

Field Experiments in Labor Economics[☆]

John A. List^{*,1}, Imran Rasul^{**,2}

^{*} Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, USA

^{**} Department of Economics, University College London, Drayton House, 30 Gordon Street, London WC1E 6BT, United Kingdom

Contents

1. Introduction	104
1.1. The experimental approach in science	106
1.1.1. <i>An experimental cornerstone</i>	107
1.1.2. <i>Early labor market field experiments at the Hawthorne plant</i>	108
1.1.3. <i>Large-scale social experiments</i>	111
1.1.4. <i>Potential shortcomings of social experiments</i>	114
1.2. Field experiments	117
1.2.1. <i>What is a field experiment?</i>	118
1.2.2. <i>A more detailed typology of field experiments</i>	120
1.2.3. <i>Simple rules of thumb for experimentation</i>	123
1.2.4. <i>Further considerations</i>	126
1.2.5. <i>Limits of field experiments</i>	128
1.3. Research in labor economics	131
1.3.1. <i>How have labor economists used field experiments?</i>	134
1.4. Summary	140
2. Human Capital	140
2.1. Measuring the effects of direct inputs	141
2.2. Teacher quality	144
2.3. Measuring the effects of policies governing the system	147
3. Labor Market Discrimination	149
3.1. Data patterns in labor markets	150
3.2. Theories of discrimination	156
3.2.1. <i>Taste-based discrimination</i>	157
3.2.2. <i>Statistical discrimination</i>	157
3.2.3. <i>Optimal employer behavior</i>	160
3.3. Empirical tests	161
3.3.1. <i>Observational data</i>	161
3.3.2. <i>Field experiments</i>	169

[☆] We gratefully acknowledge financial support from ELSE . We thank the editors, Orley Ashenfelter and David Card for comments. We thank Alec Brandon, David Herberich, Dana Krueger, Richard Murphy, Yana Peysakhovich and László Sándor for excellent research assistance. All errors remain our own.

¹ Tel.: +1-773-702-9811; fax: +1-773-702-8490.

² Tel: +44-207-679-5853; fax: +44-207-916-2775.

E-mail addresses: jlist@uchicago.edu (John A. List), i.rasul@ucl.ac.uk (Imran Rasul).

4. Firms	177
4.1. Monetary incentives	178
4.1.1. <i>Theoretical framework</i>	180
4.1.2. <i>Evidence from the field</i>	184
4.2. Non-monetary incentives	186
4.2.1. <i>Theoretical framework</i>	189
4.2.2. <i>Evidence from the field</i>	191
4.3. The employment relationship	202
4.3.1. <i>Gift exchange</i>	202
4.3.2. <i>Shirking</i>	205
4.4. Moving forward	207
5. Households	208
5.1. Efficiency	210
5.2. Moving forward	212
6. Concluding Remarks	213
References	213

Abstract

We overview the use of field experiments in labor economics. We showcase studies that highlight the central advantages of this methodology, which include: (i) using economic theory to design the null and alternative hypotheses; (ii) engineering exogenous variation in real world economic environments to establish causal relations and learn about the underlying mechanisms; and (iii) engaging in primary data collection and often working closely with practitioners. To highlight the potential for field experiments to inform issues in labor economics, we organize our discussion around the individual life cycle. We therefore consider field experiments related to the accumulation of human capital, the demand and supply of labor, and behavior within firms, and close with a brief discussion of the nascent literature of field experiments related to household decision making.

JEL classification: C93; J01

Keywords: Field experiments; Labor economics

1. INTRODUCTION

This chapter overviews the burgeoning literature in field experiments in labor economics. The essence of this research method involves researchers engineering carefully crafted exogenous variation into real world economic environments, with the ultimate aim of identifying the causal relationships and mechanisms underlying them. This chapter describes this approach and documents how field experiments have begun to yield new insights for research questions that have long been studied by labor economists.

Given our focus on such long-standing questions, in no way do we attempt to do justice to the enormous literature to which field experiments are beginning to add. Our aim is rather to showcase specific field experiments that highlight what we view to be the central advantages of this methodology: (i) using economic theory to design the null

and alternative hypotheses; (ii) engineering exogenous variation in real world economic environments to establish causal relations and learn the mechanisms behind them; and (iii) engaging in primary data collection and often working closely with practitioners.

As with any research methodology in economics, of course, not every question will be amenable to field experiments. Throughout our discussion, we will bring to the fore areas in labor economics that remain relatively untouched by field experiments. In each case, we try to distinguish whether this is simply because researchers have not had opportunities to design field experiments related to such areas, or whether the nature of the research question implies that field experiments are not the best approach to tackle the problem at hand. Finally, even among those questions where field experiments can and have provided recent insights, we will argue that such insights are most enhanced when we are able to combine them with, or take inspiration from, other approaches to empirical work in economics. For example, a number of studies we describe take insights from laboratory environments to show the importance of non-standard preferences or behaviors in real world settings. A second class of papers combine field experimentation with structural estimation to measure potential behavioral responses in alternative economic environments, and ultimately to help design optimal policies to achieve any given objective.

The bulk of the chapter is dedicated to documenting the use and insights from field experiments in labor economics. To emphasize the relevance of field experiments to labor economics, we organize this discussion by following an individual over their life cycle. More specifically, we begin by considering field experiments related to early childhood interventions and the accumulation of human capital in the form of education and skills. We then consider research questions related to the demand and supply of labor, and labor market discrimination. We then move on to consider research questions related to behavior within firms: how individuals are incentivized within firms, and other aspects of the employment relationship. Finally, we conclude with a brief discussion of the nascent literature on field experiments related to household decision making.

We have chosen these topics selectively on the basis of where carefully designed field experiments have already been conducted. Within each field, we have decided to discuss a small number of papers using field experiments that, in our opinion, showcase the best of what field experiments can achieve. In no way is our discussion meant to be an exhaustive survey of the literature on field experiments in labor economics. The literature arising just in the past decade has grown too voluminous for even a tome to do it justice.

In each stage of the life cycle considered, wherever appropriate, we try to discuss: (i) the link between the design of field experiments and economic theory; (ii) the benefits of primary data collection that is inherent in field experimentation to further probe theory and distinguish between alternative hypotheses; and (iii) how reduced form effects identified from a field experiment can be combined with theory and structural modelling to make sample predictions and inform policy design.

In the remainder of this section, we place our later discussion into historical context by describing the experimental approach in science, and arguing how among economists, labor economists have for decades been at the forefront of exploiting and advancing experimental approaches to identify causal relationships. We then lay the groundwork for the discussion in later sections. We define the common elements at the heart of all field experiments, and then present a more detailed typology to highlight the subtle distinctions between various sub-types of field experiments. This approach allows us to discuss the advantages and disadvantages of field experiments over other forms of experimentation, such as large scale social experiments and laboratory based experiments.³

Our final piece of groundwork is to identify key trends in published research in labor economics over the past decade. This allows us to organize our later discussion more clearly in two dimensions. First, we think of nearly all research questions in labor economics as mapping to particular stages of an individual's life cycle. We therefore roughly organize research questions in labor economics into those relating to the accumulation of human capital, labor market entry and labor supply choices, behavior within firms, and household decision making. Second, we are able to focus in on those sub-fields in labor economics where extant field experiments have already begun to make inroads and provide new insights. In turn, this helps us make precise the types of research question field experiments are most amenable to, areas in which field experiments have been relatively under supplied, and those research questions that are better suited to alternative empirical methods.

1.1. The experimental approach in science

The experimental approach in scientific inquiry is commonly traced to Galileo Galilei, who pioneered the use of quantitative experiments to test his theories of falling bodies. Extrapolating his experimental results to the heavenly bodies, he pronounced that the services of angels were not necessary to keep the planets moving, enraging the Church and disciples of Aristotle alike. For his efforts, Galileo is now viewed as the Father of Modern Science. Since the Renaissance, fundamental advances making use of the experimental method in the physical and biological sciences have been fast and furious.⁴

Taking the baton from Galileo, in 1672 Sir Isaac Newton used experimentation to show that white light is equal to purity, again challenging the preachings of Aristotle. The experimental method has produced a steady stream of insights. Watson and Crick used data from Rosalind Franklin's X-ray diffraction experiment to construct a theory of the chemical structure of DNA; Rutherford's experiments shooting charged particles at a piece of gold foil led him to theorize that atoms have massive, positively charged

³ In this Handbook, Charness and Kuhn (2011, 2010) provide a useful discussion of extant laboratory studies in the area of labor economics.

⁴ For a more complete discussion see List and Reiley (2008).

nuclei; Pasteur rejected the theory of spontaneous generation with an experiment that showed that micro-organisms grow in boiled nutrient broth when exposed to the air, but not when exposed to carefully filtered air. Even though the experimental method produced a steady flow of important facts for roughly 400 years, the proper construction of a counterfactual control group was not given foundations until the early twentieth century.

1.1.1. An experimental cornerstone

In 1919, Ronald Fisher was hired at Rothamsted Manor to bring modern statistical methods to the vast experimental data collected by Lawes and Gilbert (Levitt and List, 2009). The data collection methods at Rothamsted Manor were implemented in the standard way to provide practical underpinnings for the ultimate purpose of agricultural research: to provide management guidelines. For example, one of the oldest questions in the area of agricultural economics relates to agricultural yields: what is the optimal application rate of fertilizer, seed, and herbicides?

In an attempt to modernize the experimental approach at Rothamsted, Fisher introduced the concept of randomization and highlighted the experimental tripod: the concepts of replication, blocking, and randomization were the foundation on which the analysis of the experiment was based (Street, 1990). Of course, randomization was the linchpin, as the validity of tests of significance stems from randomization theory.

Fisher understood that the goal of any evaluation method is to construct the proper counterfactual. Without loss of generality, define y_{i1} as the outcome for observational unit i with treatment, y_{i0} as the outcome for unit i without treatment. The treatment effect for plot i can then be measured as $\tau_i = y_{i1} - y_{i0}$. The major problem, however, is one of a missing counterfactual—plot i is not observed in both states. Fisher understood that methods to create the missing counterfactual to achieve identification of the treatment effect were invaluable, and his idea was to use randomization.

As Levitt and List (2009) discuss, Fisher's fundamental contributions were showcased in agricultural field experiments, culminating with the landmark 1935 book, *The Design of Experiments*, which was a catalyst for the actual use of randomization in controlled experiments. At the same time, Jerzy Neyman's work on agricultural experimentation showcased the critical relationship between experiments and survey design and the pivotal role that randomization plays in both (Splawa-Neyman, 1923a,b). Neyman's work continued in the area of sampling and culminated in his seminal paper, published in 1934. As Rubin (1990) notes, it is clear that randomization was "in the air" in the early 1920s, and the major influences of the day were by scholars doing empirical research on agricultural related issues. Clearly, such work revolutionized the experimental approach and weighs on experimental designs in all sciences today.

As emphasized throughout, we view field experimenters as being engaged in data generation, primary data collection, and data evaluation. Labor economists in particular

have been at the forefront of the use of experimental designs, as highlighted by the following two historic examples.

1.1.2. Early labor market field experiments at the Hawthorne plant

In the 1920s the Western Electric Company was the monopoly supplier of telephone equipment to AT&T. Western opted to have its main factory, the Hawthorne plant located in the suburbs of Chicago, be the main supplier for this important contract. The Hawthorne plant was considered to be one of the most advanced manufacturing facilities in America at the time, and employed roughly 35,000 people, mainly first- and second-generation immigrants (Gale, 2004). Always open to new techniques to improve efficiency and profitability, officials of Western were intrigued when the National Academy of Sciences expressed interest in a hypothesis put forth by electrical suppliers, who claimed that better lighting improved productivity.

The experimental exercises that resulted have few parallels within social science. The indelible footprint of these experiments laid the groundwork for a proper understanding of social dynamics of groups and employee relations in the workplace. Indeed, the data drawn from this research became the thrust of the human relations movement of the twentieth century, and represent the underpinnings of contemporary efforts of industry to motivate and deal with workers. In academia, the Hawthorne data spawned the development of a new field of study—Industrial Psychology—and remains an important influence on the manner in which scientists conduct experimental research today. Many of the issues raised in these studies are considered part of mainstream personnel economics, as discussed in the chapter in this Handbook on Human Resource Management practices by Bloom and Van Reenen (2011). In a later section of this chapter we review how a new generation of field experiments have provided new insights into these age old questions of behavior within firms.⁵

The first experiments executed at the Hawthorne plant have been famously denoted the “illumination experiments” because they varied the amount of light in the workplace. More specifically, between 1924 and 1927 the level of lighting was systematically changed for experimental groups in different departments (Mayo, 1933, pp. 55–56, Roethlisberger and Dickson, 1939, pp. 14–18, provide a more complete account). Workers in these departments were women who assembled relays and wound coils of wire, and their output was measured as units completed per unit of time.⁶

Discussions of these data have been widespread and have been an important influence on building the urban legend. For instance, Franke and Kaul (1978, p. 624) note that “Inexplicably worker output. . . generally increased regardless of increase or decrease in

⁵ Frederick Taylor’s seminal book, *The Principles of Scientific Management*, published in 1911, which creatively considered techniques to shorten task time, was also an important stimulus for the Industrial Psychology field.

⁶ A relay was a switching device activated in the telephone exchange as each number was dialed, and was a fairly mind-numbing task: assemble a coil, armature, contact springs, and insulators by fastening them to a fixture with four screws. On average, it was roughly one minute’s worth of work.

illumination.” Yet, the only account of these experiments published at the time is [Snow \(1927\)](#), published in an engineering newsletter, and he argues that “The corresponding production efficiencies by no means followed the magnitude or trend of the lighting intensities. The output bobbed up and down without direct relation to the amount of illumination.” Unfortunately, the article does not present data or any statistical analysis. Ever since, the literature has remained at a state of question since people thought that the data were lost. Indeed, an authoritative voice on this issue, [Rice \(1982\)](#) notes that “the original research data somehow disappeared.” [Gale \(2004, p. 439\)](#) expresses similar thoughts concerning the illumination experiments: “these particular experiments were never written up, the original study reports were lost, and the only contemporary account of them derives from a few paragraphs in a trade journal” ([Roethlisberger and Dickson, 1939](#); [Gillespie, 1991](#)).⁷

Using data preserved in two library archives [Levitt and List \(2010\)](#) dug up the original data from the illumination experiment, long thought to be destroyed. Their analysis of the newly found data reveals little evidence to support the existence of a Hawthorne effect as commonly described. Namely, there is no systematic evidence that productivity jumped whenever changes in lighting occurred. Alternatively, they do uncover some weak evidence consistent with more subtle manifestations of Hawthorne effects in the data. In particular, output tends to be higher when experimental manipulations are ongoing relative to when there is no experimentation. Also consistent with a Hawthorne effect is that productivity is more responsive to experimenter manipulations of light than naturally-occurring fluctuations.

As mysterious and legendary as the illumination experiments have become, it is fair to say that the second set of experiments conducted at the plant—the relay assembly experiments—have kept academics busy for years. Using an experimental area constructed for the illumination experiments, beginning in April 1927, researchers began an experiment meant to examine the effect of workplace changes upon productivity. In this case, the task was relay assembly.

The researchers began by secretly observing the women in their natural environment for two weeks, and then used various treatments, including manipulating the environment in such a way to increase and decrease rest periods, over different temporal intervals. While their design certainly did not allow easy assessment of clean treatment effects, the experimenters were puzzled by the observed pattern: output seemingly rose regardless of the change implemented. When output remained high after the researchers returned conditions to the baseline—output had risen from 2400 relays per week to nearly 3000 relays per week—management became interested in identifying the underlying mechanisms at work.

⁷ There are many other evaluations as well. For example, a controversial article written by [Bramel and Friend \(1981\)](#), heavily laced with Marxist ideology, takes a conspiratorial view of industrial psychologists and argues that the Hawthorne effect is simply the result of “capitalist bias among modern industrial psychologists.”

Western Electric subsequently brought in academic consultants, including Elton Mayo, in 1928. With Mayo's assistance, the experiments continued and by February of 1929, when productivity was at a startling rate of a new relay dropped down the chute every 40-50 seconds, the company besieged the five women with attention, besides "a new test room supervisor, an office boy, and a lady who helped with the statistics" others could be added: "an intermittent stream of other visitors or consultants: industrialists, industrial relations experts, industrial psychologists, and university professors." (Gale, 2004, p. 443). The experiment lasted until June 1932, when the women in the test room received their notices (except the exceptional worker, Jennie Sirchio, who worked in the office for a few months before being let go) after the stock market crash of October 24, 1929. The crash induced one in ten US phones to be disconnected in 1932, leading to a decrease in Western Electric's monopoly rents of more than 80%.

The five year experiment provided a wealth of data, and much of the Hawthorne Effect's statistical underpinnings are a direct result of the relay assembly experiment. Mayo's (1933) results concluded that individuals would be more productive when they knew they were being studied.⁸ For this insight, Mayo came to be known as the "father of the Hawthorne effect", and his work led to the understanding that the workplace was, importantly, a system that was first and foremost social, and composed of several interdependent parts. When we present a detailed typology of field experiments later in this section, we make precise a distinction between those field experiments in which agents are aware of their participation in an experiment, and those in which they are unaware of the exogenous manipulation of their economic environment.

Mayo stressed that workers are not merely at work to earn an honest wage for an honest day's effort, rather they are more prominently influenced by social demands, their need for attention, input to decision making, and by the psychological factors of the environment. The notion that workers effort and behavior are driven by more than the monetary rewards of work, is an idea that has received close scrutiny among the most recent generation of field experiments in firms, as reviewed later.

Clearly, Mayo argued, being the object of attention with the study induced a sense of satisfaction among workers that made them feel proud and part of a cohesive unit, generating greater productivity levels than could ever be imagined. Mayo's disciples, Leta and Leta Roethlisberger and William Dickson, another engineer at Western Electric, produced a detailed assessment that focused mainly on the relay experimental data (Roethlisberger and Dickson, 1939) and generated similar conclusions. Industrial psychology would soon

⁸ Derivative of this path-breaking experiment were two experiments run alongside the relay experiment. Both were started in August of 1928; one was a second relay experiment, the other a mica splitting experiment. In the second relay experiment, five women workers were subjected to variations in a small group incentive program from August 1928 to March 1929. In the mica splitting experiment, the researchers began by secretly monitoring the output of five women at their regular department workstations. Their job was to split, measure, and trim mica chips that were to be used for insulation. After observing the workers secretly, they moved the women to a special test room where, unlike their cohorts, they received 10-minute rest breaks at 9:30 a.m. and 2:30 p.m.

find an important place in undergraduate and graduate curricula. Again in later sections, we provide examples of where field experiments have taken insights from psychology and laboratory environments to check for the existence and quantitative importance of such behaviors that are not encompassed within neoclassical economic models.

It is difficult to understate the importance of these findings, as they have served as the paradigmatic foundation of the social science of work (Franke and Kaul, 1978), providing a basis for an understanding of the economics of the workplace, and dramatically influenced studies in organizational development and behavior, leadership, human relations, and workplace design. The results also provide an important foundation for experimental work within the social sciences, including economics, where one must constantly be aware of the effects argued to be important in the Hawthorne relay experiment.⁹

1.1.3. Large-scale social experiments

A second period of interest directly related to field experiments in labor economics is the latter half of the twentieth century, during which government agencies conducted a series of large-scale social experiments.¹⁰ In the US, social experiments can be traced to Heather Ross, an MIT economics doctoral candidate working at the Brookings Institution. As Levitt and List (2009) discuss, Ross wrote a piece titled “A Proposal for Demonstration of New Techniques in Income Maintenance”, in which she suggested a randomly assigned social experiment to lend insights into the policy debate.

The experiment that resulted began in 1968 in five urban communities in New Jersey and Pennsylvania: Trenton, Paterson, Passaic, and Jersey City in NJ, and Scranton, PA and eventually became Ross’ dissertation research (“An Experimental Study of the Negative Income Tax”, which cost more than \$5 million—exceeding \$30 million in today’s dollars). The idea behind the experiment was to explore the behavioral effects of negative income taxation, a concept first introduced by Milton Friedman and Robert Lampman, who was at the University of Wisconsin’s poverty institute.¹¹ The experiment, which targeted roughly 1300 male-headed households who had at least one employable

⁹ The success of the relay assembly experiments led to in-depth surveys (from 1928-1931) and one final experiment in the Hawthorne plant—the “bank wiring” experiment, designed by Mayo and others from 1931-1932. The researchers began by examining the productivity of 14 men who assembled telephone terminals. They then moved these men to a special test room, without introducing any other changes in work or pay conditions. Despite the move to a separate experimental setting, the men’s output did not increase.

¹⁰ This, and the subsequent subsections, draw from Harrison and List (2004), List (2006), and Levitt and List (2009). There are many definitions of social experiments in the economics literature. Ferber and Hirsch (1982, p. 7) define a social experiment in economics as “... a publicly funded study that incorporates a rigorous statistical design and whose experimental aspects are applied over a period of time to one or more segments of a human population, with the aim of evaluating the aggregate economic and social effects of the experimental treatments.” Greenberg and Shroder (2004) define a social experiment as having at least the following four features: (i) random assignment, (ii) policy intervention, (iii) follow-up data collection, and (iv) evaluation.

¹¹ As the Editors pointed out, the basic idea of a negative income tax was a part of the liberal party platform in the 1940s, and it is usually argued that it was designed by Juliet Rhys-Williams, an amazing advocate for women in that period.

person, experimentally varied both the guaranteed level of income and the negative tax rate (Ross, 1970). The guaranteed level of income ranged from 50% to 125% of the estimated poverty line income level for a family of four (\$1650–\$4125 in 1968 dollars) while the negative income tax rate ranged from 30% to 70%.¹² The experiment lasted three years. Families in both the control and treatment groups were asked to respond to questionnaires every three months during this time span, with the questions exploring issues such as family labor supply, consumption and expenditure patterns, general mobility, dependence on government, and social integration.

The most interesting outcome for labor economists involved labor supply. Strong advocates of the negative income tax program argued that the program would provide positive, or at least no negative, work incentives. Many economists, however, were skeptical, hypothesizing that the results would show some negative effect on work effort. Early experimental results discussed in Ross (1970), argued that work effort did not decline for the treatment groups. In fact, as Ross (1970, p. 568) indicates “there is, in fact, a slight indication that the participants’ overall work effort increased during the initial test period.”

Since this initial exploration, other scholars have re-examined the experimental design and data, coming to a less optimistic appraisal. An excellent elucidation is Ashenfelter (1990), who notes that because of attrition it is not actually possible to simply tabulate the results. In this sense, and from the experimenters point of view, the experiments were flawed in part because the design took little advantage of the inherent advantages of randomization. Of course, the ultimate policy test is whether the income maintenance programs increased work incentives relative to the existing welfare system, which as Moffitt (1981) notes at that time had large benefit–reduction rates that may have discouraged work. In certain cases, the new approach did outperform existing incentive schemes, in others it did not.

More importantly for our purposes, the New Jersey income maintenance experiment is generally considered to be the first large-scale social experiment conducted in the US, for which Ross is given credit (Greenberg et al., 1999; Greenberg and Shroder, 2004).¹³

¹² The negative income tax rate works as follows. Assume that John is randomly inserted into the 100% guaranteed income (\$3300), 50% negative tax rate treatment. What this means is that when the policy binds, for each \$1 that John’s family earns on its own, they receive \$0.50 less in federal benefits. Thus, if John’s family earns \$2000 in year one, they would receive \$1000 less in program benefits, or \$2300, resulting in a total income of \$4300. In this case, if in any year John’s family earns \$6600 or more, program benefits are zero.

¹³ We emphasize large scale because there were a handful of other social experiments—such as the Perry Preschool Project begun in 1962—that preceded the New Jersey Income Maintenance experiment (Greenberg et al., 1999), and that are still being evaluated today (Heckman et al., forthcoming). A prevalent type of social experimentation in recent years is the paired-audit experiments to identify and measure discrimination. These involve the use of “matched pairs” of individuals, who are made to look as much alike as possible apart from the protected characteristics. These pairs then confront the target subjects, which are employers, landlords, mortgage loan officers, or car salesmen. The majority of audit studies conducted to date have been in the fields of employment discrimination and housing discrimination (Riach and Rich, 2002).

