

Dynamic allocation of scarce resources under supply uncertainty

Sarang Deo

Kellogg School of Management, Northwestern University, Evanston, IL 60208, s-deo@kellogg.northwestern.edu

Charles J. Corbett

UCLA Anderson School of Management, Los Angeles, CA 90095, charles.corbett@anderson.ucla.edu

We present a model of dynamic resource allocation in a setting where continuity of service is important and future resource availability is uncertain. The paper is inspired by the challenges faced by HIV clinics in resource-limited settings in the allocation of scarce HIV treatment among a large pool of eligible patients. Many clinics receive insufficient supply to treat all patients and the supply they do receive is highly uncertain. This supply uncertainty, combined with the clinical importance of an uninterrupted treatment throughout patients' life, requires the clinics to make a trade-off between providing access to treatment for new patients and ensuring continuity of treatment for current patients. Setting aside other aspects of the treatment rationing problem, we model the decisions of a clinic facing this trade-off using stochastic dynamic programming. We derive sufficient conditions under which the optimal policy coincides with the clinically preferred policy of prioritizing previously enrolled patients. We use numerical examples to investigate the impact of supply uncertainty on the performance of enrollment policies used in practice. We also discuss how our model applies to other intertemporal resource allocation decisions such as that faced by non-profit organizations where continuity of service is crucial to meeting the organization's social objective, or that faced by an entrepreneur who wants to attract new customers without reducing service quality to existing customers.

Key words: inventory rationing, HIV, supply uncertainty

1. Introduction

Many organizations have to strike a balance between offering their services to more customers and maintaining quality of service for existing customers. This trade-off becomes particularly acute when the organization faces uncertainty in the supply of a key resource. The specific example of this trade-off that inspired this paper is that faced by HIV clinics in resource-constrained settings (specifically in sub-Saharan Africa) related to allocation of antiretroviral drugs (ARVs). The challenges arise

not only because of aggregate shortage of ARVs (WHO, 2005b) but also because of the uncertainty associated with the drug supply (ITPC, 2005; BBC News, 2004; IRINNews.org, 2005). The resulting stock-outs at clinics cause interruption of treatment for patients, which can lead to adverse clinical outcomes such as treatment failure and drug resistance (IOM, 2005; WHO, 1998) and increased mortality (El-Sadr et al., 2006). The clinics have to decide between starting treatment for new patients today but incurring a higher risk of treatment interruptions in future and minimizing the risk of future treatment interruptions for current patients by restricting access for new patients today.

In this paper, we explore this trade-off by modeling the clinics' treatment allocation decision using a stochastic dynamic programming model. At the beginning of each period, the clinic receives a drug shipment of uncertain quantity over which it has little or no control. The clinic faces a deterministic demand from patients who have been treated in previous periods. We use the term "current patients" to refer to this group. In addition, the clinic can initiate treatment for individuals from a large pool of previously untreated individuals, who then become "new patients" in this period and join the pool of "current patients" thereafter. Knowing the available inventory of drugs and the demand from current patients at the beginning of each period, the clinic decides how many current and new patients to treat in each period to maximize the total expected quality adjusted life years (QUALYs) of its patients over the planning horizon.

We prove the existence of an optimal policy for the general formulation of our resource allocation problem by showing its equivalence with a multi-location multi-period inventory problem. We then use the characteristics of the resource-constrained settings (demand for drugs is much larger than the supply) to simplify the problem formulation and characterize the structure of the optimal policy. We derive sufficient conditions under which it is optimal to prioritize treatment for current patients over new patients, a recommended policy in many scale up guidelines (WHO, 2003). We also show, under these conditions, that the optimal policy is characterized by a threshold that corresponds to an upper bound on new enrollments in each period. Any excess inventory, after treating all the previously enrolled (current) patients and enrolling new patients up to this bound,

is carried to the next period as a safety stock. For the finite horizon formulation, this enrollment bound (or equivalently the safety stock) is state-dependent and dynamic. We also provide numerical illustrations to compare the performance of this enrollment policy with those used in practice.

The primary objective of our model is to explore the core trade-off described above and to characterize the structure of the optimal policy and the impact of supply uncertainty on it. Hence, for analytical tractability, we make certain simplifying assumptions regarding the flow of patients. Although these assumptions combined with the paucity of accurate data limit the applicability of our model as a decision support tool for clinics in practical settings, it does help to provide insight into the resource allocation problem, which in turn can help in building a more detailed simulation based decision-support tool.

While we use the context of HIV treatment in resource-constrained settings to motivate the development of the model and discuss the main insights, the core trade-off is present in several other contexts. For instance, nonprofit organizations that provide community services related to health-care (drug rehabilitation centers), education (after-school programs), housing (homeless shelters) etc., often face substantial uncertainty about the quantity and timing of future funding. At the same time, continuity of service is often important: withdrawing these types of social services once might make the intended beneficiaries less receptive these services in future. These nonprofit organizations have to balance their desire to expand programs with the need to maintain uninterrupted service to existing beneficiaries.

Similarly, consider the dilemma faced by many entrepreneurs: They want to attract new customers to grow their business but they also need to maintain a constant quality of service to their existing customers. For instance, a catering firm is usually better off not taking an order than taking it and then underperforming by not having enough capacity to do the job well. But if they never take a new customer, their business will never grow. Entrepreneurs often face considerable uncertainty with respect to the resources (both staff and financial) they have available at any point in time. When should such an entrepreneur decide to initiate a long-term relationship with a new customer, in light of this uncertainty?

This paper is organized as follows. In section 2, we describe the operational challenges of delivering ARVs in resource-constrained settings in some more detail. Section 3 provides a brief review of the various related streams of literature and outlines our contribution to them. The model formulation is described in section 4. The optimal policy and its properties are derived in section 5. Section 6 describes heuristics which are either used in practice or have practical appeal. We provide numerical illustrations to compare these heuristics with the optimal policy in section 7. Section 8 provides concluding remarks. Proofs for all the theoretical results are provided in the technical appendix in the e-companion to this paper.

2. Operational challenges in HIV drug supply

Antiretroviral drugs (ARVs) can neither cure nor prevent HIV infection and AIDS but can considerably reduce mortality and morbidity in HIV+ patients (Palella et al., 1998; Walensky et al., 2006). Approximately 20% of the eligible patients in Sub-Saharan Africa and other developing regions of the world were receiving ARVs in 2006 despite a multifold increase in long-term funding by donor agencies, dramatic reduction in drug prices and increased awareness as a result of the WHO's "3 by 5 initiative".

Operational bottlenecks such as limited capacity for inventory control and storage, quantification and reporting, and security of commodities have been cited among the most important reasons for this slow progress (GAO, 2006). A major consequence of these bottlenecks is the uncertainty in the supply of drugs received by the clinics. This supply uncertainty leads to periodic stock-outs of drugs as reported in various parts of the world including India, Russia, Dominican Republic (ITPC, 2005), Nigeria (Ekong et al., 2004), South Africa (BBC News, 2004), Kenya (IRINNews.org, 2002) and Swaziland (IRINNews.org, 2005). In addition to this anecdotal evidence, logistics assessment surveys commissioned by the United States Agency for International Aid (USAID) and conducted by John Snow Inc. (JSI) provide systematic evidence of stock-outs and supply uncertainty in Zimbabwe (Nyenwa et al., 2005) and Tanzania (Amenyah et al., 2005). These stock-outs cause interruption of treatment for patients which could lead to drug resistance and / or treatment failure (Bartlett,

2006). Oyugi et al. (2007); van Oosterhout et al. (2005) provide systematic clinical evidence of this phenomenon in Uganda and Malawi respectively. Drug shortages due to supply interruptions have also resulted in the death of patients in South Africa (Health Systems Trust, 2005).

However, this underlying supply uncertainty has received minimal attention in the quantification and forecasting tools used by clinics or in the academic literature. An exception is Yadav (2007) who discusses the uncertainty in procurement lead-times for essential commodities in Zambia and its impact on drug stock-outs in the presence of budget constraints. Current practice relies on informal guidelines for deciding a safety stock to manage this uncertainty and there is an urgent need for formal models to quantify the safety stock required to optimally manage the underlying supply uncertainty while scaling up treatment (Daniel, 2006). Holding excess safety stock would mean blocking scarce funds in nonproductive assets and slowing treatment expansion, while too little safety stock could result in extremely undesirable stock-outs and treatment interruptions.

Deo (2007) presents a framework for the broader issues involved in treatment scale up including how treatment, prevention and diagnoses are interlinked via patient behavior and disease epidemiology. Here we focus only on the impact of supply uncertainty on treatment programs as this effect itself is poorly understood; future work can extend this to include the impact of treatment on prevention through modified patient behavior, reduced viral load and increased patient willingness to be tested. We also do not consider the impact of current program outcomes on future resource availability.

3. Literature Review

There is a vast operations research literature on dynamic allocation of scarce resources with applications in diverse areas such as inventory rationing, yield management and new product development. However, to our knowledge, this is the first model of dynamic resource allocation that explicitly considers the issue of service continuity in the face of uncertainty regarding future resource availability. Specifically, our model extends existing inventory rationing models by explicitly modeling the conversion of customers from one segment (previously unserved) to the other (previously served)

and by including supply uncertainty. This paper contributes to the literature on resource allocation for HIV / AIDS interventions, which has predominantly focused on prevention. In the context of allocating ARVs in resource-constrained settings, it complements the existing qualitative discussion by providing a quantitative framework for rationing treatment between new and current patients at the clinics.

Our model is related to the models of inventory rationing among customer classes of differing priorities (Topkis, 1968; Evans, 1969; Nahmias and Demmy, 1981; Ha, 1997a, 1997b; de Véricourt et al., 2002; Deshpande et al., 2003). The optimal allocation policy in these models consists of a threshold or reservation level for each segment such that it is optimal to stop serving a segment if the on-hand inventory drops below the threshold associated with that segment. Frank et al. (2003), Zhang and Sobel (2001) and Gupta and Wang (2007) study inventory rationing schemes where demand from one segment has to be met while demand from the other segment can be either backlogged or lost at a penalty. The customer segments in these models are unrelated, i.e., customers do not move from one segment to the other as a result of receiving service. In contrast, in our model, the two customer segments are inherently related as customers from one pool (previously untreated) are moved permanently to another pool (previously treated) as a result of the treatment decisions. Olsen and Parker (2006) model flows of customers across segments but do not consider rationing.

Considerable work has been done in combining epidemiological models and optimal control theory to study dynamic allocation of resources in the case of HIV epidemics (Lasry et al., 2006; Zaric and Brandeau, 2001; Richter et al., 1999; Kaplan and Pollack, 1998). However, our work differs from these papers on several dimensions. First, the focus of these models is on prevention interventions whereas we focus on treatment programs for HIV. Second, these papers do not explicitly consider uncertainty in resource availability (drug supply). Third, the key trade-off faced in these models is between efficiency and equity while the key trade-off in our model is between access (enrolling more patients) and quality (providing uninterrupted treatment to enrolled patients) which is exacerbated by the uncertainty in the future supply of drugs. Lastly, most of these models (with the exception

of Lasry et al., 2006) focus on developed countries whereas our model is most relevant to resource-constrained settings.

Recent qualitative discussions on rationing strategies for ARVs in developing countries focus on the issue of “which” new patients to enroll (Rosen et al., 2005; Bennet and Chanfreau, 2005). However, it pays inadequate attention to two important characteristics of treatment scale-up: (i) patients once enrolled have to be treated continuously through their life and (ii) there is a variability in supply of drugs in addition to the aggregate shortage. We complement this literature by incorporating these characteristics in a quantitative model that to help clinics decide “how many” new patients to enroll when accurate information about the future supply of drugs is not available.

