

RIF Regression via Sensitivity Curves

Javier Alejo (Universidad de la República)
Gabriel Montes Rojas (Universidad de Buenos Aires, CONICET)
Walter Sosa Escudero (Universidad de San Andrés, CONICET)

Esta versión: Diciembre 2021

Documento de Trabajo Nº 159
Departamento de Economía
Universidad de San Andrés

RIF Regression via Sensitivity Curves

Javier Alejo*
Gabriel Montes-Rojas†
Walter Sosa-Escudero‡

December 3, 2021

Abstract

This paper proposes an empirical method to implement the recentered influence function (RIF) regression of Firpo, Fortin and Lemieux (2009), a relevant method to study the effect of covariates on many statistics beyond the mean. In empirically relevant situations where the influence function is not available or difficult to compute, we suggest to use the *sensitivity curve* (Tukey, 1977) as a feasible alternative. This may be computationally cumbersome when the sample size is large. The relevance of the proposed strategy derives from the fact that, under general conditions, the sensitivity curve converges in probability to the influence function. In order to save computational time we propose to use a cubic splines non-parametric method for a random subsample and then to interpolate to the rest of the cases where it was not computed. Monte Carlo simulations show good finite sample properties. We illustrate the proposed estimator with an application to the polarization index of Duclos, Esteban and Ray (2004).

JEL classification: J01, J31

Keywords: recentered influence function, sensitivity, inequality, polarization

*IECON-Universidad de la República, Gonzalo Ramírez 1926, C.P. 11200, Montevideo, Uruguay. E-mail: javier.alejo@fcea.edu.uy

†IIEP-BAIRES - Universidad de Buenos Aires and CONICET. Email: gabriel.montes@fce.uba.ar

‡Universidad de San Andrés and CONICET, Buenos Aires, Argentina. Email: wsosa@udesa.edu.ar

1 Introduction

The *recentered influence function* (RIF) regression, as proposed by Firpo, Fortin and Lemieux (2009), is a powerful tool to study the impact of changes in covariates on the unconditional distribution of a given outcome variable. Let Y be a random variable with cumulative distribution function F , and $v(F)$ any ‘functional’ of interest related to F . For example, if Y is income, $v(F)$ can be the mean, the Gini index, a quantile, or the poverty rate. The RIF is defined as $\text{RIF}(y, v, F) = v(F) + \text{IF}(y, v, F)$, where $\text{IF}(y, v, F)$ is the *influence function* (IF) (Hampel, 1974) that measures the marginal impact of a particular data point in the support of F in the value of $v(F)$. Influence functions play a key role in the robust statistics literature.

Firpo et al. (2009, 2018) note that since $E[\text{RIF}(Y, v, F)] = v(F)$, by the law of iterated expectations $E_X [E_{Y|X} \text{RIF}(Y, v, F)] = v(F)$, and show that the effect on $v(F)$ that arises from shifting a scalar covariate from X to $X + t$, where $t \downarrow 0$, is given by:

$$\int \frac{dE[\text{RIF}(Y, v, F)|X = x]}{dx} dF(x).$$

Hence, by properly modelling $E[\text{RIF}(Y, v, F)|X = x]$ in a regression fashion, the effect of X on v can be recovered as an ‘average derivative’ of regressing $\text{RIF}(Y, v, F)$ on X . The implementation of the method requires to construct $\text{RIF}(Y, v, F)$ analytically for the functional of interest v and then to regress it on X . In many relevant cases the IF required to obtain $\text{RIF}(Y, v, F)$ is immediately available; Fortin, Lemieux and Firpo (2011) present a useful ‘catalog’ that includes the mean, the quantiles, the variance and the Gini index (see also Essama-Nssah and Lambert (2015) and Cowell and Flachaire (2015)). However, there are many examples where this is not the case. Our paper proposes an alternative in these situations.

In this paper we propose a practical computation method based on the *sensitivity curve* (SC) (Tukey, 1977). This procedure consists in comparing the full sample functional v with that computed when the j –th observation is left out; this is the influence of this particular observation on the empirical version of v . The relevance of the proposed strategy derives from the fact that, under general conditions, the SC converges in probability to the IF (see Nasser and Alam (2006) for a discussion). We provide an intuitive proof of this result.

The SC has some practical advantages over the IF. First, even when analytically available, in many cases the estimation of the IF involves dealing with the problem of selection of the meta-parameters, like bandwidths, which may add further complications. Second, in some relevant cases the IF may be difficult when not impossible to derive analytically. As an example of this case we study the Duclos, Esteban, and Ray (2004) polarization index, where for the general case there is no analytical functional form of the IF (see Appendix A2 for a summary of the construction and motivation of this index). Finally, many relevant examples where the IF can be easily derived involve additive or quasi-additive measures that do not apply to many important situations.

This paper is organized as follows. Section 2 presents the main statistical derivations. Section 3 discusses the cubic spline method to interpolate the SC and considerably reduce computation time. Section 4 provides finite sample Monte Carlo simulations. Section 5 discusses an empirical exercise that shows that the performance of the SC is close to that of the analytical IF.

