appendix A.8.3. □

Figure 1.5 also suggests that the restriction in (1.25) can be satisfied for all $q \leq \frac{1}{2}$. Indeed, this claim can be shown formally:

Proposition 1.11. *For all $q \in (0, \frac{1}{2}]$ it is true that*

$$(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big].$$

Proof. See Appendix A.8.4. □

Case 3: Low expected valuation

Recall that in Case 3, both of the incentive compatibility constraints of a high type buyer, i.e. (HighIC-up) and (HighIC-down), bind at the optimum of the revenue maximization problem $\mathcal{RM}$.

Full surplus extraction cPPE, Case 3.

- *Equilibrium conditions:*

$$\begin{aligned}
&\text{(HighIC-up)} && (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}} = (1-\delta)(\overline{\theta}-\overline{b}^*),\\
&\text{(HighIC-down)} && (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}} = (1-\delta)q^{n-1}(\overline{\theta}-\underline{b}^*),\\
&\text{(Eq-payoff)} && v^*_{\text{fse}} = \frac{1}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big].
\end{aligned}$$

- *Parameter restriction*

$$(1-q^n)(1-q) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big].$$

The full solution to the system of equilibrium conditions is provided in Appendix A.5.3. Here I present the equilibrium bids and equilibrium payoffs. I will derive the restriction on the parameters in the course of the proof of Proposition 1.12. Observe that the restriction on the parameters has the following implication:

Lemma 1.6. *For any $q \in (0,1)$ and $n \geq 2$*

$$(1-q^n)(1-q) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \Rightarrow q \geq \frac{1-q^n}{n(1-q)}.$$

Proof. See Appendix A.8.5 □

To write down the expressions for equilibrium bids and payoffs, define $D(\delta)$ as:

$$D(\delta) = (1-q^n)(1-\delta q) + \delta q(1-q) - n(1-\delta)(1-q).$$

A low-type buyer, who wins the auction with the low equilibrium bid $\underline{b}^*$, gets:

$$\underline{\theta} - \underline{b}^* = -\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](\overline{\theta} - \underline{\theta}). \tag{1.26}$$

A high-type buyer, who wins the auction with the high equilibrium bid $\overline{b}^*$, gets:

$$\overline{\theta} - \overline{b}^* = \frac{1}{D(\delta)}\delta q^n(1-q)(\overline{\theta} - \underline{\theta}), \tag{1.27}$$

and a high-type buyer, who wins the auction with the low equilibrium bid $\underline{b}^*$, gets:

$$\overline{\theta} - \underline{b}^* = \frac{1}{D(\delta)}\delta q(1-q)(\overline{\theta} - \underline{\theta}). \tag{1.28}$$

The resulting *ex ante* equilibrium payoff of the buyers is given by:

$$v^*_{\text{fse}} = \frac{1}{nD(\delta)}(1-\delta)q^n\big[n(1-q) - (1-q^n)\big](\overline{\theta} - \underline{\theta}). \tag{1.29}$$

Note that as δ goes to 1, $D(\delta)$ goes to:

$$D(1) = (1-q^n)(1-q) + q(1-q) > 1,$$

hence $D(\delta)$ is strictly positive for δ sufficiently close to 1[7] by continuity of $D(\cdot)$.

The equilibrium bids of each type can be immediately obtained from the respective payoffs

[7]For values of δ satisfying

$$\delta > \frac{n(1-q) - (1-q^n)}{n(1-q) - q^2(1-q^{n-1})}.$$

in (1.26) and (1.27):

$$\underline{b}^* = \underline{\theta} + \frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](\overline{\theta} - \underline{\theta}) \tag{1.30}$$

$$\overline{b}^* = \overline{\theta} - \frac{1}{D(\delta)}\delta q^n(1-q)(\overline{\theta} - \underline{\theta}) \tag{1.31}$$

As in the previous two cases, I first establish the following lemma:

Lemma 1.7. *Suppose δ is sufficiently close to 1, then $\underline{\theta} < \underline{b}^* < \overline{b}^*$.*

Proof. (i) To see that $\underline{\theta} < \underline{b}^*$ for sufficiently high values of δ, observe that:

$$\underline{\theta} - \underline{b}^* \xrightarrow[\delta \to 1]{} -\frac{1}{D(1)}(1-q^n)(1-q)(\overline{\theta} - \underline{\theta}) < 0.$$

(ii) $\underline{b}^* < \overline{b}^*$ is equivalent to $\overline{\theta} - \underline{b}^* > \overline{\theta} - \overline{b}^*$ which is equivalent to:

$$\frac{1}{D(\delta)}\delta q(1-q)(\overline{\theta} - \underline{\theta}) > \frac{1}{D(\delta)}\delta q^n(1-q)(\overline{\theta} - \underline{\theta}),$$

which is clearly true since $D(\delta) > 0$ for δ high enough, and $q > q^n$ for all $n \geq 2$ and $q \in (0,1)$.

□

I now show that the bidding profile in (1.30) and (1.31) can be supported in a *collusive public perfect equilibrium* of the repeated auction game:

Proposition 1.12. *Suppose that $(1-q^n)(1-q) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]$. Suppose further that $\underline{b}^*$ and $\overline{b}^*$ are defined by (1.30) and (1.31) respectively, then there exists a critical discount factor δ^*, such that for all $\delta \in [\delta^*, 1)$ the high-revenue strategy profile corresponding to $(\overline{b}^*, \underline{b}^*)$ (as defined by 1.11) is a collusive public perfect equilibrium of the repeated auction game in the sense of Definition 1.7. Moreover, the seller achieves full surplus extraction in the limit as δ goes to 1.*

Proof sketch. The complete proof is provided in Appendix A.7.3, I briefly sketch the main arguments here. Just as in the previous two cases, by Lemma 1.2 and Lemma 1.7, it is enough to check that $\mathcal{R}^*_{\text{fse}} \geq (1-q^n)\overline{\theta}$ and that the remaining constraints in the revenue maximization problem $\mathcal{RM}$ are satisfied at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$ for high values of δ. I start with the seller's revenue.

Seller's revenue. The seller's revenue is given by:

$$\mathcal{R}^*_{\text{fse}} = (1 - q^n)\overline{\theta} + q^n\underline{\theta} - nv^*_{\text{fse}}.$$

nv^*_{fse} is given by:

$$nv^*_{\text{fse}} = \frac{1}{D(\delta)}(1 - \delta)q^n\big[n(1 - q) - (1 - q^n)\big](\overline{\theta} - \underline{\theta}).$$

Observe that $\lim_{\delta\to 1} nv^*_{\text{fse}} = 0$, which means $\mathcal{R}^*_{\text{fse}} \approx (1 - q^n)\overline{\theta} + q^n\underline{\theta}$ for δ close enough to 1, which clearly exceeds $(1 - q^n)\overline{\theta}$.

Incentive constraints. As in Cases 1 and 2, the on-schedule incentive compatibiliy constraint (HighIC-on-sch) is satisfied. The two off-schedule incentive compatibility constraints (HighIC-up) and (HighIC-down) are satisfied by construction. Hence it remains to check that the low-type incentive compatibility constraint (LowIC) is satisfied. Recall that (LowIC), evaluated at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$, is given by:

$$(1 - \delta)\frac{q^{n-1}}{n}(\underline{\theta} - \underline{b}^*) + \delta v^*_{\text{fse}} \geq 0.$$

Plugging the payoffs from (1.26) and (1.29) in, I get:

$$-(1 - \delta)\frac{q^{n-1}}{n}\frac{1}{D(\delta)}\big[(1 - q^n)(1 - \delta q) - n(1 - \delta)(1 - q)\big](\overline{\theta} - \underline{\theta}) \\ + \delta\frac{1}{nD(\delta)}(1 - \delta)q^n\big[n(1 - q) - (1 - q^n)\big](\overline{\theta} - \underline{\theta}) \geq 0,$$

which is equivalent to:

$$\delta \leq \frac{1}{1 - q} - \frac{1 - q^n}{n(1 - q)^2},$$

which is true whenever $\frac{1}{1-q} - \frac{1-q^n}{n(1-q)^2} \geq 1$ or $q \geq \frac{1-q^n}{n(1-q)}$, which is in turn true in this case by Lemma 1.6.

No-collusion constraints. Recall that in Cases 1 and 2 the no-collusion constraint (No-col-sep-1) was binding and thus the joint deviation to bidding $(\underline{b}^*, \emptyset)$ could not benefit the buyers by construction. Since we have relaxed (No-col-sep-1) here in Case 3, we must now make sure that it is satisfied at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$. Recall that the payoff from bidding $(\underline{b}^*, \emptyset)$ is given by

$$v(\underline{b}^*, \emptyset) = \frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{b}^*)}{n\big(1-\delta(1-q)^n\big)} = \frac{(1-\delta)(1-q^n)\delta q(1-q)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta(1-q)^n\big)}.$$

The equilibrium payoff v^*_{fse} exceeds $v(\underline{b}^*, \emptyset)$ as long as

$$q^{n-1}\big[n(1-q) - (1-q^n)\big] \geq \frac{(1-q^n)\delta(1-q)}{\big(1-\delta(1-q)^n\big)}$$
$$\Leftrightarrow \big(1-\delta(1-q)^n\big)q^{n-1}\big[n(1-q)-(1-q^n)\big] \geq \delta(1-q^n)(1-q),$$

which can be satisfied for any $\delta \in (0,1)$ as long as q and n satisfy

$$\big(1-(1-q)^n\big)q^{n-1}\big[n(1-q)-(1-q^n)\big] \geq (1-q^n)(1-q).$$

which is true by assumption.

Just as in Cases 1 and 2, we must check whether the constraints (No-col-pool) and (No-col-sep-2) are satisfied at $(\overline{b}^*, \underline{b}^*, v^*_{\text{fse}})$, or, equivalently, whether the buyers would lose from pooling at $\underline{b}^*$ or bidding $(\underline{b}^* + \epsilon, \underline{b}^*)$ whenever the state is w^l in the buyer-game. Suppose the buyers coordinate on bidding $(\underline{b}^* + \epsilon, \underline{b}^*)$, then their payoff is:

$$\begin{aligned} v(\underline{b}^* + \epsilon, \underline{b}^*) &= \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\underline{b}^*) + q^n(\underline{\theta}-\underline{b}^*)\big]}{n\big(1-\delta q^n\big)} \\ &= \frac{(1-\delta)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta q^n\big)}\Bigg[(1-q^n)\delta q(1-q) - q^n\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big]\Bigg]. \end{aligned}$$

We must show that $v^*_{\text{fse}} \geq v(\underline{b}^* + \epsilon, \underline{b}^*)$ for sufficiently high values of δ, i.e. that

$$\begin{aligned} &q^n\big[n(1-q) - (1-q^n)\big] \\ &\qquad \geq \frac{1}{\big(1-\delta q^n\big)}\Bigg[(1-q^n)\delta q(1-q) - q^n\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big]\Bigg], \end{aligned}$$

which holds for δ sufficiently close to 1 whenever it holds as a strict inequality at $\delta = 1$, i.e.

whenever

$$
\begin{aligned}
&q^n\big[n(1-q)-(1-q^n)\big] > \frac{1}{\big(1-q^n\big)}\big[(1-q^n)q(1-q)-q^n(1-q^n)(1-q)\big] \\
&\Leftrightarrow q^{n-1}\big[n(1-q)-(1-q^n)\big] > (1-q)(1-q^{n-1}).
\end{aligned}
$$

Now the last line is true since:

$$
(1-q)(1-q^{n-1}) < (1-q)(1-q^n) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big],
$$

where the strict is inequality is obviously true, and the weak inequality holds true in Case 3 by assumption.

If the buyers attempt to coordinate on pooling at $\underline{b}^*$ instead, then their payoff will become:

$$
\begin{aligned}
v(\underline{b}^*,\underline{b}^*) &= \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big] \\
&= \frac{\overline{\theta}-\underline{\theta}}{nD(\delta)}\Big[(1-q)\delta q(1-q)-q\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Big].
\end{aligned}
$$

As in Cases 1 and 2, I show that $v(\underline{b}^*,\underline{b}^*)$ converges to a strictly negative number as δ goes to 1. In the limit the payoff from pooling at $\underline{b}^*$ is:

$$
\frac{\overline{\theta}-\underline{\theta}}{nD(1)}\big[(1-q)q(1-q)-q(1-q^n)(1-q)\big] = \frac{q(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\big[q^n-q\big] < 0.
$$

Since v^*_{fse} is weakly positive, the payoff from pooling at $\underline{b}^*$ cannot exceed the equilibrium payoff in Case 3 for values of δ sufficiently close to 1. □

Case 3: the restriction on the parameters

The range of parameters, where Case 3 applies, equilibrium construction is defined by the following inequality:

$$
(1-q^n)(1-q) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]
$$

The pairs of q and n satisfying the above inequality are illustrated by Figure 1.5. Recall that in Lemma 1.6 we have established that this parameter restriction implies $q \geq \frac{1-q^n}{n(1-q)}$. Recall also that $q \geq \frac{1-q^n}{n(1-q)}$ implies that $n \geq 4$ because it cannot be satisfied for any q as long as $n = 2$ or $n = 3$. Combined with the result of Proposition 1.9, it implies that Case 3 applies to some values of q for all $n \geq 4$, and does not apply to any values of q for $n = 2$ or $n = 3$ (see Figure 1.5 for an illustration).

1.7 Revenue-maximizing reserve prices

The reserve prices along the equilibrium path of the full-surplus-extracting *collusive public perfect equilibria* (in the limit as δ goes to 1) are given by:

$$r^* = \begin{cases} \underline{\theta} + \frac{q(1-q^n)(\overline{\theta}-\underline{\theta})}{q(1-q^n)+q^n\left(1-(1-q)^n\right)} & \text{in Case 1} \\ \underline{\theta} + \frac{\left[(1-q^n)\right]^2(1-q)(\overline{\theta}-\underline{\theta})}{q^n\left(1-(1-q)^n\right)\left[n(1-q)-(1-q^n)\right]+\left[(1-q^n)\right]^2(1-q)} & \text{in Case 2} \\ \underline{\theta} + \frac{(1-q^n)(\overline{\theta}-\underline{\theta})}{1-q^n+q} & \text{in Case 3} \end{cases}$$

They are illustrated by Figure 1.6. The reserve prices in the full-surplus-extracting *cPPE* of the repeated auction game are decreasing in q, going to $\overline{\theta}$ as q goes to 0 and going to $\underline{\theta}$ as q goes to 1. Indeed, since q is the probability of the low type, when q is close to zero, the buyers all have high valuations with a very high probability, and when q is close to 1, the buyers all have low valuations with a very high probability. Recall that the optimal reserve prices in the one-shot auction problem are also decreasing in q, but the optimal decision is essentially a cutoff rule (for fixed values of other parameter values): for relatively low values of q the optimal reserve price is $\overline{\theta}$, while for relatively high values of q it is $\underline{\theta}$. Thus, even though the direction of dependence is the same, the functional form of this dependence is much less trivial in the repeated auction setting with collusion. Similarly, the optimal reserve prices in the one-shot auction problem are increasing in the number of buyers, but the dependence takes the form of a cutoff rule (again, when the other parameter values are fixed), where the optimal reserve price is equal to $\underline{\theta}$ when the number of buyers is relatively low, and is equal to $\overline{\theta}$ when it is relatively high. In contrast to the one-shot setting, the reserve prices in the full-surplus-extracting *cPPE*, even though also increasing in n, depend

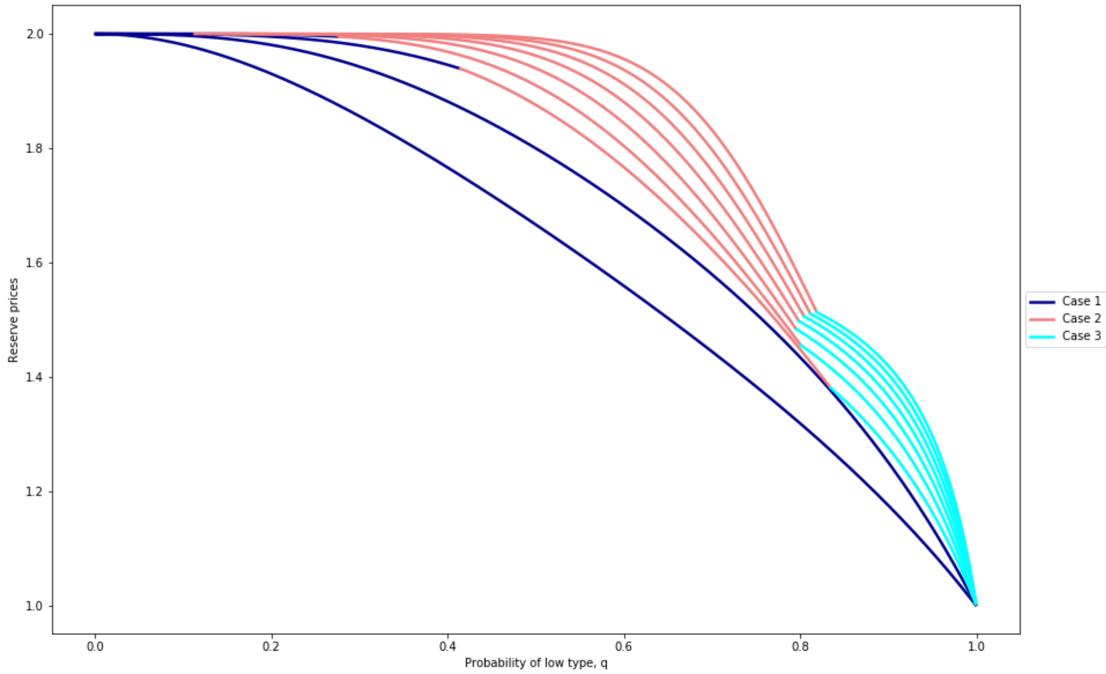

Figure 1.6: *Reserve prices in the full-surplus-extracting cPPE of the repeated auction game in the limit as* δ *goes to* 1. *Valuations are* $\underline{\theta} = 1$ *and* $\overline{\theta} = 2$. *The curves illustrate the limiting reserve prices for all probabilities of the low type* $q \in (0, 1)$, *and for each* $n \in \{2, ..., 10\}$ *moving from the southwest to the northeast as the number of buyers grows, i.e. the southwesternmost curve illustrates the reserve prices for* $n = 2$, *and the northeasternmost curve illustrates the reserve prices for* $n = 10$. *In the dark-blue, red, and light-blue segments, Cases 1, 2, and 3 apply respectively.*

on n in a much less trivial way.

This non-trivial dependence of the reserve prices on q and n can to a certain extent be explained by their very different role in the repeated setting with collusion. In the one-shot auction problem, the role of the reserve prices is to exclude certain valuation types from participation with the purpose of increasing competition among the remaining types. In the repeated setting with colluding buyers, the full-surplus-extracting *cPPE* is efficient and the reserve prices play two crucial roles. First, in the off-path component of the seller's strategy, the reserve prices are chosen to punish the buyers for deviating from the equilibrium path bidding. Second, and more importantly, the on-path component of the reserve prices makes sure that the buyers pay "upfront" for the continuation of favorable terms of trade and at the same time do not have an incentive to collude on a lower bidding profile, resolving the fundamental conflict between revenue-maximization and fighting collusion.

1.8 Concluding remarks

In this paper, I have considered a repeated first-price auction model with a non-committed seller who dynamically adjusts reserve prices to fight collusion among buyers. To model the interaction between the seller and the colluding buyers, I have proposed the solution concept of *collusive public perfect equilibrium*. A collusive public perfect equilibrium is a public perfect equilibrium that additionally requires that the buyers be unable to improve their equilibrium payoff in the "buyer-game" induced by the seller's equilibrium strategy. Studying the outcomes as the discount factor goes to 1, I find a collusive public perfect equilibrium which allows the seller to extract the entire surplus from the colluding buyers. This result suggests that the problem of collusion in repeated auctions is perhaps less severe than is commonly understood: it turns out that a sufficiently sophisticated seller can come up with rather effective strategies for fighting collusion, even when she has to publicly disclose all the bids in the end of every period.

The buyers in this paper are assumed to have access to symmetric collusive schemes. Such collusive schemes are particularly simple and thus might require no explicit communication among the buyers in practice, which makes them virtually impossible to detect for an antitrust authority. These hard-to-detect collusive schemes must therefore be addressed as part of the repeated auction design problem itself. My results imply that it can be done quite successfully. It is however well-known (see e.g. Mailath and Samuelson (2006)) that more sophisticated asymmetric collusive schemes might allow the buyers to collude more effectively, especially when they can communicate before the start of each auction. Even though asymmetric collusive schemes are often dealt with by conventional means of antitrust policy, it is worth studying if they could also be addressed via more sophisticated auction design.

Chapter 2

First Best Implementation with Costly Information Acquisition

with Hien Pham, Takuro Yamashita, and Shuguang Zhu

2.1 Introduction

In most mechanism design problems, there is a collection of agents who have exogenously given private information, and there is a principal who desires to implement a social choice rule, by designing a mechanism which incentivizes the agents to reveal their information.

In many problems in practice, however, the agents' private information is often a consequence of their own (possibly costly) *information acquisition.* For example, bidders in an oil-tract auction (Wilson, 1969) may conduct test drills; bidders in a spectrum auction may conduct market research; voters in a presidential election may investigate the candidates' past political activities; members of a hiring committee may study the job applicant's background in order to see whether he is fit for the job.

Importantly, in such situations, a mechanism in place does not only affect each agent's incentive of reporting the acquired information truthfully, but also affects his choice of *what kind of information* to acquire. In this sense, the properties of desirable mechanisms could potentially be very different from those which only guarantee truth-telling incentives for a given information structure.

Although this issue is already relevant in single-agent environments,[1] the degree of complexity is even higher in a multi-agent environment: in principle, flexibility of each agent's information acquisition action does not only mean flexibility in terms of his signal's informativeness about the payoff-relevant state (e.g., the amount of oil in a tract), but also means flexibility in terms of his signal's informativeness about his opponents' signals. This issue of higher-order information / beliefs distinguishes multi-agent from single-agent environments. Modeling the dependence of the cost of information acquisition on higher-order information is a challenging task. In order to address this issue, we make what we believe to be a reasonable first-step assumption: an agent's information acquisition cost only depends on his signal's informativeness about the payoff-relevant state, but not about the others' signals. Arguably, not all information acquisition environments satisfy this assumption[2], but some clearly do. Consider a situation in which agents (e.g., telecom companies who buy spectrum) have to acquire information from data providers (e.g., market research firms) operating on a competitive market for data. Each data provider generates signals about the payoff-relevant state (e.g., demand conditions in the mobile services market). Competition among the data providers forces them to price their data at the cost of production, which in turn depends on the informativeness of their data. The agents can then decide to make their signals perfectly correlated by strategically choosing the same data provider, or less than perfectly correlated by choosing different data providers. In both cases, the agents will pay the same price for the same informativeness, and hence the cost of information acquisition will be independent of the correlation structure among signals.

We consider a model with four or more agents, and assume that it is costless for each agent to acquire a signal that is independent from the state of the world.[3] The principal and agents share a common prior, and none of them has any private information at the beginning. We show that there exists a mechanism which allows the principal to implement any social choice rule at zero information acquisition cost to the agents. The mechanism recommends each agent to choose a special information acquisition action, which satisfies the *individually-uninformative-but-aggregately-revealing* property of Zhu (2021) (and each

[1] Mensch (2022) studies a mechanism design problem with a single agent. See also Section 2.1.1.

[2] We discuss a range of possible alternative assumptions in Section 2.5.1.

[3] The literature on cost of information proposes and discusses a variety of possible cost functions (see Section 2.1.1), but it seems to be universally accepted that uninformative signals about the state of the world are costless.

agent finds it optimal to obey this recommended action). The *individually-uninformative* part means that each agent's signal on its own is independent from the payoff-relevant state, which guarantees that his information cost is zero. The *aggregately-revealing* part means that the principal, by observing all the agents' reports — in fact, any two of them — can correctly identify the true payoff-relevant state. The fact that only two are enough, together with the fact that there are four or more agents, enables the principal to detect any unilateral deviation. It thus establishes the incentive compatibility of the mechanism.

2.1.1 Related Literature

In the literature of information acquisition in mechanism design, we usually consider restricted and/or less flexible spaces of information (see, for example, Bergemann and Välimäki (2002) for efficient mechanism design, Shi (2012) and Bikhchandani and Obara (2017) for optimal auction design, and Persico (2004), Gerardi and Yariv (2008), Gershkov and Szentes (2009), and Zhao (2016) for committee design with information acquisition.[4])

Mensch (2022) studies mechanism design with a single agent's flexible and costly information acquisition, building on the rational inattention framework (Sims (2003))[5]. Flexible and costly information acquisition is also considered by Gleyze and Pernoud (2020) who study a mechanism design problem with transferable utility and private values, in which agents acquire costly information on their own preferences and the preferences of other agents, and by Ravid et al. (2020) who study a bilateral trade model with costly information acquisition by the buyer. Flexible information but not costly information acquisition is considered by Roesler and Szentes (2017) in the context of buyer-optimal information in monopoly pricing,[6] by Bergemann et al. (2017) and Brooks and Du (2021) in the context of seller-pessimal information in common-value auction, and by Yamashita (2018) in private-value auction. All these papers have a common feature that there is a single entity, "nature", who chooses the information structure (of one or multiple agents). In contrast, in our model,

[4]Restricting to the class of conservative rules, Li (2001) solves for the optimal degree of conservatism in committee design. The optimally chosen conservative rule outperforms the *ex post* optimal rule.

[5]Mensch (2022) also considers a multiple-agent extension of his model, but restricts attention to symmetric mechanisms in an independent private values setting, in which agents can acquire information about their own values, but cannot acquire any information about others' values.

[6]See also Condorelli and Szentes (2020), though they also consider non-information changes of the agent's private information distribution.

each of the agents acquires information in a decentralized manner, leading to a very different conclusion.

The information structure we employ was proposed in the context of mechanism design by Zhu (2021), who studies information disclosure by a mechanism designer. It builds on the idea of the one-time pad, an unbreakable encryption method (Shannon, 1949).[7]

This key information structure makes the agents' acquired information statistically dependent. In quasi-linear environments, Crémer and McLean (1988) show that the principal can extract full surplus from the agents. Although the extreme positivity of the results is a common feature of our paper and theirs, the two problems are quite different. First, our paper does not assume quasi-linearity. Second, their side-bet mechanism exploits an exogenously given correlated signal structure, and it is not clear if such a signal structure can be induced in equilibrium given some reasonable space of information acquisition actions.[8] In our case, the above mentioned information structure is indeed an equilibrium outcome, even though each agent can potentially acquire information independent from the others' signals.

In non-quasi-linear environments, such as collective decision-making in committees, the first best outcome is generally not implementable under the commonly imposed restrictions on information acquisition technologies. For example, Li (2001) and Persico (2004), assuming that the agents have access to conditionally independent signals, show that the first best outcome is not attainable. In contrast to the previous results, we show that correlated information acquisition helps to implement the first best outcome.

There is a growing literature on the cost of flexible information in decision environments (see for example Sims (2003), Matejka and McKay (2015), Caplin and Dean (2015), and Pomatto et al. (2020)). Usually the main focus is on the cost of acquiring more or less precise information about a payoff-relevant state, and its relationship with a single decision-maker's optimal choice. The framework, however, has been applied in multi-player problems (i.e., in games), such as in coordination games (Yang (2015); Morris and Yang (2021);

[7]See also Krähmer (2020) and Krähmer (2021) in the context of information disclosure in mechanism design and strategic communication respectively where the randomization of information structures is allowed to keep the single agent (sender) uninformative; Kalai et al. (2010), Renou and Tomala (2012), Renault et al. (2014) in the context of games of communication network. Peters and Troncoso-Valverde (2013) apply this idea in mechanism-design games with multiple principals, and Liu (2015) applies it in his concept of individually uninformative correlating device. Our construction is most directly related to Zhu (2021).

[8]Bikhchandani (2010) shows that, indeed, an agent in the Crémer-McLean mechanism may have a strong incentive of acquiring information about others.

Denti (2020)). In particular, Denti (2020) proposes a model of unrestricted information acquisition in games, in which, as in our paper, the players can endogenously learn about a payoff-relevant state and actions of other players.

2.2 Model

2.2.1 Setup

There is a principal and $I \geq 4$ agents, and a finite set of payoff-relevant states Θ. Each agent i's payoff is denoted $u_i(d, \theta)$, when a social decision $d \in D$ is selected in state θ.[9] For example, in an auction, d is a vector of bidders' winning probabilities and their expected payments, and each u_i is quasi-linear in the payment part. Later, each agent's payoff *net* his information acquisition cost is considered as his objective.

At the beginning, neither the principal nor any of the agents know θ. Instead, we assume that there is a set of unobservable fundamental states of nature $X = [0, 1]$ with a typical element x, equipped with a Borel σ-algebra and a uniform probability measure $\mathbb{P}$[10]. Taking a richer space of the fundamental states of nature would not change our results. There is also a commonly known measurable function $\boldsymbol{\Theta} : X \to \Theta$ mapping the fundamental states of nature to the payoff-relevant states. This function induces a common prior on the payoff-relevant states as follows: $\mu_0(\theta) \equiv \int_0^1 1_{\{\boldsymbol{\Theta}(x)=\theta\}} dx =$ for each $\theta \in \Theta$.

The agents can acquire *costly* information about θ by generating private and possibly correlated signals, whereas the principal cannot acquire any information about θ. Each agent has access to a large set of possible signal realizations S_i. In principle, S_i (in particular, its size) may be a part of i's choice, but assuming exogenous S_i is without loss of generality as long as $|S_i| \geq |\Theta|$. Agent i's information acquisition action is a measurable function $\sigma_i : [0, 1] \to S_i$, such that, once x (and hence $\theta = \boldsymbol{\Theta}(x)$) is realized, then i observes $s_i = \sigma_i(x)$. Let Σ_i denote the set of all such measurable functions, defining i's information acquisition action space. We assume that information acquisition is fully private in the sense that neither the principal nor any other agent observes which information acquisition action i

[9] We can endow the principal with his own payoff function $u_0(d, \theta)$, though it is not necessary.

[10] See Gentzkow and Kamenica (2017) who adapt this approach in the context of multi-sender Bayesian persuasion.

takes and which signal realization is observed by agent i. Agent i's objective is the *net* payoff $u_i(d, \theta) - c_i(\sigma_i)$. We assume the information acquisition cost function of agent i has the following properties:

Assumption 2.1. Properties of information acquisition cost.

1. $c_i(\sigma_i) \geq 0$ *for any* $\sigma_i \in \Sigma_i$.

2. $c_i(\sigma_i) = 0$ *if* σ_i *and* $\boldsymbol{\Theta}$ *are stochastically independent.*

It is a quite general cost function specification, which includes many well-known cost functions as special cases: entropy-based cost functions, more general posterior-separable cost functions, and even some posterior-non-separable ones.[11] Observe that the second property makes sure that agent i pays nothing as long as he learns nothing about the payoff-relevant state θ even if he learns something about the fundamental state x. We believe that it is a reasonable first-step assumption but have to admit that our main result hinges on it.

Any profile of information acquisition actions σ induces a joint distribution over payoff-relevant states and signal realizations, we denote this distribution by $\alpha \in \Delta(\Theta \times S)$. When we want to make its dependence on σ more explicit, we write α_σ.

2.2.2 Mechanism

The principal faces both hidden action and hidden information of each agent. The principal commits to a mechanism at the *ex ante* stage in order to control the agents' incentives. More specifically, following the literature, we let the principal (i) send a message privately to each agent before his information acquisition action, and (ii) collect a message privately from each agent after the agent has observed a signal realization. Formally, a mechanism comprises $(R, \rho; M, \delta)$ where $R = (R_i)_{i=1}^I$ and $M = (M_i)_{i=1}^I$; R_i denotes the set of messages that the principal can send to each agent i; M_i denotes the set of messages that each agent i can send to the principal; $\rho \in \Delta(R)$ is a distribution over the principal's messages, and $\delta : R \times M \to D$ denotes the decision rule.

The timing of the game is summarized as follows:

[11] See the literature of cost of information, such as Sims (2003), Matejka and McKay (2015), Caplin and Dean (2015), and Pomatto et al. (2020).

$t = 0$: $x \sim U(0, 1)$ is drawn but no one observes it.

$t = 1$: The principal designs a mechanism $(R, \rho; M, \delta)$.

$t = 2$: After observing the mechanism and receiving $r_i \in R_i$, each agent i privately chooses his information acquisition action $\sigma_i \in \Sigma_i$.

$t = 3$: Each agent i privately observes $s_i = \sigma_i(x)$, and privately sends $m_i \in M_i$ to the principal.

$t = 4$: The principal executes $d = \delta(r, m)$ where $m = (m_i)_{i=1}^I$.

Because no agent observes the other agents' actions or information (even noisily) at all, we consider Nash equilibrium as a solution concept. Then, applying the revelation principle of Forges (1986), we focus on *direct* mechanisms where (i) the principal directly recommends an information-acquisition action to each agent, and each agent directly reports a signal to the principal, and (ii) each agent finds it optimal to obey the recommended action and truthfully report his signal.[12]

Formally, a direct mechanism comprises $((\sigma_i)_{i=1}^I, (S_i)_{i=1}^I, \delta)$, where the principal recommends $\sigma_i \in \Sigma_i$ privately to each agent i,[13] and executes $\delta(s) \in D$ if the agents report $s = (s_i)_{i=1}^I \in S = \times_{i=1}^I S_i$. A direct mechanism is *incentive compatible* if it satisfies the following constraints: for any $\sigma_i' \in \Sigma_i$ and $\tau_i : S_i \to S_i$,

$$\sum_{\theta, s_i, s_{-i}} (u_i(\delta(s_i, s_{-i}), \theta)\alpha_{\sigma_i, \sigma_{-i}}(\theta, s_i, s_{-i})) - c_i(\sigma_i)$$
$$\geq \sum_{\theta, s_i, s_{-i}} (u_i(\delta(\tau_i(s_i), s_{-i}), \theta)\alpha_{\sigma_i', \sigma_{-i}}(\theta, s_i, s_{-i})) - c_i(\sigma_i').$$

That is, each i must find it optimal to obey the recommended σ_i and report the realized s_i truthfully.

[12]The proof proceeds as follows. First, imagine an auxiliary game where there is no principal, but instead, there is a fictitious player ("player 0") who is indifferent across all decisions in any state. At first, each agent i plays σ_i privately, and then observes the realized signal s_i privately. Then, (without any communication), player 0 chooses $d \in D$. Interpreting this as a baseline extensive-form game, it is easy to see that our current game (with the principal) is the mediated communication game of this auxiliary game in the sense of Forges (1986) (see also Myerson (1986)). Thus, her revelation principle applies.

[13]We focus on a deterministic recommendation of σ, rather than any stochastic recommendation. Accordingly, δ is denoted simply by $\delta(s)$ instead of $\delta(r, s)$. Since first best implementation is achieved with pure recommendations, our focus on them is without loss of generality.

Although the constraints are concisely summarized by the inequalities above, they are actually rather complicated. First, changing σ_i affects the joint distribution α of (θ, s) and the agent's cost in a non-trivial way since agent i cannot affect agent $-i$'s information structure. Second, an agent may potentially want to make a double deviation, that is, change σ_i and at the same time change his reporting strategy.

Remark 2.1. *Here, we do not explicitly impose individual rationality / participation constraints. It is not difficult to accommodate these constraints: let us require that any feasible direct mechanism must have an extra message $m_i^{\emptyset}$ (a "non-participation" message) so that i's message space is now $S_i \cup \{m_i^{\emptyset}\}$, and $\delta(m_i^{\emptyset}, m_{-i})$ is some specific allocation (a "non-participation allocation") for agent i, for any given m_{-i}. When the non-participation message is included into the set of messages for each agent, the individual rationality constraints, both at the ex ante and ex interim stages, are captured by the above incentive compatibility constraints.*[14]

2.3 Main result

Fix any function $d^* : \Theta \to D$, which describes all the economically relevant outcomes in this environment except for the information acquisition costs. If the principal could observe θ, then any d^* is attainable without any information acquisition cost on the agents' side. In this sense, one may interpret this d^* together with zero cost for the agents as the *first-best* outcome.[15]

In this section, for any given d^*, we explicitly construct a mechanism that implements d^* at zero cost for the agents. That is, the first best outcome can be attained *even though the principal cannot directly observe θ.*

Theorem 2.1. *Fix any $d^* : \Theta \to D$. Under Assumption 2.1, there exists a mechanism (σ, S, δ) such that (i) $\sum_s \alpha(\theta, s) 1_{\{\delta(s)=d^*(\theta)\}} = \mu_0(\theta)$ for all θ, and (ii) $c_i(\sigma_i) = 0$ for all i.*

Proof. The theorem is proved by construction.

[14] Ex interim individual rationality is guaranteed because agent i can always deviate to $\tau_i(\cdot) \equiv m_i^{\emptyset}$. Ex ante individual rationality is guaranteed because agent i can always deviate to a costless σ_i and then to $\tau_i(\cdot) \equiv m_i^{\emptyset}$.

[15] For example, one may assume that $d^*(\theta)$ is the best decision of the principal given his own preferences in state θ.

