Explaining Disparities at Quantiles: An Augmented Kitagawa-Oaxaca-Blinder Decomposition

Derick S. Baum and Xiang Zhou

14 September, 2023

1 Introduction

Much social science research is devoted to studying inequalities between groups. These works can point to potential policy interventions when group membership is manipulable — for example, determined by assignment to training programs (Ashenfelter 1978) — and unveil patterns and mechanisms underlying "durable inequalities" when groups are based on inherent traits like race and sex (Tilly 1998). A variety of statistical methods have lent themselves to research on group inequalities. Prominent among these is the Kitagawa-Oaxaca-Blinder (KOB) decomposition method (Kitagawa 1955; Blinder 1973; Oaxaca 1973). Its primary motivation is to partition the mean group disparity in an outcome into a component that can be accounted for by group differences in explanatory variables and an unexplained component. Researchers have used it to investigate gender disparities in earnings (Mandel and Semyonov 2014), differences in leisure time between married and non-married mothers (Pepin, Sayer, and Casper 2018), and inequality in unemployment rates in the early months of the COVID-19 pandemic across groups based on race, sex, education, and age (Montenovo et al. 2022) — among many other applications.

While the original KOB method rested on linear models to decompose the mean differences, extensions have proposed non-parametric alternatives (DiNardo, Fortin, and Lemieux 1996; Barsky et al. 2002) and approaches to examining distributional statistics other than the mean, such as the variance (Lemieux 2006). More recently, Firpo, Fortin, and Lemieux (2009, 2018) introduced a general framework for performing a decomposition on any distributional statistic using the recentered influence function (RIF). Their method encompasses the traditional decomposition for the mean while promptly enabling the study of other statistics, such as the variance, the Gini coefficient, and quantiles.

Researchers have so far focused on gaps in group-specific quantiles (e.g., Firpo, Fortin, and Lemieux 2018; Rios-Avila and Hirsch 2014). Here, interest lies in explaining the disparity between a quantile of the distribution among members of one group and the corresponding quantile of the distribution for the other group. However, one of the main applications of the RIF is its use in unconditional quantile regressions (UQRs) to estimate the effect of variables across quantiles of the overall outcome distribution (Killewald and Bearak 2014). We leverage this application to examine overall quantile gaps. As shown below, these can be understood as group disparities in the probabilities of reaching benchmarks of the overall distribution. Studying them may reveal heterogeneous stratification patterns across the distribution, which analyses focused on mean differences or group-specific quantile disparities may overlook. For example, while policy interventions may promote equality in the lower and middle wage brackets, they may not address the advantages that make privileged groups disproportionately likely to reach the top of the distribution. We decompose the overall quantile gaps into explained and unexplained components to investigate such heterogeneity.

We also use the quantile decompositions to extend the classical KOB decomposition for mean differences. In particular, we show that the mean gap and its explained and unexplained components are averages of the corresponding quantities from the quantile decompositions across the overall distribution. This leads to the augmented KOB (aKOB) method. It partitions the mean gap in two dimensions: horizontally into explained and unexplained components — as intended by the classical method; and vertically across quantiles. The newly added vertical dimension can reveal heterogeneity in the relative importance of covariates in accounting for disparities across the distribution, which the unidimensional classical method cannot capture.

In discussing estimation strategies for the quantities involved in the decomposition, we start by reviewing the commonly used weighting and regression-imputation approaches. We then introduce a doubly robust estimator that combines the other two. To the best of our knowledge, this estimator has not been proposed before in the decomposition literature.

Our empirical example revisits the Current Population Survey (CPS) data on union-nonunion wage differences among a sample of U.S. males, used by Firpo, Fortin, and Lemieux (2009) to illustrate UQRs. The effect of unions on wages has drawn considerable interest among economists and sociologists (e.g., Card, Lemieux, and Riddell 2004; Freeman 1980; Western and Rosenfeld 2011). Our findings corroborate many of those from previous studies. Namely, unions appear to suppress inequality by strengthening low- and middle-wage workers but have negative effects on wages at the top of the distribution.

When we examine the net contributions of different variables, the aKOB decomposition reveals substantial heterogeneity across the wage distribution. At the low end, the union premium is small, and the higher experience levels of union workers explain much of their higher earnings. In the middle of the distribution, differences in covariates account for little of the gap, indicating a large union premium — partly because experience and education have offsetting effects. To some extent, the positive association between experience and wages helps explain higher wages among the more experienced union workers. This is counterbalanced, however, by the positive impact of education on wages and the higher education levels among nonunion workers. The dynamics shift at the top of the distribution, where union workers earn lower wages than nonunion workers. There, the advantages of higher education overshadow the union premium. Although unions may help level the playing field among middle-wage workers, educational inequality prevents equal access to high wages.

In the following section, we explain the logic of the classical KOB decomposition for the mean differences and discuss the three estimators above. We then draw on UQRs to present the concept of overall quantile disparities. Next, we combine the mean and quantile decompositions to propose the aKOB decomposition. The empirical section illustrates our method using the CPS data. The final section concludes.

2 Overview of the Classical KOB Decomposition

We begin with an overview of the classical KOB decomposition. We present a nonparametric formulation throughout, making no assumptions about the functional form of the relationship between the variables. While analysts often assume a linear specification — as we do in our empirical illustration below — the nonparametric formulation demonstrates that the decomposition is amenable to different modeling strategies.

2.1 Nonparametric Formulation of the Mean Decomposition

Let Y denote some outcome of interest, A a binary indicator for group membership, and X a set of explanatory variables that may account for the mean difference between the groups. The decomposition relies on the law of iterated expectations, which allows us to calculate the expected value (or mean, denoted by the $\mathbb{E}[.]$ operator) of a variable using information about another variable. Here, we incorporate information about X to compute the mean of Y: $\mathbb{E}[Y] = \mathbb{E}\left[\mathbb{E}[Y|X]\right]$. In words, the mean of Y equals the expected value of the conditional mean of Y given X. Since this concept is essential to the logic of decomposition, we illustrate it below with a simple example.

Suppose we only have five observations and that X comprises a single explanatory variable whose possible values are 1, 2, and 3. X equals 1 for the first two observations, 2 for the next two, and 3 for the last one. Further assume that the Y values for the observations are 1, 2, 3, 4, and 5, respectively. Among the observations with $X = 1$, the mean of Y is $(1 + 2)/2 = 1.5$. Among those with $X = 2$, the mean is $(3 + 4)/2 = 3.5$. Since there is only one observation with $X = 3$, the mean for this stratum is the outcome value for that observation (5). We can compute the overall mean of Y by taking the weighted average of those means, where the weights are the proportions of observations with the corresponding X value. $X = 1$

for two observations, $X = 2$ for another two, and $X = 3$ for one observation, so the weights are 2/5, 2/5, and 1/5. Therefore, the mean of Y is given by $(2/5)1.5 + (2/5)3.5 + (1/5)5 = 3$. By conditioning on $A = a$ in the inner and outer expectations of the original formula, we can adapt the law of iterated expectations to the group means: $\mathbb{E}[Y|A = a] = \mathbb{E}\left[\mathbb{E}[Y|A = a, X]|A = a\right]$, where $a = 1$ for members and $a = 0$ for non-members.

The mean gap (Δ) equals $\mathbb{E}[Y|A = 1] - \mathbb{E}[Y|A = 0]$. Using the law of iterated expectations, we can decompose it as follows:

$$
\begin{aligned}
\Delta &= \mathbb{E}[Y|A = 1] - \mathbb{E}[Y|A = 0] \\
&= \mathbb{E}[\mathbb{E}[Y|A = 1, X]|A = 1] - \mathbb{E}[\mathbb{E}[Y|A = 0, X]|A = 0] \\
&= (\mathbb{E}[\mathbb{E}[Y|A = 1, X]|A = 1] - \mathbb{E}[\mathbb{E}[Y|A = 1, X]|A = 0]) \\
&+ (\mathbb{E}[\mathbb{E}[Y|A = 1, X]|A = 0] - \mathbb{E}[\mathbb{E}[Y|A = 0, X]|A = 0]) \\
&= \underbrace{\int \mu_1(x) \left[dF_{X|A=1}(x) - dF_{X|A=0}(x)\right]}_{\Delta_X \;\; \text{(Explained)}} + \underbrace{\int \left[\mu_1(x) - \mu_0(x)\right] dF_{X|A=0}(x)}_{\Delta_U \;\; \text{(Unexplained)}} .
\end{aligned} \tag{1}
$$

The last line uses integrals to represent the expectations for a more concise exposition. $\mu_a(x)$ is the conditional mean of Y given $A = a$ and $X = x$, and $F_{X|A=a}(.)$ is the cumulative distribution function of X given $A = a$. The derivation above decomposes the gap into a component explained by differences in the covariate distributions ($dF_{X|A=1}(x) - dF_{X|A=0}(x)$) and a component attributable to differences in the conditional mean functions ($\mu_1(x) - \mu_0(x)$). We call them explained (Δ_X) and unexplained (or residual) components (Δ_U).

The key quantity enabling the decomposition is $\mathbb{E}\left[\mathbb{E}\left[Y|A = 1, X\right]|A = 0\right]$, which we add and subtract to the original expression for the mean difference to perform the partition into explained and unexplained components. It proposes a hypothetical scenario in which the distribution of observed characteristics among group members becomes the same as the distribution for non-members; or, equivalently, one in which non-members have the same returns to their observed characteristics as group members, determined by the conditional mean function $\mu_1(x)$. Because it represents a fiction, we refer to this quantity as a "counterfactual."[1]

The explained component takes out $\mu_1(x)$ as the common factor between the mean among group members and the counterfactual quantity, yielding a expression that depends on the difference in the covariate distributions between the two groups. By contrast, the unexplained component factors out the distribution of observed characteristics among non-members, leaving as the remaining quantity in the expression the difference in conditional mean functions. Note that we could have performed the decomposition using the quantity $\mathbb{E}\left[\mathbb{E}\left[Y|A = 0, X\right]|A = 1\right]$ instead. The thought exercise would be similar to the one above, except that the groups would have their roles reversed.

