



Statistical Learning for Individualized Asset Allocation

Yi Ding^a, Yingying Li^b, and Rui Song^c

^aFaculty of Business Administration, University of Macau, Taipa, Macau; ^bDepartment of ISOM and Department of Finance, Hong Kong University of Science and Technology, Kowloon, Hong Kong; ^cDepartment of Statistics, North Carolina State University, Raleigh, NC

ABSTRACT

We establish a high-dimensional statistical learning framework for individualized asset allocation. Our proposed methodology addresses continuous-action decision-making with a large number of characteristics. We develop a discretization approach to model the effect of continuous actions and allow the discretization frequency to be large and diverge with the number of observations. We estimate the value function of continuous-action using penalized regression with our proposed generalized penalties that are imposed on linear transformations of the model coefficients. We show that our proposed Discretization and Regression with generalized folded concaVe penalty on Effect discontinuity (DROVE) approach enjoys desirable theoretical properties and allows for statistical inference of the optimal value associated with optimal decision-making. Empirically, the proposed framework is exercised with the Health and Retirement Study data in finding individualized optimal asset allocation. The results show that our individualized optimal strategy improves the financial well-being of the population. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received March 2021 Accepted August 2022

KEYWORDS

Continuous-action decision-making; High-dimensional statistical learning; Individualization; Penalized regression

1. Introduction

With the rapid development of artificial intelligence, particularly machine learning, a revolution is underway in individualization. In precision medicine, thousands or millions of genetic characteristics are taken into consideration to determine the optimal treatment for an individual patient. Large retail corporations analyze massive datasets of the behavior and personal characteristics of customers to tailor their offerings to individual customers.

Individualization is not new in asset allocation. For example, before providing financial advice, investment companies ask their customers to answer questionnaires in order to determine their customers' risk profiles. An asset allocation strategy would then be recommended in the form of, for instance, a given proportion of stocks versus bonds that the customer should hold in their portfolio. However, often in these questionnaires, many of the questions are subjective and/or hypothetical. Incorporating objective data of the investors can ameliorate the weakness of the reliance of such personalization on subjective opinion.

In this study, we develop a statistical learning framework for individualized asset allocation. We focus on finding an individualized optimal proportion of wealth that should be invested in stocks for a consumption-based utility optimization problem. The essence of the problem is decision-making, where the action can take value on a continuous set. More broadly, our study provides a theoretical basis and practical showcase for continuous-action decision-making research, such as dose deci-

sion in precision medicine, movement angle in robotic control, and campaign duration in personalized marketing.

1.1. Related Literature

Methodologically, our study is closely related to the studies seeking the optimal individualized treatment rule (ITR) in fields such as personalized medicine. Q-learning and A-learning are the most popular methods for finding the optimal ITR. Q-learning (Watkins 1989) models the treatment responses, and A-learning (Murphy 2003; Robins 2004) models the contrast (regret) function; see, for example, Qian and Murphy (2011), Zhao et al. (2012), Shi et al. (2018), Athey and Wager (2019), and Zhu, Zeng, and Song (2019). These studies focus primarily on the problem in which the treatment comes from a fixed number of discrete levels, typically binary treatment.

There is a growing interest in the study of continuous-action decision-making where the treatment comes from a continuous set. Laber and Zhao (2015) propose a direct tree-based optimal rule searching method. Chen, Zeng, and Kosorok (2016) propose a direct nonparametric method by extending the outcomeweighted learning method of Zhao et al. (2012). More recently, Cai et al. (2020) propose an off-policy evaluation method by adaptively discretizing the action space using a deep jump Q-learning; Zhu et al. (2020) study a kernel assisted optimal dose rule method; while Zhou, Zhu, and Zeng (2021) propose a dimension reduced kernel approximation method. These stud-

ies focus on the estimation of optimal policy or the effect of continuous policy with low-dimensional covariates.

Theoretically, our study relates to the literature on penalized regression, especially in high-dimensional settings. The wellknown approaches of penalized regression include the least absolute shrinkage and selection operator (LASSO, Tibshirani 1996), smoothly clipped absolute deviation (SCAD, Fan and Li 2001), and minimax concave penalty (MCP, Zhang 2010). The theoretical properties and implementation of the penalized regression estimator under high-dimensional settings are investigated by Meinshausen and Bühlmann (2006), Zhao and Yu (2006), Zhang and Huang (2008), Lv and Fan (2009), Fan and Lv (2011), Wang, Kim, and Li (2013), among others. These studies investigate the problem with penalties imposed directly on model coefficients. For the generalized penalties that are imposed on a linear transformation of coefficients, Tibshirani and Taylor (2011) and Arnold and Tibshirani (2016) investigate the generalized penalty problems using lasso penalties (generalized lasso) and their studies focus mainly on the computational aspect. She (2010) discusses the sparsity recovery property of the generalized lasso estimator when the number of variables is

In economics and finance, our study relates to the literature on household finance for the elderly and studies that use Health and Retirement Study (HRS) and Consumption and Activities Mail Survey (CAMS) data. Using HRS and CAMS data, the "retirement savings puzzle" (Banks, Blundell, and Tanner 1998; Palumbo 1999) is studied by Hurd and Rohwedder (2003) and Haider and Stephens, Jr (2007). Rosen and Wu (2004), Hong, Kubik, and Stein (2004), and De Nardi and Yang (2014) study the problem of asset allocation modeling. By analyzing HRS and CAMS data, Engen et al. (1999) and Munnell, Webb, and Golub-Sass (2012) underscore a universal inadequacy of postretirement savings.

1.2. Main Challenges and Our Contributions

Developing methodologies for continuous-action decision making is challenging. Today, it is common to collect a large number of characteristics and incorporate them into the decision-making process. When there are many action levels, the dimensionality of the problem becomes even larger, posing challenges in identifying informative characteristics and estimating the model. The policy value, especially the optimal value, is an important target of inference for policy evaluation. For continuous-action decision-making, there are infinite possible options in the search for the optimal policy, posing substantial challenges to the statistical inference of the optimal value.

Empirically, substantial challenges come from the data. The HRS and CAMS survey data contain noisy and incomplete observations. The observed stock ratio data are skewed and concentrated around zero. In addition, the important utility variable is unobserved. Data preparation needs to be carefully done in order to obtain sufficient high-quality data for the statistical learning.

The main contributions of our paper are as follows.

First, in the field of individualized decision-making, we develop a high-dimensional statistical learning framework to study continuous-action decision-making. Specifically, we develop a discretization approach to model the effect of continuous actions and allow the discretization level to be large and diverge with the number of observations. The size of the discretization level is carefully analyzed in order to balance the approximation accuracy of the discretization with the compatibility in the penalized regression framework.

Second, we propose a Discretization and Regression with generalized fOlded concaVe penalty on Effect discontinuity (DROVE) method to estimate the model for continuous-action decision-making, whose design accommodates the continuous nature of the actions. Different from the standard penalized regression that penalizes the coefficients directly, DROVE penalizes the effect differences between adjacent action levels. Our estimator enjoys model coefficient estimation consistency.

Third, our novel approach provides valid statistical inference for the optimal value of continuous action when there are a large number of characteristics. We obtain the central limit theorem for the proposed DROVE estimator of the model coefficients as well as the value that is associated with optimal decisionmaking. To the best of our knowledge, our study is the first to achieve optimal value inference for continuous action, especially under a setting with high-dimensional characteristics.

Last but not least, in the fields of economics and finance, as a pioneer work, our study provides an individualized asset allocation using a high-dimensional statistical learning method that processes personal characteristic information. To address the data challenges, we use trajectory path models, which allow us to generate pseudo consumption and income for randomly assigned stock ratios. We then obtain the utility from the pseudo consumption paths and individual risk aversion models. Our empirical exercise with HRS and CAMS data shows that our individualized optimal asset allocation strategy substantially improves the financial well-being of the population and surpasses benchmark strategies that assign fixed stock ratios to all households. The superior performance of our method demonstrates the importance of individualization in asset allocation.

The article proceeds as follows. We present the statistical learning framework in Section 2 and develop the theoretical properties in Section 3. Sections 4 and 5 are devoted to simulation and empirical studies, respectively. Section 6 contains concluding remarks.

2. Statistical Learning Framework for **Continuous-Action Decision-Making**

2.1. Model Setup

Suppose that we have *n* observations $(X_i, A_i, Y_i)_{i=1,...,n}$, where $X_i \in \mathcal{X}$ is a length-d vector of covariates, A_i is the action that comes from a continuous support A, and Y_i is a random outcome. Under the potential outcome framework (Rubin 1974), we denote Y^a as the potential outcome that would have been observed under action level $a \in \mathcal{A}$. Following the literature of causal inference (e.g., Robins 2004), we make the stable unit treatment value assumption (SUTVA) that $Y = Y^a$ if the action A = a (consistency). In addition, we consider the randomiza-



tion assumption that A is independent of X and the potential outcomes (e.g., Murphy 2005), and the positivity assumption on the density function of $A \in \mathcal{A}$ (e.g., Chen, Zeng, and Kosorok 2016). A deterministic policy π maps the features space to the action space, $\pi: \mathcal{X} \to \mathcal{A}$. Under the SUTVA assumption, the optimal policy π^* determines the optimal action that maximizes the expected reward given the characteristics X, that is, $\pi^*(x) = \arg\max_a E(Y^a|X=x,A=a) = \arg\max_a E(Y|X=x,A=a)$.

Under the above framework, a particular example of interest is finding the optimal individualized asset allocation. For individualized asset allocation, **X** is individual characteristics, such as the financial and demographic status; Y is the economic reward or utility; A is the proportion of total wealth invested in stocks, that is, stock ratio, which can be continuous between 0 and 1; and $\pi^*(\cdot): \mathbb{R}^d \to [0,1]$ is the investment decision rule that yields the optimal stock ratio given the individual characteristics.

