

PROGRAM EVALUATION WITH HIGH-DIMENSIONAL DATA

A. BELLONI, V. CHERNOZHUKOV, I. FERNÁNDEZ-VAL, AND C. HANSEN

ABSTRACT. In this paper, we provide efficient estimators and honest confidence bands for a variety of treatment effects including local average (LATE) and local quantile treatment effects (LQTE) in data-rich environments. We can handle *very many* control variables, *endogenous* receipt of treatment, *heterogeneous* treatment effects, and *function-valued* outcomes. Our framework covers the special case of exogenous receipt of treatment, either conditional on controls or unconditionally as in randomized control trials. In the latter case, our approach produces efficient estimators and honest bands for (functional) average treatment effects (ATE) and quantile treatment effects (QTE). To make informative inference possible, we assume that key reduced form predictive relationships are approximately sparse. This assumption allows the use of regularization and selection methods to estimate those relations, and we provide methods for post-regularization and post-selection inference that are uniformly valid (honest) across a wide-range of models. We show that a key ingredient enabling honest inference is the use of orthogonal or doubly robust moment conditions in estimating certain reduced form functional parameters. We illustrate the use of the proposed methods with an application to estimating the effect of 401(k) eligibility and participation on accumulated assets.

The results on program evaluation are obtained as a consequence of more general results on honest inference in a general moment condition framework, where we work with possibly a continuum of moments. We provide results on honest inference for (function-valued) parameters within this general framework where modern machine learning methods are used to fit the nonparametric/high-dimensional components of the model. These include a number of supporting new results that are of major independent interest: namely, we (1) prove uniform validity of a multiplier bootstrap, (2) offer a uniformly valid functional delta method, and (3) provide results for sparsity-based estimation of regression functions for function-valued outcomes.

1. INTRODUCTION

The goal of many empirical analyses in economics is to understand the causal effect of a treatment such as participation in a government program on economic outcomes. Such analyses are often complicated by the fact that few economic treatments or government policies are randomly assigned. The lack of true random assignment has led to the adoption of a variety of quasi-experimental approaches to estimating treatment effects that are based on observational data. Such approaches include instrumental variable (IV) methods in cases where treatment is not randomly assigned

Date: First version: April 2013. This version: September 20, 2015. We gratefully acknowledge research support from the NSF. We are very grateful to the co-editor, three anonymous referees, Alberto Abadie, Stephane Bonhomme, Matias Cattaneo, Jinyong Hahn, Michael Jansson, Toru Kitagawa, Roger Koenker, Simon Lee, Yuan Liao, Oliver Linton, Blaise Melly, Whitney Newey, Adam Rosen, and seminar participants at the Bristol Econometric Study Group, CEMFI, Cornell-Princeton Conference on "Inference on Non-Standard Problems", Winter 2014 ES meeting, Semi-Plenary Lecture at of ES Summer Meeting 2014, ES World Congress 2015, 2013 Summer NBER Institute, University of Montreal, UIUC, UCL, and University of Warwick for helpful comments. We are especially grateful to Andres Santos for many useful comments.

Key words and phrases. machine learning, endogeneity, local average and quantile treatment effects, instruments, local effects of treatment on the treated, propensity score, Lasso, inference after model selection, moment condition models, moment condition models with a continuum of target parameters, Lasso and Post-Lasso with functional response data, randomized control trials.

but there is some other external variable, such as eligibility for receipt of a government program or service, that is either randomly assigned or the researcher is willing to take as exogenous conditional on the right set of control variables (or simply controls). Another common approach is to assume that the treatment variable itself may be taken as exogenous after conditioning on the right set of controls which leads to regression or matching based methods, among others, for estimating treatment effects.¹

A practical problem empirical researchers face when trying to estimate treatment effects is deciding what conditioning variables to include. When the treatment variable or instrument is not randomly assigned, a researcher must choose what needs to be conditioned on to make the argument that the instrument or treatment is exogenous plausible. Typically, economic intuition will suggest a set of variables that might be important to control for but will not identify exactly which variables are important or the functional form with which variables should enter the model. While less crucial to identifying treatment effects, the problem of selecting controls also arises in situations where the key treatment or instrumental variables are randomly assigned. In these cases, a researcher interested in obtaining precisely estimated policy effects will also typically consider including additional controls to help absorb residual variation. As in the case where including controls is motivated by a desire to make identification of the treatment effect more plausible, one rarely knows exactly which variables will be most useful for accounting for residual variation. In either case, the lack of clear guidance about what variables to use presents the problem of selecting controls from a potentially large set including raw variables available in the data as well as interactions and other transformations of these variables.

In this paper, we consider estimation of the effect of an *endogenous* binary treatment, D , on an outcome, Y , in the presence of a binary instrumental variable, Z , in settings with very many potential controls, $f(X)$. Allowing many potential controls expressly covers both the case where there are simply many controls (where $f(X) = X$) and the case where there are many technical controls $f(X)$ generated as transformations such as powers, b-splines, or interactions of raw controls,² X , along with combinations of the two cases. The notation $f(X)$ naturally accommodates these cases, and we call $f(X)$ the *controls* regardless of the case. We allow for fully *heterogeneous* treatment effects and thus focus on estimation of causal quantities that are appropriate in heterogeneous effects settings such as the local average treatment effect (LATE) or the local quantile treatment effect (LQTE). We focus our discussion on the *endogenous* case where identification is obtained through the use of an instrumental variable, but all results carry through to the *exogenous* case where the treatment is taken as exogenous unconditionally or after conditioning on sufficient controls by simply replacing the instrument with the treatment variable in the estimation and inference methods and in the formal results. In the latter case LATE reduces to the average treatment effect (ATE) and LQTE to the quantile treatment effect (QTE).

¹There is a large literature about estimation of treatment effects. See, for example, the textbook treatments in Angrist and Pischke (2008) or Wooldridge (2010) and the references therein for discussion from an economic perspective.

²See, e.g., Koenker (1988), Newey (1997), Wasserman (2006), Chen (2007), and Tsybakov (2009).

The methodology for estimating treatment effects we consider allows for cases where the number of potential controls, $p := \dim f(X)$, is much larger than the sample size, n . Of course, informative inference about causal parameters cannot proceed allowing for $p \gg n$ without further restrictions. We impose sufficient structure through the assumption that reduced form relationships such as the conditional expectations $E_P[D|X]$, $E_P[Z|X]$, and $E_P[Y|X]$ are approximately sparse. Intuitively, approximate sparsity imposes that these reduced form relationships can be represented up to a small approximation error as a linear combination, possibly inside of a known link function such as the logistic function, of a number $s \ll n$ of the variables in $f(X)$ whose identities are *a priori* unknown to the researcher. This assumption allows us to use methods for estimating models in high-dimensional sparse settings that are known to have good prediction properties to estimate the fundamental reduced form relationships. We may then use these estimated reduced form quantities as inputs to estimating the causal parameters of interest. Approaching the problem of estimating treatment effects within this framework allows us to accommodate the realistic scenario in which a researcher is unsure about exactly which confounding variables or transformations of these confounds are important and so must search among a broad set of controls.

Valid inference following model selection is non-trivial. Direct application of usual inference procedures following model selection does not provide valid inference about causal parameters even in low-dimensional settings, such as when there is only a single control, unless one assumes sufficient structure on the model that perfect model selection is possible. Such structure can be restrictive and seems unlikely to be satisfied in many economic applications. For example, a typical condition that allows perfect model selection is the “beta-min” condition, which requires that all but a small number of coefficients are exactly zero and that the non-zero coefficients are all large enough that they can be distinguished from zero with probability very near one in finite samples. Such a condition rules out the possibility that there may be some variables which have moderate, but non-zero, partial effects. Ignoring such variables may result in large omitted variables bias that has a substantive impact on estimation and inference regarding individual model parameters; see Leeb and Pötscher (2008a; 2008b); Pötscher (2009); and Belloni, Chernozhukov, and Hansen (2013a; 2014a).

The first main contribution of this paper is providing inferential procedures for key parameters used in program evaluation that are theoretically valid within approximately sparse models allowing for imperfect model selection. Our procedures build upon Belloni et al. (2010) and Belloni et al. (2012), who were the first to demonstrate in a highly specialized context, that valid inference can proceed following model selection allowing for model selection mistakes under two conditions. We formulate and extend these two conditions to a rather general moment-condition framework (e.g., Hansen (1982) and Hansen and Singleton (1982)) as follows. First, estimation should be based upon “orthogonal” moment conditions that are first-order insensitive to changes in the values of nuisance parameters that will be estimated using high-dimensional methods. Specifically, if the target parameter value α_0 is identified via the moment condition

$$E_P \psi(W, \alpha_0, h_0) = 0, \tag{1.1}$$

where h_0 is a function-valued nuisance parameter estimated via a post-model-selection or regularization method, one needs to use a moment function, ψ , such that the corresponding moment condition is orthogonal with respect to perturbations of h around h_0 . More formally, the moment condition should satisfy

$$\partial_h[\mathbb{E}_P\psi(W, \alpha_0, h)]_{h=h_0} = 0, \quad (1.2)$$

where ∂_h is a functional derivative operator with respect to h restricted to directions of possible deviations of estimators of h_0 from h_0 . Second, one needs to ensure that the model selection mistakes occurring in the estimation of nuisance parameters are uniformly “moderately” small with respect to the underlying model. Specifically, we will require that the nuisance parameter h_0 is estimated at the rate $o(n^{-1/4})$, which ensures small bias, and that the estimator takes values in a space whose entropy does not grow too fast, which ensures no overfitting. In this paper, we establish that building estimators based upon moment conditions with the orthogonality condition (1.2) holding ensures that crude estimation of h_0 via post-selection or other regularization methods has an asymptotically negligible effect on the estimation of α_0 in general frameworks. It then follows that we can form a regular, root- n consistent estimator of α_0 , uniformly with respect to the underlying model.

In the general endogenous treatment effects setting, we build the moment conditions satisfying (1.2) from the efficient influence functions for certain reduced form parameters, building upon Hahn (1998). We illustrate how orthogonal moment conditions coupled with methods developed for forecasting in high-dimensional approximately sparse models can be used to estimate and obtain valid inferential statements about a wide variety of structural/treatment effects. We formally demonstrate the uniform validity of the resulting inference within a broad class of approximately sparse models including models where perfect model selection is theoretically impossible. An important feature of our main theoretical results is that they cover the use of variable selection for *functional response data* using ℓ_1 -penalized methods. Functional response data arises, for example, when one is interested in the LQTE at not just a single quantile but over a range of quantile indices; this in turn necessitates looking at the functional dependent variable $u \mapsto 1(Y \leq u)$, where u denotes various levels that Y can cross. Considering such functional response data allows us to provide a unified inference procedure for interesting quantities such as the (local) distributional and quantile effects of the treatment, including simpler important parameters such as LQTE at a given quantile as a special case.

The second main contribution of this paper, which is not only used to establish the properties of treatment effects estimators but is also of independent interest, is providing a general set of results for uniformly valid estimation and inference methods in modern moment-condition problems in econometrics allowing for both smooth and non-smooth moment functions. A key element to our establishing uniform validity of post-regularization inference is again the use of orthogonal moment conditions. In the general framework we consider, we may have a continuum of target parameters identified via a continuum of moment conditions that involve a continuum of nuisance functions that will be estimated via modern high-dimensional methods such as Lasso, Post-Lasso, or their variants. These results contain the results on treatment effects relevant for program evaluation, particularly

the results for distributional and quantile effects, as a leading special case. These results are also immediately useful in other contexts such as nonseparable quantile models as in Chernozhukov and Hansen (2005), Chernozhukov and Hansen (2006), Chesher (2003), and Imbens and Newey (2009); semiparametric and partially identified models as in Escanciano and Zhu (2013); and many others. In our results, we first establish a functional central limit theorem for the continuum of target parameters and show that this functional central limit theorem holds uniformly in a wide range of data-generating processes P with approximately sparse continua of nuisance functions. Second, we establish a functional central limit theorem for the multiplier bootstrap that resamples the first order approximations to the standardized estimators and demonstrate its uniform-in- P validity. These uniformity results build upon and complement those given in Romano and Shaikh (2012) for the empirical bootstrap. Third, we establish a functional delta method for smooth functionals of the continuum of target parameters and a functional delta method for the multiplier bootstrap of these smooth functionals, both of which hold uniformly in P , using an appropriately strengthened notion of Hadamard differentiability. All of these results are new and are of independent interest outside of the treatment effects focus of this paper.

We illustrate the use of our methods by estimating the effect of 401(k) eligibility and 401(k) participation on measures of accumulated assets as in Chernozhukov and Hansen (2004).³ Similar to Chernozhukov and Hansen (2004), we provide estimates of ATE and QTE of 401(k) eligibility and of LATE and LQTE of 401(k) participation. We differ from this previous work by using the high-dimensional methods developed in this paper to allow ourselves to consider a broader set of controls than has previously been considered. We find that 401(k) participation has a moderate impact on accumulated financial assets at low quantiles while appearing to have a much larger impact at high quantiles. Interpreting the quantile index as “preference for savings” as in Chernozhukov and Hansen (2004), this pattern suggests that 401(k) participation has little causal impact on the accumulated financial assets of those with low desire to save but a much larger impact on those with stronger preferences for saving. It is interesting that these results are similar to those in Chernozhukov and Hansen (2004) despite allowing for a much richer set of controls.

Links to the literature. The orthogonality condition embodied in (1.2) has a long history in statistics and econometrics. For example, this type of orthogonality was used by Neyman (1979) in low-dimensional settings to deal with crudely estimated parametric nuisance parameters. See also Newey (1990), Andrews (1994b), Newey (1994), Robins and Rotnitzky (1995), and Linton (1996) for the use of this condition in semi-parametric problems.

To the best of our knowledge, Belloni et al. (2010) and Belloni et al. (2012) were the first to use the orthogonality (1.2) to expressly address the question of the uniform post-selection inference without imposing the “beta-min” conditions, either in high-dimensional settings with $p \gg n$ or in low-dimensional settings with $p \ll n$. They applied it to the specific problem of the linear instrumental variables model with many instruments where the nuisance function h_0 is the optimal

³See also Poterba, Venti, and Wise (1994; 1995; 1996; 2001), Abadie (2003), Benjamin (2003), and Ogburn et al. (2015) among others.

instrument estimated by Lasso or Post-Lasso methods and α_0 is the coefficient of the endogenous regressor. Belloni et al. (2013a) and Belloni et al. (2014a) also exploited this approach to develop a double-selection method that yields valid post-selection inference on the parameters of the linear part of a partially linear model and on average treatment effects when the treatment is binary and *exogenous* conditional on controls in both the $p \gg n$ and the $p \ll n$ setting.⁴ Subsequently, Farrell (2013) extended the results of Belloni et al. (2013a) and Belloni et al. (2014a) to estimation of ATE when the treatment is multivalued and exogenous conditional on controls using group penalization for selection. Note that this previous work on treatment effects covers only the exogenous case and does not allow for functional responses which are necessary, for example, for working with distributional or quantile treatment effects.

Our work also contributes to the line of research on obtaining \sqrt{n} -consistent and asymptotically normal estimates for low-dimensional components within traditional semiparametric frameworks, specifically in the important work by Bickel (1982), Robinson (1988), Newey (1990), van der Vaart (1991), Andrews (1994b), Newey (1994), Ai and Chen (2003, 2012), and Chen et al. (2003). The major difference is that we allow for the possibility of using modern high-dimensional methods, a.k.a. machine learning methods, for modeling and fitting the non-parametric (or high-dimensional) components of the model. In contrast to the former literature, we expressly allow for data-driven choice of the approximating model for the high-dimensional component, which addresses a crucial problem that arises in empirical work. Moreover, recent methods based on ℓ_1 -penalization, upon which we focus in this paper, allow for much more flexible modeling of the non-parametric/high-dimensional parts of the model.⁵ Our general theory in Section 5 also allows, in principle, for a wide variety of traditional and machine learning methods, including those that do not rely on approximate sparsity, as long as the methods have good approximation ability and their entropy does not increase too rapidly with the sample size - i.e. they do not overfit.

The paper also generates a number of new results on sparse estimation with functional response data. These results are of independent interest in themselves, and they build upon the work of Belloni and Chernozhukov (2011) who provided rates of convergence for variable selection when one is interested in estimating the quantile regression process with exogenous variables. More generally, this theoretical work complements and extends the rapidly growing set of results for ℓ_1 -penalized estimation methods; see, for example, Frank and Friedman (1993); Tibshirani (1996); Fan and Li (2001); Zou (2006); Candès and Tao (2007); van de Geer (2008); Huang et al. (2008); Bickel et al. (2009); Meinshausen and Yu (2009); Bach (2010); Huang et al. (2010); Belloni and Chernozhukov (2011); Kato (2011); Belloni et al. (2012); Belloni and Chernozhukov (2013); Belloni et al. (2013b); Belloni et al. (2013c); Caner and Zhang (2014); and the references therein.

⁴Note that these results as well as results of this paper on the uniform post-selection inference in moment-condition problems are new for either $p \ll n$ or $p \gg n$ settings. The results also apply to arbitrary model selection devices, such as the Dantzig selector, Square-Root-Lasso, or Adaptive Lasso, that are able to select good sparse approximating models, and “moderate” model selection errors are explicitly allowed in the paper.

⁵See, for instance, Belloni et al. (2012) and Belloni et al. (2014b) for a formalization of this claim in terms of rearranged Sobolev spaces where it is shown that traditional methods can fail to be consistent while ℓ_1 -penalized methods remain consistent and have good rates of convergence.

Plan of the Paper. Section 2 introduces the structural parameters for policy evaluation and relates these parameters to reduced form functions. Section 3 describes a three step procedure to estimate and make inference on the structural parameters and functionals of these parameters, and Section 4 provides asymptotic theory. Section 5 generalizes the setting and results to moment-condition problems with a continuum of structural parameters and a continuum of reduced form functions. Section 6 derives general asymptotic theory for the Lasso and post-Lasso estimators for functional response data used in the estimation of the reduced form functions. Section 7 presents the empirical application. We provide notation, proofs of key results, and details about implementation of the methods in the empirical example in Appendices A–E. An on-line Supplementary Appendix provides all remaining proofs, additional technical material, and results from a small Monte Carlo simulation (Belloni et al., 2015).

2. THE TREATMENT EFFECTS SETTING AND TARGET PARAMETERS

2.1. Observables and Reduced Form Parameters. The observed random variables consist of $((Y_u)_{u \in \mathcal{U}}, X, Z, D)$. The outcome variable of interest Y_u is indexed by $u \in \mathcal{U}$. We give examples of the index u below. The variable $D \in \mathcal{D} = \{0, 1\}$ is a binary indicator of the receipt of a treatment or participation in a program. It will typically be treated as endogenous; that is, we will typically view the treatment as assigned non-randomly with respect to the outcome. The instrumental variable $Z \in \mathcal{Z} = \{0, 1\}$ is a binary indicator, such as an offer of participation, that is assumed to be randomly assigned conditional on the observable covariates X with support \mathcal{X} .⁶ For example, we argue that 401(k) eligibility can be considered exogenous only after conditioning on income and other individual characteristics in the empirical application. The notions of exogeneity and endogeneity we employ are standard and thus omitted.⁷

The indexing of the outcome Y_u by u is useful to analyze functional data. For example, Y_u could represent an outcome falling short of a threshold, namely $Y_u = 1(Y \leq u)$, in the context of distributional analysis; Y_u could be a height indexed by age u in growth charts analysis; or Y_u could be a health outcome indexed by a dosage u in dosage response studies. Our framework is tailored for such functional response data. The special case with no index is included by simply considering \mathcal{U} to be a singleton set.

We make use of two key types of reduced form parameters for estimating the structural parameters of interest – (local) treatment effects and related quantities. These reduced form parameters are defined as

$$\alpha_V(z) := E_P[g_V(z, X)] \quad \text{and} \quad \gamma_V := E_P[V], \quad (2.1)$$

⁶Of course, by “randomly assigned” we mean independently of outcomes conditional on the covariates.

⁷For completeness, we provide a review of these conditions as well as restate standard conditions that are sufficient for a causal interpretation of the target parameters in the Supplementary Appendix.

where $z = 0$ or $z = 1$ are the fixed values of Z .⁸ The function g_V maps $\mathcal{Z}\mathcal{X}$, the support of the vector (Z, X) , to the real line \mathbb{R} and is defined as

$$g_V(z, x) := E_P[V|Z = z, X = x]. \quad (2.2)$$

We use V to denote a target variable whose identity may change depending on the context such as $V = \mathbf{1}_d(D)Y_u$ or $V = \mathbf{1}_d(D)$ where $\mathbf{1}_d(D) := 1(D = d)$ is the indicator function.

All the structural parameters we consider are smooth functionals of these reduced-form parameters. In our approach to estimating treatment effects, we estimate the key reduced form parameter $\alpha_V(z)$ using modern methods to deal with high-dimensional data coupled with orthogonal estimating equations. The orthogonality property allows us to deal with the “non-regular” nature of penalized and post-selection estimators which do not admit linearizations except under very restrictive conditions. The use of regularization by model selection or penalization is in turn motivated by the desire to accommodate high-dimensional data.

2.2. Target Structural Parameters – Local Treatment Effects. The reduced form parameters defined in (2.1) are key because the structural parameters of interest are functionals of these elementary objects. The local average structural function (LASF) defined as

$$\theta_{Y_u}(d) = \frac{\alpha_{\mathbf{1}_d(D)Y_u}(1) - \alpha_{\mathbf{1}_d(D)Y_u}(0)}{\alpha_{\mathbf{1}_d(D)}(1) - \alpha_{\mathbf{1}_d(D)}(0)}, \quad d \in \{0, 1\} \quad (2.3)$$

underlies the formation of many commonly used treatment effects. Under standard assumptions, the LASF identifies average potential outcomes for the group of *compliers*, individuals whose treatment status may be influenced by variation in the instrument, in the treated and non-treated states; see, e.g. Abadie (2002; 2003). The local average treatment effect (LATE) of Imbens and Angrist (1994) corresponds to the difference of the two values of the LASF:

$$\theta_{Y_u}(1) - \theta_{Y_u}(0). \quad (2.4)$$

The term local designates that this parameter does not measure the effect on the entire population but rather measures the effect on the subpopulation of compliers.

When there is no endogeneity, formally when $D \equiv Z$, the LASF and LATE become the average structural function (ASF) and average treatment effect (ATE) on the entire population. Thus, our results cover this situation as a special case where the ASF and ATE simplify to

$$\theta_{Y_u}(z) = \alpha_{Y_u}(z), \quad \theta_{Y_u}(1) - \theta_{Y_u}(0) = \alpha_{Y_u}(1) - \alpha_{Y_u}(0). \quad (2.5)$$

We also note that the impact of the instrument Z itself may be of interest since Z often encodes an offer of participation in a program. In this case, the parameters of interest are again simply the reduced form parameters

$$\alpha_{Y_u}(z), \quad \alpha_{Y_u}(1) - \alpha_{Y_u}(0).$$

⁸The expectation that defines $\alpha_V(z)$ is well-defined under the standard support condition $0 < P_P(Z = 1 | X) < 1$ a.s. This condition is standard in treatment effects estimation; see, e.g., the supplementary appendix. We impose this condition in Assumption 4.1.

Thus, the LASF and LATE are primary targets of interest in this paper, and the ASF and ATE are subsumed as special cases.

2.2.1. *Local Distribution and Quantile Treatment Effects.* Setting $Y_u = Y$ in (2.3) and (2.4) provides the conventional LASF and LATE. An important generalization arises by letting $Y_u = 1(Y \leq u)$ be the indicator of the outcome of interest falling below a threshold $u \in \mathbb{R}$. In this case, the family of effects

$$(\theta_{Y_u}(1) - \theta_{Y_u}(0))_{u \in \mathbb{R}}, \quad (2.6)$$

describe the local distribution treatment effects (LDTE). Similarly, we can look at the quantile left-inverse transform of the curve $u \mapsto \theta_{Y_u}(d)$,

$$\theta_Y^{\leftarrow}(\tau, d) := \inf\{u \in \mathbb{R} : \theta_{Y_u}(d) \geq \tau\}, \quad (2.7)$$

and examine the family of local quantile treatment effects (LQTE):

$$(\theta_Y^{\leftarrow}(\tau, 1) - \theta_Y^{\leftarrow}(\tau, 0))_{\tau \in (0,1)}. \quad (2.8)$$

The LQTE identify the differences of quantiles between the distribution of potential outcomes in the treated and non-treated states for compliers.

2.3. Target Structural Parameters – Local Treatment Effects on the Treated. We may also be interested in local treatment effects on the treated. The key object in defining these effects is the local average structural function on the treated (LASF-T) which is defined by its two values:

$$\vartheta_{Y_u}(d) = \frac{\gamma_{\mathbf{1}_d(D)Y_u} - \alpha_{\mathbf{1}_d(D)Y_u}(0)}{\gamma_{\mathbf{1}_d(D)} - \alpha_{\mathbf{1}_d(D)}(0)}, \quad d \in \{0, 1\}. \quad (2.9)$$

The LASF-T identifies average potential outcomes for the group of *treated compliers* in the treated and non-treated states under standard assumptions. The local average treatment effect on the treated (LATE-T) introduced in Hong and Nekipelov (2010) and Frölich and Melly (2013) is the difference of two values of the LASF-T:

$$\vartheta_{Y_u}(1) - \vartheta_{Y_u}(0). \quad (2.10)$$

The LATE-T may be of interest because it measures the average treatment effect for *treated compliers*, namely the subgroup of compliers that actually receive the treatment.

When the treatment is assigned randomly given controls so we can take $D = Z$, the LASF-T and LATE-T become the average structural function on the treated (ASF-T) and average treatment effect on the treated (ATE-T). In this special case, the ASF-T and ATE-T simplify to

$$\vartheta_{Y_u}(1) = \frac{\gamma_{\mathbf{1}_1(D)Y_u}}{\gamma_{\mathbf{1}_1(D)}}, \quad \vartheta_{Y_u}(0) = \frac{\gamma_{\mathbf{1}_0(D)Y_u} - \alpha_{Y_u}(0)}{\gamma_{\mathbf{1}_0(D)} - 1}, \quad \vartheta_{Y_u}(1) - \vartheta_{Y_u}(0); \quad (2.11)$$

and we can use our results to provide estimation and inference methods for these quantities.

2.3.1. *Local Distribution and Quantile Treatment Effects on the Treated.* Local distribution treatment effects on the treated (LDTE-T) and local quantile treatment effects on the treated (LQTE-T) can also be defined. As in Section 2.2.1, we let $Y_u = 1(Y \leq u)$ be the indicator of the outcome of interest falling below a threshold u . The family of treatment effects

$$(\vartheta_{Y_u}(1) - \vartheta_{Y_u}(0))_{u \in \mathbb{R}} \quad (2.12)$$

then describes the LDTE-T. We can also use the quantile left-inverse transform of the curve $u \mapsto \vartheta_{Y_u}(d)$, namely $\vartheta_Y^{\leftarrow}(\tau, d) := \inf\{u \in \mathbb{R} : \vartheta_{Y_u}(d) \geq \tau\}$, and define the LQTE-T:

$$(\vartheta_Y^{\leftarrow}(\tau, 1) - \vartheta_Y^{\leftarrow}(\tau, 0))_{\tau \in (0,1)}. \quad (2.13)$$

Under conditional exogeneity LQTE and LQTE-T reduce to the quantile treatment effects (QTE) and quantile treatment effects on the treated (QTE-T) (Koenker, 2005, Chap. 2).

3. ESTIMATION OF REDUCED-FORM AND STRUCTURAL PARAMETERS IN A DATA-RICH ENVIRONMENT

The key objects used to define the structural parameters in Section 2 are the expectations

$$\alpha_V(z) = \mathbb{E}_P[g_V(z, X)] \text{ and } \gamma_V = \mathbb{E}_P[V], \quad (3.1)$$

where $g_V(z, X) = \mathbb{E}_P[V|Z = z, X]$ and V denotes a variable whose identity will change with the context. Specifically, we shall vary V over the set \mathcal{V}_u :

$$V \in \mathcal{V}_u := \{V_{uj}\}_{j=1}^5 := \{Y_u, \mathbf{1}_0(D)Y_u, \mathbf{1}_0(D), \mathbf{1}_1(D)Y_u, \mathbf{1}_1(D)\}. \quad (3.2)$$

It is clear that $g_V(z, X)$ will play an important role in estimating $\alpha_V(z)$. A related function that will also play an important role in forming a robust estimation strategy is the propensity score $m_Z : \mathcal{Z}\mathcal{X} \mapsto \mathbb{R}$ defined by

$$m_Z(z, x) := \mathbb{P}_P[Z = z|X = x]. \quad (3.3)$$

We will denote other potential values for the functions g_V and m_Z by the parameters g and m , respectively. We can then estimate $\alpha_V(z)$ by estimating g_V and m_Z using high-dimensional modeling and estimation methods.⁹

In the rest of this section, we describe the estimation of the reduced-form and structural parameters. The estimation method consists of 3 steps:

- 1) Estimate the predictive relationships m_Z and g_V using high-dimensional nonparametric methods with model selection.
- 2) Estimate the reduced form parameters α_V and γ_V using orthogonal estimating equations to immunize the reduced form estimators to imperfect model selection in the first step.
- 3) Estimate the structural parameters and effects via the plug-in rule.

⁹Note that there is an alternative approach based on decomposing g_V as $g_V(z, x) = \sum_{d=0}^1 e_V(d, z, x)l_D(d, z, x)$ where the regression functions e_V and l_D map the support of (D, Z, X) , $\mathcal{D}\mathcal{Z}\mathcal{X}$, to the real line and are defined by $e_V(d, z, x) := \mathbb{E}_P[V|D = d, Z = z, X = x]$ and $l_D(d, z, x) := \mathbb{P}_P[D = d|Z = z, X = x]$. We provide some discussion of this approach in the supplementary appendix.

3.1. First Step: Modeling and Estimating g_V and m_Z . In this section, we discuss estimation of the conditional expectation functions g_V and m_Z . Since these functions are unknown and potentially complicated, we use a generalized linear combination of a large number of control terms

$$f(X) = (f_j(X))_{j=1}^p, \quad (3.4)$$

to approximate g_V and m_Z . Specifically, we use

$$g_V(z, x) =: \Lambda_V[f(z, x)' \beta_V] + r_V(z, x), \quad (3.5)$$

$$f(z, x) := ((1-z)f(x)', zf(x)')', \quad \beta_V := (\beta_V(0)', \beta_V(1)')', \quad (3.6)$$

$$m_Z(1, x) =: \Lambda_Z[f(x)' \beta_Z] + r_Z(x), \quad m_Z(0, x) = 1 - \Lambda_Z[f(x)' \beta_Z] - r_Z(x). \quad (3.7)$$

In these equations, $r_V(z, x)$ and $r_Z(x)$ are approximation errors, and the functions $\Lambda_V(f(z, x)' \beta_V)$ and $\Lambda_Z(f(x)' \beta_Z)$ are generalized linear approximations to the target functions $g_V(z, x)$ and $m_Z(1, x)$. The functions Λ_V and Λ_Z are taken to be known link functions Λ . The most common example is the linear link $\Lambda(u) = u$. When the response variable is binary, we may also use the logistic link $\Lambda(u) = \Lambda_0(u) = e^u / (1 + e^u)$ and its complement $1 - \Lambda_0(u)$ or the probit link $\Lambda(u) = \Phi(u) = (2\pi)^{-1/2} \int_{-\infty}^u e^{-z^2/2} dz$ and its complement $1 - \Phi(u)$. For clarity, we use links from the finite set $\mathcal{L} = \{\text{Id}, \Phi, 1 - \Phi, \Lambda_0, 1 - \Lambda_0\}$ where Id is the identity (linear) link.