2 Influence via sensitivity curves

Let $v(F)$ be a real-valued functional, where $v : \mathcal{F}_v \rightarrow \mathbb{R}$ and $\mathcal{F}_v$ is a class of distribution functions such that $F \in \mathcal{F}_v$ if $|v(F)| < \infty$. Consider two cumulative distribution functions (CDFs), F and G , and let $H_{t,F,G} = tG + (1-t)F$, $t \in [0, 1]$. Then, using the Von Mises (1947) expansion:

$$v(H) = v(F) + t \partial v(H_{t,F,G}) / \partial t |_{t=0} + r(t, F, G), \quad (1)$$

with

$$\begin{aligned} \partial v(H_{t,F,G}) / \partial t |_{t=0} &= \lim_{t \downarrow 0} \frac{v(H_{t,F,G}) - v(F)}{t} \\ &= \int \psi(y) d(G - F)(y). \end{aligned} \quad (2)$$

When $G = \Delta_y$ and Δ_y is the CDF of a random variable with probability mass of 1 at y , $\psi(y) = \partial v(H_{t,F,\Delta_y}) / \partial t |_{t=0}$ is the *influence function (IF)* of the functional v , labeled as $IF(y, v, F)$ (see Huber and Ronchetti (2009) for a general discussion; here we are following the derivation in Firpo et al. (2009, p.956)).

Consider now the last term in eq. (1). Following Von Mises (1947):

$$r(t, F, G) = \frac{\tilde{t}^2}{2} \partial^2 v(H_{t,F,\Delta_y}) / \partial t^2 \big|_{t=0}, \quad (3)$$

for some $\tilde{t} \in [0, t]$, where

$$\partial^2 v(H_{t,F,G}) / \partial t^2 \big|_{t=0} = \iint \phi(y, z) d(G - F)(y) d(G - F)(z), \quad (4)$$

with $\phi(y, z)$ a symmetric function; again, see Von Mises (1947, p. 325) for details. Note that if $v(F) = v(cF)$ for all $c > 0$ (scale invariance) then:

$$\begin{aligned} \text{(i)} \quad & \int \psi(y) dF(y) = 0 \\ \text{(ii)} \quad & \iint \phi(y, z) dF(y) dF(z) = 0 \end{aligned}$$

The proof of (i) and (ii) follows from Jaeckel (1972).¹

The *recentered influence function (RIF)*, is defined as $RIF(y, v, F) \equiv v(F) + IF(y, v, F)$, where, trivially, $E[IF(y, v, F)] = v(F)$, from property (i) above. Firpo et al. (2009) develop a RIF-regression framework that is similar to a standard regression except that the dependent variable, Y , is replaced by the IF of the statistic of interest, which allows to estimate the effects of covariates X on $v(F)$.

Unfortunately, not all indicators have an IF with a specific analytical form and thus the RIF-regression may not be practically feasible. Our proposal consists of replacing the IF by the SC.

Let $\{y_i\}_{i=1}^n$ be an iid sample and define $v_n = v(F_n)$ as the sample counterpart of $v(F)$, and let $v_n^{(j)} = v(F_n^{(j)})$ denote the case where j -th observation is left out, then:

$$\begin{aligned} F_n(y) &= \frac{1}{n} \sum_{i=1}^n 1(y_i \leq y) \\ F_n^{(j)}(y) &= \frac{1}{n-1} \sum_{i \neq j} 1(y_i \leq y). \end{aligned}$$

¹Let $G = 2F$, then $H_{t,F,G} = (1+t)F = cF$ and then by the invariance to scale

$$\partial v(H_{t,F,G}) / \partial t \big|_{t=0} = \lim_{t \downarrow 0} \frac{v(cF) - v(F)}{t} = \lim_{t \downarrow 0} \frac{0}{t} = 0.$$

Moreover,

$$\partial^2 v(H_{t,F,\Delta_y}) / \partial t^2 \big|_{t=0} = 0.$$

The *sensitivity curve (SC)* is defined as

$$SC(y_j, v_n, F_n) \equiv n \cdot \left[v(F_n) - v\left(F_n^{(j)}\right) \right]. \quad (5)$$

The key property that links the IF to the SC is the following:

Proposition 1. *Assume that $v(F)$ is twice continuously differentiable with respect to F and $\psi(y)$ and $\phi(y, z)$ exist, and that v is invariant to scale (i.e., $v(F) = v(cF)$ for $c > 0$). Then, $SC(y_j, v_n, F_n) \xrightarrow{P} IF(y_j, v, F)$ as $n \rightarrow \infty$.*

Proof. See the Appendix A1. $\square$

Consequently, if the functional v is smooth enough, the SC can be used instead of the analytical IF. Nasser and Alam (2006) show that Fréchet differentiability is sufficient for consistency. Of course smoothness may be considered a strong requirement. For example, for the case of quantiles $IF(y, Q_\tau, F) = (y - 1[y \leq Q_\tau(F)]) / f_y(Q_\tau(F))$, where $1[\cdot]$ is an indicator function, $f_y(Q_\tau(F))$ is the density of the marginal distribution of y evaluated at the τ -quantile, and $Q_\tau(F)$ is the population τ -quantile of the unconditional distribution of y . The indicator function makes it non twice differentiable.

