

WRITTEN EXAM WINTER 2020-21

ECONOMICS OF GENDER

SUGGESTED ANSWERS

This document provides an outline of suggested answers to the exam. The solutions are a guide to answering the questions, and they are not meant as exhaustive. The written solutions would be worked out more completely.

To receive the top grade, the student must with no or only very few minor weaknesses be able to demonstrate an excellent performance displaying a high level of command of all aspects of the relevant material.

The very good student demonstrates a deep understanding of the course material, and is able to connect, combine or adapt general ideas and concepts to specific problems under consideration. The student is able to refer to and extract relevant points from the academic papers in the syllabus, and relate them to the problem under consideration.

For successful completion of the course, the student demonstrates a reasonable (textbook level) insight into the course material, presents a discussion of these issues in a fairly clear and organized way, and relates presented problems to empirical findings.

THE GENDER WAGE GAP AND ITS DEVELOPMENT

“When a man earns \$100, a woman earns \$77 in the U.S., \$78.5 in Germany, \$79 in the UK, and \$83.8, on average, across EU countries” (Bennedsen et al., 2020: 2).

The gender wage gap has been and remains subject to large scientific and public debate. Both its origins, its development over time and across settings, as well as whether there is any need for policy action (and if so of what type) are central topics in labor economics.

Question 1

Describe which *supply and demand side* factors have *traditionally* been considered when examining the gender wage gap in the US and other developed countries. Shortly explain the economic reasoning for how those factors matter for the gender pay gap.

Suggested answer

Relevant references for the answer to this question are BW (2018), chapters 7-10 and course readings from part 3.2. Traditional analyses of the gender wage gap have primarily focused on the impact of economic variables and differences in those across gender: human capital, experience (and training), occupational choice and labor market discrimination. Other factors that may be mentioned are gender differences in unionization, non-standard employment, and institutional/legal factors (such as marriage bars) that all may contribute to gender differences in wages.

- The gender wage gap has traditionally mainly been related to gender differences in human capital. Students should outline, based on course discussion, how according to this reasoning, gender differences in human capital arise from differential investment decisions across m/f, who make those investment decisions as a function of total costs and expected returns. Human capital theory has focused on the fact that given differential expectations on m/f labor market careers (and thus expectations about future returns), women on average invest less in human capital. Furthermore, students should discuss how factoring in the choices beyond length of education (type of education), gender gaps in choice of field have been documented (related to expectations about future employment). These differences may also contribute to an observed gender wage gap.
- Differences in work experience (and related to this labor market training) is another factor that has been shown to contribute to the gender wage gap. Females typically

have more interrupted work careers resulting in less general labor market experience and tenure (at the same employer) relative to men. Related to the human capital channel described above, expecting these differences may impact incentives to invest in labor market (and firm-specific) training and in turn m/f wages.

- Occupational segregation has been documented in many contexts and refers to the concentration of m/f workers in different occupations. Typical findings include that women are more represented in service and professional occupations, while men dominate blue collar occupations. Moreover, women tend to be concentrated in occupations that require less education and/or stable career trajectories, indicating that women may select into occupations that offer temporal flexibility but on average pay lower wages (compensating differentials). Thus occupational characteristics related to occupational segregation, contributing to the gender wage gap. While gender segregation by occupation has been declining in developed countries, still today segregation by occupation is an important factor when studying the gender wage gap.
- Turning to the demand side of the labor market, labor market discrimination has been studied as one main explanation for the gender wage gap. Lm discrimination refers to the situation where equally qualified m/f are treated differentially on the basis of their gender (which we think of in this case as a trait unrelated to productive skills). Thus lm discrimination can contribute to the gender wage gap. In analyses of the gender wage gap, both discrimination resulting from preferences (taste-based discrimination) and expectations about future productivity (statistical discrimination) are considered and are often hard to distinguish empirically. Students should briefly highlight main features of these explanations as considered in the course: Explanations focusing on employer taste-based discrimination highlight that discriminatory employers paying lower wages to female employees than equally productive male employees may contribute to the gender wage gap (in the basic model this result depends on the strength of employers' discriminatory taste, the share of discriminatory employers among all employers, and the number of women seeking employment). Approaches focusing on statistical discrimination suggest that the gender wage gap may partly be a result of employers making decisions based on imprecise signals about individual (f/m) workers' productivity and therefore rely on easily-observable traits that—according to their experience—are predictive of average group productivity. Thus statistical discrimination can contribute to the gender wage gap through perceptions on average differences in productivity across m/f. (One consequence of this type of discrimination may be that females invest less in education/training given that they expect low returns to these investments, turning statistical discrimination into a “self-fulfilling prophecy”.)

