Decomposition Methods in the Social Sciences

GESIS Training Course

January 29 – February 1, 2024, Cologne

Johannes Giesecke (Humboldt University Berlin)
Ben Jann (University of Bern)

1. Introduction and course overview

Contents

Introduction

- Decomposition methods are used to analyze **distributional differences** in an outcome variable **between groups** or time points.
- In particular, the methods decompose the observed difference into a component that is due to **compositional differences** between the groups, and a component that is due to **differential mechanisms**.
- Example question: How can the difference in average wages between men and women (the gender wage gap) be explained? Is the difference due to ...
  - ...group differences in wage determinants (i.e. in characteristics that are relevant for wages, such as education)? (compositional differences)
  - ...differential compensation for these determinants (e.g. different returns to education for men an women, or wage discrimination against women)? (differential mechanisms)

Introduction

- Similar questions can be asked in **other contexts** (for different groups, for different outcome variables) or also when analyzing changes **over time**.
- Example question: How can the increase of earnings inequality (e.g. measured by the Gini coefficient or the D9/D1 ratio) be explained? Is the increase due to ...
  - ...changes in the distribution of characteristics that determine earnings? (changes in composition)
  - ...changes in how these characteristics affect earnings? (changes in mechanisms)

Introduction

- Conceptually, decomposition methods are closely related to the **counterfactual** model of causality (although results from decomposition analyses are often not interpreted causally).
- Hence, estimation techniques from the causal inference literature (e.g. matching or inverse probability weighting) can sometimes be useful for decomposition analyses.
- In essence, decomposition methods work by creating counterfactuals such as "How would the outcome distribution in group A look like if it had the same distribution of determinants as group B?"
- Example questions:
  - How high would the mortality rate in country A be if it had the demographic composition of country B?
  - How do average test scores between different schools compare after taking into account the socio-economic composition of the schools' pupils?

Historic development

- Decomposition methods have their origins in the seminal works of Oaxaca (1973) and Blinder (1973), who analyzed mean wage differences between groups (males vs. females, whites vs. blacks).
  - An even earlier reference is Winsborough and Dickenson (1971).
  - Similar methods have also been developed in other disciplines (for example, direct and indirect standardization in demography/epidemiology; see Kitagawa 1955, Das Gupta 1978).
- Pronounced increase in US earnings inequality since the end of the 1970s fostered various methodological innovations in labor economics since the mid 1990s.
- These more recent developments focus on topics such as ...
  - ... providing methods for distributional measures other than the mean
  - ... taking account of non-linearities and providing methods for categorical and other types of variables
  - ... clarifying the basic assumptions made by these procedures
  - ... solving statistical inference

Example

Blau, Francine D., Lawrence M. Kahn (2017). The Gender Wage Gap: Extent, Trends, and Explanations. Journal of Economic Literature 55(3):789–865.

TABLE 4
Decomposition of Gender Wage Gap, 1980 and 2010 (PSID)

	1980		2010	
	Effect of gender gap in explanatory variables		Effect of gender gap in explanatory variables	
Variables	log points	Percent of gender gap explained	log points	Percent of gender gap explained
Panel A. Human-capital specification				
Education variables	0.0129	2.7	–0.0185	–7.9
Experience variables	0.1141	23.9	0.0370	15.9
Region variables	0.0019	0.4	0.0003	0.1
Race variables	0.0076	1.6	0.0153	6.6
Total explained	0.1365	28.6	0.0342	14.8
Total unexplained gap	0.3405	71.4	0.1972	85.2
Total pay gap	0.4770	100.0	0.2314	100.0
Panel B. Full specification				
Education variables	0.0123	2.6	–0.0137	–5.9
Experience variables	0.1005	21.1	0.0325	14.1
Region variables	0.0001	0.0	0.0008	0.3
Race variables	0.0067	1.4	0.0099	4.3
Unionization	0.0298	6.2	–0.0030	–1.3
Industry variables	0.0457	9.6	0.0407	17.6
Occupation variables	0.0509	10.7	0.0762	32.9
Total explained	0.2459	51.5	0.1434	62.0
Total unexplained gap	0.2312	48.5	0.0880	38.0
Total pay gap	0.4770	100.0	0.2314	100.0

Notes: Sample includes full time nonfarm wage and salary workers ages 25–64 with at least twenty-six weeks of employment. Entries are the male–female differential in the indicated variables multiplied by the current year male log wage coefficients for the corresponding variables. The total unexplained gap is the mean female residual from the male log wage equation.

