

A Causal Detailed Decomposition

Identifying the sources of differences in distributions

Philipp Ketz¹ Blaise Melly² Miguel Salema²

September 23, 2026

¹Paris School of Economics

²University of Bern

Motivation: why do groups differ?

Why do groups differ is one of the most common questions in economics.

- Why did GDP per capita grow from 1909 to 1949? Solow (1957)
- Why are people living longer? Arriaga (1984)
- Why did the wage distribution widen in the 1980s? Juhn et al. (1993)

Answering them requires splitting a total difference between factors $\Rightarrow$ **Decomposition**.

Same research question repeatedly appears in applied work

Go beyond *how much* of a gap there is $\Rightarrow$ What are the sources of group differences?

This is a causal question!

The workhorse: detailed Oaxaca-Blinder decomposition

Study	Outlet	Question asked
<i>Economics</i>		
Altonji and Blank (1999)	<i>Handbook of Labor Economics</i>	Which factors produce the black-white and male-female wage gaps?
Blau and Kahn (2017)	<i>JEL</i>	How much of the gender wage gap is due to education, experience, occupation, industry, unionisation?
Currie et al. (2023)	<i>AER</i>	What explains the fall in racial disparities in particulate exposure? (Use RIF-OB)
Fryer and Levitt (2004)	<i>REStat</i>	Why do black children score lower on standardised tests?

Mind the Gap

Aim: Researchers are after the **sources** of the gap.

Problem: the tool they use (naive detailed OB) does not answer the question.

Contributions

(i) *The critique.*

- The naive detailed decomposition does not identify causal sources of group differences.
- The same applies to its distributional analogues (RIF).

(ii) *The solution.*

- Develop a detailed decomposition that has a clear causal interpretation.
- Requires a stand on a recursive ordering of the covariates.
- Simple to compute using a sequence of standard aggregate decompositions.

(iii) *General conditions beyond the mean.*

- Nonparametric identification, estimation and inference.
- Available for any estimand calculated from a CDF, such as quantiles.
- Extends the mediation literature, stated for means, to entire distributions.

(iv) *Application to the black-white wage gap.*

- Decompose between family background, skills and job characteristics.
- Share attributed to family moves 10% → 40% from naive to causal decomposition.

Running example: black-white wage gap

- $G \in \mathcal{G} = \{b, w\}$; black and white, respectively.
- Outcome Y is log hourly wages.
- Covariates $X = (1, P, S)'$, where
 - P are parental/family characteristics
 - S is schooling.

Questions this literature asks

- How much of the gap is set *before* labour market entry? (Neal and Johnson, 1996)
 - Factors causing it: education, experience, occupation, industry? (Altonji and Blank, 1999)
 - Why do black children already score lower at school entry? (Fryer and Levitt, 2004)
 - Why does the gap persist across generations? (Chetty et al., 2020)
-
- $\hat{\beta}_g$ denotes the OLS estimator of regressing Y on X at subsample $g \in \mathcal{G}$.
 - $\bar{Y}_g$, $\bar{X}_g$ and are sample means at subsample $g \in \mathcal{G}$.

The linear case

Oaxaca-Blinder decomposition

Kitagawa (1955); Oaxaca (1973); Blinder (1973).

Aggregate decomposition

OLS will always give $\bar{Y}_g = \bar{X}_g' \hat{\beta}_g$, so:

$$\bar{Y}_w - \bar{Y}_b = \underbrace{\bar{X}_b'(\hat{\beta}_w - \hat{\beta}_b)}_{\text{unexplained / structure}} + \underbrace{(\bar{X}_w - \bar{X}_b)' \hat{\beta}_w}_{\text{explained / composition}}.$$

Under A.1-A.2, causal interpretation is **Solved**. (Fortin et al., 2011):

- Structure is ATT of G among black workers.
- Composition is the causal effect of moving characteristics from $\bar{X}_b$ to $\bar{X}_w$.

Source: Cahuc et al. (2014), ch. 8, and Altonji and Blank (1999) for textbook treatments of the detailed decomposition.

