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, Jerusalem 91904, Israel

William E. Walsh

IBM T. J. Watson Research Center, 19 Skyline Drive, Hawthorne, New York 10532, USA

Abstract

We address the mechanism design problem of supply chain formation—the problem of negotiation mechanisms to coordinate the buying and selling of goods in multiple markets across a supply chain. Because effective negotiation strategies can be difficult to design for supply chains, we focus on incentive compatible auctions, in which the agents' dominant strategy is to simply report their private information truthfully. Unfortunately, with two-sided negotiation, characteristic of supply chains, it is impossible to simultaneously achieve perfect efficiency, budget balance, and individual rationality with incentive compatibility. To resolve this problem we introduce auctions that explicitly discard profitable trades, thus giving up perfect efficiency to maintain budget balance, incentive compatibility and individual rationality. We use a novel payment rule based on 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.

Key words: auction, supply chain formation, mechanism design, combinatorial exchange, incentive compatible, budget balance

Email addresses: `mosheb@cs.huji.ac.il` (Moshe Babaioff), `wwalsh1@us.ibm.com` (William E. Walsh).

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 [1,2], and was the subject of the 2003 Trading Agent Competition [3].

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 individual rationality, much recent effort has focused on combinatorial auctions [4] 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. [5] analyzed an IR and budget-balanced 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 [6]. In contrast, the well-known Vickrey-Clarke-Groves (VCG) auction [7–9] (also called the Generalized Vickrey Auction [10]) is incentive compatible 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 [11]. In response to this impossibility, Parkes et al. [12] explored double auction rules that minimize agents' incentives to misreport their values, but maintains BB and high (but not perfect) efficiency.

We exploit the fact that, despite the impossibility theorem, it *is* possible to attain incentive compatibility with any two of the three desirable properties (efficiency, individual rationality, budget balance) 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 [13] 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 [14] to obtain the properties.

In this paper, we extend ideas from Babaioff and Nisan's approach to introduce auctions that are incentive compatible, budget balanced, and individually rational for a broader class of supply chain formation problems. Our auctions use a novel Vickrey Trade Reduction pricing scheme, analogous to the classic VCG pricing, but giving BB with the Trade Reduction allocation our auctions produce. 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 2 we describe our model of the supply chain formation problem. In Section 3 we present the rules for the basic Trade Reduction auction. In Section 4 we discuss the computational issues. Although computing the auction is NP-hard, we describe an algorithm for computing the auction in polynomial time, given constraints on the consumer preference structure. We also present a distributed implementation of the algorithm. In Section 5 we show that, in the case when the auction correctly knows all agents' bundles of interest, but the monetary valuations are private to the agents, the auction is incentive compatible, individually rational, budget balanced, and has a good competitive ratio for allocative efficiency. In Section 6 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 7.

2 Supply Chain Formation Problem

In this section we describe the supply chain formation problem. First we describe our model of agent preferences, then we define allocations in a supply chain. Our goal in the next section is then to define a rule that chooses an allocation in a way that supports desirable economic properties. We are concerned with the one-time problem of computing the allocation. Readers interested in managing the dynamics of existing supply chains may refer to Kjenstad [15] for an extensive review.

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 with each good we indicate its discretization. For each market, the quantity of inputs needed by one agent are indicated next to the respective arrows. We assume that an agent can provide one unit (at the appropriate discretization) of its output good but may require multiple units of an input good. Borrowing a term from Lehmann et al. [16], 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 that agent. This can often be a reasonable assumption, for companies typically have an established way to produce a product.
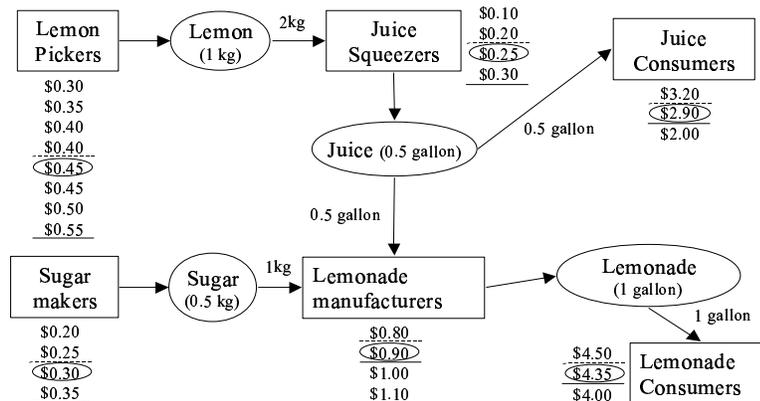

Fig. 1. An example supply chain graph in the lemonade industry.

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

4

and 0.5 gallons of juice to produce one gallon of lemonade, and a lemonade con-
sumer 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 [13] but is subsumed by the model described by Walsh and
Wellman [2].

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 bun-
dle $q = (q^1, \ldots, q^{|G|})$ indicates the quantity q^g of each good g exchanged by an
agent. Positive quantity indicates acquisition of a good (input), and negative quan-
tity 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 $\tilde{q} \geq q$ when $\tilde{q}^g \geq q^g$ for all g, and assert $\tilde{q} > q$ when $\tilde{q} \geq q$ and
$\tilde{q}^k > q^k$ 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** $\check{q}_i$
that it tries to obtain. The composition of a bundle of interest and an agent's valu-
ation thereof, depend on the class of the agent, as we detail below. We assume for
all agents i that $V_i(\check{q}_i) \in \mathbf{Z}$. For convenience, we subsequently denote $V_i(\check{q}_i)$ as $\check{v}_i$.
We denote as $\check{\mathbf{v}}$ the vector of agents' values and denote as $\check{\mathbf{q}}$ the vector of agents'
bundles of interest. The **market** $K(i)$ of agent i is the set of agents with exactly the
same bundle of interest, formally defined as $K(i) = \{j \mid \check{q}_j = \check{q}_i\}$.

---

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

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

- If $q \geq \check{q}_i$, then $V_i(q) = \check{v}_i$ (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: $\check{v}_i \leq 0$ and $\check{q}_i^g = -1$ 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 $q^g = \check{q}_i^g = -1$ and $q \geq \check{q}_i$, then $V_i(q) = \check{v}_i$ (single minded and weakly monotonic).
- Else, if $q^g = \check{q}_i^g = -1$ and $q^k < \check{q}_i^k$ where $\check{q}_i^k > 0$ 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 $\check{q}_i^k \geq 0$ 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).

We denote by *CM* the set of consumer markets and *PM* the set of producer markets.

Finally, we consider only supply chains with **unique manufacturing technologies (UMT)** property, in which there is only one market that produces any good. Formally, this means that if two producers have the same output, then they also have the same bundle of interest. 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 budget balance and our efficiency competitive ratio in our auction. However, the auction is incentive compatible and individually rational 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.*

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 $\mathbf{q}$ is ***feasible*** if and only if 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.

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: $A' = \bigcup_{i \in A | q_i \neq \emptyset} i$. We refer to the agents that receive their bundle of interest as the ***winning*** agents. The ***trade size*** $T_m(A')$ in market m of allocation A' is the number of agents in $A' \cap m$. We say that there is ***trade*** in market m in the allocation A', if the trade size in market m is positive.

The ***value*** $\mathbf{V}(A')$ ***of an allocation*** A' is the sum the agent values in A': $\mathbf{V}(A') \equiv \sum_{j \in A'} \breve{v}_j$. The value of an allocation A' excluding the value of agent i is $\mathbf{V}_{-i}(A') \equiv \sum_{j \in A', j \neq i} \breve{v}_j$. 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 $\breve{A}$ is ***efficient*** (optimal) if it is feasible and maximizes the value over all feasible allocations. The ***efficiency*** of allocation A' is $\frac{\mathbf{V}(A')}{\mathbf{V}(\breve{A})}$.

The efficient allocation $\breve{A}$ 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 $\breve{A}$ receives its bundle of interest. For instance, each of the two lemonade manufacturers in $\breve{A}$ require 1kg of sugar to produce its output, and there are four sugar makers in $\breve{A}$ 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.

3 Trade Reduction Auction

Here we present the rules for the Trade Reduction (TR) auction. The rules consist of an allocation rule, defined in Section 3.1 and a payment rule, defined in Section 3.2. We include some observations on the payment rule in Section 3.3.

3.1   Trade Reduction Auction Allocation

The auction is one-shot and sealed-bid. Each agent reports a bundle of interest q_i and a value v_i, either of which may not be truthful. The auction then computes an allocation, which assigns, for each agent, either its reported 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 rules in a manner based on an auction introduced by Babaioff and Nisan [13], but for a more general supply chain model. Conceptually, 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 incentive compatibility and budget balance, the auction removes some beneficial trades from the optimal allocation.

We define the ***bid-optimal allocation*** $A^*(\mathbf{v})$, as the feasible allocation that maximizes the sum of values with respect to the reported bids $\mathbf{v}$ (in the case that the reported bids are the true values, the bid-optimal allocation is efficient).

Definition 2 (Trade Reduction allocation) *The **Trade Reduction allocation** $A^{TR}(\mathbf{v})$ is a feasible, bid-optimal allocation but constrained to have strictly fewer winners than the bid-optimal allocation $A^*(\mathbf{v})$ in each market with non-zero trade.*

We say that the agents in $A^*(\mathbf{v})$ but not in $A^{TR}(\mathbf{v})$ are ***reduced***. (We explicitly include the bid valuations because it will be necessary to refer to different sets of bid valuations for computing the payments and proving incentive compatibility. When the actual set of bids is unambiguous or irrelevant, we simply denote the bid-optimal and TR allocation by A^* and A^{TR}, respectively.)

In the definition, we implicitly assume that there is exactly one bid-optimal allocation and exactly one Trade Reduction 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 incentive compatible 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. We show how to do so in a computationally efficient way in Section 4.