As discussed in the Introduction, the dictionary of controls, denoted by $f(X)$, can be “rich” in the sense that its dimension $p = p_n$ may be large relative to the sample size. Specifically, our results require only that $\log p = o(n^{1/3})$ along with other technical conditions. We also note that the functions f forming the dictionary can depend on n , but we suppress this dependence.

Having very many controls $f(X)$ creates a challenge for estimation and inference. A useful condition that makes it possible to perform constructive estimation and inference in such cases is termed approximate sparsity or simply sparsity. Sparsity imposes that there exist approximations of the form given in (3.5)-(3.7) that require only a small number of non-zero coefficients to render the approximation errors small relative to estimation error. More formally, sparsity relies on two conditions. First, there must exist β_V and β_Z such that, for all $V \in \mathcal{V} := \{\mathcal{V}_u : u \in \mathcal{U}\}$,

$$\|\beta_V\|_0 + \|\beta_Z\|_0 \leq s, \quad (3.8)$$

where $\|x\|_0$ is the number of non-zero components of vector x and all other norms we use are defined in Appendix A. That is, there are at most $s = s_n \ll n$ components of $f(Z, X)$ and $f(X)$ with nonzero coefficient in the approximations to g_V and m_Z . Second, the sparsity condition requires that the size of the resulting approximation errors is small compared to the conjectured size of the estimation error; namely, for all $V \in \mathcal{V}$,

$$\{\mathbb{E}_P[r_V^2(Z, X)]\}^{1/2} + \{\mathbb{E}_P[r_Z^2(X)]\}^{1/2} \lesssim \sqrt{s/n}. \quad (3.9)$$

Note that the size of the approximating model $s = s_n$ can grow with n just as in standard series estimation, subject to the rate condition

$$s^2 \log^2(p \vee n) \log^2 n/n \rightarrow 0.$$

These conditions ensure that the functions g_V and m_Z are estimable at $o(n^{-1/4})$ rate and are used to derive asymptotic normality results for the structural and reduced-form parameter estimators. They could be relaxed through the use of sample splitting methods as in Belloni et al. (2012).

The high-dimensional-sparse-model framework outlined above extends the standard framework in the program evaluation literature which assumes both that the identities of the relevant controls are known and that the number of such controls s is small relative to the sample size.¹⁰ Instead, we assume that there are many, p , potential controls of which at most s controls suffice to achieve a desirable approximation to the unknown functions g_V and m_Z ; and we allow the identity and number of these controls to be unknown. Relying on this assumed sparsity, we use selection methods to choose approximately the right set of controls.

Current estimation methods that exploit approximate sparsity employ different types of regularization aimed at producing estimators that theoretically perform well in high-dimensional settings while remaining computationally tractable. Many widely used methods are based on ℓ_1 -penalization. The Lasso method is one such commonly used approach that adds a penalty for the weighted sum of the absolute values of the model parameters to the usual objective function of an M-estimator. A related approach is the Post-Lasso method which performs re-estimation of the model after selection of variables by Lasso. These methods are discussed at length in recent papers and review articles; see, for example, Belloni et al. (2013a).

In the following, we outline the general features of the Lasso and Post-Lasso methods focusing on estimation of g_V . Given the data $(\tilde{Y}_i, \tilde{X}_i)_{i=1}^n = (V_i, f(Z_i, X_i))_{i=1}^n$, the Lasso estimator $\hat{\beta}_V$ solves

$$\hat{\beta}_V \in \arg \min_{\beta \in \mathbb{R}^{\dim(\tilde{X})}} \left(\mathbb{E}_n[M(\tilde{Y}, \tilde{X}'\beta)] + \frac{\lambda}{n} \|\hat{\Psi}\beta\|_1 \right), \quad (3.10)$$

where $\hat{\Psi} = \text{diag}(\hat{l}_1, \dots, \hat{l}_{\dim(\tilde{X})})$ is a diagonal matrix of data-dependent penalty loadings, $M(y, t) = (y-t)^2/2$ in the case of linear regression, and $M(y, t) = -\{1(y=1) \log \Lambda(t) + 1(y=0) \log(1-\Lambda(t))\}$ in the case of binary regression. The penalty level, λ , and loadings, \hat{l}_j , $j = 1, \dots, \dim(\tilde{X})$, are selected to guarantee good theoretical properties of the method. We provide further discussion of these methods for estimation of a continuum of functions in Section 6, and we specify detailed implementation algorithms used in the empirical example in Appendix F. A key consideration in this paper is that the penalty level needs to be set to account for the fact that we will be simultaneously estimating potentially a *continuum* of Lasso regressions since our V varies over the list \mathcal{V}_u with u varying over the index set \mathcal{U} .

The Post-Lasso method uses $\hat{\beta}_V$ solely as a model selection device. Specifically, it makes use of the labels of the regressors with non-zero estimated coefficients, $\hat{I}_V := \text{supp}(\hat{\beta}_V)$. The Post-Lasso estimator is then a solution to

$$\tilde{\beta}_V \in \arg \min_{\beta \in \mathbb{R}^{\dim(\tilde{X})}} \left(\mathbb{E}_n[M(\tilde{Y}, \tilde{X}'\beta)] : \beta_j = 0, j \notin \hat{I}_V \right). \quad (3.11)$$

¹⁰For example, one would select a set of basis functions, $\{f_j(X)\}_{j=1}^\infty$, such as power series or splines and then use only the first $s \ll n$ terms in the basis under the assumption that $s^C/n \rightarrow 0$ for some number C whose value depends on the specific context in a standard nonparametric approach using series.

A main contribution of this paper is establishing that the estimator $\hat{g}_V(Z, X) = \Lambda(f(Z, X)' \bar{\beta}_V)$ of the regression function $g_V(Z, X)$, where $\bar{\beta}_V = \hat{\beta}_V$ or $\bar{\beta}_V = \tilde{\beta}_V$, achieves the near oracle rate of convergence $\sqrt{(s \log p)/n}$ and maintains desirable theoretic properties while allowing for a *continuum* of response variables.

Estimation of m_Z proceeds similarly. The Lasso estimator $\hat{\beta}_Z$ and Post-Lasso estimator $\tilde{\beta}_Z$ are defined analogously to $\hat{\beta}_V$ and $\tilde{\beta}_V$ using the data $(\tilde{Y}_i, \tilde{X}_i)_{i=1}^n = (Z_i, f(X_i))_{i=1}^n$. The estimator $\hat{m}_Z(1, X) = \Lambda_Z(f(X)' \bar{\beta}_Z)$ of $m_Z(X)$, with $\bar{\beta}_Z = \hat{\beta}_Z$ or $\bar{\beta}_Z = \tilde{\beta}_Z$, also achieves the near oracle rate of convergence $\sqrt{(s \log p)/n}$ and has other good theoretic properties. The estimator of $\hat{m}_Z(0, X)$ is then formed as $1 - \hat{m}_Z(1, X)$.

3.2. Second Step: Robust Estimation of the Reduced-Form Parameters $\alpha_V(z)$ and γ_V .

Estimation of the key quantities $\alpha_V(z)$ will make heavy use of orthogonal moment functions as defined in (1.2). These moment functions are closely tied to efficient influence functions, where efficiency is in the sense of locally minimax semi-parametric efficiency. The use of these functions will deliver robustness with respect to the non-regularity of the post-selection and penalized estimators needed to manage high-dimensional data. The use of these functions also automatically delivers semi-parametric efficiency for estimating and performing inference on the reduced-form parameters and their smooth transformations – the structural parameters.

The efficient influence function and orthogonal moment function for $\alpha_V(z)$, $z \in \mathcal{Z} = \{0, 1\}$, are given respectively by

$$\psi_{V,z}^\alpha(W) := \psi_{V,z,g_V,m_Z}^\alpha(W, \alpha_V(z)) \quad \text{and} \quad (3.12)$$

$$\psi_{V,z,g,m}^\alpha(W, \alpha) := \frac{1(Z=z)(V - g(z, X))}{m(z, X)} + g(z, X) - \alpha. \quad (3.13)$$

This efficient influence function was derived by Hahn (1998); it has recently been used by Cattaneo (2010) in the series context (with $p \ll n$) and Rothe and Firpo (2013) in the kernel context. The efficient influence function and the moment function for γ_V are trivially given by

$$\psi_V^\gamma(W) := \psi_V^\gamma(W, \gamma_V), \quad \text{and} \quad \psi_V^\gamma(W, \gamma) := V - \gamma. \quad (3.14)$$

We then define estimators of the reduced-form parameters $\alpha_V(z)$ and $\gamma_V(z)$ as solutions $\alpha = \hat{\alpha}_V(z)$ and $\gamma = \hat{\gamma}_V$ to the equations

$$\mathbb{E}_n[\psi_{V,z,\hat{g}_V,\hat{m}_Z}^\alpha(W, \alpha)] = 0, \quad \mathbb{E}_n[\psi_V^\gamma(W, \gamma)] = 0, \quad (3.15)$$

where \hat{g}_V and \hat{m}_Z are constructed as in Section 3.1. We apply this procedure to each variable name $V \in \mathcal{V}_u$ and obtain the estimator¹¹

$$\hat{\rho}_u := (\{\hat{\alpha}_V(0), \hat{\alpha}_V(1), \hat{\gamma}_V\})_{V \in \mathcal{V}_u} \quad \text{of} \quad \rho_u := (\{\alpha_V(0), \alpha_V(1), \gamma_V\})_{V \in \mathcal{V}_u}. \quad (3.16)$$

The estimator and the parameter are vectors in \mathbb{R}^{d_ρ} with dimension $d_\rho = 3 \times \dim \mathcal{V}_u = 15$.

¹¹By default notation, $(a_j)_{j \in \mathcal{J}}$ returns a column vector produced by stacking components together in some consistent order.

In the next section, we formally establish a principal result which shows that

$$\sqrt{n}(\widehat{\rho}_u - \rho_u) \rightsquigarrow N(0, \text{Var}_P(\psi_u^\rho)), \quad \psi_u^\rho := (\{\psi_{V,0}^\alpha, \psi_{V,1}^\alpha, \psi_V^\gamma\})_{V \in \mathcal{V}_u}, \quad \text{uniformly in } P \in \mathcal{P}_n, \quad (3.17)$$

where \mathcal{P}_n is a rich set of data generating processes P . The notation “ $Z_{n,P} \rightsquigarrow Z_P$ uniformly in $P \in \mathcal{P}_n$ ” is defined formally in Appendix A and can be read as “ $Z_{n,P}$ is approximately distributed as Z_P uniformly in $P \in \mathcal{P}_n$.” This usage corresponds to the usual notion of asymptotic distribution extended to handle uniformity in P . Here \mathcal{P}_n is a “rich” set of data generating processes P which includes cases where perfect model selection is impossible theoretically.

We then stack all the reduced form estimators and parameters over $u \in \mathcal{U}$ as

$$\widehat{\rho} = (\widehat{\rho}_u)_{u \in \mathcal{U}} \quad \text{and} \quad \rho = (\rho_u)_{u \in \mathcal{U}},$$

giving rise to the empirical reduced-form process $\widehat{\rho}$ and the reduced-form function-valued parameter ρ . We establish that $\sqrt{n}(\widehat{\rho} - \rho)$ is asymptotically Gaussian: In $\ell^\infty(\mathcal{U})^{d_\rho}$,

$$\sqrt{n}(\widehat{\rho} - \rho) \rightsquigarrow Z_P := (\mathbb{G}_P \psi_u^\rho)_{u \in \mathcal{U}}, \quad \text{uniformly in } P \in \mathcal{P}_n, \quad (3.18)$$

where \mathbb{G}_P denotes the P -Brownian bridge (van der Vaart and Wellner, 1996, p. 81–82). This result contains (3.17) as a special case and again allows \mathcal{P}_n to be a “rich” set of data generating processes P that includes cases where perfect model selection is impossible theoretically. Importantly, this result verifies that the functional central limit theorem applies to the reduced-form estimators in the presence of possible model selection mistakes.

Since some of our objects of interest are complicated, inference can be facilitated by a multiplier bootstrap method Giné and Zinn (1984). We define $\widehat{\rho}^* = (\widehat{\rho}_u^*)_{u \in \mathcal{U}}$, a bootstrap draw of $\widehat{\rho}$, via

$$\widehat{\rho}_u^* = \widehat{\rho}_u + n^{-1} \sum_{i=1}^n \xi_i \widehat{\psi}_u^\rho(W_i). \quad (3.19)$$

Here $(\xi_i)_{i=1}^n$ are i.i.d. copies of ξ which are independently distributed from the data $(W_i)_{i=1}^n$ and whose distribution P_ξ does not depend on P . We also impose that

$$\mathbb{E}[\xi] = 0, \quad \mathbb{E}[\xi^2] = 1, \quad \mathbb{E}[\exp(|\xi|)] < \infty. \quad (3.20)$$

Examples of ξ include (a) $\xi = \mathcal{E} - 1$, where \mathcal{E} is a standard exponential random variable, (b) $\xi = \mathcal{N}$, where \mathcal{N} is a standard normal random variable, and (c) $\xi = \mathcal{N}_1/\sqrt{2} + (\mathcal{N}_2^2 - 1)/2$, where \mathcal{N}_1 and \mathcal{N}_2 are mutually independent standard normal random variables.¹² The choices of (a), (b), and (c) correspond respectively to the Bayesian bootstrap (e.g., Hahn (1997) and Chamberlain and Imbens (2003)), the Gaussian multiplier method (e.g. Giné and Zinn (1984) and van der Vaart and Wellner (1996, Chap. 3.6)), and the wild bootstrap method Mammen (1993).¹³ $\widehat{\psi}_u^\rho$ in (3.19) is an estimator of the influence function ψ_u^ρ defined via the plug-in rule:

$$\widehat{\psi}_u^\rho = (\widehat{\psi}_V^\rho)_{V \in \mathcal{V}_u}, \quad \widehat{\psi}_V^\rho(W) := \{\psi_{V,0,\widehat{g}_V,\widehat{m}_Z}^\alpha(W, \widehat{\alpha}_V(0)), \psi_{V,1,\widehat{g}_V,\widehat{m}_Z}^\alpha(W, \widehat{\alpha}_V(1)), \psi_V^\gamma(W, \widehat{\gamma}_V)\}. \quad (3.21)$$

¹²We do not consider the nonparametric bootstrap, which corresponds to using multinomial multipliers ξ , to reduce the length of the paper; but we note that the conditions and analysis could be extended to cover this case.

¹³The motivation for method (c) is that it is able to match 3 moments since $\mathbb{E}[\xi^2] = \mathbb{E}[\xi^3] = 1$. Methods (a) and (b) do not satisfy this property since $\mathbb{E}[\xi^2] = 1$ but $\mathbb{E}[\xi^3] \neq 1$ for these approaches.

Note that this bootstrap is computationally efficient since it does not involve recomputing the influence functions $\widehat{\psi}_u^\rho$.¹⁴ Each new draw of $(\xi_i)_{i=1}^n$ generates a new draw of $\widehat{\rho}^*$ holding the data and the estimates of the influence functions fixed. This method simply amounts to resampling the first-order approximations to the estimators. Here we build upon prior uses of this or similar methods in low-dimensional settings such as Hansen (1996) and Kline and Santos (2012).

We establish that the bootstrap law of $\sqrt{n}(\widehat{\rho}^* - \widehat{\rho})$ is uniformly asymptotically consistent: In the metric space $\ell^\infty(\mathcal{U})^{d_\rho}$, conditionally on the data,

$$\sqrt{n}(\widehat{\rho}^* - \widehat{\rho}) \rightsquigarrow_B Z_P, \text{ uniformly in } P \in \mathcal{P}_n,$$

where \rightsquigarrow_B denotes weak convergence of the bootstrap law in probability, as defined in Appendix B.

3.3. Third Step: Robust Estimation of the Structural Parameters. All structural parameters we consider take the form of smooth transformations of the reduced-form parameters:

$$\Delta := (\Delta_q)_{q \in \mathcal{Q}}, \text{ where } \Delta_q := \phi(\rho)(q), \quad q \in \mathcal{Q}. \quad (3.22)$$

The structural parameters may themselves carry an index $q \in \mathcal{Q}$ that can be different from u ; for example, the LQTE is indexed by a quantile index $q \in (0, 1)$. This formulation includes as special cases all the structural functions of Section 2. We estimate these quantities by the plug-in rule. We establish the asymptotic behavior of these estimators and the validity of the bootstrap as a corollary from the results outlined in Section 3.2 and the functional delta method (extended to handle uniformity in P).

For the application of the functional delta method, we require that the functional $\rho \mapsto \phi(\rho)$ be Hadamard differentiable *uniformly* in $\rho \in \mathbb{D}_\rho$, where \mathbb{D}_ρ is a set that contains the true values $\rho = \rho_P$ for all $P \in \mathcal{P}_n$, tangentially to a subset that contains the realizations of Z_P for all $P \in \mathcal{P}_n$ with derivative map $h \mapsto \phi'_\rho(h) = (\phi'_\rho(h)(q))_{q \in \mathcal{Q}}$.¹⁵ We define the estimators of the structural parameters and their bootstrap versions via the plug-in rule as

$$\widehat{\Delta} := (\widehat{\Delta}_q)_{q \in \mathcal{Q}}, \quad \widehat{\Delta}_q := \phi(\widehat{\rho})(q), \quad \text{and} \quad \widehat{\Delta}^* := (\widehat{\Delta}_q^*)_{q \in \mathcal{Q}}, \quad \widehat{\Delta}_q^* := \phi(\widehat{\rho}^*)(q). \quad (3.23)$$

We establish that these estimators are asymptotically Gaussian

$$\sqrt{n}(\widehat{\Delta} - \Delta) \rightsquigarrow \phi'_\rho(Z_P), \text{ uniformly in } P \in \mathcal{P}_n, \quad (3.24)$$

and that the bootstrap consistently estimates their large sample distribution:

$$\sqrt{n}(\widehat{\Delta}^* - \widehat{\Delta}) \rightsquigarrow_B \phi'_\rho(Z_P), \text{ uniformly in } P \in \mathcal{P}_n. \quad (3.25)$$

These results can be used to construct simultaneous confidence bands and test functional hypotheses on Δ using the methods described for example in Chernozhukov and Fernández-Val (2005) and Chernozhukov et al. (2013).

¹⁴Chernozhukov and Hansen (2006) and Hong and Scaillet (2006) proposed a related computationally efficient bootstrap scheme that resamples the influence functions.

¹⁵We give the definition of uniform Hadamard differentiability in Definition B.1 of Appendix B.

4. THEORY: ESTIMATION AND INFERENCE ON LOCAL TREATMENT EFFECTS FUNCTIONALS

Consider fixed sequences of numbers $\delta_n \searrow 0$, $\epsilon_n \searrow 0$, $\Delta_n \searrow 0$, at a speed at most polynomial in n (for example, $\delta_n \geq 1/n^c$ for some $c > 0$), $\ell_n := \log n$, and positive constants c , C , and $c' < 1/2$. These sequences and constants will not vary with P . The probability P can vary in the set \mathcal{P}_n of probability measures, termed “data-generating processes”, where \mathcal{P}_n is typically a set that is weakly increasing in n , i.e. $\mathcal{P}_n \subseteq \mathcal{P}_{n+1}$. Note that definitions of norms and other notation are collected in Appendix A.

Assumption 4.1 (Basic Assumptions). (i) Consider a random element W with values in a measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$ and law determined by a probability measure $P \in \mathcal{P}_n$. The observed data $((W_{ui})_{u \in \mathcal{U}})_{i=1}^n$ consist of n i.i.d. copies of a random element $(W_u)_{u \in \mathcal{U}} = ((Y_u)_{u \in \mathcal{U}}, D, Z, X)$, where \mathcal{U} is a Polish space equipped with its Borel sigma-field and $(Y_u, D, Z, X) \in \mathbb{R}^{3+d_X}$. Each W_u is generated via a measurable transform $t(W, u)$ of W and u , namely the map $t : \mathcal{W} \times \mathcal{U} \mapsto \mathbb{R}^{3+d_X}$ is measurable, and the map can possibly depend on P . Let

$$\mathcal{V}_u := \{V_{uj}\}_{j \in \mathcal{J}} := \{Y_u, \mathbf{1}_0(D)Y_u, \mathbf{1}_0(D), \mathbf{1}_1(D)Y_u, \mathbf{1}_1(D)\}, \quad \mathcal{V} := (\mathcal{V}_u)_{u \in \mathcal{U}},$$

where $\mathcal{J} = \{1, \dots, 5\}$. (ii) For $\mathcal{P} := \cup_{n=n_0}^{\infty} \mathcal{P}_n$, the map $u \mapsto Y_u$ obeys the uniform continuity property:

$$\limsup_{\epsilon \searrow 0} \sup_{P \in \mathcal{P}} \sup_{d_{\mathcal{U}}(u, \bar{u}) \leq \epsilon} \|Y_u - Y_{\bar{u}}\|_{P,2} = 0, \quad \sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{u \in \mathcal{U}} |Y_u|^{2+c} < \infty,$$

where the second supremum in the first expression is taken over $u, \bar{u} \in \mathcal{U}$, and \mathcal{U} is a totally bounded metric space equipped with a semi-metric $d_{\mathcal{U}}$. The uniform covering entropy of the set $\mathcal{F}_P = \{Y_u : u \in \mathcal{U}\}$, viewed as a collection of maps $(\mathcal{W}, \mathcal{A}_{\mathcal{W}}) \mapsto \mathbb{R}$, obeys

$$\sup_Q \log N(\epsilon \|F_P\|_{Q,2}, \mathcal{F}_P, \|\cdot\|_{Q,2}) \leq C \log(e/\epsilon) \vee 0$$

for all $P \in \mathcal{P}$, where $F_P(W) = \sup_{u \in \mathcal{U}} |Y_u|$, with the supremum taken over all finitely discrete probability measures Q on $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$. (iii) For each $P \in \mathcal{P}$, the conditional probability of $Z = 1$ given X is bounded away from zero or one, namely $c' \leq m_Z(1, X) \leq 1 - c'$ P -a.s., the instrument Z has a non-trivial impact on D , namely $c' \leq |P_P[D = 1|Z = 1, X] - P_P[D = 1|Z = 0, X]|$ P -a.s., and the regression function g_V is bounded, $\|g_V\|_{P, \infty} < \infty$ for all $V \in \mathcal{V}$.

Assumption 4.1 is stated to deal with the measurability issues associated with functional response data. This assumption also implies that the set of functions $(\psi_u^\rho)_{u \in \mathcal{U}}$, where

$$\psi_u^\rho := (\{\psi_{V,0}^\alpha, \psi_{V,1}^\alpha, \psi_V^\gamma\})_{V \in \mathcal{V}_u},$$

is P -Donsker uniformly in \mathcal{P} . That is, it implies

$$Z_{n,P} \rightsquigarrow Z_P \quad \text{in } \ell^\infty(\mathcal{U})^{d_\rho}, \quad \text{uniformly in } P \in \mathcal{P}, \quad (4.1)$$

$$Z_{n,P} := (\mathbb{G}_n \psi_u^\rho)_{u \in \mathcal{U}} \quad \text{and} \quad Z_P := (\mathbb{G}_P \psi_u^\rho)_{u \in \mathcal{U}}, \quad (4.2)$$

with \mathbb{G}_P denoting the P -Brownian bridge (van der Vaart and Wellner, 1996, p. 81–82) and with Z_P having bounded, uniformly continuous paths uniformly in $P \in \mathcal{P}$:

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{u \in \mathcal{U}} \|Z_P(u)\| < \infty, \quad \lim_{\varepsilon \searrow 0} \sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{d_{\mathcal{U}}(u, \tilde{u}) \leq \varepsilon} \|Z_P(u) - Z_P(\tilde{u})\| = 0. \quad (4.3)$$

We work with the sequence of constants defined prior to Assumption 4.1.

Assumption 4.2 (Approximate Sparsity). *Under each $P \in \mathcal{P}_n$ and for each $n \geq n_0$, uniformly for all $V \in \mathcal{V}$: (i) The approximations (3.5)–(3.7) hold with the link functions Λ_V and Λ_Z belonging to the set \mathcal{L} , the sparsity condition $\|\beta_V\|_0 + \|\beta_Z\|_0 \leq s$ holding, the approximation errors satisfying $\|r_V\|_{P,2} + \|r_Z\|_{P,2} \leq \delta_n n^{-1/4}$ and $\|r_V\|_{P,\infty} + \|r_Z\|_{P,\infty} \leq \epsilon_n$, and the sparsity index s and the number of terms p in the vector $f(X)$ obeying $s^2 \log^2(p \vee n) \log^2 n \leq \delta_n n$. (ii) There are estimators $\bar{\beta}_V$ and $\bar{\beta}_Z$ such that, with probability no less than $1 - \Delta_n$, the estimation errors satisfy $\|f(Z, X)'(\bar{\beta}_V - \beta_V)\|_{\mathbb{P}_{n,2}} + \|f(X)'(\bar{\beta}_Z - \beta_Z)\|_{\mathbb{P}_{n,2}} \leq \delta_n n^{-1/4}$, $K_n \|\bar{\beta}_V - \beta_V\|_1 + K_n \|\bar{\beta}_Z - \beta_Z\|_1 \leq \epsilon_n$; the estimators are sparse such that $\|\bar{\beta}_V\|_0 + \|\bar{\beta}_Z\|_0 \leq Cs$; and the empirical and population norms induced by the Gram matrix formed by $(f(X_i))_{i=1}^n$ are equivalent on sparse subsets, $\sup_{\|\delta\|_0 \leq \ell_n s} \left| \|f(X)' \delta\|_{\mathbb{P}_{n,2}} / \|f(X)' \delta\|_{P,2} - 1 \right| \leq \epsilon_n$. (iii) The following boundedness conditions hold: $\|f(X)\|_{\infty} \|V\|_{P,\infty} \leq K_n$ and $\|V\|_{P,\infty} \leq C$.*

Comment 4.1. Assumption 4.2 imposes simple intermediate-level conditions which encode both the approximate sparsity of the models as well as some reasonable behavior of the sparse estimators of m_Z and g_V . These conditions significantly extend and generalize the conditions employed in the literature on adaptive estimation using series methods. The boundedness conditions are made to simplify arguments, and they could be removed at the cost of more complicated proofs and more stringent side conditions. Sufficient conditions for the equivalence between empirical and population norms and primitive examples of functions admitting sparse approximations are given in Belloni et al. (2014a). We provide primitive conditions for Lasso estimators to satisfy the bounds above while addressing the problem of estimating continua of approximately sparse nuisance functions in Section 6. We expect that other sparsity-based estimators, such as the Dantzig selector or adaptive Lasso, could be used in the present context as well. ■

Under the stated assumptions, the empirical reduced form process $\widehat{Z}_{n,P} = \sqrt{n}(\widehat{\rho} - \rho)$ defined by (3.16) obeys the following relations. We recall definitions of convergence uniformly in $P \in \mathcal{P}_n$ in Appendix A.

Theorem 4.1 (Uniform Gaussianity of the Reduced-Form Parameter Process). *Under Assumptions 4.1 and 4.2, the reduced-form empirical process admits a linearization; namely,*

$$\widehat{Z}_{n,P} := \sqrt{n}(\widehat{\rho} - \rho) = Z_{n,P} + o_P(1) \quad \text{in } \ell^\infty(\mathcal{U})^{d_\rho}, \text{ uniformly in } P \in \mathcal{P}_n. \quad (4.4)$$

The process $\widehat{Z}_{n,P}$ is asymptotically Gaussian, namely

$$\widehat{Z}_{n,P} \rightsquigarrow Z_P \quad \text{in } \ell^\infty(\mathcal{U})^{d_\rho}, \text{ uniformly in } P \in \mathcal{P}_n, \quad (4.5)$$

where Z_P is defined in (4.2) and its paths obey the property (4.3).

Another main result of this section shows that the bootstrap law of the process

$$\widehat{Z}_{n,P}^* := \sqrt{n}(\widehat{\rho}^* - \widehat{\rho}) := \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \widehat{\psi}_u^\rho(W_i),$$

where $\widehat{\psi}_u^\rho$ is defined in (3.21), provides a valid approximation to the large sample law of $\sqrt{n}(\widehat{\rho} - \rho)$.

Theorem 4.2 (Validity of Multiplier Bootstrap for Inference on Reduced-Form Parameters). *Under Assumptions 4.1 and 4.2, the bootstrap law consistently approximates the large sample law Z_P of $Z_{n,P}$ uniformly in $P \in \mathcal{P}_n$, namely,*

$$\widehat{Z}_{n,P}^* \rightsquigarrow_B Z_P \quad \text{in } \ell^\infty(\mathcal{U})^{d_\rho}, \text{ uniformly in } P \in \mathcal{P}_n. \quad (4.6)$$

Next we consider inference on the structural functionals Δ defined in (3.22). We derive the large sample distribution of the estimator $\widehat{\Delta}$ in (3.23), and show that the multiplier bootstrap law of $\widehat{\Delta}^*$ in (3.23) provides a consistent approximation to that distribution. We rely on the functional delta method in our derivations, which we modify to handle uniformity with respect to the underlying dgp P . Our argument relies on the following assumption on the structural functionals.

Assumption 4.3 (Uniform Hadamard Differentiability of Structural Functionals). *Suppose that for each $P \in \mathcal{P}$, $\rho = \rho_P \in \mathbb{D}_\rho$, a compact metric space. Suppose $\varrho \mapsto \phi(\varrho)$, a functional of interest mapping $\mathbb{D}_\phi \subset \mathbb{D} = \ell^\infty(\mathcal{U})^{d_\rho}$ to $\ell^\infty(\mathcal{Q})$, where $\mathbb{D}_\rho \subset \mathbb{D}_\phi$, is Hadamard differentiable in ϱ tangentially to $\mathbb{D}_0 = UC(\mathcal{U})^{d_\rho}$ uniformly in $\varrho \in \mathbb{D}_\rho$, with the linear derivative map $\phi'_\varrho : \mathbb{D}_0 \mapsto \mathbb{D}$ such that the mapping $(\varrho, h) \mapsto \phi'_\varrho(h)$ from $\mathbb{D}_\rho \times \mathbb{D}_0$ to $\ell^\infty(\mathcal{Q})$ is continuous.*

The definition of uniform Hadamard differentiability is given in Definition B.1 of Appendix B. Assumption 4.3 holds for all examples of structural parameters listed in Section 2.

The following corollary gives the large sample law of $\sqrt{n}(\widehat{\Delta} - \Delta)$, the properly normalized structural estimator. It also shows that the bootstrap law of $\sqrt{n}(\widehat{\Delta}^* - \widehat{\Delta})$, computed conditionally on the data, approaches the large sample law $\sqrt{n}(\widehat{\Delta} - \Delta)$. It follows from the previous theorems as well as from a more general result contained in Theorem 5.3.

Corollary 4.1 (Limit Theory and Validity of Multiplier Bootstrap for Smooth Structural Functionals). *Under Assumptions 4.1, 4.2, and 4.3,*

$$\sqrt{n}(\widehat{\Delta} - \Delta) \rightsquigarrow T_P := \phi'_{\rho_P}(Z_P), \quad \text{in } \ell^\infty(\mathcal{Q}), \text{ uniformly in } P \in \mathcal{P}_n, \quad (4.7)$$

where T_P is a zero mean tight Gaussian process, for each $P \in \mathcal{P}$. Moreover,

$$\sqrt{n}(\widehat{\Delta}^* - \widehat{\Delta}) \rightsquigarrow_B T_P, \quad \text{in } \ell^\infty(\mathcal{Q}), \text{ uniformly in } P \in \mathcal{P}_n. \quad (4.8)$$

5. THEORY: A GENERAL PROBLEM OF INFERENCE ON FUNCTION-VALUED PARAMETERS WITH NUISANCE FUNCTIONS ESTIMATED BY MACHINE LEARNING METHODS

In this section, we consider a general setting where possibly a continuum of target parameters is of interest and we use modern machine learning methods, with Lasso-type or Post-Lasso-type

methods being the chief example, to estimate a continuum of high-dimensional nuisance functions. This setting covers a rich variety of modern moment-condition problems in econometrics including the treatment effects problem. We establish a functional central limit theorem for the estimators of the continuum of target parameters that holds uniformly in $P \in \mathcal{P}$, where \mathcal{P} includes a wide range of data-generating processes with well-approximable continuums of nuisance functions. We also derive a functional central limit theorem for the multiplier bootstrap that resamples the first order approximations to the standardized estimators of the continuum of target parameters and establish its uniform validity. Moreover, we establish the uniform validity of the functional delta method and the functional delta method for the multiplier bootstrap for smooth functionals of the continuum of target parameters using an appropriate strengthening of Hadamard differentiability.

