



Atasu, A., Toktay, B., Yeo, W. M. and Zhang, C. (2017) Effective medical surplus recovery. *Production and Operations Management*, 26(6), pp. 1142-1162.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

This is the peer reviewed version of the following article: Atasu, A., Toktay, B., Yeo, W. M. and Zhang, C. (2017) Effective medical surplus recovery. *Production and Operations Management*, 26(6), pp. 1142-1162, which has been published in final form at <http://dx.doi.org/10.1111/poms.12641>. This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#).

<http://eprints.gla.ac.uk/155616/>

Deposited on: 19 January 2018

Effective Medical Surplus Recovery

We analyze not-for-profit Medical Surplus Recovery Organizations (MSROs) that manage the recovery of surplus (unused or donated) medical products to fulfill the needs of underserved healthcare facilities in the developing world. Our work is inspired by an award-winning North American non-governmental organization (NGO) that matches the uncertain supply of medical surplus with the receiving parties' needs. In particular, this NGO adopts a recipient-driven resource allocation model, where it grants recipients access to an inventory database, and each recipient selects products of limited availability to fill a container based on its preferences. We first develop a game theoretic model to investigate the effectiveness of this approach. This analysis suggests that the recipient-driven resource allocation model may induce competition for MSRO supplies among recipients and lead to a loss in value provision through premature orders. Accordingly, we provide a number of operational mechanisms that can help MSROs deal with this problem. These mechanisms are: (i) appropriately selecting container capacities while limiting the inventory availability visible to recipients and increasing the acquisition volumes of supplies, (ii) eliminating recipient competition through exclusive single-recipient access to the MSRO inventory, and (iii) focusing on learning about recipient needs as opposed to providing them with supply information, and switching to a provider-driven resource allocation model. We use real data from the NGO by which the study was inspired and show that the proposed improvements can substantially increase the value an MSRO provides to its recipients.

Key words: Sustainable Operations, Humanitarian Supply Chains, Reverse Logistics, Medical Surplus

1. Introduction

Healthcare facilities in the United States throw away 5.9 million tons of reusable medical products every year (Berry 2014). Also known as medical surplus, these reusable products include leftovers from post-surgical procedures, unopened clinical kits that are discarded due to regulatory requirements (EPA 40 CFR259.30 Regulated Medical Waste Act), and surplus or used hospital equipment. Recognizing that these products can be used to provide healthcare for the underserved (Perry and Malkin 2011), their recovery and shipment to healthcare facilities in developing countries is a promising opportunity for the humanitarian aid world. This opportunity has resulted in the formation of a unique reverse supply chain concept, where the recovery and redistribution of medical surplus (excluding perishables such as pharmaceuticals) is undertaken by Medical Surplus Recovery Organizations (MSROs), who not only leverage unused medical products in hospitals in developed countries but also work with corporate donation programs of medical equipment manufacturers such as McKesson and Kimberly Clark to acquire unsold or returned medical equipment.

While the opportunity to allocate medical surplus to the underserved regions of the world is immense, field studies on MSROs suggest that the lack of operational expertise to match recipient demand and the uncertain supply is a major concern for MSROs (Compton 2012a, Howitt et al. 2012, Kotsi et al. 2014). MSRO supply chains face unique operational challenges that differ from traditional for-profit supply chains: First, the MSRO objective is not profit maximization (recipients do not pay for the goods supplied by the MSROs). Rather, MSROs aim to maximize value provision to their recipients. Second, due to their non-profit nature that heavily relies on donations, MSROs operate in cash-constrained environments (Beta 2012), and cost effectiveness is a key objective for an MSRO. Hence, MSROs aim to maximize the value provided to recipients by shipping full containers because container shipments are a major operational cost category. Third, medical surplus flow into MSROs is often limited, uncertain and not controllable; MSROs do not have a formal procurement or production function as in a traditional supply chain. Fourth, MSROs serve large recipient bases with a diversity of needs that dynamically change, where anticipating the exact needs of recipients can be challenging.

In sum, the problem faced by MSROs is to maximize value provision to their recipients in the absence of controllable supply, with limited information as to recipient needs, in a cash-constrained environment. This is a challenging environment where innovative operating models tailored to these conditions are needed to maximize value provision. The limited information about recipient needs is a particularly important barrier, resulting in major loss of effectiveness in practice: The World Health Organization (WHO) estimates that over seventy percent of donated medical equipment was inappropriate for the beneficiaries (Howitt et al. 2012). The inspiration for this study, a Southern US-based MSRO, referred to as Beta (name disguised for confidentiality), has developed a recipient-driven resource allocation model to address this concern: Beta allows its recipients to determine the content of shipments they are to receive. In particular, once Beta secures funding for container shipments to particular recipients through monetary (or in-kind service) donations, it provides them with online access to inventory information. The recipients then determine their container content based on their preferences and needs. Essentially, the Beta model relies on recipients' need assessment by providing them with the opportunity to select their own supplies to minimize the potential supply and demand mismatch in this challenging environment. Beta has recovered more than 2 million pounds of medical surplus to date through an extensive network of partnerships with a variety of entities in the healthcare industry and has shipped more than one thousand 40-foot containers to a diverse set of recipients. These containers have typically served distinct medical facilities in developing countries and humanitarian organizations running development (and occasionally, disaster response) programs in different parts of the world, including the Middle East, South America and Africa. Beta is recognized for its operational excellence, as demonstrated by

numerous operational excellence awards it has received, thanks to its focus on the recipient-driven approach.

Our objective in this paper is to provide an operational roadmap for MSROs by building and improving upon Beta’s operational model. In particular, we aim to investigate: (i) When a recipient-driven resource allocation model is a good operational practice for an MSRO, and (ii) Whether any operational mechanisms can be used to improve upon the effectiveness of recipient-driven MSRO operations. We address these questions in two steps: We first model the recipient-driven resource allocation model of Beta using a game-theoretic model, which takes into account the perspectives of recipients that are given simultaneous access to Beta’s inventory. This analysis suggests that when recipients’ needs are time sensitive, the recipient-driven allocation model works at its best. Otherwise, i.e., when the recipients can afford to wait longer for MSRO supplies, we identify a potential source of effectiveness loss under the recipient-driven resource allocation model. In particular, we find that competition among recipients induced by simultaneous multi-recipient access to an MSRO’s inventory can result in a reduction in value provision.

This finding is supported by an analysis of order data obtained from Beta, which suggests that recipient-driven resource allocation may lead the recipients to fill their containers very fast with lower value items even if they can afford to wait longer for MSRO supplies. Our analysis also demonstrates that certain operational mechanisms can be used to overcome this loss of effectiveness under the recipient-driven resource allocation model. These mechanisms are (i) limiting the inventory availability visible to recipients, (ii) appropriately selecting container capacities, (iii) increasing the acquisition volume of supplies, and (iv) eliminating recipient competition through exclusive single recipient access to the MSRO inventory. Alternatively, provided that the MSRO can build the capability to assess recipient needs, focusing on learning about recipient needs as opposed to providing them with supply information and switching to a provider-driven resource allocation model can improve the MSRO’s value provision capability. Our investigation regarding the value of these operational mechanisms using the data from Beta suggests that they can help MSROs substantially improve their value provision capabilities.

These results collectively suggest that well-established supply chain management principles may not translate to the MSRO context or may need adaptation to generate value. For example, it is known that a “pull” model has inventory management and demand-supply matching advantages in a traditional supply chain. However, this concept cannot be fully applied to the MSRO setting because the MSRO does not produce or buy items to recipient specifications. Instead, it can be applied by giving the recipient control of which items it orders from existing inventory by sharing inventory information. In implementing this idea, common wisdom from supply chains would suggest full inventory visibility is ideal. Yet full information visibility accelerates premature

ordering in this context. Moreover, the level of simultaneous access to information needs to be controlled because the potential recipient competition prevents recipients from waiting for the arrival of preferable items. Rather, information provision in the reverse direction appears to be more valuable in the MSRO setting: That is, learning about recipient needs (as opposed to providing supply information to the recipients) can substantially improve an MSRO's value provision.

2. Literature Review

This work draws on and contributes to the operations management (OM) literature that deals with allocation of scarce resources in the presence of uncertainty. Contrary to a conventional operational setting, in which the uncertainty in demand can often be mitigated by inventory or capacity as hedges, our work looks at a particular setting in which the demand is ample while the supply is often limited, uncertain and uncontrollable, a distinguishing characteristic of the MSRO environment.

Our work is closely related to the recently growing literature on medical surplus allocation, which typically documents best practices in the field through white papers produced within the MSRO community (Compton 2011, 2012b). This work has so far focused on measuring recipients' medical surplus utilization rates and how surplus recovery data can improve supply chains for the operating room (MacRae 1997, Rosenblatt and Silverman 1994, Rosenblatt et al. 1997). Yet, an investigation regarding the effectiveness of different resource allocation practices has not been explored to date in the medical surplus context (Kotsi et al. 2014), despite recipient-driven allocation being touted as an exemplary approach. In this paper, we fill this void by exploring the effectiveness of the recipient-driven resource allocation approach, and identifying a number of operational mechanisms that can be used to further improve the welfare provision capability of MSROs.

Our contribution also relates to the stream of research that focuses on resource allocation in healthcare, particularly those that deal with organ transplantation. In this context, organ scarcity and large recipient bases are key concerns, and the fact that potential organ recipients have decision rights in the organ allocation process (Su and Zenios 2004, 2005, Zenios et al. 2000, Akan et al. 2012, Bertsimas et al. 2013) resembles the characteristics of the recipient-driven resource allocation model in an MSRO setting. A key distinction between the two settings is that organ recipients have organ refusal rights, which may lead the recipients to decline an available offer and wait for a better one, leading to a loss in value provision stemming from the perishable nature of the donations. In the recipient-driven allocation model of MSROs, however, the main source of allocation effectiveness loss stems from recipients' premature ordering behavior, i.e., not waiting for items that may be a better fit for their needs. This is because recipient-driven medical surplus allocation implies a constrained environment where recipients compete to be the first to fill a container using a mix of multiple product categories before others, leading to container mixes that may contain products of little value for the recipients.

Our work also relates to the recently growing stream of research that deals with resource allocation problems in the broader humanitarian logistics field (see Altay and Green (2006), Van Wassenhove (2006), Van Wassenhove and Pedraza Martinez (2011), Kunz and Reiner (2012) for overviews) that relate to disaster preparedness and last-mile distribution. To name a few studies in the context of disaster preparedness, Salmeron and Apte (2010) and Duran et al. (2011) explore the pre-disaster establishment of capacity and resources to maximize post-disaster humanitarian aid; Özdamar et al. (2004) use a vehicle routing framework to develop models that optimally construct resource dispatch plans for humanitarian platforms to minimize unmet humanitarian demand; Bhattacharya et al. (2014) focus on coordinating multiple aid programs to maximize social welfare through the design of procurement and allocation systems for expensive resources; Ergun et al. (2013) analyze the value of collaboration across multiple humanitarian agencies; Pedraza Martinez et al. (2010) and Besiou et al. (2014) investigate vehicle and fleet management issues. To name a few studies in the context of last-mile distribution (see de la Torre et al. (2012) for a recent overview), Huang et al. (2012) propose and develop heuristics to explore the joint problem of vehicle routing and supply allocation decisions under cost, speed and equity trade-offs; and Balcik et al. (2008) have a cost efficiency focus in developing resource delivery schedules to minimize transportation costs in a centralized decision making system (unlike our work, which has a value provision/effectiveness focus and compares decentralized and centralized decision making approaches); while Natarajan and Swaminathan (2014) and Taylor and Xiao (2014) analyze inventory control and distribution channel management problems under funding constraints. We differ from these two streams of work in two dimensions. First, they typically focus on managing controllable resources to maximize humanitarian benefits, whereas we focus on an environment with uncertain and uncontrollable supplies. Second, they focus on the NGO perspective, i.e., generally NGOs are the only and centralized decision makers. In our analysis, recipient decisions matter and we evaluate the effectiveness of empowering recipients in resource allocation.

We also note that the provider-driven resource allocation model we propose relates to the vendor managed inventory concept that has recently been popular in the humanitarian context. Also known as *pousse-pousse* or *informed-push* model, the contraceptive distribution program in Senegal (Daff et al. 2014) develops supply chain competencies among upstream policy-makers, allowing doctors to focus on clinical issues. To reduce stockouts, operational decisions such as warehouse item-mix, delivery schedules, and quantity decisions at each point-of-sale are built upon consumption and utilization data collected from drivers delivering medical supplies. In a similar approach, the USAID-Deliver project piloted the Direct Delivery and Information Capture system to distribute items to selected service points in the Ebonyi and Bauchi states of Nigeria. Serving as mobile warehouses, trucks carrying pre-determined quantities of health supplies are accompanied

by an expert who inspects storage space, counts on-hand inventories and enters the data into a specifically-designed database for decision making. It is nevertheless important to note that this setting considers controllable supply to a large extent, while the problem faced by MSROs is different due to the uncertainty in the volume and arrival times of supplies. Nevertheless, our analysis supports the view that keeping the decision rights on the NGO side can help significantly improve welfare provision.

