



Bidding games and efficient allocations[☆]

Reshef Meir^{b,*}, Gil Kalai^a, Moshe Tennenholtz^b

^a Hebrew University, Jerusalem, Israel

^b Technion—Israel Institute of Technology, Haifa, Israel

ARTICLE INFO

Article history:

Received 4 February 2016

Available online 24 August 2018

JEL classification:

C72

C78

Keywords:

Extensive form games

Richman games

Combinatorial games

Bargaining

ABSTRACT

Richman games are zero-sum games, where in each turn players bid in order to determine who will play next (Lazarus et al., 1999). We extend the theory to impartial general-sum two player games called *bidding games*, showing the existence of pure subgame-perfect equilibria (PSPE). In particular, we show that PSPEs form a semilattice, with a unique and natural *Bottom Equilibrium*.

Our main result shows that if only two actions available to the players in each node, then the Bottom Equilibrium has additional properties: (a) utilities are monotone in budget; (b) every outcome is Pareto-efficient; and (c) any Pareto-efficient outcome is attained for some budget.

In the context of combinatorial bargaining, we show that a player with a fraction of $X\%$ of the total budget prefers her allocation to $X\%$ of the possible allocations. In addition, we provide a polynomial-time algorithm to compute the Bottom Equilibrium of a binary bidding game.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Are game-theorists better at playing games? Game-theory and game-playing are often considered to have distinct purposes and different sets of required skills. Exceptional points of intersection are games that are both fun to play and possess interesting theoretical properties; among those, a particularly exquisite example is *Bidding Chess*.¹ Unlike standard Chess where turns are alternating among white and black, in the bidding version players each start with a fixed budget, and *bid* at the beginning of each turn. The higher bidder has the right to play next after paying her bid to the other player. It is thus possible for example that a player will move 3 pieces before the other player makes a single move. Yet one needs to carefully balance the bids for the different moves, or else she could lose too much of her budget, leaving her unable to take advantage of her position in the game. Indeed, even the bidding variation of the simple Tic-Tac-Toe game is non-trivial to play, and our personal experience shows that often a game where one player seems to lead gets turned over.²

These bidding variations of popular recreational games are special cases of a wider class of games known as *Richman Games*: A Richman game is very similar to a *combinatorial game* (Albert et al., 2007), in that it consists of a directed graph

[☆] A preliminary version of this paper appeared in the 16th ACM Conference on Economics and Computation (EC-2015). The first author is supported by the Israel Science Foundation (ISF) under Grant #773/16. The second author is supported by ERC grant 320924 and ISF grant 1612/17. The third author is supported by EU project 740435-MDDS. Part of the research was performed when all of the authors were affiliated with MSR Herzlia, Israel.

* Corresponding author.

E-mail addresses: reshefm@ie.technion.ac.il (R. Meir), kalai@math.huji.ac.il (G. Kalai), moshet@ie.technion.ac.il (M. Tennenholtz).

¹ For rules and historical background, see Bhat and Payne (2009).

² Many hours of play have been dedicated in the sake of promoting science.

of states (e.g. board states of the Chess game), where states without outgoing edges assign victory either to the black player or to the white player. In contrast to standard combinatorial games where turns are alternating, a bidding game starts with an *initial budget allocation*, and in each turn players *bid for the right to play next*. The higher bidder pays her bid to the lower bidder, and chooses a directed edge to follow, until either a white or a black terminal is reached. Richman games have been formally studied by Lazarus et al. (1999) (and later by Develin and Payne, 2010), whose main contribution was a surprising connection with random-turn games. See Section 8.1 for more details on Richman games and other variations of bidding games.

However, a zero-sum game with bidding phases is still a zero-sum game, and to the best of our knowledge, bidding games with general utility functions have not been previously studied. In this work, we extend the model to general-sum two player games by adding bidding phases to an arbitrary extensive form game. Whereas zero-sum games have a *value* that each player can guaranty, when extending the definition to general-sum games we can only talk about *equilibrium*. The standard solution concept studied in extensive-form game is *pure subgame-perfect equilibrium* (PSPE), meaning that each player is playing a best-response to the other player at any subtree of the game. Indeed, the purpose of this work is to study the existence and properties of PSPE in general-sum bidding games. We highlight that this is a natural extension of previous work on zero-sum games, since the value of a zero-sum game is a special case of a PSPE outcome.

Other than the pure theoretical interest in such results, many social and economic interactions can be formalized as two-player extensive form games. For example, we can consider two agents bargaining over a set of heterogeneous items. Then any sequential bargaining protocol is in fact an extensive form game (with complete information), where the utilities for players are the respective values of the subsets of items they get in the end. Indeed, several bargaining and arbitration protocols have been suggested in the literature, some of them of a sequential nature (see Section 8.1). We can thus use bidding phases as above to design a new bargaining mechanism called *sequential scrip bargaining* (SSB): The agents each start with some fixed budget (reflecting their relative power or entitlement), and use this budget to bid over items auctioned in some order. As mechanism design often deals with efficiency and fairness of proposed mechanisms, it is natural to ask e.g. whether the equilibrium outcomes of SSB correspond to efficient and/or fair allocations.

Paper structure and contribution. The definition of bidding games and some examples are provided in Section 2. In Section 3, we show that under a mild assumption on tie-breaking, every bidding game admits at least one PSPE. Further, the set of all PSPEs forms a meet-semilattice, whose lowest point is termed the *Bottom Equilibrium*.

Our main result is given in Section 4, and it states that in any bidding game over a *binary tree* (when at every stage of the game there are two possible actions), the Bottom Equilibrium can be computed by simulating an ascending auction. Further, it admits three desirable properties: (a) each player's utility is weakly monotone in her budget; (b) a Pareto-efficient outcome is reached for any initial budget; and (c) for any Pareto-efficient outcome there is an initial budget such that this outcome is attained. We complement this result by showing that none of the above properties is guaranteed in non-binary trees. A key lemma required for the main theorem shows that the Bottom Equilibrium has a simple structure and is not sensitive to small budget fluctuations.

In Section 6 we analyze the properties of the SSB mechanism informally described above, when the agents have arbitrary combinatorial valuations over items. As a direct corollary from our main theorem, we get that the mechanism guarantees a Pareto-efficient allocation for any valuation functions agents have over subsets of items. We further show that the allocation is fair in the sense that an agent with a fraction of $X\%$ of the total budget prefers her allocation to $X\%$ of the possible allocations, a criterion known as *minimal satisfaction test* (De Clippel et al., 2012).

Finally, we prove in Section 7 that any binary bidding game can be solved in time that is polynomial in the number of states, and show how such an algorithm can be used to efficiently find the Bottom Equilibrium in SSBs with additive valuations and other succinct representations of valuations.

2. Model

For an integer k , we denote the set $\{1, 2, \dots, k\}$ by $[k]$. Unless explicitly mentioned otherwise, we assume that $N = \{1, 2\}$, i.e. that there are only two competing agents (sometimes called *players* or *bidders*). We use the notation $-i$ (instead of $3 - i$) to denote the player that is not i . We name player 1 the *white* player, and player 2 the *black* player.

2.1. Bidding games

We first define the game structure that is common to combinatorial games, random-turn games and bidding games. The bidding process which is unique to bidding games is defined below.

Definition 1. An *game structure* is a tuple $G = \langle S, s_0, T, g_1, g_2, u_1, u_2 \rangle$, where:

- S is a set of game states;
- $s_0 \in S$ is the initial state;
- $T \subseteq S$ is a set of terminal states;
- $g_i : (S \setminus T) \rightarrow 2^S$ defines the optional moves of player i in state $s \in S \setminus T$;




Fig. 1. The game G_{maj} . White wins in white terminals, black wins in black. Double arrows mark the game play in the example.

- $u_i : T \rightarrow \mathbb{R}$ defines the utility for player i in each terminal $t \in T$.

We make the following assumptions about the transition functions g_i . First, unless specified otherwise, we assume that there are no cycles. Formally, there is no sequence s_1, s_2, \dots, s_k where $s_k = s_1$ and $s_{j+1} \in g_1(s_j) \cup g_2(s_j)$ for all $j < k$. Second, unless specified otherwise, we assume that in every non-terminal state both players can play. Formally, that $g_i(s) \neq \emptyset$ for all $i \in N, s \in S \setminus T$. The utility function u_i induces a complete preference order over terminals T (see also Sec. 3.2). We denote $t \succeq_i t'$ whenever $u_i(t) \geq u_i(t')$. Given a game structure G and $s \in S$, we denote by $G|_s$ the subgame of G rooted in s . The height of s is the maximal distance between s and a leaf $t \in T(G|_s)$. In particular, $\text{height}(G) = \text{height}(s_0)$.

The following two properties are structural properties of the underlying game tree/DAG:

Definition 2. A game structure G is *impartial*, if $g_1(s) = g_2(s)$ for all $s \in S$. That is, if the same set of moves is available to both players in every state.

Unless specified otherwise, all games in this paper are impartial. In impartial games we only need to specify one transition function $g(s)$.³

Definition 3. A game structure G is *binary*, if $|g(s)| \leq 2$ for all $s \in S \setminus T$.

Playing a bidding game. In order to complete the definition of a game, we also need a method to determine who plays at every turn. Traditionally, there is a turn function that assigns the current player for every state (e.g., alternating turns in Nim and other combinatorial games). Another way is to randomly select the current player in each turn, as in Peres et al. (2007). In this work we follow the bidding framework of Lazarus et al. (1999); Develin and Payne (2010), where each player has an initial budget B_i that is used for the bidding. The higher bidder pays her bid to the lower bidder, so that the total budget $B_1 + B_2 = 1$ is constant throughout the game.

We begin with an informal description of the strategies and an example.

In each turn (suppose at state s), each player submits a bid $b_i \leq B_i$, and a “next state” $s_i^* \in g(s)$, which is realized in case i wins the round. We break ties in favor of player 1. Then the winner (say, i^*) pays her bid to the loser $-i^*$ and the game proceeds from state $s_{i^*}^*$ with the modified budgets $B_{i^*} - b_{i^*}, B_{-i^*} + b_{i^*}$.

We consider both discrete and continuous bids. Finite resolution means that all bids and budgets are multiples of $\epsilon = \frac{1}{R'}$ for some integer R' . We say that ϵ is *high resolution* if $R' = R \cdot 2^k$ where $k = \text{height}(G)$ and $R \geq 4$ is an even integer. In particular it means that $\frac{1}{\epsilon}$ is at least exponential in the height of the tree. Throughout most of the paper, we will assume a high finite resolution. Later in Sections 4.3 and 5 we consider the consequences of small and infinite resolutions, respectively.

For a given ϵ , we denote by $\mathcal{B}_\epsilon = \{(B_1, B_2) \text{ s.t. } B_1 + B_2 = 1, \text{ and } \frac{B_1}{\epsilon}, \frac{B_2}{\epsilon} \in \mathbb{N}\}$ the set of possible budget partitions. We omit the subscript ϵ when clear from context. Since in either case the total budget is fixed, B_2 can always be inferred from B_1 and vice versa. We therefore identify \mathcal{B}_ϵ with $\{0, \epsilon, 2\epsilon, \dots, 1\}$ and use either B_1 or B_2 to denote a particular budget partition.

Definition 4. A (discrete) *bidding game* is a pair $\langle G, \epsilon \rangle$, where G is a game structure as per Definition 1, and $\epsilon = \frac{1}{R'}$ (for some integer R') is the resolution of budgets and bids.

Note that a bidding game $\langle G, \epsilon \rangle$ is in fact a collection of $\frac{1}{\epsilon} + 1$ different games, one for every initial budget partition $B_1 = 0, \epsilon, 2\epsilon, \dots, 1$.

Example—majority. Consider the following (zero-sum) game structure G_{maj} , depicted in Fig. 1. In this game there are at most three turns, and the winner is the player which plays at least twice. Formally, $S = (\{0, 1, 2\} \times \{0, 1, 2\}) \setminus \{(2, 2)\}$, and an element in S is simply the number of times that each player played; $s_0 = (0, 0)$; $T = \{(s_1, s_2) \in S : s_1 = 2 \vee s_2 = 2\}$; $g(s_i, s_{-i}) = \{(s_i + 1, s_{-i}), (s_i, s_{-i} + 1)\}$ for all $s \in S \setminus T$; and $u_i(t) = 1$ if $t_i = 2$ and 0 otherwise.

³ Chess for example is not an impartial game since different moves are available to black and white. Partial zero-sum games are discussed in Section 8.1, however for acyclic zero-sum games partiality does not matter (Develin and Payne, 2010).

Note that if G_{maj} is played with alternating turns and without budgets or bidding, the white player can always win. Consider $\langle G_{\text{maj}}, \epsilon \rangle$ with $\epsilon = \frac{1}{800}$, and $B_1 = B_2 = 0.5$, i.e. where each player has an initial budget of 0.5.

The following is a possible game play:

- In turn 1, white bids 0.2 and black bids 0.15. White gains the turn, plays the round, and the new budgets are 0.3 and 0.7. $s = (1, 0)$.
- In turn 2, white bids 0.14 and black bids 0.26. Black plays and pays 0.26. $s = (1, 1)$, and budgets are updated to (0.56, 0.44).
- In the last turn, both players exhaust their budgets, so the player with higher remaining budget (white with 0.56 vs. black with only 0.44) plays and wins the game. $s = (2, 1) \in T$.

Note that a higher budget makes it easier to win in this game. For example, if $B_i > 3(B_{-i} + \epsilon)$ for some player, then player i can always gain the first two rounds by bidding more than the other player's budget.

Strategies. We next define the allowed strategies for players in bidding games, and the main solution concept we apply.

Definition 5. A strategy profile in a bidding game $\langle G, \epsilon \rangle$ is a mapping

$$\gamma : S \times \mathcal{B} \rightarrow N \times (\mathcal{B})^2 \times \mathcal{B} \times S^2 \times T,$$

where $\mathcal{B} = \mathcal{B}_\epsilon$ and $\gamma(s, B_1)$ specifies the actions taken in state s under budget partition (B_1, B_2) (recall that $B_1 + B_2 = 1$ so the partition is determined by a single parameter).

More specifically, for all $s \in S, B_1 \in \mathcal{B}, \gamma(s, B_1) = (i^*, b_1, b_2, B_{i^*}^*, s_1^*, s_2^*, t)$, where: i^* is the winner; b_1 and b_2 are the bids of both players; $B_{i^*}^*$ is the remaining budget of the winner; s_i^* is the next state chosen by each i ; and $t \in T$ (also denoted by $\gamma(\cdot)_T$) is the outcome. While some of these elements can be inferred from others, the ability to reference them directly will become useful later on.

Definition 6. γ is a valid strategy profile if for all $s \in S \setminus T$ and $B_1 \in \mathcal{B}, \gamma(s, B_1) = (i^*, b_1, b_2, B_{i^*}^*, s_1^*, s_2^*, t)$ holds the following consistency constraints:

- $b_1 \leq B_1$ and $b_2 \leq B_2$.
- $i^* = 1$ if $b_1 \geq b_2$ and otherwise $i^* = 2$. That is, highest bidder wins the round with ties broken in favor of white.
- $B_{i^*}^* = B_{i^*} - b_{i^*}$. That is, the winner pays her bid to the other player.
- $s_1^*, s_2^* \in g(s)$.
- $t = \gamma(s_{i^*}^*, B_{i^*}^*)_T$. That is, players reach to the same terminal from state s and from the next state $s_{i^*}^*$.

For a terminal $t \in T$ we define $\gamma(t, B_1) = (1, 0, 0, B_1, t, t, t)$ in every valid strategy.

We denote by $\Gamma(G, \epsilon)$ the set of all valid profiles in the bidding game $\langle G, \epsilon \rangle$.

2.2. Subgame perfect equilibria

Definition 7. A strategy profile γ is a pure subgame perfect equilibrium (PSPE) in a bidding game $\langle G, \epsilon \rangle$, if it is a valid strategy profile, and each player plays a best-response strategy to the other player at any subtree, for any budget.

We denote by $\Gamma^*(G, \epsilon) \subseteq \Gamma(G, \epsilon)$ the set of all PSPEs in bidding game $\langle G, \epsilon \rangle$.

Since there are many possible responses of a player in a given strategy profile, we use the following lemma to specify all of them explicitly. The proof is immediate, by noticing that the lemma covers all possible selections of bids and states.

Lemma 1. A valid strategy profile γ is a PSPE in the bidding game $\langle G, \epsilon \rangle$ if and only if the following incentive constraints hold for all $s \in S \setminus T$ and $B_1 \in \mathcal{B}$ (where player i^* is the winner in $\gamma(s, B_1)$):

- For any $s'_{i^*} \in g(s)$ and any $b'_{i^*} \in \mathcal{B}$ s.t. $b'_{i^*} \leq B_{i^*}, b'_{i^*} > b_{-i^*}$ (weak inequality when $i^* = 1$), it holds that $u_{i^*}(t) \geq u_{i^*}(\gamma(s'_{i^*}, B_{i^*} - b'_{i^*})_T)$. That is, the winner i^* cannot gain by still winning the turn and use a different bid and/or action.
- For any $s'_{-i^*} \in g(s)$ and any $b'_{-i^*} \in \mathcal{B}$ s.t. $b'_{-i^*} \leq B_{-i^*}, b'_{-i^*} > b_{i^*}$ (weak inequality when $i^* = 2$), it holds that $u_{-i^*}(t) \geq u_{-i^*}(\gamma(s'_{-i^*}, B_{i^*} + b'_{-i^*})_T)$. That is, the loser $-i^*$ cannot gain by increasing his bid and win the turn.
- $u_{i^*}(t) \geq u_{i^*}(\gamma(s_{-i^*}^*, B_{i^*} + b_{-i^*})_T)$. That is, the winner i^* cannot gain by lowering her bid and losing the turn (does not apply if $i^* = 1$ and $b_2 = 0$). Note that the next state selected by the loser ($s_{-i^*}^*$) must be part of the description of the strategy.

We emphasize that a strategy profile determines the actions and outcome for any budget partition in any internal node, and in particular for any initial budget partition. For each $s \in S \setminus T$ denote $\mu_s(B_1) = \gamma(s, B_1)_T$. Therefore, every profile γ for $\langle G, \epsilon \rangle$ induces a mapping μ_γ from \mathcal{B}_ϵ to outcomes $T(G)$. That is, $\mu_\gamma(B_1) = \gamma(s_0, B_1)_T$. We say that the profiles γ, γ' are *outcome-equivalent* in $\langle G, \epsilon \rangle$ if $\gamma(s, B_1)_T = \gamma'(s, B_1)_T$ for all $s \in S, B_1 \in \mathcal{B}_\epsilon$.

3. Existence and structure of equilibria

We start by showing the following basic result. For $s \in S \setminus T$, denote by $\gamma|_s$ the profile γ on the subtree rooted in s . Clearly if γ is a PSPE then so is $\gamma|_s$.

