



Sequential bargaining with common values[☆]

Paul Schweinzer*

Department of Economics, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom

ARTICLE INFO

Article history:

Received 18 April 2008
 Received in revised form 17 August 2009
 Accepted 1 September 2009
 Available online 8 September 2009

JEL classification:

C73
 C78
 D44
 D82
 J12

Keywords:

Sequential bargaining
 Common values
 Incomplete information
 Repeated games
 Cake cutting

ABSTRACT

We study the alternating-offers bargaining problem of assigning an indivisible and commonly valued object to one of two players who jointly own this object. The players are asymmetrically informed about the object's value and have veto power over any settlement. There is no depreciation during the bargaining process which involves signalling of private information. We characterise the perfect Bayesian equilibrium outcome of this game which is unique if offers are required to be strictly increasing. Equilibrium agreement is reached gradually and non-deterministically. The better informed player obtains a rent.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

What is the value of private information in a common value bargaining environment? We study this question in an alternating-offers bargaining problem where a privately informed player signals his information to an uninformed opponent through his bidding behaviour. In our setup, there is an indivisible object that is of either high or low value. Both players know these possible values. Only one player (P1) knows the true value of the object, while all the other player (P2) knows is a probability distribution over the possible values. The players are infinitely patient and—as both sides make offers—possess similar bargaining power. In the *unique* equilibrium outcome, the informed player obtains an information rent even when the object is worthless.

Our study of bargaining with non-depreciating common values complements the analysis of depreciating private values initiated by Rubinstein (1982). The assumption of depreciation typically leads to immediate or temporally finely tuned agreement in subgame perfect equilibrium. This phenomenon, however, is not always observed in bargaining situations

[☆] Thanks to Abraham Neyman, Elchanan Ben-Porath, Alex Gershkov, Sergiu Hart, Hamid Sabourian, two anonymous referees, the associate editor, and, in particular, to Robert J. Aumann and Gil Kalai. I am grateful for the hospitality of the Center for the Study of Rationality, Jerusalem, and financial support from the German Science Foundation through SFB/TR 15.

* Correspondence address: University of Manchester, Department of Economics, Oxford Road, Manchester M13 9PL, UK. Tel.: +44 0161 275 7424; fax: +44 0161 275 4812.

E-mail address: Paul.Schweinzer@manchester.ac.uk.

involving signalling of private information.¹ Incomplete information bargaining models are typically plagued with many equilibria. One might expect that the signalling aspect introduced by the common value nature of the object further accentuates this problem. This, surprisingly, turns out to be incorrect. Indeed, the game where players are restricted to increasing offers has a unique perfect Bayesian equilibrium outcome.² It involves minimally increasing offers and, for the low type P1 and the uninformed P2, stochastically accepting or increasing offers. Not only is the equilibrium outcome unique, moreover, all equilibria of this repeated signalling game are identical (modulo some technical condition applicable only to the final stage). With minor modifications, this equilibrium—and, particularly, its outcome—prevails in the extended game allowing for non-increasing bids. The possibility of indefinite waiting, however, destroys all hope for equilibrium uniqueness.

The remarkable outcome uniqueness—and thus easy testability of our predictions—results from the conflict nature of the game and the present incomplete information. The signals embedded into the informed player's offers are interpreted by the uninformed player P2 as potentially emanating from both types of P1. The informed P1's strategy must be (semi-)pooling because no full separation would be credible or indeed possible in equilibrium. This makes only a very small set of P2's beliefs compatible with equilibrium behaviour: After observing any on- or off-equilibrium-path action by the informed P1, the uninformed P2's *equilibrium* response must entail indifference between accepting the current offer and continuing to make a higher own offer.³ This indifference defines P2's beliefs and implies that the eventual equilibrium agreement is reached gradually and stochastically over a whole sequence of offers and counteroffers. The same logic applies to the extended game allowing for non-increasing offers but there, inserting any finite or infinite sequence of non-increasing offers at any stage, may still constitute an equilibrium.

Very little is known about the performance of any non-trivial bidding game with pure common values under incomplete information. The present paper attempts to make a step towards closing this gap for alternating-offer bargaining. Of course, it is unlikely that the value of whatever asset jointly owned is exactly common to the players. However, treating that value as approximately common and focusing on the settlement dynamics that players can agree on under these circumstances establishes a potentially useful benchmark case. The involved communication game illustrates through which offers the informed party can secure a profit while retaining some of his informational advantage until the end of the game. Agreement in this game is reached both gradually and non-deterministically seems to coincide with reality.

Our framework can be used to model other repeated interactions where incomplete information plays a role, for instance, repeated moral hazard or incomplete contracts negotiations. Finally, the model describes an explicit bidding game which may be useful, for example, for the dissolution of a commonly valued partnership.

1.1. Related literature

Most existing contributions to the rich literature on bargaining with incomplete information are concerned with a player's incomplete information over the opponents' discount rate and not over the pie itself for which valuations are usually taken to be independent. Hence incomplete information bargaining is typically very sensitive to discounting considerations. If a (sequential) bargaining model considers private information on the pie, then the signalling problem is usually avoided by allowing offers only by the uninformed players.⁴ Ausubel and Deneckere (1989) are an exception in characterising the equilibria of alternating-offer bargaining games with one-sided incomplete information. They analyse a private value setup and impose ex post individual rationality on the set of outcomes of their mechanism design study. The present paper, by contrast, studies an explicit alternating offers bidding game featuring incomplete information on a common value and imposing only interim individual rationality in order to understand the underlying communication game. An insight from this bidding game is that incomplete information can be used as a refinement device. While many equilibria are present in the complete information version of the game, the obtained outcome is unique under incomplete information—something which usually has to be bought by the application of some equilibrium refinement concept.⁵

Two thematically related papers discussing further applications are Morgan (2004) and Brooks et al. (2009). The first concerns itself mainly with the analysis of fair mechanisms and the second concentrates on the performance of the 'Texas Shootout' procedure in which one partner names a single buy-sell price and the other partner has the option to buy or sell at that price. The applicable bargaining literature is surveyed by Ausubel et al. (2002) and, to the best of the author's knowledge, contains no full analysis of bargaining under incomplete information with pure common values and alternating offers by both players.

¹ Consider a 'partnership dissolution' problem where pre-divorce Daimler and Chrysler jointly hold patents on, say, the Hydrogen engine. When the firms go separate ways, suppose that these rights need to be concentrated for the technology to be marketable. The firms agree on the potential value of the machine but one firm has private information about marketing conditions. The procedure proposed here provides a (potentially automated) bargaining procedure for resolving this issue. The final section lists further applications.

² This requirement for strictly increasing offers corresponds to the 'bargaining in good faith' stipulation present in many countries' labour codes. It is also prescribed by the Jewish Torah (Mishnah, Bava Metzia 4:10). Combining a weaker non-decreasing offers restriction with the usually made assumption of depreciation of the pie amounts to our requirement. A recent example with private values is Schwartz and Wen (2007).

³ There are only exceptions to this rule when such beliefs cannot exist.

⁴ Examples are Sobel and Takahashi (1983); Cramton (1984); Fudenberg et al. (1987); Evans (1989) and Deneckere and Liang (2006).

⁵ If our setting is taken as a general model of common value trade under asymmetric information, one might wonder why it does not run afoul of the no-trade theorems, e.g. Tirole (1982). The reason is our bargaining assumption that payoffs only realise if agreement is reached through the acceptance of one player. Hence the initial allocation is not Pareto efficient and the no-trade theorems do not apply.

We share our interest in strategic information transmission in repeated bidding games with several recent papers. Deneckere and Liang (2006) analyse an infinite horizon bargaining game with an information structure similar to ours but offers by the uninformed player only. They allow for interdependent valuations and analyse the properties of their equilibrium outcome in comparison to the literature on the Coase conjecture. Calcagno and Lovo (2006) are concerned with information transmission between informed and uninformed market makers through a sequence of bid-ask quotes and the resulting inventory optimisation problem. In our setup, this corresponds to sharing the undiscounted common value of a set of stocks between two asymmetrically informed players. In strategic terms their result is similar to ours: the informed player profits from equilibria in semi-separating strategies. Hörner and Jamison (2008) analyse a sequence of first-price auctions where a common value object is sold at every stage and players maximise discounted average payoffs. Under incomplete information on one side, they find that all information is gradually revealed in finite time and, surprisingly, that the uninformed player is doing better than the informed bidder. This result cannot be readily compared to ours because their bidders' payoffs do not depend on transfers but only on the value of the objects.

Turning to complete information bargaining, Compte and Jehiel (2004) study gradual bargaining and contribution games where players can opt out at each stage of bargaining. Their players' outside options depend on the players' previous bargaining concessions. Stacchetti and Smith (2003) introduce a remarkable model of 'aspirational bargaining' in which behavioural, time-independent strategies in a procedure-free environment deliver mixing dynamics which are similar to our incomplete information while theirs are governed by behavioural Markov assumptions. The first set of results of the present paper requires that bids are strictly increasing over time.⁶ This requirement for a party to (strictly) improve over previously made offers was introduced by Ståhl (1972) and relates to the idea of 'bargaining in good faith' as recently discussed by Schwartz and Wen (2007). The latter's contribution applies to complete information wage-negotiations in a private values setup and gives a short overview of the good faith bargaining literature. As described by Deck et al. (2007) and Winoto et al. (2008), other areas assuming strictly improving offers include the analysis of arbitration mechanisms and the design of communication protocols in computer science.

