Managing Inventories of Reusable Containers for Food Take-Away at a Restaurant

Nicole Perez Becker[a], Joachim Arts[a,*], Sven Reichardt[b], Anne Lange[c]

[a] *Luxembourg Center for Logistics and Supply Chain Management, University of Luxembourg, 6, Rue Richard Coudenhove – Kalergi, L-1359 Luxembourg, Luxembourg*

[b] *Department of Physics and Materials Science, University of Luxembourg, 162a, Avenue de la Faïencerie, L-1511 Luxembourg, Luxembourg*

[c] *Frankfurt University of Applied Sciences, Nibelungenplatz 1, 60318 Frankfurt am Main, Germany*

Abstract

Single-use packaging in the food services sector accounts for a substantial amount of waste, leading some restaurants to offer customers reusable containers for take-away or delivery orders. Once used, customers may return these containers to the same restaurant or another restaurant in the reusable container network or they may not return them at all. The restaurant hence faces both uncertain demand and returns for reusable containers and needs to decide on the number of containers to stock to serve its customers. We formulate this problem by modeling it as a continuous-time Markov Decision Process. Through a numerical study, we investigate the effect that different balances of demand and return intensities and their coupling have on the average total cost for the restaurant. We find that greater demand and return coupling reduces average costs, but the effects are most beneficial when the overall demand and returns of the restaurant are balanced. The restaurant can reduce costs by optimizing the supplier visit frequency in addition to the inventory level of clean containers after the supplier visit. The supplier's choice of the level of the visit cost is important as smaller scale restaurants may be penalized by a larger supplier visit cost, dissuading them from participating in reusable container systems.

Keywords: Closed-loop supply chain management, Inventory, Waste, Sustainability, Sustainable consumption, Reusable containers

1. Introduction

Every minute McDonald's uses 2.8 tons of single-use packaging to serve its customers worldwide (Zero Waste France, 2017). Every year Starbucks uses approximately seven billion disposable cups worldwide (Lucas, 2022). The food services sector generates a substantial amount of single-use packaging waste. In

*Corresponding author

Email addresses: `nicole.perez@alumni.duke.edu` (Nicole Perez Becker), `joachim.arts@uni.lu` (Joachim Arts), `sven.reichardt@uni.lu` (Sven Reichardt), `anne.lange@fb3.fra-uas.de` (Anne Lange)

the United Kingdom, lunches-to-go are estimated to produce almost 11 billion items of single-use packaging waste annually (Smithers, 2019). The rapid growth of take-away and food delivery services over the past years, facilitated by platforms and amplified by the COVID-19 pandemic, has only exacerbated this problem. Since 2017, the food delivery sector alone has more than tripled its revenues and is currently worth $150 billion USD globally (Ahuja et al., 2021).

Most single-use packaging waste cannot be recycled as it consists mainly of plastic products or plastic-coated paper products (used in most disposable coffee cups) that cannot be handled by the regular paper recycling process. As a result, much of this waste is landfilled (Schupak, 2021). In the United States, single-use containers and packaging in general account for more than 23% of the waste in landfills (US Environmental Protection Agency, 2015). This waste accounts for a significant amount of carbon emissions. Single-use packaging waste that is not recycled or landfilled is discarded into the environment, where it is left to degrade and interact with animal life.

Given the detrimental environmental effects of single-use packaging in the food sector, many jurisdictions are passing regulations to reduce its use. One type of regulation is charging customers a fee for purchasing a product in single-use packaging. For example, customers in the cities of Berkeley, California, and Vancouver, Canada, incur a surcharge of $0.25 USD ($0.25 CAD, respectively) per beverage purchased in a disposable cup (Peters, 2020; City of Vancouver, 2022). In the Netherlands, starting in July 2023, a similar measure that extends a surcharge to all single-use plastic food packaging will come into effect (Netherlands Chamber of Commerce, 2022). Some countries are going even further in their efforts to limit single-use packaging. For example, from January 2023, restaurants in Germany will be required to provide a reusable packaging alternative at no extra cost to the customer for products currently offered in single-use packaging (Germany Federal Ministry for the Environment, 2021). In Luxembourg, from January 2025, take-out and delivery meals will only be served in reusable containers (Gouvernement du Grand-Duché de Luxembourg, 2022).

The implementation of reusable container systems in the food take-away and delivery sector has become an active area of work for start-ups, public entities, and restaurant chains. While a few restaurant chains run their own reusable container systems, most reusable container systems are operated by a third-party supplier for a network of restaurants. The business models behind these reusable container systems are rapidly evolving and highly diverse. Important questions in the design of a reusable container system are how to ensure that the system is efficient, effective, and economically sustainable. Much of the diversity in business models for reusable container systems stems from the way different systems aim to address these questions. Most third-party suppliers and restaurants, for example, charge a deposit refundable upon return per clean container. This deposit-based approach aims to minimize shrinkage in the inventory of reusable containers by incentivizing returns. Some third-party suppliers operate more technologically integrated systems, deploying an app or QR codes, that enable restaurants and customers to easily track containers and deposits.

Setting system design and incentive questions to the side, however, the mere introduction of reusable containers at a restaurant brings about multiple operational challenges that may already be enough to dissuade a restaurant from participating in such a system altogether. One of the main decisions for a restaurant using reusable containers is the number of reusable containers to have on-hand. In particular, a restaurant faces a variety of customers that affect the restaurant's reusable container inventory level in different ways. Some customers, for example, only demand a clean reusable container with their order, reducing the number of clean containers in the restaurant's inventory. Other customers both demand a clean container with their order and return a dirty container, effectively having a net zero effect on inventory (after a lag time for cleaning the container before it can be used again). To make returns easier for customers, many reusable container systems are designed so that restaurants can also serve as drop-off points, meaning that some customers only return dirty containers to the restaurant, increasing the inventory level of dirty containers. Hence, a restaurant that uses reusable containers faces both uncertain demand and uncertain returns, making it difficult to control inventory levels.

In as much as the restaurant faces uncertain demands and returns, inventory management of a reusable container system resembles a number of well-studied systems in the closed-loop supply chain (e.g., repairable item and remanufacturing inventory systems) or sharing economy (e.g., bike-sharing systems) literature. However, these systems differ from the systems in our setting in two notable ways. First, in the closed-loop supply chain literature, unfulfilled demand is typically backlogged and not lost, unlike in our setting. Second, and most notably, our setting includes customers that both demand a clean container with their order and return a dirty container, resulting in a *coupled demand and return*. These customers are present in addition to customers that only demand a clean container for their meal (similar to a traditional forward flow supply chain setting) and customers that only return a dirty container (similar to a traditional reverse flow supply chain setting). Having a larger base of customers with coupled demand and returns is likely beneficial to a restaurant as these customers each have a net zero effect on the restaurant's inventory, making the restaurant more internally sustainable in terms of inventory levels and enabling the restaurant to reduce costs. Such high coupling of demand and returns may occur, for instance, in a restaurant with a relatively large loyal base of frequent customers who regularly order a meal in a reusable container and return a previously used container at the same time. In this sense, the existence of customers with coupled demand and returns may give a restaurant using reusable containers an advantage over systems with uncertain, yet independent demands and returns. However, only a few works in the closed-loop supply chain literature (i.e., Van der Laan et al., 1999; Kiesmüller, 2003) study systems in which customers with coupled demand and returns co-exist with customers that only demand or only return a product.

Our objective in this paper is to study the inventory decisions of a manager of a restaurant that participates in a reusable container system and faces customers that have different effects on reusable container inventory levels, including those that generate a coupled demand and return. In particular, we address the following questions:

(i) What is the optimal inventory policy for the restaurant?,

(ii) How does the degree of demand and return coupling affect this policy and the restaurant's costs?, and

(iii) How do other system characteristics affect the restaurant's inventory decisions and costs?

To answer these questions, we use a continuous-time Markov Decision Process to model the inventory decisions of a restaurant participating in a reusable container system. In this system, every visit of the reusable container supplier to the restaurant is an opportunity for the restaurant to rebalance its inventory. If returns or demands become too high, excess clean containers can be given to the supplier or additional clean containers can be taken from the supplier. We determine the restaurant's *optimal rebalancing policy* for the supplier visits and the supplier's optimal visit frequency from the restaurant's perspective if the restaurant is able to decide on this frequency. We model the degree to which customers' demand and returns are coupled by defining three customer streams: a stream of customers that only demand a clean container, a stream of customers that only return a dirty container, and a stream of customers that both demand a clean container and return a dirty container. Through a numerical study, we investigate the sensitivity of the optimal inventory balancing policy, the optimal supplier visit frequency, and the restaurant's costs to changes in parameters including the lost sales penalty (a cost parameter that is highly influenced by government policy for single-use containers), ratio of demand to returns, proportion of demand coupled with returns, scale of the restaurant, supplier visit costs, and dishwasher utilization.

We find that the optimal rebalancing policy is a state-dependent policy in which the optimal rebalancing level depends on the number of dirty containers at the restaurant. We also find that the restaurant's costs of operating a reusable container system can be decreased by optimizing both the rebalancing policy and the supplier visit frequency. In terms of the effect of customers with coupled demand and returns on the restaurant's performance, our results support the intuition that greater coupling of demand to returns allows the restaurant to decrease its expected total costs. This result holds irrespective of the overall balance of demand and returns at the restaurant. However, the effect of greater coupling of demand to returns is more substantial when overall demand and returns at the restaurant are more balanced. This finding highlights the greater importance of maintaining an overall balance of demands and returns as a restaurant. We also find that the third-party supplier visit cost is an important lever in making the system viable for a restaurant. Specifically, if a restaurant is mostly a collector of dirty containers or a dispenser of clean containers, a third-party supplier can make participating in the reusable container system more viable for the restaurant by reducing its visit costs. A higher supplier visit cost may also disproportionately penalize restaurants with lower demand for reusable containers. Through our modeling of this setting and our findings, we contribute to the literature on sustainable inventory systems by deriving insights into the conditions that make it easier for a restaurant to participate in a reusable container system and factors that policy-makers and reusable container suppliers can influence to make participation more economically sustainable and appealing for a restaurant.

The rest of this paper is structured as follows. Section 2 briefly reviews the literature. Section 3 describes the modeling approach. Section 4 formulates and solves the restaurant's optimal inventory balancing and optimal balancing frequency decision problems. Section 5 describes the performance metrics for the restaurant. Section 6 investigates the effect of varying different parameters on these performance metrics through a numerical study. Section 7 relaxes two of the assumptions in our base model and shows the robustness of our results with respect to these assumptions. Section 8 concludes with our main findings and future research directions.

2. Literature Review

The rapid growth in the food take-away and delivery sector has sparked an interest in restaurant operations and take-away / delivery platform operations within the Operations Management (OM)/ Operations Research (OR) community. Mao et al. (2022) focus on the delivery challenges that platforms face, highlight opportunities for future research, and provide a dataset that includes two months of orders from an online meal delivery platform operating in Hangzhou, China. Although platforms boost the restaurant's visibility and outsource delivery, these benefits also come at the cost of greater congestion in the kitchen and a cut from the restaurant's margins. Feldman et al. (2022) and Chen et al. (2022) model the relationship between platforms and restaurants and identify contract types that can coordinate the supply chain. Unlike these papers, we focus on the operational challenges of managing a restaurant that uses reusable containers (as opposed to disposable containers) to reduce its environmental impact.