The contribution of Ross, along with the excellent early summaries of the virtues of social experimentation (Orcutt and Orcutt, 1968), appears to have been instrumental in stimulating the explosion in social experiments in the ensuing decades.^{14, 15}

Such large-scale social experiments have continued in the US, and have included employment programs, electricity pricing, and housing allowances (see Hausman and Wise, 1985, for a review). While this early wave of social experiments tended to focus on testing new programs, more recent social experiments tended to be “black box” in the sense that packages of services and incentives were proffered, and the experiments were meant to test incremental changes to existing programs.¹⁶ This generation of social experiments had an important influence on policy, contributing, for instance, to the passage of the Family Support Act of 1988, which overhauled the AFDC program. Indeed, as Manski and Garfinkel (1992) note, in Title II, Section 203, 102 Stat. 2380, the Act even made a specific recommendation on evaluation procedures: “a demonstration project conducted . . . shall use experimental and control groups that are composed of a random sample of participants in the program.”

Much like the experimental contributions of the agricultural literature of the 1920s and 1930s, the large-scale social experiments conducted in the twentieth century influenced the economics literature immensely. Since the initial income maintenance social experiment, there have been more than 235 known completed social experiments (Greenberg and Shroder, 2004), each exploring public policies in health, housing, welfare, and the like. The early social experiments were voluntary experiments typically

¹⁴ The original negative income tax experiment led to three other early experiments on income maintenance, which drew samples from rural areas of North Carolina and Iowa (1970-72); Seattle and Denver (1970-78); and Gary, Indiana (1971-74). These experiments went beyond studying urban husband-wife couples that were studied in the New Jersey income maintenance experiment. For instance, the North Carolina/Iowa study was conducted by the Institute of Research on Poverty to explore behavior among the rural poor. Only one and two parent black households were studied in the Gary, IN test. The Seattle-Denver study represented the most comprehensive, including blacks, Chicanos, and whites who had either one or two parents in the household. By and large, the evidence gathered in these studies reinforced the main result in the New Jersey study, but these new studies highlighted additional insights that were important for policy making, such as in differences between male and female labor force participation, unemployment duration, and welfare participation.

¹⁵ An early social experiment in Europe was the study of Intensified Employment Services in Eskilstuna, Sweden. In 1975, a small-town employment office received a personnel reinforcement for three months and split a group of 410 unemployed job seekers who had been registered at the office for at least three months into a treatment group ($n = 216$) and a control group ($n = 194$). The control group received normal service and used the services of the office for an average of 1.5 hours over the course of the experiment, while the treatment group used office services for an average of 7.5 hours, allowing office personnel to work more intensely on the individual problems of the treatment subjects. The findings were that the percent of workers with a job at the end of the experiment, unemployment spells during the experiment, and earnings were all favorably influenced by the employment services studied. A discussion of this study, as well as other European social experiments in labor market policy can be found in Bjorklund and Regner (1996) and the various Digests of Social Experiments due to Greenberg, and Shroder. Two of the more famous examples are the Norwegian Training Experiment (Raaum and Torp, 1993) and the Restart Programme in the United Kingdom (White and Lakey, 1992).

¹⁶ For example, whereas over 80% of social experiments from 1962-74 tested new programs, since 1983 only roughly 33% did so (Greenberg et al., 1999).

designed to measure basic behavioral relationships, or deep structural parameters, which could be used to evaluate an entire spectrum of social policies. Optimists even believed that the parameters could be used to evaluate policies that had not even been conducted. As Heckman (1992) notes, this was met with deep skepticism along economists and non-economists alike, and ambitions have since been much more modest.

As Manski and Garfinkel (1992) suggest, this second wave of social experiments had a methodological influence within academic circles, as it provided an arena for the 1980s debate between experimental advocates and those favoring structural econometrics using naturally-occurring data. Manski and Garfinkel (1992) provide an excellent resource that includes insights on the merits of the arguments on both sides, and discusses some of the important methodological issues. Highlighting some of the weaknesses of social experiments helps to clarify important distinctions we draw between social experiments and the generation of field experiments which has followed.

1.1.4. Potential shortcomings of social experiments

One potential problem arising in social experiments is “randomization bias”, a situation wherein the experimental sample is different from the population of interest because of randomization. It is commonly known in the field of clinical drug trials that persuading patients to participate in randomized studies is much harder than persuading them to participate in non-randomized studies (Kramer and Shapiro, 1984). The same problem applies to social experiments, as evidenced by the difficulties that can be encountered when recruiting decentralized bureaucracies to administer the random treatment (Hotz, 1992).¹⁷

Doolittle and Traeger (1990) provide a description of the practical importance of randomization bias when describing their experience in implementing the Job Training Partnership Act. Indeed, in almost any social experiment related to job training programs, it is a concern that those most likely to benefit from the program select into the program. Moreover, as Harrison and List (2004) discuss, in social experiments, given the open nature of the political process, it is almost impossible to hide the experimental objective from the person implementing the experiment or the subject, opening up the possibility of such self-selection. As Heckman (1992) puts it, comparing social experiments to agricultural experiments: “plots of ground do not respond to anticipated treatments of fertilizer, nor can they excuse themselves from being treated.”

To see this more formally, we follow the notation above and assume that $\tau_i = y_{i1} - y_{i0}$ is the treatment effect for individual i . Figure 1 shows a hypothetical density of τ_i in the population, a density assumed to have mean, τ^* . In this case, the parameter τ^* is equivalent to the average treatment effect; this is the treatment effect of interest if the analyst is pursuing an estimate of the average effect in this population.

¹⁷ There is a growing body of evidence from laboratory settings on how individuals self-select into treatments when allowed to do so. See Lazear et al. (2009) for a recent such study, and the discussion in Charness and Kuhn (2011, 2010).




Figure 1 *Simple illustration of the selection problem.*

The concern is that selection into the experiment is not random, but might occur with a probability related to τ . Using this notion to formulate the selection rule leads to positive selection: subjects with higher τ values are more likely to participate if offered. In Fig. 1, we denote the cutoff value of τ_i as τ^+ : people above τ^+ participate, those below do not.

In this case, the treatment effect on the treated is what is measured in the social experiment: τ_{TT} . τ_{TT} is equal to $E(\tau_i | \tau_i > \tau^+)$, which represents the estimate of the treatment effect for those who select to participate. A lack of recognition of selection causes the analyst to mis-measure the treatment effect for the population of interest. Figure 1 also shows the treatment effect on the untreated, τ_{TU} . This τ_{TU} is equal to $E(\tau_i | \tau_i < \tau^+)$, which represents the unobserved estimate of the treatment effect for those who chose not to participate. Therefore, the population parameter of interest, τ^* , is a mixture of these two effects: $\tau^* = \text{Pr} \times \tau_{TT} + (1 - \text{Pr}) \times \tau_{TU}$, where Pr represents the probability of $\tau_i > \tau^+$. Even if one assumes that the population density of τ_i among participants is isomorphic to the population density of inferential interest, such selection frustrates proper inference. A related concern is whether the density of τ_i in the participant population exactly overlaps with the population of interest.

A second issue stems from Heckman (1992), Heckman and Smith (1995), and Manski (1995), who contend that participants in small-scale experiments may not be representative of individuals who would participate in ongoing, full-scale programs. Such non-representativeness of the experimental sample could occur because of a lack of information diffusion, the reluctance of some individuals to subject themselves to random assignment, or resource constraints in full-scale programs that result in program administrators restricting participants to people meeting certain criteria. As a result, making inference on how individuals would respond to the same intervention were they to be scaled up is not straightforward.

A third set of concerns stem from the supply side of those implementing the social experiment as it is scaled up. For example, the quality of those administering the intervention might be very different from the quality of personnel selected to take part in the original social experiment. Moreover, the ability of administrative agencies to closely

monitor those charged with the actual implementation of the program might also vary as programs are scaled-up. These concerns might also apply to field experiments unless they are explicitly designed to allow for such possibilities. In general, the role played by program implementers in determining program outcomes remains poorly understood and is a rich area for future study both for field experiments and researchers in general.

A fourth concern that arises in social experiments is attrition bias. Attrition bias refers to systematic differences between the treatment and control groups because of differential losses of participants. As Hausman and Wise (1979) note, a characteristic of social experiments is that individuals are surveyed before the experiment begins as well as during the experiment, which in many cases is several years. This within-person experimental design permits added power compared to a between-person experimental design—because of the importance of individual effects. But, there are potential problems, as they note (p. 455): “the inclusion of the time factor in the experiment raises a problem which does not exist in classical experiments—attrition. Some individuals decide that keeping the detailed records that the experiments require is not worth the payment, some move, some are inducted into the military.”¹⁸

Beyond sampling and implementation shortcomings, social experiments also run the risk of generating misleading inference out of sample due to the increased scrutiny induced by the experiment. If experimental participants understand their behavior is being measured in terms of certain outcomes, some of them might attempt to succeed along these outcomes. Such effects have been deemed “John Henry” effects for the control sample because such participants work harder to show their worth when they realize that they are part of the control group. More broadly, some studies denote such effects as “Hawthorne” effects; if these Hawthorne effects do not operate equally on the treatment and control group, bias is induced.¹⁹

Another factor that might lead to incorrect inference in a social experiment is control group members seeking available substitutes for treatment. This is denoted “substitution bias” in the literature, a bias that can result in significant understatement of the treatment effect. Substitution bias can occur if a new program being tested experimentally absorbs resources that would otherwise be available to members of the control group or, instead, if as a result of serving some members of a target group, the new program frees up resources available under other programs that can now be used to better serve members of the

¹⁸ Problems of attrition are well known and detailed discussions can be found in Hausman and Wise (1979) and the various chapters in Manski and Garfinkel (1992).

¹⁹ Note that the development field experiments that have arisen recently often have to confront this issue directly when making inference from their studies—even though subjects might not know that they are randomized, a survey is used to measure the outcomes so repeated interactions are a certainty. One paper that attempts to quantify the effects is due to Gine et al. (2007). In a similar spirit, Muralidharan and Sundararaman (2007) present evidence from a randomized control trial on educational interventions in India. They also present evidence to distinguish the effects of the intervention from the mere effects of being part of an observational study *per se*.

control group. The practical importance of substitution bias is provided in [Puma et al. \(1990\)](#) and [Heckman and Smith \(1995\)](#).

Although these concerns, as well as others not discussed here, need always to be accounted for, social experiments continue to be an important and valuable tool for policy analysis, as evidenced by two recent and notable large scale undertakings: *Moving To Opportunity* ([Katz et al., 2001](#)) and *PROGRESA* ([Schultz, 2004](#)), as well as the more recent social experiments documented in [Greenberg and Shroder \(2004\)](#).

1.2. Field experiments

Following from the first two periods of experimentation discussed above, the third distinct period of field experimentation is the most recent surge of field experiments in economics. [Harrison and List \(2004\)](#), [List \(2006\)](#) and [List and Reiley \(2008\)](#) provide recent overviews of this literature. The increased use of this approach reflects a long running trend in labor economics, and applied microeconomics more generally, to identify causal effects. This is not surprising given that nearly all of the central research questions in labor economics are plagued by econometric concerns related to the simultaneous determination of individual decisions related to the accumulation of human capital, self-selection into labor markets and careers. Furthermore, many of the key variables that underlie behavior in labor markets—such as motivation or talent—are either simply unmeasured or measured with error in standard surveys.

Field experiments form the most recent addition to the wave of empirical strategies to identify causal effects that have entered mainstream empirical research in labor economics since the mid 1980s. For example, these are based on fixed effects, difference-in-differences, instrumental variables, regression discontinuities, and natural experiments. Comprehensive reviews of these developments are provided in [Angrist and Krueger \(1999\)](#). At the same time as these research strategies have developed, greater emphasis has been placed on econometric methods that are robust to functional form and distributional assumptions. These include the development of semi-parametric and non-parametric estimation techniques. Reviews of these developments are provided in [Moffitt \(1999\)](#).

We view the increased use of field experiments to have its origins in the last decade in part because of an acceleration of three long-standing trends in how applied economic research is conducted: (i) the increased use of research designs to uncover credible causal effects; (ii) the increased propensity to engage in primary data collection; and (iii) the formation of ever closer interactions with practitioners and policy makers more generally.

Similar to the experiments at the Hawthorne plant and social experiments, but unlike the first-generation agricultural studies, the most recent field experiments typically apply randomization to human subjects to obtain identification. In contrast to social experiments, however, recent field experiments strive to carry out this randomization on naturally occurring populations in naturally occurring settings, often without the

research subjects being aware that they are part of an experiment. As a consequence, these more recent studies tend to be carried out opportunistically, and on a smaller scale than social experiments.²⁰

This current generation of field experiments oftentimes has more ambitious theoretical goals than social experiments, which largely aim to speak to policy makers and identify whether a package of measures leads to some desired change in outcomes. Modern field experiments in many cases are designed to test economic theory, collect facts useful for constructing a theory, and organize primary data collection to make measurements of key parameters, assuming a theory is correct. Field experiments can also help provide the necessary behavioral principles to permit sharper inference from laboratory or naturally-occurring data. Alternatively, field experiments can help to determine whether lab or field results should be reinterpreted, defined more narrowly than first believed, or are more general than the context in which they were measured. In other cases, field experiments might help to uncover the causes and underlying conditions necessary to produce data patterns observed in the lab or the field.

Since nature in most cases does not randomize agents into appropriate treatment and control groups, the task of the field experimental researcher is to develop markets, constructs, or experimental designs wherein subjects are randomized into treatments of interest. The researcher faces challenges different from those that arise either in conducting laboratory experiments or relying on naturally occurring variation. The field experimenter does not exert the same degree of control over real markets as the scientist does in the lab. Yet, unlike an empiricist who collects existing data, the field experimenter is in the data generating business, as opposed to solely engaging in data collection or evaluation. Consequently, conducting successful field experiments demands a different set of skills from the researcher: the ability to recognize opportunities for experimentation hidden amidst everyday phenomena, an understanding of experimental design and evaluation methods, knowledge of economic theory to motivate the research, and the interpersonal skills to manage what are often a complex set of relationships involving parties to an experiment.

1.2.1. What is a field experiment?

Harrison and List (2004) propose six factors that can be used to determine the field context of an experiment: the nature of the subject pool, the nature of the information that the subjects bring to the task, the nature of the commodity, the nature of the task or trading rules applied, the nature of the stakes, and the environment in which the

²⁰ In this sense, field experiments parallel the research approach that exploits “natural experiments” (Meyer, 1995; Rosenzweig and Wolpin, 2000; Angrist and Krueger, 2001), the difference being that in a field experiment the researcher actually controls the randomization herself, whereas in the natural experiment approach the researcher attempts to find sources of variation in existing data that are “as good as randomly assigned.” In addition, the close involvement of the researcher from the outset allows for primary data collection to perhaps directly help shed light on the underlying mechanisms driving causal effects.

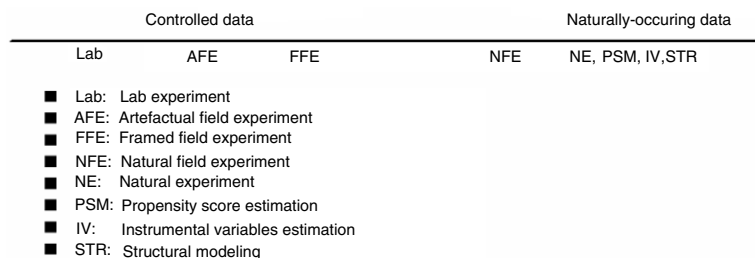


Figure 2 *A field experiment bridge.*

subjects operate. Using these factors, they discuss a broad classification scheme that helps to organize one's thoughts about the factors that might be important when moving from non-experimental to experimental data.

They classify field experiments into three categories: artefactual, framed, and natural. Figure 2 shows how these three types of field experiments compare and contrast with laboratory experiments and approaches using naturally occurring non-experimental data. On the far left in Fig. 2 are laboratory experiments, which typically make use of randomization to identify a treatment effect of interest in the lab using a subject pool of students. In this Handbook, Charness and Kuhn (2011, 2010) discuss extant laboratory studies in the area of labor economics.

The other end of the spectrum in Fig. 2 includes empirical models that make necessary identification assumptions to identify treatment effects from naturally-occurring data. For example, identification in simple natural experiments results from a difference in difference regression model: $Y_{it} = X_{it}\beta + \tau T_{it} + \eta_{it}$, where i indexes the unit of observation, t indexes years, Y_{it} is the outcome, X_{it} is a vector of controls, T_{it} is a binary treatment variable equal to one if unit i is treated and zero otherwise, $\eta_{it} = \alpha_i + \lambda_t + \varepsilon_{it}$, and τ is measured by comparing the difference in outcomes before and after for the treated group with the before and after outcomes for the non treated group. A major identifying assumption in this case is that there are no time-varying, unit-specific shocks to the outcome variable that are correlated with T_{it} , and that selection into treatment is independent of a temporary individual specific effect.

Useful alternatives include the method of propensity score matching (PSM) developed in Rosenbaum and Rubin (1983). Again, if both states of the world were observable, the average treatment effect, τ , would equal $\bar{y}_1 - \bar{y}_0$. However, given that only y_1 or y_0 is observed for each observation, unless assignment into the treatment group is random, generally $\tau \neq \bar{y}_1 - \bar{y}_0$. The solution advocated by Rosenbaum and Rubin (1983) is to find a vector of covariates, Z , such that $y_1, y_0 \perp T|Z$, $\Pr(T = 1|Z) \in (0, 1)$, where \perp denotes independence. This assumption is called the "conditional independence assumption" and intuitively means that given Z , the non-treated outcomes are what the treated outcomes would have been had they not

been treated. Or, likewise, that selection occurs only on observables. If this condition holds, then treatment assignment is said to be “strongly ignorable” (Rosenbaum and Rubin, 1983, p. 43). To estimate the average treatment effect (on the treated), only the weaker condition $E[y_0|T = 1, Z] = E[y_0|T = 0, Z] = E[y_0|Z] \Pr(T = 1|Z) \in (0, 1)$ is required. Thus, the treatment effect is given by $\tau = E[\bar{y}_1 - \bar{y}_0|Z]$, implying that conditional on Z , assignment to the treatment group mimics a randomized experiment.²¹

Other more popular methods of estimating treatment effects include the use of instrumental variables (Rosenzweig and Wolpin, 2000) and structural modeling. Assumptions of these approaches are well documented and are not discussed further here (Angrist and Krueger, 1999; Blundell and Costa-Dias, 2002). Between the two extremes in Fig. 2 are various types of field experiment. We now turn to a more patient discussion of these types.

1.2.2. A more detailed typology of field experiments

Following Harrison and List (2004), we summarize the key elements of each type of field experiment in Table 1. This also makes precise the differences between field and laboratory experiments.

Harrison and List (2004) argue that a first useful departure from laboratory experiments using student subjects is simply to use “non-standard” subjects, or experimental participants from the market of interest. In Table 1 and Fig. 2, these are denoted as “artefactual” field experiments. This type of field experiment represents a potentially useful type of exploration outside of traditional laboratory studies because it affords the researchers with the control of a standard lab experiment but with the realism of a subject pool that are the natural actors from the market of interest. In the past decade, artefactual field experiments have been used in financial applications (Alevy et al., 2007; Cipriani and Guarino, 2009), to test predictions of game theory (Levitt et al., 2009), and in applications associated with labor economics (Cooper et al., 1999).²²

²¹ Several aspects of the approach are not discussed in this discussion. For example, for these conditions to hold the appropriate conditioning set, Z , should be multi-dimensional. Second, upon estimation of the propensity score, a matching algorithm must be defined in order to estimate the missing counterfactual, y_0 , for each treated observation. The average treatment effect on the treated (TT) is given by,

$$\tau_{TT} = E[E[y_1|T = 1, p(Z)] - E[y_0|T = 0, p(Z)]] = E[E[y_1 - y_0|p(Z)]],$$

where the outer expectation is over the distribution of $Z|T = 1$. These and other issues are discussed in List et al. (2003).

²² Harrison and List (2004) discuss in detail whether student subjects exhibit different behaviors in laboratory environments that individuals drawn from other subject pools. A parallel trend in laboratory settings has been the use of “real-effort” experiments, as discussed in Charness and Kuhn (2011, 2010).

Table 1 A typology of field experiments.

	Non experimental	Natural experiments	Natural field experiments	Framed field experiment	Artefactual field experiment	Laboratory
In the field	×	×	×	×		
Real incentives	×	×	×	×		
Real task or information	×	×	×	×		
Aware of experiment				×	×	×
Appropriate people	×	×	×	×	×	
Researcher intervenes				×	×	×
Exogenous change		×				

Another example of the use of artefactual field experiments is to explain or predict non-experimental outcomes. An example of this usage is [Barr and Serneels \(2009\)](#), who correlate behavior in a trust game experiment with wage outcomes of employees of Ghanaian manufacturing enterprises. They report that a one percent increase in reciprocity in these games is associated with a fifteen percent increase in wages. Another example is [Attanasio et al. \(2009\)](#) who combine household data on social networks with a field experiment conducted with the same households in Colombia, to investigate who pools risk with whom when risk pooling arrangements are not formally enforced. Combining non-experimental and experimental research methods in this way by conducting an artefactual field experiment among survey respondents provides an opportunity to study the interplay of risk attitudes, pre-existing networks, and risk-sharing.

Moving closer to how naturally-occurring data are generated, [Harrison and List \(2004\)](#) denote a “framed field experiment” as the same as an artefactual field experiment, except that it incorporates important elements of the context of the naturally occurring environment with respect to the commodity, task, stakes, and information set of the subjects. Yet, it is important to note that framed field experiments, like lab experiments and artefactual field experiments, are conducted in a manner that ensures subjects understand that they are taking part in an experiment, with their behavior subsequently recorded and scrutinized. Framed field experiments include the Hawthorne plant experiments, the social experiments of the twentieth century, and two related experimental approaches.

One related approach might be considered a cousin of social experiments: the collection of studies done in developing countries that use randomization to identify causal effects of interventions in settings where naturally-occurring data are limited. The primary motivation for such experiments is to inform public policy. These studies typically use experimental treatments more bluntly than the controlled treatments discussed above, in that the designs often randomly introduce a package of several interventions. On the other hand, this package of measures is directly linked to a menu of actual public policy alternatives. A few recent notable examples of this type of work are the studies such as [Kremer et al. \(2009\)](#) and [Duflo et al. \(2006\)](#).

Framed field experiments have also been done with a greater eye towards testing economic theory, for instance several framed field experiments of this genre have been published in the economics literature, ranging from further tests of auction theory ([Lucking-Reiley, 1999](#); [Katkar and Reiley, 2006](#)), tests of the theory of private provision of public goods ([Bohm, 1984](#); [List, 2004a](#)), tests that examine the relative predictive power of neoclassical theory versus prospect theory ([List, 2003b, 2004b](#)), tests that explore issues in cost/benefit analysis and preference elicitation ([List, 2001, 2002a, 2003a](#); [Lusk and Fox, 2003](#); [Rozan et al., 2004](#); [Ding et al., 2005](#)), tests that explore

competitive market theory in the field List, 2002b, 2004c; List and Price, 2005), and tests of information assimilation among professional financial traders (Alevy et al., 2007).²³

Unlike social experiments, this type of framed field experiment does not need to worry about many of the shortcomings discussed above. For example, since subjects are unaware that the experiment is using randomization, any randomization bias should be eliminated. Also, these experiments tend to be short-lived and therefore attrition bias is not of major importance. Also, substitution bias should not be a primary concern in these types of studies. The cost of not having these concerns is that rarely can the long run effects of experimentally introduced interventions be assessed. This might limit therefore the appropriateness of such field experiments for questions in labor economics in which there are long time lags between when actions are made and outcomes realized.

As Levitt and List (2007a,b) discuss, the fact that subjects are in an environment in which they are keenly aware that their behavior is being monitored, recorded, and subsequently scrutinized, might also cause generalizability to be compromised. Decades of research within psychology highlight the power of the role obligations of being an experimental subject, the power of the experimenter herself, and the experimental situation (Orne, 1962). This leads to our final type of field experiment—“natural field experiments,” which complete Table 1 and Fig. 2.

Natural field experiments are those experiments completed in cases where the environment is such that the subjects naturally undertake these tasks and where the subjects do not know that they are participants in an experiment. Therefore, they neither know that they are being randomized into treatment nor that their behavior is subsequently scrutinized. Such an exercise is important in that it represents an approach that combines the most attractive elements of the lab and naturally-occurring data: randomization and realism. In addition, it is difficult for people to respond to treatments they do not necessarily know are unusual, and of course they cannot excuse themselves from being treated. Hence, many of the limitations cited above are not an issue when making inference from data generated by natural field experiments. As we document in later sections, natural field experiments have already been used to answer a wide range of traditional research questions in labor economics.