The *recentered sensitivity curve (RSC)* is defined as:

$$RSC(y_j, v_n, F_n) \equiv v_n + SC(y_j, v_n, F_n)$$

Trivially $RSC \xrightarrow{P} RIF$. Hence, our proposal is to replace RIF with RSC. An then, to apply regression models to approximate distributional effects as in Firpo et al. (2009) RIF regression method.

3 Computation of the SC via cubic splines

The RSC method described above requires to compute the functional v $n+1$ times, that is, the original using the entire sample plus all of the leave-one-out cases (i.e. n). This may be computationally cumbersome when the sample size is large. In order to save computational time we propose to compute the SC for a random sub-sample and then to interpolate to the rest of the domain of y . A widely used method to perform this type of adjustment is splines since they are implemented through a flexible functional form that is linear in parameters. In particular we use the restricted cubic splines

method to interpolate for the value of the RSC for the cases where it was not computed.

Cubic splines are piecewise-polynomial line segments whose function values and first and second derivatives agree at the boundaries where they join. The boundaries of these segments are called knots, and the fitted curve is continuous and smooth at the knot boundaries (see Smith, 1979; Wegman and Wright, 1983; Harrell, 2001, ch. 2). Let k_j , $j = 1, \dots, K$, be the knot values defined in the support of y , then the equation of the cubic spline is

$$J(y) = \beta_0 + \beta_1 y + \beta_2 y^2 + \beta_3 y^3 + \sum_{j=1}^K \gamma_j (y - k_j)_+^3,$$

where $u_+ := \max(u, 0)$. A common problem with cubic spline is that it fits poorly in the tails. One way to deal with this is by restricting $J(y)$ to be linear for $y < k_1$ and $y > k_K$. This requirement is satisfied when $\beta_2 = \beta_3 = 0$, $\sum_{j=1}^K \gamma_j = 0$ and $\sum_{j=1}^K \gamma_j k_j = 0$. Replacing this in the $J(y)$ equation, Durrleman and Simon (1989) show that the restricted cubic spline is then

$$J_R(y) = \beta_0 + \beta_1 y + \sum_{j=1}^{K-2} \gamma_j h_j(y),$$

where

$$h_j(y) = (y - k_j)_+^3 + \frac{k_K - k_j}{k_K - k_{K-1}} (y - k_{K-1})_+^3 + \frac{k_{K-1} - k_j}{k_K - k_{K-1}} (y - k_j)_+^3$$

for $j = 1, \dots, K-2$. Note that $J_R(y)$ is linear in parameters and therefore can be estimated by ordinary least-squares (OLS) methods using y and the $\{h_j\}_{j=1}^{K-2}$ auxiliary variables.

The interpolation of the RSC function proceeds in three steps.

In the first step, we select the knots on the full sample of $\{y_i\}_{i=1}^n$ and create the auxiliary variables $\{h_{1i}, \dots, h_{(K-2)i}\}_{i=1}^n$ that corresponds to the restricted cubic spline method.

In the second step, we consider a random sample without replacement of $\{y_i, h_{1i}, \dots, h_{(K-2)i}\}_{i=1}^n$ denoted by $\{y_{i^*}, h_{1i^*}, \dots, h_{(K-2)i^*}\}_{i^*=1}^{n^*}$, where $n^* < n$. For this random sample we compute the RSC for each of the n^* observations, say $\{\widehat{SC}_{i^*}\}_{i^*=1}^{n^*}$. This is the step that significantly reduces the

computation time (see the simulations in the Monte Carlo section). Moreover, we estimate the parameters $(\beta_0, \beta_1, \gamma_1, \dots, \gamma_{K-2})$ by fitting an OLS regression of $\widetilde{SC}_{i*}$ as a function of $y_{i*}, h_{1i*}, \dots, h_{(K-2)i*}^*$.

Finally, in the third step, we apply the estimated linear regression coefficients to compute the cubic spline interpolation for the full sample,

$$\widehat{SC}_i = \hat{\beta}_0 + \hat{\beta}_1 y_i + \hat{\gamma}_1 h_{1i} + \dots + \hat{\gamma}_{K-2} h_{(K-2)i}.$$

The $\widehat{SC}_i$ interpolated values are then used in the RSC method to compute the effect of covariates on the given functional $v(F)$ (see Orsini and Greenland, 2011, and Newson, 2012, for a discussion of how this interpolation works).

