

## TECHNICAL NOTE

# Technical note: Optimal allocation of limited proton therapy resources using model-based patient selection

Dávid Papp<sup>1</sup> | Jan Unkelbach<sup>2</sup>

<sup>1</sup>Department of Mathematics, North Carolina State University, North Carolina State University, Raleigh, North Carolina, USA

<sup>2</sup>Department of Radiation Oncology, University Hospital Zürich, University Hospital Zürich, Zürich CH 8091, Switzerland

**Correspondence**

Jan Unkelbach, University Hospital Zürich, Zürich CH 8091, Switzerland.  
Email: [jan.unkelbach@usz.ch](mailto:jan.unkelbach@usz.ch)

**Funding information**

National Science Foundation, Grant/Award Number: DMS-1847865; Swiss Cancer Research Foundation, Grant/Award Number: KFS-4528-08-2018

**Abstract**

**Purpose:** We consider the following scenario: A radiotherapy clinic has a limited number of proton therapy slots available each day to treat cancer patients of a given tumor site. The clinic's goal is to minimize the expected number of complications in the cohort of all patients of that tumor site treated at the clinic, and thereby maximize the benefit of its limited proton resources.

**Methods:** To address this problem, we extend the normal tissue complication probability (NTCP) model-based approach to proton therapy patient selection to the situation of limited resources at a given institution. We assume that, on each day, a newly diagnosed patient is scheduled for treatment at the clinic with some probability and with some benefit  $\Delta NTCP$  from protons over photons, which is drawn from a probability distribution. When a new patient is scheduled for treatment, a decision for protons or photons must be made, and a patient may wait only for a limited amount of time for a proton slot becoming available. The goal is to determine the  $\Delta NTCP$  thresholds for selecting a patient for proton therapy, which optimally balance the competing goals of making use of all available slots while not blocking slots with patients with low benefit. This problem can be formulated as a Markov decision process (MDP) and the optimal thresholds can be determined via a value-policy iteration method.

**Results:** The optimal  $\Delta NTCP$  thresholds depend on the number of available proton slots, the average number of patients under treatment, and the distribution of  $\Delta NTCP$  values. In addition, the optimal thresholds depend on the current utilization of the facility. For example, if one proton slot is available and a second frees up shortly, the optimal  $\Delta NTCP$  threshold is lower compared to a situation where all but one slot remain blocked for longer.

**Conclusions:** MDP methodology can be used to augment current NTCP model-based patient selection methods to the situation that, on any given day, the number of proton slots is limited. The optimal  $\Delta NTCP$  threshold then depends on the current utilization of the proton facility. Although, the optimal policy yields only a small nominal benefit over a constant threshold, it is more robust against variations in patient load.

**KEYWORDS**

Markov decision process, patient selection, proton therapy

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Medical Physics* published by Wiley Periodicals LLC on behalf of American Association of Physicists in Medicine.

## 1 | INTRODUCTION

Proton therapy is widely considered a superior form of radiotherapy compared to photons due to the favorable depth–dose curve of proton beams. As a rule of thumb, protons allow for reducing the integral dose to normal tissues surrounding the tumor by a factor of 2–3.<sup>1,2</sup> Although the number of proton therapy centers is increasing worldwide,<sup>1</sup> it remains a limited resource due to cost and size of proton facilities and not all patients who may benefit from proton therapy have access to it.<sup>3</sup> Consequently, concepts are being developed to optimally select patients for proton therapy to maximize its benefit for the health care system as a whole.

One approach for that is model-based patient selection, which has been promoted in the Netherlands and has been implemented clinically for several disease sites.<sup>4,5</sup> Following this approach, comparative treatment planning for protons and photons is performed in combination with normal tissue complication probability (NTCP) modeling. Patients are selected for proton therapy if the NTCP reduction of protons over photons exceeds a threshold. As an example, head and neck squamous cell carcinoma (HNSCC) is a treatment site that lends itself to the NTCP model-based patients selection because (1) the incidence of HNSCC is too high to treat all patients with protons; (2) many patients suffer from xerostomia and dysphagia,<sup>6,7</sup> leaving room for clinically significant NTCP reductions through protons; (3) patients benefit to varying degrees from protons due to the relative location of tumor and relevant organs at risk (OAR).<sup>8–10</sup> Below, we use HNSCC as an example, but we would like to stress that none of the methodology presented is specific to HNSCC and would equally apply to other sites.