We study the optimal continuous-action decision-making by modeling the conditional expected reward as a value function Q(x, a) = E(Y|X = x, A = a). We consider the following model

$$Y = Q_n(X, A) + \varepsilon, \quad Q_n(X, A) = \psi_0^T X + \sum_{k=2}^{L_n} \psi_k^T X \cdot 1_{\{A \in [A_{(k)}, A_{(k+1)})\}},$$
(2.1)

where $Q_n(\mathbf{X},A)$ is the value function and ε is the noise term that is independent of \mathbf{X} , $\{0=A_{(1)}<\cdots< A_{(L_n)}=1\}$ is a series of grids, $L_n\leq n$ is the number of grids that can be large and increase with sample size n, and $\psi_0,\psi_2,\cdots,\psi_{L_n}$, are d-dimensional vectors, $\psi_1=\mathbf{0}$ and $A_{(L_n+1)}>1$. The two components, $\psi_0^T\mathbf{X}$ and $\sum_{k=2}^{L_n}\psi_k^T\mathbf{X}\cdot\mathbf{1}_{\{A\in[A_{(k)},A_{(k+1)})\}}$, represent the main effect of characteristics and the treatment effect of the action on individual characteristics, respectively.

We write $\check{\mathbf{X}}_i = \left((\mathbf{X}_i)^T, (\mathbf{X}_i)^T \cdot \mathbf{1}_{\{A_i \in [A_{(2)}, A_{(3)})\}}, \dots, (\mathbf{X}_i)^T \cdot \mathbf{1}_{\{A_i \in [A_{(L_n)}, A_{(L_n+1)})\}} \right)^T$, and $\beta_n = (\beta_1, \dots, \beta_p)^T = (\psi_0^T, \psi_2^T, \dots, \psi_{L_n}^T)^T$ as a length-p vector of the coefficients, where $p = d \times L_n$. Hence, the value function can be represented as $Q_n(\mathbf{X}_i, A_i) = \beta_n^T \check{\mathbf{X}}_i$. We allow p to be large, such that $\log p = o(n)$. In particular, we include the case where characteristics space can be large, so that $\log d = o(n)$. As to the discretization level L_n , it can grow slowly with n so that the approximation error in the working model diminishes to zero as $n \to \infty$. We discuss the growth rate of L_n in more detail in Section 2.2. Under the working model (2.1), we denote the true population parameter as $\beta_n^* = (\psi_0^{*T}, \psi_2^{*T}, \dots, \psi_{L_n}^{*T})^T$.

2.2. Penalized Regression with Generalized Penalties

In order to estimate the coefficients β_n^{\star} in the high-dimensional Q-function (2.1), we develop a novel penalized regression formulation, which accommodates the fact that decisions are made from a continuum. Intuitively, when the discretization level L_n is large, the distance between adjacent action levels $|A_{(k)} - A_{(k+1)}|$ becomes small and so is the difference in their effects, $(|\psi_{k,j} - \psi_{k+1,j}|)_{1 \leq j \leq d}$. We impose penalties on $|\psi_{k,j} - \psi_{k+1,j}|$ to shrink the difference in effect between two adjacent decisions, $A_{(k)}$ and $A_{(k+1)}$, on each covariate $j=1,\ldots,d$.

Formally, we write the following penalized regression:

$$\min_{\boldsymbol{\beta}_{n} = (\boldsymbol{\psi}_{0}^{T}, \dots, \boldsymbol{\psi}_{L_{n}}^{T})^{T}} \sum_{i=1}^{n} (Y_{i} - \boldsymbol{\beta}_{n}^{T} \check{\mathbf{X}}_{i})^{2} + n \sum_{k=0}^{L_{n}} p_{\lambda_{n}}(\boldsymbol{\psi}_{k}) + n \sum_{k=2}^{L_{n}-1} \sum_{i=1}^{d} p_{\lambda_{n}}(\boldsymbol{\psi}_{k,j} - \boldsymbol{\psi}_{k+1,j}),$$
(2.2)

where the penalty function is $p_{\lambda_n}(\xi) := \sum_{i=1}^h p_{\lambda_n}(\xi_i)$ for any vector $\xi = (\xi_1, \xi_2, \dots, \xi_h)^T$, and $p_{\lambda_n}(\cdot)$ is a penalty function with tuning parameter λ_n . The first penalties $\sum_{k=0}^{L_n} p_{\lambda_n}(\psi_k)$ penalize the main effect and the treatment effect, similar to the binary case; see, for example, Zhu, Zeng, and Song (2019). The second penalties $\sum_{k=2}^{L_n-1} \sum_{j=1}^d p_{\lambda_n}(\psi_{k,j} - \psi_{k+1,j})$ penalize the discontinuity.

Remark 2.1. The penalties, $(p_{\lambda_n}(\psi_{k,j}-\psi_{k+1,j}))_{2\leq k\leq L_n-1, 1\leq j\leq d'}$ generalize the idea of fused lasso (Tibshirani et al. 2005), which imposes penalties on the difference in adjacent coordinates $|\beta_i-\beta_{i+1}|$. She (2010), Tibshirani and Taylor (2011) and Arnold and Tibshirani (2016) study the generalized lasso problem, which is a penalized regression as formulated in (2.3) with the lasso penalty function for $p_{\lambda}(\cdot)$. In contrast to these works, we propose a generalized folded concave penalty and investigate the statistical properties of the estimator under the high-dimensional setting. In the literature on individualized decision making, to the best of our knowledge, it is the first time that regression with a generalized folded concave penalized is formulated in studying the effect of continuous actions.

The proposed penalized regression (2.2) can be categorized as one with *generalized* penalties, where penalization is imposed on linear transformations of the coefficients, $\mathbf{D}\boldsymbol{\beta}_n$, for some $K \times p$ matrix $\mathbf{D} = (\mathbf{d}_1, \dots, \mathbf{d}_K)^T$:

$$\min_{\boldsymbol{\beta}_n} \frac{1}{n} \sum_{i=1}^n (Y_i - \boldsymbol{\beta}_n^T \check{\mathbf{X}}_i)^2 + p_{\lambda_n}(\mathbf{D}\boldsymbol{\beta}_n). \tag{2.3}$$

More generally, we consider the following generalized linear model (Fan and Lv 2011) where the density function $f(Y; \check{X}, \beta_n)$ satisfies

$$f(\mathbf{Y}; \check{\mathbf{X}}, \boldsymbol{\beta}_n) = \prod_{i=1}^n f_0(Y_i; \boldsymbol{\beta}_n^T \check{\mathbf{X}}_i)$$

$$= \prod_{i=1}^n \exp\left\{\frac{Y_i \boldsymbol{\beta}_n^T \check{\mathbf{X}}_i - b(\boldsymbol{\beta}_n^T \check{\mathbf{X}}_i)}{\phi}\right\} c(Y_i, \phi), \quad (2.4)$$

where $\beta_n = (\beta_1, \dots, \beta_p)^T$ is a $p \times 1$ vector of regression coefficients, $\phi \in (0, \infty)$ is the nuisance parameter of dispersion, $\check{\mathbf{X}} = (\check{\mathbf{X}}_1, \check{\mathbf{X}}_2, \dots, \check{\mathbf{X}}_n)^T$, $b(\cdot)$ is twice continuously differentiable with $b''(\cdot) > 0$, and $c(Y_i, \phi)$ is the base measure that represents the density function for Y_i when $\beta_n = \mathbf{0}$, for example, $c(Y_i, \phi) = (\pi\phi)^{-1/2} \exp(-Y_i^2/\phi)$ for normal density. For a given matrix \mathbf{D} , the penalized likelihood function with generalized penalties is

$$\mathcal{L}_{n}(\boldsymbol{\beta}_{n}, \lambda_{n}, \mathbf{D}) = l_{n}(\boldsymbol{\beta}_{n}) - \sum_{k=1}^{K} p_{\lambda_{n}}(\mathbf{d}_{k}^{T}\boldsymbol{\beta}_{n})$$

$$= \frac{1}{n} (\mathbf{Y}^{T} \check{\mathbf{X}} \boldsymbol{\beta}_{n} - \mathbf{1}^{T} \mathbf{b} (\check{\mathbf{X}} \boldsymbol{\beta}_{n})) - \sum_{k=1}^{K} p_{\lambda_{n}}(\mathbf{d}_{k}^{T} \boldsymbol{\beta}_{n}), \quad (2.5)$$

where $l_n(\boldsymbol{\beta}_n) = (\mathbf{Y}^T \check{\mathbf{X}} \boldsymbol{\beta}_n - \mathbf{1}^T \mathbf{b} (\check{\mathbf{X}} \boldsymbol{\beta}_n))/n$ and $\mathbf{b} (\check{\mathbf{X}} \boldsymbol{\beta}_n) = (b(\boldsymbol{\beta}_n^T \check{\mathbf{X}}_1), \dots, b(\boldsymbol{\beta}_n^T \check{\mathbf{X}}_n))^T$.

 $(b(\boldsymbol{\beta}_n^T \check{\mathbf{X}}_1), \dots, b(\boldsymbol{\beta}_n^T \check{\mathbf{X}}_n))^T$. We next write $\mathbf{D} = (\mathbf{D}_{\text{signal}}^T, \mathbf{D}_{\text{null}}^T)^T$, where $\mathbf{D}_{\text{signal}}$ is a $K_1 \times p$ matrix, \mathbf{D}_{null} is a $K_0 \times p$ matrix, and $K = K_0 + K_1$. Suppose that the true coefficients, $\boldsymbol{\beta}_n^{\star} = (\boldsymbol{\beta}_1^{\star}, \dots, \boldsymbol{\beta}_p^{\star})^T$, satisfy $\mathbf{D}_{\text{null}} \boldsymbol{\beta}_n^{\star} = \mathbf{0}$, that is, $\mathbf{d}_k^T \boldsymbol{\beta}_n^{\star} = \mathbf{0}$, for $k = K_1 + 1, K_1 + 2, \dots, K$. In terms of $\mathbf{D}_{\text{signal}}$, we impose no constraint on its shape or rank.

For any positive semidefinite matrix $\mathbf{A} = (a_{ij})$, we define $\|\mathbf{A}\|_2 = \max_{\|\mathbf{x}\|_2 \le 1} \|\mathbf{A}\mathbf{x}\|_2$ and $\|\mathbf{x}\|_2 = \sqrt{\sum x_i^2}$ for any vector $\mathbf{x} = (x_i)$.

We impose the following assumptions on **D** and $\rho(t,\lambda) = p_{\lambda}(t)/\lambda$.

