UTILITARIANISM AND TRANSFERS: BIDDING FOR EFFICIENCY

ANDREW KOSENKO
Department of Economics and Finance, Marist College

NATHANIEL NELIGH
Department of Economics, University of Tennessee, Knoxville

Given any perfect information dynamic game with utilities that are continuous at infinity, allowing players to transfer utility to the moving player at every history ("bid" for actions) results in a utilitarian-efficient outcome, maximizing the sum of all players' utilities. Sequentially pivotal bidding, in which players bid just enough to change an action, taking into account previous bids and utilities of other bidders, plays a key role. The payoff distributions are generally non-unique (we provide a condition for uniqueness) and exhibit weak first-bidder advantage. If players are given veto power over the bidding process, participation in the bidding mechanism can be made individually rational. We also provide an extension to a setting of imperfectly transferable utility.

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

1. INTRODUCTION

It is perhaps a truism in economics that with perfect information, complete contracts, and no transaction costs, efficient outcomes will always be achievable. Perfect information rules out signaling, screening, and moral hazard problems, while complete contracts rule out the hold-up problem. Isn't it "obvious," therefore, that complete contracts without transaction costs will ensure that outcomes of dynamic games will be, in some sense, efficient? Until recently, however, reasonably general and precise game-theoretic results that clarify this intuition have been unavailable. Even the scope of this insight—when is it valid?—and the necessary assumptions have not been formally investigated. The work of Dutta and Siconolfi (2019) represents an advance in this area, showing 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) can be achieved as long as players can contract sequentially on the next action of the other player. This result provides an indication that the efficiency conjecture may be broadly true and the required degree of completeness in contracting may be minimal. In their setting, only local, self-enforcing payoff contracts and sequential play (as opposed to simultaneous play, which would be equivalent to imperfect information) are necessary to enforce efficiency. A result of this type is rather unusual in the game theory literature; here, the solution to a game with bidding on actions exists under very permissive conditions and is unique as long as there is a unique efficient outcome. Such a combination of broad existence, weak assumptions, and relatively frequent uniqueness is quite rare.

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 extensive-form game with perfect information will result in a utilitarian-efficient outcome. More precisely, we show that if, at

Andrew Kosenko: kosenko.andrew@gmail.com
Nathaniel Neligh: nneligh@utk.edu
We thank Navin Kartik, Joseph Stiglitz, and seminar audiences at the University of Tennessee, Knoxville for valuable comments.

each node, non-moving players are able to offer action-contingent transfers ("bids") to the moving player, the equilibrium outcome of this bidding-augmented game will always correspond to a utilitarian-efficient outcome in the original game. We find that players in equilibrium employ a novel *sequential-pivotal bidding* strategy, which 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* dynamic game (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, and transferable utility. This extends the original result of Dutta and Siconolfi (2019) in a number of ways. First, our result applies to arbitrary dynamic 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). This requires that payoffs converge in an absolute sense, not a relative sense; a repeated game with geometric discounting satisfies continuity at infinity, but so too do games with many forms of dynamically consistent non-stationary discounting, such as square discounting. There also exist games with payoffs that are continuous at infinity but that do not have a repeated game structure. In section 3 we present an example of an infinite game with continuous at infinity payoffs that does not have a repeated game structure or geometrically discounted payoffs—an "infinite centipede."

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. 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. The sequential pivot in our proof resembles the Vickrey-Clarke-Groves (VCG) mechanism, but rather than being operated by a principal, it arises 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 differently to those who bid before and after them. That is, they react to the bids of earlier bidders but to the valuations of later bidders. While these two quantities match up in some cases, they do not always fully coincide. This asymmetry means that earlier bidders are able to shift the burden of changing the action to later bidders, thus leading to a weak early bidder advantage. One effect related to this sequential difference in pivotality is that bidders will often be indifferent over a range of bids, particularly those bidders who cannot enforce their preferred actions. This means that, while efficiency is guaranteed, the distribution of the payoffs is often not unique.

We also remove the requirement from Dutta and Siconolfi (2019) that players act in a fixed order along every history. This is not a true extension, however, since it can be shown that the assumption is without loss as long as trivial actions are allowed, which they are in both the original paper and in our own.

The proof of our main result (Proposition 1) proceeds in three lemmas. Lemma 1 shows that players will bid pivotally relative to their value functions. We define the *leading action* during a player's bid as the action that will be implemented if they do not bid at all. The *pivotal action*—the one that maximizes the sum of the existing bids and future valuations plus the current player's valuation—is the one the player will implement if they bid optimally. The amount the player bids for the pivotal action is equal to the difference between the current bids plus future bidders' valuations for the leading action and the same quantity for the pivotal

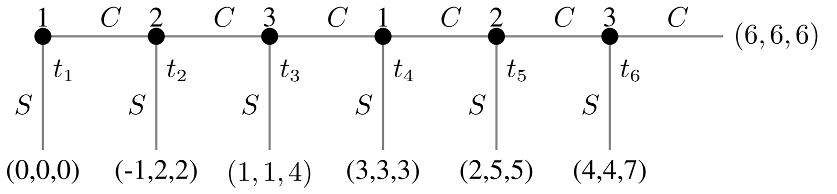

FIGURE 1.—Finite centipede with three players

action. If the pivotal action and leading action coincide, the player will not bid in a way that changes the result; this is because their presence in the sequence does not change the leading action. Lemma 2 shows that pivotal bidding will lead to the action that maximizes total value across all players. 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 the valuations of the other players. Finally, Lemma 3 shows that continuity at infinity and one-step-ahead optimality (Lemma 2) guarantee the efficiency of the outcome of the whole game.

