

## RESEARCH ARTICLE

# Estimating Average Treatment Effects With Support Vector Machines

Alexander Tarr<sup>1</sup> | Kosuke Imai<sup>1,2,3</sup> <sup>1</sup>Institute for Quantitative Social Science, Harvard University, Massachusetts, USA | <sup>2</sup>Department of Government, Harvard University, Massachusetts, USA |<sup>3</sup>Department of Statistics, Harvard University, Massachusetts, USA**Correspondence:** Kosuke Imai ([imai@harvard.edu](mailto:imai@harvard.edu))**Received:** 20 May 2024 | **Revised:** 27 November 2024 | **Accepted:** 10 January 2025**Funding:** Imai thanks the Sloan Foundation (# 2020–13946) for financial support.**Keywords:** causal inference | covariate balance | matching | subset selection | weighting

## ABSTRACT

Support vector machine (SVM) is one of the most popular classification algorithms in the machine learning literature. We demonstrate that SVM can be used to balance covariates and estimate average causal effects under the unconfoundedness assumption. Specifically, we adapt the SVM classifier as a kernel-based weighting procedure that minimizes the maximum mean discrepancy between the treatment and control groups while simultaneously maximizing effective sample size. We also show that SVM is a continuous relaxation of the quadratic integer program for computing the largest balanced subset, establishing its direct relation to the cardinality matching method. Another important feature of SVM is that the regularization parameter controls the trade-off between covariate balance and effective sample size. As a result, the existing SVM path algorithm can be used to compute the balance-sample size frontier. We characterize the bias of causal effect estimation arising from this trade-off, connecting the proposed SVM procedure to the existing kernel balancing methods. Finally, we conduct simulation and empirical studies to evaluate the performance of the proposed methodology and find that SVM is competitive with the state-of-the-art covariate balancing methods.

## 1 | Introduction

Estimating causal effects in an observational study is complicated by the lack of randomization in treatment assignment, which may lead to confounding bias. The standard approach is to weight observations such that the empirical distribution of observed covariates is similar between the treatment and control groups [1–4]. Researchers then estimate the causal effects using the weighted sample while assuming the absence of unobserved confounders. Recently, a large number of weighting methods have been proposed to directly optimize covariate balance for causal effect estimation [5–16].

This paper provides new insight into this fast-growing literature on covariate balancing by demonstrating that support vector machine (SVM), which is one of the most popular classification algorithms in machine learning [17, 18], can be used to balance covariates and estimate the average treatment effect under the standard unconfoundedness assumption. We adopt the dual coefficients for the soft-margin SVM classifier as kernel balancing weights that minimize the maximum mean discrepancy [19] between the treatment and control groups while simultaneously maximizing effective sample size. The resulting weights are bounded, leading to stable causal effect estimation. As SVM has been extensively studied and widely used, we can

exploit its well-known theoretical properties and highly optimized implementations.

All matching and weighting methods face the same trade-off between effective sample size and covariate balance with a better balance typically leading to a smaller sample size. We show that SVM directly addresses this fundamental trade-off. Specifically, the dual optimization problem for SVM computes a set of balancing weights as the dual coefficients while yielding the support vectors that comprise the largest balanced subset. In addition, the regularization parameter of SVM controls the trade-off between sample size and covariate balance. We use existing path algorithms [20, 21] to efficiently characterize the balance-sample size frontier [22]. We analyze how this trade-off between sample size and covariate balance affects causal effect estimation.

We are not the first to realize the connection between SVM and covariate balancing. Ratkovic [23] points out that the hinge-loss function of the primal optimization problem for SVM has a first-order condition, which leads to balanced covariate *sums* amongst the support vectors. Instead, we show that the dual form of the SVM optimization problem leads to the covariate *mean* balance. In addition, Ghosh [24] notes the relationship between the SVM margin and the region of covariate overlap. The author argues that the support vectors correspond to observations lying in the intersection of the convex hulls for the treated and control samples [25]. In contrast, we show that the SVM dual coefficients can be used to obtain weights for causal effect estimation. Furthermore, neither of these previous works studies the relationship between the regularization parameter of SVM and the trade-off between covariate balance and effective sample size.

The proposed methodology is also related to several covariate balancing methods. First, we establish that the SVM dual optimization problem can be seen as a continuous relaxation of the quadratic integer program for computing the largest balanced subset, which is related to cardinality matching [26]. Second, weighting with SVM can be interpreted as a kernel-based covariate balancing method. Several researchers have recently developed weighting methods to balance functions in a reproducing kernel Hilbert space (RKHS) [11, 13, 14]. SVM shares the advantage of these methods that it can balance a general class of functions and easily accommodate non-linearity and non-additivity in the conditional expectation functions for the outcomes. In particular, we show that SVM fits into the kernel optimal matching framework [14]. Unlike these covariate balancing methods, however, we can exploit the existing path algorithms for SVM to compute the set of solutions over the entire regularization path with comparable complexity to computing a single solution. Finally, we also show that a variant of soft-margin SVM is related to stable balancing weights [7].

The rest of the paper is structured as follows. In Section 2, we present our methodological results. In Section 3, we conduct simulation studies to compare the performance of SVM with that of the related covariate balancing methods. Lastly, in Section 4, we apply SVM to the data from the right heart catheterization study [27].

## 2 | Methodology

In this section, we establish several properties of SVM as a covariate balancing method. First, the SVM dual can be viewed as a regularized optimization problem that minimizes the maximum mean discrepancy (MMD). Second, we show that SVM is a relaxation of the largest balanced subset selection problem. Third, the regularization path algorithm for SVM can be used to obtain the balance-sample size frontier. Lastly, we discuss how to use SVM for causal effect estimation and compare SVM to existing kernel balancing methods.

### 2.1 | Setup and Assumptions

Suppose that we observe a simple random sample of  $N$  units from a super-population of interest,  $\mathcal{P}$ . Denote the observed data by  $D = \{\mathbf{X}_i, Y_i, T_i\}_{i=1}^N$  where  $\mathbf{X}_i \in \mathcal{X}$  represents a  $p$ -dimensional vector of covariates,  $Y_i$  is the outcome variable, and  $T_i$  is a binary treatment assignment variable that is equal to 1 if unit  $i$  is treated and 0 otherwise. We define the index sets for the treatment and control groups as  $\mathcal{T} = \{i : T_i = 1\}$  and  $\mathcal{C} = \{i : T_i = 0\}$  with the group sizes equal to  $n_T = |\mathcal{T}|$  and  $n_C = |\mathcal{C}|$ , respectively. Finally, we define the observed outcome as  $Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0)$  where  $Y_i(1)$  and  $Y_i(0)$  are the potential outcomes under treatment and control conditions, respectively.

This notation implies the Stable Unit Treatment Value Assumption (SUTVA)—no interference between units and the same version of the treatment [28]. Furthermore, we maintain the following standard identification assumptions throughout this paper.

**Assumption 1.** (Unconfoundedness). The potential outcomes  $\{Y_i(1), Y_i(0)\}$  are independent of the treatment assignments  $T_i$  conditional on the covariates  $\mathbf{X}_i$ . That is, for all  $\mathbf{x} \in \mathcal{X}$ , we have

$$\{Y_i(1), Y_i(0)\} \perp\!\!\!\perp T_i | \mathbf{X}_i = \mathbf{x}$$

**Assumption 2.** (Overlap). For all  $\mathbf{x} \in \mathcal{X}$ , the propensity score  $e(\mathbf{x}) = \Pr(T_i = 1 | \mathbf{X}_i = \mathbf{x})$  is bounded away from 0 and 1, i.e.,

$$0 < e(\mathbf{x}) < 1$$

Following the convention of classification methods, it is convenient to define the following transformed treatment variable, which is equal to either  $-1$  or  $1$ ,

$$W_i = 2T_i - 1 \in \{-1, 1\} \quad (1)$$

In addition, for  $t = 0, 1$ , we define the conditional expectation functions, disturbances, and conditional variance functions as,

$$\begin{aligned} \mathbb{E}(Y_i(t) | \mathbf{X}_i) &= f_t(\mathbf{X}_i), \quad \epsilon_i(t) = Y_i(t) - f_t(\mathbf{X}_i), \\ \sigma_i^2(\mathbf{X}_i) &= \mathbb{V}(Y_i(t) | \mathbf{X}_i) \end{aligned}$$

Thus, we have  $\mathbb{E}(\epsilon_i(t) | \mathbf{X}_i) = 0$  for  $t = 0, 1$ . Lastly, let  $\mathcal{H}_K$  denote a reproducing kernel Hilbert space (RKHS) with norm  $\|\cdot\|_{\mathcal{H}_K}$  and kernel  $K(\mathbf{X}_i, \mathbf{X}_j) = \langle \phi(\mathbf{X}_i), \phi(\mathbf{X}_j) \rangle_{\mathcal{H}_K}$ , where  $\phi : \mathbb{R}^p \mapsto \mathcal{H}_K$  is a feature mapping of the covariates to the RKHS.

## 2.2 | Support Vector Machines

Support vector machines (SVMs) are a widely-used methodology for two-class classification problems [17, 18]. SVM aims to compute a separating hyperplane of the form,

$$f(\mathbf{X}_i) = \boldsymbol{\beta}^\top \phi(\mathbf{X}_i) + \beta_0 \quad (2)$$

where  $\boldsymbol{\beta} \in \mathcal{H}_K$  is the normal vector for the hyperplane and  $\beta_0$  is the offset, and classification is done based on which side of the hyperplane  $\mathbf{X}_i$  lies on.

In this paper, we use SVM for the classification of treatment status. For non-separable data,  $\boldsymbol{\beta}$  and  $\beta_0$  are computed according to the following soft-margin SVM problem,

$$\begin{aligned} \min_{\boldsymbol{\beta}, \beta_0, \xi} \quad & \frac{\lambda}{2} \|\boldsymbol{\beta}\|_{\mathcal{H}_K}^2 + \sum_{i=1}^N \xi_i \\ \text{s.t.} \quad & W_i f(\mathbf{X}_i) \geq 1 - \xi_i, \quad i = 1, \dots, N \\ & \xi_i \geq 0, \quad i = 1, \dots, N \end{aligned} \quad (3)$$

where  $\{\xi_i\}_{i=1}^N$  are the so-called slack variables, and  $\lambda$  is a regularization parameter controlling the trade-off between the margin width of the hyperplane and the margin violation of the samples. Note that  $\lambda$  is related to the traditional SVM cost parameter  $C$  via  $\lambda = 1/C$ .

In cases where  $\boldsymbol{\beta}$  is high-dimensional, the dual form of Equation (3) is often preferable to work with. Defining the matrix  $\mathbf{Q}$  with elements  $Q_{ij} = W_i W_j K(\mathbf{X}_i, \mathbf{X}_j)$  and the vector  $\mathbf{W}$  with elements  $W_i$ , the dual form is given by,

$$\begin{aligned} \min_{\boldsymbol{\alpha}} \quad & \frac{1}{2\lambda} \boldsymbol{\alpha}^\top \mathbf{Q} \boldsymbol{\alpha} - \mathbf{1}^\top \boldsymbol{\alpha} \\ \text{s.t.} \quad & \mathbf{W}^\top \boldsymbol{\alpha} = 0 \\ & 0 < \boldsymbol{\alpha} < 1 \end{aligned} \quad (4)$$

where  $\boldsymbol{\alpha}$  are called the dual coefficients,  $\mathbf{1}$  represents a vector of ones, and  $<$  denotes an element-wise inequality.

We begin by providing an intuitive explanation of how the SVM dual given in Equation (4) can be viewed as a covariate balancing procedure. First, note that the quadratic term in the above dual objective function can be written as a weighted measure of covariate discrepancy between the treatment and control groups,

$$\boldsymbol{\alpha}^\top \mathbf{Q} \boldsymbol{\alpha} = \left\| \sum_{i \in \mathcal{T}} \alpha_i \phi(\mathbf{X}_i) - \sum_{i \in \mathcal{C}} \alpha_i \phi(\mathbf{X}_i) \right\|^2 \quad (5)$$

while the constraint  $\mathbf{W}^\top \boldsymbol{\alpha} = 0$  ensures that the sum of weights is identical between the treatment and control groups,

$$\mathbf{W}^\top \boldsymbol{\alpha} = 0 \Leftrightarrow \sum_{i \in \mathcal{T}} \alpha_i = \sum_{i \in \mathcal{C}} \alpha_i$$

Lastly, the second term in the objective,  $\mathbf{1}^\top \boldsymbol{\alpha}$ , is proportional to the sum of weights for each treatment group, since the above constraint implies  $\sum_{i \in \mathcal{T}} \alpha_i = \sum_{i \in \mathcal{C}} \alpha_i = \mathbf{1}^\top \boldsymbol{\alpha} / 2$ . Thus, SVM simultaneously minimizes the covariate discrepancy and maximizes the

effective sample size, which in turn, leads to the minimization of the weighted difference-in-means in the transformed covariate space. In addition, the weights are bounded as represented by the constraint  $0 \leq \alpha_i \leq 1$  for all  $i$ , leading to stable causal effect estimation [29].

