Econometric Theory, **0**, 2024, 1–24.
doi:10.1017/S0266466624000227

WELFARE ANALYSIS VIA MARGINAL TREATMENT EFFECTS

YUYA SASAKI
Vanderbilt University

TAKUYA URA
University of California, Davis

We consider a causal structure with endogeneity, i.e., unobserved confoundedness, where an instrumental variable is available. In this setting, we show that the mean social welfare function can be identified and represented via the marginal treatment effect as the operator kernel. This representation result can be applied to a variety of statistical decision rules for treatment choice, including plug-in rules, Bayes rules, and empirical welfare maximization rules. Focusing on the application of the empirical welfare maximization framework, we provide convergence rates of the worst-case average welfare loss (regret).

1. INTRODUCTION

One of the most important economists' goals is to advise policymakers on assigning heterogeneous individuals to treatment under consideration subject to budgetary, legal, and ethical constraints, based on evidence from available data. To achieve this goal, it is crucial to identify the social welfare function in observational data settings. For many observational datasets used by empirical researchers, treatments are likely to be endogenously selected by individuals, rather than randomly assigned. Furthermore, the effects of these treatments are often heterogeneous across individuals, even after controlling for their observable attributes. In this light, we propose a novel method to identify the mean social welfare function in the presence of unobserved heterogeneity in treatment effects while accounting for endogenous treatment selection in observational data.

It is well known today that the marginal treatment effects (MTEs) can measure heterogeneous treatment effects (Björklund and Moffitt, 1987), and the MTEs can be identified with an instrumental variable under endogenous treatment selection (Heckman and Vytlacil, 2001, 2005, 2007). It is, therefore, a natural idea to use the MTEs as a building block for the identification of the social welfare function in the presence of unobserved heterogeneity and endogeneity. In this paper, we show

We thank Toru Kitagawa, Simon Lee, two anonymous referees, and the seminar participants at LMU Munich for their very useful comments. The usual disclaimers apply. Sasaki thanks Brian and Charlotte Grove Chair for research support. Address correspondence to Takuya Ura, Department of Economics, University of California, Davis, One Shields Avenue, Davis, CA 95616, USA; e-mail: takura@ucdavis.edu

that the mean social welfare function can indeed be identified and represented via the MTEs as the operator kernel. Since the identification and estimation of the MTEs have been well established in the existing literature (e.g., Heckman and Vytlacil, 2001, 2005, 2007; Carneiro and Lee, 2009; Carneiro, Heckman, and Vytlacil, 2010; Brinch, Mogstad, and Wiswall, 2017; Lee and Salanié, 2018), our result thus paves the way for these existing theories and methods for MTEs to be directly applied to welfare analysis.

The average treatment effects (ATEs) conditional on observables are the key components of the social welfare functions. Under endogeneity, however, the ATEs generally fail to be identified due to a lack of sufficient variations in the treatment choice probabilities in response to instrumental variations. On the other hand, one of the main advantages of the MTEs is that the ATEs can be identified via their functional-form restrictions on the MTEs (e.g., Moffitt, 2008; French and Song, 2014). Consequently, the use of the MTEs facilitates the identification of the social welfare functions under endogeneity with commonly available instruments.

Once the mean social welfare function has been identified via the MTEs, we can apply it to a variety of policymakers' statistical decision problems of treatment choice, including those based on plug-in rules, Bayes rules, and empirical welfare maximization rules (see Hirano and Porter, 2020, Sect. 2.3). Focusing on the empirical welfare maximization rules in particular, we can take advantage of the technology developed by Kitagawa and Tetenov (2018) to analyze the properties of the empirical welfare maximization method in the spirit of Manski (2004). Specifically, under both heterogeneity and endogeneity, we can derive convergence rates of the worst-case average welfare loss (regret) from the maximum empirical welfare. As such, our result contributes to the literature by extending the scope of applicability of the empirical welfare maximization framework of Kitagawa and Tetenov (2018), which is originally based on the assumption of selection on observables or unconfoundedness, to the framework that allows for unobserved confoundedness or endogeneity.

There are a few papers that study policy learning under endogeneity (e.g., Kallus and Zhou, 2018; Cui and Tchetgen Tchetgen, 2020; Athey and Wager, 2021; Byambadalai, 2021; Liu, 2022). The most closely related in terms of the empirical welfare maximization is the recent paper by Athey and Wager (2021). The result that we propose in this paper neither nests nor is nested by that of Athey and Wager (2021)—these two papers play rather complementary roles. On the one hand, Athey and Wager (2021, eqn. (16)) assume homogeneous treatment effects,[1] which implies a constant marginal treatment effect, while our framework can allow for unobserved heterogeneity in treatment effects. On the other hand, the applicability of our proposed method hinges on the identification of the MTEs, while Athey and Wager (2021) do not need to identify the MTEs for their objective. In other words, our framework can accommodate unobserved heterogeneity but requires the MTEs

[1] Athey and Wager (2021) identify the mean welfare for compliers. It requires additional assumptions to identify the mean welfare for the entire population under heterogeneity.

to be identified. This tradeoff illustrates a complementary relationship between our result and the result developed by Athey and Wager (2021). Also closely related is a more recent paper by Liu (2022)—we discuss more details about the relation later in Section 3.

More broadly, this paper aims to contribute to the literature on statistical decisions in econometrics—see the recent survey by Hirano and Porter (2020) for a comprehensive review of this subject. In particular, we focus on an application of our representation theorem to bounding the worst-case average welfare loss in the spirit of Manski (2004) with the recent technology developed by Kitagawa and Tetenov (2018).[2] Also closely related is the literature on the MTEs. The applicability of our representation theorem relies on the identification and estimation of the MTEs from the aforementioned papers. Finally, this paper also complements the literature on policy relevant treatment effects (e.g., Heckman and Vytlacil, 2001, 2005, 2007; Carneiro, Lokshin, and Umapathi, 2017; Mogstad, Santos, and Torgovitsky, 2018; Sasaki and Ura, 2023) which develops methods of identification, estimation, and inference for average welfare gains under counterfactual policies based on the MTEs.

The rest of this paper is organized as follows. Section 2 introduces the model. Section 3 presents the main result of representing the mean social welfare via the MTEs. Section 4 introduces applications to three statistical decision rules. Section 5 demonstrates the use of the representation result in the empirical welfare analysis. Section 6 presents an empirical application. Section 7 concludes. Appendixes A–C contain mathematical proofs and additional empirical results.

2. MODEL

We first set up the model and notations. Consider the model

$$Y = DY_1 + (1-D)Y_0, \tag{1}$$

$$D = 1\{\tilde{\nu}(Z) - \tilde{U} \geq 0\}, \tag{2}$$

where Y denotes an observed outcome variable, D denotes an observed binary treatment variable, Z denotes a vector of observed exogenous variables, Y_0 and Y_1 denote unobserved potential outcomes under no treatment and under treatment, respectively, and $\tilde{U}$ denotes an unobserved factor of the treatment selection. The first equation (1) models the outcome production through the potential outcome

[2] Although the literature on policy choices and welfare analysis is quickly growing, the use of instrumental variables is still scarce when identifying the average welfare loss. Many of the existing papers (e.g., Kitagawa and Tetenov, 2018) use an experimental or observational dataset in which the unconfoundedness assumption holds (i.e., the counterfactual outcomes are mean independent of the treatment variable given the covariates). The welfare analysis based on other identifying assumptions than the instrumental variables are found in Manski (2004, 2009), Dehejia (2005), Schlag (2007), Bhattacharya (2009, 2013), Hirano and Porter (2009), Stoye (2009, 2012), Chamberlain (2011), Bhattacharya and Dupas (2012), Tetenov (2012), Armstrong and Shen (2015), Kock and Thyrsgaard (2017), Kitagawa and Tetenov (2018, 2021), Rai (2018), Viviano (2019, 2022), Han (2023, 2020), Qiu et al. (2020), Sun (2020), Mbakop and Tabord-Meehan (2021), and Sakaguchi (2021).

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

framework, and the second equation (2) models the treatment selection via a threshold-crossing model. The function $\tilde{\nu}$ in this threshold-crossing treatment assignment model (2) is nonparametric and is unknown to the econometrician.

This model allows for endogeneity (unobserved confoundedness) in the sense that (Y_0, Y_1) and $\tilde{U}$ may be statistically dependent even conditional on Z. For the purpose of identification, therefore, we require the vector Z to contain excluded exogenous variables (i.e., excluded instruments) as well as included exogenous variables, as formally stated in Assumption 1 below. The next assumption is standard in the recent literature on MTEs (e.g., Brinch et al., 2017; Mogstad et al., 2018).[3]

Assumption 1 (Model restrictions). (1) and (2) hold, and the random vector Z can be written as $(Z_0', X')'$, where:

(i) $\tilde{U}$ and Z_0 are independent given X;
(ii) $E[Y_d \mid Z, \tilde{U}] = E[Y_d \mid X, \tilde{U}]$ and $E[Y_d^2] < \infty$;
(iii) $\tilde{U}$ is continuously distributed with a convex support conditional on X.

