Optimal Arbitration*

Tymofiy Mylovanov Andriy Zapechelnyuk†

Abstract

We study common arbitration rules for disputes of two privately informed parties, final offer and conventional arbitration. Conventional arbitration is shown to be an optimal arbitration rule in environments with transferable utility, while final offer arbitration is optimal if utility is non transferable and the parties' interests are not too aligned. These results explain prevalence of both arbitration rules in practice.

JEL classification: D74, D82

Keywords: Communication; asymmetric information; optimal arbitration; final offer arbitration; conventional arbitration; constant-threat arbitration

*We thank Ian Eckhout and two anonymous referees for comments and suggestions that greatly improved the paper. We are also grateful to Ethem Akyol, Attila Ambrus, Mark Amstrong, Vincent Crawford, Peter Esö, Christian Hellwig, Johannes Hörner, Navin Kartik, Daniel Krähmer, Vijay Krishna, Stephan Lauermann, Gilat Levy, Ming Li, David Martimort, Rony Razin, Larry Samuelson, and Satoru Takahashi for helpful comments.

†Manuscript received May 31, 2011; final revision received August 15, 2012.

1 Introduction

A common method of resolving disputes is compulsory arbitration. Arbitration is used in labor contracts, international business transactions, divorce and child custody, security regulations, and general commerce. Two commonly used arbitration rules are final offer arbitration and conventional arbitration. Under final offer arbitration, the parties who cannot agree on a solution to a dispute submit their proposals to the arbitrator who then chooses one of these proposals as the binding solution. Under conventional arbitration, the arbitrator is unrestricted in her choice of a solution given the parties' proposals.[1]

The early proponents of conventional arbitration in the United States viewed it as a superior alternative to costly strikes in labor disputes.[2,3] Nevertheless, conventional arbitration has been criticized for providing incentives for the parties to exaggerate their proposals in order to influence the arbitration outcome and the final offer arbitration procedure was suggested (Stevens 1966) as a means to pressure the disputants to make more reasonable offers. Both the critique and the solution were later questioned by the

[1] For a review of dispute resolution mechanisms, see Roberts (2007).

[2] An example is the enactment by the US Congress in 1963 of a compulsory arbitration statute to avoid a nationwide railroad strike (Stevens 1966).

[3] Arbitration is not a novel tool of dispute resolution. In Ancient Greece, final offer arbitration, for instance, was used during the trial of Socrates (Ashenfelter, Currie, Farber and Spiegel 1992), while conventional arbitration was prescribed, although not followed, as the method of conflict resolution in the Thirty Years Peace treaty between Athens and Sparta. See Roebuck (2001) for an account of arbitration practice in Ancient Greece.

results in Crawford (1979), Farber (1980), Brams and Merrill (1983), and Gibbons (1988) that conventional arbitration has a higher degree of proposal convergence than final offer arbitration, leaving an open question of the rationale for the widespread use of final offer arbitration in practice (Roberts 2007, Spier 2007).

In this paper, we provide an explanation for the prevalence of both arbitration rules: We show that conventional arbitration is optimal if the utility of the disputing parties is transferable, whereas final offer arbitration can be optimal if this is not the case. These results are in line with the observation that final offer arbitration is often used in labor market disputes, where one of the parties is either an employee or a labor union and thus might be risk-averse or wealth constrained, while conventional arbitration is commonly used for dispute resolutions between commercial companies who are more likely to be risk-neutral and unconstrained in payments to each other.

Our approach differs from that in the existing literature in two respects. First, we focus on welfare maximization rather than minimization of the disagreement rate. Second, we study environments with and without transferable utility.[4]

The superiority of final offer arbitration in environments *without* transferable utility is consistent with the fact that it has a higher degree of proposal divergence at the arbitration stage. The divergence of proposals and the restriction that the arbitrator must

[4]The arbitration literature has predominantly focused on the environment without transferable utility. See Crawford (1979), Farber (1980), Brams and Merrill (1983), and Gibbons (1988) and the references therein. For an example of analysis with transfers, see Brams and Merrill (1991).

pick a side and cannot modify the proposals creates inefficiency and stochasticity in the arbitration award. Crucially, this provides incentives for the parties to agree on a socially optimal outcome prior to the arbitration. By contrast, under conventional arbitration, the arbitrator is free to assign any award based on her inference from the parties' proposals, creating room for counterproductive attempts at strategic manipulation of the outcome by the disputing parties.

The ranking of conventional and final offer arbitration is reversed in the environments *with* transferable utility because transfers constitute an additional incentive tool. The flexibility in providing incentives through transfers and the flexibility in the choice of the arbitration outcome are complementary, rendering conventional arbitration a more attractive mechanism.

In our model, the parties have conflicting preferences over the arbitration award and are strategic. The socially optimal action is represented by an uncertain state. We model arbitration as a two-stage game. In the first stage, the parties privately observe some information about the state and make simultaneous proposals about the arbitration award. If the proposals coincide, it is implemented. Otherwise, the parties enter arbitration by making two new proposals to an arbitrator. The arbitrator observes the new proposals and chooses an arbitration award; we assume that the arbitrator is not a part of and does not observe the proposals made in the first stage. Under conventional arbitration, the arbitrator's choice is unrestricted. Under final offer arbitration, the action awarded must coincide with one of the proposals. The arbitrator is benevolent and her payoff is given by a utility function that is maximized at the socially optimal action. This model captures environments in which the disputing parties have the decision rights over the action but have conceded, in case of a disagreement, to use one of the arbitration procedures.

Before turning to the analysis of final offer and conventional arbitration, we take a look at the benchmark of optimal arbitration rules that maximize the expected welfare. By the revelation principle, optimal arbitration rules can be sought for among direct rules in which the parties report their information truthfully and the rule implements a lottery over actions and, possibly, a transfer contingent on the reports. Proposition 1 establishes that in the environment with transferable utility there exists an arbitration rule that implements a welfare maximizing action in each state. The construction of transfers, which ensures incentive compatibility of this rule, is done using standard mechanism design methods.[5]

A priori, implementing an optimal arbitration rule requires commitment. Providing incentives for the parties to report their information truthfully might require, at least sometimes, implementing an outcome that would not be optimal given the reports. By contrast, conventional arbitration assumes no commitment – the arbitrator is free to choose any award she considers desirable after observing the proposals. In Proposition 2, we observe that the optimal arbitration rule can be replicated through conventional arbitration in an equilibrium in which the parties report their information to the arbitra-

[5]The (ex-post) efficient outcome is feasible in our model because there is no outside option for the parties subjected to the arbitration rule and thus there are no individual rationality constraints.

tor, who imposes the optimal decision and the corresponding transfers. Commitment has no additional value in this environment because the arbitrator is impartial about monetary transfers between the parties. Given the optimal transfer schedule, it is optimal for the parties to report their information truthfully and it is optimal for the arbitrator to implement the socially optimal action.

The analysis of the environment without transfers is more difficult. We distinguish two environments, with complete and incomplete information. In complete information environments, the parties have identical information. This is the standard assumption made in the literature (e.g., Gibbons 1988).[6] This assumption is meant to capture environments in which there is a close relationship between the disputing parties as, for example, in case of a union and a company arguing about wage, business partners arguing over an intended interpretation of a contract, or a family in divorce proceedings deciding on the distribution of custody rights. In this environment, an optimal arbitration rule provides incentives for the parties to tell the truth through punishment of disagreements. In Proposition 3,

[6] The assumption is also standard in the literature on cheap talk communication between a decision maker and two informed agents in payoff environments similar to the one in this paper. It has been made, for example, in Gilligan and Krehbiel (1989), Krishna and Morgan (2001a, 2001b), Battaglini (2002), Levy and Razin (2007), Ambrus and Takahashi (2008), and Li (2008, 2010). The agents are imperfectly informed in the models of Austen-Smith (1993), Wolinsky (2002), and Battaglini (2004). Ambrus and Lu (2010) construct fully revealing equilibria robust to noise in cheap talk environments. See also Li and Suen (2009) for a survey of work on decision making in committees; this literature often assumes that different members of the committee hold distinct pieces of information.

we show that an optimal arbitration rule can be found among "constant-threat" rules in which every disagreement in parties' reports is punished by the *same* (stochastic) action with a two-point support. This result is deceptively simple; the surprising part is that the punishment is constant and independent of the exact nature of parties' reports to the arbitrator. The proof relies in a curious way on a minmax inequality and concavity of the parties' payoff functions. It proceeds by showing that in truthtelling equilibrium of any arbitration rule there exists a best deviation for each of the parties and that there exists a stochastic action that delivers a lower payoff than the best deviation simultaneously for both parties. Hence, truthful reporting is also optimal in the modified constant-threat rule in which any disagreement is followed by this action.

