# EFFICIENCY IN DYNAMIC GAMES WITH SEQUENTIAL TRANSFERS

ANDREW KOSENKO
Department of Economics and Finance, Marist College

NATHANIEL NELIGH
Department of Economics, University of Tennessee, Knoxville

Given any perfect information sequential game, if players can offer action-conditional utility transfers to the moving player sequentially at every history ("bid" for actions), every Markov perfect equilibrium results in a utilitarian-efficient outcome, maximizing the sum of all players' utilities. In equilibrium players bid "pivotally," offering just enough to change which action is the most valuable, taking into account the bids and utilities of other players. Payoff distributions are generally non-unique (we provide a condition for uniqueness) and exhibit weak first-bidder advantage.

KEYWORDS: transferable utility, dynamic games, utilitarian efficiency, perfect information, common agency.

## 1. INTRODUCTION

In this paper we show that if players can commit to one-step-ahead action-contingent transfers and utility is transferable, every equilibrium of an arbitrary sequential game will result in a utilitarian-efficient outcome. More precisely, if at each node, non-moving players are able to sequentially offer action-contingent transfers ("bids") to the moving player, any Markov-perfect equilibrium outcome of this bidding-augmented game will always correspond to the outcome that maximizes the sum of the players' payoffs in the original game. We find that in equilibrium players employ a *sequential-pivotal bidding* strategy, which

---

Andrew Kosenko: kosenko.andrew@gmail.com

Nathaniel Neligh: nneligh@utk.edu

We thank Marina Halac, Navin Kartik, Juuso Toikka, Joseph Stiglitz, three anonymous referees, and seminar audiences at the University of Tennessee, Knoxville, Stony Brook Game Theory Festival/International Conference, and the 2024 North American Summer Meeting of the Econometric Society for valuable comments. We also thank Prajit Dutta, Andrea Prat, and Paolo Siconolfi for very helpful discussions.

may be of independent interest: In our construction, players bid just enough to change the mover's action, aware of the players who have bid already and the players who will bid later. Rather than relying on fixed-point arguments, we provide a constructive approach that derives the necessary bids from the value function.

Using this approach, we show that bidding on actions is sufficient to guarantee utilitarian efficiency in any *arbitrary* sequential game of perfect information (of finite or infinite horizon, with or without a repeated game structure) with a finite but arbitrary number of players, assuming perfect information, continuity at infinity of utilities, and transferable utility. This result can be viewed as a formalization of the Coase conjecture (Coase (1960)) for sequential games with perfect information, or a version of the first welfare theorem (in that we show that an efficient outcome can be achieved through simple "decentralized" transactions). The difference between our work and the classic conversation around the Coase conjecture ((Medema (2020) provides an updated discussion) is that we use a stronger efficiency concept (utilitarian as opposed to Pareto), and a game theoretic setting.

There are two main antecedents for our results: the work of Dutta and Siconolfi (2019) and Prat and Rustichini (1998). Dutta and Siconolfi (2019) show that in any (finitely or infinitely) repeated sequential two-player game with perfect information and transferable utility, strong utilitarian efficiency (in the sense of maximizing the sum of the players' utilities) will be achieved as long as players can contract sequentially on the next action of the other player. We extend this result in two main ways. First, our result applies to arbitrary sequential games, provided that the difference between the highest and lowest possible eventual total payoffs for an individual, given an action history, must converge to zero as the length of the action history grows to infinity (a standard "continuity at infinity" assumption). A repeated game structure is not required, nor is geometric discounting.

Our second contribution is in extending the original result to any finite number of players. Dutta and Siconolfi (2019) hypothesize that this is possible, but they do not go beyond two players in their paper. While it is relatively straightforward to extend the arguments of Dutta and Siconolfi (2019) to show that an efficient Markov perfect equilibrium (MPE) exists with arbitrary numbers of players, additional nuance arises when characterizing *all* Markov perfect equilibria in the many player case. This complexity arises from the presence of bids for actions which are not realized on the path of play, but do have knock-on effects on both the future bids and the future payoffs of other players. These unrealized transfers mean that, while efficiency is guaranteed, the distribution of the payoffs is often not unique, due to indifferences during bidding.

Prat and Rustichini (1998) study a sequential bids common agency problem where the principals offer conditional transfers to the agent in sequence, and the agent then takes a single action. They find that all equilibria are efficient and that the payoff vector is not, in general, unique. We extend this result by moving from the single action case to arbitrary sequential games and characterizing the conditions for a unique payoff vector.

Bernheim and Whinston (1986a), introduced the original (simultaneous bids) common agency model as a principal-agent model with multiple principals who simultaneously propose action-contingent transfer schemes, and an agent who, upon observing these offers, chooses an action. They show that coalition-proof equilibria (a subset of all equilibria) are efficient. They find that efficient MPE always exist in this setting. Bergemann and Välimäki (2003) extend the (one-shot) Bernheim and Whinston (1986a) to arbitrary sequential games. Their results show that, again, only a subset of MPE are efficient.

In light of this history, our result can be viewed conceptually as extending the work of Prat and Rustichini (1998) in a conceptually similar way to how Bergemann and Välimäki (2003) extend the work of Bernheim and Whinston (1986a) (in the sense that Prat and Rustichini's (1998) model is a special case of ours, with a single action). One should note, however, that the methods we use in the extension and results we come to are very different from those of Bergemann and Välimäki (2003). For example, in the sequential bid space, we prove that all MPEs are efficient rather than the existence of an efficient MPE. In addition, our result allows players to "take turns" being the agent in the common agency game rather than having the same player be the agent each time, and we allow the available actions for the agent change arbitrarily with the history (in Bergemann and Välimäki (2003) available actions change based on a Markovian state of the world with finite realizations). In addition, our uniqueness result is completely different both in what exactly is unique (payoff vectors in our work, as opposed to strategy profiles, in existing work) and to what this uniqueness applies (all MPE outcomes in our work, as opposed to "truthful" or coalition-proof equilibria, in existing work).

In summary, Dutta and Siconolfi (2019) show the efficiency result for two players in a dynamic game, while Prat and Rustichini (1998) show the result for many players in a one-shot game. Our work generalizes both approaches by guaranteeing efficiency in a multiplayer sequential setting. In addition, our model also delivers an additional result: under certain conditions, the *outcome* in all of these equilibria is unique (up to indifferences), but these conditions are rare and fragile. Finally, our approach is different from the approaches explored earlier: instead of relying on (relatively more opaque) fixed point theorems, we explicitly derive the necessary transfers from the value function.

The extension to non-repeated infinite games is of major importance in its own right but it should also be noted that combining games with more than two players (extending Dutta and Siconolfi (2019)) and games with multiple sequential actions (extending Prat and Rustichini (1998)) introduces a complexity to the bidding that is not present in either alone. In particular, the combination creates potential sunspots issues even in finite games, as we demonstrate with an example in section 4.3. For this reason we focus on strategies that are Markovian with respect to past bids. The model of Prat and Rustichini (1998) and finite-length games in Dutta and Siconolfi (2019) only need subgame perfection to guarantee efficiency.

In the process of extending the result to many players we find an interesting effect. The structure of the backward induction reasoning creates a type of naturally occurring sequential pivot effect where players who want to induce an action must bid for that action based on the amount that the change in the implemented action alters the valuations of later bidding players and realized bids.[1] The sequential pivot in our proof resembles the Vickrey-Clarke-Groves (VCG) mechanism, but rather than being operated by a principal, it arises endogenously from the optimizing behavior of sequentially bidding players. The pivot also has sequential features that are absent from the VCG mechanism—for instance, players react to the *bids* of earlier bidders but to the *valuations* of later bidders. This asymmetry means that earlier bidders are able to shift the burden of changing the action to later bidders, leading to weak early bidder advantage.

The proof of our main result (Theorem 1) proceeds in three Lemmata. Lemma 2 shows by induction that, during a bidding phase preceding an action, players will bid pivotally relative to their value functions. Lemma 3 shows that in every action period, pivotal bidding will lead to the action that maximizes total value across all players, as result we call "one-

[1] This type of strategy implicitly occurs in Prat and Rustichini (1998) and one case of it is constructed explicitly in their "thrifty equilibrium." In contrast, we construct it explicitly for all equilibria.

step-ahead optimality." We show this by considering the first bidder: their implemented pivotal action maximizes the sum of the valuations. This emerges because the first bidder does not face any existing bids, so they are only reacting to their own valuation and the valuations of the other players. Finally, Lemma 4 shows that continuity at infinity and one-step-ahead optimality (Lemma 3) guarantee the efficiency of the outcome of the whole game.

