

EPO Dosage Optimization for Anemia Management: Stochastic Control under Uncertainty Using Conditional Value at Risk

Authors:

Jayson McAllister, Zukui Li, Jinfeng Liu, Ulrich Simonsmeier

Date Submitted: 2018-07-31

Keywords: Conditional Value at Risk, Model Predictive Control, hemoglobin level control, anemia management

Abstract:

Due to insufficient endogenous production of erythropoietin, chronic kidney disease patients with anemia are often treated by the administration of recombinant human erythropoietin (EPO). The target of the treatment is to keep the patient's hemoglobin level within a normal range. While conventional methods for guiding EPO dosing used by clinicians normally rely on a set of rules based on past experiences or retrospective studies, model predictive control (MPC) based dosage optimization is receiving attention recently. The objective of this paper is to incorporate the hemoglobin response model uncertainty into the dosage optimization decision making. Two methods utilizing Conditional Value at Risk (CVaR) are proposed for hemoglobin control in chronic kidney disease under model uncertainty. The first method includes a set-point tracking controller with the addition of CVaR constraints. The second method involves the use of CVaR directly in the cost function of the optimal control problem. The methods are compared to set-point tracking MPC and Zone-tracking MPC through computer simulations. Simulation results demonstrate the benefits of utilizing CVaR in stochastic predictive control for EPO dosage optimization.

Record Type: Published Article

Submitted To: LAPSE (Living Archive for Process Systems Engineering)

Citation (overall record, always the latest version):

LAPSE:2018.0352

Citation (this specific file, latest version):

LAPSE:2018.0352-1

Citation (this specific file, this version):



LAPSE:2018.0352-1v1

DOI of Published Version: <https://doi.org/10.3390/pr6050060>

License: Creative Commons Attribution 4.0 International (CC BY 4.0)

Article

EPO Dosage Optimization for Anemia Management: Stochastic Control under Uncertainty Using Conditional Value at Risk

Jayson McAllister ¹ , Zukui Li ^{1,*}, Jinfeng Liu ¹  and Ulrich Simonsmeier ²

¹ Department of Chemical and Materials Engineering, University of Alberta, Edmonton, AB T6G 1H9, Canada; jmcalli@ualberta.ca (J.M.); jinfeng@ualberta.ca (J.L.)

² Cybernius Medical Ltd., St. Alberta, AB T8N 2T7, Canada; ufs@cyberren.com

* Correspondence: zukui@ualberta.ca; Tel.: +1-780-492-1107

Received: 8 April 2018; Accepted: 15 May 2018; Published: 21 May 2018



Abstract: Due to insufficient endogenous production of erythropoietin, chronic kidney disease patients with anemia are often treated by the administration of recombinant human erythropoietin (EPO). The target of the treatment is to keep the patient's hemoglobin level within a normal range. While conventional methods for guiding EPO dosing used by clinicians normally rely on a set of rules based on past experiences or retrospective studies, model predictive control (MPC) based dosage optimization is receiving attention recently. The objective of this paper is to incorporate the hemoglobin response model uncertainty into the dosage optimization decision making. Two methods utilizing Conditional Value at Risk (CVaR) are proposed for hemoglobin control in chronic kidney disease under model uncertainty. The first method includes a set-point tracking controller with the addition of CVaR constraints. The second method involves the use of CVaR directly in the cost function of the optimal control problem. The methods are compared to set-point tracking MPC and Zone-tracking MPC through computer simulations. Simulation results demonstrate the benefits of utilizing CVaR in stochastic predictive control for EPO dosage optimization.

Keywords: anemia management; hemoglobin level control; model predictive control; Conditional Value at Risk

1. Introduction

One of the major side effects of chronic kidney disease (CKD) is the reduced ability to produce endogenous erythropoietin (EPO), which is a hormone that regulates the production of red blood cells in the body. When the natural production of EPO drops significantly, these patients suffer from a condition called anemia, which is characterized by a reduced mass of red blood cells and hemoglobin level. Recombinant human erythropoietin has become the standard for treating anemia in chronic kidney disease [1]. Recently, optimization of the EPO dosage has been considered to achieve reduced deviation of patient hemoglobin levels from a normal range (zone), while also reducing drug use and expense [2].

Zone model predictive controllers have become a popular area of research in the biomedical field [3]. They have been used successfully in clinical trials in the control of blood glucose in Type 1 Diabetes Mellitus [4,5]. They have gained significant popularity over set-point tracking model predictive control (MPC) for biomedical applications due to their ability to reject sensor noise, and maintain a higher degree of stability in the presence of large plant-model mismatch [6]. They also seem more appropriate for these applications as many of them do not have a target state, but rather a target zone.

The EPO dosing algorithms that have been developed and tested previously are open-loop devices. The patient will typically have their blood labs completed once every 2–4 weeks, and the EPO dose

will be optimized for the time period in between sampling times. The EPO medications are approved manually by a clinician and administered during regular dialysis. It has been recognized by the clinical community that, while low Hb level leads to anemia, too high Hb levels can increase the risk of mortality for the patient [7]. Hence, effective methods are needed to determine the appropriate dose of EPO to maintain the target Hb level.

To develop an automatic decision support system and achieve effective anemia management, some of the first published papers on the use of advanced process control in CKD utilize set-point tracking MPC combined with neural network hemoglobin response modeling and report promising results in clinical trials [8–10]. The MPC algorithm in this case was an L1-Norm of the state deviation from the target. In another case [11], a pharmacokinetic and pharmacodynamic model was used to describe the patient’s hemoglobin response and then designed controllers based on Quantitative Feedback Theory. In the small pilot study, these researchers also showed promising results. Modeling of hemoglobin response to EPO administration was also done under uncertainty via a recent method called semi-blind Robust Identification [12]. The work presented in this paper differs significantly from these researchers. One of the more challenging aspects of hemoglobin control is the time-varying nature of the patient’s system dynamics. Patient’s can become EPO resistant [13], as well as have their health improve drastically, if endogenous EPO production increases [14]. To address the time-varying nature of the hemoglobin response, constrained recursive system identification is used. There exists a great deal of uncertainty in the model parameters, and the system is often subjected to acute disturbances, such as infection [15] and blood loss [16].

The aim of this work is to develop an advanced control system to calculate optimal EPO dosages for anemia patients, while considering process uncertainty in the control method. The overall block diagram of the control system is presented in Figure 1. The method uses an autoregressive with exogenous inputs (ARX) model with a model predictive controller. To address the process uncertainty, it is represented in the form of different scenarios and further utilized within a model predictive control formulation, with the use of Conditional Value at Risk (CVaR) technique. CVaR was introduced by Rockafellar and Uryasev and is widely used in the finance industry [17]. CVaR is a popular tool used in many risk management applications. For example, it has been shown to work well in heating, ventilation and air conditioning control systems [18]; portfolio/asset optimization [19]; and quantifying flood damage [20]. However, it has been studied in very few applications related to model predictive control. The existing methods for model predictive control under uncertainty mostly include robust MPC [21] and stochastic MPC based methods [22]. The two methods address the worst-case performance and the expected performance, respectively. CVaR is a method that addresses both of these criteria. In addition, CVaR provides an effective convex approximation of a probabilistic constraint used in stochastic MPC. Thus, in this paper, the applications of CVaR for handling both risk averse performance and the probabilistic constraint are studied.

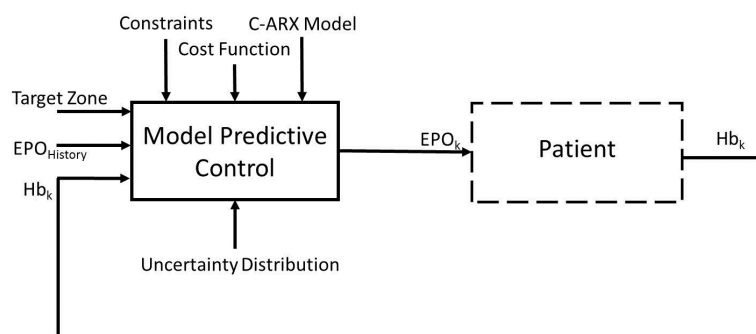


Figure 1. Overall control block diagram.