In this note, we further investigate the problem of proton patient selection in the context of limited availability of proton therapy slots. The Netherlands use a nationwide proton patient selection scheme with a fixed NTCP reduction threshold, and it is assumed that every patient above the threshold can receive proton therapy timely. In this note, we instead consider proton patient selection in a single institution, for example, a photon radiotherapy department with an integrated single-room proton facility. Such an institution may, on any given day, only have a limited number of proton slots available to treat HNSCC patients, which is smaller than the number of HNSCC patients under treatment. We assume that the goal of the institution is to maximize the benefit of its proton therapy resources, that is, to minimize the expected total number of complications in the population of HNSCC patients treated at the department. In this context, proton patient selection faces the following trade-off: A high  $\Delta NTCP$  threshold may reject too many patients and lead

to proton slots being unused. A low  $\Delta NTCP$  threshold may admit patients with mediocre benefit, who block proton slots for patients who would have a larger benefit.

In this note, we present a method to determine the optimal proton patient selection threshold, depending on the number of currently available proton slots and the number of remaining fractions for proton patients currently under treatment.

## 2 | METHODS

### 2.1 | NTCP modeling

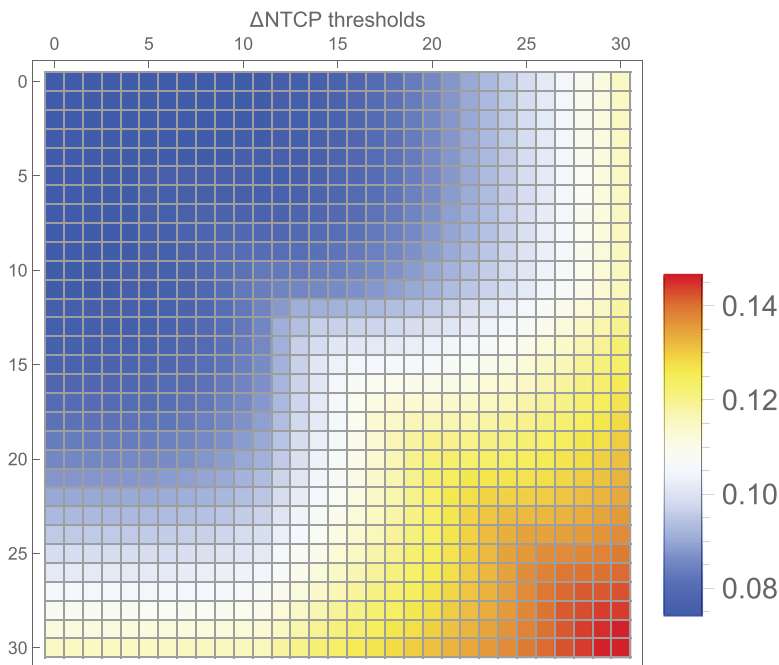
The approach discussed in this work is not limited to any particular treatment site or side effect. It only assumes a given probability distribution over  $\Delta NTCP$  for the patient population, that is, the NTCP reduction that protons yield over photons. For our motivating application, HNSCC, treatment planning studies have evaluated the expected dose reduction from proton treatments.<sup>8</sup> Additionally, several dose–response models have been proposed for the two most common side effects, xerostomia and dysphagia, based on clinical data.<sup>5,7</sup> Based on these studies, we model the distribution of the NTCP reduction by a Gaussian distribution with a mean NTCP reduction of 10% and a standard deviation of 5%. This is also consistent with the initial experience in the Dutch proton patient selection system.<sup>5</sup> Thereby, the model contains the possibility that the photon plan is superior in rare cases.

### 2.2 | Model of proton facility operation

We consider the operation of a department over discrete time steps corresponding to one working day, ignoring days when no treatments are performed. We assume that, on any given day, there is a probability  $q$  that a new HNSCC patient is scheduled for treatment at the department and a decision is made whether the patient is assigned to protons or photons. On that day, the patient has completed the standard diagnostic process, and comparative proton–photon treatment planning has been performed. The parameter  $q$  represents the patient load. As an example, we consider  $q = 0.4$  in this work, which would correspond to a clinic treating approximately 100 HNSCC patients per year, and thus, two new patients starting treatment each week on average. We further assume a 30-fraction treatment over 6 weeks, meaning that on average 12 patients are under treatment on any day. The facility is assumed to have  $N$  proton slots available for HNSCC patients, where we assume  $N$  to be smaller than the number of patients under treatment.

