Adjustable robust treatment-length optimization in radiation therapy

S.C.M. ten Eikelder*   A. Ajdari†   T. Bortfeld†   D. den Hertog‡

February 12, 2021

Traditionally, optimization of radiation therapy (RT) treatment plans has been done before the initiation of RT course, using population-wide estimates for patients' response to therapy. However, recent technological advancements have enabled monitoring individual patient response during the RT course, in the form of biomarkers. Although biomarker data remains subject to substantial uncertainties, information extracted from this data may allow the RT plan to be adapted in a biologically informative way. We present a mathematical framework that optimally adapts the treatment-length of an RT plan based on the acquired mid-treatment biomarker information, while accounting for the inexact nature of this information. We formulate the *adaptive treatment-length optimization problem* as a 2-stage problem, wherein the information about the model parameters gathered during the first stage influences the decisions in the second stage. Using Adjustable Robust Optimization (ARO) techniques we derive explicit optimal decision rules for the stage-2 decisions and solve the optimization problem. The problem allows for multiple worst-case optimal solutions. To discriminate between these, we introduce the concept of Pareto Adjustable Robustly Optimal ($\mathcal{PARO}$) solutions.

In numerical experiments using lung cancer patient data, the ARO method is benchmarked against several other static and adaptive methods. In the case of exact biomarker information, there is sufficient space to adapt, and numerical results show that taking into account both robustness and adaptability is not necessary. In the case of inexact biomarker information, accounting for adaptability and inexactness of biomarker information is particularly beneficial when robustness (w.r.t. organ-at-risk (OAR) constraint violations) is of high importance. If minor OAR violations are allowed, a nominal folding horizon approach (NOM-FH) is a good performing alternative, which can outperform ARO. Both the difference in performance and the magnitude of OAR violations of NOM-FH are highly influenced by the biomarker information quality.

Keywords: radiation therapy, adjustable robust optimization, decision rules, optimal stopping

*Department of Econometrics and Operations Research, Tilburg University, The Netherlands
†Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, USA
‡Department of Operations Management, University of Amsterdam, The Netherlands

1 Introduction

In radiation therapy (RT), the goal is to deliver a curative amount of radiation dose to the target volume(s), while keeping the dose to all organs-at-risk (OARs) within tolerable limits. As the radiation beam delivers energy to all tissues that are on its path, the OARs will (often) inevitably receive some dose as well. Treatment planning has a spatial component, where the optimal dose distribution is determined, and a temporal component, where the optimal number of treatment sessions, or fractions, is determined. Whereas the former is predominantly a geometric problem, the latter involves radiobiological effects.

Technological advances in treatment monitoring through imaging and other forms of data acquisition allow for a more accurate assessment of an individual's radiation response (Baumann et al. 2016). Biologically-based adaptive treatments aim to monitor the treatment, acquire mid-treatment biomarker information, and adapt the remainder of the treatment course accordingly. Many approaches to adaptive treatment planning have been studied in the literature. To the best of our knowledge, all existing approaches assume that all information acquired mid-treatment is exact, i.e., gives a perfect representation of the current state of treatment response. Unfortunately, the limited availability and accuracy of required biomarkers pose a primary challenge for adaptive treatments (Baumann et al. 2016). Any information from biomarker data acquired during treatment remains subject to uncertainties, stemming from both measurement errors and the inexactness in the translation of measured data to usable model parameters. Therefore, it is crucial that any adaptive treatment planning approach takes this into account. Ajdari et al. (2019) provide an overview of relevant mathematical (optimization) tools. We present an approach to optimally adapt the treatment length of RT using inexact mid-treatment biomarker information.

Specifically, we take an adjustable robust optimization (ARO, Ben-Tal et al. 2004, Yanıkoğlu et al. 2019) approach that accounts for the inexactness of biomarker information. ARO is an extension of robust optimization (RO) that takes into account the flow of information over time and exploits the fact that some decisions need to be taken only after the data has (partially) revealed itself. By using ARO, we ensure that the treatment plan delivered in the initial treatment stage (prior to obtaining biomarker information) is 'adaptation aware'. That is, it is designed with adaptation in mind, which may yield more flexibility at the time of treatment adaptation. In the standard paradigm, ARO assumes that the revealed information is exact; we employ the ARO methodology developed in De Ruiter et al. (2017) for the case when revealed information is not exact, but provides only an estimate of the true parameters.

Contributions

We consider a stylized two-stage ARO model to optimally adapt the treatment-length based on inexact biomarker information acquired mid-treatment. Although the stylized model makes several simplifying assumptions to aid the analysis, we believe it captures several important characteristics of realistic adaptive treatment planning, and it enables a precise analysis of the influence of uncertainty in biomarker information. Our main contributions are:

- We develop mathematical tools based on ARO that enable us to (i) optimally adapt the dose per fraction and treatment length after acquiring mid-treatment biomarker information, (ii) analyze the influence of biomarker information uncertainty.

- We present explicit optimal decision rules for a difficult (non-convex, mixed-integer) yet practically relevant ARO problem.

- We show that there are multiple optimal solutions for the worst-case scenario, and that these differ in performance in non-worst-case scenarios. To handle this, we introduce the concept of Pareto Adjustable Robustly Optimal ($\mathcal{PARO}$) solutions, a generalization

of Pareto Robustly Optimal ($\mathcal{PRO}$) solutions (Iancu and Trichakis 2014) to two-stage robust optimization problems. In case the acquired biomarker information is exact, $\mathcal{PARO}$ solutions are obtained.

- We perform a computational study using real lung cancer patient data to determine the optimal timing of acquiring biomarker information in case biomarker quality improves over time. Later biomarker acquisition also limits adaptation possibilities, and the optimal balance depends on the improvement rate.

Literature review

There is a large body of adaptive treatment planning research in RT, the majority of which focusses not on biologically-based uncertainties but on geometric uncertainties and anatomical changes. Chan and Mišić (2013), Mar and Chan (2015), Böck et al. (2017), Böck et al. (2019) and Lim et al. (2020) present adaptive treatment planning approaches that starts with delivery of the original treatment plan, often derived using RO methods. At given adaptation moments, the 'state' of the patient (e.g., anatomical changes, tumor shrinkage, breathing motion pattern) is observed, and the treatment is re-optimized for the remainder of the treatment plan. In RO terminology, these approaches are known as folding horizon (FH) methods. They disregard adaptation possibilities initially, and re-optimize the updated model once mid-treatment information is acquired.

Several treatment plan adaptation approaches have been proposed for adapting to biological information, differing in considered biomarker information, adapted treatment plan decisions and used methodology. Ghate (2011) and Kim et al. (2012) propose a theoretical stochastic control framework to optimally adapt the dose distribution over a fixed number of fractions, based on hypothetically-observed tumor states. Saberian et al. (2016b) concretize this theoretical framework, using simulated hypoxia (insufficient oxygen supply at cellular level) status as biomarker. Long (2015) consider a model with a constraint on the probability of radiation-induced lung toxicity, which depends on an a priori unknown model parameter. The problem is formulated as a two-stage model (before and after parameter observation), and the optimal dose distribution is determined for each stage. They consider a finite scenario set for the parameter, and the lung toxicity constraint either has to hold in expectation or has to hold for all considered scenarios.

Nohadani and Roy (2017) consider a two-stage model to adapt to hypoxia information, where the uncertainty is time-dependent. As the hypoxia information ages the uncertainty grows, until it resets at the observation moment. In each stage the dose distribution is optimized; for the second stage a finite adaptability approach is taken. It is shown that total information degradation is minimized if the observation moment is set mid-treatment. Dabadghao and Roy (2020) consider a similar time-dependent uncertainty set, for adapting to hypoxia information in a multi-stage setting. In each stage the mean tumor dose per fraction is optimized using a folding horizon approach. It is shown that total information degradation is minimized if the observation moments are set equidistant. They introduce a cost of observation (additional dose due to mid-treatment positron emission tomography (PET) scans), and determine the optimal number of observations. Both papers assume that the hypoxia state varies over time, and is exactly observed at the observation moment(s). In contrast, we assume that uncertain parameters are constant in time, and consider inexact biomarker information. Moreover, they solely adapt the dose, whereas we also adapt the total number of treatment fractions itself.

Adapting the treatment length (i.e., the fractionation schedule) based on observed radiation response has been studied before in the literature. Saka et al. (2014) consider a two-stage model where after a fixed number of treatment fractions hypoxia information is acquired. Based on this information both the remaining number of treatment fractions and the dose distribution are re-optimized, in order to maximize average hypoxia-corrected tumor dose. They focus on

maintaining hypoxia-corrected fraction size requirements. A similar approach is taken by Ajdari et al. (2018), where after each treatment fraction the tumor cell density in each voxel is observed, and adaptations are made after each treatment fractions instead of only once. The objective is to minimize the total number of tumor cells remaining (TNTCR) at the end of the treatment course. Both approaches can be considered FH methods. In contrast to our approach, it is assumed that any information acquired mid-treatment is exact.

Iancu et al. (2020) propose a conceptual robust monitoring and stopping model. They consider a system with a state $x(t)$, and after each observation moment the uncertainty in the system state $x(t)$ grows as t increases. At a new observation moment, the uncertainty reduces to zero, i.e., an exact observation is made. They consider multiple observation moments, and the goal is to time these optimally. At each observation moment, the (state-dependent) direct stopping reward is compared to the worst-case continuation reward, and the according action is taken. Their model does not allow for controls that influence state variables, i.e., applying their model to RT optimization problems would not allow to adjust the dose distribution or the mean dose per fraction.

Notation and organization

All variables and constants are 1-dimensional (belong to $\mathbb{R}$ or $\mathbb{N}$) unless indicated otherwise. In functions, a semicolon (;) is used to separate variables and constant arguments from uncertain parameters. Optimal solutions to optimization problems are indicated with an asterisk (*). Properties of optimal solutions to optimization problems have calligraphic font (e.g., $\mathcal{ARO}$) to distinguish them from methods with the same or similar abbreviations.

The remainder of this paper is organized as follows. Section 2 introduces the used biological models, background information on biomarkers and states modeling choices. Section 3 introduces the adjustable treatment-length optimization problem under the assumption of exact biomarker information and solves this using ARO techniques. Section 4 generalizes this to inexact biomarker information. Section 5 presents and discusses results of numerical experiments on a lung cancer data set. Finally, Section 6 concludes the paper.

2 Adaptive fractionation

2.1 The fractionation problem

Spatial optimization exploits the fact that, by mounting the beam head on a gantry, the tumor can be targeted from various angles. It aims to find the combination of beam angles and weights that gives the best trade-off between tumor dose conformity and healthy tissue sparing. There is a large body of literature on this topic, see for example Shepard et al. (1999), Ehrgott et al. (2008) for reviews. The result of the spatial optimization problem is a combination of beam angles and weights that gives the best trade-off between tumor dose conformity and healthy tissue sparing. The resulting dose distribution gives the dose to each voxel (3-dimensional subvolume) of the tumor and OARs. Figure 1 gives an example of a slice of a dose distribution. The beam angles and weights are chosen such that the target (contoured in black) receives a high dose and a nearby organ-at-risk (OAR, contoured in red) receives a low dose.

Typically, this dose is not delivered in a single treatment session, but spread out over multiple treatment fractions (fx). The underlying idea is that compared to tumor cells, healthy tissues often have better repair capabilities between fractions (Fowler 1989, Withers 1985). On the other hand, a treatment spread-out over a large number of treatment fractions may not deliver sufficient damage to the target volume, and increases the risk of tumor proliferation. The effect of fractionation differs per healthy tissue type and per tumor site, see, e.g., Hall and Giaccia (2012) for further details. Determining the optimal number of treatment fractions is known as the fractionation problem. Treatments with a higher number of fractions and a lower dose per
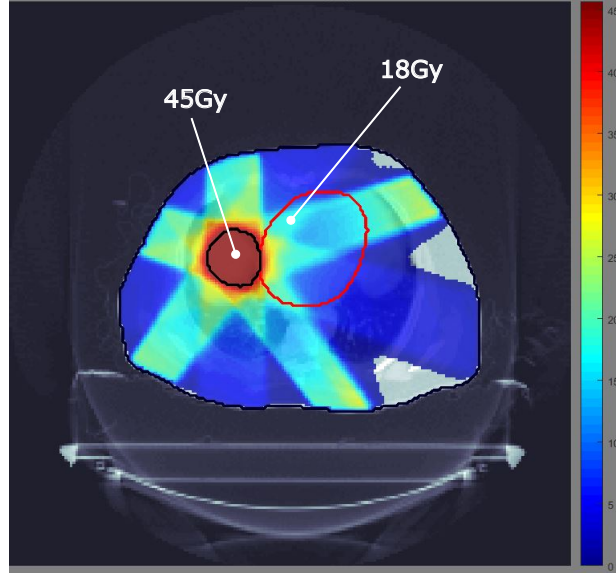
Figure 1: Slice of a 3D dose distribution (illustration). Target is contoured in black, OAR is contoured in red. The target volume has a mean dose of 45Gy. The indicated OAR voxel with dose 18Gy has a dose sparing factor of 0.4.

fraction than the conventional regimen are known as *hyperfractionated* treatments. Treatments with a lower number of fractions and a higher dose per fraction than conventional are known as *hypofractionated* treatments.

In each treatment fraction, a scaled version of the dose distribution is administered. Scaling the dose distribution influences the absolute dose delivered to each voxel, but the relative dose remains unchanged. That is, for each voxel we can define a *dose sparing factor*: the dose of this voxel as a fraction of the mean target dose. In Figure 1, the target volume receives a uniform dose of 45Gy over the entire treatment course. The indicated OAR voxel receives 18Gy and thus has a dose sparing factor of 0.4. Consequently, if in an individual treatment fraction we administer a mean target dose of 3Gy, the OAR voxel receives 1.2Gy. Thus, with a fixed dose distribution, the sole decision in treatment fraction t is d_t, the mean tumor dose in that fraction. The corresponding dose to any voxel i with dose sparing factor s_i is $s_i d_t$. Altogether, with a fixed dose distribution fractionation problems can be described using a low number of decision variables. Typically, the target dose is homogeneous, so for simplicity we assume that each target voxel receives dose d_t (i.e., has dose sparing factor 1). Nevertheless, we emphasize that this modeling approach allows for both heterogeneous dose distributions in target and OAR volumes.

The BED model (Fowler 1989, 2010, Hall and Giaccia 2012) states that the biological effect of an N-fraction dose sequence $d = (d_1, \ldots, d_N)$ (in Gray (Gy)) delivered to a tumor volume is given by

$$\text{BED}_T = \sum_{t=1}^{N} d_t + \frac{1}{\alpha/\beta} \sum_{t=1}^{N} d_t^2, \tag{1}$$

which is a model governed by a single parameter, the α/β ratio, which signifies the fractionation sensitivity of the tumor tissue. The BED to the OAR can be described by

$$\text{BED}_{\text{OAR}} = \sum_{t=1}^{N} \sigma d_t + \frac{1}{\alpha/\beta} \sum_{t=1}^{N} \sigma d_t^2, \tag{2}$$

where σ is the generalized dose-sparing factor. For the maximum BED in the OAR volume, σ is the maximum of the individual dose sparing factors s_i. In order to describe a mean BED

5

constraint or dose-volume BED constraint other choices for σ can be used (Saberian et al. 2016a, Perkó et al. 2018).

For notational convenience, let τ (for tumor) and ρ (for risk) denote the inverse α/β ratio of the tumor and OAR volume, respectively. Mizuta et al. (2012) consider the problem of minimizing OAR BED subject to a lower bound $\text{BED}_{\text{T}}^{\text{pres}}$ on tumor BED. The number of fractions N is restricted to be at most $N^{\max}$. The problem reads

$$\min_{d,N} \quad \sigma \sum_{t=1}^{N} d_t + \rho \sigma^2 \sum_{t=1}^{N} d_t^2 \tag{3a}$$

$$\text{s.t.} \quad \sum_{t=1}^{N} d_t + \tau \sum_{t=1}^{N} d_t^2 \geq \text{BED}_{\text{T}}^{\text{pres}} \tag{3b}$$

$$d_1, \ldots, d_N \geq 0, N \in \{1, \ldots, N^{\max}\}. \tag{3c}$$

Let (d^*, N^*) denote the optimal solution to (3). A simple analysis in Mizuta et al. (2012) reveals the following important result:

$$N^* = \begin{cases} 1 & \text{if } \tau \geq \sigma\rho \\ N^{\max} \text{ and } d_1^* = \ldots = d_{N^{\max}}^* & \text{otherwise.} \end{cases} \tag{4}$$

In both cases the optimal dose d^* is such that (3b) is active. It can be shown that the same result holds if we maximize tumor BED subject to an upper bound on OAR BED (Bortfeld et al. 2015). A similar result has been derived for the case with multiple OARs (Saberian et al. 2016a). There is a large body of research that optimizes the number of treatment fractions for different model formulations (see Saberian et al. (2017) and references therein).

In the current paper, we restrict to one dose-limiting OAR. For many tumor sites, there is a single OAR that restricts the doses that can be delivered, and other OARs are much less restrictive. For example, for lung cancer the mean lung dose is an important indicator of toxicity. On the other hand, for head and neck cancer many OARs must be accounted for. We emphasize that other OARs are not completely disregarded. They are taken into account implicitly, because the original dose distribution was planned with all relevant OARs taken into consideration. Moreover, by restricting the minimum and maximum (mean target) dose per fraction, extreme deviations from the standard fractionation schedule are avoided, which is also designed whilst taking multiple OARs into account.

2.2 Adaptive fractionation using biomarkers

Most fractionation optimization methods assume the tumor and OAR fractionation sensitivity parameters τ and ρ are known exactly. There is much research on the α/β ratios for different tumor sites (Van Leeuwen et al. 2018) and OAR sites (Kehwar 2005), but they remain subject to considerable uncertainties. We assume box uncertainty of the form:

$$Z := \{(\rho, \tau) : \rho_L \leq \rho \leq \rho_U, \tau_L \leq \tau \leq \tau_U\}, \tag{5}$$

with $0 < \rho_L < \rho_U$ and $0 < \tau_L < \tau_U$. It is assumed that there is a nominal scenario, e.g., parameter values $\bar{\tau}$ and $\bar{\rho}$ derived from literature. There are two reasons for assuming a box uncertainty set. First, to the best of our knowledge there is little evidence that there is any correlation between the α/β ratios of target volumes and normal tissues. Second, box uncertainty leads to simpler models, which allow for a more detailed analysis of optimal fractionation decisions.

Ajdari and Ghate (2017) also consider a box uncertainty set, and determine a *robustly optimal* fractionation scheme, i.e., one that is that is feasible for all possible realizations and that is optimal for the worst-case realization. If biomarker information acquired during treatment

provides more accurate information on fractionation sensitivity than what was available at the start of the treatment, such a static RO approach may be overly conservative.

Somaiah et al. (2015) give an overview of various mechanisms for determining fractionation sensitivity. Using blood samples, one can quantify the involvement of non-homologous end-joining (NHEJ) and homologous recombination (HR) in cells. For details on how to measure these, we refer to Bindra et al. (2013) and Barker and Powell (2010), respectively. Change in the expression of epidermal growth factor receptor (EGFR) genes can also give some hints regarding the fractionation sensitivity (Somaiah et al. 2015), which can also be measured mid-treatment. Lastly, Somaiah et al. (2015) mention that there is a close link between proliferation index and hypoxia, both of which can be measured during RT using different PET tracers. We note that there is evidence that some of these mechanisms could be subject to change during RT, depending on, amongst others, the delivered dose, hypoxia, and immune system interaction. However, as a first study to adapting to inexact biomarker information, we make the assumption that fractionation sensitivity is static throughout treatment, i.e., there is a static 'true' (ρ, τ). In Section 3 we assume to observe (measure) the true (ρ, τ) exactly, and in Section 4 we assume to observe only an estimation/approximation $(\hat{\rho}, \hat{\tau})$.

The quality of the observed parameter estimates depends amongst others on the suitability of the biomarker itself, the measurement accuracy and when the biomarker measurement is taken during the treatment course. The relationship between the data quality and the moment of biomarker observation is complex, and it is impossible to exactly quantify this. For some biomarkers the data quality may greatly improve in the first few fractions, with a diminishing improvement in later fractions[1]. For others (e.g., functional imaging such as PET and magnetic resonance imaging (MRI)), the data quality is poor at the first couple of fractions and only increases substantially in later fractions[2]. In practice, some biomarkers, e.g., radiographic information, may also exhibit a decreasing data quality for very late observation moments (due to, for instance, interference from acute inflammation in the lung). Such behavior is rare, and as such not considered here. We impose a minimum dose per fraction, and make the assumption that biomarker data quality increases in the number of treatment fractions. In this way, the change in biomarker quality is influenced by both the dose delivered and the time passed. We will investigate several functional forms for this relationship in the numerical experiments.

2.3 Modeling choices

In order to establish a meaningful model for the adjustable robust optimization approach, we restrict the dose sequence $d = (d_1, \ldots, d_{N^{\max}})$ in several ways. In addition to a maximum number of fractions, we also set a minimum $N^{\min}$. Furthermore, we assume there is a single moment N_1 where we can adapt the treatment. Under the assumption that the uncertain parameters remain constant over time, more than one observation moment is not useful if the parameter is observed exactly. With inexact observations, there can be value in multiple observations, but given the patient burden (logistically and the delivery of additional dose) and financial cost of PET scans this is not considered here.

The dose per fraction in the first N_1 fractions is assumed to be the same, denote this by d_1. Variable N_2 denotes the number of fractions after adaptation; also these fractions have equal dose, denoted by d_2. In current clinical practice, uniform fractionated treatments are the standard. By restricting to only two different dose levels, extreme deviations from standard

---

[1]This is especially true in the case of certain blood biomarkers of innate immune status (such as interleukin (IL)-6 or tumor necrosis factor (TNF)-α) which are also the markers of inflammation. In these biomarkers, as the biomarker acquisition is shifted towards later in the RT course, the information regarding the immune status gets mixed with the RT-induced inflammation and loses its specificity to immune system condition.

[2]This is because the effect of RT on tissue is cumulative and is morphologically manifested only after a certain amount of dose (which depends on the underlying tissue threshold) is delivered.
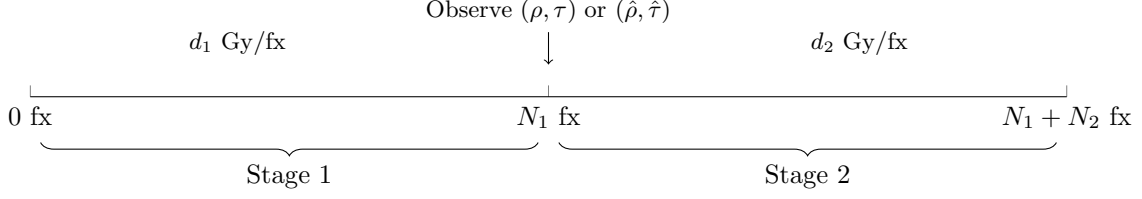
Figure 2: Schematic overview of the considered model. There are 3 variables: d_1, d_2 and N_2. First, we deliver N_1 fractions of dose d_1 per fraction. After this, we observe (ρ, τ) or $(\hat{\rho}, \hat{\tau})$. Subsequently, we deliver N_2 fractions of dose d_2 per fraction.

protocols are prevented. The above implies

$$N_2 \in \{N_2^{\min}, \ldots, N_2^{\max}\}, \tag{6}$$

with $N_2^{\min} = \max\{1, N^{\min} - N_1\}$ and $N_2^{\max} = N^{\max} - N_1$. We additionally set the constraint that $d_1, d_2 \geq d^{\min}$, for some predetermined value $d^{\min}$. Aside of preventing an unrealistically low dose per fraction, the minimum dose serves a modeling purpose for stage 1. As noted in Section 2.2, the biomarker quality can depend on both dose and time. The current model implicitly makes the assumption that an early response can only be observed via biomarkers once N_1 fractions of dose at least $d_{\min}$ have been delivered. Thus, in our models, this can be interpreted as a threshold. In the numerical experiments we investigate several temporal relationships between N_1 and biomarker quality. Lastly, we set a maximum dose per fraction $d_1^{\max}$ in stage 1, to avoid delivering dosages in stage 1 that severely restrict adaptation possibilities in stage 2. We will later impose some restrictions on the allowed combinations of $d^{\min}$, $d_1^{\max}$ and $N_2^{\max}$.