Finally, our research relates to traditional operations management research with respect to rationing policies under limited supply (Cachon and Lariviere 1999), the value of information sharing (Gavirneni et al. 1999, Cachon and Fisher 2000), and push versus pull comparison perspectives (Hopp and Spearman 2008). The problem of supply rationing among multiple downstream customers and the gaming between suppliers and customers is well-studied in the operations community (Cachon and Lariviere 1999). The problem faced by MSROs, however, is substantially different from the traditional for-profit setting due to the welfare maximization perspective that an NGO takes facing uncertain supplies with limited control and demand that is always larger than supplies. As opposed to the strategically increased order sizes in the traditional setting (Cachon and Lariviere 1999), the gaming under the MSRO's recipient-driven resource allocation model leads to premature orders with suboptimal container mixes. It is also well known that pull models have inherent advantages over push models in for-profit supply chains as they establish limits on the work-in-progress inventory levels (Hopp and Spearman 2008) and that information sharing across the supply chain can help significantly improve the supply-demand match (Gavirneni et al. 1999, Cachon and Fisher 2000). What we observe in the MSRO context is that these basic insights in traditional supply chain settings do not necessarily apply in the allocation of medical surplus. Rather, the two following deviations from for-profit supply chains appear to be key: (i) Relying on recipients' wisdom through information sharing can result in loss of effectiveness due to the inherent recipient competition this practice may induce, and (ii) For MSROs, a push model may outperform a pull model, provided that recipient need assessment is cost-effective.

The rest of the paper is organized as follows. In §3, we present a stylized model that allows us to characterize the fundamental trade-offs in a recipient-driven MSRO setting and propose a number of operational improvement opportunities for MSROs. §4 identifies two new approaches to eliminate the drawbacks of recipient competition under the recipient-driven allocation model, and §5 uses real-life data to quantify the impact of proposed improvement approaches. §6 concludes with a discussion of insights, implications for practice and future research.

3. Recipient-Driven Resource Allocation

We first develop and analyze a basic recipient-driven allocation model (R), which mimics the operational infrastructure of Beta and allows us to capture the key trade-offs faced by an MSRO in

a recipient-driven setting. In practice, Beta first secures funding for container shipments through donations. After that, recipients chosen by Beta are given online access to Beta’s inventory information (typically two recipients have access at a time). These recipients then select their container content based on their needs. Once a full container order is completed by a recipient and confirmed by Beta, the inventory database is updated, and the other recipient in the system has to make its selections from the updated database. Accordingly, to mimic Beta’s business model in a discrete-time setting, we consider an MSRO that serves two recipients by maintaining inventories of two product categories (referred to as items) over two periods $t = 0, 1$ (see Appendix A for a multiple item generalization). The two recipients are provided simultaneous access at time $t = 0$. New items arrive at time $t = 1$.

Specification of the Game: We consider a two-stage dynamic simultaneous move game in which two recipients in distinct locations have access to the MSRO inventory. At $t = 0$, the inventory levels of both items (μ_1 and μ_2) and the container capacity K are known to all parties. To avoid trivial scenarios, we assume that $\mu_1 + \mu_2 \geq K$ (a container can be filled with the available inventory); and $\mu_1, \mu_2 < K$ (the container mix can be improved for both recipients by waiting until $t = 1$). At $t = 1$, a new batch of supplies arrives and increases the availability of items by $\epsilon = (\epsilon_1, \epsilon_2)$. We assume that the inventory is non-perishable, hence the supplies only increase unless a container is shipped to a recipient. This assumption reflects Beta’s acquisition guidelines, where products are acquired only if their time to expiration exceeds 12 to 18 months, during which they are typically depleted due to Beta’s rapid turnover rates.

Recipient Valuations of Items: Let $v_{i,j}$ denote the value of item j for recipient i . We say that item $j = 1, 2$ is high-value for recipient $i = 1, 2$ if $v_{i,j} > v_{i,-j}$. We assume that the recipients are heterogeneous in their preferences, and without loss of generality recipient i prefers item i , i.e., $v_{1,1} > v_{1,2}$ and $v_{2,1} < v_{2,2}$. For simplicity, we assume that the recipients’ valuations are public information.¹

Action Space and Strategy Profiles. We denote the action space of recipient i at time $t = 0$ to be $\{O, W\}$, where O stands for ordering at $t = 0$, and W for waiting (i.e., ordering at $t = 1$). We denote the set of pure strategy profiles at $t = 0$ to be $S = \{(S_1, S_2)\}$, where S_1 stands for the action chosen by recipient 1, and S_2 for the action chosen by recipient 2. As such, S consists of four possibilities: (O, O) , (O, W) , (W, O) and (W, W) . Given K and $\boldsymbol{\mu} = (\mu_1, \mu_2)$ at $t = 0$, recipients simultaneously choose either to order or to wait. If only one of the recipients (call it i) orders at $t = 0$, the game

¹ Results available from the authors show that the main insights continue to hold in an extension where we consider a game of incomplete information in which recipient valuations are private information.

stops. Recipient i chooses order volume $y_i = (y_{i,1}, y_{i,2})$ so as to maximize the container value it receives, π_i^0 , where

$$\begin{aligned} \pi_i^0 \doteq \max_{y_{i,1}, y_{i,2}} & v_{i,1}y_{i,1} + v_{i,2}y_{i,2} \\ \text{s.t. } & y_{i,1} \leq \mu_1, y_{i,2} \leq \mu_2, y_{i,1} + y_{i,2} \leq K. \end{aligned} \quad (1)$$

Given the assumptions above, it is straightforward to see that $\pi_i^0 = v_{i,i}\mu_i + v_{i,-i}(K - \mu_i) = \Delta_i\mu_i + v_{i,-i}K$, where $\Delta_i \doteq v_{i,i} - v_{i,-i}$. Meanwhile, recipient $-i$ does not receive a container and obtains zero payoff. This assumption is a simplification for tractability, and allows us to study the main effect of recipient competition.²

If both recipients order at $t = 0$, the MSRO chooses whom to ship to according to a tie-breaking rule (see the Expected Payoff Description below for details). If both recipients wait, ϵ realizes at $t = 1$, both recipients order and the game stops. The MSRO selects a recipient according to the same tie-breaking rule. If recipient i is selected, it orders $y_{i,j}$, $j = 1, 2$ that maximizes the container value from ordering in the second period based on the following:

$$\begin{aligned} \pi_i^1 \doteq \max_{y_{i,1}, y_{i,2}} & v_{i,1}y_{i,1} + v_{i,2}y_{i,2} \\ \text{s.t } & y_{i,1} \leq \mu_1 + \epsilon_1, y_{i,2} \leq \mu_2 + \epsilon_2, y_{i,1} + y_{i,2} \leq K. \end{aligned} \quad (2)$$

It is easy to show $\pi_i^1 = \Delta_i \min(\mu_i + \epsilon_i, K) + v_{i,-i}K$. Again, we assume the other recipient obtains zero payoff.

Expected Payoff Description. For a given strategy profile $(S_1, S_2) \in S$, we denote $\Pi_i^{S_1, S_2}$ as recipient i 's total expected discounted payoff. To complete the payoff description, let p represent the recipients' belief of the probability that recipient 1 will be chosen according to the tie-breaking rule, and that this belief is consistent with the MSRO's actions. Then, under simultaneous ordering, recipient i expects to receive a container with probability p_i and to not receive a container otherwise, where $p_1 = p$ and $p_2 = 1 - p$. This formulation is inspired by the MSRO we work with. Beta has a priority score for each recipient that is based on multiple factors such as the health needs assessment and the security risk of the country the recipient is from. Some of these factors, however, may not be known in advance. For instance, the political or civil stability of a recipient's country or customs clearance conditions can change the MSRO's preference between these recipients. The allocation probability p represents such randomness in the MSRO's decision heuristic (which we refer to as a tie-breaking rule). If there were no uncertainty in these factors, a deterministic priority score could be used to determine who gets the container when the two recipients order at the same time. In this

² Results available from the authors show that the main insights continue to hold in an extension where a second container can be shipped to the second recipient in the second period.

case, the allocation probability p would be either zero or one. In sum, through the probability p we allow a more general (although stylized) representation of the recipient selection process subject to random factors that are not known in advance but determine who gets the container.

Table 1 shows the total expected discounted payoffs of all strategy profiles. The table is constructed as follows: If both recipients order, $\Pi_i^{O,O} = p_i \pi_i^0$ represents the expected payoff of recipient i , since recipient i expects to get his desired shipment value π_i^0 with probability p_i and zero payoff otherwise. If one of the recipients waits, then the ordering recipient is allocated his order ($\Pi_1^{O,W} = \pi_1^0$, $\Pi_2^{W,O} = \pi_2^0$) and the other recipient's payoff is 0. Finally, $\Pi_i^{W,W} = \delta p_i E[\pi_i^1]$ represents the expected discounted payoff for recipient i whenever both recipients wait at $t = 0$. Here δ is a discount factor that captures the time value of the value provided to the recipients. This parameter can depend on a variety of factors. For instance, it would be substantially smaller in a disaster response setting when compared to a development setting. Likewise, it could depend on the baseline capabilities of a recipient: A hospital with a substantially higher budget could have a higher δ or another hospital that needs to deal with an epidemic could have a lower δ . In essence, δ is an abstraction of such factors into one single parameter in our stylized model. For ease of exposition, we assume that both recipients have the same δ . Our results can be extended to different discount factors without modifying the insights. Accordingly, for any given strategy profile (S_1, S_2) , we define the total welfare provision of the MSRO under the recipient-driven model as $W_R^{S_1, S_2} = \Pi_1^{S_1, S_2} + \Pi_2^{S_1, S_2}$, representing the total expected welfare provided to the two recipients.

Table 1 Payoff matrix under recipient-driven resource allocation.

		Recipient 2	
		Order	Wait
Recipient 1	Order	$\Pi_1^{O,O}, \Pi_2^{O,O}$	$\Pi_1^{O,W}, 0$
	Wait	$0, \Pi_2^{W,O}$	$\Pi_1^{W,W}, \Pi_2^{W,W}$

Given the model set-up, we now define a concept that we will use throughout our analysis: The (*conditional*) *value of waiting* for a recipient at $t = 0$ is the difference between the expected discounted value of waiting to order at $t = 1$ and the value of ordering at $t = 0$ given that the other recipient waits. The value of waiting (as inferred from Table 1) for recipient 1 and recipient 2 can be respectively written as:

$$\Pi_1^{W,W} - \Pi_1^{O,W} \text{ and } \Pi_2^{W,W} - \Pi_2^{W,O}. \quad (3)$$

We next present a useful benchmark with a single recipient i that has exclusive access to the MSRO inventory. In that case, it is straightforward to show that the recipient waits to order at $t = 1$ if and only if $\pi_i^0 < \delta E[\pi_i^1]$, or

$$\delta > \bar{\delta}_i \doteq \frac{\pi_i^0}{E[\pi_i^1]} = \frac{\Delta_i \mu_i + v_{i,-i} K}{E[\Delta_i \min(\mu_i + \epsilon_i, K) + v_{i,-i} K]}. \quad (4)$$

In this case, it is clear that if they were given exclusive access to the MSRO inventory, both recipients would wait if $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$. We next characterize the equilibrium under recipient-driven allocation, for which the condition $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$ serves as an important benchmark. A summary of key notation for the model variables is provided in Table 2.

Table 2 Summary of key notation.

Notation	Description
K	Container capacity.
μ_i	Inventory level of item i at $t = 0, i = 1, 2$.
ϵ_i	Amount of new arrivals of item i at $t = 1, i = 1, 2$.
$v_{i,j}$	Value of item j for recipient $i, i = 1, 2, j = 1, 2$.
p_i	Recipients' belief of the probability that recipient i will be chosen for tie-breaking.
π_i^0	Optimal container value for recipient i from receiving a container at $t = 0, i = 1, 2$.
π_i^1	Optimal container value for recipient i from receiving a container at $t = 1, i = 1, 2$.
$\bar{\delta}_i$	A threshold for recipient i , which is equal to $\pi_i^0/E[\pi_i^1], i = 1, 2$.
δ	Discount factor.
$\Pi_i^{S_1, S_2}$	Recipient i 's total expected discounted payoff for a given strategy profile (S_1, S_2)
$W_R^{S_1, S_2}$	Total expected discounted value provision for a given strategy profile (S_1, S_2) under R model.