If the true values shown in Figure 1 are reported to the auction, then $\mathbf{V}(A^*(\mathbf{v})) = \mathbf{V}(\breve{A}) = \$7.90$ and $A^*(\mathbf{v})$ contains all agents with values and costs above the solid lines. All agents above the dashed lines are in the Trade Reduction allocation $A^{TR}(\mathbf{v})$ and $\mathbf{V}(A^{TR}(\mathbf{v})) = \$4.70$, giving an efficiency of 0.59. The TR rules require that we remove at least one agent from $A^*(\mathbf{v})$ for each market (all markets have winners), 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.

An equivalent, and perhaps more illuminating way of defining the Trade Reduction allocation is described below. We begin with an important definition.

Definition 3 (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 allocation.*

For example, in Figure 1, the following constitutes one procurement set: the juice consumer bidding \$2.90, the juice squeezer bidding \$0.25 and the two lemon pickers bidding \$0.50 and \$0.55. Note that since a producer can produce exactly one unit of one good, any procurement set has exactly one consumer. Clearly, any non-empty feasible allocation can be partitioned into procurement sets. In Figure 1, the reduced agents (indicated by bids between the solid and dashed lines) can be partitioned to two procurement sets, one including the reduced juice consumer and one including the reduced lemonade consumer.

The following lemma provides an alternate view of the Trade Reduction allocation. As we show in Section 4.2, it also allows us to construct an efficient algorithm for computing the TR allocation from the bid-optimal allocation.

Lemma 4 *The Trade Reduction allocation A^{TR} can be obtained from the bid-optimal allocation A^* by reducing a disjoint set of procurement sets as follows. For each consumer market m with trade in A^*, reduce exactly one procurement set containing one consumer in m. In each market, the agents reduced must have no higher value than the agents in A^{TR}.*

Proof. The lemma directly follows from Lemma 39. Note that Lemma 39 relies on the unique manufacturing technologies property. □

Corollary 5 (to Lemma 4) *Given A^*, the number of agents in each market in A^{TR} is uniquely defined and the number of agents reduced from each market is uniquely defined.*

3.2   Trade Reduction Auction Payments

Here we describe a new payment scheme to obtain incentive compatibility in our model. Losing agents pay zero. Each winning agent i pays the Vickrey Trade Re-

duction (VTR) value $VTR_i(\mathbf{v})$. We first present several definitions necessary to describe the VTR values of the agents.

We denote as $\mathbf{v} \equiv (v_i, \mathbf{v}_{-i})$ the vector of values reported by all agents, where $\mathbf{v}_{-i}$ is the vector of values reported by all agents except i. Let $A^*(\mathbf{v})$ be the bid-optimal allocation with respect to $\mathbf{v}$, and $A^*(\mathbf{v}_{-i})$ be the bid-optimal allocation with respect to $\mathbf{v}_{-i}$.

The ***Vickrey-Clark-Groves (VCG) payment*** from agent i with respect to the bids $\mathbf{v}$ is defined as

$$VCG_i(\mathbf{v}) \equiv \mathbf{V}(A^*(\mathbf{v}_{-i})) - \mathbf{V}_{-i}(A^*(\mathbf{v})) \tag{1}$$

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

As mentioned above, VCG payments are not budget balanced for supply chains, but we can extend the VCG idea to obtain BB payments in the Trade Reduction auction. We have done this with our new Vickrey Trade Reduction payment rule. The VTR payment rule is a normalized rule, so losing agents pay zero. Below we define the payment from a winning agent.

Definition 6 (Vickrey Trade Reduction (VTR) payment) *The Vickrey Trade Reduction payment from a winning agent $i \in A^{TR}$, with respect to the bids $\mathbf{v}$, is:*

$$VTR_i(\mathbf{v}) \equiv \mathbf{V}(A^{TR}(VCG_i(\mathbf{v}), \mathbf{v}_{-i})) - \mathbf{V}_{-i}(A^{TR}(\mathbf{v})) \tag{2}$$

where $A^{TR}(VCG_i(\mathbf{v}), \mathbf{v}_{-i})$ is the TR allocation obtained when the bid of i is replaced by $VCG_i(\mathbf{v})$.

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

3.3   Observations on the Payment Rule

The VTR payment scheme is a generalization of the payment scheme Babaioff and Nisan used in their supply chain formation auction [13]. Here, we refer to
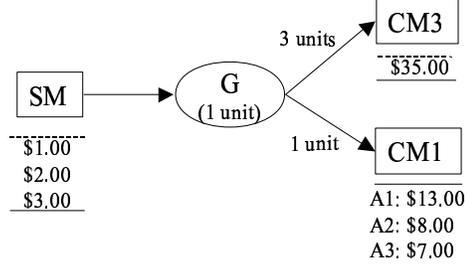
Fig. 2. Agent A1 does not win the Trade Reduction auction if it reports its true value of $13.00. It would win and pay $20.00 if it reports any value above $20.00.

the payments in their auction as the price bounding values (PBVs). The **price bounding value** PBV_i for i is the value reported by its price bounding agent, that is v_{PBA_i}. The **price bounding agent** $PBA_i(\mathbf{v})$ for winning agent i and bids $\mathbf{v}$ is $\arg\max_{j \in (A^*(\mathbf{v}) \backslash A^{TR}(\mathbf{v})) \cap K(i)} v_j$, that is the reduced agent with the highest bid in i's market. By Lemma 17 and Lemma 19, $PBA_i(\mathbf{v})$ is independent of i's bid when it wins, so we denote $PBA_i = PBA_i(\mathbf{v})$. We use the term "price bounding agent" because i pays at least PBV_i (Lemma 40) in our auction. As we show in Lemma 41, PBV payments are budget balanced, which means our auction is BB since i pays at least PBV_i. So, in effect, the price bounding agents serve as "cutoff points" to ensure that the payments from all agents above these points constitute BB.

Babaioff and Nisan's auction computes the Trade Reduction 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 incentive compatibility 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 $PBV_{A1} = \$7.00$.

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

The Vickrey Trade Reduction payment of any winning agent can be calculated from the VCG payment and the Price Bounding Value as follows.[2]

Lemma 7 $VTR_i(\mathbf{v}) = \max\{PBV_i, VCG_i(\mathbf{v})\}$ *for any agent* $i \in A^{TR}$.

Proof. By Lemma 20, the Vickrey Trade Reduction Payment for agent $i \in A^{TR}$ is the critical value for i to be in the Trade Reduction allocation. That is, if i bids above $VTR_i(\mathbf{v})$ then it wins the auction and if it bids below $VTR_i(\mathbf{v})$ it loses. Since there can be only one critical value for i, to prove the present lemma it is sufficient to show that $\max\{PBV_i, VCG_i(\mathbf{v})\}$ is also the critical value for i.

---

[2] We thank Zuo-Jun Shen for this observation.

If i bids below $VCG_i(\mathbf{v})$ it loses the auction, because by Lemma 18 it is not in the bid-optimal allocation $A^*(\mathbf{v})$, and therefore not in the Trade Reduction allocation.

If, on the other hand, i bids above $VCG_i(\mathbf{v})$, then no matter how high it bids, by Lemma 17, the bid-optimal allocation is the same. Assume then that i bids above $VCG_i(\mathbf{v})$. By Corollary 5, the number of reduced agents in i's market is uniquely defined for the bid-optimal allocation. This means that as long as agent i bids more than PBV_i, i will not be reduced. If i bids below PBV_i then it will be replaced by the price bounding agent in its market, and it will loses the auction.

We have shown that i is not in the bid-optimal allocation if it bids below $VCG_i(\mathbf{v})$ but is in the bid-optimal allocation if it bids above $VCG_i(\mathbf{v})$. We have also shown that, given that i is in the bid-optimal allocation, i is in the Trade Reduction allocation if it bids above PBV_i, but not in the TR allocation if it bids below PBV_i. Therefore, we have shown that $\max\{PBV_i, VCG_i(\mathbf{v})\}$ is the critical value for i to win the auction, and the lemma is proven. $\square$

Note that this Lemma relies on the unique manufacturing technologies property because its proof relies on Corollary 5, which also relies on the property.

4 Computing the Trade Reduction Auction

In this section we consider the computational aspects of the Trade Reduction auction. Computing the Trade Reduction auction requires computation of the Trade Reduction allocation and the Vickrey Trade Reduction payment for each agent. We show that the general problems of computing the bid-optimal and TR allocations are NP-hard. We show that given the bid-optimal allocation, we can find the TR allocation in polynomial time, so given the VCG allocation and payments, we can calculate the Trade Reduction allocation and payments in polynomial time. We also show that if the number of consumer markets is bounded by a constant, we can calculate the auction efficiently. Finally we show that the TR auction can be implemented as a protocol on a distributed markets infrastructure. We note that the computationally efficient algorithms rely crucially on Lemma 4, which in turn relies on the unique manufacturing technologies property.

Before continuing, recall that, when computing the bid-optimal and Trade Reduction allocation, the auction must break ties randomly between alternate optimal allocations, independent of reported valuations. An obvious, but very computationally inefficient algorithm for this tie breaking would be to enumerate all optimal allocations and randomly choose one. We adopt a more computationally efficient method which slightly changes the agents' bids in a preprocessing stage. 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 gives us 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 $\sum_{i\in N} 2^{-i} \neq \sum_{j\in M} 2^{-j}$, hence exactly one allocation is bid-optimal with respect to the modified bids. Similarly, this procedure obtains a unique, random TR allocation. Note that we do not include the 2^{-i} components in the agent payments.

4.1 Computational Complexity

To prove NP-hardness, we present a reduction transformation from the winner determination problem for the Combinatorial Auction with Single-Minded agents. We refer to this problem as CASM. In an instance of CASM there is a seller with a set G of heterogeneous goods, one item for each good. Each buyer i has a unique bundle of interest q_i with reported value $v_i \geq 0$ (the value of any bundle not containing q_i is zero). The winner determination problem is the problem of assigning the goods to the buyers such that the sum of the buyers reported values is maximized. Lehmann at el. [16] showed that CASM is NP-hard.