In reality, comparative treatment planning for protons and photons would be performed for each new

<sup>1</sup> Particle Therapy Co-Operative Group. Particle therapy facilities in clinical operation. Available at: <https://www.ptcog.ch/index.php/facilities-in-operation>



**FIGURE 1** A cross-section of the optimal policy for 3 proton slots assuming that one of them is immediately available. The numbers on the two axes represent the number of days until the other two slots become available, the colors indicate the optimal  $\Delta NTCP$  threshold that a patient needs to exceed to be assigned a proton slot

patient to obtain NTCP values  $NTCP^\gamma$  and  $NTCP^p$  for photon and proton plans, respectively. To simulate this process, we randomly draw a value for the benefit  $\Delta NTCP = NTCP^\gamma - NTCP^p$  from a Gaussian distribution, as outlined above.

When a patient is assigned to proton therapy, the patient blocks one proton slot for the following 30 workdays. We assume that each new patient has to start treatment shortly after they are assigned to a treatment slot, that is, a patient can wait only for up to  $T$  days for a proton slot to become available. The value of  $T$  is a clinical decision to be made by the department to limit the adverse effects of treatment delays. For this work we assume  $T = 10$ , meaning that a patient must be assigned to photons if all proton slots are blocked for the next 2 weeks.

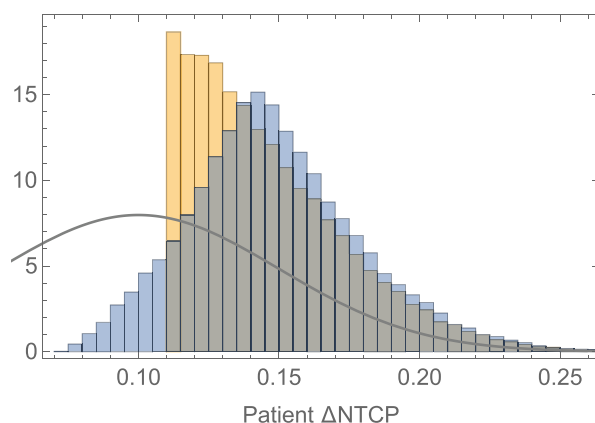
### 2.3 | Determining the optimal NTCP thresholds

We now present the methodology to determine the optimal patient selection thresholds that minimize the expected number of complications. To that end, the problem of proton patient selection is formulated as an (infinite horizon) Markov decision process (MDP).<sup>11</sup> MDP is a mathematical framework that enables optimal sequential decision making under uncertainty in situations where the goal is to maximize expected (long-term average) reward in a system whose state evolves partly at random and partly as a function of the sequential decisions. The optimal policy (course of action) of the system is state dependent. The details of the MDP

formulation of the problem of proton patient selection are as follows:

#### State.

The state of the MDP represents the state of the proton facility, which can be described by one integer  $n \in \{0, \dots, 40\}$  per slot indicating for how many additional days the slot is blocked. (The largest number is, thus, the length of the treatment plus the maximum wait time  $T$ .) For the example, if  $N = 3$ , then the state  $(5, 14, 27)$  would indicate that currently all 3 slots are occupied, and



**FIGURE 2** The distribution of the  $\Delta NTCP$  of patients assigned to proton therapy by the optimal policy (blue histogram) and the best constant-threshold policy, with  $\Delta NTCP$  threshold of 11% (yellow histogram). Also shown: the  $\Delta NTCP$  distribution of the patient population (solid curve). The constant-threshold policy favors patients with  $\Delta NTCP$  right above the threshold

the proton patients under treatment have to complete another 5, 14, and 27 fractions, respectively.

*Actions and policy.*

An action amounts to the decision whether or not to assign a new patient to proton therapy. A policy is given by a  $\Delta NTCP$  threshold for every facility state; if a new patient's  $\Delta NTCP$  is above the threshold corresponding to the state of the proton facility at the time the patient presents, then the patient is assigned a proton slot. Only those states have an assigned  $\Delta NTCP$  threshold in which one of the slots are occupied for less than  $T$  remaining days.