Figure 2 provides a schematic overview of the situation. We wish to maximize the tumor BED, subject to the constraint that the BED to the OAR is below the generalized tolerance level BED_{tol}, given by

$$\text{BED}_{\text{tol}}(\rho) = \varphi D \Big(1 + \frac{\varphi D}{T} \rho \Big), \tag{7}$$

i.e., the OAR is known to tolerate a total dose of D Gy if delivered in T fractions under dose shape factor φ. The dose shape factor is a parameter characterizing the spatial heterogeneity of a dose distribution, for more details see Saberian et al. (2016a) and Perkó et al. (2018). Note that $\text{BED}_{\text{tol}}(\rho)$ is a function of uncertain parameter ρ, the inverse α/β ratio of the OAR.

We emphasize that the model resulting from our modeling choices and assumptions does not directly represent a realistic decision-making problem in radiation therapy treatment planning. Nevertheless, it captures several important aspects of fractionation optimization. Moreover, using ARO on this stylized model, we gain insight into optimal decision rules and the role of uncertainty in adaptive fractionation optimization.

3 ARO: Biomarkers provide exact information

We present an adjustable robust optimization approach that optimally adjusts the remainder of the treatment once biomarker information has provided the true value of parameters τ and ρ. This serves as a stepping stone to the inexact data model.

3.1 Problem formulation

After delivering N_1 fractions we observe the true (ρ, τ). At the time that the stage-1 dose per fraction d_1 has to be decided, it is only known that (ρ, τ) belongs to uncertainty set Z. Put

together, the Exact Data Problem (EDP) reads:

$$\max_{d_1, d_2(\rho,\tau), N_2(\rho,\tau)} \quad \min_{(\rho,\tau) \in Z} N_1 d_1 + N_2(\rho,\tau) d_2(\rho,\tau) + \tau (N_1 d_1^2 + N_2(\rho,\tau) d_2(\rho,\tau)^2) \tag{8a}$$

$$\text{s.t.} \quad \sigma(N_1 d_1 + N_2(\rho,\tau) d_2(\rho,\tau)) + \rho \sigma^2 (N_1 d_1^2 + N_2(\rho,\tau) d_2(\rho,\tau)^2)$$
$$\leq \text{BED}_{\text{tol}}(\rho), \quad \forall (\rho,\tau) \in Z \tag{8b}$$

$$N_2(\rho,\tau) \in \{N_2^{\min}, \ldots, N_2^{\max}\}, \quad \forall (\rho,\tau) \in Z \tag{8c}$$

$$d_2(\rho,\tau) \geq d^{\min}, \quad \forall (\rho,\tau) \in Z \tag{8d}$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{8e}$$

The value for the stage-1 dose d_1 has to be decided before the value of (ρ,τ) is revealed; in ARO this is also commonly referred to as a *here-and-now* variable or decision. The values stage-2 dose d_2 and stage-2 number of fractions N_2 need to be decided only after (ρ,τ) is revealed as they may depend on the values of these parameters. Hence, they are written as functions $d_2(\rho,\tau)$ and $N_2(\rho,\tau)$ of the uncertain parameters (ρ,τ). In ARO such variables are also commonly referred to as *wait-and-see* variables or decisions. In this paper, we will adhere to the terms *stage 1* and *stage 2*, however.

Before we solve (8), we need some definitions and assumptions. The remaining BED tolerance level of the OAR, if N' fractions with dose d' have been administered, is given by

$$B(d', N'; \rho) = \text{BED}_{\text{tol}}(\rho) - \sigma d' N' - \rho \sigma^2 (d')^2 N'. \tag{9}$$

Subsequently, define the function

$$g(d', N', N''; \rho) := \frac{-1 + \sqrt{1 + \frac{4\rho}{N''} B(d', N'; \rho)}}{2\sigma\rho}. \tag{10}$$

The value of g can be interpreted as the maximum dose that can be delivered in N'' fractions if already N' fractions of dose d' are (scheduled to be) delivered. It is obtained by solving the equality version of (8b) for d_1 or d_2. Functions of this form will be used frequently throughout the remainder of this paper.

The following assumption on the relation between $d^{\min}$, $d_1^{\max}$ and the bounds on N_2 makes sure that for a given optimal number of fractions, it is feasible to deliver that number of fractions with minimum dose.

Assumption 1. *It holds that*

$$d^{min} \leq d_1^{max} \leq \min\left\{ g(d^{min}, N_2^{min}, N_1; \rho_L), g(d^{min}, N_2^{max}, N_1; \tfrac{\tau_L}{\sigma}), g(d^{min}, N_2^{max}, N_1; \rho_U) \right\}. \tag{11}$$

The particular form of the upper bound on $d_1^{\max}$ will become clear later. Numerical experiments indicate that results are not very sensitive to the choices of $d^{\min}$ and $d_1^{\max}$.

We continue by formally defining several properties of solutions. Let $X(\rho,\tau)$ denote the feasible region defined by constraints (8b)-(8e) for fixed (ρ,τ).

Definition 1 (Adjustable robustly feasible)**.** A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is adjustable robustly feasible ($\mathcal{ARF}$) to (8) if $(d_1, d_2(\rho,\tau), N_2(\rho,\tau)) \in X(\rho,\tau)$ for all $(\rho,\tau) \in Z$. ∎

Definition 2 (Adjustable robustly optimal)**.** A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is adjustable robustly optimal ($\mathcal{ARO}$) to (8) if it is $\mathcal{ARF}$ and there does not exist an $\mathcal{ARF}$ tuple $(\bar{d}_1, \bar{d}_2(\cdot), \bar{N}_2)$ such that

$$\min_{(\rho,\tau) \in Z} N_1 d_1 + N_2(\rho,\tau) d_2(\rho,\tau) + \tau (N_1 d_1^2 + N_2(\rho,\tau) d_2(\rho,\tau)^2)$$
$$< \min_{(\rho,\tau) \in Z} N_1 \bar{d}_1 + \bar{N}_2(\rho,\tau) \bar{d}_2(\rho,\tau) + \tau (N_1 \bar{d}_1^2 + \bar{N}_2(\rho,\tau) \bar{d}_2(\rho,\tau)^2). \tag{12}$$

∎

We also define the $\mathcal{ARO}$ property for the stage-1 decisions d_1 individually.

Definition 3 (Adjustable robustly optimal d_1). A stage-1 decision d_1 is $\mathcal{ARO}$ to (8) if there exist decision rules $d_2(\cdot)$ and $N_2(\cdot)$ such that $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{ARO}$ to (8). ∎

Lastly, we define optimality of a decision rule.

Definition 4 (Optimal decision rule). For a given d_1, a decision rule pair $(d_2(\cdot), N_2(\cdot))$ is optimal to (8) if $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{ARF}$ and for any $(\rho, \tau) \in Z$ it holds that

$$
\begin{aligned}
&N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau(N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \\
&\geq N_1 d_1 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau) + \tau(N_1 d_1^2 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau)^2),
\end{aligned}
\tag{13}
$$

for every $(\bar{d}_2(\cdot), \bar{N}_2(\cdot))$ such that $(d_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ is $\mathcal{ARF}$. ∎

The first observation we make in (8) is that any solution $(d_1, d_2(\cdot), N_2(\cdot))$ feasible for scenario $(\rho, \tau) \in Z$ is also feasible for $(\rho, \tau + \epsilon)$ with a higher objective value, if $\epsilon > 0$. Therefore, in any worst-case realization it will hold that $\tau = \tau_L$ (see (5)). This observation has consequences for what uncertainty sets Z need to be considered. Due to the result (4), one can in general distinguish three cases for uncertainty set Z and parameter σ:

Case 1) $\sigma \rho_U \leq \tau_L$: According to (4), for any realization (with $\tau = \tau_L$) it is optimal to deliver the minimum number of fractions in stage 2.

Case 2) $\sigma \rho_L \geq \tau_L$: According to (4), for any realization (with $\tau = \tau_L$) it is optimal to deliver the maximum number of fractions in stage 2.

Case 3) $\sigma \rho_L < \tau_L < \sigma \rho_U$: In the scenario (ρ_L, τ_L), it is optimal to deliver the maximum number of fractions in stage 2 according to (4). In the scenario (ρ_U, τ_L), it is optimal to deliver the minimum number of fractions in stage 2 according to (4).

In Cases 1 and 2, (8) is easily solved by plugging in the (worst-case) optimal value for N_2, and solving the resulting 2-variable optimization problem. Therefore, only Case 3 is of interest and in the remainder of this paper we make the following assumption.

Assumption 2. *It holds that $\sigma \rho_L < \tau_L < \sigma \rho_U$.*

In our numerical experiments, we use a lung cancer data set. Recent evidence suggests that for lung cancer Assumption 2 can indeed hold, i.e., the optimal number of treatment fractions is not always known prior to treatment. Further details are provided in Section 5.2. For other tumor sites, such as liver cancer, the tumor α/β is generally assumed to be 10 or higher, whereas the α/β of normal liver tissue is typically assumed to be 3 or 4. Thus, for such tumor sites Assumption 2 generally does not hold, and hyperfractionation is optimal.

3.2 Optimal decision rules and worst-case solution

Problem (8) is a 2-stage non-convex mixed-integer ARO problem, which are generally hard to solve. Nevertheless, due to the small number of variables the problem can be solved to optimality. In order to solve (8), we take two steps:

Step 1) Determine optimal decision rules $d_2(\cdot)$ and $N_2(\cdot)$ for fixed d_1.

Step 2) Plug in optimal decision rules and solve for d_1.
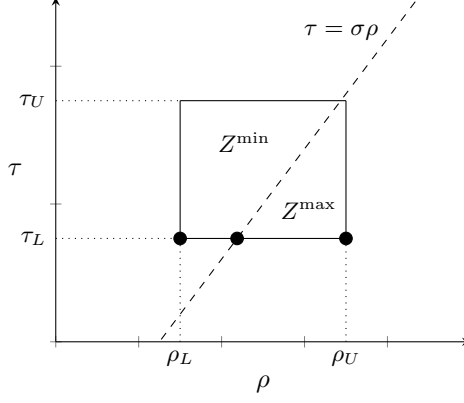
Figure 3: Split of uncertainty set Z according to (16). The circles indicate the locations of the candidate worst-case scenarios for (18).

In what follows, we give a detailed explanation of both steps. Let $(d_1^*, d_2^*(\rho, \tau), N_2^*(\rho, \tau))$ denote an $\mathcal{ARO}$ solution to (8).

Step 1: Determine optimal decision rules $d_2(\cdot)$ and $N_2(\cdot)$ for fixed d_1

Fix stage-1 variable d_1. Similar to the result (4), we will show that it is optimal to deliver either the minimum or the maximum number of fractions in stage-2. Moreover, (8b) is the only OAR dose-limiting constraint, so it will hold with equality if this does not violate variable bounds (8d) and (8e). We will show that the latter is not the case. The theorem below summarizes the result.

Theorem 1. *Let d_1 be the stage-1 decision of* (8). *The decision rules*

$$N_2^*(\rho, \tau) = \begin{cases} N_2^{min} & \text{if } \tau \geq \sigma\rho \\ N_2^{max} & \text{otherwise,} \end{cases} \tag{14}$$

and

$$d_2^*(d_1; \rho, \tau) = \begin{cases} g(d_1, N_1, N_2^{min}; \rho) & \text{if } \tau \geq \sigma\rho \\ g(d_1, N_1, N_2^{max}; \rho) & \text{otherwise} \end{cases} \tag{15}$$

are optimal to (8) *for the given d_1. These provide the unique optimal decisions unless $\tau = \sigma\rho$.*

Proof. See Appendix B.1. $\qquad\qquad\square$

Clearly, these decision rules are non-linear, and in fact split the uncertainty region in two parts: one where it is optimal to deliver the minimum number of fractions $N_2^{\min}$ in stage 2, and one where it is optimal to deliver the maximum number of fractions $N_2^{\max}$ in stage 2. This suggests splitting the uncertainty set as follows:

$$Z^{\min} := Z \cap \{\tau \geq \sigma\rho\} \tag{16a}$$
$$Z^{\max} := Z \cap \{\tau < \sigma\rho\}. \tag{16b}$$

An illustration is provided in Figure 3.

Step 2: Plug in optimal decision rules and solve for d_1

In order to find an $\mathcal{ARO}$ d_1, we introduce the following objective function for given (ρ, τ):

$$f(d_1, N_2; \rho, \tau) := \begin{cases} N_1 d_1 + N_2 g(d_1, N_1, N_2; \rho) \\ +\tau\big(N_1 d_1^2 + N_2 g(d_1, N_1, N_2; \rho)^2\big) & \text{if } d_1 \in [0, g(0, 0, N_1; \rho)] \\ -\infty & \text{otherwise,} \end{cases} \tag{17}$$

11

where, for given ρ, the value $g(0, 0, N_1; \rho)$ is the maximum dose that can be delivered in stage 1 due to the nonnegativity restriction on the stage-2 dose. From Assumption 1 it follows that $g(0, 0, N_1; \rho) \geq d_1^{\max}$ for all $(\rho, \tau) \in Z$, so f is finite for all feasible d_1. According to Lemma C.1 in Appendix C, function f is either convex, concave or constant in d_1.

Plugging in (14) and (15) and using definition (17) allows us to rewrite (8) to a problem of only variable d_1:

$$\max_{d_1} \quad \min_{(\rho, \tau) \in Z} \quad f(d_1, N_2^*(\rho, \tau), \rho, \tau) \tag{18a}$$

$$\text{s.t.} \quad d^{\min} \leq d_1 \leq d_1^{\max}. \tag{18b}$$

As noted in Section 3.1, in any worst-case realization it will hold that $\tau = \tau_L$, so it is sufficient to consider only those observations with $\tau = \tau_L$.

In order to reformulate (18), we make use of the properties of g and f in Lemma C.3 in Appendix C. In particular, Lemma C.3b states that function f is either increasing or decreasing in ρ for fixed d_1. Hence, if we move (18a) to a constraint and split according to (16), for both $Z^{\min}$ and $Z^{\max}$ it is sufficient to consider the constraint for the highest and lowest value of ρ in the uncertainty set. With $\tau = \tau_L$, this yields the scenarios (ρ_L, τ_L) and $(\frac{\tau_L}{\sigma}, \tau_L)$ for $Z^{\min}$ and $(\frac{\tau_L}{\sigma}, \tau_L)$ and (ρ_U, τ_L) for $Z^{\max}$.

Therefore, the three candidate worst-case scenarios are (ρ_L, τ_L), (ρ_U, τ_L) and $(\frac{\tau_L}{\sigma}, \tau_L)$; their locations are indicated in Figure 3. By Lemma C.1, the objective value in the third scenario is equal to K, with

$$K = \frac{1}{\sigma} B\left(0, 0, \frac{\tau_L}{\sigma}\right). \tag{19}$$

This is the maximum target BED that can be attained if radiation sensitivity parameters are exactly such that fractionation has no influence on the optimal objective value. It is equal to the maximum tolerable OAR BED for these radiation sensitivity parameters, divided by the generalized OAR dose sparing factor σ.

Putting everything together, we conclude that if $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is $\mathcal{ARO}$ to the EDP (8) then there exists a $q^* \in \mathbb{R}_+$ such that (d_1^*, q^*) is an optimal solution to

$$\max_{d_1, q} \quad q \tag{20a}$$

$$\text{s.t.} \quad q \leq f(d_1, N_2^{\min}; \rho_L, \tau_L) \tag{20b}$$

$$q \leq f(d_1, N_2^{\max}; \rho_U, \tau_L) \tag{20c}$$

$$q \leq K \tag{20d}$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{20e}$$

Conversely, if (d_1^*, q^*) is an optimal solution to (20) and $N_2^*(\cdot)$ and $d_2^*(\cdot)$ are given by (14) and (15), respectively, then $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is $\mathcal{ARO}$ to (8). Hence, (20) and EDP (8) are equivalent.

According to Lemma C.1, the RHS of (20b) and (20c) is convex and concave in d_1, respectively. Hence, (20) asks to find the value of d_1 that maximizes the minimum of a univariate convex (20b), concave (20c) and constant (20d) function on a closed interval (20e). Lemma C.2 in Appendix C provides information on the intersection points of the functions (20b)-(20d). Consequently, the optimal solution(s) to (20) is/are easily found.

Figure 4 illustrates a possible instance of (20), displaying constraints (20b)-(20d). In this case, the set of optimal solutions is the union of the two intervals for d_1 where constraint (20d) is active. This is indicated in red. Dose constraints (20e) may cut off part of these intervals. If due to constraint (20e) both these intervals are infeasible, the optimum is at one of the boundaries for d_1.
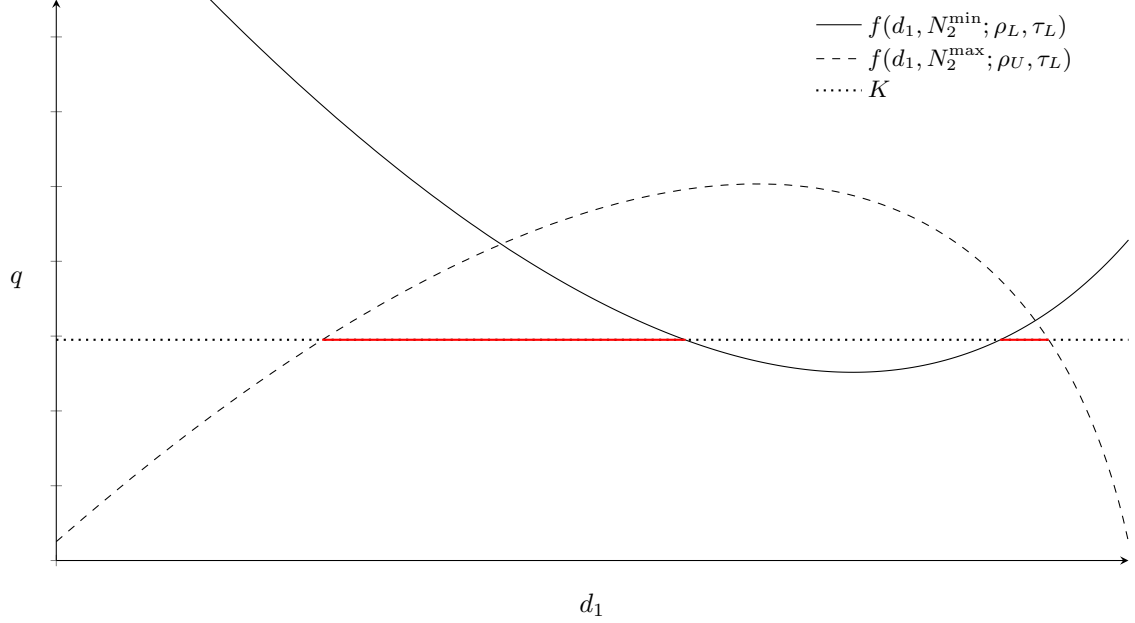
Figure 4: Schematic illustration of (20). The solid and dashed curves represent constraints (20b) and (20c), respectively, and the dotted line represent constraint (20d). The optimal intervals are indicated in red.

3.3 Pareto adjustable robustly optimal solutions

Figure 4 illustrates that it is possible that there are multiple optimal solutions to (20). These solutions are $\mathcal{ARO}$ stage-1 solutions to the EDP (8). In general, in case there are multiple $\mathcal{ARO}$ solutions these may perform vastly different if a non-worst-case scenario realizes (De Ruiter et al. 2016). Iancu and Trichakis (2014) studies static robust optimization problems with multiple robustly optimal solutions, and introduce the concept of Pareto robustly optimal ($\mathcal{PRO}$) solutions. A robustly optimal solution is called $\mathcal{PRO}$ if there is no other robustly feasible solution that has equal or better objective value for all scenarios in the uncertainty set, while being strictly better for at least one scenario. Non-$\mathcal{PRO}$ solutions are dominated by at least one $\mathcal{PRO}$ solution and are therefore not desired[3]. The concept has previously been applied to breast cancer treatment planning by Mahmoudzadeh (2015).

Iancu and Trichakis (2014) study $\mathcal{PRO}$ solutions solely for static RO problems; we generalize the concept to 2-stage adjustable robust optimization problems.

Definition 5 (Pareto adjustable robustly optimal). An $\mathcal{ARO}$ tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is Pareto adjustable robustly optimal ($\mathcal{PARO}$) to (8) if there is no tuple $(\bar{d}_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ that is $\mathcal{ARO}$ to (8) and satisfies

$$
\begin{aligned}
N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau(N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \\
\leq N_1 \bar{d}_1 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau) + \tau(N_1 \bar{d}_1^2 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau)^2) \quad \forall (\rho, \tau) \in Z
\end{aligned}
\tag{21a}
$$

$$
\begin{aligned}
N_1 d_1 + N_2(\bar{\rho}, \bar{\tau}) d_2(\bar{\rho}, \bar{\tau}) + \tau(N_1 d_1^2 + N_2(\bar{\rho}, \bar{\tau}) d_2(\bar{\rho}, \bar{\tau})^2) \\
< N_1 \bar{d}_1 + \bar{N}_2(\bar{\rho}, \bar{\tau}) \bar{d}_2(\bar{\rho}, \bar{\tau}) + \bar{\tau}(N_1 \bar{d}_1^2 + \bar{N}_2(\bar{\rho}, \bar{\tau}) \bar{d}_2(\bar{\rho}, \bar{\tau})^2) \quad \text{for some } (\bar{\rho}, \bar{\tau}) \in Z.
\end{aligned}
\tag{21b}
$$

■

We also define the concept $\mathcal{PARO}$ for the stage-1 decision d_1 individually.

---

[3]The concept of Pareto robust optimality closely resembles the concept of Pareto efficiency in multi-criteria optimization (MCO). In MCO, Pareto efficient solutions can only be improved in one criteria at the cost of a deterioration in another criteria. Only Pareto efficient solutions are of interest, and the overall goal in MCO is to compute this set of solutions (known as the Pareto surface).

Definition 6 (Pareto adjustable robustly optimal d_1)**.** A stage-1 decision d_1 is $\mathcal{PARO}$ to (8) if there exist decision rules $N_2(\cdot)$ and $d_2(\cdot)$ such that $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{PARO}$ to (8). ∎

If there are multiple $\mathcal{ARO}$ solutions, we wish to pick one that is $\mathcal{PARO}$. In general, finding PARO solutions is hard, because it requires comparing the performance of both stage-1 decisions and stage-2 decision rules on multiple scenarios simultaneously. However, for the current problem (14) and (15) are optimal decision rules. Plugging these in conditions (21) reduces the problem of finding a PARO solution to solely comparing the performance of $\mathcal{ARO}$ stage-1 decisions d_1 in non-worst-case scenarios.