Theorem 8 *Computing the bid-optimal allocation in a supply chain is NP-hard.*

Proof. We prove the theorem by a polynomial reduction transformation from CASM. The seller is transformed to G producers, each producing a single item of a unique good with zero cost. Each buyer is transformed to a consumer with bundle of interest q_i and reported value v_i. Clearly, this reduction can be done in polynomial time. Also, it is clear that A^* is an optimal allocation for this supply chain if and only if it is a solution to CASM. Thus, an algorithm for computing the optimal allocation in supply chain can be used to solve CASM, hence our problem is at least as hard as CASM. $\square$

Although we defined the Trade Reduction allocation in terms of the optimal allocation, it is possible that the TR allocation could be computed without computing the optimal allocation first. Nevertheless, we can show directly that computing the TR allocation is NP-hard.

Theorem 9 *Computing the Trade Reduction allocation in a supply chain is NP-hard.*

Proof. We prove the theorem by a polynomial reduction transformation from CASM. The transformation is similar to the one described in Theorem 8, but it has additional fictive consumers, as well as producers that produces their inputs. Let B be the set of buyers for an instance of CASM and let $L = \sum_{i\in B} v_i + 1$. Observe that L is higher then the value of any allocation. Each buyer is transformed to a consumer

with bundle of interest q_i and value v_i. For each consumer market with bundle q_i, we add another (fictive) bid with value L. For each good $g \in G$ let n_g be the total number of items of good g that are demanded by the fictive agents. For each good g we add $n_g + 1$ sellers of the good g with zero cost. Clearly, this reduction can be done in polynomial time.

An optimal allocation for this supply chain must include all fictive agents, since all of them can be in the bid-optimal allocation A^* and each has a value greater than all the real consumers together. With goods assigned to all the fictive agents, one item of each good is left to be allocated to the real consumers. These remaining goods should be assigned to maximize the total value of the real consumers, hence the allocation A' to the real consumers is a solution to the instance of CASM. However, since we are actually computing the Trade Reduction allocation A^{TR}, and not necessarily A^*, we show how to recover A' from A^{TR} in polynomial time.

From Lemma 4, A^{TR} corresponds to A^*, except without the lowest price consumer in A^* in each consumer market with trade. Observe also that there can be at most two agents in any consumer market in A^*: one fictive consumer and possibly one real consumer. We conclude that a real consumer is in A' if and only if a fictive consumer in the same market is in A^{TR}. Thus, an algorithm for computing a Trade Reduction allocation in supply chain can be used to solve CASM, hence our problem is at least as hard as CASM. □

Despite these intractability results, we show in the next section that there are cases (e.g. when the number of consumer markets is constant) in which the auction can be calculated in time polynomial in the number of agents.

4.2 Algorithms

One approach to computing the bid-optimal and Trade Reduction allocations is to apply advanced integer programming (IP) techniques from operations research (e.g., by using commercial software such as CPLEX). It is straightforward to encode the problem of computing the bid-optimal allocation as an IP. For each agent i we have a variable $e_i \in \{0,1\}$ which indicates whether i receives its bundle of interest in the chosen bid-optimal allocation. The IP is then

$$
\begin{aligned}
&\text{maximize} \quad \sum_{i \in A} v_i e_i \\
&\text{such that} \sum_{i \in A} q_i^g e_i = 0, \ \text{ for each good } g.
\end{aligned} \tag{3}
$$

The second line in Eq.(3) constrains the goods to be in material balance. Agent feasibility is ensured because agents receive either their entire bundle of interest or nothing at all.

Based on our first definition of the Trade Reduction allocation, it would appear natural to use another integer program to compute the TR allocation, given the bid-optimal allocation. Instead, we describe a polynomial-time algorithm to compute the TR allocation A^{TR} from the bid-optimal allocation A^* by removing the reduced agents as described in Lemma 4. Procedure `Trade-Reduction(A*)` (Figure 3) implements the algorithm by computing the trade size $T_m(A^{TR})$ for each market m in A^{TR}. Procedure `Producer-Market-Trade-Sizes` (Figure 4) actually computes and returns the trade sizes for the producers in reverse topological order (ordered from the consumers). Its argument $\{t_m\}_{m \in CM}$ is a **configuration**, which specifies, for each consumer market m, the trade size in m. When the trade size $T_m(A^{TR})$ is computed for market m, the $T_m(A^{TR})$ highest value agents are included in A^{TR}. In the procedures, we denote as I_g the set of markets that desire good g as an input. Recall that we denote as $\breve{q}_n^g$ the number of items of good g desired by each agent in market n, so $\breve{q}_n^g \cdot t_n$ is the total number of items of good g desired by all agents in market n.

Theorem 10 *Procedure `Trade-Reduction(A*)` computes the Trade Reduction allocation A^{TR} from A^* in time polynomial in the number of agents $|A|$.*

Proof. Since the supply chain graph is acyclic, we can order the markets in a topological order. By an inductive argument on the market in reverse topological order, the number of winners in each market producing good g equals the number of items required by the markets consuming g. Therefore the allocation is feasible. The procedure sets the trade size in each consumer market with trade to the maximal possible trade size (one less than the trade in the bid-optimal allocation). From Lemma 4 we know that these trade sizes maximize efficiency under the feasibility constraint, while ensuring that at least one agent is reduced from any market with positive trade in A^*. The procedure then maximizes efficiency by picking the highest value agent in each market, according to the trade sizes.

The procedure runs in polynomial time in $|A|$, since sorting the agents in markets (to choose the highest value agents) can be performed in time $O(|A| \lg |A|)$ and the rest of the procedure operations can all be performed in total time $O(|A|)$. $\square$

From the above theorem and Lemma 7 it is clear that if we can calculate the VCG mechanism (allocation and payments) then from the VCG allocation and payments we can calculate the Trade Reduction allocation and payments in time polynomial in the number of agents.

Observation 11 *The Trade Reduction mechanism can be calculated from the VCG mechanism in time polynomial in the number of agents.*

```
Trade-Reduction(A*)
    FOR EACH market m, Set T_m(A^{TR}) ← 0.
    FOR EACH consumer market m ∈ CM
        IF T_m(A*) > 0 THEN set T_m(A^{TR}) ← T_m(A*) − 1.
    {T_m(A^{TR})}_{m∈PM} ← Producer-Market-Trade-Sizes({T_m(A^{TR})}_{m∈CM}).
    A^{TR} ← ∅.
    FOR EACH market m, add the T_m(A^{TR}) highest value agents to A^{TR}.
    RETURN A^{TR}.
```

Fig. 3. Procedure `Trade-Reduction`.

```
Producer-Market-Trade-Sizes({t_m}_{m∈CM})
    For each producer market m ∈ PM producing good g, in reverse
        topological order
        For each market n in I_g, set t_m ← t_m + q̆_n^g · t_n.
    RETURN {t_m}_{m∈PM}.
```

Fig. 4. Procedure `Producer-Market-Trade-Sizes`.

Still, because computing the bid-optimal allocation is NP-hard, the Trade Reduction auction remains NP-hard in general. Andersson et al. [17] show that the commercial integer programming software CPLEX can be an effective method for computing the allocation for a combinatorial auction specified as an integer program, and can be faster and more expedient than special purpose algorithms. However, for our auction we can exploit the structure of a unique manufacturing technologies supply chain to compute the bid-optimal and the Trade Reduction allocations in polynomial time, if certain constraints hold on the structure of consumer preferences. Denote as N_{max} the maximal number of consumers in any consumer market.

Theorem 12 *If there are constants c,k such that $N_{max}^{|CM|} \leq c|A|^k$ then the Trade Reduction auction is polynomial-time computable in $|A|$.*

Proof. For brevity, we present only the proof concept. For any given configuration, we can apply procedure `Producer-Market-Trade-Sizes` on the configuration and calculate the size of trade in every market in a feasible allocation with this configuration (this can be done time polynomial 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 $N_{max}^{|CM|}$, which is polynomial in $|A|$ by the assumption that $N_{max}^{|CM|} \leq c|A|^k$. Therefore, finding the optimal allocation can be done by enumerating all configurations and picking the highest value feasible allocation in time polynomial in $|A|$. Given the bid-optimal allocation,

we can then compute the Trade Reduction allocation in polynomial time using the `Trade-Reduction` procedure, as described above. Finally, because calculating payments requires computing a polynomial number of optimal and TR allocations, these calculations are also polynomial-time computable in $|A|$. $\square$

Note that $N_{max} \leq |A|$, therefore if the number of consumer markets is bounded by some (fixed) constant k, then $N_{max}^{|CM|} \leq |A|^k$ and by the theorem the algorithm is polynomial-time computable in $|A|$. So we have the following corollary.

Corollary 13 *For the family of supply chains with the number of consumer markets bounded by a fixed constant, the Trade Reduction auction is polynomial-time computable in $|A|$.*

Also note that if the maximal number of consumers in a market N_{max} is bounded by a constant k_1, and there is a constant k_2 such that $|CM| \leq k_2 log|A|$ (the number consumer market is logarithmic in the number of agents) then $N_{max}^{|CM|} \leq k_1^{k_2 log|A|} \leq k_1^{k_2}|A|^{|log|k_1|}$ and the algorithm is polynomial-time computable in $|A|$.

For some supply chains it is quite reasonable to assume that each good is manufactured from at least two units of its input goods. In this case the number of agents in a procurement set grows exponentially with the depth of the supply chain. In such supply chains, the maximal number of consumers in a market that can feasibly be in an allocation is indeed significantly smaller then the total number of agents, and we can reasonably expect the assumptions of Theorem 12 to hold.