Since Θ is a finite set, we assume without loss of generality that $\Theta = \{1, \dots, T\}$. Let $K > \max\{I, T\}$ be a prime number. Because $\mathbb{P}$ is a uniform measure on $X = [0, 1]$, we can find a partition of X, denoted by $\{X_{\theta\psi}\}_{(\theta,\psi)\in\{1,\dots,T\}\times\{1,\dots,K\}}$, satisfying $\int_0^1 1_{\{x\in X_{\theta\psi}\}}dx = \frac{1}{K}\mu_0(\theta)$ for any θ and ψ. Define a measurable function $\mathbf{\Psi} : [0, 1] \to \{1, \dots, K\}$ such that, if $x \in \cup_{\theta\in\Theta} X_{\theta\psi}$, then $\mathbf{\Psi}(x) = \psi$. Immediately, $\mathbf{\Psi}$ is uniformly distributed on $\{1, \dots, K\}$ conditional on any realization θ of $\mathbf{\Theta}$, hence $\mathbf{\Psi}$ is independent of $\mathbf{\Theta}$.

Now consider the following information acquisition action profile: for each $i \in \{1, \dots I\}$, $S_i = \{1, \dots, K\}$, and $\sigma_i(x) = \mathbf{\Theta}(x) + i \cdot \mathbf{\Psi}(x) \mod K$ for any $x \in [0, 1]$. Note that the residual is calculated as in standard modular arithmetic except when $\mathbf{\Theta}(x) + i \cdot \mathbf{\Psi}(x)$ is divisible by K, in which case we set $\sigma_i(x) = K$ instead of 0. The following lemma gives the properties of (S, σ) that we need to prove the theorem.

Lemma 2.1. *The above (S, σ) satisfies:*

(i) For any $i \in \{1, \dots, I\}$, σ_i is independent of $\mathbf{\Theta}$.

(ii) Conditional on any realization of (s_i, s_j) such that $i \neq j$, the joint distribution of $\mathbf{\Theta}$ and $(\sigma_k)_{k\neq i,j}$ is degenerate.

Proof of the lemma. By definition of (S, σ), for each i, we have $s_i = \theta + i \cdot \psi \mod K$, where the random variables $\mathbf{\Psi}$ and $\mathbf{\Theta}$ are independent. Thus the signal profile $s = (s_i)_{i=1}^I$ is defined in the same way as in Zhu (2021).[16] Thus, this lemma is directly implied by Lemma 2 in Zhu (2021). □

The first property says that θ, s_i are independent, implying $c_i(\sigma_i) = 0$. The second property says that, given s_i, s_j with $i \neq j$, we can identify the true payoff-relevant state θ and any signal realization s_k without error, that is, there exist $\hat{\theta}(s_i, s_j)$ and $\hat{s}_k(s_i, s_j)$ such that:

$$\Pr(\theta = \hat{\theta}(s_i, s_j)|s_i, s_j) = \Pr(s_k = \hat{s}_k(s_i, s_j)|s_i, s_j) = 1.$$

Let the principal recommend the above σ, and offer the decision rule δ as follows: $\delta(s) =$

[16]In fact, our signal profile s coincides with what Zhu (2021) calls "the IUAR disclosure policy, where IUAR is short for *individually uninformative but aggregately revealing*.

$d^*(\theta)$ if (i) for any i, j with $i \neq j$, we have

$$\theta = \hat{\theta}(s_i, s_j);$$

or if (ii) there is i such that, for any j, k where i, j, k are all different, we have

$$\theta = \hat{\theta}(s_j, s_k).$$

In any other case, $\delta(s)$ is arbitrary.

Clearly, if the agents obey the recommendation and report their signals truthfully, then the first best outcome is attained. Therefore, we complete the proof by showing that the proposed mechanism satisfies incentive compatibility. Take any agent i, and suppose that he deviates to any σ_i' and reports $\tau_i(s_i)$ when s_i is realized. First, his cost of information acquisition increases weakly. Second, his reporting decision does not affect the social decision at all, because the principal executes $\delta(s) = d^*(\hat{\theta}(s_j, s_k))$ for an arbitrary pair (j, k) which does not include i. Therefore, the mechanism is incentive compatible. □

2.4 Applications

2.4.1 Full-surplus extraction in common value auctions

Consider the following common value auction environment. The seller (principal) has a single indivisible good, and there are $I \geq 4$ bidders. The value of the good is common to all the bidders, denoted by $\theta \in \Theta$, where Θ is finite. In fact, the analysis of this section can be straightforwardly extended to the case of "non-pure" common values where each i's valuation is $v_i(\theta)$. Let $\mu_0(\theta)$ denote the probability that θ is the bidders' common value.

Each bidder i's payoff is $\theta q_i - t_i - c_i(\sigma_i)$ if he wins the good with probability q_i, pays t_i to the seller, and spends $c_i(\sigma_i)$ as his information acquisition cost. In case he does not participate in the mechanism, his outside-option payoff is 0. The seller's payoff is revenue, $\sum_{i=1}^{I} t_i$.

The first-best expected surplus of this society is the expected common value:

$$\sum_{\theta \in \Theta} \mu_0(\theta)\theta = \mathbb{E}[\theta].$$

There are several cases where the seller can easily earn $\mathbb{E}[\theta]$. First, if the seller *knows* θ, then he can simply post price θ. Even if the seller does not know θ, if the bidders know θ as their common knowledge (i.e., as *free* information), then again the seller can earn $\mathbb{E}[\theta]$. Conversely, if all the bidders are completely *uninformed* (so that each only knows the common prior μ_0), then again, the seller can post price $\mathbb{E}[\theta]$.

Notice that, with costly information acquisition as considered in our paper, neither of the above ideas would work. First, although it might be possible to make every bidder fully learn θ in some equilibrium, it does not yield $\mathbb{E}[\theta]$ as long as full information is strictly costly. Second, if the seller posts price $\mathbb{E}[\theta]$, then each bidder has a strong incentive of knowing whether the true θ is below $\mathbb{E}[\theta]$ or not: If i finds that $\mathbb{E}[\theta|s_i] < \mathbb{E}[\theta]$ given some signal s_i, he would not buy the good. As long as such information is not too costly, the bidder would be better off by acquiring it.

Therefore, with a general information acquisition cost function, the equilibrium information should be somewhere between full and no information, and it is *a priori* unclear how the seller should find the optimal balance of information and rent extraction. Nevertheless, as long as the cost functions satisfy Assumption 2.1, Theorem 2.1 implies that the full-surplus extraction is possible.

Corollary 2.1. *Under Assumption 2.1, there is a mechanism which yields $\mathbb{E}[\theta]$ as the seller's expected revenue (and each bidder earns 0).*

It is worth emphasizing that the logic here is very different from that of Crémer and McLean (1988). In their paper, the seller exploits an *exogenously given* correlated signal structure, in order to construct a side-bet scheme that extracts the entire surplus. In our case, each bidder can choose any information structure. Indeed, if he prefers, a bidder can choose an information structure such that his information is *independent* from all the other bidders' signals. The Crémer-McLean lottery scheme, therefore, does not work here. Also, in their mechanism, each bidder's payoff can be strictly negative *ex post*, while in our case, it is zero *ex post*. Indeed, if a mechanism offers a negative ex post payoff in our environment, a bidder would have a strong incentive of getting a signal indicative of that event and then report differently, in order to avoid such a negative payoff.

2.4.2 First-best implementation in collective decision-making

Consider a committee with a designer (principal) and $I \geq 4$ members (agents) deciding whether to hire or not to hire a job market candidate. Formally, $d \in D = \{h, nh\}$. The quality of the candidate is $\theta \in \Theta$, which is unobserved *ex ante*. The designer and all members of the committee hold a common prior belief $\mu_0 \in \Delta(\Theta)$ about the candidate's quality.

The utility that each member obtains from hiring / not hiring the candidate is defined as follows:

$$u_i(d, \theta) = \begin{cases} u_i(\theta), & \text{if } d = h \\ 0, & \text{if } d = nh \end{cases}$$

Without loss of generality, we assume that $u_i(\theta) = k_i\theta$.

Only the committee members can acquire information about the candidate at cost $c_i(\sigma_i)$. The designer aims to maximize the expected sum of all members' gross utilities.[17] That is, ideally, he wants to hire the candidate if and only if $\sum_i k_i\theta \geq 0$. The first best expected surplus of all committee members is given by:

$$\sum_{\theta | \sum_i k_i\theta \geq 0} \mu_0(\theta) \sum_i k_i\theta \equiv W^{FB}$$

It is useful to note that the existing literature (see e.g. Li (2001) and Gerardi and Yariv (2008)) typically assumes that the committee members have access to information structures whose realized signals are independently distributed across them, conditional on the state of the world. Under these restrictions, the first best outcome cannot be implemented. There are two main forces that prevent the committee from implementing the first best with these restricted information structures: free-riding problem and conflict of interest. First, when committee members have a conflict of interest, they may prefer not to report their own acquired information truthfully. Second, even if all members share a common preference, information could be underprovided relative to the social optimum, because it is essentially

[17] The result extends to the case where the designer maximizes expected sum of members' net utilities (taking into account the information acquisition costs).

a public good used to make a collective decision. For example, Li (2001) suggests that distorting the decision rule away from the *ex post* optimal rule (which is optimal under exogenous information) could help to alleviate the free-riding issue.

In contrast to the previous literature, our results show that with more flexible (even though still costly) information acquisition the designer can implement the first best outcome. On the one hand, having access to a wider range of information acquisition technologies enlarges the set of feasible deviations for the agents. On the other hand, the principal now has more flexibility in designing information structures recommended to the agents. Given these two opposing effects, it is not immediately clear *a priori* whether the first best outcome becomes more or less difficult to attain. It turns out that the second effect dominates: with a larger set of feasible mechanisms, the principal is able to incentivize the agents to acquire and report their information truthfully, no matter what social choice rule the principal is trying to implement. Indeed, Theorem 2.1 implies that the first best is implementable as long as Assumption 2.1 about the cost functions holds.

Corollary 2.2. *Under Assumption 2.1, there is a mechanism which yields W^{FB} as the total expected surplus (the decision is made under full information with no cost).*

Our construction helps to resolve both of the issues that prevent first best implementation with conditionally independent signals. Recall that it costs nothing for an agent to acquire an *"individually uninformative"* signal which is assigned to him under the optimal mechanism. Therefore, the distorted provision of information is resolved. Moreover, even if committee members have a conflict of interest, under our mechanism, they cannot do better than being truthful since any unilateral deviation can be detected by the designer.

2.5 Concluding remarks

It is quite natural that agents may desire to refine their information in response to the mechanism. This paper proposes one possible framework, based on a class of information acquisition cost functions, such that the cost of information depends on the informativeness of each agent's signal about the state of the world, but on its informativeness about other agents' signals. We show that such a specification leads to an extremely positive result. We conclude the paper with two remarks.

2.5.1 Interdependent cost functions

One natural criticism may be that our mechanism induces full information if the signals are aggregated, even though any single signal is completely uninformative: would it be reasonable to assume that such σ_i is costless? Because the answer is necessarily yes under Assumption 2.1, the question is essentially whether Assumption 2.1 itself is reasonable. Assumption 2.1 is satisfied, in particular, in any information acquisition environment, in which the cost of information acquisition does not depend on the correlation structure among signals but only depends on their individual informational content. We argued in Section 2.1 that there are information acquisition environments, for which this assumption is indeed a reasonable one. In general, however, the correlation structure might affect the cost of information acquisition in various ways. On the one hand, there seem to be cases where more positive correlation is more expensive. For example, fix agent 1's private information, and consider agent 2. If acquiring a positively correlated information necessarily means that agent 2 must *steal* (perhaps a part of) agent 1's information, more positive correlation will be more costly. On the other hand, there are opposite situations. Strulovici (2021), for example, considers the environment where hard evidence is scarce in the sense that, if one agent "picks up" a piece of evidence, then it becomes difficult for the others to get the same or similar evidence. On top of that, there also seem to be cases where *less correlated* signals are more costly: for example, imagine that there exist 2 newspapers, and σ_i corresponds to the decision of which newspapers to buy (including the options of buying both or none). Assume that the information in those newspapers is highly positively correlated, and assume that agent 1 buys (only) newspaper 1 for sure. In this case, agent 2 receives a positively correlated signal regardless of the newspaper he chooses, while in order to obtain something less correlated, he must buy *both* newspapers and orthogonalize newspaper 2's information from newspaper 1's information. Thus, less correlation is more costly.

This discussion suggests that we must think more about modeling the microstructure of information acquisition, in order to determine which correlation structures are more costly. Mechanism design with such more specific information acquisition cost structures would certainly be an interesting future direction, and we hope this article could serve as a first step towards that direction.

2.5.2 Two or three agents

Obviously, the individually-uninformative-but-aggregately-revealing signal structures require four or more agents to implement the first best at zero cost for the agents. With three agents, although it is possible to determine whether *some* agent has unilaterally deviated or not, it is not possible to identify who the deviator is (and hence not possible to identify the true θ). If there exists a social decision $d \in D$ that can serve as a severe punishment for all agents for any given θ, then the same first best implementation result obtains.

With two agents, each agent has much more freedom. With two agents, even under Assumption 2.1, an extremely positive result similar to the one obtained in this paper does not hold. The optimal mechanism might involve some costly information acquisition, and hence, the specification of the cost functions matters even more.

Chapter 3

Bilateral Trade with Costly Information Acquisition

with Takuro Yamashita

3.1 Introduction

Traders dealing with complex objects often do not have enough relevant information to correctly estimate the object's value at the outset, and therefore may take potentially costly actions to acquire more information. Consider for example a landowner (a seller) who owns a plot of land which is known to likely have a commercially viable amount of oil under its surface. Suppose this landowner is not in the oil business and is thus considering selling the mineral rights to an oil company (a buyer). At the beginning neither party has a precise estimate of the amount of oil under the surface, but each party could order exploratory drilling to obtain better estimates. The outcomes of the two exploratory drilling studies could be more or less correlated depending on how much coordination between studies the landowner and the oil company achieve. It is possible that the parties decide to order a single study together, in which case the outcomes will be perfectly correlated, or two independent studies in different locations, in which case the outcomes can conceivably be independent conditional on the amount of oil under the surface.

We are interested in the problem of a third party who intermediates trade between the seller and the buyer, and is interested in maximizing her own revenue. The parties communicate with each other via the intermediary who determines the communication protocol and the resulting allocation and payments. The possibility of information acquisition by the parties presents a considerable challenge for the intermediary: in our example the landowner and the oil company may hide some aspects of their exploratory studies from the intermediary and each other, and thus will have to be incentivized to disclose what studies have been performed and what results these studies produced.

In order to better understand the problem of such an intermediary, we build a model with two players: a buyer and a seller who can trade an indivisible object, and a mechanism designer who intermediates trade between them. The object's quality (*payoff-relevant state*) determines its value for the players. We assume that conditional on knowing the true quality the players would always like to trade. Our mechanism designer is only interested in money and values the object at zero irrespective of quality.

At the outset, neither the players nor the mechanism designer have any information about the object's quality beyond a commonly known prior. In the beginning, the mechanism designer commits to a *mechanism* which consists of messages to be sent by the players later and the allocation and payment functions defined on the messages. Once the mechanism designer has selected a mechanism, the players simultaneously generate *signals* to acquire more information about the quality of the object. To model the information acquisition process, we assume that there is a probability space of fundamental states of nature and that every random variable in the model is a measurable mapping form the sample space of fundamental states to another measurable space (e.g. the object's quality is a random variable that maps the fundamental states to the space of possible qualities). Before the game starts, nature draws a fundamental state but nobody observes it. A player's *signal* is a pair consisting of a finite space of possible *signal realizations* and a random variable that maps the fundamental states to the signal realizations. The signals generated by the players are *costly*. The cost of a signal is proportional to the expected reduction of entropy achieved by the player generating the signal (i.e. as in rational inattention, see Sims (2003) and Matejka and McKay (2015)). Information acquisition is thus costly but *flexible*, allowing for arbitrary correlations across signals and the object's quality. It is also *hidden* as neither player observes the signal chosen by the other player and the intermediary does not observe the signal chosen by either player.

After the players have chosen their signals, they *privately* observe the *signal realization* corresponding to the fundamental state chosen by nature. Having observed their signal realization, they select a message to report to the mechanism designer who then announces the allocations and payments. The quality of the object corresponding to the fundamental state is then revealed and the player's payoffs are determined. The players are interested in maximizing their payoffs net of information acquisition costs.

We consider *Nash equilibria* in pure strategies[1] of the resulting mechanisms and, in the case of multiplicity, select an equilibrium that maximizes the mechanism designer's revenue. One might wonder whether choosing an equilibrium concept that takes into account the dynamic nature of our environment (e.g. perfect Bayesian equilibrium) would change our results, but, fortunately, it is not the case. Intuitively, since players only observe their own signal realizations, they obtain no information on the signal chosen by the other player. Hence, an off-equilibrium information set can only be achieved following a player's own deviation, which makes sure that every Nash equilibrium has an outcome-equivalent perfect Bayesian equilibrium (Proposition 3.2 proven in Appendix B.4 formalizes this argument).

We establish a *revelation principle* (Proposition 3.1), which allows us to restrict attention to truthful-revelation equilibria of *direct mechanisms*. Direct mechanisms ask the players to report one of the signal realizations from the support of their equilibrium signal. Signals chosen by the players induce a joint distribution over object's qualities and signal realizations (*an information structure*) whose marginal on the set of qualities is equal to the prior. Moreover, we notice that any such information structure can be induced by a pair of signals (see Lemmas 3.1 and 3.2, which directly follow form Theorem 1 in Yang (2020)). This equivalence of signals and information structures allows us to state our implementability conditions in terms of information structures, which considerably simplifies the problem.

Simplifying the problem further, we show that we can consider a restricted class of deviations for each player without loss of generality (Lemma 3.4). Intuitively, if along the equilibrium path of a truthful-revelation Nash equilibrium a player has n signal realizations, then he only has $n + 1$ available reporting deviations: he can choose to report a different signal realization or abstain from participation altogether. With only $n+1$ available actions

[1] Whether the restriction to pure strategies is without loss of generality or not is an open problem that appears to be non-trivial. Appendix B.1 discusses this issue in more detail.

the player would never want to deviate to an alternative signal with more than $n + 1$ signal realizations since additional information would be learned but essentially wasted otherwise. Thus, when we want to check if a particular information structure with a given number of signal realizations can be induced in a truthful-revelation Nash equilibrium of a direct mechanism, it is enough to consider information structures which have one additional signal realization for each player.

Considering this restricted class of deviations, we derive implementability conditions. In order to show that a given information structure paired with truthful reporting is implementable, we divide the restricted class of deviations into two subclasses. In the first subclass, we consider deviation-induced information structures that preserve the set of signal realizations for each player. To show that the deviations in the first subclass are unprofitable, we solve a finite-dimensional payoff-maximization problem for each player, where the maximum is taken over information structures which preserve the set of signal realizations for that player. In the second subclass, we consider deviation-induced information augmented with an additional signal realization. We explicitly solve for the best deviation in this class (Lemma 3.6) and derive an unprofitability condition. Combining the unprofitability conditions from the two subclasses, we obtain our implementability conditions (Proposition 3.3).

Once the number of signal realizations for each player is fixed, the implementability conditions are given by a finite-dimensional system of equations and inequalities. One can therefore directly use them to find an optimal mechanism with a given number of signal realizations. We, however, take a different approach and investigate whether full surplus extraction can be achieved by any mechanism. To achieve full surplus extraction, the mechanism designer would have to use a mechanism that is *efficient* (otherwise trade will generate less than full surplus in the first place) and individually uninformative (otherwise a chunk of generated surplus will have to be spent on paying for the players' information acquisition). We show that individually uninformative efficient mechanisms cannot achieve full surplus extraction (Proposition 3.4).

3.1.1 Related literature

Our paper is most closely related to Larionov et al. (2022) who study a general implementation problem in the same information acquisition environment but with more than two players.

They show that any social choice function can be implemented in any transferable environment with three players or more. We show that (at least as long as one restricts attention to pure strategies equilibria) having three or more players is not only sufficient but also necessary for unrestricted implementation.

The literature on information acquisition in mechanism design goes back to Bergemann and Välimäki (2002) who study efficient implementation in transferable environments with exogenously restricted information acquisition. They show that the VCG mechanism achieves both *ex ante* and *ex post* efficiency if agents have private values, but not necessarily when they have common values. Bikhchandani (2010) points out that the full surplus extraction mechanism of Crémer and McLean (1988) may not be robust to information acquisition because agents presented with a Crémer-McLean lottery may have incentives to acquire additional information about their competitors' valuations. Bikhchandani and Obara (2017) study a mechanism design problem, in which (similarly to our paper) agents can acquire costly signals about a payoff-relevant state of nature. The space of signals available to each agent is, however, exogenously restricted. Bikhchandani and Obara (2017) provide conditions under which full surplus extraction is possible in their setting.

More recently, some consideration has been given to flexible information acquisition. In a paper closely related to ours, Mensch (2022) solves for a revenue-maximizing auction among buyers who, like the players in our paper, can acquire costly and hidden information about the value of the object sold in the auction. The cost of information acquisition is assumed to belong to the posterior-separable class, which contains, among others, the entropy cost we employ in our paper. Unlike in our paper, however, the agents in Mensch (2022) have private values and are exogenously restricted to acquire information about their own values. Terstiege and Wasser (2022) solve for a revenue-maximizing auction with private values and flexible information acquisition but assume that information acquistion is *costless* and *public*. In their environment, the bidders' choice of signals is publicly revealed before the bidders privately observe a signal realization. Having observed the signals chosen by the bidders, the seller proposes a mechanism to maximize her revenue. Gleyze and Pernoud (2020) study a mechanism design problem with costly flexible information acquisition, transferable utility, and private values but allow the agents to acquire information both on their own preferences and the preferences of the other agents. Gleyze and Pernoud (2020) are interested in *informationally simple* mechanisms, i.e. those in which the participating agents have no incentive to acquire information about anyone's preferences but their own.

In our paper, the players choose an information structure to maximize their own payoffs, hence their choice may not necessarily be desirable from the mechanism designer's perspective. This feature of our model makes our paper somewhat close to the literature on "adverse" choice of information structures. Yamashita (2018) studies a private-value auction, in which for any mechanism proposed by the seller, nature chooses an information structure that minimizes the seller's revenue. Bergemann et al. (2017) and Brooks and Du (2021) study analogous models of common-value auctions. Roesler and Szentes (2017) study a bilateral trade model, in which the buyer can acquire costless information about the good's value and the seller best-responds by setting a revenue-maximizing take-it-or-leave-it price. Ravid et al. (2020) consider the same setting as Roesler and Szentes (2017) but make the buyer's information acquisition costly.

We model information acquisition by giving the players access to a large space of signals which partition an underlying set of fundamental states of nature. This way of modeling signals is introduced by Green and Stokey (1978) and is also used by Gentzkow and Kamenica (2017) in the context of Bayesian persuasion with multiple senders.

3.1.2 Roadmap

The rest of the paper is organized as follows. Section 3.2 introduces the model, Section 3.3 derives the global implementability conditions. In Section 3.4 we show that full surplus extraction is impossible. Finally, Section 3.5 concludes. All proofs are relegated to the Appendix.

3.2 Model

3.2.1 Setup

An indivisible good, whose quality v is drawn from a finite set of payoff-relevant states of the world V, can be traded between two players: a seller and a buyer. The buyer's valuation for the good of quality $v \in V$ is given by $u^b(v)$, the seller's valuation for the good of quality

$v \in V$ is given by $u^s(v)$. We assume that gains from trade always exist, i.e $u^b(v) > u^s(v)$ for any v. To model information acquisition, we assume that there is a set of fundamental states of the world $x \in X = [0,1]$ with an associated Borel σ-algebra $\mathcal{F}$ and the uniform measure $\mathbb{P}$, and a random variable $\mathbf{V} : X \to V$. At the beginning of the game, this structure is commonly known. Observe that $\mathbf{V}$ induces a common prior μ_0 on the set of qualities such that the probability of quality being equal to v is given by $\mu_0(v) \equiv \int_0^1 \mathbf{1}_{\{\mathbf{V}(x)=v\}} dx$. We assume that μ_0 has full support on V.

The players can acquire costly information about the good's quality by generating *signals*. We assume that each player p has access to a countably infinite set of possible *signal realizations*. Since the labels of signal realizations do not have any particular meaning in our setup, we assume that the set of signal realizations is the set of all natural numbers $\mathbb{N} \equiv \{1, 2, 3, \dots\}$. We use $\mathscr{P}(\mathbb{N})$ to denote the collection of all *finite non-empty* subsets of $\mathbb{N}$. A *signal* is a pair $\sigma^p = (S^p, \mathbf{S}^p)$, where $S^p \in \mathscr{P}(\mathbb{N})$ and $\mathbf{S}^p$ is a random variable that maps fundamental states of nature to signal realizations in S^p, i.e. $\mathbf{S}^p : X \to S^p$. If the fundamental state is x, then player p observes the signal realization $s^p = \mathbf{S}^p(x)$. We use Σ^p to denote the set of all signals for player p. Signals are costly, the cost of a signal σ^p, denoted by $C(\sigma^p)$, is proportional to the reduction of entropy for the player who generates that signal. We introduce the cost function formally below. The players are interested in maximizing their utilities net of information acquisition costs.

There is a mechanism designer who intermediates trade between the seller and the buyer. The designer commits to a *mechanism* at the *ex ante* stage. A mechanism is a tuple (M, q, t), where $M = M^b \times M^s$ with M^p being a finite set of messages available to player p. q is a tuple of allocation functions (q^b, q^s), where $q^p : M \to [0,1]$ determines the allocation for player p. t is a tuple of payment functions (t^b, t^s), where $t^b : M \to \mathbb{R}$ is a payment made by the buyer to the mechanism designer, and $t^s : M \to \mathbb{R}$ is a payment made by the mechanism designer to the seller. The mechanism designer is interested in maximizing her revenue.

To summarize, the timing of our game is as follows:

1. Nature draws $x \in X$ uniformly, but nobody observes it.
2. The mechanism designer commits to a mechanism (M, q, t), which is publicly observed.
3. Each player p privately chooses $\sigma^p = (S^p, \mathbf{S}^p)$.
4. Each player p privately observes $s^p = \mathbf{S}^p(x)$ and privately sends $m^p \in M^p$ to the

designer.

5. Allocations and transfers are determined according to (q, t).

The buyer gets $q^b(m)u^b(v) - t^b(m) - C(\sigma^b)$, the seller gets $t^s(m) - q^s(m)u^s(v) - C(\sigma^s)$, and the mechanism designer gets $t^b(m) - t^s(m)$, where $v = \mathbf{V}(x)$ and $m = (m^b, m^s)$.

We choose Nash equilibrium in pure strategies as our solution concept but this choice has little affect on our analysis as we show in Proposition 3.2 that any Nash equilibrium has an outcome-equivalent perfect Bayesian equilibrium (whether the restriction to pure strategies leads to loss of generality or not is an open problem; we elaborate on this issue in Appendix B.1). If there are multiple Nash equilibria, we select one that maximizes the mechanism designer's revenue. Naturally, we assume that the mechanism designer is restricted to choose mechanisms that satisfy physical feasibility, i.e. $0 \leq q^b(m) \leq q^s(m) \leq 1$ for any $m \in M$, and allow for voluntary participation both *ex ante* and *ex interim*. i.e. we assume that there exists a message $m_\emptyset \in M^p$ for any player p, such that $q^p(m_\emptyset, m^{-p}) = t^p(m_\emptyset, m^{-p}) = 0$ for any message $m^{-p} \in M^{-p}$ sent by by the other player.

3.2.2 Information structures

Each signal σ^p chosen by player p induces a joint distribution on $S^p \times V$. We use α^p to denote this joint distribution and write $\alpha^p(s^p; v)$ for the probability of player p observing the signal realization s^p and the state of the world being v. When we want to emphasize the dependence of α^p on σ^p, we write $\alpha^p[\sigma^p]$. Likewise a pair of signals (σ^b, σ^s) induces a joint distribution on $S^b \times S^s \times V$. We use α to denote this joint distribution, and use $\alpha(s^b, s^s; v)$ to denote the joint probability of the buyer observing the signal realization s^b, the seller observing the signal realization s^s, and state of the world being v. When we want to emphasize the dependence of α on (σ^b, σ^s), we write $\alpha[\sigma^b, \sigma^s]$. Clearly, we have $\text{marg}_{S^p \times V}\alpha = \alpha^p$ for any player p. In what follows, we refer to α as *information structure*. Following Kamenica and Gentzkow (2011), we introduce the following definition:

Definition 3.1. *Information structure α is Bayes-plausible, if* $\text{marg}_V \alpha = \mu_0$.

Clearly, any information structure induced by a pair of information acquisition actions must be Bayes-plausible. The following lemma shows that the converse is also true:

Lemma 3.1. *For every Bayes-plausible information structure α there exists a profile of signals that induces α.*

Now suppose that player p deviates to an alternative signal $\tilde{\sigma}^p$. Which alternative information structures $\tilde{\alpha}$ can this deviation induce? Clearly, we must have $\text{marg}_{S^{-p}\times V}\ \tilde{\alpha} = \text{marg}_{S^{-p}\times V}\ \alpha$, i.e. a deviation by player p cannot change the joint distribution of player $-p$'s signal realizations and states of the world. The following lemma shows that the converse is also true:

Lemma 3.2. *Fix a signal profile $\left(\sigma^p, \sigma^{-p}\right)$ and the associated information structure α. Consider any joint distribution $\tilde{\alpha}$ on $\tilde{S}^p \times S^{-p} \times V$ such that $\text{marg}_{\tilde{S}^{-p}\times V}\ \tilde{\alpha} = \text{marg}_{S^{-p}\times V}\ \alpha$. There exists $\tilde{\sigma}^p \in \Sigma^p$ such that $\left(\tilde{\sigma}^p, \sigma^{-p}\right)$ induces $\tilde{\alpha}$.*

Lemmas 3.1 and 3.2 both follow immediately from Theorem 1 in Yang (2020), hence their proof is omitted. Lemmas 3.1 and 3.2 will allow us to rewrite the mechanism designer's problem in terms of information structures, and thus avoid having to explicitly model players' signal choices. We return to this issue below in Subsection 3.2.4 after we discuss the cost of information acquisition.

3.2.3 Cost of information acquisition

Consider a signal σ^p chosen by player p and the distribution $\alpha^p[\sigma^p]$ induced by that signal. Having chosen σ^p, player p observes signal realization s^p with probability $\tau[\sigma^p](s^p) \equiv \sum_{v\in V} \alpha[\sigma^p](s^p; v)$. If $\tau[\sigma^p](s^p) > 0$, then the signal realization s^p induces a posterior distribution over states of the world $\mu^p[\sigma^p](s^p)$. The posterior probability of state v is given by:

$$\mu^p[\sigma^p](s^p; v) \equiv \frac{\alpha^p[\sigma^p](s^p; v).}{\tau[\sigma^p](s^p)}.$$

The cost of information acquisition action σ^p is proportional to the expected reduction of entropy achieved by $\alpha^p[\sigma^p]$:

$$C(\sigma^p) = c(\alpha^p[\sigma^p]) \equiv \chi\Bigg(H(\mu_0) - \sum_{s^p|\tau[\sigma^p](s^p)>0} \tau[\sigma^p](s^p) H\big(\mu^p[\sigma^p](s^p)\big)\Bigg),$$

where $H(\mu) = -\sum_{v\in V} \mu(v) \log\big(\mu(v)\big)$ with the standard convention $0 \log 0 = 0$.

In what follows, we normalize χ to 1. Whenever we consider a pair of information acquisition choices $\left(\sigma^p, \sigma^{-p}\right)$ inducing a joint distribution α, we find it convenient to work with a cost function defined directly on information structures as follows: $c^p(\alpha) \equiv c(\text{marg}_{S^p \times V}\alpha)$. Clearly, we have $C(\sigma^p) = c^p(\alpha[\sigma^p, \sigma^{-p}])$ for each player p. The following lemma will be helpful later in our analysis:

Lemma 3.3. *$c^p(\alpha)$ is convex in α for any player $p \in \{b, s\}$.*

Proof. See Appendix B.2. $\square$

3.2.4 Strategies, equilibria, and direct mechanisms

Consider an arbitrary mechanism $\left(M, q, t\right)$. A strategy is a tuple $\left(\sigma^p, \left\{\mathbf{m}^p[\hat{\sigma}^p]\right\}_{\hat{\sigma}^p \in \Sigma^p}\right)$, where σ^p is a signal chosen by player p on path, and $\left\{\mathbf{m}^p[\hat{\sigma}^p]\right\}_{\hat{\sigma}^p \in \Sigma^p}$ is a family of reporting functions, one for each $\hat{\sigma}^p \in \Sigma^p$, mapping signal realizations from $\hat{\sigma}^p$ to the mechanism's messages, i.e. $\mathbf{m}^p[\hat{\sigma}^p] : \hat{S}^p \to M^p$ for each $\hat{\sigma}^p \in \Sigma^p$. As player p chooses σ^p at the information acquisition stage, his on-path reports are given by $\mathbf{m}^p[\sigma^p]$, i.e. when player p observes a signal realization $s^p \in S^p$, then he sends the message $\mathbf{m}^p[\sigma^p](s^p)$. In what follows, we omit the dependence of the on-path reports on σ^p and simply write $\mathbf{m}^p(s^p)$ for the message sent by player p who has observed s^p.

We focus our attention on *direct mechanisms*. In a direct mechanism, each player p takes an information acquisition action $\sigma^p = (\mathbf{S}^p, S^p)$, and the mechanism designer asks the players to report their signal realizations or send an abstention message, thus $M^p = S^p \cup \{m_\emptyset\}$ for each player p. Let α be the information structure induced by a pair of information acquisition choices $\left(\sigma^b, \sigma^s\right)$ and let $\mathbf{m}_T^p$ be the truthful reporting function for player p, i.e. for any $s^p \in S^p$ we have $\mathbf{m}_T^p(s^p) = s^p$, and consider truthful-revelation Nash equilibria in direct mechanisms. Using Lemmas 3.1 and 3.2, we can write the truthful-revelation Nash equilibrium conditions in a direct mechanism in terms of the information structure $\alpha \in \Delta\left(S^b \times S^s \times V\right)$ and the truthful reporting function $\mathbf{m}_T^p$. α and $\mathbf{m}_T^p$ can arise in a truthful-revelation Nash equilibrium of a direct mechanism if and only if

- They are *ex ante* incentive compatible for the buyer[2]:

$$(\text{IC}^b_{\text{A}}) \quad (S^b, \alpha, \mathbf{m}^b_T) \in \underset{\tilde{\alpha}, \tilde{S}^b, \tilde{\mathbf{m}}^b}{\operatorname{argmax}} \sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big) - c^b(\tilde{\alpha}),$$

$$\text{s.t.} \quad (1) \;\; \tilde{S}^b \in \mathscr{P}(\mathbb{N}), \;\; \tilde{\alpha} \in \Delta\big(\tilde{S}^b \times S^s \times V\big), \;\; \tilde{\mathbf{m}}^b : \tilde{S}^b \to S^b \cup \{m_\emptyset\};$$

$$(2) \;\; \operatorname{marg}_{S^s \times V} \tilde{\alpha} = \operatorname{marg}_{S^s \times V} \alpha.$$

- They are *ex ante* incentive compatible for the seller[2]:

$$(\text{IC}^s_{\text{A}}) \quad (S^s, \alpha, \mathbf{m}^s_T) \in \underset{\tilde{\alpha}, \tilde{S}^s, \tilde{\mathbf{m}}^s}{\operatorname{argmax}} \sum_{s^b \in S^b} \sum_{s^s \in \tilde{S}^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(t^s(s^b, \tilde{\mathbf{m}}^s(s^s)) - q^s(s^b, \tilde{\mathbf{m}}^s(s^s))u^b(v)\big) - c^s(\tilde{\alpha}),$$

$$\text{s.t.} \quad (1) \;\; \tilde{S}^s \in \mathscr{P}(\mathbb{N}), \;\; \tilde{\alpha} \in \Delta\big(S^b \times \tilde{S}^s \times V\big), \;\; \tilde{\mathbf{m}}^s : \tilde{S}^s \to S^s \cup \{m_\emptyset\};$$

$$(2) \;\; \operatorname{marg}_{S^b \times V} \tilde{\alpha} = \operatorname{marg}_{S^b \times V} \alpha.$$

- α is Bayes-plausible:

$$(\text{BP}) \qquad \sum_{s^b \in S^b} \sum_{s^s \in S^s} \alpha(s^b, s^s; v) = \mu_0(v).$$

We now show that our focus on direct mechanisms is without loss of generality by establishing a revelation principle for Nash equilibria:

Proposition 3.1. *For any Nash equilibrium of an indirect mechanism there exists an outcome-equivalent truthful-revelation Nash equilibrium in a direct mechanism.*

Proof. See Appendix B.3. □

One could argue that we should have chosen perfect Bayesian equilibrium as our solution concept since our environment has dynamic structure. The following proposition shows, however, that the two equilibrium concepts are outcome equivalent in our setting.