In Iancu and Trichakis (2014) it is shown for linear optimization that, if we optimize over the robustly optimal solutions for a second criterion (scenario) that is in the relative interior of the uncertainty set, the resulting solution(s) are $\mathcal{PRO}$. In a similar fashion $\mathcal{PARO}$ solutions to the current problem can be found. Let X^{ARO} denote the set of $\mathcal{ARO}$ stage-1 solutions to (8). It turns out that consecutively optimizing over an auxiliary scenario where hyperfractionation is optimal and an auxiliary scenario where hypofractionation is optimal yields a set of $\mathcal{PARO}$ solutions. Let $(\rho^{\mathrm{aux\text{-}min}}, \tau^{\mathrm{aux\text{-}min}}) \in \mathrm{int}(Z^{\mathrm{min}})$, where $\mathrm{int}(\cdot)$ is the interior operator. Define the auxiliary optimization problem for the hypofractionation scenario:

$$\max_{d_1 \in X^{\mathrm{ARO}}} f\big(d_1, N_2^{\mathrm{min}}; \rho^{\mathrm{aux\text{-}min}}, \tau^{\mathrm{aux\text{-}min}}\big). \tag{22}$$

Denote the set of optimal solutions to (22) by $X^{\mathrm{aux\text{-}min}}$. Similarly, let $(\rho^{\mathrm{aux\text{-}max}}, \tau^{\mathrm{aux\text{-}max}}) \in \mathrm{int}(Z^{\mathrm{max}})$. Define the auxiliary optimization problem for the hyperfractionation scenario:

$$\max_{d_1 \in X^{\mathrm{aux\text{-}min}}} f\big(d_1, N_2^{\mathrm{max}}; \rho^{\mathrm{aux\text{-}max}}, \tau^{\mathrm{aux\text{-}max}}\big). \tag{23}$$

Note that it uses $X^{\mathrm{aux\text{-}min}}$ as input, i.e., we solve the auxiliary problems consecutively. Denote the set of optimal solutions to (23) by X^{PARO}.

Theorem 2. *All solutions in X^{PARO} are $\mathcal{PARO}$ to* (8).

Proof. See Appendix B.2. □

Solving (22) or (23) entails maximizing a strictly convex or strictly concave function over a feasible set consisting of a small number of intervals or points. Hence, these auxiliary problems are easily solved. Note that the second auxiliary problem is only relevant if the first auxiliary problem has multiple optimal solutions. Switching their order, and optimizing (22) over the set $X^{\mathrm{aux\text{-}max}}$ may lead to different solutions, and these are also $\mathcal{PARO}$. Thus, in general X^{PARO} does not contain all $\mathcal{PARO}$ solutions. The two-step approach is necessary; numerical results show that optimizing over only one auxiliary scenario may indeed yield non-$\mathcal{PARO}$ solutions.

4 ARO: Biomarkers provide inexact information

In this section we present an adjustable robust optimization approach to solve a more realistic version of the adaptive treatment-length problem. Because in practice it is impossible to exactly determine the α/β parameters from biomarker data, any values for the α/β parameters obtained during treatment are inexact. This section presents a model that accounts for uncertainty in biomarker information.

4.1 Problem formulation

The setup for the ARO problem with inexact data is based on De Ruiter et al. (2017). After N_1 fractions we obtain an estimate $(\hat{\rho}, \hat{\tau})$ for (ρ, τ), the inverse α/β parameters for the OAR and

the tumor. It is still assumed that Assumption 2 holds for uncertainty set Z. Furthermore, we assume that (ρ, τ), $(\hat{\rho}, \hat{\tau}) \in Z$ (as defined in (5)) and that $(\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}$, with

$$\hat{Z} = \{(\varepsilon^\rho, \varepsilon^\tau) \in \mathbb{R}^2 : |\varepsilon^\rho| \leq r^\rho, |\varepsilon^\tau| \leq r^\tau\}. \tag{24}$$

Here r^ρ and r^τ are parameters that define the accuracy of the estimate/observation $(\hat{\rho}, \hat{\tau})$. Set $\hat{Z}$ is the uncertainty set around the inexact observation. This can also be written as $(\rho, \tau) \in \{(\hat{\rho}, \hat{\tau})\} + \hat{Z}$, which is the Minkowski sum of a singleton and a set. This new set need not be contained in the original uncertainty set Z. Define

$$U := \{(\rho, \tau, \hat{\rho}, \hat{\tau}) : (\rho, \tau), \ (\hat{\rho}, \hat{\tau}) \in Z, \ (\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}\}, \tag{25}$$

and

$$Z_{(\hat{\rho}, \hat{\tau})} := (\{(\hat{\rho}, \hat{\tau})\} + \hat{Z}) \cap Z. \tag{26}$$

The set U contains all possible observation-realization pairs and $Z_{(\hat{\rho}, \hat{\tau})}$ contains all possible realizations after observation of $(\hat{\rho}, \hat{\tau})$. For given observation $(\hat{\rho}, \hat{\tau})$, the new upper and lower bounds for (ρ, τ) are given by

$$\hat{\tau}_L = \max\{\tau_L, \hat{\tau} - r^\tau\}, \ \ \hat{\tau}_U = \min\{\tau_U, \hat{\tau} + r^\tau\} \tag{27a}$$

$$\hat{\rho}_L = \max\{\rho_L, \hat{\rho} - r^\rho\}, \ \ \hat{\rho}_U = \min\{\rho_U, \hat{\rho} + r^\rho\}. \tag{27b}$$

Compared to Section 3, we remove Assumption 1 and impose a different (slightly stricter) assumption on the relation between $d^{\min}$, $d_1^{\max}$ and the bounds on N_2.

Assumption 3. *It holds that* $d^{min} \leq d_1^{max}$ *and*

$$d_1^{max} \leq \min\left\{g(d_1^{min}, N_2^{min}, N_1; \rho_L), g(d_1^{min}, N_2^{max}, N_1; \max\{\rho_L, \tfrac{\tau_L}{\sigma} - 2r^\rho\}), g(d_1^{min}, N_2^{max}, N_1; \rho_U)\right\}. \tag{28}$$

The inexact data problem (IDP) analogous to (8) is given by

$$\max_{d_1, d_2(\hat{\rho}, \hat{\tau}), N_2(\hat{\rho}, \hat{\tau})} \ \min_{(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U} N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2), \tag{29a}$$

$$\begin{aligned}\text{s.t.} \quad &\sigma(N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})) + \rho\sigma^2(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2) \\ &\leq \text{BED}_{\text{tol}}(\rho), \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U\end{aligned} \tag{29b}$$

$$N_2(\hat{\rho}, \hat{\tau}) \in \{N_2^{\min}, \dots, N_2^{\max}\}, \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U \tag{29c}$$

$$d_2(\hat{\rho}, \hat{\tau}) \geq d^{\min}, \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U \tag{29d}$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{29e}$$

For stage-2 variables $d_2(\hat{\rho}, \hat{\tau})$ and $N_2(\hat{\rho}, \hat{\tau})$ it is indicated that they are a function of the observations $(\hat{\rho}, \hat{\tau})$ instead of the uncertain parameters (ρ, τ). Similar to Section 3, we formally define several properties of solutions. Let $X(\rho, \tau, \hat{\rho}, \hat{\tau})$ denote the feasible region defined by constraints (29b)-(29e) for fixed $(\rho, \tau, \hat{\rho}, \hat{\tau})$.

Definition 7 (Adjustable robust feasibility)**.** A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is adjustable robustly feasible ($\mathcal{ARF}$) to (29) if $(d_1, d_2(\hat{\rho}, \hat{\tau}), N_2(\hat{\rho}, \hat{\tau})) \in X(\rho, \tau, \hat{\rho}, \hat{\tau})$ for all $(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U$. ∎

Optimality of a decision rule is defined as follows.

Definition 8 (Optimal decision rule)**.** For a given d_1, a decision rule pair $(d_2(\cdot), N_2(\cdot))$ is optimal to (29) if $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{ARF}$ and for any $(\hat{\rho}, \hat{\tau}) \in Z$ it holds that

$$\begin{aligned}&\min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2) \\ &\geq \min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} N_1 d_1 + \bar{N}_2(\hat{\rho}, \hat{\tau}) \bar{d}_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + \bar{N}_2(\hat{\rho}, \hat{\tau}) \bar{d}_2(\hat{\rho}, \hat{\tau})^2)\end{aligned} \tag{30}$$

for every $(\bar{d}_2(\cdot), \bar{N}_2(\cdot))$ such that $(d_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ is $\mathcal{ARF}$. ∎

Note that for exact data, an optimal decision rule yields the optimal decision for any *realization* in the uncertainty set Z (given d_1). For inexact data, we call a decision rule optimal if it yields the maximum worst-case (guaranteed) objective value for any *observation* in the uncertainty set Z.

4.2 Optimal decision rules and conservative approximation

Depending on both the observed $(\hat{\rho}, \hat{\tau})$ and the quality of the biomarker information (i.e., r^ρ and r^τ), we may be able to immediately determine the optimal value for N_2. Therefore, we split the uncertainty set for the observations $(\hat{\rho}, \hat{\tau})$. Define

$$Z_{\text{ID}}^{\min} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \geq \sigma \hat{\rho}_U\} \tag{31a}$$

$$Z_{\text{ID}}^{\text{int}} = \{(\hat{\rho}, \hat{\tau}) \in Z : \sigma \hat{\rho}_L < \hat{\tau}_L < \sigma \hat{\rho}_U\} \tag{31b}$$

$$Z_{\text{ID}}^{\max} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \leq \sigma \hat{\rho}_L\}, \tag{31c}$$

so that $Z = Z_{\text{ID}}^{\min} \cup Z_{\text{ID}}^{\text{int}} \cup Z_{\text{ID}}^{\max}$. Figure 5 provides an illustration. The split is such that if $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\min}$ or $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\max}$ only $N_2^{\min}$ resp. $N_2^{\max}$ fractions can be optimal in stage 2. Subset $Z_{\text{ID}}^{\text{int}}$ is the area between the dash-dotted lines. If $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$ both $N_2^{\min}$ and $N_2^{\max}$ fractions in stage 2 may be optimal for the true (ρ, τ). The following theorem states the optimal stage-2 decision rules for a given value of d_1.

Theorem 3. *Let d_1 be the stage-1 decision of* (29). *The decision rules*

$$N_2^*(d_1; \hat{\rho}, \hat{\tau}) = \begin{cases} N_2^{min} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{min} \\ \displaystyle\operatorname*{arg\,max}_{N_2 \in \{N_2^{min}, \dots, N_2^{max}\}} \min\{f(d_1, N_2; \hat{\rho}_L, \hat{\tau}_L), f(d_1, N_2; \hat{\rho}_U, \hat{\tau}_L)\} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{int} \\ N_2^{max} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{max}, \end{cases} \tag{32}$$

and

$$d_2^*(d_1; \hat{\rho}, \hat{\tau}) = \min\{g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_L), g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_U)\}, \tag{33}$$

are optimal to (29) *for the given d_1.*

Proof. See Appendix B.3. $\qquad\qquad\square$

The worst-case optimal decision rule (32) may yield a value unequal to $N_2^{\min}$ and $N_2^{\max}$ if $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$. If r^ρ and r^τ are zero, i.e., we have exact data, then it holds that $\hat{\tau}_L = \tau$ and $\hat{\rho}_L = \hat{\rho}_U = \rho$. Hence, the two functions f in the RHS of (32) are equal, and the optimal N_2^* is the one that maximizes the resulting function. One can verify that this does not depend on d_1. Hence, in case of exact data Theorem 3 reduces to Theorem 1.

It turns out that, after plugging in (32) and (33), and splitting the uncertainty set according to (31), it is not apparent how to determine the optimal stage-1 decision d_1^* for (29). In Appendix B.4 the following lower bound problem to (29) is derived, named the Approximate Inexact Data Problem (AIDP):

$$\max_{d_1, q} \quad q \tag{34a}$$

$$\text{s.t.} \quad q \leq f(d_1, N_2^{\min}; \rho_L, \tau_L) \tag{34b}$$

$$q \leq f(d_1, N_2^{\max}; \rho_U, \tau_L) \tag{34c}$$

$$q \leq K \tag{34d}$$

$$q \leq p(d_1) \tag{34e}$$
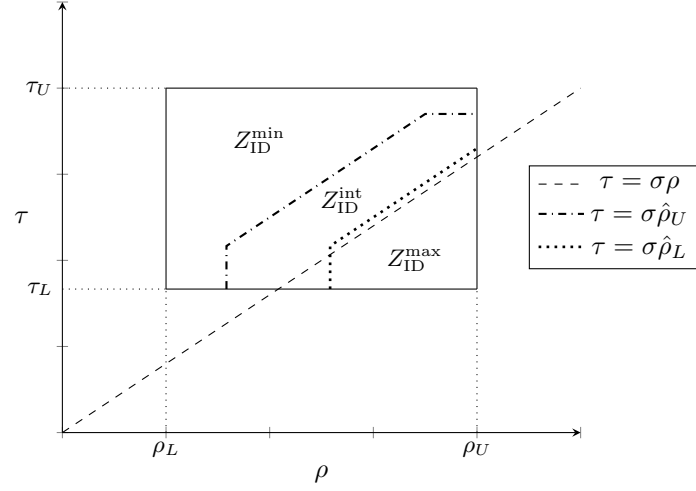
$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{34f}$$

Figure 5: The uncertainty set Z (solid lines) for $(\hat{\rho}, \hat{\tau})$ is split into $Z_{\text{ID}}^{\min}$, $Z_{\text{ID}}^{\text{int}}$, $Z_{\text{ID}}^{\max}$, according to (31). Subset Z^{int} is the area between the dotted and dash-dotted curves. If $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$ both $N_2^{\min}$ and $N_2^{\max}$ fractions in stage 2 may be optimal for the true (ρ, τ). If $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\min}$ or $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\max}$ only $N_2^{\min}$ resp. $N_2^{\max}$ fractions can be optimal in stage 2.
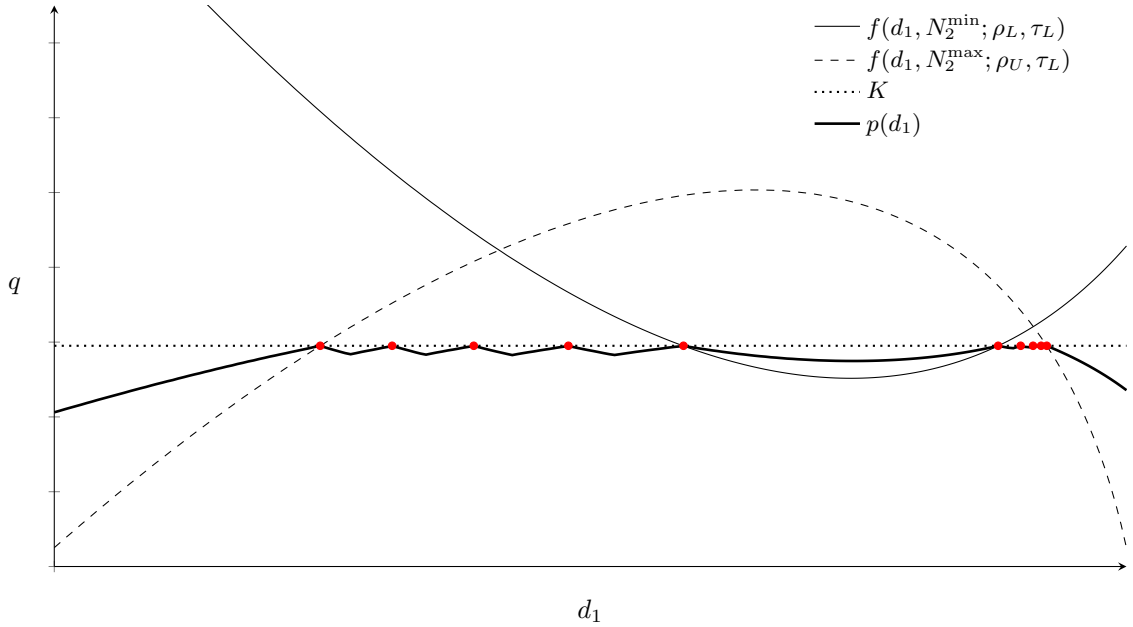

Figure 6: Schematic illustration of (34). Compared to the case with exact data (Figure 4), the thick black curve (constraint (34e)) is extra. The solid, dashed and dotted lines/curves represent constraints (34b), (34c) and (34d), respectively. Optimal solutions are indicated by red circles.

17

The AIDP is best explained using an example. Figure 6 illustrates a possible instance of (34), displaying constraints (34b)-(34e). Compared to (20) for exact biomarker information (see Figure 4), problem (34) has the added constraint (34e); a piecewise convex-concave function $p(d_1)$ defined by (C.32) in Lemma C.5 (Appendix C). It can be interpreted as follows. If $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$, the optimal number of stage-2 fractions can be inbetween $N_2^{\min}$ and $N_2^{\max}$, as shown in Theorem 3. In those cases, the optimal number of fractions also depends on the already delivered stage-1 dose d_1. The upper kinks (black dots) in the piecewise convex-concave function $p(d_1)$ in Figure 6 indicate values of d_1 where the worst-case scenario changes. The lower kinks indicate values where the optimal number of stage-2 fractions changes. The exact expression for $p(d_1)$ does not provide additional insight and is therefore omitted here.

In Figure 6, optimal solutions are locations where (34d) and (34e) are both binding, indicated by red circles. Dose constraints (34f) may cut off (some of) these points. If due to constraint (34f) none of the circles are feasible, the optimum is at one of the boundaries for d_1. Constraint (34e) is the only conservative constraint in (34). Hence, only if the feasible values for d_1 are such that none of the circles in Figure 6 are feasible and constraint (34e) is binding, it is possible that the optimal objective value of (34) is strictly worse than that of (29).

Lemmas C.1, C.2 and C.5 in Appendix C provide information on the shape and intersection points of constraint functions (34b)-(34e). Consequently, the optimal solution(s) of (29) is/are easily obtained. If (d_1^*, q^*) is optimal to AIDP (34), and $N_2^*(\cdot)$ and $d_2^*(\cdot)$ are given by (32) and (33) then $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is $\mathcal{ARF}$ to the original IDP (29). It is $\mathcal{ARF}$ because AIDP is a conservative approximation of IDP.

4.3 Pareto robustly optimal solutions to conservative approximation

Figure 6 also illustrates that it is possible that there are multiple optimal solutions to the AIDP (34). Because the AIDP provides a conservative approximation to IDP (29), optimizing over auxiliary scenario(s) as in Section 3 does not necessarily yield a stage-1 decision d_1 that is $\mathcal{PARO}$ to the original IDP. It turns out that a $\mathcal{PRO}$ solution to AIDP is obtained from the set of robustly optimal solutions to AIDP if we consecutively optimize for two auxiliary observations such that any worst-case realization is in the interior of set $Z^{\min}$ resp. $Z^{\max}$. Two important remarks are in place here. First, a $\mathcal{PRO}$ solution to AIDP need not be a $\mathcal{PARO}$ solution to IDP, even if it is $\mathcal{ARO}$ to IDP. Second, the required auxiliary scenarios need not exist; their existence depends on the values of r^{ρ} and r^{τ}. Hence, further details are omitted.

5 Numerical results

This section presents numerical results of the methods presented in Sections 3 and 4. First, Section 5.1 describes the benchmark methods against which we compare the ARO method for EDP and IDP, and Section 5.2 describes the setup of the numerical experiments.

5.1 Benchmark static and folding horizon methods

We analyze the performance of the static and folding horizon nominal method (NOM and NOM-FH), the static and folding horizon robust optimization method (RO and RO-FH) and the adjustable robust optimization method (ARO). In the folding horizon approaches only the stage-1 decisions are implemented, and the model is re-optimized for the second stage once the biomarker information is revealed.

The static method NOM optimizes for the nominal parameter values $(\bar{\rho}, \bar{\tau})$ and disregards any uncertainty and adaptability. This method is the same for both EDP and IDP. In stage 2, NOM-FH solves the nominal problem under the assumption that the obtained biomarker estimate is exact (which is an invalid assumption for IDP). This method does not guarantee robustly feasible solution (feasible for all $(\rho, \tau) \in Z$) nor a robustly optimal solution ($\mathcal{RO}$;

Problem	Method				
	NOM	NOM-FH	RO	RO-FH	ARO
EDP	-	-	$\mathcal{RO}$	$\mathcal{RO}$	$\mathcal{PARO}$
IDP	-	-	$\mathcal{RO}$	$\mathcal{RO}$	$\mathcal{ARF}$

Table 1: Guaranteed solution properties of the five methods.

(static) optimal for the worst-case realization $(\rho, \tau) \in Z$). The static method RO optimizes for the worst-case realization of (ρ, τ) in the uncertainty set Z, and disregards adaptability. For EDP the method RO-FH solves the same nominal problem as NOM-FH in stage 2; for IDP it solves a static robust optimization problem in stage 2, for which the uncertainty set is determined by the accuracy of the biomarker information. RO and RO-FH both guarantee an $\mathcal{RO}$ solution.

One may add a folding horizon component to ARO (for either EDP or AIDP). This may improve the results in case a suboptimal stage-2 decision rule is used. However, as shown in Sections 3.2 and 4.2, the used stage-2 decision rules are optimal for any realized scenario (and for given stage-1 decision d_1 in case of inexact information). Hence, adding a folding horizon component will not change results.

Table 1 provides an overview of the guaranteed solution properties of the methods. It is important to note that in case of inexact biomarker information (IDP) the methods RO and RO-FH guarantee an $\mathcal{RO}$ solution, whereas ARO guarantees only an $\mathcal{ARF}$ solution via solving the approximate problem AIDP. Depending on the approximation quality, the $\mathcal{ARF}$ solution may be close or equal to an $\mathcal{ARO}$ solution. Next to these five methods, we also report the results for the perfect information optimum (PI). This is the attainable optimum if from the start of the first fraction the true (ρ, τ) is exactly known. It can be formulated by taking the nominal problem and replacing the nominal parameter values by their true values. While in practice not a viable strategy, PI provides information on the value of perfect information, and allows us to put the performance of and differences between the other methods in perspective.

5.2 Study setup

We use a data set of 30 non-small cell lung cancer (NSCLC) patients, treated with either photon or proton therapy. The mathematical models in Sections 3 and 4 are based on the assumption that there is a single dose restricting OAR. We assume that the single dose restricting OAR is the normal lung itself[4]. For the models in Sections 3 and 4, an instance is defined by a tuple $(\sigma, \varphi, D, T, N_1, N^{\min}, N^{\max}, d^{\min}, d_1^{\max})$ and the relevant uncertainty sets.