4 Monte Carlo experiments

In this section we run some numerical simulation exercises to evaluate the computational and statistical performance of the proposed method. Throughout this section we use the following baseline model,

$$Y = 20 + X + W,$$

where $X \sim \text{Uniform}(0, 1)$ is the observable covariate and W is the unobservable variable. Then we use the two alternative models:

1. Location-scale model: $W = (1 + X)U$, with $U \sim N(0, 1)$.
2. Location-bimodal model: $W = (D(-4 + U(2 - X)) + (1 - D)(4 + Z(2 - X)))/5$, with (U, Z, D) independent and with distributions $U \sim N(0, 1)$, $Z \sim N(0, 1)$ and D Bernoulli with $\Pr(D = 1) = 0.50$.

In all the exercises in this section we use STATA version 14.1 MP (64-bit) installed on a computer with 16 GB of RAM, an Intel Core i7 processor and Windows 10 operating system.

4.1 Finite sample performance

We use 1000 Monte Carlo simulations to evaluate the estimators' performance and compute Bias, Variance and MSE (mean-squared error). We consider two sample sizes of $n = 500$ and $n = 5000$. To compute the

Table 1: Variance

	n = 500			n = 5000		
	RIF	RSC	RSC(sp)	RIF	RSC	RSC(sp)
(i): location-scale model						
Population	3.003	3.003	3.003	3.003	3.003	3.003
Mean	2.999	3.017	3.189	2.995	2.997	2.981
Bias	-0.003	0.015	0.187	-0.008	-0.006	-0.021
Var	0.369	0.374	0.460	0.035	0.035	0.035
MSE	0.369	0.374	0.494	0.035	0.035	0.036
(ii): location-bimodal model						
Population	-0.120	-0.120	-0.120	-0.120	-0.120	-0.120
Mean	-0.117	-0.118	-0.224	-0.119	-0.120	-0.131
Bias	0.003	0.002	-0.104	0.000	0.000	-0.011
Var	0.012	0.012	0.014	0.001	0.001	0.001
MSE	0.012	0.012	0.025	0.001	0.001	0.001

Note: own calculations using 1000 Monte Carlo simulations.

population parameter we use the DGPs with 10 million observations and where we compute the numerical derivative of a change $x' = x + \epsilon$, that is, $\frac{v(F) - v(F')}{\epsilon}$ where F and F' are the induced distribution functions of the corresponding DGP with X and $X + \epsilon$, respectively, and with $\epsilon = 0.0001$.

We evaluate 3 different functionals: variance (Table 1), Gini coefficient (Table 2) and DER polarization index with $\alpha = 0.5$ (Tables 3 and 4). For the variance and Gini we have an analytical formula of the IF. As such we compute the RIF effect together with the RSC proposed method. For the DER polarization we can only report the RSC effect (see the Appendix A2 for a description of this index). In all cases, for the RSC computation we report the full sample RSC method and the splines approximation (RSC(sp)). For the latter we use 100 points for $n = 500$ and 1000 for $n = 5000$.

Table 1 shows the performance of the proposed method for computing the marginal effect on the variance. The simulations show that the proposed RSC method has a similar performance to that of RIF, which is close to the population parameter in terms of Bias and MSE. The spline approximation has a weaker performance of $n = 500$ but it is similar to the full sample RSC for $n = 5000$.

Table 2 shows the performance of the proposed method for computing

Table 2: Gini

	n = 500			n = 5000		
	RIF	RSC	RSC(sp)	RIF	RSC	RSC(sp)
(i): location-scale model						
Population	2.473	2.473	2.473	2.473	2.473	2.473
Mean	2.517	2.521	2.430	2.512	2.512	2.550
Bias	0.044	0.048	-0.043	0.039	0.039	0.077
Var	0.252	0.254	0.277	0.023	0.023	0.024
MSE	0.254	0.257	0.278	0.025	0.025	0.030
(ii): location-bimodal model						
Population	-0.322	-0.322	-0.322	-0.322	-0.322	-0.322
Mean	-0.355	-0.358	-0.479	-0.359	-0.360	-0.342
Bias	-0.033	-0.035	-0.156	-0.037	-0.037	-0.020
Var	0.031	0.031	0.036	0.003	0.003	0.003
MSE	0.032	0.032	0.061	0.004	0.004	0.003

Note: own calculations using 1000 Monte Carlo simulations.

the marginal effect on the Gini coefficient. The results are in line with those for the variance: the RSC has a good performance relative to the RIF. For this case, however, the RSP(sp) approximation is much closer to the full sample RSC, and as such there is a minimum loss in efficiency for using the Spline interpolation.

Finally we consider the analysis of the DER polarization index with $\alpha = 0.5$. As discussed above there is no analytical IF for this model, and therefore the use of the RSC is the only alternative to evaluate the effect of the covariates on the DER index. For this case, we also evaluate alternative models for the RSC regression models. Given that we will not be able to derive the functional form of the conditional model of the IF conditional on X , we compute three different alternatives: linear ($\mathbf{a} + \mathbf{b}X$), quadratic ($\mathbf{a} + \mathbf{b}X + \mathbf{c}X^2$) and cubic polynomials ($\mathbf{a} + \mathbf{b}X + \mathbf{c}X^2 + \mathbf{d}X^3$). Then we compute the average partial effects, that is, $\mathbf{b}$ for the linear case, $\mathbf{b} + 2\mathbf{c}\bar{X}$ for quadratic and $\mathbf{b} + 2\mathbf{c}\bar{X} + 3\mathbf{d}\bar{X}^2$ for the cubic polynomial case.

