PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT
Vol. 23, No. 6, June 2014, pp. 951–964
ISSN 1059-1478 | EISSN 1937-5956 | 14 | 2306 | 0951
DOI 10.1111/poms.12108

Vehicle Procurement Policy for Humanitarian Development Programs

Mahyar Eftekhar
W. P. Carey School of Business, Department of Supply Chain Management, Arizona State University, Tempe, Arizona 85287-4706, USA
eftekhar@asu.edu

Andrea Masini
HEC Paris, Department of Operations Management and Information Technology, 1, rue de la Liberation, Jouy en Josas Cedex 78351, France, masini@hec.fr

Andreas Robotis
ALBA Graduate Business School, Operations Management Area, 6–8 Xenias St., Athens 11528, Greece, arobotis@alba.edu.gr

Luk N. Van Wassenhove
INSEAD, Technology and Operations Management Area, Boulevard de Constance, Fontainebleau Cedex 77305, France
luk.van-wassenhove@insead.edu

This article aims to identify optimal vehicle procurement policies for organizations engaged in humanitarian development programs and to derive general insights on the characteristics of these policies. Toward that end, we follow an inductive approach. First, we study the operations of the International Committee of the Red Cross (ICRC) in three representative countries: Sudan, Afghanistan, and Ethiopia. Using a linear programming (LP) model primed with field data provided by the ICRC, we calculate the optimal vehicle fleet size and compare it with the policies actually implemented. Second, drawing from results of the LP model, we develop a stylized quadratic control model and use it to characterize the general structure of the optimal policy under different demand scenarios and operational constraints. After demonstrating that the results of the control model are consistent with those of the LP model in the specific context analyzed, we discuss the optimal policies and the applicability of the former as a practical tool for strategic asset planning.

Key words: fleet management; humanitarian logistics; development programs; procurement
History: Received: January 2012; Accepted: February 2013 by Martin Starr, after 3 revisions.

1. Introduction

International humanitarian organizations (HOs) operate worldwide to run relief operations and development programs. Relief operations address emergency situations created by such disasters as famine, earthquakes, and floods, all of which require immediate response on a short-term basis. In contrast, development programs aim to provide long-term support and to improve living standards in poor countries and may run for years after a disaster (Van Wassenhove 2006).

Transportation plays a pivotal role in both types of programs. It is a vital element of the demand fulfillment process (Van Wassenhove 2006) and, after personnel, the second-greatest HO operating expense (Disparte 2007). Much of the transportation-related costs are associated with fleet management—that is, the acquisition, maintenance, use, and disposal of vehicles. It is not surprising that international HOs have begun to pay more attention to this issue, given the cost savings and performance improvements expected from optimized fleet management processes.

Decisions on fleet size are especially challenging for HOs because of the unusual environments in which they must operate (Pedraza Martinez et al. 2010). Security problems, poor infrastructure, and lack of reliable routes make vehicle usage patterns in HO contexts much different from those in commercial supply chains (Kovacs and Spens 2007, Tomasini and Van Wassenhove 2009). Security problems in conflict zones (Van Wassenhove 2006) affect the allocation of vehicles to missions. In some areas, humanitarian operators may be in danger if a vehicle breaks down; hence, only new vehicles can be used for field trips, while older vehicles must be used for administrative purposes in safer zones (Stapleton et al. 2008). Humanitarian organizations also face long procurement lead times. Operating in areas with poor infrastructure, they need specially equipped vehicles that must be ordered directly from manufacturers. Maintenance, too, is a challenge: in developing countries, auto dealers may not be trustworthy, outsourcing options are limited, and obtaining spare parts can be difficult. Finally, poor coordination and the lack of an adequate information technology (IT) infrastructure create extra costs and may lower performance even further (Pedraza Martinez et al. 2011).

The peculiar decision-making process of HOs creates additional challenges. Most HOs have a three-level structure: *headquarters* (HQ), which is usually established in a developed country, is responsible for strategic planning and budgeting decisions; *national delegations* coordinate activities within a country and are usually established in that country's capital city; and *subdelegations* are operational units, spread over many locations in the country, that run projects in the field. Although vehicles are deployed at the local level, fleet planning decisions are made by HQ at the aggregate level. In a centralized procurement procedure, such as the one used by the International Committee of the Red Cross (ICRC), subdelegations send requests to the national delegation, which in turn uses subdelegation data to derive an estimate of total demand that is submitted to HQ. Headquarters then uses these national estimates to decide how many vehicles should be purchased and shipped to delegations worldwide, often without having complete and accurate information upon which to make decisions.

Given these idiosyncratic characteristics, standard fleet management practices derived in commercial supply chains are unlikely to be easily applicable to HOs. Furthermore, an HO's unusual procurement process creates certain problems for decision makers. At the HQ level, central planners are responsible for making fleet size decisions on behalf of national delegations (and, by extension, their subdelegations). Yet, HQ seldom has the detailed demand and cost data needed to optimize decisions. Hence, fleet sizes are often established qualitatively by using simple, *ad hoc* heuristics. Central planners would therefore benefit from simple and parsimonious tools to guide vehicle procurement and allocation decisions—tools that use stylized aggregate demand information to infer the general structure of optimal policies. Of course, the validity of policies shaped by these tools should be compared with the detailed policies computable from detailed models when data are available.

Despite the relevance of fleet management from both an academic and a managerial perspective, the literature on aggregate-level fleet sizing in development programs is sparse. This article aims to fill that void. We seek to identify optimal vehicle procurement policies for HOs engaged in development programs and also to derive general insights on the characteristics of these policies under various operational conditions. The focus is on development programs because—despite representing a significant part of an HO's activities—they have received less attention than relief operations. We take an inductive approach in two steps. In the first step, we study ICRC operations in three representative countries (Sudan, Afghanistan, and Ethiopia) for which detailed data are available; given these data, we empirically estimate vehicle cost and capacity parameters. We next apply a linear programming (LP) model to calculate the optimal fleet size in each of these countries and then compare this size with the ICRC's implemented policy. In the second step, we draw upon the results of this LP model to develop a stylized quadratic control (QC) model that, while preserving the properties of the LP solution, is more parsimonious and has data requirements compatible with the data typically available at HO headquarters. We use this QC model to characterize the optimal policy's structure under different demand scenarios, and we discuss its applicability to strategic asset planning.

Our analysis yields several interesting insights. The LP model suggests that, in sharp contrast to the policies adopted by most HOs, the optimal fleet size remains relatively stable (under the specific demand scenario observed in the three sample countries) even when demand fluctuates. The QC model increases our understanding of the optimal policy's general structure and illustrates how that structure varies with demand characteristics, desired service levels, and the minimum and maximum thresholds for vehicle replacement.

The rest of the article is organized as follows. In section 2, we position our research with respect to the extant literature. Section 3 describes our research setting. Section 4 describes the LP model and its application to ICRC operations in Sudan, Ethiopia, and Afghanistan; section 5 presents the QC model and discusses the general properties of optimal policies. Finally, section 6 concludes the article, points out its limitations, and indicates some avenues for future research.

2. Literature Review

The humanitarian operations literature has studied fleet management in the context of both relief operations and development programs. The focus of this literature varies to reflect the different transportation requirements of the two undertakings (Pedraza Martinez et al. 2011): maximizing responsiveness and demand coverage for relief operations; and reducing costs and increasing fleet utilization for development programs.