The rest of the paper is organized as follows. Section 2 presents the hemoglobin response modeling through constrained optimization. Deterministic set-point and zone model predictive control algorithms are introduced in Section 3, for comparison purposes to the CVaR methods introduced later. An overview of Conditional Value at Risk in general is introduced in Section 4, which is a technique for addressing uncertainty. A set-point tracking model predictive controller using CVaR constraints on the upper and lower zone boundaries is then introduced in Section 5 and later compared to the set-point tracking MPC. A second controller, a zone-tracking model predictive controller using CVaR directly in the cost function, is introduced in Section 6 and later compared to the zone-tracking model predictive controller. In Section 7, the controllers are tested in computer simulations. The experimental designs are outlined for a simple ARX model based simulator as well as for a physiological model based simulator, using both Gaussian and uniform uncertainty distributions. Section 8 concludes the paper.

2. Hemoglobin Response Modeling

The system model used in the following derivations has an ARX model structure. This kind of model is one type that can be estimated through classic System Identification methods. The goal of System Identification is to map the response of a system's output, to that of its input. The internal dynamics of the system are not known or estimated. This type of modeling is also referred to as black box modeling. In this work, the model parameters are obtained through constrained optimization. The model uses a weekly sampling time ($\Delta t = 1$ week), the last Hb measurement y_k , and several past EPO doses u_k (eight-week history of dosages in this work), to predict the future Hb measurement. The order of the $b(z)$ polynomial is 8 and was determined empirically [23]. The order of the $a(z)$ polynomial in the model is set to 1. Using a higher order for the $a(z)$ polynomial may lead to inaccuracies in the model estimation and predictions because of the large amount of measurement error, mixed with the infrequent measurements typically used in the medical field. It is possible that, with more accurate measurements, a larger number of a_k parameters could yield better modeling results, but this work does not address this aspect.

Iron is necessary for the creation of red blood cells, but the effect of iron levels in the patient are ignored in this model. It is assumed that, if the treating clinician were to maintain stable and adequate iron levels within the patient, the response of the hemoglobin should be mostly related to the administration of EPO, and not related to fluctuating or low values of iron within the patient. If a patient were to have low iron levels, they will likely become EPO resistant which could effect the quality of the current model, should the change (from adequate to low levels) occur quickly. Future work should try to include the iron dynamics within the model, as there certainly are some dynamic interactions between these two systems.

The ARX model structure that maps the EPO administration events to the hemoglobin concentration is

$$y_{k+1} = a_1 y_k + b_1 u_k + \dots + b_8 u_{k-7} \quad (1)$$

where the input u represents the EPO doses divided by 5000 international unit (IU), and y represents the hemoglobin concentration in g/dL. Clinical data for an example patient are shown in Figure 2. The hemoglobin measurements (top figure) are shown as the actual measurements recorded in approximately two-week intervals. The EPO doses (bottom figure) are presented as one-week dose totals. Typically, a dialysis patient with late stage renal disease will receive EPO medications 1–3 times per week, depending on dose size. With traditional dosing protocols, the oscillatory behavior observed in the hemoglobin is quite common [24].

To estimate a patient's hemoglobin response model, the cost function in Equation (2) is used. It is the weighted least squares of the error between model predictions and actual hemoglobin values.

The parameter λ is used as a weighting value ($0 < \lambda < 1$), to put a higher priority on newer measurements. The output y represents the hemoglobin measurements.

$$\min_{a_1, b_1, \dots, b_8} \sum_{i=1}^{t_f} \lambda^{(t_f-i)} (\hat{y}_{k+i} - y_{k+i})^2 \tag{2}$$

where \hat{y}_{k+i} and y_{k+i} represent the predicted and actual hemoglobin value, respectively. Equation (3) represents the predicted hemoglobin using the historical data and the estimated model parameters (b_1, \dots, b_8, a_1) for a time horizon of 1 to t_f .

$$\begin{bmatrix} \hat{y}_{k+1} \\ \hat{y}_{k+2} \\ \vdots \\ \hat{y}_{k+t_f} \end{bmatrix} = \begin{bmatrix} u_k & \dots & u_{k-7} & y_k \\ u_{k+1} & \dots & u_{k-6} & y_{k+1} \\ \vdots & & & \vdots \\ u_{k+t_f-1} & \dots & u_{k+t_f-8} & y_{k+t_f-1} \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_8 \\ a_1 \end{bmatrix} \tag{3}$$

which can be compactly written as

$$\hat{Y} = X\theta \tag{4}$$

where θ is a vector of the optimized model parameters, Y is the measured hemoglobin, \hat{Y} is the predicted hemoglobin, and X is a matrix of the EPO dosage and measured hemoglobin concentration. Then, the objective function in Equation (2) can be written as

$$\min \theta^T H \theta + 2 \theta^T f + c$$

where $H = X^T R X$, $f = -X^T R Y$, and $c = Y^T R Y$. The measurement weight in Equation (2) is represented by the weighting matrix R , which is a diagonal matrix of the weighting values λ for each measurement.




Figure 2. Clinical data for an anemia patient.

Recursive modeling has been shown to work well in cases where the model parameters are time varying, such as in Type 1 Diabetes Mellitus [25]. Recursive modeling is the process of estimating a new patient model after each updated measurement. In this work, the patient's model parameters are relearned after each measurement, based on a moving window of estimation data. In the absence of validation data, constraints are used in the modeling formulation to ensure model stability, as well as to enforce that the model parameters follow patterns that are well described by many patients [26]. The complete constrained ARX (C-ARX) model parameter estimation problem can be represented by Equation (5). Equation (5) assumes that the delay of the system is approximately two weeks. The peak time b_k parameter is shown here with a delay of two weeks (Equation (5f)), and represents the largest contribution to the system of all the b_k parameters. The peak time b_k parameter can be shifted to the third or fourth week by reformulating constraints (Equation (5f,g)). Individual patient hemoglobin response times can vary. Some patients may begin to observe a substantial increase of hemoglobin in as little as two weeks, where some take much longer before a significant effect on the hemoglobin can be observed. The most appropriate way of dealing with the system's unknown delay is to solve the same problem where the peak time b_k parameter is at the second, third and fourth delay, and then choose the solution with the lowest cost function. Equation (5b) is used as a minimum and maximum range for the model estimates, as almost all the measurements should be within this range. Measurements that exist outside of this range may not be realistic, or may coincide with some abnormal event. Equation (5c) is used to enforce model stability. Equation (5d) is used to enforce a minimum response delay of two weeks. Equation (5e) enforces that the model parameters are always positive, because the response of the system to the inputs is always positive. Equations (5f)–(5i) are defined through experience, and are used to enforce a particular parameter shape learned in [26]. Finally, Equation (5i) is used to ensure the model relies somewhat on the inputs to the system, and was determined through trial and error.

$$\min \quad \theta^T H \theta + 2 \theta^T f + c \quad (5a)$$

$$s.t. \quad 7.0 \leq \hat{y}_{k+i} \leq 15.0 \quad i = 1, \dots, t_f \quad (5b)$$

$$0.7 \leq a_1 \leq 0.99 \quad (5c)$$

$$b_1 = 0 \quad (5d)$$

$$b_k \geq 0 \quad k = 2, \dots, 8 \quad (5e)$$

$$b_2 \geq 0.05 \quad (5f)$$

$$b_k < 0.8 b_{k-1} \quad k = 3, \dots, 8 \quad (5g)$$

$$b_8 < 0.005 \quad (5h)$$

$$\sum_{k=1}^8 b_k \geq 0.1 \quad (5i)$$

The above model is a constrained quadratic optimization problem and it is solved using the *quadprog* solver in MATLAB.