*State transitions.*

Every state has up to two possible successor states, depending on whether the state of the proton facility allows a new patient to be assigned a proton slot. For the example state (5, 14, 27) there are two possible state transitions. If on that day no new patient presents or the patient is assigned to photons, the state transitions to (4, 13, 26). If a new patient presents, it is possible to assign the patient to the slot that becomes available in 4 days, and in that case the state transitions to (34, 13, 26). State transitions are stochastic because the sequence of incoming patients is. The probability for the latter transition is given by the probability of seeing a new patient  $q$  multiplied with the probability that the patient's  $\Delta NTCP$  exceeds the selection threshold.

*Reward.*

Our goal is to minimize the expected number of complications over time, or equivalently, to maximize the (long-term) average  $\Delta NTCP$  realized by the patients. We may consider the situation that all patients are treated with photons as the baseline; if a patient is assigned to protons, an immediate benefit of  $\Delta NTCP$  is realized. The optimal policy is the collection of state-dependent  $\Delta NTCP$  thresholds that maximizes the average (per patient) realized  $\Delta NTCP$  benefit.

*Computing the optimal policy.*

The optimal policy of an MDP is characterized by the *Bellman equation*. A general algorithm to compute the optimal policy of an MDP is the iterated value and policy iteration, which can be briefly summarized as follows<sup>11</sup>: Start with any policy, for example, with every threshold being 0 (which means assigning every patient to protons if a slot is available), and then iterate the following two steps until no change is observed in the policy:

1. **Policy evaluation.** Compute the long-term expected reward for the current policy and the value of each state  $s$ , defined as the expected reward given that the system starts at state  $s$ . This can be accomplished by solving a system of linear equations, as follows.

Suppose that  $S$  is the number of states, that  $\mathbf{r} \in \mathbb{R}^S$  is the vector whose  $i$ th component is the immediate reward realized by the given policy in state  $i$ , and that  $\mathbf{P} \in \mathbb{R}^{S \times S}$  is the state transition matrix whose  $(i, j)$ th component is the probability that after the action determined by the current policy the system's next state will be  $j$  assuming that its current state is  $i$ . Finally, let  $\alpha < 1$  be a positive discount factor. Then the vector of values of each state  $\mathbf{v} \in \mathbb{R}^S$  is the (unique) solution of the linear system

$$\mathbf{v} = \mathbf{r} + \alpha \mathbf{Pv}. \tag{1}$$

In our computation, we used  $\alpha = 0.99999$ , as the benefit realized by a patient arriving later is considered to be the same as for earlier patients, but  $\alpha < 1$  is necessary for the convergence of the algorithm.

2. **Value iteration.** For each state  $s$ , compute the optimal policy assuming that the long-term rewards starting from each state are the values computed in the previous step. In our problem, this is a straight-forward optimization problem with a closed-form solution: For every state where a decision is to be made (there is room in the proton facility for a new patient), assign the new patient to a proton slot if and only if their  $\Delta NTCP$  is at least as large as the difference between the values of the states and the proton facility after assigning the patient to protons or photons, respectively. In other words, the optimal  $\Delta NTCP$  threshold is the difference between the values of the two possible subsequent states.


It can be shown that this iterative process converges to the optimal policy. We remark that the optimal policy can also be computed using linear programming, although that approach was not used in our study.

**2.4 | Comparison to a constant  $\Delta NTCP$  threshold**

The optimal policy is compared to a simpler patient selection scheme using a constant  $\Delta NTCP$  threshold that is independent of current facility utilization but optimal for the assumed values of  $N$ ,  $q$ , and  $T$  and the assumed distribution over  $\Delta NTCP$  in the patient population. The (approximately) optimal constant threshold was computed by simulating the operation of the clinic over an extended period of time for constant thresholds between 0% and 30% in 1% increments; see Figure 3.

**3 | RESULTS**

We computed the optimal policies assuming that a radiotherapy facility has up to three concurrent proton therapy



**FIGURE 3** Top panel: observed benefit using a constant  $\Delta NTCP$  threshold independent of the current utilization of the facility, with the number of slots  $N = 3$  and maximum delay  $T = 10$  days fixed. Bottom panel: the fraction of patients assigned to proton therapy. A threshold of 0% corresponds to the first-come-first-served policy. The optimal constant threshold is 11%

slots for HNSCC patients, who receive 30-fraction treatments that begin after a delay of no more than 10 days in order to wait for an available proton therapy slot. We also investigated the effect of delayed treatments, and the robustness of policies against uncertain patient arrival rates.