Scholars studying fleet management for relief operations have dealt mostly with problems of victim evacuation and aid distribution, usually with the objective of minimizing response time. The victim evacuation literature has examined various trade-offs in a number of contexts, including optimal routing (Han et al. 2006), shelter locations (Sherali et al. 2006), scheduling helicopters (Barbarosoglu et al. 2002), and pre-positioning assets (Salmeron and Apte 2010). Research on aid distribution has focused on optimizing the delivery of aid to beneficiaries by minimizing travel time or maximizing demand coverage. A variety of contexts have been studied, including the dispatching of commodities (Yi and Ozdamar 2007), routing vehicles (Campbell et al. 2008), planning distribution (De Angelis et al. 2007), and optimizing facility location (Jia et al. 2007).

Fleet management has received comparatively less attention in development programs than in relief operations (Altay and Green 2006), even though the former make up a substantial part of any HO's operations. In development programs, scholars have mostly studied other themes: earmarked budgets (Besiou et al. 2012), vehicle reliability (McCoy 2013), or incentive mechanisms to guarantee data sharing between delegations and HQ (Pedraza Martinez et al. 2010). Of the two primary goals of fleet management cited by Pedraza Martinez et al. (2011), optimizing vehicle routing and optimizing fleet sizing, only the first one has been adequately studied, mostly using the same approaches adopted by the relief operations literature (Ingolfsson et al. 2008).

The few studies that deal with vehicle procurement and replacement in development programs have taken a "micro-level" perspective. That is, they have examined the replacement of each vehicle on an individual basis (Pedraza Martinez and Van Wassenhove 2013). Studies at the micro-level are useful to local managers in subdelegations, who handle a limited number of vehicles. However, such studies are less applicable at the HQ level, where central planners can decide on fleet sizes only at the aggregate level and are constrained by the lack of detailed data or by budget limitations. Given these limitations, central planners need macro-level models that can help them determine the optimal fleet size at the aggregate level over time (Vemuganti et al. 1989). Yet, we are not aware of any paper that addresses this problem for development programs.

As the literature on commercial supply chains has treated fleet sizing too, it is useful to assess whether insights from these studies can be applied to humanitarian development programs. Although this literature has predominantly taken a deterministic approach (see Vemuganti et al. 1989), some scholars have incorporated different sources of uncertainty to study a variety of problems such as repositioning empty vehicles in a network (Song and Earl 2008), determining the optimal number of outside carriers (Klincewicz et al. 1990), allocating trucks in a hub-and-spoke network (Du and Hall 1997), or determining optimal fleet size (List et al. 2003). These models have the clear merit of treating uncertainty. However, they are not easily applicable to humanitarian development programs primarily because they require detailed data on the variance of the demand process, which are nearly impossible to obtain in the humanitarian sector. Also, in some cases, these models consider solutions that are not easily implementable in the humanitarian sector (Klincewicz et al. 1990) or analyze operational contexts and sources of uncertainty that are not a primary concern for development programs (Du and Hall 1997, Song and Earl 2008). More importantly, they do not easily allow for the incorporation of certain identifying characteristics of humanitarian operations, such as age-dependent vehicle usage patterns and procurement and reselling constraints (List et al. 2003).

In light of these limitations, our article aims to contribute to the literature on humanitarian logistics by taking a deterministic approach and a macro-level perspective to the analysis of fleet sizing problems in development programs. We develop a LP model and a stylized QC model, both of which minimize total costs over time (subject to the typical operational constraints of a large HO). Our deterministic and macro-level approach is useful to derive general insights on a problem that is relatively novel in the humanitarian operations literature and to pave the way for the development of more sophisticated (but more data-intensive) models. It also enables our consideration of such constraints as meeting demand and not exceeding monthly budgets, constraints that do not apply at the individual vehicle level yet are important at the fleet level. In that respect, this article complements and extends the work of Pedraza Martinez and Van Wassenhove (2013) on individual vehicle replacement by determining how many vehicles are needed at the aggregate level over a given period of time, under a particular demand curve, and subject to certain budget constraints. Our work differs from most research on commercial fleet management in that it accounts for two identifying characteristics of humanitarian operations: the fact that a vehicle's usage decreases with age and the existence of replacement thresholds for individual vehicles.

In summary, our article contributes to the literature on humanitarian fleet management by (a) considering constraints that are specific to humanitarian operations at the subdelegation level, (b) incorporating constraints (e.g., budget limitations) related to vehicle procurement, and (c) constructing a model that preserves the dynamics imposed by these constraints and thereby helps us describe how those dynamics interact with demand parameters to determine which fleet management decisions are optimal at the aggregate level. Note also that, unlike many studies in this area, we posit a model that describes the *general* structure of optimal policies—in other words, irrespective of the empirical context in which they are implemented.

3. Research Setting

In order to ground the study, we analyzed the operations of the ICRC between 2000 and 2007 in Ethiopia, Sudan, and Afghanistan. We received two data sets from ICRC as part of the collaboration between the INSEAD Humanitarian Research Group and the ICRC Fleet Management Unit. This information was complemented by several interviews with representatives of the ICRC and other HOs that attended the 2011 Annual Fleet Forum in Geneva. Over the period of our analysis, ICRC had established 29 subdelegations in the countries we selected. The ICRC suggested that the three countries make for an ideal research field because each is highly representative of their operations and of a typical operating environment with regard to climate, geography, infrastructure, and mission. These are also the countries in which the ICRC had its largest fleets (on average, 265 vehicles in Sudan, 187 in Afghanistan, and 139 in Ethiopia). All subdelegations were equipped with the most frequently used vehicle for humanitarian missions, the 4 × 4 Toyota Land Cruiser.

Similar to what we observed at other HOs, the ICRC's vehicle procurement process is centralized. Subdelegations periodically provide estimates of their demand for transportation services to their national delegation, which aggregates the information and submits requests to HQ. Headquarters then uses these estimates in determining how many vehicles to purchase from the manufacturer and how they should be allocated to subdelegations.

Like other HOs, the ICRC is tied (with this vehicle's manufacturer) to a commercial agreement that dictates specific constraints related to procurement. For example, the ICRC purchases vehicles at below-market prices, but it is not allowed to resell any vehicle before it is 3 years old. The manufacturer may also impose minimum and maximum purchasing quantities. Combined with budget limitations, such requirements further constrain an HO's procurement process.

The most important characteristics of the ICRC fleets in our three sample countries are described in Table 1, which summarizes the two data sets we

Table 1 ICRC Fleets, 2000–2007: Descriptive Statistics

	Panel data set														
	Afghanistan (Obs. = 89)					Ethiopia (Obs. = 116)					Sudan (Obs. = 106)				
	Mean	Median	S.D.	Min.	Max.	Mean	Median	S.D.	Min.	Max.	Mean	Median	S.D.	Min.	Max.
Age (months)	35.48	35	17.15	1	89	35.08	32	21.92	1	113	21.20	20	12.18	1	63
Avg. distance traveled (km)	1,167.76	1,029	957.05	0	9,800	2,113.69	1,870	1,528.49	0	9,997	1,181.82	858	1,282.37	0	9,903
	Cross-sectional data set														
	Afghanistan (Obs. = 187; 101 active)					Ethiopia (Obs. = 139; 95 active)					Sudan (Obs. = 265; 189 active)				
	Mean	Median	S.D.	Min.	Max.	Mean	Median	S.D.	Min.	Max.	Mean	Median	S.D.	Min.	Max.
Age active (years)	4.59	5	1.75	1	8	4.61	5	2.41	1	14	3.47	3	1.41	1	10
Age sold (years)	7.70	8	1.95	4	12	7.75	8	1.92	2	14	7.12	8	2.82	1	12
Residual value (CHF)	7,197	6,510	2,809	1,287	13,536	16,555	16,486	6,305	6,349	27,612	12,470	11,000	7,305	3,030	45,139

received. The first, an unbalanced panel data set with monthly observations, covers all the vehicles used by the ICRC in Sudan, Afghanistan, and Ethiopia for the period 2000–2007. For each vehicle and time period (month), this data set reports the country of operation, the distance traveled, the vehicle's age, and the subdelegation that used the vehicle. This data set is an extension of the one used in Pedraza Martinez and Van Wassenhove (2013); it differs in that (a) it covers a longer time period and (b) it has been cleared of potential outliers with unrealistic odometer readings. The second data set that we received is a cross-sectional one covering the ICRC fleets for the same countries and time period. This data set includes information on the country of operation, vehicle age while in use and when it was sold, the subdelegation that used the vehicle, and its residual value.