Our model could also contribute to the non-profit literature as non-profit organizations critically depend on external funding sources which are known to be highly unreliable and variable (Gronbjerg, 1992). Also, in organizations such as homeless shelters and drug rehabilitation programs, it is critical to maintain continuity of service to current beneficiaries while expanding service to new beneficiaries (Scott, 2003). de Véricourt and Lobo (2006) study the allocation of an organization’s assets among “mission” and “revenue” customers so as to maximize the total social benefit. However, not all non-profit organizations can engage in for-profit activities due either to lack of requisite skills or the domain of their activities (Dees, 1998; Foster and Bradach, 2005). Our model could be adapted to complement the model in de Véricourt and Lobo (2006) by incorporating the uncertainty in future funding and penalty of service interruption.

4. Model formulation (Finite horizon)

In this section, we present the formal model for the decision problem of an individual clinic in a resource-constrained setting that wants to maximize the expected total discounted quality adjusted life years (QALYs) of its patients. The specific objective function is general enough to apply to the non-profit and entrepreneurial contexts referred to earlier. Let N denote the length of the problem horizon consisting of discrete decision making epochs $n = N, N - 1, \dots, 3, 2, 1$ where $n = 0$ denotes

the end of the horizon. For the most part, we focus our analysis on the finite horizon formulation ($N < \infty$). We discuss the infinite horizon formulation and prove the existence of the optimal policy in Appendix A and C respectively.

4.1. Drug supply

Current distribution systems for ARTs in resource-constrained countries consist of central medical depots in the capital city from where the drugs are “pushed” to the sites of health care delivery (WHO, 2005a; WHO, 2003). The ultimate goal is to move to a more formal system where clinics order drugs based on their requirements. However, inadequate inventory management skills at the clinics make this transition from “push” to “pull” both difficult and slow (JSI, 2006; WHO, 2003). Also, due to a weak transport infrastructure, the drug supply actually received at the clinics is uncertain.

To reflect this situation, we model the supply of drugs as exogenous but stochastic; order quantity is not a decision variable for the clinic. Extending the model to include an ordering decision by clinics would be interesting but appears to be analytically intractable, in particular as the link between the orders placed and quantity received in practice is very unclear. Let \tilde{Z}_n be independently (not necessarily identically) distributed random variables that denote the supply of drugs received by the clinic in period n with cumulative distribution $\Phi_n(\cdot)$ and support on $[z_n^L, z_n^U]$. Thus at the time of deciding the number of new and current patients to treat in period n , the clinic does not know the actual quantity of drugs it will receive in the future periods ($1 \leq i < n$) but only knows the cumulative distribution $\Phi_i(\cdot)$. Let I_n and W_n denote the inventory of drugs before and after receiving the supply in period n so that $W_n = I_n + \tilde{Z}_n$.

4.2. Patients

The demand for drugs comes from patients who have been diagnosed as HIV+ and are eligible for treatment based on the national guidelines. We model the demand at the clinic to be composed of two pools of patients: $y_{n,t}$ denotes the number of current (previously treated) patients and $y_{n,u}$ denotes the number of previously untreated patients who are eligible for treatment at the time of

deciding treatment allocations. We further divide each of these pools into two subcategories based on whether they receive treatment in the current period or not and assign a quality of life (QOL) score to each of these four categories. Thus s_{tt} denotes the QOL score for previously treated patients who receive treatment in the current period, and s_{tu} denotes the QOL score for previously treated patients whose treatment is interrupted in the current period. Similarly s_{ut} denotes the QOL score for previously untreated patients who receive treatment for the first time in the current period and s_{uu} denotes the QOL score for previously untreated patients who do not receive treatment in the current period. Since we are not modeling the health status of individual patients within each subcategory, s_{tt}, s_{tu}, s_{ut} and s_{uu} could be considered as average QOL scores for each of the four subcategories.

The decision on which segments of patients to prioritize, based on socioeconomic and clinical characteristics (such as CD4+ count which reflects the state of the immune system of patients) is made at a national level (Bennet and Chanfreau, 2005), so the clinic faces demand from patients which are relatively homogenous on these attributes. Moreover, previous research has shown that the health status (as captured by the CD4+ count and QOL scores) of patients receiving treatment becomes reasonably homogenous after around six months of treatment (Cleary et al., 2006). Hence we model the pool of new and current patients for an individual clinic to be homogenous along these attributes and focus on the average health status of each pool for our analysis. An alternative formulation that models the health status of individual patients would be unamenable to analytical treatment.

4.3. System dynamics

In each period n , knowing the available inventory W_n and the demand from previously treated and untreated patients $y_{n,t}$ and $y_{n,u}$ respectively, the clinic decides on the number of current and new patients to treat, denoted by $x_{n,t}$ and $x_{n,u}$ respectively. After the treatment decisions, the inventory of drugs drops to $I_{n-1} = W_n - x_{n,t} - x_{n,u}$ and the pool of previously untreated patients reduces to $y_{n,u} - x_{n,u}$. At the end of each period, a deterministic fraction β_t of all current patients and

β_u of all new patients survive through to period $n - 1$ and the remaining patients exit the system (through death, relocation or otherwise). New patients enter the system at a rate α proportional to the number of new patients remaining at the end of the period. Thus $\beta_u x_{n,u}$ denotes the number of patients who were initiated on treatment in period n and survived, thus adding to the pool of current patients in period $n - 1$, and $\alpha(y_{n,u} - x_{n,u})$ denotes the number of patients who enter the pool of eligible patients at the beginning of the next period $n - 1$. The resulting system dynamics are given by the following set of equations:

$$y_{n-1,t} = \beta_t y_{n,t} + \beta_u x_{n,u} \quad (1)$$

$$y_{n-1,u} = (\beta_u + \alpha)(y_{n,u} - x_{n,u}) \quad (2)$$

$$W_{n-1} = W_n - x_{n,t} - x_{n,u} + \tilde{Z}_{n-1} = I_{n-1} + \tilde{Z}_{n-1} \quad (3)$$

We make several assumptions regarding the patient dynamics in our model to make the model tractable. First, the average survival rates β_t and β_u and the rate of entry in to the pool of new patients α are assumed to be deterministic. Including uncertainty in the survival rates is non-trivial but our model can be adapted to include stochastic survival rates that are independent of the uncertainty in the drug supply. Second, β_t is assumed to be exogenous and constant over time. It does not depend on the proportion of current patients who receive treatment in each period. Including this dependence would imply modeling the health status of individual patients, which as discussed earlier is analytically intractable. Third, we assume that the rate of entry in to the pool of untreated patients depends only on the size of the untreated pool and not on the treated pool. This is reasonable if the size of the treated pool is much smaller than the untreated pool and/or if treated patients do not contribute to new infections due to psycho-social (reduced risky behavior) or physiological (reduced viral load) reasons. A more exact formulation would need to include the epidemiological dynamics of new infections, which is beyond the scope of the current model. With respect to the drug supply, we assume that the drugs are not perishable. This is true for all the drugs used in the first line of treatment, which is our focus here.

4.4. Objective function

As mentioned earlier, the objective of the clinic is to maximize the total quality adjusted life years (QALYs) for the patient population over the planning horizon N . While QALYs have been traditionally used for clinical decision making at an individual level, there has been a recent trend to use QALYs at a population level to evaluate alternate policy measures (Zenios et al., 2000; Richter et al., 1999). See Loomes and McKenzie (1985) for a detailed discussion of the related issues.

Furthermore, since we assume that the underlying composition of these categories does not change over the problem horizon, our QOL parameters are time invariant. Thus the objective of the clinic for a finite horizon N is given by

$$\max_{x_{n,u} \geq 0, x_{n,t} \geq 0} E \left[\sum_{n=1}^N \delta^{N-n} h_n(x_{n,t}, x_{n,u}) \right] \quad (4)$$

where δ is a single period discount factor and h_n is the single period reward function given by:

$$\begin{aligned} h_n(x_{n,t}, x_{n,u}) &= s_{tt}x_{n,t} + s_{tu}(y_{n,t} - x_{n,t}) + s_{ut}x_{n,u} + s_{uu}(y_{n,u} - x_{n,u}) \\ &= (s_{tt} - s_{tu})x_{n,t} + (s_{ut} - s_{uu})x_{n,u} + s_{tu}y_{n,t} + s_{uu}y_{n,u} \end{aligned} \quad (5)$$

An analog of this objective function in continuous time has been used previously to analyze resource allocation decisions in the context of epidemics, for instance Brandeau et al. (2003). To denote that treatment in the current period has positive benefit for both current and new patients, it is reasonable to assume $(s_{tt} - s_{tu}) > 0$ and $(s_{ut} - s_{uu}) > 0$.

5. Analysis

Using the above building blocks, we now formulate the decision problem for the clinic as a finite horizon stochastic dynamic program. Let $V_n(W_n, y_{n,t})$ denote the maximum net benefit from the clinic's treatment decisions for the remaining n periods until the end of the horizon. Then the recursive equation corresponding to the clinic's decision problem in period n is given by:

$$\begin{aligned} V_n(W_n, y_{n,t}, y_{n,u}) &= \max_{x_{n,u} \geq 0, x_{n,t} \geq 0} \{h_n(x_{n,t}, x_{n,u}) + \delta E[V_{n-1}(W_{n-1}, y_{n-1,t}, y_{n-1,u})]\} \\ & \text{s.t. (1), (2) and (3)} \end{aligned} \quad (6)$$

$$x_{n,t} \leq y_{n,t}$$

$$x_{n,u} \leq y_{n,u}$$

$$x_{n,u} + x_{n,t} \leq W_n$$

where $h_n(x_{n,t}, x_{n,u})$ is as defined in 5.

Equations (1), (2) and (3) are the system dynamics described earlier. The next two constraints state that the number of current and new patients treated cannot be more than the total number of current and new patients in that period respectively. The last constraint states that the total number of treatments delivered in period n is limited by the available inventory. Let $f_n(x_{n,t}, x_{n,u}) \triangleq h_n(x_{n,t}, x_{n,u}) + \delta E[V_{n-1}(W_{n-1}, y_{n-1,t}, y_{n-1,u})]$ denote the maximand in (6). The existence of the optimal policy is shown in the following Proposition.

PROPOSITION 1. (i) $V_n(W_n, y_{n,t}, y_{n,u})$ is jointly concave in its arguments.
(ii) $f_n(x_{n,t}, x_{n,u})$ is jointly concave in its arguments.
(iii) The optimal policy for (6) is characterized by the existence of a vector $S^* = (I_n^*, y_{n,t}^*, y_{n,u}^*)$ such that it is optimal to move from the state vector $(W_{n+1}, y_{n+1,t}, y_{n+1,u})$ to S^* at the beginning of period n before supply is received, provided such a move is feasible. If this move is not feasible, then it is optimal to reach to the vertex of the feasible region that is closest to vector S^* .

The proof for the first two parts of Proposition 1 follows from the usual inheritance properties of dynamic programs. Thus in each period n , (6) is a concave optimization problem over a set of linear constraints. The third part follows directly from the multi-location multi-period inventory model of Karmarkar (1981). Therefore the optimal policy for our problem is similar to the modified base-stock policy for inventory problems with side-constraints. However, this structure of the optimal policy does not directly lead to a computational method and hence does not permit comparisons with the treatment allocation policies followed in practice, which is our primary objective. Next, we use the characteristic of the resource-constrained settings to reduce the dimensionality of the state-space of our problem, make the problem analytically more tractable and provide sharper characterization of the optimal policy.