[2]Observe that this formulation of *ex ante* incentive compatibility takes care of *ex ante* individual rationality as well. Consider e.g. a deviation for the buyer $(\tilde{S}^b, \tilde{\alpha}, \tilde{\mathbf{m}}^b)$, where $\tilde{S}^b = \{1\}$, $\tilde{\alpha}(1, s^s; v) = \sum_{s^b \in S^b} \alpha(s^b, s^s; v)$, and $\tilde{\mathbf{m}}^b(1) = m_\emptyset$. Observe that $\operatorname{marg}_{S^s \times V} \tilde{\alpha} = \operatorname{marg}_{S^s \times V} \alpha$ The payoff from this deviation is $\sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(1, s^s; v)0 - c^b(\tilde{\alpha}) = -c^b(\tilde{\alpha})$. Bayes-plausibility implies $\alpha^b(1, v) = \sum_{s^s \in S^s} \tilde{\alpha}(1, s^s; v) = \mu_0(v)$, hence $c^b(\tilde{\alpha}) = H(\mu_0) - H(\mu_0) = 0$. By *ex ante* incentive compatibility we have:

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\alpha) \geq 0.$$

Proposition 3.2. *Every truthful-revelation Nash equilibrium of a direct mechanism has an outcome-equivalent perfect Bayesian equilibrium in this direct mechanism.*

Proof. See Appendix B.4. □

The proof of Proposition 3.2 is somewhat tedious but the intuition is straightforward. Since players' signals are chosen simultaneously, player p has no information about the signal chosen by player $-p$. Moreover, the information that player p gets at the signal realization stage does not reveal any information about the signal chosen by player $-p$ either, hence a player can achieve an off-equilbrium information set only by deviating to a different signal himself. Then, if truthful reporting along the equilibrium path is not sequentially rational , it will not be optimal from the *ex ante* perspective either. In other words, if a player suddenly finds it profitable to misreport after observing a signal realization s^p, then he must have contingently planned to misreport following s^p from the start, but then of course truthful revelation cannot be a Nash equilibrium in the first place.

3.2.5 Revenue maximization problem

The mechanism designer maximizes her revenue subject to α-feasibility, which makes sure that the mechanism designer chooses Bayes-plausible information structures over Cartesian products of finite nonempty subsets of $\mathbb{N}^2$ and V; q-feasibility, which makes sure that the mechanism satisfies physical feasibility; and *ex ante* incentive compatibility for both players:

$$\max_{\alpha, S^b, S^s; q, t} \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(t^b(s^b, s^s) - t^b(s^b, s^s)\big), \text{ s.t.}$$

$$(\alpha\text{-F}) \quad S^b, S^s \in \mathscr{P}\big(\mathbb{N}\big), \;\; \alpha \in \Delta\big(S^b \times S^s \times V\big), \;\text{BP};$$

$$(q\text{-F}) \quad 0 \le q^b(s^b, s^s) \le q^s(s^b, s^s) \le 1 \quad \forall (s^b, s^s) \in S^b \times S^s;$$

$$(\text{IC}_\text{A}) \quad \text{IC}^b_\text{A}, \; \text{IC}^s_\text{A}.$$

3.3 Implementability

The *ex ante* incentive compatibility constraints in the designer's revenue maximization problem (Subsection 3.2.5) are rather complicated. They prevent players from deviating to a possibly different information structure and misreporting their signal realizations at the same time. The class of such deviations is extremely large. In this section, we show that it is without loss of generality to consider a much smaller class of *ex ante* deviations.

3.3.1 Restricted *ex ante* deviations

We first show that it is without loss of generality to restrict attention to those *ex ante* deviations, in which a player augments his information structure with an additional signal realization $s^p_\emptyset$ and chooses a new joint distribution on the augmented signal realization space. The player abstains from participation after observing $s^p_\emptyset$ and reports truthfully otherwise. This idea is captured by *restricted ex ante incentive compatibility constraints.*

- The restricted *ex ante* incentive compatibility constraint for the buyer is given by:

$$\text{(R-IC}^b_\text{A}) \quad (S^b, \alpha) \in \operatorname*{argmax}_{\tilde{\alpha}, \tilde{S}^b} \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\tilde{\alpha}),$$

$$\text{s.t.} \quad (1)\ \tilde{S}^b = S^b \cup \{s^b_\emptyset\},\ \ \tilde{\alpha} \in \Delta\big(\tilde{S}^b \times S^s \times V\big);$$

$$(2)\ \operatorname{marg}_{S^s \times V} \tilde{\alpha} = \operatorname{marg}_{S^s \times V} \alpha.$$

- The restricted *ex ante* incentive compatibility constraint for the seller is given by:

$$\text{(R-IC}^s_\text{A}) \quad (S^s, \alpha) \in \operatorname*{argmax}_{\tilde{\alpha}, \tilde{S}^s} \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(t^s(s^b, s^s) - q^s(s^b, s^s)u^b(v)\big) - c^s(\tilde{\alpha}),$$

$$\text{s.t.} \quad (1)\ \tilde{S}^s = S^s \cup \{s^s_\emptyset\},\ \ \tilde{\alpha} \in \Delta\big(S^b \times \tilde{S}^s \times V\big);$$

$$(2)\ \operatorname{marg}_{S^b \times V} \tilde{\alpha} = \operatorname{marg}_{S^b \times V} \alpha.$$

The following lemma establishes formally establishes that restricted *ex ante* incentive compatibility implies *ex ante* incentive compatibility.

Lemma 3.4. R-IC^p_A $\Rightarrow$ IC^p_A *for both players* $p \in \{b, s\}$.

Proof. See Appendix B.5. □

The argument at the core of Lemma 3.4's proof is straightforward. In any direct mechanism, a player, whose on-path signal has n possible signal realizations, can choose between $n+1$ possible actions: this player can report one of the signal realizations (possibly misreporting) or abstain from participation altogether. Suppose this player has a profitable unrestricted *ex ante* deviation, i.e. there is a pair consisting of an information structure and a reporting function that gives this player a strictly larger expected payoff. If the information structure in this unrestricted deviation has more than $n+1$ signal realizations, than at least two signal realizations will lead to the same action. If we scramble all signal realizations leading to the same action, we will obtain an information structure with a one-to-one mapping between signal realizations and actions. Since the labels of signal realizations do not have any specific meaning, we can always relabel them to ensure that those that do not lead to an abstention are reported truthfully. In that way, we can construct a *restricted ex ante* deviation whose information structure is less informative in the Blackwell sense and is therefore less costly. Since the rest of the payoff is exactly the same, the restricted *ex ante* deviation is even more profitable than the unrestricted one.

Suppose now that the mechanism designer hopes that a particular information structure $\alpha \in \Delta\big(S^b \times S^s \times V\big)$ will be induced in the truthful-revelation Nash equilibrium of his direct mechanism $\big((S^b \cup \{m_\emptyset\}) \times (S^s \cup \{m_\emptyset\}), q, t\big)$. If $|S^b| = k$ and $|S^s| = n$, then α is essentially a collection of $k \times n$ matrices, one for each state (we adopt the convention that the buyer is a *row player* and the seller is a *column player*):

State v	s^s_1	s^s_2	$\dots$	s^s_n
s^b_1	$\alpha_{11}(v)$	$\alpha_{12}(v)$	$\dots$	$\alpha_{1n}(v)$
s^b_2	$\alpha_{21}(v)$	$\alpha_{22}(v)$	$\dots$	$\alpha_{2n}(v)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s^b_k	$\alpha_{k1}(v)$	$\alpha_{k2}(v)$	$\dots$	$\alpha_{kn}(v)$

where $\alpha_{ij}(v)$ is the joint probability of the buyer observing the signal realization s^b_i and the seller observing the signal realization s^s_j in state v. We also denote $\mu^b_i(v)$ the posterior probability of state v as evaluated by the buyer who receives the signal realization s^b_i, and $\mu^s_j(v)$ the posterior probability of state v as evaluated by the seller who receives the signal realization s^s_j.

Lemma 3.4 shows that to make sure that his desired α will indeed be induced, the mechanism designer should only check whether (R-IC$^s_{\text{A}}$) and (R-IC$^s_{\text{A}}$) constraints are satisfied, i.e. should only check deviations that augment α by not more than one signal realization. To analyze these deviations, we find it useful to split them into two classes. The first class of deviations consists of possibly different joint distributions over the same signal realizations while the deviations in the second class augment the set of the signal realizations by exactly one more realization. The usefulness of this approach will become clear by the end of this section. Let us deal with the first class of the restricted *ex ante* deviations first.

Class 1 of restricted *ex ante* deviations

- The deviations in the first class are unprofitable for the buyer as long as α satisfies:

$$(\text{R-IC}^b_{\text{A}}\text{-1}) \qquad \alpha \in \underset{\tilde{\alpha}}{\operatorname{argmax}} \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V} \tilde{\alpha}_{ij}(v)\big(q^b_{ij}u^b(v) - t^b_{ij}\big) - c^b(\tilde{\alpha}), \quad \text{s.t.}$$
$$(1)\ \ \tilde{\alpha} \in \Delta\big(S^b \times S^s \times V\big);$$
$$(2)\ \ \text{marg}_{S^s\times V}\tilde{\alpha} = \text{marg}_{S^s\times V}\alpha.$$

- The deviations in the first class are unprofitable for the seller as long as α satisfies:

$$(\text{R-IC}^s_{\text{A}}\text{-1}) \qquad \alpha \in \underset{\tilde{\alpha}}{\operatorname{argmax}} \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V} \tilde{\alpha}_{ij}(v)\big(t^s_{ij} - q^s_{ij}u^b(v)\big) - c^s(\tilde{\alpha}), \quad \text{s.t.}$$
$$(1)\ \ \tilde{\alpha} \in \Delta\big(S^b \times S^s \times V\big);$$
$$(2)\ \ \text{marg}_{S^b\times V}\tilde{\alpha} = \text{marg}_{S^b\times V}\alpha.$$

Notice that (R-IC$^b_{\text{A}}$-1) and (R-IC$^s_{\text{A}}$-1) are finite dimensional maximization problems with concave objectives and affine constraints. Moreover, observe that a Bayes-plausible information structure that allows a player to put a posterior probability of zero on any of the states of the world can never be a solution to the maximization problems in (R-IC$^b_{\text{A}}$-1) and (R-IC$^s_{\text{A}}$-1). This is due to the properties of the entropy cost function, which makes sure that marginal costs of information acquisition go to infinity as soon as any of the posteriors approaches zero. The following lemma establishes the claim formally:

Lemma 3.5. *Suppose α is Bayes-plausible, and satisfies* (R-IC$^p_{\text{A}}$-1) *for both players $p \in \{b, s\}$, then for any $v \in V$ and for any $s^b_i \in S^b$ we have $\mu^b_i(v) > 0$, likewise for any $v \in V$ and for any $s^s_j \in S^s$ we have $\mu^s_j(v) > 0$.*

Proof. See Appendix B.6. ∎

Lemma 3.5 makes sure that the objective functions in (R-IC$^b_{\mathrm{A}}$-1) and (R-IC$^s_{\mathrm{A}}$-1) are differentiable at any optimum, hence all deviations in the first class are unprofitable if and only if α satisfies the Karush-Kuhn-Tucker optimality conditions in both problems.

Class 2 of restricted *ex ante* deviations

Recall that the deviations in the second class involve augmentation of the set of the signal realizations by exactly one signal realization.

- The deviations in the second class are unprofitable for the buyer as long as α satisfies:

$$\text{(R-IC}^b_{\mathrm{A}}\text{-2)} \quad \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\alpha_{ij}(v)\big(q^b_{ij}u^b(v)-t^b_{ij}\big)-c^b(\tilde\alpha)\geq \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\tilde\alpha_{ij}(v)\big(q^b_{ij}u^b(v)-t^b_{ij}\big)-c^b(\tilde\alpha),$$

for all $\tilde\alpha$ such that:

(1) $\tilde\alpha\in\Delta\big(\tilde S^b\times S^s\times V\big)$, where $\tilde S^b=S^b\cup\{s^b_\emptyset\}$ and $\exists s^s_j\in S^s, v\in V$ s.t. $\alpha_{\emptyset,j}(v)>0$;

(2) $\text{marg}_{S^s\times V}\tilde\alpha=\text{marg}_{S^s\times V}\alpha$.

- The deviations in the second class are unprofitable for the seller as long as α satisfies:

$$\text{(R-IC}^s_{\mathrm{A}}\text{-2)} \quad \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\alpha_{ij}(v)\big(t^s_{ij}-q^s_{ij}u^b(v)\big)-c^s(\tilde\alpha)\geq \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\tilde\alpha_{ij}(v)\big(t^s_{ij}-q^s_{ij}u^b(v)\big)-c^s(\tilde\alpha)$$

for all $\tilde\alpha$ such that:

(1) $\tilde\alpha\in\Delta\big(S^b\times\tilde S^s\times V\big)$, where $\tilde S^s=S^s\cup\{s^s_\emptyset\}$ and $\exists s^b_i\in S^b, v\in V$ s.t. $\alpha_{i,\emptyset}(v)>0$;

(2) $\text{marg}_{S^b\times V}\tilde\alpha=\text{marg}_{S^b\times V}\alpha$.

The usefulness of splitting the deviations into these two classes is illustrated by the next lemma, which shows that if α satisfies the constraints (R-IC$^b_{\mathrm{A}}$-1) and (R-IC$^s_{\mathrm{A}}$-1), then the constraints (R-IC$^b_{\mathrm{A}}$-2) and (R-IC$^b_{\mathrm{A}}$-2) can be considerably simplified.

Lemma 3.6. *Suppose α satisfies* R-IC$^p_{\mathrm{A}}$-1 *for both players $p\in\{b,s\}$, then*

- *α satisfies* R-IC$^b_{\mathrm{A}}$-2 *if and only if* $\sum_{v\in V}\exp\big(-y^b(v)\big)\leq 1$, *where*

$$y^b(v)\equiv\min_{(i,j)|\alpha_{ij}(v)>0}\Big\{q^b_{ij}u^b(v)-t^b_{ij}-\log\big(\mu^b_i(v)\big)\Big\};$$

- α *satisfies* R-IC^s_A-2 *if and only if* $\sum_{v\in V} \exp\big(-y^s(v)\big) \leq 1$, *where*

$$y^s(v) \equiv \min_{(i,j)|\alpha_{ij}(v)>0} \Big\{ t^s_{ij} - q^s_{ij} u^s(v) - \log\big(\mu^s_j(v)\big) \Big\}.$$

Proof. See Appendix B.7. □

The full proof of Lemma 3.6 is relegated to the Appendix, but we illustrate the main ideas of the proof below using a simple example with two states and two signal realizations for each player. Suppose there are indeed two payoff-relevant states of the world, i.e. $V = \{\underline{v}, \overline{v}\}$, and we would like to find out whether the following information structure (denoted α) satisfies the buyer's constraint (R-IC^b_A-2) for a given mechanism assuming that it satisfies the constraint (R-IC^b_A-1) for the same mechanism.

State $\underline{v}$	s^s_1	s^s_2
s^b_1	$\underline{\alpha}_{11}$	$\underline{\alpha}_{12}$
s^b_2	$\underline{\alpha}_{21}$	$\underline{\alpha}_{22}$

State $\overline{v}$	s^s_1	s^s_2
s^b_1	$\overline{\alpha}_{11}$	$\overline{\alpha}_{12}$
s^b_2	$\overline{\alpha}_{21}$	$\overline{\alpha}_{22}$

Let us suppose that it does not actually satisfy the constraint (R-IC^b_A-2), then we must be able to find a profitable deviation, which induces a different information structure, which transfers some probability mass from the existing signal realizations to $s^b_\emptyset$, after which the buyer abstains. For some $\epsilon > 0$ we can write down the information structure induced by this deviation as follows:

State $\underline{v}$	s^s_1	s^s_2
s^b_1	$\underline{\alpha}_{11} - \epsilon\underline{\beta}_{11}$	$\underline{\alpha}_{12} - \epsilon\underline{\beta}_{12}$
s^b_2	$\underline{\alpha}_{21} - \epsilon\underline{\beta}_{21}$	$\underline{\alpha}_{22} - \epsilon\underline{\beta}_{22}$
$s^b_\emptyset$	$\epsilon\sum_{i=1}^2 \underline{\beta}_{i1}$	$\epsilon\sum_{i=1}^2 \underline{\beta}_{i2}$

State $\overline{v}$	s^s_1	s^s_2
s^b_1	$\overline{\alpha}_{11} - \epsilon\overline{\beta}_{11}$	$\overline{\alpha}_{12} - \epsilon\overline{\beta}_{12}$
s^b_2	$\overline{\alpha}_{21} - \epsilon\overline{\beta}_{21}$	$\overline{\alpha}_{22} - \epsilon\overline{\beta}_{22}$
$s^b_\emptyset$	$\epsilon\sum_{i=1}^2 \overline{\beta}_{i1}$	$\epsilon\sum_{i=1}^2 \overline{\beta}_{i2}$

We denote the gain from this deviation $G_\alpha(\epsilon\beta)$. By assumption, $G_\alpha(\epsilon\beta) > 0$ for some $\epsilon > 0$. First of all, we notice that the payoff function of the buyer is concave and hence for any global profitable deviation there is a local deviation with a marginal gain $MG_\alpha(\beta) \equiv \lim_{\epsilon\to 0} \frac{1}{\epsilon} G_\alpha(\epsilon\beta)$ that is also strictly positive. Moreover, once we consider local deviations, we can without loss of generality take all β_{ij}'s to be weakly positive. The last claim is true because, locally, any direction of improvement of α can be represented as a linear combination of a direction that is feasible in (R-IC^b_A-1) and another direction, in which all β_{ij}'s are weakly positive. Since α solves (R-IC$^b_{\mathrm{A}}$-1) by assumption, any improvement must come from the second component of this linear combination.

Calculating the marginal gain $MG_\alpha(\beta)$, we obtain:

$$\begin{aligned} MG_\alpha(\beta) = &- \sum_{i=1}^{2}\sum_{j=1}^{2}\Big(\underline{\beta}_{ij}\big[q^b_{ij}u^b(\underline{v}) - t^b_{ij} - \log\big(\mu^b_i(\underline{v})\big)\big] + \overline{\beta}_{ij}\big[q^b_{ij}u^b(\overline{v}) - t^b_{ij} - \log\big(\mu^b_i(\overline{v})\big)\big]\Big) \\ &- \left[\underline{B}\log\left(\frac{\underline{B}}{\underline{B}+\underline{B}}\right) + \overline{B}\log\left(\frac{\overline{B}}{\underline{B}+\underline{B}}\right)\right], \end{aligned}$$

where $\underline{B} \equiv \sum_{i=1}^{2}\sum_{j=1}^{2}\underline{\beta}_{ij}$ and $\overline{B} \equiv \sum_{i=1}^{2}\sum_{j=1}^{2}\overline{\beta}_{ij}$.

Since all β_{ij}'s are weakly positive, we can obtain an better direction of payoff improvement, whose marginal gain will be equal to:

$$-\underline{B}y^b(\underline{v}) - \overline{B}y^b(\overline{v}) - \left[\underline{B}\log\left(\frac{\underline{B}}{\underline{B}+\underline{B}}\right) + \overline{B}\log\left(\frac{\overline{B}}{\underline{B}+\underline{B}}\right)\right] > 0$$

Defining $P \equiv \frac{\underline{B}}{\underline{B}+\overline{B}}$, we can write

$$-Py^b(\underline{v}) - (1-P)y^b(\overline{v}) - P\log\big(P\big) + (1-P)\log\big(1-P\big) > 0.$$

Maximizing over P, we can identify an even better direction of payoff improvement, moreover

$$\begin{aligned} &\max_P\Big\{-Py(\underline{v}) - (1-P)y(\overline{v}) - P\log\big(P\big) + (1-P)\log\big(1-P\big)\Big\} > 0 \\ \Leftrightarrow\ &\exp\big(-y^b(\underline{v})\big) + \exp\big(-y^b(\overline{v})\big) > 1, \end{aligned}$$

which establishes the "if" direction of Lemma 3.6 by contraposition. To see why the "only if" direction also holds, observe that if $\exp\big(-y^b(\underline{v})\big) + \exp\big(-y^b(\overline{v})\big) > 1$, then we can construct a profitable local Class 2 deviation by taking away some probability mass from those (i,j) in each state v, for which $q^b_{ij}u^b(v) - t^b_{ij} - \log\big(\mu^b_i(v)\big)$ is minimal, and putting this probability mass on $s^b_\emptyset$

3.3.2 Implementability conditions

Combining the Karush-Kuhn-Tucker optimality conditions from the maximization problems in (R-IC^s_A-1) and (R-IC^s_A-1) with the optimality conditions from Lemma 3.6, we obtain our main result:

Proposition 3.3. *The tuple (α, S^b, S^s) satisfies restricted ex ante incentive compatibility* R-IC^p_A *for both players $p \in \{b, s\}$ if and only if there are multipliers $\lambda^b_i(v), \lambda^s_j(v)$ for all i and j respectively, and $\phi^b_{ij}(v), \phi^s_{ij}(v)$ for all pairs (i, j) such that the following conditions are satisfied:*

$$
\begin{aligned}
&(\mathrm{ST}^b) \quad q^b_{ij}u^b(v) - t^b_{ij} - \log\left(\mu^b_i(v)\right) - \lambda^b_j(v) + \phi^b_{ij}(v) = 0 \ \forall(i,j),\ v \in V;\\
&(\mathrm{ST}^s) \quad t^s_{ij} - q^s_{ij}u^s(v) - \log\left(\mu^s_j(v)\right) - \lambda^s_i(v) + \phi^s_{ij}(v) = 0 \ \forall(i,j),\ v \in V;\\
&(\mathrm{DF}) \quad \phi^b_{ij}(v) \geq 0,\ \phi^s_{ij}(v) \geq 0 \ \forall(i,j),\ v \in V;\\
&(\mathrm{CS}) \quad \alpha_{ij}(v)\phi^b_{ij}(v) = 0,\ \alpha_{ij}(v)\phi^s_{ij}(v) = 0 \ \forall(i,j),\ v \in V;\\
&(\mathrm{NA}^b) \quad \sum_{v\in V} \exp\left(-\min_j\{\lambda^b_j(v)\}\right) \leq 1;\\
&(\mathrm{NA}^s) \quad \sum_{v\in V} \exp\left(-\min_i\{\lambda^s_i(v)\}\right) \leq 1.
\end{aligned}
$$

Proof. See Appendix B.8. □

Conditions (ST^b) and (ST^s) are stationarity conditions in the problems (R-IC^b_A-1) and (R-IC^s_A-1) respectively. (DF) are dual feasibility conditions, which make sure that the multipliers on non-negativity constraints on joint probabilities. (CS) are complementary slackness conditions. To get the no-abstention condition (NA^b), observe that, e.g. for the buyer, $q^b_{ij}u^b(v) - t^b_{ij} - \log\left(\mu^b_i(v)\right) = \lambda^b_j(v)$ whenever $\alpha_{ij}(v) > 0$ by stationarity and complementary slackness. Moreover, Lemma 3.5 implies that all seller's posteriors must be strictly positive, hence in each column j there will be at least one i such that $\alpha_{ij}(v) > 0$ in each state $v \in V$, implying that the minimum in (NA^b) can be taken over columns. (NA^b) can be obtained using a similar argument.

Proposition 3.3 allows us to restate the revenue maximization problem of the mechanism designer as follows:

$$\max_{\alpha,k,n;q,t;\phi,\lambda} \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V} \alpha_{ij}(v)\big(t_{ij}^{b} - t_{ij}^{b}\big), \text{ s.t.}$$

$$(\alpha\text{-F}) \quad S^b = \{1,\dots,k\},\ S^s = \{1,\dots,n\},\ \alpha \in \Delta\big(S^b \times S^s \times V\big),\ \text{BP};$$

$$(q\text{-F}) \quad 0 \leq q_{ij}^{b} \leq q_{ij}^{s} \leq 1 \quad \forall (i,j);$$

$$(\text{R-IC}_A) \quad \text{ST}^b,\ \text{ST}^s, \text{DF},\ \text{CS},\ \text{NA}^b,\ \text{NA}^s.$$

Since the specific labels of the signal realizations do not have any particular meaning in our analysis, we can without loss of generality assume that the mechanism designer decides on the number of signal realizations for each player only, and recommends $S^b = \{1,\dots,k\}$ for some $k \geq 1$ to the buyer and $S^s = \{1,\dots,n\}$ for some $n \geq 1$ to the seller.

Once k and n are fixed, the above maximization problem is a finite-dimensional *mathematical program with complementarity constraints*. The *complementarity constraints* (DF) and (CS) introduce non-convexities into the feasible set and make sure that most standard constraint qualifications do not hold at an optimum. Weaker constraint qualifications, however, have been shown to hold for these problems, and algorithms for numerical optimization have been developed (see e.g. Outrata et al. (1998)). Hence, one can take our implementability conditions from Proposition 3.3 and solve for an optimal mechanism with a given number of signal realizations numerically.

3.4 Impossibility of full surplus extraction

We take an analytical approach, however, and investigate whether the mechanism designer can extract full surplus from the players. We show below that full surplus extraction is impossible even in the simplest setting with a binary payoff-relevant state of the world, i.e. when $V = \{\underline{v}, \overline{v}\}$.

3.4.1 Individually uninformative efficient mechanisms

To extract full surplus, the mechanism designer must make sure that the mechanism he offers is *individually uninformative* (i.e. such that each individual player acquires no information about the payoff-relevant state) and *efficient*. If she offers an inefficient mechanism, then less than full trading surplus will be generated in the first place. If she offers an individually informative mechanism, then she will have to leave a chunk of the generated surplus to the agents to cover their information acquisition cost. Hence, both conditions are clearly necessary for full surplus extraction.

It is important to distinguish *individually uninformative* mechanisms from *uninformative* mechanisms. *Uninformative* mechanisms are such that no information is generated at all, even when the players' signals are combined. The mechanism designer, who offers an individually uninformative mechanism, can acquire some information about the payoff-relevant state by virtue of observing both players' reports, even when each individual player acquires no information at all.

In fact, Larionov et al. (2022) construct a mechanism and an information structure, such that the mechanism designer fully learns the payoff-relevant state of the world by observing just two realizations of individually uninformative signals. Larionov et al. (2022) show that full surplus extraction is implementable by this mechanism in any transferable environment with three or more players. We can first ask whether the same mechanism and the same information structure can help the mechanism designer achieve full surplus extraction in our bilateral trade setting. The information structure of Larionov et al. (2022), adapted to our setting, is as follows:

State $\underline{v}$	s_1^s	s_2^s
s_1^b	$\frac{\mu_0(\underline{v})}{2}$	0
s_2^b	0	$\frac{\mu_0(\underline{v})}{2}$

State $\overline{v}$	s_1^s	s_2^s
s_1^b	0	$\frac{\mu_0(\overline{v})}{2}$
s_2^b	$\frac{\mu_0(\overline{v})}{2}$	0

It is easy to verify that the proposed information structure is Bayes-plausible and individually uninformative: the posterior distribution of player p on the payoff-relevant states is always equal to the prior, no matter which signal signal has been observed by player p.

The allocation and the transfer functions are:

q^p	s_1^s	s_2^s
s_1^b	1	1
s_2^b	1	1

t^p	s_1^s	s_2^s
s_1^b	$u^p(\underline{v})$	$u^p(\overline{v})$
s_2^b	$u^p(\overline{v})$	$u^p(\underline{v})$

Observe that the signal realization profiles are arranged in such a way that allows the mechanism designer to fully learn the payoff-relevant state of the world from observing two *truthful* reports. Hence if the players had an incentive to report truthfully this mechanism would allow the mechanism designer to extract full surplus from the players. Unfortunately for the mechanism designer the players have an incentive to deviate. Consider for example the incentives of the buyer. If the buyer reports truthfully, then his payoff is zero since full surplus is then extracted by the seller. On the other hand, if the buyer deviates to the following information structure[3]:

State $\underline{v}$	s_1^s	s_2^s
s_1^b	$\frac{\mu_0(\underline{v})}{2}$	0
s_2^b	0	$\frac{\mu_0(\underline{v})}{2}$

State $\overline{v}$	s_1^s	s_2^s
s_1^b	$\frac{\mu_0(\overline{v})}{2}$	0
s_2^b	0	$\frac{\mu_0(\overline{v})}{2}$

,

then his payoff will be:

$$\mu_0(\underline{v})\big(u^b(\underline{v}) - u^b(\underline{v})\big) + \mu_0(\overline{v})\big(u^b(\overline{v}) - u^b(\underline{v})\big) = \mu_0(\overline{v})\big(u^b(\overline{v}) - u^b(\underline{v})\big) > 0,$$

making the deviation profitable. Hence the mechanism constructed by Larionov et al. (2022) does not help in this setting. One might still wonder whether it is possible to extract full surplus via a more sophisticated mechanism. We show below that it is impossible: the revenue achieved by any *individually uninformative efficient* mechanism cannot exceed the revenue achieved by the revenue-maximizing *uninformative* mechanism.

3.4.2 Revenue-maximizing uninformative mechanism

To establish our impossibility result, let us consider *uninformative* mechanisms. In an uninformative mechanism, the players generate *trivial* signals, i.e. signals with singleton signal realization sets ($S^p = \{s^p\}$ for each player p). Using the global implementability conditions derived in Proposition 3.3, we can write the mechanism designer's revenue maximization

[3]The new information structure is clearly also individually uninformative, it also respects the marginal distribution on $S^s \times V$ and hence can be induced by a buyer's deviation to a different signal.

problem as:

$$\max_{q,t,\lambda} \; t^b - t^s, \quad \text{s.t.}$$

$$(\text{ST}^b) \quad q^b u^b(\underline{v}) - t^b - \log\big(\mu_0(\underline{v})\big) - \lambda^b(\underline{v}) = 0, \quad q^b u^b(\overline{v}) - t^b - \log\big(\mu_0(\overline{v})\big) - \lambda^b(\overline{v}) = 0;$$

$$(\text{ST}^s) \quad t^s - q^s u^s(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \lambda^s(\underline{v}) = 0, \quad t^s - q^s u^s(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda^s(\overline{v}) = 0;$$

$$(\text{NA}^b) \quad \exp(-\lambda^b\big(\underline{v})\big) + \exp(-\lambda^b\big(\overline{v})\big) \leq 1;$$

$$(\text{NA}^s) \quad \exp(-\lambda^s\big(\underline{v})\big) + \exp(-\lambda^s\big(\overline{v})\big) \leq 1;$$

$$(q\text{-F}) \quad 0 \leq q^b \leq q^s \leq 1.$$

The solution to this revenue-maximization problem is summarized in the next lemma:

Lemma 3.7. *The mechanism that sets $q^b = q^s = q^*$, where q^* solves*

$$\max_q \Big\{ -\log\big[\mu_0(\underline{v})\exp(-qu^b(\underline{v})) + \mu_0(\overline{v})\exp(-qu^b(\overline{v}))\big] - \log\big[\mu_0(\underline{v})\exp(qu^s(\underline{v})) + \mu_0(\overline{v})\exp(qu^s(\overline{v}))\big],$$
$$\text{s.t } 0 \leq q \leq 1\Big\},$$

and the optimal transfers are given by $t^b(q^*) \equiv -\log\big[\mu_0(\underline{v})\exp(-q^*u^b(\underline{v})) + \mu_0(\overline{v})\exp(-q^*u^b(\overline{v}))\big]$ and $t^s(q^*) \equiv \log\big[\mu_0(\underline{v})\exp(q^*u^s(\underline{v})) + \mu_0(\overline{v})\exp(q^*u^s(\overline{v}))\big]$ is a revenue-maximizing uninformative mechanism with two payoff-relevant states.

Proof. See Appendix B.9. □

Observe that the following is true for the buyer due to strict convexity of $\exp(\cdot)$:

$$t^b(q^*) \leq -\log\big[\mu_0(\underline{v})\exp(-u^b(\underline{v})) + \mu_0(\overline{v})\exp(-u^b(\overline{v}))\big] < \mu_0(\underline{v})u^b(\underline{v}) + \mu_0(\overline{v})u^b(\overline{v}).$$

Likewise, the following is true for the seller due to strict convexity of $\exp(\cdot)$:

$$t^s(q^*) \geq \log\big[\mu_0(\underline{v})\exp(u^s(\underline{v})) + \mu_0(\overline{v})\exp(u^s(\overline{v}))\big] > \mu_0(\underline{v})u^s(\underline{v}) + \mu_0(\overline{v})u^s(\overline{v}).$$

We can therefore conclude that the revenue from the revenue-maximizing uninformative mechanism of Lemma 3.7 is always strictly below full surplus:

$$t^b(q^*) - t^s(q^*) < \mu_0(\underline{v})\big[u^b(\underline{v}) - u^s(\underline{v})\big] + \mu_0(\overline{v})\big[u^b(\overline{v}) - u^s(\overline{v})\big] = \text{Full surplus}$$

3.4.3 Impossibility

In order to show that full surplus extraction is not implementable, it remains to show that the revenue from any *individually uninformative efficient* mechanism cannot exceed the revenue of the revenue-maximizing uninformative mechanism of Lemma 3.7. The following proposition accomplishes this task.

Proposition 3.4. *Suppose $V = \{\underline{v}, \overline{v}\}$, and consider any individually uninformative efficient mechanism. The revenue from such a mechanism does not exceed the revenue of the revenue-maximizing uninformative mechanism of Lemma 3.7.*

Proof. See Appendix B.10. □

The proof of Proposition 3.4 is relegated to the Appendix, but the core argument can be summarized as follows. First of all, the stationarity and dual feasibility conditions of the buyer imply that for every j the transfer of the buyer is given by $t^b_{ij} = u^b(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \lambda_j(\underline{v}) = u^b(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda_j(\overline{v})$ for all i such that $\alpha_i j(\underline{v}) + \alpha_{ij}(\overline{v})$, i.e. for those pairs of signal realizations (s^b_i, s^s_j) that occur with positive probability in at least one state of the world. The no-abstention condition (NA^b) then implies that the buyer's transfer can be bounded from above:

$$t^b_{ij} \leq -\log\left[\mu_0(\underline{v})\exp(-u^b(\underline{v})) + \mu_0(\overline{v})\exp(-u^b(\overline{v}))\right]$$

for all (i, j) such that (s^b_i, s^s_j) occurs with positive probability.

A similar argument for the seller allows us to bound his transfer from below:

$$t^s_{ij} \geq \log\left[\mu_0(\underline{v})\exp(u^s(\underline{v})) + \mu_0(\overline{v})\exp(u^s(\overline{v}))\right]$$

for all (i, j) such that (s^b_i, s^s_j) occurs with positive probability.

The revenue from any *individually uninformative efficient* mechanism can be then bounded from above as follows:

$$\begin{aligned}&\sum_{i=1}^{k}\sum_{j=1}^{n}\big(\alpha_{ij}(\underline{v}) + \alpha_{ij}(\overline{v})\big)(t^b_{ij} - t^s_{ij})\\ &\leq -\log\left[\mu_0(\underline{v})\exp(-u^b(\underline{v})) + \mu_0(\overline{v})\exp(-u^b(\overline{v}))\right] - \log\left[\mu_0(\underline{v})\exp(u^s(\underline{v})) + \mu_0(\overline{v})\exp(u^s(\overline{v}))\right],\end{aligned}$$

where the right-hand side clearly does not exceed the revenue generated by the mechanism of Lemma 3.7.

3.5 Concluding remarks

We have considered a mechanism design problem with information acquisition in a bilateral trade environment. At the beginning, the buyer, the seller, and the mechanism designer have no information about the good's quality beyond a common prior. The buyer and the seller can generate signals from a large signal space to acquire more information about the good's quality. The mechanism designer commits to a mechanism taking information acquisition by the players into account.

We characterize the set of implementable mechanisms in this environment. To check whether a particular tuple of allocations, transfers, and signals is implementable, one has to check whether these allocations, transfers, and the information structure induced by the signals satisfy a finite-dimensional system of equations and inequalities.

We use our implementability conditions to show that the mechanism designer cannot extract full surplus from the players. We leave the characterization of revenue-maximizing mechanisms for future research.