Theorem 2. Any bidding game $\langle G, \epsilon \rangle$ has a PSPE, i.e. $\Gamma^*(G, \epsilon) \neq \emptyset$.

Further, if for all $s_j \in g(s_0)$ there is a PSPE $\gamma^{(j)}$ in $\langle G|_{s_j}, \epsilon \rangle$, then there is a PSPE γ' in $\langle G, \epsilon \rangle$ s.t. for all $s_j \in g(s_0), \gamma'|_{s_j} = \gamma^{(j)}$.

Intuitively, we construct such an equilibrium by traversing the game tree bottom-up, where in each node the players simulate an ascending auction (for every possible budget) to determine the equilibrium bids. We emphasize that in the actual bidding game bids are simultaneous, but we use an ascending auction in order to determine the equilibrium bids.

3.1. Existence of PSPE

Let a *first-price game* be a two-player one-shot game of the following form: each player submits a bid $b_i \in \mathcal{B}$, and the utilities of both players (denoted u_1 and u_2) are an arbitrary function of the higher bid.

Given a bid profile (b_1, b_2) , a *better-reply* of player i is a bid b'_i that results in a strictly better outcome, i.e., such that $u_i(\max\{b'_i, b_{-i}\}) > u_i(\max\{b_i, b_{-i}\})$.

Lemma 3. Consider any sequence of alternating better-replies in a first-price game. Then after the first time a bidder raises her bid, no bidder ever lowers her bid.

Proof. Recall that tie-breaking is in favor of bidder 1. Suppose that at some step bidder 1 increases her bid from (b_1, b_2) to b'_1 . Note that $b'_1 \geq b_2$, as otherwise $b_1 < b'_1 < b_2$ and $(b_1, b_2); (b'_1, b_2)$ lead to the same outcome (if the first increase is by bidder 2 then $b'_2 > b_1$). Any further step of bidder 1 must keep $b'_1 \geq b_2$, otherwise we just go back to the outcome of (b_1, b_2) (which is weakly worse for bidder 1).

Consider the next reply by bidder 2—he must also increase his bid to some $b'_2 > b'_1 \geq b_2$ in order to have any effect. Then bidder 1 must increase again and so on. Thus no bidder ever reduces her bid (after the first increase). \square

Let $\langle G, \epsilon \rangle$ be a bidding game, (s, B_1) be some state and budget, and $\boldsymbol{\gamma} = (\gamma_j)_{s_j \in g(s)}$ be some action profile in all the subgames $G|_{s_j}$. We observe that $\langle G, \epsilon \rangle, (s, B_1)$ and $\boldsymbol{\gamma}$ together induce a first-price game played in state (s, B_1) :

- In each s_j consider the mapping $\mu_j : \mathcal{B} \rightarrow T(G)$;
- For every $B_i^* \in \mathcal{B}_\epsilon$, let $s_i^*(B_i^*) \in \operatorname{argmax}_{s_j \in g(s)} u_i(\mu_j(B_i^*))$, i.e., the best selected state for i under remaining budget B_i^* ;
- For every $b_1 \geq b_2$, the outcome is $t^* = \mu_{s_1^*(B_1 - b_1)}(B_1 - b_1) \in T(G)$, and thus utilities $u_1(t^*), u_2(t^*)$ are only determined by the higher bid b_1 . Similarly when $b_2 > b_1$.

We denote this first-price game by $G(\epsilon, s, B_1, \boldsymbol{\gamma})$.

Proof of Theorem 2. Let $k = \operatorname{height}(G)$. We prove by induction on the height of the tree. The base case is trivial since in a tree of height 0, $S \setminus T$ is empty.

We first set the next states players select in s . Let $g(s) = \{s_1, \dots, s_q\}$. By the induction hypothesis, for each subgame $G|_{s_j}$ there is some PSPE $\gamma^{(j)}$. This defines a first-price game $G(\epsilon, s, B_1, \boldsymbol{\gamma})$ played in state (s, B_1) , where $\boldsymbol{\gamma} = (\gamma^{(j)})_{s_j \in g(s)}$.

Suppose we initialize both bids to 0. By Lemma 3, bids can only increase. Thus there is only a finite number of bid increments, from any initial state.⁴ We then set the bids in $\gamma'(s, B_1)$ to be the equilibrium bids of the first-price game, and action $s_i^* = \operatorname{argmax}_{s_j \in g(s)} u_i(\mu_{s_j}(B_i - b_i))$ for $i \in \{1, 2\}$. We complete the definition of γ' by setting $\gamma'|_{s_j} = \gamma^{(s_j)}$. \square

3.2. Genericity and uniqueness

Generic games are games where agents have strict preferences over all outcomes. In classical extensive-form games (without bidding), it is known that genericity entails the existence of a unique PSPE. However in our game there are simultaneous steps, and thus genericity may not be sufficient for uniqueness.

Our next example shows that in the general case, there might be multiple PSPEs that are not outcome-equivalent.

⁴ This uses the fact that there is only a finite number of possible bids. In the continuous case there are some caveats, see Section 5.




Fig. 2. The game G_{two} . Note that utilities are generic.

Proposition 4. *There is a generic bidding game with PSPEs that lead to different outcomes under the same budget.*

Proof. Consider the game G_{two} in Fig. 2, with some high budget resolution.

We first describe a single equilibrium $\gamma|_x$ in the subtree of node x . Any player with budget more than 0.5 (or ≥ 0.5 for white) after bidding at x clearly goes to y and gets 9, while the other player gets 1. The player i with the higher budget at x bids the difference $b_i = B_i(x) - 0.5$ if $i = 1$ and $b_i = B_i(x) - 0.5 - \epsilon$ if $i = 2$, and selects y . The other player $-i$ bids $\min\{B_{-i}, b_i\}$ if $i = 2$ and $\min\{B_{-i}, b_i + \epsilon\}$ if $i = 1$, and selects $(5, 5)$. Thus if the budget is $B_1(x) \in [0.25, 0.75 - \epsilon]$, the outcome is $(5, 5)$; if $B_1(x) < 0.25$ the outcome is $(1, 9)$; and if $B_1(x) \geq 0.75$ the outcome is $(9, 1)$. Note that this is a PSPE at x : if y is selected then the loser at x cannot raise, and if $(5, 5)$ is selected and the loser raises the bid to win, they will get to y with not enough budget to win again.

Next, we set the initial budget at s_0 to $(0.5, 0.5)$, and describe two distinct equilibrium bids at s_0 .

In the first PSPE γ^* , the bids at s_0 are $b_1 = b_2 = 0$, and both players select x . White wins and remains with budget 0.5. According to $\gamma|_x$, the game ends at terminal $(5, 5)$ (marked with $*$). This is clearly a PSPE.

In the second PSPE γ^{**} , the bids at s_0 are $b_1 = b_2 = 0.5$, and both players select $(2, 2)$ (marked with $**$). This is an equilibrium, since white remains with budget 0 after winning at s_0 , and selecting x means a utility of $1 < 2$ for her. Black selects $(2, 2)$ for the same reason, and cannot deviate by changing his bid. Finally, if white lowers her bid the outcome does not change. \square

Enforcing generic preferences. In the general case, it is possible that a player is indifferent between two outcomes $t, t' \in T$. In such cases, we will assume throughout the paper that the player has a strict preference towards the outcome that is also better for the other player.⁵ Thus every player has a strict preference order over all distinct outcomes. Recall that $t \succ_i t'$ means that player i strictly prefers t over t' . In the remainder of the paper, $t \succ_i t'$ if either $u_i(t) > u_i(t')$, or $u_i(t) = u_i(t')$ and $u_{-i}(t) > u_{-i}(t')$. We highlight that this assumption only increases the number of potential deviations from a given state, and thus cannot add new equilibria.

3.3. The PSPE semilattice

In general, $\Gamma(G, \epsilon)$ may contain multiple PSPEs. However, PSPEs can be partially ordered according to the bids of the players. We define a partial order over all valid profiles. Intuitively, profiles with lower winning bids will be lower in this order, however since each profile consists of many bids at different states, we need to construct the partial order carefully. In particular, profiles with smaller gaps between bids will be lower, and the order is first decided over nodes closer the bottom. Only if both profiles are equivalent in all nodes at some level or deeper, we check the bids at a higher level for a difference.

Given strategy profile γ and state $s \in S$, the function γ_s maps any budget $B_1 \in \mathcal{B}$ to bids b_1 and b_2 , where $\gamma_s(B_1) = (\gamma(s, B_1)_{b_1}, \gamma(s, B_1)_{b_2})$ (do not confuse with $\gamma|_s$, which is the restriction of γ on the subtree rooted in s).

Definition 8. The partial order π over $\Gamma(G, \epsilon)$ is defined as follows.

- For a given budget level B_1 , $\gamma'(s, B_1) \succ_\pi \gamma(s, B_1)$ if either:
 1. the winning bid in γ' is higher, i.e. $\max\{b'_1, b'_2\} > \max\{b_1, b_2\}$;
 2. the winning bids are equal, and the losing bid in γ' is lower;
 3. $b'_1 = b_2 > b_1 = b'_2$ (i.e. bids are the same), but white wins in $\gamma'(s, B_1)$.
 Thus for every (s, B_1) , π is a linear order over profiles, and in particular over PSPEs.
- For $s \in S \setminus T$, we define $\gamma'_s \succ_\pi \gamma_s$ if $\gamma'(s, B_1) \geq_\pi \gamma(s, B_1)$ for all $B_1 \in \mathcal{B}$, with at least one strict inequality. Thus for each node s , π is only a partial order over profiles (see Fig. 3 for an example).
- For $s \in S$, we define the partial order $\sigma(s)$ recursively (see Fig. 4 for an example):
 - $\gamma' =_{\sigma(t)} \gamma$ for all $t \in T$;
 - $\gamma' \succ_{\sigma(s)} \gamma$ if both:
 1. $\gamma' \geq_{\sigma(s')} \gamma$ for all $s' \in g(s)$, and
 2. either $\gamma' \succ_{\sigma(s')} \gamma$ for some $s' \in g(s)$, or $\gamma'_s \succ_\pi \gamma_s$.
- Finally, $\gamma' \succ_\pi \gamma$ if $\gamma' \succ_{\sigma(s_0)} \gamma$.

⁵ To see why this assumption is necessary, note that without it even the most simple game with a single decision node and two terminals may result in a Pareto-dominated outcome if the player with a higher budget is indifferent between the two terminals.

B_1	0	ϵ	2ϵ	...	$1 - \epsilon$	1
γ^1	–	(0, 4 ϵ)	(2 ϵ , 2 ϵ)	...	–	–
γ^2	–	(ϵ , 4 ϵ)	(ϵ , 2 ϵ)	...	–	–
γ^3	–	(0, 5 ϵ)	(0, 5 ϵ)	...	–	–
γ^4	–	(0, 4 ϵ)	(2 ϵ , 6 ϵ)	...	–	–

Fig. 3. The table shows the bids under four different profiles at state s , for various budget levels. Suppose that the bids for all $B_1 \notin \{\epsilon, 2\epsilon\}$ are the same under all profiles. From the table we can infer the following linear orders over profiles at given budget levels: $\gamma^2(s, \epsilon) <_{\pi} \gamma^1(s, \epsilon) =_{\pi} \gamma^4(s, \epsilon) <_{\pi} \gamma^3(s, \epsilon)$; and $\gamma^1(s, 2\epsilon) <_{\pi} \gamma^2(s, 2\epsilon) <_{\pi} \gamma^3(s, 2\epsilon) <_{\pi} \gamma^4(s, 2\epsilon)$. Next, we can infer a partial order over profiles at state s : $\gamma_s^1, \gamma_s^2 <_{\pi} \gamma_s^3, \gamma_s^4$. However the relation e.g. between γ_s^1 and γ_s^2 is not defined.

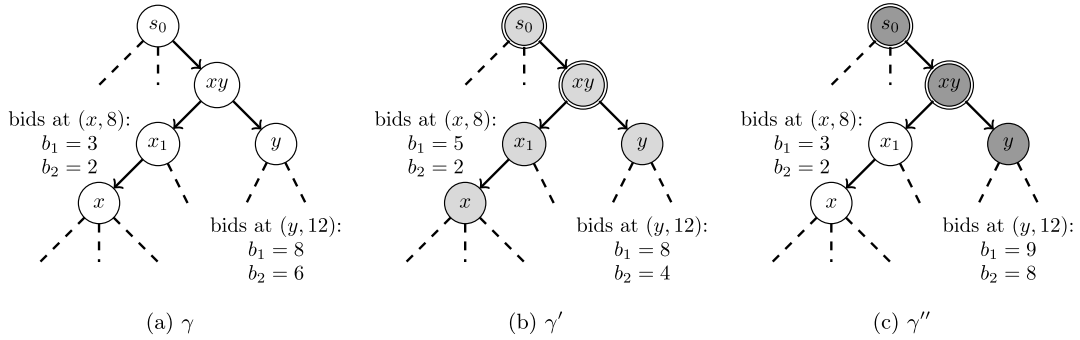


Fig. 4. The figure shows a part of three different profiles of the same game. The bids in nodes x and y under particular budgets are shown for each profile. For either x, y we assume that the bids are the same across all profiles for any other budget. We also assume that in all nodes that are not shown the bids are the same across all three profiles under any budget. For the other nodes $\{x_1, xy, s_0\}$ we do not assume anything about the bids, and they may differ between profiles. In each of $s \in \{x, y\}$ we get a different order π over the three profiles: $\gamma_x' >_{\pi} \gamma_x =_{\pi} \gamma_x''$ and $\gamma_y' >_{\pi} \gamma_y' >_{\pi} \gamma_y$. The shading shows the order $\sigma(s)$, where for any s , darker color means the respective equilibrium is higher in the order (unless comparing two nodes with double lines). For example: $\gamma' >_{\sigma(x_1)} \gamma =_{\sigma(x_1)} \gamma''$. To see why, note that by our construction $\gamma' \geq_{\sigma(s)} \gamma$ in all $s \in g(x_1)$, and the relation is strict for at least one child (x) . Thus we can ignore the bids in x_1 and still determine the relation $\sigma(x_1)$. Another example is $\gamma'' >_{\sigma(xy)} \gamma$. Observe however that γ', γ'' cannot be compared according to $\sigma(xy)$ and $\sigma(s_0)$ (nodes with double line). For xy this is because there are conflicting relations in the children of xy , and for s_0 this is since the relation in at least one child (xy) is undefined. We thus get that $\gamma' >_{\pi} \gamma$ and $\gamma'' >_{\pi} \gamma$, whereas γ', γ'' are incomparable.

Any partial order on $\Gamma(G, \epsilon)$ is in particular a partial order on $\Gamma^*(G, \epsilon)$. Note that π ignores the part of the strategy profile that specifies the selection of the next state. We remark that once bids are set, due to the genericity assumption there is only one way to specify state selection *in equilibrium*. Thus this would not pose a problem when comparing two PSPEs from $\Gamma^*(G, \epsilon)$.

Theorem 5. $(\Gamma^*(G, \epsilon), \pi)$ is a meet-semilattice.

Proof. Given two PSPEs γ, γ' for the same game $\langle G, \epsilon \rangle$, we define their meet $\hat{\gamma}$ as the highest PSPE $\hat{\gamma}$ such that $\hat{\gamma} \leq_{\pi} \gamma$ and $\hat{\gamma} \leq_{\pi} \gamma'$. The meet $\hat{\gamma} = \gamma \wedge \gamma'$ is unique, and is computed by Algorithm 1. Intuitively, the algorithm traverses both PSPEs from the bottom up, and takes the “lower” bids from each pair $\gamma(s, B_1), \gamma'(s, B_1)$ according to the order π .

Let $\hat{\gamma} = \hat{\gamma}(\gamma, \gamma')$ be the output of Algorithm 1 on input γ, γ' . To prove the theorem, we need to show that:

- the meet operation is well defined (i.e., that Algorithm 1 never fails);
- $\hat{\gamma}$ is a PSPE;
- $\hat{\gamma} = \gamma \wedge \gamma'$;

The algorithm never fails since by induction by the time we reach node s , we already have that all of $\hat{\gamma}|_{s'}$ for $s' \in g(s)$ are PSPEs.

To see why $\hat{\gamma}$ is a PSPE, we assume by induction that this holds for all nodes below s , and go over all cases of the construction in Algorithm 1. In Cases 1 and 3, $\hat{\gamma}|_s = \gamma|_s$ (or $\gamma'|_s$) and is thus a PSPE. In Case 2, for every $B_1 \in \mathcal{B}$ both $\gamma_s(B_1)$ and $\gamma'_s(B_1)$ are equilibrium bids, so it does not matter which one we select (note that for all $s' \in g(s)$, $\hat{\gamma}_{s'} = \gamma_{s'} = \gamma'_{s'}$). Since $\hat{\gamma}_s(B_1)$ are equilibrium bids for all $B_1 \in \mathcal{B}$, $\hat{\gamma}|_s$ is a PSPE. In Case 4 (when γ, γ' are incomparable by their children), by Proposition 2 it is possible to complete the bids for any budget level B_1 such that $\hat{\gamma}(s, B_1)$ is a PSPE. As for every state (s, B_1) the relation π is a linear order over all profiles, $\hat{\gamma}(s, B_1)$ is uniquely defined. In fact, it is composed of the highest winning bid $b_{i^*} \leq B_{i^*}$ and the lowest losing bid b_{-i^*} such that (b_{i^*}, b_{-i^*}) is an equilibrium of the first-price game $G(\epsilon, s, B_1, \gamma)$.

To show that $\hat{\gamma} = \gamma \wedge \gamma'$ we need to show that $\hat{\gamma}$ is the “largest” PSPE that is still weakly smaller than both γ, γ' . That is, (i) $\hat{\gamma} \leq_{\pi} \gamma, \gamma'$; and (ii) for any $\hat{\gamma}' \leq_{\pi} \gamma, \gamma'$, it also holds that $\hat{\gamma}' \leq_{\pi} \hat{\gamma}$.

