



Bandit algorithms for policy learning: methods, implementation, and welfare-performance

Toru Kitagawa¹ · Jeff Rowley²

Received: 15 June 2024 / Revised: 31 August 2024 / Accepted: 31 August 2024
© Japanese Economic Association 2024

Abstract

Static supervised learning—in which experimental data serves as a training sample for the estimation of an optimal treatment assignment policy—is a commonly assumed framework of policy learning. An arguably more realistic but challenging scenario is a dynamic setting in which the planner performs experimentation and exploitation simultaneously with subjects that arrive sequentially. This paper studies bandit algorithms for learning an optimal individualised treatment assignment policy. Specifically, we study applicability of the EXP4.P (Exponential weighting for Exploration and Exploitation with Experts) algorithm developed by Beygelzimer et al. (Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings, pp 19–26, 2011) to policy learning. Assuming that the class of policies has a finite Vapnik–Chervonenkis dimension and that the number of subjects to be allocated is known, we present a high probability welfare-regret bound of the algorithm. To implement the algorithm, we use an incremental enumeration algorithm for hyperplane arrangements. We perform extensive numerical analysis to assess the algorithm’s sensitivity to its tuning parameters and its welfare-regret performance. Further simulation exercises are calibrated to the National Job Training Partnership Act (JTPA) Study sample to determine how the algorithm performs when applied to economic data. Our findings highlight various computational challenges and suggest that the limited welfare gain from the algorithm is due to substantial heterogeneity in causal effects in the JTPA data.

Keywords Empirical Welfare Maximisation · Hyperplane arrangements · Reinforcement learning · Treatment choice

✉ Toru Kitagawa
toru_kitagawa@brown.edu

¹ Department of Economics, Brown University, Providence, USA

² Department of Economics, University College London, London, UK

1 Introduction

How to make use of evidence for policymaking is a topic of great importance. The growing literature on statistical treatment choice (Dehejia, 2005; Manski, 2004) and on policy learning (Athey & Wager, 2021, Kitagawa & Tetenov, 2018; 2021, Kitagawa et al., 2023, Mbakop & Tabord-Meehan, 2021, etc.) develops formal frameworks and methods for combining causal inference with the social planner's decision, making use of statistical decision theory and machine learning methods. A commonly assumed setting is that of static supervised learning, wherein experimental data serves as training data and learning an optimal treatment assignment policy happens only once. The attraction of this setting is its simplicity; it ignores, however, the important dynamic aspects of learning and of implementation of treatment assignment policies. Subjects that are assigned to treatments and that contribute to causal evidence often appear sequentially over time. Accordingly, operations to accumulate evidence, learn causal effects, and assign treatments can run simultaneously over multiple time periods. In such a dynamic setting: Which strategy should the social planner implement to maximise a welfare criterion? Can we expect this strategy to deliver substantial welfare gains in public policy applications? How does this strategy compare with the static supervised learning approach, in terms of the assignment policy that it implements and the welfare that it generates? And, is it feasible to compute and implement a sophisticated dynamic sampling and assignment strategy in practice?

This paper aims to answer these questions by formulating dynamic policy learning as a multi-arm bandit problem with contextual information. Multi-arm bandit problems capture dynamic environments in which a new subject arrives in every period, with the decision-maker assigning the subject to one of the many candidate treatments that are available. The decision-maker learns the effects of these treatments from the subjects' responses, and updates the rules by which subjects are assigned to treatments sequentially. The goal of the analysis is to find rules that maximise the sum of the subjects' treatment outcomes by effectively balancing exploration and exploitation, for which there is a trade-off. In the bandit problem setting, contextual information refers to the observable pre-treatment characteristics of the sequentially arriving subjects, upon which the decision-maker can discriminate (i.e., by ascribing different treatment rules to subjects with different pre-treatment characteristics). Individualised assignment based upon contextual information outperforms non-individualised assignment if treatment effects are heterogeneous across the pre-treatment covariates.

Our approach is to associate an arm in a bandit problem with an individualised treatment assignment rule that maps a subject's pre-treatment covariates to an assigned treatment. We then propose a bandit algorithm to learn an optimal individualised treatment assignment policy from a class of policies. We assume that the class of individualised treatment assignment policies is finite or is infinite with finite Vapnik–Chervonenkis (VC) dimension, as is considered in Kitagawa and Tetenov (2018). Hence, we deal with the challenging situation of a bandit problem in which there is an infinite number of arms—something that arises if

a subject's pre-treatment covariates include a continuous variable. In contrast to contextual linear bandit problems—like those studied in Auer (2002) and Abbasi-Yadkori et al. (2011), to list but a few relevant papers—our approach does not impose functional form restrictions on the conditional average treatment effect.

Our approach builds upon the influential work of Auer et al. (2002) that proposes the EXP4 (Exponential weighting for Exploration and Exploitation with Experts) algorithm for an adversarial bandit. Extending the original EXP4 algorithm, Beygelzimer et al. (2011) proposes the EXP4.P algorithm for learning classification rules with a finite VC dimension—which we denote by $D < \infty$. Both algorithms consider a set of experts and probability distributions over them. Each expert gives a recommendation (i.e., which bandit arm to pull given contextual information). The decision-maker makes their decision by aggregating the experts' recommendations according to a probability distribution. Every time a decision is made, the decision-maker receives feedback on each expert's performance, updating the probability distribution over the experts that they use based upon this feedback. Viewing an assignment rule as an expert, Beygelzimer et al. (2011) shows that iterative implementation over a number of periods—which we denote by $T < \infty$ —attains a high-probability upper bound on the distribution of cumulative regret of order $O(\sqrt{T \cdot D \cdot \ln(T)})$, implying that the average welfare-regret converges at a rate of $O(\sqrt{D/T \cdot \ln(T)})$.

We modify the EXP4.P algorithm to develop a version of the algorithm that is tailored to policy learning. Specifically, we associate the aforementioned set of experts with a class of individualised treatment assignment rules that is finite or of finite VC dimension. Our modified algorithm accommodates treatment choice given an additive welfare criterion and a general bounded outcome variable, and we show that, for suitable choices of tuning parameters, it achieves a high-probability upper bound on the cumulative welfare-regret of order $O(\sqrt{T \cdot D \cdot \ln(T)})$.

An important step in implementing the EXP4.P algorithms is to coarsen the class of assignment rules to a finite class using information gathered from the first τ observations. This step is computationally non-trivial and, to our knowledge, development of computationally efficient methods for this step is an open question. Focusing on the class of Linear Eligibility Score (LES) rules, we show that the coarsening problem is equivalent to a hyperplane arrangement problem in the geometry literature. Various cell enumeration procedures have been proposed in that literature, and we suggest using the Incremental Enumeration algorithm introduced by Rada and Černý (2018) (and improved upon by Gu & Koenker, 2022) as one way to conduct this step.

To illustrate implementation of our modified EXP4.P algorithm and to assess its welfare-performance, we perform extensive simulation studies. In one simulation design, we modify the variances of the potential outcomes (i.e., the heterogeneity of individual causal effects) whilst holding fixed the magnitude of conditional average treatment effects. We find that the (median and variance of the) welfare-performance of the EXP4.P algorithm is sensitive to the heterogeneity of individual causal effects, indicating a challenge in those policy applications where subjects' treatment responses are highly heterogeneous in unobservables.

We also assess the sensitivity of this welfare-performance to the tuning parameters that govern the exploration and exploitation trade-off of the algorithm.

Through a novel simulation design that we calibrate to actual economic data, we also investigate how our modified EXP4.P algorithm performs in public policy applications. We use data from the National Job Training Partnership Act (JTPA) Study to do this. The original JTPA sample contains the time stamp of when each experimental subject entered the Study. Maintaining the order of individual entry, we perform a counterfactual analysis of the welfare-level that the social planner would attain if assignment were based upon the EXP4.P algorithm. We estimate potential outcome regressions via random forest methods using the JTPA sample, and construct distributions of potential outcomes from these and their residuals. By applying our modified EXP4.P algorithm to a multitude of samples drawn from these distributions, we obtain a distribution over the average welfare that the algorithm attains.

Recent research in policy learning extends to and intersects with the machine learning literature on bandit algorithms (see Lattimore & Szepesvári, 2020, for a monograph on bandit algorithms). The welfare-performance criterion that we consider concerns the cumulative welfare for sequentially arriving units rather than for the super-population. The latter corresponds to the welfare objective in the best-arm identification problems as studied in Ariu et al. (2021), Kasy and Sautmann (2021), Russo and Roy (2016), among others. Athey et al. (2022) and Qin and Russo (2024) assess the trade-off between the in-sample welfare and the super-population welfare of best-arm identification and study how to balance them out. Recent advances in bandit algorithms in the econometrics literature include those studied in Adusumilli (2021), Dimakopoulou et al. (2017), Kock et al. (2022), Kuang and Wager (2023), to list but a few relevant papers. Application and feasible implementation of EXP4.P algorithms to policy learning have, to our knowledge, not been studied in the policy learning literature.

Preceding the EXP4.P algorithm, the literature of online learning has studied exponential weighting of experts in prediction and portfolio choice problems; see Cesa-Bianchi and Lugosi (2006), Littlestone and Warmuth (1994), and Vovk (1990) for a monograph and early works of the topic. Chen (2023) and Viviano and Bradic (2023) apply the perspective and methods of online aggregation of experts to causal inference with panel data.

This paper's dynamic approach to policy learning is different from the approach considered in the literature on dynamic treatment regimes (Han, 2023; Ida et al., 2024; Ko et al., 2022; Murphy, 2003; Robins, 1986; Zhang et al., 2018). That literature considers estimation of adaptive allocation of treatments to the same individual over multiple time periods using exogenous training data. For a similar reason, this paper differs from the time-series Empirical Welfare Maximisation (EWM)-approach proposed in Kitagawa et al. (2024). Exponential weighting over classification rules or individualised treatment rules based upon their empirical performance appears, however, in Probably Approximately Correct (PAC) Bayes analysis present in the supervised learning literature (see Bégin et al., 2014, Kitagawa et al., 2022, McAllester, 2003, and references therein).

2 Framework

A utilitarian social planner is faced with a sequential allocation problem of K treatment arms over T periods—both of which are finite and are known to the social planner. The allocation problem involves the repeated interaction of the social planner with a stable and passive environment (Nature and subjects). The social planner observes subjects' responses to realised treatment arms and bases her choice of treatment in subsequent periods, in part, upon these responses.

2.1 Timing and the flow of information

A subject—here, subject t —is characterised by a collection of J features together with a collection of K counterfactual responses (i.e., by the subject's response to each treatment arm). We denote the collection of features by $\mathbf{x}_t \in \mathcal{X} \subset \mathbb{R}^J$ and refer to these as covariates, with $x_{t,j} \in \mathbb{R}$ denoting the j th covariate. We denote the collection of counterfactual responses by $\mathbf{y}_t = (y_t(1), \dots, y_t(K))^T \in \mathbb{R}^K$ and refer to these as potential outcomes, with $y_t(k) \in \mathbb{R}$ denoting the potential outcome corresponding to the k th treatment arm. We denote the population by P , which constitutes a joint distribution over covariates and potential outcomes that we concatenate as $(\mathbf{y}_t^T, \mathbf{x}_t^T)$. We maintain the following assumptions on the population throughout.

Assumption 1 (*Stationary population*) A sequence of random variables $(\mathbf{y}_t^T, \mathbf{x}_t^T)$ is independently and identically distributed to P .

Remark 1 Assumption 1 is standard in stochastic bandit problems. We impose this assumption in our numerical analysis and so make this assumption here too. A performance guarantee for welfare-regret under a nonstationary or adversarial environment is, however, known for EXP-type algorithms (see Lattimore & Szepesvári, 2020).

Assumption 2 (*Bounded outcomes*) \mathbf{y}_t is such that $0 \leq y_t(k) \leq M$ for all $k = 1, \dots, K$ and $t = 1, \dots, T$.

Remark 2 Assumption 2 requires that outcome data is both bounded and non-negative. Economic data can always be bounded, but may not be guaranteed to be non-negative. Via the addition of a positive constant, economic data can be made non-negative to satisfy Assumption 2. As is discussed in Kitagawa and Tetenov (2018) in the context of static policy learning, an estimated optimal policy is not generally invariant to the addition of constants to outcomes. See Remark 3 below for further discussion.

At the beginning of each period, a single subject is randomly selected from the population. We adopt the convention of labelling subjects according to the order in which they are selected (i.e., Nature selects subject t in period t). Once a subject

is selected, their covariates are revealed to the social planner (i.e., the social planner observes \mathbf{x}_t). The social planner then—directly or as the outcome of some randomisation—administers a treatment, which we denote by $k_t \in \{1, \dots, K\}$, and we assume that the subject fully complies with the assigned treatment. The social planner then observes the realised outcome (i.e., the social planner observes $y_t(k_t)$) without delay. On the other hand, the social planner does not observe the potential outcomes corresponding to unrealised treatment arms. At the end of each period, the social planner retains their accumulated knowledge up to that point (i.e., the social planner adds $(\mathbf{x}_t, k_t, y_t(k_t))$ to the information that she possesses at the beginning of period t) and carries this through to the next period. Nature continues to select subjects until the population is exhausted—we reiterate that T , in addition to K , is finite and is known.

2.2 The social planner's actions

We denote the information set of the social planner at the beginning of period t by \mathcal{I}_t . The information set of the social planner is empty at the beginning of the first period, and otherwise evolves according to

$$\mathcal{I}_{t+1} = (\mathcal{I}_t, k_t, \mathbf{x}_t^\top, y_t(k_t)), \quad (2.1)$$

as per Sect. 2.1. We emphasise that a subject's covariates are revealed to the social planner in advance of her choice (i.e., the information set of the social planner in period t at the time that she makes her choice comprises \mathcal{I}_t and \mathbf{x}_t).

Following the terminology of the online learning literature (e.g., Cesa-Bianchi & Lugosi, 2006), we define an *expert* as $f : \mathcal{X} \rightarrow \Delta^K$ —a mapping from covariates to a probability distribution over the set of treatments. Each expert specifies a time-invariant individualised assignment rule (as a function of covariates) that allows for randomised allocation. That is, $f(\mathbf{x}_t)$ constitutes a K -vector whose k th entry—which we denote by $f_k(\mathbf{x}_t)$ —specifies the probability that subject t is assigned to treatment k . For instance, when there are only two treatment arms, the single index policy $f(\mathbf{x}_t) = (1((1, \mathbf{x}_t^\top)\beta < 0), 1((1, \mathbf{x}_t^\top)\beta \geq 0))^\top$ that is often considered in the context of static treatment choice (e.g., Kitagawa & Tetenov, 2018) corresponds to an expert that allocates one of two treatment arms deterministically according to a fixed threshold rule for the sign of $(1, \mathbf{x}_t^\top)\beta$. We denote the set of experts by \mathfrak{F} .