Question 2

One way of analyzing the determinants of the gender wage gap is the Oaxaca Blinder decomposition (OB).

- (a) Describe its formal calculation and the intuition behind this approach.
- (b) What have studies based on OB decompositions concluded about the development of the US gender wage gap since the 1980s? What do those studies conclude about the importance of different factors for the gender wage gap then and now?
- (c) Shortly discuss what the OB composition allows us to conclude about the role of discrimination in explaining wage gaps in countries such as the US and Denmark.

Suggested answer

(a) See lecture 10 material. The OB composition separates differences in *mean* wages across groups into explained (the impact of differences in characteristics) and unexplained (the impact of differences in returns to those characteristics) components. Intuitively the approach asks “how much of the average wage gap can be explained away using a set of observable differences across m/f”. The OB composition makes use of well-known properties of OLS, stating that the regression coefficients in a multiple regression analysis capture the effect of a change in the mean of a given X on the unconditional mean of Y. Thus the coefficient on X (let’s say education) in a wage regression can be interpreted as “if the average length of education increases with two years, by how much would average wages increase?”. In the OB decomposition we make use of this property and that we can look at linear expectation functions to analyze mean differences. Importantly, (being sequential in nature), across studies that use OB, the order/choice of characteristics matters for conclusions on “how much of the wage gap can be explained”. In other words, an important point is that the specific characteristics included in the analysis shape our conclusions about the share of explained variation across studies. Decomposing the gender wage gap we can start by writing the wages of males and females as

$$\bar{W}_m = \alpha_m + \bar{X}_m \beta_m \quad (1)$$

$$\bar{W}_f = \alpha_f + \bar{X}_f \beta_f \quad (2)$$

The \bar{X} s are average male/female characteristics included (such as education, age, occupa-

tion). We can then write the gender gap as:

$$\Delta^w = \alpha_m + \beta_m \bar{X}_m - \alpha_f - \beta_f \bar{X}_f \quad (3)$$

Rewrite (add and subtract $\beta_f \bar{X}_m$):

$$\Delta^w = \alpha_m + \beta_m \bar{X}_m - \alpha_f - \beta_f \bar{X}_f + \beta_f \bar{X}_m - \beta_f \bar{X}_m \quad (4)$$

Collect terms:

$$\Delta^w = (\alpha_m - \alpha_f) + \bar{X}_m(\beta_m - \beta_f) + \beta_f(\bar{X}_m - \bar{X}_f) \quad (5)$$

First 2 terms: differences in α, β s: “unexplained” part of the wage gap (the differential return to the X’s for m/f); last term: differences in the included characteristics (Xs, averages): “explained” part of the wage gap. Students should demonstrate their understanding of the approach by commenting on the fact that—apart from the included X’s—our way of constructing the decomposition (choosing the weighting scheme) also impacts the numerical results.

(b) Students should summarize the evidence on the US gender wage gap, see lecture 10 and reference to Blau and Kahn (2017) and Blau and Winkler (2018), chapters 7 and 10; Since the 1980s, the gender wage gap has been narrowing in the US. Research points to three broad factors in explaining this trend: 1) changes in human capital and other qualifications; 2) changes in the extent of labor market discrimination against women (role of equal pay policies); and 3) changes in the wage structure/differential returns in different occupations.