As mentioned, the payoff vector generated by this process is generally not unique due to indifferences during the bidding process. To understand when the payoffs are uniquely determined, we present necessary and sufficient conditions for the uniqueness of the payoff vector. These conditions can be checked by using a simple algorithm. The conditions are quite strict; for example, if in the bidding sequence there are two players who have different leading actions in equilibrium, then the payoffs are guaranteed to be non-unique.

In addition to the payoff distribution result, we also provide several results on the properties of this mechanism: i) participation in the game with bids is individually rational, ii) all equilibria have the same total payoff, and iii) an extension to the case of imperfectly transferable utility shows that the result holds with two players and two actions but not more generally.

An Example: The Finite Centipede

To illustrate the workings of the bidding mechanism,[1] consider one well-known dynamic game—the centipede (Rosenthal (1981)). We first illustrate the mechanism and the result in this relatively simple, three-player setting. In section 4, we present an infinite-horizon example with two players.

Consider the finite centipede with three players and with the payoffs illustrated in figure 1.[2] 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, the unique equilibrium of the bidding-augmented centipede becomes one where all players continue at every opportunity, yielding the utilitarian-efficient outcome.

More explicitly, equilibrium bidding (suppose that nonmoving players bid in numerical order) prescribes that at t_6, P2 (the 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 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.

[1]While we use the term "bidding mechanism," this is not a true "mechanism" in the sense of mechanism design.
[2]Neither of the examples we develop are covered by the work of Dutta and Siconolfi (2019); one of our examples has three players and the other one has a non-repeated game structure.

Context and Connections

The ideas we explore are not entirely new. Ronald Coase (Coase (1960)) in "The Problem of Social Cost" tells a vivid story in which the noise from a confectionary maker next door is preventing a physician from practicing in a newly constructed examination room. The doctor sues, and the owner of the confectionary is forced to stop using their machinery. Coase goes on to observe that a better outcome (though bargaining) is possible:

> The doctor would have been willing to waive his right and allow the machinery to continue in operation if the confectioner would have paid him a sum of money which was greater than the loss of income [from moving, or building an insulating wall]. The confectioner would have been willing to do this if the amount he would have to pay the doctor was less than the fall in income he would suffer if he had to change his mode of operation [...]. The solution of the problem depends essentially on whether the continued use of the machinery adds more to the confectioner's income than it subtracts from the doctor's.

Coase's key insight—that bargaining restores efficiency—is at the core of our analysis. "Bidding for efficiency" may be but one method for attaining utilitarian optimality, stated for a particular setting (transferable utility, contractible actions, and utility having the property of continuity at infinity). We hope this setup may contribute to both the literature on efficiency in non-cooperative games and the literature on bilateral contracts.

From a theoretical perspective, our result may perhaps be viewed in at least three ways: $i)$ as an analogue of the first fundamental theorem of welfare economics, but stated for games with transferable utility, $ii)$ as a contribution to the Nash program, and $iii)$ as a contribution to the program on implementing jointly optimal decisions using transfers.

Viewed as a version of the fundamental theorems of welfare economics, rephrased for games, our result shows that an augmented type of Markov (with respect to bids) Nash equilibrium is utilitarian-efficient in arbitrary dynamic games with perfect information and that a utilitarian-efficient outcome can be "decentralized" in such games, using bilateral contracts instead of endowment reallocations. Our result is, to some degree, a variant of the first fundamental theorem but it focuses on stronger utilitarian efficiency, as opposed to the usual Pareto sense of efficiency, which is weaker.

Taking a Nash program (Nash (1953)) perspective, this may be viewed as a link between non-cooperative and cooperative games: an explicit non-cooperative game mechanism that illustrates how a cooperative outcome may be attained. One difference in our approach, of course, is that we do not axiomatize the utilitarian-efficient outcome, taking it as a primitive.

Finally, our result has a lot of commonalities with the long program of research attempting to overcome individually profitable deviations to implement a socially optimal outcome—Pigouvian taxation, VCG mechanisms, storable voting (Casella (2005)), and various methods of "overcoming incentive constraints by linking decisions," as was lucidly expressed by Jackson and Sonnenschein (2007).

Broadly, our result illustrates how a simple modification to a strategic situation—that is, the introduction of bids—may dramatically improve outcomes. Extensive-form games of perfect information are used throughout economics, particularly in applied game theory, behavioral economics, 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. 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.

The rest of the paper is organized as follows: Section 2 presents the notation, definitions, and equilibrium concept. Section 3 discusses two additional examples in detail. Section 4 states and proves the main result, while section 5 discusses features of the "bidding for efficiency" mechanism—(non)uniqueness of payoffs, weak first-mover advantage, and individual rationality. Section 6 extends the model to an imperfectly transferable utility setting with two players. The literature review is in section 7, and section 8 briefly concludes. Proofs are relegated to the appendix, with the exception of the proof of the fact that players bid pivotally (Lemma 1), which is illustrative and, thus, appears in the main text.

2. SETTING

We now formally introduce our setup.

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 function $A_i(h_t)$, 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 where no strict superhistory exists or infinite histories. Note that we treat any two infinite histories as the same history when they only diverge after an infinite number of actions. By continuity at infinity (see upcoming definition 1), this is without loss of generality. 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. Let U be the set of all possible payoffs in a game.