2.2 Estimating the Quantities in the Decomposition

Having explained the logic of the decomposition, we now turn to estimation strategies for the quantities involved. We can use the sample means within each group to estimate the group means, weighted when appropriate (e.g., by survey weights). Various approaches are, however, available for estimating the counterfactual.

2.2.1 Regression-Imputation Estimator

We begin with what is perhaps the most intuitive approach, as it attempts to estimate the counterfactual quantity directly. We call it the regression-imputation (RI) estimator:

$$
\hat{\theta}^{RI} = \widehat{\mathbb{E}}\left[\widehat{\mathbb{E}}\left[Y|A = 1, X\right]|A = 0\right]. \tag{2}
$$

It implies two alternative strategies:

[1] Despite the use of the term, this quantity does not have a causal interpretation. In other words, it refers to a "descriptive" counterfactual rather than a causal counterfactual.

(1) Fit an outcome model among group members, yielding $\widehat{\mathbb{E}}\,[Y|A=1,X]$. Then use this to obtain predicted values based on the covariate distribution for non-members. Finally, average these predictions to evaluate the outer expectation.
(2) Fit a pooled outcome model including the group dummy and the covariates among all units: $\widehat{\mathbb{E}}\,[Y|A,X]$. Then obtain imputed values of the inner expectation ($\widehat{\mathbb{E}}\,[Y|A=1,X]$) by setting A to 1. Finally, evaluate the outer expectation by averaging these imputed values only among non-members.

To be consistent, the RI estimator depends on the correct specification of the outcome model. In addition to the linear form often assumed by analysts, logit or probit models are available for binary outcomes, Poisson or negative binomial models for count outcomes, or even complex machine learning models when a high-dimensional covariate vector is present.

2.2.2 Weighting Estimator

The second strategy proposes an indirect way to estimate the counterfactual quantity by manipulating its mathematical expression (DiNardo, Fortin, and Lemieux 1996; Barsky et al. 2002):

$$\begin{aligned}\mathbb{E}[\mathbb{E}[Y|A=1,X]|A=0] &= \int \mu_1(x)dF_{X|A=0}(x) \\ &= \int \mu_1(x)\omega(X)dF_{X|A=1}(x),\end{aligned} \tag{3}$$

where $\omega(X)$ is a weighting factor, defined as

$$\omega(X)=\frac{\Pr(A=1)\Pr(A=0|X)}{\Pr(A=0)\Pr(A=1|X)}.$$

The equivalence of the two lines in Equation (3) follows from Bayes's rule:

$$f_{X|A=0}(x)=\underbrace{\frac{\Pr(A=1)\Pr(A=0|X)}{\Pr(A=0)\Pr(A=1|X)}}_{\omega(X)}\underbrace{\frac{\Pr(A=1|X)f_X(x)}{\Pr(A=1)}}_{f_{X|A=1}(x)},$$

where $f_X(.)$ is the marginal density function and $f_{X|A=a}(.)$ is the conditional density function given group membership.

Equation (3) shows that the counterfactual can be seen as a weighted average of the outcome among group members, where the weights eliminate the differences in the covariate distributions between the two groups — thus accomplishing the same objective as the RI estimator but through an alternative route. Since we take the average among group members only, we can replace $\Pr(A=1)$ with a group indicator in the numerator of $\omega(X)$, leading to the following weighting estimator:

$$\hat{\theta}^W=\frac{\sum_i W_iY_i}{\sum_i W_i}, \text{ where } W_i=\frac{\mathbb{I}[A_i=1]}{\widehat{\Pr}[A_i=0]}\frac{\widehat{\Pr}[A_i=0|X_i]}{\widehat{\Pr}[A_i=1|X_i]}. \tag{4}$$

$\mathbb{I}(.)$ is an indicator function, and $\widehat{\Pr}[A_i=a]$ is the sample proportion of individuals in group a. The estimator requires fitting a propensity score model for A given X, used to obtain the predicted probabilities $\widehat{\Pr}[A_i=a|X_i]$. It thus avoids functional form assumptions about the outcome model at the cost of relying on the correct specification of the propensity score function in order to be consistent.

2.2.3 Doubly Robust Estimator

We now combine the RI and weighting estimators to introduce a doubly robust (DR) estimator (Tsiatis 2006):

$$
\begin{aligned}
\hat{\theta}^{DR} = \frac{1}{n}\sum_{i=1}^{n} \Big\{ & W_i \left(Y_i - \widehat{\mathbb{E}}\left[Y_i|A_i = 1, X_i\right]\right) \\
& + \frac{\mathbb{I}[A_i = 0]}{\widehat{\Pr}[A_i = 0]} \left(\widehat{\mathbb{E}}\left[Y_i|A = 1, X\right] - \widehat{\mathbb{E}}\left[\widehat{\mathbb{E}}\left[Y_i|A_i = 1, X_i\right]|A = 0\right]\right) \\
& + \widehat{\mathbb{E}}\left[\widehat{\mathbb{E}}\left[Y_i|A_i = 1, X_i\right]|A = 0\right] \Big\},
\end{aligned}
\tag{5}
$$

where n is the sample size, and all other terms are defined as before. The DR estimator involves fitting both outcome and propensity score models. It is doubly robust because it will consistently estimate the counterfactual quantity if either of these models is correctly specified, but not necessarily both (Robins, Rotnitzky, and Zhao 1994). Appendix A derives the proof of this double robustness.

Despite the widespread use of doubly robust estimators in the causal inference literature (e.g., Tchetgen and Shpitser 2012; Chernozhukov et al. 2018), little has been said about their application to the KOB decomposition. An exception is Kline's (2011) demonstration that the KOB estimator based on a linear model for the outcome is consistent for the counterfactual if either the outcome model is correctly specified or if $\Pr(A = 0|X)/\Pr(A = 1|X)$ is a linear function of X. Due to its reliance on linearity, this double robustness is much more restrictive than that of $\hat{\theta}^{DR}$, which requires no assumptions about the functional forms of the outcome or propensity score models.

3 Disparities at Quantiles of the Overall Distribution

We now adapt the logic of the classical KOB decomposition for mean differences to study overall quantile disparities. This concept relies on the recentered influence function (RIF) and unconditional quantile regressions (UQRs), which have drawn considerable interest from economists and sociologists over the last decade (e.g., Firpo, Fortin, and Lemieux 2009; Rios-Avila and Maroto 2022; Killewald and Bearak 2014; Borgen, Haupt, and Wiborg 2023). This section starts with an overview of the RIF and UQRs. It then shows how we can use them to study overall quantile gaps.

3.1 RIF and UQRs

Much like traditional regressions, UQRs involve an outcome and a set of explanatory variables. The outcome here, however, is a dependent variable transformed by the RIF. The RIF for an arbitrary distributional statistic $\nu(F_Y)$ (e.g., mean, variance, quantiles, etc.) is defined as $\mathrm{RIF}(y; \nu(F_Y)) = \mathrm{IF}(y; \nu(F_Y)) + \nu(F_Y)$. $\mathrm{IF}(y; \nu(F_Y))$ is the influence function for $\nu(F_Y)$, capturing how the addition of an individual observation affects that distributional statistic (Rios-Avila 2020). Aggregating $\nu(F_Y)$ back to $\mathrm{IF}(y; \nu(F_Y))$ ensures that the mean of $\mathrm{RIF}(Y; \nu(F_Y))$ is the distributional statistic itself: $\mathbb{E}[\mathrm{RIF}(Y; \nu(F_Y))] = \nu(F_Y)$.

The RIF for the p-quantile of a distribution is given by

$$
\mathrm{RIF}(Y; Q_p) = Q_p + \frac{p - \mathbb{I}(Y \leq Q_p)}{f_Y(Q_p)}, \tag{6}
$$

where Q_p is the p-quantile of the distribution, p is the corresponding probability for that quantile, $f_Y(Q_p)$ is the probability density function of the dependent variable Y evaluated at Q_p, and $\mathbb{I}(Y \leq Q_p)$ is an indicator function that returns 1 if a value of Y is smaller than the quantile and 0 otherwise.

UQRs are designed to assess the impact of explanatory variables on the overall distribution. Without further assumptions, they only enable an interpretation at the population level, asking how a change in the distribution of the explanatory variables affects the quantile of interest (Borgen, Haupt, and Wiborg 2023). Consider the simple case of a linear regression of $\mathrm{RIF}(Y; Q_p)$ on a single explanatory variable X, which Firpo, Fortin, and Lemieux (2009) call RIF-OLS:[2]

[2] More flexible strategies, which do not require a linearity assumption, also exist for estimating the effect of variables on quantiles (Firpo, Fortin, and Lemieux 2009).

$$\mathbb{E}[\mathrm{RIF}(Y; Q_p)|X] = \beta_0 + \beta_1 X. \tag{7}$$

Taking the mean of (7) yields $\mathbb{E}[\mathbb{E}[\mathrm{RIF}(Y; Q_p)|X]] = \beta_0 + \beta_1 \mathbb{E}[X]$. By the law of iterated expectations, the left-hand side is equivalent to $\mathbb{E}[\mathrm{RIF}(Y; Q_p)]$. But recall that the mean of the RIF for a distributional statistic is the statistic itself, so $\mathbb{E}[\mathrm{RIF}(Y; Q_p)] = Q_p$. The coefficient β_1 thus describes the impact of a one-unit change in the mean of X on the quantile of interest. UQRs allow us to estimate, for example, how the 0.10-quantile of the wage distribution among women changes if the average number of children increases by one (Rios-Avila and Maroto 2022).

If, instead of X in Equation (7), we had a binary variable A indicating group membership, the coefficient β_1 would represent the difference $\mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = 1] - \mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = 0]$. That is, it would yield the quantile disparity between a scenario where no one is a member of group A and one where everyone is (Borgen, Haupt, and Wiborg 2023). Rios-Avila and Maroto (2022) suggest making the coefficient more interpretable by considering changes in the incidence rate instead, i.e., the percentage point increase in the share of members of A. This involves multiplying the estimated coefficient by the chosen change in incidence rate (e.g., an analyst interested in assessing the effect of a 10% increase in the proportion of group members would multiply the coefficient by 0.10).

For an individual-level interpretation (e.g., the expected change in wages with a one-unit increase in the number of children among women at the 0.10-quantile of the wage distribution), we must assume that observations retain their positions in the distribution regardless of their values on the explanatory variables (rank-invariance), or, more flexibly, that no systematic changes in ranks occurs with changes in the explanatory variables (rank-similarity) (Wenz 2019).

