Incentive-Compatible, Budget-Balanced, yet Highly Efficient Auctions for Supply Chain Formation

Moshe Babaioff
School of Computer Science and Engineering
The Hebrew University of Jerusalem, Israel
mosheb@cs.huji.ac.il

William E. Walsh
IBM T. J. Watson Research Center
19 Skyline Dr., Hawthorne, NY, 10532, USA
wwalsh1@us.ibm.com

ABSTRACT

Engineering automated negotiation across the supply chain is a central research challenge for the important problem of supply chain formation. The difficult problem of designing negotiation strategies is greatly simplified if the negotiation mechanism is incentive compatible, in which case the agents' dominant strategy is to simply report their private information truthfully. Unfortunately, with two-sided negotiation it is impossible to simultaneously achieve perfect efficiency, budget balance, and individual rationality with incentive compatibility. This bears directly on the mechanism design problem for supply chain formation—the problem of designing auctions to coordinate the buying and selling of goods in multiple markets across a supply chain. We introduce incentive compatible, budget balanced, and individually rational auctions for supply chain formation inspired by previous work of Babaioff and Nisan, but extended to a broader class of supply chain topologies. The auctions explicitly discard profitable trades, thus giving up perfect efficiency to maintain budget balance and individual rationality. We use a novel payment rule analogous to Vickrey-Clarke-Groves payments, but adapted to our allocation rule. The first auction we present is incentive compatible when each agent desires only a single bundle of goods, the auction correctly knows all agents' bundles of interest, but the monetary valuations are private to the agents. We introduce extensions to maintain incentive compatibility when the auction does not know the agents' bundles of interest. We establish a good worst case bound on efficiency when the bundles of interest are known, which also applies in some cases when the bundles are not known. Our auctions produce higher efficiency for a broader class of supply chains than any other incentive compatible, individually rational , and budget-balanced auction we are aware of.

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: Electronic Commerce—*Payment schemes*
; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems, Coherence and Coordination*

General Terms

Algorithms, Economics

Keywords

Mechanism Design, Auctions, Supply Chain Formation

1. INTRODUCTION

Supply chain formation is the problem of determining the production and exchange relationships across a supply chain. Whereas typical research in supply chain management focuses on optimizing production and delivery in a fixed supply chain structure, we are concerned with ad hoc establishment of supply chain relationships in response to varying needs, costs, and resource availability. These individual relationships cannot be established in isolation because a functioning supply chain requires a complete sequence of production through the supply chain. As business relationships become ever more flexible and dynamic, there is an increasing need to automate this supply chain formation process. Automated supply chain formation is being recognized as an important research challenge [11, 13], and has been chosen as the subject of the upcoming 2003 Trading Agent Competition.

Because procurement and supply contracts in a supply chain can involve significant production commitments and large monetary exchanges, it is important for an agent to negotiate effectively on behalf of a business. However, strategic analysis can be very complex when agents must negotiate contracts for outputs and multiple inputs simultaneously across a supply chain. Fortunately, careful design of the negotiation mechanism can simplify the agents' strategic problem enormously. We can effectively engineer away the agents' strategic problem by designing an auction to be ***incentive compatible*** (IC), in which case an agent's dominant strategy is to simply report its private information truthfully. Other properties are also important in a business setting. An auction should be ***individually rational*** (IR), that is no agent would pay more than its valuation for the goods it receives. The auction should be ***budget balanced*** (BB) (the auction does not lose money), else there would typically be little incentive to run the auction. Additionally, it is desirable that the auction be efficient (maximize total agent value) to ensure that all gains from trade are realized.

To date, there has been relatively little mechanism design work that meets the needs of automated supply chain formation. To address the problem of IR, much recent effort has focused on combinatorial auctions [3] which, by allowing agents to place indivisible bids for bundles of goods, ensure that agents do not buy partial bundles of no value. Much of this work has been on one-sided auctions. To address the two-sided negotiation necessary in a supply chain, Walsh et al. [14] analyzed an IR and BB auction that avoids negotiation miscoordination by allowing combinatorial bids across the supply chain. They found the strategic analysis challenging, and were able to derive Bayes-Nash equilibria for only restricted network topologies [13]. In contrast, the well-known Vickrey-Clarke-Groves (VCG) auction [2, 4, 12] (also called the Generalized Vickrey Auction [6]) is IC and efficient, but not BB with the two-sided bidding needed in a supply chain. Myerson and Satterthwaite showed that, in two-sided negotiation, it is, unfortunately, impossible to simultaneously achieve perfect efficiency, BB, and IR from an IC mechanism [8]. In response to this impossibility, Parkes et al. [10] explored double auction rules that minimize agents' incentives to misreport their values, but maintains BB and high (but not perfect) efficiency.

In this work, we exploit the fact that, despite the impossibility theorem, it *is* possible to attain IC with any two of the three desirable properties (efficiency, IR, BB) in an auction for supply chain formation. To ensure IC, IR and BB, we develop auctions that produce inefficient allocations by design. If this approach seems misguided, we note that the Myerson-Satterthwaite theorem actually states more strongly that the three properties cannot be obtained even in Bayes-Nash equilibrium. Thus, since efficiency loss is inevitable in supply chain formation (assuming BB and IR), we focus on simplifying the agents' strategic problem by ensuring IC. Still, it is important that we do not ignore efficiency altogether, for a highly inefficient auction would likely be unacceptable for business negotiations. Indeed, a trivial way to get IC, IR, and BB is to perform no allocation, which is clearly unacceptable. Babaioff and Nisan [1] presented a novel approach to obtaining IC, BB, and IR and high efficiency in linear supply chains by structuring auctions in terms of production markets, rather than directly as goods exchanges. This allowed them to use a variant of McAfee's double auction [7] to obtain the properties.

In this paper, we use ideas from Babaioff and Nisan's approach to introduce auctions that are IC, BB, and IR for a broader class of supply chain formation problems. We provide good worst case bounds on efficiency when the auction knows the agents' bundles of interest, and in some cases when it does not. Our auctions produces higher efficiency for a broader class of supply chains than any other IR, IC, and BB auction we are aware of.

In Section 1 we describe our model of the supply chain formation problem. In Section 3 we present an auction that is IC, IR, and BB when each agent desires only a single bundle of goods, there is only one way to produce each good, the auction correctly knows all agents' bundles of interest, but the monetary valuations are private to the agents. We establish a good competitive ratio for allocative efficiency. We also outline an algorithm for computing the auction in polynomial time, given fixed consumer preference structures. In Section 4 we introduce extensions to maintain incentive compatibility when the auction does not know the agents' bundles of interest. We conclude and suggest avenues for future work in Section 5.

2. SUPPLY CHAIN FORMATION PROBLEM

2.1 Supply Chain Model

Before describing the formal details, we illustrate a supply chain with a stylized example in a small lemonade industry, as shown in Figure 1. The figure shows in a ***supply chain graph*** how the lemon juice and lemonade can be manufactured from lemons and sugar by agents in the supply chain. In the figure, an oval indicates a good in the supply chain. A box indicates a market, which is a set of agents who desire exactly the same set of input and output goods. The arrows indicate the input/output relationships between the agents and the goods. The goods are traded in discrete quantities, and under each good we indicate the discretization of the goods. For each market, the quantity of inputs needed by one agent and outputs that can be produced by one agent are indicated next to the respective arrows. We assume that an agent can provide one unit of its output good but may require multiple units of an input good. Borrowing a term from Lehmann et al. [5], we say the agents are single minded to identify the property that each agent has a single bundle of input and output goods that is of interest to the agent. This can often be a reasonable assumption, for companies typically have an established way to produce a product.

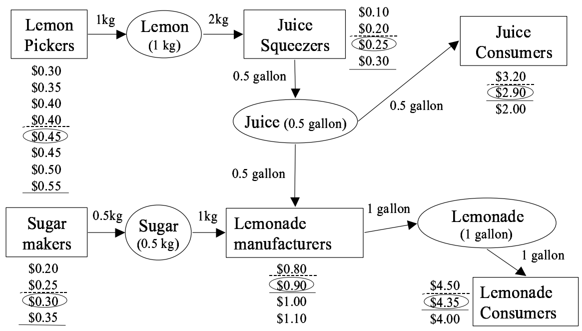

Figure 1: An example supply chain graph in the lemonade industry.

An agent with an output is a producer, and an agent with only inputs is a consumer. For instance, a lemonade manufacturer (a producer) requires 1kg of sugar and 0.5 gallons of juice to produce one gallon of lemonade, and a lemonade consumer wishes to buy 1 gallon of lemonade. A consumer obtains a monetary value from acquiring its bundle of interest, and a producer incurs a monetary cost from producing a good. The values and costs of the individual agents are indicated in a list adjacent to each market (observe that values are sorted from highest to lowest and costs are sorted from lowest to highest).

The formal model we describe subsumes the linear supply chain model described by Babaioff and Nisan [1] but is subsumed by the model described by Walsh [13].

Formally, we have a set A of agents and a set G of goods, with agents indicated by integers in $[1, \ldots, |A|]$ and goods indicated by integers in $[1, \ldots, |G|]$. A bundle $q = (q^1, \ldots, q^{|G|})$ indicates the quantity q^g of each good g exchanged by an agent. Positive quantity indicates acquisition of a good (input), and negative quantity indicates provision of a good (output). We restrict our attention to quantities $q^g \in \{-1, 0\} \cup \mathbf{Z}^+$. In other words, we consider agents that can require multiple units of an input, but produce at most one unit of an output. We further restrict our attention to ***single output agents*** that supply at most a single unit of a single good. That is $q^g = -1$ for at most one good g. When comparing quantities in bundles of goods, we assert $\tilde{q}_i \geq q$ when $\tilde{q}_i^g \geq q^g$ for all g, and assert $\tilde{q}_i > q$ when $\tilde{q}_i \geq q$ and $\tilde{q}_i^k > q^k$ for some good k.

