ECAI 2023
K. Gal et al. (Eds.)

© 2023 The Authors.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/FAIA230506

Truthful and Equitable Lateral Transshipment in Multi-Retailer Systems

Garima Shakya[a;*], Sai Koti Reddy Danda[b], Swaprava Nath[c], Pankaj Dayama[d] and Surya Shravan Kumar Sajja[d]

[a]Kyushu University, Fukuoka, Japan
[b]Avesha Inc., Bangalore, India
[c]Indian Institute of Technology Bombay, Mumbai, India
[d]IBM Research, Bangalore, India

Abstract. We consider a multi-retailer system where the sellers are connected with each other via a transportation network and the transactions with the consumers happen on a platform. Each consumer is serviced by only one retailer. Since the demands to the sellers (i.e., the retailers on the platform) are stochastic in nature, supplies can be either in excess or in deficit. Transshipping these items laterally among the retailers benefits both, the platform and the retailers. For retailers, excess supply leads to wastage and deficit to a loss of revenue, while via transshipment, they get a better outcome. The platform can also earn some revenue in facilitating this process. However, only the sellers know their excess (which can be salvaged at a price or transshipped to another seller) or the deficit (which can be directly procured from a supplier or transshipped from another seller), both of which have multiple information that is private. We propose a model that allows lateral transshipment at a price and design mechanisms such that the sellers are incentivized to *voluntarily participate* and be *truthful*. Experimenting on different types of network topologies, we find that the sellers at more *central* locations in the network get an unfair advantage in the classical mechanism that aims for *economic efficiency*. We, therefore, propose a modified mechanism with tunable parameters which can ensure that the mechanism is more *equitable* for non-central retailers. Our synthetic data experiments show that such mechanisms do not compromise too much on efficiency, and also reduce *budget imbalance*.

1 Introduction

Online and in-store inventories are integrated in modern markets, providing a platform for retail partners and consumers to interact. Consumers can browse, compare, and purchase products on these platforms. Each retailer has stores and warehouses at a few specific geographical locations to service in a limited area. Most of the warehouses are legally not permitted to service beyond their jurisdiction. Due to the uncertainty in demand, these stores order in advance and keep the supplies in their inventory. After demand realization, the stores may face stock-out or have excess supplies. If the product is not consumed and remains on the shelf, the excess inventory results in increased inventory holding costs, and shortage of supply results in poor service level and lost revenue.

For perishable goods like dairy products, seafood, meat [9, 24], baked goods [12], fruits, flowers, medical supplies [34, 36] etc., the loss due to excess supply is enormous due to the limited shelf life. Similarly, for products with expensive downtime costs, like spare parts, the firms maintain sufficient inventory, thus incurring high inventory handling costs to quickly respond to a breakdown of a system to avoid lengthy downtime [32]. There exist a variety of such contexts, e.g., chemical plants, airline industry, power-generating plants.

In multi-retailer systems, a better planning and collaboration among the retailers can improve efficiency, alleviate loss, and reduce inventory management costs. We consider the use of *lateral transshipment* between strategic parties of the same echelon, e.g., retailers and wholesalers. We use the term transshipment to mean lateral transshipment. Competitive retailers aim to optimize their objectives, making collaboration challenging and leading to supply chain inefficiency. Another challenge is to distribute the profit after transshipment among the retailers of different firms. The monetary amount paid by a retailer to the other retailer on the exchange of the transshipped product is known as the *transshipment price*.

Finding the transshipment plan that maximizes the total value of the agents requires access to the private information (consisting of multiple components called the *type*) of each agent. This makes the setting a competitive game between the retailers where they may strategically choose to lie and misreport actual types to maximize their payoff. Incentivizing competitive retailers to reveal their true types to the planner is one of the main challenges. Therefore, finding the allocation that maximizes the total profit is a *mechanism design* [4], as well as a *multi-agent and distributed planning* problem.

If the private information is one-dimensional, standard mechanism design techniques, such as the popularly known **VCG** family of mechanisms [6, 17, 33], can be applied [31]. However, lateral transshipment requires knowing multiple parameters from retailers, e.g., the quantity of the excess demand or supply for each retailer and the salvation price, which makes it a non-trivial multi-dimensional mechanism design problem. Certain impossibility results [13, 27, 29] tell us that monetary transfers are necessary for revealing private information under such settings. Further, the mechanism must ensure that every retailer always gets a non-negative benefit from participating in the transshipment process. For trades between the distributed retailers, the more central ones on the transshipment network get an unfair advantage over those at non-central locations. In this paper, we ex-

* Corresponding Author. Email: garima@inf.kyushu-u.ac.jp.

Please check ArXiv or contact the authors for any appendices or supplementary material mentioned in the paper.

plore ways to find an *equitable* lateral transshipment, which provides more *equal opportunities to the equal-sized retailers* to compete and survive in the market irrespective of their position on the network.

The main objectives of this paper is to model the lateral transshipment problem in a strategic environment and find a mechanism to compute the transshipment plan (allocation and pricing) such that: (a) the agents have the incentive to report their private information truthfully *(Strategyproofness)*, (b) the agents always prefer to participate in the transshipment *(Individual Rationality)*, (c) the mechanism requires minimum amount of resources to compute the outcome *(Computational Tractability)*, (d) the collaborative valuation is maximum *(Efficiency)*, and (e) the agents are treated equitably *(Equitability)*. Achieving equitability and efficiency is hard to satisfy together [5, 14]. Hence we aim to get more equitable solutions that do not lead to a significant loss of efficiency.

1.1 Our contributions

This paper provides a model for the modern online market. Motivated by the interaction between the retailers and the consumers, we consider a platform where consumers select the product and place an order from a retailer. We propose a game-theoretic model for the single-shot interaction of retailers. The retailers individually estimate the demand, decide the amount of inventory to order, and privately place the order from outside sources. Due to uncertainty in consumer demand, the retailers may face excess or shortage of the products. We propose the **W**eighted **V**alue **T**ransshipment (**WVT**) mechanism that has a pre-defined and publicly known contract for allocation and profit distribution in the single-shot interaction. The platform asks the retailers to announce their multi-dimensional types and decides the transshipment allocation and price using the **WVT** mechanism. The main contributions of this paper are as follows:

- We show that the **WVT** mechanism incentivizes the retailers to report their multi-dimensional private information truthfully (Theorem 1).
- **WVT** mechanism ensures that the retailer can never get worse off by participating in the transshipment hence guarantees their voluntary participation (Theorem 2).
- We prove that there exists a strongly polynomial-time algorithm for the **WVT** mechanism to compute the transshipment allocation and pricing (Theorem 3).
- The experimental analysis on a classical mechanism for truthful resource allocation considering different network topologies as the transshipment network shows that the classical mechanism leads to an unfair advantage to more central retailers in the transshipment network (Figure 2).
- The **WVT** mechanism has tunable parameters that can ensure a more equitable outcome for smaller or non-central retailers, as shown in the experiments (Figure 3(b)).
- We experimentally compare the outcomes of **WVT** with that of **VCG** on synthetic data. The results show that the **WVT** mechanism does not compromise too much on efficiency (Figure 3(a)) and reduces the budget imbalance (Figure 3(c)).