3.2 Using the Rank Transformation to Define Overall Quantile Gaps

The idea of manipulating group proportions, on which the interpretation of coefficients for binary variables in UQRs usually rests, can be difficult to conceive — especially if groups are based on inherent traits like race or sex. Next, we introduce a more substantively useful interpretation involving the rank transformation of the dependent variable.

We begin by rewriting Equation (6):

$$\mathrm{RIF}(Y; Q_p) = c_{1,p}\mathbb{I}(Y > Q_p) + c_{2,p}, \tag{8}$$

where $c_{1,p} = 1/f_Y(Q_p)$ and $c_{2,p} = Q_p - c_{1,p}(1 - p)$. The conditional expectation of $\mathrm{RIF}(Y; Q_p)$ given group membership is thus:

$$\mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = a] = c_{1,p}\Pr[Y > Q_p|A = a] + c_{2,p}.$$

This is the conditional probability of placing above the quantile of interest, rescaled by a factor $c_{1,p}$ that reflects the relative importance of the quantile to the distribution and recentered by a constant $c_{2,p}$ (Firpo, Fortin, and Lemieux 2018). The group difference is:

$$\begin{aligned}\Delta(p) &= \mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = 1] - \mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = 0] \\ &= c_{1,p}\left(\Pr[Y > Q_p|A = 1] - \Pr[Y > Q_p|A = 0]\right).\end{aligned} \tag{9}$$

When the distribution of Y is uniform, all of its quantiles have the same relative importance, so $c_{1,p}$ becomes a constant. If it is uniform over $(0, 1)$, $c_{1,p}$ is unity for all p. We can transform Y to have a uniform distribution over $(0, 1)$ by plugging its values into the empirical cumulative distribution function (ECDF) to recover their ranks. The rank of some value t on the Y distribution is given by the proportion of observations below it:

$$\hat{F}_n(t) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{I}(y_i < t). \tag{10}$$

Let Y^{rank} be the transformed variable storing the ranks of the Y values. If we substitute Y^{rank} for Y in Equation (9), $\Delta(p)$ yields the group disparity in the probability of reaching a benchmark (defined by the quantile) of the overall distribution. We can estimate it by regressing $\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}}) = \mathbb{I}(Y^{\text{rank}} > Q_p^{\text{rank}}) + c_{2,p}^{\text{rank}}$ on the group indicator. The constant $c_{2,p}^{\text{rank}}$ equals $Q_p^{\text{rank}} - (1-p)$. When Y is continuous and the sample is reasonably large, Q_p^{rank} approximates p, so $c_{2,p}^{\text{rank}} \approx 2p - 1$.

In short, we can define the overall quantile gap as the difference in probabilities of reaching benchmarks of the overall distribution if we use $\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}})$ instead of $\text{RIF}(Y; Q_p)$ as the outcome in the UQR.

4 Using Quantile Disparities to Augment the KOB Decomposition

The classical KOB method decomposes the mean gap horizontally into explained and unexplained components. This section introduces a vertical dimension, which decomposes the total mean disparity and its explained and unexplained components across quantiles of the distribution.

We begin by partitioning the gap at the p-quantile following the same logic outlined in Equation (1) for the case of the mean:

$$\begin{aligned}
\Delta(p) &= \mathbb{E}[\text{RIF}(Y; Q_p)|A=1] - \mathbb{E}[\text{RIF}(Y; Q_p)|A=0] \\
&= \mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=1, X]|A=1] - \mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=0, X]|A=0] \\
&= (\mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=1, X]|A=1] - \mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=1, X]|A=0]) \\
&+ (\mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=1, X]|A=0] - \mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A=0, X]|A=0]) \\
&= \underbrace{\int m_1(x, Q_p)\left[dF_{X|A=1}(x) - dF_{X|A=0}(x)\right]}_{\Delta_X(p) \ \ \text{(Explained)}} + \underbrace{\int \left[m_1(x, Q_p) - m_0(x, Q_p)\right] dF_{X|A=0}(x)}_{\Delta_U(p) \ \ \text{(Unexplained)}},
\end{aligned} \tag{11}$$

where $m_a(x, Q_p)$ is the conditional mean of $\text{RIF}(Y; Q_p)$ given $A = a$ and $X = x$, and $\Delta_X(p)$ and $\Delta_U(p)$ are the explained and unexplained components.

If we use Y^{rank} instead of Y as the dependent variable, $\mathbb{E}[\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}})|A = a]$ represents the sum of two terms: (1) the conditional probability of reaching the corresponding quantile of the overall distribution and (2) the constant term $c_{2,p}^{\text{rank}}$, which approximates $2p - 1$, as shown at the end of the previous section. We can then subtract $c_{2,p}^{\text{rank}}$ from $\mathbb{E}[\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}})|A = a]$ to recover the conditional probability.

Moreover, the counterfactual $\mathbb{E}[\mathbb{E}[\text{RIF}(Y; Q_p)|A = 1, X]|A = 0]$ represents the mean of $\text{RIF}(Y; Q_p)$ among group members under the fiction that their covariate distribution is the same as that of non-members. With Y^{rank} substituting for Y as the dependent variable, it represents the sum of the counterfactual probability and $c_{2,p}^{\text{rank}}$. As before, we can subtract the constant term to compute the counterfactual probability.

The mean and quantile decompositions are inextricably linked. In particular, the total mean disparity and its explained/unexplained components are averages of the corresponding quantities from the quantile decompositions across the overall distribution. This holds for any outcome with a finite lower bound, such as income, wealth, and their monotone transformations (e.g., logarithm, ECDF, etc.). The result derives from the following lemma:

$$\int_0^1 \text{RIF}(y; Q_p) dp = y. \tag{12}$$

In words, the average RIF for a given y value across all quantiles of the distribution is y itself. Lemma (12) implies a series of results that allow us to state the relationships in Table 1. We prove the lemma and the other results in Appendix B.

Table 1: Horizontal and Vertical Decompositions of the Mean Disparity

	Component				
Dimension	Total		Explained		Unexplained
Horizontal	$\underbrace{\Delta}$	=	$\underbrace{\Delta_X}$	+	$\underbrace{\Delta_U}$
Vertical	$\int_0^1 \Delta(p)dp$	=	$\int_0^1 \Delta_X(p)dp$	+	$\int_0^1 \Delta_U(p)dp$

Table 1 describes the augmented KOB (aKOB) decomposition. It partitions the mean disparity horizontally and vertically. The horizontal dimension (the first row of the table) decomposes it into explained and unexplained components, as intended by the classical method. The vertical dimension (the second row) decomposes it across quantiles of the overall distribution: the total mean disparity is the average of the overall quantile gaps across the overall distribution; similarly, the explained and unexplained components from the mean decomposition are the averages of the corresponding quantities from the quantile decompositions. Since the importance of explanatory variables in accounting for group disparities may vary at different points of the distribution, the aKOB decomposition reveals heterogeneity that the classical method overlooks. The empirical example below illustrates one such situation. But before presenting it, we discuss implementation strategies for our methods.

5 Implementation

The RI, weighting, and DR estimators presented above for partitioning the mean gap are directly applicable to quantile decompositions by replacing Y with $\text{RIF}(Y; Q_p)$ (or $\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}})$) in their formulas.

Several algorithms exist for estimating the quantiles needed to compute $\text{RIF}(Y; Q_p)$ (Hyndman and Fan 1996), and most statistical software will implement at least one of them. The $\text{RIF}(Y; Q_p)$ quantity also requires that the marginal density function of Y be estimated. Firpo, Fortin, and Lemieux (2009) recommend the kernel density estimator for this. Statistical software can also estimate the ECDF used to construct Y^{rank} — analysts should use the weighted ECDF if weights are present. As discussed above, Y^{rank} follows a uniform distribution over $(0, 1)$, so the density function is known for $\text{RIF}(Y^{\text{rank}}; Q_p^{\text{rank}})$, obviating the need to estimate it.

OLS and logistic/probit models, as well as more flexible alternatives (e.g., local regression, generalized additive models, etc.), are available for fitting the outcome and propensity score models in the RI, weighting, and DR estimators. The DR estimator, in particular, is amenable to machine learning outcome and propensity score models, which lead to theoretically justified standard errors and may improve robustness and efficiency (Chernozhukov et al. 2018). The empirical example relies on linear models to simplify the illustration.[3]

To draw inferences from the decomposition analyses, we use bootstrapping to construct standard errors around the counterfactual estimates (Fortin, Lemieux, and Firpo 2011).

6 Empirical Illustration

The decline in unionization in the 1980s among male workers coincided with the rise in wage inequality, leading to considerable research on the connection between the two (Freeman 1993; Western and Rosenfeld 2011). Several of these studies have explored potential heterogeneous effects across the wage distribution (e.g., Card 1996; DiNardo, Fortin, and Lemieux 1996; Chamberlain 1994). The findings here point to an increasing union premium from the low end to the middle of the distribution, where it reaches its peak, followed by a sharp decline that culminates in the negative effect of unions on wages in the upper tail — partly because the wage compression effects of unions (Freeman and Medoff 1984) limit the earning potential of high-wage union

[3]In analyzing the same data, Firpo, Fortin, and Lemieux (2009) found that a linear model produces results similar to those of more flexible alternatives.

workers relative to their nonunion counterparts (Cai and Liu 2008). As a result, while the fall of organized labor reduced inequality at the bottom of the distribution (although this was offset by inequality-engendering phenomena like declining minimum wages; see Lemieux (2008) for a review), it widened the gap between the middle brackets and the top quantiles (DiNardo, Fortin, and Lemieux 1996; Firpo, Fortin, and Lemieux 2018).

In addition to confirming several of these findings, our empirical illustration uses the aKOB decomposition to compare the relative importance of different variables in accounting for the union gap across the wage distribution. We revisit the data set examined by Firpo, Fortin, and Lemieux (2009) to illustrate UQRs, which comprises a sample of 266,956 U.S. males from the 1983-1985 Outgoing Rotation Group supplement of the Current Population Survey.[4] Below we describe the variables included in our analyses.

