Policy learning with U-statistics: Inequality, Inequality of Opportunity and Intergenerational Mobility*
PRELIMINARY DRAFT DO NOT CITE

Joël Terschuur

Universidad Carlos III de Madrid

Abstract

I study a general setting for policy learning with Social Welfare Functions (SWFs) defined by semiparametric U-statistics in experimental and observational settings. I use orthogonal scores to bound the regret at parametric rate whenever the propensity score and the nuisance parameters are unknown. This work expands previous results to welfare functions defined by general orthogonal scores. Three main applications of the general theory guide the paper: (i) Inequality aware SWFs, (ii) Inequality of Opportunity aware SWFs and (iii) Intergenerational Mobility SWFs. I use the Panel Study of Income Dynamics (PSID) to asses the effect of attending preschool on adult earnings and estimate optimal policy rules based on parental years of education and parental income.

1 Welfare economics for inequality, IOp and rank correlations

The policy learning literature is at the intersection of welfare economics and econometrics. Before we delve into the econometric problem of computing optimal rules and evaluating their statistical performance I present in this section the main welfare objects we are going to be interested in. The most basic welfare function is that of the average outcome. Suppose we have some continuous random outcome $Y_i \in \mathbb{R}^+$. A utilitarian planner cares about

$$W = \mathbb{E}[Y_i].$$

*Research funded by Ministerio de Ciencia e Innovación, grant ECO2017-86675-P, MCI/AEI/FEDER/UE, grant PGC 2018-096732-B-100, and Comunidad de Madrid, grants EPUC3M11 (VPRICIT) and H2019/HUM-589. I also thank the supervision of Juan Carlos Escanciano and the meetings with Toru Kitagawa during my visit to Brown University in the Spring semester of 2023.

The above welfare does not care about other distributional aspects asides from the average outcome. A first approach to include distributional concerns in our analysis is to follow Dalton (1920) and Atkinson et al. (1970) and consider increasing and concave transformations $u(\cdot)$ of the outcome[1].

$$W = \mathbb{E}[u(Y_i)].$$

This welfare function will already rank two outcome distributions in the same way for all increasing and concave $u(\cdot)$ if the Lorenz curve of one of the distributions is everywhere above the Lorenz curve of the other distribution and has equal or higher mean; equivalently, if one distribution second order stochastically dominates the other. However, if we want to obtain a complete ordering we need to specify $u(\cdot)$ further. One popular choice is

$$u(y) = \begin{cases} \frac{y^{1-\theta}}{1-\theta} & \text{if } \theta \in (0,1) \\ \log(y) & \text{if } \theta = 1. \end{cases}$$

θ captures the concavity of $u(\cdot)$ and can therefore be interpreted as an inequality aversion parameter. This paper also focuses on welfare which is aware of Inequality of Opportunity (IOp). IOp is the part of total inequality which can be explained by circumstances, i.e. by variables that are outside the control of the individual such as parental education or parental income. Let $X_i \in \mathbb{R}^k$ be such a random vector of circumstances. Let also $\gamma(X_i) = \mathbb{E}[Y_i|X_i]$, i.e. the best predictor of the outcome Y_i given the circumstances X_i. By looking at the distribution of $\gamma(X_i)$ instead of that of the outcome Y_i we get IOp averse welfare functions. For instance,

$$W = \mathbb{E}[u(\gamma(X_i))].$$

If there is no IOp, circumstances are unable to predict the outcome and we have that the best predictor is the unconditional mean: $\gamma(X_i) = \mathbb{E}[Y_i]$. In this case we have that $W = u(\mathbb{E}[Y_i])$ so we only care about the average income (with a different scale due to $u(\cdot)$). If we have maximum IOp, the outcome is a deterministic function of the circumstances and $\gamma(X_i) = Y_i$. Then, $W = \mathbb{E}[u(Y_i)]$. Since all inequality is IOp, we are back at the inequality averse welfare function.

Another option to take into account distributional concerns is to weight differently different parts of the distribution. Let F_Y be the distribution of the outcome and F_Y^{-1} be the quantiles. Then, for some weights $w(\cdot)$ a planner might have the following welfare in mind

$$W = \int_0^1 F_Y^{-1}(\tau) w(\tau) d\tau.$$

This welfare has been used in Mehran (1976), Donaldson and Weymark (1980), Weymark (1981), Donaldson and Weymark (1983) or Aaberge et al. (2021). If we let $w_k(\tau) = (k-1)(1-\tau)^{k-2}$

[1] With abuse of notation we call W to all welfare functions as they appear.

we get what is known as the extended Gini family of social welfare functions. In this paper we focus on $k = 3$ which is known as the standard Gini social welfare function and can be shown to be

$$\begin{aligned} W &= \mathbb{E}[Y_i](1 - G(Y_i)) \\ &= (1/2)\mathbb{E}[Y_i + Y_j - |Y_i - Y_j|], \end{aligned}$$

where $G(Y_i)$ is the Gini coefficient of the distribution of Y_i, the second equality follows from the fact that we can write the Gini of Y_i as $G(Y_i) = \mathbb{E}[|Y_i - Y_j]/\mathbb{E}[Y_i + Y_j]$ where Y_j is a copy of Y_i (i.e. the Gini can be interpreted as a normalized absolute distance between the outcomes of two individuals taken at random). The welfare above is utilitarian as long as there is no inequality ($G(Y_i) = 0$) and penalizes positive values of the Gini coeffcient. Again, if we do not care about inequality but only about IOp we can look at the distribution of $\gamma(X_i)$ instead of the distribution of Y_i. In that case we have

$$\begin{aligned} W &= \mathbb{E}[Y_i](1 - G(\gamma(X_i))) \\ &= (1/2)\mathbb{E}[\gamma(X_i) + \gamma(X_j) - |\gamma(X_i) - \gamma(X_j)|]. \end{aligned}$$

If there is no IOp, then $G(\gamma(X_i)) = 0$ and we are back in the utilitarian case. If there is full IOp, then $G(\gamma(X_i)) = G(Y_i)$ and we are back to the standard Gini social welfare function of outcome Y_i.