The literature on transshipment is rich and extensive. In the following subsection, we review only those works that are closely related.

1.2 Related work

Researchers have considered transshipment allocation and profit distribution between multiple retailers under a *parent firm* as a cooperative game and proposed the effective use of profit distribution methods such as the Shapley value, the nucleolus, and the τ-value from cooperative game theory [20, 21].

The setting with two non-cooperative retailers is well studied [30], and there exist solution for optimal transshipment [18, 28]. In a setting with more than two retailers, Anupindi, Bassok, and Zemel [2] gave a two-stage framework where in the first stage, the players are non-cooperative and decide the number of items to order individually. Assuming that each retailer's actual residual supply and demand is complete knowledge, [2] provide a dual price allocation for transshipments and prove that the solution is in the *core*[1]. In a following work, Granot and Sošić [15] proposed a three-stage model, where each retailer has the opportunity to decide how much of her residual supply/demand she would like to share with others and strategically reports the residuals before the allocation is determined. They show an impossibility result that says, for transshipment games, there are no allocation rules based on dual prices that can ensure complete sharing of the residuals. Yan and Zhao [35] proposed a model and mechanism for coordination among the manufacturer and retailers. At first, each retailer decides whether to participate in the transshipment in the future and pays a participation fee accordingly. After demand realization, the retailers strategically choose and report the amount of the residual they want to share. Next, the efficient allocation is computed, and the net profit is given to both the parties instead of distributing it among both of them. The mechanism leads to only a grand coalition inducing complete residual sharing with an appropriately set participation fee.

To make a moral decision concerning the distribution of residuals, one way is to consider the well-known Aristotle's principle of distributive justice [22], "Equals should be treated equally, and unequals unequally, in proportion to the relevant *similarities* and *differences*." Hornibrook, Fearne, and Lazzarin [19] developed a behavioral theoretical approach to fill the gap between equitability, justice, and the supply chain relationships between the buyers and sellers to achieve a fair allocation of resources such as time, effort, and money. A study by Fearne et al. [10] measures fairness in supply chain trading relationships and shows the importance of understanding equitability in sustainable supply chains.

It is worth noting that there is a dearth of models in the extant literature that can simultaneously consider the efficiency and profit distribution goals in transshipment allocation between multiple non-cooperative retailers. Some of the literature assumes that the retailers share their complete private information and then examine methods for a cooperative game to find transshipment allocation and price distribution between the retailers, e.g., [2]. The other kind of literature considers additional participation fees, e.g., [35], which may result in a negative payoff for retailers and, therefore, retailers' lack of interest in participating in transshipment. We describe our problem setting and contributions in the following section.

2 Model Descriptions and Assumptions

Define $\mathcal{R} = \{1, 2, \cdots, n\}$ to be the set of retailers of a product[2] available on the platform. The minimum trade volume is referred to as *one unit*. We consider a single-shot interaction of the retailers. The retailers are myopic and want to maximize their payoff. Each

[1] An allocation is said to be in *core* if it is efficient and provides coalitional rationality, which means no group of retailers can collude and get more benefit than that in the given allocation [8].

[2] This model easily generalizes to multiple products with additive valuations for the retailer.

retailer independently estimates the future demand that she will receive, checks the stock she has, and decides the quantity of product to buy from the manufacturers. Retailer $i \in \mathcal{R}$ orders Q_i units from a manufacturer at per unit cost b_i. Each consumer submits their demand to one of the retailers. Each retailer $i \in \mathcal{R}$ receives the demand D_i on the platform and sells the product at per unit selling price p_i to the consumers. After satisfying demand with her inventory, each retailer strategically decides how much of her excess supply she wants to share with others. If a retailer faces stock-out due to excess demand, she strategically decides how much of the excess demand she wants to report to the platform. The platform computes the allocation of the residuals by transshipment among retailers. Each retailer i faces a per unit penalty cost ρ_i if it has unmet demand. If there are unsold products in i's inventory, they can be monetized at a salvage value of s_i per unit.

The order quantity Q_i, buying cost b_i, and the salvage value s_i are the private information of the retailer i and are represented as a tuple $\mathcal{Z}_i = (Q_i, b_i, s_i)$, which we call the *type* of i. Let $\mathcal{Z} = [\mathcal{Z}_i]_{i\in\mathcal{R}}$ denote the vector of the agents' types. The set of all possible private information of an agent i is denoted by Θ_i. The tuple containing Θ_i of each $i \in \mathcal{R}$ is denoted by $\Theta = \times_{i\in N}\Theta_i$, where Θ_{-i} denotes the set of all type profiles excluding the type of i. The tuple $\mathcal{Z}_i$ is unknown to the platform and the other retailers in $\mathcal{R}$. As the interaction with the customers happens on the platform, the platform knows the realized demands D_is and selling prices p_is. We assume that the platform also knows the penalty costs ρ_i for all $i \in \mathcal{R}$, and the locations of the retailers on the transshipment network, hence the per unit transportation costs between them. We represent the per unit transportation cost between retailers i and k as τ_{ik}. Let $\mathbb{A}$ be the set of all possible transshipment allocations for a product. The transshipment allocation, $\mathcal{A} \in \mathbb{A}$ can be represented as a matrix $[a_{ik}]$ for $i, k \in \mathcal{R}$, s.t., $a_{ik} \in \mathbb{R}_{\geqslant 0}$ is the quantity of products to be transshipped from retailer i to k. The retailer k earns p_k by selling one unit of product transshipped from i to her and satisfying the previously unmet demand. The platform computes the share of transshipment profit, $\mathfrak{p}_{ik}$ paid by k to i, for each unit of transshipment from i to k. We assume that the receiving retailer k pays τ_{ik} per unit as the transportation cost. Each $i \in \mathcal{R}$ has a value v_i for the transshipment allocation, representing the revenue she gets after the transshipment happens.[3] The value v_i is the total earning through transshipment from and to i plus the total salvage value i gets from unsold inventory minus the total penalty cost for the unmet demand. Mathematically,

$$v_i(\mathcal{A}, \mathfrak{p}, \mathcal{Z}_i) = \sum_{k\in\mathcal{R}\setminus\{i\}} a_{ik}\mathfrak{p}_{ik} + \sum_{l\in\mathcal{R}\setminus\{i\}} a_{li}(p_i - \mathfrak{p}_{li} - \tau_{li}) + \Big((Q_i - D_i)^+ - \sum_{\substack{l\in\mathcal{R}\\ l\neq i}} a_{ik}\Big)s_i - \Big((D_i - Q_i)^+ - \sum_{\substack{l\in\mathcal{R}\\ l\neq i}} a_{li}\Big)\rho_i \tag{1}$$

The total revenue i gets is the sum of the direct revenue (before transshipment) and her valuations for a given allocation and transshipment prices $(\mathcal{A}, \mathfrak{p})$, $\text{Revenue}_i(\mathcal{A}, \mathfrak{p}, \mathcal{Z}_i) = (\min\{D_i, Q_i\}\, p_i - b_i Q_i) + v_i(\mathcal{A}, \mathfrak{p}, \mathcal{Z}_i)$. All the notations are summarized in a table in the appendix.