Using OB decompositions Blau and Kahn point to (i) changes in the importance of different factors explaining the wage gap across m/f, and (ii) changes in the size of the unexplained part of the gap: While differences in human capital and experience across m/f could explain a significant share of the gender wage gap in the early 1980s (around 25pct), increases in women’s lm experience and educational attainment diminished the importance of these factors for explaining remaining wage gaps by 2010 (women still have slightly less experience on average but outperform men wrt education). On the contrary, female occupational/industry status plays an important role in explaining the remaining gender wage gap in 2010, i.e. has gained relative importance as an explanatory factor. This fact is closely linked to changes in the wage structure (favoring male dominated occupations). Finally, while the wage gap has narrowed in the period, as has the unexplained part of the wage gap (partly due to anti-discrimination efforts or changes in unmeasured skills or the composition of the labor force), the unexplained gap in the decomposition still remains important in 2010.

(c) The OB composition is a useful tool to decompose *mean* differences across groups. Often times the “unexplained part” of the variation is termed as the impact of discrimination in

the public debate. However, this conclusion is not an adequate label: First, the unexplained part contains potentially omitted characteristic that relate to productivity and thus wages, such as motivation or willingness to compete. It is rarely ever possible to include all relevant characteristics in a study of wages. Second, the role of pre-labor market discrimination may bias our results as it may impact observables that we include such as educational attainment or occupation, i.e. make it inadequate to term the unexplained part in the OB decomposition as discrimination.

THE GENDER WAGE GAP TODAY: KIDS AND POLICIES

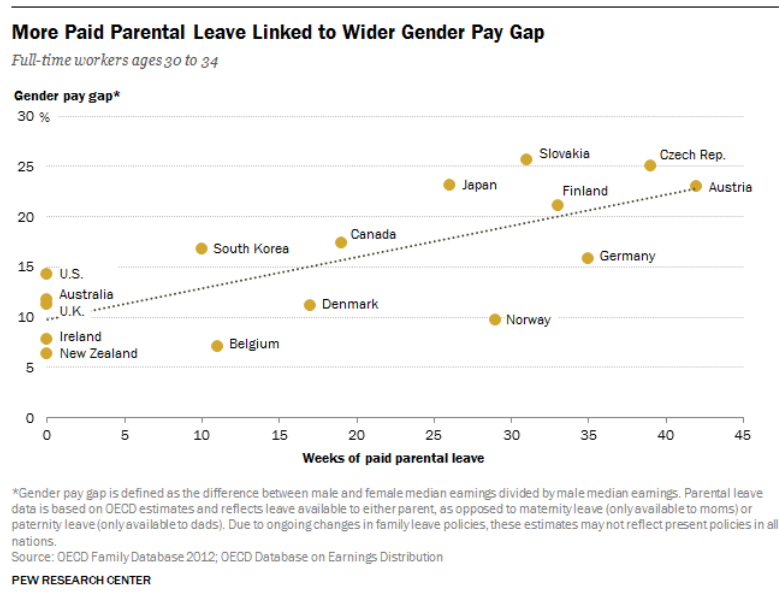
Question 3

In the light of remaining wage inequalities across men and women, a growing literature has focused on the impact of childbearing for female market outcomes and gender inequality.

(a) Shortly explain why ex ante predictions about the impact of *access* to maternity leave programs for mothers' labor market outcomes are ambiguous.

(b) Describe the below graph and its central message. Explain why one should be careful when making causal statements based on the graph.

Figure 1: Paid leave duration and the gender pay gap in the OECD, 2012.



(c) Kleven et al (2019) analyze data from Denmark and conclude that remaining gender inequality in wages is due to the consequences of childbearing for mothers. Why is it difficult to estimate the impact of children on female wages? Describe how Kleven et al (2019) tease out the causal impact of children on female labor market outcomes and gender inequality (what is their empirical approach and its assumptions).

Suggested answer

(a) Students should discuss this point based on course discussions: Access to birth-related leave for mothers may both impact maternal labor supply and wages. On one side, maternity leave may allow new mothers to maintain a connection to employers and thus be able to return to a job with good match quality. This may positively impact both mothers' return decisions and their wages. On the other side, if leave prologues the period that *all* mothers spent outside the labor market, leave will decrease maternal labor supply and potentially impact wages negatively (fx due to a loss of human capital). The impact of maternity leave may thus critically depend on design features such as length, compensation levels, or interactions with other policies.