1.2.3. *Simple rules of thumb for experimentation*

Scholars have produced a variety of rules of thumb to aid in experimental design. Following List et al. (2010), we provide a framework to think through these issues. Suppose that a single treatment T results in (conditional) outcomes Y_{i0} if $T = 0$, where $Y_{i0}|X_i \sim N(\mu_0, \sigma_0^2)$, and Y_{i1} if $T = 1$, where $Y_{i1}|X_i \sim N(\mu_1, \sigma_1^2)$. Since the experiment has not yet been conducted, the experimenter must form beliefs about the variances of outcomes across the treatment and control groups, which may, for example,

²³ Of course, this is just a select sampling of the work of this sort, for a more comprehensive list please see www.fieldexperiments.com.

come from theory, prior empirical evidence, or a pilot experiment. The experimenter also has to make a decision about the minimum detectable difference between mean control and treatment outcomes, $\mu_1 - \mu_0 = \delta$, that the experiment is meant to be able to detect. In essence, δ is the minimum average treatment effect, $\bar{\tau}$, that the experiment will be able to detect at a given significance level and power. Finally, we assume that the significance of the treatment effect will be determined using a t -test.

The first step in calculating optimal sample sizes requires specifying a null hypothesis and a specific alternative hypothesis. Typically, the null hypothesis is that there is no treatment effect, i.e. that the effect size is zero. The alternative hypothesis is that the effect size takes on a specific value (the minimum detectable effect size). The idea behind the choice of optimal sample sizes in this scenario is that the sample sizes have to be just large enough so that the experimenter: (i) does not falsely reject the null hypothesis that the population treatment and control outcomes are equal, i.e. commit a Type I error; and, (ii) does not falsely accept the null hypothesis when the actual difference is equal to δ , i.e. commit a Type II error. More formally, if the observations for control and treatment groups are independently drawn and $H_0 : \mu_0 = \mu_1$ and $H_1 : \mu_0 \neq \mu_1$, we need the difference in sample means $\bar{Y}_1 - \bar{Y}_0$ (which are of course not yet observed) to satisfy the following two conditions related to the probabilities of Type I and Type II errors.

First, the probability α of committing a Type I error in a two-sided test, i.e. a significance level of α , is given by,

$$\frac{\bar{Y}_1 - \bar{Y}_0}{\sqrt{\frac{\sigma_0^2}{n_0} + \frac{\sigma_1^2}{n_1}}} = t_{\alpha/2} \Rightarrow \bar{Y}_1 - \bar{Y}_0 = t_{\alpha/2} \sqrt{\frac{\sigma_0^2}{n_0} + \frac{\sigma_1^2}{n_1}}, \quad (1)$$

where σ_T^2 and n_T for $T \in \{0, 1\}$ are the conditional variance of the outcome and the sample size of the control and treatment groups. Second, the probability β of committing a Type II error, i.e. a power of $1 - \beta$, in a one-sided test, is given by,

$$\frac{(\bar{Y}_1 - \bar{Y}_0) - \delta}{\sqrt{\frac{\sigma_0^2}{n_0} + \frac{\sigma_1^2}{n_1}}} = -t_\beta \Rightarrow \bar{Y}_1 - \bar{Y}_0 = \delta - t_\beta \sqrt{\frac{\sigma_0^2}{n_0} + \frac{\sigma_1^2}{n_1}}. \quad (2)$$

Using (1) to eliminate $\bar{Y}_1 - \bar{Y}_0$ from (2) we obtain,

$$\delta = (t_{\alpha/2} + t_\beta) \sqrt{\frac{\sigma_0^2}{n_0} + \frac{\sigma_1^2}{n_1}}. \quad (3)$$

It can easily be shown that if $\sigma_0^2 = \sigma_1^2 = \sigma^2$, i.e. $\text{var}(\tau_i) = 0$, then the smallest sample sizes that solve this equality satisfy $n_0 = n_1 = n$ and then,

$$n_0^* = n_1^* = n^* = 2 (t_{\alpha/2} + t_{\beta})^2 \left(\frac{\sigma}{\delta} \right)^2. \quad (4)$$

If the variances of the outcomes are not equal this becomes,

$$N^* = \left(\frac{t_{\alpha/2} + t_{\beta}}{\delta} \right)^2 \left(\frac{\sigma_0^2}{\pi_0^*} + \frac{\sigma_1^2}{\pi_1^*} \right), \quad (5)$$

$$\pi_0^* = \frac{\sigma_0}{\sigma_0 + \sigma_1}, \quad \pi_1^* = \frac{\sigma_1}{\sigma_0 + \sigma_1},$$

where $N = n_0 + n_1$, $\pi_0 + \pi_1 = 1$, $\pi_0 = \frac{n_0}{n_0 + n_1}$.

If sample sizes are large enough so that the normal distribution is a good approximation for the t-distribution, then the above equations provide a closed form solution for the optimal sample sizes. If sample sizes are small, then n must be solved by using successive approximations. Optimal sample sizes increase proportionally with the variance of outcomes, non-linearly with the significance level and the power, and decrease proportionally with the square of the minimum detectable effect. The relative distribution of subjects across treatment and control is proportional to the standard deviation of the respective outcomes. This suggests that if the variance of outcomes under treatment and control are fairly similar—namely, in those cases when there are expected to be homogeneous treatment effects—there should not be a large loss in efficiency from assigning equal sample sizes to each.

In cases when the *outcome* variable is dichotomous, under the null hypothesis of no treatment effect, $\mu_0 = \mu_1$, one should always allocate subjects equally across treatments. Yet, if the null is of the form $\mu_1 = k\mu_0$, where $k > 0$, then the sample size arrangement is dictated by k in the same manner as in the continuous case. If the cost of sampling subjects differs across treatment and control groups, then the ratio of the sample sizes is inversely proportional to the square root of the relative costs. Interestingly, differences in sampling costs have exactly the same effect on relative sample sizes of treatment and control groups as differences in variances.

In those instances where the unit of randomization is different from the unit of observation, special considerations must be paid to the correlation in outcomes between units in the same treated cluster. Specifically, the number of observations required is multiplied by $1 + (m - 1)\rho$, where ρ is the intracluster correlation coefficient and m is the size of each cluster. The optimal size of each cluster increases with the ratio of the within to between cluster standard deviation, and decreases with the square root of the ratio of the cost of sampling a subject to the fixed cost of sampling from a new cluster. Since the

optimal sample size is independent of the available budget, the experimenter should first determine how many subjects to sample in each cluster and then sample from as many clusters as the budget permits (or until the optimal total sample size is achieved).²⁴

A final class of results pertains to designs that include several levels of treatment, or more generally when the treatment variable itself is continuous, but we assume homogeneous treatment effects. The primary goal of the experimental design in this case is to simply maximize the variance of the treatment variable. For example, if the analyst is interested in estimating the effect of treatment and has strong priors that the treatment has a linear effect, then the sample should be equally divided on the endpoints of the feasible treatment range, with no intermediate points sampled. Maximizing the variance of the treatment variable under an assumed quadratic, cubic, quartic, etc., relationship produces unambiguous allocation rules as well: in the quadratic case, for instance, the analyst should place half of the sample equally distributed on the endpoints and the other half on the midpoint. More generally, optimal design requires that the number of treatment cells used should be equal to the highest polynomial order of the anticipated treatment effect, plus one.

1.2.4. Further considerations

In light of the differences between field experimentation and other empirical methods—lab experiments and using observational data—it is important to discuss some perceived differences and potential obstacles associated with this research agenda. One shortcoming of field experiments is the relative difficulty of replication vis-à-vis lab experiments.²⁵ As Fisher (1926) emphasized, replication is an important advantage of the experimental methodology. The ability of other researchers to reproduce quickly the experiment, and therefore test whether the results can be independently verified, not only serves to generate a deeper collection of comparable data but also provides incentives for the experimenter to collect and document data carefully.

There are at least three levels at which replication can operate. The first and most narrow of these involves taking the actual data generated by an experiment and reanalyzing the data to confirm the original findings. A second notion of replication is to run an experiment which follows a similar protocol to the first experiment to determine whether similar results can be generated using new subjects. The third and most general

²⁴ Relatedly, there is a recent but steadily expanding literature in statistics and economics on how experimental evidence on treatment effect heterogeneity may be used to maximize gains from social programs. One example is Bhattacharya and Dupas (2010) who study the problem of allocating a binary treatment among a target population based on observables, to maximize the mean social welfare arising from an eventual outcome distribution, when a budget constraint limits what fraction of the population can be treated.

²⁵ This is especially so if we compare field experiments to laboratory experiments that utilize student subject pools. Even by changing the subject pool slightly, as in artefactual field experiments, replicability becomes an issue as more still needs to be understood on the self-selection into experiments of such non-standard subjects (Charness and Kuhn, 2011, 2010).

conception of replication is to test the hypotheses of the original study using a new research design.

Lab experiments and many artefactual and framed field experiments lend themselves to replication in all three dimensions: it is relatively straightforward to reanalyze existing data, to run new experiments following existing protocols, and (with some imagination) to design new experiments testing the same hypotheses.

With natural field experiments, the first and third types of replication are easily done (i.e. reanalyzing the original data or designing new experiments), but the second type of replication (i.e. re-running the original experiment, but on a new pool of subjects) is more difficult. This difficulty arises because by their very nature, many field experiments are opportunistic and might be difficult to replicate because they require the cooperation of outside entities or practitioners, or detailed knowledge and the ability to manipulate a particular market.

Another consideration associated with field experiments relates to ethical guidelines (Dunford, 1990; Levitt and List, 2009). The third parties that field experimenters often need to work with can be concerned by the need to randomize units of observation into treatments. The benefits of such an approach need to be conveyed, as well as a practical sense of how to achieve this. For example, given resource constraints, practitioners are typically unable to roll out interventions to all intended recipients immediately. The field experimenter can intervene to randomly assign the order in which individuals are treated (or offered treatment), not whether they eventually receive the treatment or not.

With the onset of field experiments, new issues related to informed consent naturally arise. Ethical issues surrounding human experimentation is of utmost import. The topic of informed consent for human experimentation were recognized as early as the nineteenth century (Vollmann and Winau, 1996), but the principal document to provide guidelines on research ethics was the Nuremberg Code of 1947. The Code was a response to malfeasance of Nazi doctors, who performed immoral acts of experimentation during the Second World War. The major feature of the Code was that voluntary consent became a requirement in clinical research studies, where consent can be voluntary only if subjects: (i) are physically able to provide consent; (ii) are free from coercion; and, (iii) can comprehend the risks and benefits involved in the experiment.

What is right for medical trials need not be appropriate for the social sciences. To thoughtlessly adopt the Nuremberg Code whole cloth for field experiments without considering the implications would be misguided. In medical trials, it is sensible to have informed consent as the default because of the serious risk potential in most clinical studies. In contrast, the risks posed in some natural field experiments in economics are small or nonexistent, although such risks are almost certain to become more heterogeneous across field experiments as this research method becomes more prevalent. Hence while there might be valid arguments for making informed consent the exception, rather than the rule, in a field experimental context, it is true to say that covert

experimentation remains hotly debated in the literature. For more detailed discussions, the interested reader should see [Dingwall \(1980\)](#) and [Punch \(1985\)](#).

There are certain cases in which seeking informed consent directly interferes with the ability to conduct the research ([Homan, 1991](#)). For example, for years economists have been interested in measuring and detecting discrimination in the marketplace. Labor market field studies present perhaps the deepest amount of work in the area of discrimination. The work in this area can be parsed into two distinct categories, personal approaches and written applications.

Personal approaches include studies that have individuals either attend job interviews or apply for employment over the telephone. In these studies, the researcher matches two testers that are identical along all relevant employment characteristics except the comparative static of interest (e.g., race, gender, age), and after appropriate training the testers approach potential employers who have advertised a job opening. Researchers “train” the subjects simultaneously to ensure that their behavior and approach to the job interview are similar.

Under the written application approach, which can be traced to [Jowell and Prescott-Clarke \(1970\)](#), carefully prepared written job applications are sent to employers who have advertised vacancies. The usual approach is to choose advertisements in daily newspapers within some geographic area to test for discrimination. Akin to the personal approaches, great care is typically taken to ensure that the applications are similar across several dimensions except the variable of interest.

It strikes us as unusually difficult to explore whether, and to what extent, race or gender influence the jobs people receive, or the wages they secured, if one had to receive informed consent from the discriminating employer. For such purposes, it makes sense to consider executing a natural field experiment. This does not suggest that in the pursuit of science, moral principles should be ignored. Rather, in those cases Local Research Ethics Committees and Institutional Review Boards (IRBs) in the US serve an important role in weighing whether the research will inflict harm, gauging the extent to which the research benefits others, and determining whether experimental subjects selected into the environment on their own volition and are treated justly in the experiment.

1.2.5. Limits of field experiments

Clearly, labor economists rarely have the ability to randomize variables directly related to individual decisions such as educational attainment, the choice to migrate, the minimum wage faced, or retirement ages or benefits. This might in part reflect why some active research areas in labor economics have been relatively untouched by field experiments, as described in more detail below. However, field experiments allow the researcher scope to randomize key elements of the economic environment faced that determine such outcomes. For example in the context of educational attainment, it is plausible to design field experiments that create random variation over the monetary costs of acquiring education, information on the potential returns to education, knowledge of

the potential costs and benefits of education, or changes in the quality of inputs into the educational production function. Given the early and close involvement of researchers and the fact that primary data collection effort is part of a field experiment, there is always the potential to mitigate measurement error and omitted variables problems that are prevalent in labor economics (Angrist and Krueger, 1999).

Social experiments and field experiments are relatively easy for policy makers to understand. When designed around the evaluation of a particular policy or intervention, it is more straightforward to conduct a cost benefit analysis of the policy than would be possible through other empirical methods. As discussed before, a concern of using social experiments relates to sample attrition. While such attrition is less relevant in many field experiments, it is important to be clear that this often comes at the cost of field experiments evaluating relatively short run impacts of any given intervention. How outcomes evolve over time—in the absence of the close scrutiny of the experimenter, or how interventions should be scaled up to other units and other implementers, remain questions that field experimenters will have to always confront directly. Along these lines, we will showcase a number of field experiments in which researchers have combined random variation they have engineered to identify reduced form causal effects, with structural modeling to make out of sample predictions.

A second broad category of concerns for field experiments relate to sample selection. These can take a number of forms relating to the non-random selection of individuals, organizations, and interventions. At the individual level and in cases in which written consent is required, as for social and laboratory experiments, the self selection of individuals into the field experiment needs to be accounted for. Relatedly, the timing of decisions over who is potentially eligible to participate are critical, and potentially open to manipulation or renegotiation.

At the organizational level, there exists concerns related to whether we observe a non-random selection of organizations, or practitioners self-select to be subject to a field experiments. Similar concerns arise for social experiments often from political economy considerations.

Finally, at the intervention level, a concern is that practitioners, with whom field experimenters typically need to work, might only be willing to consider introducing interventions along dimensions they *a priori* expect to have beneficial effects. On the one hand this begs the question of why such practices have not been adopted already. On the other hand, one benefit of field experimentation might be that through closer ties between researchers and practitioners, the latter are prompted to think and learn about how they might change their behavior in privately optimal ways, and can be assured they will be able to provide concrete evidence of any potential benefits of such changes.

A third category of concerns relate to how unusual is the intervention. Although many parameters can be experimentally varied, it is important to focus on those parameters that would naturally vary across economic environments, and to calibrate

the magnitude of induced variations based on the range of parameter values actually observed in similar economic environments. Introducing unusual types of variation, or variations of implausible or unusual magnitude, or those that do not accord with theory, will be hard to make generalizations from and will not easily map back to an underlying theory. At the very least, care needs to be taken to separately identify whether responses to interventions reflect changes in equilibrium behavior that will persist in the long run, or agent's short run learning how to behave in new or unusual circumstances induced by the experimenter.

Fourth, there can sometimes be concerns that the third parties researchers collaborate with, might be under resource constraints that lead to the same set of implementers simultaneously or sequentially dealing with treated and control populations. Such implementation might lead to contamination effects and some of the other biases discussed above in relation to social experiments. This can lead to the use of within subject designs, where the researchers engineer an exogenously timed change to the economic environment, rather than between subject designs. The field experiments we discuss in later sections utilize both approaches.

Taken together, most of these concerns can be summarized as relating to the “external validity” of any field experiment—namely the ability to extrapolate meaningfully outside of the specific economic environment considered. This feature remains key to the worth of many field experiments. Field experiments almost inevitably face a trade-off between understanding the specifics of a given context and the generalizability of their findings. This trade-off can be eased by implementing a field experiment that considers the sources of heterogenous effects, or that combines reduced form estimates based on exogenous variation with structural modelling to predict responses to alternative interventions or to the same intervention in a slightly different economic environment.

Finally, it is worth reiterating that although primary data collection is a key element of field experimentation, this raises the costs of entry and might limit the number of experimenters relative to other purely lab based approaches. As will be apparent in the remainder of this chapter, there remain many issues in labor economics in which field experiments have yet to penetrate. In part these limits might be due to lack of opportunities, in some cases it might be because the activities under study are clandestine or illegal, although we will discuss carefully crafted field experiments to explore issues of racial discrimination for example. However, in some cases it is because the nature of the research question is simply not amenable to field experimentation. For example, questions relating to the design of labor market institutions are likely to remain outside the realm of field experimentation. In these and other cases, the controlled environment of the laboratory is the ideal starting point for economic research. Indeed, in this volume, [Charness and Kuhn \(2011, 2010\)](#) discuss the large laboratory-based literature on multiple aspects of the design of labor markets—such as market clearing mechanisms and contractual incompleteness. More generally, they discuss in detail the relative merits

of laboratory and field experiments. We share their view that no one research method dominates the other, and that in many scenarios using a combination of methods is likely to be more informative.

1.3. Research in labor economics

An enormous range of research questions are addressed by labor economists today. While the core issues studied by labor economists have always related to labor supply, labor demand, and the organization of labor markets, to focus our discussion, we limit attention to a select few topics. These reflect long-standing traditional areas of work in labor economics.²⁶

First, since the seminal contributions of Gary Becker and Jacob Mincer, research in labor economics, particularly related to labor supply, has placed much emphasis on understanding individual decision making with regards to the accumulation of human capital. This emphasis has widened the traditional purview of labor economists to include all decision making processes that affect human capital accumulation. These decisions are as broad as those taken in the marriage market, within the household, and those on the formation of specific forms of human capital such as investments into crime. By emphasizing the role of individual decision making, subfields in labor related to the accumulation of human capital might be especially amenable to the use of field experiments.

Second, the empirical study of labor demand has been similarly revolutionized by the rapid increase in the availability of panel data on individuals, the personnel records of firms, and matched employer-employee data.²⁷ This has driven and fed back into research on various aspects of labor demand such as labor mobility, wage setting, rent sharing, and more generally, on the provision of incentives within organizations. This set of questions are again motivated by understanding the behavior of individuals and firms, there are rich possibilities to advance knowledge in related subfields through the use of carefully crafted field experiments. Field experiments offer the potential for researchers to lead data collection efforts.

To cover these broad areas, we loosely organize the discussion so as to follow an individual as they make important labor related decisions over their life cycle. Hence we discuss the role of field experiments in answering questions relating to early childhood interventions and the accumulation of human capital in the form of education and skills. We then consider research questions related to the demand and supply of labor, and labor market discrimination. We then move on to consider research questions related to

²⁶ More detailed discussions of how the study of labor economics has evolved over time can be found in [Freeman \(1987\)](#) and [Taber and Weinberg \(2008\)](#).

²⁷ To understand the magnitude of this change, we note that [Stafford \(1986\)](#) finds that among the 759 papers published in six leading journals between 1965 and 1983, virtually none was based on microdata with individual firms or establishments as the unit of analysis.

behavior within firms: how individuals are incentivized within firms, and other aspects of the employment relationship. Finally, we end with a brief discussion of the nascent literature on field experiments related to household decision making.

Table 2 shows the number of published papers in selected subfields of labor economics in the decade prior to the last volume of the *Handbook of Labor Economics* (1990–99), and over the last decade (2000–09). The table is based on all published papers in the leading general interest journals of *The American Economic Review*, *Econometrica*, *The Journal of Political Economy*, *The Quarterly Journal of Economics* and the *Review of Economic Studies*.²⁸ We use the *Journal of Economic Literature* classifications to place journal articles into one subfield within labor economics.²⁹

Table 2 highlights a number of trends in published research in labor economics. First, the number of labor economics papers published in the top-tier general interest journals has not changed much over time. There were 278 published between 1990 and 1999, and 315 published between 2000 and 2009. Some of this increase probably reflects an increased numbers of papers in these journals as a whole, rather than changes in the relative importance of labor economics to economists. Examining the data by subfield, we do see changes in the composition of published papers in labor. There are large increases in the number of papers relating to: (i) education and the formation of human capital; (ii) firm behavior and personnel economics; (iii) household behavior; (iv) crime. Some of these increases reflect the wider available of data described above, such as personnel data from firms and matched employer–employee data sets, and primary data collected on households. Field experiments—an important component of which is primary data collection—are well placed to reinforce these trends. Indeed, below we discuss how field experiments have contributed to the first three of these areas in which there has been an increase in labor economics papers.

We observe a decline in papers on the organization of labor markets—an area in which not many field experiments have been conducted, in part because these questions are not well suited to field experimentation. Finally, the remaining subfields on the demand and supply of labor and on ageing and retirement remain relatively stable over the last two decades, and here field experiments remain scarce, but there might be particularly high returns from such research designs being utilized.

Second, the balance between theoretical and empirical work has remained relatively constant over the two decades. In both time periods, there have been approximately double the number of empirical as theoretical papers published in labor economics. We do not know whether for other areas of economics approximately a third of published papers are theoretical, but as will be emphasized throughout, labor economics has no shortage of theories that carefully designed field experiments can help determine the

²⁸ The numbers do not include papers and proceedings volumes.

²⁹ Earlier reviews of trends in published papers in labor economics include Stafford (1986), Manser (1999), and Moffitt (1999).

Table 2 Published research in labor economics by decade.

Subfield	Year of publication: 1990-99			Year of publication: 2000-09		
	Theory	Empirics	Total	Theory	Empirics	Total
Demand for education/formation of human capital	7	29	36	10	46	56
The demand and supply of labor	10	58	68	20	53	73
Organization of labor markets	28	59	87	20	48	68
Firm behaviour/personnel economics	26	23	49	33	28	61
Household economics	8	20	28	12	27	39
Aging and Retirement	0	6	6	0	6	6
Crime	2	2	4	1	11	12
Total	81	197	278	96	219	315

The table is based on all published papers in the leading general interest journals of the American Economic Review, Econometrica, the Journal of Political Economy, the Quarterly Journal of Economics, and the Review of Economic Studies. The numbers do not include papers and proceedings volumes. We use the Journal of Economic Literature classifications to place journal articles into one subfield within labor economics.

empirical relevance of. Within each subfield there are nearly always more empirical papers published than theoretical, with the exception of research into firm behavior and personnel economics, a pattern that holds across both decades. In other subfields, the ratios of theory to empirical papers vary considerably. Some areas such on the demand for education and formation of human capital have four to five times as many empirical papers, and the subfield of crime has been largely empirically driven.

1.3.1. How have labor economists used field experiments?

Table 3 presents evidence on the approach used by published papers in labor economics over the last decade.³⁰

Three factors stand out. First, field experiments have been widely used in labor economics over the past decade, with there being 25 published papers using this research methodology in some form. For example, despite the surge in papers using laboratory experiments, over the last decade more papers published in the top-tier journals have employed field experiments. However, the number of empirical papers employing field experiments is still dwarfed by other empirical methodologies—there are 25 papers employing field experiments compared to 60 utilizing natural experiments, and 129 using non-experimental methods.

Second, the use of field experiments has thus far been concentrated to address research questions in a relatively small number of subfields in labor economics. Of the 25 published field experiments, three framed field experiments have been concerned with investments into education early in the life cycle, three natural field experiments have focused on the evaluation of specific labor market programs, and five natural field experiments have focused on incentives within firms.

In other subfields, such as on the determinants of wages and labor market discrimination, currently only one field experiment has been published, in contrast to four laboratory experiments. We view many research questions on discrimination in labor markets to be particularly amenable to study using field experiments. Hence this is one area in which field experiments have been relatively under supplied. Finally, the subfield of crime, which as documented in Table 2, has grown due almost exclusively to empirical papers, remains completely untouched by field experiments.

The third major fact to emerge from Table 3 is that there is a large supply of theory in labor economics, as previously noted in Table 2. Table 3 shows that this supply of theory is across all the subfields in labor economics. As we view carefully crafted field experiments to be able to potentially test between different theories, it would seem as if many areas of study of labor—across the life cycle from birth to retirement—are amenable to this method, and can give feedback on directions for future theoretical advancements.

