A Simple Bargaining Mechanism that Elicits Truthful Reservation Prices

Steven J. Brams, Todd R. Kaplan & D. Marc Kilgour

Group Decision and Negotiation
Published in cooperation with the Institute for Operations Research and the Management Sciences and its Section on Group Decision and Negotiation

ISSN 0926-2644

Group Decis Negot
DOI 10.1007/s10726-014-9395-5

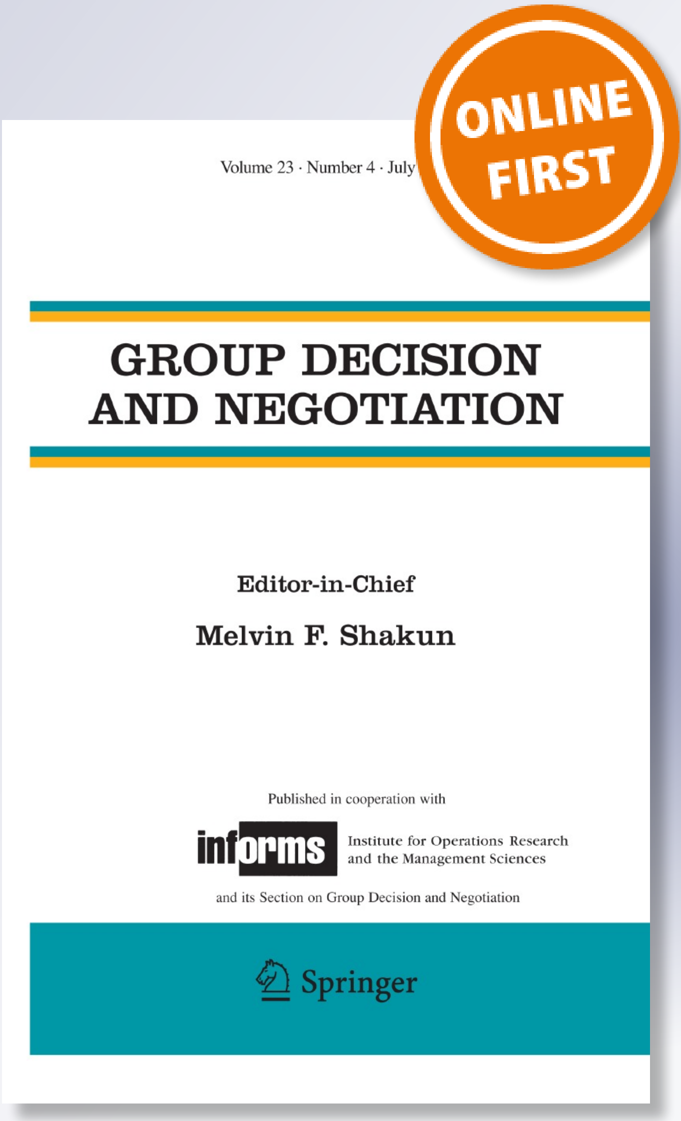

Group Decis Negot
DOI 10.1007/s10726-014-9395-5

A Simple Bargaining Mechanism that Elicits Truthful Reservation Prices

Steven J. Brams · Todd R. Kaplan · D. Marc Kilgour

Abstract We describe a simple 2-stage mechanism whereby for two bargainers, a Buyer and a Seller, it is a weakly dominant strategy to report their true reservation prices in the 1st stage. If the Buyer reports a higher reservation price than the Seller, then the referee announces that there is a possibility for trade, and the bargainers proceed to make offers in a 2nd stage. The average of the 2nd-stage offers becomes the settlement if they both fall into the interval between the reported reservation prices; if only one offer falls into this interval, it is the settlement, but it is implemented with probability $\frac{1}{2}$; if neither offer falls into the interval, there is no settlement. Comparisons are made with other bargaining mechanisms.

Keywords Bargaining · Truth-telling mechanisms · Probabilistic implementation · Incomplete information

We gratefully acknowledge the helpful comments of two referees as well as those from seminar participants in the Tel Aviv International Workshop on Game Theory, Bargaining Theory and Applications Summer School at the University of Turku, Theoretical Aspects of Rationality and Knowledge—TARK meetings at the University of Groningen, and the Royal Economic Society Conference at Cambridge.

S. J. Brams
Department of Politics, New York University, New York, NY, USA
e-mail: steven.brams@nyu.edu

T. R. Kaplan
Department of Economics, University of Haifa, Haifa, Israel

T. R. Kaplan
Department of Economics, University of Exeter Business School, Exeter, UK
e-mail: Dr@ToddKaplan.com

D. M. Kilgour (✉)
Department of Mathematics, Wilfrid Laurier University, Waterloo, ON, Canada
e-mail: mkilgour@wlu.ca

Published online: 01 June 2014

JEL Classifications C72 · C78

1 Introduction

How to get players to go to their "bottom lines" in bargaining is an age-old problem. Multiple parties on the same side can be induced using a sealed-bid second-price auction, or Vickrey (1961) auction, whereby (with multiple Buyers) the high bidder pays the second-highest bid, rendering what the winner pays independent of what he or she bid. The extension of this idea to Vickrey (1961), Clarke (1971) and Groves (1973) (VCG) mechanisms in more general two-sided bargaining likewise induces honesty, because the settlement does not depend directly on what a player offers. Similarly, Brams and Kilgour (2001) show that when players bid for rooms in a house in which they share the rent, the "gap procedure" creates a kind of partial independence, motivating the housemates to make truthful bids that sum to the total rent of the house.