(a) The graph shows a positive cross-country correlation of the gender pay gap with the number of paid weeks of parental leave provided in the set of OECD countries considered. The graph is based on a measure of the gender gap among full-time workers in a relevant (but selected) age group: around the time where families have their first child. While the graph suggests a positive effect of longer leave options on the gender pay gap in that age group for the given countries, we should be careful with causal statements about this relationship: The graph is based on a selected sample of countries (e.g. omitting countries like Sweden with long leave entitlements but wage gaps at the lower end of the spectrum) and only considers the gender pay gap in a specific age group (at least this fact may constrain our ability to extrapolate the finding). Most importantly, however, there may be other factors that covary with leave entitlements, such as other family-related policies (childcare availability), discrimination of mothers on the labor market, or norms around the adequate time that mothers should spend at home. All of these may impact our conclusions about the causal impact of leave duration on the gender wage gap.

(c) Lecture 15 material and Kleven et al (2019) reading; an important constrain in empirical research on the impact of children on labor market outcomes is selection into motherhood (students should spell concrete issues). Women, who become mothers may be selected on characteristics (such as ability or motivation) that remain unobserved for the researchers but that impact both, their decisions on fertility and their outcomes on the labor market. Empirical studies have confronted this issue in different ways mainly focusing on the inten-

sive margin (studying exogenous changes in family size: students should refer to a literature using sibling sex-composition and twin births as instruments). Kleven et al (2019) focus on changes of lm outcomes around the birth of first children (in the longer run their estimates capture the impact of all children born in a family). They use an event study: This means that they measure labour market outcome over time for prospective parents both prior to and after the birth of a child. Their analyses are based separately on data for m/f who become parents for the first time. Thus they compare outcomes of f/m (such as wages) over time and relative to each group's outcomes at time=-1, i.e., the year before the birth of the child. The analysis then identifies the impact of the first child (and all consecutive children) on outcomes separately for f/m. They quantify the impact of children in terms of percentage changes in outcomes relative to the counterfactual development of outcomes in the situation without a child for each group. For their event graphs to identify the causal impact of children, they make two main assumptions: first, they assume that outcomes would evolve smoothly over time. This assumption means that in the year of childbirth there is no other reason for an abrupt change in lm outcomes. Thus in the absence of the birth, those outcomes (such as earnings) would not change abruptly. Second, they assume that there is arbitrary variation in the timing of the first birth. This assumption means that parents cannot fully predict/plan the timing of their first birth. In a second step they use the estimates for the impact of children on wages for f/m to compute a measure of gender inequality at each time t (to show how much women are penalized relative to men).

Question 4

A recent empirical literature focuses on the role of gender differences in *psychological attributes* for explaining gender wage gaps. These factors—if relevant—may lead to other policy implications than traditionally advocated.

One policy that has gained attention is mandated transparency in job application or work place settings. Bennedsen et al (2020) study the impact of wage transparency rules in private firms in Denmark on the gender wage gap in those firms.

(a) What main challenges do Bennedsen et al (2020) confront when attempting to identify the impact of wage transparency policies on the gender wage gap? Explain how the authors address these challenges in their (main) design and choice of sample. (You may think about the “ideal experiment” and why a comparison of outcomes across firms that choose/do not choose to have transparency rules likely not reflects the effect of transparency rules.)

(b) Table 3 presents the main finding in Bennedsen et al (2020) for the impact of wage transparency rules in firms on the gender wage gap. What is the main conclusion from the table? Which challenge does the estimation presented in column 3 address?