The choice of kernel function  $K(\mathbf{X}_i, \mathbf{X}_j)$  and its corresponding feature map  $\phi$  determine the type of covariate balance enforced by SVM, as shown in Equation (5). In this paper, we focus on the linear, polynomial, and radial basis function (RBF) kernels. The linear kernel  $K(\mathbf{X}_i, \mathbf{X}_j) = \mathbf{X}_i^\top \mathbf{X}_j$  corresponds to a feature map  $\phi(\mathbf{X}_i) = \mathbf{X}_i$ , and hence the quadratic term  $\boldsymbol{\alpha}^\top \mathbf{Q} \boldsymbol{\alpha}$  measures the discrepancy in the original covariates. The general form for the degree  $d$  polynomial kernel with scale parameter  $c$  is  $K(\mathbf{X}_i, \mathbf{X}_j) = (\mathbf{X}_i^\top \mathbf{X}_j + c)^d$ . For example, the quadratic kernel with  $d = 2$  leads to a discrepancy measure of the original covariates, their squares, and all pairwise interactions. The final kernel considered in this paper is the RBF kernel with scale parameter  $\gamma$ :  $K(\mathbf{X}_i, \mathbf{X}_j) = \exp(-\gamma \|\mathbf{X}_i - \mathbf{X}_j\|^2)$ . This kernel can be viewed as a generalization of the polynomial kernel in the limit  $d \rightarrow \infty$ .

In addition, the Karush–Kuhn–Tucker (KKT) conditions for soft-margin SVM lead to the following useful characterization for a solution  $\boldsymbol{\alpha}$ :

$$\begin{aligned} W_i f(\mathbf{X}_i) = 1 &\Rightarrow 0 \leq \alpha_i \leq 1 \\ W_i f(\mathbf{X}_i) < 1 &\Rightarrow \alpha_i = 1 \\ W_i f(\mathbf{X}_i) > 1 &\Rightarrow \alpha_i = 0 \end{aligned} \quad (6)$$

The set of units that satisfy  $W_i f(\mathbf{X}_i) > 1$  are easy to classify and receive zero weight. These units occur in regions of little or no overlap and are most difficult to balance. The set of units that satisfy  $W_i f(\mathbf{X}_i) = 1$  are referred to as marginal support vectors, while those that meet  $W_i f(\mathbf{X}_i) < 1$  are the non-marginal support vectors. Collectively, these two sets correspond to the units that the optimal hyperplane has the most difficulty classifying. As we see in the next section, the support vectors comprise a balanced subset.

In sum, the SVM dual problem finds the bounded weights that minimize the covariate discrepancy between the treatment and control groups while simultaneously maximizing the effective sample size. The regularization parameter  $\lambda$  controls which of these two components receives more emphasis. SVM chooses optimal weights such that difficult-to-balance units are given zero weight. Given this intuition, we now establish more formally a connection between SVM, covariate balancing, and causal effect estimation.

## 2.3 | SVM as a Maximum Mean Discrepancy Minimizer

We now show that SVM minimizes the maximum mean discrepancy (MMD) of covariate distribution between the treatment and control groups. The MMD is a commonly used measure of distance between probability distributions [19] that was recently proposed as a metric for balance assessment in causal inference [30]. Specifically, we show that the SVM dual problem given in

Equation (4) can be viewed as a regularized optimization problem for computing weights that minimize the MMD.

The MMD, which is also called the kernel distance, is a measure of distance between two probability distributions based on the difference in mean function values for functions in the unit ball of an RKHS. The MMD has found use in several statistical applications [31–33]. Given the unit ball RKHS  $\mathcal{F}_K = \{f \in \mathcal{H}_K : \|f\|_{\mathcal{H}_K} \leq 1\}$  and two probability measures  $F$  and  $G$ , the MMD is defined as

$$\gamma_K(F, G) := \sup_{f \in \mathcal{F}_K} \left| \int f dF - \int f dG \right| \quad (7)$$

An important property of the MMD is that when  $K$  is a characteristic kernel (e.g., the Gaussian radial basis function kernel and Laplace kernel), then  $\gamma_K(F, G) = 0$  if and only if  $F = G$  [34].

The computation of  $\gamma_K(F, G)$  requires the knowledge of both  $F$  and  $G$ , which is typically unavailable. In practice, an estimate of  $\gamma_K(F, G)$  using the empirical distributions  $\hat{F}_m$  and  $\hat{G}_n$  can be computed as

$$\gamma_K(\hat{F}_m, \hat{G}_n) = \left\| \frac{1}{m} \sum_{i: X_i \sim F} \phi(X_i) - \frac{1}{n} \sum_{j: X_j \sim G} \phi(X_j) \right\|_{\mathcal{H}_K} \quad (8)$$

where  $m$  and  $n$  are the size of the samples drawn from  $F$  and  $G$ , respectively. The properties of this statistic are well-studied [35]. In causal inference, the empirical MMD can be used to assess the balance between the treated and control samples [30]. This is done by setting  $F = P(X_i | T_i = 1)$  and  $G = P(X_i | T_i = 0)$ . Then, the quantity  $\gamma_K(\hat{F}_m, \hat{G}_n)$  gives a measure of independence between the treatment assignment  $T_i$  and the observed pre-treatment covariates  $X_i$ .

Equation (8) suggests a weighting procedure that balances the covariate distributions between the treatment and control groups by minimizing the empirical MMD. We define a weighted variant of the empirical MMD as

$$\gamma_K(\hat{F}_\alpha, \hat{G}_\alpha) = \left\| \sum_{i \in \mathcal{T}} \alpha_i \phi(X_i) - \sum_{j \in \mathcal{C}} \alpha_j \phi(X_j) \right\| = \sqrt{\alpha^\top Q \alpha} \quad (9)$$

where  $\hat{F}_\alpha$  and  $\hat{G}_\alpha$  denote the reweighted empirical distributions under weights  $\alpha$ , which are restricted to the simplex set,

$$\mathcal{A}_{\text{simplex}} = \left\{ \alpha \in \mathbb{R}^N : 0 < \alpha < 1, \sum_{i \in \mathcal{T}} \alpha_i = \sum_{j \in \mathcal{C}} \alpha_j = 1 \right\} \quad (10)$$

The optimization problem for minimizing the empirical MMD is therefore formulated as,

$$\begin{aligned} \min_{\alpha} \quad & \sqrt{\alpha^\top Q \alpha} \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{simplex}} \end{aligned} \quad (11)$$

Computing weights according to this problem is generally not preferable due to the lack of regularization, which leads to overfitting and sparse  $\alpha$ , resulting in many discarded samples. The following theorem establishes that the SVM dual problem can be

viewed as a regularized version of the optimization problem in Equation (11).

**Theorem 1.** (SVM Dual Problem as Regularized MMD Minimization). *Let  $\alpha_*(\lambda)$  denote the solution to the SVM dual problem under  $\lambda$ , defined in Equation (4). Consider the normalized weights  $\tilde{\alpha}_*(\lambda) = 2\alpha_*(\lambda)/\mathbf{1}^\top \alpha_*(\lambda)$  such that  $\tilde{\alpha}_*(\lambda) \in \mathcal{A}_{\text{simplex}}$ . Then,*

- (i) *There exists  $\lambda_*$  such that  $\tilde{\alpha}_*(\lambda_*)$  is a solution to the MMD minimization problem, defined in Equation (11).*
- (ii) *The quantity  $\tilde{\alpha}_*(\lambda)^\top Q \tilde{\alpha}_*(\lambda)$  is a monotonically increasing function of  $\lambda$ .*

The proof is given in Appendix A. Theorem 1 shows that the regularization parameter  $\lambda$  controls the trade-off between the covariate imbalance, measured as the MMD, and the effective sample size, measured as the sum of the support vector weights  $\mathbf{1}^\top \alpha$ . Thus, a greater size of the support vector set may lead to a worse covariate balance within that set.

Finally, although we do not consider such extensions in this paper, the MMD can be defined for non-RKHS function classes [19], see Section 2.5. See Huling and Mak [36] who consider an alternative non-parametric balance metric based on energy distance.

## 2.4 | SVM as a Relaxation of the Largest Balanced Subset Selection

SVM can also be seen as a continuous relaxation of the quadratic integer program (QIP) for computing the largest balanced subset. We modify the optimization problem in Equation (4) by replacing the continuous constraint  $0 < \alpha < 1$  with the binary integer constraint  $\alpha \in \{0, 1\}^N$ . This problem, which we refer to as SVM-QIP, is given by:

$$\begin{aligned} \min_{\alpha} \quad & \frac{1}{2\lambda} \alpha^\top Q \alpha - \mathbf{1}^\top \alpha \\ \text{s.t.} \quad & W^\top \alpha = 0 \\ & \alpha \in \{0, 1\}^N \end{aligned} \quad (12)$$

Interpreting the variables  $\alpha_i$  as indicators of whether or not unit  $i$  is selected into the optimal subset, we see that the objective is a trade-off between the subset balance in the projected features (first term) and subset size (second term), where balance is measured by a difference in sums. However, the constraint  $W^\top \alpha = 0$  requires the optimal subset to have an equal number of treated and control units, so balancing the feature sums also implies balancing the feature means.

Thus, the SVM dual in Equation (4) is a continuous relaxation of the largest balanced subset problem represented by SVM-QIP in Equation (12), with the set of support vectors comprising an approximation to the largest balanced subset, as these are the units for which  $\alpha_i > 0$ . The quality of the approximation is difficult to characterize generally since differences between the two computed subsets are influenced by a number of factors, including separability in the data, the value of  $\lambda$ , and the choice of

kernel. However, in our own experiments, we observe a significant overlap between the two solutions. This result suggests that SVM uses non-integer weights to augment the SVM-QIP solution without compromising the balance in the selected subset. This highlights the advantages of weighting over matching.

In Appendix B, we further motivate SVM as a balancing procedure by connecting it to two existing covariate balancing methods, cardinality matching [26] and stable balancing weights [7]. Specifically, we show that SVM can be viewed as a weighting analog to cardinality matching. In addition, a variant of the SVM problem defined in Equation (3), called L2-SVM, forms an optimization problem that simultaneously minimizes balance and weight dispersion. SVM differs from these methods primarily in the way that balance is enforced.

## 2.5 | Regularization Path as a Balance-Sample Size Frontier

Another important advantage of using SVM to perform covariate balancing is the existence of path algorithms that can efficiently compute the set of solutions to Equation (4) over different values of  $\lambda$ . Since Theorem 1 establishes that  $\lambda$  controls the trade-off between the MMD and subset size, the path algorithm for SVM can be viewed as the weighting analog to the balance-sample size frontier [22]. Below, we briefly discuss the path algorithm and explain how the path can be interpreted as a balance-sample size frontier.

Path algorithms for SVM were first proposed by Hastie et al. [20], who showed that the weights  $\alpha$  and scaled intercept  $\alpha_0 := \lambda\beta_0$  are piecewise linear in  $\lambda$  and presented an algorithm for computing the entire path of solutions with a comparable computation cost to finding a single solution. However, their algorithm was prone to numerical problems and would fail in the presence of singular submatrices of  $\mathbf{Q}$ . Recent work on SVM path algorithms has addressed these issues. In our analysis, we use the path algorithm presented in Sentelle, Anagnostopoulos, and Georgiopoulos [21], which we briefly describe here.

The regularization path for SVM is characterized by a sequence of breakpoints, representing the values of  $\lambda$  at which either one of the support vectors on the margin  $W_i f(X_i) = 1$  exits the margin, or a non-marginal observation reaches the margin. Between these breakpoints, the coefficients of the marginal support vectors  $\alpha_i$  change linearly in  $\lambda$ , while the coefficients of all other observations stay fixed as  $\lambda$  is changed. Since the KKT conditions must be met for any solution  $\alpha$ , we can form a linear system of equations to compute how each  $\alpha_i$  and  $\alpha_0$  changes with respect to  $\lambda$ .

Based on this idea, beginning with an initial solution corresponding to some large initial value of  $\lambda$ , the path algorithm first computes how the current marginal support vectors change with respect to  $\lambda$ . Given this quantity, the next breakpoint in the path is computed by decreasing  $\lambda$  until a marginal support vector exits the margin, i.e.,  $\alpha_i = 0$  or  $\alpha_i = 1$ , or a non-marginal observation enters the margin, i.e.,  $W_i f(X_i) = 1$ . At this point, the marginal support vector set is updated, and the changes in  $\alpha_i$  and  $\alpha_0$ , as well as the next breakpoint, are computed. This procedure repeats until the terminal value of  $\lambda$  is reached.

The initial solution in the SVM regularization path corresponds to the solution at  $\lambda_{\max}$  such that for any  $\lambda > \lambda_{\max}$ , the minimizing weight vector  $\alpha$  does not change. We assume without a loss of generality that  $n_T \leq n_C$ . Then initially,  $\alpha_i = 1$  for all  $i \in \mathcal{T}$ , and the remaining weights are computed according to,

$$\begin{aligned} \arg \alpha \min \quad & \alpha^\top \mathbf{Q} \alpha \\ \text{s.t.} \quad & \sum_{i \in \mathcal{C}} \alpha_i = n_T \\ & \alpha_i = 1, \quad i \in \mathcal{T} \\ & \alpha \in [0, 1]^N \end{aligned} \quad (13)$$

The initial solution computes control weights to minimize the weighted empirical MMD while fixing the renormalized weights  $\tilde{\alpha}_i = n_T^{-1}$  for the treated units. This corresponds to the largest subset, as measured by  $\sum_i \alpha_i$ , amongst all solutions on the regularization path.