³⁰ The total number of papers reported in Table 2 is not quite reflected in the totals recorded in Table 3. This is because in Table 3 we sometimes record a paper in more than one column if it utilizes a range of empirical techniques. For example, the total number of non-theory papers by subfield and method in Table 3 is greater than total non-theory papers found in Table 2 (224 > 219) because five papers used multiple methods and so were counted twice.

Table 3 Published papers 2000-9 by subfield and empirical method.

	Theoretical	Non experimental	Natural experiments	Natural field experiments	Framed field experiment	Artefactual field experiment	Laboratory
Demand for education/formation of human capital							
Early childhood interventions on human capital accumulation	2	3	3	1	0	0	0
(RCT) conditional cash transfer programs	1	2	1	1	1	0	0
Educational production function	3	6	9	0	3	0	0
Educational spillovers	1	3	2	0	0	0	0
Returns to education	3	8	3	0	0	0	0
The demand and supply of labor							
Wage and tax sensitivities	6	7	1	1	0	0	0
Determinants of wages/discrimination	12	28	4	1	0	0	4
Segmented labor markets	0	1	2	0	0	0	0
Demand for labor/skills	2	3	0	0	0	0	1

(continued on next page)

Table 3 (continued)

	Theoretical	Non experimental	Natural experiments	Natural field experiments	Framed field experiment	Artefactual field experiment	Laboratory
Organization of labor markets							
Unions, minimum wages and other labor market institutions	1	7	3	0	0	0	1
Labor market programs	0	1	3	3	0	0	0
Public sector labor markets	0	0	0	0	0	0	0
Occupational choice/intergenerational mobility/labor market segmentation	7	9	0	0	0	0	2
Immigration	0	6	2	0	0	0	0
Unemployment	12	13	3	0	0	0	0
Firm behavior/personnel economics							
Employee and executive incentives	16	4	4	5	0	0	2
The employment relationship/gift exchange	1	0	0	2	0	1	0
Peer effects	1	2	2	0	0	0	0
Workplace organization	15	4	0	2	0	0	0

Table 3 (continued)

	Theoretical	Non experimental	Natural experiments	Natural field experiments	Framed field experiment	Artefactual field experiment	Laboratory
Household economics							
Family size	1	3	2	0	0	0	0
Marital bargaining	2	0	0	0	1	0	0
The marriage market	6	4	3	1	0	0	0
Child labor	2	2	0	0	0	0	0
Household labor supply	0	2	1	0	0	0	0
Female participation	1	3	5	0	0	0	0
Retirement							
Decision to retirement	0	2	1	1	0	1	0
Health and retirement	0	1	0	0	0	0	0
Crime							
Crime	1	5	6	0	0	0	0
Total	96	129	60	18	5	2	10

The table is based on all published papers in the leading general interest journals of the American Economic Review, Econometrica, the Journal of Political Economy, the Quarterly Journal of Economics, and the Review of Economic Studies. The numbers do not include papers and proceedings volumes. We use the Journal of Economic Literature classifications to place journal articles into one subfield within labor economics.

To develop this point further, [Table 4](#) provides a breakdown of how theory and evidence have been combined in labor economics, broken down by empirical method.

Two factors stand out. First, non-experimental papers are slightly more likely to use no theory than field experiments. Second, testing between more than one theory remains scarce, irrespective of the empirical approach. Although not all empirical papers should necessarily test theory, it is as important to establish facts on which future theory can be built. When testing between theories, it is important to both establish the power of these tests, to provide refutability or falsification checks, and to present evidence of the internal validity of the results. Natural field experiments might have a comparative advantage along such dimensions. Given such settings relate to real world behaviors, individuals are typically not restricted in how they respond to a change in their economic environment, which opens up the possibility of detecting behavior consistent with multiple theories.

Mirroring the discussion in [Moffitt \(1999\)](#), a second feature of on how best to use field experiments, that we aim to emphasize throughout, is the need to combine the use of field experiments with other research methodologies. For example, they might be combined with structural estimation, utilize a combination of evidence from the laboratory and the field, or draw inspiration from lab findings to establish plausible null and alternative hypotheses to be tested between.

Applying the full spectrum of approaches in trying to answer a single question can yield extra insights. A first example of such research relates to the importance of social preferences, which have been documented in numerous lab and field settings. To explore social preferences using a variety of approaches, [List \(2006\)](#) conducts artefactual, framed, and natural field experiments analyzing gift exchange. The games have buyers making price offers to sellers, and in return sellers select the quality level of the good provided to the buyer. Higher quality goods are costlier for sellers to produce than lower quality goods, but are more highly valued by buyers. The artefactual field experimental results mirror the typical findings with other subject pools: strong evidence for social preferences was observed through a positive price and quality relationship. Similarly constructed framed field experiments provide similar insights. Yet, when the environment is moved to the marketplace via a natural field experiment, where dealers are unaware that their behavior is being recorded as part of an experiment, little statistical relationship between price and quality emerges.

A second example comes from the series of field experiments presented in [List \(2004b\)](#)—from artefactual to framed to natural—in an actual marketplace to help distinguish between the major theories of discrimination: animus and statistical discrimination. Using data gathered from bilateral negotiations, he finds a strong tendency for minorities to receive initial and final offers that are inferior to those received by majorities in a natural field experiment. Yet, much like the vast empirical literature documenting discrimination that exists, these data in isolation cannot pinpoint the nature of discrimination. Under certain plausible scenarios, the results are consonant with at

Table 4 Testing theory by empirical method, published papers 2000-9.

Method	Research method	Theoretical	Non Experimental	Natural experiments	Natural field experiments	Framed field experiments	Artefactual field experiments	Laboratory
Pure theory	Theory	96	0	0	0	0	0	0
Theory and evidence	Test one theory Test between theories	0	61	5	4	0	0	1
Pure evidence	Pure empirics	0	58	49	9	4	1	4
		96	129	60	18	5	2	10

The table is based on all published papers in the leading general interest journals of the American Economic Review, Econometrica, the Journal of Political Economy, the Quarterly Journal of Economics, and the Review of Economic Studies. The numbers do not include papers and proceedings volumes. We use the Journal of Economic Literature classifications to place journal articles into one subfield within labor economics.

least three theories: (i) animus-based or taste-based discrimination, (ii) differences in bargaining ability, and (iii) statistical discrimination. By designing allocation, bargaining, and auction experiments, List (2004b) is able to construct an experiment wherein the various theories provide opposing predictions. The results across the field experimental domains consistently reveal that the observed discrimination is not due to animus or bargaining differences, but represents statistical discrimination.

1.4. Summary

We now move to describe, by various stages of the life cycle, how field experiments have been utilized in labor economics and the insights they have provided. Where appropriate, we discuss how these results have complemented or contradicted evidence using alternative research methods, and potential areas for future field experiments. We begin with the individual at birth and the accumulation of human capital before they enter the labor market. We then consider research questions related to the demand and supply of labor, and labor market discrimination. We then move on to consider research questions related to behavior within firms: how individuals are incentivized within firms, and other aspects of the employment relationship. Finally, we end with a brief discussion of the nascent literature on field experiments related to household decision making.

2. HUMAN CAPITAL

The literature associated with human capital acquisition prior to labor market entry is vast, and there is no room here to do it justice. As in the sections that follow, we therefore focus our discussion on a select few strands of this research and describe how field experiments can and have advanced knowledge within these strands. Even within this narrower branch of work, we are limited to focusing on select studies of inputs into the educational production function, where these inputs might be supplied by the school system, students, or their families.³¹

To see the issues, we follow Glewwe and Kremer's (2006) presentation of a framework for the education production function with the following reduced form representation,

$$S = f(C, H, Q, P), \quad (6)$$

$$A = h(C, H, Q, P), \quad (7)$$

where S is years of schooling, A is skills learned (achievement), C is a vector of child characteristics (including "innate ability"), H is a vector of household characteristics, Q is a vector of school and teacher characteristics (quality), and P is a vector of prices related to schooling. Q and P are both functions of education policies (EP) and local

³¹ For related reviews of the literature, see the excellent work of Card (1999), and in this Handbook, the chapter by Fryer (2011).

community characteristics (L), which can be substituted into Eqs (6) and (7) to yield the following reduced form,

$$S = f(C, H, L, EP), \quad (8)$$

$$A = h(C, H, L, EP). \quad (9)$$

Similar to Mincerian human capital earnings functions, this framework estimates the partial equilibrium effects of educational inputs and policies, rather than general equilibrium effects that alter returns to education and thereby demand. Broadly, there are two approaches to estimating the production function.

The first focuses on measuring the effect of direct inputs, such as per pupil expenditure, class size, teacher quality and family background (i.e. estimating Eqs (6) and (7)). The second examines the effects of educational policies governing the structure of the school system (i.e., estimating Eqs (8) and (9)). In both cases, non-experimental and experimental estimates have shed insights into the relationships in the education production function (for a literature survey see [Hanushek \(1986\)](#)). To help place field experiments in this area in a wider context, we now turn to a non-exhaustive discussion of select work using both approaches, but not based on field experiments.

2.1. Measuring the effects of direct inputs

An early measurement study focusing on the effect of direct inputs is the report due to [Coleman et al. \(1966\)](#), who explored what fraction of the variation in student achievement could be explained by direct inputs. The Coleman report found only a weak association between school inputs and outputs. Subsequent regression based approaches largely replicated the findings in the Coleman report. Yet, one remarkably consistent result did emerge from these early studies: students situated in classrooms with a larger number of students outperformed children in smaller classes on standardized tests. This result is robust to inclusion of several conditioning variables, such as key demographic variables.

One aspect that this robust empirical finding highlights is the care that should be taken to ensure reverse causality and omitted variable bias do not frustrate proper inference. Given that the simple regression approach potentially suffers from biases due to endogeneity of policy placement, omitted variables, and measurement error—i.e., it is almost always the case that some unobserved element of the vectors C , H , Q , P or L will be correlated with both the outcome and the observed variables of interest—researchers have sought out other means to explore the parameters of the production function.

One such approach uses natural experiments. One neat example is the work of [Angrist and Lavy \(1999\)](#), who use legal rules to estimate the effect of class size on student performance. Assume legal limits on class size prevent the number of students in a

classroom from exceeding 25. Then consider a particular school that has cohorts ranging from 70–100. Thus, if a cohort includes 100 children, we would have four classrooms of size 25, whereas if the cohort includes 76 children, we end up with 4 classrooms with 19 children occupying each. Angrist and Lavy (1999) compare standardized test scores across students placed in different sized classrooms and find that a ten-student reduction raises standardized test scores by about 0.2 to 0.3 standard deviations.

As Keane (forthcoming) points out, this type of approach has similar drawbacks associated with the simple regression framework in the Coleman report. For instance, incoming cohort sizes might not be determined randomly because high performing schools attract more students. Likewise, cohort size might be affected by parents reacting to large class sizes by sending their kids elsewhere for schooling. Similar issues revolve around teacher assignment to small and large classrooms, which might not be randomly determined.

In this way, the Angrist and Lavy (1999) estimates should be viewed as a first step in understanding the importance of class size on student performance. The next step is to deepen our understanding by exploring the robustness of these results. One approach is to look for more observational data, another is to use randomization directly—similar to the accidental randomization of natural experiments, purposeful randomization can aid the scientific inquiry.

A central figure in using randomization in the area of education is William McCall, an education psychologist at Columbia University who, at odds with his more philosophical contemporaries, insisted on quantitative measures to test the validity of education programs. For his efforts, McCall is credited as an early proponent of using randomization rather than matching as a means to exclude rival hypotheses, and his work continues to influence the field experiments conducted in education today.³²

A landmark social experiment measuring the effects of classroom size is the Tennessee STAR experiment. In this intervention, more than 10,000 students were randomly assigned to classes of different sizes from kindergarten through third grade. Similar to the social experiments discussed in the first Section, the STAR experiment had both attrition bias and selection problems in that some students changed from larger to smaller classrooms after the assignment had occurred. Nevertheless, even after taking these problems into account, Krueger (1999) put together a detailed analysis that suggests there are achievement gains from studying in smaller classes.

Combined, these two examples indicate that there very well might be a statistically meaningful relationship between class sizes and academic achievement, but the broader literature has not concluded that to be necessarily true. Scanning the entire set of

³² Rockoff (2009) presents an overview of a substantial, but overlooked, body of field experiments class size that developed prior to World War II.

estimates from natural experiments and field experiments one is left with mixed evidence on the effects of class size at various tiers of the education system (Angrist and Lavy, 1999; Case and Deaton, 1999; Hoxby, 2000a; Kremer, 2003; Krueger, 2003; Hanushek, 2007; Bandiera et al., 2010, forthcoming).

Lazear (2001) theorizes that class size is dependent upon the behavior of students. As disruptive students are a detriment to the learning of their entire class, he proposes that the optimal class size is larger for better-behaved students. In his model, larger classes may be associated with higher student achievement, and may in part explain the mixed results in previous studies. This is one area where a natural field experiment might be able to help. One can envision that a test of Lazear's (2001) theory is not difficult if the researcher takes the data generation process in her own hands: designing experimental treatments that interact class size with student behavior would permit an estimation of parameters of interest for measures of both class size and peer inputs into the educational production function.

The results from this literature, more generally, make it clear how one could move forward with a research agenda based on field experimentation. For instance, are there critical non-linearities in the relationship between class sizes and academic performance, as suggested for university class sizes in Bandiera et al. (2010, forthcoming)? One might argue that the effects of smaller class sizes drop to zero at some critical threshold due to lost peer effects. What about the composition of classrooms? Even though the effects of peer composition are mixed (Hoxby, 2000b; Zimmerman, 2003; Angrist and Lang, 2004; Hoxby and Weingarth, 2006; Lavy et al., 2008; De Giorgi et al., 2009; Duflo et al., 2009), it might be the case that gender balance plays a key role in the classroom.

Even if we were to find strong evidence that class size matters for academic performance and answer the questions posed above, Eqs (6)–(9) highlight other features that we must be aware of before pushing such estimates too far. What is necessary is proper measurement of the estimates of the parameters of the production function, as well as an understanding of the decision rules of school administrators and parents. The next step is to deepen our understanding by exploring whether other more cost effective approaches to improve student achievement exist, say by understanding the optimal investment stream in students: at what age level are resources most effective in promoting academic achievement?

One line of work that addresses this question is the set of social experiments that explore achievement interventions before children enter school. Given that Fryer (2011) presents a lucid description of such interventions, we only briefly mention them here. The landmark social experiment in this area is the Perry Preschool program, which involved 64 students in Michigan who attended the Perry Preschool in 1962. Since then, dozens of other programs have arisen that explore what works with early childhood

intervention, including Head Start, the Abecedarian Project, Educare, Tulsa's universal pre-kindergarten program, and several others too numerous to list (see Fryer's Table 5).³³

As Fryer (2011) notes, outcomes in these programs exhibit substantial variance. And, even in those cases that were met with great success, the achievement gains faded through time. Indeed, in many cases once school started the students in these programs gave back all academic gains (Currie and Thomas, 1995, 2000; Anderson, 2008). Another fact with the bulk of these programs is that they exhibit much homogeneity, mostly following from the general design in the Perry Preschool program. Much has been learned about early childhood development in the previous several decades, and this presents the field experimenter with a unique opportunity to make large impacts on children's lives. As Fryer (2011) notes, incorporating new insights from biology and developmental psychology represent opportunities for future research.

Such estimates cause us to pause and ask whether resource expenditures affect academic performance at all. In this spirit, there is a large literature that explores how direct school inputs, such as school expenditures, influence student performance. As a whole, the early literature found only a weak relationship between overall school expenditures and student achievement, primarily because resources tend to be allocated inefficiently (see Hanushek, 2006; Glewwe and Kremer, 2006 for a review of the recent literature). In response to these findings, a growing area of research uses both natural experiments and field experiments to examine a wide range of targeted investments in order to identify the effects and compare the cost-effectiveness of various interventions. For example, experiments have been carefully designed to identify the returns (in terms of schooling or achievement) to inputs such as school supplies, additional teachers, remedial education, or computer programs (Banerjee et al., 2001; Angrist and Lavy, 2002; Kremer et al., 2002; Glewwe et al., 2004, 2003; Banerjee et al., 2007).

2.2. Teacher quality

While the evidence on the effect of per pupil expenditure, class size, and peer composition is mixed, teacher quality has been found to be clearly important. Hanushek (2007) finds that the differences between schools can be attributed primarily to teacher quality differences. Little of this variation, however, can be explained by either teacher salaries or observable characteristics such as education and experience (Rivkin et al., 2005; Hanushek, 2006). Just as it is difficult to identify high quality teachers, little is

³³ There is evidence that the first five years of life are critical for lifelong development. Hence resource poor or unstimulating environments early in life are likely to detrimentally impact children's cognitive, motor, social-emotional development, and their health status (Grantham-McGregor et al., 1991; Heckman and Masterov, 2005; Engle et al., 2007; Grantham-McGregor et al., 2007). As adults, they are more likely to have high fertility rates and are less likely to provide adequate stimulation and resources for their own children, thus contributing to the intergenerational transmission of poverty and economic inequality (Sen, 1999). The current debate, to which social experiments have contributed, focuses on understanding the types of intervention that might be effective for the child and their families, and cost-effective from society's viewpoint.

known about how to improve teacher quality and performance. Given the evidence that education and professional development are largely ineffective, there is a growing interest in the use of performance-based incentives to improve teacher quality and effort.

The design and implementation of such incentives raises several areas of future study that observational data and field experiments can adequately fill, including: (i) what are the performance effects on the incentivized tasks and how can incentives be designed to cost-effectively maximize these effects; (ii) what are the effects on non-incentivized tasks, and how can incentives be designed to avoid diversion of effort in multitasking; (iii) how do teachers (of different quality) sort into different incentive and pay structures; and (iv) how does sorting affect general equilibrium teacher quality.

Evidence from non-experimental studies, natural experiments, and field experiments suggests that incentives can improve teacher performance (Lavy, 2002; Glewwe et al., 2003; Figlio and Kenny, 2007; Muralidharan and Sundararaman, 2007; Lavy, 2009; Duflo et al., 2009). Clearly, tighter links can be established between this literature and the larger labor literature on incentive design (Prendergast, 1999) to which, as discussed below, field experiments are also beginning to contribute.

More broadly, field experiments exploring mechanism design issues, such as comparing piece rate and tournament incentives, are rare. Also, these programs generally load incentives onto a single performance measure such as teacher attendance or student test scores, raising concerns that teachers might divert effort away from non-incentivized tasks (Holmstrom and Milgrom, 1991). Here, the evidence is mixed with some studies finding broad improvements in teacher effort (Duflo et al., 2009; Lavy, 2009) and others finding evidence of narrow efforts, such as teaching to the test, that divert effort from other tasks and do not improve long term student achievement (Kremer, 2006; Jacob, 2005).

Similarly, teacher sorting into incentive and pay structures is largely unexplored. Lazear (2001) applies his analysis of performance pay and productivity in a company (discussed in further detail below) to teacher incentives, suggesting that the effects of incentives on sorting could be comparable to effects on teacher effort. A well-designed field experiment could explore whether and how teacher sorting on incentives occurs. In general, field experiments that apply theories about incentive design, sorting and selection from other areas of labor could make a large contribution to the teacher incentives literature. Many of these issues arise in a later section when we discuss the role of field experiments in understanding behavior within firms.

Along with the school inputs, the primary inputs into the educational production function come from students and their families. A large literature models the effect of individual characteristics, family background and parental resources on schooling and achievement (Cameron and Heckman, 2001; Cameron and Taber, 2004). While it is impossible to randomly assign characteristics to individuals or to randomly assign children to families, quasi-experimental studies have exploited variation due to adoption in order

to separately identify genetic inputs (“nature”) from parental inputs (“nurture”) (Plug and Vijverberg, 2003). Other studies focus on potential barriers to individual investment in human capital production. These include high costs to education, perhaps due to credit constraints or high discount rates, and low marginal returns to education due, for example, to poor health or lack of human capital investment prior to entering school.

Estimates from non-experimental studies and natural experiments suggest that credit constraints are of limited importance in schooling decisions (Cameron and Taber, 2004; Stanley, 2003). However, estimates from natural experiments and field experiments of Conditional Cash Transfer programs (CCTs) find largely positive and significant effects, suggesting that (at least among the population targeted by CCTs) reducing present costs of education can affect schooling and human capital investment decisions. Non-experimental studies, natural experiments and field experiments have also found positive and significant effects from conditional cash transfer programs based on enrollment, attendance, and performance (Cornwell et al., 2006; Barrera-Osorio et al., 2008; Angrist and Lavy, 2009; Maxfield et al., 2003; Kremer et al., 2009). Few of these experiments, however, explore how conditional cash transfers can be most effectively designed.

Berry (2009) develops a model of household education production in which parents’ ability to motivate their children is dampened by moral hazard. He then designs incentives to test several predictions of the model including the ability of parents to commit and the relative efficacy of incentives awarded to parents or to children based on the relative productivity of the two parties. Similarly, Levitt et al. (2010) implement a field experiment that compares both the incentive recipient (parent or student) and the incentive mechanism (piece rate or lottery). They also compare a year long broad-based incentive program that motivates sustained effort to an immediate one-time incentive aimed solely at increasing effort on a single standardized test. This design allows the authors to test a model of family investment, responsiveness to incentive mechanisms, and human capital returns from varying levels of effort. Both of these field experiments illustrate that researchers can design instruments that build on and test economic theory.

While conditional cash transfers aim to induce improvements in achievement by motivating greater effort and investment, a second strand of interventions attempts to directly improve abilities that can improve achievement. A growing of interest in this area focuses on investment in early childhood. Researchers argue that improving the abilities of young children can have long run returns on educational achievement, attainment and other outcomes such as employment, crime, fertility and health (Cunha and Heckman, 2009). Evidence from non-experimental studies, natural experiments and field experiments suggest that early education interventions can have significant effects on lifetime outcomes (Currie and Thomas, 1995; Currie, 2001; Garces et al., 2002; Behrman et al., 2004; Todd and Wolpin, 2006; Ludwig and Miller, 2007; Heckman et al., 2010).

Most of these studies require econometric techniques, such as matching, to correct for lack of valid randomization. And all of them are limited to identifying the effect of the intervention as a whole. They are not able to explore, for example, the relative importance of educational interventions compared to interventions that increase parental investments in early childhood. Given the evidence that early childhood is a key period of development and the relatively sparse body of empirical work, field experiments could address open questions related to: (i) the short and long run returns of the various inputs of the educational production function; (ii) to collect primary data and design field experiments to help decompose overall changes in outcomes from any given intervention into those arising from the behavioral responses of children, parents and teachers. Akin to the literature on public and private transfers to households (Albarran and Attanasio, 2003), this second strand of research can help shed light on whether altering some inputs leads to other inputs in the educational production function to be crowded in or out.

A final strand of the literature focuses on improving child health as a means of increasing school attendance rates. Estimates from natural experiments and field experiments find that health interventions have a positive and significant effect on school attendance (Bleakley, 2007; Bobonis et al., 2006; Miguel and Kremer, 2004). Miguel and Kremer (2004) expand beyond identification of individual returns to health interventions, modeling the positive externalities of deworming ignored in previous estimations. They use a field experiment randomized over schools to estimate positive externalities on the health and school attendance of untreated children in treated schools and schools neighboring treated schools. They also examine effects on test performance and estimate the health care and educational cost effectiveness of the program. As the authors argue, studies that ignore positive externalities in the comparison groups will underestimate the effect of the intervention by missing the external effects of deworming and underestimating the direct effect in comparison with an inflated baseline, biasing treatment effects towards zero. They point out that this identification problem is well recognized in the labor literature estimating the effects of job training programs on both participants and non-participants. The authors suggest an extension of their study that randomizes treatment at various levels such as within schools, across schools, and within clusters of schools.

2.3. Measuring the effects of policies governing the system

Recent studies of educational policy exploit natural experiments with randomized lotteries and variation in school district density to estimate the effects of school competition, school choice, school vouchers, school accountability and the presence of relatively autonomous public schools, such as charter schools (Clark, 2009; Cullen et al., 2006; Hoxby, 2000c; Jacob, 2004; Rouse, 1998; Angrist et al., 2002; Abdulkadiroglu et al., 2009). While proponents of expanding school choice argue that, as in other

markets, choice and competition will improve overall school quality and efficiency, the empirical studies find somewhat mixed evidence on these educational policies.³⁴

For example, non-experimental studies, natural experiments and field experiments finding that vouchers improve educational achievement include Peterson et al. (2003), Krueger and Zhu (2004), Angrist et al. (2002, 2006). On the other hand, using randomized school lotteries, Cullen et al. (2006) find that school choice programs have little or no effect on academic achievement, and they suggest that this result may be due to parents making poor choices. Hastings and Weinstein (2008) explore this hypothesis using both a natural experiment and a natural field experiment to examine how reducing information costs affects parental choices. In the natural experiment, parents listed their preferences for schools within a district both before and after receiving information mandated by No Child Left Behind (NCLB). The natural field experiment randomized distribution of a simplified version of the NCLB information to parents who had also received NCLB information and to parents who had received no information.