In Proposition 4, we show that the optimal arbitration rule can be implemented through final offer arbitration. This result requires an additional assumption that the preferences of the disputing parties are not too aligned. The idea is that the parties behave spitefully at the arbitration stage and make extreme proposals that minimize the payoff of their opponent. The arbitrator then randomizes between these proposals replicating the constant-threat punishment in the optimal arbitration rule and providing incentives for the parties to agree on an outcome prior to arbitration. Hence, final offer arbitration weakly outperforms conventional arbitration. Furthermore, in the environments in which the parties always prefer opposing extreme decisions, conventional arbitration performs especially poorly and is strictly inferior to final offer arbitration (Corollary 1).[7]

[7]For some recent studies of conventional and final offer arbitration in different environments, see

Intuitively, final offer arbitration provides the arbitrator with more commitment power to impose punishments for disagreements, as compared to conventional arbitration. This difference becomes relevant in the environments without transfers where it may not be feasible to implement socially optimal outcome in each state.

We study incomplete information environments in Section 4.4. Unfortunately, the analysis of optimal arbitration rules and final offer arbitration rule for arbitrary noise structures has proven intractable. We, therefore, focus on a restricted class of information structures. Specifically, we consider environments in which the support of parties' information is discrete and study performance of final offer arbitration in the limit as the information structure converges to that of complete information. In Proposition 6, we construct a constant-threat arbitration rule that is incentive compatible in the noisy environment and implements an outcome that converges to that of the optimal arbitration rule under complete information as the noise vanishes. In Proposition 7, we show that in spite of the noise this rule can be implemented through final offer arbitration. The crucial part of the proof is a construction of the incentives for the arbitrator to randomize between the parties' proposals after a disagreement. Unlike in the environment with complete information, disagreement is not an out-of-equilibrium event and we do not have the freedom of assigning out-of-equilibrium beliefs that make randomization optimal.[8] In

Hanany, Kilgour and Gerchak (2007), Olszewski (2011), and Yildiz (2011).

[8]As pointed out by Battaglini (2002) in the context of cheap talk communication between two perfectly informed agents and a decision maker, the equilibria in complete information environments might contain implausible out-of-equilibrium beliefs. This issue is avoided in the noisy environments.

equilibrium, the disputing parties randomize over their proposals in a manner that induce beliefs which make the arbitrator indifferent about approving of the proposals. We also need to provide incentives for the parties to randomize; this is possible by requiring the parties to randomize only in a subset of extreme states.

Our model of arbitration is related to that in Gibbons (1988), who studies conventional and final offer arbitration in a similar environment and, by contrast to our model, shows that conventional arbitration is superior to final offer arbitration. In Gibbons (1988), the arbitrator observes a noisy private signal about the state; the parties are risk-neutral, and their preferences are state-independent (we do not require the latter assumption, except in Section 4.3). Our results are different because we model final offer arbitration as a two-period dynamic interaction in which arbitration is preceded by negotiation, whereas there is no negotiation stage in Gibbons (1988). This allows separating disagreement at the negotiation stage from the punishment at the arbitration stage and gives the arbitrator more flexibility in providing incentives for the parties to agree on an optimal outcome.

A number of papers explore arbitration procedures different from final offer and conventional arbitration: combined arbitration (Brams and Merrill 1986), final offer arbitration with a bonus (Brams and Merrill 1991), double-offer arbitration (Zeng, Nakamura and Ibaraki 1996), amended final offer arbitration (Zeng 2003), closest-offer principle arbitration (Armstrong and Hurley 2002), etc.[9] This literature focuses on incremental improvement of the existing arbitration procedures and their relative performance; char-

[9]See Armstrong and Hurley (2002) for a review.

acterization of optimal arbitration rules is left open.

There is a connection between the model in this paper and the cheap talk literature with two privately informed agents and a decision maker (Krishna and Morgan 2001b, Battaglini 2002, Ambrus and Takahashi 2008, Ambrus and Lu 2010).[10] As we discuss it in Section 4, any outcome of cheap talk communication between the disputing parties and the arbitrator can be implemented through conventional arbitration, whereas the converse is not necessarily true, because the arbitrator cannot overrule the outcome if the parties agree prior to arbitration. The literature on cheap talk has focused on establishing conditions under which the decision maker can achieve the *first best* outcome of implementing the optimal action in each state (Krishna and Morgan 2001a, Krishna and Morgan 2001b, Battaglini 2002, Ambrus and Takahashi 2008).[11] We are interested in the performance of the specific arbitration procedures against the benchmark case in which the arbitrator has full commitment power, regardless of whether the first best outcome is implementable. The problem of optimal decision rules for two agents with private information has been studied in Martimort and Semenov (2008). Our models and approaches are quite different. In particular, they focus on agents who are biased in the same direction and consider dominant strategy implementation. Our paper is also related

[10]Crawford and Sobel (1982) is the seminal reference on cheap talk communication with one agent. For models of cheap talk communication with two agents see also Krishna and Morgan (2004), Battaglini (2004), Li (2008, 2010).

[11]The models in this literature are predominantly static. An exception is Esö and Fong (2010), who show that the first best outcome can be implemented in a dynamic cheap talk environment.

to Battaglini (2004) who considers a multidimensional environment with multiple agents and noisy signals. Battaglini shows that minimal commitment power is sufficient to implement an outcome arbitrarily close to the first best as the number of agents becomes sufficiently high.

The remainder of the paper is organized as follows. Section 2 describes the model. We study the environment with transfers in Section 3 and the environment without transfers in Section 4. The proof omitted in the text is in the Appendix.

2 The Model

There are two agents $i = 1, 2$ and a benevolent arbitrator $i = 0$. The agents have decision rights over an action from set $Y = [0, 1]$. In addition, the agents might be able to transfer utility by making a payment; let t denote the net payment from agent 1 to agent 2. We consider two environments, with unrestricted transfers, $t \in \mathbb{R}$, and without transfers, $t \equiv 0$.[12]

The parties have conflicting preferences and are strategic. In addition, they have unverifiable private information. For instance, in labor contract disputes over wages

[12] Transferable utility is a good assumption for environments in which the disputing parties are risk-neutral, liquidity unconstrained, and able to exchange monetary payments. In practice, there are multiple reasons that might make utility only partially transferable, such as risk-aversion, legal restrictions on payments, and limited liability. The environments with and without transferable utility, thus, represent the two polar benchmark cases.

and international trade disputes over quotas, their private information can describe their understanding of the original agreement upon wage or quota, while in a dispute about child custody the private information might reflect the parents' opinion about the relative allocation of custody rights that would be optimal for the well-being of the child.

Formally, each agent observes a private signal $x_i \in X_i \subseteq Y$. The signals are distributed according to a joint cumulative probability distribution F with the support on a subset of $X_1 \times X_2 \subseteq X = [0,1]^2$. Let $F_i(\cdot)$ denote the marginal of x_i and $F_{-i}(\cdot|x_i)$ denote the posterior cdf of x_{-i} conditional on x_i. We will also consider an environment with complete information in which the agents' signals are perfectly correlated, $x_1 = x_2$ for all $(x_1, x_2) \in \text{supp}(F)$.