Finally, I also consider the problem of intergenerational mobility. Let $X_{1i} \in \mathbb{R}$ be the first component of X_i. A measure of rank correlation between Y_i and X_{1i} is the Kendall-τ

$$\tau = \mathbb{E}[sgn(Y_i - Y_j)sgn(X_{1i} - X_{1j})].$$

This parameter is popular in the intergenerational mobility literature (see Chetty et al. (2014) or Kitagawa et al. (2018)) where X_{1i} is parental income and Y_i is the child's income. For some target rank correlation $t \in [-1, 1]$ an intergenerational mobility aware welfare function is

$$W = -\left|\mathbb{E}\left[sgn(Y_i - Y_j)sgn(X_{1i} - X_{1j})\right] - t\right|.$$

2 Policy learning with general orthogonal scores

Consider random variables $(Y_i(1), Y_i(0), D_i, X_i) \sim F_0$ where $(Y_i(1), Y_i(0)) \in \mathcal{Y} \times \mathcal{Y}$ are real-valued potential outcomes, i.e. $Y_i(1)$ is the outcome of individual i under treatment and $Y_i(0)$ is the outcome of individual i in the absence of treatment. D_i is a binary treatment and $X_i \in \mathcal{X}$ is a vector of pre-treatment covariates. Let $\gamma^{(j)}(X_i) = \mathbb{E}[Y_i(j)|X_i] \in \Gamma$ for $j = 0, 1$ be potential predictions, i.e. the predictions of the potential outcomes given X_i. We observe an i.i.d. sample $(Z_1, ..., Z_n)$ with $Z_i = (Y_i, D_i, X_i) \in \mathcal{Z}$ and $Y_i = Y_i(1)D_i + Y_i(0)(1 - D_i) \in \mathcal{Y}$. Let $\pi : \mathcal{X} \mapsto \{0, 1\}$ be a treatment rule which indicates who receives treatment and Π be a collection of such treatment rules. We are interested in choosing a policy $\pi \in \Pi$ so as to maximize the following welfare function

$$W(\pi) = \mathbb{E}[g(Y_i(1), X_i, \gamma^{(1)})\pi(X_i) + g(Y_i(0), X_i, \gamma^{(0)})(1 - \pi(X_i))]. \tag{2.1}$$

Example 1 (IOp Atkinson) *If we were interested in an inequality aware SWF we could use Atkinson SWFs,* $W(\pi) = \mathbb{E}[u(Y_i(1))\pi(X_i) + u(Y_i(0))(1 - \pi(X_i))]$ *with* $u(\cdot)$ *a concave function. In this case the optimal policy can be estimated using the methods in Kitagawa and Tetenov (2018) and Athey and Wager (2021). If we want an IOp aware SWF we can simply look at the distribution of* $\gamma(X_i)$ *instead of at the distribution of* Y_i*:*

$$W(\pi) = \mathbb{E}[u(\gamma^{(1)}(X_i))\pi(X_i) + u(\gamma^{(0)}(X_i))(1 - \pi(X_i))].$$

■

Importantly, (2.1) is not observable since for a given individual we do not observe both potential outcomes. In order to identify (2.1) we first need our sample to come from an experimental or observational experiment where the policy has already been implemented and where the following holds.

Assumption 1 *i)* $(Y_i(1), Y_i(0)) \perp D_i | X_i$,

ii) There exists $\kappa \in (0, 1/2]$ *such that* $e(x) \in [\kappa, 1 - \kappa]$.

The next proposition states the first identification result. The identification results depends on whether g depends on the potential outcomes. If it does one can use a direct method based on regression functions or use inverse propensity score weighting.

Proposition 2.1 *Under Assumption 1,* $W(\pi)$ *is identified in the following ways*

$$W(\pi) = \mathbb{E}[m_1(Z_i, \gamma, \nu)\pi(X_i) + m_0(Z_i, \gamma, \nu)(1 - \pi(X_i))],$$

where

$$m_1(Z_i, \gamma, \nu) = \begin{cases} g(X_i, \gamma_1) & \text{if } \nu = g \\ \varphi(1, X_i, \gamma_1) & \text{if } \nu = \varphi \\ \frac{g(Y_i, X_i, \gamma_1)D_i}{e(X_i)} & \text{if } \nu = e \end{cases}, \quad m_0(Z_i, \gamma, \nu) = \begin{cases} g(X_i, \gamma_0) & \text{if } \nu = g \\ \varphi(0, X_i, \gamma_0) & \text{if } \nu = \varphi \\ \frac{g(Y_i, X_i, \gamma_0)(1 - D_i)}{1 - e(X_i)} & \text{if } \nu = e \end{cases},$$

where $\gamma(D_i, X_i) = \mathbb{E}[Y_i | D_i, X_i]$, $\gamma_j(X_i) = \gamma(j, X_i)$ *for* $j = 0, 1$ *and* $\varphi(D_i, X_i) = \mathbb{E}[g(Y_i, X_i, \gamma) | D_i, X_i]$.

The proposition above distinguishes three cases based on $\nu \in \{g, \varphi, e\}$. When $\nu = g$ we are in the case in which g only depends on potential nuisance parameters but not on actual potential outcomes, i.e. $g(u, X_i, \gamma^{(j)}) = g(t, X_i, \gamma^{(j)}) \equiv g(X_i, \gamma^{(j)})$ for all $u, t \in \mathcal{Y}$ and $u \neq t$. $\nu = \varphi$ is the case in which the researcher chooses to follow a direct method approach instead of an inverse propensity score weighting approach and $\nu = e$ is the inverse propensity score weighting approach.

Hence, depending on which case we are, we will have either γ or (γ, φ) or (γ, e) as nuisance parameters. To enjoy double robustness properties and local robustness to high dimensional and ML first steps I provide orthogonal scores in the next result. First I need the following assumption to take care of the nuisance parameter γ.

Assumption 2 *There exist (α_1, α_0) such that for any $\tilde{\gamma} \in L_2$ and $j = 0, 1$ and $\tau >= 0$*

$$\frac{d}{d\tau}\mathbb{E}[m_j(Z_i, \bar{\gamma}_\tau, \nu)]\bigg|_{\tau=0} = \frac{d}{d\tau}\mathbb{E}[\alpha_j(D_i, X_i, \nu)\bar{\gamma}_\tau(D_i, X_i)]\bigg|_{\tau=0},$$

where $\bar{\gamma}_\tau = \gamma + \tau\tilde{\gamma}$.

Since γ enters m_j only through g a sufficient conditon is to assume a similar result for the function g instead of m_j and then it will be straightforward to find the α for each $\nu \in \{g, \varphi, e\}$.

Proposition 2.2 *The orthogonal score is given by*

$$\Gamma_i(\pi) = \Gamma_{1i}\pi(X_i) + \Gamma_{0i}(1 - \pi(X_i)),$$

where if $\nu \neq g$ we have

$$\Gamma_{1i} = \varphi(1, X_i, \gamma) + \frac{D_i}{e(X_i)}(Y_i - \varphi(1, X_i, \gamma)) + \alpha_1(D_i, X_i, \nu)(Y_i - \gamma(D_i, X_i)),$$
$$\Gamma_{0i} = \varphi(0, X_i, \gamma) + \frac{1 - D_i}{1 - e(X_i)}(Y_i - \varphi(0, X_i, \gamma)) + \alpha_0(D_i, X_i, \nu)(Y_i - \gamma(D_i, X_i)).$$

and if $\nu = g$ we have $\Gamma_{1i} = \varphi(1, X_i, \gamma) + \alpha_1(D_i, X_i, g)(Y_i - \gamma(D_i, X_i))$ and $\Gamma_{0i} = \varphi(0, X_i, \gamma) + \alpha_0(D_i, X_i, g)(Y_i - \gamma(D_i, X_i))$.