Assumption 2.1. $\zeta_{\min}^+(\mathbf{D}_{\text{null}}\mathbf{D}_{\text{null}}^T) \geq c$ and $\max_{1 \leq k \leq K} \|\mathbf{d}_k\|_2 \leq C$ for some constants c, C > 0, where $\zeta_{\min}^+(\mathbf{D}_{\text{null}}\mathbf{D}_{\text{null}}^T)$ denotes the smallest nonzero eigenvalue of $\mathbf{D}_{\text{null}}\mathbf{D}_{\text{null}}^T$.

Remark 2.2. Assumption 2.1 is met by our design of D in (2.2) as $\zeta_{\min}^+(\mathbf{D}_{\text{null}}\mathbf{D}_{\text{null}}^T) \geq 1$ and $\max_{1 \leq k \leq K} \|\mathbf{d}_k\|_2 \leq \sqrt{2}$.

Assumption 2.2. (1) $\rho(t,\lambda)$ is increasing and concave in $t \in [0,\infty)$; (2) $\rho(t,\lambda)$ is differentiable in $t \in (0,\infty)$ with $\rho'(0+,\lambda) > 0$; and (3) if $\rho'(t,\lambda)$ is dependent on λ , $\rho'(t,\lambda)$ is increasing in $\lambda \in (0,\infty)$ and $\rho'(0+)$ is independent of λ .

Remark 2.3. Assumption 2.2 describes the characteristics of a folded concave function class; see, for example, Lv and Fan (2009) and Fan and Lv (2011). Popular examples of folded concave functions include SCAD (Fan and Li 2001) and MCP (Zhang 2010). Our numerical examples uses SCAD as representative of the folded concave penalty function family.

We define $s_n = p - \text{rank}(\mathbf{D}_{\text{null}})$. We regard s_n as the nonsparsity coefficient, which is a natural extension of the nonsparsity coefficient for standard penalized regression when D is an identity matrix. The definition of s_n is in line with that of the degrees of freedom for generalized lasso (Tibshirani and Taylor 2011, 2012). The true model is considered to be sparse in the sense that $s_n \ll n$. The sparsity of the model comes from two sources. One source of sparsity is the large covariate space in which there can be many irrelevant variables. In real applications, for example, investment agencies collect a large number of covariates, although only a small proportion of the variables are useful. The sparsity that comes from high-dimensional covariates is the usual notion of model sparsity discussed in the literature. In addition, we assume sparsity on the treatment effect difference between adjacent action levels. This is essentially to assume that the effect of continuous action exhibits smoothness; hence, when L_n is large, many adjacent action levels have roughly the same treatment effect. This is sensible in real applications. For example, in asset allocation, when there are small changes in stock ratio, utility would not vary drastically.

Let $g_n = 2^{-1} \min\{|\mathbf{d}_j^T \boldsymbol{\beta}_n^*|, \mathbf{d}_j^T \boldsymbol{\beta}_n^* \neq 0\}$ be half of the minimum signal. We impose the following assumptions.

Assumption 2.3. $g_n \gg \lambda_n \gg \max(\sqrt{s_n/n}, \sqrt{(\log p)/n}), \max(s_n, \log p) = o(n), \text{ and } p'_{\lambda_n}(g_n) = o\left(\min(n^{-1/2}s_n^{-1/2}, n^{-1/2}K_1^{-1}s_n^{1/2})\right)$

Assumption 2.3 states that the minimal signal should be sufficiently large to be distinguishable from the noise. If g_n decreases as L_n grows, the minimal signal condition would constrain the number of discretization levels that we are able to handle. Asymptotically, L_n can be $O(n^{\varsigma})$ for some $\varsigma < 1/3$. In Appendix A of the supplementary materials, we discuss the rate of L_n under a varying coefficient example. We illustrate the choice of L_n with our practical example in Section 4.1. Additional regularity conditions (Assumptions C.1–C.4) are in Appendix C of the supplementary materials.

We summarize the proposed approach for continuous-action decision-making: Discretization and Regression with generalized fOlded concaVe penalty on Effect discontinuity (DROVE) as follows.

Algorithm DROVE

Step I. Discretize the action support into a series of grids $A_{(1)} < \cdots < A_{(L_n)}$ with L_n growing with n, and $L_n = O(n^{\varsigma})$ for some $\varsigma < 1/3$.

Step II. Perform the penalized regression with the generalized folded concave penalty on effect discontinuity, (2.2), and obtain the estimated value function $\widehat{Q}_n(X, A)$. The GLLN algorithm to be introduced in Section 3.1 can be used to solve (2.2).

Step III. The estimated optimal decision making is $\widehat{\pi}_n^*(\mathcal{X}) := \min\{A: \widehat{Q}_n(\mathcal{X}, A) = \max_{A_{(1)} \leq A \leq A_{(L_n)}} \widehat{Q}_n(\mathcal{X}, A)\}$ for an individual with characteristics \mathcal{X} , and the optimal value function is $\widehat{Q}_n^*(\mathcal{X}) := \widehat{\psi}_0^T \mathcal{X} + \max\left(0, (\widehat{\psi}_k^T \mathcal{X})_{k=2,\dots,L_n}\right)$.

3. Statistical Properties

In this section, we present the statistical properties of the proposed DROVE methodology for continuous-action decision-making. The proofs are in Appendix E of the supplementary materials.

3.1. Theoretical Properties of Coefficient Estimation

The following theorems give the statistical properties of our coefficient estimator.

Theorem 3.1. Under Assumptions 2.1–2.3, C.1, and C.2, there exists a strict local maximizer $\widehat{\boldsymbol{\beta}}_n = (\widehat{\boldsymbol{\psi}}_0^T, \widehat{\boldsymbol{\psi}}_2^T, \dots, \widehat{\boldsymbol{\psi}}_{L_n}^T)^T$ of the penalized likelihood function \mathcal{L}_n , which satisfies

$$P(\widehat{\psi}_{Z_0} = 0) \to 1 \text{ as } n \to \infty, \quad \text{and} \quad \|\widehat{\beta}_n - \beta_n^{\star}\|_2 = O_p(\sqrt{s_n/n}),$$

where $\widehat{\psi}_{Z_0} = (\widehat{\psi}_{z_0})_{(i,i) \in Z_0}, Z_0 = \{(i,i) : \psi_{z_0}^{\star} = 0, i = 0, i$

where $\widehat{\psi}_{Z_0}=(\widehat{\psi}_{i,j})_{(i,j)\in Z_0}, Z_0=\{(i,j): \psi^{\star}_{i,j}=0, i=0,2,\dots,L_n, 1\leq j\leq d\}.$ In addition,

$$P(D_{\text{null}}\widehat{\boldsymbol{\beta}}_n = \mathbf{0}) \to 1 \text{ as } n \to \infty, \quad \text{and} \quad \|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_n^*\|_2 = O_p(\sqrt{s_n/n}),$$
(3.1)

where $\widehat{\theta}_n = \mathbf{M}^{-1}\widehat{\boldsymbol{\beta}}_n$ and $\boldsymbol{\theta}_n^* = \mathbf{M}^{-1}\boldsymbol{\beta}_n^*$ for the transformation matrix \mathbf{M} defined in (B.2) in Appendix B.

Theorem 3.2. Under the assumptions of Theorem 3.1 and Assumption C.3, with probability tending to 1 as $n \to \infty$, the local maximizer in Theorem 3.1 satisfies

$$\sqrt{n}\Omega_n(\widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_n^{\star}) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \phi G),$$

where Ω_n is a $q \times p$ matrix, $q \leq s_n$ and fixed, $n\Omega_n U_0 B_n^{-1} U_0^T \Omega_n^T \rightarrow$ G, G is a $q \times q$ positive definite matrix, $||G||_2 = O(1)$, U_0 and \mathbf{B}_n are defined in (B.1) in Appendix B and Assumption C.3, respectively.

In practice, the covariance matrix ϕG needs to be estimated. Following the conventional technique (e.g., that of Fan and Li 2001), we estimate the variance using the following sandwich formula: $n\Omega_n \widehat{\mathbf{U}}_0 \widehat{\mathbf{B}}_n^{-1} \Big(\sum_{i=1}^n (\widetilde{\mathbf{z}}_i \widetilde{\mathbf{z}}_i^T \widehat{\boldsymbol{\varepsilon}}_i^2) \Big) \widehat{\mathbf{B}}_n^{-1} \widehat{\mathbf{U}}_0^T \Omega_n^T$, where $\widehat{\mathbf{U}}_0$ is the orthogonal matrix that spans the null space of $\widehat{\mathbf{D}}_{null}$, $\widehat{\mathbf{D}}_{\mathrm{null}}$ is the sub-matrix of D, which satisfies $\widehat{\mathbf{D}}_{\mathrm{null}}\widehat{\boldsymbol{\beta}}_n = \mathbf{0}$, $\widehat{\mathbf{B}}_{n} = \widetilde{\mathbf{X}}_{\text{signal}}^{T} \mathbf{\Sigma} (\widecheck{\mathbf{X}} \widehat{\boldsymbol{\beta}}_{n}) \widetilde{\mathbf{X}}_{\text{signal}}, \ \widetilde{\mathbf{X}}_{\text{signal}} = (\widetilde{\mathbf{z}}_{1}, \dots, \widetilde{\mathbf{z}}_{n})^{T} = \widecheck{\mathbf{X}} \widehat{\mathbf{U}}_{0},$ and $\Sigma(\delta) = \operatorname{diag}(b''(\delta_1), \dots b''(\delta_n))$ for any $\delta \in \mathbb{R}^n$.

According to Theorems 3.1 and 3.2, our coefficient estimator achieves the oracle property in that it identifies the true model with probability tending to one and enjoys strong consistency property.

Next, with regard to implementation, we introduce the following generalized local linear approximation (GLLA) algorithm, which is a generalization of the LLA algorithm (Zou and Li 2008; Fan, Xue, and Zou 2014).