In the settings where the agents' valuations are private, and the mechanism does not have any additional structures (e.g., payments in our context), only dictatorial mechanisms are truthful [13, 29]. This negative result holds irrespective of whether agents' preferences are *ordinal* (representable as an order relation over the outcomes) or *cardinal* (agents have a real number to represent the intensity of the preference). Note that in our setup, the agents' preferences are cardinal. A complementary analysis by Roberts [27, Thm 7.2] shows that a dictatorship result reappears under certain mild conditions in a *quasi-linear* setting (which is our current setting) unless transfers (of utility) are allowed. Therefore, the use of transfers in some form is inevitable to ensure truthfulness of the agents. In this paper, the mechanism decides the allocation $\mathcal{A}$, the transshipment price $\mathfrak{p}$, which directly affects the valuations of retailers; and determines payments or transfers $\mathcal{P} = (\mathcal{P}_i, i \in \mathcal{R})$ for each of the retailers. We assume that every retailer wants to maximize their valuation and also wants to pay less. The net payoff or utility[4] of a retailer is assumed to follow a standard *quasi-linear form* [31]:

$$u_i((\mathcal{A}, \mathfrak{p}, \mathcal{P}), \mathcal{Z}_i) = v_i(\mathcal{A}, \mathfrak{p}, \mathcal{Z}_i) - \mathcal{P}_i. \tag{2}$$

As the utility function depends on $\mathcal{Z}_i$, which is the private information of the retailer i, the platform needs the retailers to report the $\mathcal{Z}_i$s to the mechanism designer who decides the allocation to achieve a certain objective. This leaves an opportunity for a retailer to misreport her private information and get a better allocation. A mechanism designer needs to carefully design the allocations and payments in the face of such strategic behavior of the retailers. We use the notation $\mathcal{X}$ to denote the pair $(\mathcal{A}, \mathfrak{p})$ and the set of all possible $\mathcal{X}$ as $\mathbb{X}$. To distinguish, we denote $\hat{\mathcal{Z}}_i$ as the announced information and $\mathcal{Z}_i$ as the true information of i. Therefore, a mechanism in this setting is defined as a function formally defined as follows.

DEFINITION **1 (Transshipment Mechanism)** *A* Transshipment Mechanism (TM) *is a mapping $f : \Theta \to \mathbb{X} \times \mathbb{R}^n$ that maps the reported type vector to an allocation, transshipment price and payment for each retailer. Hence, $f(\mathcal{Z}) = (\mathcal{X}(\mathcal{Z}), \mathcal{P}(\mathcal{Z}))$, where $\mathcal{X}$ is the function which computes the allocation and the transshipment price, and $\mathcal{P}$ is the payment function.*[5]

The TM defines two payments as its output: the transshipment price indicating the price at which the transaction between the source and destination retailers happen, and the payment indicating the side-payment to satisfy other desirable properties, e.g., truthfulness, individual rationality (defined in the next section). We use $v_i(\mathcal{X}(\mathcal{Z}))$ or $v_i(\mathcal{X})$ as a shorthand for $v_i(\mathcal{X}(\mathcal{Z}), \mathcal{Z}_i)$ when the arguments are obvious from the context. In the supplementary material, we formally define the desirable properties namely, *efficiency*, *truthfulness* and *individual rationality* of a TM.

3 The Proposed Mechanism

Equitability among retailers can be addressed in many ways. We consider the setting where each agent i is given a potentially different weight w_i, which can be based on externalities (e.g., their position in the network). We use the notation w to denote the weight vector, $w = [w_i]_{i\in\mathcal{R}}$. The allocation and payment decisions in the proposed mechanism resemble the affine maximizer rule [27]. However, due to the multi-dimensional types of the agents, the computational complexity of the transshipment allocation, and the transshipment pricing, the proposed mechanism becomes significantly different in the

[3] The valuation of a retailer for different products is independent. The total valuation for all products is assumed to be the sum of the values of each product.

[4] We will only consider the valuation component of the revenue in the utility of the agent since the direct revenue $(min\{D_i, Q_i\}\, p_i - b_i Q_i)$ is insensitive to the mechanism, which decides the allocation and payments.

[5] We overload the notation $\mathcal{X}$ and $\mathcal{P}$ to denote both functions and values of those functions, since their use will be clear from the context.

current setup than the classical affine maximizer. We describe the distinctions after we present the proposed mechanism.

Allocation function: The allocation function is a weighted utilitarian function, with the objective to maximize the weighted sum of valuations given by the following optimization problem (OP).

$$\begin{aligned} \underset{\mathcal{X}=(\mathcal{A},\mathfrak{p})\in\mathbb{X}}{\operatorname{argmax}} & \sum_{i\in\mathcal{R}} w_i\, v_i(\mathcal{X},\mathcal{Z}_i) \\ \text{s.t.} & \sum_{k\in\mathcal{R}\setminus\{i\}} a_{ik} \leqslant (Q_i - D_i)^+ \quad \forall i\in\mathcal{R} \\ & \sum_{l\in\mathcal{R}\setminus\{i\}} a_{li} \leqslant (D_i - Q_i)^+ \quad \forall i\in\mathcal{R} \\ & s_i \leqslant \mathfrak{p}_{ik} \leqslant p_k + \rho_k - \tau_{ik},\ a_{ik}\geqslant 0, \quad \forall i,k\in\mathcal{R}, \end{aligned} \tag{3}$$

The operator $(.)^+$ denotes $max(.,0)$. The first set of constraints in (3) ensures that the total transshipment from every retailer i to others is not more than the excess supply $(Q_i - D_i)^+$. Similarly, the second set of constraints ensures that the total transshipment to i from other retailers is not more than the unmet demand $(D_i - Q_i)^+$. The third set of constraints bound the transshipment price $\mathfrak{p}_{ik}$ to make sure that it is beneficial to every i and k. For every unit of the transshipment from i to k, k does not face the unmet demand and hence is not charged with the penalty cost ρ_k, which would have been charged in the absence of transshipment. Retailer k also earns p_k from the sale of the transshipped product, but pays the transportation cost τ_{ik}. Hence k earns a total profit of $p_k + \rho_k - \tau_{ik}$ from per unit a_{ik}. The third set of constraints ensures that $\mathfrak{p}_{ik}$ is not more than the profit earned by k if the transshipment happens; otherwise, it is better for her not to buy this unit of transshipment. At the same time, every retailer i gets $\mathfrak{p}_{ik}$ from the transshipment, which is at least as much as she earns if the transshipment did not happen. As i gets per unit salvage value s_i in the absence of transshipment, the lower bound for $\mathfrak{p}_{ik}$ is s_i. Note that OP (3) is solved by the mechanism designer who can only access the reported types $\hat{\mathcal{Z}}$. Denote the optimal solution of OP (3) by $\mathcal{X}^*(\hat{\mathcal{Z}})$ and the solution of a similar optimization problem when i does not participate in the transshipment by ${\mathcal{X}_{-i}}^*(\hat{\mathcal{Z}}_{-i})$.