Algorithm 1 COMPUTE-MEET(γ, γ').

```

Initialize all nodes as 'untouched';
for every node  $s$  in post-order do
  if all of  $g(s)$  are untouched, and  $\gamma_s = \gamma'_s$  // Case 1
  then
    set  $\hat{\gamma}_s = \gamma_s$ ;
  if all of  $g(s)$  are untouched, and  $\gamma_s \neq \gamma'_s$  // Case 2
  then
    for all  $B_1 \in \mathcal{B}$  do
      set  $\hat{\gamma}(s, B_1) = \min_{\pi} \{\gamma(s, B_1), \gamma'(s, B_1)\}$ ;
    change  $s$  to 'touched';
  else
    // some  $s_j \in g(s)$  are touched
    Let  $\boldsymbol{\gamma} = (\hat{\gamma}|_{s_j})_{s_j \in g(s)}$ ;
    if  $\hat{\gamma}|_{s_j} = \gamma|_{s_j}$  for all  $s_j \in g(s)$  // Case 3.1
    then
       $\hat{\gamma}_s = \gamma_s$ ;
    else if  $\hat{\gamma}|_{s_j} = \gamma'|_{s_j}$  for all  $s_j \in g(s)$  // Case 3.2
    then
       $\hat{\gamma}_s = \gamma'_s$ ;
    else if all of  $\hat{\gamma}|_{s_j}$  are PSPEs // Case 4
    then
      for all  $B_1 \in \mathcal{B}$  do
        set  $\hat{\gamma}(s, B_1)$  to the maximal equilibrium of the first-price game  $G(\epsilon, s, B_1, \boldsymbol{\gamma})$  according to  $\pi$ ;
    else
      FAIL
    change  $s$  to 'touched';
  
```

We prove (i) and (ii) together. Suppose that $\hat{\gamma} \leq_{\pi} \gamma, \gamma'$, then in particular $\hat{\gamma} \leq_{\sigma(s)} \gamma, \gamma'$ at all $s \in S$. Assume by induction that $\hat{\gamma} \leq_{\sigma(s_j)} \hat{\gamma} \leq_{\sigma(s_j)} \gamma, \gamma'$ for all $s_j \in g(s)$ (this vacuously holds at the leaves).

Thus, condition 1 in Definition 8 holds, and it remains to show that one of the two alternatives for condition 2 also holds, i.e., either the inequality in some children is strict, or there is (weak) inequality in s as well.

Specifically, if some of the inequalities $\hat{\gamma} \leq_{\sigma(s_j)} \gamma$ are strict (the first alternative), then $\hat{\gamma} <_{\sigma(s)} \gamma$ (and same for γ') and we are done with (i). Similarly, if some of the inequalities $\hat{\gamma} \leq_{\sigma(s_j)} \gamma'$ are strict, then $\hat{\gamma} <_{\sigma(s)} \gamma'$ and we are done with (ii).

To complete (i), note that if $\hat{\gamma}_s$ is constructed in Cases 1 or 2, then $\hat{\gamma}_s \leq_{\pi} \gamma_s, \gamma'_s$ and the second alternative of condition 2 holds. If Case 3.1 holds then $\hat{\gamma}|_{s_j}, \gamma'|_{s_j}$ must differ at some $s_j \in g(s)$, which means the first alternative holds for γ' and the second alternative (with equality) holds for γ (in Case 3.2 we get the opposite). If Case 4 holds, then $\hat{\gamma}$ differs from both γ, γ' in at least one child s_j (may be in different children). In either case we get that $\hat{\gamma} \leq_{\sigma(s)} \gamma, \gamma'$. When applied to the root s_0 we get $\hat{\gamma} \leq_{\pi} \gamma, \gamma'$.

To complete (ii), note that when there are no strict inequalities in the children of s , both of $\hat{\gamma}_s$ and $\hat{\gamma}'_s$ are equilibrium bids in the same first-price game $G(\epsilon, s, B_1, \boldsymbol{\gamma})$. We show that $\hat{\gamma}'_s \leq_{\pi} \hat{\gamma}_s$ in each of the four cases by which $\hat{\gamma}_s$ was constructed. In Cases 1 and 3 this is holds since $\hat{\gamma}_s$ equals one of γ_s, γ'_s , and $\hat{\gamma}'_s \leq_{\pi} \gamma_s, \gamma'_s$. In Cases 2 and 4, by construction for all B_1 , and since π is a linear order over all bidding profiles in (s, B_1) , we have $\hat{\gamma}(s, B_1) \leq_{\pi} \hat{\gamma}'(s, B_1)$, and thus $\hat{\gamma}'_s \leq_{\pi} \hat{\gamma}_s$ as well. This means that the second alternative for condition 2 holds (with weak inequality), and thus $\hat{\gamma} \leq_{\sigma(s)} \hat{\gamma}'$. In particular $\hat{\gamma} \leq_{\pi} \hat{\gamma}'$ as required. \square

We demonstrate the computation of the meet through an example. Suppose that we compute the meet of γ^3, γ^4 from Fig. 3 at node s . Assuming that all nodes below s are untouched, we are in Case 2. Then at each budget level we take the pair of bids that is lower, getting $\hat{\gamma}(s, \epsilon) = \gamma^4(s, \epsilon) = (0, 4\epsilon)$, $\hat{\gamma}(s, 2\epsilon) = \gamma^3(s, 2\epsilon) = (0, 5\epsilon)$.

Next, suppose that the three profiles in Fig. 4 are PSPEs, and consider the meet $\hat{\gamma} = \gamma' \wedge \gamma''$. We have that $\hat{\gamma}|_x = \gamma''|_x$ due to Case 2. Similarly, $\hat{\gamma}|_y = \gamma'|_y$. We then complete $\hat{\gamma}$ in the other nodes x_1, xy, s_0 as specified. Note that as $\gamma', \gamma'' >_{\pi} \gamma$, their meet $\hat{\gamma} \geq_{\pi} \gamma$ as well (in this example it holds with a strict relation).

Since $\Gamma(G, \epsilon)$ is finite, $\Gamma^*(G, \epsilon)$ is also finite and thus by Theorem 5, π has a unique minimum, which we refer to as the *Bottom Equilibrium*. We denote the Bottom Equilibrium of the bidding game (G, ϵ) by $BE(G, \epsilon)$.

4. Efficiency, monotonicity, and binary trees

Having established the existence of PSPEs, and in particular the Bottom Equilibrium, we next turn to study the properties of this outcome, how it depends on the budget, and whether it is desired from the perspective of the players. This is the primary section of this work, where we lay out our main positive result.

Definition 9. A profile $\gamma \in \Gamma(G, \epsilon)$ is *monotone*, if $\mu_{\gamma}(B'_1) \geq_1 \mu_{\gamma}(B_1)$ whenever $B'_1 > B_1$, and likewise for player 2.




Fig. 5. The game G_{bad} .

Let $T_P(G) \subseteq T(G)$ be the set of Pareto-efficient outcomes in G . That is, $t \in T_P(G)$ if there is no $t' \in T(G)$ s.t. $t' \succ_1 t$, $t' \succ_2 t$.⁶

Definition 10. A profile $\gamma \in \Gamma(G, \epsilon)$ is *Pareto-optimal*, if $\mu_\gamma(B_1) \in T_P(G)$ for all $B_1 \in \mathcal{B}$.

Definition 11. A profile $\gamma \in \Gamma(G, \epsilon)$ is *Pareto-surjective*, if for all $t \in T_P(G)$, there is some budget $B_1 \in \mathcal{B}$, s.t. $\mu_\gamma(B_1) = t$.

Intuitively, our main result is that *binary* bidding games always have a highly desirable PSPE (the Bottom Equilibrium), which has all the above properties. We will first show that if the game is not binary, none of these properties is guaranteed. Violating Pareto-surjective is trivial: just consider a root with three children, whose utilities are $(1, 3)$, $(2, 2)$, and $(3, 1)$. There is no equilibrium where $(2, 2)$ is selected.

Proposition 6. *There is a (non-binary) bidding game, where every PSPE violates monotonicity and Pareto-optimality.*

Proof. Consider G_{bad} (Fig. 5) with some high resolution, and initial budget $B_2 = 0$. Clearly white can play at will, and thus the only PSPE leads to the outcome $(10, 7)$ (marked). Denote by $B_1(x)$ the budget at node x . Note that if players reach x , then the outcome will be $(10, 7)$ if $B_1(x) \geq 0.75$ and otherwise $(0, 9)$. Now suppose that initial budgets are $B_1 = 0.8, B_2 = 0.2$. If black takes the first round, he will choose $(1, 8)$, since by choosing x we will have $B_1(x) \geq B_1 = 0.8 > 0.75$. In order to get to $(10, 7)$, white has to get to x with at least 0.75 , that is to bid $b_1 \leq 0.05$. However if $b_1 < 0.2 = B_2$, then black will overbid and take the turn. Thus in every PSPE white must bid at least $b_1 = 0.2$ in s_0 , and play $(2, 1)$. This contradicts both monotonicity and Pareto: the utility of black dropped from 7 to 1 when his budget increased from 0 to 0.2 , and $(2, 1)$ is Pareto-dominated by $(10, 7)$. \square

Budget resolution. Without sufficient budget (high resolution), the bidding process may not be powerful enough to get the efficient outcome (see Section 4.3). However we argue that with high resolution, we can think of the budget interval as continuous, and this will be useful (in fact crucial) in our proofs later on. This idea can be demonstrated by looking at small game trees.

Consider a game tree with a single decision node s (height 1). It is clear that regardless of the budget resolution ϵ , it only matters whether $B_1 \geq 0.5$, since this will determine who can win the turn. Similarly, if we consider a game of height 2, then there are four relevant “budget intervals,” namely $[0, 0.25)$, $[0.25, 0.5)$, $[0.5, 0.75)$ and $[0.75, 1]$: in the first one the black player can win both turns, in the second black can choose at least one victory, and so on. We next provide a general definition of this property.

Recall that by our high resolution assumption, $\epsilon = \frac{1}{R \cdot 2^k}$ for some $R \geq 4$. For each $j < 2^k$, we define $B_1^j(k) = \frac{j}{2^k}$ and

$$\mathcal{B}^j(k) = \{B_1 \in \mathcal{B}_\epsilon : B_1^j \leq B_1 < B_1^{j+1}\} = \{\frac{j}{2^k} + r\epsilon : r = 0, 1, \dots, R - 1\}.$$

That is, \mathcal{B}^j is the j 'th budget subinterval, which is “closed” from below (includes $\frac{j}{2^k}$) and “open” from above.⁷ We omit the parameter k when clear from the context.

Definition 12. A profile $\gamma \in \Gamma(G, \epsilon)$ is an *interval profile* if for all $s \in S \setminus T$, all $j < 2^k$ and all $\hat{B}_1 \in \mathcal{B}^j$, $\mu_\gamma(s, \hat{B}_1) = \mu_\gamma(s, B_1^j)$.

When an equilibrium is not monotone, the equilibrium bids at both ends of the interval $\mathcal{B}^j(k)$ may reach different intervals at the $k - 1$ level. Thus as we increase the resolution ϵ , the structure of equilibria may become more and more fragmented. For both practical and technical reasons we would prefer equilibria that are interval profiles. Unfortunately, in the general case the Bottom Equilibrium may not be an interval profile due to lack of monotonicity. However for *binary trees*, we will see that the Bottom Equilibrium is both monotone and an interval profile.

⁶ Recall that by our assumption \succ_1 and \succ_2 are strict orders.

⁷ The budget $B_1 = 1$ is included in the last interval \mathcal{B}^{2^k-1} .

Algorithm 2 FIND-PSPE(G, R).

```

Set  $\epsilon \leftarrow \frac{1}{k} 2^{-\text{height}(G)}$ ; // Global variable
 $\mathcal{B} \leftarrow \{c \cdot \epsilon : c \in \{0, 1, \dots, \lceil 1/\epsilon \rceil\}\}$ ;
Initialize tables  $A_1, A_2, T^*, I$  of size  $|S| \times |\mathcal{B}|$ ;
In every leaf  $t$ , assign  $T^*(t, B_1) = t$  for all  $B_1 \in \mathcal{B}$ ;
for every node  $s \in S \setminus T$  in post-order do
  for every budget partition  $B_1 \in \mathcal{B}$ : do
     $b_1 \leftarrow 0; b_2 \leftarrow 0$ ; // Initialize both bids
     $i^* \leftarrow 1$ ; // Initially player 1 wins
     $s_1 \leftarrow \text{argmax}_{s' \in g(s)} u_1(T^*(s', B_1))$ ;
     $s_2 \leftarrow \text{argmax}_{s' \in g(s)} u_2(T^*(s', B_1))$ ;
     $t^* \leftarrow T^*(s_1, B_1)$ ;
    if  $s_1 = s_2$  then
      // No conflict. White wins with bid 0 and plays  $s_1$ 
      Break;
  repeat
    // Try deviation by loser  $-i^*$ :
     $b'_{-i^*} \leftarrow b_{-i^*} + \epsilon$ ;
     $t' \leftarrow T^*(s_{-i^*}, B_{-i^*} - b'_{-i^*})$ ;
    if  $b'_{-i^*} > B_{-i^*}$  or  $t' \leq_{-i^*} t^*$  then
      Break; // Deviation impossible or not successful
    else
       $i^* \leftarrow -i^*$ ; // Loser becomes winner
       $b_{i^*} \leftarrow b'_{i^*}$ ; // Deviation becomes the new bid
       $t^* \leftarrow t'$ ;
    // Write down the equilibrium strategies in  $\gamma(s, B_1)$ :
     $I(s, B_1) \leftarrow i^*$ ; // winner
     $A_i(s, B_1) \leftarrow (b_i, s_i)$  for  $i \in \{1, 2\}$ ; // actions of players (bid+selection)
     $T^*(s, B_1) \leftarrow t^*$ ; // terminal
  until;
return  $\gamma \leftarrow (I, A_1, A_2, T^*)$ ;

```

4.1. Binary bidding games

The remainder of this section is dedicated to the proof of our main theorem:

Theorem 7. Let G be a binary bidding game with high resolution ϵ , and let $\gamma = BE(G, \epsilon)$. Then:

- (A) At any point $\gamma(s, B_1)$, either $b_1 = b_2$ (if white wins) or $b_2 = b_1 + \epsilon$ (if black wins);
- (B) γ is an interval profile;
- (C) γ is Pareto-surjective;
- (D) γ is Pareto-optimal;
- (E) γ is monotone.

Recall that we denote by $\gamma|_s$ the projection of γ on a subgame of G that is rooted in s . Similarly, $\mu|_s = \mu_{\gamma|_s}$. We will sometimes use the budget of player 2 as the input for μ , rather than the budget of player 1. This should be clear from the notation. That is, $\mu(B_1)$ and $\mu(B_2)$ always refer to the same outcome in $T(G)$.

Denote the two children of s_0 by s_l, s_r . As the proof of all parts is by induction on the height of G , we assume that $\gamma = BE(G, \epsilon)$ has been computed on both subtrees $G_l = G|_{s_l}$ and $G_r = G|_{s_r}$, and denote $\gamma^{(l)} = \gamma|_{s_l}, \gamma^{(r)} = \gamma|_{s_r}$ and similarly for μ_l, μ_r . By construction, we have that $\gamma^{(l)} = BE(G_l, \epsilon), \gamma^{(r)} = BE(G_r, \epsilon)$. Thus by the induction hypothesis both of $\gamma^{(l)}, \gamma^{(r)}$ have all properties (A)–(E).

Proof of Theorem 7(A) (near-tie). By induction in each subtree we have that $\gamma^{(l)}, \gamma^{(r)}$ are monotone. Suppose (b_1, b_2) are the Bottom Equilibrium bids at some state $\gamma_s(B_1)$, and assume towards a contradiction that $b_1 > b_2$. Consider the bid profile $\gamma'_s(B_1) = (b_1, b_2 + \epsilon)$. Either $(b_1, b_2 + \epsilon)$ are equilibrium bids, or there is a deviation. However the only deviation that is not also available in (b_1, b_2) is by the white player. W.l.o.g. white can deviate by bidding $b'_1 = (b_2 + \epsilon) - \epsilon = b_2$, since as long as she drops below $b_2 + \epsilon$ the outcome is the same (this is a first-price game). Now either $(b'_1, b_2 + \epsilon)$ are equilibrium bids, or black has a deviation. By monotonicity of $\gamma^{(l)}$ and $\gamma^{(r)}$, black cannot gain by bidding $b'_2 > b_2 + \epsilon$, since $\mu_l(B_2 - b'_2) \leq \mu_l(B_2 - (b_2 + \epsilon))$ and likewise for μ_r . Thus the only deviation is to decrease his bid and drop the round, w.l.o.g. to $b'_2 = b'_1$. Again either (b'_1, b'_2) is an equilibrium, or white drops to $b''_1 = b'_2 - \epsilon$ and so on until no player has a deviation at some (b^*_1, b^*_2) . Thus we get a PSPE γ^* where $\gamma^*_s(B_1) = (b^*_1, b^*_2)$. If there were at least two deviations, then $b^*_1 < b_1$. If there are no deviations, then $b^*_1 = b_1$ and $b^*_{-i^*} > b_2$. If there is exactly one deviation then $b^*_{i^*} = b^*_2 = b_1$ and $b^*_1 = b_2$. In either case, we get that $(b^*_1, b^*_2) <_{\pi} (b_1, b_2)$, and thus $\gamma^* <_{\pi} \gamma$, which is a contradiction to γ being the Bottom

Equilibrium. The proof for $b_2 > b_1 + \epsilon$ is the same except that there are only two cases for (b_1^*, b_2^*) (no deviations, or at least one). \square

Lemma 8. Algorithm 2 computes γ .

Proof. Recall that $k = \text{height}(G)$. Note that the algorithm effectively goes over all states $s \in S \setminus T$ from the bottom up, and for any possible budget allocation, runs an auction where players iteratively modify their bids by the smallest possible increment ϵ , until none of them can gain by further modifying her bid. Note that the players never change their selected state when increasing their bid. Indeed, if player 2 has an incentive to raise at the initial bids of $(0, 0)$ (and select child s_2) only if he strictly prefers $t' = \mu_2(B_2 - \epsilon)$ over $t^* = \mu_1(B_2)$, and $\mu_1(B_2) \geq \mu_1(B_2 - \epsilon)$ by monotonicity of $\gamma^{(2)}$. Similarly, now player 1 has an incentive to raise only if she prefers $t' = \mu_1(B_1 - \epsilon)$ over $t^* = \mu_2(B_2 - \epsilon)$, which means that $t' >_1 \mu_2(B_2 - \epsilon) = \mu_2(B_1 + \epsilon) \geq_1 \mu_2(B_1 - \epsilon)$ by monotonicity of $\gamma^{(1)}$. This continues with every raise of bid, so player 1 always prefers s_1 and player 2 always prefers s_2 .