Algorithm Generalized local linear approximation (GLLA)

Step I. Initialize
$$\widehat{\boldsymbol{\beta}}_{n}^{(0)} = \widehat{\boldsymbol{\beta}}_{n}^{\text{init}}$$
, and $\widehat{\boldsymbol{\Theta}}^{(0)} = (\widehat{\boldsymbol{\vartheta}}_{1}^{(0)}, \widehat{\boldsymbol{\vartheta}}_{2}^{(0)}, \dots, \widehat{\boldsymbol{\vartheta}}_{K}^{(0)})^{T}$, where $\widehat{\boldsymbol{\vartheta}}_{k}^{(0)} = \widehat{\boldsymbol{w}}_{k}^{(0)} \mathbf{d}_{k}$, and $\widehat{\boldsymbol{w}}_{k}^{(0)} = \rho_{\lambda_{n}}'(|\mathbf{d}_{k}^{T}\widehat{\boldsymbol{\beta}}_{n}^{(0)}|)$ for $1 < k < K$.

Step II. For m=1, 2, ..., repeat the following till convergence.

- 1. Solve $\widehat{\boldsymbol{\beta}}_{n}^{(m)} = \arg\max_{\boldsymbol{\beta}_{n}} l_{n}(\widehat{\boldsymbol{\beta}}_{n}) \lambda_{n} \|\widehat{\boldsymbol{\Theta}}^{(m-1)}\widehat{\boldsymbol{\beta}}_{n}\|_{1}$, where $\|\cdot\|_{1}$ denotes the ℓ_{1} vector norm, such that $\|\boldsymbol{w}\|_{1} = \sum_{i=1}^{n} |w_{i}|$ for any $\boldsymbol{w} = (w_{1}, \dots, w_{n})^{T}$.

 2. Update $\widehat{\boldsymbol{\Theta}}^{(m)} = (\widehat{\boldsymbol{\vartheta}}_{1}^{(m)}, \widehat{\boldsymbol{\vartheta}}_{2}^{(m)}, \dots, \widehat{\boldsymbol{\vartheta}}_{K}^{(m)})^{T}$, where $\widehat{\boldsymbol{\vartheta}}_{k}^{(m)} = \widehat{\boldsymbol{w}}_{k}^{(m)} \mathbf{d}_{k}$, and $\widehat{\boldsymbol{w}}_{k}^{(m)} = \rho_{\lambda_{n}}' (|\mathbf{d}_{k}^{T} \widehat{\boldsymbol{\beta}}_{n}^{(m)}|)$ for $1 \leq k \leq K$.

Assumption 2.2'. $\rho'(0+,\lambda) \geq a_1$; $\rho'(t,\lambda) \geq a_1$ for $t \in (0,a_2\lambda)$; $\rho'(t,\lambda) = 0$ for $t > a\lambda$ with constants $a > a_2 > 0$, and $a_1 > 0$.

Remark 3.1. Assumption 2.2' holds for the folded concave penalty, such as SCAD and MCP.

Proposition 3.1 gives the property of the GLLA algorithm.

Proposition 3.1. Under the assumptions of Theorem 3.1 and Assumption 2.2', assume that the oracle estimator $\widehat{\beta}_n$ in Theorem 3.1 is unique. In addition, assume that

$$\max_{s_n+1\leq j\leq p}|\widehat{\theta}_{nj}^{\text{init}}|\leq a_2\lambda_n, \quad \text{and} \quad \min_{1\leq k\leq K_1}|\mathbf{d}_k^T\widehat{\boldsymbol{\beta}}_n^{\text{init}}|\geq a\lambda_n,$$

for a_2 and a defined in Assumption 2.2' and $\widehat{\theta}_n^{\text{init}} = \mathbf{M}^{-1} \widehat{\beta}_n^{\text{init}}$. Then with probability tending to 1, the GLLA algorithm initialized by $\widehat{\beta}_n^{\rm init}$ finds the oracle estimator $\widehat{\beta}_n$ after one iteration.

Remark 3.2. One potential choice of the initial value is the generalized lasso estimator, which uses ℓ_1 penalty. The algorithm

for solving generalized lasso problem is discussed in Tibshirani and Taylor (2011) and Arnold and Tibshirani (2016). When D = I, the error bound exists for standard lasso estimator under proper designs; see (C1) and Corollary 3 in Fan, Xue, and Zou (2014).

About the computation cost of the algorithm, in particular the computation complexity with respect to L_n , heuristically, the generalized lasso needs $O(\max(K^2n, Kn^2))$ operations; see, for example, Tibshirani and Taylor (2011). The total computation costs of Steps I and II, therefore, are $O(\max(dL_n, n)dL_n n)$ and $O(\max(dL_n, n)dL_n n)$, respectively. In practice, the maximum iterations can be set as $O(\log n)$ and the algorithm converges fast within a few iterations. The total computation cost of the algorithm is $O((\log n) \max(dL_n, n)dL_n n)$, which is polynomial in L_n , implying that the algorithm is scalable to high dimensions.

3.2. Optimal Value Estimation and Inference

One important advantage of our DROVE approach is that it allows for proper inference of the value associated with a decision rule, in particular, the optimal decision that achieves the maximum value.

Let \mathcal{X} be the personal characteristic vector that belongs to the testing population. We use a different notation, \mathcal{X} , to distinguish the testing population from the estimation sample X and stress that \mathcal{X} is independent of Y. Suppose that the value function (2.1) holds for \mathcal{X} ; thus, it follows that $E(R|\mathcal{X},A) = Q_n(\mathcal{X},A)$. Let $\pi_n^*(\cdot)$ denote the optimal decision, and let $Q_n^*(\cdot)$ denote the optimal Q-function associated with $\pi_n^*(\cdot)$. More formally,

$$\pi_n^*(\mathcal{X}) := \min\{A : Q_n(\mathcal{X}, A) = \max_{A_{(1)} \le A \le A_{(L_n)}} Q_n(\mathcal{X}, A)\}, (3.3)$$

$$Q_n^*(\mathcal{X}) := \psi_0^{\star T} \mathcal{X} + \max\left(0, (\psi_k^{\star T} \mathcal{X})_{k=2,\dots,L_n}\right). \tag{3.4}$$

Suppose that $\mathcal{X} \sim F$, and the optimal value is $E(Q_n^*(\mathcal{X})) =$ $\int Q_n^*(\mathcal{X})dF(\mathcal{X})$. The expectation is taken over the distribution of the testing population \mathcal{X} given the working model Q_n . More generally, for any given decision rule $\pi(\cdot): \mathbb{R}^d \to [0,1]$ and for $\mathcal{X} \in \mathbb{R}^d$, we use $\check{\mathcal{X}}_{\pi}$ to denote a length-p vector

$$\check{\mathcal{X}}_{\pi} = \left((\mathcal{X})^{T}, (\mathcal{X})^{T} \cdot \mathbf{1}_{\{\pi(\mathcal{X}) \in [A_{(2)}, A_{(3)})\}}, \dots, (\mathcal{X})^{T} \cdot \mathbf{1}_{\{\pi(\mathcal{X}) \in [A_{(L_{n})}, A_{(L_{n}+1)})\}} \right)^{T}.$$
(3.5)

Using the notation in (3.5), the optimal Q-function is $Q_n^*(\mathcal{X}) =$ $\beta_n^{\star T} \mathring{\mathcal{X}}_{n,*}$, and the optimal value is: $E(Q_n^*(\mathcal{X})) = \int Q_n^*(\mathcal{X}) dF(\mathcal{X})$ $=\beta_n^{\star T}\int \check{\mathcal{X}}_{\pi_n^*}dF(\mathcal{X})=\beta_n^{\star T}E(\check{\mathcal{X}}_{\pi_n^*}).$ Let $\mathbb{P}_N(\cdot)$ denote the empirical mean measure for a sample of

size N: $\mathbb{P}_N(\omega) = \sum_{i=1}^N \omega_i/N$ for any $\omega = (\omega_i)_{1 \le i \le N}$. Given a testing sample of size N, $(\mathcal{X}_i)_{1 \le i \le N} \sim F$, we estimate the optimal value by

$$\mathbb{P}_{N}\widehat{Q}_{n}^{*}(\mathcal{X}) = \frac{1}{N} \sum_{i=1}^{N} \widehat{Q}_{n}^{*}(\mathcal{X}_{i}) = \widehat{\boldsymbol{\beta}}_{n}^{T} \left(\frac{1}{N} \sum_{i=1}^{N} \check{\mathcal{X}}_{i,\widehat{\pi}_{n}^{*}}\right) = \widehat{\boldsymbol{\beta}}_{n}^{T} \mathbb{P}_{N}(\check{\mathcal{X}}_{\widehat{\pi}_{n}^{*}}),$$
(3.6)

where

$$\widehat{\pi}_n^*(\mathcal{X}) := \min\{A : \widehat{Q}_n(\mathcal{X}, A) = \max_{A_{(1)} \le A \le A_{(L_n)}} \widehat{Q}_n(\mathcal{X}, A)\}, (3.7)$$

$$\widehat{Q}_n^*(\mathcal{X}) := \widehat{\boldsymbol{\psi}}_0^T \mathcal{X} + \max\left(0, (\widehat{\boldsymbol{\psi}}_k^T \mathcal{X})_{k=2,\dots,L_n}\right). \tag{3.8}$$

Remark 3.3. To estimate and make inference about the optimal value, we use the testing sample that is distinct from the Qfunction estimation sample. This approach is similar in spirit to the sample-splitting method in machine learning literature and is advocated by many recent studies on treatment effect evaluation, such as Chernozhukov et al. (2017), Wager and Athey (2018), and Athey and Wager (2019). The samplesplitting reduces bias and facilitates the valid inference of the optimal value. A discussion about the results of optimal value estimation using the estimation sample is in Appendix D of the supplementary materials. On the other hand, because we are mostly interested in predicting the decision-making effect on a broad population that extends beyond the estimation sample for which the decision-making effects are observable, it is also practically reasonable to use the testing sample to evaluate the effect of optimal decision-making.

Assumption 3.1. Assume

- 1. $s_n = o(n^{\delta_1})$ and $L_n = o(n^{\delta_2})$ for some $\delta_1, \delta_2 > 0$, and $2\delta_1 + \delta_2 < 1$.
- 2. $(\mathcal{X}_j)_{1 \leq j \leq N}$ are iid and independent of Y, $E(Q_n^*(\mathcal{X})^2) < \infty$, n = O(N), and $N = O(n^M)$ for some M > 1. There exist some constants r_0 , r_1 and $r_2 > 0$, for any t > 0, $P(|\chi_j| > t) \leq r_0 \exp(-r_1 t^{r_2})$ for all $1 \leq j \leq d$, where $\mathcal{X} =: (\chi_1, \ldots, \chi_d)^T$.