This design allows the authors to measure the effect of each piece of information alone as well as their interaction. They find that information on school-level academic performance pushes parents to choose higher scoring schools (with no differences across the types of information received). Using IV estimation, they also argue that these choices lead to increased academic achievement.

Similarly, a growing body of research has begun to identify the right tail of the distribution of treatment effects among heterogeneous charter schools (Dobbie and Fryer, 2009; Hoxby and Muraka, 2009; Angrist et al., 2010). These studies rely on randomized lotteries in oversubscribed schools and can only identify the effect of a school (or school system) as a whole. They have reported suggestive evidence, however, on specific features that correlate with successful schools, such as longer days, longer school years, highly academic environments and so on. Field experiments could be used to complement this work by separately identifying the effects of charter school innovations, such as length of school day, school time, and general environmental conditions on the educational production function.

The field experiments discussed in this section highlight several important advantages of their usage for labor economists. For example, they can address biases in previous empirical estimates, including those from non-experimental studies. They are able to build in empirical and theoretical literature from several fields, such as education, health, and labor. Finally, they can be used to identify parameters beyond the direct return of an input into an individual educational production function and explore mechanism design issues.

³⁴ Several theoretical papers suggest that school vouchers will lead to overall welfare gains, increased stratification, and efficiency gains (Epple and Romano, 1998; Ferreyra, 2007; Nechyba, 2000; Rouse, 1998; Figlio and Rouse, 2006; Hsieh and Urquiola, 2006; Epple et al., 2006; Arcidiacono, 2005).

In the end, it is clear that empirical explorations into human capital acquisition prior to labor market entry is invaluable, and that there are several approaches that can be used in concert to learn more about the important parameters of interest. We argue that in this area field experiments can usefully add to the knowledge gained from naturally-occurring data, and the many low apples that are left to be picked give us great confidence that field experiments will only grow in importance in tackling particulars in the educational production function.

3. LABOR MARKET DISCRIMINATION

Philosophers as far removed as Arcesilaus, Heraclitus, and Plato have scribed of injustice and extolled upon the virtues of removing it for the betterment of society. Perhaps taking a lead from these scholars, social scientists have studied extensively gender, race and age based discrimination in the marketplace. In this section we explore the stage of the life cycle where individuals are entering the labor market. We focus mainly on discrimination in labor markets and how field experiments can lend insights into this important social issue.

We begin with a statistical overview of the data patterns in labor market outcomes across minority and majority agents. To make precise how field experiments might be carefully designed, we need to discuss theories for why such discrimination exists. The two major economic theories of discrimination that we discuss are: (i) certain populations having a general “distaste” for minorities (Becker, 1957) or a general “social custom” of discrimination (Akerlof, 1980); (ii) statistical discrimination (Arrow, 1972; Phelps, 1972), which is third-degree price discrimination as defined by Pigou (1920)—marketers using observable characteristics to make statistical inferences about productivity or reservation values of market agents.

Empirically testing for marketplace discrimination has taken two quite distinct paths: regression-based methods and field experiments. The former technique typically tests for a statistical relationship between an outcome measure, such as wage or price, and a group membership indicator. By and large, regression studies find evidence of discrimination against minorities in the marketplace.³⁵ Field experimental studies, which have arisen over the past 35 years, typically use matched pairs of transactors to test for discrimination. Due to the control that field studies offer the experimenter, they have become quite popular and have by now been carried out in at least ten countries (Riach and Rich, 2002). Across several heterogeneous labor markets, as well as product markets as diverse as home insurance and new car sales, field studies have made a strong case that systematic discrimination against minorities is prevalent in modern societies.

³⁵ A comprehensive summary of the regression-based literature on discrimination are contained in Altonji and Blank (1999) and Yinger (1998).

While regression-based empirical studies have served to provide an empirical foundation that indicates discrimination is prevalent in the marketplace, they have been less helpful in distinguishing the causes of discrimination. As Riach and Rich (2002) note, findings from field studies appear to be more consistent with the majority white populations having a general “distaste” for minorities in the sense of Becker (1957) or a general “social custom” of discrimination in line with Akerlof (1980); but statistical discrimination (Arrow, 1972; Phelps, 1972), or marketers using observable characteristics to make statistical inference about productivity or reservation values of market agents, for example, cannot be ruled out, *ex ante* or *ex post*.

Before one can even begin to discuss social policies to address discrimination, it is critical to understand the causes of the underlying preferential treatment that certain groups receive. As has been emphasized throughout, the potential for field experiments to be explicitly designed to test between theories, is a key advantage of this approach over other methodologies. In this section, we provide a framework for how field experiments can be used to advance our understanding of not only the extent of discrimination in the marketplace but also the nature of discrimination observed.

3.1. Data patterns in labor markets

As Altonji and Blank (1999) noted, researchers have observed labor market differences across race and gender lines for decades. Yet, the magnitude of market differences, and hence what a new generation of field experiments seek to explain, has changed substantially over time. For example, there was convergence in the black/white wage gap during the 1960s and early 1970s, but such convergence lost steam in the two decades afterwards. In addition, the Hispanic/white wage gap has risen among both males and females in the 1980s and 1990s. Of course, the world has not remained stagnant since the 1990s, and this section is meant to update the results in Altonji and Blank (1999).

Table 5 presents the labor outcomes of whites, blacks, and Hispanics by gender in 2009. Table 5 includes a set of labor market outcomes by race and gender that labor economists have studied for decades. The data are based on tabulations from the Current Population Survey (CPS) from May 2009. Row 2 of Table 5 indicates that white men earn 13% (21%) more than white women (black and Hispanic men) on an hourly basis. Black and Hispanic women earn less than minority men and majority women.

When one focuses on annual earnings, row 3 of Table 5, the differential between white men continues: they earn more than 20% higher wages than minority men. Yet, for women the racial difference becomes markedly higher—50% for white women to black and 30% for white women to Hispanic. The differentials remain when we focus on full-time employees—rows 7 and 8 of Table 5. In general Table 5 tells a story that has been told often before: white men earn more money for hours worked than other groups, and white women earn more than their female counterparts.

Table 5 Labor market data by race and gender.

	White males	Black males	Hispanic males	White males	Black females	Hispanic females
All workers (2009)						
(1) Share of all workers	0.355	0.054	0.086	0.315	0.061	0.059
(2) Hourly wage	\$17 (10.3)	\$14 (7.4)	\$14 (6.9)	\$15 (8.8)	\$13 (6.8)	\$12 (6.7)
(3) Annual earnings	\$64,642 (52182.5)	\$53,252 (48829.9)	\$50,573 (36977.0)	\$62,293 (55483.4)	\$41,533 (42769.5)	\$40,266 (39587.2)
(4) Weeks worked	48.01 (9.8)	49.29 (8.5)	48.13 (9.4)	47.32 (10.6)	48.03 (9.9)	46.67 (11.9)
(5) Hours worked per week	40.2	39	38.3	35	36.9	34.9
(6) Share part time	0.132	0.139	0.132	0.279	0.188	0.267
Full-time-full year (2009)						
(7) Hourly wage	\$19 (10.3)	\$15 (7.6)	\$14 (7.1)	\$16 (8.1)	\$14 (6.6)	\$13 (6.2)
(8) Annual earnings	\$66,928 (48880.6)	\$56,855 (50287.3)	\$51,517 (36652.8)	\$61,948 (47603.9)	\$44,510 (44894.0)	\$52,756 (41008.6)
All persons						
(9) Share over employed	0.812	0.726	0.829	0.720	0.679	0.695
(10) Unemployment rate (Jan 2010)	11.40%	20.70%	14.00%	7.50%	14.20%	12.30%
(11) Employment rate (Jan 2010)	63.50%	51.50%	66.10%	54.30%	51.20%	50.00%

Standard deviations are in parentheses.

Source: Current population survey, May 2009.




Figure 3 Median weekly earnings of male workers. Year 2000 dollars. (Source: Bureau of Labor Statistics.)

Figures 3 and 4 complement these wage data by showing for each gender, the time series of annual median weekly income from 1969 to present for whites and blacks and 1986 to present for Hispanics.³⁶ These figures bring to light some interesting trends. Regardless of racial or ethnic group, wage rates for women continue to grow faster than for men. Within each gender, though, the 2000s did very little for racial or ethnic differentials. In fact, for both genders any local trend of convergence is reversed by the mid-2000s. In part this could be a function of the well documented rise in wage inequality during the second half of the 2000s.

Another important set of data points on Table 5 is the extent to which whites face lower unemployment rates. Figures 5 and 6 extend this information by showing, by gender, the time series of unemployment rates for whites, blacks, and Hispanics. One interesting element is the magnitude of unemployment changes for whites versus blacks and Hispanics. The mid-2000s saw no change to this trend, with the impact of recessions falling harder on blacks and Hispanics relative to whites. This trend does not seem to depend strongly on gender either, even though neither gender nor any racial or ethnic group seems to be immune from being hit by the 2009/10 recession.

Wages and unemployment rates are a function of labor force participation rates as well. Figure 7 shows the time series of labor force participation. The convergence in participation rates from the 70s through 90s continued into the 2000s, although

³⁶ Weekly earnings figures are taken from the Current Population Survey. They are for all employed people over age 25 that reported weekly earnings above zero. Data before 1979 is taken from the May supplement of the CPS. After 1979 data is taken from the CPS Annual Earnings File. Earnings from the May supplement for 1969–1972 were reported in ranges. The midpoint of each range was assumed to be the actual earnings for each individual.




Figure 4 Median weekly earnings of female workers. Year 2000 dollars. (Source: Bureau of Labor Statistics.)

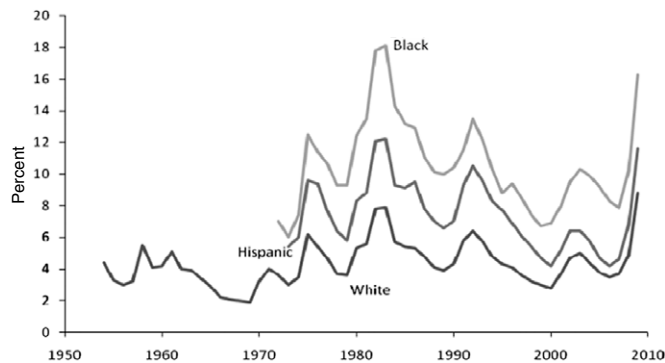


Figure 5 Male unemployment rates (annual averages) for men over 20. (Source: Bureau of Labor Statistics.)

the pace of that convergence has slowed. Men of every race/ethnicity have dropped out of the labor force at a very slow rate while Hispanic and white females have increased participation. Interestingly, African American women have higher labor force participation than white women.

In considering the causes for these labor market disparities, economists have explored whether the workers themselves bring heterogeneous attributes to the workplace. To shed insights into this issue, we provide [Table 6](#), which shows educational differences, family differences, and regional composition.

Rows 2 through 6 in [Table 6](#) shows that whites obtain more years of education than blacks and Hispanics. Interestingly, white women are almost uniformly more educated than their ethnic/racial counterparts. This result is also reflected in the years




Figure 6 *Female unemployment rates (annual averages) for women over 20. (Source: Bureau of Labor Statistics.)*




Figure 7 *Labor force participation rates, 20 years and older. (Source: Bureau of Labor Statistics)*

of experience variable. Rows 8 through 10 in Table 6 give a sense of the different family choices (marriage and fertility) that are made by whites, blacks, and Hispanics—another important input to wages, especially for women. Row 8 shows that whites are more likely to be married (and perhaps enjoy the efficiencies of household trade), but row 10 shows that white women are likely to have fewer children to spend time caring for than black and Hispanic women. Rows 11 through 20 show the geographic breakdown of each race/ethnic group. Local labor market opportunities are surely influential for wages and in general, whites are from higher earning regions like New England and the Pacific.

Table 6 Personal characteristics by race and gender.

	All	White males	Black males	Hispanic males	White females	Black females	Hispanic females
(1) Share of all persons	1.00	0.32	0.05	0.07	0.35	0.07	0.07
Education							
(2) Less than high school	14.25%	9.88%	19.57%	38.58%	8.85%	16.96%	35.23%
(3) High school	30.92%	31.26%	36.86%	31.29%	31.11%	33.40%	28.94%
(4) Some post-HS training	19.66%	19.95%	20.21%	14.25%	20.76%	22.49%	16.77%
(5) College degree	26.11%	28.03%	18.55%	12.47%	29.58%	21.68%	15.82%
(6) More than college	9.07%	10.89%	4.81%	3.44%	9.70%	5.47%	3.24%
(7) Potential experience (Age-educ-5)	27.5 (18.3)	28.1 (17.8)	25.2 (17.5)	21.6 (16.3)	29.6 (18.9)	26.4 (18.3)	23.2 (17.3)
(8) Share married	0.545	0.597	0.409	0.52	0.557	0.3	0.497
(9) No. children age less than 6	0.148 (0.439)	0.128 (0.417)	0.101 (0.363)	0.199 (0.489)	0.139 (0.429)	0.16 (0.445)	0.26 (0.547)
(10) Total no. of children (age < 18)	0.809 (1.120)	0.707 (1.070)	0.804 (1.170)	1.18 (1.250)	0.747 (1.080)	0.985 (1.200)	1.244 (1.220)
(11) Share in SMSA*	0.789	0.751	0.888	0.912	0.752	0.891	0.918
Region							
(12) New England	11.2%	13.5%	4.2%	4.2%	13.7%	3.8%	5.0%
(13) Middle Atlantic	9.6%	9.3%	10.2%	9.7%	9.5%	11.5%	10.6%
(14) East-North Central	11.8%	13.1%	12.2%	6.0%	13.2%	12.5%	5.6%
(15) West-North Central	11.6%	14.1%	5.1%	4.1%	14.0%	4.5%	3.5%
(16) South Atlantic	17.9%	16.2%	38.7%	14.5%	16.5%	39.3%	13.9%
(17) East-South Central	5.0%	5.2%	9.4%	1.4%	5.3%	9.2%	0.9%
(18) West-South Central	8.3%	6.6%	10.6%	17.4%	6.6%	11.0%	17.9%
(19) Mountain	10.2%	10.8%	3.1%	14.4%	10.4%	2.7%	14.3%
(20) Pacific	14.5%	11.1%	6.5%	28.1%	10.8%	5.5%	28.4%

Source: Current population survey, March 2009. Standard deviations in parenthesis.

* Defined as residing in SMSA with at least one million inhabitants.

Overall, these data are in line with Altonji and Blank (1999), who find large educational differences among these groups, with race and ethnicity mattering much more than gender. Of course, what these education differences represent is difficult to parse. On the one hand, they might be mostly due to different preferences. Alternatively, they might reflect behavior of agents who expect to face discrimination later on in the labor market—referred to as “pre-market” discrimination. As Altonji and Blank (1999) note, there is evidence that some minorities have been denied market opportunities, perhaps leading to less than efficient levels of schooling investment.

While the labor market outcomes disparities observed in Table 5 and Figs 3–7 might represent differences mainly due to these individual investment choices, perhaps investment varies because of preferences, comparative advantage, and the like. For example, another hypothesis put forth is that such outcomes are at least partly due to discrimination in the labor market. The remainder of this section briefly discusses theories of discrimination and attempts to test these theories, contrasting regression-based approaches to field experiments.

3.2. Theories of discrimination

We follow the literature and define labor market discrimination as a situation in which persons who provide labor market services and who are equally productive in a physical or material sense are treated unequally in a way that is related to an observable characteristic such as race, ethnicity, or gender. By “unequal” we mean these persons receive different wages or face different demands for their services at a given wage.

We consider two main economic models: entrepreneurs are willing to forego profits to cater to their “taste” for discrimination, as first proposed by Becker (1957). The second model is “statistical” discrimination: in an effort to maximize profits, firm owners discriminate based on a set of observables because they have imperfect information. This could be as simple as employers having imperfect information on the relative skills or productivity of minority versus majority agents. The models in both of these literatures are deep, and rich with good intuition. We do not have the space to do them justice, but strive simply to provide a sketch of each model to give the reader a sense of how one can test between them. We urge the reader to see Altonji and Blank (1999) for a more detailed presentation of these models and their implications.³⁷

³⁷ As far as the law is concerned, both types of discrimination—taste based and statistical—are illegal. For example, in credit markets, the Equal Credit Opportunity Act (Sec. 701, as amended in March 1976) states that it “shall be unlawful for any creditor to discriminate against any applicant, with respect to any aspect of the credit transaction. . . on the basis of race, color, religion, national origin, sex or marital status, or age. . .”. The law implies that while it is allowed to differentiate among customers based on characteristics of the customer (e.g., credit history) or the product that are linked to the expected return of the transaction, it is illegal to use the customer’s membership in a group to distinguish among customers. In other words, firms should make decisions about the customer as if they had no information regarding the customer’s race, sex, etc. This, for example, is true regardless of whether race is or is not a good proxy for risk factors in the credit market (Ladd, 1998).

3.2.1. Taste-based discrimination

In his doctoral dissertation, [Becker \(1957\)](#) modeled prejudice or bigotry as a “taste” for discrimination among employers. Becker modeled employers as maximizing a utility function that is the sum of profits plus a disutility term from employing minorities,

$$U = PF(N_{NM} + N_M) - W_{NM}N_{NM} - W_M N_M - dN_M, \quad (10)$$

where P is product price, F is the production function, which takes on two arguments: the number of employees that are non-minority N_{NM} and minority N_M . The second term is the wage bill and the final term is the disutility from employing minorities, dN_M . For prejudiced employers, the marginal cost of employment of a minority worker is $W_M N_M + dN_M$. Accordingly, d is the “coefficient of discrimination,” or the level of distaste of the employer for employing a minority worker. The higher d , the more likely the employer will hire non-minority workers, even if they are less productive than minority workers.

The Becker model then shows that the wage premium for non-minority workers is determined by the preferences of the least prejudiced employer who hires minority workers. Several extensions to this model have been proposed in the literature, including the possibility that d is a function of the job type, wage level, or the extent of segregation in the labor market. For example, [Coate and Loury \(1993\)](#) develop a model that restricts all employers to have identical preferences, but makes d a factor only when the employer hires minority workers for skilled jobs; an important consideration in the model then becomes the ratio of minority and non-minority people working in skilled jobs.

A logical conclusion of many of the studies in this area is that with certain assumptions—in many cases free entry, constant returns to scale, segmenting, etc.—in the long run non-discriminating employers will increase to the point that it is no longer necessary for minority workers to work for prejudiced employers, eliminating any wage discrepancies between minority and non-minority workers. This is a testable implication.

3.2.2. Statistical discrimination

[Arrow \(1972\)](#) and [Phelps \(1972\)](#) discuss discrimination that is consistent with the notion of profit-maximization, or [Pigou’s \(1920\)](#) “third-degree price discrimination.” In this class of model, in their pursuit of the most profitable transactions, marketers use observable characteristics to make statistical inference about reservation values of market agents. The underlying premise implicit in this line of work is that employers have incomplete information and use observables to guide their behavior. For example, if they believe that women might be more likely to take time out of the labor force, employers with high adjustment costs might avoid those expected to have higher attrition rates. Firms then have an incentive to use gender to “statistically discriminate” among workers if gender is correlated with attrition.

Of course, employers can discriminate along second moments of observable distributions too. Sobel and Takahashi (1983) develop a model along such lines, and their model is reconsidered in List and Livingstone (2010), which we closely follow here. In this framework, employers look at second moments and use prior beliefs about the productivity of group members to influence hiring and wage outcomes.

In the case where workers approach employers in an effort to sell their labor services, the employer proposes the wage (price) in each period. The worker can accept or reject the offer. If the offer is rejected, the employer makes another offer. If the offer is rejected in the terminal period, no exchange occurs. To keep the analysis simple and without losing focus on the critical incentives, we consider a two-period model. The results can all be extended to an n -period model.

Consider the situation where the employer's reservation value is public information and denoted v_b , where $v_b \in [0, 1]$, and the employer knows only the distribution from which the worker's reservation valuation is drawn. An employer confronts a potential worker, who has reservation value v_s , which is drawn from a distribution $F(v)$ on support $[0, 1]$. It is assumed that this c.d.f. is continuously differentiable, and that the resulting p.d.f. $f(v)$ is positive for all $v_s \in [0, 1]$. The employer discounts future payoffs by q , $q \in [0, 1]$. Further assume that the worker discounts future payoffs at the rate of p , $p \in [0, 1]$. p and q can be thought of as the costs of bargaining and are known by both players.

The bargaining process proceeds as follows: the employer proposes a price (wage) to the worker in period 1. The worker can accept or reject the offer. If the offer is rejected, the employer proposes a new price. It is assumed that the new proposal must be a wage (price) that is no lower than the original offer. The worker can accept or reject this proposition. If it is rejected, the game ends and no transaction occurs.

Following a no-commitment equilibrium, the employer is assumed to make the period 1 offer at the beginning of period 1, and subsequently chooses the period 2 offer using the information gained from the worker's rejection of the period 1 offer. Let x_1 be the employer's offer in period 1, and let x_2 be the employer's offer in period 2. Also, define, for period $i = 1, 2$,

$$S_i = \begin{cases} 0 & \text{if } i = 0 \\ (x_1 - px_2) & \text{if } i = 1 \\ \frac{(1-p)}{x_2} & \text{if } i = 2. \end{cases} \quad (11)$$

A worker whose reservation value is v_s will prefer accepting in period i to accepting in period $i + 1$ if $x_1 - v_s > p(x_{i+1} - v_s)$, or if $v_s < S_i$. A worker's most preferred time to accept is period i if $S_{i-1} < v < S_i$, so $F(S_i) - F(S_{i-1})$ is the employer's *ex ante* probability of hiring in the i th period and $(v_b - x_i)[F(S_i) - F(S_{i-1})]$ is the employer's *ex ante* undiscounted expected profit in period i .

The employer's maximization problem can be stated in terms of his choice of the period 2 offer, x_2 , and of S_1 , which implies a choice of x_1 , since $S_1 = \frac{(x_1 - px_2)}{(1-p)}$. The employer's optimal strategies are found via backwards induction, starting with his period 2 decision. If the period 1 offer is rejected, then the employer knows $v_s \geq S$. The employer chooses an offer x_2 to maximize his expected profits. Let $\pi(S)$ be this maximum value,

$$\pi(S) = \max_{x_2} \frac{(v_b - x_2)[F(x_2) - F(S)]}{1 - F(S)}. \quad (12)$$

Let $x_2(S)$ be the unique value of x_2 that solves (12). The first order condition of this problem, which implicitly defines $x_2(S)$, implies that,

$$(v_b - x_2(S))f(x_2(S)) = F(x_2(S)) - F(S). \quad (13)$$

Since the offer $x_2(S)$ must be less than the employer's valuation v_b , the left-hand side of (13) is positive, so the right-hand side must also be positive. For this to be the case, it must be true that,

$$v_b > x_2(S) > S. \quad (14)$$

In other words, in equilibrium, the second period offer must be greater than the first period offer, and both offers must be less than v_b .

The no-commitment equilibrium is fully characterized by $x_2(S)$ and a first period price, \hat{x}_1 , that solves,

$$\max_x (v_b - x)[F(S(x))] + q\pi(S(x))[1 - F(S(x))], \quad (15)$$

subject to,

$$S(x) = \begin{cases} S, & \text{where } S = \frac{x - px_2(S)}{1 - p} \text{ if such as } S \in [0, 1] \text{ exists} \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

Substituting in the constraint and the definition of $\pi(S)$, the problem becomes,

$$\begin{aligned} \max_{0 \leq S \leq 1} & v_b F(S) - (1 - p)SF(S) \\ & - px_2(S)F(S) + q(v_b - x_2(S))[F(x_2(S)) - F(S)]. \end{aligned} \quad (17)$$

If \widehat{S} solves (17), then $\widehat{x}_1 = (1 - p)\widehat{S} + px_2(\widehat{S})$ and $\widehat{x}_2 = x_2(\widehat{S})$ are the no-commitment equilibrium offers. The first-order condition of (17) implies,

$$(1 - q)v_b f(S) - (1 - p)[F(S) - Sf(S)] + (q - p)x_2(S)f(S) + px_2'(S)F(S) = 0. \quad (18)$$

3.2.3. Optimal employer behavior

Within this framework one can analyze differences in how an employer will behave when he confronts members of the various groups. To obtain insights on the impact of changes in the variance of the worker's reservation value on both offers that the employer may make, we consider a simple example where the worker's value is drawn from a uniform distribution. There are two groups of potential workers. Members of group 1 draw their reservation value from a uniform distribution with lower bound a_1 and upper bound b_1 . Members of group 2 draw their value from a uniform distribution with lower bound a_2 and upper bound b_2 . Assume $a_1 > a_2$ and $b_1 < b_2$, so the variance of group 2's distribution is larger than the variance of group 1's distribution. Without loss of generality, further assume that the bounds are such that the distributions have equivalent means. Now, consider the employer's equilibrium offers, when confronting a worker who is a member of group i , $i = 1, 2$.