5.1. Model reformulation

At present, drug supply in many developing countries is enough to reach only a small fraction of all the eligible patients. WHO does not expect to reach its target of universal access until at least 2010 (WHO, 2006) and PEPFAR targets to put only 2 million people on treatment by 2009 (PEPFAR, 2006). Moreover, due to increasing new infections, according to the epidemiological update by WHO (2005b): “*Indications are that some of the treatment gaps will narrow further in the immediate years ahead, but not at the pace required to effectively contain the epidemic*”. This resource-constrained condition is also symptomatic of most non-profit organizations and entrepreneurial firms. To reflect this situation, we assume that $y_{n,u} > W_n \forall n$, i.e., the demand for drugs will outstrip supply over the planning horizon. In Appendix B, we provide a sufficient condition in the form of an upper bound on the support of distributions $\Phi_n(\cdot)$ for this to hold. If $y_{n,u} > W_n \forall n$, the constraint $x_{n,u} \leq y_{n,u}$ is never tight and the feasible set $x_{n,u}$ does not depend on $y_{n,u}$, hence reducing the dimensionality of the state-space. Thus, substituting $y_{n,u} = (\beta_u + \alpha)^{N-n} y_{N,u} - \sum_{i=n+1}^N (\beta_u + \alpha)^{i-n} x_{i,u}$ and leaving out the constant term involving $y_{N,u}$ in (4), the objective function becomes

$$\max_{x_{n,u} \geq 0, x_{n,t} \geq 0} E \left[\sum_{n=1}^N \delta^{N-n} \left((s_{tt} - s_{tu}) x_{n,t} + \left(s_{ut} - s_{uu} \sum_{i=0}^{n-1} (\delta (\beta_u + \alpha))^i \right) x_{n,u} + s_{tu} y_{n,t} \right) \right]$$

Defining $s_{uu}^n \triangleq s_{uu} \sum_{i=0}^{n-1} (\delta (\beta_u + \alpha))^i$ and $\hat{h}_n(x_{n,t}, x_{n,u}) = (s_{tt} - s_{tu}) x_{n,t} + (s_{ut} - s_{uu}^n) x_{n,u} + s_{tu} y_{n,t}$, the formulation in (6) can be modified as

$$\begin{aligned} V_n(W_n, y_{n,t}) = & \max_{x_{n,u} \geq 0, x_{n,t} \geq 0} \left\{ \hat{h}_n(x_{n,t}, x_{n,u}) + \delta E[V_{n-1}(W_{n-1}, y_{n-1,t})] \right\} \\ & s.t. (1) \text{ and } (3) \\ & x_{n,t} \leq y_{n,t} \\ & x_{n,u} + x_{n,t} \leq W_n \end{aligned} \quad (7)$$

Thus, we have used the resource-constrained condition to reduce the dimensionality of the state-space, but at the expense of introducing non-stationarity in one of the problem parameters, s_{uu}^n . For the remainder of the paper, we shall focus on this formulation.

5.2. Two-period model

First, we solve the most simple, yet still nontrivial instance of (7) with $N = 2$ to highlight some of the difficulties associated with the formulation. Specifically, we show that for some parameter

values the optimal solution can appear inconsistent with the current clinical practice of prioritizing current patients over previously untreated patients. This can mean that clinical practice is not optimal, or that the parameters prevailing in practice do not allow this inconsistency. In the absence of accurate information about these parameter values in practice, we cannot determine whether current practice is suboptimal or not. We derive additional conditions on the parameters that are sufficient to ensure that the optimal policy is structurally consistent with clinical recommendations. We use the same notation as in section 4. The only source of uncertainty in this two-period problem is the quantity of drugs received at $n = 1$ whose cumulative distribution is denoted by $\Phi_1(\cdot)$. At $n = 2$, the available inventory W_2 and size of the current patient pool $y_{2,t}$ are known. The clinic has to decide the number of new and current patients to be treated in period 2 and 1 denoted by $x_{2,t}$, $x_{2,u}$, $x_{1,t}$, $x_{1,u}$. We need the following definitions in order to characterize the optimal policy:

$$k_1 = \frac{(s_{tu} - s_{uu}(1 - \delta\alpha)) - (1 - \delta\beta_u)(s_{tt} - s_{ut})}{\delta\beta_u((s_{tt} - s_{tu}) - (s_{ut} - s_{uu}))}$$

$$k_2 = \frac{s_{ut} - s_{uu}(1 + \delta(\beta_u + \alpha)) + \delta(s_{tt} - (s_{ut} - s_{uu})(1 + \beta_u))}{((s_{tt} - s_{tu}) - (s_{ut} - s_{uu}))(1 + \delta\beta_u)}$$

PROPOSITION 2. *The optimal policy for the two period problem is given by:*

Case I $s_{ut} - s_{uu} < s_{tt} - s_{tu}$:

$$x_{2,t}^* = \min\{y_{2,t}, [W_2 - \eta]^+\}, \quad x_{2,u}^* = \min\{W_2 - \min\{y_{2,t}, [W_2 - \eta]^+\}, \theta\}, \quad x_{1,t}^* = \min\{y_{1,t}, W_1\},$$

$$x_{1,u}^* = [W_1 - y_{1,t}]^+ \text{ where}$$

$$\eta = \min\left\{x_{2,u} \geq 0: \frac{\partial f_2}{\partial x_{2,u}} \leq \frac{\partial f_2}{\partial x_{2,t}} \Big|_{x_{2,u} + x_{2,t} = W_2}\right\} = \left[\frac{\Phi_1^{-1}(k_1) - \beta_t y_{2,t}}{\beta_u}\right]^+$$

$$\theta = \min\left\{x_{2,u} \geq 0: \frac{\partial f_2}{\partial x_{2,u}} \Big|_{x_{2,t} = y_{2,t}} \leq 0\right\} = \left[\frac{\Phi_1^{-1}(k_2) + W_2 - (1 + \beta_t)y_{2,t}}{1 + \beta_u}\right]^+$$

Case IIA $s_{ut} - s_{uu} > s_{tt} - s_{tu}$ and $s_{uu} - s_{tu} < \frac{(s_{ut} - s_{uu}) - (s_{tu} - s_{tt})}{\delta\beta_u}$:

$$x_{2,t}^* = 0, \quad x_{2,u}^* = W_2, \quad x_{1,t}^* = [W_1 - y_{1,u}]^+, \quad x_{1,u}^* = \min\{W_1, y_{1,u}\}$$

Case IIB $s_{ut} - s_{uu} > s_{tt} - s_{tu}$ and $s_{uu} - s_{tu} > \frac{(s_{ut} - s_{uu}) - (s_{tt} - s_{tu})}{\delta\beta_u}$:

$$x_{2,t}^* = \min\{W_2, y_{2,t}\}, \quad x_{2,u}^* = [W_2 - y_{2,t}]^+, \quad x_{1,t}^* = 0, \quad x_{1,u}^* = W_1$$

Proposition 2 shows that even for the simple two period problem, and even with the reduced state-space that follows from the resource constraint, the structure of the optimal policy is quite complicated. The optimal policy depends not only on the relative values of the QOL parameters but also on the current system state.

In case I, $s_{tt} - s_{tu} > s_{ut} - s_{uu}$, i.e., the single period value of treatment is higher for previously treated patients. Then, as expected, it is optimal to prioritize current patients for the last period $n = 1$. However, the prioritization is not unambiguous for $n = 2$ and depends on the relative values of the state variables (W_2 and $y_{2,t}$) and the thresholds (η and θ). For instance, consider the case when $W_2 - y_{2,t} < \eta < W_2$ and $\eta < \theta$, where η and θ also depend on QOL parameters. In this intermediate case, $x_{2,t}^* = W_2 - \eta$ and $x_{2,u}^* = \eta$ i.e., new patients are enrolled before all the current patients have been treated since this allocation equalizes the marginal return from the treatment of both patient segments.

In cases IIA and IIB, the condition $s_{ut} - s_{uu} > s_{tt} - s_{tu}$ implies that the single period value of treatment is higher for previously untreated patients. In both these cases, it is optimal to prioritize treatment for new patients in the last period $n = 1$. However, this condition is not sufficient to maintain the prioritization for both periods. Prioritization for new patients is maintained earlier for $n = 2$ only if s_{uu} is not too much greater than s_{tu} (case IIA). However, in case IIB, if s_{uu} is sufficiently greater than s_{tu} , the prioritization is reversed and it is optimal to prioritize current patients for $n = 2$.

5.3. Prioritization of current patients

Recent clinical studies have clearly shown that even structured treatment interruptions can drastically increase the mortality and morbidity in HIV+ patients (El-Sadr et al., 2006). Hence, continuous treatment for life is the recommended practice once a patient is enrolled for treatment (IOM, 2005). In terms of the two-period model described above, this is equivalent to requiring $x_{2,u}^* > 0$ only if $x_{2,t}^* = y_{2,t}$, i.e., previously untreated patients $y_{2,u}$ are treated after all previously treated patients $y_{2,t}$ have been treated, irrespective of the state variables W_2 and $y_{2,t}$. Here we investigate

what additional conditions on the parameter values will guarantee that the optimal policy has this structure. It is clear that $s_{tt} - s_{tu} > s_{ut} - s_{uu}$ is required for prioritization of current patients in the last period, $n = 1$. However, Proposition 2 noted that this is not sufficient to ensure prioritization of current patients for $n = 2$. Note that if $\eta = 0$ then the optimal policy in Case I would be equivalent to prioritizing current patients for $n = 2$. In the next proposition, we build on this idea to derive sufficient conditions on the parameter values so that prioritization of current patients is optimal in all periods.

PROPOSITION 3. *It is optimal to prioritize current patients over new patients in every period if the following conditions are satisfied:*

$$(C1) \quad (s_{tt} - s_{tu}) > (s_{ut} - s_{uu})$$

$$(C2) \quad s_{tu} (1 - \delta(\beta_t - \beta_u)) < (s_{tt} - s_{ut})(1 - \delta\beta_t) + s_{uu} (1 - \delta(\beta_u + \alpha - \beta_t))$$

Moreover, the optimal solution under these conditions is given by

$$x_{n,t}^* = \min \{y_{n,t}, W_n\} \text{ and } x_{n,u}^* = \min \left\{ \theta_n, [W_n - y_{n,t}]^+ \right\} \text{ where}$$

$$\theta_n = \min \left\{ x_{n,u} \geq 0 : \left. \frac{\partial f_n}{\partial x_{n,u}} \right|_{x_{n,t}=y_{n,t}} \leq 0 \right\}.$$

Condition (C1) states that, on average, the one-period health benefit from treating a previously treated patient is higher than the one-period health benefit from treating a new patient. This is reasonable since not treating a previously treated patient can lead to development of drug resistance and failure of the treatment regimen while the one-period health benefit from treating a new patient might be limited as the main benefits accrue after continued treatment. Condition (C2) is less easy to interpret, but it helps to consider a few special cases. For $\delta = 0$, (C2) reduces to (C1) confirming that for a single period problem (C1) is sufficient to guarantee prioritization of current patients. For $\beta_t = \beta_u$ and $s_{tt} = s_{ut}$, (C2) reduces to $s_{tu} < s_{uu} (1 + \delta\alpha)$ which implies that the average QOL score of patients with interrupted treatment should not be too much higher than the average QOL score of unenrolled patients. Intuitively, if (C2) is not satisfied then the penalty from treatment interruption is not high enough to warrant prioritization of current patients over new patients. The

fact that it is not known when this condition holds in practice combined with the the substantial effect the status of this condition has on optimal treatment policies, highlights the need for clinical and epidemiological research to obtain more insight into these factors.