For the statistical inference of the optimal value, we further impose regularity condition; see Assumption C.4 in Appendix C of the supplementary materials.

Theorem 3.3 gives the asymptotic distribution of the estimated optimal value.

Theorem 3.3. Under the assumptions of Theorem 3.1 and Assumption 3.1, with probability tending to 1 as $n, N \to \infty$, the $\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X})$ defined in (3.6) satisfies

(i)

$$\mathbb{P}_{N}\widehat{Q}_{n}^{*}(\mathcal{X}) - E(Q_{n}^{*}(\mathcal{X})) = o_{p}(1). \tag{3.9}$$

(ii) If, in addition, $2\delta_1 + \delta_2 < 1/2$, Assumptions C.3 and C.4 hold, $\lim_{n\to\infty} nE\check{\mathcal{X}}_{\pi_n^*}^T \mathbf{U}_0 \mathbf{B}_n^{-1} \mathbf{U}_0^T E\check{\mathcal{X}}_{\pi_n^*} \phi = \sigma_1^2$ and $\lim_{n,N\to\infty} \text{var}(Q_n^*(\mathcal{X})) n/N = \sigma_2^2$ for some constants $\sigma_1 > 0$, $\sigma_2 \ge 0$. Let $\sigma_*^2 = \sigma_1^2 + \sigma_2^2$, then

$$\sqrt{n}\Big(\mathbb{P}_N\widehat{Q}_n^*(\mathcal{X}) - E\Big(Q_n^*(\mathcal{X})\Big)\Big) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^2).$$
 (3.10)

In practice, in order to apply Theorem 3.3 to perform feasible statistical inference, we need to estimate the variance. We estimate the variance using the following sandwich formula:

$$\widehat{\sigma}_{*}^{2} = n \mathbb{P}_{N} \check{\mathcal{X}}_{\widehat{\pi}_{n}^{*}}^{T} \widehat{\mathbf{U}}_{0} \widehat{\mathbf{B}}_{n}^{-1} \Big(\sum_{i=1}^{n} (\widetilde{\mathbf{z}}_{i} \widetilde{\mathbf{z}}_{i}^{T} \widehat{\boldsymbol{\varepsilon}}_{i}^{2}) \Big) \widehat{\mathbf{B}}_{n}^{-1} \widehat{\mathbf{U}}_{0}^{T} \mathbb{P}_{N} \check{\mathcal{X}}_{\widehat{\pi}_{n}^{*}} + \widehat{\text{var}} (\widehat{Q}_{n}^{*}(\mathcal{X})) n/N,$$

$$(3.11)$$

where $\mathbb{P}_N \check{\mathcal{X}}_{\widehat{\pi}_n^*}$ and $\widehat{\mathrm{var}}(\widehat{Q}_n^*(\mathcal{X}))$ are the sample mean of $\check{\mathcal{X}}_{\widehat{\pi}_n^*}$ and sample variance of $\widehat{Q}_n^*(\mathcal{X})$, respectively, based on the testing sample $(\mathcal{X}_j)_{1 \leq j \leq N}$.

Remark 3.4. Our Theorems 3.1 and 3.3 are established for the continuous-action decision making in which the discretization level $L_n \to \infty$ and hence require intrinsically different methodological design and novel mathematical treatment than existing studies on the binary case (e.g., Shi, Song, and Lu 2016). We device a innovative DROVE approach, and obtain a new result (3.1) in Theorem 3.1, which ensures that the adjacent treatment levels can be identified. Unlike Shi, Song, and Lu (2016), we work under the assumptions for transformed design matrices and transformed parameters (Assumptions C.1–C.3). More essentially, we consider a more relaxed constraint on the general penalty matrix **D**. We use a partial reparameterization technique to show the statistical properties of DROVE; see Appendix B of the supplementary materials for more details.

Remark 3.5. For continuous action, there exists no nonparametric/semiparametric approach that yields a \sqrt{n} -consistent estimator of effect curve without imposing parametric assumptions; see, for example, the discussion in Kennedy et al. (2017). In this study, on the other hand, we focus on the high-dimensional parametric model (2.1), in which the parameter space is sparse in the general sense, hence, our approach recovers the true sparse model consistently. In addition, we allow the discretization level L_n to diverge slowly with n, such that $L_n s_n^2 = o(\sqrt{n})$, under which setting, inference of the optimal value is obtainable.

In addition to the optimal value, our approach also allows for inference of the value difference between the optimal decision and a given decision rule. For a given decision rule, say, $\pi(\cdot)$: $\mathbb{R}^d \to [0,1]$, the associated Q-function and the value of $\pi(\cdot)$ are $Q_n(\mathcal{X},\pi) = \psi_0^{\star T} \mathcal{X} + \sum_{k=2}^{L_n} \psi_k^{\star T} \mathcal{X} \cdot \mathbf{1}_{\{\pi(\mathcal{X}) \in [A_{(k)}, A_{(k+1)})\}}$, and

$$E(Q_n(\mathcal{X},\pi)) = \int Q_n(\mathcal{X},\pi) dF(\mathcal{X}) = \beta_n^{\star T} E(\check{\mathcal{X}}_{\pi}).$$
 (3.12)

The value difference between π_n^* and π is $E(Q_n^*(\mathcal{X})) - E(Q_n(\mathcal{X}, \pi))$, which is also called the regret of the policy π ; see, for example, Athey and Wager (2019).

We estimate the value associated with $\pi(\cdot)$ using $\mathbb{P}_N\widehat{Q}_n(\mathcal{X},\pi)$, that is.

$$\mathbb{P}_{N}\widehat{Q}_{n}(\mathcal{X},\pi) = \frac{1}{N} \sum_{i=1}^{N} \widehat{Q}_{n} \Big(\mathcal{X}_{i}, \pi(\mathcal{X}_{i}) \Big) = \widehat{\boldsymbol{\beta}}_{n}^{T} \mathbb{P}_{N}(\check{\mathcal{X}}_{\pi}), \quad (3.13)$$

and we estimate the value difference between π_n^* and π using $\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X}) - \mathbb{P}_N \widehat{Q}_n(\mathcal{X}, \pi)$.

Proposition 3.2 gives the asymptotic distribution of the estimated value difference $\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X}) - \mathbb{P}_N \widehat{Q}_n(\mathcal{X}, \pi)$.

Proposition 3.2. Under the assumptions of Theorem 3.1 and Assumption 3.1, given a decision rule $\pi(\cdot)$: $\mathbb{R}^d \to [0,1]$, with probability tending to 1 as $n,N\to\infty$, the $\mathbb{P}_N\widehat{Q}_n^*$ and $\mathbb{P}_N\widehat{Q}_n(\mathcal{X},\pi)$ defined in (3.6) and (3.13) satisfy

(

(i) $\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X}) - \mathbb{P}_N \widehat{Q}_n(\mathcal{X}, \pi) - \left(E Q_n^*(\mathcal{X}) - E Q_n(\mathcal{X}, \pi) \right) = o_p(1).$

(ii) If, in addition, $2\delta_1 + \delta_2 < 1/2$, Assumptions C.3 and C.4 hold, $\lim_{n\to\infty} nE(\check{\mathcal{X}}_{\pi_n^*} - \check{\mathcal{X}}_{\pi})^T \mathbf{U}_0 \mathbf{B}_n^{-1} \mathbf{U}_0^T E(\check{\mathcal{X}}_{\pi_n^*} - \check{\mathcal{X}}_{\pi}) \phi = \sigma_1^2$, and $\lim_{n,N\to\infty} \text{var} \Big(Q_n^*(\mathcal{X}) - Q_n(\mathcal{X},\pi) \Big) n/N = \sigma_2^2$ for some constants $\sigma_1 > 0$, and $\sigma_2 \ge 0$. Let $\sigma_{*,\pi}^2 = \sigma_1^2 + \sigma_2^2$, then

$$\sqrt{n} \Big(\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X}) - \mathbb{P}_N \widehat{Q}_n(\mathcal{X}, \pi) - \Big(EQ_n^*(\mathcal{X}) - EQ_n(\mathcal{X}, \pi) \Big) \Big) \\
\xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{*\pi}^2).$$

Similar to (3.11), we can esimate $\sigma_{*,\pi}^2$ using the following sandwich formula

$$\widehat{\sigma}_{*,\pi}^{2} = n \mathbb{P}_{N} (\check{\mathcal{X}}_{\widehat{\pi}_{n}^{*}} - \check{\mathcal{X}}_{\pi})^{T} \widehat{\mathbf{U}}_{0} \widehat{\mathbf{B}}_{n}^{-1}$$

$$\left(\sum_{i=1}^{n} (\widetilde{\mathbf{z}}_{i}^{*} \widetilde{\mathbf{z}}_{i}^{T} \widehat{\boldsymbol{\varepsilon}}_{i}^{2}) \right) \widehat{\mathbf{B}}_{n}^{-1} \widehat{\mathbf{U}}_{0}^{T} \left(\mathbb{P}_{N} (\check{\mathcal{X}}_{\widehat{\pi}_{n}^{*}} - \check{\mathcal{X}}_{\pi}) \right)$$

$$+ \widehat{\text{var}} \left(\widehat{\mathbf{Q}}_{n}^{*} (\mathcal{X}) - \widehat{\mathbf{Q}}_{n} (\mathcal{X}, \pi) \right) n / N,$$
(3.14)

where $\mathbb{P}_N \check{\mathcal{X}}_{\pi}$ and $\widehat{\text{var}}(\widehat{Q}_n^*(\mathcal{X}) - \widehat{Q}_n(\mathcal{X}, \pi))$ are the sample mean of $\check{\mathcal{X}}_{\pi}$ and sample variance of $\widehat{Q}_n^*(\mathcal{X}) - \widehat{Q}_n(\mathcal{X}, \pi)$, respectively.