Agent i has a ***valuation function*** V_i that assigns a value to any bundle q, and $V_i(q) \in \{-\infty, \mathbf{Z}\}$. Agent i obtains utility $U_i(q, M) = V_i(q) - M$, for exchanging bundle q and paying M monetary units. We assume that the agents are rational and try to maximize their utility over all possible outcomes. We refer to $V_i(q)$ as agent i's ***value*** for the bundle of goods q, and we denote the vector of all agents' valuation functions by $\mathbf{V} = (V_1, \ldots, V_{|A|})$. We interpret negative values as ***costs*** (e.g., cost of production, or opportunity cost of providing a good). We assume the valuation functions are normalized at $V_i(\mathbf{0}) = 0$ and the value is ***weakly monotonic***[1] in the quantity of goods, that is $V_i(\tilde{q}_i) \geq V_i(q)$ for all $\tilde{q}_i$ such that $\tilde{q}_i > q$. When $V_i(q) = -\infty$, we say that the bundle q is ***infeasible*** for agent i, and when $V_i(q) \in \mathbf{Z}$ we say the bundle is ***feasible*** for the agent.

Agents are ***single minded*** in that each agent i has a unique ***bundle of interest*** $\hat{q}_i$ that it tries to obtain. The composition of a bundle of interest and an agent's valuation thereof, depend on the class of the agent, as we detail below. We assume for all agents i that $V_i(\hat{q}_i) \in \mathbf{Z}$. For convenience, we subsequently denote $V_i(\hat{q}_i)$ as v_i. The ***market*** $K(i)$ of agent i is the set of agents with exactly the same bundle of interest, formally defined as $K(i) = \{j \mid \hat{q}_j = \hat{q}_i\}$.

There are two classes of agents, defined by further constraints on agents' bundles of interest and values. A ***consumer*** i obtains positive value ($v_i \geq 0$) for acquiring all goods in its bundle of interest ($\hat{q}_i > \mathbf{0}$), but cannot produce any goods . The consumer's value $V_i(q)$ for bundle q is such that:

- If $q \geq \hat{q}_i$, then $V_i(q) = v_i$ (single minded and weakly monotonic).
- Else, if $q^k < 0$ for some good k, then $V_i(q) = -\infty$ (a consumer cannot feasibly produce any good).
- Otherwise $V_i(q) = 0$ (a consumer has zero value for any feasible bundle not containing its bundle of interest).

A ***producer*** i can produce a single unit of a single output from a specific (possibly empty) set of inputs, while incurring a cost: $v_i \leq 0$ and $\hat{q}_i^g = -1$ for exactly one good g. A producer cannot feasibly produce its output without all inputs, nor can it feasibly produce any other output. The producer's value $V_i(q)$ for bundle q is such that:

- If $q^g = \hat{q}_i^g = -1$ and $q \geq \hat{q}_i$, then $V_i(q) = v_i$ (single minded and weakly monotonic).
- Else, if $q^g = \hat{q}_i^g = -1$ and $q^k < \hat{q}_i^k$ where $\hat{q}_i^k > 0$ for some good k, then $V_i(q) = -\infty$ (a producer needs all inputs to feasibly produce its output).
- Else, if $q^k < 0$ where $\hat{q}_i^k \geq 0$ for some good k, then $V_i(q) = -\infty$ (a producer can feasibly produce only one good).
- Otherwise $V_i(q) = 0$ ($q \geq 0$ and a producer has zero cost if it does not produce any good).

Finally, we consider only supply chains with ***Unique Manufacturing Technologies (UMT)***, in which there is only one market that produces any good. Note that, although a good can be made in only one way, there can be multiple producers in any market, and multiple markets that require the good as an input. As we will show, UMT is necessary to ensure BB and our efficiency competitive ratio in our auction. However, the auction is IC and IR without the UMT restriction.

The relationship between markets and goods can be represented as a supply chain graph, as illustrated in Figure 1 and described above. We assume that any graph is directed acyclic, but can have undirected cycles. The market structure defines the supply chain topology.

DEFINITION 1 (SUPPLY CHAIN TOPOLOGY). *A **supply chain topology** is a set of markets.*

[1]Weak monotonicity is equivalent to free disposal for agents.

2.2 Allocations

Given a set of agents, we want to determine the production and exchange of goods that constructs a supply chain. An ***allocation*** $\mathbf{q}$ specifies how much of each good is bought and sold by each agent. Let the allocation of good g to agent i be q_i^g, with $q_i^g > 0$ meaning that i buys $|q_i^g|$ units of g, and $q_i^g < 0$ meaning that i sells $|q_i^g|$ units of g in the allocation. Allocation q is ***feasible*** iff each agent is feasible and each good is in ***material balance***, that is $\sum_{i \in A} q_i^g = 0$ for each good g.[2]

Throughout this paper, we will consider only allocations that give an agent either all or none of its bundle of interest. Since each agent has one bundle of interest, it will be convenient to identify an allocation $\mathbf{q}$ by the set of agents A' that receive their bundle in the allocation: $A' = \bigcup_{i \in A | q_i \neq \emptyset} i$. The ***value*** $\mathbf{V}(A')$ ***of an allocation*** A' is the sum the agent values in A': $\mathbf{V}(A') \equiv \sum_{j \in A'} v_j$. The value of an allocation A' excluding the value of agent i is $\mathbf{V}_{-i}(A') \equiv \sum_{j \in A', j \neq i} v_j$. In the auction below, the true values are not known, so the allocation values are computed with respect to the values reported in the agents' bids. When it is necessary to specify the values explicitly, we denote the value of allocation A' with respect to specific values v as $\mathbf{V}^v(A')$. An allocation A^* is ***efficient*** if it is feasible and maximizes the value over all feasible allocations. The ***efficiency*** of allocation A' is $\frac{\mathbf{V}(A')}{\mathbf{V}(A^*)}$.

The efficient allocation A^* for the supply chain graph shown in Figure 1 has value \$7.90, and contains the agents whose costs and values are specified above the solid line in each market. The reader can verify that all goods are in material balance and that each agent in A^* receives its bundle of interest. For instance, each of the two lemonade manufacturers in A^* require 1kg of sugar to produce its output, and there are four sugar makers in A^* to provide the 2kg required in total. Similarly, there are two lemonade consumers to buy each of the 1 gallons of lemonade produced by the lemonade manufacturers.

The following definition is useful in proving our theorems.

DEFINITION 2 (PROCUREMENT SET). *A **procurement set** $S\{A'\}$ in allocation A' is a set of agents constituting a non-empty feasible allocation that contains no other non-empty feasible allocations.*

Clearly, any non-empty feasible allocation can be partitioned into procurement sets.

3. AUCTION FOR THE KNOWN SINGLE-MINDED MODEL

Here we present an auction for a known single-minded model of agent utility. We say "known" because we assume that it is common knowledge that the auction correctly knows the bundle of interest of all agents, but an agents' monetary valuation for its bundle of interest is private and independent of other agents' values. The "known" assumption can be plausible in established industries where production technologies are well known.

For obvious reasons, we call the auction KSM-TR (Known Single-Minded Trade Reduction). Under the KSM model, an auction is IC iff each agent has the incentive to report its true valuation for its desired bundle. We show that KSM-TR is IC, IR, and BB, and has a good competitive ratio for efficiency.

[2]Since the agents can produce at most a single unit of a single good, the set of allocations would be the same if we required only that supply weakly exceed demand.

3.1 KSM-TR Auction Mechanism

Each agent reports a value $\check{v}_i$, which may or may not be v_i, to the auction. The auction then computes an allocation, which assigns, for each agent, either its bundle of interest or the zero bundle. It also computes payments to be made by each agent. The auction is a centralized mechanism that uses trade reduction (TR) rules in a manner based on an auction introduced by Babaioff and Nisan [1], but for a more general supply chain model. The auction first computes an optimal allocation, based on the reported values, and uses this to compute a TR allocation and the agents payments. To ensure IC and BB, the auction then removes some beneficial trades from the optimal allocation.

We specify KSM-TR as mixed-integer-linear programs (MIPs), which can be represented in standard MIP format (including the binary max and min functions, with a simple transformation). We can apply advanced integer programming techniques from operations research to solve the MIPs, but computing the solutions are intractable for sufficiently large problems. Our problem is a generalization of winner determination in a combinatorial auction with single-minded preferences, which has been shown to be NP-hard [5]. However, as we show in Section 3.4, we can compute the auction in polynomial time for a fixed number of consumer markets.

Descriptions of major variables:

- $e_i \in \{0,1\}$ indicates whether agent i receives its bundle of interest in the chosen bid-optimal allocation.
- $x_i \in \{0,1\}$ indicates whether agent i receives its bundle of interest in the TR allocation.
- $\bar{x}_i \in \{0,1\}$ indicates whether agent i is not in the TR allocation.
- $r_i \in \{0,1\}$ indicates whether agent i is the price bounding agent (explained below) in its market.

First, we compute a ***bid-optimal allocation*** with respect to the reported values $\check{v} = (\check{v}_1, \ldots, \check{v}_{|A|})$:

$$\begin{aligned} &\text{maximize} && \textstyle\sum_{i\in A} \check{v}_i e_i \\ &\text{such that} && \textstyle\sum_{i\in A} q_i^g e_i = 0, \quad \text{for each good } g. \end{aligned} \tag{1}$$

Given the bid-optimal allocation $\tilde{A}(\check{v}) = \{i \mid e_i = 1\}$ with respect to the reported values $\check{v}$, we compute the TR allocation $A^{TR}(\check{v}) = \{i \mid x_i = 1\}$ according to the TR rules, given by a solution to the equations below. Given a solution, if $x_i = 0$, then agent ***i loses*** the auction, but if $x_i = 1$, then i ***wins*** the auction. A winner pays at least the value bid by the price bounding agent in the same market (Lemma 28).

The MIP equations are:

$$\text{maximize} \sum_{i\in\tilde{A}(\check{v})} \check{v}_i x_i. \tag{2}$$

Subject to the following Trade Reduction constraints ((3)–(9)) .