To estimate the welfare for a given $\pi \in \Pi$ we employ cross-fitting as in Escanciano and Terschuur (2022). Let the data be split in L groups $I_1, ..., I_l$, then

$$\hat{W}_n(\pi) = \frac{1}{n}\sum_{l=1}^{L}\sum_{i \in I_l}\hat{\Gamma}_{1i,l}\pi(X_i) + \hat{\Gamma}_{0i,l}(1 - \pi(X_i)),$$

where

$$\hat{\Gamma}_{1i,l} = \hat{\varphi}_l(1, X_i, \hat{\gamma}_l) + \frac{D_i}{\hat{e}_l(X_i)}(Y_i - \hat{\varphi}_l(1, X_i, \hat{\gamma}_l)) + \hat{\alpha}_{1,l}(D_i, X_i, \nu)(Y_i - \hat{\gamma}_l(D_i, X_i)),$$
$$\hat{\Gamma}_{0i,l} = \hat{\varphi}_l(0, X_i, \hat{\gamma}_l) + \frac{1 - D_i}{1 - \hat{e}_l(X_i)}(Y_i - \hat{\varphi}_l(0, X_i, \hat{\gamma}_l)) + \hat{\alpha}_{0,l}(D_i, X_i, \nu)(Y_i - \hat{\gamma}_l(D_i, X_i)),$$

and $(\hat{\varphi}_l, \hat{e}_l, \hat{\gamma}_l, \hat{\alpha}_{j,l})$, $j = 0, 1$, are estimators of the nuisance functions which do not use observations in I_l. Again, whenever g does not depend on the potential outcomes the middle term in both expressions is zero. This is the case in the example of Atkinson welfare IOp.

Example 1 (IOp Atkinson (cont.)) *For $\theta \in (0,1]$ let*

$$U(\gamma(x)) = \begin{cases} \frac{\gamma(x)^{1-\theta}}{1-\theta} & \text{if } \theta \in (0,1) \\ \log(\gamma(x)) & \text{if } \theta = 1. \end{cases}$$

θ controls the concavity of U and therefore is a parameter capturing inequality aversion which can be picked by the policy maker. In this case $g = U$ which only depends on the nuisance parameters, so $\nu = g$. The orthogonal score for $\theta \in (0,1]$ is

$$\begin{aligned}\Gamma_i(\pi) = U(\gamma(1, X_i)) &+ \frac{\gamma(D_i, X_i)^{-\theta} D_i}{e(X_i)}(Y_i - \gamma(D_i, X_i))\pi(X_i) \\ + U(\gamma(0, X_i)) &+ \frac{\gamma(D_i, X_i)^{-\theta}(1 - D_i)}{1 - e(X_i)}(Y_i - \gamma(D_i, X_i))(1 - \pi(X_i)),\end{aligned}$$

i.e. $\alpha_1(D_i, X_i, g) = e(X_i)^{-1}\gamma(D_i, X_i)^{-\theta} D_i$ *and* $\alpha_0(D_i, X_i, g) = (1 - e(X_i))^{-1}\gamma(D_i, X_i)^{-\theta}(1 - D_i)$.
■

The estimator of the optimal treatment rule among a class of rules Π is

$$\hat{\pi} = \arg\max_{\pi \in \Pi} \hat{W}_n(\pi).$$

Before analysing the statistical performance of such rule let us first extend the results in this section to welfare functions based on U-statistics. This will allow us to consider inequality and IOp aware SWFs based on the Gini coefficient and also a rank correlation aware SWF which is relevant for intergenerational mobility.

3 Policy learning with U-statistics

Let now $\pi_{ab}(X_i, X_j) = \mathbb{1}(\pi(X_i) = a) \times \mathbb{1}(\pi(X_j) = b)$ with $a, b \in \{0, 1\}$. Now we consider the following SWFs

$$W(\pi) = \mathbb{E}\Big[\sum_{(a,b)\in\{0,1\}^2} g(Y_i(a), X_i, Y_j(b), X_j, \gamma^{(a)}, \gamma^{(b)})\pi_{ab}(X_i, X_j)\Big]. \tag{3.1}$$

Example 2 (Inequality) *We can accommodate the standard Gini welfare function with*

$$g(Y_i(a), Y_j(b)) = (1/2)(Y_i(a) + Y_j(b) - |Y_i(a) - Y_j(b)|).$$

Kitagawa and Tetenov (2021) analyse treatment allocation for this welfare function starting from another representation of the Gini coefficient which depends on the c.d.f. of Y_i. Our representation is more useful whenever the population among which the treatment is allocated and the population over which the welfare is computed differs. ■

Example 3 (Inequality of Opportunity IOp) *We can apply the standard Gini welfare function to the distribution of the predictions to get $\mathbb{E}[\gamma(X_i)](1-G(\gamma(X_i)))$. This fits our setting by letting*

$$g(X_i, X_j, \gamma^{(a)}, \gamma^{(b)}) = (1/2)(\gamma^{(a)}(X_i) + \gamma^{(b)}(X_j) - |\gamma^{(a)}(X_i) - \gamma^{(b)}(X_j)|).$$

■

Example 4 (Rank correlation) *If we want to allocate a treatment targeting a specific Kendall-τ, say $t \in \mathbb{R}$, we have to extend our setting to transformations of the RHS of 3.1. We can define*

$$g(Y_i(a), X_{1i}, Y_j(b), X_{1j}) = sgn(Y_i(a) - Y_j(b))sgn(X_{1i} - X_{1j}),$$

where $sgn(a) = \mathbb{1}(a > 0) - \mathbb{1}(a < 0)$ and let

$$W(\pi) = -\left|\mathbb{E}\left[\sum_{(a,b)\in\{0,1\}^2} g(Y_i(a), X_{1i}, Y_j(b), X_{1j})\pi_{ab}(X_i, X_j)\right] - t\right|.$$

■

Proposition 3.1 *Under Assumption 1, $W(\pi)$ in (3.1) is identified in the following ways*

$$W(\pi) = \mathbb{E}\left[\sum_{(a,b)\in\{0,1\}^2} m_{ab}(Z_i, Z_j, \gamma, \nu)\pi_{ab}(X_i, X_j)\right],$$

where

$$m_{ab}(Z_i, Z_j, \gamma, \nu) = \begin{cases} g(X_i, X_j, \gamma_a, \gamma_b) & \text{if } \nu = g \\ \varphi(a, X_i, b, X_j, \gamma_a, \gamma_b) & \text{if } \nu = \varphi \\ g(Y_i, X_i, Y_j, X_j, \gamma_a, \gamma_b)D_iD_j/e_{ab}(X_i, X_j) & \text{if } \nu = e \end{cases},$$

where $\varphi(a, X_i, b, X_j, \gamma_a, \gamma_b) = \mathbb{E}[g(Y_i, X_i, Y_j, X_j, \gamma_a, \gamma_b)|D_i, X_i, D_j, X_j]$ and $e_{ab}(X_i, X_j) = \mathbb{P}(D_i = a|X_i)\mathbb{P}(D_j = b|X_j)$.