The agents' payoffs are quasilinear functions

$$u_1(x_1, x_2, y) - t, \quad u_2(x_1, x_2, y) + t.$$

The arbitrator is benevolent and maximizes the utility function $u_0(x_1, x_2, y)$. The utility function u_0 captures the preferences over the state-dependent socially optimal action and can include interests of the third parties and the society at large. For instance, it might reflect the welfare of the child, ignoring the preferences of the parents, or it may be equal to a weighted sum of the functions u_i, where $i = 1, 2$, of the disputing parties. Note the assumption that the arbitrator is impartial about the allocation of transfers. This assumption is innocuous if there are no transfers, if the arbitrator is unconcerned about the payoffs of the disputing parties, or if their payoffs are given equal weight in the consideration of the arbitrator (e.g., $u_0(x_1, x_2, y) = u_1(x_1, x_2, y) - t + u_2(x_1, x_2, y) + t$).

If, however, the arbitrator were to care disproportionally about the utility of one of the parties, the assumption would impose a restriction.[13]

For the environments with complete information, we denote the payoff functions, with some abuse of notation, by $u_i(x, y)$. We assume that $u_i(x_1, x_2, y)$ is continuously differentiable on $X \times Y$, strictly concave in y, $i = 0, 1, 2$, and that the agent's payoff satisfies the single-crossing condition

$$\frac{\partial^2 u_i(x_1, x_2, y)}{\partial x_i \partial y} \geq 0, \quad i = 1, 2.$$

In addition, for each i function u_i has a unique maximizer $y_i^*(x_1, x_2)$ that is continuous and non-decreasing in its arguments.

Finally, the agents have opposing biases relative to the arbitrator,[14]

$$y_1^*(x_1, x_2) < y_0^*(x_1, x_2) < y_2^*(x_1, x_2), \quad \text{for almost all } (x_1, x_2) \in X \tag{OP$_1$}$$

For example, in the labor dispute one party may prefer the lowest possible wage while the other may prefer the highest possible wage, whereas in the child custody setting the parents may be biased in the direction of having more access to the child than would be in the interest of the child.

[13] With unequal weights, the arbitrator effectively has redistribution concerns and the optimal outcome is simply to assign an infinite transfer from one party to another, trivializing the role of choice of y.

[14] The standard assumption in the arbitration literature is that conflict of preferences is extreme: in each state, one party prefers the lowest action, whereas the other party prefers the highest action. Our model is more general and allows for non-trivial state-dependent preferences.

For some results we will also assume that the agents have a sufficient conflict of preferences in the sense that their preferences over extreme actions are opposite for all signal realizations,

$$u_1(x_1, x_2, 0) \geq u_1(x_1, x_2, 1) \;\text{ and }\; u_2(x_1, x_2, 0) \leq u_2(x_1, x_2, 1) \;\text{ for all } (x_1, x_2) \in X \quad (\mathrm{OP}_2)$$

Let $\mathcal{Y}$ denote the set of distributions on Y (randomized actions). A *direct arbitration rule with transfers* is a pair (μ, τ), where $\mu : X \to \mathcal{Y}$ is an action rule and $\tau : X \to \mathbb{R}$ is a transfer rule. A direct arbitration rule induces a game, in which after observing x_1 and x_2 the agents simultaneously make reports $\hat{x}_i \in X_i$ and action $\mu(\hat{x}_1, \hat{x}_2)$ and transfer $\tau(\hat{x}_1, \hat{x}_2)$ are implemented.

Denote by $U_i^\mu(x_i, \hat{x}_i)$ the expected utility of agent i under action rule μ if her signal is x_i and her report is $\hat{x}_i$, provided the other agent reports the truth, $\hat{x}_{-i} = x_{-i}$,

$$U_i^\mu(x_i, \hat{x}_i) \;=\; \int\limits_{x_{-i} \in Y} u_i(x_i, x_{-i}, \mu(\hat{x}_i, x_{-i})) \mathrm{d}F_{-i}(x_{-i}|x_i).$$

Similarly, let $\tau_i(\hat{x}_i)$ be the expected transfer from agent i to her opponent under transfer rule τ if she reports $\hat{x}_i$ and the other agent reports the truth,

$$\tau_1(\hat{x}_1) = \int_{x_2 \in Y} \tau(\hat{x}_1, x_2) dF_2(x_2|\hat{x}_1) \quad \text{and} \quad \tau_2(\hat{x}_2) = -\int_{x_1 \in Y} \tau(x_1, \hat{x}_2) dF_1(x_1|\hat{x}_2)$$

We consider Bayesian incentive compatible arbitration rules in which truthtelling, $\hat{x}_i = x_i$, is optimal for all realizations of signals, provided the opponent's reports are also truthful, i.e., for all $x, \hat{x} \in X_i$ and $i = 1, 2$,

$$U_i^\mu(x, x) - \tau_i(x) \geq U_i^\mu(x, \hat{x}) - \tau_i(\hat{x}). \tag{IC}$$

By the revelation principle, any equilibrium outcome of the agents' interaction in a game whose space of outcomes is a space of probability distributions over Y and transfers can be represented by the truthtelling equilibrium outcome in some incentive compatible arbitration rule.

A direct arbitration rule μ is *optimal* if it maximizes the expected payoff of the arbitrator,

$$v^{\mu} = \mathbb{E}u_0(x, \mu(x_1, x_2)),$$

among all incentive compatible direct arbitration rules. Since the set of incentive compatible direct arbitration rules is compact in weak topology and v^{μ} is continuous in μ, an optimal direct arbitration rule exists.

The revelation principle justifies our focus on truthtelling equilibria. Nevertheless, optimal arbitration rules could permit multiple equilibria.

3 Arbitration with Transfers

We start with an environment in which transfers are allowed, $\tau(x_1, x_2) \in \mathbb{R}$. Consider the following example with complete information environment, $x_1 = x_2 = x \in X_1 = X_2 = \bar{X}$, in which the arbitrator's payoff function is a weighted sum of the agents' payoffs net of transfers,

$$u_0(x, y) \equiv \gamma u_1(x, y) + (1 - \gamma) u_2(x, y), \quad \gamma \in (0, 1). \tag{1}$$

In this environment we can implement the arbitrator's most preferred action $y_0^*(x) =$

$y_0^*(x,x)$ for each $x \in \bar{X}$. To see the intuition, imagine that we pick some action $\hat{y}$ and assign it after any disagreement between the agents. We can now provide incentives to one of the agents to report his information truthfully by paying him the difference in his payoff from the arbitrator's optimal action $y_0^*(x)$ and what he can obtain by the deviation to $\hat{y}$. This payment must be charged to the other agent, making him the residual claimant and implying that this agent prefers to report the truth if and only if the *sum* of the agents' payoffs $u_1(x,y) + u_2(x,y)$ is greater under $y_0^*(x)$ rather than $\hat{y}$. If $\gamma = 1/2$, the result follows. For other values of γ, we need to find an appropriate $\hat{y}$. This can be done, as the following argument shows.

Let $\mu(x,x) = y_0^*(x)$ and $\mu(x_1,x_2) = \hat{y}$ whenever $x_1 \neq x_2$. The transfer is such that agent 2 is always indifferent whether to report the truth or not, specifically, we set $\tau(x,x) = u_2(x,\hat{y}) - u_2(x,y_0^*(x))$ and $\tau(x_1,x_2) = 0$ whenever $x_1 \neq x_2$. The only relevant incentive constraint is that of agent 1,

$$u_1(x,\mu(x,x)) - \tau(x,x) \geq u_1(x,\mu(\hat{x}_1,x)) - \tau(\hat{x}_1,x), \quad x, x_1 \in \bar{X},\ \hat{x}_1 \neq x.$$

Under this rule it is equivalent to

$$u_1(x,y_0^*(x)) + u_2(x,y_0^*(x)) \geq u_1(x,\hat{y}) + u_2(x,\hat{y}), \quad x \in \bar{X}. \tag{2}$$

Note that all this constraint requires is that the surplus from implementing the arbitrator's optimal action exceeds the surplus from threat action $\hat{y}$ in each state. By choosing $\hat{y}$ sufficiently extreme, we ensure that (2) is satisfied. Let $\gamma \leq 1/2$. Choose $\hat{y} = \max_{x \in \bar{X}} y_0^*(x)$ and denote by $\bar{y}(x)$ the maximizer of $u_1(x,y) + u_2(x,y)$. Observe that $y_0^*(x)$ coincides with $\bar{y}(x)$ for $\gamma = 1/2$ and with $y_2^*(x) = \arg\max u_2(x, y)$ for $\gamma = 0$. Then, by (OP$_1$) and continuity of u_0 w.r.t. y and γ and by the Maximum Theorem,

$$\bar{y}(x) \leq y_0^*(x) \leq y_2^*(x) \text{ for all } x \in \bar{X}.$$

But then (2) is satisfied, since $u_1(x, y) + u_2(x, y)$ is concave in y and is maximized at $\bar{y}(x) \leq y_0^*(x) \leq \hat{y} = \max_{x \in \bar{X}} y_0^*(x)$. For the case of $\gamma > 1/2$, we set the threat action to be $\hat{y} = \min_{x \in \bar{X}} y_0^*(x)$ and repeat the argument.

