Labour Economics 30 (2014) 32–40

Contents lists available at ScienceDirect

Labour Economics

journal homepage: www.elsevier.com/locate/labeco

Gender and the labor market: What have we learned from field and lab experiments?☆

Ghazala Azmat *, Barbara Petrongolo

Queen Mary, University of London, School of Economics Finance, Mile End Road, London E1 4NS, United Kingdom
Centre for Economic Performance (LSE), United Kingdom

ARTICLE INFO

Article history:
Received 11 January 2014
Received in revised form 29 April 2014
Accepted 30 April 2014
Available online 20 June 2014

JEL classification:
J16
J24
J71
C91
C92
C93

Keywords:
Gender
Field experiments
Lab experiments
Discrimination
Gender preferences

ABSTRACT

We discuss the contribution of the experimental literature to the understanding of both traditional and previously unexplored dimensions of gender differences and discuss their bearings on labor market outcomes. Experiments have offered new findings on gender discrimination, and while they have identified a bias against hiring women in some labor market segments, the discrimination detected in field experiments is less pervasive than that implied by the regression approach. Experiments have also offered new insights into gender differences in preferences: women appear to gain less from negotiation, have lower preferences than men for risk and competition, and may be more sensitive to social cues. These gender differences in preferences also have implications in group settings, whereby the gender composition of a group affects team decisions and performance. Most of the evidence on gender traits comes from the lab, and key open questions remain as to the source of gender preferences—nature versus nurture, or their interaction—and their role, if any, in the workplace.

1. Introduction

Women have made major inroads in labor markets throughout the past century, resulting in clear convergence in human capital investment and employment prospects and outcomes relative to men (Goldin, 2006). However, while the gender gap in schooling has closed—and even reversed—in most rich countries, there are remaining gender differences in pay and employment levels, as well as in the types of activities that men and women perform in the labor market (OECD, 2002). Women's progress in the labor market has also led to major advances in labor economics, reflecting women's changing role in the economy and identifying the factors behind the remaining disparities with respect to men. The development of novel empirical methods to identify gender differences has been accompanied by broadening perspectives on the gender dimensions of interest.

This paper reviews recent advances in the economics of gender that have been achieved via the experimental approach. We discuss how the experimental literature contributes to a deeper understanding of recurrent questions on gender, as well as to the broadening of research questions towards previously unexplored dimensions of gender differences, and we examine their bearings on labor-market outcomes.[1]

The factors driving gender differences in the labor market can be broadly categorized into three forces, which might be interconnected: productivity, preferences and discrimination. By the end of the 1990s, the state-of-the-art work on gender inequalities, summarized in Altonji and Blank's (1999) chapter in the *Handbook of Labor Economics*, had focused mainly on productivity differences related to human-capital accumulation and discrimination as the main sources of gender gaps in wages and hours. At the same time, Altonji and Blank (1999)

☆ Paper prepared for the EALE 25 Memorial Issue of Labour Economics. We thank Andrea Ichino for very useful comments as a discussant.

* Corresponding author at: Queen Mary, University of London, School of Economics and Finance, Mile End Road, London E1 4NS, United Kingdom.
E-mail addresses: g.azmat@qmul.ac.uk (G. Azmat), b.petrongolo@qmul.ac.uk (B. Petrongolo).

[1] A number of recent papers survey these literatures extensively. Bertrand (2011) discusses new advances in the gender literature, based on both experimental and non-experimental evidence, with an emphasis on gender differences in preferences, gender identity, and women's well-being. Croson and Gneezy (2009) cover lab experiments on gender preferences in detail. The growing usage of field and lab experiments in labor economics is discussed by List and Rasul (2011) and Charness and Kuhn (2011), respectively.

http://dx.doi.org/10.1016/j.labeco.2014.06.005
0927-5371/

also indicated that a lack of direct evidence on discrimination and gender preferences was a key obstacle to cleanly differentiate among the effects of these three forces. Progress in these areas was hindered by the presence of unmeasured, confounding factors in the study of gender discrimination; and the difficulty of extracting clean information on psychological traits from naturally-occurring data in the study of gender preferences. By providing data explicitly suited to addressing the questions of interest, and allowing tight control on the environment, the experimental approach provides a valuable source of evidence on these and other gender issues.

The experimental approach, both in the field and the lab, represents the most recent addition to the labor economist's toolkit, in the quest for identification of causal effects of interests. As a clear indication of this trend, two chapters in the latest *Handbook* volume are devoted entirely to discussing lessons drawn from experiments in labor economics (List and Rasul, 2011; Charness and Kuhn, 2011). Various factors have contributed to the increased use of experiments in labor research. First, the growing influence of the behavioral literature on economic research has expanded labor economists' perspectives to largely unexplored fields at the border between economics and psychology (see Bertrand, 2011, for a recent survey on new perspectives on gender disparities). Second, labor economists have become more clearly aware of the limitations of observational data in answering old and new questions and have at the same time set higher standards for empirical inference. Finally, economists have become increasingly able and willing to engage in the collection of original data. Given these advances, the experimental approach has made available new data on traditional labor questions and enabled labor economists to address new questions by "letting questions determine the data to be obtained, instead of the data determining the questions that can be asked" (Duflo, 2006).

