Designing Individualized Therapy for Patients with Hypertension

M. Zargoush, M. Gumus, V. Verter, S. Daskalopoulou

Abstract

Hypertension has not been well studied by operations researchers from a clinical decision support perspective. Moreover, little personalized (i.e. patient-centric) guidance is available regarding the number and combination of antihypertensive medications. To fill this gap, we develop a Markov Decision Process (MDP) to characterize the optimal sequence (and combination) of antihypertensive medications under the standard medication dose. Our model is patient-centric as it takes into account a set of relevant patient characteristics such as age, gender, blood pressure level, smoking habits, diabetes status, and cholesterol level. Based on a set of intuitive assumptions, we prove that our model yields a series of structured optimal policies. Having calibrated our model based on real data and medical literature, we analyze these optimal policies and discuss their insights to the real practice. We also compare the benefits, in terms of quality adjusted life expectancy, QALE, obtained from our results with those obtained from British Hypertension Society (BHS) guideline.

1 Introduction and Literature Review

High blood pressure or hypertension (HTN) is a chronic medical condition in which the blood pressure (BP) in vessels elevates to an abnormally high level. More than 1 billion people in the world suffer from hypertension. In the US, alone, nearly one in three adults aged 40 to 59, and two in three in older adults have some degrees of high blood pressure (Go et al., 2014; Martin, 2008). According to (NNHS, 2004), HTN is the most frequent chronic condition among long-term stay residents aged 65 and more, and a major factor for many other chronic diseases such as myocardial infarction, stroke, heart failure, renal failure, and retinopathy. HTN is the most common reason for a primary-care consultation due to a chronic disease complication (Lovibond et al., 2011). In 2010, the direct and indirect cost of HTN has been estimated as $46.4 billion (Go et al., 2014). Forecasts also show that by 2030, the overall cost of HTN could amount to more than $274 billion (Lovibond et al., 2011; Mancia et al., 2007). Hypertension, also, is a key risk factor for the development of serious cardiovascular diseases (mainly heart attack and stroke) which are the leading causes of death: approximately 77% of the first stoke events occur among patients with high blood pressure (Mancia et al., 2007). The clinical benefits of controlling hypertension have been documented in numerous studies (Collins et al., 1990; MacMahon et al., 1990; Neaton et al., 1993). For example, every 10 mmHg reduction in BP alone has been estimated to contribute to 22% reduction in coronary heart disease (CHD) and 41% reduction in stoke (ST) (Law et al., 2009).

Despite the numerous serious health issues which are caused directly or indirectly, by hypertension and availability of effective treatments, the proportion of hypertensive patients who achieve satisfactory blood pressure control is still low. Only about one-half of all hypertensives even know that they have high blood pressure, only about 1/4 to 1/3 are currently under any treatment regimen, and only about 1/6 receive adequate blood pressure-lowering therapy (Wilber and Barrow, 1972). Lifestyle change can improve blood pressure control. The use of medications (i.e. antihypertensive therapy), however, is necessary for patients for whom the lifestyle change fails to control hypertension. Unfortunately around 2/3 of hypertensive patients do not achieve optimal blood pressure levels with a single drug, hence most patients are usually required to take multiple antihypertensive medications simultaneously to the rest of their lives on a daily basis (Saito et al., 2005). The main objective of antihypertensive therapy, hence, is to decide on the combination of medications so as to minimize both the likelihood of serious cardiovascular diseases (CVD), and the drug-related side-effects. In doing so, the key patient characteristics, i.e. diabetes status, blood cholesterol, smoking habits, age and gender, in addition to the clinical measurements of blood pressure level must be taken into account.

Most of medical guidelines on hypertension management include recommendations for antihypertensive prescriptions. However, there is a little agreement among them on the sequence with which the antihypertensive medications should be prescribed. For example, little agreement can be seen among them on which antihypertensive should be prescribed as the first-line medication (Saito et al., 2005): the US 8th Report of the Joint National Committee, JNC8, (James et al., 2014) and the revised World Health Organization and International Society of Hypertension (WHO/ISH, 2003) recommend diuretics (or thiazide) as the first-line drug, primarily for the cost considerations. The European Society of Hypertension, instead, does not provide any specific recommendation as the first-line treatment; according to this guideline, the selection of antihypertensives should be exclusively based on pathological conditions of patients, medication effectiveness, their side-effects and tolerability, but not based on price considerations (Saito et al., 2005). Hence the choice of the first-line and subsequent antihypertensives is a debatable medical decision issue.

When it comes to the literature search, hypertension management is one of the topics in chronic care management which has been less studied by operations researchers from a clinical decision support perspective. Excluding Mason et al. (2012), which is not entirely focused on hypertension management, the most recent paper in our literature search dates back to 1993, which is surprising given the existing challenges in the antihypertensive prescription policies discussed above. The majority of OR/MS literature on hypertension management has focused on screening (i.e. diagnosis) rather than treatment and control (i.e. chronic care). However, the common evidence that most hypertensive patient either do not pursue care or eventually discontinue it suggests that it may be reasonable to shift our attention from

screening to increasing the likelihood that patients achieve and maintain blood pressure control. (Schechter, 1988, 1990a, 1990b) studied a series of related problems in hypertension management taking the real blood pressure as unknown. These studies are based on Bayesian updating as the learning stage of the analysis. In the optimization stage and based on utility theory, the author seeks to identify, within a very theoretical context, the time when the measurements should be stopped, which also indicates the medication initiation time. In these studies, however, the author doesn't consider the issue of multiple antihypertensive agents. Nor does he take into account the risk of serious events as a consequence of elevated blood pressure. Finally, the dynamic nature of this problem is overlooked in these studies. The author acknowledges these limitations and particularly calls upon a dynamic framework as an appropriate model for the optimization task.

Optimal antihypertensive selection was also studied in (Moyé and Roberts, 1982). Using decision trees, the paper determines, from a cost perspectives, what protocol of antihypertensives is optimal as the therapeutic regimen. Recently, Mason et al. (2012) have studied management of diabetes using Markov Decision Process (MDP) framework. While their focus is primarily on diabetes, they have also considered as a secondary issue, the management of hypertension as it usually co-occurs with diabetes. Their analysis is mainly from cost perspectives, and they consider equal side-effects for the antihypertensive agents. In their analysis, they don't discuss the optimal policies for medication prescriptions.

The key issue is that little personalized (i.e. patient-centric) guidance is available regarding the number and combination of the five classes of medications currently being used to control hypertension. NICE (2006) warrants further research in this area. Moreover, as discussed earlier, hypertension has not been well studied by operations researchers from a clinical decision support perspective. To fill this gap, we develop a stochastic dynamic programming model to characterize the optimal sequence (and combination) of antihypertensive medications. Our model is patient-centric as it takes into account the set of relevant patient characteristics listed above. Along these lines, we explore the following research questions:

1. Do the optimal medication prescription policies have a certain structure that can provide insights to the clinicians?
2. Do the optimal policies depend on the patient characteristics, and, if yes, how?
3. What is the value of the personalized treatment policies?

The remainder of the paper is organized as follows: will be completed after finishing the manuscript.

2 Problem Statement

We define the **Hypertension Management Problem (HTNMP)** as whether to prescribe a new antihypertensive medication, considering the patient's clinical and demographic characteristics, in order to maximize his/her quality of life before a terminal event (first nonfatal CVD or death) during a problem horizon (e.g. 40-100 years of age). At each visit the physician decides whether to add a medication to the list of medications the patients is currently taking (if any). Following the common practice and the clinical guidelines, at each visit only one medication can be added, and once it is added, it cannot be removed afterwards (due to the higher side-effects associated with the removal). Since hypertension is chronic, the patients visit their physicians on a regular basis to ensure their blood pressure is under control. Figure 1 illustrates the timeline of decisions and events.