Payment function: For every agent i, the payment computed by the mechanism is as follows,

$$\mathcal{P}_i := \begin{cases} \frac{1}{w_i}\Big(\sum\limits_{\substack{\ell\in\mathcal{R}\\ \ell\neq i}} w_\ell v_\ell({\mathcal{X}_{-i}}^*) - \sum\limits_{\substack{\ell\in\mathcal{R}\\ \ell\neq i}} w_\ell v_\ell(\mathcal{X}^*)\Big) & w_i > 0 \\ 0 & w_i = 0 \end{cases} \tag{4}$$

Assume that each agent i reveals her $\mathcal{Z}_i$ truthfully. If $w_i \neq 0$, the second term of $\mathcal{P}_i$ is paid to i and is the sum of weighted values of each agent except i for the allocation given by Eq. (3). The first term is paid by i, the sum of weighted values of each agent except i for the allocation by Eq. (3) that would have been made if i would not have been present. Thus, i has to pay the difference in the values of the objective function when i is present and absent, concerning w_i. Notice that, similar to v_is, the $\mathcal{P}_i$s can be negative as well. If so, the platform pays $\mathcal{P}_i$ to i. We discuss the budget imbalance in Section 5.3. The mechanism is succinctly presented in Algorithm 1.

Two significant observations make the proposed mechanism different in the current setup.

- The allocation decision here considers not only the volume of the item transshipped but also the price at which the transshipment occurs - each of which has its own constraints to be satisfied.
- The types of agents in the classic affine maximizer rule are single-dimensional, i.e., the value of that agent. In our setting, the type of the agent i has three components: ordered quantity Q_i, purchase price b_i, and salvage price s_i. Hence, to show that properties truthfulness and individual rationality hold in such a *multidimensional* setting is a non-trivial exercise [7, 31].

Algorithm 1 **W**eighted **V**alue **T**ransshipment (**WVT**) Mechanism for a single-shot interaction of the retailers

1: Every retailer $i \in \mathcal{R}$ reports her type $\hat{\mathcal{Z}}_i$ to the platform, where $\mathcal{Z}_i = (Q_i, b_i, s_i)$.
2: Using the realized demand D, reported information $\hat{\mathcal{Z}}$, and weights w, the platform computes the TM $f(\hat{\mathcal{Z}}) = (\mathcal{X}^*(\hat{\mathcal{Z}}), \mathcal{P}^*(\hat{\mathcal{Z}}))$ where $\mathcal{X}^*$ and $\mathcal{P}^*$ are given by Equation 3 and (4) respectively.
3: **Output:** $\mathcal{X}^*(\hat{\mathcal{Z}})$ and the payment vector $\mathcal{P}^*$.

From the first two constraints in OP (3), the mechanism ensures that no retailer is asked to transship more than the residuals she reports to the platform. We assume that the platform can verify the sale of the product and the selling price for each retailer as these transactions are performed on the platform. So, we assume that the mechanism can be implemented as a contract between the platform and the retailers, which they cannot break. This implies:

1. If the mechanism assigns a transshipment, then both the sender and receiver retailers are bound to follow it. Suppose the retailer i had initially misreported her type information and does not have the product in her inventory to sell to k. In that case, the retailer will have to buy the product from the outside market (assumed to be at a much higher price) to fulfill the commitment.
2. At the end of the transshipment, the platform can verify if the receiving retailer sells the transshipped units of the products or not. If not, then it is assumed that she has violated the contract.

4 Theoretical Guarantees

In this section, we summarize the theoretical guarantees of **WVT**. Due to paucity of space, the proofs are deferred to the appendix.

Theorem 1 *The* **WVT** *mechanism is dominant strategy truthful.*

Theorem 1 shows that reporting private information truthfully maximizes a retailer's utility, regardless of what other retailers report. Next result proves that retailers are motivated to join the mechanism on their own.

Theorem 2 *The* **WVT** *mechanism is individually rational.*

The forthcoming results show that **WVT** is strongly polynomial. To show this, we reduce OP (Equation 3) to the b-matching problem, which is known to be strongly polynomial [1].

DEFINITION **2** (***b*-Matching Problem [1])** *Consider a graph $G = (V, E)$, where V is the set of nodes and E is the set of undirected edges. Each edge $e_{u,v} \in E$ between any two nodes $u, v \in V$, has a cost $c_{u,v}$. Let $b = (b_1, b_2, \ldots, b_{|V|})$. A b-matching problem for G is to find the non-negative integer edge weights $y_{u,v}$ which maximises the total cost, $\sum_{u,v\in V} c_{u,v}\ y_{u,v}$ where the sum of weight on edges connected to a node u is no more than b_u, $\forall u \in V$. The b-matching problem can be written as follows.*

$$\begin{aligned} \underset{y}{\operatorname{argmax}} & \sum_{u,v\in V} c_{u,v}\ y_{u,v} \\ \textit{s.t.} & \sum_{v\in V\setminus\{u\}} y_{u,v} \leqslant b_u \quad \forall u\in V, \quad y_{u,v}\geqslant 0, \quad \forall u,v\in V \end{aligned} \tag{5}$$

Theorem 3 *There exists a strongly polynomial algorithm for computing transshipment allocation and pricing in* **WVT**.

Proof: To prove the above statement, we construct a graph and show that finding an optimal b-matching in that graph is computationally equivalent to finding the optimal solution for OP in equation 3.