There is a war of attrition flavour to the perfectly competitive nature of our game and the stage by stage mixing of the players involved leading to a unique equilibrium. In some formulations of the war of attrition, a unique Bayesian equilibrium is obtained as well. As detailed by Ponsatí and Sákovics (1995), uniqueness in these formulations stems from the restriction of the set of possible outcomes to either losing or winning—while the only restriction which we impose on outcomes comes through the bidding grid. Ponsatí (1997) introduces the two-outcome restriction into a bargaining problem with both-sided incomplete information over reserve values and allows for a third possible outcome which she interprets as a mediated compromise. Compared to our analysis, this outcome restriction simplifies the signalling problem and her depreciation assumption concentrates attention on the role of delay in finding agreement.

Our game can alternatively be interpreted as an ascending auction where the highest bidder wins the object and pays his bid to the loser. From this point of view, our analysis addresses questions similar to those explored by Milgrom and Weber (1982) and Engelbrecht-Wiggans et al. (1983) who characterise the incentives for additional information acquisition (or communication) in standard two-bidder, sealed-bid auctions.⁷ In contrast to their setting, players have mutual veto power here as results, for instance, from joint ownership. Introducing such considerations into the auction context leads to effects reminiscent of the study of auctions with toeholds by Bulow et al. (1999). Both there and in our model, a bidder with a toehold bids more aggressively than without because his bid is at the same time an offer for the remaining part of the object and a demand for his own holdings.⁸

2. The model

We consider two identical, risk-neutral players P1, P2 and two possible common values for an indivisible object $\theta \in \{\underline{\theta}, \bar{\theta}\}$. We normalise $\underline{\theta} = 0$ and assume $\bar{\theta} \geq 3$ to avoid trivialities. P1 is assumed to know the realisation of θ . We denote the high-type informed player by $\bar{P}1$ and the low-type by $\underline{P}1$. P2's only information is her public prior $\varphi^0 = \text{pr}(\theta = \bar{\theta})$ which we set, for simplicity, to 1/2 although nothing is special about this number. The prior is refined into P2's beliefs φ^i on the basis of P1's observed bids. We denote players by $i \in \{1, 2\}$, $j = 3 - i$.

⁶ This restriction of strategies, in effect, makes the game finite and allows for a backward solution. As shown subsequently, however, a version of this equilibrium prevails when this monotonicity restriction is dropped.

⁷ Similarly, our mechanism relates to ascending-clock auctions. The main problem in formulating these in our framework is to capture the non-bargaining event of both players accepting simultaneously. Sákovics (1993) illustrates that even in the one-shot, complete-information case, simultaneous offers permit any outcome. If this problem is resolved by the auctioneer alternately giving preference to one of the players, we are back in the current setup (without jump-bids).

⁸ The dynamic bargaining game studied in this paper can be formally interpreted as a repeated game of incomplete information as studied by Aumann and Maschler (1966) or Mertens et al. (1994). Indeed, the communication side of our model poses questions similar to those addressed there. That literature, however, typically derives average payoffs from long interactions which do not arise naturally in our context. Nevertheless we use repeated game terminology throughout. Our game also relates to the 'cake-cutting' problem of fairly dividing some pie. However, this literature does not usually consider asymmetric information or indivisible pies. An overview presenting a wealth of examples is Brams and Taylor (1996).

The game starts with P1 offering a payment o_1^1 (subscripts are players, superscripts time periods) to P2 for sole ownership of the object.⁹ Pure offers $o_i^t, t > 0$ are restricted to the set of admissible bids $B = \{0, 1, \dots, \bar{B}\} \subset \mathbb{N}$ where $\bar{B} > \bar{\theta}$ ('all the money in the world'). This defines a minimal offer increment of 1 (currency unit). The terms offers and bids are used interchangeably. If Pj accepts Pi's offer, Pi pays the offered amount to Pj, Pi gets the object and the game is over. If Pj does not accept Pi's offer, nothing is paid, and Pj makes an own offer. Players go on making alternating offers until one player accepts an offer. This must happen at the latest after the highest possible bid \bar{B} is reached. Adding the option for players to terminate negotiations with zero payoffs at each stage changes neither the analysis nor the results as this choice is always dominated by accepting the previous offer.

We set $o_2^0 = o_1^{-1} = 0$ equal to the low value of the object and, for the finite game \mathcal{Q} , we require offers to be strictly increasing over time $o_i^t - o_i^{t-2} > 0$. This requirement is relaxed in the extended game \mathcal{E} once the equilibrium is characterised. Accepting the opponent's last offer is denoted 'q'. Mixed offers attach probability α_i^t to the pure continuation bid o_i^t and the complementary probability to accepting. We denote such mixtures by $\beta_i^t = [\alpha_i^t : o_i^t, q]$.¹⁰

Pi's strategy β_i consists of a sequence of potentially mixed stage actions for every possible opponent action at each of the opponent's information sets and for each possible type of player i . Players observe the opponents' realised offers and enjoy perfect recall. The players' final payoffs are written $u_i(\beta|\theta)$ and given by the object's value minus the payment made for the winner and the agreed payment to the loser. Occasionally, we use the notation $\lfloor x \rfloor$ to denote the largest integer less than or equal to some $x \in \mathbb{R}$.

To sum up, our game is an alternating-offer bargaining game with incomplete information over common values and no discounting. The intuition and dynamics of our equilibrium are developed through example in the following section. The results are presented in Section 4 where we first analyse the game with strictly increasing bids and then extend our results by allowing for non-increasing offers. Proofs and details can be found in Appendix.

3. An example

We start with the case of complete information to provide a benchmark. Consider the example values $\underline{\theta} = 0$ and $\bar{\theta} = 3.5$ in the game with strictly increasing offers. With complete information, obviously, there cannot exist an equilibrium in which players use the same observable actions in both states of the world. Keeping the discussion of the complete information case informal, an equilibrium for the case of $\theta = 0$ is for P1 to accept straight away (and everywhere else) and for P2 to do likewise. This leads to the unique equilibrium outcome (0, 0). A 'minimally increasing' equilibrium path (thus involving delay) for the case of $\theta = 3.5$ is for P1 to offer 1, then P2 to offer 1, then P1 to offer 2 which P2 accepts leading to the outcome (1.5, 2).¹¹ The outcome is determined by θ , of course, which determines the bidding grid. For large $\bar{\theta}$, both players can achieve close to half the object's value and the second-mover advantage vanishes.

We now turn to the more interesting case of incomplete information and stay with the above values of $\underline{\theta} = 0$ and $\bar{\theta} = 3.5$. This is the simplest example providing the intuition of the equilibrium while illustrating the emergence of P1's information rent.¹² Recall that admissible bids are the non-negative integers bounded by some large number $\bar{B} > 3.5$. In this case, backward induction from \bar{B} shows that no player bids in excess of 2. Hence only equilibrium offers of {0, 1, 2} are considered below. In principle (and similar to the above complete information case), there may be fully revealing 'separating' equilibria in which, for instance, $\bar{P1}$ bids 1 and $\underline{P1}$ always accepts. If these were equilibrium strategies, $\underline{P1}$ could use the same action as $\bar{P1}$ in the low-value state and thus fool P2 into believing to be in the high-value state and exploit her deterministically by accepting. As this cannot be equilibrium behaviour, no strategy which deterministically reveals the value of the object can be an equilibrium as long as P2 can still condition her response on this information.

Applied to the present example, the equilibrium sketched in the introduction takes the following form: P1 starts by making an offer of 1—with probability one for $\bar{P1}$, and with probability α_1 for $\underline{P1}$ who accepts with the complementary probability $(1 - \alpha_1)$. If an offer is observed, P2 updates her belief $\varphi = \text{pr}(\theta = \bar{\theta})$ and responds to offers of 1 or 2 by mixing between bidding 1 herself with probability α_2 and accepting. She accepts higher offers. $\underline{P1}$ accepts and $\bar{P1}$ bids 2 revealing the true value if an offer of 1 is observed and accepts otherwise. If P2 gets an offer, she accepts it. All higher offers are accepted. We formalise the example equilibrium candidate $\beta^* = ((\beta_1^1(\theta), \beta_1^2(\theta), \beta_1^3(\theta)), (\beta_2^2, \beta_2^4))$ as $\beta_1^*(\underline{\theta}) = ([\alpha_1 = (1/2) : 1, q], q$ for any other bid),

$$\beta_1^*(\bar{\theta}) = \begin{cases} (1, 2, q) & \text{if } o_2^2 = 1 \\ (1, q) & \text{if } o_2^2 > 1 \end{cases}, \quad \beta_2^*(\theta) = \begin{cases} ([\alpha_2 = \frac{1}{2} : 1, q], q), \varphi = \frac{2}{3} & \text{if } o_1^1 = 1 \\ ([\alpha_2' = \frac{2}{3} : 1, q], q), \varphi' = \frac{3}{\bar{\theta}} & \text{if } o_1^1 = 2 \\ (q), \varphi = 1 & \text{if } o_1^1 > 2. \end{cases} \tag{1}$$

⁹ Since only P1 holds private information, having P2 start the game just inserts a trivial stage at the beginning.