Because some of the parameters we estimate for the LP model are time independent, we converted the first data set (which had a panel data structure) into a cross-sectional format by computing the averages (over time) of all the time-dependent variables. We then merged the two data sets, and we used the new, combined database to estimate the parameters of the LP model.

4. Linear Programming Model for Vehicle Fleet Sizing

4.1. Model Structure

We propose a LP model that optimizes vehicle fleet sizing at the national delegation level, where detailed data on vehicles and estimated demand are typically available. In the interest of space, we only discuss the general logic of the model, the methods adopted to estimate its input parameters, as well as the results of its application to the ICRC fleets in Afghanistan, Sudan, and Ethiopia (further details on the model formulation are available from the authors).

The decision maker's objective is to identify the number of vehicles of age a in period t that minimizes total fleet management cost over the decision horizon, subject to some operational constraints. Vehicle allocation decisions are revised monthly. In any period t (i.e., every month), the optimal fleet size is determined by choosing the number of new vehicles to be purchased and the number of vehicles of age a that should be sold. In accordance with studies that have addressed the same problem for commercial fleets (Vemuganti et al. 1989), we consider three major costs: purchasing cost, maintenance costs, and the residual value of vehicles at the end of their operational lives (this is an opportunity cost because HOs can recover some of that value by reselling the vehicles locally at a price r_a). We discount all cost functions using a 2% annual interest rate and express all budgets and cash flows in real (not nominal) terms. Since all our cost variables are country dependent, the model is effectively solved at the national delegation level.

The model includes the ICRC's actual operational constraints: a sales constraint accounting for the prohibition against ICRC reselling a vehicle in the first 36 months of its life; a service level constraint guaranteeing that the fleet planned usage (i.e., the maximum distance that vehicles in the fleet can travel in a given amount of time and in a given operating environment) is sufficient to meet estimated demand in each period; and a budget constraint B_t capturing the limited budgets of ICRC delegations for purchasing new vehicles and maintaining their fleets.

4.2. Parameter Estimation

4.2.1. Fleet Planned Usage. Unlike most papers on disaster management, which use synthetic data, we use field data to estimate the model parameters empirically. The proprietary data set obtained from the ICRC was used to calculate two variables directly: the planned usage of individual vehicles, d_a, and the total transportation demand during each period, D_t. The planned usage of individual vehicles was then used to compute total fleet planned usage and to estimate maintenance costs and the residual value of individual vehicles.

In the humanitarian sector a vehicle's planned usage, d_a (i.e., the maximum distance a vehicle can travel in a given period), decreases with a vehicle's age. For safety reasons, vehicles are not assigned to risky field missions after a specified age $\bar{a}$ (typically 24 months); vehicles older than $\bar{a}$ are used only for safer and shorter trips in urban areas (Stapleton et al. 2008). This policy of switching the use (mission type) of vehicles after a critical age threshold implies that the parameter d_a is best approximated by a two-step function:

$$d_a = \begin{cases} b_0^- + b_1^- a & \text{if } a \leq \bar{a}, \\ b_0^+ + b_1^+ a & \text{if } a > \bar{a}. \end{cases} \quad (1)$$

Although ICRC recommends switching mission types after 24 months, in practice the critical age threshold $\bar{a}$ depends on the country of operations. In order to estimate the coefficients b_1^- and b_1^+ and then determine the age threshold $\bar{a}$ at which switching takes place, we used two different approaches: a cross-sectional analysis and a panel data analysis.

In the former approach we estimated Equation (1) for each country separately using the cross-sectional database. The models were estimated for age thresholds $\bar{a}$ below and above the ICRC policy of $\bar{a} = 24$. For each run, we used a Chow test to check for equality between the coefficients b_1^- and b_1^+, and we retained the value of $\bar{a}$ with the highest significance level in the test.

Results are presented in Table 2. In the age interval $[0, \bar{a}]$, planned usage does not depend on vehicle age; but after the age threshold $\bar{a}$, planned usage decreases sharply with age. The estimated thresholds were 18 months for Afghanistan and 23 months for Sudan and Ethiopia.

We next used these results to calculate vehicle planned usage as follows. In the interval $(\bar{a}, A]$, we calculated usage via $d_a = \hat{b}_0 + \hat{b}_1^+ a$ (where $\hat{b}_0$ and $\hat{b}_1^+$ are, respectively, the estimated values of b_0 and b_1^+). In the interval $[0, \bar{a}]$, where d_a is independent of a, we followed the approach of Lapre et al. (2000) and calculated d_a as the maximum monthly distance traveled in the interval $[0, \bar{a}]$. These results were validated by reestimating the relationship between planned usage and vehicle age using a panel data approach with subdelegation fixed effects. The estimates of $\bar{a}$ in the panel data analysis were consistent with those from the cross-sectional approach.

4.2.2. Demand. In development programs, the demand for transportation services is relatively stable and predictable (Pedraza Martinez et al. 2010). Most of the demand fluctuations in these programs occur with significantly longer cycles than the time interval in our analysis and are not difficult to predict because they are due to such assignable causes as the HO increasing or decreasing its level of activity in a country. Furthermore, unpredictable demand variation in development programs is inherently lower than in relief operations because development programs are *long-term* and repeated endeavors for which fleet managers can generate relatively accurate forecasts. And even if some inherent variability is observed at the daily or weekly level, its magnitude is greatly attenuated at the monthly level because of pooling effects and because most journeys are non-critical and can easily be backlogged.

Subdelegations do not record demand for transportation services in each period, so we used the total distance traveled by the whole fleet over a given time period as a reasonable proxy for aggregated demand. Since using this proxy may create endogeneity problems due to the correlation between fleet size and number of journeys completed, we conducted additional tests using censored data models and *2SLS* estimation methods to assess the potential error induced by our approach. The results of the tests (available from the authors) provided no evidence of endogeneity.

4.2.3. Cost Functions and Budget. Purchasing cost data were obtained through interviews with the ICRC fleet manager. Over the period of our analysis, Toyota charged a fixed and constant price of CHF 28,000 for each vehicle purchased by ICRC (inclusive of vehicle shipment costs). Maintenance costs (which include both preventive maintenance costs and miscellaneous costs) and residual value were calculated using a method developed by Pedraza Martinez and Van Wassenhove (2013). Preventive maintenance costs are a function of the cost and the number of components that need to be replaced every period. The replacement schedule of components depends, in turn, on a vehicle's odometer reading (except for batteries, whose replacement schedule is based on age). Miscellaneous costs and residual value also depend on a vehicle's odometer reading. Since Pedraza Martinez and Van Wassenhove (2013) estimate those costs empirically in the same research setting as the one of this article, we could use their approach and estimate all costs using coefficients from their study. For further details on the estimation procedure we refer the interested reader to their article.

The ICRC did not disclose data on its actual budgets, so we took an alternative approach to estimating the parameter B_t. Following advices from the managers we interviewed, we used the actual fleet cost as a conservative estimate of B_t and computed nominal monthly budgets simply by dividing the total fleet cost for a given time period by the number of months in that period.