Solving through backwards induction, the employer first calculates his period 2 offer, as in (12). Substituting in the uniform distribution and simplifying, the problem becomes,

$$\max_{x_2} (v_b - x) \left(\frac{x - a}{b - S} \right) - (v_b - x) \left(\frac{S - a}{b - S} \right). \quad (19)$$

The solution of this problem is,

$$\widehat{x}_2 = x_2(S) = \frac{S + v_b}{2}. \quad (20)$$

Hence, the period 2 decision is a function of S , which is chosen in period 1. The employer's period 1 problem is to solve (17). Substituting in the distribution and the solution $x_2(S)$, the first order condition is given by,

$$\left(1 - \frac{1}{2}p\right)2S - \left(1 - \frac{1}{2}p\right)a + \frac{1}{2}q(S - v_b) = 0, \quad (21)$$

which implies the solution,

$$\widehat{S} = \frac{(1 - \frac{p}{2})a_i + \frac{1}{2}qv_b}{2 - p + \frac{q}{2}}, \quad (22)$$

making the optimal period 2 offer,

$$\hat{x}_2 = \frac{1}{4 - 2p + q} \left[\left(1 + \frac{q}{2}\right) v_b + \left(1 - \frac{p}{2}\right) a_i \right], \quad (23)$$

and the optimal period 1 offer,

$$\hat{x}_1 = (1 - p) \frac{\left(1 - \frac{p}{2}\right) a_i + \frac{1}{2} q v_b}{2 - p + \frac{q}{2}} + p \frac{1}{4 - 2p + q} \left[\left(1 + \frac{q}{2}\right) v_b + \left(1 - \frac{p}{2}\right) a_i \right]. \quad (24)$$

Note that the optimal offers \hat{x}_1 and \hat{x}_2 are both increasing in a_i , and therefore decreasing in the variance of reservation values. Group 2's reservation value is drawn from a distribution with a larger variance, hence $a_1 > a_2$. In this example, then, the analysis shows that when the employer believes he is dealing with a member of a group whose reservation value is widely distributed (group 2), he will offer to hire at a lower wage than he would if the worker were a member of a group with a lower variance (group 1). This is true despite the fact that the first moments of the distributions are identical. This prediction provides one means to test the statistical discrimination model against the taste-based discrimination model. We return to this notion below.

3.3. Empirical tests

Scholars have concerned themselves primarily with the question “is there discrimination in market X?” and much less time has been spent on answering the question “why do firms discriminate?” As economists interested in public policy, however, we should be interested in not only the extent of discrimination but also the source of discrimination. Conditional on the existence of discrimination, it is imperative to understand the source of discrimination, since one cannot begin to craft social policies to address discrimination if its underlying causes are ill-understood. We now turn to an overview of a select set of studies that measure discrimination.

3.3.1. Observational data

One of the most important means to empirically test for marketplace discrimination in labor markets is to use regression-based methods. The focus using this approach has ranged from measuring labor force participation to modeling wage determination. Within the line of work that explores wages, the overarching theme is to decompose wage differentials, using an Oaxaca decomposition, between groups into what can be

explained by observables and what cannot be explained by observables. More specifically, consider a simple model that makes wages for minorities as follows,

$$W_M = \beta_M X_M + e_M, \quad (25)$$

and wages for non-minority agents as,

$$W_{NM} = \beta_{NM} X_{NM} + e_{NM}, \quad (26)$$

where W represent wages, X is a vector of individual specific observables that affect wages, and e is a classical error term. The wage difference between minority and non-minority agents can be computed by differencing these equations as follows,

$$\begin{aligned} W_M - W_{NM} &= [\beta_M X_M + e_M] - [\beta_{NM} X_{NM} + e_{NM}] \\ &= [X_M - X_{NM}] \beta_M + [\beta_M - \beta_{NM}] X_{NM}. \end{aligned}$$

The first term in the right most expression, $[X_M - X_{NM}] \beta_M$, is the component of the wage difference that is explained: it arises because of differences in the average characteristics of group members, such as region of residence, experience, or education level. The second term, $[\beta_M - \beta_{NM}] X_{NM}$, is the part of the wage difference that is not explained by the regression model—the differences in the response coefficients of the regression, or the rate of return differences across minorities and non-minorities. This last term encompasses differences in wages due to differences in the returns to similar characteristics between groups. For example, returns to education may differ across minorities and non-minorities. The fraction of wage difference due to this second term is typically called the “share” of wage differences due to discrimination.

Before discussing some of the general results from various regression-based approaches, it is important to qualify the results. First, the approach of assuming that the entire second component, $[\beta_M - \beta_{NM}] X_{NM}$, is due solely to discrimination is likely not correct. For example, for this to be true the wage equation must be well specified. If omitted variable bias exists, then the response coefficients will be biased. Second, this equation captures only discrimination in the labor market as measured today. That is, even if no discrimination is found to exist in such a model in today’s wages, that does not imply discrimination is unimportant. For example, if women are constantly denied market opportunities for skilled jobs, they might not invest optimally to obtain such positions. In the literature, such under investment is denoted as market discrimination before, or “pre” market discrimination. Clearly, it is difficult to parse the effects of years past with the current effects of discrimination, and this should be kept in mind when interpreting the empirical results below—both those from the regression based model as well as from field experiments.

The regression models can be applied to the data discussed above from the Current Population Survey (CPS) from May 2009. Yet, given that [Altonji and Blank \(1999\)](#) summarize a series of regression results from such wage equations that do not differ markedly from ours, we simply restate the main results. First, we find white men receive significantly higher wages than black men, even after controlling for education, job experience, region of residence, and occupation. Following the letter of the model, this is evidence of discrimination in the data.

Second, even after controlling for key factors, Hispanic men and minority female workers have lower wages than their non-minority counterparts. Once again, if one sticks to the interpretation of the model, this is suggestive evidence that discrimination exists between these groups. One should highlight, however, that there are certain difficulties in using CPS data for such an exercise—such as the problem of not having individual ability measures, such as cognitive and non-cognitive abilities. [Altonji and Blank \(1999\)](#) extend the CPS results by modeling data from the National Longitudinal Survey of Youth (NSLY). In general, their results with NSLY data confirm that an improved specification reduces the unexplained effects for blacks and for women.

While this line of work is suggestive that discrimination exists in the labor market, due to productivity unobservables the nature of discrimination is not discernible without rather strong assumptions: are minority men receiving lower wages because of tastes or because of statistical discrimination?

Some headway has been made in these regards recently in several clever studies. One such study is the ingenuous paper of [Goldin and Rouse \(2000\)](#), who use audition notes from a series of auditions among national orchestras in order to determine whether or not blind auditions—those in which musicians auditioned behind a screen—help women relatively more than men. The authors use a panel data set and identify discrimination by the change in hiring practices toward blind auditions that occurred in the 1970s. Goldin and Rouse study the actual audition records obtained from orchestra personnel managers and orchestra archives from eight major symphony orchestras from the late 1950's to 1995. These records contain lists of everyone auditioning (first and last name) with notation around the names of those who advance. There are three rounds of auditions considered: preliminary, semifinal, and final. The gender of the participants is determined by their name (96% of the records are distinctly masculine or feminine).

Eighty-four percent of all preliminary rounds were blind, seventy-eight percent of all semifinal rounds were blind, and seventeen percent of all final rounds were blind. In addition, the authors have personnel rosters that describe final assignments (members of the orchestra). There is variation in hiring practices over time, so that within one orchestra, the same audition may be blind or non-blind over time and across categories (preliminary, semifinal, and final). In addition, since success is rare, the same musician sometimes auditions more than once.

In the authors' data, 42 percent of individuals competed in more than one round and 24 percent competed in more than one audition. Including musician fixed-effects, the authors identify the effect of a screen to hide gender from those individuals who auditioned both with and without a screen. Without this "ability" control (individual fixed-effects) the data suggests that women are worse off with blind auditions. However, controlling for individual fixed effects, the authors find that for women who make it to the finals, a blind audition increases their likelihood of winning by 33 percentage points.

In their main specification, the authors find that women are significantly less likely to advance from semifinals when auditions are blind, but significantly more likely to advance from preliminary auditions and final auditions when they audition behind a screen. Turning to the final outcome space—what is the effect of the screen on the hiring of women?—the authors estimate that though they are unable to obtain a statistically significant effect (since the likelihood of winning an audition is less than three percent), women are five percentage points more likely to be hired than men when auditions are completely blind and there is no semifinal round. There is no difference between the likelihood that women are hired relative to men when there is a semifinal round and auditions are blind.

Ultimately, the effects discussed give pause to reported "traditional" orchestra practices. In particular, "a strong presumption exists that discrimination has limited the employment of female musicians." Before the implementation of blind auditions, committees were instituted to overthrow the biased hiring practices of conductors (who reportedly hired select males from a small set of well known instructors). However, sex-based biases seemed to dominate hiring, even in the face of "democratization." As the authors demonstrate, the institution of blind hiring significantly increased the success rate of women in most auditions.

However, it is difficult for the authors to parse whether or not the discrimination is taste-based or statistical. The authors note that an orchestra is a team, which requires constant improvement and study together. In this sense, female-specific absences—maternity leave—can impact the quality of the orchestra significantly and may motivate statistical discrimination against women. Using their data, the authors note that the average female musician took 0.067 leaves of absence per year, compared to the average males' 0.061 leaves. The length of leave was negligibly different between genders. These statistics imply that taste-based discrimination, assuming no performance differences between hired males and females, are at least in part the cause of the discrimination against female musicians. Again, without the strong assumption that conditional on being hired, women and men of the same audition caliber perform indistinguishably in their careers, it is difficult for this innovative work to parse the type of discrimination observed.

A second clever piece of work based on the regression approach is due to [Altonji and Pierret \(2001\)](#), who create a model that generates strict predictions on the effect of

race on wages over time under a hypothesis of statistical discrimination based on race by employers. Notably, they conclude that if firms do not statistically discriminate based on race (if they follow the law), but race is negatively related to productivity, then: (i) the race gap will widen with experience, and, (ii) adding a favorable variable that the hiring firm cannot observe will reduce the race difference in the experience profile. The authors find that the data satisfy these predictions: the race gap widens with experience and the addition of a “skill” variable reduces the race gap in experience slopes. Thus, the authors conclude that employers “do not make full use of race as information.”

Fundamentally, the authors’ model studies the differential effect on wages of “easy to observe” s variables and “hard to observe” z variables that predict worker productivity. While s variables such as schooling should have a smaller and smaller effect on wage over time, since an employer’s experience with the worker reveals far more important predictors of productivity, those variables that are difficult to observe such as skill have a relatively larger effect on wages as time goes on. This implies that the authors can identify whether or not the easily observable characteristic of race is acting as an s variable, or if employers are ignoring it. If employers are ignoring race, but race remains negatively correlated with productivity, then race acts as a z variable, appearing more important—more predictive of wage—over time. Again, the authors find support for the latter case.

The authors estimate their model using NLSY 1979 data—a panel study of men and women aged 14–21 in 1978 that have been surveyed annually since 1979. The data on white and black men with eight or more years of education forms the basis of their empirical analysis. The authors use AFQT (the Armed Forces Qualification Test) scores as a variable that employers do not observe, but that predicts productivity. In addition, the authors control for the first job held by all subjects in order to ensure that their results are not driven by the effect that a high AFQT may have on a worker’s access to jobs in which skill is observed, rather than “dead-end jobs” where skill is never observed. Because the authors control for secular shifts in the wage structure, their identification of the interactions between time and observable (s) characteristics and unobservable or ignored (z) characteristics comes from variation across age cohorts.

The authors find that a one standard deviation shift in AFQT rises from having no effect on wages when experience is zero to increasing log wages by 0.0692 when experience is 10. This supports the result that employers learn about productivity. The coefficient on education interacted with experience declines from 0.0005 to -0.0269 when the variable $AFQT \times experience$ is added. With an intercept of 0.0832 with the addition, we can conclude that the effect of an extra year of education declines from 0.829 to 0.0595 over ten years. This suggests that employers statistically discriminate on the basis of education because they have limited information about labor market entrants. In short, the effect of easy-to-observe variables like education dwindles as hard-to-observe variables like ability become more available—as time goes on and the employer becomes more familiar with the quality of the worker. The authors find similar effects with their

other hard-to-observe variables that correlate with productivity such as father's education and sibling wage rate: as experience increases, these variables become more and more predictive of higher wages (though the effect of father's education is never significant).

The main analysis is on whether or not employers statistically discriminate based on race. If firms use race as information—that is, as easily-observable predictors of performance similar to education—then the effect of race over time on wages should decline as hard-to-observe variables like skill (predicted by the AFQT) become more transparent over time. If firms ignore race, however, the initial (experience = 0) race gap should be small, and should widen with experience if race is negatively related to productivity. Also, when race is ignored (a z variable) adding another z variable like AFQT*experience will reduce the race gap in experience slopes. The authors note that the effect of a “black” dummy will not necessarily be zero even if firms do not statistically discriminate on the basis of race, since race may be correlated with legally usable information available to the employer but not to the econometrician.

Empirical analysis shows that the effect of adding AFQT*experience decreases the race gap in experience slopes (from -0.1500 to -0.0861); this is the opposite of what we would expect if employers fully used race as a predictor of performance (as they do with schooling—recall, the addition of AFQT*experience increases the amount by which the impact of education changes over time). Using another prediction of their model, that the effect of learning on the s variables will equal the effect of learning on the z variables times the relationship between the s and z variables—that there are spillover effects from learning—the authors are able to reject race as an s variable but not able to reject race as a z variable.

A few points are of note. First, if the quantity of training is influenced by the employer's beliefs about a worker's productivity, effects of training cannot be separated from the effects of statistical discrimination with learning. In addition, if taste-based discrimination becomes more prevalent at higher level positions, a widening of the race gap based on experience may be a reflection of increased taste-based discrimination rather than employer learning. Finally, the authors model the effect of statistical discrimination on wages, but not on the extended hiring decision. Based on these considerations, the authors note that any of their results on race-based discrimination should be interpreted cautiously.

To summarize, Altonji and Pierret test for statistical discrimination in a very reasonable way: they argue that if firms statistically discriminate, an observable characteristic such as race will be very important in predicting wages early in the employment history—before productivity is well observed—but becomes less important in predicting wages as time goes on and the worker accumulates experience. In the data the opposite is true, suggesting that under the assumption that the model is well specified firms attempt to ignore race in their hiring decisions, but that race is correlated with productivity (which is revealed) and so it becomes more and more predictive of wages as time goes on.

A third innovative regression-based study is due to Charles and Guryan (2008), who use state-level variation in historical wage and survey data to empirically test the impacts of discrimination on the labor market, focusing on taste for discrimination. The main theoretical result from Becker's work explored by Charles and Guryan is the assertion that black workers are hired by the least prejudiced employers in the market due to sorting in the labor market. Furthermore, they examine whether racial wage gaps are determined by the prejudice of a marginal employer, not the average. This sorting mechanism provides Charles and Guryan with two empirical regularities to verify Becker's work: (i) the level of prejudice observed by the employers displaying large amounts of prejudice (in the upper tail of a distribution of prejudice) should not impact wages; (ii) holding prejudice constant, wages should be lower with more blacks in the labor market.

Although they do not target the question of taste versus statistical based discrimination directly, they do include a variable for the skill difference between blacks and whites in regressions run as robustness checks. This and other robustness checks do not alter the main results which find support for Becker's theory of marginal prejudice affecting wages: marginal and low percentile prejudice levels negatively impact the black white wage gap while higher percentile and average prejudice levels have no impact; also the percent of the population that is black has a negative impact on the wage gap.

Charles and Guryan (2008) begin by empirically motivating the relationship between the black-white wage gap and prejudice by displaying the correlation between wage data from the CPS and white survey responses to questions concerning racial sentiments from the General Social Survey (GSS). After displaying the positive wage gap to prejudice relationship, Charles and Guryan review the theoretical findings to clarify the hypotheses of interest and then discuss the data. The data being used for prejudice is a non-uniform (the same questions are not asked every year) nationally representative survey with state-level data from 1972-2004. The survey questions used in this analysis are those from white responders and are vetted to reflect prejudice as much as possible (for example a question on whether "the government was obligated to help blacks" was not used due to the possible response aimed at the government). The survey responses were used to formulate a prejudice index relative to the responses given in 1977 and a prejudice distribution and the data on prejudice is combined with CPS May monthly supplement from 1977 and 1978 and CPS Merged Outgoing Rotation Group (MORG) for analysis.

The empirical results come from a hedonic wage regression. The regressions are run at the state level under the assumptions that employment markets are at the state level and interstate moves are costly. Because the prejudice measure they have is at the state level, Charles and Guryan take an additional step to allow for more reasonable standard errors than ones that would come from a full regression with observations at the individual level. This additional step comprises removing the prejudice index but including a black dummy variable for each state (state-black dummy interaction) in the first stage wage hedonic, and then using the coefficient from this interaction term as the

dependent variable in a second stage regression which includes the prejudice index. Five main measures of the prejudice index are analyzed: average, marginal, 10th percentile, median and 90th percentile. The marginal level of prejudice is calculated as the “ p th percentile of the prejudice distribution, where p is the percentage of the state workforce that is black” and the prejudice distribution is calculated from the GSS data. Additionally, the fraction of the population that is black in the state is included in the second stage.

The second stage regression results all support Becker’s theory. The first result of a negative impact on the black–white relative wages (negative means lower wages for blacks) attributed to the average level of prejudice is not significant and becomes positive when the marginal level of prejudice is included. The impact of the marginal prejudice measure is always negative and significant. This is also the case for the coefficient on the measurement of the fraction of the state population that is black (always negative and significant). These first results are taken as indication that the average prejudice measures fail to explain the wage gap, while the marginal and fraction of black have the assumed relationship from Becker’s work.

The additional prejudice measurements: 10th percentile, median and 90th percentile, provide further support for Becker’s theory. When included together in a regression, both with and without the percent of the state’s population that is black, the 10th percentile is the only variable of significance (it is negative). This result is taken as further support of Becker’s theory because of the indication that higher measurements of prejudice do not affect the wage gap (note that when the proportion of the state’s population that is black is included, the 10th percentile increases in both absolute magnitude and significance).

Various robustness checks are completed such as the inclusion of variables to indicate skill as mentioned above. Two skill measures are used: (i) separate reading and math variables which measure the difference between black and white test scores at the state level from a National Assessment of Educational Progress–Long Term Trend (NAEP–LTT) test, and, (ii) black–white relative school quality measures used and [Card and Krueger \(1992\)](#) (for which they reduce the sample to just southern states). In both cases the results are similar to when the skill proxies are not included. Although this identification strategy does not disentangle the impact of taste and statistical based discrimination, the inclusion of skill level measures does suggest that this is taste-based discrimination under the assumption that the skill measures accurately reflect the difference in work–place abilities between races and that these differences in abilities are known by the employers. In a best case scenario, identifying statistical discrimination would require some measure of employee productivity by race and employment.

Further robustness checks investigate other possible endogeneity issues. An instrument of the proportion black in the state workforce in 1920 was used to account for possible endogeneity issues with the percent of the state’s current population that is black. No difference in results was found. Finally, [Charles and Guryan \(2008\)](#) include a measure from the National Education Longitudinal Survey of 1988 (NELS) to account

for the fraction of co-workers that were of the same race. The results again supported Becker's theory that market sorting results in blacks being more segregated towards lower prejudiced employers: the wage gap is larger when the co-workers are more mixed when accounting for racial prejudice and the black proportion of the population.

The overall result is best restated directly from the last paragraph in the paper: "Our various results suggest that racial prejudice among whites accounts for as much as one-fourth of the gap in wages between blacks and whites... a present discounted loss in annual earnings for blacks between \$34,000 and \$115,000, depending on the intensity of the prejudice of the marginal white in their states."

Similar to the above studies, making an assumption on the regression specification, allows Charles and Guryan (2008) to begin to parse the type of discrimination observed. As such, as all of these incredibly insightful studies illustrate, one can go a long way in detecting discrimination, and its sources, but pinpointing exactly the extent that taste based and statistical discrimination is the underlying motive, is only possible with additional assumptions.

3.3.2. Field experiments

A complementary approach to measuring and disentangling the nature of discrimination is to use field experiments. Although a very recent study thoroughly catalogues a variety of field experiments that test for discrimination in the marketplace (Riach and Rich, 2002), a brief summary of the empirical results is worthwhile to provide a useful benchmark. Labor market field studies present perhaps the broadest line of work in the area of discrimination. The work in this area can be parsed into two distinct categories: personal approaches and written applications.

Personal approaches include studies that have individuals either attend job interviews or apply for employment over the telephone. In these studies, the researcher matches two testers who are identical along all relevant employment characteristics except the comparative static of interest (e.g., race, gender, age). Then, after appropriate training, the testers approach potential employers who have advertised a job opening. Researchers "train" the subjects simultaneously to ensure that their behavior and approach to the job interview are similar.

Under the written application approach, which can be traced to Jowell and Prescott-Clarke (1970), carefully prepared written job applications are sent to employers who have advertised vacancies. The usual approach is to choose advertisements in daily newspapers within some geographic area to test for discrimination. Akin to the personal approaches, great care is typically taken to ensure that the applications are similar across several dimensions except the variable of interest.

It is fair to say that this set of studies, including both personal and written approaches, has provided evidence that discrimination against minorities across gender, race, and

age dimensions exists in the labor market. But due to productivity unobservables, the nature or cause of discrimination is not discernible. This point is made quite starkly in Heckman and Siegelman (1993, p. 224), who note that “audit studies are crucially dependent on an unstated hypothesis: that the distributions of unobserved (by the testers) productivity characteristics of majority and minority worker are identical.” They further note (p. 255): “From audit studies, one cannot distinguish variability in unobservables from discrimination.” Accordingly, while these studies provide invaluable insights into documenting that discrimination exists, care should be taken in making inference about the type of discrimination observed.

Much like the labor market regression studies discussed above, the literature examining discrimination in product markets has yielded important insights. Again, rather than provide a broad summary of the received results, we point the reader to Yinger (1998) and Riach and Rich (2002), who provide nice reviews of the product market studies.³⁸ We would be remiss, however, not to at least briefly discuss the flavor of this literature.

One often cited, recent study is the careful work due to Bertrand and Mullainathan (2004), who utilize a natural field experiment to determine whether or not blacks are discriminated against by employers. By sending resumes with randomly assigned white- or black-sounding names to want-ads advertised in Boston and Chicago newspapers, Bertrand and Mullainathan find that white names receive 50% more callbacks for an interview than black names. This racial gap is uniform across occupation, industry, and employer size. Additionally, whites receive greater benefits to a higher-quality resume than blacks. Although Bertrand and Mullainathan are unable to test the type of discrimination, whether taste-based or statistical, as it is uncertain what information the employer is utilizing from the resumes, the authors use the results to suggest an alternate theory be considered, such as one based on lexicographic searches.

To choose names that are distinctly white-sounding or black-sounding, Bertrand and Mullainathan use name frequency data calculated from birth certificates of all babies born in Massachusetts between 1974 and 1979. Distinctiveness of a name is calculated as having a sufficiently high ratio of frequency in one racial group to that of the other racial group. The 9 most distinct male and 9 most distinct female names for each racial group, along with corresponding white- or black-sounding last names, are used. To verify this method of distinction, a brief survey was conducted in Chicago asking respondents to identify each name as “White”, “African-American”, “Other”, or “Cannot Tell.” Names that were not readily identified as white or black were discarded.

The authors sampled resumes posted more than six months prior to the start of the experiment on two job search websites to use as a basis for experimental resumes. The

³⁸ The interested reader should also see the recent special Symposium issue on Discrimination in Product, Credit, and Labor Markets that appeared in the *Journal of Economic Perspectives* Spring (1998).

resumes sampled were restricted to people seeking employment in sales, administrative support, clerical services, and customer service in Boston and Chicago, and were purged of the original owner's name and address. To minimize similarities to actual job seekers, Chicago resumes are used in Boston and Boston resumes are used in Chicago (after names of previous employers and schools are changed appropriately). The quality of the resumes were sorted into two groups (high and low), with high-quality resumes having some combination of more labor market experience; fewer gaps in employment history; being more likely to have an e-mail address, certification degree, or foreign language skills; or been awarded honors of some kind. Education is not varied between high- and low-quality resumes to ensure each resume qualifies for the position offered, and approximately 70% of all resumes included a college degree of some kind.