¹⁰ We refrain from a general definition of a mixed stage action over a larger support of pure actions because, as it turns out, nothing more complicated is needed. We do not, however, artificially restrict stage actions.

¹¹ This is just one equilibrium path and there are other equilibria in the general complete information game—some of them may also include jump bids.

¹² This payoff structure is generalised by (A.11) for higher $\bar{\theta}$. In the present example, the extreme coarseness of the bidding grid makes this approximation very inaccurate.

In detail, this profile β^* prescribes the following sequence of play:

t=0: Nature decides on θ and sends a fully revealing signal to P1 and no signal to P2.

t=1: P1's minimal continuation offer not ending the game is $o_1^1 = 1$. Depending on the object's value, $\underline{P1}$ uses the type-dependent lottery $[\alpha_1 : 1, q]$ to mix between offering 1 and accepting in case of $\theta = 0$. $\overline{P1}$ bids 1 with probability one in case of $\theta = 3.5$. Because $\underline{P1}$ plays a mixed action, he must be indifferent between the payoffs of all pure actions in the support of this mixed action, and, in particular, his acceptance payoff of zero. This translates into

$$0 = (1 - \alpha_2)(-1) + \alpha_2(1) \Leftrightarrow \alpha_2 = \frac{1}{2}. \tag{2}$$

After observing an offer of 1, P2 uses the conditional mixture probability embedded in P1's equilibrium strategy (1) to compute her posterior. These are $\text{pr}(o_1^1 = 1 | \theta = 0) = \alpha_1$ and $\text{pr}(o_1^1 = 1 | \theta = 3.5) = 1$. Thus $\underline{P1}$ should bid $o_1^1 = 1$ with probability α_1 (and accept otherwise) and for $\overline{P1}$ to bid 1 with probability one. Observing P1's bid of 1, P2 updates her prior $\varphi^0 = 1/2$ to

$$\text{pr}(\theta = 3.5 | o_1^1 = 1) = \frac{\text{pr}(o_1^1 = 1 | \theta = 3.5)\text{pr}(\theta = 3.5)}{\text{pr}(o_1^1 = 1)} = \frac{1}{1 + \alpha_1} = \varphi. \tag{3}$$

t=2: P2's minimal continuation offer is 1. Given her posterior of φ , P2 plays the mixed action $[\alpha_2 : 1, q]$ for any $\alpha_2 \in [0, 1]$ because—through the appropriately chosen mixture α_1 —she is made indifferent between her acceptance payoff of one and her continuation payoff $\varphi(2) + (1 - \varphi)(-1)$. In particular, she is willing to play the mixed action $[\alpha_2 = (1/2) : 1, q]$ which makes $\underline{P1}$ indifferent between accepting and bidding 1 as required by (2). For the α_1 -generated φ to be optimal, it has to satisfy P2's equality of payoffs for all pure actions in the support of her mixed action in addition to φ resulting from Bayes' rule (3). This implies $(\alpha_1, \varphi) = ((1/2), (2/3))$. Notice that φ is calculated backwards from the last stage.

t=3: Observing P2's offer of $o_2^2 = 1$, $\underline{P1}$ accepts. $\overline{P1}$ bids the fully revealing $o_1^3 = 2$.¹³

t=4: Observing $\overline{P1}$'s continuation, P2 accepts.

Expected payoffs from β^* in (1) are $u(\beta^* | \theta = 3.5) = (2, 1.5)$ and $u(\beta^* | \theta = 0) = (0, 0)$ with ex ante expectation of $(1, 0.75)$. For observed off-equilibrium-path bids by P1, P2's response (i.e. her action complete with equilibrium mixture probability and belief) is part of her equilibrium strategy (1). Her beliefs are undetermined by the solution concept of weak perfect Bayesian Nash equilibrium alone. But we know that P2 must be indifferent between accepting and continuing or else she could deduce the object's value from P1's choice of action: If P1 chooses a (mixed) action which makes her prefer accepting, this is beneficial *only* to $\overline{P1}$ while continuation with probability one is beneficial *only* to $\underline{P1}$. Since P1 cannot credibly reveal the state in this fashion, P2 must be indifferent between continuing and accepting in equilibrium. This determines her unique equilibrium belief and, in turn, the conviction that an observed jump-bid was taken by $\underline{P1}$ with the probability α'_1 determined from her belief through Bayes' rule. In case of observing $o_1^1 = 2$ in the example, this gives $\varphi' = 3/3.5$, $\alpha'_2 = 2/3$ and $\alpha'_1 = 1/30$.

Thus an observed jump bid affects P2's beliefs such that her indifference is restored on the deviation path of play. Her involved mixing leaves $\overline{P1}$ with a payoff expectation of $3.5/3 < 2$ which renders his jump unprofitable.¹⁴ Part of the example's extensive form is shown in Fig. 1 where greyed triangles symbolise the range of mixed actions. The tree to the left of Nature's move shows the deviation path after P1's $o_1^1 = 2$. There are no profitable deviations.

4. Results

To the end of solving the game, the first three lemmata determine properties of any equilibrium of the game. The key insight is Lemma 2 which says that, in any equilibrium, P2 cannot attribute an observed action to a particular type of P1 with certainty. In order to avoid exploitation, she must mix between accepting an offer and making a higher counteroffer. Lemmata 4 and 5 construct a unique path through the game based on a terminal accepting condition for the players supplied by Lemma 1. Proposition 1 shows that all equilibria of the finite game \mathcal{Q} share the same outcome and are characterised by this unique path (modulo a technical condition). Proposition 2 determines when the informed player starts signalling his type in equilibrium and calculates a simple payoff approximation. Proposition 3 relaxes the assumption of strictly increasing offers and shows that a variant of the above equilibria still constitutes an equilibrium of the generalised game \mathcal{E} . In the following, we use the abbreviations wp1 and wpp for 'with probability one' and 'with positive probability,' respectively.

Lemma 1. (i) In any equilibrium of \mathcal{Q} , the player indexed i moving at t accepts wp1 if

$$o_i^{t-2} + o_j^{t-1} > \bar{\theta} - 1. \tag{4}$$

(ii) In no equilibrium of \mathcal{Q} , $\overline{P1}$ accepts wpp if $o_1^{t-2} + o_2^{t-1} < \bar{\theta} - 1$. (iii) If $o_i^{t-2} + o_j^{t-1} = \bar{\theta} - 1$, both $\overline{P1}$ and a P2 believing that $\theta = \bar{\theta}$ are indifferent between (any mixture of) $o_i^{t-2} + 1$ and accepting.

¹³ If $\bar{\theta} = 3$, $\overline{P1}$ is indifferent between his minimal-increment $o_1^3 = 2$, accepting, or any mixture between the two. This is one reason for the non-uniqueness of β^* : the last potential move by P1 is arbitrary for integer $\bar{\theta}$.

¹⁴ Because $o_1^{t-1} + o_2^t > \bar{\theta}$ after any higher jumps $o_1^{t-1} \geq 3$, P2 accepts with probability one for any belief. Thus her off-equilibrium-path beliefs after observing such a jump-bid are undefined by the game alone.




Fig. 1. A partial game tree of the example with $\theta \in \{0, 3.5\}$.

At this point, $\bar{P}1$'s stage action $\beta_1^t(\bar{\theta})$ may still be mixed over several offers $o_1^t \in [o_1^{t-1} + 1, \dots, \bar{\theta} - o_2^{t-1}]$. As lemma 4 shows that $\beta_1^t(\bar{\theta})$ is pure in equilibrium, we abuse notation for the sake of simplicity and let o_1^t denote both a single action and a set of actions up to Lemma 4.

Lemma 2 defines P2's equilibrium beliefs for any possible sequence of play and as long as her continuation action is not ruled out by (4). Whenever there exists a belief which allows P2 to be indifferent between accepting and bidding up, P2 is indeed made indifferent by P1's equilibrium action. This is the key insight required for the rest of the analysis.

Lemma 2. (i) No time- $(t - 1)$ action by P 1 revealing his type wp 1 can be part of an equilibrium as long as $o_2^{t-2} + o_1^{t-1} < \bar{\theta} - 1$. (ii) In any equilibrium of \mathcal{Q} , for P2 moving at t with $o_2^{t-2} + o_1^{t-1} < \bar{\theta} - 1$ and equilibrium continuation payoffs $u_2^t(\cdot|\bar{\theta}), u_2^t(\cdot|\underline{\theta})$, P2 is exactly indifferent between accepting and bidding up if $\exists \varphi^t \in [\varphi^0, 1)$ such that

$$o_1^{t-1} = \varphi^t u_2^t(\cdot|\bar{\theta}) + (1 - \varphi^t) u_2^t(\cdot|\underline{\theta}). \tag{5}$$

(iii) If the belief $\tilde{\varphi}^t$ solving (5) is strictly smaller than the prior φ^0 , then P2 continues wp 1.