3. Deterministic Control Formulation: Setpoint MPC and Zone-MPC

The stochastic control model described in later sections is compared to the traditional method of implementing model predictive control (MPC), including a set-point tracking MPC and a zone-tracking MPC. First, the set-point tracking model is presented in Equation (6). The target hemoglobin level is represented by y_{sp} . Q is the state tuning parameter. N is the prediction horizon.

$$\min_u \quad \frac{1}{2} \sum_{i=1}^N Q \|\hat{y}_{k+i} - y_{sp}\| + \frac{1}{2} \sum_{i=0}^{N-2} \|\Delta u_{k+i}\| \quad (6a)$$

$$s.t. \quad u_L \leq u_{k+i} \leq u_H \quad i = 0, \dots, N-2 \quad (6b)$$

$$\Delta u_L \leq \Delta u_{k+i} \leq \Delta u_H \quad i = 0, \dots, N-2 \quad (6c)$$

Figure 3a is an illustration of set-point tracking MPC, where the output is driven to the set-point. Control input chatter is often observed in set-point tracking MPC.




Figure 3. Comparison of setpoint and zone tracking MPC.

The zone-tracking MPC formulation is presented in Equation (7). Zone control is realized through the constraints set on δ_{k+i} in Equation (7b). Set-point tracking MPC is a special case of zone control, where the zone limits y_H and y_L are equal.

$$\min_{u, \delta} \quad \frac{1}{2} \sum_{i=1}^N Q \|\hat{y}_{k+i} - \delta_{k+i}\| + \frac{1}{2} \sum_{i=0}^{N-2} \|\Delta u_{k+i}\| \quad (7a)$$

$$s.t. \quad y_L \leq \delta_{k+i} \leq y_H \quad i = 1, \dots, N \quad (7b)$$

$$u_L \leq u_{k+i} \leq u_H \quad i = 0, \dots, N-2 \quad (7c)$$

$$\Delta u_L \leq \Delta u_{k+i} \leq \Delta u_H \quad i = 0, \dots, N-2 \quad (7d)$$

In zone-MPC, it is common to see the output oscillate slightly between the upper and lower limits while the system stabilizes to a steady state within the control zone limits. Figure 3b is an illustration of zone-MPC.

It is worth pointing out that the cost function can be modeled using either absolute value or a quadratic function for both controllers presented above.

4. Conditional Value at Risk

The uncertainty considered in this work is in the form of process uncertainty on the output and can be represented by the variable w_k in Equation (8).

$$\hat{y}_{k+1} = a_1 y_k + b_1 u_k + \dots + b_8 u_{k-7} + w_k \quad (8)$$

To address the uncertainty in the hemoglobin response to EPO dosage, this work studies two Conditional Value at Risk (CVaR) techniques for control under uncertainty. In general, Value at Risk (VaR) at confidence level β (VaR_β) is defined as the maximum loss value that is assigned to some desired probability, β . That is to mean that a loss less than or equal to VaR_β will occur $\beta \times 100\%$ of the time. What VaR_β fail to account for, is the losses that occur $(1 - \beta) \times 100\%$ of the time with a loss larger than VaR_β . CVaR at confidence level β ($CVaR_\beta$) aims to minimize this loss and it is defined as the average of the loss exceeding VaR_β . For instance, a slightly higher VaR_β may be tolerated, if the average value of the β -tail distribution is lower. This would ensure that when a loss greater than VaR_β does occur, it will likely be smaller. A diagram is included in Figure 4 to facilitate a better understanding of the value in which CVaR seeks to minimize.

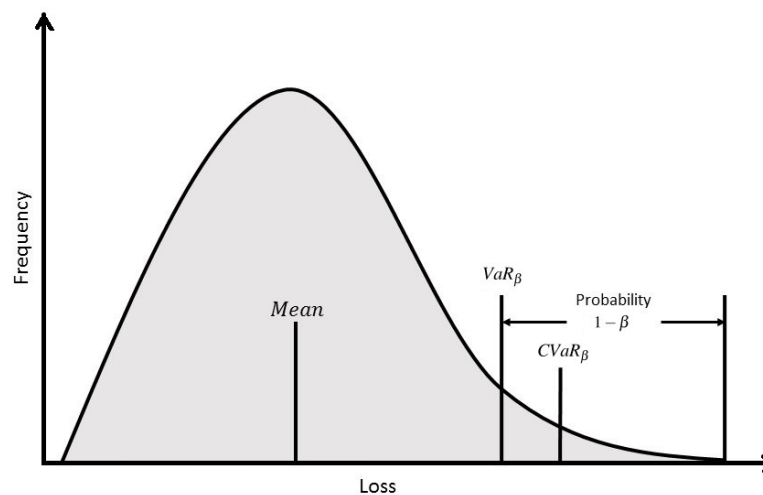


Figure 4. Depiction of the definition of Conditional Value at Risk.

Another major advantage of using CVaR constraints over other types of chance constraints, is that the uncertainty can follow any distribution. With a known distribution, different scenarios are generated for each sampling instant by sampling the random distribution. In this fashion, the controller optimization problem can take into account multiple scenarios, which help to approximate the different possibilities of measurement and process uncertainties, to ultimately provide a more robust solution.

Next, the CVaR technique is incorporated into the MPC formulation in two different ways. The first method handles chance constraints through CVaR approximation, whereas the second method directly optimizes CVaR in the cost function. Both controllers were tested in computer simulations against traditional zone-MPC and set-point tracking MPC to show the benefits that the two methods may provide.

5. Stochastic Control Using CVaR Constraints

The uncertainty is located within the output, \hat{y}_k , in the form of a random variable w_k . w_k is assumed to follow a certain distribution. The controller is designed to take into account a process uncertainty from a known distribution. The optimization problem uses two chance constraints (Equations (9b) and (9c)) on the output of the system, one on the lower bound and the other one on the upper bound. Due to the conservative approximation of the chance constraint that the CVaR constraint results in, they will only be used for $k = 3$, as shown below. If more constraints are used, the problem will become infeasible. $k = 2$ is not used because $b_1 = 0$, which means that, if the hemoglobin were to fall outside of the zone, the controller would only have a single input to try and move the hemoglobin into the zone to avoid infeasibility. The result would be a very aggressive control move in these scenarios that should be avoided. It should also be noted that the cost function is computed with Equation (1), where \hat{y}_k represents the predicted hemoglobin without uncertainty. Chance constraints are used in solving optimization problems under uncertainty. There exists uncertainty in the output prediction \tilde{y}_k , which is computed through Equation (8). The meaning of the chance constraint is that utilizing some knowledge of the uncertainty in \tilde{y}_k , the constraint will be infeasible no more than ϵ amount of the time. ϵ is the probability of violation. The chance constrained set-point MPC formulation is proposed as following:

$$\min \quad Q \sum_{k=1}^N |\hat{y}_k - y_{sp}| + \sum_{k=0}^{N-2} |\Delta u_k| \quad (9a)$$

$$s.t. \quad Pr\{\tilde{y}_k > y_H\} \leq \epsilon \quad k = 3 \quad (9b)$$

$$Pr\{\tilde{y}_k < y_L\} \leq \epsilon \quad k = 3 \quad (9c)$$

$$u_L \leq u_{k+i} \leq u_H \quad i = 0, \dots, N-2 \quad (9d)$$

$$\Delta u_L \leq \Delta u_{k+i} \leq \Delta u_H \quad i = 0, \dots, N-2 \quad (9e)$$

Due to the conservatism of the CVaR constraints, only a single point along the prediction horizon will utilize the constraints. Note that the Δu_k term is summed from 0 to $N-2$ because b_1 is always estimated as zero in the ARX model. The following derivations use a prediction horizon, N , of eight weeks, and the ARX model introduced previously. Equation (10) is used to demonstrate the derivation, which is the upper bound constraint [17].

$$Pr\{\tilde{y}_k > y_H\} \leq \epsilon \quad (10)$$

In the derivation below, the indicator function $1_{(0,\infty)}(u)$ is used. This function holds a value of 1 if $u > 0$ and a value of 0 if $u \leq 0$. For any positive parameter α , we have $1_{(0,\infty)}(u) = 1_{(0,\infty)}(\frac{1}{\alpha}u)$. Defining an upper bounding function $\phi(u)$ for the indicator function, the following inequality can be written.