The social planner's goal is to efficiently learn a best expert (i.e., individualised treatment assignment) in \mathfrak{F} from sequentially arriving subjects. We define a *treatment assignment policy* (of the social planner) as a dynamic strategy for assigning subjects to treatment arms that learns the best-performing experts in \mathfrak{F} using information contained in \mathcal{I}_t . Specifically, a treatment assignment policy (constructed upon available data) is given by a sequence $\{\mathbf{p}_t : \mathcal{I}_t \times \mathcal{X} \rightarrow \Delta^K | t = 1, \dots, T\}$, where \mathbf{p}_t is an \mathcal{I}_t -measurable map that maps subject t 's covariates \mathbf{x}_t to a probability for each treatment arm. As such, the treatment of subject t is allocated according to $k_t | (\mathcal{I}_t, \mathbf{x}_t) \sim \mathbf{p}_t(\mathbf{x}_t)$.

The learning algorithms that we consider in this paper build an assignment policy by aggregating experts in \mathfrak{F} according to their performance up to period t . As

we present in the next section, what drives the aggregation formula for \mathbf{p}_t in the EXP algorithm and derivative algorithms is a probability distribution over \mathfrak{F} that we denote by \mathbf{q}_t . It is convenient to separate the construction of \mathbf{q}_t according to the cardinality of \mathfrak{F} .

Assumption 3 (Complexity) Either:

- (a) Finite experts— $K < \infty$ and \mathfrak{F} has finite cardinality equal to N ; or
- (b) Experts with complexity controlled by a finite VC dimension— $K = 2$ and \mathfrak{F} consists of deterministic rules of the form $f(\mathbf{x}) = 1(\mathbf{x} \in G)$, $G \subset \mathcal{X}$, and the class $\mathcal{G} = \{G\}$ that spans \mathfrak{F} has finite VC dimension equal to D .¹

In Sect. 3.1, we study the case of Assumption 3a, where the number of treatment arms can be more than two and some experts in \mathfrak{F} are allowed to be probabilistic but, crucially, the number of experts is finite. In Sect. 3.2, we study the case of Assumption 3b, where the number of treatment arms is allowed to be infinite but, crucially, there are only two treatment arms—further extension to more than two treatment arms is beyond the scope of this paper—and the set of experts comprises deterministic rules whose complexity is controlled by a finite VC dimension. We present the EXP4.P algorithm for each case, respectively.

2.3 The social planner’s preferences and regret

The objective of the social planner is to maximise some welfare criterion. Following the literature on bandit algorithms, we assume that the objective function of the social planner is the sum of the subjects’ outcomes, and that the social planner is risk-neutral.

We define the empirical and average welfare attained by implementing the recommendations of a single expert f throughout the entirety of the time horizon by

$$\hat{W}_T(f) \doteq \sum_{t=1}^T f(\mathbf{x}_t)^\top \mathbf{y}_t, \tag{2.2}$$

$$W_T(f) \doteq E_p \left(\sum_{t=1}^T f(\mathbf{x}_t)^\top \mathbf{y}_t \right) = T \cdot E_p(f(\mathbf{x}_t)^\top \mathbf{y}_t), \tag{2.3}$$

respectively, where the equality on the right-hand side of Eq. (2.3) follows by Assumption 1. Here, $f(\mathbf{x}_t)^\top \mathbf{y}_t$ is the welfare contribution of subject t weighted according to the randomisation over treatment arms that is induced by f . We note

¹ We let $\mathbf{X}^\ell \doteq \{\mathbf{x}^1, \dots, \mathbf{x}^\ell\}$ be a non-empty finite set with ℓ points in \mathcal{X} . Given \mathcal{G} —a class of subsets in \mathcal{X} —we define $N(\mathbf{X}^\ell) = |\{\mathbf{X}^\ell \cap G : G \in \mathcal{G}\}|$ as the number of different subsets of \mathbf{X}^ℓ picked out by $G \in \mathcal{G}$. The VC dimension of \mathcal{G} is the largest ℓ such that $\sup_{\mathbf{X}^\ell} N(\mathbf{X}^\ell) = 2^\ell$ holds. See Vapnik (1998) for more detailed explanation and numerous examples.

that the optimal welfare defined by $\sup_{f \in \mathfrak{F}} W_T(f)$ agrees with the social planner's target welfare of the super-population in the static setting of treatment choice in Kitagawa and Tetenov (2018).

Given $(k_t)_{t=1}^T$ —a sequence of treatment arms assigned to each subject—we define the *empirical regret* and *regret* as

$$\hat{R}_T \doteq \sup_{f \in \mathfrak{F}} \hat{W}_T(f) - \sum_{t=1}^T y_t(k_t), \quad (2.4)$$

$$R_T \doteq \sup_{f \in \mathfrak{F}} W_T(f) - \sum_{t=1}^T y_t(k_t), \quad (2.5)$$

respectively. \hat{R}_T is defined relative to the benchmark of the maximal empirical welfare, while R_T is defined relative to the maximal average welfare. Since covariates, potential outcomes (and so realised outcomes), and realised treatments are all random, both \hat{R}_T and R_T are random. We assess the performance of a treatment assignment policy generating $(k_t)_{t=1}^T$ by some distributional features of empirical regret or regret. In particular, the performance guarantees that we provide in Sect. 3 for the EXP4.P algorithm are stated in terms of a uniform high-probability upper bound for the right-tail of the distribution of \hat{R}_T or of R_T . The reason that we focus on a high-probability upper bound rather than the mean of the regret distribution is that bounding the mean can overlook the thick-tail phenomenon of the regret distribution for EXP-type algorithms (see, for instance, 2020, §Chapter12). Conversely, by noting that the mean of regret can be written as the integration of tail probabilities, we can obtain an upper bound for the mean regret based upon the high-probability upper bounds for the tail probabilities (as long as the implementation of the algorithm does not depend upon the confidence level). We reiterate that the regret criterion that we consider in this paper is defined in terms of the T subjects in the sample rather than in terms of the super-population of subjects, which is in contrast to the literature on best-arm identification in bandit problems.

3 The EXP4.P algorithm and performance guarantees

The EXP4.P algorithm is a learning algorithm that aims to find a best expert by efficiently balancing exploration and exploitation. It incorporates information about the effectiveness of previously undertaken interventions into the choice of current assignment by ascribing more weight to experts that have previously performed well. In this section, we introduce the EXP4.P algorithm and provide high-probability performance guarantees for the regret that it incurs. Our exposition depends upon whether Assumption 3a or Assumption 3b holds, and we separate our analysis of these two cases—modifying the EXP4.P algorithm accordingly to operate in each environment.

3.1 The F-EXP4.P variant

In this section, we assume that Assumption 3a holds such that there are a finite number of experts in \mathfrak{F} , and adapt the EXP4.P algorithm to this setting—we refer to this variant of the algorithm as F-EXP4.P.

Algorithm 1 (F-EXP4.P; Beygelzimer et al., 2011) Input the following objects.

- i. A collection of tuning parameters, denoted by (β, γ, η) , such that $0 \leq \beta \leq 1$, $0 \leq \gamma \leq 1$, and $\eta \geq 0$;
- ii. a finite class of experts, denoted by \mathfrak{F} , such that $\mathfrak{F} = \{f^1, \dots, f^N\}$ with $N = |\mathfrak{F}| < \infty$; and
- iii. a maximum outcome value, denoted by M , such that $0 < M < \infty$.

Let $q_1^i = 1/N$ for $i = 1, \dots, N$. For each $t = 1, \dots, T$, iterate

1. For each $k = 1, \dots, K$, calculate a policy weight via

$$p_t(k) = [1 - \gamma] \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i + \gamma/K. \tag{3.1}$$

2. Sample k_t from $\{1, \dots, K\}$ according to $(p_t(1), \dots, p_t(K))$.
3. For each $k = 1, \dots, K$, calculate an estimate of the associated potential outcome via

$$\tilde{y}_t(k) = [\beta \cdot M^2 + y_t(k_t) \cdot 1(k_t = k)]/p_t(k). \tag{3.2}$$

4. For each $i = 1, \dots, N$, calculate a score via

$$\tilde{s}_t^i = \sum_{k=1}^K f_k^i(\mathbf{x}_t) \cdot \tilde{y}_t(k), \tag{3.3}$$

and a cumulative score via²

$$\tilde{S}_t^i = \begin{cases} \tilde{s}_t^i & \text{if } t = 1, \\ \tilde{S}_{t-1}^i + \tilde{s}_t^i & \text{if } t = 2, \dots, T. \end{cases} \tag{3.4}$$

5. Update the probability weights over the experts via

$$q_{t+1}^i = \exp(\eta \cdot \tilde{S}_t^i) / \sum_{\ell=1}^N \exp(\eta \cdot \tilde{S}_t^\ell). \tag{3.5}$$

END □

² The exponential terms in this formula can grow quickly and so, to avoid computational overflow, it is practical to normalise \tilde{S}_t^i by subtracting $\max_{\ell} \tilde{S}_t^\ell$.

Similar to the EXP-type algorithms proposed in the machine learning literature, F-EXP4.P is premised around three key ideas.

First, F-EXP4.P administers treatment randomly according to \mathbf{p}_t , which we recall is a probability distribution over the treatment arms. This probability distribution evolves according to the outcomes that are realised in earlier periods. Specifically, the evolution of \mathbf{p}_t is tied to the evolution of \mathbf{q}_t , which we recall is a probability weighting over experts; \mathbf{p}_t assigns more weight to treatment arms that experts weighted heavily by \mathbf{q}_t recommend. The regularisation term of γ/K in Eq. (3.1) ensures a desirable amount of exploration, with the tuning parameter of γ set so as to converge to zero as $T \rightarrow \infty$. Conversely, \mathbf{q}_t is constructed based upon estimates of the average cumulative score for each expert—captured by the dependence of Eq. (3.5) upon $(\tilde{S}_t^i)_{i=1}^N$. Experts with higher cumulative scores are given more weight in \mathbf{q}_t . The combined implication of \mathbf{p}_t and \mathbf{q}_t is that a treatment arm whose average outcome is estimated to be higher is more likely to be assigned.

Second, to construct estimates of expert performance, F-EXP4.P enacts an inverse probability weighting. Ignoring the regularisation term of $\beta \cdot M^2$ in Eq. (3.2), the remaining term is the observed outcome inverse weighted by the probability of assigned treatment, which provides an unbiased estimate for the average potential outcome conditional on \mathcal{I}_t and \mathbf{x}_t . The regularisation term of $\beta \cdot M^2$ provides a default imputation value for the missing potential outcomes. Theorem 1 suggests a value for β that converges to zero—alongside the regularisation term of which it is a part—as $T \rightarrow \infty$.

Third, F-EXP4.P constructs the probability weighting over the experts so that it is proportional to the exponential tilting of estimates of their cumulative scores. The sensitivity of \mathbf{q}_t to realisations of data is controlled by η , which we emphasise is positive; the larger η is, the more sensitive are the weights to experts' cumulative scores, implying that the algorithm leans more towards exploitation than exploration. The choice of η balances the relative strength of the exploration and exploitation motives and so affects the concentration rate of the algorithm.

We now present a welfare-regret guarantee for F-EXP4.P in terms of the empirical regret. The choices of the tuning parameters presented in the next theorem optimises the rate of high-probability regret upper bound.

Theorem 1 (Beygelzimer et al., 2011, § Theorem 2) *Assume that Assumption 3a holds alongside Assumptions 1 and 2. Let $0 < \delta < 1$ and $\omega \doteq \sqrt{\ln(N/\delta)/K}$, and set*

$$\begin{aligned} \beta &= \omega \cdot \sqrt{1/T} \cdot 1/M, \\ \gamma &= \omega \cdot \sqrt{1/T} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot K, \\ \eta &= \omega \cdot \sqrt{1/T} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot 1/2M, \end{aligned} \tag{3.6}$$

as the parameters of the EXP4.P algorithm. Provided that \mathfrak{F} includes a randomising expert f^{random} that ascribes equal probability to each treatment arm (or includes a subset of experts that can be linearly combined to mimic one) and $\max(\omega^2, 4K \cdot \ln(N)) \leq T$, then

$$\hat{R}_T \leq 7M \cdot \omega \cdot \sqrt{K^2 \cdot T} = 7M \sqrt{K \cdot T \cdot \ln(N/\delta)} \tag{3.7}$$

holds with probability at least $1 - \delta$.

We present proof of Theorem 1 in Sect. A. Theorem 1 extends Beygelzimer et al. (2011, §Theorem2), which requires that the potential outcomes be contained in the unit interval—rather than bounded and non-negative, as we assume. This extension modifies the choice of tuning parameters from that of Beygelzimer et al. (2011, §Theorem2). The maximal regret bound that we obtain then differs from that of Beygelzimer et al. (2011, §Theorem2) in that the regret that we obtain is scaled by M , and we include this as an additional parameter. We emphasise that our maximal regret bound—and the algorithm itself—coincides with that of Beygelzimer et al. (2011, §Theorem2) if the potential outcomes are contained in the unit interval.

Remark 3 We note that the assignment policies that are delivered by F-EXP4.P (with tuning parameters chosen according to Eq. (3.6)) are not invariant to the addition of a constant to the outcome variable. This is because the inverse probability weighted estimate of the expert’s score (denoted by \tilde{s}_t^i) is not equivariant to additive constants. Accordingly, in those situations in which outcomes are observed to take negative values, the allocations produced by Algorithm 1 can be sensitive to the choice of constant that is added to outcomes so as to satisfy the nonnegativity requirement of Assumption 2.

3.2 The VC-EXP4.P variant

The standard setting in policy learning is one in which \mathfrak{F} contains infinitely many treatment assignment rules. F-EXP4.P restricts the class of treatment assignment rules to be finite (i.e., only a finite number of individualised assignment rules are allowed), and so is incompatible with the standard setting. In this section, we assume that Assumption 3b holds such that there are infinitely many deterministic (i.e., non-random) experts in \mathfrak{F} with complexity controlled by a finite VC dimension, and adapt the EXP4.P algorithm to this setting—we refer to this variant of the algorithm as VC-EXP4.P.

Algorithm 2 (VC-EXP4.P; Beygelzimer et al., 2011) Input the following objects.

- i. A duration for the coarsening phase, denoted by τ , such that $\tau \in \mathbb{R}_{++}$ and $\tau < T$;
- ii. a class of experts, denoted by \mathfrak{F} , of finite VC dimension; and
- iii. a maximum outcome value, denoted by M , such that $0 < M < \infty$.