We now consider environments with imperfectly correlated signals and more general arbitrator's preferences. Let $\mu^*(x_1, x_2) = y_0^*(x_1, x_2)$ and

$$\tau^*(x_1, x_2) = \bar{\tau}_1(x_1) - \bar{\tau}_2(x_2),$$

where

$$\bar{\tau}_i(x) = U_i^\mu(x, x) - \int_0^x \left. \frac{\partial U_i^\mu(s, z)}{\partial s} \right|_{z=s} \mathrm{d}s - \int_{\tilde{x} \in Y} \left[U_i^\mu(\tilde{x}, \tilde{x}) - \int_0^{\tilde{x}} \left. \frac{\partial U_i^\mu(s, z)}{\partial s} \right|_{z=s} \mathrm{d}s \right] \mathrm{d}F_i(\tilde{x}).$$

Note that the expected transfer is normalized to be equal to 0,

$$\mathbb{E}\left[\bar{\tau}_i(x)\right] \equiv \int_{x \in Y} \bar{\tau}_i(x) \mathrm{d}F_i(x) = 0.$$

Proposition 1 *In the environments with transferable utility, there exists an incentive compatible arbitration rule that implements the arbitrator's preferred outcome, $y_0^*(x_1, x_2)$ for all $(x_1, x_2) \in X$.*

Proof. Substituting

$$U_i^\mu(x, \hat{x}) \quad = \quad U_i^\mu(\hat{x}, \hat{x}) + \int_{\hat{x}}^x \frac{\partial U_i^\mu(s, \hat{x})}{\partial s} \mathrm{d}s$$

and the expressions for transfers in (IC), we get that (IC) holds if and only if for all $x, \hat{x} \in X_i, i = 1, 2,$

$$\int_{\hat{x}}^{x} \left(\left. \frac{\partial U_i^{\mu}(s,z)}{\partial s} \right|_{z=s} - \frac{\partial U_i^{\mu}(s,\hat{x})}{\partial s} \right) \mathrm{d}s \;\; \geq \;\; 0.$$

Expanding this condition, we get

$$\int_{\hat{x}}^{x} \int_{x_{-i} \in Y} \left(\int_{\hat{x}}^{s} \frac{\partial^2 u_i(s, x_{-i}, y_0^*(z, x_{-i}))}{\partial s \partial y} \frac{\partial y_0^*(z, x_{-i})}{\partial z} \mathrm{d}z \right) \mathrm{d}F_{-i}(x_{-i}|x_i) \mathrm{d}s \;\; \geq \;\; 0,$$

which holds by the single-crossing property of the agents' utility functions and monotonicity of y_0^*. ■

The arbitration rule (μ^*, τ^*) that implements the arbitrator's most preferred action for each realization of signals can be implemented via conventional arbitration. We define *conventional arbitration* as a game in which both parties simultaneously and publicly make negotiation proposals $y_i \in Y$. If the parties agree on some action y, $y_1 = y_2 = y$, then that action is implemented. Otherwise, the arbitrator chooses an action y and a transfer τ; this choice should be sequentially rational given her equilibrium beliefs. For simplicity, we assume that the arbitrator does not observe the proposals made in the first stage; this assumption is not essential.[15]

In equilibrium of conventional arbitration game, the agents communicate their signals truthfully by proposing $y_i = x_i$, $i = 1, 2$. The arbitrator believes that the state is given by the agents' reports, implements action $\mu^*(y_1, y_2)$, and sets transfer $\tau^*(y_1, y_2)$. Truthful reporting is optimal by incentive compatibility of (μ^*, τ^*). Thus:

[15]The solution concept for this and other games considered in this paper is perfect Bayesian equilibrium.

Proposition 2 *In the environments with transferable utility, conventional arbitration can implement the arbitrator's preferred outcome, $y_0^*(x_1, x_2)$ for all $(x_1, x_2) \in X$.*

4 Arbitration without transfers

We now consider environments in which transfers are not allowed, $\tau(x_1, x_2) \equiv 0$ for all $(x_1, x_2) \in X$. We first present the results for the environment with complete information, $x_1 = x_2 = x \in X_1 = X_2 = \bar{X}$ and then consider environments with incomplete information. Throughout this section we assume that conditions (OP_1) and (OP_2) hold. Condition (OP_1) imposes structure on the direction of the biases of the disputing parties and is important for the results. Condition (OP_2) requires that the interests of the parties are not too aligned. It simplifies the characterization of the optimal arbitration rules in Section 4.1; the results, however, extend to the environments in which this condition does not hold. The condition is important for the results on the superiority of the final offer arbitration in Section 4.2. It is used to ensure optimality of the parties' "spiteful" behavior at the arbitration stage.

4.1 Constant-Threat Arbitration

Absent transfers, we need a different means to provide incentives for the agents to be truthful. In complete information environments, the agents have the same information about the state, and the truthful reports must be identical. In order to motivate each agent to agree with the other agent in a truthtelling equilibrium under an arbitration rule, the rule must punish disagreements. The difficulty here is that if a disagreement is observed, it is unclear which agent, if any, tells the truth. As a result, a punishment after a disagreement may depend non-trivially on the agents' reports.

Consider the following arbitration procedure called *constant-threat arbitration*. The agents learn the state and then simultaneously propose actions (y_1, y_2). If the agents agree on an action, $y = y_1 = y_2$, then y is implemented; if the agents disagree, then the arbitrator implements a constant *threat lottery* with support on extreme actions 0 and 1 which is *independent* of the proposed actions. We show that any optimal direct arbitration rule without transfers can be implemented through *constant-threat arbitration* with a properly chosen threat lottery.

Proposition 3 *Consider the environment with complete information and no transfers. Then, an optimal arbitration rule can be implemented via constant-threat arbitration.*

To see why this is true, consider an optimal direct arbitration rule μ. Recall that μ is incentive compatible, that is, for every state $x \in \bar{X}$, reporting the truth is a Nash equilibrium, $\hat{x}_1 = \hat{x}_2 = x$. First, observe that concavity of the agents' payoff functions implies that any lottery over actions implemented after a disagreement, $\mu(x_1, x_2)$, $x_1 \neq x_2$, can be replaced by some lottery $p_{\mu(x_1,x_2)}$ with support on $\{0, 1\}$ without affecting incentive compatibility.

If it so happens that the *same* replacement lottery p^* can be used for all lotteries

$\mu(x_1, x_2)$, the construction of the equivalent constant-threat arbitration rule is straightforward: We assign lottery p^* to any disagreement and ask the agents to propose the action $y_1(x) = y_2(x) = \mu(x, x)$. Then, incentive compatibility of μ implies that proposed strategies constitute an equilibrium in the constructed constant-threat arbitration rule.

We now prove that such a p^* exists. We have just argued that without loss of generality we can assume that any disagreeing reports x_1 and x_2 result in a lottery with support on $\{0, 1\}$. Let $P(x_1, x_2)$ denote the probability this lottery assigns on action 1. By (OP_2), agent 1's utility from a lottery on $\{0, 1\}$ that assigns probability p on 1 is decreasing in p, whereas agent 2's utility is increasing in p. Let $\underline{p} = \sup_x \inf_{x_1} P(x_1, x)$ and $\overline{p} = \inf_x \sup_{x_2} P(x, x_2)$. For any state x the minimal value of $P(x_1, x)$ that can be secured by agent 1's deviation in μ is weakly greater than $\underline{p}$, and by incentive compatibility of μ, agent 1 prefers action $\mu(x, x)$ to that lottery. Hence agent 1 prefers $\mu(x, x)$ to any constant-threat lottery that assigns probability $p^* \geq \underline{p}$ on action 1. The symmetric argument holds for agent 2 with $p^* \leq \overline{p}$. Finally, since maximin does not exceed minimax, there exists p^* such that $\underline{p} \leq p^* \leq \overline{p}$. The full proof is deferred to the Appendix.