Now we apply Proposition 3.1 to identify the welfare in each of our three main examples.

Example 2 (Inequality (cont.)) *In this example the welfare is identified by*

$$W(\pi) = \mathbb{E}\left[\frac{1}{2}(Y_i + Y_j - |Y_i - Y_j|) \sum_{(a,b)\in\{0,1\}^2} \frac{D_iD_j}{e_{ab}(X_i, X_j)}\pi_{ab}(X_i, X_j)\right].$$

■

Example 3 (IOp (cont.)) *In this example the welfare is identified by*

$$W(T) = \frac{1}{2}\mathbb{E}\Big[\sum_{(a,b)\in\{0,1\}^2}\Big(\gamma_a(X_i)+\gamma_b(X_j)-|\gamma_a(X_i)-\gamma_b(X_j)|\Big)\pi_{ab}(X_i,X_j)\Big].$$

■

Example 4 (Intergenerational mobility (cont.)) *In this example the welfare is identified by*

$$W(T) = -\Big|\mathbb{E}\Big[sgn(X_{1i}-X_{1j})sgn(Y_i-Y_j)\sum_{(a,b)\in\{0,1\}^2}\frac{D_iD_j}{e_{ab}(X_i,X_j)}\pi_{ab}(X_i,X_j)\Big]-t\Big|.$$

■

In Examples 2 and 4 we have used $\nu = e$ since we will see that it makes the estimation simpler. Example 3 does not depend on the potential outcomes ($\nu = g$). In order to compute the orthogonal scores we need to assume a similar linearization property as that in Assumption 2.

Assumption 3 *There exist α_{ab}, $P < \infty$, and constants (c_{1p}, c_{2p}) for $p = 1, ..., P$, such that for all $(a,b) \in \{0,1\}^2$ the following linearization holds*

$$\frac{d}{d\tau}\mathbb{E}[m_{ab}(Z_i,Z_j,\bar{\gamma}_\tau,\nu)] = \mathbb{E}\Big[\sum_{p=1}^{P}\alpha^{\gamma}_{ab,p}(D_i,X_i,D_j,X_j,\nu)(c_{1p}\bar{\gamma}_\tau(D_i,X_i)+c_{2p}\bar{\gamma}_\tau(D_j,X_j))\Big],$$

where $\bar{\gamma}_\tau$ is defined as in Assumption 2.

Again, γ enters m_{ab} only through g so a sufficient condition that would allow to compute the α_{ab} for each $\nu \in \{g, \varphi, e\}$ is to assume a linearization like the above for g instead of for m_{ab}. Now we are ready to present the result of the orthogonal scores for welfare functions defined with U-statistics.

Proposition 3.2 *The orthogonal scores are given by*

$$\Gamma_{ij}(\pi) = \sum_{(a,b)\in\{0,1\}^2}\Gamma^{ab}_{ij}\pi_{ab}(X_i,X_j),$$

where

$$\Gamma^{ab}_{ij} = \begin{cases} m_{ab}(Z_i,Z_j,\gamma,g) + \phi^{\gamma}_{ab}(D_i,X_i,D_j,X_j,\gamma,\alpha^{\gamma}) & \text{if } \nu = g \\ m_{ab}(Z_i,Z_j,\gamma,\varphi) + \phi^{m}_{ab}(D_i,X_i,D_j,X_j,\varphi,\alpha^{e}) + \phi^{\gamma}_{ab}(D_i,X_i,D_j,X_j,\gamma,\alpha^{\gamma}) & \text{if } \nu = \varphi \\ m_{ab}(Z_i,Z_j,\gamma,e) + \phi^{e}_{ab}(D_i,X_i,D_j,X_j,e,\alpha^{e}) + \phi^{\gamma}_{ab}(D_i,X_i,D_j,X_j,\gamma,\alpha^{\gamma}) & \text{if } \nu = e, \end{cases}$$

where

$$\phi^{\gamma}_{ab}(D_i, X_i, D_j, X_j, e, \alpha^{\gamma}) = \sum_{p=1}^{P} \alpha^{\gamma}_{ab,p}(D_i, X_i, D_j, X_j, e)(c_{1p}Y_i + c_{2p}Y_j - c_{1p}\gamma(D_i, X_i) - c_{2p}\gamma(D_j, X_j)),$$

$$\phi^{m}_{ab}(D_i, X_i, D_j, X_j, \varphi, \alpha^{m}) = \alpha^{m}_{ab}(D_i, X_i, D_j, X_j, \varphi)(g(Y_i, X_i, Y_j, X_j, \gamma_a, \gamma_b) - \varphi(D_i, X_i, D_j, X_j, \gamma_a, \gamma_b)),$$

$$\phi^{e}_{ab}(D_i, X_i, D_j, X_j, e, \alpha^{e}) = \alpha^{e}_{ab,1}(X_i)(\mathbb{1}(D_i = a) - e_a(X_i)) + \alpha^{e}_{ab,2}(X_j)(\mathbb{1}(D_j = b) - e_b(X_j)),$$

and

$$\alpha^{m}_{ab}(D_i, X_i, D_j, X_j, \varphi) = \frac{D_i D_j - e_{ab}(X_i, X_j)}{e_{ab}(X_i, X_j)(1 - e_{ab}(X_i, X_j))},$$

$$\alpha^{e}_{ab,1}(X_i) = -\mathbb{E}\left[\frac{g(Y_i, X_i, Y_j, X_j, \gamma_a, \gamma_b) D_i D_j}{e_a(X_i)^2 e_b(X_j)}\middle| X_i\right],$$

$$\alpha^{e}_{ab,2}(X_j) = -\mathbb{E}\left[\frac{g(Y_i, X_i, Y_j, X_j, \gamma_a, \gamma_b) D_i D_j}{e_a(X_i) e_b(X_j)^2}\middle| X_j\right].$$

Now we can see how Proposition 3.2 applies to our examples. We

Example 2 (Inequality (cont.)) *In this example we have that*

$$\Gamma_{ij} = \frac{1}{2}(Y_i + Y_j - |Y_i - Y_j|) \sum_{(a,b)\in\{0,1\}^2} \frac{D_i D_j}{e_{ab}(X_i, X_j)} \pi_{ab}(X_i, X_j)$$
$$+ \sum_{(a,b)\in\{0,1\}^2} \alpha^{e}_{ab,1}(X_i)(\mathbb{1}(D_i = a) - e_a(X_i)) + \alpha^{e}_{ab,2}(X_j)(\mathbb{1}(D_j = b) - e_b(X_j)).$$