Corollary 1. P 2's beliefs φ^t are uniquely defined on and off the equilibrium path whenever $o_2^{t-2} + o_1^{t-1} \leq \bar{\theta} - 1$ because, at all these periods, the indifference (5) must hold. Otherwise P 2's equilibrium behaviour could be exploited by at least one of the types of P 1 (which is impossible in equilibrium). Off the equilibrium path, P 2's beliefs are undefined if $o_2^{t-2} + o_1^{t-1} > \bar{\theta} - 1$ as this must lead to P 2 accepting at t wp1 for any belief.

Thus, in equilibrium, P2 ascribes any observed continuation bid to either type P1 wpp unless any possible equilibrium continuation payoff for P2 is lower than this offer. Hence, any fully separating equilibrium action of this game are followed immediately by P2's accepting. P2 responds to any other observed continuation by P1 with her (completely specified) equilibrium response entailing the belief uniquely defined by Lemma 2. Put the other way, no fully separating equilibria can exist in which different types of P1 set different observable continuation actions if P2's subsequent continuation is not ruled out by (4).

Corollary 2. In equilibrium and at each stage, P1 bids only actions in the support of $\bar{P}1$'s equilibrium offers at that stage, accepts or mixes among these.

Next, P1's mixture probability α_1^t is determined using the above results implying that P2's posterior beliefs φ^t only depend on her prior and the most recently observed offer o_1^{t-1} . As beliefs stem from Bayes' rule whenever possible, after observing any o_1^{t-1} , P2 updates using P1's ($\bar{P}1$'s known equilibrium continuation probability $\underline{\alpha}_1^{t-1}$ ($\bar{\alpha}_1^{t-1}$ which is later shown to equal 1) to

$$\varphi^t = \text{pr}(\bar{\theta}|o_1^{t-1}) = \frac{\text{pr}(o_1^{t-1}|\bar{\theta})\text{pr}(\bar{\theta})}{\text{pr}(o_1^{t-1}|\bar{\theta})\text{pr}(\bar{\theta}) + \text{pr}(o_1^{t-1}|\underline{\theta})\text{pr}(\underline{\theta})} = \frac{\bar{\alpha}_1^{t-1}\varphi^{t-2}}{\bar{\alpha}_1^{t-1}\varphi^{t-2} + \underline{\alpha}_1^{t-1}(1 - \varphi^{t-2})}.$$

Lemma 3. In any equilibrium of \mathcal{Q} , at $t - 1$: (i) $\underline{P1}$ bids with probability α_1^{t-1} $\overline{P1}$'s equilibrium action o_1^{t-1} and accepts with $(1 - \alpha_1^{t-1})$ iff $\exists \varphi^t \in (\varphi^0, 1)$ solving (5). (ii) $\underline{P1}$ accepts wp1 if P2 accepts wp1 at t . (iii) $\underline{P1}$ continues wp1 if the φ^t solving (5) is smaller or equal to P2's prior φ^0 .

The following two lemmata make the case against jump bidding, i.e. bids which increase the previous own bid by more than the minimal amount. They start from the acceptance condition developed in Lemma 1 and construct a unique path backwards to the initial node. The intuition against jump bids is that (i) P2 will not attribute an observed jump to a specific type of P1 and (ii) the acceptance condition (4) implies that, in equilibrium, an own jump reduces the highest offer forthcoming from the opponent while (iii) increasing the payment to the opponent if he accepts this increased offer. Intuitively, notice that if P1 chooses any action which makes the still uninformed P2 strictly prefer accepting, this is beneficial only to $\overline{P1}$ while continuation with probability one is beneficial only to $\underline{P1}$. Thus any such action would be fully and instantly revealing and 'faking' the other type cannot be equilibrium behaviour. Therefore, since P1 cannot bluff the opponent into a deterministic response, a jump bid will (in expectation) hurt P1's payoffs against a rational and stage-by-stage mixing and minimally increasing P2. The only way of actually benefitting from his informational advantage is for P1 to use this information truthfully, i.e. by making it gradually known to P2 at minimal cost through minimally increasing own offers.

Lemma 4. In any equilibrium of \mathcal{Q} , $\overline{P1}$ starts by offering $o_1^1 = 1$ and increases his previous bid o_1^{t-2} minimally wp1 whenever $o_2^{t-1} + o_1^t < \bar{\theta}$.

Since there is a unique mixture probability α_1^{t-1} generating the equilibrium belief φ^t for any given history, this is the only mixture probability compatible with P2's equilibrium behaviour. We remark that $\underline{P1}$'s mixture probability α_1^t derived in the proof of Lemma 3 is decreasing in t and P2's mixture probability α_2^t from the proof of Lemma 4 is increasing in t . Moreover, P2's continuation probability α_2^t is increasing in P1's observed bid o_1^{t-1} .

Lemma 5. In any equilibrium of \mathcal{Q} , at all t where $o_2^{t-2} + o_1^{t-1} \leq \bar{\theta} - 1$, P2 mixes with probability $\alpha_2^t \in [0, 1]$ between accepting and increasing her previous offer o_2^{t-2} minimally.

Definition 1. Among the strategy profiles satisfying Lemmata 1–5, we define the profile β^* in which (i) $\overline{P1}$ increases his offer at t wp1 if $o_1^{t-2} + o_2^{t-1} = \bar{\theta} - 1$ and (ii) P2's belief φ^t is set to 1 whenever P2 accepts wp1 for any belief. We summarise the properties of β^* as:

- (a) Players accept wp1 at stage t if $o_1^{t-2} + o_2^{t-1} > \bar{\theta} - 1$.
- (b) At each t , players $\underline{P1}$ and P2 increase their previous offer o_i^{t-2} by the minimal increment of 1 wpp α_1^t and α_2^t , respectively, as long as $o_1^{t-2} + o_2^{t-1} < \bar{\theta} - 1$ and accept with the complimentary probability. $\overline{P1}$ bids $o_1^{t-2} + 1$ wp1 if $o_1^{t-2} + o_2^{t-1} \leq \bar{\theta} - 1$.
- (c) P2's belief φ^t is uniquely determined through backward induction as making her exactly indifferent between accepting and offering $o_2^{t-2} + 1$. If this indifference belief is smaller than φ^0 , P2 continues wp1 and accepts wp1 (setting $\varphi^t = 1$) if the required belief is at least 1.
- (d) $\underline{P1}$'s continuation probability α_1^{t-1} relates to P2's belief φ^t uniquely through Bayes' rule. There is a unique continuation probability α_2^t compatible with $\underline{P1}$'s mixing with α_1^{t-1} .
- (e) At some stage $t_s - 1$, the backward induction process defining beliefs and mixture probabilities either reaches the initial node or the belief required for P2's indifference is lower than P2's prior. Then, for $t < t_s$, players increase their offers minimally wp1, i.e. $\alpha_i^t = 1$.

The next proposition summarises that the above lemmata describe all equilibria of the bargaining game \mathcal{Q} . Since the profiles only differ in payoff irrelevant beliefs and, potentially, one mixed terminal stage action over identical outcomes, the proposition also shows outcome uniqueness.

Proposition 1. Any profile satisfying Lemmata 1–5 is a weak perfect Bayesian Nash equilibrium of \mathcal{Q} .¹⁵ All equilibria are outcome equivalent to β^* .

Having established outcome uniqueness, we use the equilibrium properties to approximate the (prior dependent) period at which P2 accepts wp1 unless $\underline{P1}$ starts to mix in equilibrium.

Proposition 2. In β^* , $\underline{P1}$ starts mixing at t_s^* , the first odd period following $t_s = (\lfloor \bar{\theta} \rfloor - 3)/4$.

Using this result, Corollary 3 (in Appendix) approximates equilibrium payoffs as¹⁶

$$u(\beta^*|\underline{\theta}) \approx \left(\frac{\bar{\theta} - 7}{8}, -\frac{\bar{\theta} - 7}{8} \right), \quad u(\beta^*|\bar{\theta}) \approx \left(\frac{5\bar{\theta} - 11}{8}, \frac{3\bar{\theta} + 11}{8} \right).$$

¹⁵ β^* is a weak perfect Bayesian Nash equilibrium because all constraints on off-equilibrium-path beliefs are made through the structure of the game and not through the equilibrium concept.

¹⁶ This payoff approximation is for odd $\lfloor \bar{\theta} \rfloor$. It is imprecise for low $\bar{\theta}$ and stated for convenience only.

Thus the informed player benefits from a refinement of the bidding grid (equivalent to increasing $\bar{\theta}$) and obtains an information rent even when the object is worthless. Normalising $\bar{\theta} = 1$ and letting the implied bidding grid becoming very fine, one obtains the limit payoffs

$$u(\beta^*|\underline{\theta}) \approx \left(\frac{1}{8}, -\frac{1}{8}\right), \quad u(\beta^*|\bar{\theta}) \approx \left(\frac{5}{8}, \frac{3}{8}\right).$$

We finally show that a variant of β^* remains an equilibrium of the extended game \mathcal{E} allowing for general offers. The proof requires assumptions on infinite deviation paths and their payoffs which we relegate to the appendix. We denote a non-increasing bid by ‘0’ and by $b_i^t(\beta)$ the continuation part o_i^t of $\beta_i^t = [\alpha_i^t : o_i^t, q]$ (recall that $o_2^0 = o_1^{-1} = 0$). The following definition gives an algorithm providing a pair of (mixed) stage action and belief for any history in \mathcal{E} .