#### Calibrating the expected benefit.

The overall benefit of the patient population is measured by the long-term average benefit, that is, the expected value of NTCP reduction, which is equivalent to measuring the average number of prevented side effects. Based on our  $\Delta NTCP$  model, the expected benefit if everyone receives proton treatment is 10%. By definition, the baseline policy of everyone receiving photon treatment corresponds to a benefit of 0. Since, with two new patients presenting each week, we have 12 patients concurrently receiving treatment on average, but only  $N \leq 3$  slots, a first-come-first-served assignment of patients (where an available proton slot is assigned to the next arriving patient) can only yield a benefit of around  $10/(12/N)\%$ , that is, around 2.5% for  $N = 3$ . At the other extreme, if we could treat the quarter of the patients with the highest benefit (corresponding

to a  $\Delta NTCP$  threshold of 13.4%), the average  $\Delta NTCP$  for the proton patients would be 16.4%, increasing the overall benefit of protons for the population to 4.1%. However, when the number of proton patients is constrained to  $N \leq 3$  on any given day, the 4.1% represents an unachievable upper bound and the benefit of the optimal proton patient selection scheme will lie between 2.5% and 4.1%.

### 3.1 | The structure of the optimal policy


The optimal  $\Delta NTCP$  threshold depends on the current utilization of the proton facility. For  $N = 3$  slots, the optimal  $\Delta NTCP$  threshold varied greatly, from 0.074 in the state when all proton slots are immediately available, to 0.16 in the state when one slot will be available in  $T = 10$  days and the others have just been taken. Figure 1 shows a two-dimensional cross section of this policy for the states when one proton slot is immediately available. Unsurprisingly, the threshold is monotone in the availability: When a second slot is available sooner, then admitting a patient with less benefit carries a lower risk of blocking all slots from a patient with greater benefit.

The long-term benefit of this policy is  $\Delta NTCP = 3.57\%$ ; closing 67% of the gap between the first-come-first-served policy and the unattainable outcome of filling all treatment slots with patients of the highest benefit. Simulating the facility operation using this policy provides further insight into the utilization of proton slots. Recall that with  $N = 3$  treatment slots and 30-day treatments the best one can hope for is that 25% of the patients may receive proton therapy. The policy in Figure 1 assigns 24% of the patients to proton therapy and thus uses most of the proton slots. Approximately 5% of all proton slots remain unused and approximately 14% of the treatment days leave at least one proton slot empty. The blue histogram in Figure 2 shows the distribution of  $\Delta NTCP$  values for the patients treated with protons.

### 3.2 | Comparison to simpler policies with constant threshold

Figure 3 investigates the simplified strategy of using a constant  $\Delta NTCP$  threshold independent of the current utilization of the facility, for the nominal parameter values  $N = 3$ ,  $q = 0.4$ , and  $T = 10$ . A threshold of 0% corresponds to random patient selection, that is, giving an available proton slot to the next patient who benefits from proton treatment. This yields an average NTCP reduction of 2.5%, reflecting that 25% of patients are treated with protons. The optimal constant threshold of 11% yields an average NTCP reduction of 3.4% (compared to 3.57% for the optimal facility state-dependent thresholds). A constant  $\Delta NTCP$  threshold





**FIGURE 4** The dependence of the average benefit of optimal proton slot allocation on the maximum allowed delay  $T$  until treatment must begin