Appendix A

Appendix to Chapter 1

A.1 Solution of the one-shot auction problem

Let us first consider the choices made by the buyers who face a reserve price r. Depending on the reserve price chosen by the seller, there are four possible cases to consider:

Case i: $r \leq \underline{\theta}$

In this case both types of each buyer will be willing to participate in the auction. The low types will bid their own valuation $\underline{\theta}$ and receive the payoff of 0. The high types will randomize on $(\underline{\theta}, \overline{b}]$ where $\overline{b} = (1 - q^{n-1})\overline{\theta} + q^{n-1}\underline{\theta}$ according to

$$G(b) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta} - \underline{\theta}}{\overline{\theta} - b}\right)^{\frac{1}{n-1}} - 1\right],$$

and will get the payoff of $q^{n-1}(\overline{\theta} - \underline{\theta})$. The *ex ante* equilibrium payoff of the buyers is:

$$v^*_{r \leq \underline{\theta}} = (1-q)q^{n-1}(\overline{\theta} - \underline{\theta}).$$

The seller generates revenue:

$$\mathcal{R}^*_{r \leq \underline{\theta}} = (1 - q^n)\overline{\theta} + q^n\underline{\theta} - nv^*_{r \leq \underline{\theta}} = (1 - q^n)\overline{\theta} + q^n\underline{\theta} - n(1-q)q^{n-1}(\overline{\theta} - \underline{\theta}).$$

Case ii: $\underline{\theta} < r < \overline{\theta}$

In this case only the high types are willing to participate in the first price auction. The high types will randomize on $(\underline{\theta}, \overline{b}]$ where $\overline{b} = (1 - q^{n-1})\overline{\theta} + q^{n-1}r$ according to

$$G(b) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta} - r}{\overline{\theta} - b}\right)^{\frac{1}{n-1}} - 1\right],$$

and will get the payoff of $q^{n-1}(\overline{\theta} - r)$, which leads to the ex ante equilibrium payoff of:

$$v^*_{\underline{\theta}<r<\overline{\theta}} = (1 - q)q^{n-1}(\overline{\theta} - r).$$

The resulting revenue of the seller who chooses a reserve price $r \leq \underline{\theta}$:

$$\mathcal{R}^*_{\underline{\theta}<r<\overline{\theta}} = (1 - q^n)\overline{\theta} - nv^*_{\underline{\theta}<r<\overline{\theta}} = (1 - q^n)\overline{\theta} - n(1 - q)q^{n-1}(\overline{\theta} - r).$$

Case iii: $r = \overline{\theta}$

In this case only high types are willing to participate, and they of course have no choice but to bid $\overline{\theta}$ in equilibrium, and the resulting revenue will be:

$$\mathcal{R}^*_{r=\overline{\theta}} = (1 - q^n)\overline{\theta}.$$

Case iv: $r > \overline{\theta}$

In this case neither type wants to participate, so every buyer will choose to abstain and the seller will get zero revenue.

Revenue achieved in **Case ii** is clearly inferior to that achieved in **Case iii**, so setting $\underline{\theta} < r < \overline{\theta}$ cannot be part of any subgame-pefect equilibrium of the static auction game. The reserve prices $r \leq \underline{\theta}$ and $r = \overline{\theta}$ could however be optimal for the seller.

Case i: $r \leq \underline{\theta}$

In this case it is clear that both types will participate will be willing to participate. It can be easily shown that there is no Nash equilibrium in pure strategies. It is also immediately clear that the low types will never place a bid higher than their own valuation because winning with such a high bid would lead to a negative payoff. But low types should not place a bid that is lower than their valuation even if they have an opportunity to do so. Suppose low type bidders do place a bid $r < \underline{b} < \underline{\theta}$ in equilirium, then one of them could deviate to $\underline{b} + \epsilon$ and guarantee winning the auction for sure if his competitor is of low type as well, hence there is a profitable deviation.

Suppose $\Phi(b)$ is the unconditional distribution of equilibrium bids for every player. The expected payoff of a bidder with type $\overline{\theta}$ is given by:

$$\Phi^{n-1}(b)(\overline{\theta} - b). \tag{A.1}$$

Assuming that only low types bid $\underline{\theta}$ we must have $\Phi(\underline{\theta}) = q$ hence by indifference we have:

$$\Phi^{n-1}(b)(\overline{\theta} - b) = q^{n-1}(\overline{\theta} - \underline{\theta}). \tag{A.2}$$

hence $\Phi(b) = q\left(\frac{\overline{\theta}-\underline{\theta}}{\overline{\theta}-b}\right)^{\frac{1}{n-1}}$. To find the upper bound of the support we solve $q\left(\frac{\overline{\theta}-\underline{\theta}}{\overline{\theta}-\overline{b}}\right)^{\frac{1}{n-1}} = 1$, which leads to $\overline{b} = (1 - q^{n-1})\overline{\theta} + q^{n-1}\underline{\theta}$. Hence the high type player randomizes over $\big(\underline{\theta}, \overline{b}\big]$. Since $\Phi(b)$ is the unconditional distribution of equilibrium bids, the actual mixed strategy of the high type is:

$$G(b) \equiv \Phi(b|\theta_i = \overline{\theta}) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta} - \underline{\theta}}{\overline{\theta} - b}\right)^{\frac{1}{n-1}} - 1\right]. \tag{A.3}$$

The above analysis naturally leads to the following lemma:

Lemma A.1. *If $r \leq \underline{\theta}$,*

(i) the low type bids his own valuation in equilibrium: $\underline{b} = \underline{\theta}$,

(ii) the high type randomizes his bids on $\big(\underline{\theta}, (1 - q^{n-1})\overline{\theta} + q^{n-1}\underline{\theta}\big]$ according to

$$G(b) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta} - \underline{\theta}}{\overline{\theta} - b}\right)^{\frac{1}{n-1}} - 1\right].$$

The low type expected equilibrium payoff is 0, the high type expected equilibrium payoff is $q^{n-1}(\overline{\theta} - \underline{\theta})$, which leads to the ex ante equilibrium payoff of:

$$v^*_{r\leq\underline{\theta}} = (1 - q)q^{n-1}(\overline{\theta} - \underline{\theta}).$$

The equilibrium in Lemma A.1 is efficient, hence it leads to the total surplus given by: $(1-q^n)\overline{\theta}+q^n\underline{\theta}$. The resulting revenue of the seller who chooses a reserve price $r \leq \underline{\theta}$:

$$\begin{aligned}\mathcal{R}^*_{r\leq\underline{\theta}} &= (1-q^n)\overline{\theta}+q^n\underline{\theta}-nu_i^* \\ &= (1-q^n)\overline{\theta}+q^n\underline{\theta}-n(1-q)q^{n-1}(\overline{\theta}-\underline{\theta}).\end{aligned} \tag{A.4}$$

Case ii: $\underline{\theta} < r < \overline{\theta}$

In this case only the high types are willing to participate in the first price auction. It can also be shown that there is no equilibrium in pure strategies. Hence we will be looking for an equilibrium in mixed strategies. Suppose that a high type buyer randomizes his bids according to the distribution function $G(b)$. The payoff of a high type buyer who is bidding b is given by:

$$\begin{aligned}&\Big(q^{n-1}+(n-1)(1-q)q^{n-2}G(b)+\ldots+(1-q)^{n-1}G^{n-1}(b)\Big)(\overline{\theta}-b) \\ &= \big(q+(1-q)G(b)\big)^{n-1}(\overline{\theta}-b).\end{aligned} \tag{A.5}$$

Assuming that r is the lower bound of the support of $G(b)$ and that $G(b)$ has no mass points we get $G(r)=0$. By indifference we get for every b in the support:

$$\big(q+(1-q)G(b)\big)^{n-1}(\overline{\theta}-b) = \big(q+(1-q)G(r)\big)^{n-1}(\overline{\theta}-r) = q^{n-1}(\overline{\theta}-r), \tag{A.6}$$

which immediately gives us:

$$G(b) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta}-r}{\overline{\theta}-b}\right)^{\frac{1}{n-1}}-1\right]. \tag{A.7}$$

To find the upper bound of the support $\overline{b}$ we solve $\frac{q}{1-q}\big[\big(\frac{\overline{\theta}-r}{\overline{\theta}-\overline{b}}\big)^{\frac{1}{n-1}}-1\big]=1$ which leads to $\overline{b}=(1-q^{n-1})\overline{\theta}+q^{n-1}r$. Hence the following lemma:

Lemma A.2. *If* $\underline{\theta} < r < \overline{\theta}$,

(i) the low type chooses to abstain from participation $\underline{b}=\emptyset$,

(ii) the high type randomizes his bids on $\left[r, (1-q^{n-1})\overline{\theta} + q^{n-1}r\right]$ *according to*

$$G(b) = \frac{q}{1-q}\left[\left(\frac{\overline{\theta}-r}{\overline{\theta}-b}\right)^{\frac{1}{n-1}} - 1\right].$$

The low type expected equilibrium payoff is 0*, the high type expected equilibrium payoff is* $q^{n-1}(\overline{\theta}-r)$*, which leads to the ex ante equilibrium payoff of:*

$$v^*_{\underline{\theta}<r<\overline{\theta}} = (1-q)q^{n-1}(\overline{\theta}-r).$$

Since only the high types trade with the seller in the equilibrium in Lemma A.2, the resulting total surplus is given by: $(1-q^n)\overline{\theta}$. The resulting revenue of the seller who chooses a reserve price $r \leq \underline{\theta}$:

$$\mathcal{R}^*_{\underline{\theta}<r<\overline{\theta}} = (1-q^n)\overline{\theta} - nu^*_i = (1-q^n)\overline{\theta} - n(1-q)q^{n-1}(\overline{\theta}-r). \tag{A.8}$$

Case iii: $r = \overline{\theta}$

In this case only high types are willing to participate, and they of course have no choice but to bid $b = \overline{\theta}$ in equilibrium, and the resulting revenue will be:

$$\mathcal{R}^*_{r=\overline{\theta}} = (1-q^n)\overline{\theta}. \tag{A.9}$$

Case iv: $r > \overline{\theta}$

In this case neither type wants to participate, so every buyer will choose to abstain and the seller will get zero revenue.

A.2 Separating equilibrium payoffs

Suppose that in every period along the equilibrium path a low type buyer bids $\underline{b}$, and a high type buyer bids $\bar{b}$. Then a low type bidder wins with probability $1/n$ only if all his competitors are of low type as well, hence his equilibrium payoff is given by:

$$\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}). \tag{A.10}$$

A high type bidder may win in several different cases: whenever $k-1$ of his competitors are also high type buyers, he wins with probability $1/k$, hence his winning probability is equal to:

$$
\begin{aligned}
&(1-q)^{n-1}\frac{1}{n}+(n-1)q(1-q)^{n-2}\frac{1}{n-1}+\frac{(n-1)(n-2)}{2}q^2(1-q)^{n-3}\frac{1}{n-2}+...+q^{n-1}1\\
&=(1-q)^{n-1}\frac{1}{n}+q(1-q)^{n-2}+\frac{(n-1)}{2}q^2(1-q)^{n-3}+...+q^{n-1}\\
&=\frac{1}{n}\left[(1-q)^{n-1}+nq(1-q)^{n-2}+\frac{n(n-1)}{2}q^2(1-q)^{n-3}+...+nq^{n-1}\right]\\
&=\frac{1}{n(1-q)}\left[(1-q)^{n}+nq(1-q)^{n-1}+\frac{n(n-1)}{2}q^2(1-q)^{n-2}+...+nq^{n-1}(1-q)\right]\\
&=\frac{1}{n(1-q)}\left[\underbrace{(1-q)^{n}+nq(1-q)^{n-1}+\frac{n(n-1)}{2}q^2(1-q)^{n-2}+...+nq^{n-1}(1-q)+q^n}_{=(1-q+q)^n=1}-q^n\right]\\
&=\frac{1}{n(1-q)}(1-q^n).
\end{aligned}
$$

The expected payoff of a high type's buyer then is:

$$\frac{1-q^n}{n(1-q)}(\bar{\theta}-\bar{b}). \tag{A.11}$$

The resulting *ex ante* equilibrium payoff of each buyer is then

$$
\begin{aligned}
v_i &= (1-q)\frac{1-q^n}{n(1-q)}(\bar{\theta}-\bar{b})+q\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}) \\
&=\frac{1}{n}\big[(1-q^n)(\bar{\theta}-\bar{b})+q^n(\underline{\theta}-\underline{b})\big].
\end{aligned} \tag{A.12}
$$

The resulting revenue of the seller is:

$$\mathcal{R}_s = (1 - q^n)\bar{b} + q^n \underline{b}. \tag{A.13}$$

A.3 Proof of Proposition 1.4

Proof. [1] Note first that both in the low-revenue separating and zero-revenue pooling equilibrium, the buyer-game induced by the seller's equilibrium strategy is the repeated first-price auction game with zero reserve price. Denote $\mathcal{V}$ the set of strongly symmetric public perfect equilibrium payoffs of this buyer-game. Denote $\hat{v} = \sup \mathcal{V}$. We have to distinguish two classes of strongly symmetric public perfect equilibria: (i) equilibria in which a separating bidding profile is played *in the first period*, and (ii) equilibria in which a pooling bidding profile is played *in the first period*.

(i) *A separating bidding profile is played in the first period*

Suppose first that the optimal payoff $\hat{v}$ is achieved by a symmetric public perfect equilibrium in which the buyers separate in the first period. Suppose $b(\cdot)$ is the equilibrium action taken in the first period. Denote $\underline{b}$ and $\overline{b}$ the bids placed in the first period by a low-type buyer and a high-type buyer respectively. Suppose that the equilibrium continuation value after the first period is given by $v^* : \mathbb{R}^n_+ \to \mathbb{R}$, then the equilibrium payoff of a high-type buyer i is given by:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}) + \delta\,\mathbb{E}\left(v^*(\overline{b}, b(\theta_{-i}))\right).$$

The equilibrium payoff of a low-type buyer i is given by:

$$(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}) + \delta\,\mathbb{E}\left(v^*(\underline{b}, b(\theta_{-i}))\right).$$

The on-schedule incentive compatibility constraint of a low-type buyer is then given by:

$$(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}) + \delta\,\mathbb{E}\left(v^*(\underline{b}, b(\theta_{-i}))\right) \geq (1-\delta)\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b}) + \delta\,\mathbb{E}\left(v^*(\overline{b}, b(\theta_{-i}))\right).$$

[1] See a similar argument in Chapter 11.2 of Mailath and Samuelson (2006) in the context of a repeated price competition game with adverse selection.

Subtract $\delta\hat{v}$ and divide both sides by $(1-\delta)$:

$$\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b})+\frac{\delta}{1-\delta}\,\mathbb{E}\left(v^*(\underline{b},b(\theta_{-i}))-\hat{v}\right)\geq\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b})+\frac{\delta}{1-\delta}\,\mathbb{E}\left(v^*(\overline{b},b(\theta_{-i}))-\hat{v}\right),$$

and define $\overline{x}\equiv\frac{\delta}{1-\delta}\,\mathbb{E}\left(v^*(\overline{b},b(\theta_{-i}))-\hat{v}\right)$ and $\underline{x}\equiv\frac{\delta}{1-\delta}\,\mathbb{E}\left(v^*(\underline{b},b(\theta_{-i}))-\hat{v}\right)$. The incentive compatibility constraint of a low-type buyer can then be written as:

$$\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b})+\underline{x}\geq\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b})+\overline{x}. \tag{A.14}$$

Recall that the continuation payoffs in any strongly symmetric public perfect equilibrium must be strongly symmetric public perfect equilibrium payoffs themselves, hence we must have $\overline{x}\leq 0$ and $\underline{x}\leq 0$ since $\hat{v}=\sup\mathcal{V}$.

The *ex ante* equilibrium payoff is given by:

$$\hat{v}=(1-\delta)\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\overline{b})+q^n(\underline{\theta}-\underline{b})\right]+(1-q)\delta\,\mathbb{E}\left(v^*(\overline{b},b(\theta_{-i}))\right)+q\delta\,\mathbb{E}\left(v^*(\underline{b},b(\theta_{-i}))\right).$$

Subtracting $\delta\hat{v}$ and dividing by $(1-\delta)$ on both sides, we obtain:

$$\hat{v}=\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\overline{b})+q^n(\underline{\theta}-\underline{b})\right]+(1-q)\overline{x}+q\underline{x}.$$

Combining this expression with the low-type incentive compatibilty constraint in (A.14) and our observation that $\underline{x},\overline{x}\leq 0$, we must conclude that[2]:

$$\hat{v}\leq\max_{\underline{b},\overline{b};\underline{x},\overline{x}}\frac{1}{n}\left[(1-q^n)(\overline{\theta}-\overline{b})+q^n(\underline{\theta}-\underline{b})\right]+(1-q)\overline{x}+q\underline{x}\quad\text{subject to} \tag{A.15}$$

$$\text{(IC)}\quad\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b})+\underline{x}\geq\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b})+\overline{x},$$

$$\text{(Feas)}\quad\underline{x},\overline{x}\leq 0.$$

Let us consider the maximization problem in (A.15). Clearly the (IC) constraint must be binding at the optimum: suppose not, i.e. suppose $\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b})+\underline{x}>\frac{1-q^n}{n(1-q)}(\underline{\theta}-\overline{b})+\overline{x}$,

[2] The solution to this maximization problem provides an upper bound on strongly symmetric equilibrium payoffs since all the other incentive compatibility constraints are ignored, and the constraint $\underline{x},\overline{x}\leq 0$ is necessary for feasibility of continuation values but not sufficient.

then choose $\overline{b}' < \overline{b}$ such that the constraint is still satisfied, and this will clearly improve the value of the objective. Hence, at the optimum of (A.15), we must have

$$\frac{q^{n-1}}{n}(\underline{\theta} - \underline{b}) + \underline{x} = \frac{1 - q^n}{n(1-q)}(\underline{\theta} - \overline{b}) + \overline{x},$$

which we can solve for $(1 - q^n)(\underline{\theta} - \overline{b})$, to obtain:

$$(1 - q^n)(\underline{\theta} - \overline{b}) = (1 - q)q^{n-1}(\underline{\theta} - \underline{b}) + n(1 - q)(\underline{x} - \overline{x}),$$

which then implies:

$$(1 - q^n)(\overline{\theta} - \overline{b}) = (1 - q^n)(\overline{\theta} - \underline{\theta}) + (1 - q)q^{n-1}(\underline{\theta} - \underline{b}) + n(1 - q)(\underline{x} - \overline{x}). \qquad \text{(A.16)}$$

Plugging (A.16) into the objective function in (A.15), we get:

$$\begin{aligned}
&\frac{1}{n}\big[(1 - q^n)(\overline{\theta} - \underline{\theta}) + (1 - q)q^{n-1}(\underline{\theta} - \underline{b}) + n(1 - q)(\underline{x} - \overline{x}) + q^n(\underline{\theta} - \underline{b})\big] + (1 - q)\overline{x} + q\underline{x} \\
&= \frac{1}{n}(1 - q^n)(\overline{\theta} - \underline{\theta}) + \frac{1}{n}\big[(1 - q)q^{n-1} + q^n\big](\underline{\theta} - \underline{b}) + (1 - q)(\underline{x} - \overline{x}) + (1 - q)\overline{x} + q\underline{x} \\
&= \frac{1}{n}(1 - q^n)(\overline{\theta} - \underline{\theta}) + \frac{1}{n}\big[(1 - q)q^{n-1} + q^n\big](\underline{\theta} - \underline{b}) + \underline{x} \\
&= \frac{1}{n}(1 - q^n)(\overline{\theta} - \underline{\theta}) + \frac{1}{n}q^{n-1}(\underline{\theta} - \underline{b}) + \underline{x},
\end{aligned}$$

which implies that:

$$\hat{v} \leq \max_{\underline{b},\underline{x}} \frac{1}{n}(1 - q^n)(\overline{\theta} - \underline{\theta}) + \frac{1}{n}q^{n-1}(\underline{\theta} - \underline{b}) + \underline{x} \;\; \text{subject to } \overline{x} \leq 0.$$

The optimum is clearly achieved when $\underline{b} = 0$ and $\underline{x} = 0$, which means that:

$$\hat{v} \leq \frac{1}{n}(1 - q^n)(\overline{\theta} - \underline{\theta}) + \frac{1}{n}q^{n-1}\underline{\theta} = v^*_{\text{lrs}}.$$

Hence, if the buyers play a separating bidding profile in the first period in an optimal strongly symmetric equilibrium of this buyer-game, then the optimal equilibrium payoff cannot exceed the equilibrium payoff of the low-revenue separating equilibrium.

(ii) *A pooling bidding profile is played in the first period*

Consider now a class of strongly symmetric public perfect equilibria in which the buyers pool in the first period, and denote b the equilibrium action of both types in the first period. Suppose that the optimal payoff $\hat{v}$ is achieved by an equilibrium in this class. Suppose that $v^* : \mathbb{R}^n_+ \to \mathbb{R}$ is the equilibrium continuation value after the first period. The *ex ante* equilibrium payoff is given by:

$$\hat{v} = (1-\delta)\frac{1}{n}\big(\mathbb{E}(\theta) - b\big) + \delta v^*(b, \ldots, b)$$

Subtracting $\delta\hat{v}$ and dividing by $(1-\delta)$ on both sides, we obtain:

$$\frac{\hat{v} - \delta\hat{v}}{1-\delta} = \frac{1}{n}\big(\mathbb{E}(\theta) - b\big) + \frac{\delta}{1-\delta}(v^*(b, \ldots, b) - \hat{v})$$

Denote $x = \frac{\delta}{1-\delta}(v^*(b, \ldots, b) - \hat{v})$ and rewrite the above expression as:

$$\hat{v} = \frac{1}{n}\big(\mathbb{E}(\theta) - b\big) + x$$

Since continuation values must be strongly symmetric equilibrium payoffs themselves, we have $x \leq 0$, and therefore:

$$\begin{aligned} \hat{v} &\leq \max_{b,x} \frac{1}{n}\big(\mathbb{E}(\theta) - b\big) + x \quad \text{subject to } x \leq 0 \\ &= \frac{1}{n}\,\mathbb{E}(\theta) = v^*_{\text{zrp}} \end{aligned}$$

Hence, if the buyers play a pooling bidding profile in the first period in an optimal strongly symmetric equilibrium of this buyer-game, then the optimal equilibrium payoff cannot exceed the equilibrium payoff of the zero-revenue pooling equilibrium.

We can now conclude that there are only two candidates for the optimal strongly symmetric public perfect equilibrium payoff of the buyer-game: the payoff from the low-revenue separating equilibrium and the payoff from the zero-revenue pooling equilibrium. The result then follows from the analysis in the main text.

□

A.4 Proof of the Monotonicity lemma

Proof. Consider first the high reserve price state ω^h. Clearly in any public perfect equilibrium the payoff in this state must be zero, hence we can without loss of generality assume that $b_{\omega^h}(\overline{\theta}) = \overline{\theta}$ and $b_{\omega^h}(\underline{\theta}) = \emptyset$.

Consider now the low reserve price state ω^l, in which the buyer-game starts. Consider any strongly symmetric public perfect equilibrium of the buyer game. Pick any history that leads to state ω^l and suppose any high-type buyer bids according to $b_{\omega^l}(\overline{\theta}) = \overline{b}$ and any low-type buyer bids according to $b_{\omega^l}(\underline{\theta}) = \underline{b}$ after that history, and the equilibrium continuation value is given by $v^*_{\omega^l}(b) : A^n(\omega^l) \to \mathbb{R}$. The equilibrium payoff of a high-type buyer is given by:

$$(1-\delta)p(\overline{b})(\overline{\theta} - \overline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\overline{b}, b_{\omega^l}(\theta_{-i}))\right), \tag{A.17}$$

where $p(\overline{b})$ is the winning probability from bidding $\overline{b}$ in the current period. Analogously the equilibrium payoff of a low-type buyer i is equal to:

$$(1-\delta)p(\underline{b})(\underline{\theta} - \underline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\underline{b}, b_{\omega^l}(\theta_{-i}))\right), \tag{A.18}$$

where $p(\underline{b})$ is the winning probability from bidding $\underline{b}$ in the current period.

Since the above are assumed to be public perfect equilibrium payoffs, the following incentive compatibility must be satisfied, for a high type buyer:

$$(1-\delta)p(\overline{b})(\overline{\theta} - \overline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\overline{b}, b_{\omega^l}(\theta_{-i}))\right) \geq (1-\delta)p(\underline{b})(\overline{\theta} - \underline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\underline{b}, b_{\omega^l}(\theta_{-i}))\right), \tag{A.19}$$

and for a low type buyer:

$$(1-\delta)p(\underline{b})(\underline{\theta} - \underline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\underline{b}, b_{\omega^l}(\theta_{-i}))\right) \geq (1-\delta)p(\overline{b})(\underline{\theta} - \overline{b}) + \delta\,\mathbb{E}\left(v^*_{\omega^l}(\overline{b}, b_{\omega^l}(\theta_{-i}))\right). \tag{A.20}$$

Adding inequalities (A.19) and (A.20) together and canceling the continuation values on both sides, we obtain:

$$
\begin{aligned}
&(1-\delta)p(\bar{b})(\bar{\theta}-\bar{b})+(1-\delta)p(\underline{b})(\underline{\theta}-\underline{b}) \geq (1-\delta)p(\underline{b})(\bar{\theta}-\underline{b})+(1-\delta)p(\bar{b})(\underline{\theta}-\bar{b})\\
&\Leftrightarrow p(\bar{b})(\bar{\theta}-\bar{b})+p(\underline{b})(\underline{\theta}-\underline{b}) \geq p(\underline{b})(\bar{\theta}-\underline{b})+p(\bar{b})(\underline{\theta}-\bar{b})\\
&\Leftrightarrow p(\bar{b})\bar{\theta}+p(\underline{b})\underline{\theta} \geq p(\underline{b})\bar{\theta}+p(\bar{b})\underline{\theta}\\
&\Leftrightarrow p(\bar{b})(\bar{\theta}-\underline{\theta})+p(\underline{b})(\underline{\theta}-\bar{\theta}) \geq 0\\
&\Leftrightarrow \big(p(\bar{b})-p(\underline{b})\big)(\bar{\theta}-\underline{\theta}) \geq 0\\
&\Leftrightarrow p(\bar{b})-p(\underline{b}) \geq 0,
\end{aligned}
$$

which implies that $\bar{b} \geq \underline{b}$. □

A.5 Solutions of equilibrium conditions

A.5.1 Solution of Case 1

Recall that the equilibrium conditions in Case 1 are:

$$v^*_{\text{fse}} = \frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{b}^*)}{n(1-\delta(1-q)^n)}, \tag{A.21}$$

and:

$$(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta} - \underline{b}^*) + \delta v^*_{\text{fse}} = 0, \tag{A.22}$$

where $v^*_{\text{fse}} = \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\underline{\theta} - \underline{b}^*)\big]$.

Combining the equations (A.21) and (A.22), we get

$$(1-\delta)\frac{q^{n-1}}{n}(\underline{\theta} - \underline{b}^*) + \delta\frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{b}^*)}{n(1-\delta(1-q)^n)} = 0, \tag{A.23}$$

which we can solve for the equilibrium value of $\underline{b}$:

$$\underline{b}^* = \frac{\delta q\big(1-q^n\big)\overline{\theta} + q^n\big(1-\delta(1-q)^n\big)\underline{\theta}}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)}, \tag{A.24}$$

which we can now use to compute the payoff of each type conditional upon winning with $\underline{b}^*$, for a low type buyer we have:

$$\begin{aligned}
\underline{\theta} - \underline{b}^* &= \underline{\theta} - \frac{\delta q\big(1-q^n\big)\overline{\theta} + q^n\big(1-\delta(1-q)^n\big)\underline{\theta}}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} \\
&= \frac{\delta q\big(1-q^n\big)\underline{\theta} + q^n\big(1-\delta(1-q)^n\big)\underline{\theta} - \delta q\big(1-q^n\big)\overline{\theta} - q^n\big(1-\delta(1-q)^n\big)\underline{\theta}}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} \\
&= \frac{-\delta q\big(1-q^n\big)(\overline{\theta} - \underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} < 0;
\end{aligned} \tag{A.25}$$

and for a high type buyer we have:

$$\begin{aligned}\overline{\theta}-\underline{b}^{*} &= \overline{\theta}-\frac{\delta q\big(1-q^{n}\big)\overline{\theta}+q^{n}\big(1-\delta(1-q)^{n}\big)\underline{\theta}}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)} \\ &= \frac{\delta q\big(1-q^{n}\big)\overline{\theta}+q^{n}\big(1-\delta(1-q)^{n}\big)\overline{\theta}-\delta q\big(1-q^{n}\big)\overline{\theta}-q^{n}\big(1-\delta(1-q)^{n}\big)\underline{\theta}}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)} \\ &= \frac{q^{n}\big(1-\delta(1-q)^{n}\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)}>0,\end{aligned} \tag{A.26}$$

which combined with (A.21) gives us the resulting equilibrium payoff:

$$\begin{aligned}v^{*}_{\text{fse}} &= \frac{(1-\delta)(1-q^{n})(\overline{\theta}-\underline{b}^{*})}{n(1-\delta(1-q)^{n})}= \\ &= \frac{(1-\delta)(1-q^{n})}{n(1-\delta(1-q)^{n})}\times\frac{q^{n}\big(1-\delta(1-q)^{n}\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)} \\ &= \frac{1}{n}\frac{(1-\delta)q^{n}(1-q^{n})(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)}.\end{aligned} \tag{A.27}$$

Recall that the *ex ante* equilibrium payoff in a separating equilibrium is equal to $\frac{1}{n}\big[(1-q^{n})(\overline{\theta}-\overline{b})+q^{n}(\underline{\theta}-\underline{b})\big]$, we must therefore have:

$$\frac{1}{n}\big[(1-q^{n})(\overline{\theta}-\overline{b}^{*})+q^{n}(\underline{\theta}-\underline{b}^{*})\big]=\frac{1}{n}\frac{(1-\delta)q^{n}(1-q^{n})(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)},$$

which, knowing $\underline{\theta}-\underline{b}^{*}$ from (A.25), we can solve for $\overline{\theta}-\overline{b}^{*}$ to obtain:

$$\begin{aligned}\overline{\theta}-\overline{b}^{*} &= \frac{(1-\delta)q^{n}(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)}-\frac{q^{n}(\underline{\theta}-\underline{b}^{*})}{1-q^{n}} \\ &= \frac{(1-\delta)q^{n}(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)}+\frac{q^{n}\delta q(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)} \\ &= \frac{q^{n}(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^{n}\big)+q^{n}\big(1-\delta(1-q)^{n}\big)},\end{aligned} \tag{A.28}$$

from which we can now compute $\underline{\theta} - \overline{b}^*$:

$$\begin{aligned} \underline{\theta} - \overline{b}^* &= \underline{\theta} - \overline{\theta} + \overline{\theta} - \overline{b}^* = \\ &= \overline{\theta} - \overline{b}^* - (\overline{\theta} - \underline{\theta}) \\ &= \frac{q^n(1-\delta(1-q))(\overline{\theta} - \underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} - (\overline{\theta} - \underline{\theta}) \\ &= \frac{\delta q\big(2q^n - q^{n-1} - 1 + q^{n-1}(1-q)^n\big)(\overline{\theta} - \underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)}. \end{aligned} \tag{A.29}$$

We can now use expression (A.28) to determine $\overline{b}^*$:

$$\overline{b}^* = \overline{\theta} - \frac{q^n(1-\delta(1-q))(\overline{\theta} - \underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)}. \tag{A.30}$$

A.5.2 Solution of Case 2

Recall that in Case 2 the equilibrium conditions are given by:

$$v_{\text{fse}}^* = \frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{b}^*)}{n(1-\delta(1-q)^n)}, \tag{A.31}$$

and:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta} - \overline{b}^*) + \delta v_{\text{fse}}^* = (1-\delta)(\overline{\theta} - \overline{b}^*), \tag{A.32}$$

where $v_{\text{fse}}^* = \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\underline{\theta} - \underline{b}^*)\big]$.

The equilibrium condition in (A.31) implies that:

$$(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\underline{\theta} - \underline{b}) = \frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{b}^*)}{1-\delta(1-q)^n}, \tag{A.33}$$

which can in turn be rewritten as:

$$(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\underline{\theta} - \underline{b}^*) = \frac{(1-\delta)(1-q^n)(\overline{\theta} - \underline{\theta})}{1-\delta(1-q)^n} + \frac{(1-\delta)(1-q^n)(\underline{\theta} - \underline{b})}{1-\delta(1-q)^n}. \tag{A.34}$$

Collecting terms, we get:

$$(1-q^n)(\overline{\theta}-\overline{b}^*)+\left[\frac{q^n-\delta q^n(1-q)^n-(1-\delta)(1-q^n)}{1-\delta(1-q)^n}\right](\underline{\theta}-\underline{b}^*)=\frac{(1-\delta)(1-q^n)(\overline{\theta}-\underline{\theta})}{1-\delta(1-q)^n}. \tag{A.35}$$

Recall that the biding incentive compatibility constraint in (A.32) implies

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\frac{\delta}{n}\big[(1-q^n)(\overline{\theta}-\overline{b}^*)+q^n(\underline{\theta}-\underline{b})\big]=(1-\delta)(\overline{\theta}-\overline{b}^*). \tag{A.36}$$

This condition can be rewritten as:

$$\frac{\delta q^n}{n}(\underline{\theta}-\underline{b}^*)=(1-\delta)(\overline{\theta}-\overline{b}^*)-(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)-\frac{\delta}{n}(1-q^n)(\overline{\theta}-\overline{b}^*) \tag{A.37}$$

$$=(1-\delta)(\overline{\theta}-\overline{b}^*)-\frac{1-q^n}{n}(\overline{\theta}-\overline{b}^*)\left(\frac{1-\delta}{1-q}+\delta\right) \tag{A.38}$$

$$=(\overline{\theta}-\overline{b}^*)\left[(1-\delta)-\frac{1-q^n}{n(1-q)}(1-\delta q)\right] \tag{A.39}$$

$$=\frac{\big[n(1-q)(1-\delta)-(1-q^n)(1-\delta q)\big](\overline{\theta}-\overline{b}^*)}{n(1-q)}. \tag{A.40}$$

Using equations (A.35) and (A.40), the system of equilibrium conditions can now be written as:

$$(1-q^n)(\overline{\theta}-\overline{b}^*)+\left[\frac{q^n-\delta q^n(1-q)^n-(1-\delta)(1-q^n)}{1-\delta(1-q)^n}\right](\underline{\theta}-\underline{b}^*)=\frac{(1-\delta)(1-q^n)(\overline{\theta}-\underline{\theta})}{1-\delta(1-q)^n},$$

$$\delta q^n(\underline{\theta}-\underline{b}^*)=\frac{\big[n(1-q)(1-\delta)-(1-q^n)(1-\delta q)\big](\overline{\theta}-\overline{b}^*)}{1-q},$$

which can be solved for optimal payoffs $\overline{\theta}-\overline{b}^*$ and $\underline{\theta}-\underline{b}^*$:

$$\underline{\theta}-\underline{b}^*=-\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta}), \tag{A.41}$$

$$\overline{\theta}-\overline{b}^*=\frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta}-\underline{\theta}), \tag{A.42}$$

where $D(\delta)$ is given by:

$$D(\delta)=q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]+(1-q^n)\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big].$$

The *ex ante* equilibrium payoff can be found from:

$$nv^*_{\text{fse}} = (1-q^n)(\bar\theta - \bar b^*) + q^n(\underline\theta - \underline b^*) \tag{A.43}$$

$$= \frac{q^n(1-q^n)(\bar\theta-\underline\theta)}{D(\delta)}\big[\delta(1-q^n)(1-q) - (1-q^n)(1-\delta q) + n(1-\delta)(1-q)\big] \tag{A.44}$$

$$= \frac{q^n(1-q^n)(\bar\theta-\underline\theta)}{D(\delta)}\big[(1-q^n)(\delta - \delta q - 1 + \delta q) + n(1-\delta)(1-q)\big] \tag{A.45}$$

$$= \frac{q^n(1-q^n)(\bar\theta-\underline\theta)}{D(\delta)}(1-\delta)\big[-(1-q^n) + n(1-q)\big]. \tag{A.46}$$

$$\tag{A.47}$$

Hence the *ex ante* equilibrium payoff is:

$$v^*_{\text{fse}} = \frac{1}{nD(\delta)}(1-\delta)q^n(1-q^n)\big[n(1-q) - (1-q^n)\big](\bar\theta - \underline\theta). \tag{A.48}$$

We can now determine the payoff of the high type who wins with a low bid, i.e. $\bar\theta - \underline b^*$. Combining the expression for the *ex ante* equilibrium payoff in (A.48) and the equilibrium condition in A.31 we get

$$\frac{(1-\delta)(1-q^n)(\bar\theta - \underline b^*)}{n(1-\delta(1-q)^n)} = \frac{1}{nD(\delta)}(1-\delta)q^n(1-q^n)\big[n(1-q) - (1-q^n)\big](\bar\theta - \underline\theta), \tag{A.49}$$

which can be solved for $\bar\theta - \underline b^*$:

$$\bar\theta - \underline b^* = \frac{1}{D(\delta)}q^n(1-\delta(1-q)^n)\big[n(1-q) - (1-q^n)\big](\bar\theta - \underline\theta). \tag{A.50}$$

A.5.3 Solution of Case 3

Recall that in Case 3 the equilibrium conditions are given by:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\bar\theta - \bar b^*) + \delta v^*_{\text{fse}} = (1-\delta)(\bar\theta - \bar b), \tag{A.51}$$

and:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\bar\theta - \bar b^*) + \delta v^*_{\text{fse}} = (1-\delta)q^{n-1}(\bar\theta - \underline b), \tag{A.52}$$

where $v^*_{\text{fse}} = \frac{1}{n}\big[(1-q^n)(\bar\theta - \bar b) + q^n(\underline\theta - \underline b)\big]$.