Potential outcomes and the aggregate decomposition

- $Y(g, x)$ be the potential outcome when $G = g$, $X = x$.
- $X(g)$ is the potential covariate value for $G = g$.

Observation rule: $X = X(G)$, $Y = Y(G, X)$.

Definition: $Y(g) := Y(g, X)$.

Assumption 1 (linearity, within-group strict exogeneity)

For $g \in \mathcal{G}$, $Y(g, x) = x' \beta_g + U_g$, where $\mathbb{E}[U_g \mid G = g, X] = 0$.

Assumption 2 (cross-group mean independence)

For $g \in \mathcal{G}$, $\mathbb{E}[U_g \mid G] = 0$.

Then by A1, A2 and standard OLS regularity conditions:

$$\begin{aligned}\bar{X}'_b(\hat{\beta}_w - \hat{\beta}_b) &\xrightarrow{P} \mathbb{E}[Y(w) - Y(b) \mid G = b], \\ (\bar{X}_w - \bar{X}_b)' \hat{\beta}_w &\xrightarrow{P} \mathbb{E}[Y(w) \mid G = w] - \mathbb{E}[Y(w) \mid G = b].\end{aligned}$$

Naive detailed decomposition

- Mainstream: economics journals, handbooks, textbooks, policy reports.
- Accounting identity holds by construction.

$$(\bar{X}_w - \bar{X}_b)' \hat{\beta}_w = \underbrace{(\bar{P}_w - \bar{P}_b) \hat{\beta}_w^P}_{\text{"due to" family}} + \underbrace{(\bar{S}_w - \bar{S}_b) \hat{\beta}_w^S}_{\text{"due to" schooling}}$$

The causal experiment


Change the mechanism generating P , and let the remainder of the covariate system respond naturally.

$\neq$

The naive detailed OB

A *joint* intervention: move P **and** simultaneously prevent S from responding, which breaks the system that generates the covariates.

Why the naive detailed fails: simplified running example



Takeaway

All of the composition effect originates in family background.

- Eg. worker's parents were discriminated against.

Problem

- The naive detailed decomposition attributes 100% of the gap to S .
- Technical reason: $\hat{\beta}_w^P \approx 0$.

Formally: a simultaneous system

For each $g \in \mathcal{G}$, let the covariates be generated by

$$\begin{aligned}P(g, s) &= \gamma_g^P + \gamma_g^{S \rightarrow P} s + U_g^P, \\S(g, p) &= \gamma_g^S + \gamma_g^{P \rightarrow S} p + U_g^S.\end{aligned}$$

Then the components of the naive detailed decomposition converge to


$$(\bar{P}_w - \bar{P}_b) \hat{\beta}_w^P \xrightarrow{P} \left[\frac{\gamma_w^P + \gamma_w^{S \rightarrow P} \gamma_w^S}{1 - \gamma_w^{S \rightarrow P} \gamma_w^{P \rightarrow S}} - \frac{\gamma_b^P + \gamma_b^{S \rightarrow P} \gamma_b^S}{1 - \gamma_b^{S \rightarrow P} \gamma_b^{P \rightarrow S}} \right] \beta_w^P, \quad \text{Ditto for } S.$$

A simultaneous system is not identified without further assumptions ([Working, 1927](#)).

Any assumption that identifies the system makes progress

- We impose a triangular structure. (When causal/temporal ordering allows it.)
Standard in econometrics: [Newey et al. \(1999\)](#), [Chesher \(2003\)](#), [Imbens and Newey \(2009\)](#).
- Instruments are hard to find in typical applications.

Triangularity is quite natural in the example



Assumption 4 (triangularity and linearity)

For $g \in \mathcal{G}$, the covariates follow a recursive structure:

$$\begin{aligned}P(g, s) &= P(g) = \gamma_g^P + U_g^P, \\S(g, p) &= \gamma_g^S + \gamma_g^{P \rightarrow S} p + U_g^S,\end{aligned}$$

with $\mathbb{E}[U_g^P \mid G = g] = 0$ and $\mathbb{E}[U_g^S \mid G = g, P] = 0$.