Our result illustrates how a simple modification to a strategic situation—that is, the introduction of bids for actions—may dramatically improve outcomes. More broadly, the "bidding for efficiency" approach elucidates the limits of how efficiency in games may be reached, by using contracts that are "simple" in the following senses: $i)$ bilateral; $ii)$ one-period-ahead; $iii)$ decentralized and uncoordinated; and $iv)$ explicit (no "black box"), at least relative to directly contracting on outcomes. Such bilateral payments may arise anytime transfers are possible and actions are contractible (i.e. under the minimal requirement that bilateral contracts on actions are enforceable).

We present three examples illustrating our results. In section 2 we present a three-player "centipede" example to build intuition. In section 3 we go through an infinite game with continuous at infinity payoffs that does not have a repeated game structure or geometrically discounted payoffs—an "infinite centipede," that is covered by our theory. Neither of these examples is covered by existing approaches. In section 4.3, we show by a counterexample how efficiency may fail without Markovian bids; this example also illustrates how bids for actions which are not taken — a feature that is absent from two-player settings — matter, and complicate the construction.

We also provide several results on the properties of this mechanism: there is weak early bidder advantage, and payoffs are generically non-unique. We provide necessary and sufficient conditions, checkable using an algorithm, for uniqueness of the payoff vector. Lastly, we consider what happens when utility is imperfectly transferable in the sense that beneficial "money" can be exchanged, but the value of money is potentially non-homogenous and non-linear. We find that efficiency is guaranteed only in the narrowest two-player one-shot case.

The paper is organized as follows: Section 2 presents the notation, definitions, and equilibrium concept. Section 3 discusses an infinite-horizon example in detail. Section 4 states and proves the main result, while section 5 discusses features of the "bidding for actions"

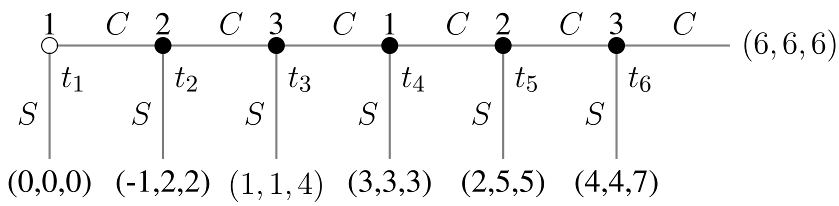


FIGURE 1.—Finite centipede with three players

mechanism—(non)uniqueness of payoffs, weak first-mover advantage, and the Pareto properties of the bidding-augmented game relative to the original game. Section 6 probes the limit of the assumption of transferable utility, and studies a one-period model with imperfectly transferable utility. The literature review is in section 7. Proofs are relegated to the appendix, with the exception of the proof of the fact that players bid pivotally (Lemma 2), which is illustrative and, thus, appears in the main text.

## 2. MODEL

Before diving into definitions, we first illustrate the workings of the bidding mechanism,[2] in a relatively simple, three-player setting of one well-known dynamic game—the centipede (Rosenthal (1981)). In section 3 we present an infinite-horizon example with two players.

Consider the finite centipede with three players and with the payoffs illustrated in figure 1.[3] As usual, all subgame perfect equilibria involve stopping immediately. However, the utilitarian-efficient outcome is to continue at all nodes, which yields $(6, 6, 6)$ for the players. Illustrating our main result, if the players are allowed to bid for actions, all equilibria of the bidding-augmented centipede become ones where all players continue at every opportunity, yielding the utilitarian-efficient outcome.

More explicitly, suppose that nonmoving players bid in numerical order. In one equilibrium, at $t_6$, P2 (a non-moving player) will bid one unit of utility to P3, to incentivize them to play $C$. The payoffs are now $(6, 5, 7)$, inducing P3 to choose $C$. At $t_5$ and earlier no bids

[2]While we use the term "bidding mechanism," this is not a true "mechanism" in the sense of mechanism design.

[3]Neither of the examples we develop are covered by the work of Dutta and Siconolfi (2019) or Prat and Rustichini (1998); one of our examples has three players and the other one has a non-repeated game structure, and both have sequential moves.

will take place since P2 (and the other players) already has enough of an incentive to play $C$. In fact, the bid at $t_6$ is the *only* necessary bid: Here, a single transfer is enough to get to efficiency.

We now formally introduce our setup.

## 2.1. *Notation*

Let $\Gamma = \{N, H, P, A, \pi_i\}$ be a given extensive-form game with perfect information, finite actions at each node, and no chance moves, where

1. The set of players is $N = \{1, 2, \ldots, \overline{N}\}$.
2. We work directly with action histories. Let $h_t = \{a^1, a^2, \ldots, a^t\}$ denote the history of the actions until time $t$. We interpret the number of actions in a history as a "time period." Let $H_t$ denote the set of all histories with $t$ elements, and let $\mathcal{H} = \cup H_t$ denote the set of all possible histories.
3. There is a player function $P : \mathcal{H} \to N$ specifying the player who moves at $h_t$, and we refer to $P(h)$ if the time period is arbitrary or clear from the context.
4. The set of actions for player $i$ after action history $h_t$ is given by a correspondence $A_i(h_t)$, or simply $A(h)$, if the moving player and time period are clear. After an action $a_{t+1}$ we write the evolution of the action history as $h_{t+1} = (h_t, a^{t+1})$.
5. Denote by $Z$ the set of *terminal histories*—that is, either finite histories that have no further successors, or infinite histories. For every terminal history, there is a vector of payoffs $\pi_i : Z \to \mathbb{R}$ for each player if that history is reached; thus, we assign payoffs to all terminal histories *ex-ante*. We assume that $\pi_i(z)$ is uniformly bounded in magnitude. Note that we treat any two infinite histories as the same history when they only diverge after an infinite number of actions.[4] Given continuity at infinity (see upcoming definition 1), this assumption does not impose any restrictions on payoffs.

We also assume that the utility function is continuous at infinity:

[4] One can think of this assumption as imposing a notion of distance on the space of action histories that ensures that the distance between histories that only differ in actions that are "far into the future" converges to 0 as the differences get farther into the future. This assumption is needed because if $(A, A, A, ...) \neq (A, A, A, ..., A, B)$, then history $(A, A, A, ...)$ would have a successor history and would not be terminal.

**Definition** 1—Continuity at infinity**.** A utility function is *continuous at infinity* if, given an $\epsilon > 0$ there exists $t(\epsilon)$ such that for action histories $h_t$ and $h'_t$ that agree up until time $t$, we have

$$\max_{z \in Z(h_t), z' \in Z(h'_t)} |\pi(z) - \pi(z')| < \epsilon \tag{1}$$

where $Z(h)$ is the set of all terminal histories that follow action history $h$. This is a rephrasing of the standard "continuity at infinity" assumption for our setting; the meaning and implications (i.e., that payoffs "far" into the future are not too important) are standard.

### 2.2. *The Bidding-Augmented Game*

Given a dynamic game with perfect information, we augment it such that immediately before each action is taken, each non-moving player (in a sequence) may offer the player moving at that action history a set of non-negative transfers, contingent on the action the moving player takes. We call this the *bidding phase*. We refer to the players in the bidding phase other than the moving player as *bidders*; of course, bidders are also players, but we make the distinction to emphasize where a player is in the process. We let utility be quasi-linear and transferable, so bids are in terms of utility.

For any $\Gamma$, we construct the version with bids as follows: $\Gamma^{BA} = \{N, \hat{\mathcal{H}}, \hat{P}, \hat{A}, \pi_i\}$, which is a *bidding-augmented* game of $\Gamma$. The set of players is the same. The histories of $\Gamma^{BA}$ are elements of the set $\hat{\mathcal{H}}$ and are constructed by taking the histories of $\Gamma$ and adding to each action history a bidding phase that precedes the action phase.

The augmented player function $\hat{P}$ is constructed from $P$ by allowing players to offer action-contingent utility transfers to the mover specified by $P$ when it is their turn to do so. These bids occur in some fixed order during the augmented histories that immediately precede the action. For simplicity of notation, we assume that at each action history $h_t$ players bid in the order $\{1, 2, 3, P(h_t) - 1, P(h_t) + 1, \ldots\}$; this assumption may be generalized to an arbitrary bid order without changing any of the logic in this paper. The order can even be stochastic as long as a bidding player knows who has already bid and who will bid after them.