While a VCG mechanism induces truthful reporting, it cannot be both ex-post budget balanced and ex-post individually rational.[1] Many intuitive mechanisms that are budget balanced do not induce honest reporting. When two bargainers haggle over the price of some good or service, averaging their offers ("splitting the difference") does not induce honesty, because it gives them incentives to exaggerate in opposite directions. When reservation prices are uniformly distributed, Chatterjee and Samuelson (1983) (CS) show that this mechanism has a simple symmetric equilibrium in which exaggeration is piecewise linear.[2] While Myerson and Satterthwaite (1983) show the CS mechanism maximizes efficiency for the symmetric uniform case, there is also an infinity of asymmetric equilibria (Leininger et al. 1989). Hence, the CS mechanism cannot form the basis of a scheme that always induces truth-telling.[3]

In this paper, we give a simple two-stage mechanism whereby two bargainers have as a weakly dominant strategy to truthfully reveal their reservation prices to a referee in stage 1. If the 1st-stage offers overlap, there is the *potential* for an agreement, which is realized—at the mean of the 2nd-stage offers—if *both* bargainers' offers fall in the overlap interval; if *only one* bargainer's offer falls in this interval, this offer becomes the agreement with probability $\frac{1}{2}$, but otherwise not; if *neither* bargainer's offer falls in the overlap interval, there is no agreement.

[1] The dominant-strategy mechanisms of Manelli and Vincent (2010), Gershkov and Goeree (2013) also lack ex-post balanced budgets.

[2] More precisely, in their game, the final price is an average of the offers if they overlap (i.e., if the Buyer's offer equals or exceeds the Seller's, thereby creating an *overlap interval*); otherwise, there is no agreement, and the players get nothing.

[3] The Revelation Principle (see Myerson 1979) implies that any equilibrium (such as CS) can be converted into a truth-telling mechanism by creating for each player a "robot" programmed to play the equilibrium strategy corresponding to the value it receives from the player. For instance, in a first-price, two-player auction with a uniform distribution of values, it is an equilibrium to bid half one's value. One can induce truth-telling by changing the rules so that the highest-bidder-wins but pays half his bid. To modify the CS mechanism with uniform values in a similar fashion, one can use the average of three numbers: the two bids and the expected value. However, if the Revelation Principle is used to create schemes, truth-telling may not be a weakly dominant strategy or even the unique equilibrium, as we see with the CS mechanism. For background information on mechanism design, see Nisan (2007).

Besides historical reasons, why do we care about making truth-telling a weakly dominant strategy? The designer may wish to gain information about the distributions of reservation prices, which could be useful to similar agents in reaching an agreement or could be used to create a better mechanism.[4] In the absence of truth-telling, gaining this information would be impossible if there were multiple equilibria, and it could be difficult even if there were not. A designer would need to learn the reservation prices from the bids by inverting equilibrium strategies. If the designer knew less than the agents about the distributions of values, then he would not know the equilibrium strategy to invert. In practice, the problem of acquiring distributional information is aggravated when agents have heterogeneous knowledge or make mistakes. This is ameliorated by employing a mechanism that has a weakly dominant strategy to truth-tell since then such strategies would likely be deployed.[5]

As we will show, our procedure, like another probabilistic mechanism (Brams and Kilgour 1996) that we discuss in Sect. 5, is not maximally efficient (Myerson and Satterthwaite 1983). In particular, if the players have uniform distributions over each others' reservation prices, it is not as efficient as the the mechanism of Chatterjee and Samuelson (1983), in part because of the random draw when exactly one offer falls in the overlap interval.

But even if no agreement is *realized*, our procedure does reveal—if it continues to stage 2—that the reservation prices of the bargainers *allow* for a mutually profitable agreement. This information may be useful for other parties that bargain under similar circumstances. Again, a VCG can obtain this information by inducing truthful reporting (which is weakly dominant), but it cannot be both ex-post budget balanced or ex-post individually rational.[6]

Another benefit of our mechanism is that there is a positive probability of settlement even for extreme reservation values; in contrast, bargainers with extreme values have no incentive to bargain under the CS mechanism, because there is no possibility of any agreement. Finally, the simplicity of both the rules and the equilibrium should appeal to both experimentalists and practitioners.

2 The Mechanism

We consider the possible sale of an object by a Seller to a Buyer. If they can agree on a price p, the object will be transferred from Seller to Buyer, and the Seller will receive p as compensation. Of course, Seller prefers a higher p, whereas Buyer prefers a lower p. If they cannot agree on a p, there is no sale.

[4] Distributional information is quite important for Google and Yahoo in helping set their reserve prices for advertisements. Also, releasing information about toxic loan bundles may aid in their sale.

