Submitted to *Operations Research*
manuscript (Please, provide the manuscript number!)

The Nutritious Supply Chain: Optimizing Humanitarian Food Aid

Koen Peters, Sérgio Silva, Rui Gonçalves, Mirjana Kavelj
Supply Chain Planning Unit, World Food Programme, koen.peters@wfp.org, sergio.silva@wfp.org, rui.goncalves@wfp.org, mirjana.kavelj@wfp.org

Hein Fleuren, Dick den Hertog
Department of Econometrics and Operations Research, Tilburg University, h.fleuren@tilburguniversity.edu, d.denhertog@tilburguniversity.edu

Ozlem Ergun
Department of Mechanical and Industrial Engineering, Northeastern University, o.ergun@northeastern.edu

Mallory Freeman
Advanced Technology Group, UPS, malloryfreeman@ups.com

The World Food Programme (WFP) is the largest humanitarian agency fighting hunger worldwide, reaching around 90 million people with food assistance in 80 countries each year. To deal with the operational complexities inherent in its mandate, WFP has been developing tools to assist its decision makers with integrating supply chain decisions across departments and functional areas. This paper describes a mixed integer linear programming model that simultaneously optimizes the food basket to be delivered, the sourcing plan, the routing plan, and the transfer modality of a long-term recovery operation for each month in a pre-defined time horizon. By connecting traditional supply chain elements to nutritional objectives, we have made significant breakthroughs in the operational excellence of WFP's most complex operations, such as Iraq, Yemen, and El Niño. We show how we used optimization to reduce the operational costs in Iraq by 12%, without compromising the nutritional value supplied. Additionally, we show how we used optimization in Yemen to manage the scaling-up of the existing operation from three to six million beneficiaries.

Key words: supply chain; nutrition; MILP; humanitarian logistics; WFP; case study

1. Introduction

Humanitarian organizations are currently facing an unprecedented amount of high-level crises. Conflicts in countries such as Syria, South Sudan, and Yemen have been unprecedentedly long and large in scale, and many African countries are suffering from droughts and poor harvests. These crises result in rapidly deteriorating living conditions for everyone in the vicinity, threatening millions of innocent people with hunger, malnutrition, and worse. For decades, humanitarian organizations such as the United Nations (UN), Médecins sans Frontières (MSF), and the International Committee of the Red Cross (ICRC) have been doing everything in their power to provide assistance to those in need.

There are about 821 million undernourished people in the world today (FAO et al. 2018). This means that just over one in nine people on Earth do not have access to enough food to be healthy and lead an active life (FAO et al. 2015b). Hunger and malnutrition are in fact the number one risk to health worldwide—greater than AIDS, malaria, and tuberculosis combined. The UN considers hunger one of the greatest solvable problems, and this gave rise to the Zero Hunger Challenge (FAO et al. 2015a). The Zero Hunger Challenge is a global initiative launched by the previous UN Secretary General, Ban Ki-moon, and calls on everyone to do their part to eliminate hunger in our lifetimes.

One of the key players in responding to emergencies and eliminating hunger is the United Nations World Food Programme (WFP). WFP is the world's largest humanitarian organization fighting hunger worldwide, with more than 16,000 employees across the globe. In emergencies it distributes food where it is needed, saving the lives of victims of war, civil conflict, and natural disasters. After the cause of an emergency has passed, WFP uses food assistance to help communities rebuild their lives and return to a semblance of normality. In 2017, a total of 91.4 million people spread across 83 countries were reached directly by WFP assistance, with 80% of WFP resources allocated to supporting beneficiaries in conflict environments (WFP 2018).

Delivering food to this number of people in these kinds of environments requires a supply chain that is agile, adaptable, and aligned (Lee 2004). Every year, WFP procures, transports, and distributes almost 4 million metric tonnes (4 billion kilograms) of food to people in need. The poor infrastructure conditions and high levels of insecurity in conflict areas necessitate creativity and flexibility in delivering this food, and WFP has been known to employ every variety of transportation method—from elephants and camels to airplanes and barges. In 2017, on any given day WFP had 5,000 chartered trucks on the road, 92 chartered planes in the air, and 20 chartered ships at sea (WFP 2018). By virtue of its excellent logistical performance, WFP is even mandated to lead logistics operations whenever a humanitarian agency requires a joint response from UN agencies and the humanitarian community.

1.1. Literature

There are several phases in responding to a disaster. In the literature we generally see that disaster timelines are split into four stages: mitigation, preparedness, response, and recovery (Van Wassenhove 2006; Altay and Green 2006; Ergun et al. 2011). Mitigation focuses on the prevention of a disaster and the reduction of its intensity, for instance by setting up alert systems that warn against floods or by defining building guidelines for areas that are vulnerable to earthquakes or hurricanes. Preparedness is more concerned with setting up the appropriate resources (both physical and human), such as the building and stocking of strategic warehouses or the training of personnel. Response starts once the disaster has occurred; it includes activities such as the delivery of food and services to those in need, evacuation of the affected region, and the collection of debris. Lastly, the recovery phase aims to return a semblance of normality to the area affected.

Holguín-Veras et al. (2012) split up the recovery phase in short-term and long-term recovery, where short-term recovery is a transitional stage covering damage assessments, repairs, housing, etc. Long-term recovery may span multiple years and includes the rebuilding of infrastructure and distribution of medical and food supplies to prevent disease and malnutrition. This long-term recovery is one of the focus areas of WFP. Most literature on humanitarian logistics focuses on the preparedness and response side of a disaster, whereas long-term recovery is a topic that does not receive much attention.

Many researchers have characterized and discussed the challenges and opportunities of humanitarian logistics. Van Wassenhove (2006) discusses how the private sector can learn from the agility and adaptability inherent in humanitarian supply chains, and how the humanitarian sector can learn from the established Supply Chain Management (SCM) best practices in the private sector. In a follow-up paper, Van Wassenhove and Pedraza Martinez (2012) illustrate the potential of Operations Research (OR) in particular for adapting such SCM best practices to humanitarian logistics. Apte (2010) discusses research issues and potential actions surrounding the field of humanitarian logistics, and reviews analytical models from the literature to understand the state-of-the-art in humanitarian logistics. Apte mentions the sustaining of long-term developmental aid when discussing research gaps. Çelik et al. (2012) also mention that there is a research gap in relation to long-term recovery—using the term Long-Term Humanitarian Development instead. Additionally, they highlight that there is a lack of good implementation of decision support tools in humanitarian operations. Holguín-Veras et al. (2012) pinpoint research gaps that need to be filled to enhance both the efficiency of humanitarian logistics and the realism of the mathematical models designed to support it. They argue that humanitarian logistics is too broad a field to fit neatly into a single definition of operational conditions, and urge researchers to treat these different operational conditions separately.

Despite the qualitative attention to long-term recovery, there are few mathematical formulations available that cover the entire scope of a humanitarian supply chain. Most existing research is focused on (a combination of) three sub-problems, namely that of facility location (Balcik and Beamon 2008), distribution (Haghani and Oh 1996; Özdamar and Demir 2012; Rottkemper et al. 2012; Rancourt et al. 2015), and inventory control (Beamon and Kotleba 2006; Pérez-Rodríguez and Holguín-Veras 2015). Humanitarian researchers have extended the traditional models for these three sub-problems with constraints, objectives, and solution methods to facilitate the special requirements of a humanitarian supply chain. These extensions include (but are not limited to) research on the appropriate objective function (Holguín-Veras et al. 2013; Gralla et al. 2014; Gutjahr and Nolz 2016) and modeling uncertain demand, prices, and capacities (Bozorgi-Amiri et al.

2013; Rawls and Turnquist 2012; Ben-Tal et al. 2011). For an in-depth discussion of what is necessary to make a traditional (i.e. commercial) supply chain model work in a humanitarian context, we refer to Holguín-Veras et al. (2012).

For food aid in particular, the design of food baskets is an important topic. Many papers illustrate how mathematical models can be used to generate food baskets or nutritious products that satisfy all nutritional requirements (Ryan et al. 2014; Briend et al. 2003; Baldi 2017; Fleige et al. 2010). For example, Carlson et al. (2003) develop a quadratic optimization model that, for each age-gender group, selects the optimal food plan that meets the dietary standards, adheres to the budget constraints, and resembles the reported food consumption for that specific age-gender group (therefore making it more likely that the food plan "sticks"). Similarly, Chastre et al. (2007) develop linear programming routines to generate hypothetical diets using a combination of foods that will enable a family to meet their energy and nutrient requirements as recommended by the World Health Organisation (WHO) and the Food and Agriculture Organisation (FAO) at the lowest possible cost. As the software (Cost of the Diet) can account for the frequency with which foods are eaten, for example by specifying that a particular food is eaten three times a day, the food baskets can be adjusted to reflect typical dietary patterns.

One aspect of humanitarian aid that is not yet covered in the Operations Research literature is that humanitarian organizations have several methods for delivering assistance at their disposal. Whereas in the past transfers were done exclusively in-kind (i.e. the organization buys, transports, and distributes commodities), there has been a recent trend in providing beneficiaries with cash or vouchers instead, allowing them to purchase their own commodities at local markets or selected retailers. Lentz et al. (2013) discuss the rise of these new food assistance instruments. There are multiple interacting effects to be considered, and it is important to weigh the benefits (reduced transportation costs, increase in the beneficiary's dignity) against the dangers (the influx of cash or vouchers may distort the local economy). They state that no single tool is always and everywhere preferable, and seek to educate the reader on their appropriate use. In particular, the economic repercussions of choosing one transfer modality over another are notoriously hard to measure, making it difficult to select the appropriate instrument. Ryckembusch et al. (2013) discuss an analytical tool that is able to compare the cost-effectiveness of transfer modalities. Their tool, the "Omega Value", considers the trade-off between total costs (procurement, transportation, services, etc.) and the "Nutrient Value Score" (NVS). The NVS is a weighted score function that shows to what extent all nutritional requirements are met (for nutrients such as energy, protein, vitamins, etc.). The Omega Value shows the nutritional value per dollar spent, and can assist policy makers in making the right choice of food basket and transfer modality.

1.2. Research questions

The literature survey above highlighted several research gaps. In particular, we found that:

a) long-term humanitarian assistance has received little attention,
b) most papers focus on sub-problems (e.g. facility location),
c) alternative transfer modalities (e.g. cash, vouchers) have not been addressed yet, and
d) there is a lack of models that are actually being implemented.

In this paper we try to close these four gaps by addressing our main research question:

1. Is it possible to integrate all of WFP's key supply chain decisions (i.e. food basket design, transfer modality selection, and sourcing & delivery plan) into a tractable mathematical model?

Additionally, we try to assess the impact of the model through three additional research questions:

2. Is the model implementable given current data availability and can it be incorporated into WFP's decision-making process?
3. What lessons can we learn with regards to food aid policy from the model's results?
4. Can we generalize the approach that has been developed to other domains?

The rest of the paper is structured as follows: in Section 2 we provide a brief description of the key components of WFP's supply chain. In Section 3 we develop a Mixed Integer Linear Programming (MILP) model that optimizes the food basket design, transfer modality selection, sourcing plan, and delivery plan of a long-term recovery operation for each month in a pre-defined time horizon. In Section 4 we apply this model to three WFP use cases: Iraq, Yemen, and El Niño, demonstrating how we can use the optimization model to reduce the operational costs and improve the effectiveness of complex operations. In Section 5 we discuss some insights from the model as relevant for food aid policy and decision making. In Section 6 we extend the core model with additional transfer modalities. In Section 7 we explore other possible applications (humanitarian and non-humanitarian) of the model and the approach we are describing in this paper. In Section 8 we provide concluding remarks.

2. WFP's Supply Chain

Before defining the mathematics, it would be prudent to offer a brief discussion of WFP's supply chain and the components of our model. We consider the food basket, transfer modalities, sourcing, and logistics network.

2.1. Food basket

WFP's food aid depends on the context of the crisis and the beneficiaries' health, demographics, and access to non-WFP food. The bulk of food assistance is provided through General Food Distribution (GFD)—a monthly parcel that provides enough food for a standard-sized family. Additional food assistance is supplied to more vulnerable beneficiaries, such as those suffering from malnutrition, young children, and pregnant or nursing women. Nutritionists track the nutrients that WFP's food baskets provide, and make sure that they align with the needs of the different beneficiary types (e.g. a young child may receive a fortified cereal on top of the generic ration). Traditionally demand was defined in terms of pre-determined food baskets per beneficiary type, but in this paper we take one step back and define demand as nutritional requirements—allowing us to optimize the food baskets, rather than using the pre-defined ones. For the applications in this paper we consider eleven nutrients (3 macronutrients and 8 micronutrients), but the developed methodology works for any number of nutrients.

2.2. Transfer modalities

There are multiple ways of satisfying the demand—be it a pre-defined food basket or a nutritional requirement. Whereas WFP used to supply the necessary commodities itself, the last decade there has been a transition towards cash-based transfers (also known as Cash & Vouchers). Traditional food transfers require WFP to source, store, transport, and distribute the commodities, whereas a cash-based transfer allows beneficiaries to obtain the commodities themselves from local markets and shops. Food transfers incur additional supply chain costs, but allow WFP more flexibility in their sourcing (making use of bulk purchases, price seasonality, tax waivers, etc.). Cash-based transfers (CBT) are more direct (so more of the donation is spent on procurement), but they are dependent on local market availability and their impact is harder to measure up-front. We consider three types of CBT assistance:

1. Commodity Vouchers; these vouchers entitle a beneficiary to a specific ration of one or more commodities (e.g. 15kg of maize), to be purchased at WFP-contracted stores.
2. Value Vouchers; these vouchers entitle a beneficiary to purchase a range of commodities from WFP-contracted stores, up to a certain value (e.g. 20 USD). The beneficiary can choose what combination of commodities best fits their needs.
3. Cash; this is an unrestricted form of assistance, where the beneficiary has full control over where they spend their money and what they spend it on.

In this paper we develop and apply a model that considers Commodity Vouchers. In Section 6 we develop a model that also incorporates the Value Vouchers and Cash, and include some first results.

2.3. Sourcing

Depending on the transfer modality, multiple sources are available. Cash-based transfers (such as e-Vouchers and direct cash transfers to debit cards) allow beneficiaries to purchase commodities at Local Markets (LM). For food transfers WFP has more purchasing options. We distinguish three supplier types: International Suppliers (IS), Regional Suppliers (RS), and Local Suppliers (LS).

Local Suppliers (wholesale) can be found in the recipient country; Regional Suppliers can be found in neighboring countries, where transport from the supplier to the recipient country is usually done by land. Procuring from International Suppliers involves shipping the commodities to a Discharge Port (DP).