Unlike in the environments with transferable utility, optimal arbitration rules might be unable to implement the arbitrator's most preferred action in each of the states. Nevertheless, optimal arbitration rules share some qualitative properties.

Let us normalize the arbitrator's bliss point to satisfy

$$y_0^*(x, x) = x, \quad x \in \bar{X}.$$

Consider an optimal constant-threat rule in which after a disagreement action 1 is implemented with probability p. By concavity of payoff functions, in state $x = p$, both agents prefer action $y = p$ to the threat lottery,

$$u_i(p,p) > pu_i(p,1) + (1-p)u_i(p,0).$$

This implies that an optimal rule implements the most preferred alternative for the arbitrator, $\mu(x,x) = x$, at least in state $x = p$. In addition, since the agents' payoff functions are strictly concave, we obtain $\mu(x,x) = x$ whenever x belongs to a proper interval containing $p \in (0,1)$.

Observation 1 *An optimal arbitration rule implements the socially optimal alternative of the arbitrator on an interval in Y.*

We now describe the structure of an optimal arbitration rule in states where the outcome differs from the arbitrator's most preferred action. For a given probability p of action 1, let $\tilde{X}_i^p$ be the set of states in which agent i strictly prefers the threat lottery to the socially optimal action,

$$\tilde{X}_i^p = \{x \in [0,1] \, : \, u_i(x,x) < \bar{u}_i(x,p)\},$$

where $\bar{u}_i(x,p)$ is agent i's expected payoff from the threat lottery p,

$$\bar{u}_i(x,p) = (1-p)u_i(x,0) + pu_i(x,1).$$

Hence, $\tilde{X}_1^p \cup \tilde{X}_2^p$ is the set of states where implementing arbitrator's most preferred action is not incentive compatible.

Observation 2 *For any state x in $\tilde{X}_1^p \cup \tilde{X}_2^p$, the incentive constraint of only one of the agents is violated, i.e., $\tilde{X}_1^p \cap \tilde{X}_2^p = \varnothing$.*

Proof. By (OP_1), $y_1^*(x,x) < y_0^*(x,x) \equiv x < y_2^*(x,x)$ for almost all $x \in Y$. If $p > x$, then agent 1 prefers action x to action $y = p$ and hence to the threat lottery. Otherwise, agent 2 prefers x to the threat lottery. Hence, at least one agent prefers x to the threat lottery. ■

Thus, an optimal constant-threat rule stipulates to choose action $\mu(x,x)$ that is the "closest" point to x (from the perspective of the arbitrator) subject to the incentive constraints for the agents. Since at every state $x \in \tilde{X}_i^p$ only agent i's incentive constraint is relevant, we obtain

$$\mu(x,x) \in \underset{y\,:\,u_i(x,y) \geq \bar{u}_i(x,p)}{\arg\max} \; u_0(x,y).$$

That is to say, the arbitrator will distort the implemented action, $\mu(x,x)$ in favor of the agent whose incentive constraint is binding, such that this agent is indifferent between $\mu(x,x)$ and the threat lottery.

4.2 Final Offer Arbitration

We now show that in the environment without transfers the optimal arbitration rule can be implemented via final offer arbitration. Let the agents simultaneously propose actions (y_1, y_2). If the agents agree, $y_1 = y_2$, then the agreed action is implemented. Otherwise, they enter arbitration by proposing actions (z_1, z_2) to the arbitrator who then chooses the one of these two actions that maximizes her utility w.r.t. her ex-post beliefs about the state. We assume that the original proposals (y_1, y_2) are not observed by the arbitrator.[16]

Final offer arbitration replicates the optimal constant-threat arbitration rule as follows. At every state x, in equilibrium the agents propose $y_1 = y_2 = \mu(x, x)$. After a disagreement agents behave spitefully and propose extreme actions, $z_1, z_2 \in \{0, 1\}$ that are least preferred by their opponent: agent 1 proposes action 0 and agent 2 proposes action 1 to the arbitrator. Then, if $(z_1, z_2) = (0, 1)$, the arbitrator implements the optimal threat lottery p^* on $\{0, 1\}$; otherwise she chooses the more extreme of the two proposed decisions with probability one. By Proposition 3, this strategy makes truthful reports at the settlement stage incentive compatible. Disagreement is out of equilibrium, hence we choose the beliefs of the arbitrator such that the above strategy is sequentially rational.

Proposition 4 *Consider the environment with complete information and no transfers. Then, an optimal arbitration rule can be implemented via final offer arbitration.*

Proof. Let μ be an optimal constant-threat arbitration rule, where p^* denotes the threat lottery with support on $\{0, 1\}$ that puts probability p^* on action 1. For every pair of arbitration proposals (z_1, z_2), let $z^* = \max\{z_1, z_2\}$ and $z_* = \min\{z_1, z_2\}$. Consider the

[16] This assumption makes the game and analysis simpler by restricting the structure of the arbitrator's information sets. The restriction is relevant for the environment with noise in which arbitration occurs on the equilibrium path. We expect that the results continue to hold in the environment in which the arbitrator observes the first stage proposals at the cost of the increased complexity of the agents' mixing strategy after observing extreme states (see Section 4.4).

following strategies. If $z^* \neq z_*$, the arbitrator chooses action z^* with probability

$$\pi_p(z_*, z^*) = \begin{cases} 0, & \text{if } z_* < 1 - z^*, \\ 1, & \text{if } z_* > 1 - z^*, \\ p, & \text{if } z_* = 1 - z^*. \end{cases}$$

In the negotiation stage, the parties propose $y_1 = y_2 = \mu(x, x)$. In the arbitration stage, the parties propose $z_1 = 0$ and $z_2 = 1$. By Proposition 3, these strategies implement the outcome of the optimal arbitration rule.

Furthermore, since arbitration is off the equilibrium path, there is freedom in assigning beliefs about x after a disagreement. To make the arbitrator's behavior a best response, we construct her beliefs by assigning probability q on 1 and probability $1-q$ on 0, in such a way that z_* is preferred to z^* (e.g., $q = 0$) if $z_* < 1 - z^*$ and z^* is preferred to z_* (e.g., $q = 1$) if $z_* > 1 - z^*$. Furthermore, if $z_* = 1 - z^*$, the concavity of the arbitrator's payoff function implies that there exists q such that

$$(1-q)u_0(0, z_*) + qu_0(1, z_*) = (1-q)u_0(0, z^*) + qu_0(1, z^*),$$

in which case the arbitrator is indifferent between choosing z_* and z^*, and hence any lottery is a best response.

To establish optimality of the agents' behavior, note that at the arbitration stage a deviation to any non-extreme action in $(0, 1)$ is ignored, and a deviation to the opposite extreme will lead to implementation of that extreme action with certainty, making the deviant weakly worse off by (OP_2) (say, if agent 1 deviates to action $z_1' = 1$, then $(z_1', z_2) =$

$(1, 1)$, so the arbitrator must implement 1, the least preferred outcome of agent 1). Finally, note that in the constant-threat arbitration rule each agent prefers $y = \mu(x, x)$ to the lottery outcome. ■

The key difference between this model and those in Farber (1980) and Gibbons (1988) is that we explicitly introduce the negotiation stage. Hence, the negotiation and the arbitration proposals become separated, which allows final offer arbitration to implement the outcome of the optimal arbitration rule.

4.3 Conventional Arbitration

We now consider conventional arbitration, in which both parties simultaneously and publicly make negotiation proposals $y_i \in Y$ and the arbitrator chooses an action y if and only if their proposals disagree. Note that conventional arbitration can implement any equilibrium of cheap talk game in which two parties simultaneously send messages to the arbitrator about their information, who then chooses an action that is sequentially rational given his posterior beliefs (Krishna and Morgan (2001b), Battaglini (2002)). The converse need not be true because the arbitrator cannot overrule an outcome if the parties' proposals agree.