[5] See Bergemann and Morris (2005), Hagerty and Rogerson (1987) for additional arguments in favor of dominant-strategy mechanisms, which are called robust or strategy-proof mechanisms.

[6] Another mechanism that induces truth-telling in weakly dominant strategies asks Buyer and Seller to submit their reservation prices and independently draws a possible sale price at random from a given distribution. If the price is satisfactory to both Buyer and Seller, then a sale occurs at that price. This mechanism is similar to the Becker et al. (1964) procedure, but the latter is less efficient since any price is determined independently of values.

For definiteness, our discussion is phrased in terms of a possible sale, but our mechanism has many applications, such as to the settlement of a claim by an insured party against an insurer. In that case, the insured party, preferring a higher settlement, plays the role of Seller, with the insurer in the role of Buyer.

We model the Seller's reservation price for the object, S, as the value of a random variable with cumulative distribution function F_S. The Buyer's reservation price, B, is the value of a random variable with cumulative distribution function F_B. Both F_S and F_B have support $[C, D]$. Both players' reservation prices are private information, and their utilities are quasi-linear, so that if a sale takes place at price p, Buyer will receive $B - p$ and Seller will receive $p - S$. If there is no sale, both players receive 0. The players are risk-neutral.

The mechanism is a two-stage procedure:

Stage 1 The players submit *reserves* to the referee: Seller submits $\widehat{S}$ and Buyer submits $\widehat{B}$. The reserves may or may not equal the corresponding reservation prices (i.e., the 1st-stage submissions are not necessarily truthful). If $\widehat{S} \leq \widehat{B}$, the *overlap interval* is $[\widehat{S}, \widehat{B}]$, and the procedure moves to stage 2. If $\widehat{S} > \widehat{B}$, the reserves do not overlap, there is no settlement, and the procedure ends.

Stage 2 The players are told that they reached stage 2.[7] The players submit *offers* to the referee: Seller submits $s \geq \widehat{S}$, and Buyer submits $b \leq \widehat{B}$. If both s and b fall in the overlap interval defined in stage 1, there is a sale at price $p = \frac{s+b}{2}$. If only one of s and b falls in the overlap interval, the name of one player is selected at random; if the selected player's offer is the one in the overlap interval, then it is sale price; if not, there is no sale. If neither offer is in the overlap interval, there is no sale.

This mechanism determines (i) whether there is a sale and (ii) if there is a sale, at what price. As usual, we model each player as privately learning its own (true) reservation price (S or B) prior to stage 1, and using this information to choose its strategy: $(\widehat{S}, s)$ for Seller; $(\widehat{B}, b)$ for Buyer. Thus, a strategy for Seller is a pair of functions $\widehat{S}(S)$ and $s(S)$ that give the values of its strategic variables as a function of its reservation price. Similarly, Buyer's strategy can be thought of as two functions, $\widehat{B}(B)$ and $b(B)$.

One strategy for a player *weakly dominates* another strategy for that player if the first yields an expected utility that is at least as great as the second, no matter what strategy is chosen by the opponent. A (Bayesian-Nash) *equilibrium* is a profile of strategies with the property that each player's equilibrium strategy maximizes the player's expected utility, given that the opponent plays according to its equilibrium strategy.

To represent our mechanism, we use two functions,

$$t : \mathbb{R}^4 \longrightarrow [0, 1] \quad \text{and} \quad p : \mathbb{R}^4 \longrightarrow \mathbb{R} \tag{1}$$

with the interpretation that $t(\widehat{S}, s, \widehat{B}, b)$ is the probability that an agreement is reached if the 1st-stage reserves are $\widehat{S}$ and $\widehat{B}$, and the 2nd-stage offers are s and b; similarly,

[7] This knowledge does affect strategies since there cannot be a transaction if stage 2 is not reached.

$p(\widehat{S}, s, \widehat{B}, b)$ is the price. Note that both $\widehat{S}$ and s are functions of Seller's true reservation price S; we have written $\widehat{S}$ instead of $\widehat{S}(S)$, and s instead of $s(S)$, for notational simplicity. Observe that the value of p is irrelevant if the value of t is 0; that is, if there is no transaction. Using the functions t and p, we can describe our mechanism formally, as follows:

$$(t, p) = \begin{cases} (1, \frac{s+b}{2}) & \text{if } \widehat{S} \leq s, b \leq \widehat{B}, \\ (\frac{1}{2}, b) & \text{if } \widehat{S} \leq b \leq \widehat{B} < s, \\ (\frac{1}{2}, s) & \text{if } b < \widehat{S} \leq s \leq \widehat{B}, \\ (0, 0) & \text{otherwise.} \end{cases} \tag{2}$$

We can construct a strategically equivalent mechanism by retaining stage 1 and replacing stage 2 by

Stage $2'$ One player, Seller or Buyer, is chosen at random. If Seller is chosen, and if Seller's 2nd-stage offer s satisfies $s \leq \widehat{B}$, then the transaction takes place at price $p = s$; if $s > \widehat{B}$, then there is no sale. Similarly, if the player chosen is Buyer, and if Buyer's 2nd-stage offer b satisfies $\widehat{S} \leq b$, then the transaction takes place at price $p = b$; if $b < \widehat{S}$, there is no transaction.

