



SAIIE29 Proceedings, 24th - 26th of October 2018, Spier, Stellenbosch, South Africa © 2018 SAIIE

## A CONCEPT DEMONSTRATOR FOR SELF-ORGANISING DEMAND-DRIVEN INVENTORY MANAGEMENT IN PHARMACEUTICAL SUPPLY CHAINS

M. du Plessis<sup>1</sup>, JH van Vuuren<sup>2\*</sup> & J. van Eeden

<sup>1</sup>Stellenbosch Unit for Operations Research in Engineering (SUnORE)  
Department of Industrial Engineering  
University of Stellenbosch  
[18314937@sun.ac.za](mailto:18314937@sun.ac.za)

<sup>2</sup>Stellenbosch Unit for Operations Research in Engineering (SUnORE)  
Department of Industrial Engineering  
University of Stellenbosch  
[vuuren@sun.ac.za](mailto:vuuren@sun.ac.za)

<sup>3</sup>Department of Industrial Engineering  
University of Stellenbosch  
[jveeden@sun.ac.za](mailto:jveeden@sun.ac.za)

### ABSTRACT

Perennial stock-outs of essential medicines are commonplace in the pharmaceutical supply chains of developing countries. Stock-outs are mainly attributed to a general lack of collective information sharing in pharmaceutical supply chains. In this paper, a computerised agent-based simulation model concept demonstrator is proposed and demonstrated hypothetically as part of a larger drive to establish the value of leveraging information sharing in pharmaceutical supply chains with a view to enhance decision-making. The objective of this paper is to outline the prerequisite research inputs, design requirements and hypothetical implementation of the aforementioned demonstrator. The work reported on in this paper remains a work in progress.

---

<sup>1</sup> The author is enrolled for an M Eng (Industrial) degree in the Department of Industrial Engineering, University of Stellenbosch

<sup>2</sup> The author is a professor in Operations Research at the Department of Industrial Engineering, University of Stellenbosch

<sup>3</sup> The author is a senior lecturer at the Department of Industrial Engineering, University of Stellenbosch

\*Corresponding author

## 1. INTRODUCTION:

Developing nations carry a considerable burden in terms of life-threatening diseases while the treatment of these diseases is significantly complicated by stock-outs and shortages of critical medicines. Stock-outs are preventable, but to successfully thwart medicine stock-outs and their damaging consequences demands a major overhaul in the management of traditional pharmaceutical supply chains of developing countries.

Recent statistics underline the scale of the global medicine stock-outs dilemma. The Global AIDS Response Progress Reporting programme [1], for example, reported that 38 of 108 low- and middle-income countries experienced stock-outs of antiretroviral medicines in 2013. In South Africa, a survey conducted in 2015 by the *Stop Stock-outs Project* consortium [2] revealed that approximately one in four health care facilities suffered from stock-outs of either antiretroviral or tuberculosis medicines during the three-month period preceding the survey. Furthermore, 70% of these stock-outs lasted longer than one month, underlining the supply chain's failure to resolve the root causes of stock-outs rapidly.

The consequences of medicine stock-outs are pervasive and are the most severe on the subsequently untreated patients. Increased drug resistance, aggravation of, or transmission of, disease and even death are some of the harrowing consequences associated with treatment failures [2,3,4]. The impact of stock-outs is particularly harsh on impoverished communities in rural areas which depend solely on public health care services. These poor patients are forced to pay frequent, and costly, visits to their local health care facilities. Regrettably, if they are confronted with stock-outs at these facilities they are turned away and compelled to visit even farther facilities, with no guarantee of medicine availability at these facilities either [5].

The prevailing reasons for pharmaceutical supply chain under-performance in public health sectors include fragmented accountability amongst stakeholders [6], superfluous supply chain complexity [6,7], funding complexities and inadequacies [6,8,9], as well as insufficient inventory management in the face of information shortages and incompetent distribution systems [9,10]. A lack of data capturing and data sharing is, however, attributed as one of the predominant obstacles toward pharmaceutical supply chain improvement in developing nations [6].

Developing countries may not have access to the resources required to implement proper information technology systems in their pharmaceutical supply chains, but the irrefutable advantages of information sharing are plain to see. Sharing supply chain information, such as demand forecasts and inventory levels, across an entire supply network allows organisations to proactively plan for disruptive events, instead of reacting (belatedly) to these events [11]. As a result, supply chains are able to better balance supply and demand, improve stakeholder accountability and ameliorate overall supply chain performance at a reduced cost [6,12,13]. It may be argued that information sharing is a suitable starting point for supply chain reform, because it allows organisations to collaborate to their mutual benefit.

Initiatives utilising the benefits of information sharing in pharmaceutical supply chains have successfully been introduced in some African countries in recent years. A study conducted in 2011 disclosed that at least 60% of stock-outs in the Senegalese contraceptive supply chain occurred at warehouses and health care facilities, despite stock availability at a national level. These problems sprouted from dismal inventory management and poor distribution practices. Upon the implementation of a new system according to which dedicated logisticians actively utilise stock data to manage inventory and curb stock-outs, these stock-outs declined to less than 2% across 140 health care facilities during the first six months [14].