$$1_{(0,\infty)}(\frac{1}{\alpha}u) \leq \phi(\frac{1}{\alpha}u) \quad (11)$$

Replacing u with $\tilde{y}_k - y_H$, the following inequality can be written.

$$Pr\{\tilde{y}_k - y_H > 0\} = E[1_{(0,\infty)}(\tilde{y}_k - y_H)] \leq E[\phi(\frac{1}{\alpha}(\tilde{y}_k - y_H))] \leq \epsilon \quad (12)$$

If $\tilde{y}_k - y_H > 0$, the constraint is violated and the indicator function holds a value of 1, and vice versa. $\phi(u)$ is the upper bounding function for the indicator function and it is a conservative estimate of the probability of violation, ϵ . From this, any solution that satisfies the inequality will also satisfy the constraint. The focus remains on the right side of this inequality. $(u)^+$ is defined as the maximum operator, which holds a maximum value between 0 and the input u . Applying the upper bounding function $\phi(u) = (u + 1)^+$, a conservative approximation of the chance constraint is defined as

$$E\left[\left(\frac{1}{\alpha}(\tilde{y}_k - y_H) + 1\right)^+\right] \leq \epsilon \quad (13)$$

multiply by $\frac{\alpha}{\epsilon}$

$$\frac{1}{\epsilon} E\left[\left((\tilde{y}_k - y_H) + \alpha\right)^+\right] \leq \alpha \quad (14)$$

moving α to the left hand side

$$-\alpha + \frac{1}{\epsilon} E\left[\left((\tilde{y}_k - y_H) + \alpha\right)^+\right] \leq 0 \quad (15)$$

To reduce conservatism, it is necessary to find the minimum α that satisfies the inequality. That is, it is desired to find the smallest upperbound. This is a CVaR constraint, but it is difficult to calculate its value in its present form.

$$\min_{\alpha} -\alpha + \frac{1}{\epsilon} E\left[\left((\tilde{y}_k - y_H) + \alpha\right)^+\right] \leq 0 \quad (16)$$

The expectation operator is removed with the introduction of sampling to approximate the CVaR constraint. M is the total number of scenarios. For instance, if there were a single random variable per scenario, and the prediction horizon was 1, it would be necessary to sample the random variables' distribution M times. π_j is the probability that a single scenario will occur, and is equal to $\frac{1}{M}$. Note that \tilde{y} changes to $y_{k,j}$ because the output is now an expected value with the uncertainty added in.

$$-\alpha + \frac{1}{\epsilon} \sum_{j=1}^M \pi_j \left[\left((y_{k,j} - y_H) + \alpha\right)^+\right] \leq 0 \quad (17)$$

The expression inside of the max operator can be replaced by two separate constraints. Equation (18) replaces the max operator term with the variable $v_{j,k}$, and constraints are set on $v_{j,k}$ to satisfy the original max operator's functionality. The combination of all three constraints in Equation (18) can be used to provide a conservative estimate of the original chance constraint in Equation (10) through the use of sampling [18]. These constraints are solved using many different scenarios of the random variable distribution, leading to it being part of a class of controllers called Scenario-Based MPC.

$$-\alpha_k + \frac{1}{\epsilon} \sum_{j=1}^M \pi_j v_{j,k} \leq 0 \quad k = 3 \quad (18a)$$

$$v_{j,k} \geq (y_{k,j} - y_H) + \alpha_k \quad j = 1, \dots, M, \quad k = 3 \quad (18b)$$

$$v_{j,k} \geq 0 \quad j = 1, \dots, M, \quad k = 3 \quad (18c)$$

Similarly, for the lower bound constraints, we have the following approximation

$$-\gamma_k + \frac{1}{\epsilon} \sum_{j=1}^M \pi_j \omega_{j,k} \leq 0 \quad k = 3 \quad (19a)$$

$$\omega_{j,k} \geq (y_L - y_{k,j}) + \gamma_k \quad j = 1, \dots, M, \quad k = 3 \quad (19b)$$

$$\omega_{j,k} \geq 0 \quad j = 1, \dots, M, \quad k = 3 \quad (19c)$$

Finally, notice that the cost function in Equation (9) contains absolute terms. This equation can be manipulated to remove the absolute sign and the resulting linear program with all the constraints is outlined in Equation (20).

$$\min \quad Q \sum_{k=1}^N h_k + \sum_{k=0}^{N-2} g_k \quad (20a)$$

$$\text{s.t.} \quad h_k \geq \hat{y}_k - y_{sp} \quad k = 1, \dots, N \quad (20b)$$

$$h_k \geq -(\hat{y}_k - y_{sp}) \quad k = 1, \dots, N \quad (20c)$$

$$g_k \geq \Delta u_k \quad k = 0, \dots, N-2 \quad (20d)$$

$$g_k \geq -\Delta u_k \quad k = 0, \dots, N-2 \quad (20e)$$

where \hat{y}_k is the Hb prediction from the ARX model without process uncertainty term.

Summarizing the above derivations, the complete deterministic approximation to the chance constrained set-point MPC formulation is given as follows:

$$\min \quad Q \sum_{k=1}^8 h_k + \sum_{k=0}^6 g_k \quad (21a)$$

$$\text{s.t.} \quad -\alpha_k + \frac{1}{\epsilon} \sum_{j=1}^M \pi_j v_{j,k} \leq 0 \quad k = 3 \quad (21b)$$

$$v_{j,k} \geq (y_{k,j} - y_H) + \alpha_k \quad j = 1, \dots, M, \quad k = 3 \quad (21c)$$

$$v_{j,k} \geq 0 \quad j = 1, \dots, M, \quad k = 3 \quad (21d)$$

$$-\gamma_k + \frac{1}{\epsilon} \sum_{j=1}^M \pi_j \omega_{j,k} \leq 0 \quad k = 3 \quad (21e)$$

$$\omega_{j,k} \geq (y_L - y_{k,j}) + \gamma_k \quad j = 1, \dots, M, \quad k = 3 \quad (21f)$$

$$\omega_{j,k} \geq 0 \quad j = 1, \dots, M, \quad k = 3 \quad (21g)$$

$$h_k \geq \hat{y}_k - y_{sp} \quad k = 1, \dots, N \quad (21h)$$

$$h_k \geq -(\hat{y}_k - y_{sp}) \quad k = 1, \dots, N \quad (21i)$$

$$g_k \geq \Delta u_k \quad k = 0, \dots, N - 2 \quad (21j)$$

$$g_k \geq -\Delta u_k \quad k = 0, \dots, N - 2 \quad (21k)$$

$$u_L \leq u_k \leq u_H \quad k = 0, \dots, N - 2 \quad (21l)$$

$$\Delta u_L \leq \Delta u_k \leq \Delta u_H \quad k = 0, \dots, N - 2 \quad (21m)$$

There are $4M + 76$ constraints and the optimization vector consists of $2M + 34$ variables. The controller optimization problem involves the use of M scenarios, where M typically holds a value larger than 100 and is selected by the user. A single scenario is simply the trajectory of the hemoglobin response while including uncertainty values drawn randomly from the known uncertainty distribution. It should quickly become apparent that scenario-based constraints such as these lead to a very large number of constraints, with $4M + 2$ constraints related to the CVaR constraints alone. With this in mind, the cost function used in this design is a linear program and MATLAB's *linprog* function was used along with its dual-simplex optimization algorithm. This controller is referred to as CVaR₁ below.