This course

- This seminar gives a detailed introduction to the basic Oaxaca-Blinder (OB) decomposition and also covers some of the newer developments.
- We will focus on **counterfactual decompositions** in the sense described above. There are also other types of “decompositions”, such as the factor decomposition of inequality measures, that will not be covered.

Course overview

- Conceptually, the course is divided into two thematic blocks.
- Block I: Oaxaca-Blinder decomposition
  - Basics; post-estimation, tables and graphs; problems such as the index problem, arbitrary transformations, and the base category problem; extensions to non-linear models; difference-in-difference decompositions
- Block II: Beyond the mean
  - Reweighting, RIF regression approach, Juhn-Murphy-Pierce 1993, approaches based on quantile regression or distributional regression
- Each block contains cycles of theoretical inputs, examples, and exercises (using Stata).

Program

Day 1 Introduction
The basic Oaxaca-Blinder decomposition

Day 2 Problems and solutions in the OB decomposition
Decomposition methods for nonlinear models

Day 3 Difference-in-differences decompositions
Decompositions based on reweighting

Day 4 Decompositions based on recentered influence functions
Decompositions based on quantile or distribution regression

Class materials / readings

- A comprehensive (but for applied researchers sometimes not very accessible) review is provided by:
  - Fortin, Nicole, Thomas Lemieux, Sergio Firpo (2011). Decomposition Methods in Economics. Pp. 1–102 in: O. Ashenfelter and D. Card (eds.). Handbook of Labor Economics. Amsterdam: Elsevier.
- For a concise and easy to understand overview of the basic Oaxaca-Blinder decomposition, see:
  - Jann, Ben (2008). The Blinder-Oaxaca decomposition for linear regression models. The Stata Journal 8(4):453–479.
- Apart from that: read the specialized literature on the various decomposition procedures. References will be provided throughout the slides.
- Slides, literature, and solutions to exercises will be provided on Ilias.

General idea of decomposition methods

- There are two groups, $g = \{0, 1\}$, e.g. males and females. Variable G denotes group membership, e.g. $G = 0$ if male and $G = 1$ if female.
- Of interest is the overall difference between the groups with respect to some functional ν of Y (e.g. the mean, variance, or Gini coefficient):

  $$\Delta^{\nu} = \nu\big(F_{Y|G=0}\big) - \nu\big(F_{Y|G=1}\big)$$

  where $F_{Y|G=g}$ is shorthand notation for $F_{Y|G}(y|g)$.
- Assume that Y is determined by covariates X and an error term ϵ:

  $$Y = m^{g}(X, \epsilon)$$

  where $m^{g}()$ is some group-specific structural function. For example, $m()$ can be a linear function as in linear regression:

  $$m(X, \epsilon) = \beta_0 + \beta_1 X_1 + \cdots + \beta_K X_K + \epsilon$$

  However, $m()$ can also be something much more complicated.

General idea of decomposition methods

- Furthermore, assume (for now) that there are no distributional differences in ϵ between the groups. The overall difference Δ^ν can then be due to:
  - group differences in the distribution of X
  - group differences in structural function $m()$
- The goal of decomposition methods now is to partition the overall difference into these components:

$$\Delta^\nu = \Delta^\nu_X + \Delta^\nu_S$$

where Δ^ν_X denotes the component due to differences in the distribution of X and Δ^ν_S denotes the component due to differences in $m()$.