2.4. Logistics network

The hand-over between suppliers and WFP is very flexible and dependent on the Incoterms (Ramberg 2011). Usually WFP takes charge of the commodities at a loading port (for international purchases) or at one of its hubs inside the recipient country (for local and regional purchases). From Discharge Ports WFP moves commodities into the country to so-called Extended Delivery Points (EDP), transshipment points where commodities can be stored, packaged, consolidated, etc. From the EDPs the commodities are transported (usually by truck) to the Final Delivery Points (FDP), where they are handed over to WFP's Cooperating Partners: local NGOs that will take care of what is called the Last Mile Distribution (Balcik et al. 2008) in order to reach the final distribution points (e.g. schools, villages, hospitals, etc.). Note that the logistics network may vary between countries and is very context dependent, WFP may take charge of the entire delivery network (from pick-up at the supplier to last mile distribution), or it may outsource the network partially or entirely to local logistics providers or sometimes even the government. For example, WFP has about 5,000 trucks on the road every day, but its own fleet consists of "only" a thousand trucks, so much of the transportation is outsourced.

3. Model

In this paper we extend a capacitated, multi-commodity, multi-period network flow model with nutritional components. For an introduction to network flow models, we refer to Ball et al. (1995); an example tailored to a humanitarian supply chain can be found in Haghani and Oh (1996). This type of model is commonly used to optimize sourcing and delivery strategies, two of the four components that we aim to integrate. The supply chain network is sketched as per Fig. 1. The

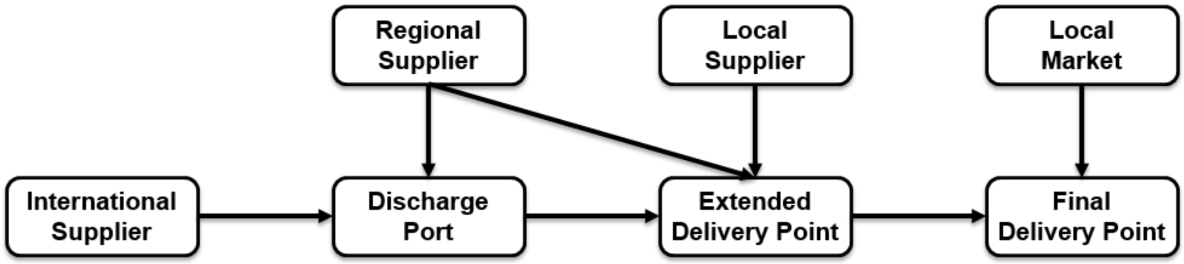

Figure 1 High-level overview of the modeled supply chain network.

suppliers are modeled as source nodes, with the discharge ports and WFP warehouses acting as transshipment nodes. The demand (or sink) nodes are the Final Delivery Points (FDPs), where demand is dependent on the (variable) food basket and the number of beneficiaries.

To integrate the food basket design into the network flow model, we define a ration variable that governs the commodities flowing into an FDP, ensuring that the food WFP sends addresses the nutrient gap, is distributable, and is palatable. This link between nutrients, commodities, and rations resembles the flexible bill of materials sometimes found in the manufacturing industry (Ram et al. 2006), where the end-product is a (monthly) food basket and the number of beneficiaries in the FDP the demand for this end-product. An interesting distinction is that in the manufacturing industry the end-product is fixed and the fulfilled demand is variable, whereas here we usually consider the fulfilled demand fixed (100% of beneficiaries receive a food basket) and make the quality of the end-product variable. So, when there is a funding shortfall WFP prefers to supply a less nutritious food basket to all beneficiaries rather than supplying the full food basket to fewer beneficiaries. Product design and sourcing are traditionally done separately, and there is much potential to improve end-to-end performance through joint decision making (Novak and Eppinger 2001).

Integrating the transfer modality selection is possible but not straight-forward. The easiest approach is to model local markets as source nodes that are linked directly to FDPs. Beneficiaries can then receive some or all of their commodities from these local markets through Commodity

Vouchers (one of the transfer modalities). We discuss a more advanced approach that also covers Value Vouchers and Cash in Section 6.

Note that the model we are using at WFP is vast and needs to be able to handle a plethora of (mathematically trivial) constraints, such as sourcing restrictions, capacity utilization, beneficiary preferences, funding allocation, etc. Additionally, the core of the model is based on traditional network flow model specifications. Because of this, in the main body of the paper we discuss only the unique components. A more comprehensive description of the model that was used for the case studies can be found in Appendix A.

3.1. Defining demand

Typically in multi-commodity network flows the demand is defined as a parameter dem_{ikt} for every location i, commodity k, and period t. Here, we break up the demand using several parameters:

dem_{it} = Number of beneficiaries at location i in period t

$days_t$ = Feeding days in period t

$nutreq_l$ = Daily requirements for nutrient l

$nutval_{kl}$ = Contribution to nutrient l per gram of commodity k.

We define variables R_k (continuous, non-negative) to optimize the food basket, where R_k represents the gram per person per day of commodity k that WFP supplies. These ration variables ensure that beneficiaries receive the same food basket, and allow us to impose restrictions on the composition of the ration. For example, this formulation makes it easy to specify that beneficiaries should receive at least 300 grams of rice per day, or that split peas are not accepted. Generally speaking, R_k should be such that the nutrient gap is closed:

$$\sum_k nutval_{kl} \times R_k \geq nutreq_l \ \forall l. \tag{1}$$

Rather than a fixed demand parameter dem_{ikt} we end up with a variable demand:

$$D_{ikt} = dem_{it} \times days_t \times R_k \ \forall i, k, t. \tag{2}$$

Note that the ration variable R_k can easily be extended to allow more flexibility, e.g. defining rations as R_{kt} allows for a food basket that is dynamic over time. This helps WFP to make the most of seasonal price windows (e.g. the harvest season basket could be different from the lean season basket). Similarly we can define the rations as R_{bk} to cover the needs of different beneficiary types (e.g. different nutritional requirements). Differentiation by location (R_{ik}) is less common for in-kind food baskets, but could open up opportunities to diversify the basket using commodities that are only available in a handful of locations (e.g. procurement from local markets). An example of location differentiation for CBT is discussed in Section 6.

3.2. Objective and goals

Conventional network flow models usually revolve around profit maximization or cost minimization, but for humanitarian operations there are often additional considerations. We refer to Holguín-Veras et al. (2013) for a comprehensive treatise on this subject.

In this paper, we consider four main classes of objectives:

1. **Efficiency** relates to resource utilization, and covers objectives such as the cost of the operation and the utilization of port capacity.
2. **Effectiveness** relates to the nutritional impact of WFP's assistance, and covers objectives such as the remaining nutrient gap and the dietary diversity score.
3. **Development** relates to the impact on the local economy, and covers objectives such as the dollars spent on local procurement and CBT assistance.
4. **Agility** relates to the responsiveness of the solution, and covers objectives such as the maximum and average lead time of the supply chain.

Typically, users of the model specify a range of programmatic goals (e.g. they request a solution that covers 100% of the nutrient gap, covers 30% of the needs through local procurement, with a lead time of at most 3 months). The model then identifies the solution that meets these targets at the lowest possible cost (i.e. pareto-optimal). Some examples of common goals and their mathematical formulation are provided in Appendix A.4.

Depending on the goals and objective chosen, additional variables and constraints are added to the base model. Many goals can be formulated as linear constraints, even if the measure itself is non-linear (e.g. percentage local procurement). Note that not all combinations of objectives and goals are possible in a linear model. For example, maximizing the number of beneficiaries would conflict with a variable food basket, because in the demand constraint of the network flow model we multiply the number of beneficiaries with their demand—if both are variable, the resulting model is non-linear.

3.3. Solution approach

The mathematical model is formulated in Python, and solved using the COIN-OR solver (Saltzman 2002). We connect to the COIN-OR solver using the PuLP module (Mitchell et al. 2011). PuLP does not support multi-threading, so all calculations are performed using a single processor. Parallel processing is supported however, so it is possible to run multiple scenarios simultaneously. The coding language, solver, and module are all freeware, making the tool easy to implement.

The mathematical optimization model is first pre-processed to keep the problem size as small as possible. By cleaning and filtering the input data we prevent redundant variables and constraints, resulting in significant boosts to the solution speed. Additionally, we split up the constraints of the mathematical model into core constraints and situational constraints. This allows us to quickly set up new instances of the optimization model when doing scenario analyses. In practice, we find that the model setup time is reduced by approximately 80% for sub-sequent scenarios.

Users can interact with the mathematical model through a web application, coded in Python and hosted using Django. The web app allows them to add operational constraints, specify goals for the key outputs, and gives them access to a range of automated analyses. Results from their analyses are displayed in interactive graphics that provide quick insights into all of the important decision variables (food basket composition, sourcing plan, delivery plan, etc.) and the resulting performance (through pivot tables, charts, etc.).

3.4. Verification and validation

When using analytics to (re)design or manage an operation, it is of paramount importance that the model is verified and validated. In the humanitarian field in particular, many people are analytically averse - there used to be a general consensus that humanitarian operations are too difficult to model and optimize. This made extensive and continual verification and validation crucial. We developed the optimization model and tool over the course of several years, in close collaboration with the end users (WFP decision makers in the field). Through iterative development we were able to identify the required capabilities, model them, and ensure that they were implemented correctly.

When collaborating with WFP's Country Offices, we ensure that we mirror their budgeted costs when imposing the current solution on the optimization model. This allows us to quickly compare the model's calculations with theirs, which immediately highlights any components that are estimated differently or where some unmodeled costs are incurred. The latter triggers a series of in-depth discussions where we identify the unmodeled costs and find a way to introduce them into the mathematical model.

An other form of validation we use regularly is comparing the current solution to unrestricted solutions from the optimization tool, allowing us to rapidly identify operational constraints and preferences. This approach also generates strong buy-in from the local experts, because they start appreciating the scope, flexibility, and speed of the optimization tool.

The frequent rounds of verification and validation contributed significantly to the adoption rate of the tool, ensured that the model was built on data that was actually available, and that its results were intuitive and usable for WFP decision makers.

3.5. Assumptions and limitations

Every optimization tool of this size comes with a long list of assumptions and limitations; we highlight the most important ones.

We assume that every beneficiary (of a certain type) receives the same food basket, and that no beneficiary goes unfed. This assumption is based on the humanitarian principles of humanity and impartiality—it would be unethical to feed only those beneficiaries that are conveniently located for instance.

We do not model the transportation network between suppliers and loading ports. WFP floats tenders in specific countries to ensure competitive prices, and this means that price forecasts are done on a country level. As a result, the suppliers may differ every time WFP orders, making the cost from a supplier to the loading port difficult to estimate. Generally WFP requests FOB (Free On Board) offers, which means that the supplier takes care of the transportation up to the loading of the ship at a major port in their country.

Costs are captured using \$/mt rates. Given WFP's economies of scale and the type of contracts that it usually has with suppliers and contractors, these linear rates are representative of the actual costs.

The modeled nutrition measure, the Nutritional Value Score, is a simple average across all nutrients. We are working with WFP's Nutrition department to come to better measures, reflecting the importance of some nutrients over others. For general food distribution for example, a shortfall in macro nutrients (energy, protein, fat) is often considered more severe than a shortfall in minerals or vitamins. In programs that prevent or treat malnutrition on the other hand, we want to favor these micro nutrients more heavily.

The current model is not robust against uncertain parameters (such as costs, lead times, and capacities). While we recognize the importance of robustness in optimizing humanitarian supply chains (see for example the El Niňo application in the next section), the introduction of robust optimization has been put on hold until WFP improves its data and becomes more analytically mature. In parallel to the roll out of the optimization tool described in this paper, we are developing advanced forecasting algorithms for WFP's uncertain data. In absence of high-quality forecasts, it is difficult to get traction for robust optimization. In the meantime we verify with WFP's experts that the parameter values used are the current "best guess", and base our analyses on that.

4. Applications

The mathematical model presented in this paper is the result of iterative development over more than four years. Through regular pilots we were able to identify what data is reliably available, what kind of analyses are the most impactful, and how the optimization tool fits into WFP's business processes. Additionally, we use these pilots to continually verify and validate the mathematical model underlying the tool by scheduling in-depth sessions with experts throughout the organization and by comparing the tool's outputs to historical performance.

Since we started applying the tool, optimization gained significant traction within WFP. At present, the software is used continually to provide WFP's biggest and most complex operations with optimization support, such as Syria (three million beneficiaries), Ethiopia (five million), South Sudan (two million), and Yemen (eight million). Together, these four operations account for some eighteen million beneficiaries—a fifth of WFP's average yearly workload. In this section we highlight some of the results we achieved in Iraq, Yemen, and the South African region, and show how different types of analyses are being used to support strategic decision making in WFP's most complex operations.

4.1. Iraq

Years of conflict have hindered Iraq's economic development. Since 2014, the occupation of the Islamic State of Iraq and the Levant (ISIL) in Iraq has resulted in the displacement of more than three million people. When, two years later, the Iraqi Security Forces (ISF) launched a military offensive to regain control, fighting deepened insecurity, rolled back development and exacerbated vulnerabilities. Many Iraqis sought refuge in neighbouring countries and in Europe. Beset by violence, social disruption and economic hardship, thousands of Iraqi families were left in desperate need of food assistance.

Although many are now returning home every month, around 700,000 Iraqis are still living in camps with few possibilities to earn an income enabling them to put food on the table. In addition, an estimated quarter of a million Syrian refugees have sought refuge in northern Iraq, placing additional pressure on limited resources. At the end of 2017, around 800,000 people still require some sort of food assistance in Iraq every month; some 10 million need humanitarian assistance in general. The security landscape remains volatile, sometimes posing access challenges for humanitarian actors. (WFP 2015a).

WFP has been operating in Iraq since 1991, and has provided food assistance in the country since April 2014 through emergency operations to assist those hundreds of thousands Iraqis forced from their homes by recent violence. In the face of further mass displacement from major Iraqi cities, such as Mosul and Ramadi, WFP scaled up activities to reach an average of 1.5 million people per month in all 18 governorates, including hard-to-reach areas. WFP is assisting people through monthly family food parcels for those with access to cooking facilities, food vouchers that can be redeemed at local shops, and ready-to-eat food known as Immediate Response Rations that provide a family of five with food for three days.

WFP's operations in Iraq often face funding shortfalls, so it is vital that the operation's design is as cost-effective as possible. In October 2015 Iraq's Country Office requested optimization support to redesign their food basket for different levels of funding, allowing them to supply as many kilocalories as possible from the donations that they receive.