6. Stochastic Control Using a CVaR Cost Function

The second method discussed in this work optimizes the control inputs using conditional value at risk directly within the cost function. In contrast to the first method, the solution to this optimization problem is always feasible. The controller derivation follows the work of [17,19] closely, but with the exception that the algorithm is modified to use a zone-MPC formulation. The optimal solution can be thought of as the solution that minimizes the average loss that occurs in the β -tail distribution of the probability density function (pdf) of the loss function. The controller formulation in Equation (22) uses the distribution of the random variable, w_k , to generate several different plausible scenarios of process uncertainty.

$$\min_{u, \delta} \quad CVaR_{\beta} f(u, \delta, w) \quad (22a)$$

$$s.t. \quad y_L \leq \delta_{k+i} \leq y_H \quad i = 1, \dots, N \quad (22b)$$

$$u_L \leq u_{k+i} \leq u_H \quad i = 0, \dots, N - 2 \quad (22c)$$

$$\Delta u_L \leq \Delta u_{k+i} \leq \Delta u_H \quad i = 0, \dots, N - 2 \quad (22d)$$

where the cost function is given as

$$f(u, \delta, w) = Q \sum_{i=1}^N |\hat{y}_{k+i} - \delta_{k+i}| + \sum_{i=0}^{N-2} |\Delta u_{k+i}| \quad (23)$$

The loss function is associated with decision variables u_k and auxiliary variables δ_k and random variables w_k . Note that Q is the tuning parameter in this equation and zone control is facilitated through the use of the variable δ_k . The conditional expectation of loss in the bad tail leads to the idea of Conditional Value at Risk. $CVaR_{\beta}$ is defined as the conditional expectation of the loss above ℓ_{β}

$$CVaR_{\beta} f = \phi_{\beta} = \frac{1}{1 - \beta} \int_{f \geq \ell_{\beta}} f(u, \delta, w) p(w) dw \quad (24)$$

Rockafellar and Uryasev [17] showed that the $CVaR_{\beta}$ can be determined as follows

$$\phi_{\beta} = \min_{\ell} \ell + \frac{1}{1 - \beta} E [f(u, \delta, w) - \ell]^+ = \min_{\ell} \ell + \frac{1}{1 - \beta} \int_w [f(u, \delta, w) - \ell]^+ p(w) dw \quad (25)$$

$[f]^+$ denotes the max of $[f, 0]$. ℓ is an auxiliary variable (can be viewed as VAR_β) to be optimized. The integral can be approximated by sampling of the random variable: w_1, \dots, w_M with probability π_1, \dots, π_M

$$\phi_\beta = \min_{\ell} \ell + \frac{1}{1-\beta} \sum_{j=1}^M \pi_j [f(u, \delta, w_j) - \ell]^+ \quad (26)$$

The absolute terms from Equation (23) can be linearized by introducing variables h_k and g_k . Note that $y_{k+1+i,j}$ represents the expected output for the scenario.

$$\min_{u_k, \delta_k, \ell, v_j} \quad Q \sum_{i=1}^N h_{k+i,j} + \sum_{i=0}^{N-2} g_{k+i} \quad (27a)$$

$$s.t. \quad h_{k+i,j} \geq y_{k+1+i,j} - \delta_{k+i,j} \quad i = 1, \dots, N, j = 1, \dots, M \quad (27b)$$

$$h_{k+i,j} \geq -(y_{k+1+i,j} - \delta_{k+i,j}) \quad i = 1, \dots, N, j = 1, \dots, M \quad (27c)$$

$$g_{k+i} \geq \Delta u_{k+i} \quad i = 0, \dots, N-2 \quad (27d)$$

$$g_{k+i} \geq -\Delta u_{k+i} \quad i = 0, \dots, N-2 \quad (27e)$$

Then, by combining Equations (26) and (27), the CVaR cost can be written as

$$\min_{u_k, \delta_k, \ell} \ell + (1-\beta)^{-1} \sum_{j=1}^M \pi_j \left[Q \sum_{i=1}^N h_{k+i,j} + \sum_{i=0}^{N-2} g_{k+i} - \ell \right]^+ \quad (28)$$

This is equivalent to

$$\min_{u_k, \delta_k, \ell, v_j} \quad \ell + (1-\beta)^{-1} \sum_{j=1}^M \pi_j v_j \quad (29a)$$

$$s.t. \quad v_j \geq Q \sum_{i=0}^{N-1} h_{k+i,j} + \sum_{i=0}^{N-2} g_{k+i} - \ell \quad j = 1, \dots, M \quad (29b)$$

$$v_j \geq 0 \quad j = 1, \dots, M \quad (29c)$$

The variable v_j is introduced to allow the calculation of the max operator. The final form of the stochastic MPC problem is a linear program outlined in Equation (30) that minimizes the $CVaR_\beta$, where k, j, i represent the sampling instant, scenario number and sampling instant along the prediction horizon, respectively.

$$\min_{u_k, \delta_k, \ell, v_j} \quad \ell + (1-p)^{-1} \sum_{j=1}^M \pi_j v_j \quad (30a)$$

$$s.t. \quad v_j \geq Q \sum_{i=0}^{N-1} h_{k+i,j} + \sum_{i=0}^{N-2} g_{k+i} - \ell \quad j = 1, \dots, M \quad (30b)$$

$$v_j \geq 0 \quad j = 1, \dots, M \quad (30c)$$

$$h_{k+i,j} \geq y_{k+1+i,j} - \delta_{k+i,j} \quad i = 1, \dots, N, j = 1, \dots, M \quad (30d)$$

$$h_{k+i,j} \geq -(y_{k+1+i,j} - \delta_{k+i,j}) \quad i = 1, \dots, N, j = 1, \dots, M \quad (30e)$$

$$g_{k+i} \geq \Delta u_{k+i} \quad i = 0, \dots, N-2 \quad (30f)$$

$$g_{k+i} \geq -\Delta u_{k+i} \quad i = 0, \dots, N-2 \quad (30g)$$

$$y_L \leq \delta_{k+i,j} \leq y_H \quad i = 1, \dots, N, j = 1, \dots, M \quad (30h)$$

$$u_L \leq u_{k+i} \leq u_H \quad i = 0, \dots, N-2 \quad (30i)$$

$$\Delta u_L \leq \Delta u_{k+i} \leq \Delta u_H \quad i = 0, \dots, N-2 \quad (30j)$$

The above Equation results in $4MN + 2M + 6N$ constraints and an optimization vector length of $2MN + 2N + M + 1$. This controller is referred to as CVaR₂ below.

7. Computer Simulation Results

The following section outlines the experimental design and the simulation results comparing the CVaR methods to set-point tracking MPC and zone-MPC for both Gaussian and uniform distributions, for both an ARX model based patient simulator and a pharmacokinetics and pharmacodynamics (PK/PD) model based simulator.

Both simulator cases use the same Gaussian and uniform distributions for the random variable w_k . The histograms are outlined in Figure 5. Figure 5a follows a Gaussian distribution of $N(0, 0.2^2)$, while Figure 5b is a uniform distribution of $-0.4 \leq w_k \leq 0.4$. In all cases, the controllers that use the CVaR technique were given the exact distribution of w_k to calculate the scenarios for each controller iteration. For these simulations, each scenario would draw random variables for each sampling time in the future that is used in the controller design. For instance, in the first CVaR controller, the uncertainty is only needed up to $k + 3$ for each sampling instant, meaning there are only three values drawn from the uncertainty distribution for each scenario. In the second CVaR controller, it draws the full eight weeks of values for each scenario.




Figure 5. Uncertainty distributions.

The performance of the controllers are measured based on four metrics. The state performance metrics are the integrated output error (IOE) outside of the control zone, and the percent of points in the zone (PIZ). The input metrics are average weekly EPO dose (EPO/week) and the average weekly change in dose (avg ΔEPO). The four performance metrics can be calculated using Equation (31).

$$IOE = \sum_{k=1}^{tf} f_k, \quad f_k = \begin{cases} Hgb_k - y_H, & \text{if } Hgb_k > y_H \\ y_L - Hgb_k, & \text{if } Hgb_k < y_L \\ 0, & \text{if } y_L \leq Hgb_k \leq y_H \end{cases} \quad (31a)$$

$$PIZ = 100 \frac{\sum_{k=1}^{tf} c_k}{tf}, \quad c_k = \begin{cases} 1, & \text{if } y_L \leq Hgb_k \leq y_H \\ 0, & \text{if } y_L > Hgb_k \text{ or } Hgb_k > y_H \end{cases} \quad (31b)$$

$$EPO/week = \frac{\sum_{k=1}^{tf} EPO_k}{tf} \quad (31c)$$

$$Avg \Delta EPO = \frac{\sum_{k=2}^{tf} |EPO_k - EPO_{k-1}|}{tf} \quad (31d)$$

7.1. Test under An ARX Model Based Simulator

The ARX model based simulations were performed without time-varying parameters and recursive modeling. These simulations represent a nominal case, where the patient model does not change. A single constrained-ARX model was used. A simulation time of 2000 weeks was used, to obtain better statistical power of the results. A sampling time of one week was used. The controller tuning parameters are shown in Table 1. The tuning parameters were defined empirically, to limit the Average Change in EPO parameters to be low. In practice, EPO is often administered in 1000 IU increments, so the controllers were tuned to have the average weekly change less than this value.