4.3. Results of the LP Model

The LP model was run for Afghanistan, Sudan, and Ethiopia separately. With this model, we seek to

Table 2 Individual Vehicle's Planned Usage: Cross-sectional Analysis

		Chow test results			Age interval regressions					
Country of operations	Knot	Fcrit.	F	Null hy.	$\hat{b}_0^-$, $\hat{b}_0^+$	$\hat{b}_1^-$, $\hat{b}_1^+$	p-value	t	Adj. R^2	Planned usage (d_a)
Afghanistan	$a \leq 18$	3.11	17.37	Rej.	1676.98	2.26	0.86	0.18	−0.10	2124.77
	$a > 18$					−11.57	0.00	−12.56	0.69	$1676.98 - 11.57a$
Ethiopia	$a \leq 23$	3.10	4.77	Rej.	—	2.44	0.86	0.18	−0.06	3281.15
	$a > 23$				3180.19	−24.76	0.00	−27.39	0.91	$3180.19 - 24.76a$
Sudan	$a \leq 23$	3.20	21.34	Rej.	—	−17.14	0.23	−1.26	0.03	2509.65
	$a > 23$				2175.62	−23.51	0.00	−5.89	0.52	$2175.62 - 23.51a$

obtain a "reference" optimal policy that could be used to generate insights for the more general control model; therefore, it was run for the entire time span over which we had data. The results are plotted in Figures 1 and 2, which compare (respectively) optimal with actual usage and optimal with actual fleet size.

Figure 2 suggests that the policy recommended by our LP model smooths out demand variations and keeps fleet size fluctuations to a minimum (i.e., it confirms that optimal fleet usage is less variable than demand). A series of *F*-tests confirmed that optimal fleet usage exhibits significantly lower variation than does demand. The null hypothesis of equality of variances was rejected for all three countries in our analysis with $F = 2.10$ ($p < 0.01$), $F = 1.56$ ($p < 0.05$) and $F = 1.49$ ($p < 0.05$) for, respectively, Afghanistan, Sudan, and Ethiopia.

To obtain further insights into the differences between the optimal policy and the actual policy, we compared the cost of the optimal policy identified by the LP model to the cost of the policy actually implemented for all three countries. Over the period of our analysis, using the optimal policy would have reduced costs by 7.9% in Afghanistan (from CHF 5.3 million to CHF 4.9 million), by 19.2% in Sudan (from CHF 10.3 million to CHF 8.3 million), and by 26.2% in Ethiopia (from CHF 3.8 million to CHF 2.8 million)—a cumulative cost savings exceeding CHF 3.7 million.

5. Quadratic Control Model for Determining Optimal Policies under General Demand Scenarios

5.1. Model Rationale

The LP model generates a detailed vehicle procurement policy and yields interesting insights. However, these results are valid only for the particular realization of the demand function observed in our three sample countries from 2000 to 2007. In order to derive optimal policies for different demand scenarios, decision makers at the HQ level would want to rerun the LP model for those scenarios; however this is both time consuming and data intensive—and thus less useful in the humanitarian context, where data are not readily available.

We offer an alternative and more practical approach to guiding HQ decisions: a stylized model that draws on the results of the LP to describe the structure of optimal policies not only in the specific case considered but also under different and more general demand scenarios. Because the governance as well as the procurement and maintenance policies of any subdelegation are imposed centrally by ICRC headquarters, they all face constraints that contribute to usage being less variable than demand. This fact allows us to generalize the results derived from our LP model and to use them to construct the QC model. This model can be used to analyze the optimal policy's general properties, to conduct validity checks on

Figure 1 Optimal Capacity vs. Demand (kilometers)

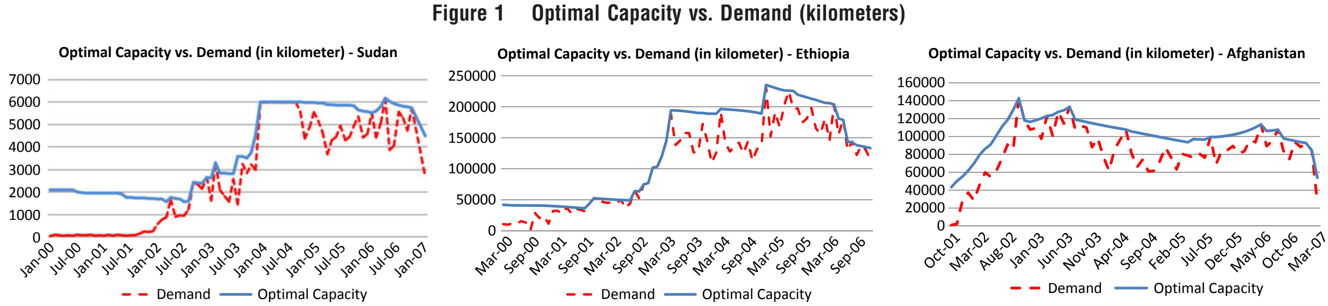

(a) Sudan **(b)** Ethiopia **(c)** Afghanistan

Figure 2 Optimal vs. Actual Fleet Size (number of vehicles)

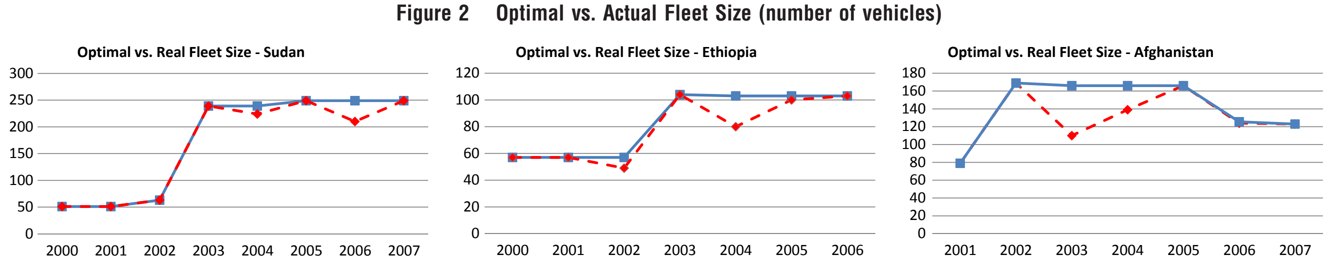

(a) Sudan **(b)** Ethiopia **(c)** Afghanistan the heuristics currently used by HOs, and to run sensitivity analyses that elucidate the impact of changes in input parameters on the optimal policy. Our model establishes links among budget dynamics, planned usage dynamics, and the operational context—relations that have not been examined in the literature.

We make two assumptions to develop the QC model. The first (based on results of the LP model) is that high fleet size variability is not optimal. We therefore construct a QC model that accounts—through the term $u^2(t)$, as explained in what follows—for such variability by penalizing the rate at which fleet size increases or decreases over time. The second assumption is that efficient operations ought to match supply with demand to a level specified by management. Our model allows the decision maker to choose this level (through the parameter q). Thus, the model penalizes time-averaged overstocking and understocking costs as well as procurement decisions that increase fleet size variability over the planning horizon.

This model is constructed to optimize fleet size at an aggregate level and works with data that are only approximate; for these reasons, we use the average number of vehicles (instead of the planned usage) to represent demand. Observe that the average demand per period in number of vehicles can be easily determined by dividing the total demand (in kilometers) by the average distance that a vehicle is driven per period. The model's inputs are demand requirements and operational constraints; its output is a fleet size time trajectory that satisfies those conditions while minimizing fleet size variation. In addition, the model allows the decision maker to decide how much of a penalty to impose for stocking costs and fleet size variability. In section 5.2, we introduce the model. In section 5.3, we use it to analyze the structure of optimal policies for three general demand cases. In section 5.4, we compare the optimal policy based on the QC model for a specific realization of demand with its counterpart based on the LP model.