By property (A), the winning and losing bids in γ are nearly tied or tied. The algorithm goes over all such pairs of bids in increasing order and thus must reach (b_1, b_2) as in $\gamma(s, B_1)$ unless it stops at some lower bids. Suppose there are lower bids $\gamma'_s(B_1) = (b'_1, b'_2)$ where the loser $-i^*$ cannot deviate by ϵ and the algorithm stops. Then by monotonicity of $\gamma^{(1)}, \gamma^{(2)}$, neither i^* nor $-i^*$ can deviate by further increasing their bid. Further, by Lemma 3, bids can only increase so there is no deviation by decreasing a bid either, meaning that γ' is also a PSPE which is lower than γ . A contradiction. \square

Proof of Theorem 7(B) (Intervals). Let s be a node at height k and $j < 2^k$. Consider $\gamma_s(B_1^j) = (b_1, b_2)$ and $\hat{B}_1 \in \mathcal{B}^j(k)$. Note that the size of each level k interval is R , and the size of level $k - 1$ intervals is $2R$. Thus $\hat{B}_1 = B_1^j + r\epsilon$ for some $r < R$. Suppose first that $b_1 \geq b_2$, i.e. that white wins. By property (A), $b_1 = b_2$. Let $j', j'' < 2^{k-1}$ be the level $k - 1$ budget intervals reached when white wins and if white drops, respectively. Formally, $B_1^j - b_1 \in \mathcal{B}^{j'}(k - 1)$ and $B_1^j + b_1 = B_1^j + b_2 \in \mathcal{B}^{j''}(k - 1)$. We argue that there are 3 cases:

Case 0 $b_1 = b_2 = 0$ (no conflict). Both agents prefer the same node $s' \in g(s)$ at the current budget B_1 . So all of interval $\mathcal{B}^{j'}(k - 1)$ is (weakly) preferred to all of interval $\mathcal{B}^{j''}(k - 1)$.

For any $\hat{B}_1 \in \mathcal{B}^j(k)$ with $\hat{b}_1 = \hat{b}_2 = 0$ this still holds, as $\mathcal{B}^j(k)$ is contained in both $\mathcal{B}^{j'}(k - 1)$ and $\mathcal{B}^{j''}(k - 1)$ (which are actually the same interval but in different nodes). Thus $\gamma(s, \hat{B}_1) = (0, 0)$ as well.

Case 1L $B_1^j + b_1$ is the maximal point in $\mathcal{B}^{j''}(k - 1)$, i.e. $B_1^j + b_1 = B_1^{j''+1} - \epsilon = B_1^{j''} + (2R - 1)\epsilon$. This means that $B_1 - b_1 = B_1^j + \epsilon$. By the progress of Algorithm 2 we can conclude due to Lemma 8 that white prefers outcome j' (and winning the turn) over outcome j'' (and losing the turn), otherwise she would not have made the last increase from b_1 . Similarly, black prefers j'' over j' , but prefers j' over $j'' + 1$, as otherwise he would deviate to $b_2 + \epsilon$.

Let $\hat{B}_1 = B_1^j + r\epsilon$ (where $r < R$). Denote $\hat{b}_1 = b_1 - r\epsilon, \hat{b}_2 = \hat{b}_1 = b_2 - r\epsilon$. Equilibrium bids are presented graphically in Fig. 6. We argue that

- $\hat{B}_1 + \hat{b}_2 \in \mathcal{B}^{j''}(k - 1)$;
- $\hat{B}_1 - \hat{b}_1 \in \mathcal{B}^{j'}(k - 1)$; and
- $\gamma_s(\hat{B}_1) = (\hat{b}_1, \hat{b}_2)$

The first claim is obvious, since $\hat{B}_1 + \hat{b}_2 = B_1^j + b_2 \in \mathcal{B}^{j''}(k - 1)$. To see why the second claim holds, note that since $r < R$,

$$\hat{B}_1 - \hat{b}_1 = B_1^j + r\epsilon - (b_1 - r\epsilon) = B_1^j - b_1 + 2r\epsilon < B_1^j - b_1 + (2R - 1)\epsilon,$$

which is still in interval $\mathcal{B}^{j'}(k - 1)$.

For the third claim, suppose the equilibrium bids are lower than \hat{b}_1 . If white wins then the outcome is either in interval j' , or in a higher interval which is even worse for black. This means black would have a deviation to \hat{b}_2 . Similarly, if black wins with $b'_2 \leq \hat{b}_2$ then the outcome is wither in interval j'' or in a lower interval which is worse for white, and thus white would have a deviation to \hat{b}_1 . By Lemma 8, it is left to show that Algorithm 2 stops at (\hat{b}_1, \hat{b}_2) . Indeed, the next deviation by black would reach interval $j'' + 1$, and we know black prefers the current outcome j' . Intuitively what we showed is that $B_1^j + b_2 = B_1^{j''} + (2R - 1)\epsilon$ is the “critical point” where the auction stops, as long as $B_1 \in \mathcal{B}^j(k)$ (see rightmost dotted line in Fig. 6).

Case 1W $B_1^j + b_1$ is not the maximal point in its interval $\mathcal{B}^{j''}(k - 1)$. Then we argue that $B_1^j - b_1$ is the highest point of $\mathcal{B}^{j'}(k - 1)$, i.e. $B_1^j(k) - b_1 = B_1^{j'}(k - 1) + (2R - 1)\epsilon$. This is since both agents prefer interval j' (where white wins) over j'' (where black wins), as otherwise black would raise by ϵ . Then since $b_1 = b_2$, we get that $B_1^j + b_2 = B_1^{j''}(k - 1) + \epsilon$, i.e. that the losing bid reaches the second-lowest point on its interval.




Fig. 6. The middle row shows intervals j' and j'' of level $k - 1$ (each of size $2R \cdot \epsilon$). The top row shows an interval j of level k (of size $R \cdot \epsilon$) and the equilibrium bids in Case 1L at (s, B_1^j) . The bottom row shows the equilibrium bids at $\hat{B}_1 = B_1^j + r\epsilon$. Note that $b_1 = b_2$ and $\hat{b}_1 = \hat{b}_2$.




Fig. 7. Equilibrium bids at (s, B_1^j) and (s, \hat{B}_1) in Case 1W.

Now, we continue in a similar way to Case 1L, except that here critical point is $B_1^j - b_1 = B_1^{j'}(k - 1) + (2R - 1)\epsilon$, i.e. determined by the winner $i = 1$ (leftmost dotted line in Fig. 7). Thus define $\hat{b}_1 = \hat{b}_2 = b_1 + r\epsilon$, and we will show that

- $\hat{B}_1 - \hat{b}_1 \in \mathcal{B}^{j'}(k - 1)$;
- $\hat{B}_1 + \hat{b}_2 \in \mathcal{B}^{j''}(k - 1)$; and
- $\gamma_s(\hat{B}_1) = (\hat{b}_1, \hat{b}_2)$

The first claim follows since $\hat{B}_1 - \hat{b}_1 = B_1^j - b_1 \in \mathcal{B}^{j'}(k - 1)$, and the second since

$$\hat{B}_1 + \hat{b}_2 = B_1^j(k) + r\epsilon + (b_1 + r\epsilon) = B_1^j(k) + b_1 + 2r\epsilon = B_1^{j''}(k - 1) + \epsilon + 2r\epsilon \leq B_1^{j''}(k - 1) + (2R - 1)\epsilon.$$

As in Case 1L, the algorithm would not stop at lower bids than \hat{b}_1 . Any increase by black would reach interval j' or higher (worse for black), but we know black prefers j' over j'' so would not deviate.

In either case, we get by the induction hypothesis that

$$\mu_\gamma(s, \hat{B}_1) = \mu_\gamma(s_i, \hat{B}_1 - \hat{b}_1) = \mu_\gamma(s_i, B_1^j - b_1) = \mu_\gamma(s, B_1^j).$$

If $b_2 > b_1$, then Case 0 is impossible, and we get two cases similar to the above analysis (see Fig. 8): Case 2L where the critical point is $B_1^j - b_1 = B_1^{j'}(k - 1)$ and $B_1^j + b_2 = B_1^{j''}(k - 1) + \epsilon$; and case 2W where the critical point is $B_1^j + b_2 = B_1^{j''}(k - 1)$, and $B_1^j - b_1 = B_1^{j'}(k - 1) + \epsilon$. In either case we get that there is some critical player \hat{i} , for which $\hat{B}_i - \hat{b}_i = B_1^j - b_i$, and that $\hat{B}_{-\hat{i}} - \hat{b}_{-\hat{i}}, B_1^j - b_{-\hat{i}}$ reach the same level $k - 1$ budget interval (details omitted but very similar to cases 1L,1W). Then again, by the induction hypothesis

$$\mu_\gamma(s, \hat{B}_1) = \mu_\gamma(s_i, \hat{B}_2 - \hat{b}_2) = \mu_\gamma(s_i, B_2^j - b_2) = \mu_\gamma(s, B_1^j). \quad \square$$

Note that in particular we proved the following corollary (phrased as a lemma since will be used later):




Fig. 8. Equilibrium bids at (s, B_1^j) in Case 2L and 2W. Note that in both cases $b_2 = b_1 + \epsilon$. The critical point in each case is marked with a double dotted line.

Lemma 9. For any $s \in S \setminus T$ and $j < 2^k$, one of the following five cases applies for the Bottom Equilibrium $\gamma_s(B_1^j) = (b_1, b_2)$.

Case	Winner	Critical	Bid of player 1	Bid of player 2
0	$i^* = 1$	–	0	0
1L	$i^* = 1$	$\hat{i} = 2$	$B_1^j(k) - b_1 = B_1^{j'}(k - 1) + \epsilon$	$B_1^j(k) + b_2 = B_1^{j''} + (2R - 1)\epsilon$
1W	$i^* = 1$	$\hat{i} = 1$	$B_1^j(k) - b_1 = B_1^{j'}(k - 1) + (2R - 1)\epsilon$	$B_1^j(k) + b_2 = B_1^{j''} + \epsilon$
2L	$i^* = 2$	$\hat{i} = 1$	$B_1^j(k) - b_1 = B_1^{j'}(k - 1)$	$B_1^j(k) + b_2 = B_1^{j''}(k - 1) + \epsilon$
2W	$i^* = 2$	$\hat{i} = 2$	$B_1^j(k) - b_1 = B_1^{j'}(k - 1) + \epsilon$	$B_1^j(k) + b_2 = B_1^{j''}(k - 1)$

Proof of Theorem 7(C) (Pareto-surjective). We prove by induction on the height of the tree G . For height 0 it is obvious. For height 1 either one leaf weakly Pareto-dominates the other, in which case this leaf is reached in μ regardless of budgets, or each player has a favorite leaf in which case both leaves are Pareto-efficient.

Let $t^* \in T_P(G)$, we need to show that there is some initial budget $B \in \mathcal{B}$ s.t. that $\mu(B) = t^*$. W.l.o.g. $t^* \in T_P(G_1)$.

By induction on property (C), let B_1^* s.t. $\mu_l(B_1^*) = t^*$ (must exist), and let $t_r^* = \mu_r(B_1^*)$. As γ is an interval profile, w.l.o.g. $B_1^* = B_1^j(k) + 2\epsilon$ for some $j < 2^k$ (i.e. it an even multiple of ϵ , and is not on the edge of its respective interval). This is possible by property (B) and since $R \geq 4$. Thus by slightly increasing or decreasing the budget, the outcome will not change: $\mu_l(B_1^* - 2\epsilon) = \mu_l(B_1^* + \epsilon) = t^*$ [also equal to $\mu_l(B_2^* - \epsilon)$]; and likewise for μ_r .

We know that t_r^* does not strictly Pareto dominate t^* (as $t^* \in T_P(G)$), thus either $t_r^* <_1 t^*$ or $t_r^* <_2 t^*$.

Case I: Both players prefer t^* over t_r^* at budget B_1^* . We claim that the equilibrium bids are $b_1 = b_2 = 0$.

Suppose that black raises to $b' > b_2 = 0$, takes the round and goes right. then

$$\mu_r(B_2^* - b') \leq \mu_r(B_2^*) = t_r^* <_2 t^*,$$

and thus black does not gain. Therefore $b_1 = b_2 = 0$, and $\mu(B_1^*) = \mu_l(B_1^*) = t^*$.

Case II: Suppose that $t_r^* >_2 t^*$, $t_r^* <_1 t^*$, i.e. that black strictly prefers child s_r at budget B_2^* , whereas white strictly prefers s_l . Let $\hat{B}_2 \leq B_2^*$ be the lowest budget s.t. for every $\beta_2 \in [\hat{B}_2, B_2^*]$, $\mu_r(\beta_2) >_2 t^*$ (it is possible that $\hat{B}_2 = 0$). Set $\hat{t}_r = \mu_r(\hat{B}_2)$. By Pareto-optimality of t^* it follows that for every $\beta_2 \in [\hat{B}_2, B_2^*]$, white prefers s_l (since $\mu_r(\beta_1) <_1 t^* = \mu_l(B_1^*) \leq \mu_l(\beta_1)$).

We set $B_2 = \frac{B_2^* + \hat{B}_2}{2}$, $B_1 = 1 - B_2$, and argue that $\mu(B_1) = t^*$. Let $b^* = \frac{B_2^* - \hat{B}_2}{2}$. We will show that the equilibrium in s_0 is $b_1 = b_2 = b^*$. Note that \hat{B}_2 is the lower bound of a budget interval and thus also an even multiple of ϵ , so B_2 and b^* are integer multiples of ϵ .

As long as bids are strictly below b^* , the resulting budget is in the “conflict zone” $[\hat{B}_2, B_2^*]$, and the lowest bidder will raise to take the turn from the other bidder. Suppose now that $b_1 = b_2 = b^*$. If $\hat{B}_2 = 0$, then $b_2 = b^* = B_2$ and black cannot raise. Otherwise, by definition of b^* if black raises his bid to $b'_2 = b^* + \epsilon$, then

$$t^* \geq_2 \mu_r(\hat{B}_2 - \epsilon) = \mu_r(B_2 - (b^* + \epsilon)) = \mu_r(B_2 - b'_2),$$

so black cannot gain by selecting s_r . Thus $b_1 = b_2 = b^*$ is an equilibrium, where $\mu(B_1) = \mu_l(B_1^*)$.

Case III: $t_r^* >_1 t^*$, i.e. white strictly prefers s_r at budget B_1^* . Quite expectantly, the proof of case III is similar to case II, where white and black change roles. However, there are some fine issues due to tie-breaking, so we lay out the full proof.

Let $\hat{B}_1 \leq B_1^*$ be the lowest budget s.t. for every $\beta_1 \in [\hat{B}_1, B_1^*]$, $\mu_r(\beta_1) >_1 t^*$. Set $\hat{t}_r = \mu_r(\hat{B}_1)$. As in case II, whenever the budget after the bid is in the range $[\hat{B}_1, B_1^*]$, white prefers s_r whereas black prefers s_l . We define $B_1 = \frac{\hat{B}_1 + B_1^*}{2}$, $b^* = \frac{B_1^* - \hat{B}_1}{2}$, and argue that the equilibrium reached in s_0 is $b_1 = b^*$, $b_2 = b^* + \epsilon$.

Indeed, if $b_1 < b_2 \leq b^*$, then black selects s_l under the remaining budget $B_2 - b_2 \geq B_2^*$, and we have $\mu_r(B_1 - b_2) >_1 t^* = \mu_l(B_1^*) = \mu_l(B_2 - b^*) \geq_1 \mu_l(B_2 - b_2)$. I.e., white will raise her bid to $b'_1 = b_2$.

Similarly, if $b_1 = b_2 \leq b^*$, then white selects s_r , but black wants to raise to $b'_2 = b_2 + \epsilon$: since $B_2 - b'_2 \geq B_2 - b^* - \epsilon = B_2^* - \epsilon$, we have⁸

$$\mu_l(B_2 - b'_2) \geq_2 \mu_l(B_2^* - \epsilon) = \mu_l(B_2^*) = t^* >_2 \hat{t}_r = \mu_r(B_2 + b^*) \geq_2 \mu_r(B_2 + b_1).$$

It remains to show that $b_2 = b_1 + \epsilon = b^* + \epsilon$ is an equilibrium, i.e. that white will not raise. This is exactly as in case II. Either white cannot raise, or $b'_1 > b^*$. In the latter case, $\mu_r(B_1 - b'_1) \leq_1 \mu_r(B_1 - b^* - \epsilon) = \mu_r(\hat{B}_1 - \epsilon) \leq_1 t^*$, where the last inequality is by definition of \hat{B}_1 . Thus white cannot gain by selecting s_r , and clearly not by keeping s_l at a lower budget. \square

Proof of Theorem 7(D) (Pareto-optimal). Let b_1, b_2 be equilibrium bids in s_0 reached by the above process. We use throughout the proof the fact that by monotonicity of both subgames, it only makes sense for a player to raise her bid if she selects the branch that is not currently played. Note that we may not use the fact that μ itself is monotone, since this will be the last thing we prove.

Assume w.l.o.g. that $b_1 = b_2$ (white takes) and that under budget $B_1 - b_1$, white prefers s_l . We argue that $t_l^* = \mu_l(B_1 - b_1)$ is Pareto-efficient in G . Assume, toward a contradiction, that there is some $t_r^* \in T_P(G_r)$ that Pareto dominates t_l^* (i.e., $t_r^* >_{1,2} t_l^*$). Since $\gamma^{(r)}$ has property (C) by induction, $t_r^* = \mu_r(B'_1)$ for some budget B'_1 .

Case 1: Suppose first that $B'_1 > B_1$, i.e. that white must drop the turn in order to reach s_r with sufficient budget. Denote $b' = B'_1 - B_1$ (note that $b' = B_2 - B'_2 \leq B_2$). We argue that $b' > b_1$. If $b_1 = 0$, then clearly $b' > 0 = b_1$, thus suppose $b_1 > 0$. Then by construction of the equilibrium (b_1, b_2) , t_l^* is strictly preferred by white over the outcome of $(b_1 - \epsilon, b_2)$, which is $\mu_r(B_1 + b_1)$ (as $b_2 = b_1$). Thus,

$$\mu_r(B_1 + b') = t_r^* >_1 t_l^* >_1 \mu_r(B_1 + b_1).$$

By monotonicity, we must have $b' > b_1 = b_2$. Next, recall that $t_r^* >_2 t_l^*$. Thus black could raise his bid to $b' > b_1 = b_2$ and strictly gain. This shows that (b_1, b_2) are not an equilibrium in s_0 , in contradiction to the construction of γ .

Case 2: $B'_1 = B_1 - b'$ for some $b' \geq 0$. If $b_1 = 0$, then $\mu_r(B_1) \geq_1 \mu_r(B'_1) = t_r^* >_1 t_l^* = \mu_l(B_1)$. That is, white prefers s_r but chose s_l . A contradiction.

Thus suppose $b_1 > 0$, and we will show that this will lead to a contradiction.