We assume that the random selection of a player in stage $2'$ is independent of the players' reservation prices. The equivalence of the two mechanisms arises because, if stage $2'$ is followed, the players' expected utilities are exactly as in (2). For example, if $\widehat{S} \leq s, b \leq \widehat{B}$, then Seller's expected utility is

$$\frac{1}{2}(s - S) + \frac{1}{2}(b - S) = \frac{s + b}{2} - S = p - S, \tag{3}$$

where p is determined by the first condition of (2). A similar relation holds for Buyer. The verification is immediate if the 2nd-stage offer of only one player, or none, falls in the overlap interval. Below, we will use the stage 2 and stage $2'$ formulations interchangeably.

Finally, we assume that the players' reservation prices are independently distributed. Moreover, we assume that their distributions, F_S and F_B, have identical support $[C, D]$, are continuous, and have strictly positive densities.

3 A Simple Truth-Telling Mechanism

We define truth-telling as follows:

Definition 1 Seller's strategy $(\widehat{S}, s)$ is **truth-telling** if $\widehat{S}(S) = S$ for all $S \in [C, D]$. Buyer's strategy $(\widehat{B}, b)$ is **truth-telling** if $\widehat{B}(B) = B$ for all $B \in [C, D]$. A strategy profile $(\widehat{S}, s; \widehat{B}, b)$ is a **truth-telling equilibrium** if it is an equilibrium and both players' strategies are truth-telling.

Note that truth-telling refers to the players' reserve strategies (1st stage), not their offer strategies (2nd stage).

Buyer's monotone hazard rate condition is satified iff Buyer's cumulative distribution function, $F_B(x)$, satisfies $\frac{d}{dx}\frac{F_B'(x)}{1-F_B(x)} \geq 0$ for all $x \in [C, D]$. Similarly, *Seller's monotone hazard rate* condition is satisfied iff Seller's cumulative distribution function, $F_S(x)$, satisfies $\frac{d}{dx}\frac{F_S'(x)}{F_S(x)} \leq 0$ for all $x \in [C, D]$. If both Buyer's and Seller's monotone hazard rate conditions are satisfied, our procedure has a unique truth-telling equilibrium, as shown next.

Proposition 1 *Truth-telling is a weakly dominant strategy for both Buyer and Seller and the players' 2nd-stage offers, s^* and b^*, are solutions of $1 - F_B(s) = (s - S)F_B'(s)$ and $F_S(b) = (B - b)F_S'(b)$, respectively. Moreover, if $F_S(\cdot)$ and $F_B(\cdot)$ satisfy the monotone hazard rate conditions, then the solutions s^* and b^* are unique, and the only truth-telling equilibrium is $(S, s^*; B, b^*)$.*

Proof Seller knows the value of S and determines his strategy $(s, \widehat{S})$ using the functions $s(S)$ and $\widehat{S}(S)$. Seller's expected utility is taken with respect to Buyer's value, B, and the random choice of Seller or Buyer. Therefore, Seller's expected utility given S (calculated according to the procedure of stage $2'$) is

$$\frac{1}{2}E\left[I_{b(B)\geq\widehat{S}} \cdot (b(B) - S)\right] + \frac{1}{2}E\left[I_{\widehat{B}(B)\geq s} \cdot (s - S)\right] \tag{4}$$

where I_C is an indicator function that takes the value of 1 or 0 according to whether condition C is true or false. The first expression is associated with the random selection of Buyer (so the price is b), and the second expression with the selection of Seller (so the price is s). Note that $\widehat{S}$ only appears in the expected utility as part of the indicator function in the first expression. Also, note that the first expression is maximized when the indicator function is 1 whenever $b(B) - S > 0$ and is 0 otherwise. This is achieved by setting $\widehat{S}$ equal to S. This is truth-telling and is true for any $b(B)$. Hence, it is a weakly dominant strategy to tell the truth. Likewise, one can show the same for truth-telling of Buyer.

From (2), if there is a positive probability of trade, then $s \geq \widehat{S}$ and $b \leq \widehat{B}$. Next, consider stage $2'$, and suppose that Seller is chosen. Then the two strategy functions, $s(S)$ and $\widehat{B}(B)$, can be proven to be differentiable and strictly increasing, using method of Theorem 1 of Chatterjee and Samuelson (1983). Similarly, assuming that Buyer is chosen leads to a proof that $\widehat{S}(S)$ and $b(B)$ are differentiable and strictly increasing. Therefore we can rewrite (4) as

$$\frac{1}{2}\int_{b^{-1}(\widehat{S})}^{D} (b(B) - S)\, dF_B(B) + \frac{1}{2}\int_{\widehat{B}^{-1}(s)}^{D} (s - S)\, dF_B(B). \tag{5}$$

Consider the information provided by (5) about Seller's choice of strategy functions. The first integral of (5) depends on $\widehat{S}$ but not s, and the second integral of (5) depends on s but not $\widehat{S}$. Therefore, Seller maximizes its expected utility by choosing $\widehat{S}$ to maximize the first integral, and s to maximize the second.