Each allocated and price-bounding agent is in $\tilde{A}(\check{v})$:

$$x_i \le e_i; \quad r_i \le e_i \tag{3}$$

Each agent $i \in \tilde{A}(\check{v})$ is either in the allocation or not in the allocation:

$$x_i + \bar{x}_i = 1 \tag{4}$$

All goods g are in material balance:

$$\sum_{i\in A} q_i^g x_i = 0 \tag{5}$$

For all agents $i \in \tilde{A}(\check{v})$, an agent can't be simultaneously in the allocation and bound the price:

$$r_i + x_i \le 1 \tag{6}$$

For every market k, if there is at least one agent in this market that is in the allocation, then exactly one agent bounds the price in the market, otherwise no agent bounds the price:

$$\sum_{i\in k\cap\tilde{A}(\check{v})} r_i = \max \bigcup_{i\in k\cap\tilde{A}(\check{v})} x_i \tag{7}$$

For every market k, the price-bounding agent must have a lower value than all allocated agents in the market. Thus, for all pairs $i, j \in k\cap\tilde{A}(\check{v})$, $i \ne j$:

$$m_{i,j} = \min(r_i, x_j); \quad \check{v}_i m_{i,j} \le \check{v}_j m_{i,j} \tag{8}$$

For every market k, the price-bounding agent must have a higher value than all non-allocated agents in the market. Thus, for all pairs $i, j \in k\cap\tilde{A}(\check{v})$, $i \ne j$:

$$t_{i,j} = \min(r_i, \bar{x}_j); \quad \check{v}_i t_{i,j} \ge \check{v}_j t_{i,j} \tag{9}$$

In the above, we have implicitly assumed that there is exactly one bid-optimal allocation and exactly one TR allocation that satisfy the equations. In general, we need a rule to break ties between multiple bid-optimal and TR allocations. It can be shown that the auction is not IC if we break ties between alternate bid-optimal allocations in favor of the one that gives the maximum bid-value TR allocation. However, we maintain IC if we break ties randomly, independent of reported valuations. A computationally efficient way to perform the random, value-independent tie breaking is as follows. First, we require that all valuations be reported to the auction as integers. The auction randomly maps the integers $[1, \ldots, |A|]$ to agents, one-to-one. The value 2^{-i} is added to the reported value of an agent assigned to the number i. To see that this modification to the bids makes (1) compute a unique, optimal allocation, observe that $\sum_{i=1}^{|A|} 2^{-i} < 1$, hence any allocation computed is optimal with respect to the bids as they are submitted. Observe also that for any two disjoint sets of positive, integers N and M, we have $\sum_{i\in N} 2^{-i} \ne \sum_{j\in M} 2^{-j}$, hence exactly one allocation satisfies (1) with respect to the modified bids. Similarly, there is a unique TR allocation that satisfies (2). We do not include the 2^{-i} components in the agent payments (described below).

If the true values shown in Figure 1 are reported to the auction, then $\mathbf{V}(\tilde{A}(\check{v})) = \mathbf{V}(A^*) = \$7.90$ and $\tilde{A}(\check{v})$ contains all agents with values and costs above the solid lines. All agents above the dashed lines are in the trade reduction allocation $A^{TR}(\check{v})$ and $\mathbf{V}(A^{TR}(\check{v})) = \$4.70$, giving an efficiency of 0.59. The TR rules require that we remove at least one agent from $\tilde{A}(\check{v})$ for each market, hence we reduce one agent from each of the following markets: juice consumers, lemonade manufacturers, and lemonade consumers. Since one agent is removed from the juice consumers and lemonade manufacturers markets, we have to remove *two* agents from the juice squeezer market to maintain material balance of the juice good. Because each juice squeezer require 2kg of lemons, but each lemon picker provides only 1kg of lemons, we must remove four agents from the lemon pickers market to maintain material balance of the lemon good. Similarly, we must remove two agents from the sugar markers market to maintain material balance of the sugar good.

The ***price bounding agent*** $PBA_i(\check{v})$ for i and bids $\check{v}$ is agent j such that $r_{i,j} = 1$, which is the highest bidding agent reduced from $\tilde{A}(\check{v})$ in i's market. By Lemma 6 and Lemma 8, $PBA_i(\check{v})$ is independent of i's bid when it wins, so we denote $PBA_i = PBA_i(\check{v})$. The

price bounding value PBV_i for i is $\check{v}_{PBA_i}$. We use the term "price bounding agent" because i pays at least PBV_i (Lemma 28) in our auction. As we show in Lemma 29, PBV payments are BB, which means our auction is BB since i pays at least PBV_i. So, in effect, the PBAs serve as "cutoff points" to ensure that the payments from all agents above these points constitute BB.

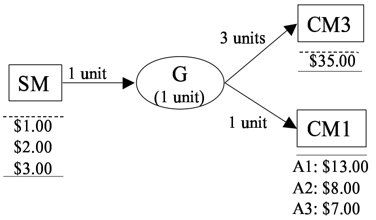

Figure 2: A supply chain for which PBV payments are not incentive compatible.

Babaioff and Nisan's auction [1] computes the TR allocation (in a computationally efficient, distributed fashion for linear supply chains) and requires agent i to pay PBV_i. In Figure 1, the values reported by price-bounding agents are just below the dashed lines and circumscribed by ovals. Although PBV_i payments give IC for linear supply chains, they do not give IC in our more general model, as demonstrated in Figure 2. If agent A1 bids \$13.00, as indicated, it does not win because it is not in the bid-optimal allocation. A1 has an incentive to bid any value above \$20.00 because then it would win but pay only $PBV_{\mathsf{A1}} = \$8.00$.

Here we describe a new payment scheme to obtain IC in our model. Each agent pays the ***Vickrey Trade Reduction (VTR)*** value $VTR_i(\check{v})$, which we construct in the following. We denote as $\check{v} \equiv (\check{v}_i, \check{v}_{-i})$ the vector of values reported by all agents, where $\check{v}_{-i}$ is the vector of values reported by all agents except i. Let $\tilde{A}(\check{v})$ be the bid-optimal allocation with respect to $\check{v}$, and $\tilde{A}(\check{v}_{-i})$ be the bid-optimal allocation with respect to $\check{v}_{-i}$.

The VCG value of i with respect to the bids $\check{v}$ is defined as

$$VCG_i(\check{v}) \equiv \mathbf{V}(\tilde{A}(\check{v}_{-i})) - \mathbf{V}_{-i}(\tilde{A}(\check{v})) \qquad (10)$$

Intuitively, $VCG_i(\check{v})$ is the "harm" done by agent i to the other agents by bidding $\check{v}_i$. Observe that $VCG_i(\check{v}) \leq \check{v}_i$ and that $VCG_i(\check{v}) = 0$ if i is not in a bid-optimal allocation. Consider A1 in Figure 2. If it bids as shown in the figure, it is not in the bid-optimal allocation and $VCG_i(\check{v}) = 0$. If instead it bids \$100, then A1 is in the bid-optimal allocation and $VCG_{\mathsf{A1}} = 29 - 9 = 20$. Observe that i would be in the bid-optimal allocation if it bids any value above \$20.

As mentioned above, VCG payments are not BB, but we can extend the VCG idea to obtain BB payments in the TR auction. The Vickrey Trade Reduction (VTR) value for agent i with respect to the bids $\check{v}$ is defined as:

$$VTR_i(\check{v}) \equiv \mathbf{V}(A^{TR}(VCG_i(\check{v}), \check{v}_{-i})) - \mathbf{V}_{-i}(A^{TR}(\check{v})) \qquad (11)$$

where $A^{TR}(VCG_i(\check{v}), \check{v}_{-i})$ is the TR allocation obtained when the bid of i is replaced by $VCG_i(\check{v})$. The values are computed with respect to the bids used to compute the TR allocations. If $i \in A^{TR}(\check{v})$ and tie breaking is necessary, we use the $A^{TR}(VCG_i(\check{v}), \check{v}_{-i})$ and $A^{TR}(\check{v})$ allocations containing i in the VTR computation. Observe that if $i \notin A^{TR}(\check{v})$, then since $VCG_i(\check{v}) \leq \check{v}_i$, $\mathbf{V}(A^{TR}(VCG_i(\check{v}), \check{v}_{-i})) = \mathbf{V}_{-i}(A^{TR}(\check{v}))$ and $VTR_i(\check{v}) = 0$. Consider A1 in Figure 2. If it bids as shown in the figure, it is does not win and $VTR_i(\check{v}) = 0$. If instead it bids \$100, then A1 wins and $VTR_i(\check{v}) = 25 - 5 = 20$. In fact, if i bids any value above \$20 it would win and pay \$20.

We note that $VTR_i(\check{v}) = PBV_i$ for linear supply chains, hence our auction is equivalent to Babaioff and Nisan's auction when applied to linear supply chains.

3.2 Incentive Compatibility, Individual Rationality, and Budget Balance

The main theorem we prove for the KSM-TR mechanism is:

THEOREM 3. *The KSM-TR auction produces a feasible allocation, and is incentive compatible in dominant strategies, individually rational, and budget balanced.*

PROOF. The auction produces a feasible allocation because it treats the bids as all-or-nothing and ensures that allocated supply equals allocated demand (Equation (5)). Incentive compatibility is proven in Lemma 11, individual rationality is proven in Lemma 10, and budget balance is proven in Lemma 29. ☐

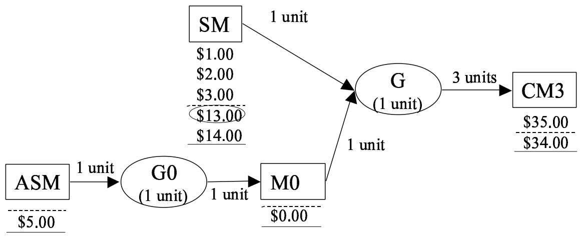

Figure 3: A supply chain without the UTM property and for which KSM-TR is not BB.

We note that the proofs of IC and IR do not depend on the UMT property. However, KSM-TR is not BB if the UMT property does not hold, as shown in Figure 3. There, $VTR = PBV$ for all agents, and it is easy to see that the total payments are $34 - 3 \cdot 13 < 0$.

OBSERVATION 4. *The KSM-TR auction is IC and IR, but not necessarily BB, if the UMT property does not hold.*

It is possible to regain BB by adding an explicit BB constraint to the TR auction, but this at the expense of IC.

LEMMA 5. *To prove incentive compatibility of the KSM-TR auction, we can assume, without loss of generality, that there is one bid-optimal and one possible TR allocation for any set of bids.*

PROOF. To prove the lemma, we use the concept of a ***randomized mechanism***, that is a probability distribution over a family of deterministic mechanisms, as introduced by Nisan and Ronen [9]. First, we show that KSM-TR is a randomized mechanism. Recall that KSM-TR assigns unique integers to all the agents, which in turn uniquely specifies which bid-optimal and TR allocation will be chosen from the alternates with the same bid value. Then every possible assignment of integers specifies a deterministic mechanism, and the probability distribution over integer assignments is a randomized mechanism.

Nisan and Ronen showed that if each deterministic mechanism is incentive compatible, then the randomized mechanism is incentive compatible. Therefore, we need only show that every deterministic instance of KSM-TR is incentive compatible. But showing this is equivalent proving incentive compatibility under the assumption that there is a single bid-optimal and single TR allocation for any given set of bids. ☐

With Lemma 5, we assume in the sequel that there is a single bid-optimal allocation and single TR allocation with respect to any set of reported values.