4. Simulation Study

4.1. Simulation Setup

We generate data from the model $Y = \psi_0^T \mathbf{X} + \sum_{k=2}^{L_n} \psi_k^T \mathbf{X} \cdot \mathbf{1}_{\{A=A_{(k)}\}} + \varepsilon$ and calibrate the parameters based on the empirical data as described in Section 5. Specifically, the covariates \mathbf{X} is a vector of length 9 and randomly drawn from the personal characteristics that are used in the Q-learning estimation in the empirical study, which includes three binary variables, two categorical variables, three continuous variables and one intercept term.

We divide the continuous action interval [0, 1] into 11 discrete levels, that is, $L_n = 11$. The stock ratio A_i for each X_i is randomly assigned from $\{0, 0.1, ..., 1\}$ with equal probability. The choice of L_n is the same as that in the empirical studies. Empirically, in order to balance the applicability of the methodology with the approximation accuracy of the working model, the choice of L_n takes both the theoretical rate and the common practice into consideration. For example, in our personalized asset allocation study using the HRS data, the training sample size is n = 2000, $n^{1/3} \approx 13$, and we choose $L_n = 11$. On the other hand, it is a common practice for mutual funds to use a 10% incrementation and make 0%, 10%, 20%, ..., 100% recommendations. We illustrate the case when $L_n = 20$ in Appendix G of the supplementary materials. When a larger n is available, nice numerical properties can be expected for a larger range of L_n .

The Q-function parameters are learned from the empirical study and with hard thresholding (see Appendix F of the supplementary materials for a full description). The number

of coefficients is therefore $p=11\times 9=99$. The noise, ε , is generated from $\mathcal{N}(0,0.5^2)$. We design the generalized penalty matrix $\mathbf{D}=(\mathbf{D}_{\text{signal}}^T,\mathbf{D}_{\text{null}}^T)^T$ according to (2.2). The $\boldsymbol{\beta}_n=(\boldsymbol{\psi}_0^T,\boldsymbol{\psi}_2^T,\ldots,\boldsymbol{\psi}_{L_n}^T)^T$ has 55 zeros, and rank $(\mathbf{D}_{\text{null}})=76$, thus, the nonsparsity coefficient is $s_n=99-76=23$. The minimal signal is $g_n=0.4$.

4.2. Simulation Results

We estimate the model using the DROVE method. For comparison, we also evaluate the results using the standard lasso (std-lasso) and standard scad² (std-scad). We also present the results of the infeasible oracle estimator. The oracle solution is obtained by performing a least squares regression of the response Y over $\widetilde{\mathbf{X}}_{\text{signal}}$, as defined in (C.2) in Appendix C. The regression yields the oracle transformed estimator $\widehat{\boldsymbol{\theta}}_n$, and the oracle estimator $\widehat{\boldsymbol{\beta}}_n$ is obtained by the transformation $\widehat{\boldsymbol{\beta}}_n = \mathbf{M}\widehat{\boldsymbol{\theta}}_n$. The tuning parameters are chosen by minimizing the validation error.

We conduct 500 independent replications with the sample sizes n=2000 and n=3000. We evaluate the coefficient estimation accuracy of various methods by the ℓ_2 and ℓ_1 errors, and measure the parameter selection accuracy by the false positive rate (FP/N, the number of all false positives divided by that of all total negatives in $\widehat{\beta}_n$), and false negative rate (FN/P, the number of all false negatives divided by that of all total positives in $\widehat{\beta}_n$). The total negatives (N) and the total positives (P) in β_n are 55 and 44, respectively. For counting the false positives and false negatives, we set the threshold level to be 10^{-4} . The results are in Table 1. The results show that our DROVE estimator has the lowest estimation error among all compared methods.

Next, we estimate the optimal value and construct its confidence interval (CI). According to Theorem 3.3, the confidence interval of the optimal value at a significance level $\alpha > 0$ is $(\mathbb{P}_N \widehat{Q}_n^*(\mathcal{X})) \pm \widehat{\sigma}(\mathbb{P}_N \widehat{Q}_n^*) \times \mathcal{N}^{-1}(1 - 0.5\alpha)$, where $\mathcal{N}(\cdot)$ is the cumulative distribution function of the standard normal distribution, $\widehat{\sigma}(\mathbb{P}_N Q_n^*) = \widehat{\sigma}_n^* / \sqrt{n}$. The feasible variance estimator in (3.11) is used. The estimation sample sizes are n = 2000 and 3000, and the testing sample sizes are N = 5000 and 15,000. The testing sample is randomly generated independently based on empirical data of the personal characteristics. Table 2 reports the estimated optimal and empirical coverages of 90%, 95%, and 99% CI from 500 replications. The results suggest that the estimated optimal value is close to the actual one and the estimated CIs show good accuracy. When the estimation sample size increases from 2000 to 3000, the empirical coverage goes closer to the nominal one.

Finally, we construct confidence intervals for the difference between the optimal value and the value of a pre-specified decision rule $\pi(\cdot)$. We build confidence intervals according to Proposition 3.2 and use the feasible standard deviation estimator in (3.14). We calculate the empirical coverage from 500 replications. The results are in Table 3. The results show that the

¹See: https://institutional.vanguard.com/assets/pdf/vrpa/InvestorQuestionn aireAssetAllocationInsert.pdf.

²Existing approaches for the decision-making problem under the high-dimensional setting focus on binary problems, among which one of the most comparable to ours is Shi, Song, and Lu (2016), who use the SCAD estimator to fit Q-function. The std-scad estimator we evaluate, therefore, can be considered as an ad-hoc extension of Shi, Song, and Lu (2016) by discretization to the continuous-action setting.



Table 1. Coefficient estimation and parameter selection accuracy of various methods.

Error	Oracle	DROVE	std-scad	std-lasso
n = 2000				
$\ \widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_n\ _2$	1.683(0.322)	4.370(0.986)	4.844(0.788)	5.348(0.904)
$\ \widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_n\ _1$	8.158(1.628)	24.054(7.083)	29.113(6.089)	37.195(6.034)
FP/N	_	0.148(0.083)	0.242(0.109)	0.740(0.082)
FN/P	-	0.135(0.070)	0.177(0.069)	0.073(0.046)
n = 3000				
$\ \widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_n\ _2$	1.348(0.239)	3.221(0.862)	3.909(0.723)	4.344(0.699)
$\ \widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_n\ _1$	6.558(1.225)	16.486(5.648)	23.033(5.555)	30.466(4.964)
FP/N	_	0.095(0.076)	0.237(0.118)	0.772(0.074)
FN/P	-	0.094(0.068)	0.124(0.062)	0.042(0.035)

NOTE: Coefficient estimation accuracy is measured by the ℓ_2 and ℓ_1 errors, and parameter selection accuracy is measured by the false positive rate (FP/N) and false negative rate (FN/P). We report the mean and standard deviation (in parentheses) from 500 replications.

Table 2. Estimated optimal value and empirical coverage of confidence intervals for the optimal value.

(n, N)	(2000,5000)	(2000,15,000)	(3000,5000)	(3000,15,000)	
Estimates	$E(Q_n^*) = 0.667$				
$\mathbb{P}_N\widehat{Q}_n^*$	0.665	0.663	0.664	0.667	
$\widehat{\sigma}(\mathbb{P}_N\widehat{Q}_n^*)$	0.024	0.023	0.020	0.019	
Coverage					
90% CI	0.838	0.828	0.874	0.864	
95% CI	0.898	0.896	0.932	0.928	
99% CI	0.958	0.952	0.988	0.990	

NOTE: The estimation sample sizes are n=2000 and 3000, and the testing sample sizes are N=5000 and 15,000. The results are the empirical coverage of 90%, 95% and 99% CI from 500 replications, together with the mean of the estimated optimal values, $\mathbb{P}_N(\widehat{Q}_n^*)$, and the mean of the estimated standard deviations, $\widehat{\sigma}(\mathbb{P}_N\widehat{Q}_n^*)$.

estimated value difference is close to the true one for different estimation sample sizes and testing sample sizes. Moreover, the estimated CIs exhibit good accuracy. When the estimation sample size increases from 2000 to 3000, the empirical coverage of the confidence interval goes closer to the nominal level.

More simulation examples under different simulation settings with various nonsparsity degree s_n and various choices of L_n are in Appendix G of the supplementary materials. The results show robustness in the performance of our estimator under various settings.

In summary, the simulation results show that our proposed methodology gives rise to better model estimation accuracy than the standard penalized regression when sparsity occurs to the linear projection of the coefficients. Our approach also allows for valid inferences of the optimal value and value differences using the feasible variance estimator, which performs well in terms of empirical coverage.

5. Empirical Study

5.1. Individualized Asset Allocation Optimization

We study the asset allocation optimization problem under a consumption-based utility framework. Specifically, we adopt the additive utility with constant relative risk aversion (CRRA, e.g., Hall 1978). The utility has the form $U(C_1, C_2, ..., C_T) = C_1$

Table 3. Estimated value differences and the empirical coverage of confidence intervals for the value difference between the optimal decision and pre-specified decision rules.

(n, N)	(2000,5000)	(2000,15,000)	(3000,5000)	(3000,15,000)
Estimates		$E(Q_n^*) - E(Q_n^*)$	n_{n,π_1}) = 1.037	
$\mathbb{P}_N \widehat{Q}_n^* - \mathbb{P}_N \widehat{Q}_{n,\pi_1}$	1.034	1.030	1.039	1.040
$\widehat{\sigma}(\mathbb{P}_N\widehat{Q}_n^* - \mathbb{P}_N\widehat{Q}_{n,\pi_1})$	0.032	0.030	0.026	0.025
Coverage				
90% CI	0.830	0.816	0.856	0.844
95% CI	0.896	0.888	0.914	0.902
99% CI	0.958	0.954	0.978	0.978
Estimates		$E(Q_n^*) - E(Q_{n,\pi_2}) = 1.734$		
$\mathbb{P}_N\widehat{Q}_n^* - \mathbb{P}_N\widehat{Q}_{n,\pi_2}$	1.736	1.732	1.730	1.733
$\widehat{\sigma}(\mathbb{P}_{N}\widehat{Q}_{n}^{*}-\mathbb{P}_{N}\widehat{Q}_{n,\pi_{2}}^{2})$	0.036	0.036	0.030	0.029
Coverage				
90% CI	0.846	0.846	0.866	0.862
95% CI	0.902	0.908	0.928	0.920
99% CI	0.962	0.958	0.988	0.988

NOTE: The estimation sample sizes are n=2000 and 3000, and the testing sample sizes are N=5000 and 15,000. The results are the empirical coverage of 90%, 95% and 99% CI from 500 replications, together with the mean of the estimated value difference $\mathbb{P}_N\widehat{Q}_n^* - \mathbb{P}_N\widehat{Q}_{n,\pi}$ and the mean of the estimated standard deviation $\widehat{\sigma}(\mathbb{P}_N\widehat{Q}_n^* - \mathbb{P}_N\widehat{Q}_{n,\pi})$.