We are interested in function-valued target parameters indexed by $u \in \mathcal{U} \subset \mathbb{R}^{d_u}$. We denote the true value of the target parameter by

$$\theta^0 = (\theta_u)_{u \in \mathcal{U}}, \text{ where } \theta_u \in \Theta_u \subset \Theta \subset \mathbb{R}^{d_\theta}, \text{ for each } u \in \mathcal{U}.$$

We assume that for each $u \in \mathcal{U}$, the true value θ_u is identified as the solution to the following moment condition:

$$\mathbb{E}_P[\psi_u(W_u, \theta_u, h_u(Z_u))] = 0, \quad (5.1)$$

where W_u is a random vector that takes values in a Borel set $\mathcal{W}_u \subset \mathbb{R}^{d_w}$ and contains as a subcomponent the vector Z_u taking values in a Borel set \mathcal{Z}_u , the moment function

$$\psi_u : \mathcal{W}_u \times \Theta_u \times T_u \mapsto \mathbb{R}^{d_\theta}, \quad (w, \theta, t) \mapsto \psi_u(w, \theta, t) = (\psi_{uj}(w, \theta, t))_{j=1}^{d_\theta} \quad (5.2)$$

is a Borel measurable map, and the function

$$h_u : \mathcal{Z}_u \mapsto \mathbb{R}^{d_t}, \quad z \mapsto h_u(z) = (h_{um}(z))_{m=1}^{d_t} \in T_u(z), \quad (5.3)$$

is another Borel measurable map that denotes the possibly infinite-dimensional nuisance parameter. The sets $T_u(z)$ are assumed to be convex for each $u \in \mathcal{U}$ and $z \in \mathcal{Z}_u$. Finite-dimensional nuisance parameters that do not depend on Z_u are treated as part of h_u as well.

We assume that the continuum of nuisance functions $(h_u)_{u \in \mathcal{U}}$ is well-approximable and can be well estimated by the modern generation of statistical and machine learning methods. In particular, our regularity conditions allow for approximately sparse nuisance functions, which can be modeled and estimated using modern regularization and post-selection methods such as Lasso and Post-Lasso. We let $\hat{h}_u = (\hat{h}_{um})_{m=1}^{d_t}$ denote the estimator of h_u , which we assume obeys the conditions in Assumption 5.3. The estimator $\hat{\theta}_u$ of θ_u is constructed as any approximate ϵ_n -solution in Θ_u to a sample analog of the moment condition (5.1), i.e.,

$$\|\mathbb{E}_n[\psi_u(W_u, \hat{\theta}_u, \hat{h}_u(Z_u))]\| \leq \inf_{\theta \in \Theta_u} \|\mathbb{E}_n[\psi(W_u, \theta, \hat{h}_u(Z_u))]\| + \epsilon_n, \text{ where } \epsilon_n = o(n^{-1/2}). \quad (5.4)$$

Comment 5.1 (Handling Over-identified Cases). We do not analyze over-identified cases explicitly, but it is helpful to note that they can be handled within the current framework. Let $\psi_u^o(W_u, \theta, h_u^o(Z_u))$ be the original over-identifying moment function. Let A_u denote the point-wise optimal matrix of linear combinations of the moments, so that the final moment function $\psi_u(W_u, \theta, h(Z_u)) = A_u \psi_u^o(W_u, \theta, h_u^o(Z_u))$ has the same dimension as θ_u . Here $h_u(Z_u) = (\text{vec}(A_u))'$,

$h'_u(Z_u)$); that is, we simply treat A_u as part of the nuisance function h_u being estimated. We do not analyze the preliminary estimation of A_u in the present paper in order to maintain the focus on exactly identified cases as in Section 4. Note that A_u depends on Z_u only trivially — it is a constant function of Z_u . ■

The key condition needed for regular estimation of θ_u is an orthogonality or immunization condition. The simplest to explain, yet strongest, form of this condition can be expressed as follows:

$$\partial_t \mathbb{E}_P[\psi_u(W_u, \theta_u, t) | Z_u] \Big|_{t=h_u(Z_u)} = 0, \text{ a.s.}, \quad (5.5)$$

subject to additional technical conditions such as continuity (5.6) and dominance (5.7) stated below, where we use the symbol ∂_t to abbreviate $\frac{\partial}{\partial t}$. This condition holds in the previous setting of inference on relevant treatment effects after interchanging the order of the derivative and expectation. The formulation here also covers certain non-smooth cases such as structural and instrumental quantile regression problems.

In the formal development, we use a more general form of the orthogonality condition.

Definition 5.1 (Orthogonality for Moment Condition Models, General Form). *For each $u \in \mathcal{U}$, suppose that (5.1)–(5.3) hold. Consider \mathcal{H}_u , a set of measurable functions $z \mapsto h(z) \in T_u(z)$ from \mathcal{Z}_u to \mathbb{R}^{d_t} such that $\|h(Z_u) - h_u(Z_u)\|_{P,2} < \infty$ for all $h \in \mathcal{H}_u$. Suppose also that the set $T_u(z)$ is a convex subset of \mathbb{R}^{d_t} for each $z \in \mathcal{Z}_u$. We say that ψ_u obeys a general form of orthogonality with respect to \mathcal{H}_u uniformly in $u \in \mathcal{U}$ if the following conditions hold: For each $u \in \mathcal{U}$, the derivative*

$$t \mapsto \partial_t \mathbb{E}_P[\psi_u(W_u, \theta_u, t) | Z_u] \text{ is continuous on } t \in T_u(Z_u) \text{ P-a.s.}; \quad (5.6)$$

is dominated,

$$\left\| \sup_{t \in T_u(Z_u)} \left\| \partial_t \mathbb{E}_P[\psi_u(W_u, \theta_u, t) | Z_u] \right\| \right\|_{P,2} < \infty; \quad (5.7)$$

and obeys the orthogonality condition:

$$\mathbb{E}_P \left[\partial_t \mathbb{E}_P[\psi_u(W_u, \theta_u, h_u(Z_u)) | Z_u] (h(Z_u) - h_u(Z_u)) \right] = 0 \quad \text{for all } h \in \mathcal{H}_u. \quad (5.8)$$

The Supplementary Appendix provides a slightly more general definition. The orthogonality condition (5.8) reduces to (5.5) when \mathcal{H}_u can span all measurable functions $h : \mathcal{Z}_u \mapsto T_u$ such that $\|h\|_{P,2} < \infty$ but is more general otherwise.

Comment 5.2 (Orthogonalization typically expands the number of nuisance parameters). It is important to use a moment function ψ_u that satisfies the orthogonality property given in (5.8). Generally, if we have a moment function $\tilde{\psi}_u$ which identifies θ_u but does not have this property, we can construct a moment function ψ_u that identifies θ_u and has the required orthogonality property by projecting the original function $\tilde{\psi}_u$ onto the orthocomplement of the tangent space for the original set of nuisance functions h_u^0 ; see, for example, van der Vaart and Wellner (1996), van der Vaart (1998, Chap. 25), Kosorok (2008), Belloni et al. (2013b), and Belloni et al. (2014a).

Note that the projection typically depends on P , which gives rise to additional nuisance parameters h_u^n , which are then incorporated together with the original nuisance parameters into the new parameter $h_u = (h_u^0, h_u^n)$. Note that this is a feature of all of the examples we consider. For example, the orthogonal moment functions in the exogenous case of the treatment effects framework depend on both the regression function and the propensity score function. This point is clarified further by considering the classical linear model as demonstrated in the next remark. ■

Comment 5.3 (A High-Dimensional Linear Regression Example). To illustrate the orthogonality condition, let us consider the linear model:

$$Y = D\theta_0 + X'\beta_0 + \epsilon, \quad \mathbb{E}_P[\epsilon X] = 0, \quad D = X'\pi_0 + \nu, \quad \mathbb{E}_P[\nu X] = 0, \quad (5.9)$$

where D is the treatment and X are the controls of high dimension $p \gg n$. Call the first equation the regression equation, and the second equation the propensity score equation. The orthogonal moment condition that identifies the projection coefficient θ_0 is the Frisch-Waugh-Lovell partialling out interpretation of θ_0 :

$$\mathbb{E}_P(U - \nu\theta_0)\nu = 0, \quad (5.10)$$

where U is the population residual left after projecting out the controls X from the outcome, i.e. $Y = X'\delta_0 + U$, $\mathbb{E}_P U X = 0$; and ν is the population residual left after projecting out controls from the treatment, as defined in the propensity score equation. The high-dimensional nuisance function is $h(Z) = (X'\delta, X'\pi)'$, for $Z = X$, with true value denoted by $h_0(Z) = (X'\delta_0, X'\pi_0)'$. Now the moment function

$$\psi(W, \theta, h) = \{(Y - X'\delta) - (D - X'\pi)\theta\}(D - X'\pi), \quad (5.11)$$

has the required orthogonality property (5.8), since by the law of iterated expectations and some simple algebra:

$$\begin{aligned} & \mathbb{E}_P \left[\partial_t \mathbb{E}_P [\psi_0(W, \theta_0, h_0(Z)) | Z] (h(Z) - h_0(Z)) \right] \\ &= \left(\mathbb{E}_P [(D - X'\pi_0) X' a], \mathbb{E}_P [\{(Y - X'\delta_0) + 2(D - X'\pi_0)\theta_0\} X' b] \right) = 0, \end{aligned} \quad (5.12)$$

for $a = \delta - \delta_0$ and $b = \pi - \pi_0$. In fact $\psi(W, \theta_0, h_0)$ is the semi-parametrically efficient score for θ_0 . The resulting estimator of θ_0 is root- n consistent and asymptotically normal, uniformly within a class of approximately sparse models as follows from the general results of this section, and is also semi-parametrically efficient. See also Belloni et al. (2014a) which deals with the partially linear model in detail and thus covers this linear example as a special case.

Note that the orthogonal moment function contains two nuisance functions – the regression function and the propensity score – $X'\delta_0$ and $X'\pi_0$. We could also identify θ_0 through non-orthogonal moment conditions, containing single nuisance functions:

$$\mathbb{E}_P\{[Y - D\theta_0 - X'\beta_0]D\} = 0 \quad \text{or} \quad \mathbb{E}_P\{[Y - D\theta_0](D - X'\pi_0)\} = 0.$$

The first moment condition corresponds to the regression method, while the second to the so-called covariate balancing method. Importantly, the use of these non-orthogonal moment conditions generally does not produce an estimator for θ_0 that is \sqrt{n} -consistent and asymptotically normal uniformly in the class of approximately sparse models. This failure occurs because we are forced to

use highly non-regular estimators to estimate the nuisance functions $X'\delta_0$ and $X'\pi_0$ in the $p \gg n$ setting. In fact, this failure would also occur with a low number of controls, including having only $p = 1$, whenever selection procedures that exclude irrelevant variables with very high probability are used to estimate the regression parameter δ_0 or the propensity score parameter π_0 . For more discussion and documentation of this failure, see Leeb and Pötscher (2008a; 2008b); Pötscher (2009); and Belloni, Chernozhukov, and Hansen (2013a; 2014a). By contrast, constructing orthogonal moment conditions – involving the projection of both the outcome and the treatment onto the controls and thereby combining the regression and covariate balancing methods – makes it possible to achieve \sqrt{n} consistency and asymptotic normality uniformly within a class of approximately sparse models. ■

In what follows, we shall denote by δ , c_0 , c , and C some positive constants. For a positive integer d , $[d]$ denotes the set $\{1, \dots, d\}$. We shall impose the following regularity conditions.

Assumption 5.1 (Moment condition problem). *Consider a random element W , taking values in a measure space $(\mathcal{W}, \mathcal{A}_W)$, with law determined by a probability measure $P \in \mathcal{P}_n$. The observed data $((W_{ui})_{u \in \mathcal{U}})_{i=1}^n$ consist of n i.i.d. copies of a random element $(W_u)_{u \in \mathcal{U}}$ which is generated as a suitably measurable transformation with respect to W and u . Uniformly for all $n \geq n_0$ and $P \in \mathcal{P}_n$, the following conditions hold: (i) The true parameter value θ_u obeys (5.1) and is interior relative to $\Theta_u \subset \Theta \subset \mathbb{R}^{d_\theta}$, namely there is a ball of radius δ centered at θ_u contained in Θ_u for all $u \in \mathcal{U}$, and Θ is compact. (ii) For $\nu := (\nu_k)_{k=1}^{d_\theta+d_t} = (\theta, t)$, each $j \in [d_\theta]$ and $u \in \mathcal{U}$, the map $\Theta_u \times T_u(Z_u) \ni \nu \mapsto \mathbb{E}_P[\psi_{uj}(W_u, \nu) | Z_u]$ is twice continuously differentiable a.s. with derivatives obeying the integrability conditions specified in Assumption 5.2. (iii) For all $u \in \mathcal{U}$, the moment function ψ_u obeys the orthogonality condition given in Definition 5.1 for the set $\mathcal{H}_u = \mathcal{H}_{un}$ specified in Assumption 5.3. (iv) The following identifiability condition holds: $\|\mathbb{E}_P[\psi_u(W_u, \theta, h_u(Z_u))]\| \geq 2^{-1}(\|J_u(\theta - \theta_u)\| \wedge c_0)$ for all $\theta \in \Theta_u$, where the singular values of $J_u := \partial_\theta \mathbb{E}[\psi_u(W_u, \theta_u, h_u(Z_u))]$ lie between c and C for all $u \in \mathcal{U}$.*

The conditions of Assumption 5.1 are mild and standard in moment condition problems. Assumption 5.1(iv) encodes sufficient global and local identifiability to obtain a rate result. The suitably measurable condition, defined in Appendix A, is a mild condition satisfied in most practical cases.

Assumption 5.2 (Entropy and smoothness). *The set $(\mathcal{U}, d_{\mathcal{U}})$ is a semi-metric space such that $\log N(\epsilon, \mathcal{U}, d_{\mathcal{U}}) \leq C \log(e/\epsilon) \vee 0$. Let $\alpha \in [1, 2]$, and let α_1 and α_2 be some positive constants. Uniformly for all $n \geq n_0$ and $P \in \mathcal{P}_n$, the following conditions hold: (i) The set of functions $\mathcal{F}_0 = \{\psi_{uj}(W_u, \theta_u, h_u(Z_u)) : j \in [d_\theta], u \in \mathcal{U}\}$, viewed as functions of W is suitably measurable; has an envelope function $F_0(W) = \sup_{j \in [d_\theta], u \in \mathcal{U}, \nu \in \Theta_u \times T_u(Z_u)} |\psi_{uj}(W_u, \nu)|$ that is measurable with respect to W and obeys $\|F_0\|_{P, q} \leq C$, where $q \geq 4$ is a fixed constant; and has a uniform covering entropy obeying $\sup_Q \log N(\epsilon \|F_0\|_{Q, 2}, \mathcal{F}_0, \|\cdot\|_{Q, 2}) \leq C \log(e/\epsilon) \vee 0$. (ii) For all $j \in [d_\theta]$ and $k, r \in [d_\theta + d_t]$, and $\psi_{uj}(W) := \psi_{uj}(W_u, \theta_u, h_u(Z_u))$,*

- (a) $\sup_{u \in \mathcal{U}, (\nu, \bar{\nu}) \in (\Theta_u \times T_u(Z_u))^2} \mathbb{E}_P[(\psi_{uj}(W_u, \nu) - \psi_{uj}(W_u, \bar{\nu}))^2 | Z_u] / \|\nu - \bar{\nu}\|^\alpha \leq C$, P -a.s.,
- (b) $\sup_{d_{\mathcal{U}}(u, \bar{u}) \leq \delta} \mathbb{E}_P[(\psi_{uj}(W) - \psi_{\bar{u}j}(W))^2] \leq C\delta^{\alpha_1}$, $\sup_{d_{\mathcal{U}}(u, \bar{u}) \leq \delta} \|J_u - J_{\bar{u}}\| \leq C\delta^{\alpha_2}$,

- (c) $\mathbb{E}_P \sup_{u \in \mathcal{U}, \nu \in \Theta_u \times T_u(Z_u)} |\partial_{\nu_r} \mathbb{E}_P [\psi_{uj}(W_u, \nu) | Z_u]|^2 \leq C,$
(d) $\sup_{u \in \mathcal{U}, \nu \in \Theta_u \times T_u(Z_u)} |\partial_{\nu_k} \partial_{\nu_r} \mathbb{E}_P [\psi_{uj}(W_u, \nu) | Z_u]| \leq C, P\text{-a.s.}$

Assumption 5.2 imposes smoothness and integrability conditions on various quantities derived from ψ_u . It also imposes conditions on the complexity of the relevant function classes.

In what follows, let $\Delta_n \searrow 0$, $\delta_n \searrow 0$, and $\tau_n \searrow 0$ be sequences of constants approaching zero from above at a speed at most polynomial in n (for example, $\delta_n \geq 1/n^c$ for some $c > 0$).

Assumption 5.3 (Estimation of nuisance functions). *The following conditions hold for each $n \geq n_0$ and all $P \in \mathcal{P}_n$. The estimated functions $\hat{h}_u = (\hat{h}_{um})_{m=1}^{d_t} \in \mathcal{H}_{un}$ with probability at least $1 - \Delta_n$, where \mathcal{H}_{un} is the set of measurable maps $Z_u \ni z \mapsto h = (h_m)_{m=1}^{d_t}(z) \in T_u(z)$ such that*

$$\|h_m - h_{um}\|_{P,2} \leq \tau_n, \quad \tau_n^2 \sqrt{n} \leq \delta_n,$$

and whose complexity does not grow too quickly in the sense that $\mathcal{F}_1 = \{\psi_{uj}(W_u, \theta, h(Z_u)) : j \in [d_\theta], u \in \mathcal{U}, \theta \in \Theta_u, h \in \mathcal{H}_{un}\}$ is suitably measurable and its uniform covering entropy obeys

$$\sup_Q \log N(\epsilon \|F_1\|_{Q,2}, \mathcal{F}_1, \|\cdot\|_{Q,2}) \leq s_n(\log(a_n/\epsilon)) \vee 0,$$

where $F_1(W)$ is an envelope for \mathcal{F}_1 which is measurable with respect to W and satisfies $F_1(W) \leq F_0(W)$ for F_0 defined in Assumption 5.2. The complexity characteristics $a_n \geq \max(n, e)$ and $s_n \geq 1$ obey the growth conditions:

$$n^{-1/2} \left(\sqrt{s_n \log(a_n)} + n^{-1/2} s_n n^{\frac{1}{q}} \log(a_n) \right) \leq \tau_n \text{ and } \tau_n^{\alpha/2} \sqrt{s_n \log(a_n)} + s_n n^{\frac{1}{q} - \frac{1}{2}} \log(a_n) \log n \leq \delta_n,$$

where q and α are defined in Assumption 5.2.

Assumption 5.3 imposes conditions on the estimation rate of the nuisance functions h_{um} and on the complexity of the functions sets that contain the estimators \hat{h}_{um} . This allows for a wide variety of modern modeling assumptions and regularization methods for function fitting, both the traditional methods and the new statistical and machine learning methods. Within the approximately sparse framework, the index s_n corresponds to the maximum of the dimension of the approximating models and of the size of the selected models; and $a_n = p \vee n$. Under other frameworks, these parameters could be different; yet if they are well-behaved, then our results still apply. Thus, these results potentially cover other frameworks, where structured assumptions other than approximate sparsity are used to make the estimation and modeling problem manageable. It is important to point out that the class \mathcal{F}_1 generally won't be Donsker because its entropy is allowed to increase with n . Allowing for non-Donsker classes is crucial for accommodating modern high-dimensional estimation methods for the nuisance functions as we have seen in the previous section. This feature makes the conditions imposed here very different from the conditions imposed in various classical references on dealing with nonparametrically estimated nuisance functions; see, for example, van der Vaart and Wellner (1996), van der Vaart (1998), Kosorok (2008), and other references listed in the introduction.

The following theorem is one of the main results of the paper:

Theorem 5.1 (Uniform Functional Central Limit Theorem for a Continuum of Target Parameters). *Under Assumptions 5.1, 5.2, and 5.3, for an estimator $(\widehat{\theta}_u)_{u \in \mathcal{U}}$ that obeys equation (5.4), $\sqrt{n}(\widehat{\theta}_u - \theta_u)_{u \in \mathcal{U}} = (\mathbb{G}_n \bar{\psi}_u)_{u \in \mathcal{U}} + o_P(1)$ in $\ell^\infty(\mathcal{U})^{d_\theta}$, uniformly in $P \in \mathcal{P}_n$, where $\bar{\psi}_u(W) := -J_u^{-1} \psi_u(W_u, \theta_u, h_u(Z_u))$, and*

$$(\mathbb{G}_n \bar{\psi}_u)_{u \in \mathcal{U}} \rightsquigarrow (\mathbb{G}_P \bar{\psi}_u)_{u \in \mathcal{U}} \text{ in } \ell^\infty(\mathcal{U})^{d_\theta}, \text{ uniformly in } P \in \mathcal{P}_n,$$

where the paths of $u \mapsto \mathbb{G}_P \bar{\psi}_u$ are a.s. uniformly continuous on $(\mathcal{U}, d_{\mathcal{U}})$ and

$$\sup_{P \in \mathcal{P}_n} \mathbb{E}_P \sup_{u \in \mathcal{U}} \|\mathbb{G}_P \bar{\psi}_u\| < \infty \text{ and } \lim_{\delta \rightarrow 0} \sup_{P \in \mathcal{P}_n} \sup_{d_{\mathcal{U}}(u, \bar{u}) \leq \delta} \|\mathbb{G}_P \bar{\psi}_u - \mathbb{G}_P \bar{\psi}_{\bar{u}}\| = 0.$$

Comment 5.4. It is important to mention here that this result on a continuum of parameters solving a continuum of moment conditions is completely new. The prior approaches dealing with continua of moment conditions with infinite-dimensional nuisance parameters, for example, the ones given in Chernozhukov and Hansen (2006) and Escanciano and Zhu (2013), impose Donsker conditions on the class of functions, following Andrews (1994a), that contain the values of the estimators of these nuisance functions. This approach is precluded in our setting because the resulting class of functions in our case has entropy that grows with the sample size and therefore is not Donsker. Hence, we develop a new approach to establishing the results which exploits the interplay between the rate of growth of entropy, the biases, and the size of the estimation error. In addition, the new approach allows for obtaining results that are uniform in P . ■

We can estimate the law of Z_P with the bootstrap law of

$$\widehat{Z}_{n,P}^* := \sqrt{n}(\widehat{\theta}_u^* - \widehat{\theta}_u)_{u \in \mathcal{U}} := \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \widehat{\psi}_u(W_i) \right)_{u \in \mathcal{U}}, \quad (5.13)$$

where $(\xi_i)_{i=1}^n$ are i.i.d. multipliers as defined in equation (3.20), $\widehat{\psi}_u(W_i)$ is the estimated score

$$\widehat{\psi}_u(W_i) := -\widehat{J}_u^{-1} \psi_u(W_{ui}, \widehat{\theta}_u, \widehat{h}_u(Z_{ui})),$$

and \widehat{J}_u is a suitable estimator of J_u .¹⁶ The bootstrap law is computed by drawing $(\xi_i)_{i=1}^n$ conditional on the data.

The following theorem shows that the multiplier bootstrap provides a valid approximation to the large sample law of $\sqrt{n}(\widehat{\theta}_u - \theta_u)_{u \in \mathcal{U}}$.

Theorem 5.2 (Uniform Validity of Multiplier Bootstrap). *Suppose Assumptions 5.1, 5.2, and 5.3 hold, the estimator $(\widehat{\theta}_u)_{u \in \mathcal{U}}$ obeys equation (5.4), and that, for the constant α defined in Assumption 5.2 and some positive constant α_3 , uniformly in $P \in \mathcal{P}_n$ with probability $1 - \delta_n$,*

$$(u \mapsto \widehat{J}_u) \in \mathcal{J}_n = \{u \mapsto \bar{J}_u : \|\bar{J}_u - \bar{J}_{\bar{u}}\| \leq C \|u - \bar{u}\|^{\alpha_3}, \|\bar{J}_u - J_u\| \leq \tau_n^{\alpha/2}, \text{ for all } (u, \bar{u}) \in \mathcal{U}^2\}.$$

Then, $\widehat{Z}_{n,P}^* \rightsquigarrow_B Z_P$ in $\ell^\infty(\mathcal{U})^{d_\theta}$, uniformly in $P \in \mathcal{P}_n$.

¹⁶We do not discuss the estimation of J_u since it is often a problem-specific matter. In Section 3, J_u was equal to minus the identity matrix, so we did not need to estimate it.

We next derive the large sample distribution and validity of the multiplier bootstrap for the estimator $\widehat{\Delta} := \phi(\widehat{\theta}) := \phi((\widehat{\theta}_u)_{u \in \mathcal{U}})$ of the functional $\Delta := \phi(\theta^0) = \phi((\theta_u)_{u \in \mathcal{U}})$ using the functional delta method. The functional $\theta^0 \mapsto \phi(\theta^0)$ is defined as a uniformly Hadamard differentiable transform of $\theta^0 = (\theta_u)_{u \in \mathcal{U}}$. The following result gives the large sample law of $\sqrt{n}(\widehat{\Delta} - \Delta)$, the properly normalized estimator. It also shows that the bootstrap law of $\sqrt{n}(\widehat{\Delta}^* - \widehat{\Delta})$, computed conditionally on the data, is consistent for the large sample law of $\sqrt{n}(\widehat{\Delta} - \Delta)$. Here $\widehat{\Delta}^* := \phi(\widehat{\theta}^*) = \phi((\widehat{\theta}_u^*)_{u \in \mathcal{U}})$ is the bootstrap version of $\widehat{\Delta}$, and $\widehat{\theta}_u^* = \widehat{\theta}_u + n^{-1} \sum_{i=1}^n \xi_i \widehat{\psi}_u(W_i)$ is the multiplier bootstrap version of $\widehat{\theta}_u$ defined via equation (5.13).

Theorem 5.3 (Uniform Limit Theory and Validity of Multiplier Bootstrap for Smooth Functionals of θ). *Suppose that for each $P \in \mathcal{P} := \cup_{n \geq n_0} \mathcal{P}_n$, $\theta^0 = \theta_P^0$ is an element of a compact set \mathbb{D}_θ . Suppose $\theta \mapsto \phi(\theta)$, a functional of interest mapping $\mathbb{D}_\phi \subset \mathbb{D} = \ell^\infty(\mathcal{U})^{d_\theta}$ to $\ell^\infty(\mathcal{Q})$, where $\mathbb{D}_\theta \subset \mathbb{D}_\phi$, is Hadamard differentiable in θ tangentially to $\mathbb{D}_0 = UC(\mathcal{U})^{d_\theta}$ uniformly in $\theta \in \mathbb{D}_\theta$, with the linear derivative map $\phi'_\theta : \mathbb{D}_0 \mapsto \mathbb{D}$ such that the mapping $(\theta, h) \mapsto \phi'_\theta(h)$ from $\mathbb{D}_\theta \times \mathbb{D}_0$ to $\ell^\infty(\mathcal{Q})$ is continuous. Then,*

$$\sqrt{n}(\widehat{\Delta} - \Delta) \rightsquigarrow T_P := \phi'_{\theta_P^0}(Z_P) \quad \text{in } \ell^\infty(\mathcal{Q}), \text{ uniformly in } P \in \mathcal{P}_n, \quad (5.14)$$

where T_P is a zero mean tight Gaussian process, for each $P \in \mathcal{P}$. Moreover,

$$\sqrt{n}(\widehat{\Delta}^* - \widehat{\Delta}) \rightsquigarrow_B T_P \quad \text{in } \ell^\infty(\mathcal{Q}), \text{ uniformly in } P \in \mathcal{P}_n. \quad (5.15)$$

To derive Theorem 5.3, we strengthen the usual notion of Hadamard differentiability to a uniform notion introduced in Definition B.1. Theorems B.3 and B.4 show that this uniform Hadamard differentiability is sufficient to guarantee the validity of the functional delta uniformly in P . These new uniform functional delta method theorems may be of independent interest.

6. THEORY: LASSO AND POST-LASSO FOR FUNCTIONAL RESPONSE DATA

In this section, we provide results for Lasso and Post-Lasso estimators with function-valued outcomes and linear or logistic links. As these results are of interest beyond the context of estimation of nuisance functions for moment condition problems or treatment effects estimation, we present this section in a way that leaves it autonomous with respect to the rest of the paper.

6.1. The generic setting with function-valued outcomes. Consider a data generating process with a functional response variable $(Y_u)_{u \in \mathcal{U}}$ and observable covariates X satisfying for each $u \in \mathcal{U}$,

$$E_P[Y_u | X] = \Lambda(f(X)' \theta_u) + r_u(X), \quad (6.1)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}^p$ is a set of p measurable transformations of the initial controls X , θ_u is a p -dimensional vector, r_u is an approximation error, and Λ is a fixed known link function. The notation in this section differs from the rest of the paper with Y_u and X denoting a generic response and a generic vector of covariates to facilitate the application of these results to other contexts. We only consider the linear link function, $\Lambda(t) = t$, and the logistic link function, $\Lambda(t) = \exp(t)/\{1 + \exp(t)\}$, in detail.

Considering the logistic link is useful when the functional response is binary, though the linear link can be used in that case as well under some conditions. For example, it is useful for estimating a high-dimensional generalization of the distributional regression models considered in Chernozhukov et al. (2013) where the response variable is the continuum $(Y_u = 1(Y \leq u))_{u \in \mathcal{U}}$. Even though we focus on these two cases we note that the principles discussed here apply to many other M -estimators with convex (or approximately convex) criterion functions. In the remainder of the section, we discuss and establish results for ℓ_1 -penalized and post-model selection estimators of $(\theta_u)_{u \in \mathcal{U}}$ that hold uniformly over $u \in \mathcal{U}$.

Throughout the section, we assume that $u \in \mathcal{U} \subset [0, 1]^{d_u}$ and that we have n i.i.d. observations from dgps where (6.1) holds, $\{(Y_{ui})_{u \in \mathcal{U}}, X_i\}_{i=1}^n$, available for estimating $(\theta_u)_{u \in \mathcal{U}}$. For each $u \in \mathcal{U}$, penalty level λ , and diagonal matrix of penalty loadings $\widehat{\Psi}_u$, we define the Lasso estimator as

$$\widehat{\theta}_u \in \arg \min_{\theta \in \mathbb{R}^p} \mathbb{E}_n[M(Y_u, f(X)' \theta)] + \frac{\lambda}{n} \|\widehat{\Psi}_u \theta\|_1 \quad (6.2)$$

where $M(y, t) = \frac{1}{2}(y - \Lambda(t))^2$ in the case of linear regression, and $M(y, t) = -\{1(y = 1) \log \Lambda(t) + 1(y = 0) \log(1 - \Lambda(t))\}$ in the case of the logistic link function for binary response data. For each $u \in \mathcal{U}$, the Post-Lasso estimator based on a set of covariates \widetilde{T}_u is then defined as

$$\widetilde{\theta}_u \in \arg \min_{\theta \in \mathbb{R}^p} \mathbb{E}_n[M(Y_u, f(X)' \theta)] \quad : \quad \text{supp}(\theta) \subseteq \widetilde{T}_u \quad (6.3)$$

where the set \widetilde{T}_u contains $\text{supp}(\widehat{\theta}_u)$ and possibly additional variables deemed as important.¹⁷ We will set $\widetilde{T}_u = \text{supp}(\widehat{\theta}_u)$ unless otherwise noted.