The *SMS for Life* programme, established in 2009, is a web-based reporting system that allows health facility workers to report stock levels on a weekly basis by means of simple SMS messages. This practice of stock level reporting has subsequently alleviated the stock-out predicaments in Kenya and Tanzania and the system is geared for roll-out in more African countries [15,16].

The value of mobile technology in respect of information sharing in pharmaceutical supply chains is also underlined in South Africa's *Stock Visibility Solution* (SVS) programme. The SVS is a mobile phone-based reporting system that allows dispensing clinic staff to report stock levels at regular intervals [17]. The periodic capturing of stock level data allows health care facilities to purposefully manage inventory in a drive to thwart stock-outs.

This paper reports on work in progress that is aimed at, amongst others, the utilisation of information sharing in pharmaceutical supply chains with a view to enhance decision-making.

## 2. PROBLEM DESCRIPTION AND RESEARCH METHODOLOGY

The problem considered in this research involves the performance of conventional pharmaceutical supply chains in developing countries and how these may be improved by the adoption of demand-driven supply chain management principles. In particular, the practice of supply chain information sharing is investigated in order to establish its potential value in respect of effective inventory management. An agent-based simulation model concept demonstrator is developed for use as a test bed to evaluate the efficacy of various inventory replenishment policies within a pharmaceutical supply chain context. The simulation model also accommodates the possibility of modelling user-specified demand scenarios in order to investigate their influence on the effectiveness of inventory replenishment regimes.

The concept demonstrator embraces two modelling paradigms. The first is a *descriptive* paradigm where the model is employed to evaluate the effectiveness of a pre-specified, traditional inventory management policy explicitly embedded in the model. The second paradigm, on the other hand, follows a *prescriptive* approach. According to this paradigm, the user does not select a pre-defined policy as in the case of the descriptive paradigm. Instead, the simulation model is employed to *discover* effective inventory management protocols for the simulated pharmaceutical supply chain network. In other words, an effective inventory management policy is *prescribed* to the user.

The execution of research toward this paper is segmented into three distinct stages. The first stage comprises a brief review of the academic literature relevant to this research project. Thereafter a conceptual framework for capturing the structure of a pharmaceutical supply chain network in a format suitable for use in a simulation modelling environment is established. Finally, a hypothetical example of applying the proposed simulation model is proffered in the third and final stage.

## 3. LITERATURE REVIEW

The literature review in this section consists of four disparate parts, namely a review of the notion of demand-driven supply chain management (in §3.1), a review of the various basic concepts in inventory management (in §3.2), a brief overview of the concepts of self-organisation and emergence (in §3.3) and finally, a review of the machine learning paradigm of reinforcement learning (in §3.4).

### 3.1 Demand-driven supply chain management

A common denominator in the traditional management of supply chains is an emphasis placed on the activities involved with the downstream movement of commodities along a supply chain [18]. Organisations will, for example, streamline their production processes and distribution operations in order to improve the efficiency with which goods are moved downstream in a supply chain. Despite acknowledging the importance of these downstream management activities, advocates of the so-called *demand chain management* (DCM) notion suggest that the focus of this traditional approach is misplaced. DCM is a relatively new concept supporting the notion that end user demand should drive the upstream processes (such as manufacturing and distribution) in a supply chain [18]. As such, the end user is considered as the starting point in a supply chain as opposed to being viewed as the final destination. This particular school of thought arises from the idea that a supply chain ultimately serves to fulfil the needs of the end user. Although products flow downstream toward the end user in a supply chain, it is the end user's demand that should govern the nature of the upstream activities.

Fisher [19] proposed that any supply chain performs two distinct functions. The first is the *physical function* which embodies the physical transformation of raw materials to finished products, and the movement of these goods along a supply chain. The physical function determines a supply chain's efficiency. Manufacturing, delivery and inventory storage outlays are classified as incurring physical costs since they are part of the physical function. The second function is the *market mediation* function and its purpose is to ensure that customer demand is successfully satisfied. Market mediation costs are incurred when supply exceeds demand, or the other way around. In the case of oversupply, excessive stocks may be sold at a loss or even discarded in the case of perishables. Undersupply of stock, on the other hand, reflects lost sales opportunities. In other words, the market mediation function embodies the idea that neither a surplus nor a shortage of stock is desirable in a supply chain.

Fisher furthermore suggested that organisations may prioritise one function at the expense of the other. An organisation subject to predictable demand can, for example, deliberately plan to avoid both a surplus of stock as well as a shortage of stock. Such a position enables a firm to devote its attention to enhancing supply chain efficiency because the market mediation costs are not considered as significant. Organisations faced with unpredictable demand, on the other hand, typically prioritise market mediation costs over physical costs, because they prioritise customer satisfaction irrespective of their attained level of supply chain efficiency.