Part (i) concerns the treatment assignment model (2) solely, and this is the only independence assumption to be imposed on the model, implying that we can allow for an arbitrary statistical dependence between the potential outcomes (Y_0, Y_1) and $\tilde{U}$, even conditional on Z. Part (ii) states the exclusion restriction of the random subvector Z_0 of Z, and bounded second moments of the potential outcomes (Y_0, Y_1). Part (iii) rules out point masses and holes in the conditional distribution of $\tilde{U}$ given X.

For ease of analysis, by following the literature on the MTEs, we apply normalizing transformations, $U \equiv F_{\tilde{U}|X}(\tilde{U})$ and $\nu(Z) \equiv F_{\tilde{U}|X}(\tilde{\nu}(Z))$, in the threshold crossing model (2). The following lemma confirms convenient properties to be used throughout the rest of the paper, as a result of these normalizing transformations under Assumption 1.

LEMMA 1 (Normalization). *Suppose that Assumption 1(i) and (iii) hold. Then (i) $D = 1\{\nu(Z) - U \geq 0\}$, and (ii) U is distributed uniformly over* $[0,1]$ *conditional on Z.*

A proof of this lemma is provided in Appendix A.1. Consequently, we can rewrite the threshold-crossing treatment selection model (2) without loss of generality as

$$D = 1\{\nu(Z) - U \geq 0\} \text{ with } U|Z \sim \text{Uniform}(0,1). \tag{3}$$

We will hereafter consider the model (3) in place of the original model (2) by Assumption 1.

[3] We do not require the support condition for Z at the moment in this section. See the discussions in Section 5.2.2.

3. MAIN RESULT

In this section, we show that the social welfare function can be identified and represented via the MTEs as the operator kernel. To this end, we first introduce and define the two key ingredients of this result, namely the social welfare function and the MTEs.

A policymaker assigns individuals with certain observed attributes Z to a treatment $D = 1$. Thus, a treatment assignment rule is represented by a decision set $G \subset \mathcal{Z}$, where $\mathcal{Z}$ is a set of values that Z may take. Specifically, the decision set G represents the policy in which individuals with $Z \in G$ are assigned to treatment $D = 1$, while those with $Z \notin G$ are not. Let $\mathcal{G}$ denote the collection of all the decision sets G under consideration subject to the policymakers' constraints. It is worthwhile mentioning that, in many cases, the policymakers cannot, or do not want to, use all the variables in Z, particularly the excluded instrument Z_0. We can incorporate these policymakers' constraints in our framework by making $\mathcal{G}$ to be a proper subset of the powerset of $\mathcal{Z}$.

With these notations, the social welfare function $W : \mathcal{G} \to \mathbb{R}$ is defined by

$$W(G) = E[1\{Z \in G\}Y_1 + 1\{Z \notin G\}Y_0].$$

It represents the value of the mean welfare for each possible treatment assignment rule indexed by G. We can also define the MTE (Björklund and Moffitt, 1987) by

$$MTE(u,x) = E[Y_1 - Y_0 \mid U = u, X = x].$$

Recall that X is the included sub-vector of the random vector Z of exogenous variables that affect the treatment assignment (Assumption 1) and that U is the normalized unobserved factor of the treatment selection (Lemma 1 or Equation (3)). With these definitions of the social welfare function and the MTE, we now state the following theorem as the main result of this paper.

THEOREM 1 (Representation). *Under Assumption 1,*

$$W(G) = E[Y_0] + E\left[1\{Z \in G\} \int_0^1 MTE(u,X)du\right] \qquad \text{for every } G \in \mathcal{G}. \tag{4}$$

Proof. Since Y_1 and Y_0 are integrable under Assumption 1(ii), we have

$$W(G) = E[1\{Z \in G\}Y_1 + 1\{Z \notin G\}Y_0] = E[1\{Z \in G\}(Y_1 - Y_0)] + E[Y_0].$$

Now, the statement of this theorem follows from

$$\begin{aligned} E[1\{Z \in G\}(Y_1 - Y_0)] &= E[1\{Z \in G\}E[Y_1 - Y_0 \mid Z, U]] \\ &= E[1\{Z \in G\}MTE(U,X)] \\ &= E[1\{Z \in G\}E[MTE(U,X) \mid Z]] \\ &= E\left[1\{Z \in G\} \int_0^1 MTE(u,X)du\right], \end{aligned}$$

where the first equality follows from the law of iterated expectations, the second equality follows from Assumption 1(ii) and the definition of $MTE(u,x)$, the third equality follows from another application of the law of iterated expectations, and the fourth equality follows from Lemma 1 under Assumption 1(i) and (iii). □

The representation (4) of the mean social welfare via the MTEs is the key result of this paper. Because there is an existing literature on identification and estimation for the MTEs (e.g., Heckman and Vytlacil, 2001, 2005, 2007; Carneiro and Lee, 2009; Carneiro et al., 2010), our result (4) paves the way for empirical welfare analysis under the potential endogeneity or unobserved confoundedness based on the existing identification and estimation methods of the MTEs. We present a few examples of such applications in Sections 4 and 5.

There are a few remarks in order. First, we want to emphasize the distinctive roles played by X and Z that can be observed in (4). Note that only X affects the social welfare function through the heterogeneity of the MTEs. On the other hand, Z affects the social welfare function only through the proportion of treatment but not through the heterogeneity of the MTEs.

Second, we remark on relations to and differences from Kitagawa and Tetenov (2018), who use the representation

$$W(G) = E[Y_0] + E[1\{Z \in G\}\tau(X)] \text{ with } \tau(x) \equiv E[Y_1 - Y_0|X = x].$$

Our representation (4) is closely related to this representation. Under the unconfoundedness assumption, Kitagawa and Tetenov (2018) use the identification of $\tau(x)$ by $E[Y|D = 1, X = x] - E[Y|D = 0, X = x]$. On the other hand, under the unobserved confoundedness in our setup, the corresponding operator kernel $\tau(x)$ is not identified by $E[Y|D = 1, X = x] - E[Y|D = 0, X = x]$ in general. Instead, we propose to take advantage of the identification and estimation of $MTE(u,x)$ from the literature on the MTEs.

Third, our definition of W measures the welfare that will be achieved when almost all the individuals with $Z \in G$ are assigned to the treatment. This welfare measure presumes full compliance under the treatment assignment which will be rationalized by strong legal power or a large amount of resources held by policymakers. Relaxing this requirement, a more recent paper by Liu (2022) studies the social welfare that will be realized by "encouragement" as opposed to assignment.

Finally, we close this section with a discussion on a generalization of the above identification result. While we focus on the case of binary treatments, our argument straightforwardly extends to the cases of multivalued treatments. Suppose that there are $\bar{d} + 1$ treatment options indexed by $d \in \{0, 1, \ldots, \bar{d}\}$. Let Y_d denote the potential outcome under treatment d. With $d = 0$ taken as the default option, one can define the social welfare function of assigning individuals with $Z \in G_d$ to treatment d for each $d \in \{1, \ldots, \bar{d}\}$ by

$$W(G_1, \ldots, G_{\bar{d}}) = E\left[\sum_{d=1}^{\bar{d}} 1\{Z \in G_d\} Y_d + 1\left\{Z \notin \bigcup_{d=1}^{\bar{d}} G_d\right\} Y_0\right].$$

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

Using the same argument as in the proof of Theorem 1, we can rewrite it into

$$W(G_1,\ldots,G_{\bar{d}}) = \sum_{d=1}^{\bar{d}} E[1\{Z\in G_d\}(Y_d - Y_0)] + E[Y_0].$$

Using the same argument as in the proof of Theorem 1 again, each summand in the above expression can be written in turn as

$$E[1\{Z\in G_d\}(Y_d - Y_0)] = E\left[1\{Z\in G_d\}\int_{[0,1]^{\bar{d}}} MTE^{(0,d)}(u,X)du\right],$$

where $MTE^{(0,d)}(u,X) = E[Y_d - Y_0|U = u,X]$ denotes the MTE identified by Lee and Salanié (2018).

4. APPLICATIONS TO STATISTICAL DECISION RULES

Once we obtain the representation (4) of the mean social welfare via $MTE(u,x)$, we may apply it to a variety of policymakers' statistical decision problems for treatment choice. In this section, following Hirano and Porter (2020, Sect. 2.3), we introduce applications to the three popular statistical decision rules: plug-in rules, Bayes rules, and empirical welfare maximization rules. In Section 5, we discuss the empirical welfare maximization rules in further detail based on recent technologies.