Generate treatment assignments by implementing the following algorithm:

1. For $t = 1, \dots, \lceil \tau \rceil$, sample k_t randomly from the uniform distribution on $\{1, \dots, K\}$.

2. Coarsen \mathfrak{F} to $\mathfrak{G} \subset \mathfrak{F}$ such that for any $f \in \mathfrak{F}$, there exists exactly one $g \in \mathfrak{G}$ satisfying $f(\mathbf{x}_t) = g(\mathbf{x}_t)$ for all $t = 1, \dots, \lceil \tau \rceil$.
3. Let $N = |\mathfrak{G}| - 1$ and implement F-EXP4.P over $t = \lceil \tau \rceil + 1, \dots, T$ using $\mathfrak{G} \cup f^{\text{random}}$ as the set of experts.

END □

The key innovation of VC-EXP4.P is the initial refinement of \mathfrak{F} to \mathfrak{G} , which occurs during an initial coarsening phase. Once \mathfrak{G} is formed, the algorithm proceeds to a run phase during which F-EXP4.P is enacted on $\mathfrak{G} \cup f^{\text{random}}$. The reason that we append f^{random} to the coarsened class is to meet the assumption for the regret guarantee of F-EXP4.P shown in Theorem 1. Whether this coarsening is straightforward to implement or not depends upon how \mathfrak{F} is specified. For instance, when \mathfrak{F} comprises assignment rules based upon linear eligibility scores, we present a feasible algorithm for hyperplane arrangements to refine \mathfrak{F} to \mathfrak{G} .

We now present a welfare-regret guarantee for VC-EXP4.P that holds for particular choices of the tuning parameters and duration of the coarsening phase.

Theorem 2 (Beygelzimer et al., 2011§ Theorem 5) *Assume that Assumption 3b holds alongside Assumptions 1 and 2. Let $0 < \delta < 1$ and set*

$$\begin{aligned} \beta &= \omega \cdot \sqrt{1/[T - \lceil \tau \rceil]} \cdot 1/M, \\ \gamma &= \omega \cdot \sqrt{1/[T - \lceil \tau \rceil]} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot K, \\ \eta &= \omega \cdot \sqrt{1/[T - \lceil \tau \rceil]} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot 1/2M, \end{aligned} \tag{3.8}$$

as the parameters of the EXP4.P algorithm given

$$\tau = \sqrt{T \cdot [2D \cdot \ln(T \cdot e/D) + \ln(3/\delta)]}, \tag{3.9}$$

as the duration of the coarsening phase. Provided that $D \ll T$,³ then, with a universal constant $c > 0$,

³ For the EXP4.P algorithm to be well-defined, \mathfrak{F} or \mathfrak{G} cannot be too rich. The condition that $\max(\omega^2, 4K \cdot \ln(N)) \leq T$ is necessary to ensure this when \mathfrak{F} is finite. When \mathfrak{F} is infinite, however, we require not only that this condition holds for the \mathfrak{G} that is obtained from it but that $\tau < T$. Since $\omega^2 \leq \tau$, it is necessary that

$$\max(8 \ln(N), 2 \ln(N) + \ln(2/\delta)) \leq T \tag{3.12}$$

holds, which, under $N = |\mathfrak{G}|$, translates to

$$T \in \{t : \max(8D \cdot \ln(\tau \cdot e/D), 2D \cdot \ln(\tau \cdot e/D) + \ln(2/\delta)) \leq t \mid \tau = \sqrt{t \cdot [2D \cdot \ln(t \cdot e/D) + \ln(2/\delta)]}\} \tag{3.13}$$

by Sauer’s Lemma. For Eq. (3.10) to hold generally then, the condition that $D \ll T$ is sufficient (i.e., the VC dimension is small relative to the time horizon).

$$R_T \leq cM \cdot \sqrt{\lceil \tau \rceil^2 + T \cdot \ln(3/\delta)} \tag{3.10}$$

holds with probability at least $1 - \delta$.

We present proof of Theorem 2 in Sect. A. An important implication of controlling the complexity of (deterministic) experts in \mathfrak{F} by limiting the VC dimension is that, given $(\mathbf{x}_t)_{t=1}^{\lceil \tau \rceil}$, we can construct $\mathfrak{G} = \{g : \mathcal{X} \rightarrow \{0, 1\}\}$ to represent the remaining experts in \mathfrak{F} in the sense that there exists exactly one $g \in \mathfrak{G}$ satisfying $f(\mathbf{x}_t) = g(\mathbf{x}_t)$ for all $t = 1, \dots, \lceil \tau \rceil$. The cardinality of \mathfrak{G} can be much smaller than 2^τ , the cardinality of binary functions mapping τ points. Theorem 2 implies that choosing τ to be of the order of $O(\sqrt{T \cdot D \cdot \ln(T)})$ gives a sufficient degree of coarsening for the run phase to have a high-probability regret bound of the order of $O(\sqrt{T \cdot D \cdot \ln(T)})$.

An important and practically relevant class of experts is the Linear Eligibility Score (LES) class (Kitagawa and Tetenov, 2018):

$$\mathfrak{L}_J \doteq \left\{ (1((1, \mathbf{x}^T)\beta < 0), 1((1, \mathbf{x}^T)\beta \geq 0))^T : \beta \in \mathbb{R}^{J+1} \right\}, \tag{3.11}$$

which corresponds to the class of 0-1 assignments spanned by hyperplane partitions in $\mathcal{X} \subset \mathbb{R}^J$. If \mathfrak{F} is an LES class, we can improve upon Eq. (3.12) by providing a tighter bound on the cardinality of \mathfrak{G} . Define, for $t \geq 2$

$$\text{Harding}(t, J) \doteq 2 \sum_{j=0}^J \binom{t-1}{j}, \tag{3.14}$$

as per Harding (1967). We note that $\text{Harding}(t, J)$ is the maximum number of distinct signed partitions of t points in \mathbb{R}^{J+1} —the ambient space—that can be induced by hyperplanes.⁴ As such, Eq. (3.14) corresponds to the maximum number of linear assignment rules that induce distinct allocations of t observations over the two treatment arms. It is straightforward to show that Eq. (3.14) is less than the corresponding bound on the cardinality of \mathfrak{G} that is implied by Sauer’s Lemma (Sauer, 1972; Shelah, 1972; Vapnik and Chervonenkis, 1971)—which is central to our derivation of Theorem 2—when $\mathfrak{F} = \mathfrak{L}_J$.

To implement improvement of Eq. (3.12) then, we leverage Eq. (3.14) to derive a tighter bound on the maximal cardinality of \mathfrak{F} when this coincides with the LES class.

Theorem 3 *Assume that Assumption 3b holds alongside Assumptions 1 and 2 and, moreover, that $\mathfrak{F} = \mathfrak{L}_J$. Let $0 < \delta < 1$ and set*

⁴ Harding (1967) in fact determines the maximum number of distinct partitions of t points in \mathbb{R}^{J+1} that can be induced by hyperplanes; the maximum number of distinct signed partitions is obtained trivially by doubling the maximum number of distinct partitions.

$$\begin{aligned}
 \beta &= \omega \cdot \sqrt{1/[T - \tau]} \cdot 1/M, \\
 \gamma &= \omega \cdot \sqrt{1/[T - \tau]} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot K, \\
 \eta &= \omega \cdot \sqrt{1/[T - \tau]} \cdot \sqrt{\ln(N)/\ln(N/\delta)} \cdot 1/2M,
 \end{aligned}
 \tag{3.15}$$

as the parameters of the EXP4.P algorithm given

$$\tau = \sqrt{T \cdot [2 \ln \circ \text{Harding}(T, J) + \ln(3/\delta)]},
 \tag{3.16}$$

as the duration of the coarsening phase. Provided that $J + 1 \ll T$,

$$R_T \leq cM \cdot \sqrt{\lceil \tau \rceil^2 + T \cdot \ln(3/\delta)}
 \tag{3.17}$$

holds with probability at least $1 - \delta$ with a univesal constant $c > 0$.

Coarsening the LES class can be formulated as a cell enumeration problem and achieved using an incremental enumeration algorithm—outlined below, alongside some practical recommendations.

3.3 Coarsening the class of linear eligibility assignment rules

Suppose that $\mathfrak{F} = \mathfrak{Q}_J$ so that $f \in \mathfrak{F}$ can be indexed by $\beta \in \mathbb{R}^{J+1}$. The fundamental insight that we build upon to coarsen \mathfrak{F} to \mathfrak{G} is that, for $f^i \in \mathfrak{G}$ and $f^\ell \in \mathfrak{G}$, there must exist some \mathbf{x}_t in $(\mathbf{x}_t)_{t=1}^{\lceil \tau \rceil}$ such that

$$\text{Sign}\left((1, \mathbf{x}_t^\top)\beta^i\right) \neq \text{Sign}\left((1, \mathbf{x}_t^\top)\beta^\ell\right).
 \tag{3.18}$$

In words, that we can find a subject (during the coarsening phase) for whom f^i and f^ℓ recommend different treatments.

Mathematically, Eq. (3.18) is equivalent to β_i and β_ℓ being located on opposite sides of the hyperplane defined by

$$H_t \doteq \{\beta : (1, \mathbf{x}_t^\top)\beta = 0\}.
 \tag{3.19}$$

The problem of constructing the coarse class for the LES class can be seen as a cell enumeration problem. That is, the problem of determining which cells—identified by a label (i.e., the concatenation of a sign for each subject in the sample according to whether the left- or right-hand side of Eq. (3.18) is true) and equivalent to a partition of the parameter space of β —are compatible with a given hyperplane arrangement, and calculating a point in the interior of each cell.

The cell enumeration formulation highlights that \mathfrak{G} is not unique. If β^i and β^ℓ both induce the same label in the sample then \mathfrak{G} can include β^i or β^ℓ but not both, and are said to belong to a common equivalence class—or cell. The invariance of treatment allocations in the same equivalence class is also exploited by Pinkse (1993, see also Rosen & Ura, 2019) in the closely-related problem of maximum score estimation.

The incremental enumeration algorithm introduced by Rada and Černý (2018, building upon earlier work by Avis & Fukuda, 1996; Sleumer, 1998) and

improved upon by Gu and Koenker (2022) is an efficient way to construct \mathfrak{G} . It considers the enumeration problem one hyperplane at a time, exploiting the similarity of neighbouring cells. The number of operations that the algorithm involves is proportional to the number of cells that are compatible with a hyperplane arrangement, which Buck (1943) establishes is less than τ^J , and far fewer than the 2^τ partitions that a naïve *brute force* algorithm operates over and that is infeasible for even moderately-sized samples.

Before discussing how the incremental enumeration algorithm works, however, we clarify what we mean by *cell* and by *label*. To facilitate this clarification, imagine that we (are able to and do) plot a hyperplane of the form that is given in Eq. (3.19) for each subject in the sample. A cell is the intersection of half-spaces defined by a given hyperplane arrangement. A label is a consistent identifier of whether a point is located in the positive or negative half-spaces defined by a hyperplane arrangement. For instance, by recording whether a cell is above or below each hyperplane in the sense of Eq. (3.18), we can identify it by a sequence of pluses and minuses, which we refer to as a label. A naïve consideration of the problem then is to consider a label—of which there are 2^τ possible permutations of pluses and minuses—and to determine whether such a label identifies a non-empty cell. Given that each cell is defined as the intersection of half-spaces, this amounts to solving a linear programme in which the objective is to maximise a slackness variable $0 \leq r \leq 1$ subject to the constraints of the form;

$$\text{Constraint}_\ell = \begin{cases} (1, \mathbf{x}_\ell^\top)\beta_- - (1, \mathbf{x}_\ell^\top)\beta_+ + r \cdot \|(1, \mathbf{x}_\ell^\top)\|_2 \leq 0, & \text{if } \text{Sign}((1, \mathbf{x}_\ell^\top)\beta) = +, \\ (1, \mathbf{x}_\ell^\top)\beta_+ - (1, \mathbf{x}_\ell^\top)\beta_- + r \cdot \|(1, \mathbf{x}_\ell^\top)\|_2 \leq 0, & \text{if } \text{Sign}((1, \mathbf{x}_\ell^\top)\beta) = -, \end{cases} \tag{3.20}$$

where $\beta_+ \in \mathbb{R}_+^{J+1}$ and $\beta_- \in \mathbb{R}_+^{J+1}$ are the positive and negative parts of β , respectively, in the sense that $\beta = \beta_+ - \beta_-$ holds, and (β_+, β_-) are choice variables in the linear programme.

The incremental enumeration algorithm exploits the fact that if there does not exist a feasible solution to the linear programme for t subjects (i.e., the linear programme with constraints of the form that is given in Eq. (3.20)—one constraint for each $\ell = 1, \dots, t$) for a given label, then there is no feasible solution to the linear programme for $t + 1$ subjects that appends this label by a plus or a minus. Put simply, empty cells cannot be divided.

We now present a (deliberately) simple description of the steps of the incremental enumeration algorithm (see Rada & Černý 2018 for a more detailed description).

Algorithm 3 (Incremental enumeration; Rada & Černý, 2018) The user is required to input the following objects.

- i. A duration for the coarsening phase, denoted by τ , such that $\tau \in \mathbb{R}_{++}$ and $\tau < T$; and
- ii. a LES class of experts \mathfrak{Q}_J of finite VC dimension $D = J + 1$.

For $t = 1, \dots, \lceil \tau \rceil$, iterate.

- a. If $t = 1$ then propose + and – as labels; else append + and – separately to any existing labels comprising $t - 1$ pluses and minuses.
- b. Construct and solve a linear programme with constraints of the form that is given in Eq. (3.20)—one constraint for each $\ell = 1, \dots, t$ —for each proposed label.
- c. If the linear programme associated with a proposed label has a feasible solution then keep this label and store the solution; else disregard the proposed label.

For each stored label, construct and solve a linear programme with constraints of the form that is given in Eq. (3.20)—one constraint for each $\ell = 1, \dots, t$ —and return \mathcal{G} as the collection of $\beta = \beta_+ - \beta_-$ that are found as part of the solution.

END

□

Since the total number of labels that can be proposed at any step of Algorithm 3 is twice the number of labels that are carried forward from the previous step, the total number of labels that need to be checked is given by Eq. (3.14). Provided that a label is admitted, then a natural witness to the associated cell is its Chebyshev centre⁵; the solution to the linear programme for t subjects is the pseudo-Chebyshev centre of a cell and is, incidentally, also the witness to its division (i.e., any non-empty cell that is obtained from it upon the addition of another constraint).

Figure 1 illustrates the output of the incremental enumeration algorithm for four randomly generated points. Every admissible cell is the mirror image of another; it is, therefore, sufficient to characterise the full set of admissible cells by excluding all cells that are mirror images of another cell in the characterisation. We exploit this property in Fig. 1 to limit the number of treatment rules that we need to state. The generated points are compatible with a total of 14 cells, which is fewer than the 16 cells that naïve calculation suggests.

4 Linear experiments

To provide some insight into the practical implementation and performance of VC-EXP4.P,⁶ we create several artificial datasets, run VC-EXP4.P with the LES class of assignments, and assess its welfare-performance.