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

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

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

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.

The Bidding-Augmented Game

Given a dynamic game with perfect information, we can augment it such that immediately before each action is taken, each non-moving player may offer the player moving at that action history a set of non-negative transfers, contingent on the player's action taken. 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 Γ, 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 Γ. The set of players is the same and the histories of Γ^{BA} all exist in the set $\hat{\mathcal{H}}$ and are constructed by taking the histories of Γ and adding to each action history a transfer phase that precedes the action phase. During the transfer phase, each non-mover in sequence gets the opportunity to offer action-contingent transfers to the mover.

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 for, each action history, 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 $\{1, 2, 3, P(h_t) - 1, P(h_t) + 1, \dots\}$; 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 bid function, $\hat{A}$, is constructed similarly, giving the bidders options to bid before the actions. We keep the 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 $b_i(a_j; b_{-j<i}, V_{-j>i})$. 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. Player i's bid depends on the bids that have already been made (which we denote by $b_{-j<i}$) and the value functions (which we define below, in equation 5) of the later bidders as well as the moving player's value function (we denote these values by $V_{-j>i}$). In light of the complexity of such complete notation, we drop the 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 $r_t = (h_\tau, \{b_i\}_{i \in N \setminus P(h_t)})$, where $t = \tau N + |\{b_i\}_{i \in N \setminus P(h_t)}|$. The space of all such histories is denoted by R.

Note that (with a small abuse of notation) the histories of Γ 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

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

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, 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 on earlier bids. 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 $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ be a strategy profile. A strategy profile σ in Γ^{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

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

be the value function for player i at time t under strategy profile σ, with the understanding that the actions $a_i(r_t)$ and the bids $b_j(r_t)$ are determined according to the strategy profile σ. 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:

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

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

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.

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 (SEH)* and the outcome of said history a *strongly efficient outcome (SEO)*.

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.

3. EXAMPLE: AN INFINITE CENTIPEDE

Consider now an infinite version of the two-player centipede (see figure 2). Here we illustrate in detail the workings of the bidding mechanism and the results. 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, the payoffs are continuous at infinity. Thus, the fundamental tradeoff present in the centipede is preserved.

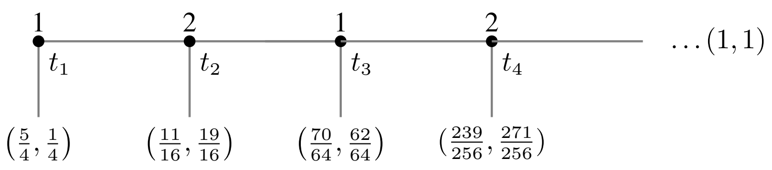

FIGURE 2.—An infinite centipede

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

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

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 $\frac{1}{4}$ at t_1 to player 1, player 1 will play C at t_1, instead of stopping. Onward, at t_2 P1 will bid $\frac{3}{16}$, at t_3 P2 will bid $\frac{6}{64}$, and so on.

The general bid structure is as follows:

$$b_i^t=\begin{cases}1-0.25^{t_i}+0.5^{t_i}+\sum_{i=1}^{t_i-1}t_i=1.25 & \text{for } t_i \text{ even}\\ 1-0.25^{t_i}+0.5^{t_i}-\sum_{i=1}^{t_i-1}t_i=0.75 & \text{for } t_i \text{ odd}\end{cases} \tag{9}$$

Note that only the non-moving player gets to bid. 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.

Importantly, players gain enough utility from the efficient actions to bid for the needed bids for those actions (this is apparent from inspecting the bid functions in equation 2). 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.

PROPOSITION 1: *The outcome of every MPBE of* Γ^{BA} *results in a strongly efficient outcome* z^* *of* Γ.

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, Proposition 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 (using Lemma 1) that players have an incentive to bid enough to implement the efficient action ("pivotally") in equilibrium. Then we show (in Lemma 2) 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 3) we show that one-period-ahead efficiency along with continuity at infinity guarantees overall efficiency. This is efficiency "across" periods. Taken together, Lemmas 2 and 3 prove Proposition 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 portions 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, and suppress it in the notation of the value function as well. This allows us to treat the value of various h_{t+1}s as essentially end points with defined values.

Before we prove Lemma 1, 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, which implies the result.

Given a fixed set of bids $\{b_k\}_{k=1,2,\ldots,i-1}$ and a strategy profile σ, 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 utility 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.

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 1—Pivotal bidding in equilibrium: *In any MPBE of* Γ^{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

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

the solution of which is identically equal to the a_N^* defined by action-pivotality because there are no future players 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 2—One-step-ahead optimality: *In any MPBE of the bidding-augmented game Γ^{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 2 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.

Second Step: One-Step-Ahead Optimality is Equivalent to Global Optimality

We finish the proof with the following lemma:

LEMMA 3—One-step-ahead Optimality is Equivalent to Global Optimality: *If the payoffs are continuous at infinity and*

$$\bar{V}(h_t, \sigma^*) = \max_{a \in A(h_t)} \bar{V}((h_t, a), \sigma^*) \tag{23}$$

then

$$\bar{V}^*(\emptyset, \sigma^*) = \max_{z} \bar{\pi}(z) \tag{24}$$

For the proof, see the appendix. The proof works by using continuity at infinity and Lemma 2 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.