In the following, we denote by $\check{v}^1$ the set of reported values when i bids $\check{v}_i^1$, that is $\check{v}^1 = (\check{v}_i^1, \check{v}_{-i})$. Similarly $\check{v}^2 = (\check{v}_i^2, \check{v}_{-i})$.

LEMMA 6. *If $i \in \tilde{A}(\check{v}^1)$ and $i \in \tilde{A}(\check{v}^2)$ for some agent i, then $VCG_i(\check{v}^1) = VCG_i(\check{v}^2)$ and $\tilde{A}(\check{v}^2) = \tilde{A}(\check{v}^1)$.*

PROOF. That $VCG_i(\check{v}^1) = VCG_i(\check{v}^2)$ follows directly from the well-known fact that VCG payments are IC when agents receive the bid-optimal allocation. It follows then that $\tilde{A}(\check{v}^2) = \tilde{A}(\check{v}^1)$ because there is only one optimal allocation. □

Denote as $\tilde{A}_i$ the optimal allocation containing i. When such an allocation exists, we define $VCG_i = \mathbf{V}(\tilde{A}(\check{v}_{-i})) - \mathbf{V}_{-i}(\tilde{A}_i)$. ($\tilde{A}_i$ is uniquely defined by Lemma 6.)

LEMMA 7. *If there exists a feasible allocation containing i, agent i is in $\tilde{A}(\check{v})$ if $\check{v}_i > VCG_i$ and i is not in $\tilde{A}(\check{v})$ if $\check{v}_i < VCG_i$.*

PROOF. If $\check{v}_i > VCG_i$ then $\check{v}_i > VCG_i = \mathbf{V}(\tilde{A}(\check{v}_{-i})) - \mathbf{V}_{-i}(\tilde{A}_i)$, giving us $\mathbf{V}(\tilde{A}(\check{v})) = \check{v}_i + \mathbf{V}_{-i}(\tilde{A}_i) > \mathbf{V}(\tilde{A}(\check{v}_{-i}))$. Thus $\tilde{A}_i$ must be optimal, hence $i \in \tilde{A}(\check{v})$ when $\check{v}_i > VCG_i$. If $\check{v}_i < VCG_i$ then $\check{v}_i < \mathbf{V}(\tilde{A}(\check{v}_{-i})) - \mathbf{V}_{-i}(\tilde{A}_i)$. Thus $\tilde{A}_i$ is not optimal, hence $i \notin \tilde{A}(\check{v})$ when $\check{v}_i < VCG_i$. □

LEMMA 8. *If $i \in A^{TR}(\check{v}^1)$ and $i \in A^{TR}(\check{v}^2)$ for some agent i, then $VTR_i(\check{v}^1) = VTR_i(\check{v}^2)$ and $A^{TR}(\check{v}^1) = A^{TR}(\check{v}^2)$.*

PROOF. Assume, wlog, that $\check{v}_i^2 > \check{v}_i^1$. By Lemma 6, $VCG_i = VCG_i(\check{v}^1) = VCG_i(\check{v}^2)$, hence $\mathbf{V}(A^{TR}(VCG_i(\check{v}^1), \check{v}_{-i})) = \mathbf{V}(A^{TR}(VCG_i(\check{v}^2), \check{v}_{-i}))$. It remains to show that $\mathbf{V}_{-i}(A^{TR}(\check{v}^2)) = \mathbf{V}_{-i}(A^{TR}(\check{v}^1))$, to prove $VTR_i(\check{v}^1) = VTR_i(\check{v}^2)$.

Since i is in the TR allocation, it is also in the bid optimal allocation. So by Lemma 7, $\tilde{A}(\check{v}^2) = \tilde{A}(\check{v}^1)$, hence both TR allocations are chosen from the same bid optimal allocation. Then, since $A^{TR}(\check{v}^1)$ clearly satisfies all the auction constraints when i bids $\check{v}_i^2$, and since $A^{TR}(\check{v}^2)$ is the optimal TR allocation when i bids $\check{v}_i^2$

$$\mathbf{V}(A^{TR}(\check{v}^2)) \geq \mathbf{V}^{\check{v}^2}(A^{TR}(\check{v}^1)) = \mathbf{V}(A^{TR}(\check{v}^1)) - \check{v}_i^1 + \check{v}_i^2.$$

Subtracting $\check{v}_i^2$ from both sides, we have $\mathbf{V}_{-i}(A^{TR}(\check{v}^2)) \geq \mathbf{V}_{-i}(A^{TR}(\check{v}^1))$. Now we need to show that $\mathbf{V}_{-i}(A^{TR}(\check{v}^2)) \leq \mathbf{V}_{-i}(A^{TR}(\check{v}^1))$.

Assume, contrary to which we wish to prove, that $\mathbf{V}_{-i}(A^{TR}(\check{v}^2)) > \mathbf{V}_{-i}(A^{TR}(\check{v}^1))$. If $A^{TR}(\check{v}^2)$ satisfies all the auction constraints when i bids $\check{v}_i^1$ then since $A^{TR}(\check{v}^1)$ is optimal with respect to $\check{v}^1$

$$\mathbf{V}(A^{TR}(\check{v}^1)) \geq \mathbf{V}^{\check{v}^1}(A^{TR}(\check{v}^2)) = \mathbf{V}(A^{TR}(\check{v}^2)) - \check{v}_i^2 + \check{v}_i^1.$$

Subtracting $\check{v}_i^1$ from both sides gives us $\mathbf{V}_{-i}(A^{TR}(\check{v}^1)) \geq \mathbf{V}_{-i}(A^{TR}(\check{v}^2))$ which is a contradiction. If, on the other hand, $A^{TR}(\check{v}^2)$ does not satisfy all the auction constraints, then the only constraint that could be violated is that $\check{v}_j > \check{v}_i^1$, where j is the price bounding agent in $K(i) \cup A^{TR}(\check{v}^2)$. Consider allocation $A' = (A^{TR}(\check{v}^2) \setminus \{i\}) \cup \{j\}$. A' satisfies all the auction constraints when i bids $\check{v}_i^1$. So

$$\begin{aligned}\mathbf{V}(A^{TR}(\check{v}^1)) \geq \mathbf{V}^{\check{v}^1}(A') &= \mathbf{V}(A^{TR}(\check{v}^2)) - \check{v}_i^2 + \check{v}_j \\ &> \mathbf{V}(A^{TR}(\check{v}^2)) - \check{v}_i^2 + \check{v}_i^1.\end{aligned}$$

Subtracting $\check{v}_i^1$ from both sides gives us $\mathbf{V}_{-i}(A^{TR}(\check{v}^1)) \geq \mathbf{V}_{-i}(A^{TR}(\check{v}^2))$, contradicting our assumption. Thus $\mathbf{V}_{-i}(A^{TR}(\check{v}^2)) = \mathbf{V}_{-i}(A^{TR}(\check{v}^1))$, giving us $VTR_i(\check{v}^1) = VTR_i(\check{v}^2)$. Also, since there is only one TR allocation, we have $A^{TR}(\check{v}^1) = A^{TR}(\check{v}^2)$.

□

Denote as A_i^{TR} the optimal TR allocation containing agent i. When such an allocation exists, we define $VTR_i = \mathbf{V}(A^{TR}(\check{v}^{VCG_i})) - \mathbf{V}_{-i}(A_i^{TR})$. ($VTR_i$ is uniquely defined by Lemma 8.)

LEMMA 9. *If there exists a feasible TR allocation containing i, agent i wins the TR auction if $\check{v}_i > VTR_i$, and loses the TR auction if $\check{v}_i < VTR_i$.*

PROOF. First, we establish that $VTR_i \geq VCG_i$. Assume, to the contrary, that $VTR_i < VCG_i$, then $\mathbf{V}(A^{TR}(\check{v}^{VCG_i})) - \mathbf{V}_{-i}(A^{TR}(\check{v})) < VCG_i$. So $\mathbf{V}(A^{TR}(\check{v}^{VCG_i})) < \mathbf{V}_{-i}(A^{TR}(\check{v})) + VCG_i$. Because $A^{TR}(\check{v}^{VCG_i})$ is optimal when i bids VCG_i, we have $\mathbf{V}(A^{TR}(\check{v}^{VCG_i})) \geq \mathbf{V}_{-i}(A^{TR}(\check{v})) + VCG_i$, which is a contradiction.

Now, we prove that if $\check{v}_i > VTR_i$ then $i \in A^{TR}(\check{v})$. By the above, $\check{v}_i > VTR_i \geq VCG_i$, so $i \in \tilde{A}(\check{v})$, and since the value of A_i^{TR} must be weakly monotonic with the bid of i, $\mathbf{V}(A_i^{TR}) \geq \mathbf{V}(A^{TR}(\check{v}^{VCG_i}))$. Assume, contrary to which we wish to prove, that i loses the auction. So $\mathbf{V}(A_i^{TR}) < \mathbf{V}(A^{TR}(\check{v}))$. Since $\check{v}_i > VCG_i$, i must also lose with bid VCG_i, hence $\mathbf{V}(A^{TR}(\check{v})) = \mathbf{V}(A^{TR}(\check{v}^{VCG_i}))$. It follows that $\mathbf{V}(A_i^{TR}) < \mathbf{V}(A^{TR}(\check{v}^{VCG_i}))$, which is a contradiction.

Finally, we prove that if $\check{v}_i < VTR_i$ then $i \notin A^{TR}(\check{v})$. By Lemma 7, if $\check{v}_i < VCG_i$ then $i \notin \tilde{A}(\check{v})$ and therefore $i \notin A^{TR}(\check{v})$. Now, consider the case where $VCG_i \leq \check{v}_i < VTR_i$. Assume, contrary to which we wish to prove, that $i \in A^{TR}(\check{v})$. Then $A^{TR}(\check{v}) = A_i^{TR}$, hence $\check{v}_i < VTR_i = \mathbf{V}(A^{TR}(\check{v}^{VCG_i})) - \mathbf{V}_{-i}(A_i^{TR})$, giving us $\check{v}_i + \mathbf{V}_{-i}(A_i^{TR}) < \mathbf{V}(A^{TR}(\check{v}^{VCG_i}))$. Also, since $VCG_i \leq \check{v}_i$, we have $\mathbf{V}(A^{TR}(\check{v}^{VCG_i})) \leq \mathbf{V}(A^{TR}(\check{v}))$. Therefore

$$\mathbf{V}(A^{TR}(\check{v})) = \check{v}_i + \mathbf{V}_{-i}(A_i^{TR}) < \mathbf{V}(A^{TR}(\check{v}^{VCG_i})) \leq \mathbf{V}(A^{TR}(\check{v}))$$

which is a contradiction.