Table 1. Controller settings.

Controller Identifier	MPC	Zone-MPC	CVaR ₁	CVaR ₂
Q	0.8	10	0.8	0.3
M	-	-	500	500
Δu_H (IU)	20,000	20,000	20,000	20,000
β	-	-	-	0.95
ϵ	-	-	0.1	-
y_L	10.5	10	10.5	10
y_H	10.5	11	10.5	11
Target	10.5	10–11	10.5	10–11
Constraint Limits	-	-	9.5–11.5	-

The results for the Gaussian distribution simulation are outlined in Table 2. The CVaR formulations offer significant improvements over traditional zone-MPC, but their performance is only marginally better than set-point tracking MPC. The time to solve is included in the table. Due to the large number of added variables introduced for sampling in the CVaR controllers, the time per iteration also increases.

Table 2. Simulation results under Gaussian distribution.

Performance Statistic	MPC	Zone-MPC	CVaR ₁	CVaR ₂
IOE ($\frac{g}{dL \text{ week}}$)	112.1	174.3	113.2	108.0
PIZ (%)	74.2	64.8	74.0	74.0
EPO/week (IU)	5449	5450	5450	5448
Avg ΔEPO	603	489	938	541
Time per Iteration (sec)	0.005	0.005	0.091	0.588

The results for the uniform distribution are outlined in Tables 3 and 4. CVaR₁ has a large increase in the state performance over set-point tracking MPC, but it comes at the cost of an aggressive control action. The overall statistics show the controller to be more aggressive, but in fact it is actually quite stable with regards to its control input action.



Table 3. Setpoint MPC and CVaR₁ model performance under uniform distribution.

Performance Statistic	MPC	CVaR ₁
IOE ($\frac{g}{dL \text{ week}}$)	146.7	115.5
PIZ (%)	69.9	73.4
EPO/week (IU)	5435	5438
Avg ΔEPO	701	1413

Table 4. Zone MPC and CVaR₂ model performance under uniform distribution.

Performance Statistic	Zone-MPC	CVaR ₂
IOE ($\frac{g}{dL \cdot week}$)	189.2	140.8
PIZ (%)	65.3	71.0
EPO/week (IU)	5433	5435
Avg ΔEPO	592	698

Figure 6 depicts a small window of the full 2000 week simulation results comparing MPC to CVaR₁. CVaR₁ behaves similar to zone-MPC, but the controller will be extra aggressive outside of the zone boundary as the hemoglobin approaches the constraint boundaries. The overall aggressiveness of the controller is much larger due to these larger moves, even though the doses are typically more stable. Figure 7 compares CVaR₂ to zone-MPC. As compared to zone-MPC, the response is much better, as the CVaR is able to pre-emptively change the dose in anticipation of the hemoglobin leaving the zone, whereas zone-MPC only reacts to the hemoglobin leaving the zone, causing it to traverse beyond the zone boundaries more often.

**Figure 6.** Comparisons to CVaR₁ for the uniform distribution.**Figure 7.** Comparisons to CVaR₂ for the uniform distribution.

7.2. Test under PK/PD Model Based Simulator

For this test, the patient simulator shown in Figure 8 was used which was designed in detail in [23]. The patient simulator represents a more realistic scenario to test the controllers with. The patient simulator uses a system of nonlinear delayed differential equations (DDE) based on pharmacokinetics and pharmacodynamics to model hemoglobin response.

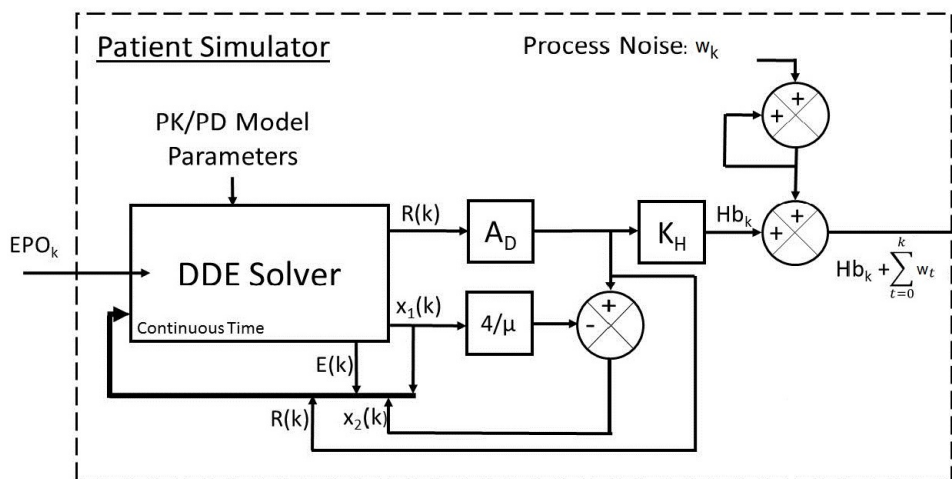


Figure 8. Block diagram of the patient simulator.

The DDE model was proposed by Chait et al. [11]. Pharmacokinetics is the study of the movement of drugs within the body while pharmacodynamics is related to the mechanisms by which the drugs affect the body. The PK/PD model is described below.

$$E_{en} = \frac{CH_{en}}{\mu K_H S - H_{en}} \tag{32a}$$

$$\frac{dE(t)}{dt} = \frac{-V E(t)}{K_m + E(t)} - \alpha E(t) + dose(t) \tag{32b}$$

$$\frac{dR(t)}{dt} = \frac{S (E_{en} + E(t - D))}{(C + E_{en} + E(t - D))} - 4 \frac{x_1(t)}{\mu^2} \tag{32c}$$

$$\frac{dx_1(t)}{dt} = x_2(t) \tag{32d}$$

$$\frac{dx_2(t)}{dt} = \frac{S (E_{en} + E(t - D))}{(C + E_{en} + E(t - D))} - 4 \frac{x_1(t)}{\mu^2} - 4 \frac{x_2(t)}{\mu} \tag{32e}$$

In Equation (32), states $E(t)$ and $R(t)$ represent the pool of exogenous erythropoietin and the population of red blood cells (RBC) within the body, respectively. States $x_1(t)$ and $x_2(t)$ are internal states that aid in calculating $R(t)$. E_{en} is the endogenous erythropoietin naturally produced by the body. K_H is the average amount of hemoglobin per RBC (also known as the mean corpuscular hemoglobin, MCH). The value used here is fixed at 29.5 pg/cell, which is within the reference range of 27–33 pg/cell [11]. The hemoglobin value is directly proportionate to the RBC population; the hemoglobin value can be attained by multiplying the RBC population estimate by the MCH value. The function $dose(t)$ is a train of impulses, representing the EPO injections. The model has an initial condition as represented in Equation (33) and requires two measurements of hemoglobin

(Hb_1 and Hb_2) and the time in between those measurements (t_1 and t_2) to estimate. The prior history of the exogenous EPO is unknown, and assumed to be zero for all $t \leq 0$.

$$\dot{R}_0 = \frac{(Hb_2 - Hb_1)}{K_H(t_2 - t_1)} \quad (33a)$$

$$E(0) = 0 \quad (33b)$$

$$R(0) = \frac{Hb_1}{K_H} \quad (33c)$$

$$x_1(0) = \frac{\mu(H_{en} - \mu K_H R_0)}{4K_H} \quad (33d)$$

$$x_2(0) = R(0) - \frac{4 x_1(0)}{\mu} \quad (33e)$$

Table 5 contains parameters that are estimated for an individual patient using clinical data and nonlinear least-squares regression. The *dde23* and *lsqnonlin* functions in MATLAB were used for DDE solution and parameter estimation, respectively.