We say that conventional arbitration is (weakly) *inferior* to final offer arbitration if the arbitrator's maximal expected payoff is (weakly) lower under conventional arbitration than under final offer arbitration. By Proposition 4:

Proposition 5 *Consider the environment with complete information and no transfers. Then, conventional arbitration is weakly inferior to final offer arbitration.*

Under conventional arbitration, only deterministic actions can be sequentially rational for the arbitrator, since her utility function is strictly concave. That is, punishment by randomized actions is impossible. So the ability of the arbitrator to provide incentives is substantially limited as compared to final offer arbitration, where the incentives are provided by a lottery over extreme actions.

We now consider a class of environments where conventional arbitration is strictly inferior. Suppose that

$$u_1(x, y) \text{ is strictly decreasing and } u_2(x, y) \text{ is strictly increasing in } y \text{ for all } x \in \bar{X}. \quad \text{(M)}$$

That is, irrespective of the state, the most preferred actions of the agents are the opposite extremes. Then:

Observation 3 *In complete information environment without transfers, where (M) holds, conventional arbitration implements a constant action in Y.*

Proof. W.l.o.g. we can consider equilibria in pure strategies in which the agents agree with probability one. Indeed, let $\tilde{y}$ be a stochastic outcome of some equilibrium at some state x. Modify the strategies in state x by making the agents propose the expected value of $\tilde{y}$ and keeping the rest of the strategies intact. Then, by concavity of payoff functions, the modified profile of strategies constitutes an equilibrium.

We now show that in every equilibrium the implemented action is state-independent. Consider an equilibrium where $y(z_1, z_2)$ denotes the arbitrator's action after disagreeing proposals (z_1, z_2). We now construct an auxiliary zero-sum game. The payoffs of agents 1 and 2 in this game are given by $-y(z_1, z_2)$ and $y(z_1, z_2)$, respectively. By (M), these payoffs represent the same ordinal preferences of the agents as their real payoffs at every state. Let y_p be the value of this game; its existence is implied by the existence of equilibrium in conventional arbitration. Then, y_p is the best action that can be secured by each agent i. Agent 1 can agree only on actions $y \geq y_p$ and agent 2 can agree only on actions $y \leq y_p$. Consequently, the only implementable action is y_p.[17] ■

Corollary 1 *Consider the environment with complete information and no transfers and let (M) hold. Then, conventional arbitration is strictly inferior to final offer arbitration.*

Proof. The proof follows from Observation 1 and Observation 3. ■

4.4 The Environment with Noisy Signals

We now return to the model in which the agents' information is incomplete. We assume that $X_1 = X_2$ is the discrete grid with step $1/K$ for some integer K:

$$X_1 = X_2 = \bar{X} = \{x^0, x^1, \ldots, x^{K-1}, x^K\} \tag{G}$$

[17] Note that every $y_p \in Y$ can be supported in equilibrium. As disagreement is out of equilibrium, we can set the arbitrator's posterior beliefs such that the agents' messages are ignored and y_p is the optimal action conditional on disagreement, so $y(z_1, z_2) = y_p$ for all (z_1, z_2).

where $x^0 = 0, x^K = 1, x^k - x^{k-1} = 1/K$ for all $k = 1, \ldots, K$, and $K \geq 1$.

We consider the case where the agents' signals are correlated and study the limit to complete information environment. The amount of noise in the agents' signals is measured by

$$\delta = \inf \{\varepsilon > 0 \,:\, \Pr [|x_1 - x_2| > \varepsilon] < \varepsilon\}.$$

The following notations are in order. Let $f_\delta(x_1, x_2)$ be the joint probability distribution over $\bar{X}^2$, where δ indicates the amount of noise. We study the limit $f_\delta(x_1, x_2) \to f_0(x_1, x_2)$ as $\delta \to 0$.[18] Let $\bar{f}_{i,\delta}(x_i) = \sum_{x_j \in \bar{X}} f_\delta(x_i, x_j)$ be the marginal probability that i receives signal x_i, and let $f_{i,\delta}(x_i|x_j) = f_\delta(x_i, x_j)/\bar{f}_{j,\delta}(x_j)$ be the probability that signal of i is x_i conditional on j having received x_j.

Denote by $U_\delta(\mu)$ the ex ante expected payoff of the arbitrator under rule μ in the environment with the amount of noise $\delta \geq 0$ assuming that the agents report their information truthfully,

$$U_\delta(\mu) = \sum_{x_1, x_2 \in \bar{X}} u_0(x_1, x_2, \mu(x_1, x_2)) f_\delta(x_1, x_2).$$

The next proposition asserts that there is no discontinuity in the performance of constant-threat arbitration rules in noiseless environment and noisy environment with small noise.

Proposition 6 *For every K, there exists a sequence of incentive compatible constant-threat rules $\hat{\mu}_\delta$ such that $U_\delta(\hat{\mu}_\delta) \to U_0(\mu_0)$ as $\delta \to 0$.*

[18] That is, we consider a convergent sequence of joint probability distributions $\{f_{\delta_n}\}$ such that $\delta_n \to 0$ and $f_{\delta_n} \to f_0$ pointwise as $n \to \infty$.

Proof. Let μ_0 be an optimal constant threat rule for the environment with zero noise, $\delta = 0$, and let p be the probability of action 1 in the constant threat lottery in this rule. Next, let $\hat{\mu}_\delta$ be an incentive compatible constant-threat rule in the environment with noise that has the same constant threat lottery, p, and

$$\hat{\mu}_\delta \in \arg\min_{\mu} \sum_{x \in \bar{X}} (\mu_0(x,x) - \mu(x,x))^2.$$

Such a rule exists since the feasible set of rules satisfying the constraints of the problem is not empty and contains, in particular, the constant threat rule that $\mu_\delta(x,x) = p$.

The incentive constraints for $\hat{\mu}_\delta$ are

$$f_{i,\delta}(x_i|x_i) u_i(x_i, x_i, \hat{\mu}_\delta(x_i, x_i)) + \sum_{x_j \neq x_i} f_{i,\delta}(x_j|x_i) \bar{u}_i(x_i, x_j, p) \tag{3}$$

$$\geq f_{i,\delta}(x_i'|x_i) u_i(x_i, x_i', \hat{\mu}_\delta(x_i', x_i')) + \sum_{x_j \neq x_i'} f_{i,\delta}(x_j|x_i) \bar{u}_i(x_i, x_j, p), \quad x_i, x_i' \in \bar{X}, \ i = 1, 2,$$

where $\bar{u}_i(x_1, x_2, p)$ denotes agent i's expected payoff from the threat lottery,

$$\bar{u}_i(x_1, x_2, p) = (1-p) u_i(x_1, x_2, 0) + p u_i(x_1, x_2, 1).$$

They can be, equivalently, rewritten as

$$u_i(x_i, x_i, \hat{\mu}_\delta(x_i, x_i)) - \bar{u}_i(x_i, x_i, p) \geq \frac{f_{i,\delta}(x_i'|x_i)}{f_{i,\delta}(x_i|x_i)} \left(u_i(x_i, x_i', \hat{\mu}_\delta(x_i', x_i')) - \bar{u}_i(x_i, x_i', p) \right)$$

for all $x_i, x_i' \in \bar{X}$.

Recall that μ_0 satisfies $u_i(x, x, \mu_0(x,x)) - \bar{u}_i(x, x, p) \geq 0$; furthermore, by Observation 2, this constraint is satisfied with slack for at least one agent. Therefore, since $\frac{f_{i,\delta}(x_i'|x_i)}{f_{i,\delta}(x_i|x_i)} \leq \frac{\delta}{1-\delta} \to 0$ as $\delta \to 0$, we have $\hat{\mu}_\delta \to \mu_0$ uniformly.