Assumption 5 (cross-group mean independence; schooling equation)

For $g \in \mathcal{G}$, $\mathbb{E}[U_g^S \mid G] = 0$.


The building blocks

$$\Theta_{r,h} \equiv \mathbb{E} [Y(w, P(h), S(r, P(h))) \mid G = h] \quad \forall r, h \in \mathcal{G}.$$

$$\Delta_P \equiv \Theta_{w,w} - \Theta_{w,b},$$

$$\Delta_S \equiv \Theta_{w,b} - \Theta_{b,b},$$

$$\Delta_P + \Delta_S = \mathbb{E}[Y(w) \mid G = w] - \mathbb{E}[Y(w) \mid G = b]$$



The detailed components

Under Assumptions 1, 2, 4 and 5

$$\Delta_P = (\gamma_w^P - \gamma_b^P) \left(\underbrace{\beta_w^P}_{\text{direct}} + \underbrace{\gamma_w^{P \rightarrow S} \beta_w^S}_{\text{through } S} \right),$$
$$\Delta_S = (\gamma_w^S - \gamma_b^S) \beta_w^S + \gamma_b^P (\gamma_w^{P \rightarrow S} - \gamma_b^{P \rightarrow S}) \beta_w^S.$$

- Δ_P : move the *distribution* of P from b to w , letting S adjust. Direct *and* indirect effect.
- Δ_S : move the *process* of S from b to w , at the black distribution of P .

Source: A similar problem can be found in ([Daniel et al., 2015](#)).

Simple implementation

Downside of direct plug-in

Requires the conditional distribution of each covariate given its predecessors.

A sequence of aggregate decompositions

1. Run the aggregate OB with P only. The explained component is $\hat{\Delta}_P$.
2. Run it again with (P, S) . The *increment* in the explained component is $\hat{\Delta}_S$.

Proposition 1 (equivalence of plug-in and sequential decompositions)

For the mean the two approaches are finite-sample numerically identical.

- This gives a causal foundation to sequential covariate introduction (DiNardo et al., 1996; Butcher and DiNardo, 2002; Altonji et al., 2012).
- Path dependence was the standard objection to that practice (Fortin et al., 2011; Gelbach, 2016). *The sensitivity is a feature.*

Nonparametric identification of distributional decompositions

From means to distributions

We now decompose

$$F_Y(y \mid G = w) - F_Y(y \mid G = b) = \underbrace{F_{Y(w)}(y \mid G = b) - F_{Y(b)}(y \mid G = b)}_{\text{unexplained}} + \underbrace{F_{Y(w)}(y \mid G = w) - F_{Y(w)}(y \mid G = b)}_{\text{explained}}$$

No linearity, no parametric model for the dependence between covariates.

Assumption 7 (Sequential independence assumption)

For $g \in \mathcal{G}$, $x \in \mathcal{X}_b \cup \mathcal{X}_w$: (i) $Y(g, x) \perp\!\!\!\perp G$, (ii) $Y(g, x) \perp\!\!\!\perp X \mid G$.

Assumption 8 (common support)

$\mathcal{X}_b = \mathcal{X}_w$.

Proposition 2

Under Assumptions 7–8, for any $g, g' \in \mathcal{G}$,

$$F_{Y(g)}(y \mid G = g') = F_{Y\langle g|g' \rangle}(y) := \int_{\mathcal{X}_{g'}} F_Y(y \mid G = g, X = x) dF_X(x \mid G = g'),$$

which is observable.

- Adding $X(g) \perp\!\!\!\perp G$ delivers the *unconditional* components,
 $F_{Y(g, X(g'))}(y) = F_{Y\langle g|g' \rangle}(y)$.
- This is the distributional analogue of the aggregate OB result; nothing new so far ([Chernozhukov et al., 2013](#)).