■

Example 3 (IOp (cont.))

$$\Gamma_{ij} = \frac{1}{2} \sum_{(a,b)\in\{0,1\}^2} \Big(\gamma_a(X_i) + \gamma_b(X_j) - |\gamma_a(X_i) - \gamma_b(X_j)|\Big) \pi_{ab}(X_i, X_j)$$
$$+ \sum_{(a,b)\in\{0,1\}^2} \Big(\varphi_{ai}(1 - \delta^{ab}_{ij})(Y_i - \gamma(D_i, X_i)) + \varphi_{bj}(1 - \delta^{ab}_{ij})(Y_j - \gamma(D_j, X_j))\Big).$$

Assuming either (i) $\mathbb{P}(\gamma_T(X_i) - \gamma_T(X_j) = 0) = 0$ or that (ii) $x_i \neq x_j \implies \gamma_T(X_i) - \gamma_T(X_j) \neq 0$, it follows that Assumption 3 holds with $P = 2$, $(c_{11}, c_{12}, c_{21}, c_{22}) = (1, 0, 0, 1)$, $\alpha_{ab,1} = \varphi_{ai}(1 - \delta_{ij})$, $\alpha_{ab,2} = \varphi_{bi}(1 - \delta_{ij})$ and $\varphi_{ai} = \mathbb{1}(D_i = a)/e_a(X_i)$ and $\delta^{ab}_{ij} = sgn(\varphi_{ai}\gamma(D_i, X_i) - \varphi_{bj}\gamma(D_j, X_j))$ (see proof of Proposition 1 in Escanciano and Terschuur (2022)). These assumptions deal with the point of non-differentiability of the absolute value. (i) is satisfied if $\gamma_T(X_i)$ is absolutely continuous, for example if γ_T is strictly monotone on a circumstance which is absolutely continuous given all the other circumstances. Assumption (ii) says that two observations with different circumstances must have different fitted values. ■

Example 4 (Intergenerational mobility (cont.)) *The orthogonal score is given by*

$$\Gamma_{ij} = \frac{1}{2}(sgn(X_{1i} - X_{1j})sgn(Y_i - Y_j)) \sum_{(a,b)\in\{0,1\}^2} \frac{D_i D_j}{e_{ab}(X_i, X_j)} \pi_{ab}(X_i, X_j)$$
$$+ \sum_{(a,b)\in\{0,1\}^2} \alpha^e_{ab,1}(X_i)(\mathbb{1}(D_i = a) - e_a(X_i)) + \alpha^e_{ab,2}(X_j)(\mathbb{1}(D_j = b) - e_b(X_j)).$$

■

To estimate the welfare for a given $\pi \in \Pi$ I use an adaptation to U-statistics of the cross-fitting used before (see Escanciano and Terschuur (2022)). I split the pairs $\{(i, j) \in \{1, ..., n\}^2 : i < j\}$in L groups $I_1, ..., I_l$, then

$$\hat{W}_n(\pi) = \frac{1}{n} \sum_{l=1}^{L} \sum_{(i,j)\in I_l} \hat{\Gamma}_{ij,l}, \tag{3.2}$$

were $\hat{\Gamma}_{ij,l}$ is the same as Γ_{ij} but with all nuisance parameters replaced by estimators which do not use observations in the pairs in I_l. As before, the estimator of the optimal treatment rule among a class of rules Π is

$$\hat{\pi} = \arg\max_{\pi\in\Pi} \hat{W}_n(\pi).$$

4 Asymptotic statistical guarantees

In general we have that for a given treatment rule π, orthogonal scores are given by

$$\Gamma_{ij}(\pi) = \sum_{(a,b)\in\{0,1\}^2} \psi_{ab}(Z_i, Z_j, \gamma, \nu, \alpha)\pi_{ab}(X_i, X_j).$$

ψ_{ab} is the sum of a function whose expectation identifies a term in the welfare plus other functions which are correction terms needed to achieve orthogonality with respect to the nuisance parameter γ, nuisance parameters which appear in the identification result ν and additional nuisance parameters which appear in the process of computing the orthogonal scores α

$$\psi_{ab}(Z_i, Z_j, \gamma, \nu, \alpha) = m_{ab}(Z_i, Z_j, \gamma, \nu) + \phi^\gamma_{ab}(Z_i, Z_j, \gamma, \alpha^\gamma) + \phi^\nu_{ab}(Z_i, Z_j, \nu, \alpha^\nu).$$

This framework accomodates also the welfare functions which are not defined as U-statistics if $\psi_{ab}(Z_i, Z_j, \gamma, \nu, \alpha)$ does not depend on Z_j and only depends on a so that we could rewrite it as $\psi_a(Z_i, \gamma, \nu, \alpha)$ for $a \in \{0, 1\}$. For this reason we stick to this notation and do not state all conditions and results for welfare functions which are not U-statistics and those which are. In the next subsections we give conditions on the convergence of the nuisance parameters and on the complexity of the policy class Π which will allow us to prove asymptotical statistical guarantees for our estimated treatment rules.

4.1 Conditions on the nuisance parameter estimators

I give high level conditions for the estimators of all nuisance parameters which have to be used to estimate the welfare. These conditions have been showed to hold for a variety of non-parametric estimators such as kernels or sieve estimators. The assumptions below are analogous to those in Escanciano and Terschuur (2022).

Assumption 4 $\mathbb{E}[|\psi(Z_i, Z_j, \gamma, \nu, \alpha)|^2] < \infty$ *and for* $(a, b) \in \{0, 1\}^2$, $\omega \in \{\gamma, \nu\}$ *and for some* $a(n) \to 0$

(i) $n^{\lambda_\gamma}\sqrt{\mathbb{E}(|m_{ab}(z_i, z_j, \hat{\gamma}_l, \nu) - m_{ab}(z_i, z_j, \gamma, \nu)|^2)} \leq a(n)$;

(ii) $n^{\lambda_\nu}\sqrt{\mathbb{E}(|m_{ab}(z_i, z_j, \gamma, \hat{\nu}_l) - m_{ab}(z_i, z_j, \gamma, \nu)|^2)} \leq a(n)$;

(iii) $n^{\lambda_\gamma}\sqrt{\mathbb{E}(|\phi^\gamma_{ab}(z_i, z_j, \hat{\gamma}_l, \alpha^\gamma) - \phi^\gamma_{ab}(z_i, z_j, \gamma, \alpha^\gamma)|^2)} \leq a(n)$;

(iv) $n^{\lambda_\nu}\sqrt{\mathbb{E}(|\phi^\nu_{ab}(z_i, z_j, \hat{\nu}_l, \alpha^\nu) - \phi^\nu_{ab}(z_i, z_j, \nu, \alpha^\nu)|^2)} \leq a(n)$;

(v) $n^{\lambda_\alpha}\sqrt{\mathbb{E}(|\phi^\omega_{ab}(z_i, z_j, \omega, \hat{\alpha}^\omega_l) - \phi^\omega_{ab}(z_i, z_j, \omega, \alpha^\omega)|^2)} \leq a(n)$,

where $0 < \lambda_\gamma, \lambda_\nu, \lambda_\alpha < 1/2$.