Construct an edge-weighted bipartite graph $G = (H, J, E)$, where H and J are two sets of vertices such as $H = J = \mathcal{R}$. Let the cost of every edge $e_{i,j} \in E$ is $c_{i,j}$. Set $c_{i,j} = w_i(\mathfrak{p}_{ij} - s_i) + w_j(p_j - \mathfrak{p}_{ij} - \tau_{ij} + \rho_j)$. Define b as, $b_i = (Q_i - D_i)^+$ for every $i \in H$, and $b_j = (D_j - Q_j)^+$ for every $j \in J$. The objective function of the b-matching problem of G (as shown in Figure 1), can be written as

$$\underset{y}{\operatorname{argmax}} \sum_{i \in H} \sum_{j \in J} c_{i,j}\, y_{i,j} \tag{6}$$

Notice that, a retailer can either have excess supply or excess

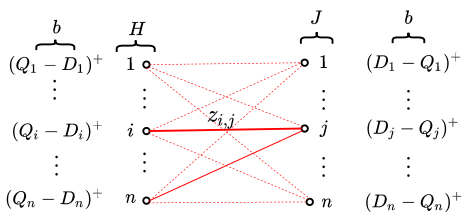

Figure 1: Constructed edge-weighted bipartite graph $G = (H, J, E)$ and b

demand, which means at least one of b-values $(D_j - Q_j)^+$ and $(Q_i - D_i)^+$ is equal to zero. Therefore, $y_{i,i} = 0$, for all $i \in R$. Mathematically, the b-matching problem can be written as follows.

$$\begin{aligned}
&\underset{y}{\operatorname{argmax}} \left(\sum_{i \in H} \sum_{j \in J \setminus \{i\}} w_i y_{i,j} (\mathfrak{p}_{ij} - s_i) \right. \\
&\qquad \left. + \sum_{i \in H} \sum_{j \in J \setminus \{i\}} w_j y_{i,j} (p_j - \mathfrak{p}_{ij} - \tau_{ij} + \rho_j) \right) \\
&s.t. \sum_{i \in H} \sum_{j \in J \setminus \{i\}} y_{i,j} \leq (Q_i - D_i)^+ \quad \forall i \in H \\
&\sum_{i \in H \setminus \{j\}} y_{i,j} \leq (D_j - Q_j)^+, \quad y_{i,j} \geqslant 0 \quad \forall i \in H, j \in J
\end{aligned} \tag{7}$$

Finding optimal solution (say y^*) of OP (Equation 7) is computationally equivalent to finding optimal allocation (say $\mathfrak{a}^*_{ij}$) for **WVT** in the following form.

$$\begin{aligned}
&\underset{A \in \mathcal{A}}{\operatorname{argmax}} \left(\sum_{i \in \mathcal{R}} \sum_{k \in \mathcal{R} \setminus \{i\}} w_i a_{ik} (\mathfrak{p}_{ik} - s_i) \right. \\
&\qquad \left. + \sum_{i \in \mathcal{R}} \sum_{l \in \mathcal{R} \setminus \{i\}} w_i a_{li} (p_i - \mathfrak{p}_{li} - \tau_{li} + \rho_i) \right) \\
&\text{s.t.} \sum_{k \in \mathcal{R} \setminus \{i\}} a_{ik} \leqslant (Q_i - D_i)^+ \quad \forall i \in \mathcal{R} \\
&\sum_{l \in \mathcal{R} \setminus \{i\}} a_{li} \leqslant (D_i - Q_i)^+, \quad a_{ik} \geqslant 0, \quad \forall i, k \in \mathcal{R}
\end{aligned} \tag{8}$$

Hence, Theorem 3 is proved. ■

The following section considers the performance of **WVT** for certain metrics that are not captured theoretically.

5 Experimental Results

The theoretical results ensure truthfulness and participation guarantees of **WVT**. However, other social welfare metrics, e.g., equitability, surplus in the budget, and efficiency, have not been theoretically captured. We carry out an experimental study using synthetic data to understand how the **WVT** performs on those metrics.

Our first experiment shows the need to choose appropriate weights in **WVT**.

5.1 Network position effect on utility

When we discuss the weights in the **WVT**, a natural question arises: "*why do we need different weights?*". To answer this, we consider the average utility of the retailers partitioned w.r.t. their network centrality measure. Consider the particular case when the weight for every retailer is unity. It reduces **WVT** to the **VCG** mechanism [6, 17, 33] which provides an allocatively efficient (one that maximizes the sum of the valuations of all agents) transshipment allocation.

We consider an Erdős–Rényi (ER) graph with 10 retailers having the edge forming probability of 0.7 to emulate the transshipment network. We consider the *closeness centrality* [3] as the measure of the retailers' positional impact on the network. The retailers in the network are partitioned into bins based on the closeness centrality, and the average value of the utilities of the retailers in each bin is computed.

For every $i \in \mathcal{R}$, we chose the following parameters: $b_i = 15$, $p_i = 30$, $\tau_{ik} = 15$, $\rho_i = 10$, $s_i = 10$. We generate D_i and Q_i from Normal distribution with mean $(\mu) = 500$ and standard deviation $(\sigma) = 50$, which shows the uncertainty in demand, and the estimated quantity of products to order from the manufacturer.

We randomly generate 1200 ER networks, and for each network, generate 200 instances of D_i and Q_i for every retailer i. Figure 2 shows the average utility plot w.r.t. the centers of the centrality bins. Notice that the utility increases with closeness centrality even though all the retailers have identical statistical and parametric properties (plots with other centrality measures are also similar). This experiment clearly shows the inequality introduced in the net payoff of the retailers due to their network positions and serves as the motivation for the design of the weights. The weights need to be decreasing in the network centrality so that the utilities earned by the retailers having identical statistical and parametric properties are more equalized. This implies that the *equal* retailers are treated more *equally*.

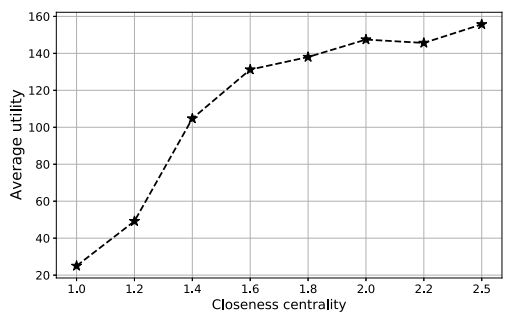

Figure 2: Utility vs centrality plot under **VCG**.