4.3 Distributed Implementation

The Trade Reduction auction algorithm described in Theorem 12 can also be implemented as a distributed protocol between markets, generalizing the protocol presented in Babaioff and Nisan [13] (for a general overview of distributed algorithmic mechanism design, we refer the reader to Feigenbaum and Shenker [18]). Again, for a fixed number of consumer markets for example, this protocol will run in time polynomial in the number of agents.

Each agent communicates with a mediator representing its market. Each market communicates with its input and output markets, and consumer markets also communicate with a single coordinator. The mediators and coordinator are assumed to be obedient and run the protocol as specified (only the agents are strategic).

The protocol works as follows. Each consumer first sends its bid to its respective mediator. When a consumer mediator receives all bids from its consumers, it sends the *number* of bids in that market to the coordinator. When the coordinator receives the numbers of bids from each consumer mediator, it enumerates all configurations,

each specifying the trade size for each consumer market. Then, for each configuration, it sends the respective trade size to each consumer mediator. Recall that we denote the set of markets that desire good g as an input by I_g. When a mediator receives the number of units it needs to produce for each market in I_g (a consumer mediator treats the coordinator as demanding a virtual output from it) it performs `Market-Trade-Size-Propagate` (Figure 5). This procedure sums the number of output items demanded to determine the number of winners in the market m. It then propagates to each input market of m the number of input units needed by m from that input market. In the following, we denote by Υ_m the set of markets that produce input goods used by market m.

```
Market-Trade-Size-Propagate()
    WHEN q̆ⁿᵍ · tₙ is received from each market n in I_g, Set t_m ← Σ_{n∈I_g} q̆ⁿᵍ · tₙ.
    IF m has no input
        THEN perform Market-Value-Propagate().
    ELSE BEGIN
        FOR EACH market n ∈ Υ_m producing good k, send q̆_m^k · t_m to n.
        WHEN a value V_n is received from each market n ∈ Υ_n
            Perform Market-Value-Propagate().
    END
```

Fig. 5. Procedure `Market-Trade-Size-Propagate` executed by the mediator for market m with output good g.

When the trade sizes are computed, producers with no inputs initiate `Market-Value-Propagate` (Figure 6) propagates back to the coordinator the total value of the agents that will trade. The coordinator receives all the values from the consumer markets and sum them to find the total value of the allocation.

```
Market-Value-Propagate()
    Set V_m ← the sum of the bid values of the t_m highest agents in m.
    FOR EACH market n ∈ Υ_m, Set V_m ← V_m + V_n.
    Send V_m to one arbitrary market in I_g and send 0 to all other markets in I_g.
```

Fig. 6. Procedure `Market-Value-Propagate` executed by the mediator for market m with output good g.

The entire protocol described above is performed for each configuration, and the coordinator chooses the configuration with the highest value. Then, to determine the Trade Reduction allocation, the coordinator subtracts one from the optimal trade size for each consumer market and sends these trade sizes to the consumer mediators. The same protocol then computes the TR allocation in one round.

By the same argument as in Theorem 12 and from the above protocol we conclude the following observation.

Observation 14 *If there are constants c, k such that $N_{max}{}^{|CM|} \leq c|A|^k$ then the auction can be implemented as a distributed protocol with running time polynomial in $|A|$.*

5   Auction Properties With the Known Single-Minded Model

We first consider the properties of the Trade Reduction 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.

Under this model, we call the auction KSM-TR (Known Single-Minded Trade Reduction). Because the auction knows the bundles of interest, each agent only reports its valuation v_i, which may or may not be its true value $\check{v}_i$, to the auction. Under the KSM model, an auction is incentive compatible (IC) if and only if each agent has the incentive to report its true valuation for its desired bundle.

In this section we show certain desirable economic properties of the KSM-TR auction. First, we show that the auction is incentive compatible, individually rational, and budget balanced. Then, we show that it has a good competitive ratio for efficiency. Finally, we enumerate the properties that do and do not hold if we relax the unique manufacturing technologies constraint.

5.1   Incentive Compatibility, Individual Rationality, and Budget Balance

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

Theorem 15 *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 by definition. Incentive compatibility is proven in Lemma 22, individual rationality is proven in Lemma 21, and budget balance is proven in Lemma 41. □

Lemma 16 *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 Trade Reduction 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 intro-

duced by Nisan and Ronen [19]. Recall we can implement the KSM-TR auction by assigning 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. Thus, we need only show that every deterministic instance of KSM-TR is incentive compatible. But showing this is equivalent to proving incentive compatibility under the assumption that the bid-optimal and TR allocations are unique for any given set of bids. □

With Lemma 16, 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 $\mathbf{v}^1$ the set of reported values when i bids v_i^1, that is $\mathbf{v}^1 = (v_i^1, \mathbf{v}_{-i})$. Similarly $\mathbf{v}^2 = (v_i^2, \mathbf{v}_{-i})$.

Lemma 17 *If $i \in A^*(\mathbf{v}^1)$ and $i \in A^*(\mathbf{v}^2)$ for some agent i, then $VCG_i(\mathbf{v}^1) = VCG_i(\mathbf{v}^2)$ and $A^*(\mathbf{v}^2) = A^*(\mathbf{v}^1)$.*

Proof. That $VCG_i(\mathbf{v}^1) = VCG_i(\mathbf{v}^2)$ follows directly from the well-known fact that VCG payments are incentive compatible when agents receive the bid-optimal allocation (if $VCG_i(\mathbf{v}^1) \neq VCG_i(\mathbf{v}^2)$ then i could manipulate its payment by changing its bid, in which case VCG payments would not be incentive compatible). It follows then that $A^*(\mathbf{v}^2) = A^*(\mathbf{v}^1)$ because there is only one optimal allocation. □

Denote as A_i^* the optimal allocation containing i. When such an allocation exists, we define $VCG_i = \mathbf{V}(A^*(\mathbf{v}_{-i})) - \mathbf{V}_{-i}(A_i^*)$. ($A_i^*$ is uniquely defined by Lemma 17.)

Lemma 18 *If there exists a feasible allocation containing i, agent i is in $A^*(\mathbf{v})$ if $v_i > VCG_i$ and i is not in $A^*(\mathbf{v})$ if $v_i < VCG_i$.*

Proof. If $v_i > VCG_i$ then $v_i > VCG_i = \mathbf{V}(A^*(\mathbf{v}_{-i})) - \mathbf{V}_{-i}(A_i^*)$, giving us $\mathbf{V}(A^*(\mathbf{v})) = v_i + \mathbf{V}_{-i}(A_i^*) > \mathbf{V}(A^*(\mathbf{v}_{-i}))$. Thus A_i^* must be optimal, hence $i \in A^*(\mathbf{v})$ when $v_i > VCG_i$. If $v_i < VCG_i$ then $v_i < \mathbf{V}(A^*(\mathbf{v}_{-i})) - \mathbf{V}_{-i}(A_i^*)$. Thus A_i^* is not optimal, hence $i \notin A^*(\mathbf{v})$ when $v_i < VCG_i$. □

Lemma 19 *If $i \in A^{TR}(\mathbf{v}^1)$ and $i \in A^{TR}(\mathbf{v}^2)$ for some agent i, then $VTR_i(\mathbf{v}^1) = VTR_i(\mathbf{v}^2)$ and $A^{TR}(\mathbf{v}^1) = A^{TR}(\mathbf{v}^2)$.*

Proof. Assume, wlog, that $v_i^2 > v_i^1$. By Lemma 17, $VCG_i = VCG_i(\mathbf{v}^1) = VCG_i(\mathbf{v}^2)$, hence $\mathbf{V}(A^{TR}(VCG_i(\mathbf{v}^1), \mathbf{v}_{-i})) = \mathbf{V}(A^{TR}(VCG_i(\mathbf{v}^2), \mathbf{v}_{-i}))$. It remains to show that $\mathbf{V}_{-i}(A^{TR}(\mathbf{v}^2)) = \mathbf{V}_{-i}(A^{TR}(\mathbf{v}^1))$, to prove $VTR_i(\mathbf{v}^1) = VTR_i(\mathbf{v}^2)$.

Since i is in the TR allocation, it is also in the bid optimal allocation. So by Lemma 18, $A^*(\mathbf{v}^2) = A^*(\mathbf{v}^1)$, hence both TR allocations are chosen from the same bid optimal allocation. Then, since $A^{TR}(\mathbf{v}^1)$ clearly satisfies all the auction constraints when i bids v_i^2, and since $A^{TR}(\mathbf{v}^2)$ is the optimal TR allocation when i bids v_i^2

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

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

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

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

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

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

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

Denote as A_i^{TR} the optimal TR allocation containing agent i. When such an allocation exists, we define $VTR_i = \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i})) - \mathbf{V}_{-i}(A_i^{TR})$. ($VTR_i$ is uniquely defined by Lemma 19.)

Lemma 20 *If there exists a feasible TR allocation containing i, agent i wins the TR auction if $v_i > VTR_i$, and loses the TR auction if $v_i < VTR_i$.*

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

Now, we prove that if $v_i > VTR_i$ then $i \in A^{TR}(\mathbf{v})$. By the above, $v_i > VTR_i \geq VCG_i$, so $i \in A^*(\mathbf{v})$ by Lemma 18. Then, since $v_i > VTR_i = \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i})) - \mathbf{V}_{-i}(A_i^{TR})$ it follows that $\mathbf{V}(A_i^{TR}) > \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i}))$. Assume, contrary to which we wish to prove, that i loses the auction, giving us $\mathbf{V}(A_i^{TR}) \leq \mathbf{V}(A^{TR}(\mathbf{v}))$. Since $v_i > VCG_i$, i

21

must also lose with bid VCG_i, hence $\mathbf{V}(A^{TR}(\mathbf{v})) = \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i}))$. It follows that $\mathbf{V}(A_i^{TR}) \leq \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i}))$, which is a contradiction.