□

LEMMA 10. *The KSM-TR auction is individually rational.*

PROOF. We must prove that an agent receives non-negative utility from bidding truthfully. If i loses, it pays zero and has zero utility. If i wins the auction by bidding truthfully, then by Lemma 9, $v_i \geq VTR_i$, hence its utility is $v_i - P_i = v_i - VTR_i \geq 0$ □

LEMMA 11. *The KSM-TR auction is incentive compatible in dominant strategies.*

PROOF. Consider the case in which agent i wins the auction by bidding its true value. If i bids untruthfully and loses, then it gets zero utility, which by Lemma 10 cannot be better than its utility with a truthful bid. If i bids untruthfully and wins the auction, then by Lemma 8 its payment, and hence it's utility remains the same.

Now consider the case in which i loses the auction by bidding truthfully. Its utility is zero and $v_i \leq VTR_i$ by Lemma 9. If i bids untruthfully and loses, its utility remains zero. If i bids untruthfully and wins, its utility is $v_i - P_i = v_i - VTR_i \leq 0$.

In both cases, we have shown that an agent cannot improve its utility by bidding truthfully, thus proving the lemma. □

3.3 Efficiency Analysis

We have established that KSM-TR is IC, IR, and BB, but we also want acceptable efficiency. In this section, we establish a good worst-case bound on the efficiency of the auction. This bound is such that, as the minimum number of trades in any consumer market grows in a fixed topology with the property, the TR allocation converges to perfect efficiency.

DEFINITION 12 (EFFICIENCY OF AN AUCTION). *The* ***efficiency*** $\text{Eff}^{\text{AUC}}(v)$ ***of auction*** *AUC producing allocation A^{AUC} for agents with valuations v and efficient allocation A^* is*

$$\text{Eff}^{\text{AUC}}(v) = \frac{\mathbf{V}(A^{\text{AUC}})}{V(A^*)}.$$

.

If the auction can produce alternate allocations due to randomization, then the efficiency is the minimum over all possible allocations.

DEFINITION 13 (EFFICIENCY COMPETITIVE RATIO). *An* ***efficiency competitive ratio function*** *of auction AUC is a function* $\text{Ratio}^{\text{AUC}}(v)$ *such that* $\text{Eff}^{\text{AUC}}(v) \geq \text{Ratio}^{\text{AUC}}(v)$ *for any vector of valuations v.*

Because KSM-TR generates only positive-value allocations, the efficiency is always in the range $[0, 1]$, hence we establish a competitive ratio in this range also. The closer the competitive ratio is to one, the more efficient the auction.

For market m, we denote as $T_m(A^*)$ the number of winning agents (trade size) in market m in the allocation A^*. We denote by CM^* the set of consumer markets with non zero trade size in A^*.

THEOREM 14. *The following function is an efficiency competitive ratio function for the KSM-TR auction:*

$$Ratio^{KSM-TR}(v) = \min_{m \in CM^*} \frac{T_m(A^*) - 1}{T_m(A^*)}$$

if $A^* \neq \emptyset$ and

$$Ratio^{KSM-TR}(v) = 1$$

if $A^* = \emptyset$.

PROOF. Refer to Appendix A. □

Note that Theorem 14 gives a worse case bound which holds for any valuations of the agents, therefore it holds for any distribution of valuations. The bound is dependent only on the number of trades in the optimal allocation.

Recall that the efficiency of A^{TR} for Figure 1 is 0.59. In this supply chain, there are two trades in each consumer market in A^*, giving us $\text{Ratio}^{\text{KSM-TR}}(v) = 1/2$, which is indeed less than the actual efficiency.

Typically, our auction achieves higher efficiency than the competitive ratio. The efficiency can be significantly higher when there is a large difference between the value of the agents in the auction allocation and the value of the agents reduced (recall that the low-valued agents are reduced). For instance, consider a supply chain with two markets: a producer market $M1$ with no inputs and an output desired by consumers in market $M2$. If $T_{M2}(A^*) = 2$, then $\text{Ratio}^{\text{KSM-TR}}(v) = 1/2$. But if both producers in $M1$ have a value of 0, c_1 is the highest-value consumer, and c_2 is the second-highest-value consumer in $M2$, then $\text{Eff}^{\text{KSM-TR}}(v) = v_{c_1}/(v_{c_1} + v_{c_2})$. Clearly then, $\text{Eff}^{\text{KSM-TR}}(v) \to 1$ as $v_{c_1}/v_{c_2} \to \infty$.

Nevertheless, the competitive ratio is a tight worst-case bound, in the following sense. Given an optimal allocation, there exists a set of bids supporting the allocation that give efficiency arbitrarily close to the competitive ratio. We first give a lemma to help prove this claim.

LEMMA 15. *The number of agents in each market in the TR allocation is uniquely defined. For each consumer market m, exactly one agent in m and its associated procurement set is removed from A^* to obtain the TR allocation.*

PROOF. The lemma directly follows from Lemma 27. □

THEOREM 16. *Let A^* be the efficient allocation for agents A and some set of values. Then for any $\varepsilon > 0$, there exists a vector of values v for agents A with the same optimal allocation that gives the bound*

$$\text{Eff}^{\text{KSM-TR}}(v) \leq \text{Ratio}^{\text{KSM-TR}}(v) + \varepsilon.$$

PROOF. Let $\bar{m} = \arg\min_{m \in CM^*} (T_m(A^*) - 1)/T_m(A^*)$. Intuitively we can see the theorem is true when the value of the consumers in $\bar{m}$ is much higher than in all other consumer markets, making the consumers in $\bar{m}$ dominate the efficiency. More formally, we can construct the desired v as follows.

- All consumers not in A^* have zero value.
- All producers not in A^* have a cost of 1 (any cost that is larger then the value of the above consumers will do).
- All producers in A^* have zero value.
- For all consumers $c \in CM^* \setminus \bar{m}$ we set $v_c \leftarrow 1$.
- For all consumers $c \in \bar{m}$ we set $v_c \leftarrow w$ for some value w to be defined (i.e., all such consumers have the same value).

Note that any agent that was not in A^*, is not in the efficient allocation with the new vector of values v, and any agent in A^* remains in the efficient allocation. By Lemma 15 exactly one consumer need be reduced from each market in CM^* in order to satisfy all the conditions of the KSM-TR mechanism. Therefore the efficiency is:

$$\text{Eff}^{\text{KSM-TR}}(v) = \frac{(T_{\bar{m}}(A^*) - 1)w + \sum_{m \in CM^* \setminus \bar{m}} (T_m(A^*) - 1)}{T_{\bar{m}}(A^*)w + \sum_{m \in CM^* \setminus \bar{m}} T_m(A^*)}$$

Hence,

$$\lim_{w \to \infty} \text{Eff}^{\text{KSM-TR}}(v) = \frac{T_{\bar{m}}(A^*) - 1}{T_{\bar{m}}(A^*)} = \text{Ratio}^{\text{KSM-TR}}(v)$$

The theorem follows immediately. □

The Myseron-Satterthwaite impossibility theorem [8] (discussed in Section 1) holds, in particular, for the case of a single producer with no inputs wishing to sell one good to a single consumer. In this case, the impossibility theorem implies that no trade can occur if we want BB, IR, and IC. With this in mind, and using reasoning similar to that in the proof of Theorem 16, we can conclude that, when any consumer market has only one consumer in the efficient allocation, no auction can have better than a zero efficiency competitive ratio. Thus, KSM-TR gives the best possible competitive ratio in this case.

The competitive ratio for Theorem 14 does not hold when the UMT property does not hold. Consider a topology with two markets that supply the same good g. Market 1 has one agent that can produce g with a low cost L from some zero cost good k, and Market 2 has two agents that can each produce g with a high cost H. There are three agents, each with value $H+1$ in a consumer market that desire g. In the efficient allocation, all three items will be traded and the value of the efficient allocation is $3*(H+1)-(L+2*H)=(H-L+3)$. In the TR allocation one agent in each market will be reduced and the allocation value is $H+1-H=1$. The efficiency is then $1/(H-L+3)$, which can be arbitrarily close to zero as H grows. This violates the efficiency ratio of $2/3$ from the theorem for this non-UMT supply chain.

3.4 Computational Complexity and Distributed Implementation

Although computing the optimal allocation (and hence the TR allocation and VTR payments) is NP-hard, we can compute it in polynomial time for a *fixed* number of consumer markets. Denote as CM the set of consumer markets. A ***configuration*** specifies, for each $m \in CM$, the trade size in m.

THEOREM 17. *For a fixed number $|CM|$ of consumer markets, the TR auction is polynomial time computable in $|A|$.*

PROOF. For brevity, we present only the proof concept. By an argument similar to the one presented in Lemma 15, we claim that for a fixed configuration, the size of trade in every market is uniquely decided and polynomial time computable in $|A|$. Given the trade sizes, we simply pick the highest value agents in each market, which can clearly be done in time polynomial in $|A|$. Thus, for a fixed configuration, we can find the highest value feasible allocation in time polynomial in $|A|$.

The number of configurations is at most $|A|^{|CM|}$, which is polynomial in $|A|$ for fixed $|CM|$. Therefore, finding the optimal allocation can be done by enumerating all configurations and picking the highest value feasible allocation time polynomial in $|A|$. Given the optimal allocation, by Lemma 15, the TR allocation A^{TR} can be calculated by removing one procurement set for any consumer market with non empty trade in A^*. This is clearly polynomial time computable in $|A|$. Finally, because calculating payments requires computing a polynomial number of optimal and TR allocations, these calculations are also polynomial time computable in $|A|$. $\square$

The TR auction can also be implemented as a distributed protocol between markets, generalizing the protocol presented in Babaioff and Nisan [1]. Again, for a fixed number of consumer markets, this protocol will run in time polynomial in the number of agents. Each agent sends its bid to a mediator representing its market. Each market communicates with its input and output markets, and consumer markets also communicate with a single coordinator.

