Contents lists available at ScienceDirect

European Economic Review

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

Discrimination in the laboratory: A meta-analysis of economics experiments

Tom Lane

Department of Economics, University Park, University of Nottingham, NG7 2 RD, United Kingdom

A B S T R A C T

Economists are increasingly using experiments to study and measure discrimination between groups. In a meta-analysis containing 441 results from 77 studies, we find groups significantly discriminate against each other in roughly a third of cases. Discrimination varies depending upon the type of group identity being studied: it is stronger when identity is artificially induced in the laboratory than when the subject pool is divided by ethnicity or nationality, and higher still when participants are split into socially or geographically distinct groups. In gender discrimination experiments, there is significant favouritism towards the opposite gender. There is evidence for both taste-based and statistical discrimination; tastes drive the general pattern of discrimination against out-groups, but statistical beliefs are found to affect discrimination in specific instances. Relative to all other decision-making contexts, discrimination is much stronger when participants are asked to allocate payoffs between passive in-group and out-group members. Students and non-students appear to discriminate equally. We discuss possible interpretations and implications of our findings.

© 2015 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Meta-analysis – a commonplace technique in medical science, psychology and, to a growing extent, economics – holds advantages over literature review in terms of objectivity and analytical rigour. In recent years, the experimental economics literature appears to have reached a critical mass at which researchers are finding meta-analyses useful.[1] The benefit of these works is that, by aggregating data across a large number of experiments and exploiting natural between-study design variation, they pinpoint behavioural regularities and the variables that modify them more precisely than could be done through qualitative review.

We run a meta-analysis on the body of studies investigating discrimination in lab and lab-in-the-field experiments, a sub-literature which has certainly reached the necessary critical mass for such a venture. Economists' interest in discrimination has been strong ever since Becker (2010), and with the growth of experimental economics in the last two decades, experiments have emerged as a popular complement to survey-based econometric studies.

These experiments create a controlled environment and therefore allow much cleaner measurements of discrimination than the analysis of naturally-occurring data, avoiding such problems as omitted variable bias and reverse causality. Furthermore, by testing for a very fundamental and general form of discrimination – simply, whether subjects treat others differently depending on which group those others belong to – experimental economists can produce findings of interest not only to their own discipline but also across the social sciences. Also, through the use of incentives, experiments hold a key advantage over questionnaire-based measures of discrimination, in that they elicit revealed rather than reported discrimination.

Psychologists had already been studying discrimination in the lab for decades, and experimental economists have drawn on their knowledge, particularly regarding the minimal group paradigm. This technique was first introduced by Tajfel et al. (1971) and has spawned a huge body of experiments wherein group identity is artificially induced in the laboratory. This is often done by, in a preliminary phase of an experiment, asking subjects to state their preference for one artist over another, or to randomly draw a colour. The experimenter then splits the subject pool into groups according to their art preference, or the colour they have drawn, and makes it known to participants that the division is based on these differences. Subsequent stages of such experiments involve interaction tasks between the groups and find discrimination surprisingly (at least to the early researchers) often.

To study discrimination, experimental economists set up games such as the dictator game, the trust game or the prisoner's dilemma, and invite a subject pool segregated along the lines of a particular identity-based characteristic (or else generate this segregation with artificial groups). They make subjects aware of the group affiliation of those they interact with, and then measure how their behaviour varies according to whether individuals they are interacting with share their identity (are in-group) or do not (are out-group).

The number of economics experiments of this type has grown rapidly since the turn of the century and now encompasses substantial diversity across several dimensions. Even after omitting many papers which investigate discrimination but do not meet our inclusion criteria devised to ensure a consistent approach (see Section 2), we are left with a dataset consisting of 441 experimental results (significant and null) from 77 studies – more data than most of the other experimental economics meta-analyses have had. In order to aid the progression of this literature, it is worth taking stock of what has been found to date, particularly as casual inspection reveals non-uniformity in the results; the strength of discrimination found against out-groups varies considerably, and some experiments even find discrimination in the opposite direction, i.e. against the in-group.

The aim of this meta-analysis is both to yield broad insights on discrimination and to inform the designers of future experiments testing for it. We first investigate the average strength of discrimination across the literature. We then inquire how it tends to vary according to specific experimental characteristics.

In particular, we are interested in whether the strength of discrimination depends on the type of identity being investigated. Comparing the level of discrimination between artificial (i.e. minimal) groups and various types of natural groups (such as those based on ethnicity, nationality, religion, gender and social/geographical affiliation) is particularly interesting. One might expect 'minimal' groups to yield minimal levels of discrimination. However, it is also conceivable that artificial identity inducement confers an experimenter demand effect in favour of discrimination, or that the experimental priming of sensitive natural identities reduces subjects' desire to discriminate owing to a preference not to engage in socially unacceptable behaviour. Evidence for these possibilities, in the form of relatively strong discrimination in artificial group experiments, could have implications for the external validity of certain experiments.

A further interesting question is whether the strength of discrimination varies according to the type of decision subjects are asked to make. This has implications in terms of the real-world circumstances in which discrimination can be most expected to appear and for the generalisability of findings.

We further ask whether experiments with students reveal greater or lesser discrimination than those with non-students. This is also important for the external validity of findings, and is a question worth pursuing as some studies (e.g. Bellemare and Kröger, 2007; Anderson et al., 2013) have found students are not entirely representative of wider populations in economics experiments.

This meta-analysis also aims to shed light on the motivations behind discrimination. Some experiments have been designed specifically to distinguish between taste-based discrimination and statistical discrimination – the two models that continue to dominate the theoretical literature in economics. The taste-based model, proposed by Becker (2010), entails individuals gaining direct utility from the act of discriminating against out-groups. Meanwhile, according to theories of statistical discrimination – beginning with Arrow (1972) – individuals aim to maximise their own payoffs given their beliefs and expectations about others' characteristics and behaviour, and discrimination occurs when those beliefs and expectations vary depending on the group to which the others belong. Understanding the relative importance of these two motivations will improve the focus of future research and the design of policies aimed at combating discrimination.

Finally, we include a subsection on experiments investigating gender discrimination. Gender is unique amongst the identity types in having the same two groups in each experiment. It is therefore simple to make a clean comparison between male-to-female discrimination and female-to-male discrimination.

In summary, the meta-analysis presented below aims to address the following questions: (1) What is the general pattern of discrimination across the literature? (2) How does the level of discrimination vary according to the type of identity groups are based upon? (3) How does the level of discrimination depend upon the decision-making context? (4) Do students discriminate any more or less than non-students? (5) Does the experimental literature provide more support for taste-based