Reusable containers have been studied in the context of food production and distribution systems in the past. Glock & Kim (2014) provide a literature review for models with returnable packaging materials. Much of this literature uses deterministic demand/return models. Glock (2017) is a recent example for food packaging specifically. Accorsi et al. (2022) do a transport study of food packaging materials over a network with known demand and provide a robust formulation.Taheri et al. (2021) study reusable container systems for consumer goods products and focus on the role of incentives, particularly in terms of the trade-off between durability of the containers and levels of deposit. These studies examine a reusable container system between a wholesaler and a retailer. Demand that cannot be met immediately is backordered in these settings. By contrast, demand that cannot be met immediately is lost in our setting and our modeling set-up reflects this complexity.

In terms of modeling, our work relates to the literature on inventory management in closed-loop supply chains and in sharing economy applications. Closed-loop supply chains are characterized by a reverse flow of products from the customer to the manufacturer in addition to the forward flow from the manufacturer to the customer. Reviews of closed-loop supply chain research include Guide & Van Wassenhove (2009), Fleischmann et al. (1997), and Souza (2013). Within the closed-loop supply chain literature, repairable item inventory systems and remanufacturing / hybrid production inventory systems have been extensively studied. In a repairable item inventory system, the breakdown of an item in use by a customer triggers a return of the item to the manufacturer for repair and an immediate demand for a new working item to replace the defective item. If replacement items are not in stock, demand is backordered. Because demand and returns are perfectly correlated, there is no uncertainty about the quantity and timing of returns compared to the demand. A review of research on repairable item inventory systems is provided by Guide Jr & Srivastava (1997).

In a remanufacturing / hybrid production inventory system, a manufacturer can produce a new product from scratch or remanufacture the product using recovered materials. Product is recovered from the market once the customer has no further use for the product, upgrades the product, or when the product reaches its end-of-life. A common assumption in this literature is that demand is independent of returns (e.g., Fleischmann et al., 1997; DeCroix, 2006; DeCroix et al., 2005). A few works, however, study systems in which uncertain demand and returns are correlated either in the same time period or across time periods. For a periodic review hybrid production system operating on a finite horizon, Kiesmüller & Van der Laan (2001) model the relationship between returns that occur in a current period and demands that occurred in a previous period. They find that ignoring this time-dependence of demand and returns can result in higher costs. For a continuous review hybrid production system, Van der Laan et al. (1999) model a correlation between demand and returns in the same period as the probability that a return will trigger an immediate demand. They evaluate the system under two types of control policies and find that increased correlation between demand and returns results in lower costs under either policy, but is especially beneficial when the return rate is high.

While the literature on these closed-loop supply chain inventory systems provides a background for our study, these systems are fundamentally different from the system we study in several ways. First, demand that is not fulfilled is backordered. This assumption is reasonable for specialized or expensive products such as manufacturing equipment or in a business-to-business context, but not for consumer products that can easily be substituted by another product. Second, either customers generate perfectly coupled demand and returns (as in the repairable item literature) or customers generate demand and returns that are not explicitly related (as in the remanufacturing literature excluding the above-mentioned exceptions). As such, it is not necessary to consider the arrival of customers that are more diverse in terms of their impact on inventory levels (namely the co-existence of customers that only demand a product or return a product with customers that demand and return a product at the same time) and how the balance of each of those customer types affects the overall system.

Sharing economy inventory systems – and in particular bike-sharing systems (see Kabra et al., 2016, for an overview) – do resemble our problem setting more, in that demand can be lost. In a bike-sharing system, customers collect a bike at one station and return it to the same or another station after a period of time, resulting in uncertainty in both the timing and location of returns. If a bike is not available at a station, the user may decide to take a bike from another station (resulting in a spillover of demand to another location) or substitute biking altogether with another means of transport (resulting in a lost sale). However, bike-sharing systems have other features that do not translate well into our setting. First, demand and returns are never coupled (i.e., a return of a bike does not trigger a demand for a bike). Second, bike-sharing system operators, in addition to managing the inventory of bikes at each station, also must manage the inventory of available docks for bikes to be returned. The number of docks available limits the number of returns that can be accepted at each station. To deal with imbalances between demand and availability for bikes and docks at individual stations throughout the day, bike-sharing system operators reposition the bikes. The inventory repositioning decisions in addition to the inventory placement decisions have been the focus of many studies. For realistically sized applications, solving for these decisions can be computationally challenging under different types of objectives or focusing on different problem features (e.g., Raviv & Kolka, 2013; Datner et al., 2019; Shu et al., 2013). In summary, while bike-sharing inventory systems provide insights into how to manage loss inventory systems with uncertain demand and returns, the literature on this area is also limited in terms of its applicability to our setting.

Given that related inventory systems are different in several key ways from the inventory system in our problem setting, our contribution is to model the unique aspects of a restaurant that uses reusable containers instead of single-use containers to serve its take-away and delivery orders. Specifically, we model a lost-sales system with uncertain demand and returns. In this system, customers that generate coupled demand for a new product and returns of a used product co-exist with customers that only demand a new product or customers that only return a used product. By studying the restaurant's inventory problem, we develop an understanding of how the level of the cost parameters and the degree of demand and return coupling affect the overall costs of operating such a system. Aside from assisting the restaurant, this understanding can help policy-makers in assessing to what degree they should encourage the use of reusable containers or third-party suppliers in designing systems with the operational challenges of different types of restaurants in mind.

3. Model Description

We consider a restaurant that stocks reusable containers for take-away and delivery food products. The restaurant manages two types of inventories: an inventory of clean containers (denoted by x_C) and an inventory of dirty containers (denoted by x_D). Three different types of customers visit the restaurant. Type 1 customers only demand a meal in a reusable container. Type 2 customers both demand a meal in a reusable container and return a dirty container, hence demand and returns are coupled. If no clean container is available, the type 2 customer still returns the dirty container. Type 3 customers only return a dirty container. Customers of type $i \in \{1, 2, 3\}$ arrive to the restaurant according to a Poisson Process with rate λ_i.

All returns of dirty containers are processed by the restaurant's dishwasher. For simplicity of the model, we assume that the time to wash one container is exponentially distributed with mean μ_I^{-1}. For stability we require $\lambda_2 + \lambda_3 < \mu_I$. In Section 7.1, we discuss both the validity of this choice of dishwashing time distribution in the context of more realistic dishwasher operating times and extend our model to batch dishwashing.

A third-party supplier of clean containers operates the system. When the supplier visits, he can either provide more clean containers or collect excess clean containers from the restaurant, thereby helping the restaurant control its inventory of clean containers. The supplier visits the restaurant according a Poisson process with rate μ_E. To make sure that this choice of distribution does not lead to model-dependent artifacts in our results, we relax this assumption in Section 7.2 and allow the time between supplier visits to be Erlang distributed. The visit frequency μ_E will be a decision variable for the restaurant. Figure 1 illustrates the dynamics of the system.

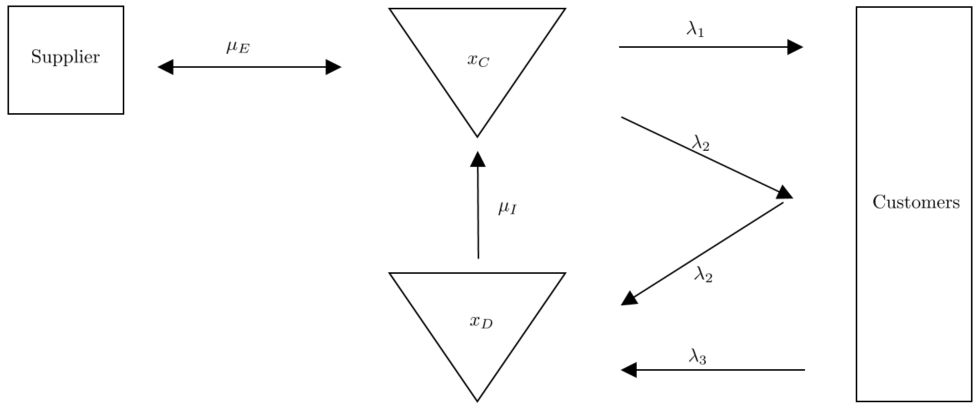

Figure 1: Dynamics of the reusable container system at the restaurant

The restaurant incurs a fixed cost c per supplier visit and a holding cost h per container held per time unit. If the restaurant does not have any clean containers on hand when a customer demands a container, the restaurant incurs a penalty p for not being able to satisfy the demand for a reusable container. This penalty represents the negative consequences of a stock-out of clean reusable containers at the moment of demand (i.e., a "lost sale"). In practice, this penalty can represent the cost of the single-use container used to pack the meal instead of the reusable one, a loss-of-goodwill costs such as the cost of decreased customer trust in the restaurant's commitment to sustainability, or the cost of customer dissatisfaction in having to pay a surcharge for a disposable container. However, in jurisdictions where single-use packaging is banned, it could be as much as the cost of losing the sale of the food product altogether. In fact, the level of this penalty is the main parameter that the policy-maker can influence in this model through regulations. The system operates in continuous time over an infinite time horizon. The restaurant's objective is to minimize the long-run average cost rate by deciding on *(i)* the number of clean containers that it should have in inventory after the supplier visits and *(ii)* the supplier visit frequency.

4. Optimal Inventory Balancing Policy and Balancing Frequency

4.1. Optimal Inventory Balancing Policy

The decision of how much inventory of clean containers to take from or give to the supplier can be modeled as a Markov Decision Process (MDP) for any given supplier visit frequency μ_E. The state of this MDP is the tuple $(x_C, x_D) \in \mathbb{N}_0^2$ where x_C and x_D denote the inventory level of clean and dirty containers at the restaurant, respectively. The decision in this MDP is the number of clean containers to have in inventory at the restaurant after a visit from the supplier, where the action space is $\mathbb{N}_0$. Note that, if we let y denote the number of clean containers at the restaurant after a visit from the supplier and if the restaurant is in state (x_C, x_D) before the visit, then $y - x_C$ is the number of clean containers taken from the supplier. A negative value of $y - x_C$ indicates that clean containers were returned to the supplier.

To transform this continuous time MDP to a discrete time MDP, we use uniformization with $\gamma = \sum_{i=1}^{3} \lambda_i + \mu_I + \mu_E$ as the uniform transition rate and scale time such that $\gamma = 1$. The probability of transitioning from state (x_C, x_D) to state (x'_C, x'_D) under decision $y \in \mathbb{N}_0$ is given by:

$$p((x_C, x_D), y, (x'_C, x'_D)) = \begin{cases} \lambda_1, & \text{if } x'_C = (x_C - 1)^+ \text{ and } x'_D = x_D \\ \lambda_2, & \text{if } x'_C = (x_C - 1)^+ \text{ and } x'_D = x_D + 1 \\ \lambda_3, & \text{if } x'_C = x_C \text{ and } x'_D = x_D + 1 \\ \mu_I, & \text{if } x'_C = x_C + 1,\ x'_D = x_D - 1, \text{ and } x_D > 0 \\ \mu_I, & \text{if } x'_C = x_C,\ x'_D = x_D, \text{ and } x_D = 0 \\ \mu_E, & \text{if } x'_C = y \text{ and } x'_D = x_D \\ 0, & \text{otherwise} \end{cases}$$

The direct cost of being in state (x_C, x_D) is given by $h(x_C + x_D) + p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0)$, where $\mathbb{I}(\cdot)$ is the indicator function. Note that the number of dirty containers at the restaurant x_D is not affected by the choice of y but only by the rate of returns $\lambda_2 + \lambda_3$ and dishwasher capacity μ_I. As such, the holding costs associated with dirty containers are sunk costs and the relevant direct cost function can also be written as $h(x_C) + p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0)$. However, we keep the holding cost from the dirty containers in the relevant direct cost function for completeness.

