

# Optimal dynamic treatment allocation

Anders Bredahl Kock and Martin Thyrsgaard\*

University of Oxford, Aarhus University and CREATES

October 4, 2017

## Abstract

In a treatment allocation problem the individuals to be treated often arrive gradually. Initially, when the first treatments are made, little is known about the effect of the treatments but as more treatments are assigned the policy maker gradually learns about their effects by observing the outcomes. Thus, there is a tradeoff between exploring the available treatments to learn about their merits and exploiting the best treatment, i.e. administering it as often as possible, in order to maximise the cumulative welfare of all the assignments made. Furthermore, a policy maker may not only be interested in the expected effect of the treatment but also its riskiness. Thus, we allow for the welfare function to depend on the first and second moments of the distribution of treatment outcomes. We propose a dynamic treatment policy which attains the minimax optimal regret relative to the unknown best treatment in this dynamic setting. We allow for the data to arrive in batches as, say, unemployment programs only start once a month or blood samples are only sent to the laboratory for investigation in batches. Furthermore, we show that the minimax optimality does not come at the price of overly aggressive experimentation as we provide upper bounds on the expected number of times any suboptimal treatment is assigned. We also consider the case where the outcome of a treatment is only observed with delay as it may take time for the treatment to work. Thus, a doctor faces a tradeoff between getting imprecise information quickly by making the measurement soon after the treatment is given or getting precise information later at the expense of less information for the individuals who are treated in the meantime. Finally, using Danish register data, we show how our treatment policy can be used to assign unemployed to active labor market policy programs in order to maximise the probability of ending the unemployment spell.

*Keywords:* Treatment allocation, Outcomes observed with delay, Batched data, Welfare function, Bandit framework.

*JEL classifications:* C18, C22, J68

---

\*Address for correspondence: anders.kock@economics.ox.ac.uk. This work was supported by CREATES which is funded by the Danish National Research Foundation (DNRF78).

# 1 Introduction

A policy maker must often assign treatments gradually as the individuals to be treated do not arrive simultaneously. For example, people become unemployed gradually throughout the year and assignment to one of several unemployment programs is often made shortly thereafter. Similarly, patients with too high blood pressure arrive gradually to a medical clinic and the doctor assigns one of several treatments to each of them. The policy maker or doctor gradually accrues information by observing the outcome of previous treatments prior to the next assignment. Throughout the paper we shall use these two examples as illustrations of our results and be particularly concerned with how treatments should be assigned in order to maximise welfare. In doing so one faces a tradeoff between exploring which treatment works best and exploiting the information gathered so far from previous assignments in order to assign the best treatment to as many individuals as possible. Here exploring amounts to administering different treatments in order to gain information about their effectiveness.

The above setup is in stark contrast to typical estimation of treatment effects where one presupposes the existence of a data set of a certain size  $N$ . Thus, in the typical setting, the size and composition of the data set are determined prior to estimation. Based on this exogenously given data set, treatment effects are estimated and assignments are made. We consider the case where the observed data that the treatment assignments must be based on is a part of the policy in the sense that it depends on the previous choices of the the policy maker. Thus, the policy maker enters already in the design phase of the treatment program and he can adjust the experiment as data accumulates. In other words, he decides *how* to draw the sample and thus its composition by the allocations he makes. Furthermore, the sample size itself may be a random variable unknown to the policy maker as he does not know a priori how many will become unemployed in the course of the year that the program is scheduled to run and the exact shape of a good treatment rule will depend on the expected number of individuals to be treated: if many individuals are expected to become unemployed in the course of the year it might be beneficial to experiment more in the beginning to harvest the benefits of increased information later on.

It should be noted that the goal of this paper is not to test whether one treatment is better than the other ones at the end of the treatment period. This would amount to a pure exploration problem where the sole purpose of the sampling is to maximize the amount of information at the end of the sample without regard to the welfare of the treated individuals. While this problem is interesting in its own right it is often not viable for ethical reasons in the social sciences. Instead, the problem under investigation here is how to sample (assign treatments) in order to maximise the expected cumulative welfare of the treated individuals.

We consider a setting where the policy maker has to choose one of several treatments whose outcome distribution may vary across individuals due to differences in observable characteristics

(covariates). We show how to take this heterogeneity across individuals into account in a dynamic treatment setting where the individuals to be treated arrive gradually and establish the minimax optimality of the proposed rule. Furthermore, we allow for a setting where individuals, and thus information, may arrive in batches. For example, people do not get assigned to an unemployment program on the exact day they become unemployed as new programs might only start once a month. Thus, people are pooled and as a result data arrives in batches. This setup strikes a middle ground between the *bandit* framework where individuals arrive one-by-one and the classic treatment effect framework where a data set of size  $N$  is presupposed. We also consider a related practically relevant situation where the outcomes of previous treatments are observed only with delay. In a medical trial, for example, one may choose to delay the measurement of the outcome of the treatment in order to obtain more precise information of the effect of a certain drug. As it takes time for the effect of a drug to set in the delaying of the measurement can lead to more precise information on the effect of the drug. The price of this delay is that less information is available when treating other patients prior to the measurement being made. Thus, there is a trade-off between obtaining imprecise information quickly (by making the measurement shortly after the treatment) and obtaining more precise information later (by postponing the measurement).

We also argue that the desirability of a treatment cannot be measured only by its expected outcome. A sensible welfare function must take into account the risk of a treatment. For example, it may well be that drug A is expected to lower the blood pressure slightly more than drug B but A might still not be preferred if it is much more risky than B. In this paper we shall measure the *risk* of a treatment by its variance and take into account that mean as well as variance may be relevant in determining the most desirable treatment.

Furthermore, our approach easily accommodates practical policy concerns restricting the type of treatment rules that are feasible. For instance, the policy maker may want rules that depend on the individual's characteristics in a simple way due to political or ethical reasons.

We show that our policy can obtain the minimax optimal regret compared to the infeasible policy which for each individual knows the optimal treatment a priori. We also provide an upper bound on the expected number of times that our treatment policy assigns any suboptimal treatment. This is an important ethical guarantee since it ensures that the low regret is not obtained at the price of wild experimentation or maltreatment of many individuals in order to achieve a greater welfare in the long run.

Finally, we illustrate our dynamic treatment policy on Danish register data by showing how unemployed workers can be assigned to active labor market policy programs such that the sum of the probabilities of ending the unemployment spell is maximised. We consider two treatments: i) extra education and ii) job training and find that our dynamic treatment policy eliminates job training (and retains extra education) for older individuals irrespective of their gender. A potential explanation for this is that older workers benefit more from updating their skill set than younger

ones who have more recently finished their education.

## 1.1 Related literature

Our paper is related to the literature on statistical treatment rules in econometrics. Manski (2004) proposed *conditional empirical success* (CES) rules which take a finite partition of the covariate space and on each set of this partition dictate to assign the treatment with the highest sample average. When implementing CES rules one must decide on how fine to choose the partition of the covariate space and thus faces a tradeoff between using highly individualized rules and having enough data to accurately estimate the treatment effects for each group in the partition. Among other things, Manski (2004) provides sufficient conditions for full individualization to be optimal. The tradeoff between full individualization of treatments and having sufficient data to estimate the treatment effects accurately is also found in our dynamic treatment setting.

Stoye (2009) showed that if one does not restrict how outcomes vary with covariates then full individualization is always minimax optimal. Thus, if age is a covariate, information on treatment effects for 30 year olds should not be used when making treatment decisions for 31 year olds. This result relies on the fact that without any restrictions on how the outcome distribution varies with covariates, this relationship could be infinitely wiggly such that even similar individuals may carry no information about how treatments affect the other person. Also, as the support of the covariate vector grows, these "no-cross-covariate" rules become no-data rules as for many values of the covariates there will be no observations. This is certainly the case for continuous covariates. Our assumptions rule out such wiggleness as no practical policy can be expected to work well in such a setting.

Furthermore, our work is related to the recent work of Kitagawa and Tetenov (2015) who consider treatment allocation through an *empirical welfare maximization* lens. The authors take the view that realistic policies are often constrained to be simple due to ethical, legislative or political reasons. Using techniques from empirical risk minimization they show how their procedure is minimax optimal within the considered class of realistic policies. Our approach is related to theirs in that we also allow the policy maker to focus on simple rules in the dynamic framework. Furthermore, Athey and Wager (2017) have used concepts from semiparametric efficiency theory to establish regret bounds that scale with the semiparametrically efficient variance.

Other important papers on statistical treatment rules in econometrics focusing on the case where the sample is given include Chamberlain (2000), Dehejia (2005), Hirano and Porter (2009), Bhattacharya and Dupas (2012), Stoye (2012), Tetenov (2012) and Kasy (2014).

Our paper is also related to the vast literature on bandit problems as the modeling framework adopted here is that of a multi-armed bandit with covariates. In the classical bandit problems one seeks to maximize the expected cumulative reward from pulling arms with unknown means one by one. In a seminal paper Robbins (1952) introduced a class of bandit problems and proposed

some initial solutions guaranteeing that the average reward will converge to the mean of the best arm. The first paper which considered bandit problems where one observes a covariate prior to making an allocation decision was Woodroffe (1979) who made a parametric assumption on how covariates affect outcomes. The first work allowing covariates to affect the distribution of outcomes in a nonparametric way was Yang et al. (2002). For an excellent review of the literature on bandit problems we refer to Bubeck and Cesa-Bianchi (2012).

Our dynamic treatment policy is a variant of the successive elimination (SE) policy of Perchet and Rigollet (2013) which is in turn related to the work of Even-Dar et al. (2003). The successive elimination policy gets its name from its successive elimination of suboptimal arms. It trades off the desire to quickly eliminate suboptimal arms with the desire not to eliminate the best arm. Perchet and Rigollet (2013) provided adaptive and distribution independent upper bounds on the regret of the SE policy in the setting without covariates. They next proposed the binned successive elimination (BSE) which groups individuals with similar characteristics into bins (groups). For each group of individuals the SE policy is run separately. Perchet and Rigollet (2013) showed that by grouping individuals in a certain way this leads to minimax optimal regret bounds which are sublinear in the number of assignments made. Our treatment policy also builds on successive elimination of suboptimal treatments applied separately to a number of groups defined by the policy maker. There are several differences, however, to the setting studied in Perchet and Rigollet (2013). First, we consider welfare functions depending not only on the mean treatment outcome but also on the variance. Furthermore, we do not require that the welfare function is linear. This allows us to capture the dynamic risk-return tradeoff facing a policy maker when deciding which treatment to assign. Second, we consider the case where the outcome of treatments is only observed with delay. As already explained this creates a tradeoff between obtaining imprecise information quickly and obtaining precise information later. Third, we provide upper bounds on the regret of our policy for any choice of grouping individuals. These upper bounds depend on the geometry of the chosen grouping. Allowing for groups of general shapes is important to inform policy makers about how exactly their choice of grouping individuals affects regret since the choice of groups achieving minimax regret may not always be politically or ethically feasible. Fourth we extend our method to also allow for discrete covariates such as gender which may affect the treatment outcome. Finally, we allow the data to arrive in batches (note that batched bandit problems were considered in Perchet et al. (2016) in the setting with only two arms and in the absence of covariates).

The multi-armed bandit setup has also been used in the context of social learning and strategic experimentation in the works of e.g. Bolton and Harris (1999), Keller et al. (2005) and Klein and Rady (2011). Here several agents have to choose the amount of experimentation (pulling a risky arm) taking into account that the information obtained will be available to all other players as well. While the agents have an incentive to free ride they also want to experiment in order to bring forward the time where extra information is generated.

The term *optimal dynamic treatment allocation* in the bandit framework as discussed in this paper should not be confused with similar terms in the medical statistics literature. In that literature *adaptive treatment strategies/adaptive interventions* and *dynamic treatment regimes* refer to a setting where the same individual is observed repeatedly over time and the level as well as the type of the treatment is adjusted according to the individual’s needs. References to this setting include Robins (1997), Lavori et al. (2000), Murphy et al. (2001), Murphy (2003) and Murphy (2005).

The remainder of the paper is organised as follows. Section 2 considers a setting where the treatment outcomes do not depend on observable individual specific characteristics. Next, Section 3 introduces covariates and establishes regret bounds for the dynamic treatment policy. When grouping individuals in a specific way, these bounds are minimax optimal. It is also shown that the expected number of sub-optimal assignments increases slowly and we investigate how to handle discrete covariates. Section 4 investigates the effect of outcomes being observed with delay and Section 5 implements our method on the register based Danish DREAM data by assigning unemployed to various treatments. Finally, Section 6 concludes while 7 contains all proofs.

## 2 The treatment problem without covariates

We begin by considering the dynamic treatment allocation problem where the outcome of the treatments do not depend on observable individual specific characteristics. The distribution of treatment outcomes does not vary across individuals. While this setting may often be too restrictive, the bounds established in this section will be used as important ingredients in establishing the properties of our treatment rules in the setting where covariates are observed on each individual prior to the treatment assignment.

Consider a setting with  $K + 1$  different treatments and  $N$  assignments.  $N$  is a random variable whose value need not be known to the policy maker at the beginning of the treatment assignment problem <sup>1</sup>. For example, he does not know at the beginning of the year how many will become unemployed during the year. Let  $Y_t^{(i)} \in [0, 1]$  denote the outcome from assigning treatment  $i$ ,  $i = 1, \dots, K + 1$  to individual  $t$ ,  $t = 1, \dots, N$  where the subscript  $t$  indicates the order in which individuals are treated. It is merely for technical reasons that we assume the treatment outcomes to take values in  $[0, 1]$  and this interval can be generalized to any interval  $[-I_1, I_2]$  for some  $I_1, I_2 \in \mathbb{R}$  or  $Y_t^{(i)}$  being sub-gaussian without qualitatively changing our results. The framework accommodates treatments with different costs since, whenever it makes sense,  $Y_t^{(i)}$  can be defined net of costs <sup>2</sup>.

---

<sup>1</sup>We consider a setting with  $K + 1$  treatments for purely notational reasons since it is the number of suboptimal treatments,  $K$ , which will enter our regret bounds as well as many of the arguments in the appendix.

<sup>2</sup>We add the proviso “whenever it makes sense” since sometimes the outcome of a treatment is hard or impossible to measure on a monetary scale, for example whether a patient survives a surgery or not.

We allow for the data to arrive in  $M$  batches of sizes  $m_j$ ,  $j = 1, \dots, M$ , such that the total number of assignments  $N = \sum_{j=1}^M m_j$ . If an unemployment program is run for twelve months and new programs start every month then  $M = 12$  and the  $m_j$  indicate how many individuals become unemployed in the  $j$ th month. The  $m_j$  are allowed to be random variables as the policy maker does not know a priori how many will become unemployed each month. This is in contrast to typical treatment allocation problems where the size as well as the composition of the data set are taken as given. Every individual  $t$  belongs to exactly one of the batches. For each batch the outcomes of the assignments are only observed at the end of the batch. Thus, the treatment assignments for individuals belonging to batch  $\tilde{j}$  can only depend on the outcomes observed from previous batches  $j = 1, \dots, \tilde{j} - 1$ . This is reasonable as information gained from treating persons who have become unemployed prior to person  $t$ , yet in the same month/batch cannot be used to inform the treatment allocation of person  $t$  as all persons from the same batch start their programs at the same time. Note also that  $m_j = 1$  for all  $j$  amounts to the traditional bandit problem without batched data (the individuals arrive in batches of size one).

The vectors  $Y_t = (Y_t^{(1)}, \dots, Y_t^{(K+1)})$ , are assumed to be i.i.d. However, the joint distribution of the entries of  $Y_t$  is left unspecified. This in accordance with real applications where an unemployed individual's response to two types of job training programs may be highly correlated. We are interested in a general *welfare function*  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$  of the mean  $\mu^{(i)} = \mathbb{E}Y_t^{(i)}$  and variance  $(\sigma^2)^{(i)} = \mathbb{E}(Y_t^{(i)} - \mu^{(i)})^2$  of the treatment outcome  $Y_t^{(i)}$ . Thus, define  $f^{(i)} = f(\mu^{(i)}, (\sigma^2)^{(i)})$ . The *welfare maximizing* (best) treatment is denoted by  $*$  and satisfies  $f^{(*)} = \arg \max_{1 \leq i \leq K+1} f^{(i)}$ <sup>3</sup>. The welfare maximizing treatment strikes the optimal balance between expected treatment outcome and the riskiness of the treatment. Let  $\Delta_i = f^{(*)} - f^{(i)} \geq 0$  be the difference between the best and the  $i$ th treatment and assume that  $\Delta_1 \geq \dots \geq \Delta_K > \Delta_{K+1} = 0$ . The ranking of the  $\Delta_i$  is without loss of generality and does not necessarily imply a ranking of neither the  $\mu^{(i)}$  nor the  $(\sigma^2)^{(i)}$ .