3.1. Equilibrium Characterization of the R Model

PROPOSITION 1. *There exist two possible equilibria under the R model: (O, O) and (W, W) . (O, O) is always an equilibrium, and (W, W) is an equilibrium if $\delta > \bar{\delta} \doteq \max(\frac{\bar{\delta}_1}{p_1}, \frac{\bar{\delta}_2}{p_2})$. Furthermore, $W_R^{W, W} > W_R^{O, O}$ if $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$.*

Proposition 1 states that a Nash equilibrium always exists under recipient-driven allocation and is sustained by strategy profiles where both recipients choose the same actions.³ Both players ordering at $t = 0$ is always an equilibrium, while waiting can be sustained in equilibrium only if $\delta > \bar{\delta} \geq \max(\bar{\delta}_1, \bar{\delta}_2)$. This characterization combined with the single-recipient benchmark allows us to identify an important insight regarding the effectiveness of the R model: When $\delta \leq \max(\bar{\delta}_1, \bar{\delta}_2)$ (as defined in Equation (4)) both recipients ordering is the only equilibrium, and it is preferred from a welfare point of view when δ is sufficiently small. That is, if the discount rate δ is sufficiently small, i.e., recipient needs are very urgent, the R model is effective.

On the other hand, when $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$, the effectiveness of the R model can be undermined by recipient competition. This is because the (W, W) equilibrium with higher welfare provision may not exist in this range and more specifically when $\delta \in (\max(\bar{\delta}_1, \bar{\delta}_2), \bar{\delta})$, although the recipients would

³ All proofs are provided in Appendix B.

have preferred waiting if they were given exclusive access to the MSRO inventory. Accordingly, in what follows we restrict our analysis to the case with $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$ to explore the implications of recipient competition under the R model.

ASSUMPTION 1. $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$.

Given this assumption and the equilibrium characterization in Proposition 1, the MSRO faces two scenarios: (i) $\delta \in (\max(\bar{\delta}_1, \bar{\delta}_2), \bar{\delta})$ and (ii) $\delta > \bar{\delta}$. In the first scenario, although waiting is optimal for a recipient that does not face competition (per Assumption 1), only the (O, O) equilibrium exists. In other words, the first scenario is an example of the well-known prisoner's dilemma, and the unique (O, O) equilibrium is an undesirable one from the MSRO perspective. The second scenario presents an example of the stag-hunt game, where the MSRO observes two possible equilibria that represent risk-dominance (the (O, O) equilibrium) and payoff-dominance (the (W, W) equilibrium that is welfare improving). The key observation from the MSRO perspective is that if $\delta > \bar{\delta}$, the Pareto-efficient payoff dominant (W, W) equilibrium can exist. In this case, we assume that the MSRO can serve as a coordination medium between recipients to make sure that the (W, W) equilibrium is chosen.⁴ As such, our focus in what follows is to identify how the MSRO can offset $\bar{\delta}$ to expand its ability to induce the (W, W) equilibrium.

3.2. Inducing Recipient Waiting under the R Model

We start this section by noting that $\bar{\delta}$ depends on $\boldsymbol{\mu}$, K and $\boldsymbol{\epsilon}$, hereafter denoted by $\bar{\delta}(\boldsymbol{\mu}, K, \boldsymbol{\epsilon})$ where appropriate. That is, the MSRO can organize its operational infrastructure to induce $\delta > \bar{\delta}$ in the recipient-driven allocation setting as long as there exist a set of $(\boldsymbol{\mu}, K, \boldsymbol{\epsilon})$ such that $\delta > \bar{\delta}(\boldsymbol{\mu}, K, \boldsymbol{\epsilon})$. In other words, the existence of the Pareto-efficient (W, W) equilibrium depends on inventory levels $\boldsymbol{\mu}$, container capacity K , and the arrival stream of items into the MSRO inventory (characterized by the distribution of $\boldsymbol{\epsilon}$). Accordingly, we next investigate whether and how an MSRO can use these operational levers to increase its welfare provision through enabling the waiting equilibrium. To do so, we first expand and rewrite Equation (3), i.e., the value of waiting. We present this equation for recipient 1; the case for recipient 2 can be written in a similar fashion:

$$\Pi_1^{W,W} - \Pi_1^{O,W} = \delta p E[\pi_1^1] - \pi_1^0 = -(1 - p\delta)(K v_{1,2} + \mu_1 \Delta_1) + \delta p \Delta_1 E[\min(K - \mu_1, \epsilon_1)]. \quad (5)$$

Note that the first term on the right-hand side of the equality in Equation (5) is the value lost by delaying the optimal container order at $t = 0$. The second term on the other hand is the value gained by waiting for the arrival of the preferred items at $t = 1$. If the value gain is higher than the value loss for both recipients, the (W, W) equilibrium exists. Let $G(\boldsymbol{\epsilon}) \doteq \{(K, \boldsymbol{\mu}) : \delta p_i E[\pi_i^1] > \pi_i^0, \forall i\}$ be the region where the (W, W) equilibrium exists.

⁴ See Skyrms (2004) for a detailed discussion regarding coordination in the stag-hunt game.

PROPOSITION 2. $G(\epsilon) \neq \emptyset$ is characterized by (i) $K \in (\hat{K}(\mu), \bar{K}(\mu))$ with respect to capacity K , and (ii) $\mu \in (0, \mathbf{I}(K))$ with respect to inventory μ . Furthermore, if $\epsilon \leq_{st} \epsilon'$, then $G(\epsilon) \subseteq G(\epsilon')$.

Proposition 2 suggests that the recipient-driven allocation model R can only achieve Pareto-efficiency if the container capacity is neither too high nor too low, the available inventory of either item category is not too high, and the supply arrivals for either item category is sufficiently high. A sketch of the proof for the capacity condition in Proposition 2 can be provided by examining Equation (5). Note that the first term on the right hand side of this equation, $-(1 - p_1\delta)(Kv_{1,2} + \mu_1\Delta_1)$ is decreasing in K , and the second term $\delta p_1\Delta_1(E[\min(K - \mu_1, \epsilon_1)])$ is concave increasing in K . Hence, the summation of the two terms results in an inverse U-shaped function of K , unless the slope of the first term is too large, in which case $G(\epsilon) = \emptyset$. Otherwise, it is straightforward to show that (W, W) is an equilibrium in the range of capacity where the inverse U-shaped function takes positive values for both recipients. In turn, this analysis suggests that the container capacity should be between a lower and upper bound such that both recipients can be induced to wait. The intuition behind the capacity condition is that the value loss from delaying the shipment of non-preferred items would be prohibitively high for larger capacity levels (i.e., the term $-(1 - p_1\delta)Kv_{1,2}$ is much larger in absolute terms if K is high), while waiting brings little value if the container capacity is small, because the value increase through the arrival of the incoming preferred items is limited by the capacity itself. Accordingly, waiting can be induced under the R model only if the capacity level is within a certain range.

The inventory condition in Proposition 2 can also be explained by Equation (5). The first term on the right-hand side of this equation is decreasing in the inventory level μ componentwise, and the second term $\delta p_1\Delta_1(E[\min(K - \mu_1, \epsilon_1)])$ is a concave decreasing function of μ componentwise. Hence, it is straightforward to show that the value of waiting is strictly decreasing in μ componentwise. Consequently, if the inventory levels visible to the recipients are below certain thresholds, then both recipients will wait. This observation suggests that one way for an MSRO to induce recipient waiting could be to strategically censor the inventory availability information provided to the recipients. That is, inventory information censoring, in the form of presenting lower than actual inventories in the online database, could help induce recipient waiting.

The effect of the arrival distribution presented in Proposition 2 follows a stochastic ordering argument, which suggests that a “larger” random arrival is more likely to induce waiting. The intuition behind this result is conceptually the opposite of the inventory condition discussed above. The second term in Equation (5) is larger when the arrival distribution is stochastically larger, i.e., if the chance of larger arrivals into the MSRO inventory is higher, the value of waiting will be higher and the recipients will be more likely to wait. The managerial implication of this result is

that increasing the number of supply partners (which would increase the arrival volume but also make it more variable) may have an additional benefit for an MSRO beyond the items' inherent value: increasing the value of waiting for the recipients and resulting in higher welfare provision by inducing the waiting equilibrium.

PROPOSITION 3. *Let $\boldsymbol{\mu} \leq \boldsymbol{\mu}'$, and $\hat{K}'(\boldsymbol{\mu})$, $\bar{K}'(\boldsymbol{\mu})$ and $\mathbf{I}'(K)$ be the bounds defining $G(\boldsymbol{\epsilon}') \neq \emptyset$ where $\boldsymbol{\epsilon}' \geq_{st} \boldsymbol{\epsilon}$. Then $(\hat{K}(\boldsymbol{\mu}'), \bar{K}(\boldsymbol{\mu}')) \subseteq (\hat{K}(\boldsymbol{\mu}), \bar{K}(\boldsymbol{\mu}))$ and $(\hat{K}(\boldsymbol{\mu}), \bar{K}(\boldsymbol{\mu})) \subseteq (\hat{K}'(\boldsymbol{\mu}), \bar{K}'(\boldsymbol{\mu}))$. $\mathbf{I}(K)$ is component-wise concave in K and $\mathbf{I}'(K) \geq \mathbf{I}(K)$.*

Proposition 3 shows that the feasible capacity range to induce waiting by both recipients is larger if the inventory levels are lower and the arrival distribution is stochastically larger. The inventory levels that will induce waiting on the other hand are higher with a stochastically larger arrival distribution, but non-monotonic in the capacity level.

In sum, this analysis implies that an aligned combination of the three operational levers (lower inventory, higher arrival rates and intermediate ranges of dedicated container capacity) can avoid premature ordering and improve value provision in the recipient-driven model. Of course, the benefits of these choices would have to be weighed against their opportunity cost (e.g. transportation cost increase if shipping in smaller containers, partial value loss from not providing full inventory visibility, the cost of the effort of working with more partners). It is also important to note that inducing waiting may not even be possible under the R model, as the existence of $\hat{K}(\boldsymbol{\mu})$ and $\bar{K}(\boldsymbol{\mu})$ within feasible capacity bounds and a non-negative $\mathbf{I}(K)$ is not guaranteed. To shed light on why this is the case, Proposition 4 provides a simple sufficient condition.

PROPOSITION 4. *(W, W) cannot be induced if $\frac{v_{i,-i}}{v_{i,i}} \geq p_i \delta$ for any i .*

The condition in this proposition suggests that if the value of the non-preferred item for either recipient is close to the same for the preferred item, waiting cannot be induced. The intuition behind this result follows from the value of waiting characterization in Equation (5). In particular, note that the absolute magnitude of the first term on the right-hand side of the equality will be higher if the value of the non-preferred item is sufficiently high, implying that the value of waiting can be negative. In sum, Proposition 4 suggests that the recipient-driven allocation model, due to the inherent competition it induces between recipients, cannot achieve the welfare improving equilibrium if the recipient valuations of different product categories are too close to each other, irrespective of the capacity, inventory or arrival rate levels observed by the MSRO.

This observation implies that the R model has limited potential to maximize an MSRO's value provision due to the inherent recipient competition it induces. Accordingly, under the circumstances that prevent a welfare-improving waiting equilibrium under the R model, an MSRO that aims to

maximize its value provision may need to eliminate recipient competition. We posit that there are two ways to achieve this: One is to provide inventory access to one recipient at a time. Alternatively, the MSRO could go beyond the boundaries of the R model and reverse the flow of information, i.e., as opposed to providing inventory information to the recipients, the MSRO could ask recipients to provide information regarding their needs, and make container decisions on their behalf (i.e., switch to a push-based operational model). We investigate these two options in the next section.

4. Eliminating Recipient Competition

4.1. Single-Recipient Access (SR Model)

Consider the first option discussed at the end of §3, i.e., providing inventory access to one recipient at a time, which is the simplest way to eliminate recipient competition under the R model. In this case, it is straightforward to show that any recipient's payoff will be at least as high as her equilibrium payoff under the R Model (for any K , μ , ϵ and δ), if the recipient has exclusive access to the MSRO inventory. This argument is an immediate consequence of Proposition 1, which states that the discount rate threshold above which any chosen recipient prefers to wait is lower than the same under recipient competition. This proposition, along with the results in Proposition 2 also directly implies that the ranges of capacity, inventory, and arrival volume for which the recipient waits are larger under single-recipient access.

Now, assume that the MSRO selects the single recipient to provide inventory access to according to the same tie-breaking rule as mentioned in §3, consistent with the belief p of the recipients. We refer to this approach as the SR model. In this case, defining the chosen recipient i 's value provision with exclusive MSRO inventory access as $\Pi_i^{SR} \doteq \max(\pi_i^0, \delta E[\pi_i^1])$, the total expected welfare provided by the MSRO under the single-recipient benchmark can be written as $W_{SR} = p\Pi_1^{SR} + (1-p)\Pi_2^{SR}$. Defining W_R as the equilibrium total expected welfare under the R model, the following result formally states that the SR Model will increase the MSRO's expected welfare provision compared to the same under the R model.