Next, we analyze the implications of the optimal solution for the case when (C2) holds. Treating an additional new patient in the current period has three effects. First, there is an immediate social benefit of the treatment, s_{ut} . Second, it reduces the available inventory to be carried over to the next period. Third, it increases the pool of current patients in the next period by β_u (adjusting for exit of patients). The uncertainty regarding the supply of drugs in the future periods implies that this new patient increases the chance that a previously enrolled patient might go untreated in the future. Thus the optimal policy balances the expected penalty of interrupting the treatment of previously treated patients in the future periods with the immediate benefit of treating an additional new patient. The quantity $[W_n - \theta_n - y_{n,t}]^+$ could be interpreted as the safety stock to be carried over to protect current patients against the future supply uncertainty. Before moving further, it is important to also clarify what the optimal policy does not imply. The optimal policy (and the others discussed later) is based on population averages and provides guidelines to clinics on the rate at which they can enroll new patients. It cannot help the clinic make a decision to treat or not treat a specific patient.

It is instructive to contrast Proposition 3 with the results from traditional models of ordering and rationing inventory across multiple customer classes. The optimal policy in these models usually involves thresholds, one corresponding to each customer segment, such that it is optimal to not serve a particular segment if the on-hand inventory falls below the corresponding threshold (Topkis, 1969; Ha, 1997a, 1997b; de Véricourt et al., 2001). This enables the decision maker to carry enough safety stock to protect against the uncertain demand from high value customers in future periods. However, in our model, since the supply is uncertain and beyond the control of the clinic, a safety stock is built and maintained by restricting the enrollment of new patients. This serves to protect the “higher value” patients (previously treated patients) from any unanticipated supply interruptions in future periods. The next proposition provides more insight into the structure of the threshold θ_n

for the special case of $\beta_t = \beta_u$. This is a reasonable condition during the initial phase of treatment scale-up when the full impact of treatment on survival of patients has not been realized.

PROPOSITION 4. *Let $\beta_t = \beta_u$. If $\theta_t > 0$, then*

(i) $\theta_n = \psi_n(W_n) - y_{n,t}$ and

(ii) $F(z_{n-1}) \succ_{fsd} G(z_{n-1})$ implies that $\theta_n(F) > \theta_n(G)$.

When $\beta_t = \beta_u$ and $\theta_n > 0$, the system states in period $n - 1$ depend only on $y_{n,t} + \theta_n$. Hence, what matters is how much total inventory was dispensed rather than how this was divided between new and current patients. Part (i) of Proposition 4 shows that in this case, the optimal policy is equivalent to carrying over a fraction of the available inventory to the next period as a safety stock. This fraction is given by $\frac{W_n - (\theta_n + y_{n,t})}{W_n} = \frac{W_n - \psi_n(W_n)}{W_n}$. Since the supply is stochastic and dynamic, this fraction is not a constant but depends on the available inventory in that period. Part (ii) shows that if $\beta_t = \beta_u$ and the next period's drug supply is stochastically greater, then everything else being equal, the safety stock would be reduced, or equivalently, the enrollment cap θ_n would increase.

6. Enrollment heuristics

Proposition 3 describes that under conditions (C1) and (C2) it is optimal to prioritize the treatment for current patients and to enroll new patients up to a threshold θ_n . While prioritization of current patients is generally followed in practice, the enrollment policies that are actually implemented have much simpler structure than this threshold policy, which involves solving the recursive dynamic program (7). In this section, we describe two such heuristics that have practical appeal and contrast them with the optimal prioritization policy from Proposition (3); in the next section we report numerical illustrations to evaluate the heuristics.

6.1. Safety-stock policy

A common approach recommended in real life scale up situations is to maintain a safety stock equivalent to a few months of demand to buffer against supply uncertainty and probable treatment interruptions in the future periods (Chandani and Muwonge, 2003; WHO, 2003; WHO, 2004; Harries

et al., 2007). We assume that even under this policy, current patients are always prioritized over new patients. Thus, using our previous notation, a safety stock policy can be denoted by

$$x_{n,t}^S = \min \{y_{n,t}, W_n\} \text{ and } x_{n,u}^S = [W_n - (a + 1)y_{n,t}]^+; a > 0 \quad (8)$$

where the superscript S denotes the safety stock policy and $ay_{n,t}$ is the amount of the safety stock, equivalent to a periods of demand from current patients. The popularity of this approach is largely due to its simplicity and intuitive appeal and widespread use in traditional inventory systems. However, among organizations that carry out logistics and supply chain implementations related to ART scale-up, there is a recognition that this simple approach might not be optimal and a more scientific approach is needed especially due to the non-stationary nature of scale-up situations (Daniel, 2006). We shall compare the performance of this policy with that of the optimal policy in Section 7.

6.2. Myopic policy

As seen from Proposition 3, the optimal policy involves the possibility of holding on to scarce drugs even though there is an inexhaustible pool of new patients that could be enrolled for treatment. This aspect of the optimal policy (and any policy that involves keeping safety stock) could be unappealing to health care practitioners for ethical reasons. Moreover, many health care programs including WHO's 3-by-5 campaign and PEPFAR programs have explicitly focused on the number of patients enrolled rather than the number of patients receiving uninterrupted treatment as a measure of program success. Our interaction with supply chain consultants working in this area revealed that there is a lot of political pressure to put as many people on treatment as possible without fully considering the potential future impact of these enrollment decisions.

An extreme policy that focuses only on the current period and completely ignores the impact of new enrollments on the ability to continue treatment in the future is obtained by solving the single period problem. This myopic policy is given by $x_{n,t}^m = \min \{y_{n,t}, W_n\}$ and $x_{n,u}^* = [W_n - y_{n,t}]^+$. The next proposition provides a sufficient condition for such a myopic policy to be optimal.

PROPOSITION 5. *A myopic policy is optimal if (C1) and (C2) and the following condition are satisfied:*

$$(C3) \quad (s_{ni} - s_{nr}) \geq \delta(s_{ci} - s_{cr}) + [s_{ni} - s_{ci}]^+ \sum_{u=1}^{T-1} (\delta(\beta_2 + \alpha))^u$$

First note that for $\delta = 0$, (C3) reduces to $s_{ut} \geq s_{uu}$ which we have assumed to be true. Thus if the future is completely discounted, the myopic policy is optimal, as expected. Now for $\delta > 0$, if $s_{tu} > s_{uu}$, (C3) reduces to $(s_{ut} - s_{uu}) \geq \delta(s_{tt} - s_{tu})$. This is because enrolling a new patient transfers the patient from a pool of low QOL score and survival rate into a pool of high QOL score and high survival rate thus increasing the total QALY score of the clinic. Hence the only relevant comparison is between improving the QOL score of a new patient today and improving the QOL score of a current patient tomorrow. On the other hand, if $s_{tu} < s_{uu}$, it implies that the average QOL score of patients with treatment interruptions is worse than that of the new patient pool. Hence a myopic policy would be optimal only if the benefit from treating a new patient today outweighs the cost from interrupting treatment for a current patient in all the future periods.

7. Numerical illustrations

In this section, our primary objective is to examine the impact of the supply uncertainty on the performance of the two enrollment heuristics (myopic policy and safety stock policy with different values of the safety stock parameter a) described in Section 6 and compare their performance to the optimal enrollment policy. Ideally we would perform these comparisons using actual data, but such data do not exist. Hence, although the experiments are suggestive, we cannot draw conclusions about the actual suboptimality of the policies used in practice.

7.1. Setting parameter values

Our model relies on QOL parameters for patient segments that are defined on the basis of current and past treatment status. These values are not directly available in the existing literature. However, we have tried our best to use the literature as a source of guidance in choosing parameter values for our experiments. Table 1 shows the parameter values chosen for the numerical illustrations and the source for each of them. The first two studies (Tengs and Lin, 2002; Holtgrave and Pinkerton, 1997)

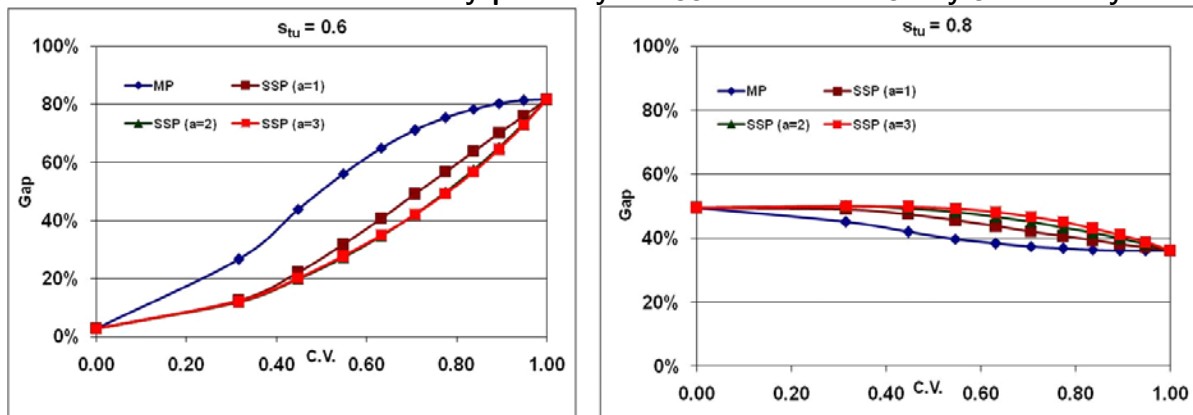
Table 1 Quality of Life estimates

	Tengs and Lin (2002)	Holtgrave and Pinkerton (1997)	and Jelsma et al. (2005); Cleary et al. (2006)	Values chosen
s_{tt}	0.93 (Asymptomatic HIV infection)	0.94 (Asymptomatic HIV infection)	0.85 (ART >12 months)	0.90
s_{tu}	0.81 (Symptomatic HIV infection)	No appropriate estimate	No appropriate estimate	0.60-0.80
s_{ut}	0.81 (Symptomatic HIV infection)	0.70-0.80 (200 < CD4 < 499)	0.71 (ART 0-3 months)	0.75
s_{uu}	0.60-0.70 (CD4 < 200 or Clinical AIDS)	0.60-0.65 (CD4 < 200 or Clinical AIDS)	0.71 (HIV+; no ART)	0.65

are meta-analyses of various studies conducted in the U.S. Jelsma et al. (2005) examined the health status of HIV+ patients in South Africa using a visual analog scale (VAS), the results of which were then converted into utilities using the time-tradeoff method. These scores and the methodology are reported in Cleary et al. (2006). According to the national guidelines in many sub-Saharan African countries, patients become eligible for HAART when their CD4+ counts drop below 200 or if they show clinical symptoms of AIDS. Hence based on the three sources, s_{uu} was chosen to be 0.65. The value we chose for s_{ut} was based on the assumption that after one month of treatment, CD4+ count of patients would, on average be between 200 and 399. For s_{tt} , we assumed that previously enrolled patients who receive uninterrupted treatment in this period would be asymptomatic. However, since Cleary et al. (2006) reported lower QOL for patients on HAART for a year than the QOL value for patients with asymptomatic infection, we adjusted our estimate of s_{tt} downwards from 0.94 to 0.90. It was relatively difficult to choose a value for s_{tu} using the available data as it depends not only on the CD4+ count but also on the severity of adverse outcomes such as drug resistance and fraction of patients with these outcomes. Hence, we decided to use values between 0.60 and 0.80 and examine the sensitivity of our results to this variation.