We address Seller's choice of s to maximize the second integral of (5). With the substitution $\widehat{B}(B) = B$, this integral becomes

$$W(s) = \frac{1}{2}\int_{s}^{D} (s - S)\, dF_B(B) \tag{6}$$

$$= \frac{1}{2}(s - S)(1 - F_B(s)). \tag{7}$$

Because $S < D$, the maximum of $W(s)$ must be interior, as $W(S) = W(D) = 0$ but $W(\frac{S+D}{2}) > 0$. Hence, the maximum must occur at a value of s that solves the first-order condition $1 - F_B(s) = (s - S)F'_B(s)$. This equation always has at least one solution. Moreover, under the monotone hazard rate condition for Buyer, this solution is unique; we denote it $s^*(S)$.

Buyer's optimal offer, and its relation to the monotone hazard rate condition for Seller, are analogous. In particular, if both monotone hazard rate conditions are satisfied, there is a unique truth-telling equilibrium, which we denote $(S, s^*; B, b^*)$.[8]

□

Our proof that $(S, s^*; B, b^*)$ is an equilibrium thus relies on the fact that strategies that are not truth-telling are weakly dominated by strategies that are, so there must be a truth-telling equilibrium. Then the offer strategies are obtained by maximizing the players' expected utilities under the assumption of truth-telling. To understand why truth-telling dominates, note that each player benefits from maximizing the width of the overlap interval, up to its reservation price—Seller from below and Buyer from above—in order to ensure, insofar as possible, that the 2nd-stage bids, s and b, fall in the interval, thereby meeting a necessary condition for an agreement.

In fact, truthfully reporting one's reservation price in stage 1 is analogous to bidding one's reservation price in a Vickrey auction: Just as a player cannot win in a Vickrey auction without being the highest bidder, a bargainer cannot reach a settlement unless there is an overlap interval, leading to stage 2. In each case, a player goes to its bottom line for two reasons: (i) failing to do so in stage 1 could preclude a favorable outcome in stage 2 (or cause an unfavorable outcome) and (ii) once in stage 2, the outcome does not depend on what the player reported in stage 1.

A player's utility, if positive, does not depend on the reserves, $\widehat{S}$ and $\widehat{B}$, submitted in stage 1 but, instead, on its bid, s or b, submitted in stage 2. The independence between a player's reserve and its offer implies that it can "afford" to be truthful in stage 1. In fact, a player cannot do worse by reporting its reservation price truthfully in stage 1, and may do better, so we say that under our mechanism each player has an incentive to report its reservation prices truthfully.

The story is different, however, in stage 2: Each player will have an incentive to shade its offer, depending on the distribution of the opponent's reservation price. In the

8 Riley and Zeckhauser (1983), among others, solve a problem similar to finding s to maximize $W(s)$, which can be interpreted as a monopolist's optimal take-it-or-leave-it offer. The problem is also equivalent to finding a Buyer's optimal take-it-or-leave-it offer.

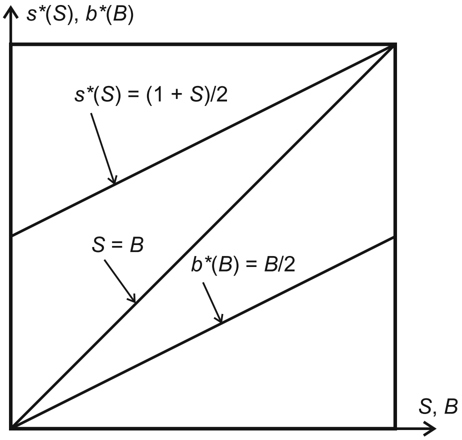

Fig. 1 Offer strategies in Example 1

next section, we illustrate, for two specific distribution functions, how much shading is optimal.

4 Examples

For our examples, we assume $C = 0$ and $D = 1$, so that $0 \leq S, B \leq 1$.

Example 1 Uniform distribution: $F_S(x) = F_L(x) = x$.

The players' optimal offers, $s^*(S) = \frac{1+S}{2}$ and $b^*(B) = \frac{B}{2}$, are shown in Fig. 1above. Notice that each of these strategies halves the distance between the reservation prices, S and B (shown as the line $S = B$), and the endpoints, 1 and 0, respectively, of the bargaining range. In particular, Seller never offers below $\frac{1}{2}$, and Buyer never offers above $\frac{1}{2}$.

It is clear that truthfully reporting one's reservation price in the 1st stage is weakly but not strictly dominant. For example, if Seller's true value is $S = \frac{3}{4}$, then, as can be seen from Fig. 1, Seller's 2nd-stage bid will be $s = s^*(\frac{3}{4}) = \frac{7}{8}$ at equilibrium. If $B > \frac{7}{8}$, there will be a sale with probability $\frac{1}{2}$; otherwise, there is no possibility of a sale. Hence, in the 1st stage, Seller will be indifferent between reporting $\frac{3}{4}$ and, say, $\frac{5}{8}$ (provided the 2nd-stage bid remains $s = \frac{7}{8}$).