4.1. Plug-in Rules

Suppose that the distribution of Z is parametrized by δ and we have an estimator $(\hat{\delta},\widehat{MTE})$ for (δ,MTE). In this case, we can estimate the maximizer for the population social welfare by maximizing

$$\int 1\{z\in G\}\int_0^1 \widehat{MTE}(u,x)dud\mu_{\hat{\delta}}(z)$$

over $G\in\mathcal{G}$, where $\mu_\delta(\cdot)$ is the probability measure of Z indexed by δ.[4] Here, we may ignore the term $E[Y_0]$ in (4) since it does not affect the maximization problem over $G\in\mathcal{G}$. Note that $(\hat{\delta},\widehat{MTE})$ is an estimator, and so a maximizer of the above objective function is a statistical decision rule, i.e., it is a function of the observed data.

4.2. Bayes Rules

Suppose that the distribution of Z is parametrized by δ and MTE is parametrized by η, and that we have a prior probability measure of (δ,η). We can construct

[4]When $\mathcal{G}$ is the power set of $\mathcal{Z}$, the plug-in rule is a set G such that $\{x : \int_0^1 \widehat{MTE}(u,x)du > 0\} \subset G$ and $\{x : \int_0^1 \widehat{MTE}(u,x)du < 0\} \cap G = \emptyset$. In other words, the plug-in rule is basically the set $\{x : \int_0^1 \widehat{MTE}(u,x)du \geq 0\}$ as in Kitagawa and Tetenov (2018, eqn. (1.13)). When $\mathcal{G}$ is a proper subset of the power set of $\mathcal{Z}$, however, the plug-in rule is not equal to the set $\{x : \int_0^1 \widehat{MTE}(u,x)du \geq 0\}$ in general.

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

the posterior distribution, denoted by $\pi_{\text{posterior}}$, of (δ,η) by Bayesian updating. Given the posterior probability measure of (δ,η) and our representation (4), we can construct the Bayes welfare

$$\int \left(\int 1\{z \in G\} \int_0^1 MTE_\eta(u,x)dud\mu_\delta(z) \right) d\pi_{\text{posterior}}(\delta,\eta).$$

The Bayes rule is the maximizer of the Bayes welfare over $G \in \mathcal{G}$.

4.3. Empirical Welfare Maximization Rules

The empirical welfare maximization rule uses the empirical distribution of Z and an estimator $\widehat{MTE}$ for MTE. Namely, with a random sample $\{Z_1,\ldots,Z_n\}$ of size n, we can define the empirical welfare by

$$E_n\left[1\{Z \in G\} \int_0^1 \widehat{MTE}(u,X)du \right],$$

where E_n denotes the sample average operator, i.e., $E_n f(Y,D,Z) = n^{-1} \sum_{i=1}^n f(Y_i,D_i,Z_i)$ for any measurable function f. The empirical welfare maximization rule selects the maximizer of this empirical welfare over $G \in \mathcal{G}$. Section 5 presents more detailed analyses of the asymptotic properties of the maximum of this empirical welfare relative to the population mean welfare under the oracle action.

Remark 1. In the plug-in rule and the empirical welfare maximization rule, one may use an estimator $\widehat{MTE}$ for MTE based on observations from a different population from the policy target population, as long as this estimator needs to satisfy external validity. In other words, we can use observations from a different population to obtain $\widehat{MTE}$ provided that the two populations share the same MTEs.

Remark 2. The three rules are not mutually exclusive. The empirical welfare maximization rule can be a special case of the plug-in rule where an empirical distribution of Z is used as $\mu_{\hat{\delta}}$. The plug-in rule is a Bayes rule where the posterior distribution $\pi_{\text{posterior}}(\delta,\eta)$ is degenerate at $(\hat{\delta},\hat{\eta})$ with $\widehat{MTE} = MTE_{\hat{\eta}}$. When the posterior distribution $\pi_{\text{posterior}}(\delta,\eta)$ satisfies the independence between δ and η, the Bayes rule can be interpreted as a plug-in rule where the posterior means for (MTE_η,μ_δ) are used as an estimator for (δ, MTE).

5. APPLICATIONS TO EMPIRICAL WELFARE MAXIMIZATION

We demonstrate applications of the representation result (4) to empirical welfare maximization in this section. For the purpose of exposition of the core idea, we first focus on the case where the mapping $(u,x) \mapsto MTE(u,x)$ is known by a researcher in Section 5.1. We then present the case where the mapping $(u,x) \mapsto MTE(u,x)$ is unknown by a researcher and thus needs to be estimated in Section 5.2.

5.1. Empirical Welfare Maximization with Known Marginal Treatment Effects

In this section, we assume that we know the mapping $(u,x) \mapsto MTE(u,x)$. The empirical welfare maximizer in this setting is given by

$$\hat{G}_{EWM} \in \arg\max_{G\in\mathcal{G}} E_n\left[1\{Z\in G\}\int_0^1 MTE(u,X)du\right].$$

The population mean social welfare under this $\hat{G}_{EWM}$ is $W(\hat{G}_{EWM})$. We present a uniform asymptotic analysis of $W(\hat{G}_{EWM})$ relative to $\sup_{G\in\mathcal{G}} W(G)$, i.e., the population mean social welfare under the oracle action. To this end, consider the following assumption.

Assumption 2. (i) $|\int_0^1 MTE(u,X)du| \leq \bar{M} < \infty$ a.s. for the class $\mathcal{P}(\bar{M})$ of distributions of (Y_0,Y_1,D,Z). (ii) $\mathcal{G}$ has a finite VC-dimension $v < \infty$ and is countable.

The above assumption is a modification of Assumption 2.1(BO)–(VC) in Kitagawa and Tetenov (2018) tailored to our framework with the MTEs.[5] Assumption 2(i) requires a bounded integral of the marginal treatment effect function. As a sufficient condition, it holds when the outcome variable is bounded by some constant which is naturally satisfied in some applications. Assumption 2(ii) restricts the complexity of the class of treatment functions $Z \mapsto 1\{Z\in G\}$. As a sufficient condition, when X has a finite support, this assumption will automatically hold where v is the cardinality of the power set for the support of X. Kitagawa and Tetenov (2018, p. 598) collects a few examples of $\mathcal{G}$ with finite VC dimensions.

The following corollary to Theorem 1 provides a convergence rate of the worst-case average welfare loss (regret) by the empirical welfare maximization.

COROLLARY 1. *Under Assumptions 1 and 2, one has*

$$\sup_{P\in\mathcal{P}(\bar{M})} E_{P^n}\left[\sup_{G\in\mathcal{G}} W(G) - W(\hat{G}_{EWM})\right] \leq 2C_1\bar{M}\sqrt{\frac{v}{n}},$$

where C_1 is a universal constant.

A proof is provided in Appendix A.2. Corollary 1 shows that the worst-case average welfare loss (regret) by the empirical welfare maximization converges to zero at rates no slower than $n^{-1/2}$ uniformly over data generating processes. This corollary extends and is a counterpart of Theorem 2.1 of Kitagawa and Tetenov (2018).

[5] For the countability assumption, see (cf. Kitagawa and Tetenov, 2018, fn 4).

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

5.2. Empirical Welfare Maximization with Unknown Marginal Treatment Effects

In this section, we consider the case where the mapping $(u,x) \mapsto MTE(u,x)$ is unknown by a researcher and thus needs to be estimated from empirical data. The empirical welfare maximizer in this setting is given by

$$\hat{G}_{hybrid} \in \arg\max_{G \in \mathcal{G}} E_n \left[1\{Z \in G\} \int_0^1 \widehat{MTE}(u,X)du \right],$$

where $\widehat{MTE}$ is an estimator for MTE. The existing literature on MTEs provides a list of alternative estimators $\widehat{MTE}$ for MTE. We therefore first provide a general sufficient condition in terms of $\widehat{MTE}$ that accommodates a wide range of possible estimators in Section 5.2.1. This will be followed up by a specific estimator $\widehat{MTE}$ with lower-level primitive conditions tailored to it in Section 5.2.2.

5.2.1. *A Sufficient Condition.* We consider the following general high-level assumption about an estimator $\widehat{MTE}$ for MTE.

Assumption 3. For a class of data generating processes $\mathcal{P}_m$, there exists a sequence $\psi_n \to \infty$ such that

$$\limsup_{n\to\infty} \sup_{P \in \mathcal{P}_m} \psi_n E_{P^n} \left[E_n \left[\left| \int_0^1 (\widehat{MTE}(u,X) - MTE(u,X))du \right| \right] \right] < \infty.$$

We can interpret the above condition as the ψ_n^{-1}-consistency of the estimator $\widehat{MTE}$, where ψ_n^{-1} can be slower than the parametric convergence rate of $n^{-1/2}$. This condition leads to the rate, $\psi_n^{-1} \vee n^{-1/2}$, of convergence for the worst-case average welfare loss (regret), as formally stated as a corollary to Theorem 1 below.