We worked intensively with their supply chain management team to gather all the necessary data and to identify operational constraints (with respect to procurement, logistics, transfer modalities, beneficiary preferences, etc.). Once we had a good grasp of the situation we ran hundreds of analyses using the optimization tool to identify alternative designs that could improve the performance of the operation. Many of the solutions were pushing the envelope, so there was much going back and forth to verify the feasibility of these solutions. This intense collaboration and the iterative process resulted in their supply chain management team feeling a strong sense of ownership of the final outcomes, resulting in a rapid implementation of the final recommendation.

One of the main deliverables to Iraq's management team was Fig. 2, containing our recommendations for Iraq's Family Food Parcel (FFP)—supplied to half a million Iraqis every month. The

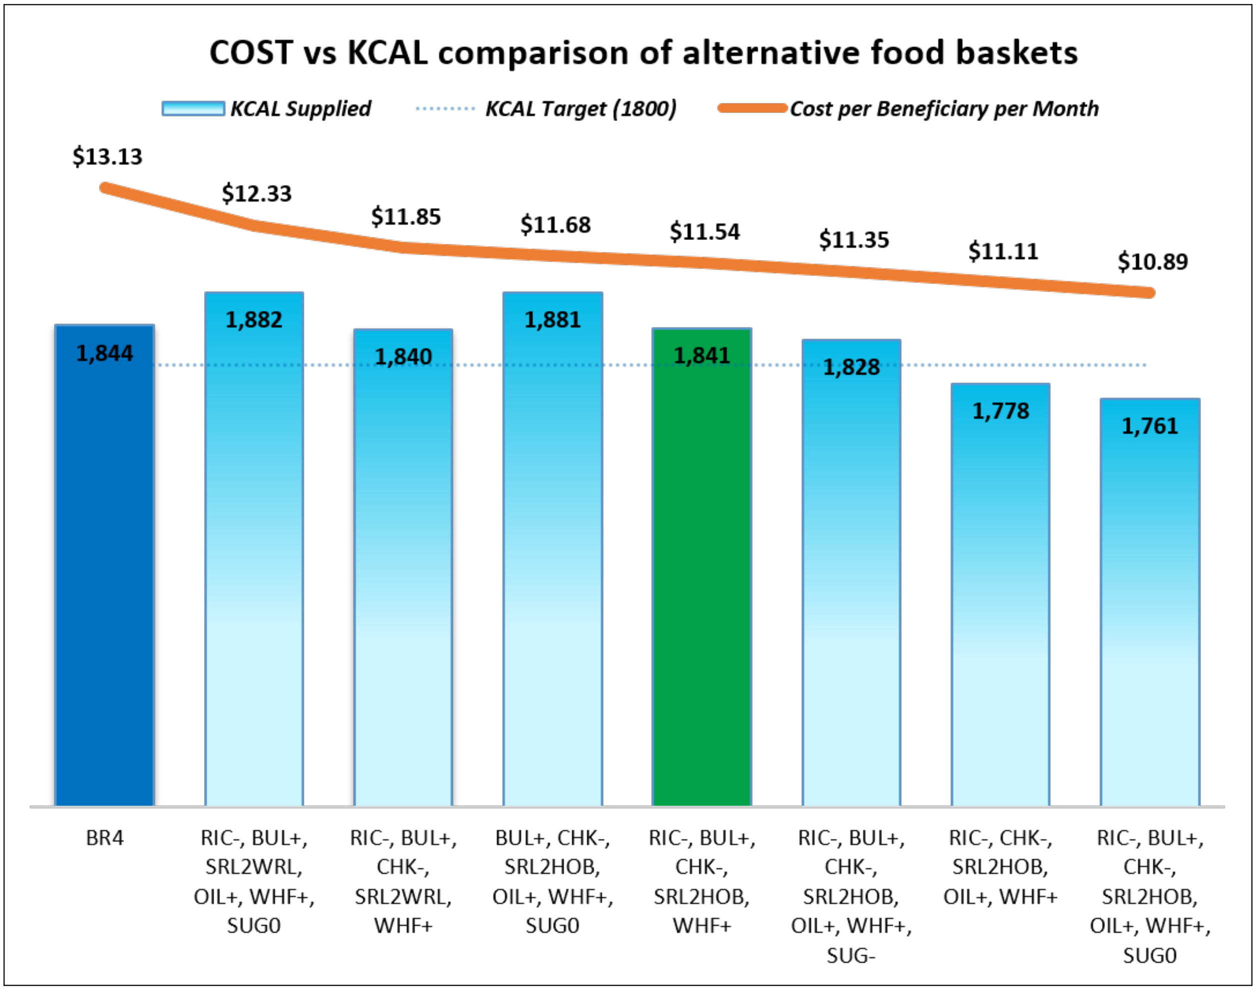

Figure 2 An overview of alternative food baskets for Iraq's Family Food Parcel (FFP). We compare the kilocalories and cost per beneficiary per month of each option (as given by the acronyms on the x-axis). The acronyms show how the food basket changed compared with the original food basket (BR4). RIC=Rice, BUL=Bulgur Wheat, WHF=Wheat Flour, SRL=Split Red Lentils, WRL=Whole Red Lentils, WIB=White Beans, CHK=Chickpeas, HOB=Horse Beans, OIL=Sunflower Oil, SUG=White Sugar. A "+" or "-" denotes an increase or decrease in the current ration size. A "0" denotes a removal of the commodity, whereas the "2" denotes that we change commodity a to commodity b using the same ration size. Ration size increments are pre-defined and based on commercial packaging types for that specific commodity. The green option has been implemented in practice, representing a cost reduction of 12% while supplying only 3 kilocalories less.

first column represents the latest official food basket (BR4), supplying 1844 kilocalories (daily) at a cost of 13.13 USD per beneficiary per month. Each subsequent column is an optimized basket that uses different interventions (commodity swaps, ration adjustments, etc.) to improve the cost-effectiveness of the operation. Iraq officially adopted the food basket represented by the green column, supplying 1841 kilocalories (only 3 kcal less) and 69% of the nutrient gap (same as BR4) at a cost of 11.54 USD per beneficiary per month (12% reduction). At the time of implementation the FFP was being supplied to 800,000 beneficiaries every month, so this adoption corresponded to 1.3M USD monthly savings. This meant that WFP could supply the same amount of beneficiaries with a nutritious food basket in case of up to 1.3M USD funding shortfall per month, or these 12% savings could be used to supply an additional 109,000 (14%) beneficiaries with a FFP every month if funding remained at the same level. Note that in this analysis all savings result from changing the food basket, we did not identify any significant savings on the existing supply chain network. The optimized food basket was distributed throughout 2016 and 2017, representing total savings of more than 25M USD, with positive feedback from the beneficiaries regarding the changes in commodities and ration sizes compared to the original food basket.

4.2. Yemen

Even before fighting broke out in early 2015, Yemen was one of the poorest countries in the Arab world. With an average life expectancy below 64, the nation is ranked 178th out of 188 on the Human Development Index - a composite statistic of life expectancy, education, and per capita income indicators, that is used to rank countries into four tiers of human development.

Nearly four years of conflict have left thousands of civilians dead and 3 million internally displaced. Its impact on the country's infrastructure has been devastating, with major overland routes and airports severely damaged. Despite ongoing humanitarian assistance, 15.9 million people wake up hungry every day, which is more than half of Yemen's population of 28 million. It is estimated that, in the absence of food assistance, this number would go up to 20 million (WFP 2015b).

WFP has been active in Yemen since 1967. In 2019, WFP is scaling up to provide 12 million people with monthly food assistance through direct food distributions or vouchers that people can use at retailers in areas where the markets are functioning. Each family of six gets a monthly ration of wheat flour, pulses, vegetable oil, sugar and salt. Beneficiaries include internally displaced persons and returnees, vulnerable populations in the most food insecure areas, people affected by transient crises, infants, pregnant and nursing women affected by acute and chronic malnutrition, and school-age children.

Since the start of the main operation WFP has gradually been scaling up towards the twelve million beneficiaries, and this is proving to be a herculean task in light of the limited resources available and the escalation of conflict within Yemen. In December 2015 they were still reaching three million beneficiaries with in-kind food assistance, with the aim of reaching five million in 2016. Cash-based transfers were being scaled up to one million beneficiaries in parallel, allowing WFP to reach 6 million people with life-saving assistance. In November 2015 Yemen's management team requested optimization support for their pending scale-up from three to six million beneficiaries.

Similar to our approach in Iraq (and all other applications), we worked intensively with local experts and management to gather the necessary data and identify all operational constraints from a supply chain and donor/beneficiary preference perspective. We were requested to keep the commodities in the food basket as it was, and focus mainly on the optimization of ration sizes. Seeing as distribution is carried out in active war zones, we designed the food baskets in such a way that the monthly ration for each family (of six) corresponds to an industry-standard packaging size—making distribution as fast and seamless as possible. We presented a breakdown of our recommendations in the form of a column chart similar to the one used in Iraq, but the most powerful deliverable was a Beneficiary Matrix.

The Beneficiary Matrix (Fig. 3) shows the interdependence of funding levels, nutritional targets, and beneficiary numbers. It shows decision makers how much it costs to supply different numbers of beneficiaries with the current food basket, and how alternative food baskets allow WFP to

Basket Name	USD per Beneficiary	KCAL	NVS	20M USD	30M USD	40M USD	45M USD	50M USD	55M USD	60M USD	70M USD	80M USD	90M USD	100M USD
Current basket	$ 13.89	100%	95%	1.4	2.2	2.9	3.2	3.6	4.0	4.3	**5.0**	5.8	6.5	7.2
WHE75 SYP10 OIL5 WSB10	$ 13.69	98%	94%	1.5	2.2	2.9	3.3	3.7	4.0	4.4	5.1	5.8	6.6	7.3
WHE75 SYP10 OIL7 WSB05	$ 13.28	97%	91%	1.5	2.3	3.0	3.4	3.8	4.1	4.5	5.3	6.0	6.8	7.5
WHE75 SYP10 OIL5 WSB05	$ 12.87	93%	89%	1.6	2.3	3.1	3.5	3.9	4.3	4.7	5.4	6.2	7.0	7.8
WHE50 SYP10 OIL7 WSB15	$ 11.92	85%	94%	1.7	2.5	3.4	3.8	4.2	4.6	**5.0**	5.9	6.7	7.6	8.4
WHE50 SYP10 OIL6 WSB15	$ 11.72	83%	94%	1.7	2.6	3.4	3.8	4.3	4.7	5.1	6.0	6.8	7.7	8.5
WHE50 SYP10 OIL7 WSB10	$ 11.10	80%	90%	1.8	2.7	3.6	4.1	4.5	**5.0**	5.4	6.3	7.2	8.1	9.0
WHE50 SYP10 OIL5 WSB10	$ 10.70	76%	88%	1.9	2.8	3.7	4.2	4.7	5.1	5.6	6.5	7.5	8.4	9.3
WHE50 SYP10 OIL6 WSB05	$ 10.08	73%	84%	2.0	3.0	4.0	4.5	**5.0**	5.5	6.0	6.9	7.9	8.9	9.9

Figure 3: This Beneficiary Matrix (Yemen) shows the interdependence of funding levels, nutritional targets, and beneficiary numbers. Columns 2-4 show the performance of a food basket (column 1). We display the cost per beneficiary per month and two effectiveness measures: the kilocalories and the percentage of all nutrients supplied (as percentages of the minimum daily requirements). The remaining columns show how many beneficiaries we can supply (in million beneficiaries per month) with these food baskets under different funding scenarios (ranging from 20 to 100M USD per month). For example, the current food basket can supply 4M beneficiaries if WFP receives 55M USD per month, and 5M beneficiaries in case WFP receives 70M USD per month.

either cope with funding shortfalls (without cutting down on beneficiary numbers) or to increase the number of beneficiaries reached with the available funding (by cutting down on nutrition). For example, suppose we are supplying 4 million beneficiaries with the current food basket and we would like to scale up to 5 million. The matrix shows us that supplying 5 million beneficiaries with the current basket would cost 70M USD per month (15M USD (27%) more than the current cost). Alternatively, we could reach the 5 million beneficiaries by slightly reducing the nutrition levels of the food basket. In this case, we can observe that there is a food basket (WHE50 SYP10 OIL7 WSB10) that supplies 5 million beneficiaries with 80% kilocalories and 90% NVS for the same level of funding (55M USD). Insight into this trichotomy allows WFP to manage their scale-up properly, and enables it to communicate clearly to donors what is required from them if WFP is to reach all people at risk within Yemen with nutritious food. Every quarter the matrix is updated to reflect the latest data changes, to the point where WFP is now reaching around eight million beneficiaries, and planning to reach its target of twelve million in 2019.

4.3. El Niño

In 2016-2017, the El Niño climate pattern, which is strongly linked to weather fluctuations around the globe, was fueling an international food security crisis for millions of people. By disturbing rainfall and temperature patterns, El Niño affected agriculture, water supplies and the spread of disease, and was threatening the food security and livelihoods for some 60 million people worldwide. Particular areas of concern included nearly all of Southern Africa which was the hardest hit region; Ethiopia and its neighbors Somalia and Sudan in East Africa; Central America's 'dry corridor', nearby Haiti and the northern region of South America; and many of Asia's island nations including Indonesia, Papua New Guinea, and Philippines.

Between October 2015 and January 2016, El Niño conditions caused the lowest recorded rainfall in at least 35 years across many regions of Southern Africa. The same period also recorded the hottest temperatures in the past 10 years. El Niño's impact on rain-fed agriculture is severe. Poor rainfall combined with excessive temperatures created conditions that were unfavorable for crop growth in many areas. In Lesotho, South Africa, Swaziland, Zambia, and Zimbabwe, planting was delayed by up to two months or more, which severely impacted maize yields. It became evident that the 2015-16 maize harvest would be insufficient to cover full cereal needs for the region without significant importation.

As more information about the impact became available, we were asked to provide support by drafting various contingency analyses. Given that the local cereal availability would be too limited to meet all needs, we used the model to explore global sourcing strategies under different capacity and export assumptions. For example, if some of the countries barely have enough to meet their own needs, they may impose export bans to ensure availability for their own population. We found that many countries appeared to be self-sufficient, but WFP would need to source about 20,000 metric tonnes of white maize from international markets, more than a quarter of the total needs. Given the market prices and the strong GMO restrictions in the South African region, Argentina and Mexico were identified as the optimal sourcing locations by our model. In Fig. 4 we show the global sourcing plan.