De Treville *et al.* [20] subsequently defined a *demand chain* as a supply chain in which the market mediation function predominates a supply chain's function to optimise its physical efficiency (the physical function). The adoption of a demand chain approach may seem suitable for a pharmaceutical supply chain because the need to successfully fulfil patient demand is of paramount importance. Pharmaceutical supply chains are, in fact, compelled to pursue a patient service level of 100% because failure to do so would signify the occurrence of stock-outs [21]. Organisations in pharmaceutical supply chains may, however, pursue conflicting objectives. Consider a primary health care facility, such as a clinic, which seeks to minimise medicine stock-outs so as to fully satisfy patient demand. A drug manufacturer upstream, on the other hand, may solely pursue profit maximisation with little regard for the downstream clinic's service level target. This example illustrates that a progression from a conventional pharmaceutical supply chain to a *pharmaceutical demand chain* which prioritises market mediation may not be as simple as it would seem at first.

The practice of information sharing is a powerful enabler of demand-driven supply chain management because it allows organisations to better understand customer demand and to collaborate effectively. The concept of information sharing, also called *supply chain visibility*, refers to the degree according to which supply chain organisations share information that is pivotal to their own activities and which they consider to be of mutual benefit to themselves and other firms in the supply chain [22]. Inventory levels, demand forecasts, order tracking and sales data are examples of information shared in supply chains in order to enhance their collective performance [23]. In the case of a sudden disease outbreak, for example, patient demand for a particular drug may increase considerably over a short period of time. If health care facilities do not carry enough stock to fulfil this increased demand, they set off a reverberating chain of belated, large orders along the supply chain. If the rapid demand increase is, however, made known to upstream facilities promptly through information sharing, they can increase their operations accordingly in anticipation of larger orders.

### 3.2 Inventory replenishment

A significant trade-off faced by inventory managers during their decision-making processes is the trade-off between supply chain responsiveness and efficiency [24]. Carrying large inventories and shortening lead times generally make a supply chain more responsive. The increased responsiveness is, however, traded for significant inventory holding costs and large transport costs, respectively [24].

*Inventory replenishment policies* are typically employed by inventory managers to determine reorder points and reorder quantities. Simchi-Levi *et al.* [25] identified six supply chain variables that play a role in the formulation of an inventory replenishment policy. *Customer demand* is arguably the most significant factor because organisations ultimately strive to fulfil their customers' demand. Secondly, *ordering costs* and *inventory holding costs* are of obvious financial importance. And to ensure the timely receipt of ordered goods, the reorder point should be informed by the *replenishment lead time*, which may not be deterministic. Furthermore, the order quantity may be based on the *current inventory level* of the product in question. Additionally, the *length* of the *planning period* shapes the scope and the nature of inventory management decisions. Finally, the *service level* target may be a determinant of the reorder point and the reorder quantity.

The dynamic nature of supply chains suggests that the parameters of an inventory replenishment policy should be informed by the current state of the supply chain environment with a view to making better decisions. In other words, inventory replenishment protocols should not be too rigid, for otherwise they may fail in the face of changes in the supply chain environment. Owing to the large degree of variability and uncertainty in a supply chain, the inventory management process remains an intricately complex task.

### 3.3 Self-organisation and emergence

The concepts of *self-organisation* and *emergence* are reviewed in order to explore their potential application to inventory management protocols in pharmaceutical supply chains.

De Wolf and Holvoet [26] describe self-organisation as a continuous process in which coordinated organisation manifests itself through the independent behaviour of systems, without any control instructions being imposed from outside the system. 'Organisation' here refers to the presence of a so-called 'structure' that can be of a spatial, temporal or functional nature. Although a self-organising process is void of external control, it does not preclude data inputs from outside the system. A fundamental property of self-organising systems is that they are considered extremely robust and adaptable because they can reproduce 'organisation' in the face of environmental changes [26,27].

The presence of self-organisation may give rise to the related phenomenon of *emergence*. Emergence materialises in a system when the local interactions between its individual constituents culminate at a higher level in the development of a structure (called 'coherent emergents') that is not explicitly represented at a lower level [26,28]. An example of self-organisation and emergence in nature is illustrated in Figure 1. When a colony of ants arrive at a gap in their path, they often use their bodies to build a living bridge without any external supervision or instructions. Each ant follows two simple rules. First, it slows down as it reaches the gap and secondly, it freezes when it feels another ant walking over it. The ants continue in this fashion until they have successfully bridged the gap. Through the ants' self-organising behaviour, a living bridge *emerges*. The bridge may be classified as an *emergent* because no individual ant is representative of the bridge. The bridge is only formed at a higher level through the local interactions between the ants at a lower level.