We choose $N = 24$ to reflect a time horizon of 2 years which is quite natural for resource-limited setting, i.e., each period can be thought of as equivalent to a month. We also choose initial state variables to be $y_{N,t} = I_N = 0$. Use of discounting in health economics is not free of controversy (Krahn and Gafni, 1993). We follow the conventional approach (Shepard and Thompson, 1979; Drummond et al., 1980; Drummond, 1980) and set the monthly discount rate $\delta = 0.995$, which is

Figure 1 Performance of heuristics as a function of the coefficient of variation of the supply distribution. MP denotes the Myopic Policy and SSP denotes the Safety-Stock Policy.



approximately equivalent to an annual discount rate of 5%. However, due to the short horizon our results are not sensitive to the actual choice of the discount rate. We model supply as a three-point distribution with support over the set $\{0, 6, 12\}$. We consider symmetric probability distributions of the form $\Pr(z = 0) = \Pr(z = 12) = p$ and $\Pr(z = 6) = 1 - 2p$. The coefficient of variation for this supply distribution is given by $C.V. = \sqrt{2p}$. Using this form of the supply distribution allows us to change the variance of the distribution without changing the mean. Also since the coefficient of variation is independent of the mean, our results do not depend on the absolute value of the mean. We consider three versions of the safety stock policy depending on the level of safety stock a in (8): $a = 1$, $a = 2$ and $a = 3$ periods.

7.2. Results

The performance of each heuristic is evaluated using the formulation in (7). This captures the increase in QALY score over the baseline of no treatment. Then the performance of each heuristic relative to the optimal enrollment policy (of Proposition 3) is calculated as: $\% \text{ gap} = \frac{V(\text{optimal}) - V(\text{heuristic})}{V(\text{optimal})}$. Figure 1 shows the gap as a function of the coefficient of variation of the supply distribution for two extreme values of s_{tu} , which leads to several observations.

First, the behavior of the performance gap is different for values of s_{tu} greater than s_{uu} and lower than s_{uu} . For lower values of s_{tu} (implying a high penalty for treatment interruption), performance of all the heuristics relative to the optimal policy worsens as the supply uncertainty increases. In

other words, when interruptions are costly, the value of using the optimal policy increases with supply uncertainty. Also, since it is valuable to avoid interruptions, a higher level of safety stock yields better performance. On the other hand, as expected, the myopic policy performs better than the safety stock policy for higher values of s_{tu} (implying relatively low penalty for treatment interruption) since there is no need to avoid treatment interruptions. Also a higher safety stock level results in worse performance, i.e., higher performance gap.

However, note that myopic policy is not necessarily optimal for higher values of s_{tu} . This is because one can do better than the myopic policy by not enrolling any patient in the earlier periods, building stock and then increasing the enrolled patient pool in a future period. In other words, for cases with low interruption penalty, the optimal policy results in higher overall enrollment levels compared to the myopic policy. Note that increasing uncertainty actually improves the performance of the heuristics in this case. This is because given the structure of the supply distribution, higher coefficient of variation implies higher probability of $z = 12$ in certain periods which enables higher enrollment numbers in those periods (even though that might lead to treatment interruptions later).

8. Conclusion and future research

We study the problem of optimally allocating scarce and unreliable supply of resources between previously served and previously unserved customers when continuity of service for previously served customers is desired, with the particular application to the case of a clinic allocating an uncertain supply of antiretroviral drugs to HIV+ individuals. We use dynamic programming to derive the optimal policy of the clinic with the objective of maximizing the total discounted quality adjusted life years of its patients. Our analysis shows that under certain conditions, the optimal policy results in prioritization of current patients, an accepted standard of care. But it also leads to restricting access to treatment for new patients. In our numerical illustrations the optimal enrollment policy (with enforced prioritization of current patients) performs substantially better than enrollment heuristics followed in practice for a wide range of parameter values. We find that supply uncertainty can greatly exacerbate the suboptimality gap of these heuristics. However, as mentioned earlier,

our model is a simplified representation of the resource-constrained context which abstracts from various links between diagnosis, prevention and treatment. An explicit inclusion of these links would be required before the findings from this model could be used in practical settings.

Our work can be extended in several different directions. The model could be made more elaborate, at the expense of tractability, by considering patient flows in a larger system which includes prevention, treatment and diagnosis. One could also allow the clinics to determine order quantities themselves, which is not often the case yet but will gradually become more common. It would be interesting to analyze how a central depot should allocate drugs to multiple clinics. Another extension would be the empirical determination of the actual rationing policies followed by clinics and the impact of supply uncertainty on these policies. We are currently in the process of preparing this, including attempting a limited empirical validation of the model proposed here.

Acknowledgements

The authors are grateful to members of Project USAID DELIVER for several insights into the workings of ARV supply chains in resource-constrained settings. The paper benefitted from several helpful discussions with Felipe Caro, Scott Carr, and Kevin McCardle of the UCLA Anderson School of Management and Drs. Thomas Coates, John Fahey and Martin Shapiro of the David Geffen School of Medicine at UCLA. The authors would also like to thank the seminar participants at UCLA, ISB, Georgia Tech, Northwestern, NUS, UT Austin, INSEAD and Dartmouth for many insightful comments. The second author is grateful to the Technical University Eindhoven, The Netherlands, where he was on sabbatical during part of this study.

References

- Abrams, R.A. and U. S. Karmarkar. 1979. Infinite horizon investment-consumption policies. *Management Sci.* **25**. 1005-1013.
- Amenyah, J., B. Chovitz, E. Hasselberg, A. Karim, D. Mmari, S. Nyinondi, T. Rosche. 2005. Tanzania: Integrated Logistics System Pilot-Test Evaluation: Using the Logistics Indicator Assessment Tool. Arlington, Va.: DELIVER, for the U.S. Agency for International Development.
- Bartlett, J.G. 2006. Ten years of HAART: Foundation for the future. 13th Conference on Retroviruses and Opportunistic Infections. Denver, CO.

- BBC News. 2004. Outrage at SA AIDS drug shortage. June 1.
- Bennet, S. and C. Chanfreau. 2005. Approaches to rationing antiretroviral treatment: ethical and equity implications. *Bulletin of the World Health Organization*. **83**. 541-547.
- Brandeau, M.L., G.S. Zaric, A. Richter. 2003. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *Journal of Health Economics*. **22**. 575-598.
- Chandani, Y. and M. Muwonge. 2003. Logistics and procurement decisions and issues for consideration for introducing and expanding access to ARVs in Uganda. JSI
- Cleary, S.M., D. McIntyre. A. Boule. 2006. The cost-effectiveness of antiretroviral treatment in Khayelitsha, South Africa - a primary data analysis. *Cost effectiveness and Resource Allocation*. **4:20**.
- Daniel, G. 2006. Improving ARV Medicines and Information Management in Ethiopia, January 8–March 11, 2006: Technical Assistance Update and Trip Report. Submitted to the U.S. Agency for International Development by the Rational Pharmaceutical Management Plus Program. Arlington, VA: Management Sciences for Health.
- de Véricourt, F., F. Karaesman, Y. Dallery. 2002. Optimal stock allocation for a capacitated supply system. *Management Sci.* **48**. 1486-1501.
- de Véricourt, F. and M. S. Lobo. 2006. Resource and revenue management in nonprofit operations. Working paper. Fuqua School of Business.
- Dees, J.G. 1998. Enterprising nonprofits. *Harvard Business Review*. Jan-Feb 1998. 58-67.
- Denardo, E. 1965. Contraction mappings in the theory underlying dynamic programming. *SIAM Review*. **9**. 165-177.
- Deo, S. 2007. Three essays in healthcare operations management. Ph.D. Dissertation. UCLA Anderson School of Management.
- Deshpande, V., M. A. Cohen, K. Donhoue. 2003. A threshold inventory rationing policy for service-differentiated demand classes. *Management Sci.* **49**. 683-703.

- Deuermeyer, B.L., W.P. Pierskalla. 1978. A by-product production system with an alternative. *Management Sci.* **24**. 1373-1383.
- Drummond, M.F. 1980. *Studies in Economic Appraisal in Health Care*. New York: Oxford University Press.
- Drummond, M.F., G. Stoddart, G.W. Torrance. 1980. *Principles of Economic Appraisal in Health Care*. Oxford: Oxford University Press.
- Ekong, E., V. Idemyor, O. Akinlade, A. Uwah. 2004. Challenges to antiretroviral (ARV) drug therapy in resource-limited settings - Progress and challenges in the Nigerian initiative. *11th Conference on Retroviruses and Opportunistic Infection*. San Francisco, USA
- El-Sadr W.M., J.D. Lundgren, J.D. Neaton et al. 2006. CD4+count-guided interruption of anti-retroviral treatment *New Eng J of Med* **22**. 2283-2296
- Evans, R.V. 1967. Inventory control of a multi-product system with a limited production resource. *Naval. Res. Logist. Quart.* **14**. 173-184.
- Evans, R.V. 1969. Sales and restocking policies in a single item inventory system. *Management Sci.* **14**. 463-472.
- Foster, W. and J. Bradach. 2005. Should nonprofits seek profits? *Harvard Business Review*. February 2005. 92-100.
- Frank, K.C., R.Q. Zhang, I. Duenyas. 2003. Optimal policies for inventory systems with priority demand classes. *Oper. Res.* **51**. 993-1002.
- Gronbjerg, K.A. 1992. *Nonprofit human service organizations: Funding strategies and patterns of adaptation*. In Y. Hasenfeld (Ed.), *Human Services as Complex Organizations*. 73-97. Newbury Park, CA: Sage.
- Ha, A. 1997a. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Sci.* **43** 1093-1103.
- Ha, A. 1997b. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Res. Logist.* **44** 458-472.

- Harries, A.D., E.J. Schouten, S.D. Makombe, E. Libamba, H.N. Neufville, E. Some, G. Kadewered, D. Lungue. Ensuring uninterrupted supplies of antiretroviral drugs in resource-poor settings: an example from Malawi. *Bulletin of the World Health Organization*. **85**. 152-155.
- Health Systems Trust. 2005. Implementing the Comprehensive Care and Treatment Programme for HIV and AIDS patients in the Free State: Sharing experiences. Conference report.
- Holtgrave, D.R. and S.D. Pinkerton. 1997. Updates of cost illness and quality of life estimates for use in economic evaluations of HIV prevention programs. *JAIDS*. **16**. 54-62.
- IOM. 2005. *Scaling Up Treatment for the Global AIDS Pandemic: Challenges and Opportunities*. The National Academies Press.
- IRINNews.org. 2005. Swaziland: HIV positive Swazis take govt to task over ARV supply. December 6.
- IRINNews.org. 2002. Coalities decries "Regular" shortage of HIV drugs. April 16.
- ITPC. 2005. Missing the target: A report on HIV/AIDS treatment access from the frontlines.
- Jelsma, J., E. Maclean, J. Hughes, X. Tinise, M. Darder. 2005. An investigation into the health-related quality of life of individuals living with HIV who are receiving HAART. *AIDS Care*. **17**. 579-588.
- JSI. 2006. Personal communication.
- Kaplan, E.H. and H. Pollack. 1998. Allocating HIV prevention resources. *Socio-Econ. Plann. Sci*. **32**. 257-263.
- Karmarkar, U.S. 1981. The multiperiod multilocation inventory problem. *Oper. Res*. **29**. 215-228.
- Lasry, A., G.S. Zaric, M.W. Carter. 2007. Multi-level resource allocation for HIV prevention: A model for developing countries. *Eur J of Oper Res*. **180**. 786-799.
- Lippman, S.A. 1974. Semi-markov decision processes with unbounded rewards. *Management Sci*. **19**. 717-731.
- Macklin, R. 2004. Ethics and equity in access to HIV treatment - 3 by 5 initiative. Background paper for the consultation on equitable access to treatment and care for HIV/AIDS.