A treatment allocation is a rule  $\pi = \{\pi_t\}$  assigning a treatment from the set  $\{1, \dots, K+1\}$  to every individual  $t = 1, \dots, N$ . This allocation can only depend on the outcomes from previous batches.

Our goal is to maximize expected cumulated welfare over the  $N$  treatments. This is equivalent to minimizing the expected difference to the infeasible welfare one would have obtained from always assigning the best treatment, i.e. minimizing the expected value of the *regret*

$$R_N(\pi) = \sum_{j=1}^M \sum_{i=1}^{m_j} \left( f^{(*)} - f^{(\pi_{j,i})} \right) = \sum_{t=1}^N \left( f^{(*)} - f^{(\pi_t)} \right). \quad (2.1)$$

where the second inequality is due to the fact that each individual  $t$  can be uniquely identified with an assignment  $(i)$  made in a batch  $(j)$ .

<sup>3</sup>We assume without loss of generality that the best treatment is unique.

## 2.1 Examples

Throughout this paper we assume that  $f$  is Lipschitz continuous from  $[0, 1]^2$  equipped with the  $\ell_1$ -norm to  $\mathbb{R}$  with Lipschitz constant  $\mathcal{K}$  independent of all other parameters, i.e.

$$|f(u_1, u_2) - f(v_1, v_2)| \leq \mathcal{K} (|u_1 - v_1| + |u_2 - v_2|)$$

By making concrete choices for  $f$  our framework contains the following instances as special cases.

1.  $f(\mu, \sigma^2) = \mu$  (implying  $\mathcal{K} = 1$ ) amounts to the classical bandit problem where one only targets the mean welfare. However, unlike this paper, the classical setting does not consider batched data or outcomes that are observed with delay. On the other hand, the classical setting does include the case where one is interested in maximizing cumulated expected utility: one can simply define  $Y_t^{(i)} = u(\tilde{Y}_t^{(i)})$  as the utility received from assigning treatment  $i$  at time  $t$  with outcome  $\tilde{Y}_t^{(i)}$ .
2.  $f(\mu, \sigma^2) = \frac{\mu}{\sigma}$  amounts to Sharpe ratios which are frequently used in financial applications to measure risk-return tradeoffs. If  $(\sigma^2)^{(i)} \geq c$  for all  $i = 1, \dots, K + 1$  then one has by the mean value theorem that  $\mathcal{K} = \frac{1}{2c^{3/2}}$  works.
3.  $f(\mu, \sigma^2) = \mu - \frac{\alpha}{2}\sigma^2$  for a risk aversion parameter  $\alpha > 0$  is another typical way of measuring the tradeoff between expected outcomes and their variance. Here  $\mathcal{K} = \max(1, \alpha/2)$ .
4.  $f(\mu, \sigma^2) = -\sigma^2$  (implying  $\mathcal{K} = 1$ ) amounts to the case where one is interested only in minimizing the variance.
5. The theory developed in this paper can be extended to the case where one is interested in dynamically maximizing the sum of welfare functions of any number of moments, i.e.  $f(\mu_1^{(i)}, \dots, \mu_d^{(i)})$  for some  $d \geq 1$ , where  $\mu_k^{(i)} = \mathbb{E}[(Y_t^{(i)})^k]$  is the  $k$ 'th moment of  $Y_t^{(i)}$ . Higher moments than the second one may be relevant if the the policy maker has, say, skewness aversion. This is relevant in dynamic portfolio allocation problems and finance as in Harvey and Siddique (2000). To keep the exposition simple, we have chosen to focus on the case where  $f$  depends on the first two moments only as the extension to welfare functions of strictly more than two moments is mainly technical.

Informally, the dynamic treatment policy works by eliminating treatments that are deemed to be inferior based on the outcomes observed so far. We then take turns assigning each of the remaining treatments in the next batch. This is the exploration step. After this step elimination takes place again.

To describe the policy more formally, let  $m_{i,j}$  be the number of times treatment  $i$  is assigned in batch  $j$ . Thus,  $m_j = \sum_{i=1}^{K+1} m_{i,j}$  and we define  $B_i(b) = \sum_{j=1}^b m_{i,j}$  as the number of times



treatment  $i$  has been assigned up to and including  $b$  batches,  $b = 1, \dots, M$ . Next, let  $\hat{\mu}_s^{(i)} = \frac{1}{s} \sum_{t=1}^s Y_t^{(i)}$  and  $(\hat{\sigma}_s^2)^{(i)} = \frac{1}{s} \sum_{t=1}^s (Y_t^{(i)} - \hat{\mu}_s^{(i)})^2$  be estimators of  $\mu^{(i)}$  and  $(\sigma^2)^{(i)}$ , respectively.

**Dynamic treatment policy:** Denote by  $\hat{\pi}$  the dynamic treatment policy. Let  $\mathcal{I}_b \subseteq \{1, \dots, K+1\}$  be the set of remaining treatments before batch  $b$  and let  $\underline{B}(b) = \min_{i \in \mathcal{I}_b} B_i(b)$  be the number of times that each remaining treatment at least has been assigned up to and including batch  $b$ .

1. In each batch  $b = 1, \dots, M$  we take turns assigning each remaining treatment. Thus, the difference between the number of times that any pair of remaining treatments has been assigned at the end of a batch is at most one.
2. At the end of batch  $b$  eliminate treatment  $i$  if

$$\max_{i \in \mathcal{I}_b} f(\hat{\mu}_{\underline{B}(b)}^{(i)}, (\hat{\sigma}_{\underline{B}(b)}^2)^{(i)}) - f(\hat{\mu}_{\underline{B}(b)}^{(i)}, (\hat{\sigma}_{\underline{B}(b)}^2)^{(i)}) \geq 32\gamma \sqrt{\frac{2}{\underline{B}(b)} \overline{\log} \left( \frac{T}{\underline{B}(b)} \right)}$$

where  $\gamma > 0$ ,  $T \in \mathbb{N}$  and  $\overline{\log}(x) = \log(x) \vee 1$ .

The dynamic treatment policy uses the sample counterparts of  $\mu^{(i)}$  and  $(\sigma^2)^{(i)}$  to evaluate whether treatment  $i$  is inferior to the best of the remaining treatments. The parameter  $\gamma$  controls how aggressively treatments are eliminated. Small values of  $\gamma$  make it easier to eliminate inferior treatments but also induce a risk of potentially eliminating the best treatment. The parameter  $T$ , which will often be set equal to the expected sample size  $n = \mathbb{E}(N)$ , is needed exactly to ensure that we are cautious eliminating treatments after the first couple of batches where  $\hat{\mu}_{\underline{B}(b)}^{(i)}$  and  $(\hat{\sigma}_{\underline{B}(b)}^2)^{(i)}$  could be based on few observations and thus need not be precise estimates of  $\mu^{(i)}$  and  $(\sigma^2)^{(i)}$ , respectively. From a technical point of view, this ensures that we can uniformly (over treatments) control the probability of eliminating the best treatment. Note that eliminating the best treatment is very costly as regret will accumulate linearly after such a mistake<sup>4</sup>. Furthermore, the dynamic treatment policy need neither to know sample size  $N$ , nor the number of batches  $M$  in order to run. It can be stopped at any point in time with regret bounds as outlined in Theorem 2.1 below.

**Remark** *In practice one may also consider a policy which allows treatments to reenter the treatment set even after they have been eliminated. On the other hand, there is no reason for this from a theoretical point of view as the rates in Corollary 3.1.1 below are minimax optimal such that one can at most expect to improve the constant entering the upper bound on regret. Heuristically, the dynamic treatment policy is constructed in such a way that treatments are only eliminated if we are very certain that they are suboptimal. Thus, there is no need to reintroduce previously eliminated treatments.*

<sup>4</sup>If the best treatment is eliminated then the regret from each subsequent treatment is  $f^{(*)} - f^{(\hat{\pi}_t)} \geq \Delta_K > 0$

## 2.2 Optimal treatment assignment without covariates

Without an upper bound on the size of the batches it is clear that no non-trivial upper bound on regret can be established. For example, the data could arrive in one batch of size  $N$  implying that feedback is never received prior to any assignment. Thus, we shall assume that no batch is larger than  $\bar{m}$  where  $\bar{m}$  is non-random, i.e.  $m_j \leq \bar{m}$  for  $j = 1, \dots, M$ . Our first result provides an upper bound on the regret incurred by the dynamic treatment policy.

**Theorem 2.1** *Consider a treatment problem with  $(K + 1)$  treatments and unknown number of assignments  $N$  with expectation  $n$  that is independent of the treatment outcomes. By implementing the dynamic treatment policy with parameters  $\gamma = \mathcal{K}$  and  $T = n$  one obtains the following bound on the expected regret*

$$\mathbb{E} [R_N(\hat{\pi})] \leq C \min \left( \bar{m} \mathcal{K}^2 \sum_{j=1}^K \frac{1}{\Delta_j} \overline{\log} \left( \frac{n \Delta_j^2}{\mathcal{K}^2} \right), \sqrt{n \mathcal{K}^3 \bar{m} K \log(\bar{m} K / \mathcal{K})} \right) \quad (2.2)$$

for a positive constant  $C$ .

The upper bound in Theorem 2.1 consists of two parts. The first part is adapting to the unknown distributional characteristics  $\Delta_j$ . Note that the regret in this part only increases logarithmically in the the expected number of treatments  $n$ . This logarithmic rate is unimprovable in general since it is known to be optimal even in the case where one only targets the mean (which in our setting corresponds to  $f(x, y) = x$ ) such that  $\mathcal{K} = 1$ ) and the treated individuals arriving one-by-one ( $\bar{m} = 1$ ), see e.g. Theorem 2.2 in Bubeck and Cesa-Bianchi (2012). On the other hand, the first part of (2.2) can be made arbitrarily large by letting e.g.  $\Delta_1 \rightarrow 0$ . Thus, the bound is not uniform in the underlying distribution of the data. The second part is uniform and in fact yields the minimax optimal rate up to a factor of  $\sqrt{\log(K)}$  even in the case where only the welfare function  $f(x, y) = x$  is considered. It is reasonable that both parts of the upper bound in (2.2) are increasing in  $\bar{m}$  since as the maximum batch size increases the time between potential elimination of suboptimal treatments increases implying that these are assigned more often. Similarly, more experimentation between treatments takes place when the number of these,  $K$ , is increased which results in increased regret.

Note that the implementation of the dynamic treatment algorithm requires knowledge of the expected number of individuals that are going to be treated. In medical experiments the total number of individuals participating is often determined a priori making  $N$  known and deterministic (and equal to  $n$ ). On the other hand, when allocating unemployed to treatments the total number of individuals becoming unemployed in the course of the year is unknown. However, one often has a good estimate on the expected value  $n$  which is what matters for the treatment policy. For example, one may use averages of the number of individuals who have become unemployed in previous years

to estimate  $n$ . Alternatively, one can use the doubling trick which resets the treatment policy at prespecified times in order to avoid any assumptions on the size of  $N$  or  $n$ . Usage of the doubling trick would imply that eliminated treatments reappear and get another chance every time the policy is reset thus allowing for the efficiency of treatments to vary over time. For further details on the doubling trick and its implementation we refer to Shalev-Shwartz et al. (2012).

## 2.3 Suboptimal treatments

Theorem 2.1 showed that the expected cumulated welfare of the dynamic treatment policy will not be much larger than the one from the infeasible policy that always assigns the best treatment. However, for an assignment rule to be ethically and politically viable it is important that it does not yield high welfare at the cost of sacrificing certain individuals by wild experimentation. For example, it may not be ethically defensible for a doctor to assign a suboptimal treatment to a patient in order to gain more certainty for future treatments. The following theorem shows that our treatment rule does not suffer from such a problem in the sense that the expected number of times any suboptimal treatment is assigned only increases logarithmically in the sample size.

**Theorem 2.2** *Suppose the dynamic treatment policy is implemented with parameters  $T = n$  and  $\gamma = \mathcal{K}$ . Let  $T_i(t)$  denote the number of times treatment  $i$  is assigned up to and including observation  $t$ . Then*

$$\mathbb{E} [T_i(N)] \leq C \left( \mathcal{K}^2 K \frac{\log \left( \frac{n}{\mathcal{K}^2} \right)}{\Delta_i^2} + K \bar{m} + \mathcal{K}^2 \right),$$

for any suboptimal treatment  $i \in \{1, \dots, K\}$  and a positive constant  $C$ .

The important ethical guarantee on the treatment rule is that it only assigns very few persons to a suboptimal treatment (logarithmic growth rate in the sample size). It is in line with intuition that the closer any suboptimal treatment is to being optimal ( $\Delta_i$  closer to zero) the more difficult it is to guarantee that this treatment is rarely assigned. The reason is that this treatment must be assigned more often before it can confidently be concluded that it is suboptimal and thus eliminated.

## 3 Treatment outcomes depending on covariates

So far we have considered the case where the outcome of a treatment does not depend on the characteristics of the individual it is assigned to. In reality, however, persons react differently to the same type of treatment: while a certain medicine may work well for one person it may be outright dangerous to assign it to another person if this person is allergic to some of the substances. Similarly, the effect of further education on the probability of an unemployed individual finding a job may also depend on, e.g., the age of the individual: individuals close to the retirement age may

benefit more from short courses updating their skill set while young individuals may benefit more from going back to school for an extended period of time.

Prior to assigning individual  $t$  to a treatment we observe a vector  $X_t \in [0, 1]^d$  of covariates with distribution  $P_X$ . In the case of assigning unemployed persons to various unemployment programs  $X_t$  could include age, length of education and years of experience. It is merely for technical convenience that we assume the variables to take values in  $[0, 1]$  and the assumption of bounded support can be replaced by tail conditions on the distribution of  $X_t$ .  $P_X$  is assumed to be absolutely continuous with respect to the Lebesgue measure with density bounded from above by  $\bar{c} > 0$ . As we now observe covariates on each individual prior to the treatment assignment we condition on these as, for example, the risk of a treatment may be individual specific and depend on, e.g., whether the person has an allergy or not. Thus, in close analogy to the setting without covariates, we now define the conditional means and variances  $\mu^{(i)}(X_t) = \mathbb{E}(Y^{(i)}|X_t)$  and  $(\sigma^2)^{(i)}(X_t) = \mathbb{E}[(Y^{(i)} - \mu^{(i)}(X_t))^2|X_t]$  as well as  $f^{(i)}(X_t) = f(\mu^{(i)}(X_t), (\sigma^2)^{(i)}(X_t))$ . As  $\mu^{(i)}(X_t)$  and  $(\sigma^2)^{(i)}(X_t)$  are unknown to the policy maker they must be gradually learned by experimentation. In the presence of covariates a policy  $\pi = \{\pi_t\}$  is a sequence of random functions  $\pi_t : [0, 1]^d \rightarrow \{1, \dots, K + 1\}$  where  $\pi_t$  can only depend on treatment outcomes from previous batches. For any  $X_t$ , a social planner (oracle) who knows the conditional mean and variance functions and wishes to maximize welfare assigns the treatment<sup>5</sup>

$$\pi^*(X_t) \in \arg \max_{i=1, \dots, K+1} f^{(i)}(X_t)$$

and receives  $f^{\pi^*(X_t)}(X_t) = \max_{i=1, \dots, K+1} f^{(i)}(X_t) =: f^*(X_t)$ . Thus,  $f^*(x)$  is the pointwise maximum of the  $f^{(i)}(x)$ ,  $i = 1, \dots, K + 1$ . The goal of a treatment policy is to get as close to the oracle solution as possible in terms of welfare. The welfare loss (regret) of a policy  $\pi$  compared to the oracle is

$$R_N(\pi) = \sum_{t=1}^N (f^{\pi^*(X_t)}(X_t) - f^{\pi_t(X_t)}(X_t)) = \sum_{t=1}^N (f^*(X_t) - f^{\pi_t(X_t)}(X_t)) \quad (3.1)$$

It is important to note the difference between (2.1) and (3.1). (2.1) considers the difference between unconditional means while (3.1) considers the difference between conditional means. The latter is more ambitious as we consider each individual separately through  $X_t$  and seek to minimize the distance to the reward of the outcome that would have been optimal for this specific person (with covariates  $X_t$ ). On the other hand, in the setting without covariates, we only seek to get as close to the outcome of the treatment that is best on average.

In order to prove upper bounds on the regret we restrict the  $\mu^{(i)}(X_t)$  and  $(\sigma^2)^{(i)}(X_t)$  to be reasonably smooth. This is a sensible property to impose since individuals with similar characteristics can be expected to react similarly to the same treatment. In particular, we assume that  $\mu^{(i)}(X_t)$  and

---

<sup>5</sup>If there are several treatments achieving the maximal welfare the oracle assigns any of these.