**Figure 1: A living bridge emerges from the self-organising behaviour of ants [29,30].**

It may be argued that effective, externally coordinated inventory management in a pharmaceutical supply chain is extremely difficult, or even impossible, given the myriad of supply chain variables that influence inventory management decisions. Self-organisation (a process void of external control) is therefore explored as an alternative means of coordinating inventory management. A self-organising supply chain, by implication, is void of any form of centralised control and each facility manages its own inventory exclusively. In this research project, we investigate the conjecture that local coordination between facilities may lead to the *emergence* of a greater structure where the global supply chain functions as a coordinated system in respect of inventory management.

Consider a simple example of a self-organising pharmaceutical supply chain in which each facility in the chain 'organises' itself with a view to prevent stock-outs locally. These facilities, in other words, are autonomous and actively manage their own inventory in pursuit of an 'organisation' in which stock-outs are prevented. Additionally, in an information sharing supply chain, these facilities may utilise the available information to inform their inventory management decisions accordingly. There is, however, no explicit coordination between the facilities in the supply chain. If a storage depot is, for example, perturbed by a drastic demand increase, the facility may 'reorganise' itself by increasing its order quantities. Emergence may subsequently occur in the supply chain as a set of management policies prescribing reorder points and reorder quantities in pursuit of effective inventory management.

### 3.4 Reinforcement learning

*Reinforcement learning* is a branch of machine learning where a learning agent learns behaviour in an environment through interaction with the environment [31]. The premise of reinforcement learning rests on the idea that if a particular action yields desirable results, the inclination to repeat the same action is *reinforced* [32]. This closely relates to the learning process followed by humans and animals. A new-born elephant, for example, tries many strategies and fails often before it can stand upright. Over time, the baby elephant learns

to avoid the actions that caused it to fall down and it hones the skills that proved more fruitful in pursuit of its goal to stand upright.

A fundamental characteristic of reinforcement learning is that a learning agent can evaluate the desirability of its actions according to a numerical reward signal, but it is not told which actions to take in order to improve its performance [33,34]. The reward signal is expressed in terms of a pre-specified goal that is pursued by the agent. Hence, a learning agent has to attempt many different strategies by itself in order to learn what behaviour maximises its reward signal. This learning process can informally be described as learning through *trial-and-error*. The new-born elephant, for example, is said to learn through trial-and-error which strategies prove to be more successful.

The reinforcement learning approach is commonly described in terms of an *agent* and an *environment* [33,34]. The *agent* is the learning actor which interacts with its *environment* in order to learn about the environment. A *state* describes the situation in the environment at any given time instant. At discrete time steps, the agent is presented with an array of actions from which it can choose. The selected action influences the environment and the environment provides feedback to the agent in the form of a reward signal and by transitioning into a new state. The reward signal is employed to evaluate the immediate reward received for the selected action. Notably, the agent's objective in reinforcement learning is to maximise its cumulative reward and, occasionally, these rewards may be significantly delayed [33,34]. The reinforcement learning cycle repeats itself many times and, over time, the agent learns to map different situations to particular actions that yield desirable results. A schematic of this learning paradigm is shown in Figure 2. The outcome of a reinforcement learning process may be described as a *look-up table* [33]. This table maps all possible environment states to appropriate actions that have proved to maximise the agent's reward during the learning process.

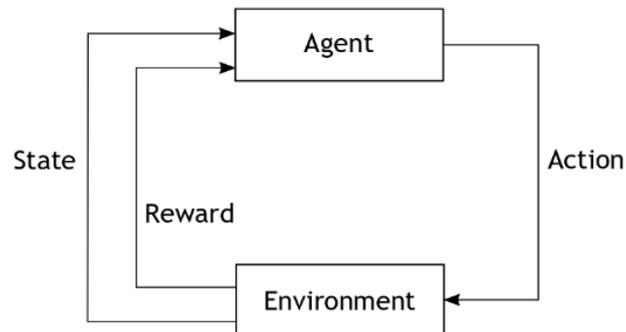


Figure 2: The reinforcement learning cycle (adapted from [33]).

Reinforcement learning forms the cornerstone of the prescriptive paradigm of the proposed concept demonstrator, as discussed in §1. Each facility type in a pharmaceutical supply chain (*i.e.* manufacturer, warehouse, clinic, *etc.*), is trained as a reinforcement learner. A unique look-up table of state-action pairs is subsequently generated for each facility type that can be utilised by the facility to inform decision-making on a daily basis aimed at improving performance indicators aligned with various management objectives.

#### 4. CONCEPTUAL INPUT FRAMEWORK FOR PHARMACEUTICAL SUPPLY CHAIN MODELLING

As discussed in §1, an agent-based pharmaceutical supply chain simulation model is put forward in this paper. The simulation model follows a generic design so as to enhance its flexibility and potential value for decision makers in pharmaceutical supply chains. According to this generic design, the simulation model receives a user-specified supply chain structure as input. This input structure, called an *input framework*, should sufficiently capture the level of abstraction required to model a pharmaceutical supply chain mathematically, as per the purposes of this research. This section is devoted to a description of the conceptual design of such an input framework. This framework is not, however, exhaustive, but simply serves as a point of departure for the development of a comprehensive, well-rounded input framework.