General idea of decomposition methods

- Potential outcomes (note the close relation to the counterfactual model of causality):
  - given are the *potential* outcomes $Y^0 = m^0(X, \epsilon)$ and $Y^1 = m^1(X, \epsilon)$
  - for the observed Y we have

$$Y = \begin{cases} Y^0 & \text{if } G = 0 \\ Y^1 & \text{if } G = 1 \end{cases}$$

  - that is, potential outcome Y^0 is observed in group 0 and potential outcome Y^1 is observed in group 1
  - for decompositions, however, we are also interested in the (unobserved) counterfactuals
    - How would Y look like in group 0 if it were generated by $m^1()$ instead of $m^0()$ (and vice versa)?
    - Example: How much would a man with given characteristics earn if he were paid like a women?

General idea of decomposition methods

- That is, to identify Δ^ν_X and Δ^ν_S, we need a counterfactual distribution $F_{Y^g|G\neq g}$. We use red color to emphasize counterfactuals.
- For example, let $F_{Y^0|G=1}$ be the counterfactual distribution of Y in group 1 if we assume that Y is determined in group 1 according to group 0's structural function $m^0()$.
- By adding and subtracting $\nu\big(F_{Y^0|G=1}\big)$, the decomposition can then be written as:

$$
\begin{aligned}
\Delta^\nu &= \nu\big(F_{Y|G=0}\big) - \nu\big(F_{Y|G=1}\big) \\
&= \big\{\nu\big(F_{Y|G=0}\big) - \nu\big(F_{Y^0|G=1}\big)\big\} + \big\{\nu\big(F_{Y^0|G=1}\big) - \nu\big(F_{Y|G=1}\big)\big\} \\
&= \Delta^\nu_X + \Delta^\nu_S
\end{aligned}
$$

General idea of decomposition methods

- The main goal of different decomposition methods is to find good ways to determine $\nu\left(F_{Y^0|G=1}\right)$ or, likewise, $\nu\left(F_{Y^1|G=0}\right)$, so that Δ_X^ν and Δ_S^ν can be estimated.
- Depending on context, further distinctions are made, e.g. by allowing the distribution of ϵ to differ by groups (in which case the decomposition has a third component Δ_ϵ^ν) or by partitioning $m()$ into a part related to X and a part related to ϵ (to subdivide Δ_S^ν into a component related to observables and a component related to unobservables).
- Similarly, detailed decompositions that separate the contributions of the different X variables are often of interest.
  - Example: how much of the gender wage gap can be explained by differences in education, how much by differences in work experience?

Distributions of random variables

- Random variables X, Y, Z ...
- The variables can be:
  - continuous: any value within a given interval is possible
  - discrete: only a fixed set of values is possible
- Discrete variables are often categorical in the sense that distances between values have no meaning (e.g. group membership).
- Distribution function $F_Y(y)$ (CDF)
  - displays the probability that Y is smaller than or equal to some given value y

$$F_Y(y) = \Pr(Y \leq y)$$

- Density function $f_Y(y)$ (PDF)
  - displays how the probability mass is distributed along the values of Y
  - for continuous variables, the PDF is defined as

$$f_Y(y) = F'_Y(y) \quad \text{such that} \quad F_Y(y) = \int_{-\infty}^{y} f_Y(z)\, dz$$

Distributions of random variables

- the integral of the PDF between two points is equal to the probability mass between these two points

$$\Pr(a \leq Y \leq b) = \int_a^b f_Y(y)\,dy$$

- for discrete variables with possible values $y_1, y_2, \ldots, y_J$, $f_Y(y)$ is a probability mass function

$$f_Y(y) = \Pr(Y = y) \quad \text{such that} \quad F_Y(y) = \sum_{y_j \leq y} \Pr(Y = y_j)$$

- below we will sometimes use integrals of $f_Y(y)$ even if Y is a discrete variable; this is an abuse of notation for sake of simplicity; think of the integral being a sum in these cases