With slow-onset disasters such as the El Niño crisis, the assumptions and forecasts are very prone to change as the crisis develops. Because of this, it is important to perform sensitivity analyses so WFP can ensure its plans are still valid if the forecasts change. At the time of the analysis, South Africa was procuring significant quantities locally, in particular for cereals and pulses. If the availability of commodities reduces due to poor harvests, we often see significant price spikes in local markets, so it was important to evaluate whether a local procurement strategy still made sense if the local prices increased drastically. In Fig. 5 we show an optimized sourcing strategy under different price scenarios. We find that for the current prices the optimal breakdown is to purchase 29.7% International, 31.9% Local, and 38.4% Regional. Price increases as little as 10% already increase the optimal International procurement ratio to 40%, up to 78% when Local/Regional prices double. If we consider the cost of the baseline plan (i.e. optimal plan under the assumption that the prices remain stable) under different price scenarios, we see that it is up to 700k USD

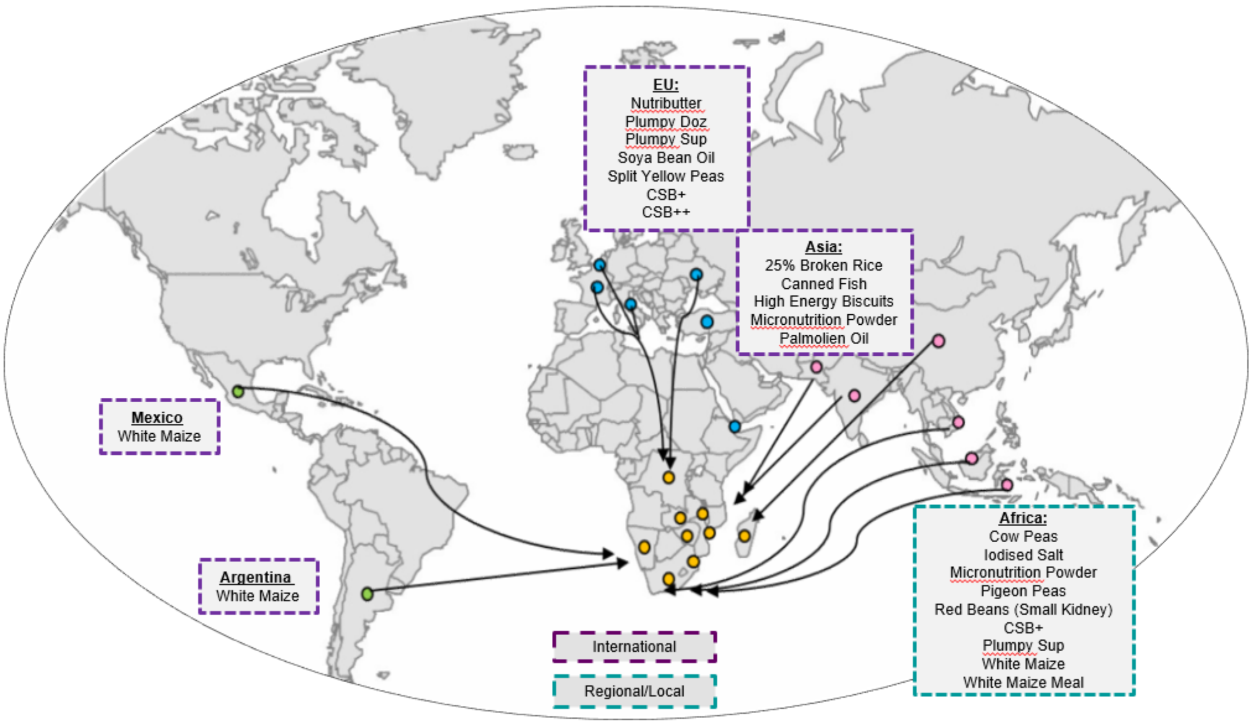

Figure 4 An optimized global sourcing plan in case local cereal production is not sufficient to meet all needs in the most affected South African countries.

more expensive per month (7.8%) than the optimized plan for scenarios where prices increase. We see that depending on El Niño's effect on the local markets, which is likely to increase local prices, we may have to consider transitioning to a more international-heavy sourcing plan. Considering that such a shift would increase lead times by two months for many commodities, it is important to start shifting straight away.

5. Insights

Applications in countries such as Iraq, Yemen, and South Africa have demonstrated the added value of integrating key decisions into one model, allowing WFP to significantly improve the efficiency and effectiveness of its operations. In this section we explore some of the insights that can be obtained from this integrated optimization approach, and what this means for decision making at WFP and for food aid policy in general.

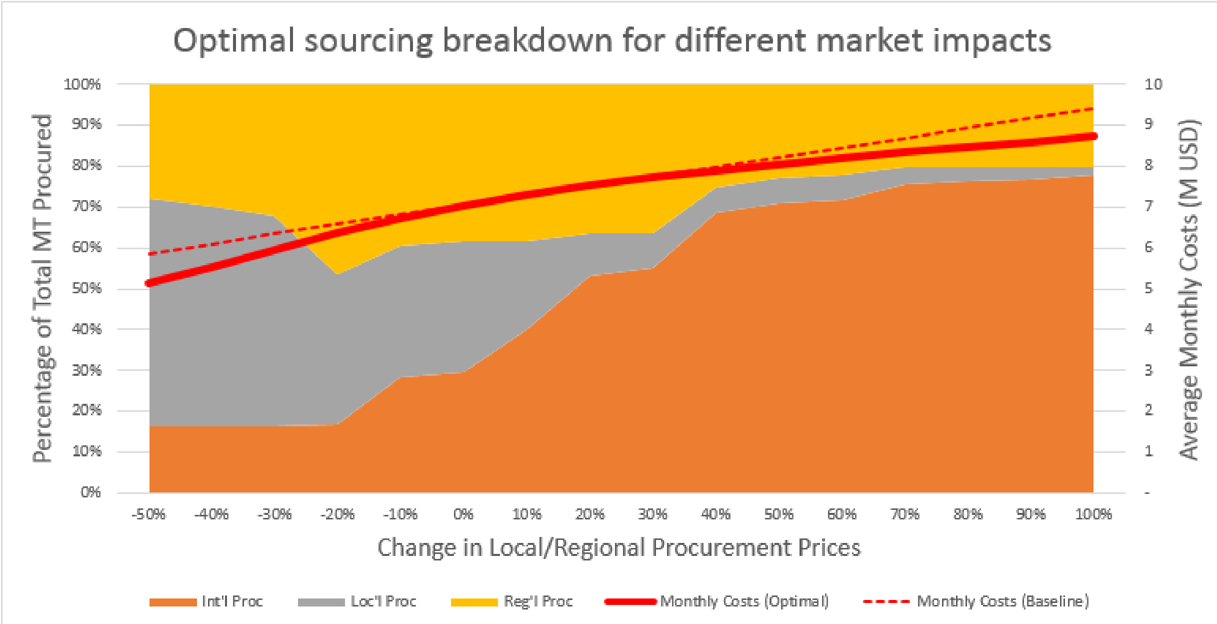

Figure 5 Optimal sourcing strategies under different funding scenarios. On the left axis we show the percentage of the total need supplied through the three sourcing types (local, regional, and international). On the x-axis we show the different price scenarios (from very favorable at -50% to very unfavorable at +100%). The red line is plotted against the right axis, and shows how the optimal monthly costs differ for these scenarios (7M USD currently, 8.7M USD if local/regional prices increase by 100%). The dotted line shows the costs under these scenarios if we implement the optimized plan as per the current price estimates (i.e. the 0% plan).

5.1. Trade-offs

The mathematical model we defined has a single objective (usually costs), but we tend to evaluate solutions using multiple metrics (efficiency, effectiveness, agility, and development). We use a goal programming approach to find Pareto-efficient curves for these metrics, providing insight into what kind of outcomes are achievable and how this would affect the cost of the operation. We highlight two of these trade-off analyses in particular: one for the Nutritional Value Score (effectiveness) and one for the Cash-Based Transfer Ratio (development).

The Nutritional Value Score Trade-off curve (Fig. 6) shows the lowest cost at which WFP can attain different levels of nutrition (effectiveness). Each plot represents a solution to the mathematical model, i.e. it corresponds to an optimized food basket, sourcing plan, and delivery plan that

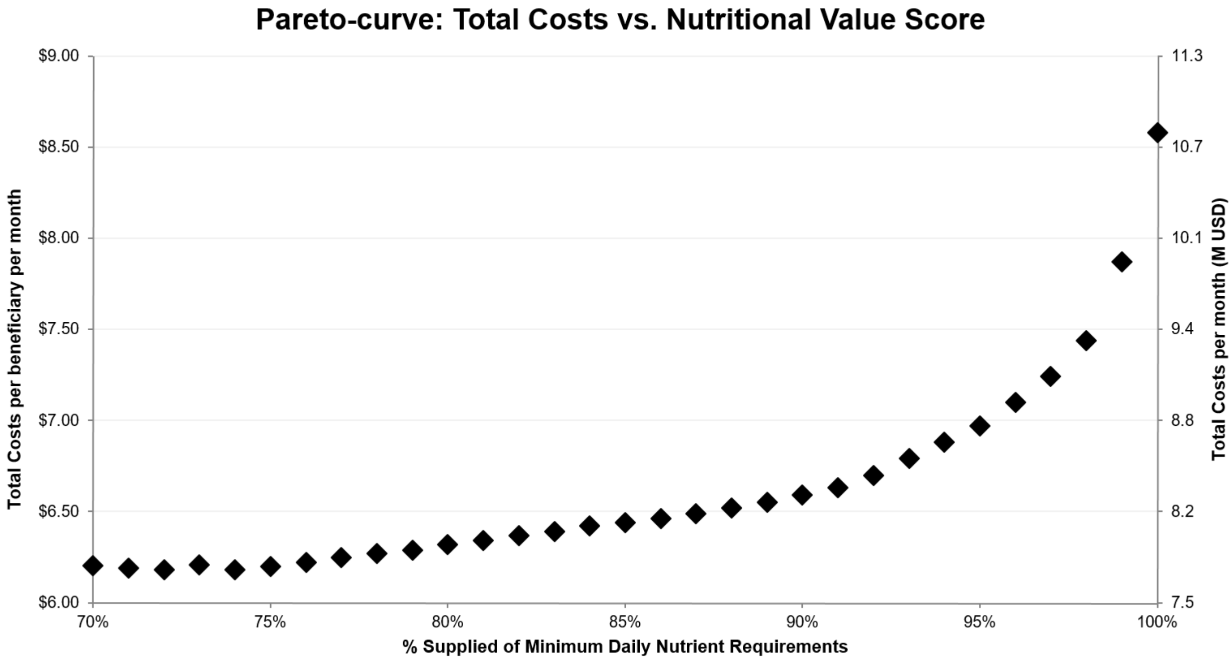

Figure 6 The Pareto-efficient curve for the Nutritional Value Score (NVS). For different levels of NVS the optimization tool finds the cheapest food basket, sourcing plan, and delivery plan. On the y-axes we have the cost per beneficiary per month (left) and the cost of the entire operation per month (right). Note that the cost per beneficiary per month is calculated for beneficiaries that receive this food basket, whereas the operation costs also include every other activity in that country.

takes into account all user-added constraints. Ideally WFP supplies 100% of the Recommended Daily Intake (RDI), but depending on the context (available funding, number of beneficiaries in need, access to nutritious commodities) this may not be realistic. With this trade-off graph we show decision makers the price tag of supplying different levels of nutrition. In Fig. 6 for instance, we can see that supplying 95% of the required nutrients reduces the cost of the operation from 10.8M USD to 8.8M USD (18.7% reduction). This means that we could still supply all beneficiaries with 95% of their requirements if our funding is 18.7% short, and alternatively that we could reach 15-20% more beneficiaries with 95% of their requirements if there is no funding shortfall. Because of this big difference in cost and potential outreach, it is important to understand what level of nutritional requirements food aid policy recommends, and what these requirements are based on. If supplying 95% is still “OK” for the majority of beneficiaries, it may have a big impact on how humanitarian organizations design their programs.

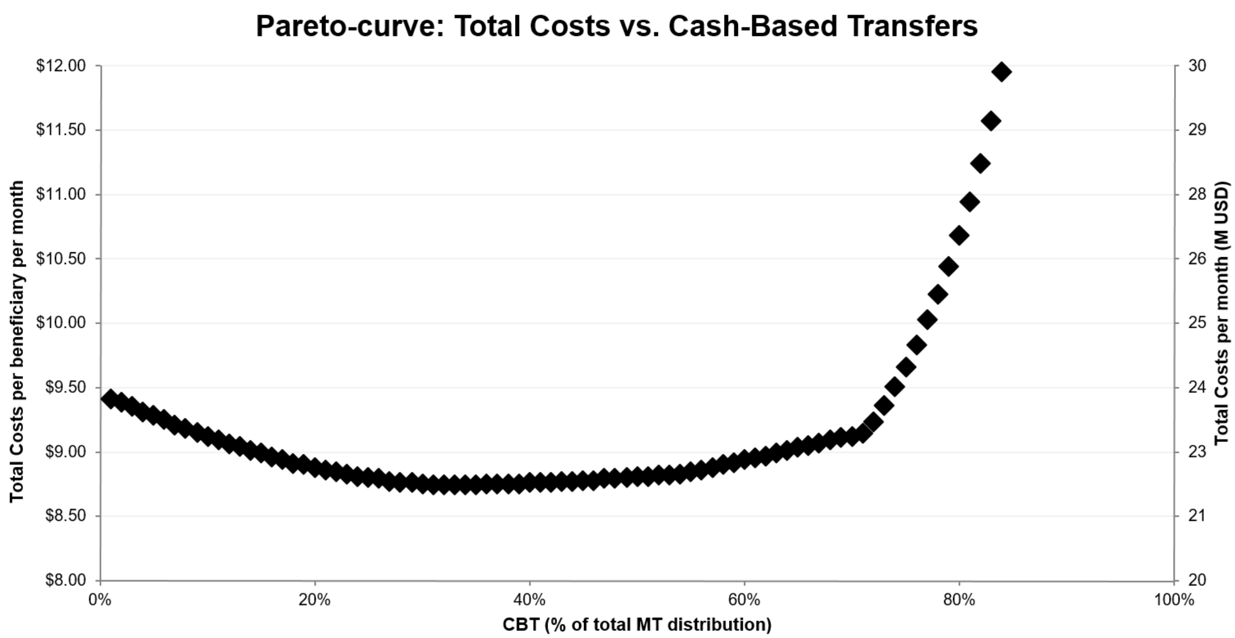

Figure 7 The Pareto-efficient curve for the Cash-Based Transfers metric. For different levels of CBT (as percentage of the total metric tonnes distributed) the optimization tool finds the cheapest food basket, sourcing plan, and delivery plan. On the y-axes we have the cost per beneficiary per month (left) and the cost of the entire operation per month (right). Note that the cost per beneficiary per month is calculated for beneficiaries that receive this food basket, whereas the operation costs also include every other activity in that country.