Figure 1: Timeline of Decisions and Events

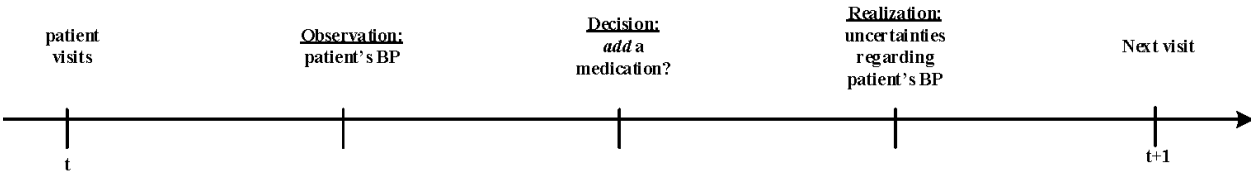

The information we use to determine optimal decisions includes systolic blood pressure, age, gender, diabetes status, blood cholesterol level, and smoking habits. We measure the quality of life as the patient's Quality Adjusted Life Years, QALYs, which is influenced by the presence of events (CVD or death) and the medication side-effects. With this definition, therefore, **HTNMP** can be seen as a primary intervention problem in which we maximize the patients' quality of life by minimizing their risk of facing serious CVD events and medications side-effects through controlling blood pressure level. The antihypertensive medications are selected based on their effectiveness (in reducing blood pressure) and side-effects (in reducing quality of life). Although blood pressure measurement might include noise, in this study, we analyze the noise-free measurement scheme as a benchmark. In this paper, we take into account the following most common antihypertensive categories: Beta Blockers (BB), Angiotensin-Converting-Enzyme Inhibitor (ACEI), Angiotensin II Receptor Blockers (ARBs), Diuretic (with the most common type being Thiazide), Calcium Channel Blockers (CCBs).

3 Problem Formulation

In this section, we present the mathematical formulation for the different components of the MDP model. First, we provide the key modeling assumptions. Next, we formulate the various MDP elements, which will lead us to present the Bellman Optimality Equation (i.e. the MDP model).

Our modeling is based on the following two fundamental assumptions regarding the treatment policies, which as discussed earlier, are based on medical guidelines and current practice:

Assumption 1 at each decision epoch, at most one medication can be added to the medication(s) the patient is taking.

Assumption 2 the medication prescription decision is non-reversible, meaning that once a medication is prescribed, cannot be removed until the problem terminates.

3.1 State Space

The state space at time t, denoted by x_t, contains all patient's information which are necessary for decision making at each visit. The state space includes two elements: patient's health status at time t, b_t, and her medication status, $\mathbf{m_t}$. Therefore:

$$x_t = (b_t, \mathbf{m_t})$$

- **Patient's health status:** We represent the patient's health status at each period by $b_t \in \mathcal{B} = \{B_0, B_0 + 1, ..., B, B + 1, B + 2, B + 3\}$ which includes information on her blood pressure status, her first nonfatal CVD events and death. In the above representation, B_0 and B are the minimum and maximum blood pressure levels for a hypertensive patient. Since the most serious consequences of high blood pressure are Cardiovascular Diseases (CVD) which mainly include Coronary Heart Disease (CHD) and Stroke (ST), we incorporate these two nonfatal CVD events in our model through the states $\{B + 1\}$ and $\{B + 2\}$, respectively. Finally, $\{B + 3\}$ represents all-but-CVD mortality. We don't include CVD mortality, because we explicitly consider their nonfatal occurrence before the fatal occurrence. For the ease of representation, we divide the patient's health states into two sets of states: non-terminal health state (i.e. event-free blood pressure level) $\tilde{b}_t \in \{B_0 ...., B\}$ and the terminal states $\tilde{\tilde{b}}_t \in \{B + 1,\ B + 2,\ B + 3\}$. In our model, $\{B+1\}$ and $\{B+2\}$ are semi-absorbing states as they will *virtually* be transitioned to the absorbing state $\{B+3\}$ (with zero reward) with certainty. We do so in order to avoid multiple occurrences of their nonzero reward in our optimization process.
- **Patient's medication status**: The vector $\mathbf{m_t}$ stores the list of antihypertensives the patient is taking at period t. Hence for a problem with M possible medications, we will need $\mathbf{m_t} = (m_{1t}, m_{2t}, \ldots, m_{Mt})$, where $m_{it} \in \{0, 1\}$ indicates whether the patient is taking medication i at time t.

3.2 Action Space

At each decision epoch t, the physician decides which one of the medications (if any) should be added to the list of medications the patients is taking. This corresponds to the classical stopping time problem where, in our case, it is defined as when to stop waiting until a new medication should be prescribed. For $M \geq 2$, the problem turns to a multiple or nested stopping time problem (i.e. one stopping time problem per each medication). Considering the two assumptions discussed earlier, the action space at each decision epoch t is defined as the following:

$$\mathrm{A}(x_t) = \mathrm{A}(b_t, \mathbf{m}_t) = \{0, i\} \setminus \{i'\} \ \forall i, i' : m_{it} = 0, \ m_{i't} = 1$$

Hence the possible action at each decision epoch is $a(x_t) \in \mathrm{A}(x_t)$, where $a(x_t) = i$ means adding medication i , and $a(x_t) = 0$ means adding no medication (i.e. "do-nothing" or "wait" in the stopping time problem terminology) .

3.3 Transition between States

For the two state variables b_t and $\mathbf{m}_t$, we need to characterize the transition functions which determine the state space in the next period as a function of the current states and decision. While the transition between medication states is deterministic, the transition between health states is stochastic, governed by a Markov chain $P_t[b_{t+1} | b_t, a_t]$.

3.3.1 Medication State Transition Function

The transition function for the medication state, $\mathbf{G}(\mathbf{m}_t, a_t)$ is a deterministic function which identifies the next medication state given the current medication state and decision. Mathematically, the elements of vector $\mathbf{G}$, denoted g, can be identified in the following way:

$$g(m_{it}, a_t) = \max\left(m_{it}, I_{(a_t = i)}\right)$$

where $I_{(\cdot)}$ is the indicator function defined as:

$$I_{(a_t = i)} = \begin{cases} 1 & i > 0 \\ 0 & \text{otherwise} \end{cases}$$

it assigns 1 if the statement is true, 0 otherwise.

3.3.2 Health State Transition Function

The transition function for the non-terminal health state $\tilde{b}_t \in \{B_0, \ldots, B\}$ is governed by the Markov chain $P[\tilde{b}_{t+1}|\tilde{b}_t, a_t]$ which depends on the health state and action in period t. The elements of this transition matrix are $p_t(\tilde{b}_{t+1}|\tilde{b}_t, a_t)$, in other words $P_t(\tilde{b}_{t+1}|\tilde{b}_t, a_t) = [p_t(\tilde{b}_{t+1}|\tilde{b}_t, a_t)]$.

In the absence of any medication treatment, an individual's blood pressure in the next stage, $\tilde{b}_{t+1}$, transitions approximately according to a normal distribution with mean β_{t+1} and variance σ^2 (O'Sullivan et al., 1999; Whelton et al., 2002):

$$\tilde{b}_{t+1} \sim N\{\beta_{t+1}, \sigma^2\}$$

To incorporate the effect of aging, we consider an upward linear trend for the change in blood pressure level as the result of aging. Given these considerations, the mean of the normally distributed blood pressure in the next stage can be represented as:

$$\beta_{t+1} = \beta_t + \gamma \Delta t$$

where β_t is the mean (or observed) blood pressure at time t, γ is the slope of the annual blood pressure change with age, and Δt is the length of time between two successive decision epochs. In order to incorporate the role of medication prescription in the transition, we modeled the medication effectiveness as the reduction in the mean of the blood pressure distribution (USDHHS, 2011). Figure 2 illustrates the concept more clearly.

Fig1: Impact of Medication Effectiveness on Transition Probability

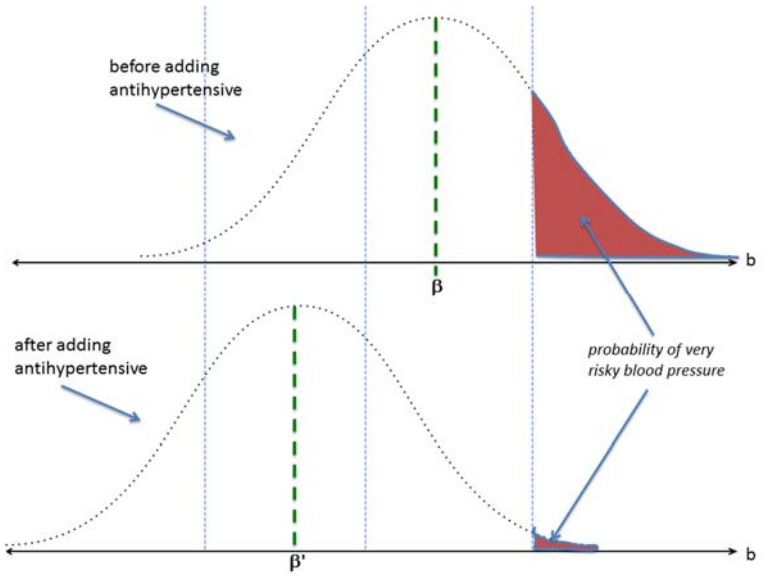

We assume that each medication i, immediately and deterministically, reduces the mean of blood pressure by ϵ_i mmHg. Hence the mean of the blood pressure distribution in the next period, β_{t+1}, when medication i is prescribed to the patient whose current blood pressure level is β_t, will be computed as:

$$\beta_{t+1} = \beta_t - \epsilon_i I_{(a=i)} + \gamma \Delta t$$

Hence in the absence of any medication prescription, we get back to $\beta_{t+1} = \beta_t + \gamma \Delta t$.