These are mild mean-square consistency conditions for $\hat{\gamma}_l$, $\hat{\nu}_l$ and $\hat{\alpha}_l$ separately. Assumption 4 often follows from the L2 convergence rates of the nuisance estimators. Define also the following interaction terms for $\omega \in \{\gamma, \nu\}$

$$\hat{\xi}^\omega_l(w_i, w_j) = \phi(z_i, z_j, \hat{\omega}_l, \hat{\alpha}^\omega_l) - \phi(z_i, z_j, \omega, \hat{\alpha}^\omega_l) - \phi(z_i, z_j, \hat{\omega}_l, \alpha^\omega) + \phi(z_i, z_j, \omega, \alpha^\omega).$$

Assumption 5 *either*

(i) $\sqrt{n}\mathbb{E}(\hat{\xi}^\omega_l(w_i, w_j)) \leq a(n), \mathbb{E}(|\hat{\xi}^\omega_l(w_i, w_j)|^2) \leq a(n)$ *or,*

(ii) $\sqrt{n}\binom{n}{2}^{-1}\sum_{(i,j)\in I_l}|\hat{\xi}^\omega_l(W_i, W_j)| \leq 0$.

These are rate conditions on the remainder terms $\hat{\xi}^\omega_l(w_i, w_j)$. Often, Assumption 5 follows if $\sqrt{n}||\hat{\alpha}^\omega_l - \alpha||||\hat{\omega}_l - \omega|| \leq a(n)$, where $||\cdot||$ denotes the L2 norm.

4.2 Conditions on the complexity of the policy class

The complexity of the policy class must also be restricted. If all sorts of subsets of $\mathcal{X}$ are allowed to decide who should be treated then we get overfitted policy rules. As in Athey and Wager (2021) we measure the policy class complexity with its VC dimension (see for instance Wainwright (2019)) which we allow to grow with the sample size. Hence, for now on we subscript the policy class by n, Π_n.

Assumption 6 *There are constants $0 < \beta < 1/2$ and $n^* \geq 1$ such that for all $n \geq n^*$, $VC(\Pi_n) < n^\beta$.*

Examples of finite VC-dimension classes are linear eligibility scores or generalized eligibility scores (see Kitagawa and Tetenov (2018)). Policy classes which increase with the sample size are for example decision trees (see Athey and Wager (2021)).

4.3 Upper bounds

Let now

$$\begin{aligned} W(\pi) &= \mathbb{E}\Big[\sum_{(a,b)\in\{0,1\}^2} \psi_{ab}(Z_i, Z_j, \gamma, \nu, \alpha)\pi_{ab}(X_i, X_j)\Big], \\ \widetilde{W}_n(\pi) &= \binom{n}{2}^{-1} \sum_{i<j} \Big[\sum_{(a,b)\in\{0,1\}^2} \psi_{ab}(Z_i, Z_j, \gamma, \nu, \alpha)\pi_{ab}(X_i, X_j)\Big], \\ \hat{W}_n(\pi) &= \binom{n}{2}^{-1} \sum_{l=1}^{L} \sum_{(i,j)\in I_l} \Big[\sum_{(a,b)\in\pi} \psi_{ab}(Z_i, Z_j, \hat{\gamma}_l, \hat{\nu}_l, \hat{\alpha}_l)\pi_{ab}(X_i, X_j).\Big], \end{aligned}$$

$W(\pi)$ and $\widetilde{W}_n(\pi)$ are the welfare at policy rule π and the unfeasible estimator of the welfare when all nuisance parameters are known. $\hat{W}_n(\pi)$ is the feasible estimator which we already introduced in (3.2). Let $W^*_{\Pi_n} = \sup_{\pi\in\Pi_n} W(\pi)$ be the best possible welfare. We want to give upper bounds to the regret: $\mathbb{E}[W^*_{\Pi_n} - W(\hat{\pi})]$, i.e. the expected difference between the best possible welfare and the the welfare evaluated at the estimated policy. As usual I start bounding the regret as follows

$$\begin{aligned} \mathbb{E}[W^*_{\Pi_n} - W(\hat{\pi})] &\leq 2\mathbb{E}\Big[\sup_{\pi\in\Pi_n} |\hat{W}_n(\pi) - W(\pi)|\Big] \\ &\leq 2\mathbb{E}\Big[\sup_{\pi\in\Pi_n} |\hat{W}_n(\pi) - \widetilde{W}_n(\pi)|\Big] + 2\mathbb{E}\Big[\sup_{\pi\in\Pi_n} |\widetilde{W}_n(\pi) - W(\pi)|\Big], \end{aligned} \tag{4.1}$$

where in the second inequality we have added and subtracted $\widetilde{W}_n(\pi)$ and used the triangle inequality. The second term above is just a standard centered U-process indexed by $\pi \in \Pi_n$. We start as in Athey and Wager (2021) by showing the rate of convergence of this second term. We work for some fixed $(a, b) \in \{0, 1\}^2$ and we define the following set

$$\Pi_{ab,n} = \{\pi_{ab} : \pi \in \Pi_n\}.$$

The first step is to bound it by the Rademacher complexity which we define as

$$\mathcal{R}_n(\Pi_{ab,n}) = \mathbb{E}_\varepsilon\Bigg(\sup_{\pi\in\Pi_n}\Bigg|\lfloor n/2\rfloor^{-1} \sum_{i=1}^{\lfloor n/2\rfloor} \varepsilon_i \Gamma^{ab}_{i,\lfloor n/2\rfloor+i}\pi_{ab}(X_i, X_{\lfloor n/2\rfloor+i})\Bigg|\Bigg).$$

Lemma 1

$$\mathbb{E}\left[\sup_{\pi\in\Pi_n}|\widetilde{W}_n(\pi)-W(\pi)|\right]\leq\mathbb{E}[2\mathcal{R}_n(\Pi_{ab,n})].$$

Now we want an asymptotic upper bound for $\mathbb{E}[\mathcal{R}_n(\Pi_{ab})]$. Importantly, we want the bound to depend on the following variance

$$S_{ab}=\mathbb{E}[\Gamma_{i,j}^{2\,ab}].$$

While Kitagawa and Tetenov (2018) and others provide bounds in terms of the max of the scores, Athey and Wager (2021) provide bounds based on the variance and the efficient variance. The next result provides a bound on the Rademacher complexity based on S_{ab}.