A policy $\pi : \mathbb{N}_0^2 \to \mathbb{N}_0$ is a decision rule that prescribes an action y to every possible system state (x_C, x_D). Policy π induces a stochastic process $(X_C^\pi(t), X_D^\pi(t))$ where $t \in \mathbb{N}_0$ is the indexed time unit in the horizon. Let Π be the set of Markovian policies. The average cost rate of a given rebalancing policy π is given by

$$g(\pi) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T h(X_C^\pi(t) + X_D^\pi(t)) + p(\lambda_1 + \lambda_2)\mathbb{I}(X_C^\pi(t) = 0)dt\right].$$

We seek to minimize this average cost rate. We let $g^* = \inf_{\pi\in\Pi} g(\pi)$ denote the optimal average cost rate and π^* denote the optimal policy that achieves this average cost rate, thus $g^* = g(\pi^*)$.

For g^* to exist and be finite, the Markov Chain induced by a Markovian policy must have a stationary distribution. Puterman (2014) outlines conditions to establish the existence of a stationary distribution. The present MDP is unichain and aperiodic because state $(0, 0)$ can be reached from any other state under any policy and because it has a self-transition. Furthermore, observe that the number of dirty containers in the restaurant (X_D) behaves as the number of customers in an M/M/1 queue with arrival rate $\lambda_2+\lambda_3$ and service rate μ_I, regardless of any inventory balancing policy, because returns are always processed. The stability condition $\lambda_2 + \lambda_3 < \mu_I$ ensures that X_D has a stationary geometric distribution with dishwasher utilization $\rho_I := (\lambda_2 + \lambda_3)/\mu_I$ as the parameter for any policy π (see Gross et al. (2008) for general results on M/M/1 queues). When demand exceeds returns, X_C will have a stationary distribution under any Markovian policy. When returns exceed demand, a Markovian policy needs to return clean containers to the supplier at a rate of at least $\lambda_3 - \lambda_1$. Otherwise, X_C will have positive drift and build up. A policy that reduces clean container inventory to any finite number during supplier visits will avoid this drift and so an optimal policy will too. Thus, there is a large class of Markovian policies that includes an optimal Markovian policy such that X_C will have a stationary distribution. For the remainder of the paper we will only look at such policies. Given this discussion, there exists an optimal Markovian policy π^* and an optimal cost rate g^* that satisfy the Bellman optimality equations:

$$
\begin{aligned}
V(x_C, x_D) + g^* = {} & h(x_C + x_D) + p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0) + \lambda_1 V((x_C - 1)^+, x_D) + \lambda_2 V((x_C - 1)^+, x_D + 1) \\
& + \lambda_3 V(x_C, x_D + 1) + \mu_I \mathbb{I}(x_D = 0) V(x_C, x_D) + \mu_I \mathbb{I}(x_D > 0) V(x_C + 1, x_D - 1) \\
& + \mu_E \min_y V(y, x_D) \qquad \forall (x_C, x_D) \in \mathbb{N}_0^2
\end{aligned}
\tag{1}
$$

where $V(x_C, x_D)$ is the relative value function.

Inspection of the last term on the right hand side of the Bellman equations (1) reveals that for each possible number of dirty containers at the restaurant (x_D), there is an optimal number of clean containers that the restaurant wishes to have when the supplier visits the restaurant. This observation is important and stated in the following proposition.

Proposition 1. *There exist state-dependent rebalancing levels $y^*(x_D)$ for each number of dirty containers in the restaurant (x_D) such that it is average-optimal to increase (decrease) the number of clean containers to $y^*(x_D)$ when the supplier visits the restaurant.*

We solve this MDP numerically using the value iteration algorithm. The numerical evaluation requires us to set bounds for the state space variables x_C and x_D. Given that the dishwasher is an $M/M/1$ queue and the number of dirty containers in the restaurant X_D is a geometrically distributed random variable with parameter ρ_I, the upper bound for X_D is set at the value on the support of X_D that corresponds to the 99th-percentile of this distribution. We need a different approach to bound X_C since there can be a positive or negative flow of containers between the supplier and the restaurant and unfulfilled demand is lost. As a proxy, we use the demand between supplier visits, D_S, to compute an upper bound on the number of clean containers. The distribution of D_S is a geometric distribution with parameter $\frac{\lambda_1+\lambda_2}{\lambda_1+\lambda_2+\mu_E}$. The derivation of this distribution can be found in the Appendix. Similar to the upper bound for X_D, we set the upper bound for X_C at the 99th-percentile of D_S.

4.2. Optimal Inventory Balancing Frequency

For a supplier visit frequency μ_E, $g^*(\mu_E)$ is the average lost sales penalty and holding cost rate under an optimal balancing policy π^*. Given a supplier visit cost per visit of c, the restaurant incurs total costs with a rate of $c\mu_E$ for the supplier visits, meaning that the restaurant's relevant costs are

$$C(\mu_E) = c\mu_E + g^*(\mu_E).$$

Whereas $g^*(\mu_E)$ decreases in μ_E, the supplier visit cost term increases in μ_E, creating a trade-off for the restaurant. The restaurant's optimal supplier visit frequency is $\mu_E^* = \text{argmin}_{\mu_E} C(\mu_E)$. We compute μ_E^* using a golden section search. $C(\mu_E^*)$ are the restaurant's relevant costs under the optimal supplier visit frequency μ_E^*.

5. Performance Metrics

In steady state, the average total cost rate $g^* = h(\mathbb{E}[X_C] + \mathbb{E}[X_D]) + p(\lambda_1 + \lambda_2)\mathbb{P}(X_C = 0)$. This cost rate can be decomposed into the average total holding cost rate g_H^*, where $g_H^* = h(\mathbb{E}[X_C] + \mathbb{E}[X_D])$, and the average total lost sales cost rate g_L^*, where $g_L^* = p(\lambda_1 + \lambda_2)\mathbb{P}(X_C = 0)$. The average total supplier visit cost rate, g_S^*, is given by $g_S^* = c\mu_E^*$.

In addition to the average cost rates, several other metrics are used to assess the system's performance. Because unfulfilled demand is lost, the expected lost sales rate and the fill rate are important metrics to track. To clarify, we broadly use the term "lost sale" to refer to a demand that cannot be met by providing a meal in a clean reusable container, regardless of whether the meal itself can be sold in a disposable container or not. We use that Poisson Arrivals See Time Averages (PASTA) (Wolff, 1982) to evaluate the following performance metrics. The expected *lost sales rate* is given by:

$$\mathbb{E}[\text{Loss}] = \mathbb{P}(X_C = 0)(\lambda_1 + \lambda_2).$$

We compute this rate recursively. Details are provided in the Appendix. The *fill rate*, denoted by β, is the ratio of the fulfilled demand to total demand in steady state and is given by:

$$\beta = \frac{\lambda_1 + \lambda_2 - \mathbb{P}(X_C = 0)(\lambda_1 + \lambda_2)}{\lambda_1 + \lambda_2} = 1 - \mathbb{P}(X_C = 0).$$

The restaurant can be a net receiver or giver of clean containers when the supplier visits. We define the *flow* as the long term demand for clean containers from the supplier. A negative flow indicates a net outflow of clean containers from the restaurant to the supplier whereas a positive flow indicates a net inflow of clean containers from the supplier to the restaurant. The expected flow is given by:

$$\mathbb{E}[\text{Flow}] = \lambda_1 + \lambda_2 - \mathbb{P}(X_C = 0)(\lambda_1 + \lambda_2) - (\lambda_2 - \lambda_3) = \lambda_1 - \mathbb{P}(X_C = 0)(\lambda_1 + \lambda_2) + \lambda_3.$$

The final two performance metrics we track are the expected number of clean and dirty containers at the restaurant. Using general results on the number of items in a $M/M/1$ queueing system, the expected *number of dirty containers* at the restaurant is:

$$\mathbb{E}[X_D] = \frac{\rho_I}{1 - \rho_I} = \frac{\lambda_2 + \lambda_3}{\mu_I - \lambda_2 - \lambda_3}.$$

Rearranging the terms in the expression for the average total cost rate g^* in steady state, the expected *number of clean containers* at the restaurant can be derived as:

$$\mathbb{E}[X_C] = \frac{g^* - p(\lambda_1 + \lambda_2)\mathbb{P}(X_C = 0)}{h} - \mathbb{E}[X_D].$$

6. Numerical Study

As mentioned in Section 1, the uncertainty in both the demand for and returns of reusable containers makes it more difficult for the restaurant to control inventory levels than if it were just facing uncertainty in demand (i.e., as in the single-use containers case). If the restaurant faces on average too much demand relative to returns, it mostly uses the supplier's visit to obtain additional clean containers. If the restaurant faces too many returns relative to demand, it mostly uses the supplier's visit to offload excess containers. The overall balance of demand to returns at the restaurant is a key characteristic of the restaurant's reusable container operations and it affects the restaurant's inventory decisions and costs. To control for possible differences in inventory decisions and costs driven by differences in the balance of demand to returns, we define the average demand-to-returns ratio, τ, as a parameter that we vary systematically in our numerical study, where

$$\tau = \frac{\lambda_1 + \lambda_2}{\lambda_2 + \lambda_3}. \tag{2}$$

A restaurant's demand and returns are *balanced* when $\tau = 1$. When $\tau > 1$ (respectively $\tau < 1$), the restaurant has on average more (respectively less) demand for containers than returns.

Since the restaurant faces customers that are heterogeneous in terms of their effects on the inventory levels of reusable containers and one of the customer types it faces generates coupled demand and returns, one question we set out to investigate is how the proportion of coupled demand and returns out of total demand affects the restaurant's inventory decisions and costs. To do so, we define the average proportion of demand that is coupled with returns, η, as a second control parameter in our numerical study, where

$$\eta = \frac{\lambda_2}{\lambda_1 + \lambda_2}. \tag{3}$$

Scaling effects, such as economies of scale or congestion effects, may result in inappropriate comparisons as they may drive differences in inventory decisions and costs across systems. To control for such effects, we define a third control parameter for the scale of the average demand for reusable containers at the restaurant, denoted by Λ, where

$$\Lambda = \lambda_1 + \lambda_2. \tag{4}$$

To generate the instances in our numerical study, we calculate $\lambda_i, i \in \{1, 2, 3\}$ for fixed levels of τ, η, Λ. Solving for the system of equations consisting of equations (2), (3), and (4) we calculate $\lambda_i, i \in \{1, 2, 3\}$ as $\lambda_1 = (1 - \Lambda)\eta$, $\lambda_2 = \Lambda\eta$, and $\lambda_3 = \frac{\Lambda}{\tau} - \lambda_2$.

Note that the proportion of coupled demand and returns, η, is in fact limited by the demand-to-returns ratio, τ. That is, the greater the demand for reusable containers relative to returns (the higher $\tau > 1$), the more that customers that only demand a container but not return one (type 1 customers) dominate the overall demand relative to customers that have coupled demand and returns (type 2 customers). Type 2 customers help balance the restaurant's inventory whereas type 1 customers reduce inventory, tipping the overall demand and returns ratio so that demand outstrips returns. As a result, for a given demand-to-returns ratio $\tau > 1$, some levels of η may not be feasible. Mathematically, for such inconsistent τ and η values, λ_3 becomes negative. For example, for a restaurant with $\tau = 1.5$, type 1 customers dominate to the extent that the highest proportion of coupled demand and returns possible is 0.67 and any higher value has $\lambda_3 < 0$. Therefore, in our selection of parameters, we restrict our settings for η and τ to values that are feasible and consistent for all τ.