To compute an optimal allocation, each consumer market first sends the number bids in its market to the coordinator, which enumerates all configurations. For a given configuration, the trade size is propagated from the coordinator, to the consumer markets, then through the producer markets. In each market, if the trade size is t, the t highest bidding agents are chosen. The total value of these bids are propagated and summed along the paths from the producer markets to the consumer markets, then to the coordinator. A market propagates these values to at most one (arbitrarily chosen) market that uses its output. When the total value reaches the coordinator, it knows the optimal value that can be obtained for the configuration. This procedure is performed for each configuration, allowing the coordinator to choose the optimal allocation. The TR allocation is computed similarly, but, by Lemma 15, with one fewer trade in each consumer market for each configuration. The VCG and VTR payments are computed by repeated applications of the above procedure.

OBSERVATION 18. *For a fixed number $|CM|$ of consumer markets, the auction can be implemented as a distributed protocol with running time polynomial in $|A|$.*

4. AUCTIONS FOR THE UNKNOWN SINGLE-MINDED MODEL

In many situations it may not be reasonable to assume that an auction knows the bundle of interest of the agents. Now we consider the case where both an agent's bundle of interest and monetary valuation are private and independent of other agents. With this model, which we call the Unknown Single-Minded (USM) model, the auctions must elicit the bundle of interest information from the agents. An auction is ***incentive compatible in dominant strategies*** iff each agent has the incentive to report its bundle of interest, and its valuation thereof, truthfully. For the USM model, we sometimes need to use a weaker solution concept. An auction is ***Nash incentive compatible*** iff each agent has the incentive to report its bundle of interest, and its valuation thereof, truthfully, given that all other agents do so also.

4.1 USM-TR Auction Mechanisms

In a ***USM-TR auction***, each agent i reports a value $\check{v}_i$ and bundle of interest $\check{q}_i$, either of which may not be true, and uses the TR rules, but possibly with additional rules. We call USM-TR-Base the auction that simply executes KSM-TR after receiving the bids. Unfortunately, USM-TR-Base is not generally incentive compatible because, due to weak monotonicity of preferences, an agent may be able to gain by (untruthfully) reporting a bundle that contains its bundle of interest. For instance, consider the case in which we have some consumer a with $\hat{q}_a^g = 1$ for good g only, and we have another consumer b with the same bundle of interest except that $\hat{q}_b^k = 1$ for some good k such that $\hat{q}_a^k = 0$. Assume that a is the only agent in its true market. If a bids truthfully, it gets reduced if it is in the optimal allocation, hence gets zero utility. Assume further that $v_a > v_b$. Then if b is winning in its own market, a would win by reporting the bundle $\check{q}_a = \hat{q}_b$ with value $V_a(\hat{q}_b)$ to the auction. Since $V_a(\hat{q}_a) = V_a(\hat{q}_b)$, agent a would obtain a higher utility by misreporting its bundle of interest than by reporting truthfully. Nevertheless, we have established necessary and sufficient conditions for a USM-TR auction to be incentive compatible in dominant strategies, as specified in the next theorem.

THEOREM 19. *A USM-TR auction is incentive compatible in dominant strategies iff no agent i can improve its utility by reporting any other $\check{q}_i$ such that $\check{q}_i > \hat{q}_i$.*

PROOF. *Case if:* The proof of Lemma 11 holds in the USM-TR auction, hence we know that no agent has an incentive to misreport its bundle valuation, assuming that it reports its true bundle of interest. Thus we need only establish that no agent has an incentive to misreport its bundle of interest.

If $\check{q}_i^k < \hat{q}_i^k$ for agent i and any good k, then $V_i(\check{q}_i) \leq 0$. But by individual rationality, $V_i(\hat{q}_i) \geq 0$. Thus, no agent i can gain by reporting a bundle with any component smaller than in $\hat{q}_i$. Therefore, it is sufficient to establish that i cannot increase its utility by reporting $\check{q}_i > \hat{q}_i$ to establish incentive compatibility.

Case only if: True by definition of incentive compatibility. $\square$

OBSERVATION 20. *In a USM-TR auction, no producer can improve its utility by misrepresenting itself as a consumer and no* *consumer can improve its utility by misrepresenting itself as a producer.*

With this observation, in the sequel we treat consumers and producers separately.

By the UTM assumption, if a producer i misrepresents its bundle of interest unilaterally, i will be the only agent in its market and will lose. Thus, we can get Nash IC by limiting Theorem 19 to consumers only.

OBSERVATION 21. *A USM-TR auction is Nash incentive compatible, iff no consumer can improve its utility by unilaterally reporting any other $\check{q}_i$ such that $\check{q}_i > \hat{q}_i$.*

4.2 Nash Incentive Compatibility by Removing Bids

Consider an auction USM-TR-RB (for USM-TR Remove Bigger) which removes from consideration every bid q_a such that $q_a > q_b$ for some other bid q_b, where q_a and q_b are both consumer bids. At first glance it seems that, since q_a is removed, the auction satisfies Observation 21. However, if b reports q_a instead of its true bundle of interest $\hat{q}_b$, and if there is no other bid q_c such that $q_a > q_c$, then no q_a bids get removed. We can guarantee that all q_a bids get removed if there exists a bid q_c other than $\hat{q}_b$ such that $q_a > q_c$. If this holds, then the auction satisfies Observation 21.

OBSERVATION 22. *USM-TR-RB is Nash incentive compatible, if for any consumer b, if there exists a bid q_a such that $q_a > \hat{q}_b$ then there exists a bid q_c such that $q_a > q_c$.*

4.3 Dominant Strategies Incentive Compatibility by Merging Markets

We can ensure incentive compatibility in dominant strategies by *merging* markets, rather than removing bids. The USM-TR-Merge is a USM-TR auction that first accepts bids, and then, before performing trade reduction, merges the consumer bids as follows: for each consumer i and good g, replace its reported bundle $\check{q}_i$ with $\tilde{q}$ such that $\tilde{q}^k = \max_j \check{q}^k_j$, where j is a consumer. We call $\tilde{q}$ the consumers' joint bundle of interest. The auction similarly merges the reported bundles of all producers with the same output.

THEOREM 23. *USM-TR-Merge is incentive compatible in dominant strategies.*

PROOF. By Theorem 19, we need only show that an agent cannot increase its utility by bidding for a bundle greater than its bundle of interest. The bundle union rules and single-mindedness ensure that an agent does not gain a higher value by bidding for a greater bundle. It remains to prove that an agent's payment does not decrease by reporting a greater bundle.

First we consider consumers. Let $\hat{\tilde{q}}$ the consumers' joint bundle of interest when i bids $\hat{q}_i$ and assume that i reports $\check{q}_i > \hat{q}_i$. If $\tilde{q} = \hat{\tilde{q}}$, then the payment by i does not change. Now consider the case where $\tilde{q} > \hat{\tilde{q}}$. Since all consumers share the same joint bundle of interest, the winning consumers are simply those with the highest reported values. By Lemma 9, if i wins, its payment is the minimal bid value necessary to be in the TR allocation. Since $\tilde{q} > \hat{\tilde{q}}$, more inputs are needed than if i would report $\hat{q}_i$. Since the additional inputs incur additional cost, the minimal value for i to win is at least as high with report $\check{q}_i$ as with $\hat{q}_i$. We conclude that the payment of a consumer does not decrease by bidding a larger bundle. Using similar reasoning, we also conclude that a producer does not decrease its payment by reporting a larger bundle. □

Recall that, by the UTM assumption, one market produces any good. Thus, since producers bid truthfully in USM-TR-Merge, no producers will actually be merged. Still, although no producer markets are actually merged, the merging rule is still necessary to ensure IC in dominant strategies.

In general, it is ambiguous whether USM-TR-Merge would give higher or lower efficiency than USM-TR-Base with agents reporting truthfully. If all true consumer markets contain only one consumer, then there would be no trade without merging, hence merging could not make the allocation worse and might improve it. But if consumer markets contain multiple consumers, then merging markets could increase the costs of an allocation, giving it a lower value than without merging.

If each consumer desires exactly k units total of any goods, we can gain IC without merging any consumer markets. In USM-TR-Merge-kIC (USM-TR-Merge k Input Consumers) we merge producer markets but not consumer markets and reject all consumer bids for other than k units. With the k-unit restriction, no consumer can feasibly misrepresent itself as any other consumer.

OBSERVATION 24. *USM-TR-Merge-kIC is incentive compatible in dominant strategies.*

Since no merging is actually performed, our competitive ratio holds for USM-TR-Merge-kIC.

OBSERVATION 25. *The efficiency competitive ratio from Theorem 14 holds for USM-TR-Merge-kIC.*

5. DISCUSSION AND FUTURE WORK

We have presented auctions for supply chain formation that are incentive compatible, individually rational, and budget balanced. We are not aware of any other auctions with these properties and with comparably high efficiency for as broad a class of supply chain topologies we consider. Nevertheless, we believe there may be further opportunities for improving efficiency of the KSM-TR auction while maintaining the properties. Our current approach relies on the existence of multiple agents with the same bundles of interest to obtain high efficiency. We hope to find methods for lessening the dependence. It is also our hope that further study will provide insights into obtaining incentive compatibility and budget balance with higher efficiency in the unknown single minded model.

We are also interested in developing auctions for a broader class of agent utility functions, namely without the single minded restriction. Consider the following obvious variant of our auction to allow OR or XOR bids. We change the auction to allow agents to place OR or XOR bids, and include the OR and XOR constraints in the auction. We also change the VTR_i payments so that i's payment does not depend on its own bids. With these changes, an agent can manipulate the allocation in its favor by changing one of its bids, thus violating IC. Consider the case in Figure 2 where consumer A1 true preferences contain XOR components $13 in market CM1 and $35 in market CM3. If A1 bids truthfully, it will win non of its bids. If instead, A1 bids less than $28 in the CM3 market, it will win one unit of the good in market CM1 and pay less than $13, giving it a positive utility. We get the same phenomenon if the bid is OR instead or XOR. In either case, the auction is not incentive compatible. It seems that obtaining incentive compatibility for these and other more general utility functions will present interesting mechanism design challenges.

6. REFERENCES

[1] M. Babaioff and N. Nisan. Concurrent auctions across the supply chain. In *Third ACM Conference on Electronic Commerce*, pages 1–10, 2001.

[2] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, Fall 1971.

[3] S. de Vries and R. Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, to appear.

[4] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[5] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):1–26, 2002.