Fictitious addresses were created and randomly assigned to the resumes based on real streets in Boston and Chicago. The authors selected up to three addresses in each 5-digit zip code in both cities using the White Pages. Virtual phone lines with voice mailboxes were assigned to applicants in each race/sex/city/resume quality cell to track callbacks. The outgoing message for each line was recorded by someone of the appropriate race and gender, and did not include a name. Additionally, four e-mail addresses were created for each city, and were applied almost exclusively to the high-quality resumes.

The field experiment was carried out between July 2001 and January 2002 in Boston and between July 2001 and May 2002 in Chicago. In most cases, two each of the high- and low-quality resumes were sent to each sales, administrative support, and clerical and customer services help-wanted ad in the Sunday editions of *The Boston Globe* and *The Chicago Tribune* (excluding ads asking applicants to call or appear in person to apply). The authors logged the name and contact information for each qualifying employer, along with information on the position advertised and any specific requirements applicants must have. Also recorded was whether or not the ad explicitly stated that the employer is an "Equal Opportunity Employer."

For each ad, one high-quality resume and one low-quality resume were randomly assigned a black-sounding name (with the remaining two resumes receiving white-sounding names). Male and female names were randomly assigned for sales jobs, while primarily female names were used for administrative and clerical jobs to increase the rates of callbacks. Addresses were also randomly assigned, and appropriate phone numbers were added before formatting the resumes (with randomly chosen fonts, layout, and cover letters) and faxing or mailing them to the employer. A total of 4870 resumes were sent to over 1300 employment ads. Of these, 2446 were of high-quality while 2424 were of low-quality.

Results are measured by whether a given resume elicits a callback or an e-mail back for an interview. Resumes with white-sounding names have a 9.65% chance of receiving a callback compared to 6.45% for black-sounding names, a 3.2 percentage point difference. This difference can only be attributed to name manipulation. According to

these results, whites are 49% (50%) more likely to receive a callback for an interview in Chicago (Boston). This gap exists for both males and females, with a larger, though statistically insignificant, racial gap among males in sales occupations. An additional year of workforce experience increases the likelihood of a callback by approximately 0.4 percentage point, thus the return to a white name is equivalent to 8 additional years of experience. High-quality resumes receive significantly more callbacks for whites (11% compared to 8.5%, $p = 0.0557$), while blacks only see a 0.51% increase (from 6.2% to 6.7%). Whites are favored (defined as more whites than blacks being called back for a specific job opening) by 8.4% of employers, where blacks are favored by only 3.5% of employers, a very statistically significant difference ($p = 0.0000$). The remaining 88% of employers treat both races equally, with 83% of employers contacting none of the applicants.

A probit regression of the callback dummy on resume characteristics (college degree, years experience, volunteer experience, military experience, e-mail address, employment holes, work in school, honors, computer skills, special skills, fraction of high school dropouts in the neighborhood, fraction of neighborhood attending college or more, fraction of neighborhood that is white, fraction of neighborhood that is black, and log median per capita income) is created from a random subsample of one-third of the resumes. The remaining resumes are ranked using the estimated coefficients by predicted callback. Under this classification, blacks do significantly benefit from high-quality resumes, but they benefit less than whites (callback rates for high versus low are 1.6 for blacks and 1.89 for whites). The presence of an e-mail address, honors, or special skills have a positive significant effect on the likelihood of a callback. Interestingly, computer skills negatively predict callback and employment holes positively predict callback. Additionally, there is little systematic relationship between job requirements and the racial gap in callback.

Applicants living in whiter, more educated, or higher-income neighborhoods have a higher probability of receiving a callback, and there is no evidence that blacks benefit any more than whites from living in a whiter, more educated zip code. There is, however, a marginally significant positive effect of employer location on black callbacks.

Of all employers, 29% state that they are “Equal Opportunity Employers” and 11% are federal contractors, however these two groups are associated with a larger racial gap in callback. The positive white/black gap in callbacks was found in all occupation and industry categories except for transportation and communication. No systematic relationship between occupation earnings and the racial gap in callback was found.

Bertrand and Mullainathan did not design their study specifically test the two theories of discrimination, statistical and taste-based, and do not believe that either of the two can fully explain their findings. While both models explain the average racial gap, their results do not support animus. There is no evidence of a larger racial gap among jobs that explicitly require communication skills and jobs for which customer or

co-worker contacts are more likely to be higher, which would be expected by theory. Further, as blacks' credentials increase the cost of discrimination should increase, but this doesn't explain why blacks get relatively lower returns to a higher-quality resume. This, combined with the uniformity of the racial gap across occupations, casts doubt on statistical discrimination theories as well.

The authors suggest that other models may do a better job than statistical or taste models at explaining these particular findings. For example, a lexicographic search by employers may result in resumes being rejected as soon as they see a black name, thus experience and skills are not rewarded because they are never seen. This theory may explain the uniformity of the race gap if this screening process is similar across jobs. The results could also follow from employers having coarser stereotypes for blacks. In any case, Bertrand and Mullainathan acknowledge the need for a theory beyond statistical discrimination and taste to explain their findings in full.

Another nice example of a natural field experiment is due to [Riach and Rich \(2006\)](#), who extend the literature by using carefully matched written applications made to advertised job vacancies in England to test for sexual discrimination in hiring. They find statistically significant discrimination against men in the "female occupation" and against women in the "male occupation." This is important evidence to begin to uncover the underlying causes for labor market discrimination. This study is also careful to point out that it is difficult to parse the underlying motivation for why such discrimination exists. Even without such evidence, however, the paper is powerful in that it provides a glimpse of an important phenomenon in a significant market, and provocatively leads to questions that need to be addressed before strong policy advice can be given.

There are a number of other studies that examine discrimination and differential earnings in labor markets based on sexual orientation ([Arabsheibani et al., 2005](#); [Weichselbaumer, 2003](#); [Berg and Lien, 2002](#)), but like these two natural field experiments, they also have difficulties parsing the type of discrimination observed.

One might then ask, if field experiments have similar difficulties as regression based methods in parsing the nature of discrimination, why bother with this approach. Our answer is that field experiments in labor economics have the potential to parse both the nature and extent of discrimination observed in markets.

As a starting point, consider [List \(2004b\)](#), who made use of several settings in a naturally-occurring marketplace (the sports card market) to show that a series of field experiments can parse the two forms of discrimination. More specifically, after first demonstrating that dealers treat "majority" (white men) and "minority" (older white men, nonwhite men and white women) buyers and sellers in the marketplace differently, List provides evidence suggesting that sportscard dealers knowingly statistically discriminate. By executing a variety of field experiments, the evidence provided parses statistical discrimination from taste based discrimination and an agent's ability to bargain when interacting with a dealer. The experiments conducted by List demonstrate a

framework for potentially parsing the two forms of discrimination which could be utilized and moved forward to inform discrimination discussions in other markets.

The first experiment discussed in List (2004b) is similar to an audit study in that dealers are approached by buyers from both majority and minority groups with an offer of buying or selling a sportscard (unlike most audit studies, the subjects do not know that they are part of a study on discrimination, just that it's an economic study). The results from this first experiment are highly suggestive that dealers base offers on group membership: buyers in the minority groups of white women and older white men received initial offers that were 10–13% greater than white male buyers when buying cards and minority groups received 30% lower initial offers when selling their cards.

Further, this initial framed field experiment shows that the gap between minority and majority subjects' offers remain from the initial to final offers for inexperienced subjects but to a large part converges for subjects with experience. But, this convergence comes at a cost of time: subjects in the minority group having to invest a significantly larger portion of time to achieve better final offers.

This result provides support for non-taste based discrimination due to the convergence of the gap in offers through bargaining, a result that would not hold under a theory of taste based discrimination where the dealer would simply hold to one price. Finally, by surveying dealers in addition to subjects that were buying and selling, List controls for dealer experience in the marketplace and finds a positive relationship between dealer experience and discrimination as measured by the difference between a dealer's average majority and average minority offers. Suggesting that statistical discrimination may be evident unless one believes that taste based discrimination increases with experience as a dealer at sportscard shows.

Although this initial experiment can measure discrimination, more treatments are necessary to parse statistical discrimination from alternative explanations. In total, List runs four experiments in addition to the framed field experiment described in the previous paragraph: (i) a dictator game artefactual field experiment with dealers as the dictator and four descriptions as the receiver: white men, non-white men, white women and white mature men; (ii) two framed field experiment treatments that are bilateral exchange markets with dealers selling to agents with randomly drawn reservation values, where in one market dealers know that the reservation value is random and in a second it is ambiguous; (iii) a Vickrey second price auction that is a natural field experiment; and, (iv) a framed field experimental game designed to determine dealers' perceptions of the reservation value distributions of sportscard market participants. Each additional experiment helps parse the two forms of discrimination and the bargaining ability of the subjects and the results of all the experiments are necessary for List to suggest that dealers knowingly statistically discriminate.

First, the relatively uniform offers made to receivers across majority and minority groups in the dictator game suggests that dealers do not display taste based discrimination,

at least in artefactual field experiments. Second, through the bilateral exchange markets, three results are found which each point towards statistical discrimination by testing hypotheses drawn directly from the theories of taste based and statistical discrimination. First, experienced dealers are found to lose less surplus than inexperienced dealers. Second, minority and majority buyers perform similarly with the randomly set reservation prices but not in the treatment where dealers think that reservation values are “homegrown values.” Finally, experienced dealers perform worse when it is ambiguous whether the reservation value is drawn randomly—suggesting that they are utilizing inferences which are not performing well (i.e. their statistical discrimination rubric fails due to the randomly set reservation value). These two additional experiments point toward statistical discrimination.

Yet, it is only through the final two experiments that sufficient evidence is provided for statistical discrimination through a discovery of a variation in reservation value distributions of sportscard market participants and dealer knowledge of the variation. The results from the Vickrey second price auction are used for two purposes: (i) to determine whether the reservation value distributions of the majority and minority are indeed different and (ii) to provide distributions to determine the dealers’ abilities to accurately assign distributions in the final experimental game. The results from the Vickrey auction do show that the reservation values for the minority group have a larger variance than the reservation values for the majority group, suggesting that statistical discrimination could be utilized for profit maximization—see the above model. Further, when different reservation value distributions are shown to dealers in the final experiment, a majority of all dealers are able to determine which distributions are from which groups and experienced dealers are able to correctly assign distributions more often than inexperienced dealers.

Although List focuses on a market that every consumer does not necessarily approach, the framework of multiple field experiments to move towards identifying the form of statistical discrimination is one that should be considered for use elsewhere. Most importantly, this study highlights that a series of field experiments can be used to uncover the causes and underlying conditions necessary to produce data patterns observed in the lab or in uncontrolled field data.

This study shows highlights that a deeper economic understanding is possible by taking advantage of the myriad settings in which economic phenomena present themselves. In this case, field experimentation in a small-scale field setting is quite useful in developing a first understanding when observational data is limited or experimentation in more “important” markets is not possible. Yet, it is important to extend this sort of analysis to more distant domains.

This is exactly what is offered in [Gneezy et al. \(2010\)](#), who explore the incidence of discrimination against the disabled by examining actual behavior in a well-functioning marketplace—the automobile repair market. This study uses a traditional audit study, but

combines it with a specific field experimental treatment to allow the authors to parse the type of discrimination observed.

The audit portion of the study was standard: the assignment given to subjects is clear: approach body shop j to receive a price quote to fix automobile i . The authors included subjects from two distinct groups—disabled white males age 29–45 and non-disabled white males age 29–45—who each visited six body shops. The disabled subjects in this experiment were all confined to a wheelchair and drove a specialized vehicle. All of the automobiles, which were personally owned by our disabled subjects, had visible body problems. Importantly, both testers in any given pair approached body shops with the identical car.

The authors find that overall, the disabled received considerably higher average price quotes, \$1425, than the non-disabled, \$1212. Inference as to why this disparate treatment exists, of course, is an open issue. Several clues provide potential factors at work: (i) access—many body shops are not easily approachable via wheelchairs; this considerably restricts the set of price offers the disabled can receive; and (ii) time—while the non-disabled can easily park and proceed to the front desk, the process is much more complex for the disabled. First, he must find a suitable parking place: it is very uncommon to have designated places for the disabled in body shops. As a result, the disabled must have special parking which permits the use of a wheelchair. Moreover, it must be a space that will be unoccupied when he returns to pick up the repaired vehicle. After finding an appropriate parking space, the disabled must commit much more effort and time to approach the service desk. An additional related problem which makes the expected search cost higher for the disabled person is that in some cases it is necessary to leave the car for the day in order to obtain a price quote. Using a taxi is much more complex for the disabled than for the non-disabled.

To investigate the search cost explanation further, the authors obtained data on search effort at the tester level and perceived search effort at the body shop level to examine if realizations of these variables are consistent with the pattern of discrimination observed. From this survey, the authors find that the non-disabled typically consult far fewer body shops: on average, the non-disabled visit 3.5 different mechanics whereas the disabled visit only 1.67 mechanics, a difference that is statistically significant. Concerning the supply side, the authors asked body shops questions revolving around body shop perceptions of the degree of search among the disabled and non-disabled. The results are consonant with the consumer-side statements observed above: the disabled are believed to approach 1.85 body shops for price quotes while the non-disabled are expected to approach 2.85, a difference of more than 50 percent and one that is significant. This evidence is consistent with statistical discrimination based on mechanics' beliefs about relative search costs and how they map into reservation value distributions. Yet, the survey evidence alone is only suggestive and further investigation is necessary to pinpoint the underlying mechanism at work.

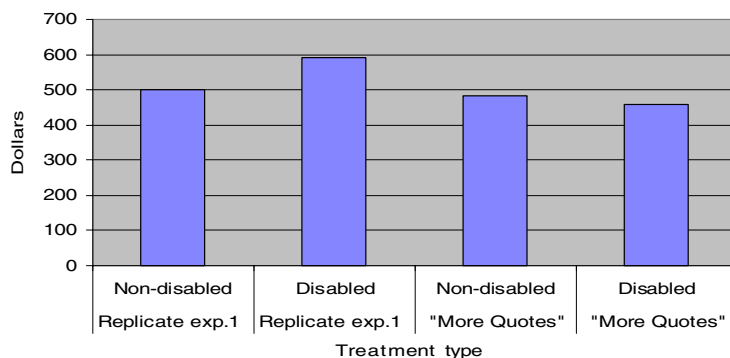


Figure 8 *Complementary experiment I summary.*

For these purposes, the authors provide a sharper focus on the underlying reason for discrimination by running a complementary field experiment. In this experiment, the authors not only replicated the initial results, with new testers and new vehicles in need of repair, but in another treatment had these exact same agents approach body shops explicitly noting that [they're] “getting a few price quotes” when inquiring about the damage repair estimate. If differential search costs cause discrimination, then the authors should observe the offer discrepancies disappearing in this treatment.

This is exactly the result they observe. Although in the replication treatment the disabled continued to receive higher asks, when both agent types noted that they were “getting a few price quotes,” the disabled agents were able to secure offers that were not statistically distinguishable from the offers received by the non-disabled. We provide support of this insight in Fig. 8, which highlights the discrepancies observed when search is believed to be heterogeneous across the disabled and non-disabled. In this case, the first two bars show that the differences are nearly 20%. Yet when both agents clearly signal that this particular mechanic visit is just one part of their entire search process, these disparities are attenuated and indeed change signs.

While these two examples are not directly related to labor market outcomes, they display the power of the field experimental method to test important theories within labor economics, and especially the theories of discrimination discussed earlier. In this regard, we believe that similar treatments can be carried out in labor markets to explore wage differences, job offer differences, and other labor market outcomes.

4. FIRMS

In the third stage of the life cycle, most individuals work within firms or some other hierarchical organization. In this section we describe how field experiments have contributed to knowledge on how workers and employees behave in such settings. Following the research areas described in Tables 2 and 3, we frame our discussion on the

following research themes: (i) the effects on monetary incentives on worker behavior; (ii) the interplay between monetary and non-monetary incentives; (iii) aspects of the employment relationship, such as gift-exchange between employers and employees, and the link between employer monitoring and employee shirking behavior.^{39, 40}

4.1. Monetary incentives

A core principle in economics is that incentives matter. The role of monetary incentives within firms and organizations has been long studied in sociology and management literatures. With the application of contract theory to behavior within firms (Hart and Holmstrom, 1987) and the development of personnel economics (Lazear, 1995), such questions are now integrated within mainstream labor economics. For economists, the basic questions have been: (i) how do workers respond to a given set of incentives?; (ii) what are the optimal set of incentives an employer should provide?⁴¹

An earlier generation of empirical studies exploited firm's personnel data to measure the productivity effects of compensation schemes on individual workers. An econometric challenge facing these studies is that observed incentive contracts might well be endogenous to firm's performance (Prendergast, 1999; Chiappori and Salanie, 2003). In other words, identifying causal effects of incentives on behavior is confounded by the presence of unobservables, such as managerial practices, that determine both which compensation schemes are chosen, and worker productivity. In earlier research this concern has been addressed in between firm studies using instrumental variables approaches (Groves et al., 1994). However, this concern applies even if such effects are identified from a within worker or within firm comparison as incentives change over time (Jones and Kato, 1995; Ichniowski et al., 1997; Paarsch and Shearer, 1999, 2000; Lazear, 2000). A related concern is that such changes in incentives might be reflective of a wider package of changes in management practices. Hence, akin to

³⁹ The field experiment approach shares many of the characteristics of the insider econometrics approach to understand the causes and consequences of behavior within a firm (Ichniowski and Shaw, forthcoming). However a key distinction is that field experiments explicitly rely on exogenous variation created with the specific influence of researchers in order to identify causal effects. Clearly, not every intervention that a researcher could design and implement is socially useful—there is little value added in implementing practices that firms are never otherwise observed engaging in. However, this does not preclude the fact that carefully designed interventions can help researchers to uncover causal relations and the mechanisms behind them.

⁴⁰ Our discussion focuses predominantly on natural field experiments within firms. There also exists a separate branch of artefactual field experiments where subject pools are drawn from manufacturing workers (Barr and Serneels, 2009), fishermen (Carpenter and Seki, 2010) and employees in large firms (Charness and Villeval, 2009).

⁴¹ Many of the wider literature related to the research questions we touch upon, such as incentive pay and teams, are discussed in greater detail in the Chapter on Human Resource Management by Bloom and Van Reenen (2011), also in this Handbook. They summarize the evidence from across countries showing the increasing use of performance pay over time. In the Chapter on Personnel Economics in this Volume by Oyer and Schaefer (2011), further issues related to incentive pay and firm hires is discussed at greater length.

social experiments, what is actually being evaluated is potentially the sum total of many concomitant changes in the firm's organization rather than an isolated change in worker incentives all else equal. This is of particular concern given the view that there exist complementarities between organizational practices so that firms are better off choosing a package of practices rather than in isolation (Milgrom and Roberts, 1990; Ichniowski et al., 1997). With such multiple underlying changes, mapping the evidence, however cleanly identified is the change in behavior, to any underlying theory is less clear cut.⁴²

Field experiments introduce exogenously timed variation in incentive structures that are orthogonal to other management practices. This opens up the possibility to identify the causal impact of monetary incentives on the behavior of individual workers, and on firm performance as a whole. Combining personnel files from human resource departments within the firm, with primary data collection that is inherent in field experimentation, allows researchers to examine the effect of monetary incentives on a range of margins of worker behavior, capturing both the intended and unintended consequences of incentive provision.

There are good theoretical reasons for collecting such extensive information on worker behaviors when evaluating the response to incentives. For example, multi-tasking theory suggests that when monetary incentives are provided based on a subset of tasks that the firm can directly measure performance in, workers may reallocate their effort away from other tasks they are engaged in, their employer is affected by, but their compensation is not based on (Holmstrom and Milgrom, 1991). Similarly, if the provision of incentives alters the distribution of pay across workers in the same tier of the firm hierarchy, this might alter worker's behavior towards co-workers, say through cooperation or sabotage (Lazear, 1989). Finally, there might be ways in which workers can game against any incentive scheme. All such unintended consequences of monetary incentives need to be accounted for to both accurately understand how workers respond to incentives and to begin to think through the optimal incentive design.⁴³

Employers might not collect such information *ex ante*. Hence the need to engage in primary data collection efforts to complement the rich information available in firm's personnel files. Field experiments—that involve close cooperation between researchers and firm management—are well placed to advance in this direction. Ultimately, as witnessed in some of the field experiments described below, this allows a closer mapping between the evidence and underlying theory, and to draw implications for optimal incentive provision.

⁴² Due to these empirical challenges, it is not surprising that much of the early evidence testing theories in personnel economics originated from laboratory environments. For example, Bull et al. (1987) provide evidence from the lab on the predictions of rank order tournament theory; Fehr and Fischbacher (2002) review the experimental evidence on social preferences in workplace environments. The wider availability of personnel data and ever closer links being forged between researchers and firms has allowed the literature in field experiments within firms to flourish.

⁴³ Charness and Kuhn (2011, 2010) review the extensive evidence from laboratory settings on sabotage.

4.1.1. Theoretical framework

To understand some of the theoretical questions and empirical challenges faced in this literature, it is instructive to first reconsider Lazear's (2000) original analysis of the Safelite Glass Corporation, a large auto-glass firm in which the primary task of worker's at the bottom-tier of the firm's hierarchy is to install automobile windshields. Lazear used non-experimental methods to estimate the productivity effects of the firm moving from a compensation scheme in which workers were paid an hourly wage scheme, to one in which they were paid a piece rate for each windshield installed, with a minimum guarantee. This pioneering work brings to the fore many of the issues that have influenced all the subsequent literature, and allows us to highlight the specific issues that field experiments help address.

The model is as follows. Worker's utility depends on income Y and effort e , $U(Y, e)$ with $U_1 > 0$, $U_2 < 0$.⁴⁴ Worker's output q depends on effort and her ability, θ , so $q = q(e, \theta)$ with output assumed to be observable and $q_1, q_2 > 0$. For any given output q_0 , there is a unique effort level that achieves this, denoted $e_0(\theta)$. It is then straightforward to see that $\frac{\partial e}{\partial \theta} = -\frac{q_2}{q_1} < 0$ so that higher ability workers need exert less effort to achieve a given output. If workers choose not to work at any firm, their outside option from leisure is denoted $U(0, 0)$. Hence the lowest ability worker that would accept employment at a firm with a required output level and wage W , is denoted θ_0 and is such that,

$$U(W, e_0(\theta_0)) = U(0, 0). \quad (27)$$

All workers of higher ability earn rents from employment over leisure. Similarly, suppose a worker of a given ability could take up employment at another firm offering a wage-minimum effort pair $(\widehat{W}, \widehat{e})$. Hence with inter-firm competition there might exist an upper cutoff in ability, θ_h , such that,

$$U(W, e_0(\theta_h)) = U(\widehat{W}(\theta_h), \widehat{e}(\theta_h)), \quad (28)$$

where workers of ability higher than θ_h prefer to take the alternative employment contract.

⁴⁴ There is a long-standing idea in psychology that rewards may hinder performance (Kruglanski, 1978). There is some evidence on this from laboratory settings where offering small amounts of monetary compensation is found to decrease effort relative to paying nothing (Gneezy and Rustichini, 2000), and where explicit incentives sometimes result in worse compliance than incomplete labor contracts (Fehr and Falk, 1999; Fehr and Schmidt, 2000). This might either be because small monetary incentives crowd out intrinsic motivation, an idea formalized by Benabou and Tirole (2000), or because the individual is reluctant to signal his willingness to accept low wages. We do not know of field evidence that examines such non-monotonic effects of monetary incentives on effort.




Figure 9 Compensation before and after at Safelite.

This framework makes clear that incentive structures will affect two types of behavior, an idea developed in more detail in Lazear (2005). First, there will be change in effort e exerted by individual workers in response to monetary incentives. This is referred to as the “incentive effect”. Second, the compensation scheme will induce a differential composition of workers within the firm over time. Some workers will prefer to join this firm from other employers. These changes in workforce composition can be thought of as the “selection effect” of monetary incentives.