Table 3: Gender Pay Gap Disclosure and Employee Wages

This table reports the effects of gender pay gap disclosure on employee wages. Each column presents coefficients from a different multivariate regression. Columns 1 and 5 (2 and 6) are estimated using only male (female) employees. All other columns use the entire sample. The dependent variable is employee annual wage (log-transformed). *Treated* is a dummy that takes the value 1 for individuals working in firms with 35-50 employees before the introduction of the law and 0 for employees in firms with 20-34 employees. *Post* takes a value of 1 for 2006, 2007, and 2008, and a value of 0 for years 2003, 2004, and 2005. Individual controls include employee work experience and age. Firm controls include sales (log-transformed). Detailed descriptions of the variables are given in Table A1. Standard errors are double clustered at the individual and the firm level and reported in parentheses. ***, **, and * correspond to statistical significance at the 1%, 5%, and 10% levels, respectively.

	Male	Female	All	All	Male	Female	All	All
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated x Post	-0.0167*** (0.0039)	0.0028 (0.0045)	0.0028 (0.0044)		-0.0152*** (0.0037)	0.0036 (0.0044)	0.0036 (0.0043)	
Male x Post			-0.0022 (0.0034)	-0.0041 (0.0036)			-0.0033 (0.0034)	-0.0046 (0.0035)
Treated x Post x Male			-0.0195*** (0.0052)	-0.0148*** (0.0053)			-0.0190*** (0.0052)	-0.0144*** (0.0053)
Person-Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	Yes	Yes	Yes	No
Firm x Year FE	No	No	No	Yes	No	No	No	Yes
Firm Controls	No	No	No	No	Yes	Yes	Yes	No
Individual Controls	No	No	No	No	Yes	Yes	Yes	Yes
R^2	0.868	0.827	0.866	0.884	0.868	0.826	0.866	0.883
N	145,852	79,532	225,384	223,624	145,262	79,027	224,289	222,529

Source: Bennedsen et al (2020)

(c) Transparency policies may impact how firms/managers behave, but also impact how female and male applicants and employees behave. As Bennedsen et al describe, this impact may be due to “differences in risk aversion or bargaining power” (p.8).

Describe based on course readings what empirical studies conclude about *gender differences in psychological attributes*. Make explicit references to at least two relevant studies and their main findings. Critically discuss whether and, if yes, how those findings may help us to understand remaining gender wage gaps in developed countries and what policy proposals (such as transparency rules) may be relevant?

Suggested answer

(a) Bennedsen et al state that empirical analyses (such as their own) face two main challenges: first, lack of (panel) data on f/m employees’ wages, and second, endogeneity concerns when comparing firms that do or do not introduce transparency rules (students should spell out the issues rather than just mentioning the need for exogenous variation): In an observational setting with transparency rules being implemented by some firms, a comparison of

employee and firm outcomes across treated/control firms is likely biased. This bias comes from underlying differences across firms that choose to implement transparency vs firms that choose not to (Differences in characteristics of the firm and employees or other policies in the firm that are not observed to the researcher but are correlated to both, wage transparency in the firm and outcomes of employees, introduce omitted variable bias).

An ideal experiment would allocate mandatory transparency rules at random among firms and track wage development for employees in those firms, as well as hiring/promotion/quitting of employees over time. Bennedsen et al (2020) exploit a natural experiment in Denmark, which mimics this experiment: DK introduced mandatory wage transparency rules in firms with more than 35 employees. The authors exploit the introduction of transparency rules in a sample of DK firms that are just below or above the described and arbitrary cutoff (of 35 employees) where the transparency rules bite. In that way, they make more plausible that they examine the development of the wage gap in comparable firms (they argue that they have a good counterfactual for treated firms).

The Danish setting also allows the authors to also confront the data issue described above because they can track employees' wages (and other outcomes) over time. Additionally, the data provides the authors with firm identifiers (as well as firm-level characteristics and outcome measures) and a firm-employee match.

Thus in their main empirical strategy the authors compare the development of employee outcomes (e.g., wages) over time (pre and post law) across treated and control firms (defined by firm size) in a DiD framework.