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), generic non-uniqueness of payoffs (in Proposition 2 we give a necessary and sufficient condition for the payoff to be unique), and, finally, individual rationality (if we include a biddable veto process by which any player can veto participating in Γ^{BA}), which we discuss in Proposition 4,

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 make general statements about the payoff distribution, so we consider slightly narrower statements in this section. These results are, however, rich in intuition.

Consider, first, a case where the bidder's preferences are aligned in the sense that there is an action that they all prefer over all other outcomes. 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 Proposition 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 section. 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 Proposition 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.

Uniqueness of Payoffs

In this section we establish necessary and sufficient conditions for realized bids to be unique. Changing the bidding order will generally change the set of possible distributions of the payoffs. 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, an equilibrium σ^*(with specific properties we will define soon), 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 σ^* 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:

$$m_i = \hat{T}_i(\hat{a}_i) + F_i(\hat{a}_i) \tag{25}$$

where

$$\hat{T}_i(\hat{a}_i) = \begin{cases} V_{P(h_t)}(\hat{a}_i) + \sum_{j=1}^{i-1} \hat{b}_i & \hat{a}_i = a^* \\ V_{P(h_t)}(\hat{a}_i) & \hat{a}_i \neq a^* \end{cases} \tag{26}$$

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:

$$\bar{T}_i(a) = \max_{j<i} m_j - F_j(a) \tag{27}$$

We have the following result:

PROPOSITION 2: *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. This proof shows that checking one specific equilibrium is sufficient to determine whether the realized bids are unique. This works because the specific equilibrium checked has the maximal allowance for non-uniqueness in a tight, achievable way.

In dynamic settings with finite time, one can start at terminal nodes and work backwards to check for the overall uniqueness of the payoffs. In infinite time settings this is not generally possible.

As corollaries we have two additional results, which may, perhaps, be useful in applications.

COROLLARY 3: *Suppose at each history there are only two actions and that all players are value-pivotal. Then the equilibrium vector of the realized bids is unique.*

Corollary 3 provides an easy way to check the sufficient condition for payoff uniqueness.
By contrast, Corollary 4 provides another easy way to check a sufficient condition for payoffs to *not* be unique:

COROLLARY 4: *Suppose the leading actions are different along the equilibrium path:* $\tilde{a}_i \neq \tilde{a}_{i+1}$*. Then the equilibrium vector of the realized bids is not unique.*

Pareto Efficiency and Individual Rationality

Let us now discuss how payoffs in Γ are related to payoffs in Γ^{BA}, and, therefore, whether players are incentivized to participate in the bidding-augmented game. Suppose that before the players play Γ, they are given the option of playing Γ^{BA} instead. Would they wish to do so? Does it depend on the game, on the equilibria of the game, or on the order in which they get to bid?

Weak Pareto Efficiency

We begin by noting that if a player finds themselves in Γ^{BA}, bidding (any weakly positive amount) is preferred to not bidding at all. This is so simply because any player can always bid zero for all actions and for all other players. With strictly positive bids, a player may improve their lot; thus, a player always weakly prefers to bid.

Importantly, it is *not* true that every equilibrium of Γ^{BA} Pareto dominates every equilibrium of Γ—there exist situations in which while the *total* payoff in Γ^{BA} is greater than the total payoff in Γ, the *individual* payoffs in Γ^{BA} may be lower that those in Γ for some (though not all) players. The following counterexample illustrates this: Suppose there are three players, 1, 2, and 3. Only player 3 has an action to take (the other two players take no non-bidding actions in this simple game), choosing between action A, with payoffs $(1, 1, 1)$, and action B, with payoffs $(10, 0, 0)$. In this Γ (i.e., without bids), the only "equilibrium" is A. With bids present (and supposing that player 1 bids before player 2 for player 3's action), player 1 will bid 2 units for action B, yielding payoffs $(8, 0, 2)$. Player 2's payoff is thus lower in Γ^{BA} than it was in Γ.

Therefore, we define a modified game where we explicitly model each player's decision whether to participate in Γ^{BA} (following, *inter alia*, Jackson and Sonnenschein (2007)).

Consider the following situation: Suppose that before playing either Γ or Γ^{BA}, a) each player can choose whether to veto allowing bids during the main game[3] and b) players can bid for each other player's action during the veto stage. In other words, there is a pre-game stage at which players decide, in some fixed sequential order, which game to play, each player holds veto power (over playing Γ^{BA} as opposed to the default Γ), and players can also bid for other players' actions during this pre-game stage. We denote this game as Γ^{BAV} for "bidding-augmented with veto."

PROPOSITION 3—Weak Pareto efficiency: *Fix* Γ*. The payoff of each player in* Γ^{BAV} *is weakly greater than their payoff in* Γ.

Thus, while the equilibrium payoff vector in Γ^{BA} does not necessarily Pareto dominate the equilibrium payoff vector in Γ, the equilibrium payoffs in Γ^{BAV} do Pareto dominate the payoffs in the corresponding equilibrium of Γ.

The intuition behind this result—and its proof—is simple: We apply Proposition 1 twice. The first application (to the Γ^{BA} subgame of Γ^{BAV}) yields a utilitarian-efficient outcome in that

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

subgame. Applying the proposition to Γ^{BAV} guarantees that Γ^{BAV} will be played. Combining this with the fact that any player can adopt a strategy of not bidding and vetoing in the initial veto phase yields the result in Proposition 3.