The input framework should capture facility-specific information, product-specific information, as well as inventory management parameters and product demand profiles. A list of attributes that captures the high-

level, facility-specific information is shown in Table 1. This information set describes the size of the supply chain network and its prevailing facility types.

**Table 1: Facility-specific information provided as input.**

Attribute	Description
Facility name	Common name used to identify the facility
Location	Spatial information
Facility type	The nature of a facility's operations. For example: Manufacturer, storage facility, hospital, clinic, etc.
Tier	Specification of the relevant echelon in the supply chain
Storage capacity	The facility's total storage capacity for the purposes of the simulation

The simulation model's generic design allows for the inclusion of user-specified pharmaceutical products. The product-specific information required extends only to the name of the product and its shelf life, as shown in Table 2.

**Table 2: Product-specific information provided as input.**

Attribute	Description
Product name	Commonly used product name
Shelf life (if perishable)	Shelf life duration (from date of manufacture)

A traditional *from-to* matrix can be used to capture the connections between facilities in a supply chain. These connections indicate the flow of goods between facilities. The matrix is of size  $n \times n$  where  $n$  denotes the total number of facilities in the supply chain. If product units flow *from* facility  $i$  to facility  $j$ , the  $(i, j)^{\text{th}}$  entry in the matrix adopts a value of 1, or a value of 0 otherwise. The facility names are obtained from Table 1. An example of a *from-to* matrix for a supply chain comprising three facilities is shown in Table 3. Facility A, for example, distributes goods to Facilities B and C. Facility B, on the other hand, distributes only to Facility C and Facility C, in turn, does not distribute any inventory to other facilities.

**Table 3: An example of a  $3 \times 3$  from-to matrix provided as input.**

	Facility A	Facility B	Facility C
Facility A	-	1	1
Facility B	0	-	1
Facility C	0	0	-

As discussed in §1, the descriptive paradigm of the concept demonstrator allows the user to evaluate the performance of pre-specified inventory replenishment policies. In order to facilitate this paradigm, the user is required to specify the parameters of these policies as part of the input framework. The relevant inventory management parameters required to model the inventory management processes are presented in Table 4. According to the table, the user can specify parameters for a continuous review policy, or for a periodic review policy. Table 4 may, of course, be extended to include more inventory management policies.

**Table 4: Inventory management parameters provided as input.**

Attribute	Description
Ordering facility	Facility name from Table 1
Product	Product name from Table 2
Starting inventory	Product quantities available at the start of the simulation
Minimum order quantity	If applicable
Maximum order quantity	If applicable
Reorder point	For a continuous review policy
Reorder quantity	For a continuous review policy
Review interval	For a periodic review policy
Order-up-to level	For a periodic review policy
Lead time (days)	As a function of order size, may be stochastic

In order to model the manufacturing operations of manufacturers in the supply chain, the production process characteristics have to be captured in the format as shown in Table 5. This information is applicable to manufacturers only.

**Table 5: Manufacturing information provided as input.**

Attribute	Description
Manufacturer	Facility name from Table 1
Product	Product name from Table 2
Starting inventory	Product quantities available at the start of the simulation
Production trigger	Signal that triggers the initiation of the production process
Production rate	Expressed in number of batches per time unit
Batch size	The number of units in a single batch

Finally, forecasted demand data and actual demand data for the simulation period can be provided as input. Users may provide either synthesised data or actual data as input. The nature of the demand data required is elucidated in Table 6.

**Table 6: Demand data provided as input.**

Attribute	Description
Facility	Name from Table 1
Product	Product name from Table 2
Actual demand	Daily demand for each simulated day
Forecasted demand	Daily forecasted demand for each simulated day

A potential implementation of this input framework is demonstrated by means of a hypothetical example in the following section.

## 5. HYPOTHETICAL EXAMPLE

The objective of this section is to integrate the salient elements of the literature review in §3 with the input framework of §4 in order to demonstrate how it may be applied in practice by means of a small hypothetical example.

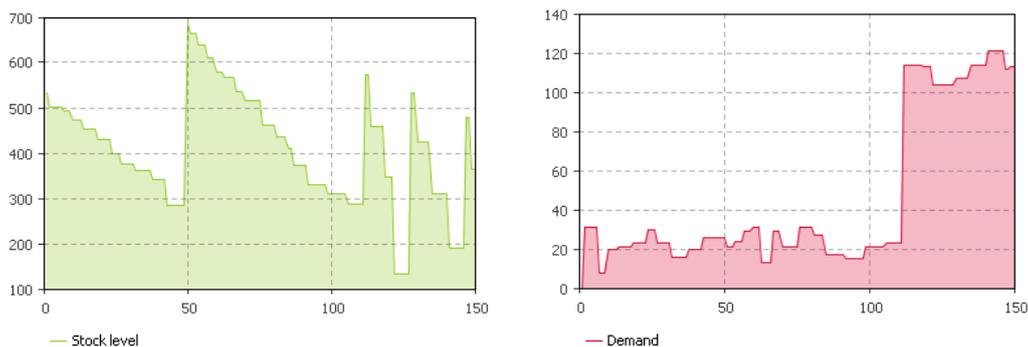
Consider a simple pharmaceutical supply chain comprising a single manufacturer, a single warehouse and two clinics. The supply chain facilitates the flow of *Painstill* drugs from the manufacturer to the clinics, via the warehouse. Currently all four facilities in the supply chain employ traditional continuous review replenishment policies.