(b) The main findings in the table are based on DiD estimations as described above (that should be explained and not just mentioned). Column 3 presents the main “triple difference” results that shows that male wage growth is -1.95pp lower than female wage growth in treated firms. The triple difference strategy addresses the following challenge: In a DiD analysis we are concerned about a main problem—differential trends of outcomes (wages) across small/large (control/treated) firms. Thus comparing outcomes across treated/control firms over time could lead to biased results if we focused on all employees of smaller/larger firms. By adding a comparison of m/f wages within treated/control firms for the pre/post period, confounding factors would have to impact m/f wage development *differentially within the same firm* in order to bias the results in the study.

Results: Columns (1) and (2) present DiD results comparing wages in treated/control firms prior to and after the policy introduction, separately. Column (3) combines the data for m/f and adds a third difference: The separate m/f analyses show that wage growth for male employees is significantly lower in treated than in control firms (1.67pp difference).

For female employees, the introduction of wage transparency rules has no significant effect on wage growth and the relevant the point estimate is positive but very small (0.28pp). Column (3) shows that male wage growth is -1.95pp lower than female wage growth in treated firms. Thus the results suggest that transparency rules predominantly impacted male wage growth negatively. The remaining columns of Table 3 (columns 5-8) test the robustness of the main finding of slower male wage growth to inclusion of control variables (at firm and individual level) and conclude that the main message of the analysis remains unchanged.

(c) Student answers to this question will vary depending on students' choice of articles from the syllabus (and they will draw on lecture 12 and lecture 21). We have discussed a number of studies looking into gender differences in psychological attributes and the relevance of those differences for explaining lm outcomes (such as the gender wage gap) in the course. Examples include studies on attitudes towards competition and negotiation, risk attitudes, or social preferences. Importantly, students should assess the relevance of these studies for understanding the gender wage gap and critically discuss implications for policies:

Examples for relevant studies include: Students may refer to Leibbrandt and List (2015) on gender differences in negotiating behavior and the transparency of negotiations. The results of this study suggest that transparency in job ads (through its impact on negotiation behavior of w/m) may help address remaining gender pay differences. Niederle and Vesterlund (2007) study gender differences in competition and find that qualified women (in a lab experiment) enter into competitions at lower rates than men. This finding may help explain gender pay gaps that may arise through selection into competitive (and highly rewarding) environments. In a follow-up study based on lab experiments, Niederle, Segal and Vesterlund (2013) find that affirmative action may be one potential way to address suboptimal entry decisions of qualified females into competitive environments. Students could also summarize findings presented in the Handbook chapter by Bertrand (2010) that reviews a literature on gender differences in psychological attributes, mainly from lab experiments but also field work.

Do these findings help us understand gender differences in labor market outcomes? Students should **discuss** this question drawing on course material. As concluded by Bertrand (2010), while an increasing number of predominantly lab studies has shown gender differences in psychological attributes, there is much less work documenting that those differences matter "in the real world". Additionally, students should refer to the discussion in Nelson (2015) and the Bertrand Ely lecture (covered in the final lecture of the course) of across group vs within group differences (gender similarity argument). This view on gender differences makes a different point: maybe gender differences in psychological attributes are not as

large and important as we tend to believe. Nelson presents evidence to suggest that within group differences tend to be larger than between group differences, and she argues that publication bias may confound our understanding of the importance and size of gender differences in psychological attributes. If this is the case, policies that are motivated by (potential small) between group gender differences may not effectively tackle the underlying causes for differences in lm outcomes. On the contrary, persistent beliefs about and focus on differences in gender gaps in psychological attributes may lead to gender stereotypes, as discussed in the Bertrand HB chapter and the Bertrand Ely lecture.

If gender differences in psychological attributes are influential for remaining gender gaps on the labor market, an important question with respect to policies relates to their **determinants**: Are they biologically determined? Are they a result of our upbringing in the family? Are they shaped by norms? Do they emerge because of prescriptive gender stereotypes (women as pro-social and caring, men as competitive)? In the latter case, policies that address those stereotypes (such as counter-stereotypical behavior by role models (as in Scott, Page and West (2010), who show that access to role models has the potential to change “gendered choices” during education) may be relevant and contribute to narrowing gender gaps on the labor market.