This normal transition function, however, governs only the transition between blood pressure states $\tilde{b_t}$, and not the transition from blood pressure states to the terminal state $\tilde{\tilde{b}}_t$, nor does it consider the occurrence of these terminal events. In order to identify the transition between all health states, however, we need three other elements: we denote by $\alpha_{1t}(\tilde{b}, a)$, $\alpha_{2t}(\tilde{b}, a)$ and $\alpha_{3t}(\tilde{b}, a)$ the probabilities of moving from blood pressure states $\tilde{b}_t$ to the terminal states $\{B+1\}$, $\{B+2\}$ and $\{B+3\}$ at time $t+1$, respectively. Taking into account the probability of death α_{3t}, one can identify these probabilities of facing $\{B+1\}$ (i.e. CHD) and $\{B+2\}$ (i.e. ST) as $\alpha'_{1t} = (1-\alpha_{3t})\alpha_{1t}$ and $\alpha'_{2t} = (1-\alpha_{3t})\alpha_{2t}$, respectively. Given the above considerations and noting that the nonfatal events are semi-absorbing, we can now fully characterize the health state transition probability function as:

$$p_t(b_{t+1} \mid b_t, a_t) = \begin{cases} \left(1 - \sum_{j=1}^{3} \alpha'_{jt}(b_t, a_t)\right) \widetilde{p_t}\left(b_{t+1} \mid b_t, a_t\right) & b_t, b_{t+1} \in \{B_0, ..., B\} \\ \alpha'_{1t}(b_t, a_t) & b_t \in \{B_0, ..., B\},\ b_{t+1} = \{B+1\} \\ \alpha'_{2t}(b_t, a_t) & b_t \in \{B_0, ..., B\},\ b_{t+1} = \{B+2\} \\ \alpha_{3t}(b_t, a_t) & b_t \in \{B_0, ..., B\},\ b_{t+1} = \{B+3\} \\ 1 & b_t = \{B+j\} : j = 1,2,3,\ b_{t+1} = \{B+3\} \\ 0 & \text{otherwise} \end{cases}$$

with $\alpha'_{3t} = \alpha_{3t}$.

3.4 Reward Functions

3.4.1 Immediate Reward Function

The immediate reward function $r_t(b_t, \mathbf{m}_t, a_t)$ captures the one-period quality adjusted life years (QALYs) of the patient who is in state $x_t = (b_t, \mathbf{m}_t)$ and follows decision a_t. QALY captures both quantity and quality of life years. It is equal to 1 for a patient with perfect health (for a full year) and 0 for death. In our

problem, hence, it is affected by the burden (i.e. the side-effects) of medication, denoted by δ_i, which captures the decrements in the quality of life as a consequence of using medication i. For a patient who will not experience any terminal events, we express the event-free immediate reward function as:

$$r'_t(b_t, \mathbf{m_t}, a_t) = \Delta t\left(1 - \sum_i \delta_i m_{i,t+1}\right)$$

where $m_{i,t+1}$ is the i^{th} element of the updated medication state $\mathbf{m_{t+1}}$.

This formulation, however, assumes no event will occur on $[t, t+1)$. In order to account for the possibility of reaching a terminal event during period t, we assign QALYs to these events and consider their probability of occurrence during that period. Since the occurrence of such events during a period is unknown, we assume they occur in the middle of the period. For instance, for the terminal state $\{B+3\}$ with QALY=0, the patient experiences ½ year of event-free quality of life, i.e.:

$$Q(B+3) = \frac{1}{2}(r'_t(.) + 0) = \frac{r'_t(.)}{2}$$

Denoting by ω_j, the quality of life for a patient with event $B+j$: $j \in \{1, 2\}$, then:

$$Q(B+j) = \frac{1}{2}\left(r'_t(.) + \omega_j\right), j \in \{1,2\}$$

Summarizing all the above considerations, we can now represent the expected immediate reward function as:

$$\begin{aligned} r_t(b_t, \mathbf{m_t}, a_t) &= Pr(alive)\,Quality(alive) + \sum_{j=1}^{3} Pr(B+j)\,Quality(B+j) \\ &= \left(1 - \alpha'_{1t}(b,a) - \alpha'_{2t}(b,a) - \alpha_{3t}(b,a)\right) r'_t(.) + \sum_{j=1}^{3} \alpha'_{jt}(b_t, a_t) Q(B+j) \end{aligned}$$

It follows that $r_t = r'_t$ in the absence of any terminal event (i.e. $\alpha_{1t} = \alpha_{2t} = \alpha_{3t} = 0$).

3.4.2 Lump-Sum Terminal Reward Function

We denote by R_{jt}, the total expected terminal reward for a patient who experiences a nonfatal terminal event j (i.e. CHD or ST) in period t. In other words, R_{jt} is the life-expectancy of a patient who experiences his first event j in period t: It follows that $R_{3t} = 0$ for death. By taking the nonfatal events as semi-absorbing, we make sure that $R_{jt}, j=\{1,2\}$, occur only once for a patient who faces the nonfatal events.

3.5 The MDP Formulation

Denoting by $v_t(x_t)$, the maximum total expected quality-adjusted life before the terminal events for a patient on state x_t in period t, the optimal solution to **HTNMP** can be obtained by solving the following Bellman Optimality Equation:

$$v_t\left(b_t, \mathbf{m_t}\right) = \begin{cases} \max\limits_{a_t \in A_t\left(b_t, \mathbf{m_t}\right)} \left\{ r_t\left(b_t, \mathbf{m_t}, a_t\right) + \lambda^{\Delta t} \sum\limits_{b_{t+1}} p\left(b_{t+1} \mid b_t, a_t\right) v_{t+1}\left(b_{t+1}, \mathbf{m_{t+1}}\right) \right\} & b \in \left\{B_0, ..., B\right\} \\ R_{jt}\left(b_t, \mathbf{m_t}\right) & b \in \left\{B + j\right\} : j \in \left\{1, 2, 3\right\} \\ R_T & t = T \end{cases}$$

where:

- λ is the annual discount factor ($0 \le \lambda \le 1$),
- R_T is the life expectancy of a patient at the end of the decision horizon. For instance $R_{100} = 0$, assuming that patients die at the age of 100, otherwise R_T should be estimated from the available life tables.

4 Structural Properties

In this section, we prove that threshold policy (aka monotone policy or control-limit policy) is an optimal solution to the **Hypertension Management Problem (HTNMP)**. Without loss of generality, we present the proofs for $M = 1$ (namely the hypertension management problem with one medication, hence the decision becomes choosing between prescribing the single medication $a = 1$, or waiting $a = 0$). This is due to the existing order between medications (both in terms of effectiveness and side-effects): the higher the effectiveness is, the higher will be the side-effects, and vice versa. Otherwise one decision will be dominated by another, hence will be removed from the list of medications. We ordered the medications in the ascending rank of their effectiveness (or side-effects). Hence the proof is true for any pairs of $a \in \{i^+, i^-\} : i^+ > i^-$. For any different ordering, our requirements for optimality of the threshold policy (which will be discussed in this section) will not be fulfilled.

First, we need to review some technical definitions and assumptions which are relevant to our problem.

4.1 Technical Definitions

1. A transition probability matrix P (and the equivalent Markov chain) is called IFR (increasing failure rate) if its rows are in increasing stochastic order. In other words:

$$\rho(b) = \sum_{k=b'}^{B+3} P(k \mid b, a)$$

is nondecreasing in b for all $b' \in \{B_0, B_0 + 1, ..., B + 2,\ B + 3\}$. This definition is equivalent to the well-known notion of first-order stochastic dominance which has an intuitive interpretation in the context of our problem: the higher the level of blood pressure, the more likely is moving to even higher blood pressure or terminal events.