⁵ The Chebyshev centre of a bounded convex set is defined as the origin of the largest possible circle that can be drawn that is entirely enclosed within that set; by specifying an upper limit on the radius of the circle, we can compute something akin to the Chebyshev centre for a possibly unbounded convex set. We emphasise that cells need not be unbounded, with notable examples being + and – (i.e., the cells that are compatible with a single hyperplane).

⁶ We use three Intel Core i7-8700 CPU @ 3.20GHz computers equipped with 32GB RAM running Ubuntu 20.04.6 LTS (Focal Fossa) or Windows 10 Enterprise LTSC and R 4.3.0.

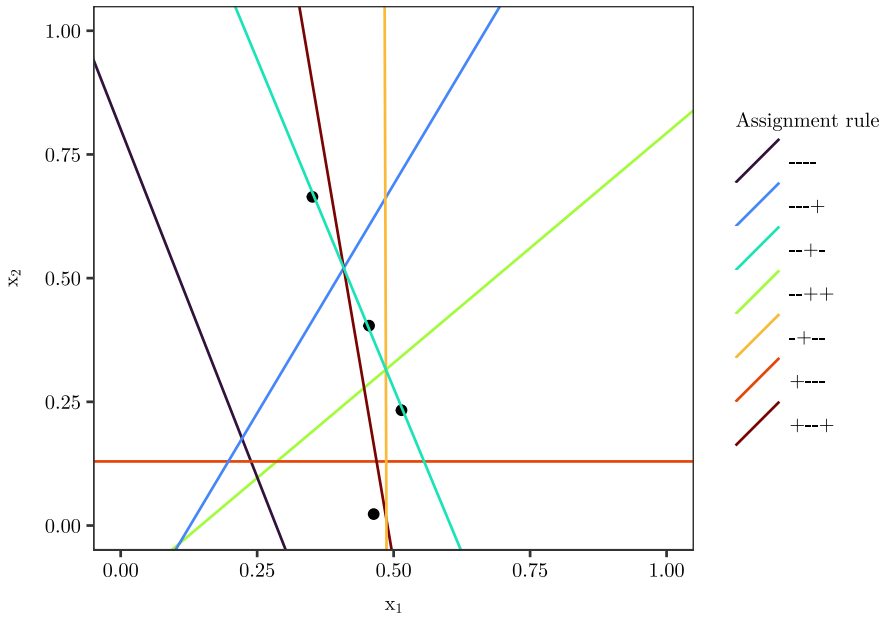


Fig. 1 A sufficient characterisation of admissible treatment rules for $\tau = 4$ observations

4.1 Design

We construct 1000 artificial datasets, each comprising 1000 realisations of a collection of two covariates—which we denote by $\mathbf{x}_t = (x_{t,1}, x_{t,0})^T \in \mathbb{R}^2$, as before—and of a collection of two private shocks—which we denote by $\mathbf{u}_t = (u_{t,1}, u_{t,0})^T \in \mathbb{R}^2$. The covariates and private shocks—which play the role of underlying heterogeneity—together generate a subject’s potential outcomes. We fix the realisations of the covariates across the datasets (i.e., \mathbf{x}_t is the same in every dataset) but allow the realisations of the private shocks to vary; by also fixing τ , we can use the same \mathcal{G} throughout our simulations to significantly reduce computation time. We relate the covariates and unobserved heterogeneity to the potential outcomes via a log-normal specification. We maintain log-normality throughout the remainder of Sect. 4 only, assuming that the covariates are generated from a uniform distribution on the unit square. Motivating our choice of specification is the appropriateness of a log-normal distribution as an approximation of income or other wealth-related indices—common measures of welfare in treatment choice applications—and the non-negativity of log-normal random distributions—in accordance with the non-negativity requirement of Assumption 2. The log-normality assumption does not, however, satisfy the boundedness requirement of Assumption 2. To implement VC-EXP4.P, we use the first τ observations to both compute \mathcal{G} and to inform our choice of the tuning parameters of the F-EXP4.P-step that occurs during the run phase.

Assumption 4 (*Log-normality*) The relationship between the covariates and unobserved heterogeneity, and the potential outcomes is governed by

$$\begin{aligned} y_{t,1} &= \exp(x_{t,1} - x_{t,2} + u_{t,1} - \sigma^2/2), \\ y_{t,0} &= \exp(u_{t,0} - \sigma^2/2), \end{aligned} \quad (4.1)$$

where $\mathbf{u}_t \sim \text{Normal}((0, 0)^T, \sigma^2 \cdot \mathbb{I}_2)$.

We note that, aside from its interpretation as a standard deviation, σ controls the signal-to-noise ratio under Assumption 4, and is, therefore, related to what we regard as the degree of difficulty of learning the optimal policy of an artificial dataset.

The first-best optimal assignment rule under Assumption 4 is given by $1(x_{t,1} \geq x_{t,2})$. This rule belongs to the LES class.

A property of the EXP4.P algorithm—and other exponential tilting procedures—is that it more heavily updates the current policy when the realised outcome is large. Compressing the distribution of effects (i.e., the differences between the potential outcomes) is, in theory, likely to decrease the convergence rate of the EXP4.P algorithm to whatever assignment rule is best-in-class.

Our intention in normalising the potential outcomes by $\sigma^2/2$ is to ensure comparability of the means of the potential outcomes for different standard deviations of the unobserved heterogeneity. To guarantee that Assumption 2 holds, we normalise the potential outcomes by dividing Eq. (4.1) by the maximum value of the potential outcomes that is realised for a given standard deviation of the unobserved heterogeneity,⁷ but reverse this normalisation prior to presenting the results of our experiments to maintain comparability.

To vary the difficulty of learning, we manipulate the variance of the unobserved heterogeneity in Eq. (4.1). We then measure the corresponding difficulty via the population probability of misclassification under the optimal rule.

$$\text{Difficulty} = \Pr(\text{Sign}(y_{t,1} - y_{t,0}) \neq \text{Sign}(x_{t,1} - x_{t,2})). \quad (4.2)$$

Figure 2 plots the relationship between the standard deviation of the unobserved heterogeneity and the measure of difficulty; the black line is the theoretical difficulty implied by Eq. (4.2).

4.2 Implementation

We restrict attention to LES rules with two covariates and set $\mathfrak{F} = \mathfrak{L}_2$. To coarsen \mathfrak{L}_2 , we utilise the incremental enumeration algorithm described in Sect. 3.3. We find that 31,154 assignment rules are compatible with the first 177 realisations of the covariates (the duration of the coarsening phase suggested by Eq. (3.16) given

⁷ For these experiments, we leverage our position as the oracle to set M equal to the maximum value of both potential outcomes across all of the artificial datasets. In practice, this is not feasible and so we suggest using the first τ observations to estimate some maximal value, or else to rely on insight or economic theory.

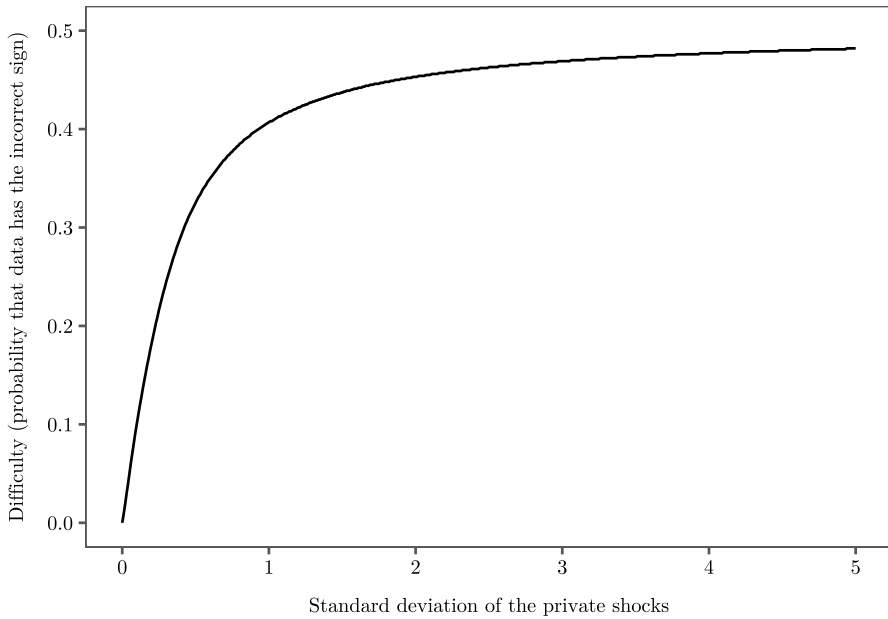


Fig. 2 The difficulty of Eq. (4.1) computed using Monte Carlo approximation

Eq. (3.6)),⁸ which we emphasise are constant across all 1000 artificial datasets—we only vary the unobserved heterogeneity and so the potential outcomes. For comparison, this is the same number of assignment rules that Harding (1967) suggests could be admitted.

To assess whether the coarsening phase generate a sizeable welfare loss, it is useful to know whether the coarse class contains the optimal rule, or assignment rules that are close to it. We find that the coarse class does not contain the optimal rule, but does contain several assignment rules that are close to it. In particular, the coarse class contains an assignment rule that allocates only one individual (out of 1000) differently to the oracle rule.

To draw meaningful conclusions about the performance of the EXP4.P algorithm, we also implement two other simple estimators whose performance we can use as a benchmark for comparison. We implement (i.) the optimal rule through $t = 1, \dots, 1000$ —what we refer to as the oracle rule; and (ii.) randomisation over the first 177 observations followed by the best-performing deterministic assignment rule in \mathcal{G} according to the empirical welfare criterion of Kitagawa and Tetenov (2018)—what we refer to as τ -EWM.

⁸ As a guide, our implementation of the incremental enumeration algorithm completes in seven minutes on our computer system.

4.3 Results

We conduct numerous experiments that vary the standard deviation of the unobserved heterogeneity or the tuning parameters of the EXP4.P algorithm. We set δ to the standard 95%-level

We capture the behaviour of the EXP4.P algorithm by computing the probability that the social planner's chosen policy enacts the same intervention as the optimal assignment rule. This probability can be computed for every observation and is induced by \mathbf{p}_t —the probability distribution over the treatment arms—underlying the assignment probability of the EXP4.P algorithm. A higher chance of correct classification (i.e., coincidence of the treatment assignment of VC-EXP4.P and that of the optimal rule) leads to a lower regret; how much and how quickly the probability of correct classification increases over time is informative as to the performance of the EXP4.P algorithm. We plot the probability of correct classification as Figs. 3, 4, 5 and 6.

In Figs. 3, 4, 5 and 6, the black ribbon captures the sample probability that the policy enacts the same intervention as the optimal rule in each of the 1000 artificial datasets that we construct; the white line that bisects the ribbon is the median probability amongst these datasets. In Figs. 4, 5 and 6, the orange line that dissects the ribbon is the corresponding median probability when the standard deviation of the unobserved heterogeneity is set to zero and is included for comparison.

We observe that the planner's chosen policy is more likely to coincide with the treatment assignment of the optimal rule over time but that the magnitude and

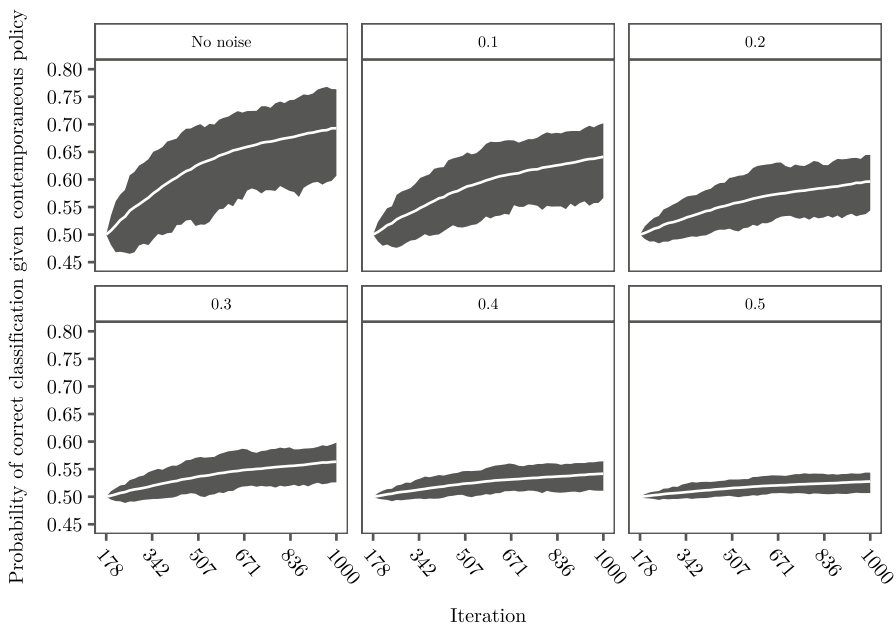


Fig. 3 The performance of VC-EXP4.P as the standard deviation of the unobserved heterogeneity increases

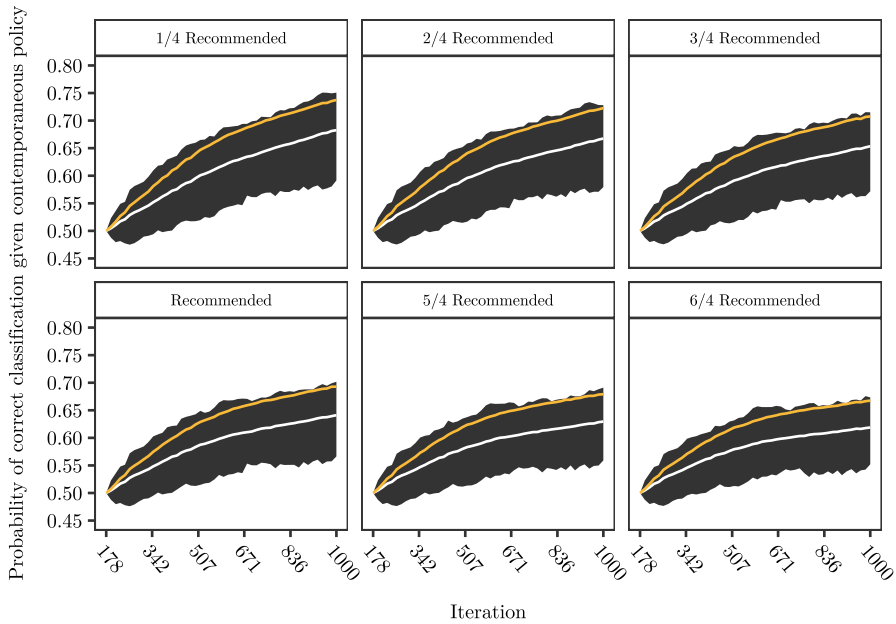


Fig. 4 The performance of VC-EXP4.P as β varies given $\sigma = 0.1$

rate of this improvement largely depends upon the difficulty and the choice of tuning parameters.