The last control parameter we define is the utilization of the dishwasher, ρ_I. Using the fact that the dishwasher behaves like an $M/M/1$ queue with utilization $\rho_I = \frac{\lambda_2 + \lambda_3}{\mu_I}$, we calculate μ_I for fixed levels of ρ_I from $\mu_I = \frac{\lambda_2 + \lambda_3}{\rho_I}$. Setting values for $\rho_I < 1$ trivially ensures that the stability condition $\lambda_2 + \lambda_3 < \mu_I$ is met. Table 1 summarizes the distinct parameter values considered in our numerical study.

In most service settings, the cost of a lost sale substantially outweighs the holding cost of a unit of inventory. Additionally, in our setting, depending on how stringent the policy-maker's policy is in discouraging the use of single-use containers, not having a clean reusable container for a customer's order could either mean that the restaurant needs to pack the order in a single-use container or cannot make the sale at all (in case of a ban on single-use containers). Since reusable containers are a response to policy interventions to minimize or eliminate the use of single-use packaging and we are interested in studying the optimal inventory

Input parameter	No. of values	Values
Holding cost (per time unit per unit), h	1	1
Underage penalty (per unit underage), p	4	50, 100, 150, 200
Supplier visit cost, c	4	100, 200, 300, 400
Ratio of demand to returns, τ	5	0.8, 0.9, 1.0, 1.1, 1.2
Proportion of demand that is coupled with returns, η	5	0, 0.2, 0.4, 0.6, 0.8
Scale of demand, Λ	3	50, 75, 100
Dishwasher utilization, ρ_I	4	0.6, 0.7, 0.8, 0.9
Total number of instances	4,800	

Table 1: Input parameter values for the full factorial test bed.

policies of a system under such packaging constraints, we set the penalty p for using the discouraged item of single-use packaging or not being able to sell an order relatively high in comparison to the holding cost h. We normalize $h = 1$ for ease of comparison of the relative differences to the other cost parameter values.

The supplier visit cost c captures the cost of the operator servicing the restaurant. As reusable container systems in food services are a relatively new business model that is still evolving, there is substantial variance in the ways in which suppliers charge a service fee. Our choice of the parameter values for the supplier visit cost c is therefore intended to cover a wide range, in which the supplier visit cost is from 0.5-8 times of the lost sales penalty. For example, suppose that a restaurant does not have a reusable container to make a sale in a jurisdiction in which single-use containers are banned and suppose the value of this sale is \$15. The supplier visit cost would then be anything in the range of $\$7.5 - \$120$. We believe this range is reasonable in reflecting the service fee, especially when taken in comparison to the lost sales penalty.

Generating instances in the way described yields a full factorial test bed of $1 \times 4 \times 4 \times 5 \times 5 \times 3 \times 4 = 4,800$ instances. We note that our approach as outlined in Section 4 allows us to obtain the optimal average cost rate and other performance metrics over an infinite time horizon *directly* through value iteration and does not require an event-by-event simulation. Table 2 summarizes the results of the numerical study.

Before delving into the effects of varying each of the parameters, two general observations can be made about Table 2. First, across the various parameter values, the expected loss $\mathbb{E}[\text{Loss}]$ is very low relative to the average demand. As a result, the fill rate β is almost always 1.00, hence we are operating in a high service regime. Second, the expected number of dirty containers $\mathbb{E}[X_D]$ is relative unaffected and insensitive to most parameters except for the dishwasher utilization ρ_I, which means that it is mostly the number of clean containers $\mathbb{E}[X_C]$ that is driving the holding cost rate in the system.

Parameter	Value	Count	μ_E^*	$C(\mu_E^*)$	g_L^*	g_H^*	g_S^*	$\mathbb{E}$[Loss]	$\mathbb{E}[X_C]$	$\mathbb{E}[X_D]$	$\mathbb{E}$[Flow]	β
τ	0.8	960	0.29	149.84	2.85	82.24	64.75	0.03	78.03	4.21	-18.78	1.00
	0.9	960	0.19	115.10	5.07	67.45	42.58	0.05	63.24	4.21	-8.39	1.00
	1.0	960	0.12	102.37	17.09	57.05	28.22	0.18	52.84	4.21	-0.18	1.00
	1.1	960	0.28	174.24	30.14	83.76	60.33	0.33	79.56	4.21	6.49	1.00
	1.2	960	0.41	227.72	35.08	104.94	87.70	0.38	100.73	4.21	12.12	0.99
η	0	960	0.26	163.97	20.66	86.25	57.05	0.22	82.04	4.21	-1.78	1.00
	0.2	960	0.26	159.21	19.47	82.91	56.83	0.21	78.70	4.21	-1.76	1.00
	0.4	960	0.26	154.19	18.19	79.34	56.66	0.20	75.13	4.21	-1.75	1.00
	0.6	960	0.26	148.86	16.75	75.57	56.54	0.18	71.36	4.21	-1.73	1.00
	0.8	960	0.26	143.05	15.17	71.38	56.50	0.16	67.17	4.21	-1.72	1.00
Λ	50	1,600	0.21	128.40	15.82	66.22	46.36	0.17	62.01	4.21	-1.21	1.00
	75	1,600	0.26	155.23	18.20	79.82	57.21	0.20	75.61	4.21	-1.75	1.00
	100	1,600	0.31	177.94	20.13	91.23	66.58	0.22	87.03	4.21	-2.29	1.00
ρ_I	0.6	1,200	0.26	150.30	17.97	75.58	56.75	0.19	74.08	1.50	-1.75	1.00
	0.7	1,200	0.26	151.45	17.99	76.71	56.75	0.19	74.38	2.33	-1.75	1.00
	0.8	1,200	0.26	153.68	18.03	78.91	56.75	0.19	74.91	4.00	-1.75	1.00
	0.9	1,200	0.26	159.99	18.21	85.16	56.62	0.20	76.16	9.00	-1.75	1.00
p	50	1,200	0.23	139.15	19.90	68.67	50.58	0.40	64.46	4.21	-1.95	0.99
	100	1,200	0.26	152.32	18.15	78.01	56.16	0.18	73.80	4.21	-1.73	1.00
	150	1,200	0.27	159.52	17.33	83.11	59.08	0.12	78.90	4.21	-1.67	1.00
	200	1,200	0.28	164.43	16.82	86.57	61.04	0.08	82.36	4.21	-1.64	1.00
c	100	1,200	0.39	113.97	12.74	62.64	38.59	0.14	58.43	4.21	-1.69	1.00
	200	1,200	0.26	145.10	16.96	75.99	52.15	0.18	71.78	4.21	-1.74	1.00
	300	1,200	0.21	168.40	19.98	85.18	63.24	0.22	80.98	4.21	-1.77	1.00
	400	1,200	0.18	187.95	22.52	92.54	72.89	0.24	88.33	4.21	-1.80	1.00
All		4,800	0.26	153.85	18.05	79.09	56.72	0.19	74.88	4.21	-1.75	1.00

Table 2: Results of numerical study with parameters from Table 1

6.1. Ratio of Demand to Returns

The intuition that a system in which demand and returns are more balanced has lower costs is reflected in the numerical study results. The restaurant indeed minimizes its average total costs when demand and returns are balanced, i.e., when the demand-to-returns ratio $\tau = 1.0$. We illustrate how the different cost components vary with τ in Figure 2 and use this figure to explain the mechanism behind this result.

As the restaurant's demand-to-return ratio decreases from $\tau = 1.0$ to $\tau = 0.8$ (i.e., moving from the center to the left of Figure 2), returns are higher than demand. Although the restaurant has fewer lost sales as can be seen in the decrease in $\mathbb{E}$[Loss] in Table 2, the restaurant's inventory of containers (in particular clean containers $\mathbb{E}[X_C]$) increases and the holding cost rate g_H^* increases. This higher holding cost rate results in an overall increase in total cost rate. When the supplier visits, the restaurant uses this opportunity to offload clean containers to the supplier as can be seen by the increasingly negative expected flow $\mathbb{E}$[Flow] between the restaurant and the supplier.

In the opposite direction from the center, as τ increases from $\tau = 1.0$ to $\tau = 1.2$, the restaurant faces more demand than returns. The restaurant is penalized from both the underage and holding costs perspective. Not only do lost sales increase because of the additional demand but holding cost also increases as the restaurant stocks more to avoid the high underage penalty. The supplier's visit is an opportunity for the restaurant to receive an inflow of clean containers, as can be seen from the increasing positive expected flow.

The restaurant also benefits from a more balanced demand-to-returns ratio in one more way: a balanced

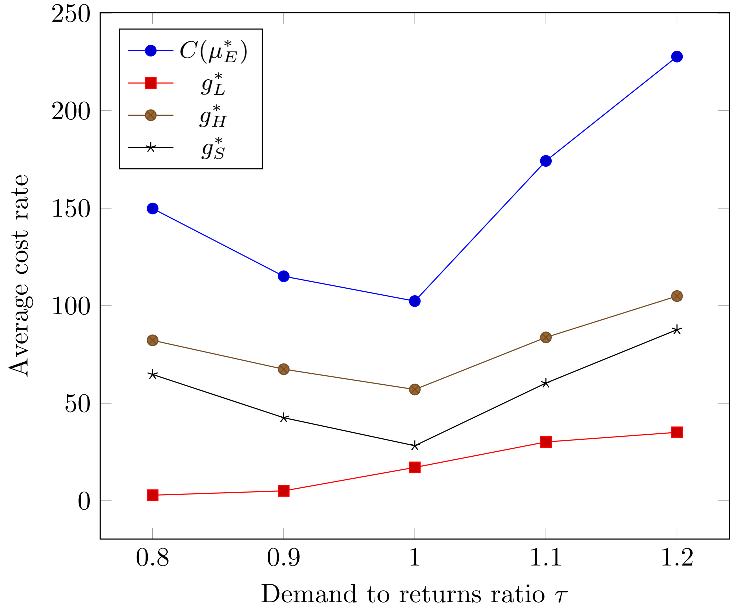

Figure 2: Average cost rate as a function of the ratio of demand-to-returns ratio τ.

demand-to-return ratio allows the restaurant to reduce the supplier visit costs. As τ moves away from 1.0 in either direction, the optimal supplier visit frequency μ_E^* increases. The restaurant is less able to manage its inventory in a cost efficient manner internally – that is, from its demand and returns alone – and is more dependent on the supplier for rebalancing. For the supplier operating the container system, it is crucial that this service is designed in a way that restaurants can rely on the rebalancing, especially when demand and returns are not balanced.

In practice, it may be difficult for the restaurant to have balanced demand and returns. If there is an imbalance between demand and returns, however, it is preferable for the restaurant to have more returns than demand instead of more demand than returns. This observation is likely a result of the relatively high underage penalty, which is in line with reality.

Notice that, even in a balanced system, there is an average net outflow of clean containers from the restaurant to the supplier when the supplier visits as can be seen from the negative $\mathbb{E}[\text{Flow}]$. This observation is a consequence of the fact that unsatisfied demand is lost. The supplier, despite serving the role of an inventory balancer in the system, is not able to balance demands and returns. All returns are collected but not all demand is fulfilled, so that, in steady state, fulfilled demand is lower than returns and the restaurant generates a net outflow of containers to the supplier.