Corollary 2. *Under Assumptions 1–3,*

$$\sup_{P \in \mathcal{P}_m \cap \mathcal{P}(\bar{M})} E_{P^n} \left[\sup_{G \in \mathcal{G}} W(G) - W(\hat{G}_{hybrid}) \right] = O(\psi_n^{-1} \vee n^{-1/2}).$$

A proof of this corollary is provided in Appendix A.3. It serves as a counterpart of Theorem 2.5 of Kitagawa and Tetenov (2018). Different estimators $\widehat{MTE}$ for MTE in general entail different convergence rates. We next present a concrete estimator with lower-level sufficient conditions for the high-level condition in Assumption 3.

5.2.2. *Parametric Estimation for the Marginal Treatment Effects.* By Heckman and Vytlacil (1999, 2001, 2005), the MTEs can be identified from data via

$$MTE(u,x) = \frac{\partial E[Y \mid \nu(Z) = u, X = x]}{\partial u}.$$

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

As mentioned in the introductory section, one of the main advantages of the MTEs is that the ATEs and thus the social welfare functions can be identified via functional-form restrictions on the MTEs even if instruments induce insufficient variations in the treatment choice probabilities. In this light, we consider a parametric model for $E[Y \mid \nu(Z) = u, X = x]$:

$$E[Y \mid \nu(Z) = u, X = x] = \xi(x)'\beta_0 + \xi(x)'(\beta_1 - \beta_0)u + \sum_{k=1}^{K} \alpha_k \phi_k(u), \tag{5}$$

where $\xi(x)$ is a vector-valued function of x (e.g., $\xi(x) = (1, x, x^2, x^3)'$ for a cubic specification in scalar x), K is a fixed integer, and $\phi_1, \ldots, \phi_K$ are functions that a researcher chooses. Let $\hat{\nu}(Z)$ denote some estimator of the propensity score $\nu(Z)$, and define

$$\begin{aligned}
\mathcal{X} &= ((1-\nu(Z))\xi(X)', \nu(Z)\xi(X)', \phi_1(\nu(Z)), \ldots, \phi_K(\nu(Z)))', \\
\hat{\mathcal{X}} &= ((1-\hat{\nu}(Z))\xi(X)', \hat{\nu}(Z)\xi(X)', \phi_1(\hat{\nu}(Z)), \ldots, \phi_K(\hat{\nu}(Z)))', \\
\theta &= (\beta_0', \beta_1', \alpha_1, \ldots, \alpha_K).
\end{aligned}$$

Let $\hat{\theta} = (\hat{\beta}_0', \hat{\beta}_1', \hat{\alpha}_1, \ldots, \hat{\alpha}_K)'$ be the OLS estimator for θ by regressing Y on $\hat{\mathcal{X}}$, that is,

$$\hat{\theta} = E_n\left[\hat{\mathcal{X}}\hat{\mathcal{X}}'\right]^{-1} E_n\left[\hat{\mathcal{X}} Y\right].$$

Then, the MTEs can be simply estimated by the following linear functional of $\hat{\theta}$:

$$\widehat{MTE}(u,x) = \xi(x)'(\hat{\beta}_1 - \hat{\beta}_0) + \sum_{k=1}^{K} \hat{\alpha}_k \frac{d}{du}\phi_k(u).$$

Therefore, the operator kernel in our representation (4) can be estimated by the simple linear expression

$$\int_0^1 \widehat{MTE}(u,x)du = \xi(x)'(\hat{\beta}_1 - \hat{\beta}_0) + \sum_{k=1}^{K} \hat{\alpha}_k(\phi_k(1) - \phi_k(0)).$$

For this concrete estimator, we provide a set of lower-level conditions in the proposition below that guarantee the aforementioned high-level condition in Assumption 3 to be satisfied.

PROPOSITION 1. *Let C and c be positive constants, and let ψ_n be a sequence with $\psi_n \le n^{1/2}$. Suppose that ϕ_k is differentiable with $\sup_{u\in[0,1]} \left|\frac{d}{du}\phi_k(u)\right| \le C$, for every $k = 1, \ldots, K$, and that the parameter space for θ is compact so that for sufficiently large n,*

$$\|\hat{\theta}\| + \|\theta\| \le C \text{ almost surely.} \tag{6}$$

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

Furthermore, suppose that $\mathcal{P}_m$ is a class of data generating processes such that

$$\limsup_{n\to\infty} \sup_{P\in\mathcal{P}_m} \psi_n E_{P^n}\left[\max_{i=1,\dots,n} |\hat{\nu}(Z_i)-\nu(Z_i)|^2\right]^{1/2} < \infty, \quad (7)$$

$$\max\{E\left[\|\xi(X)\|^4\right], E\left[|Y|^4\right]\} < C, \quad (8)$$

$$\lambda_{\min}\left(E\left[\mathcal{X}\mathcal{X}'\right]\right) \geq c. \quad (9)$$

Then, Assumption 3 is satisfied.

A proof is provided in Appendix A.4. The condition in (7) requires a convergence rate for an estimator $\hat{\nu}(z)$ of the propensity score $\nu(z)$ uniformly over the data generating processes. This condition can be checked with specific propensity score estimators. For example, we can use the local polynomial estimator $\hat{\nu}(z)$ for $\nu(z)$, for which Kitagawa and Tetenov (2018, Appendix H) derive a uniform convergence rate. Specifically, the convergence in (7) follows directly from their Lemma E.4(ii). For another example, one could consider a linear propensity score model $\hat{\nu}(z) = p(z)'\hat{\gamma}$ and its least squares estimator $\hat{\nu}(z) = p(z)'\hat{\gamma}$. In this case, (7) can be satisfied with $\psi_n = \sqrt{n}$, so that the convergence rate for the worst-case average welfare loss (regret) in (2) holds with the parametric root n rate.

The condition in (9) is the standard identification condition to rule out the multicollinearity among the elements in $\mathcal{X}$. Nonparametric global identification of the MTEs over the entire domain requires a continuous instrument Z with the full support of the distribution of $\nu(Z)$, i.e., identification at infinity. With the parametric specification (5), on the other hand, the MTEs are globally identified via the parametric extrapolation along with the identification condition (9) to rule out the collinearity. In light of the rare availability of continuous instruments with full support, this approach is suggested by Cornelissen et al. (2016, Sect. 4.3) in their survey article for labor economists.

6. EMPIRICAL ILLUSTRATION

We present an empirical application of our proposed method to data from the National Job Training Partnership Act Study. In this experimental study, applicants were randomly assigned to treatment and control groups, where those individuals assigned to the treatment group were eligible to receive job training programs for 18 months. Through this experiment along with information about observed attributes of the individuals, researchers can evaluate the benefits and costs of this job training program in terms of labor market outcomes. We refer readers to Bloom et al. (1997) for details of this dataset, as well as Heckman, Ichimura, and Todd (1997) and Abadie, Angrist, and Imbens (2002) for program evaluation studies using this dataset.

It is generally infeasible to identify and estimate the ATEs when a binary instrument induces only partial variations in the treatment choice probabilities.

If one employs the MTE with parametric functional-form restrictions, however, we can identify and estimate the ATEs and thus the social welfare as well. We motivate our framework with the practical advantage of the MTEs in the Introduction and in Section 5.2.2 where we focus on a parametric model of the MTEs. The current real data example with a binary instrument of eligibility illustrates this case in point.

Kitagawa and Tetenov (2018) and Byambadalai (2021) conduct welfare analyses using the same data set. Kitagawa and Tetenov (2018) consider the empirical welfare maximization with the 30-month earnings after the random assignment as the outcome, the random assignment as the treatment, and years of education and pre-program annual earnings as two controls. Since some individuals assigned to the treatment group did not participate in the training program, the welfare measure of Kitagawa and Tetenov (2018) is defined and interpreted from the intention-to-treat perspective. Using an instrumental variable, Byambadalai (2021) studies welfare gains and losses based on the actual exposure to the training program as the treatment, where the random assignment in turn plays the role of an instrument. Following up with these papers, we consider the empirical welfare maximization (as in Kitagawa and Tetenov, 2018) but with the treatment defined by the actual exposure to the training program for which we use the random assignment as an instrument (as in Byambadalai, 2021).

Following these two benchmark studies, we focus on the subsample of adults for whom information about the 30-month earnings after the random assignment (Y), the actual exposure to the training program (D), the random assignment to treatment and control groups (Z_0), years of education (X_1), and pre-program annual earnings (X_2) are available. We thus obtain a sample of 9,223, exactly as in Kitagawa and Tetenov (2018) and Byambadalai (2021). Our summary statistics including the instrument also coincide with those presented in Byambadalai (2021).