2. A function $g(x, a)$ is *superadditive* if for $x^+ \geq x^-$ and $a^+ \geq a^-$, then:

$$g(x^+, a^+) + g(x^-, a^-) \geq g(x^+, a^-) + g(x^-, a^+)$$

If the reverse inequality holds, $g(x, a)$ is said to be *subadditive*.

A common variation of the statement of a superadditive function is the following equivalent condition:

$$g(x^+, a^+) - g(x^-, a^+) \geq g(x^+, a^-) - g(x^-, a^-)$$

This expression means that the incremental change in x for a larger value of a is greater than for a smaller value of a. In other words, $g(x^+, a) - g(x^-, a)$ is nondecreasing in a or $g(x, a^+) - g(x, a^-)$ is nondecreasing in x.

4.2 Technical Assumptions

In this study, we make the following technical assumptions which are necessary for the proof of optimal threshold policy.

As1. $\alpha'_{1t}(b_t, a_t)$, $\alpha'_{2t}(b_t, a_t)$ and $\alpha_{3t}(b_t, a_t)$ are (i) nondecreasing in b_t and (ii) subadditive.

As2. (i) $r_t(b_t, \mathbf{m}_t, a_t)$ and (ii) $r_T(b_t, \mathbf{m}_t)$ are nonincreasing in b_t.

As3. $q(k \mid b,a) = \sum_{b'=k}^{\infty} p(b' \mid b,a)$ (the reverse cumulative distribution function for the transition matrix) is (i) IFR and (ii) nondecreasing in b.

As4. $r_t(b_t, \mathbf{m}_t, a_t)$ is superadditive.

As5. $q(k|b, a)$ is subadditive.

Note that threshold policy is only expressed for a state variable with physical interpretation and clear ordering. Since the medication state **m** (i.e. the list, and not the number, of medications) doesn't have such property, we present our proofs for a fixed medication state.

Note also that all the above assumptions have intuitive interpretations in our problem: **As1(i)** means as the blood pressure increases, the risk of encountering terminal events does not decreases. **As1(ii)** states that such risk doesn't decrease if we do nothing (i.e. $a = 0$) than if we prescribe a medication (i.e. $a = 1$). In other words $\alpha'_{jt}(b_t, 0) - \alpha'_{jt}(b_t, 1)$ is nondecreasing in b_t, $\forall j \in \{1, 2, 3\}$ and $\alpha'_{3t} = \alpha_{3t}$. **As2** means that none of the reward functions (i.e. immediate and terminal rewards) increase with blood pressure level. We previously discussed the intuitive interpretation of IFRness of $q(k|b, a)$ in **As3(i)**. **As3(ii)** implies that as the blood pressure increases, the risk of reaching a certain level of blood pressure and higher (including terminal events) does not decrease. Note that this is similar to the IFRness property of the transition function. **As4** means the increase in quality of life as a result of taking a medication versus doing nothing increases with blood pressure level (in other words, blood pressure reduction has a greater influence at higher blood pressure levels), or $r(b, 1) - r(b, 0)$ is nondecreasing in b:

$$r(b^+, 1) - r(b^+, 0) \geq r(b^-, 1) - r(b^-, 0)$$

As5 implies:

$$q(k|b^+, 1) - q(k|b^+, 0) \leq q(k|b^-, 1) - q(k|b^-, 0)$$

since $q(k|b, 1) \leq q(k|b, 0)$, we can rewrite the above equation as:

$$q(k|b^+, 0) - q(k|b^+, 1) \geq q(k|b^-, 0) - q(k|b^-, 1)$$

which means the risk of transitioning to a certain blood pressure level or higher, including the terminal events (i.e. $b \geq k$) if we do nothing versus if we prescribe the medication increases with blood pressure level which is a quite intuitive assumption, again based on the greater blood pressure lowering effects in higher blood pressure levels. Note that since α'_{jt} is an element of transition probability matrix (hence an element of $q(k|b,a)$), **As1(i)** and **As1(ii)** are specific cases of **As3(i)** and **As(5)**, respectively, hence it might be dropped from the list of assumptions.

4.3 Optimality of Threshold Policy for the Hypertension Management Problem

In this section we prove optimality of threshold policy for the **HTNPM** given the set of definitions and assumptions discussed earlier. The proofs are available in the Appendix.

Define:

$$H_t(b_t, a_t) = r_t(b_t, a_t) + \lambda^{\Delta t} \sum_{b_{t+1}} p_t(b_{t+1} \mid b_t, a_t) \upsilon_{t+1}(b_{t+1})$$

Hence:

$$\upsilon_t(b_t) = \max_{a_t} \left[H_t(b_t, a_t) \right]$$

Proposition 1 Suppose that $\max_{a_t} \left[H_t(b_t, a_t) \right]$ is attained, **As2** and **As3(ii)** hold. Then $\upsilon_t(b_t)$ is nonincreasing in b_t for $t = 1, ..., T$.

Lemma 1 Suppose $H_t(b_t, a_t)$ is a superadditive function on $\mathcal{B} \times \mathrm{A}$ and for each $b \in \mathcal{B}$, $\max_{a_t} \left[H_t(b_t, a_t) \right]$ exists. Then

$$a_t^*(b_t) = \max \left\{ a'_t \in \arg\max \left[H_t(b_t, a_t) \right] \right\}$$

is monotone nondecreasing in b.

Lemma 2 Let two reverse cumulative distribution functions be such that:

$$\sum_{j=k}^{\infty} p_t(j \mid b) \leq \sum_{j=k}^{\infty} p_{t'}(j \mid b)$$

for all k. For a nonincreasing function $v(j)$; $j = 0, 1, ...$ then

$$\sum_{j=k}^{\infty} p_t(j \mid b) v(j) \geq \sum_{j=k}^{\infty} p_{t'}(j \mid b) v(j)$$

Proposition 2 Suppose the maximum in $H_t(b_t, a_t)$ is attained. Given **As2** and **As3**, $v_t(b_t)$ is nonincreasing in b_t for $t = 1, 2, ..., T$.

Theorem 1 Suppose for $t = 1, 2, ..., T$, **As2-As5** hold. Then there exist an optimal policy $a_t^*(b_t)$ which is nondecreasing in b_t for $t = 1, 2, ..., T - 1$. In other words, at each period t, there is a threshold b_t^* such that

$$a_t^*(b) = \begin{cases} 0, & b < b_t^* \\ 1, & b \geq b_t^* \end{cases}$$

5 Parameter Estimation

In this section, we describe the procedure for estimating the parameters in our MDP model. The parameters in our model were estimated from two sources: the MGH Clinic dataset as the primary data, and the medical literature as the secondary data. Our model includes the following set of parameters to be estimated for the numerical analysis:

1. The variance of the normal probability distribution which characterizes the transition between blood pressure states (i.e. σ^2)
2. Medication effectiveness (i.e. ϵ_i)
3. Medication side-effects (i.e. δ_i)
4. Probabilities of facing terminal events (i.e. $\alpha_{jt}(.)$)
5. Slope of the blood pressure trend over time (i.e. γ)
6. Disutility of the nonfatal events (i.e. ω_j)
7. Lump-sum terminal reward (i.e. R_{jt})
8. Discount factor (i.e. λ)

5.1 The variance of the Normal transition function between health states

We estimated σ^2 using the data that we collected in Montreal General Hospital (MGH) Clinic. In doing so, we stratified the data on gender and removed those patients with less than 3 visits. Table 1 provides more details regarding the baseline characteristics in this dataset.

Table 1: Baseline Characteristics for MGH Data

Factor	Value
total number of patients	109
male	54 (50%)
female	55 (50%)
min age	21
max age	84
min number of visits	1
max number of visits	8
number of visits mean	3.2
number of visits median	3
total number of visits	347
earliest visit	Aug. 2008
latest visit	Feb. 2013
total number of valid cases (i.e. patients with more than 2 visits)	57 (52%)
total number of valid male cases	26 (46%)
total number of valid female cases	31 (54%)
total number of valid visits	278

Since σ^2 is related to the transitions between health states in the absence of medications, we needed to *de-medicate* the observations which were collected when the patient was on medication. Since we assumed a deterministic medication effectiveness, hence we performed the de-medication by reverting the measurements to their medication-free status. For example, for a patient with observation $\beta = 150$ mmHg who is on a medication with 10mmHg reduction in blood pressure (i.e. $\epsilon = 10$), the de-medicated observation will be $\beta'' = 160$.