6.2. Proportion of Coupled Demand and Returns out of Total Demand

As can be seen from Table 2, the more demand and returns are coupled (i.e., the higher η), the easier it is for the restaurant to balance its inventory levels locally. Both lost sales and holding costs decrease, resulting in an overall decrease in the average total cost rate. The restaurant mostly uses the supplier visits to reduce its inventory of clean containers, but it does not reduce this outflow by much as the proportion of coupled demand increases.

The fact that μ_E^* is not more sensitive to changes in the proportion of coupled demand is a more unexpected result. The restaurant's improved ability to sustain its operations internally when more of its demands are coupled with returns would suggest that the restaurant does not need the supplier as much for rebalancing, leading to a lower optimal supplier visit frequency. The results in Table 2 support this intuition, but the decrease in μ_E^* as η increases is modest.

One plausible explanation is that the benefit of an increased proportion of coupled demand and returns depends on the overall balance of demand to returns at the restaurant. To investigate this relationship, we examine the restaurant's performance metrics for each level of the demand-to-returns ratio τ as the proportion of coupled demand η increases. The results are displayed in Table 3.

τ	η	Count	μ_E^*	$C(\mu_E^*)$	g_L^*	g_H^*	g_S^*	$\mathbb{E}$[Loss]	$\mathbb{E}[X_C]$	$\mathbb{E}[X_D]$	$\mathbb{E}$[Flow]	β
	0	192	0.29	153.7	3.89	85.31	64.50	0.04	81.1	4.21	-18.79	1.00
	0.2	192	0.29	151.52	3.29	83.59	64.64	0.03	79.38	4.21	-18.78	1.00
0.8	0.4	192	0.29	149.53	2.74	82.01	64.78	0.03	77.8	4.21	-18.78	1.00
	0.6	192	0.29	147.84	2.3	80.66	64.89	0.02	76.45	4.21	-18.77	1.00
	0.8	192	0.29	146.61	2.04	79.64	64.92	0.02	75.43	4.21	-18.77	1.00
	0	192	0.19	123.48	7.40	73.62	42.46	0.08	69.41	4.21	-8.41	1.00
	0.2	192	0.19	119.11	6.2	70.47	42.44	0.06	66.26	4.21	-8.40	1.00
0.9	0.4	192	0.19	114.81	4.98	67.30	42.53	0.05	63.09	4.21	-8.38	1.00
	0.6	192	0.19	110.76	3.84	64.27	42.66	0.04	60.06	4.21	-8.37	1.00
	0.8	192	0.19	107.36	2.96	61.60	42.80	0.03	57.39	4.21	-8.36	1.00
	0	192	0.14	121.99	22.09	69.98	29.93	0.23	65.77	4.21	-0.23	1.00
	0.2	192	0.13	113.6	20.05	64.51	29.03	0.21	60.31	4.21	-0.21	1.00
1.0	0.4	192	0.12	104.12	17.7	58.23	28.19	0.18	54.02	4.21	-0.18	1.00
	0.6	192	0.12	93.03	14.76	50.92	27.35	0.15	46.71	4.21	-0.15	1.00
	0.8	192	0.11	79.11	10.87	41.62	26.61	0.11	37.42	4.21	-0.11	1.00
	0	192	0.29	185.31	32.88	91.65	60.78	0.36	87.44	4.21	6.46	0.99
	0.2	192	0.28	180.14	31.7	87.97	60.47	0.35	83.76	4.21	6.47	1.00
1.1	0.4	192	0.28	174.64	30.33	84.12	60.19	0.33	79.91	4.21	6.48	1.00
	0.6	192	0.28	168.75	28.79	79.89	60.08	0.32	75.68	4.21	6.50	1.00
	0.8	192	0.28	162.37	27.02	75.19	60.15	0.3	70.98	4.21	6.52	1.00
	0	192	0.41	235.34	37.06	110.69	87.59	0.41	106.49	4.21	12.09	0.99
	0.2	192	0.41	231.67	36.13	107.98	87.55	0.4	103.78	4.21	12.10	0.99
1.2	0.4	192	0.41	227.87	35.21	105.05	87.61	0.39	100.84	4.21	12.11	0.99
	0.6	192	0.41	223.92	34.06	102.12	87.74	0.37	97.92	4.21	12.13	0.99
	0.8	192	0.41	219.8	32.96	98.84	88.00	0.36	94.63	4.21	12.14	0.99
All		4800	0.26	153.85	18.05	79.09	56.72	0.19	74.88	4.21	-1.75	1.00

Table 3: Effect of increasing the proportion of coupled demand η for different values of the demand-to-returns ratio τ.

The largest benefit of increased coupling of demand and returns is realized when demand and returns are balanced, i.e., $\tau = 1.0$. In this situation, increased coupling of demand and returns enables the restaurant to be more internally sustainable and to reduce the supplier visit frequency significantly. The restaurant needs to hold very little inventory of clean containers when most customers bring the same number of dirty containers as they demand when ordering take-away or delivery meals. In fact, the restaurant only needs to hold enough clean containers to bridge the time required to clean returned containers while serving customers with the clean ones it has.

The benefit of increased coupling of demand and returns decreases as the ratio of demand to returns becomes more imbalanced in either direction. Because of the imbalance of overall demand to returns, the restaurant holds higher levels of reusable containers. In this case, the benefit of supply and demand being coupled is less pronounced.

Another observation from Table 3 is that, even when controlling for differences in the effect of coupled demand and returns at various levels of the demand-to-returns ratio, μ_E^* remains relatively stable, except for when overall demand and returns are balanced. This observation suggests that, for systems that are not balanced, μ_E^* is robust to changes in coupled demand and returns and most of the adjustment occurs in the optimal inventory balancing policy. Therefore, the value of optimizing μ_E as the coupling of demand and returns varies is limited when demand and returns are imbalanced. The managerial insight is that a one-time optimization of μ_E for a given τ may suffice.

6.3. Scale of Demand and Supplier Visit Cost

More frequent supplier visits are required to rebalance a system with a higher scale of demand Λ. The reason for the more frequent rebalancing is that, as the restaurant's demand increases, the restaurant stocks more containers to meet the additional demand but it also incurs more lost sales, making it harder for the restaurant to balance its inventory through its local demand and returns.

As the scale of demand increases, all costs increase, but they increase by less compared to the increase in scale. This observation suggests that larger restaurants may benefit from economies of scale. One of the key decisions in designing a reusable container system is the cost per supplier visit. A higher supplier visit cost, c, decreases the optimal supplier visit frequency, μ_E^*. With less frequent supplier visits, lost sales increase and the restaurant holds more inventory on average, either because it cannot offload excess inventory as frequently or because it carries more, in light of the less frequent visits to satisfy demand.

Whether a higher supplier visit cost affects restaurants of different scales equally is not immediately clear. Larger-scale restaurants incur higher supplier visit costs from more frequent supplier visits and higher underage and holding costs, but they can also divide these costs among more customers. To better understand the effect of scale on costs, we examine the average cost per demand fulfilled (i.e., per meal served in a reusable container), which we compute as $C(\mu_E^*)/(\lambda_1 + \lambda_2)(1 - \mathbb{P}(X_C = 0))$. Figure 3 shows the average cost/demand fulfilled as a function of scale for every level of supplier visit cost c. The difference between the average cost/demand fulfilled when the supplier visit cost $c = 400$ (green curve) compared to when $c = 100$ (blue curve) is greater for a smaller scale restaurant than for a larger scale restaurant. More specifically, the increase in demand at a larger restaurant is able to offset the increase in the supplier visit cost by more than the increase in demand at a smaller restaurant. This observation suggests that smaller restaurants are penalized more with a higher supplier visit cost and highlights the importance of the choice of the supplier visit c and its potential impact on the decision of smaller restaurants to participate in a reusable container system.

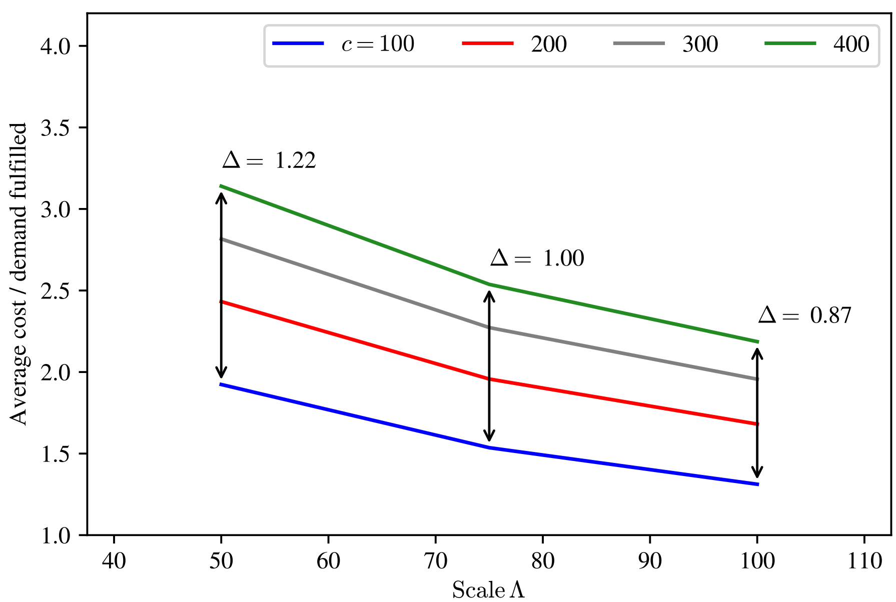

Figure 3: Scale effects on average cost/demand fulfilled for a given supplier visit cost c.

6.4. Utilization of Dishwasher

The dishwasher utilization has a negligible effect on the cost of the system. Since X_D behaves as the number of customers in an M/M/1 queue regardless of the rebalancing policy, the holding cost of dirty containers is a sunk cost that is not affected by the choice of y or μ_E. However, as the dishwasher utilization increases, the number of dirty containers $\mathbb{E}[X_D]$ increases, too, and thus the holding cost of dirty containers increases. The buildup of dirty containers is a result of the dishwasher becoming busier and not being able to take in as many dirty containers as arrive.

The number of clean containers, on the other hand, is not significantly affected by the utilization of the dishwasher. This result is consistent with Burke's theorem (Burke, 1956) which says that the departure process from an M/M/1 queue is a Poisson Process. Hence the arrival process of clean containers from the dishwasher is a Poisson process with rate $\lambda_2 + \lambda_3$ in steady state regardless of the utilization.

6.5. Underage Penalty

A higher underage penalty p results in a higher average total cost rate. To avoid lost sales, both a higher supplier visit frequency and a higher inventory level are optimal. The increase in the average total cost rate is driven by these higher supplier visit and holding costs.

Policy-makers can influence the level of the underage penalty, p, through regulations. For example, in jurisdictions that ban single-use containers altogether, the underage penalty for stocking-out on a reusable container is high, as it means that the restaurant is not able to sell the food item. In jurisdictions that charge a customer a surcharge for the use of a disposable container, the underage penalty for the restaurant is likely much lower. The restaurant's decision to hold higher inventory levels of reusable containers when p increases in response to such regulations shows that these regulations can help promote the availability of reusable containers in restaurants. Our model can be extended to study the sensitivity of the number of single-use containers displaced as a function of the level of the underage penalty. Such an analysis can inform policy-makers in terms of realistic target-setting for reductions of single-packaging waste from the food services sector.