Finally, we prove that if $v_i < VTR_i$ then $i \notin A^{TR}(\mathbf{v})$. By Lemma 18, if $v_i < VCG_i$ then $i \notin A^*(\mathbf{v})$ and therefore $i \notin A^{TR}(\mathbf{v})$. Now, consider the case where $VCG_i \leq v_i < VTR_i$. Assume, contrary to which we wish to prove, that $i \in A^{TR}(\mathbf{v})$. Then $A^{TR}(\mathbf{v}) = A_i^{TR}$, hence $v_i < VTR_i = \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i})) - \mathbf{V}_{-i}(A_i^{TR})$, giving us $v_i + \mathbf{V}_{-i}(A_i^{TR}) < \mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i}))$. Also, since $VCG_i \leq v_i$, we have $\mathbf{V}(A^{TR}(\mathbf{v}^{VCG_i})) \leq \mathbf{V}(A^{TR}(\mathbf{v}))$. Therefore

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

which is a contradiction. $\square$

Lemma 21 *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 by bidding truthfully, then by Lemma 20, $\check{v}_i \geq VTR_i$, hence its utility is $\check{v}_i - P_i = \check{v}_i - VTR_i \geq 0$ $\square$

Lemma 22 *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 21 cannot be better than its utility with a truthful bid. If i bids untruthfully and wins the auction, then by Lemma 19 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 $\check{v}_i \leq VTR_i$ by Lemma 20. If i bids untruthfully and loses, its utility remains zero. If i bids untruthfully and wins, its utility is $\check{v}_i - P_i = \check{v}_i - VTR_i \leq 0$.

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

5.2 Efficiency Analysis

We have established that Known Single-Minded Trade Reduction is incentive compatible, individually rational, and budget balanced, 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 Trade Reduction allocation converges to perfect efficiency.

Definition 23 (Efficiency of an auction) *The **efficiency** $\mathrm{Eff}^{\mathrm{AUC}}(\check{\mathbf{v}})$ **of an auction** AUC producing allocation A^{AUC} in equilibrium, for agents with valuations $\check{\mathbf{v}}$ is the*

efficiency of A^{AUC}. If the auction can produce alternate allocations due to randomization, then the efficiency is the minimum over all possible allocations.

Because our auction is incentive compatible, efficiency is measured with respect to an allocation produced by truthful bidding.

Definition 24 (Efficiency competitive ratio) *An **efficiency competitive ratio function** of auction AUC is a function $\text{Ratio}^{\text{AUC}}(\check{\mathbf{v}})$ such that $\text{Eff}^{\text{AUC}}(\check{\mathbf{v}}) \geq \text{Ratio}^{\text{AUC}}(\check{\mathbf{v}})$ for any vector of valuations $\check{\mathbf{v}}$.*

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.

We denote by CM^* the set of consumer markets with non zero trade size in $\check{A}$.

Theorem 25 *The following function is an efficiency competitive ratio function for the KSM-TR auction:*

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

if $\check{A} \neq \emptyset$ and

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

if $\check{A} = \emptyset$.

Proof concept. Here we present the basic intuition for the theorem. Refer to Appendix A for the full proof. Consider a simple supply chain with one consumer market wanting a single unit of a single good from one producer market. The Trade Reduction allocation reduces the lowest value buyer and seller, that is, the lowest value procurement set. If the trade size in the optimal allocation is $n > 0$, then the trade size in the TR allocation is $n - 1$. The efficiency is the lowest when all procurement sets (seller-buyer pairs) have the same value, in which case the efficiency is $(n-1)/n$.

More generally, in a supply chain with unique manufacturing technologies property, any two consumers of the same market belong to the same topology of procurement sets. That is, the same number of the same goods need to be produced for each consumer. Again, the lowest value procurement set is reduced in each consumer market. The consumer market m with the lowest trade size has the largest fraction of agents reduced. If the procurement sets of m also constitute most of the value in A^*, then the greatest value will be lost here and this reduction would dictate the efficiency loss. So, the guaranteed efficiency is the minimal over all consumer markets, of the ratio between the trade size in the TR allocation and the trade size in the optimal allocation. $\square$

23

Note that Theorem 25 gives a worst 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 $\check{A}$, giving us $\text{Ratio}^{\text{KSM-TR}}(\check{\mathbf{v}}) = 1/2$, 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 m_1 with no inputs and an output desired by consumers in market m_2. If $T_{m_2}(\check{A}) = 2$, then $\text{Ratio}^{\text{KSM-TR}}(\check{\mathbf{v}}) = 1/2$. But if both producers in m_1 have a value of 0, c_1 is the highest-value consumer, and c_2 is the second-highest-value consumer in m_2, then $\text{Eff}^{\text{KSM-TR}}(\check{\mathbf{v}}) = \check{v}_{c_1}/(\check{v}_{c_1} + \check{v}_{c_2})$. Clearly then, $\text{Eff}^{\text{KSM-TR}}(\check{\mathbf{v}}) \to 1$ as $\check{v}_{c_1}/\check{v}_{c_2} \to \infty$.

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.

Theorem 26 *Let $\check{A}$ be the efficient allocation for agents A and some set of values. Then for any $\varepsilon > 0$, there exists a vector of values $\check{\mathbf{v}}$ for agents A with the same optimal allocation that gives the bound*

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

Proof. Let $\bar{m} = \arg\min_{m \in CM^*}(T_m(\check{A}) - 1)/T_m(\check{A})$. 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 $\check{\mathbf{v}}$ as follows.

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

Note that any agent that was not in $\check{A}$, is not in the efficient allocation with the new vector of values $\check{\mathbf{v}}$, and any agent in $\check{A}$ remains in the efficient allocation. By Lemma 4, exactly one consumer need to be reduced from each market in CM^* in order to satisfy all the conditions of the KSM-TR mechanism. Therefore the

24

efficiency is:

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

Hence,

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

The theorem follows immediately. $\square$

The Myseron-Satterthwaite impossibility theorem [11] (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 budget balance, individual rationality, and incentive compatibility. With this in mind, and using reasoning similar to that in the proof of Theorem 26, 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.

5.3 Economic Properties without the Unique Manufacturing Technologies Constraint

The Trade Reduction auction allocation and payments rules can be applied even when the unique manufacturing technologies property does not hold. However, the alternate definition of the auction allocation described in Lemma 4 requires the unique manufacturing technologies (UMT) property. Additionally, the characterization of the VTR payments in Lemma 7 does not hold when the UMT property does not hold. Figure 7 shows an example for which this is the case. For any agent winning i in market SM, $PBV_i = -\$13.00$ and $VCG_i(\mathbf{v}) = -\$16.00$, but $VTR_i = -\$4.00 \neq -\$13.00 = \max\{PBV_i, VCG_i(\mathbf{v})\}$.
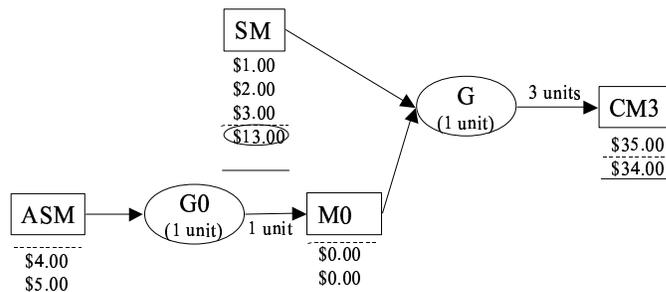

Fig. 7. A supply chain for which $VTR_i(\mathbf{v}) \neq \max\{PBV_i, VCG_i(\mathbf{v})\}$ (i.e., Lemma 7 does not hold because the unique manufacturing technologies property does not hold).

We proved NP-hardness for computing the auction for supply chains with the UMT property. The auction is then necessarily NP-hard when the UMT property does not hold. Unfortunately, because Lemma 4 (and hence Corollary 5) depends on the UMT property, the centralized and decentralized polynomial (given constraints on the consumers) algorithms cannot be used for non-UMT supply chains.

The auction is incentive compatible and individually rational for non-UMT supply chains because the proofs do not depend on the UMT property. However, the Trade Reduction auction is not budget balanced if the UMT property does not hold, as shown in Figure 8 (a variation of Figure 7). There, $VTR = PBV$ for all agents, and it is easy to see that the total payment is $34 - 3 \cdot 13 < 0$.

Observation 27 *The Trade Reduction auction is incentive compatible and individually rational, but not necessarily budget balanced, if the unique manufacturing technologies property does not hold.*

It is possible to regain budget balance by adding an explicit budget balance constraint to the TR auction, but this at the expense of incentive compatibility.
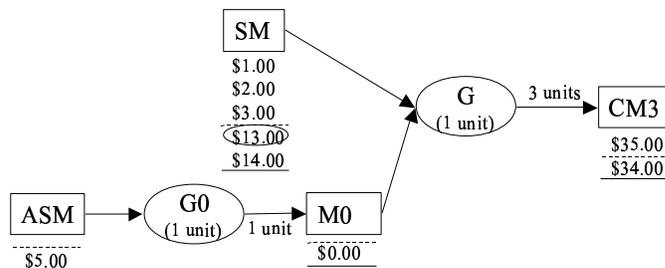

Fig. 8. A supply chain without the unique manufacturing technologies property and for which Trade Reduction auction is not budget balanced.

The competitive ratio for Theorem 25 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.

26

6 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*** if and only if each agent has the incentive to report its bundle of interest, and its valuation thereof, truthfully, regardless of what the other agents report. For the USM model, we sometimes need to use a weaker solution concept. An auction is ***ex post Nash incentive compatible*** if and only if each agent has the incentive to report its bundle of interest, and its valuation thereof, truthfully, given that all other agents also do so (independent of the distribution of the values of the other agents).