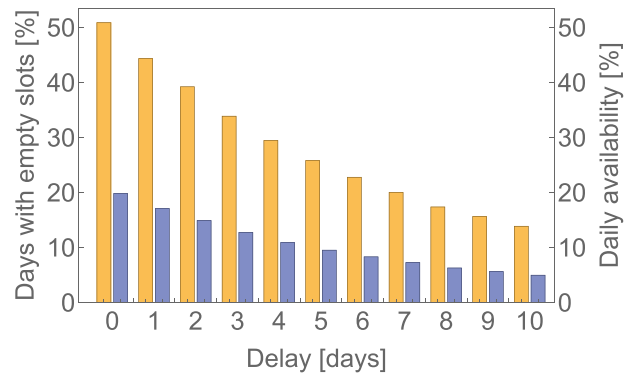
higher than the optimal leads to many unused proton slots (Figure 3, bottom panel), which rapidly reduces the benefit below even that of the first-come-first-served policy. The lower average NTCP reduction of the constant threshold policy can be explained by two factors. At the optimal constant threshold of 11%, 23% of patients are treated with protons (compared to 24% for the facility state-dependent policy). In addition, the facility state-dependent policy results in a distribution of realized  $\Delta NTCP$  reductions with higher mean (Figure 2).

### 3.3 | Dependence on parameters

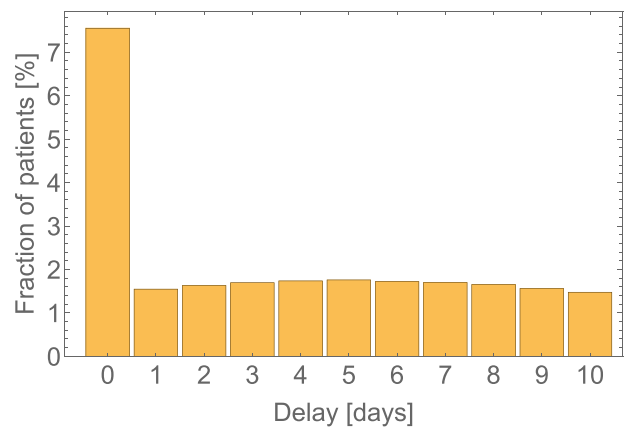
The optimal policy depends on various parameters. Naturally, the  $\Delta NTCP$  thresholds scale with parameters of the probability distribution over  $\Delta NTCP$ . Doubling both mean and standard deviation would simply double the  $\Delta NTCP$  thresholds and the benefit. Furthermore, increasing  $N$  or decreasing  $q$  would increase the percentage of patients that can be treated with protons, and thus corresponds to lower  $\Delta NTCP$  thresholds and higher average benefit from protons.

#### 3.3.1 | The benefit of treatment delays

We also investigated how the utilization of proton slots and the overall benefit is impacted by the maximum allowed delay  $T$ . We fixed the number of proton slots at  $N = 3$  and considered a maximum delay  $T$  ranging from 0 to 10 days. Figure 4 shows the dependence of the benefit on the maximum delay  $T$  until treatment must begin. In our example it is clear that, as long as the wait does not compromise the treatment outcome, allowing longer delays is beneficial. This is explained by the increasing utilization of the proton facility allowed by the delayed treatment (Figure 5). When no treatment delay is allowed ( $T = 0$ ), approximately 20% of all pro-



**FIGURE 5** Availability of proton slots as a function the maximum delay, with the number of slots  $N = 3$  fixed. Yellow bars (left): the fraction of days with an immediately available proton slot. Blue bars (right): the average fraction of proton slots available



**FIGURE 6** The frequency of patients who received proton treatment delayed by  $t$  days ( $t = 0, \dots, 10$ ). Although the maximum allowed delay is 10 days, a large number of patients can be assigned a proton slot immediately

ton slots remain vacant, which is reduced to about 5% when a delay of 2 weeks ( $T = 10$ ) is allowed (blue bars in Figure 5). This is because patients can be assigned a proton slot even if they present when all slots are taken. An interesting feature of the optimal policy is that, even with a relatively long maximum delay of 10 days, most patients have very little wait; see Figure 6.

#### 3.3.2 | Robustness of patient selection policies

It is also instructive to compare the performances of the optimal state-dependent threshold policy and the optimal fixed-threshold policy that are computed assuming a fixed known patient load  $q$  under the assumption of *different* patient loads than the estimated one. This is indicative of the policies' performance under incorrectly specified or time-varying patient load. Table 1 shows the computed long-term expected benefit ( $\Delta NTCP$ ) of

**TABLE 1** Investigating the robustness of the  $\Delta NTCP$  benefit of the state-dependent slot allocation policy versus the constant-threshold policy against unknown or varying patient load. Off-diagonal entries represent scenarios when the assumed patient load  $q$  over- or underestimates the true patient load  $q^*$ . Note that, for constant  $N = 3$ , higher patient load  $q^*$  means that a lower percentage of patients is treated with protons, leading to a reduced average benefit