The augmented action and bid correspondence $\hat{A}$ is constructed similarly, giving the bidders options to bid before the actions. We keep the set of available actions compact by assuming that a player's bids must be weakly less than the difference between their supremal and infimal potential payoffs. This assumption is purely technical and without loss of generality, as no player will ever want to make bids greater than this amount.

We allow players to decline transfers, although because the transfers are non-negative (and only strictly positive transfers that are "large enough" will play a role), we do not explicitly incorporate this choice into the analysis.

*Histories, Strategies, and Equilibrium*

The bid of player $i$ at an action history $h_t$ for action $a_j$ is denoted as $b_i(a_j; h_t, b_{<i})$. This bid is a contingent payment offered by the bidding player, $i$, to the mover, $P(h_t)$, to be paid out if the mover takes action $a_j$. Since Player $i$'s bid depends on the action history and the bids that have already been made (which we denote by $b_{<i}$), we generally drop these arguments and refer to $b_i^t(a_j)$ to improve readability.

The profile of player $i$'s bids for all actions at a history is $b_i^t$. Bidding players may bid positive amounts for multiple actions. Denote by $a_i(h_t) = \{a(h_t)|P(h_t) = i\}$ the action of player $i$ at action history $h_t$ if that player is the mover at that node.

Bids on previous actions are sunk from the point of view of the current mover or bidder. As such, we focus on histories that do not include these older bids.

Formally, the *relevant history* in the bidding-augmented game is a set $r_t = (h_\tau, \{b_i\}_{i \in o^\tau_{t-\tau N}})$, where $t = \tau N + |\{b_i\}_{i \in N \setminus P(h_t)}|$ and $o^\tau_{t-\tau N}$ are the first $t - \tau N$ elements of a re-ordering, $o^\tau$, of $\{1, 2, ..., N\} \setminus P(h_t)$ (specifically it is the order of bidders at history $h_\tau$). The space of all such histories is denoted by $R$.

Note that (with a small abuse of notation) the histories of $\Gamma$ are relevant histories of $\hat{\Gamma}$. This means that we can continue to use the same payoff function, $\pi_i(z)$, in our discussion of the bidding-augmented game. It also means that any function that can be applied to a relevant history can also be applied to a history composed of only actions.

We define the *net realized bid* function for a given player as

$$n_i(h_t) = \begin{cases} \sum_j b_j^t & P(h_t) = i \\ -b_i^t & P(h_t) \neq i \end{cases} \tag{2}$$

If player $i$ is moving at $h_t$, the net realized payoff function gives the total bids that other players have given to player $i$ for the realized actions they took, and if player $i$ is not moving, $n_i(h_t)$ is the total amount $i$ paid out to the mover.

A player's final payoff from a terminal history $z$ of length $t$ is given by

$$\pi_i(z) + \sum_{s \le t} n_i(h_s) \tag{3}$$

where here the $h_s$s are sub-histories of $z$.

**Definition** 2—Strategies**.** A *strategy* for player $i$ in the augmented game is

$$\sigma_i = \sigma_i(r_t) = (b_i(r_t), a_i(r_t)), \forall r_t \in R \tag{4}$$

A strategy specifies a player's bid every time they get to bid and a player's move every time they get to move (as a function of relevant histories), with the convention that player $i$ bids $0$ during their move and a bidder takes a null action during another player's move.

Due to the structure of relevant histories, the strategy $\sigma_i(r_t)$ may depend on the action history, $h_t$, in $r_t$ and on the bids "within" a period $(b_{-i}(h_t))$ but not bids from before the last action. By using relevant histories, we implicitly assume that bids and actions are independent of the (payoff irrelevant) history of the bids made before the previous action. As such, our strategies are Markovian with respect to older bids.

Let $\sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_N\}$ be a strategy profile. A strategy profile $\sigma$ in $\Gamma^{BA}$ generates a distribution over actions and bids, $\gamma(\sigma)$, leading to a distribution over realized bids and terminal histories. Before we define an equilibrium, we formally define the value function for each player. Let

$$V_i(r_t, \sigma) = \mathbb{E}_{\gamma(\sigma)} \left[ \pi_i(z) + \sum_{s \ge \tau} n_i(h_s) | r_t \right] \tag{5}$$

be the *value function for player* $i$ *at time* $t$ *under strategy profile* $\sigma$, with the understanding that the actions $a_i(r_t)$ and the bids $b_j(r_t)$ are determined according to the strategy profile $\sigma$. The *joint value function* is

$$\overline{V}(r_t, \sigma^*) = \sum_{i=1}^{N} V_i(r_t, \sigma^*) \tag{6}$$

Now we can define the equilibrium concept:

**Definition** 3—Equilibrium. A *Markov-perfect-bidding equilibrium* (MPBE) is a strategy profile $\sigma^*$ such that for each player $i$, for each $t$, and for every action or bid $c^*$ that occurs with positive probability, we have

$$c^* \in \arg\max_{c \in \hat{A}(r_t)} V_{\hat{P}(r_t)}((r_t, c), \sigma^*) \tag{7}$$

where $(r_t, c)$ is understood to mean the relevant history $r_t$ followed by (action or bid) $c$.

This type of equilibrium is Markovian with respect to older bids because, as previously mentioned, our definition of strategies uses histories that discard those elements. This prevents any potential sunspots based on older bids, although action-based sunspots are still allowed. The importance of being Markov-perfect with respect to bidding is demonstrated by a counterexample in section 4.3. As a concept, Markov perfect bidding equilibrium is more general than traditional Markov perfection, because it allows for sunspots in non-bid actions and recent bids, but it is still less general than subgame perfect equilibrium.

In addition to the equilibrium, we are also concerned with the efficiency of the outcome.

**Definition** 4—Efficiency. Our notion of efficiency is $\bar{\pi}(z) = \sum_{i=1}^{N} \pi_i(z)$. Call a history $z^* \in Z$ with the property that $\bar{\pi}(z^*) \geq \bar{\pi}(z'), \forall z' \in Z$ a *strongly efficient history* and the outcome of said history a *strongly efficient outcome*.

Such an outcome $z^*$ is the outcome that maximizes the sum of the payoffs of the players—that is, it is the best outcome in the utilitarian sense. We suppose that the utilitarian outcome is desirable and we work under this assumption to illustrate the efficiency result.

## 3. EXAMPLE: AN INFINITE CENTIPEDE

Here we illustrate in detail the workings of the bidding mechanism and the results. Consider an infinite version of the two-player centipede, depicted in figure 2. We modify the payoffs to satisfy several key properties from the finite version. Namely, the payoff from stopping is always greater for the moving player than any possible payoff from continuing; in addition, the sum of the payoffs is increasing as the game proceeds, the sum of the payoffs from continuing forever remains greater than the sum of the payoffs from stopping, and the payoffs from stopping fall (thus preserving the incentive to end the game at first opportunity); finally, payoffs are continuous at infinity. Thus, the fundamental tradeoff present in the centipede is preserved.

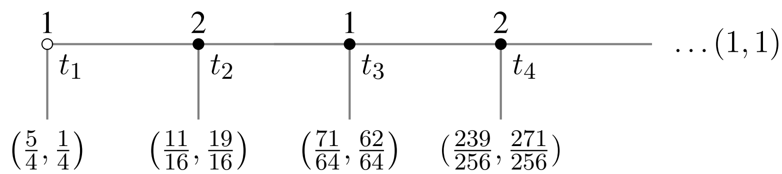


FIGURE 2.—An infinite centipede

In general, the payoff structure for terminal nodes is as follows:

$$(\pi_1(t),\pi_2(t)) = \begin{cases} \left(\left(1-\left(\frac{1}{4}\right)^t\right)+\left(\frac{1}{2}\right)^t\right),\left(\left(1-\left(\frac{1}{4}\right)^t\right)-\left(\frac{1}{2}\right)^t\right) & \text{for } t \text{ odd} \\ \left(\left(1-\left(\frac{1}{4}\right)^t\right)-\left(\frac{1}{2}\right)^t\right),\left(\left(1-\left(\frac{1}{4}\right)^t\right)+\left(\frac{1}{2}\right)^t\right) & \text{for } t \text{ even} \end{cases} \tag{8}$$

We assign payoffs $(1,1)$ to the infinite history $(C,C,C,C...)$. If we allow players to bid for actions, the outcome changes drastically (here all Nash equilibria of the game without bidding also involve stopping immediately). If player 2 transfers $\frac{7}{16}$ at $t_1$ to player 1, player 1 will play $C$ at $t_1$, instead of stopping. Onward, at $t_2$ P1 will bid $\frac{19}{64}$, at $t_3$ P2 will bid $\frac{43}{256}$, and so on.