The regularization path completes when the resulting solution has no non-marginal support vectors, or in the case of non-separable data, when  $\lambda = 0$ . In practice, however, path algorithms run into numerical issues when  $\lambda$  is small, so we terminate the path at  $\lambda_{\min} = 1 \times 10^{-3}$ , which appears to work well in our experiments. This value is often greater than  $\lambda^*$ , which corresponds to the MMD-minimizing solution defined in Theorem 1. In our experience, however, the differences in balance between these two solutions are negligible.

Summarizing the regularization path, we see that the initial solution at  $\lambda_{\max}$  has the largest weight sum  $\sum_{i=1}^N \alpha_i = 2 \min\{n_T, n_C\}$  and can be viewed as the largest balanced subset retaining all observations in the minority class. As we move through the path, the SVM dual problem imposes greater restrictions on balance in the subset, which leads to smaller subsets with better balance, until we reach the terminal value  $\lambda_{\min}$ , at which the weighted empirical MMD is smallest on the path.

## 2.6 | Causal Effect Estimation

Theorem 1 establishes that the SVM dual problem can be viewed as a regularized optimization problem for computing balancing weights to minimize the MMD. However, achieving a high degree of balance often requires significant pruning of the original sample, especially in scenarios where the covariate distributions for the treated and control groups have limited overlap. In this section, we provide a characterization of this trade-off between subset size and subset balance and discuss its impact on the bias of causal effect estimates.

Recent work by Kallus [14] established that several existing matching and weighting methods are minimizing the dual norm of the bias for a weighted estimator, a property called error dual norm minimizing. The author proposes a new method, kernel optimal matching (KOM), which minimizes the dual norm of the bias when the conditional expectation functions are embedded in an RKHS. The idea is that the RKHS encapsulates a general class of functions that can accommodate non-linearity and non-additivity in the conditional expectation functions, leading to balancing weights which are robust to model misspecification. Below, we show that SVM also fits into this KOM framework.

We restrict our attention to the following weighted difference-in-means estimator,

$$\hat{\tau} = \sum_{i \in \mathcal{T}} \alpha_i Y_i - \sum_{i \in \mathcal{C}} \alpha_i Y_i \quad (14)$$

where  $\alpha \in \mathcal{A}_{\text{simplex}}$  is computed via the application of SVM to the data  $\{\mathbf{X}_i, T_i\}_{i=1}^N$ . Below, we derive the form for the conditional bias with respect to two estimands, the sample average treatment effect (SATE),  $\tau_{\text{SATE}}$ , and the sample average treatment effect for the treated (SATT),  $\tau_{\text{SATT}}$ . We then discuss how to compute this bias when the conditional expectation functions  $f_0$  and  $f_1$  are unknown.

As shown in Appendix C, under Assumptions 1 and 2, the conditional bias with respect to  $\tau_{\text{SATE}}$  and  $\tau_{\text{SATT}}$  for the estimator above is given by,

$$\mathbb{E}(\hat{\tau} - \tau \mid \{\mathbf{X}_i, T_i\}_{i=1}^N) = \sum_{i=1}^N \alpha_i W_i f_0(\mathbf{X}_i) + \sum_{i=1}^N (\alpha_i T_i - v_i) \tau(\mathbf{X}_i) \quad (15)$$

where

$$v_i = \begin{cases} 1/N & \text{if } \tau = \tau_{\text{SATE}} \\ T_i/n_T & \text{if } \tau = \tau_{\text{SATT}} \end{cases} \quad (16)$$

and  $\tau(\mathbf{X}_i) := \mathbb{E}(Y_i(1) - Y_i(0) \mid \mathbf{X}_i) = f_1(\mathbf{X}_i) - f_0(\mathbf{X}_i)$ .

The first term in Equation (15) represents the bias due to the imbalance of prognostic score [37]. Unfortunately, computation of this quantity is difficult since  $f_0$  is typically unknown. However, we can embed  $f_0$  in a unit-ball RKHS,  $\mathcal{F}_K$ , and consider an  $f_0$  that maximizes the square of this bias term. As shown in Appendix D, this strategy leads to the following optimization problem that is of the same form as that of the MMD minimization problem given in Equation (11),

$$\begin{aligned} \min_{\alpha} \quad & \gamma_K^2(\hat{F}_{\alpha}, \hat{G}_{\alpha}) \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{simplex}} \end{aligned} \quad (17)$$

Thus, the SVM dual problem can also be viewed as a method for minimizing bias due to prognostic imbalance.

However, SVM does not address the second term of the conditional bias in Equation (15). This term represents the bias due to extrapolation outside of the weighted treatment group to the population of interest. For example, if the SATT is the target quantity, the term represents the bias due to the difference between the weighted and unweighted CATE for the treatment group. Thus, the prognostic balance achieved by SVM may induce this CATE bias. This is a direct consequence of the trade-off between balance and effective sample size, which is controlled by the regularization parameter as shown earlier.

## 2.7 | Relation to Kernel Balancing Methods

We now discuss the relations between SVM and a class of closely related methods, known as kernel balancing [11, 13, 38].

Consider the following alternative decomposition of conditional bias,

$$\begin{aligned} \mathbb{E}(\hat{\tau} - \tau \mid \{\mathbf{X}_i, T_i\}_{i=1}^N) &= \sum_{i=1}^N (\alpha_i T_i - v_i) f_1(\mathbf{X}_i) \\ &+ \sum_{i=1}^N \{v_i - \alpha_i(1 - T_i)\} f_0(\mathbf{X}_i) \end{aligned} \quad (18)$$

To minimize this bias, kernel balancing methods restrict  $f_0$  and  $f_1$  to a RKHS  $\mathcal{F}_K := \{(f_0, f_1) \in \mathcal{H}_{K_0} \times \mathcal{H}_{K_1} : \sqrt{\|f_0\|_{\mathcal{H}_{K_0}}^2 + \|f_1\|_{\mathcal{H}_{K_1}}^2} \leq c\}$  for some  $c$  and consider minimizing the largest bias under the pair  $(f_0, f_1) \in \mathcal{F}_K$ . This problem is given by,

$$\begin{aligned} \min_{\alpha} \quad & \sup_{f_0, f_1 \in \mathcal{F}_K} \left[ \sum_{i=1}^N (\alpha_i T_i - v_i) f_1(\mathbf{X}_i) - \sum_{i=1}^N \{v_i - \alpha_i(1 - T_i)\} f_0(\mathbf{X}_i) \right]^2 \\ \text{s.t.} \quad & \alpha \in \mathcal{A} \end{aligned} \quad (19)$$

where  $\mathcal{A}$  denotes the constraints on the weights. For example, Kallus, Pennicooke, and Santacatterina [38] restricts them to  $\mathcal{A}_{\text{simplex}}$  whereas Wong and Chan [11] essentially uses  $\alpha_i \geq N^{-1}$ , though their formulation is slightly different than that given above. Note that approaches targeting ATT for estimation, such as the one considered by Hazlett [13], fix all treated weights to  $\alpha_i = n_T^{-1}$  and focus only on the term involving  $f_0(\mathbf{X}_i)$ .

As shown in Kallus, Pennicooke, and Santacatterina [38], the problem of minimizing this worst-case conditional bias amounts to computing weights that balance the treatment and control covariate distributions with respect to the empirical distribution for the population of interest. Let  $\hat{F}_{\alpha}$  and  $\hat{G}_{\alpha}$  denote the weighted empirical covariate distributions for the treatment and control groups, respectively, while having  $\hat{H}_v$  represent the empirical covariate distribution corresponding to the population of interest. Then if  $\mathcal{F}_K$  is also restricted to the unit-ball RKHS (fixing the size of  $f_0, f_1$  is necessary since the bias scales linearly with  $\|f_0\|_{\mathcal{H}_{K_0}}$  and  $\|f_1\|_{\mathcal{H}_{K_1}}$ ), i.e.,  $c = 1$ , the optimization problem in Equation (19) can be written in terms of the minimization of the empirical MMD statistic:

$$\begin{aligned} \min_{\alpha} \quad & \gamma_{K_1}^2(\hat{F}_{\alpha}, \hat{H}_v) + \gamma_{K_0}^2(\hat{G}_{\alpha}, \hat{H}_v) \\ \text{s.t.} \quad & \alpha \in \mathcal{A} \end{aligned} \quad (20)$$

This objective does not contain a measure of distance between the conditional covariate distributions  $\hat{F}_{\alpha}$  and  $\hat{G}_{\alpha}$ . Instead, balance between these two distributions is indirectly encouraged through balancing each one individually with respect to the target distribution  $\hat{H}_v$ . This is in contrast with SVM, which directly balances the covariate distribution between the treatment and control groups.

## 3 | Simulations

In this section, we examine the performance of SVM in ATE estimation under two different simulation settings. We also examine the connection between SVM and the QIP for the largest balanced subset.

### 3.1 | Setup

We consider two simulation setups used in previous studies. Simulation A comes from Lee, Lessler, and Stuart [39] who use a slightly modified version of the simulations presented in Setoguchi et al. [40]. We adopt the exact setup corresponding to their “scenario G,” which is briefly summarized here. We refer readers to the original article for the exact specifications. For each simulated dataset, we generate 10 covariates  $\mathbf{X}_i = (X_{i1}, \dots, X_{i10})^\top$  from the standard normal distribution, with correlation introduced between four pairs of variables. Treatment assignment is generated according to  $P(T_i = 1 | \mathbf{X}_i) = \text{expit}(\boldsymbol{\beta}^\top f(\mathbf{X}_i))$ , where  $\boldsymbol{\beta}$  is some coefficient vector, and  $f(\mathbf{X}_i)$  controls the degree of additivity and linearity in the true propensity score model. This scenario uses the true propensity score model with a moderate amount of non-linearity and non-additivity. The outcome model was specified to be linear in the observed covariates with a constant, additive treatment effect:  $Y_i(T_i) = \gamma_0 + \boldsymbol{\gamma}^\top \mathbf{X}_i + \tau T_i + \epsilon_i$ , with  $\tau = -0.4$  and  $\epsilon_i \sim \mathcal{N}(0, 0.1)$ .

Simulation B comes from Wong and Chan [11] and represents a more difficult scenario where both the propensity score and outcome regression models are non-linear in the observed covariates. For each simulated data set, we generate a ten-dimensional random vector  $\mathbf{Z}_i = (Z_{i1}, \dots, Z_{i10})^\top$  from the standard normal distribution. The observed covariates are the non-linear functions of these variables,  $\mathbf{X}_i = (X_{i1}, \dots, X_{i10})^\top$ , where  $X_1 = \exp(Z_1/2)$ ,  $X_2 = Z_2/[1 + \exp(Z_1)]$ ,  $X_3 = (Z_1 Z_3/25 + 0.6)^3$ ,  $X_4 = (Z_2 + Z_4 + 20)^2$ , and  $X_j = Z_j$ ,  $j = 5, \dots, 10$ . Treatment assignment follows  $P(T_i = 1 | \mathbf{Z}_i) = \text{expit}(-Z_1 - 0.1Z_4)$ , which corresponds to Model 1 of Wong and Chan [11]. Finally, the outcome model is specified as  $Y(T_i) = 200 + 10T_i + (1.5T_i - 0.5)(27.4Z_1 + 13.7Z_2 + 13.7Z_3 + 13.7Z_4) + \epsilon_i$ , with  $\epsilon_i \sim \mathcal{N}(0, 1)$ . Note that the true PATE is 10 under this model.

### 3.2 | Comparison Between SVM and SVM-QIP

We begin by examining the connection between SVM and SVM-QIP by comparing solutions obtained using one simulated dataset of  $N = 500$  units under Simulation A. Specifically, we first compute the SVM path using the path algorithm described

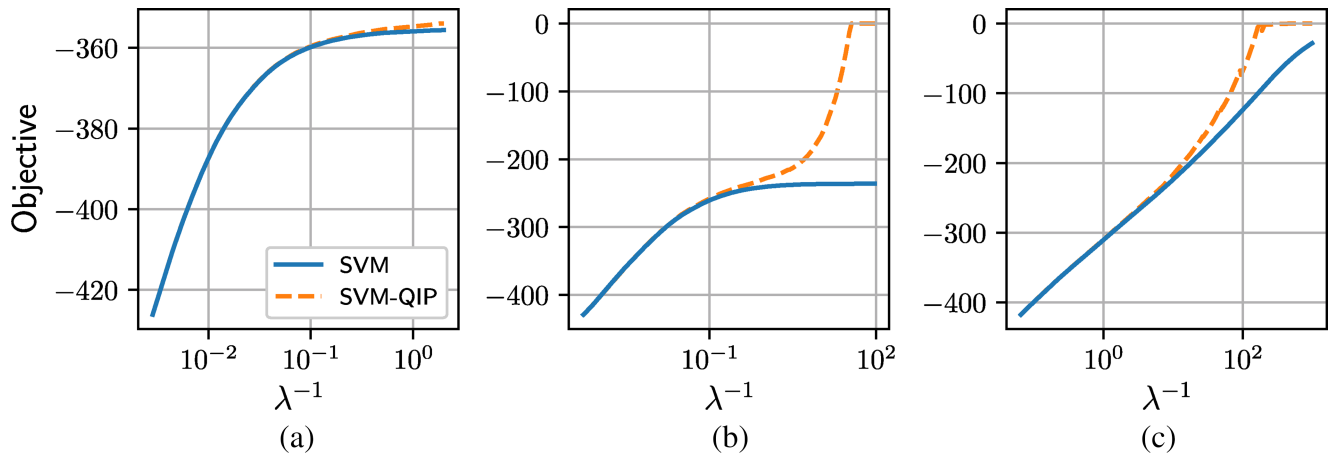
in Section 2.5, obtaining a set of regularization parameter breakpoints  $\lambda$ . Next, we compute the SVM-QIP solution for each of these breakpoints using the Gurobi optimization software [41]. We limit the solver to spending 5 min of runtime for each problem. Finding the *exact* integer-valued solution under a given  $\lambda$  requires a significant amount of time, but a good approximation can typically be found in a few seconds.