Definition 2. We define the strategy profile $^*\beta_\infty$ and beliefs φ_∞^t , for all $t > 0$ and any history of play $h^t = (o_1^1, o_2^2, \dots, o_j^{t-3}, o_i^{t-2}, o_j^{t-1})$, as the sequence of (mixed) action, belief pairs

$$^*(\beta_{i,\infty}^t, \varphi_\infty^t)|h^t = \begin{cases} ^*(\beta_i^t, \varphi^t)|h^t & \text{if } t = 1 \vee o_j^{t-1} \geq b_j^{t-1}(\beta^*), \\ ^*(\beta_i^{t-2}, \varphi^{t-2})|h^{t-2} & \text{if } \underbrace{(o_j^{t-1} = o_i^{t-2} = \text{‘0’}) \wedge o_j^{t-3} \geq b_j^{t-3}(\beta^*)}_{\text{complete ‘0’-cycle}}; t = t - 2, \\ o_i^{t-2} & \text{otherwise} \end{cases}$$

where $^*(\beta_i^t, \varphi^t)$ are the equilibrium stage action and belief prescribed by β^* in \mathcal{Q} after h^t .

Proposition 3. $^*\beta_\infty$ is an equilibrium of the extended game \mathcal{E} allowing for non-increasing bids. The outcome from \mathcal{E} following $^*\beta_\infty$ equals that from β^* in \mathcal{Q} .

5. Concluding remarks

Further direct applications of our framework include partnership dissolution problems where two asymmetrically informed players (e.g. one active and one silent partner) jointly own a firm, agreeing on a profit sharing rule between two firms involved in a joint venture and deciding whether to spin-off some yet-to-be-proven innovation (‘selling the project to the manager’) or developing it inside the firm. It can be used to solve the problems of splitting an inheritance (e.g. an Amish farm or company) under the provision of maintaining it as a unit or—as the sign on a pie’s value is of no consequence—of agreeing on payments for hosting some airport or waste disposal site in one community although more communities profit from it. ‘Fair Buy-Sell’ (<http://www.fairoutcomes.com/>) offers a commercial procedure for settling related disputes and describes further applications.

Our assumptions on priors and possible bids extend easily. Allowing for incomplete information on both sides changes the problem fundamentally and the general case requires separate study. Preliminary studies indicate, however, that outcome uniqueness will not necessarily be lost. Generalising the type space to a larger set retains the result in the sense that some informed types will always probabilistically mimic the highest type. We depart from the standard model by analysing a commonly valued pie without depreciation during the bargaining process. The first assumption focusses attention on the signalling of private information. From the point of view of the communication game, introducing a heterogenous, commonly known valuation function which only depends on the common value parameter does not change the equilibrium dynamics—the player’s signalling would still be on the common value part. We make the second assumption because we are interested in situations such as cake-cutting or automated bargaining where the object’s value does not change during the negotiation period.

Appendix A. Proofs

Proof of Lemma 1.

- (i) As we are interested in the highest possible bids, we only need to consider $\theta = \bar{\theta}$. Since admissible bids are bounded from above and given that no player has accepted yet, offers in the finite game must eventually reach the upper bound \bar{B} . Suppose that P_i makes the last admissible continuation bid $o_i^{\hat{t}} = \bar{B} > \theta$ at some time \hat{t} . Then, at $\hat{t} + 1$, P_j must accept P_i ’s offer through accepting wp1. P_i ’s payoff at $\hat{t} + 1$ is then $\bar{\theta} - o_i^{\hat{t}}$ and P_j ’s payoff is $o_i^{\hat{t}}$. Since $\bar{\theta} - o_i^{\hat{t}} = \bar{\theta} - \bar{B} < 0$, however, P_i can do better by accepting at \hat{t} and accepting P_j ’s offer $o_j^{\hat{t}-1} > 0$. Knowing that, P_j will also accept if her time $t - 1$ acceptance payoff o_i^{t-2} exceeds her time t continuation payoff $\bar{\theta} - o_j^{t-1}$, we obtain P_i ’s acceptance condition wp1 at

general t as $o_i^{t-1} + o_i^t > \bar{\theta}$ or, for a minimally increasing o_i^t ,

$$o_i^{t-2} + o_j^{t-1} > \bar{\theta} - 1.$$

Therefore, P_i will not make a bid $o_i^t > \bar{\theta} - o_j^{t-1}$ because this would lead to P_j accepting at $t + 1$ wp1 giving a payoff of $\bar{\theta} - o_i^t < o_j^{t-1}$ while accepting at t secures o_j^{t-1} .

- (ii) Since $o_1^{t-2} + o_2^{t-1} < \bar{\theta} - 1$, $\bar{P1}$ can get $\bar{\theta} - o_1^t$ if $P2$ accepts at the following stage or $o_2^{t+1} > o_2^{t-1}$ if she does not and $\bar{P1}$ accepts the stage after. Thus, there is a feasible bid $o_1^t < \bar{\theta} - o_2^{t-1}$ which secures $\bar{P1}$ at least $\min\{\bar{\theta} - o_1^t, o_2^{t+1}\}$ which is strictly higher than her payoff from accepting at t of o_2^{t-1} .
- (ii) Since $\bar{P1}$'s time t acceptance and continuation payoffs are identical for $o_1^t = o_1^{t-2} + 1$ if $o_1^{t-2} + o_2^{t-1} = \bar{\theta} - 1$, both actions (and any mixture between) are equally good for both $\bar{P1}$ and $P2$. The same applies to $P2$ if she knows that $\theta = \bar{\theta}$ wp1 (i.e. after full revelation).

□

Proof of Lemma 2.

- (i) We show first that in every candidate separating equilibrium not followed by $P2$ accepting wp1 the following stage, there exists a profitable deviation for $\underline{P1}$.

$\underline{\theta}$: The only equilibrium of the complete information, zero-sum game for $\theta = 0$ is for both players to accept whatever the opponent does. The same is true in the continuation game after the revelation of $\theta = 0$. Thus any strategy which reveals $P1$'s type as $\theta = 0$ must end with $\underline{P1}$ accepting at $t - 1$ wp1 giving $o_2^{t-2} > 0$ because, otherwise, $P2$ would accept wp1 at t giving $-o_1^{t-1} < 0$ to $\underline{P1}$. Therefore, the only possible continuation separating action must reveal the value of the object as high.

$\bar{\theta}$: Let $\bar{P1}$ continue at $t - 1$ wp1 while $\underline{P1}$ accepts wp1 and reveals $\theta = \bar{\theta}$ through his observed equilibrium bid o_1^{t-1} while there is an admissible successive o_2^t such that $o_1^{t-1} + o_2^t < \bar{\theta}$. Then $P2$ will continue at t wp1 because by offering o_2^t , she can get at least $\bar{\theta} - o_2^t > o_1^{t-1}$. Thus after the high-value state is revealed, $P2$'s continuation payoff is necessarily higher than her acceptance payoff at that stage. But given that $P2$ continues wp1, $\underline{P1}$ will deviate to $\bar{P1}$'s action to obtain $o_2^t > o_2^{t-2}$ and thereby destroy the candidate separating equilibrium.

The above argument forbids any partially revealing action not leading to indifference by $P2$ at the following stage because, in the first case of $\underline{\theta}$, $P2$ will accept wp1 for any belief which does not make her at least indifferent causing $\underline{P1}$ to accept the period before. In the second case of $\bar{\theta}$, $P2$ continues wp1 for any belief not making her at most indifferent. This presents a certain way for $\underline{P1}$ to improve over his mixing payoffs granting him only o_2^{t-2} .

- (ii) Whenever $P2$'s continuation at t is not ruled out by (4), a belief $\tilde{\varphi}^t$ such that $o_1^{t-1} \neq \tilde{\varphi}^t u_2^t(\cdot|\underline{\theta}) + (1 - \tilde{\varphi}^t) u_2^t(\cdot|\bar{\theta})$ cannot be part of an equilibrium. Thus (5) must hold.
- (iii) When the prior-based continuation payoff is strictly higher than $P1$'s initial offers, $P2$ cannot possibly be made indifferent for any observed o_1^{t-1} . Hence she continues wp1. □

Proof of Lemma 3.