The general bid to the moving player (only the non-moving player gets to bid) is given by:

$$b_t = 1\frac{1}{2}\left(\frac{1}{2}\right)^t - 1\frac{1}{4}\left(\frac{1}{4}\right)^t \tag{9}$$

With these bids, the play proceeds forever, with each player continuing when they get a chance to move.

Besides illustrating the fact that allowing players to bid will result in the utilitarian-efficient outcome, this example has several notable features. First, the bidding mechanism is manifestly nontrivial—there are an infinite number of on-path transfers that are determined by the payoffs in the underlying game. It can be checked that the value from continuing is $\frac{5}{4}$ for player 1 and $\frac{3}{4}$ for player 2; this value remains constant as the play proceeds. Second, players gain enough utility from the efficient actions to make up for the costs of the bids needed to achieve them. Thus, not only does there exist a sequence of bids that alter the actions to achieve efficiency but this sequence is optimally attainable based on the payoffs from the efficient outcome.

## 4. MAIN RESULT: TRANSFERS IMPLEMENT THE UTILITARIAN OUTCOME

With the preliminaries out of the way, we now present our main result.

THEOREM 1: *The outcome of every MPBE of* $\Gamma^{BA}$ *results in a strongly efficient outcome* $z^*$ *of* $\Gamma$.

In other words, strikingly, allowing for conditional transfers results in utilitarian efficiency in a large class of games. For instance in a sequential perfect information version of the prisoner's dilemma, Theorem 1 shows that with transfers (and without communication) the outcome would be to cooperate. Similarly, the outcome in the centipede with transfers would be to continue for as long as necessary.

We prove this result in several steps. First, we show (in Lemma 2) that players have an incentive to bid enough to implement the efficient action ("pivotally") in equilibrium. Then we show (in Lemma 3) that this style of bidding guarantees one-period-ahead efficiency with respect to the value function. This is efficiency "within" a period. Finally (in Lemma 4) we show that one-period-ahead efficiency along with continuity at infinity guarantees overall efficiency. This is efficiency "across" periods. Taken together, Lemmas 3 and 4 prove Theorem 1.

### 4.1. *First Step: Pivotal Bidding and One-step-ahead Optimality*

In this section we show that non-moving players will have an incentive to bid for the utilitarian-efficient action at every action history.

To that end, fix an action history, $h_t$, and consider the incentives of the bidding players between $h_t$ and $h_{t+1}$. For the purposes of this section, the action components of the history will remain fixed and, as such, we suppress the dependence of the various objects on $h_t$ whenever possible (specifically in the value function).

We also fix $\sigma^{h_{t+1}}$, the strategy profile continuing after the next action (the dependence of the value function implicitly depends on $\sigma^{h_{t+1}}$ but we do not include it in the notation since this is always implicitly true for value functions). This allows us to treat the value of various $h_{t+1}$s as endpoints with defined values.

To state and prove Lemma 2, we first need to define the running total of the bids function, the future bidder value function, and the notion of *pivotal* bidding. Then we show that the players do, in fact, bid pivotally.

Given a fixed set of bids $\{b_k\}_{k=1,2,\ldots,i-1}$ and a strategy profile $\sigma$, we define the *running total function during the bid of player $i$ for action $a$*

$$T_i(a) = V_{P(h_t)}(a) + \sum_{k=1}^{i-1} b_k(a) \tag{10}$$

Note that the mover's (player $P(h_t)$) value function $V_{P(h_t)}$ is included in the running total because the mover's value of the action contributes to their preference for the action similar to the way the bids do. Also, we define the future bidder value function during the bid of player $i$

$$F_i(a) = \sum_{k=i+1}^{N-1} V_k(a) \tag{11}$$

as the sum of the value functions for all future bidding players; note that the running total function includes the value function of the mover and the *bids* of the preceding bidders, while the future bidder value function includes the *value functions* (as opposed to the bids) of the future bidders.

Let $\tilde{a}_i = \tilde{a}_i(T_i, F_i) = \arg\max_a T_i(a) + F_i(a)$ be the *leading action during the move of player $i$*.

**Definition** 5—Action-pivotality**.** Player $i$ is *action-pivotal* if there is an action $a^* \neq \tilde{a}_i$ such that

$$a^* \in \arg\max_a V_i(a) + T_i(a) + F_i(a) \tag{12}$$

Thus, a player is action-pivotal if the leading action without the value of this player is different from the leading action with this player included. If there are multiple possible $a^*$s, the pivotal player picks one arbitrarily; note that a tie breaking assumption at this step is *not* needed.

**Definition** 6—Pivotal bidding**.** Player $i$ *bids pivotally* if they bid

$$b_i(a^*) = T_i(\tilde{a}_i) + F_i(\tilde{a}_i) - (T_i(a^*) + F_i(a^*)) \tag{13}$$

for action $a^*$, if it exists, and

$$b_i(a) < T_i(\tilde{a}_i) + F_j(\tilde{a}_i) - (T_i(a) + F_i(a)) \tag{14}$$

for $a \neq a^*$.

Pivotal bidding plays a key role in our approach. It allows us to explicitly construct the bids that are optimal in equilibrium.

LEMMA 2—Pivotal bidding in equilibrium: *In any MPBE of* $\Gamma^{BA}$ *all players bid pivotally.*

This means that all players bid enough to shift the leading action if their preferences make them pivotal in determining the leading action. They can bid any amount for options that are not the leading actions as long as they do not bid so much that the action and bid are realized and they do not increase the amount they must bid to implement the pivotal action.

Note that under this Lemma multiple bidders may be action-pivotal but it is not possible for multiple players to be pivotal for different actions.

PROOF: We prove this by (backward) induction on the set of bidders, beginning with the last bidder in a bidding phase. Let $N$ be the index of this player.

**Mover:** The mover will pick the action that maximizes

$$V_{P(h_t)}(a) + \sum_{k \neq P(h_t)} b_k(a) \tag{15}$$

**Last Bidder**: Player $N$'s utility only depends on the action they implement and the required bid. Suppose $\tilde{a}_N$ is the leading action before player $N$'s bid. Player $N$ can implement an action $a$ by bidding

$$b_N(a) = T_N(\tilde{a}_N) - T_N(a) \tag{16}$$

So Player $N$'s optimization problem becomes

$$\begin{aligned} &\max_a V_N(a) - \underbrace{T_N(\tilde{a}_N) + T_N(a)}_{=b_N(a)} = \\ &= V_N(a) - V_{P(h_t)}(\tilde{a}_N) - \textstyle\sum_{k=1}^{N-1} b_i(\tilde{\alpha}_N) + V_{P(h_t)}(a) + \textstyle\sum_{k=1}^{N-1} b_i(a) \end{aligned} \tag{17}$$

the solution of which is identically equal to the $a_N^*$ defined by action-pivotality because there are no future bidders and, therefore, no $F_i$ term and because

$$\arg\max_a V_N(a) - V_{P(h_t)}(\tilde{a}_N) - \sum_{k=1}^{N-1} b_i(\tilde{\alpha}_N) + V_{P(h_t)}(a) + \sum_{k=1}^{N-1} b_i(a) \tag{18}$$

$$= \arg\max_a V_N(a) + V_{P(h_t)}(a) + \sum_{k=1}^{N-1} b_i(a)) \tag{19}$$

Player $N$ can freely bid for other actions as long as they are not implemented and do not raise the cost of implementing $a_N^*$. As such, player $N$ will not bid for $\tilde{a}_N$, nor will they bid such that $T_N(a_j) + b_N(a_j) \geq T_N(\tilde{a}_N)$ for any $a_j \neq a_N^*$. Thus, player $N$ will bid pivotally.

**Inductive step**: Given that all future players will bid pivotally (by the induction assumption), we show that player $j = 1, \ldots, N-1$ will also bid pivotally in the sense that all optimal actions coincide with the actions determined by pivotal bidding.