[6] J. K. MacKie-Mason and H. R. Varian. Generalized Vickrey auctions. Technical report, Dept. of Economics, Univ. of Michigan, July 1994.

[7] R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:434–450, 1992.

[8] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[9] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, April/May 2001.

[10] D. C. Parkes, J. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Seventeenth International Joint Conference on Artificial Intelligence*, pages 1161–1168, 2001.

[11] N. M. Sadeh, D. W. Hildum, D. Kjenstad, and A. Tseng. MASCOT: An agent-based architecture for dynamic supply chain creation and coordination in the internet economy. *Production Planning and Control*, 12(3), 2001.

[12] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[13] W. E. Walsh. *Market Protocols for Decentralized Supply Chain Formation*. PhD thesis, University of Michigan, 2001.

[14] W. E. Walsh, M. P. Wellman, and F. Ygge. Combinatorial auctions for supply chain formation. In *Second ACM Conference on Electronic Commerce*, pages 260–269, 2000.

APPENDIX

A. PROOFS

We prove Theorem 14 in this appendix. Before doing so, we present a number of definitions and lemmas necessary for the proof.

We denote by $S\{A'\}(m)$ the number of agents in market m in the procurement set $S\{A'\}$ and denote by $Markets(S\{A'\})$ the set of markets M such that $S\{A'\}(m) \neq 0$ for all $m \in M$.

DEFINITION 26 (PROCUREMENT SET TOPOLOGY). *Two procurement sets $S\{A'\}_1$, $S\{A'\}_2$ of the allocation A' are of the same* ***procurement set topology*** $\mathbf{S}$*, if for every market m, $S\{A'\}_1(m) = S\{A'\}_1(m)$. In this case we write $S\{A'\}_1, S\{A'\}_2 \in \mathbf{S}$. We say that procurement set topology* $\mathbf{S}$ *is* ***in the allocation*** *A' if there exist a procurement set $S\{A'\}$ of the topology* $\mathbf{S}$ *in A'. We denote as $\hat{\mathbf{S}}(\mathbf{A'})$ the set of all procurement set topologies* $\mathbf{S}$ *that are in A'.*

We denote as $\hat{\mathbf{S}}_T$ the set of all procurement set topologies that are possible with supply chain topology T.

We note that procurement sets topologies are independent of the actual allocation, given the topology of the supply chain. For a fixed supply chain topology, the set of procurement set topologies that are possible in that supply chain topology is fixed.

For any procurement set topology $\mathbf{S} \in \hat{\mathbf{S}}_\mathbf{T}$, we denote as $\mathbf{S}(m)$ the number of winners in market m in any procurement set of topology $\mathbf{S}$. $\mathbf{S}(m)$ is well defined. Since a procurement set topology is minimal, there is exactly one consumer market m for which $\mathbf{S}(m) > 0$. For that market $\mathbf{S}(m) = 1$, and we call m the ***market of the procurement set***.

For any procurement set topology $\mathbf{S} \in \hat{\mathbf{S}}_\mathbf{T}$, let $N^*(\mathbf{S})$ be the maximal number of disjoint procurement sets with the same topology $\mathbf{S}$ in the optimal allocation A^*. For any procurement set topology $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$, we have $N^*(\mathbf{S}) > 0$, and for any other topology $N^*(\mathbf{S}) = 0$. Similarly, let $N^{TR}(\mathbf{S})$ be the maximal number of disjoint procurement sets with the same topology $\mathbf{S}$ in the reduction allocation.

For any market m we define $T_m(A^*)$ to be the number of winning agents (trade size) in market m in the efficient allocation A^*, and $R_m(A^{TR})$ to be the number of winning agents in A^* that are losers in the KSM-TR reduction allocation in market m. The size of trade in market m in the KSM-TR allocation is therefore $T_m(A^{TR}) = T_m(A^*) - R_m(A^{TR})$.

We denote as CM^* the set of consumer markets with non zero trade size in the efficient allocation.

LEMMA 27. *Let v be any vector of agents values. Let $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$ be any procurement set topology in the efficient allocation A^* for v. Then:*

- *There is a one-to-one mapping of all $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$ to all consumer markets in CM^*. So we can mark those procurement set topologies by $\mathbf{S}^i$ for $i = 1, \ldots, |CM^*|$, and we assume that it is mapped to market m_i.*
- *For each procurement set topology $\mathbf{S}^i$, $N^*(\mathbf{S}^i) = N^{TR}(\mathbf{S}^i) + 1$.*
- *For each procurement set topology $\mathbf{S}^i$ and its mapped consumer market m_i, $N^*(\mathbf{S}^i) = T_{m_i}(A^*)$ and $N^{TR}(\mathbf{S}^i) = T_{m_i}(A^*) - 1$.*
- $$A^* = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)} S^i_j \text{ and } A^{TR} = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)-1} S^i_j$$

  where S^i_j is the j procurement set of topology $\mathbf{S}^i$, and $S^i_{N^*(\mathbf{S}^i)}$ is the lowest valuation agents of all agents in procurement sets of topology $\mathbf{S}^i$.

PROOF. First we show that there is a one-to-one mapping of all $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$ to all consumer markets in CM^*. Let m be the consumer market of the procurement set topology $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$. We map $\mathbf{S}$ to its consumer market m, and we need to show that there is only one procurement set topology $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$ that is mapped to m.

The Unique Manufacturing Technologies property of the supply chain causes that there is only one procurement set topology that has an agent in market m as we show below. The proof is by induction on the markets in reverse topological order. If a procurement set topology has one agent in consumer market m, it uniquely produces the number of goods that are needed to satisfy this agent, and since a single market produces each good, it uniquely sets the number of winners in each of those markets. This can be continued till the number of agents in each market is uniquely specified, hence there is only one procurement set topology with agents in market m. The proof is by induction on the markets in reverse topological order. Clearly the claim is true for By our construction, the desired mapping exists. We conclude that, in order to reduce one agent in market m, a procurement set of topology $\mathbf{S}$ must be reduced in the reduction allocation, therefore for any $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$, we have $N^*(\mathbf{S}) \geq N^{TR}(\mathbf{S}) + 1$.

To prove that for any $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$, $N^*(\mathbf{S}) \leq N^{TR}(\mathbf{S})+1$ we use the following observation. We claim that, for any procurement set S of topology $\mathbf{S}$ in A^*, we have $V(S) > 0$ from the point of view of the KSM-TR auction. To see this, assume to the contrary that $V(S) < 0$. Then, by removing this procurement set we increase the efficient allocation value, which is a contradiction. The auction never observes $V(S) = 0$ because it adds the 2^{-i} values to the bids, hence no subset of agents can have value exactly zero.

Now we prove that for any $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$, $N^*(\mathbf{S}) \leq N^{TR}(\mathbf{S})+1$. Assume that this is not true, then for some $\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)$, $N^*(\mathbf{S}) > N^{TR}(\mathbf{S})+1$, or equivalently $N^*(\mathbf{S}) \geq N^{TR}(\mathbf{S})+2$. Then it must be that at least two procurement sets of the same topology are reduced by the auction. Since $V(S) > 0$ for any procurement set S in A^*, we can add one of the reduced procurement sets of the same topology and increase the value of the allocation from the point of view of the auction. We can add one of the reduced procurement sets while ensuring that every market has a price bounding agent, since we still have one procurement set reduced and both procurement sets share the same set of markets. Also, by Lemma 29, we can maintain the budget balance constraint. Thus, we can add one of the reduced procurement sets while maintaining the constraints on a TR allocation, thus contradicting the requirement that KSM-TR maximizes the allocation value, subject to the constraints. Therefore the assumption is not true, and we have proven that $N^*(\mathbf{S}) \leq N^{TR}(\mathbf{S})+1$ and therefore $N^*(\mathbf{S}) = N^{TR}(\mathbf{S})+1$.

Any procurement set topology $\mathbf{S}$ contains a single agent in its consumer market m, hence $N^*(\mathbf{S}) = T_m(A^*)$. Since we have shown that $N^*(\mathbf{S}) = N^{TR}(\mathbf{S})+1$, then $N^{TR}(\mathbf{S}) = N^*(\mathbf{S})-1 = T_m(A^*)-1$.

From all the above we conclude that

$$A^* = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)} S^i_j \text{ and } A^{TR} = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)-1} S^i_j$$

where S^i_j is the j procurement set of topology $\mathbf{S}^i$, and $S^i_{N^*(\mathbf{S}^i)}$ is the lowest valuation agents of all agents in procurement sets of topology $\mathbf{S}^i$. □

LEMMA 28. *For any winning agent i, $VTR_i \geq PBV_i$.*

PROOF. Assume, to the contrary, that $VTR_i < PBV_i$. If i bids any value $\check{v}_i$ such that $VTR_i < \check{v}_i$, then i wins the auction by Lemma 9. In particular, i wins if it bids $VTR_i < \check{v}_i < PBV_i$. But by the auction rule $\check{v}_i \geq PBV_i$, which is a contradiction. □

LEMMA 29. *Let v be any vector of agents values. The KSM-TR allocation for v*

$$A^{TR} = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)-1} S^i_j$$

with KSM-TR payments is budget balanced.

PROOF. We denote by $V(S^i_j)$ the sum of valuations of all agents in procurement set S^i_j, and by $P(S^i_j)$ the sum of payments from all those agents. We denote by $Pay(A^{TR})$ the sum of payments of all agents in the reduction allocation.

Since $Pay(A^{TR}) = \sum_{i=1}^{|CM^*|} \sum_{j=1}^{N^*(\mathbf{S}^i)-1} P(S^i_j)$, we must show that $Pay(A^{TR}) \geq 0$ to prove that A^{TR} with KSM-TR payments is budget-balanced. To do this, it is sufficient to show that $P(S^i_j) \geq 0$ for $i = 1, \ldots, |CM^*|$ and $j = 1, \ldots, N^*(\mathbf{S}^i)-1$.

We can build a one-to-one mapping of agents from procurement set S^i_j to agents from procurement set $S^i_{N^*(\mathbf{S}^i)}$, since both procurement sets are of the same topology and have the same number of agents in each market.

Since $S^i_{N^*(\mathbf{S}^i)}$ is in the efficient allocation, it must be that $V(S^i_{N^*(\mathbf{S}^i)}) \geq 0$, else it could be removed from the efficient allocation to get a better allocation, which is a contradiction.