6.1 Variables

The group membership variable identifies whether the respondent was a member of a union or similar employee association for his current job. In keeping with the discussion above, we use wage rank (given by the ECDF) as the dependent variable to enable the interpretation of overall quantile gaps as disparities in probabilities of reaching benchmarks of the overall distribution. Additionally, we follow Firpo, Fortin, and Lemieux (2009) in considering these variables as potential contributors to the union wage gap:

- Race indicates if the respondent was white or nonwhite.
- Education has six categories: elementary (less than nine years of education), high school dropout (between nine and eleven years), high school (twelve years), some college (between thirteen and fifteen years), college (sixteen years), and post-graduate (more than sixteen years).
- Work experience has nine categories: less than five years, between five and nine, and six other five-year interval categories up to the last category of forty years or more.
- Marital status indicates if the respondent was married or not married.

Table 2 compares the percentages of union and nonunion workers in each category of the variables above. Particularly notable are the disparities in educational attainment, experience, and marital status. While a larger fraction of nonunion workers (45%) has some college education or higher (versus 34% of union workers), they are largely inexperienced, with 21% of them reporting less than five years of work (compared to only 7% of union workers). Additionally, a larger proportion of nonunion workers (40%) is unmarried (versus 27% of union workers).

[4] The data set is available at https://sites.google.com/view/nicole-m-fortin/data-and-programs?authuser=0.

Table 2: Percentages in Categories of Explanatory Variables by Union Membership

Variable	Union	Nonunion
Race		
White	85.41	88.65
Nonwhite	14.59	11.35
Education		
Elementary	7.16	6.53
High school dropout	12.34	12.84
High school	46.94	35.33
Some college	18.14	19.92
College	7.54	15.00
Post-graduate	7.87	10.37
Years of Experience		
Less than 5	6.89	21.33
5-9	13.03	18.49
10-14	16.02	15.64
15-19	14.72	11.30
20-24	11.78	8.47
25-29	10.35	7.23
30-34	9.49	6.27
35-39	8.22	5.47
40 or more	9.50	5.79
Marital Status		
Married	73.46	60.50
Not married	26.54	39.50

Results weighted by survey weights.

6.2 Decomposition Analyses

Before presenting the results for the decomposition analyses, a note on the choice of the counterfactual is warranted. Barsky et al. (2002) recommend selecting the group whose covariate distribution spans a broader range of values. The idea is to avoid extrapolations out of the observed range in estimating the counterfactual. For example, to study the extent to which income differences account for the racial wealth gap, the authors construct a counterfactual from the perspective of whites, proposing a scenario in which their earnings distribution shifted downward to match that of blacks. Since the black earnings distribution is concentrated at the low end, the alternative counterfactual for the scenario in which the black distribution shifted upward to the levels of whites would require extrapolating out of the observed black earnings range.

Given that the proportion of nonunion workers far surpasses that of union workers in our data set (74% versus 26%), we find evaluating a counterfactual from the standpoint of nonunion workers the natural choice — indeed, the RI method yields similar results using the pooled approach and selecting nonunion workers as the reference group (not shown).

6.2.1 Horizontal Decomposition for the Mean Difference

We first partition the mean wage rank gap into a component explained by differences in race, education, experience, and marital status and an unexplained component. This is the horizontal decomposition intended by the classical method. Table 3 shows the results.

Table 3: Mean Decomposition

		RI (OLS)	Weighting (Logit)	DR
Mean Rank (Union):	0.602			
Mean Rank (Nonunion):	0.464			
Union-Nonunion Difference:	0.138			
Counterfactual Mean Rank		0.495 (0.00070)	0.494 (0.00078)	0.493 (0.00076)
Explained Component		0.031	0.030	0.029
Percentage Explained		22.34	21.66	21.01

Bootstrapped standard errors for counterfactual estimates in parentheses (250 iterations).

Union and nonunion workers have mean ranks of 0.602 and 0.464, a difference of 0.138. The different methods estimate that the mean rank among nonunion workers would increase to a little over 0.490 if their covariate distribution became the same as that of union workers. The RI method yields the most precise counterfactual estimate, whereas the standard error around the DR estimate is between its RI and weighting counterparts. The large sample size, however, allows us to draw inferences with considerable precision under all methods.

We partition the total disparity as follows: $\left(\bar{Y}^{\text{rank}}_{\text{union}} - \text{counterfactual}\right) + \left(\text{counterfactual} - \bar{Y}^{\text{rank}}_{\text{nonunion}}\right)$. The first and second terms represent the unexplained and explained components, respectively. Using the DR estimate, we find that differences in the explanatory variables account for $(0.493 - 0.464)/0.138 = 21\%$ of the gap. We can interpret the unexplained component as the union premium — i.e., gains from affiliation with unions not attributable to observed characteristics.

6.2.2 Quantile Decompositions

Table 4 presents the decompositions for the 0.10, 0.25, 0.50, 0.75, and 0.90 quantiles. The interpretation here follows a similar logic to the one above for the mean decomposition, except that, instead of the mean ranks, we now consider the probabilities of reaching quantiles of the distribution.

Differences in the explanatory variables explain a considerable portion of the gap (between 30% and 35% depending on the estimator) at the 0.10 quantile, implying a relatively small union premium there. If nonunion workers had the same distribution of covariates as union workers, their probability of reaching the 0.10 quantile of the wage distribution would increase from the actual value of 0.871 to approximately 0.91.

The union premium increases as we move up the wage distribution. It finds its zenith at the 0.50 quantile, where the explanatory variables account for only 15-17% of the gap.

Further along, we find that union workers have a lower probability of reaching the 0.90 quantile of the wage distribution than nonunion workers (0.070 versus 0.111). The percentage of the gap explained by the covariates is negative for this quantile because the models indicate a slight increase in the probability for nonunion workers if their distribution of observed characteristics matched that of union workers.

The estimators differ the most in the upper tail of the distribution: the RI estimator indicates larger counterfactual estimates for the 0.75 and 0.90 quantiles than the other two estimators. But they all lead to the same substantive conclusions. As in the horizontal decomposition for the mean gap, the DR estimates lie between their RI and weighting counterparts in terms of precision. The large sample size, however, results in minimal inferential uncertainty across all estimators.

These findings corroborate previous studies showing that unions benefit middle-wage earners the most.

Table 4: Quantile Decompositions

		RI (OLS)	Weighting (Logit)	DR
0.10-quantile				
Prob. (Union):	0.981			
Prob. (Nonunion):	0.871			
Union-Nonunion Difference:	0.109			
Counterfactual Prob.		0.905	0.910	0.909
		(0.00062)	(0.00065)	(0.00064)
Explained Component		0.034	0.039	0.038
Percentage Explained		30.74	35.40	34.74
0.25-quantile				
Prob. (Union):	0.918			
Prob. (Nonunion):	0.690			
Union-Nonunion Difference:	0.227			
Counterfactual Prob.		0.742	0.747	0.746
		(0.00096)	(0.00106)	(0.00104)
Explained Component		0.051	0.057	0.056
Percentage Explained		22.63	25.01	24.51
0.50-quantile				
Prob. (Union):	0.687			
Prob. (Nonunion):	0.433			
Union-Nonunion Difference:	0.254			
Counterfactual Prob.		0.477	0.473	0.472
		(0.00112)	(0.00126)	(0.00123)
Explained Component		0.044	0.040	0.039
Percentage Explained		17.15	15.84	15.35
0.75-quantile				
Prob. (Union):	0.301			
Prob. (Nonunion):	0.232			
Union-Nonunion Difference:	0.069			
Counterfactual Prob.		0.251	0.245	0.244
		(0.00097)	(0.00107)	(0.00105)
Explained Component		0.019	0.013	0.012
Percentage Explained		27.20	18.47	17.17
0.90-quantile				
Prob. (Union):	0.070			
Prob. (Nonunion):	0.111			
Union-Nonunion Difference:	-0.040			
Counterfactual Prob.		0.116	0.112	0.112
		(0.00060)	(0.00067)	(0.00066)
Explained Component		0.006	0.001	0.001
Percentage Explained		-13.89	-3.64	-2.49

Bootstrapped standard errors for counterfactual estimates in parentheses (250 iterations).

6.2.3 aKOB Decomposition

The three estimators lead to similar findings in our remaining analyses, so we hereafter only show the results based on the DR estimator. The classical KOB method partitions the mean gap horizontally into explained and unexplained components. The aKOB approach also decomposes it across quantiles to reveal heterogeneity. Figure 1 displays the horizontal (the horizontal lines) and the vertical (the colored lines) dimensions of the aKOB decomposition. The horizontal lines represent the decomposition in Table 3, and the colored lines represent the decompositions for all quantiles between 0.05 and 0.95 in 0.01 increments. The long-dashed lines reflect the extent of the union premium.