 $\sum_{t=1}^{T} \gamma^t u(C_t)$, where C_t is the consumption at time t, $u(C) = C^{1-\rho}/(1-\rho)$, ρ is the risk aversion parameter that reflects the sensitivity of utility to the randomness in the income, and γ is the discount rate (set as 0.96; see, e.g., Gourinchas and Parker 2002). The consumption evolves according to the net income I_t and wealth W_t constraint $I_t - C_t = W_{t+1} - W_t$. The objective is to maximize the expected utility by choosing the optimal proportion of the total financial asset invested in stocks, that is, the stock ratio, A. The utility optimization problem spans the preceding five years (i.e., T = 5).

Formally, the optimal stock ratio solves the following optimization problem³: $\pi^*(X) = \operatorname{argmax}_{A \in [0,1]} E(U(C_1, C_2, \dots C_5 | X, A))$, where X is the personal characteristics. The asset allocation decision A determines the distribution of income and consequently influences the dynamics of wealth and consumption. We consider randomness in income from not only the financial returns but also from other sources such as wages and medical expenses, the distributions of which vary from individual to individual. Furthermore, the risk aversion parameter γ also differs between individuals and relates to their personal characteristics.

5.2. Data

We use the Health and Retirement Study (HRS) data from 1992 to 2014. HRS is a national-level longitudinal survey of more than 22,000 U.S. residents over the age of 50. We include key

³In this exploratory study, we consider a simplified situation where utilities are only determined by the consumption within the next five years. This does not mean, however, that individuals do not care what happens after five years. In fact, in our model, consumption is considered optimal with respect to wealth, in particular, C₅ depends on W₅. An individual's preference for a high C₅ is typically consistent with a high W₅, which means that one does care about having enough savings for five years later. In future works where multi-stage decision-making is studied, this would become even less an issue.



variables on households finance (e.g., total financial wealth, value of stocks held and income) and individual characteristics (e.g., age, marital status, education level, health condition and working status) in our analysis. We also incorporate the Consumption and Activities Mail Survey (CAMS) data that contain consumption information, which began in 2002, complementing the HRS. The CAMS data cover approximately 6800 people, which is a subset of respondents covered in the HRS data. The relevant variable in CAMS that we incorporate in our analysis is the nondurable consumption. The raw HRS and CAMS data are noisy and contain many missing observations. In addition, the observed stock ratios are concentrated around zero. Instead of using the raw data, we use pseudo observations that are randomly generated from models of trajectory paths for consumption and income. We also learn risk aversion parameters by matching observed investment profiles with individual characteristics. Details of the model-based random experiment generation are in Appendix H of the supplementary materials.

5.3. Statistical Learning Implementation and Results

We conduct the analysis based on three samples: estimation sample, evaluation sample, and testing sample. For the observations that are covered by both CAMS and HRS, we split them into an estimation sample and an evaluation sample. The estimation of the Q-function is performed on the estimation sample, which contains 2000 observations. The model fitting is evaluated on the evaluation sample, which includes 1000 observations. We put all the 15,000+ observations that are not used for the model estimation in the testing sample, which includes those that are not covered by CAMS. The testing sample is used in evaluating the value of strategies.

5.3.1. Q-function Estimation

We include the following covariates in the Q-function (2.1): gender, education, unhealthiness, working status, marital status, age, household wealth, income, and an intercept term, thus, d = 9. We follow the common practice in asset allocation applications and adopt the 0.1 incrementation in stock ratios. The stock ratios are divided into 11 levels $0, 0.1, \ldots, 1$ (i.e., $L_n = 11$). The total number of coefficients is therefore $p = 9 \times 11 = 99$. All covariates are standardized to have a mean of 0 and standard deviation of 0.1. For each individual *i* in the estimation sample, we randomly assign a stock ratio A_i (0, 0.1, ..., 1) and perform the model-based random experiments. We set the utility reward variable Y_i to be a Gaussianized score based on the rank of the average utility achieved by A_i within all stock ratio levels; see Appendix H.3. We estimate the value function (2.1) using these 2000 observations in the estimation sample. The details of the personal characteristics included in the model and the estimated model coefficients are in Appendix I of the supplementary materials.

We then check the goodness of fit of the model on the evaluation sample of size 1000. The evaluation sample is different from the estimation sample but within the subset covered by the CAMS that has consumption information, and hence the utility scores are obtainable. The out-of-sample R^2 is 31.4% (insample R^2 : 45.0%), suggesting a good fit.

Table 4. Performance of various strategies.

Strategy	$\mathbb{P}_{N}\widehat{Q}_{n,\pi}$	$\mathbb{P}_{N}\widehat{Q}_{n}^{*}-\mathbb{P}_{N}\widehat{Q}_{n,\pi}$	95% CI of ($EQ_n^* - EQ_{n,\pi}$)	
π_{obs} π_0	0.216	0.403	0.328	0.479
	0.303	0.316	0.227	0.406
$\pi_{0.1}$ $\pi_{0.2}$	0.238	0.381	0.268	0.493
	0.244	0.375	0.273	0.477
$\pi_{0.3}$ $\pi_{0.4}$	0.272	0.347	0.255	0.439
	0.267	0.352	0.250	0.454
$\pi_{0.5}$ $\pi_{0.6}$	-0.382	1.001	0.857	1.145
	-0.334	0.953	0.816	1.091
$\pi_{0.7}$ $\pi_{0.8}$	-0.362	0.981	0.841	1.121
	-0.374	0.993	0.850	1.137
$\pi_{0.9}$ $\pi_{1.0}$	-1.089	1.708	1.546	1.870
	-1.079	1.698	1.541	1.855
$\widehat{\pi}_{n}^{*}$	0.619	-	-	-

NOTE: This table gives the results of the value comparison among the estimated optimal strategy, $(\widehat{\pi}_n^*)$, observed strategy, (π_{obs}) , and fixed strategies, $(\pi_j = j, j = 0, 0.1, \ldots, 1)$. The values are the estimated value (i.e., the average estimated utility reward) for various strategies, $\mathbb{P}_N \widehat{Q}_{n,\pi}$, the estimated value improvement when using the optimal strategy over the benchmark strategies, $\mathbb{P}_N \widehat{Q}_n^* - \mathbb{P}_N \widehat{Q}_{n,\pi}$, and the 95% CIs of the value difference $EQ_n^* - EQ_{n,\pi}$. The model estimation uses a random estimation sample of 2000. The testing sample comprises the remaining 15,000+ observations in the HRS from distinct individuals.

5.3.2. Strategy Performance

We evaluate the performance of our estimated individualized optimal asset allocation on the testing sample ($N \ge 15,000$). The testing sample contains the individuals in HRS who are not included in the estimation sample. Because most observations of the testing sample are not covered by CAMS, consumption information is scarce. Nevertheless, because our final Q-function does not include the consumption variable, using the estimated Q-function, we can evaluate the performance of various strategies on the testing sample.

We compare the performance of our individualized optimal asset allocation strategy⁴ $\widehat{\pi}_n^*$ with the following benchmark strategies: (a) we assign the same fixed (0, 0.1, ..., 1) stock ratio to all individuals, a strategy denoted by $\pi_j = j, j = 0, 0.1, 0.2, \ldots, 1$; (b) we assign the originally observed stock ratio, which is denoted by π_{obs} .

For each strategy, we compute the estimated policy value $\mathbb{P}_N\widehat{Q}_{n,\pi} = \sum_{i=1}^N \widehat{Q}_n(\mathcal{X}_i,\pi(\mathcal{X}_i))/N$. We construct CIs for $E(Q_n^*) - E(Q_n(\mathcal{X},\pi))$ according to Proposition 3.2. The results are in Table 4.

The results in Table 4 suggest that our individualized optimal asset allocation decisions substantially improve over the benchmark strategies in terms of policy values. Compared to the observed stock ratios, the individualized strategy improves the average utility reward from 0.216 to 0.619. It also yields an average utility reward that is 0.316–1.708 greater, relative to the fixed strategies. The 95% CIs of the value differences are away from zero, suggesting that the individualized optimal strategy achieves a significantly higher value compared with the values of the observed and fixed strategies. This result demonstrates the importance of individualization in asset allocation.

⁴The proposed strategy is based on the DROVE estimator we develope. We also evaluated the strategy based on standard SCAD estimator (std-scad) used in Shi, Song, and Lu (2016) by direct discretization to the continuous-action setting. Compared to the std-scad strategy, our DROVE approach achieves a higher average value (0.619 vs. 0.597), and the difference between the two is statistically significant.



A low proportion of households participate in the equity market; few of them hold stocks (e.g., Hong, Kubik, and Stein 2004; Campbell 2006). We also find a low participation in equities. Notably, when adopting the optimal strategy, stock market participation increases from 38% to 63%; see Appendix J of the supplementary materials for more details. These results show that our optimal individualized strategy improves the financial well-being of the population and it entails a higher stock market participation on average, which indicates a healthier and more active stock market.

6. Conclusion and Discussion

We develop a high-dimensional statistical learning methodology for continuous-action decision-making with an important application in individualized asset allocation. We show that our DROVE approach enjoys consistency in the model coefficients estimation. Moreover, our approach achieves valid statistical inference on the optimal value. Empirically, we apply the proposed methodology to study the individualized asset allocation problem using HRS and CAMS data. Under a consumptionbased utility framework, our individualized optimal asset allocation strategy substantially improves the financial well-being of the population. The outperformance of our individualized optimal strategy over the fixed stock ratio strategies highlights the benefit of individualization.