PROPOSITION 5. $W_{SR} \geq W_R$.

The practical implication of Proposition 5 is that an MSRO operating under the R model can benefit from reducing the number of recipients that have simultaneous access to its inventory database. Nevertheless, it is important to note that reducing the number of recipients with access to MSRO inventory may also imply fewer recipients being served over a fixed period of time, which is not captured by our stylized two-period model. There may also be organizational reasons beyond the scope of our analysis (such as more donations, or better resource/personnel use of the MSRO) that favor simultaneous multi-recipient access as well. In that case, choosing the R model in conjunction with the operational mechanisms discussed above will help improve its welfare provision.

4.2. Provider-Driven Resource Allocation (P) Model

We now turn our attention to a scenario where the MSRO changes the direction of the information flow, i.e., the MSRO no longer provides inventory information to the recipients, but acquires necessary information regarding their needs, i.e., $v_{i,j}, i, j = 1, 2$. Accordingly, at $t = 0$, given $v_{i,j}, i, j = 1, 2$, the MSRO has two options in our model: (i) place an order (and simultaneously select the recipient with the higher payoff) or (ii) wait until the arrival of the next batch of supplies, which increases the available inventory by ϵ . If the MSRO chooses to order at $t = 0$, it determines the volume of items $y_{i,j}$ that maximizes the container value for each recipient i and selects the recipient with a higher container value according to:

$$\begin{aligned} \hat{\pi}^0 &\doteq \max_{i=1,2} \max_{y_{i,1}, y_{i,2}} v_{i,1}y_{i,1} + v_{i,2}y_{i,2} \\ \text{s.t. } &y_{i,1} \leq \mu_1, y_{i,2} \leq \mu_2, y_{i,1} + y_{i,2} \leq K. \end{aligned} \quad (6)$$

As before, if the MSRO waits, an additional volume $\epsilon = (\epsilon_1, \epsilon_2)$ of supply is received. Then, the MSRO determines the volume of items $y_{i,j}$ that maximizes the container value for each recipient i and selects the recipient with a higher container value according to:

$$\begin{aligned} \hat{\pi}^1 &\doteq \max_{i=1,2} \max_{y_{i,1}, y_{i,2}} v_{i,1}y_{i,1} + v_{i,2}y_{i,2} \\ \text{s.t. } &y_{i,1} \leq \mu_1 + \epsilon_1, y_{i,2} \leq \mu_2 + \epsilon_2, y_{i,1} + y_{i,2} \leq K. \end{aligned} \quad (7)$$

Then, the total expected discounted welfare provision under the P model can be written as $W_P = \max(\hat{\pi}^0, \delta E[\hat{\pi}^1])$. While we omit the technical details that characterize the optimal decisions under the P model for brevity (results available from the authors), we next present a result that formally establishes the dominance of the P model over the R model, which proves the value of information flow reversal for MSROs.

PROPOSITION 6. $W_P \geq W_{SR}$.

The intuition behind Proposition 6 is that while the SR model (which is readily better than the R model) selects a recipient at $t = 0$ according to the tie-breaking rule that is independent of inventory information and allows the recipient to choose when to order, the P model determines which recipient to order for and what to order based on inventory information. As such, the P model dominates the SR model, which follows from Jensen's inequality. Given that welfare provision under the SR approach dominates the same under the R model, the P model provides the highest value to the recipient base. In sum, this analysis suggests that information flow from the recipients to the MSRO combined with MSRO-driven resource allocation is the ideal scenario, provided that recipient needs assessment is feasible. In particular, the benefit of the P model stems from the fact

that it eliminates the potential for premature orders induced by recipient competition under the pull-based R model.

A natural question at this point is whether gathering recipient needs information is feasible for the MSRO to effectively exercise the P model. The WHO finding regarding the high level of mismatched donations suggests this is not easy, but neither is it infeasible. For example, Beta tries to evaluate recipient needs before a recipient qualifies for service so they have a basic understanding of the categories of value to each recipient. They also have biomedical engineers that periodically visit some recipients to help them install/maintain donated equipment and help validate the wish lists provided by the recipients. Face The Challenge (FTC) is an organization highly regarded as an ideal steward of provider-driven resource allocation. To identify the items that best meet their recipients' medical needs, FTC periodically sends doctors to visit recipients. Nevertheless, acquiring such information can be costly, as it requires having personnel on the ground, and the cost will increase with the size of the recipient base. Therefore, the information acquisition capability may vary between MSROs depending on their resources and the size of their recipient base. In sum, while reversing the information flow can be costly (especially if recipient needs evolve dynamically over time and the recipient base is large and geographically dispersed as for Beta), our analysis provides a basis for engaging in conversations with donors around the value of a provider-driven allocation model to include such information gathering.

5. Numerical Analysis

In this section, we explore the effectiveness of the R model and assess the welfare improvement potential of the mechanisms we propose by using medical surplus recovery data collected from Beta. This data allows us to estimate model parameters (e.g., supply arrival rate, recipient-specific item valuations, inventory levels and capacity), which we then use to investigate the potential value of the operational mechanisms we propose in §3-4. The data is also useful for comparing the recipients' ordering behavior with that of Beta: It was collected starting from January 2013 when Beta started assisting recipients by filling the containers on the recipients' behalf. This initiative was based on Beta's observation that recipient-driven orders were often rushed, with recipients frequently placing seemingly useless items in their containers (Beta 2014). Accordingly, unlike the original operating regime where recipients had full control over what was to be placed in their dedicated containers, Beta terminated recipient access to inventory before the container was completely filled by the recipient (by having the inventory access passwords expire), and filled the rest of the container on the recipient's behalf. This new approach allows us to compare Beta and the recipients' ordering behaviors at a container level.

We describe the data in more detail below. After that, we compare Beta and recipients with respect to their container fill behaviors and carry out an impact assessment of the proposed improvements developed in §3-4.

5.1. Data

The data set contains shipment and packing list data from 36 containers shipped by Beta between January 2013 and October 2014. During this period, Beta tracked information regarding 14,900 distinct items in its inventory database, spanning sixty-seven product categories. The shipment data consists of details such as donor values of products (which represent the market value of those products), weights of the selected products, quantities of each product type for each container shipment, recipient names, transportation modes, container capacities, and shipment dates. The data set was downloaded and provided by Beta in December 2014 for our analysis.

The packing list data consists of information regarding boxes of unique product categories, which we consider as our unit of analysis, used to fill a container. In particular, this data contains three date-related fields where every box goes through the following sequence: “Created”, “AddedToOrderDate”, and “ConfirmedDate”. The “Created” time-stamp refers to the time when a box became available in Beta’s inventory. At this time, recipients that have access to the inventory database (or Beta) can place an order to place this box in the recipient’s container. Once a recipient or Beta orders the box for the container, the ‘AddedToOrderDate’ time-stamp is generated. Finally, the “ConfirmedDate” time-stamp represents the date at which the container is filled up and the boxes in the packing list are secured for a shipment to the specific recipient. In addition to those, we also have information as to each recipient’s inventory database access permission issue and end dates, and information regarding recipient characteristics. This data set allows us to consolidate four pieces of key information that we utilize in our empirical investigation:

Arrival Rates: The first piece of information regards the arrival rate of products in each product category. We determine the arrival rate (supply volume per day-between-arrivals) of each box type using the “Created” time-stamps, and denote the arrival rate of a product category as the total arrival rate for all box types in that category.

Inventory Levels: The second piece of information regards the inventory availability of each product category. We determine the inventory levels by comparing the inflow of boxes through the “Created” time-stamps, and outflow of the same by the “ConfirmedDate” time-stamp, along with associated volumes. In our data, the earliest record of a box created comes from 2006 that was picked by some recipient three years later. Therefore, the initial inventory levels used in our analysis (as of January 2013) are retrospectively determined by supplies that have arrived since 2006 but that have not yet been ordered by any recipient.

Valuations: The third piece of information is the value of each product category in Beta’s inventory for recipients. A natural measure for this value is the donor value, i.e., the donor’s estimate of the market value of a donated product. That is, one may expect higher market value items to be picked more frequently by the recipients. An examination of the recipient orders reveals that recipients do not necessarily place the highest market value items in their containers, which is a clear indication that recipient needs may favor ordering lower market value items as well. To account for recipient-specific, needs-based valuations of different product categories, we apply a recipient-choice based adjustment to donor values, which provides an estimated inherent value for each recipient.

Our adjustment procedure ensures that the recipients’ actual ordering behavior in the data matches the solution to the optimization problem in §3, and proceeds as follows: First, we analyze the Beta inventory before and after a recipient’s order process is complete, and observe box types whose inventories were depleted via the recipient’s order. We consider these as “critical” for the recipient. The remaining box types in Beta’s inventory (with positive inventory level after the recipient order), are then called “non-critical” for this recipient. We calculate the average donor value of non-critical boxes that were picked by the recipient, and denote it as v_{non} . If the donor value of a critical box is above v_{non} , we set the recipient valuation of that box to be its donor value. Otherwise, we set the recipient valuation of that box to be slightly above (5 cents per lb.) v_{non} . For a non-critical box that was picked by the recipient, we set the recipient valuation to be equal to v_{non} . For a non-critical box that was not picked by the recipient, if the donor value is below v_{non} , we set the recipient valuation to be its donor value. Otherwise, we set the recipient valuation to be equal to v_{non} .⁵ Once the valuations are obtained for every product category for each recipient, we use them to estimate the value of every container filled by either Beta or its recipient. While this approach is not unique, it is one that keeps recipient-specific values as close as possible to the donor values (at most 2.91% more at a container level), while allowing our models to reproduce the specific container mix selected by the recipients.

Order Placement Speed: The fourth piece of information regards the speed at which boxes were placed by the recipients and Beta. For each recipient order, we obtain the recipient’s inventory

⁵ Consider an example where box types 1-5 are available, with donor values 25¢, 20¢, 40¢, 50¢, and 20¢/lb. For simplicity, assume each box weighs 1lb. Assume recipient A orders boxes 2-5 and only boxes 4 and 5 are critical, i.e., their inventory is depleted via recipient A’s order. In this case, the average value of non-critical boxes picked by the recipient (i.e., boxes 2 and 3) is $v_{non} = (20 + 40)/2 = 30¢$. Box type 4 has a donor value 50¢ > 30¢. As such, we set recipient A’s valuation for box type 4 to its donor value 50¢. For box type 5, however, this is not the case. Because 20¢ < 30¢, we set recipient A’s valuation for box type 5 to be $30 + 5 = 35¢$. Second, we set A’s valuation for non-critical boxes picked by the recipient (i.e., boxes 2 and 3) to be equal to $v_{non} = 30¢$. Finally, for box type 1, since the donor value 25¢ < 30¢, we set recipient A’s valuation for box type 1 to be its donor value 25¢. Therefore, recipient A’s valuations for boxes 1-5 are adjusted to 25¢, 30¢, 30¢, 50¢, and 35¢/lb, respectively. Under this set of valuations, depleting boxes 4 and 5 and filling the rest of the container with boxes 2 and 3, which matches recipient A’s actual ordering behavior, is an optimal solution to recipient A’s value maximization problem.

database access permission issue and end dates, which give us the recipient’s length of involvement, i.e., the total time it took a recipient to place the chosen set of boxes in its container. We then divide this number by the number of boxes placed by the recipient to calculate the rate at which a recipient ordered boxes. We estimate Beta’s length of involvement, i.e., the total time it took for Beta to place the chosen set of boxes in the recipient’s container, by calculating the time between “AddedToOrderDate” values for the first and last boxes that were placed in a container by Beta. Dividing this number by the number of boxes placed by Beta in that container gives us the rate at which Beta placed boxes in containers.

5.2. Comparing Beta and Recipient Orders

Among the 36 containers filled during the period of analysis, 34 were filled jointly by Beta and recipients, while two containers were filled only by Beta. We focus only on the jointly-filled shipments, which allows us to compare Beta’s and recipients’ ordering behaviors matched at a container level. Furthermore, we note that all these shipments were subject to simultaneous inventory access by exactly two recipients. Table 3 provides the summary statistics for these containers.

Table 3 Summary statistics for the 34 shipments that were jointly filled by recipients and Beta, under simultaneous two recipient access. Standard deviations are provided in parentheses.