Clinically, the number of treatment fractions varied between 33 and 37 fractions, with the majority of patients receiving 37 fractions. We set $N^{\min} = 30$ and $N^{\max} = 40$, to allow for slight deviations from the clinical standard. We assume the biomarker acquisition is made once $N_1 = 10$ fractions have been administered. This implies $N_2^{\min} = 20$ and $N_2^{\max} = 30$. Mean lung dose tolerance is $D = 20$Gy, and we set $T = 37$ as that is the clinically standard regimen. The patients differ in (σ, φ), which characterize the spatial dose distribution. Using the clinically delivered dose distribution, we derive for each normal lung voxel its dose sparing factor s_i (see Section 2.1). The dose shape factor φ and the generalized dose sparing factor σ for mean OAR

---

[4]This is in line with clinical practice wherein normal lung is treated as the most important normal tissue and the treatment is designed as to minimize the radiation exposure to normal lung.

BED are given by

$$\varphi = \frac{n \sum_{i=1}^n s_i^2}{\left(\sum_{i=1}^n s_i\right)^2}, \tag{35a}$$

$$\sigma = \frac{\sum_{i=1}^n s_i^2}{\sum_{i=1}^n s_i}, \tag{35b}$$

see Perkó et al. (2018) for details.

Cox (1986) estimate normal lung tissue α/β to be between 2.4 and 6.3. We set the nominal value at the midpoint 4.35. The α/β of NSCLC lung tumors has traditionally been assumed to be above 10 Gy. However, recent NSCLC hypofractionation trials show promising results, indicating that NSCNC cells are more sensitive to fraction size than previously assumed, i.e., have a lower α/β than 10. Santiago et al. (2016) find values between 2.2 and 9.0. We set the nominal value at the midpoint 5.6. Put together, we get the following uncertainty set for the inverse α/β ratios:

$$Z = \{(\rho, \tau) : 1/6.3 \le \rho \le 1/2.4, \ 1/9.0 \le \tau \le 1/2.2\}, \tag{36}$$

and the nominal scenario is $(\bar{\rho}, \bar{\tau}) = (1/4.35, 1/5.6)$. With this uncertainty set, 20 out of 30 patient cases satisfy Assumption 2: these are used in the numerical experiments. For the remaining ten patients the optimal number of treatment fractions can be determined prior to treatment, so these are removed.

To discriminate between multiple $\mathcal{ARO}$ solutions, we follow the procedure detailed in Section 3.3 in the case of exact biomarker information. The auxiliary scenarios are sampled uniformly from $\text{int}(Z^{\min})$ and $\text{int}(Z^{\max})$. In the case of inexact biomarker information, the procedure discussed in Section 4.3 is followed if the required auxiliary observations exist. If such observations exist, we sample uniformly from Z until we have found two auxiliary observations for which any worst-case realization is in $\text{int}(Z^{\min})$ resp. $\text{int}(Z^{\max})$. If such observations do not exist, the robustly optimal solution to AIDP with lowest stage-1 dose is selected. The method RO (and therefore also RO-FH) may also find multiple robustly optimal solutions. For the obtained set of robustly optimal solutions we again follow the procedure detailed in Section 3.3. It turns out that for RO, the robustly optimal solutions often perform identical in non-worst-case scenarios. We optimize over the auxiliary scenarios consecutively; the first auxiliary scenario is the scenario corresponding to $\text{int}(Z^{\min})$.

The minimum dose per fraction is $d^{\min} = 1.5$ Gy and the maximum stage-1 dose per fraction is $d_1^{\max} = 3$. This satisfies Assumption 1 (for EDP) and 3 (for IDP). Using these parameter values, it is feasible to deliver $N_2^{\max}$ fractions with dose $d^{\min}$ in *all* scenarios in Z. This means that stage-1 decisions cannot render stage 2 infeasible for RO, NOM (and their FH counterparts) or PI. Numerical results indicate that results are not sensitive to the choice of $d^{\min}$ and $d_1^{\max}$.

First, Section 5.3 presents and discusses the results for the problem with exact biomarker information (EDP) of Section 3. After that, Section 5.4 presents and discusses the results for the problem with inexact biomarker information (IDP) of Section 4. Lastly, Section 5.5 again considers the inexact biomarker information case, and varies parameter N_1 in order to determine the optimal moment of biomarker acquisition.

We consider a sample of 200 scenarios for (ρ, τ) from Z. For each scenario, we compute the the average tumor BED over 20 patients is computed, thus creating a tumor BED distribution for the 'average' patient. For this tumor BED distribution we report the mean, 5% quantile and worst-case value. Next to this, we report the true worst-case tumor BED over Z (averaged over 20 patients). Note that the true worst-case scenario can differ per patient, so the true worst-case BED is is typically not attained in the sample. For OAR violations, we report the percentage by which the OAR BED tolerance is exceeded (i.e., percentage overdose). The maximum violation is the maximum value found over all patients and scenarios. All reported decision variable statistics are averaged over all patients and scenarios.

5.3 Results exact biomarker information

Table 2 presents the results. Altogether, the results indicate that the value of exact information is high. NOM-FH performs very similar to ARO. This illustrates that ignoring uncertainty and adaptability in stage 1 neither compromises worst-case or mean performance, nor does it lead to OAR constraint violations if treatment can be adapted based on exact biomarker information. In fact, NOM-FH outperforms RO-FH, indicating that accounting for uncertainty in stage-1 is overly conservative.

NOM is the only method that is not worst-case optimal, but yields the highest mean tumor BED accross the sample. However, it is the only method that results in OAR constraint violations. In the nominal scenario $(\bar{\rho}, \bar{\tau})$ it is optimal to hypofractionate for all patients, so the mean N_2 equals 20 for NOM. The other static method, RO, is worst-case optimal, but yields lower tumor BED accross the entire sample. It delivers significantly more fractions on average, i.e., it decides to hyperfractionate more often.

NOM-FH adds a folding horizon component to NOM, and this results in zero violations and worst-case optimality. It does have slightly lower sample mean tumor BED. RO-FH adds a folding horizon component to RO, and this results in improved performance accross the entire sample. It chooses to hypofractionate more often than RO. ARO is worst-case optimal and performs very similar to NOM-FH. Excluding NOM (for OAR constraint violations) and PI (not implementable), NOM-FH, ARO, RO-FH and RO yield the (possibly joint) highest objective value in 83.4%, 76.4%, 22.1% and 0.6% of all (scenario, patient) instances, respectively.

The results of Table 2 show that the methods have different stage-1 decisions d_1; this indicates the existence of multiple worst-case optimal stage-1 solutions. As indicated in Section 5.2, RO, RO-FH and ARO optimize over auxiliary scenarios in this case. According to Theorem 2, ARO finds a $\mathcal{PARO}$ solution this way. Overall, methods that deliver a relatively low dose in stage-1 perform better than the methods that deliver a higher dose. This may be data set-specific. From PI we see that for the majority of patients and scenarios hypofractionation is optimal (average $N_2 = 22.2$). We emphasize that for different data sets, where for the majority of scenarios and patients hyperfractionation is optimal, a higher stage-1 dose (which allows for $N^{\max}$ constant-dose fractions) may perform better, such as the result of RO and RO-FH.

	Method					
	NOM	NOM-FH	RO	RO-FH	ARO	PI
Tumor BED - sample mean (Gy)	162.75	161.44	156.57	160.14	161.40	161.49
Tumor BED - sample 5% quantile (Gy)	151.98	150.94	147.57	150.04	150.90	151.04
Tumor BED - sample wc (Gy)	145.98	146.33	142.53	145.58	146.32	146.39
Tumor BED - wc over Z (Gy)	**114.72**	**116.19**	**116.19**	**116.19**	**116.19**	**116.19**
OAR violation - mean (%)	1.25	0	0	0	0	0
OAR violation - max (%)	4.22	0	0	0	0	0
Stage-1 dose d_1 (Gy)	1.50	1.50	2.29	2.29	1.51	1.66
Stage-2 dose d_2 (Gy)	3.45	3.24	2.48	2.95	3.24	3.19
Stage-2 fractions N_2	20.0	22.2	27.2	22.2	22.2	22.2

Table 2: Results for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z (200 scenarios). For each scenario, results are averaged over 20 patients*. All methods optimize for worst-case tumor BED in Z, which is displayed in bold.
*: the maximum OAR violation is computed over all patients and scenarios.

Figure 7 shows the complete cumulative scenario-tumor BED graph. A point (x, y) in Figure 7 can be interpreted as follows: for the average patient, in $y\%$ of scenarios the tumor BED is at least x Gy. The results clearly demonstrate that RO and RO-FH are outperformed by the other methods. Both NOM-FH and ARO are visually almost indistinguishable from PI. NOM performs even better accross the entire sample (except the first percent of the sample), at the
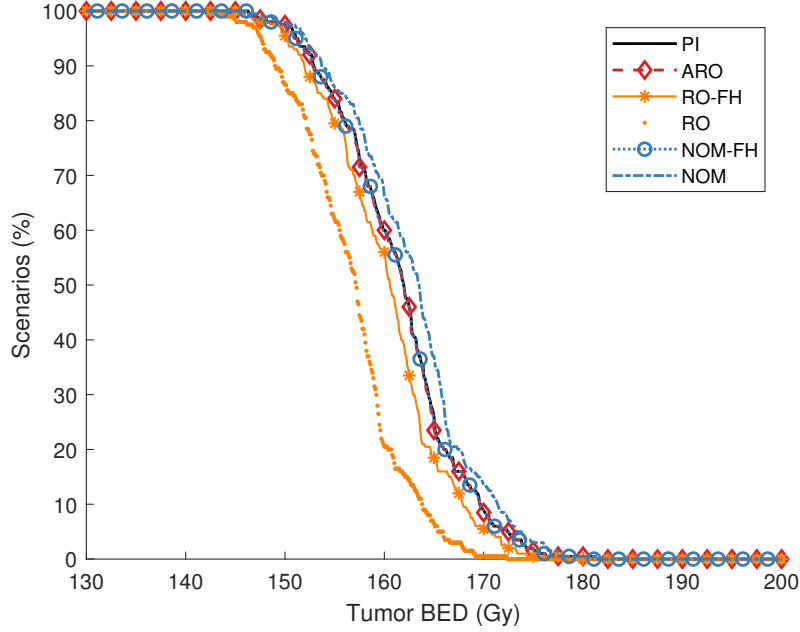
Figure 7: Cumulative scenario-tumor BED graph for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z (200 scenarios). A point (x, y) indicates that in $y\%$ of scenarios the tumor BED (averaged over 20 patients) is at least x Gy. ARO and NOM-FH are very close to PI.

cost of OAR BED violations.

To see the difference in mean performance between the multiple worst-case optimal solutions, we compare $\mathcal{PARO}$ solution found by the ARO method to the $\mathcal{ARO}$ solution that performs worst in the two auxiliary scenarios. Table 3 shows the results, OAR constraint violations are zero in all cases. The worst-performing $\mathcal{ARO}$ solution has a considerably higher stage-1 dose. This implies that (for the current parameter settings) delivering a high stage-1 dose does not allow as much adjustment possibilities in stage 2 as a low stage-1 dose, but it does allow for adjustments to reach the worst-case optimum. Relative to the results of Table 2, the difference between the best and worst ARO solution is considerable: the worst-performing ARO solution performs worse than the RO-FH solution.

	Method	
	$\mathcal{ARO}_{\text{worst}}$	$\mathcal{ARO}_{\text{best}}$
Tumor BED - sample mean (Gy)	159.87	161.40
Tumor BED - sample 5% quantile (Gy)	149.86	150.90
Tumor BED - sample worst-case (Gy)	145.34	146.32
Tumor BED - wc over Z (Gy)	**116.19**	**116.19**
Stage-1 dose d_1 (Gy)	2.67	1.51
Stage-2 dose d_2 (Gy)	2.79	3.24
Stage-2 fractions N_2	22.2	22.2

Table 3: Comparison between the best ($\mathcal{PARO}$) and worst performing $\mathcal{ARO}$ solutions, for uniform sampling of (ρ, τ) over Z (200 scenarios). For each scenario, results are averaged over 20 patients. All methods optimize for worst-case tumor BED in Z, which is displayed in bold. OAR constraint violations are zero in all cases.

Appendix A reports the results of an auxiliary experiment where the (ρ, τ) samples are drawn from a superset of Z, to compare the out-of-sample performance of the methods. NOM remains the only method with OAR constraint violations. Compared to Table 2, static methods NOM

22

and RO have poor performance. The relative performance of the adaptive methods remains mostly unchanged.

Altogether, the results of Section 5.3 demonstrate that if exact biomarker information is available mid-treatment, most stage-1 decisions allow for sufficient adaptation space in stage 2, also with realizations outside of Z. Different stage-1 decisions yield the worst-case optimum, have good performance on the scenario sample and have no OAR BED violations. We note that all presented differences in tumor BED are of relatively small magnitude. One reason for this is that the number of stage-2 fractions is restricted to $[N_2^{\min}, N_2^{\max}] = [20, 30]$. If the minimum number of fractions represents a 'true' hypofractionation case, the dose per fraction can vary more, and the difference in performance between hypo- and hyperfractionation strategies is amplified.

5.4 Results inexact biomarker information

In case of in inexact biomarker information (IDP), we do not obtain the true parameter values (ρ, τ) after $N_1 = 10$ fractions, but only an estimate $(\hat{\rho}, \hat{\tau})$. As discussed in Section 4, we specify a new uncertainty set $\hat{Z}$ such that $(\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}$. Let DQ $\in [0, 1]$ indicate the data quality. Then we set $\hat{Z}$ such that the width of the new uncertainty intervals for τ and ρ is $(1 - \text{DQ})$ times the width of the original intervals $[\tau_L, \tau_U]$ and $[\rho_L, \rho_U]$. That is, DQ$\cdot 100\%$ can be interpreted as the percentage by which the uncertainty intervals can be reduced due to the observation. The relation with the accuracy parameter r^ρ (or similarly r^τ) is given by

$$r^\rho = \frac{1}{2}(\rho_U - \rho_L)(1 - \text{DQ}). \tag{37}$$

Note that even DQ $= 0$ has some value as the new interval is centered around the observation, which already cuts off part of the original uncertainty set Z. We pick DQ $= 2/3$, so the obtained information after fraction N_1 reduces the size of the interval by 66.7% around the new observation. Variations for DQ are considered in Section 5.5. For all 20 patients the required auxiliary scenarios for the method of Section 4.3 can be found.

Table 4 shows the results. The robust methods RO, RO-FH and ARO are all worst-case optimal. This indicates that, although not theoretically guaranteed, ARO finds an ARO solution in all considered scenarios. The mean performance of RO-FH and ARO is further away from PI than in the case with exact biomarker information (Table 2). This is as expected, as due to inexact observations the possibility for ARO and RO-FH to make adjustments is less valuable, whereas PI is not in influenced by this. On the other hand, NOM and NOM-FH are not worst-case optimal, but have better performance on the sample of scenarios, at the cost of OAR violations.

ARO is the only method (together with PI) that has a different stage-1 decision than in the case with exact biomarker information. This is because it is the only method that takes inexactness of biomarker information into account at the start of stage 1. The average stage-1 dose d_1 differs considerably between ARO and RO-FH, whereas their worst-case performance is equal on average (and equal to PI). This demonstrates the existence of multiple worst-case optimal solutions. Whereas optimizing worst-case optimal solutions for ARO over two auxiliary scenarios does not guarantee a PARO solution (Section 4.3), results in Table 4 indicate that it does yield solutions that perform slightly better on average than RO-FH.

For ARO, it is noteworthy that the average number of stage-2 fractions (24.3 fx) differs from that of PI (22.1 fx). Although ARO uses optimal decision rules for stage 2, these are optimal for the worst-case scenario in the new uncertainty set $Z_{(\hat{\rho}, \hat{\tau})}$, and need not be optimal for the 'true' realization in this set. In fact, NOM-FH treats the inexact biomarker information as the 'true' parameter values, and administers 22.0 fractions, on average, which is closer than that of PI. Although the fractionation decision of NOM-FH is not worst-case optimal, Table 4 shows that it performs better on the samples of scenarios.

	Method					
	NOM	NOM-FH	RO	RO-FH	ARO	PI
Tumor BED - sample mean (Gy)	162.52	161.03	156.38	158.76	159.46	161.16
Tumor BED - sample 5% quantile (Gy)	151.36	150.12	147.04	148.51	148.94	150.21
Tumor BED - sample wc (Gy)	147.79	146.04	144.01	145.02	145.33	146.18
Tumor BED - wc over Z (Gy)	**114.72**	**115.96**	**116.19**	**116.19**	**116.19**	**116.19**
OAR violation - mean (%)	1.25	0.16	0	0	0	0
OAR violation - max (%)	4.23	1.49	0	0	0	0
Stage-1 dose d_1 (Gy)	1.50	1.50	2.29	2.29	1.79	1.65
Stage-2 dose d_2 (Gy)	3.45	3.25	2.48	2.78	2.92	3.20
Stage-2 fractions N_2	20.0	22.0	27.2	24	24.3	22.1

Table 4: Results for experiments with inexact biomarker information (data quality DQ = 2/3) and uniform sampling of (ρ, τ) over Z (200 scenarios). All results are averages over a sample of 20 patients. For each scenario, results are averaged over 20 patients*. All methods optimize for worst-case tumor BED in Z, which is displayed in bold.
*: the maximum OAR violation is computed over all patients and scenarios.

Figure 8 shows the complete cumulative scenario-tumor BED graph for the 'average patient'. Whereas in case of exact biomarker information (Figure 7), the ARO line was very close to PI, here a clear difference can be observed. NOM and NOM-FH outperform ARO (and RO and RO-FH) over the entire distribution.
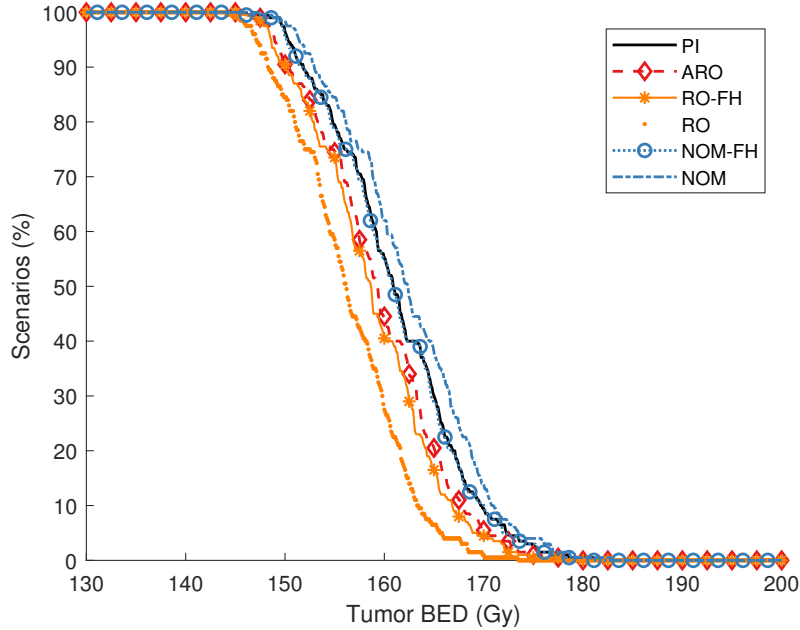

Figure 8: Cumulative scenario-tumor BED graph for experiments with inexact biomarker information (data quality DQ = 2/3) and uniform sampling of (ρ, τ) over Z (200 scenarios). A point (x, y) indicates that in $y\%$ of scenarios the tumor BED (averaged over 20 patients) is at least x Gy. NOM-FH is very close to PI.

The good performance of NOM and NOM-FH in terms of sample mean tumor BED does come at the cost of OAR violations. However, these are relatively minor. The reason for this is that the number of stage-2 fractions is relatively high (between 20 and 30fx), so any method delivers reasonably low dose per fraction in stage 2. Consequently, the quadratic term in the BED model is smaller, and so is the influence of the α/β parameters. With higher dose per fraction, the use of incorrect (e.g., nominal) α/β parameter values may yield higher OAR
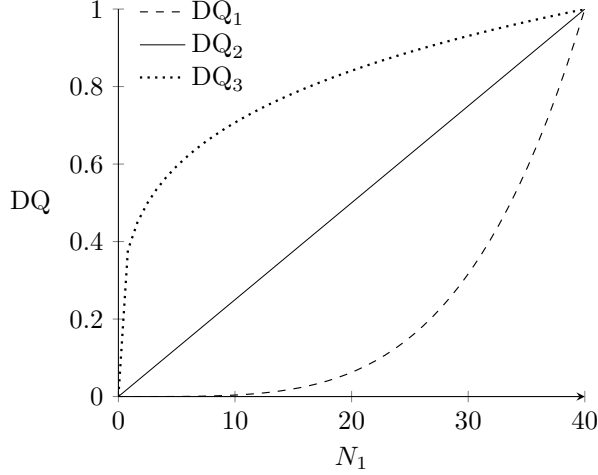
Figure 9: The biomarker data quality is a function of the number of treatment fractions N_1 after which it is acquired. We consider three functions $\text{DQ}_i(N_1)$, $i = 1, 2, 3$.

constraint violations. Preliminary experiments for stereotactic body radiation therapy (SBRT, an RT modality that uses around five high dose fractions) indeed result in slightly higher OAR constraint violations for NOM and NOM-FH. In any case, a trade-off can be observed between higher tumor BED attained by NOM and NOM-FH and associated OAR constraint violations.

5.5 Optimal moment of biomarker acquisition

The moment of biomarker observation, need not be fixed. Part of the decision making process then involves choosing this observation moment such that it maximally improves treatment quality. Late observation may result in limited possibilities for treatment adaptation, whereas with too early observation one cannot yet reliably observe the true individual patient response. Although one can incorporate N_1 as a decision variable in the mathematical model, the small decision space allows to simply vary its value in numerical experiments. We assume a (hypothetical) mathematical relationship between information point N_1 and the data quality parameter DQ. With $N^{\max}$ the maximum number of fractions, we consider the following three data quality functions:

$$\text{DQ}_1(N_1) = \left(\frac{N_1}{N^{\max}} \right)^4 \tag{38a}$$

$$\text{DQ}_2(N_1) = \frac{N_1}{N^{\max}} \tag{38b}$$

$$\text{DQ}_3(N_1) = \left(\frac{N_1}{N^{\max}} \right)^{1/4}. \tag{38c}$$

Hence, DQ_1, DQ_2 and DQ_3 describe a convex, linear and concave relationship between observation moment and data quality, respectively. Figure 9 shows the graphs of the three functions. Whether DQ_1, DQ_2 or DQ_3 is most realistic depends on the specific biomarker(s) that is/are used, see Section 2.2 for details.