The Cash-Based Transfer Ratio Trade-off (Fig. 7) shows the lowest cost at which WFP can attain different levels of Cash & Voucher transfers. Cash-based transfers (CBT) are very popular among WFP donors at the moment, because of the positive effects they have on local markets and the dignity of beneficiaries. Depending on the operation, CBT may be more expensive than supplying the food in-kind however, so it is important to assess whether the premium price weighs up against the non-quantifiable benefits. With this graph we show the most cost-effective level of CBT, and how the cost increases as one deviates further from the optimal ratio. We can observe (Fig. 7) that initially the cost decreases when increasing CBT, as the model will start supplying the most remote FDPs and the most cost-effective local commodities through commodity vouchers. Increasing CBT further means that the model has to start using less cost-effective local commodities in order to supply the required nutrients. The example shows that the optimal ratio for this operation was

33%, but that we could scale up to 70% CBT with only a small increase in cost (0.9M USD; 4.4%). Decision makers have to weigh the (non-quantifiable) added value of increasing CBT for local markets and beneficiaries against this increase in cost. Insight into the relative cost of CBT compared to in-kind, and to what extent the local markets can absorb the new demand is very important when deciding on the optimal transfer modality. Many food aid policies recommend one transfer modality in favor of the others, but we find that there is no single optimal transfer modality.

5.2. Cost-effectiveness of commodities

Optimal food baskets are highly dependent on the available commodities, their nutritional profile, and the nutrient gap that we aim to close. For example, let us consider four cereals that are viable for the Nigeria operation: rice, sorghum, wheat, and (white) maize.

Table 1 displays the nutritional contents of these four cereals for 11 nutrients. By optimizing the sourcing strategy for these four cereals we can obtain their minimum cost (end-to-end). Combining the optimal costs with the nutritional contents allows us to calculate the relative cost-effectiveness of each of these commodities. Table 2 shows how much USD we have to spend on each cereal to supply 100% of the Recommended Daily Intake (RDI). For example, we can observe that we need to purchase $0.66 worth of rice for each beneficiary if we want to supply a days' worth of kilocalories (2100), as opposed to $0.54 if we opt for sorghum. Similarly, we can learn that wheat is the most cost-effective when it comes to supplying proteins, white maize for raboflavin, etc. Across all nutrients, sorghum is the most cost-effective cereal for Nigeria, and rice the least cost-effective. Of course, a food basket is composed of more than just cereals, so the availability of commodities like oils and pulses and their respective cost-effectiveness will decide what the optimal cereal is. All commodities have a unique nutritional profile, and the optimization model described in this paper can help us balance these nutritional values when designing food baskets while taking into account the resulting costs.

Cereal	Energy (kcal)	Protein (g)	Fat (g)	Calcium (mg)	Iron (mg)	Iodine (μg)	Vit. A (μg)	Thiamine (mg)	Riboflavin (mg)	Niacin (mg)	Vit. C (mg)
Rice	360.00	7.00	0.50	9.00	1.70	0.00	0.00	0.10	0.03	5.58	0.00
Sorghum	335.00	11.00	3.00	26.00	4.50	0.00	0.00	0.34	0.15	5.00	0.00
Wheat	330.00	12.30	1.50	36.00	4.00	0.00	0.00	0.30	0.07	8.92	0.00
White Maize	350.00	10.00	4.00	7.00	2.71	0.00	0.00	0.39	0.20	2.20	0.00
RDI	2100.00	52.50	40.00	450.00	22.00	150.00	500.00	0.90	1.40	13.86	28.00

Table 1: This table contains the nutritional value of four cereals (per 100g), with the Recommended Daily Intake (RDI) for reference. Note that the units differ by nutrient (e.g. Energy is measured in kilocalories, whereas Iron is measured in mg).

Cereal	Cost ($/kg)	Energy	Protein	Fat	Calcium	Iron	Iodine	Vit. A	Thiamine	Riboflavin	Niacin	Vit. C
Rice	$1.13	$0.66	$0.84	$9.01	$5.63	$1.46	inf	inf	$1.01	$5.25	$0.28	inf
Sorghum	$0.86	$0.54	$0.41	$1.14	$1.48	$0.42	inf	inf	$0.23	$0.80	$0.24	inf
Wheat	$0.87	$0.55	$0.37	$2.32	$1.09	$0.48	inf	inf	$0.26	$1.74	$0.14	inf
White Maize	$0.93	$0.56	$0.49	$0.93	$5.96	$0.75	inf	inf	$0.22	$0.65	$0.58	inf

Table 2: This table contains the relative cost-effectiveness of four cereals with respect to supplying different nutrients. Each value indicates how much USD we have to spend on the cereal every day to supply 100% of the nutritional requirement, so the lower the value the more cost-effective the commodity is for that particular nutrient.

Rather than redesigning food baskets from scratch using these insights, in practice we try to stay close to the current design of the operation (making any recommendations much easier to implement). We highlight two types of analyses that allow us to use slight redesigns of the current food basket to improve the performance of an operation: a commodity swap analysis and an optimization of ration sizes.

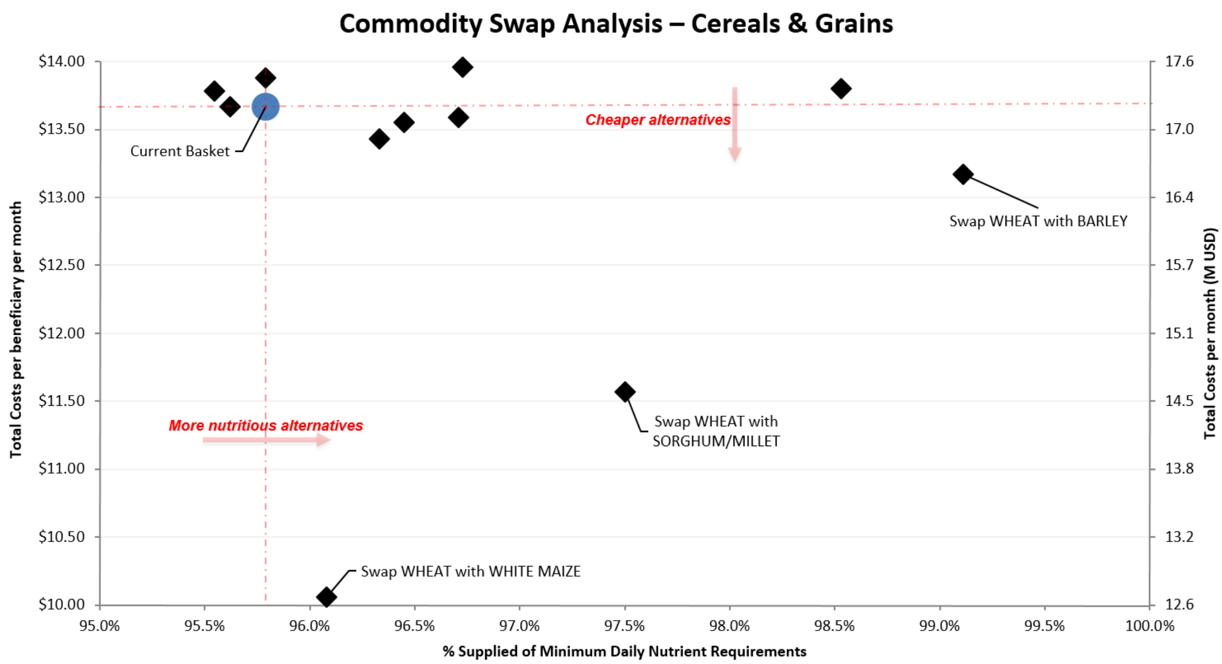

Figure 8 This Commodity Swap Analysis shows the impact on nutrition and cost of swapping one commodity (from the Cereals & Grains food group). The y-axes show the efficiency, with the cost per beneficiary per month on the left and the monthly operation costs on the right. On the x-axis we show the effectiveness, as a percentage of the minimum daily nutrient requirements. Each plot corresponds to a unique food basket (incl. its optimal sourcing and delivery plan), with the blue circle representing the current food basket. All plots that are below and to the right of the circle are therefore strict improvements (from a cost and nutrition perspective respectively).

The Commodity Swap Analysis (Fig. 8) shows the impact of swapping one commodity in the current food basket on the performance of the operation—in this case on the supplied nutrients (effectiveness) and operation costs (efficiency). The graph shows an analysis of cereals and grains for one of WFP's projects. Each iteration, we replace one of the cereals/grains in the current food basket with a cereal/grain that is not currently included (using the same ration size). We can observe, for example, that replacing the wheat in the current basket with white maize would increase the NVS by 0.3%-points, while reducing the cost of the operation by 4.5M USD (26.5%). Decision makers can combine these quantitative insights with context-specific constraints (such as donor and beneficiary preferences) to find the best food basket design. For instance, changing to sorghum/millet makes no sense if the beneficiaries have no idea how to prepare this grain, no matter how cost-effective it is.

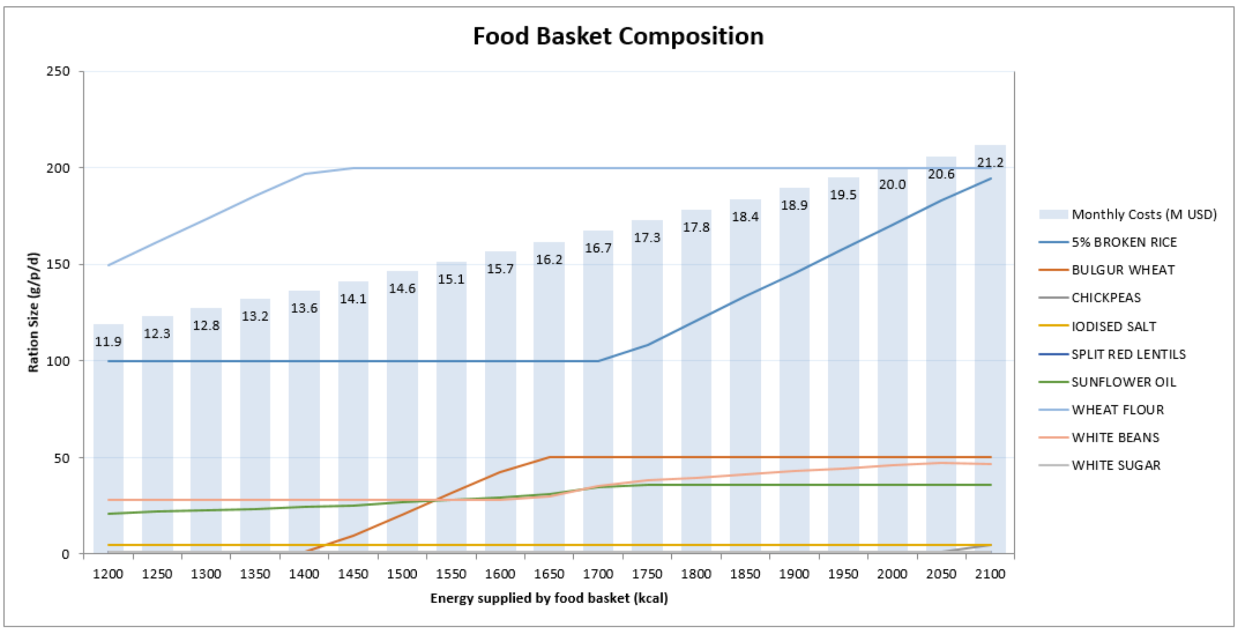

Figure 9 This graph displays the optimal ration sizes for different levels of nutrition. On the y-axis we display the ration size in grams per person per day, while the x-axis shows different effectiveness levels (measured in kilocalories). Each column corresponds to a food basket (composed of a combination of the nine commodities with ration sizes as given on the y-axis) and displays the total cost per month.

The Optimal Ration Size Analysis (Fig. 9) shows how the current food basket can be adjusted to become as cost-effective as possible for different performance measures (in this case we consider a nutritional target—kilocalories). When doing this type of analysis we include additional constraints that guide the composition of the food basket, to make sure that all optimized food baskets are in fact implementable. The graph shows clearly which commodities are the most cost-effective (the ones that are increased first) when it comes to supplying energy - in this case the wheat flour is better than the bulgur wheat which in turn is better than the rice. The columns show the cost (Millions USD per month) for each of the food baskets. We regularly use this kind of analysis to help WFP decide which ration sizes to cut in case of a funding shortfall, as it helps to provide as much nutrition as possible given the available funding.

6. Extension to Value Vouchers and Cash

The model was initially developed to manage two transfer modalities: in-kind transfers and commodity vouchers. Since then, the humanitarian community and WFP in particular have seen a steady shift towards less restricted forms of assistance, such as Value Vouchers and Cash. In this section, we describe how the model can be extended to deal with these transfer modalities.

6.1. Cash-Based Transfers

Even though Cash-Based Transfers were introduced more than a decade ago (May 2007), it is still difficult to properly evaluate and compare the various transfer modalities from a cost-effectiveness perspective. The main reason for this is that for two of the three CBT approaches (Value Vouchers and Cash) the effectiveness of the transfer is dependent on the beneficiary's purchasing behavior and local market conditions (cost, availability). For Commodity Vouchers we know exactly what the beneficiary will receive, and base our effectiveness on that. In practice, we consider two main approaches for evaluating the effectiveness of CBT assistance: Ex-Ante and Ex-Post.

In an Ex-Ante analysis we make a best guess about the purchasing behavior of the beneficiary (e.g. 15kg maize, 5kg pulses, 2L oil, etc.) and calculate the effectiveness based on the assumption that all beneficiaries purchase and consume this basket. This "best guess" basket is usually based on local expert knowledge and household surveys. This projection of the effectiveness requires little data and can provide some initial insights into the potential cost-effectiveness of CBT assistance.

In practice however, the expected purchasing behavior is usually very different from the actual purchasing behavior. In an Ex-Post analysis we analyze the effectiveness of CBT assistance based on actual expenditure data, that is increasingly being collected by several branches within WFP (e.g. Point of Sales level SKU data). Each purchase we can link to a nutritional contribution and cost, and this allows us to define the actual expenditure pattern of beneficiaries and the effectiveness of WFP assistance for any location and period where actual expenditure data is available. This is a very data-heavy approach (SKUs have to be converted to weights and commodities, and all entries have to be translated to English), and can only be applied after CBT assistance has been implemented.