7. Extensions

In the previous sections, we formulated and solved a model that captures the most essential building blocks of the inventory management problem of a restaurant that serves food in reusable containers. Various simplifying assumptions were made in this base model to study some of the unique features of this system and keep the model simple and transparent. The base model, however, can be adjusted and extended in ways that are more aligned with practice. In this section, we discuss two ways in which some of the assumptions can be relaxed. We also show that our main insight that demand coupling is always beneficial and that it is more important for the restaurant to have a balanced demand-to-supply ratio is robust with respect to the relaxing of our assumptions.

7.1. Generalized Dishwashing Process

In the base model, the time that it takes to wash one container is exponentially distributed and the dishwasher only washes a single container at a time. First, the processing times of the dishwasher can be represented by more suitable distributions than the exponential distribution (including distributions with less variation). The dishwashing model is intended to be a simple model that captures the effects of congestion due to finite processing capacity. While one could model the processing time of the dishwasher with a more general distribution, congestion scales in fundamentally the same way for $M/M/1$ queues as for $M/G/1$ queues, i.e., it scales as $\rho_I/(1-\rho_I)$ Gross et al. (2008). Our results demonstrate that the dishwasher utilization ρ_I is not a substantial driver of the costs of the system. As such, all other queueing characteristics (including processing times with a smaller coefficient of variation) will have a significantly smaller effect.

Second, in practice, many restaurants have dishwashers that wash in batches. To extend our model to cover this scenario, we can use a model for the dishwasher in which it waits for a full batch of $B \in \mathbb{N}$ dirty containers to be present before it starts washing. Since each individual dirty container arrives according to a Poisson process, the inter-arrival time of a batch has an Erlang-B distribution. Thus, the number of dirty container batches in the restaurant behaves as the number of customers in an $E_B/M/1$ queue. The state of this modified MDP is the tuple (x_C, x_D, x_W) where $x_C \in \mathbb{N}_0$ denotes the inventory level of clean containers at the restaurant, $x_D \in \mathbb{N}_0$ denotes the inventory level of dirty containers that are waiting to be put in the dishwasher, and $x_W \in \{0, 1\}$ denotes the status of the dishwasher, where $x_W = 0$ means that the dishwasher is not running and $x_W = 1$ means that the dishwasher is running. The optimal Markovian re-supply policy π^* and optimal cost rate g^* satisfy the following Bellman optimality equations:

$$
\begin{aligned}
V(x_C, x_D, x_W) + g^* = h(x_C + x_D) &+ p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0) + \lambda_1 V((x_C - 1)^+, x_D, x_W) \\
&+ \lambda_2 \mathbb{I}(x_W = 0 \wedge x_D \geq B - 1) V((x_C - 1)^+, x_D + 1 - B, 1) \\
&+ \lambda_2 \mathbb{I}(x_W = 1 \vee x_D < B - 1) V((x_C - 1)^+, x_D + 1, x_W) \\
&+ \lambda_3 \mathbb{I}(x_W = 0 \wedge x_D \geq B - 1) V(x_C, x_D + 1 - B, 1) \\
&+ \lambda_3 \mathbb{I}(x_W = 1 \vee x_D < B - 1) V(x_C, x_D + 1, x_W) + \mu_I \mathbb{I}(x_W = 0) V(x_C, x_D, x_W) \\
&+ \mu_I \mathbb{I}(x_W = 1 \wedge x_D < B) V(x_C + B, x_D, 0) \\
&+ \mu_I \mathbb{I}(x_W = 1 \wedge x_D \geq B) V(x_C + B, x_D - B, 1) \\
&+ \mu_E \min_y V(y, x_D, x_W) \qquad \forall (x_C, x_D, x_W) \in \mathbb{N}_0^2 \times \{0, 1\}
\end{aligned}
$$

where $V(x_C, x_D, x_E)$ is the relative value function. In addition to the parameter settings from our base model, we add a parameter for batch size $B \in \{5, 10, 15\}$. As a result, we have a full factorial test bed of $1 \times 4 \times 4 \times 5 \times 5 \times 3 \times 4 \times 3 = 14,400$ instances. The results of our numerical study are summarized in Table 4.

Parameter	Value	Count	μ_E^*	$C(\mu_E^*)$	g_L^*	g_H^*	g_S^*	$\mathbb{E}$[Loss]	$\mathbb{E}[X_C]$	$\mathbb{E}[X_D]$	$\mathbb{E}[Flow]$	β
B	5	4800	0.26	178.07	19.82	102.25	56.00	0.21	88.83	13.42	-1.77	1.00
	10	4800	0.25	210.86	23.16	132.36	55.33	0.25	107.41	24.95	-1.80	1.00
	15	4800	0.26	243.40	26.68	161.43	55.29	0.28	124.95	36.48	-1.84	1.00
τ	0.8	2880	0.29	207.40	9.28	134.41	63.70	0.10	109.46	24.95	-18.85	1.00
	0.9	2880	0.19	179.30	12.63	124.57	42.10	0.13	99.62	24.95	-8.46	1.00
	1.0	2880	0.13	167.85	22.54	116.45	28.86	0.23	91.50	24.95	-0.23	1.00
	1.1	2880	0.27	226.08	33.76	134.35	57.97	0.37	109.40	24.95	6.45	0.99
	1.2	2880	0.40	273.26	37.90	150.29	85.06	0.41	125.34	24.95	12.09	0.99
η	0	2880	0.26	216.16	24.47	135.33	56.36	0.26	110.38	24.95	-1.81	1.00
	0.2	2880	0.26	213.55	23.87	133.74	55.93	0.26	108.79	24.95	-1.81	1.00
	0.4	2880	0.26	210.87	23.25	132.10	55.52	0.25	107.15	24.95	-1.80	1.00
	0.6	2880	0.25	208.10	22.60	130.37	55.12	0.24	105.42	24.95	-1.79	1.00
	0.8	2880	0.25	205.20	21.91	128.53	54.76	0.23	103.57	24.95	-1.79	1.00
Λ	50	4800	0.21	183.80	21.00	117.32	45.47	0.22	92.37	24.95	-1.26	1.00
	75	4800	0.26	212.35	23.39	132.97	55.99	0.25	108.02	24.95	-1.80	1.00
	100	4800	0.30	236.18	25.28	145.75	65.15	0.27	120.80	24.95	-2.34	1.00
ρ_I	0.6	3600	0.26	179.96	19.69	104.36	55.90	0.21	94.08	10.29	-1.76	1.00
	0.7	3600	0.25	191.70	21.06	115.34	55.30	0.23	100.64	14.70	-1.78	1.00
	0.8	3600	0.25	212.28	23.68	133.86	54.75	0.25	110.13	23.72	-1.81	1.00
	0.9	3600	0.26	259.17	28.46	174.50	56.21	0.30	123.40	51.10	-1.86	1.00
p	50	3600	0.23	191.96	25.02	117.52	49.42	0.50	92.57	24.95	-2.05	0.99
	100	3600	0.25	208.71	23.35	130.41	54.95	0.23	105.46	24.95	-1.79	1.00
	150	3600	0.27	218.01	22.52	137.59	57.90	0.15	112.64	24.95	-1.70	1.00
	200	3600	0.28	224.42	21.99	142.54	59.89	0.11	117.58	24.95	-1.66	1.00
c	100	3600	0.38	171.69	17.82	115.61	38.25	0.19	90.66	24.95	-1.74	1.00
	200	3600	0.25	202.29	22.17	129.13	50.99	0.24	104.17	24.95	-1.79	1.00
	300	3600	0.21	225.03	25.21	138.12	61.70	0.27	113.16	24.95	-1.82	1.00
	400	3600	0.18	244.10	27.68	145.21	71.21	0.30	120.25	24.95	-1.85	1.00
All		14,400	0.26	210.78	23.22	132.02	55.54	0.25	107.06	24.95	-1.80	1.00

Table 4: Results of the numerical study for the extended model with batch dishwashing.

Intuitively, one would expect the overall costs of the system to increase when operating a batch dishwashing process as opposed to a unit dishwashing process. In a batch processing setting, dirty containers build up until a full batch size is reached and the dishwasher does not run continuously but rather has a waiting time until it has a full batch. This increase in costs $C(\mu_E^*)$ is exactly what we see in the numerical results in Table 4 as the batch size B increases. The restaurant not only holds more clean containers $\mathbb{E}[X_C]$ but also more dirty containers $\mathbb{E}[X_C]$ than when the dishwasher does not clean in batches. Therefore, the holding cost rate g_H^* increases significantly.

Despite this increase in the number of containers held, because the demand process remains unchanged, we would still expect more demand coupling in the system to be beneficial and for the positive effects of demand coupling to be most pronounced when the demand-to-returns ratio is balanced. To check the robustness of our original finding with respect to the batch size, we show in Figure 4 the total average relevant cost rate $C(\mu_E^*)$ as a function of the demand coupling η for each level of the demand-to-returns ratio τ across batch sizes $B \in \{1, 5, 10, 15\}$. While we see the expected general increase in the total cost rate when the batch size increases, we observe that our finding that increased demand coupling reduces the total cost rate and is most helpful if the restaurant has a balanced demand-to-returns ratio τ is robust with respect to the batch size B. In other words, the base model with $B = 1$ already captures the essential dynamics of the system.

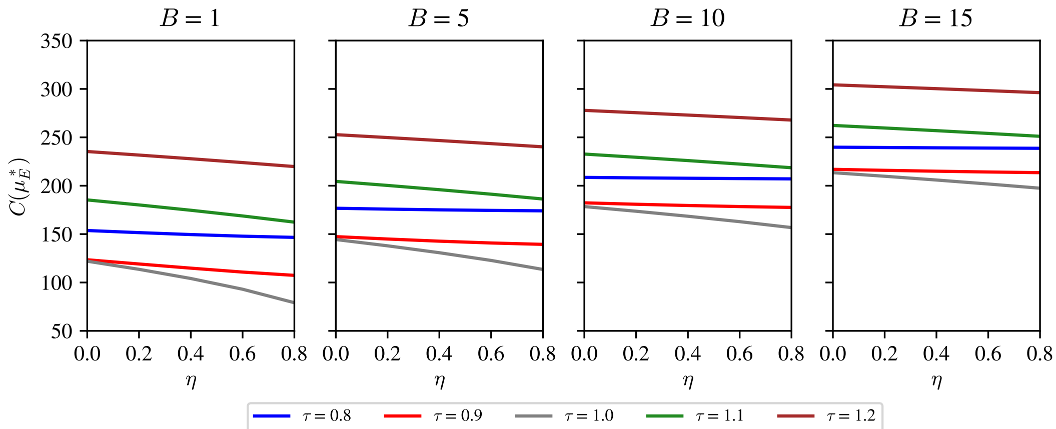

Figure 4: $C(\mu_E^*)$ for each level of the demand-to-returns ratio τ as demand coupling η increases for different batch sizes B. The case $B = 1$ is equivalent to the base model.

7.2. *Generalized Resupply Intervals*

The time between supplier visits has an exponential distribution in the base model, i.e., the supplier visit process is memoryless. In practice, it is reasonable to expect more regularity in the time between supplier visits. Within the Markovian framework, the base model can be extended to better reflect this by assuming that the times between supplier visits follow an Erlang distribution with $k \in \mathbb{N}$ phases. The parameter k determines the shape of the inter-arrival time distribution; $k = 1$ reduces the model to the base case and as $k \to \infty$ the inter-arrival times become more predictable, i.e., the coefficient of variation scales as $1/\sqrt{k}$.