Statistics	Beta	Recipient
# Observations	34	34
% Boxes filled	17	83
Total time to order, days	22.76 (18.33)	15.50 (14.78)
Average time to order (per box), days	0.32 (0.31)	0.025 (0.03)
Average \$ value per box	262.56 (176.93)	136.70 (33.37)

Using a paired t-test, we find that the rate at which recipients placed boxes in containers is significantly faster than Beta (the average time to order a box is almost 11.8 times more for Beta, $t(33) = 4.84, p < 0.001$).⁶ We also find that the average value of boxes placed by Beta is significantly (92%) higher than that of a recipient ($t(33) = 4.15, p < 0.001$). In other words, there is a significant difference between the container fill behavior of Beta and the recipients. Beta waits longer and places boxes with higher estimated recipient value in recipient containers than the recipients themselves. A natural question is: To what extent is the higher value in boxes placed by Beta driven by the extended period of time taken to fill the container? We find that 95.30% of the boxes placed by Beta were placed after the recipients’ inventory database access period, and

⁶ Note also that even the total time of involvement of Beta is significantly higher than that of the recipients, even though Beta places significantly fewer boxes in containers: The total time of involvement for Beta (at a container level) is on the average 7.3 days longer ($t(33) = 1.82, p < 0.05$) than the same for recipients.

among those boxes 29.79% were of types that had zero inventory after recipient orders, and arrived after the recipient’s inventory access period. This suggests that an important fraction of the value added from waiting appears to be driven by the longer wait that allowed Beta to leverage new arrivals into its inventory. Furthermore, the average estimated value of those boxes was \$226.13, significantly higher than that of an average box placed by the recipients (\$136.70).

Collectively, these observations suggest that recipients order faster than Beta and do not wait for the arrival of potentially more valuable items in the near future. Combined with Beta officials’ statements regarding the inefficiencies in the recipient orders, these observations give us reason to believe that reducing or eliminating the pressure to place premature orders will create higher value provision by the MSRO, which we investigate next. These observations also point to the value of an in-depth empirical investigation of drivers of recipient ordering behavior based on a more comprehensive data set, which is beyond the scope of this study.

5.3. Impact Analysis

We next perform a numerical analysis based on Beta data to quantify the value of the operational improvements we proposed in §3 and 4. In particular, we first analyze the potential of the operational improvements to the R model in inducing recipient waiting. We then quantify the welfare provision potential of the SR and P models.

5.3.1. Inducing the Waiting Equilibrium: Improvements to the R Model In this section, we consider all 34 containers that were shipped during the period of interest. For each shipment, we follow the model set-up used in §3 and consider a two-period analysis, where the period length is half a month (based on the average recipient shipment period length of 15.5 days in Table 1). We investigate whether the capacity adjustment, inventory visibility and arrival rate change levers could be used to drive the recipient to wait longer if the traditional recipient-driven model were in place. We let $\delta = 0.99$,⁷ and assume egalitarian tie-breaking (i.e., $p = 0.5$).

To perform this analysis, we first estimate the arrival volume for each product category using the historical arrival data on hand as described in §5.1, and assume that the arrivals are i.i.d. in every period. The inventory level for each product category μ_j at the beginning of January 2013 is also calculated using the arrival and shipment dates of each product category. Each recipient i ’s valuation of each product category j ($v_{i,j}$) is calculated using the procedure described in §5.1. We consider the two capacity levels available to Beta in practice, 40- and 20-ft containers, which correspond to approximately 10000 and 5000 lb. in a container, respectively.

⁷ This corresponds to 23% compounded annual discount. Additional analysis available from the authors shows that all results presented in this section directionally hold under two alternative discount levels with $\delta = 0.95$ and $\delta = 0.999$. These values were chosen to represent a range of δ values, as δ is difficult to estimate without having direct access to the recipients’ particular conditions at the time of allocation.

Before we proceed with a detailed analysis, we note that Beta’s realized container shipments are not paired with perfect overlap as our model in §3 assumes. That is, Beta typically ships containers to recipients one by one; after one recipient leaves, a new recipient is given access into the system. In this way, although Beta consistently maintains two-recipient access throughout the observed periods, the password issue and end dates differ between recipients. Accordingly, an equilibrium analysis of the sort we perform in §3 is difficult to exercise, and we turn our attention to a proxy measure, the value of waiting, to analyze the R model’s potential for inducing recipients to wait. In particular, we calculate the value of waiting as per Equation (5) for all recipients in our data set as a proxy for their likelihood of waiting.

With this data, we first consider the value of waiting under a hypothetical exclusive single-recipient scenario (as in Equation (4)), which allows us to predict that 73.5% (i.e., 25 of 34) of the recipients would have preferred to wait for another period to complete their container orders (i.e., the value of waiting would be positive), had they been given exclusive inventory access. This observation is in line with the results presented in §5.2, and suggests that there is room for improvement through the operational mechanisms introduced in §3.

Accordingly, we next look for the potential of the operational mechanisms in inducing recipient waiting. To do so, we first define a status-quo benchmark under the R model, which considers the actual container shipments that realized, and calculates the value of waiting using the actual shipment data, i.e., 40-ft containers and the actual inventory levels in Beta’s system at the time the container fill process is completed. At this benchmark, the value of waiting averaged across all realized shipments is given as \$-54,512. To assess the change in the value of waiting when the operational mechanisms we propose are exercised (i.e., different container capacity, inventory level and arrival rates), we use the theoretically optimal container contents based on the model formulation in §3, assuming the same inventory availability as when the container mix was selected under the status-quo benchmark. A comparison between the status-quo benchmark and these scenarios allows us to quantify the potential of these operational mechanisms in inducing recipient waiting as follows.

We first look at the effect of capacity reduction on the average value of waiting across all recipients by assuming a 20-ft container, which is the only practically viable option per Beta officials. We find that such a capacity reduction leads to a 80.3% increase in the average value of waiting (-\$10,749) across all containers. The value of waiting with a 40-(20-) ft container is positive only for 11.8 (29.4) percent of the recipients, i.e., the number of recipients with positive value of waiting more than doubles. We next investigate the value of inventory level adjustment and identify the range of inventory level reduction that can induce waiting, which is illustrated in Figure 1. An immediate observation from Figure 1 is that, the average value of waiting is positive at 38% visibility and

Figure 1 The effect of inventory adjustment at 32 – 42% of the benchmark inventory on the value of waiting (in \$1000s).

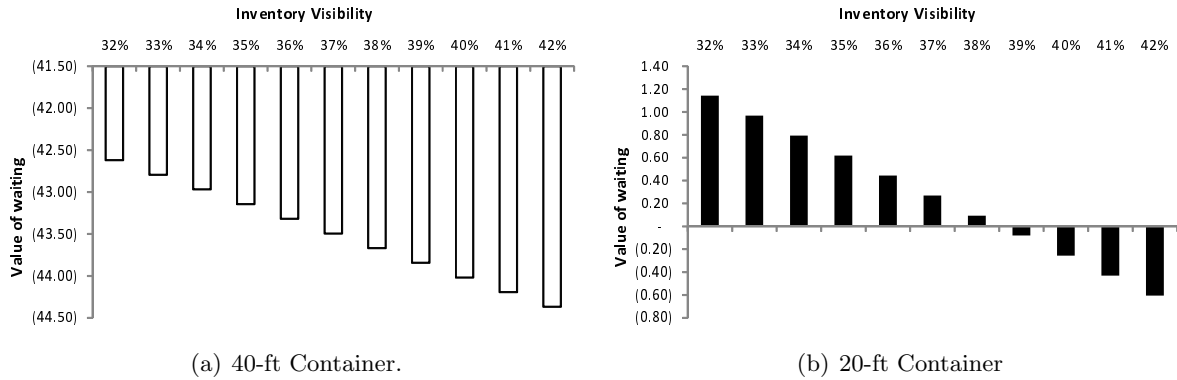
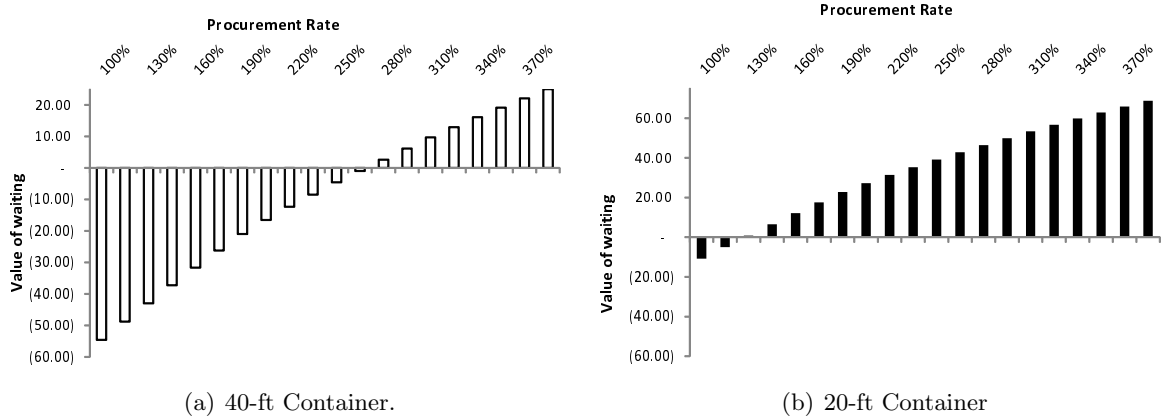


Figure 2 The effect of increasing the procurement rate on the value of waiting (in \$1000s).



below if a 20-ft container is utilized. In addition, at 38% of the benchmark inventory level, the value of waiting with a 20-ft container is positive for 32.5 percent of the recipients. Next, we consider the effect of increasing the procurement volume of products by Beta, i.e., we analyze whether further effort by Beta to increase the inflow of goods by signing up more hospitals or corporate donors to provide more medical surplus is an effective approach. Figure 2 shows that with a 20-ft container, a positive average value of waiting can be obtained by increasing the procurement rate by 30%. In addition, with a 30% increase in the procurement rate, the value of waiting is positive for 32.5 percent of the recipients. Consequently, the data collected from Beta suggests that the value of the simple operational improvements we identified in §3 appear to be capable of affecting the recipients' ordering behavior. Essentially, by using a combination of capacity adjustment, visible inventory level decrease and procurement rate increase, Beta could decrease the competitive ordering pressure on the recipients.

We also note that these operational improvements not only increase the value of waiting but also change the average value in a container, which we define as the recipient value of a container

in our model per lb. of capacity (i.e., π_i^0/K) given the actual inventory levels at the time the actual container fill process is completed. Similarly to the value of waiting calculation above, we first calculate this measure at the status-quo benchmark for each container that shipped using the actual shipment data (i.e., 40-ft container and the actual inventory levels). We note that this value coincides with that in our formulation of the R model described in §3, because of the item valuation estimation described in §5.1. We then calculate the same under the operational improvement scenarios discussed above to observe the following: Capacity reduction leads to an increased average value in a container, which is intuitive given the decreasing returns to scale due to limited inventory availability of higher value items. In particular, a switch from the benchmark with a 40-ft container to a 20-ft container increases the average value in the container by 32.3% (from \$12.20/lb to \$16.14/lb). However, a reduction in the level of visible inventory reduces the value provided by higher-value items, the lack of which leads to average value loss in the container. In particular, the average value in a 40-ft container goes down to \$10.4/lb at 38% of the current inventory level, where the average value of waiting switches to a positive number.

5.3.2. The Value of Eliminating Recipient Competition We next estimate the effect of eliminating recipient competition through the SR and P models on the MSROs value provision capability. To do so, we adopt heuristic approaches (described below) that extend our two-period two recipient models to multi-period multi-recipient analyses.

We start with the SR model heuristic, in which the MSRO provides inventory access to a single recipient at any point in time by rescheduling the 34 shipments in consideration using the same data as above over 42 ordering periods (i.e., 21 months). In this case, we assume that the MSRO randomly determines which recipient to provide inventory access to, where the chosen recipient is granted a 2-period access to the MSRO inventory and may choose to order at one of the two ordering periods. Once it orders, inventory is reduced accordingly. In every period over the 42 periods, inventory is also increased by the actual arrival volume in that period. We continue this procedure on a rolling horizon basis, and if all 34 shipments have not been completed over the 42 periods for which we have data, we continue the procedure based on inventory updates with simulated supply arrival volumes until all recipients have been scheduled. Comparing the net present average container value under this heuristic allocation with the same under the status-quo R model benchmark (i.e., the actual container shipments) we find an increase of 168% in the net present average container value (from \$12.20/lb to \$32.73/lb) under the SR model.

Next, we evaluate the value provision improvement under a P model heuristic. To do so, we consider all 34 recipients in our data set and we derive an improved value provision (i.e., increased average container value across all recipients) by having the MSRO reschedule the 34 shipments in

consideration using the same data as above. As before, we use a heuristic approach on a two-period rolling horizon basis that determines the provider-driven resource allocation in every period per the model described in §4. We first use the available inventory and arrival data to compute the optimal container order composition for each recipient at time t , starting from $t = 1$, and identify the recipient that receives the highest payoff from ordering immediately. Next, we estimate the expected discounted payoff from waiting for another period (i.e., 15.5 days) for each recipient. We then compare the highest payoff from ordering immediately and the highest expected discounted payoff from ordering in the next period. If the former exceeds the latter, we ship the corresponding recipient's container with its optimal container mix, update the inventory level of each product category accordingly, and remove the particular recipient from the database. Otherwise, no order is placed in that period, the inventory level is updated based on the realized supply arrival volume in period t , and we move to period $t + 1$. This algorithmic allocation procedure, which utilizes the exact historical inventory and product arrival data from Beta, shows that the net present average container value under this heuristic allocation is 265% higher than the same under the R model (from \$12.20/lb to \$44.59/lb). From Beta's perspective, this result clearly indicates the importance of recipient needs assessment. It can also be helpful in comparing the value provision potential of donor dollars spent towards container shipments versus learning about recipient needs, should Beta consider implementing the provider-driven allocation approach.