Outputs from these Ex-Ante and Ex-Post analyses are generally compared against the cost-effectiveness of in-kind food assistance to decide on the most appropriate transfer modality. In practice, the transfer modality decision is made for each FDP (i.e. either the FDP receives an in-kind basket or CBT assistance), and ideally decision makers would like to switch flexibly between the modalities during the planning horizon (e.g. CBT during harvest season, in-kind during the lean season). Note that Commodity Vouchers are often considered as in-kind assistance, because the beneficiary does not have the freedom to choose what to spend their voucher on. We decided to incorporate the transfer modality selection more explicitly in the model, seeing as the initial model covered only the Commodity Voucher assistance.

6.2. Location Differentiation approach

The idea behind this extension (Location Differentiation) is to decide for each FDP individually whether it should use an in-kind food basket (may include Commodity Vouchers) or a CBT contribution (Value Voucher or Cash), where the in-kind basket is optimized as before and the CBT contribution would result in performance (cost, effectiveness) as per our best guess (i.e. data input). The CBT basket (which may differ by location and by month), would ideally be based on an Ex-Post evaluation of prior CBT distributions so that we can compare the actual performance of CBT in each location against the actual performance of an optimized in-kind basket. The mathematical formulation for this extension can be found in Appendix B.

The main outputs of the Location Differentiation approach are an optimized food basket for the selected activity, and an indication for each FDP for each month whether they should opt for CBT assistance (for which the effectiveness may differ by location and by month) or the In-Kind assistance (for which the effectiveness may differ by month but not by location). The model will choose for CBT in those FDPs where it makes the most sense (e.g. well-established markets), and/or during the periods that make the most sense (e.g. post-harvest season). This allows a very flexible approach, and a very elegant way to trade-off between CBT and In-Kind while utilizing all of the functionality of the core model (e.g. specifying programmatic goals, funding restrictions, etc.).

We implemented the Location Differentiation approach in our web app and tested it using data from WFP's Mali operation. We found that the optimal cost was 20.22 USD per beneficiary per month if we allow for In-Kind only. When we allowed for unrestricted CBT (Location Differentiation approach) as per an ex-post evaluation, we found that the cost decreased to 17.65 USD (13% reduction). In the optimised scenario, about 53% of the beneficiaries received unrestricted CBT, with an average NVS of 73% for FDPs that receive CBT and 100% for those that receive In-Kind. We found that the optimal transfer modality is heavily dependent on the location, as prices (and expenditure patterns) may vary significantly in different cities. In practice we usually found that very urbanized areas favor CBT because there is competition between vendors, whereas in the more remote areas (where there may be only one or two vendors), a WFP-owned supply chain is generally more efficient.

7. Other Possible Applications

7.1. Humanitarian applications

WFP supplies life-saving assistance to more than 80 million beneficiaries every year, but it is not the only actor when it comes to providing food aid. Local governments often have their own school feeding programs in place, supplying children in primary and secondary education with nutritious meals. Sometimes they also provide take-away rations to incentivize parents to continue sending their children to school. Consider that India alone is already providing some 120 million children with school lunches through its Mid Day Meal Scheme. Applying the methods developed in this paper to such school meal programs has a huge potential, and we have started engaging with several countries to help them improve the cost-effectiveness of their school feeding initiatives (Benin, Colombia, Egypt, India, etc.).

Another humanitarian area where the methodology developed has merit is in the supply chain for non-foods (e.g. UNHRD, WASH, vaccines, etc.). Currently, beneficiary needs are defined in terms of nutrients, but needs for non-food items such as mosquito nets and vaccines could be modeled in very similar ways. This becomes particularly interesting when some non-food items supply multiple needs, e.g. you could provide a kit instead of individual items if it makes sense from a cost perspective. Additionally, it would allow for coordination across multiple humanitarian actors and avoids duplication of efforts (e.g. there is need for one mosquito net, regardless of how many agencies aim to provide one). We are currently exploring this extension together with the Logistics Cluster.

7.2. Non-humanitarian applications

One of the core elements of our approach is that we do not supply a fixed product, but that we optimize the product while looking into the resulting supply chain. This train of thought may also have merit for non-humanitarian problems. Below we discuss some examples.

In the cattle industry there are many options for feedstuffs (Arosemena et al. 1995). If a farmer specifies the nutrients that they would like their cattle to consume, the model developed here could be used by producers to design an optimal mix of feedstuffs given the availability, costs, and nutritional contents of relevant components (barley, silage, fodder beet, etc.). For example, if there is access to hay or by-products from the brewing industry, through optimization it is possible to design feedstuffs that rely more heavily on these components.

Similarly, we could consider the optimal blend and production cost of different coffee types (Vitzthum 1999, Andueza et al. 2007). Given some required specs (bitterness, foam, aroma, etc.), we could design the optimal blend (Arabica vs Robusta bean, conventional vs torrefacto roast, bean origin) and brewing process (decoction, infusion, gravitational feed, or pressurized percolation). Simultaneously we could look into the necessary supply chain and manufacturing costs to create the final blend.

8. Conclusion

In this paper we have explored the possibilities of integrating key supply chain decisions, such as the food basket design, the transfer modality selection, and the sourcing & delivery plan, into a single model. We defined a mixed integer linear programming model based on a multi-commodity, multi-period network flow that covers the full spectrum of decisions at WFP.

Applications in some of WFP's most complex emergencies (Iraq, Yemen, and El Niño), have demonstrated the added value of optimization in managing humanitarian operations. In Yemen, optimization enabled the supply chain management team to scale up its operation from three to six million beneficiaries through rapid evaluation of the trichotomy between beneficiary numbers, food basket quality, and available funding. This insight allowed the supply chain management team to scale up the operation responsibly, while providing donors with evidence-based information about the resources necessary to reach all people in need and the impact of lack and / or untimeliness of resources on beneficiaries and operational costs. In Iraq, optimization enabled the supply chain management team to reduce the cost of general food distribution (supplying 800,000 beneficiaries with their daily nutritional needs) by twelve percent (representing more than 25M USD savings to date), while only providing three kilocalories less. In South Africa, we used optimization to generate sourcing and delivery plans for different contingency scenarios, such as limited local availability and significant price spikes due to El Niño.

The main challenge that was identified is properly covering all forms of cash-based transfers (CBT). Two transfer modalities (Cash and Value Vouchers) require a complex mathematical model that is significantly harder to solve and set up. The initial results are promising however, so in the near future we will develop better solver strategies and/or heuristics to make the model tractable for larger problem sizes. First results indicate that unrestricted Cash and Value Vouchers are best utilized in urbanized areas, whereas a WFP-owned supply chain is more efficient for remote locations.

Next to the WFP applications, we explored to what extent the developed models can support food aid policy and whether they can be applied to other domains. We find that the models can help evaluate the cost of various nutritional targets, and that optimization can help local decision makers to design baskets that address the nutrient gap in the most cost-efficient way. We find that there is no clear dominant strategy (e.g. local always better than international, CBT always better than in-kind, etc.) that is applicable in every context, and recommend that humanitarians are given the flexibility to change their strategy on a case-by-case basis.

With regards to other domains, we find similar challenges in both the humanitarian and non-humanitarian world. National school feeding programs and supply chains for non-food relief items have many similarities with the approach developed for WFP, and in the commercial sector we find large overlaps with the composition of feedstuffs and coffee blends.

Through the development of the models presented in this paper, optimization has become an enabler for cross-functional collaboration at WFP. By integrating key decisions, performance measures, and operational constraints for all major components of WFP's supply chain into one model, WFP is now able to rapidly assess the impact of major decisions on its overall performance. By connecting the composition of the food basket and the supply chain that is necessary to deliver it, we are able to increase the efficiency and effectiveness of WFP operations, ensuring that, with the funds available, WFP can continue to supply life-saving assistance to as many people in need as possible.

Appendix A: Full Model Specification

In Section 3 we focus on describing the differences with standard network flow models. For the sake of completeness, here we provide a full breakdown of parameters and constraints that were used for the applications in Iraq, Yemen, and South Africa. This means that the formulation below is based on Commodity Vouchers, but not yet on Value Vouchers and Cash transfer modalities.

A.1. Sets

We define multiple sub-sets for the nodes in particular to ease the definition of constraints and statistics.

$\mathcal{N}$ = Set of nodes ($i,j \in \mathcal{N}$)
$\mathcal{N}_S$ = Set of Source nodes
$\mathcal{N}_{IS}$ = Set of International Suppliers
$\mathcal{N}_{RS}$ = Set of Regional Suppliers
$\mathcal{N}_{LS}$ = Set of Local Suppliers
$\mathcal{N}_{LM}$ = Set of Local Markets
$\mathcal{N}_P$ = Set of Procurement nodes ($\mathcal{N}_{IS} \cup \mathcal{N}_{RS} \cup \mathcal{N}_{LS} \cup \mathcal{N}_{LM}$)
$\mathcal{N}_{DP}$ = Set of Discharge Ports
$\mathcal{N}_{EDP}$ = Set of Extended Delivery Points
$\mathcal{N}_T$ = Set of Transshipment nodes ($\mathcal{N}_P \cup \mathcal{N}_{DP} \cup \mathcal{N}_{EDP}$)
$\mathcal{N}_{FDP}$ = Set of Final Delivery Points
$\mathcal{K}$ = Set of commodities ($k \in \mathcal{K}$)
$\mathcal{G}$ = Set of food groups ($g \in \mathcal{G}$)
$\mathcal{L}$ = Set of nutrients ($l \in \mathcal{L}$)
$\mathcal{B}$ = Set of beneficiary types ($b \in \mathcal{B}$)
$\mathcal{T}$ = Set of months ($t \in \mathcal{T}$).

Note that the suppliers ($\mathcal{N}_{IS}$,$\mathcal{N}_{RS}$,$\mathcal{N}_{LS}$) provide access to commodities that WFP ships to Final Delivery Points ($\mathcal{N}_{FDP}$) using its network of Discharge Ports ($\mathcal{N}_{DP}$) and Extended Delivery Points ($\mathcal{N}_{EDP}$). Local Markets ($\mathcal{N}_{LM}$) provide beneficiaries with direct access to commodities when they receive a Cash-based transfer. Each Procurement node ($\mathcal{N}_P$) may be linked to multiple Source nodes ($\mathcal{N}_S$), which we use to model procurement restrictions; Procurement nodes represent physical handover points between WFP and a supplier, whereas in our implementation the Source node represents the country of origin for the commodity (which may differ from the country where WFP takes ownership).

A.2. Parameters

Due to the unique nature of long-term food distribution we require some non-conventional parameters, such as nutritional values and feeding days.

α = Conversion rate from Metric Tonnes (mt) to gram (g) (10^6)
ben = The beneficiary type to be optimized ($ben \in \mathcal{B}$)
cap^H_{it} = Handling capacity (in mt) of node $i \in \mathcal{N}_{EDP} \cup \mathcal{N}_{DP}$ in month t
cap^T_{ijt} = Transportation capacity (in mt) from node i to node j in month t
cap^P_{ikt} = Procurement capacity (in mt) of commodity k from source $i \in \mathcal{N}_S$ in month t
$cost_{ijkt}$ = Cost (in \$/mt) of moving commodity k from node i to node j in month t
$days_b$ = Current number of feeding days per month for beneficiaries of type b
dem_{bit} = Number of beneficiaries of type b at node $i \in \mathcal{N}_{FDP}$ in month t
dur_{ij} = Duration (in days) of moving from node i to node j
$group_k$ = The food group that commodity k belongs to ($group_k \in \mathcal{G}$)
hc_i = Handling costs (in \$/mt) at node $i \in \mathcal{N}_{DP} \cup \mathcal{N}_{EDP} \cup \mathcal{N}_{FDP}$
inv_{ikt} = Incoming arrivals (in mt) of commodity k at node $i \in \mathcal{N}_{EDP} \cup \mathcal{N}_{DP}$ in month t
ltp_{ij} = Maximum duration (in days) to supply all FDPs if we procure from source $i \in \mathcal{N}_S$ at node $j \in \mathcal{N}_{\mathcal{P}}$
$nutreq_{bl}$ = Nutritional requirement of beneficiaries of type b for nutrient l
$nutval_{kl}$ = Nutritional value per gram of commodity k for nutrient l
$odoc^{CV}$ = Other Direct Operational Costs rate (in %) for C&V transfers
$odoc^{F}$ = Other Direct Operational Costs rate (in \$/mt) for food transfers
rat_{bk} = Current ration size (gram/person/day) of commodity k for beneficiaries of type b
sc_i = Storage costs (in \$/mt/month) at node $i \in \mathcal{N}_{DP} \cup \mathcal{N}_{EDP}$.

Handling/storage capacities (cap^H_{it}) are only tracked for Discharge Ports (DPs) and Extended Delivery Points (EDPs), although we can easily extend it to other nodes. Basically, the suppliers take care of any handling/storage taking place at Procurement nodes and WFP's cooperating partners do the same at the FDP level. Costs ($cost_{ijkt}$) are captured for all movements through the supply chain network; movements between source nodes and procurement nodes incur procurement costs, other movements are transportation costs. Costs may change over time due to seasonality and stock market movements for procurement and due to rainy seasons for transportation. We split demand (dem_{bit}) per location and beneficiary type, and allow it to change over time so we can phase out or scale up operations. It is important to distinguish between beneficiary types, because their nutritional requirements are vastly different. The needs of a child are very different from those of a newly arrived refugee or a nursing woman for example. The food baskets currently used (rat_{bk}) and their (monthly) feeding days ($days_b$) are pre-defined for each beneficiary type. The model optimizes one of these baskets at a time (ben), and supplies the other beneficiary types with their pre-defined basket. Optimizing all food baskets simultaneously is mathematically possible (and tractable), but rarely done in practice because prioritization across beneficiary types is difficult to quantify. We use inv_{ikt} to model initial inventories and incoming shipments, allowing us to easily integrate the model's decisions with the current status of an operation. Other Direct Operational Costs (ODOC) ($odoc^F$ and $odoc^{CV}$) are surcharges per transfer modality that take care of additional costs that are incurred beyond the procurement, transport, storage, and handling of the commodities (such as milling, packaging, monitoring, reporting, etc.), where $odoc^F$ is usually defined in \$/mt and $odoc^{CV}$ as a percentage surcharge for every \$ transferred to beneficiaries.

A.3. Variables

The two major variables are the ones governing commodity flows (F_{ijkt}) and ration sizes (R_{kt}), the others are mainly used to calculate statistics.