Figure 1: Four examples of partitioning  $[0, 1]^d$  for  $d = 2$ . The two leftmost ways of grouping individuals correspond to simple rules where group membership is determined by checking whether  $x_1$  and  $x_2$  are above or below certain values. The third rule corresponds to the intersection of two linear eligibility scores  $a_i + b'_i x \geq c_i$ ,  $i = 1, 2$ . The fourth grouping, though not very practically applicable, serves to illustrate that in principle our theory allows for very general ways of grouping individuals.

$\sigma^{(i)}(X_t)$  are  $(\beta, L)$ -Hölder continuous. To be precise, letting  $\|\cdot\|$  denote the Euclidean norm on  $[0, 1]^d$ , we assume that  $\mu^{(i)}, (\sigma^2)^{(i)} \in \mathcal{H}(\beta, L)$  for all  $i = 1, \dots, K+1$ , where for every  $g \in \mathcal{H}(\beta, L)$

$$|g(x) - g(y)| \leq L \|x - y\|^\beta \quad \text{for all } x, y \in [0, 1]^d.$$

### 3.1 Grouping individuals

In the presence of covariates the idea of the dynamic treatment policy is to group individuals into groups according to the values of the covariates. Thus, we define a partition of  $[0, 1]^d$  which consists of Borel measurable sets  $B_1, \dots, B_F$ , called groups/bins, such that  $P_X(B_j) > 0$ ,  $\cup_{j=1}^F B_j = [0, 1]^d$ , and  $B_j \cap B_k = \emptyset$  for  $j \neq k$ . The policy maker groups individuals according to the value of their covariates and seeks to treat each group in a welfare maximizing way. However, the policy maker may be constrained by political or ethical considerations. For example, a realistic unemployment policy cannot group individuals into overly many groups and the rules determining which group an individual belongs to cannot be too complicated. Most realistic policies would choose the groups in such a way that individuals with similar characteristics belong to the same group as it can be expected that the same policy is best for similar individuals. Figure 1 illustrates various ways of grouping individuals.

For any group  $B_j$  define

$$\bar{\mu}_j^{(i)} = \mathbb{E}(Y_t^{(i)} | X_t \in B_j) = \frac{1}{P_X(B_j)} \int_{B_j} \mu^{(i)}(x) dP_X(x)$$

and

$$(\bar{\sigma}^2)_j^{(i)} = \text{Var}(Y_t^{(i)} | X_t \in B_j) = \mathbb{E}(Y_t^{(i)2} | X_t \in B_j) - [\mathbb{E}(Y_t^{(i)} | X_t \in B_j)]^2$$

as the mean and variance of  $Y_t^{(i)}$  given  $X_t$  falls in  $B_j$ . We apply the dynamic treatment policy without covariates separately to each group. To do so, define the groupwise counterpart of the welfare pertaining to treatment  $i$  from the setting without covariates in Section 2 as  $f_j^{(i)} = f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$ . As  $\bar{\mu}_j^{(i)}$  and  $(\bar{\sigma}^2)_j^{(i)}$  vary across groups one can target different optimal treatments for each group,  $j = 1, \dots, F$ . We use the dynamic treatment policy without covariates of Section 2 to target  $\max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  for each group. By the smoothness assumptions on  $f, \mu^{(i)}(x)$  and  $(\sigma^2)^{(i)}(x)$ ,  $\max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  will not be very far from the "fully individualized" target  $f^*(x) = \max_{1 \leq i \leq K+1} f(\mu^{(i)}(x), (\sigma^2)^{(i)}(x))$  as formalized in the appendix. At this stage one may ask why one does not simply use a treatment policy which directly targets  $\max_{1 \leq i \leq K+1} f(\mu^{(i)}(x), (\sigma^2)^{(i)}(x))$ . First, full individualization/discrimination is often not possible due to ethical or legislative constraints. Second, the regret bound obtained in Corollary 3.1.1 for the proposed policy is minimax rate optimal. Thus, even though for each group we target the policy which is best on average for that group nothing is lost (up to a multiplicative constant) even when we compare our performance to the fully individualized optimal policy. Third, a high degree of individualization is only useful for very large data sets as very small groups (in terms of the Lebesgue measure of the group) would otherwise imply very few individuals belonging to each group. This would result in only exploration being carried out for each group as no treatments can be eliminated based on very few assignments. We shall provide an example of how to optimally (in the sense of minimax regret) handle this tradeoff in Corollary 3.1.1 below.

Let  $N_{B_j}(t) = \sum_{s=1}^t 1_{\{X_s \in B_j\}}$  denote the number of individuals who have been assigned to group  $B_j$  when  $t$  individuals have been treated. Furthermore,  $\bar{B}_j = \lambda_d(B_j)$  denotes the Lebesgue measure of group  $j$ . Let  $\hat{\pi}_{B_j, N_{B_j}(t)}$  be the assignment made by the dynamic treatment policy without covariates applied only individuals who belong to group  $B_j$ . This policy is implemented with parameters  $\gamma = \mathcal{K}L$  and  $T = n\bar{B}_j$ . The dynamic treatment policy  $\bar{\pi}$  with covariates is then a sequence of mappings  $\bar{\pi}_t : [0, 1]^d \rightarrow \{1, \dots, K+1\}$  where

$$\bar{\pi}_t(x) = \hat{\pi}_{B_j, N_{B_j}(t)}, \quad x \in B_j$$

Thus, when  $X_t \in B_j$ , the dynamic treatment policy with covariates makes the assignment dictated by the dynamic treatment policy without covariates when applied only to individuals belonging to group  $B_j$ .

### 3.2 Upper and lower bounds on regret

Denote by  $\mathcal{S} = \mathcal{S}(\beta, L, \mathcal{K}, d, \bar{c}, \bar{m})$  a treatment problem where  $f$  is Lipschitz continuous with constant  $\mathcal{K}$ ,  $X_t \in [0, 1]^d$  has distribution  $P_X$  which is absolutely continuous with respect to the Lebesgue measure with density bounded from above by  $\bar{c} > 0$ , maximal batch size  $\bar{m}$  and  $\mu^{(i)}, (\sigma^2)^{(i)} \in \mathcal{H}(\beta, L)$  for all  $i = 1, \dots, K$ . Unless stated otherwise we will consider problems in  $\mathcal{S}$  in the sequel.

The performance of our policy depends critically on the way the policy maker chooses to group individuals. To characterize this grouping, define  $V_j = \sup_{x,y \in B_j} \|x - y\|$  as the maximal possible difference in the characteristics of any two individuals assigned to group  $j$ . The next result provides an upper bound on the regret compared to the infeasible oracle which knows  $\mu^{(i)}(x)$  and  $(\sigma^2)^{(i)}(x)$  and thus which treatment is optimal for an individual with characteristics  $x \in [0, 1]^d$ .

**Theorem 3.1** Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$ . Then, for a grouping characterized by  $\{V_1, \dots, V_F\}$  and  $\{\bar{B}_1, \dots, \bar{B}_F\}$ , expected regret is bounded by

$$\mathbb{E} [R_N(\bar{\pi})] \leq C \sum_{j=1}^F \left[ \sqrt{\bar{m}K \log(\bar{m}K) n \bar{B}_j} + n \bar{B}_j V_j^\beta \right] \quad (3.2)$$

for a positive constant  $C$ . In particular, (3.2) is valid uniformly over  $\mathcal{S}$ .

Theorem 3.1 provides an upper bound on the regret of the dynamic treatment policy for *any* type of grouping of individuals that the policy maker may choose. Allowing for groups with arbitrary characteristics is useful since the policy maker may be constrained in such a way that choosing the groups such that the right hand side of (3.2) is minimized over groups is not possible. The size of the regret depends on the characteristics  $\bar{B}_j$  and  $V_j$  of the grouping. Note that the upper bound on the regret is increasing in these two quantities. However, choosing the groups such that  $\bar{B}_j$  and  $V_j$  are small implies that the number of groups,  $F$ , must be large. In general the upper bound in (3.2) cannot be improved since by choosing the groups as in Corollary 3.1.1 below one achieves the minimax rate of regret. We elaborate further on this below.

The first part of the upper bound in (3.2) is the regret accumulated from implementing the dynamic treatment policy without covariates on each group separately targeting  $\max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  for group  $j = 1, \dots, F$ . Notice the similarity to the second part of Theorem 2.1 once the Lipschitz constant  $\mathcal{K}$  in that result is merged into  $c$  and  $n = E(N)$  is replaced by  $n \bar{B}_j$  which up to the constant  $\bar{c}$  bounds the expected number of observations falling in bin  $j$  from above. This similarity is an implication of using the dynamic treatment policy on each group separately. The second part of the bound in (3.2) is the approximation error resulting from targeting  $\max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  instead of  $\max_{1 \leq i \leq K+1} f(\mu^{(i)}(x), (\sigma^2)^{(i)}(x))$ . Clearly, the larger the groups are chosen (as measured by  $\bar{B}_j$  and  $V_j$ ) the more dissimilar could the treatment that is best for the average individual of the group be from the treatment which is best for any given individual in the group.

A particular type of groups are the quadratic ones which use hard thresholds for each entry of  $X_t$  to create hypercubes that partition  $[0, 1]^d$ . These are particularly relevant in practice due to their simplicity and an example of these bins is given in the second display of Figure 1. More precisely, fix  $M > 0$  and define

$$B_k = \left\{ x \in \mathcal{X} : \frac{k_l - 1}{M} \leq x_l \leq \frac{k_l}{M}, l = 1, \dots, d \right\} \quad (3.3)$$

for  $k = (k_1, \dots, k_d) \in \{1, \dots, M\}^d$ . Thus,  $M$  is the number of splits along each dimension of  $X_t$  creating a partition of  $M^d$  smaller hypercubes  $B_1, \dots, B_{M^d}$  with side lengths  $1/M$ .

**Corollary 3.1.1** Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$ . Set  $M = \lfloor \left( \frac{n}{\bar{m}K \log(\bar{m}K)} \right)^{1/(2\beta+d)} \rfloor$ . Then, expected regret is bounded by

$$\mathbb{E} [R_N(\bar{\pi})] \leq Cn \left( \frac{\bar{m}K \log(\bar{m}K)}{n} \right)^{\frac{\beta}{2\beta+d}}. \quad (3.4)$$

for a positive constant  $C$ . In particular, (3.4) is valid uniformly over  $\mathcal{S}$ .

Note that the larger the number of covariates  $d$ , the smaller will the number of splits  $M$  in each dimension be as it must be ensured that enough observations fall in each group. The larger the number of potential treatments  $K + 1$  is the more experimentation will take place and hence the regret compared to the infeasible oracle policy.

The bound in (3.4) is, as a function of  $n$ , optimal in a minimax sense and cannot be improved by more than multiplicative constants. To see this consider the case of  $\bar{m} = 1$  and  $K = 1$  (two treatments are available) in which case (3.4) reduces to  $E [R_N(\bar{\pi})] \leq cn^{1-\frac{\beta}{2\beta+d}}$ .

**Theorem 3.2** Let  $\bar{m} = 1$  and  $K = 1$ . For any policy  $\pi$

$$\sup_{\mathcal{S}} \mathbb{E} [R_N(\pi)] \geq Cn^{1-\frac{\beta}{2\beta+d}}$$

for some positive constant  $C$ .

Theorem 3.2 shows that up to multiplicative constants no treatment policy can have a lower maximal regret over  $\mathcal{S}$  than the dynamic treatment policy.

### 3.3 Exogenously given groups

Sometimes the groups  $B_1, \dots, B_F$  are dictated exogenously upon the policy maker and thus can not be chosen to maximize welfare as in the previous section. This could be the case due to legislative reasons stating that, say, individuals with similar characteristics must be treated equally. As a result we can no longer target the welfare from the fully individualized policy,  $f^*(x)$ . In our context this means that we must find *one* treatment which best suits all individuals in each of the prespecified groups. For individuals in group  $B_j$ ,  $j = 1, \dots, F$  a candidate for the omnibus best treatment is  $f_j^* = \arg \max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$ , i.e. the treatment maximizing the welfare of a person with average characteristics  $\bar{\mu}_j^{(i)}$  and  $(\bar{\sigma}^2)_j^{(i)}$ . Recalling that  $f_j^{(i)} = f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  and introducing the *modified regret*  $\tilde{R}_j(\bar{\pi}) = \sum_{t=1}^{N_{B_j}(N)} \left( f_j^* - f_j^{(\hat{\pi}_{B_j,t})} \right)$  of group  $B_j$  of the dynamic treatment policy we seek to upper bound

$$\tilde{R}_N(\bar{\pi}) = \sum_{j=1}^F \tilde{R}_j(\bar{\pi}). \quad (3.5)$$



Note that (3.5) differs from the regret in (3.1) in that we no longer target the outcome of the fully individualized treatment.

**Theorem 3.3** Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$ . Then, for a grouping characterized by  $\{V_1, \dots, V_F\}$  and  $\{\bar{B}_1, \dots, \bar{B}_F\}$ , one has

$$\mathbb{E} \left[ \tilde{R}_N(\bar{\pi}) \right] \leq C \sum_{j=1}^F \left[ \sqrt{\bar{m}K \log(\bar{m}K) n \bar{B}_j} \right] \quad (3.6)$$

for a positive constant  $C$ . In particular, (3.6) is valid uniformly over  $\mathcal{S}$ .

The upper bound on modified regret is identical to the one in Theorem 3.1 except for the absence of the term  $c \sum_{j=1}^F \bar{B}_j V_j^\beta$  which previously served as an upper bound on the approximation error  $f^*(x) - f_j^*$  for all  $x \in B_j$ . However, as the groups are now exogenously given, this approximation error is unavoidable and it no longer makes sense to target  $f^*(x)$  – we can no longer choose the characteristics of the groups  $\bar{B}_j$  and  $V_j^\beta$  such that  $n \bar{B}_j V_j^\beta$  is small.

### 3.4 Ethical considerations

We next show that even in the presence of covariates the dynamic treatment policy does not make many suboptimal assignments. Our first result is a straightforward consequence of Theorem 2.2. On any bin  $1 \leq j \leq F$  the following result bounds the number of times that a treatment  $1 \leq i \leq K + 1$  which does not maximize  $f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  is assigned. Let  $T_{i,j}(N)$  be the number of times treatment  $i$  is assigned on bin  $j$  in the course of a total of  $N$  assignments. Calling treatment  $i$  *suboptimal on bin  $B_j$*  if  $\Delta_i := \max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)}) - f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)}) > 0$  we have the following result.

**Theorem 3.4** Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$ . Then, for group  $B_j$  characterized by  $V_j$  and  $\bar{B}_j$

$$\mathbb{E} [T_{i,j}(N)] \leq C \left( \mathcal{K}^2 K \frac{\log \left( \frac{n \bar{B}_j}{\mathcal{K}^2} \right)}{\Delta_i^2} + K \bar{m} + \mathcal{K}^2 \right),$$

for any treatment  $i$  that is suboptimal on bin  $B_j$  and a positive constant  $C$ .

Theorem 3.4 guarantees that any treatment whose combination of mean and variance over  $B_j$  does not maximize  $f$  will only rarely be assigned. In fact, the number of times a treatment that is suboptimal on bin  $B_j$  is assigned only grows logarithmically in the expected number of individuals falling in bin  $B_j$ . Notice also the similarity to Theorem 2.2 where  $n$  has now been replaced by  $n \bar{B}_j$  which up to the constant  $\bar{c}$  is an upper bound on the expected number of individuals falling in group  $j$ .

A potential drawback of Theorem 3.4 is that for each group  $B_j$  the maximizer of  $f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  depends on the way the policy maker has chosen  $B_j$ . A different way of assessing the number of suboptimal treatments assigned is to directly consider each individual and check whether the optimal treatment was assigned or not. We say that treatment  $i$  is *suboptimal for individual  $t$*  if  $f^*(X_t) > f^{(i)}(X_t)$ . Therefore, another way of declaring the fairness of a policy  $\pi$  is by an upper bound on the number of individuals to whom a suboptimal treatment was assigned:

$$S_N(\pi) = \sum_{t=1}^N \mathbb{1}_{\{f^*(X_t) \neq f^{\pi_t}(X_t)\}}$$

It is sensible that a nontrivial upper bound on  $\mathbb{E}(S_N(\pi))$  (a bound less than  $n$ ) can only be established if the best treatment is sufficiently much better than the second best – otherwise these cannot be distinguished from each other. To formalize this notion let

$$f^\sharp(x) = \begin{cases} \max_{i=1, \dots, K+1} \{f^{(i)}(x) : f^{(i)}(x) < f^*(x)\} & \text{if } \min_{i=1, \dots, K+1} f^{(i)}(x) < f^*(x) \\ f^*(x) & \text{otherwise} \end{cases}$$

denote the value of the second best treatment for an individual with characteristics  $x \in [0, 1]^d$ .