- McGough, L.J., S. J. Reynolds, T.C. Quinn, J.M. Zenilman. 2005. Which patients first? Setting priorities for antiretroviral therapy where resources are limited. *American Journal of Public Health*. **95**. 1173-1180.
- Miller, B. 1974. Optimal consumption with stochastic income stream. *Econometrica*. **42**. 253-266.
- Nahmias, S., W. S. Demmy. 1981. Operating characteristics of an inventory system with rationing. *Management Sci.* **27** 1236–1245.
- Nyenwa, J., D. Alt, A. Karim, T. Kufa, J. Mboyane, Y. Ouedraogo, T. Simoyi. 2005. Zimbabwe HIV & AIDS Logistics System Assessment. Arlington, Va.: John Snow, Inc./DELIVER, for the U.S. Agency for International Development.
- Olsen, T.L. and R.P. Parker. 2006. Customer behavior in inventory management. *Working paper*. Yale School of Management.
- van Oosterhout, J.J., N. Bodasing, J.J. Kunwenda, C. Nyirenda, J. Mallewa, P.R. Cleary, M.P. de Baar, R. Schuurman, D.M. Burger, E.E. Zijlstra. 2005. Evaluation of antiretroviral therapy results in a resource-poor setting in Blantyre, Malawi. *Tropical Medicine and International Health*. **10**. 464-470.
- Oyugi, J.H., J. Byakika-Tusiime, K. Ragland, O. Laeyendecker, R. Mugerwa, C. Kityo, P. Mugenyi, T.C. Quinn, D.R. Bangsberg. 2007. Treatment interruptions predict resistance in HIV-positive individuals purchasing fixed-dose combination antiretroviral therapy in Kampala, Uganda. *AIDS*. **21**. 965-971.
- Palella et al. 1998. Declining morbidity and mortality among patients with advanced Human Immunodeficiency Virus infection. *N. Engl. J. Med.* **338**. 853-860.
- Richter, A., M.L. Brandeau, and D.K. Owens. 1999. An analysis of optimal resource allocation for HIV prevention among injection drug users and nonusers. *Med. Decis. Making*. **19**. 167–179.
- Rosen, S. I. Sanne, A. Collier, J.L. Simon. 2005. Rationing antiretroviral therapy for HIV / AIDS in Africa: Choices and consequences. *PLoS Medicine*. **2**. e303.
- Scott, K. 2003. Funding matters: The impact of Canada's new funding regime on nonprofit and voluntary organizations. *A report by Canadian Council on Social Development*.

- Shepard D.S., M.S. Thompson. 1979. First Principles of Cost-Effectiveness in Health. *Public Health Reports*. **6**. 535-543.
- Simpson, V.P. 1978. Optimum Solution Structure for a Repairable Inventory Problem. *Oper. Res.* **26**. 270-281.
- Stokey, N.L., R.E. Lucas, Jr., E.C. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Topkis, D. M. 1968. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Sci.* **15** 160–176.
- Van Nunen, J.A.E.E. and J. Wessels. 1975. A note on dynamic programming with unbounded rewards. *Management Sci.* **24**. 576-580.
- Wagstaff, A. 1991. QALYs and equity-efficiency tradeoff. *Journal of Health Economics*. **10**. 21-41.
- Walensky R.P., A.D. Paltiel, E. Losina. Three million years of life saved: The survival benefits of AIDS therapy in the United States. *J Infect Dis*. 2006. forthcoming.
- WHO. 1998. Guidance modules on antiretroviral treatments. Module 8: ARVs - Regulation, distribution and control.
- WHO. 2003. Emergency scale up of antiretroviral therapy in resource-limited settings: technical and operational recommendations to achieve 3 by 5.
- WHO. 2004. Patient monitoring guidelines for HIV care and antiretroviral therapy (ART).
- WHO. 2005a. Progress on global access to HIV antiretroviral therapy: An update on “3 by 5”.
- WHO. 2005b. AIDS epidemic update: December 2005.
- WHO 2006. Progress on global access to HIV antiretroviral therapy: A report on “3 by 5” and beyond.
- Yadav, P. 2007. Analysis of the public, private and mission sector supply chains for essential drugs in Zambia. Draft version.
- Zaric, G.S., M.L. Brandeau. 2001. Optimal investment in a portfolio of HIV prevention programs. *Medical Decision Making*. **21**. 391-408.

Zenios, S.A., L.M. Wein, G.M Chertow. 2000. Dynamic allocation of kidneys to candidates on the transplant waiting list. *Oper. Res.* **49**. 549-569

Zhang, R.Q. and M. Sobel. 2001. Inventory policies for systems with stochastic and deterministic demand. *Oper. Res.* **49**. 157-162.

Appendix A: Infinite horizon formulation

The model discussed so far is for a horizon of finite length denoted by N . However, analysis of an infinite horizon model could be appropriate if N is not known with certainty or if N is large enough so that the infinite horizon problem can be considered as an approximation to the finite horizon problem. The infinite horizon problem corresponding to (7) is stated below:

$$V^* = \max_{x_{n,t} \geq 0, x_{n,u} \geq 0} \sum_{n=1}^{\infty} \delta^{n-1} \hat{h}_n(x_{n,t}, x_{n,u}) \quad (9)$$

$$s.t. \quad x_{n,t} \leq y_{n,t} \quad \forall n$$

$$x_{n,u} + x_{n,t} \leq W_n \quad \forall n$$

The corresponding recursive equation in the infinite horizon case is given by:

$$V(W, y_c) = \max_{x_n \geq 0, x_c \geq 0} \left\{ \hat{h}(x_c, x_n) + \delta E[V(W - x_c - x_n + Z, \beta(y_c + x_n))] \right\} \quad (10)$$

$$s.t. \quad x_c \leq y_c$$

$$x_n + x_c \leq W$$

However, for the infinite horizon formulation to be meaningful, the resource-constrained condition needs to be satisfied for all periods. A sufficient condition for this to happen is provided in the appendix. Other technical challenges in our formulation which make the infinite horizon problem difficult are (i) the single period reward function $h(\cdot)$ and hence the value function $V(\cdot)$ is unbounded since W is not uniformly bounded from above, and (ii) the underlying state-space is continuous.

Following the approach by Lippman (1974) and Van Nunen and Wessels (1978) among others, we define a modified sup norm that bounds V . Also, to resolve the issue of continuous state-space, we allow only Borel measurable policies to ensure that the underlying transition functions have the Feller property (Stokey et al., 1989). Using these modifications, we can redefine a Banach space

over which the contraction mapping approach (Denardo, 1967) can be applied to show that the equation (10) has a unique fixed point. This result is summarized in the following Proposition and the details of our approach are given in the appendix. A similar approach has been used for consumption investment problems by Abrams and Karmarkar (1979) and Miller (1974).

PROPOSITION 6. *The recursive equation (10) has a unique solution \hat{V} , which satisfies $\hat{V} = V^* = \lim_{t \rightarrow \infty} V_t$ and there exists a unique optimal policy such that V^* is attained.*

Appendix B: Resource-constrained condition

In our model described in (7), we assumed that $y_{n,u} > W_n \forall n$. Since $y_{n,u}$ and W_n are both random variables, this is true if certain restrictions are placed on the supply distributions $\Phi_n(\cdot)$. Here, we derive one such restriction in the form of an upper bound on the support of $\Phi_n(\cdot)$. Consider the finite horizon problem with initial conditions $y_{N,n}$ and I_N before the shipment in period N is received. Then $y_{N,u} > W_N$ if $z_N < y_{N,u} - I_N$. Suppose this is true. Then for period $N - 1$ under any feasible solution $x_{N,u}$ and $x_{N,t}$; $y_{N-1,u} = (y_{N,u} - x_{N,u})(\beta_u + \alpha)$ and $W_{N-1} = W_N - x_{N,u} - x_{N,t} + z_{N-1}$. Now

$$\begin{aligned} y_{N-1,u} > W_N &\iff \\ (y_{N,u} - x_{N,u})(\beta_u + \alpha) > W_N - x_{N,u} - x_{N,t} + z_{N-1} &\iff \\ z_{N-1} < y_{N,u}(\beta_u + \alpha) - (\beta_u + \alpha - 1)x_{N,u} + x_{N,t} - W_N &\quad (11) \end{aligned}$$

Since (11) has to be true for all feasible $x_{N,u}$, $x_{N,t}$ and $\beta_u + \alpha - 1 > 0$ we substitute $x_{N,t} = 0$ and $x_{N,u} = W_N$ to obtain a lower bound on RHS. Thus (11) is satisfied for all feasible $x_{N,u}$, $x_{N,t}$ if $(\beta_u + \alpha)z_N + z_{N-1} < (\beta_u + \alpha)(y_{N,u} - I_N)$. Continuing this inductively, we find that a sufficient condition to ensure $y_{n,u} > W_n \forall n$ is given by

$$\sum_{i=n}^N (\beta_u + \alpha)^i z_i^U < (\beta_u + \alpha)^N (y_{N,u} - I_N) \quad \forall n \quad (12)$$

where z_i^U is the upper bound on the support of z_i . A less tight bound is obtained by replacing each z_i^U by $\max_{n \leq i \leq N} z_i^U$ in (12) to obtain

$$\max_{n \leq i \leq N} z_i^U < \frac{(\beta_u + \alpha)^N (y_{N,u} - I_N)}{\sum_{i=n}^N (\beta_u + \alpha)^i} = \frac{(y_{N,u} - I_N) \left(1 - \frac{1}{(\beta_u + \alpha)}\right)}{\left(1 - \frac{1}{(\beta_u + \alpha)^{N-n+1}}\right)} \quad \forall n \quad (13)$$

Now since (13) has to be satisfied $\forall n$, we substitute $n = 1$ to obtain

$$\max_{1 \leq i \leq N} z_i^U < \frac{(y_{N,u} - I_N) \left(1 - \frac{1}{(\beta_u + \alpha)}\right)}{\left(1 - \frac{1}{(\beta_u + \alpha)^N}\right)} \quad (14)$$

Note that for $N \rightarrow \infty$, RHS of (14) $\rightarrow (y_{N,u} - I_N) \left(1 - \frac{1}{(\beta_u + \alpha)}\right)$ and max operator in the LHS has to be replaced by sup. Thus the equivalent condition for infinite horizon problem is

$$\sup_{1 \leq i \leq N} z_i^U < (y_{N,u} - I_N) \left(1 - \frac{1}{(\beta_u + \alpha)}\right) \quad (15)$$

While analyzing the infinite horizon problem in Section 6, we assume that (15) is satisfied.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

Proofs of Statements

Proof of Proposition 1:

(i) We use induction to prove this. Let $\mathcal{S}_n = \{(x_{n,u}, x_{n,t}) : x_{n,u} + x_{n,t} \leq W_n, 0 \leq x_{n,t} \leq y_{n,t}, 0 \leq x_{n,u}\}$.

Note that \mathcal{S}_n is a convex set. Using this notation, the recursive equation for $n = 1$ is given by:

$$V_1(W_1, y_{1,t}) = \max_{(x_{1,t}, x_{1,u}) \in \mathcal{S}_1} (s_{tt} - s_{tu})x_{1,t} + (s_{ut} - s_{uu})x_{1,u} + s_{tu}y_{1,t} + s_{uu}y_{1,u}$$

Thus $V_1(W_1, y_{1,t})$ is jointly concave in its arguments. Now assume that the result holds for $n - 1$.