- (i) $P2$ is willing to mix in equilibrium if there is a belief $\varphi^t \in (\varphi^0, 1)$ solving (5). Thus, for a rational $\varphi^t \in (\varphi^0, 1)$, Bayes' rule determines a particular continuation probability α_1^{t-1} for both types of $P1$ on an observed equilibrium bid o_1^{t-1} . Hence $\underline{P1}$ mixes at $t - 1$ if he is indifferent between accepting and offering (any pure) o_1^{t-1}

$$o_2^{t-2} = (1 - \alpha_2^t)(-o_1^{t-1}) + \alpha_2^t(o_2^t) \Leftrightarrow \alpha_2^t = \frac{o_2^{t-2} + o_1^{t-1}}{o_2^t + o_1^{t-1}} \in (0, 1) \tag{A.1}$$

implying that $P2$ must mix if $\underline{P1}$ mixes.¹⁷

- (ii) If $P2$ accepts at t wp1, $\underline{P1}$ gets $-o_1^{t-1}$ by continuing but o_2^{t-2} by accepting and thus accepts wp1.
- (iii) If $P2$ bids o_2^t at t wp1, $\underline{P1}$ gets o_2^{t-2} by accepting wpp but gets o_2^t wp1 by continuing and thus continues wp1. □

Proof of Lemma 4. Consider the first period t in any sequence of play where Lemma 1 prescribes accepting wp1 for either $P2$ or $\bar{P1}$, that is, $o_i^{t-2} + o_j^{t-1} > \bar{\theta} - 1$. There are two cases:

¹⁷ As all pure actions in the support of a mixed action give the same payoff, it is immaterial whether $\bar{P1}$'s equilibrium bid is mixed or pure. Using the observed o_1^{t-1} is enough to determine $\underline{P1}$'s mixture condition.

1. P2 accepts at t wp1 with terminal payoffs of $(\bar{\theta} - o_1^{t-1}, o_1^{t-1})$ ('odd' $[\bar{\theta}]$ case).

(i) At $t - 1$, $\bar{P1}$ maximises his payoff of $\bar{\theta} - o_1^{t-1}$ by minimising his bid o_1^{t-1} , that is, by minimally increasing his previous bid $o_1^{t-3} + 1 = o_1^{t-1}$. He is indifferent between $o_1^{t-3} + 1$ and accepting if exactly $o_1^{t-3} + 1 + o_2^{t-2} = \bar{\theta}$. In this case, any mixture gives the same outcome to both players (in β^* , $\bar{P1}$ bids $o_1^{t-3} + 1$ wp1)

$$\bar{P1} : \bar{\theta} - o_1^{t-3} - 1 = o_2^{t-2}, \quad P2 : o_1^{t-3} + 1 = \bar{\theta} - o_2^{t-2}. \tag{A.2}$$

($\bar{P1}$ accepts wp1.) Combined time t accepting $o_1^{t-1} + o_2^{t-2} > \bar{\theta} - 1$ and time $t - 1$ continuation $o_2^{t-2} + o_1^{t-1} \leq \bar{\theta}$ imply that, in equilibrium, P2 accepts at stage t when

$$o_1^{t-1} + o_2^{t-2} = [\bar{\theta}]. \tag{A.3}$$

(ii) At $t - 2$, P2 is indifferent between accepting and continuing (Lemma 2). As $\bar{P1}$ mixes at $t - 3$ (Lemma 3), she mixes with probability α_2^{t-2} . Hence at $t - 3$, $\bar{P1}$ gets

$$(1 - \alpha_2^{t-2})(\bar{\theta} - o_1^{t-3}) + \alpha_2^{t-2}(\bar{\theta} - (o_1^{t-3} + 1)).$$

Since $\bar{\theta} - o_1^{t-3} > \bar{\theta} - o_1^{t-3} - 1$, $\bar{P1}$ maximises his payoff by bidding $o_1^{t-3} = o_1^{t-5} + 1$.

(iii) At $t - 4$, P2 is again indifferent between accepting and continuing. Since $\bar{P1}$ mixes at $t - 5$, she continues with $\alpha_2^{t-4} \in (0, 1)$. Thus $\bar{P1}$ expects at $t - 5$

$$(1 - \alpha_2^{t-4})(\bar{\theta} - o_1^{t-5}) + \alpha_2^{t-4}\alpha_2^{t-2}(\bar{\theta} - (o_1^{t-5} + 2)).$$

A jump bid of $o_1^{t-5} > o_1^{t-7} + 1$ decreases both $(\bar{\theta} - o_1^{t-5})$ and $(\bar{\theta} - (o_1^{t-5} + 2))$. P2's continuation probability α_2^{t-4} in (A.1) is increasing in o_1^{t-5} and that $\bar{\theta} - o_1^{t-5} > \bar{\theta} - o_1^{t-5} - 2$. Thus both a and the probability of getting a decrease. Similarly, the probability for the lower payoff b increases in o_1^{t-5} . Thus, $\bar{P1}$ bids $o_1^{t-5} = o_1^{t-7} + 1$.

(iv) The same argument applies until either P2 stops mixing or the initial node is reached. When P2 stops mixing at some \tilde{t} , the above mixture probabilities $\alpha_2^t = 1$ for all $t \leq \tilde{t}$.

2. $\bar{P1}$ accepts at t wp1 with terminal payoffs of $(o_2^{t-1}, \bar{\theta} - o_2^{t-1})$ ('even' $[\bar{\theta}]$ case).

(i) At $t - 1$, P2 is indifferent between accepting and continuing (Lemma 2). Combined time t accepting $o_1^{t-2} + o_2^{t-1} > \bar{\theta} - 1$ and time $t - 1$ continuation $o_1^{t-2} + o_2^{t-1} \leq \bar{\theta}$ imply

$$o_2^{t-1} + o_1^{t-2} = [\bar{\theta}] \tag{A.4}$$

at t . As $\bar{P1}$ mixes at $t - 2$ (Lemma 3), P2 mixes with probability $\alpha_2^{t-1} \in (0, 1)$ and $\bar{P1}$ faces the time $t - 2$ payoff

$$(1 - \alpha_2^{t-1})(\bar{\theta} - o_1^{t-2}) + \alpha_2^{t-1}(o_2^{t-1}), \quad \text{where } o_2^{t-1} < \bar{\theta} - o_1^{t-2}.$$

Both $(\bar{\theta} - o_1^{t-2})$ and the accepting probability $(1 - \alpha_2^{t-1})$ are lowered by a jump bid as under (iii) above and $o_1^{t-2} < \bar{\theta} - o_2^{t-1}$ because otherwise P2 would accept at $t - 1$. Since (A.4) holds when P1 accepts at t in equilibrium, an increase in o_1^{t-2} decreases o_2^{t-1} in equilibrium. As P2's continuation probability (A.1) of getting $o_2^{t-1} < \bar{\theta} - o_1^{t-2}$ increases in a jump bid o_1^{t-2} , $\bar{P1}$ maximises his payoff by bidding $o_1^{t-2} = o_1^{t-4} + 1$.

(ii) The same reasoning as under 1 applies.

Having shown that $\bar{P1}$'s continuation action is pure and since $\bar{P1}$ mixes because P2 does (Lemmata 2 and 3), P2's indifference condition fully determines $\bar{P1}$'s continuation probability α_1^{t+1}

$$o_1^{t-1} = (1 - \varphi^t)[(1 - \alpha_1^{t+1})(-o_2^t) + \alpha_1^{t+1}(o_1^{t+1})] + \varphi^t(o_1^{t+1}) \Leftrightarrow \alpha_1^{t+1} = \frac{o_1^{t-1} - o_1^{t+1}\varphi^t + (1 - \varphi^t)o_2^t}{(1 - \varphi^t)(o_1^{t+1} + o_2^t)} \in (0, 1). \quad \square \tag{A.5}$$

Proof of Lemma 5. We consider the same two cases as in the proof of Lemma 4:

1. P2 accepts at t wp1 with terminal payoffs of $(\bar{\theta} - o_1^{t-1}, o_1^{t-1})$ ('odd' $[\bar{\theta}]$ case).

- (i) At $t - 1$, $\overline{P1}$ bids $o_1^{t-1} = o_1^{t-3} + 1$ wp1 and $\underline{P1}$ accepts wp1. At $t - 2$, P2 expects $(1 - \varphi^{t-2})(-o_2^{t-2}) + \varphi^{t-2}(o_1^{t-1})$ where o_1^{t-1} decreases in equilibrium through an increased o_2^{t-2} through the acceptance condition (A.3). Hence P2 maximises her payoff by bidding $o_2^{t-2} = o_2^{t-4} + 1$. Moreover, since her equilibrium belief φ^{t-2} ensures that this payoff expectation exactly equals her acceptance payoff of o_1^{t-3} , she is willing to mix between accepting and bidding $o_2^{t-2} = o_2^{t-4} + 1$.
- (ii) At $t - 3$, $\overline{P1}$ bids o_1^{t-3} wp1 and $\underline{P1}$ mixes between continuing and accepting. At $t - 4$, P2 faces a payoff of

$$(1 - \varphi^{t-4})[(1 - \alpha_1^{t-3})(-o_2^{t-4}) + \alpha_1^{t-3}\{(1 - \alpha_2^{t-2})(-o_1^{t-3}) + \alpha_2^{t-2}(-o_2^{t-2})\}] + \varphi^{t-4}[(1 - \alpha_2^{t-2})o_1^{t-3} + \alpha_2^{t-2}(o_1^{t-1})] \quad (A.6)$$

where $o_1^{t-3} = o_1^{t-1} - 1$ is linked to P2's current bid o_2^{t-4} through the final acceptance condition $o_2^{t-2} + o_1^{t-1} = \lceil \bar{\theta} \rceil$ and thus both o_1^{t-3} and o_1^{t-1} decrease indirectly through a jump bid at $t - 4$. Finally, $o_2^{t-4} > o_2^{t-6} + 1$ reduces the subsequent time $t - 3$ and time $t - 1$ payoffs directly. Showing that the continuation probabilities α_1 (A.5) and α_2 (A.1) are reinforcing this effect is more involved but can be easily seen from the fact that $\underline{P1}$ is indifferent between any equilibrium jump bids by P2 but $\overline{P1}$'s payoff

$$(1 - \alpha_2^{t-2})(\bar{\theta} - o_1^{t-3}) + \alpha_2^{t-2}(\bar{\theta} - o_1^{t-1})$$

increases strictly through the final acceptance condition $o_2^{t-2} + o_1^{t-1} = \lceil \bar{\theta} \rceil$. As the game is constant sum, a jump bid cannot increase P2's payoff (A.6).