Again, $j$'s utility only depends on the action they implement and the required bid. If (by the inductive hypothesis) all future players bid pivotally, $j$ can implement an action $a$ by offering

$$b_j(a) = T_j(\tilde{a}_j) + F_j(\tilde{a}_j) - T_j(a) - F_j(a) \tag{20}$$

Thus, $j$'s optimization problem becomes

$$\max_a V_j(a) - b_j(a) = \max_a V_j(a) - T_j(\tilde{a}_j) - F_j(\tilde{a}_j) + T_j(a) + F_j(a) \tag{21}$$

which is optimized at $a_j^*$, again, as before, by the definitions of $T$ and $F$. Player $j$ can freely bid for other outcomes as long as they are not implemented and do not raise the cost of implementing $a_j^*$. As such, they will not bid for $\tilde{a}_j$, nor will they bid such that

$T_j(a_j) + F_j(a_j) + b_j(a_j) \geq T_j(\tilde{a}_N) + F_j(\tilde{a}_j)$ for any $a_j \neq a_N^*$. Therefore player $j$ will bid pivotally. *Q.E.D.*

We have thus shown that all players bid pivotally in equilibrium. Given such pivotal bidding, we establish the next result—the fact that in equilibrium individual optimization will implement the action that maximizes the one-step-ahead *joint* value function, a result we call "one-step-ahead optimality."

LEMMA 3—One-step-ahead optimality: *In any MPBE of the bidding-augmented game $\Gamma^{BA}$, the action implemented by the bidders is the action maximizing the joint value function at every step:*

$$\bar{V}(h_t, \sigma^*) = \max_{a \in A(h_t)} \bar{V}(\underbrace{(h_t, a)}_{h_{t+1}}, \sigma^*) \tag{22}$$

For the proof, see the appendix. The proof works by looking at the first bidder and showing that they will implement the action that maximizes the total value function. Lemma 3 shows that in equilibrium pivotal bidding results in bidders (and movers) acting in a way that maximizes the one-step-ahead joint value function. Thus, pivotal bidding keeps the implemented actions on track to implement the utilitarian outcome.

### 4.2. *Second Step: One-Step-Ahead Optimality is Equivalent to Global Optimality*

We finish the proof with the following Lemma:

LEMMA 4—One-step-ahead Optimality is Equivalent to Global Optimality: *Suppose that payoffs are continuous at infinity and equation (22) holds. Then we have*

$$\bar{V}^*(\emptyset, \sigma^*) = \max_z \bar{\pi}(z) \tag{23}$$

For the proof, see the appendix. The proof works by using continuity at infinity and Lemma 3 to bound the value function in a way that converges as $t$ goes to infinity. This Lemma guarantees the efficiency of any outcome of the MPBE and finishes the proof of Theorem 1.

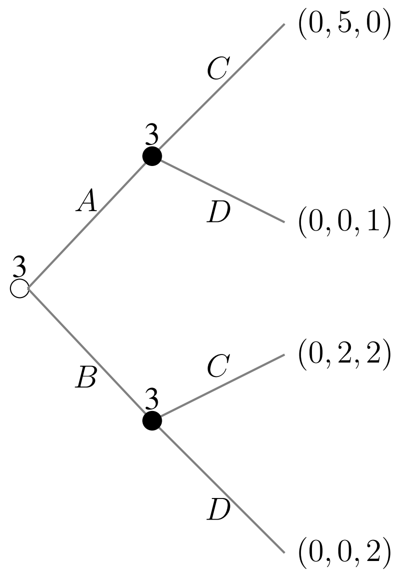


FIGURE 3.—A Counterexample with Non-Markov Strategies

### 4.3. *A Counterexample with Non-Markov Strategies*

Let us now illustrate how things can go wrong if bidding strategies can be conditioned on bids in previous bidding periods - that is, if the bidding strategy is not Markov with respect to past bids. Consider the following three-player game with two periods, where player 3 takes an action at both periods. At each node player 1 bids first, then player 2.

Consider the following bidding strategies:

Player 1:

- If player 2 bids a non-zero amount in the first bidding phase, and action $A$ is chosen in period 1, bid 3 for $D$ in period 2.
- Otherwise, bid nothing.

Player 2:

- Bid nothing in in the first bidding phase.
- Bid 1 for $C$ in the second bidding phase after $A$ if they (player 2) did not bid in period 1
- Bid 4 for $C$ in the second bidding phase after $A$ if they (player 2) bid a non-zero amount in period 1
- Bid 2 for $C$ in the second bidding phase after $B$ regardless of previous bids.

Player 3:

- Play the value maximizing action in both periods.

The result is $(B, C)$ with a payoff of $(0, 2, 2)$. This is less efficient than $(A, C)$ which has a total payoff of $5$. To see that this is a subgame-perfect equilibrium note that player 1's bid is not realized, so he is doing as well as he can (0 payoff). Player 2 and player 3 are both bidding and acting in accordance with their value function.

The trick here is that $V(B) = (0, 2, 2)$ consistently but the value of $V(A)$ is inconsistent. $V(A) = (0, 4, 1)$ if player 2 does not bid in the first bidding phase but $V(A) = (0, 1, 4)$ if player 2 does bid. Player 2 would like to bid enough to cause Player 3 to play $A$ under normal circumstances. However, thanks to player 1's bid-based sunspot-style strategy, the process of bidding makes $A$ worse than $B$ for Player 2. There is no way for player 2 to access his desired value from $A$ without "destroying" it. Since $B$ is achievable by Player 2 without bids, no bidding takes place in period 1.

This example contrasts significantly with the model of Dutta and Siconolfi (2019) where uniqueness breaks down for non-Markovian-in-bids (what they call "action perfect equilibria") strategies only in the infinite horizon case. It also contrasts with Prat and Rustichini (1998) where all SPE are efficient. Only in environments with more than two players and more than one period does this difficulty arise.

The difference comes from the fact that bringing in more than two players causes non-unique optimal moves to be common, often with potential to influence the payoffs of other player. If there are two or more periods, this multiplicity can allow current bids to impact the value function indirectly, breaking down the incentives for efficient bidding.

## 5. PROPERTIES OF THE BIDDING MECHANISM

We turn now to the features of the sequential pivot bidding mechanism and show that it satisfies a number of important properties: weak first-mover advantage (in corollaries 1 and 2), and generic non-uniqueness of payoffs (in section 5.2). Proposition 1 characterizes when payoffs *are* unique.

### 5.1. *Order of Bids and Distribution of Payoffs*

We start by discussing the distribution of the payoffs and particularly how it is impacted by the exogenously specified order in which the players get to bid. While every order of bidders yields the utilitarian-efficient outcome, the order does influence the distribution of the bids (and therefore, of the final payoffs).

The non-uniqueness of the payoffs makes it difficult to give general statements about the payoff distribution, so we consider slightly narrower statements in this subsection. The intuition from these results does apply broadly.

Consider, first, a case where the bidders' preferences are aligned in the sense that there is an action that they all prefer over all other actions. Assume that the mover wants a different action to avoid a trivial outcome with no bids. In this case, there will be (weak) first-mover advantage—bidding earlier rather than later is better. The intuition is that because the preferences are aligned, the earlier bidder(s) can "shift" the burden of implementing the preferred outcome to later bidders.

COROLLARY 1: *Consider a single-move game where the mover (player $0$) wants one action $a$ and all bidders want $a' \neq a$. Moving a player to an earlier bidding position while keeping the order of the bids otherwise identical will weakly decrease that player's bid.*

This corollary follows from Theorem 1 and guarantees a weak first-mover advantage. If the incentives of the bidders are misaligned, the bid order has a more complex effect of changing the degree and type of non-uniqueness. We discuss uniqueness in more detail in the following subsection. Here, we present a simple case:

COROLLARY 2: *Consider a single-move game where the each player only receives a payoff from one action and no two players receive a payoff from the same action. The only bidder to get a payoff is the one whose preferred action benefits them the most. If they are the first bidder, they will pay the value of the player with the second highest value. If they are the last bidder, their payoff could take any value between zero and their value for their preferred option.*

This corollary also follows from Theorem 1. The result shows that, under certain conditions, the sequential bidding mechanism can resemble a second price auction, but changing the order of the bids can introduce a great deal of non-uniqueness.

Payoffs *are* unique only under stringent conditions, which we turn to now.

### 5.2. *Uniqueness of Payoffs*

In this section we establish necessary and sufficient conditions for realized bids to be unique. This result shows that uniqueness of payoffs is a very fragile property.