For both methods, we compute the objective function value at each of the breakpoints as well as the coverage of the SVM-QIP solution by the SVM solution. The latter represents the proportion of units with non-zero SVM weights that are included in the largest balanced subset identified by SVM-QIP. Formally, the coverage is defined as

$$\text{cvg}(\lambda) = \frac{|\lceil \boldsymbol{\alpha}_{\text{SVM}}(\lambda) \rceil \cap \boldsymbol{\alpha}_{\text{SVM-QIP}}(\lambda)|}{|\boldsymbol{\alpha}_{\text{SVM-QIP}}(\lambda)|}$$

To examine the effects of separability on the quality of the approximation, we perform the above analysis using three different types of features. Specifically, we use a linear kernel with the untransformed covariates (linear), a linear kernel with the degree-2 polynomial features formed by concatenating the original covariates with all two-way interactions and squared terms (polynomial), and the Gaussian RBF with scale parameter chosen according to the median heuristic (RBF). In all cases, we scale the input feature matrix such that the columns have 0 mean and standard deviation of 1 before performing the kernel computation.

Figure 1 shows that the objective values for the SVM and SVM-QIP solutions are close when the penalty on the balance  $\lambda^{-1}$  is small, with divergence between the two methods occurring towards the end of the regularization path. In the linear case, we see that the paths for the two methods are nearly identical, suggesting that their solutions are essentially the same. Divergence in the polynomial and RBF settings is more pronounced due to greater separability in the transformed covariate space, which is more difficult to balance without non-integer weights. When  $\lambda$  is very small, we also find that SVM-QIP returns  $\boldsymbol{\alpha} = \mathbf{0}$ , indicating that the penalty on balance is too great. Lastly, the effects of



**FIGURE 1** | Comparison of the objective value between SVM and SVM-QIP. The blue line denotes the objective value of the SVM solution, and the orange dotted line denotes the objective value of the SVM-QIP solution. (a) Linear; (b) Polynomial; (c) RBF.

approximating the SVM-QIP solution are reflected in the RBF setting, where upon close inspection the objective value appears to be somewhat noisy and non-monotonic.

Interestingly, the coverage plots in Figure 2 show that even when the objective values of the two methods are divergent, the SVM solution still predominantly covers the SVM-QIP solution. The regions with zero coverage in the polynomial and RBF settings correspond to instances where the balance penalty is so significant that a non-trivial solution cannot be found for the SVM-QIP. This result illustrates that SVM approximates one-to-one matching by augmenting a well-balanced matched subsample with some non-integer weights. This leads to an increased subset size while preserving the overall balance within the subsample.

### 3.3 | Performance of SVM

Next, we evaluate the performance of SVM in estimating the ATE for Simulations A and B. For each scenario, we generate 1,000 datasets with  $N = 500$  samples. For each simulated dataset, we compute the ATE estimate over a fixed grid of 100  $\lambda$  values chosen based on the simulation scenario and input feature. As described in Section 3.2, we use the linear, polynomial, and RBF-induced features, standardizing the covariate matrix before passing it to the kernel in all cases.

Figure 3 plots the distribution of ATE estimates over Monte Carlo simulations against the regularization parameter  $\lambda$ . The results for Simulation A (top panel) show that the bias approaches zero as the penalty on balance increases ( $\lambda$  decreases). This is because the conditional bias in the estimate under the outcome model for Simulation A is given by,

$$\mathbb{E}[\hat{\tau} - \tau | X_{1:N}, T_{1:N}] = \gamma^\top \left( \sum_{i=1}^N \alpha_i T_i X_i \right)$$

This implies that all bias comes from prognostic score imbalance. This quantity becomes the smallest when minimizing  $\|\sum_{i=1}^N \alpha_i T_i X_i\|$ , which is controlled by the regularization

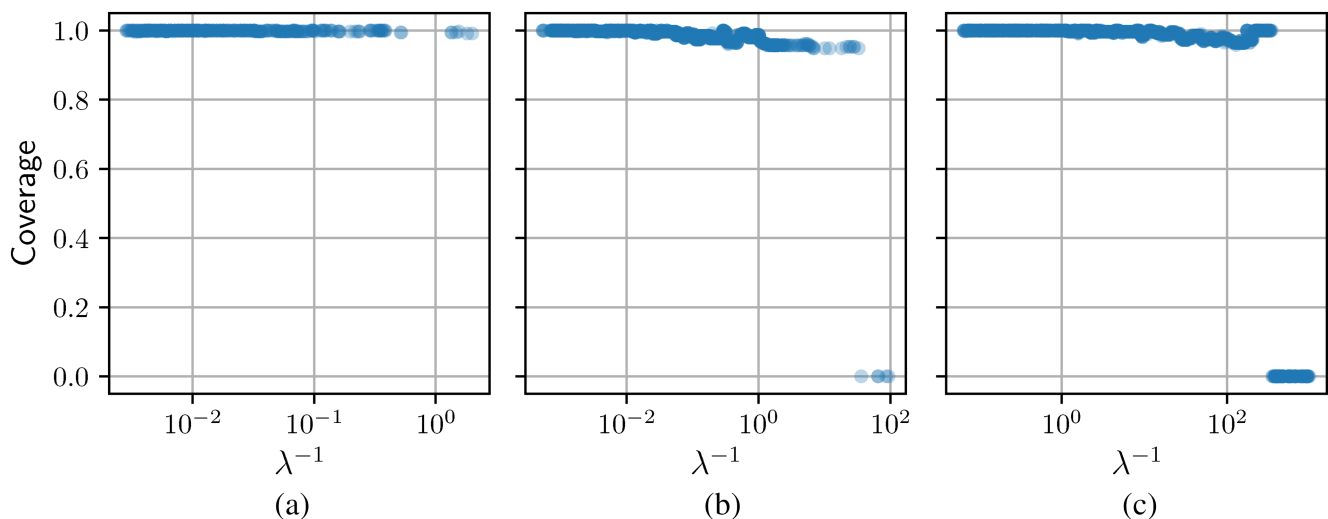
parameter  $\lambda$  in the SVM dual objective under the linear setting. Note that this quantity is also small under both the polynomial and RBF input features.

We also find that under all three settings, there is relatively little change in the variance of the estimates along most of the path, suggesting that the variance gained from trimming the sample is counteracted by the variance decreased from correcting for heteroscedasticity. The exception to this observation occurs at the beginning of the linear case, where the reduction in bias also reduces the variance, and at the end of the RBF path, where the amount of trimming is so substantial relative to the balance gained that the variance increases.

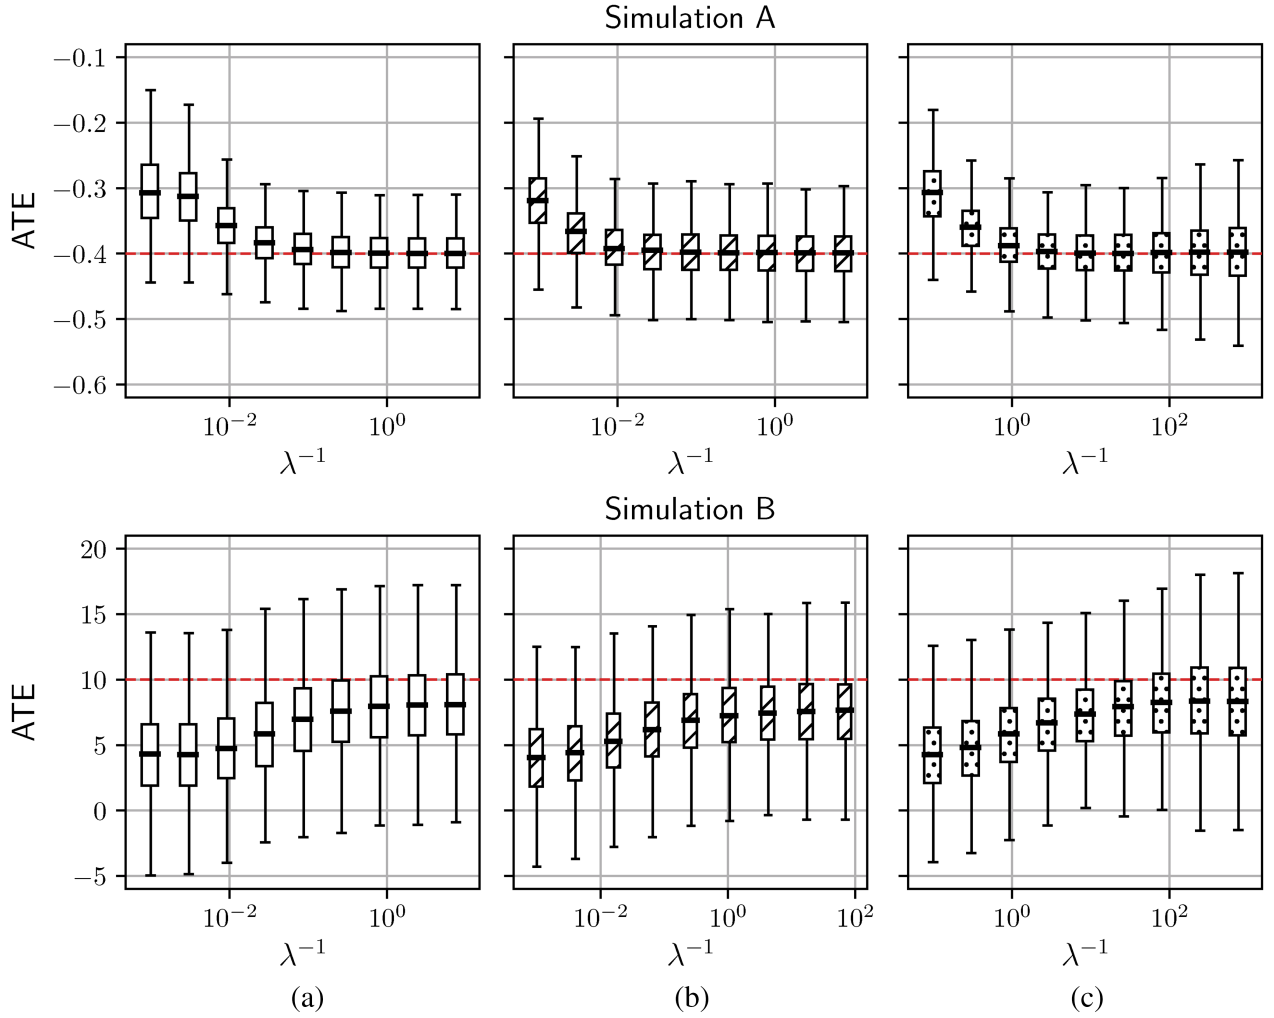
For Simulation B (bottom panel), the bias decreases as the penalty on balance increases. However, due to misspecification, non-linearity, and treatment effect heterogeneity in the outcome model, the bias never decays to zero as shown in Section 2.6. We also find that the SVM with the linear kernel can reduce bias as well as the other kernels, suggesting that SVM is robust to misspecification and non-linearity in the outcome model. Similar to Simulation A, we observe relatively small changes in the variance as the constraint on balance increases, except at the end of the RBF path where there is substantial sample pruning.

### 3.4 | Comparison With Other Methods

Next, we compare the performance of SVM with that of other methods. Our results below show that the performance of SVM is comparable to that of related state-of-the-art covariate balancing methods available in the literature. In particular, we consider kernel optimal matching KOM; [38], kernel covariate balancing KCB; [11], cardinality matching CARD; [26], and inverse propensity score weighting (IPW) based on logistic regression (GLM) and random forest (RFRST), both of which were used in the original simulation study by Lee, Lessler, and Stuart [39]. For SVM, we compute solutions using  $\lambda^{-1} = 0.42$ ,  $\lambda^{-1} = 0.10$ , and  $\lambda^{-1} = 2.60$  for Simulation A under the linear, polynomial, and RBF settings,



**FIGURE 2** | The proportion of samples in the SVM-QIP largest balanced subset covered by the SVM solution. The instances of zero coverage in the polynomial and RBF settings represent the cases where the SVM-QIP fails to find a non-trivial solution. (a) Linear; (b) Polynomial; (c) RBF.



**FIGURE 3** | ATE estimates for Simulations A (top) and B (bottom) over the SVM regularization path. The boxplots represent the distribution of the ATE estimates over Monte Carlo simulations. The red dashed line corresponds to the true ATE. (a) Linear; (b) Polynomial; (c) RBF.

respectively. For Simulation B, we use  $\lambda^{-1} = 1.07$ ,  $\lambda^{-1} = 1.92$ , and  $\lambda^{-1} = 10.48$ . These values are taken from the grid of  $\lambda$  values used in the simulation based on visual inspection of the path plots in Figure 3 around where the estimate curve flattens out.