Figure 2 graphs the results of the equilibrium strategies we have identified for all possible values of S and B. A sale occurs with certainty when $B < 2S$ and $B > \frac{1+S}{2}$; these two conditions define the region with darker shading in Fig. 2. This region has small values of S and large values of B, with the difference between them so great that the offers, s^* and b^*, fall in the overlap interval.

A transaction occurs with probability $\frac{1}{2}$ when $2S < B < \frac{1+S}{2}$ and when $\frac{1+S}{2} < B < \min\{2S, 1\}$, which are the two regions with lighter shading in Fig. 2. In the first of these regions (lower left), $s^* > B$ but $b^* > S$, so there is a sale at $p = b^*$ when Buyer's name is drawn in stage $2'$, and no sale otherwise. Similarly, in the upper right

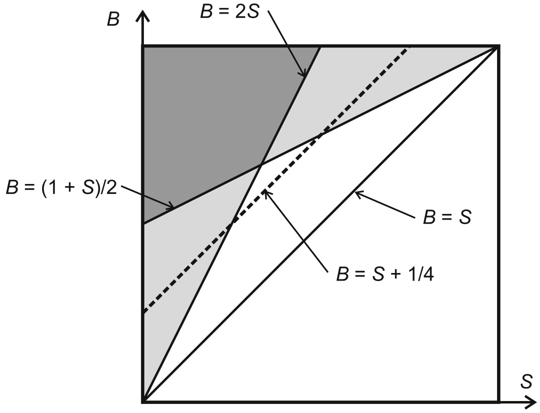

Fig. 2 Conditions for a transaction in Example 1

region, $s^* < B$ and $b^* < S$, so there is a sale at $p = s^*$ when Seller's name is drawn in stage $2'$, and no sale otherwise.

It is instructive to compare our mechanism with the Chatterjee and Samuelson (1983) mechanism, which produces a transaction, for certain, if and only if $B \geq S + \frac{1}{4}$. This inequality defines the area above the dashed line in Fig. 2. We compare the two mechanisms using the expected surplus they produce, which because of our assumptions equals the total expected utility of Buyer and Seller after the transaction, if any. For an "ideal" procedure, which produces a settlement whenever the players' reservation prices overlap, the total surplus is

$$\int_0^1 \int_S^1 (B - S)\, dS dB = \frac{1}{6}. \tag{8}$$

Myerson and Satterthwaite (1983) demonstrated that no mechanism can produce a larger surplus than the CS mechanism, which gives

$$\int_0^{\frac{3}{4}} \int_{S+\frac{1}{4}}^1 (B - S)\, dB dS = \frac{9}{64}. \tag{9}$$

The surplus from our mechanism is

$$\frac{1}{2} \int_0^1 \int_{\frac{1+S}{2}}^1 (B - S)\, dB dS + \frac{1}{2} \int_0^{\frac{1}{2}} \int_{2S}^1 (B - S)\, dB dS = \frac{1}{8}, \tag{10}$$

which is $\frac{8}{9} = 88.9\ \%$ of the maximum possible surplus.

But there are other ways to compare mechanisms. One positive aspect of ours is the potential for trade at all possible values of S and all possible values of B, a feature that the CS mechanism does not share. For instance, if $S = 0.8$ and $B \geq 0.9$, a sale occurs with probability 0.5 under our mechanism, but probability 0 under the CS mechanism.

Example 2 Power distribution: $F_S(x) = x^\alpha$, $F_B(x) = 1 - (1 - x)^\beta$, for $\alpha, \beta > 0$.

It is easy to verify that these distributions satisfy the monotone hazard rate conditions. Buyer's optimal offer is $b^* = \frac{\alpha B}{1+\alpha}$ and Seller's is $s^* = \frac{1+\beta S}{1+\beta}$, in agreement with Example 1, which corresponds to $\alpha = \beta = 1$. For example, when $\alpha = \beta = 2$, $b^*(B) = \frac{2}{3}B$ and $s^*(S) = \frac{1}{3} + \frac{2}{3}S$, and when $\alpha = \beta = \frac{1}{2}$, $b^*(B) = \frac{1}{3}B$ and $s^*(S) = \frac{2}{3} + \frac{1}{3}S$.

5 Discussion and Conclusions

We have demonstrated a simple and elegant 2-stage mechanism that induces two bargainers to be truthful in reporting their reservation prices in the 1st stage; if these prices criss-cross, the referee reports that there is an overlap interval, and the bargainers make offers in a 2nd stage. The mean of these offers becomes the settlement if they both fall in the overlap interval. If only one offer does, it becomes the settlement price, but it is implemented with probability $\frac{1}{2}$; if it is not implemented, or if neither offer falls in the overlap interval, there is no settlement.[9]