Table 3 shows the simulation results for the full-sample RSC computation and Table 4 for the Spline interpolation. The location-scale model works similarly across methods, with large reduction in Bias and MSE when the largest sample size is used. The location-bimodal model, however, shows

Table 3: DER($\alpha = 0.5$) - RSC

	n = 500			n = 5000		
	Linear	Quadratic	Cubic	Linear	Quadratic	Cubic
(i): location-scale model						
Population	2.153	2.153	2.153	2.153	2.153	2.153
Mean	2.249	2.249	2.193	2.226	2.226	2.146
Bias	0.096	0.096	0.040	0.073	0.073	-0.007
Var	0.217	0.219	0.687	0.020	0.019	0.066
MSE	0.226	0.228	0.689	0.025	0.025	0.066
(ii): location-bimodal model						
Population	0.291	0.291	0.291	0.291	0.291	0.291
Mean	0.441	0.443	0.106	0.657	0.657	0.229
Bias	0.149	0.151	-0.186	0.366	0.366	-0.063
Var	0.190	0.186	0.605	0.024	0.023	0.082
MSE	0.212	0.209	0.640	0.158	0.157	0.086

Note: own calculations using 1000 Monte Carlo simulations.

considerable heterogeneity across models. In most cases the quadratic approximation seems to correctly capture the effect of a marginal effect of X on the DER index. For the Splines, the sample size requirement seems to be more demanding than in previous models.

4.2 Computing time

We analyze the goodness of fit of the spline interpolation by simulating a random realization of $n = 1000$ using the location-scale model. Figure 1 shows the RSC computed with the complete sample together with the spline interpolation RSC(sp) using a random 10% of the original sample. Although the RSC of the DER(0.5) seems to be quite complicated to approximate compared to those of Gini and variance, the adjustment of the spline seems to be reasonable for the three indicators analyzed.

Table 5 shows the average computation time of the RSC and RSC(sp) with different sample sizes. For this exercise we use 50 random samples generated with the location-scale model. In all cases, a subsample of 10% of the original sample was considered for the RSC(sp) interpolation.


As expected, for all sample sizes, the fastest RSC to compute is for the variance, while the slowest is the DER(0.5) index, since it involves a non-

Table 4: DER($\alpha = 0.5$) - RSC (sp)

	n = 500			n = 5000		
	Linear	Quadratic	Cubic	Linear	Quadratic	Cubic
(i): location-scale model						
Population	2.153	2.153	2.153	2.153	2.153	2.153
Mean	2.004	2.004	1.970	2.271	2.271	2.202
Bias	-0.149	-0.149	-0.183	0.118	0.119	0.049
Var	0.270	0.271	0.798	0.019	0.019	0.067
MSE	0.292	0.293	0.832	0.033	0.033	0.069
(ii): location-bimodal model						
Population	0.291	0.291	0.291	0.291	0.291	0.291
Mean	-0.971	-0.970	-1.152	0.343	0.343	0.084
Bias	-1.263	-1.262	-1.443	0.052	0.052	-0.208
Var	0.162	0.161	0.484	0.027	0.026	0.076
MSE	1.756	1.753	2.568	0.029	0.029	0.119

Note: own calculations using 1000 Monte Carlo simulations.

Figure 1: Comparison of RSC and RSC(sp) fit



Notes: own calculations using a random draw with sample size $n = 1000$, $n^* = 100$, STATA 14.1 MP (64-bit), 16 GB of RAM, an Intel Core i7 processor and Windows 10 operating system.

Table 5: Average computing time (seconds)

Index	Sample	RSC	RSC(sp)	RSC(sp)/RSC
Variance	$500 \leq n \leq 1100$	0.17	0.03	19.7%
	$1100 < n \leq 1800$	0.43	0.04	9.4%
	$1800 < n \leq 2500$	0.85	0.06	6.8%
Gini	$500 \leq n \leq 1100$	6.14	0.61	9.9%
	$1100 < n \leq 1800$	12.11	1.13	9.3%
	$1800 < n \leq 2500$	18.4	1.69	9.2%
DER(0.5)	$500 \leq n \leq 1100$	20.2	0.66	3.3%
	$1100 < n \leq 1800$	63	1.42	2.3%
	$1800 < n \leq 2500$	128	2.49	1.9%

Note: own calculations using 50 random draws with **STATA** 14.1 MP (64-bit), 16 GB of RAM, an Intel Core i7 processor and Windows 10 operating system.


parametric estimate of a density. The time required to compute the complete RSC increases markedly with the sample size; however, the estimate based on the spline RSC(sp) increases only slightly. For example, for samples between 500 and 1100 observations, computing the RSC of the variance using splines represents 19.7% of the time it takes with the complete sample, while with larger sample sizes this percentage represents just under 10%. This saving in computational time is similar for the Gini index (9.5% average) and definitely more noticeable for the DER (2.5% average). Figure 2 clearly shows the relative computational advantage of using spline interpolation as larger samples are used.