The next challenge in estimating σ^2 from data is the missing observations. Fixing the decision epochs at quarterly intervals, we used smoothing spline technique to estimate the missing observations (Alagoz, 2004). More specifically, we used cubic spline technique to interpolate the missing observations after the de-medication procedure. For instance, for a patient with the de-medicated observations 140, 150 and 130 at quarters 1, 3 and 4, the data is missing at $t = 2$. Using these data, the Cubic spline estimates $\beta_2 = 153$ mmHg (Figure 3).

Figure 3 Estimating missing observations using Cubic Spline Technique

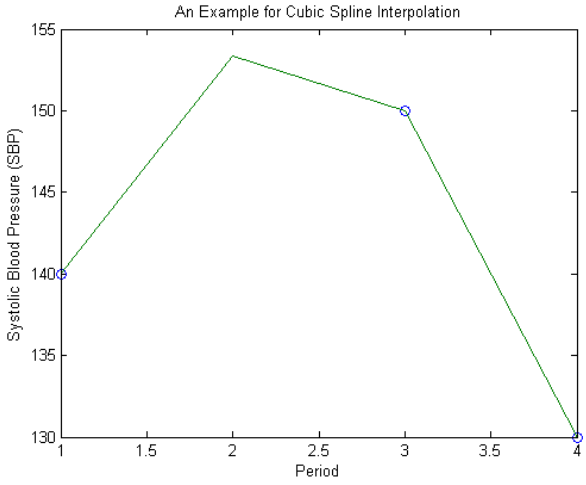

In our study, we also assumed that the volatility (i.e. variance) of the transition functions does not depend on age (we didn't find any literature which claims otherwise). Finally, for each gender-specific stratification, we estimated σ^2 from the changes in the de-medicated blood pressure measurements. The results are summarized as follows:

$$\sigma^2_{MALE} = 17.23$$
$$\sigma^2_{FEMALE} = 9.61$$

These estimated parameters ascertain that male patients experience a lager fluctuation in their blood pressure than females.

5.2 Medication Effectiveness

Medication effectiveness ϵ_i is defined as the extent by which the blood pressure is reduced after taking antihypertensive i. (Law et al., 2003) have provided estimation for ϵ_i for the five categories of medications in our study. They have estimated ϵ_i (for the standard dose) as a linear function of the patient's systolic blood pressure in the following way:

$$\epsilon_1(SBP) = 0.1SBP - 6.9$$
$$\epsilon_2(SBP) = 0.1SBP - 6.6$$
$$\epsilon_3(SBP) = 0.1SBP - 6.5$$
$$\epsilon_4(SBP) = 0.1SBP - 6.3$$
$$\epsilon_5(SBP) = 0.1SBP - 5.1$$

The above parameters clarify that among the five categories, m_1 and m_5 are the least and most effective medications in reducing blood pressure, respectively. These parameters also identify the order of

the antihypertensive medications (from lowest effectiveness to the highest effectiveness or from the lowest side-effects to the highest side-effects) as the following:

1. Angiotensin-Converting-Enzyme Inhibitor (ACEI)
2. Diuretics (D)
3. Calcium Channel Blockers (CCB)
4. Beta Blockers (BB)
5. Angiotensin II Receptor Antagonists (ARBs)

5.3 Probability of Nonfatal Events

We estimated $\alpha_{1t}(.)$ and $\alpha_{2t}(.)$ from the Framingham Study, one of the most well-known epidemiological studies conducted to date (D'Agostino et al., 2008). The study provides simple equations for the 10-year risk of CVD, for both female and male, as a function of age (t), total cholesterol (CLT), High-Density Lipoprotein cholesterol (HDL), Systolic Blood pressure (SBP), Smoking Status (SMK) and Diabetes Status (DIAB). The study also provides calibration factors in order to quantify the risk of CVD-specific events being CHD or ST from the computed CVD risk (which are interpreted as the proportions of CVD events being CHD or ST). Using these equations, we not only estimate the risk of nonfatal events, but also add further personalization to our problem by incorporating the effect of the two major chronic diseases which usually co-occur with hypertension (i.e. diabetes and blood cholesterol), in addition to other risk factors (i.e. gender and smoking) which contribute to the risk of the CVD events. We used the following four equations to estimate $\alpha_{1t}(.)$ and $\alpha_{2t}(.)$ for male (M) and female (F):

$$\alpha_{1t}^{M} = k_1^{M}\left(1-0.95021^{r_M}\right)$$
$$\alpha_{2t}^{M} = k_2^{M}\left(1-0.95021^{r_M}\right)$$
$$\alpha_{1t}^{F} = k_1^{F}\left(1-0.88936^{r_F}\right)$$
$$\alpha_{2t}^{F} = k_2^{F}\left(1-0.88936^{r_F}\right)$$

where $k_j^h,\ h \in \{M, F\},\ j \in \{1, 2\}$ are the gender-specific calibration factors for the two nonfatal events which were estimated as:

$$k_1^M = 0.7174$$
$$k_2^M = 0.1590$$
$$k_1^F = 0.6086$$
$$k_2^F = 0.2385$$

$$r_F = exp\{2.3288Log(t) + 1.20904Log(CLT) - 0.70833Log(HDL) + 2.67157Log(SBP) + 0.52873SMK + 0.69154DIAB - 26.1931\}$$

$$r_M = exp\{3.06117Log(t) + 1.12370Log(CLT) - 0.93263Log(HDL) + 1.93303Log(SBP) + 0.65451SMK + 0.57367DIAB - 23.9802\}$$

5.4 Probability of Non-CVD-Mortality

We used (Lewington et al., 2002), another seminal paper in the medical literature, to estimate the risk of non-CVD-mortality which excludes the mortalities caused by the CVD events (i.e. CVD-Specific-Mortality) from All-Cause-Mortality. In this paper, the probability of stroke mortality was estimated as a function of age and SBP, using the following exponential risk equations:

$$1000p^{ST}(t, SBP) = exp\{a_t + b_t SBP\}$$

The numerical values of a_t and b_t were provided for different age-groups as in Table 2.

Table 2 Numerical values of the parameters of the stroke mortality risk equations

Age Group	a_t	b_t
40-49	-14.647	0.051
50-59	-12.617	0.048
60-69	-10.821	0.042
70-79	-8.502	0.035
80-89	-5.159	0.020

The probability of CVD-mortality was estimated as:

$$p^{CVD-Mortality}(t, SBP) = \frac{1}{k_2} p^{ST}(t, SBP)$$

with $k_2 = 0.20$ (the average of k_2^M and k_2^F in the Framingham study, as the CVD-specific mortality in this study was not gender-specific).

We also assumed that at each age, the Non-CVD-mortality differs from the CVD-mortality by $\varphi(t)$ which was estimated in [9] using the following regression equation:

$$\varphi(t) = 2\times10^{-8}t^4 - 3\times10^{-6}t^3 + 2\times10^{-4}t^2 - 0.59\times10^{-2}t + 0.0642;\ R^2 = 0.9999$$

Finally, the Non-CVD-Mortality (or All-But-CVD-Mortality) were estimated using the following risk equations:

$$p^{Non-CVD-Mortality}(t, SBP) = p^{CVD-Mortality}(t, SBP) + \varphi(t)$$

5.5 Post Events Life Expectancy

We defined the Lump-Sum-Terminal Rewards for the non-fatal events, namely $R_1(t)$ and $R_2(t)$, as the Quality Adjusted Life Expectancy (QALE) after facing CHD and ST, respectively. These elements quantify both quantity and quality of years being lived after facing the two nonfatal events. Hannerz and Nielsen (2001) have estimated the gender-specific Life Expectancy (LE) for CVD events. Using those equations, we quantified the gender-specific QALE for both nonfatal events as:

$$R_j^h(t) = \omega_j LE^h(t) \qquad h \in \{M, F\}, j \in \{1, 2\}$$

where ω_j is the disutility of the nonfatal event $\{B + j\}$, and $LE(t)$ was estimated for male and female using the following equations:

$$LE^F(t) = 0.0073t^2 - 1.4696t + 76.425;\ R^2 = 0.9999$$
$$LE^M(t) = 0.0073t^2 - 1.3673t + 68.683;\ R^2 = 0.9994$$

The above equations ascertain that female patients experience a longer life after CVD events than male. Multiplying the above equations by the quality of life of being on CHD or ST (i.e. ω_1 and ω_2), we quantified QALE for both CHD and ST (i.e. $R_j^h(t)$).