There are several mechanisms equivalent to, or related to, ours. For example, stage 1 and stage 2 could be simultaneous. Or they could be reversed. In the latter case, the 1st stage would be identical to the CS mechanism; if there is no trade, the 2nd stage would ask for reserves. Then a player would be chosen with probability $\frac{1}{2}$; that player's offer would be the transaction price if it is within the opponent's reserve.[10] There is also the possibility of combining our mechanism with that of CS. The designer would collect the information $(\widehat{B}, b; \widehat{S}, s)$ and run the CS mechanism with probability α and our mechanism with probability $1-\alpha$. For all $\alpha \in (0, 1)$, it would be weakly dominant to tell the truth, i.e., $\widehat{B} = B$, $\widehat{S} = S$. Moreover, this scheme would approach CS in efficiency as $\alpha \to 1$.[11]

Our mechanism is less efficient than the CS mechanism for uniformly distributed reservation prices, but it does have several features to recommend it, such as the possibility of a transaction even for an extreme reservation price. If the monotone hazard rate condition holds, there is a unique truth-telling equilibrium (and, even without this condition, it is almost always unique). Under our mechanism, truth-telling in the 1st

[9] There is always a settlement under final-offer arbitration, in which two parties make so-called final offers, and an arbitrator, who cannot split the difference, chooses one of the other. (Among other applicatons, it is used in salary disputes in major league baseball.) Although this mechanism does not induce convergence to the median of the arbitrator's perceived probability distribution over fair settlements (Brams and Merrill 1983), there are several mechanisms that do induce convergence to the median (Brams and Merrill 1986; Zeng 2003; Zeng et al. 1996). Analogous to being truthful about one's reservation price under our bargaining mechanism, proposing the perceived median under these alternative arbitration procedures is a dominant strategy, enabling the two parties to reach a settlement.

[10] If $\widehat{B} \geq s > b \geq \widehat{S}$, the selling price would be $(s + b)/2$ in the original mechanism. In this case, it would be s or b, each with probability $1/2$.

[11] The gains from truth-telling would be much smaller, but still positive.

stage is always optimal, but other behavior may produce the same results. However, for a player whose value falls within the range of 2nd-stage offers of the opponent, the equilibrium strategy of truth-telling is strictly optimal. Of course, our mechanism admits a no-trade equilibrium (a trade occurs with probability zero) in addition to the truth-telling equilibrium, but this feature plagues any mechanism that satisfies individual rationality and balances the budget.[12] We feel that such an equilibrium would be unlikely to occur in practice when the alternative, truth-telling, has positive expected utility. In fact, in our mechanism the designer would know that, if a trade occurred, there must have been truth-telling in stage 1. Moreover, any reserve that lies in the interior of the support of the opponent's reservation prices must be truthful.

Part of the inefficiency of our mechanism stems from the random implementation of a 2nd-stage offer, s or b, as the exchange price when only one offer falls in the overlap interval. Randomizing the implementation of a single inside offer is the penalty one pays to render a player's reserve independent of its offer in the expected-payoff calculation, thereby making it optimal for the player to report truthfully its reservation price. This independence would be broken, and it would be suboptimal for a player to truthfully report her reservation price, if single inside offers were implemented with certainty.

Hagerty and Rogerson (1987) show that the only dominant-strategy mechanism in this environment is a posted price, that is, where a price is chosen in advance and a sale occurs at that price only if the Seller's value is below the price and the Buyer's value is above the price. Our mechanism is more efficient since we are able to get around this restriction by only requiring that the first stage is in dominant strategies.[13]

Brams and Kilgour (1996) analyze other mechanisms that induce two bargainers to be truthful, including a "bonus procedure" in which a third party induces the bargainers to be truthful by paying them a bonus when their bids criss-cross. But it is their "penalty procedure" that is closest to the present mechanism in inducing truth-telling behavior.[14]

Under it, the bargainers make simultaneous offers in a single stage, with the proviso that the probability of implementation of a settlement is a function of the *degree* of overlap, if any, of the bids: the greater the overlap, the higher this probability.[15] This procedure yields a surplus of $\frac{1}{12}$, which is 50 % of the maximum possible, falling far short of the 88.9 % achieved by the present mechanism. Moreover, unlike

[12] More precisely, when Buyer's and Seller's values are drawn from distributions with the same non-trivial support, any mechanism that satisfies individual rationality and balances budgets must admit a no-trade equilibrium. To see this, let C be the lower bound of the common support and let $D > C$ be the upper bound. Suppose that $B = C$ or $S = D$ at any equilibrium of the mechanism. Then no trade is possible, as there can be no surplus for the agents. It follows that the strategy profile $(\widehat{B}(C), b(C); \widehat{S}(D), s(D))$, where Buyer behaves as if his value is C and Seller as if her value is D, is always an equilibrium, because the surplus is zero, and must remain so even if one player changes strategy unilaterally.

[13] Note that our requirement is slightly stronger in that we require truth-telling in the first-stage to be weakly dominant.

[14] More details on the comparison of the different honesty-inducing bargaining procedures are given in Kilgour et al. (2011).

[15] The probability of a *certain* settlement in the present mechanism also increases as the overlap of the 1st-stage reserves increases, because a greater overlap increases the likelihood that *both* players' 2nd-stage offers will fall into the overlap interval, ensuring a settlement.

the present mechanism, the players never learn whether their failure to settle was because (i) their bids did not criss-cross (as in stage 1), or (ii) they did criss-cross but probabilistic implementation prevented a settlement. In principle, however, they could be told whether (i) or (ii) prevented a settlement; if (ii), they might be motivated to try again, but not using the same mechanism (see discussion below).