5.2. Model Formulation

Equations (2)–(6) describe the objective function of the QC model and its constraints. $x(t)$ is a state variable representing the number of vehicles available in period t, and $u(t)$ is the control variable accounting for the number of vehicles either purchased or sold in period t. Note that $u(t)$ can be either positive (when vehicles are purchased) or negative (when they are sold). The rate at which fleet size changes is $\dot{x}(t)$ and it is the summation of the control variable $u(t)$ and the number of old vehicles (per period) that should be replaced owing to safety or maintenance concerns. This number is given by $x(t)/\tau$, where τ is the average time before a vehicle is replaced. We remark that the need to replace vehicles from time to time is another source of variability in fleet size. According to Pedraza Martinez and Van Wassenhove (2013), the individual replacement policy should be set as a function of a vehicle's odometer reading; hence, τ depends on the number of kilometers a vehicle is driven before it needs to be replaced. We estimate this number as $\frac{\text{optimal replacement odometer}}{\text{average monthly usage}}$. For example, if vehicles are optimally replaced at an odometer reading of 100,000 km and if vehicles average 2,000 km per month, then τ is 50 months. M is a constant that reflects maximum purchasing and selling quantities, whereas δ represents the least proportion of demand to be covered. Finally, $D(t)$ is the demand in period t (in number of vehicles).

Objective function:

$$\min_{u(t)} J = \frac{1}{2}\int_0^T (q[x(t) - D(t)]^2 + r[u(t)]^2)dt \tag{2}$$

subject to the following expressions.

State equation:

$$\dot{x}(t) = u(t) - \frac{x(t)}{\tau}. \tag{3}$$

Demand fulfillment constraint:

$$x(t) \geq \delta D(t). \tag{4}$$

Purchasing/selling constraint:

$$-M \leq u(t) \leq M. \tag{5}$$

Boundary condition:

$$x(0) = 0. \tag{6}$$

The term $r[u(t)]^2$ is used to penalize fleet size variation, and the term $q[x(t) - D(t)]^2$ is used to penalize the mismatch between fleet size and demand. Here, q and r are, respectively, the penalty cost due to fleet size–demand mismatches and the penalty cost due to purchasing or selling vehicles. Thus, if the ratio q/r is high, then the model forces fleet size to match demand; but if q/r is low, the model focuses on minimizing fleet size variation and allows for a greater mismatch between fleet size and demand. The value of q could be set by specifying a value for the cost of the average mismatch between fleet size and demand for a given time period. This parameter is a proxy for customer service and can be estimated empirically because it depends on the cost a subdelegation could incur for wasting perishable food products or for not being able to deliver drugs to beneficiaries due to an insufficient fleet size. In order both to simplify notation and to limit solution complexity, we assume that overstocking and understocking are equally expensive. (Note that in relief operations understocking can be more costly: it may result in a mission being canceled and, ultimately, in lives being lost.) The parameters τ and δ can also be used to account for differences between programs (health programs may require both shorter replacement intervals and higher service levels than do food programs).

Several scenarios can be represented by altering the parameters q and r, whose values can be set based on past experience. Constraint (3) is the state equation representing the fleet size variation in each period. The term $u(t)$ is the decision variable, which indicates how many vehicles should be added to (or removed from) the fleet; $x(t)/\tau$ is the number of vehicles that must be replaced in period t. Note that for a given number of missions, $x(t)/\tau$ is greater in countries that are (geographically) larger because their vehicles must cover more ground in a given period to accomplish the same missions. Note also that, since HOs cannot purchase partial vehicles, $u(t)$ is rounded to the nearest integer when the model is used in practice. Constraint (4) guarantees that, in any period, a portion δ of demand will be satisfied. The parameter δ is positive, and it can take values greater than 1 when some vehicles are used as a safety stock. Constraint (5) captures a limit that may be placed on the number of vehicles bought or sold each period. As in the LP case, it reflects annual order quotas imposed by vehicle manufacturers as a condition of offering discounts. The selling constraint reflects transaction costs, which arise because it is not economical for an HO to resell just a few used vehicles. To simplify notation, the maximum selling and purchasing levels are assumed to be equal. Finally, constraint (6) is the boundary condition. Without loss of generality, we assume that the initial number of vehicles is zero.

5.3. Analysis of Demand Scenarios

In this section, we use the QC model just developed to study the structure of optimal policies under three different demand scenarios: (i) a general demand function that consists of an increasing (or decreasing) term and of a term that varies around a constant; (ii) a demand function that consists only of an increasing (or decreasing) term; and (iii) a demand function that consists only of a term that varies around a constant. These functions reflect well the structure of the real demand observed in the three countries that we analyze in this study. Between 2000 and 2007, for example, demand in Sudan can be subdivided into three distinct patterns: constant from $t = 0$ to $t = 22$ (i.e., from 2000 to 2002), linearly increasing from $t = 23$ to $t = 33$ (2002–2003), and oscillating around a constant mean from $t = 34$ to $t = 84$ (2003–2007).

5.3.1. General Demand Function. We consider the general nonmonotonic demand function

$$D(t) = \alpha t + \beta \sin(\omega t), \tag{7}$$

which represents the most general case. The first term in (7) represents the increasing (or decreasing) trend; the second term represents the demand oscillation. The parameter ω is the frequency of the demand oscillation, and β is its magnitude. For development programs (unlike relief operations), we can assume that β is small because typically the demand for humanitarian services in these programs is strongly correlated with the size of the population in the affected area, which is relatively constant over time in the short to medium run. The term αt captures a common situation. The ICRC often starts operating in a country with a few small projects and then expands its operations over time; demand then follows an upward trend and varies with time. A symmetrical situation occurs (now with $\alpha < 0$) as the ICRC withdraws from a country.

Whereas relief operations encounter a lot of unpredictable variability, the demand variability in development programs is fairly cyclical. It can therefore be accurately and conveniently approximated by a sinusoidal function, which reflects well the seasonality of some demand drivers. For instance, organizations running food programs typically face demand peaks during or right after drought periods, which are seasonal events. Likewise, the demand for health programs increases in tropical regions after floods, which are also seasonal events. Our interviews with executive managers in large HOs (including the ICRC, the World Food Program, and World Vision International) confirmed that demand variation exists but that its magnitude β is usually a small percentage of the total demand. The model's demand parameters can also be used to accommodate differences among different program types—for example, the demand for food programs may oscillate more frequently but at a lower amplitude than does the demand for health programs.

Our first proposition follows from solving the QC model for the demand function (7). (All proofs are available upon request).