Figure 10 shows the results for the numerical experiments where we vary the information point N_1 from 0 to $N^{\max} - 1$ for data quality functions $\text{DQ}_i(N_1)$, $i = 1, 2, 3$. The left vertical axis indicates the mean tumor BED (averaged over all patients and scenarios), the right vertical axis indicates the maximum OAR tolerance violation for NOM and NOM-FH. It is important to note that as N_1 increases past $N^{\min} = 30$, this also increases the minimum number of fractions
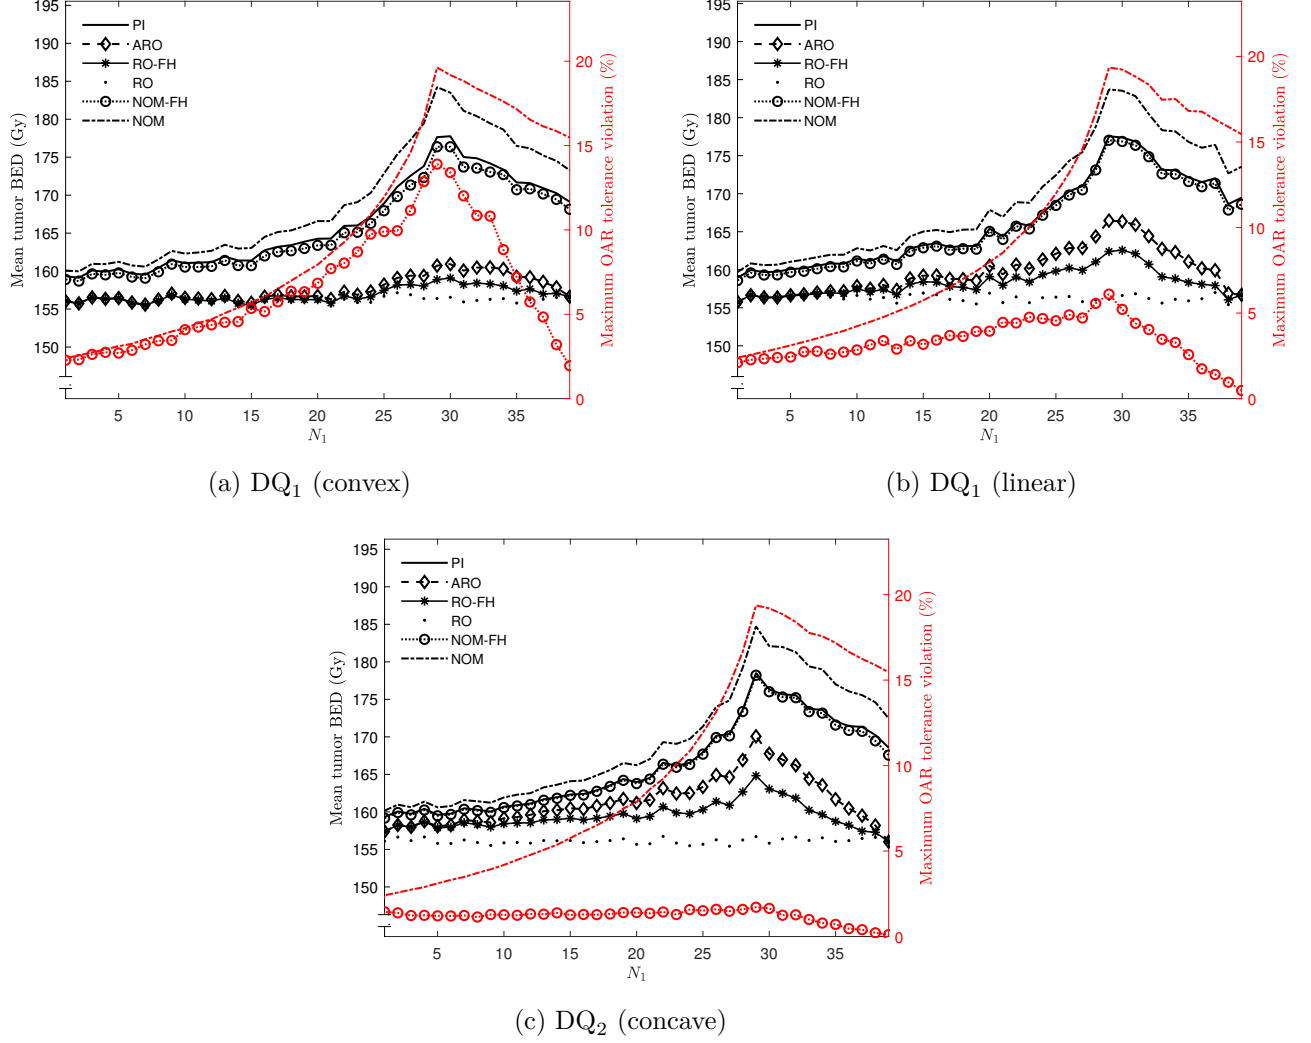
(a) DQ$_1$ (convex)

(b) DQ$_1$ (linear)

(c) DQ$_2$ (concave)

Figure 10: Results for varying the information point N_1 from 1 to $N^{\max} - 1$, for data quality functions $DQ_i(N_1)$, $i = 1, 2, 3$. The maximum OAR BED constraint violation (%) of NOM (dash-dotted) and NOM-FH (dotted, circle marker) is measured against the right vertcial axis. Note that the left vertical axis measures displays the mean tumor BED (averaged over all patients and scenarios), while the methods maximize the worst-case tumor BED per patient.

correspondingly. Moreover, the dose per fraction is constant per treatment stage, so the choice of N_1 also influences the types of treatments that can be delivered.

For these reasons the curve for PI is not constant, even though it does not actually use biomarker information. The optimal moment of biomarker acquisition for PI is $N_1 = 29$. This is because the minimum number of fractions is $N^{\min} = 30$. Hence, if hypofractionation is optimal we can deliver one more fraction with high dose, and deliver a low dose in stage 1. If hyperfractionation is optimal we can deliver 11 more fractions (and get the total maximum of 40) with low dose. Having $N_1 > 29$ forces the use of more than $N_2^{\min} = 30$ fractions, which is disadvantageous for those (patients, scenario) cases where hypofractionation is optimal.

For the latter reason the NOM curve is also not constant. It results in a higher tumor BED than PI for any value of N_1, at the cost of OAR violations of up to 20%. NOM-FH yields a sample mean tumor BED close to PI for all three DQ functions, and the OAR violations depend on the DQ function. With poor data quality (Figure 10a) and an observation moment close to $N_1 = 29$, OAR violations over 10% are possible, despite the fact that NOM-FH is an adaptive method. On the other hand, with good (Figure 10a) data quality, the violations remain

below 2%. The OAR tolerance violations are highest near $N_1 = 29$. This is because in case of hypofractionation in stage 2 the influence of the α/β parameters is highest, as was noted in Section 5.4.

The robust methods RO, RO-FH and ARO do not result in any OAR violations, by construction. The better the data quality, the larger the differences between RO, RO-FH and ARO. This implies that, if robustness is required, there is value in (i) adapting based on inexact information, (ii) taking adaptability into account when planning the stage-1 dose. The good performance of NOM-FH shows that this value diminishes if OAR violations are allowed. NOM-FH does not account for adaptability, and does not take inexactness of biomarker information into account. Nevertheless, it results in higher sample mean tumor BED for any N_1, and the difference increases from poor (convex) to good (concave) data quality. Thus, Figure 10c illustrates the trade-off between higher sample mean tumor BED and possible OAR violations that was also observed in Section 5.4 (for the entire sample distribution).

The shape of the data quality function influences the optimal moment of biomarker observation only slightly. For all adaptive methods, we find that the peak is more pronounced for high data quality (concave) than low data quality (convex), but it is centered around $N_1 = 29$. In case of convex data quality the peak is relatively flat, indicating a trade-off between observing at $N_1 = 29$ (giving maximum adaptation flexibility) and postponing (waiting for higher data accuracy).

6   Concluding remarks

In this paper we have presented an ARO approach to optimally adapt the treatment length of radiation therapy treatments, using mid-treatment biomarker information. Using an ARO approach, adaptability is taken into account prior to treatment and it provides insight into the optimal stage-2 decisions.

In the case of exact biomarker information, there is sufficient space to adapt, and numerical results show that taking into account both robustness and adaptability is not necessary. In the case of inexact biomarker information, adaptive strategies can use only parameter estimates instead of true parameter values, and may still yield violations if this uncertainty is not accounted for. Accounting for adaptability and inexactness of biomarker information is particularly beneficial when robustness (w.r.t. OAR violations) is of high importance. If minor OAR violations are allowed, NOM-FH is a good performing alternative, which can outperform ARO. NOM-FH and ARO thus yield a trade-off between higher performance and OAR violations. Both the difference in performance and the magnitude of OAR violations of NOM-FH are highly influenced by the data quality (i.e., accuracy of parameter estimates).

The current setting can be extended in several ways. In practice the tumor and OAR α/β values would have to be estimated from actual biomarkers (e.g., imaging, blood-based biomarkers, genotyping), which can be incorporated in the model. Furthermore, the approach can be extended to heterogeneous tumor response (different α/β ratios for different tumor subvolumes), or time-dependent response parameters. Other RT applications may also benefit from ARO, such as re-optimization to account for organ motion or setup errors, optimization using the MR-linac or combining RT with chemotherapy.

References

Ajdari, A. and Ghate, A. (2017). Robust fractionation in radiotherapy. Unpublished paper.

Ajdari, A., Ghate, A., and Kim, M. (2018). Adaptive treatment-length optimization in spatiobiologically integrated radiotherapy. *Phys. Med. Biol.*, 63(7):075009.

Ajdari, A., Niyazi, M., Nicolay, N., Thieke, C., Jeraj, R., and Bortfeld, T. (2019). Towards optimal stopping in radiation therapy. *Radiother. Oncol.*, 134:96–100.

Barker, C. A. and Powell, S. N. (2010). Enhancing radiotherapy through a greater understanding of homologous recombination. *Semin. Radiat. Oncol.*, 20(4):267–273.

Baumann, M., Krause, M., Overgaard, J., Debus, J., Bentzen, S. M., Daartz, J., ..., and Bortfeld, T. (2016). Radiation oncology in the era of precision medicine. *Nat. Rev. Cancer*, 16(4):234–249.

Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Math. Program.*, 99:351–376.

Bindra, R. S., Goglia, A. G., Jasin, M., and Powell, S. N. (2013). Development of an assay to measure mutagenic non-homologous end-joining repair activity in mammalian cells. *Nucleic Acids Res.*, 41(11):e115.

Böck, M., Eriksson, K., and Forsgren, A. (2019). On the interplay between robustness and dynamic planning for adaptive radiation therapy. *Biomed. Phys. Eng. Express*, 5(4):045004.

Böck, M., Eriksson, K., Forsgren, A., and Hårdemark, B. (2017). Toward robust adaptive radiation therapy strategies. *Med Phys.*, 44(6):2054–2065.

Bortfeld, T., Ramakrishnan, J., Tsitsiklis, J. N., and Unkelbach, J. (2015). Optimization of radiation therapy fractionation schedules in the presence of tumor repopulation. *INFORMS J. Comput.*, 27(4):788–803.

Chan, T. C. Y. and Mišić, V. V. (2013). Adaptive and robust radiation therapy optimization for lung cancer. *Eur. J. Oper. Res*, 231:745–756.

Cox, J. D. (1986). Presidential address: Fractionation: A paradigm for clinical research in radiation oncology. *Int. J. Radiat. Oncol. Biol. Phys.*, 13:1271–1281.

Dabadghao, S. and Roy, A. (2020). Optimal interventions for adaptive robust optimization under time-dependent uncertainty with application to radiotherapy. Available at SSRN.

De Ruiter, F. J. C. T., Ben-Tal, A., Brekelmans, R. C. M., and den Hertog, D. (2017). Robust optimization of uncertain multistage inventory systems with inexact data in decision rules. *Comput. Manag. Sci.*, 14(1):45–77.

De Ruiter, F. J. C. T., Brekelmans, R. C. M., and den Hertog, D. (2016). The impact of the existence of multiple adjustable robust solutions. *Math. Program.*, 160:531–545.

Ehrgott, M., Güler, Ç., Hamacher, H. W., and Shao, L. (2008). Mathematical optimization in intensity modulated radiation therapy. *4OR-Q. J. Oper. Res.*, 6:199–262.

Fowler, J. F. (1989). The linear-quadratic formula and progress in fractionated radiotherapy. *Brit. J. Radiol.*, 62(740):679–694.

Fowler, J. F. (2010). 21 Years of biologically effective dose. *Brit. J. Radiol.*, 83:554–568.

Ghate, A. (2011). Dynamic optimization in radiotherapy. In Geunes, J., Gray, P., and Greenberg, H., editors, *Tutorials in Operations Research: transforming research into action*, chapter 4. INFORMS.

Hall, E. J. and Giaccia, A. J. (2012). *Radiobiology for the radiologist.* Lippincott Williams & Wilkins. Philadelphia, Pennsylvania, USA.

Iancu, D. A. and Trichakis, N. (2014). Pareto efficiency in robust optimization. *Manag. Sci.*, 60(1):130–147.

Iancu, D. A., Trichakis, N., and Yoon, D. Y. (2020). Monitoring with limited information. *Manag. Sci.* Articles in Advance.

Kehwar, T. S. (2005). Analytical approach to estimate normal tissue complication probability using best fit of normal tissue tolerance doses into the NTCP equation of the linear quadratic model. *J. Canc. Res. Ther.*, 1(3):168–179.

Kim, M., Ghate, A., and Phillips, M. H. (2012). A stochastic control formalism for dynamic biologically conformal radiation therapy. *Eur. J. Oper. Res.*, 219:541–556.

Lim, G. J., Kardar, L., Ebrahimi, S., and Cao, W. (2020). A risk-based modeling approach for radiation therapy treatment planning under tumor shrinkage uncertainty. *Eur. J. Oper. Res.*, 280(1):266–278.

Long, T. (2015). *Optimization problems in radiation therapy treatment planning.* PhD thesis, University of Michigan. Chapter 5: Adaptive treatment planning for lung cancer.

Mahmoudzadeh, H. (2015). *Robust optimization methods for breast cancer radiation therapy.* PhD thesis, University of Toronto. Chapter 4: Pareto robust optimization in breast cancer RT.

Mar, P. A. and Chan, T. C. Y. (2015). Adaptive and robust radiation therapy in the presence of drift. *Phys. Med. Biol.*, 60(9):3599–3615.

Mizuta, M., Takao, S., Date, H., Kishimoto, N., Sutherland, K. L., Onimaru, R., and Shirato, H. (2012). A mathematical study to select fractionation regimen based on physical dose distribution and the linear-quadratic model. *Int. J. Radiat. Oncol. Biol. Phys.*, 84(3):829–833.

Nohadani, O. and Roy, A. (2017). Robust optimization with time-dependent uncertainty in radiation therapy. *IISE Trans. Healthc. Syst. Eng.*, 7(2):81–92.

Perkó, Z., Bortfeld, T., Hong, T., Wolfgang, J., and Unkelbach, J. (2018). Derivation of mean dose tolerances for new fractionation schemes and treatment modalities. *Phys. Med. Biol.*, 63(3):035038.

Saberian, F., Ghate, A., and Kim, M. (2016a). Optimal fractionation in radiotherapy with multiple normal tissues. *Math. Med. Biol.*, 33(2):211–252.

Saberian, F., Ghate, A., and Kim, M. (2016b). A theoretical stochastic control framework for adapting radiotherapy to hypoxia. *Phys. Med. Biol.*, 61(19):7136–7161.

Saberian, F., Ghate, A., and Kim, M. (2017). Spatiotemporally optimal fractionation in radiotherapy. *INFORMS J. Comput.*, 29(3):422–437.

Saka, B., Rardin, R. L., and Langer, M. P. (2014). Biologically guided intensity modulated radiation therapy planning optimization with fraction-size dose constraints. *J. Oper. Res. Soc.*, 65:557–571.

Santiago, A., Barczyk, S., Jelen, U., Engenhart-Cabillic, R., and Wittig, A. (2016). Challenges in radiobiological modeling: can we decide between LQ and LQ-L models based on reviewed clinical NSCLC treatment outcome data. *Radiat. Oncol.*, 11:67.

Shepard, D. M., Ferris, M. C., Olivera, G. H., and Mackie, T. R. (1999). Optimizing the delivery of radiation therapy to cancer patients. *SIAM Rev.*, 41(4):721–744.

Somaiah, N., Rothkamm, K., and Yarnold, J. (2015). Where do we look for markers of radiotherapy fraction size sensitivity? *Clin. Oncol.*, 27:570–578.

Van Leeuwen, C. M., Oei, A., Crezee, J., Bel, A., Franken, N. A. P., Stalpers, L. J. A., and Kok, H. P. (2018). The alpha and beta of tumours: a review of parameters of the linear-quadratic model, derived from clinical radiotherapy studies. *Radiation. Oncol.*, 13:96.

Withers, H. R. (1985). Biological basis for altered fractionation schemes. *Cancer*, 55:2086–2095.

Yanıkoğlu, İ., Gorissen, B. L., and den Hertog, D. (2019). A survey of adjustable robust optimization. *Eur. J. Oper. Res.*, 277:799–813.

Appendix

A   Results exact biomarker information: out-of-sample performance

To investigate the out-of-sample performance of the methods, we assume a uniform distribution for (ρ, τ) over a larger set than Z. We can write Z as

$$Z = \{(\rho, \tau) : \rho_L \le \rho \le \rho_U, \tau_L \le \tau \le \tau_U\} = \{(\rho, \tau) : |\bar{\rho} - \rho| \le \varepsilon_\rho, \ |\bar{\tau} - \tau| \le \varepsilon_\tau\}, \qquad (A.1)$$

where $(\varepsilon_\rho, \varepsilon_\tau)$ is the maximum deviation from the nominal scenario $(\bar{\rho}, \bar{\tau})$. This allows us to define

$$Z_c = \{(\rho, \tau) : |\bar{\rho} - \rho| \le c\varepsilon_\rho, \ |\bar{\tau} - \tau| \le c\varepsilon_\tau\}, \qquad (A.2)$$

where $c > 0$ is a parameter. We assume a uniform distribution over the new set Z_c. If $c = 1$, we have $Z_c = Z$, so we sample exactly from Z. If $c > 1$, we sample from an interval that is c^2 times as large as Z (c times larger for both ρ and τ). For $c = 2$ we obtain the results in Table A.1. The stage-1 dose d_1 is the same as in Table 2 for all methods except PI, because PI is the only method that is aware that the samples are not taken from uncertainty set Z but from Z_2. For NOM, the maximum violation percentage has increased slightly. All other methods are able to deal with the out-of-sample realizations and do not have any OAR constraint violations.

Due to the larger sampling space (the area of Z_2 is four times the area of Z), the difference between sample mean and sample worst-case performance is much larger than in Table 2 for all methods. The true worst-case objective value in Z is still lower than the sample worst-case in Z_2. The reason for this is that the true worst-case scenario can differ per patient. The relative performance of the adaptive methods remains mostly unchanged.

| | Method | | | | | |
	NOM	NOM-FH	RO	RO-FH	ARO	PI
Tumor BED - sample mean (Gy)	156.49	158.92	151.30	156.83	158.87	159.06
Tumor BED - sample 5% quantile (Gy)	140.14	142.11	137.55	141.11	142.11	142.32
Tumor BED - sample wc (Gy)	134.06	136.80	132.56	136.09	136.79	137.12
Tumor BED - wc over Z (Gy)	**114.72**	**116.19**	**116.19**	**116.19**	**116.19**	**116.19**
OAR violation - mean (%)	1.16	0	0	0	0	0
OAR violation - max (%)	5.30	0	0	0	0	0
Stage-1 dose d_1 (Gy)	1.50	1.50	2.29	2.29	1.51	1.73
Stage-2 dose d_2 (Gy)	3.45	3.21	2.48	2.92	3.21	3.15
Stage-2 fractions N_2	20	22.9	27.2	22.9	22.9	22.9

Table A.1: Results for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z_2. For each scenario, results are averaged over 20 patients*. All methods optimize for worst-case tumor BED in Z, which is displayed in bold.

*: the maximum OAR violation is computed over all patients and scenarios
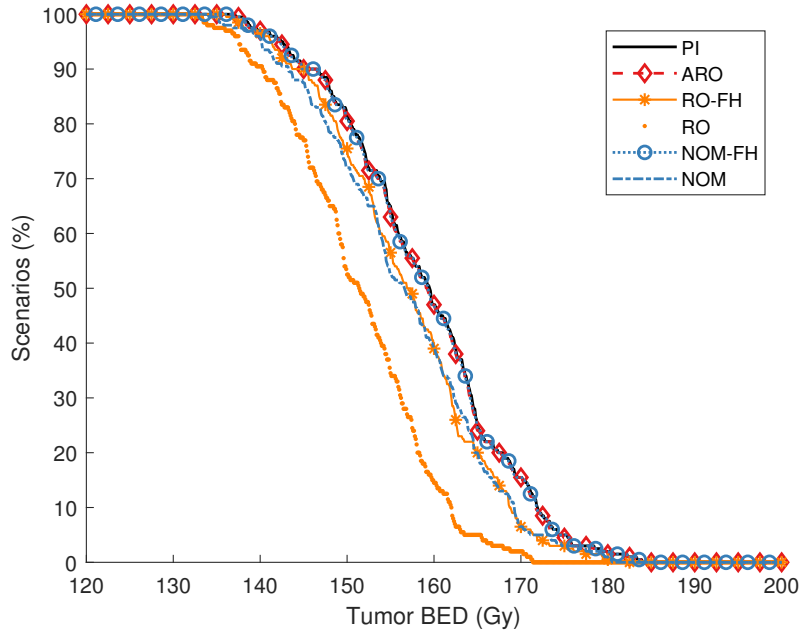

Figure A.1: Cumulative scenario-tumor BED graph for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z_2 (200 scenarios). A point (x, y) indicates that in $y\%$ of scenarios the tumor BED (averaged over 20 patients) is at least x Gy. ARO and NOM-FH are very close to PI.

NOM-FH and ARO are near optimal for the worst-case sample scenario, and are also close to PI in the sample 5% quantile and sample mean. The relative performance of the adaptive methods remains mostly unchanged, RO-FH performs slightly worse than NOM-FH and ARO, similar to Table 2. Compared to Table 2, NOM and RO have poor performance across the sample. This indicates bad performance of the static methods on scenarios outside of Z.

Figure A.1 shows the complete cumulative scenario-tumor BED graph for the 'average patient'. Compared to Figure 7, the main difference is the decrease in performance of NOM. Naturally, the performance of static nominal optimization is directly related to the magnitude of possible deviations from the nominal scenario, which is higher in Z_2 than in Z.

B   Extra analyses and proofs

For convenience, we repeat the definitions of functions B, g and f:

$$B(d', N'; \rho) := \varphi D \Big(1 + \frac{\varphi D}{T}\rho\Big) - \sigma d_1 N_1 - \sigma^2 \rho d_1^2 N_1 \tag{B.1a}$$

$$g(d', N', N''; \rho) := \frac{-1 + \sqrt{1 + \frac{4\rho}{N''} B(d', N'; \rho)}}{2\sigma\rho} \tag{B.1b}$$

$$f(d_1, N_2; \rho, \tau) := \begin{cases} N_1 d_1 + N_2 g(d_1, N_1, N_2; \rho) + \tau\big(N_1 d_1^2 + N_2 g(d_1, N_1, N_2; \rho)^2\big) & \text{if } d_1 \in [0, g(0, 0, N_1; \rho)] \\ -\infty & \text{otherwise,} \end{cases} \tag{B.1c}$$

see (9), (10) and (17).

B.1   Proof Theorem 1

First, we show that for fixed d_1, feasible to (8), and given (ρ, τ), it is optimal to minimize the number of stage-2 fractions if $\tau \geq \sigma\rho$, and it is optimal to maximize the number of stage-2 fractions otherwise. After that, we show that with stage-2 dose d_2 such that (8b) holds with equality, $N_2(\rho, \tau) = N_2^{\min}$ is feasible if $\tau \geq \sigma\rho$ and $N_2(\rho, \tau) = N_2^{\max}$ is feasible otherwise.