The chief departure between the analysis when \mathcal{U} is a singleton and the functional response case is that the penalty level needs to be set to control selection errors uniformly over $u \in \mathcal{U}$. To do so, we will set λ so that with high probability

$$\frac{\lambda}{n} \geq c \sup_{u \in \mathcal{U}} \left\| \widehat{\Psi}_u^{-1} \mathbb{E}_n [\partial_\theta M(Y_u, f(X)' \theta_u)] \right\|_\infty, \quad (6.4)$$

where $c > 1$ is a fixed constant. When \mathcal{U} is a singleton the strategy above is similar to Bickel et al. (2009), Belloni and Chernozhukov (2013), and Belloni et al. (2011), who use an analog of (6.4) to derive the properties of Lasso and Post-Lasso. When \mathcal{U} is not a singleton, this strategy was first employed in the context of ℓ_1 -penalized quantile regression processes by Belloni and Chernozhukov (2011).

To implement (6.4), we propose setting the penalty level as

$$\lambda = c\sqrt{n}\Phi^{-1}(1 - \gamma/\{2pn^{d_u}\}), \quad (6.5)$$

¹⁷The total number of additional variables \widehat{s}_a should also obey the same growth conditions that s obeys. For example, if the additional variables are chosen so that $\widehat{s}_a \lesssim \|\widehat{\theta}_u\|_0$ the growth condition is satisfied with probability going to one for the designs covered by Assumptions 6.1 and 6.2. See also Belloni et al. (2014a) for a discussion on choosing additional variables.

where d_u is the dimension of \mathcal{U} , $1 - \gamma$ with $\gamma = o(1)$ is a confidence level associated with the probability of event (6.4), and $c > 1$ is a slack constant.¹⁸ When implementing the estimators, we set $c = 1.1$ and $\gamma = .1/\log(n)$, which is theoretically motivated and practically tested in an extensive set of simulation experiments in Belloni et al. (2014a). In addition to the penalty parameter λ , we also need to construct a penalty loading matrix $\widehat{\Psi}_u = \text{diag}(\{\widehat{l}_{uj}, j = 1, \dots, p\})$. This loading matrix can be formed according to the following iterative algorithm.

Algorithm 6.1 (Estimation of Penalty Loadings). *Choose $\gamma \in [1/n, \min\{1/\log n, pn^{d_u-1}\}]$ and $c > 1$ to form λ as defined in (6.5), and choose a constant $K \geq 1$ as an upper bound on the number of iterations. (0) Set $k = 0$, and initialize $\widehat{l}_{uj,0}$ for each $j = 1, \dots, p$. For the linear link function, set $\widehat{l}_{uj,0} = \{\mathbb{E}_n[f_j^2(X)(Y_u - \bar{Y}_u)^2]\}^{1/2}$ with $\bar{Y}_u = \mathbb{E}_n[Y_u]$. For the logistic link function, set $\widehat{l}_{uj,0} = \frac{1}{2}\{\mathbb{E}_n[f_j^2(X)]\}^{1/2}$. (1) Compute the Lasso and Post-Lasso estimators, $\widehat{\theta}_u$ and $\widetilde{\theta}_u$, based on $\widehat{\Psi}_u = \text{diag}(\{\widehat{l}_{uj,k}, j = 1, \dots, p\})$. (2) Set $\widehat{l}_{uj,k+1} := \{\mathbb{E}_n[f_j^2(X)(Y_u - \Lambda(f(X)'\widetilde{\theta}_u))^2]\}^{1/2}$. (3) If $k > K$, stop; otherwise set $k \leftarrow k + 1$ and go to step (1).*

6.2. Properties of a Continuum of Lasso and Post-Lasso: Linear Link. We provide sufficient conditions for establishing good performance of the estimators discussed above when the linear link function is used. In the statement of the following assumption, $\delta_n \searrow 0$ and $\Delta_n \searrow 0$ are fixed sequences approaching zero from above at a speed at most polynomial in n (for example, $\delta_n \geq 1/n^c$ for some $c > 0$), $\ell_n := \log n$, and c, C, κ', κ'' and $\nu \in (0, 1]$ are positive finite constants.

Assumption 6.1. *Consider a random element W taking values in a measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$, with law determined by a probability measure $P \in \mathcal{P}_n$. The observed data $((Y_u)_{u \in \mathcal{U}}, X_i)_{i=1}^n$ consist of n i.i.d. copies of random element $((Y_u)_{u \in \mathcal{U}}, X)$, which is generated as a suitably measurable transformation of W and u . The model (6.1) holds with linear link $t \mapsto \Lambda(t) = t$ for all $u \in \mathcal{U} \subset [0, 1]^{d_u}$, where d_u is fixed and \mathcal{U} is equipped with the semi-metric $d_{\mathcal{U}}$. Uniformly for all $n \geq n_0$ and $P \in \mathcal{P}_n$, the following conditions hold. (i) The model (6.1) is approximately sparse with sparsity index obeying $\sup_{u \in \mathcal{U}} \|\theta_u\|_0 \leq s$ and the growth restriction $\log(p \vee n) \leq \delta_n n^{1/3}$. (ii) The set \mathcal{U} has uniform covering entropy obeying $\log N(\epsilon, \mathcal{U}, d_{\mathcal{U}}) \leq d_u \log(1/\epsilon) \vee 0$, and the collection $(\zeta_u = Y_u - \mathbb{E}_P[Y_u | X], r_u)_{u \in \mathcal{U}}$ are suitably measurable transformations of W and u . (iii) Uniformly over $u \in \mathcal{U}$, the moments of the model are boundedly heteroscedastic, namely $c \leq \mathbb{E}_P[\zeta_u^2 | X] \leq C$ a.s., and $\max_{j \leq p} \mathbb{E}_P[|f_j(X)\zeta_u|^3 + |f_j(X)Y_u|^3] \leq C$. (iv) For a fixed $\nu > 0$ and a sequence K_n , the dictionary functions, approximation errors, and empirical errors obey the following regularity conditions: (a) $c \leq \mathbb{E}_P[f_j^2(X)] \leq C$, $j = 1, \dots, p$; $\max_{j \leq p} |f_j(X)| \leq K_n$ a.s.; $K_n^2 s \log(p \vee n) \leq \delta_n n$. (b) With probability $1 - \Delta_n$, $\sup_{u \in \mathcal{U}} \mathbb{E}_n[r_u^2(X)] \leq Cs \log(p \vee n)/n$; $\sup_{u \in \mathcal{U}} \max_{j \leq p} |(\mathbb{E}_n - \mathbb{E}_P)[f_j^2(X)\zeta_u^2]| \vee |(\mathbb{E}_n - \mathbb{E}_P)[f_j^2(X)Y_u^2]| \leq \delta_n$; $\log^{1/2}(p \vee n) \sup_{d_{\mathcal{U}}(u, u') \leq 1/n} \max_{j \leq p} \{\mathbb{E}_n[f_j(X)^2(\zeta_u - \zeta_{u'})^2]\}^{1/2} \leq \delta_n$, and $\sup_{d_{\mathcal{U}}(u, u') \leq 1/n} \|\mathbb{E}_n[f(X)(\zeta_u - \zeta_{u'})]\|_{\infty} \leq \delta_n n^{-1/2}$. (c) With probability $1 - \Delta_n$, the empirical minimum and maximum sparse eigenvalues are bounded from zero and above, namely $\kappa' \leq \inf_{\|\delta\|_0 \leq s \ell_n} \|f(X)'\delta\|_{\mathbb{P}_n, 2} \leq \sup_{\|\delta\|_0 \leq s \ell_n} \|f(X)'\delta\|_{\mathbb{P}_n, 2} \leq \kappa''$.*

¹⁸When the set \mathcal{U} is a singleton, one can use the penalty level in (6.5) with $d_u = 0$. This choice corresponds to that used in Belloni et al. (2014a).

Assumption 6.1 is only a set of sufficient conditions. The finite sample results in the Supplementary Appendix allow for more general conditions (for example, d_u can grow with the sample size). We verify that the more technical conditions in Assumption 6.1(iv)(b) hold in a variety of cases, see Lemma J.2 in Appendix J in the Supplementary Appendix. Under Assumption 6.1, we establish results on the performance of the estimators (6.2) and (6.3) for the linear link function case that hold uniformly over $u \in \mathcal{U}$ and $P \in \mathcal{P}_n$.

Theorem 6.1 (Rates and Sparsity for Functional Responses under Linear Link). *Under Assumption 6.1 and setting the penalty and loadings as in Algorithm 6.1, for all n large enough, uniformly for all $P \in \mathcal{P}_n$ with \mathbb{P}_P probability $1 - o(1)$, for some constant \bar{C} , the Lasso estimator $\hat{\theta}_u$ is uniformly sparse, $\sup_{u \in \mathcal{U}} \|\hat{\theta}_u\|_0 \leq \bar{C}s$, and the following performance bounds hold:*

$$\sup_{u \in \mathcal{U}} \|f(X)'(\hat{\theta}_u - \theta_u)\|_{\mathbb{P}_{n,2}} \leq \bar{C} \sqrt{\frac{s \log(p \vee n)}{n}} \quad \text{and} \quad \sup_{u \in \mathcal{U}} \|\hat{\theta}_u - \theta_u\|_1 \leq \bar{C} \sqrt{\frac{s^2 \log(p \vee n)}{n}}.$$

For all n large enough, uniformly for all $P \in \mathcal{P}_n$, with \mathbb{P}_P probability $1 - o(1)$, the Post-Lasso estimator corresponding to $\hat{\theta}_u$ obeys

$$\sup_{u \in \mathcal{U}} \|f(X)'(\tilde{\theta}_u - \theta_u)\|_{\mathbb{P}_{n,2}} \leq \bar{C} \sqrt{\frac{s \log(p \vee n)}{n}}, \quad \text{and} \quad \sup_{u \in \mathcal{U}} \|\tilde{\theta}_u - \theta_u\|_1 \leq \bar{C} \sqrt{\frac{s^2 \log(p \vee n)}{n}}.$$

We note that the performance bounds are exactly of the type used in Assumption 4.2 (see also Assumption H.1 in the Supplementary Appendix). Indeed, under the condition $s^2 \log^2(p \vee n) \log^2 n \leq \delta_n n$, the rate of convergence established in Theorem 6.1 yields $\sqrt{s \log(p \vee n)/n} \leq o(n^{-1/4})$.

6.3. Properties of Lasso and Post-Lasso Estimators: Logistic Link. We provide sufficient conditions to state results on the performance of the estimators discussed above for the logistic link function. Consider the fixed sequences $\delta_n \searrow 0$ and $\Delta_n \searrow 0$ approaching zero from above at a speed at most polynomial in n , $\ell_n := \log n$, and the positive finite constants c , C , κ' , κ'' , and $\underline{c} \leq 1/2$.

Assumption 6.2. *Consider a random element W taking values in a measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$, with law determined by a probability measure $P \in \mathcal{P}_n$. The observed data $((Y_{ui})_{u \in \mathcal{U}}, X_i)_{i=1}^n$ consist of n i.i.d. copies of random element $((Y_u)_{u \in \mathcal{U}}, X)$, which is generated as a suitably measurable transformation of W and u . The model (6.1) holds with $Y_{ui} \in \{0, 1\}$ with the logistic link $t \mapsto \Lambda(t) = \exp(t)/\{1 + \exp(t)\}$ for each $u \in \mathcal{U} \subset [0, 1]^{d_u}$, where d_u is fixed and \mathcal{U} is equipped with the semi-metric $d_{\mathcal{U}}$. Uniformly for all $n \geq n_0$ and $P \in \mathcal{P}_n$, the following conditions hold. (i) The model (6.1) is approximately sparse with sparsity index obeying $\sup_{u \in \mathcal{U}} \|\theta_u\|_0 \leq s$ and the growth restriction $\log(p \vee n) \leq \delta_n n^{1/3}$. (ii) The set \mathcal{U} has uniform covering entropy obeying $\log N(\epsilon, \mathcal{U}, d_{\mathcal{U}}) \leq d_u \log(1/\epsilon) \vee 0$, and the collection $(\zeta_u = Y_u - \mathbb{E}_P[Y_u | X], r_u)_{u \in \mathcal{U}}$ is a suitably measurable transformation of W and u . (iii) Uniformly over $u \in \mathcal{U}$ the moments of the model satisfy $\max_{j \leq p} \mathbb{E}_P[|f_j(X)|^3] \leq C$, and $\underline{c} \leq \mathbb{E}_P[Y_u | X] \leq 1 - \underline{c}$ a.s. (iv) For a sequence K_n , the dictionary functions, approximation errors, and empirical errors obey the following boundedness and empirical regularity conditions: (a) $\sup_{u \in \mathcal{U}} |r_u(X)| \leq \delta_n$ a.s.; $c \leq \mathbb{E}_P[f_j^2(X)] \leq C$, $j = 1, \dots, p$; $\max_{j \leq p} |f_j(X)| \leq K_n$ a.s.; and $K_n^2 s^2 \log^2(p \vee n) \leq \delta_n n$. (b) With*

probability $1 - \Delta_n$, $\sup_{u \in \mathcal{U}} \mathbb{E}_n[r_u^2(X)] \leq Cs \log(p \vee n)/n$; $\sup_{u \in \mathcal{U}} \max_{j \leq p} |(\mathbb{E}_n - \mathbb{E}_P)[f_j^2(X)\zeta_u^2]| \leq \delta_n$; $\sup_{u, u' \in \mathcal{U}, d_{\mathcal{U}}(u, u') \leq 1/n} \max_{j \leq p} \{\mathbb{E}_n[f_j(X)^2(\zeta_u - \zeta_{u'})^2]\}^{1/2} \leq \delta_n$, and $\sup_{u, u' \in \mathcal{U}, d_{\mathcal{U}}(u, u') \leq 1/n} \|\mathbb{E}_n[f(X)(\zeta_u - \zeta_{u'})]\|_{\infty} \leq \delta_n n^{-1/2}$. (c) With probability $1 - \Delta_n$, the empirical minimum and maximum sparse eigenvalues are bounded from zero and above: $\kappa' \leq \inf_{\|\delta\|_0 \leq s\ell_n} \|f(X)' \delta\|_{\mathbb{P}_{n,2}} \leq \sup_{\|\delta\|_0 \leq s\ell_n} \|f(X)' \delta\|_{\mathbb{P}_{n,2}} \leq \kappa''$.

The following result characterizes the performance of the estimators (6.2) and (6.3) for the logistic link function case under Assumption 6.2.

Theorem 6.2 (Rates and Sparsity for Functional Response under Logistic Link). *Under Assumption 6.2 and setting the penalty and loadings as in Algorithm 6.1, for all n large enough, uniformly for all $P \in \mathcal{P}_n$ with \mathbb{P}_P probability $1 - o(1)$, the following performance bounds hold for some constant \bar{C} :*

$$\sup_{u \in \mathcal{U}} \|f(X)'(\hat{\theta}_u - \theta_u)\|_{\mathbb{P}_{n,2}} \leq \bar{C} \sqrt{\frac{s \log(p \vee n)}{n}} \quad \text{and} \quad \sup_{u \in \mathcal{U}} \|\hat{\theta}_u - \theta_u\|_1 \leq \bar{C} \sqrt{\frac{s^2 \log(p \vee n)}{n}}.$$

and the estimator is uniformly sparse: $\sup_{u \in \mathcal{U}} \|\hat{\theta}_u\|_0 \leq \bar{C}s$. For all n large enough, uniformly for all $P \in \mathcal{P}_n$, with \mathbb{P}_P probability $1 - o(1)$, the Post-Lasso estimator corresponding to $\hat{\theta}_u$ obeys

$$\sup_{u \in \mathcal{U}} \|f(X)'(\tilde{\theta}_u - \theta_u)\|_{\mathbb{P}_{n,2}} \leq \bar{C} \sqrt{\frac{s \log(p \vee n)}{n}}, \quad \text{and} \quad \sup_{u \in \mathcal{U}} \|\tilde{\theta}_u - \theta_u\|_1 \leq \bar{C} \sqrt{\frac{s^2 \log(p \vee n)}{n}}.$$

Comment 6.1. The performance bounds derived in Theorem 6.2 satisfy the conditions of Assumption 4.2 (see also Assumption H.1 in the Supplementary Material). Moreover, since in the logistic case the link function is 1-Lipschitz and the approximation errors are assumed to be small, the results above establish the same rates of convergence for estimators of the conditional probabilities; for example,

$$\sup_{u \in \mathcal{U}} \|\mathbb{E}_P[Y_u | X] - \Lambda(f(X)' \hat{\theta}_u)\|_{\mathbb{P}_{n,2}} \leq \bar{C} \sqrt{\frac{s \log(p \vee n)}{n}}.$$

7. APPLICATION: THE EFFECT OF 401(K) PARTICIPATION ON ASSET HOLDINGS

As a practical illustration of the methods developed in this paper, we consider the estimation of the effect of 401(k) eligibility and participation on accumulated assets as in Abadie (2003) and Chernozhukov and Hansen (2004). Our goal here is to illustrate the estimation results and inference statements and to make the following points that underscore our theoretical findings: 1) In a *low-dimensional setting*, where the number of controls is low and therefore there is no need for selection, our robust post-selection inference methods perform well. That is, the results of our methods agree with the results of standard methods that do not employ any selection. 2) In a *high-dimensional setting*, where there are (moderately) many controls, our post-selection inference methods perform well, producing well-behaved estimates and confidence intervals compared to the erratic estimates and confidence intervals produced by standard methods that do not employ selection as a means of regularization. 3) Finally, in a *very high-dimensional setting*, where the number of controls is comparable to the sample size, the standard methods break down completely, while our methods

still produce well-behaved estimates and confidence intervals. These findings are in line with our theoretical results about uniform validity of our inference methods.

The key problem in determining the effect of participation in 401(k) plans on accumulated assets is saver heterogeneity coupled with the fact that the decision to enroll in a 401(k) is non-random. It is generally recognized that some people have a higher preference for saving than others. It also seems likely that those individuals with high unobserved preference for saving would be most likely to choose to participate in tax-advantaged retirement savings plans and would tend to have otherwise high amounts of accumulated assets. The presence of unobserved savings preferences with these properties then implies that conventional estimates that do not account for saver heterogeneity and endogeneity of participation will be biased upward, tending to overstate the savings effects of 401(k) participation.

To overcome the endogeneity of 401(k) participation, Abadie (2003) and Chernozhukov and Hansen (2004) adopt the strategy detailed in Poterba, Venti, and Wise (1994; 1995; 1996; 2001) and Benjamin (2003), who used data from the 1991 Survey of Income and Program Participation and argue that eligibility for enrolling in a 401(k) plan in this data can be taken as exogenous after conditioning on a few observables of which the most important for their argument is income. The basic idea of their argument is that, at least around the time 401(k)'s initially became available, people were unlikely to be basing their employment decisions on whether an employer offered a 401(k) but would instead focus on income. Thus, eligibility for a 401(k) could be taken as exogenous conditional on income, and the causal effect of 401(k) eligibility could be directly estimated by appropriate comparison across eligible and ineligible individuals.¹⁹ Abadie (2003), Chernozhukov and Hansen (2004), and Ogburn et al. (2015) use this argument for the exogeneity of eligibility conditional on controls to argue that 401(k) eligibility provides a valid instrument for 401(k) participation and employ IV methods to estimate the effect of 401(k) participation on accumulated assets.

As a complement to the work cited above, we estimate various treatment effects of 401(k) participation on financial wealth using high-dimensional methods. A key component of the argument underlying the exogeneity of 401(k) eligibility is that eligibility may only be taken as exogenous after conditioning on income. Both Abadie (2003) and Chernozhukov and Hansen (2004) adopt this argument but control only for a small number of terms. One might wonder whether the small number of terms considered is sufficient to adequately control for income and other related confounds. At the same time, the power to learn anything about the effect of 401(k) participation decreases as one controls more flexibly for confounds. The methods developed in this paper offer one resolution to this tension by allowing us to consider a very broad set of controls and functional forms under the assumption that among the set of variables we consider there is a relatively low-dimensional set that adequately captures the effect of confounds. This approach is more general than that pursued

¹⁹Poterba, Venti, and Wise (1994; 1995; 1996; 2001) and Benjamin (2003) all focus on estimating the effect of 401(k) eligibility, the intention to treat parameter. Also note that there are arguments that eligibility should not be taken as exogenous given income; see, for example, Engen et al. (1996) and Engen and Gale (2000).

in previous research which implicitly assumes that confounding effects can adequately be controlled for by a small number of variables chosen *ex ante* by the researcher.

We use the same data as Chernozhukov and Hansen (2004). The data consist of 9,915 observations at the household level drawn from the 1991 SIPP. We use net financial assets as the outcome variable, Y , in our analysis. Our treatment variable, D , is an indicator for having positive 401(k) balances; and our instrument, Z , is an indicator for being eligible to enroll in a 401(k) plan. The vector of raw covariates, X , consists of age, income, family size, years of education, a married indicator, a two-earner status indicator, a defined benefit pension status indicator, an IRA participation indicator, and a home ownership indicator. Further details can be found in Chernozhukov and Hansen (2004).

We present detailed results for three different sets of controls $f(X)$. The first specification uses indicators of marital status, two-earner status, defined benefit pension status, IRA participation status, and home ownership status, second order polynomials in family size and education, a third order polynomial in age, and a quadratic spline in income with six break points²⁰ (Quadratic Spline specification). The second specification augments the Quadratic Spline specification by interacting all the non-income variables with each term in the income spline (Quadratic Spline Plus Interactions specification). The final specification forms a larger set of potential controls by starting with all of the variables from the Quadratic Spline specification and forming all two-way interactions between all of the non-income variables. The set of main effects and interactions of all non-income variables is then fully interacted with all of the income terms (Quadratic Spline Plus Many Interactions specification).²¹ The dimensions of the set of controls are thus 35, 311, and 1756 for the Quadratic Spline, Quadratic Spline Plus Interactions, and Quadratic Spline Plus Many Interactions specification, respectively. For methods that do not use variable selection, we use 32, 272, and 1526 variables resulting from removing terms that are perfectly collinear. We refer to the specification without interactions as *low-p*, to the specification with only income interactions as *high-p*, and to the specification with all two-way interactions further interacted with income as *very-high-p*.

We report a variety of results for each specification. Under the maintained assumption that 401(k) eligibility may be taken as exogenous after controlling for the variables defined in the preceding paragraph, we can estimate intention to treat effects of 401(k) eligibility using the methods of this paper using 401(k) eligibility as $D = Z$. We report the estimated average intention to treat and average intention to treat on the treated as the ATE and ATE-T, and we report estimates of quantile intention to treat and quantile intention to treat on the treated effects as QTE and QTE-T. We also directly apply the results of this paper to estimate effects of 401(k) participation, reporting

²⁰Specifically, we allow for income, income-squared, and then interact these two variables with seven dummies for the categories formed by the cut points.

²¹The specifications are motivated by the original specification used in Abadie (2003), Benjamin (2003), and Chernozhukov and Hansen (2004) allowing for data-dependent accommodation of nonlinearity. We report results based on the exact specification used in previous papers in the Supplementary Appendix.

estimates of the LATE, LATE-T, LQTE, and LQTE-T for each specification.²² For comparison, we also report estimates of the eligibility effect from the linear model without selection and with selection using the approach of Belloni et al. (2014a) and estimates of the participation effect from linear instrumental variables estimation without selection and with selection as in Chernozhukov et al. (2015).

Estimation of all these treatment effects depends on first-stage estimates of reduced form functions as detailed in Section 3. We estimate reduced form functions where the outcome is continuous using ordinary least squares when no model selection is used or Post-Lasso when selection is used. We estimate reduced form functions where the outcome is binary by logistic regression when no model selection is used or Post- ℓ_1 -penalized logistic regression when selection is used. We only report selection-based estimates in the very-high- p setting.²³ We refer to Appendix F for detailed discussion of implementing our approach in this example.

Estimates of the ATE, ATE-T, LATE and LATE-T as well as the coefficient on 401(k) eligibility from the linear model and coefficient on 401(k) participation in the linear IV model are given in Table 1. In this table, we provide point estimates for each of the three sets of controls with and without variable selection. We report conventional heteroscedasticity consistent standard error estimates for the linear model and linear IV coefficient. For the ATE, ATE-T, LATE, and LATE-T, we report both analytic and multiplier bootstrap standard errors. The bootstrap standard errors are based on 500 bootstrap replications with Mammen (1993) weights as multipliers.

Looking first at the two sets of standard error estimates for the average treatment effect estimates, we see that the bootstrap and analytic standard are quite similar and that one would not draw substantively different conclusions from using one versus the other. We also see that estimates of the effect of 401(k) eligibility using the linear model and estimates of the effect of 401(k) participation using the linear IV model are broadly consistent with each other across all specifications and regardless of whether or not variable selection is done. We also have that the estimates of the ATE, ATE-T, LATE, and LATE-T are very similar regardless of whether selection is used in the low- p Quadratic Spline specification. The ATE and ATE-T both indicate a positive and significant average effect of 401(k) eligibility; and the LATE and LATE-T suggest positive and significant effects of 401(k) participation for compliers. The similarity in the low- p case is reassuring as it illustrates that there is little impact of variable selection relative to simply including everything in a low-dimensional setting.²⁴

²²We note that because of one-sided compliance the local effects for the treated actually coincide with population effects for the treated; see Frölich and Melly (2013).

²³The estimated propensity score shows up in the denominator of the efficient moment conditions. As is conventional, we use trimming to keep the denominator bounded away from zero with trimming set to 10^{-12} . Trimming occurs in the Quadratic Spline Plus Interactions (12 observations trimmed) and Quadratic Spline Plus Many Interactions specifications (9915 observations trimmed) when selection is not done. Trimming never occurs in the selection-based estimates in this example. We choose not to report unregularized estimates in the very-high- p specification since all observations are trimmed and, in fact, have estimated propensity scores of either 0 or 1.

²⁴In the low-dimensional setting, using all available controls is semi-parametrically efficient and allows uniformly valid inference. Thus, the similarity between the results in this case is an important feature of our method which results

We observe somewhat different results in the Quadratic Spline Plus Interactions specification. For both the ATE and the LATE in the Quadratic Spline Plus Interactions case, we see a substantially larger point estimate without selection than with selection, with the selection results being similar to those obtained in the low- p case. Along with the larger point estimate, we also see that the estimated standard errors in the no selection case for the ATE and LATE are roughly three times larger than the standard errors in the selection case. For the ATE-T and LATE-T in the Quadratic Spline Plus Interactions case, point estimates following selection are notably smaller than without selection but estimated standard errors after selection are somewhat larger. We note that one might suspect estimated standard errors for all of the estimators without selection to be substantially downward biased in this case similar to estimated standard errors in the model with many controls as demonstrated in Cattaneo et al. (2010) though this issue has not been explored for the treatment effects estimators we employ. Finally, we see a large difference in the Orthogonal Polynomials Plus Many Interactions Specifications as estimates cannot even be computed reliably without selection due to severe overfitting: The estimated propensity score is either 0 or 1 for every observation.

We provide estimates of the QTE and QTE-T in Figure 1 and estimates of the LQTE and LQTE-T in Figure 2. The left column Figure 1 gives results for the QTE, and the right column displays the results for the QTE-T. Similarly, the left and right columns of Figure 2 provide the LQTE and LQTE-T respectively. We give the results for the Quadratic Spline, Quadratic Spline Plus Interactions, and Quadratic Spline Plus Many Interactions specification in the top row, middle row, and bottom row respectively. In each graphic, we use solid lines for point estimates and report uniform 95% confidence intervals with dashed lines.

Looking across the figures, we see a similar pattern to that seen for the estimates of the average effects in that the selection-based estimates are stable across all specifications and are very similar to the estimates obtained without selection from the baseline low- p Quadratic Spline specification. In the more flexible Quadratic Spline plus Interactions specification, the estimates that do not make use of selection behave somewhat erratically. This erratic behavior is especially apparent in the estimated LQTE of 401(k) participation where we observe that small changes in the quantile index may result in large swings in the point estimate of the LQTE and estimated standard errors are quite large. Again, this erratic behavior is likely due to overfitting due to the large set of variables considered. As with the average effects, estimated quantile effects without selection in the Quadratic Spline Plus Many Interactions specification are not reported as the estimated propensity score is always 0 or 1.

If we focus on the LQTE and LQTE-T estimated from variable selection methods, we find that 401(k) participation has a small impact on accumulated net total financial assets at low quantiles while appearing to have a larger impact at high quantiles. Looking at the uniform confidence intervals, we can see that this pattern is statistically significant at the 5% level and that we would reject the hypothesis that 401(k) participation has no effect and reject the hypothesis of a constant treatment effect more generally.

from our reliance on low-bias moment functions and sensible variable selection devices to produce semi-parametrically efficient estimators and uniformly valid inference statements *following* model selection.

It is also worth discussing the results of the variable selection briefly as well. Due to the number of models and variable selection steps taken, especially in computing quantile effects, it is not practical to give a complete accounting of the selected variables here. Rather, we note that for the linear model, linear IV, ATE, and LATE results, we select between two and 22 variables depending on the specification of controls and left-hand-side variable. The median number of variables selected for the QTE and LQTE results, where the median is taken across index values u , across the different specifications of controls and left-hand-side variables varies between one and 11. There is considerable variability in the number of variables selected across u though, ranging from a minimum of no variables selected to a maximum of 237 selected variables.²⁵ The selected variables themselves mostly correspond to capturing the effect of income. For example, the union of the variables selected in forming each of the reduced form quantities used for estimating the LATE in the Quadratic Spline Plus Many Interactions specification consists of 36 variables, only four of which do not include income.²⁶ This pattern of largely selecting terms that are direct income effects or interactions of income with other variables holds up across the specifications considered.

It is interesting that our results are similar to those in Chernozhukov and Hansen (2004) despite allowing for a much richer set of controls. The fact that we allow for a rich set of controls but produce similar results to those previously available lends further credibility to the claim that previous work controlled adequately for the available observables.²⁷ Finally, it is worth noting that this similarity is not mechanical or otherwise built in to the procedure. For example, applications in Belloni et al. (2012) and Belloni et al. (2014a) use high-dimensional variable selection methods and produce sets of variables that differ substantially from intuitive baselines.

APPENDIX A. NOTATION

A.1. Overall Notation. We consider a random element $W = W_P$ taking values in the measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$, with probability law $P \in \mathcal{P}$. Note that it is most convenient to think about P as a parameter in a parameter set \mathcal{P} . We shall also work with a bootstrap multiplier variable ξ taking values in $(\mathbb{R}, \mathcal{A}_{\mathbb{R}})$ that is independent of W_P , having probability law P_{ξ} , which is fixed throughout. We consider $(W_i)_{i=1}^{\infty} = (W_{i,P})_{i=1}^{\infty}$ and $(\xi_i)_{i=1}^{\infty}$ to be i.i.d. copies of W and ξ , which are

²⁵Having more than 100 variables selected occurs in the very high dimensional setting when the outcome in the penalized regression is $\mathbf{1}_0(D)Y_u$ for the six lowest values of u among the subset of households eligible for 401(k)'s and for the six highest values of u among the subset of households that are not eligible for 401(k)'s.