Changing the bidding order will generally change the set of possible distributions of the payoffs, as discussed above. In this section, we fix the order of the bids at each action history and only consider the payoff non-uniqueness arising from multiple equilibria with different realized bids.

We begin with some definitions: Take any history $h_t$, a MPBE $\sigma^*$(with specific properties we will define shortly), and an associated set of value functions $V_i\left(h_{t+1},\sigma^*\right)$. Consider the bids between $h_t$ and $h_{t+1}$. Suppose that in this equilibrium $\sigma^*$ each player only bids as required and makes no optional bids (this equilibrium exists in all cases). We denote the resulting bids for the optimal action ($a^*$) as $\hat{b}_i$, for convenience. We call the resulting leading actions $\hat{a}_i$. Next, we define the value of player $i$'s leading action in this equilibrium:

$$m_i = \hat{T}_i(\hat{a}_i) + F_i(\hat{a}_i) \tag{24}$$

where

$$\hat{T}_i(\hat{a}_i) = \begin{cases} V_{P(h_t)}(\hat{a}_i) + \sum_{j=1}^{i-1} \hat{b}_i & \hat{a}_i = a^* \\ V_{P(h_t)}(\hat{a}_i) & \hat{a}_i \neq a^* \end{cases} \tag{25}$$

and, as before, $V_{P(h_t)}$ is the value function of the moving player. Note that $F_i(a)$ is independent of the strategy profile.

Finally, we define a running total limit:

$$\bar{T}_i(a) = \max_{j<i} m_j - F_j(a) \tag{26}$$

We have the following result:

PROPOSITION 1: *The realized bids are unique between $h_t$ and $h_{t+1}$ if and only if $\bar{T}_i(a) + F_i(a) \leq m_i, \forall i, a$.*

For the proof, see the appendix. Payoffs are unique in a sequential game if this condition holds every pair of consecutive sub-histories along the efficient history, since it's an "if and only if" condition.

Essentially the result works as follows: a player sometimes has slack that lets them make unrealized bids, as long as they can do so without influencing the amount that they have to pay to implement their optimal action. This slack is what creates non-uniqueness. The proof looks at how much slack it is possible for each player to have and then checks whether there is enough to cause uniqueness to fail. The slack comes from the fact that players are willing to bid for the optimal action up until the running total plus their own result reaches a certain threshold. If the running total exceeds this threshold then that means that the given player does not need to bid for the leading action in that equilibrium and so they also have some freedom to influence the payoffs of others without influencing their own payoff.

As mentioned, this result shows that the uniqueness of payoffs is very fragile. For example, it implies that if the leading actions are ever different along the equilibrium path ($\tilde{a}_i \neq \tilde{a}_{i+1}$) then the equilibrium vector of the realized bids is not unique. This, in turn, means that if any player is not action-pivotal, uniqueness immediately fails.

In contrast, it is very difficult to guarantee unique payoffs. If there are only two actions and all players are action-pivotal, that is sufficient, but this is a very special case.

### 5.3. *Ranking of Equilibria*

Let us now briefly discuss how payoffs in $\Gamma$ are related to payoffs in $\Gamma^{BA}$. While is it trivially weakly optimal to participate in bidding in the bidding-augmented game (given that one can bid zero), it is not true that a player will always want to participate in the bidding augmented game (versus playing the underlying game without bids) because the outcome of the bidding-augmented game does not always Pareto-dominate the equilibrium outcome of the underlying game.

However, participation in the bidding-augmented game can be guaranteed to provide Pareto improvement in the following way: suppose that before playing either $\Gamma$ or $\Gamma^{BA}$, $i)$ each player can choose whether to veto allowing bids during the main game[5] and $ii)$ players can bid for each other player's action during the veto stage. It can be shown that the outcomes of this version of the game with bidding on vetoes will Pareto dominate the outcomes of $\Gamma$. We omit the proof as it is largely trivial (applying Theorem 1 twice).

[5]Thus, if any one (or more) player(s) chooses to not participate in $\Gamma^{BA}$, all players play $\Gamma$.

### 5.4. *Discussion of Applications*

Our main focus in on the theoretical result, but we turn now - briefly - to applications. The first is obvious enough to be almost trivial: Extensive-form games of perfect information are used throughout economics, particularly in applied game theory and empirical industrial organization (for instance, Selten's "chain store" game, Stackelberg competition, dictator games, and public good provision games). Often the equilibria of these games are inefficient (as would be the case, for instance, in an extensive-form perfect information analogue of a prisoner's dilemma), and the question becomes how to get to the efficient equilibrium. Our result implies that if transfers and contractible actions are available, this is all that is necessary. For instance, in a repeated game of entry with an incumbent and a series of potential entrants, if transfers are available, all the modeler has to do is determine what the utilitarian-efficient outcome is; our result ensures that the efficient outcome will prevail, regardless of other features (such as the number of entrants, periods or actions).

The main application of Dutta and Radner (2004) is enforcement of climate change treaties and the main application of Prat and Rustichini (1998) is contributions to a politician who will make a policy decision (i.e. "lobbying"). Our result immediately applies to these two settings as well. With regard to the latter setting, our result applies to a broader class of lobbying settings, including those where decisions might be taken sequentially (say, first in a lower house of a legislature, then in an upper house, and then endorsed by the executive), repeated instances of lobbying efforts, or lobbying different politicians.

## 6. IMPERFECTLY TRANSFERABLE UTILITY: MONEY IN THE UTILITY FUNCTION

The discussion so far has focused on a setting of transferable utility. While realistic in many applications, this assumption may not hold in certain important settings. This naturally leads to the question, will bidding for actions guarantee efficiency when utility is not perfectly transferable? To address this question we now present a version of our main result for a setting of *imperfectly* transferable utility (ITU), where agents' utilities include transfers ("money") potentially non-linearly. We state the result for the case of two players—one mover and one bidder. The restriction to two players and one action is necessary—counterexamples exist with two or more bidders or multiple sequential actions.

In this case we only consider a simplified game that has one actor who acts once and whose actions impact themselves and one other player. Augmenting this game gives the

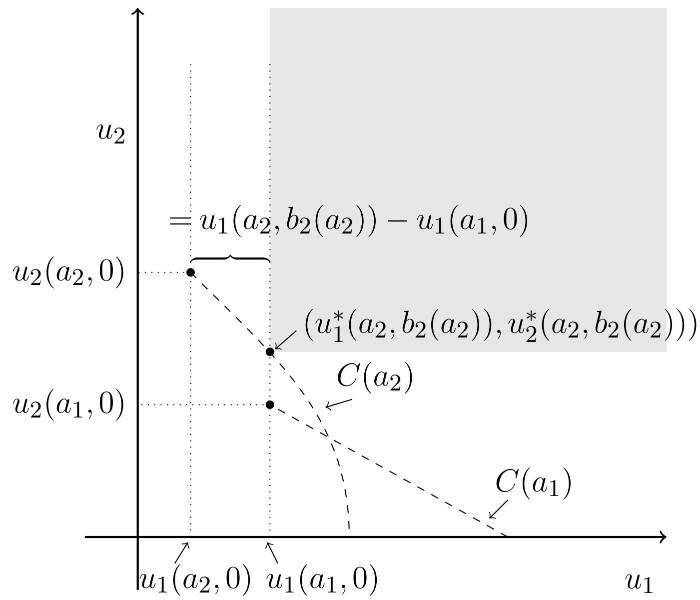


FIGURE 4.—The ITU case with two players. The equilibrium of the bidding-augmented game remains Pareto efficient. In this figure, action $a_1$ is the default action, which we denote in text by $a^d$.

other player an opportunity to bid on actions using money; we (naturally) assume that utilities are increasing in money.

PROPOSITION 2: *The outcome of the bidding-augmented version of the simplified game with ITU is Pareto efficient.*

PROOF: The argument is geometric, as developed in figure 4. Without loss, suppose that the mover is player 1 (with a utility function $u_1(a,b)$ and the bidding player is player 2 (with utility $u_2(a,b)$. The first player's utility is increasing in $b$ for $b > 0$, while the second player's utility is decreasing in $b$ over the same range.

Consider the utility space curves for each action, $C(a)$, which contain all points $(u_1, u_2)$ such that $u = u_1(a,b)$ and $u_2 = u_B(a,b)$ for some $y$. All of the curves $C(a)$ are downward sloping, since transferring money improves one player's utility at the cost of the other. Call the collection of all points on all curves $P$. Define the default action $a^d = \arg\max_a u_1(a,0)$ and further define $u_1^* = u_1(a^d, 0)$. Then, the implemented option will be $(u_1^*, u_2^*)$, where $u_2^* = \max_a \{u_2^a : (u_1^*, u_2^a) \in P\}$.