First, we investigate how changing the standard deviation of the unobserved heterogeneity affects performance. In each panel of Fig. 3 (from left to right) we increase the standard deviation of the unobserved heterogeneity, holding fixed the tuning parameters of the EXP4.P algorithm at the level recommended by Theorem 1. We emphasise that although each dataset reveals identical sequences of the covariates, at every iteration both the intervention (even for an identical policy) and the potential outcomes that are realised differ and it is this attribute that generates the distribution of probabilities that is captured by the black ribbon. As is to be expected, as the standard deviation of the unobserved heterogeneity increases, the EXP4.P algorithm finds it increasingly difficult to learn an optimal policy. The difficulty corresponding to each panel is zero, 10.1%, 18.2%, 24.5%, 29.1%, and 32.5%, respectively. The EXP4.P algorithm struggles to show even minimal signs of convergence for comparatively high standard deviations of the unobserved heterogeneity.

Second, we investigate how changing β affects the performance of the EXP4.P algorithm. In each panel of Fig. 4 (from left to right) we scale β relative to its recommended value in Theorem 1, holding fixed the other tuning parameters of the EXP4.P algorithm at the level recommended by Theorem 1. The exploration motive is increasing in β , which governs how much to score unrealised treatment arms (and so too the experts that recommend these treatment arms). Fig. 4 suggests that the EXP4.P algorithm converges more slowly to the optimal rule as β increases.

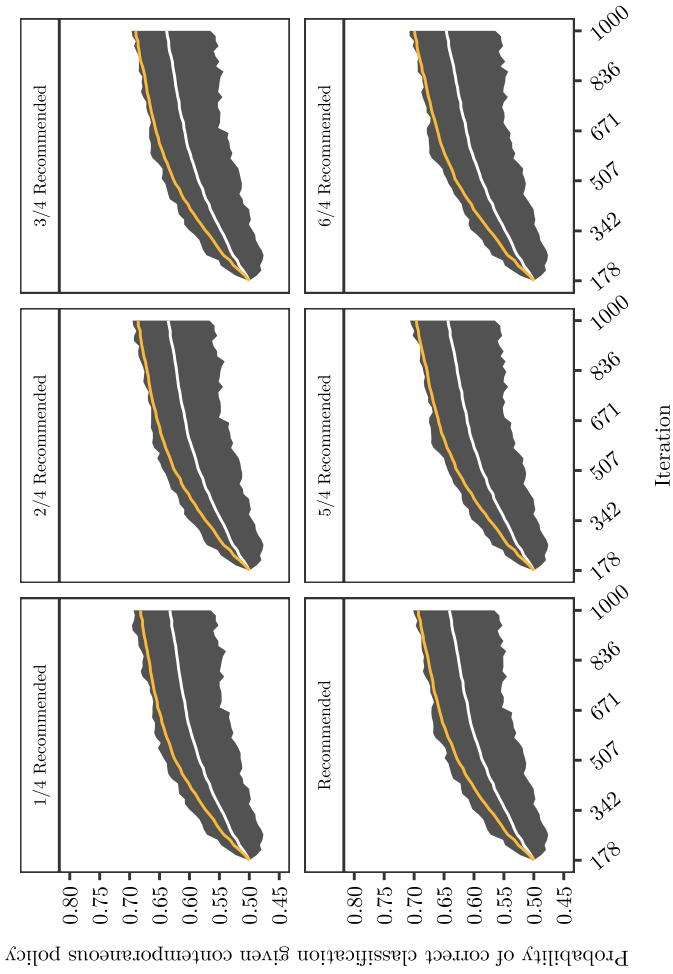


Fig. 5 The performance of VC-EXP4.P as γ varies given $\sigma = 0.1$

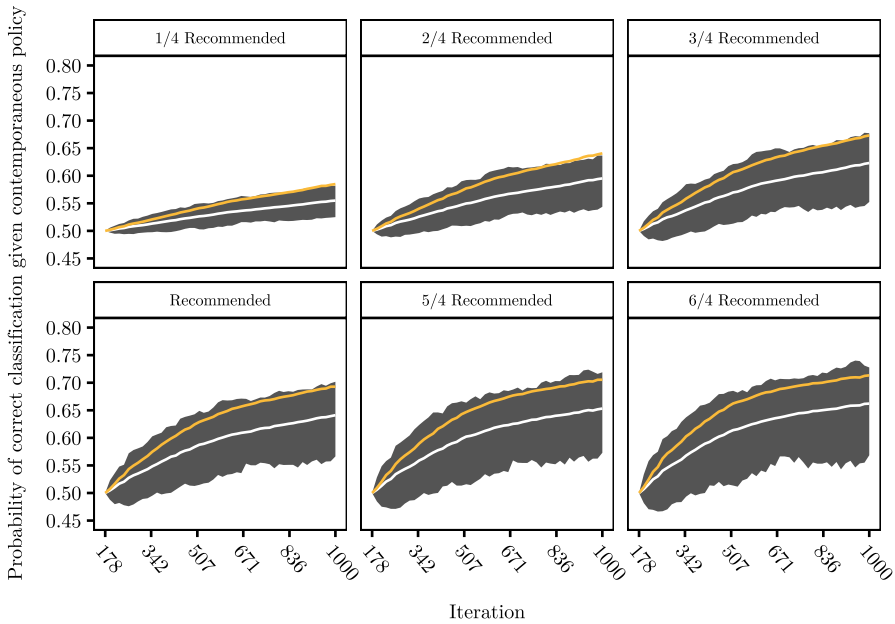


Fig. 6 The performance of VC-EXP4.P as η varies given $\sigma = 0.1$

Third, we investigate how changing γ affects the performance of the EXP4.P algorithm. In each panel of Fig. 5 (from left to right) we scale γ relative to its recommended value in Theorem 1, holding fixed the other tuning parameters of the EXP4.P algorithm at the level recommended by Theorem 1. The exploration motive is increasing in γ , which governs the willingness to randomise versus adopting expert recommendations (weighted by the expert weights). Fig. 5 suggests that the EXP4.P algorithm is not particularly sensitive to the choice of γ .

Fourth, we investigate how changing η affects the performance of the EXP4.P algorithm. In each panel of Fig. 6 (from left to right) we scale η relative to its recommended value in Theorem 1, holding fixed the other tuning parameters at the level recommended by Theorem 1. The exploitation motive is increasing in η , which governs the sensitivity of the assignment probability to realised outcomes. Fig. 6 suggests that the EXP4.P algorithm is particularly sensitive to the choice of η —more so than to the values of the other tuning parameters. The median probability of correct classification is increasing in η (i.e., the white line reaches a higher maximum), although this is accompanied by an increase in the variance of the correct classification sample probability (i.e., the black ribbon widens). This trade-off between the average and variance of performance resembles the bias-variance trade-off that commonly appears in nonparametric estimation for the choice of a smoothing parameter.

The patterns that we observe in Figs. 3, 4, 5 and 6 mirror differences in the average welfare-level that the EXP4.P algorithm attains that are evident in Tables 1 and 2—experiments in which VC-EXP4.P converges slowly to the oracle rule correspond to those experiments in which it attains a lower average welfare-level. Theorems 1 to

Table 1 The average welfare-level attained by various estimators as the standard deviation of the unobserved heterogeneity increases

Estimator	σ					
	No noise	0.1	0.2	0.3	0.4	0.5
Oracle rule	2.001	2.001	2.000	2.000	2.000	1.999
τ -EWM	1.891	1.891	1.891	1.890	1.886	1.884
VC-EXP4.P	1.782	1.761	1.745	1.733	1.726	1.721
Coarsening phase	1.707	1.707	1.708	1.708	1.709	1.709
Run phase	1.798	1.772	1.753	1.739	1.729	1.723

Table 2 The average welfare-level attained by various estimators as the parameters of the EXP4.P algorithm vary given $\sigma = 0.1$

Estimator	Tuning parameter value					
	1/4 Rec	2/4 Rec	3/4 Rec	Rec	5/4 Rec	6/4 Rec
Oracle rule	2.001	2.001	2.001	2.001	2.001	2.001
τ -EWM	1.890	1.890	1.890	1.890	1.890	1.890
VC-EXP4.P (β)	1.771	1.767	1.764	1.761	1.758	1.755
Coarsening phase	1.707	1.707	1.707	1.707	1.707	1.707
Run phase	1.784	1.780	1.776	1.772	1.769	1.766
VC-EXP4.P (γ)	1.765	1.764	1.762	1.761	1.759	1.758
Coarsening phase	1.707	1.707	1.707	1.707	1.707	1.707
Run phase	1.777	1.776	1.774	1.772	1.771	1.769
VC-EXP4.P (η)	1.728	1.742	1.752	1.761	1.767	1.773
Coarsening phase	1.707	1.707	1.707	1.707	1.707	1.707
Run phase	1.733	1.749	1.762	1.772	1.780	1.787

3 and their proofs necessitate only fairly weak constraints on the values of the parameters of the EXP4.P algorithm with the definitions contained therein serving only, in practice, as recommendations. Theorems 1 to 3 and their proofs can, for instance, be suitably adjusted to facilitate more aggressive updating, such as Figs. 3, 4, 5 and 6 and Tables 1 and 2 suggest is necessary for the EXP4.P algorithm to attain a higher average welfare-level. How much more aggressively the EXP4.P algorithm would need to update to attain an average welfare-level that is similar to that attained by τ -EWM is, however, unclear but would appear substantial. Not only does τ -EWM outperform VC-EXP4.P across all of our experiments but does so whilst incurring far less computational expense.

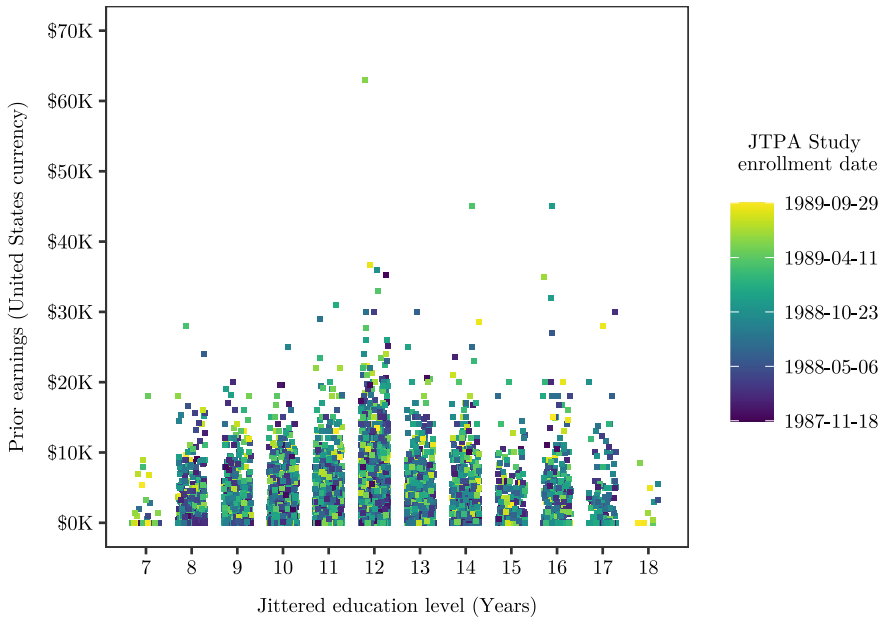


Fig. 7 Date of enrollment into the National JTPA Study

5 National job training partnership act study

The National Job Training Partnership Act (JTPA) Study comprises 9223 observations of subjects, their education level and prior earnings. Applicants were enrolled onto the Study throughout the years 1987–1989, and were randomly allocated to one of two treatment arms—some applicants were extended training, job search assistance and other services provided by the JTPA, with the remaining applicants denied this support. Along with information collected prior to enrollment, the Study also collected administrative and survey data relating to applicants' earnings in the 30 months following its start. Further details about the data and the Study can be found elsewhere (see, for instance, Bloom et al., 1997). We restrict attention to a sample of 9,223 observations for which data on years of education and pre-programme earnings amongst the sample of adults (aged 22 years and older) used in the original evaluation of the programme and in subsequent studies (Abadie et al., 2002; Bloom et al., 1997; Heckman et al., 1997) is available. Applicants in the Study were assigned to the two treatment arms in the ratio of two to one. Like Kitagawa and Tetenov (2018), we define the intervention to be the initial assignment of treatment, rather than the actual take-up due to the presence of non-compliance in the experiment.

5.1 Design

We exploit the sequential arrival of applicants in the JTPA Study to test the performance of VC-EXP4.P. We sort applicants by the date of their enrollment onto the Study, randomly ordering applicants with the same date of enrollment to mitigate any possible systematic issues arising from data entry—although it seems apparent from Fig. 7 that there are no systematic time trends in the distribution of observables during the Study. We maintain this order throughout. A restrictive feature of our analysis is that we assume that the social planner observes the realised outcome of an applicant *before* the next applicant arrives. This is certainly an unrealistic assumption and presumes a favourable scenario. The purpose of our analysis, however, is to study whether the EXP4.P algorithm performs well even in such an ideal situation. A more realistic analysis would replace the actual outcome by a short-run surrogate, as considered in Athey et al. (2019). We leave this extension for future research.

The Study data reports applicants' earnings only for the treatment arm that they were allocated to—we do not observe their counterfactual earnings for obvious reasons. Although implementation of the EXP4.P algorithm does not require the social planner to observe the outcomes associated with other treatment arms, our simulation exercises do since the treatment implemented by the EXP4.P algorithm need not coincide with the treatment arm that applicants were allocated to. We use random forest methods (Breiman, 2001) to regress applicants' earnings on their education level and prior earnings separately for each group of applicants—the applicants that were offered support, and the applicants that were denied support. We subtract the average cost \$774 from applicants' earnings for those subjects that were offered support, prior to running the random forests. Given that treatment was randomised in the JTPA study, we take these random forest regressions as estimators for the potential outcome regressions of $E[y_t(1)|\mathbf{x}_t]$ and $E[y_t(0)|\mathbf{x}_t]$. We use these regressions to form a prediction about applicants' earnings for each of the two treatment arms—with the applicant's education level and prior earnings as predictors—regardless of the treatment arm that they were allocated to. We also use these regressions to construct an empirical distribution of residual earnings for each treatment arm, by subtracting the difference between an applicant's earnings and their predicted earnings for the treatment arm that they were allocated to. By randomly drawing a residual earning from each empirical distribution and adding these to an applicant's predicted earnings for the corresponding treatment arm, we are able to construct a simulated outcome for each applicant for each of the two treatment arms, including the counterfactual one. We repeat this process to generate many simulated samples and apply VC-EXP4.P (setting δ to the standard 95%-level, and setting the values of the tuning parameters as per Theorem 3) to them to obtain a distribution of welfare-performance.