- (iii) The same argument applies until either P2 stops mixing or the initial node is reached. When P2 stops mixing at some \tilde{t} , the above mixture probabilities $\alpha_2^t = 1$ for all $t \leq \tilde{t}$.

2. $\overline{P1}$ accepts at t wp1 with terminal payoffs of $(o_2^{t-1}, \bar{\theta} - o_2^{t-1})$ ('even' $\lceil \bar{\theta} \rceil$ case). The analysis is identical to the above when the final time t , high-state payoff is replaced by $(o_2^{t-1}, \bar{\theta} - o_2^{t-1})$. All other expressions and arguments remain unchanged. \square

Proof of Proposition 1. Existence is ensured as the equilibrium is explicitly constructed. Uniqueness follows from the fact that, in equilibrium, \underline{Q} ends wp1 only when (A.3 and A.4) are triggered. Although there are many possible histories leading to these conditions, Lemmata 4 and 5 show that only the minimal-increase profile is compatible with both the terminal acceptance conditions and maximisation at each stage. Lemma 2 supplies the unique beliefs through arguing that P1 cannot credibly reveal his information as long as this benefits only one of his possible types. The unique mixture probabilities which turn the minimal-increase profile from Lemmata 4 and 5 into an equilibrium are supplied by Lemma 3. The only source of non-uniqueness of equilibrium actions is at $\overline{P1}$'s terminal equilibrium continuation action where mixture probabilities may be arbitrary between minimal increase and accepting for integer $\bar{\theta}$. Any mixture gives the same outcomes (A.2). The only case where beliefs are arbitrary is when they do not matter, namely, when an out of equilibrium action leads to P2 accepting immediately for any belief. \square

Proof of Proposition 2. Denote the (odd-valued) period where $\underline{P1}$ starts mixing by t_s . In an equilibrium β^* , $t_s + 1$ is the first period in which P2's prior-based payoff expectation from β^* is lower than her sure payoff from accepting $o_1^{t_s}$. Hence, $\underline{P1}$ must mix at t_s in order to manipulate P2's beliefs or else she will accept at $t_s + 1$. Thus, on the equilibrium path, P2 accepts at $t_s + 1$ if¹⁸

$$o_1^{t_s} = \frac{t_s + 1}{2} > \varphi^0 u_2^{t_s+3}(\beta^* | \bar{\theta}) - (1 - \varphi^0) o_2^{t_s+1} = \frac{1}{2} u_2^{t_s+3}(\beta^* | \bar{\theta}) - \frac{t_s + 1}{4}. \quad (A.7)$$

The equilibrium continuation payoff in (A.7) consists of P2's low value payoff of $o_1^{t_s}$ given through $\underline{P1}$'s mixing the stage before and P2's high value continuation payoff at the stage following $t_s + 1$ (which is $t_s + 3$ because $\overline{P1}$ does not accept wpp). We write P2's acceptance payoff at t as $u_2^t(q)$. P2's expected continuation payoff $u_2^{t_s+3}(\beta^* | \bar{\theta})$ given a high valued object is then obtained recursively as

$$u_2^{t_s+3}(\beta^* | \bar{\theta}) = (1 - \alpha_2^{t_s+3}) u_2^{t_s+3}(q) + \alpha_2^{t_s+3} u_2^{t_s+5}(\beta^* | \bar{\theta}).$$

It turns out to be more convenient to rewrite this, for odd $\lceil \bar{\theta} \rceil$, in the form¹⁹

$$u_2^{t_s+1}(\beta^* | \bar{\theta}) = \sum_{\tau=1}^{(\lceil \bar{\theta} \rceil - t_s)/2} \left(\prod_{t=1}^{\tau-1} \alpha_2^{2t+t_s-1} \right) (1 - \alpha_2^{2\tau+t_s-1}) o_1^{2\tau+t_s-2} + \left(\prod_{t=1}^{(\lceil \bar{\theta} \rceil - t_s)/2} \alpha_2^{2t+t_s-1} \right) o_1^{\lceil \bar{\theta} \rceil}$$

where the summation adds $\overline{P1}$'s stage payoff from P2 accepting for every second period between $t_s + 1$ and $\lceil \bar{\theta} \rceil - 1$. The first product is the probability of reaching P2's choice node and the second product is the probability of reaching the final node

¹⁸ Choosing a higher prior than $\varphi^0 = 1/2$ will increase P2's equilibrium payoff expectation, increase t_s^* and thus decrease $\overline{P1}$'s information rent.

¹⁹ The expression for even $\lceil \bar{\theta} \rceil$ is similar and given by replacing $\lceil \bar{\theta} \rceil$ by $\lceil \bar{\theta} \rceil + 1$ in (A.8) and changing the last term from $(\lceil \bar{\theta} \rceil + 1)/2$ to $(3\lceil \bar{\theta} \rceil - 2\bar{\theta} + 2)/2$.

where P2 accepts wp1. Since P1's equilibrium path offer is $o_1^t = (t + 1)/2$ and α_2^t is given by (A.1), this simplifies to

$$\begin{aligned}
 u_2^{t_s+1}(\beta^*|\bar{\theta}) &= \frac{1}{2} \sum_{\tau=1}^{(\bar{\theta}_J-t_s)/2} \left(\prod_{t=1}^{\tau-1} \frac{2t+t_s-2}{2t+t_s-1} \right) + \left(\prod_{t=1}^{(\bar{\theta}_J-t_s)/2} \frac{2t+t_s-2}{2t+t_s-1} \right) \frac{|\bar{\theta}_J+1}{2} \\
 &= \frac{1-t_s}{2} + 2^{t_s-|\bar{\theta}_J-1} \frac{(|\bar{\theta}_J+1)!}{(((|\bar{\theta}_J+1)/2)!)^2} \frac{(((t_s-1)/2)!)^2}{(t_s-1)!} \frac{|\bar{\theta}_J+1}{2} \\
 &= \frac{1-t_s}{2} + 2^{t_s-|\bar{\theta}_J-1} \left(\frac{|\bar{\theta}_J+1}{2} \right) \left(\frac{t_s-1}{2} \right)^{-1} \frac{|\bar{\theta}_J+1}{2} \\
 &= \frac{1-t_s}{2} + 2 \frac{\Gamma((|\bar{\theta}_J+2)/2)}{\Gamma((|\bar{\theta}_J+1)/2)} \frac{\Gamma((t_s+1)/2)}{\Gamma(t_s)/2}
 \end{aligned} \tag{A.8}$$

where we use $\prod_{t=1}^{(\bar{\theta}_J-t_s)/2} t = ((|\bar{\theta}_J-t_s)/2)!$ together with the definition of the central binomial coefficient

$$\binom{2n}{n} = \frac{(2n)!}{(n!)^2}. \tag{A.9}$$

The Euler Gamma function Γ is used in the last step to obtain a more compact result (the intermediate steps are tedious and thus omitted). Since no closed form representation of (A.9) is known, we use Stirling's approximation $n! \approx \sqrt{2\pi n} n^n e^{-n}$ to approximate (A.9) as $4^n / (\sqrt{n}\sqrt{\pi})$.²⁰

We use this, in turn, to approximate (A.8) as

$$u_2^{t_s+1}(\beta^*|\bar{\theta}) \approx \frac{1-t_s}{2} + \frac{2\sqrt{t_s-1}}{\sqrt{|\bar{\theta}_J+1}} \frac{|\bar{\theta}_J+1}{2}. \tag{A.10}$$

Plugging this approximation for $u_2^{t_s+3}(\beta^*|\bar{\theta})$ back into (A.7) gives

$$\frac{t_s+1}{2} > \frac{1}{2} \left(\frac{1-(t_s+2)}{2} + \frac{2\sqrt{(t_s+2)-1}}{\sqrt{|\bar{\theta}_J+1}} \frac{|\bar{\theta}_J+1}{2} \right) - \frac{1}{2} \left(\frac{t_s+1}{2} \right) \Leftrightarrow t_s^* > \frac{|\bar{\theta}_J-3}{4}.$$

As P2 will accept at $t_s + 1$ if her beliefs are not adjusted by a mixed action of P1 at the first possible t_s , the equilibrium value of t_s is given by the first odd period after the above t_s^* . \square