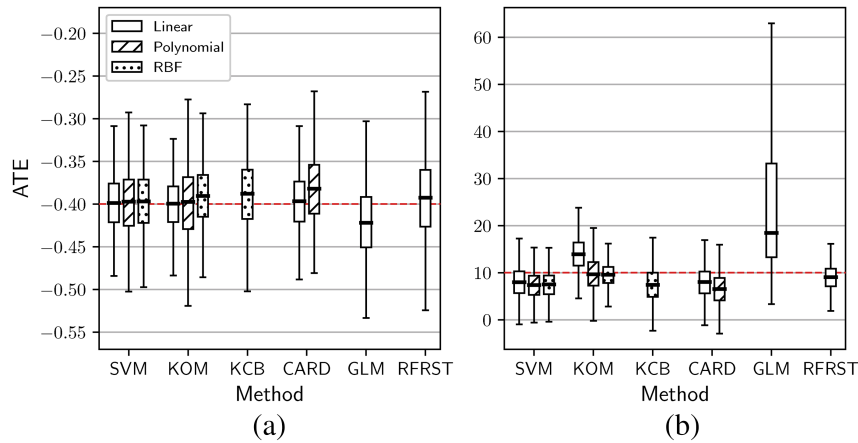
For KOM, we compute weights under the linear, polynomial, and RBF settings described earlier with the default settings for the provided code. For KCB, we compute weights using the RBF kernel and use its default settings. While KCB allows for other kernel functions, it was originally designed for the use of RBF and Sobolev kernels. We find its results to be poor when using the linear and polynomial features. For CARD, we used a threshold of 0.01 and 0.1 times the standardized difference-in-means for the linear- and polynomial-induced features, respectively, and we set the search time for the algorithm to 3 min. For GLM and RFRST, we used the linear-induced features and the default algorithm settings described in Lee, Lessler, and Stuart [39].

Figure 4 plots the distributions of the effect estimates over 1,000 simulated datasets for both scenarios. Simulation A (left panel) shows comparable performance across all methods, with SVM and KOM having the best performance in terms of both bias and variance. In particular, SVM achieves near zero bias under all

three input features. The results for KCB show that it performs slightly worse in comparison to the other kernel methods, with greater bias and variance under the RBF setting.

The results for CARD show near identical performance with SVM under the linear setting, however, results under the polynomial setting are notably worse. The reason for this comes from the choice of balance threshold, which was set to 0.1 times the standardized difference-in-means of the input feature matrix. Although decreasing the scalar below 0.1 would lead to a more balanced matching, we found that the algorithm was unable to consistently find a solution for all datasets with scalar multiples smaller than 0.1. This result highlights the main issue with defining balance dimension-by-dimension, which makes it difficult to enforce a small overall balance without information on the underlying geometry of the data. Lastly, the propensity score methods show the worst performance. This is somewhat expected as the true propensity score model is more complicated than the true outcome model under this simulation setting.

We note that further reduction in the variance of the SVM solution while preserving bias is likely possible with a more principled method of choosing the solution for each simulated dataset. In



**FIGURE 4** | Boxplots for ATE estimates for Simulations A (left) and B (right). The hatch pattern denotes the input feature (Linear, Polynomial, or RBF) — kernel optimal matching KOM; [38], kernel covariate balancing KCB; [11], cardinality matching CARD; [26], and IPW with propensity score modeling via logistic regression (GLM) and random forest (RFRST). The red dashed line corresponds to the true ATE. (a) Simulation A; (b) Simulation B.

general, a value of  $\lambda$  that works well for one dataset may not work for another. A better approach would examine estimates over the path and balance-sample size curves for each dataset individually. Nevertheless, our heuristic procedure for selecting a solution produced high-quality results.

The results for Simulation B (right panel) show a slightly more varying performance across methods. Amongst the kernel methods, we find that KOM has the best performance under the polynomial and RBF settings, achieving near zero bias under these scenarios, while SVM has the best performance under the linear setting. The discrepancy under the linear setting is due to misspecification, which leads to a poor regularization parameter choice and consequently poor balance and bias under the KOM procedure.

We also find that SVM is unable to drive the bias to zero, which is due to the treatment effect heterogeneity in the outcome model. As discussed in Section 2.6, SVM ignores the second term in the conditional bias decomposition in Equation (15), which is zero under a constant additive treatment effect in Simulation A but is non-zero in Simulation B. In contrast, KOM targets both bias terms in its formulation, which leads to greater bias reduction.

In comparison to the other kernel methods, KCB has a comparable bias to SVM but greater variance. For CARD, we observe comparable results to SVM under the linear setting, but worse performance under the polynomial setting due to the reasons mentioned above. Lastly, we find mixed results between the two propensity score methods. Logistic regression (GLM) has the worst performance while Random forest (RFRST) exhibits the second-best performance. This result is likely due to the simple structure of the true propensity score model, whose non-linearity can only be accurately modeled by RFRST.

## 4 | Empirical Application: Right Heart Catheterization Study

In this section, we apply the proposed methodology to the right heart catheterization (RHC) data set originally analyzed by

Connors et al. [27]. This observational data set was used to study the effectiveness of right heart catheterization, a diagnostic procedure, for critically ill patients. The key result from the study was that after adjusting for a large number of pre-treatment covariates, right heart catheterization appeared to reduce survival rates. This finding contradicts the existing medical perception that the procedure is beneficial.

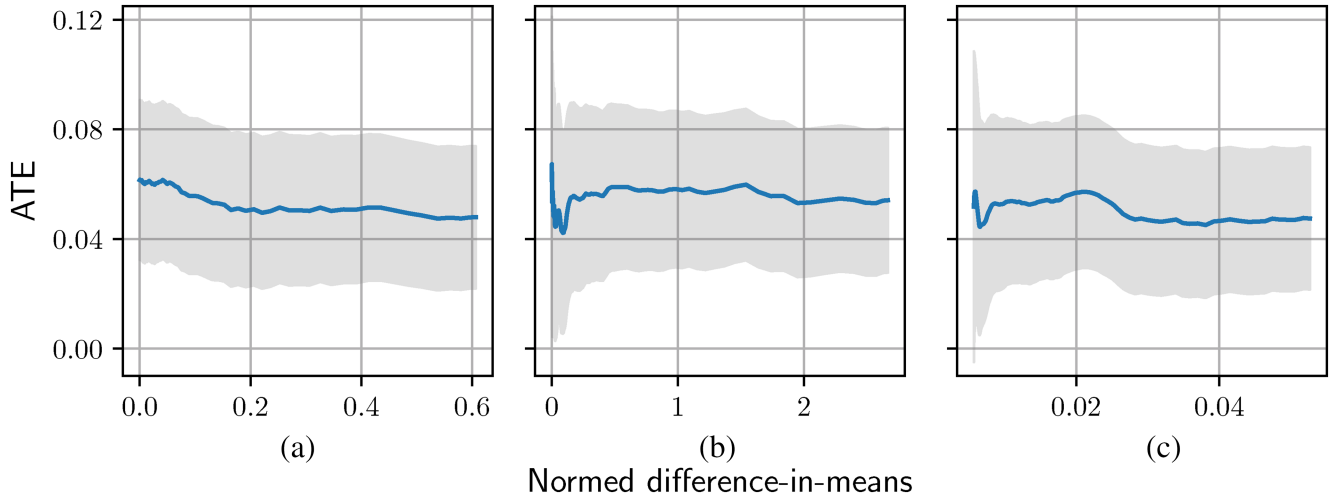
### 4.1 | Data and Methods

The data set consists of 5,735 patients, with 2,184 of them assigned to the treatment group and 3,551 assigned to the control group. For each patient, we observe the treatment status, which indicates whether or not he/she received catheterization within 24 h of hospital admission. The outcome variable represents death within 30 days. Finally, the dataset contains a total of 72 pre-treatment covariates that are thought to be related to the decision to perform the right heart catheterization. These variables include background information about the patient, such as age, sex, and race, indicator variables for primary/secondary diseases and comorbidities, and various measurements from the medical test results.

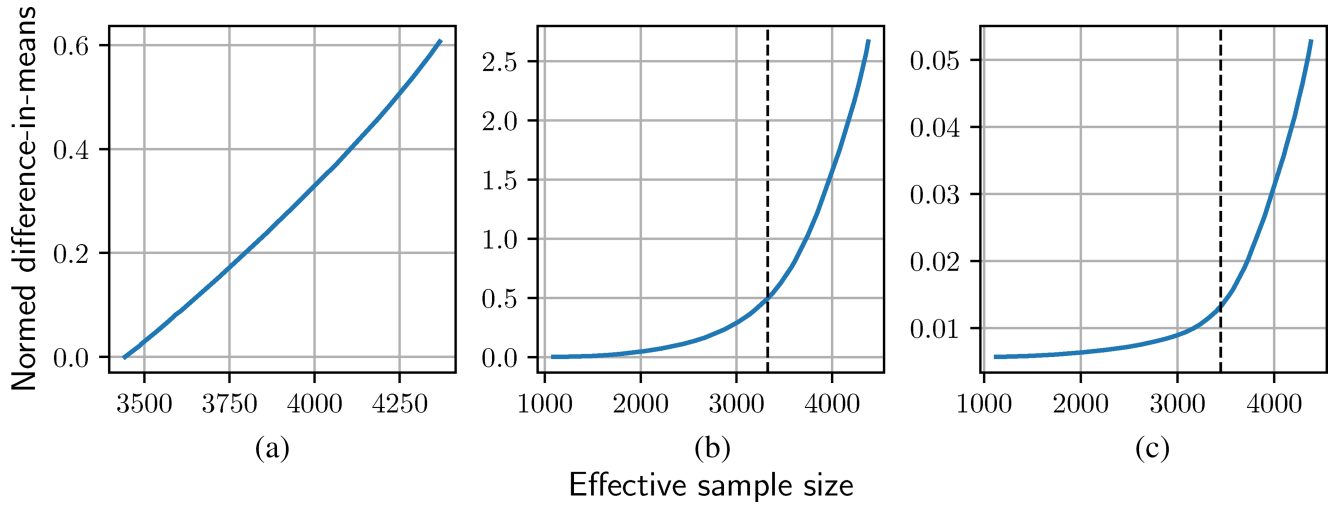
We compute the full SVM regularization paths under the linear, polynomial, and RBF settings described in Section 3.1. For the polynomial features, we exclude all trivial interactions (e.g., interactions between categories of the same categorical variable) and squares of binary-valued covariates. We also compute the KOM weights under all three settings, the KCB weights under the RBF setting, and the CARD weights under the linear and polynomial settings with a threshold set to 0.1 times the standardized difference-in-means.

### 4.2 | Results

Figure 5 plots the ATE estimates over the SVM regularization paths with the pointwise 95% confidence intervals based on the weighted Neyman variance estimator [42], Chapter 19. The horizontal axis represents the normed difference-in-means within



**FIGURE 5** | ATE estimates for the RHC data over the SVM regularization path. The horizontal axis represents the normed difference-in-means in covariates within the weighted subset. The solid blue line denotes the average estimate, and the solid gray background denotes the pointwise 95% confidence intervals. (a) Linear; (b) Polynomial; (c) RBF.



**FIGURE 6** | Trade-off between balance and effective sample size. The black dashed-line indicates the estimated elbow point. (a) Linear; (b) Polynomial; (c) RBF.

the weighted subset as a covariate balance measure. For all three settings, we find that the estimated ATE slightly increases as the weighted subset becomes more balanced, supporting the results originally reported in [27] that right heart catheterization decreased survival rates.

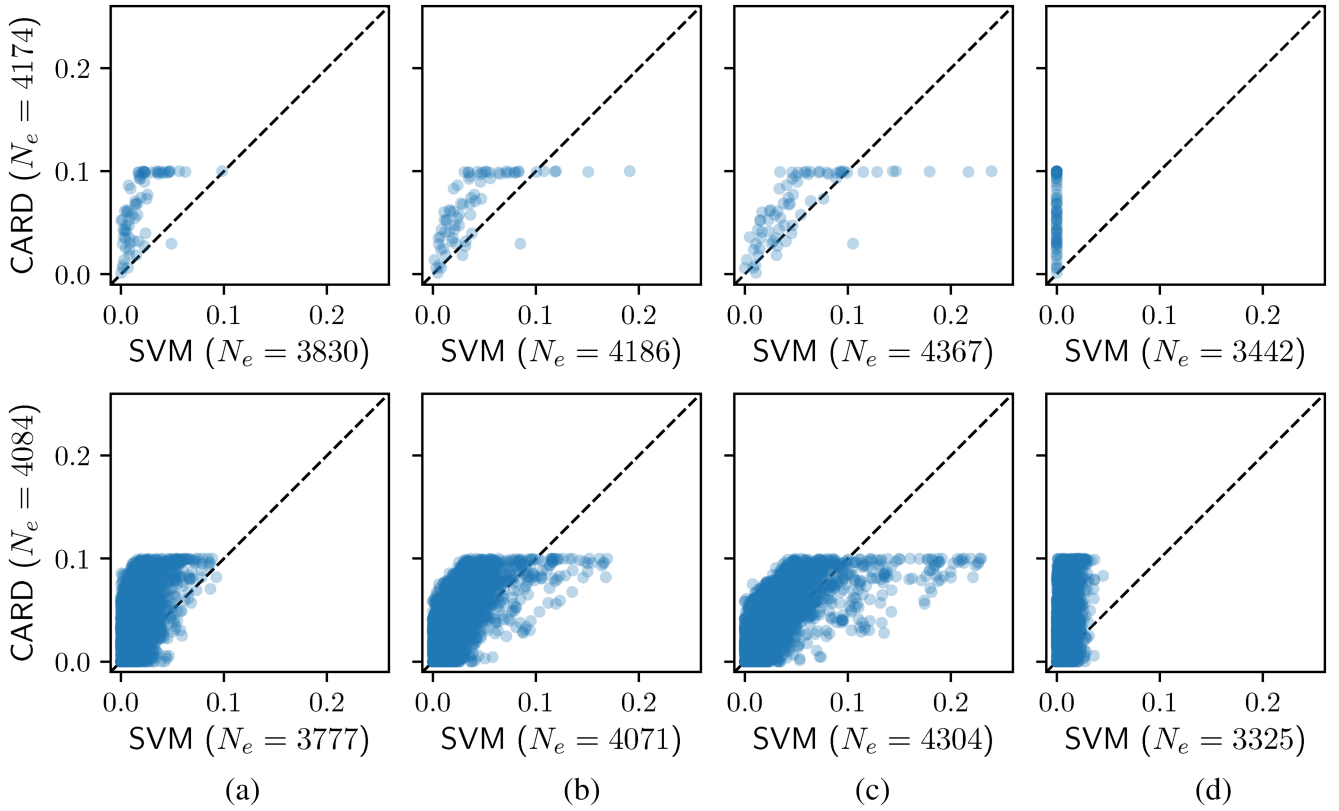
Figure 6 illustrates the trade-off between the balance measure (the normed difference-in-means in covariates within the weighted subset) and effective subset size, as the balance-sample size frontier. Such graphs can be useful to researchers in selecting a solution along the regularization path for estimating the ATE. Across all cases, we achieve a good amount of balance improvement once the data set is pruned to about 3,500, which occurs around where the trade-off between subset size balance becomes less favorable.