Let

$$V_n(y_{n,t}^i, W_n^i) = f_n(x_{n,t}^i, x_{n,u}^i) \quad i = 1, 2$$

$$\text{and } V_n(y_{n,t}^\lambda, W_n^\lambda) = f_n(x_{n,t}^*, x_{n,u}^*)$$

where $(y_{n,t}^\lambda, W_n^\lambda) \triangleq \lambda(y_{n,t}^1, W_n^1) + (1 - \lambda)(y_{n,t}^2, W_n^2)$. Also define $(x_{n,t}^\lambda, x_{n,u}^\lambda) \triangleq \lambda(x_{n,t}^1, x_{n,u}^1) + (1 - \lambda)(x_{n,t}^2, x_{n,u}^2)$. Since \mathcal{S}_n is a convex set $(x_{n,t}^\lambda, x_{n,u}^\lambda) \in \mathcal{S}_n$. Using this notation

$$\begin{aligned} \lambda V_t(y_{n,t}^1, W_n^1) + (1 - \lambda) V_n(y_{t,c}^2, W_n^2) &= \lambda h_t(x_{n,t}^1, x_{t,n}^1) + (1 - \lambda) h_n(x_{n,t}^2, x_{n,u}^2) \\ &\quad + \delta E [\lambda V_{n-1}(\beta(y_{n,t}^1 + x_{t,n}^1), W_n^1 - x_{n,t}^1 - x_{t,n}^1)] \\ &\quad + \delta E [(1 - \lambda) V_{n-1}(\beta(y_{t,c}^2 + x_{n,t}^2), W_n^2 - x_{n,t}^2 - x_{n,u}^2)] \\ &\leq h_n(x_{n,t}^\lambda, x_{n,u}^\lambda) + \delta E [V_{n-1}(\beta(y_{n,t}^\lambda + x_{n,u}^\lambda), W_n^\lambda - x_{n,t}^\lambda - x_{n,u}^\lambda)] \\ &= f_n(x_{n,t}^\lambda, x_{n,u}^\lambda) \leq f_n(x_{n,t}^*, x_{n,u}^*) = V_n(y_{n,t}^\lambda, W_n^\lambda) \end{aligned}$$

Thus $V_n(\cdot)$ is jointly concave in its arguments.

(ii) From (i), V_{n-1} is jointly concave in its arguments for all realizations of Z_{n-1} . Thus, V_{n-1} is also jointly concave in $(x_{n,t}, x_{n,u})$ since W_{n-1} and $y_{n-1,c}$ are linear transformations of $(x_{n,t}, x_{n,u})$ as seen from (3) for all realizations of Z_{n-1} . Since the expectation operator preserves concavity, $E[V_{n-1}]$ is also jointly concave in $(x_{n,t}, x_{n,u})$. $h_n(x_{n,t}, x_{n,u})$ is linear and hence jointly concave in its arguments. Since $f_n(x_{n,t}, x_{n,u})$ is a sum of two concave functions, it is also jointly concave in $(x_{n,t}, x_{n,u})$.

(iii) We show that our problem can be reformulated in the form described in Karmarkar (1981).

Introduce slack variables $\rho_n^1, \rho_n^2, \rho_n^3$ in the constraints in (6), and $u_{n,t} = y_{n,t} + \frac{\beta_u}{\beta_c} x_{n,u}$. Then define

$$\mathbf{y}_n = \begin{bmatrix} W_n \\ y_{n,t} \\ y_{n,u} \\ y_{n,t} \\ y_{n,u} \\ W_n \end{bmatrix}; \mathbf{x}_n = \begin{bmatrix} x_{n,t} \\ x_{n,u} \\ \rho_n^1 \\ \rho_n^2 \\ \rho_n^3 \end{bmatrix}; \mathbf{u}_n = \begin{bmatrix} I_n \\ u_{n,t}^1 \\ u_{n,t}^2 \\ 0 \\ 0 \\ 0 \end{bmatrix}; \mathbf{A} = \begin{bmatrix} -1 & -1 & 0 & 0 & 0 \\ 0 & \frac{\beta_u}{\beta_t} & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ -1 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & -1 & 0 \\ -1 & -1 & 0 & 0 & -1 \end{bmatrix}$$

and a transition function ω such that

$$\omega(\mathbf{u}_n, \tilde{z}_{n-1}) = \begin{bmatrix} I_n - \tilde{z}_{n-1} \\ \beta_t u_{n,t}^1 \\ (\beta_u + \alpha) u_{n,t}^2 \\ \beta_t u_{n,t}^1 \\ (\beta_u + \alpha) u_{n,t}^2 \\ I_n - \tilde{z}_{n-1} \end{bmatrix}$$

With these definitions, the problem formulation (7) becomes

$$\begin{aligned} V_n(W_n, y_{n,t}) &= \max_{\mathbf{x}_n \geq \mathbf{0}} \left\{ \hat{h}_n(x_{n,t}, x_{n,u}) + \delta E[V_{n-1}(W_{n-1}, y_{n-1,c})] \right\} \\ \text{s.t. } & \mathbf{u}_n = \mathbf{A}\mathbf{x}_n + \mathbf{y}_n \\ & \mathbf{y}_{n-1} = \omega(\mathbf{u}_n, \tilde{z}_{n-1}) \end{aligned}$$

This formulation is equivalent to the formulation (MP) in Karmarkar (1981) and hence Proposition 8 applies completing the result.

Proof of Proposition 2: First consider the case when $s_{ut} - s_{uu} > s_{tt} - s_{tu}$. For $n = 1$, the optimal solution is trivial. Thus

$$E_{z_1}[V_1(W_1, y_1)] = s_{tu}y_{1,c} + (s_{ut} - s_{uu})(I_1 + E[z_1]) \quad (\text{EC.1})$$

Then the objective function for $n = 2$ becomes:

$$f_2(x_{2,t}, x_{2,u}) = ((s_{tt} - s_{tu}) - \delta(s_{ut} - s_{uu}))x_{2,t} + (s_{ut} - s_{uu}^2) + \delta(s_{tu}\beta_u - (s_{ut} - s_{uu})) + (s_{ut} - s_{uu})E[z_1] \quad (\text{EC.2})$$

Since (EC.2) is linear in its arguments, comparing the coefficients gives the desired result.

Now consider the case of $s_{tt} - s_{tu} > s_{ut} - s_{uu}$. Again the optimal policy for $n = 1$ is straightforward. Thus evaluating $E[V_1(I_1, y_{1,c})]$ and adding the single period reward function, the objective function for $n = 2$ becomes

$$f_2(x_{2,t}, x_{2,u}) = (s_{tt} - s_{tu})x_{2,t} + (s_{ut} - s_{uu}^2)x_{2,u} + \delta \left((s_{tt} - (s_{ut} - s_{uu}))y_{1,c} + (s_{ut} - s_{uu})I_1 - ((s_{tt} - s_{tu}) - (s_{ut} - s_{uu})) \int_{z_1^L}^{y_1 - I_1} \Phi_1(\tilde{z}_1) d\tilde{z}_1 \right)$$

Since this is a constrained optimization problem with concave objective function and linear constraints, KKT conditions are both necessary and sufficient for optimality. Let λ_1 and λ_2 be the lagrangean multipliers associated with the constraints in (7). Then forming the largrangean in the usual manner and taking the first order derivatives

$$\frac{\partial L_2}{\partial x_{2,t}} = (s_{tt} - s_{tu}) - \delta(s_{ut} - s_{uu}) - \lambda_1 - \lambda_2 = 0 \quad (\text{EC.3})$$

$$\begin{aligned} \frac{\partial L_2}{\partial x_{1,u}} &= (s_{ut} - s_{uu}^1) + \delta(- (s_{ut} - s_{uu}) - ((s_{tt} - s_{tu}) - (s_{ut} - s_{uu})) \Phi_1(y_1 - I_1)) \\ &+ \beta_u ((s_{tt} - (s_{ut} - s_{uu})) - ((s_{tt} - s_{tu}) - (s_{ut} - s_{uu}))) \Phi_1(y_1 - I_1) - \lambda_2 = 0 \quad (\text{EC.4}) \end{aligned}$$

Then depending on which of the constraints are binding or slack we get following cases:

Case I: $\lambda_2 > 0; \lambda_1 = 0$

$x_{2,t}^* + x_{2,u}^* = W_2$ and $I_1 = 0$. Substituting $\lambda_2 > 0; \lambda_1 = 0$ and equating (EC.3) and (EC.4) we obtain

$$\Phi(y_1^*) = \frac{(s_{tu} - s_{uu}(1 - \delta\alpha)) - (1 - \delta\beta_2)(s_{tt} - s_{ut})}{\delta\beta_2((s_{tt} - s_{tu}) - (s_{ut} - s_{uu}))} = k_1. \text{ Thus } x_{2,u}^* = \frac{\Phi_1^{-1}(k_1) - \beta_t y_{2,t}}{\beta_u} \text{ and } x_{2,t}^* = W_2 - \frac{\Phi_1^{-1}(k_1) - \beta_t y_{2,t}}{\beta_u}.$$

For this to be feasible, we need the following conditions to be satisfied: (i) $\Phi_1^{-1}(k_1) - \beta_t y_{2,t} \geq 0$ (ii)

$$\beta_u W_2 + \beta_t y_{2,t} \geq \Phi_1^{-1}(k_1) \text{ and (iii) } \beta_u W_2 + (\beta_t - \beta_u) y_{2,t} \leq \Phi_1^{-1}(k_1).$$

Case II: $\lambda_1 > 0; \lambda_2 = 0$

$$x_{2,t}^* = y_{2,t}. \text{ Substituting in (EC.4) we obtain } \Phi(y_1^* - I_1^*) = \frac{s_{ut} - s_{uu}(1 + \delta(\beta_u + \alpha)) + \delta(s_{tt} - (s_{ut} - s_{uu}))(1 + \beta_u)}{((s_{tt} - s_{tu}) - (s_{ut} - s_{uu}))(1 + \delta\beta_2)} =$$

$$k_2. \text{ Simplifying and substituting the value of } x_{2,t}^*, \text{ we get } x_{2,u}^* = \frac{\Phi_1^{-1}(k_1) + W_2 - (1 + \beta_t)y_{2,t}}{(1 + \beta_u)}. \text{ Again for this}$$

to be feasible, we need the following conditions to be satisfied: (i) $(1 + \beta_u)y_{2,t} - W_2 \leq \Phi_1^{-1}(k_2)$ (ii)

$$\beta_u W_2 + (\beta_t - \beta_u)y_{2,t} \geq \Phi_1^{-1}(k_2).$$

Case III: $\lambda_1 > 0; \lambda_2 > 0$

$x_{2,t}^* = y_{2,t}$ and $x_{2,u}^* = W_2 - y_{2,t}$ and. Substituting $\lambda_1 > 0; \lambda_2 > 0$ in (EC.3) and (EC.4) we get $\beta_u W_2 +$

$$(\beta_t - \beta_u)y_{2,t} \geq \Phi_1^{-1}(k_1). \text{ This is feasible if (i) } W_2 \geq y_{2,t}.$$

Case IV: $\lambda_1 > 0; \lambda_2 > 0$

This is not possible since $(s_{tt} - s_{tu}) - \delta(s_{ut} - s_{uu}) > 0$.

Thus combining all the three cases, we get the desired form of the optimal policy, where η and θ are defined appropriately.