5.6 Other parameters

Table 3 summarizes the estimated parameters in our model.

Table 3 Numerical values of the remaining parameters

Parameter	Estimated Value	Source
Slope of the linear blood pressure trend over time (γ)	γ_{male}=0.30	(Wilkins et al., 2010)
	γ_{female}= 0.37	
Disutility of the nonfatal events (ω_j)	ω_1=0.23	(Lovibond et al., 2011)
	ω_2=0.37	
Discount Factor (λ)	0.996	(Lovibond et al., 2011)
Side-effects (disutility) of medications (δ_i)	δ_1 =0.010	(Lawrence et al., 1996; Tengs and Wallace, 2000)
	δ_2 =0.015	
	δ_3 =0.016	
	δ_4 =0.018	
	δ_5 =0.022	

6 Computational Results and Conclusions

In this section, we provide further details for solving our MDP model base on the threshold property of its optimal solutions. We illustrate the optimal policies for a select few examples, based on our estimated parameters, and we discuss how to interpret these policies. We also compare the benefits, in terms QALE, obtained from our policies with those obtained from British Hypertension Society (BHS) guideline. Finally, we discuss the general insights, which can be derived from our results.

6.1 Properties of the Optimal Policies

We proved in section 3 that the optimal solutions to our problem follow a threshold structure (aka *threshold*, *control-limit* or *monotone* policy). This is of interest, particularly because of the ease of interpretation and implementation of such policies. A finite horizon MDP model admitting threshold policies can be solved using *Monotone Backward Induction* (Puterman, 2014). The importance of this algorithm is that the maximization is carried out over the decision set $A'(b)$ which becomes smaller with b. Hence, if at some b', $A'(b')$ contains only one element, then no further computation is needed since the action will be optimal for all $b \geq b'$. In the worst case $A'(b) = A(b)$, where the number of iterations will equal those of the traditional backward induction.

6.2 Numerical Example

Assuming the set of intuitive properties for the components of our MDP model (discussed before), we characterized the structure of optimal prescription policy. Accordingly, the optimal policy consists of a series of threshold levels for each medication. Furthermore, threshold levels are nested within each other, which imposes an optimal sequence in which medications are prescribed. Also, since thresholds are both time-dependent and patient-specific, both the optimal sequence and duration between two consecutive prescriptions will be customized to the individual characteristics of the patient.

In this section, we select, as a case study, a risk free male patient (i.e. non-smoker, non-diabetic with low cholesterol level) and present the optimal antihypertensive medication policy for such a patient. Figure 4, for example, illustrates the optimal policy for such a patient given he is not taking any medication. In other words, this figure exhibits the antihypertensive initiation policy for such patient proposed by our model.

Figure 4 Optimal Medication Prescription for a Risk Free Male Patient Taking No Medication

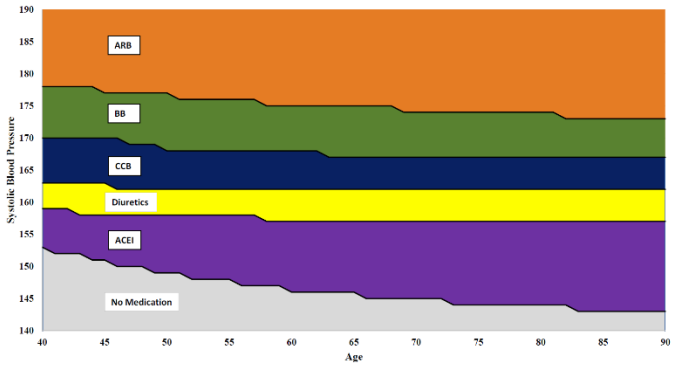

The optimal policy for this case ascertains that the higher the blood pressure, the more aggressive (i.e. more effective, although with greater side-effects) should be the medication initiation decision (in order to maximize his quality-adjuster life years). If this patient has the age of 40, the optimal policy is not to prescribe any medication if his systolic blood pressure (SBP) is less than 152 mmHg. If his systolic blood pressure is between 152 and 158, then we need to prescribe ACEI, and so on. As can be realized, these simple policies provide informative insights to the sequence of medication prescriptions. In all cases, the threshold levels for the medications were ranked, from low to high, based on their effectiveness (hence their side-effect). It follows that for $SBP_1 \leq SBP_2$, it is never optimal to prescribe a medication with higher effectiveness (hence higher side-effect) at SBP_1 while prescribing another medication with lower effectiveness at SBP_2. Note that this doesn't mean that we never start with a medication of higher effectiveness. As Figure 4 depicts, for a patient with the above characteristics, the optimal policy recommends to start with the medication of the highest effectiveness, if the patient's blood pressure exceeds 180 mmHg. The optimal policies for all case-mixes are not only of threshold type, but their thresholds also are non-increasing in age: as the patient gets older, the blood pressure level above which the medications should be prescribed never increases, which is intuitive given that age is a risk factor for CVD events. The observation is especially more evident for the first medication in the list as it shows a steeper initiation threshold.

For an instance of a problem with M medications, the complete optimal policy consists of $\sum_{i=1}^{M-1}\binom{M}{i}$ threshold policies (one per each medication state). Hence, Figure 4 provides the optimal policy only for patients who have not started taking any medication. In the next visit, we need to update the patient's medication state, hence reading his optimal policy from a different graph. Take, for instance, a risk free male patient with systolic blood pressure=160 mmHg. The optimal decision for such a patient is to take ACEI. The optimal decision for the same patient one year later should be read from Figure 5 which pertains to patients who are taking ACEI.

Figure 5 Optimal Medication Prescription for a Risk Free Male Patient on ACEI

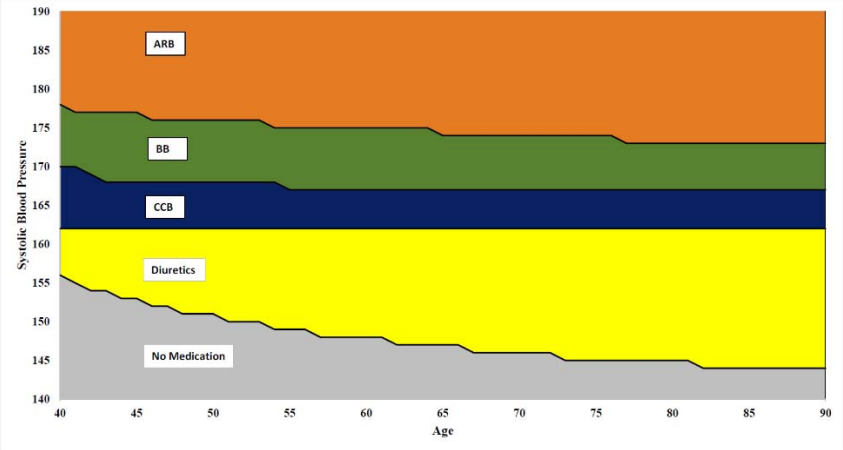

Figures 6-8 depict the optimal policies for patients who are taking medications {ACEI and Diuretics}, medications {ACEI, Diuretics and CCB}, and medications {ACEI, Diuretics, CCB and BB}, respectively. As can be seen in Figure 8, the only available option is either to wait (i.e. doing nothing) or to prescribe the remaining medication in the list (i.e. ARB). An interesting observation is that as the number of medications gets larger, the waiting decision area becomes larger too (i.e. the blood pressure threshold for adding the medication shifts upward). This means that the value of adding a new medication reduces with the number of previously-added medications. This is a medically-established observation (known as the diminishing marginal effectiveness of medications) that our modeling paradigm was able to capture (Timbie et al., 2010).