As we do for our previous simulation exercises, we compare the performance of VC-EXP4.P with several benchmarks. We implement (i.) the optimal rule, which

we determine from the random forest⁹; (ii.) the infeasible quintic rule, which is an optimal assignment rule among the class of single indices formed by a fifth-order polynomial of the two covariates; (iii.) an optimal assignment rule in the LES class; (iv.) to treat everyone; (v.) to treat no-one; (vi.) τ -EWM with the first 898 applicants used for coarsening \mathfrak{F}^{10} ; and (vii.) τ -EWM over the first 609 applicants. Given that VC-EXP4.P involves a coarsening phase, we can also compare the average welfare-level that is accrued during this sequence of observations to the average welfare-level during the subsequent run phase.

5.2 Implementation

Aside from enrollment date, Fig. 7 also reports education level and prior earnings in the sample. We note two things about the distribution of covariates. First, education level is recorded on a discrete scale (although data is artificially jittered for legibility), in years. An implication is that any assignment rule that allocates an applicant to a particular treatment arm must also allocate applicants with the same education level and lower prior earnings (or higher, depending upon the *sign* of the rule) to the same treatment arm. Second, prior earnings are highly concentrated for every education level. In particular, many applicants have zero prior earnings. An implication is that the Study can, in practice, be reduced to a far smaller number of unique observations (of education level and prior earnings) for the purpose of refining the assignment rules.¹¹

To ensure that the potential outcomes are non-negative, we add \$34,027 to every observation, subsequently subtracting this amount before reporting our results. Specifically, we add \$34,027 so as to ensure Assumption 2, since applicants' simulated earnings can otherwise be negative due to the lack of restriction that we impose on the random forests (predicted earnings and residual earnings can both be negative). We set M to \$257,300, which is the sum of the maxima of the predicted earnings and residual earnings that are returned by the random forests following this normalisation.¹²

To reduce the LES class, we utilise the incremental enumeration algorithm described in Sect. 3.3. We find that 104,252 assignment rules are compatible with the education level and prior earnings of the first 609 applicants to enroll on the

⁹ The infeasible optimal rule allocates is to offer support whenever the random forest predicts that the difference between the potential outcomes plus the difference in the average residuals between the two treatment arms—which is not zero—is positive, and denies support otherwise.

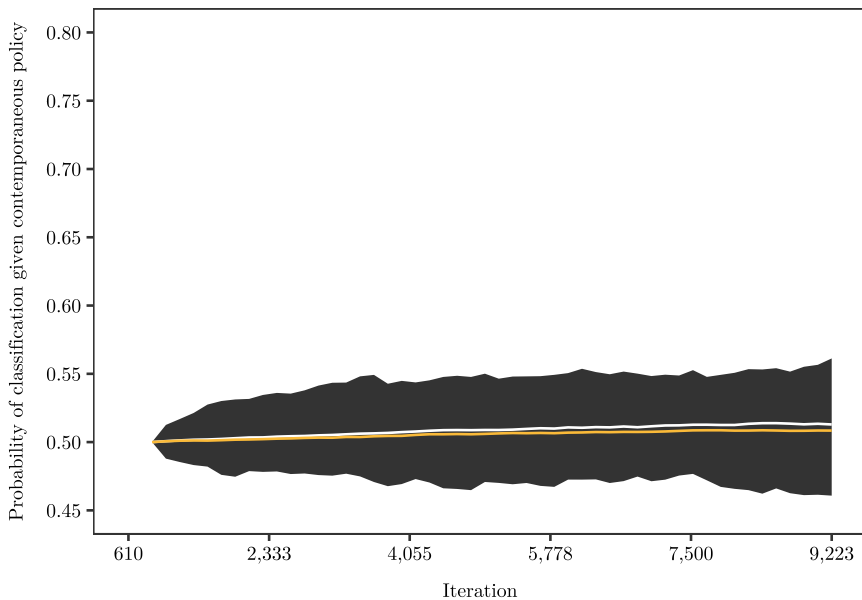
¹⁰ 898 roughly corresponds to the number of applicants in the study that were enrolled during the first 90 days (i.e., one quarter). For comparison, the suggested duration of the coarsening phase as per Theorem 2 is 609 applicants.

¹¹ The 609 applicants that are considered during the coarsening phase reduce to 365 unique observations.

¹² We note that the maximum potential outcome that could be observed in each simulation during the first 609 periods ranges from as little as \$98,023 to as much as \$198,337; the standard deviation of these observations is less than \$20,000 in each simulation. How to set M in the absence of full information is an open question.

Table 3 Distributional features of the welfare-level attained by various estimators

Estimator	Quantile							Mean
	0%	10%	25%	50%	75%	90%	100%	
Infeasible optimal rule	16,821	17,144	17,238	17,338	17,433	17,529	17,876	17,337
Infeasible quintic rule	16,312	16,615	16,714	16,815	16,911	17,001	17,275	16,812
Optimal LES rule	16,187	16,486	16,579	16,683	16,776	16,873	17,120	16,676
Treat everyone	16,000	16,314	16,406	16,513	16,609	16,705	16,996	16,508
Treat no-one	14,865	15,122	15,218	15,319	15,414	15,493	15,861	15,315
τ -EWM with $\tau = 898$	15,007	15,598	15,776	16,003	16,220	16,393	16,841	16,000
τ -EWM with $\tau = 609$	15,015	15,539	15,719	15,996	16,215	16,404	16,868	15,979
VC-EXP4.P	15,430	15,726	15,822	15,924	16,020	16,117	16,366	15,922
Coarsening phase	14,172	15,193	15,555	15,971	16,399	16,746	17,814	15,967
Run phase	15,395	15,711	15,817	15,918	16,018	16,122	16,363	15,918

**Fig. 8** The simulated performance of VC-EXP4.P based on the samples in the National JTPA Study

Study. For comparison, Harding (1967) suggests that up to 370,274 assignment rules could be admitted if the Study data did not include duplicates or applicants with the same prior earnings (particularly, zero prior earnings) and different education levels, or if education level were not discrete. As such, the EXP4.P algorithm can incur substantial computational expense—especially in terms of memory usage.

5.3 Results

We construct 1000 simulated samples that we use to assess the performance of VC-EXP4.P and our benchmark procedures. We report quantiles and the mean of the welfare-level that each method attains in Table 3. We plot the probability of classification as Fig. 8. The black ribbon captures the probabilities that the policy enacts the same intervention as the optimal LES rule at every iteration in each of the 1000 artificial datasets that we construct; the white line that bisects the ribbon is the median probability amongst these simulated datasets; the orange line that dissects the ribbon is the corresponding median probability when the comparison is the intervention dictated by the infeasible oracle rule—rather than the optimal LES rule.

We plot the allocation that is prescribed for the Study as Fig. 9. We contrast the prescription of the infeasible optimal rule with the prescription of the optimal LES rule that is to offer support to those subjects to the left of the downwards-sloping line. We note several findings.

First, on average, applicants benefit from training, job search assistance and other services provided by the JTPA, as is illustrated by the higher welfare-level attained by treating everyone versus treating no-one. The magnitude and sign of this effect is, however, known to be heterogeneous; individualised assignment can be beneficial, as is shown in Kitagawa and Tetenov (2018). Comparison of the welfare-level attained by the infeasible optimal rule (the plug-in rule based on the random forest estimates) and by the treat everyone policy are in line with heterogeneity in the sign of the treatment effect. The difference in welfare-level shrinks, however, if we

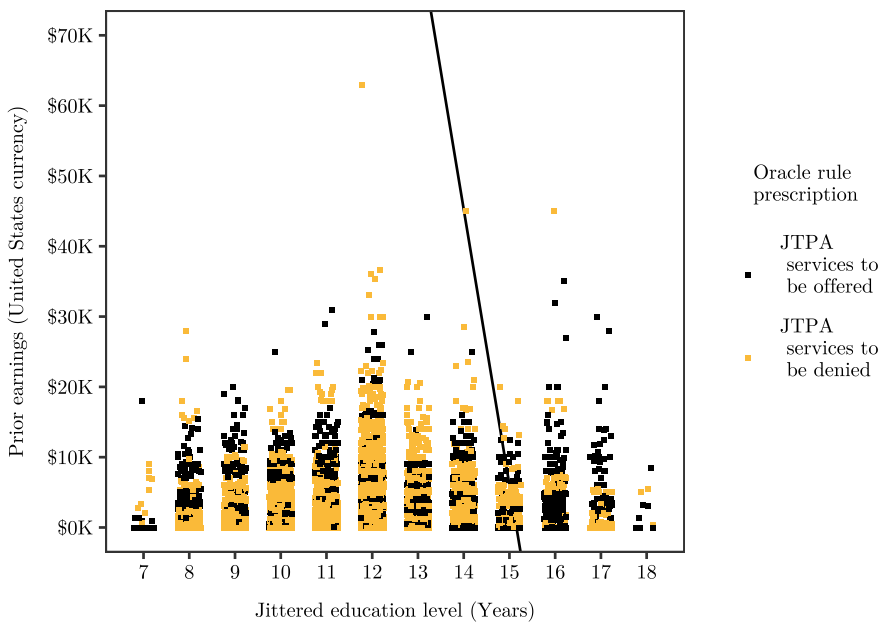


Fig. 9 The allocation that is prescribed by the infeasible oracle rule for the National JTPA Study

compare the results for the optimal quintic rule or optimal LES rule with the treat everyone policy.

Second, in terms of the average welfare for in-sample subjects, neither the EXP4.P algorithm nor EWM performs well, as is evidenced by the inferior mean welfare-levels that are attained by VC-EXP4.P and τ -EWM versus the treat everyone policy. Theorem 2 suggests that the welfare-level of VC-EXP4.P converges to the welfare-level of the optimal LES rule, but we do not observe this happening. This result is also indicative of the Study having a high degree of difficulty. The counterpart to Eq. (4.1) that we employ here is the probability that the sign of the difference in potential outcomes—as predicted by the random forest or by the closest assignment rule in the LES class—is different from the sign of the difference in potential outcomes that is realised. We find that the Study has a difficulty rating of 44.0% and 47.6% (relative to a maximum of 50.0%) according to these measures, respectively. This conclusion is also supported by Fig. 8, which shows that the probability of correct classification does not grow over time and remains close to one half—the same as can be achieved by pure randomisation using an unbiased coin. The infeasible optimal rule is highly non-linear, and rules in the LES class are unable to effectively replicate it—resulting in the large welfare gap between the infeasible optimal rule and the optimal LES rule.

Third, VC-EXP4.P attains a similar welfare-level to τ -EWM. This indicates that, in the JTPA Study sample, the welfare gain from performing sequential learning using VC-EXP4.P is small compared with the simpler two-stage learning procedure of τ -EWM, despite a substantial difference in computational cost.

In summary, the results of our simulations—based upon the JTPA Study sample—indicate limited benefits to implementing the EXP4.P algorithm, even in an unrealistically ideal scenario in which the social planner can update the assignment policies of as many as 8000 subjects and in which treatment response can be observed immediately after assignment. We attribute this limited benefit to substantial heterogeneity and non-linearity in treatment effect in the JTPA Study sample (i.e., the JTPA Study sample is a difficult sample). Although the effect of non-linearity can be mitigated somewhat by increasing the VC dimension of \mathfrak{F} —evidenced by our results for the infeasible quintic rule—increasing the complexity of \mathfrak{F} leads to a substantial increase in the computational cost of implementing the EXP4.P algorithm, particularly during the coarsening phase of VC-EXP4.P.

6 Conclusions

This paper studies applicability of the existing EXP4.P algorithm to the problem of treatment choice, where the social planner's goal is to learn an optimal individualised treatment assignment policy within a constrained class of assignment rules. For general bounded outcome variables, we provide a welfare-regret guarantee for the EXP4.P algorithm, and a practical method for coarsening the class of assignment rules—a necessary step when this class is infinite with complexity controlled by a finite VC dimension—that exploits the similarity of the coarsening problem to a hyperplane arrangement and cell enumeration problem.

We study the suitability of the EXP4.P algorithm to treatment choice problems through numerical analysis. Using a novel simulation design that is based upon the JTPA Study sample and that mimics the empirical context of assignment to job training, we assess the performance of the EXP4.P algorithm relative to oracle assignment rules and a non-adaptive benchmark policy based upon static EWM. Our main finding is somewhat discouraging. Specifically, we find that the EXP4.P algorithm can perform poorly in situations in which the standard deviation of unobserved heterogeneity is large relative to the size of conditional average treatment effects or in which non-linearity of these effects limits the welfare gain of simple (linear) assignment policies.¹³ The JTPA Study sample exhibits these properties and, as a result, implementing the EXP4.P algorithm does not deliver a notable welfare gain relative to the policy recommended by non-sequential EWM. Moreover, although our theoretical results establish uniform convergence of welfare-regret, we do not observe this convergence in practice.

We regard several issues as open questions and interesting topics for future research. First, implementation of the EXP4.P algorithm, importantly, requires that the time horizon (i.e., the number of periods over which allocation is to occur) is fixed and known in advance. This assumption is restrictive. Second, our analysis of the EXP4.P algorithm when the class of assignment rules is infinite with complexity controlled by a finite VC dimension is limited to the case of two treatment arms. Extending this analysis to more treatment arms complicates the coarsening step that must be undertaken, and the link to hyperplane arrangements needs to be properly modified. Whether information about treatment response yielded during the coarsening step can be incorporated (it is currently discarded) to improve the performance of the EXP4.P algorithm is unclear, and is pertinent even in the case of two treatment arms. Finally, our analysis of the EXP4.P algorithm does not consider budget or capacity constraints. Such constraints are important in many public policy applications.

Proofs

Proof of Theorem 1 Our proof closely follows the proof given in Beygelzimer et al. (2011) except that our proof explicitly introduces the upper bound of the outcome variable M . Throughout, we subscript the expectation operator by time, to denote the conditional expectation that is contingent upon the information set possessed by the social planner in a given period; we subscript the expectation operator by i and time, to denote the probability distribution over treatment arms that is induced by f^i in a given period. We otherwise maintain the notation defined in the main text, additionally defining $s_t^i = \sum_{k=1}^K y_t(k) \cdot f_k^i(\mathbf{x}_t)$ and $S_t^i = \sum_{\ell=1}^t s_{\ell}^i$. The provisos of the theorem dictate that \mathfrak{F} includes a randomising expert (ascribing equal probability to each

¹³ We conduct further experiments as to the effect of non-linearity of conditional average treatment effects on the performance of the EXP4.P algorithm supporting this statement, the results of which are available upon request.

treatment arm), and $T \geq \omega^2 \vee [4K \cdot \ln(N)]$. These requirements translate directly, to $0 \leq \beta \leq 1/M$ and $0 \leq \gamma < 1/2$.