²⁶Let i_1 be the indicator for income in the first income category, and define $i_2 - i_7$ similarly. Let db be the defined benefit dummy, ira be the IRA dummy, $hown$ be the home ownership dummy, mar be the married dummy, te be the two-earner household dummy, ed be years of schooling, and $fsize$ be family size. The exact identities of the variables selected for modeling any reduced form quantity used in estimating the LATE in the very-high-dimensional case are $i_1, i_2, i_3, income * i_3, income^2 * i_6, db, ira * hown, age * ira, ed * ira, i_1 * fsize, i_1 * fsize^2 * db, i_2 * fsize, i_2 * fsize^2 * db, i_3 * age^3, i_3 * fsize, i_3 * mar, i_3 * fsize * te, i_3 * fsize * mar, i_4 * fsize, i_4 * te, i_4 * fsize * te, i_4 * mar * te, i_4 * ed * fsize, i_5 * ed^2 * te, i_5 * fsize * te, income * ira, income * hown, income * mar * hown, income * te * hown, income * fsize * ira, income * fsize^2 * ira, income * i_1 * fsize, income^2 * ira * hown, income^2 * ed^2 * te, income^2 * i_3 * fsize,$ and $income^2 * i_6 * hown$.

²⁷Of course, the estimates are still not valid causal estimates if one does not believe that 401(k) eligibility can be taken as exogenous after controlling for income and the other included variables.

also independent of each other. The data will be defined as some measurable function of W_i for $i = 1, \dots, n$, where n denotes the sample size.

We require the sequences $(W_i)_{i=1}^\infty$ and $(\xi_i)_{i=1}^\infty$ to live on a probability space $(\Omega, \mathcal{A}_\Omega, P_P)$ for all $P \in \mathcal{P}$; note that other variables arising in the proofs do not need to live on the same space. It is important to keep track of the dependence on P in the analysis since we want the results to hold uniformly in P in some set \mathcal{P}_n which may be dependent on n . Typically, this set will increase with n ; i.e. $\mathcal{P}_n \subseteq \mathcal{P}_{n+1}$.

Throughout the paper we signify the dependence on P by mostly using P as a subscript in P_P , but in the proofs we sometimes use it as a subscript for variables as in W_P . The operator \mathbb{E} denotes a generic expectation operator with respect to a generic probability measure P , while \mathbb{E}_P denotes the expectation with respect to P_P . Note also that we use capital letters such as W to denote random elements and use the corresponding lower case letters such as w to denote fixed values that these random elements can take.

We denote by \mathbb{P}_n the (random) empirical probability measure that assigns probability n^{-1} to each $W_i \in (W_i)_{i=1}^n$. \mathbb{E}_n denotes the expectation with respect to the empirical measure, and $\mathbb{G}_{n,P}$ denotes the empirical process $\sqrt{n}(\mathbb{E}_n - P)$, i.e.

$$\mathbb{G}_{n,P}(f) = \mathbb{G}_{n,P}(f(W)) = n^{-1/2} \sum_{i=1}^n \{f(W_i) - P[f(W)]\}, \quad P[f(W)] := \int f(w) dP(w),$$

indexed by a measurable class of functions $\mathcal{F} : \mathcal{W} \mapsto \mathbb{R}$; see van der Vaart and Wellner (1996, chap. 2.3). We shall often omit the index P from $\mathbb{G}_{n,P}$ and simply write \mathbb{G}_n . In what follows, we use $\|\cdot\|_{P,q}$ to denote the $L^q(P)$ norm; for example, we use $\|f(W)\|_{P,q} = (\int |f(w)|^q dP(w))^{1/q}$ and $\|f(W)\|_{\mathbb{P}_n,q} = (n^{-1} \sum_{i=1}^n |f(W_i)|^q)^{1/q}$. For a vector $v = (v_1, \dots, v_p)' \in \mathbb{R}^p$, $\|v\|_1 = |v_1| + \dots + |v_p|$ denotes the ℓ_1 -norm of v , $\|v\| = \sqrt{v'v}$ denotes the Euclidean norm of v , and $\|v\|_0$ denotes the ℓ_0 -“norm” of v which equals the number of non-zero components of v . For a positive integer k , $[k]$ denotes the set $\{1, \dots, k\}$. For x_n, y_n denoting sequences in \mathbb{R} , the statement $x_n \lesssim y_n$ means that $x_n \leq Ay_n$ for some constant A that does not depend on n .

We say that a collection of random variables $\mathcal{F} = \{f(W, t), t \in T\}$, where $f : \mathcal{W} \times T \rightarrow \mathbb{R}$, indexed by a set T and viewed as functions of $W \in \mathcal{W}$, is *suitably measurable* with respect to W if it is image admissible Suslin class, as defined in Dudley (1999, p 186). In particular, \mathcal{F} is suitably measurable if $f : \mathcal{W} \times T \rightarrow \mathbb{R}$ is measurable and T is a Polish space equipped with its Borel sigma algebra, see Dudley (1999, p 186). This condition is a mild assumption satisfied in practical cases.

A.2. Notation for Stochastic Convergence Uniformly in P . All parameters, such as the law of the data, are indexed by P . This dependency is sometimes kept implicit. We shall allow for the possibility that the probability measure $P = P_n$ can depend on n . We shall conduct our stochastic convergence analysis uniformly in P , where P can vary within some set \mathcal{P}_n , which itself may vary with n .

The convergence analysis, namely the stochastic order relations and convergence in distribution, uniformly in $P \in \mathcal{P}_n$ and the analysis under all sequences $P_n \in \mathcal{P}_n$ are equivalent. Specifically, consider a sequence of stochastic processes $X_{n,P}$ and a random element Y_P , taking values in the

normed space \mathbb{D} , defined on the probability space $(\Omega, \mathcal{A}_\Omega, P_P)$. Through most of the Appendix $\mathbb{D} = \ell^\infty(\mathcal{U})$, the space of uniformly bounded functions mapping an arbitrary index set \mathcal{U} to the real line, or $\mathbb{D} = UC(\mathcal{U})$, the space of uniformly continuous functions mapping an arbitrary index set \mathcal{U} to the real line. Consider also a sequence of deterministic positive constants a_n . We shall say that

- (i) $X_{n,P} = O_P(a_n)$ uniformly in $P \in \mathcal{P}_n$, if $\lim_{K \rightarrow \infty} \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} P_P^*(|X_{n,P}| > Ka_n) = 0$,
- (ii) $X_{n,P} = o_P(a_n)$ uniformly in $P \in \mathcal{P}_n$, if $\sup_{K > 0} \lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} P_P^*(|X_{n,P}| > Ka_n) = 0$,
- (iii) $X_{n,P} \rightsquigarrow Y_P$ uniformly in $P \in \mathcal{P}_n$, if $\sup_{P \in \mathcal{P}_n} \sup_{h \in \text{BL}_1(\mathbb{D})} |\mathbb{E}_P^* h(X_{n,P}) - \mathbb{E}_P h(Y_P)| \rightarrow 0$.

Here the symbol \rightsquigarrow denotes weak convergence, i.e. convergence in distribution or law, $\text{BL}_1(\mathbb{D})$ denotes the space of functions mapping \mathbb{D} to $[0, 1]$ with Lipschitz norm at most 1, and the outer probability and expectation, P_P^* and \mathbb{E}_P^* , are invoked whenever (non)-measurability arises.

Lemma A.1. *The above notions (i), (ii) and (iii) are equivalent to the following notions (a), (b), and (c), each holding for every sequence $P_n \in \mathcal{P}_n$:*

- (a) $X_{n,P_n} = O_{P_n}(a_n)$, i.e. $\lim_{K \rightarrow \infty} \lim_{n \rightarrow \infty} P_{P_n}^*(|X_{n,P_n}| > Ka_n) = 0$;
- (b) $X_{n,P_n} = o_{P_n}(a_n)$, i.e. $\sup_{K > 0} \lim_{n \rightarrow \infty} P_{P_n}^*(|X_{n,P_n}| > Ka_n) = 0$;
- (c) $X_{n,P_n} \rightsquigarrow Y_{P_n}$, i.e. $\sup_{h \in \text{BL}_1(\mathbb{D})} |\mathbb{E}_{P_n}^* h(X_{n,P_n}) - \mathbb{E}_{P_n} h(Y_{P_n})| \rightarrow 0$.

The claims follow straightforwardly from the definitions, so the proof is omitted. We shall use this equivalence extensively in the proofs of the main results without explicit reference.

APPENDIX B. KEY TOOLS I: UNIFORM IN P DONSKER THEOREM, MULTIPLIER BOOTSTRAP, AND FUNCTIONAL DELTA METHOD

B.1. Uniform in P Donsker Property. Let $(W_i)_{i=1}^\infty$ be a sequence of i.i.d. copies of the random element W taking values in the measure space $(\mathcal{W}, \mathcal{A}_\mathcal{W})$ according to the probability law P on that space. Let $\mathcal{F}_P = \{f_{t,P} : t \in T\}$ be a set of suitably measurable functions $w \mapsto f_{t,P}(w)$ mapping \mathcal{W} to \mathbb{R} , equipped with a measurable envelope $F_P : \mathcal{W} \mapsto \mathbb{R}$. The class is indexed by $P \in \mathcal{P}$ and $t \in T$, where T is a fixed, totally bounded semi-metric space equipped with a semi-metric d_T . Let $N(\epsilon, \mathcal{F}_P, \|\cdot\|_{Q,2})$ denote the ϵ -covering number of the class of functions \mathcal{F}_P with respect to the $L^2(Q)$ seminorm $\|\cdot\|_{Q,2}$ for Q a finitely-discrete measure on $(\mathcal{W}, \mathcal{A}_\mathcal{W})$. We shall use the following result.

Theorem B.1 (Uniform in P Donsker Property). *Work with the set-up above. Suppose that for $q > 2$*

$$\sup_{P \in \mathcal{P}} \|F_P\|_{P,q} \leq C \text{ and } \lim_{\delta \searrow 0} \sup_{P \in \mathcal{P}} \sup_{d_T(t,\bar{t}) \leq \delta} \|f_{t,P} - f_{\bar{t},P}\|_{P,2} = 0. \quad (\text{B.1})$$

Furthermore, suppose that

$$\lim_{\delta \searrow 0} \sup_{P \in \mathcal{P}} \int_0^\delta \sup_Q \sqrt{\log N(\epsilon \|F_P\|_{Q,2}, \mathcal{F}_P, \|\cdot\|_{Q,2})} d\epsilon = 0. \quad (\text{B.2})$$

Let \mathbb{G}_P denote the P -Brownian Bridge, and consider

$$Z_{n,P} := (Z_{n,P}(t))_{t \in T} := (\mathbb{G}_n(f_{t,P}))_{t \in T}, \quad Z_P := (Z_P(t))_{t \in T} := (\mathbb{G}_P(f_{t,P}))_{t \in T}.$$

(a) Then, $Z_{n,P} \rightsquigarrow Z_P$ in $\ell^\infty(T)$ uniformly in $P \in \mathcal{P}$, namely

$$\sup_{P \in \mathcal{P}} \sup_{h \in BL_1(\ell^\infty(T))} |\mathbb{E}_P^* h(Z_{n,P}) - \mathbb{E}_P h(Z_P)| \rightarrow 0.$$

(b) The process $Z_{n,P}$ is stochastically equicontinuous uniformly in $P \in \mathcal{P}$, i.e., for every $\varepsilon > 0$,

$$\lim_{\delta \searrow 0} \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} \mathbb{P}_P^* \left(\sup_{d_T(t, \bar{t}) \leq \delta} |Z_{n,P}(t) - Z_{n,P}(\bar{t})| > \varepsilon \right) = 0.$$

(c) The limit process Z_P has the following continuity properties:

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{t \in T} |Z_P(t)| < \infty, \quad \lim_{\delta \searrow 0} \sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{d_T(t, \bar{t}) \leq \delta} |Z_P(t) - Z_P(\bar{t})| = 0.$$

(d) The paths $t \mapsto Z_P(t)$ are a.s. uniformly continuous on (T, d_T) under each $P \in \mathcal{P}$.

Comment B.1. [Important Feature of the Theorem] This is an extension of the uniform Donsker theorem stated in Theorem 2.8.2 in van der Vaart and Wellner (1996), which allows for the function classes \mathcal{F} to be **dependent on P** themselves. This generalization is crucial and is required in all of our problems.

B.2. Uniform in P Validity of Multiplier Bootstrap. Consider the setting of the preceding subsection. Let $(\xi)_{i=1}^n$ be i.i.d multipliers whose distribution does not depend on P , such that $\mathbb{E}\xi = 0$, $\mathbb{E}\xi^2 = 1$, and $\mathbb{E}|\xi|^q \leq C$ for $q > 2$. Consider the multiplier empirical process:

$$Z_{n,P}^* := (Z_{n,P}^*(t))_{t \in T} := (\mathbb{G}_n(\xi f_{t,P}))_{t \in T} := \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i f_{t,P}(W_i) \right)_{t \in T}.$$

Here \mathbb{G}_n is taken to be an extended empirical processes defined by the empirical measure that assigns mass $1/n$ to each point (W_i, ξ_i) for $i = 1, \dots, n$. Let $Z_P = (Z_P(t))_{t \in T} = (\mathbb{G}_P(f_{t,P}))_{t \in T}$ as defined in Theorem B.1.

Theorem B.2 (Uniform in P Validity of Multiplier Bootstrap). *Assume the conditions of Theorem B.1 hold. Then (a) the following unconditional convergence takes place, $Z_{n,P}^* \rightsquigarrow Z_P$ in $\ell^\infty(T)$ uniformly in $P \in \mathcal{P}$, namely*

$$\sup_{P \in \mathcal{P}} \sup_{h \in BL_1(\ell^\infty(T))} |\mathbb{E}_P^* h(Z_{n,P}^*) - \mathbb{E}_P h(Z_P)| \rightarrow 0,$$

and (b) the following conditional convergence takes place, $Z_{n,P}^* \rightsquigarrow_B Z_P$ in $\ell^\infty(T)$ uniformly in $P \in \mathcal{P}$, namely uniformly in $P \in \mathcal{P}$

$$\sup_{h \in BL_1(\ell^\infty(T))} |\mathbb{E}_{B_n} h(Z_{n,P}^*) - \mathbb{E}_P h(Z_P)| = o_P^*(1),$$

where \mathbb{E}_{B_n} denotes the expectation over the multiplier weights $(\xi_i)_{i=1}^n$ holding the data $(W_i)_{i=1}^n$ fixed.

B.3. Uniform in P Functional Delta Method and Bootstrap. We shall use the functional delta method, as formulated in van der Vaart and Wellner (1996, Chap. 3.9). Let \mathbb{D}_0 , \mathbb{D} , and \mathbb{E} be normed spaces, with $\mathbb{D}_0 \subset \mathbb{D}$. A map $\phi : \mathbb{D}_\phi \subset \mathbb{D} \mapsto \mathbb{E}$ is called *Hadamard-differentiable* at $\rho \in \mathbb{D}_\phi$ tangentially to \mathbb{D}_0 if there is a continuous linear map $\phi'_\rho : \mathbb{D}_0 \mapsto \mathbb{E}$ such that

$$\frac{\phi(\rho + t_n h_n) - \phi(\rho)}{t_n} \rightarrow \phi'_\rho(h), \quad n \rightarrow \infty,$$

for all sequences $t_n \rightarrow 0$ in \mathbb{R} and $h_n \rightarrow h \in \mathbb{D}_0$ in \mathbb{D} such that $\rho + t_n h_n \in \mathbb{D}_\phi$ for every n .

We now define the following notion of the uniform Hadamard differentiability:

Definition B.1 (Uniform Hadamard Tangential Differentiability). Consider a map $\phi : \mathbb{D}_\phi \mapsto \mathbb{E}$, where the domain of the map \mathbb{D}_ϕ is a subset of a normed space \mathbb{D} and the range is a subset of the normed space \mathbb{E} . Let \mathbb{D}_0 be a normed space, with $\mathbb{D}_0 \subset \mathbb{D}$, and \mathbb{D}_ρ be a compact metric space, a subset of \mathbb{D}_ϕ . The map $\phi : \mathbb{D}_\phi \mapsto \mathbb{E}$ is called *Hadamard-differentiable uniformly in $\rho \in \mathbb{D}_\rho$ tangentially to \mathbb{D}_0 with derivative map $h \mapsto \phi'_\rho(h)$* , if

$$\left| \frac{\phi(\rho_n + t_n h_n) - \phi(\rho_n)}{t_n} - \phi'_\rho(h) \right| \rightarrow 0, \quad \left| \phi'_{\rho_n}(h_n) - \phi'_\rho(h) \right| \rightarrow 0, \quad n \rightarrow \infty,$$

for all convergent sequences $\rho_n \rightarrow \rho$ in \mathbb{D}_ρ , $t_n \rightarrow 0$ in \mathbb{R} , and $h_n \rightarrow h \in \mathbb{D}_0$ in \mathbb{D} such that $\rho_n + t_n h_n \in \mathbb{D}_\phi$ for every n . As a part of the definition, we require that the derivative map $h \mapsto \phi'_\rho(h)$ from \mathbb{D}_0 to \mathbb{E} is linear for each $\rho \in \mathbb{D}_\rho$. ■

Comment B.2. Note that the definition requires that the derivative map $(\rho, h) \mapsto \phi'_\rho(h)$, mapping $\mathbb{D}_\rho \times \mathbb{D}_0$ to \mathbb{E} , is continuous at each $(\rho, h) \in \mathbb{D}_\rho \times \mathbb{D}_0$. ■

Comment B.3 (Important Details of the Definition). Definition B.1 is different from the definition of uniform differentiability given in van der Vaart and Wellner (1996, p. 379, eq. (3.9.12)), since our definition allows \mathbb{D}_ρ to be much smaller than \mathbb{D}_ϕ and allows \mathbb{D}_ρ to be endowed with a much stronger metric than the metric induced by the norm of \mathbb{D} . These differences are essential for infinite-dimensional applications. For example, the quantile/inverse map is uniformly Hadamard differentiable in the sense of Definition B.1 for a suitable choice of \mathbb{D}_ρ : Let $T = [\epsilon, 1 - \epsilon]$, $\mathbb{D} = \ell^\infty(T)$, $\mathbb{D}_\phi =$ set of cadlag functions on T , $\mathbb{D}_0 = \text{UC}(T)$, and \mathbb{D}_ρ be a compact subset of $C^1(T)$ such that each $\rho \in \mathbb{D}_\rho$ obeys $\partial\rho(t)/\partial t \geq c > 0$ on $t \in T$, where c is a positive constant. However, the quantile/inverse map is not Hadamard differentiable uniformly on \mathbb{D}_ρ if we set $\mathbb{D}_\rho = \mathbb{D}_\phi$ and hence is not uniformly differentiable in the sense of the definition given in van der Vaart and Wellner (1996) which requires $\mathbb{D}_\rho = \mathbb{D}_\phi$. It is important and practical to keep the distinction between \mathbb{D}_ρ and \mathbb{D}_ϕ since the estimated values $\hat{\rho}$ may well be outside \mathbb{D}_ρ unless explicitly imposed in estimation even though the population values of ρ are in \mathbb{D}_ρ by assumption. For example, the empirical cdf is in \mathbb{D}_ϕ , but is outside \mathbb{D}_ρ . ■

Theorem B.3 (Functional delta-method uniformly in $P \in \mathcal{P}$). Let $\phi : \mathbb{D}_\phi \subset \mathbb{D} \mapsto \mathbb{E}$ be Hadamard-differentiable uniformly in $\rho \in \mathbb{D}_\rho \subset \mathbb{D}_\phi$ tangentially to \mathbb{D}_0 with derivative map ϕ'_ρ . Let $\hat{\rho}_{n,P}$ be a sequence of stochastic processes taking values in \mathbb{D}_ϕ , where each $\hat{\rho}_{n,P}$ is an estimator of the parameter $\rho_P \in \mathbb{D}_\rho$. Suppose there exists a sequence of constants $r_n \rightarrow \infty$ such

that $Z_{n,P} = r_n(\widehat{\rho}_{n,P} - \rho_P) \rightsquigarrow Z_P$ in \mathbb{D} uniformly in $P \in \mathcal{P}_n$. The limit process Z_P is separable and takes its values in \mathbb{D}_0 for all $P \in \mathcal{P} = \cup_{n \geq n_0} \mathcal{P}_n$, where n_0 is fixed. Moreover, the set of stochastic processes $\{Z_P : P \in \mathcal{P}\}$ is relatively compact in the topology of weak convergence in \mathbb{D}_0 , that is, every sequence in this set can be split into weakly convergent subsequences. Then, $r_n(\phi(\widehat{\rho}_{n,P}) - \phi(\rho_P)) \rightsquigarrow \phi'_{\rho_P}(Z_P)$ in \mathbb{E} uniformly in $P \in \mathcal{P}_n$. If $(\rho, h) \mapsto \phi'_\rho(h)$ is defined and continuous on the whole of $\mathbb{D}_\rho \times \mathbb{D}$, then the sequence $r_n(\phi(\widehat{\rho}_{n,P}) - \phi(\rho_P)) - \phi'_{\rho_P}(r_n(\widehat{\rho}_{n,P} - \rho_P))$ converges to zero in outer probability uniformly in $P \in \mathcal{P}_n$. Moreover, the set of stochastic processes $\{\phi'_{\rho_P}(Z_P) : P \in \mathcal{P}\}$ is relatively compact in the topology of weak convergence in \mathbb{E} .

The following result on the functional delta method applies to any bootstrap or other simulation method obeying certain conditions. This includes the multiplier bootstrap as a special case. Let $D_{n,P} = (W_{i,P})_{i=1}^n$ denote the data vector and $B_n = (\xi_i)_{i=1}^n$ be a vector of random variables, used to generate bootstrap or simulation draws (this may depend on the particular method). Consider sequences of stochastic processes $\widehat{\rho}_{n,P} = \widehat{\rho}_{n,P}(D_{n,P})$, where $Z_{n,P} = r_n(\widehat{\rho}_{n,P} - \rho_P) \rightsquigarrow Z_P$ in the normed space \mathbb{D} uniformly in $P \in \mathcal{P}_n$. Also consider the bootstrap stochastic process $Z_{n,P}^* = Z_{n,P}(D_{n,P}, B_n)$ in \mathbb{D} , where $Z_{n,P}$ is a measurable function of B_n for each value of D_n . Suppose that $Z_{n,P}^*$ converges conditionally given D_n in distribution to Z_P uniformly in $P \in \mathcal{P}_n$, namely that

$$\sup_{h \in \text{BL}_1(\mathbb{D})} |\mathbb{E}_{B_n}[h(Z_{n,P}^*)] - \mathbb{E}_P h(Z_P)| = o_P^*(1),$$

uniformly in $P \in \mathcal{P}_n$, where \mathbb{E}_{B_n} denotes the expectation computed with respect to the law of B_n holding the data $D_{n,P}$ fixed. This is denoted as “ $Z_{n,P}^* \rightsquigarrow_B Z_P$ uniformly in $P \in \mathcal{P}_n$.” Finally, let $\widehat{\rho}_{n,P}^* = \widehat{\rho}_{n,P} + Z_{n,P}^*/r_n$ denote the bootstrap or simulation draw of $\widehat{\rho}_{n,P}$.

Theorem B.4 (Uniform in P functional delta-method for bootstrap and other simulation methods). *Assume the conditions of Theorem B.3 hold. Let $\widehat{\rho}_{n,P}$ and $\widehat{\rho}_{n,P}^*$ be maps as indicated previously taking values in \mathbb{D}_ϕ such that $r_n(\widehat{\rho}_{n,P} - \rho_P) \rightsquigarrow Z_P$ and $r_n(\widehat{\rho}_{n,P}^* - \widehat{\rho}_{n,P}) \rightsquigarrow_B Z_P$ in \mathbb{D} uniformly in $P \in \mathcal{P}_n$. Then, $X_{n,P}^* = r_n(\phi(\widehat{\rho}_{n,P}^*) - \phi(\widehat{\rho}_{n,P})) \rightsquigarrow_B X_P = \phi'_{\rho_P}(Z_P)$ uniformly in $P \in \mathcal{P}_n$.*

B.4. Proof of Theorem B.1. Part (a) and (b) are a direct consequence of Lemma B.2. In particular, Lemma B.2(a) implies stochastic equicontinuity under arbitrary subsequences $P_n \in \mathcal{P}$, which implies part (b). Part (a) follows from Lemma B.2(b) by splitting an arbitrary sequence $n \in \mathbb{N}$ into subsequences $n \in \mathbb{N}'$ along each of which the covariance function $(t, s) \mapsto c_{P_n}(t, s) := P_n f_{s,P_n} f_{t,P_n} - P_n f_{s,P_n} P_n f_{t,P_n}$ converges uniformly and therefore also pointwise to a uniformly continuous function on (T, d_T) . This is possible because $\{(t, s) \mapsto c_P(t, s) : P \in \mathcal{P}\}$ is a relatively compact set in $\ell^\infty(T \times T)$ in view of the Arzela-Ascoli Theorem, the assumptions in equation (B.1), and total boundedness of (T, d_T) . By Lemma B.2(b) pointwise convergence of the covariance function implies weak convergence to a tight Gaussian process which may depend on the identity \mathbb{N}' of the subsequence. Since this argument applies to each such subsequence that split the overall sequence, part (b) follows.

Part (c) is immediate from the imposed uniform covering entropy condition and Dudley’s metric entropy inequality for expectations of suprema of Gaussian processes (Corollary 2.2.8 in van der

Vaart and Wellner (1996)). Claim (d) follows from claim (c) and a standard argument, based on the application of the Borel-Cantelli lemma. Indeed, let $m \in \mathbb{N}$ be a sequence and $\delta_m := 2^{-m} \wedge \sup \left\{ \delta > 0 : \sup_{P \in \mathcal{P}} \mathbb{E}_P \sup_{d_T(t, \bar{t}) \leq \delta} |Z_P(t) - Z_P(\bar{t})| < 2^{-2m} \right\}$, then by the Markov inequality $\mathbb{P}_P \left(\sup_{d_T(t, \bar{t}) \leq \delta_m} |Z_P(t) - Z_P(\bar{t})| > 2^{-m} \right) \leq 2^{-2m+m} = 2^{-m}$. This sums to a finite number over $m \in \mathbb{N}$. Hence, by the Borel-Cantelli lemma, for almost all states $\omega \in \Omega$, $|Z_P(t)(\omega) - Z_P(\bar{t})(\omega)| \leq 2^{-m}$ for all $d_T(t, \bar{t}) \leq \delta_m \leq 2^{-m}$ and all m sufficiently large. Hence claim (d) follows. \blacksquare

B.5. Proof of Theorem B.2. Claim (a) is verified by invoking Theorem B.1. We begin by showing that $Z_P^* = (\mathbb{G}_P \xi f_{t,P})_{t \in T}$ is equal in distribution to $Z_P = (\mathbb{G}_P f_{t,P})_{t \in T}$, in particular, Z_P^* and Z_P share identical mean and covariance function, and thus they share the continuity properties established in Theorem B.1. This claim is immediate from the fact that multiplication by ξ of each $f \in \mathcal{F}_P = \{f_{t,P} : t \in T\}$ yields a set $\xi \mathcal{F}_P$ of measurable functions $\xi f : (w, \xi) \mapsto \xi f(w)$, mapping $\mathcal{W} \times \mathbb{R}$ to \mathbb{R} . Each such function has mean zero under $P \times P_\xi$, i.e. $\int s f(w) dP_\xi(s) dP(w) = 0$, and the covariance function $(\xi f, \xi \tilde{f}) \mapsto P f \tilde{f} - P f P \tilde{f}$. Hence the Gaussian process $(\mathbb{G}_P(\xi f))_{\xi f \in \xi \mathcal{F}_P}$ shares the zero mean and the covariance function of $(\mathbb{G}_P(f))_{f \in \mathcal{F}_P}$.

We are claiming that $Z_{n,P}^* \rightsquigarrow Z_P^*$ in $\ell^\infty(T)$ uniformly in $P \in \mathcal{P}$, where $Z_{n,P}^* := (\mathbb{G}_n \xi f_{t,P})_{t \in T}$. We note that the function class \mathcal{F}_P and the corresponding envelope F_P satisfy the conditions of Theorem B.1. The same is also true for the function class $\xi \mathcal{F}_P$ defined by $(w, \xi) \mapsto \xi f_P(w)$, which maps $\mathcal{W} \times \mathbb{R}$ to \mathbb{R} and its envelope $|\xi| F_P$, since ξ is independent of W . Let Q now denote a finitely discrete measure over $\mathcal{W} \times \mathbb{R}$. By Lemma L.1 multiplication by ξ does not change qualitatively the uniform covering entropy bound:

$$\log \sup_Q N(\epsilon \| |\xi| F_P \|_{Q,2}, \xi \mathcal{F}_P, \|\cdot\|_{Q,2}) \leq \log \sup_Q N(2^{-1} \epsilon \| F_P \|_{Q,2}, \mathcal{F}_P, \|\cdot\|_{Q,2}).$$

Moreover, multiplication by ξ does not affect the norms, $\|\xi f_P(W)\|_{P \times P_\xi, 2} = \|f_P(W)\|_{P, 2}$, since ξ is independent of W by construction and $\mathbb{E} \xi^2 = 1$. The claim then follows.

Claim (b). For each $\delta > 0$ and $t \in T$, let $\pi_\delta(t)$ denote a closest element in a given, finite δ -net over T . We begin by noting that

$$\begin{aligned} \Delta_P &:= \sup_{h \in \text{BL}_1} |\mathbb{E}_{B_n} h(Z_{n,P}^*) - \mathbb{E}_P h(Z_P)| \\ &\leq I_P + II_P + III_P := \sup_{h \in \text{BL}_1} |\mathbb{E}_P h(Z_P \circ \pi_\delta) - \mathbb{E}_P h(Z_P)| \\ &\quad + \sup_{h \in \text{BL}_1} |\mathbb{E}_{B_n} h(Z_{n,P}^* \circ \pi_\delta) - \mathbb{E}_P h(Z_P \circ \pi_\delta)| + \sup_{h \in \text{BL}_1} |\mathbb{E}_{B_n} h(Z_{n,P}^* \circ \pi_\delta) - \mathbb{E}_{B_n} h(Z_{n,P}^*)|, \end{aligned}$$

where here and below BL_1 abbreviates $\text{BL}_1(\ell^\infty(T))$.

First, we note that $I_P \leq \mathbb{E}_P \left(\sup_{d_T(t, \bar{t}) \leq \delta} |Z_P(t) - Z_P(\bar{t})| \wedge 2 \right) =: \mu_P(\delta)$ and $\lim_{\delta \searrow 0} \sup_{P \in \mathcal{P}} \mu_P(\delta) = 0$. The first assertion follows from

$$I_P \leq \sup_{h \in \text{BL}_1} \mathbb{E}_P |h(Z_{n,P}^* \circ \pi_\delta) - h(Z_{n,P}^*)| \leq \mathbb{E}_P \left(\sup_{t \in T} |Z_P \circ \pi_\delta(t) - Z_P(t)| \wedge 2 \right) \leq \mu_P(\delta),$$

and the second assertion holds by Theorem B.1 (c).

Second, $E_P^* III_P \leq E_P^* \left(\sup_{d_T(t, \bar{t}) \leq \delta} |Z_{n,P}^*(t) - Z_{n,P}^*(\bar{t})| \wedge 2 \right) =: \mu_P^*(\delta)$ and $\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} |\mu_P^*(\delta) - \mu_P(\delta)| = 0$. The first assertion follows because $E_P^* III_P$ is bounded by

$$E_P^* \sup_{h \in \text{BL}_1} E_{B_n} |h(Z_{n,P}^* \circ \pi_\delta) - h(Z_{n,P}^*)| \leq E_P^* E_{B_n} \left(\sup_{t \in \bar{T}} |Z_{n,P}^* \circ \pi_\delta(t) - Z_{n,P}^*(t)| \wedge 2 \right) \leq \mu_P^*(\delta).$$

The second assertion holds by part (a) of the present theorem.