Since $\mu_r(B_1 - b') = t_r^* >_1 t_l^* \geq_1 \mu_r(B_1 - b_1)$, we must have by monotonicity of μ_r that $0 \leq b' \leq b_1$. By the construction of the equilibrium bids, white strictly preferred $t_l^* = \mu_l(B_1 - b_1)$ over the previous state $\mu_r(B_1 + b_2) = \mu_r(B_1 + b_1)$ (where the bid of player 1 was strictly lower). Thus $t_r^* \geq_1 t_l^* >_1 \mu_r(B_1 + b_1) \geq_1 \mu_r(B_1 - b') = t_r^*$, which is a contradiction. \square

Proof of Theorem 7(E) (Monotonicity). By the induction hypothesis, both of μ_l, μ_r are monotone. Let $B'_1 = B_1 + \epsilon$, and let b_1, b_2 be the equilibrium bids in $\gamma(s_0, B_1)$. Likewise, b'_1, b'_2 are the equilibrium bids in $\gamma(s_0, B'_1)$, where $B'_1 = B_1 + \epsilon$.

Case 1: Suppose white takes under $\gamma(s_0, B_1)$, thus $b_1 = b_2$ by property (A). By our assumption, player 1 selects s_l at $\gamma(s_0, B_1)$, and the outcome is $t_l = \mu_l(B_1 - b_1) \geq_1 \mu_r(B_1 + b_2)$ (in fact the inequality is strict if $b_1 > 0$, or else player 1 would never bid b_1).

We argue that $b'_2 \leq b_1 + \epsilon$. Otherwise, it means that Algorithm 2 reached bids $(b_1 + \epsilon, b_1 + \epsilon)$ under budget B'_1 and then black raised again. However, at bid $b_1 + \epsilon$ white remains with $B'_1 - (b_1 + \epsilon) = B_1 - b_1$ and thus we know the outcome is $t'_l = \mu_l(B_1 - b_1) \geq_2 \mu_r(B_2 - (b_1 + \epsilon))$ (since black does not raise to $b_1 + \epsilon$ under B_1). By raising under B'_1 , the outcome for black will be

$$\mu_r(B'_2 - (b_1 + 2\epsilon)) = \mu_r(B_2 - b_1 - 3\epsilon) \leq_2 \mu_r(B_2 - (b_1 + \epsilon)) \leq_2 t'_l,$$

that is, black does not gain.

We conclude that $b'_1, b'_2 \leq b_1 + \epsilon$. Then, by bidding $b'_1 = b_1 + \epsilon$ at budget B'_1 and selecting s_l , white can guarantee a value of $\mu_l(B'_1 - b'_1) = \mu_l(B_1 + \epsilon) - (b_1 + \epsilon) = \mu_l(B_1 - b_1)$. Thus its current value $\mu(B'_1)$ is at least the same, and we get $\mu(B'_1) \geq_1 \mu_l(B_1 - b_1) = \mu(B_1)$.

Case 2: Next, we prove for the case where black takes under $\gamma(s_0, B_1)$ (i.e., $b_2 = b_1 + \epsilon$). Note that in particular $b_2 \geq \epsilon$. W.l.o.g. black selects s_r , and the outcome is $t_r = \mu(B_2) = \mu_r(B_2 - b_2) = \mu_r(B_1 + b_2)$.

⁸ The equality is by our selection of B_2^* so that it is not on the edge of its respective interval. This is the only place where our assumption of high resolution is applied. However this seemingly innocuous assumption is critical, as we show in the next section.

By Lemma 8 and the way Algorithm 2 works, player 1 selects the same child of s_0 (w.l.o.g s_l) at any bid at or below the equilibrium bids. Thus, if $b_1 > 0$ then for every $\hat{b}_1 \leq b_1$, $\mu_l(B_1 - \hat{b}_1) >_1 \mu_r(B_1 - \hat{b}_1)$. Player 2 always prefers the other child (s_r), unless both bid 0 in equilibrium, thus for every $\hat{b}_2 < b_2$, $\mu_r(B_2 - \hat{b}_2) >_2 \mu_l(B_2 - \hat{b}_2)$.

We claim that $b'_2 > b_2 - 2\epsilon$. Assume, towards a contradiction, that $b'_2 \leq b_2 - 2\epsilon$ (in particular $b_1, b_2 > 0$). Suppose that black wins, i.e. that $b'_2 = b'_1 + \epsilon$ and $t' = \mu_r(B'_2 - b'_2)$. Note that b'_2 (and reaching $t_r^* = \mu_r(B_2 - b'_2)$) is not a winning equilibrium bid under B_1 , i.e., by bidding $b_1^* = b'_2$ white remains with $B_1 - b_1^*$ and gets a strictly better outcome $\mu_l(B_1 - b'_2) >_1 t_r^*$. Therefore (the second inequality is by monotonicity of μ_r),

$$\mu_l(B_1 - (b'_2 + \epsilon)) = \mu_l(B_1 - b'_2) >_1 t_r^* = \mu_r(B_2 - b'_2) \geq_1 \mu_r(B'_2 - b'_2) = t',$$

i.e., white has a deviation from (b'_1, b'_2) in $\gamma(s_0, B'_1)$.

Now suppose that white wins, i.e. that $b'_2 = b'_1$ and $t' = \mu_l(B'_1 - b'_1)$. We use the fact that (b_1, b_1) is not an equilibrium in $\gamma(s_0, B_1)$, as black raised to $b_2 = b_1 + \epsilon$. Note also that $b'_2 < b_2 - \epsilon = b_1$.

$$\begin{aligned} \mu_r(B'_2 - b_1) &= \mu_r((B_2 - \epsilon) - b_1) = \mu_r(B_2 - (b_1 + \epsilon)) = \mu_r(B_2 - b_2) \\ &>_2 \mu_l(B_2 + b_1) && \text{(by the last increase of black in } \gamma(s_0, B_1)) \\ &\geq_2 \mu_l(B_2 + b'_1) \geq_2 \mu_l(B'_2 + b'_1) && \text{(by monotonicity of } \mu_l) \\ &= \mu_l(B'_1 - b'_1) = t'. \end{aligned}$$

This means that black can strictly gain at $\gamma(s_0, B'_1)$ by raising his bid from b'_2 to b_1 . In either case (black or white win), we get a contradiction to (b'_1, b'_2) being an equilibrium, which proves the claim.

Thus $b'_2 \geq b_2 - \epsilon$. We next lower-bound the utility of the new outcome $t' = \mu(B'_1)$ to player 1: $t' \geq_1 \mu_r(B'_1 + b'_2)$, since if black wins in $\gamma(s_0, B'_1)$ then the right-hand side is t' , and if white wins then the outcome is strictly better for her, or otherwise she would not raise. Then,

$$t' \geq_1 \mu_r(B'_1 + b'_2) = \mu_r(B_1 + b'_2 + \epsilon) \geq_1 \mu_r(B_1 + b_2) = t_r,$$

where the second inequality is by monotonicity of μ_r .

Finally, the above proof only shows monotonicity for the white player. Monotonicity for the black player now follows immediately from property (D). Since $\mu(B_2) = \mu(B_1)$ is (weakly) better for white than $\mu(B_2 + \epsilon) = \mu(B_1 - \epsilon)$, it cannot be strictly better for black. Otherwise $\mu(B_2 + \epsilon)$ would be Pareto dominated by $\mu(B_2)$. Thus $\mu(B_2 + \epsilon) \geq_2 \mu(B_2)$. \square

4.2. Why bottom equilibrium?

While there may be multiple equilibria, we argue that the Bottom Equilibrium, when exists, is the most natural one. One justification is that players may want to maintain as much budget as they can in each round, given that their utility in the rest of the game is monotone in the remaining budget. In other words, if we focus on the single bidding round, each player is inclined to bid the minimal amount.

A second informal justification is that in the Bottom Equilibrium each player is nearly indifferent between winning and losing the round: small deviations (i.e. higher or lower bid) of the other player $-i$ cannot substantially hurt i . Further, the loser of the round in the Bottom Equilibrium never regrets if the other player bids too low and becomes the loser. On the other hand in other PSPEs the loser has a “risk” of becoming a winner with very low budget. Thus it is safer in some sense to play the Bottom Equilibrium.

Our third justification is that we can maintain the Bottom Equilibrium as the *unique* PSPE by slightly changing the rules of the auction.⁹ Indeed, suppose that instead of a sealed-bid auction in every step, we hold an ascending auction (similar to an English auction), where in each step the price rise by ϵ and the loser has a chance to accept the new price and become a winner. Since this is essentially equivalent to the auction being simulated in Algorithm 2, the unique outcome is the Bottom Equilibrium.

Corollary 10. Any binary bidding game $\langle G, \epsilon \rangle$ with the above ascending auction rule has a unique PSPE, which is the Bottom Equilibrium of $\langle G, \epsilon \rangle$.

4.3. Low budget resolution

We showed in Section 3 that a PSPE exists regardless of the budget resolution. However for our main theorem we also assumed “high resolution” that is at least exponential in the height of the game tree. We next show that high resolution is necessary for Pareto Efficiency (main part of Theorem 7) to hold.

⁹ We thank David Parkes for this observation.


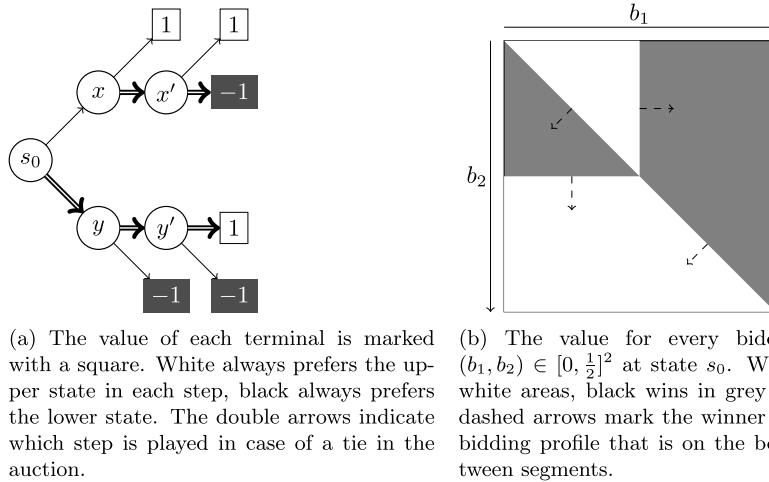


Fig. 9. The game G_k . Pareto-optimal terminals are marked with *.



(a) The value of each terminal is marked with a square. White always prefers the upper state in each step, black always prefers the lower state. The double arrows indicate which step is played in case of a tie in the auction.

(b) The value for every bidding profile $(b_1, b_2) \in [0, \frac{1}{2}]^2$ at state s_0 . White wins in white areas, black wins in grey areas. The dashed arrows mark the winner in case of a bidding profile that is on the boundary between segments.

Fig. 10. A game with continuous budgets, inconsistent tie-breaking, and no SPE.

Proposition 11. For any $k \geq 2$ there is a binary bidding game structure G_k of height k , s.t. for any $1 < R' < 2^{k-1}$, the game $(G_k, \frac{1}{R'})$ with $B_1 < 0.5$ has no Pareto-optimal PSPEs.

Proof. Consider the binary game G_k , described in Fig. 9. In order to reach one of the Pareto-optimal leaves (marked with *), the same player must win $k - 1$ times in a row after a branch is selected (as each branch is an ordinal 0-sum game). Since $2^{k-1}B_2 > 1$, and the budget of the loser doubles with each loss, player 1 cannot win in all $k - 1$ times unless $B_2(y_1) = 0$. Similarly, since after every win of player 2, the budget of player 1 will be at least $\frac{1}{R'} > 2^{1-k}$, player 2 cannot win in all $k - 1$ times either (even if initially $B_2(x_1) = 1$).

Therefore, if after the first turn the selected child is x_1 , the outcome will be $(8, 1)$; and if the selected child is y_1 then the outcome will be $(1, 8)$ (again, unless $B_2(y_1) = 0$). The only way to reach $B_2(y_1) = 0$ is if black bids all of his budget $b_2 = B_2$ at s_0 and then selects y_1 . However, this is not an equilibrium: since $b_1 \leq B_1 < b_2 - \frac{1}{R'}$, black can decrease his bid to $B_1 + \frac{1}{R'}$ and gain 8 instead of 7.

Note that due to our tie-breaking assumption it is impossible for agent 1 to similarly surrender her budget, since agent 2 must bid at least $\frac{1}{R'}$ to win each turn. \square

The example above does not cover the case of $R' = 1$ (“single coin”). In that case, consider the game G'_k (for some $k \geq 2$) and a root node s' whose children are s_0 and $(1, 8)$. If player 2 starts with the coin ($B_2 = 1$) then in the only equilibrium he wins the first turn and selects $(1, 8)$, as he cannot get more than 7 by reaching s_0 .

5. Continuous bids

We denote by $(G, -)$ the game G with continuous bids, where the set of allowed budget partitions is $\mathcal{B} = [0, 1]$ and $B_1 + B_2 = 1$. Note that the strategy space in such games is infinite, and thus it is not a-priori clear that these games have a subgame perfect equilibrium, let alone a PSPE. Indeed, we show that if the tie-breaking method can depend on the node, then there are games, even zero-sum games, where no equilibrium exists.

Proposition 12. There is a zero-sum bidding game with node-specific tie breaking, that has no subgame-perfect equilibrium.




Fig. 11. Equilibrium bids at $\bar{\gamma}_s(\hat{B}_1)$.

Proof. Consider the game depicted in Fig. 10a. This is a zero-sum game where the white player is the maximizer. By the tie-breaking rule, white wins in x' if her budget is strictly above 0.5, and wins in y' if her budget is at least 0.5. In state x , white needs to take at least one turn to win the game. Thus a budget of $B_1(x) > 0.25$ is sufficient. Similarly, in state y white has to take twice, and thus must have a budget of $b_1(y) \geq 0.75$ to win the game.

Clearly, if the initial budget partition in s_0 is $B_1 \geq 0.5 + \delta$ for some $\delta > 0$, then white can win by bidding $b_1 = 0.25 + \delta/2$: either she takes the turn and reaches x with $0.25 + \delta/2$, or loses the turn and reaches to y with $B_1 + b_2 \geq 0.5 + \delta + (0.25 + \delta/2) = 0.75 + 3\delta/2 > 0.75$. Similarly, if $B_1 < 0.5 + \delta$ then black can bid $0.25 + \delta/2$ and force white to lose.

Now, suppose $B_1 = B_2 = 0.5$. We first show that there is no value in pure strategies, by showing that for any bid of one player, the other player has a bid that leads to victory. The outcome for any pair of deterministic bids (b_1, b_2) is displayed in Fig. 10b.

Let $b_1 \in [0, B_1]$. If $b_1 < 0.25$ then set $b_2 = b_1$. Thus black takes the turn and gets to y with $B_2(y) > 0.25$, which is sufficient for victory. If $b_1 \geq 0.25$, then black can bid $b_2 < b_1$. Now white remains with $B_1(x) = 0.5 - b_1 \leq 0.25$, which means black wins. Thus white cannot guarantee any value above -1 .

We next consider the black player. Let $b_2 \in [0, B_2]$. If $b_2 < 0.25$ then denote $b_2 = 0.25 - \delta$, and set $b_1 = b_2 + \delta/2 < 0.25$. Then white takes and gets to x with $B_1(x) = 0.5 - b_1 > 0.25$, which is sufficient for victory. If $b_2 \geq 0.25$, then set $b_1 < b_2$. Now black remains with $B_2(y) = 0.5 - b_2 \leq 0.25$, which means white wins. Thus black cannot guarantee any value below 1.

It is not hard to generalize the above argument to see that there is no value even in mixed strategies. Indeed, let $b_2 \in \Delta([0, B_2])$ be any mixed strategy of black. There must be an open interval $A = (0.25 - \delta, 0.25)$ s.t. b_2 assigns a probability of less than 0.01 to A . Then white can bid $b_1 = 0.25 - \delta/2$ and win w.p. of at least 0.99: by our arguments on pure strategies, the only bids of black that beat b_1 are in the range $[b_1, 0.25) \subseteq A$. Therefore the value of the game under equal initial budgets is at least $0.99(1) + 0.01(-1) = 0.98$.

Similarly, let $b_1 \in \Delta([0, B_1])$ be some mixed strategy of white. Again there is some open interval $A' = (0.25 - \delta, 0.25)$ that is played by white w.p. less than 0.01. Then by playing $b_2 = 0.25 - \delta/2$ black can win w.p. of at least 0.99, and the value cannot be above -0.98 . \square

Consider a game structure G . With continuous bids, the Bottom Equilibrium is not well-defined, as positive equilibrium bids may be arbitrarily small. However, we next show that playing with continuous bids is essentially the same as playing with high budget resolution. Denote by $\bar{B}^j(k) = [B_1^j(k), B_1^{j+1}(k))$ the j 'th continuous budget interval of level k .

Theorem 13. Let $\gamma = BE(G, \epsilon)$ for some binary game G and high resolution ϵ . There is a PSPE $\bar{\gamma}$ of $\langle G, - \rangle$ such that for all $j < 2^k$ and for all $\hat{B}_1 \in \bar{B}^j(k)$, $\bar{\mu}(\hat{B}_1) = \bar{\mu}(B_1^j) = \mu(B_1^j)$ (where $\bar{\mu} = \mu_{\bar{\gamma}}$).

In other words, there is a PSPE of $\langle G, - \rangle$ that for any budget in the interval yields exactly the same outcome we would get in the Bottom Equilibrium of $\langle G, \epsilon \rangle$ for the same budget interval.

Proof. The proof is (as usual) by induction. Let $(b_1, b_2) = \gamma_s(B_1^j)$ be the equilibrium bids in the discrete game. We denote by j' and j'' the respective level $k - 1$ intervals of $B_1^j - b_1$ and $B_1^j + b_2$. By induction, $\bar{\mu}(B_1) = \bar{\mu}(B_1^{j'}) = \mu(B_1^{j'})$ for all $B_1 \in \bar{B}^{j'}(k - 1)$ and likewise for j'' . So that in level $k - 1$ it only matters in which budget interval we arrive.

For each $\hat{B}_1 \in \bar{B}^j(k)$ we need to define equilibrium bids $\hat{b}_1 \leq \hat{B}_1, \hat{b}_2 \leq \hat{B}_2$ so that: (a) $\hat{B}_1 - \hat{b}_1 \in \bar{B}^{j'}(k - 1)$; (b) $\hat{B}_2 - \hat{b}_2 \in \bar{B}^{j''}(k - 1)$; (c) $\hat{b}_1 \geq \hat{b}_2 \iff b_1 \geq b_2$. This will guarantee that any deviation to a different interval at $\bar{\gamma}_s(\hat{B}_1)$ entails a deviation in $\gamma_s(B_1^j)$. In addition, we have to make sure that (d) no agent has a new deviation to a different point inside the interval (i.e. one that was not previously available due to the discrete budgets).

Let $\delta = B_1^{j+1} - \hat{B}_1 > 0$. All five cases of Lemma 9 map into only three cases.