Experiments in economics have evolved in a number of directions, and, to organize our discussion, it is helpful to classify empirical strategies in labor economics according to the degree of control allowed to the researcher. At one end of the spectrum, there are the traditionally used, naturally-occurring data, over which the researcher has no control—over the information elicited or the economic environment—and for which identifying assumptions are needed to estimate the causal effect of treatment. At the other end of the spectrum, there are laboratory experiments, which use randomization to identify the effect of treatment in the lab (most typically on a subject pool of students) and which allow the researcher to fully control the environment. Thanks to randomization, the causal effect of treatment is identified simply by the difference in mean treatment and control outcomes. Somewhere between these two extremes are (various typologies of) field experiments, which use randomization in a natural-occurring environment—typically on relevant sample pools that may not be aware of their participation in an experiment—and thus allow for a combination of control and realism (Harrison and List, 2004).

In recent decades, the empirical literature on gender has progressed along this conceptual spectrum. Economists have long been interested in the causes and consequences of gender discrimination in the marketplace. Early work on discrimination extensively used the regression approach and decomposition techniques on observational data. However, the increased awareness of this approach's limitations has gradually shifted the emphasis of empirical work on this topic towards field experiments such as audit and correspondence studies, which aim to compare outcomes in the same job for two individuals who are identical in all respects other than gender. While experiments have been used more extensively to study race, rather than gender, discrimination, the experimental approach has in some cases provided clean evidence on gender discrimination in hiring, and represents a promising path for future research. We discuss findings from this approach in Section 2.

More recently, growing emphasis on potential differences in psychological attributes between men and women has shifted the attention of experimental work towards the study of various dimensions of gender preferences, including preferences towards risk, competition, negotiation and other-regarding preferences (see Croson and Gneezy, 2009, for an exhaustive review of the experimental work on gender preferences). Potential differences in preferences and psychological attributes might offer additional insight into gender gaps in participation to the labor market, in the types of jobs held, and in the performance in a given job. Information on gender preferences is typically elicited in a lab environment, which best isolates one factor of decision, say the attitude towards risk. The recent literature contains numerous examples from this approach, and is discussed in Section 3.

Finally, the emphasis on gender differences at an individual level has led to a recent interest in the role of these differences in collective settings. Higher female representation in high-profile jobs in politics and the corporate sector—partly prompted by regulation such as the introduction of gender quotas in several countries—has led academics and policymakers alike to question the consequences of teams' gender composition for collective decision making. Section 4 links recent empirical evidence from the field and the lab to team work and discusses evidence on the impact of the gender composition of teams on decision making and firm performance.

Section 5 concludes this survey by summarizing the state of the art and discussing current open issues and directions for future research.

2. Discrimination

The study of discrimination, encompassing concepts, measurement and impact, has featured prominently in the gender literature since Becker's (1957) seminal work.[2] Gender discrimination in the labor market is defined as a situation in which equally productive men and women are rewarded differently, making it necessary to correctly measure differences in productivity in order to pin down the discrimination residual. The early literature has used regression-based methods on observational data—typically labor force or household survey data—to test for discrimination in the labor market. The most common approach (Oaxaca, 1973; Blinder, 1973) consists in decomposing wage (or participation) differentials between men and women into an 'explained' gap, driven by gender differences in observable worker and, sometimes, job characteristics; and an 'unexplained' gap, driven by different returns to given characteristics that are in turn associated with discrimination. Results from this literature are summarized by Altonji and Blank (1999), and point to large unexplained gaps in gender wages and participation rates.

While the existence of an unexplained gap in wages is certainly consistent with discrimination, this measure suffers from two main drawbacks. First, most observational data inevitably lack information on some of the determinants of a worker's productivity, which are nevertheless observed and valued by employers. Thus the unexplained gap is contaminated by unobserved differences in productivity, and whether it provides an upward or downward bias of the true extent of discrimination depends on the sign of differences in such unobservables.[3] Second, if pre-labor market investment in human capital is affected by expectations of future discrimination, part of the impact of discrimination is captured by observable productivity differences, and the resulting unexplained gap would underestimate the true extent of discrimination. In the first case, the regression approach would control for 'too little,' while in the second case, it would control for 'too much.'

Experiments are a natural response to some of the weaknesses inherent in the regression approach to discrimination on conventional survey data. As "discrimination is a causal effect defined by a hypothetical ceteris paribus conceptual experiment" (Heckman, 1998), the experimental approach allows researchers to approximate the ceteris paribus condition by comparing outcomes for otherwise identical men

[2] See Altonji and Blank (1999, Sections 3 and 4) and references therein.

[3] For this purpose, employer–employee matched data provide an improvement over survey data, as they allow researchers to better extract information on productivity of individuals, see Hellerstein and Neumark (2006).