Note that conditions (A.51) and (A.52) together imply $\overline{\theta} - \overline{b}^* = q^{n-1}(\overline{\theta} - \underline{b}^*)$. Hence the equilibrium payoff becomes:

$$v^*_{\text{fse}} = \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\underline{\theta} - \underline{b})\big] \tag{A.53}$$
$$= \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\overline{\theta} - \overline{\theta} + \underline{\theta} - \underline{b}^*)\big] \tag{A.54}$$
$$= \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q^n(\overline{\theta} - \underline{b}^*) - q^n(\overline{\theta} - \underline{\theta})\big] \tag{A.55}$$
$$= \frac{1}{n}\big[(1-q^n)(\overline{\theta} - \overline{b}^*) + q(\overline{\theta} - \overline{b}^*) - q^n(\overline{\theta} - \underline{\theta})\big] \tag{A.56}$$
$$= \frac{1}{n}\big[(1-q^n+q)(\overline{\theta} - \overline{b}^*) - q^n(\overline{\theta} - \underline{\theta})\big]. \tag{A.57}$$

The upward incentive compatibility constraint in (A.51) can then be written as:

$$(1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta} - \overline{b}^*) + \delta\frac{1}{n}\big[(1-q^n+q)(\overline{\theta} - \overline{b}^*) - q^n(\overline{\theta} - \underline{\theta})\big] = (1-\delta)(\overline{\theta} - \overline{b}^*). \tag{A.58}$$

which can then be solved for $\overline{\theta} - \overline{b}^*$:

$$\overline{\theta} - \overline{b}^* = \frac{\delta q^n(1-q)(\overline{\theta} - \underline{\theta})}{(1-q^n)(1-\delta q) + \delta q(1-q) - n(1-\delta)(1-q)}. \tag{A.59}$$

We can now introduce shorthand notation for the denominator:

$$D(\delta) = (1-q^n)(1-\delta q) + \delta q(1-q) - n(1-\delta)(1-q). \tag{A.60}$$

The *ex ante* equilibrium payoff can now be calculated from (A.57):

$$
\begin{aligned}
nv^*_{\text{fse}} &= (1-q^n+q)(\overline{\theta}-\overline{b}^*)-q^n(\overline{\theta}-\underline{\theta}) && \text{(A.61)}\\
&= (1-q^n+q)\frac{\delta q^n(1-q)(\overline{\theta}-\underline{\theta})}{(1-q^n)(1-\delta q)+\delta q(1-q)-n(1-\delta)(1-q)}-q^n(\overline{\theta}-\underline{\theta})\\
&= \frac{q^n(\overline{\theta}-\underline{\theta})}{D(\delta)}\big[(1-q^n+q)\delta(1-q)-(1-q^n)(1-\delta q)-\delta q(1-q)+n(1-\delta)(1-q)\big]\\
&= \frac{q^n(\overline{\theta}-\underline{\theta})}{D(\delta)}\big[(1-q^n)\big(\delta(1-q)-(1-\delta q)\big)+n(1-\delta)(1-q)\big]\\
&= \frac{(1-\delta)q^n(\overline{\theta}-\underline{\theta})}{D(\delta)}\big[n(1-q)-(1-q^n)\big]\\
&= \frac{1}{D(\delta)}(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}).
\end{aligned}
$$

The *ex ante* equilibrium payoff is then given by:

$$
v^*_{\text{fse}} = \frac{1}{nD(\delta)}(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}). \qquad \text{(A.62)}
$$

The payoff of a high type buyer who wins with the low bid can be calculated from A.59 and the fact that $\overline{\theta}-\underline{b}^* = \frac{1}{q^{n-1}}(\overline{\theta}-\overline{b}^*)$, and is therefore given by:

$$
\begin{aligned}
\overline{\theta}-\underline{b}^* &= \frac{\delta q(1-q)(\overline{\theta}-\underline{\theta})}{(1-q^n)(1-\delta q)+\delta q(1-q)-n(1-\delta)(1-q)} && \text{(A.63)}\\
&= \frac{1}{D(\delta)}\delta q(1-q)(\overline{\theta}-\underline{\theta}). && \text{(A.64)}
\end{aligned}
$$

A low type buyer payoff can be calculated from $nv^*_i = (1-q^n)(\overline{\theta}-\overline{b})+q^n(\underline{\theta}-\underline{b}^*)$:

$$
\begin{aligned}
q^n(\underline{\theta}-\underline{b}^*) &= nv^*_{\text{fse}}-(1-q^n)(\overline{\theta}-\overline{b}^*) && \text{(A.65)}\\
&= \frac{1}{D(\delta)}(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})-(1-q^n)\frac{1}{D(\delta)}\delta q^n(1-q)(\overline{\theta}-\underline{\theta}),
\end{aligned}
$$

which implies:

$$
\begin{aligned}
\underline{\theta}-\underline{b}^* &= \frac{1}{D(\delta)}\Big[(1-\delta)\big[n(1-q)-(1-q^n)\big]-(1-q^n)\delta(1-q)\Big](\overline{\theta}-\underline{\theta}) && \text{(A.66)}\\
&= \frac{1}{D(\delta)}\big[n(1-q)(1-\delta)-(1-q^n)(1-\delta q)\big](\overline{\theta}-\underline{\theta}) && \text{(A.67)}
\end{aligned}
$$

A.6 Proof of Lemma 1.5

Proof. We have shown $\underline{\theta} < \underline{b}^*$ in the main text. To show $\underline{b}^* < \overline{b}^*$, consider the payoffs defined by (1.20) and (1.21). It suffices to show that $\overline{\theta} - \underline{b}^* > \overline{\theta} - \overline{b}^*$, which is equivalent to:

$$\frac{1}{D(\delta)} q^n \big(1 - \delta(1-q)^n\big)\big[n(1-q) - (1-q^n)\big](\overline{\theta} - \underline{\theta}) > \frac{1}{D(\delta)} \delta q^n (1-q^n)(1-q)(\overline{\theta} - \underline{\theta}) \tag{A.68}$$

$$\Leftrightarrow \big(1 - \delta(1-q)^n\big)\big[n(1-q) - (1-q^n)\big] > \delta(1-q^n)(1-q). \tag{A.69}$$

It is easy to see that the above inequality holds for all δ whenever:

$$\big(1 - (1-q)^n\big)\big[n(1-q) - (1-q^n)\big] > (1-q^n)(1-q) \tag{A.70}$$

since the left-hand side of the inequality is decreasing in δ and the right-hand side of the inequality is increasing in δ.

Recall now that we assume that $q \geq \frac{1-q^n}{n(1-q)}$ which is equivalent to:

$$n(1-q)^2 q \geq (1-q^n)(1-q), \tag{A.71}$$

and in particular implies that $n \geq 4$

I now show that A.71 implies A.70 by showing that:

$$\big(1 - (1-q)^n\big)\big[n(1-q) - (1-q^n)\big] > n(1-q)^2 q \tag{A.72}$$

for $n \geq 4$.

Observe first that $\big(1-(1-q)^n\big) > \big(1-(1-q)^2\big) = q(2-q)$ for $n \geq 4$. Since $n(1-q)-(1-q^n)$ is strictly positive it suffices to show that:

$$q(2-q)\big[n(1-q) - (1-q^n)\big] > n(1-q)^2 q, \tag{A.73}$$

which is equivalent to:

$$
\begin{aligned}
&(2-q)\big[n(1-q)-(1-q^n)\big] > n(1-q)^2 \qquad (A.74)\\
&(2-q)n(1-q) - n(1-q)^2 > (2-q)(1-q^n)\\
&n(1-q)(2-q-1+q) > (2-q)(1-q^n)\\
&n(1-q) > (2-q)(1-q^n)\\
&n(1-q) > (2-q)(1-q)\sum_{k=0}^{n-1} q^k\\
&n > (2-q)\sum_{k=0}^{n-1} q^k = (1-q)\sum_{k=0}^{n-1} q^k + \sum_{k=0}^{n-1} q^k = 1-q^n+\sum_{k=0}^{n-1} q^k.
\end{aligned}
$$

Consider the function $f(q) = 1 - q^n + \sum_{k=0}^{n-1} q^k$. Differentiating $f(q)$ with respect to q I get:

$$
\begin{aligned}
f'(q) &= -nq^{n-1} + \sum_{k=1}^{n-1} kq^{k-1} > -nq^{n-1} + \sum_{k=1}^{n-1} kq^{n-1}\\
&= -nq^{n-1} + q^{n-1}\sum_{k=1}^{n-1} k = q^{n-1}\left[\sum_{k=1}^{n-1} k - n\right]\\
&= q^{n-1}\left[\frac{(1+n-1)(n-1)}{2} - n\right]\\
&= q^{n-1}n\left[\frac{(n-1)}{2} - 1\right] = q^{n-1}n\frac{(n-3)}{2} > 0,
\end{aligned}
$$

where the last inequality is true since $n \geq 4$ by assumption.

Hence we can conclude that $f(q)$ is strictly increasing on $(0,1)$. Computing $f(1)$ we obtain:

$$f(1) = 1 - 1^n + \sum_{k=0}^{n-1} 1^k = n, \qquad (A.75)$$

therefore $f(q) < n$ for all $q \in (0,1)$. □

A.7 Proofs of Propositions 1.6, 1.8, 1.12 (Full-surplus-extracting cPPE)

A.7.1 Proof of Proposition 1.6

Proof. Full surplus extraction and $\mathcal{R}^* \geq (1-q^n)\overline{\theta}$ are shown in the main text, hence by Lemma 1.2 it remains to check the incentive constraints and the no-collusion constraints. I start by checking incentive compatibility.

(I) *On-schedule incentive compatibility of the buyers*

Consider a high-type buyer, his equilibrium payoff must be higher than the payoff he could obtain by mimicking the behavior of a low type buyer:

$$\text{(HighIC-on-sch)} \qquad \frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*) \geq \frac{q^{n-1}}{n}(\overline{\theta}-\underline{b}^*). \tag{A.76}$$

Plugging the respective payoffs in, we obtain:

$$\frac{1-q^n}{1-q}\frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)} \geq q^{n-1}\frac{q^n\big(1-\delta(1-q)^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}, \tag{A.77}$$

which simplifies to:

$$\frac{1-q^n}{1-q}(1-\delta(1-q)) \geq q^{n-1}\big(1-\delta(1-q)^n\big) \tag{A.78}$$

$$\Leftrightarrow \frac{1-q^n}{1-q} - \delta(1-q^n) \geq q^{n-1} - q^{n-1}\delta(1-q)^n \tag{A.79}$$

$$\Leftrightarrow \frac{(1-q)\sum_{k=0}^{n-1} q^k}{1-q} - \delta(1-q^n) \geq q^{n-1} - q^{n-1}\delta(1-q)^n \tag{A.80}$$

$$\Leftrightarrow \sum_{k=0}^{n-2} q^k \geq \delta(1-q^n) - q^{n-1}\delta(1-q)^n \tag{A.81}$$

$$\Leftrightarrow \frac{1}{\delta}\sum_{k=0}^{n-2} q^k \geq (1-q^n) - q^{n-1}(1-q)^n. \tag{A.82}$$

Since $\frac{1}{\delta}\sum_{k=0}^{n-2} q^k > \sum_{k=0}^{n-2} q^k$, it is enough to show that:

$$\sum_{k=0}^{n-2} q^k \geq (1-q^n) - q^{n-1}(1-q)^n \tag{A.83}$$

$$\Leftrightarrow 1 + \sum_{k=1}^{n-2} q^k \geq 1 - q^n - q^{n-1}(1-q)^n \tag{A.84}$$

$$\Leftrightarrow \sum_{k=1}^{n-2} q^k + q^n \geq -q^{n-1}(1-q)^n, \tag{A.85}$$

which is clearly true since the left-hand side of the above inequality in (A.85) is strictly positive, and the right-hand side is strictly negative.

I now turn to off-schedule incentive compatibility for both types.

(II) *Off-schedule incentive compatibility of the buyers*

Consider first a high-type buyer who deviates to $\underline{b}^* + \epsilon$. The associated incentive compatibility constraint is given by:

$$\text{(HighIC-down)} \qquad (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta} - \overline{b}^*) + \delta v^*_{\text{fse}} \geq (1-\delta)q^{n-1}(\overline{\theta} - \underline{b}^*). \tag{A.86}$$

Plugging the respective payoffs in, we obtain:

$$\frac{(1-\delta)(1-q^n)}{n(1-q)} \frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} + \frac{\delta}{n}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}$$
$$\geq (1-\delta)q^{n-1}\frac{q^n\big(1-\delta(1-q)^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)},$$

which simplifies to:

$$\frac{1-q^n}{n(1-q)}(1-\delta(1-q))+\frac{\delta}{n}(1-q^n)\geq q^{n-1}\big(1-\delta(1-q)^n\big) \tag{A.87}$$
$$\Leftrightarrow \frac{1-q^n}{1-q}-\delta(1-q^n)+\delta(1-q^n)\geq nq^{n-1}\big(1-\delta(1-q)^n\big) \tag{A.88}$$
$$\Leftrightarrow \frac{1-q^n}{1-q}\geq nq^{n-1}\big(1-\delta(1-q)^n\big) \tag{A.89}$$
$$\Leftrightarrow \frac{1-q^n}{1-q}-nq^{n-1}\geq -nq^{n-1}\delta(1-q)^n \tag{A.90}$$
$$\Leftrightarrow \sum_{k=0}^{n-1}q^k-nq^{n-1}\geq -nq^{n-1}\delta(1-q)^n. \tag{A.91}$$

which is true since the left-hand side of (A.91) is strictly positive and the right-hand side of (A.91) is strictly negative.

Now consider a high-type buyer who deviates to $\overline{b}^*+\epsilon$. The associated incentive compatibility constraint is given by:

$$\text{(HighIC-up)} \qquad (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}}\geq(1-\delta)(\overline{\theta}-\overline{b}^*), \tag{A.92}$$

which is equivalent to:

$$\delta v^*_{\text{fse}}\geq(1-\delta)\left(1-\frac{1-q^n}{n(1-q)}\right)(\overline{\theta}-\overline{b}^*). \tag{A.93}$$

Plugging the respective payoffs in, we get:

$$\frac{\delta}{n}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)}\geq(1-\delta)\left(1-\frac{1-q^n}{n(1-q)}\right)\frac{q^n(1-\delta(1-q))(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)},$$

which is equivalent to:

$$\frac{\delta}{n}(1-q^n) \geq \left(1 - \frac{1-q^n}{n(1-q)}\right)(1-\delta(1-q)) \tag{A.94}$$

$$\Leftrightarrow \frac{\delta}{n}(1-q^n) \geq 1-\delta(1-q) - \frac{1-q^n}{n(1-q)} + \frac{\delta}{n}(1-q^n) \tag{A.95}$$

$$\Leftrightarrow 0 \geq 1-\delta(1-q) - \frac{1-q^n}{n(1-q)} \tag{A.96}$$

$$\Leftrightarrow \delta \geq \frac{1}{1-q} - \frac{1-q^n}{n(1-q)^2}. \tag{A.97}$$

The condition on δ identified in (A.97) can only be satisfied if:

$$\frac{1}{1-q} - \frac{1-q^n}{n(1-q)^2} < 1 \tag{A.98}$$

$$\Leftrightarrow 1 - \frac{1-q^n}{n(1-q)} < 1-q \Leftrightarrow nq < \frac{1-q^n}{1-q}, \tag{A.99}$$

which is true by assumption.

(III) *No-collusion constraints*

Suppose $\overline{\theta}$ bids off schedule and $\underline{\theta}$ bids on schedule. The associated constraint is:

$$\text{(No-col-sep-2)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^* + \epsilon, \underline{b}^*) = \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\underline{b}^*) + q^n(\underline{\theta}-\underline{b}^*)\big]}{n(1-\delta q^n)}. \tag{A.100}$$

Computing $(1-q^n)(\overline{\theta}-\underline{b}^*) + q^n(\underline{\theta}-\underline{b}^*)$, we get

$$(1-q^n)(\overline{\theta}-\underline{b}^*) + q^n(\underline{\theta}-\underline{b}^*) = \tag{A.101}$$

$$= (1-q^n)\frac{q^n\big(1-\delta(1-q)^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)} + q^n\frac{-\delta q\big(1-q^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)}$$

$$= \frac{q^n(1-q^n)\big(1-\delta(1-q)^n - \delta q\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)}, \tag{A.102}$$

which then implies that the payoff from this bidding profile is equal to:

$$v'(\underline{b}^* + \epsilon, \underline{b}^*) = \frac{(1-\delta)q^n(1-q^n)\big(1-\delta(1-q)^n - \delta q\big)(\overline{\theta}-\underline{\theta})}{n(1-\delta q^n)\big(\delta q\big(1-q^n\big) + q^n\big(1-\delta(1-q)^n\big)\big)}. \tag{A.103}$$

We need to establish that $v^*_{\text{fse}} \geq v'(\underline{b}^* + \epsilon, \underline{b}^*)$. i.e.

$$\frac{1}{n}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)} \geq \frac{(1-\delta)q^n(1-q^n)\big(1-\delta(1-q)^n-\delta q\big)(\overline{\theta}-\underline{\theta})}{n(1-\delta q^n)\big(\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)\big)},$$

which simplifies to:

$$1 \geq \frac{1-\delta(1-q)^n-\delta q}{1-\delta q^n} \tag{A.104}$$

$$\Leftrightarrow 1-\delta q^n \geq 1-\delta(1-q)^n-\delta q \tag{A.105}$$

$$\Leftrightarrow -\delta q^n \geq -\delta(1-q)^n-\delta q \tag{A.106}$$

$$\Leftrightarrow -q^n \geq -(1-q)^n-q \tag{A.107}$$

$$\Leftrightarrow (1-q)^n \geq -q+q^n, \tag{A.108}$$

which is true since the right-hand side of A.108 is strictly negative, and the left-hand side is strictly positive.

Suppose both types pool at $\underline{b}^*$. The associated constraint is:

$$\text{(No-col-pol)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^*, \underline{b}^*) = \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big], \tag{A.109}$$

where

$$\begin{aligned} v(\underline{b}^*, \underline{b}^*) &= \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big] = \\ &= \frac{(1-q)}{n}\frac{q^n\big(1-\delta(1-q)^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)} + \frac{q}{n}\frac{-\delta q\big(1-q^n\big)(\overline{\theta}-\underline{\theta})}{\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)} \\ &= \frac{\big((1-q)q^n\big(1-\delta(1-q)^n\big)-\delta q^2(1-q^n)\big)(\overline{\theta}-\underline{\theta})}{n\big(\delta q\big(1-q^n\big)+q^n\big(1-\delta(1-q)^n\big)\big)}. \end{aligned} \tag{A.110}$$

Consider the numerator of (A.110) in the limit as δ goes to 1:

$$(1-q)q^n\big(1-(1-q)^n\big) - q^2(1-q^n) \tag{A.111}$$

$$= (1-q)\left[q^n\big(1-(1-q)^n\big) - q^2\sum_{k=0}^{n-1} q^k\right] \tag{A.112}$$

$$= (1-q)\left[q^n - q^n(1-q)^n - q^2\sum_{k=0}^{n-3} q^k - q^n - q^{n+1}\right] \tag{A.113}$$

$$= (1-q)\left[-q^n(1-q)^n - q^2\sum_{k=0}^{n-3} q^k - q^{n+1}\right] < 0. \tag{A.114}$$

Recall that v^*_{fse} is weakly positive, whereas the payoff in A.110 goes to a negative value. By continuity there is a δ^* in the neighborhood of 1 such that for all $\delta > \delta^*$ the equilibrium payoff v^*_{fse} exceeds the payoff in A.110.

□

A.7.2 Proof of Proposition 1.8

Proof. Full surplus extraction and $\mathcal{R}^* \geq (1-q^n)\overline{\theta}$ are shown in the main text, hence by Lemma 1.2 it remains to check the incentive constraints and the no-collusion constraints. I start by checking incentive compatibility.

(I) *On-schedule incentive compatibility of the buyers*

Consider a high-type buyer on-schedule incentive compatibility condition:

$$\text{(HighIC-on-sch)} \qquad \frac{1-q^n}{n(1-q)}(\overline{\theta} - \overline{b}^*) \geq \frac{q^{n-1}}{n}(\overline{\theta} - \underline{b}^*). \tag{A.115}$$

Plugging the payoffs defined in (1.20) and (1.21) I get:

$$\frac{1-q^n}{n(1-q)}\frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta}-\underline{\theta}) \tag{A.116}$$
$$\geq \frac{q^{n-1}}{n}\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}),$$

which is equivalent to:

$$\delta\big(1-q^n)\big)^2 \geq q^{n-1}\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]. \tag{A.117}$$

which is in particular true whenever

$$\delta(1-q^n)(1-q^n) \geq q^{n-1}\big[n(1-q)-(1-q^n)\big], \tag{A.118}$$

i.e. for all δ satisfying $\delta \geq \frac{q^{n-1}\left[n(1-q)-(1-q^n)\right]}{(1-q^n)(1-q^n)}$. Note that such δ exist in (0,1) since

$$\begin{aligned}
&(1-q^n)(1-q^n) > q^{n-1}\big[n(1-q)-(1-q^n)\big] \qquad \text{(A.119)}\\
&\Leftrightarrow (1-q^n)(1-q^n)+q^{n-1}(1-q^n) > nq^{n-1}(1-q)\\
&\Leftrightarrow (1-q^n)(1-q^n+q^{n-1}) > nq^{n-1}(1-q)\\
&\Leftrightarrow (1-q^n)(1+q^{n-1}(1-q)) > nq^{n-1}(1-q)\\
&\Leftrightarrow (1+q^{n-1}(1-q))\sum_{k=0}^{n-1} q^k > nq^{n-1},
\end{aligned}$$

where the last inequality is true since $\sum_{k=0}^{n-1} q^k > nq^{n-1}$ and $1+q^{n-1}(1-q) > 1$. Thus the high type on-schedule incentive compatibility constraint is satisfied for a sufficiently high δ.

(II) *Off-schedule incentive compatibility of the buyers*

Let us now turn to the off-schedule incentive compatibility constraints of the buyers. Consider first a low-type buyer. He must be willing to participate in the bidding with the bid $\underline{b}^*$ as opposed to abstaining and getting a zero payoff:

$$\text{(LowIC)} \qquad (1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}^*)+\delta v^*_{\text{fse}} \geq 0. \tag{A.120}$$

Plugging the payoffs defined in (1.19) and (1.22), I obtain:

$$\begin{aligned}
&-(1-\delta)\frac{q^{n-1}}{n}\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta}) \qquad \text{(A.121)}\\
&\qquad +\delta\frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \geq 0,
\end{aligned}$$

which simplifies to:

$$-\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]+\delta q\big[n(1-q)-(1-q^n)\big]\geq 0 \tag{A.122}$$
$$\Leftrightarrow n(1-\delta)(1-q)-(1-q^n)(1-\delta q)+\delta qn(1-q)-\delta q(1-q^n)\geq 0$$
$$\Leftrightarrow n(1-\delta+\delta q)(1-q)-(1-q^n)\geq 0$$
$$\Leftrightarrow 1-\delta+\delta q\geq \frac{1-q^n}{n(1-q)}\Leftrightarrow \delta\leq \frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2},$$

which is true since $\frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2}\geq 1$ by assumption that $q\geq \frac{1-q^n}{n(1-q)}$.

Consider a high type buyer who attempts a downward deviation to $\underline{b}^*+\epsilon$. The associated incentive compatibility condition is given by:

$$\text{(HighIC-down)} \qquad (1-\delta)\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*)+\delta v^*_{\text{fse}}\geq (1-\delta)q^{n-1}(\overline{\theta}-\underline{b}^*). \tag{A.123}$$

Plugging the payoffs defined in (1.20), (1.21), and (1.22) into the above inequality, I obtain:

$$\begin{aligned}&(1-\delta)\frac{1-q^n}{n(1-q)}\frac{1}{D(\delta)}\delta q^n(1-q^n)(1-q)(\overline{\theta}-\underline{\theta})\\&\quad+\delta\frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})\\&\quad\geq (1-\delta)q^{n-1}\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}),\end{aligned} \tag{A.124}$$

which simplifies to:

$$\begin{aligned}&(1-q^n)\frac{1}{n}\delta(1-q^n)+\delta\frac{1}{n}(1-q^n)\big[n(1-q)-(1-q^n)\big]\\&\qquad\geq q^{n-1}\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big],\end{aligned} \tag{A.125}$$

which can be further simplified to:

$$\delta(1-q^n)(1-q)\geq q^{n-1}\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big], \tag{A.126}$$

i.e. for all discount factors δ such that:

$$\delta \geq \frac{q^{n-1}\big[n(1-q)-(1-q^n)\big]}{(1-q^n)(1-q)+q^{n-1}(1-q)^n\big[n(1-q)-(1-q^n)\big]} \tag{A.127}$$

which can only be satisfied when:

$$\frac{q^{n-1}\big[n(1-q)-(1-q^n)\big]}{(1-q^n)(1-q)+q^{n-1}(1-q)^n\big[n(1-q)-(1-q^n)\big]} < 1, \tag{A.128}$$

or:

$$(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big], \tag{A.129}$$

which is true by assumption.

(III) *No-collusion constraints*

Suppose $\overline{\theta}$ bids off schedule and $\underline{\theta}$ bids on schedule. The associated no-collusion constraint is given by:

$$\text{(No-col-sep-2)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^*+\epsilon, \underline{b}^*) = \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\overline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big]}{n(1-\delta q^n)}. \tag{A.130}$$

Plugging the payoffs in, we can rewrite the right-hand side as:

$$\begin{aligned}\frac{(1-\delta)}{n(1-\delta q^n)}\Bigg[&(1-q^n)\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \\ &-q^n\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta})\Bigg],\end{aligned} \tag{A.131}$$

which simplifies to:

$$\begin{aligned}\frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{n(1-\delta q^n)D(\delta)}\Bigg(&\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \\ &-\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Bigg).\end{aligned} \tag{A.132}$$

We now have to make sure that it is below v^*_{fse}, i.e.

$$\begin{aligned}\frac{1-\delta}{nD(\delta)}q^n(1-q^n)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \geq \\ \geq \frac{(1-\delta)q^n(1-q^n)(\overline{\theta}-\underline{\theta})}{n(1-\delta q^n)D(\delta)}\bigg(\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \\ -\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\bigg),\end{aligned} \tag{A.133}$$

which is equivalent to:

$$\begin{aligned}(1-\delta q^n)\big[n(1-q)-(1-q^n)\big] \geq \big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \\ -\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big],\end{aligned} \tag{A.134}$$

which in turn simplifies to:

$$\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big] \geq \delta\big(q^n-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big], \tag{A.135}$$

which can only be satisfied if:

$$\delta \geq \frac{n(1-q)-(1-q^n)}{n(1-q)-q(1-q^n)-\big(q^n-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]}, \tag{A.136}$$

which in turn can only be satisfied for a high enough $\delta \in (0,1)$ only if:

$$(1-q^n)(1-q) > \big(q^n-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]. \tag{A.137}$$

It is easy to show that the above inequality is implied by the parameter restriction of Case 2. Recall that the parameter restriction is given by:

$$(1-q^n)(1-q) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]. \tag{A.138}$$

Observe that $q^n-(1-q)^n < q^{n-1}\big(1-(1-q)^n\big)$, which establishes the result.

Suppose now the buyers pool at $\underline{b}^*$, the associated no-collusion constraint is:

$$\text{(No-col-pool)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^*,\underline{b}^*) = \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big], \tag{A.139}$$

where

$$\begin{aligned}
v(\underline{b}^*, \underline{b}^*) &= \frac{1}{n}\big[(1-q)(\overline{\theta} - \underline{b}^*) + q(\underline{\theta} - \underline{b}^*)\big] \\
&= \frac{1}{n}\bigg[(1-q)\frac{1}{D(\delta)}q^n\big(1-\delta(1-q)^n\big)\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \\
&- q\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](1-q^n)(\overline{\theta}-\underline{\theta})\bigg].
\end{aligned} \tag{A.140}$$

I show that $v(\underline{b}^*, \underline{b}^*)$ converges to a strictly negative number as δ goes to 1. Indeed in the limit $v(\underline{b}^*, \underline{b}^*)$ is given by:

$$\frac{(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\big[q^n\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] - q(1-q^n)(1-q^n)\big]. \tag{A.141}$$

I now verify that:

$$\begin{aligned}
&\frac{(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\big[q^n\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] - q(1-q^n)(1-q^n)\big] < 0 \\
&\Leftrightarrow q^n\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] - q(1-q^n)(1-q^n) < 0 \\
&\Leftrightarrow (1-q^n)(1-q^n) > q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big].
\end{aligned}$$

It suffices to show that $(1-q^n)(1-q^n) > q^{n-1}\big[n(1-q)-(1-q^n)\big]$, which has already been established above. I repeat the argument here for completeness:

$$\begin{aligned}
&(1-q^n)(1-q^n) > q^{n-1}\big[n(1-q)-(1-q^n)\big] \\
&\Leftrightarrow (1-q^n)(1-q^n) + q^{n-1}(1-q^n) > nq^{n-1}(1-q) \\
&\Leftrightarrow (1-q^n)(1-q^n+q^{n-1}) > nq^{n-1}(1-q) \\
&\Leftrightarrow (1-q^n)(1+q^{n-1}(1-q)) > nq^{n-1}(1-q) \\
&\Leftrightarrow (1+q^{n-1}(1-q))\sum_{k=0}^{n-1} q^k > nq^{n-1},
\end{aligned}$$

where the last inequality is true since $\sum_{k=0}^{n-1} q^k > nq^{n-1}$ and $1+q^{n-1}(1-q) > 1$.

$\square$

A.7.3 Proof of Proposition 1.12

Proof. Full surplus extraction and $\mathcal{R}^* \geq (1-q^n)\overline{\theta}$ are shown in the main text, hence by Lemma 1.2 it remains to check the incentive constraints and the no-collusion constraints. Let us now check on-schedule incentive compatibility.

(I) *On-schedule incentive compatibility of the buyers*

Consider a high type buyer who contemplates an on-schedule deviation. The associated on-schedule incentive compatibility condition is given by:

$$\text{(HighIC-on-sch)} \qquad \frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*) \geq \frac{q^{n-1}}{n}(\overline{\theta}-\underline{b}^*).$$

Note that both $\overline{\theta}-\overline{b}^*$ and $\overline{\theta}-\underline{b}^*$ are strictly positive for δ high enough. Recall that by construction of this public perfect equilibrium $\overline{\theta}-\overline{b}^* = q^{n-1}(\overline{\theta}-\underline{b}^*)$ and therefore we obtain:

$$\frac{1-q^n}{n(1-q)}(\overline{\theta}-\overline{b}^*) > \frac{1}{n}(\overline{\theta}-\overline{b}^*) = \frac{q^{n-1}}{n}(\overline{\theta}-\underline{b}^*).$$

where the first inequality is true since $1-q^n > 1-q$ for $n \geq 2$ and $q \in (0,1)$, implying that the high-type on-schedule incentive compatibility is satisfied.

(II) *Off-schedule incentive compatibility of the buyers*

Having dealt with the on-schedule incentive compatibility constraint of the buyers, I now establish that the off-schedule incentive compatibility constraints of the buyers are satisfied. Consider first a low type buyer. A low type buyer must prefer participating in the auction with the bid $\underline{b}^*$ as opposed to abstaining and getting zero forever:

$$\text{(LowIC)} \qquad (1-\delta)\frac{q^{n-1}}{n}(\underline{\theta}-\underline{b}^*) + \delta v^*_{\text{fse}} \geq 0. \tag{A.142}$$

Plugging the payoffs from (1.26) and (1.29) I get:

$$-(1-\delta)\frac{q^{n-1}}{n}\frac{1}{D(\delta)}\big[(1-q^n)(1-\delta q) - n(1-\delta)(1-q)\big](\overline{\theta}-\underline{\theta}) \tag{A.143}$$
$$+\,\delta\frac{1}{nD(\delta)}(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \geq 0,$$

which is equivalent to:

$$\begin{aligned}
&-\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]+\delta q\big[n(1-q)-(1-q^n)\big]\geq 0 \qquad \text{(A.144)}\\
&\Leftrightarrow -(1-q^n)(1-\delta q)+n(1-\delta)(1-q)+\delta qn(1-q)-\delta q(1-q^n)\geq 0\\
&\Leftrightarrow n(1-q)(1-\delta+\delta q)-(1-q^n)\geq 0\\
&\Leftrightarrow 1-\delta+\delta q\geq \frac{1-q^n}{n(1-q)}\\
&\Leftrightarrow 1-\frac{1-q^n}{n(1-q)}\geq \delta-\delta q \Leftrightarrow \delta\leq \frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2},
\end{aligned}$$

which is true since $\frac{1}{1-q}-\frac{1-q^n}{n(1-q)^2}\geq 1$ by assumption that $q\geq \frac{1-q^n}{n(1-q)}$.

(III) *No-collusion constraints*

Suppose $\overline{\theta}$ bids on schedule and $\underline{\theta}$ bids off schedule. The associated no-collusion constraint is is given by:

$$\text{(No-col-sep-1)} \qquad v^*_{\text{fse}}\geq v(\underline{b}^*,\emptyset)=\frac{(1-\delta)(1-q^n)(\overline{\theta}-\underline{b}^*)}{n\big(1-\delta(1-q)^n\big)}. \qquad \text{(A.145)}$$

Recall the formula for $\overline{\theta}-\underline{b}^*$ in (1.28), plugging it into the payoff formula above, I get

$$v'(\underline{b}^*,\emptyset)=\frac{(1-\delta)(1-q^n)\delta q(1-q)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta(1-q)^n\big)}. \qquad \text{(A.146)}$$

The goal is to show that for δ sufficiently high $v^*_{\text{fse}}\geq v'(\underline{b}^*,\emptyset)$, i.e.

$$\frac{1}{nD(\delta)}(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta})\geq \frac{(1-\delta)(1-q^n)\delta q(1-q)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta(1-q)^n\big)}, \qquad \text{(A.147)}$$

which is equivalent to:

$$\begin{aligned}
&q^{n-1}\big[n(1-q)-(1-q^n)\big]\geq \frac{(1-q^n)\delta(1-q)}{\big(1-\delta(1-q)^n\big)} \qquad \text{(A.148)}\\
&\Leftrightarrow \big(1-\delta(1-q)^n\big)q^{n-1}\big[n(1-q)-(1-q^n)\big]\geq \delta(1-q^n)(1-q),
\end{aligned}$$

which can be satisfied for any $\delta \in (0,1)$ as long as it is true that[3]

$$\big(1-(1-q)^n\big)q^{n-1}\big[n(1-q)-(1-q^n)\big] \geq (1-q^n)(1-q),$$

which is assumed is Case 3.

Suppose $\overline{\theta}$ bids off schedule and $\underline{\theta}$ bids on schedule. The associated no-collusion constraint is

$$\text{(No-col-sep-2)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^*+\epsilon,\underline{b}^*) = \frac{(1-\delta)\big[(1-q^n)(\overline{\theta}-\underline{b}^*)+q^n(\underline{\theta}-\underline{b}^*)\big]}{n\big(1-\delta q^n\big)}. \tag{A.149}$$

Recall the formulas for $\underline{\theta}-\underline{b}^*$ and $\overline{\theta}-\underline{b}^*$ in (1.26) and (1.28) respectively, the above payoff can then be written as

$$v(\underline{b}^*+\epsilon,\underline{b}^*) = \frac{(1-\delta)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta q^n\big)}\Big[(1-q^n)\delta q(1-q)-q^n\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Big].$$

Our goal is to show that $v^*_{\text{fse}} \geq v'(\underline{b}^*+\epsilon,\underline{b}^*)$, i.e.

$$\begin{aligned}\frac{1}{nD(\delta)}&(1-\delta)q^n\big[n(1-q)-(1-q^n)\big](\overline{\theta}-\underline{\theta}) \\ &\geq \frac{(1-\delta)(\overline{\theta}-\underline{\theta})}{nD(\delta)\big(1-\delta q^n\big)}\Big[(1-q^n)\delta q(1-q)-q^n\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Big],\end{aligned} \tag{A.150}$$

which is equivalent to:

$$\begin{aligned}q^n&\big[n(1-q)-(1-q^n)\big] \\ &\geq \frac{1}{\big(1-\delta q^n\big)}\Big[(1-q^n)\delta q(1-q)-q^n\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Big],\end{aligned} \tag{A.151}$$

which holds for δ sufficiently close to 1 whenever it holds as a strict inequality at

[3] Note that it is required that δ satisfy

$$\delta \leq \frac{q^{n-1}\big[n(1-q)-(1-q^n)\big]}{(1-q^n)(1-q)+q^{n-1}(1-q)^n\big[n(1-q)-(1-q^n)\big]}.$$

The restriction on the parameters assumed in Case 3 makes sure that the right-hand side of this inequality is weakly above 1.

$\delta = 1$, i.e. whenever

$$q^n\big[n(1-q)-(1-q^n)\big] > \frac{1}{(1-q^n)}\big[(1-q^n)q(1-q)-q^n(1-q^n)(1-q)\big] \quad \text{(A.152)}$$
$$\Leftrightarrow q^n\big[n(1-q)-(1-q^n)\big] > q(1-q)-q^n(1-q)$$
$$\Leftrightarrow q^{n-1}\big[n(1-q)-(1-q^n)\big] > (1-q)(1-q^{n-1}).$$

Now the last line is true since:

$$(1-q)(1-q^{n-1}) < (1-q)(1-q^n) \leq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big].$$

where the first inequality is evidently true, and the second inequality holds true in Case 3 by assumption. The result follows by the fact that $q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] < q^{n-1}\big[n(1-q)-(1-q^n)\big]$, which in turn is true because $1-(1-q)^n < 1$.