Consider problem (8). At the start of stage 2, we have delivered N_1 fractions with dose d_1 per fraction. Let (ρ, τ) be the realization of the uncertain parameters. The stage-2 problem reads

$$N_1 d_1 + \tau N_1 d_1^2 + \max_{d_2, N_2}\ N_2 d_2 + \tau N_2 d_2^2 \tag{B.2a}$$

$$\text{s.t.}\ \ \sigma N_2 d_2 + \rho\sigma^2 N_2 d_2^2 \leq B(d_1, N_1, \rho) \tag{B.2b}$$

$$d_2 \geq d_{\min} \tag{B.2c}$$

$$N_2 \in \{N_2^{\min}, \dots, N_2^{\max}\}. \tag{B.2d}$$

This is a static fractionation problem. Constraint (B.2b) will hold with equality at the optimum, because it is the only dose-limiting constraint. This yields

$$d_2^*(d_1, N_2; \rho) = g(d_1, N_1, N_2; \rho). \tag{B.3}$$

Secondly, this allows us to rewrite the objective to

$$\max_{d_2, N_2}\ N_2 d_2 \Big(\frac{\sigma\rho - \tau}{\sigma\rho}\Big) + \frac{\tau B(d_1, N_1, \rho)}{\sigma^2\rho}, \tag{B.4}$$

which implies that if $\tau > \sigma\rho$ it is optimal to minimize $d_2 N_2$. If $\tau < \sigma\rho$ it is optimal to maximize $d_2 N_2$, and if $\tau = \sigma\rho$ the objective value is independent of the value of N_2. Similar results are obtained in Mizuta et al. (2012), Bortfeld et al. (2015). As given in Section 3, at the optimum

$$N_2 d_2^*(d_1, N_2; \rho) = N_2 g(d_1, N_1, N_2; \rho) = \frac{-N_2 + \sqrt{N_2^2 + 4N_2\rho B(d_1, N_1; \rho)}}{2\sigma\rho}, \tag{B.5}$$

and it is straightforward to show that

$$\frac{\partial N_2 g(d_1, N_1, N_2; \rho)}{\partial N_2} \geq 0. \tag{B.6}$$

Hence, if $\tau > \sigma\rho$, it is optimal to minimize the number of fractions, and if $\tau < \sigma\rho$ it is optimal to maximize the number of fractions. If $\tau = \sigma\rho$, every feasible number of fractions is optimal.

For the second part, we must show that for any $(\rho, \tau) \in Z \cap \{\tau \geq \sigma\rho\}$ resp. $(\rho, \tau) \in Z \cap \{\tau < \sigma\rho\}$, it is indeed possible to deliver $N_2^{\min}$ resp. $N_2^{\max}$ fractions with dose according to (B.3) in stage 2. That is, we must show

$$g(d_1, N_1, N_2^{\min}; \rho) \geq d^{\min}, \quad \forall(\rho, \tau) \in Z \cap \{\tau \geq \sigma\rho\} \tag{B.7a}$$

$$g(d_1, N_1, N_2^{\max}; \rho) \geq d^{\min}, \quad \forall(\rho, \tau) \in Z \cap \{\tau < \sigma\rho\}, \tag{B.7b}$$

which is equivalent to

$$d_1 \leq g(d^{\min}, N_2^{\min}, N_1; \rho), \quad \forall(\rho, \tau) \in Z \cap \{\tau \geq \sigma\rho\} \tag{B.8a}$$

$$d_1 \leq g(d^{\min}, N_2^{\max}, N_1; \rho), \quad \forall(\rho, \tau) \in Z \cap \{\tau < \sigma\rho\}. \tag{B.8b}$$

Lemma C.3a states that g is increasing or decreasing in ρ for a fixed first argument. Hence, it is sufficient to consider only the largest and smallest value of ρ in either subset of Z. Therefore, (B.8) is equivalent to

$$d_1 \leq g(d^{\min}, N_2^{\min}, N_1; \rho_L) \tag{B.9a}$$

$$d_1 \leq g(d^{\min}, N_2^{\min}, N_1; \min\{\frac{\tau_U}{\sigma}, \rho_U\}) \tag{B.9b}$$

$$d_1 \leq g(d^{\min}, N_2^{\max}, N_1; \frac{\tau_L}{\sigma}) \tag{B.9c}$$

$$d_1 \leq g(d^{\min}, N_2^{\max}, N_1; \rho_U). \tag{B.9d}$$

From (B.1b) we see that function g is decreasing in its second argument, so (B.9b) is redundant. The remaining three conditions in (B.9) hold true due to Assumption 1. Hence, an optimal decision rule for $N_2(\cdot)$ is given by

$$N_2^*(\rho, \tau) = \begin{cases} N_2^{\min} & \text{if } \tau \geq \sigma\rho \\ N_2^{\max} & \text{otherwise,} \end{cases} \tag{B.10}$$

and

$$d_2^*(d_1; \rho, \tau) = \begin{cases} g(d_1, N_1, N_2^{\min}; \rho) & \text{if } \tau \geq \sigma\rho \\ g(d_1, N_1, N_2^{\max}; \rho) & \text{otherwise.} \end{cases} \tag{B.11}$$

are optimal decision rules for $N_2(\cdot)$ and $d_2(\cdot)$, respectively. For $\tau \neq \sigma\rho$, these give the unique optimal decisions. For $\tau = \sigma\rho$ any $N_2 \in \{N_2^{\min}, \ldots, N_2^{\max}\}$ is optimal, and the corresponding optimal d_2 follows according to (B.3).

B.2   Proof Theorem 2

Due to Theorem 1 a stage-1 decision d_1 is $\mathcal{PARO}$ according to Definition 6 if conditions (21) hold with $(d_2^*(\cdot), N_2^*(\cdot))$ plugged in. Thus, we must show that for any $d_1 \in X^{\text{PARO}}$ there is no $\mathcal{ARO}$ $\bar{d}_1$ such that

$$f(d_1, N_2^*(\rho, \tau); \rho, \tau) \leq f(\bar{d}_1, N_2^*(\rho, \tau); \rho, \tau) \quad \forall(\rho, \tau) \in Z \tag{B.12a}$$

$$f(d_1, N_2^*(\bar{\rho}, \bar{\tau}); \bar{\rho}, \bar{\tau}) < f(\bar{d}_1, N_2^*(\bar{\rho}, \bar{\tau}); \bar{\rho}, \bar{\tau}) \quad \text{for some } (\bar{\rho}, \bar{\tau}) \in Z. \tag{B.12b}$$

If $|X^{\text{PARO}}| = 1$, then the single element yields a strictly better objective value than all other elements in X^{ARO} in either scenario $(\rho^{\text{aux-min}}, \tau^{\text{aux-min}})$ or $(\rho^{\text{aux-max}}, \tau^{\text{aux-max}})$ or both, so it is $\mathcal{PARO}$. For the remainder of this proof we assume $|X^{\text{PARO}}| \geq 2$.

Consider $X^{\text{aux-min}}$. By construction of $(\rho^{\text{aux-min}}, \tau^{\text{aux-min}})$ it holds that $\tau^{\text{aux-min}} \neq \sigma\rho^{\text{aux-min}}$. Hence, according to Lemma C.4, there can be at most two values for d_1 in $X^{\text{aux-min}}$ that yield the same objective value f in scenario $(\rho^{\text{aux-min}}, \tau^{\text{aux-min}})$. Hence, $|X^{\text{aux-min}}| = |X^{\text{PARO}}| = 2$. Denote the two elements of X^{PARO} by d_1' and d_1'', let $d_1' < d_1''$. Solutions d_1' and d_1'' are both optimal to (22) and (23). Hence, according to Lemma C.4, it holds that

$$d_1'' = t(d_1'; \rho^{\text{aux-min}}, \tau^{\text{aux-min}}) \tag{B.13a}$$

$$d_1'' = t(d_1'; \rho^{\text{aux-max}}, \tau^{\text{aux-max}}). \tag{B.13b}$$

From the definition of t (see (C.21)) we derive for $\sigma\rho \neq \tau$:

$$\frac{\partial t(d_1; \rho, \tau)}{\partial \rho} = \frac{2N_2^*(\rho, \tau)}{N_1 + N_2^*(\rho, \tau)} \frac{\partial g(d_1, N_1, N_2^*(\rho, \tau); \rho)}{\partial \rho}, \tag{B.14}$$

because $N_2^*(\rho, \tau)$ is constant in ρ unless $\sigma\rho = \tau$. According to Lemma C.3a, if for given N_2 it holds that $d_1 \neq d_1^-(N_2)$ and $d_1 \neq d_1^+(N_2)$ (defined in (C.3)), then function $g(d_1, N_1, N_2, \rho)$ is strictly increasing or decreasing in ρ. By construction, it holds that $d_1^+(N_2) = t(d_1^-(N_2); \rho, \tau)$ for any ρ. According
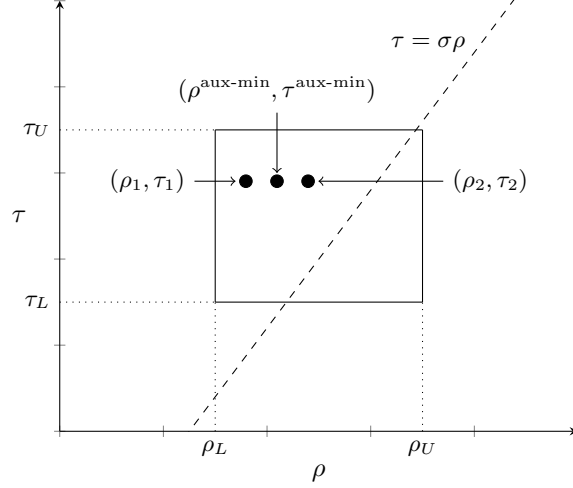
Figure B.1: Case $d_1' \neq d_1^-(N_2^{\min})$. Construction of two new scenarios (ρ_1, τ_1) and (ρ_2, τ_2) from scenario $(\rho^{\text{aux-min}}, \tau^{\text{aux-min}})$. Solution d_1' outperforms d_1'' at one scenario, vice versa at the other.

to Lemma C.2b, we have $d_1^-(N_2^{\min}) \neq d_1^-(N_2^{\max})$, so d_1' cannot be equal to both. Additionally, it cannot hold that $d_1' = d_1^+(N_2^{\min})$ or $d_1' = d_1^+(N_2^{\max})$, because it would imply $d'' \leq d'$. Hence, either $d_1' \notin \{d_1^-(N_2^{\min}), d_1^+(N_2^{\min})\}$ or $d_1' \notin \{d_1^-(N_2^{\max}), d_1^+(N_2^{\max})\}$ holds (or both).

We show that in either case, we can construct two new scenarios where d_1' outperforms d_1'' in one scenario, and vice versa in the other. Suppose $d_1' \notin \{d_1^-(N_2^{\min}), d_1^+(N_2^{\min})\}$. In this case, it holds that

$$\frac{\partial t(d_1'; \rho^{\text{aux-min}}, \tau^{\text{aux-min}})}{\partial \rho} \neq 0. \tag{B.15}$$

We consider two new scenarios. Let $\varepsilon > 0$ be a sufficiently small number and define

$$(\rho_1, \tau_1) := (\rho^{\text{aux-min}} - \varepsilon, \tau^{\text{aux-min}}) \qquad (\in \text{int}(Z^{\min})) \tag{B.16a}$$

$$(\rho_2, \tau_2) := (\rho^{\text{aux-min}} + \varepsilon, \tau^{\text{aux-min}}). \qquad (\in \text{int}(Z^{\min})) \tag{B.16b}$$

This is visualized in Figure B.1. Due to (B.15) and (B.13a), it holds that

$$\left(t(d_1'; \rho_1, \tau_1) > d'' \wedge t(d_1'; \rho_2, \tau_2) < d''\right) \vee \left(t(d_1'; \rho_1, \tau_1) < d'' \wedge t(d_1'; \rho_2, \tau_2) > d''\right). \tag{B.17}$$

If the first clause is is true, we obtain

$$f(d_1', N_2^{\min}; \rho_1, \tau_1) > f(d_1'', N_2^{\min}; \rho_1, \tau_1) \tag{B.18a}$$

$$f(d_1', N_2^{\min}; \rho_2, \tau_2) < f(d_1'', N_2^{\min}; \rho_2, \tau_2), \tag{B.18b}$$

where we used convexity of $f(d_1, N_2^{\min}; \rho, \tau)$ for $(\rho, \tau) \in \text{int}(Z^{\min})$. Similarly, if the second clause of (B.17) is true, we obtain

$$f(d_1', N_2^{\min}; \rho_1, \tau_1) < f(d_1'', N_2^{\min}; \rho_1, \tau_1) \tag{B.19a}$$

$$f(d_1', N_2^{\min}; \rho_2, \tau_2) > f(d_1'', N_2^{\min}; \rho_2, \tau_2). \tag{B.19b}$$

In either case, there is a scenario in $Z^{\min}$ where d_1' outperforms d_1'' and a scenario in $Z^{\min}$ where d_1'' outperforms d_1'. Hence, both d_1' and d_1'' are $\mathcal{PARO}$. Using similar arguments, we can show that in case $d_1' \notin \{d_1^-(N_2^{\max}), d_1^+(N_2^{\max})\}$ also both d_1' and d_1'' are $\mathcal{PARO}$.

B.3 Proof Theorem 3

Consider problem (29). At the start of stage 2, we have delivered N_1 fractions with dose d_1 per fraction. Let $(\hat{\rho}, \hat{\tau})$ be the observation. The resulting stage-2 problem for (29) reads

$$\max_{d_2, N_2} \quad \min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} (N_1 d_1 + N_2 d_2) + \tau(N_1 d_1^2 + N_2 d_2^2) \tag{B.20a}$$

$$\text{s.t.} \quad \sigma N_2 d_2 + \rho \sigma^2 N_2 d_2^2 \leq B(d_1, N_1, \rho) \quad \forall (\rho, \tau) \in Z_{\hat{\rho}, \hat{\tau}} \tag{B.20b}$$

$$d_2 \geq d_{\min} \tag{B.20c}$$

$$N_2 \in \{N_2^{\min}, \dots, N_2^{\max}\}. \tag{B.20d}$$

33

This is a static robust optimization problem. Constraint (B.20b) will hold with equality at the optimum, because it is the only dose-limiting constraint. Solving for d_2 yields the constraint

$$d_2 = g(d_1, N_1, N_2; \rho), \quad \forall (\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}, \tag{B.21}$$

and this is used to rewrite (B.20a) and (B.20b) in terms of functions f and g. Problem (B.20) is equivalent to

$$\max_{N_2} \quad \min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} \quad f(d_1, N_2, \rho, \tau) \tag{B.22a}$$

$$\text{s.t.} \quad g(d_1, N_1, N_2; \rho) \geq d^{\min}, \quad \forall (\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})} \tag{B.22b}$$

$$N_2 \in \{N_2^{\min}, \ldots, N_2^{\max}\}. \tag{B.22c}$$

Similar to the exact case (Section 3), in any worst-case realization it will hold that τ is at its lowest value, so it is sufficient to consider only those observations $(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}$ with $\tau = \hat{\tau}_L$. Additionally, according to Lemma C.3 functions f and g are increasing or decreasing in ρ. Hence, there are two candidate worst-case scenarios: $(\hat{\rho}_L, \hat{\tau}_L)$ and $(\hat{\rho}_U, \hat{\tau}_L)$. We can rewrite (B.22) to

$$\max_{N_2} \quad \min \left\{ f(d_1, N_2, \hat{\rho}_L, \hat{\tau}_L), \ f(d_1, N_2, \hat{\rho}_U, \hat{\tau}_L) \right\} \tag{B.23a}$$

$$\text{s.t.} \quad g(d_1, N_1, N_2; \hat{\rho}_L) \geq d^{\min} \tag{B.23b}$$

$$g(d_1, N_1, N_2; \hat{\rho}_U) \geq d^{\min} \tag{B.23c}$$

$$N_2 \in \{N_2^{\min}, \ldots, N_2^{\max}\}. \tag{B.23d}$$

We distinguish three cases:

- Case $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\max}$: Analogous to the proof of Theorem 1, one can show that for any realization $(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}$ it is optimal to maximize the number of fractions in stage 2. We plug in $N_2^*(\rho, \tau) = N_2^{\max}$ and show that it is feasible. Constraints (B.23b) and (B.23c) reduce to

$$\min \left\{ g(d_1, N_1, N_2^{\max}; \hat{\rho}_L), g(d_1, N_1, N_2^{\max}; \hat{\rho}_U) \right\} \geq d^{\min}, \tag{B.24}$$

which is equivalent to

$$d_1 \leq \min \left\{ g(d^{\min}, N_2^{\max}, N_1; \hat{\rho}_L), g(d^{\min}, N_2^{\max}, N_1; \hat{\rho}_U) \right\}. \tag{B.25}$$

It holds that $\hat{\rho}_L \geq \frac{\hat{\tau}_L}{\sigma} \geq \frac{\tau_L}{\sigma}$, and $\hat{\rho}_U \leq \rho_U$. According to Lemma C.3a function g is either increasing or decreasing in ρ for other arguments fixed. Hence, by Assumption 3 condition (B.25) holds. Hence, $N_2^*(\rho, \tau) = N_2^{\max}$ is feasible and optimal. Thus, (B.22) equals

$$\min \left\{ f(d_1, N_2^{\max}, \hat{\rho}_L, \hat{\tau}_L), \ f(d_1, N_2^{\max}, \hat{\rho}_U, \hat{\tau}_L) \right\}. \tag{B.26}$$

By definition of f, this implies

$$d_2 = \min \left\{ g(d_1, N_1, N_2^{\max}; \hat{\rho}_L), g(d_1, N_1, N_2^{\max}; \hat{\rho}_U) \right\}. \tag{B.27}$$

- Case $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\min}$: Similar to the previous case. Analogous to the proof of Theorem 1, one can show that for any realization $(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}$ it is optimal to minimize the number of fractions in stage 2. We plug in $N_2^*(\rho, \tau) = N_2^{\min}$ and show that it is feasible. Similar to the previous case, constraints (B.23b) and (B.23c) reduce to

$$d_1 \leq \min \left\{ g(d^{\min}, N_2^{\min}, N_1; \hat{\rho}_L), g(d^{\min}, N_2^{\min}, N_1; \hat{\rho}_U) \right\}. \tag{B.28}$$

It holds that $\hat{\rho}_L \geq \rho_L$, and $\hat{\rho}_U \leq \rho_U$. Hence, by Assumption 3, Lemma C.3a and using the fact that function g is decreasing in its second argument, condition (B.28) holds. Hence, $N_2^*(\rho, \tau) = N_2^{\min}$ is feasible and optimal. Similar to the previous case, we find

$$d_2 = \min \left\{ g(d_1, N_1, N_2^{\min}; \hat{\rho}_L), g(d_1, N_1, N_2^{\min}; \hat{\rho}_U) \right\}. \tag{B.29}$$

- Case $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$: The optimal number of fractions in stage-2 is not known a priori. By definition of $Z_{\text{ID}}^{\text{int}}$, it holds that $\hat{\rho}_L \geq \max\{\rho_L, \frac{\tau_L}{\sigma} - 2r^\rho)\}$ and $\hat{\rho}_U \leq \rho_U$. By Assumption 3 it holds that

$$d_1 \leq \min \left\{ g(d^{\min}, N_2^{\max}, N_1; \max\{\rho_L, \frac{\tau_L}{\sigma} - 2r^\rho)\}, g(d^{\min}, N_2^{\max}, N_1; \rho_U) \right\}. \tag{B.30}$$

Lemma C.3a, the fact that function g is decreasing in its third argument and (B.30) together imply that (B.23b) and (B.23c) hold for any feasible N_2. Hence, from problem (B.23) we derive

$$N_2^*(d_1; \hat{\rho}, \hat{\tau}) = \underset{N_2 \in \{N_2^{\min}, \ldots, N_2^{\max}\}}{\arg\max} \min \left\{ f(d_1, N_2, \hat{\rho}_L, \hat{\tau}_L), \ f(d_1, N_2, \hat{\rho}_U, \hat{\tau}_L) \right\}, \tag{B.31}$$

and by definition of f the corresponding value for d_2 is

$$d_2 = \min \left\{ g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_L), g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_U) \right\}. \tag{B.32}$$

Combining the above three cases, we arrive at the optimal decision rules (32) and (33) for fixed d_1.

B.4 Extra analysis to Section 4

This analysis makes use of the lemmas in Appendix C. Consider problem (29). For given d_1, the optimal stage-2 decision rules are given by Theorem 3. As stated in Section 4, we split the uncertainty set Z into three subsets. This enables us to exploit the fact that depending on $(\hat{\rho}, \hat{\tau})$ the value $N_2^*(d_1; \hat{\rho}, \hat{\tau})$ may be known in advance. The split (31) is repeated here for convenience

$$Z_{\mathrm{ID}}^{\min} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \geq \sigma \hat{\rho}_U\} \tag{B.33a}$$

$$Z_{\mathrm{ID}}^{\mathrm{int}} = \{(\hat{\rho}, \hat{\tau}) \in Z : \sigma \hat{\rho}_L < \hat{\tau}_L < \sigma \hat{\rho}_U\} \tag{B.33b}$$

$$Z_{\mathrm{ID}}^{\max} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \leq \sigma \hat{\rho}_L\}, \tag{B.33c}$$

so that $Z = Z_{\mathrm{ID}}^{\min} \cup Z_{\mathrm{ID}}^{\mathrm{int}} \cup Z_{\mathrm{ID}}^{\max}$. The associated sets of observation-realization pairs $(\rho, \tau, \hat{\rho}, \hat{\tau})$ are given by

$$U^i = U \cap \{(\rho, \tau, \hat{\rho}, \hat{\tau}) : (\hat{\rho}, \hat{\tau}) \in Z_{\mathrm{ID}}^i\}, \quad i \in \{\min, \mathrm{int}, \max\}, \tag{B.34}$$

so it holds that $U = U^{\min} \cup U^{\mathrm{int}} \cup U^{\max}$. Set U^i can be interpreted as the set of observation-realization pairs for which the observation $(\hat{\rho}, \hat{\tau})$ is in set Z_{ID}^i. Figure 5 illustrates the subsets Z_{ID}^i. Set $U^{\min}$ consists of those observation-realization pairs $(\rho, \tau, \hat{\rho}, \hat{\tau})$ for which $N_2^*(d_1; \hat{\rho}, \hat{\tau}) = N_2^{\max}$. If $(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U^{\mathrm{int}}$, then based on the observation $(\hat{\rho}, \hat{\tau})$ it is unclear what fractionation is worst-case optimal. Last, if $(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U^{\max}$ we know $N_2^*(d_1; \hat{\rho}, \hat{\tau}) = N_2^{\min}$. Problem (29) is equivalent to

$$\max_{d_1, q} \quad q \tag{B.35a}$$

$$\text{s.t.} \quad q \leq f(d_1, N_2^{\min}; \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U^{\min} \tag{B.35b}$$

$$q \leq f(d_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U^{\mathrm{int}} \tag{B.35c}$$

$$q \leq f(d_1, N_2^{\max}; \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U^{\max} \tag{B.35d}$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{B.35e}$$

Similar to the exact case (Section 3), in any worst-case realization it will hold that $\tau = \tau_L$. Therefore, any observation with $\hat{\tau} - r^\tau > \tau_L$ cannot yield the worst-case realization. Define

$$U_L^i = U_i \cap \{(\rho, \tau, \hat{\rho}, \hat{\tau}) : \hat{\tau} - r^\tau \leq \tau_L\}, \quad i \in \{\min, \mathrm{int}, \max\}, \tag{B.36}$$

which is the subset of U^i of observation-realization pairs for which τ_L is a possible realization of τ. Constraints (B.35b)-(B.35e) can be replaced by

$$q \leq f(d_1, N_2^{\min}; \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\min} \tag{B.37a}$$

$$q \leq f(d_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\mathrm{int}} \tag{B.37b}$$

$$q \leq f(d_1, N_2^{\max}; \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\max} \tag{B.37c}$$

For (B.37a) and (B.37c) it remains to find the worst-case realization of ρ for which the observation-realization pair is in $U_L^{\min}$ and $U_L^{\max}$, respectively. According to Lemma C.3b, function f is increasing or decreasing in ρ for fixed d_1, so it is sufficient to check the maximum and minimum realization of ρ for which the observation-realization pair is in those sets. These are

$$\min\{\rho : (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\min}\} = \rho_L, \quad \max\{\rho : (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\min}\} = \frac{\tau_L}{\sigma} \tag{B.38a}$$

$$\min\{\rho : (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\max}\} = \frac{\tau_L}{\sigma}, \quad \max\{\rho : (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{\max}\} = \rho_U. \tag{B.38b}$$

Plugging in $\rho = \frac{\tau_L}{\sigma}$ in (B.37a) and (B.37c) yields $q \leq K$, with K defined by (19). Lemma C.5 provides a conservative approximation of constraint (B.37b). Putting everything together, the optimum of the following problem is a lower bound to the optimum of (B.35) (or, equivalently, (29)):

$$\max_{d_1, q} \quad q \tag{B.39a}$$

$$\text{s.t.} \quad q \leq f(d_1, N_2^{\min}; \rho_L, \tau_L) \tag{B.39b}$$

$$q \leq f(d_1, N_2^{\max}; \rho_U, \tau_L) \tag{B.39c}$$

$$q \leq K \tag{B.39d}$$

$$q \leq p(d_1) \tag{B.39e}$$

$$d_{\min} \leq d_1 \leq d_1^{\max}, \tag{B.39f}$$

with $p(d_1)$ defined by (C.32) in Appendix C. Constraint (B.39e) is the only conservative constraint, all other constraints are exact reformulations. In particular, this means that if for a solution the objective value equals K, it is certain that this is an optimal solution. It is easy to obtain other straightforward conservative approximations of (B.35c). For instance, a policy that delivers $N_2^{\min}$ or $N_2^{\max}$ fractions (or any number in between, for that matter) for any observation $(\hat{\rho}, \hat{\tau}) \in Z_{\text{ID}}^{\text{int}}$ is a conservative approximation. However, these perform less good and do not use all available information, as explained in the proof of Lemma C.5.