	State-dependent threshold			Constant threshold		
	$q^* = 0.2$	$q^* = 0.4$	$q^* = 0.6$	$q^* = 0.2$	$q^* = 0.4$	$q^* = 0.6$
$q = 0.2$	5.72%			5.62%		
$q = 0.4$	5.56%	3.57%	2.56%	5.07%	3.40%	2.42%
$q = 0.6$	5.19%		2.60%	3.99%		2.49%

both policies assuming different combinations of the assumed patient loads  $q \in \{0.2, 0.4, 0.6\}$  used to compute the optimal policies and the true arrival rate  $q^* \in \{0.2, 0.4, 0.6\}$  (corresponding to 1, 2, 3 new patients per week).

The table reveals that the reduction in the expected  $\Delta NTCP$  benefit that results from a misspecified  $q$  is considerably smaller for the state-dependent threshold policy than it is for the fixed-threshold policy, particularly when the true patient load  $q^*$  is overestimated. This is explained by the fact that the state-dependent threshold adapts to the lack of new patient arrivals by lowering the threshold of admittance to the proton facility, whereas the fixed-threshold strategy is too rigid and results in a greater underutilization of the proton facility. Figure 3 also explains why overestimating  $q^*$  is worse than underestimating it: The  $\Delta NTCP$  benefit rapidly declines to zero if the threshold is higher than necessary.

## 4 | DISCUSSION

With the limited availability of proton therapy, efforts should be made to optimally select proton patients to maximize the overall benefit of the existing facilities. NTCP model-based patient selection represents one approach to this problem.<sup>4,5</sup> Currently, the approach is implemented with a fixed  $\Delta NTCP$  threshold in countries with a nation-wide referral system, where one may assume that a proton slot can be made available to an eligible patient in a timely manner. In this note, we instead consider the situation that the number of proton slots on any given day is limited. This may, for example, be the case for an individual clinic with a single room proton machine. The optimal  $\Delta NTCP$  threshold then depends on the number of available proton slots and their current utilization.

We compared patient selection for the optimal, facility state-dependent  $\Delta NTCP$  threshold to the optimal constant threshold, that is, the  $\Delta NTCP$  threshold that is optimal for the number of available slots, the assumed probability distribution over  $\Delta NTCP$  in the patient pop-

ulation, and the average number of patients under treatment, but independent of the current slot utilization. It was found that the additional benefit of facility state-dependent thresholds was quite modest if both policies are evaluated for parameters they were optimized for. Nevertheless, there are arguments in favor of a state-dependent policy. (1) Using a state-dependent policy does neither cause any technical difficulties for implementation, nor does it increase the workload for the clinic compared to a constant threshold as the optimal policy needs to be computed only once. (2) A state-dependent policy is in line with the intuition that one should lower the  $\Delta NTCP$  threshold of the proton facility if underutilized and increase the threshold if the patient load is high. (3) The latter aspect makes the patient selection more robust against variations in the patient load.

Calculating the optimal  $\Delta NTCP$  thresholds (both state-dependent or constant) depends on parameter values and model assumptions that may not fully represent the complexity of clinical reality. In this work, we assumed that a fixed number of proton slots  $N$  is available each day for treatments of a given disease site such as HNSCC. In reality, this number would likely fluctuate because the number of patients of other tumor sites routinely referred to proton therapy will fluctuate. This situation could be modeled via extensions of the MDP model. For example, one could assume that patients with certain tumor sites are guaranteed a proton slot, and that each day there is some probability that such a patient presents. Then, the next available proton slot would have to be assigned to this patient, temporarily reducing the number of proton slots available to HNSCC.

In this note, we assumed that the benefit of protons is described by a single  $\Delta NTCP$  value. In reality, different toxicities of varying severity may be relevant. The patient selection method presented in this note may then be applied with  $\Delta NTCP$  representing a weighted sum of  $\Delta NTCP$  values for different side effects. Although the relative importance of distinct side effects may be debatable, this difficulty is not specific to the method presented here but is inherent to any NTCP model-based selection scheme. For example, in the Dutch scheme,<sup>5</sup> grade  $\geq 2$  xerostomia and grade  $\geq 2$  dysphagia are treated as equally important, corresponding to equal weights for both toxicities. A reduction of 5% in tube feeding dependence is considered as worthwhile as a 10% reduction in grade  $\geq 2$  toxicity, which could be represented by double the weight in the sum of  $\Delta NTCP$  values.