5 Empirical illustration

This section presents empirical applications. We first compare the empirical performance of RIF and RSC for the variance and the Gini index, for which the IF can be obtained analytically. Then we add the DER polarization index (Duclos, Esteban, and Ray, 2004) where an explicit analytical closed-form solution for IF is not available (see the Appendix A2 for a description of this index).

We use an extract from the Merged Outgoing Rotation Group of the Current Population Survey of 1983, 1984 and 1985 for males only. More details about the data can be found in Lemieux (2006). The variable of in-

Figure 2: Comparison of computing times



Note: own calculations using 50 random draws with **STATA** 14.1 MP (64-bit), 16 GB of RAM, an Intel Core i7 processor and Windows 10 operating system.

Table 6: Wage Inequality

	RIF	Variance RSC	RSC(sp)	RIF	Gini RSC	RSC(sp)
Union	-15.51*** (0.189)	-15.51*** (0.162)	-16.42*** (0.157)	-6.62*** (0.055)	-6.62*** (0.046)	-6.89*** (0.046)
Education	1.48*** (0.030)	1.48*** (0.037)	1.22*** (0.035)	-0.49*** (0.009)	-0.49*** (0.010)	-0.56*** (0.010)
Experience	0.15*** (0.007)	0.15*** (0.008)	0.10*** (0.008)	-0.10*** (0.002)	-0.10*** (0.002)	-0.11*** (0.002)
Married	-10.13*** (0.190)	-10.13*** (0.185)	-10.98*** (0.181)	-5.24*** (0.056)	-5.24*** (0.057)	-5.44*** (0.057)
Non-white	0.91*** (0.262)	0.91*** (0.248)	1.33*** (0.246)	1.53*** (0.077)	1.53*** (0.079)	1.65*** (0.080)
Observations	266,956	266,956	266,956	266,956	266,953	266,956

Source: Extract from the Merged Outgoing Rotation Group of the Current Population Survey of 1983, 1984 and 1985. Notes: Standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; (sp) indicates that the RSC was estimated using a cubic spline with a random subsample of 1000 points; all estimates are multiplied by 100.

terest is Y , the hourly wage, and the covariates X are an indicator of whether the individual is unionized, years of education, whether he is married, non-white, his experience. We use a linear specification in all regressions but, given the results of the previous section, in the case of the polarization index we add a more flexible specification that incorporates the squares and non-trivial interactions of all the covariates.

Obtaining the RSC for each observation using the leave-one-out method can be computationally intensive if n is too large since it requires a separate calculation for each observation. Therefore, we also consider computing the RSC by interpolating an estimated spline using 1000 random points in the distribution of Y (this is denoted as $RSC(sp)$).

Table 6 shows results for the variance and the Gini index. Remarkably, the differences between the RIF and RSC regressions are negligible. Interestingly, the approximation obtained through the spline intrapolation seems to be accurate, suggesting that it is a convenient computational strategy relative to the leave-one-out method.

Table 7 also shows results for the DER polarization indexes, for which the Gini columns correspond to a particular case ($\alpha = 0$), and for proper polarization we set $\alpha = 0.5$ following Duclos et al. (2004). We stress the fact that the IF function is not available for this case, hence we obtain results based on the RSC solely. Note that in this case the coefficients of the linear

Table 7: Wage Polarization: DER($\alpha = 0.50$)

	RSC		RSC(sp)	
	Linear	Quad. Form.	Linear	Quad. Form.
Union	-3.38*** (0.027)	-3.43*** (0.030)	-3.43*** (0.025)	-3.45*** (0.027)
Education	-0.35*** (0.006)	-0.35*** (0.006)	-0.40*** (0.005)	-0.39*** (0.005)
Experience	-0.07*** (0.001)	-0.06*** (0.001)	-0.08*** (0.001)	-0.06*** (0.001)
Married	-3.28*** (0.034)	-2.42*** (0.035)	-3.10*** (0.031)	-2.36*** (0.033)
Non-white	0.99*** (0.050)	1.04*** (0.049)	0.96*** (0.044)	1.01*** (0.043)
Observations	266,956	266,956	266,956	266,956

Source: Extract from the Merged Outgoing Rotation Group of the Current Population Survey of 1983, 1984 and 1985. Notes: Standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; (sp) indicates that the RSC was estimated using a cubic spline with a random subsample of 1000 points; all estimates are multiplied by 100.

model give similar results to the average partial effects of the more flexible model. Considering the results of the simulations in the previous section, this is probably due to the large sample size since the RSC approximation to the RIF is more precise. Again, the computationally convenient spline approximation produces similar results than when RSC is computed directly. Even though a detailed study of the effects on inequality and polarization exceeds the scope of this note, we remark that all factors reduce both measures (i.e., higher levels education predict less *unconditionally* inequality and polarization), and that effects are stronger for inequality.