Figure 6 Optimal Medication Prescription for a Risk Free Male Patient on Medications ACEI and Diuretics

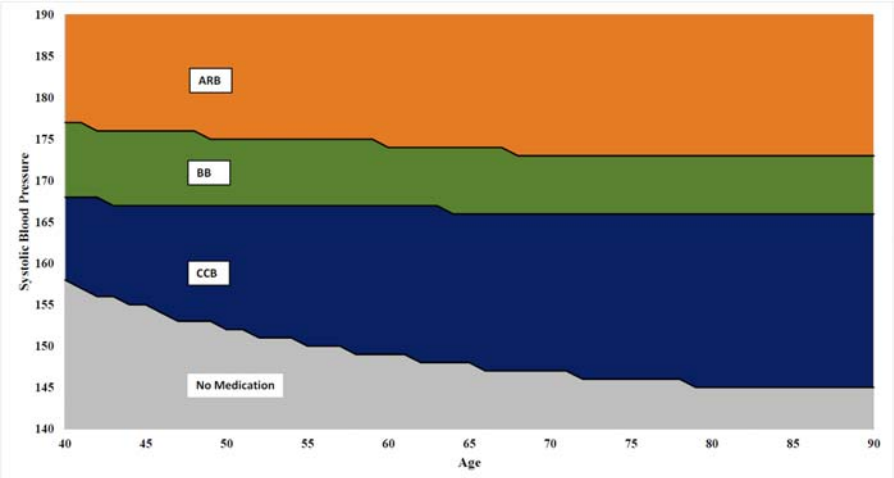

Figure 7 Optimal Medication Prescription for a Risk Free Male Patient on Medication ACEI, Diuretics, and CCB

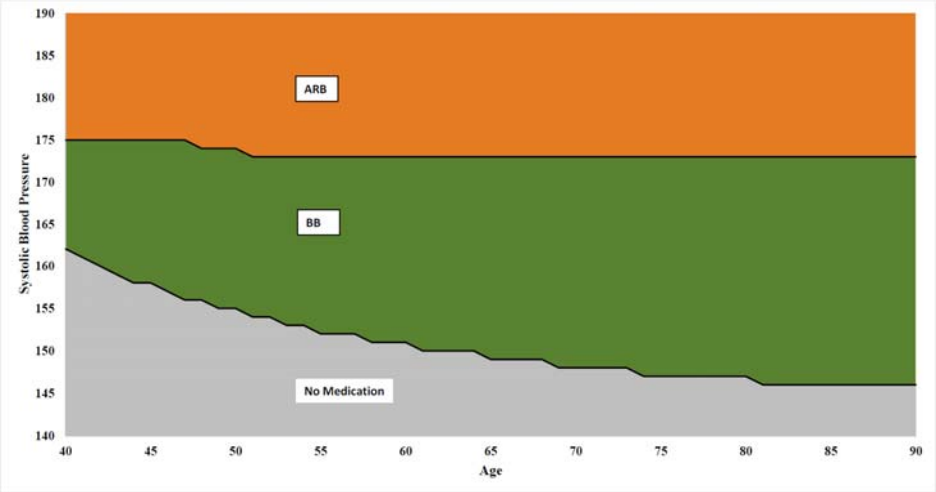

Figure 8 Optimal Medication Prescription for a Risk Free Male Patient on Medications ACEI, Diuretics, CCB, and BB

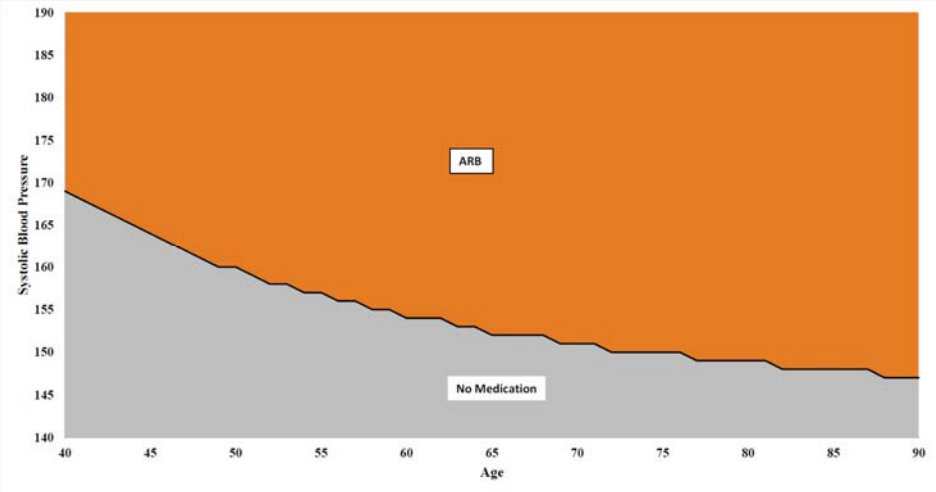

Our results are also consistent with a series of medical observations, such as:

1. The medications need to be prescribed at a lower threshold for a male than for a female. This is consistent with the fact that being male alone is a risk factor for CVD events [ask Dr. Daskalopoulou for the ref].
2. The medication need to be prescribed at a lower threshold for a male who is smoker than when he is diabetic. This suggests that for a male, being smoker is riskier than being diabetic. We have observed the opposite results for a female, inferring that for a female, being diabetic is riskier than being smoker [ask Dr. Daskalopoulou for the ref].

6.2.1 Comparison with Guidelines

Among several prevailing guidelines, our optimal policies are closest, in terms of structure, to the British Hypertension Society (BHS) guideline. Although the optimal sequence of medication induced by our policy is in alignment with BHS guidelines in principle, it differs from the guidelines in the sense that the sequence of medications recommended by BHS guideline is not as individualized as our recommendations (National Clinical Guideline Centre (UK), 2011). In order to quantify the value of this individualized treatment, we compared the Quality Adjusted Life Expectancy (QALE) of a patient under two scenarios: when the patient follows BHS recommendations, and when the patient follows the recommendations derived from our optimal policies. Figure 9 illustrates the improvements in QALE, as a result of following our optimal procedure versus the BHS recommendations. In this figure, risk-free category corresponds to being non-smoker, non-diabetic, and having low cholesterol level.

Figure 9: Improvement in QALE when a patient follows our optimal policy rather than following BHS guidelines

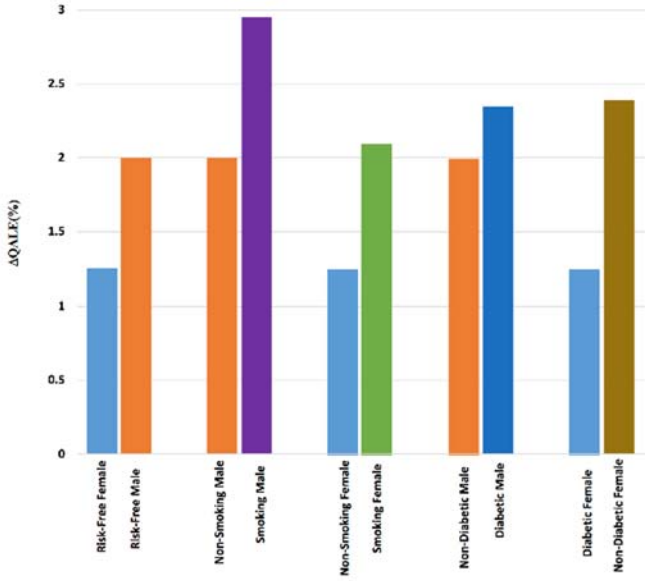

Evidently, the improvements in QALE will increase when the patient has an additional risk factor (e.g. being male versus female, being smoker for male, or being diabetic for female, as discussed above). For example, a male patient benefits more than a female patient from following our recommendations than following that of BHS, and a male patient who is smoker benefits more than when he is non-smoker. The difference in the benefit due to being smoker is less for a female patient. A female patient, however, benefits more when she is diabetic than being non-diabetic. The difference in the benefit due to being diabetic is less for a male patient. As summarized in Table 4, our results help to identify the structure of patient-centric treatment policies as well as the specific conditions under which they yield better outcomes for the patients.