Detailed identification: the main result

Assumption 10 (Triangularity)

$$P(g, s) = P(g) \quad \forall g \in \mathcal{G}.$$

Assumption 11 (sequential independence; detailed)

For any $g, g' \in \mathcal{G}, p, s : (i) \{Y(g, p, s), S(g', p)\} \perp\!\!\!\perp G;$

$(ii) \{Y(g, p, s), S(g', p)\} \perp\!\!\!\perp P \mid G;$

$(iii) Y(g, p, s) \perp\!\!\!\perp S \mid G, P.$

Proposition 3 (main result)

Under Triangularity, Assumption 11 and common support

$$\begin{aligned} F_{Y(g, P(g'), S(g'', P(g')))}(y \mid G = g') &= F_{Y\langle g \mid g', g'' \rangle}(y) \equiv \\ &\iint F_Y(y \mid G = g, P = p, S = s) dF_S(s \mid G = g'', P = p) dF_P(p \mid G = g'), \\ &\quad \forall g, g', g'' \in \mathcal{G} \end{aligned}$$

The distributional decomposition

Identified by observable objects: (Using $F_{Y\langle w|b,b\rangle}(y) = F_{Y\langle w|b\rangle}(y)$.)

$$\underbrace{F_Y(y \mid G = w) - F_{Y\langle w|b,w\rangle}(y)}_{\Delta_P} \quad \text{and} \quad \underbrace{F_{Y\langle w|b,w\rangle}(y) - F_{Y\langle w|b\rangle}(y)}_{\Delta_S},$$

Two routes to the same estimand

1. **Direct:** plug in estimates of $F_Y(y \mid P, S, G)$, $F_S(s \mid P, G)$ and $F_P(p \mid G)$.
 - Hits a computational wall fast: five continuous regressors and a grid of 100 quantile regressions per distribution gives 100^5 pseudo-observations at the last step.
2. **Sequential:** taking differences of successive explained components. $F_{Y\langle w|b,w\rangle}$ is exactly the aggregate decomposition run with P alone. Feasible.

Equivalence of plug-in and sequential decompositions

The counterfactual distributions coincide exactly in the population. The mean estimators coincide in finite samples.

Estimation and inference

Plug-in, following Chernozhukov, Fernández-Val and Melly (2013)

1. Linear QR by group, on a grid $\tau_1, \dots, \tau_T$ uniform on $(0, 1)$:

$$Q_Y(\tau \mid X = x, G = g) = x' \beta_g(\tau) \implies \hat{\beta}_g(\tau_t), t = 1, \dots, T.$$

2. Invert the fitted quantile process into a conditional CDF:

$$\hat{F}_Y(y \mid X = x, G = w) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}\{x' \hat{\beta}_w(\tau_t) \leq y\}.$$

3. Integrate over the *other* group's covariates:

$$\widehat{F}_{Y\langle w|b \rangle}(y) = \frac{1}{n_b} \sum_{i: G_i=b} \hat{F}_Y(y \mid X = X_i, G = w).$$

White wage structure, black characteristics.

Sequential: one decomposition per block

Run steps 1–3 with P alone, then with (P, S) :

$$\hat{\Delta}_P(y) = \hat{F}_Y(y \mid G = w) - \hat{F}_{Y_{\langle w|b,w \rangle}}(y), \quad \hat{\Delta}_S(y) = \hat{F}_{Y_{\langle w|b,w \rangle}}(y) - \hat{F}_{Y_{\langle w|b \rangle}}(y).$$

- Each step is one call to a standard aggregate decomposition.
- The *increment* in the explained component is the block's contribution.
- The whole sequence is bootstrapped jointly.

The parametric assumption is not reasonable at every step

The parametric assumptions are not reasonable here

- Identification (Propositions 2–3) is *fully nonparametric*. Linear QR is an implementation convenience, **no part of the theory asks for it, practitioners do**.
- Imposing linear QR at *every* τ , and at *every* step of the sequence $\implies$ misspecification (at least) in the early steps because linear conditional quantiles are not generally preserved by marginalisation $\implies$ we lose consistency/AN for the intended counterfactual distributions.


The fix: impose no family. (Where we are heading.)