We also examine differences in dimension-by-dimension balance between SVM and CARD and between SVM and KOM under

the linear and polynomial settings. We do not conduct such a comparison for RBF, which is infinite-dimensional. Here, we consider four different SVM solutions: the largest subset size solution whose standardized difference-in-means in covariates was below 0.1 for all dimensions, the solution whose effective sample size was nearest the subset size for the other method, the solution whose normed difference-in-means in covariates were closest to that of the other method, and the solution occurring at the kneedle estimate for the elbow of the balance-weight sum curve. We take the minimum-balance solution when no elbow exists, as in the linear case. The effective sample size is computed according to the following Kish's formula:

$$N_e = \frac{(\sum_{i \in \mathcal{T}} \alpha_i)^2}{\sum_{i \in \mathcal{T}} \alpha_i^2} + \frac{(\sum_{i \in \mathcal{C}} \alpha_i)^2}{\sum_{i \in \mathcal{C}} \alpha_i^2} \quad (21)$$

Figure 7 presents the covariate balance comparisons between SVM and CARD for both linear and polynomial settings.



**FIGURE 7** | Comparison of covariate standardized difference-in-means between SVM and cardinality matching (CARD) under the linear (top) and polynomial (bottom) settings with different SVM solutions: (a) standardized difference-in-means less than 0.1 in all covariates, (b) effective sample size closest to that of CARD, (c) normed difference-in-means closest to that of CARD, and (d) elbow of the regularization path. The effective sample size for each method is given as  $N_e$  in the parentheses. Note that darker areas correspond to higher concentrations of points. (a) Small difference-in-means; (b) Closest effective sample size; (c) Closest normed difference-in-means; (d) Elbow.

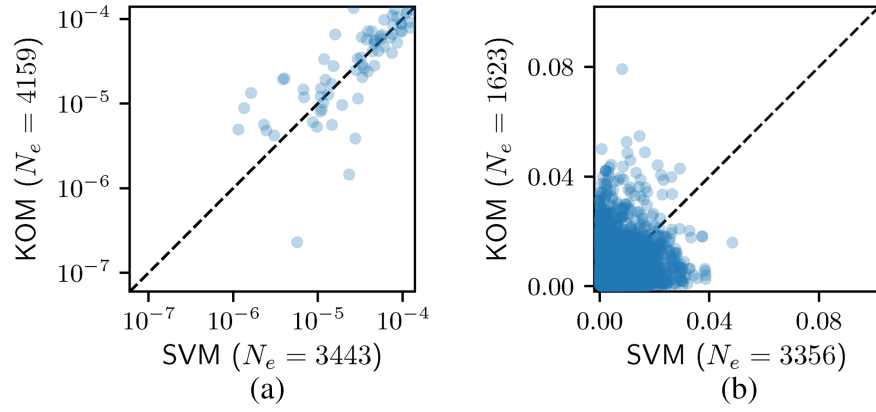
Compared against the small difference-in-means solution (leftmost column) for which the standardized difference-in-means for all covariates are below 0.1, CARD retains more observations in its selected subset, but SVM achieves a better covariate balance than CARD for most dimensions although there are some large imbalances. This is expected because SVM minimizes the overall covariate imbalance without a constraint on each dimension as in CARD. We observe a similar result when comparing CARD with the SVM solutions based on the closest effective sample size solution (left-middle) and the closest normed difference-in-means solution (right-middle). It is notable that the latter generally achieves a better covariate balance while retaining more observations than CARD. Finally, the results for the elbow SVM solution (rightmost column) show that tight balance is attainable with a moderate amount of sample pruning, with near exact balance in the linear setting. This level of covariate balance is difficult to achieve with CARD due to the infeasibility of optimization particularly in high-dimensional settings.

Figure 8 shows the dimensional balance comparisons for the KOM solution against the SVM solution. Here, we see that under the linear setting, KOM retains significantly more units than the SVM solution while attaining the same balance. This is due to the  $\sum_i \alpha W_i = 0$  constraint of SVM, which encourages the selected subset to have a roughly equal proportion of treated and control units, while the KOM solution allows the resulting subset

to be more imbalanced. Under the polynomial setting, however, we find that SVM retains significantly more units than KOM while achieving a similar degree of covariate balance, which is likely due to poor regularization parameter choice by the KOM algorithm.

Lastly, we compare the point estimates of the ATE, the weighted Neyman standard error, and the effective sample size for SVM, CARD (with linear and polynomial), KOM, and KCB (with RBF) in Table 1. We considered three different solutions from the SVM path:  $\text{SVM}_{\text{imbalance}}$ , which corresponds to the initial solution for which the balance constraint is most relaxed and  $\alpha_i = 1$ ,  $i \in \mathcal{T}$ ,  $\text{SVM}_{\text{balance}}$ , which corresponds to the most regularized solution with the best covariate balance on the path, and  $\text{SVM}_{\text{elbow}}$ , which corresponds to the solution occurring at the elbow of the balance-weight sum curves shown in Figure 6. When computing standard errors, we follow the convention and treat the weights as fixed [3, 6, 7, 14]. See Abadie and Spiess [43] for a recent investigation of post-matching standard errors.

The results show that SVM leads to a positive estimate in all cases, which agrees with the original finding reported by Connors et al. [27]. We also find that the three SVM solutions differ most significantly in their standard errors, which increases as the constraint on balance becomes stronger and the subset is more pruned, as shown in the effective sample size column. In particular, the heavily balanced  $\text{SVM}_{\text{balance}}$  solution under the RBF setting leads



**FIGURE 8** | Comparison of covariate standardized difference-in-means between SVM and kernel optimal matching (KOM) under the linear (left) and polynomial (right) settings. (a) Linear; (b) Polynomial.

**TABLE 1** | The estimated effect of right heart catheterization on death within 30 days after treatment. We compare the results based on cardinality matching (CARD) and Kernel Optimal Matching (KOM) with those based on three SVM solutions – the solution with the best covariate balance on the regularization path in terms of normed difference-in-means (SVM<sub>balance</sub>), the elbow solution (SVM<sub>elbow</sub>), and the solution to the worst covariate balance on the path (SVM<sub>imbalance</sub>).

Feature	Method	Estimate	Effective	
			Standard error	sample size
Linear	SVM <sub>balance</sub>	0.0615	0.0148	3442
	SVM <sub>elbow</sub>	—	—	—
	SVM <sub>imbalance</sub>	0.0479	0.0132	4367
	CARD	0.0335	0.0138	4174
	KOM	0.0656	0.0147	4159
Polynomial	SVM <sub>balance</sub>	0.0634	0.0279	1087
	SVM <sub>elbow</sub>	0.0588	0.0154	3325
	SVM <sub>imbalance</sub>	0.0541	0.0135	4375
	CARD	0.0313	0.0139	4084
	KOM	0.0452	0.0251	1623
RBF	SVM <sub>balance</sub>	0.0518	0.0289	1123
	SVM <sub>elbow</sub>	0.0527	0.0148	3444
	SVM <sub>imbalance</sub>	0.0474	0.0132	4378
	KCB	0.0337	0.0173	3306
	KOM	0.0582	0.0166	4185

to a 95% confidence interval which overlaps with zero. This is in contrast with the less balanced SVM<sub>elbow</sub> solution, which has both a larger effect estimate and a smaller confidence interval. This result demonstrates the necessity of computing the regularization path so that researchers may avoid low-quality solutions due to poor parameter choice.

Compared to other methods, we observe that KOM yields greater estimates of positive effects with comparable standard errors in the linear and RBF settings. However, under the polynomial setting, the standard error is much larger and the sample is

significantly more pruned than the modestly balanced SVM<sub>elbow</sub> solution. Both CARD and KCB produce smaller positive effect estimates, with the standard error for KCB leading to a 95% confidence interval which overlaps with zero.

## 5 | Concluding Remarks

In this paper, we show how support vector machines (SVMs) can be used to compute covariate balancing weights and estimate causal effects. We establish a number of interpretations of SVM as a covariate balancing procedure. First, the SVM dual problem computes weights that minimize the MMD while simultaneously maximizing effective sample size. Second, the SVM dual problem can be viewed as a continuous relaxation of the largest balanced subset problem. Lastly, similar to existing kernel balancing methods, SVM weights minimize the worst-case bias due to prognostic score imbalance. Additionally, path algorithms can be used to compute the entire set of SVM solutions as the regularization parameter varies, which constitutes a balance-sample size frontier.

Our work suggests several possible directions for future research. On the algorithmic side, a disadvantage of the proposed methodology is that it encourages a roughly equal effective number of treated and control units in the optimal subset, which can lead to unnecessary sample pruning. One could use weighted SVM [44] to address this problem, but existing path algorithms are applicable only to unweighted SVM. On the theoretical side, our results suggest a fundamental connection between the support vectors and the set of overlap. Steinwart [45] shows that the fraction of support vectors for a variant of the SVM discussed here asymptotically approaches the measure of this overlap set, suggesting that SVM may be used to develop a statistical test for the overlap assumption.

Finally, on the theoretical side, it is of interest to further investigate the choice of regularization parameter for SVM. We have shown that the SVM's regularization path corresponds to a balance-sample size frontier. One possible approach is to choose the value of the regularization parameter that minimizes the mean squared error of the estimated ATE. However, this approach requires the specification of an outcome model to

characterize the variance. On the practical side, the exploration of possible regression adjustment after SVM matching can be of interest as such an approach may lead to more robust causal inference [3, 43].

## Acknowledgments

We thank Brian Lee for making his simulation code available and Francesca Dominici, Chad Hazlett, Gary King, Jose Zubizarreta, and seminar participants at the Institute for Quantitative Social Science, Harvard University for helpful discussions. An anonymous reviewer of the Alexander and Diviya Magaro Peer Pre-Review Program for insightful comments.

## Conflicts of Interest

The authors declare no conflicts of interest.