First, we establish a large deviation inequality for \tilde{S}_T^i relative to S_T^i and, in particular, for a best performing expert. Let E_t denote the conditional expectation given $(\mathcal{I}_t, \mathbf{x}_t)$ and consider

$$E_t(\exp(\beta \cdot s_t^i - \beta \cdot \tilde{s}_t^i)) = E_t\left(\exp\left(\sum_{k=1}^K (\beta \cdot y_t(k) - \beta \cdot \tilde{y}_t(k)) f_k^i(\mathbf{x}_t)\right)\right) \tag{A.1}$$

$$\leq E_t\left(\sum_{k=1}^K \exp(\beta \cdot y_t(k) - \beta \cdot \tilde{y}_t(k)) \cdot f_k^i(\mathbf{x}_t)\right) \tag{A.2}$$

$$= \sum_{k=1}^K E_t(\exp(\beta \cdot y_t(k) - \beta \cdot [\beta \cdot M^2 + y_t(k) \cdot 1(k_t = k)]/p_t(k))) \cdot f_k^i(\mathbf{x}_t), \tag{A.3}$$

where Eq. (A.2) makes use of Jensen’s Inequality. Since $\beta \cdot y_t(k) \leq 1$ by Assumption 2 together with our requirement upon the magnitude of β as assumed in Theorem 1, an exponential bound $\exp(z) \leq 1 + z + z^2$ whenever $z \leq 1$ implies

$$E_t(\exp(\beta \cdot y_t(k) - \beta \cdot \tilde{y}_t(k))) = E_t(\exp(\beta \cdot y_t(k) - \beta \cdot y_t(k) \cdot 1(k_t = k)/p_t(k))) \cdot \exp(p_t(k)/[\beta \cdot M]^2) \tag{A.4}$$

$$\leq E_t(1 + [\beta \cdot y_t(k) - \beta \cdot y_t(k) \cdot 1(k_t = k)]/p_t(k))^2 \cdot \exp(p_t(k)/[\beta \cdot M]^2) \tag{A.5}$$

$$\leq [1 + \beta^2 \cdot E_t(y_t(k)^2)/p_t(k)] \cdot \exp(p_t(k)/[\beta \cdot M]^2) \tag{A.6}$$

$$\leq 1. \tag{A.7}$$

We note that Eqs. (A.4)–(A.5) relies on the fact that the realised treatment is drawn according to \mathbf{p}_t , with the expectation of the linear term in the exponential bound then equals to zero, and Eqs. (A.6)–(A.7) makes use of the alternative exponential bound that is $1 + z \leq \exp(z)$, and Assumption 2 together with our requirement upon the magnitude of β being bounded by $1/M$ as implied by the assumptions of Theorem 1. It then follows that

$$E_p\left(\exp(\beta \cdot S_T^i - \beta \cdot \tilde{S}_T^i)\right) \leq \prod_{t=1}^T E_t\left(\exp(\beta \cdot s_t^i - \beta \cdot \tilde{s}_t^i)\right) \leq 1. \tag{A.8}$$

By Markov’s inequality and Eq. (A.8), we obtain

$$\begin{aligned}
 &P\left(S_T^i - \tilde{S}_T^i \geq [\beta \cdot K \cdot M^2] \cdot T\right) \\
 &= P\left(\exp(\beta \cdot S_T^i - \beta \cdot \tilde{S}_T^i) \geq \exp\left(\beta \cdot [\beta \cdot K \cdot M^2] \cdot T\right)\right) \tag{A.9}
 \end{aligned}$$

$$\leq E_P\left(\exp(\beta \cdot S_T^i - \beta \cdot \tilde{S}_T^i)\right) \cdot \exp\left(-\beta \cdot [\beta \cdot K \cdot M^2] \cdot T\right) \tag{A.10}$$

$$\leq \exp(-\beta \cdot A \cdot T), \tag{A.11}$$

where $A = \beta KM^2$. By the choice of β as per Theorem 1, we obtain

$$P\left(S_T^i - \tilde{S}_T^i \geq A \cdot T\right) \leq \delta/N \tag{A.12}$$

Note that Eq. (A.12) holds for an arbitrary expert, building on which we can obtain a large deviation inequality of $\max_i(S_T^i)$ relative to $\max_i(\tilde{S}_T^i)$. Let $i^* \doteq \arg \max_i S_T^i$ and, note that

$$\begin{aligned}
 &\left\{ \max_i(S_T^i) \geq \max_i(\tilde{S}_T^i) + A \cdot T \right\} \subset \left\{ S_T^{i^*} \geq (\tilde{S}_T^{i^*}) + A \cdot T \right\} \\
 &\subset \left\{ \exists i S_T^i \geq \tilde{S}_T^i + A \cdot T \right\} = \bigcup_{i=1}^N \left\{ S_T^i \geq \tilde{S}_T^i + A \cdot T \right\}.
 \end{aligned}$$

Hence, by applying the union bound, we obtain

$$P\left(\max_i(S_T^i) - A \cdot T \geq \max_i(\tilde{S}_T^i)\right) \leq \sum_{i=1}^N P\left(S_T^i - \tilde{S}_T^i \geq A \cdot T\right) \leq \delta \tag{A.13}$$

Next, we establish a high probability bound for the welfare regret using Eq. (A.13). Define

$$C_t \doteq \begin{cases} N & \text{for } t = 0 \\ \sum_{i=1}^N \exp(\eta \cdot \tilde{S}_t^i) & \text{for } t = 1, \dots, T \end{cases} \tag{A.14}$$

which we can relate to a (prospective) policy weight using $q_t^i = \exp(\eta \cdot \tilde{S}_{t-1}^i)/C_{t-1}$, where we interpret $\tilde{S}_0^i = 0$ for $i = 1, \dots, N$. Hence,

$$C_t/C_{t-1} = \sum_{i=1}^N \exp(\eta \cdot \tilde{s}_t^i) \cdot \exp(\eta \cdot \tilde{S}_{t-1}^i)/C_{t-1} \tag{A.15}$$

$$= \sum_{i=1}^N \exp(\eta \cdot \tilde{s}_t^i) \cdot q_t^i \tag{A.16}$$

$$\leq 1 + \eta \cdot \underbrace{\sum_{i=1}^N \tilde{s}_t^i \cdot q_t^i}_{B_{1t}} + \eta^2 \cdot \underbrace{\sum_{i=1}^N [\tilde{s}_t^i]^2 \cdot q_t^i}_{B_{2t}}. \tag{A.17}$$

We note that Eqs. (A.16)–(A.17) relies on the exponential bound previously used in Eq. (A.5), which is applicable here by Assumption 2 together with the given specifications of β , γ and η in the current theorem. The two labelled terms in Eq. (A.17) can then be bounded as follows.

$$B_{1t} = \sum_{i=1}^N \sum_{k=1}^K \tilde{y}_i(k) \cdot f_k^i(\mathbf{x}_t) \cdot q_t^i \tag{A.18}$$

$$= \sum_{k=1}^K [\beta \cdot M^2 + y_t(k_t) \cdot 1(k_t = k)] / p_t(k) \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i \tag{A.19}$$

$$\leq 1/[1 - \gamma] \cdot \sum_{k=1}^K [\beta \cdot M^2 + y_t(k_t) \cdot 1(k_t = k)] \tag{A.20}$$

$$= (A + y_t(k_t)) / [1 - \gamma] \tag{A.21}$$

We note that Eqs. (A.19)–(A.20) is by the construction of \mathbf{p}_t .¹⁴

$$B_{2t} = \sum_{i=1}^N [\tilde{y}_i]^i f^i(\mathbf{x}_t)]^2 \cdot q_t^i \tag{A.22}$$

$$\leq \sum_{i=1}^N \sum_{k=1}^K \tilde{y}_i(k)^2 \cdot f_k^i(\mathbf{x}_t) \cdot q_t^i \tag{A.23}$$

$$= \sum_{k=1}^K [[\beta \cdot M^2 + y_t(k_t) \cdot 1(k_t = k)] / p_t(k)]^2 \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i \tag{A.24}$$

$$\leq 2 \sum_{k=1}^K [[\beta \cdot M^2]^2 + y_t(k_t)^2 \cdot 1(k_t = k)] / p_t(k) \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i / p_t(k) \tag{A.25}$$

$$\leq 2/[1 - \gamma] \cdot \sum_{k=1}^K [[\beta \cdot M^2]^2 + y_t(k_t)^2 \cdot 1(k_t = k)] / p_t(k) \tag{A.26}$$

$$\leq 2/(1 - \gamma) \cdot [A^2/\gamma + M \cdot y_t(k_t) / p_t(k_t)] \tag{A.27}$$

We note that Eqs. (A.22)–(A.23) makes use of the Cauchy–Schwarz Inequality and the fact that f is constrained to the unit simplex¹⁵; Eqs. (A.24)–(A.25) makes

¹⁴ Specifically, $p_t(k) = [1 - \gamma] \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i + \gamma/K$, or $p_t(k) \geq [1 - \gamma] \cdot \sum_{i=1}^N f_k^i(\mathbf{x}_t) \cdot q_t^i$. Similarly, $p_t(k) \geq \gamma/K$.

¹⁵ We use the expectation version of the Cauchy–Schwarz Inequality here, with f^i playing the role of probability weights.

use of the quadratic inequality $[a + b]^2 \leq 2[a^2 + b^2]$; and Eq. (A.25)–(A.26) and Eqs. (A.26)–(A.27) are by the construction of \mathbf{p}_t and Assumption 2. Applying a logarithmic telescope to Eq. (A.15), we obtain

$$\begin{aligned} & \ln(C_T/C_0) \\ &= \sum_{t=1}^T \ln(C_t/C_{t-1}) \end{aligned} \tag{A.28}$$

$$\leq \sum_{t=1}^T \ln(1 + \eta \cdot B_{1t} + \eta^2 \cdot B_{2t}) \tag{A.29}$$

$$\leq \sum_{t=1}^T [\eta \cdot B_{1t} + \eta^2 \cdot B_{2t}] \tag{A.30}$$

$$\begin{aligned} & \leq \eta \cdot T \cdot [1 + 2\eta \cdot A/\gamma] \cdot A/[1 - \gamma] \\ & \quad + \eta \cdot \sum_{t=1}^T [1 + 2\eta \cdot M/p_t(k_t)] \cdot y_t(k_t)/[1 - \gamma] \end{aligned} \tag{A.31}$$

$$= \eta/[1 - \gamma] \cdot \underbrace{[1 + \beta \cdot M] \cdot A \cdot T}_{A'} + \sum_{t=1}^T y_t(k_t) + \gamma/K \cdot \sum_{t=1}^T y_t(k_t)/p_t(k_t) \tag{A.32}$$

We note that Eq. (A.28)–(A.29) implements Eq. (A.17); we note that Eqs. (A.29)–(A.30) makes use of the logarithmic inequality $\ln(1 + z) \leq z$; we note that Eqs. (A.30)–(A.31) implements Eqs. (A.21) and (A.27); and we note that Eqs. (A.31)–(A.32) is by rearrangement and by the specifications of β, γ and η as per Theorem 1. Rewriting Eq. (A.32) via multiplication by the inverse of the common term, we obtain the following upper bound.

$$\begin{aligned} & [1 - \gamma]/\eta \cdot \ln(C_T/C_0) \\ & \leq A' + \sum_{t=1}^T y_t(k_t) + \gamma/K \cdot \sum_{t=1}^T y_t(k_t)/p_t(k_t) \end{aligned} \tag{A.33}$$

$$\leq A' + \sum_{t=1}^T y_t(k_t) + \gamma \cdot \sum_{t=1}^T \sum_{k=1}^K \tilde{y}_t(k)/K \tag{A.34}$$

$$= A' + \sum_{t=1}^T y_t(k_t) + \gamma \cdot \tilde{S}_T^{\text{random}} \tag{A.35}$$

$$\leq A' + \sum_{t=1}^T y_t(k_t) + \gamma \cdot \max_i(\tilde{S}_T^i), \tag{A.36}$$

where $\tilde{S}_T^{\text{random}}$ is the cumulative score of the pure randomizing expert that enacts each possible action with equal probability.

We also obtain the following lower bound.

$$\begin{aligned}
 & [1 - \gamma]/\eta \cdot \ln(C_T/C_0) \\
 & \geq [1 - \gamma] \cdot \max_i(\tilde{S}_T^i) - [1 - \gamma]/\eta \cdot \ln(C_0)
 \end{aligned} \tag{A.37}$$

$$\geq [1 - \gamma] \cdot \max_i(\tilde{S}_T^i) - [1 - \gamma]/\eta \cdot \ln(N), \tag{A.38}$$

where Eq. (A.37) follows since $\sum_{i=1}^N \exp(\eta \tilde{S}_T^i) \geq \exp(\eta \max_i(\tilde{S}_T^i))$ and Eq. (A.38) follows by the uniformity of the initial weights as specified in Theorem 1. Combining Eqs. (A.36) and (A.38) yields

$$[1 - 2\gamma] \cdot \max_i(\tilde{S}_T^i) - \sum_{t=1}^T y_t(k_t) \leq A' + \underbrace{[1 - \gamma]/\eta \cdot \ln(N)}_{A''}. \tag{A.39}$$

Each of the two terms in the right-hand side of Eq. (A.39) can be bounded as follows:

$$A' \leq 2 \cdot A \cdot T = 2M \cdot \sqrt{K \cdot T \cdot \ln(N/\delta)}, \tag{A.40}$$

where this inequality follows by $\beta \leq 1/M$ as implied by the assumptions of Theorem 1.

$$A'' = 2M \cdot \sqrt{K \cdot \ln(N)} \cdot [\sqrt{T} - \sqrt{K \cdot \ln(N)}] \leq 2M \cdot \sqrt{K \cdot T \cdot \ln(N/\delta)}, \tag{A.41}$$

where we note that $0 < \delta < 1$. Applying Eq. (A.13)—more precisely, its complement—to the left-hand side of Eq. (A.39) and rearranging, we determine that

$$\begin{aligned}
 \hat{R}_T = \max_i(S_T^i) - \sum_{i=1}^T y_t(k_t) & \leq 2\gamma \cdot \max_i S_T^i + [1 - 2\gamma] \\
 & \cdot A \cdot T + A' + A''
 \end{aligned} \tag{A.42}$$

holds with probability at least $1 - \delta$. Utilising the results of Eqs. (A.40) and (A.41), making use of the support condition $\max_i(S_T^i) \leq M \cdot T$, and noting $(1 - 2\gamma) \cdot A \cdot T \leq M\sqrt{K \cdot T \cdot \ln(N/\delta)}$, we obtain the conclusion of Theorem 1. \square

Proof of Theorem 2 Let $f^* \in \arg \sup_{f \in \mathfrak{F}} W_T(f)$, and $g^* \in \mathfrak{G}$ be a best expert in the run phase in the sense of $g^* \in \arg \max_{g \in \mathfrak{G}} \sum_{t=[\tau]+1}^T g(\mathbf{x}_t)^\top \mathbf{y}_t$. For the class of experts with non-randomised assignments, Sauer’s lemma bounds the cardinality of $\{(f(x_1), \dots, f(x_T)) : f \in \mathfrak{F}\}$ from above by $\left(\frac{eT}{D}\right)^D$. Consider the following decomposition of R_T :

$$\begin{aligned}
 R_T &= \sup_{f \in \mathfrak{F}} E_P \left(\sum_{t=1}^T f(\mathbf{x}_t)^\top \mathbf{y}_t \right) - \sum_{t=1}^T y_t(k_t) \\
 &= \underbrace{[\tau] E_P(f^*(\mathbf{x}_t)^\top \mathbf{y}_t)}_{(i)} - \underbrace{\left[\sum_{t=[\tau]+1}^T f^*(\mathbf{x}_t)^\top \mathbf{y}_t - E_P \left(\sum_{t=[\tau]+1}^T f^*(\mathbf{x}_t)^\top \mathbf{y}_t \right) \right]}_{(ii)}
 \end{aligned} \tag{A.43}$$

$$\begin{aligned}
 &+ \underbrace{\sum_{t=[\tau]+1}^T (f^*(\mathbf{x}_t) - g^*(\mathbf{x}_t))^\top \mathbf{y}_t}_{(iii)} + \underbrace{\sum_{t=[\tau]+1}^T (g^*(\mathbf{x}_t)^\top \mathbf{y}_t - y_t(k_t))}_{(iv)}.
 \end{aligned} \tag{A.44}$$

Here, term (i) corresponds to the error accumulated during the coarsening phase for the first $[\tau]$ periods. The sum of terms (ii) and (iii) capture the error accumulated by using \mathfrak{G} instead of \mathfrak{F} . Term (iv) corresponds to the regret associated with the run phase when the EXP4.P-F variant is implemented for the next $T - [\tau]$ periods with the coarsened class \mathfrak{G} .