- Case 0: $b_1 = b_2 = 0$. We set $\hat{b}_1 = \hat{b}_2 = 0$ as well. (a)+(b)+(c)+(d) trivially hold.
- Case 1 (both L and W): $b_1 = b_2 > 0$. Note that according to both Cases 1L and 1W of Lemma 9, $B_1^{j'} + B_1^{j''} - 2B_1^j + 2^{-(k-1)} = 0$.

We define $\hat{b}_1 = B_1^{j''+1} - \hat{B}_1$, and $\hat{b}_2 = \max\{0, \hat{B}_1 - B_1^{j'+1} + \delta\}$ (see Fig. 11). Thus (a)+(b) hold by construction. (c) holds even when $\hat{b}_2 > 0$ since

$$\begin{aligned} \hat{b}_1 - \hat{b}_2 &= B_1^{j''+1} - \hat{B}_1 - (\hat{B}_1 - B_1^{j'+1} + \delta) \\ &= B_1^{j''+1} - 2(B_1^{j'+1} - \delta) + B_1^{j'+1} - \delta = B_1^{j''+1} - 2B_1^{j'+1} + B_1^{j'+1} + \delta \\ &= B_1^{j''} + 2^{-(k-1)} + B_1^{j'} + 2^{-(k-1)} - 2B_1^j - 2 \cdot 2^{-k} + \delta = \delta > 0. \end{aligned}$$

As for (d), for any $b'_2 > \hat{b}_1$, we have

$$\hat{B}_1 + b'_2 > \hat{B}_1 + \hat{b}_1 = B_1^{j''+1},$$

i.e. we get a point outside the interval $\overline{\mathcal{B}}^{j''}$, so black cannot gain by bidding higher. White cannot gain by bidding $b'_1 \in [\hat{b}_2, \hat{b}_1)$ since $\hat{B}_1 - b'_1 \leq \hat{B}_1 - \hat{b}_2 = B_1^{j'+1} - \delta \in \overline{\mathcal{B}}^{j'}$, i.e. in the same budget interval.

- Case 2 (both L and W): $b_2 = b_1 + \epsilon$. We set $\hat{b}_1 = \hat{B}_1 - B_1^{j'}$ and $\hat{b}_2 = B_1^{j''+1} - B_1^{j'+1}$. Clearly (a) + (b) holds. (d) also holds since (as in Case 1) it is impossible for $-i^* = 1$ to deviate and stay in the same interval, and lowering the winning bid of $i^* = 2$ still reaches interval $\mathcal{B}^{j''}(k-1)$. It remains to show (c).

In both Cases 2L and 2W of Lemma 9 we have

$$\text{Case 2L: } B_1^j - B_1^{j'} + \epsilon = b_1 + \epsilon = b_2 = (B_1^{j''} + \epsilon) - B_1^j \Rightarrow B_1^{j''} + B_1^{j'} - 2B_1^j = 0;$$

$$\text{Case 2W: } B_1^j - B_1^{j'} = b_1 + \epsilon = b_2 = B_1^{j''} - B_1^j \Rightarrow B_1^{j''} + B_1^{j'} - 2B_1^j = 0.$$

Then we get that similarly to Case 1,

$$\begin{aligned} \hat{b}_2 - \hat{b}_1 &= B_1^{j''+1} - B_1^{j'+1} - (\hat{B}_1 - B_1^{j'}) \\ &= B_1^{j''+1} - B_1^{j'+1} - ((B_1^{j'+1} - \delta) - B_1^{j'}) = B_1^{j''+1} + B_1^{j'} - 2B_1^{j'+1} + \delta \\ &= B_1^{j''} + 2^{-(k-1)} + B_1^{j'} - 2B_1^j - 2 \cdot 2^{-k} + \delta = \delta > 0, \end{aligned}$$

as required. \square

One immediate corollary is that as long as the resolution is high, the Bottom Equilibrium induces the same mapping from budget intervals to outcomes (since for any ϵ the mapping is the same one as in the continuous case). In the continuous case we may have (infinitely) many PSPEs, all with the mapping $\overline{\mu}$, but neither of them can be referred to as “the Bottom Equilibrium.” We thus say that a PSPE $\overline{\gamma} \in \Gamma^*(G, -)$ is a *Bottom Equilibrium* if $\mu_{\overline{\gamma}} = \overline{\mu}$, i.e. if it induces the same outcome as $BE(G, \epsilon)$ for some high resolution ϵ .

6. From bidding games to combinatorial bargaining

There are two ways we can think of bidding games. First, we can look for games that inherently have budgets and sequential bidding. While there are some examples of scenarios that arguably fit into this model, typically the incentive structure and/or the bidding structure is quite different (see Section 6.4 below, and the Related work section). A more promising way is to think of the sequential bidding process as a *mechanism* that is designed to increase cooperation and welfare among a pair of players in various scenarios.

6.1. Sequential scrip bargaining

Definition 13. A 2-player *sequential scrip bargaining* (SSB) is a tuple $F = \langle K, v_1, v_2, \tau \rangle$, where: K is the set of items ($k = |K|$); $v_i : 2^K \rightarrow \mathbb{R}$ is the value function of agent $i \in N$; and τ is a permutation over items K .

Intuitively, items K are offered for sale according to order τ , and in each turn agents bid for the current item. To complete the game we need to describe the bidding rules. As expected, we assign a budget B_i to each agent; the highest bidder in each turn pays her bid to the other agent, and decides who gets the item. In case of a tie we treat agent 1 as the winner.

This bargaining mechanism seems like a close reminiscent of *sequential auctions* (Gale and Stegeman, 2001; Bae et al., 2007). However, typically in the auction literature items are auctioned in a first- or second-price auction where the highest bidder pays the (first or second) bid to the seller; and utilities are quasi-linear. In particular, an agent’s utility depends not just on the allocation, but also on the amount of money the agent paid for her items (see Section 8.1).

SSBs differ from these standard auctions in two ways. First, while we use the first-price auction rule, the higher bidder pays the other player, rather than paying to the seller, highlighting that this is a bargaining mechanism rather than an auction. Second and more importantly, the budget B_i is “scrip money” and has no value outside the game, thus the utility of i from a bundle $S_i \subseteq K$ is $v_i(S_i)$, regardless of how much of the budget the agent spent.

6.2. SSBs are binary bidding games

Proposition 14. Every SSB with k items has an equivalent bidding game structure over a binary tree of height k , and vice versa.

By “equivalent to” we mean that for any resolution ϵ , there is a one-to-one mapping between strategy profiles in the original and induced game, that preserves utilities. Intuitively, we can describe any SSB as a binary tree where any internal node at depth k' corresponds to an allocation of the first k' items. Note that this mapping is independent of the initial budget allocation, the budget resolution, and even of the bidding rule.

Proof. In the first direction, given $F = \{K, v_1, v_2, \tau\}$, let $K_d \subseteq K$ be the set of first d elements of K (according to τ). We construct a bidding game $G = \langle S, s_0, T, g, u_1, u_2 \rangle$ as follows. Let S_d be the set of all 2^d possible bi-partitions of K_d for $d \in \{0, 1, \dots, k\}$, and $S = \bigcup_{d \leq k} S_d$. Note that these are partial allocations of K to the two players. s_0 is the empty allocation (the unique member of K_0). Every state can be written as $s = (K^1, K^2) \in S_d$, where K^i is the set of items held by player i in state s . There are two states accessible from s : either player 1 or player 2 get item $\tau(d+1)$. Thus $g(s)$ contains exactly two states in S_{d+1} . The set of terminals T coincides with S_k , i.e. all full partitions of K . Finally, for every $t = (K^1, K^2) \in T$, we set $u_i(t) = v_i(K^i)$. Thus G is an impartial binary tree of height k , which is clearly equivalent to F .

In the other direction, suppose we are given a binary bidding game G of height k . W.l.o.g. G is a balanced tree: if it is in DAG form, we can clone every node with several incoming edges, along with its subtree; if some branches are shorter than k , we extend them in the trivial way: replace each terminal t at depth $k' < k$ with a balanced binary subtree of height $k - k'$, all of whose leaves are identical to t .

Now we have a balanced tree of height k . For every internal node, arbitrarily label one child as “left” and the other as “right”. We construct an SSB $F = \langle k, v_1, v_2, \tau \rangle$ as follows. Let τ be the identity permutation over $K = [k]$. We identify any path from the root s_0 to a terminal t with a partition (K^1, K^2) of K where K^1 contains all levels d s.t. the left child was selected, and K^2 contains all other levels (i.e., where the right child was selected). Note that any $K' \subseteq K$ appears in exactly one terminal t'_1 as K^1 , and in exactly one terminal t'_2 as K^2 . We set $v_i(K') = u_i(t'_i)$. Thus the set of allocations in F coincides with the set of paths in the original game G . \square

For example, the game G_{maj} (Sec. 2.1) coincides with a sequential scrip bargaining over three identical items, where each agent assigns a value of 1 to bundles of size two or more, and 0 otherwise.

Since SSBs are essentially equivalent to bidding games on binary trees, an immediate corollary of Theorem 7 and Proposition 14 is the following. For a set of items K , let $T_P(K)$ be the set of Pareto-efficient allocations of K among two players.

Theorem 15. Let $F = \langle K, v_1, v_2, \tau \rangle$ be a sequential scrip bargaining, and let ϵ be some high budget resolution. Then $\langle F, \epsilon \rangle$ has a Bottom Equilibrium γ , s.t. $\mu = \mu_\gamma$ is a monotone and surjective mapping from B_ϵ to $T_P(K)$.

As we mentioned in the introduction, the set of items K need not be a collection of physical items. We next consider some applications of SSBs.

6.3. Applications for SSB mechanisms

Multi-issue voting. Consider a set of voters, voting over k binary issues (Winter, 1997; Lacy and Niou, 2000; Lang and Xia, 2009). Here the voters (who are the players) have a complete (weak) preference order over the 2^k outcomes, but need not attribute cardinal utilities to them. It is known that when preferences are unrestricted, either truthful or strategic voting may lead to the selection of a Pareto-dominated outcome. Voting on the issues sequentially provides a partial solution, but Pareto-dominated outcomes may still be selected (Lacy and Niou, 2000). While our model only applies for two voters, it provides strong guaranties of Pareto-efficiency, regardless of the agenda or the structure of preferences. In particular, a Condorcet loser is never selected, and the Condorcet winner is always selected if it exists.

Arbitration. For two self-interested parties, a bidding game can be used as an arbitration mechanism that reaches an efficient outcome, and also allows a natural way to take into account parties of different importance, or weight. As a concrete example, we can think of the Democrats and the GOP bargaining over the various clauses of the Healthcare Act, or some other reform. The initial budget of each party can be set, say, based on its number of seats. Arbitration mechanisms for two parties (not necessarily in a combinatorial setting) have been widely studied in the literature (Sprumont, 1993; Anbarci, 2006; De Clippel et al., 2012).

The *minimal satisfaction test (MST)* (De Clippel et al., 2012) is an ordinal fairness criterion, that is satisfied if the selected outcome is above the median outcome in every player’s preference order. That is, every player weakly prefers the selected outcome to at least half of all possible outcomes. De Clippel et al. studied several sequential mechanisms, and showed that three of them—The *Alternate-Strike* mechanism (Anbarci, 2006), the *Voting by Alternating Voters and Vetoes* mechanism (Sprumont, 1993), and the *Shortlisting* mechanism—each implement some Pareto-efficient outcome that satisfies MST.

SSB and fair arbitration mechanisms. Consider the following parametrized extension of the MST requirement. We say that an outcome t satisfies the (α_1, α_2) -satisfaction test, if each player i weakly prefers t to at least $\lceil \alpha_i |T| \rceil$ outcomes.

Proposition 16. *Let (G, ϵ) be a bidding game over a full binary tree, then any PSPE outcome satisfies the (B_1, B_2) -satisfaction test.*

In particular, for equal budget we get that the outcome satisfies the minimal satisfaction test of De Clippel et al. (2012).

Proof. We will prove that there is a strategy for player 1 that *guarantees* an outcome weakly preferred to $\lceil B_1 |T| \rceil$ outcomes, and likewise for player 2 and $\lceil B_2 |T| \rceil$. Since any equilibrium outcome is at least as good as the maximin outcome, this entails the proposition.

Given G and B_1 , we define a zero-sum game structure G' , that is equivalent to G except for the utility functions. We set $u'_1(t) = 1$ if t is weakly preferred to at least $\lceil B_1 |T| \rceil$ terminals, and -1 otherwise; and set $u'_2 = -u'_1$. Thus each player can either ‘win’ or ‘lose’ in every terminal. Note that if $B_1 |T|$ is an integer, then there are exactly $|T| - B_1 |T| + 1 = (1 - B_1) |T| + 1$ winning nodes, and $B_1 |T| - 1$ losing nodes. Otherwise there are $\lfloor B_1 |T| \rfloor$ losing nodes, since the next one is weakly preferred to all of them plus itself, i.e. to $\lceil B_1 |T| \rceil$ terminals. Thus the number of losing nodes is exactly $\lceil B_1 |T| \rceil - 1$.

According to Lazarus et al. (1996), the *Richman value* $R(s)$ of a node in a zero-sum win/lose game, is the minimal value s.t. any $B_1 > R(s)$ guarantees a victory for player 1. Lazarus et al. show that $R(s)$ equals *exactly* to the probability that player 1 loses in the corresponding “spinner game” of G' (the same game only turns are taken at random rather than by bidding).

This probability is easy to compute. Since G' is a zero-sum game, at each node, white (weakly) prefers one child and black prefers the other. W.l.o.g., white always prefers the right child and black always prefers the left. Thus if players take random turns, they reach a terminal that is selected uniformly at random from T . We conclude that

$$R(s_0) = Pr_{t \sim \text{Uniform}(T)}(u'_1(t) = -1) = \frac{|\{t : u'_1(t) = -1\}|}{|T|} = \frac{\lceil B_1 |T| \rceil - 1}{|T|} < B_1,$$

which means that a budget of B_1 is sufficient to guarantee a victory to player 1 in G' . By definition of G' , player 1 can guarantee an outcome in G that is weakly preferred to $\lceil B_1 |T| \rceil$ outcomes. The proof for player 2 is symmetric. \square

The SSB mechanism has the additional property of implementing the entire set of Pareto-efficient outcomes. As the proposition shows, every such outcome that is attained under a given budget allocation, satisfies the parametrized MST.

Communication complexity. Suppose that Alice and Bob each hold private information $a \in A$ and $b \in B$, respectively. They can exchange arbitrary messages sequentially in some predefined order. Given a function $f : A \times B \rightarrow \{0, 1\}$, the *communication complexity* of f is the total number of bits Alice and Bob must exchange so that at least one of them knows $f(a, b)$ with certainty (Yao, 1979).

Communication complexity in games was studied mainly in the context of equilibrium computation, where a and b are the utility functions of players in some game, and f is some property of the equilibrium of this game (Conitzer and Sandholm, 2004; Hart and Mansour, 2010). However, for an SSB with k items, the size of the valuation function of each player is already exponential in k (in the general case), so there is no much hope to compute an equilibrium with little communication.

We have a much more modest requirement. Let $H = (\Gamma, \Psi, h)$ be a 2-player game form, where $\Gamma = \Gamma_1 \times \Gamma_2$ is the set of valid pure strategy profiles, and Ψ is the set of possible outcomes, and $h(\gamma_1, \gamma_2) \in \Psi$ is the outcome of playing strategies γ_1, γ_2 in H .

Definition 14. The *playing communication complexity* of H is the number of bits players need to communicate so that at least one of them knows $h(\gamma_1, \gamma_2)$.

In our case, $\Psi = 2^K$ contains all possible allocations of K among the two players, and H is some mechanism the players use to determine the allocation.

For example, in the simple mechanism where players each choose sequentially an item from K until all items are selected, there are k rounds, and the communication required in each round is $\log(k)$ bits to specify which item. Thus the playing communication complexity of this mechanism is $O(k \log(k))$ (even though the strategies players use may be complicated).

On the other hand, consider the Shortlisting mechanism (De Clippel et al., 2012). The mechanism has two rounds where player 1 first marks the 2^{k-1} outcomes that she prefers to all others, and then player 2 selects among them. Clearly revealing the full strategy γ_1 to player 2 takes a number of bits that is exponential in k , and there does not seem to be a more efficient way to find $h(\gamma_1, \gamma_2)$. The other mechanisms in De Clippel et al. (2012) suffer from a similar problem.¹⁰

¹⁰ Recall that these mechanisms were not designed for combinatorial valuations so this is not surprising.

Note that given an order τ over items, the SSB $F(K, \tau)$ is a valid allocation mechanism that maps any pair of strategies $\gamma = (\gamma_1, \gamma_2)$ to an allocation $\mu_\gamma \in 2^K$.

Proposition 17. For any order τ , the playing communication complexity of SSB is $O(k^2)$.

Proof. An SSB has k rounds. In each round, each player needs to submit a bid, and the winner decides on whether to take or give the current item. For any constant resolution R , there are $R2^k$ possible budgets (and thus possible bids). So $\log(R2^k) = O(k)$ bits are sufficient to specify the bid. One additional bit is sufficient to specify the decision on the item. Thus in total each player is required to transfer $O(k^2)$ bits in order to play her strategy. \square

SSB is the first mechanism that guarantees efficiency, fairness, and low playing communication complexity. Thus SSB may be considered as a feasible and desirable mechanism for combinatorial bargaining/arbitration when the playing parties are asymmetric.

Zero-sum games. In a zero-sum game, player 1 gets an outcome that beats $B_1|T|$ outcomes, if and only if player 2 gets an outcome that beats $B_2|T|$ outcomes. Thus by Proposition 16 the set of outcomes and the utility functions define the (unique) PSPE outcome completely, regardless of the tree structure.

6.4. Other games

Consider any (impartial) game in tree-form. If the tree is not binary, we can modify it by recursively breaking decision nodes to a subtree of binary decisions. Then we can divide a large or continuous budget among the players, and let them play the new bidding game. If players adhere to the Bottom Equilibrium, then the initial budget partition will determine the final outcome (from $T_P(G)$). Thus we can get a good outcome even in games where playing by turns would lead to a poor outcome. In particular, with proper initial budgets we can implement the outcome with maximum social welfare, maximum Egalitarian welfare, etc.

Centipede games. One class of games that falls under the conditions of Theorem 7, is that of Centipede games (Rosenthal, 1981). Such games are a notorious example to how rationality leads players to end up in poor outcomes. Under random turn there are still Centipede games where players always choose to finish early, even though staying in the game is eventually significantly better for both.

Nevertheless, in our game setting there is a PSPE where players are guaranteed to continue until they reach one of the Pareto-efficient outcomes (the last two leaves).