Table 5. Parameters for the PK/PD Model.

Parameter	Description
H_{en}	Hemoglobin Level due to Endogenous Erythropoietin
μ	Mean RBC life span
V	Maximal Exogenous Erythropoietin clearance rate
K_m	Exogenous Erythropoietin level that produces half maximal clearance rate
α	Linear clearance constant
S	Maximal RBC production rate stimulated by E_p
C	Amount of EPO that produces half maximal RBC production rate
D	Time required for EPO-stimulated RBCs to start forming

The DDE model is solved in continuous time and sampled every two weeks. The output of the DDE model is the red blood cell population (R_k), which is multiplied by the mean corpuscular hemoglobin (K_H) to get the hemoglobin measurement. The hemoglobin measurement is then added to an integrating process uncertainty (w_k). It is important to note the difference between the ARX model simulator and PK/PD model simulator in regards to the way w_k is used. In the ARX simulator case, w_k is filtered with the patient model. In the PK/PD simulator case, the disturbance is integrating. The magnitudes are drawn from the same distributions in both cases.

The simulator also includes the possibility of acute step and ramp disturbances to simulate infections and blood losses frequently observed in actual patient data. Acute disturbances are facilitated through the use of the variable A_D , which holds a value of one unless a disturbance has occurred. If an acute disturbance occurs, it will hold a value less than 1, and the existing pool of red blood cells in the patient simulator model is multiplied by the fractional value to reduce it. A step disturbance is then modeled by an impulse in A_D , while an infection is modeled by a step function in A_D . Plant-model mismatch exists as the controller model is linear. The patient simulator was shown to mimic the real life hemoglobin system dynamics of patients well [23]. The total simulation time was 500 weeks. Recursive modeling is necessary to capture the time-varying nature of the system model. Although the patient model does not change significantly from sampling interval to sampling interval, the patient model is still re-estimated to provide continuity between the models at each sampling interval. A hemoglobin sampling time of two weeks is used, and two weekly dose total inputs solved by the controller are used during each sampling interval. The recursive modeling method applies a data re-sampling algorithm to get one week measurements (through linear interpolation), and weekly dose totals. The one-week re-sampled hemoglobin history is also smoothed, using a five-point moving average filter, before modeling occurs. The weighting parameter λ is set to 0.965. The tuning parameters for

each of the controllers are outlined in Table 6. The tuning parameters were again defined empirically. For these simulations, the Avg ΔEPO was tuned to be approximately equal in most cases to better compare the controller results to one other.

Table 6. Controller settings.

Controller Identifier	MPC	Zone-MPC	CVaR ₁	CVaR ₂
Q	3	10	3	1.5
M	-	-	500	500
Δu_H (IU)	20,000	20,000	20,000	20,000
u_H (IU)	30,000	30,000	30,000	30,000
β	-	-	-	0.95
ϵ	-	-	0.3	-
y_L	11	10	10	10
y_H	11	12	12	12
Target	11	10–12	10–12	10–12
Constraint Limits	-	-	10–12	-

For reference, more information about the actual disturbances is shown in Figures 9 and 10. The top figures are the summation of the disturbance over the course of the simulation. The center figures are the magnitude of the disturbance at each point. The third figures are the acute disturbances that occur during the simulation, which are the same in both cases. In the simulation, the acute disturbances include seven blood loss events, and three infection events in both cases. The PK/PD parameters used in the simulator are as follows: $\alpha = 0.2718$, $C = 22.45$, $D = 6.33$, $H_{en} = 7.81$, $K_m = 76.05$, $\mu = 91.54$, $S = 0.00554$, and $V = 1655$.




Figure 9. Uncertainty in the PK/PD based simulations: Gaussian uncertainty plus acute disturbance.

The results for the Gaussian uncertainty case are presented in Table 7. The four controllers' performance statistics are shown in Figure 11 for a portion of the full 500-week simulation. The tuning parameters were chosen to try and keep the controller average ΔEPO approximately equal, so it is easier to compare the state performance between the controllers. However, the zone-MPC controller is unable to reach the same level of aggressiveness as the other three without becoming relatively unstable. Using the CVaR constraints improves the performance drastically over zone-MPC, and offers

modest improvements over traditional MPC. Using CVaR in the cost function, offers similar state performance to that of set-point tracking MPC, with the same average control action. Using CVaR in the cost function allows the medication dosing to remain stable when the hemoglobin is within the zone. As it nears the borders, the controller will typically make large moves in the medication dose. Overall, the medication dose remains very stable, even though the average ΔEPO does not portray this feature in the statistics.




Figure 10. Uncertainty in the PK/PD based simulations: uniform uncertainty plus acute disturbance.




Figure 11. Results under Gaussian process uncertainty.


Table 7. Simulation results under Gaussian distribution.

Performance Statistic	MPC	Zone-MPC	CVaR ₁	CVaR ₂
IOE ($\frac{\text{g}}{\text{dL week}}$)	25.99	30.47	25.04	39.05
PIZ (%)	80.0	76.5	81.3	80.0
EPO/week (IU)	20,044	20,059	20,399	19,795
Avg ΔEPO	1333	805	1267	1343

For the uniform distribution case, the overall results are shown in Table 8. A short window of the overall simulation can be seen in Figure 12. The CVaR constraints allow the controller to significantly improve the PIZ metric over the other controllers, while remaining with a similar average controller aggressiveness. Using CVaR in the cost function has an average change in dose similar to MPC and CVaR₁, but, looking at Figure 12, it is again easy to see that the controller rarely changes the medication dose, but, when it does so, it is typically a significant move.

Table 8. Simulation results under uniform distribution.

Performance Statistic	MPC	Zone-MPC	CVaR ₁	CVaR ₂
IOE ($\frac{\text{g}}{\text{dL week}}$)	44.58	55.12	40.7	45.10
PIZ (%)	70.4	65.2	77.8	71.7
EPO/week (IU)	20,429	20,548	20,793	20,779
Avg ΔEPO	1223	838	1236	1289

**Figure 12.** Results under uniform process uncertainty.

8. Conclusions

This work explores the use of Conditional Value at Risk (CVaR) in addressing process uncertainty in hemoglobin level control in chronic kidney disease patients. CVaR is able to utilize any type of uncertainty distribution and does not rely on Gaussian assumptions. Two methods of utilizing CVaR were introduced, one where CVaR constraints are used, and the other where the cost function is the CVaR. Simulations were performed under both Gaussian and uniform type of uncertainty distributions. A simple ARX based and a more complex PK/PD based patient simulator were used in the control performance test. Simulation results show that adding CVaR constraints to the setpoint tracking MPC problem increased the state performance over all of the other controllers. Using CVaR directly in the cost function resulted in improvements over traditional zone-MPC. Moreover, it also offers the benefit of having the medication doses remain more stable over the course of the treatment, which may be a highly sought after feature of a control method by clinicians. Based on these results, using Conditional Value at Risk is a promising tool in improving control performance in hemoglobin concentration control and EPO dosage optimization. Finally, this work has only focused on the process uncertainty and the model parameter uncertainty is skipped, which will be investigated in future work.

Author Contributions: Z.L., J.M. and J.L. participated in the design and analysis of control methods and experiments. J.M. wrote the MATLAB code and performed all the simulations. J.M. wrote the paper. Z.L., J.M. and J.L. participated in the revisions and editing of the paper. U.S. provided the clinical data for modeling purposes.