6. Conclusion

According to the World Health Organization (WHO), the majority of donated equipment was inappropriate for the beneficiaries in the health sector (Howitt et al. 2012), and only eight percent of recipients were estimated to work with MSROs in ensuring quality and sufficient knowledge transfer to realize the full potential of the donated equipment (Compton 2012a). Hence, the need for operations management research to improve the supply-demand match in this domain is clear (Kotsi et al. 2014), and our objective in this paper is to provide operational guidelines to improve resource allocation practices of MSROs.

Building on the recipient-driven resource allocation model of an award-winning MSRO, our analysis provides the following insights for MSROs: When facing an impatient recipient base with urgent needs, a recipient-driven model appears to be ideal. Otherwise, i.e., if recipients can afford to wait for better supplies, the following operational improvement path should be considered: (i) the supply-demand match for an MSRO can be significantly improved by a provider-driven resource allocation model that relies on understanding the recipient needs. That is, a reversal in the direction of information flow in the MSRO supply chain, which eliminates inventory information provision to recipients (that the recipient-driven model relies on) and utilizes recipient needs information in the

MSRO-driven container fill processes can help substantially improve the MSRO's value provision. However, given that MSROs operate in cash-constrained environments that rely on donations and the fact that keeping up to date with detailed and potentially time-varying recipient needs can be costly, this approach may require a strategic assessment by MSROs. In particular, it may require engaging in conversations with donors around the value of investing in such information gathering to manage donor expectations. (ii) If for whatever reason an MSRO is bound to remain in a recipient-driven setting (e.g., recipient needs assessment is prohibitively expensive or the donors require recipient-driven allocations), recipient competition for products in the MSRO inventory may lead recipients to prematurely order a potentially suboptimal mix of supplies. If this is the case, the first possible action an MSRO can take is to reduce the number of recipients that have simultaneous inventory access to minimize the impact of recipient competition. However, this may reduce the number of recipients that an MSRO can serve over a fixed period of time, which implies that an MSRO that secures donations for many recipients at a time may be forced to provide simultaneous multiple-recipient access. In that case, our analysis suggests that the potential loss of effectiveness driven by competition can be partially avoided by a combination of operational mechanisms such as reducing the visibility of available inventory, determining the appropriate container capacity allocated to recipients, and increasing the arrival volume of supplies. We believe the insights from our analysis can be valuable for MSROs operating in a development context specifically, where the operational characteristics involve uncontrollable, time sensitive (but non-emergency) and constrained multi-product supplies used to meet the uncertain needs of a large recipient base.

The analysis in this paper can be extended in two important directions. First, many of the insights derived from our stylized model can be extended to multi-period models with multiple recipients. A key difference of the multi-period case from our stylized model with two periods and a single container is that under the multi-period case, recipients who do not receive the first container have the opportunity to receive a container in future periods; their payoff is not zero. We verified that the main insights continue to hold in an extension to the two-period model where a second container can be shipped to the second recipient in the second period, mimicking the positive future payoff implicit in a multi-period setting. In other words, even with the possibility of receiving a container in the future, the premature ordering behavior induced by competition still exists, and the operational levers we propose, such as adjusting the inventory visibility, container size and the arrival rates of products would still help induce waiting. However, a full analysis of the multi-period case for the centralized problem (i.e., the provider-driven model) is intractable since the decisions involve the sequencing of recipients under uncertain arrivals of multiple product categories. Even the deterministic version of this problem with two product categories is at least

as challenging as the Hamiltonian path problem. Analytically characterizing the welfare loss of the recipient-driven model and the improvement of the proposed operational mechanisms for the multi-period case is an interesting topic that we leave for future research.

Second, under circumstances where recipient needs assessment in the field is prohibitively expensive, in order to implement a provider-driven model, one can investigate whether truth-inducing mechanisms can be used to elicit the valuations or preferences of each recipient for different product categories. It would be interesting to investigate whether a truth-inducing mechanism can achieve the first best value provision (i.e., global optimal assuming the valuations of recipients are known), and if not, to what extent a truth-inducing mechanism can improve the value provision over the recipient-driven model, and what the welfare loss caused by the information gap would be.

References

- Akan, M., O. Alagoz, B. Ata, F. Erenay, and A. Said. 2012. A broader view of designing the liver allocation system. *Operations Research*, 60 (4), 757-770.
- Altay, N. and W.G. Green. 2006. OR/MS research in disaster operations management. *European Journal of Operational Research*, 175 (1), 475-493.
- Balcik, B., B. Beamon and K. Smilowitz. 2008. Last mile distribution in humanitarian relief. *Journal of Intelligent Transportation Systems*, 12 (2), 51-63.
- Berry, A. The Power of One, MedShare 2014 Annual Report.
- Bertsimas, D, V.F. Farias and N. Trichakis. 2013. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61 (1), 73-87.
- Besiou, M., Pedraza-Martinez, A.J., and Van Wassenhove, L.N. 2014. Vehicle supply chains in humanitarian operations: decentralization, operational mix, and earmarked funding. *Production and Operations Management*, 23 (11), 1950-1965.
- Beta. 2012. Personal communications with Chief Operating Officer of Beta.
- Beta. 2014. Personal communications with Shipment Manager of Beta.
- Bhattacharya, S., S. Hasija and L.N. Van Wassenhove. 2014. Designing efficient infrastructural investment and asset transfer mechanisms in humanitarian supply chains. *Production and Operations Management*, 23 (9), 1511-1521.
- Cachon, G.P. and M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Science*, 48 (6), 1032-1048.
- Cachon, G. and M. Lariviere. 1999. Capacity choice and allocation: strategic behavior and supply chain performance. *Management Science*, 45 (8), 1091-1108.
- Compton, B. 2011. First do no harm: assessing and selecting high-quality medical surplus recovery organizations. The Catholic Health Association.

-
- Compton, B. 2012a. Let's make donations more effective: don't waste these important efforts. *Health progress (Saint Louis, Mo.)*, 93 (2), 84.
- Compton, B. 2012b. Responsible redistribution of medical supplies & equipment. The Catholic Health Association.
- Daff, B.M., C. Seck, H. Belkhatat and P. Sutton. 2014. Informed push distribution of contraceptives in senegal reduces stockouts and improves quality of family planning services. *Global Health: Science and Practice*, 2 (2), 245-252.
- de la Torre, L.E., I.S. Dolinskaya and K.R. Smilowitz. 2012. Disaster relief routing: Integrating research and practice. *Socio-Economic Planning Sciences*, 46 (1), 88-97.
- Duran, S., M.A. Gutierrez and P. Keskinocak. 2011. Pre-positioning of emergency items for CARE international. *Interfaces*, 41 (3), 223-237.
- Ergun, Ö., L. Gyi, J. Heier-Stamm, P. Keskinocak and J. Swann. 2013. Improving humanitarian operations through technology-enabled collaboration. *Productions and Operations Management*, 23 (6), 1002-1014.
- Gavirneni, S., R. Kapuscinski and S. Tayur. 1999. Value of information in capacitated supply chains. *Management Science*, 45 (1), 16-24.
- Hopp, W. and M. Spearman. 2008. *Factory Physics*. McGraw-Hill, Irwin.
- Howitt, P., A. Darzi, G.Z. Yang, et al. 2012. Technologies for global health. *The Lancet*, 380 (9840), 507-535.
- Huang, M., K. Smilowitz and B. Balcik. 2012. Models for relief routing: Equity, efficiency and efficacy. *Transportation Research Part E*, 48 (1), 2-18.
- Kotsi, T., L.N. Van Wassenhove and M. Hensen. 2014. Medicine donations: Matching demand with supply in broken supply chains. *INSEAD Working Paper*.
- Kunz, N. and G. Reiner. 2012. A meta-analysis of humanitarian logistics research. 2012. *Journal of Humanitarian Logistics and Supply Chain Management*, 2 (2), 116-147.
- MacRae, J.W. 1997. The REMEDY utilization study: A model for donation evaluations. PhD Thesis. Yale University.
- Natarajan, K. V. and J.M. Swaminathan. 2014. Inventory management in humanitarian operations: Impact of amount, schedule, and uncertainty in funding. *Manufacturing & Service Operations Management*, 16 (4), 595-603.
- Ozdamar, L., E. Ekinici and B. Kucukyazici. 2004. Emergency logistics planning in natural disasters. *Annals of Operations Research*, 129 (1-4), 217-245.
- Pedraza Martinez, A., S. Hasija and L.N. Van Wassenhove. 2010. An operational mechanism design for fleet management coordination in humanitarian operations. *INSEAD Working Paper*.
- Perry, L. and R. Malkin. 2011. Effectiveness of medical equipment donations to improve health systems: How much medical equipment is broken in the developing world?. *Medical & Biological Engineering Computing*, 49 (7), 719-722.

- Rosenblatt, W.H. and D.G. Silverman. 1994. Cost-effective use of operating room supplies based on the REMEDY database of recovered unused materials. *Journal of Clinical Anaesthesia*, 6 (5), 400-404.
- Rosenblatt, W.H., A. Chavez, D. Tenney and D.G. Silverman. 1997. Assessment of the economic impact of an overage reduction program in the operating room. *Journal of Clinical Anaesthesia*, 9 (6), 478-481.
- Salmeron, M. and A. Apte. 2010. Stochastic optimization for natural disaster asset prepositioning. *Production and Operations Management*, 19 (5), 561-574.
- Shaked, M. and A. Shanthikumar. 2007. *Stochastic Orders*. Springer Series in Statistics.
- Skyrms, B. 2004. *The stag hunt and the evolution of social structure*. Cambridge University Press.
- Su, X. and S. Zenios. 2004. Patient choice in kidney allocation: The role of queueing discipline. *Manufacturing & Service Operations Management*, 6 (4), 280-301.
- Su, X. and S. Zenios. 2005. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations Research*, 53 (3), 443-455.
- Taylor, T. A. and W. Xiao. 2014. Subsidizing the distribution channel: Donor funding to improve the availability of malaria drugs. *Management Science*, 60 (10), 2461-2477.
- Van Wassenhove, L.N. 2006. Humanitarian aid logistics: Supply chain management in high gear. *Journal Operational Research Society*, 57 (5), 475-489.
- Van Wassenhove, L.N. and A.J. Pedraza Martinez. 2011. Using OR to adapt supply chain management best practices to humanitarian logistics. *International Transactions in Operational Research*, 19 (1-2), 307-322.
- Zenios, S.A., G.M. Chertow and L.M. Wein. 2000. Dynamic allocation of kidneys to candidates on the transplant waiting list. *Operations Research*, 48 (4), 549-569.

Appendix

A. Multiple-Item Model Extension

In this section, we extend the basic stylized model in §3 to a multiple-item model. Our main goal is to obtain payoff expressions for every recipient that are used as inputs for representing recipient-driven resource allocation in normal-form when more than two items are provided. To do so, we consider the following modifications. First, the available inventory consists of m items with μ_j of item j at stage 1 ($t = 0$). Second, $v_{i,j}$ denotes the value of item j for recipient i for $1 \leq j \leq m$.

Action Space and Strategy Profiles. The action space for recipient i remains as $\{O, W\}$ with the set of pure strategy profiles S consisting of four possibilities: (O, O) , (O, W) , (W, O) and (W, W) . We define the appropriate linear programs contingent on observing K , $\boldsymbol{\mu} = (\mu_1, \dots, \mu_m)$ and p . At stage 1, recipient i simultaneously chooses either to order or to wait. If recipient i orders, it simultaneously selects $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,m})$ to maximize its container value according to (8):

$$\begin{aligned} & \max_{\mathbf{y}_i \in \mathbf{R}_+^m} v_{i,1}y_{i,1} + v_{i,2}y_{i,2} + \dots + v_{i,m}y_{i,m} \\ & \text{s.t. } y_{i,1} \leq \mu_1, \dots, y_{i,m} \leq \mu_m, \sum_{j=1}^m y_{i,j} \leq K. \end{aligned} \quad (8)$$

If either one or both recipients order, the game stops. If both recipients wait, an additional volume $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_m)$ of supply is received over the period. At stage 2, both recipients order and recipient i selects $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,m})$ that maximizes its container value based on the new inventory levels:

$$\begin{aligned} & \max_{\mathbf{y}_i \in \mathbf{R}_+^m} v_{i,1}y_{i,1} + v_{i,2}y_{i,2} + \dots + v_{i,m}y_{i,m} \\ & \text{s.t. } y_{i,1} \leq \mu_1 + \epsilon_1, \dots, y_{i,m} \leq \mu_m + \epsilon_m, \sum_{j=1}^m y_{i,j} \leq K. \end{aligned} \quad (9)$$

The game stops. The solutions to (8) and (9) are respectively denoted by $\pi_i^t(\boldsymbol{\mu}^t)$, $t = 0, 1$ where $\boldsymbol{\mu}^0 = \boldsymbol{\mu}$ and $\boldsymbol{\mu}^1 = \boldsymbol{\mu} + \boldsymbol{\epsilon}$. To develop closed-form expressions to recipient i 's payoff $\pi_i^t(\boldsymbol{\mu}^t)$, we use bijections on the set of item indexes to recipient-dependent ordering of $v_{i,j}$ so as to generalize recipient i 's preference. In particular, the valuation of item $\tau^{-1}(j)$ is ranked j for recipient i whenever there exists some bijection $\tau: \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ such that $v_{i,\tau^{-1}(1)} \geq v_{i,\tau^{-1}(2)} \geq \dots \geq v_{i,\tau^{-1}(m)}$.