- DML: cross-fitting and Neyman-orthogonal moments at each step (Belloni et al., 2017; Hsu et al., forthcoming).
- Sieve QR (Belloni et al., 2019): nonparametric asymptotically, so the failure vanishes while staying inside the QR framework.

Application: the black–white wage gap

Setting

NLSY 1979; outcome is log hourly wages in 1990; white workers are the reference group.




1. **Family background:** parents' education, age, gender, urban/rural upbringing, living with both parents at 18.
2. **Skills:** years of schooling, AFQT.
3. **Job characteristics:** occupation (blue/white collar), sector (services/manufacturing).

Sign convention: black minus white, so negative numbers mean lower black wages.

The decomposition at the mean

	Causal		Naive detailed	
	log hourly wage	Prop.	log hourly wage	Prop.
	(1)	(2)	(3)	(4)
Wage gap	-.224 (.022)	100.0%	-.224 (.022)	100.0%
Explained	-.191 (.021)	85.1%	-.191 (.021)	85.1%
Structure	-.033 (.026)	14.9%	-.033 (.026)	14.9%
Detailed:				
Family background	-.089 (.010)	39.8%	-.023 (.009)	10.5%
Skills	-.076 (.017)	34.0%	-.111 (.018)	49.5%
Job characteristics	-.025 (.006)	11.3%	-.056 (.006)	25.2%

Across the distribution



Conclusion

Conclusion

What we show

1. The naive detailed decomposition does not identify causal sources.
2. Under a triangular ordering and sequential independence, a detailed decomposition of the *whole distribution* is identified.
3. It is implemented as a sequence of standard aggregate decompositions.
4. In the black–white wage gap, the share attributed to family goes from 10% to 40%.

Limitations

1. Requires a credible causal ordering of covariates;
2. Sequential independence is strong.

Next

- Analyze the setting of the Gelbach decomposition.
- Estimation without a parametric family at each step (DML or sieve QR)

Thank you.

References i

- Altonji, J. G. and Blank, R. M. (1999). Race and gender in the labor market. *Handbook of Labor Economics*, 3C, 3143–3259.
- Altonji, J. G., Bharadwaj, P. and Lange, F. (2012). Changes in the characteristics of American youth. *Journal of Labor Economics*, 30(4), 783–828.
- Angrist, J., Chernozhukov, V. and Fernández-Val, I. (2006). Quantile regression under misspecification. *Econometrica*, 74(2), 539–563.
- Arriaga, E. E. (1984). Measuring and explaining the change in life expectancies. *Demography*, 21(1), 83–96.
- Avin, C., Shpitser, I. and Pearl, J. (2005). Identifiability of path-specific effects. *IJCAI*, 357–363.
- Belloni, A., Chernozhukov, V., Fernández-Val, I. and Hansen, C. (2017). Program evaluation and causal inference with high-dimensional data. *Econometrica*, 85(1), 233–298.
- Belloni, A., Chernozhukov, V., Chetverikov, D. and Fernández-Val, I. (2019). Conditional quantile processes based on series or many regressors. *Journal of Econometrics*, 213(1), 4–29.
- Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: extent, trends, and explanations. *Journal of Economic Literature*, 55(3), 789–865.
- Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 8(4), 436–455.
- Blom, E., Mabel, Z. and Monarrez, T. (2020). Understanding equity gaps in college graduation. Urban Institute.

- Bollen, K. A. and Pearl, J. (2013). Eight myths about causality and structural equation models. *Handbook of Causal Analysis for Social Research*, 301–328.
- Butcher, K. F. and DiNardo, J. (2002). The immigrant and native-born wage distributions. *ILR Review*, 56(1), 97–121.
- Cahuc, P., Carcillo, S. and Zylberberg, A. (2014). *Labor Economics*. MIT Press, 2nd ed.
- Chernozhukov, V., Fernández-Val, I. and Melly, B. (2013). Inference on counterfactual distributions. *Econometrica*, 81(6), 2205–2268.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1), C1–C68.
- Chesher, A. (2003). Identification in nonseparable models. *Econometrica*, 71(5), 1405–1441.
- Chetty, R., Hendren, N., Jones, M. R. and Porter, S. R. (2020). Race and economic opportunity in the United States. *Quarterly Journal of Economics*, 135(2), 711–783.
- Choi, J. H., McCargo, A., Neal, M., Goodman, L. and Young, C. (2019). Explaining the black–white homeownership gap. Urban Institute.
- Currie, J., Voorheis, J. and Walker, R. (2023). What caused racial disparities in particulate exposure to fall? *American Economic Review*, 113(1), 71–97.