Acknowledgments: This work was supported by the Natural Sciences and Engineering Research Council of Canada.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ARX	Autoregressive with exogenous inputs
CKD	Chronic Kidney Disease
CVaR	Conditional Value at Risk
DDE	Delayed Differential Equation
EPO	Erythropoietin
IOE	Integrated Output Error
IU	International Units
PIZ	Percent of Points in the Zone
PK/PD	Pharmacokinetic and pharmacodynamic
MPC	Model Predictive Control
RBC	Red Blood Cells
VaR	Value at Risk

References

- Hayat, A.; Haria, D.; Salifu, M. Erythropoietin stimulating agents in the management of anemia of chronic kidney disease. *Patient Preference Adherence* **2008**, *2*, 195–200. [[PubMed](#)]
- Miskulin, D.; Weiner, D.; Tighiouart, H.; Ladik, V.; Servilla, K.; Zager, P.; Martin, A.; Johnson, H.; Meyer, K. Computerized decision support for EPO dosing in hemodialysis patients. *Am. J. Kidney Dis.* **2009**, *54*, 1081–1088. [[CrossRef](#)] [[PubMed](#)]
- Grosman, B.; Dassau, E.; Zisser, H.; Jovanovic, L.; Doyle, F. Zone Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. *J. Diabetes Sci. Technol.* **2010**, *4*, 961–975. [[CrossRef](#)] [[PubMed](#)]
- Batora, V.; Tarnik, M.; Murgas, J.; Schmidt, S.; Nørgaard, K.; Poulsen, N.; Madsen, H.; Boiroux, D.; Jørgensen, J. The Contribution of Glucagon in an Artificial Pancreas for People with Type 1 Diabetes. In Proceedings of the 2015 American Control Conference (ACC), Chicago, IL, USA, 1–3 July 2015; pp. 5097–5102.

5. Rivadeneira, P.; Ferramosca, A.; Gonzalez, A. Impulsive Zone Model Predictive Control with Application to Type I Diabetic Patients. In Proceedings of the 2016 IEEE Conference on Control Applications (CCA), Buenos Aires, Argentina, 19–22 September 2016; pp. 544–549.
6. Lee, J.; Gondhalekar, R.; Doyle, F. Design of an Artificial Pancreas using Zone Model Predictive Control with a Moving Horizon State Estimator. In Proceedings of the 2014 IEEE 53rd Annual Conference on Decision and Control (CDC), Los Angeles, CA, USA, 15–17 December 2014; pp. 6975–6980.
7. Bradbury, B.D.; Danese, M.D.; Gleeson, M.; Critchlow, C.W. Effect of Epoetin alfa dose changes on hemoglobin and mortality in hemodialysis patients with hemoglobin levels persistently below 11 g/dL. *Clin. J. Am. Soc. Nephrol.* **2009**, *4*, 630–637. [[CrossRef](#)] [[PubMed](#)]
8. Gaweda, A.; Jacobs, A.; Aronoff, G.; Brier, M. Model predictive control of erythropoietin administration in the anemia of ESRD. *Am. J. Kidney Dis.* **2008**, *51*, 71–79. [[CrossRef](#)] [[PubMed](#)]
9. Brier, M.; Gaweda, A.; Dailey, A.; Jacobs, A.; Aronoff, G. Randomized trial of model predictive control for improved anemia management. *Clin. J. Am. Soc. Nephrol.* **2010**, *5*, 814–820. [[CrossRef](#)] [[PubMed](#)]
10. Gaweda, A.; Jacobs, A.; Aronoff, G.; Rai, S.; Brier, M. Individualized anemia management reduces hemoglobin variability in hemodialysis patients. *J. Am. Soc. Nephrol.* **2014**, *25*, 159–166. [[CrossRef](#)] [[PubMed](#)]
11. Chait, Y.; Horowitz, J.; Nichols, B.; Shrestha, R.P.; Hollot, C.V.; Germain, M.J. Control-Relevant Erythropoiesis Modeling in End-Stage Renal Disease. *IEEE Trans. Biomed. Eng.* **2014**, *61*, 658–664. [[CrossRef](#)] [[PubMed](#)]
12. Akabua, E.; Inanc, T.; Gaweda, A.; Brier, M.; Kim, S.; Zurada, J. Individualized Model Discovery: The Case of Anemia Patients. *J. Comput. Methods Programs Biomed.* **2015**, *118*, 23–33. [[CrossRef](#)] [[PubMed](#)]
13. Alves, M.; Vilaca, S.; Carvalho, M.; Fernandes, A.; Dusse, L.; Gomes, K. Resistance of dialyzed patients to erythropoietin. *Rev. Bras Hematol. Hemoter.* **2015**, *37*, 190–197. [[CrossRef](#)] [[PubMed](#)]
14. Weis, L.; Metzger, M.; Haymann, J.; Thervet, E.; Flamant, M.; Vrtovsniak, F.; Gauci, C.; Houillier, P.; Froissart, M.; Letavernier, E.; et al. Renal Function Can Improve at Any Stage of Chronic Kidney Disease. *PLoS ONE* **2013**, *8*, e81835 [[CrossRef](#)] [[PubMed](#)]
15. Malhotra, V.; Beniwal, P.; Pursnani, L. Infections in Chronic Kidney Disease. *Clin. Queries Nephrol.* **2012**, *1*, 253–258. [[CrossRef](#)]
16. Saeed, F.; Agrawal, N.; Greenberg, E.; Holley, J. Lower Gastrointestinal Bleeding in Chronic Hemodialysis Patients. *Int. J. Nephrol.* **2011**, *2011*, 272535. [[CrossRef](#)] [[PubMed](#)]
17. Rockafellar, R.; Uryasev, S. Optimization of Conditional Value-at-Risk. *J. Risk* **2000**, *2*, 21–42. [[CrossRef](#)]
18. Parisio, A.; Molinari, M.; Varagnolo, D.; Johansson, K. A Scenario-based Predictive Control Approach to Building HVAC Management Systems. In Proceedings of the 2013 IEEE International Conference on Automation Science and Engineering (CASE), Madison, WI, USA, 17–20 August 2013; pp. 428–435.
19. Bemporad, A.; Puglia, L.; Gabbriellini, T. A Stochastic Model Predictive Control Approach to Dynamic Option Hedging with Transaction Costs. In Proceedings of the 2011 American Control Conference (ACC), San Francisco, CA, USA, 29 June–1 July 2011; pp. 3862–3867.
20. Zhang, X.; Liu, P.; Xu, C.; Ming, B.; Xie, A.; Feng, M. Conditional Value-at-Risk for Nonstationary Streamflow and Its Application for Derivation of the Adaptive Reservoir Flood Limited Water Level. *J. Water Resour. Plan. Manag.* **2018**, *144*, 04018005. [[CrossRef](#)]
21. Bemporad, A.; Morari, M. Robust model predictive control: A survey. In *Robustness in Identification and Control*; Springer: London, UK, 1999; pp. 207–226.
22. Mesbah, A. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Syst.* **2016**, *36*, 30–44. [[CrossRef](#)]
23. McAllister, J.; Li, Z.; Liu, J.; Simonsmeier, U. Erythropoietin Dose Optimization for Anemia in Chronic Kidney Disease Using Recursive Zone Model Predictive Control. *IEEE Trans. Control Syst. Technol.* **2018**, in press. [[CrossRef](#)]
24. Thanakitcharuand, P.; Jirajan, B. Prevalence of hemoglobin cycling and its clinical impact on outcomes in Thai end-stage renal disease patients treated with hemodialysis and erythropoiesis-stimulating-agents. *J. Med. Assoc. Thailand.* **2016**, *99*, 28–37.

25. Turksoy, K.; Bayrak, E.; Quinn, L.; Littlejohn, E.; Cinar, A. Multivariable Adaptive Closed-Loop Control of an Artificial Pancreas without Meal and Activity Announcement. *Diabetes Technol. Ther.* **2013**, *15*, 386–400. [[CrossRef](#)] [[PubMed](#)]
26. Ren, J.; McAllister, J.; Li, Z.; Liu, J.; Simonsmeier, U. Modeling of Hemoglobin Response to Erythropoietin Therapy through Constrained Optimization. In Proceedings of the 6th International Symposium on Advanced Control of Industrial Processes, Taipei, Taiwan, 28–31 May 2017; pp. 245–250.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).