Payoff Description. For a given strategy profile $(s_1, s_2) \in S$, capacity K , and volume of items $\boldsymbol{\mu}$, we denote $\Pi_{i,t}^{s_1, s_2}(\boldsymbol{\mu}^t)$ as recipient i 's payoff at stage $t + 1$. To complete Table 1, the result below allows us to specify $\pi_i^0(\boldsymbol{\mu})$ and $E[\pi_i^1(\boldsymbol{\mu} + \boldsymbol{\epsilon})]$.

PROPOSITION 7. *Let $\mathbf{y}_i^* = (y_{i,1}^*, \dots, y_{i,m}^*)$ and suppose there exists some τ on $\{1, \dots, m\}$ such that $v_{i,\tau^{-1}(j)}$ is non-increasing in j . If recipient i orders, then the optimal solution to (8) is given by*

$$\mathbf{y}_{i,j}^* = \begin{cases} \mu_j, & \text{if } j \in A^\tau; \\ K - \sum_{j=1}^{\tau^{-1}(\tilde{m}-1)} \mu_j, & \text{if } j = \tau^{-1}(\tilde{m}); \\ 0, & \text{otherwise.} \end{cases}$$

If the recipient waits, then the optimal solution to (9) is given by

$$\mathbf{y}_{i,j}^* = \begin{cases} E \left[\min \left(\left[K - \sum_{l=1}^{j-1} (\mu_{\tau^{-1}(l)} + \epsilon_{\tau^{-1}(l)}) \right]^+, \mu_j + \epsilon_j \right) \right], & \text{if } j \in A^\tau; \\ K - \sum_{j=1}^{\tau^{-1}(\tilde{m}-1)} y_{i,j}^*, & \text{if } j = \tau^{-1}(\tilde{m}); \\ 0, & \text{otherwise.} \end{cases}$$

where $S_k(\tau) = \sum_{j=1}^k \mu_{\tau^{-1}(j)}$, $\tilde{m} = \inf\{k \leq m \mid S_k(\tau) > K\}$ and $A^\tau = \{\tau^{-1}(1), \dots, \tau^{-1}(\tilde{m}-1)\}$.

Given that $\mu_{\tau^{-1}(1)} + \dots + \mu_{\tau^{-1}(\tilde{m})} \geq K$, then the additional arrival of $\epsilon_{\tilde{m}}$ would provide no value-add should recipient j chooses to wait. In fact, ordering any item ranked $\tilde{m} + 1$ onwards is suboptimal for recipient i . Therefore, the expressions $\pi_i^0(\boldsymbol{\mu})$ or $E[\pi_i^1(\boldsymbol{\mu} + \boldsymbol{\epsilon})]$ are given by the respective optimal solutions from Proposition 7, i.e., $\mathbf{v}_i \cdot \mathbf{y}_i^*$ whenever $\mathbf{v}_i = (v_{i,1}, \dots, v_{i,m})$ that can be used as inputs for completing Table 1.

B. Proofs

Proposition 1 (i). (O, O) is an equilibrium if and only if $\Pi_1^{O,O}(\boldsymbol{\mu}) \geq \Pi_1^{W,O}(\boldsymbol{\mu})$ and $\Pi_2^{O,O}(\boldsymbol{\mu}) \geq \Pi_2^{O,W}(\boldsymbol{\mu})$. These equations are always satisfied since $\Pi_1^{W,O}(\boldsymbol{\mu}) = \Pi_2^{O,W}(\boldsymbol{\mu}) = 0$, therefore (O, O) is always an equilibrium. The (W, O) and (O, W) strategies are dominated by (O, O) for both players.

(ii). The condition $\delta > \max\left(\frac{\bar{\delta}_1}{p_1}, \frac{\bar{\delta}_2}{p_2}\right)$ implies that $\delta > \frac{\bar{\delta}_i}{p_i} = \frac{\pi_i^0}{p_i E[\pi_i^1]} \Leftrightarrow \delta p_i E[\pi_i^1] > \pi_i^0$ for $i = 1, 2$. Under the R -model, these are equivalent to $\Pi_1^{W,W}(\boldsymbol{\mu}) > \Pi_1^{O,W}(\boldsymbol{\mu})$ and $\Pi_2^{W,W}(\boldsymbol{\mu}) > \Pi_2^{O,W}(\boldsymbol{\mu})$. Hence, (W, W) is an equilibrium.

(iii). Our final objective is to compare $W_R^{O,O}$ and $W_R^{W,W}$ for $\delta > \max(\bar{\delta}_1, \bar{\delta}_2)$. Then $W_R^{O,O} = p_1 \pi_1^0 + p_2 \pi_2^0 < \delta(p_1 E[\pi_1^1] + p_2 E[\pi_2^1]) = W_R^{W,W}$. This is because $\delta > \bar{\delta}_i$ is equivalent to $\delta E[\pi_i^1] > \pi_i^0$ for $i = 1, 2$. \square

Proposition 2 (i) We want to characterize $G(\boldsymbol{\epsilon}) \doteq \{(K, \boldsymbol{\mu}) : \delta p_i E[\pi_i^1] - \pi_i^0 > 0, \forall i\}$ in terms of capacity decision. For our purpose, we analyze (5) by letting

$$g_i(K) = -(1 - p_i \delta)(K v_{i,-i} + \mu_i \Delta_i) + \delta p_i \Delta_i E[\min(K - \mu_i, \epsilon_i)].$$

(W, W) is an equilibrium for those K for which $g_1(K)$ and $g_2(K)$ are simultaneously positive. To characterize $G(\boldsymbol{\epsilon})$, we characterize the g_1 and g_2 functions.

We note that $g_i(0) < 0$ and $g_i(\mu_i) < 0$ since ϵ_i is a non-negative random variable. Let F_i and f_i be the c.d.f. and p.d.f. for ϵ_i , respectively. We have

$$\frac{d}{dK} g_i(K) = -(1 - p_i \delta) v_{i,-i} + \delta p_i \Delta_i (1 - F_i(K - \mu_i)). \quad (10)$$

$$\frac{d^2}{dK^2} g_i(K) = \begin{cases} 0 & \text{if } K \leq \mu_i; \\ -f_i(K - \mu_i) \delta p_i \Delta_i & \text{o/w.} \end{cases}$$

Case 1: Let $\frac{v_{1,2}}{v_{1,1}} \geq \delta p_1$ or $\frac{v_{2,1}}{v_{2,2}} \geq \delta p_2$. We will show that at least one of $\frac{d}{dK} g_i(K) < 0$ for all $K > 0$, which implies that $G(\boldsymbol{\epsilon}) = \emptyset$. Without loss of generality, let us assume that $\frac{v_{1,2}}{v_{1,1}} \geq \delta p_1$. This condition implies that $\frac{d}{dK} g_1(K) = -(1 - p_1 \delta) v_{1,2} + \delta p_1 \Delta_1 = -v_{1,2} + \delta p_1 v_{1,1} < 0$ over the interval $K \in (0, \mu_1)$. Let us observe what happens for $K \in [\mu_1, \infty)$. We define $\varphi_i = F_1^{-1}\left(\frac{\delta p_i v_{i,i} - v_{i,-i}}{\delta p_i (v_{i,i} - v_{i,-i})}\right)$, then the sign of the derivative in (10) becomes

$$\frac{d}{dK} g_i(K) = \begin{cases} \geq 0, & \text{if } \mu_i \leq K \leq \mu_i + \varphi_i; \\ < 0, & \text{if } K > \mu_i + \varphi_i. \end{cases}$$

Now $\frac{v_{1,2}}{v_{1,1}} \geq \delta p_1$ also implies that $\varphi_1 = 0$ and so, we have $\frac{d}{dK}g_1(\mu_1) = 0$. For $K > \mu_1$, the sign of $\frac{d}{dK}g_1(K)$ is again negative. In other words, the slope of $g_1(K)$ is always negative over all non-negative range of capacity. Together with $g(\mu_1) < 0$, we conclude that (5) is negative for recipient 1 and thus, $G(\epsilon) = \emptyset$.

Case 2: Let $\frac{v_{i,-i}}{v_{i,i}} < \delta p_i$. For $K \in (0, \mu_i)$, $\frac{v_{i,-i}}{v_{i,i}} < \delta p_i$ implies that $\frac{d}{dK}g_i(K) = -v_{i,-i} + \delta p_i v_{i,i} > 0$. Combining with the slope of $g_i(K)$ over $[\mu_i, \infty)$, we obtain

$$\frac{d}{dK}g_i(K) = \begin{cases} \geq 0, & \text{if } 0 \leq K \leq \mu_i + \varphi_i; \\ < 0, & \text{if } K > \mu_i + \varphi_i. \end{cases}$$

Since $g_i(K)$ is concave on $[\mu_i, \infty)$, $g_i(\mu_i) < 0$, $\lim_{K \rightarrow \infty} g_i(K) < 0$, and g_i attains its maximum at $\mu_i + \varphi_i$, if $g_i(\mu_i + \varphi_i) > 0$, then there exist two capacity levels $\hat{K}_i(\boldsymbol{\mu})$ and $\bar{K}_i(\boldsymbol{\mu})$ such that $\mu_i < \hat{K}_i(\boldsymbol{\mu}) < \bar{K}_i(\boldsymbol{\mu})$, $g_i(\hat{K}_i(\boldsymbol{\mu})) = 0$ and $g_i(\bar{K}_i(\boldsymbol{\mu})) = 0$. We choose $\hat{K}(\boldsymbol{\mu}) = \max(\hat{K}_1(\boldsymbol{\mu}), \hat{K}_2(\boldsymbol{\mu}))$ and $\bar{K}(\boldsymbol{\mu}) = \min(\bar{K}_1(\boldsymbol{\mu}), \bar{K}_2(\boldsymbol{\mu}))$. If $g_i(\mu_i + \varphi_i) \leq 0$, then $G(\epsilon) = \emptyset$.

(ii) We want to characterize $G(\epsilon) = \{(K, \boldsymbol{\mu}) : \delta p_i E[\pi_i^1] - \pi_i^0 > 0, \forall i\}$ in terms of the inventory level. In this case, (5) can be rewritten as:

$$\begin{aligned} h_i(\mu_i) &= -(1 - p_i \delta)(K v_{i,-i} + \mu_i \Delta_i) + \delta p_i \Delta_i E[\min(K - \mu_i, \epsilon_i)] \\ &= -(1 - p_i \delta)K v_{i,-i} + \delta p_i \Delta_i E[\min(K - \mu_i, \epsilon_i)] - (1 - p_i \delta)\mu_i \Delta_i. \end{aligned}$$

The term $-(1 - p_i \delta)K v_{i,-i}$ is constant and $\delta p_i \Delta_i E[\min(K - \mu_i, \epsilon_i)] - (1 - p_i \delta)\mu_i \Delta_i$ is strictly decreasing in μ_i . This is because $\frac{d}{d\mu_i} h_i(\mu_i) = -\Delta_i(1 - \delta p_i F_i(K - \mu_i)) < 0$. As $\lim_{y \rightarrow \infty} h_i(y) < 0$, if $h_i(0) > 0$, then there exists $I_i(K)$ such that $g_i(I_i(K)) = 0$. If $h_i(0) > 0$, then let $I_i(K) = 0$. Then

$$\text{recipient } i\text{'s optimal action} = \begin{cases} \text{orders,} & \text{if } \mu_i \geq I_i(K); \\ \text{waits,} & \text{if } \mu_i < I_i(K). \end{cases}$$

Thus, we choose $\mathbf{I}(K) = (I_1(K), I_2(K))$ and both recipients wait whenever $\boldsymbol{\mu} \leq \mathbf{I}(K)$.