**Assumption 1 (Margin condition)** *We say that the margin condition is satisfied with parameter  $\alpha > 0$  if there exists a constant  $C > 0$  and a  $\delta_0 \in (0, 1)$  such that*

$$P(0 < f^*(X_t) - f^\sharp(X_t) < \delta) \leq C\delta^\alpha \quad \forall \delta \in [0, \delta_0]$$

Essentially, the margin condition limits the probability that the best and the second best treatment are very close to each other. Larger values of  $\alpha$  mean that it is easier to distinguish the best and second best treatment from each other<sup>6</sup>. To the best of our knowledge, the margin condition was first used in the literature on statistical treatment rules by Kitagawa and Tetenov (2015) to improve the rates of their empirical welfare maximization classifier. Before this, similar assumptions had been used in the literature on classification analysis, Mammen et al. (1999), Tsybakov (2004b). Perchet and Rigollet (2013) have used the margin condition in the context of bandits. The margin condition is satisfied if, for example,  $f^*(X_t) - f^\sharp(X_t)$  has a density with respect to the Lebesgue measure which is bounded from above by a constant  $a > 0$ . In that case we may set  $C = a$  and  $\alpha = 1$ . We refer to Kitagawa and Tetenov (2015) for more examples of when the margin condition is satisfied.

**Theorem 3.5** *Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$  which also satisfies the margin condition. Then for any policy  $\pi$ ,*

$$\mathbb{E}(S_N(\pi)) \leq Cn^{\frac{1}{1+\alpha}} \mathbb{E} [R_N(\pi)]^{\frac{\alpha}{1+\alpha}} \quad (3.7)$$

<sup>6</sup>Or all treatments are exactly equally good meaning that no sub-optimal treatments exist.

for some positive constant  $C$ . Using the dynamic treatment policy  $\bar{\pi}$  and grouping individuals as in (3.3) yields

$$\mathbb{E}(S_N(\pi)) \leq Cn \left[ \frac{\bar{m}K \log(\bar{m}K)}{n} \right]^{\frac{\alpha\beta}{(1+\alpha)(2\beta+d)}}. \quad (3.8)$$

(3.7) provides an upper bound on the expected number of times a policy  $\pi$  assigns a treatment which is suboptimal for individual  $t$ . This is done in terms of the regret incurred by the policy. (3.8) considers the case of the successive elimination policy with a particular group structure. Note that  $\mathbb{E}(S_N(\pi))$  is guaranteed to grow only sublinearly in  $n$ . However, as  $\alpha$  approaches 0, which amounts to relaxing the margin condition and making the best and second best treatments indistinguishable, the upper bound on  $\mathbb{E}(S_N(\pi))$  becomes almost linear in  $n$ .

### 3.5 Discrete covariates

Until now we have assumed  $P_X$  to be absolutely continuous with respect to the Lebesgue measure. However, many covariates that may influence the identity of the optimal treatment are discrete. For example, gender may affect the outcome of an unemployment program. Furthermore, we may not always observe a continuous variable perfectly as data might only be informative about which of finitely many wealth groups an individual belongs without providing the exact, continuously scaled, wealth.

In order to accommodate discrete covariates, partition  $X_t = (X'_{t,D}, X'_{t,C})'$  where  $X_{t,D} \in A = A_1 \times \dots \times A_{d_D}$  contains the measurements of the  $d_D$  discrete covariates. Each  $A_l \subseteq \mathbb{N}$ ,  $l = 1, \dots, d_D$  is finite with cardinality  $|A_l|$ . For the continuous covariates we assume  $X_{t,C} \in [0, 1]^{d_C}$  such that  $X_t$  is  $(d_D + d_C)$ -dimensional. As in (3.1) the regret of our treatment policy is measured against the infeasible target  $f^*(X_t) = \max_{1 \leq i \leq K+1} f(\mu^{(i)}(X_t), (\sigma^2)^{(i)}(X_t))$ . On the other hand, it does not make sense to assume  $\mu^{(i)}(x) = \mu^{(i)}(x_D, x_C)$  or  $(\sigma^2)^{(i)}(x) = (\sigma^2)^{(i)}(x_D, x_C)$  to be  $(\beta, L)$ -Hölder continuous in  $x_D$ . Thus, discrete covariates must be handled differently from continuous ones. Instead we shall now assume that for each *fixed*  $a \in A$  one has that  $\mu_a^{(i)}(x_C) := \mu^{(i)}(a, x_C)$  and  $(\sigma^2)_a^{(i)}(x_C) := (\sigma^2)^{(i)}(a, x_C)$  belong to  $\mathcal{H}(\beta, L)$ . Since  $a$  can only take  $F_D = |A| = |A_1| \cdot \dots \cdot |A_{d_D}|$  possible values it is without loss of generality to assume  $\beta$  and  $L$  not to depend on  $a$ .

Our treatment policy now works by fully individualizing treatments across the discrete covariates. In other words, for every of the  $F_D$  possible values of the vector of discrete covariates we implement the dynamic treatment policy  $\bar{\pi}$  by constructing groups only based on the continuous variables just as in Section 3.1. For each value of the discrete covariate we allow for different ways of grouping based on the continuous covariates. For example, one may want to construct different wealth groups for men and women in order to obtain, eg, groups with equally many individuals. For each  $a \in A$  let  $B_{a,j}$ ,  $j = 1, \dots, F_a$  be the partition of  $[0, 1]^{d_C}$  used.

Formally, for each  $a \in A$ , let  $\bar{\pi}_{t,a}$  be the dynamic treatment policy with continuous covariates applied to the grouping  $B_{a,j}$ ,  $j = 1, \dots, F_a$ . Thus, the dynamic treatment in the presence of discrete covariates is a mapping  $\tilde{\pi} : A_1 \times \dots \times A_{d_D} \times [0, 1]^{d_C} \rightarrow \{1, \dots, K + 1\}$  where

$$\tilde{\pi}_t(x) = \bar{\pi}_{t,a}(x_C) = \hat{\pi}_{(\{a\} \times B_{a,j}), N_{a,j}(t)}, \quad x_D = a \text{ and } x_C \in B_{a,j}$$

with  $N_{a,j}(t) = \sum_{s=1}^t 1_{(X_{s,D}=a, X_{s,C} \in B_{a,j})}$ . Denote by  $\tilde{\mathcal{S}} = \tilde{\mathcal{S}}(\beta, L, \mathcal{K}, d_C, \bar{m})$  a treatment problem where  $f$  is Lipschitz continuous with constant  $\mathcal{K}$ ,  $X_{t,D} \in A$  is discrete,  $X_{C,t} \in [0, 1]^d$  has distribution  $P_X$  which is absolutely continuous with respect to the Lebesgue measure, maximal batch size  $\bar{m}$  and  $\mu_a^{(i)}, (\sigma^2)_a^{(i)} \in \mathcal{H}(\beta, L)$  for all  $i = 1, \dots, K + 1$ . Letting  $V_{a,j} = \sup_{x,y \in B_{a,j}} \|x - y\|$  we have that  $\tilde{\pi}$  enjoys the following upper bound on regret.

**Theorem 3.6** Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\tilde{\mathcal{S}}$ . Then, if for each  $a \in A$  individuals are grouped as  $\{B_{a,1}, \dots, B_{a,F_a}\}$ , expected regret is bounded by

$$\begin{aligned} \mathbb{E} [R_N(\tilde{\pi})] \leq & C \sum_{a \in A} \sum_{j=1}^{F_a} \left( \sqrt{\bar{m}K \log(\bar{m}K) n \mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})} \right. \\ & \left. + n \mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j}) V_{a,j}^\beta \right). \end{aligned} \quad (3.9)$$

for a positive constant  $C$ . In particular, (3.9) is valid uniformly over  $\tilde{\mathcal{S}}$ .

The upper bound on regret in (3.9) generalizes the upper bounds in Theorems 2.1 (no covariates) and 3.1 (continuous covariates only). For example, the latter follows from (3.9) by letting  $|A| = 1$  and assuming that  $X_{t,C}$  is absolutely continuous with respect to the Lebesgue measure with density bounded from above by  $\bar{c}$ . Also, the case of purely discrete covariates is covered as a special case of (3.9). In that case the approximation error vanishes as  $V_{a,j} = 0$

## 4 Treatment outcomes observed with delay

Oftentimes the outcome of a treatment is only observed with delay. For example, a medical doctor may choose not to measure the effect of a treatment immediately after it has been assigned as it takes time for the treatment to work. However, delaying the measurement for an extended period of time also implies that many new patients will arrive before the outcome of the previous treatment is known. Thus, the type of treatment assigned to these patients must be decided based on less information. Put differently, there is a tradeoff between getting imprecise information now and obtaining precise information later. A similar tradeoff exists when assigning unemployed to job training programs as it takes time to find a job. Therefore, it may not be advisable to measure the effect of a job training program right after its termination.

To formalize this intuition consider the following model of delay. For simplicity, we focus first on the setting without covariates. For every  $Y_t^{(i)}$ , we can decompose

$$Y_t^{(i)} = \mu^{(i)} + \eta_t^{(i)}$$

where  $\mathbb{E}(\eta_t^{(i)}) = 0$ . Since  $Y_t^{(i)}, \mu^{(i)} \in [0, 1]$  it follows that  $\eta_t^{(i)} = Y_t^{(i)} - \mu^{(i)} \in [-1, 1]$ . Thus, without further assumptions, the deviations of  $Y_t^{(i)}$  around its mean are in  $[-1, 1]$ . We shall model the notion of measurements becoming more precise if they are delayed by restricting this interval. To be precise, we assume that

$$\eta_t^{(i)} = Y_t^{(i)} - \mu^{(i)} \in [-\bar{a}_l, \bar{a}_u] \quad (4.1)$$

where  $\bar{a}_l, \bar{a}_u \in [0, 1]$ . In this section we let  $\bar{a}(D) = \bar{a}_u - \bar{a}_l$  be a function of the number of batches  $D$  the measurements are delayed by. Thus, if  $\bar{a}(D)$  is a decreasing function, increasing the delay results in  $Y_t^{(i)}$  being a less noisy measure of  $\mu^{(i)}$ . Restricting the support of  $\eta_t^{(i)}$  is not the only way of modeling that measurements become more precise if they are delayed. One could also let the variance of the  $\eta_t^{(i)}$  be a decreasing function of  $D$ . In fact, any assumption which implies stronger concentration of sample averages around the population means will suffice. As the welfare function  $f$  also depends on the second moment  $\mu_2^{(i)} = \mathbb{E}[Y_t^{(i)2}]$  and since  $Y_t^{(i)2}, \mu_2^{(i)} \in [0, 1]$  arguments identical to the ones above mean that we can model increased measurement precision due to delay as

$$\eta_t^{(i)2} = Y_t^{(i)2} - \mu_2^{(i)} \in [-\bar{a}_l, \bar{a}_u]$$

First, we establish an upper bound on regret of the dynamic treatment policy when treatment outcomes are observed with delay but no covariates are observed.

**Dynamic treatment policy** Denote by  $\hat{\pi}$  the dynamic treatment policy. Let  $\mathcal{I}_b \subseteq \{1, \dots, K+1\}$  be the set of remaining treatments before batch  $b$  and let  $\underline{B}(b) = \min_{i \in \mathcal{I}_b} B_i(b)$  be the number of observations that we have at least made of each of the remaining treatments after batch  $b$ .

1. In each batch  $b = 1, \dots, D-1$  we take turns assigning the treatments  $\{1, \dots, K+1\}$ . No elimination takes place as no outcomes are observed.
2. In each batch  $b = D, \dots, M$  we take turns assigning each remaining treatment (treatments in  $\mathcal{I}_b$ ).
3. At the end of batch  $b = D, \dots, M$  eliminate treatment  $i$  if

$$\max_{i \in \mathcal{I}_b} f(\hat{\mu}_{\underline{B}(b)}^{(i)}, (\hat{\sigma}_{\underline{B}(b)}^2)^{(i)}) - f(\hat{\mu}_{\underline{B}(b)}^{(i)}, (\hat{\sigma}_{\underline{B}(b)}^2)^{(i)}) \geq 16\gamma \sqrt{\frac{2\bar{a}^2}{\underline{B}(b)} \overline{\log} \left( \frac{T}{\underline{B}(b)} \right)}$$

where  $\gamma > 0$ ,  $T \in \mathbb{N}$  and  $\overline{\log}(x) = \log(x) \vee 1$ .

Notice how the dynamic treatment policy in the presence of delay differs from the one without delay. First, no elimination takes place after the first  $D - 1$  batches as no treatment outcomes are observed after these. Second, the elimination rule has been slightly modified as we can now eliminate more aggressively if  $\bar{a}$  is small, i.e. the treatment outcomes are less noisy measurements of the population parameters.

**Theorem 4.1 (No covariates)** *Consider a treatment problem with  $(K + 1)$  treatments and an unknown number of assignments  $N$  with expectation  $n$  that is independent of the treatment outcomes. The treatment outcomes are observed with a delay of  $D$  batches as outlined above. By implementing the dynamic treatment policy with parameters  $\gamma = \mathcal{K}$  and  $T = n$  one obtains the following bound on the expected regret*

$$\mathbb{E} [R_N(\bar{\pi})] \leq C \min \left( \mathcal{K}^2 \bar{a}^2 \sum_{i=1}^K \frac{1}{\Delta_i} \overline{\log} \left( \frac{n \Delta_i^2}{\bar{a}^2} \right) + \bar{m}(K + D), \sqrt{\mathcal{K}^3 \bar{a}^3 \bar{m} K \overline{\log} (\bar{m} K / \mathcal{K} \bar{a})} n + \bar{m}(K + D) \right), \quad (4.2)$$

where  $C$  is a positive constant.

Assume that  $\bar{a} = \bar{a}(D)$  is a decreasing function. Then Theorem 4.1 illustrates the tradeoff between getting imprecise information now and precise information later. This tradeoff is found in the adaptive (first part) as well as the uniform part (second part) of the upper bound on regret of the dynamic treatment policy. Increasing  $D$  directly increases the upper bound on regret since information is obtained later but indirectly decreases the regret via a reduced  $\bar{a}$ . It can also be shown that the bound in Theorem 4.1 reduces to the one in Theorem 2.1 when  $D = 0$  and  $\bar{a} = 1$ .

We turn next to the setting with continuous covariates and treatment outcomes being observed with delay. As in the setting without delay, we implement the dynamic treatment policy separately for each group  $B_1, \dots, B_F$  with parameters  $\gamma = \mathcal{K}L$  and  $T = n\bar{B}_j$ ,  $j = 1, \dots, F$ .

**Theorem 4.2** *Fix  $\beta \in (0, 1]$ ,  $\mathcal{K}, L > 0$ ,  $d \geq 2$  and consider a treatment problem in  $\mathcal{S}$  where the outcomes are observed with a delay of  $D$  batches. Then, for a grouping characterized by  $\{V_1, \dots, V_F\}$  and  $\{\bar{B}_1, \dots, \bar{B}_F\}$ , expected regret is bounded by*

$$\mathbb{E} [R_N(\bar{\pi})] \leq C \left( \sum_{j=1}^F \left[ \sqrt{\bar{m} K \bar{a}^3 \log(\bar{m} K / \bar{a}) n \bar{B}_j} + n \bar{B}_j V_j^\beta + K \bar{m} \right] + \bar{m} D \right). \quad (4.3)$$

for a positive constant  $C$ . In particular, (4.3) is valid uniformly over  $\mathcal{S}$ .

The first part of the upper bound on expected regret in (4.3) (the sum over the  $F$  groups) is identical to the upper bound in Theorem 3.1 except for the presence of  $\bar{a}$ . The smaller  $\bar{a}$  is the

smaller will this part be as the observed outcomes of the treatments will be very close to the population counterparts and the treatment that is best for each group is quickly found. As  $\bar{a}$  is usually a decreasing function in  $D$ , the upper bound in (4.3) clearly illustrates the tradeoff between postponing the measurement to get precise information later and getting (imprecise) information quickly. The term under the square root holds the key to the benefit from delaying as it corresponds to the regret of a treatment problem which starts only after  $D$  batches but where measurements are observed more precisely. On the other hand, the term  $\bar{m}D$  is an upper bound on the regret incurred from assigned individuals blindly for  $D$  batches each of which contains no more than  $\bar{m}$  individuals.

## 5 Assigning unemployed workers to active labor market policies

We illustrate the dynamic treatment policy by using it to assign unemployed individuals to job training programs. Our theory predicts that, even though the optimal treatment is person specific, we can get as close as possible to the welfare that could be attained had we known a priori which treatment is best for each individual. In the present example the welfare that we target is the sum of the probabilities of ending an unemployment spell. Thus, in our general context we are targeting the first moment only, amounting to  $f(\mu, \sigma^2) = \mu$ . Our results are based on the Danish DREAM register data set which contains data on every person who has received certain transfer payments or has participated in one of the active labor market policy programs (treatments). We consider the period from January 2007 to December 2015 and focus on individuals who have received one of the following two treatments: i) additional education type treatments<sup>7</sup> (ii) on the job training<sup>8</sup> This leaves us with 15347 individuals. For each week we observe whether an individual is participating in one of the two treatment programs. From the first time this is the case we wait 18 months before we check whether the individual has found a job. We do so since it takes time for a treatment to work<sup>9</sup>. For each individual we observe the gender and the age and it is sensible to conjecture that the nature of the optimal treatment depends on these. Gender is coded as a discrete variable while age is continuous.