5.2 Equitability and efficiency

Following the definition of *egalitarian* allocation, where the objective is to maximize the utility of the most unfortunate individuals in

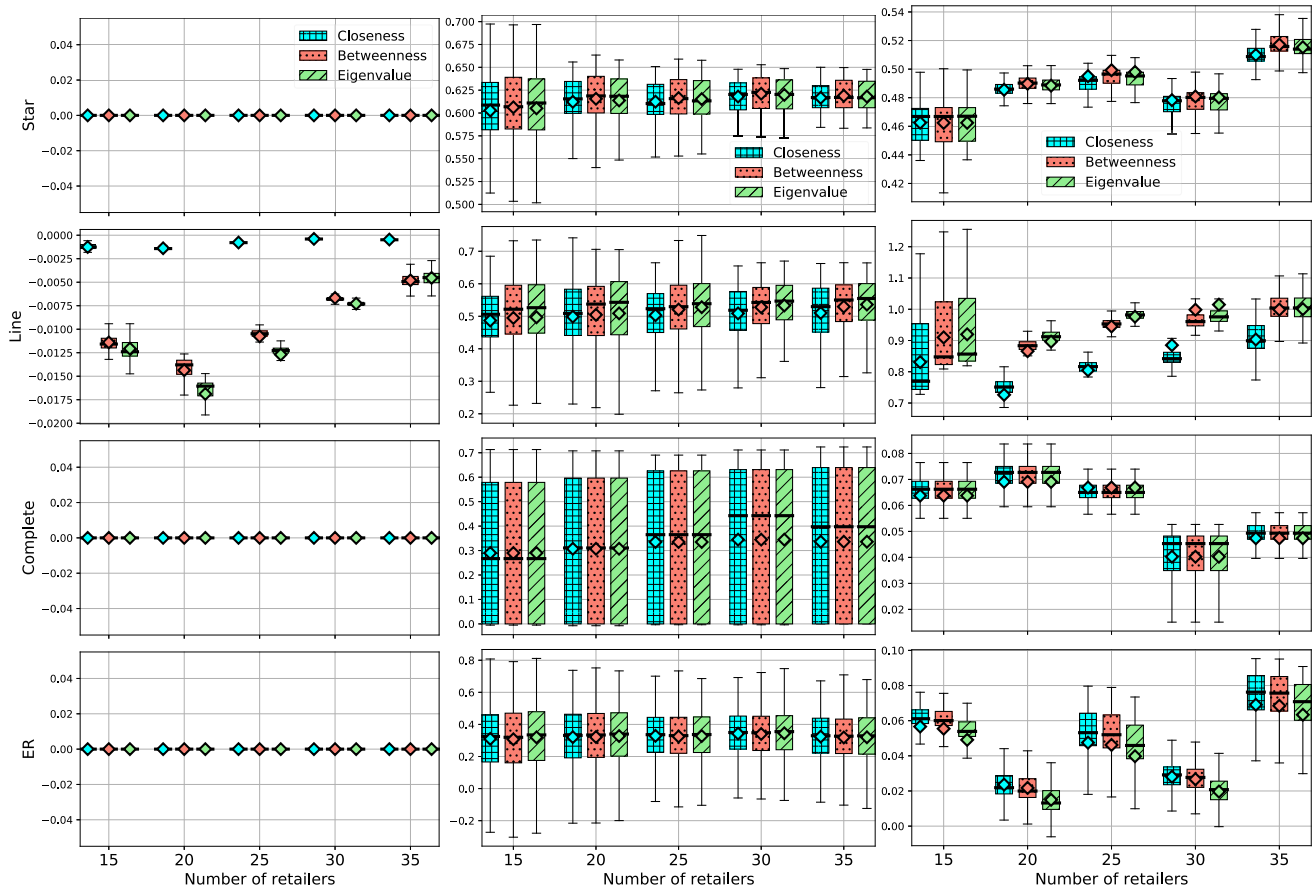

(a) Inefficiency Factor (IF) (b) Equitability Factor (EF) (c) Fractional Budget Surplus (FBS)

Figure 3: Performance of **WVT** under different metrics. Horizontal black lines and the diamond markers inside the boxes denote the median and the mean, respectively.

society so that every agent gets the same welfare level [23, 26], we define the *inequitability* ($\mathcal{I}$) of a TM f for an input instance $\mathcal{Z}$ as the variance[6] of the utilities of the retailers. Mathematically,

$$\mathcal{I}(f, \mathcal{Z}) = \texttt{var}([u_i(f(\mathcal{Z}), \mathcal{Z}_i)]_{i \in \mathcal{R}}) \tag{9}$$

A large value of $\mathcal{I}$ indicates that utilities of the retailers are significantly different from each other, whereas a small $\mathcal{I}$ indicates the opposite. $\mathcal{I} = 0$ means that the retailers get equal utilities.

While it is clear from the discussions in the previous section that the weights need to decrease with the centrality measures, it is unclear how the decreasing function should look. In this section, we attempt to heuristically choose a function and learn the parameters to reduce the inequitability (Equation 9) to a certain extent. The problem of finding an optimal weight vector that minimizes inequitability remains an open problem. The weight function we choose is

$$w_i = e^{-\alpha c_i} + \beta, \tag{10}$$

where c_i is the centrality of retailer i in the network. For the experiments, we consider *three* widely used centrality measures: *closeness*, *betweenness* [11], and *eigenvector* [25]. The mechanisms we consider in the **WVT** class will use the weights corresponding to these centralities using Equation (10).

To capture the equitability introduced by a mechanism, we first define the *equitability factor* (EF) of a TM f for an input instance $\mathcal{Z}$ as follows.

$$\text{EF}(f, \mathcal{Z}) = 1 - \frac{\mathcal{I}(f, \mathcal{Z})}{\mathcal{I}(\textbf{VCG}, \mathcal{Z})} \tag{11}$$

A mechanism with a higher value of EF is more equitable.

Since the weights of the **WVT** mechanisms can be different, it may not always yield an efficient outcome like the **VCG**. The following metric captures the *inefficiency factor* (IF) of a TM f for an input instance $\mathcal{Z}$.

$$\text{IF}(f, \mathcal{Z}) = \frac{\sum_{i \in \mathcal{R}} v_i(\textbf{VCG}(\mathcal{Z}), \mathcal{Z}_i) - \sum_{i \in \mathcal{R}} v_i(f(\mathcal{Z}), \mathcal{Z}_i)}{\sum_{i \in \mathcal{R}} v_i(\textbf{VCG}(\mathcal{Z}), \mathcal{Z}_i)} \tag{12}$$

VCG provides an efficient outcome; therefore IF can never be greater than zero. A larger negative value will imply that the mechanism is more inefficient than the **VCG**. For brevity, we will omit the arguments of the above factors wherever they are clear from the context.

We compare the results given by the above two metrics for *three* **WVT** mechanisms for the choices of weights corresponding to three centrality measures on *four* standard network structures, viz., star, line, complete, and ER networks (with edge forming probability to be 0.5). Figures 3(a) and 3(b) show the IF and EF plots respectively for different network structures with the increasing number of retailers.

[6] Note that another plausible inequitability notion can be $\max_i(u_i) - \min_i(u_i)$. We do not use that notion because in the presence of an outlier retailer dealing with a very large or very small number of demand/supply, $\max_i(u_i) - \min_i(u_i)$ can be a significantly large value, even if all the other retailers have equal utilities and, therefore, $\max_i(u_i) - \min_i(u_i)$ captures less information about the inequity than the variance of the utilities.