- Quantile function $Q_Y(p)$
  - is equal to the value of Y for which $\Pr(Y \leq y)$ is equal to p
  - $Q_Y(p)$ is the inverse of the distribution function: $Q_Y(p) = F_Y^{-1}(p)$

Example: the normal distribution

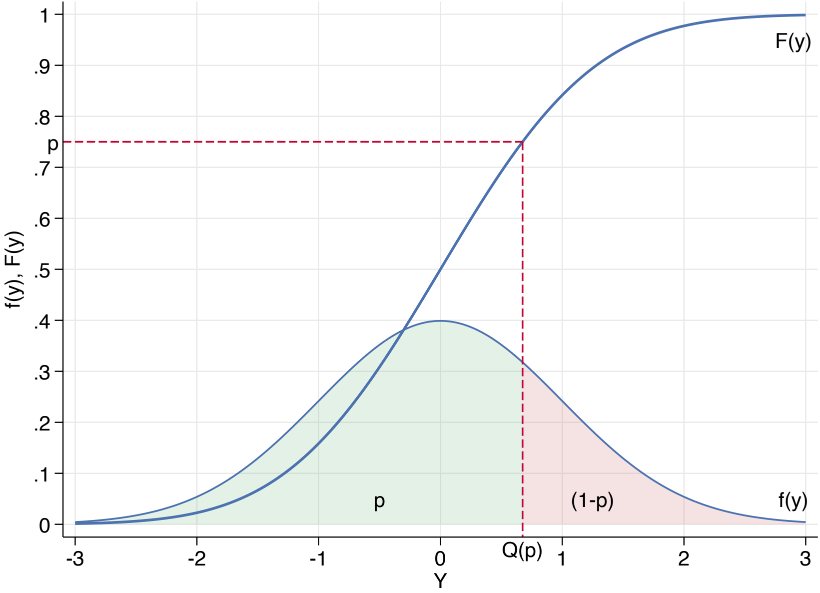

Example: the normal distribution

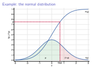

```
local p = .75
local Q75 = invnormal(`p')
twoway (function normalden(x), range(-3 `Q75') psty(p3) recast(area) lcolor(%0) fcolor(%20)) //
       (function normalden(x), range(`Q75' 3) psty(p4) recast(area) lcolor(%0) fcolor(%20)) ///
       (function normalden(x), range(-3 3) psty(p1)) ///
       (function normal(x), range(-3 3) psty(p1) lw(*1.5)) ///
       (pci `p' -3.1 `p' `Q75'  `p' `Q75' -.025 `Q75', lsty(xyline) lp(-)) ///
    , xlabel(#10) ylabel(#10) yti("f(y), F(y)") legend(off) plotr(margin(zero)) ///
    xti(Y) xscale(range(-3.1 3.1)) yscale(range(-.025 1.025)) ///
    ymlabel(`p' "p", add notick nogrid labsize(medsmall)) ///
    xmlabel(`Q75' "Q(p)", add notick nogrid labsize(medsmall)) ///
    text(.05 2.9 "f(y)" .95 2.9 "F(y)" 0.05 -0.5 "p" 0.05 1.25 "(1-p)")
```