F_{ijkt} = Flow of commodity k from node i to node j in month t (metric tonnes)
R_{kt} = Ration of commodity k in the food basket of month t (gram/person/day)
S_{lt} = Shortfall (slack variable) of nutrient l in month t (fraction)
O_{lt} = Overshoot (slack variable) of nutrient l in month t (fraction)
SFI_{lt} = Binary indicator, 1 if there is a shortfall for nutrient l in month t
P_{ijt} = Binary indicator, 1 if food is procured from source i at procurement node j in month t
LT_t = Maximum lead time of commodities purchased in month t (days).

A.4. Objectives

In order to adequately track and constrain the model's outputs we define dozens of statistics (most of them linear combinations of the variables). Here we focus on the two most important statistics for each of the four goal classes (efficiency, effectiveness, development, and agility):

TOC = Total Operation Cost (Efficiency 1)
FCR = Full Cost Recovery Rate (Efficiency 2)
CAL = Kilocalories (Effectiveness 1)
NVS = Nutritional Value Score (Effectiveness 2)
LOC = Local Procurement (Development 1)
CBT = Cash-Based Transfers (Development 2)
ALT = Average Lead Time (Agility 1)
MLT = Maximum Lead Time (Agility 2).

A multi-objective approach that allowed users to optimize any of the statistics was included in earlier versions of the model, but we discarded it for the sake of solvability. We found that the solution times are significantly higher (3-20x) when incorporating performance measures other than costs into the objective function. For cost objectives the model is highly linear and little branch & bound is necessary to find the optimal solution; this is not the case for other objectives. Because of this big difference in solution speed, it is more practical to use the other statistics as constraints (akin to a goal programming approach).

We define the eight main measures as follows:

$$TOC = \sum_{i\in\mathcal{N}_S}\sum_{j,k,t} cost_{ijkt} \times F_{ijkt} \tag{3}$$

$$+ \sum_{i\notin\mathcal{N}_S}\sum_{j\neq i}\sum_{k,t} cost_{ijkt} \times F_{ijkt} \tag{4}$$

$$+ \sum_{i\in\mathcal{N}_{DP}\cup\mathcal{N}_{EDP}}\sum_{k,t} sc_i \times F_{iikt} \tag{5}$$

$$+ \sum_{i\in\mathcal{N}_{DP}\cup\mathcal{N}_{EDP}}\sum_{j\neq i,k,t} hc_i \times F_{ijkt} + \sum_{j\in\mathcal{N}_{FDP}}\sum_{i,k,t} hc_i \times F_{ijkt} \tag{6}$$

$$+ \sum_{i\in\mathcal{N}_S}\sum_{j\notin\mathcal{N}_{LM}}\sum_{k,t} odoc^F \times F_{ijkt} \tag{7}$$

$$+ \sum_{j \in \mathcal{N}_{LM}} \sum_{i,k,t} cost_{ijkt} \times odoc^{CV} \times F_{ijkt} \tag{8}$$

$$FCR = \frac{\sum_{i,j,k,t} cost_{ijkt} \times F_{ijkt}}{\sum_{i,j,k,t} F_{ijkt}} \tag{9}$$

$$CAL = \frac{\sum_{k} nutval_{k,energy} \times R_{kt}}{|\mathcal{T}|} \tag{10}$$

$$NVS = \frac{\sum_{l,t}(1 - S_{lt})}{|\mathcal{L}|\,|\mathcal{T}|} \times 100\% \tag{11}$$

$$LOC = \frac{\sum_{j \in \mathcal{N}_{LS}} \sum_{i,k,t} F_{ijkt}}{\sum_{j \in \mathcal{N}_{P}} \sum_{i,k,t} F_{ijkt}} \times 100\% \tag{12}$$

$$CBT = \frac{\sum_{j \in \mathcal{N}_{LM}} \sum_{i,k,t} F_{ijkt}}{\sum_{j \in \mathcal{N}_{P}} \sum_{i,k,t} F_{ijkt}} \times 100\% \tag{13}$$

$$ALT = \frac{\sum_{i,j,k,t} dur_{ijk} \times F_{ijkt}}{\sum_{i,j,k,t} F_{ijkt}} \tag{14}$$

$$MLT = \max_{t} LT_t. \tag{15}$$

The usual optimization objective, the Total Operation Cost, consists of the following components: the procurement costs (3), the transportation costs (4), the storage costs (5), the handling costs (6), the ODOC costs for food (7), and the ODOC costs for C&V (8). For WFP this efficiency measure makes the most sense, because a lower cost per beneficiary means that it is able to supply more beneficiaries from the available funding. The other measures usually are used as constraints (goals) in this paper to ensure that the resulting solutions are in line with the objectives of the operation that is being optimized.

A second efficiency measure is the Full Cost Recovery Rate, the average \$/mt that WFP spends on food. The FCR rate varies wildly between countries, and is particularly high when countries are landlocked (e.g. more transport necessary to get the commodities to the beneficiaries) or when WFP is forced to adopt sub-optimal strategy (e.g. sourcing, transfer modality) by donors or the local government. The global average FCR is about a dollar per kg of food.

We measure effectiveness using kilocalories and the Nutritional Value Score (as developed by Ryckembusch et al. (2013)). The number of kilocalories (10) that a food basket provides is very easy to capture, and is still one of the most important and easy to understand performance measures when it comes to measuring effectiveness. For a more holistic view of all macro- and micro-nutrients we use the NVS (11). We know the nutritional profile of the targeted beneficiaries and the nutritional contents of our food basket; the NVS measures how close the current food basket is to the required nutrients. It is defined as the sum of the delivered percentages for each nutrient l, where we truncate the delivered nutrients at 100% (i.e. the score does not improve if you supply more than necessary). In practice, this means that we can just sum the (reversed) shortfalls. The maximum value for this statistic is then the amount of nutrients that are being tracked (11 for the applications in Iraq and Yemen), which is not a very intuitive measure. We adjust the measure slightly by dividing the original NVS by the amount of nutrients. To measure the NVS performance adequately across the time horizon we also average it over the months. The NVS defined here therefore measures the average requirements supplied across all nutrients, with a maximum value of 100% when we supply at least 100% of the requirements for each of the tracked nutrients in each of the months in the time horizon. Note that the NVS does not favor some nutrients over others, and the NVS score may hide the fact that one nutrient is not supplied at all. Because of these limitations, we are working with WFP's Nutrition department to come up with better measures.

Humanitarian organizations prefer using local businesses to source and supply food, as this contributes to the development of the country. As a performance measure we consider the percentage of commodities purchased locally, either through wholesale procurement (12) or commodity vouchers (13). Note that even though the 'percentage purchased' measure is non-linear, upper and lower bounds (goals) for this measure can be included as linear constraints.

For the lead times (which are notoriously hard to capture in linear programming models) we use two proxy measures. One is a maximum lead time proxy variable (15), which is used when WFP needs to respond to a disaster fast (i.e. they need to get a food basket inside the country within x days). In practice however, we generally apply this model to long-term recovery operations. This means that there is always food on its way to the recipient country, there are inventories in the warehouses, commodities with a long lead time are being pre-positioned, etc. Maximum lead times are of less concern in such a scenario. Additionally, if one commodity has a long lead time but the others can be supplied quickly, the maximum lead time does not reflect the agility of the supply chain well. The second proxy measure is therefore the average number of days that it takes for a metric tonne of food to arrive at its destination after being ordered (14). Note that this is again a non-linear measure that can be modeled as a linear constraint.

A.5. Constraints

In the implementation of the optimization model we make a distinction between constraints that always hold and optional constraints. For the minimum working example, we focus on the former. Through a graphical user interface, WFP officers may impose additional context-specific constraints such as beneficiary preferences, long-term supplier agreements, funding restrictions, etc. Most of those (optional) constraints are mathematically trivial though.

$$\sum_j F_{ijkt} = inv_{ikt} + \sum_{(j,i,k,t^*)\in\mathcal{A}(i,k,t)} F_{jikt^*} \;\; \forall i \in \mathcal{N}_T \;\; \forall k,t. \tag{16}$$

Constraint (16) is the traditional 'flow in = flow out' constraint for the transshipment nodes. The parameter inv_{ikt} contains initial inventories and incoming shipments that are already underway when running the tool. $\mathcal{A}(i,k,t)$ are pre-generated sets that contain the indices of all flow variables F_{ijkt} that arrive in transshipment node i with commodity k at time t, generated based on their lead times (dur_{ij}). They are currently defined as follows:

$$\mathcal{A}(i,k,t) = \{(j,i,k,t^*) : t^* + f(dur_{j,i}) = t\} \ \forall i,k,t$$
$$f(t) = \left\lfloor \frac{t}{30} \right\rfloor + \mathbb{1}_{\{t \bmod 30 > 20\}},$$

so we look at the duration of each arc and transform the lead time in days to a lead time in months, using a cut-off point of 20 days. This cut-off point is very context-specific and may need to be revised when applying this algorithm to non-WFP supply chains. In practice we use this formula to generate the parameter initially, but users can override the 'rounded' lead time for any arc themselves if the default cut-off results in poor behavior.

$$\sum_j F_{jikt} \times \alpha = dem_{ben,it} \times days_{ben} \times R_{kt} + \sum_{b \neq ben} dem_{bit} \times days_b \times rat_{bk} \ \forall i \in \mathcal{N}_{FDP} \ \forall k,t. \tag{17}$$

Constraint (17) states that the flow into an FDP must equal its demand (number of beneficiaries times feeding days times daily ration for commodity k plus any additional requirements for that commodity from other beneficiary types). Note that we multiply incoming flows by $\alpha = 10^6$, converting the metric tonnes values to grams (which is how food baskets are specified).

$$F_{ijkt} \leq cap^P_{ikt} \times P_{ijt} \ \forall i \in \mathcal{N}_S \ \forall j,k,t. \tag{18}$$

Constraint (18) specifies that the procurement flow must remain lower than its capacity. Multiplying the capacity by the binary variable P_{ijt} allows us to track maximum lead times, as follows:

$$LT_t \geq P_{ijt} \times ltp_{ij} \ \forall i \in \mathcal{N}_S \ \forall j,t. \tag{19}$$

Recall that ltp_{ij} is a proxy for the lead time when purchasing commodities in node j from source i. We currently generate this value by applying Dijkstra's Shortest Path Algorithm to get the fastest route to each of the FDPs and taking the maximum of those values. The lead time proxy then reflects the shortest time it takes to reach all FDPs when a commodity is procured this way. If this is the only source for this commodity, the proxy value is accurate. If multiple sources are used to supply the same commodity in some month, this proxy is an upper bound. In practice however, it is uncommon to have multiple sources for the same commodity in the same month (the sources do change over time to benefit from price seasonality).

$$\sum_k F_{ijkt} \leq cap^T_{ijt} \ \forall i,j,t. \tag{20}$$

Constraint (20) bounds the flow between nodes, so that the flow cannot exceed the (transportation) capacity of an arc. Note that arc capacities may change over time because of hazardous conditions (security, rainy seasons, etc.).

$$\sum_{jk} F_{jikt} \leq cap^H_{it} \; \forall i \in \mathcal{N}_{DP} \cup \mathcal{N}_{EDP} \; \forall t. \tag{21}$$

Constraint (21) bounds the flow through a node (based on handling capacity rather than transportation capacity).

$$\sum_k nutval_{kl} \times R_{kt} = nutreq_{ben,l} \times (1 - S_{lt} + O_{lt}) \; \forall l, t. \tag{22}$$

Constraint (22) states that the supplied nutrients must cover the requirements. To ensure that the slack variables for shortfalls (S_{lt}) and overshoots (O_{lt}) are working as intended, we need to make sure that they are never positive at the same time. For this we use the shortfall indicator:

$$\begin{aligned} S_{lt} &\leq SFI_{lt} && \forall l, t \\ O_{lt} &\leq (1 - SFI_{lt}) \times 10 && \forall l, t. \end{aligned} \tag{23}$$

Note that S_{lt} is a continuous variable between 0 and 1, whereas the overshoot is only bounded from below (maximum nutrient intakes are not as well defined as minimum nutrient intakes).

When leaving all food basket choices up to the model, it will choose the most cost-effective way to supply all nutrients, regardless of the palatability of the resulting food basket. This may mean that a daily ration consists of half a kg of peas, that it includes large amounts of fortified food (not very palatable and with limited supply), or even 200 milliliters of oil. All these instances, and more, were found during test scenarios. Intuitively, these solutions make no sense and are therefore not credible (despite being the *cost*-optimal way of delivering the nutritional requirements). Through interviews with nutrition experts and food basket designers, the “unwritten rules” for food baskets were brought to light and quantified. The following constraints result in sensible food basket outputs:

$$\begin{aligned} \sum_{k:group_k=g} R_{kt} &\geq minrat_g && \forall g, t \\ \sum_{k:group_k=g} R_{kt} &\leq maxrat_g && \forall g, t \\ R_{\text{Iodised Salt},t} &\geq 5 && \forall t \\ S_{\text{Energy},t} &\leq 0.1 && \forall t \\ S_{\text{Protein},t} &\leq 0.1 && \forall t \\ S_{\text{Fat},t} &\leq 0.1 && \forall t, \end{aligned} \tag{24}$$

Food Group	$minrat_g$	$maxrat_g$
Cereals & Grains	250	500
Pulses & Vegetables	30	130
Oils & Fats	15	40
Mixed & Blended Foods	0	60
Meat & Fish & Dairy	0	40

Table 3 The sensible food basket constraints (in grams per person per day) for each food group. These ration sizes are tailored for General Food Distribution (GFD), which is the bulk of WFP distribution.

with $minrat_g$ and $maxrat_g$ values as given in Table 3. Following these constraints, the rations consist of plenty of staple foods (grains and pulses) to make two or three dishes per day, oil to satisfy fat requirements and for cooking, some fortified foods to prevent malnutrition, and iodised salt to satisfy iodine and sodium requirements. Additionally, we make sure that the three most important nutrients have at most a very small shortfall. Naturally, extra constraints should be added to capture the preferences of the targeted beneficiaries (that may differ a lot between different tribes or countries of origin). Note that these constraints are geared towards so-called General Food Distribution (GFD), which is the ration that we supply to most (if not all) targeted beneficiaries in a country. Beneficiaries with additional needs (children, pregnant and nursing women, newly arrived refugees, etc.) receive the required supplements on top of this GFD basket. These constraints are disabled when we optimize non-GFD baskets.