References

- Cowell, F.A., Flachaire, E. 2015. Statistical Methods for Distributional Analysis. In Anthony B. Atkinson and Francois Bourguignon (eds.), *Handbook of Income Distribution*. Amsterdam: Elsevier.
- Davies, J.B., Fortin, N.M., Lemieux, T. 2017. Wealth inequality: Theory, Measurement and Decomposition. *Canadian Journal of Economics/Revue Canadienne d'Économique* 50(5): 1224-1261.
- DiNardo, J., Fortin, N.M., Lemieux, T. 2017. Labor Market Institutions

- and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica* 64(5): 1001-1044.
- Duclos, J.-Y., Esteban, J., Ray, D. 2004. Polarization: Concepts, Measurement, Estimation. *Econometrica* 72(6): 1737-1772.
- Durrleman, S., Simon, R. 1989. Flexible Regression Models with Cubic Splines. *Statistics in Medicine* 8(5): 551-561.
- Essama-Nssah, B., Lambert, P.J. 2015. Chapter 6: Influence Functions for Policy Impact Analysis. In John A. Bishop and Rafael Salas (eds.), *Inequality, Mobility and Segregation: Essays in Honor of Jacques Silber*, pp.135-159. Bigley, UK: Emerald Group Publishing Limited.
- Firpo, S.P., Fortin, N.M., Lemieux, T. 2009. Unconditional Quantile Regressions. *Econometrica* 77(3): 953-973.
- Firpo, S.P., Fortin, N.M., Lemieux, T. 2018. Decomposing Wage Distributions Using Recentered Influence Function Regressions. *Econometrics* 6(3): 41.
- Fortin, N.M., Lemieux, T., Firpo, S.P. 2011. Decomposition Methods in Economics. In Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics*. Amsterdam: Elsevier.
- Gasparini, L., Horenstein, M., Molina, E., Olivieri, S. 2008. Income Polarization in Latin America: Patterns and Links with Institutions and Conflict, *Oxford Development Studies*, 36: 461-484.
- Hampel, F. 1974. The Influence Curve and its Role in Robust Estimation. *Journal of the American Statistical Association* 69(346): 383-393.
- Harrell, F. E., Jr. 2001. *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. New York: Springer.
- Huber, P., Ronchetti, E.M. 2009. *Robust Statistics* (2nd edition). Wiley.
- Jaekel, L.A. 1972. Estimating regression coefficients by minimizing the

dispersion of the residuals. *Annals of Mathematical Statistics* 43, 1449-1458.

Lemieux, T. 2006. Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *American Economic Review* 96(3): 461-498.

Nasser, M., Alam, M. 2006. Estimators of Influence Function. *Communications in Statistics - Theory and Methods*, 35(1), 21-32.

Newson, R. B. 2012. Sensible parameters for univariate and multivariate splines. *Stata Journal*, 12: 479–504

Orsini, N., and S. Greenland. 2011. A procedure to tabulate and plot results after flexible modeling of a quantitative covariate. *Stata Journal*, 11, 1–29.

Smith, P. L. 1979. Splines as a useful and convenient statistical tool. *American Statistician*, 33, 57–62.

Tukey, J.W. 1977. Exploratory Data Analysis, Addison-Wesley, Reading, MA.

von Mises, R. 1947. On the Asymptotic Distribution of Differentiable Statistical Functions. *Annals of Mathematical Statistics* 18(3): 309-348.

Wegman, E. J., and I. W. Wright. 1983. Splines in statistics. *Journal of the American Statistical Association*, 78, 351–365.

Appendix A1

Proof of Proposition 1.

Using eq. (1) with F_n and $F_n^{(j)}$ for the case of $t = 1$:

$$v(F_n) = v(F_n^{(j)}) + \int \psi_n(y) d(F_n - F_n^{(j)})(y) + r(\tilde{t}, F_n, F_n^{(j)}), \quad (6)$$

for some $\tilde{t} \in [0, 1]$. Note that $\psi_n(y) = IF(y, v, F_n) \xrightarrow{P} \psi(y)$ by continuity of the probability limit.

Now note that $n[F_n - F_n^{(j)}] = 1(y_j < y) + O_p(1)$ because

$$\begin{aligned} F_n(y) &= \frac{1}{n} 1(y_j \leq y) + \frac{n-1}{n} F_n^{(j)}(y), \\ F_n(y) - F_n^{(j)}(y) &= \frac{1}{n} 1(y_j \leq y) + \frac{n-1}{n} F_n^{(j)}(y) - F_n^{(j)}(y), \\ F_n(y) - F_n^{(j)}(y) &= \frac{1}{n} 1(y_j \leq y) - \frac{1}{n} F_n^{(j)}(y). \end{aligned}$$

That is,

$$n[F_n(y) - F_n^{(j)}] = 1(y_j \leq y) - a_n, \quad (7)$$

with $a_n = F_n^{(j)}(y) \xrightarrow{P} F(y)$ by the Law of Large Numbers.