We estimate the propensity score by

$$\hat{\nu}(Z) = \begin{cases} 0, & \text{if } Z_0 = 0, \\ \widehat{E}[D|X, Z_0 = 1], & \text{if } Z_0 = 1, \end{cases}$$

where $\widehat{E}[D|X, Z_0 = 1]$ is the least squares predictor of the projection of D on the pre-program annual earnings and L powers of the years of education (cf., Kitagawa and Tetenov, 2018). We will compute the empirical welfare maximizing policies and welfare gains under various values of L, and will demonstrate that the results are insensitive to the choice of L. Note that we set $\hat{\nu}(Z) = 0$ when $Z_0 = 0$, because those individuals assigned to the control group have no eligibility to participate in the training program according to the rule of the experiment.

We estimate the marginal treatment effect using the procedure introduced in Section 5.2.2. Specifically, we consider

$$E[Y \mid \nu(Z) = u, X = x] = \xi(x)'\beta_0 + \xi(x)'(\beta_1 - \beta_0)u + \sum_{k=1}^{K} \alpha_k \phi_k(u),$$

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

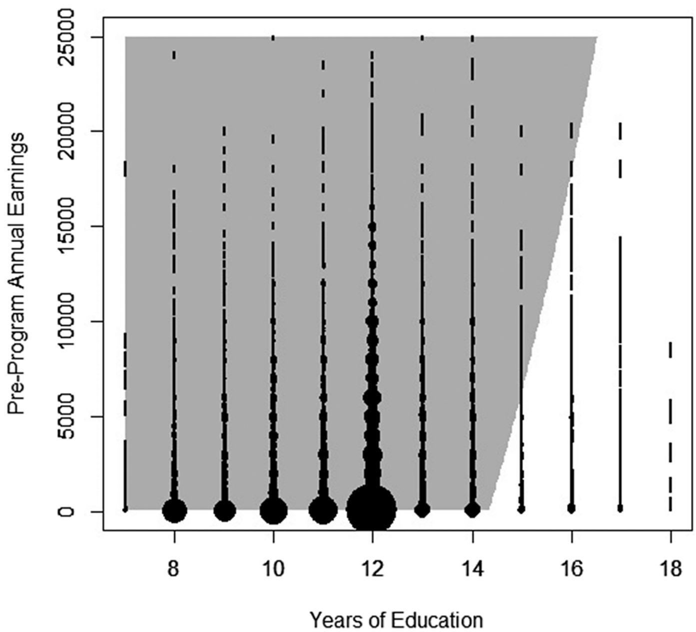

FIGURE 1. Empirical welfare maximization rule with no treatment cost. The area of black circles represents the density of the data at the location of the shape. We use $K = 10$ for the degree of nonlinearity in estimating the marginal treatment effects, and $L = 10$ for the degree of the years of education in estimating the propensity score.

where $\xi(X) = (1, X_1, X_1^2, X_1^3, X_2)'$, $\{\phi_k\}_{k=1}^K$ consists of a finite subset of the Fourier basis, i.e., $\phi_1(u) = \sin(2\pi u)$, $\phi_2(u) = \cos(2\pi u)$, $\phi_3(u) = \sin(\pi u)$, $\phi_4(u) = \cos(\pi u)$, $\phi_5(u) = \sin(2\pi u/3)$, $\phi_6(u) = \cos(2\pi u/3)$, and so on. We will compute the empirical welfare maximizing policies and welfare gains under various values of K, and will demonstrate that the results are insensitive to the choice of K as well as that of L.

The gray shade in Figure 1 represents the empirical welfare maximization rule $\hat{G}_{hybrid}$ over the power set $\mathcal{G}$ of $\{7, ..., 18\} \times \{0, ..., 25,000\}$ based on the full sample.[6] The area of a black circle represents the density of the data at the location of the shape. We use $K = 10$ and $L = 10$ to generate this particular graph, but the results are insensitive to the choice of K and L—see Appendix C for sensitivity analyses. Since the training program costs \$774 for each participant, we also repeat the same computational procedure with the outcome variable Y defined as the 30-month earnings after the random assignment subtracted by this cost, \$774. The gray shade

[6]In this problem, we do not need to search for all the $2^{275,000}$ subsets. Note that, without a constraint, a maximizing set can be characterized by the set of all the positive points. The gray shade represents this set.

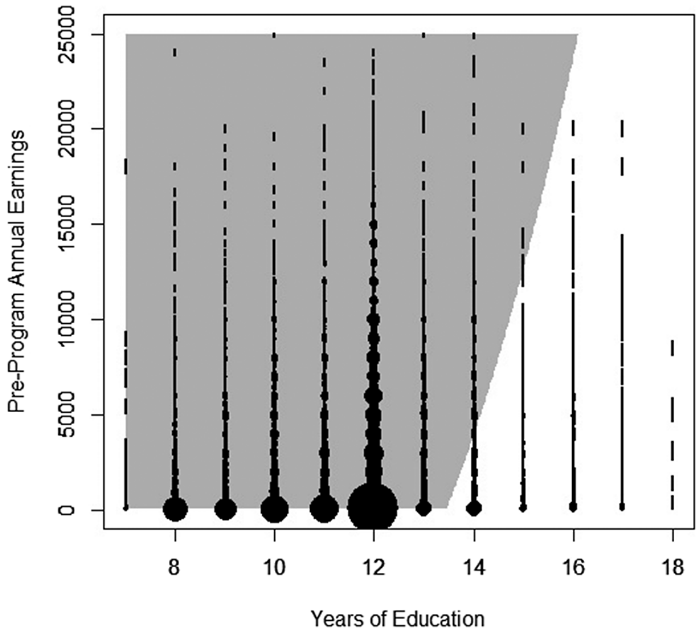

FIGURE 2. Empirical welfare maximization rule with the treatment cost of \$774 per assignee. The area of black circles represents the density of the data at the location of the shape. We use $K = 10$ for the degree of nonlinearity in estimating the marginal treatment effects, and $L = 10$ for the degree of the years of education in estimating the propensity score.

in Figure 2 illustrates the empirical welfare maximization rule accounting for these program costs.

In each of Figures 1 and 2, observations in the top-left part of the figure are selected by the empirical welfare maximizing treatment assignment rule. These results are analogous to those of Kitagawa and Tetenov (2018), and thus reinforce their conclusions even if we consider the actual exposure to the training program as the outcome based on which to measure the welfare. These results are intuitive, as adults in this part of the graph tend to have higher levels of unobserved abilities as reflected by the pre-program annual earnings despite not having received higher levels of education.

Table 1 summarizes the estimated welfare gains from empirical welfare maximization treatment assignment rules.[7] As a benchmark, we also display the results under the simplistic policy of treating everybody. Under this benchmark policy, the estimated welfare gains per population member are \$1,857–\$1,875 without accounting for the training costs, and \$1,121–\$1,239 after accounting for the

[7]Using the full sample to estimate both the policy and the welfare gains will incur over-fitting biases. To avoid such biases, we use the twofold sample splitting with a fixed seed.

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

Table 1. Estimated welfare gains from empirical welfare maximization (EWM) with assignment policies based on years of education and pre-program annual earnings. K denotes the degree of nonlinearity in estimating the marginal treatment effects, and L denotes the degree of the years of education in estimating the propensity score.

Outcome variable:			30-month post-program earnings No treatment cost		30-month post-program earnings $774 cost for each assigned treatment	
Treatment rule:	K	L	Share of population to be treated	Estimated welfare gain per population member	Share of population to be treated	Estimated welfare gain per population member
Treat everyone	5	5	1.00	$1,861	1.00	$1,125
Treat everyone	5	10	1.00	$1,875	1.00	$1,139
Treat everyone	10	5	1.00	$1,857	1.00	$1,121
Treat everyone	10	10	1.00	$1,871	1.00	$1,135
EWM	5	5	0.91	$1,996	0.85	$1,339
EWM	5	10	0.91	$1,996	0.85	$1,339
EWM	10	5	0.91	$1,992	0.85	$1,336
EWM	10	10	0.91	$1,992	0.85	$1,336

training costs. On the other hand, the empirical welfare maximizing treatment assignment rule would engender the estimated welfare gains per population of $1,992–$1,996 without accounting for the training costs, and $1,336–$1,339 after accounting for the training costs.

A few remarks are in order regarding these estimation results. First, these welfare gains are larger than those reported by Kitagawa and Tetenov (2018). It is probably because we analyze welfare gains as a consequence of assigning individuals to actual training, whereas Kitagawa and Tetenov (2018) analyze welfare gains as a consequence of granting eligibility to participate in the training program. Second, the results are insensitive to the choice of K and L in terms of both the shares of population to be treated and the estimated welfare gains. This feature demonstrates the robustness of the results against choices of specific estimation procedures in this particular application. Third, we use the power set for $\mathcal{G}$ of feasible decision rules G for this illustration. In general, however, we may also incorporate policy-relevant constraints in $\mathcal{G}$. One may also use regularizing constraints, such as the set of linear eligibility score rules considered by Kitagawa and Tetenov (2018).