Suppose both types pool at $\underline{b}^*$. The associated no-collusion constraint is:

$$\text{(No-col-pool)} \qquad v^*_{\text{fse}} \geq v(\underline{b}^*, \underline{b}^*) = \frac{1}{n}\big[(1-q)(\overline{\theta}-\underline{b}^*)+q(\underline{\theta}-\underline{b}^*)\big]. \quad \text{(A.153)}$$

Recall again the formulas for $\underline{\theta}-\underline{b}^*$ and $\overline{\theta}-\underline{b}^*$ in (1.26) and (1.28) respectively, the pooling payoff is then given by:

$$v(\underline{b}^*, \underline{b}^*) = \frac{\overline{\theta}-\underline{\theta}}{nD(\delta)}\Big[(1-q)\delta q(1-q) - q\big[(1-q^n)(1-\delta q)-n(1-\delta)(1-q)\big]\Big]. \quad \text{(A.154)}$$

As in Cases 1 and 2, I show that $\lim_{\delta\to 1} v(\underline{b}^*, \underline{b}^*) < 0$ implying that $v(\underline{b}^*, \underline{b}^*) < 0$ for any δ sufficiently close to 1:

$$\begin{aligned}\lim_{\delta\to 1} v'(\underline{b}^*, \underline{b}^*) &= \frac{\overline{\theta}-\underline{\theta}}{nD(1)}\big[(1-q)q(1-q)-q(1-q^n)(1-q)\big] \\ &= \frac{q(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\big[(1-q)-(1-q^n)\big] \\ &= \frac{q(1-q)(\overline{\theta}-\underline{\theta})}{nD(1)}\big[q^n-q\big] < 0.\end{aligned}$$

□

A.8 Proofs of Propositions 1.7, 1.9, 1.10, 1.11, and Lemma 1.6 (Parameter regions)

A.8.1 Proof of Proposition 1.7

Proof. Both sides of the equation can be divided by $1-q$ to obtain: $\sum_{k=0}^{n-1} q^k - nq = 0$, which can again be divided by $1-q$ to obtain: $1 - \sum_{k=1}^{n-2}(n-1-k)q^k = 0$. Define the function:

$$g(q) = 1 - \sum_{k=1}^{n-2}(n-1-k)q^k.$$

Clearly $g(0) = 1$, and $g(1)$ is given by:

$$\begin{aligned} g(1) &= 1 - \sum_{k=1}^{n-2}(n-1-k) = 1 - (n-1)(n-2) + \sum_{k=1}^{n-2} k \\ &= 1 - (n-1)(n-2) + \frac{(n-1)(n-2)}{2} = 1 - \frac{(n-1)(n-2)}{2} = \frac{n}{2}(3-n) < 0. \end{aligned}$$

hence the equation has a solution on $(0,1)$ for every $n \geq 4$ by the Intermediate Value Theorem.

Consider now the derivative of $g(\cdot)$:

$$g'(q) = -\sum_{k=1}^{n-2}(n-1-k)kq^{k-1} < 0.$$

which implies that the solution q^* is unique and that $q < \frac{1-q^n}{n(1-q)}$ for all $q < q^*$ and vice versa. □

A.8.2 Proof of Proposition 1.9

Proof. Consider the equation:

$$\begin{aligned}
&(1-q^n)(1-q) = q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \\
&\Leftrightarrow (1-q^n) = q^{n-1}\big(1-(1-q)^n\big)\left[n-\sum_{k=0}^{n-1} q^k\right] \\
&\Leftrightarrow (1-q)\sum_{k=0}^{n-1} q^k = q^{n-1}\big(1-(1-q)^n\big)(1-q)\sum_{k=0}^{n-2}(n-1-k)q^k \\
&\Leftrightarrow \sum_{k=0}^{n-1} q^k = q^{n-1}\big(1-(1-q)^n\big)\sum_{k=0}^{n-2}(n-1-k)q^k.
\end{aligned}$$

and consider the function:

$$g(q) = q^{n-1}\big(1-(1-q)^n\big)\sum_{k=0}^{n-2}(n-1-k)q^k - \sum_{k=0}^{n-1} q^k.$$

Clearly $g(0) = -1$ and $g(1)$ is computed as:

$$\begin{aligned}
g(1) &= \sum_{k=0}^{n-2}(n-1-k)1^k - \sum_{k=0}^{n-1} 1^k \\
&= (n-1)^2 - \sum_{k=0}^{n-2} k - n \\
&= (n-1)^2 - \frac{(n-1)(n-2)}{2} - n = n\frac{n-3}{2} > 0.
\end{aligned}$$

The result follows by continuity of $g(q)$.

□

A.8.3 Proof of Proposition 1.10

Proof. Note that the expression can be rewritten as:

$$\underbrace{(1-q^n)(1-q)}_{\to 1-q \text{ as } n\to\infty} - nq^{n-1}\underbrace{(1-q)\big(1-(1-q)^n\big)}_{\to 1-q \text{ as } n\to\infty} + \underbrace{q^{n-1}(1-q^n)\big(1-(1-q)^n\big)}_{\to 0 \text{ as } n\to\infty}.$$

It thus remains to check that $\lim_{n\to\infty} nq^{n-1} = 0$. Taking logs, I get:

$$\log\left(nq^{n-1}\right) = \log(n) + (n-1)\log(q) \leq \sqrt{n-1} + (n-1)\log(q)$$
$$= (n-1)\left(\frac{1}{\sqrt{n-1}} + \log(q)\right).$$

Note that since $\log(q)$ is strictly negative and $\frac{1}{\sqrt{n-1}}$ goes to 0 as n goes to infinity, we have for a large enough n:

$$(n-1)\left(\frac{1}{\sqrt{n-1}} + \log(q)\right) \leq (n-1)\frac{\log(q)}{2}.$$

Since $\log(q) < 0$ we have $\lim_{n\to\infty}(n-1)\frac{\log(q)}{2} = -\infty$, but then $\lim_{n\to\infty} \log\left(nq^{n-1}\right) = -\infty$, which establishes the claim.

□

A.8.4 Proof of Proposition 1.11

Proof. The parameter restriction can be rewritten as:

$$\frac{1-q^n}{1-(1-q)^n} > q^{n-1}\left[n - \sum_{k=0}^{n-1} q^k\right].$$

Observe that $\frac{1-q^n}{1-(1-q)^n} \geq 1$ for all $q \leq \frac{1}{2}$ since $1-q^n \geq 1-(1-q)^n$ is equivalent to $1-q \geq q$. It thus suffices to show that $1 \geq nq^{n-1}$ for all $q \in (0, \frac{1}{2}]$. Define the function $f(q) = nq^{n-1} - 1$. It is clearly strictly increasing in q since $f'(q) = n(n-1)q^{n-2}$. It thus suffices to check that the claim is true for $q = \frac{1}{2}$ or $1 \geq n\frac{1}{2^{n-1}}$ which is equivalent to $2^{n-1} \geq n$, which is true for all $n \geq 2$. □

A.8.5 Proof of Lemma 1.6

Proof. We can rewrite the two inequalities as:

$$q^{n-1}\left(1-(1-q)^n\right)\left[n(1-q) - (1-q^n)\right] - (1-q^n)(1-q), \geq 0. \tag{A.155}$$

$$n(1-q)q - (1-q^n) \geq 0 \tag{A.156}$$

Our goal is to show that the inequality in (A.155) implies the inequality in (A.156). It suffices to show that

$$n(1-q)q-(1-q^n)\geq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big]-(1-q^n)(1-q), \quad \text{(A.157)}$$

which can be rewritten as:

$$n(1-q)q-(1-q^n)+(1-q^n)(1-q)\geq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \quad \text{(A.158)}$$

$$\Leftrightarrow n(1-q)q-q(1-q^n)\geq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \quad \text{(A.159)}$$

$$\Leftrightarrow q\big[n(1-q)-(1-q^n)\big]\geq q^{n-1}\big(1-(1-q)^n\big)\big[n(1-q)-(1-q^n)\big] \quad \text{(A.160)}$$

$$\Leftrightarrow 1\geq q^{n-2}\big(1-(1-q)^n\big). \quad \text{(A.161)}$$

which is clearly true for any $n\geq 2$ and $q\in(0,1)$.

□

Appendix B

Appendix to Chapter 3

B.1 Mixed and correlated strategies

In this appendix, we provide an argument suggesting that the treatment of mixed and correlated strategies in our environment might require altogether different methods. In particular, we explore one rather natural approach one could take to prove that mixed and correlated strategies are outcome-equivalent to pure strategies, and show, by providing a counterexample, that this approach does not yield the desired result.

Suppose that the players randomize over the sets of signals $R^b = \{\sigma_1^b, \sigma_2^b, \ldots, \sigma_K^b\}$ and $R^s = \{\sigma_1^s, \sigma_2^s, \ldots, \sigma_N^s\}$. Their strategy profile gives rise to the following joint distribution over information acquisition actions

	σ_1^s	σ_2^s	$\ldots$	σ_N^s
σ_1^b	$\mathcal{P}[\sigma_1^b, \sigma_1^s]$	$\mathcal{P}[\sigma_1^b, \sigma_2^s]$	$\ldots$	$\mathcal{P}[\sigma_1^b, \sigma_1^s]$
σ_2^b	$\mathcal{P}[\sigma_2^b, \sigma_1^s]$	$\mathcal{P}[\sigma_2^b, \sigma_2^s]$	$\ldots$	$\mathcal{P}[\sigma_2^b, \sigma_N^s]$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
σ_K^b	$\mathcal{P}[\sigma_K^b, \sigma_1^s]$	$\mathcal{P}[\sigma_K^b, \sigma_2^s]$	$\ldots$	$\mathcal{P}[\sigma_K^b, \sigma_N^s]$

Note that these randomizations could in principle be correlated if we enriched our setup with an additional communication stage at the beginning of the game, in which the mechanism designer would issue correlated recommendations to the players. We show below, however, that even independent randomizations cause difficulties.

If one wanted to prove that our restriction to pure strategies is without loss of generality, one could define a new information structure by finding the average over the information structures given above as follows

$$\hat{\alpha}(s_i^b, s_j^s; v) \equiv \sum_{\sigma^b \in R^b} \sum_{\sigma^s \in R^s} \mathcal{P}[\sigma^b, \sigma^s] \alpha[\sigma^b, \sigma^s](s_i^b, s_j^s; v),$$

and notice that, due to Bayes-plausibility of the new information structure, the new information structure can be induced by a pure strategy profile $(\hat{\sigma}^b, \hat{\sigma}^s)$. One could then hope that if the original distribution of the information structures arises in some equilibrium, then the new information structure can also arise in an outcome equivalent equilbrium of a possibly different mechanism. The next counterexample shows that this strategy will not work: it is possible to construct a deviation from the resulting pure strategy profile $(\hat{\sigma}^b, \hat{\sigma}^s)$ that induces an information structure that cannot be induced by a deviation from the original mixed/correlated strategy profile (see Gentzkow and Kamenica (2017) and Li and Norman (2018) who point out a similar issue in the context of multisender Bayesian persuasion).

B.1.1 Counterexample

Consider the following strategy profile:

	σ_1^s	σ_2^s
σ^b	$\frac{1}{2}$	$\frac{1}{2}$

In words, the seller mixes between σ_1^s and σ_2^s with equal probabilities. The buyer plays σ^b with probability 1. The strategies are defined as follows:

- $\sigma^b = (S^b, \mathbf{S}^b)$, where $S^b = \{s_1^b, s_2^b\}$ and $\mathbf{S}^b$ is given by:

$$\mathbf{S}^b(x) = \begin{cases} s_1^b & \text{if } x \in [0, 0.25] \cup (0.5, 0.75], \\ s_2^b & \text{if } x \in (0.25, 0.5] \cup (0.75, 1]. \end{cases}$$

The corresponding partition of $X = [0, 1]$ is illustrated by:

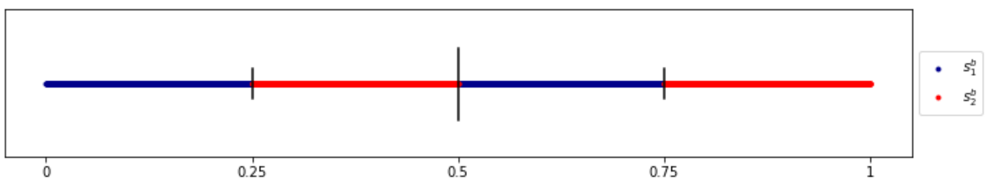

- $\sigma_1^s = (S_1^s, \mathbf{S}_1^s)$, where $S_1^s = \{s_1^s, s_2^s\}$ and $\mathbf{S}_1^s$ is given by:

$$\mathbf{S}_1^s(x) = \begin{cases} s_1^s & \text{if } x \in [0, 0.25] \cup (0.5, 0.75], \\ s_2^s & \text{if } x \in (0.25, 0.5] \cup (0.75, 1]. \end{cases}$$

The corresponding partition of $X = [0, 1]$ is illustrated by:

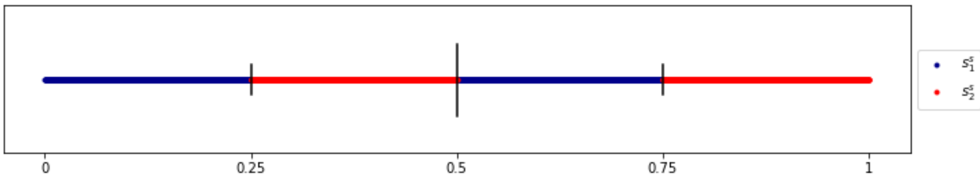

- $\sigma_2^s = (S_2^s, \mathbf{S}_2^s)$, where $S_2^s = \{s_1^s, s_2^s\}$ and $\mathbf{S}_2^s$ is given by:

$$\mathbf{S}_2^s(x) = \begin{cases} s_1^s & \text{if } x \in (0.25, 0.5] \cup (0.75, 1], \\ s_2^s & \text{if } x \in [0, 0.25] \cup (0.5, 0.75]. \end{cases}$$

The corresponding partition of $X = [0, 1]$ is illustrated by:

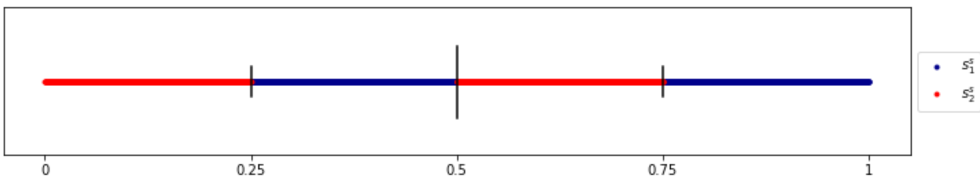

Observe that if the players play the signal profile (σ^b, σ_1^s), they induce the information structure $\alpha[\sigma^b, \sigma_1^s]$ given by:

State $\underline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{4}$	0
s_2^b	0	$\frac{1}{4}$

State $\overline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{4}$	0
s_2^b	0	$\frac{1}{4}$

Likewise, if the players play the signal profile (σ^b, σ_2^s), they induce the information structure $\alpha[\sigma^b, \sigma_2^s]$ given by

State $\underline{v}$	s_1^s	s_2^s
s_1^b	0	$\frac{1}{4}$
s_2^b	$\frac{1}{4}$	0

State $\overline{v}$	s_1^s	s_2^s
s_1^b	0	$\frac{1}{4}$
s_2^b	$\frac{1}{4}$	0

The average over the two information structures $\frac{1}{2}\alpha[\sigma^b,\sigma_1^s]+\frac{1}{2}\alpha[\sigma^b,\sigma_2^s]$ is given by:

State $\underline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{8}$	$\frac{1}{8}$
s_2^b	$\frac{1}{8}$	$\frac{1}{8}$

State $\overline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{8}$	$\frac{1}{8}$
s_2^b	$\frac{1}{8}$	$\frac{1}{8}$

Lemma 3.1 in the main text ensures that $\frac{1}{2}\alpha[\sigma^b,\sigma_1^s]+\frac{1}{2}\alpha[\sigma^b,\sigma_2^s]$ can be induced by a profile of pure signals. Use $\left(\hat{\sigma}^b,\hat{\sigma}^s\right)$ to denote this profile of pure signals. Lemma 3.2 in the main text shows that, by deviating to some $\tilde{\sigma}^b$ (i.e. to the pure signals profile $\left(\tilde{\sigma}^b,\hat{\sigma}^s\right)$), the buyer can induce any information structure that has the same seller-marginals as $\frac{1}{2}\alpha[\sigma^b,\sigma_1^s]+\frac{1}{2}\alpha[\sigma^b,\sigma_2^s]$. In particular, there exists $\tilde{\sigma}^b$ such that $\left(\tilde{\sigma}^b,\hat{\sigma}^s\right)$ induces $\alpha[\sigma^b,\sigma_1^s]$ since $\alpha[\sigma^b,\sigma_1^s]$ has the same seller-marginals as $\frac{1}{2}\alpha[\sigma^b,\sigma_1^s]+\frac{1}{2}\alpha[\sigma^b,\sigma_2^s]$. The following proposition, however, shows that it's impossible to obtain $\alpha[\sigma^b,\sigma_1^s]$ by taking averages over the information structures induced by any deviation from σ^b when the seller plays his original mixed strategy $\frac{1}{2}\sigma_1^s+\frac{1}{2}\sigma_s^2$:

Proposition B.1. *There is no $\tilde{\sigma}^b$ such that $\alpha[\sigma^b,\sigma_1^s]=\frac{1}{2}\alpha[\tilde{\sigma}^b,\sigma_1^s]+\frac{1}{2}\alpha[\tilde{\sigma}^b,\sigma_2^s]$.*

Proof. Suppose for a contradiction that there is such a $\tilde{\sigma}^b$ and recall that $\alpha[\sigma^b,\sigma_1^s]$ is given by:

State $\underline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{4}$	0
s_2^b	0	$\frac{1}{4}$

State $\overline{v}$	s_1^s	s_2^s
s_1^b	$\frac{1}{4}$	0
s_2^b	0	$\frac{1}{4}$

Since only signal realizations s_1^b and s_2^b occur with positive probability under $\tilde{\sigma}^b$, it is without loss of generality to restrict attention to $\tilde{\sigma}^b=(\tilde{S}^b,\tilde{\mathbf{S}}^{\mathbf{b}})$ such that $\tilde{S}^b=\{s_1^b,s_2^b\}$ and $\tilde{\mathbf{S}}^{\mathbf{b}}:X\to\tilde{S}^b$. To obtain a contradiction, we make the following observations:

- $\alpha[\sigma^b,\sigma_1^s](s_1^b,s_2^s;\underline{v})=0$, hence it must be true that $\alpha[\tilde{\sigma}^b,\sigma_1^s](s_1^b,s_2^s;\underline{v})=\alpha[\tilde{\sigma}^b,\sigma_2^s](s_1^b,s_2^s;\underline{v})=0$. Given the above definitions of σ_1^s and σ_2^s these imply that $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_1^b)\cap(0.25,0.5]=\emptyset$ and $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_1^b)\cap[0,0.25]=\emptyset$ respectively, which in turns means that $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_1^b)\cap[0,0.5]=\emptyset$.

- $\alpha[\sigma^b,\sigma_1^s](s_2^b,s_1^s;\underline{v})=0$, hence it must be true that $\alpha[\tilde{\sigma}^b,\sigma_1^s](s_2^b,s_1^s;\underline{v})=\alpha[\tilde{\sigma}^b,\sigma_2^s](s_2^b,s_1^s;\underline{v})=0$. Given the above definitions of σ_1^s and σ_2^s these imply that $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_2^b)\cap[0,0.25]=\emptyset$ and $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_2^b)\cap(0.25,0.5]=\emptyset$ respectively, which in turns means that $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s_2^b)\cap[0,0.5]=\emptyset$.

Hence $\left([\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s^b_1) \cup [\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s^b_2) \right) \cap [0, 0.5] = \emptyset$ implying that $[\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s^b_1) \cup [\tilde{\mathbf{S}}^{\mathbf{b}}]^{-1}(s^b_2) \neq X$, implying in turn that $\tilde{\mathbf{S}}^{\mathbf{b}}$ cannot be a function from X to $\tilde{S}^b$.

□

B.2 Proof of Lemma 3.3

Proof. We prove the statement for the buyer only as the proof for the seller is analogous. Suppose the set of payoff-relevant states of the world is given by $V = \{\underline{v}, \ldots, \overline{v}\}$ and suppose that the proposed information structure has k signal realizations for the buyer and n signal realizations for the seller. If $|S^b| = k$ and $|S^s| = n$, then the information structure on is a collection of $k \times n$ matrices, one for each state (we adopt the convention that the buyer is a *row player* and the seller is a *column player*):

State v	s_1^s	s_2^s	$\ldots$	s_n^s
s_1^b	$\alpha_{11}(v)$	$\alpha_{12}(v)$	$\ldots$	$\alpha_{1n}(v)$
s_2^b	$\alpha_{21}(v)$	$\alpha_{22}(v)$	$\ldots$	$\alpha_{2n}(v)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s_k^b	$\alpha_{k1}(v)$	$\alpha_{k2}(v)$	$\ldots$	$\alpha_{kn}(v)$

The cost of this information structure for the buyer is given by:

$$c^b(\alpha) = H(\mu_0) + \sum_{i=1}^{k} \sum_{v \in V} \left[\left(\sum_{j=1}^{n} \alpha_{ij}(v) \right) \log \left(\frac{\sum_{j=1}^{n} \alpha_{ij}(v)}{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v})} \right) \right].$$

Define $f_i(\alpha) \equiv \sum_{v \in V} \left[\left(\sum_{j=1}^{n} \alpha_{ij}(v) \right) \log \left(\frac{\sum_{j=1}^{n} \alpha_{ij}(v)}{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v})} \right) \right]$, the expected entropy component of signal realization s_i^b. The cost function can then be written as $c^b(\alpha) = H(\mu_0) + \sum_{i=1}^{m} f_i(\alpha)$. We first show the following:

Lemma B.1. *$f_i(\alpha)$ is convex for every i.*

Proof. We first find $\nabla f_i(\alpha)$. To do that, note that the partial derivative of $f_i(\alpha)$ with respect to any $\alpha_{il}(v)$ is the same across all l and is given by:

$$\frac{\partial f_i(\alpha)}{\partial \alpha_{il}(v)} = \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right) + \left(\sum_{j=1}^n \alpha_{ij}(v)\right) \frac{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}{\sum_{j=1}^n \alpha_{ij}(v)} \frac{\partial}{\partial \alpha_{il}(v)} \left(\frac{\sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right)$$
$$+ \sum_{\hat{v}\neq v} \left(\sum_{j=1}^n \alpha_{ij}(\hat{v})\right) \frac{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}{\sum_{j=1}^n \alpha_{ij}(\hat{v})} \frac{\partial}{\partial \alpha_{il}(v)} \left(\frac{\sum_{j=1}^n \alpha_{ij}(\hat{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right),$$

which can be rewritten as:

$$\begin{aligned}
\frac{\partial f_i(\alpha)}{\partial \alpha_{il}(v)} &= \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right) + \left(\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})\right) \frac{\partial}{\partial \alpha_{il}(v)} \left(\sum_{\hat{v}\in V} \frac{\sum_{j=1}^n \alpha_{ij}(\hat{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right) \\
&= \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right) + \left(\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})\right) \frac{\partial}{\partial \alpha_{il}(v)} 1 \\
&= \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right).
\end{aligned}$$

The gradient of $f_i(\alpha)$ is therefore given by:

$$(\nabla f_i(\alpha))^T = \left[\log\left(\frac{\sum_{j=1}^n \alpha_{ij}(\underline{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right), \ldots, \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(\underline{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right), \ldots, \ldots, \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(\overline{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right), \ldots, \log\left(\frac{\sum_{j=1}^n \alpha_{ij}(\overline{v})}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}\right)\right].$$

To determine the Hessian of $f_i(\alpha)$ we have to take second-order derivatives. Note that for any l and r and for any state v the following is true:

$$\frac{\partial^2 f_i(\alpha)}{\partial \alpha_{il}(v) \partial \alpha_{ir}(v)} = \frac{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}{\sum_{j=1}^n \alpha_{ij}(v)} \frac{1 \sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v}) - 1 \sum_{j=1}^n \alpha_{ij}(v)}{\left[\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})\right]^2} = \frac{1}{\sum_{j=1}^n \alpha_{ij}(v)} \frac{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v}) - \sum_{j=1}^n \alpha_{ij}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{ij}(\hat{v})}$$

Defining $A_i(v) \equiv \sum_{j=1}^{n} \alpha_{ij}(v)$, we can write:

$$\frac{\partial^2 f_i(\alpha)}{\partial \alpha_{il}(v) \alpha_{ir}(v)} = \frac{1}{A_i(v)} \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(v)}{\sum_{\hat{v} \in V} A_i(\hat{v})}$$

Note also that for every l and r and for any pair of states $v \neq \tilde{v}$ the following is true:

$$\frac{\partial^2 f_i(\alpha)}{\partial \alpha_{il}(\tilde{v}) \partial \alpha_{ir}(v)} = \frac{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v})}{\sum_{j=1}^{n} \alpha_{ij}(v)} \frac{0 \sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v}) - 1 \sum_{j=1}^{n} \alpha_{ij}(v)}{\left[\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v}) \right]^2} = \frac{-1}{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{ij}(\hat{v})} = \frac{-1}{\sum_{\hat{v} \in V} A_i(\hat{v})}$$

The Hessian of $f_i(\alpha)$ can then be written as $\nabla^2 f_i(\alpha) = \frac{1}{\sum_{\hat{v} \in V} A_i(\hat{v})} \mathcal{H}_i(\alpha)$, where $\mathcal{H}_i(\alpha)$ is the following matrix:

$$\begin{array}{c|cccccccc}
 & \alpha_{i1}(\underline{v}) & \ldots & \alpha_{in}(\underline{v}) & \ldots & \ldots & \alpha_{i1}(\overline{v}) & \ldots & \alpha_{in}(\overline{v}) \\
\hline
\alpha_{i1}(\underline{v}) & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\underline{v})}{A_i(\underline{v})} & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\underline{v})}{A_i(\underline{v})} & \ldots & \ldots & -1 & \ldots & -1 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
\alpha_{in}(\underline{v}) & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\underline{v})}{A_i(\underline{v})} & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\underline{v})}{A_i(\underline{v})} & \ldots & \ldots & -1 & \ldots & -1 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots \\
\vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots \\
\alpha_{i1}(\overline{v}) & -1 & \ldots & -1 & \ldots & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\overline{v})}{A_i(\overline{v})} & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\overline{v})}{A_i(\overline{v})} \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
\alpha_{in}(\overline{v}) & -1 & \ldots & -1 & \ldots & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\overline{v})}{A_i(\overline{v})} & \ldots & \frac{\sum_{\hat{v} \in V} A_i(\hat{v}) - A_i(\overline{v})}{A_i(\overline{v})}
\end{array}$$

We now show that $\nabla^2 f_i(\alpha)$ is positive semi-definite. To do that consider an arbitrary vector $x \in \mathbb{R}^{n|V|}$ and evaluate $x^T \nabla^2 f_i(\alpha) x$. Let $x(v) \in \mathbb{R}^n$ for states $v \in V$ be such that x can be obtained by concatenating vectors $x(v)$ across all $v \in V$. Let e denote the vector consisting of n ones, i.e. $e^T = [1, \dots, 1] \in \mathbb{R}^n$. We then have:

$$
\begin{aligned}
x^T \nabla^2 f_i(\alpha) x &= \frac{1}{\sum_{\hat{v}\in V} A_i(\hat{v})} \sum_{v\in V} \left(\frac{\sum_{\hat{v}\in V} A_i(\hat{v}) - A_i(v)}{A_i(v)} \big(e^T x(v)\big)^2 - e^T x(v) \sum_{\hat{v}\neq v} e^T x(\hat{v}) \right) \\
&= \frac{1}{\sum_{\hat{v}\in V} A_i(\hat{v})} \sum_{v\in V} \left(\frac{\sum_{\hat{v}\in V} A_i(\hat{v})}{A_i(v)} \big(e^T x(v)\big)^2 - \big(e^T x(v)\big)^2 - e^T x(v) \sum_{\hat{v}\neq v} e^T x(\hat{v}) \right) \\
&= \sum_{v\in V} \frac{1}{A_i(v)} \big(e^T x(v)\big)^2 - \frac{1}{\sum_{\hat{v}\in V} A_i(\hat{v})} \sum_{v\in V} \left(\big(e^T x(v)\big)^2 + e^T x(v) \sum_{\hat{v}\neq v} e^T x(\hat{v}) \right) \\
&= \sum_{v\in V} \frac{1}{A_i(v)} \big(e^T x(v)\big)^2 - \frac{1}{\sum_{v\in V} A_i(v)} \left(\sum_{v\in V} e^T x(v) \right)^2
\end{aligned}
$$

Defining $X(v) \equiv e^T x(v)$ for every $v \in V$, we can write:

$$
x^T \nabla^2 f_i(\alpha) x = \sum_{v\in V} \frac{1}{A_i(v)} X^2(v) - \frac{1}{\sum_{v\in V} A_i(v)} \left(\sum_{v\in V} X(v) \right)^2
$$

To show that $\nabla^2 f_i(\alpha)$ is positive semi-definite, we have to show that the above expression is weakly positive for all $\{X(v)\}_{v\in V}$. In order to do that, we show that

$$
\min_{\{X(v)\}_{v\in V}} \left\{ \sum_{v\in V} \frac{1}{A_i(v)} X^2(v) - \frac{1}{\sum_{v\in V} A_i(v)} \left(\sum_{v\in V} X(v) \right)^2 \right\} \geq 0
$$

To that end, consider the restricted problem for some $\breve{X} \in \mathbb{R}$ given by:

$$
\min_{\{X(v)\}_{v\in V}} \left\{ \sum_{v\in V} \frac{1}{A_i(v)} X^2(v) - \frac{1}{\sum_{v\in V} A_i(v)} \breve{X}^2 \text{ s.t. } \sum_{v\in V} X(v) = \breve{X} \right\}.
$$

The restricted problem is clearly convex in $\{X(v)\}_{v\in V}$, hence the first order conditions are necessary and sufficient for minimization. The Lagrangian of this restricted problem is given by:

$$\mathcal{L}\big(X;\eta\big)=\sum_{v\in V}\frac{1}{A_i(v)}X^2(v)-\frac{1}{\sum_{v\in V}A_i(v)}\check{X}^2-2\eta\Big(\sum_{v\in V}X(v)-\check{X}\Big).$$

The optimality conditions are given by:

$$\begin{cases}\frac{1}{A_i(v)}2X^*(v)-2\eta^*=0 & \forall v\in V,\\ \sum_{v\in V}X^*(v)=\check{X}.\end{cases}$$

The minimum is achieved at $X^*(v)=\frac{A_i(v)\check{X}}{\sum_{\hat{v}\in V}A_i(\hat{v})}$, and the value of the objective achieved at the minimum is given by:

$$\begin{aligned}\sum_{v\in V}\frac{1}{A_i(v)}\frac{A_i^2(v)\check{X}^2}{\big(\sum_{\hat{v}\in V}A_i(\hat{v})\big)^2}-\frac{1}{\sum_{v\in V}A_i(v)}\check{X}^2&=\sum_{v\in V}\frac{A_i(v)\check{X}^2}{\big(\sum_{\hat{v}\in V}A_i(\hat{v})\big)^2}-\frac{1}{\sum_{v\in V}A_i(v)}\check{X}^2\\&=\check{X}^2\left[\frac{\sum_{v\in V}A_i(v)}{\big(\sum_{v\in V}A_i(v)\big)^2}-\frac{1}{\sum_{v\in V}A_i(v)}\right]=0.\end{aligned}$$

implying that the minimal value achieved in the restricted problem is zero for every $\check{X}\in\mathbb{R}$, implying in turn that the minimal value achieved by the unrestricted problem is also zero, hence $x^T\nabla^2 f_i(\alpha)x\geq 0$ for every $x\in\mathbb{R}^{n|V|}$ and that $\nabla^2 f_i(\alpha)$ is positive semi-definite, which means that $f_i(\alpha)$ is convex. □

Recall that $c(\alpha)=H(\mu_0)+\sum_{i=1}^m f_i(\alpha)$ and hence is a sum of convex functions, implying that $c(\alpha)$ is convex. □

B.3 Proof of Proposition 3.1

Proof. Consider an indirect mechanism $\big(M_{\text{IN}}, q_{\text{IN}}, t_{\text{IN}}\big)$. Let $\big[\big(\sigma^b, \{\mathbf{m}^b_{\text{IN}}[\hat{\sigma}^b]\}_{\hat{\sigma}^b\in\Sigma^b}\big),\ \big(\sigma^s, \{\mathbf{m}^s_{\text{IN}}[\hat{\sigma}^s]\}_{\hat{\sigma}^s\in\Sigma^s}\big)\big]$, where $\sigma^b = \big(S^b, \mathbf{S}^b\big)$ and $\sigma^s = \big(S^s, \mathbf{S}^s\big)$, be its Nash equilibrium. Let α be the information structure induced by the information acquisition choices $\big(\sigma^b, \sigma^s\big)$. Using Lemmas 3.1 and 3.2, we can write the equilibrium conditions as follows.

- For the buyer:

$$(S^b, \alpha, \mathbf{m}^b_{\text{IN}}[\sigma^b]) \in \underset{\tilde{\alpha}, \tilde{S}^b, \tilde{\mathbf{m}}^b_{\text{IN}}}{\operatorname{argmax}} \sum_{s^b\in\tilde{S}^b}\sum_{s^s\in S^s}\sum_{v\in V} \tilde{\alpha}(s^b, s^s; v)\Big(q^b_{\text{IN}}\big(\tilde{\mathbf{m}}^b_{\text{IN}}(s^b), \mathbf{m}^s_{\text{IN}}[\sigma^s](s^s)\big)u^b(v) - t^b_{\text{IN}}\big(\tilde{\mathbf{m}}^b_{\text{IN}}(s^b), \mathbf{m}^s_{\text{IN}}[\sigma^s](s^s)\big)\Big) - c^b(\tilde{\alpha}),$$

$$\text{s.t.}\quad (1)\ \ \tilde{S}^b \in \mathscr{P}\big(\mathbb{N}\big),\ \ \tilde{\alpha} \in \Delta\big(\tilde{S}^b \times S^s \times V\big),\ \ \tilde{\mathbf{m}}^b_{\text{IN}} : \tilde{S}^b \to M^b_{\text{IN}};$$

$$(2)\ \ \operatorname{marg}_{S^s\times V}\tilde{\alpha} = \operatorname{marg}_{S^s\times V}\alpha.$$

- For the seller:

$$(S^s, \alpha, \mathbf{m}^s_{\text{IN}}[\sigma^s]) \in \underset{\tilde{\alpha}, \tilde{S}^s, \tilde{\mathbf{m}}^s_{\text{IN}}}{\operatorname{argmax}} \sum_{s^b\in S^b}\sum_{s^s\in\tilde{S}^s}\sum_{v\in V} \tilde{\alpha}(s^b, s^s; v)\Big(t^s_{\text{IN}}\big(\mathbf{m}^b_{\text{IN}}[\sigma^b](s^b), \tilde{\mathbf{m}}^s_{\text{IN}}(s^s)\big) - q^s_{\text{IN}}\big(\mathbf{m}^b_{\text{IN}}[\sigma^b](s^b), \tilde{\mathbf{m}}^s_{\text{IN}}(s^s)\big)u^s(v)\Big) - c^s(\tilde{\alpha}),$$

$$\text{s.t.}\quad (1)\ \ \tilde{S}^s \in \mathscr{P}\big(\mathbb{N}\big),\ \ \tilde{\alpha} \in \Delta\big(S^b \times \tilde{S}^s \times V\big),\ \ \tilde{\mathbf{m}}^s_{\text{IN}} : \tilde{S}^s \to M^s_{\text{IN}};$$

$$(2)\ \ \operatorname{marg}_{S^b\times V}\tilde{\alpha} = \operatorname{marg}_{S^b\times V}\alpha.$$

Consider the following direct mechanism $\big(M_{\text{D}}, q_{\text{D}}, t_{\text{D}}\big)$, where the message space is given by $M_{\text{D}} \equiv \big(S^b \cup \{m_\emptyset\}\big) \times \big(S^s \cup \{m_\emptyset\}\big)$; the allocation function is defined as $q^p_{\text{D}}(s^b, s^s) \equiv q^p_{\text{IN}}\big(\mathbf{m}^b_{\text{IN}}[\sigma^b](s^b), \mathbf{m}^s_{\text{IN}}[\sigma^s](s^s)\big)$, and the transfer function is given by $t^p_{\text{D}}(s^b, s^s) \equiv t^p_{\text{IN}}\big(\mathbf{m}^b_{\text{IN}}[\sigma^b](s^b), \mathbf{m}^s_{\text{IN}}[\sigma^s](s^s)\big)$ for both players $p \in \{b, s\}$. We claim that $\big[\big(\sigma^b, \{\mathbf{m}^b_{\text{D}}[\hat{\sigma}^b]\}_{\hat{\sigma}^b\in\Sigma^b}\big),\ \big(\sigma^s, \{\mathbf{m}^s_{\text{D}}[\hat{\sigma}^s]\}_{\hat{\sigma}^s\in\Sigma^s}\big)\big]$, where $\mathbf{m}^p_{\text{D}}[\hat{\sigma}^p] = \mathbf{m}^p_T$ for all $\hat{\sigma}^p \in \Sigma^p$ is a Nash equilibrium in the direct mechanism.