In the extended model, the state of the MDP is the tuple (x_C, x_D, x_E) where $x_C, x_D \in \mathbb{N}_0$ denote the inventory levels of clean containers and dirty containers at the restaurant, respectively, and $x_E \in \{1, ..., k\}$ denotes the phase of the arrival process of the supplier. The optimal Markovian rebalancing policy π^* and optimal cost rate g^* satisfy the Bellman optimality equations:

$$
\begin{aligned}
V(x_C, x_D, x_E) + g^* = h(x_C + x_D) &+ p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0) + \lambda_1 V((x_C - 1)^+, x_D, x_E) \\
&+ \lambda_2 V((x_C - 1)^+, x_D + 1, x_E) + \lambda_3 V(x_C, x_D + 1, x_E) + \mu_I \mathbb{I}(x_D = 0) V(x_C, x_D, x_E) \\
&+ \mu_I \mathbb{I}(x_D > 0) V(x_C + 1, x_D - 1, x_E) + \mu_E \mathbb{I}(x_E < k) \min_y V(x_C, x_D, x_E + 1) \\
&+ \mu_E \mathbb{I}(x_E = k) \min_y V(y, x_D, 1) \qquad \forall (x_C, x_D, x_E) \in \mathbb{N}_0^2 \times \{1, \ldots, k\}
\end{aligned}
$$

where $V(x_C, x_D, x_E)$ is the relative value function.

To study this extended model, the parameter set of the base model is expanded to include the phase parameter k and we focus on the cases $k \in \{1, 2, 3\}$. As a result, we have a full factorial test bed of

$1 \times 4 \times 4 \times 5 \times 5 \times 3 \times 4 \times 3 = 14,400$ instances. The results of our numerical study for this extended model are summarized in Table 5.

Compared to the total cost rate in the base model (where $k = 1$), the total cost rate when when the supplier visits become more regular (where $k > 1$) is lower. This is in line with the intuition that less uncertainty about the time between supplier visits enables the restaurant to better plan its inventory levels. As a result, compared to the base model, the restaurant can carry less inventory and reduce the optimal supplier visit frequency μ_E^*, hence lowering its total cost rate.

However, in this extended model, we still find that our main insight about the beneficial effect of demand coupling on the total cost rate holds. As we did for the extended model with an increased dishwasher batch size, we check the robustness of our main finding by plotting in Figure 5 the total average relevant cost rate $C(\mu_E^*)$ as a function of the demand coupling η for each level of the demand-to-returns ratio τ for different numbers of phases $k \in \{1, 2, 3\}$. Figure 5 shows that, for every level of k, increased demand coupling reduces the total cost rate and its effect is most helpful when the restaurant has a balanced demand-to-returns ratio τ. We can thus conclude that already the simple and efficient base model, in which the inter-arrival times are exponentially distributed, yields the same insight as different shapes of supplier time distributions.

Parameter	Value	Count	μ_E^*	$C(\mu_E^*)$	g_L^*	g_H^*	g_S^*	$\mathbb{E}[\text{Loss}]$	$\mathbb{E}[X_C]$	$\mathbb{E}[X_D]$	$\mathbb{E}[Flow]$	β
k	1	4800	0.26	153.85	18.05	79.09	56.72	0.19	74.88	4.21	-1.75	1.00
	2	4800	0.23	139.18	25.93	63.6	49.66	0.27	59.39	4.21	-1.82	1.00
	3	4800	0.23	141.09	35.31	55.93	49.85	0.36	51.72	4.21	-1.91	1.00
τ	0.8	2880	0.26	136.12	2.10	76.10	57.91	0.02	71.89	4.21	-18.77	1.00
	0.9	2880	0.17	106.19	4.34	63.84	38.01	0.04	59.63	4.21	-8.38	1.00
	1.0	2880	0.10	97.27	21.31	54.35	21.62	0.21	50.14	4.21	-0.21	1.00
	1.1	2880	0.27	166.56	45.69	62.75	58.12	0.48	58.55	4.21	6.34	0.99
	1.2	2880	0.39	217.42	58.71	73.99	84.72	0.61	69.78	4.21	11.89	0.99
η	0	2880	0.25	154.93	29.31	72.94	52.69	0.30	68.73	4.21	-1.85	1.00
	0.2	2880	0.24	150.14	28.02	69.82	52.30	0.29	65.62	4.21	-1.84	1.00
	0.4	2880	0.24	145.09	26.56	66.47	52.07	0.27	62.26	4.21	-1.83	1.00
	0.6	2880	0.24	139.68	25.00	62.89	51.79	0.26	58.68	4.21	-1.81	1.00
	0.8	2880	0.24	133.71	23.26	58.92	51.53	0.24	54.71	4.21	-1.79	1.00
Λ	50	4800	0.20	120.67	22.22	56.25	42.21	0.23	52.04	4.21	-1.26	1.00
	75	4800	0.24	146.01	26.76	66.84	52.41	0.28	62.63	4.21	-1.83	1.00
	100	4800	0.28	167.45	30.32	75.53	61.61	0.31	71.32	4.21	-2.38	1.00
ρ_I	0.6	3600	0.24	141.16	26.37	62.65	52.14	0.27	61.15	1.50	-1.82	1.00
	0.7	3600	0.24	142.32	26.39	63.79	52.14	0.27	61.45	2.33	-1.83	1.00
	0.8	3600	0.24	144.54	26.39	65.99	52.15	0.27	61.99	4.00	-1.83	1.00
	0.9	3600	0.24	150.82	26.56	72.4	51.86	0.27	63.40	9.00	-1.83	1.00
p	50	3600	0.21	130.32	25.69	58.73	45.90	0.51	54.52	4.21	-2.07	0.99
	100	3600	0.24	143.16	26.37	65.47	51.33	0.26	61.26	4.21	-1.82	1.00
	150	3600	0.25	150.25	26.68	69.07	54.50	0.18	64.86	4.21	-1.73	1.00
	200	3600	0.26	155.12	26.99	71.56	56.57	0.13	67.35	4.21	-1.69	1.00
c	100	3600	0.36	107.97	18.80	52.76	36.41	0.19	48.55	4.21	-1.75	1.00
	200	3600	0.24	136.87	24.78	63.6	48.50	0.26	59.39	4.21	-1.81	1.00
	300	3600	0.19	158.19	29.24	71.2	57.75	0.30	66.99	4.21	-1.85	1.00
	400	3600	0.16	175.82	32.90	77.27	65.65	0.34	73.06	4.21	-1.89	1.00
All		14,400	0.24	144.71	26.43	66.21	52.08	0.27	62.00	4.21	-1.83	1.00

Table 5: Numerical study results

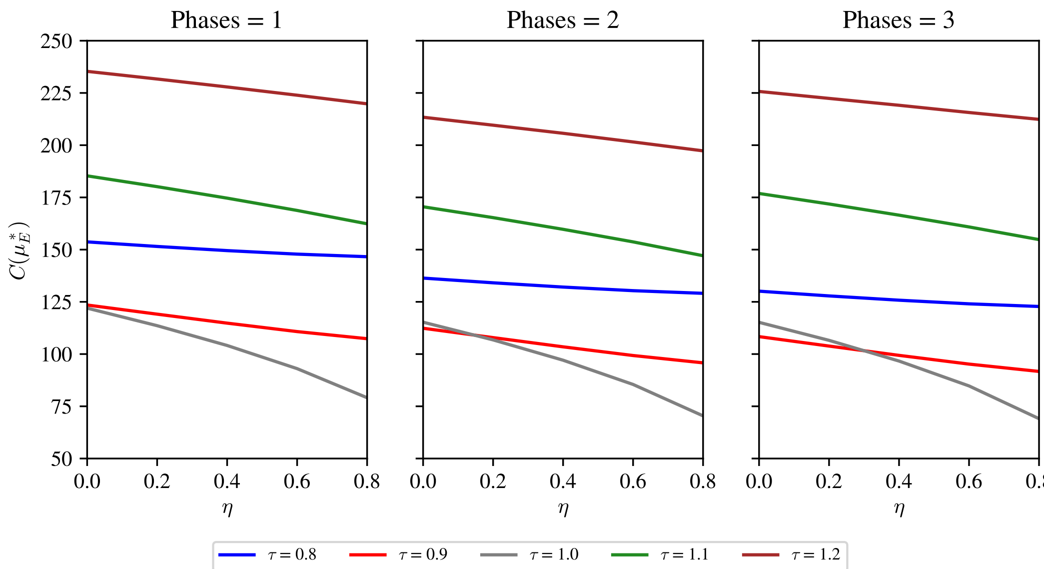

Figure 5: $C(\mu_E^*)$ for each level of demand-to-return ratio τ as demand coupling η increases for different numbers of phases k. The case $k = 1$ is equivalent to the base model.

8. Conclusion

The increased popularity and availability of food take-away and delivery options has contributed to a growth in single-use packaging waste. Policy-makers, start-ups, and restaurants are introducing reusable container systems to reduce the demand for and production of single-use packaging.

The design of reusable container systems is still evolving and many strategic and operational issues remain to be resolved. At the very core of these systems, individual restaurants need to decide on how many reusable containers to have on-hand and how frequently the supplier should visit to help balance inventory. The uncertainty in both demand and returns of reusable containers makes it difficult for the restaurant to control its inventory levels. Additionally, restaurants cater to multiple customer types that have different effects on the inventory level of reusable containers at the restaurant. Some customers only demand a clean container with their order (forward flow of products), other customers only return used containers (reverse flow of products), and some customers do both at the same time. This last set of customers have coupled demand and returns and, because they have a net zero effect on inventory levels, have the potential to balance the system. Our problem setting differs from related settings previously studied in that it is a lost sales system where traditional forward flow customers and reverse flow customers co-exist with customers with coupled demand and returns.

We find that the restaurant's optimal inventory rebalancing policy is one that depends only on the number of dirty containers at the restaurant when the supplier visits and not on the number of clean containers on hand. We also find that the restaurant can minimize costs by optimizing the supplier visit frequency, but the optimal supplier visit frequency is relatively insensitive to changes in many of the parameters. As a result, we suggest for the restaurant to optimize the visit frequency only when major changes affect the business and rather focus on the container inventory policy.

Having a greater proportion of customers with coupled demand and returns is always beneficial to a restaurant in that it allows the restaurant to better control the inventory of containers internally. The restaurant is less dependent on the supplier to balance its inventory and can reduce costs. However, just how beneficial depends on the overall balance of demand and returns at the restaurant. The cost reduction from increased coupling of demand and returns is most substantial when overall demand and returns are balanced. Such a balance is likely to emerge, for example, in a restaurant that has a base of loyal customers that frequently order from the restaurant. Our findings about the benefit of coupled demand are robust to changes in modeling assumptions, such as batch dishwashing and more regular supplier visits.

From a system-design perspective, our study highlights the importance of two decisions: (i) the supplier's choice of visit cost and (ii) the policy-maker's choice of the underage penalty cost. A higher supplier visit cost can disproportionately penalize restaurants with a smaller scale of demand for reusable containers, making it financially difficult for smaller restaurants to use reusable containers. Similarly, a higher underage penalty set by the policy-maker for not having a clean reusable container can play a role in increasing the amount of reusable containers in use for food take-away and delivery.

Many questions remain about how to design reusable container systems and how to incentivize customers, restaurants, and the third-party supplier to participate in the system. In this study, we put these questions to the side and focused on a single restaurant's inventory decisions to derive insights into how a restaurant can manage its inventory of reusable containers. However, for a reusable container system to be commercially viable, it must be beneficial to all participating parties. For example, a future analysis of the inventory and visit-frequency decisions for the reusable container supplier would highlight the key trade-offs faced by the supplier, in addition to those of the restaurant that we have highlighted in this paper, and how a supplier may operate this system.