Nash bargaining game. The Nash bargaining game (Nash, 1950) for two players is typically described by a (convex) set of feasible outcomes \mathcal{F} in the plane, whose boundary forms the Pareto-efficient frontier, and a status-quo point q . Given \mathcal{F} and q , we can think of an extensive-form bargaining game, where players start from q and at each state (point of the grid) they can either “go right” (increase the utility of player 1) or “go up” (increase the utility of player 2). The terminal states are the outcomes along the Pareto frontier. In this case it is clear that every outcome is Pareto-efficient (as this is how we defined the terminal states), but the properties of monotonicity and Pareto-surjective are still interesting. The (unique) PSPE induces a generalized solution concept for the Nash bargaining game (a point on the Pareto frontier for any budget allocation).

A natural question is how the PSPE solution—in particular for the case of equal budgets—compares with other solution concepts such as the Nash bargaining solution and the Kalai–Smorodinsky bargaining solution (Kalai and Smorodinsky, 1975). We note that there is much interest in non-cooperative implementations of various bargaining solutions, e.g. (Moulin, 1984).

To answer our question, we first modify the game tree described above, by complementing each leaf (a point on the Pareto frontier) to a subtree with identical leaves, such that all leaves are at the same depth. Clearly this does not change the equilibria of the game—it just means the game continues for a few more redundant rounds after the outcome is already determined. Since the SSB induced by a Nash bargaining game is ordinally-zero-sum (as all terminals are on the Pareto frontier), binary, and complete, it follows from Proposition 16 that when the budget is equally partitioned, the realized outcome is reached when players make the same number of moves. In other words, if we normalize q to the origin $(0, 0)$, then both players have the same utility (or the closest approximation to same utility), which coincides with the *Egalitarian*, or *proportional*, solution (Kalai, 1977).

7. Computational issues

In this section we elaborate on the algorithmic considerations in computing the Bottom Equilibrium. The key contribution in this section is Algorithm 3, presented in page 188. We also implemented the algorithm in Matlab (along with code to generate simple binary bidding games), and will share the code upon request.

Proposition 18. *Let G be a binary bidding game structure. Algorithm 3 computes a profile γ that is outcome-equivalent to a Bottom Equilibrium of $(G, -)$. Moreover, Algorithm 3 runs in time $|S| \cdot \text{poly}(|T|) = \text{poly}(|S|)$.*

Before showing the full proof, we explain the intuition. How can we even represent the equilibrium efficiently, if at any node there are infinitely many budgets to consider (or exponentially many in k , with discrete high resolution)? The answer lies in *monotonicity*. Since we can order all reachable terminals, and the outcome is monotone in the budget of player 1, we only need to find the cutoff points in the interval $[0, 1]$ where the next terminal becomes the outcome. There can be at most $|T| - 1$ such cutoff points, and the main challenge is to show: (a) that they can be computed efficiently from the cutoff points of both children; and (b) that for every budget between two given cutoff points, we can efficiently compute some bids that reach the correct equilibrium outcome.

Proof. We saw that at any level k budget interval $\bar{B}^j(k)$, the equilibrium bids in all $B_1 \in \bar{B}_1^j$ reach the same level $k - 1$ intervals (which we previously denoted by $\bar{B}^{j'}(k - 1), \bar{B}^{j''}(k - 1)$), although they may reach different budget points. Our algorithm only computes this mapping, i.e. from j to j' and j'' , and then assigns arbitrary bids that reach those intervals from B_1^j . These bids may not be stable themselves, but they contain enough information to reconstruct the equilibrium bids from any $B_1 \in \bar{B}_1^j$ according to Theorem 13.

There are three differences between Algorithm 2 (see page 175) and Algorithm 3: (a) the fast algorithm only computes the equilibrium in specific budget points (z_{lr}), rather than for any possible budget $B \in \mathcal{B}_\epsilon$ (or even any $B_1^j(k)$); (b) when black deviates, the fast algorithm increments the bid by a quantity larger than ϵ ; (c) after computing the equilibrium at all critical budget points, the fast algorithm “filters” some points out, and does not add them to the description of the PSPE. In contrast to ϵ in Algorithm 2, the exact size of δ is not important, as long as it is sufficiently small (anything below $2^{-\text{height}(G)-1}$ would suffice).

We need to show that despite those differences we get the same equilibrium outcome, i.e., the same mapping $\mu : [0, 1] \rightarrow T(G)$.

The a 'th budget interval is the union of several consequent subintervals of size 2^{-k} . Formally, $[F_s(a), F_s(a + 1)) = \bigcup_{2^k F_s(a) \leq j < 2^k F_s(a+1)} \bar{B}^j$.

We start by showing (b), i.e. that by bidding $b'_2 = b_1 + \delta$ or $b''_2 = \min\{B_1 - F_{s_1}(a_1 - 1), F_{s_2}(a'_2 + 1) - B_1 - \delta\}$, the same equilibrium is reached. We first argue that $b'_2 > b_1$ (so black indeed takes). Since $B_1 - b_1 \geq F_{s_1}(a_1)$, and F_{s_1} is strictly increasing, clearly $B_1 - F_{s_1}(a_1 - 1) > b_1$.

Similarly, $B_1 + b'_2 \in [F_{s_2}(a'_2), F_{s_2}(a'_2 + 1))$ by definition of f_{s_2} . Thus $F_{s_2}(a'_2 + 1) - B_1 - \delta > b'_2 - \delta = b_1$. We conclude that $b''_2 = \min\{B_1 - F_{s_1}(a_1 - 1), F_{s_2}(a'_2 + 1) - B_1 - \delta\} > b_1$.

Next, note that by bidding b'_2 or b''_2 , black reaches the same budget interval a'_2 in s_2 . The problem with just setting $b''_2 = F_{s_2}(a'_2 + 1) - B_1 - \delta$, is that this may prevent a further deviation by white, that would have been possible after a lower increment. That is, we may reach an equilibrium that is not the lowest-bids equilibrium. However, if white has a deviation after black, then we only care if she has a deviation that brings her all the way to the next budget interval $a_1 - 1$ in s_1 —otherwise black would deviate back to a'_2 with another δ increment and so on. Therefore constraining $b''_2 \leq B_1 - F_{s_1}(a_1 - 1)$ would prevent this problem by always allowing white to bid the same. In other words, it guarantees that $a'_1 = a'_1$.

As a concrete example, consider such an ascending auction in Fig. 12, for $B_1 = 0.4$. Whenever black increases his bid b''_2 , it should be to the next point among $b''_2 = F_{s_2}(2) - B_1 - \delta = 0.2 - \delta$, $b''_2 = B_1 - F_{s_1}(1) = 0.4$, and $b''_2 = F_{s_2}(3) - B_1 - \delta = 0.6 - \delta$, since if there is any deviation, the auction will eventually reach this point. The corresponding remaining budget $B_1 + b''_2$ for each of the three options is marked by d_1, d_2 and d_3 , respectively. Note that d_2 and d_3 both belong to the same budget interval of s_2 , but skipping d_2 may result in looking over an equilibrium where white bids all of B_1 and wins.

We now turn to show (a).

By induction, the budget between any consecutive critical points $[a_l, a_l + 1)$ yields the same outcome in s_l , and likewise for s_r . Fix some r, l and let $B_1 \in [z_{lr}, z_{lr+})$, where $z_{lr+} = \min\{z_{lr'} : z_{lr'} > z_{lr}\}$. We argue that whether white or black win in (s, z_{lr}) , the winner reaches the same next state s' and with the same remaining budget as from (s, B_1) . The reason is the same as in the proof of Theorem 7(B) and of Theorem 13 and we do not repeat it.

It is left to show (c), i.e. that no information is lost when filtering the $(|T| + 1)^2$ critical points. Due to monotonicity, \bar{T} is sorted, thus U contains the first index of every unique entry in \bar{T} . The filtering removes redundant parts of the equilibrium, as there are at most $|T|$ points where the outcome changes. Thus we only consider these points.

Finally, we consider the complexity of the fast algorithm. After filtering, the size of the lists F_s, T_s^*, A_s^1, A_s^2 is at most $|T|$ each. With every deviation, either a_1 strictly decreases or a_2 strictly increases, and thus there can be at most $2|T|$ deviations. The heaviest part in computing a deviation is calculating f_{s_i} . Since $|F_{s_i}| \leq |T|$ a naïve computation is linear in $|T|$. Therefore *Ascending - Auction()* runs in $O(|T|^2)$. It is called at most $|T|^2$ times in every internal node, thus the total complexity is $O(|S| \cdot |T|^4) = O(|S|^5)$. \square




Fig. 12. The top and bottom lines show the budget intervals in both children, where $F_{s_1} = (0, 0.2, 0.7)$ (say, for the left child), and $F_{s_2} = (0, 0.6)$. Any budget maps to a unique interval in each child. E.g., $f_{s_1}(0.3) = 2$ since $0.3 \in [F_{s_1}(2), F_{s_1}(3))$. Note that the rightmost interval always ends at 1. The middle line shows all potential cutoff points (before filtering) in node s : in increasing order, $(z_{lr})_{lr} = (0, 0.1, 0.3, 0.35, 0.4, 0.5, 0.6, 0.65, 0.8, 0.85, 1)$. For example, $z_{lr} = 0.4 = \frac{0.2+0.6}{2}$ is obtained for $l=2, r=2$, as shown in the figure. The points marked by d_j are the only deviations that black needs to consider, where d_1, d_3 correspond to the right edge of s_2 intervals (minus δ), and d_2 corresponds to the left edge of s_1 interval (marked with *).

Algorithm 3 FIND-PSPE-FAST(G).

```

Initialize( $\bar{F}, \bar{T}^*, \bar{A}^1, \bar{A}^2, \bar{I}$ );
// See description of Initialize() for explanation of variable roles.
for every leaf  $t$ : do
    // In the leaves there is only one budget interval, which is all the range  $[0, 1]$ .
     $F_t \leftarrow (0)$ ; // An array of length one
     $T_t^* \leftarrow (t)$ ;
// Traverse nodes bottom up:
for every node  $s$  in post-order do
    Let  $s_l, s_r$  be the two children of  $s$ ;
    For every  $a_l = 1, \dots, |F_{s_l}| + 1$  and  $a_r = 1, \dots, |F_{s_r}| + 1$ , set  $z_{lr} = \frac{F_{s_l}(a_l) + F_{s_r}(a_r)}{2}$ ;
    Sort  $\{z_{lr}\}_{lr}$  in increasing order;
    //  $z_{lr}$  are the only "critical points" of the budget in state  $s$ . The equilibrium bids for any point in the
    // interval  $[z_{lr}, z_{lr+})$  reach the same budget interval in the next state (either  $s_l$  or  $s_r$ ).
    Initialize global empty lists  $F, T, A_1, A_2$ ;
    for every  $l, r$  (in increasing order of  $z_{lr}$ ) do
        Set  $B_1 = z_{lr}$ ; // Budget of player 1
        // Compute the Bottom equilibrium of the current step for budget  $B_1$  (and thus for all its interval)
         $((s_1, a_1), (s_2, a_2), t) \leftarrow \text{Ascending-Auction}(G, s, B_1)$ ;
        Append( $F, B_1$ ); Append( $T, t$ );
        Append( $A_1, (s_1, a_1)$ ); Append( $A_2, (s_2, a_2)$ )
     $U \leftarrow T(1)$ ;
    for  $\ell = 2, 3, \dots, |T|$  do
        if  $T(\ell) \neq T(\ell - 1)$  then
            Append( $U, \ell$ );
    // Write down the equilibrium strategies in  $s$ :
     $F_s \leftarrow \text{Filter}(F, U)$ ;  $T_s^* \leftarrow \text{Filter}(T, U)$ ;
     $A_s^1 \leftarrow \text{Filter}(A_1, U)$ ;  $A_s^2 \leftarrow \text{Filter}(A_2, U)$ ;
return  $\gamma \leftarrow (\bar{F}, \bar{T}^*, \bar{A}^1, \bar{A}^2)$ ;

```

Function Initialize($\bar{F}, \bar{T}^*, \bar{A}^1, \bar{A}^2$).

```

for every node  $s \in S$ : do
    Initialize global empty lists  $F_s, T_s^*, A_s^1, A_s^2$ ;
    // A variable  $\bar{X}$  is a list with one entry per state, i.e.  $\bar{X} = (X_s)_{s \in S}$ .
    // Intuitively,  $F_s$  contains all budget cutoff points.  $F_s(1) = 0$  always. For every  $a \in \{1, 2, \dots, |F_s|\}$ , all budgets
    //  $B_1 \in [F_s(a), F_s(a + 1))$  are equivalent, where  $F_s(a') = 1$  for all  $a' > |F_s|$ . We use  $f_s: [0, 1] \rightarrow \{1, 2, \dots, |F_s|\}$  as the
    // inverse of  $F_s$ , where  $f_s(B_1)$  is the index of the budget interval in  $s$  that contains  $B_1$ . Formally,
    //  $f_s(B_1) = \max\{a: B_1 \geq F_s(a)\}$ .
    //  $T_s^*(a) \in T$  is the outcome that will be reached from state  $(s, F_s(a))$ .
    //  $A_s^i$  is the strategy of  $i$ , where  $A_s^i(a) = \{s', a'\}$ .

```

7.1. Compiling small DAGs from succinct valuations

Suppose we are given a sequential scrip bargaining, with some succinct representation of the value functions v_i . While we can apply Proposition 14 to construct an equivalent bidding game, the naïve construction would yield a complete binary tree of height k . In particular, the number of nodes $|S|$ would be exponential in k . In the worst case, there is not much we

Function Ascending-Auction(G, s, B_1).

```

 $b_1, b_2 \leftarrow 0$  ; // Initialize bids
 $a_l \leftarrow f_{s_l}(B_1)$ ;
 $a_r \leftarrow f_{s_r}(B_1)$ ;
 $s_1 \leftarrow \operatorname{argmax}_{j \in \{l, r\}} u_1(T_{s_j}^*(a_j))$ ; // Child preferred by white
 $s_2 \leftarrow \operatorname{argmax}_{j \in \{l, r\}} u_2(T_{s_j}^*(a_j))$ ; // Child preferred by black
if  $s_1 = s_2$  then
    // White wins with bid 0, and plays  $s_1$ . Reaches budget interval indexed by  $a_1$  (either  $a_l$  or  $a_r$ ).
     $a_1 \leftarrow f_{s_1}(B_1)$ ;
     $t_1 \leftarrow T_{s_1}^*(a_1)$ ;
    return  $\langle (s_1, a_1), (s_1, a_1), t_1 \rangle$ ;
repeat
    // Compute outcome under current bids  $b_1, b_2$ :
     $a_1 \leftarrow f_{s_1}(B_1 - b_1)$ ;  $t_1 \leftarrow T_{s_1}^*(a_1)$ ; // Current outcome
     $a_2 \leftarrow f_{s_2}(B_1 + b_2)$ ;
     $b'_2 \leftarrow b_1 + \delta$ ; // Try a deviation of black
     $a'_2 \leftarrow f_{s_2}(B_1 + b'_2)$ ;  $t'_2 \leftarrow T_{s_2}^*(a'_2)$ ; // Outcome after deviation
    if  $b'_2 > 1 - B_1$  or  $t'_2 \leq t_1$  then
        // Black cannot or would not raise. In this budget interval (at least  $B_1 = z_{lr}$  and below  $z_{lr+}$ ), white
        // wins. The bid  $b_1$  is set so that  $B_1 - b_1$  falls in the budget interval indexed by  $a_1$  in  $s_1$ .
        return  $\langle (s_1, a_1), (s_2, a_2), t_1 \rangle$ ;
    // A higher bid of black with the same outcome as  $b'_2$ :
     $b''_1 \leftarrow \min\{B_1 - F_{s_1}(a_1 - 1), F_{s_2}(a'_2 + 1) - B_1 - \delta\}$ ;
     $b'_1 \leftarrow b'_2$ ; // Try a deviation of white
     $a'_1 \leftarrow f_{s_1}(B_1 - b'_1)$ ;
     $t'_1 \leftarrow T_{s_1}^*(a'_1)$ ;
    if  $t''_1 \leq t'_2$  then
        // White would not raise (note that  $b'_1 \leq B_1$  by definition of  $b'_2$  so white can always raise). In this
        // budget interval, black wins and plays  $s_2$ . Reaches budget interval indexed by  $a'_2$ .
        return  $\langle (s_1, a_1), (s_2, a'_2), t'_2 \rangle$ ;
     $b_1 \leftarrow b'_1$ ;  $b_2 \leftarrow b'_2$ ;
until;
```

Function Filter(L, U).

```

Initialize  $L^0$  as an empty list;
for  $j = 1, \dots, |U|$  do
     $\ell \leftarrow U(j)$ ;
    Append( $L^0, L(\ell)$ );
return  $L^0$ 
```

can do. For example, if buyers assign distinct utilities to exponentially many bundles of items, then any bidding game must contain this number of terminal states. Also, if $|T|$ is fixed, but merely computing an optimal partition of K is NP-complete, then it must also be NP-hard to construct an equivalent succinct bidding game (if such a game even exists). This is since constructing the tree would in particular require us to efficiently generate all $|T|$ outcomes, and one them is the optimal partition (we just check all of them).

Additive valuations. We first solve a simple case, where valuations are additive, i.e., $v_i(S) = \sum_{j \in S} v_i(j)$. Suppose that we sold the first 4 items out of k , whose values to agent 1 are 2, 2, 4, and 3. Then the agent does not care if she currently has the first two items or just the third. If the other agent is also indifferent between these partial partitions, then we can represent them with a single state.

We can generalize this observation using a dynamic programming technique, similar to the one used for Knapsack problems (Bellman and Dreyfus, 1962).

Proposition 19. *Given an additive SSB F with integer valuations, there is an equivalent binary bidding game with $|S| \leq (k + 1) \cdot M$ states, where $M = (v_1(K) + 1) \times (v_2(K) + 1)$. Further, such a bidding game can be constructed efficiently.*

Proof. We assign a state s_{j, m_1, m_2} for every $j \in \{0, \dots, k\}$ and every pair $m_1 \in \{0, 1, \dots, v_1(K)\}, m_2 \in \{0, 1, \dots, v_2(K)\}$. The initial state is $s_{0,0,0}$. For every $j = 1, \dots, k$, we go over all states in level $j - 1$. We connect each such state s_{j-1, m_1, m_2} with two children: $s_{j, m_1 + v_1(\tau(j)), m_2}$, and $s_{j, m_1, m_2 + v_2(\tau(j))}$. That is, either agent 1 or agent 2 gets the j 'th item (according to order τ). We then add all new children to level j .

Note that any state at level j corresponds to a partition of the first j items. Finally, we identify the terminals T with the last level k (full partitions), and assign $u_i(s_{k, m_1, m_2}) = m_i$. \square

Using an implementation of Algorithm 3 in Matlab, we computed the PSPE prices for an SSB with k identical items. When the initial budget is the same (0.5 for each player), we witness a clear pattern *increasing prices* over time. This pattern breaks when we start from different initial budget allocations.