7. CONCLUSION

An important research goal for empirical economists is to provide policymakers with guidance on how heterogeneous individuals can be assigned to a treatment

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

under consideration based on evidence from empirical data. To this goal, it is essential to identify a social welfare function from observational data. For many observational datasets used in empirical economic research, treatments are likely to be endogenously selected by rational agents. Furthermore, the effects of these treatments can be heterogeneous even after controlling for observed attributes. In this light, given the abilities of the MTEs to measure heterogeneous treatment effects, we propose the usage of the MTEs for identifying the mean social welfare function in the presence of unobserved heterogeneity in treatment effects while accounting for endogenous treatment selection in the empirical data.

Our main result, Theorem 1, establishes that the mean social welfare can be represented via the MTEs as the operator kernel. We apply this main result to a few of policymakers' statistical decision problems, such as the plug-in rule, the Bayes rule, and the empirical welfare maximization rule. Focusing on the empirical welfare maximization in particular, we derive convergence rates of the worst-case average welfare loss (regret) from the maximum empirical welfare under alternative scenarios. The proposed representation in Theorem 1 has an immediate benefit for empirical welfare analyses, because we can apply the existing machinery developed for the MTEs to the variety of empirical welfare analysis.

Our theoretical results in Section 5 focus on the empirical welfare maximization rules. We do not have results for a general class of plug-in rules or a general class of Bayes rules, which we left for future research.

APPENDICES

A. PROOFS

A.1. Proof of Lemma 1

Proof. The first statement of this lemma follows because Assumption 1(iii) implies that $F_{\tilde{U}|X=x}$ is strictly increasing and therefore

$$D = 1\{\tilde{\nu}(Z) - \tilde{U} \geq 0\} = 1\{F_{\tilde{U}|X}(\tilde{\nu}(Z)) - F_{\tilde{U}|X}(\tilde{U}) \geq 0\} = 1\{\nu(Z) - U \geq 0\}.$$

The second statement follows because Assumption 1(i) and (iii) implies

$$P(U \leq u \mid Z) = P(F_{\tilde{U}|X}(\tilde{U}) \leq u \mid Z) = P(\tilde{U} \leq F^{-1}_{\tilde{U}|X}(u) \mid Z) = P(\tilde{U} \leq F^{-1}_{\tilde{U}|X}(u) \mid X) = u.$$ □

A.2. Proof of Corollary 1

Proof. Define the function f by

$$f(Z;G) = 1\{Z \in G\} \int_0^1 MTE(u,X)du.$$

Let $\mathcal{F} = \{f(\,\cdot\,;G) : G \in \mathcal{G}\}$. Then, $\mathcal{F}$ is a class of uniformly bounded functions with $\|f\|_\infty \leq \bar{M}$ for all $f \in \mathcal{F}$. From Assumption 2, it follows that $\mathcal{F}$ is of a VC subgraph class with VC

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

dimension $v < \infty$. By Kitagawa and Tetenov (2018, Lem. A.4),

$$E_{P^n}\left[\sup_{f\in\mathcal{F}} |E_n[f]-E_P[f]|\right] \leq C_1\bar{M}\sqrt{\frac{v}{n}},$$

where C_1 is a universal constant defined in Kitagawa and Tetenov (2018, Lem. A.4). Defining

$$W_n(G) = E[Y_0] + E_n\left[1\{Z\in G\}\int_0^1 MTE(u,X)du\right],$$

we have

$$\sup_{P\in\mathcal{P}(\bar{M})} E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n(G)-W(G)|\right] = \sup_{P\in\mathcal{P}(\bar{M})} E_{P^n}\left[\sup_{f\in\mathcal{F}} |E_n[f]-E_P[f]|\right] \leq C_1\bar{M}\sqrt{\frac{v}{n}}. \tag{A.1}$$

Following the derivations in Kitagawa and Tetenov (2018, eqn. (2.2)), we have, for any $\tilde{G}\in\mathcal{G}$, that

$$\begin{aligned} W(\tilde{G}) - W(\hat{G}_{EWM}) &= W(\tilde{G}) - W_n(\hat{G}_{EWM}) + W_n(\hat{G}_{EWM}) - W(\hat{G}_{EWM}) \\ &\leq W(\tilde{G}) - W_n(\tilde{G}) + W_n(\hat{G}_{EWM}) - W(\hat{G}_{EWM}) \\ &\leq 2\sup_{G\in\mathcal{G}} |W_n(G)-W(G)|, \end{aligned} \tag{A.2}$$

where the first inequality uses $W_n(\hat{G}_{EWM}) \geq W_n(\tilde{G})$. Therefore,

$$\sup_{P\in\mathcal{P}(\bar{M})} E_{P^n}\left[\sup_{G\in\mathcal{G}} W(G) - W(\hat{G}_{EWM})\right] \leq 2C_1\bar{M}\sqrt{\frac{v}{n}}$$

follows from (A.1) and (A.2). □

A.3. Proof of Corollary 2

Proof. Define

$$\hat{W}_n(G) = E[Y_0] + E_n\left[1\{Z\in G\}\int_0^1 \widehat{MTE}(u,X)du\right].$$

Following the derivations in Kitagawa and Tetenov (2018, eqn. (A.29)), we obtain, for any $\tilde{G}\in\mathcal{G}$, that

$$\begin{aligned} W(\tilde{G}) - W(\hat{G}_{hybrid}) &= W_n(\tilde{G}) - \hat{W}_n(\tilde{G}) - W_n(\hat{G}_{hybrid}) + \hat{W}_n(\hat{G}_{hybrid}) \\ &\quad + W(\tilde{G}) - W_n(\tilde{G}) + W_n(\hat{G}_{hybrid}) - W(\hat{G}_{hybrid}) \\ &\quad + \hat{W}_n(\tilde{G}) - \hat{W}_n(\hat{G}_{hybrid}) \\ &\leq W_n(\tilde{G}) - \hat{W}_n(\tilde{G}) - W_n(\hat{G}_{hybrid}) + \hat{W}_n(\hat{G}_{hybrid}) \\ &\quad + W(\tilde{G}) - W_n(\tilde{G}) + W_n(\hat{G}_{hybrid}) - W(\hat{G}_{hybrid}) \end{aligned}$$

$$\leq E_n\left[(1\{Z\in\hat{G}_{hybrid}\}-1\{Z\in\tilde{G}\})\int_0^1(\widehat{MTE}(u,X)-MTE(u,X))du\right]$$
$$+2\sup_{G\in\mathcal{G}}|W_n(G)-W(G)|, \quad \textbf{(A.3)}$$

where the first inequality uses $\hat{W}_n(\tilde{G})\leq\hat{W}_n(\hat{G}_{hybrid})$. By Assumption 3, we have

$$\sup_{P\in\mathcal{P}_m\cap\mathcal{P}(\bar{M})}E_{P^n}\left[E_n\left[(1\{Z\in\hat{G}_{hybrid}\}-1\{Z\in\tilde{G}\})\int_0^1(\widehat{MTE}(u,X)-MTE(u,X))du\right]\right]$$
$$=O(\psi_n^{-1}). \quad \textbf{(A.4)}$$

By (A.1) in the proof of Theorem 1, we have

$$\sup_{P\in\mathcal{P}_m\cap\mathcal{P}(\bar{M})}E_{P^n}\left[\sup_{G\in\mathcal{G}}|W_n(G)-W(G)|\right]=O(n^{-1/2}). \quad \textbf{(A.5)}$$

The claim of the theorem now follows from (A.3)–(A.5). □

A.4. Proof of Proposition 1

Proof. From on the structure of $MTE(u,x)$, we have

$$E_{P^n}\left[E_n\left[\left|\int_0^1(\widehat{MTE}(u,X)-MTE(u,X))du\right|\right]\right]$$
$$=E_{P^n}\left[E_n\left[\left|\xi(X)'(\hat{\beta}_1-\beta_1)-\xi(X)'(\hat{\beta}_0-\beta_0)+\sum_{k=1}^K(\hat{\alpha}_k-\alpha_k)(\phi_k(1)-\phi_k(0))\right|\right]\right]$$
$$\leq E_{P^n}\left[E_n[\|\xi(X)\|](\|\hat{\beta}_1-\beta_1\|+\|\hat{\beta}_0-\beta_0\|)\right]+C\sum_{k=1}^K E_{P^n}\left[|\hat{\alpha}_k-\alpha_k|\right]$$
$$\leq E_{P^n}\left[(E_n-E)[\|\xi(X)\|](\|\hat{\beta}_1-\beta_1\|+\|\hat{\beta}_0-\beta_0\|)\right]$$
$$+E[\|\xi(X)\|]E_{P^n}\left[(\|\hat{\beta}_1-\beta_1\|+\|\hat{\beta}_0-\beta_0\|)\right]+C\sum_{k=1}^K E_{P^n}\left[|\hat{\alpha}_k-\alpha_k|\right],$$

where C is used to bound $|\phi_k(1)-\phi_k(0)|$. By (6) and (8), it suffices to show