6.1 USM-TR Auction Mechanisms

Consider the Trade Reduction auction under the USM model, in which each agent i reports a value v_i and bundle of interest q_i, either of which may not be true. Unfortunately, the auction 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 $\breve{q}_a^g = 1$ for good g only, and we have another consumer b with the same bundle of interest except that $\breve{q}_b^k = 1$ for some good k such that $\breve{q}_a^k = 0$. 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 $\breve{v}_a > \breve{v}_b$. Then if b is winning in its own market, a would win by reporting the bundle $q_a = \breve{q}_b$ with value $V_a(\breve{q}_b)$ to the auction. Since $V_a(\breve{q}_a) = V_a(\breve{q}_b)$, agent a would obtain a higher utility by misreporting its bundle of interest than by reporting truthfully.

Since agents cannot obtain positive value by bidding untruthfully for smaller bundles, it appears that agents' incentive to untruthfully bid for larger bundles is the root cause of lack of incentive compatibility for the TR auction under the USM model. We formalize a variant of this notion by introducing a new class of auctions based on the Trade Reduction auction. In the next sections we present specific auctions from this class that remove the incentive for agents to untruthfully bid for larger bundles, giving incentive compatibility.

A ***USM-TR auction*** accepts bids from all agents, then applies a publicly known ***conversion function F*** to the bundles reported in the bids, to get a ***converted*** bundle for each agent. The auction then applies the Trade Reduction auction rules on the

converted bundles with the submitted values. The i's output of F is the converted bundle for agent i and is denoted by $F_i(q_i, \mathbf{q}_{-i})$. The first argument q_i is the bundle reported by i and the second argument $\mathbf{q}_{-i}$ is the vector of bundles reported by all agents other than i. Because we do not want F to discriminate between agents, we require that $F_i(q, \mathbf{q}_{-i}) = F_j(q, \mathbf{q}_{-j})$ when $\mathbf{q}_{-i}$ equals $\mathbf{q}_{-j}$ as *sets* of bundles (without considering the order of the bundles). For any agent i and for any $\mathbf{q}_{-i}$, F must also satisfy the following conditions for any bundle q_i:

(1) $q_i \leq F_i(q_i, \mathbf{q}_{-i})$.
(2) The original and converted bundles have the same output: for any good g, $F_i(q_i, \mathbf{q}_{-i})^g = -1$ if and only if $q_i^g = -1$.

Note that the above conditions allow a function that converts to bundles with a null good $\bar{g}$ that is not produced by any producer. This allows us to effectively remove the bid because an agent can never win a bundle containing $\bar{g}$. We use this notion in the next two sections.

Also note that any agent i has the same valuations for its desired bundle and for the converted bundle of its desired bundle, that is $V_i(\check{q}_i) = V_i(F_i(\check{q}_i, \mathbf{q}_{-i}))$. Since the Trade Reduction auction is individually rational, it follows that a USM-TR auction is individually rational.

An agent has no incentive to misreport its output good, that is, it has no incentive to report an output not in its true bundle or interest or to not report an output that is in its true bundle of interest.

Lemma 28 *In a USM-TR auction, for any $\mathbf{q}_{-i}$, and any reported values $\mathbf{v}$, an agent receives non-positive utility by misreporting its output good.*

Proof. Recall that a consumer cannot feasibly produce any goods, a producer cannot feasibly produce any good other than other than its true output, and infeasibility results in negative utility, regardless of the payment. Thus, since F does not change the output reported by any agent, no agent will have positive utility if it reports an output that is not its true output. It remains to prove that no producer has positive utility if it misrepresents itself as a consumer (by reporting no output). This is true since a producer has non positive value for obtaining any goods, and a USM-TR auction will not give it a positive payment unless it provides an output good (since this is true for to Trade Reduction auction applied on the converted bids). Thus, any agent receives non-positive utility by misreporting its output good. □

With the above lemma, we can assume that consumers will not misrepresent themselves as producers, and producers will not misrepresent themselves as consumers or as having any output good other than specified in their bundles of interest.

An agent has an incentive to report only bundles that convert to a bundle that contains its bundle of interest.

Lemma 29 *In a USM-TR auction, for any agent i, any $\mathbf{q}_{-i}$ and any reported values $\mathbf{v}$, if i does not report q_i such that $F_i(q_i, \mathbf{q}_{-i}) \geq \check{q}_i$ then i receives non-positive utility.*

Proof. First we consider consumers. By Lemma 28 we can assume that a consumer bids for a consumer's bundle, and the USM-TR rules specify that the bundle remains a consumer's bundle after applying the conversion function. By definition, consumer i receives positive value only for bundles q such that $q \geq \check{q}_i$, and the Trade Reduction auction never gives a consumer positive payment, hence the lemma holds for consumers.

Now, consider a producer. If it is not the case that $F_i(q_i, \mathbf{q}_{-i}) \geq \check{q}_i$, then it reports outputs it cannot produce, or it does not report all of the input units it needs to produce its true output. In the former case, it receives non-positive utility by Lemma 28, and in the latter case, it cannot be feasible, hence it receives non-positive utility, no mater how high is the payment it receives. Thus the lemma holds for producers. □

Now we can show that an agent has no incentive to misreport its value.

Lemma 30 *In a USM-TR auction, for any agent i, any $\mathbf{q}_{-i}$, any $\mathbf{v}_{-i}$, and any reported bundle q_i with the same output as $\check{q}_i$ such that $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$, i cannot improve its utility by misreporting its value.*

Proof. If q_i has the properties required by the lemma, we have $V_i(\check{q}_i) = V_i(F_i(q_i, \mathbf{q}_{-i}))$. Thus, once q_i is fixed, the incentive compatibility in dominant strategies of the Trade Reduction auction implies that i's best strategy is to report its value truthfully, regardless of $\mathbf{q}_{-i}$ and $\mathbf{v}_{-i}$. □

With this and the preceding lemmas, we need only care about misrepresentations that result in a agent receiving a bundle that contains the agent's bundle of interest and has the same output. We now present the main theorems of this section.

Theorem 31 *A USM-TR auction is incentive compatible in dominant strategies if and only if, for any agent i, any $\mathbf{q}_{-i}$ and any $\mathbf{v}_{-i}$, i has no incentive to report any bundle q_i with the same output as $\check{q}_i$, such that $q_i \neq \check{q}_i$ and $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$.*

Proof. Case if: Since $\check{q}_i \leq F_i(\check{q}_i, \mathbf{q}_{-i})$ and since F does not change the reported output, we have $V_i(\check{q}_i) = V_i(F_i(\check{q}_i, \mathbf{q}_{-i}))$. Thus, since the Trade Reduction auction is individually rational, i receives non-negative utility from reporting its bundle of interest and value truthfully.

By Lemma 29, i has non-positive utility if it is not the case that $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$. So i can obtain positive utility only if $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$. By Lemma 28, i has non-positive utility only if it reports its output truthfully. Hence, by Lemma 30, i would report its value truthfully. Then, the auction is incentive compatible in dominant strategies if any agent i has no incentive to report $q_i \neq \check{q}_i$ with its true output, such that $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$.

Case only if: With Lemmas 28–30, the misrepresentation stated in this Theorem is the only possible way an agent could improve its utility. Thus the "only if" case is true by definition of incentive compatibility. □

Theorem 32 *A USM-TR auction is ex post Nash incentive compatible if and only if, for any agent i, any q̌, and any v̌, i has no incentive to report any bundle q_i with the same output as $q̌_i$, such that $q_i \neq q̌_i$ and $q̌_i \leq F_i(q_i, \mathbf{q}_{-i})$, when all other agents report truthfully.*

Proof. The proof is similar to the proof of Theorem 31. □

6.2 Ex Post *Nash Incentive Compatibility by Removing Bids*

Consider a particular USM-TR auction we call the **USM-TR-RB** (for USM-TR Remove Bigger) auction, with the conversion function F defined as follows:

(1) For any producer i we have $F_i(q_i, \mathbf{q}_{-i}) = q_i$.
(2) For any consumer i, if there exists an agent j such that $q_i > q_j$, then $F_i(q_i, \mathbf{q}_{-i}) = q_i \cup \bar{g}$ (the bid is removed), otherwise $F_i(q_i, \mathbf{q}_{-i}) = q_i$.

At first glance it might appear that, because consumer bundle q_i is removed when $q_i > q_j$ for another consumer j, the auction satisfies Theorem 32 for consumers. However, if j reports q_i instead of q_j, and there is no other reported bundle q_k such that $q_i > q_k$, then no q_i bundles get removed. We can guarantee that all such q_i bids get removed if there exists a reported bundle q_k other than q_j such that $q_i > q_k$. If this holds, then the auction satisfies Theorem 32.

Theorem 33 *The USM-TR-RB auction is* ex post *Nash incentive compatible, if, for any consumer j, if there is a consumer i such that $q̌_i > q̌_j$ then there exists a consumer k such that $q̌_i > q̌_k$.*

Proof. We show that the conditions of Theorem 32 hold for all agents. The auction does not change the producers' reported bundles. If a producer reports its bundle of interest untruthfully, but with the same output, while all other agents report truthfully, the unique manufacturing technologies property ensures that the producer will be the only agent in its market and will lose the auction. Hence, Theorem 32 implies *ex post* Nash incentive compatibility for producers.

Now, consider consumer j as in the theorem, and assume that all other agents bid truthfully but j bids such that $q_j > q̌_j$. If there is no consumer i such that $q_j = q̌_i$, then j will be the only agent in its reported market and will lose the auction. If, on the other hand, there is such a consumer i, and also a consumer k such that $q̌_i > q̌_k$, then the bid of j (and also i) will be removed. In either case Theorem 32 implies *ex post* Nash incentive compatibility for consumers and the theorem is proven. □

Note that the condition of the theorem above always holds if there are at least two agents in each market. Also note that a consumer would be forced out of the auction if another consumer has a strictly smaller bundle of interest. This could happen even if the forced-out consumer has a very high value, potentially causing the efficiency to be very low. In fact, the efficiency can be arbitrarily close to zero, hence, in general, we can give no positive competitive ratio for efficiency. However, if there are no consumers i and j such that $\check{q}_i > \check{q}_j$, then Theorem 33 holds and there is a good competitive ratio for efficiency.