An advantage of the present mechanism is that the players *always* learn if stage 2 is reached and, therefore, that there is an overlap interval and the *potential* for a mutually profitable settlement. While our mechanism does not reveal the amount of regret—for example, how close the 2nd-stage offers are to the overlap interval—we see no reason why the values of $\widehat{S} = S$, $\widehat{B} = B$, s, and b could not be revealed by the referee, making public the reason why implementation failed in the 2nd stage.

The knowledge that the optimality of shading one's "bottom line" in stage 2 was all that prevented a settlement might motivate other bargainers in a similar situation to try to find an agreement *by other means* (e.g., face-to-face informal bargaining, mediation, etc.). These are discussed in Brams and Mitts (2013), who suggest several applications, especially to contract, regulatory, and corporate law. We stress, however, that under our model, the bargainers must assign probability 0 to the possibility that they could benefit from the procedure when it produces no agreement.[16]

References

Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. Behav Sci 9(3):226–232

Bergemann Dirk, Morris Stephen (2005) Robust mechanism design. Econometrica 73(6):1771–1813

Brams SJ, Kilgour DM (1996) Bargaining procedures that induce honesty. Group Decis Negot 5(3):239–262

Brams SJ, Kilgour DM (2001) Competitive fair division. J Polit Econ 109(2):418–443

Brams SJ, Merrill S III (1983) Equilibrium strategies in final-offer arbitration: there is no median convergence. Manag Sci 29(8):927–941

Brams SJ, Merrill S III (1986) Binding versus final-offer arbitration: a combination is best. Manag Sci 32(10):1346–1355

Brams SJ, Mitts JR (2013) Law and mechanism design: procedures to induce honest bargaining. N Y Univ Annu Surv Am Law 68(4):729–790

Clarke E (1971) Multipart pricing of public goods. Public Choice 8:19–33

Chatterjee K, Samuelson W (1983) Bargaining under incomplete information. Oper Res 31(5):835–851

Gershkov A, Jacob K, Goeree JK, Kushnir A, Moldovanu B (2013) On the equivalence of Bayesian and dominant strategy implementation. Econometrica 81(1):197–220

Groves T (1973) Incentives in teams. Econometrica 41(4):617–631

Hagerty KM, Rogerson WP (1987) Robust trading mechanisms. J Econ Theory 42(1):94–107

Kilgour DM, Brams SJ, Kaplan TR (2011) Three procedures for inducing honesty in bargaining. In: ACM proceedings of the 13th conference on theoretical aspects of rationality and knowledge, pp 170–176

16 Thus, even if the procedure finds an overlap in stage 1, its failure to produce a settlement in stage 2 implies that its further use is foreclosed. (If this were not the case, truthfulness would not be optimal in stage 1.) But if the procedure is unsuccessful, there is nothing in the model that prevents the bargainers from continuing to negotiate with each other—in effect, to transcend the limitations of the procedure. It would be interesting to run an experiment that gives subjects the choice of walking away or continuing to negotiate if there is no settlement in stage 2. How many would elect to continue, and how successful would they be? Besides answering this question, the simple linear offers recommended by our mechanism could be tested in an experiment that assumes two-sided incomplete information; see, for example Radner and Schotter (1989), Rapoport et al. (2008) for tests of other bargaining mechanisms that make this assumption.

Leininger W, Linhart PB, Radner R (1989) The sealed-bid mechanism for bargaining with incomplete information. J Econ Theory 48(1):63–106

Manelli AM, Vincent DR (2010) Bayesian and dominant-strategy implementation in the independent private values model. Econometrica 78(6):1905–1938

Myerson RB (1979) Incentive compatibility and the bargaining problem. Econometrica 47:61–73

Myerson RB, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. J Econ Theory 29(2): 265–281

Nisan N (2007) Introduction to mechanism design (for computer scientists). In: Nisan N, Roughgarden T, Tardos É, Vazirani VV (eds) Algorithmic game theory. Cambridge University Press, Cambridge, pp 209–241

Radner R, Schotter A (1989) The sealed-bid mechanism: an experimental study. J Econ Theory 48(1): 179–220

Rapoport A, Daniel TE, Seale DA (2008) Asymmetric two-person bargaining under incomplete information: strategic play and adaptive learning. In: Plott CR, Smith VL (eds) Handbook of experimental economics results, vol 1. Elsevier North Holland, pp 560–571

Riley J, Zeckhauser R (1983) Optimal selling strategies: when to haggle, when to hold firm. Q J Econ 98(2):267–289

Vickrey W (1961) Counterspeculation, auctions, and competitive sealed bid tenders. J Financ 26(1):8–37

Zeng D-Z (2003) An amendment to final-offer arbitration. Math Soc Sci 46(1):9–19

Zeng D-Z, Nakamura S, Ibaraki T (1996) Double-offer arbitration. Math Soc Sci 31(3):147–170