Define $\epsilon(\delta) := \delta \vee \sup_{P \in \mathcal{P}} \mu_P(\delta)$. Then, by Markov's inequality, followed by $n \rightarrow \infty$,

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P_P^* \left(III_P > \sqrt{\epsilon(\delta)} \right) \leq \limsup_{n \rightarrow \infty} \frac{\sup_{P \in \mathcal{P}} \mu_P^*(\delta)}{\sqrt{\epsilon(\delta)}} \leq \frac{\sup_{P \in \mathcal{P}} \mu_P(\delta)}{\sqrt{\epsilon(\delta)}} \leq \sqrt{\epsilon(\delta)}.$$

Finally, by Lemma B.1, for each $\varepsilon > 0$, $\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P_P^* (II_P > \varepsilon) = 0$.

We can now conclude. Note that $\epsilon(\delta) \searrow 0$ if $\delta \searrow 0$, which holds by the definition of $\epsilon(\delta)$ and the property $\sup_{P \in \mathcal{P}} \mu_P(\delta) \searrow 0$ if $\delta \searrow 0$ noted above. Hence for each $\varepsilon > 0$ and all $0 < \delta < \delta_\varepsilon$ such that $3\sqrt{\epsilon(\delta)} < \varepsilon$,

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P_P^* (\Delta_P > \varepsilon) \leq \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P_P^* \left(I_P + II_P + III_P > 3\sqrt{\epsilon(\delta)} \right) \leq \sqrt{\epsilon(\delta)}.$$

Sending $\delta \searrow 0$ gives the result. ■

B.6. Auxiliary Result: Conditional Multiplier CLT in \mathbb{R}^d uniformly in $P \in \mathcal{P}$. We rely on the following lemma, which is apparently new. (An analogous result can be derived for almost sure convergence from well-known non-uniform multiplier central limit theorems, but this strategy requires us to put all the variables indexed by P on the single underlying probability space, which is much less convenient in applications.)

Lemma B.1 (Conditional Multiplier Central Limit Theorem in \mathbb{R}^d uniformly in $P \in \mathcal{P}$). *Let $(Z_{i,P})_{i=1}^\infty$ be i.i.d. random vectors on \mathbb{R}^d , indexed by a parameter $P \in \mathcal{P}$. The parameter P represents probability laws on \mathbb{R}^d . For each $P \in \mathcal{P}$, these vectors are assumed to be independent of the i.i.d. sequence $(\xi_i)_{i=1}^\infty$ with $E\xi_1 = 0$ and $E\xi_1^2 = 1$. There exist constants $2 < q < \infty$ and $0 < M < \infty$, such that $E_P Z_{1,P} = 0$ and $(E_P \|Z_{1,P}\|^q)^{1/q} \leq M$ uniformly for all $P \in \mathcal{P}$. Then, for every $\varepsilon > 0$*

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}} P_P^* \left(\sup_{h \in \text{BL}_1(\mathbb{R}^d)} \left| E_{B_n} h \left(n^{-1/2} \sum_{i=1}^n \xi_i Z_{i,P} \right) - E_P h \left(N(0, E_P Z_{1,P} Z'_{1,P}) \right) \right| > \varepsilon \right) = 0,$$

where E_{B_n} denotes the expectation over $(\xi_i)_{i=1}^n$ holding $(Z_{i,P})_{i=1}^n$ fixed.

Proof of Lemma B.1. Let X and Y be random variables in \mathbb{R}^d , then define $d_{BL}(X, Y) := \sup_{h \in \text{BL}_1(\mathbb{R}^d)} |Eh(X) - Eh(Y)|$. It suffices to show that for any sequence $P_n \in \mathcal{P}$ and $N^* \sim n^{-1/2} \sum_{i=1}^n \xi_i Z_{i,P_n} \mid (Z_{i,P_n})_{i=1}^n$, $d_{BL} \left(N^*, N(0, E_{P_n} Z_{1,P_n} Z'_{1,P_n}) \right) \rightarrow 0$ in probability (under P_{P_n}).

Following Bickel and Freedman (1981), we shall rely on the Mallow's metric, written m_r , which is a metric on the space of distribution functions on \mathbb{R}^d . For our purposes it suffices to recall that given a sequence of distribution functions $\{F_k\}$ and a distribution function F , $m_r(F_k, F) \rightarrow 0$ if and only

if $\int g dF_k \rightarrow \int g dF$ for each continuous and bounded $g : \mathbb{R}^d \rightarrow \mathbb{R}$, and $\int \|z\|^r dF_k(z) \rightarrow \int \|z\|^r dF(z)$. See Bickel and Freedman (1981) for the definition of m_r .

Under the assumptions of the lemma, we can split the sequence $n \in \mathbb{N}$ into subsequences $n \in \mathbb{N}'$, along each of which the distribution function of Z_{1,P_n} converges to some distribution function F' with respect to the Mallow's metric m_r , for some $2 < r < q$. This also implies that $N(0, E_{P_n} Z_{1,P_n} Z'_{1,P_n})$ converges weakly to a normal limit $N(0, Q')$ with $Q' = \int z z' dF'(z)$ such that $\|Q'\| \leq M$. Both Q' and F' can depend on the subsequence \mathbb{N}' .

Let F_k be the empirical distribution function of a sequence $(z_i)_{i=1}^k$ of constant vectors in \mathbb{R}^d , where $k \in \mathbb{N}$. The law of $N_{F_k}^* = k^{-1/2} \sum_{i=1}^k \xi_i z_i$ is completely determined by F_k and the law of ξ (the latter is fixed, so it does not enter as the subscript in the definition of $N_{F_k}^*$). If $m_r(F_k, F') \rightarrow 0$ as $k \rightarrow \infty$, then $d_{BL}(N_{F_k}^*, N(0, Q')) \rightarrow 0$ by Lindeberg's central limit theorem.

Let \mathbb{F}_n denote the empirical distribution function of $(Z_{i,P_n})_{i=1}^n$. Note that $N^* = N_{\mathbb{F}_n}^* \sim n^{-1/2} \sum_{i=1}^n \xi_i Z_{i,P_n} \mid (Z_{i,P_n})_{i=1}^n$. By the law of large numbers for arrays, $\int g d\mathbb{F}_n \rightarrow \int g dF'$ and $\int \|z\|^r d\mathbb{F}_n(z) \rightarrow \int \|z\|^r dF'(z)$ in probability along the subsequence $n \in \mathbb{N}'$. Hence $m_r(\mathbb{F}_n, F') \rightarrow 0$ in probability along the same subsequence. We can conclude that $d_{BL}(N_{\mathbb{F}_n}^*, N(0, Q')) \rightarrow 0$ in probability along the same subsequence by the extended continuous mapping theorem (van der Vaart and Wellner, 1996, Theorem 1.11.1).

The argument applies to every subsequence \mathbb{N}' of the stated form. The claim in the first paragraph of the proof thus follows. \blacksquare

B.7. Donsker Theorems for Function Classes that depend on n . Let $(W_i)_{i=1}^\infty$ be a sequence of i.i.d. copies of the random element W taking values in the measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$, whose law is determined by the probability measure P , and let $w \mapsto f_{n,t}(w)$ be measurable functions $f_{n,t} : \mathcal{W} \rightarrow \mathbb{R}$ indexed by $n \in \mathbb{N}$ and a fixed, totally bounded semi-metric space (T, d_T) . Consider the stochastic process

$$(\mathbb{G}_n f_{n,t})_{t \in T} := \left\{ n^{-1/2} \sum_{i=1}^n (f_{n,t}(W_i) - P f_{n,t}) \right\}_{t \in T}.$$

This empirical process is indexed by a class of functions $\mathcal{F}_n = \{f_{n,t} : t \in T\}$ with a measurable envelope function F_n . It is important to note here that the dependence on n allows us to have *the class itself* be possibly dependent on the law P_n .

Lemma B.2 (Donsker Theorem for Classes Changing with n). *Work with the set-up above. Suppose that for some fixed constant $q > 2$ and every sequence $\delta_n \searrow 0$:*

$$\|F_n\|_{P_n, q} = O(1), \quad \sup_{d_T(s,t) \leq \delta_n} \|f_{n,s} - f_{n,t}\|_{P_n, 2} \rightarrow 0, \quad \int_0^{\delta_n} \sup_Q \sqrt{\log N(\epsilon \|F_n\|_{Q, 2}, \mathcal{F}_n, \|\cdot\|_{Q, 2})} d\epsilon \rightarrow 0.$$

(a) *Then the empirical process $(\mathbb{G}_n f_{n,t})_{t \in T}$ is asymptotically tight in $\ell^\infty(T)$ i.e. stochastically equicontinuous. (b) For any subsequence such that the covariance function $P_n f_{n,s} f_{n,t} - P_n f_{n,s} P_n f_{n,t}$ converges pointwise on $T \times T$, $(\mathbb{G}_n f_{n,t})_{t \in T}$ converges in $\ell^\infty(T)$ to a Gaussian process with covariance function given by the limit of the covariance function along that subsequence.*

Proof. The use of Theorem 2.11.1 in van der Vaart and Wellner (1996), which does allow for the probability space to depend on n , allows us to establish claim (a), by repeating the proof (verbatim) of Theorem 2.11.22 in van der Vaart and Wellner (1996, p. 220-221), except that the probability law is allowed to depend on n . (For the sake of completeness, the Supplementary Appendix, provides the complete proof). The proof of claim (b) follows by a standard argument from the stochastic equicontinuity established in claim (a) and finite-dimensional convergence along the indicated subsequences. \blacksquare

B.8. Proof of Theorems B.3 and B.4. The proof consists of two parts, each proving the corresponding theorem.

Part 1. We can split \mathbb{N} into subsequences $\{\mathbb{N}'\}$ along each of which $Z_{n,P_n} \rightsquigarrow Z' \in \mathbb{D}_0$ in \mathbb{D} , $\rho_{P_n} \rightarrow \rho'$ in \mathbb{D}_ρ ($n \in \mathbb{N}'$), where Z' and ρ' can possibly depend on \mathbb{N}' . It suffices to verify that for each \mathbb{N}' :

$$r_n(\phi(\widehat{\rho}_{n,P_n}) - \phi(\rho_{P_n})) \rightsquigarrow \phi'_{\rho'}(Z') \quad (n \in \mathbb{N}') \quad (\text{B.3})$$

$$r_n(\phi(\widehat{\rho}_{n,P_n}) - \phi(\rho_{P_n})) - \phi'_{\rho_{P_n}}(r_n(\widehat{\rho}_{n,P_n} - \rho_{P_n})) \rightsquigarrow 0 \quad (n \in \mathbb{N}'), \quad (\text{B.4})$$

$$r_n(\phi(\widehat{\rho}_{n,P_n}) - \phi(\rho_{P_n})) - \phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n} - \rho_{P_n})) \rightsquigarrow 0 \quad (n \in \mathbb{N}'), \quad (\text{B.5})$$

where the last two claims hold provided that $(\rho, h) \mapsto \phi'_\rho(h)$ is defined and continuous on the whole of $\mathbb{D}_\rho \times \mathbb{D}$. The claim (B.5) is not needed in Part 1, but we need it for the Part 2.

The map $g_n(h) = r_n(\phi(\rho_{P_n} + r_n^{-1}h) - \phi(\rho_{P_n}))$, from $\mathbb{D}_n = \{h \in \mathbb{D} : \rho_{P_n} + r_n^{-1}h \in \mathbb{D}_\phi\}$ to \mathbb{E} , satisfies $g_n(h_n) \rightarrow \phi'_{\rho'}(h)$ for every subsequence $h_n \rightarrow h \in \mathbb{D}_0$ (with $n \in \mathbb{N}'$). Application of the extended continuous mapping theorem (van der Vaart and Wellner, 1996, Theorem 1.11.1) yields (B.3).

Similarly, the map $m_n(h) = r_n(\phi(\rho_{P_n} + r_n^{-1}h) - \phi(\rho_{P_n})) - \phi'_{\rho_{P_n}}(h)$, from $\mathbb{D}_n = \{h \in \mathbb{D} : \rho_{P_n} + r_n^{-1}h \in \mathbb{D}_\phi\}$ to \mathbb{E} , satisfies $m_n(h_n) \rightarrow \phi'_{\rho'}(h) - \phi'_{\rho'}(h) = 0$ for every subsequence $h_n \rightarrow h \in \mathbb{D}_0$ (with $n \in \mathbb{N}'$). Application of the extended continuous mapping theorem (van der Vaart and Wellner, 1996, Theorem 1.11.1) yields (B.4). The proof of (B.5) is completely analogous and is omitted.

To establish relative compactness, work with each \mathbb{N}' . Then $\phi'_{\rho_{P_n}}(h)$ mapping \mathbb{D}_0 to \mathbb{E} satisfies $\phi'_{\rho_{P_n}}(h_n) \rightarrow \phi'_{\rho'}(h)$ for every subsequence $h_n \rightarrow h \in \mathbb{D}_0$ (with $n \in \mathbb{N}'$). Application of the extended continuous mapping theorem (van der Vaart and Wellner, 1996, Theorem 1.11.1) yields that $\phi'_{\rho_{P_n}}(Z_P) \rightsquigarrow \phi'_{\rho'}(Z')$.

Part 2. We can split \mathbb{N} into subsequences $\{\mathbb{N}'\}$ as above. Along each \mathbb{N}' ,

$$r_n(\widehat{\rho}_{n,P_n}^* - \rho_{P_n}) \rightsquigarrow Z'' \in \mathbb{D}_0 \text{ in } \mathbb{D}, \quad r_n(\widehat{\rho}_{n,P_n} - \rho_{P_n}) \rightsquigarrow Z' \in \mathbb{D}_0 \text{ in } \mathbb{D}, \quad \rho_{P_n} \rightarrow \rho' \text{ in } \mathbb{D}_\rho \quad (n \in \mathbb{N}'),$$

where Z'' is a separable process in \mathbb{D}_0 (which is given by Z' plus its independent copy \bar{Z}'). Indeed, note that $r_n(\widehat{\rho}_{n,P_n}^* - \rho_{P_n}) = Z_{n,P_n}^* + Z_{n,P_n}$, and (Z_{n,P_n}^*, Z_{n,P_n}) converge weakly unconditionally to (\bar{Z}', Z') by a standard argument.

Given each \mathbb{N}' the proof is similar to the proof of Theorem 3.9.15 of van der Vaart and Wellner (1996). We can assume without loss of generality that the derivative $\phi'_{\rho'} : \mathbb{D} \rightarrow \mathbb{E}$ is defined and continuous on the whole of \mathbb{D} . Otherwise, if $\phi'_{\rho'}$ is defined and continuous only on \mathbb{D}_0 , we can extend it to \mathbb{D} by a Hahn-Banach extension such that $C = \|\phi'_{\rho'}\|_{\mathbb{D}_0 \rightarrow \mathbb{E}} = \|\phi'_{\rho'}\|_{\mathbb{D} \rightarrow \mathbb{E}} < \infty$; see van der Vaart

and Wellner (1996, p. 380) for details. For each \mathbb{N}' , by claim (B.5), applied to $\widehat{\rho}_{n,P_n}$ and to $\widehat{\rho}_{n,P_n}^*$ replacing $\widehat{\rho}_{n,P_n}$,

$$\begin{aligned} r_n(\phi(\widehat{\rho}_{n,P_n}) - \phi(\rho_{P_n})) &= \phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n} - \rho_{P_n})) + o_{P_n}^*(1), \\ r_n(\phi(\widehat{\rho}_{n,P_n}^*) - \phi(\rho_{P_n})) &= \phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n}^* - \rho_{P_n})) + o_{P_n}^*(1). \end{aligned}$$

Subtracting these equations conclude that for each $\varepsilon > 0$

$$\mathbb{E}_{P_n} \mathbf{1} \left(\left\| r_n(\phi(\widehat{\rho}_{n,P_n}^*) - \phi(\widehat{\rho}_{n,P_n})) - \phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n}^* - \widehat{\rho}_{n,P_n})) \right\|_{\mathbb{E}}^* > \varepsilon \right) \rightarrow 0 \quad (n \in \mathbb{N}'). \quad (\text{B.6})$$

For every $h \in \text{BL}_1(\mathbb{E})$, the function $h \circ \phi'_{\rho'}$ is contained in $\text{BL}_C(\mathbb{D})$. Moreover, $r_n(\widehat{\rho}_{n,P}^* - \widehat{\rho}_{n,P}) \rightsquigarrow_B Z_P$ in \mathbb{D} uniformly in $P \in \mathcal{P}_n$ implies $r_n(\widehat{\rho}_{n,P}^* - \widehat{\rho}_{n,P}) \rightsquigarrow_B Z'$ along the subsequence $n \in \mathbb{N}'$. These two facts imply that

$$\sup_{h \in \text{BL}_1(\mathbb{E})} \left| \mathbb{E}_{B_n} h \left(\phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n}^* - \widehat{\rho}_{n,P_n})) \right) - \mathbb{E} h(\phi_{\rho'}(Z')) \right| = o_{P_n}^*(1) \quad (n \in \mathbb{N}').$$

Next for each $\varepsilon > 0$ and along $n \in \mathbb{N}'$

$$\begin{aligned} & \sup_{h \in \text{BL}_1(\mathbb{E})} \left| \mathbb{E}_{B_n} h \left(r_n(\phi(\widehat{\rho}_{n,P_n}^*) - \phi(\widehat{\rho}_{n,P_n})) \right) - \mathbb{E}_{B_n} h \left(\phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n}^* - \widehat{\rho}_{n,P_n})) \right) \right| \\ & \leq \varepsilon + 2\mathbb{E}_{B_n} \mathbf{1} \left(\left\| r_n(\phi(\widehat{\rho}_{n,P_n}^*) - \phi(\widehat{\rho}_{n,P_n})) - \phi'_{\rho'}(r_n(\widehat{\rho}_{n,P_n}^* - \widehat{\rho}_{n,P_n})) \right\|_{\mathbb{E}}^* > \varepsilon \right) = o_{P_n}(1), \end{aligned}$$

where the $o_{P_n}(1)$ conclusion follows by the Markov inequality and by (B.6). Conclude that

$$\sup_{h \in \text{BL}_1(\mathbb{E})} \left| \mathbb{E}_{B_n} h \left(r_n(\phi(\widehat{\rho}_{n,P_n}^*) - \phi(\widehat{\rho}_{n,P_n})) \right) - \mathbb{E} h(\phi_{\rho'}(Z')) \right| = o_{P_n}^*(1) \quad (n \in \mathbb{N}'). \blacksquare$$

APPENDIX C. KEY TOOLS II: PROBABILISTIC INEQUALITIES

Let $(W_i)_{i=1}^n$ be a sequence of i.i.d. copies of random element W taking values in the measure space $(\mathcal{W}, \mathcal{A}_{\mathcal{W}})$ according to probability law P . Let \mathcal{F} be a set of suitably measurable functions $f : \mathcal{W} \mapsto \mathbb{R}$, equipped with a measurable envelope $F : \mathcal{W} \mapsto \mathbb{R}$.

The following maximal inequality is due to Chernozhukov et al. (2012).

Lemma C.1 (A Maximal Inequality). *Work with the setup above. Suppose that $F \geq \sup_{f \in \mathcal{F}} |f|$ is a measurable envelope with $\|F\|_{P,q} < \infty$ for some $q \geq 2$. Let $M = \max_{i \leq n} F(W_i)$ and $\sigma^2 > 0$ be any positive constant such that $\sup_{f \in \mathcal{F}} \|f\|_{P,2}^2 \leq \sigma^2 \leq \|F\|_{P,2}^2$. Suppose that there exist constants $a \geq e$ and $v \geq 1$ such that $\log \sup_Q N(\epsilon \|F\|_{Q,2}, \mathcal{F}, \|\cdot\|_{Q,2}) \leq v(\log a + \log(1/\epsilon))$, $0 < \epsilon \leq 1$. Then*

$$\mathbb{E}_P [\|\mathbb{G}_n\|_{\mathcal{F}}] \leq K \left(\sqrt{v\sigma^2 \log \left(\frac{a\|F\|_{P,2}}{\sigma} \right)} + \frac{v\|M\|_{P,2}}{\sqrt{n}} \log \left(\frac{a\|F\|_{P,2}}{\sigma} \right) \right),$$

where K is an absolute constant. Moreover, for every $t \geq 1$, with probability $> 1 - t^{-q/2}$,

$$\|\mathbb{G}_n\|_{\mathcal{F}} \leq (1 + \alpha) \mathbb{E}_P [\|\mathbb{G}_n\|_{\mathcal{F}}] + K(q) \left[(\sigma + n^{-1/2} \|M\|_{P,q}) \sqrt{t} + \alpha^{-1} n^{-1/2} \|M\|_{P,2} t \right], \quad \forall \alpha > 0,$$

where $K(q) > 0$ is a constant depending only on q . In particular, setting $a \geq n$ and $t = \log n$, with probability $> 1 - c(\log n)^{-1}$,

$$\|\mathbb{G}_n\|_{\mathcal{F}} \leq K(q, c) \left(\sigma \sqrt{v \log \left(\frac{a \|F\|_{P,2}}{\sigma} \right)} + \frac{v \|M\|_{P_{P,q}}}{\sqrt{n}} \log \left(\frac{a \|F\|_{P,2}}{\sigma} \right) \right), \quad (\text{C.1})$$

where $\|M\|_{P_{P,q}} \leq n^{1/q} \|F\|_{P,q}$ and $K(q, c) > 0$ is a constant depending only on q and c .

APPENDIX D. PROOF FOR SECTION 4

These results follow from the application of results given in Section 5. The details are given in the Supplementary Appendix.

APPENDIX E. PROOFS FOR SECTION 5

E.1. Proof of Theorem 5.1. In the proof $a \lesssim b$ means that $a \leq Ab$, where the constant A depends on the constants in Assumptions 5.1–5.3, but not on n once $n \geq n_0$, and not on $P \in \mathcal{P}_n$. Since the argument is asymptotic, we can assume that $n \geq n_0$ in what follows. In order to establish the result uniformly in $P \in \mathcal{P}_n$, it suffices to establish the result under the probability measure induced by any sequence $P = P_n \in \mathcal{P}_n$. In the proof we shall use P , suppressing the dependency of P_n on the sample size n . Also, let

$$B(W) := \max_{j \in [d_\theta], k \in [d_\theta + d_t]} \sup_{\nu \in \Theta_u \times T_u(Z_u), u \in \mathcal{U}} \left| \partial_{\nu_k} \mathbb{E}_P[\psi_{uj}(W_u, \nu) \mid Z_u] \right|, \quad (\text{E.1})$$

Step 1. (A Preliminary Rate Result). In this step we claim that wp $1 - o(1)$, $\sup_{u \in \mathcal{U}} \|\widehat{\theta}_u - \theta_u\| \lesssim \tau_n$. By definition

$$\|\mathbb{E}_n \psi_u(W_u, \widehat{\theta}_u, \widehat{h}_u(Z_u))\| \leq \inf_{\theta \in \Theta_u} \|\mathbb{E}_n \psi_u(W_u, \theta, \widehat{h}_u(Z_u))\| + \epsilon_n \text{ for each } u \in \mathcal{U},$$

which implies via triangle inequality that uniformly in $u \in \mathcal{U}$ with probability $1 - o(1)$

$$\left\| P[\psi_u(W_u, \widehat{\theta}_u, h_u(Z_u))] \right\| \leq \epsilon_n + 2I_1 + 2I_2 \lesssim \tau_n, \quad (\text{E.2})$$

for I_1 and I_2 defined in Step 2 below. The \lesssim bound in (E.2) follows from Step 2 and from the assumption $\epsilon_n = o(n^{-1/2})$. Since by Assumption 5.1(iv), $2^{-1}(\|J_u(\widehat{\theta}_u - \theta_u)\| \wedge c_0)$ does not exceed the left side of (E.2) and $\inf_{u \in \mathcal{U}} \text{mineig}(J'_u J_u)$ is bounded away from zero uniformly in n , we conclude that $\sup_{u \in \mathcal{U}} \|\widehat{\theta}_u - \theta_u\| \lesssim (\inf_{u \in \mathcal{U}} \text{mineig}(J'_u J_u))^{-1/2} \tau_n \lesssim \tau_n$.

Step 2. (Define and bound I_1 and I_2) We claim that with probability $1 - o(1)$:

$$\begin{aligned} I_1 &:= \sup_{\theta \in \Theta_u, u \in \mathcal{U}} \left\| \mathbb{E}_n \psi_u(W_u, \theta, \widehat{h}_u(Z_u)) - \mathbb{E}_n \psi_u(W_u, \theta, h_u(Z_u)) \right\| \lesssim \tau_n, \\ I_2 &:= \sup_{\theta \in \Theta_u, u \in \mathcal{U}} \left\| \mathbb{E}_n \psi_u(W_u, \theta, h_u(Z_u)) - P \psi_u(W_u, \theta, h_u(Z_u)) \right\| \lesssim \tau_n. \end{aligned}$$

To establish this, we can bound $I_1 \leq 2I_{1a} + I_{1b}$ and $I_2 \leq I_{1a}$, where with probability $1 - o(1)$,

$$\begin{aligned} I_{1a} &:= \sup_{\theta \in \Theta_u, u \in \mathcal{U}, h \in \mathcal{H}_{un} \cup \{h_u\}} \left\| \mathbb{E}_n \psi_u(W_u, \theta, h(Z_u)) - P \psi_u(W_u, \theta, h(Z_u)) \right\| \lesssim \tau_n, \\ I_{1b} &:= \sup_{\theta \in \Theta_u, u \in \mathcal{U}, h \in \mathcal{H}_{un} \cup \{h_u\}} \left\| P \psi_u(W_u, \theta, h(Z_u)) - P \psi_u(W_u, \theta, h_u(Z_u)) \right\| \lesssim \tau_n. \end{aligned}$$

These bounds in turn hold by the following arguments. In order to bound I_{1b} we employ Taylor's expansion and the triangle inequality. For $\bar{h}(Z, u, j, \theta)$ denoting a point on a line connecting vectors $h(Z_u)$ and $h_u(Z_u)$, and t_m denoting the m th element of the vector t ,

$$\begin{aligned} I_{1b} &\leq \sum_{j=1}^{d_\theta} \sum_{m=1}^{d_t} \sup_{\theta \in \Theta_u, u \in \mathcal{U}, h \in \mathcal{H}_{un}} \left| P \left[\partial_{t_m} P \left[\psi_{uj}(W_u, \theta, \bar{h}(Z, u, j, \theta)) \mid Z_u \right] (h_m(Z_u) - h_{um}(Z_u)) \right] \right| \\ &\leq d_\theta d_t \|B\|_{P,2} \max_{u \in \mathcal{U}, h \in \mathcal{H}_{un}, m \in [d_t]} \|h_m - h_{um}\|_{P,2}, \end{aligned}$$

where the last inequality holds by the definition of $B(W)$ given earlier and Hölder's inequality. By Assumption 5.2(ii)(c), $\|B\|_{P,2} \leq C$, and by Assumption 5.3, $\sup_{u \in \mathcal{U}, h \in \mathcal{H}_{un}, m \in [d_t]} \|h_m - h_{um}\|_{P,2} \lesssim \tau_n$, hence we conclude that $I_{1b} \lesssim \tau_n$ since d_θ and d_t are fixed.

In order to bound I_{1a} we employ the maximal inequality of Lemma C.1 to the class

$$\mathcal{F}_1 = \{\psi_{uj}(W_u, \theta, h(Z_u)) : j \in [d_\theta], u \in \mathcal{U}, \theta \in \Theta_u, h \in \mathcal{H}_{un} \cup \{h_u\}\},$$

defined in Assumption 5.3 and equipped with an envelope $F_1 \leq F_0$, to conclude that with probability $1 - o(1)$,

$$I_{1a} \lesssim n^{-1/2} \left(\sqrt{s_n \log(a_n)} + n^{-1/2} s_n n^{\frac{1}{q}} \log(a_n) \right) \lesssim \tau_n.$$

Here we use that $\log \sup_Q N(\epsilon \|F_1\|_{Q,2}, \mathcal{F}_1, \|\cdot\|_{Q,2}) \leq s_n \log(a_n/\epsilon) \vee 0$ by Assumption 5.3; $\|F_0\|_{P,q} \leq C$ and $\sup_{f \in \mathcal{F}_1} \|f\|_{P,2}^2 \leq \sigma^2 \leq \|F_0\|_{P,2}^2$ for $c \leq \sigma \leq C$ by Assumption 5.2(i); $a_n \geq n$ and $s_n \geq 1$ by Assumption 5.3.

Step 3. (Linearization) By definition

$$\sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \hat{\theta}_u, \hat{h}_u(Z_u))\| \leq \inf_{\theta \in \Theta_u} \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \theta, \hat{h}_u(Z_u))\| + \epsilon_n n^{1/2}.$$

Application of Taylor's theorem give that for all $u \in \mathcal{U}$

$$\begin{aligned} \sqrt{n} \mathbb{E}_n \psi_u(W_u, \hat{\theta}_u, \hat{h}_u(Z_u)) &= \sqrt{n} \mathbb{E}_n \psi_u(W_u, \theta_u, h_u(Z_u)) \\ &\quad + J_u \sqrt{n} (\hat{\theta}_u - \theta_u) + D_{u,0} (\hat{h}_u - h_u) + II_1(u) + II_2(u), \end{aligned}$$

where the terms II_1 and II_2 are defined in Step 4. Then by the triangle inequality for all $u \in \mathcal{U}$ and Steps 4 and 5 we have

$$\begin{aligned} &\left\| \sqrt{n} \mathbb{E}_n \psi_u(W_u, \theta_u, h_u(Z_u)) + J_u \sqrt{n} (\hat{\theta}_u - \theta_u) + D_{u,0} (\hat{h}_u - h_u) \right\| \\ &\leq \epsilon_n \sqrt{n} + \sup_{u \in \mathcal{U}} \left(\inf_{\theta \in \Theta_u} \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \theta, \hat{h}_u(Z_u))\| + \|II_1(u)\| + \|II_2(u)\| \right) = o_P(1), \end{aligned}$$

where the $o_P(1)$ bound follows from Step 4, $\epsilon_n \sqrt{n} = o(1)$ by assumption, and Step 5.