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m}\psi_n E_{P^n}\left[\left\|\hat{\theta}-\theta\right\|\right]<\infty.$$

Since $E_n\left[\hat{\mathcal{X}}(Y-\hat{\mathcal{X}}'\hat{\theta})\right]=0$ and $E\left[\mathcal{X}(Y-\mathcal{X}'\theta)\right]=0$, we can write

$$E\left[\mathcal{X}\mathcal{X}'\right](\hat{\theta}-\theta)=-(E_n-E)\left[\mathcal{X}\mathcal{X}'\right]\hat{\theta}+(E_n-E)[\mathcal{X}Y]-E_n\left[\hat{\mathcal{X}}\hat{\mathcal{X}}'-\mathcal{X}\mathcal{X}'\right]\hat{\theta}+E_n\left[(\hat{\mathcal{X}}-\mathcal{X})Y\right].$$

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

Therefore,

$$\begin{aligned}\left\|\hat{\theta}-\theta\right\| &\leq \lambda_{\min}(E[\mathcal{X}\mathcal{X}'])^{-1}\left\|(E_n-E)[\mathcal{X}\mathcal{X}']\right\|_2\left\|\hat{\theta}\right\| \\ &+\lambda_{\min}(E[\mathcal{X}\mathcal{X}'])^{-1}\left\|(E_n-E)[\mathcal{X}Y]\right\| \\ &+\lambda_{\min}(E[\mathcal{X}\mathcal{X}'])^{-1}\left\|E_n\left[\hat{\mathcal{X}}\hat{\mathcal{X}}'-\mathcal{X}\mathcal{X}'\right]\right\|_2\left\|\hat{\theta}\right\| \\ &+\lambda_{\min}(E[\mathcal{X}\mathcal{X}'])^{-1}\left\|E_n\left[(\hat{\mathcal{X}}-\mathcal{X})Y\right]\right\|.\end{aligned}$$

Therefore, the statement of this theorem follows from (6), (9), Lemma 2, and Lemma 3, where the last two lemmas are stated and proved in Appendix B. □

B. AUXILIARY LEMMAS FOR THE PROOF OF PROPOSITION 1

LEMMA 2. *Under the assumptions of Proposition 1, one has*

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} n^{1/2}E_{P^n}\left[(E_n-E)\left[\|\xi(X)\|\right]\right]<\infty,$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} n^{1/2}E_{P^n}\left[\|(E_n-E)[\mathcal{X}Y]\|\right]<\infty,$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} n^{1/2}E_{P^n}\left[\left\|(E_n-E)\left[\mathcal{X}\mathcal{X}'\right]\right\|_2\right]<\infty,$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} n^{1/2}E_{P^n}\left[\left|(E_n-E)\left[\|\xi(X)\|^2\right]\right|\right]<\infty,$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} n^{1/2}E_{P^n}\left[\left|(E_n-E)\left[|Y|^2\right]\right|\right]<\infty.$$

Proof. The statements follow from evaluating the second moment for E_n-E for each random variable. The second moments for these variables are bounded uniformly over $P\in\mathcal{P}_m$ by (8). □

LEMMA 3. *Under the assumptions of Proposition 1, one has*

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} \psi_n E_{P^n}\left[\left\|E_n\left[\hat{\mathcal{X}}\hat{\mathcal{X}}'-\mathcal{X}\mathcal{X}'\right]\right\|_2\right]<\infty,$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} \psi_n E_{P^n}\left[\left\|E_n\left[(\hat{\mathcal{X}}-\mathcal{X})Y\right]\right\|\right]<\infty.$$

Proof. By definition of $\mathcal{X}$ and $\hat{\mathcal{X}}$, we have

$$\|\hat{\mathcal{X}}\hat{\mathcal{X}}'-\mathcal{X}\mathcal{X}'\|_2\leq C_2(\|\xi(X)\|^2+1)|\hat{\nu}(Z)-\nu(Z)|$$

for a positive constant $C_2<\infty$. Moreover, we have

$$\left\|\hat{\mathcal{X}}-\mathcal{X}\right\|\leq C_3|\hat{\nu}(Z)-\nu(Z)|(2\|\xi(X)\|+K-1)$$

for a positive constant $C_3<\infty$, and therefore $\left\|(\hat{\mathcal{X}}-\mathcal{X})Y\right\|\leq C_3|\hat{\nu}(Z)-\nu(Z)|(2\|\xi(X)\|+K-1)|Y|$. By (7) and Lemma 2, the statement of this lemma holds. □

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

C. ADDITIONAL EMPIRICAL RESULTS

In this appendix, we present additional empirical results that are not presented in Section 6 in the main text. The empirical welfare maximization rules depicted by Figures 1 and 2 in the main text are based on the choice of $K = 10$ and $L = 10$ in the estimation procedure. In this appendix, we demonstrate that the results are insensitive to the choice of these parameters, K and L.

Figure C1 illustrates the empirical welfare maximization rules with no treatment cost based on $K \in \{5, 10\}$ and $L \in \{5, 10\}$. Figure C2 illustrates the empirical welfare maximization rules with the treatment cost of \$774 per assignee based on $K \in \{5, 10\}$ and $L \in \{5, 10\}$. Observe that the shaded areas are fairly insensitive to the choice of K and L in each of these two figures. The estimated welfare gains reported for various choices of K and L in Table 1

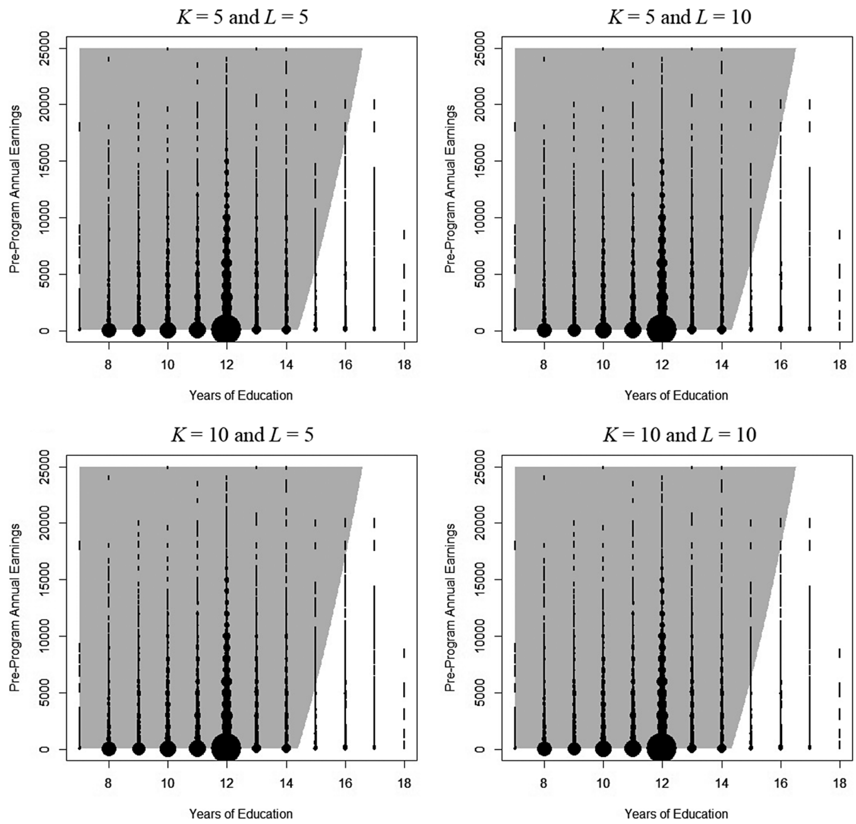

FIGURE C1. Empirical welfare maximization rule with no treatment cost. The area of black circles represents the density of the data at the location of the shape. We use $K \in \{5, 10\}$ for the degree of nonlinearity in estimating the marginal treatment effects and $L \in \{5, 10\}$ for the degree of the years of education in estimating the propensity score.

https://doi.org/10.1017/S0266466624000227 Published online by Cambridge University Press