All the other parameters except the network positions are chosen identical for every retailer to analyze the effect of network positions. Therefore, the vertices having the same network position receive statistically identical utilities. For every retailer i the parameters are given by: $\rho_i = 10$, $b_i = 20$, $p_i = 50$, $s_i = 5$. For every edge (i, k) in the network $\tau_{ik} = 10$. We generate D_i and Q_i from $\mathcal{N}(500, 50)$. The line network has the highest diameter amongst any connected graph. For a sufficiently large number of retailers on a line network with high edge costs of $\tau_{ik} = 10$, the number of transshipment is insignificant. This is because every retailer has very few retailers with whom the transshipment is beneficial. To allow for the possibility of sufficient transshipment for analysis in a large line network, we choose a low transportation cost per edge ($\tau_{ik} = 1$). The weights are computed via Equation 10 (with $\alpha = 0.5, \beta = 1$) for the three chosen centrality measures as depicted in Figure 3.

The experiments are repeated 500 times for star, line, and complete networks, by generating random instances of D_i and Q_i $\forall i \in \mathcal{R}$. For ER network, we repeat the experiments for 2500 times, by generating 50 (D_i, Q_i) pairs $\forall i \in \mathcal{R}$, and for each such instance generating 50 ER random networks.

From the results of the different number of retailers shown in Figure 3(b), we find that the **WVT** reduces the inequitability (from **VCG**) by about 60% in star, 50% in line, and 30% in both complete and ER networks in the case of closeness centrality. The results are similar for betweenness and eigenvector centralities as well. It is interesting to note that this does not come at a big sacrifice in efficiency. For the chosen parameters, only in line networks, **WVT** compromises up to 2% of the efficiency for eigenvector and betweenness centrality and no significant loss in case of closeness centrality. For star, complete and ER networks, **WVT** makes no compromise in efficiency.

From these results, we conclude that it is possible to transship among the retailers reducing the inequitability due to network positions in a truthful, self-participatory manner without a significant compromise in the efficiency.

Discussion on the IF plots (Figure 3(a)): It is interesting to note that the IF is very close to zero, which is an effect of the event that the social welfare, i.e., the sum of the valuations, was almost the same in different graphs for most of the random instances. This happens even though the allocations in **WVT** and **VCG** were not the same always. The different allocations by **WVT** and **VCG** ($\mathcal{A}^{\texttt{WVT}}$ and $\mathcal{A}^{\texttt{VCG}}$ respectively) do change the individual utilities of the retailers. However, in the experiment with the chosen parameters, we found that the total quantity of the transshipment, i.e., $\sum_{i,k\in\mathcal{R}} a_{ik}$ is almost same in $\mathcal{A}^{\texttt{WVT}}$ and $\mathcal{A}^{\texttt{VCG}}$. As all the individual parameters (except the transportation costs, which are identical over the edges) are identical for every retailer in both the allocations, the difference in the social welfare is only due to the difference in total transportation cost $\sum_{i,k\in\mathcal{R}} \tau_{ik} a^{\texttt{WVT}}_{ik}$ and $\sum_{i,k\in\mathcal{R}} \tau_{ik} a^{\texttt{VCG}}_{ik}$. Since these values are insignificant in comparison to the optimal social welfare ($\sum_{i\in\mathcal{R}} v_i(\texttt{VCG}(\mathcal{Z}), \mathcal{Z}_i)$), the IF looks arbitrarily close to zero in the figure. We could have chosen a larger value of τ_{ik}, which needs to be large enough to be comparable to the optimal social welfare. But such a large value of τ_{ik} reduces the quantity of transshipment significantly, making the need of the **WVT** insignificant. Hence, even if the allocation by **WVT** and **VCG** are very distinct, the change in the social welfare is insignificant.

5.3 Budget surplus

While the monetary transfers in these mechanisms serve as an instrument to ensure truthfulness, it is desirable that the mechanism designer do not earn a significant surplus of these payments or run into a large deficit to run the mechanism. Ideally, one would like to have the sum of all these payments to be zero (which means the money is only redistributed), and we call such mechanisms to be *budget balanced*. However, in mechanisms with monetary transfers, ensuring both efficiency and budget balance is not generically possible [16]. In the **WVT**, there are two components of the monetary transfer: (a) the transshipment prices computed by **WVT**, which are one-to-one transactions between the retailers, and hence, are budget balanced ($\sum_{i,k\in\mathcal{R}} \mathfrak{p}_{ik} = 0$) by design, and (b) the side-payments, $\mathcal{P}_i$s, which exist to ensure certain desirable properties of the mechanism, e.g., truthfulness. However, in this setup, the positive surplus of $\sum_{i\in\mathcal{R}} \mathcal{P}_i$ has an advantage since it can be easily distributed to the customers on the platform (who are not the players in this mechanism) as gift coupons or monetary discounts and the mechanism can be budget balanced. If the surplus is negative, i.e., resulting in a deficit, we need a larger value of $\sum_{i\in\mathcal{R}} \mathcal{P}_i$ so that the deficit can be minimized, for any mechanism f and input instance $\mathcal{Z}$.

Therefore, a larger value of $\sum_{i\in\mathcal{R}} \mathcal{P}_i$ is more preferred. We capture how much a transshipment function f increases the surplus over **VCG** using *fractional budget surplus* (FBS) factor defined as follows.

$$\text{FBS}(f, \mathcal{Z}) = \frac{\sum_{i\in\mathcal{R}} \mathcal{P}^f_i(\mathcal{Z}) - \sum_{i\in\mathcal{R}} \mathcal{P}^{\texttt{VCG}}_i(\mathcal{Z})}{\sum_{i\in\mathcal{R}} v_i(\texttt{VCG}(\mathcal{Z}), \mathcal{Z}_i)} \tag{13}$$

Therefore, an FBS factor of 0.05 implies that the surplus under **WVT** increases by 5% of the optimal welfare than that of **VCG** under the same instance $\mathcal{Z}$. The optimal welfare in all experiments was always positive. If the TM is IR (satisfied by both **WVT** and **VCG**), then the payments should always be at most the valuation. Hence the $\sum_{i\in N} P^f_i \leq \sum_{i\in\mathcal{R}} v_i(f(\mathcal{Z}), \mathcal{Z}_i)$ for any IR TM f. This allows us to compare them and the ratio FBS in terms of percentage.

Figure 3(c) shows that for star and line networks, the budget surplus increases by $45-52\%$ and $70-100\%$ respectively. For complete and ER networks, the increase is of $4-7\%$ and $2-7\%$ respectively. The results for the three centrality measures are similar.

6 Conclusions and Future Work

In this paper, we study the excess inventory or unmet demands faced in the supply chain by the retailers who interact with the consumers on a shopping platform. The retailers are connected by a transportation network and have multi-dimensional *private* information. We provided a mechanism called **WVT** that satisfies truthfulness and individual rationality. Experimentally, we show that it does not compromise on the efficiency too much and has a better equitability factor and budget surplus than the classic **VCG** mechanism. Therefore, the **WVT** mechanism provides an engineering solution to almost all desirable properties required in a practical multi-retailer supply chain networks. Future work includes studying other aspects of fairness such as reward or compensation, and finding group-strategyproof mechanisms to avoid coalitional deviations.