Moreover, by the orthogonality condition:

$$D_{u,0}(\widehat{h}_u - h_u) := \left(\sum_{m=1}^{d_t} \sqrt{n} P \left[\partial_{t_m} P[\psi_{uj}(W_u, \theta_u, h_u(Z_u)) | Z_u] (\widehat{h}_m(Z_u) - h_{um}(Z_u)) \right] \right)_{j=1}^{d_\theta} = 0.$$

Conclude using Assumption 5.1(iv) that

$$\sup_{u \in \mathcal{U}} \left\| J_u^{-1} \sqrt{n} \mathbb{E}_n \psi_u(W_u, \theta_u, h_u(Z_u)) + \sqrt{n} (\widehat{\theta}_u - \theta_u) \right\| \leq o_P(1) \sup_{u \in \mathcal{U}} (\text{mineg}(J_u' J_u)^{-1/2}) = o_P(1).$$

Furthermore, the empirical process $(-\sqrt{n} \mathbb{E}_n J_u^{-1} \psi_u(W_u, \theta_u, h_u(Z_u)))_{u \in \mathcal{U}}$ is equivalent to an empirical process \mathbb{G}_n indexed by $\mathcal{F}_P := \{\bar{\psi}_{uj} : j \in [d_\theta], u \in \mathcal{U}\}$, where $\bar{\psi}_{uj}$ is the j -th element of $-J_u^{-1} \psi_u(W_u, \theta_u, h_u(Z_u))$ and we make explicit the dependence of \mathcal{F}_P on P . Let $\mathcal{M} = \{M_{ujk} : j, k \in [d_\theta], u \in \mathcal{U}\}$, where M_{ujk} is the (j, k) element of the matrix J_u^{-1} . \mathcal{M} is a class of uniformly Hölder continuous functions on $(\mathcal{U}, d_{\mathcal{U}})$ with a uniform covering entropy bounded by $C \log(e/\epsilon) \vee 0$ and equipped with a constant envelope C , given the stated assumptions. This result follows from the fact that by Assumption 5.2(ii)(b)

$$\begin{aligned} \max_{j, k \in [d_\theta]} |M_{ujk} - M_{\bar{u}jk}| &\leq \|J_u^{-1} - J_{\bar{u}}^{-1}\| = \|J_u^{-1}(J_u - J_{\bar{u}})J_{\bar{u}}^{-1}\| \\ &\leq \|J_u - J_{\bar{u}}\| \sup_{\bar{u} \in \mathcal{U}} \|J_{\bar{u}}^{-1}\|^2 \lesssim \|u - \bar{u}\|^{\alpha_2}, \end{aligned} \quad (\text{E.3})$$

and the constant envelope follows by Assumption 5.1(iv). Since \mathcal{F}_P is generated as a finite sum of products of the elements of \mathcal{M} and the class \mathcal{F}_0 defined in Assumption 5.2, the properties of \mathcal{M} and the conditions on \mathcal{F}_0 in Assumption 5.2(ii) imply that \mathcal{F}_P has a uniformly well-behaved uniform covering entropy by Lemma L.1, namely

$$\sup_{P \in \mathcal{P} = \cup_{n \geq n_0} \mathcal{P}_n} \log \sup_Q N(\epsilon \|CF_0\|_{Q,2}, \mathcal{F}_P, \|\cdot\|_{Q,2}) \lesssim \log(e/\epsilon) \vee 0,$$

where $F_P = CF_0$ is an envelope for \mathcal{F}_P since $\sup_{f \in \mathcal{F}_P} |f| \lesssim \sup_{u \in \mathcal{U}} \|J_u^{-1}\| \sup_{f \in \mathcal{F}_0} |f| \leq CF_0$ by Assumption 5.2(i). The class \mathcal{F}_P is therefore Donsker uniformly in P because $\sup_{P \in \mathcal{P}} \|F_P\|_{P,q} \leq C \sup_{P \in \mathcal{P}} \|F_0\|_{P,q}$ is bounded by Assumption 5.2(ii), and $\sup_{P \in \mathcal{P}} \|\bar{\psi}_u - \bar{\psi}_{\bar{u}}\|_{P,2} \rightarrow 0$ as $d_{\mathcal{U}}(u, \bar{u}) \rightarrow 0$ by Assumption 5.2(b) and (E.3). Application of Theorem B.1 gives the results of the theorem.

Step 4. (Define and Bound II_1 and II_2). Let $II_1(u) := (II_{1j}(u))_{j=1}^{d_\theta}$ and $II_2(u) = (II_{2j}(u))_{j=1}^{d_\theta}$, where

$$\begin{aligned} II_{1j}(u) &:= \sum_{r,k=1}^{d_\nu} \sqrt{n} P [\partial_{\nu_k} \partial_{\nu_r} P[\psi_{uj}(W_u, \bar{\nu}_u(Z_u, j)) | Z_u] \{\widehat{\nu}_{ur}(Z_u) - \nu_{ur}(Z_u)\} \{\widehat{\nu}_{uk}(Z_u) - \nu_{uk}(Z_u)\}], \\ II_{2j}(u) &:= \mathbb{G}_n(\psi_{uj}(W_u, \widehat{\theta}_u, \widehat{h}_u(Z_u)) - \psi_{uj}(W_u, \theta_u, h_u(Z_u))), \end{aligned}$$

where $\nu_u(Z_u) := (\nu_{uk}(Z_u))_{k=1}^{d_\nu} := (\theta_u', h_u(Z_u)')'$, $\widehat{\nu}_u(Z_u) := (\widehat{\nu}_{uk}(Z_u))_{k=1}^{d_\nu} := (\widehat{\theta}_u', \widehat{h}_u(Z_u)')'$, $d_\nu = d_\theta + d_t$, and $\bar{\nu}_u(Z_u, j)$ is a vector on the line connecting $\nu_u(Z_u)$ and $\widehat{\nu}_u(Z_u)$.

First, by Assumptions 5.2(ii)(d) and 5.3, the claim of Step 1, and the Hölder inequality,

$$\begin{aligned} \max_{j \in [d_\theta]} \sup_{u \in \mathcal{U}} |II_{1j}(u)| &\leq \sup_{u \in \mathcal{U}} \sum_{r,k=1}^{d_\nu} \sqrt{n} P [C |\widehat{\nu}_{ur}(Z_u) - \nu_{ur}(Z_u)| |\widehat{\nu}_{uk}(Z_u) - \nu_{uk}(Z_u)|] \\ &\leq C \sqrt{n} d_\nu^2 \max_{k \in [d_\nu]} \sup_{u \in \mathcal{U}} \|\widehat{\nu}_{uk} - \nu_{uk}\|_{P,2}^2 \lesssim_P \sqrt{n} \tau_n^2 = o(1). \end{aligned}$$

Second, we have that with probability $1 - o(1)$, $\max_{j \in [d_\theta]} \sup_{u \in \mathcal{U}} |II_{2j}(u)| \lesssim \sup_{f \in \mathcal{F}_2} |\mathbb{G}_n(f)|$, where, for $\Theta_{un} := \{\theta \in \Theta_u : \|\theta - \theta_u\| \leq C\tau_n\}$,

$$\mathcal{F}_2 = \left\{ \psi_{uj}(W_u, \theta, h(Z_u)) - \psi_{uj}(W_u, \theta_u, h_u(Z_u)) : j \in [d_\theta], u \in \mathcal{U}, h \in \mathcal{H}_{un}, \theta \in \Theta_{un} \right\}.$$

Application of Lemma C.1 with an envelope $F_2 \lesssim F_0$ gives that with probability $1 - o(1)$

$$\sup_{f \in \mathcal{F}_2} |\mathbb{G}_n(f)| \lesssim \tau_n^{\alpha/2} \sqrt{s_n \log(a_n)} + n^{-1/2} s_n n^{\frac{1}{q}} \log(a_n), \quad (\text{E.4})$$

since $\sup_{f \in \mathcal{F}_2} |f| \leq 2 \sup_{f \in \mathcal{F}_1} |f| \leq 2F_0$ by Assumption 5.3; $\|F_0\|_{P,q} \leq C$ by Assumption 5.2(i); $\log \sup_Q N(\epsilon \|F_2\|_{Q,2}, \mathcal{F}_2, \|\cdot\|_{Q,2}) \lesssim (s_n \log a_n + s_n \log(a_n/\epsilon)) \vee 0$ by Lemma L.1 because $\mathcal{F}_2 = \mathcal{F}_1 - \mathcal{F}_0$ for the \mathcal{F}_0 and \mathcal{F}_1 defined in Assumptions 5.2(i) and 5.3; and σ can be chosen so that $\sup_{f \in \mathcal{F}_2} \|f\|_{P,2} \leq \sigma \lesssim \tau_n^{\alpha/2}$. Indeed,

$$\begin{aligned} \sup_{f \in \mathcal{F}_2} \|f\|_{P,2}^2 &\leq \sup_{j \in [d_\theta], u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} P \left(P[(\psi_{uj}(W_u, \nu(Z_u)) - \psi_{uj}(W_u, \nu_u(Z_u)))^2 | Z_u] \right) \\ &\leq \sup_{u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} P(C \|\nu(Z_u) - \nu_u(Z_u)\|^\alpha) \\ &= \sup_{u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} C \|\nu - \nu_u\|_{P,\alpha}^\alpha \leq \sup_{u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} C \|\nu - \nu_u\|_{P,2}^\alpha \lesssim \tau_n^\alpha, \end{aligned}$$

where the first inequality follows by the law of iterated expectations; the second inequality follows by Assumption 5.2(ii)(a); and the last inequality follows from $\alpha \in [1, 2]$ by Assumption 5.2, the monotonicity of the norm $\|\cdot\|_{P,\alpha}$ in $\alpha \in [1, \infty]$, and Assumption 5.3.

Conclude using the growth conditions of Assumption 5.3 that with probability $1 - o(1)$

$$\max_{j \in [d_\theta]} \sup_{u \in \mathcal{U}} |II_{2j}(u)| \lesssim \tau_n^{\alpha/2} \sqrt{s_n \log(a_n)} + n^{-1/2} s_n n^{\frac{1}{q}} \log(a_n) = o(1). \quad (\text{E.5})$$

Step 5. In this step we show that $\sup_{u \in \mathcal{U}} \inf_{\theta \in \Theta_u} \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \theta, \widehat{h}_u(Z_u))\| = o_P(1)$. We have that with probability $1 - o(1)$

$$\inf_{\theta \in \Theta_u} \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \theta, \widehat{h}_u(Z_u))\| \leq \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \bar{\theta}_u, \widehat{h}_u(Z_u))\|,$$

where $\bar{\theta}_u = \theta_u - J_u^{-1} \mathbb{E}_n \psi_u(W_u, \theta_u, h_u(Z_u))$, since $\bar{\theta}_u \in \Theta_u$ for all $u \in \mathcal{U}$ with probability $1 - o(1)$, and, in fact, $\sup_{u \in \mathcal{U}} \|\bar{\theta}_u - \theta_u\| = O_P(1/\sqrt{n})$ by the last paragraph of Step 3.

Then, arguing similarly to Step 3 and 4, we can conclude that uniformly in $u \in \mathcal{U}$:

$$\sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \bar{\theta}_u, \widehat{h}_u(Z_u))\| \leq \sqrt{n} \|\mathbb{E}_n \psi_u(W_u, \theta_u, h_u(Z_u)) + J_u(\bar{\theta}_u - \theta_u) + D_{u,0}(\widehat{h}_u - h_u)\| + o_P(1)$$

where the first term on the right side is zero by definition of $\bar{\theta}_u$ and $D_{u,0}(\widehat{h}_u - h_u) = 0$. \blacksquare

E.2. Proof of Theorem 5.2. STEP 0. In the proof $a \lesssim b$ means that $a \leq Ab$, where the constant A depends on the constants in Assumptions 5.1–5.3, but not on n once $n \geq n_0$, and not on $P \in \mathcal{P}_n$. In Step 1, we consider a sequence P_n in \mathcal{P}_n , but for simplicity, we write $P = P_n$ throughout the proof, suppressing the index n . Since the argument is asymptotic, we can assume that $n \geq n_0$ in what follows.

Let \mathbb{P}_n denote the measure that puts mass n^{-1} at the points (ξ_i, W_i) for $i = 1, \dots, n$. Let \mathbb{E}_n denote the expectation with respect to this measure, so that $\mathbb{E}_n f = n^{-1} \sum_{i=1}^n f(\xi_i, W_i)$, and \mathbb{G}_n denote the corresponding empirical process $\sqrt{n}(\mathbb{E}_n - P)$, i.e.

$$\mathbb{G}_n f = \sqrt{n}(\mathbb{E}_n f - P f) = n^{-1/2} \sum_{i=1}^n \left(f(\xi_i, W_i) - \int f(s, w) dP_\xi(s) dP(w) \right).$$

Recall that we define the bootstrap draw as:

$$Z_{n,P}^* := \sqrt{n}(\hat{\theta}^* - \hat{\theta}) = \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \hat{\psi}_u(W_i) \right)_{u \in \mathcal{U}} = \left(\mathbb{G}_n \xi \hat{\psi}_u \right)_{u \in \mathcal{U}},$$

where $\hat{\psi}_u(W) = -\hat{J}_u^{-1} \psi_u(W_u, \hat{\theta}_u, \hat{h}_u(Z_u))$.

STEP 1. (Linearization) In this step we establish that

$$\zeta_{n,P}^* := Z_{n,P}^* - G_{n,P}^* = o_P(1) \quad \text{in } \mathbb{D} = \ell^\infty(\mathcal{U})^{d_\theta}, \quad (\text{E.6})$$

where $G_{n,P}^* := (\mathbb{G}_n \xi \bar{\psi}_u)_{u \in \mathcal{U}}$, and $\bar{\psi}_u(W) = -J_u^{-1} \psi_u(W_u, \theta_u, h_u(Z_u))$.

With probability $1 - \delta_n$, $\hat{h}_u \in \mathcal{H}_{un}$, $\hat{\theta}_u \in \Theta_{un} = \{\theta \in \Theta_u : \|\theta - \theta_u\| \leq C\tau_n\}$, and $J_u \in \mathcal{J}_n$, so that $\|\zeta_{n,P}^*\|_{\mathbb{D}} \lesssim \sup_{f \in \mathcal{F}_3} |\mathbb{G}_n[\xi f]|$, where

$$\mathcal{F}_3 = \left\{ \tilde{\psi}_{uj}(\bar{\theta}_u, \bar{h}_u, \bar{J}_u) - \bar{\psi}_{uj} : j \in [d_\theta], u \in \mathcal{U}, \bar{\theta}_u \in \Theta_{un}, \bar{h}_u \in \mathcal{H}_{un}, \bar{J}_u \in \mathcal{J}_n \right\},$$

where $\tilde{\psi}_{uj}(\bar{\theta}_u, \bar{h}_u, \bar{J}_u)$ is the j -th element of $-\bar{J}_u^{-1} \psi_u(W_u, \bar{\theta}_u, \bar{h}_u(Z_u))$, and $\bar{\psi}_{uj}$ is the j -th element of $-J_u^{-1} \psi_u(W_u, \theta_u, h_u(Z_u))$. By the arguments similar to those employed in the proof of the previous theorem, under Assumption 5.3 and the additional conditions stated in the theorem, \mathcal{F}_3 obeys

$$\log \sup_Q N(\epsilon \|F_3\|_{Q,2}, \mathcal{F}_3, \|\cdot\|_{Q,2}) \lesssim (s_n \log a_n + s_n \log(a_n/\epsilon)) \vee 0,$$

for an envelope $F_3 \lesssim F_0$. By Lemma L.1, multiplication of this class by ξ does not change the entropy bound modulo an absolute constant, namely

$$\log \sup_Q N(\epsilon \|\xi\|_{Q,2} \|F_3\|_{Q,2}, \xi \mathcal{F}_3, \|\cdot\|_{Q,2}) \lesssim (s_n \log a_n + s_n \log(a_n/\epsilon)) \vee 0.$$

Also $\mathbb{E}[\exp(|\xi|)] < \infty$ implies $(\mathbb{E}[\max_{i \leq n} |\xi_i|^2])^{1/2} \lesssim \log n$, so that, using independence of $(\xi_i)_{i=1}^n$ from $(W_i)_{i=1}^n$ and Assumption 5.2(i),

$$\left\| \max_{i \leq n} \xi_i F_0(W_i) \right\|_{P_{P,2}} \leq \left\| \max_{i \leq n} \xi_i \right\|_{P_{P,2}} \left\| \max_{i \leq n} F_0(W_i) \right\|_{P_{P,2}} \lesssim n^{1/q} \log n.$$

Applying Lemma C.1,

$$\sup_{f \in \xi \mathcal{F}_3} |\mathbb{G}_n(f)| = O_P \left(\tau_n^{\alpha/2} \sqrt{s_n \log(a_n)} + \frac{s_n n^{1/q} \log n}{\sqrt{n}} \log(a_n) \right) = o_P(1),$$

for $\sup_{f \in \mathcal{F}_3} \|f\|_{P,2} = \sup_{f \in \mathcal{F}_3} \|f\|_{P,2} \lesssim \sigma_n \lesssim \tau_n^{\alpha/2}$, where the details of calculations are similar to those in the proof of Theorem 5.1. Indeed, with probability $1 - o(\delta_n)$,

$$\begin{aligned} \sup_{f \in \mathcal{F}_3} \|f\|_{P,2}^2 &\lesssim \sup_{u \in \mathcal{U}} \|J_u^{-1}\|^2 \sup_{j \in [d_\theta], u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} P(P[(\psi_{uj}(W_u, \nu(Z_u)) - \psi_{uj}(W_u, \nu_u(Z_u)))^2 | Z_u]) \\ &+ \sup_{u \in \mathcal{U}} \|\bar{J}_u^{-1} - J_u^{-1}\|^2 \sup_{j \in [d_\theta], u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} P(P[\psi_{uj}(W_u, \nu(Z_u))^2 | Z_u]) \\ &\lesssim \sup_{u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} \|\nu - \nu_u\|_{P,\alpha}^\alpha + \tau_n^\alpha \lesssim \sup_{u \in \mathcal{U}, \nu \in \Theta_{un} \times \mathcal{H}_{un}} \|\nu - \nu_u\|_{P,2}^\alpha + \tau_n^\alpha \lesssim \tau_n^\alpha, \end{aligned}$$

where the first inequality follows from the triangle inequality and the law of iterated expectations; the second inequality follows by Assumption 5.2(ii)(a), Assumption 5.2(i), and $\sup_{u \in \mathcal{U}} \|\bar{J}_u^{-1} - J_u^{-1}\|^2 \lesssim \tau_n^\alpha$ by the assumptions of the theorem and the continuous mapping theorem; the third inequality follows from $\alpha \in [1, 2]$ by Assumption 5.2, the monotonicity of the norm $\|\cdot\|_{P,\alpha}$ in $\alpha \in [1, \infty]$, and Assumption 5.3; and the last inequality follows from $\|\nu - \nu_u\|_{P,2} \lesssim \tau_n$ by the definition of Θ_{un} and \mathcal{H}_{un} . The claim of Step 1 follows.

STEP 2. Here we are claiming that $Z_{n,P}^* \rightsquigarrow_B Z_P$ in $\mathbb{D} = \ell^\infty(\mathcal{U})^{d_\theta}$, under any sequence $P = P_n \in \mathcal{P}_n$, where $Z_P = (\mathbb{G}_P \bar{\psi}_u)_{u \in \mathcal{U}}$. By the triangle inequality and Step 1,

$$\sup_{h \in \text{BL}_1(\mathbb{D})} \left| \mathbb{E}_{B_n} h(Z_{n,P}^*) - \mathbb{E}_P h(Z_P) \right| \leq \sup_{h \in \text{BL}_1(\mathbb{D})} \left| \mathbb{E}_{B_n} h(G_{n,P}^*) - \mathbb{E}_P h(Z_P) \right| + \mathbb{E}_{B_n} (\|\zeta_{n,P}^*\|_{\mathbb{D}} \wedge 2),$$

where the first term is $o_P^*(1)$, since $G_{n,P}^* \rightsquigarrow_B Z_P$ by Theorem B.2, and the second term is $o_P(1)$ because $\|\zeta_{n,P}^*\|_{\mathbb{D}} = o_P(1)$ implies that $\mathbb{E}_P(\|\zeta_{n,P}^*\|_{\mathbb{D}} \wedge 2) = \mathbb{E}_P \mathbb{E}_{B_n}(\|\zeta_{n,P}^*\|_{\mathbb{D}} \wedge 2) \rightarrow 0$, which in turn implies that $\mathbb{E}_{B_n}(\|\zeta_{n,P}^*\|_{\mathbb{D}} \wedge 2) = o_P(1)$ by the Markov inequality. ■

E.3. Proof of Theorem 5.3. This is an immediate consequence of Theorems 5.1, 5.2, B.3, and B.4. ■

APPENDIX F. IMPLEMENTATION DETAILS

In this section, we provide details about how we implemented the methodology developed in the main body of the paper in the empirical example. We first discuss estimation of local average treatment effects (LATE) and then extend this discussion to estimation of local quantile treatment effects (LQTE). Estimation of all other quantities proceeds in a similar fashion and so is not discussed.

F.1. Local Average Treatment Effects. Recall that the LATE of treatment D on outcome Y is defined as

$$\Delta_{LATE} = \theta_Y(1) - \theta_Y(0) = \frac{\alpha_{1_1(D)Y}(1) - \alpha_{1_1(D)Y}(0)}{\alpha_{1_1(D)}(1) - \alpha_{1_1(D)}(0)} - \frac{\alpha_{1_0(D)Y}(1) - \alpha_{1_0(D)Y}(0)}{\alpha_{1_0(D)}(1) - \alpha_{1_0(D)}(0)}$$

for $\alpha_V(z)$ and $\theta_Y(d)$ defined in equations (2.1) and (2.3) respectively. It then follows by plugging in the definition of $\alpha_V(z)$ that we can express the LATE as

$$\Delta_{LATE} = \frac{\alpha_Y(1) - \alpha_Y(0)}{\alpha_{1_1(D)}(1) - \alpha_{1_1(D)}(0)}.$$

To obtain an estimate of the LATE, we thus need estimates of $\alpha_Y(z)$ and $\alpha_{1_1(D)}(z)$. Using the low-bias moment function given in equation (3.13), estimates of these key quantities can be constructed from estimates of $E_P[Y|Z = 1, X]$, $E_P[Y|Z = 0, X]$, $E_P[D|Z = 1, X]$, $E_P[D|Z = 0, X]$, and $E_P[Z|X]$ where Z is the binary instrument (401(k) eligibility); D is the binary treatment (401(k) participation); X is the set of raw covariates discussed in the empirical section; and Y is net financial assets. In our application, we have $E_P[D|Z = 0, X] = 0$ since one cannot participate unless one is eligible. We estimate the remainder of the functions using post-LASSO to estimate $E_P[Y|Z = 1, X]$ and $E_P[Y|Z = 0, X]$ and post- ℓ_1 -penalized logistic regression to estimate $E_P[D|Z = 1, X]$ and $E_P[Z|X]$.

To estimate $E_P[Y|Z = 1, X]$, we postulate that $E_P[Y|Z = 1, X] \approx f(X)' \beta_Y(1)$, where $f(X)$ is one of the pre-specified sets of controls discussed in the empirical section with dimension p . Let \mathcal{I}_1 denote the indices of observations that have $z_i = 1$. To estimate the coefficients $\beta_Y(1)$, we apply the formulation of the post-LASSO estimator given in Belloni et al. (2012) with outcomes $\{y_i\}_{i \in \mathcal{I}_1}$ and covariates $\{f(x_i)\}_{i \in \mathcal{I}_1}$. We set $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2(2p)))$ where $\Phi(\cdot)$ is the standard normal distribution function. We calculate penalty loadings according to Algorithm A.1 of Belloni et al. (2012) using post-LASSO coefficient estimates at each iteration and with a the maximum number of iterations set to 15.²⁸ Let $\widehat{\beta}_Y(1)$ denote the resulting post-LASSO estimates of the coefficients using λ given above and the final set of penalty loadings. We then estimate $E_P[Y|Z = 1, X = x_i]$ as $f(x_i)'\widehat{\beta}_Y(1)$ for each $i = 1, \dots, n$. We follow the same procedure to obtain estimates of $E_P[Y|Z = 0, X = x_i]$ as $f(x_i)'\widehat{\beta}_Y(0)$ for each $i = 1, \dots, n$ where $\widehat{\beta}_Y(0)$ are the post-LASSO estimates using only the observations with $z_i = 0$.

Estimation of $E_P[D|Z = 1, X]$ and $E_P[Z|X]$ proceed similarly replacing post-LASSO estimation with post- ℓ_1 -penalized logistic regression. Specifically, we assume that $E_P[D|Z = 1, X] \approx \Lambda_0(f(X)'\beta_D(1))$ where $\Lambda_0(\cdot)$ is the logistic link function. We then obtain estimates of $\beta_D(1)$ by using the post- ℓ_1 -penalized estimator defined in equations (3.10) and (3.11) based on the logistic link function and with outcomes $\{d_i\}_{i \in \mathcal{I}_1}$ and covariates $\{f(x_i)\}_{i \in \mathcal{I}_1}$ for \mathcal{I}_1 defined as above. We set $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2(2p)))$ where $\Phi(\cdot)$ is the standard normal distribution function. We calculate penalty loadings using Algorithm 6.1 of the main text with a maximum of 15 iterations. Let $\widehat{\beta}_D(1)$ denote the resulting post- ℓ_1 -penalized estimates of the coefficients using λ given above and the final set of penalty loadings. We estimate $E_P[D|Z = 1, X = x_i]$ as $\Lambda_0(f(x_i)'\widehat{\beta}_D(1))$ for each $i = 1, \dots, n$. We follow this procedure to obtain estimates of $E_P[Z|X = x_i]$ as $\Lambda_0(f(x_i)'\widehat{\beta}_Z)$ for each $i = 1, \dots, n$ where $\widehat{\beta}_Z$ are the post- ℓ_1 -penalized coefficient estimates obtained with $\{z_i\}_{i=1}^n$ as the outcome and $\{f(x_i)\}_{i=1}^n$ as covariates using $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2p))$.

²⁸Here and in all following instances, we stop iterating before reaching the maximum number of iterations if the ℓ_2 -norm of the difference in penalty loadings calculated across consecutive iterations is less than 10^{-6} .

Using these baseline quantities, we obtain estimates

$$\begin{aligned}\hat{\alpha}_Y(1) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{z_i(y_i - f(x_i)'\hat{\beta}_Y(1))}{\Lambda_0(f(x_i)'\hat{\beta}_Z)} + f(x_i)'\hat{\beta}_Y(1) \right) = \frac{1}{n} \sum_{i=1}^n \psi_{1,i} \\ \hat{\alpha}_Y(0) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{(1-z_i)(y_i - f(x_i)'\hat{\beta}_Y(0))}{1 - \Lambda_0(f(x_i)'\hat{\beta}_Z)} + f(x_i)'\hat{\beta}_Y(0) \right) = \frac{1}{n} \sum_{i=1}^n \psi_{0,i} \\ \hat{\alpha}_{11(D)}(1) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{z_i(d_i - \Lambda_0(f(x_i)'\hat{\beta}_D(1)))}{\Lambda_0(f(x_i)'\hat{\beta}_Z)} + \Lambda_0(f(x_i)'\hat{\beta}_D(1)) \right) = \frac{1}{n} \sum_{i=1}^n v_{1,i} \\ \hat{\alpha}_{11(D)}(0) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{(1-z_i)d_i}{1 - \Lambda_0(f(x_i)'\hat{\beta}_Z)} \right) = \frac{1}{n} \sum_{i=1}^n v_{0,i} = 0.\end{aligned}$$

We then plug these estimates in to obtain

$$\hat{\Delta}_{LATE} = \frac{\hat{\alpha}_Y(1) - \hat{\alpha}_Y(0)}{\hat{\alpha}_{11(D)}(1) - \hat{\alpha}_{11(D)}(0)}.$$

We report both analytic and bootstrap standard error estimates for the LATE. The analytic standard errors are calculated as

$$\sqrt{\frac{1}{n-1} \sum_{i=1}^n \left(\frac{\psi_{1,i} - \psi_{0,i}}{\hat{\alpha}_{11(D)}(1) - \hat{\alpha}_{11(D)}(0)} - \hat{\Delta}_{LATE} \right)^2 / n}.$$

We use wild bootstrap weights for obtaining the multiplier bootstrap estimates of the standard errors with 500 bootstrap replications. Specifically, for each $b = 1, \dots, 500$, we calculate a bootstrap estimate of the LATE as

$$\hat{\Delta}_{LATE}^b = \frac{\frac{1}{n} \sum_{i=1}^n (\psi_{1,i} - \psi_{0,i}) \xi_i^b}{\frac{1}{n} \sum_{i=1}^n (v_{1,i} - v_{0,i}) \xi_i^b}$$

where $\xi_i^b = 1 + r_{1,i}^b/\sqrt{2} + ((r_{2,i}^b)^2 - 1)/2$ is the bootstrap draw for multiplier weight for observation i in bootstrap repetition b where $r_{1,i}^b$ and $r_{2,i}^b$ are random numbers generated as iid draws from two independent standard normal random variables. The bootstrap standard error estimate is then the bootstrap interquartile range rescaled with the normal distribution: $[q_{LATE}(.75) - q_{LATE}(.25)]/[q_N(.75) - q_N(0.25)]$, where $q_{LATE}(p)$ is the p th quantile of $\{\hat{\Delta}_{LATE}^b\}_{b=1}^{500}$ and $q_N(p)$ is the p th quantile of the $N(0, 1)$.

F.2. Local Quantile Treatment Effects. Calculation and inference for LQTE is more cumbersome than for the LATE. We begin by choosing the set over which we would like to look at the LQTE. In our example, we chose to look at quantiles in the interval $[0.1, 0.9]$.