The statistical learning framework developed in this article has broad implications. Methodologically, our approach can be extended to the study of stochastic policies, which can have advantages for problems with partially observed states (e.g., Singh, Jaakkola, and Jordan 1994) and be of interest for applications such as mobile health under infinite horizon settings (e.g., Luckett et al. 2019; Liao, Klasnja, and Murphy 2021). Another important direction is to extend our framework from the single period to multiple-period decision-making. For multiple-stage problems, one potential approach would be to model the Qfunction for each stage using the method developed in this article and then perform a backward recursive procedure to obtain the optimal dynamic decisions. About estimation and inference of the optimal value, challenges arise in the multi-stage problem given that the estimation errors in the optimal rule and the associated optimal value in the latter stage carry over to the former stages. In terms of application, for example, not limited to the setting exercised in this article, the framework can be readily applied to achieving other wealth management objectives, such as post-retirement saving adequacy. The framework can also be extended to the study of multi-class asset allocation, and to incorporate dynamic prediction models of mean and volatility of financial returns.

Supplementary Materials

The online supplementary materials contain the additional assumptions, proofs of the main results in the article, and additional discussions on the theoretical and numerical results.

Acknowledgments

We thank the editor, the associate editor, and two anonymous referees for very constructive comments.

Funding

Research is partially supported by startup fund of Hong Kong Polytechnic University, RGC GRF 15302321, RGC GRF16502118 and T31-604/18N of the HKSAR, and NSF-DMS-1555244.

Disclosure Statement

The authors report there are no competing interests to declare.

ORCID

Rui Song https://orcid.org/0000-0003-1875-2115

References

- Arnold, T. B., and Tibshirani, R. J. (2016), "Efficient Implementations of the Generalized Lasso Dual Path Algorithm," Journal of Computational and Graphical Statistics, 25, 1-27. [640,641,643]
- Athey, S., and Wager, S. (2019), "Efficient Policy Learning," arXiv preprint arXiv:1702.02896. [639,644]
- Banks, J., Blundell, R., and Tanner, S. (1998), "Is There a Retirement-Savings Puzzle?" American Economic Review, 88, 769-788. [640]
- Cai, H., Shi, C., Song, R., and Lu, W. (2020), "Deep Jump q-Evaluation for Offline Policy Evaluation in Continuous Action Space," arXiv preprint arXiv:2010.15963. [639]
- Campbell, J. Y. (2006), "Household Finance," The Journal of Finance, 61, 1553-1604. [648]
- Chen, G., Zeng, D., and Kosorok, M. R. (2016), "Personalized Dose Finding Using Outcome Weighted Learning," Journal of the American Statistical Association, 111, 1509-1521. [639,641]
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., and Newey, W. (2017), "Double/Debiased/Neyman Machine Learning of Treatment Effects," American Economic Review, 107, 261-65. [644]
- De Nardi, M., and Yang, F. (2014), "Bequests and Heterogeneity in Retirement Wealth," European Economic Review, 72, 182-196. [640]
- Engen, E. M., Gale, W. G., Uccello, C. E., Carroll, C. D., and Laibson, D. I. (1999), "The Adequacy of Household Saving," Brookings Papers on Economic Activity, 1999, 65-187. [640]
- Fan, J., and Li, R. (2001), "Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties," Journal of the American Statistical Association, 96, 1348-1360. [640,642,643]
- Fan, J., and Lv, J. (2011), "Nonconcave Penalized Likelihood with np-Dimensionality," IEEE Transactions on Information Theory, 57, 5467-5484. [640,641,642]
- Fan, J., Xue, L., and Zou, H. (2014), "Strong Oracle Optimality of Folded Concave Penalized Estimation," The Annals of Statistics, 42, 819-849. [643]
- Gourinchas, P.-O., and Parker, J. A. (2002), "Consumption Over the Life Cycle," Econometrica, 70, 47-89. [646]
- Haider, S. J., and Stephens, Jr, M. (2007), "Is There a Retirement-Consumption Puzzle? Evidence using Subjective Retirement Expectations," The Review of Economics and Statistics, 89, 247–264. [640]
- Hall, R. E. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, 86, 971-987. [646]
- Hong, H., Kubik, J. D., and Stein, J. C. (2004), "Social Interaction and Stock-Market Participation," The Journal of Finance, 59, 137–163. [640,648]
- Hurd, M., and Rohwedder, S. (2003), "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement," Technical report, National Bureau of Economic Research. [640]
- Kennedy, E. H., Ma, Z., McHugh, M. D., and Small, D. S. (2017), "Nonparametric Methods for Doubly Robust Estimation of Continuous Treatment Effects," Journal of the Royal Statistical Society, Series B, 79, 1229-1245.
- Laber, E., and Zhao, Y. (2015), "Tree-based Methods for Individualized Treatment Regimes," Biometrika, 102, 501-514. [639]
- Liao, P., Klasnja, P., and Murphy, S. (2021), "Off-policy Estimation of Long-Term Average Outcomes with Applications to Mobile Health," Journal of the American Statistical Association, 116, 382-391. [648]

(

- Luckett, D. J., Laber, E. B., Kahkoska, A. R., Maahs, D. M., Mayer-Davis, E., and Kosorok, M. R. (2019), "Estimating Dynamic Treatment Regimes in Mmobile Hhealth using v-learning," *Journal of the American Statistical Association*, 115, 692–706. [648]
- Lv, J., and Fan, Y. (2009), "A Unified Approach to Model Selection and Sparse Recovery using Regularized Least Squares," The Annals of Statistics, 37, 3498–3528. [640,642]
- Meinshausen, N., and Bühlmann, P. (2006), "High-Dimensional Graphs and Variable Selection with the Lasso," *The Annals of Statistics*, 34, 1436–1462. [640]
- Munnell, A. H., Webb, A., and Golub-Sass, F. (2012), "The National Retirement Risk Index: An Update," Center for Retirement Research at Boston College, 1, 719–744. [640]
- Murphy, S. A. (2003), "Optimal Dynamic Treatment Regimes," Journal of the Royal Statistical Society, Series B, 65, 331–355. [639]
- ——— (2005), "An Experimental Design for the Development of Adaptive Treatment Strategies," Statistics in Medicine, 24, 1455–1481. [641]
- Palumbo, M. G. (1999), "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," The Review of Economic Studies, 66, 395–421. [640]
- Qian, M., and Murphy, S. A. (2011), "Performance Guarantees for Individualized Treatment Rules," The Annals of Statistics, 39, 1180–1210. [639]
- Robins, J. M. (2004), "Optimal Structural Nested Models for Optimal Sequential Decisions," in *Proceedings of the Second Seattle Symposium* in *Biostatistics*, pp. 189–326, Springer. [639,640]
- Rosen, H. S., and Wu, S. (2004), "Portfolio Choice and Health Status," Journal of Financial Economics, 72, 457–484. [640]
- Rubin, D. B. (1974), "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 66, 688–701. [640]
- She, Y. (2010), "Sparse Regression with Exact Clustering," Electronic Journal of Statistics, 4, 1055–1096. [640,641]
- Shi, C., Fan, A., Song, R., and Lu, W. (2018), "High-Dimensional a-learning for Optimal Dynamic Treatment Regimes," *The Annals of Statistics*, 46, 925–957. [639]
- Shi, C., Song, R., and Lu, W. (2016), "Robust Learning for Optimal Treatment Decision with np-dimensionality," *Electronic Journal of Statistics*, 10, 2894–2921. [644,645,647]
- Singh, S. P., Jaakkola, T., and Jordan, M. I. (1994), "Learning Without State-Estimation in Partially Observable Markovian Decision Processes," in Machine Learning Proceedings 1994, pp. 284–292, Elsevier. [648]

- Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," Journal of the Royal Statistical Society, Series B, 58, 267–288. [640]
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005), "Sparsity and Smoothness via the Fused Lasso," *Journal of the Royal Statistical Society*, Series B, 67, 91–108. [641]
- Tibshirani, R. J., and Taylor, J. (2011), "The Solution Path of the Generalized Lasso," *The Annals of Statistics*, 39, 1335–1371. [640,641,642,643]
- ——— (2012), Degrees of Freedom in Lasso Problems. The Annals of Statistics, 40(2):1198–1232. [642]
- Wager, S., and Athey, S. (2018), "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests," *Journal of the American Sta*tistical Association, 113, 1228–1242. [644]
- Wang, L., Kim, Y., and Li, R. (2013), "Calibrating Non-convex Penalized Regression in Ultra-High Dimension," *The Annals of Statistics*, 41, 2505–2536. [640]
- Watkins, C. (1989), Learning from Delayed Rewards. PhD thesis. King's College, Cambridge. [639]
- Zhang, C.-H. (2010), "Nearly Unbiased Variable Selection Under Minimax Concave Penalty," *The Annals of Statistics*, 38, 894–942. [640,642]
- Zhang, C.-H., and Huang, J. (2008), "The Sparsity and Bias of the Lasso Selection in High-Dimensional Linear Regression," *The Annals of Statistics*, 36, 1567–1594. [640]
- Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," Journal of Machine Learning Research, 7, 2541–2563. [640]
- Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012), "Estimating Individualized Treatment Rules Using Outcome Weighted Learning," *Journal of the American Statistical Association*, 107, 1106–1118. [639]
- Zhou, W., Zhu, R., and Zeng, D. (2021), "A Parsimonious Personalized Dose-Finding Model via Dimension Reduction," *Biometrika*, 108, 643–659. [639]
- Zhu, L., Lu, W., Kosorok, M. R., and Song, R. (2020), "Kernel Assisted Learning for Personalized Dose Finding," in Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 56–65. [639]
- Zhu, W., Zeng, D., and Song, R. (2019), "Proper Inference for Value Function in High-Dimensional q-Learning for Dynamic Treatment Regimes," Journal of the American Statistical Association, 114, 1404– 1417. [639,641]
- Zou, H., and Li, R. (2008), "One-Step Sparse Estimates in Nonconcave Penalized Likelihood Models," *The Annals of Statistics*, 36, 1509–1533.
 [643]