The Painstill supply chain has suffered from large-scale stock-outs in recent months and it has been decided to investigate avenues for improving its inventory management practices. In particular, the value of supply chain information sharing is of interest and how it may inform effective inventory management in the supply chain with a view to minimise stock-outs. The management team has turned to the simulation model proposed in this paper to support their decision-making processes. After populating the input framework of §4, the management team decides to employ both the descriptive and prescriptive modelling paradigms.

### Descriptive paradigm

According to the descriptive paradigm, a pre-defined replenishment policy is selected for each facility from a list of possible policies. The management team decides to continue with a continuous review policy for each facility. Using the simulation model, the management team can now experiment with different parameter values (reorder points and reorder quantities) for each facility in order to determine how they may improve the effectiveness of the continuous review policies. The operation of the supply chain is now simulated according to the specified parameters. The movement of Painstill units through the supply chain, and charts denoting information such as inventory levels may be displayed during the simulation model execution. At the end of the simulation run, a set of key performance indicator values are provided as output. Examples of suitable key performance indicators may include attained service levels, the number of stock-outs per facility, the average stock-out duration per facility, as well as the procurement costs and inventory holding costs incurred by each facility. An example of a graphic denoting a facility's stock level data and demand data are shown in Figure 3.

Suppose that the sudden demand increase observed at day 111 is attributed to an unexpected disease outbreak. It is evident from the stock level graph that the facility carried enough stock to fulfil the increased demand initially. The stock level has, however, declined to a minimum of 140 units on day 122. The relevant decision-makers may therefore infer that it is best to increase the facility's order quantities in the face of a similar demand increase in order to negate the possibility of stock-outs.



**Figure 3: A graph of a facility's stock level over time (left) and the same facility's corresponding demand over time (right). A sudden demand increase is observed at day 111 due to a sudden disease outbreak.**

#### Prescriptive paradigm

The prescriptive paradigm integrates the concepts of demand-driven supply chain management, self-organisation and reinforcement learning as a means to prescribe effective replenishment policies for the modelled supply chain network. The output of this paradigm is therefore a set of key performance indicator values, as well as a set of policies prescribing reorder points and reorder quantities for each facility. 'Self-organisation' in this context means that each facility makes its own inventory decisions, as alluded to in §3.3. The prescriptive paradigm lends itself to a large degree of scalability because of the supply chain's self-organising property. The effects of the local interactions between facilities may ripple outward until a form of organisation is achieved and maintained across the entire supply network. The size of the supply chain therefore has little influence on the model complexity.

The user should explicitly define the level and degree of information that may be shared and used by other facilities in order to facilitate their decision-making processes. The management team can, for example, explore the effect of sharing both clinics' stock level data with the warehouse. This level of visibility may presumably allow the warehouse to increase its inventory proactively should the clinics' stock levels start to decline rapidly in response to a sudden demand increase. The management team can, however, also investigate the value of sharing the clinics' stock level data with both the warehouse and the manufacturer, provided that the manufacturer also has visibility over the warehouse's stock level data. Intuitively, it may be argued that the increased level of supply chain visibility may be accompanied by improved overall supply chain performance. Comparing these two scenarios at the hand of the simulation model may elucidate whether the larger investment in information technology required for the second scenario is, in fact, justified in respect of the simulation results. It may, for example, be that the increased level of information sharing of the second scenario does not significantly improve on the effectiveness delivered in the first scenario. The simulation model may prove extremely useful for similar comparison analyses.

Once the information sharing structure has been configured in the model, reinforcement learning may be applied as a mechanism to discover self-organising replenishment heuristics. Since facilities of the same type share common operational characteristics, all the instances of a particular type of facility can use the same look-up table. Therefore, only one agent can represent each facility type, be trained and only generate one look-up table per agent. For the Painstill supply chain, a manufacturer agent, a warehouse agent and a clinic agent have to be trained as reinforcement learners, respectively. Notably, both clinics use the single look-up table generated by the clinic agent. The learning process is expected to be computationally expensive, but this is an offline process which is executed *a priori*.

The respective reward functions of these agents may be specified by the user. A typical reward function should reward desirable actions, such as the successful fulfilment of customer demand. For undesirable scenarios, such as the occurrence of stock-outs, product expirations or large inventory carrying costs, agent punishment should result in the form of a negative reward signal. Notably, the reward function need not be particular to one agent, but rewards may be shared amongst agents. Rewards may be shared between a warehouse and a clinic (the clinic orders from this warehouse), for example, in an attempt to enhance their collective performance.