Suppose for a contradiction that this is not the case, then one of the players has a profitable deviation to untruthful reporting, a different information acquisition action, or both. Let us suppose that it is the buyer who can profitably deviate (the argument for the seller is identical), then the tuple $(S^b, \alpha, \mathbf{m}_T^b)$ violates the constraint IC_A^b for the direct mechanism $\big(M_\text{D}, q_\text{D}, t_\text{D}\big)$, i.e. there exists an information acquisition action $\tilde{\sigma}^b = \big(\tilde{S}^b, \tilde{\mathbf{S}}^b\big)$ inducing a new joint distribution $\tilde{\alpha} \in \Delta\big(\tilde{S}^b \times S^s \times V\big)$ and $\text{marg}_{S^s \times V} \tilde{\alpha} = \text{marg}_{S^s \times V} \alpha$, and a (possibly but not necessarily truthful) reporting function $\tilde{\mathbf{m}}_\text{D}^b : \tilde{S}^b \to S^b \cup \{m_\emptyset\}$ such that for signal realization s^b:

$$\sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q_\text{D}^b(\tilde{\mathbf{m}}_\text{D}^b(s^b), s^s)u^b(v) - t_\text{D}^b(\tilde{\mathbf{m}}_\text{D}^b(s^b), s^s)\big) - c^b(\tilde{\alpha})$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q_\text{D}^b(s^b, s^s)u^b(v) - t_\text{D}^b(s^b, s^s)\big) - c^b(\alpha),$$

which implies by definition of allocation and transfer functions in the direct mechanism

$$\sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q_\text{IN}^b\big(\mathbf{m}_\text{IN}^b[\sigma^b](\tilde{\mathbf{m}}_\text{D}^b(s^b)), \mathbf{m}_\text{IN}^s[\sigma^s](s^s)\big)u^b(v) - t_\text{IN}^b\big(\mathbf{m}_\text{IN}^b[\sigma^b](\tilde{\mathbf{m}}_\text{D}^b(s^b)), \mathbf{m}_\text{IN}^s[\sigma^s](s^s)\big)\big) - c^b(\tilde{\alpha})$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q_\text{IN}^b\big(\mathbf{m}_\text{IN}^b[\sigma^b](s^b), \mathbf{m}_\text{IN}^s[\sigma^s](s^s)\big)u^b(v) - t_\text{IN}^b\big(\mathbf{m}_\text{IN}^b[\sigma^b](s^b), \mathbf{m}_\text{IN}^s[\sigma^s](s^s)\big)\big) - c^b(\alpha),$$

which in turn means that the tuple $(S^b, \alpha, \mathbf{m}_\text{IN}^b[\sigma^b])$ violates the equilibrium condition, hence a contradiction.

□

B.4 Proof of Proposition 3.2

Let $\left[\left(\sigma^b, \left\{\mathbf{m}^b[\hat{\sigma}^b]\right\}_{\hat{\sigma}^b\in\Sigma^b}\right),\ \left(\sigma^s, \left\{\mathbf{m}^s[\hat{\sigma}^s]\right\}_{\hat{\sigma}^s\in\Sigma^s}\right)\right]$ be a truthful-revelation Nash equilibrium of a direct mechanism. It implies in particular $\mathbf{m}^b[\sigma^b] = \mathbf{m}^b_T$ and $\mathbf{m}^s[\sigma^s] = \mathbf{m}^s_T$. Use $\alpha \in \Delta(S^b \times S^s \times V)$ to denote the joint distribution of signal realizations and states of the world induced by the on-path information acquisition choice $\left(\sigma^b, \sigma^s\right)$. By assumption $\left(\alpha, S^p, \mathbf{m}^p_T\right)$ satisfies *ex ante* incentive compatibility IC^p_{A} for player p and α is Bayes-plausible.

Consider a new profile $\left[\left(\sigma^b, \left\{\mathbf{m}^b_{\text{PBE}}[\hat{\sigma}^b]\right\}_{\hat{\sigma}^b\in\Sigma^b}\right),\ \left(\sigma^s, \left\{\mathbf{m}^s_{\text{PBE}}[\hat{\sigma}^s]\right\}_{\hat{\sigma}^s\in\Sigma^s}\right)\right]$ where $\mathbf{m}^p_{\text{PBE}}[\sigma^p] \equiv \mathbf{m}^p_T$ and $\mathbf{m}^p_{\text{PBE}}[\hat{\sigma}^p]$ for $\hat{\sigma} \neq \sigma^p$ are to be defined below.

By construction, this strategy profile is outcome-equivalent to the original Nash equilibrium strategy profile. We are now going to show that it can be a perfect Bayesian equilibrium profile of the same direct mechanism. To do that, let us first specify the players' beliefs. Let $\mathcal{I}^p\left(\hat{\sigma}^p, s^p\right)$ denote the information set achieved by player p who has played $\hat{\sigma}^p \in \Sigma^p$ and observed a signal realization $s^p \in \mathbb{N}$. Let $\gamma^p\left(\hat{\sigma}^{-p}, s^{-p}, v|\mathcal{I}^p\left(\hat{\sigma}^p, s^p\right)\right)$ denote the belief of player p that player $-p$ has played $\hat{\sigma}^{-p} \in \Sigma^{-p}$, has observed the signal realization $s^{-p} \in \mathbb{N}$; and the state of the world is $v \in V$. We specify the players' beliefs as follows:

1. The beliefs at information sets $\mathcal{I}^p\left(\hat{\sigma}^p, s^p\right)$ such that $\hat{\sigma}^p \neq \sigma^p$ are derived using Bayes rule for the buyer from $\alpha[\hat{\sigma}^b, \sigma^s]$ and for the seller from $\alpha[\sigma^b, \hat{\sigma}^s]$. These beliefs are given by:

$$\gamma^b\left(\hat{\sigma}^s, s^s; v|\mathcal{I}^b\left(\sigma^b, s^b\right)\right) = \begin{cases} \frac{\alpha[\hat{\sigma}^b,\sigma^s](s^b,s^s;v)}{\sum_{i=1}^{+\infty}\alpha[\hat{\sigma}^b,\sigma^s](i,s^s;v)} & \text{for } \hat{\sigma}^s = \sigma^s. \\ 0 & \text{otherwise.} \end{cases}$$

$$\gamma^s\left(\hat{\sigma}^b, s^b; v|\mathcal{I}^s\left(\sigma^s, s^s\right)\right) = \begin{cases} \frac{\alpha[\sigma^b,\hat{\sigma}^s](s^b,s^s;v)}{\sum_{j=1}^{+\infty}\alpha[\sigma^b,\hat{\sigma}^s](s^b,j;v)} & \text{for } \hat{\sigma}^b = \sigma^b, \\ 0 & \text{otherwise.} \end{cases}$$

2. The beliefs at $\mathcal{I}^p\big(\sigma^p, s^p\big)$ are derived using Bayes rule from α. These beliefs are given by:

$$\gamma^b\big(\hat{\sigma}^s, s^s; v|\mathcal{I}^b\big(\sigma^b, s^b\big)\big) = \begin{cases} \frac{\alpha(s^b, s^s; v)}{\sum_{i=1}^{+\infty}\alpha(i, s^s; v)} & \text{for } \hat{\sigma}^s = \sigma^s. \\ 0 & \text{otherwise.} \end{cases}$$

$$\gamma^s\big(\hat{\sigma}^b, s^b; v|\mathcal{I}^s\big(\sigma^s, s^s\big)\big) = \begin{cases} \frac{\alpha(s^b, s^s; v)}{\sum_{j=1}^{+\infty}\alpha(s^b, j; v)} & \text{for } \hat{\sigma}^b = \sigma^b, \\ 0 & \text{otherwise.} \end{cases}$$

Let us now show that the strategy profile $\big[\big(\sigma^b, \big\{\mathbf{m}^b_{\text{PBE}}[\hat{\sigma}^b]\big\}_{\hat{\sigma}^b\in\Sigma^b}\big),\ \big(\sigma^s, \big\{\mathbf{m}^s_{\text{PBE}}[\hat{\sigma}^s]\big\}_{\hat{\sigma}^s\in\Sigma^s}\big)\big]$ is sequentially rational given the beliefs specified above.

B.4.1 Reporting after off-path information acquisition

Let us start with off-path information acquisition actions. Suppose the buyer has arrived at the information set $\mathcal{I}^b\big(\hat{\sigma}^b, s^b\big)$ with $\hat{\sigma}^b \neq \sigma^b$, obtain the report following $\big(\hat{\sigma}^b, s^b\big)$ by solving (in case there are many solutions, pick any):

$$\mathbf{m}^b_{\text{PBE}}[\hat{\sigma}^b](s^b) \equiv \arg\max_{m\in S^b\cup\{m_\emptyset\}} \sum_{s^s\in S^s}\sum_{v\in V} \frac{\alpha[\hat{\sigma}^b, \sigma^s](s^b, s^s; v)}{\sum_{i=1}^{+\infty}\alpha[\hat{\sigma}^b, \sigma^s](i, s^s; v)}\big(q^b(m, s^s)u^b(v) - t^b(m, s^s)\big)$$

The resulting reporting function $\mathbf{m}^b_{\text{PBE}}[\hat{\sigma}^b]$ is sequentially rational by construction.

Likewise suppose the seller has arrived at the information set $\mathcal{I}^s\big(\hat{\sigma}^s, s^s\big)$ with $\hat{\sigma}^s \neq \sigma^s$, obtain the report following $\big(\hat{\sigma}^s, s^s\big)$ by solving:

$$\mathbf{m}^s_{\text{PBE}}[\hat{\sigma}^s](s^s) \equiv \arg\max_{m\in S^s\cup\{m_\emptyset\}} \sum_{s^b\in S^b}\sum_{v\in V} \frac{\alpha[\sigma^b, \hat{\sigma}^s](s^b, s^s; v)}{\sum_{j=1}^{+\infty}\alpha[\sigma^b, \hat{\sigma}^s](s^b, j; v)}\big(t^s(s^b, m) - q^s(s^b, m)u^b(v)\big)$$

The resulting reporting function $\mathbf{m}^{s}_{\mathrm{PBE}}[\hat{\sigma}^{s}]$ is sequentially rational by construction.

It remains to show sequential rationality of truthful reporting after taking the on-path information acquisition action.

B.4.2 Reporting after on-path information acquisition

Let us now move on to on-path information acquisition actions. Suppose player p has arrived at the information set $\mathcal{I}^{p}\big(\sigma^{p}, s^{p}\big)$. At this information set player p believes that player $-p$ has taken his on-path action as well with probability 1, and player p's beliefs about signal realizations are derived from α using Bayes' rule. The proposed perfect Bayesian equilibrium strategy prescribes truthful reporting after playing the on-path information acquisition action. There are two ways, in which player p could deviate from truthful reporting: he could misreport a particular signal realization, or he could abstain following a particular signal realization. In what follows, we show that these deviations are not profitable.

Misreporting a signal realization

If a signal realization s^{b} occurs with positive probability given α, then the buyer is willing to report it truthfully as long as the following *ex interim* incentive compatibility condition is satisfied:

$$(\mathrm{IC}^{b}_{\mathrm{I}}) \qquad \sum_{s^{s}\in S^{s}} \sum_{v\in V} \alpha(s^{b}, s^{s}; v)\big(q^{b}(s^{b}, s^{s})u^{b}(v) - t^{b}(s^{b}, s^{s})\big) \geq \sum_{s^{s}\in S^{s}} \sum_{v\in V} \alpha(s^{b}, s^{s}; v)\big(q^{b}(\tilde{s}^{b}, s^{s})u^{b}(v) - t^{b}(\tilde{s}^{b}, s^{s})\big)$$

for all $\tilde{s}^{b} \in S^{b}$.

Likewise if a signal realization s^{s} occurs with positive probability given α, then the buyer is willing to report it truthfully as long as the following *ex interim* incentive compatibility condition is satisfied:

$$(\mathrm{IC}^s_{\mathrm{I}}) \qquad \sum_{s^b\in S^b}\sum_{v\in V}\alpha(s^b,s^s;v)\big(t^s(s^b,s^s)-q^s(s^b,s^s)u^b(v)\big) \geq \sum_{s^b\in S^b}\sum_{v\in V}\alpha(s^b,s^s;v)\big(t^s(s^b,\tilde{s}^s)-q^s(s^b,\tilde{s}^s)u^b(v)\big)$$

for all $\tilde{s}^s \in S^s$.

The following lemma shows that the *ex interim* incentive compatibility conditions are implied by *ex ante* incentive compatibility conditions:

Lemma B.2. $\mathrm{IC}^p_{\mathrm{A}} \Rightarrow \mathrm{IC}^p_{\mathrm{I}}$ *for both players* $p \in \{b, s\}$

Proof. We show that $\neg IC^b_I \Rightarrow \neg IC^b_A$. The argument for the seller is again identical. Suppose that the mechanism is not *ex interim* incentive compatible for the buyer, i.e. there exists a signal realization $x^b \in S^b$, which occurs with positive probability, and a non-truthful report $\tilde{x}^b \in S^b$ such that:

$$\sum_{s^s\in S^s}\sum_{v\in V}\alpha(x^b,s^s;v)\big(q^b(x^b,s^s)u^b(v)-t^b(x^b,s^s)\big) < \sum_{s^s\in S^s}\sum_{v\in V}\alpha(x^b,s^s;v)\big(q^b(\tilde{x}^b,s^s)u^b(v)-t^b(\tilde{x}^b,s^s)\big)$$

Consider an *ex ante* deviation to $(S^b, \alpha, \tilde{\mathbf{m}}^b)$, where:

$$\tilde{\mathbf{m}}^b(s^b) = \begin{cases} s^b, & \text{if } s^b \neq x^b; \\ \tilde{x}^b, & \text{if } s^b = x^b. \end{cases}$$

The payoff from this deviation is given by:

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big) - c^b(\alpha)$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\alpha),$$

implying that the mechanism is not *ex ante* incentive compatible for the buyer. The argument for the seller is identical. □

Abstaining instead of reporting a signal realization

If a signal realization s^b occurs with positive probability given α, the buyer is willing to report it instead of abstaining if the following *ex interim* individual rationality condition is satisfied:

$$(\mathrm{IR}^b_\mathrm{I}) \qquad \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) \geq 0.$$

Likewise if a signal realization s^s occurs with positive probability given α, the seller is willing to report it instead of abstaining if the following *ex interim* individual rationality condition is satisfied:

$$(\mathrm{IR}^s_\mathrm{I}) \qquad \sum_{s^b \in S^b} \sum_{v \in V} \alpha(s^b, s^s; v)\big(t^s(s^b, s^s) - q^s(s^b, s^s)u^b(v)\big) \geq 0.$$

The following lemma shows that the *ex interim* individual rationality conditions are implied by *ex ante* incentive compatibility conditions:

Lemma B.3. $\mathrm{IC}^p_\mathrm{A} \Rightarrow \mathrm{IR}^p_\mathrm{I}$ *for both players* $p \in \{b, s\}$

Proof. We show that $\neg IR_I^b \Rightarrow \neg IC_A^b$. The argument for the seller is again identical. Suppose that the mechanism is not *ex interim* individually rational for the buyer, i.e. there exists a signal realization $x^b \in S^b$, which occurs with positive probability, such that

$$\sum_{s^s \in S^s} \sum_{v \in V} \alpha(x^b, s^s; v)\big(q^b(x^b, s^s)u^b(v) - t^b(x^b, s^s)\big) < 0$$

Consider an *ex ante* deviation to $(S^b, \alpha, \tilde{\mathbf{m}}^b)$, where:

$$\tilde{\mathbf{m}}^b(s^b) = \begin{cases} s^b, & \text{if } s^b \neq x^b; \\ m_\emptyset, & \text{if } s^b = x^b. \end{cases}$$

The payoff from this deviation is given by:

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big) - c^b(\alpha)$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\alpha),$$

implying that the mechanism is not *ex ante* incentive compatible for the buyer. The argument for the seller is identical. □

B.5 Proof of Lemma 3.4

Proof. We show that $\neg \text{IC}^b_\text{A} \Rightarrow \neg \text{R-IC}^b_\text{A}$. The argument for the seller is identical. Suppose that the mechanism violates *ex ante* incentive compatibility for the buyer, i.e. there exists an *ex ante* deviation $(\tilde{S}^b, \tilde{\alpha}, \tilde{\mathbf{m}}^b)$ where $\tilde{S}^b \in \mathscr{P}\big(\mathbb{N}\big)$, $\tilde{\alpha} \in \Delta\big(\tilde{S}^b \times S^s \times V\big)$, $\tilde{\mathbf{m}}^b : \tilde{S}^b \to S^b \cup \{m_\emptyset\}$, and $\text{marg}_{S^s \times V}\tilde{\alpha} = \text{marg}_{S^s \times V}\alpha$, such that

$$\sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big) - c^b(\tilde{\alpha})$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\alpha).$$

We are going to show that there is a restricted deviation $(\tilde{S}^b_\text{R}, \tilde{\alpha}_\text{R})$ where $\tilde{S}^b_\text{R} = S^b \cup \{s^b_\emptyset\}$, $\tilde{\alpha}_\text{R} \in \Delta\big(\tilde{S}^b_\text{R} \times S^s \times V\big)$, and $\text{marg}_{S^s \times V}\tilde{\alpha}_\text{R} = \text{marg}_{S^s \times V}\alpha$, such that

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}_\text{R}(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\tilde{\alpha}_\text{R})$$
$$> \sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \alpha(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\alpha).$$

Define $\tilde{\mathcal{X}}^b(s^b) \equiv \{x^b \in \tilde{S}^b | \tilde{\mathbf{m}}^b(x^b) = s^b\}$ and $\tilde{\mathcal{X}}^b(m_\emptyset) \equiv \{x^b \in \tilde{S}^b | \tilde{\mathbf{m}}^b(x^b) = m_\emptyset\}$, i.e. the set of all signal realizations $x^b \in \tilde{S}^b$ such that the reports of s^b and $m_\emptyset$ are submitted respectively under $\tilde{\mathbf{m}}^b$. Define the restricted information structure as:

$$\tilde{\alpha}_R(s^b, s^s; v) \equiv \sum_{x^b \in \tilde{\mathcal{X}}^b(s^b)} \tilde{\alpha}(x^b, s^s; v) \qquad \forall\ s^b \in S^b,$$
$$\tilde{\alpha}_R(s^b_\emptyset, s^s; v) \equiv \sum_{x^b \in \tilde{\mathcal{X}}^b(m_\emptyset)} \tilde{\alpha}(x^b, s^s; v).$$

The restricted information structure respects the marginals of the seller by construction, and thus also can be a part of a feasible deviation. Indeed,

$$\sum_{s^b \in S^b} \tilde{\alpha}_R(s^b, s^s; v) + \tilde{\alpha}_R(s^b_\emptyset, s^s; v) = \sum_{s^b \in S^b} \Bigg[\sum_{x^b \in \tilde{\mathcal{X}}^b(s^b)} \tilde{\alpha}(x^b, s^s; v) + \sum_{x^b \in \tilde{\mathcal{X}}^b(m_\emptyset)} \tilde{\alpha}(x^b, s^s; v) \Bigg]$$
$$= \sum_{x^b \in \tilde{S}^b} \tilde{\alpha}(x^b, s^s; v)$$

for every $s^s \in S^s$

Clearly by construction we also obtain

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}_{\mathrm{R}}(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big)$$
$$= \sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big).$$

By construction $\tilde{\alpha}_R$ is Blackwell-less-informative than $\tilde{\alpha}$ for the buyer, which means that the expected entropy of $\tilde{\alpha}$ is lower than that of $\tilde{\alpha}_R$, implying $c^b(\tilde{\alpha}_{\mathrm{R}}) \leq c^b(\tilde{\alpha})$, which in turn implies

$$\sum_{s^b \in S^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}_{\mathrm{R}}(s^b, s^s; v)\big(q^b(s^b, s^s)u^b(v) - t^b(s^b, s^s)\big) - c^b(\tilde{\alpha}_{\mathrm{R}})$$
$$\geq \sum_{s^b \in \tilde{S}^b} \sum_{s^s \in S^s} \sum_{v \in V} \tilde{\alpha}(s^b, s^s; v)\big(q^b(\tilde{\mathbf{m}}^b(s^b), s^s)u^b(v) - t^b(\tilde{\mathbf{m}}^b(s^b), s^s)\big) - c^b(\tilde{\alpha}),$$

establishing the claim.

□

B.6 Proof of Lemma 3.5

Proof. We prove the statement of the lemma for the buyer only. The proof for the seller is analogous. We have to distinguish two cases:

1. $k = 1$: in this case the buyer can only receive one signal realization s_1^b. By Bayes-plausibility we have $\mu_1^b(v) = \mu_0(v) > 0$ for any $v \in V$, hence the statement of the lemma holds trivially.

2. $k > 1$: in this case the buyer can receive many signal realizations. Suppose for a contradiction that there exists a state $v' \in V$ such that after receiving signal realization s_1^b the buyer puts probability zero on state v', i.e. $\mu_1^b(v') = 0$. Note that since the labels of signal realizations do not have any particular meaning in our analysis, choosing s_1^b is without loss of generality. Since s_1^b leads to a zero posterior on v', the information structure at v' can be written as:

State v'	s_1^s	s_2^s	$\dots$	s_l^s	$\dots$	s_n^s
s_1^b	0	0	$\dots$	0	$\dots$	0
s_2^b	$\alpha_{21}(v')$	$\alpha_{22}(v')$	$\dots$	$\alpha_{2l}(v')$	$\dots$	$\alpha_{2n}(v')$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
s_r^b	$\alpha_{r1}(v')$	$\alpha_{r2}(v')$	$\dots$	$\alpha_{rl}(v')$	$\dots$	$\alpha_{rn}(v')$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
s_k^b	$\alpha_{k1}(v')$	$\alpha_{22}(v')$	$\dots$	$\alpha_{kl}(v')$	$\dots$	$\alpha_{kn}(v')$

The payoff from this information structure is given by:

$$\sum_{j=1}^{n}\sum_{v\in V\setminus\{v'\}}\alpha_{1j}(v)\big(q_{1j}^{b}u^{b}(v)-t_{1j}^{b}\big)+\sum_{i=2}^{k}\sum_{j=1}^{n}\sum_{v\in V}\alpha_{ij}(v)\big(q_{ij}^{b}u^{b}(v)-t_{ij}^{b}\big)-H(\mu_0)$$
$$-\sum_{v\in V\setminus\{v'\}}\left(\sum_{j=1}^{n}\alpha_{1j}(v)\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{1j}(v)}{\sum_{j=1}^{n}\sum_{\hat{v}\in V\setminus\{v'\}}\alpha_{1j}(\hat{v})}\right]-\underbrace{0\log 0}_{=0}-\sum_{i=2}^{k}\sum_{v\in V}\left(\sum_{j=1}^{n}\alpha_{ij}(v)\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{ij}(v)}{\sum_{j=1}^{n}\sum_{\hat{v}\in V}\alpha_{ij}(\hat{v})}\right].$$

Observe that at least one of the $\alpha_{ij}(v')$ for some $i\neq 1$ must be strictly positive by Bayes-plausibility. Otherwise Bayes -plausibility would imply $\mu_0(v')=0$ contradicting the full support assumption. Let us assume without loss of generality that $\alpha_{rl}(v')>0$ and consider now an alternative information structure $\tilde{\alpha}$, in which every $\tilde{\alpha}_{ij}(v)=\alpha_{ij}(v)$ for all pairs (i,j) in all states $v\neq v'$. In state v', we transfer a small probability mass from (s_r^b,s_l^s) to (s_1^b,s_l^s), the alternative information structure in state v' can then be written as:

State v'	s_1^s	s_2^s	$\dots$	s_l^s	$\dots$	s_n^s
s_1^b	0	0	$\dots$	ϵ	$\dots$	0
s_2^b	$\alpha_{21}(v')$	$\alpha_{22}(v')$	$\dots$	$\alpha_{2l}(v')$	$\dots$	$\alpha_{2n}(v')$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
s_r^b	$\alpha_{r1}(v')$	$\alpha_{r2}(v')$	$\dots$	$\alpha_{rl}(v')-\epsilon$	$\dots$	$\alpha_{rn}(v')$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
s_k^b	$\alpha_{k1}(v')$	$\alpha_{k2}(v')$	$\dots$	$\alpha_{kl}(v')$	$\dots$	$\alpha_{kn}(v')$

for some small $\epsilon>0$. Observe that $\text{marg}_{S^s\times V}\tilde{\alpha}=\text{marg}_{S^s\times V}\alpha$, and hence $\tilde{\alpha}$ can be a feasible deviation for the buyer in (R-IC$_{\text{A}}^{b}$-1).

The payoff from this deviation is given by:

$$
\begin{aligned}
&\epsilon\big(q^b_{1l}u^b(v') - t^b_{1l}\big) + \sum_{j=1}^{n}\sum_{v\in V\setminus\{v'\}} \alpha_{1j}(v)\big(q^b_{1j}u^b(v) - t^b_{1j}\big) + \sum_{i=2}^{k}\sum_{j=1}^{n}\sum_{v\in V} \alpha_{ij}(v)\big(q^b_{ij}u^b(v) - t^b_{ij}\big) - \epsilon\big(q^b_{rl}u^b(v') - t^b_{rl}\big) - H(\mu_0)\\
&- \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{1j}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{1j}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V\setminus\{v'\}}\alpha_{1j}(\hat v)+\epsilon}\right] - \epsilon\log\left[\frac{\epsilon}{\sum_{j=1}^{n}\sum_{\hat v\in V\setminus\{v'\}}\alpha_{1j}(\hat v)+\epsilon}\right]\\
&- \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{rj}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)-\epsilon}\right] - \left(\sum_{j=1}^{n}\alpha_{rj}(v') - \epsilon\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v')}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)-\epsilon}\right]\\
&- \sum_{i\neq 1,r}\sum_{v\in V} \left(\sum_{j=1}^{n} \alpha_{ij}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{ij}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{ij}(\hat v)}\right].
\end{aligned}
$$

The gain from this deviation as a function of ϵ is given by:

$$
\begin{aligned}
G(\epsilon) \equiv\ &\epsilon\big(q^b_{1l}u^b(v') - t^b_{1l}\big) - \epsilon\big(q^b_{rl}u^b(v') - t^b_{rl}\big)\\
&- \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{1j}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{1j}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V\setminus\{v'\}}\alpha_{1j}(\hat v)+\epsilon}\right] - \epsilon\log\left[\frac{\epsilon}{\sum_{j=1}^{n}\sum_{\hat v\in V\setminus\{v'\}}\alpha_{1j}(\hat v)+\epsilon}\right]\\
&+ \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{1j}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{1j}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V\setminus\{v'\}}\alpha_{1j}(\hat v)}\right]\\
&- \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{rj}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)-\epsilon}\right] - \left(\sum_{j=1}^{n}\alpha_{rj}(v') - \epsilon\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v')-\epsilon}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)-\epsilon}\right]\\
&+ \sum_{v\in V\setminus\{v'\}} \left(\sum_{j=1}^{n} \alpha_{rj}(v)\right) \log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v)}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)}\right] + \left(\sum_{j=1}^{n}\alpha_{rj}(v')\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{rj}(v')}{\sum_{j=1}^{n}\sum_{\hat v\in V}\alpha_{rj}(\hat v)}\right].
\end{aligned}
$$

Define the function $\psi(\epsilon) \equiv \frac{1}{\epsilon}G(\epsilon)$:

$$\begin{aligned}
\psi(\epsilon) = &\left(q_{1l}^b u^b(v') - t_{1l}^b\right) - \left(q_{rl}^b u^b(v') - t_{rl}^b\right) - \log\left[\frac{\epsilon}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) + \epsilon}\right] \\
&- \sum_{v\in V\setminus\{v'\}} \left[\frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{1j}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{1j}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) + \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{1j}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{1j}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v})}\right]\right] \\
&- \sum_{v\in V\setminus\{v'\}} \left[\frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v}) - \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v})}\right]\right] \\
&- \left[\frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v') - \epsilon\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v') - \epsilon}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v}) - \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v')\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v')}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v})}\right]\right].
\end{aligned}$$

Let us introduce auxiliary function to ease the notation:

$$\begin{aligned}
\rho_1(\epsilon; v) &\equiv \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{1j}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{1j}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) + \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{1j}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{1j}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v})}\right], \\
\rho_r(\epsilon; v) &\equiv \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v}) - \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v)\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v})}\right], \\
\xi(\epsilon) &\equiv \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v') - \epsilon\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v') - \epsilon}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v}) - \epsilon}\right] - \frac{1}{\epsilon}\left(\sum_{j=1}^n \alpha_{rj}(v')\right)\log\left[\frac{\sum_{j=1}^n \alpha_{rj}(v')}{\sum_{j=1}^n \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v})}\right].
\end{aligned}$$

$\psi(\epsilon)$ can then be rewritten as:

$$\psi(\epsilon) \equiv \left(q_{1l}^b u^b(v') - t_{1l}^b\right) - \left(q_{rl}^b u^b(v') - t_{rl}^b\right) - \log\left[\frac{\epsilon}{\sum_{j=1}^n \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) + \epsilon}\right] - \sum_{v\in V\setminus\{v'\}} \rho_1(\epsilon; v) - \sum_{v\in V\setminus\{v'\}} \rho_r(\epsilon; v) - \xi(\epsilon).$$

We are now going to determine the right-limit of $\psi(\epsilon)$ as ϵ approaches zero. The following lemma holds:

Lemma B.4. $\lim_{\epsilon\to 0^+} \psi(\epsilon) = +\infty$.

Proof. Observe first that since $\sum_{j=1}^{n} \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) > 0$, we have:

$$\lim_{\epsilon\to 0^+} \left(\left(q_{1l}^{b} u^{b}(v') - t_{1l}^{b}\right) - \left(q_{rl}^{b} u^{b}(v') - t_{rl}^{b}\right) - \log\left[\frac{\epsilon}{\sum_{j=1}^{n} \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v}) + \epsilon}\right]\right) = +\infty.$$

It thus remains to show that the remaining terms converge to a finite value.

- Consider $\rho_1(\epsilon; v)$ for some $v \in V \setminus \{v'\}$. There are two possibilities:

  (i) If v is such that $\sum_{j=1}^{n} \alpha_{1j}(v) = 0$, then $\rho_1(\epsilon) = \frac{1}{\epsilon}\left[0\log 0 - 0\log 0\right] = \frac{0}{\epsilon} = 0$, hence $\lim_{\epsilon\to 0^+} \rho_1(\epsilon) = 0$.

  (ii) If v is such that $\sum_{j=1}^{n} \alpha_{1j}(v) > 0$, then

  $$\lim_{\epsilon\to 0^+} \rho_1(\epsilon) = -\frac{\sum_{j=1}^{n} \alpha_{1j}(v)}{\sum_{j=1}^{n} \sum_{\hat{v}\in V\setminus\{v'\}} \alpha_{1j}(\hat{v})},$$

  which is finite.

- Consider $\rho_r(\epsilon; v)$ for some $v \in V \setminus \{v'\}$. There are again two possibilities:

  (i) If v is such that $\sum_{j=1}^{n} \alpha_{rj}(v) = 0$, then $\rho_1(\epsilon) = \frac{1}{\epsilon}\left[0\log 0 - 0\log 0\right] = \frac{0}{\epsilon} = 0$, hence $\lim_{\epsilon\to 0^+} \rho_1(\epsilon) = 0$.

  (ii) If v is such that $\sum_{j=1}^{n} \alpha_{rj}(v) > 0$, then

  $$\lim_{\epsilon\to 0^+} \rho_r(\epsilon) = \frac{\sum_{j=1}^{n} \alpha_{rj}(v)}{\sum_{j=1}^{n} \sum_{\hat{v}\in V} \alpha_{rj}(\hat{v})},$$

  which is finite.

- Consider $\xi(\epsilon)$. Recall that by assumption $\sum_{j=1}^{n} \alpha_{rj}(v') > 0$, hence we have

$$\lim_{\epsilon \to 0^+} \xi(\epsilon) = \frac{\sum_{j=1}^{n} \alpha_{rj}(v')}{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{rj}(\hat{v})} - \log \left[\frac{\sum_{j=1}^{n} \alpha_{rj}(v')}{\sum_{j=1}^{n} \sum_{\hat{v} \in V} \alpha_{rj}(\hat{v})} \right] - 1,$$

  which is finite.

$\square$

Since $\lim_{\epsilon \to 0^+} \psi(\epsilon) = +\infty$, we conclude that for all $N > 0$ there exists $\epsilon > 0$ small enough such that $\psi(\epsilon) = \frac{1}{\epsilon} G(\epsilon) > N$, implying $G(\epsilon) > \epsilon N > 0$, implying that the constructed deviation $\tilde{\alpha}$ is profitable for all ϵ small enough, and thus contradicting the optimality of α.

$\square$

B.7 Proof of Lemma 3.6

Proof. We prove the statement of the lemma for the buyer only. The proof for the seller is almost identical.

"If". To establish the "if" direction of the claim we prove the contrapositive statement. Consider the following $k \times n$ information structure α given by:

State v	s_1^s	s_2^s	$\ldots$	s_n^s
s_1^b	$\alpha_{11}(v)$	$\alpha_{12}(v)$	$\ldots$	$\alpha_{1n}(v)$
s_2^b	$\alpha_{21}(v)$	$\alpha_{22}(v)$	$\ldots$	$\alpha_{2n}(v)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s_k^b	$\alpha_{k1}(v)$	$\alpha_{k2}(v)$	$\ldots$	$\alpha_{kn}(v)$

The buyer's payoff from this information structure is given by:

$$\begin{aligned}
&\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\alpha_{ij}(v)\big(q_{ij}^b u^b(v)-t_{ij}^b\big)-c^b(\alpha)\\
&=\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\alpha_{ij}(v)\big(q_{ij}^b u^b(v)-t_{ij}^b\big)-H(\mu_0)-\sum_{i=1}^{k}\sum_{v\in V}\left(\sum_{j=1}^{n}\alpha_{ij}(v)\right)\log\left[\frac{\sum_{j=1}^{n}\alpha_{ij}(v)}{\sum_{j=1}^{n}\sum_{\hat{v}\in V}\alpha_{ij}(\hat{v})}\right].
\end{aligned}$$

Suppose that α satisfies R-IC$_{\text{A}}^b$-1 but does not satisfy R-IC$_{\text{A}}^b$-2, then there exists a profitable deviation for the buyer which involves augmenting α with a an $k+1$-st signal realization $s_\emptyset^b$. This deviation has the following form:

State v	s_1^s	s_2^s	$\dots$	s_n^s
s_1^b	$\alpha_{11}(v)-\beta_{11}(v)$	$\alpha_{12}(v)-\beta_{12}(v)$	$\dots$	$\alpha_{1n}(v)-\beta_{1n}(v)$
s_2^b	$\alpha_{21}(v)-\beta_{21}(v)$	$\alpha_{22}(v)-\beta_{22}(v)$	$\dots$	$\alpha_{2n}(v)-\beta_{2n}(v)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
s_k^b	$\alpha_{k1}(v)-\beta_{k1}(v)$	$\alpha_{k2}(v)-\beta_{k2}(v)$	$\dots$	$\alpha_{kn}(v)-\beta_{kn}(v)$
$s_\emptyset^b$	$\sum_{i=1}^k \beta_{i1}(v)$	$\sum_{i=1}^k \beta_{i2}(v)$	$\dots$	$\sum_{i=1}^k \beta_{in}(v)$

where $\sum_{i=1}^k \sum_{j=1}^n \sum_{v\in V} \beta_{ij}(v) > 0$. The payoff from this deviation is given by:

$$\sum_{i=1}^k \sum_{j=1}^n \sum_{v\in V} \big(\alpha_{ij}(v)-\beta_{ij}(v)\big)\big(q_{ij}^b u^b(v) - t_{ij}^b\big) - H(\mu_0) - \sum_{i=1}^k \sum_{v\in V} \left(\sum_{j=1}^n \big(\alpha_{ij}(v)-\beta_{ij}(v)\big)\right) \log\left[\frac{\sum_{j=1}^n \big(\alpha_{ij}(v)-\beta_{ij}(v)\big)}{\sum_{j=1}^n \sum_{\hat{v}\in V} \big(\alpha_{ij}(\hat{v})-\beta_{ij}(\hat{v})\big)}\right]$$
$$-\sum_{v\in V}\left(\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)\right) \log\left[\frac{\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)}{\sum_{i=1}^k \sum_{j=1}^n \sum_{\hat{v}\in V} \beta_{ij}(\hat{v})}\right],$$

which can be rewritten as:

$$\sum_{i=1}^k \sum_{j=1}^n \sum_{v\in V} \big(\alpha_{ij}(v)-\beta_{ij}(v)\big)\big(q_{ij}^b u^b(v) - t_{ij}^b\big) - c^b(\alpha-\beta) - \sum_{v\in V}\left(\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)\right) \log\left[\frac{\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)}{\sum_{i=1}^k \sum_{j=1}^n \sum_{\hat{v}\in V} \beta_{ij}(\hat{v})}\right].$$

We now define the gain-from-deviation function as the difference between the payoff from the deviation and the payoff from α:

$$G_\alpha(\beta) \equiv -\sum_{i=1}^k \sum_{j=1}^n \sum_{v\in V} \beta_{ij}(v)\big(q_{ij}^b u^b(v) - t_{ij}^b\big) - c^b(\alpha-\beta) + c^b(\alpha) - \sum_{v\in V}\left(\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)\right) \log\left[\frac{\sum_{i=1}^k \sum_{j=1}^n \beta_{ij}(v)}{\sum_{i=1}^k \sum_{j=1}^n \sum_{\hat{v}\in V} \beta_{ij}(\hat{v})}\right].$$

Since the deviation under consideration is profitable, we have $G_\alpha(\beta) > 0$, We now define the function $\psi(\epsilon) \equiv \frac{1}{\epsilon} G_\alpha(\epsilon\beta)$ for $\epsilon > 0$.