Lastly, we add some non-negativity constraints:

$$\begin{aligned} F_{ijkt} &\geq 0 \;\forall i,j,k,t \\ R_{kt} &\geq 0 \;\forall k,t \\ S_{lt} &\geq 0 \;\forall l,t \\ O_{lt} &\geq 0 \;\forall l,t \\ LT_t &\geq 0 \;\forall t. \end{aligned} \tag{25}$$

Appendix B: Extension: Location Differentiation approach

To track the choice between CBT and In-Kind for each FDP, we create new binary variables CBT_{it}, which are 1 if FDP i receives CBT assistance in month t, and 0 if it receives In-Kind assistance. We need to make a small adjustment to the demand constraint in the network flow model to account for this choice. For this we need an auxiliary variable (continuous) that keeps track of the In-Kind demand for the activity to be optimized:

$$SAD_{ikt} \geq dem_{it} \times R_{kt} - M \times CBT_{it} \ \forall i,k,t \tag{26}$$

$$OAD_{ikt} = \sum_{b} dem_{bit} \times rat_{bk} \ \forall i,k,t, \tag{27}$$

where M is a sufficiently large number. The Selected Activity Demand (SAD_{ikt}) will be equal to the in-kind demand as per the optimized food basket (R_{kt}) if the FDP chooses In-Kind, and 0 otherwise. The Other Activity Demand (OAD_{ikt}) keeps track of the in-kind demand for the other activities (b) that still need to be supplied. Note that we do not need a variable for this, the explicit specification is just for notational convenience. Similarly, we are ignoring the conversion factors in these formulas (MT to gram conversion and daily ration to monthly ration conversion).

With these new variables defined, the new demand constraint becomes:

$$\sum_{i} F_{ijkt} \geq SAD_{jkt} + OAD_{jkt} \ \forall j,k,t. \tag{28}$$

Note that in practice we may need to limit the number of switches between IK and CBT assistance (to make the solution more implementable), e.g. if we decide to switch it should be at least for n months. A potential solution for this is to track the decision to switch as a(n) (additional) binary variable (by location, by month), where the original CBT_{it} variables are now defined implicitly by these switch-variables. The approach can be modeled using the following constraints:

We need new binary variables ST_{it} ("Switch To") that are equal to 1 if location i swaps transfer modality (to CBT or IK) in month t:

$$ST^{CBT}_{it}, ST^{IK}_{it} \in \{0,1\} \ \forall i,t. \tag{29}$$

For any location, in any month, we can swap at most once:

$$ST^{CBT}_{it} + ST^{IK}_{it} \leq 1 \ \forall i,t. \tag{30}$$

We have to start with one of the two modalities:

$$ST^{CBT}_{it_0} + ST^{IK}_{it_0} = 1 \ \forall i. \tag{31}$$

The original CBT variables are now determined implicitly by the new swap variables:

$$CBT_{it} = CBT_{i,t-1} + ST^{CBT}_{it} - ST^{IK}_{it} \ \forall i,t \tag{32}$$

$$CBT_{it_0} = ST^{CBT}_{it_0} \ \forall i. \tag{33}$$

Ensure that swaps hold for at least n months (n would probably be something like 4-6 months):

$$ST_{it}^{IK} \leq 1 - \sum_{t_*=t-n}^{t} ST_{it_*}^{CBT} \ \forall i, t \tag{34}$$

$$ST_{it}^{CBT} \leq 1 - \sum_{t_*=t-n}^{t} ST_{it_*}^{IK} \ \forall i, t. \tag{35}$$

During implementation of the Location Differentiation approach we found that the solution speed increased drastically. The core model is highly linear in nature, with the binary variables mostly impacting capacities but not the costs (e.g. if during branch and bound a binary variable is 0.7, increasing it to 1 would not affect the cost of the node). With the new binary variables governing the CBT decisions, the problem becomes significantly more difficult to solve. With our default solver, we were unable to find solutions for more than three months of demand. We were able to define a custom solver strategy using Gurobi that makes the problem tractable again even for large planning horizons, but there is still some research and experimentation to be done to reduce the optimality gap to reasonable levels. We will be refining our approach moving forward.

Appendix C: Evolution

Historically, WFP has had a siloed approach to managing and optimizing its supply chains. With humanitarian operations becoming increasingly complex, the need arose for cross-functional innovations that consider WFP's operations holistically. Initial ad-hoc analyses by WFP's Logistics & Supply Chain Development Unit showed that there was a huge potential for savings in Syria by integrating the decisions of the Procurement and Logistics teams—a thorough cross-functional analysis of the sourcing strategy, corridor allocation, and supply chain network design led to savings of two million USD. The analysis was performed by modeling Syria's operation in Excel and using the built-in Excel Solver to find improvements.

From this initial analysis two things became clear. Firstly, while it was possible to optimize sourcing and delivery for a WFP operation using the Excel solver, it was a very labor intensive and cumbersome process. Secondly, it was noted that the sourcing and delivery solution was heavily dependent on the project design, i.e. the food basket that beneficiaries were supposed to receive and the way in which they were supposed to receive it (food, cash, or voucher). Changing even one commodity would have effects that rippled throughout the entire sourcing and delivery strategy. Integrating project design decisions into the Excel solver proved too difficult however, so WFP reached out to its partner universities to develop a model and tool that could cope with the complexity of connecting project design decisions to traditional sourcing and delivery decisions.

In 2014 students from Georgia Tech (USA) and Tilburg University (The Netherlands) developed a tractable prototype that connected the major decisions, and this was then piloted in Syria during a redesign of WFP's operation there. Through optimization we were able to identify concrete savings opportunities, and we recommended a range of food basket adjustments that led to significant savings throughout 2015 (23 million USD). Our collaborative approach to optimization, and our ability to rapidly analyze, quantify, and visualize key decisions resulted in traction for optimization initiatives. Since then we have been focusing on improving the accuracy, scope, and user-friendliness of the prototype, whilst institutionalizing optimization as a management tool and supporting WFP's biggest and most complex operations with ad-hoc analyses using the model.

Moving forward, we are investing heavily in data quality and availability, enabling us to fully reap the benefits of optimization. Future work includes the development of advanced forecasting methods for local market prices, and extending the existing model with robustness against uncertainty in procurement prices and port capacities. In parallel, we are developing a similar model that enables us to optimize the use of scarce resources (such as supplier capacities, corporate stocks, and port capacities) at global and regional planning levels, with the aim of facilitating a monthly planning process (S&OP) at WFP.

Acknowledgments

The work described in this paper is the result of more than five years of iterative development in collaboration with universities and WFP colleagues around the world. Without their enthusiasm and constructive criticism this initiative would have fallen flat a long time ago. Instead, we are able to see it grow every month, with more and more operations looking to us for optimization support when faced with complex operational challenges, and several staff having started using the web application autonomously. We would like to thank everyone who helped us better understand WFP operations, those who invited us to support their operations, and those who have contributed funding so we could continue this work.

References

Altay N, Green W (2006) OR/MS research in disaster operations management. *European Journal of Operational Research* 175(1):475–493.

Andueza S, Vila M, De Peña M, Cid C (2007) Influence of coffee/water ratio on the final quality of espresso coffee. *Journal of the Science of Food and Agriculture* 87(4):586–592.

Apte A (2010) *Humanitarian logistics: A new field of research and action* (Now Publishers Inc).

Arosemena A, DePeters E, Fadel J (1995) Extent of variability in nutrient composition within selected by-product feedstuffs. *Animal Feed Science and Technology* 54(1-4):103–120.

Balcik B, Beamon B (2008) Facility location in humanitarian relief. *International Journal of Logistics* 11(2):101–121.

Balcik B, Beamon B, Smilowitz K (2008) Last mile distribution in humanitarian relief. *Journal of Intelligent Transportation Systems* 12(2):51–63.

Baldi G (2017) Fill the Nutrient Gap analysis: An introduction. Webinar, URL `http://www.implementnutrition.org/wp-content/uploads/2017/12/Fill-the-Nutrient-Gap-SISN-webinar-Nov-2017_FINAL.pdf`.

Ball M, Magnanti T, Monma C, Nemhauser G (1995) *Handbooks in Operations Research and Management Science Volume 7: Network Models* (Elsevier Science).

Beamon B, Kotleba S (2006) Inventory modelling for complex emergencies in humanitarian relief operations. *International Journal of Logistics: Research and Applications* 9(1):1–18.

Ben-Tal A, Do Chung B, Mandala S, Yao T (2011) Robust optimization for emergency logistics planning: Risk mitigation in humanitarian relief supply chains. *Transportation Research Part B: Methodological* 45(8):1177–1189.

Bozorgi-Amiri A, Jabalameli M, Al-e Hashem S (2013) A multi-objective robust stochastic programming model for disaster relief logistics under uncertainty. *OR Spectrum* 35(4):905–933.

Briend A, Darmon N, Ferguson E, Erhardt J (2003) Linear programming: a mathematical tool for analyzing and optimizing children's diets during the complementary feeding period. *Journal of Pediatric Gastroenterology and Nutrition* 36(1):12–22.

Carlson A, Lino M, Gerrior S, Basiotis P (2003) Revision of USDA's low-cost, moderate-cost, and liberal food plans. *Family Economics and Nutrition Review* 15(2):43–51.

Çelik M, Ergun O, Johnson B, Keskinocak P, Lorca A, Pekgün P, Swann J (2012) Humanitarian logistics. *In Tutorials in Operations Research: New Directions in Informatics, Optimization, Logistics, and Production. Hanover, MD: Institute for Operations Research and the Management Sciences* .

Chastre C, Duffield A, Kindness H, LeJeune S, Taylor A (2007) The minimum cost of a healthy diet: Findings from piloting a new methodology in four study locations. *Save the Children UK, London* .

Ergun Ö, Keskinocak P, Swann J (2011) Introduction to the special issue on humanitarian applications: Doing good with good OR. *Interfaces* 41(3):215–222.

FAO, IFAD, UNICEF, WFP, WHO (2018) The state of food security and nutrition in the world 2018. building climate resilience for food security and nutrition. Licence: CC BY-NC-SA 3.0 IGO.

FAO, IFAD, WFP (2015a) Achieving zero hunger: The critical role of investments in social protection and agriculture. URL `http://www.fao.org/publications/card/en/c/de8f4e93-3062-4562-b08b-ecdc1386ec67/`.

FAO, IFAD, WFP (2015b) The state of food insecurity in the world 2015. Meeting the 2015 international hunger targets: Taking stock of uneven progress. URL `http://www.fao.org/publications/card/en/c/c2cda20d-ebeb-4467-8a94-038087fe0f6e/`.

Fleige L, Moore W, Garlick P, Murphy S, Turner E, Dunn M, Van Lengerich B, Orthoefer F, Schaefer S (2010) Recommendations for optimization of fortified and blended food aid products from the United States. *Nutrition Reviews* 68(5):290–315.

Gralla E, Goentzel J, Fine C (2014) Assessing trade-offs among multiple objectives for humanitarian aid delivery using expert preferences. *Production and Operations Management* 23(6):978–989.

Gutjahr W, Nolz P (2016) Multicriteria optimization in humanitarian aid. *European Journal of Operational Research* 252(2):351–366.

Haghani A, Oh S (1996) Formulation and solution of a multi-commodity, multi-modal network flow model for disaster relief operations. *Transportation Research Part A: Policy and Practice* 30(3):231–250.

Holguín-Veras J, Jaller M, Van Wassenhove L, Pérez N, Wachtendorf T (2012) On the unique features of post-disaster humanitarian logistics. *Journal of Operations Management* 30(7):494–506.

Holguín-Veras J, Pérez N, Jaller M, Van Wassenhove L, Aros-Vera F (2013) On the appropriate objective function for post-disaster humanitarian logistics models. *Journal of Operations Management* 31(5):262–280.

Lee H (2004) The triple-a supply chain. *Harvard Business Review* 82(10):102–113.

Lentz E, Barrett C, Gómez M, Maxwell D (2013) On the choice and impacts of innovative international food assistance instruments. *World Development* 49:1–8.

Mitchell S, O'Sullivan M, Dunning I (2011) PuLP: A linear programming toolkit for Python. URL `http://www.optimization-online.org/DB_FILE/2011/09/3178.pdf`.

Novak S, Eppinger S (2001) Sourcing by design: Product complexity and the supply chain. *Management Science* 47(1):189–204.

Özdamar L, Demir O (2012) A hierarchical clustering and routing procedure for large scale disaster relief logistics planning. *Transportation Research Part E: Logistics and Transportation Review* 48(3):591–602.

Pérez-Rodríguez N, Holguín-Veras J (2015) Inventory-allocation distribution models for postdisaster humanitarian logistics with explicit consideration of deprivation costs. *Transportation Science* .

Ram B, Naghshineh-Pour M, Yu† X (2006) Material requirements planning with flexible bills-of-material. *International Journal of Production Research* 44(2):399–415.

Ramberg J (2011) *ICC Guide to Incoterms 2010* (ICC Publishing S.A.).

Rancourt M, Cordeau J, Laporte G, Watkins B (2015) Tactical network planning for food aid distribution in Kenya. *Computers & Operations Research* 56:68–83.

Rawls C, Turnquist M (2012) Pre-positioning and dynamic delivery planning for short-term response following a natural disaster. *Socio-Economic Planning Sciences* 46(1):46–54.

Rottkemper B, Fischer K, Blecken A (2012) A transshipment model for distribution and inventory relocation under uncertainty in humanitarian operations. *Socio-Economic Planning Sciences* 46(1):98–109.

Ryan K, Adams K, Vosti S, Ordiz M, Cimo E, Manary M (2014) A comprehensive linear programming tool to optimize formulations of ready-to-use therapeutic foods: an application to Ethiopia. *The American Journal of Clinical Nutrition* 100(6):1551–1558.

Ryckembusch D, Frega R, Silva M, Gentilini U, Sanogo I, Grede N, Brown L (2013) Enhancing nutrition: A new tool for ex-ante comparison of commodity-based vouchers and food transfers. *World Development* 49:58–67.

Saltzman M (2002) COIN-OR: An open-source library for optimization. *Programming Languages and Systems in Computational Economics and Finance*, 3–32 (Springer).

Van Wassenhove L (2006) Humanitarian aid logistics: Supply chain management in high gear†. *Journal of the Operational Research Society* 57(5):475–489.

Van Wassenhove L, Pedraza Martinez A (2012) Using OR to adapt supply chain management best practices to humanitarian logistics. *International Transactions in Operational Research* 19(1-2):307–322.

Vitzthum O (1999) *Thirty Years of Coffee Chemistry Research*, 117–133 (Springer US).

WFP (2015a) The World Food Programme's operations in Iraq. URL `http://www.wfp.org/countries/iraq`.

WFP (2015b) The World Food Programme's operations in Yemen. URL `http://www.wfp.org/countries/yemen`.

WFP (2018) WFP year in review 2017. URL `https://www.wfp.org/content/year-in-review-2017`.