Observation 34 *If there are no consumers i and j such that $\check{q}_i > \check{q}_j$, then the USM-TR-RB auction is* ex post *Nash incentive compatible and the efficiency competitive ratio from Theorem 25 holds for the USM-TR-RB auction.*

6.3 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*** auction is a USM-TR auction with a conversion function F that sets the converted bundle for agent i to have the maximum number of inputs of any reported bundle with the same output as reported by i. Formally, for agent i, each good g such that $q_i^g > 0$, and with O_i being the set of agents with the same output as i, $F_i(q_i, \mathbf{q}_{-i}) = \max_{j \in O_i} q_j^g$. Observe that all producers with the same reported output are merged into a single market and all consumers are merged into a single market.

Theorem 35 *The USM-TR-Merge auction is incentive compatible in dominant strategies.*

Proof. We show that for any consumer i and any bids from the other agents, i cannot improve its utility by reporting a bundle $q_i \neq \check{q}_i$ such that $\check{q}_i \leq F_i(q_i, \mathbf{q}_{-i})$. This will establish the conditions of Theorem 31 for consumers. A similar argument holds for producers with the same output, proving the theorem.

For any such bundle q_i from i, the partition of agents into markets is the same as if i bid truthfully. For any such bundle giving $F_i(\check{q}_i, \mathbf{q}_{-i}) = F_i(q_i, \mathbf{q}_{-i})$ the converted bundles for all agents will be exactly the same as if i bid truthfully, hence i cannot improve its utility by reporting q_i instead of $\check{q}_i$ in this case.

For any bundle q_i giving $F_i(\check{q}_i, \mathbf{q}_{-i}) < F_i(q_i, \mathbf{q}_{-i})$ the partition of the agents into markets is the same and the bundles for all producer markets are the same as if i bid truthfully. However, the market containing all the consumers requires more inputs. The value to i for winning is the same as if it bid truthfully, but we will show that its payment can only increase, which means the utility of i cannot increase. By Lemma 20, if i wins, its payment is the minimal bid value necessary to be in the Trade Reduction allocation. Since $F_i(\check{q}_i, \mathbf{q}_{-i}) < F_i(q_i, \mathbf{q}_{-i})$, more inputs are needed

31

for each consumer when i bids q_i. Since the additional inputs incur additional non-negative cost for each consumer, and the ordering of the consumers in their market (by reported value) does not change, the minimal value for i to win is at least as high with report q_i as with $\breve{q}_i$. So we conclude that its payment is at least as high with report q_i as with $\breve{q}_i$. $\square$

Recall that, by the unique manufacturing technologies property assumption, at most one market can produce any good. Thus, since producers bid truthfully in USM-TR-Merge, no markets of producers will actually be merged. Still, although no producer markets are actually merged, the merging rule is still necessary to ensure incentive compatibility in dominant strategies for the producers.

In general, it is ambiguous whether the USM-TR-Merge auction would give higher or lower efficiency than the Trade Reduction auction for agents with privately known bundles and values, 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. In fact, efficiency could be arbitrarily close to zero, hence we can establish no positive competitive ratio for efficiency.

With an additional assumption on the consumers' bundle of interest, we can avoid merging consumers markets and ensure high efficiency. We say that the supply chain has the *k-**Input Consumers*** property, if for some fixed integer k (publicly known), each consumer desires exactly k units of all goods in total. That is, for any consumer i, $\sum_{g \in G} \breve{q}_i^g = k$. If for some fixed k, the k-input consumers property holds, we can gain incentive compatibility without merging any consumer markets. In the ***USM-TR-Merge-kIC*** (USM-TR-Merge k-Input Consumers) auction, F_i is defined for producers as in USM-TR-Merge. No consumer markets are merges, but the auction rejects all consumer bids for other than k units in total. Formally, for each consumer i, if $\sum_g \breve{q}_i^g = k$ then $F_i(q_i, \mathbf{q}_{-i}) = \breve{q}_i$, otherwise $F_i(q_i, \mathbf{q}_{-i}) = \breve{q}_i \cup \bar{g}$ (the bid is removed). With the k-input consumers property, no consumer can feasibly misrepresent itself as any other consumer.

Observation 36 *When the k-Input Consumers property holds, the USM-TR-Merge-kIC auction is incentive compatible in dominant strategies.*

Proof. We show that the conditions of Theorem 31 hold. The conditions hold for producers by the same argument as in the proof of Theorem 35. The conditions hold for consumers since any agent that misreports its bundle of interest is removed, and never improves its utility. Thus the USM-TR-Merge-kIC auction is incentive compatible in dominant strategies. $\square$

Since no merging is actually performed, and no agent is removed when the agents bid truthfully, so our competitive ratio bound holds for the USM-TR-Merge-kIC auction.

Observation 37 *When the k-Input Consumers property holds, the efficiency competitive ratio from Theorem 25 holds for the USM-TR-Merge-kIC auction.*

7   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 Trade Reduction 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. For instance, OR preferences allow an agent to specify that it would accept one or more bundles specified in a set of bundles, with a different value for each bundle, and it values the set of bundles it receives as the sum of the values of the bundles. XOR preferences allow an agent to specify that it wants exactly one bundle from a specified set, with a different value for each bundle, and it values the set of bundles as the maximal value of any bundle it receives. These extensions would allow agents (companies) to express, for instance, different values/costs for different quantities and alternate production technologies. 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 incentive compatibility. Consider the case in Figure 2 where the true preferences of A1 contain XOR components \$13 in market CM1 and \$35 in market CM3. If A1 bids truthfully, it will win none 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.

Future work should include more realistic features of supply chains into our basic model, including multiple unit production, time-based preferences (e.g., deadlines

and production lead time) and non-commodity multi attribute goods. We expect to be faced with interesting mechanism design challenges in adapting the auction to an extended model while maintaining desirable economic properties.

In a supply chain, agents may not wish to reveal their full valuations outright in the single-shot auction we developed. Indeed, if agents are not single minded, it could be infeasible for agents to communicate their entire preferences. To address this problem, iterative combinatorial auctions have been studied by others [20] [21]. Extending these results to a supply chain model might increase the viability of auctions for real-world supply chains [22].

References

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

[2] W. E. Walsh, M. P. Wellman, Decentralized supply chain formation: A market protocol and competitive equilibrium analysis, Journal of Artificial Intelligence Research 19 (2003) 513–567.

[3] N. Sadeh, R. Arunachalam, J. Eriksson, N. Finne, S. Janon, TAC-03: A supply chain trading competition, AI Magazine 24 (1) (2003) 92–93.

[4] S. de Vries, R. Vohra, Combinatorial auctions: A survey, INFORMS Journal on Computing 15 (3) (2003) 284–309.

[5] W. E. Walsh, M. P. Wellman, F. Ygge, Combinatorial auctions for supply chain formation, in: Second ACM Conference on Electronic Commerce, 2000, pp. 260–269.

[6] W. E. Walsh, Market protocols for decentralized supply chain formation, Ph.D. thesis, University of Michigan (2001).

[7] E. H. Clarke, Multipart pricing of public goods, Public Choice 11 (1971) 17–33.

[8] T. Groves, Incentives in teams, Econometrica (1973) 617–631.

[9] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance 16 (1961) 8–37.

[10] J. K. MacKie-Mason, H. R. Varian, Generalized Vickrey auctions, Tech. rep., Dept. of Economics, Univ. of Michigan (Jul. 1994).

[11] R. B. Myerson, M. A. Satterthwaite, Efficient mechanisms for bilateral trading, Journal of Economic Theory 29 (1983) 265–281.

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

[13] M. Babaioff, N. Nisan, Concurrent auctions across the supply chain, in: Third ACM Conference on Electronic Commerce, 2001, pp. 1–10.

[14] R. P. McAfee, A dominant strategy double auction, Journal of Economic Theory 56 (1992) 434–450.

[15] D. Kjenstad, Coordinated supply chain scheduling, Ph.D. thesis, Norwegian University of Science and Technology (1998).

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

[17] A. Andersson, M. Tenhunen, F. Ygge, Integer programming for combinatorial auction winner determination, in: Fourth International Conference on Multi-Agent Systems, 2000, pp. 39–46.

[18] J. Feigenbaum, S. Shenker, Distributed algorithmic mechanism design: Recent results and future directions, in: 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, 2002, pp. 1–13.

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

[20] L. M. Ausubel, P. Milgrom, Ascending auctions with package bidding, Tech. rep., University of Maryland and Stanford University (2001).

[21] D. C. Parkes, L. H. Ungar, Iterative combinatorial auctions: Theory and practice, in: Seventeenth National Conference on Artificial Intelligence, 2000, pp. 74–81.

[22] P. Cramton, Applications of auction theory: Ascending auctions, European Economic Review 42 (1998) 745–756.

A   Proofs

We prove Theorem 25 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 $m \in M$ such that $S\{A'\}(m) \neq 0$. We denote by $V(S\{A'\})$ the sum of valuations of all agents in procurement set $S\{A'\}$, so $V(S\{A'\}) = \sum_{i \in S\{A'\}} v_i$.

Definition 38 (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, $S\{A'\}_1(m) = S\{A'\}_2(m)$. 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 $\mathbf{S} \in \hat{\mathbf{S}}_T$, 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 **consumer market of the procurement set topology** $\mathbf{S}$.