Individual Rationality: Participation in the Bidding-Augmented Game

We now turn to the question of individual rationality. Would all players participate in Γ^{BA} relative to Γ, given the choice? More precisely, consider the following:

Definition 7—Individual rationality**.** Fix a voting order for the voting phase and an equilibrium σ^*_Γ of Γ. Participation in Γ^{BAV} is *individually rational (IR)* for player i if the payoff $\pi_i(z^*)$ (where z^* is the outcome of σ^*_Γ of Γ) is weakly less than their payoff (including the net transfers made at the pre-game stage) in any MPBE σ^* of Γ^{BAV} that contains σ^*_Γ after any veto.

Definition 7 is a standard definition of individual rationality as being weakly better than an outside option for all players. The following proposition affirms that participation in Γ^{BAV} is individually rational by this definition:

PROPOSITION 4—Individual Rationality: *Participation in Γ^{BAV} is weakly individually rational (relative to participating in Γ) for each player i.*

Proposition 4 is simply an immediate corollary of Proposition 3. Proposition 4 also shows the issue we noted above—that without additional bids, the payoffs of some players may be lower in Γ^{BA} than in Γ—does not arise if the pre-game stage has the biddable veto component described above. The reason this is so is that while some players may not get to make payoff-critical moves in Γ (and, therefore, they will never receive substantial bids), *all* players move and receive bids in Γ^{BAV}.

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 is necessary—counterexamples exist with two or more bidders.

Utilities are increasing in money. 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 other player an opportunity to bid on actions using currency.

PROPOSITION 5: *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 3. 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.

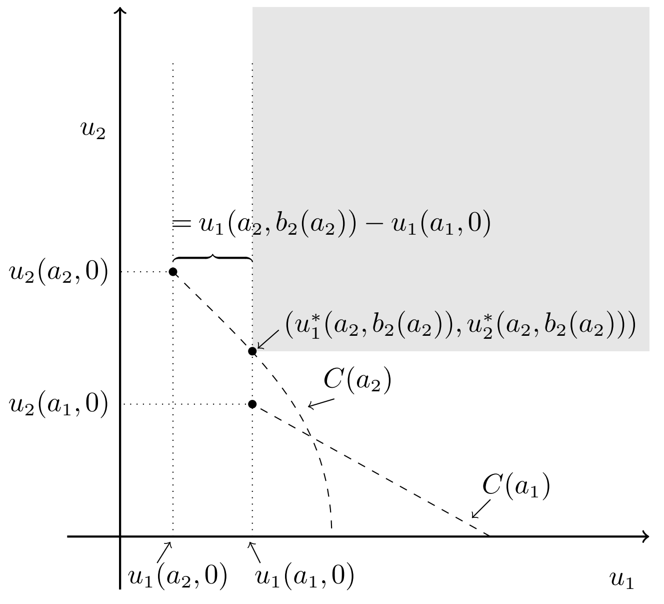

FIGURE 3.—The ITU case with two players. The equilibrium of the bidding-augmented game remains Pareto efficient. In this figure, action a_1 is the default action, which we denote in text by a_0.

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 $a_0 = \arg\max_a u_1(a, 0)$ and further define $u_1^* = u_1(a_0, 0)$. 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 for some a' 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" and the Coase conjecture; we also touch on the Nash program.

The work of Jackson and Wilkie (2005), Dutta and Siconolfi (2019), and Dutta and Radner (2023) is, perhaps, the closest to our approach.[4] Dutta and Siconolfi (2019) prove a similar result — allowing players to use transfers yields the utilitarian-efficient outcome—in the setting of asynchronous repeated games with two players. The work of Dutta and Radner (2023) is also very closely related—they also show that in an infinite horizon game with a fixed, dynamic stage game (i.e., a repeated dynamic game) adding transfers implements the utilitarian optimum. The difference is that what Dutta and Radner (2023) do for a specific important game

[4]Similar ideas have found applications in computer science: Li and Lin (2021), Kuhnle, Richley, and Perez-Lavin (2023).

with a state variable, we do for an arbitrary dynamic game (without state variables). As well, their focus is on the interpretation of their model—enforceability of climate change treaties. Reflecting this, they add an element absent in our model—the level of greenhouse gas emissions, which evolves according to a transition equation that depends on the joint choices of the players—and neither model is a special case of the other, due to the difference in the payoff specifications. An important area of application of these ideas has been the economics of climate change, as Dutta and Radner ((2004), (2023)) highlight; see also the references to the climate literature therein.

Thus, the differences between our work and that of Dutta, Siconolfi, and Radner are as follows: Dutta and Siconolfi (2019) present a two-player game with a repeated game structure; Dutta and Radner (2023) present an N-player dynamic game with a specific functional form assumption and a state variable. Our work is stated for N players, with or without a repeated game structure, and the only assumption we make on the preferences is continuity at infinity.

Jackson and Wilkie (2005) allow their players (perhaps more than two) to make such binding side payments; in their setup, this contracting takes place before the underlying game is played. Crucially, and akin to Kalai and Kalai (2013), they focus on simultaneous-move games (a situation not covered by our setup). 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 (and 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, Jackson and Wilkie (2005) rely on simultaneous transfers (as opposed to our sequential approach); this creates room for profitable deviations in the transfer phase, which undermines efficiency. Furthermore, in our work players are able to make further transfers as the game progresses (or veto participating in the voting phase). This possibility of repeated interaction, absent from Jackson and Wilkie (2005), ensures that individual optimization and utilitarian efficiency motives coincide. Table 1 illustrates the differences among the work discussed so far along two dimensions—the number of players and the horizon of the game.