Acknowledgments

Nath is supported by a MATRICS grant (MTR/2021/000367) from Science and Engineering Research Board, Govt. of India, a Technocraft Center for Applied Artificial Intelligence fund (DO/2021-TCAI002-009), and an Indian Institute of Technology Bombay fund (2022058).

Ethics Statement. This paper adheres to the principles of research ethics, research integrity, and social responsibility. The research conducted in this paper is purely theoretical, and no human or animal subjects were involved in the research. It addresses an essential aspect of resource wastage and economic loss in the supply chain due to the non-cooperation between multiple retailers or suppliers. Monetary transfer between the agents is well accepted in the market design. Therefore, there were no ethical concerns, in particular, related to the treatment of participants or the use of personal data.

References

[1] Richard P. Anstee, 'A polynomial algorithm for b-matchings: An alternative approach', *Information Processing Letters*, **24**(3), 153–157, (1987).

[2] Ravi Anupindi, Yehuda Bassok, and Eitan Zemel, 'A general framework for the study of decentralized distribution systems', *Manufacturing & Service Operations Management*, **3**(4), 349–368, (2001).

[3] Alex Bavelas, 'Communication patterns in task-oriented groups', *The journal of the acoustical society of America*, **22**(6), 725–730, (1950).

[4] Tilman Börgers, *An introduction to the theory of mechanism design*, Oxford University Press, USA, 2015.

[5] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia, *Handbook of computational social choice*, Cambridge University Press, 2016.

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

[7] Partha Dasgupta and Eric Maskin, 'Efficient Auctions*', *The Quarterly Journal of Economics*, **115**(2), 341–388, (05 2000).

[8] Morton Davis and Michael Maschler, 'The kernel of a cooperative game', *Naval Research Logistics Quarterly*, **12**(3), 223–259, (1965).

[9] Banu Yetkin Ekren, Sachin Kumar Mangla, Ecem Eroglu Turhanlar, Yigit Kazancoglu, and Guo Li, 'Lateral inventory share-based models for iot-enabled e-commerce sustainable food supply networks', *Computers & Operations Research*, **130**, 105237, (2021).

[10] Andrew Fearne, David Yawson, Abbi Buxton, and Justin Tait, 'Measuring fairness in supply chain trading relationships', (2012).

[11] Linton C Freeman, 'A set of measures of centrality based on betweenness', *Sociometry*, 35–41, (1977).

[12] Ranjan Ghosh and Mattias Eriksson, 'Food waste due to retail power in supply chains: Evidence from sweden', *Global Food Security*, **20**, 1–8, (2019).

[13] Allan Gibbard, 'Manipulation of voting schemes: a general result', *Econometrica: journal of the Econometric Society*, 587–601, (1973).

[14] Julian Le Grand, 'Equity versus efficiency: The elusive tradeoff', *Ethics*, **100**(3), 554–568, (1990).

[15] Daniel Granot and Greys Sošić, 'A three-stage model for a decentralized distribution system of retailers', *Operations research*, **51**(5), 771–784, (2003).

[16] Jerry Green and Jean-Jacques Laffont, *Incentives in public decision-making*, Elsevier North-Holland, 1979.

[17] T. Groves, 'Incentives in teams', *Econometrica*, **41**, 617–631, (1973).

[18] Behzad Hezarkhani and Wieslaw Kubiak, 'A coordinating contract for transshipment in a two-company supply chain', *European Journal of Operational Research*, **207**(1), 232–237, (2010).

[19] Sue Hornibrook, Andrew Fearne, and Maria Lazzarin, 'Exploring the association between fairness and organisational outcomes in supply chain relationships', *International Journal of Retail & Distribution Management*, (2009).

[20] Eda Kemahlıoğlu-Ziya and John J Bartholdi III, 'Centralizing inventory in supply chains by using shapley value to allocate the profits', *Manufacturing & Service Operations Management*, **13**(2), 146–162, (2011).

[21] Sebastián Lozano, Placido Moreno, Belarmino Adenso-Díaz, and Encarnacion Algaba, 'Cooperative game theory approach to allocating benefits of horizontal cooperation', *European Journal of Operational Research*, **229**(2), 444–452, (2013).

[22] Hervé Moulin, *Fair division and collective welfare*, MIT press, 2003.

[23] Roger B. Myerson, 'Utilitarianism, egalitarianism, and the timing effect in social choice problems', *Econometrica*, **49**(4), 883–897, (1981).

[24] Dilupa Nakandala, Henry Lau, and Paul K.C. Shum, 'A lateral transshipment model for perishable inventory management', *International Journal of Production Research*, **55**(18), 5341–5354, (2017).

[25] Mark EJ Newman, 'The mathematics of networks', *The new palgrave encyclopedia of economics*, **2**(2008), 1–12, (2008).

[26] John Rawls, 'Some reasons for the maximin criterion', *The American Economic Review*, **64**(2), 141–146, (1974).

[27] Kevin Roberts, 'The characterization of implementable choice rules', *Aggregation and revelation of preferences*, **12**(2), 321–348, (1979).

[28] Nils Rudi, Sandeep Kapur, and David F Pyke, 'A two-location inventory model with transshipment and local decision making', *Management science*, **47**(12), 1668–1680, (2001).

[29] Mark Allen Satterthwaite, 'Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions', *Journal of economic theory*, **10**(2), 187–217, (1975).

[30] Jing Shao, Harish Krishnan, and S Thomas McCormick, 'Incentives for transshipment in a supply chain with decentralized retailers', *Manufacturing & Service Operations Management*, **13**(3), 361–372, (2011).

[31] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2008.

[32] A.C.C. van Wijk, I.J.B.F. Adan, and G.J. van Houtum, 'Optimal lateral transshipment policies for a two location inventory problem with multiple demand classes', *European Journal of Operational Research*, **272**(2), 481–495, (2019).

[33] W. Vickrey, 'Counter speculation, auctions, and competitive sealed tenders', *Journal of Finance*, **16**(1), 8–37, (1961).

[34] Yusheng Wang, Zhijie Sasha Dong, and Shaolong Hu, 'A stochastic prepositioning model for distribution of disaster supplies considering lateral transshipment', *Socio-Economic Planning Sciences*, **74**, 100930, (2021).

[35] Xinghao Yan and Hui Zhao, 'Inventory sharing and coordination among n independent retailers', *European Journal of Operational Research*, **243**(2), 576–587, (2015).

[36] Meimei Zheng, Ningxin Du, HUI ZHAO, Edward Huang, and Kan Wu, 'A study on the optimal inventory allocation for clinical trial supply chains', *Applied Mathematical Modelling*, (2021).