Since all of the curves are downward sloping, this is a Pareto-dominant point as long as no $C(a)$ has its most upper-left point to the upper right of $(u_1^*, u_2^*)$. We show that this would lead to a contradiction.

The most upper-left point of each $C(a)$ is the point $(u_1(a,0), u_2(a,0))$. If a curve begins to the upper right of $(u_1^*(a_2, b_2(a_2), u_2^*(a_2, b_2(a_2)))$, in the region shaded light gray, this would imply that for some $a'$ such that $u_1(a',0) > u_1^* = u_1(a_0, 0)$, which is a contradiction. *Q.E.D.*

## 7. LITERATURE REVIEW

As discussed in the introduction, our work contributes to several strands of the literature, chief among them the work on dynamic games with transfers and the work on "efficiency." The work of Prat and Rustichini (1998), Dutta and Siconolfi (2019), Dutta and Radner (2023), and Jackson and Wilkie (2005) are the closest to our approach. Having already discussed the link with Prat and Rustichini (1998) and Dutta and Siconolfi (2019) in the Introduction, we now focus on the others.

The work of Dutta and Radner (2023) also shows that adding transfers to an infinitely repeated game implements the utilitarian optimum. However, Dutta and Radner (2023) focus on a specific important game based on climate change negotiations with a state variable representing greenhouse gasses. Their focus is on the interpretation of their model—enforceability of climate change treaties.

Jackson and Wilkie (2005) allow their arbitrary number player to make binding contingent side payments before the underlying game is played. Crucially, and akin to Kalai and Kalai (2013), they focus on simultaneous-move games and simultaneous offers. The equilibria in their model may be *in*efficient (in fact, transfers may destroy all Pareto-dominant Nash equilibria of the underlying game with two players). The essential reason for the divergent conclusions (setting aside the modeling differences) is that, while in both the Jackson and Wilkie (2005) model and our work players can use transfers to internalize externalities, the use of simultaneous transfers (as opposed to our sequential approach) creates room for profitable deviations in the transfer phase, which undermines efficiency.

Similarly (and in contrast to our results) Ellingsen and Paltseva (2016) show that the Coase theorem need not hold in a fixed simultaneous-move game with $N > 2$ players, pre-game agreements to participate, and endogenous transfers. The reason is that (unlike in our model) if some players commit to neither give nor receive transfers, the remaining players players still play the transfer-modified game; this creates the potential for inefficiency.

Goldlücke and Kranz (2012) study a model that is related in spirit: adding transfers (as well as potentially "burning money" - lowering one's own payoffs) to a model of infinitely repeated games with imperfect public monitoring. In their model (as in ours), the transfer and action stages alternate, although the differences are quite significant - infinitely repeated games with imperfect monitoring are (obviously) not games of perfect information, and furthermore, in their setup, players act and choose transfers simultaneously. In such a model, equilibria may fail to be efficient, which is due in part to the possibility of burning money, and in part to the limited enforceability of transfers. Bernheim and Whinston (1986b) study a common agency model with potential incomplete information (an agent, whose action is unobserved), in which all equilibria are "efficient" (in their, different, sense of implementing an action at the lowest cost; different actions may be implemented in different equilibria, and the Pareto ranking of these is ambiguous in general); under some condition there exist equilibria that are Pareto efficient (being strong Nash equilibria).

Our work also has a similarity to the literature that highlights the importance of the "pivotal" agent—the agent without whom an outcome is not obtained and with whom the outcome is obtained. See Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998) for pivotal voting, Shapley and Shubik (1954) for an index of power (also based on pivotality), and Vickrey (1961), Clarke (1996), and Groves (1973) for their VCG mechanism. While our sense of pivotality is not quite the same (our pivotal bidders respond to future *values* but past *bids*) and arises endogenously, the similarity is interesting.

## APPENDIX: PROOFS

PROOF OF LEMMA 3: By Lemma 2, players bid pivotally, thus implementing action

$$a^* \in \max_{a \in A(h_t)} V_i(a) + T_i(a) + F_i(a) \tag{27}$$

for each player $i$.

Note, first, that multiple bidders may be pivotal (with respect to the same pivotal action) but the welfare-maximizing action is unique. Furthermore, different players cannot be pivotal with respect to *different* pivotal actions.

Letting $P(h_t) = N$ to simplify the notation, consider the situation from the point of view of player 1, the first bidder at an arbitrary history $h_t$. We have $\tilde{a}_1 = \arg\max_a T_1(a) + F_1(a) = \arg\max_a \sum_{k=2}^{N-1} V_k(a)$ as the leading action before player 1 bids. If player 1 is action-pivotal, they implement the utilitarian efficient outcome.

If player 1 is not action-pivotal, they will implement $\tilde{a}_1$ and, in equilibrium, we have the following relation:

$$\tilde{a}_1 = \arg\max_a V_1(a) + T_1(a) + F_1(a) \tag{28}$$

and

$$V_1(\tilde{a}_1) + T_1(\tilde{a}_1) + F_1(\tilde{a}_1) = V_1(\tilde{a}_1) + V_N(\tilde{a}_1) + \sum_{j=2}^{N-1} V_k(\tilde{a}_1) = \overline{V}(\tilde{a}_1) = \overline{V}(a^*) \tag{29}$$

That is, $\tilde{a}_1$ maximizes the joint value function (again, if player 1 is not action-pivotal).

If player 1 is action-pivotal, they will implement an action:

$$a^1 \in \arg\max_a V_1(a) + \sum_{i=2}^{N} V_i(a) \tag{30}$$

which is again the joint value-maximizing action $a^*$. If there are no other pivotal players, we are done. If all other pivotal players are pivotal with respect to $a^*$ and not with respect to other actions, we are also done.

Furthermore, because under pivotal bidding it is impossible for multiple players to be action-pivotal with respect to different actions, player 1 effectively implements the efficient action, which stays implemented throughout bidding process.

*Q.E.D.*

PROOF OF LEMMA 4: We argue towards a contradiction and begin with two observations.

First, note that for any history $h_t$, it must be that

$$\overline{V}(h_t, \sigma^*) \in \left[\min_{z \in Z(h_t)} \overline{\pi}(z), \max_{z \in Z(h_t)} \overline{\pi}(z)\right] \tag{31}$$

where $Z(h_t)$ is the set of terminal histories that succeed $h_t$ .

Second, note that by backward induction and Lemma 3, $\overline{V}(\emptyset,\sigma^*) \geq \overline{V}(h_t,\sigma^*)$, for any finite $h_t$.

Turning now to the proof of the Lemma, suppose, toward a contradiction, that

$$\overline{V}(\emptyset,\sigma^*) = \overline{\pi}(z^*) - \epsilon \tag{32}$$

and take any history $h_t$. By the first observation, $\overline{V}(h_t,\sigma) \in \left[\min_{z\in Z(h_t)} \overline{\pi}(z), \max_{z\in Z(h_t)} \overline{\pi}(z)\right]$ (for any strategy profile, not just in equilibrium) because the joint value function is an expectation over the outcomes in this range. Now, by continuity at infinity, choose an $\epsilon$ and take $t(\epsilon)$ such that $\max_{z\in Z(h^*_{t(\epsilon)})} \pi(z) - \min_{z\in Z(h^*_{t(\epsilon)})} \pi(z) \leq \frac{\epsilon}{2}$. Here, $h^*_{t(\epsilon)}$ is the history containing the first $t(\epsilon)$ elements of $z^*$. Trivially,

$$\overline{V}(h^*_{t(\epsilon)}) \geq \overline{\pi}(z^*) - \frac{\epsilon}{2} \tag{33}$$

Furthermore, since $h^*_{t(\epsilon)}$ is finite, by the second observation we must also have

$$\overline{V}(\emptyset,\sigma^*) \geq \overline{V}(h^*_t,\sigma^*) \geq \overline{\pi}(z^*) - \frac{\epsilon}{2} \tag{34}$$

which contradicts equation 32. *Q.E.D.*

PROOF OF PROPOSITION 1: We begin with a few observations. First, if realized bids are unique, then $b^*_i = T_i(a^*) + F_i(a^*) - T_i(\tilde{a}_i) - F_i(\tilde{a}_i)$ and $T_i(a^*) = V_{P(h_t)}(a^*) + \sum_{j=1}^{i-1} b^*_j$ are fixed for all equilibria. In addition, $F_i(a^*)$ is fixed regardless of the bidding strategies, so a unique payoff vector guarantees that $T_i(\tilde{a}_i) + F_i(\tilde{a}_i)$ must be the same for all equilibria and equal to $\hat{T}_i(\hat{a}(i)) + F_i(\hat{a}(i))$.