Example: quantile function of the normal distribution

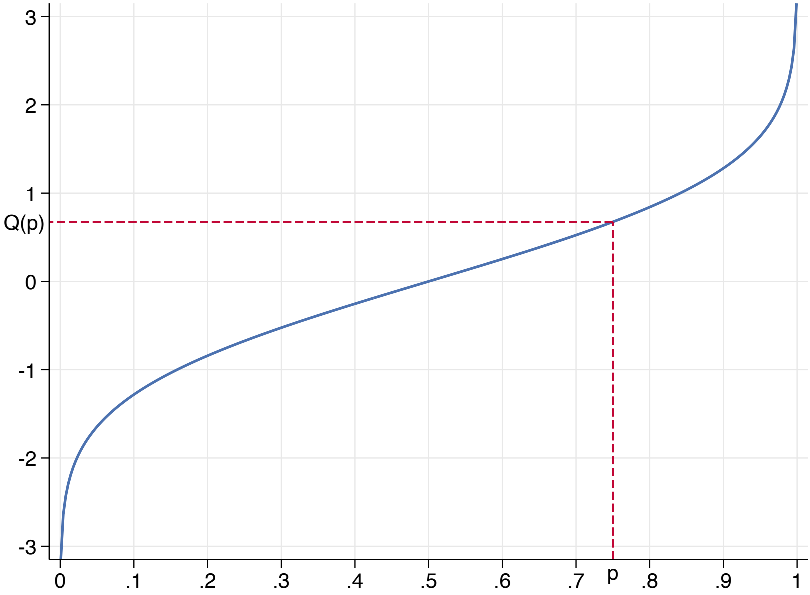

Example: quantile function of the normal distribution

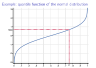

```
local p = .75
local Q75 = invnormal(`p')
local ll = 1-normal(3.15)
local ul = normal(3.15)
twoway (function invnormal(x), range(`ll' `ul') psty(p1) lw(*1.5)) ///
       (pci -3.15 `p' `Q75' `p'  `Q75' `p' `Q75' -.015, lsty(xyline) lp(-)) ///
    , xlabel(#10) ylabel(#10) yti("") legend(off) plotr(margin(zero)) ///
    xti("") yscale(range(-3.15 3.15)) xscale(range(-.015 1.015)) ///
    xmlabel(`p' "p", add notick nogrid labsize(medsmall)) ///
    ymlabel(`Q75' "Q(p)", add notick nogrid labsize(medsmall))
```

Expected value

- Expected value of Y ("the mean")
  - if Y is continuous
    $$E(Y) = \int_{-\infty}^{\infty} y\, f_Y(y)\, dy$$
  - if Y is discrete with possible values $y_1, \ldots, y_J$
    $$E(Y) = \sum_{j=1}^{J} \Pr(Y = y_i) y_i$$
- Some useful relations (see, e.g., Mood et al. 1974)
  - $E(a + bY) = a + b\,E(Y)$
  - $E(X + Y) = E(X) + E(Y)$

Program for today

- The Oaxaca-Blinder decomposition
  - Basic mechanics
  - Estimation
  - Standard errors
  - The detailed decomposition
  - Example analysis
- Exercise 1
- Post-estimation
  - Hypothesis tests
  - Linear and nonlinear combinations
  - Tables and graphs
- Exercise 2

References

Blinder, Alan S. (1973). Wage Discrimination: Reduced Form and Structural Estimates. The Journal of Human Resources 8(4):436–455.

Das Gupta, Prithwis (1978). A General Method of Decomposing a Difference Between Two Rates into Several Components. Demography 15(1):99–112.

Fortin, Nicole, Thomas Lemieux, Sergio Firpo (2011). Decomposition Methods in Economics. Pp. 1–102 in: O. Ashenfelter and D. Card (eds.). Handbook of Labor Economics. Amsterdam: Elsevier.

Jann, Ben (2008). The Blinder-Oaxaca decomposition for linear regression models. The Stata Journal 8(4):453–479.

Blau, Francine D., Lawrence M. Kahn (2017). The Gender Wage Gap: Extent, Trends, and Explanations. Journal of Economic Literature 55(3):789–865.

Kitagawa, Evelyn M. (1955). Components of a Difference Between Two Rates. Journal of the American Statistical Association 50(272): 1168–1194.

Mood, Alexander M., Franklin A. Graybill, Duane C. Boes (1974). Introduction to the Theory of Statistics. New York: McGraw-Hill.

Oaxaca, Ronald (1973). Male-Female Wage Differentials in Urban Labor Markets. International Economic Review 14(3):693–709.

Winsborough, H. H., Peter Dickinson (1971). Components of Negro-White Income Differences. Pp. 6–8 in: Proceedings of the Social Statistics Section. Washington, DC: American Statistical Association.