The result now follows from the continuity of u_0 in y, since the expected payoff of the arbitrator under rule $\hat{\mu}_\delta$ is equal to

$$U_\delta(\hat{\mu}_\delta) = \sum_{x \in \bar{X}} f_\delta(x, x) u_0(x, x, \hat{\mu}_\delta(x, x)) + \sum_{x_1 \neq x_2} f_\delta(x_1, x_2) \bar{u}_0(x_1, x_2, p).$$

■

In complete information environments, an optimal constant-threat arbitration rule μ_0 can be implemented via final offer arbitration. The difficulty of implementing the analogous incentive compatible rule $\hat{\mu}_\delta$ in the environment with noise $\delta > 0$ is that the arbitration stage is reached with a strictly positive probability, hence the arbitrator's beliefs cannot be arbitrary. Yet we can adjust strategies of the agents by letting them randomize their proposals in case of receiving extreme signals, $x_i \in \{0, 1\}$, that alter the posterior beliefs of the arbitrator such that she is indifferent between the two extreme proposals. The reason why at least one agent finds it optimal to randomize after receiving an extreme signal is that this agent's incentive compatibility constraint is binding and he is indifferent between the action after an agreement and the stochastic threat action after a disagreement.

Proposition 7 *For every K, there exists a sequence of equilibrium outcomes ρ_δ in final offer arbitration with the arbitrator's payoff $U_\delta(\rho_\delta)$ such that $U_\delta(\rho_\delta) \to U_0(\mu_0)$ as $\delta \to 0$.*

Proof. We replicate the final offer arbitration equilibrium construction of the constant-threat rules $\hat{\mu}_\delta$ in the proof of Proposition 6, with one modification. In the first stage, the agents make proposals $y_i = \hat{\mu}_\delta(x_i, x_i)$, except for $i = 1$ when $x_1 = 1$, or $i = 2$ when

$x_2 = 0$. We will describe the remainder of the agents' strategies in the first stage below.

At the second stage, on the equilibrium path, agents 1 and 2 propose $z_1 = 0$ and $z_2 = 1$, and the arbitrator randomizes between 0 and 1 with probabilities $(1-p, p)$. If either agent deviates at the second stage, this is out-of-equilibrium behavior, so the beliefs of the arbitrator are exactly as in the proof of Proposition 4. Otherwise, the beliefs of the arbitrator are given by the Bayes rule.

To construct the agents' behavior at the settlement stage for the extreme signals, consider the hypothetical environment in which the agents propose $y_i(x_i) = \hat{\mu}_\delta(x_i, x_i)$ for all $x_i \in X_i$, including the extreme signals. Let H_δ be the probability of *disagreement*, $y_1 \neq y_2$,

$$H_\delta = 1 - \sum_{x_1, x_2 \in \bar{X}} \mathbf{1}_{\{\hat{\mu}_\delta(x_1,x_1) = \hat{\mu}_\delta(x_2,x_2)\}} f_\delta(x_1, x_2) \leq 1 - \sum_{x \in \bar{X}} f_\delta(x, x) \leq \delta.$$

The posterior beliefs of the arbitrator conditional on a disagreement in this environment are given by the joint probability distribution h_δ:

$$h_\delta(x_1, x_2) = \begin{cases} \frac{f_\delta(x_1,x_2)}{H_\delta}, & \hat{\mu}_\delta(x_1, x_1) \neq \hat{\mu}_\delta(x_2, x_2) \\ 0, & \text{otherwise.} \end{cases}$$

There is no *a priori* reason for the arbitrator to be indifferent between actions 0 and 1 given these beliefs. Suppose that given h_δ the arbitrator prefers action 1 to action 0. (The construction for the other case is symmetric.)

Case 1. Assume that in $\hat{\mu}_\delta$ satisfies agent 2's incentive compatibility constraint with equality for $x_2 = 0$. Then, if $x_2 = 0$, in the first stage the agent proposes $y_2 = \hat{\mu}_\delta(0, 0)$ with probability $1-t$ and randomizes uniformly among all actions in $Y \backslash \hat{\mu}_\delta(0,0)$ with probability t. Otherwise, if $x_2 > 0$, the agent proposes $y_2(x_2) = \hat{\mu}_\delta(x_2, x_2)$. Agent 1's proposal strategy is $y_1(x_1) = \hat{\mu}_\delta(x_1, x_1)$ for all $x_1 \in X_1$, including the extreme signals. The optimality of the agent 2's behavior given signal $x_2 = 0$ follows from incentive compatibility of $\hat{\mu}_\delta$ and our assumption that agent 2's incentive constraint is binding at $x_2 = 0$.

Now, we construct the value of t that makes it optimal for the arbitrator to randomize after a disagreement on the equilibrium path. The arbitrator's beliefs after disagreement that are induced by these strategies are given by

$$\hat{h}_{\delta,t}(x_1, x_2) = \begin{cases} \frac{f_\delta(x_1,x_2)}{\hat{H}_{\delta,t}}, & \hat{\mu}_\delta(x_1, x_1) \neq \hat{\mu}_\delta(x_2, x_2) \\ \frac{t f_\delta(0,0)}{\hat{H}_{\delta,t}}, & x_1 = x_2 = 0, \\ 0, & \text{otherwise,} \end{cases}$$

where the probability of disagreement, $\hat{H}_{\delta,t}$, is equal to

$$\hat{H}_{\delta,t} = 1 - \sum_{x_1, x_2 \in \bar{X}} \mathbf{1}_{\{\hat{\mu}_\delta(x_1,x_1) = \hat{\mu}_\delta(x_2,x_2)\}} f_\delta(x_1, x_2) + t f_\delta(0,0) \leq \delta + t f_\delta(0,0).$$

Observe that for any given $t > 0$, $\frac{t f_\delta(0,0)}{\hat{H}_{\delta,t}} \to 1$ as $\delta \to 0$. Thus, for any fixed and sufficiently small $\delta > 0$ the arbitrator prefers action 0 if $t = 1$, whereas, by assumption, she prefers action 1 if $t = 0$. Consequently, there exists $t^*(\delta) \in (0,1)$ such that the arbitrator is indifferent about actions 0 and 1 if we set $t = t^*(\delta)$. Furthermore, $t^*(\delta) \to 0$ as $\delta \to 0$ since the probability of disagreement because of noise converges to 0 as noise vanishes.

Case 2. Assume that $\hat{\mu}_\delta$ satisfies agent 2's incentive compatibility constraint with strict inequality for $x_2 = 0$. By $y_0^*(x, x) = x$ and (OP$_1$), it must be that $\hat{\mu}_\delta(0,0) = 0$. Then, by

(OP_2) agent 2 cannot strictly prefer $y_2 = \hat{\mu}_\delta(0,0)$ to the threat lottery, a contradiction. $\blacksquare$

In reference to conventional arbitration, by the same argument as in Section 4.3, we have:

Proposition 8 *If (M) holds, conventional arbitration is strictly inferior to final offer arbitration for every positive amount of noise* $\delta > 0$.

Tymofiy Mylovanov: Department of Economics, University of Pennsylvania, 456 McNeil Bld, 3718 Locust Walk, Philadelphia, PA 19104, USA. Email: mylovanov@gmail.com

Andriy Zapechelnyuk: School of Economics and Finance, Queen Mary, University of London, Mile End Road, London E1 4NS, UK. Email: a.zapechelnyuk@qmul.ac.uk

Appendix: Proof of constant-threat optimality

Proof of Proposition 3. Let μ be an optimal arbitration rule. Observe that by concavity of $u_i(x,y)$ in y, $i=1,2$, for any measure λ,

$$\int u_i(x,y)\lambda(dy) \geq \left(1-\int y\lambda(dy)\right)u_i(x,0)+\left(\int y\lambda(dy)\right)u_i(x,1), \quad x\in\bar{X}.$$

Hence, replacing $\mu(x_1,x_2)$, $x_1\neq x_2$, by a lottery that puts probability $\int y\mu(x_1,x_2)(dy)$ on action 1 and the complementary probability on action 0 will not violate the incentive constraints of the agents. Therefore, there exists an equivalent arbitration rule μ' in which every threat lottery implemented after a disagreement has support on $\{0,1\}$.