C   Extra lemmas

This appendix states and proves several frequently used properties of functions g and f.

Lemma C.1 (Convexity/concavity f w.r.t. d_1). *Let $\rho > 0$, $\tau > 0$ and $N_1, N_2 \in \mathbb{N}_+$. Let $d_1 \in [0, g(0, 0, N_1; \rho)]$. The following properties hold for function f:*

- *Function $f(d_1, N_2; \rho, \tau)$ is strictly convex in d_1 if $\tau > \rho\sigma$, with unique minimizer $g(0, 0, N_1 + N_2; \rho)$;*

- *Function $f(d_1, N_2; \rho, \tau)$ is strictly concave in d_1 if $\tau < \rho\sigma$, with unique maximizer $g(0, 0, N_1 + N_2; \rho)$;*

- *Function $f(d_1, N_2; \rho, \tau)$ is constant in d_1 if $\tau = \rho\sigma$, with value $\frac{1}{\sigma}B(0, 0, \frac{\tau}{\sigma})$.*

Proof. The partial derivative of f w.r.t. d_1 is given by

$$\frac{\partial f(d_1, N_2; \rho, \tau)}{\partial d_1} = N_1 + N_2 \frac{\partial g(d_1, N_1, N_2; \rho)}{\partial d_1} + \tau\left(2N_1 d_1 + 2N_2 g(d_1, N_1, N_2; \rho)\frac{\partial g(d_1, N_1, N_2; \rho)}{\partial d_1}\right), \tag{C.1}$$

where the partial derivative of g w.r.t. d_1 is given by

$$\frac{\partial g(d_1, N_1, N_2; \rho)}{\partial d_1} = -\frac{1}{N_2}(N_1 + 2N_1 d_1 \sigma\rho)\left(1 + \frac{4\rho}{N_2}B(d_1, N_1; \rho)\right)^{-\frac{1}{2}}. \tag{C.2}$$

Define $h(d_1, N_2; \rho) = 1 + 4\frac{\rho}{N_2}B(d_1, N_1; \rho)$. Then, plugging (C.2) in (C.1), we obtain

$$
\begin{aligned}
\frac{\partial f(d_1, N_2; \rho, \tau)}{\partial d_1} = {} & (N_1 - (N_1 + 2N_1 d_1 \sigma\rho)h(d_1, N_2; \rho)^{-\frac{1}{2}} \\
& + \tau\left(2N_1 d_1 - \frac{2}{N_2}(N_1 + 2N_1 d_1 \sigma\rho)h(d_1, N_2; \rho)^{-\frac{1}{2}}N_2 \frac{-1 + h(d_1, N_2; \rho)^{\frac{1}{2}}}{2\sigma\rho}\right) \\
= {} & \frac{N_1}{\sigma\rho}\left(h(d_1, N_2; \rho)^{-\frac{1}{2}}(2\sigma\rho d_1 + 1) - 1\right)(\tau - \rho\sigma).
\end{aligned}
$$

Further elementary math shows that $h(d_1, N_2; \rho)^{-\frac{1}{2}}(2\sigma\rho d_1 + 1) - 1 = 0$ if and only if $d_1 = g(0, 0, N_1 + N_2; \rho)$. For the second derivative of f w.r.t. d_1 we obtain:

$$
\begin{aligned}
\frac{\partial^2 f(d_1, N_2; \rho, \tau)}{\partial d_1^2} = {} & \left(\frac{\tau - \rho\sigma}{\sigma\rho}N_1\right)\frac{\partial}{\partial d_1}h(d_1, N_2; \rho)^{-\frac{1}{2}}(2\sigma\rho d_1 + 1) \\
= {} & \left(\frac{\tau - \rho\sigma}{\sigma\rho}N_1\right)\left[h(d_1, N_2; \rho)^{-\frac{1}{2}}2\sigma\rho + \frac{2\rho}{N_2}(2\sigma\rho d_1 + 1)h(d_1, N_2, \rho)^{-\frac{3}{2}}(\sigma N_1 + 2\rho\sigma^2 d_1 N_1)\right],
\end{aligned}
$$

and the second part of this product is positive. Hence, its sign depends only on the term $\tau - \rho\sigma$. Combining the result for the first and second derivative, we obtain

- Function $f(d_1, N_2; \rho, \tau)$ is strictly convex in d_1 if $\tau > \rho\sigma$, with unique minimizer $g(0, 0, N_1 + N_2; \rho)$;
- Function $f(d_1, N_2; \rho, \tau)$ is strictly concave in d_1 if $\tau < \rho\sigma$, with unique maximizer $g(0, 0, N_1 + N_2; \rho)$;
- Function $f(d_1, N_2; \rho, \tau)$ is constant in d_1 otherwise.

If $\tau = \rho\sigma$, we can rewrite $f(d_1, N_2; \frac{\tau}{\sigma}, \tau)$ to

$$f(d_1, N_2; \frac{\tau}{\sigma}, \tau) = \max_{d_2} \left\{ d_1 N_1 + d_2 N_2 + \tau(d_1^2 N_1 + d_2^2 N_2) \mid \sigma(d_1 N_1 + d_2 N_2) + \rho\sigma^2(d_1^2 N_1 + d_2^2 N_2) \leq B(0, 0; \rho) \right\}$$

$$= \max_{d_2} \left\{ d_1 N_1 + d_2 N_2 + \tau(d_1^2 N_1 + d_2^2 N_2) \mid \sigma(d_1 N_1 + d_2 N_2) + \rho\sigma^2(d_1^2 N_1 + d_2^2 N_2) = B(0, 0; \rho) \right\}$$

$$= \max_{d_2} \left\{ d_1 N_1 + d_2 N_2 + \tau(d_1^2 N_1 + d_2^2 N_2) \mid d_1 N_1 + d_2 N_2 + \tau(d_1^2 N_1 + d_2^2 N_2) = \frac{1}{\sigma} B(0, 0, \frac{\tau}{\sigma}) \right\}$$

$$= \frac{1}{\sigma} B(0, 0, \frac{\tau}{\sigma}).$$

$\square$

Define

$$d_1^-(N_2) = \begin{cases} \dfrac{\varphi D - \varphi D \left(1 + (N_1 + N_2)\frac{N_2 - T}{N_1 T}\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)} & \text{if } N_1 + N_2 \geq T \wedge N_1 \leq T \\ -\infty & \text{otherwise} \end{cases} \tag{C.3a}$$

$$d_1^+(N_2) = \begin{cases} \dfrac{\varphi D + \varphi D \left(1 + (N_1 + N_2)\frac{N_2 - T}{N_1 T}\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)} & \text{if } N_1 + N_2 \geq T \wedge N_2 \leq T \\ +\infty & \text{otherwise.} \end{cases} \tag{C.3b}$$

If two functions f with equal N_2 but different ρ intersect, d_1 takes value $d_1^-(N_2)$ or $d_1^+(N_2)$. The following lemma provides information on the existence and location of these intersection points. We consider only those values for d_1 where function $f(d_1, N_2; \rho, \tau)$ is finite for all $(\rho, \tau) \in Z$. Let

$$d_{\text{UB}} = \min_{(\rho, \tau) \in Z} g(0, 0, N_1; \rho). \tag{C.4}$$

Lemma C.2 (Properties d_1^- and d_1^+). *Let $N_1, T \in \mathbb{N}_+$.*

(a) *Let $N_2 \in \mathbb{N}_+$. If $\rho_1 \neq \rho_2$, the equation*

$$f(d_1, N_2; \rho_1, \tau) = f(d_1, N_2; \rho_2, \tau), \tag{C.5}$$

has the following real roots for d_1 on the interval $[0, d_{UB}]$:

• $d_1^-(N_2)$ and $d_1^+(N_2)$	*if $N_1 + N_2 \geq T$, $N_2 \leq T$ and $N_1 \leq T$*	(C.6a)
• $d_1^-(N_2)$	*if $N_1 + N_2 \geq T$, $N_2 \leq T$ and $N_1 > T$*	(C.6b)
• $d_1^+(N_2)$	*if $N_1 + N_2 \geq T$, $N_2 > T$ and $N_1 \leq T$*	(C.6c)
• *no roots on interval*	*if $N_1 + N_2 \geq T$, $N_2 > T$ and $N_1 > T$*	(C.6d)
• *no real roots*	*otherwise.*	(C.6e)

(b) *Let $N_2^A, N_2^B \in \mathbb{N}_+$ such that $N_2^A < N_2^B$. It holds that*

   (i) *If $d_1^-(N_2^A)$ and $d_1^-(N_2^B)$ are both finite, then $d_1^-(N_2^A) > d_1^-(N_2^B)$;*

   (ii) *If $d_1^+(N_2^A)$ and $d_1^+(N_2^B)$ are both finite, then $d_1^+(N_2^A) \leq d_1^+(N_2^B)$.*

Proof. Both parts of the lemma are proved individually.

Proof Lemma C.2a

By definition of f, the equation $f(d_1, N_2; \rho_1, \tau) = f(d_1, N_2; \rho_2, \tau)$ reduces to $g(d_1, N_1, N_2; \rho_1) = g(d_1, N_1, N_2; \rho_2)$

with $d_1 \in [0, \min\{g(0, 0, N_1; \rho_1), g(0, 0, N_1; \rho_2)\}]$. By construction of g, this means we are interested in the pairs (d_1, d_2) that solve the system

$$\sigma(N_1 d_1 + N_2 d_2) + \rho_1 \sigma^2(N_1 d_1^2 + N_2 d_2^2) = \varphi D(1 + \rho_1 \frac{D}{T} \varphi) \tag{C.7a}$$

$$\sigma(N_1 d_1 + N_2 d_2) + \rho_2 \sigma^2(N_1 d_1^2 + N_2 d_2^2) = \varphi D(1 + \rho_2 \frac{D}{T} \varphi) \tag{C.7b}$$

$$d_1 \geq 0, \ d_2 \geq 0. \tag{C.7c}$$

We subtract $\frac{\rho_2}{\rho_1}$ times (C.7a) from (C.7b) and solve for d_1 to obtain

$$d_1 = \frac{\varphi D - \sigma N_2 d_2}{\sigma N_1}. \tag{C.8}$$

We know that $d_2 = g(d_1, N_1, N_2; \rho_1)$. Plugging in (C.8) in this expression and simplifying yields the following roots for d_2:

$$r_2^-(N_2) = \frac{\varphi D + \varphi D \left(1 + (N_1 + N_2)\left(\frac{N_1}{T} - 1\right)/N_2\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)} \tag{C.9a}$$

$$r_2^+(N_2) = \frac{\varphi D - \varphi D \left(1 + (N_1 + N_2)\left(\frac{N_1}{T} - 1\right)/N_2\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)}. \tag{C.9b}$$

Plugging (C.9) in (C.8) and simplifying yields the following roots for d_1:

$$r_1^-(N_2) = \frac{\varphi D - \varphi D \left(1 + (N_1 + N_2)\frac{N_2 - T}{N_1 T}\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)} \tag{C.10a}$$

$$r_1^+(N_2) = \frac{\varphi D + \varphi D \left(1 + (N_1 + N_2)\frac{N_2 - T}{N_1 T}\right)^{\frac{1}{2}}}{\sigma(N_1 + N_2)}. \tag{C.10b}$$

These roots need not be real-valued, nor in the interval $[0, \min\{g(0, 0, N_1; \rho_1), g(0, 0, N_1; \rho_2)\}]$. For both $r_1^-(N_2)$ and $r_1^+(N_2)$ to be real-valued, we require that

$$1 + (N_1 + N_2)\left(\frac{N_1}{T} - 1\right)/N_2 \geq 0,$$

which reduces to $N_1 + N_2 \geq T$. Furthermore, for nonnegativity of $r_1^-(N_2)$ and $r_1^+(N_2)$ it suffices to check nonnegativity of the former. This is equivalent to

$$\varphi D - \sigma N_2 d_2^- \geq 0,$$

which reduces to $N_2 \leq T$. Moreover, it needs to hold that $r_1^+(N_2) \leq g(0, 0, N_1; \rho_1)$ and $r_1^+(N_2) \leq g(0, 0, N_1; \rho_2)$. This is equivalent to $r_2^+(N_2) \geq 0$, which can be rewritten to

$$1 + (N_1 + N_2)\left(\frac{N_1}{T} - 1\right)/N_2 \leq 1,$$

and this reduces to $N_1 \leq T$. Parameters $d_1^-(N_2)$ resp. $d_1^+(N_2)$ (see (C.3)) take the values of $r_1^-(N_2)$ resp. $r_1^-(N_2)$ if they are a root of (C.5), and $-\infty$ resp. $+\infty$ otherwise. All together, we obtain the cases in (C.6).

It remains to show that the obtained roots are in the interval $[0, d_{\text{UB}}]$. It is already shown that, if they are (real-valued) roots to (C.5), then $d_1^-(N_2), d_1^+(N_2)$ are nonnegative. Furthermore, in that case $d_1^-(N_2) \leq d_1^+(N_2)$. It holds that

$$\frac{\partial g(0, 0, N_1; \rho)}{\partial \rho} \leq 0 \Leftrightarrow N_1 \leq T.$$

Hence, if $d_1^+(N_2)$ is a real-valued root to (C.5) it follows that

$$d_{\text{UB}} = \min_{(\rho, \tau) \in Z} g(0, 0, N_1; \rho) \geq \lim_{\rho \to +\infty} g(0, 0, N_1; \rho) = \frac{\varphi D}{\sigma \sqrt{N_1 T}} \geq d_1^+(N_2),$$

where the second equality follows from the definition of g. This implies that indeed $d_1^-(N_2), d_1^+(N_2) \in [0, d_{\text{UB}}]$.

Proof Lemma C.2b
Assume $N_2^A, N_2^B \in \mathbb{N}_+$ such that $N_2^A \leq N_2^B$, and assume $N_1 + N_2^A \geq T$. Statements (i) and (ii) are proved individually.

Proof part (i)
Assume $d_1^-(N_2^A)$ and $d_1^-(N_2^B)$ are both finite. The denominator of $d_1^-(N_2)$ (see (C.3a)) is increasing in N_2. The derivative (w.r.t. N_2) of the part within the square root in the numerator of (C.3a) is given by $(N_1 T)^{-1}(N_1 + 2N_2 - T) \geq 0$, because $N_1 + N_2 \geq T$. Hence, the numerator is decreasing in N_2, while the denominator is increasing in N_2. This implies $d_1^-(N_2^A) > d_1^-(N_2^B)$.

Proof part (ii)
Assume $d_1^+(N_2^A)$ and $d_1^+(N_2^B)$ are both finite. One can show that

$$\frac{\partial d_1^+(N_2)}{\partial N_2} = \varphi D \frac{(N_1 + N_2)(N_1 + 2N_2 - T) - 2N_1 T \sqrt{\frac{N_2(N_1 + N_2 - T)}{N_1 T}} - 2N_2(N_1 + N_2 - T)}{2N_1 T \sigma (N_1 + N_2)^2 \sqrt{\frac{N_2(N_1 + N_2 - T)}{N_1 T}}}. \tag{C.11}$$

This implies

$$\frac{\partial d_1^+(N_2)}{\partial N_2} \geq 0$$

$$\Leftrightarrow (N_1 + N_2)(N_1 + 2N_2 - T) - 2N_1 T \sqrt{\frac{N_2(N_1 + N_2 - T)}{N_1 T}} - 2N_2(N_1 + N_2 - T) \geq 0$$

$$\Leftrightarrow \frac{N_1(N_1 + N_2 - T) + TN_2}{2N_1 T} \geq \sqrt{\frac{N_2(N_1 + N_2 - T)}{N_1 T}}$$

$$\Leftrightarrow \frac{\left(N_1(N_1 + N_2 - T) + TN_2\right)^2}{4N_1^2 T^2} \geq \frac{N_2(N_1 + N_2 - T)}{N_1 T}$$

$$\Leftrightarrow \frac{(N_1 - T)^2 (N_1 + N_2)^2}{4N_1^2 T^2} \geq 0, \tag{C.12}$$

where the fourth line is obtained by using the fact that $N_1 + N_2 \geq T$, and squaring on both sides. The last line follows from simple algebraic manipulations. Condition (C.12) clearly holds, so $d_1^+(N_2^A) \leq d_1^+(N_2^B)$. □

Lemma C.3 (Derivative f and g w.r.t. ρ). *Let $(\rho, \tau) \in Z$.*

(a) *Let $N', N'' \in \mathbb{N}_+$. Let $d' \in [0, d_{UB}]$. If $N' + N'' < T$, then*

$$\frac{\partial g(d', N', N''; \rho)}{\partial \rho} < 0 \quad \text{for all } d' \in [0, d_{UB}]. \tag{C.13}$$

If $N' + N'' \geq T$, then

$$\frac{\partial g(d', N', N''; \rho)}{\partial \rho} \begin{cases} < 0 & \text{if } d' \in [0, d_1^-(N'')) \cup (d_1^+(N''), d_{UB}] \\ = 0 & \text{if } d' \in [0, d_{UB}] \cap \{d_1^-(N''), d_1^+(N'')\} \\ > 0 & \text{if } d' \in [0, d_{UB}] \cap (d_1^-(N''), d_1^+(N'')). \end{cases} \tag{C.14}$$

(b) *Let $N_1, N_2 \in \mathbb{N}_+$. Let $d_1 \in [0, d_{UB}]$. If $N_1 + N_2 < T$, then*

$$\frac{\partial f(d_1, N_2; \rho, \tau)}{\partial \rho} < 0 \quad \text{for all } d_1 \in [0, d_{UB}]. \tag{C.15}$$

If $N_1 + N_2 \geq T$, then

$$\frac{\partial f(d_1, N_2; \rho, \tau)}{\partial \rho} \begin{cases} < 0 & \text{if } d_1 \in [0, d_1^-(N_2)) \cup (d_1^+(N_2), d_{UB}] \\ = 0 & \text{if } d_1 \in [0, d_{UB}] \cap \{d_1^-(N_2), d_1^+(N_2)\} \\ > 0 & \text{if } d_1 \in [0, d_{UB}] \cap (d_1^-(N_2), d_1^+(N_2)). \end{cases} \tag{C.16}$$

Proof. We first prove Lemma C.3a, after that the result of Lemma C.3b is easily obtained.

It holds that

$$\frac{\partial g(d', N', N''; \rho)}{\partial \rho} = \frac{\sqrt{1 + \frac{4\rho}{N''}B(d', N'; \rho)} - 1 + \frac{2\rho}{N''}(d'N'\sigma - \varphi D)}{2\sigma\rho^2\sqrt{1 + \frac{4\rho}{N''}B(d', N'; \rho)}}, \tag{C.17}$$

so

$$\frac{\partial g(d', N', N''; \rho)}{\partial \rho} \geq 0 \Leftrightarrow \sqrt{(N'')^2 + 4\rho N''B(d', N'; \rho)} \geq N'' + 2\rho(\varphi D - d'N'\sigma). \tag{C.18}$$

We distinguish 2 cases:

- $\varphi D \geq d'N'\sigma$. In this case, squaring (C.18) on both sides and simplifying results in

$$-\sigma^2 N'(N' + N'')d'^2 + 2\varphi DN'\sigma d' + \left(\frac{N''}{T} - 1\right)\varphi^2 D^2 \geq 0, \tag{C.19}$$

which is a condition independent of ρ. If $N' + N'' < T$, this inequality has no roots for d', and (C.19) holds for all $d' \in [0, \frac{\varphi D}{N'\sigma}]$. If $N' + N'' \geq T$ one can verify that $d_1 = d_1^-(N'')$ and $d_1 = d_2^+(N'')$ are the roots of this concave parabola if they are finite. The smaller root, $d_1^-(N'')$, is finite if and only if $N'' \leq T$. The larger root, $d_1^+(N'')$ is finite if and only if $N' \leq T$.

- $\varphi D < d'N'\sigma$. In this case, $B(d', N'; \rho) > 0$ only if $N' > T$. In this case, the delivered dose exceeds the dose that is used to set the BED tolerance, which is only possible if the number of fractions N' is strictly larger than the reference number of fractions T. Condition (C.18) clearly holds, so $g(d', N', N''; \rho)$ is increasing in ρ. Using the fact that $N' > T$ it is easily shown that $d_1^- < \frac{\varphi D}{\sigma N'} < d'$. Additionally, it can be shown that $d_{\text{UB}} < d_1^+(N'')$. Hence this case satisfies (C.14). Putting all of the above together yields the required result for g, i.e., Lemma C.3a.