Another natural extension of our work would be to analyze the operations of a network of restaurants offering reusable containers and the reusable container supplier's inventory decisions. It is possible, for instance, that demand or returns can be pooled in ways that reduce the costs for suppliers, restaurants, and even customers. As the business model for reusable container systems in the food services sector continues to evolve, we hope that our work can serve as a building block in the development of sustainable reusable container systems.

Appendix

Distribution of Demand Between Supplier Visits

Let S denote the time between supplier visits. Then S is exponentially distributed with mean $\frac{1}{\mu_E}$. Let D_S denote the demand between supplier visits. Observe that they demand in a fixed interval of length s has a Poisson distribution with mean $(\lambda_1 + \lambda_2)s$. The demand over an interval of random length S can be found by using the law of total probability:

$$\begin{aligned}
\mathbb{P}(D_S = x) &= \int_0^\infty \mathbb{P}(D_S = x | S = s) f_S(s) ds \\
&= \int_0^\infty e^{-\lambda s} \frac{(\lambda s)^x}{x!} \mu e^{-\mu s} ds \\
&= \frac{\mu}{x!} \int_0^\infty e^{-(\lambda+\mu)s} (\lambda s)^x ds \\
&= \frac{\mu}{x!} \left(\frac{\lambda}{\lambda+\mu}\right)^x \left(\frac{1}{\lambda+\mu}\right) \underbrace{\int_0^\infty z^x e^{-z} dz}_{\Gamma(x+1)=x!} \\
&= \frac{\mu}{\lambda+\mu} \left(\frac{\lambda}{\lambda+\mu}\right)^x \\
&= \left(1 - \frac{\lambda}{\lambda+\mu}\right) \left(\frac{\lambda}{\lambda+\mu}\right)^x
\end{aligned}$$

In the fourth line, we changed the variable of integration to $z = (\lambda + \mu)s$. The integral at the end of this line yields the Gamma function $\Gamma(x + 1)$. For integer x, $\Gamma(x + 1) = x!$, so $x!$ cancels out. Observe that D_S has a geometric distribution with parameter $\frac{\lambda_1+\lambda_2}{\lambda_1+\lambda_2+\mu_E}$.

Computation of Expected Lost Sales Rate

Let $U_n(x_C, x_D)$ be the expected lost sales rate after n events starting in state (x_C, x_D), where $U_n(x_C, x_D)$ is given by:

$$\begin{aligned}
U_n(x_C, x_D) = \mathbb{I}(x_C = 0)(\lambda_1 + \lambda_2) \\
&+ \lambda_1 \big(\mathbb{I}(x_C = 0) U_{n-1}(x_C, x_D) + \mathbb{I}(x_C > 0) \lambda_1 U_{n-1}(x_C - 1, x_D)\big) \\
&+ \lambda_2 \big(\mathbb{I}(x_C = 0) U_{n-1}(x_C, x_D + 1) + \mathbb{I}(x_C > 0) U_{n-1}(x_C - 1, x_D + 1)\big) \\
&+ \lambda_3 U_{n-1}(x_C, x_D + 1) \\
&+ \mu_I \big(\mathbb{I}(x_D = 0) U_{n-1}(x_C, x_D) + \mathbb{I}(x_D > 0) U_{n-1}(x_C + 1, x_D - 1)\big) \\
&+ \mu_E U_{n-1}(y_n^*(x_D), x_D),
\end{aligned}$$

where $y_n^*(x_D) \in \operatorname{argmin}_{y \in \mathbb{N}_0} V_n(y, x_D)$ and V_n is defined recursively as

$$\begin{aligned} V_n(x_C, x_D) = h(x_C + x_D) + p(\lambda_1 + \lambda_2)\mathbb{I}(x_C = 0) + \lambda_1 V_{n-1}((x_C - 1)^+, x_D) + \lambda_2 V_{n-1}((x_C - 1)^+, x_D + 1) \\ + \lambda_3 V_{n-1}(x_C, x_D + 1) + \mu_I \mathbb{I}(x_D = 0) V_{n-1}(x_C, x_D) + \mu_I \mathbb{I}(x_D > 0) V_{n-1}(x_C + 1, x_D - 1) \\ + \mu_E \min_y V_{n-1}(y, x_D) \qquad \forall (x_C, x_D) \in \mathbb{N}_0^2. \end{aligned}$$

For an optimal policy, the expected lost sales rate is:

$$\mathbb{E}[\text{Loss}] = \lim_{n \to \infty} U_{n-1}(x_C, x_D) - U_n(x_C, X_D)$$

for any $(x_C, x_D) \in \mathbb{N}_0^2$.

References

Accorsi, R., Cholette, S., Manzini, R., & Mucci, L. (2022). Managing uncertain inventories, washing, and transportation of reusable containers in food retailer supply chains. *Sustainable Production and Consumption*, *31*, 331–345.

Ahuja, K., Chandra, V., Lord, V., & Peens, C. (2021). *Ordering in: The rapid evolution of food delivery*. Technical Report McKinsey & Company. URL: https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ordering-in-the-rapid-evolution-of-food-delivery#.

Burke, P. J. (1956). The Output of a Queuing System. *Operations Research*, *4*, 699–704.

Chen, M., Hu, M., & Wang, J. (2022). Food Delivery Service and Restaurant: Friend or Foe? *Management Science*. Articles in Advance. doi:10.1287/mnsc.2021.4245.

City of Vancouver (2022). Green Vancouver: Single-Use Item Reduction Strategy. URL: https://vancouver.ca/green-vancouver/single-use-items.aspx.

Datner, S., Raviv, T., Tzur, M., & Chemla, D. (2019). Setting Inventory Levels in a Bike Sharing Network. *Transportation Science*, *53*, 62–76.

DeCroix, G., Song, J.-S., & Zipkin, P. (2005). A series system with returns: Stationary analysis. *Operations Research*, *53*, 350–362.

DeCroix, G. A. (2006). Optimal policy for a multiechelon inventory system with remanufacturing. *Operations Research*, *54*, 532–543.

Feldman, P., Frazelle, A. E., & Swinney, R. (2022). Managing Relationships Between Restaurants and Food Delivery Platforms: Conflict, Contracts, and Coordination. *Management Science*. Articles in Advance. doi:10.1287/mnsc.2022.4390.

Fleischmann, M., Bloemhof-Ruwaard, J. M., Dekker, R., Van der Laan, E., Van Nunen, J. A., & Van Wassenhove, L. N. (1997). Quantitative models for reverse logistics: A review. *European Journal of Operational Research*, *103*, 1–17.

Germany Federal Ministry for the Environment (2021). Mehrweg wird möglich im To-Go-Bereich. URL: https://www.bmuv.de/pressemitteilung/mehrweg-wird-moeglich-im-to-go-bereich/.

Glock, C. H. (2017). Decision support models for managing returnable transport items in supply chains: A systematic literature review. *International Journal of Production Economics*, *183*, 561–569.

Glock, C. H., & Kim, T. (2014). Container management in a single-vendor-multiple-buyer supply chain. *Logistics Research*, *7*, 1–16.

Gouvernement du Grand-Duché de Luxembourg (2022). Présentation du paquet de lois dit "Offall- a Ressourcëpak"[Press Release]. gouvernement.lu. URL: `https://gouvernement.lu/fr/actualites/toutes_actualites/communiques/2022/05-mai/17-welfring-offallpak.html`.

Gross, D., Shortle, J. F., Thompson, J. M., , & Harris, C. M. (2008). *Fundamentals of Queueing Theory* volume 399. John Wiley & Sons.

Guide, V. D. R., & Van Wassenhove, L. N. (2009). OR FORUM—The Evolution of Closed-Loop Supply Chain Research. *Operations Research*, *57*, 10–18.

Guide Jr, V. D. R., & Srivastava, R. (1997). Repairable inventory theory: Models and applications. *European Journal of Operational Research*, *102*, 1–20.

Kabra, A., Belavina, E., & Girotra, K. (2016). Bike-Share Systems. In *Environmentally Responsible Supply Chains* (pp. 127–142). Springer.

Kiesmüller, G. P., & Van der Laan, E. A. (2001). An inventory model with dependent product demands and returns. *International Journal of Production Economics*, *72*, 73–87.

Kiesmüller, G. (2003). A new approach for controlling a hybrid stochastic manufacturing/remanufacturing system with inventories and different leadtimes. *European Journal of Operational Research*, *147*, 62–71.

Van der Laan, E., Salomon, M., Dekker, R., & Van Wassenhove, L. (1999). Inventory Control in Hybrid Systems with Remanufacturing. *Management Science*, *45*, 733–747.

Lucas, A. (2022). Starbucks unveils new plans to eliminate single-use cups, encourage reusable mugs. CNBC. URL: `https://www.cnbc.com/2022/03/15/starbucks-unveils-new-plans-to-eliminate-single-use-cups-encourage-reusable-mugs.html`.

Mao, W., Ming, L., Rong, Y., Tang, C. S., & Zheng, H. (2022). On-Demand Meal Delivery Platforms: Operational Level Data and Research Opportunities. *Manufacturing & Service Operations Management*. Articles in Advance. doi:`10.1287/msom.2022.1112`.

Netherlands Chamber of Commerce (2022). Ban on plastic straws, plates, and other single-use plastics. business.gov.nl. URL: `https://business.gov.nl/running-your-business/environmental-impact/making-your-business-sustainable/ban-on-plastic-straws-plates-and-other-single-use-plastics/`.

Peters, A. (2020). At a growing number of coffee shops, getting a coffee to go means checking out a cup. Fast Company. URL: `https://www.fastcompany.com/90450872/at-a-growing-number-of-coffee-shops-getting-a-coffee-to-go-means-checking-out-a-cup`.

Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.

Raviv, T., & Kolka, O. (2013). Optimal Inventory Management of a Bike-Sharing Station. *IIE Transactions*, *45*, 1077–1093.

Schupak, A. (2021). 'The cusp of a reuse revolution': startups take the waste out of takeout. The Guardian. URL: `https://www.theguardian.com/us-news/2021/oct/22/takeout-delivery-food-waste-free-reusable-containers-startups`.

Shu, J., Chou, M. C., Liu, Q., Teo, C.-P., & Wang, I.-L. (2013). Models for Effective Deployment and Redistribution of Bicycles within Public Bicycle-Sharing Systems. *Operations Research*, *61*, 1346–1359.

Smithers, R. (2019). 'Lunch on the go' habit generates 11bn items of packaging waste a year. The Guardian. URL: https://www.theguardian.com/environment/2019/may/14/lunch-on-go-habit-generates-11bn-items-packaging-waste-year-uk.

Souza, G. C. (2013). Closed-Loop Supply Chains: A Critical Review and Future Research. *Decision Sciences*, *44*, 7–38.

Taheri, M., Bouchery, Y., Transchel, S., & Fransoo, J. (2021). Managing reusable packaging via a deposit system. *Working Paper*.

US Environmental Protection Agency (2015). *Reducing Wasted Food & Packaging: A Guide for Food Services and Restaurants*. Technical Report US Environmental Protection Agency. URL: https://www.epa.gov/sites/default/files/2015-08/documents/reducing_wasted_food_pkg_tool.pdf.

Wolff, R. W. (1982). Poisson Arrivals See Time Averages. *Operations Research*, *30*, 223–231.

Zero Waste France (2017). *McDonald's: une politique déchets à contre-courant de l'économie circulaire*. Technical Report Zero Waste France. URL: https://www.zerowastefrance.org/wp-content/uploads/2018/02/rapport-mc-do-web.pdf.