Then,

$$\begin{aligned} n \cdot [v(F_n) - v(F_n^{(j)})] &= \int \psi_n(y) d(1(y_j \leq y) - a_n)(y) + n \cdot r(\tilde{t}, F_n, F_n^{(j)}) \\ n \cdot [v(F_n) - v(F_n^{(j)})] &= \int \psi_n(y) d(1(y_j \leq y))(y) - \int \psi_n(y) d(a_n)(y) + n \cdot r(\tilde{t}, F_n, F_n^{(j)}) \end{aligned} \quad (8)$$

Using the fact that $1(y_j \leq y)$ is the Dirac function, the first term of eq. (8) is

$$\int \psi_n(y) d(1(y_j \leq y))(y) = \psi_n(y_j) \xrightarrow{P} \psi(y_j).$$

Noting that $a_n \xrightarrow{P} F(y)$ and $\psi_n(y) \xrightarrow{P} \psi(y)$, by continuity of the probability limit, the second term of (8) becomes

$$\text{plim} \int \psi_n(y) d(a_n)(y) = \int \psi(y) dF(y) = 0,$$

because of property (i). Then,

$$\text{plim} \int \psi(y) d(1(y_j \leq y) + a_n)(y) = \int \psi(y) d(1(y_j \leq y))(y) = \psi(y).$$

It remains to study the third term in (8). From (3),

$$\begin{aligned} \mathbf{n} \cdot \mathbf{r} \left(\tilde{\mathbf{t}}, F_{\mathbf{n}}, F_{\mathbf{n}}^{(j)} \right) &= \mathbf{n} \cdot \frac{\tilde{\mathbf{t}}^2}{2} \iint \psi(\mathbf{y}, z) d \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) (\mathbf{y}) d \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) (z) \\ \mathbf{n} \cdot \mathbf{r} \left(\tilde{\mathbf{t}}, F_{\mathbf{n}}, F_{\mathbf{n}}^{(j)} \right) &= \frac{1}{\mathbf{n}} \cdot \frac{\tilde{\mathbf{t}}^2}{2} \iint \psi(\mathbf{y}, z) d \left[\mathbf{n} \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) \right] (\mathbf{y}) d \left[\mathbf{n} \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) \right] (z) \end{aligned}$$

for some $\tilde{\mathbf{t}} \in [0, 1]$. Then using (7) and property (ii),

$$\text{plim} \iint \phi(\mathbf{y}, z) d \left[\mathbf{n} \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) \right] (\mathbf{y}) d \left[\mathbf{n} \left(F_{\mathbf{n}} - F_{\mathbf{n}}^{(j)} \right) \right] (z) = \phi(\mathbf{y}_j, \mathbf{y}_j).$$

Then it follows that

$$\text{plim} \left\{ \mathbf{n} \cdot \mathbf{r} \left(\tilde{\mathbf{t}}, F_{\mathbf{n}}, F_{\mathbf{n}}^{(j)} \right) \right\} = \left(\text{plim} \frac{1}{\mathbf{n}} \right) \cdot \frac{\tilde{\mathbf{t}}^2}{2} \phi(\mathbf{y}_j, \mathbf{y}_j) = 0.$$

Then, the result follows,

$$\text{plim} \frac{v(F_{\mathbf{n}}) - v \left(F_{\mathbf{n}}^{(j)} \right)}{1/\mathbf{n}} = \psi(\mathbf{y}_j) = \text{IF}(\mathbf{y}_j, v, F). \text{ QED}$$

Appendix A2: Polarization index

We motivate the case of a model where the IF is not available: the DER polarization index (Duclos, Esteban, and Ray, 2004).

Polarization is an important welfare concept in economics and political science. Intuitively, it measures the tension between individuals in a society, that depends positively on how distant individuals are between groups (alienation) and how close they are within a group (identification). From this perspective, a standard measure of inequality like the Gini index focuses on just the first component. Duclos et al. (2004) provide a full axiomatic framework that leads to a logically coherent measure of polarization. For a detailed empirical study on polarization for the case of Latin America and the Caribbean, see Gasparini et al. (2008).

Let $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{\mathbf{n}}$ be an iid sample of incomes, ordered from lowest to highest. Duclos et al. (2004) propose the following empirical measure of polarization:

$$P_\alpha = \frac{1}{n} \sum_{i=1}^n \hat{f}(y_i)^\alpha \hat{a}(y_i)$$

where $\hat{a}(y_i) = \hat{\mu} + y_i (n^{-1}(2i-1) - 1) - n^{-1} \left(2 \sum_{j=1}^{i-1} y_j + y_i \right)$, $\hat{\mu}$ is the sample mean and $\hat{f}(y_i)$ is an estimate of the density of incomes. The parameter α is set exogenously and plays a key role in characterizing polarization. As a matter of fact, when $\alpha = 0$ polarization reduces to the Gini index (note that for this particular case the IF is available). Larger values of α result in the index giving relatively more importance to identification, that is, to how close individuals are ‘surrounded’ by others of similar income. The axiomatic approach of Duclos et al. (2004) imposes lower and upper bounds to the values α may take in practice.