The partial derivative of f w.r.t. ρ is given by

$$\frac{\partial f(d_1, N_2; \rho, \tau)}{\partial \rho} = \frac{\partial g(d_1, N_1, N_2; \rho)}{\partial \rho}\left(N_2 + 2\tau N_2 g(d_1, N_1, N_2; \rho)\right). \tag{C.20}$$

Hence, the sign of the partial derivative of f w.r.t. ρ is equal to the sign of the partial derivative of g w.r.t. ρ. The result of Lemma C.3b immediately follows. $\square$

For given (ρ, τ) such that $\tau \neq \sigma\rho$, define the *twin point* of $d_1 \in W(\rho, \tau)$ as

$$t(d_1; \rho, \tau) := \frac{(N_1 - N_2^*(\rho, \tau))d_1 + 2N_2^*(\rho, \tau)g(d_1, N_1, N_2^*(\rho, \tau); \rho)}{N_1 + N_2^*(\rho, \tau)}, \tag{C.21}$$

where

$$W(\rho, \tau) := \left[\max\{0, t(g(0, 0, N_1; \rho); \rho, \tau)\}, \ \min\{t(0; \rho, \tau), g(0, 0, N_1; \rho)\}\right] \backslash \{g(0, 0, N_1 + N_2; \rho)\}. \tag{C.22}$$

Figure C.1 illustrates the relation between d_1 and $t(d_1; \rho, \tau)$. Set W can be interpreted as the points d_1 for which there exists another point the graph of f that has the same value, we refer to such points as twin points. The following lemma proves that for fixed (ρ, τ) any d_1 in the set $W(\rho, \tau)$ has a twin point $t(d_1; \rho, \tau)$ that is also in the set $W(\rho, \tau)$, and their objective values are equal.

Lemma C.4. *Let $(\rho, \tau) \in Z$ such that $\tau \neq \sigma\rho$, let $N_2 = N_2^*(\rho, \tau)$ and let $d_1 \in [0, g(0, 0, N_1; \rho)]$. The equation*

$$f(d_1, N_2; \rho, \tau) = f(d_1', N_2; \rho, \tau) \tag{C.23}$$

has a solution $d_1' \in [0, g(0, 0, N_1; \rho)]$ unequal to d_1 if and only if $d_1 \in W(\rho, \tau)$. In that case, there is a unique solution $d_1' = t(d_1; \rho, \tau) \in W(\rho, \tau)$, and it holds that $d_1 = t(t(d_1; \rho, \tau); \rho, \tau)$.

Proof. Let (ρ, τ) such that $\tau \neq \sigma\rho$, let $N_2 = N_2^*(\rho, \tau)$. We first show that if there exists a solution d_1' to (C.23) unequal to d_1, then this solution is $d_1' = t(d_1; \rho, \tau)$ and that $d_1 = t(t(d_1; \rho, \tau); \rho, \tau)$. Subsequently,

(a) Case $t(0; \rho, \tau) < g(0, 0, N_1; \rho)$.
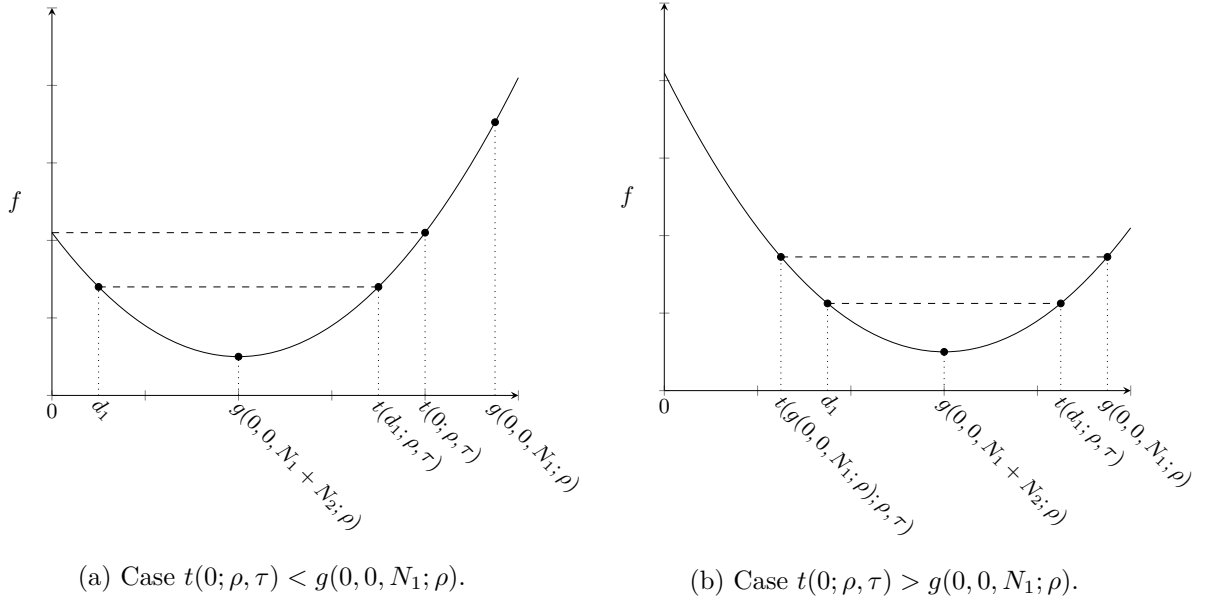
(b) Case $t(0; \rho, \tau) > g(0, 0, N_1; \rho)$.

Figure C.1: Illustration of d_1 and $t(d_1; \rho, \tau)$, for convex f.

we show that this solution exists only if $d_1 \in W(\tau, \rho)$ and that in that case also $t(d_1; \rho, \tau) \in W(\rho, \tau)$.

First part:
Let $L, Q \in \mathbb{R}_+$ denote the linear and quadratic contribution of d_1 to f, i.e.,

$$L(d_1, N_2, \rho) := N_1 d_1 + N_2 g(d_1, N_1, N_2; \rho) \tag{C.24a}$$

$$Q(d_1, N_2, \rho) := N_1 d_1^2 + N_2 g(d_1, N_1, N_2; \rho)^2. \tag{C.24b}$$

We show that for any $d_1 \in [0, g(0, 0, N_1; \rho)]$ unequal to $g(0, 0, N_1 + N_2, \rho)$ there is a unique solution d_1' such that

$$L(d_1, N_2; \rho) = L(d_1', N_2; \rho) \tag{C.25a}$$

$$Q(d_1, N_2; \rho) = Q(d_1', N_2; \rho). \tag{C.25b}$$

Any such d_1' has exactly the same objective value as d_1. Plug (C.24) in (C.25), rewrite the first equation to eliminate g, and plug this in the second equation to get

$$
\begin{aligned}
d_1' &= \frac{L(d_1, N_2; \rho) \pm \sqrt{\frac{N_2}{N_1}\left((N_1 + N_2)Q(d_1, N_2; \rho) - L^2(d_1, N_2; \rho)\right)}}{N_1 + N_2} \\
&= \frac{N_1 d_1 + N_2 g(d_1, N_1, N_2; \rho) \pm N_2(d_1 - g(d_1, N_1, N_2; \rho))}{N_1 + N_2}.
\end{aligned}
\tag{C.26}
$$

The '+' solution to (C.26) returns $d_1' = d_1$, and the '−' solution returns

$$d_1' = \frac{(N_1 - N_2)d_1 + 2N_2 g(d_1, N_1, N_2, \rho)}{N_1 + N_2}, \tag{C.27}$$

and we denote this solution by $t(d_1; \rho, \tau)$. By construction, it holds that $d_1 = t(t(d_1; \rho, \tau); \rho, \tau)$. Because $\tau \neq \sigma\rho$, function $f(d_1, N_2; \rho, \tau)$ is a strictly convex or concave function according to Lemma C.1, so $f(d_1, N_2; \rho, \tau) = z$ for some constant $z \in \mathbb{R}$ has either 0, 1 or 2 solutions. In particular, $d_1 = t(d_1; \rho, \tau)$ if and only if d_1 equals minimizer $g(0, 0, N_1 + N_2; \rho)$. Hence if there exists a solution d_1' to (C.23) unequal to d_1, then this solution is $d_1' = t(d_1; \rho, \tau)$.

Second part:

- Suppose $d_1 \notin W(\rho, \tau)$. We distinguish three cases. Case (i): '$d_1 = g(0, 0, N_1 + N_2; \rho)$'. Because this is the unique minimizer of f, there does not exist a d_1' with equal objective value. Case (ii).

41

'$d_1 > \min\{t(0; \rho, \tau),\ g(0, 0, N_1; \rho)\}$'. Because $d_1 \in [0, g(0, 0, N_1; \rho)]$, this implies $d_1 > t(0, \rho, \tau)$. As shown in the first part of the proof, it holds that $d_1 = t(t(d_1; \rho, \tau); \rho, \tau)$. Hence, $d_1 > t(0; \rho, \tau)$ is equivalent to $t(t(d_1; \rho, \tau), \rho, \tau) > t(0; \rho, \tau)$. Because $t(d_1; \rho, \tau)$ is decreasing in d_1, this implies $t(d_1; \rho, \tau) < 0$, so according to (B.1c) it holds that $f(t(d_1; \rho, \tau), N_2; \rho, \tau) = -\infty$ and we have a contradiction. Case (iii): '$d_1 < \max\{0,\ t(g(0, 0, N_1; \rho); \rho, \tau)\}$'. Similar to case (ii), one can show that $f(t(d_1; \rho, \tau), N_2; \rho, \tau) = -\infty$.

- Suppose $d_1 \in W(\rho, \tau)$. From (C.26) one can see that, because the term $d_1 - g(d_1, N_1, N_2, \rho)$ is increasing in d_1, the function $t(d_1; \rho, \tau)$ is decreasing in d_1. Consequently,

$$d_1 \leq \min\{t(0; \rho, \tau), g(0, 0, N_1; \rho)\} \Leftrightarrow d_1' \geq \max\{0, t(g(0, 0, N_1; \rho); \rho, \tau)\}. \tag{C.28}$$

Furthermore, using the same argument,

$$d_1 \geq \max\{0, t(g(0, 0, N_1; \rho); \rho, \tau)\} \Leftrightarrow d_1' \leq \min\{t(0; \rho, \tau), g(0, 0, N_1; \rho)\}. \tag{C.29}$$

Therefore, it holds that $d_1' \in W(\rho, \tau)$.

$\square$

In the following lemma, let $I(\cdot|S)$ denote the indicator function for a set S:

$$I(x|S) = \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{otherwise.} \end{cases} \tag{C.30}$$

Lemma C.5. *For given $q \in \mathbb{R}_+$ and given $d_1 \in [0, d_{UB}]$,*

$$q \leq f(d_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{int}, \tag{C.31}$$

holds if $\tau \leq p(d_1)$, with

$$p(d_1) = \sum_{\eta \in \{N_2^{min}, \ldots, N_2^{max}\}} \max\left\{f(d_1, \eta; \rho_L^{int}, \tau_L),\ f(d_1, \eta - 1; \rho_U^{int}, \tau_L)\right\} I(d_1|S_\eta) \tag{C.32}$$
$$+ f(d_1, N_2^{min}; \rho_L^{int}, \tau_L) I(d_1|S_{min}) + f(d_1, N_2^{max}; \rho_U^{int}, \tau_L) I(d_1|S_{max}),$$

where sets S_{min}, S_{max} and S_η are defined in (C.37a), (C.37d) and (C.42), respectively, and ρ_L^{int}, ρ_U^{int} are defined in (C.44).

Proof. By definition of $N_2^*(d_1; \hat{\rho}, \hat{\tau})$ and U_L^{int}, it holds that

$$q \leq f(d_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \rho, \tau), \quad \forall (\rho, \tau, \hat{\rho}, \hat{\tau}) \in U_L^{int}, \tag{C.33}$$

is equivalent to

$$q \leq \max_{\tilde{\eta} \in \{N_2^{min}, \ldots, N_2^{max}\}} \min\{f(d_1, \tilde{\eta}; \hat{\rho}_L, \hat{\tau}_L), f(d_1, \tilde{\eta}; \hat{\rho}_U, \hat{\tau}_L)\}, \quad \forall (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{int} \cap \{(\hat{\rho}, \hat{\tau}) : \hat{\tau} \leq \tau_L + r^\tau\}, \tag{C.34}$$

and because function f is increasing in τ, we need to consider only those observations $(\hat{\rho}, \hat{\tau})$ with $\hat{\tau}_L = \tau_L$. For the first part of the proof, we fix the observation $(\hat{\rho}, \hat{\tau})$, plug in $\hat{\tau}_L = \tau_L$, and rewrite (C.34) for this fixed observation.

Because $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{int}$, it holds that $\sigma\hat{\rho}_L < \tau_L < \sigma\hat{\rho}_U$. Hence, by Lemma C.1, function $f(d_1, \eta, \hat{\rho}_L, \tau_L)$ is convex and $f(d_1, \eta, \hat{\rho}_U, \tau_L)$ is concave in d_1 for any $\eta \in \mathbb{N}_+$. We make use of results of Lemma C.2. Define

$$E^- = \{\eta : N_1 + \eta \geq T, \eta \leq T\} \cap \{N_2^{min}, \ldots, N_2^{max}\} \tag{C.35a}$$
$$E^+ = \{\eta : N_1 + \eta \geq T, N_1 \leq T\} \cap \{N_2^{min}, \ldots, N_2^{max}\}, \tag{C.35b}$$

and let η_{min}^-, η_{max}^-, η_{min}^+ and η_{max}^+, denote the smallest and largest elements of E^- and E^+, respectively. If $\eta \in E^-$ respectively $\eta \in E^+$, then, according to Lemma C.2a, $d_1 = d_1^-(\eta)$ respectively $d_1 = d_1^+(\eta)$ is a nonnegative real root of

$$f(d_1, \eta; \hat{\rho}_L, \tau_L) = f(d_1, \eta; \hat{\rho}_U, \tau_L),$$

and the corresponding objective value equals K. From Lemma C.2b we know that

$$d_1^-(N_2^{\max}) < \ldots < d_1^-(N_2^{\min}) \le d_1^+(N_2^{\min}) \le \ldots \le d_1^+(N_2^{\max}). \tag{C.36}$$

We use this to split the domain $[0, d_{\mathrm{UB}}]$ as follows:

$$S_{\min} = \begin{cases} (d_1^-(\eta_{\min}^-), d_1^+(\eta_{\min}^+)) & \text{if } E^- \neq \emptyset, E^+ \neq \emptyset \\ [0, d_1^+(\eta_{\min}^+)) & \text{if } E^- = \emptyset, E^+ \neq \emptyset \\ (d_1^-(\eta_{\min}^-), d_{\mathrm{UB}}] & \text{if } E^- \neq \emptyset, E^+ = \emptyset \\ \emptyset & \text{if } N_1 + N_2^{\max} < T \\ [0, d_{\mathrm{UB}}] & \text{otherwise} \end{cases} \tag{C.37a}$$

$$S_\eta^- := \begin{cases} [d_1^-(\eta), d_1^-(\eta-1)] & \text{if } \eta_{\min}^- \le \eta - 1 < \eta \le \eta_{\max}^- \\ \emptyset & \text{otherwise} \end{cases} \quad \forall \eta \in \{N_2^{\min}+1, \ldots, N_2^{\max}\} \tag{C.37b}$$

$$S_\eta^+ := \begin{cases} [d_1^+(\eta-1), d_1^+(\eta)] & \text{if } \eta_{\min}^+ \le \eta - 1 < \eta \le \eta_{\max}^+ \\ \emptyset & \text{otherwise} \end{cases} \quad \forall \eta \in \{N_2^{\min}+1, \ldots, N_2^{\max}\} \tag{C.37c}$$

$$S_{\max} = \begin{cases} [0, d_1^-(\eta_{\max}^-)) \cup (d_1^+(\eta_{\max}^+), d_{\mathrm{UB}}] & \text{if } E^- \neq \emptyset, E^+ \neq \emptyset \\ (d_1^+(\eta_{\max}^+), d_{\mathrm{UB}}] & \text{if } E^- = \emptyset, E^+ \neq \emptyset \\ [0, d_1^-(\eta_{\max}^-)) & \text{if } E^- \neq \emptyset, E^+ = \emptyset \\ [0, d_{\mathrm{UB}}] & \text{if } N_1 + N_2^{\max} < T \\ \emptyset & \text{otherwise.} \end{cases} \tag{C.37d}$$

We will reformulate (C.34) on each interval (set) separately, assuming it is nonempty.

1. "$S_{\min}$": If $d_1 \in S_{\min}$, then $f(d_1, \eta; \hat{\rho}_L, \tau_L) < f(d_1, \eta; \hat{\rho}_U, \tau_L)$ for all $\eta \in \{N_2^{\min}, \ldots, N_2^{\max}\}$ according to Lemma C.3b, so it is optimal to deliver $N_2^{\min}$ fractions. Hence, on this interval (C.34) is equivalent to

$$q \le f(d_1, N_2^{\min}; \hat{\rho}_L, \tau_L). \tag{C.38}$$

2. "S_η^-": From Lemma C.2a we know that $f(d_1, \eta, \hat{\rho}_L, \tau_L) = f(d_1, \eta, \hat{\rho}_U, \tau_L)$ if $d_1 = d_1^-(\eta)$ or $d_1 = d_1^+(\eta)$. In this case, the objective value equals K. Furthermore, function $f(d_1, \eta, \hat{\rho}_L, \tau_L)$ is convex and $f(d_1, \eta, \hat{\rho}_U, \tau_L)$ is concave in d_1. Consider the interval $[d_1^-(\eta), d_1^-(\eta-1)]$. It holds that

$$f(d_1, \eta; \hat{\rho}_L, \tau_L) \le K \le f(d_1, \eta-1; \hat{\rho}_L, \tau_L) \tag{C.39a}$$
$$f(d_1, \eta-1; \hat{\rho}_U, \tau_L) \le K \le f(d_1, \eta; \hat{\rho}_U, \tau_L). \tag{C.39b}$$

This implies that if $d_1 \in [d_1^-(\eta), d_1^-(\eta-1)]$, it is optimal to deliver either η or $\eta - 1$ fractions. If we deliver η fractions, the restricting worst-case scenario is $(\hat{\rho}_L, \tau_L)$ and the value f is above K for the scenario $(\hat{\rho}_U, \tau_L)$. If we deliver $\eta' > \eta$ fractions, the value for the scenario $(\hat{\rho}_L, \tau_L)$ decreases, while the value for the scenario $(\hat{\rho}_U, \tau_L)$ increases even further. Hence, delivering $\eta' > \eta$ fractions cannot be optimal. Similarly, delivering less than $\eta - 1$ fractions cannot be optimal. Therefore, if $d_1 \in [d_1^-(\eta), d_1^-(\eta-1)]$ it is optimal to deliver either η or $\eta - 1$ fractions. This implies that on the interval S_η^- constraint (C.34) is equivalent to

$$q \le \max\{f(d_1, \eta; \hat{\rho}_L, \tau_L), \; f(d_1, \eta-1; \hat{\rho}_U, \tau_L)\}. \tag{C.40}$$

Note that this result does not depend on the values $\hat{\rho}_L$ and $\hat{\rho}_U$, we only use that $\hat{\rho}_L < \frac{\tau_L}{\sigma} < \hat{\rho}_U$.

3. "S_η^+": Similar to the case for S_η^+, one can show that for $d_1 \in S_\eta^+$ constraint (C.34) is equivalent to (C.40).

4. "$S_{\max}$": If $d_1 \in S_{\max}$, then $f(d_1, \eta; \hat{\rho}_L, \tau_L) > f(d_1, \eta; \hat{\rho}_U, \tau_L)$ for all $\eta \in \{N_2^{\min}, \ldots, N_2^{\max}\}$ according to Lemma C.3b, so it is optimal to deliver $N_2^{\max}$ fractions. Hence, on this interval (C.34) is equivalent to

$$q \le f(d_1, N_2^{\max}; \hat{\rho}_U, \tau_L). \tag{C.41}$$

43

For sets S_η^- and S_η^+ the reformulation is the same. Therefore, define

$$S_\eta = S_\eta^- \cup S_\eta^+. \tag{C.42}$$

Putting everything together, for $d_1 \in [0, d_{\mathrm{UB}}]$ the constraint (C.34) is equivalent to

$$q \leq \sum_{\eta \in \{N_2^{\min}, \ldots, N_2^{\max}\}} \max\{f(d_1, \eta; \hat{\rho}_L, \tau_L), \ f(d_1, \eta - 1; \hat{\rho}_U, \tau_L)\} I(d_1 | S_\eta)$$
$$+ f(d_1, N_2^{\min}; \hat{\rho}_L, \tau_L) I(d_1 | S_{\min}) + f(d_1, N_2^{\max}; \hat{\rho}_U, \tau_L) I(d_1 | S_{\max}), \quad \forall (\hat{\rho}, \hat{\tau}) \in Z_{\mathrm{ID}}^{\mathrm{int}} \cap \{(\hat{\rho}, \hat{\tau}) : \hat{\tau} \leq \tau_L + r^\tau\}. \tag{C.43}$$

In order to find a tractable conservative robust counterpart of (C.43), denote

$$\rho_L^{\mathrm{int}} = \max\{\rho_L, \frac{\tau_L}{\sigma} - 2r^\rho\} \tag{C.44a}$$

$$\rho_U^{\mathrm{int}} = \min\{\rho_U, \frac{\tau_L}{\sigma} + 2r^\rho\}, \tag{C.44b}$$

and note that $\rho_L^{\mathrm{int}} \leq \hat{\rho}_L < \frac{\tau_L}{\sigma} < \hat{\rho}_U \leq \rho_U^{\mathrm{int}}$. Only if $d_1 \in S_\eta$, the robust counterpart is conservative. By Lemma C.3b, it holds that function f is strictly decreasing, constant or strictly increasing in ρ for fixed d_1, so

$$f(d_1, \eta; \hat{\rho}_L, \tau_L) \geq \min\{f(d_1, \eta; \rho_L^{\mathrm{int}}, \tau_L), f(d_1, \eta; \frac{\tau_L}{\sigma}, \tau_L)\} = \min\{f(d_1, \eta; \rho_L^{\mathrm{int}}, \tau_L), K\} = f(d_1, \eta; \rho_L^{\mathrm{int}}, \tau_L),$$

where the second equality follows from (C.39). A similar result holds for $f(d_1, \eta - 1; \hat{\rho}_U, \tau_L)$. Furthermore, as shown before, f is increasing in ρ on $S_{\min}$ and decreasing in ρ on $S_{\max}$. Therefore, a conservative approximation of (C.31) is given by

$$q \leq \sum_{\eta \in \{N_2^{\min}, \ldots, N_2^{\max}\}} \max\{f(d_1, \eta; \rho_L^{\mathrm{int}}, \tau_L), \ f(d_1, \eta - 1; \rho_U^{\mathrm{int}}, \tau_L)\} I(d_1 | S_\eta)$$
$$+ f(d_1, N_2^{\min}; \rho_L^{\mathrm{int}}, \tau_L) I(d_1 | S_{\min}) + f(d_1, N_2^{\max}; \rho_U^{\mathrm{int}}, \tau_L) I(d_1 | S_{\max}), \tag{C.45}$$

and the RHS is $p(d_1)$. $\qquad \square$

Function $p(d_1)$ is a piece-wise function. On intervals defined by $S_{\min}$ and $S_{\max}$ it is convex and concave, respectively. On any interval S_η function $p(d_1)$ is the maximum of a concave and convex function.