Proposition 1. *For the demand function $\alpha t + \beta \sin(\omega t)$, let $x(t)$ and $u(t)$ be the unconstrained solution of problem (2) given by*

$$x(t) = \frac{\alpha q t \tau^2}{q\tau^2 + r} + \frac{\beta q \sin(\omega t)}{q + r(1/\tau^2 + \omega^2)}, \tag{8}$$

$$u(t) = \frac{q\tau^2\alpha}{r + q\tau^2} + \frac{q\tau\alpha t}{r + q\tau^2} + \frac{\tau\beta(\alpha\omega\cos(\omega t) + q\sin(\omega t))}{r + \tau^2(r\omega^2 + q)}. \tag{9}$$

Then, in each period, the optimal solution will be as given by one of the following cases:

(a) *If $x(t) \geq \delta D(t)$ and $-M \leq u(t) \leq M$ (i.e., if neither constraint (4) or (5) is binding), then the optimal solution is $u^*(t) = u(t)$.*
(b) *If $x(t) \geq \delta D(t)$ and $u(t) > M$ or $u(t) < -M$ (i.e., if constraint (4) is not binding but constraint (5) is binding), then the optimal solution is $u(t)^* = M$ or $u(t)^* = -M$.*
(c) *If $x(t) < \delta D(t)$ and $-M \leq u(t) \leq M$ (i.e., if constraint (4) is binding but constraint (5) is not), then the optimal solution is $u(t)^* = \delta(\dot{D}(t) + D(t)/\tau)$.*
(d) *If $x(t) < \delta D(t)$ and $u(t) > M$ (i.e., if both constraints (4) and (5) are binding), then $u(t)^*$ is the minimum of the solutions of cases (b) and (c).*

Observe that the optimal policy is time dependent. Only one of the optimal solutions of Proposition 1 holds for any given period, but during that period any one solution can be optimal at different times. The results of this proposition also hold when there *is* an inventory of vehicles at time $t = 0$. In that case, $u(t)$ should be set to 0 in any period t before the initial fleet size reaches $D(t)$—after which, (9) is followed.

When neither (4) nor (5) is a binding constraint, the procurement policy $u^*(t)$ given by (9) consists of a constant term, a linearly increasing term, and a third term that oscillates with the same frequency as (but with a different phase from) the demand. Purchasing/selling decisions take place before demand increases or decreases: the numerator $\tau b(\alpha\omega \cos(\omega t) + q \sin(\omega t))$ of the third term increases before the demand term $b \sin(\omega t)$ increases. It is trivial to show that, as the ratio $\beta/\tau\omega$ decreases, the third term of (9) decreases and approaches zero. Unlike relief operations, development programs do not expect to experience a high β/ω ratio and so, as Proposition 1 shows, the dynamics of fleet size variation depend mostly on the vehicle replacement frequency $1/\tau$.

Depending on the parameter values, either one or both of constraints (4) and (5) may be binding. If (4) is binding and (5) is not, then the optimal policy is $u(t) = \delta(\dot{D}(t) + D(t)/\tau)$. If both constraints are binding, then the optimal policy is to increase fleet size quickly by setting $u(t) = M$ until there are enough vehicles to cover the demand. This result is in line with what we observe from the LP model. In that model, as the budget constraint narrows, the model calls for purchasing as many vehicles in each period as are allowed (by the available budget) until the fleet size can accommodate demand. It is interesting that, depending on the parameter values, the optimal fleet size strategy changes from one that levels fleet size by smoothing out demand requirements to one that "chases" demand (Slack et al. 2006). A combination of these strategies is required when the procurement rate $u(t)$ is constrained and demand requirements are high.

The results of Proposition 1 can be summarized as follows: *Depending on the portion of demand coverage and procurement constraints, there are three main regions occupied by the procurement policy for the demand function $at + \beta \sin(\omega t)$. In the first region, the procurement policy is given by (9); a portion of the demand is procured at $t = 0$, after which the procurement policy has oscillatory characteristics similar to those of demand. In the second region, the procurement policy is to match demand. In the third region, an early buildup of fleet size to meet future demand needs is optimal.*

In order to examine the link between budget requirements and operational environment, we consider the effects of ω and τ on the budget required for the demand function (7). Given that budget requirements may vary as a function of time and mission characteristics, we consider the upper bound of those requirements associated with completing a mission. Let h denote the cost of holding vehicles (maintenance and miscellaneous costs) and p the cost of purchasing a vehicle. Then the following equation describes the maximum budget available in any period:

$$\bar{B} = \max_{t \in [0,T]} (hx(t) + pu(t)). \tag{10}$$

Our next proposition examines the effects of τ and ω on the upper limit of the budget, $\bar{B}$.

Proposition 2. *Let $\omega_R = 1/\tau$ denote the frequency of vehicle replacement. If constraints (4) and (5) are not binding, then the upper limit of the budget requirements to complete a mission are affected by the frequency of demand oscillation as follows:*

$$\frac{d\bar{B}}{d\omega} \begin{cases} < 0 & \text{if } \frac{\omega}{\omega_R} > \frac{\sqrt{r(h^2 r\tau^2 + p^2(q\tau^2 + r))} - hr\tau}{pr\tau^2}; \\ > 0 & \text{if } \frac{\omega}{\omega_R} < \frac{\sqrt{r(h^2 r\tau^2 + p^2(q\tau^2 + r))} - hr\tau}{pr\tau^2}. \end{cases} \tag{11}$$

This proposition states that completing missions for which ω/ω_R is high requires a budget with a lower upper bound. The intuition is as follows. Recall from Proposition 1 that the optimal policy oscillates with the same frequency as the demand but with a different phase. In other words, the procurement of vehicles meant to replace old vehicles occurs during the same period in which new vehicles are procured to prepare for the next demand peak. This dynamic allows us to minimize the mismatch between fleet size and demand, as fleet size is augmented in a single period to compensate for both the increase in demand and the vehicle replacement. Increasing procurement levels in one period to compensate for both vehicle replacement and future demand reduces the time-averaged variability $u^2(t)$ of purchasing/selling decisions and consequently the fleet size–demand mismatch. Hence, overall costs decrease as well. This outcome may not prevail, however, if vehicles need to be replaced often and demand increases much later in the future. In that case, frequent purchasing/selling decisions are necessary to compensate for vehicle replacement. That being said, it may not be economical to purchase more vehicles (to satisfy future demand) and thereby incur high holding costs.

Proposition 2 reveals a link between budget requirements and such operational characteristics as demand oscillation and vehicle replacement policies. It also suggests that the upper budget limit in missions for which the frequency of demand peaks is greater than the frequency of vehicle replacement is *lower* than that limit in missions for which this inequality is reversed.

5.3.2. Increasing or Decreasing Demand. We now consider a simpler version of the demand function (7) in which we concentrate only on the increasing (decreasing) term αt. This case is representative when both the magnitude β and frequency ω of demand oscillation are small, as usually occurs when the ICRC must continuously expand its actions in a country. We then have the following proposition.

PROPOSITION 3. *For the demand function αt, the optimal fleet size is*

$$x^*(t) = \frac{\alpha q \tau^2}{q\tau^2 + r} t \tag{12}$$

and the optimal control is

$$u^*(t) = \frac{\alpha q \tau^2}{q\tau^2 + r}\left(1 + \frac{t}{\tau}\right). \tag{13}$$

This proposition yields some useful intuition. Observe that Equation (12) implies $q\tau^2/(q\tau^2 + r) \leq 1$; this term approaches unity as the vehicle replacement interval τ increases. Therefore, constraint (4) is always binding when $\delta > 1$—that is, when demand must always be satisfied. Also, note that constraint (4) will either bind for the whole demand interval $[0,T]$ (when $\delta > q\tau^2/(q\tau^2 + r)$) or it will not bind at all. This explains why the optimal policy is not time dependent.

For a linearly increasing demand function, if constraint (4) is not binding then the only demand that should be satisfied is the portion $q\tau^2/(q\tau^2 + r)$, which increases in τ. To see this, note that $q\tau^2/(q\tau^2 + r) \to 1$ as $\tau \to \infty$. Thus, *optimal fleet size is closer to demand in regions with shorter travel distances and hence with high τ* (e.g., in Haiti) than in regions with longer travel distances to cover and hence with low τ (e.g., in Sudan). In regions with shorter-distance missions and hence less frequent vehicle replacement, following demand more closely is a less expensive policy than it would be in regions with longer distance missions and hence more frequent vehicle replacement.