A *state* should describe all the information that is visible to a learning agent at any given time instant during the learning process. The state may include a number of dimensions, such as the agent's current inventory level and its forecasted demand. If, for example, the Painstill manufacturer has visibility over the warehouse's stock level, the manufacturer agent's state space would include the warehouse's stock level data. In other words, the size of an agent's state should be informed by the level and degree of information sharing. At discrete time steps, each agent should decide whether to place a new order for Painstill drugs, or not. If the agent decides to order, the order quantity should also be selected. The simulation model should replicate many possible scenarios in the supply chain and the reinforcement learning algorithm should experiment with different order strategies during the process. Over time, each agent learns which inventory management decisions (actions) yield the most reward within a particular situation (state), and these are documented in a look-up table. In other words, each agent learns when to place a new order and how much to order, given a certain state.

The management team may use the results of both the descriptive and prescriptive paradigms to provide them with decision support in pursuit of their drive to improve the efficiency of the Painstill supply chain. The simulation results may, for example, provide insight as to which facilities suffer the most from stock-outs and why under different demand scenarios. The management team can also identify those facilities at which product expirations occur most frequently and subsequently learn which replenishment policies may prevent them. Additionally, cost-related key performance indicators may provide an indication of the financial feasibility of a particular policy. Finally, the prescriptive paradigm may elucidate which type of information should be shared, and with whom, for the best outcome in respect of the management team's various objectives.

## 6. CONCLUSION

The objective of this paper was fourfold. First, to establish the background context of this research project as the inventory management methodologies of demand-driven pharmaceutical supply chains. Secondly, to provide a brief overview of the relevant academic literature which serves as a basis for the work conducted in this research. Thereafter, a preliminary, conceptual input framework for pharmaceutical supply chain modelling was developed. Finally, the potential use of the planned proposed concept demonstrator was illustrated by means of a hypothetical example. It is important to stress that the work described in this paper is not concluded. This paper serves as a prelude to a larger research project in which the value of self-organisation and information sharing in the pharmaceutical supply chains of developing countries is explored and quantified. It is acknowledged that these concepts may not be readily compatible with existing supply chain infrastructure and resources in these countries. This research, however, aims to elucidate whether self-organisation is a suitable instrument for pharmaceutical supply chain reform.

## REFERENCES

- [1] World Health Organization. 2014. *Global update on the health sector response to HIV, 2014*, World Health Organization.
- [2] Doctors Without Borders (MSF), RuDASA, RHAP, TAC, SECTION27 and SAHIVSoc. 2016. *2015 Stock outs National Survey: Third annual report - South Africa*, Stop Stock Outs Project.
- [3] Harries, A.D., Schouten, E.J., Makombe, S.D., Libamba, E., Neufville, H.N., Some, E., Kadewere, G. and Lungu, D. 2007. Ensuring uninterrupted supplies of antiretroviral drugs in resource-poor settings: An example from Malawi, *Bulletin of the World Health Organization*, 85(2), pp. 152-155.
- [4] Nicholson, A., English, R.A., Guenther, R.S. and Claiborne, A.B. 2013. *Developing and strengthening the global supply chain for second-line drugs for multidrug-resistant tuberculosis*, The National Academic Press, Washington (DC).
- [5] Hodes, R., Price, I., Bungane, N., Toska, E. and Cluver, L. 2017. How front-line healthcare workers respond to stock-outs of essential medicines in the Eastern Cape Province of South Africa, *South African Medical Journal*, 107(9), pp. 738-740.