The results are even more different in Eso and Schummer (2004), who also focus on the two-player second-price auction (and, therefore, *ipso facto* in an incomplete information setting) and arrive at a different conclusion: They show that "bribes" result in inefficiency in the sense that the object is misallocated with a positive probability (i.e., the low value bidder bribes the high value bidder and the bribe is accepted with positive probability in all robust equilibria).

We also contribute to the discussion on some classic issues in bargaining, namely, the Coase (1960) conjecture, and Medema (2020), who provides an updated discussion and references. The difference between our work and the conversation around the Coase conjecture is that we use a stronger efficiency concept and a game theoretic setting. In terms of results, we show that in fact no intervention of a mechanism designer or a court is needed—utilitarian-efficient agreements are self-enforcing, given sequential bidding.

A similar insight appears in Calabresi (1970), who argues that the burden of preventing accidents lies on the "cheapest cost avoider"—the party who can most easily prevent an accident. Of course, one issue in this setting is that the "cheapest cost avoider" may not be concerned with the consequences of an accident (that need not involve the "cheapest cost avoider" at all). Our work, however, shows that if the situation is analyzed as a dynamic game, then it is in the interest of *other* parties to i) avoid an accident by having the "cheapest cost avoider" prevent it, and ii) reimburse the "cheapest cost avoider."

By contrast, Ellingsen and Paltseva (2016), in a similar setting (pre-game agreements to participate and endogenous transfers) show that the Coase theorem need not hold (they work in a setting of a fixed simultaneous-move game with N players). The reason efficiency does not emerge as the only outcome in their setting is that if some players do not agree to the proposed transfers, all players still play the transfer-modified game. This creates the possibility that some players may "free-ride" on others' agreements. In our case, of course, if any player vetoes the bidding-augmented game, all players play the underlying non-transfer-modified game.

Jackson et al (2015) present a model animated by the same spirit but in a substantially different setting. They study a setting with two players "negotiating" over a set of items; negotiation takes the form of various alternating offer processes (generalizations of the Rubenstein-Stahl protocol) with transferable utility. They provide a condition for when such negotiations are efficient (in the sense that the maximum total surplus is realized— the same sense is used here), although it is not easily intepretable in terms of our formalism. They also provide an example where an equilibrium is efficient in the case of asymmetric information about the maximum possible joint surplus. Their work is not reliant on knowledge of the distributions (i.e., it is "detail-free"), nor is it reliant on a mechanism designer.

By contrast, in the mechanism design literature, Jackson and Sonnenschein (2007) provide an important contribution. They show that given a social choice problem, with common knowledge of the possible preference distributions and without transferable utility, there exists a mechanism that, in a sequence of repetitions of the initial problem, (approximately) implements the efficient outcome. The equilibria they construct are (approximately) efficient, where efficiency is judged by certain utility levels dictated by the target social choice function. The efficiency sense here differs from ours: The informational arrangements are different (they allow for incomplete information and assume common knowledge of the type distributions), but the flavor of the result—implementing "efficient" outcomes in dynamic games among groups of players—is the same.

Our work also contributes to the Nash program (Nash (1953)) that aims to provide a foundation (in the form of noncooperative games) for cooperative games and solution concepts; Serrano ((2005) (2021)) and Durlauf and Blume (2010) provide comprehensive and up-to-date surveys of this literature. To illustrate the connection between our contribution and this program of research, we quote from Serrano (2021), who writes:

> By proposing non-cooperative games that specify the details of negotiation, the Nash program [....] will tell a story about how coalitions form and what sort of interaction among players is happening. In that process, [...] the cooperative solution will be understood as the outcome of a series of strategic problems facing individual players. Second, novel connections and differences among solutions may now be uncovered from the distinct negotiation procedures that lead to each of them. [...] Focusing on the features of the rules of negotiation that lead to different cooperative solutions takes one a long way in opening the 'black box' of how a coalition came about, and contributes to a deeper understanding of the circumstances under which one solution versus another may be more appropriate to use.

Viewed from this perspective, our work provides one such procedure—a dynamic game form, transferable utility, and contractible actions—for how a cooperative solution concept that prescribes payoffs/value that add up to a utilitarian value may be implemented in a particular situation.

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*), the similarity is marked.

8. CONCLUDING REMARKS

Throughout the paper we suppose that the utilitarian outcome is desirable and we work under this assumption to illustrate the efficiency result.

There are certainly reasons why one may not wish to implement the utilitarian outcome in some settings.[5] However as we show, participation in a bidding-augmented game can be made individually rational. Therefore, the utilitarian outcome in our setting is perhaps more attractive than in other settings and may not suffer from the same criticisms (i.e., inequality, or lack of focus on the welfare of the least well-off participants). Transfers ensure that while not all players gain equally, they do gain enough to make it worth their while.