Lemma 2 *Under Assumptions 4 and 5*

$$\mathbb{E}[\mathcal{R}_n(\Pi_{ab,n})]=\mathcal{O}\left(\sqrt{\frac{S_{ab}\cdot VC(\Pi_n)}{\lfloor n/2\rfloor}}\right).$$

Now we want to provide asymptotic upper bounds for the first term in (4.1). Escanciano and Terschuur (2022) show that for given $\pi\in\Pi_n$

$$\sqrt{n}(\hat{W}_n(\pi)-\widetilde{W}_n(\pi))\rightarrow_p 0.$$

The next result makes the above uniform in $\pi\in\Pi_n$.

Lemma 3 (Uniform coupling) *Under Assumptions 4 and 5*

$$\sqrt{n}\mathbb{E}[\sup_{\pi\in\Pi_n}|\hat{W}_n(\pi)-\widetilde{W}_n(\pi)|]=\mathcal{O}\left(1+\frac{VC(\Pi_n)}{\lfloor n/2\rfloor^{\min(\lambda_\gamma,\lambda_\nu,\lambda_\alpha)}}\right).$$

Finally, using Lemmas 2 and 3 the following holds.

Theorem 1 *Suppose Assumptions 4 and 5 hold, that Assumption 6 holds with $\beta<\min(\lambda_\gamma,\lambda_\nu,\lambda_\alpha)$. Then*

$$\mathbb{E}[W^*_{\Pi_n}-W(\hat{\pi})]=\mathcal{O}\left(\sqrt{\frac{S_{ab}\cdot VC(\Pi_n)}{\lfloor n/2\rfloor}}\right).$$

5 Empirical application

In our empirical application we study the optimal allocation of children to preschool for our leading welfare functions. We make use of the Panel Study of Income Dynamics (PSID) database which has been following families for nearly 50 years. The nature of this survey allows us to observe a rich set of circumstances and long term outcomes. In 1995, PSID asked about the participation in preschool to adults between 18-30 years hold. Hence, we are able to track long

	Estimate	se	pval	Gini	IOp	Kendall	n
Earnings 30-35	5499	1889	0.004	0.434	0.194	0.159	2730

Table 1: ATE, Gini, IOp and Kendal-τ

term outcomes of these individuals. We take as outcome the average earnings at 30 to 35 years old. We assume selection on observables holds. In particular we condition on sex,birthyear, average parental income in the 5 years before birth, mother's education, father's education and whether the individual is black.

In Table 1 we see the results of estimating the Average Causal Effect (ATE), Gini, IOp and Intergenerational mobility as captured by the rank correlation of parents and child income. To estimate the ATE I use doubly robust scores and XGBoost to estimate the regression functions and propensity scores. Under our assumption of no selection on observables we observe a sizeable and significant positive effect of attending preschool of 5,499$ of added annual earnings. We see that the Gini coefficient is 0.43, in line with official statistics and that IOp is 0.19, meaning that 44% of total inequality can be explained by the circumstances we observe. The Kendall-τ is 0.16 which indicates a positive association between parental and child incomes. We compute optimal treatment rules based on parental income X_{1i} and mother's years of education X_{2i}. For the target in the Kendall-τ welfare we use half of the actual value in Table 1. We use the following policy class

$$\Pi = \{\{X \in \mathcal{X} : X_1 \leq x_1 \,\&\, X_2 \leq x_2\} : (x_1, x_2) \in \{F_{X_1}^{-1}(p) : p = 0.2, 0.4, 0.6, 0.8, 1\} \times \{x_2 = 1, ..., 17\}\}.$$

Table 2 shows the results of the optimal policy rule compared to the situation in which no one is treated. $W0$ is the welfare when no one is treated, W^* is the welfare at the optimal rule. We also see the mean, Gini and IOp when no one is treated and at the optimal treatment rule. Wg^* is the porcentual change in welfare from not treating no one to implementing the optimal treatment.

	W0	Mean0	Gini0 or IOp0	W*	Wg*	Mean*	Gini* or IOp*
Utilitarian	43572			50109	0.15		
Gini	26891	43415	0.38	35641	0.33	49231	0.28
IOp	30814	43193	0.29	38135	0.24	49944	0.24
Kendall	-0.11			-0.01			

Table 2: Welfare, mean, Gini and IOp when no one is treated vs at the optimal treatment rule. Gini or IOp columns have the Gini for the row with the Gini welfare and IOp for the row with the IOp welfare.

Under utilitarian welfare we see that welfare increases a 15% under the optimal rule. Average earnings go from 43,572$ to 50,109$. Under the standard Gini welfare function we observe an increase of 33% in the welfare. The mean increases substantially but not as much as in the utilitarian case and the Gini strongly decreases from 0.38 to 0.28. The IOp aware welfare function is like the utilitarian one and we only see a moderate decrease in IOp from 0.29 to 0.24. Finally, for the Integenerational mobility aware welfare we see that we get much closer to the target under the optimal policy rule. If we do not treat anyone then the (negative) distance between the Kendall-τ and the target is 0.11, however, under the optimal policy rule it decreases sharply to just -0.01. Figures 1-4 show who gets treated under the different policy rules. In the x-axis we divide the parental income in quintiles and in the y-axis we haver mother's years of education.

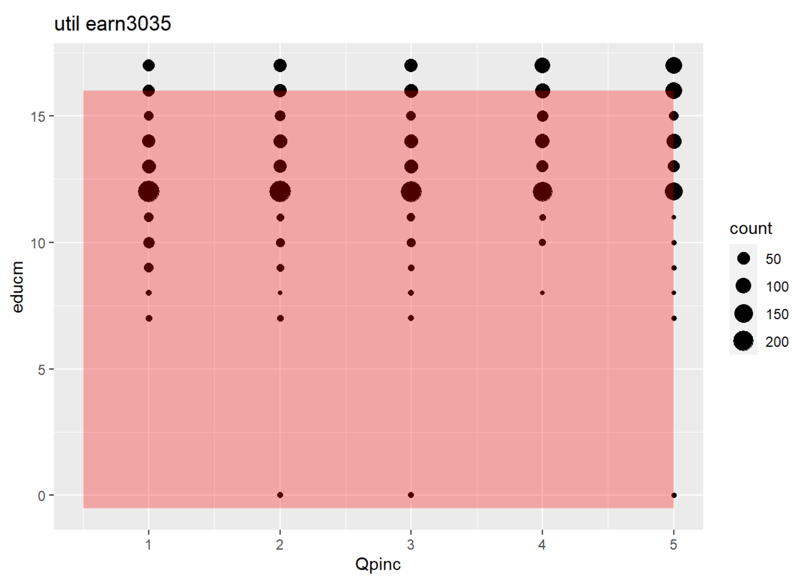

Figure 1: Utilitarian optimal policy rule

We can see that both utilitarian and IOp welfare fucntions prescribe almost everyone to treatment. Inequality averese welfare function restricts treatment to those in the poorest three lowest quintiles of the parental income distribution while the Kendall welfare restricts it to those in the 4 poorest quintiles of parental income and those with mother's education below 14 years.