We bound each of (i) - (iv) terms separately. By Assumption 2, (i) can be upper bounded by $[\tau]M$. To bound (ii), we invoke Chernoff’s inequality,

$$\begin{aligned}
 P \left(\sum_{t=[\tau]+1}^T f^*(\mathbf{x}_t)^\top \mathbf{y}_t - E_P \left(\sum_{t=[\tau]+1}^T f^*(\mathbf{x}_t)^\top \mathbf{y}_t \right) \leq M \sqrt{(T - [\tau])/2 \cdot \ln(3/\delta)} \right) \\
 \geq 1 - \delta/3.
 \end{aligned} \tag{A.45}$$

Note that (iv) corresponds to the empirical regret of the EXP4.P-F algorithm for $T - [\tau]$ periods under the coarse class \mathfrak{G} whose cardinality is bounded by $|\mathfrak{G}| \leq \left(\frac{e[\tau]}{D}\right)^D$ by Sauer’s lemma. Hence, with probability at least $1 - \delta/3$, it holds

$$(iv) \leq c_1 M \cdot \sqrt{[T - [\tau]] \cdot [2D \cdot \ln([\tau] \cdot e/D) + \ln(3/\delta)]}, \tag{A.46}$$

where c_1 is a universal constant.

Given a sequence of subjects, note that (iii) can be bounded by $M \cdot \|f^* - g^*\|$, where

$$\|f^* - g^*\| \doteq \sum_{t=[\tau]+1}^T 1(f^*(\mathbf{x}_t) \neq g^*(\mathbf{x}_t)).$$

Under Assumption 1, this sequence is as likely to occur as any other sequence of subjects, of which there are a total of $T!$ such permutations. Denote a permutation by a one-to-one function $\chi : \{1, \dots, T\} \rightarrow \{1, \dots, T\}$, and consider a uniform distribution over the $T!$ permutations denoted by P_χ . Conditional on a given sequence and the covariates associated with it, we have

$$\begin{aligned}
 P_{\chi} \left(\sum_{t=\lceil \tau \rceil + 1}^T 1(f^{g^*}(\mathbf{x}_{\chi(t)}) \neq g^*(\mathbf{x}_{\chi(t)})) > m \right) \\
 \leq P_{\chi}(\exists f \in \mathfrak{F} \text{ and } h \in \mathfrak{G} \text{ satisfying } \|f - h\| > m \\
 \text{and } f(\mathbf{x}_{\chi(t)}) = h(\mathbf{x}_{\chi(t)}) \text{ for } t \leq \lceil \tau \rceil)
 \end{aligned} \tag{A.47}$$

$$\leq |\mathfrak{F}|^2 \cdot \binom{T - m}{\lceil \tau \rceil} / T! \tag{A.48}$$

$$\leq [e \cdot T/D]^{2D} \cdot [1 - m/T]^{\lceil \tau \rceil} \tag{A.49}$$

$$\leq [e \cdot T/D]^{2D} \cdot \exp(-\lceil \tau \rceil \cdot m/T). \tag{A.50}$$

We note that Eq. (A.47) maintains the requirement that f and h share the same classification of the first $\lceil \tau \rceil$ subjects and disagree more than m times during the run phase, Eqs. (A.47)–(A.48) is the product of an upper bound on the number of possible pairs of functions and an upper bound on the probability that a random permutation produces at least m errors, Eqs. (A.48)–Eq. (A.49) makes use of Sauer’s Lemma, and relies on the properties of binomial coefficients,¹⁶ and Eqs. (A.49)–Eq. (A.50) makes use of the exponential bound, $1 + z \leq \exp(z)$. Since Eq. (A.50) does not depend upon the sequence in which subjects arrive, the same bound applies for the unconditional probability over all possible sequences. By setting

$$m = \frac{T}{\lceil \tau \rceil} \cdot [2D \cdot \ln(e \cdot T/D) + \ln(3/\delta)] \tag{A.51}$$

we determine that the error accumulated by using \mathfrak{G} in place of \mathfrak{F} is bounded from above by m with probability at least $1 - \delta/3$.

Combining Eqs. (A.43), (A.45), (A.46), and (A.51), we obtain that for a universal constant $c_2 > 0$,

$$R_T/M \leq \lceil \tau \rceil + c_2 \sqrt{[T - \lceil \tau \rceil] \cdot [2D \cdot \ln(\lceil \tau \rceil \cdot e/D) + \ln(3/\delta)]} + m \tag{A.52}$$

$$\leq \lceil \tau \rceil + c_2 \sqrt{T \cdot [2D \cdot \ln(T \cdot e/D) + \ln(3/\delta)] + T \cdot \ln(3/\delta)} + m \tag{A.53}$$

$$\leq \lceil \tau \rceil + c_2 \sqrt{m \cdot \lceil \tau \rceil + T \cdot \ln(3/\delta)} + m \tag{A.54}$$

with probability at least $1 - \delta$. We can then minimise the right-hand side of Eq. (A.54) by setting τ such that $m = \tau^2/\lceil \tau \rceil$. The statement of Theorem 2 then follows. □

¹⁶ Specifically, $(T - m)/T \geq (T - m - z)/(T - z)$.

Proof of Theorem 3 We follow the proof of Theorem 2, modifying those parts that make reference to \mathfrak{F} . Specifically, we exploit the assumption that $\mathfrak{F} = \mathfrak{L}_J$ to obtain a tighter bound on its cardinality by invoking Harding(t, J) rather than Sauer's lemma. We omit the details for brevity. \square

Acknowledgements Toru Kitagawa is honoured to accept the 2023 Nakahara prize by the Japanese Economic Association, and thanks the selection committee for their consideration. The authors thank the editor and a referee for beneficial comments, and gratefully acknowledge financial support from ERC Grant 715940, the ESRC Centre for Microdata Methods and Practice (CeMMAP) (RES-589-28-0001).

Data availability The data set and R codes to perform the numerical analysis of this article are available for downloading from the corresponding author's webpage: <https://sites.google.com/brown.edu/torukitagawa/home/publication>.

References

- Abadie, A., Angrist, J., & Imbens, G. (2002). Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings. *Econometrica*, 70, 91–117.
- Abbasi-Yadkori, Y., Pál, D., & Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24, 2312–2320.
- Adusumilli, K. (2021). Risk and optimal policies in bandit experiments. arXiv preprint [arXiv:2112.06363](https://arxiv.org/abs/2112.06363).
- Ariu, K., Kato, M., Komiya, J., McAlinn, K., & Qin, C. (2021). Policy choice and best arm identification: Asymptotic analysis of exploration sampling. arXiv preprint [arXiv:2109.08229](https://arxiv.org/abs/2109.08229).
- Athey, S., Byambadalai, U., Hadad, V., Krishnamurthy, S. K., Leung, W., & Williams, J. J. (2022). Contextual bandits in a survey experiment on charitable giving: within-experiment outcomes versus policy learning. arXiv.
- Athey, S., Chetty, R., Imbens, G. W., & Kang, H. (2019). *The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely*. Working Paper 26463, National Bureau of Economic Research.
- Athey, S., & Wager, S. (2021). Policy learning with observational data. *Econometrica*, 89, 133–161.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3, 397–422.
- Auer, P., Cesa-Bianchi, N., Freund, Y., & Schapire, R. (2002). The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32, 48–77.
- Avis, D., & Fukuda, K. (1996). Reverse search for enumeration. *Discrete Applied Mathematics*, 65, 21–46.
- Bégin, L., Germain, P., Lavoie, F., & Roy, J.-F. (2014). PAC-Bayesian theory for transductive learning. *Artificial Intelligence and Statistics*, 105–113.
- Beygelzimer, A., Langford, J., Li, L., Reyzin, L., & Schapire, R. (2011). Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, JMLR Workshop and Conference Proceedings (pp. 19–26).
- Bloom, H. S., Orr, L. L., Bell, S. H., Cave, G., Doolittle, F., Lin, W., & Bos, J. M. (1997). The benefits and costs of JTPA Title II-A programs: Key findings from the National Job Training Partnership Act study. *Journal of human resources*, 549–576.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32.
- Buck, R. C. (1943). Partition of space. *The American Mathematical Monthly*, 50, 541–544.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.
- Chen, J. (2023). Synthetic control as online linear regression. *Econometrica*, 91, 465–491.
- Dehejia, R. H. (2005). Program evaluation as a decision problem. *Journal of Econometrics*, 125, 141–173.
- Dimakopoulou, M., Zhou, Z., Athey, S., & Imbens, G. (2017). *Estimation considerations in contextual bandits*. arXiv preprint [arXiv:1711.07077](https://arxiv.org/abs/1711.07077).

- Gu, J., & Koenker, R. (2022). Nonparametric maximum likelihood methods for binary response models with random coefficients. *Journal of the American Statistical Association*, 117, 732–751.
- Han, S. (2023). Optimal dynamic treatment regimes and partial welfare ordering. *Journal of the American Statistical Association*, 1–11.
- Harding, E. F. (1967). The number of partitions of a set of N points in k dimensions induced by hyperplanes. *Proceedings of the Edinburgh mathematical society*, 15, 285–289.
- Heckman, J., Ichimura, H., & Todd, P. (1997). Matching as an econometric evaluation estimator: Evidence on its performance applied to the JTPA program, part I. *Theory and Methods, Review of Economic studies*, 64, 605–654.
- Ida, T., Ishihara, T., Ito, K., Kido, D., Kitagawa, T., Sakaguchi, S., & Sasaki, S. (2024). *Dynamic targeting: Experimental evidence from energy rebate programs*. NBER working paper.
- Kasy, M., & Sautmann, A. (2021). Adaptive treatment assignment in experiments for policy choice. *Econometrica*, 89, 113–132.
- Kitagawa, T., Lopez, H., & Rowley, J. (2022). *Stochastic treatment choice with empirical welfare updating*. arXiv preprint [arXiv:2211.01537](https://arxiv.org/abs/2211.01537).
- Kitagawa, T., Sakaguchi, S., & Tetenov, A. (2023). *Constrained classification and policy learning*. arXiv preprint [arXiv:2106.12886](https://arxiv.org/abs/2106.12886).
- Kitagawa, T., Wang, W., & Xu, M. (2024). *Policy choice in time-series by empirical welfare maximization*. arXiv preprint.
- Kitagawa, T., & Tetenov, A. (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica*, 86, 591–616.
- Kitagawa, T., & Tetenov, A. (2021). Equality-minded treatment choice. *Journal of Business & Economic Statistics*, 39, 561–574.
- Ko, R., Uetake, K., Yata, K., & Okada, R. (2022). *When to target customers? Retention management using dynamic off-policy policy learning*. Available at SSRN.
- Kock, A. B., Preinerstorfer, D., & Veliyev, B. (2022). Functional sequential treatment allocation. *Journal of the American Statistical Association*, 117, 1311–1323.
- Kuang, X., & Wager, S. (2023). Weak signal asymptotics for sequentially randomized experiments, *Management Science*.
- Lattimore, T., & Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Littlestone, N., & Warmuth, M. (1994). The weighted majority algorithm. *Information and Computation*, 108, 212–261.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72, 1221–1246.
- Mbakop, E., & Tabord-Meehan, M. (2021). Model selection for treatment choice: Penalized welfare maximization. *Econometrica*, 89, 825–848.
- McAllester, D. A. (2003). PAC-Bayesian stochastic model selection. *Machine Learning*, 51, 5–21.
- Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 331–355.
- Pinkse, C. (1993). On the computation of semiparametric estimates in limited dependent variable models. *Journal of Econometrics*, 58, 185–205.
- Qin, C., & Russo, D. (2024). *Optimizing adaptive experiments: a unified approach to regret minimization and best-arm identification*. arXiv.
- Rada, M., & Černý, M. (2018). A new algorithm for enumeration of cells of hyperplane arrangements and a comparison with Avis and Fukuda's reverse search. *SIAM Journal on Discrete Mathematics*, 32, 455–473.
- Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period-application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7, 1393–1512.
- Rosen, A. M., & Ura, T. (2019). *Finite sample inference for the maximum score estimand*. arXiv preprint [arXiv:1903.01511](https://arxiv.org/abs/1903.01511).
- Russo, D., & Roy, B. V. (2016). Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39, 1221–1243.
- Sauer, N. (1972). On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13, 145–147.
- Shelah, S. (1972). A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41, 247–261.

- Sleumer, N. (1998). Output-sensitive cell enumeration in hyperplane arrangements. In *Algorithm Theory-SWAT'98: 6th Scandinavian Workshop on Algorithm Theory Stockholm, Sweden, July 8–10, 1998 Proceedings 6* (pp. 300–309). Springer.
- Vapnik, V. N. (1998). *Statistical learning theory*. John Wiley & Sons.
- Vapnik, V. N., & Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16, 264–280.
- Viviano, D., & Bradic, J. (2023). Synthetic learner: model-free inference on treatments over time. *Journal of Econometrics*, 234, 691–713.
- Vovk, V. (1990). Aggregating strategies. In *Proceedings of the 3rd Annual Workshop on Computational Learning Theory* (pp. 372–383).
- Zhang, Y., Laber, E. B., Davidian, M., & Tsiatis, A. A. (2018). Interpretable dynamic treatment regimes. *Journal of the American Statistical Association*, 113, 1541–1549.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.