We now show that there exists a constant-threat arbitration rule μ^c equivalent to μ'. For every pair of different reports, $x_1,x_2\in\bar{X}, x_1\neq x_2$, let $P(x_1,x_2)$ be the probability that $\mu'(x_1,x_2)$ assigns to 1 after a disagreement. We extend the definition of $P(\cdot,\cdot)$ to $\bar{X}^2$ by setting $P(x,x)=\int y\mu'(x,x)(dy)$ for all $x\in\bar{X}$. Define

$$\mathcal{P}_1(x)=\{P(x',x)|x'\in\bar{X}\} \text{ and } \mathcal{P}_2(x)=\{P(x,x')|x'\in\bar{X}\}.$$

For all $x\in\bar{X}$, $p\in[0,1]$, and $i=1,2$ let

$$D_i(x,p)=\max\{0,pu_i(x,1)+(1-p)u_i(x,0)-u_i(x,\mu(x,x))\}.$$

By construction, a deviation by agent i in state x leading to a lottery in $\mathcal{Y}^*$ that assigns probability $p\in[0,1]$ to action 1 is non-profitable iff $D_i(x,p)=0$. Furthermore, by definition of $P(x,x)$,

$$D_i(x,P(x,x))=0, \quad x\in\bar{X}, i=1,2.$$

Thus, incentive constraints (IC) can be written as

$$D_i(x,p) = 0, \quad x \in \bar{X},\ p \in \mathcal{P}_i(x),\ i = 1,2. \tag{IC$'$}$$

Observe that by (OP$_2$)

$$\begin{aligned} &D_1(x,p) \text{ is non-increasing in } p \text{ for every } x \in \bar{X}; \\ &D_2(x,p) \text{ is non-decreasing in } p \text{ for every } x \in \bar{X}. \end{aligned} \tag{4}$$

Let

$$\begin{aligned} a_1(x) &= \inf \mathcal{P}_1(x), \quad x \in \bar{X}; \\ a_2(x) &= \sup \mathcal{P}_2(x), \quad x \in \bar{X}. \end{aligned}$$

By (IC) and continuity of $D_i(x,p)$ w.r.t. p, we have $D_i(x, a_i(x)) = 0$ for $x \in \bar{X}$. By (4),

$$\begin{aligned} &D_1(x,p) = 0, \quad p \geq a_1(x),\ x \in \bar{X}; \\ &D_2(x,p) = 0, \quad p \leq a_2(x),\ x \in \bar{X}. \end{aligned} \tag{5}$$

Define

$$\begin{aligned} \underline{p} &= \sup_{x\in\bar{X}} a_1(x) = \sup_{x\in\bar{X}} \inf \mathcal{P}_1(x) = \sup_{x\in\bar{X}} \inf_{x'\in\bar{X}} P(x',x); \\ \overline{p} &= \inf_{x\in\bar{X}} a_2(x) = \inf_{x\in\bar{X}} \sup \mathcal{P}_2(x) = \inf_{x'\in\bar{X}} \sup_{x\in\bar{X}} P(x',x). \end{aligned}$$

Then, there exists p^c such that $\underline{p} \leq p^c \leq \overline{p}$. By (5),

$$D_i(x,p^c) = 0, \quad x \in \bar{X}, i = 1,2.$$

The result now follows from (IC$'$). ■

References

Ambrus, Atilla and Shih-En Lu, "Robust fully revealing equilibria in multi-sender cheap talk," mimeo, Harvard University, 2010.

Ambrus, Attila and Satoru Takahashi, "The multi-sender cheap talk with restricted state spaces," *Theoretical Economics* 3 (2008), 1–27.

Armstrong, Michael J. and W. J. Hurley, "Arbitration using the closest offer principle of arbitrator behavior," *Mathematical Social Sciences* 43 (2002), 19–26.

Ashenfelter, Orley, Janet Currie, Henry S. Farber, and Matthew Spiegel, "An Experimental Comparison of Dispute Rates in Alternative Arbitration Systems," *Econometrica* 60 (1992), 1407–1433.

Austen-Smith, David, "Interested Experts and Policy Advice: Multiple Referrals under Open Rule," *Games and Economic Behavior* 5 (1993), 3–43.

Battaglini, Marco, "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica* 70 (2002), 1379–1401.

Battaglini, Marco, "Policy Advice with Imperfectly Informed Experts," *Advances in Theoretical Economics* 4 (2004), Article 1.

Brams, Steven J. and Samuel III Merrill, "Equilibrium Strategies for Final-Offer Arbitration: There is no Median Convergence," *Management Science* 29 (1983), 927–941.

Brams, Steven J. and Samuel III Merrill, "Binding Versus Final-Offer Arbitration: A Combination is Best," *Management Science* 32 (1986), 1346–1355.

Brams, Steven J. and Samuel III Merrill, "Final-offer arbitration with a bonus," *European Journal of Political Economy* 7 (1991), 79–92.

Crawford, Vincent P, "On Compulsory-Arbitration Schemes," *Journal of Political Economy* 87 (1979), 131–159.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica* 50 (1982), 1431–1451.

Esö, Peter and Yuk-fai Fong, "Wait and See: A Theory of Communication over Time," mimeo, University of Oxford, 2010.

Farber, Henry S., "An Analysis of Final-Offer Arbitration," *Journal of Conflict Resolution* 24 (1980), 683–705.

Gibbons, Robert, "Learning in Equilibrium Models of Arbitration," *American Economic Review* 78 (1988), 896–912.

Gilligan, T. and K. Krehbiel, "Asymetric information and legislative rules with a heterogeneous committee," *American Journal of Political Science* 33 (1989), 459–490.

Hanany, Eran, D. Marc Kilgour, and Yigal Gerchak, "Final-Offer Arbitration and Risk Aversion in Bargaining," *Management Science* 53 (2007), 1785–1792.

Krishna, Vijay and John Morgan, "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review* 95 (2001), 435–452.

Krishna, Vijay and John Morgan, "A Model of Expertise," *Quarterly Journal of Economics* 116 (2001), 747–775.

Krishna, Vijay and John Morgan, "The Art of Conversation: Eliciting Information from Experts Through Multi-Stage Communication," *Journal of Economic Theory* 117 (2004), 147–179.

Levy, Gilat and Ronny Razin, "On the limits of communication in multidimensional cheap talk: a comment," *Econometrica* 75 (2007), 885–893.

Li, Hao and Wing Suen, "Viewpoint: Decision-Making in Committees," *Canadian Journal of Economics* 42 (2009), 359–392.

Li, Ming, "Two (talking) heads are not better than one," *Economics Bulletin* 63 (2008), 1–8.

Li, Ming, "Advice from Multiple Experts: A Comparison of Simultaneous, Sequential, and Hierarchical Communication," *The B.E. Journal of Theoretical Economics (Topics)* 10 (2010), Article 18.

Martimort, David and Aggey Semenov, "The informational effects of competition and collusion in legislative politics," *Journal of Public Economics* 92 (2008), 1541–1563.

Olszewski, Wojciech, "A welfare analysis of arbitration," *American Economic Journal: Microeconomics* 3 (2011), 174–213.

Roberts, Gary E., "Impasse Resolution Procedures," in Evan M. Berman, ed., *Encyclopedia of Public Administration and Public Policy*, 2nd edition (CRC Press, 2007), pp. 987–991.

Roebuck, Derek, *Ancient Greek Arbitration* (Holo Books, 2001).

Spier, Kathryn E., "Litigation," in A. Mitchel Polinsky and Stevem Shavell, eds., *Handbook of Law and Economics*, Vol. I (North-Holland, 2007).

Stevens, Carl M., "Is Compulsory Arbitration Compatible With Bargaining?" *Industrial Relations* 5 (1966), 38–52.

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

Yildiz, Muhamet, "Nash meets Rubinstein in final-offer arbitration," *Economics Letters* 110 (2011), 226–230.

Zeng, Dao-Zhi, "An amendment to final-offer arbitration," *Mathematical Social Sciences* 46 (2003), 9–19.

Zeng, Dao-Zhi, Shinya Nakamura, and Toshihide Ibaraki, "Double-offer arbitration," *Mathematical Social Sciences* 31 (1996), 147–170.