Finally, there is one additional, perhaps peculiar, interpretation of our results. The main result states that with transferable utility and contractible actions, transfers (bribes) can make everyone better off under certain conditions. This may seem counter-intuitive, as one often thinks of bribery as decreasing overall welfare. There are several things that make such "bribery" welfare-improving in our context. First, the continuity-at-infinity assumption is crucial for this interpretation: This assumption implies that payoffs far into the future are not too relevant for today's decision-making. Real-world bribery may violate this. For instance, if today's bribes erode social ties and public trust in various institutions, eventually leading to a collapse of trust (which is critical for economic activity and development), continuity at infinity would fail, and our result would not apply. Secondly, the individual rationality result (Proposition 4) shows that "bribery is good"—but only provided everyone knows about this, agrees to participate in the game with bribes, and everybody may veto this option before any bribery takes place (aside from the bribery about the vetos). This is also not necessarily how real bribery takes place—often it is covert, an act of subterfuge. Finally, because real bribery takes place largely (or entirely) outside the scope of the law, actions may not be fully contractible—yet another sense in which our assumptions are not entirely reflective of the "bribery" interpretation.

APPENDIX: PROOFS

PROOF OF LEMMA 2: By Lemma 1, players bid pivotally, thus implementing action

$$a^* \in \max_{a \in A(h_t)} V_i(a) + T_i(a) + F_i(a) \tag{28}$$

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:

$$\tilde{a}_1 = \arg\max_a V_1(a) + T_1(a) + F_1(a) \tag{29}$$

[5]To wit, Rawls's (1971) *A Theory of Justice*, Dworkin's (1977) *Taking Rights Seriously*, and Nozick's (1974) *Anarchy, State, and Utopia* all reject the utilitarian approach.

and

$$V_1(\tilde{a}_1) + T_1(\tilde{a}_1) + F_1(\tilde{a}_1) = V_1(\tilde{a}_1) + V_N(\tilde{a}_1) + \sum_{j=2}^{N-1} V_k(\tilde{a}_1) = \overline{V}(\tilde{a}_1) = \overline{V}(a^*) \tag{30}$$

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 the following action:

$$a^1 \in \arg\max_a V_1(a) + \sum_{i=2}^{N} V_i(a) \tag{31}$$

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 3: We argue towards a contradiction and begin with two observations.

First, note that for any history h_t, it must be that $\overline{V}(h_t, \sigma^*) \in [\min_{z \in Z(h_t)} \overline{\pi}(z), \max_{z \in Z(h_t)} \overline{\pi}(z)]$, where $Z(h_t)$ is the set of terminal histories that succeed h_t .

Second, note that by backward induction and Lemma 2, $\overline{V}(\emptyset, \sigma^*) \geq \overline{V}(h_t, \sigma^*)$, 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, $\overline{V}(h_t, \sigma) \in [\min_{z \in Z(h_t)} \overline{\pi}(z), \max_{z \in Z(h_t)} \overline{\pi}(z)]$ (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 ϵ and take $t(\epsilon)$ such that $\max_{z \in Z(h^*_{t(\epsilon)})} \pi(z) - \min_{z \in Z(h^*_{t(\epsilon)})} \pi(z) \leq \frac{\epsilon}{4}$. Here, $h^*_{t(\epsilon)}$ is the history containing the first $t(\epsilon)$ elements of z^*. Trivially,

$$\overline{V}(h^*_{t(\epsilon)}) \geq \overline{\pi}(z^*) - \frac{\epsilon}{4} \tag{33}$$

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

$$\overline{V}(\emptyset, \sigma^*) \geq \overline{V}(h^*_t, \sigma^*) \geq \overline{\pi}(z^*) - \frac{\epsilon}{4} \tag{34}$$

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

PROOF OF PROPOSITION 2: 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

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

If the payoffs are unique,equation 33 is equal to

$$m_i - T_i(a) - F_i(a) \quad (36)$$

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

$$b_i(a) \leq m_i - T_i(a) - F_i(a) \quad (37)$$

or, rearranging,

$$b_i(a) + T_i(a) \leq m_i - F_i(a) \quad (38)$$

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

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

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

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

which implies

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

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

$$T'_j(a^*) + F_j(a^*) - T'_j\left(\tilde{a}'_j\right) - F_j\left(\tilde{a}'_j\right) = \hat{b}_j \quad (42)$$

since i is the first divergence. Thus,

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

By the definition of i, we have that

$$T'_j(a^*) = T'_j(a^*), \forall j < i \tag{44}$$

and thus equation (43) reduces to

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

By pivotal bidding, this means

$$b'_j(a) + T'_j(a) \leq m_j - F_j(a), \forall a \neq a^*, j < i \tag{46}$$

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

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

for some a. This contradicts the assumption $\overline{T}_i(a) + F_i(a) \leq m_i, \forall i, a$, and we are done.
Q.E.D.

PROOF OF COROLLARY 1: For notational convenience, say that the mover is player 0. Say also that the player of interest is player $k.5$. In other words, they move after Player k and before Player $k+1$. This notation allows us to avoid relabeling the other players after changing the bidding position.

Assume WLOG $\pi_0(a) = 0$ and $\pi_i(a') = 0 \forall i \neq 0$, so the payoffs are being expressed in terms of the relative gain. The case where $\pi_0(a') \geq \sum_{i \neq 0} \pi_i(a)$ is trivial since no bidder receives positive payoffs and no bids are realized.