## References

1. J. K. Lunceford and M. Davidian, "Stratification and Weighting via the Propensity Score in Estimation of Causal Treatment Effects: A Comparative Study," *Statistics in Medicine* 23, no. 19 (2004): 2937–2960, <https://doi.org/10.1002/sim.1903>.
2. D. B. Rubin, *Matched Sampling for Causal Effects* (Cambridge: Cambridge University Press, 2006).
3. D. E. Ho, K. Imai, G. King, and E. A. Stuart, "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference," *Political Analysis* 15, no. 3 (2007): 199–236.
4. E. A. Stuart, "Matching Methods for Causal Inference: A Review and a Look Forward," *Statistical Science* 25, no. 1 (2010): 1–21.
5. J. Hainmueller, "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis* 20 (2012): 25–46.
6. K. Imai and M. Ratkovic, "Covariate Balancing Propensity Score," *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 76, no. 1 (2014): 243–263.
7. J. R. Zubizarreta, "Stable Weights That Balance Covariates for Estimation With Incomplete Outcome Data," *Journal of the American Statistical Association* 110, no. 511 (2015): 910–922.
8. K. C. G. Chan, S. C. P. Yam, and Z. Zhang, "Globally Efficient Nonparametric Inference of Average Treatment Effects by Empirical Balancing Calibration Weighting," *Journal of the Royal Statistical Society: Series B: Methodological* 78 (2016): 673–700.
9. S. Athey, G. W. Imbens, and S. Wager, "Approximate Residual Balancing: Debiased Inference of Average Treatment Effects in High Dimensions," *Journal of the Royal Statistical Society: Series B: Methodological* 80, no. 4 (2018): 597–623.
10. F. Li, K. L. Morgan, and A. M. Zaslavsky, "Balancing Covariates via Propensity Score Weighting," *Journal of the American Statistical Association* 113, no. 521 (2018): 390–400, <https://doi.org/10.1080/01621459.2016.1260466>.
11. R. K. Wong and K. C. G. Chan, "Kernel-Based Covariate Functional Balancing for Observational Studies," *Biometrika* 105, no. 1 (2018): 199–213.
12. Q. Zhao, "Covariate Balancing Propensity Score by Tailored Loss Functions," *Annals of Statistics* 47, no. 2 (2019): 965–993.
13. C. Hazlett, "Kernel Balancing: A Flexible Non-parametric Weighting Procedure for Estimating Causal Effects," *Statistica Sinica* 30, no. 3 (2020): 1155–1189.
14. N. Kallus, "Generalized Optimal Matching Methods for Causal Inference," *Journal of Machine Learning Research* 21, no. 62 (2020): 1–54.
15. Y. Ning, S. Peng, and K. Imai, "Robust Estimation of Causal Effects via High-Dimensional Covariate Balancing Propensity Score," *Biometrika* 107, no. 3 (2020): 533–554.
16. Z. Tan, "Regularized Calibrated Estimation of Propensity Scores With Model Misspecification and High-Dimensional Data," *Biometrika* 107, no. 1 (2020): 137–158.
17. C. Cortes and V. Vapnik, "Support-Vector Networks," *Machine Learning* 20, no. 3 (1995): 273–297.
18. B. Schölkopf, A. J. Smola, and F. Bach, *Learning With Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (Cambridge, MA: MIT press, 2002).
19. A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola, "A Kernel Method for the Two-Sample-Problem," *Advances in Neural Information Processing Systems* 19 (2007): 513–520.
20. T. Hastie, S. Rosset, R. Tibshirani, and J. Zhu, "The Entire Regularization Path for the Support Vector Machine," *Journal of Machine Learning Research* 5, no. Oct (2004): 1391–1415.
21. C. Sentelle, G. Anagnostopoulos, and M. Georgiopoulos, "A Simple Method for Solving the SVM Regularization Path for Semidefinite Kernels," *IEEE Transactions on Neural Networks and Learning Systems* 27, no. 4 (2016): 709.
22. G. King, C. Lucas, and R. A. Nielsen, "The Balance-Sample Size Frontier in Matching Methods for Causal Inference," *American Journal of Political Science* 61, no. 2 (2017): 473–489.
23. M. Ratkovic, *Balancing Within the Margin: Causal Effect Estimation With Support Vector Machines* (Princeton, NJ: Department of Politics, Princeton University, 2014), <https://www.princeton.edu/~sim;ratkovic/public/BinMatchSVM.pdf>.
24. D. Ghosh, "Relaxed Covariate Overlap and Margin-Based Causal Effect Estimation," *Statistics in Medicine* 37, no. 28 (2018): 4252–4265.
25. G. King and L. Zeng, "The Dangers of Extreme Counterfactuals," *Political Analysis* 14, no. 2 (2006): 131–159.
26. J. R. Zubizarreta, R. D. Paredes, and P. R. Rosenbaum, "Matching for Balance, Pairing for Heterogeneity in an Observational Study of the Effectiveness of For-Profit and Not-For-Profit High Schools in Chile," *Annals of Applied Statistics* 8, no. 1 (2014): 204–231.
27. A. F. Connors, T. Speroff, N. V. Dawson, et al., "The Effectiveness of Right Heart Catheterization in the Initial Care of Critically Ill Patients," *Journal of the American Medical Association* 276, no. 11 (1996): 889–897.
28. D. B. Rubin, "Comments on 'on the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9' by J. Splawa-Neyman Translated From the Polish and Edited by D. M. Dabrowska and T. P. Speed," *Statistical Science* 5 (1990): 472–480.
29. Z. Tan, "Bounded, Efficient and Doubly Robust Estimation With Inverse Weighting," *Biometrika* 97, no. 3 (2010): 661–682.
30. Y. Zhu, J. S. Savage, and D. Ghosh, "A Kernel-Based Metric for Balance Assessment," *Journal of Causal Inference* 6, no. 2 (2018): 1–19.
31. A. Gretton, K. Fukumizu, C. Teo, L. Song, B. Schölkopf, and A. Smola, "A Kernel Statistical Test of Independence," *Advances in Neural Information Processing Systems* 20 (2007): 585–592.
32. A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A Kernel Two-Sample Test," *Journal of Machine Learning Research* 13, no. 1 (2012): 723–773.
33. B. K. Sriperumbudur, "Mixture Density Estimation via Hilbert Space Embedding of Measures," in *International Symposium on*

*Information Theory Proceedings* (St. Petersburg, Russia: IEEE, 2011), 1027–1030.

34. B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet, “Hilbert Space Embeddings and Metrics on Probability Measures,” *Journal of Machine Learning Research* 11 (2010): 1517–1561.

35. B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. Lanckriet, “On the Empirical Estimation of Integral Probability Metrics,” *Electronic Journal of Statistics* 6 (2012): 1550–1599.

36. J. D. Huling and S. Mak, “Energy Balancing of Covariate Distributions,” *Journal of Causal Inference* 12, no. 1 (2024): 20220029, <https://doi.org/10.1515/jci-2022-0029>.

37. B. B. Hansen, “The Prognostic Analogue of the Propensity Score,” *Biometrika* 95, no. 2 (2008): 481–488.

38. N. Kallus, B. Pennicooke, and M. Santacatterina, “More Robust Estimation of Average Treatment Effects Using Kernel Optimal Matching in an Observational Study of Spine Surgical Interventions,” *Statistics in Medicine* 40, no. 10 (2021): 2305–2320.

39. B. K. Lee, J. Lessler, and E. A. Stuart, “Improving Propensity Score Weighting Using Machine Learning,” *Statistics in Medicine* 29, no. 3 (2010): 337–346.

40. S. Setoguchi, S. Schneeweiss, M. A. Brookhart, R. J. Glynn, and E. F. Cook, “Evaluating Uses of Data Mining Techniques in Propensity Score Estimation: A Simulation Study,” *Pharmacoepidemiology and Drug Safety* 17, no. 6 (2008): 546–555.

41. O. L. Gurobi, “Gurobi Optimizer Reference Manual,” 2020.

42. G. W. Imbens and D. B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences* (New York, NY: Cambridge University Press, 2015).

43. A. Abadie and J. Spiess, “Robust Post-Matching Inference,” *Journal of the American Statistical Association* 117, no. 538 (2022): 983–995, <https://doi.org/10.1080/01621459.2020.1840383>.

44. C. F. Lin and S. D. Wang, “Fuzzy Support Vector Machines,” *IEEE Transactions on Neural Networks* 13, no. 2 (2002): 464–471.

45. I. Steinwart, “Sparseness of Support Vector Machines—Some Asymptotically Sharp Bounds,” *Advances in Neural Information Processing Systems* 16 (2004): 1069–1076.

46. W. Dinkelbach, “On Nonlinear Fractional Programming,” *Management Science* 13, no. 7 (1967): 492–498.

47. S. Schaible, “Fractional Programming. II, on Dinkelbach’s Algorithm,” *Management Science* 22, no. 8 (1976): 868–873.

## Supporting Information

Additional supporting information can be found online in the Supporting Information section.

## Appendix A

### Proof of Theorem 1

We prove this theorem by establishing several equivalents reformulations of the SVM dual problem. By equivalence, we mean that these problems differ from one another only in the scaling of the regularization parameter, implying that their regularization paths consist of the same set of solutions. More formally, given two optimization problems A1 and A2, we say that A1 and A2 are equivalent if a solution  $\alpha_1^*$  for A1 can be used to construct a solution for A2.

Denote the SVM weight set

$$\mathcal{A}_{\text{SVM}} = \left\{ \alpha \in \mathbb{R}^N : 0 < \alpha_i < 1, \sum_{i \in \mathcal{I}} \alpha_i = \sum_{j \in \mathcal{C}} \alpha_j \right\}$$

and consider a rescaled version of the SVM dual given in Equation (4), which we label as A1:

$$\begin{aligned} \min_{\alpha} \quad & \alpha^\top Q \alpha - \mu \mathbf{1}^\top \alpha \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{SVM}} \end{aligned} \quad (\text{A1})$$

Note that for a given  $\lambda$  and solution  $\alpha_*$  to the original problem defined in Equation (4) under  $\lambda$ ,  $\alpha_*$  is also a solution to the rescaled problem (A1) under  $\mu = 2\lambda$ . This establishes equivalence between these two problems.

We begin by proving the following lemma, which allows us to replace the squared seminorm term  $\alpha^\top Q \alpha$  with  $\sqrt{\alpha^\top Q \alpha}$  to obtain the problem

$$\begin{aligned} \min_{\alpha} \quad & \sqrt{\alpha^\top Q \alpha} - \nu \mathbf{1}^\top \alpha \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{SVM}} \end{aligned} \quad (\text{A2})$$

**Lemma 1.** *The problems given in Equations (A1) and (A2) are equivalent.*

*Proof.* This result follows from the strong duality of the SVM dual problem, which allows us to form the following equivalent problem in which the penalized term  $\mathbf{1}^\top \alpha$  is replaced with a hard constraint with threshold  $\epsilon$ :

$$\begin{aligned} \min_{\alpha} \quad & \alpha^\top Q \alpha \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{SVM}} \\ & \mathbf{1}^\top \alpha \geq \epsilon \end{aligned} \quad (\text{A3})$$

The solution to Equation (A3) is unchanged whether we minimize  $\alpha^\top Q \alpha$  or  $\sqrt{\alpha^\top Q \alpha}$ , so the problem

$$\begin{aligned} \min_{\alpha} \quad & \sqrt{\alpha^\top Q \alpha} \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{SVM}} \\ & \mathbf{1}^\top \alpha \geq \epsilon \end{aligned}$$

is identical to Equation (A3). By strong duality, we can again enforce the hard constraint on the term  $\mathbf{1}^\top \alpha$  through a penalized term with a new regularization parameter, which establishes the equivalence between Equations (A1) and (A2).  $\square$

Next, we consider the fractional program

$$\begin{aligned} \min_{\alpha} \quad & \frac{\sqrt{\alpha^\top Q \alpha}}{\mathbf{1}^\top \alpha / 2} \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{SVM}} \end{aligned} \quad (\text{A4})$$

The following lemma connects Equation (A4) to the reformulated SVM problem given in Equation (A2) through Dinkelbach’s method [46, 47]:

**Lemma 2.** (Dinkelbach, Theorem 1 [46]). *Suppose  $\alpha_* \in \mathcal{A}_{\text{SVM}}$  and  $\alpha_* \neq \mathbf{0}$ . Then*

$$q_* = \frac{\sqrt{\alpha_*^\top Q \alpha_*}}{\mathbf{1}^\top \alpha_* / 2} = \min_{\alpha \in \mathcal{A}_{\text{SVM}}} \frac{\sqrt{\alpha^\top Q \alpha}}{\mathbf{1}^\top \alpha / 2}$$

*if, and only if*

$$\min_{\alpha \in \mathcal{A}_{\text{SVM}}} \sqrt{\alpha^\top Q \alpha} - \frac{q_*}{2} \mathbf{1}^\top \alpha = \sqrt{\alpha_*^\top Q \alpha_*} - \frac{q_*}{2} \mathbf{1}^\top \alpha_* = 0$$

Thus, the solution to the rescaled SVM dual problem defined in Equation (A2) under  $\nu = q_*/2$  minimizes the fractional program given

in Equation (A4). Finally, we consider the MMD minimization problem defined in Equation (11),

$$\begin{aligned} \min_{\alpha} \quad & \sqrt{\alpha^\top Q \alpha} \\ \text{s.t.} \quad & \alpha \in \mathcal{A}_{\text{simplex}} \end{aligned} \quad (\text{A5})$$

The following lemma establishes equivalence between Equations (A4) and (A5) under the proper renormalization of the fractional program solution.

**Lemma 3.** Assume any solution  $\alpha_*$  to (A4) is such that  $\alpha_* \neq \mathbf{0}$ . Then the problems (A4) and (A5) are equivalent.