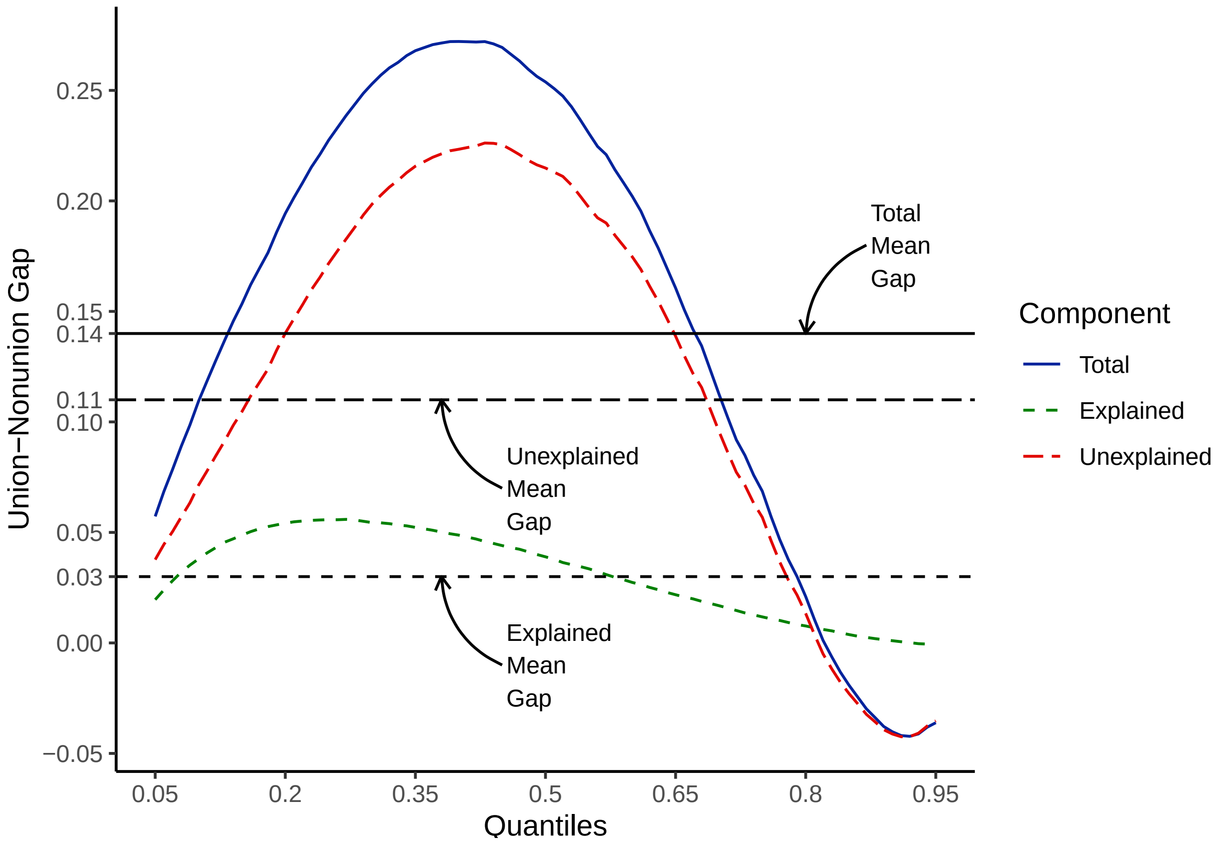

Figure 1: aKOB Decomposition (Doubly Robust Estimates)

The explained line peaks at the low end, resulting in a larger portion of the total disparity explained by differences in covariate distributions. However, the total gap grows toward the middle of the distribution as the explained component declines, generating a considerable premium for the middle quantiles. At the high end (approximately from the 0.80 quantile forward), the total gap becomes negative — i.e., nonunion workers have a higher probability of reaching the top of the distribution than union workers — while the explained component remains approximately zero.

Being averages of the corresponding quantities from the quantile decompositions (Table 1), the elements of the horizontal decomposition for the mean difference must lie within the range of those quantities. It is no surprise, then, that the horizontal lines intersect the colored ones. They summarize patterns encountered across the distribution and, in doing so, may lose relevant information. Figure 1 is a visual illustration of the key advantage of the aKOB decomposition over the classical method: it unveils heterogeneity overlooked by a decomposition focused only on the mean gap.

So far, we have considered the explanatory variables in conjunction, which may conceal their unique roles in accounting for the union wage gap. The following section overcomes this by examining their net contributions.

Table 5: Net Contributions of the Explanatory Variables

		Quantile		
	Mean	0.10	0.50	0.90
Union:	0.602	0.981	0.687	0.070
Nonunion:	0.464	0.871	0.433	0.111
Union-Nonunion Difference:	0.138	0.109	0.254	-0.040
All Covariates				
Gap Explained	0.029	0.038	0.039	0.001
Percentage Explained	21.01	34.74	15.35	-2.49
Net Contribution of Race				
Gap Explained	-0.003	-0.002	-0.004	-0.001
Percentage Explained	-2.16	-1.78	-1.69	3.13
Net Contribution of Education				
Gap Explained	-0.022	-0.004	-0.030	-0.027
Percentage Explained	-15.80	-3.21	-11.77	66.36
Net Contribution of Experience				
Gap Explained	0.032	0.022	0.046	0.017
Percentage Explained	23.25	20.34	18.14	-42.27
Net Contribution of Marital Status				
Gap Explained	0.002	0.002	0.003	0.000
Percentage Explained	1.47	1.92	1.08	-1.03

Doubly Robust Estimates.

6.2.4 Net Contribution of Each Explanatory Variable

Let us assign generic labels X_1, X_2, X_3, X_4 to our explanatory variables. To study the net contribution of one of them (say, X_4), we adapt the procedure outlined in Fortin, Lemieux, and Firpo (2011): we compare the results from a model including all explanatory variables (full model) to those from a model omitting X_4 (restricted model).[5] We repeat this for each explanatory variable.

We consider two measures of net contribution: (1) the explained component difference between the full and restricted models and (2) the difference in the percentage of the total gap explained by the covariates under the two models. The latter facilitates comparisons between quantiles since the size of the gap varies considerably across the distribution. If the differences are positive, the full model accounts for a larger fraction of the gap than the restricted one, so the omitted variable has an *explanatory* effect. If the differences are negative, the variable has a *suppressing* effect since removing it enhances the explanatory power.

Table 5 shows the results for the mean decomposition and the 0.10, 0.50, and 0.90 quantiles decompositions — points of the distribution where considerable heterogeneity is evident (see Table 4 and Figure 1). The first rows reproduce the information from Tables 3 and 4 for the reader's convenience.

We begin by noting the minimal net contributions of race and marital status for both the mean and quantile decompositions. As shown in Table 2, union and nonunion workers have similar race distributions, with whites accounting for the vast majority of the respondents in the two groups. A larger fraction of union workers (73%) is married, however (versus 61% of nonunion workers). Hence, we conjecture that the low net contribution of marital status is due to the comparatively limited effect of this variable on wages.

Table 6 helps us assess this; it displays the coefficients for an OLS regression of wage rank on the union

[5]This procedure circumvents path dependence (originating from sequential decompositions) and the need for a linear outcome model in performing the detailed decomposition (Jann 2008). Although it does not produce net explained components that sum up perfectly to the total explained component based on all covariates, it is still helpful for comparing the relative importance of explanatory variables.

indicator and the other explanatory variables (first column on the left) and logistic regressions with $\mathbb{I}(Y > Q_p)$ (Equation (8)) as the dependent variable (the three remaining columns). The logistic regression models thus yield the log-odds of placing above the corresponding quantile. The union coefficients corroborate the results from the decomposition analyses: union workers have higher average ranks and are favored throughout most of the distribution, except at the high end. Although married respondents hold an advantage over their unmarried counterparts, the marital status coefficient is relatively small compared to the coefficients for the education and experience categories.

The union/nonunion discrepancy in education and experience (Table 2) and the large effects of these variables on wage explain their large net contributions in Table 5. Moreover, their contributions have opposite signs, indicating that while one variable has an explanatory effect, the other has a suppressing effect. For the mean and the two lowest quantiles, experience has a positive net contribution, whereas education is a suppressor. The pattern reverses at the 0.90 quantile.

Recall from Table 2 that union and union workers diverge the most in the highest education and the lowest experience categories. A larger percentage (45%) of nonunion workers have some college education or higher (versus 34% of union workers), with 21% of them reporting less than five years of work experience (compared to 7% of union workers). It will be helpful, therefore, to concentrate on the coefficients for these categories in Table 6 to understand the net contributions of education and experience.

For the mean and the 0.10 and 0.50 quantiles, the positive net contribution of experience is attributable to the large negative effect of having less than five years of experience on wages (relative to having 20-24 years of experience) and the disproportionate presence of nonunion workers in this category. By contrast, wages are positively associated with having some college experience or a higher education degree (rather than only having a high school diploma), so the predominance of nonunion workers in these education categories undermines our ability to explain the union wage gap.

Note that the net contribution of experience far surpasses that of education (in magnitude) at the 0.10 quantile — the gap explained is 0.022 for experience, compared to -0.004 for education. This relates to the increasing importance of education as we move up the wage distribution. While experience overshadows education in determining the prospects of reaching the 0.10 quantile, higher education gains prominence in the middle of the distribution until it becomes a crucial determinant of earning high wages. Y-standardized coefficients (not shown) confirm this trend. Between the 0.10 and 0.90 quantiles, the coefficient for "some college" rises from 0.108 to 0.337, that for "college" grows from 0.622 to 0.833, and that for "post-graduate" doubles (from 0.493 to 1). By contrast, the effect of having less than five years of experience remains relatively stable, varying between -0.836 and -1.082. Together, education and experience have offsetting effects for explaining the disparity at the 0.90 quantile. However, an analysis of their net contributions reveals that higher education makes nonunion workers disproportionately likely to attain high wages.

6.3 Summary of Findings

The union-wage relationship varies considerably across the wage distribution. An account of this heterogeneity must take heed that workers have different stakes in the standardizing effects of collective bargaining on wages (Freeman 1980). Cai and Liu (2008) argue that due to their low skills and relatively high substitutability, lower- and middle-wage workers have limited individual bargaining power, benefiting the most from organized labor. By contrast, high skills and low substitutability grant high-wage earners substantial individual bargaining power, so "association with unions makes little difference for them in terms of bargained wage outcomes" (Cai and Liu 2008, 497). Our findings illuminate this claim by showing that higher education is the main asset enabling elevated individual bargaining power among high-earning workers. It gives nonunion workers an edge at the top of the wage distribution, even though the lack of work experience harms their prospects of attaining middle wages.