Now consider the case where $\pi_0(a') < \sum_{i \neq 0} \pi_i(a)$. Option a will be implemented. Given pivotal bidding, $k.5$ will pay $\pi_0(a') - \sum_{i=1}^{k} b_i(a) - \sum_{i=k+1}^{N-1} \pi_i(a)$. Note that $b_i(a)$ can depend on whether i is before or after $k.5$ but, otherwise, it does not depend on k. Under pivotal bidding, $b_i(a) \leq \pi_i(a)$, since a is implemented (otherwise it would not be worth it to make the bid). Therefore, the amount paid to the player of interest is weakly increasing in k. *Q.E.D.*

PROOF OF COROLLARY 2: For notational convenience, say that the mover is player 0. Assume $\pi_i(a_i) \geq 0$ and $\pi_i(a_j) = 0, \forall j \neq i$. If $\pi_0(a_0) \geq \max_i \pi_i(a_i)$ then no bidder ever gets a non-zero payoff. Now assume that $k = \arg\max_i \pi_i(a_i)$ and $k > 0$. Under pivotal bidding, a_k will be implemented. If player k goes first, they must bid $\max_{i \neq k} \pi_i(a_i)$ according to pivotal bidding.

If player k goes last then their bid is determined by $\max_i \pi_0(a_i) + \sum_{j=1}^{k-1} b_j(a_i)$. Since this value is determined by bids for non-implemented options, the bids are not constrained to be less than the bidders' valuations. Instead they are constrained to be low enough not to change the implemented option. As such, the bids from earlier players can be anywhere between 0 and $\pi_k(a_k)$. Player k will then have to bid the maximum over the other players' bids and $\pi_0(a_0)$. *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 [18]

CALABRESI, G. (1970). *The Cost of Accidents: A Legal and Economic Analysis*. Yale University Press. http://www.jstor.org/stable/j.ctt1cc2mmw [17]

CASELLA, A. (2005). Storable Votes. *Games and Economic Behavior* 51(2), 391-419. https://doi.org/10.1016/j.geb.2004.09.009 [4]

CLARKE, E. H. (1971). Multipart Pricing of Public Goods. *Public Choice* XI: 17-33 [18]

COASE, R.H. (1960). The Problem of Social Cost. *The Journal of Law and Economics* 1960 3:, 1-44 [4, 17]

DURLAUF, S.N., BLUME, L.E. (2010). Nash program. In: Durlauf, S.N., Blume, L.E. (eds) *Game Theory. The New Palgrave Economics Collection*. Palgrave Macmillan, London. https://doi.org/10.1057/9780230280847_26 [18]

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 [1, 2, 3, 16, 17]

DUTTA, PRAJIT, AND ROY RADNER. (2023). Climate Change and Transfers: A Coase Theorem. Mimeo. [16, 17]

DUTTA, PRAJIT, AND ROY RADNER. (2004). Self-enforcing climate-change treaties. *PNAS* doi/10.1073/pnas.0400489101 [17]

DWORKIN, R. (1977). *Taking Rights Seriously*. Harvard University Press, Cambridge, MA. [19]

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 [17]

ESO, P., AND J. SCHUMMER. 2004. Bribing and signaling in second price auctions. *Games and Economic Behavior*. 47(2): 299-324. https://doi.org/10.1016/j.geb.2003.06.005 [17]

GROVES, T. (1973). Incentives in Teams. *Econometrica* 61: 617-631. [18]

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 [18]

JACKSON, MATTHEW O. AND SONNENSCHEIN, HUGO AND XING, YIQING (2015). A Theory of Efficient Negotiations. Available at SSRN: https://ssrn.com/abstract=2580894 or http://dx.doi.org/10.2139/ssrn.2580894 [18]

JACKSON, MATTHEW AND WILKIE, SIMON(2005). Endogenous Games and Mechanisms: Side Payments Among Players, *Review of Economic Studies*, 72, issue 2, p. 543-566. [16, 17]

JACKSON, M.O. AND SONNENSCHEIN, H.F. (2007), Overcoming Incentive Constraints by Linking Decisions. *Econometrica*, 75: 241-257. https://doi.org/10.1111/j.1468-0262.2007.00737.x [4, 14, 18]

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 [17]

KUHNLE, A., J. RICHLEY, AND D. PEREZ-LAVIN (2023). Learning Strategic Value and Cooperation in Multi-Player Stochastic Games through Side Payments. arXiv:2303.05307, https://doi.org/10.48550/arXiv.2303.05307 [16]

LI L. AND D. LIN (2021). Three-Player Cooperative Game With Side-Payments for Discretionary Lane Changes of Connected Vehicles. *IEEE Access*, vol. 9, pp. 159848-159857, doi: 10.1109/ACCESS.2021.3130031. [16]

MEDEMA, S. G. (2020). The Coase Theorem at Sixty. *Journal of Economic Literature*, 58(4), 1045-1128. https://www.jstor.org/stable/27030503 [17]

NASH, J. (1953). Two-Person Cooperative Games. *Econometrica*, 21(1), 128-140. https://doi.org/10.2307/1906951 [4, 18]

NOZICK, R. (1974, 2006). *Anarchy, State, and Utopia*. Basic Books, New York, NY. [19]

RAWLS, J. (1971). *A Theory of Justice: Original Edition*. Harvard University Press. https://doi.org/10.2307/j.ctvjf9z6v [19]

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 [3]

SERRANO, R. (2004). Fifty years of the Nash program, 1953-2003. *Investigaciones Economicas* 29: 219-258. [18]

SERRANO, R. (2021) Sixty-seven years of the Nash program: time for retirement?. SERIEs 12, 35-48. https://doi.org/10.1007/s13209-020-00221-5 [18]

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 [18]

VICKREY, WILLIAM (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* XVI: 8-37. [18]