*Proof.* Let  $\alpha_4 \neq \mathbf{0}$  and  $\alpha_5$  be solutions to problems defined in Equations (A4) and (A5), respectively, and consider the vector-valued function  $f : \mathcal{A}_{\text{SVM}} \setminus \mathbf{0} \mapsto \mathcal{A}_{\text{simplex}}$ ,  $f(\alpha) = \alpha / (\mathbf{1}^\top \alpha / 2)$ , which normalizes the weights in the treated and control groups to each sum to 1. First note that since  $\mathcal{A}_{\text{simplex}} \subset \mathcal{A}_{\text{SVM}} \setminus \mathbf{0}$ ,  $\alpha_5$  is feasible for Equation (A4). Then by optimality of  $\alpha_4$ , we have

$$\frac{\sqrt{\alpha_4^\top Q \alpha_4}}{\mathbf{1}^\top \alpha_4 / 2} \leq \frac{\sqrt{\alpha_5^\top Q \alpha_5}}{\mathbf{1}^\top \alpha_5 / 2} = \sqrt{\alpha_5^\top Q \alpha_5}$$

Next, note that  $f(\alpha_4)$  is feasible for Equation (A5). Then by optimality of  $\alpha_5$ , we have

$$\sqrt{\alpha_5^\top Q \alpha_5} \leq \sqrt{f(\alpha_4)^\top Q f(\alpha_4)} = \frac{\sqrt{\alpha_4^\top Q \alpha_4}}{\mathbf{1}^\top \alpha_4 / 2}$$

In order for both of these inequalities to be true, we must have

$$\frac{\sqrt{\alpha_4^\top Q \alpha_4}}{\mathbf{1}^\top \alpha_4 / 2} = \frac{\sqrt{\alpha_5^\top Q \alpha_5}}{\mathbf{1}^\top \alpha_5 / 2} = \sqrt{\alpha_5^\top Q \alpha_5}$$

□

Note that the assumption  $\alpha \neq \mathbf{0}$  in Lemmas 2 and 3 holds when  $v \geq q_*/2$ , provided that the data does not consist of samples all belonging to the same class. To see this, note that when  $v = q_*/2$ , both  $\alpha = \mathbf{0}$  and any scaled version of the MMD-minimizing solution,  $\alpha = c\alpha_*$ ,  $c > 0$ , lead to an objective value of 0 in Equation (A2). Then since Equation (A2) is a strictly monotonically decreasing function of  $v$ , we must have  $\alpha \neq \mathbf{0}$  for  $v > q_*/2$ . Now since the regularization path defined in this work only considers  $v \geq q_*/2$ , we may safely assume  $\alpha \neq \mathbf{0}$ .

We are now ready to prove Theorem 1. Part (i): Lemma 1 establishes equivalence between the regularization paths for the rescaled SVM dual defined in Equations (A1) and (A2). In addition, Lemma 2 establishes the existence of  $v_*$  such that the solution to Equation (A2) under  $v_*$  is also a solution to Equation (A4). Then, it follows that there exists  $\lambda_*$  such that the solution to the rescaled SVM dual problem under  $\lambda_*$  minimizes Equation (A4). Finally, recall that Lemma 3 establishes that the minimizing solution to Equation (A4) is also a solution to the weighted MMD minimization problem. Therefore, there exists  $\lambda_*$  such that the solution to the SVM dual under  $\lambda_*$  minimizes the weighted MMD. Part (ii): The proof follows from Lemma 3 of Schaible [47].

## Appendix B

### Relation to Cardinality Matching and Stable Balancing Weights

Cardinality matching is an optimization procedure that maximizes the number of matches subject to a set of covariate balance constraints. The optimization problem for cardinality matching is given by,

$$\begin{aligned} \min_{m_{ij}} \quad & \sum_{i \in \mathcal{T}} \sum_{j \in \mathcal{C}} m_{ij} \\ \text{s.t.} \quad & \left| \sum_{i \in \mathcal{T}} \sum_{j \in \mathcal{C}} m_{ij} [f_b(X_{id}) - f_b(X_{jd})] \right| \end{aligned}$$

$$\begin{aligned} & \leq \varepsilon_{db} \sum_{i \in \mathcal{T}} \sum_{j \in \mathcal{C}} m_{ij}, & d = 1, \dots, D, \quad b = 1, \dots, B \\ & \sum_{j \in \mathcal{C}} m_{ij} \leq 1, \quad i \in \mathcal{T} \\ & \sum_{i \in \mathcal{T}} m_{ij} \leq 1, \quad j \in \mathcal{C} \\ & m_{ij} \in \{0, 1\}, \quad i \in \mathcal{T}, \quad j \in \mathcal{C} \end{aligned} \quad (\text{B1})$$

where  $m_{ij}$  are selection variables indicating whether treated unit  $i$  is matched to control unit  $j$ ,  $X_{id}$  denotes the  $d$ th element of covariate vector  $\mathbf{X}_i$ ,  $f_b$  is an arbitrary function of the covariates specifying each of the  $B$  balance conditions, and  $\varepsilon_{db}$  is a tolerance selected by a researcher. Common choices for  $f_b$  are the first- and second-order moments, and  $\varepsilon_{db}$  is typically set to a scalar multiple of the corresponding standardized difference-in-means.

To establish the connection between SVM and cardinality matching, we first note that cardinality matching need not be formulated as a matched pair optimization problem. In fact, the balance constraints between pairs, as formulated in Equation (B1), are equivalent to those between the treatment and control groups in the selected subsample. Similarly, the one-to-one matching constraints are equivalent to restricting the number of treated and control units in the selected subsample to be equal. Therefore, defining the indicator variable  $\alpha_i$  for selection into the optimal subset, we can rewrite the optimization problem for cardinality matching as,

$$\begin{aligned} \min_{\alpha} \quad & \frac{1}{2} \sum_{i=1}^N \alpha_i \\ \text{s.t.} \quad & \left| \sum_{i \in \mathcal{T}} \alpha_i f_b(X_{id}) - \sum_{j \in \mathcal{C}} \alpha_j f_b(X_{jd}) \right| \\ & \leq \frac{1}{2} \varepsilon_{db} \sum_{i=1}^N \alpha_i, & d = 1, \dots, D, \quad b = 1, \dots, B \\ & \sum_{i=0}^N \alpha_i W_i = 0, \\ & \alpha_i \in \{0, 1\}, & i = 1, \dots, N \end{aligned} \quad (\text{B2})$$

Comparing this problem to the SVM dual defined in Equation (4), we see two differences. First, cardinality matching restricts  $\alpha_i$  to be integer-valued, while SVM allows for  $\alpha_i$  to be continuous. Second, balance in the optimal subset is enforced differently in each method. Cardinality matching imposes covariate-specific balance by bounding each dimension's difference-in-means, while SVM imposes aggregated balance by penalizing the normed difference-in-means. The preference between these two measures of balance may in part depend on the dimensionality of the covariates and a priori knowledge about confounding mechanisms. If we suspect certain covariates to be strong confounders, then bounding those specific dimensions may be reasonable. However, if no such information is available and the covariate space is high-dimensional, then restricting the overall balance may be preferable.

Closely related to cardinality matching are stable balancing weights (SBW), which is a weighting method that aims to minimize the dispersion of the weights subject to a set of balance conditions. The optimization problem for SBW is given by

$$\begin{aligned} \min_{\alpha} \quad & \|\alpha\|^2 \\ \text{s.t.} \quad & \left| \sum_{i \in \mathcal{T}} \alpha_i f_b(X_{id}) - \sum_{j \in \mathcal{C}} \alpha_j f_b(X_{jd}) \right| \\ & \leq \varepsilon_{db}, \quad d = 1, \dots, D, \quad b = 1, \dots, B \\ & \alpha \in \mathcal{A}_{\text{simplex}} \end{aligned} \quad (\text{B3})$$

where  $f_b(\cdot)$  and  $\varepsilon_{db}$  denote the same quantities defined in cardinality matching.

In order to connect SBW to SVM, we consider L2-SVM, a variant of SVM obtained by replacing  $\xi_i$  with  $\xi_i^2$  in the objective function of the SVM primal problem is given in Equation (3). L2-SVM has a corresponding dual form

$$\begin{aligned} \min_{\alpha} \quad & \left( \frac{1}{2} \alpha^\top Q \alpha - \mathbf{1}^\top \alpha \right) + \frac{\lambda}{2} \|\alpha\|^2 \\ \text{s.t.} \quad & W^\top \alpha = 0 \\ & \alpha > 0 \end{aligned} \quad (\text{B4})$$

Here, the term contained in parentheses relates to the MMD, while the second term relates to the dispersion of the renormalized weights  $\alpha/(1^\top \alpha/2)$ . As before, the main difference between L2-SVM and SBW lie in the way balance is enforced, with the former encouraging aggregate balance through normed difference-in-means and the latter using dimension-specific balance through a constraint.

## Appendix C

### Conditional Bias With Respect to SATE and SATT

In this section, we derive the conditional bias for the weighted difference-in-means estimator. Note that our derivation follows the one given in [38]. Consider the problem of estimating the SATE and SATT, defined as

$$\tau_{\text{SATE}} = \frac{1}{N} \sum_{i=1}^N Y_i(1) - Y_i(0) \text{ and } \tau_{\text{SATT}} = \frac{1}{n_T} \sum_{i \in T} Y_i(1) - Y_i(0)$$

respectively. We denote the weighted estimator  $\hat{\tau}$ , which has a form

$$\hat{\tau} = \sum_{i \in T} \alpha_i Y_i - \sum_{i \in C} \alpha_i Y_i$$

where  $\alpha_i \in \mathcal{A}_{\text{simplex}}$ . The conditional bias with respect to the SATE is given by

$$\begin{aligned} \mathbb{E}[\hat{\tau} - \tau_{\text{SATE}} | \mathbf{X}_{1:N}, T_{1:N}] &= \sum_{i \in T} \mathbb{E}[\alpha_i Y_i | \mathbf{X}_{1:N}, T_{1:N}] \\ &\quad - \sum_{i \in C} \mathbb{E}[\alpha_i Y_i | \mathbf{X}_{1:N}, T_{1:N}] - \mathbb{E}[\tau_{\text{SATE}} | \mathbf{X}_{1:N}, T_{1:N}] \\ &= \sum_{i=1}^N \alpha_i [T_i - (1 - T_i)] \mathbb{E}[Y_i(T_i) | \mathbf{X}_i, T_i] \\ &\quad - \frac{1}{N} \sum_{i=1}^N \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i, T_i] \\ &= \sum_{i=1}^N \alpha_i [T_i - (1 - T_i)] \mathbb{E}[Y_i(T_i) | \mathbf{X}_i] \\ &\quad - \frac{1}{N} \sum_{i=1}^N \mathbb{E}[Y_i(1) - Y_i(0) | \mathbf{X}_i] \\ &= \sum_{i=1}^N \alpha_i T_i f_1(\mathbf{X}_i) - \sum_{i=1}^N \alpha_i (1 - T_i) f_0(\mathbf{X}_i) - \frac{1}{N} \sum_{i=1}^N \tau(\mathbf{X}_i) \\ &= \sum_{i=1}^N \alpha_i T_i [f_0(\mathbf{X}_i) + \tau(\mathbf{X}_i)] - \sum_{i=1}^N \alpha_i (1 - T_i) f_0(\mathbf{X}_i) \\ &\quad - \frac{1}{N} \sum_{i=1}^N \tau(\mathbf{X}_i) \\ &= \sum_{i=1}^N (\alpha_i T_i - N^{-1}) \tau(\mathbf{X}_i) + \sum_{i=1}^N \alpha_i [T_i - (1 - T_i)] f_0(\mathbf{X}_i) \end{aligned}$$

$$= \sum_{i=1}^N (\alpha_i T_i - N^{-1}) \tau(\mathbf{X}_i) + \sum_{i=1}^N \alpha_i W_i f_0(\mathbf{X}_i)$$

where the second equality follows from SUTVA, and the third equality follows from Assumptions 1 and 2. By a similar argument, the conditional bias with respect to the SATT is

$$\begin{aligned} \mathbb{E}[\hat{\tau} - \tau_{\text{SATT}} | \mathbf{X}_{1:N}, T_{1:N}] &= \sum_{i \in T} \mathbb{E}[\alpha_i Y_i | \mathbf{X}_{1:N}, T_{1:N}] - \sum_{i \in C} \mathbb{E}[\alpha_i Y_i | \mathbf{X}_{1:N}, T_{1:N}] \\ &\quad - \mathbb{E}[\tau_{\text{SATT}} | \mathbf{X}_{1:N}, T_{1:N}] \\ &= \sum_{i \in T} (\alpha_i - n_T^{-1}) \tau(\mathbf{X}_i) + \sum_{i=1}^N \alpha_i W_i f_0(\mathbf{X}_i) \end{aligned}$$

## Appendix D

### Worst-Case Bias in an RKHS

We consider the problem of minimizing the bias due to prognostic score imbalance, defined in Equation (15). Restricting  $f_0$  to the unit-ball RKHS, defined as  $\mathcal{F}_K$  in Section 2.3, and considering the  $f_0$  which maximizes the absolute value of this quantity, we compute the worst-case squared bias due to prognostic score imbalance as

$$B^2(\alpha; \mathbf{X}_{1:N}, T_{1:N}) = \sup_{f_0 \in \mathcal{F}_K} \left( \sum_{i=1}^N \alpha_i W_i f_0(\mathbf{X}_i) \right)^2$$

We can simplify this expression by

$$\begin{aligned} B^2(\alpha; \mathbf{X}_{1:N}, T_{1:N}) &= \sup_{f_0 \in \mathcal{F}_K} \left( \sum_{i=1}^N \alpha_i W_i f_0(\mathbf{X}_i) \right)^2 \\ &= \sup_{f_0 \in \mathcal{F}_K} \left( \sum_{i=1}^N \alpha_i W_i \langle f_0, \phi(\mathbf{X}_i) \rangle \right)^2 \\ &= \sup_{f_0 \in \mathcal{F}_K} \left( \left\langle f_0, \sum_{i=1}^N \alpha_i W_i \phi(\mathbf{X}_i) \right\rangle \right)^2 \\ &= \left\| \sum_{i=1}^N \alpha_i W_i \phi(\mathbf{X}_i) \right\|_{\mathcal{H}_K}^2, \\ &= \gamma_K^2 \left( \hat{F}_\alpha, \hat{G}_\alpha \right) \end{aligned}$$

where the second line follows from the reproducing property of the RKHS, the third line follows from the bilinearity of inner products, and the fourth line follows from the Cauchy-Schwarz inequality.