- Daniel, R. M., De Stavola, B. L., Cousens, S. N. and Vansteelandt, S. (2015). Causal mediation analysis with multiple mediators. *Biometrics*, 71(1), 1–14.
- DiNardo, J., Fortin, N. M. and Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973–1992. *Econometrica*, 64(5), 1001–1044.
- Eurostat (2021). Decomposition of the unadjusted gender pay gap. Statistical Working Paper.
- Firpo, S., Fortin, N. M. and Lemieux, T. (2018). Decomposing wage distributions using recentered influence function regressions. *Econometrics*, 6(2), 28.
- Fortin, N., Lemieux, T. and Firpo, S. (2011). Decomposition methods in economics. *Handbook of Labor Economics*, 4A, 1–102.
- Fryer, R. G. and Levitt, S. D. (2004). Understanding the black–white test score gap in the first two years of school. *Review of Economics and Statistics*, 86(2), 447–464.
- Gelbach, J. B. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, 34(2), 509–543.
- Greenland, S., Robins, J. M. and Pearl, J. (1999). Confounding and collapsibility in causal inference. *Statistical Science*, 14(1), 29–46.
- Heckman, J. J. (2008). Econometric causality. *International Statistical Review*, 76(1), 1–27.

- Hsu, Y.-C., Huber, M. and Yen, C.-Y. (forthcoming). Doubly robust estimation of distributional mediation effects. *Journal of Business & Economic Statistics*.
- Huber, M. (2015). Causal pitfalls in the decomposition of wage gaps. *Journal of Business & Economic Statistics*, 33(2), 179–191.
- Imbens, G. W. and Newey, W. K. (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica*, 77(5), 1481–1512.
- Juhn, C., Murphy, K. M. and Pierce, B. (1993). Wage inequality and the rise in returns to skill. *Journal of Political Economy*, 101(3), 410–442.
- Kino, S., Bernabé, E. and Sabbah, W. (2020). How much do preventive health behaviors explain education- and income-related inequalities in health? *Preventive Medicine*, 132, 105998.
- Kitagawa, E. M. (1955). Components of a difference between two rates. *Journal of the American Statistical Association*, 50(272), 1168–1194.
- Lundberg, I. (2024). The gap-closing estimand. *Sociological Methods & Research*, 53(3), 1098–1139.
- Neal, D. and Johnson, W. (1996). The role of pre-market factors in black–white differences. *Journal of Political Economy*, 104, 869–895.
- Newey, W. K., Powell, J. L. and Vella, F. (1999). Nonparametric estimation of triangular simultaneous equations models. *Econometrica*, 67(3), 565–603.

- Oaxaca, R. (1973). Male–female wage differentials in urban labor markets. *International Economic Review*, 14(3), 693–709.
- Quintas-Martínez, V., Bahadori, M. T., Santiago, E., Mu, J. and Heckerman, D. (2024). Multiply-robust causal change attribution. *ICML*.
- Rothe, C. (2015). Decomposing the composition effect. *Journal of Business & Economic Statistics*, 33(3), 323–337.
- Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics*, 39(3), 312–320.
- VanderWeele, T. J. and Vansteelandt, S. (2014). Mediation analysis with multiple mediators. *Epidemiologic Methods*, 2(1), 95–115.
- Working, E. J. (1927). What do statistical “demand curves” show? *Quarterly Journal of Economics*, 41(2), 212–235.
- Zhou, X. (2022). Semiparametric estimation for causal mediation analysis with multiple causally ordered mediators. *Journal of the Royal Statistical Society B*, 84(3), 794–821.