(iii) Let $\boldsymbol{\epsilon} \leq_{st} \boldsymbol{\epsilon}' = (\epsilon'_1, \epsilon'_2)$ and $G(\boldsymbol{\epsilon}') = \{(K, \boldsymbol{\mu}) : \delta p_i E[\hat{\pi}_i^1] > \pi_i^0\}$ where $\hat{\pi}_i^1 = \Delta_i \min(K, \mu_i + \epsilon'_i) + K v_{i,-i}$. We want to show that $G(\boldsymbol{\epsilon}) \subseteq G(\boldsymbol{\epsilon}')$. $\boldsymbol{\epsilon} \leq_{st} \boldsymbol{\epsilon}'$ implies $E[f(\epsilon)] \leq E[f(\epsilon')]$ for any non-decreasing function f , provided the expectation exists (see Shaked and Shanthikumar (2007)). Pick $(K, \boldsymbol{\mu}) \in G(\boldsymbol{\epsilon})$. Given that $\boldsymbol{\epsilon} \leq_{st} \boldsymbol{\epsilon}'$ and $\min(K, \mu_i + y)$ is increasing in $y \in \mathbf{R}_+$, we must have

$$\delta p_i E[\hat{\pi}_i^1] = \Delta_i E[\min(K, \mu_i + \epsilon'_i)] + K v_{i,-i} \geq \Delta_i E[\min(K, \mu_i + \epsilon_i)] + K v_{i,-i} = \delta p_i E[\pi_i^1] > \pi_i^0.$$

Thus, we have $(K, \boldsymbol{\mu}) \in G(\boldsymbol{\epsilon}')$. \square

Proposition 3 (i) Let $\boldsymbol{\mu} \leq \boldsymbol{\mu}'$. Our goal is to show that at higher inventory levels, the range of capacity that sustains the (W, W) equilibrium gets reduced. Define $g'_i(K) = -(1 - p_i \delta)(K v_{i,-i} + \mu'_i \Delta_i) + \delta p_i \Delta_i E[\min(K - \mu'_i, \epsilon_i)]$. Suppose $\hat{K}_i(\boldsymbol{\mu}')$ and $\bar{K}_i(\boldsymbol{\mu}')$ solve $g'_i(K) = 0$. It suffices to show that $\hat{K}_i(\boldsymbol{\mu}) \leq \hat{K}_i(\boldsymbol{\mu}')$ and $\bar{K}_i(\boldsymbol{\mu}) \geq \bar{K}_i(\boldsymbol{\mu}')$. These (together with Proposition 2) will imply that

$$\begin{aligned} \hat{K}(\boldsymbol{\mu}) &= \max(\hat{K}_1(\boldsymbol{\mu}), \hat{K}_2(\boldsymbol{\mu})) \leq \max(\hat{K}_1(\boldsymbol{\mu}'), \hat{K}_2(\boldsymbol{\mu}')) = \hat{K}(\boldsymbol{\mu}') \\ \bar{K}(\boldsymbol{\mu}) &= \min(\bar{K}_1(\boldsymbol{\mu}), \bar{K}_2(\boldsymbol{\mu})) \geq \min(\bar{K}_1(\boldsymbol{\mu}'), \bar{K}_2(\boldsymbol{\mu}')) = \bar{K}(\boldsymbol{\mu}') \\ &\Rightarrow (\hat{K}(\boldsymbol{\mu}'), \bar{K}(\boldsymbol{\mu}')) \subseteq (\hat{K}(\boldsymbol{\mu}), \bar{K}(\boldsymbol{\mu})). \end{aligned}$$

Since $\boldsymbol{\mu} \leq \boldsymbol{\mu}'$, the strict decreasing property of $h_i(\mu_i)$ in μ_i implies that

$$\begin{aligned} g_i(K) &= -(1-p_i\delta)(Kv_{i,-i} + \mu_i\Delta_i) + \delta p_i\Delta_i E[\min(K - \mu_i, \epsilon_i)] \\ &\geq -(1-p_i\delta)(Kv_{i,-i} + \mu'_i\Delta_i) + \delta p_i\Delta_i E[\min(K - \mu'_i, \epsilon_i)]. \end{aligned}$$

Therefore, $g_i(\bar{K}_i(\boldsymbol{\mu}')) \geq g'_i(\bar{K}_i(\boldsymbol{\mu}')) = 0$. Due to the concavity of $g_i(K)$ in $K > 0$, we have $\bar{K}_i(\boldsymbol{\mu}') \leq \bar{K}_i(\boldsymbol{\mu})$. In a similar vein, we have $g_i(\hat{K}_i(\boldsymbol{\mu}')) \geq g'_i(\hat{K}_i(\boldsymbol{\mu}')) = 0$. Therefore, $\hat{K}_i(\boldsymbol{\mu}) \leq \hat{K}_i(\boldsymbol{\mu}')$. The result is proven.

(ii) Let $\boldsymbol{\epsilon} \leq_{st} \boldsymbol{\epsilon}'$. We want to show that for a stochastically larger arrival distribution, the capacity range supporting (W, W) becomes larger, i.e., $(\hat{K}(\boldsymbol{\mu}), \bar{K}(\boldsymbol{\mu})) \subseteq (\hat{K}'(\boldsymbol{\mu}), \bar{K}'(\boldsymbol{\mu}))$. It suffices to show that $\hat{K}'(\boldsymbol{\mu}) \leq \hat{K}(\boldsymbol{\mu})$ and $\bar{K}(\boldsymbol{\mu}) \leq \bar{K}'(\boldsymbol{\mu})$. Since $\min(y, x)$ is a non-decreasing function in x (see Shaked and Shanthikumar (2007)), we have

$$g_i(K) \leq -(1-p_i\delta)(Kv_{i,-i} + \mu_i\Delta_i) + \delta p_i\Delta_i E[\min(K - \mu_i, \epsilon'_i)]. \quad (11)$$

As $\hat{K}'(\boldsymbol{\mu})$ is the lower bound of $G(\boldsymbol{\epsilon}')$ that produces zero on the RHS of (11), we have $g_i(\hat{K}'(\boldsymbol{\mu})) \leq 0$. Furthermore, the bound $\hat{K}(\boldsymbol{\mu})$ satisfies $g_i(\hat{K}(\boldsymbol{\mu})) = 0$, implying that $g_i(\hat{K}'(\boldsymbol{\mu})) \leq g_i(\hat{K}(\boldsymbol{\mu}))$. Over $K < \mu_i + \varphi_i$, $g_i(K)$ is increasing and thus, $\hat{K}'(\boldsymbol{\mu}) \leq \hat{K}(\boldsymbol{\mu})$. For $K > \mu_i + \varphi_i$, $g_i(K)$ is decreasing and thus, $\bar{K}'(\boldsymbol{\mu}) \geq \bar{K}(\boldsymbol{\mu})$.

(iii) Recall that the inventory bound defining $G(\boldsymbol{\epsilon}') \neq \emptyset$ is given by $\mathbf{I}(K) = (I_1(K), I_2(K))$, where $I_i(K)$ satisfies $-(1-p_i\delta)(Kv_{i,-i} + I_i(K)\Delta_i) + \delta p_i\Delta_i E[\min(K - I_i(K), \epsilon_i)] = 0$. Taking first and second order derivatives w.r.t K , we have

$$\begin{aligned} &-(1-p_i\delta)(v_{i,-i} + I'_i(K)\Delta_i) + \delta p_i\Delta_i(1 - I'_i(K))(1 - F_i(K - I_i(K))) = 0 \\ &-(1-p_i\delta)I''_i(K)\Delta_i + \delta p_i\Delta_i[-I''_i(K)(1 - F_i(K - I_i(K))) - (1 - I'_i(K))^2 f_i(K - I_i(K))] = 0. \end{aligned}$$

On rearranging, we have

$$\begin{aligned} I'_i(K) &= \frac{\delta p_i\Delta_i(1 - F_i(K - I_i(K))) - (1-p_i\delta)v_{i,-i}}{\Delta_i(1 - \delta p_i F_i(K - I_i(K)))} \\ I''_i(K) &= \frac{-\delta p_i\Delta_i(1 - I'_i(K))^2 f_i(K - I_i(K))}{\Delta_i(1 - \delta p_i F_i(K - I_i(K)))}. \end{aligned}$$

Therefore, $I_i(K)$ is strictly concave on $K > 0$. Therefore, $\mathbf{I}(I_1(K), I_2(K))$ is component-wise concave in $K > 0$.

(iv) Let $\boldsymbol{\epsilon} \leq_{st} \boldsymbol{\epsilon}'$. We want to show that for a stochastically larger arrival distribution, the bound on $G(\boldsymbol{\epsilon})$ defined by the inventory levels increases. It suffices to show that $I_i(K) \leq I'_i(K)$ whenever $I'_i(K)$ that satisfies $-(1-p_i\delta)(Kv_{i,-i} + \mu_i\Delta_i) + \delta p_i\Delta_i E[\min(K - \mu_i, \epsilon'_i)] = 0$. In a similar argument to (11),

$$h_i(I'_i(K)) \leq -(1-p_i\delta)(Kv_{i,-i} + \mu_i\Delta_i) + \delta p_i\Delta_i E[\min(K - \mu_i, \epsilon'_i)] = 0 = h_i(I_i(K)).$$

Given that $h_i(y)$ is decreasing in $y > 0$ and $h_i(I'_i(K)) \leq h_i(I_i(K))$, we have $I'_i(K) \geq I_i(K)$. This implies that $\mathbf{I}'(K) = (I'_1(K), I'_2(K)) \geq (I_1(K), I_2(K)) = \mathbf{I}(K)$. \square

Proposition 4 This follows from Proposition 2(i), where we show that under the stated conditions, $G(\boldsymbol{\epsilon}) = \emptyset$. \square

Proposition 5 Our goal is to show that $W_R \leq W_{SR}$. Observe that we have

$$\Pi_i^{SR} = \max(\pi_i^0, \delta E[\pi_i^1]) = \frac{1}{p_i} \max(\Pi_i^{O,O}, \Pi_i^{W,W}) \Rightarrow p_i \Pi_i^{SR} \geq \max(\Pi_i^{O,O}, \Pi_i^{W,W}).$$

The welfare generated by MSRO based on the realized equilibrium under R -model:

$$W_R = \begin{cases} \Pi_1^{O,O} + \Pi_2^{O,O}, & \text{if } \delta \leq \bar{\delta}; \\ \Pi_1^{W,W} + \Pi_2^{W,W}, & \text{if } \delta > \bar{\delta}. \end{cases}$$

Therefore, we have $W_{SR} = p_1 \Pi_1^{SR} + p_2 \Pi_2^{SR} \geq \max(\Pi_1^{O,O}, \Pi_1^{W,W}) + \max(\Pi_2^{O,O}, \Pi_2^{W,W}) \geq W_R$. \square

Proposition 6 Our goal is to show that $W_P \geq W_R$. To do so, let us define $W'_{SR} = \max(\Pi_1^{SR}, \Pi_2^{SR})$. This is the welfare generated by the MSRO that assigns inventory access to the recipient with the highest expected payoff. We proceed as follows: (i) We show that $W'_{SR} \geq W_{SR}$, i.e., the welfare by assigning access to the recipient with the highest expected payoff dominates W_{SR} , the randomized selection for exclusive access, (ii) the welfare generated by P -model dominates, i.e., $W_P \geq W'_{SR}$.

(i) To show that $W'_{SR} \geq W_{SR}$, we note that since $p_1 + p_2 = 1$, we have

$$W'_{SR} = p_1 \max(\Pi_1^{SR}, \Pi_2^{SR}) + p_2 \max(\Pi_1^{SR}, \Pi_2^{SR}) \geq p_1 \Pi_1^{SR} + p_2 \Pi_2^{SR} = W_{SR}.$$

(ii) To show $W_P \geq W'_{SR}$, we note that since

$$\begin{aligned} W_p &= \max(\max(\pi_1^0, \pi_2^0), \delta E[\max(\pi_1^1, \pi_2^1)]) \\ &\geq \delta E[\max(\pi_1^1, \pi_2^1)] \geq \delta \max(E[\pi_1^1], E[\pi_2^1]). \end{aligned}$$

The second inequality comes from applying Jensen's inequality. The definition of W_p also implies that $W_p \geq \max(\pi_1^0, \pi_2^0)$. Therefore, we have

$$\begin{aligned} W_p &\geq \max(\max(\pi_1^0, \pi_2^0), \delta \max(E[\pi_1^1], E[\pi_2^1])) \\ &= \max(\max(\pi_1^0, \delta E[\pi_1^1]), \max(\pi_2^0, \delta E[\pi_2^1])) = \max(\Pi_1^{SR}, \Pi_2^{SR}) = W'_{SR}. \end{aligned}$$

Once we establish (i) and (ii), the proof follows from Proposition 5 since $W_{SR} \geq W_R$. \square