Corollary 3. For odd-valued $|\bar{\theta}_J$, the players' expected payoffs from β^* are given by²¹

$$\begin{aligned}
 u(\beta^*|\underline{\theta}) &= \left(\frac{t_s^*-1}{2}, -\frac{t_s^*-1}{2} \right) \approx \left(\frac{\bar{\theta}-7}{8}, -\frac{\bar{\theta}-7}{8} \right), \\
 u(\beta^*|\bar{\theta}) &= \left(\bar{\theta} - u_2(\beta^*|\bar{\theta}), \underbrace{\frac{1-t_s^*}{2} + 2 \frac{\Gamma((|\bar{\theta}_J+2)/2)}{\Gamma((|\bar{\theta}_J+1)/2)} \frac{\Gamma((t_s^*+1)/2)}{\Gamma((t_s^*)/2)}}_{=u_2(\beta^*|\bar{\theta})} \right) \approx \left(\frac{5\bar{\theta}-11}{8}, \frac{3\bar{\theta}+11}{8} \right).
 \end{aligned} \tag{A.11}$$

P2's precise payoff $u_2(\beta^*|\bar{\theta})$ is (A.8). Because the high-value game is ex post $\bar{\theta}$ -sum, $\bar{P1}$'s payoff is just $u_1(\beta^*|\bar{\theta}) = \bar{\theta} - u_2(\beta^*|\bar{\theta})$. P1's payoff is given by t_s^* through $u_1(\beta^*|\underline{\theta}) = u_1^*(q) = o_2^{t_s^*-1} = (t_s^*-1)/2$ because all pure actions in the support of P1's first mixed action must give the same payoff. Since this branch of the game is ex post zero-sum, $u_2(\beta^*|\underline{\theta}) = -u_1(\beta^*|\underline{\theta})$. Thus for odd $|\bar{\theta}_J$ and t_s^* defined as the first odd period after $(|\bar{\theta}_J-3)/4$, payoffs are indeed given by (A.11). The approximations are obtained by setting $\bar{\theta} \approx |\bar{\theta}_J$, ignoring the required rounding up of t_s^* to the next odd integer period and plugging $t_s = (\bar{\theta}-3)/4$ into the approximation (A.10).

Infinite game

In order to examine robustness, we analyse the game allowing non-increasing offers. To fully define this extended game \mathcal{E} , payoffs must be assigned to infinite sequences of '0'-deviations (which we call *cycles*). We define the infinite cycle payoff

²⁰ The asymptotic error involved in Stirling's approximation is of order $1/n$, so it vanishes as $\bar{\theta}$ gets large.

²¹ The corresponding payoffs for even $|\bar{\theta}_J$ can be easily calculated from Proposition 2.

as the limit of finite-cycle payoffs.²² A finite cycle has the form: follow β^* up to some period, then cycle '0' an even finite number of times, and then continue β^* . For any finite cycle, this payoff is $u(\beta^*)$; we now extend this to infinite cycles.

Assumption 1. '0'-Deviations from any equilibrium of ε are believed to be made wp1. The payoffs from following β^* up to some period and then cycling infinitely by each player offering '0' wp1, are given by the equilibrium payoff of the finite game $u(\beta^*)$.

Proof of Proposition 3. We start by interpreting '0' as repeating the previous own offer. In order to confirm β_∞^* as an equilibrium of the infinite game, we first confirm optimality of the prescribed actions on the equilibrium path, then on any deviation path and then show that no deviations from the deviation path are profitable.

1. Given β_∞^* , bidding '0' wp1 cannot be a profitable deviation from the equilibrium path because this is (in equilibrium) followed by the opponent bidding '0' (wp1) which leaves the game in precisely the state before entering the '0'-cycle.²³ Thus bidding '0' gives the equilibrium payoffs. The same is true for any finite repetition, and, through Assumption 1, for any infinite repetition of cycles as well.
2. Conforming to β_∞^* and bidding '0' after the initial '0' is optimal because any higher bid would constitute a jump bid which was shown previously not to be profitable.
3. Conforming to β_∞^* and bidding β_i^f after a full '0'-cycle is optimal because any higher bid would constitute a jump bid and entering another '0'-cycle cannot be profitable because any cycle gives the same payoff as the equilibrium action.

As no individual cycle has any implication on payoffs or beliefs, strategies comprising more complicated deviations than the above single-stage deviations cannot have any implication either. Since β_∞^* prescribes the same stage actions as β^* in the finite game, the payoffs are the same. Reinterpreting '0' as any non-increasing bid does not change the above argument.

Existence of β_∞^* does not follow from our (backward) induction arguments for the finite game. It follows, however, from general arguments developed by Aumann and Maschler (1966); Mertens et al. (1994) and Simon et al. (1995). \square

References

- Aumann, R.J., Maschler, M., 1966. Game Theoretic Aspects of Gradual Disarmament, in Report of the U.S. Arms Control and Disarmament Agency, vol. ST-80, Chapter 5, pp. 1–55. Reprinted in Aumann, R.J., Maschler, M., 1994. Repeated Games with Incomplete Information. MIT Press, Cambridge, MA.
- Ausubel, L.M., Cramton, P., Deneckere, R.J., 2002. Bargaining with incomplete information. In: Aumann, R.J., Hart, S. (Eds.), Handbook of Game Theory, vol. 3. Elsevier Science Publishers B.V., Greenwich, Connecticut, pp. 1897–1945.
- Ausubel, L.M., Deneckere, R.J., 1989. A direct mechanism characterization of sequential bargaining with one-sided incomplete information. Journal of Economic Theory 48, 18–46.
- Brams, S., Taylor, A.D., 1996. Fair Division: From Cake-Cutting to Dispute Resolution. Cambridge University Press, Cambridge, UK.
- Brooks, R., Lando, C., Spier, K., 2009. Trigger happy or gun shy? Dissolving Common-Value Partnerships with Texas Shootouts, University of Alberta Working Paper, 2009-1.
- Bulow, J., Huang, M., Klemperer, P., 1999. Toeholds and takeovers. Journal of Political Economy 107, 427–454.
- Calcagno, R., Lovo, S., 2006. Bid-ask price competition with asymmetric information between market makers. Review of Economic Studies 73, 329–356.
- Compte, O., Jehiel, P., 2004. Gradualism in bargaining and contribution games. Review of Economic Studies 71, 975–1000.
- Cramton, P., 1984. Bargaining with incomplete information. Review of Economic Studies 51, 579–593.
- Deck, C.A., Farmer, A., Zeng, D.-Z., 2007. Amended final-offer arbitration outperforms final-offer arbitration. American Law and Economics Review 9, 384–407.
- Deneckere, R., Liang, M.-Y., 2006. Bargaining with interdependent values. Econometrica 74, 1309–1364.
- Engelbrecht-Wiggans, R., Milgrom, P., Weber, R., 1983. Competitive bidding and proprietary information. Journal of Mathematical Economics 11, 161–169.
- Evans, R.A., 1989. Sequential bargaining with correlated values. Review of Economic Studies 56, 499–510.
- Fudenberg, D., Levine, D.K., Tirole, J., 1987. Incomplete information bargaining with outside opportunities. Quarterly Journal of Economics 102, 37–57.
- Hörner, J., Jamson, J.S., 2008. Sequential common-value auctions with asymmetrically informed bidders. Review of Economic Studies 75, 475–498.
- Mertens, J.-F., Sorin, S., Zamir, S., 1994. Repeated Games. Parts A–C. CORE Discussion Paper, 9420-2.
- Milgrom, P., Weber, R., 1982. The value of information in a sealed-bid auction. Journal of Mathematical Economics 10, 105–114.
- Morgan, J., 2004. Dissolving a partnership (un)fairly. Economic Theory 23, 909–923.
- Ponsatí, C., 1997. Compromise vs capitulation in bargaining with incomplete information. Annales D'Économie et de Statistique 48, 191–210.
- Ponsatí, C., Sákovics, J., 1995. The war of attrition with incomplete information. Mathematical Social Sciences 29, 239–254.
- Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. Econometrica 50, 97–110.
- Sákovics, J., 1993. Delay in bargaining games with complete information. Journal of Economic Theory 59, 78–95.
- Schwartz, J.A., Wen, Q., 2007. Wage negotiation under good faith bargaining. International Game Theory Review 9, 551–564.
- Simon, R.S., Spież, S., Toruńczyk, H., 1995. The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk-Ulam type. Israel Journal of Mathematics 92, 1–21.
- Sobel, J., Takahashi, I., 1983. A multi-stage model of bargaining. Review of Economic Studies 50, 411–426.
- Stacchetti, E., Smith, L., 2003. Aspirational Bargaining, University of Michigan Working Paper.
- Ståhl, I., 1972. Bargaining Theory. Stockholm School of Economics, Stockholm.
- Tirole, J., 1982. On the possibility of speculation under rational expectations. Econometrica 50, 741–781.
- Winoto, P., McCalla, G., Vassileva, J., 2008. An extended alternating-offers bargaining protocol for automated negotiation in multi-agent systems. In: Meersman, R., Tari, Z. (Eds.), On the Move to Meaningful Internet Systems. Springer, Heidelberg.

²² We could alternatively view an infinite cycle as bargaining breakdown and assign some exogenous payoff of up to and including $u(\beta^*)$. This leaves our results unchanged and may fit infinite-game depreciation considerations better than the above.

²³ To be able to use the same time index and history as in the finite game, we reset the time index $t = t - 2$ in Definition 2 after each completed cycle. This drops the cycle from the history of the infinite game.