Clearly $\psi(1) = G_\alpha(\beta) > 0$. $\psi(\epsilon)$ is written as:

$$\psi(\epsilon) = -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta_{ij}(v)\big(q_{ij}^{b}u^{b}(v) - t_{ij}^{b}\big) + \frac{c^{b}(\alpha-\epsilon\beta) - c^{b}(\alpha)}{-\epsilon} - \sum_{v\in V}\left(\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)\right)\log\left[\frac{\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)}{\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{\hat{v}\in V}\beta_{ij}(\hat{v})}\right].$$

We establish the following lemma:

Lemma B.5. *$\psi(\epsilon)$ is weakly decreasing.*

Proof. It is enough to show that $\frac{c^{b}(\alpha-\epsilon\beta)-c^{b}(\alpha)}{-\epsilon}$ is weakly decreasing. To that end, take $0 < \epsilon_1 < \epsilon_2 < 1$ and observe that $\alpha - \epsilon_1\beta = \big(1 - \frac{\epsilon_1}{\epsilon_2}\big)\alpha + \frac{\epsilon_1}{\epsilon_2}(\alpha - \epsilon_2\beta)$. Recall that $c^{b}(\alpha)$ is convex by Lemma 3.3, hence

$$c^{b}(\alpha-\epsilon_1\beta) \le \left(1 - \frac{\epsilon_1}{\epsilon_2}\right)c^{b}(\alpha) + \frac{\epsilon_1}{\epsilon_2}c^{b}(\alpha-\epsilon_2\beta) \Leftrightarrow \frac{c^{b}(\alpha-\epsilon_1\beta) - c^{b}(\alpha)}{-\epsilon_1} \ge \frac{c^{b}(\alpha-\epsilon_2\beta) - c^{b}(\alpha)}{-\epsilon_2}.$$

□

We now define the marginal gain-from-deviation $MG_\alpha(\beta) \equiv \lim_{\epsilon\to 0}\frac{1}{\epsilon}G_\alpha(\epsilon\beta)$. Recall that α satisfies R-IC$_{\mathrm{A}}^{b}$-1 by assumption, hence Lemma 3.5 ensures that all the posteriors induced by α are strictly positive, which in turn makes sure that the limit $\lim_{\epsilon\to 0}\frac{1}{\epsilon}G_\alpha(\epsilon\beta)$ is well-defined and given by:

$$\begin{aligned}MG_\alpha(\beta) = &-\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta_{ij}(v)\big(q_{ij}^{b}u^{b}(v) - t_{ij}^{b}\big) + \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta_{ij}(v)\log\left[\frac{\sum_{j=1}^{n}\alpha_{ij}(v)}{\sum_{j=1}^{n}\sum_{\hat{v}\in V}\alpha_{ij}(\hat{v})}\right]\\ &-\sum_{v\in V}\left(\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)\right)\log\left[\frac{\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)}{\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{\hat{v}\in V}\beta_{ij}(\hat{v})}\right],\end{aligned}$$

which can be rewritten as:

$$MG_\alpha(\beta) = -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta_{ij}(v)\big(q_{ij}^b u^b(v) - t_{ij}^b - \log\big(\mu_i^b(v)\big)\big) - \sum_{v\in V}\left(\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)\right)\log\left[\frac{\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)}{\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{\hat{v}\in V}\beta_{ij}(\hat{v})}\right].$$

Defining $y_{ij}^b(v) \equiv q_{ij}^b u^b(v) - t_{ij}^b - \log\big(\mu_i^b(v)\big)$, we can rewrite the marginal gain-from-deviation as

$$MG_\alpha(\beta) = -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta_{ij}(v)y_{ij}^b(v) - \sum_{v\in V}\left(\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)\right)\log\left[\frac{\sum_{i=1}^{k}\sum_{j=1}^{n}\beta_{ij}(v)}{\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{\hat{v}\in V}\beta_{ij}(\hat{v})}\right].$$

The following lemma holds:

Lemma B.6. *$MG_\alpha(\beta) > 0$.*

Proof. Recall that $MG_\alpha(\beta) = \lim_{\epsilon\to 0}\frac{1}{\epsilon}G_\alpha(\epsilon\beta) = \lim_{\epsilon\to 0}\frac{1}{\epsilon}\psi(\epsilon)$. By Lemma B.5, $\psi(\epsilon) \geq \psi(1)$ for every $0 < \epsilon < 1$, hence $\lim_{\epsilon\to 0}\psi(\epsilon) \geq \psi(1) > 0$. □

Let us now decompose the marginal deviation under consideration into two parts, β' and β'' by defining:

$$\beta'_{ij}(v) \equiv \begin{cases}\beta_{ij}(v) - \frac{1}{z_j(v)}\sum_{i=1}^{k}\beta_{ij}(v) & \text{if } \alpha_{ij}(v) > 0\\ 0, & \text{otherwise}\end{cases}; \qquad \beta''_{ij}(v) \equiv \begin{cases}\frac{1}{z_j(v)}\sum_{i=1}^{k}\beta_{ij}(v), & \text{if } \alpha_{ij}(v) > 0\\ 0, & \text{otherwise}\end{cases};$$

where $z_j(v)$ is the number of zero elements in the vector $[\alpha_{1j}(v), \ldots, \alpha_{kj}(v)]$. Observe that by construction we have for every j:

$$\sum_{i=1}^{k}\beta'_{ij}(v) = \sum_{i=1}^{k}\beta_{ij}(v) - z_j(v)\frac{1}{z_j(v)}\sum_{i=1}^{k}\beta_{ij}(v) = 0 \quad \text{and} \quad \sum_{i=1}^{k}\beta''_{ij}(v) = z_j(v)\frac{1}{z_j(v)}\sum_{i=1}^{k}\beta_{ij}(v) = \sum_{i=1}^{k}\beta_{ij}(v).$$

We can now rewrite the marginal gain in terms of β' and β'' as follows:

$$\begin{aligned} MG_\alpha(\beta) &= -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta'_{ij}(v)y^b_{ij}(v) - \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta''_{ij}(v)y^b_{ij}(v) - \sum_{v\in V}\left(\sum_{i=1}^{k}\sum_{j=1}^{n}\beta''_{ij}(v)\right)\log\left[\frac{\sum_{i=1}^{k}\sum_{j=1}^{n}\beta''_{ij}(v)}{\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{\hat{v}\in V}\beta''_{ij}(\hat{v})}\right] \\ &= -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta'_{ij}(v)y^b_{ij}(v) + MG_\alpha(\beta''). \end{aligned}$$

We establish the following lemma:

Lemma B.7. $MG_\alpha(\beta'') > 0$.

Proof. Observe that $-\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta'_{ij}(v)y^b_{ij}(v)$ is the directional derivative of the objective function in the constraint R-IC$^b_{\mathrm{A}}$-1 at α in the direction $-\beta'$. Since α satisfies the constraint R-IC$^b_{\mathrm{A}}$-1 by assumption and $-\beta'$ is a feasible direction in R-IC$^b_{\mathrm{A}}$-1, we must have $-\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta'_{ij}(v)y^b_{ij}(v) \leq 0$. Since $MG_\alpha(\beta) > 0$, we have $MG_\alpha(\beta'') > 0$. □

Defining $B''(v) = \sum_{i=1}^{k}\sum_{j=1}^{n}\beta''(v)$, we can rewrite the marginal gain of β'' at α as follows:

$$MG_\alpha(\beta'') = -\sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta''_{ij}(v)y^b_{ij}(v) - \sum_{v\in V}B''(v)\log\left[\frac{B''(v)}{\sum_{\hat{v}\in V}B''(\hat{v})}\right].$$

Recall that $y^b(v) = \min_{(i,j)|\alpha_{ij}>0}\{y^b_{ij}(v)\}$. The following lemma holds:

Lemma B.8. $-\sum_{v\in V}B''(v)y^b(v) - \sum_{v\in V}B''(v)\log\left[\frac{B''(v)}{\sum_{\hat{v}\in V}B''(\hat{v})}\right] > 0$.

Proof. Recall that all $\beta''_{ij}(v)$ are weakly positive by construction, and moreover are equal to zero whenever $\alpha_{ij}(v) = 0$. We therefore have $\sum_{v\in V}B''(v)y^b(v) \leq \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V}\beta''_{ij}(v)y^b_{ij}(v)$, implying $-\sum_{v\in V}B''(v)y^b(v) - \sum_{v\in V}B''(v)\log\left[\frac{B''(v)}{\sum_{\hat{v}\in V}B''(\hat{v})}\right] \geq MG_\alpha(\beta'')$, which together with the previous lemma establishes the claim. □

Dividing the expression in the lemma above by $\sum_{\hat{v}\in V} B''(\hat{v})$, we get

$$-\sum_{v\in V} \frac{B''(v)}{\sum_{\hat{v}\in V} B''(\hat{v})} y^b(v) - \sum_{v\in V} \frac{B''(v)}{\sum_{\hat{v}\in V} B''(\hat{v})} \log\left[\frac{B''(v)}{\sum_{\hat{v}\in V} B''(\hat{v})}\right] > 0.$$

Defining $P(v) \equiv \frac{B''(v)}{\sum_{\hat{v}\in V} B''(\hat{v})}$, we can rewrite the above inequality as:

$$-\sum_{v\in V} P(v) y^b(v) - \sum_{v\in V} P(v) \log\big(P(v)\big) > 0,$$

which clearly implies:

$$0 < \max_{\{P(v)\}_{v\in V}} \Big\{ -\sum_{v\in V} P(v) y^b(v) - \sum_{v\in V} P(v) \log\big(P(v)\big) \text{ s.t. } \sum_{v\in V} P(v) = 1,\ P(v) \geq 0\ \forall v \in V \Big\}.$$

To evaluate the right-hand-side, relax the non-negativity constraints and write down the Lagrangian of the relaxed problem:

$$\mathcal{L}\big(P;\nu\big) = -\sum_{v\in V} P(v) y^b(v) - \sum_{v\in V} P(v) \log\big(P(v)\big) - \nu\Big(\sum_{v\in V} P(v) - 1\Big).$$

Observe that the objective function in the relaxed problem is strictly concave and the feasible set is convex, implying that the first order conditions are necessary and sufficient for optimality. The optimality conditions are therefore given by:

$$\begin{cases} -y^b(v) - \log\big(P^*(v)\big) - 1 - \nu^* = 0 \quad \forall v \in V, \\ \sum_{v\in V} P^*(v) = 1. \end{cases}$$

The optimum is achieved at:

$$P^*(v) = \frac{\exp\big(-y^b(v)\big)}{\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)}.$$

We then have:

$$\begin{aligned}
& -\sum_{v\in V} P^*(v)y^b(v) - \sum_{v\in V} P^*(v)\log\big(P^*(v)\big) > 0 \\
\Leftrightarrow & -\sum_{v\in V} \frac{\exp\big(-y^b(v)\big)}{\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)} y^b(v) - \sum_{v\in V} \frac{\exp\big(-y^b(v)\big)}{\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)} \log\left(\frac{\exp\big(-y^b(v)\big)}{\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)}\right) > 0 \\
\Leftrightarrow & -\sum_{v\in V} \exp\big(-y^b(v)\big)y^b(v) - \sum_{v\in V}\exp\big(-y^b(v)\big)\log\left(\frac{\exp\big(-y^b(v)\big)}{\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)}\right) > 0 \\
\Leftrightarrow & -\sum_{v\in V} \exp\big(-y^b(v)\big)y^b(v) + \sum_{v\in V} \exp\big(-y^b(v)\big)y^b(v) + \sum_{v\in V}\exp\big(-y^b(v)\big)\log\left(\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)\right) > 0 \\
\Leftrightarrow & \sum_{v\in V}\exp\big(-y^b(v)\big)\log\left(\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)\right) > 0 \\
\Leftrightarrow & \log\left(\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big)\right) > 0 \\
\Leftrightarrow & \sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big) > 1,
\end{aligned}$$

which established the contrapositive claim.

"Only if". To establish the "only if" direction, we again prove the contrapositive statement. Suppose that α satisfies R-IC$^b_{\text{A}}$-1 and $\sum_{\hat{v}\in V}\exp\big(-y^b(\hat{v})\big) > 1$. The above calculations show that it is possible to construct profitable local deviation from α that involves augmenting α with a $k+1$-st signal realization. This deviation involves transferring probability mass to the $k+1$-st signal realization from those $\alpha_{rl}(v)$ that satisfy $(r,l) = \arg\max_{(i,j)|\alpha_{ij}(v)>0}\{y^b_{ij}(v)\}$ for each state $v \in V$. The profitability of this deviation implies that α violates R-IC$^b_{\text{A}}$-2.

□

B.8 Proof of Proposition 3.3

Proof. Once again we prove the statement of the proposition for the buyer only since the proof for the seller is virtually identical. The tuple (α, S^b, S^s) satisfies (R-IC$^b_{\mathrm{A}}$) if and only if both (R-IC$^b_{\mathrm{A}}$-1) and (R-IC$^b_{\mathrm{A}}$-2) are satisfied.

Recall the optimization problem from the incentive compatibility constraint (R-IC$^b_{\mathrm{A}}$-1) is given by:

$$\text{(R-IC}^b_{\mathrm{A}}\text{-1)} \quad \alpha \in \operatorname*{argmax}_{\tilde{\alpha}} \sum_{i=1}^{k}\sum_{j=1}^{n}\sum_{v\in V} \tilde{\alpha}_{ij}(v)\big(q^b_{ij}u^b(v) - t^b_{ij}\big) - c^b(\tilde{\alpha}), \quad \text{s.t.}$$
$$(1)\ \tilde{\alpha} \in \Delta\big(S^b \times S^s \times V\big);$$
$$(2)\ \operatorname{marg}_{S^s\times V}\tilde{\alpha} = \operatorname{marg}_{S^s\times V}\alpha.$$

Lemma 3.5 shows that if α satisfies (R-IC$^b_{\mathrm{A}}$-1), all posteriors must be strictly positive, implying that the objective function of the above optimization problem must be differentiable at α. Lemma 3.3 shows that the objective function in this optimization problem is concave. Since all the constraint functions are affine in α, these observations imply that the Karush-Kuhn-Tucker conditions are necessary and sufficient for optimality of α in (R-IC$^b_{\mathrm{A}}$-1). The Karush-Kuhn-Tucker conditions are written as follows:

$$(\mathrm{ST}^b) \quad q^b_{ij}u^b(v) - t^b_{ij} - \log\left(\mu^b_i(v)\right) - \lambda^b_j(v) + \phi^b_{ij}(v) = 0 \ \forall (i,j),\ v \in V;$$
$$(\mathrm{DF}^b) \quad \phi^b_{ij}(v) \geq 0,\ \forall (i,j),\ v\in V;$$
$$(\mathrm{CS}^b) \quad \alpha_{ij}(v)\phi^b_{ij}(v) = 0,\ v \in V.$$

Here ST stands for stationarity, DF stands for dual feasibility, and CS stands for complementary slackness; ϕ's are KKT-multipliers on the non-negativity constraints and λ's are KKT-multipliers on the equality constraints.

To obtain (NA^b), recall from Lemma 3.6 that as long as (R-IC$^b_{\mathrm{A}}$-1) is satisfied, (R-IC$^b_{\mathrm{A}}$-2) is equivalent to:

$$\sum_{v\in V} \exp\left(- \min_{(i,j)|\alpha_{ij}(v)>0} \left\{ q_{ij}(v)u^b(v) - t^b_{ij} - \log\left(\mu^b_i(v)\right) \right\} \right) \leq 1,$$

which, combined with condition (ST^b) is equivalent to:

$$\sum_{v\in V} \exp\big(-\min_{(i,j)|\alpha_{ij}(v)>0}\big\{\lambda_j^b(v)+\phi_{ij}^b(v)\big\}\big)\leq 1.$$

$(\mathrm{DF})^b$ and (CS^b) together imply that $\phi_{ij}^b(v)=0$ whenever $\alpha_{ij}(v)>0$, which implies that the above inequality is equivalent to

$$\sum_{v\in V} \exp\big(-\min_{(i,j)|\alpha_{ij}(v)>0}\big\{\lambda_j^b(v)\big\}\big)\leq 1.$$

Lemma 3.5 implies also that all the seller's posteriors must be positive, which means that in every column j there is at least one strictly positive $\alpha_{ij}(v)$ in every state v. We can therefore simply minimize over columns in the above inequality, hence

$$\sum_{v\in V} \exp\big(-\min_{j}\big\{\lambda_j^b(v)\big\}\big)\leq 1.$$

□

B.9 Proof of Lemma 3.7

Proof. Recall that the revenue-maximization problem is given by:

$$\max_{q,t,\lambda} \; t^b - t^s, \quad \text{s.t.}$$

$$(\mathrm{ST}^b) \quad q^b u^b(\underline{v}) - t^b - \log\big(\mu_0(\underline{v})\big) - \lambda^b(\underline{v}) = 0, \quad q^b u^b(\overline{v}) - t^b - \log\big(\mu_0(\overline{v})\big) - \lambda^b(\overline{v}) = 0;$$

$$(\mathrm{ST}^s) \quad t^s - q^s u^s(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \lambda^s(\underline{v}) = 0, \quad t^s - q^s u^s(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda^s(\overline{v}) = 0;$$

$$(\mathrm{NA}^b) \quad \exp(-\lambda^b\big(\underline{v})\big) + \exp(-\lambda^b\big(\overline{v})\big) \leq 1;$$

$$(\mathrm{NA}^s) \quad \exp(-\lambda^s\big(\underline{v})\big) + \exp(-\lambda^s\big(\overline{v})\big) \leq 1;$$

$$(q\text{-F}) \quad 0 \leq q^b \leq q^s \leq 1.$$

From $(\mathrm{ST})^b$ we have $-\lambda^b(\underline{v}) = t^b - q^b u^b(\underline{v}) + \log\big(\mu_0(\underline{v})\big)$ and $-\lambda^b(\overline{v}) = t^b - q^b u^b(\overline{v}) + \log\big(\mu_0(\overline{v})\big)$. Plugging the two expressions into (NA^b), we obtain:

$$\exp\big(t^b - q^b u^b(\underline{v}) + \log\big(\mu_0(\underline{v})\big)\big) + \exp\big(t^b - q^b u^b(\overline{v}) + \log\big(\mu_0(\overline{v})\big)\big) \leq 1,$$

which can be rewritten as:

$$t^b \leq -\log\big[\mu_0(\underline{v})\exp(-q^b u^b\big(\underline{v})\big) + \mu_0(\overline{v})\exp\big(-q^b u^b(\overline{v})\big)\big].$$

Likewise from $(\mathrm{ST})^s$ we have $-\lambda^s(\underline{v}) = q^s u^s(\underline{v}) - t^s + \log\big(\mu_0(\underline{v})\big)$ and $-\lambda^s(\overline{v}) = q^s u^s(\overline{v}) + \log\big(\mu_0(\overline{v})\big)$. Plugging the two expressions into (NA^s), we obtain:

$$\exp\big(q^s u^s(\underline{v}) - t^s + \log\big(\mu_0(\underline{v})\big)\big) + \exp\big(q^s u^s(\overline{v}) + \log\big(\mu_0(\overline{v})\big)\big) \leq 1,$$

which can be rewritten as:

$$t^s \geq \log\big[\mu_0(\underline{v})\exp\big(q^s u^s(\underline{v})\big) + \mu_0(\overline{v})\exp\big(q^s u^s(\overline{v})\big)\big].$$

The revenue maximization problem therefore simplifies to:

$$\begin{aligned}
&\max_{q,t}\ t^b - t^s, \quad \text{s.t.} \\
(\mathrm{NA}^b)\quad & t^b \leq -\log\left[\mu_0(\underline{v})\exp(-q^b u^b(\underline{v})) + \mu_0(\overline{v})\exp\left(-q^b u^b(\overline{v})\right)\right]; \\
(\mathrm{NA}^s)\quad & t^s \geq \log\left[\mu_0(\underline{v})\exp\left(q^s u^s(\underline{v})\right) + \mu_0(\overline{v})\exp\left(q^s u^s(\overline{v})\right)\right]; \\
(q\text{-F})\quad & 0 \leq q^b \leq q^s \leq 1;
\end{aligned}$$

which clearly implies that both (NA^b) and (NA^s) are binding at the optimum. The mechanism designer then solves:

$$\begin{aligned}
\max_{q^b,q^s}\ & -\log\left[\mu_0(\underline{v})\exp(-q^b u^b(\underline{v})) + \mu_0(\overline{v})\exp\left(-q^b u^b(\overline{v})\right)\right] - \log\left[\mu_0(\underline{v})\exp\left(q^s u^s(\underline{v})\right) + \mu_0(\overline{v})\exp\left(q^s u^s(\overline{v})\right)\right], \\
\text{s.t.}\ & 0 \leq q^b \leq q^s \leq 1.
\end{aligned}$$

Since the objective is increasing in q^b, the feasibility constraint $q^b \leq q^s$ is binding. The principal therefore solves:

$$\begin{aligned}
\max_{q}\ & -\log\left[\mu_0(\underline{v})\exp(-q u^b(\underline{v})) + \mu_0(\overline{v})\exp\left(-q u^b(\overline{v})\right)\right] - \log\left[\mu_0(\underline{v})\exp\left(q u^s(\underline{v})\right) + \mu_0(\overline{v})\exp\left(q u^s(\overline{v})\right)\right], \\
\text{s.t.}\ & 0 \leq q \leq 1.
\end{aligned}$$

□

B.10 Proof of Proposition 3.4

Proof. Consider any efficient and individually uninformative mechanism. Since we assume $u^s(\underline{v}) < u^b(\underline{v})$ and $u^s(\overline{v}) < u^b(\overline{v})$, efficiency implies $q^b_{ij} = q^s_{ij} = 1$ whenever a pair of signal realizations (s^b_i, s^s_j) occurs with strictly positive probability in at least one state.

Consider now any i, r and j such that $\alpha_{ij}(\underline{v}) > 0$ and $\alpha_{rj}(\overline{v}) > 0$. The buyer's stationarity and dual feasibility conditions imply:

$$
\begin{aligned}
(\mathrm{ST}^b_{ij}) \quad & u^b(\underline{v}) - t^b_{ij} - \log\big(\underline{\mu}_0(\underline{v})\big) - \lambda^b_j(\underline{v}) = 0, \\
& u^b(\overline{v}) - t^b_{ij} - \log\big(\mu_0(\overline{v})\big) - \lambda^b_j(\overline{v}) \leq 0,
\end{aligned}
$$

which in turn implies that $u^b(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \underline{\lambda}^b_j(\underline{v}) \geq u^b(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda^b_j(\overline{v})$, and

$$
\begin{aligned}
(\mathrm{ST}^b_{rj}) \quad & u^b(\underline{v}) - t^b_{rj} - \log\big(\mu_0(\underline{v})\big) - \lambda^b_j(\underline{v}) \leq 0, \\
& u^b(\overline{v}) - t^b_{rj} - \log\big(\mu_0(\overline{v})\big) - \lambda^b_j(\overline{v}) = 0,
\end{aligned}
$$

which in turn implies that $u^b(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \lambda^b_j(\underline{v}) \leq u^b(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda^b_j(\overline{v})$.

Combining the above expressions we conclude that $u^b(\underline{v}) - \log\big(\mu_0(\underline{v})\big) - \lambda^b_j(\underline{v}) \leq u^b(\overline{v}) - \log\big(\mu_0(\overline{v})\big) - \lambda^b_j(\overline{v})$ and hence $t^b_{ij} = t^b_{rj}$. Since rows i and r were arbitrarily chosen, we conclude that all transfers in column j are the same across all rows for which $\alpha_{ij}(\underline{v}) + \alpha_{ij}(\overline{v}) > 0$.

Observe that for any j we have:

$$
\exp\big(-\lambda^b_j(\underline{v})\big) + \exp\big(-\lambda^b_j(\overline{v})\big) \leq \exp\big(-\min_l\{\lambda^b_l(\underline{v})\}\big) + \exp\big(-\min_l\{\lambda^b_l(\overline{v})\}\big) \leq 1,
$$

where the second inequality is due to (NA^b). We therefore have

$$\exp\left(t^b_{ij} - u^b(\underline{v}) + \log(\mu_0(\underline{v}))\right) + \exp\left(t^b_{ij} - u^b(\overline{v}) + \log(\mu_0(\overline{v}))\right) \leq 1$$

for any (i,j) such that $\alpha_{ij}(\underline{v}) + \alpha_{ij}(\overline{v})$. The last inequality implies that any such t^b_{ij} can be upper-bound as follows:

$$t^b_{ij} \leq -\log\left[\mu_0(\underline{v})\exp\left(-u^b(\underline{v})\right) + \mu_0(\overline{v})\exp\left(-u^b(\overline{v})\right)\right].$$

The analogous argument for the seller implies that all t^s_{ij} in row i are the same across all columns j for which $\alpha_{ij}(\underline{v}) + \alpha_{ij}(\underline{v}) > 0$. We also have for the seller:

$$\exp\left(-\lambda^s_i(\underline{v})\right) + \exp\left(-\lambda^s_i(\overline{v})\right) \leq \exp\left(-\min_l\{\lambda^s_l(\underline{v})\}\right) + \exp\left(-\min_l\{\lambda^s_l(\overline{v})\}\right) \leq 1,$$

where the second inequality is due to (NA^s). We therefore have for any such t_{ij}:

$$\exp\left(u^s(\underline{v}) - t^s_{ij} + \log\left(\mu_0(\underline{v})\right)\right) + \exp\left(u^s(\overline{v}) - t^s_{ij} + \log\left(\mu_0(\overline{v})\right)\right) \leq 1,$$

for any (i,j) such that $\alpha_{ij}(\underline{v}) + \alpha_{ij}(\overline{v})$. The last inequality implies that any such t^s_{ij} can be lower-bound as follows:

$$t^s_{ij} \geq \log\left[\mu_0(\underline{v})\exp\left(u^s(\underline{v})\right) + \mu_0(\overline{v})\exp\left(u^s(\overline{v})\right)\right].$$

We can then write the revenue from our *individually uninformative* and *efficient* mechanism as:

$$
\begin{aligned}
&\sum_{i=1}^{m}\sum_{j=1}^{n}\big(\alpha_{ij}(\underline{v})+\alpha_{ij}(\overline{v})\big)(t^b_{ij}-t^s_{ij})\\
&\leq \Bigg(-\log\big[\mu_0(\underline{v})\exp\big(-u^b(\underline{v})\big)+\mu_0(\overline{v})\exp\big(-u^b(\overline{v})\big)\big]-\log\big[\mu_0(\underline{v})\exp\big(u^s(\underline{v})\big)+\mu_0(\overline{v})\exp\big(u^s(\overline{v})\big)\big]\Bigg)\underbrace{\sum_{i=1}^{m}\sum_{j=1}^{n}\big(\alpha_{ij}(\underline{v})+\alpha_{ij}(\overline{v})\big)}_{=1}\\
&=-\log\big[\mu_0(\underline{v})\exp\big(-u^b(\underline{v})\big)+\mu_0(\overline{v})\exp\big(-u^b(\overline{v})\big)\big]-\log\big[\mu_0(\underline{v})\exp\big(u^s(\underline{v})\big)+\mu_0(\overline{v})\exp\big(u^s(\overline{v})\big)\big]\\
&\leq \max_q\Big\{-\log\big[\mu_0(\underline{v})\exp\big(-u^b(\underline{v})\big)+\mu_0(\overline{v})\exp\big(-u^b(\overline{v})\big)\big]-\log\big[\mu_0(\underline{v})\exp\big(u^s(\underline{v})\big)+\mu_0(\overline{v})\exp\big(u^s(\overline{v})\big)\big],\quad \text{s.t.}\quad 0\leq q\leq 1\Big\}
\end{aligned}
$$

□

Bibliography

Abdulkadiroglu, A. and K. Chung (2004), "Auction design with tacit collusion." Working paper.

Abreu, D., D. Pearce, and E. Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58 (5), 1041–1063.

Athey, S., K. Bagwell, and C. Sanchirico (2004), "Collusion and price rigidity." *Review of Economic Studies*, 71 (2), 317–349.

Bergemann, D., B. Brooks, and S. Morris (2017), "Informationally robust optimal auction design." *Working paper*.

Bergemann, D. and J. Hörner (2018), "Should first-price auctions be transparent?" *American Economic Journal: Microeconomics*, 10 (3), 177–218.

Bergemann, D. and J. Välimäki (2002), "Information acquisition and efficient mechanism design." *Econometrica*, 70, 1007–1033.

Bergemann, D. and J. Välimäki (2019), "Dynamic mechanism design: An introduction." *Journal of Economic Literature*, 57 (2), 235–274.

Bikhchandani, S. (2010), "Information acquisition and full surplus extraction." *Journal of Economic Theory*, 145, 2282–2308.

Bikhchandani, S. and I. Obara (2017), "Mechanism design with acquisition of correlated information." *Economic Theory*, 63, 783–812.

Blackwell, D. (1965), "Discounted dynamic programming." *The Annals of Mathematical Statistics*, 36 (1), 226–235.

Brooks, B. and S. Du (2021), "Optimal auction design with common values: An informationally-robust approach." *Econometrica*, 89, 1313–1360.

Caplin, A. and M. Dean (2015), "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105, 2183–2203.

Carroll, G. (2019), "Robustness in mechanism design and contracting." *Annual Review of Economics*, 11, 139–166.

Chassang, S., K. Kawai, J. Nakabayashi, and J. Ortner (2021), "Robust screens for non-competitive bidding in procurement auctions." *Econometrica*. Forthcoming.

Che, Y.K. and J. Kim (2009), "Optimal collusion-proof auctions." *Journal of Economic Theory*, 144 (2), 565–603.

Condorelli, Daniele and Balázs Szentes (2020), "Information design in the holdup problem." *Journal of Political Economy*, 128, 681–709.

Correia-da Silva, J. (2017), "A survey on the theory of collusion under adverse selection in auctions and oligopolies." Working paper.

Cramton, P., R. Gibbons, and P. Klemperer (1987), "Dissolving a partnership efficiently." *Econometrica*, 55, 615–632.

Crémer, J. and R.P. McLean (1988), "Full extraction of the surplus in bayesian and dominant strategy auctions." *Econometrica*, 1247–1257.

Denti, T. (2020), "Unrestricted information acquisition." *Working paper.*

Forges, F. (1986), "An approach to communication equilibria." *Econometrica*, 54, 705–731.

Fudenberg, D., D. Levine, and E. Maskin (1994), "The folk theorem with imperfect public information." *Econometrica*, 62, 997–1040.

Gentzkow, M. and E. Kamenica (2017), "Bayesian persuasion with multiple senders and rich signal spaces." *Games and Economic Behavior*, 104, 411–429.

Gerardi, Dino and Leeat Yariv (2008), "Information acquisition in committees." *Games and Economic Behavior*, 62, 436–459.

Gershkov, Alex and Balázs Szentes (2009), "Optimal voting schemes with costly information acquisition." *Journal of Economic Theory*, 144, 36–68.

Gleyze, Simon and Agathe Pernoud (2020), "Infromationally simple incentives." *Working paper.*

Green, Jerry R. and Nancy L. Stokey (1978), "Two representations of information structures and their comparisons." *Technical Report No. 271. Institute for Mathematical Studies in Social Sciences, Stanford University.*

Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet (2010), "A commitment folk theorem." *Games and Economic Behavior*, 69, 127–137.

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

Kandori, M. and I. Obara (2006), "Towards a belief-based theory of repeated games with private monitoring, an application of POMDP." Working paper.

Krähmer, Daniel (2020), "Information disclosure and full surplus extraction in mechanism design." *Journal of Economic Theory*, 187.

Krähmer, Daniel (2021), "Information design and strategic communication." *American Economic Review: Insights*, 3, 51–66.

Larionov, Daniil, Hien Pham, Takuro Yamashita, and Shuguang Zhu (2022), "First best implementation with costly information acquisition." *Working paper.*

Li, Fei and Peter Norman (2018), "On bayesian persuasion with multiple senders." *Economics Letters*, 170, 66–70.

Li, Hao (2001), "A theory of conservatism." *Journal of political Economy*, 109, 617–636.

Liu, Q. (2015), "Correlation and common priors in games with incomplete information." *Journal of Economic Theory*, 148, 49–75.

Mailath, G. J. and L. Samuelson (2006), *Repeated Games and Reputations: Long-run relationships*. Oxford University Press, New York.

Matejka, P. and A. McKay (2015), "Rational inattention to discrete choices: A new foundation for the multinomial logit." *American Economic Review*, 105, 272–298.

McAfee, R. P. and J. McMillan (1992), "Bidding rings." *American Economic Review*, 82 (3), 579–599.

Mensch, J. (2022), "Screening inattentive agents." *American Economic Review (Forthcoming).*

Morris, S. and M. Yang (2021), "Coordination and continuous stochastic choice." Working paper.

Myerson, R. (1986), "Multistage games with communication." *Econometrica*, 54, 323–358.

Ortner, J., S. Chassang, J. Nakabayashi, and K. Kawai (2020), "Screening adaptive cartels." Working paper.

Outrata, Jiři, Michal Kočvara, and Jochem Zowe (1998), *Nonsmooth Approach to Optimization Problems with Equilibrium Constraints: Theory, Applications and Numerical Results*. Springer New York, NY.

Pavan, A., I. Segal, and J. Toikka (2014), "Dynamic mechanism design: a Myersonian approach." *Econometrica*, 82 (2), 601–653.

Persico, Nicola (2004), "Committee design with endogenous information." *The Review of Economic Studies*, 71, 165–191.

Peters, M. and C. Troncoso-Valverde (2013), "A folk theorem for competing mechanisms." *Journal of Economic Theory*, 148, 953–973.

Pomatto, L., P. Strack, and Omer Tamuz (2020), "The cost of information." *Working paper.*

Ravid, Doron, Roesler Anne-Katrin, and Balázs Szentes (2020), "Learning before trading: On the inefficiency of ignoring free information." *Working paper.*

Renault, J., L. Renou, , and T. Tomala (2014), "Secure message transmission on directed networks." *Games and Economic Behavior*, 85, 1–18.

Renou, L. and T. Tomala (2012), "Mechanism design and communication networks." *Theoretical Economics*, 7, 489–533.

Roesler, Anne-Katrin and Balázs Szentes (2017), "Buyer-optimal learning and monopoly pricing." *American Economic Review*, 107, 2072–2080.

Shannon, C. (1949), "Communication theory of secrecy systems." *Bell System Technical Journal*, 28, 656–715.

Shi, X. (2012), "Optimal auctions with information acquisition." *Games and Economic Behavior*, 74, 666–686.

Sims, C. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50, 665–690.

Skrzypacz, A. and H. Hopenhayn (2004), "Tacit collusion in repeated auctions." *Journal of Economic Theory*, 114 (1), 153–169.

Strulovici, B. (2021), "Can society function without ethical agents? an informational perspective." *Working paper.*

Terstiege, Stefan and Cédric Wasser (2022), "Competitive information disclosure to an auctioneer." *American Economic Journal: Microeconomics (Forthcoming).*

Thomas, C. J. (2005), "Using reserve prices to deter collusion in procurement competition." *Journal of Industrial Economics*, 53 (3), 301–326.

Wilson, R. (1969), "Competitive bidding with disparate information." *Management Science*, 15, 446–448.

Yamashita, T. (2018), "Revenue guarantee in auction with a (correlated) common prior and additional information." *Working paper.*

Yang, Kai Hao (2020), "A note on arbitrary joint distributions using partitions." *Note.*

Yang, M. (2015), "Coordination with flexible information acquisition." *Journal of Economic Theory*, 158, 721–738.

Zhang, W. (2021), "Collusion enforcement in repeated first-price auctions." *Theoretical Economics*. Forthcoming.

Zhao, Xin (2016), "Heterogeneity and unanimity: Optimal committees with information acquisition." *Working paper*.

Zhu, Shuguang (2021), "Private disclosure with multiple agents." *Working paper*.

Eidesstattliche Erklärung

Hiermit erkläre ich, dass ich die vorliegende Dissertation selbstständig angefertigt und die benutzten Hilfsmittel vollständig und deutlich angegeben habe.

Mannheim, 11.08.2022

Daniil Larionov

Curriculum Vitae

- PhD in Economics, 2015 - 2022, *University of Mannheim*
- MSc in Economics, 2015 - 2017, *University of Mannheim*
- BSc in Economics, 2009 - 2014, *Saint Petersburg State University*