To calculate the LQTE, we first calculate the local average structural function for outcomes $Y_u = 1(Y \leq u)$ for a set of u and then invert to obtain estimates of the LQTE. In our example, we chose to look at $u \in [q_Y(.05), q_Y(.95)]$ where $q_Y(.05)$ and $q_Y(.95)$ are respectively the sample 5th and 95th percentiles of the outcome of interest Y . Since looking at the continuum of values in this interval is infeasible, we discretize the interval and look at $Y_u = 1(Y \leq u)$ for $u \in \{q_Y(.05), q_Y(.06), q_Y(.07), \dots, q_Y(.93), q_Y(.94), q_Y(.95)\}$. I.e. we set u equal to each percentile

of Y between the 5th and 95th percentiles for a total of 91 different values of u to be considered. For each value of u , we need an estimate of the local average structural function defined in (2.3) for $d \in \{0, 1\}$:

$$\theta_{1(Y \leq u)}(d) = \frac{\alpha_{1_d(D)1(Y \leq u)}(1) - \alpha_{1_d(D)1(Y \leq u)}(0)}{\alpha_{1_d(D)}(1) - \alpha_{1_d(D)}(0)}.$$

As with the LATE, we need estimates of $E_P[D|Z = 1, X]$ and $E_P[Z|X]$. We estimate these quantities as we did for the LATE but change the value of the penalty parameter used to reflect the fact that we are now interested in a large set, in theory a continuum, of model selection problems. Specifically, we assume that $E_P[D|Z = 1, X] \approx \Lambda_0(f(X)' \beta_D(1))$ where $\Lambda_0(\cdot)$ is the logistic link function and $f(X)$ is one of the pre-specified sets of controls discussed in the empirical section with dimension p . We then obtain estimates of $\beta_D(1)$ by using the post- ℓ_1 -penalized estimator defined in equations (3.10) and (3.11) based on the logistic link function and with outcomes $\{d_i\}_{i \in \mathcal{I}_1}$ and covariates $\{f(x_i)\}_{i \in \mathcal{I}_1}$ for \mathcal{I}_1 defined as above. We set $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2n(2p)))$ where $\Phi(\cdot)$ is the standard normal distribution function. We calculate penalty loadings using Algorithm 6.1 with a maximum of 15 iterations. Let $\widehat{\beta}_D(1)$ denote the resulting post- ℓ_1 -penalized estimates of the coefficients using λ given above and the final set of penalty loadings. We estimate $E_P[D|Z = 1, X = x_i]$ as $\Lambda_0(f(x_i)' \widehat{\beta}_D(1))$ for each $i = 1, \dots, n$. We follow this procedure to obtain estimates of $E_P[Z|X]$ as $\Lambda_0(f(x_i)' \widehat{\beta}_Z)$ for each $i = 1, \dots, n$ where $\widehat{\beta}_Z$ are the post- ℓ_1 -penalized coefficient estimates obtained with $\{z_i\}_{i=1}^n$ as the outcome and $\{f(x_i)\}_{i=1}^n$ as covariates and $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2np))$. We also still have $E_P[D|Z = 0, X] = 0$ in our application since one cannot participate in a 401(k) unless one is eligible. We then plug-in these estimates to obtain

$$\begin{aligned} \widehat{\alpha}_{1_1(D)}(1) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{z_i(d_i - \Lambda_0(f(x_i)' \widehat{\beta}_D(1)))}{\Lambda_0(f(x_i)' \widehat{\beta}_Z)} + \Lambda_0(f(x_i)' \widehat{\beta}_D(1)) \right) = \frac{1}{n} \sum_{i=1}^n v_{1,1,i} \\ \widehat{\alpha}_{1_1(D)}(0) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{(1 - z_i)d_i}{1 - \Lambda_0(f(x_i)' \widehat{\beta}_Z)} \right) = \frac{1}{n} \sum_{i=1}^n v_{1,0,i} = 0 \\ \widehat{\alpha}_{1_0(D)}(1) &= 1 - \widehat{\alpha}_{1_1(D)}(1), \quad \widehat{\alpha}_{1_0(D)}(0) = 1 - \widehat{\alpha}_{1_1(D)}(0). \end{aligned}$$

We also need to obtain estimates of $\alpha_{1_d(D)1(Y \leq u)}(z)$ for each value of u and for

$$(z, d) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}.$$

These estimates will depend on the propensity score, $E_P[Z|X]$, estimated above and quantities of the form $E_P[1(D = d)1(Y \leq u)|Z = z, X]$. We again approximate this function with $E_P[1(D = d)1(Y \leq u)|Z = z, X] \approx \Lambda_0(X' \beta_{1_d(D)Y_u}(z))$ and estimate the coefficients $\beta_{1_d(D)Y_u}(z)$ for each combination of d and z and each u using the post- ℓ_1 -penalized estimator defined in equations (3.10) and (3.11) based on the logistic link function. We set $\lambda = 1.1\sqrt{n}\Phi^{-1}(1 - (1/\log(n))/(2n(2p)))$ where $\Phi(\cdot)$ is the standard normal distribution function. We calculate penalty loadings using Algorithm 6.1 of the main text with a maximum of 15 iterations. We follow this procedure for each u with $\{1(y_i \leq u)1(d_i = 1)\}_{i \in \mathcal{I}_1}$ as the outcome and covariates $\{f(x_i)\}_{i \in \mathcal{I}_1}$, with $\{1(y_i \leq u)1(d_i = 0)\}_{i \in \mathcal{I}_1}$ as the outcome and covariates $\{f(x_i)\}_{i \in \mathcal{I}_1}$, and with $\{1(y_i \leq u)1(d_i = 0)\}_{i \in \mathcal{I}_0}$ as the outcome and covariates $\{f(x_i)\}_{i \in \mathcal{I}_0}$ for \mathcal{I}_1 and \mathcal{I}_0 defined as above to obtain point estimates $\widehat{\beta}_{1_1(D)Y_u}(1)$, $\widehat{\beta}_{1_0(D)Y_u}(1)$, and $\widehat{\beta}_{1_0(D)Y_u}(0)$ respectively. We then estimate $E_P[1(D = 1)1(Y \leq u)|Z = 1, X = x_i]$

as $\Lambda_0(f(x_i)' \widehat{\beta}_{1_1(D)Y_u}(1))$ for each $i = 1, \dots, n$ and obtain estimates of $E_P[1(D = 0)1(Y \leq u)|Z = 1, X = x_i]$, and $E_P[1(D = 0)1(Y \leq u)|Z = 0, X = x_i]$ analogously. As before, we have $E_P[1(D = 1)1(Y \leq u)|Z = 0, X] = 0$ since one cannot participate unless one is eligible. We then plug-in these estimates to obtain

$$\begin{aligned}\widehat{\alpha}_{1_1(D)1(Y \leq u)}(1) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{z_i(d_i 1(y_i \leq u) - \Lambda_0(f(x_i)' \widehat{\beta}_{1_1(D)Y_u}(1)))}{\Lambda_0(f(x_i)' \widehat{\beta}_Z)} + \Lambda_0(f(x_i)' \widehat{\beta}_{1_1(D)Y_u}(1)) \right) = \frac{1}{n} \sum_{i=1}^n \kappa_{u,1,1,i} \\ \widehat{\alpha}_{1_1(D)1(Y \leq u)}(0) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{(1 - z_i)(d_i 1(y_i \leq u))}{1 - \Lambda_0(f(x_i)' \widehat{\beta}_Z)} \right) = \frac{1}{n} \sum_{i=1}^n \kappa_{u,1,0,i} = 0 \\ \widehat{\alpha}_{1_0(D)1(Y \leq u)}(1) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{z_i((1 - d_i) 1(y_i \leq u) - \Lambda_0(f(x_i)' \widehat{\beta}_{1_0(D)Y_u}(1)))}{\Lambda_0(f(x_i)' \widehat{\beta}_Z)} + \Lambda_0(f(x_i)' \widehat{\beta}_{1_0(D)Y_u}(1)) \right) = \frac{1}{n} \sum_{i=1}^n \kappa_{u,0,1,i} \\ \widehat{\alpha}_{1_0(D)1(Y \leq u)}(0) &= \frac{1}{n} \sum_{i=1}^n \left(\frac{(1 - z_i)((1 - d_i) 1(y_i \leq u) - \Lambda_0(f(x_i)' \widehat{\beta}_{1_0(D)Y_u}(0)))}{1 - \Lambda_0(f(x_i)' \widehat{\beta}_Z)} + \Lambda_0(f(x_i)' \widehat{\beta}_{1_0(D)Y_u}(0)) \right) = \frac{1}{n} \sum_{i=1}^n \kappa_{u,0,0,i}.\end{aligned}$$

Estimates of the local average structural (distribution) functions are formed using the estimators defined in the previous two paragraphs as

$$\widehat{\theta}_{1(Y \leq u)}(d) = \frac{\widehat{\alpha}_{1_d(D)1(Y \leq u)}(1) - \widehat{\alpha}_{1_d(D)1(Y \leq u)}(0)}{\widehat{\alpha}_{1_d(D)}(1) - \widehat{\alpha}_{1_d(D)}(0)}.$$

To obtain LQTE estimates, we then need to invert these local average structural functions. Since we only have the estimated distribution for each d evaluated on the finite grid of points $u \in \{q_Y(.05), q_Y(.06), q_Y(.07), \dots, q_Y(.93), q_Y(.94), q_Y(.95)\}$, we do this inversion by linearly interpolating the value of the distribution function between these points to find the value of the outcome associated with each quantile in the set $q \in [0.1, 0.11, .0, 12, \dots, 0.89, .0.9]$ which we denote as $\widehat{\theta}_Y^{\leftarrow}(q, d)$. The LQTE at point q is then estimated as $\widehat{\Delta}(q) = \widehat{\theta}_Y^{\leftarrow}(q, 1) - \widehat{\theta}_Y^{\leftarrow}(q, 0)$.

For the LQTE, we only report inference based on the multiplier bootstrap using 500 bootstrap replications. For each $b = 1, \dots, 500$, we generate bootstrap weights as $\xi_i^b = 1 + r_{1,i}^b / \sqrt{2} + ((r_{2,i}^b)^2 - 1) / 2$ for observation i in bootstrap repetition b where $r_{1,i}^b$ and $r_{2,i}^b$ are random numbers generated as iid draws from two independent standard normal random variables. We then use these weights to form bootstrap estimates of the local average structural functions

$$\widehat{\theta}_{1(Y \leq u)}^b(d) = \frac{\widehat{\alpha}_{1_d(D)1(Y \leq u)}^b(1) - \widehat{\alpha}_{1_d(D)1(Y \leq u)}^b(0)}{\widehat{\alpha}_{1_d(D)}^b(1) - \widehat{\alpha}_{1_d(D)}^b(0)}$$

where

$$\begin{aligned}\widehat{\alpha}_{1_1(D)}^b(1) &= \frac{1}{n} \sum_{i=1}^n \xi_i^b v_{1,1,i}, \quad \widehat{\alpha}_{1_1(D)}^b(0) = \frac{1}{n} \sum_{i=1}^n \xi_i^b v_{1,0,i}, \\ \widehat{\alpha}_{1_0(D)}^b(1) &= 1 - \widehat{\alpha}_{1_1(D)}^b(1), \quad \widehat{\alpha}_{1_0(D)}^b(0) = 1 - \widehat{\alpha}_{1_1(D)}^b(0), \\ \widehat{\alpha}_{1_1(D)1(Y \leq u)}^b(1) &= \frac{1}{n} \sum_{i=1}^n \xi_i^b \kappa_{u,1,1,i}, \quad \widehat{\alpha}_{1_1(D)1(Y \leq u)}^b(0) = \frac{1}{n} \sum_{i=1}^n \xi_i^b \kappa_{u,1,0,i} = 0, \\ \widehat{\alpha}_{1_0(D)1(Y \leq u)}^b(1) &= \frac{1}{n} \sum_{i=1}^n \xi_i^b \kappa_{u,0,1,i}, \quad \widehat{\alpha}_{1_0(D)1(Y \leq u)}^b(0) = \frac{1}{n} \sum_{i=1}^n \xi_i^b \kappa_{u,0,0,i}.\end{aligned}$$

From these bootstrap estimates of the average structural distribution functions, we obtain bootstrap LQTE estimates as above through inversion by linearly interpolating the value of the distribution function between the finite set of points at which we have estimated values to find the value of the outcome associated with each quantile in the set $q \in [0.1, 0.11, .0, 12, \dots, 0.89, .0.9]$, denoted $(\hat{\theta}_Y^{\leftarrow}(q, d))^b$. The bootstrap estimate of the LQTE for bootstrap replication b at point q is then $\hat{\Delta}^b(q) = (\hat{\theta}_Y^{\leftarrow}(q, 1))^b - (\hat{\theta}_Y^{\leftarrow}(q, 0))^b$. We form bootstrap standard error estimates for the LQTE at each quantile q as

$$s(q) = [q_{LQTE}(.75) - q_{LQTE}(.25)]/[q_N(.75) - q_N(.25)],$$

where $q_{LQTE}(p)$ is the p th quantile of $\{\hat{\Delta}^b(q)\}_{b=1}^{500}$ and $q_N(p)$ is the p th quantile of the $N(0, 1)$.

We also use the bootstrap LQTE estimates to obtain the critical values we use when plotting the uniform confidence bands in our example. We form bootstrap t-statistics for each quantile q as $t^b(q) = (\hat{\Delta}^b(q) - \hat{\Delta}(q))/s(q)$. We then take $t_{\max}^b = \max_q\{|t^b(q)|\}$ and use the 95th percentile of the bootstrap distribution of t_{\max}^b as the critical value in constructing the confidence intervals for our figures following for example Chernozhukov et al. (2013).

REFERENCES

- ABADIE, A. (2002): "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models," *Journal of the American Statistical Association*, 97, 284–292.
- (2003): "Semiparametric Instrumental Variable Estimation of Treatment Response Models," *Journal of Econometrics*, 113, 231–263.
- AI, C. AND X. CHEN (2003): "Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions," *Econometrica*, 71, 1795–1843.
- (2012): "The semiparametric efficiency bound for models of sequential moment restrictions containing unknown functions," *Journal of Econometrics*, 170, 442–457.
- ANDREWS, D. W. (1994a): "Empirical process methods in econometrics," *Handbook of Econometrics*, 4, 2247–2294.
- ANDREWS, D. W. K. (1994b): "Asymptotics for semiparametric econometric models via stochastic equicontinuity," *Econometrica*, 62, 43–72.
- ANGRIST, J. D. AND J.-S. PISCHKE (2008): *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.
- BACH, F. (2010): "Self-concordant analysis for logistic regression," *Electronic Journal of Statistics*, 4, 384–414.
- BELLONI, A., D. CHEN, V. CHERNOZHUKOV, AND C. HANSEN (2012): "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain," *Econometrica*, 80, 2369–2429, arxiv, 2010.
- BELLONI, A. AND V. CHERNOZHUKOV (2011): " ℓ_1 -Penalized Quantile Regression for High Dimensional Sparse Models," *Annals of Statistics*, 39, 82–130.
- (2013): "Least Squares After Model Selection in High-dimensional Sparse Models," *Bernoulli*, 19, 521–547, arXiv, 2009.
- BELLONI, A., V. CHERNOZHUKOV, I. FERNANDEZ-VAL, AND C. HANSEN (2015): "Supplement to "Program Evaluation with High-Dimensional Data"," Tech. rep., ArXiv.
- BELLONI, A., V. CHERNOZHUKOV, AND C. HANSEN (2010): "LASSO Methods for Gaussian Instrumental Variables Models," 2010 arXiv:[math.ST], <http://arxiv.org/abs/1012.1297>.
- (2013a): "Inference for High-Dimensional Sparse Econometric Models," *Advances in Economics and Econometrics. 10th World Congress of Econometric Society. August 2010*, III, 245–295.
- (2014a): "Inference on Treatment Effects After Selection Amongst High-Dimensional Controls," *Review of Economic Studies*, 81, 608–650.
- BELLONI, A., V. CHERNOZHUKOV, AND K. KATO (2013b): "Uniform Post Selection Inference for LAD Regression Models," *arXiv preprint arXiv:1304.0282*.
- BELLONI, A., V. CHERNOZHUKOV, AND L. WANG (2011): "Square-Root-LASSO: Pivotal Recovery of Sparse Signals via Conic Programming," *Biometrika*, 98, 791–806, arxiv, 2010.
- BELLONI, A., V. CHERNOZHUKOV, L. WANG, ET AL. (2014b): "Pivotal estimation via square-root lasso in nonparametric regression," *The Annals of Statistics*, 42, 757–788.
- BELLONI, A., V. CHERNOZHUKOV, AND Y. WEI (2013c): "Honest Confidence Regions for Logistic Regression with a Large Number of Controls," *arXiv preprint arXiv:1304.3969*.
- BENJAMIN, D. J. (2003): "Does 401(k) eligibility increase saving? Evidence from propensity score subclassification," *Journal of Public Economics*, 87, 1259–1290.
- BICKEL, P. J. (1982): "On adaptive estimation," *The Annals of Statistics*, 647–671.
- BICKEL, P. J. AND D. A. FREEDMAN (1981): "Some asymptotic theory for the bootstrap," *The Annals of Statistics*, 1196–1217.

- BICKEL, P. J., Y. RITOV, AND A. B. TSYBAKOV (2009): "Simultaneous analysis of Lasso and Dantzig selector," *Annals of Statistics*, 37, 1705–1732.
- CANDÈS, E. AND T. TAO (2007): "The Dantzig selector: statistical estimation when p is much larger than n ," *Ann. Statist.*, 35, 2313–2351.
- CANER, M. AND H. H. ZHANG (2014): "Adaptive elastic net for generalized methods of moments," *Journal of Business and Economic Statistics*, 32, 30–47.
- CATTANEO, M., M. JANSSON, AND W. NEWEY (2010): "Alternative Asymptotics and the Partially Linear Model with Many Regressors," *Working Paper*, <http://econ-www.mit.edu/files/6204>.
- CATTANEO, M. D. (2010): "Efficient semiparametric estimation of multi-valued treatment effects under ignorability," *Journal of Econometrics*, 155, 138–154.
- CHAMBERLAIN, G. AND G. W. IMBENS (2003): "Nonparametric applications of Bayesian inference," *Journal of Business & Economic Statistics*, 21, 12–18.
- CHEN, X. (2007): "Large Sample Sieve Estimation of Semi-Nonparametric Models," *Handbook of Econometrics*, 6, 5559–5632.
- CHEN, X., O. LINTON, AND I. V. KEILEGOM (2003): "Estimation of Semiparametric Models when the Criterion Function Is Not Smooth," *Econometrica*, 71, 1591–1608.
- CHERNOZHUKOV, V., D. CHETVERIKOV, AND K. KATO (2012): "Gaussian approximation of suprema of empirical processes," *ArXiv e-prints*.
- CHERNOZHUKOV, V. AND I. FERNÁNDEZ-VAL (2005): "Subsampling inference on quantile regression processes," *Sankhyā*, 67, 253–276.
- CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND B. MELLY (2013): "Inference on counterfactual distributions," *Econometrica*, 81, 2205–2268.
- CHERNOZHUKOV, V. AND C. HANSEN (2004): "The impact of 401(k) participation on the wealth distribution: An instrumental quantile regression analysis," *Review of Economics and Statistics*, 86, 735–751.
- (2005): "An IV Model of Quantile Treatment Effects," *Econometrica*, 73, 245–262.
- (2006): "Instrumental quantile regression inference for structural and treatment effect models," *J. Econometrics*, 132, 491–525.
- CHERNOZHUKOV, V., C. HANSEN, AND M. SPINDLER (2015): "Valid Post-Selection and Post-Regularization Inference In Linear Models with Many Controls and Instruments," *American Economic Review: Papers and Proceedings*.
- CHESHER, A. (2003): "Identification in nonseparable models," *Econometrica*, 71, 1405–1441.
- DUDLEY, R. M. (1999): *Uniform central limit theorems*, vol. 63 of *Cambridge Studies in Advanced Mathematics*, Cambridge: Cambridge University Press.
- ENGEN, E. M. AND W. G. GALE (2000): "The Effects of 401(k) Plans on Household Wealth: Differences Across Earnings Groups," Working Paper 8032, National Bureau of Economic Research.
- ENGEN, E. M., W. G. GALE, AND J. K. SCHOLZ (1996): "The Illusory Effects of Saving Incentives on Saving," *Journal of Economic Perspectives*, 10, 113–138.
- ESCANCIANO, J. C. AND L. ZHU (2013): "Set inferences and sensitivity analysis in semiparametric conditionally identified models," CeMMAP working papers CWP55/13, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
- FAN, J. AND R. LI (2001): "Variable selection via nonconcave penalized likelihood and its oracle properties," *Journal of American Statistical Association*, 96, 1348–1360.
- FARRELL, M. (2013): "Robust Inference on Average Treatment Effects with Possibly More Covariates than Observations," *Working Paper*.
- FRANK, I. E. AND J. H. FRIEDMAN (1993): "A Statistical View of Some Chemometrics Regression Tools," *Technometrics*, 35, 109–135.
- FRÖLICH, M. AND B. MELLY (2013): "Identification of treatment effects on the treated with one-sided non-compliance," *Econometric Reviews*, 32, 384–414.
- GINÉ, E. AND J. ZINN (1984): "Some limit theorems for empirical processes," *Ann. Probab.*, 12, 929–998, with discussion.
- HAHN, J. (1997): "Bayesian bootstrap of the quantile regression estimator: a large sample study," *Internat. Econom. Rev.*, 38, 795–808.
- (1998): "On the role of the propensity score in efficient semiparametric estimation of average treatment effects," *Econometrica*, 66, 315–331.
- HANSEN, B. E. (1996): "Inference when a nuisance parameter is not identified under the null hypothesis," *Econometrica*, 64, 413–430.
- HANSEN, L. P. (1982): "Large sample properties of generalized method of moments estimators," *Econometrica*, 50, 1029–1054.
- HANSEN, L. P. AND K. J. SINGLETON (1982): "Generalized instrumental variables estimation of nonlinear rational expectations models," *Econometrica*, 50, 1269–1286.
- HONG, H. AND D. NEKIPELOV (2010): "Semiparametric efficiency in nonlinear LATE models," *Quantitative Economics*, 1, 279–304.
- HONG, H. AND O. SCAILLET (2006): "A fast subsampling method for nonlinear dynamic models," *J. Econometrics*, 133, 557–578.
- HUANG, J., J. L. HOROWITZ, AND S. MA (2008): "Asymptotic properties of bridge estimators in sparse high-dimensional regression models," *The Annals of Statistics*, 36, 587613.
- HUANG, J., J. L. HOROWITZ, AND F. WEI (2010): "Variable selection in nonparametric additive models," *Ann. Statist.*, 38, 2282–2313.
- IMBENS, G. W. AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475.
- IMBENS, G. W. AND W. K. NEWEY (2009): "Identification and estimation of triangular simultaneous equations models without additivity," *Econometrica*, 77, 1481–1512.
- KATO, K. (2011): "Group Lasso for high dimensional sparse quantile regression models," Preprint, ArXiv.
- KLINE, P. AND A. SANTOS (2012): "A Score Based Approach to Wild Bootstrap Inference," *Journal of Econometric Methods*, 1, 23–41.
- KOENKER, R. (1988): "Asymptotic Theory and Econometric Practice," *Journal of Applied Econometrics*, 3, 139–147.
- (2005): *Quantile regression*, Cambridge university press.
- KOSOROK, M. R. (2008): *Introduction to Empirical Processes and Semiparametric Inference*, Series in Statistics, Berlin: Springer.
- LEEB, H. AND B. M. PÖTSCHER (2008a): "Can one estimate the unconditional distribution of post-model-selection estimators?" *Econometric Theory*, 24, 338–376.
- (2008b): "Recent developments in model selection and related areas," *Econometric Theory*, 24, 319–322.

- LINTON, O. (1996): "Edgeworth approximation for MINPIN estimators in semiparametric regression models," *Econometric Theory*, 12, 30–60.
- MAMMEN, E. (1993): "Bootstrap and wild bootstrap for high dimensional linear models," *The Annals of Statistics*, 255–285.
- MEINSHAUSEN, N. AND B. YU (2009): "Lasso-type recovery of sparse representations for high-dimensional data," *Annals of Statistics*, 37, 2246–2270.
- NEWBY, W. K. (1990): "Semiparametric efficiency bounds," *Journal of Applied Econometrics*, 5, 99–135.
- (1994): "The asymptotic variance of semiparametric estimators," *Econometrica*, 62, 1349–1382.
- (1997): "Convergence Rates and Asymptotic Normality for Series Estimators," *Journal of Econometrics*, 79, 147–168.
- NEYMAN, J. (1979): " $C(\alpha)$ tests and their use," *Sankhya*, 41, 1–21.
- OGBURN, E. L., A. ROTNITZKY, AND J. M. ROBINS (2015): "Doubly robust estimation of the local average treatment effect curve," *Journal of the Royal Statistical Society: Series B*, 77, 373–396.
- POTERBA, J. M., S. F. VENTI, AND D. A. WISE (1994): "401(k) Plans and Tax-Deferred savings," in *Studies in the Economics of Aging*, ed. by D. A. Wise, Chicago, IL: University of Chicago Press.
- (1995): "Do 401(k) Contributions Crowd Out Other Personal Saving?" *Journal of Public Economics*, 58, 1–32.
- (1996): "Personal Retirement Saving Programs and Asset Accumulation: Reconciling the Evidence," Working Paper 5599, National Bureau of Economic Research.
- (2001): "The Transition to Personal Accounts and Increasing Retirement Wealth: Macro and Micro Evidence," Working Paper 8610, National Bureau of Economic Research.
- PÖTSCHER, B. (2009): "Confidence Sets Based on Sparse Estimators Are Necessarily Large," *Sankhya*, 71-A, 1–18.
- ROBINS, J. M. AND A. ROTNITZKY (1995): "Semiparametric efficiency in multivariate regression models with missing data," *J. Amer. Statist. Assoc.*, 90, 122–129.
- ROBINSON, P. M. (1988): "Root- N -consistent semiparametric regression," *Econometrica*, 56, 931–954.
- ROMANO, J. P. AND A. M. SHAIKH (2012): "On the uniform asymptotic validity of subsampling and the bootstrap," *The Annals of Statistics*, 40, 2798–2822.
- ROTHER, C. AND S. FIRPO (2013): "Semiparametric Estimation and Inference Using Doubly Robust Moment Conditions," Tech. rep., NYU preprint.
- TIBSHIRANI, R. (1996): "Regression shrinkage and selection via the Lasso," *J. Roy. Statist. Soc. Ser. B*, 58, 267–288.
- TSYBAKOV, A. B. (2009): *Introduction to nonparametric estimation*, Springer.
- VAN DE GEER, S. A. (2008): "High-dimensional generalized linear models and the lasso," *Annals of Statistics*, 36, 614–645.
- VAN DER VAART, A. W. (1991): "On differentiable functionals," *The Annals of Statistics*, 178–204.
- (1998): *Asymptotic Statistics*, Cambridge University Press.
- VAN DER VAART, A. W. AND J. A. WELLNER (1996): *Weak Convergence and Empirical Processes*, Springer Series in Statistics.
- WASSERMAN, L. (2006): *All of nonparametric statistics*, Springer New York.
- WOOLDRIDGE, J. M. (2010): *Econometric Analysis of Cross Section and Panel Data*, Cambridge, Massachusetts: The MIT Press, second ed.
- ZOU, H. (2006): "The Adaptive Lasso And Its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429.

Table 1: Estimates and standard errors of average effects

Specification		Exogenous: 401(k) Eligibility			Endogenous: 401(k) Participation			
Series approximation	Dimension	Selection	Linear Model	ATE	ATE-T	Linear IV	LATE	LATE-T
Quadratic Spline	35 (32)	N	8997 (1252)	8093 (1082) {967}	11250 (1513) {1423}	12926 (1796) {1413}	11579 (1548) {2195}	15969 (2148) {2195}
Quadratic Spline	35 (32)	Y	8967 (1270)	7614 (1224) {1234}	10257 (1776) {1676}	12890 (1821) {1709}	10937 (1758) {2576}	14560 (2520) {2576}
Quadratic Spline Plus Interactions	311 (272)	N	9019 (1258)	11775 (4202) {4202}	11740 (1779) {1757}	12973 (1804) {1709}	17529 (6256) {6249}	16664 (2526) {2558}
Quadratic Spline Plus Interactions	311 (272)	Y	8307 (1313)	7077 (1358) {1237}	8830 (2133) {2105}	11784 (1995) {1709}	10168 (1952) {1963}	12533 (3027) {2818}
Quadratic Spline Plus Many Interactions	1756 (1526)	N	8860 (1358)	- - -	- - -	12827 (1960) -	- - -	- - -
Quadratic Spline Plus Many Interactions	1756 (1526)	Y	8536 (1321)	7848 (1317) {1334}	9602 (2047) {1894}	10671 (2001) {1835}	11267 (1890) {1835}	13629 (2906) {2862}

Notes: The sample is drawn from the 1991 SIPP and consists of 9,915 observations. All the specifications control for age, income, family size, education, marital status, two-earner status, defined benefit pension status, IRA participation status, and home ownership status. Quadratic Spline uses indicators for marital status, two-earner status, defined benefit pension status, IRA participation status, and home ownership status; a third order polynomial in age; second order polynomials in education and family size; and a piecewise quadratic polynomial in income with six break points. The "Quadratic Spline Plus Interactions" specification include all first-order interactions between the income variables and the remaining variables. The specification denoted "Quadratic Spline Plus Many Interactions" takes all first-order interactions between all non-income variables and then fully interacts these interactions as well as the main effects with all income variables. Analytic standard errors are given in parentheses. Bootstrap standard errors based on 500 repetitions with Mammen (1993) multipliers are given in braces.




FIGURE 1. QTE and QTE-T estimates of the effect of 401(k) eligibility on net financial assets.

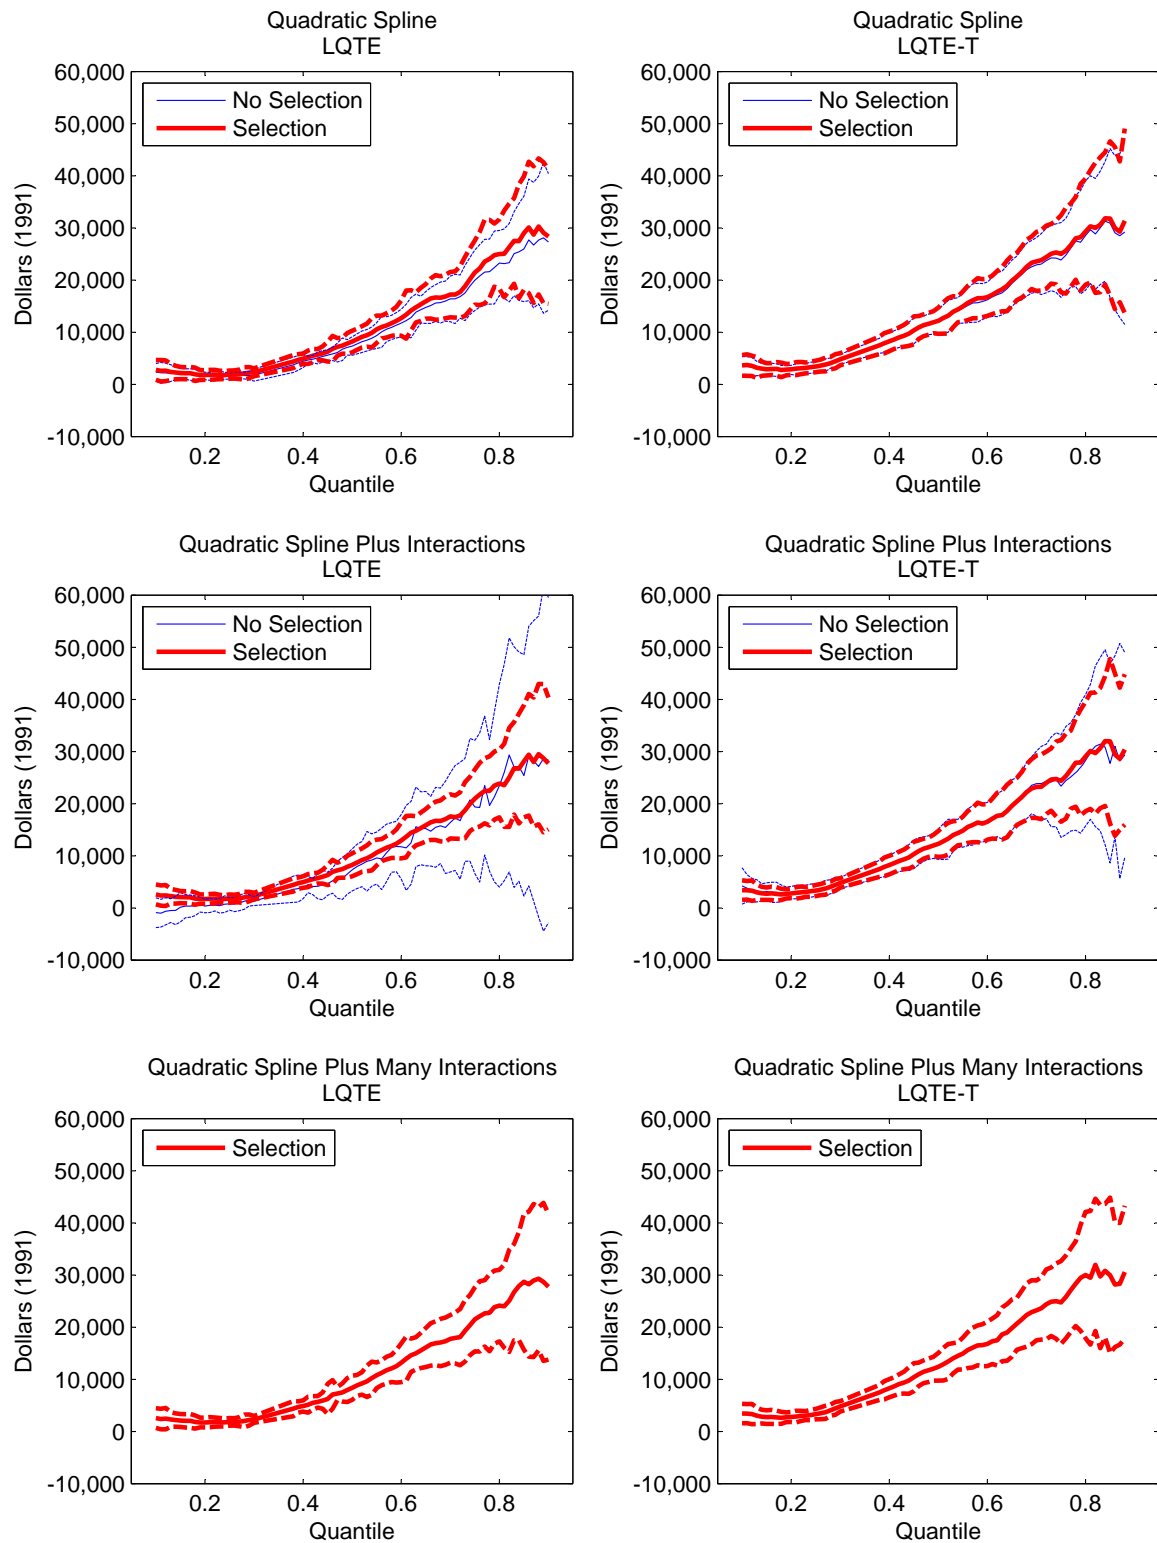


FIGURE 2. LQTE and LQTE-T estimates of the effect of 401(k) participation on net financial assets.