Proof of Proposition 3: We shall use induction to prove this result. First consider $n = 1$. Clearly since (C1) states that $s_{tt} - s_{tu} > s_{ut} - s_{uu}$, the optimal policy has the desired form since $\frac{\partial f_1}{\partial x_{1,t}} > \frac{\partial f_1}{\partial x_{1,u}}$ and $\frac{\partial f_1}{\partial x_{1,t}} > 0$. Now consider period n . The partial derivatives of the maximand are given by:

$$\frac{\partial f_n}{\partial x_{n,u}} = (s_{ut} - s_{uu}^n) + \delta E \left[\beta_u \frac{\partial V_{n-1}}{\partial y_{n-1,c}} - \frac{\partial V_{n-1}}{\partial W_{n-1}} \right] \quad (\text{EC.5})$$

$$\frac{\partial f_n}{\partial x_{n,t}} = (s_{tt} - s_{tu}) + \delta E \left[-\frac{\partial V_{n-1}}{\partial W_{n-1}} \right] \quad (\text{EC.6})$$

Now suppose that the following two results hold for $n - 1$:

$$E \left[\frac{\partial V_{n-1}}{\partial W_{n-1}} \right] < (s_{tt} - s_{tu}) \quad (\text{EC.7})$$

$$E \left[\frac{\partial V_{n-1}}{\partial y_{n-1,c}} \right] < \frac{(s_{tt} - s_{tu}) - (s_{ut} - s_{uu}^n)}{\delta \beta_2} \quad (\text{EC.8})$$

These together imply that $\frac{\partial f_n}{\partial x_{n,t}} > 0$ and $\frac{\partial f_n}{\partial x_{n,t}} > \frac{\partial f_n}{\partial x_{n,u}}$ and hence the optimal policy has desired form in period t . In order to carry forward the induction we need to show that (EC.7) and (EC.8) hold for t . Using the structure of the optimal policy in period t , we can write the following:

$$V_n(W_n, y_{n,t}) = \begin{cases} (s_{tt} - s_{tu})W_n + s_{tu}y_{n,t} + \delta E[V_{n-1}(z_{n-1}, \beta_t y_{n,t})] & \text{if } z_{n-1} < y_{n,t} - I_n \\ s_{tt}y_{n,t} + (s_{ut} - s_{uu}^n)(W_n - y_{n,t}) + \delta E[V_{n-1}(z_{n-1}, \beta_u W_n + (\beta_t - \beta_u)y_{n,t})] & \text{if } y_{n,t} - I_n < z_{n-1} < y_{n,t} - I_n + \theta_n \\ s_{tt}y_{n,t} + (s_{ut} - s_{uu}^n)\theta_n + \delta E[V_{n-1}(W_n - y_{n,t} - \theta_n + z_{n-1}, \beta_t y_{n,t} + \beta_u \theta_n)] & \text{if } z_{n-1} > y_{n,t} - I_n + \theta_n \end{cases} \quad (\text{EC.9})$$

Using the fact that $V_n(\cdot)$ is jointly concave in its arguments we can see that $E \left[\frac{\partial V_n}{\partial W_n} \right] < (s_{tt} - s_{tu})$ and $E \left[\frac{\partial V_n}{\partial y_{n,t}} \right] < s_{tt} - (s_{ut} - s_{uu}^n) + \delta(\beta_t - \beta_u) E \left[\frac{\partial V_{n-1}}{\partial y_{n-1,c}} \right]$. Clearly (EC.7) holds for t . Now to prove the remaining, consider

$$\begin{aligned} E \left[\frac{\partial V_n}{\partial y_{n,t}} \right] &< s_{tt} - (s_{ut} - s_{uu}^n) + \delta(\beta_t - \beta_u) E \left[\frac{\partial V_{n-1}}{\partial y_{n-1,c}} \right] \\ &< s_{tt} - (s_{ut} - s_{uu}^n) + \delta(\beta_t - \beta_u) \frac{(s_{tt} - s_{tu}) - (s_{ut} - s_{uu}^n)}{\delta \beta_2} \\ &= \frac{(s_{tt} - (s_{ut} - s_{uu}^n))\beta_t - (\beta_t - \beta_u)s_{tu}}{\beta_u} \end{aligned} \quad (\text{EC.10})$$

where we have used (EC.8). Now rewriting (C2) we obtain

$$\begin{aligned}
s_{tu}(1 - \delta(\beta_t - \beta_u)) &< (s_{tt} - s_{ut})(1 - \delta\beta_t) + s_{uu}(1 + \delta(\beta_u + \alpha - \beta_t)) \\
&= (s_{tt} - s_{ut})(1 - \delta\beta_t) + s_{uu}^2 - s_{uu}^1(\delta\beta_t) \\
&< (s_{tt} - s_{ut})(1 - \delta\beta_t) + s_{uu}^{n+1} - s_{uu}^n(\delta\beta_t)
\end{aligned} \tag{EC.11}$$

where we have used the definition of s_{uu}^n and that it is increasing in n if $(\beta_u + \alpha - \beta_t) > 0$. Now substituting (EC.11) in (EC.10) we obtain the desired result.

Proof of Proposition 4: θ_n is a solution to the equation $\frac{\partial f_n}{\partial x_{n,u}} = (s_{ut} - s_{uu}^n) + \delta E \left[\beta_u \frac{\partial V_{n-1}}{\partial y_{n-1,t}} - \frac{\partial V_{n-1}}{\partial W_{n-1}} \right] = 0$, where the expectation is with respect to the distribution of Z_{n-1} i.e., F_{n-1} or G_{n-1} . First, using the implicit function theorem it is clear that $\frac{\partial \theta_n}{\partial y_{n,t}} = -\frac{\frac{\partial}{\partial y_{n,t}} E \left[\beta_u \frac{\partial V_{n-1}}{\partial y_{n-1,t}} - \frac{\partial V_{n-1}}{\partial W_{n-1}} \right]}{\frac{\partial}{\partial \theta_n} E \left[\beta_u \frac{\partial V_{n-1}}{\partial y_{n-1,t}} - \frac{\partial V_{n-1}}{\partial W_{n-1}} \right]} = -1$ since $\frac{\partial W_{n-1}}{\partial y_{n,t}} = \frac{\partial W_{n-1}}{\partial \theta_n}$ and $\frac{\partial y_{n-1,t}}{\partial y_{n,t}} = \frac{\partial y_{n-1,t}}{\partial \theta_n}$. This proves the part (i). Now for part (ii) $\frac{\partial V_{n-1}}{\partial W_{n-1}}$ is a non-increasing function of W_{n-1} and hence of Z_{n-1} . Hence, $-\frac{\partial V_{n-1}}{\partial W_{n-1}}$ is

a non-decreasing function of Z_{n-1} . Also using $\beta_t = \beta_u$ from (EC.9), we have

$$\frac{\partial V_{n-1}}{\partial y_{n-1,t}} = \begin{cases} s_{tu} + \delta\beta_t E \left[\frac{\partial V_{t-2}}{\partial y_{t-2,t}} \right]; & z_{n-1} < y_{n-1,t} - I_{n-1} \\ s_{tt} - (s_{ut} - s_{uu}^n); & z_{n-1} \geq y_{n-1,t} - I_{n-1} \end{cases}$$

Now $\frac{\partial V_{n-1}}{\partial y_{n-1,t}}$ is decreasing in $y_{n-1,t}$ and hence increasing in z_{n-1} for a given $y_{n-1,t}$. Hence $E_F \left[-\frac{\partial V_{n-1}}{\partial W_{n-1}} \right] \geq E_G \left[-\frac{\partial V_{n-1}}{\partial W_{n-1}} \right]$ by first order stochastic dominance. This implies that $\frac{\partial f_n}{\partial x_{n,u}} \Big|_F > \frac{\partial f_n}{\partial x_{n,u}} \Big|_G$ and hence $\theta_n(F) > \theta_n(G)$. This proves the result.

Proof of Proposition 6: Define a borel measurable function $r : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that chooses $(x_{n,t}, x_{n,u}) \in \mathcal{S}_n$ given $(W_n, y_{n,t})$. Let \mathcal{V} denote the space of continuous real-valued functions satisfying the following boundedness property:

$$\sup_{W,y} \left\{ \frac{\|V(W,y)\|}{\max\{W,D\}} \right\} < \infty \tag{EC.12}$$

where D is a positive constant and define the distance function

$$q(V, V') = \sup_{W,y} \left\{ \frac{\|V(W,y) - V'(W,y)\|}{\max\{W,D\}} \right\}; \quad V, V' \in \mathcal{V}$$

Then (\mathcal{V}, g) is a Banach space (Miller, 1974). Define $g(\cdot)$ as a vector valued function defining the transitions in (1) and (3). Since, r is a measurable function, $g(\cdot)$ is a transition function (Theorem 8.9; Stokey et al., 1989). Moreover, since $g(\cdot)$ is continuous, it possesses the Feller Property (Stokey et al., 1989; p.237). Hence the operator

$$(H_r V)(W, y) = h(r(W, y)) + \delta E[V(g(W, y))]$$

maps into a continuous function. To show that $H_r V$ satisfies the boundedness property in (EC.12) note that

$$\begin{aligned} (H_r V)(W, y) &= h(r(W, y)) + \delta E[V(g(W, y))] \\ &\leq s_{tt}W + \delta E[V(W - x_t - x_u + z, \beta(y + x_u))] \\ &\leq s_{tt}W + \delta V(W + E[z], \beta(y + x_n)) \\ &\leq s_{tt}W + \delta M \max\{W + E[z], D\} \end{aligned} \tag{EC.13}$$

where the last inequality is due to the boundedness condition (EC.12) and the inequality before that is because $V(W, y)$ is concave and increasing in W . Clearly, (EC.13) implies that $(H_r V)(W, y)$ is bounded since $E[z]$ is bounded. Thus $H_r V : \mathcal{V} \rightarrow \mathcal{V}$. Hence Theorem 9.6 and 9.2 from Stokey et al. (1989) prove the result.

Proof of Proposition 5: Using results (i) and (ii) of Proposition 3 in (EC.5), we obtain:

$$\begin{aligned} \frac{\partial f_n}{\partial x_{n,u}} &\geq (s_{ut} - s_{uu}^n) + \delta \left(s_{tu} \beta_u \sum_{i=0}^{n-2} (\delta \beta_t)^i - (s_{tt} - s_{tu}) \right) \\ &\geq (s_{ut} - s_{uu}^n) + s_{tu} \sum_{i=1}^{n-1} (\delta \beta_2)^i - \delta (s_{tt} - s_{tu}) \\ &= (s_{tu} - s_{uu}) \sum_{i=0}^{n-1} (\delta \beta_2)^i - (s_{tu} - s_{ut}) - \delta (s_{tt} - s_{tu}) \end{aligned}$$

Now two cases are possible depending on the relative values of s_{tu} and s_{uu} . If $s_{tu} > s_{uu}$, then $(s_{tu} - s_{uu}) \sum_{i=1}^{n-1} (\delta \beta_2)^i \geq s_{tu} - s_{uu}$ and hence condition (C3) in the hypothesis implies $\frac{\partial f_n}{\partial x_{n,u}} \geq 0 \quad \forall n \leq N$. If $s_{tu} < s_{uu}$, then $(s_{tu} - s_{uu}) \sum_{i=0}^{n-1} (\delta \beta_2)^i \geq (s_{tu} - s_{uu}) \sum_{i=0}^{N-1} (\delta \beta_2)^i$ and hence condition (C4) in the hypothesis implies $\frac{\partial f_n}{\partial x_{n,u}} \geq 0 \quad \forall n \leq N$.