For an increasing demand function $D(t) = \alpha t$, we use $u(t)$ to represent the increase in fleet size during period t. Taking the derivative of (10) with respect to both q and r, we find that

$$\frac{d\bar{B}}{dq} = \frac{\alpha r \tau(t(h\tau + p) + p\tau)}{(q\tau^2 + r)^2} > 0, \tag{14}$$

$$\frac{d\bar{B}}{dr} = -\frac{\alpha q \tau(t(h\tau + p) + p\tau)}{q\tau^2 + r)^2} < 0. \tag{15}$$

These equalities imply that, for a linear demand function, the budget increases when the penalty for supply–demand mismatch is high.

Another notable interaction is that between the individual replacement policy τ and the fleet budget constraint. Although the optimal individual replacement policy aims to minimize each vehicle's operational costs, system constraints may render this policy infeasible in some situations. Despite the exogenous reasons (e.g., maintenance policies) driving these decisions, it remains possible for decision makers to modify them without affecting fleet quality. For the demand function αt, the budget is affected by τ as follows:

$$\frac{d\bar{B}}{d\tau}\begin{cases} > 0 & \text{if } \tau < \mathcal{F}, \\ < 0 & \text{if } \tau > \mathcal{F}; \end{cases} \tag{16}$$

here

$$\mathcal{F} := \frac{\sqrt{r(T^2(h^2 r + p^2 q) + 2hprT + p^2 r)} + hrT + pr}{pqT}.$$

Thus, as τ increases, the budget increases up to a threshold but then decreases afterward. The reason is that, when τ is small, a higher portion of the fleet per unit time is replaced, and so additional new vehicles are needed to satisfy the demand. The resulting increased costs are due to the increased variability of $u(t)$, which (as shown with regard to the LP model) is expensive. Even so, with higher values of τ, this variability cost is increasingly outweighed by the benefit of replacing vehicles at a lower rate.

5.3.3. Demand Fluctuating around a Constant Mean. We now consider a demand function that varies around a constant:

$$D(t) = \alpha + \beta \sin(\omega t). \tag{17}$$

PROPOSITION 4. *When constraints (4) and (5) are not binding, for the demand function $\alpha + \beta \sin(\omega t)$, the optimal fleet size is*

$$x^*(t) = \frac{\alpha q \tau^2}{r + q\tau^2} + \frac{\beta q \sin(\omega t)}{q + r\omega^2 + \frac{r}{\tau^2}} \quad (18)$$

and the optimal control is

$$u^*(t) = \frac{\alpha q \tau}{r + q\tau^2} + \frac{\beta q \tau [\sin(\omega t) + \tau \omega \cos(\omega t)]}{q\tau^2 + r\omega^2\tau^2 + r}. \quad (19)$$

Similar to the results of Proposition 1, here we have four cases for the optimal solutions depending on whether (4) and/or (5) are binding. However, we forgo discussing these cases because doing so does not enhance our intuition of the problem beyond that obtained via Proposition 1. The main difference is that, for the unconstrained case, Proposition 4 has no counterpart to the linearly increasing term in Equation (9).

We emphasize that our model suggests a level strategy when the HO is restricted in the rate at which it can increase or decrease the fleet size. Restrictions of this kind are common in HOs and result from operational constraints and bureaucratic reasons such as budget limitations, difficulties in shipping vehicles to different countries, the manufacturer's capacity restrictions, and the often vague timelines of humanitarian projects. Given that these constraints are fairly common in HOs, our model suggests that a level fleet size strategy is usually optimal in the humanitarian sector.

5.4. Comparison of LP- and QC-Based Policies

The QC model just described can be applied at the HQ level to identify optimal policies for aggregated fleet sizing. Because demand is relatively stable in development programs, past demand data can be used to estimate the parameters α, β, and ω in the demand function (7) and to generate aggregate forecasts for future periods. The parameters q and r can likewise be set based on past experience—or "reverse engineered" by applying the model to previous optimal policies obtained via the LP model. Finally, once the specific function for the optimal policy has been determined, optimal vehicle sizes and optimal purchasing/selling quantities can be found for each period of interest by computing numerical values for the state function $x(t)$ and for the control function $u(t)$ in those periods.

We shall illustrate the practical application of the QC model—and assess its validity—by comparing the optimal policy it generates for a specific demand scenario against the optimal policy generated by the LP model for the same scenario. We examine the country with the most complex and most general demand function among those we studied, that is, Sudan. The demand function in Sudan consisted of the three distinct patterns outlined in section 5.3. From $t = 0$ to $t = 22$ it was constant, with an initial value of 50 vehicles (i.e., $\alpha = 50$ without oscillation). From $t = 23$ to $t = 33$, demand increased linearly with $\alpha = 15.2$. Finally, from $t = 34$ to $t = 84$, it oscillated around a constant α at frequency $\omega \approx 0.26$. Therefore, the demand is estimated by setting $\alpha = 15.2$ for periods 23–33 and $\alpha = 250$ for the last interval (approximated as $D_t = 250 + 20 \sin 0.26t$). We calculated $\tau = \frac{100{,}000}{1{,}181.8} \approx 84$, where 100,000 km is the optimal vehicle replacement threshold (Pedraza Martinez and Van Wassenhove 2013) and 1,181.8 km is the average planned usage of vehicles in Sudan (see Table 1). Both r and q are set at HQ. We set a rather large $r = 250$ compared to $q = 2$, assuming that the per-period cost of each vehicle above (or below) the average demand is significantly higher than the costs associated with attempting to match demand.

The QC model generates the following policy. From $t = 0$ to $t = 22$, it retains the initial stock. From $t = 23$ to $t = 33$, Proposition 3 holds: new vehicles are purchased in accordance with Equation (13), and the optimal fleet size at any time during this period is given by Equation (12). Finally, Proposition 4 holds from $t = 34$ to $t = 84$; for this period, the model calls for purchasing new vehicles per Equation (19) and the optimal fleet size is given by Equation (18).

We used the policy just summarized to calculate the number of fleet vehicles specified by Equations (12) and (13) at discrete intervals $t = 0,1,\ldots,T$. We then compared these values to those generated by the LP model and calculated the absolute value of the relative difference. The comparison—also illustrated in Figure 3—indicates that the differences QC vs. LP (6.2%), QC vs. Demand (4.1%), and LP vs. Demand (5.4%) are limited and that the policies generated by the two models have a similar structure. However, the LP model typically generates a solution that exceeds demand, whereas the QC model keeps the fleet size closer to demand.

6. Conclusions, Limitations, and Future Research

In this article, we examine the properties of optimal humanitarian fleet procurement policies for development programs. After studying the operations of a large international organization (the ICRC) in Afghanistan, Sudan, and Ethiopia, we apply a LP model to calculate the optimal fleet size in each country. We then draw on these LP results to build a stylized QC model that, while preserving the properties of the LP solution, is more parsimonious and has data requirements that are better matched with the data to which HO headquarters typically has access.

Figure 3 Optimal QC Policy vs. Optimal LP Policy: Sudan

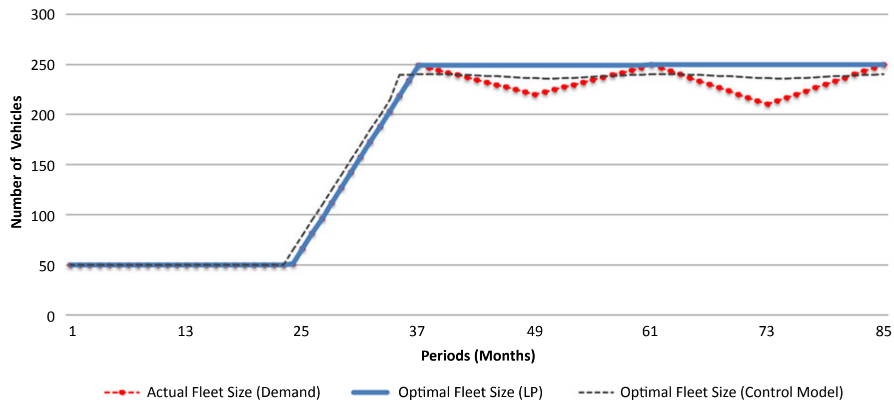

We use the model to characterize the general structure of the optimal policy under different demand scenarios and to obtain additional intuition on the trade-offs faced by HOs in their fleet management decisions. After demonstrating that the QC model's results are consistent with those of the LP model, we discuss the former's applicability to strategic asset planning.