Another approach to make optimal use of limited proton slots is combined proton-photon radiotherapy with optimal slot assignment. In such treatments, for each patient, some fractions are delivered with protons and others with photons. Loizeau et al.<sup>12</sup> suggested reassigning proton therapy slots on a daily basis by calculating the incremental NTCP reduction from one

additional proton fraction for all patients currently under treatment. Every day, the proton slots are then assigned to those patients who benefit the most from one additional proton fraction. Thereby, all proton slots can be used in a near-optimal manner.

## 5 | CONCLUSION

In the situation that, on any given day, the number of proton slots is limited, the optimal  $\Delta NTCP$  threshold for proton patient selection depends on the current utilization of the proton facility. Such a facility state-dependent policy yields a modest improvement in the average  $\Delta NTCP$  improvement from protons compared to a simplified policy using a constant threshold when evaluated for the expected patient load. However, the facility state-dependent policy is more robust against variations in patient load.

## ACKNOWLEDGMENT

This material is based on work supported by the National Science Foundation under Grant No. DMS-1847865 and the Swiss Cancer Research Foundation under Grant No. KFS-4528-08-2018.

Open access funding provided by Universitat Zurich.

## CONFLICTS OF INTEREST

The authors have no conflicts to disclose.

## REFERENCES

1. Brada M, Bortfeld T. Proton therapy: the present and the future. *Semin Radiat Oncol.* 2013;23:75-6.
2. Mohan R, Grosshans D. Proton therapy—p and future. *Adv Drug Deliv Rev.* 2017;109:26-44.
3. Glimelius B, Ask A, Bjelkengren G, et al. Number of patients potentially eligible for proton therapy. *Acta Oncol.* 2005;44:836-849. <https://doi.org/10.1080/02841860500361049>

4. Langendijk JA, Lambin P, De Ruyscher D, Widder J, Bos M, Verheij M. Selection of patients for radiotherapy with protons aiming at reduction of side effects: the model-based approach. *Radiother Oncol.* 2013;107:267-273.
5. Tambas M, Steenbakkers RJHM, Laan HP, et al. First experience with model-based selection of head and neck cancer patients for proton therapy. *Radiother Oncol.* 2020;151:206-213.
6. Christianen ME, Schilstra C, Beetz I, et al. Predictive modelling for swallowing dysfunction after primary (chemo)radiation: results of a prospective observational study. *Radiother Oncol.* 2012;105:107-114.
7. Houweling AC, Philippens MEP, Dijkema T, et al. A comparison of dose-response models for the parotid gland in a large group of head-and-neck cancer patients. *Int J Radiat Oncol Biol Phys.* 2010;76:1259-1265.
8. Jakobi A, Bandurska-Luque A, Stützer K, et al. Identification of patient benefit from proton therapy for advanced head and neck cancer patients based on individual and subgroup normal tissue complication probability analysis. *Int J Radiat Oncol Biol Phys.* 2015;92:1165-1174.
9. Water TA, Lomax AJ, Bijl HP, et al. Potential benefits of scanned intensity-modulated proton therapy versus advanced photon therapy with regard to sparing of the salivary glands in oropharyngeal cancer. *Int J Radiat Oncol Biol Phys.* 2011;79:1216-1224.
10. Laan HP, Water TA, Herpt HE, et al. The potential of intensity-modulated proton radiotherapy to reduce swallowing dysfunction in the treatment of head and neck cancer: a planning comparative study. *Acta Oncol.* 2013;52:561-9.
11. Winston WL. *Operations Research.* 4th ed. Brooks/Cole; 2004.
12. Loizeau N, Fabiano S, Papp D, et al. Optimal allocation of proton therapy slots in combined proton-photon radiation therapy. *Int J Radiat Oncol Biol Phys.* 2021;111:196-207.

**How to cite this article:** Papp D, Unkelbach J. Technical note: Optimal allocation of limited proton therapy resources using model-based patient selection. *Med Phys.* 2022;49:4980–4987. <https://doi.org/10.1002/mp.15812>