---

<sup>7</sup>This includes regular education as well as 6-week educational courses chosen by the unemployed individual.

<sup>8</sup>This includes salary subsidies in either public or private companies.

<sup>9</sup>To be precise, it is only measured on a monthly basis whether an individual has found a job or not. Thus, the weekly observations on whether one is participating in a treatment program is uniquely assigned to a month.

## 5.1 Practical implementation

Since the employment status is registered monthly this is a natural batch size. Furthermore, as we wait 18 months until we register whether an individual has found a job or not this amounts to a delay of 18 batches in the terminology introduced in Section 4. As long as no treatment has been eliminated, we ensure that each of the treatments is assigned equally often. Thus, if in the data set additional education has been assigned to  $n_{1,b}$  individuals in month (batch)  $b$  while job training has been assigned to  $n_{2,b}$  individuals in the same month with  $n_{1,b} < n_{2,b}$  then we base our calculation of sample averages on  $n_{1,b}$  individuals for each treatment. Out of the  $n_{2,b}$  individuals who have received job training in month  $b$  we choose  $n_{1,b}$  at random. Furthermore, we choose  $M = 2$  meaning that age is divided into two groups: 0-42 year olds and 43 and upwards. This particular split is chosen as 42 is the median age in the sample. Thus, there are four groups in total and we set  $T = 15000/4$  to reflect that we expect about equally many individuals to fall in each group. Note that we have implicitly assumed that we have a decent estimate of the final sample size,  $N = 15347$  since 15000 is a “good” approximation to  $N$ . We choose the threshold for eliminating treatments as  $2A\sqrt{\frac{2}{s}\ln\left(\frac{T}{s}\right)}$  for a scaling parameter  $A$ . This follows from the choice of threshold in Section 4 with  $\bar{a} = 2$  (since the treatment outcome takes the value 0 or 1, cf. the discussion around display (4.1)) and only targeting the mean such that  $\mathcal{K} = 1$  as well as inspection of the proof of Theorem 4.1 revealing that the constants in the definition of the threshold can be reduced when one only targets the mean outcome of a treatment. We choose  $A = 0.25$  in the sequel to illustrate how elimination of treatments takes place as the theoretically justified choice of  $A = 1$  resulted in both treatments being retained for all four groups.

## 5.2 Results

To get an idea of the overall efficiency of the two treatments Table 1 contains the fraction of individuals who have found a job at the time of first elimination of a treatment (left panel). The table also contains the corresponding numbers at the end of the treatment period for the often hypothetical case of none of the two treatments being eliminated (right panel). We include the latter statistic to investigate whether one might at the end of the treatment period regret a previous elimination decision. These numbers are reported separately for each of the four groups. No matter at which of the two times we measure the success rate, it is around 50% for both treatments. It is “as expected” that the difference in the success rates is larger at the time of first elimination than at the end of the treatment period since for the former we are conditioning on crossing the threshold for elimination.

Figure 2 illustrates the difference in the success rates for the two treatments for each of the four groups. It also indicates the thresholds for eliminating either one of the two treatments. In particular, if the difference in success rates becomes higher than the upper threshold then the dynamic



	Women		Men			Women		Men	
	$\leq 42$	$> 42$	$\leq 42$	$> 42$		$\leq 42$	$> 42$	$\leq 42$	$> 42$
Age					Age				
Educ	-	0.5475	0.4894	0.5433	Educ	0.5200	0.4876	0.5615	0.5000
JobTrain	-	0.4842	0.5356	0.4685	JobTrain	0.5233	0.5111	0.5804	0.4854

Table 1: Success rates for the extra education and job training treatment at the time of first elimination (left panel) of a treatment and at the time of the last treatment assignment (right panel).

treatment policy deems the job training programs to be inferior to education. If the difference drops below the lower threshold then education is eliminated and is no longer assigned. The top left panel of Figure 2 contains the results for women below the age of 42. The two treatments are almost equally successful throughout the whole assignment period which is witnessed by the absence of an elimination of either of the treatments. Furthermore, the difference in the success rates of the two treatments is never larger than 4% and after 1400 assignments have been made (end of treatment period) the difference is less than 0.5% (cf Table 1). Turning to the top right panel, which contains the results for women above 42 years of age, the picture is very different: after a little more than 200 assignments the success rate for the extra education treatment is sufficiently much higher than the one for job training in order to eliminate the latter. Thus, from this point and onwards, individuals belonging to this group are only assigned to extra education. However, we still plot the difference between the success rates of the two treatments as if none of them had been eliminated. Interestingly, we see that this difference actually crosses the threshold again shortly after job training is eliminated and it even becomes so negative at the end of the treatment period that one would have eliminated education had this stage been reached. Turning to the bottom left panel, which contains the results for men below the age of 42, we see that extra education is reasonably strongly eliminated in the sense that the difference between the success rates almost stays below the bottom threshold for the remaining treatment period once it goes below it. The bottom right panel contains the results for men which are older than 42. For these we find a markedly different picture as it is the job training treatment that is eliminated. However, the elimination is much weaker than the one for the young men as the difference in success rates quickly goes below the upper threshold again and stays there for a long period. To summarise, there is a tendency towards extra education being more successful for the persons who are older than 42 years of age irrespective of their gender. This could be due to the fact that these individuals have had a higher deterioration of their skill set than those who have recently finished their education. Hence, the older individuals will gain more from updating their skill set.

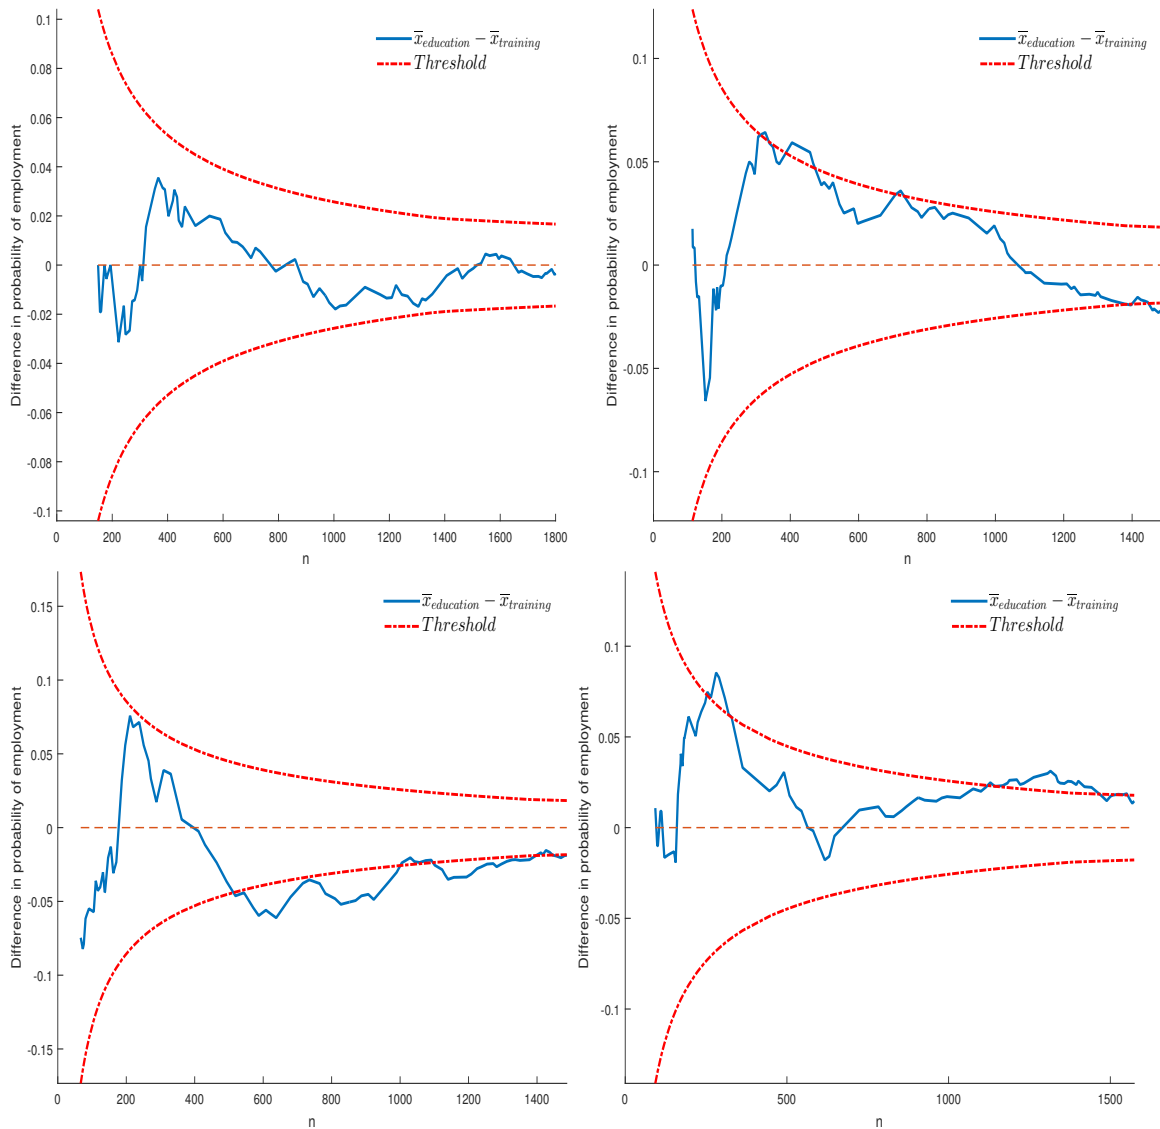


Figure 2: Top left panel: women of age less than 42, top right panel: women of age greater than 42, bottom left panel: men of age less than 42, bottom right: panel men of age greater than 42. The blue line is the success rate of education minus the success rate of job training. The red lines indicate the thresholds for when a treatment gets eliminated. If the first crossing of the blue and red lines takes the blue line above the upper red line, job training is eliminated while if the first crossing takes the blue line below the lower red line, education is eliminated.

## 6 Conclusions

This paper considers a treatment allocation problem where the individuals to be treated arrive gradually and potentially in batches. The goal of the policy maker is to maximize the welfare over the  $N$  treatment assignments made. As the policy maker does not know a priori about the virtues of the available treatments he faces an exploration-exploitation tradeoff. Prior to each assignment he observes covariates on the individual to be treated thus allowing for the optimal treatment to vary across individuals. Our setup allows the welfare function not only to depend on the expected treatment outcome but also on the risk of the treatment. We show that a variant of the dynamic treatment policy obtains the minimax optimal regret. This strong welfare guarantee does not come at the price of overly wild experimentation as we show that the number of suboptimal treatments only grows quite slowly in the total number of assignments made. We also establish upper bounds on the regret of the dynamic treatment policy when the outcome of the treatments are only observed with delay.

## 7 Appendix

Throughout the appendix we let  $C$  be a positive constant that may change from line to line.

### 7.1 Proof of Theorems 2.1 and 2.2

The following lemma will lead to Theorem 2.1.

**Lemma 7.1** *Consider a treatment problem with  $(K + 1)$  treatments and unknown number of assignments  $N$  with expectation  $n$  that is independent of the treatment outcomes. Suppose that  $f$  is Lipschitz continuous with known constant  $\mathcal{K}$ . For any  $\Delta > 0$ ,  $T > 0$  and  $\gamma \geq \mathcal{K}$  the expected regret from running the dynamic treatment policy can then be bounded as*

$$\mathbb{E} [R_N(\hat{\pi})] \leq C \left( \frac{\gamma^2 K}{\Delta} \left( 1 + \frac{n}{T} \right) \overline{\log} \left( \frac{T \Delta^2}{4608 \gamma^2} \right) + n \Delta^- + \frac{n \bar{m} K}{T} \right), \quad (7.1)$$

where  $\Delta^-$  is the largest  $\Delta_j$  such that  $\Delta_j < \Delta$  if such a  $\Delta_j$  exists, and  $\Delta^- = 0$  otherwise.

**Proof.** Define  $\epsilon_s = u(s, T) = 32 \sqrt{\frac{2}{s} \overline{\log} \left( \frac{T}{s} \right)}$ . Recall that if the optimal treatment as well as some treatment  $i$  have not been eliminated before batch  $b$  (i.e.,  $i, * \in \mathcal{I}_b$ ), then the optimal treatment will eliminate treatment  $i$  if  $\hat{\Delta}_i(\underline{B}(b)) \geq \gamma \epsilon_{\underline{B}(b)}$ , and treatment  $i$  will eliminate the optimal treatment if  $\hat{\Delta}_i(\underline{B}(b)) \leq -\gamma \epsilon_{\underline{B}(b)}$ .

To say something about when either of these two events occurs we introduce the (unknown) quantity  $\tau_i^*$  which is defined through the relation

$$\Delta_i = 48\gamma\sqrt{\frac{2}{\tau_i^*}\log\left(\frac{T}{\tau_i^*}\right)}, \quad i = 1, \dots, K.$$

Since  $\tau_i^*$  in general will not be an integer, we also define  $\tau_i = \lceil \tau_i^* \rceil$ . Next introduce the hypothetical batch  $b_i = \min\{l : \underline{B}(l) \geq \tau_i^*\}$ . It is the first batch after which we have more than  $\tau_i^*$  observations on all remaining treatment. Notice that

$$\tau_i^* \leq \underline{B}(b_i) \leq \tau_i^* + \bar{m} \leq C\frac{\gamma^2}{\Delta_i^2}\log\left(\frac{T\Delta_i^2}{4608\gamma^2}\right) + \bar{m}, \quad (7.2)$$

$$\tau_i \leq \underline{B}(b_i), \quad (7.3)$$

$$\underline{B}(b_i) \leq \tau_i + \bar{m}, \quad (7.4)$$

Notice that  $1 \leq \tau_1 \leq \dots \leq \tau_K$  and  $1 \leq b_1 \leq \dots \leq b_K$ . Define the following events:

$\mathcal{A}_i = \{\text{The optimal treatment has not been eliminated before batch } b_i\}$ ,

$\mathcal{B}_i = \{\text{Every treatment } j \in \{1, \dots, i\} \text{ has been eliminated after batch } b_j\}$ .

Furthermore, let  $\mathcal{C}_i = \mathcal{A}_i \cap \mathcal{B}_i$ , and observe that  $\mathcal{C}_1 \supseteq \dots \supseteq \mathcal{C}_K$ . For any  $i = 1, \dots, K$ , the contribution to regret incurred after batch  $b_i$  is at most  $\Delta_{i+1}N$  on  $\mathcal{C}_i$ . In what follows we fix a treatment,  $K_0$ , which we will have more to say about later. Using this we get the following decomposition of expected regret:

$$\begin{aligned} \mathbb{E}[R_N(\hat{\pi})] &= \mathbb{E}\left[R_N(\hat{\pi})\left(\sum_{i=1}^{K_0} 1_{\mathcal{C}_{i-1} \setminus \mathcal{C}_i} + 1_{\mathcal{C}_{K_0}}\right)\right] \\ &\leq n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \setminus \mathcal{C}_i) + \sum_{i=1}^{K_0} B_i(b_i) \Delta_i + n \Delta_{K_0+1}. \end{aligned} \quad (7.5)$$

For every  $i = 1, \dots, K$  the event  $\mathcal{C}_{i-1} \setminus \mathcal{C}_i$  can be decomposed as  $\mathcal{C}_{i-1} \setminus \mathcal{C}_i = (\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) \cup (\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1})$ .