We find that, if HOs are constrained by the rate at which fleet size levels can be changed, then the optimal policy is to level fleet size by smoothing out demand requirements. Yet if HOs can replace vehicles frequently, face substantial overstocking and understocking costs, and are relatively unconstrained by procurement budgets, then the optimal policy is to follow a so-called chase strategy. Because the humanitarian context seldom satisfies the latter conditions, our results also indicate that a level strategy would be optimal for most humanitarian missions. The analysis presented here can also be used to identify cases in which HOs can revert to simpler, intuitive procurement strategies that do not require data-intensive solutions. For instance: if budgets are tight, then in each period simply purchase as many vehicles as the budget allows until peak demand is eventually met. In addition, we show that lower budgets are required for missions in which the demand oscillation frequency exceeds the vehicle procurement frequency.

Our findings should be viewed in light of some limitations. First, the models are developed in the context of a fleet consisting of homogeneous vehicles. The case of heterogenous fleets could be treated either by decomposing the problem and running the models separately for different model categories or by adding model-specific indicators. Second, we have assumed that demand requirements are known before the program and do not change during the program. In the LP model, we assumed also that demand is exogenous and independent of fleet size, but when HOs increase their fleet size they may decide to run more missions simply because they can. Third, we assumed that a constant budget is available in each period. That might be true for large HOs, but not for small ones. Future research could well examine the relationships among the procurement policies and financial structure of HOs and their fleet management policies. Finally, the LP model parameters were estimated using data from a specific organization. Although the ICRC is a large and representative HO, the validity of our results should be tested also in other empirical contexts.

Acknowledgments

Mahyar Eftekhar carried out this research project while pursuing his doctoral studies at HEC Paris. He is grateful to the school for the financial support received. The associate editor, two anonymous reviewers, and participants at the POMS 2011 conference in Reno and the INFORMS 2010 annual meeting in Austin provided insightful comments. The authors also gratefully acknowledge the help of ICRC and Dario Moro, who provided invaluable assistance in the data collection process.

References

Altay, N., W. G. Green. 2006. OR/MS research in disaster operations management. *Eur. J. Oper. Res.* **175**: 475–493.

Barbarosoglu, G., L. Ozdamar, A. Cevik. 2002. An interactive approach for hierarchical analysis of helicopter logistics in disaster relief operations. *Eur. J. Oper. Res.* **140**(1): 118–133.

Besiou, M., A. Pedraza-Martinez, L. N. Van Wassenhove. 2012. Decentralization and earmarked funding in humanitarian logistics for relief and development. INSEAD Working Papers (2012/10/TOM/ISIC).

Campbell, A. M., D. Vanderbussche, W. Hermann. 2008. Routing for relief efforts. *Transport. Sci.* **42**(2): 127–145.

De Angelis, V., M. Mecoli, C. Nikoi, G. Storchi. 2007. Multiperiod integrated routing and scheduling of world food programme cargo planes in Angola. *Comput. Oper. Res.* **34**(6): 1601–1615.

Disparte, D. 2007. The postman parallel. *CarNation* **2**: 22–27.

Du, Y., R. Hall. 1997. Redistribution for center-terminal transportation networks. *Manage. Sci.* **43**(2): 145–157.

Han, L. D., F. Yuan, S. Chin, H. Hwang. 2006. Global optimization of emergency evacuation assignments. *Interfaces* **36**(6): 502–513.

Ingolfsson, A., S. Budge, E. Erkut. 2008. Optimal ambulance location with random delays and travel times. *Health Care Manage. Sci.* **11**: 262–274.

Jia, H., F. Ordonez, M. Dessouky. 2007. A modeling framework for facility location of medical services for large-scale emergencies. *IIE Trans.* **39**: 41–55.

Klincewicz, J., H. Luss, M. G. Pilcher. 1990. Fleet size planning when outside carrier services are available. *Transport. Sci.* **24**(3): 169–182.

Kovacs, G., K. Spens. 2007. Humanitarian logistics in disaster relief operations. *Int. J. Physic. Distrib. Log. Manag.* **37**(2): 99–114.

Lapre, M. A., A. S. Mukherjee, L. N. Van Wassenhove. 2000. Behind the learning curve: Linking learning activities to waste reduction. *Manage. Sci.* **46**(5): 597–611.

List, G. F., B. Wood, L. K. Nozick, M. A. Turnquist, D. A. Jones, E. A. Kjeldgaard, C. R. Lawton. 2003. Robust optimization for fleet planning under uncertainty. *Transp. Res. Part E: Log. Transp. Rev.* **39**: 209–227.

McCoy, J. 2013. Overcoming the challenges of the last mile: A model of Riders for Health. B. Denton, ed. *Handbook of Healthcare Operations Management: Methods and Applications*, Springer, NY, **184**: 483–509.

Pedraza Martinez, A., H. Sameer, L. N. Van Wassenhove. 2010. An operational mechanism design for fleet management coordination in humanitarian operations. INSEAD Working Paper (2010/87/TOM/ISIC).

Pedraza Martinez, A., O. Stapleton, L. N. Van Wassenhove. 2011. Last mile fleet management in humanitarian operations: A case-based approach. *J. Oper. Manag.* **29**(5): 404–421.

Pedraza Martinez, A., L. N. Van Wassenhove. 2013. Vehicle replacement in the International Committee of Red Cross. *Prod. Oper. Manag.* **22**(2): 365–376.

Salmeron, J., A. Apte. 2010. Stochastic optimization for natural disaster asset prepositioning. *Prod. Oper. Manag.* **19**(5): 561–574.

Sherali, H. D., E. K. Bish, X. Zhu. 2006. Airline fleet assignment concepts, models and algorithms. *Eur. J. Oper. Res.* **172**: 1–30.

Slack, N., S. Chambers, R. Johnston, A. Betts. 2006. *Operations and process management: Principles and practice for strategic impact*. Prentice Hall, Harlow, UK.

Song, D. P., C. F. Earl. 2008. Optimal empty vehicle repositioning and fleet-sizing for two-depot service systems. *Eur. J. Oper. Res.* **185**: 760–777.

Stapleton, O., A. Pedraza Martinez, L. N. Van Wassenhove. 2008. Keep on truckin': Managing vehicles at the International Committee of the Red Cross.

Tomasini, R., L. N. Van Wassenhove. 2009. *Humanitarian Logistics*. INSEAD Business Press, Fontainebleau, France.

Van Wassenhove, L. N. 2006. Humanitarian aid logistics: Supply chain management in high gear. *J. Oper. Res. Soc.* **57**: 475–489.

Vemuganti, R. R., M. Oblak, A. Aggarwal. 1989. Network models for fleet management. *Decis. Sci.* **20**(1): 182–197.

Yi, W., L. Ozdamar. 2007. A dynamic logistics coordination model for evacuation and support in disaster response activities. *Eur. J. Oper. Res.* **179**(3): 1177–1193.