Table 4: Summary of Key Results

Do the optimal medication prescription policies have a certain structure that can provide insights to the clinicians?	*Yes. The optimal policies exhibit "threshold" structure.*
Do the optimal policies depend on the patient characteristics, and how?	*Yes. The thresholds are different for different patient characteristics. As the patient characteristics get riskier (e.g. being male, smoker, or diabetic) the thresholds decrease, in a different magnitude, which means lower tolerance for high blood pressure.*
What is the value of the personalized treatment policies?	*Our personalized treatment recommendations outperform BHS guideline recommendations in terms of QALE: the riskier the patient characteristic, the higher is the benefit of following our treatment recommendations.*

7 Concluding remarks

From the theoretical perspective, the problem can be seen as a multiple stopping time problem in which we have to identify when to stop waiting for each of the medications to be prescribed depending on various patient characteristics. From the medical perspective, our problem can be seen as a preventive care model in which the risk of first serious CVD events (hence their consequence in reducing the quality of life) are minimized. We have also considered the possibility of non-CVD death associated with the patient profile, and we showed, in most cases, our policies outperform those of the BHS guideline.

References

Alagoz, O. (2004). Optimal policies for the acceptance of living-and cadaveric-donor livers. University of Pittsburgh.

Collins, R., Peto, R., MacMahon, S., Godwin, J., Qizilbash, N., Collins, R., MacMahon, S., Hebert, P., Eberlein, K.A., Taylor, J.O., et al. (1990). Blood pressure, stroke, and coronary heart disease: Part 2, short-term reductions in blood pressure: overview of randomised drug trials in their epidemiological context. The Lancet *335*, 827–838.

D'Agostino, R.B., Vasan, R.S., Pencina, M.J., Wolf, P.A., Cobain, M., Massaro, J.M., and Kannel, W.B. (2008). General cardiovascular risk profile for use in primary care: the Framingham Heart Study. Circulation *117*, 743–753.

Go, A.S., Mozaffarian, D., Roger, V.L., Benjamin, E.J., Berry, J.D., Blaha, M.J., Dai, S., Ford, E.S., Fox, C.S., Franco, S., et al. (2014). Heart disease and stroke statistics–2014 update: a report from the American Heart Association. Circulation *129*, e28–e292.

Group, S.C.R., and others (1991). Prevention of stroke by antihypertensive drug treatment in older persons with isolated systolic hypertension. JAMA J. Am. Med. Assoc. *265*, 3255–3264.

Hannerz, H., and Nielsen, M.L. (2001). Life Expectancies Among Survivors of Acute Cerebrovascular Disease. Stroke *32*, 1739–1744.

James, P.A., Oparil, S., Carter, B.L., Cushman, W.C., Dennison-Himmelfarb, C., Handler, J., Lackland, D.T., LeFevre, M.L., MacKenzie, T.D., Ogedegbe, O., et al. (2014). 2014 evidence-based guideline for the management of high blood pressure in adults: report from the panel members appointed to the Eighth Joint National Committee (JNC 8). Jama *311*, 507–520.

Law, M., Morris, J., Wald, N., and others (2009). Use of blood pressure lowering drugs in the prevention of cardiovascular disease: meta-analysis of 147 randomised trials in the context of expectations from prospective epidemiological studies. BMJ *338*.

Law, M.R., Wald, N.J., Morris, J.K., and Jordan, R.E. (2003). Value of low dose combination treatment with blood pressure lowering drugs: analysis of 354 randomised trials. BMJ *326*, 1427.

Lawrence, W.F., Fryback, D.G., Martin, P.A., Klein, R., and Klein, B.E.K. (1996). Health status and hypertension: A population-based study. J. Clin. Epidemiol. *49*, 1239–1245.

Lewington, S., Clarke, R., Qizilbash, N., Peto, R., Collins, R., and Prospective Studies Collaboration (2002). Age-specific relevance of usual blood pressure to vascular mortality: a meta-analysis of individual data for one million adults in 61 prospective studies. Lancet *360*, 1903–1913.

Lovibond, K., Jowett, S., Barton, P., Caulfield, M., Heneghan, C., Hobbs, F.R., Hodgkinson, J., Mant, J., Martin, U., Williams, B., et al. (2011). Cost-effectiveness of options for the diagnosis of high blood pressure in primary care: a modelling study. The Lancet *378*, 1219–1230.

MacMahon, S., Peto, R., Collins, R., Godwin, J., Cutler, J., Sorlie, P., Abbott, R., Neaton, J., Dyer, A., and Stamler, J. (1990). Blood pressure, stroke, and coronary heart disease: part 1, prolonged differences in blood pressure: prospective observational studies corrected for the regression dilution bias. The Lancet *335*, 765–774.

Mancia, G., De Backer, G., Dominiczak, A., Cifkova, R., Fagard, R., Germano, G., Grassi, G., Heagerty, A.M., Kjeldsen, S.E., Laurent, S., et al. (2007). 2007 Guidelines for the management of arterial hypertension The Task

Force for the Management of Arterial Hypertension of the European Society of Hypertension (ESH) and of the European Society of Cardiology (ESC). Eur. Heart J. *28*, 1462–1536.

Martin, J. (2008). Hypertension guidelines: revisiting the JNC 7 recommendations. J. Lanc. Gen. Hosp. *3*, 91–97.

Mason, J.E., England, D.A., Denton, B.T., Smith, S.A., Kurt, M., and Shah, N.D. (2012). Optimizing statin treatment decisions for diabetes patients in the presence of uncertain future adherence. Med. Decis. Making *32*, 154–166.

Moyé, L.A., and Roberts, S.D. (1982). Modeling the pharmacologic treatment of hypertension. Manag. Sci. *28*, 781–797.

National Clinical Guideline Centre (UK) (2011). Hypertension: The Clinical Management of Primary Hypertension in Adults: Update of Clinical Guidelines 18 and 34 (London: Royal College of Physicians (UK)).

Neaton, J.D., Grimm, R.H., Prineas, R.J., Stamler, J., Grandits, G.A., Elmer, P.J., Cutler, J.A., Flack, J.M., Schoenberger, J.A., McDonald, R., et al. (1993). Treatment of mild hypertension study: final results. Jama *270*, 713–724.

NICE (2006). Hypertension: management of hypertension in adults in primary care: clinical guidelines CG34 National Institute for Clinical Excellence.

NNHS (2004). National Nursing Home Survey (NNHS).

O'Sullivan, J., Derrick, G., Griggs, P., Foxall, R., Aitkin, M., and Wren, C. (1999). Ambulatory blood pressure in schoolchildren. Arch. Dis. Child. *80*, 529–532.

Puterman, M.L. (2014). Markov decision processes: discrete stochastic dynamic programming (John Wiley & Sons).

Saito, I., Kobayashi, M., Matsushita, Y., and Saruta, T. (2005). Pharmacoeconomical evaluation of combination therapy for lifetime hypertension treatment in Japan. Jpn. Med. Assoc. J. *48*, 574.

Schechter, C.B. (1988). Sequential analysis in a Bayesian model of diastolic blood pressure measurement. Med. Decis. Making *8*, 191–196.

Schechter, C.B. (1990a). Sequential Decision Making with Continuous Disease States and Measurements II. Application to Diastolic Blood Pressure. Med. Decis. Making *10*, 256–265.

Schechter, C.B. (1990b). Sequential Decision Making with Continuous Disease States and Measurements I. Theory. Med. Decis. Making *10*, 242–255.

Tengs, T.O., and Wallace, A. (2000). One thousand health-related quality-of-life estimates. Med. Care 583–637.

Timbie, J.W., Hayward, R.A., and Vijan, S. (2010). Diminishing Efficacy of Combination Therapy, Response-Heterogeneity, and Treatment Intolerance Limit the Attainability of Tight Risk Factor Control in Patients with Diabetes. Health Serv. Res. *45*, 437–456.

USDHHS (2011). JNC7:The seventh report of the joint national committee on prevention, detection, evaluation, and treatment of high blood pressur-complete report.

Whelton, P.K., He, J., Appel, L.J., Cutler, J.A., Havas, S., Kotchen, T.A., Roccella, E.J., Stout, R., Vallbona, C., Winston, M.C., et al. (2002). Primary prevention of hypertension: clinical and public health advisory from The National High Blood Pressure Education Program. JAMA *288*, 1882–1888.

WHO/ISH (2003). 2003 World Health Organization (WHO)/International Society of Hypertension (ISH) statement on management of hypertension. J. Hypertens. *21*, 1983–1992.

Wilber, J.A., and Barrow, J.G. (1972). Hypertension- a community problem. Am. J. Med. *52*, 653–663.

Wilkins, K., Campbell, N., Joffres, M.R., McAlister, F.A., Nichol, M., Quach, S., Johansen, H.L., and Tremblay, M.S. (2010). Blood pressure in Canadian adults. Health Rep. *21*, 37–46.