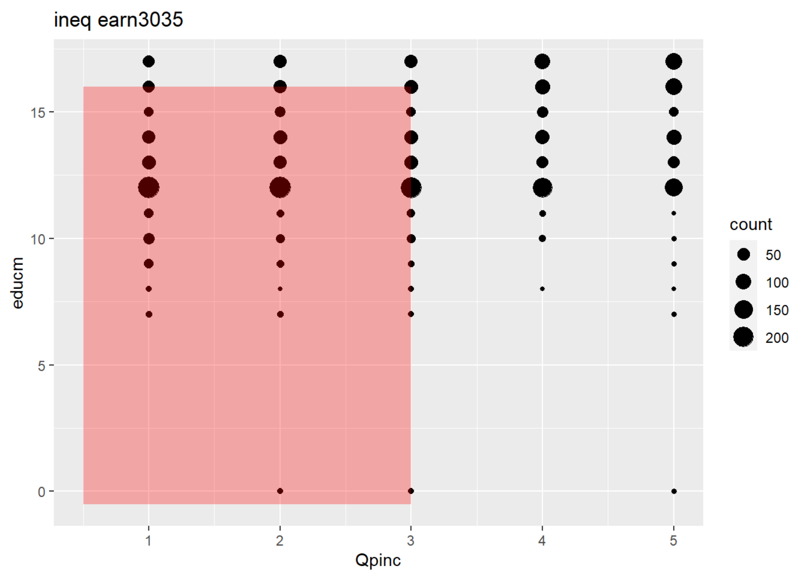

Figure 2: Standard Gini welfare optimal policy rule

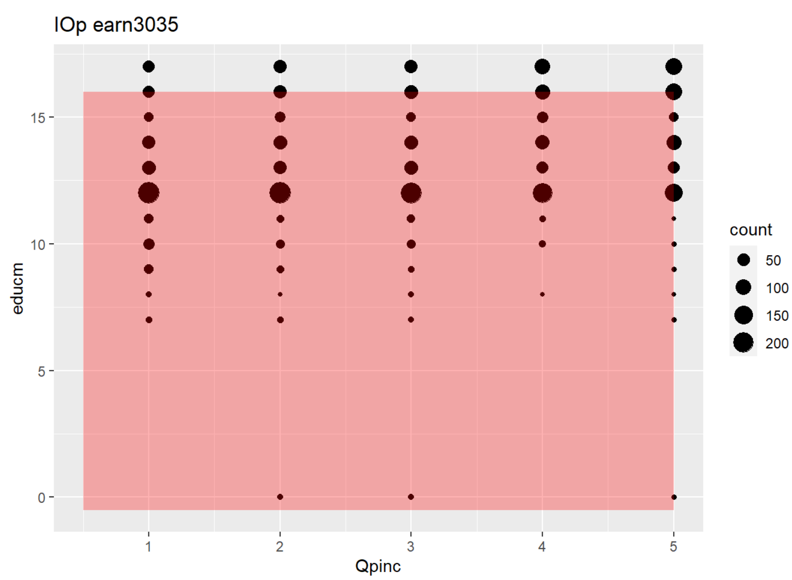

Figure 3: IOp optimal policy rule

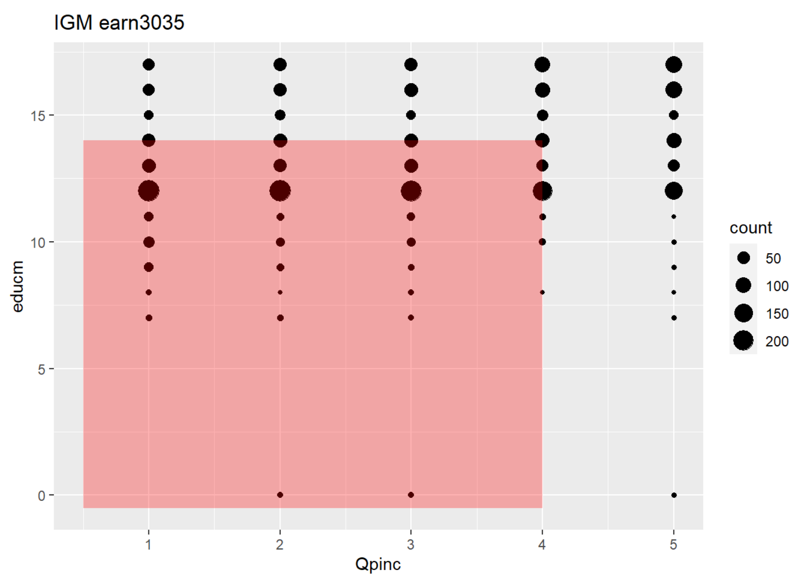

Figure 4: Integenerational mobility welfare optimal policy rule

References

Aaberge, R., T. Havnes, and M. Mogstad (2021): "Ranking intersecting distribution functions," *Journal of Applied Econometrics*, 36, 639–662.

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

Atkinson, A. B. et al. (1970): "On the measurement of inequality," *Journal of economic theory*, 2, 244–263.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014): "Where is the land of opportunity? The geography of intergenerational mobility in the United States," *The Quarterly Journal of Economics*, 129, 1553–1623.

Dalton, H. (1920): "The measurement of the inequality of incomes," *The Economic Journal*, 30, 348–361.

Donaldson, D. and J. A. Weymark (1980): "A single-parameter generalization of the Gini indices of inequality," *Journal of economic Theory*, 22, 67–86.

——— (1983): "Ethically flexible Gini indices for income distributions in the continuum," *Journal of Economic Theory*, 29, 353–358.

Escanciano, J. C. and J. R. Terschuur (2022): "Debiased Semiparametric U-Statistics: Machine Learning Inference on Inequality of Opportunity," *arXiv preprint arXiv:2206.05235*.

Kitagawa, T., M. Nybom, and J. Stuhler (2018): "Measurement error and rank correlations," Tech. rep., cemmap working paper.

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

——— (2021): "Equality-minded treatment choice," *Journal of Business & Economic Statistics*, 39, 561–574.

Mehran, F. (1976): "Linear measures of income inequality," *Econometrica: Journal of the Econometric Society*, 805–809.

Wainwright, M. J. (2019): *High-dimensional statistics: A non-asymptotic viewpoint*, vol. 48, Cambridge university press.

Weymark, J. A. (1981): "Generalized Gini inequality indices," *Mathematical Social Sciences*, 1, 409–430.