Therefore, the first term on the right-hand side of (7.5) can be written as

$$n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \setminus \mathcal{C}_i) = n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) + n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}). \quad (7.6)$$

Notice that  $\mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) = 0$  if  $b_{i-1} = b_i$ . On the event  $\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}$  the optimal treatment has not eliminated treatment  $i$  after batch  $b_i$ . Therefore, for the last term on the right hand side of equation (7.6) we find that

$$\mathbb{P}(\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}) \leq \mathbb{P}\left(\hat{\Delta}_i(\underline{B}(b_i)) \leq \gamma \epsilon_{\underline{B}(b_i)}\right)$$

$$\begin{aligned}
&\leq \mathbb{P} \left( \hat{\Delta}_i(\underline{B}(b_i)) - \Delta_i \leq \gamma \epsilon_{\tau_i} - \Delta_i \right) \\
&\leq \mathbb{E} \left[ \mathbb{P} \left( |\hat{\Delta}_i(\underline{B}(b_i)) - \Delta_i| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \mid \underline{B}(b_i) \right) \right]
\end{aligned}$$

For any  $s \geq \tau_i$  we have that

$$\begin{aligned}
&\mathbb{P} \left( |\hat{\Delta}_i(s) - \Delta_i| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\hat{\mu}_s^{(i)}, (\hat{\sigma}_s^2)^{(i)}) + f(\mu^{(i)}, (\sigma^2)^{(i)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |f(\hat{\mu}_s^{(i)}, (\hat{\sigma}_s^2)^{(i)}) - f(\mu^{(i)}, (\sigma^2)^{(i)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right). \tag{7.7}
\end{aligned}$$

Furthermore, for any  $j \in \{i, *\}$ , we have

$$\begin{aligned}
&\mathbb{P} \left( |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| + |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{4\mathcal{K}} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right). \tag{7.8}
\end{aligned}$$

By the mean value theorem we have that

$$\mathbb{P} \left( |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right) \leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{32\mathcal{K}} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |(\hat{\mu}_{2,s})^{(j)} - \mu_2^{(j)}| \geq \frac{1}{16\mathcal{K}} \gamma \epsilon_{\tau_i} \right), \tag{7.9}$$

where  $\mu_2 = \mathbb{E} [Y_1^2]$  and  $\hat{\mu}_{2,s} = \frac{1}{s} \sum_{i=1}^s Y_i^2$ . By combining equations (7.7), (7.8), (7.9), and applying Hoeffding's inequality as well as the fact that  $\gamma \geq \mathcal{K}$ , we arrive at the following bound,

$$\begin{aligned}
\mathbb{P} \left( |\hat{\Delta}_i(s) - \Delta_i| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \right) &\leq C \exp \left( -\frac{1}{1024} \epsilon_{\tau_i}^2 s \right) \\
&\leq C \exp \left( -\frac{1}{1024} \epsilon_{\tau_i}^2 \tau_i \right) \\
&= C \exp \left( -\overline{\log} \left( \frac{T}{\tau_i} \right) \right) \\
&\leq C \frac{\tau_i}{T}.
\end{aligned}$$

Thus,

$$\mathbb{P} (\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}) \leq C \frac{\tau_i}{T}$$

On the event  $\mathcal{C}_{i-1} \cap \mathcal{A}_i^c$  the optimal treatment is eliminated between batch  $b_{i-1} + 1$  and  $b_i$ . Furthermore, every suboptimal treatment  $j \leq i - 1$  has also been eliminated. Therefore the probability of this event can be bounded as follows:

$$\begin{aligned} \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) &\leq \mathbb{P}\left(\exists(j, s), i \leq j \leq K, b_{i-1} + 1 \leq s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &\leq \sum_{j=i}^K \mathbb{P}\left(\exists s, b_{i-1} + 1 \leq s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &= \sum_{j=i}^K [\Phi_j(b_i) - \Phi_j(b_{i-1})], \end{aligned}$$

where  $\Phi_j(b) = \mathbb{P}\left(\exists s \leq b; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right)$ . We now proceed to bound terms of the form  $\Phi_j(b_i)$  for  $j \geq i$ .

$$\begin{aligned} \mathbb{P}\left(\exists s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) &\leq \mathbb{P}\left(\exists s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) - \Delta_j \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &\leq \mathbb{P}\left(\exists s \leq \underline{B}(b_i); \hat{\Delta}_j(s) - \Delta_j \leq -\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; \hat{\Delta}_j(s) - \Delta_j \leq -\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{2}\gamma\epsilon_s\right) \\ &\quad + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{2}\gamma\epsilon_s\right). \end{aligned}$$

For any  $j \in \{i, \dots, K, *\}$  we find that

$$\begin{aligned} &\mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{2}\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |(\hat{\mu}_{2,s}^{(j)}) - \mu_2^{(j)}| \geq \frac{1}{8\mathcal{K}}\gamma\epsilon_s\right) \\ &\quad + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{16\mathcal{K}}\gamma\epsilon_s\right) \\ &\leq C \frac{\tau_i + \bar{m}}{T} \end{aligned}$$

where we have used equation (7.4) and Lemma A.1 in Perchet and Rigollet (2013). It follows that

$$\begin{aligned} \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) &\leq \sum_{i=1}^{K_0} \Delta_i \sum_{j=i}^K [\Phi_j(b_i) - \Phi_j(b_{i-1})] \\ &\leq \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \Phi_j(b_i) (\Delta_i - \Delta_{i+1}) + \sum_{j=K_0}^K \Delta_{K_0} \Phi_j(b_{K_0}) + \sum_{j=1}^{K_0-1} \Delta_j \Phi_j(b_j) \end{aligned}$$

$$\leq \frac{48}{T} \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} (\tau_i + \bar{m}) (\Delta_i - \Delta_{i+1}) + \frac{48}{T} \sum_{j=1}^K \Delta_{j \wedge K_0} (\tau_{j \wedge K_0} + \bar{m}). \quad (7.10)$$

Using equation (7.3) we obtain

$$\begin{aligned} \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \tau_i (\Delta_i - \Delta_{i+1}) &\leq C \gamma^2 \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \frac{(\Delta_i - \Delta_{i+1})}{\Delta_i^2} \overline{\log} \left( \frac{T \Delta_i^2}{4608 \gamma^2} \right) \\ &\leq C \gamma^2 \sum_{j=1}^K \int_{\Delta_{j \wedge K_0}}^{\Delta_1} \frac{1}{x^2} \overline{\log} \left( \frac{T x^2}{4608 \gamma^2} \right) dx \\ &\leq C \gamma^2 \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{4608 \gamma^2} \right). \end{aligned}$$

The part involving  $\bar{m}$  in equation (7.10) can be bounded by

$$\bar{m} \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} (\Delta_i - \Delta_{i+1}) + \sum_{j=1}^K \Delta_{j \wedge K_0} \bar{m} \leq \bar{m} K.$$

Bringing things together we have

$$\sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) \leq C \left( \frac{\gamma^2}{T} \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{4608 \gamma^2} \right) + \frac{K \bar{m}}{T} \right) \quad (7.11)$$

Combining this with equation (7.6) and (7.5) we obtain

$$\begin{aligned} \mathbb{E} [R_N(\hat{\pi})] &\leq C \left( \frac{\gamma^2 n}{T} \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{4608 \gamma^2} \right) + \frac{n \gamma^2}{T} \sum_{j=1}^{K_0} \frac{1}{\Delta_j} \overline{\log} \left( \frac{T \Delta_j^2}{4608 \gamma^2} \right) \right. \\ &\quad \left. + \sum_{i=1}^{K_0} B_i(b_i) \Delta_i + n \Delta_{K_0+1} + \frac{n \bar{m} K}{T} \right) \\ &\leq C \left( \left( 1 + \frac{n}{T} \right) \gamma^2 \sum_{j=1}^{K_0} \frac{1}{\Delta_j} \overline{\log} \left( \frac{T \Delta_j^2}{4608 \gamma^2} \right) + \frac{\gamma^2 n}{T} \frac{K - K_0}{\Delta_{K_0}} \overline{\log} \left( \frac{T \Delta_{K_0}^2}{4608 \gamma^2} \right) \right. \\ &\quad \left. + n \Delta_{K_0+1} + \frac{n \bar{m} K}{T} \right). \quad (7.12) \end{aligned}$$

Fix  $\Delta > 0$  and let  $K_0$  be such that  $\Delta_{K_0+1} = \Delta^-$ . Define the function

$$\phi(x) = \frac{1}{x} \overline{\log} \left( \frac{T x^2}{4608 \gamma^2} \right),$$

and notice that  $\phi(x) \leq 2e^{-1/2}\phi(x')$  for any  $x \geq x' \geq 0$ . Using this with  $x' = \Delta$  and  $x = \Delta_i$  for  $i \leq K_0$  we obtain

$$\mathbb{E} [R_N(\hat{\pi})] \leq C \left( \frac{\gamma^2 K}{\Delta} \left(1 + \frac{n}{T}\right) \overline{\log} \left( \frac{T\Delta^2}{4608\gamma^2} \right) + n\Delta^- + \frac{n\bar{m}K}{T} \right). \quad (7.13)$$

■

**Proof of Theorem 2.1.** Consider the dynamic treatment policy with  $\gamma = \mathcal{K}$  and  $T = n$ . From equation (7.12) it follows that for any  $K_0 \leq K$

$$\begin{aligned} \mathbb{E} [R_N(\hat{\pi})] &\leq C \left( \mathcal{K}^2 \sum_{j=1}^{K_0} \frac{1}{\Delta_j} \overline{\log} \left( \frac{n\Delta_j^2}{4608\mathcal{K}^2} \right) \right. \\ &\quad \left. + \mathcal{K}^2 \frac{K - K_0}{\Delta_{K_0}} \overline{\log} \left( \frac{n\Delta_{K_0}^2}{4608\mathcal{K}^2} \right) + n\Delta_{K_0+1} + \frac{n\bar{m}K}{T} \right). \end{aligned} \quad (7.14)$$

This can be used to show

$$\mathbb{E} [R_N(\hat{\pi})] \leq C \min \left( \mathcal{K}^2 \sum_{j=1}^K \frac{1}{\Delta_j} \overline{\log} \left( \frac{n\Delta_j^2}{4608\mathcal{K}^2} \right) + \bar{m}K, \sqrt{n\mathcal{K}^3\bar{m}K \log(\bar{m}K/\mathcal{K})} \right). \quad (7.15)$$

where the first part of the upper bound in (7.15) follows by using (7.14) with  $K = K_0$ . The second part follows from Lemma 7.1 by choosing  $\Delta = \sqrt{4608(23757 + \bar{m})\mathcal{K}K \log((23757 + \bar{m})K/\mathcal{K})/n}$ .

**Proof of Theorem 2.2.** The idea of the proof is similar to the one used in the proof of Theorem 1 in Auer and Cesa-Bianchi (2002). For now we will keep  $N$  fixed.<sup>10</sup> First, note that for any positive integer  $l$

$$T_i(N) = 1 + \sum_{t=K+1}^N \mathbf{1}_{\{\hat{\pi}_t=i\}} \leq l + \sum_{t=K+1}^N \mathbf{1}_{\{\hat{\pi}_t=i, T_i(t-1) \geq l\}} \leq l + \sum_{t=K+1}^N \mathbf{1}_{\{T_i(t-1) \geq l\}} \leq l + N\mathbf{1}_{\{T_i(N-1) \geq l\}}$$

It remains to bound the probability of the event  $\{T_i(N-1) \geq l\}$ . This is the probability that treatment  $i$  has not been eliminated before having been assigned at least  $l$  times. Define  $\tilde{b}_i = \max\{b : \sum_{j=1}^b m_{i,j} < l\}$  and note that if treatment  $i$  is assigned  $l$  times then it cannot have been eliminated after  $\tilde{b}_i$  batches. In particular, it cannot have been eliminated by the optimal treatment. Let

$$\mathcal{A}_i = \{\text{the optimal treatment has not been eliminated after batch } \tilde{b}_i\}$$

For any  $t$  we have that  $\{T_i(N-1) \geq l\} \subseteq (\{T_i(N-1) \geq l\} \cap \mathcal{A}_i) \cup \mathcal{A}_i^c$ . Thus,  $(\{T_i(N-1) \geq l\} \cap \mathcal{A}_i) \subseteq \{\hat{\Delta}_i(\underline{B}(\tilde{b}_i)) - \Delta_i \leq \gamma \epsilon_{\underline{B}(\tilde{b}_i)} - \Delta_i\}$  which implies

$$\mathbb{E} [T_i(N)] \leq l + N\mathbb{P}(\{T_i(N-1) \geq l\} \cap \mathcal{A}_i) + N\mathbb{P}(\mathcal{A}_i^c)$$

<sup>10</sup>In other words all calculations are done conditional on  $N$ .



$$\leq l + N\mathbb{P}\left(\hat{\Delta}_i(\underline{B}(\tilde{b}_i)) - \Delta_i \leq \gamma\epsilon_{\underline{B}(\tilde{b}_i)} - \Delta_i\right) + N\mathbb{P}(\mathcal{A}_i^c)$$

From equations (7.2) and (7.3) of Lemma 7.1 we have that (where  $\tau_i$  is defined in the said lemma)

$$\tau_i \leq \frac{C\mathcal{K}^2}{\Delta_i^2} \overline{\log}\left(\frac{n}{4609\mathcal{K}^2}\right)$$

Thus, by letting  $l = \bar{m} + \lceil \frac{4609\mathcal{K}^2}{\Delta_i^2} \overline{\log}\left(\frac{n}{4609\mathcal{K}^2}\right) \rceil$  it follows that  $\tau_i \leq l - \bar{m} \leq \underline{B}(\tilde{b}_i) < l$ . In particular, we have that  $\gamma\epsilon_{\underline{B}(\tilde{b}_i)} \leq \gamma\epsilon_{\tau_i} \leq \frac{2}{3}\Delta_i$ . Hence,

$$\begin{aligned} \mathbb{P}\left(\hat{\Delta}_i(\underline{B}(\tilde{b}_i)) - \Delta_i \leq \gamma\epsilon_{\underline{B}(\tilde{b}_i)} - \Delta_i\right) &\leq \mathbb{P}\left(|\hat{\Delta}_i(\underline{B}(\tilde{b}_i)) - \Delta_i| \geq \frac{1}{3}\Delta_i\right) \\ &\leq C\mathbb{E}\left[\exp\left(-\frac{\underline{B}(\tilde{b}_i)\Delta_i^2}{1536}\right)\right] \\ &\leq C\exp\left(-\frac{(l - \bar{m})\Delta_i^2}{1536}\right) \\ &\leq C\frac{\mathcal{K}^2}{n}. \end{aligned}$$

Next, we bound the term involving  $\mathcal{A}_i^c$ . To this end we start by noting that if the optimal treatment does not survive until batch  $\tilde{b}_i$ , then it must have been eliminated in one of the batches before  $\tilde{b}_i$ .

$$\begin{aligned} \mathbb{P}(\mathcal{A}_i^c) &\leq \sum_{j=1}^K \mathbb{P}\left(\exists s \leq \underline{B}(\tilde{b}_i) : \hat{\Delta}_j(s) \leq -\gamma\epsilon_s\right) \\ &\leq \sum_{j=1}^K \mathbb{P}\left(\exists s \leq l : \hat{\Delta}_j(s) \leq -\gamma\epsilon_s\right) \\ &\leq CK\frac{l}{n}, \end{aligned}$$

where the last inequality follows from an application of Lemma A.1 in Perchet and Rigollet (2013). Bringing things together, taken expectations with respect to  $N$  and using Jensen's inequality in order to replace  $N$  with its expectation yields the desired result.