7 Concluding Remarks

The classical KOB method has proved to be a powerful tool for investigating group inequalities. However, because it focuses on a horizontal decomposition of the mean gap into explained and unexplained components,

Table 6: Regressions of Wage on Covariates (N = 266,956)

		Log-Odds of Placing Above Quantile		
	OLS	0.10	0.50	0.90
Union Membership (Ref=Nonunion)				
Union	0.107***	1.665***	1.111***	−0.512***
	(0.001)	(0.033)	(0.012)	(0.020)
Race (Ref=White)				
Nonwhite	−0.079***	−0.606***	−0.634***	−0.534***
	(0.002)	(0.024)	(0.017)	(0.031)
Education (Ref=High School)				
Elementary	−0.189***	−1.484***	−1.473***	−1.655***
	(0.002)	(0.033)	(0.022)	(0.069)
High school dropout	−0.096***	−1.055***	−0.776***	−0.796***
	(0.001)	(0.021)	(0.017)	(0.045)
Some college	0.074***	0.261***	0.635***	0.793***
	(0.001)	(0.023)	(0.014)	(0.024)
College	0.213***	1.500***	1.621***	1.960***
	(0.002)	(0.041)	(0.018)	(0.022)
Post-graduate	0.242***	1.190***	1.847***	2.352***
	(0.002)	(0.049)	(0.021)	(0.023)
Years of Experience (Ref=20-24)				
Less than 5	−0.278***	−2.018***	−2.202***	−2.547***
	(0.002)	(0.045)	(0.023)	(0.050)
5-9	−0.140***	−0.806***	−1.091***	−1.380***
	(0.002)	(0.046)	(0.021)	(0.032)
10-14	−0.076***	−0.426***	−0.589***	−0.809***
	(0.002)	(0.049)	(0.020)	(0.029)
15-19	−0.029***	−0.211***	−0.229***	−0.313***
	(0.002)	(0.053)	(0.022)	(0.028)
25-29	0.013***	−0.125**	0.108***	0.177***
	(0.002)	(0.058)	(0.024)	(0.030)
30-34	0.016***	−0.088	0.139***	0.191***
	(0.002)	(0.060)	(0.025)	(0.032)
35-39	0.018***	−0.053	0.130***	0.253***
	(0.003)	(0.063)	(0.026)	(0.034)
40 or more	0.004	−0.216***	0.063**	0.206***
	(0.003)	(0.055)	(0.025)	(0.040)
Marital Status (Ref=Married)				
Not married	−0.072***	−0.975***	−0.522***	−0.320***
		(0.021)	(0.012)	(0.021)
Constant	0.552***	3.586***	0.273***	−2.293***
	(0.002)	(0.044)	(0.018)	(0.026)

*p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parentheses.
Models weighted by survey weights.

it may mask varied stratification patterns unfolding at different points of the distribution. In this paper, we decomposed the gap vertically across quantiles of the overall distribution, revealing heterogeneity that traditional approaches overlook. Incorporating this vertical dimension allowed us to expand the classical method while preserving its original purpose. We called this augmented version aKOB decomposition.

The vertical dimension decomposes "overall quantile gaps," which we defined as group disparities in probabilities of reaching benchmarks of the overall distribution. To develop this concept, we used the rank transformation of the dependent variable in UQRs involving the group indicator as the predictor. By doing so, we avoided the usual interpretation proposed by UQR users that the coefficient for a binary predictor refers to the impact of changes in group proportions on quantiles of the overall distribution. We find this interpretation difficult to conceive in practice, especially when groups are based on inherent traits like race or sex.

Studying overall quantile gaps enables analysts to devise interventions tailored to specific segments of the distribution — as opposed to relying on a one-size-fits-all approach to combating inequality. For example, we found in our empirical illustration that while unionization helped workers achieve middle wages in the 1980s, unequal access to higher education remained an impediment to high salaries. An analysis centered exclusively on the horizontal decomposition of the mean gap would show only that union workers earn more than nonunion workers, on average, and that this is partly due to the discrepancy in years of work experience between the two groups. By investigating different parts of the distribution, the aKOB decomposition can unveil heterogeneous inequality patterns that are not evident in the classical decomposition.

In discussing estimators for the quantities involved in the decomposition, we introduced a doubly robust alternative that combines the commonly used regression-imputation and weighting estimators. It requires that either the outcome or propensity score model be correctly specified (but not necessarily both) to produce consistent results. Among the advantages of the doubly robust estimator is its compatibility with machine learning models, which may lead to considerable gains in robustness and efficiency. We hope this paper's contributions to decomposition methods help advance research on group inequalities.

8 Appendix A: Double Robustness of the $\hat{\theta}^{DR}$ Estimator

This section demonstrates the double robustness of the $\hat{\theta}^{DR}$ estimator presented in Equation (5). The proof largely follows the steps outlined in Tsiatis (2006) and Glynn and Quinn (2010). To facilitate exposition, we define new notation for the terms introduced in the main text. Let:

$$\mu_1(X_i, \hat{\xi}_n) \equiv \widehat{\mathbb{E}}\left[Y_i | A_i = 1, X_i\right], \text{ and}$$
$$\pi(X_i, \hat{\psi}_n) \equiv \widehat{\Pr}[A_i = 1 | X_i],$$

where $\mu_1(X_i, \hat{\xi}_n)$ is the estimated outcome model in a sample of size n, governed by a finite dimensional parameter ξ. Similarly, $\pi(X_i, \hat{\psi}_n)$ is the estimated propensity score function in that sample, governed by a finite dimensional parameter ψ. We assume that $\hat{\xi}_n$ and $\hat{\psi}_n$ converge in probability to some values ξ^* and ψ^*, respectively, as the sample size approaches infinity. The outcome model and propensity score models are said to be correctly specified if $\xi^* = \xi_0$ and $\psi^* = \psi_0$.

We can then rewrite Equation (5) as follows:

$$\begin{aligned}
\hat{\theta}_n^{DR} = \frac{1}{n}\sum_{i=1}^{n}\Bigg\{ & \frac{\mathbb{I}[A_i = 1]}{1 - \widehat{\Pr}[A_i = 1]} \frac{1 - \pi(X_i, \hat{\psi}_n)}{\pi(X_i, \hat{\psi}_n)} \left[Y_i - \mu_1(X_i, \hat{\xi}_n)\right] \\
& + \frac{\mathbb{I}[A_i = 0]}{1 - \widehat{\Pr}[A_i = 1]} \left(\mu_1(X_i, \hat{\xi}_n) - \widehat{\mathbb{E}}\left[\mu_1(X_i, \hat{\xi}_n) | A = 0\right]\right) \\
& + \widehat{\mathbb{E}}\left[\mu_1(X_i, \hat{\xi}_n) | A = 0\right] \Bigg\}.
\end{aligned}$$

By the law of large numbers, $\hat{\theta}_n^{DR}$ converges in probability to:

$$\begin{aligned}
\operatorname*{plim}_{n\to\infty} \hat{\theta}_n^{DR} = \mathbb{E}\Bigg[& \underbrace{\frac{\mathbb{I}[A = 1]}{1 - \mathbb{E}[A]} \frac{1 - \pi(X, \psi^*)}{\pi(X, \psi^*)} \left[Y - \mu_1(X, \xi^*)\right]}_{1} \\
& + \underbrace{\frac{\mathbb{I}[A = 0]}{1 - \mathbb{E}[A]} \left(\mu_1(X, \xi^*) - \mathbb{E}\left[\mu_1(X, \xi^*) | A = 0\right]\right)}_{2} \\
& + \mathbb{E}\left[\mu_1(X, \xi^*) | A = 0\right] \Bigg].
\end{aligned} \tag{A.1}$$

Let us first consider the case where the outcome model is correctly specified, but no assumptions are made about the propensity score model — that is, $\xi^* = \xi_0$, but ψ^* could differ from ψ_0. To demonstrate that this implies the consistency of $\hat{\theta}^{DR}$, it suffices to show that $\mathbb{E}\left[1\right]$ and $\mathbb{E}\left[2\right]$ equal zero.

The law of iterated expectations allows us to write the following expressions: $\mathbb{E}\left[1\right] = \mathbb{E}\left[\mathbb{E}\left[1 \middle| X\right]\right]$ and $\mathbb{E}\left[1 \middle| X\right] = \mathbb{E}\left[\mathbb{E}\left[1 \middle| X, A\right] \middle| X\right]$. We will show that the latter quantity equals zero. We use the definition of expectation to inspect it further:

$$\begin{aligned}
\mathbb{E}\left[\mathbb{E}\left[1 \middle| X, A\right] \middle| X\right] &= \sum_{a \in \{0,1\}} \mathbb{E}\left[1 \middle| X, A = a\right] \Pr[A = a | X] \\
&= \mathbb{E}\left[1 \middle| X, A = 0\right] \Pr[A = 0 | X] + \mathbb{E}\left[1 \middle| X, A = 1\right] \Pr[A = 1 | X].
\end{aligned}$$

Note that the first term of the sum equals zero since 1 involves an indicator function identifying group members. Substituting $\pi(X, \psi_0)$ in for $\Pr[A = 1|X]$ and writing out ① under the assumption that $\xi^* = \xi_0$ leaves us with the following expression:

$$\begin{aligned}
&\mathbb{E}\left[\frac{1 - \pi(X, \psi^*)}{(1 - \mathbb{E}[A])\, \pi(X, \psi^*)} \left[Y - \mu_1(X, \xi_0)\right] \middle| X, A = 1 \right] \pi(X, \psi_0) \\
&= \pi(X, \psi_0) \frac{1 - \pi(X, \psi^*)}{(1 - \mathbb{E}[A])\, \pi(X, \psi^*)} \mathbb{E}\left[Y - \mu_1(X, \xi_0) | X, A = 1\right],
\end{aligned}$$

by pulling out what is known. But if the outcome model is correctly specified, the term $\mathbb{E}\left[Y - \mu_1(X, \xi_0)|X, A = 1\right]$ will average to zero asymptotically since the residual in that case will have zero conditional mean at each value of X. We have thus shown that $\mathbb{E}\left[\mathbb{E}\left[① \middle| X, A\right] \middle| X\right]$ equals zero, from which it follows that $\mathbb{E}\left[①\right] = 0$.