- [6] Yadav, P. 2015. Health product supply chains in developing countries: Diagnosis of the root causes of underperformance and an agenda for reform, *Health Systems & Reform*, 1(2), pp. 142-154.
- [7] Humphreys, G. 2011. Vaccination: Rattling the supply chain, *Bulletin of the World Health Organization*, 89(5), pp. 324-325.
- [8] Cameron, A., Ewen, M., Ross-Degnan, D., Ball, D. and Laing, R. 2009. Medicine prices, availability, and affordability in 36 developing and middle-income countries: A secondary analysis, *The Lancet*, 373(9659), pp. 240-249.
- [9] Kangwana, B.B., Njogu, J., Wasunna, B., Kedenge, S.V., Memusi, D.N., Goodman, C.A., Zurovac, D. and Snow, R.W. 2009. Malaria drug shortages in Kenya: A major failure to provide access to effective treatment, *American Journal of Tropical Medicine and Hygiene*, 80(5), pp. 737-738.
- [10] Bateman, C. 2013. Drug stock-outs: Inept supply-chain management and corruption, *South African Medical Journal*, 103(9), pp. 600-602.
- [11] Griffin, P.M., Nembhard, H.B., DeFlitch, C.J., Bastian, N.D., Kang, H. and Munoz, D.A. 2016. *Healthcare systems engineering*, John Wiley & Sons, Hoboken (NJ).
- [12] Frohlich, M.T. and Westbrook, R. 2002. Demand chain management in manufacturing and services: Web-based integration, drivers and performance, *Journal of Operations Management*, 20(6), pp. 729-745.
- [13] Ripin, D.J., Jamieson, D., Meyers, A., Warty, U., Dain, M. and Khamsi, C. 2014. Antiretroviral procurement and supply chain management, *Antiviral Therapy*, 19(3), pp. 79-89.
- [14] Daff, B.M., Seck, C., Belkhatay, H. and Sutton, P. 2014. Informed push distribution of contraceptives in Senegal reduces stockouts and improves quality of family planning services, *Global Health: Science and Practice*, 2(2), pp. 245-252.
- [15] Barrington, J., Wereko-Brobby, O., Ward, P., Mwafongo, W. and Kungulwe, S. 2010. SMS for Life: A pilot project to improve anti-malarial drug supply management in rural Tanzania using standard technology, *Malaria Journal*, 9(1), pp. 1-9.
- [16] Githinji, S., Kigen, S., Memusi, D., et al. 2013. Reducing stock-outs of life saving malaria commodities using mobile phone text-messaging: SMS for Life study in Kenya, *PLOS One*, 8(1), pp. 1-8.
- [17] Mezzanine. 2018. *Stock Visibility Solution*, [Online], [Cited June 2018], Available from: <https://www.mezzanineware.com/svs/>.
- [18] Langabeer, J.R. and Rose, J. 2002. *Creating demand driven supply chains: How to profit from demand chain management*, Spiro Press, London.
- [19] Fisher, M. 1997. What is the right supply chain for your product?, *Harvard Business Review*, 75(2), pp. 105-116.
- [20] De Treville, S., Shapiro, R.D. and Hameri, A.P. 2004. From supply chain to demand chain: The role of lead time reduction in improving demand chain performance, *Journal of Operations Management*, 21(6), pp. 613-627.
- [21] Uthayakumar, R. and Priyan, S. 2013. Pharmaceutical supply chain and inventory management strategies: Optimization for a pharmaceutical company and a hospital, *Operations Research for Health Care*, 2(3), pp. 52-64.
- [22] Barratt, M. and Oke, A. 2007. Antecedents of supply chain visibility in retail supply chains: A resource-based theory perspective, *Journal of Operations Management*, 25(6), pp. 1217-1233.
- [23] Lee, H.L. and Whang, S. 2000. Information sharing in a supply chain, *International Journal of Manufacturing Technology and Management*, 1(1), pp. 79-93.
- [24] Chopra, S. and Meindl P. 2013. *Supply chain management: Strategy, planning and operation*, 5<sup>th</sup> Edition, Pearson Education Limited, Essex.
- [25] Simchi-Levi, D., Simchi-Levi, E. and Kaminsky P. 2000. *Designing and managing the supply chain: Concepts, strategies, and cases*, McGraw-Hill, New York (NY).
- [26] De Wolf, T. and Holvoet, T. 2005. *Emergence versus self-organisation: Different concepts but promising when combined*, pp. 1-15 in Brueckner, S.A., Serugendo, G.D.M., Karageorgos, A. and Nagpal, R. (Eds). *Engineering self-organising systems: Methodologies and applications*, Springer, Berlin.
- [27] Heylighen, F. 2001. The science of self-organization and adaptativity, *Encyclopedia of Life Support Systems*, 5(3), pp. 253-280.
- [28] Odell, J. 2002. Agents and complex systems, *Journal of Object Technology*, 1(2), pp. 35-45.
- [29] Hartnett, K. 2018. *The simple algorithm that ants use to build bridges*, [Online], [Cited June 2018], Available from: <https://www.quantamagazine.org/the-simple-algorithm-that-ants-use-to-build-bridges-20180226/>.
- [30] Manohar, V. 2018. *Unity is strength*, [Online], [Cited June 2018], Available from: <https://www.shutterstock.com/image-photo/unity-strength1011406435?src=BbM2Y7v-2FfNGqgiTyr5TKw-1-74>.



SAIIE29 Proceedings, 24th - 26th of October 2018, Spier, Stellenbosch, South Africa © 2018 SAIIE

- [31] Kaelbling, L.P., Littman, M.L. and Moore, A.W. 1996. Reinforcement learning: A survey, *Journal of Artificial Intelligence Research*, 4, pp. 237-285.
- [32] Sutton, R.S., Barto, A.G. and Williams, R.J. 1992. Reinforcement learning is direct adaptive optimal control, *IEEE Control Systems*, 12(2), pp. 19-22.
- [33] Marsland, S. 2009. *Machine learning: An algorithmic perspective*, CRC Press, Boca Raton (FL).
- [34] Sutton, R.S. and Barto A.G. 1998. *Reinforcement learning: An introduction*, The MIT Press, Cambridge (MA).