## 7.2 Proof of Theorems in Section 3

**Proof of Theorem 3.1.** It is convenient to define the constant  $c_1 = 6L\mathcal{K} + 1$ , which will enter several of the bounds derived below. Furthermore, we let  $c$  denote a positive constant which may change from line to line. By the construction of the treatment policy it follows that the regret can be written as  $R_N(\bar{\pi}) = \sum_{j=1}^F R_j(\bar{\pi})$ , where

$$R_j(\bar{\pi}) = \sum_{t=1}^N \left( f^*(X_t) - f^{(\hat{\pi}_{B_j}, N_{B_j}(t))}(X_t) \right) 1_{(X_t \in B_j)}.$$

We start by providing an upper bound on the welfare lost for each group  $B_j$  due to the policy targeting  $f_j^* = \max_{1 \leq i \leq K} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)})$  instead of  $f^*(x)$ . To this end note that

$$\begin{aligned} f^*(x) &= \max_{1 \leq i \leq K+1} f(\mu^{(i)}(x), (\sigma^2)^{(i)}(x)) \\ &\leq \max_{1 \leq i \leq K+1} f(\bar{\mu}_j^{(i)}, (\bar{\sigma}^2)_j^{(i)}) + \mathcal{K} \max_{1 \leq i \leq K+1} |\mu^{(i)}(x) - \bar{\mu}_j^{(i)}| + \mathcal{K} \max_{1 \leq i \leq K+1} |(\sigma^2)^{(i)}(x) - (\bar{\sigma}^2)_j^{(i)}|. \end{aligned}$$

Fix  $x \in B_j$  and  $i \in \{1, \dots, K+1\}$ . Then, for all  $y \in B_j$ , one has by the  $(L, \beta)$ -Hölder continuity of  $\mu^{(i)}(x)$

$$\mu^{(i)}(x) \leq \mu^{(i)}(y) + |\mu^{(i)}(x) - \mu^{(i)}(y)| \leq \mu^{(i)}(y) + LV_j^\beta,$$

which upon integrating over  $y$  yields  $\mu^{(i)}(x) \leq \bar{\mu}_j^{(i)} + LV_j^\beta$ . Similarly, it holds that  $\mu^{(i)}(x) \geq \bar{\mu}_j^{(i)} - LV_j^\beta$  such that for all  $x \in B_j$  we have  $|\mu^{(i)}(x) - \bar{\mu}_j^{(i)}| \leq LV_j^\beta$ . Next, note that the map  $[0, 1] \ni z \mapsto z^2$  is Lipschitz continuous with constant 2 which implies that  $(\mu^{(i)}(x))^2$  is  $(2L, \beta)$ -Hölder. This, together with the  $(L, \beta)$ -Hölder continuity of  $(\sigma^2)^{(i)}(x) = \mathbb{E}(Y_t^{(i)2} | X_t = x) - (\mu^{(i)}(x))^2$  implies that  $\mathbb{E}(Y_t^{(i)2} | X_t = x)$  is  $(3L, \beta)$ -Hölder continuous. Thus, by similar arguments as above  $|\mathbb{E}(Y_t^{(i)2} | X_t = x) - \mathbb{E}(Y_t^{(i)2} | X_t \in B_j)| \leq 3LV_j^\beta$  for all  $x \in B_j$ . The mean value theorem also yields that  $|(\mu^{(i)}(x))^2 - \bar{\mu}_j^{(i)2}| \leq 2LV_j^\beta$  for all  $x \in B_j$ . Therefore,

$$\begin{aligned} |(\sigma^2)^{(i)}(x) - (\bar{\sigma}^2)_j^{(i)}| &= |\mathbb{E}(Y_t^{(i)2} | X_t = x) - (\mu^{(i)}(x))^2 - [\mathbb{E}(Y_t^{(i)2} | X_t \in B_j) - \bar{\mu}_j^{(i)2}]| \\ &\leq 5LV_j^\beta \end{aligned}$$

Thus, for  $x \in B_j$ ,

$$f^*(x) \leq f_j^* + c_1 V_j^\beta.$$

A similar argument to the above yields that for all  $x \in B_j$

$$f^{(\bar{\pi}_t)}(x) \geq \bar{f}_j^{(\bar{\pi}_t)} - c_1 V_j^\beta.$$

Next we define  $\tilde{R}_j(\bar{\pi}) = \sum_{t=1}^{N_{B_j}(N)} \left( f_j^* - f_j^{(\hat{\pi}_{B_j, t})} \right)$ . This corresponds to the regret associated with a treatment problem without covariates where treatment  $i$  yields reward  $\bar{f}_j^{(i)}$ , and the best treatment yields  $f_j^* = \max_i \bar{f}_j^{(i)} \leq \bar{f}_j^*$ . Therefore, we can write

$$\begin{aligned} R_j(\bar{\pi}) &= \sum_{t=1}^N \left( f^*(X_t) - f^{(\hat{\pi}_{B_j, N_{B_j}(t)})}(X_t) \right) \mathbf{1}_{(X_t \in B_j)} \leq \sum_{t=1}^N \left( f_j^* - \bar{f}_j^{(\hat{\pi}_{B_j, N_{B_j}(t)})} + 2c_1 V_j^\beta \right) \mathbf{1}_{(X_t \in B_j)} \\ &= \tilde{R}_j(\bar{\pi}) + 2c_1 V_j^\beta N_{B_j}(N), \end{aligned}$$

where  $N_{B_j}(N)$  is the number of observations falling in bin  $j$  given that there are  $N$  observations in total. Taking expectations, and using that the density of  $X_t$  is bounded from above implies that  $\mathbb{E}[N_j(N)] \leq \bar{c} n \bar{B}_j$  gives

$$\mathbb{E}[R_j(\bar{\pi})] \leq \mathbb{E}[\tilde{R}_j(\bar{\pi})] + n \bar{c} \bar{B}_j c_1 V_j^\beta. \quad (7.16)$$

Since  $\mathbb{E} [\tilde{R}_j(\bar{\pi})]$  is the expected regret of a treatment problem without covariates we can apply Theorem 7.1 with the following values

$$\Delta = \sqrt{\frac{\bar{m}K \log(\bar{m}K)}{n\bar{B}_j}}, \quad \gamma = \mathcal{K}L, \quad T = n\bar{B}_j,$$

for each bin  $j = 1, \dots, F$  to obtain the following bound on the regret accumulated across any group  $j$ :

$$\mathbb{E} [R_j(\bar{\pi})] \leq C \left[ \sqrt{\bar{m}K \log(\bar{m}K)n\bar{B}_j} + n\bar{B}_j V_j^\beta \right].$$

Thus, adding up the expected regret over all  $F$  groups yields

$$\mathbb{E} [R_N(\bar{\pi})] \leq C \sum_{j=1}^F \left[ \sqrt{\bar{m}K \log(\bar{m}K)n\bar{B}_j} + n\bar{B}_j V_j^\beta \right].$$

**Proof of Corollary 3.1.1.** The result follows from Theorem 3.1 upon noting that  $\bar{B}_j = M^{-d}$  and  $V_j = \sqrt{d}M^{-1}$  for  $j = 1, \dots, M$  (and ignoring the constant  $\sqrt{d}$ ) with  $M$  as in the theorem completes the proof.

**Proof of Theorem 3.2.** Define  $\underline{S} = \mathcal{S}(\alpha, \beta, L, \mathcal{K}, d, \bar{c}, \bar{m})$  to be the subset of  $\mathcal{S}(\beta, L, \mathcal{K}, d, \bar{c}, \bar{m}) =: \bar{S}$  which also satisfies the margin condition. Then, for  $\bar{m} = 1, K = 2$ , all  $\alpha > 0$  and any policy  $\pi$

$$\sup_{\bar{S}} \mathbb{E} [R_N(\pi)] \geq \sup_{\underline{S}} \mathbb{E} [R_N(\pi)] \geq Cn^{1-\frac{\beta+\beta\alpha}{2\beta+d}} = Cn^{1-\frac{\beta}{2\beta+d}} \cdot n^{\frac{-\beta\alpha}{2\beta+d}}$$

for a constant  $C$  not depending on  $\alpha$  and where the second inequality follows from Theorem 4.1 in Rigollet and Zeevi (2010). Now notice that for every  $n$  there exists an  $\alpha$  sufficiently small such that  $n^{\frac{-\beta\alpha}{2\beta+d}} > \frac{1}{2}$ . Thus,  $\sup_{\bar{S}} \mathbb{E} [R_N(\pi)] \geq \frac{C}{2}n^{1-\frac{\beta}{2\beta+d}}$  as desired.

**Proof of Theorem 3.3.** The proof follows from only considering the contribution to regret coming from  $\mathbb{E}[\tilde{R}_j(\bar{\pi})]$  in (7.16) of the proof of Theorem 3.1.

**Proof of Theorem 3.4.** The proof is identical to the proof of Theorem 2.2 but with  $n = \mathbb{E}(N)$  replaced by  $\bar{c}n\bar{B}_j \geq \mathbb{E}(N_{B_j}(N))$ . Thus, the expected number of assignments is replaced by an upper bound on the expected number of individuals falling in group  $j$  and the result of Theorem 2.2 is applied on each group separately.

**Proof of Theorem 3.5.** The proof is similar to that found in Tsybakov (2004a) and Rigollet and Zeevi (2010). Fix  $\delta < \delta_0$ . Then, for any policy  $\pi$

$$R_N(\pi) \geq \delta \sum_{t=1}^N 1_{\{f^*(X_t) - f^{\pi_t}(X_t) > \delta\}}$$

$$\begin{aligned} &\geq \delta \left( S_N(\pi) - \sum_{t=1}^N \mathbf{1}_{\{0 < |f^*(X_t) - f^{\pi_t}(X_t)| \leq \delta\}} \right) \\ &\geq \delta \left( S_N(\pi) - \sum_{t=1}^N \mathbf{1}_{\{0 < |f^*(X_t) - f^\#(X_t)| \leq \delta\}} \right) \end{aligned}$$

Since  $S_n(\pi) \leq N$  there exists a  $c > 0$  not depending on  $N$  such that  $\left(\frac{S_n(\pi)}{cn}\right)^{\frac{1}{\alpha}} < \delta_0$ . Thus, we can set  $\delta = \left(\frac{S_n(\pi)}{cn}\right)^{\frac{1}{\alpha}}$  and use the margin condition upon integration on both sides of the above display to get (3.7). To obtain (3.8) insert (3.4) into (3.7).

**Proof of Theorem 3.6.** The proof is similar to the proof of Theorem 3.1 once we fix a value of the discrete covariates. Let  $c$  denote a positive constant which may change from line to line. By the construction of the treatment policy it follows that the regret can be written as  $R_N(\hat{\pi}) = \sum_{a \in A} \sum_{j=1}^{F_a} R_{a,j}(\hat{\pi})$ , where

$$R_{a,j}(\hat{\pi}) = \sum_{t=1}^N \left( f^*(X_t) - f^{(\hat{\pi}_{\{a\} \times B_{a,j}}, N_{a,j}(t))}(X_t) \right) \mathbf{1}_{(X_{t,D}=a, X_{t,C} \in B_{a,j})}.$$

For any bin  $B_{a,j}$  define

$$\bar{\mu}_{a,j}^{(i)} = \mathbb{E}(Y_t^{(i)} | X_{t,D} = a, X_{t,C} \in B_{a,j}) = \frac{1}{P_X(X_{t,D} = a, X_{t,C} \in B_{a,j})} \int_{a \times B_{a,j}} \mu^{(i)}(x) dP_X(x)$$

and

$$\begin{aligned} (\bar{\sigma}^2)_{a,j}^{(i)} &= \text{Var}(Y_t^{(i)} | X_{t,D} = a, X_{t,C} \in B_{a,j}) \\ &= \mathbb{E}(Y_t^{(i)2} | X_{t,D} = a, X_{t,C} \in B_{a,j}) - [\mathbb{E}(Y_t^{(i)} | X_{t,D} = a, X_{t,C} \in B_{a,j})]^2 \end{aligned}$$

Furthermore, let  $\bar{f}_{a,j}^{(i)} = f(\bar{\mu}_{a,j}^{(i)}, (\bar{\sigma}^2)_{a,j}^{(i)})$  with  $f_{a,j}^* = \max_{1 \leq i \leq K+1} \bar{f}_{a,j}^{(i)}$ . By exactly the same arguments as in the proof of Theorem 3.1 we now get that for  $x \in \{a\} \times B_{a,j}$ ,

$$f^*(x) \leq f_{a,j}^* + cV_{a,j}^\beta.$$

as well as,

$$f^{(\hat{\pi}_{\{a\} \times B_{a,j}}, N_{a,j}(t))}(x) \geq \bar{f}_{a,j}^{(\hat{\pi}_{\{a\} \times B_{a,j}}, N_{a,j}(t))} - cV_{a,j}^\beta.$$

Next we define  $\tilde{R}_{a,j}(\tilde{\pi}) = \sum_{t=1}^{N_{a,j}} \left( f_{a,j}^* - \bar{f}_{a,j}^{(\tilde{\pi}_{\{a\} \times B_{a,j}}, t)} \right)$ . This corresponds to the regret associated with a treatment problem without covariates where treatment  $i$  yields reward  $\bar{f}_{a,j}^{(i)}$ , and the best treatment yields  $f_{a,j}^* = \max_i \bar{f}_{a,j}^{(i)} \leq \bar{f}_{a,j}^*$ . Therefore, we can write

$$R_{a,j}(\hat{\pi}) = \sum_{t=1}^N \left( f^*(X_t) - f^{(\hat{\pi}_{\{a\} \times B_{a,j}}, N_{a,j}(t))}(X_t) \right) \mathbf{1}_{(X_{t,D}=a, X_{t,C} \in B_{a,j})}$$

$$\begin{aligned}
&\leq \sum_{t=1}^N \left( f_{a,j}^* - \bar{f}_{a,j}^{\hat{\pi}(\{a\} \times B_{a,j}), N_{a,j}(t)} + 2cV_{a,j}^\beta \right) \mathbf{1}_{(X_{t,D}=a, X_{t,C} \in B_{a,j})} \\
&= \tilde{R}_{a,j}(\hat{\pi}) + 2cV_{a,j}^\beta N_{a,j}(N),
\end{aligned}$$

where  $N_{a,j}(N)$  is the number of observations for which  $x \in a \times B_{a,j}$  given that there are  $N$  observations in total. Taking expectations, and using that  $N$  is independent of all other random variables implies  $\mathbb{E} [N_{a,j}(N)] \leq n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})$  gives

$$\mathbb{E} [R_{a,j}(\hat{\pi})] \leq \mathbb{E} [\tilde{R}_{a,j}(\hat{\pi})] + n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})cV_{a,j}^\beta.$$

Since  $\mathbb{E} [\tilde{R}_{a,j}(\hat{\pi})]$  is the expected regret of a treatment problem without covariates we can apply Theorem 7.1 with the following values

$$\Delta = \sqrt{\frac{\bar{m}K \log(\bar{m}K)}{n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})}}, \quad \gamma = \mathcal{K}L, \quad T = n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j}),$$

for each  $a \in A$  and  $B_{a,j}$ ,  $j = 1, \dots, F_a$  to obtain the following bound on the regret accumulated across any group:

$$\mathbb{E} [R_{a,j}(\hat{\pi})] \leq c \left[ \sqrt{\bar{m}K \log(\bar{m}K)n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})} + n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})_j V_j^\beta \right].$$

Thus, adding up the expected regret over all groups yields

$$\mathbb{E} [R_N(\hat{\pi})] \leq c \sum_{a \in A} \sum_{j=1}^{F_a} \left[ \sqrt{\bar{m}K \log(\bar{m}K)n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})} + n\mathbb{P}(X_{t,D} = a, X_{t,C} \in B_{a,j})_j V_{a,j}^\beta \right].$$

### 7.3 Proof of Theorems in Section 4

**Proof of Theorem 4.1.** Define  $\epsilon_s = u(s, T) = 16\sqrt{\frac{2\bar{a}^2}{s} \log\left(\frac{T}{s}\right)}$ . In the following we will distinguish between two types of batches, namely batches of individuals that have to be assigned a treatment, and batches of information on the outcome of previously assigned treatments. The latter type of batches will be the key object of interest when determining whether or not to eliminate a given treatment, whereas the former will be relevant when counting the total regret from running the treatment policy. In this proof we let  $\underline{B}(s)$  denote the minimal number of observed outcomes per treatment based on  $s$  batches of information. Consider a batch  $b$  of information. Recall that if the optimal treatment as well as some treatment  $i$  have not been eliminated, then the optimal treatment will eliminate treatment  $i$  if  $\hat{\Delta}_i(\underline{B}(b)) \geq \gamma\epsilon_{\underline{B}(b)}$ , and treatment  $i$  will eliminate the optimal treatment if  $\hat{\Delta}_i(\underline{B}(b)) \leq -\gamma\epsilon_{\underline{B}(b)}$ .