Now we must show that $\mathbb{E}\left[②\right] = 0$. It is helpful to note that $\mathbb{E}\left[②\right] = \mathbb{E}\left[\mathbb{E}\left[② \middle| A\right]\right]$ and

$$\begin{aligned}
\mathbb{E}\left[\mathbb{E}\left[② \middle| A\right]\right] &= \sum_{a \in \{0,1\}} \mathbb{E}\left[② \middle| A = a\right] \Pr[A = a] \\
&= \mathbb{E}\left[② \middle| A = 0\right] \Pr[A = 0] + \mathbb{E}\left[② \middle| A = 1\right] \Pr[A = 1].
\end{aligned}$$

The second term of the sum equals zero because ② includes an indicator function identifying non-members. Replacing $\Pr[A = 0]$ with $1 - \mathbb{E}[A]$ and writing out ② yields the expression:

$$\begin{aligned}
&\mathbb{E}\left[\frac{(\mu_1(X, \xi_0) - \mathbb{E}\left[\mu_1(X, \xi_0)|A = 0\right])}{1 - \mathbb{E}[A]} \middle| A = 0 \right] (1 - \mathbb{E}[A]) \\
&= \left(\frac{1 - \mathbb{E}[A]}{1 - \mathbb{E}[A]} \right) \mathbb{E}\left[(\mu_1(X, \xi_0) - \mathbb{E}\left[\mu_1(X, \xi_0)|A = 0\right]) \middle| A = 0 \right] \\
&= \mathbb{E}\left[\mu_1(X, \xi_0) \middle| A = 0\right] - \mathbb{E}\left[\mathbb{E}\left[\mu_1\left(X, \xi_0\right)|A = 0\right] \middle| A = 0\right] \\
&= \mathbb{E}\left[\mu_1(X, \xi_0) \middle| A = 0\right] - \mathbb{E}\left[\mu_1(X, \xi_0) \middle| A = 0\right] \\
&= 0,
\end{aligned}$$

by pulling out constants and using the linearity of the expectation operator.

Having demonstrated that $\mathbb{E}\left[①\right] = \mathbb{E}\left[②\right] = 0$, we conclude that $\hat{\theta}_n^{DR}$ converges in probability to $\operatorname*{plim}_{n\to\infty} \hat{\theta}_n^{DR} = \mathbb{E}\left[\mathbb{E}\left[\mu_1(X, \xi_0)|A = 0\right]\right]$ when $\xi^* = \xi_0$. In other words, the doubly robust estimator will be consistent as long as the outcome model used in the regression-imputation estimator is correctly specified — even if the propensity score function for the weighting estimator is not.

We now turn to the case where $\psi^* = \psi_0$ — i.e., we depart from the premise that the propensity score model is correctly specified while making no assumptions about the outcome model. We rearrange the terms in Equation (A.1) to facilitate this part of the proof:

$$\begin{aligned}
\operatorname*{plim}_{n\to\infty} \hat{\theta}_n^{DR} = \mathbb{E}\Bigg[&\left(\frac{\mathbb{I}[A = 1]}{1 - \mathbb{E}[A]} \frac{1 - \pi(X, \psi^*)}{\pi(X, \psi^*)} \right) Y \\
&+ \underbrace{\left(\frac{\mathbb{I}[A = 0]}{1 - \mathbb{E}[A]} - \frac{\mathbb{I}[A = 1]}{1 - \mathbb{E}[A]} \frac{1 - \pi(X, \psi^*)}{\pi(X, \psi^*)} \right) \mu_1(X, \xi^*)}_{③} \\
&+ \underbrace{\left(1 - \frac{\mathbb{I}[A = 0]}{1 - \mathbb{E}[A]} \right) \mathbb{E}\left[\mu_1(X, \xi^*)|A = 0\right]}_{④} \Bigg].
\end{aligned} \tag{A.2}$$

To prove the consistency of $\hat{\theta}^{DR}$ when $\psi^* = \psi_0$, we must show that $\mathbb{E}\left[\textcircled{3}\right] = \mathbb{E}\left[\textcircled{4}\right] = 0$.

First, we proceed by demonstrating that $\mathbb{E}\left[\textcircled{3}\middle| X\right]$ equals zero:

$$\begin{aligned}\mathbb{E}\left[\textcircled{3}\middle| X\right] &= \mathbb{E}\left[\left(\frac{\mathbb{I}[A=0]}{1-\mathbb{E}[A]} - \frac{\mathbb{I}[A=1]}{1-\mathbb{E}[A]}\frac{1-\pi(X,\psi_0)}{\pi(X,\psi_0)}\right)\mu_1(X,\xi^*)\middle| X\right] \\ &= \mu_1(X,\xi^*)\left\{\frac{1}{1-\mathbb{E}[A]}\mathbb{E}[\mathbb{I}[A=0]|X] - \frac{1-\pi(X,\psi_0)}{(1-\mathbb{E}[A])\pi(X,\psi_0)}\mathbb{E}[\mathbb{I}[A=1]|X]\right\},\end{aligned}$$

by pulling out constants and what is known and using the linearity of the expectation operator. Also note that $\mathbb{E}[\mathbb{I}[A=1]|X] = \pi(X,\psi_0)$ and $\mathbb{E}[\mathbb{I}[A=0]|X] = 1-\pi(X,\psi_0)$. Plugging these in and canceling out the terms inside the curly brackets will lead to a zero value. Therefore, $\mathbb{E}\left[\textcircled{3}\middle| X\right]$ equals zero and, consequently, so does $\mathbb{E}\left[\textcircled{3}\right]$.

Next, we show that $\mathbb{E}\left[\textcircled{4}\right] = 0$. We can easily see this by writing out $\textcircled{4}$ and using the properties of the expectation operator, along with the fact that $\mathbb{E}[\mathbb{I}[A=0]] = 1-\mathbb{E}[A]$:

$$\begin{aligned}\mathbb{E}\left[\textcircled{4}\right] &= \mathbb{E}\left[\left(1-\frac{\mathbb{I}[A=0]}{1-\mathbb{E}[A]}\right)\mathbb{E}\left[\mu_1(X,\xi^*)|A=0\right]\right] \\ &= \mathbb{E}\left[\mu_1(X,\xi^*)|A=0\right]\mathbb{E}\left[1-\frac{\mathbb{I}[A=0]}{1-\mathbb{E}[A]}\right] \\ &= \mathbb{E}\left[\mu_1(X,\xi^*)|A=0\right]\left(1-\frac{\mathbb{E}[\mathbb{I}[A=0]]}{1-\mathbb{E}[A]}\right) \\ &= \mathbb{E}\left[\mu_1(X,\xi^*)|A=0\right]\left(1-\frac{1-\mathbb{E}[A]}{1-\mathbb{E}[A]}\right) \\ &= 0.\end{aligned}$$

Since $\mathbb{E}\left[\textcircled{3}\right] = \mathbb{E}\left[\textcircled{4}\right] = 0$, Equation (A.2) reduces to:

$$\operatorname*{plim}_{n\to\infty}\hat{\theta}_n^{DR} = \mathbb{E}\left[\left(\frac{\mathbb{I}[A=1]}{1-\mathbb{E}[A]}\frac{1-\pi(X,\psi_0)}{\pi(X,\psi_0)}\right)Y\right]$$

when $\psi^* = \psi_0$. This implies that the doubly robust estimator will be consistent when the propensity score model for the weighting estimator is correctly specified — regardless of whether the same holds for the outcome model in the regression-imputation estimator.

This completes the proof.

9 Appendix B: Proof of the Mathematical Identities in Section 4

We now prove the mathematical identities in Section 4. We begin by demonstrating Lemma (12), which we use as a stepping stone to other results presented in that section. Throughout, we assume that the dependent variable has a finite lower bound C.

$$
\begin{aligned}
&\int_0^1 \mathrm{RIF}(y; Q_p)dp \\
&= \int_0^1 \left[Q_p + \frac{p - \mathbb{I}(y \le Q_p)}{f_Y(Q_p)} \right] dp \\
&= \int_{-\infty}^{\infty} \left[t + \frac{\Pr[Y \le t] - \mathbb{I}(y \le t)}{f_Y(t)} \right] dF_Y(t) \;\; (\text{Replace } Q_p \text{ with } t = F^{-1}(p) \text{ and } p \text{ with } \Pr[Y \le t] = F_Y(t)) \\
&= \int_{-\infty}^{\infty} t dF_Y(t) + \int_{-\infty}^{\infty} \{\Pr[Y \le t] - \mathbb{I}(y \le t)\} \, dt \;\; (\text{Since } dF_Y(t)/dt = f_Y(t) \Leftrightarrow dt = dF_Y(t)/f_Y(t)) \\
&= \int_{-\infty}^{\infty} t dF_Y(t) + \int_{C}^{\infty} \{\Pr[Y \le t] - \mathbb{I}(y \le t)\} \, dt \;\; (Y \text{ has a finite lower bound } C) \\
&= \int_{-\infty}^{\infty} t dF_Y(t) + \int_{C}^{\infty} \{\mathbb{I}(y > t) - \Pr[Y > t]\} \, dt \;\; (\text{Since } \Pr[Y \le t] = 1 - \Pr[Y > t] \text{ and } \mathbb{I}(y \le t) = 1 - \mathbb{I}(y > t)) \\
&= \mathbb{E}[Y] + (y - C) - (\mathbb{E}[Y] - C) \\
&= y.
\end{aligned}
\tag{B.1}
$$

To understand the penultimate line, note that

$$
\mathbb{I}(y > t) = \begin{cases} 1, & \text{if } y > t \Leftrightarrow t \in (C, y) \\ 0, & \text{otherwise.} \end{cases}
$$

So we can rewrite the integral $\int_C^{\infty} \mathbb{I}(y > t)dt$ as $\int_C^{y} 1dt$, yielding $y - C$.