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

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

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

Lemma 39 *Let $\breve{v}$ be any vector of agent values. Let $\hat{\mathbf{S}}(\breve{A})$ be the set of procurement set topologies in the efficient allocation $\breve{A}$ for $\breve{v}$. Then:*

(1) For any consumer market $m \in CM^*$, there is only one procurement set topology in $\hat{\mathbf{S}}(\breve{A})$ for which m is the consumer market. So we can mark the consumer markets by m_i and the procurement set topologies by $\mathbf{S}^i$ for $i = 1, \dots, |CM^*|$.
(2) For each procurement set topology $\mathbf{S}^i$, $N^*(\mathbf{S}^i) = N^{TR}(\mathbf{S}^i) + 1$.
(3) For each procurement set topology $\mathbf{S}^i$ and its consumer market m_i, $N^*(\mathbf{S}^i) = T_{m_i}(\breve{A})$ and $N^{TR}(\mathbf{S}^i) = T_{m_i}(A^{TR}) = T_{m_i}(\breve{A}) - 1$.
(4) For each market m, the trade sizes in market m in the optimal and TR alloca-

tions are:

$$T_m(\check{A}) = \sum_{i=1}^{|CM^*|} \mathbf{S}^i(m) T_{m_i}(\check{A}) \quad and \quad T_m(A^{TR}) = \sum_{i=1}^{|CM^*|} \mathbf{S}^i(m)(T_{m_i}(\check{A}) - 1)$$

(5) The allocations $\check{A}$ and A^{TR} can be partitioned to procurement sets in the following way.

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

where S_j^i is the j procurement set of topology $\mathbf{S}^i$, $\bigcup_{i=1}^{|CM^*|} S_{N^*(\mathbf{S}^i)}^i$ are the reduced agents and $V(S_{N^*(\mathbf{S}^i)}^i) \leq V(S_j^i)$ for all $j = 1, \dots, N^*(\mathbf{S}^i)$.

Proof.

Proof of (1). 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 defines 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 for which m is the consumer market. Hence we have proven (1).

Proof of (2). Let $\mathbf{S} \in \hat{\mathbf{S}}(\check{A})$ be the unique procurement set topology of any consumer market $m \in CM^*$. We now show that $N^*(\mathbf{S}) = N^{TR}(\mathbf{S}) + 1$. From the above we conclude that in order to reduce one agent in market m and maintain feasibility, a procurement set of topology $\mathbf{S}$ must be reduced in the Trade Reduction allocation, therefore $N^*(\mathbf{S}) \geq N^{TR}(\mathbf{S}) + 1$.

It remains to prove that $N^*(\mathbf{S}) \leq N^{TR}(\mathbf{S}) + 1$. Assume that this is not true, then for some $\mathbf{S} \in \hat{\mathbf{S}}(\check{A})$, $N^*(\mathbf{S}) > N^{TR}(\mathbf{S}) + 1$, or equivalently $N^*(\mathbf{S}) \geq N^{TR}(\mathbf{S}) + 2$. This means that at least two procurement sets of the same topology $\mathbf{S}$ are reduced. Below we show that for any procurement set S of topology $\mathbf{S} \in \hat{\mathbf{S}}(\check{A})$, we have $V(S) > 0$. Since $V(S) > 0$ for any such procurement set S, we can add one of the reduced procurement sets of topology $\mathbf{S}$ and increase the value of the allocation while ensuring that every market with trade has at least one reduced agent. This is true since we still have one procurement set of topology $\mathbf{S}$ reduced and both procurement sets share the same set of markets. Thus, we can add one of the reduced procurement sets while maintaining the constraints on a TR allocation, 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 $N^*(\mathbf{S}) = N^{TR}(\mathbf{S}) + 1$ and we have proven (2).

We now show that for any procurement set S of topology $\mathbf{S} \in \hat{\mathbf{S}}(\check{\mathbf{A}})$, 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 integral bids (as described in Section 4), hence no subset of agents can have value exactly zero.

Proof of (3). Any procurement set topology $\mathbf{S}$ contains a single agent in its consumer market m, and by feasibility of the allocation, all inputs to each of the consumers must be manufactured, therefore $N^*(\mathbf{S}) = T_m(\check{A})$ and $N^{TR}(\mathbf{S}) = T_m(A^{TR})$. Since we have shown that $N^*(\mathbf{S}) = N^{TR}(\mathbf{S}) + 1$, then $N^{TR}(\mathbf{S}) = N^*(\mathbf{S}) - 1 = T_m(\check{A}) - 1 = T_m(A^{TR})$ and we have proven (3).

Proof of (4). The trade size $T_m(\check{A})$ for each market m is as stated in the lemma from the following observation. Feasibility of the allocation $\check{A}$ dictates that if the trade size in each consumer market m_i is $T_{m_i}(\check{A})$, then for each market m a set of agents of size $\mathbf{S}^i(m)T_{m_i}(\check{A})$ must be in the allocation (and these sets must be disjoint since each producer manufacturers a single item of one good). The trade size in market m is the sum of these quantities over all consumer markets. A similar argument works for the TR trade sizes. We have proven (4).

Proof of (5). From all the above we conclude that

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

where S_j^i is the j procurement set of topology $\mathbf{S}^i$. $\bigcup_{i=1}^{|CM^*|} S_{N^*(\mathbf{S}^i)}^i$ are the reduced agents and $V(S_{N^*(\mathbf{S}^i)}^i) \leq V(S_j^i)$ for all $j = 1, \ldots, N^*(\mathbf{S}^i)$. This is since the reduced agents are the lowest value agents in each market, so any reduced procurement set has a lower value (as sum of the agents value) then any procurement set of the same topology which is in the optimal allocation. We have seen that one procurement set of each consumer market must be reduced, and therefore all the reduced agents can be partitioned to one procurement set of each consumer market. We have proven (5). $\square$

Lemma 40 *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 v_i such that $VTR_i < v_i$, then i wins the auction by Lemma 20. In particular, i wins if it bids $VTR_i < v_i < PBV_i$. But by the definition of PBV_i, $v_i \geq PBV_i$, which is a contradiction. $\square$

Lemma 41 *Let $\check{v}$ be any vector of agent values. The KSM-TR allocation for $\check{v}$ 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 Trade Reduction allocation.

We must show that $Pay(A^{TR}) \geq 0$ to prove that A^{TR} with KSM-TR payments is budget-balanced. By Lemma 39 the Trade Reduction allocation is

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

therefore $Pay(A^{TR}) = \sum_{i=1}^{|CM^*|} \sum_{j=1}^{N^*(\mathbf{S}^i)-1} P(S^i_j)$. So it is sufficient to show that $P(S^i_j) \geq 0$ for $i = 1, \ldots, |CM^*|$ and $j = 1, \ldots, N^*(\mathbf{S}^i) - 1$, to prove the lemma.

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 40 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, by summing over all agents in the procurement set, 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. $\square$

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

Definition 42 (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 $\sum_{i \in P^{A'}_i} \check{v}_i$*

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

Definition 43 (good partition pair) *Given vector of agent values $\check{v}$, with efficient allocation $\check{A}$ and a feasible reduction allocation A^{FTR}, we say that the allocations have a **good partition pair** P^*, P^{tr} if there exists a partition P^* for the efficient allocation $\check{A}$ of size k, and a partition P^{tr} for a feasible reduction allocation A^{FTR} of size k, such that for any $i = 1, \ldots, k$:*

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

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

Lemma 44 *Given vector of agent values $\check{\mathbf{v}}$, with non-empty efficient allocation $\check{A}$ which has a good partition pair P^*, P^{tr}, we have:*

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

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

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

Since P^*, P^{tr} is a good partition pair, for every $i = 1, \ldots, k$ it is true that $V(P_i^*) \geq V(P_i^{tr}) \geq 0$. Therefore, we can apply Lemma 45 to get

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

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

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

□

Lemma 45 *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. $\square$

From Lemma 44 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 46 *Let $\check{v}$ be any vector of agents values with efficient allocation $\check{A}$. $\check{A}$ has a good partition pair.*

Proof. By Lemma 39 the efficient allocation is constructed from procurement set topologies $\mathbf{S}^i$ for $i = 1, \ldots, |CM^*|$ such that

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

where S_j^i is the j procurement set of topology $\mathbf{S}^i$, and $S_{N^*(\mathbf{S}^i)}^i$ 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 Trade 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 the two requirements for a good partition pair hold.

- $V(P_i^*) \geq V(P_i^{tr}) \geq 0$: Since $S_{N^*(\mathbf{S}^i)}^i$ is a procurement set, it has a non-negative value. Hence, $V(P_i^*) - V(P_i^{tr}) = V(S_{N^*(\mathbf{S}^i)}^i) \geq 0$. Since every procurement set has non-negative value, we also have $V(P_i^{tr}) \geq 0$.

- $\frac{V(P_i^{tr})}{V(P_i^*)} \geq \min_{\mathbf{S} \in \hat{\mathbf{S}}(\check{A})} \frac{N^{TR}(\mathbf{S})}{N^*(\mathbf{S})}$: 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 47 we get

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

$\square$

Lemma 47 *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. $\square$

Finally, we are ready to prove the theorem.

Theorem 25 *Let $\check{v}$ be any vector of agents values. The following is an efficiency competitive ratio function for the KSM-TR auction:*

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

if $\check{A} \neq \emptyset$ and

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

if $\check{A} = \emptyset$.

Proof. The second component of the competitive ratio is true by definition, hence we prove the first component. By Lemma 46 $\check{v}$ has a good partition pair. By applying Lemma 44, the efficiency of KSM-TR satisfies the following:

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

From Lemma 39 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 $N^{TR}(\mathbf{S}) = T_m(A^{TR}) = T_m(\breve{A}) - 1$ and $N^*(\mathbf{S}) = T_m(\breve{A})$.

So we conclude that

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

Therefore

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

which is what we wanted to prove. $\square$