To be able to say something about when either of these two events occurs we introduce the (unknown) quantity,  $\tau_i^*$ , which is defined through the relation

$$\Delta_i = 24\gamma \sqrt{\frac{2\bar{a}^2}{\tau_i^*} \log\left(\frac{T}{\tau_i^*}\right)}, \quad i = 1, \dots, K.$$

Since  $\tau_i^*$  in general will not be an integer, we also define  $\tau_i = \lceil \tau_i^* \rceil$ . Next introduce the hypothetical batch (of information)  $b_i = \min\{l : \underline{B}(l) \geq \tau_i^*\}$ . It is the first batch of information after which we have more than  $\tau_i^*$  observations of the outcome of treatment  $i$ . Notice that

$$\tau_i^* \leq \underline{B}(b_i) \leq C \left( \frac{\bar{a}^2 \gamma^2}{\Delta_i^2} \log \left( \frac{T \Delta_i^2}{1152 \bar{a}^2 \gamma^2} \right) + \bar{m} \right), \quad (7.17)$$

$$\tau_i \leq \underline{B}(b_i), \quad (7.18)$$

$$\underline{B}(b_i) \leq \tau_i + \bar{m}, \quad (7.19)$$

Notice that  $1 \leq \tau_1 \leq \dots \leq \tau_K$  and  $1 \leq b_1 \leq \dots \leq b_K$ . Define the following events:

$\mathcal{A}_i = \{\text{The optimal treatment has not been eliminated after batch } b_i \text{ has been observed}\},$

$\mathcal{B}_i = \{\text{Every treatment } j \in \{1, \dots, i\} \text{ has been eliminated after batch } b_j \text{ has been observed}\}.$

Furthermore, let  $\mathcal{C}_i = \mathcal{A}_i \cap \mathcal{B}_i$ , and observe that  $\mathcal{C}_1 \supseteq \dots \supseteq \mathcal{C}_K$ . For any  $i = 1, \dots, K$ , the contribution to regret incurred after batch  $b_i$  of information is at most  $\Delta_{i+1}N$  on  $\mathcal{C}_i$ . In what follows we fix a treatment,  $K_0$ , which we will be specific about later. Using this and letting  $m$  denote the expected number of observations in a batch we get the following decomposition of expected regret:

$$\begin{aligned} \mathbb{E} [R_N(\hat{\pi})] &= \mathbb{E} \left[ R_N(\hat{\pi}) \left( \sum_{i=1}^{K_0} 1_{\mathcal{C}_{i-1} \setminus \mathcal{C}_i} + 1_{\mathcal{C}_{K_0}} \right) \right] \\ &\leq n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \setminus \mathcal{C}_i) + \sum_{i=1}^{K_0} B_i(b_i) \Delta_i + n \Delta_{K_0+1} + Dm, \end{aligned} \quad (7.20)$$

where the last term is due to the fact that the delay means that the all treatment allocations during the first  $D+1$  batches have to be made without any information about the treatment outcomes. For every  $i = 1, \dots, K$  the event  $\mathcal{C}_{i-1} \setminus \mathcal{C}_i$  can be decomposed as follows

$$\mathcal{C}_{i-1} \setminus \mathcal{C}_i = (\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) \cup (\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}).$$

Therefore, the first term on the right-hand side of (7.20) can be written as

$$n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \setminus \mathcal{C}_i) = n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) + n \sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}). \quad (7.21)$$

Notice that the first term on the right-hand side will be zero if  $b_{i-1} = b_i$ . On the event  $\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}$  the optimal treatment has not eliminated treatment  $i$  at batch  $b_i$ . Therefore, for the last term on the right hand side of equation (7.21) we find that

$$\mathbb{P}(\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}) \leq \mathbb{P}(\hat{\Delta}_i(\underline{B}(b_i)) \leq \gamma \epsilon_{\underline{B}(b_i)})$$

$$\begin{aligned}
&\leq \mathbb{P} \left( \hat{\Delta}_i(\underline{B}(b_i)) - \Delta_i \leq \gamma \epsilon_{\tau_i} - \Delta_i \right) \\
&\leq \mathbb{E} \left[ \mathbb{P} \left( |\hat{\Delta}_i(\underline{B}(b_i)) - \Delta_i| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \mid \underline{B}(b_i) \right) \right]. \tag{7.22}
\end{aligned}$$

For any  $s \geq \tau_i$  we have that

$$\begin{aligned}
&\mathbb{P} \left( |\hat{\Delta}_i(s) - \Delta_i| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\hat{\mu}_s^{(i)}, (\hat{\sigma}_s^2)^{(i)}) + f(\mu^{(i)}, (\sigma^2)^{(i)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{2} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |f(\hat{\mu}_s^{(i)}, (\hat{\sigma}_s^2)^{(i)}) - f(\mu^{(i)}, (\sigma^2)^{(i)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right). \tag{7.23}
\end{aligned}$$

Furthermore, for any  $j \in \{i, *\}$ , we have

$$\begin{aligned}
&\mathbb{P} \left( |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{4} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| + |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{4\mathcal{K}} \gamma \epsilon_{\tau_i} \right) \\
&\leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right). \tag{7.24}
\end{aligned}$$

By the mean value theorem we have that

$$\mathbb{P} \left( |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{8\mathcal{K}} \gamma \epsilon_{\tau_i} \right) \leq \mathbb{P} \left( |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{32\mathcal{K}} \gamma \epsilon_{\tau_i} \right) + \mathbb{P} \left( |(\hat{\mu}_{2,s})^{(j)} - \mu_2^{(j)}| \geq \frac{1}{16\mathcal{K}} \gamma \epsilon_{\tau_i} \right), \tag{7.25}$$

where  $\mu_2 = \mathbb{E} [Y_1^2]$  and  $\hat{\mu}_{2,s} = \frac{1}{s} \sum_{i=1}^s Y_i^2$ . Combining equations (7.22), (7.23), (7.24) and (7.25) and using Hoeffding's inequality to each of the three terms as well as the fact that  $\gamma \geq \mathcal{K}$  we arrive at the following bound,

$$\begin{aligned}
\mathbb{P} \left( |\hat{\Delta}_i(s) - \Delta_i| \geq \gamma \epsilon_{\tau_i} \right) &\leq C \exp \left( -\frac{1}{1024\bar{a}} \epsilon_{\tau_i}^2 s \right) \\
&\leq C \exp \left( -\frac{1}{1024\bar{a}} \epsilon_{\tau_i}^2 \tau_i \right) \\
&= C \exp \left( -\overline{\log} \left( \frac{T}{\tau_i} \right) \right) \\
&\leq C \frac{\tau_i}{T}.
\end{aligned}$$

Thus,

$$\mathbb{P} (\mathcal{B}_i^c \cap \mathcal{A}_i \cap \mathcal{B}_{i-1}) \leq C \frac{\tau_i}{T}$$

On the event  $\mathcal{C}_{i-1} \cap \mathcal{A}_i^c$  the optimal treatment is eliminated between the time batch  $b_{i-1} + 1$  and  $b_i$  of information arrives. Furthermore, every suboptimal treatment  $j \leq i - 1$  has also been eliminated. Therefore, the probability of this event can be bounded as follows:

$$\begin{aligned} \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) &\leq \mathbb{P}\left(\exists(j, s), i \leq j \leq K, b_{i-1} + 1 \leq s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &\leq \sum_{j=i}^K \mathbb{P}\left(\exists s, b_{i-1} + 1 \leq s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &= \sum_{j=i}^K [\Phi_j(b_i) - \Phi_j(b_{i-1})], \end{aligned}$$

where  $\Phi_j(b) = \mathbb{P}\left(\exists s \leq b; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right)$ . We proceed to bounding terms of the form  $\Phi_j(b_i)$  for  $j \geq i$ .

$$\begin{aligned} \mathbb{P}\left(\exists s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) \leq -\gamma\epsilon_{\underline{B}(s)}\right) &\leq \mathbb{P}\left(\exists s \leq b_i; \hat{\Delta}_j(\underline{B}(s)) - \Delta_j \leq -\gamma\epsilon_{\underline{B}(s)}\right) \\ &\leq \mathbb{P}\left(\exists s \leq B_j(b_i); \hat{\Delta}_j(s) - \Delta_j \leq -\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; \hat{\Delta}_j(s) - \Delta_j \leq -\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{2}\gamma\epsilon_s\right) \\ &\quad + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(*)}, (\hat{\sigma}_s^2)^{(*)}) - f(\mu^{(*)}, (\sigma^2)^{(*)})| \geq \frac{1}{2}\gamma\epsilon_s\right). \end{aligned}$$

For any  $j \in \{i, \dots, K, *\}$  we find that

$$\begin{aligned} &\mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |f(\hat{\mu}_s^{(j)}, (\hat{\sigma}_s^2)^{(j)}) - f(\mu^{(j)}, (\sigma^2)^{(j)})| \geq \frac{1}{2}\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |(\hat{\sigma}_s^2)^{(j)} - (\sigma^2)^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) \\ &\leq \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{4\mathcal{K}}\gamma\epsilon_s\right) + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |(\hat{\mu}_{2,s}^{(j)}) - \mu_2^{(j)}| \geq \frac{1}{8\mathcal{K}}\gamma\epsilon_s\right) \\ &\quad + \mathbb{P}\left(\exists s \leq \tau_i + \bar{m}; |\hat{\mu}_s^{(j)} - \mu^{(j)}| \geq \frac{1}{16\mathcal{K}}\gamma\epsilon_s\right) \\ &\leq C \frac{\tau_i + \bar{m}}{T}, \end{aligned}$$

where we once more have used equation (7.19) and Lemma A.1 in Perchet and Rigollet (2013).

Using this we find that

$$\sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) \leq \sum_{i=1}^{K_0} \Delta_i \sum_{j=i}^K [\Phi_j(b_i) - \Phi_j(b_{i-1})]$$



$$\begin{aligned}
&\leq \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \Phi_j(b_i) (\Delta_i - \Delta_{i+1}) + \sum_{j=K_0}^K \Delta_{K_0} \Phi_j(b_{K_0}) + \sum_{j=1}^{K_0-1} \Delta_j \Phi_j(b_j) \\
&\leq C \left( \frac{1}{T} \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} (\tau_i + \bar{m}) (\Delta_i - \Delta_{i+1}) + \frac{1}{T} \sum_{j=1}^K \Delta_{j \wedge K_0} (\tau_{j \wedge K_0} + \bar{m}) \right).
\end{aligned} \tag{7.26}$$

Observe that, by (7.18),

$$\begin{aligned}
\sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \tau_i (\Delta_i - \Delta_{i+1}) &\leq C \gamma^2 \bar{a}^2 \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} \frac{(\Delta_i - \Delta_{i+1})}{\Delta_i^2} \overline{\log} \left( \frac{T \Delta_i^2}{1152 \bar{a}^2 \gamma^2} \right) \\
&\leq C \bar{a}^2 \gamma^2 \sum_{j=1}^K \int_{\Delta_{j \wedge K_0}}^{\Delta_1} \frac{1}{x^2} \overline{\log} \left( \frac{T x^2}{1152 \bar{a}^2 \gamma^2} \right) dx \\
&\leq C \bar{a}^2 \gamma^2 \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{1152 \bar{a}^2 \gamma^2} \right).
\end{aligned} \tag{7.27}$$

The parts involving  $\bar{m}$  in equation (7.26) can be bounded by

$$\bar{m} \sum_{j=1}^K \sum_{i=1}^{j \wedge K_0 - 1} (\Delta_i - \Delta_{i+1}) + \sum_{j=1}^K \Delta_{j \wedge K_0} \bar{m} \leq \bar{m} K. \tag{7.28}$$

Bringing together equations (7.26), (7.27) and (7.28) we see that

$$\sum_{i=1}^{K_0} \Delta_i \mathbb{P}(\mathcal{C}_{i-1} \cap \mathcal{A}_i^c) \leq C \left( \frac{\bar{a}^2 \gamma^2}{T} \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{1152 \bar{a}^2 \gamma^2} \right) + \frac{K \bar{m}}{T} \right). \tag{7.29}$$

Combining this with equation (7.21) and (7.20) we obtain

$$\begin{aligned}
\mathbb{E} [R_N(\hat{\pi})] &\leq C \left( \frac{\bar{a}^2 \gamma^2 n}{T} \sum_{j=1}^K \frac{1}{\Delta_{j \wedge K_0}} \overline{\log} \left( \frac{T \Delta_{j \wedge K_0}^2}{1152 \bar{a}^2 \gamma^2} \right) + \frac{n \bar{a}^2 \gamma^2}{T} \sum_{j=1}^{K_0} \frac{1}{\Delta_j} \overline{\log} \left( \frac{T \Delta_j^2}{1152 \bar{a}^2 \gamma^2} \right) \right. \\
&\quad \left. + \sum_{i=1}^{K_0} B_i(b_i) \Delta_i + n \Delta_{K_0+1} + \frac{n \bar{m} K}{T} + m D \right) \\
&\leq C \left( \left( 1 + \frac{n}{T} \right) \bar{a}^2 \gamma^2 \sum_{j=1}^{K_0} \frac{1}{\Delta_j} \overline{\log} \left( \frac{T \Delta_j^2}{1152 \bar{a}^2 \gamma^2} \right) + \frac{\bar{a}^2 \gamma^2 n}{T} \frac{K - K_0}{\Delta_{K_0}} \overline{\log} \left( \frac{T \Delta_{K_0}^2}{1152 \bar{a}^2 \gamma^2} \right) \right. \\
&\quad \left. + n \Delta_{K_0+1} + \frac{n \bar{m} K}{T} + \bar{m} D \right).
\end{aligned} \tag{7.30}$$

Fix  $\Delta > 0$  and let  $K_0$  be such that  $\Delta_{K_0+1} = \Delta^-$ . Define the function  $\phi(\cdot)$  by

$$\phi(x) = \frac{1}{x} \overline{\log} \left( \frac{T x^2}{1152 \bar{a}^2 \gamma^2} \right),$$

and notice that  $\phi(x) \leq 2e^{-1/2}\phi(x')$  for any  $x \geq x' \geq 0$ . Using this with  $x' = \Delta$  and  $x = \Delta_i$  for  $i \leq K_0$  we obtain the following bound on the expected regret.

$$\mathbb{E} [R_N(\hat{\pi})] \leq C \left( \frac{\bar{a}^2 \gamma^2 K}{\Delta} \left( 1 + \frac{n}{T} \right) \overline{\log} \left( \frac{T \Delta^2}{1152 \bar{a}^2 \gamma^2} \right) + n \Delta^- + \frac{n \bar{m} K}{T} + m D \right). \quad (7.31)$$

Note that we by definition we have that  $m \leq \bar{m}$ . The theorem then follows by arguments similar to those in the proof of theorem 2.1. ■

**Proof of Theorem 4.2.** Recall equation (7.16). Applying Theorem 4.1 with the following values

$$\Delta = \sqrt{\frac{\bar{m} K \bar{a} \log(\bar{m} K / \bar{a})}{n \bar{B}_j}}, \quad \gamma = \mathcal{K} L, \quad T = n \bar{B}_j,$$

for each bin  $j = 1, \dots, F$ , we obtain the following bound on the regret accumulated across the any bin  $j$ :

$$\mathbb{E} [R_j(\bar{\pi})] \leq c \left[ \sqrt{\bar{m} \bar{a}^3 K \log(\bar{m} K / \bar{a}) n \bar{B}_j} + n \bar{B}_j V_j^\beta + \bar{m} K + m_j D \right],$$

where  $m_j$  is the expected batch size associated with bin  $j$ . Note that  $m_j \leq \bar{m} \bar{B}_j$ . Thus,

$$\mathbb{E} [R_N(\bar{\pi})] \leq c \left( \sum_{j=1}^F \left[ \sqrt{\bar{m} \bar{a}^3 K \log(\bar{m} K / \bar{a}) n \bar{B}_j} + n \bar{B}_j V_j^\beta + \bar{m} K \right] + \bar{m} D \right).$$

## References

- Susan Athey and Stefan Wager. Efficient policy learning. *arXiv preprint arXiv:1702.02896*, 2017.
- Peter Auer and Nicolò Cesa-Bianchi. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 2002.
- Debopam Bhattacharya and Pascaline Dupas. Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 167(1):168–196, 2012.
- Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*, 2012.

- Gary Chamberlain. Econometrics and decision theory. *Journal of Econometrics*, 95(2):255–283, 2000.
- Rajeev H Dehejia. Program evaluation as a decision problem. *Journal of Econometrics*, 125(1): 141–173, 2005.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for reinforcement learning. In *ICML*, pages 162–169, 2003.
- Campbell R Harvey and Akhtar Siddique. Conditional skewness in asset pricing tests. *The Journal of Finance*, 55(3):1263–1295, 2000.
- Keisuke Hirano and Jack R Porter. Asymptotics for statistical treatment rules. *Econometrica*, 77 (5):1683–1701, 2009.
- Maximilian Kasy. Using data to inform policy. Technical report, 2014.
- Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005.
- Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Cemmap working paper*, 2015.
- Nicolas Klein and Sven Rady. Negatively correlated bandits. *The Review of Economic Studies*, page rdq025, 2011.
- Philip W Lavori, Ree Dawson, and A John Rush. Flexible treatment strategies in chronic disease: clinical and research implications. *Biological psychiatry*, 48(6):605–614, 2000.
- Enno Mammen, Alexandre B Tsybakov, et al. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- Charles F. Manski. Statistical treatment rules for heterogenous populations. *Econometrica*, 2004.
- Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
- Susan A Murphy. An experimental design for the development of adaptive treatment strategies. *Statistics in medicine*, 24(10):1455–1481, 2005.
- Susan A Murphy, Mark J van der Laan, and James M Robins. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.

- V. Perchet and P. Rigollet. The multi-armed bandit problem with covariates. *The Annals of Statistics*, 2013.
- V. Perchet, P. Rigollet, Sylvain Chassang, and Erik Snowberg. Batched bandit problems. *The Annals of Statistics*, 44(2):660–681, 2016.
- Philippe Rigollet and Assaf Zeevi. Nonparametric bandits with covariates. *arXiv preprint arXiv:1003.1630*, 2010.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58, 1952.
- James M Robins. Causal inference from complex longitudinal data. In *Latent variable modeling and applications to causality*, pages 69–117. Springer, 1997.
- Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- J. Stoye. Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151: 70–81, 2009.
- Jörg Stoye. Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics*, 166(1):138–156, 2012.
- Aleksey Tetenov. Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics*, 166(1):157–165, 2012.
- A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, 2004a.
- Alexandre B Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, pages 135–166, 2004b.
- Michael Woodroofe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.
- Yuhong Yang, Dan Zhu, et al. Randomized allocation with nonparametric estimation for a multi-armed bandit problem with covariates. *The Annals of Statistics*, 30(1):100–121, 2002.