Additionally, note that $\Pr[Y > t] = \mathbb{E}[\mathbb{I}(Y > t)]$. By the definition of expectation, $\mathbb{E}[\mathbb{I}(Y > t)] = \int_{-\infty}^{\infty} \mathbb{I}(y > t)dF_Y(y)$. And since we have imposed a finite lower bound C for the outcome, we can also constrain this integral to a lower bound C, leading to the following expression:

$$
\begin{aligned}
\int_C^{\infty} \Pr[Y > t]dt &= \int_C^{\infty} \mathbb{E}[\mathbb{I}(Y > t)]dt \\
&= \int_C^{\infty} \int_C^{\infty} \mathbb{I}(y > t)dF_Y(y)dt \\
&= \int_C^{\infty} \int_C^{\infty} \mathbb{I}(y > t)dtdF_Y(y) \;\; (\text{Fubini's Theorem}) \\
&= \int_C^{\infty} (y - C)dF_Y(y) \\
&= \mathbb{E}[Y - C] \\
&= \mathbb{E}[Y] - C.
\end{aligned}
$$

Next, we use the lemma above to show that $\int_0^1 m_a(x, Q_p)dp = \mu_a(x)$:

$$
\begin{aligned}
\int_0^1 m_a(x, Q_p)dp &= \int_0^1 \mathbb{E}[\mathrm{RIF}(Y; Q_p)|A = a, X = x]dp \\
&= \int_0^1 \int_{-\infty}^{\infty} \mathrm{RIF}(y; Q_p) dF_{Y|A=a,X=x}(y)dp \;\; \text{(Definition of expectation)} \\
&= \int_{-\infty}^{\infty} \int_0^1 \mathrm{RIF}(y; Q_p) dp dF_{Y|A=a,X=x}(y) \;\; \text{(Fubini's Theorem)} \\
&= \int_{-\infty}^{\infty} y dF_{Y|A=a,X=x}(y) \;\; \text{(Lemma (B.1))} \\
&= \mathbb{E}[Y|A = a, X = x] \\
&= \mu_a(x).
\end{aligned} \tag{B.2}
$$

Equation (B.2), in turn, implies that

$$
\begin{aligned}
\int_0^1 \Delta_X(p)dp &= \int_0^1 \int m_1(x, Q_p) \left[dF_{X|A=1}(x) - dF_{X|A=0}(x)\right] dp \;\; \text{(Substitute in expression from Equation (11))} \\
&= \int \int_0^1 m_1(x, Q_p)dp \left[dF_{X|A=1}(x) - dF_{X|A=0}(x)\right] \;\; \text{(Fubini's Theorem)} \\
&= \int \mu_1(x) \left[dF_{X|A=1}(x) - dF_{X|A=0}(x)\right] \;\; \text{(Equation (B.2))} \\
&= \Delta_X \;\; \text{(Equation (1))},
\end{aligned} \tag{B.3}
$$

and that

$$
\begin{aligned}
\int_0^1 \Delta_U(p)dp &= \int_0^1 \int \left[m_1(x, Q_p) - m_0(x, Q_p)\right] dF_{X|A=0}(x)dp \;\; \text{(Substitute in expression from Equation (11))} \\
&= \int \int_0^1 \left[m_1(x, Q_p) - m_0(x, Q_p)\right] dp dF_{X|A=0}(x) \;\; \text{(Fubini's Theorem)} \\
&= \int \left[\mu_1(x) - \mu_0(x)\right] dF_{X|A=0}(x) \;\; \text{(Linearity of integration and Equation (B.2))} \\
&= \Delta_U \;\; \text{(Equation (1))}.
\end{aligned} \tag{B.4}
$$

References

Ashenfelter, Orley. 1978. "Estimating the Effect of Training Programs on Earnings." *Rev. Econ. Stat.* 60 (1): 47–57. https://doi.org/10.2307/1924332.

Barsky, Robert, John Bound, Kerwin Ko'Charles, and Joseph P. Lupton. 2002. "Accounting for the Black-White Wealth Gap: A Nonparametric Approach." *J. Am. Stat. Assoc.* 97 (459): 663–73. https://doi.org/10.1198/016214502388618401.

Blinder, Alan S. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *J. Hum. Resour.* 8 (4): 436–55. https://doi.org/10.2307/144855.

Borgen, Nicolai T., Andreas Haupt, and Øyvind Nicolay Wiborg. 2023. "Quantile Regression Estimands and Models: Revisiting the Motherhood Wage Penalty Debate." *Eur. Sociol. Rev.* 39 (2): 317–31. https://doi.org/10.1093/esr/jcac052.

Cai, Lixin, and Amy Y. C. Liu. 2008. "Union Wage Effects in Australia: Is There Variation Along the Distribution?" *Econ. Rec.* 84 (267): 496–510. https://doi.org/10.1111/j.1475-4932.2008.00513.x.

Card, David. 1996. "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis." *Econometrica* 64 (4): 957–79. https://doi.org/10.2307/2171852.

Card, David, Thomas Lemieux, and W. Craig Riddell. 2004. "Unions and Wage Inequality." *J. Labor Res.*, Journal of Labor Research, 25 (4): 519–62. https://doi.org/10.1007/s12122-004-1011-z.

Chamberlain, Gary. 1994. "Quantile Regression, Censoring, and the Structure of Wages." In *Advances in Econometrics*, 171–210. Cambridge University Press. https://doi.org/10.1017/CCOL0521444594.005.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. "Double/Debiased Machine Learning for Treatment and Structural Parameters." *Econom. J.* 21 (1): C1–68. https://doi.org/10.1111/ectj.12097.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica* 64 (5): 1001–44. https://doi.org/10.2307/2171954.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77 (3): 953–73. https://doi.org/10.3982/ECTA6822.

———. 2018. "Decomposing Wage Distributions Using Recentered Influence Function Regressions." *Econometrics* 6 (2): 28–28. https://doi.org/10.3390/econometrics6020028.

Fortin, Nicole M., Thomas Lemieux, and Sergio Firpo. 2011. "Chapter 1 - Decomposition Methods in Economics." In *Handbook of Labor Economics*, edited by Orley Ashenfelter and David Card, 4:1–102. Elsevier. https://doi.org/10.1016/S0169-7218(11)00407-2.

Freeman, Richard B. 1980. "Unionism and the Dispersion of Wages." *Ind. Labor Relat. Rev.*, Industrial and Labor Relations Review, 34 (1): 3–23. https://doi.org/10.1177/001979398003400101.

———. 1993. "How Much Has De-Unionization Contributed to the Rise in Male Earnings Inequality?" In *Uneven Tides*, 133–64. Russell Sage Foundation. https://doi.org/10.7758/9781610441469.6.

Freeman, Richard B., and James L. Medoff. 1984. *What Do Unions Do?* New York: Basic Books.

Glynn, Adam N., and Kevin M. Quinn. 2010. "An Introduction to the Augmented Inverse Propensity Weighted Estimator." *Polit. Anal.* 18 (1): 36–56. https://doi.org/10.1093/pan/mpp036.

Hyndman, Rob J., and Yanan Fan. 1996. "Sample Quantiles in Statistical Packages." *Am. Stat.* 50 (4): 361–61. https://doi.org/10.2307/2684934.

Jann, Ben. 2008. "The Blinder–Oaxaca Decomposition for Linear Regression Models." *Stata J.* 8 (4): 453–79. https://doi.org/10.1177/1536867X0800800401.

Killewald, Alexandra, and Jonathan Bearak. 2014. "Is the Motherhood Penalty Larger for Low-Wage Women? A Comment on Quantile Regression." *Am. Sociol. Rev.* 79 (2): 350–57. https://doi.org/10.1177/0003122414524574.

Kitagawa, Evelyn M. 1955. "Components of a Difference Between Two Rates." *J. Am. Stat. Assoc.* 50 (272): 1168–94. https://doi.org/10.1080/01621459.1955.10501299.

Kline, Patrick. 2011. "Oaxaca-Blinder as a Reweighting Estimator." *Am. Econ. Rev.* 101 (3): 532–37. https://doi.org/10.1257/aer.101.3.532.

Lemieux, Thomas. 2006. "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *Am. Econ. Rev.* 96 (3): 461–98. https://doi.org/10.1257/aer.96.3.461.

———. 2008. "The Changing Nature of Wage Inequality." *J. Popul. Econ.* 21 (1): 21–48. https://doi.org/10.1007/s00148-007-0169-0.

Mandel, Hadas, and Moshe Semyonov. 2014. "Gender Pay Gap and Employment Sector: Sources of Earnings Disparities in the United States, 1970–2010." *Demography* 51 (5): 1597–1618. https://doi.org/10.1007/s13524-014-0320-y.

Montenovo, Laura, Xuan Jiang, Felipe Lozano-Rojas, Ian Schmutte, Kosali Simon, Bruce A. Weinberg, and Coady Wing. 2022. "Determinants of Disparities in Early COVID-19 Job Losses." *Demography* 59 (3): 827–55. https://doi.org/10.1215/00703370-9961471.

Oaxaca, Ronald. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *Int. Econ. Rev. Phila.* 14 (3): 693–709. https://doi.org/10.2307/2525981.

Pepin, Joanna R., Liana C. Sayer, and Lynne M. Casper. 2018. "Marital Status and Mothers' Time Use: Childcare, Housework, Leisure, and Sleep." *Demography* 55 (1): 107–33. https://doi.org/10.1007/s13524-018-0647-x.

Rios-Avila, Fernando. 2020. "Recentered Influence Functions (RIFs) in Stata: RIF Regression and RIF Decomposition." *Stata J.* 20 (1): 51–94. https://doi.org/10.1177/1536867X20909690.

Rios-Avila, Fernando, and Barry T. Hirsch. 2014. "Unions, Wage Gaps, and Wage Dispersion: New Evidence from the Americas." *Ind. Relat. Berkeley* 53 (1): 1–27. https://doi.org/10.1111/irel.12044.

Rios-Avila, Fernando, and Michelle Lee Maroto. 2022. "Moving Beyond Linear Regression: Implementing and Interpreting Quantile Regression Models With Fixed Effects." *Sociological Methods & Research*, February, 00491241211036165. https://doi.org/10.1177/00491241211036165.

Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao. 1994. "Estimation of Regression Coefficients When Some Regressors Are Not Always Observed." *J. Am. Stat. Assoc.* 89 (427): 846–66. https://doi.org/10.1080/01621459.1994.10476818.

Tchetgen, Eric J. Tchetgen, and Ilya Shpitser. 2012. "Semiparametric Theory for Causal Mediation Analysis: Efficiency Bounds, Multiple Robustness and Sensitivity Analysis." *Ann. Stat.* 40 (3): 1816–45. https://doi.org/10.1214/12-AOS990.

Tilly, Charles. 1998. *Durable Inequality.* Berkeley: University of California Press.

Tsiatis, Anastasios. 2006. *Semiparametric Theory and Missing Data.* 1st ed. 2006. Springer Series in Statistics. New York, NY: Springer New York : Imprint: Springer.

Wenz, Sebastian E. 2019. "What Quantile Regression Does and Doesn't Do: A Commentary on Petscher and Logan (2014)." *Child Dev.* 90 (4): 1442–52. https://doi.org/10.1111/cdev.13141.

Western, Bruce, and Jake Rosenfeld. 2011. "Unions, Norms, and the Rise in U.S. Wage Inequality." *Am. Sociol. Rev.* 76 (4): 513–37. https://doi.org/10.1177/0003122411414817.