($\Rightarrow$) We first show the necessity of the condition: By pivotal bidding, optional bids cannot exceed

$$T_i(\tilde{a}_i) + F_i(\tilde{a}_i) - T_i(a) - F_i(a) \tag{35}$$

If the payoffs are unique, equation (33) is equal to

$$m_i - T_i(a) - F_i(a) \tag{36}$$

Thus, if the payoffs are unique, each player will bid up to a certain amount for action $a$:

$$b_i(a) \leq m_i - T_i(a) - F_i(a) \tag{37}$$

or, rearranging,

$$b_i(a) + T_i(a) \leq m_i - F_i(a) \tag{38}$$

In other words, player $i$ will be willing to bid only up until the bids plus the transfers reach the specified level. This means that the maximum possible running total after the bid is effectively independent of the current running total except in cases where the current running total already exceeds the limit.

Define $\bar{b}_i(a) = m_i - F_i(a)$ as a player's maximum running total for a given action. Under uniqueness, the running total for a (non-optimal) action during period $i$ can be up to

$$\bar{T}_i(a) = \max_{j<i} \bar{b}_j(a) \tag{39}$$

Note that, if $\bar{T}_i(a) + F_i(a) > m_i$ for some $a$ and $i$, then the payoff vector is not unique, since it is possible for the running total for $a$ to be $\bar{T}_i(a)$ and this would imply that $T_i(a) + F_i(a) \geq \hat{T}_i(\hat{a}(i)) + F_i(\hat{a}(i))$ and therefore there would be a different realized bid.

If the condition is violated for $i$, there is an equilibrium where player $i$ has a different running total plus future value for their leading action compared with the equilibrium with no optional bids. This change in value implies a change in the realized bid, by the definition of pivotal bidding. Hence, *the set of realized bids is not unique*.

($\Leftarrow$) To show sufficiency we argue by contradiction. Assume $\bar{T}_i(a) + F_i(a) \leq m_i, \forall i, a$ and there is another equilibrium with bids $b_i'$, running total $T'$, and leading actions $\tilde{a}'(i)$ such that $\hat{b}_i(a^*) \neq b_i'(a^*)$ for some $i$. First consider bidder $i$, for whom this is true. Note that

$$\hat{b}_i \neq b_i'(a^*) \tag{40}$$

which implies

$$T_i'\left(\tilde{a}_j'\right) + F_i\left(\tilde{a}_j'\right) > \hat{T}_i(\hat{a}_j) + F_i(\hat{a}_j) \tag{41}$$

Note that the inequality goes in this direction because the equilibrium that gives $\hat{T}$ and $\hat{a}$ is the one with the minimal bids on all actions up until $i$.

In this situation, for all $j < i$ we have

$$T_j'(a^*) + F_j(a^*) - T_j'\left(\tilde{a}_j'\right) - F_j\left(\tilde{a}_j'\right) = \hat{b}_j \tag{42}$$

since $i$ is the first divergence. Thus,

$$T_j'\left(a^*\right)+F_j\left(a^*\right)-T_j'\left(\tilde{a}_j'\right)-F_j\left(\tilde{a}_j'\right)=\hat{T}_j\left(a^*\right)+F_j\left(a^*\right)-\hat{T}_j(\hat{a}_j)-F_j(\hat{a}_j),\forall j<i \quad (43)$$

By the definition of $i$, we have that

$$T_j'\left(a^*\right)=T_j'(a^*),\forall j<i \quad (44)$$

and thus equation (43) reduces to

$$\hat{T}_j(\hat{a}_j)-F_j(\hat{a}_j)=T_j'(\tilde{a}_j')+F_j(\tilde{a}_j'),\forall j<i \quad (45)$$

By pivotal bidding, this means

$$b_j'(a)+T_j'(a)\leq m_j-F_j(a),\forall a\neq a^*,j<i \quad (46)$$

which implies $T_j'(a)\leq\overline{T}_i(a),\forall a\neq a^*$. By the definition of $i$, $T_i'(a^*)=\hat{T}_i(a^*)$. Combining this with equation (41), we obtain

$$\overline{T}_i(a)+F_i(a)>\hat{T}_i(\hat{a}_i)+F(\hat{a}_i) \quad (47)$$

for some $a$. This contradicts the assumption $\overline{T}_i(a)+F_i(a)\leq m_i,\forall i,a$.

*Q.E.D.*

REFERENCES


AUSTEN-SMITH, D., AND BANKS, J. S. (1996). "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *The American Political Science Review*, 90(1), 34-45. https://doi.org/10.2307/2082796 [26]

BERGEMANN, DIRK AND JUUSO VÄLIMÄKI. (2003) Dynamic Common Agency. *Journal of Economic Theory* 111(1), 23-48. https://doi.org/10.1016/S0022-0531(03)00079-6

BERNHEIM, B., AND M. WHINSTON (1986a). Menu auctions, resource allocation, and economic influence. *The Quarterly Journal of Economics*, 101(1), 1-31. https://doi.org/10.2307/1884639

BERNHEIM, B. D., & WHINSTON, M. D. (1986b). Common Agency. *Econometrica*, 54(4), 923-942. https://doi.org/10.2307/1912844

CLARKE, E. H. (1971). Multipart Pricing of Public Goods. *Public Choice* XI: 17-33 [26]

COASE, R.H. (1960). The Problem of Social Cost. *The Journal of Law and Economics* 1960 3:, 1-44 [2]

DUTTA, PRAJIT AND PAOLO SICONOLFI (2019). Asynchronous games with transfers: Uniqueness and optimality. *Journal of Economic Theory*, 184, https://doi.org/10.1016/j.jet.2019.05.005 [2, 3, 4, 6, 19, 25]

DUTTA, PRAJIT, AND ROY RADNER. (2023). Climate Change and Transfers: A Coase Theorem. Mimeo. [25]

DUTTA, PRAJIT, AND ROY RADNER (2004). Self-enforcing climate-change treaties. *PNAS* doi/10.1073/pnas.0400489101 [23]

ELLINGSEN, TORE & ELENA PALTSEVA (2016). Confining the Coase Theorem: Contracting, Ownership, and Free-Riding, *The Review of Economic Studies*, 83(2), 547-586, https://doi.org/10.1093/restud/rdw001 [25]

GOLDLÜCKE, S. AND S. KRANZ (2012). Infinitely repeated games with public monitoring and monetary transfers. *Journal of Economic Theory*, 147, 1191 1221. https://doi.org/10.1016/j.jet.2012.01.008

GROVES, T. (1973). Incentives in Teams. *Econometrica* 61: 617-631. [26]

FEDDERSEN, T., AND PESENDORFER, W. (1998). Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting. *The American Political Science Review*, 92(1), 23-35. https://doi.org/10.2307/2585926 [26]

JACKSON, MATTHEW AND WILKIE, SIMON (2005). Endogenous Games and Mechanisms: Side Payments Among Players, *Review of Economic Studies*, 72, issue 2, p. 543-566. [25]

KALAI, A. AND E. KALAI (2013). Cooperation in Strategic Games Revisited, *The Quarterly Journal of Economics*, 128(2), 917-966, https://doi.org/10.1093/qje/qjs074 [25]

MEDEMA, S. G. (2020). The Coase Theorem at Sixty. *Journal of Economic Literature*, 58(4), 1045-1128. https://www.jstor.org/stable/27030503 [2]

PRAT, A. AND RUSTICHINI, A. (1998). *Sequential Common Agency*. Mimeo. [2, 3, 4, 6, 19, 23, 25]

ROSENTHAL, ROBERT W. (1981) Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory* 25(1): 92-100. https://doi.org/10.1016/0022-0531(81)90018-1 [6]

SHAPLEY, L.S., AND M. SHUBIK (1954). A method for evaluating the distribution of power in a committee system. *American Political Science Review* 48: 787-792 [26]

VICKREY, WILLIAM (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* XVI: 8-37. [26]