By Lemma 28 the payment P_k from each agent k in S^i_j is at least as high as the PBV_k. This agent has the highest value of all reduced agents in k's market. In particular PBV_k is higher then the valuation of the agent that agent k is mapped to in $S^i_{N^*(\mathbf{S}^i)}$. Hence, we conclude that $P(S^i_j) \geq \sum_{k \in S^i_j} PBV_k \geq V(S^i_{N^*(\mathbf{S}^i)}) \geq 0$, which is what we wanted to prove. □

We need some additional definitions to carry on with our proofs.

DEFINITION 30 (ALLOCATION PARTITION). *An **allocation partition** $P^{A'}$ of a feasible allocation A' is a partition $P^{A'}_1, P^{A'}_2, \ldots, P^{A'}_k$ of the agents in A'. The **size of the partition** is k. For any set $P^{A'}_i$, the **value**, $V(P^{A'}_i)$, of the set is $\sum_{i \in P^{A'}_i} v_i$*

We call an allocation A^{FTR} a ***feasible reduction allocation*** if it satisfies all constraints for a TR allocation, except that it possibly does not maximize value (Equation (2)).

DEFINITION 31 (GOOD PARTITION PAIR). *Given vector of agents values v, with efficient allocation A^*, we say that the allocation have a **good partition pair** P^*, P^{tr} if there exists a partition P^* for the efficient allocation A^* of size k, and a partition P^{tr} for a feasible reduction allocation A^{FTR} of size k, such that for any $i = 1, \ldots, k$,*

- $P^{tr}_i \subseteq P^*_i$.
- $V(P^*_i) \geq V(P^{tr}_i) \geq 0$.
- $\frac{V(P^{tr}_i)}{V(P^*_i)} \geq \min_{\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$.

For valuations with a good partition pair we can bound the efficiency of KSM-TR in the following way:

LEMMA 32. *Given vector of agents values v, with non-empty efficient allocation A^* which has a good partition pair P^*, P^{tr}, we have:*

$$Eff^{KSM-TR}(v) = \frac{V(A^{TR})}{V(A^*)} \geq \min_{\mathbf{S} \in \hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$$

PROOF. Let the good partition pair P^*, P^{tr} have size k. Since P^* is a partition of the efficient allocation, $V(A^*) = \sum_{i=1}^k V(P^*_i)$, and since P^{tr} is a partition of a feasible reduction allocation, $V(A^{TR}) \geq V(A^{FTR}) = \sum_{i=1}^k V(P^{tr}_i)$. Therefore

$$Eff^{KSM-TR}(v) = \frac{V(A^{TR})}{V(A^*)} \geq \frac{\sum_{i=1}^k V(P^{tr}_i)}{\sum_{i=1}^k V(P^*_i)}.$$

Since P^*, P^{tr} is a good partition pair, for every $i = 1, \ldots, k$ it is true that $V(P^*_i) \geq V(P^{tr}_i) \geq 0$. Therefore, we can apply Lemma 33 to get

$$Eff^{KSM-TR}(v) \geq \frac{\sum_{i=1}^k V(P^{tr}_i)}{\sum_{i=1}^k V(P^*_i)} \geq \min_{i=1}^{k} \frac{V(P^{tr}_i)}{V(P^*_i)}$$

Since P^*, P^{tr} is a good partition pair, for every $i = 1, \ldots, k$ it is true that $\frac{V(P_i^{tr})}{V(P_i^*)} \geq \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$, therefore

$$\begin{aligned} Eff^{KSM-TR}(v) &\geq \min_{i=1}^{k} \frac{V(P_i^{tr})}{V(P_i^*)} \geq \min_{i=1}^{k} \left(\min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})} \right) \\ &= \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})} \end{aligned}$$

□

LEMMA 33. *For any set of indexes m and pairs R_m and O_m such that $0 \leq R_m \leq O_m$ it is true that*

$$\frac{\sum_m R_m}{\sum_m O_m} \geq \min_m \left(\frac{R_m}{O_m} \right)$$

PROOF. Let k be the index of elements that minimize the ratio $\frac{R_m}{O_m}$. For every m $\frac{R_m}{O_m} \geq \frac{R_k}{O_k}$, therefore for every m, $O_k * R_m \geq R_k * O_m$.

Summing over m we get $O_k * (\sum_m R_m) \geq R_k * (\sum_m O_m)$. Hence, $\frac{\sum_m R_m}{\sum_m O_m} \geq \frac{R_k}{O_k} = \min_m \frac{R_m}{O_m}$, which is what we wanted to prove. □

From Lemma 32 we conclude that, if the efficient allocation has a good partition pair and there is no procurement set topology with a single procurement set of this topology in the efficient allocation, then we get a competitive ratio of at least $1/2$.

LEMMA 34. *Let v be any vector of agents values with efficient allocation A^*. A^* has a good partition pair.*

PROOF. By Lemma 27 the efficient allocation is constructed from procurement set topologies $\mathbf{S}^i$ for $i = 1, \ldots, |CM^*|$ such that

$$A^* = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)} S_j^i \text{ and } A^{TR} = \bigcup_{i=1}^{|CM^*|} \bigcup_{j=1}^{N^*(\mathbf{S}^i)-1} S_j^i$$

where S_j^i is the j procurement set of topology $\mathbf{S}^i$, and $S^i_{N^*(\mathbf{S}^i)}$ is the lowest valuation agents of all agents in procurement sets of topology $\mathbf{S}^i$. Observe that there are $N^{TR}(\mathbf{S}^i) = N^*(\mathbf{S}^i) - 1$ procurement sets of the i topology in the reduction allocation A^{TR}.

Let $P_i^* = \bigcup_{j=1}^{N^*(\mathbf{S}^i)} S_j^i$, and let $P_i^{tr} = \bigcup_{j=1}^{N^*(\mathbf{S}^i)-1} S_j^i$. We need to show that all three requirements for a good partition pair holds.

- $P_i^{tr} \subseteq P_i^*$: This is true by construction.
- $V(P_i^*) \geq V(P_i^{tr}) \geq 0$: Since $S^i_{N^*(\mathbf{S}^i)}$ is a procurement set, it has a non-negative value. Hence, $V(P_i^*) - V(P_i^{tr}) = V(S^i_{N^*(\mathbf{S}^i)}) \geq 0$. Since every procurement set has non-negative value, we also have $V(P_i^{tr}) \geq 0$.
- $\frac{V(P_i^{tr})}{V(P_i^*)} \geq \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$: First observe that $\frac{V(P_i^{tr})}{V(P_i^*)} = \frac{\sum_{j=1}^{N^*(\mathbf{S}^i)-1} V(S_j^i)}{\sum_{j=1}^{N^*(\mathbf{S}^i)} V(S_j^i)}$. Hence, by applying Lemma 35 we get

$$\frac{V(P_i^{tr})}{V(P_i^*)} \geq \frac{N^*(\mathbf{S}^i)-1}{N^*(\mathbf{S}^i)} = \frac{N^{TR}(\mathbf{S}^i)}{N^*(\mathbf{S}^i)} \geq \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}.$$

□

LEMMA 35. *Let $n \in \mathbf{Z}^+$, $m \in \{1, \ldots, n\}$, and $X_i \in \mathbf{R}^+$ for all $i \in \{1, \ldots, n\}$. If $X_i \geq X_m$ for all $i < m$ and $X_i \leq X_m$ for all $i > m$, then*

$$\frac{\sum_{i=1}^m X_i}{\sum_{i=1}^n X_i} \geq \frac{m}{n}$$

PROOF. The proof is by induction on n for any fixed m. For any n such that $n \geq m$ we prove the claim by induction on n.

If $n = m$ the claim is true since we have 1 on both sides of the inequality. Now assume that we have proven the claim for some n_0 such that $n_0 \geq m$, to prove the claim for $n_0 + 1$. By the induction hypothesis,

$$\frac{\sum_{i=1}^m X_i}{\sum_{i=1}^{n_0} X_i} \geq \frac{m}{n_0},$$

hence $n_0 \sum_{i=1}^m X_i \geq m \sum_{i=1}^{n_0} X_i$.

Since $X_i \geq X_m \geq X_{n_0+1}$ for all $i \leq m$, we have $\sum_{i=1}^m X_i \geq m X_{n_0+1}$. Using the induction hypothesis we get by summation

$$n_0 \sum_{i=1}^m X_i + \sum_{i=1}^m X_i \geq m \sum_{i=1}^{n_0} X_i + m X_{n_0+1}$$

therefore

$$\frac{\sum_{i=1}^m X_i}{\sum_{i=1}^{n_0+1} X_i} = \frac{\sum_{i=1}^m X_i}{\sum_{i=1}^{n_0} X_i + X_{n_0+1}} \geq \frac{m}{n_0+1}$$

which is what we wanted to prove. □

Finally, we are ready to prove the theorem.

Theorem 14 *Let v be any vector of agents values. The following is an efficiency competitive ratio function for the KSM-TR auction:*

$$Ratio^{KSM-TR}(v) = \min_{m\in CM^*} \frac{T_m(A^*)-1}{T_m(A^*)}$$

if $A^* \neq \emptyset$ and

$$Ratio^{KSM-TR}(v) = 1$$

if $A^* = \emptyset$.

PROOF. The second component of the competitive ratio is true by definition, hence we prove the first component. By Lemma 34 v has a good partition pair. By applying Lemma 32, the efficiency of KSM-TR satisfies the following:

$$Eff^{KSM-TR}(v) = \frac{V(A^{TR})}{V(A^*)} \geq \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$$

From Lemma 27 we know that there is a one-to-one mapping of procurement set topologies in the efficient allocation to consumer markets with non-zero trade. If procurement set topology $\mathbf{S}$ is mapped to market m, then $N^{TR}(\mathbf{S}) = T_m(A^{TR}) = T_m(A^*) - 1$ and $N^*(\mathbf{S}) = T_m(A^*)$.

So we conclude that

$$Eff^{KSM-TR}(v) = \frac{V(A^{TR})}{V(A^*)} \geq \min_{\mathbf{S}\in\hat{\mathbf{S}}(\mathbf{A}^*)} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})} = \min_{m\in CM^*} \frac{T_m(A^*)-1}{T_m(A^*)}.$$

Therefore

$$Ratio^{KSM-TR}(v) = \min_{m\in CM^*} \frac{T_m(A^*)-1}{T_m(A^*)} \leq Eff^{KSM-TR}(v),$$

which is what we wanted to prove. □