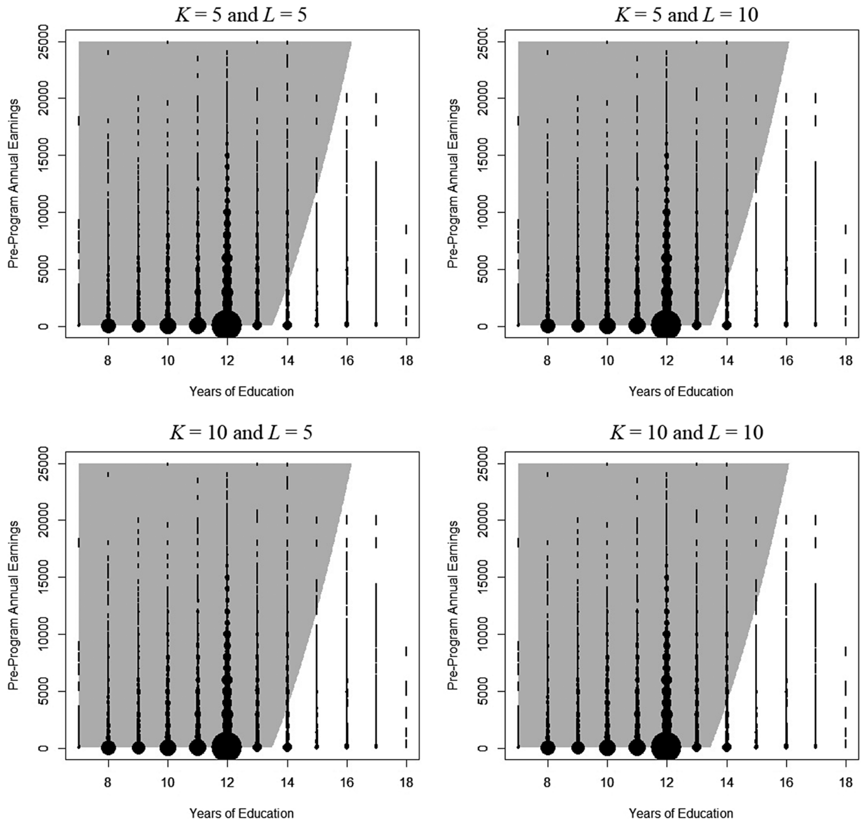

FIGURE C2. Empirical welfare maximization rule with the treatment cost of \$774 per assignee. The area of black circles represents the density of the data at the location of the shape. We use $K \in \{5, 10\}$ for the degree of nonlinearity in estimating the marginal treatment effects and $L \in \{5, 10\}$ for the degree of the years of education in estimating the propensity score.

in the main text are computed based on the empirical welfare maximization rules presented in Figures C1 and C2.

REFERENCES

Abadie, A., Angrist, J., & Imbens, G. (2002). Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings. *Econometrica*, 70, 91–117.

Armstrong, T., & S. Shen (2015). Inference on optimal treatment assignments [Working paper].

Athey, S., & Wager, S. (2021). Policy learning with observational data. *Econometrica*, 89, 133–161.

Bhattacharya, D. (2009). Inferring optimal peer assignment from experimental data. *Journal of the American Statistical Association*, 104, 486–500.

Bhattacharya, D. (2013). Evaluating treatment protocols using data combination. *Journal of Econometrics*, 173, 160–174.

Bhattacharya, D., & Dupas, P. (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 167, 168–196.

Björklund, A., & Moffitt, R. (1987). The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics*, 69, 42–49.

Bloom, H. S., Orr, L. L., Bell, S. H., Cave, G., Doolittle, F., Lin, W., & Bos, J. M. (1997). The benefits and costs of JTPA Title II-A programs: Key findings from the National Job Training Partnership Act study. *Journal of Human Resources*, 32, 549–576.

Brinch, C. N., Mogstad, M., & Wiswall, M. (2017). Beyond LATE with a discrete instrument. *Journal of Political Economy*, 125, 985–1039.

Byambadalai, U. (2021). Identification and inference for welfare gains without unconfoundedness [Working paper].

Carneiro, P., Heckman, J. J., & Vytlacil, E. (2010). Evaluating marginal policy changes and the average effect of treatment for individuals at the margin. *Econometrica*, 78, 377–394.

Carneiro, P., & Lee, S. (2009). Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality. *Journal of Econometrics*, 149, 191–208.

Carneiro, P., Lokshin, M., & Umapathi, N. (2017). Average and marginal returns to upper secondary schooling in Indonesia. *Journal of Applied Econometrics*, 32, 16–36.

Chamberlain, G. (2011). Bayesian aspects of treatment choice. In J. Geweke, G. Koop, & V. H. Dijk (Eds.), *The Oxford handbook of Bayesian econometrics* (pp. 11–39). Oxford University Press.

Cornelissen, T., Dustmann, C., Raute, A., & Schönberg, U. (2016). From LATE to MTE: Alternative methods for the evaluation of policy interventions. *Labour Economics*, 41, 47–60.

Cui, Y., & Tchetgen Tchetgen, E. (2020). A semiparametric instrumental variable approach to optimal treatment regimes under endogeneity. *Journal of the American Statistical Association*, 116, 162–173.

Dehejia, R. H. (2005). Program evaluation as a decision problem. *Journal of Econometrics*, 125, 141–173.

French, E., & Song, J. (2014). The effect of disability insurance receipt on labor supply. *American Economic Journal: Economic Policy*, 6, 291–337.

Han, S. (2023). Optimal dynamic treatment regimes and partial welfare ordering. *Journal of the American Statistical Association*, (forthcoming).

Han, S. (2020). Comment: Individualized treatment rules under endogeneity. *Journal of the American Statistical Association*, 116, 192–195.

Heckman, J. J., Ichimura, H., & Todd, P. E. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Review of Economic Studies*, 64, 605–654.

Heckman, J. J., & Vytlacil, E. (1999). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences*, 96, 4730–4734.

Heckman, J. J., & Vytlacil, E. (2001). Policy-relevant treatment effects. *American Economic Review*, 91, 107–111.

Heckman, J. J., & Vytlacil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73, 669–738.

Heckman, J. J., & Vytlacil, E. (2007). Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. In J. J. Heckman, & E. E. Leamer, *Handbook of econometrics* (Vol. 6, Chapter 71, pp. 4875–5143). Elsevier.

Hirano, K., & Porter, J. R. (2009). Asymptotics for statistical treatment rules. *Econometrica*, 77, 1683–1701.

Hirano, K., & Porter, J. R. (2020). Asymptotic analysis of statistical decision rules in econometrics. In *Handbook of econometrics* (Vol. 7, pp. 283–354). Elsevier.

Kallus, N., & A. Zhou (2018). Confounding-robust policy improvement. *Advances in Neural Information Processing Systems*, 31.

Kitagawa, T., & Tetenov, A. (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica*, 86, 591–616.

Kitagawa, T., & Tetenov, A. (2021). Equality-minded treatment choice. *Journal of Business & Economic Statistics*, 39, 561–574.

Kock, A. B., & M. Thyrsgaard (2017). Optimal sequential treatment allocation [Working paper].

Lee, S., & Salanié, B. (2018). Identifying effects of multivalued treatments. *Econometrica*, 86, 1939–1963.

Liu, Y. (2022). Policy learning under endogeneity using instrumental variables. Preprint. arXiv:2206.09883.

Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72, 1221–1246.

Manski, C. F. (2009). *Identification for prediction and decision*. Harvard University Press.

Mbakop, E., & Tabord-Meehan, M. (2021). Model selection for treatment choice: Penalized welfare maximization. *Econometrica*, 89, 825–848.

Moffitt, R. (2008). Estimating marginal treatment effects in heterogeneous populations. *Annales d'Economie et de Statistique*, 91 /92, 239–261.

Mogstad, M., Santos, A., & Torgovitsky, A. (2018). Using instrumental variables for inference about policy relevant treatment parameters. *Econometrica*, 86, 1589–1619.

Qiu, H., Carone, M., Sadikova, E., Petukhova, M., Kessler, R. C., & Luedtke, A. (2020). Optimal individualized decision rules using instrumental variable methods. *Journal of the American Statistical Association*, 116, 174–191.

Rai, Y. (2018). Statistical inference for treatment assignment policies [Working paper].

Sakaguchi, S. (2021). Estimating optimal dynamic treatment assignment rules under intertemporal budget constraints [Working paper].

Sasaki, Y., & Ura, T. (2023). Estimation and inference for policy relevant treatment effects. *Journal of Econometrics*, 234, 394–450.

Schlag, K. H. (2007). Eleven% designing randomized experiments under minimax regret [Working paper].

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics*, 151, 70–81.

Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics*, 166, 138–156.

Sun, L. (2020). Empirical welfare maximization with constraints [Working paper].

Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics*, 166, 157–165.

Viviano, D. (2019). Policy targeting under network interference [Working paper].

Viviano, D. (2022). Policy design in experiments with unknown interference [Working paper].