Weight-based games. A similar technique can be applied in games where the valuations are not additive, but still base on some additive notion of weight. We say that v_i is *weight-based* if every item has a fixed weight w_j , and there is a function $f_i : \mathbb{N} \rightarrow \mathbb{R}$, s.t. $v_i(S) = f_i(\sum_{j \in S} w_j)$. Note that without further constraints any function is weight-based, as we can make sure that each bundle has a different total weight. However if weights are bounded then we get a succinct representation for a subclass of functions.

This idea can be further generalized. We say that v_i is *multi-weight-based* if every item has a vector of fixed weights (w_j^1, \dots, w_j^q) , and there is a function $f_i : \mathbb{N} \rightarrow \mathbb{R}$, s.t. $v_i(S) = f_i(\sum_{j \in S} w_j^1, \dots, \sum_{j \in S} w_j^q)$. As a concrete example, suppose that the items are computing machines, each with some properties like storage, memory, bandwidth, etc. The value of a set of machines to a client depends only on the total storage, total memory, and total bandwidth, regardless of how these resources are allocated across among the machines.

Note that additive SSBs are a special case of multi-weight-based SSBs, where there are two dimensions. $f_1(S) = \sum_{j \in S} w_j^1, f_2(S) = \sum_{j \in S} w_j^2$. Note however that we allow f_i to be a completely arbitrary function, and it does not even need to be monotone.

Proposition 20. *Given a multi-weight-based SSB F , there is an equivalent bidding game with $|S| \leq (k + 1) \cdot M^2$ states, where $M = (w^1(K) + 1) \times \dots \times (w^q(K) + 1)$. Further, such a bidding game can be constructed efficiently.*

Proof. Construction is very similar to Proposition 19. We identify each state at level j with a tuple $\langle j, m_1^1, m_2^1, \dots, m_1^q, m_2^q \rangle$, where m_i^r now tracks the total dimension- r weight of the bundle owned by agent i . We link each state to its two children by adding all respective weights of item $\tau(j)$ to the bundle owned by one of the agents.

Finally, in every terminal $t = s_{\langle k, m_1^1, m_2^1, \dots, m_1^q, m_2^q \rangle}$ (state that corresponds to a full allocation), we set $u_1(t) = f_1(m_1^1, \dots, m_1^q)$, $u_2(t) = f_2(m_2^1, \dots, m_2^q)$. \square

As a corollary from Propositions 18 and 20, a PSPE in multi-weight-based and in additive SSBs can be computed efficiently. That is, in time that is polynomial in the number of items and in the total weight/value.

Voting with single peaked preferences. If voters have a strict preference order over all 2^k possible outcomes of the game, there cannot be a succinct representation, since we will need 2^k terminals. However consider the following scenario. Each voter as an “ideal” point $t_i^* \in \{0, 1\}^k$. The utility of any other outcome only depends on the Manhattan distance from t_i^* , i.e. on the number of coordinates in which t, t_i^* differ. We may also assign a different weight to each coordinate, so that $u_i(t) = -\sum_{j \leq k} w_{ij} |t_i^*(j) - t(j)|$.

This scenario has a simple mapping to an SSB of additive items: The value to i of every coordinate on which t_1^*, t_2^* differ is w_{ij} . If bidder i “buys” coordinate j she may set it to $t_i^*(j)$. The value of every coordinate j on which $t_1^*(j) = t_2^*(j)$ can be arbitrary. The bidders will not compete on these coordinates, as it is Pareto-optimal to set $t(j) = t_1^*(j)$.

Pathfinding. The items in a pathfinding problem over a graph, are pairs $s = \langle location, time \rangle$. In a naïve representation, all pairs would be allocated, and the value of a bundle of pairs to an agent, would be the optimal legal path from source to target that this bundle contains. If our game states would each represent a subset of such pairs, we would quickly reach an exponential blowup. We modify the auction so that a state $\langle location_1, location_2, time \rangle$ means that each agent i is in $location_i$ at $time$, after reaching there from the start. Thus we do not auction states (tuples) in an arbitrary order, but rather construct an alternative bidding game. We present an outline of the construction.

We start from $s_0 = \langle start_1, start_2, 0 \rangle$, and link each state $s = \langle location_1, location_2, time \rangle$ to two other states: the first is just a copy of s , and the second is one of the legal moves from s , for example $\langle location'_1, location_2, time + 1 \rangle$, where $location'_i$ is either a neighbor of $location_i$ or $location_i$ itself, and it does not hold that $location'_1 = location_2$. We keep enough copies of s to exhaust all legal moves (at most the number of locations, squared). The terminal states are those where $time = MaxTime$, or both agents reached their targets. After the first agent reaches her target, we record her time. The negative utility to each agent in a terminal is her time.

Thus, in total we have at most $MaxTime^2 \times |locations|^4$ states including all copies.

8. Discussion and related work

We presented a simple and intuitive mechanism—sequential scrip bargaining—that implements the full range of Pareto-efficient allocations of items to two agents with arbitrary valuation functions. Further, our mechanism can be applied to efficiently and fairly settle other combinatorial bargaining problems that involve two parties. The pure subgame-perfect equilibrium can be computed in time that is polynomial and the size of the full game tree/DAG, and for many classes of valuations, polynomial in the size of the succinct representation.

8.1. Related work

Richman games. Lazarus et al. (1996, 1999) were the first to systematically analyze bidding variations of zero-sum games. They coined the term “Richman games” in honor of David Ross Richman, the original inventor, and considered games with an infinitely divisible unit of budget. A Richman game is a directed graph (possibly with cycles) with two terminal nodes (say, black and white), and a full play is a path starting from some node and ending in a terminal node. The goal of each player is to end the game in her own terminal. The main focus of Lazarus et al. was on the following question: “in every node, what is the minimal fraction of the budget that will guarantee a victory for the white player?” The *Richman function* assigns a unique value to every node $R(s)$, which is the average of the Richman values of its lowest and highest neighbors (normalizing the values of white and black terminals to 0 and 1, respectively). Lazarus et al. show that the Richman function exists on every graph and is also unique if the graph is finite. It turns out that $R(s)$ marks the critical budget: if $B_1 > R(s)$ then white can force a victory when starting from node s ; if $B_1 < R(s)$ then black can force a victory. Moreover, $R(s)$ is also the exact probability that black wins in a game where instead of bidding, the player in each turn is selected by a fair coin toss (see Peres et al., 2007 for more details on random-turn games).

Lazarus et al. further study other variations of bidding games by applying different auction rules. In particular, in the *Poorman game* the highest bidder pays the bank rather than to the other player, so the total budget shrinks with every step.

Develin and Payne (2010) extended the theory of zero-sum bidding games in several important aspects. In most recreational games, such as Tic-Tac-Toe, Chess, and Checkers, white and black can perform different actions when in the same board state s , thus games are not necessarily impartial. Develin and Payne extended some of the previous results of Lazarus et al. to partial games such as Chess and Tic-Tac-Toe. In addition, they considered discrete bids, and the implications of various tie-breaking schemes. To demonstrate their approach, Develin and Payne showed a complete solution of bidding Tic-Tac-Toe for every possible initial budget. Interestingly, in the continuous case the Richman value of the initial state in Tic-Tac-Toe is $133/256 \cong 0.519$, which means that an advantage of $\sim 4\%$ in the initial budget is sufficient to guarantee a victory. In a related work, Payne and Robeva (2008) developed a near-optimal algorithm for playing bidding Hex. In a very recent paper, Larsson and Wästlund (2017) explore some intricacies of partiality, and in particular show examples of Chess endgames where no value exists.

Negotiations and bargaining. In typical bargaining problems there is a *conflict deal* that is implemented if agents cannot reach agreement (Nash, 1950), or utility depends on time until agreement (Rubinstein, 1982; Fatima et al., 2004). In our mechanism there is no conflict deal or discounting and thus it is hard to compare the protocols directly.

Yet, our Bottom Equilibrium can be intuitively thought of as a combinatorial version of the *minimal sufficient concession* principle, that is often employed by agents in the monotonic concession protocol (Zlotkin and Rosenschein, 1996): at every node the winner makes the minimal concession (in terms of future bargaining power) that allows her to win the round.

Sequential combinatorial auctions. “Standard” sequential auctions (with quasi-linear utilities rather than budgets) have been studied by several researchers (Gale and Stegeman, 2001; Rodriguez, 2009; Leme et al., 2012). These papers focused on particular classes of value functions (unit-demand, submodular, etc.) and generally demonstrated that while pure equilibria exist, they may be substantially inefficient for some of these classes. Inefficient outcomes occur even in 2-buyer, complete information auctions (Bae et al., 2007).

All-pay rules and Colonel Blotto. Powell (2009) models a particular sequential game between an attacker and a defender (a sequential Colonel Blotto game). This game is essentially a specific general-sum bidding game over a degenerated tree (a path), where the utility in every match—the success probability—is determined by the invested resources of each party. Our model does not cover this particular game due to several differences. First, in our model the utility is only determined in the leaves rather than accumulated over the entire path;¹¹ More importantly, the utility in every match depends exactly on the “bids” (investments) of both parties, and both parties discard their bids (an “all-pay” auction).

Another all-pay variation of bidding games was studied in Menz et al. (2015), who developed a fast algorithm for computing the optimal strategy. Note that the optimal strategy under all-pay rules must be mixed, i.e. use probabilistic bids.

Sequential auctions with budgets. Lastly and closest to our work, Huang et al. (2012) study particular bidding game, in which two agents use an initial budget to bid over identical items that are sold sequentially. This is essentially a sequential scrip auction mechanism, yet their model is not strictly a special case of ours, as they add some additional refinements.¹² While we believe that the models are very close, the focus on a particular value function (which makes the game an almost-zero-sum game) allows Huang et al. to provide an accurate characterization of the allocations under PSPE as a function of the budgets. Our results are qualitative in nature, but apply to arbitrary utility functions.

¹¹ This would be a minor difference if the number of outcomes was finite.

¹² In Huang et al. (2012) players pay to the bank rather to one another (the Poorman variation), and also try to minimize spent budget as a secondary goal.

Huang et al. prove that item prices *decrease* over time. Curiously, we observe (empirically) that prices *increase*.¹³ One interpretation is that the result of Huang et al. is an artifact of the auction rule, as in the Poorman version the total amount of money is also decreasing in each round.

8.2. Variations

By changing the auction rules, different types of bidding games arise. For example, an all-pay auction may better describe various real world scenarios like security games, sport matches, and R&D competitions. However these games typically do not have a pure equilibrium point, and are more difficult to analyze (see Related work section for specific games that have been studied). Using a second-price rather than a first-price auction to determine the winner should not have a large effect, since in equilibrium the bids will always be very close.

An important question is whether our results still hold when the “Poorman game” version is played, i.e. when the highest bidder pays the bank rather than the other player. Since in the continuous case the strategy at a given node only depends on the fraction B_1/B_2 , there is a one-to-one mapping between strategies in the continuous Richman and Poorman variations of a game. We would thus expect our main results (e.g. Pareto-optimality of the Bottom Equilibrium) to hold, but this would require a separate proof. Also computational properties may be far less convenient. The Poorman version has one important advantage over the Richman version—the game definition naturally extends to any number of players. It is an open question whether a Pareto-optimal PSPE exists in a binary game for more than two players, either when playing together, or when connected in a network of bilateral bargaining connections (Chakraborty et al., 2009). We believe that the answer to this question is negative, since there are simple Poorman games with three players that are non-monotone.

8.3. Future work

Other than studying the variations mentioned above, there are several questions that remain open. One natural direction is how changing the order in which items are offered (or the agenda in sequential voting) affects the induced equilibrium. It is known that in general bargaining problems the agenda can strongly affect the outcome (Fershtman, 1990; Inderst, 2000). While the agenda may also affect the outcome in our model, it is possible that there are classes of valuations where all orders yield a similar outcome in PSPE. For example, from Proposition 16 the order of issues in zero-sum bargaining games does not matter.

More important questions arise when non-determinism enters. For example, we conjecture that our results (both existence and optimality) also hold under random tie-breaking, and that such a randomized mechanism in fact leads to better fairness guarantees. On the other hand random tie-breaking can exponentially increase the number of possible outcomes and thus efficient computation may not be possible.¹⁴ In addition, agents' valuations themselves may be private and sampled from some distribution, and we are interested whether similar results regarding efficiency and monotonicity apply in games without complete information.

Finally, experiments with bargaining games in the spirit of Weg et al. (1990); De Clippel et al. (2012) may reveal whether the game-theoretic analysis is a good proxy to the actual behavior, and whether our mechanism can improve welfare and satisfaction in practice.

Acknowledgments

The authors thank David Parkes and Urban Larsson for valuable feedback on the manuscript.

References

- Albert, Michael, Nowakowski, Richard, Wolfe, David, 2007. *Lessons in Play: An Introduction to Combinatorial Game Theory*. A K Peters Ltd.
- Anbarci, Nejat, 2006. Finite alternating-move arbitration schemes and the equal area solution. *Theory Decis.* 61 (1), 21–50.
- Bae, Junjick, Beigman, Eyal, Berry, Randall, Honig, Michael L., Vohra, Rakesh, 2007. Efficiency bounds for sequential resource allocation auctions. In: CDC'07, pp. 765–770.
- Bellman, R.E., Dreyfus, S.E., 1962. *Applied Dynamic Programming*. Princeton University Press, Princeton, NJ.
- Bhat, Jay, Payne, Sam, 2009. Bidding chess. *Math. Intell.* 31 (4), 37–39.
- Chakraborty, Tanmoy, Kearns, Michael, Khanna, Sanjeev, 2009. Network bargaining: algorithms and structural results. In: ACM-EC'09, pp. 159–168.
- Conitzer, Vincent, Sandholm, Tuomas, 2004. Communication complexity as a lower bound for learning in games. In: *Proceedings of the Twenty-First International Conference on Machine Learning*. ACM, p. 24.
- De Clippel, Geoffroy, Eliaz, Kfir, Knight, Brian, 2012. On the Selection of Arbitrators. Technical report. Working Paper. Department of Economics Brown University.
- Develin, Mike, Payne, Sam, 2010. Discrete bidding games. *Electron. J. Comb.* 17 (1), R85.
- Fatima, Shaheen S., Wooldridge, Michael, Jennings, Nicholas R., 2004. An agenda-based framework for multi-issue negotiation. *Artif. Intell.* 152 (1), 1–45.

¹³ Using a fast implementation of our algorithm in Matlab, we computed the PSPE prices for an SSB with k identical items. When the initial budget is the same (0.5 for each player), we witness a clear pattern of *increasing prices* over time. This pattern breaks when we start from unequal initial budget allocations.

¹⁴ We thank Michael Simkin for verifying this assertion.

- Fershtman, Chaim, 1990. The importance of the agenda in bargaining. *Games Econ. Behav.* 2 (3), 224–238.
- Gale, Ian L., Stegeman, Mark, 2001. Sequential auctions of endogenously valued objects. *Games Econ. Behav.* 36 (1), 74–103.
- Hart, Sergiu, Mansour, Yishay, 2010. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games Econ. Behav.* 69 (1), 107–126.
- Huang, Zhiyi, Devanur, Nikhil R., Malec, David, 2012. Sequential auctions of identical items with budget-constrained bidders. arXiv preprint arXiv:1209.1698.
- Inderst, Roman, 2000. Multi-issue bargaining with endogenous agenda. *Games Econ. Behav.* 30 (1), 64–82.
- Kalai, Ehud, Smorodinsky, Meir, 1975. Other solutions to Nash's bargaining problem. *Econometrica*, 513–518.
- Kalai, Ehud, 1977. Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica*, 1623–1630.
- Lacy, Dean, Niou, Emerson M.S., 2000. A problem with referendums. *J. Polit.* 12 (1), 5–31.
- Lang, Jérôme, Xia, Lirong, 2009. Sequential composition of voting rules in multi-issue domains. *Math. Soc. Sci.* 57 (3), 304–324.
- Larsson, Urban, Wästlund, Johan, 2017. Endgames in bidding chess. arXiv preprint arXiv:1703.01624.
- Lazarus, Andrew J., Loeb, Daniel E., Propp, James G., Ullman, Daniel, 1996. Richman games. In: *Games of no Chance*. In: MSRI Publications, vol. 29, pp. 439–449.
- Lazarus, Andrew J., Loeb, Daniel E., Propp, James G., Stromquist, Walter R., Ullman, Daniel H., 1999. Combinatorial games under auction play. *Games Econ. Behav.* 27 (2), 229–264.
- Leme, Renato Paes, Syrgkanis, Vasilis, Tardos, Éva, 2012. Sequential auctions and externalities. In: *SODA'12*, pp. 869–886.
- Menz, Michael, Wang, Justin, Xie, Jiyang, 2015. Discrete all-pay bidding games. arXiv preprint arXiv:1504.02799.
- Moulin, Herve, 1984. Implementing the Kalai–Smorodinsky bargaining solution. *J. Econ. Theory* 33 (1), 32–45.
- Nash, John F., 1950. The bargaining problem. *Econometrica*, 155–162.
- Payne, Sam, Robeva, Elina, 2008. Artificial intelligence for bidding hex. arXiv preprint arXiv:0812.3677.
- Peres, Yuval, Schramm, Oded, Sheffield, Scott, Wilson, David B., 2007. Random-turn hex and other selection games. *Am. Math. Mon.* 114 (5), 373–387.
- Powell, Robert, 2009. Sequential, nonzero-sum “Blotto”: allocating defensive resources prior to attack. *Games Econ. Behav.* 67 (2), 611–615.
- Rodriguez, Gustavo E., 2009. Sequential auctions with multi-unit demands. *B.E. J. Theor. Econ.* 9 (1).
- Rosenthal, Robert W., 1981. Games of perfect information, predatory pricing and the chain-store paradox. *J. Econ. Theory* 25 (1), 92–100.
- Rubinstein, Ariel, 1982. Perfect equilibrium in a bargaining model. *Econometrica*, 97–109.
- Sprumont, Yves, 1993. Intermediate preferences and Rawlsian arbitration rules. *Soc. Choice Welf.* 10 (1), 1–15.
- Weg, Eythan, Rapoport, Amnon, Felsenthal, Dan S., 1990. Two-person bargaining behavior in fixed discounting factors games with infinite horizon. *Games Econ. Behav.* 2 (1), 76–95.
- Winter, Eyal, 1997. Negotiations in multi-issue committees. *J. Public Econ.* 65 (3), 323–342.
- Yao, Andrew Chi-Chih, 1979. Some complexity questions related to distributive computing (preliminary report). In: *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*. ACM, pp. 209–213.
- Zlotkin, Gilad, Rosenschein, Jeffrey S., 1996. Mechanism design for automated negotiation, and its application to task oriented domains. *Artif. Intell.* 86 (2), 195–244.