The incentive effect can be easily understood graphically. Figure 9 shows the relationship between output, q , and compensation for the two schemes relevant for Lazear’s study: (i) a fixed hourly wage subject to a minimum output requirement q_0 , resulting in total compensation W ; (ii) a linear piece rate scheme $bq - K$ with a minimum guarantee of W . As Fig. 9 shows, for output levels between q_0 and q^* the worker receives W under both compensation schemes, and for output higher than q^* earns strictly more under the piece rate scheme.⁴⁵

On the incentive effect, the model makes clear that moving from the fixed hourly wage scheme to the piece rate scheme does not cause the output of any individual to fall, and causes average output to rise. Low ability workers, indicated with a solid indifference curve in Fig. 9, remain indifferent between the two schemes and would produce output q_0 at point A under both. Higher ability workers, indicated with a dashed indifference curve in Fig. 9, would prefer to increase their effort and move to point B. This is because the piece rate scheme allows higher ability workers to raise their utility through increased compensation that more than offsets any increase in their effort. As a result, the dispersion

⁴⁵ Firms typically provide workers some insurance by allowing their output to occasionally fall below the required minimum q^* , but a worker that consistently fails to meet this performance threshold is likely to be fired or assigned to another task.

in worker effort and output rises as long as there is at least one worker that chooses to produce more than q^* .⁴⁶

On the selection effect, under plausible conditions, the average ability of workers rises with the move to piece rates. This is because low ability individuals remain indifferent between working for this firm under either incentive scheme. If they were willing to work for the firm under the fixed wage scheme, they should remain willing to do so under piece rates all else equal. On the other hand, high ability workers might be attracted to this firm from other firms that for example, have higher minimum output standards or pay piece rates but at a lower rate b . In short, theory predicts that there should be no change in the number of low ability workers who are willing to work at the firm, but that piece rates attract high ability workers so the right tail of the ability distribution in the firm should thicken.

As described in more detail below, existing field experiments have focused on identifying the incentive effects, and the research designs used have been less amenable to pin down these types of selection effect. Yet it is important to emphasize the need for future research to provide credible research designs to uncover both effects.⁴⁷

In Lazear's study, he documents the total effect of the change in monetary incentives to workers was around a 44% increase in worker productivity, defined to be the number of windshields installed by the worker per eight hour day. Around half the increase was due to incentive effects, namely a change in effort of the same worker as he moved from a fixed hourly wage to a piece rate scheme. However, the other half was entirely due to the selection effect, namely productivity changes due to endogenous changes in the composition of workers in response to change in monetary incentives. A legacy of Lazear's study is to show that both motives likely underlie why firms choose to alter their output based incentive structures in the first place. A carefully crafted field experiment that begins to measure whether and how the compensation policies of a given firm have such spillover effects on other firms that compete for similar workers, would open up a rich research agenda tying together the study of within-firm compensation policies on equilibrium wage-setting behavior in labor markets.⁴⁸

⁴⁶ Whether workers exert more or less effort in response to a higher piece rate b of course depends on the balance of income and substitution effects. Evidence from the lab and field in Gneezy and Rustichini (2000) suggested that the relationship between piece rates and effort was U-shaped with low piece rates eliciting less effort than a zero piece rate. One explanation would be that small levels of financial compensation crowd out workers' intrinsic motivation to exert effort.

⁴⁷ Laboratory experiments have begun to explore in more detail the selection effects of incentives (Dohmen and Falk, 2006; Cadsby et al., 2007; Vandegrift et al., 2007; Niederle and Vesterlund, 2007; Eriksson et al., 2008a,b). These studies are described in more detail in Charness and Kuhn (2011, 2010).

⁴⁸ This links together with recent development in structural estimation of search and matching models in labor markets. For example, Cahuc et al. (2006) develop and estimate an equilibrium model with strategic wage bargaining and on-the-job search. An important innovation on their paper is that when an employed worker receives an outside job offer, a three-player bargaining process is started between the worker, her/his initial employer and the employer which made the outside offer. They use the model to examine wage determination in France using matched employer-employee

Linking to the design of field experiments

The model provides a series of implications that have impinged on the first generation of field experiments over the last decade. First, given worker heterogeneity, changes in compensation scheme will nearly always affect average effort and output, as well as the dispersion of effort and output. Given the linkage between performance and pay, this inevitably results in changes in the distribution of earnings across workers at the same tier of the firm hierarchy. Hence linking pay to performance might have unintended negative consequences on worker and firm performance as a result of such increased earnings inequality. This might manifest itself in the form of workers reducing cooperation with co-workers (Baron and Pfeffer, 1994; Bewley, 1999; Lazear, 1989), workers sabotaging the performance of others, or workers being directly worse off in utility terms, all else equal, as a result of them being structurally averse to pay inequality (Fehr and Schmidt, 1999; Charness and Rabin, 2002). Field experiments are particularly adept at detecting and quantifying such unintended consequences because researchers are engaged in primary data collection, and the firm would have no incentive to collect such information *ex ante* as part of its personnel files, especially if pay for performance type compensation schemes have not been previously implemented. Some of the field experiments described below have collected qualitative evidence from workers to explore these channels, in addition to using personnel files to measure the direct productivity effects of incentives.

Second, given the selection effect of monetary incentives, there is inevitably a change in workers' peer group over time. The composition of peers, or their social ties with each other, play no role in the standard neoclassical model in which worker preferences are only defined over their own income and effort. There are good reasons to probe this assumption in the field. First, if peer effects determine workplace behavior because they alter the marginal return to worker's effort, then understanding how workers respond to changes in monetary incentives requires an understanding of the mechanisms underlying such peer effects. Second, extending the neoclassical model to take into account such peer effects or social concerns, as has been done in Kandel and Lazear (1992), Lazear (1989), Rotemberg (1994) and Fershtman et al. (2003), has implications for many aspects of firm behavior including the optimal design of incentives. Such concerns are a recurring theme in the series of natural field experiments conducted by Bandiera et al. that are described below.

Third, the model highlights that under fixed hourly wage schemes, there should not be much heterogeneity in workers output or effort, despite workers being heterogeneous in ability. This does not fit the evidence very well. For example, in Lazear's study there was considerable dispersion in worker productivity even under the fixed hourly wage scheme. To better explain behavior in fixed wage settings, theory suggests workers might

data from 1993 to 2000. They find that inter-firm competition is quantitatively important for wage determination, and raising wages above reservation levels.

respond to other forms of non-monetary or implicit incentives, such as gift-exchange motives where workers exert more effort in response to employers paying higher than market clearing wages, the ability to shirk (Shapiro and Stiglitz, 1984; Macleod and Malcomson, 1989), or promotion prospects and career concerns (Dewatripont et al., 1999). Such mechanisms might also predict how workers sort across firms (Stiglitz, 1975). Some of these aspects are highlighted by the field experiments discussed below on the employment relationship.

Fourth, the model implies that low ability workers should not leave the firm with the move to piece rates. If they were willing to work for the firm under the fixed wage scheme, they should remain willing to do so under piece rates all else equal. On the other hand, given output differences among workers under price rates, management can more easily identify low ability workers when pay is tied to their performance. Over a longer time period, this might lead to them being fired. This type of selection effect caused by employer learning the true ability of workers, has not been studied by field experiments.⁴⁹

Finally, given worker heterogeneity, the first best for the firm would be to set a worker specific piece rate, b_i . This would be chosen to equate the marginal benefit of effort to its marginal cost. However, we generally observe firms being constrained to offer bottom-tier workers the *same* compensation scheme. This may be because of legal, technological or informational constraints (Lazear, 1989; Bewley, 1999; Encinosa et al., 1997; Fehr et al., 2004). To overcome this, one hypothesis is that firms can get closer to this first best by linking managers' pay to the firm's performance. Managers then have greater incentives to target their effort to specific workers, and from the worker's point of view it is then as if they face an individual specific incentive scheme. This idea is developed in the natural field experiment by Bandiera et al. (2007) described below, which then links managerial incentives to pay inequality among workers.

4.1.2. Evidence from the field

In the decade following the wave of studies using personnel data and insider econometrics to understand responses to monetary incentives (Ichniowski and Shaw, forthcoming), field experiments have begun to exploiting exogenous and randomly timed variation in compensation pay to both reinforce these existing results using non-experimental methods, as well as providing new insights.

Among the first of these studies was Shearer's (2004) natural field experiment, designed to estimate the productivity gains moving from a piece rate to a fixed wage

⁴⁹ A related concern has been on the existence of ratchet effects in response to pay for performance (Gibbons, 1987), whereby workers deliberately underperform to keep the piece rate high. Such ratchet concerns have been documented in firms where productivity shocks are uncommon such as shoe making (Freeman and Kleiner, 2005) and bricklaying (Roy, 1952). Cooper et al. (1999) present evidence from an artefactual field experiment on Chinese students and managers on such ratchet concerns, that might be of particular concern in planned economies.

scheme for tree planters in British Columbia, Canada. In contrast to Lazear (2000), this setting is one in which tree planters are usually paid a piece rate with no guaranteed minimum, and fixed wages are rarely used *ex ante*. Workers were randomly assigned at the start of a work day, to plant under one of the incentive schemes. Hence a within worker comparison can be exploited to estimate the incentive effects. To reduce the likelihood of this comparison confounding other effects unrelated to the compensation schemes in place, work took place on fields of similar conditions, and there was a constant length of the work day. The incentive effect estimate is then based on a total of 120 observations on daily worker productivity, 60 under each scheme. The relatively small sample size—nine male workers were randomly selected from the firm—reflects the difficulty researchers initially faced in real world settings in convincing firms to randomly assign individuals to alternative compensation schemes. As more field experiments are conducted, some of these constraints are being eased. For example, some of the field experiments described below are based on data on hundreds of workers.⁵⁰

Shearer's field experiment reveals the incentive effect of having a piece rate rather than a fixed wage compensation schemes to be a 20% productivity increase. The magnitude of this is comparable to Lazear's (2000) findings—moving from a fixed wage scheme—based on non-experimental data, although this is wholly by chance. There is no reason *a priori* to expect the behavioral response of workers to these two incentive schemes to be of the same magnitude given the very different types of worker involved, nature of the production function, and that the piece rate b was not the same across settings. In line with theory, Shearer finds the standard deviation of output across workers was higher under piece rates. Overall, it was found that unit costs under piece rates were around 13% lower than under fixed wages.

To shed more light on how workers would have responded in a slightly different environment and to alternative compensation schemes, Shearer then develops and estimates a structural model. In terms of altering the economic environment, the structural model is used to shed light on what would have been the productivity gains if management was imperfectly informed about planting conditions. This yields comparable estimates of the productivity gains of moving to piece rates. To shed light on alternative compensation schemes, Shearer explores how workers would have responded to an efficiency wage scheme. The efficiency wage that would induce effort levels equal to those observed under fixed wages in the field experiment, is calculated and its implied unit costs are compared with those achieved under piece rates. This exercise suggests that fixed wages would lead to a 2.7% increase in unit costs relative to piece rates.

Shearer (2004) uses a close to best practice methodology in combining estimates from a field experiment with structural modelling within the same setting. This combination,

⁵⁰ The small sample size used in Shearer (2004) also reflects the nature of tree planting firms. They typically employ less than 100 planters.

following from discussions in Heckman and Smith (1995) and Keane and Wolpin (1997), first identifies the existence and magnitude of important causal effects using reduced form evidence. To then move away from such black-box findings, the researcher then uses structural modelling to posit an underlying behavioral mechanism behind the effects, assess the sensitivity of the estimates to slight alterations in the economic environment, and to make headway in understanding the optimal compensation structure. Of course, the validity of the structural model can itself be tested by exploring whether it predicts the responses observed to the exogenous variation engineered by the field experiment.

Another example of how field experiments can and should learn from other methodologies is at the heart of the natural field experiment of Hossain and List (2009). They use theoretical insights on framing effects from behavioral economics that have previously found empirical support in laboratory experiments (Kahneman and Tversky, 1979; Thaler, 1980; Samuelson and Zeckhauser, 1988; Ellingsen and Johannesson, 2008), to see if, in the field, framing manipulations affect worker responses to bonus incentives. Their setting is a high tech Chinese firm producing consumer electronics, where workers are organized into both individual and team production. They find that bonuses framed as both “losses” and “gains” increase productivity, for both individuals and teams. Teams respond more to bonuses posed as losses than as comparable bonuses posed as gains. The comparable effects for individuals are of the same sign but are not statistically significantly different to each other. Team productivity is enhanced by 1% purely due to the framing manipulation. Neither the framing nor the incentive effect lose their importance over the six month study period. Nor are there any detrimental effects on the quality of work as measured by product defect rates.

On a practical note, the results highlight that conditional on bonuses being provided, framing matters, and as framing can be adjusted almost costlessly, there are simply ways in which firms can further enhance productivity responses to monetary incentives. Theoretically, these results from the field provide an example of the prevalence of loss aversion in a natural labor market setting. As such the results provide external validity to laboratory evidence, and should be seen to provide a strong argument in favor of field and laboratory experiments being complements, not substitutes.

4.2. Non-monetary incentives

Organizations use a variety of non-pecuniary based incentives to motivate their employees. We discuss three forms of non-monetary incentive: status goods, feedback, and social incentives.⁵¹

⁵¹ This list is not meant to be exhaustive. We focus on these because field experiments have provided insights on these margins to a greater extent than for other types of non-monetary incentive such as those discussed in Francois (2000), Dixit (2002), Prendergast (2001), Benabou and Tirole (2003), Seabright (2002), Delfgaauw and Dur (2004), Akerlof and Kranton (2005), and Besley and Ghatak (2005).

Under status incentive schemes, employees are given some positional good, such as an “employee of the month” job title. The notion that individuals crave status has been long studied (Veblen, 1934; Friedman and Savage, 1948; Duesenberry, 1949; Frank, 1985) and more recently formalized in the context of organizations providing status incentives in Moldovanu et al. (2007) and Besley and Ghatak (2008). They emphasize that for status incentives to be effective, the positional good must be valued by employees, it must be scarce, and its allocation rule rewards the deserving.

Recent evidence on these effects have been found in laboratory settings (Ball et al., 2001; Brown-Kruse et al., 2007) but few field experiments in which researchers have worked closely with a firm to exogenously vary such status rewards. One exception is Greenberg (1988), who reports results from a field experiment based on 198 employees in the underwriting department of a large insurance company. These employees were randomly assigned on a temporary basis to the offices of either higher, lower, or equal-status co-workers while their own offices were being refurbished. Relative to those workers reassigned to equal-status offices, those reassigned to higher status offices raised their performance, and those reassigned to lower status offices lowered their performance. The size of these performance changes were directly related to the magnitude of the status changes encountered. The results are interpreted as providing real world evidence on equity theory for non-monetary rewards.

In the future, we envisage field experiments being designed that randomly vary the first two margins on effective status incentives described above: how valued the positional good is, and its scarcity. In contrast, a field experiment that randomly allocated such positional goods might not be as informative, unless it was clearly related to some reallocation that would have occurred in any case, such as in Greenberg’s clever study described above. Otherwise, such random allocations would not be representative of the kinds of allocation rule employers actually use, and so cloud the interpretation of any such results.⁵²

A second class of non-monetary incentive in organizations relates to the provision of feedback. While there is a long tradition in psychology on feedback effects (Thorndike, 1913), economists have only recently begun to investigate its causes and consequences. Much of this research has focused on the theory of optimal feedback provision as mid-term reviews (Lizzeri et al., 2002; Ederer, 2008). Theory indicates that feedback on past performance can affect current performance either directly if past and current performances are substitutes or complements in the agent’s utility function, labelled a

⁵² There are field experiments on charitable giving that have exogenously varied the visibility of donations to assess whether such status concerns or prestige motives drive giving behavior Soetevent (2005), something that has been found to be the case in laboratory settings of public goods games (Andreoni and Petrie, 2004; Rege and Telle, 2004). Echoing some of the results below on gift-exchange in the field and the lab, Soetevent (2005) finds evidence that for some times of charitable cause, contributions increase when they can be socially recognized, but that this effect diminishes over time.

preference effect, or indirectly by revealing information on the marginal return to current effort, labelled a *signaling* effect.⁵³ The direct preference effect is relevant if, for instance, agents are compensated according to a performance target or fixed bonus scheme, so that being informed of high levels of past performance induces the agent to reduce her current effort relative to her past effort, and still meet her overall performance target. The indirect signaling effect is relevant if, for instance, the agent's marginal return to effort depends on her ability and this would be unknown if feedback were not provided.⁵⁴

Both these mechanisms imply the effect of feedback is heterogeneous across individuals and it might increase or reduce current effort, so the socially optimal provision of feedback remains ambiguous. Research in psychology also suggests feedback effects are heterogeneous and may crowd in or crowd out intrinsic motivation (Butler, 1987; Deci et al., 2001). Consistent with such heterogeneous effects, the organizational behavior literature finds that performance feedback within firms is far from ubiquitous (Meyer et al., 1965; Beer, 1990; Gibbs, 1991).⁵⁵

While there is a growing empirical literature on the effect of feedback in laboratory settings (Eriksson et al., 2008a,b), and in natural experiments (Bandiera et al., 2009a,b; Azmat and Iriberry, 2009; Blanes-i-Vidal and Nossol, 2009), evidence from field settings remains scarce. Bandiera et al. (2009a,b) provide one such analysis, in which the focus is on the provision of feedback to teams. This study is described in detail below.⁵⁶

The third class of non-monetary incentives relate to changes in behavior induced because of the presence and identity of co-workers—namely social relations in the workplace. The idea that there exists an interplay between social relations and monetary incentives in the workplace goes back to the old Hawthorne studies mentioned earlier and have been long considered in the organizational and business sociology literatures (Mayo, 1933; Barnard, 1938; Roethlisberger and Dickson, 1939; Roy, 1952;

⁵³ The organizational behavior and psychology literatures have also emphasized the signaling effects of feedback, as well as other related comparative statics such as how individuals change strategies in response to feedback (Vollmeyer and Rheinberg, 2005), and the specific type of information that should be conveyed in feedback (Butler, 1987; Cameron and Pierce, 1994).

⁵⁴ Two strands of the economics literature have explored aspects of the signaling effect of feedback. The first strand focuses on whether individuals update their priors in response to feedback consistent with Bayes' rule (Slovic and Lichtenstein, 1971). The second strand focuses on whether agents react more to positive than negative feedback because of self serving biases such as confirmatory bias (Rabin and Schrag, 1999), or overconfidence (Malmendier and Tate, 2005; Van Den, 2004).

⁵⁵ On the heterogeneous effects of feedback, the meta-analysis of Kluger and Denisi (1996) covering 131 studies in psychology with 13,000 subjects finds that two thirds of studies report positive feedback effects. On the optimal provision of feedback, when the agent knows her ability so that there is no indirect signaling effect of feedback, whether feedback should be optimally provided or not is sensitive to the specification of the agent's cost of effort function (Lizzeri et al., 2002; Aoyagi, 2007). More general results have been derived when agents learn their ability through feedback and ability is complementary to effort (Ederer, 2008).

⁵⁶ A separate branch of the literature has focused on the strategic manipulation of feedback by the principal (Malcolmson, 1984; Gibbs, 1991; Aoyagi, 2007), of which there is anecdotal evidence from the field (Longnecker et al., 1987) and laboratory (Ederer and Fehr, 2007). Evidence from field experiments on feedback remains scarce.

Williams and O'Reilly, 1998). Such concerns have begun to be incorporated in economic theory (Kandel and Lazear, 1992; Rotemberg, 1994; Prendergast and Topel, 1996), and credible evidence on their existence, magnitude and underlying mechanisms have begun to emerge in a nascent literature using non-experimental methods in combination with personnel data (Mas and Moretti, 2009; Bandiera et al., 2010, forthcoming).

In a series of natural field experiments, Bandiera et al. provide evidence on the effect of incentives on individual and firm performance within the same firm. These field experiments engineer exogenously timed variation in the incentive structures faced by workers in the firm. The common thread running through these studies is to provide evidence on the interplay between monetary and non-monetary incentives in the workplace. The specific form of non-monetary incentives considered are those arising from social relations in the workplace, so that workers behavior, and response to monetary incentives, might differ depending on the nature of the social ties they have with co-workers, their superiors, and their subordinates. Given that this form of non-monetary incentive is what field experiments have predominantly focused on, we first develop a framework that makes precise how such incentives can be incorporated into an otherwise standard model, and then map this framework to the empirical evidence from the field.

4.2.1. Theoretical framework

Suppose worker i 's payoff depends on three components. First, she derives some benefit from exerting effort e_i towards a productive task. This benefit, $B(e_i, \cdot)$, reflects in part how her effort maps into income through the monetary compensation scheme. To cover a wide range of compensation schemes including absolute performance evaluation incentive schemes such as piece rates, relative performance evaluation schemes such as rank order tournaments, or team incentives, these benefits will in general also depend on co-workers' effort, \mathbf{e}_{-i} . Second, the worker faces a convex cost of effort, $C(\theta_i, e_i)$, where workers are of heterogeneous ability, θ_i . Finally, we assume worker i places some weight on the utility of co-worker j , π_{ij} . In turn, such social preferences π_{ij} might depend on the existence or strength of the social tie between individuals i and j . This third component of the worker's payoff function generates social incentives.⁵⁷ Workers simultaneously choose their efforts to maximize their total payoff,

$$\max_{e_i} B_i(e_i, \mathbf{e}_{-i}) - C(\theta_i, e_i) + \sum_{j \neq i} \pi_{ij} [B_j(e_j, \mathbf{e}_{-j}) - C_j(\theta_j, e_j)]. \quad (29)$$

⁵⁷ Social preferences can be thought of as a reduced form representation of a number of models. They depict behavior consistent with reciprocity or altruism (Fehr and Schmidt, 1999), or the evolutionary equilibrium of a repeated Prisoner's Dilemma game in which workers learn which strategies to play (Levine and Pesendorfer, 2002; Sethi and Somanathan, 1999). In the field experiment reported in Bandiera et al. (2005), they attempt to distinguish between models in which workers' preferences display altruism towards others, and models in which workers behave *as if* they are altruistic because, for instance, they play trigger strategies to enforce implicit collusive agreements.

The first order condition is⁵⁸,

$$\frac{\partial B_i(\cdot)}{\partial e_i} - \frac{\partial C(\theta_i, e_i)}{\partial e_i} + \sum_{j \neq i} \pi_{ij} \frac{\partial B_j(\cdot)}{\partial e_i}. \quad (30)$$

The monetary compensation scheme determines the marginal benefit of effort, $\frac{\partial B_i(\cdot)}{\partial e_i} \geq 0$. As the worker has social incentives, she takes account of the fact that on the margin, her effort also affects the benefits that accrue to others, $\frac{\partial B_j(\cdot)}{\partial e_i} \leq 0$. As mentioned above, the precise sign of this social interaction depends on the nature of peer effects between workers that are socially connected, and the monetary compensation scheme in place.

The theoretical predictions of such models generate a wide range of behavioral responses. For example, working alongside friends might make work more enjoyable, generate contagious enthusiasm among friends, provide positive role models, or generate incentives to compete to be the best in the network of friends. All such mechanisms, that effectively increase the net benefits of effort, imply workers exert more effort in the presence of their friends relative to themselves when they work in the absence of their friends. Alternatively, working with friends might create contagious malaise, or lead to low effort norms within friends or co-workers more generally. All such mechanisms, that effectively decrease the net benefits of effort, imply workers exert less effort in the presence of their friends. Finally, the presence of friends might have heterogeneous effects across workers in that some exert more effort in the presence of their friends relative to when they work solely with non-friends, and others exert less effort. For example, friends or co-workers may conform to a common norm (Bernheim, 1994), or workers might be averse to pay inequality within their network (Fehr and Schmidt, 1999; Charness and Rabin, 2002). In either case, relative to when they work only with non-friends— (i) low ability workers exert more effort in the presence of their friends, and; (ii) high ability workers exert less effort in the presence of their friends. These aspects are highlighted by the field experiments discussed below on non-monetary incentives.

The field experiments in Bandiera et al. are designed to engineer exogenous variation in the incentives faced by workers to identify $\frac{\partial B_i(\cdot)}{\partial e_i}$, corresponding to a similar reduced form parameter as in Lazear (2000) and Shearer (2004). They then combine this variation with primary data collected on social networks and plausibly exogenous variation in the assignment of friends as co-workers over time, to identify social incentives as embodied in $\frac{\partial B_j(\cdot)}{\partial e_i}$. In these experiments, the authors examine the effects of social incentives both within and across tiers of the firm hierarchy. Namely in some studies i and j are co-workers engaged in the same tasks, and in other studies the pair correspond to a manager

⁵⁸ The model would be complicated if there were also knowledge spillovers such that effort exerted by worker i reduced the cost of effort of worker j . While such knowledge spillovers have been found in workplace settings (Moretti, 2004; Ichniowski et al., 1997) we abstract from them here.

and her subordinate. Moreover, they study cases in which: (i) individual effort hurts co-workers $\frac{\partial B_j(\cdot)}{\partial e_i} < 0$, as in the case of relative incentive schemes; (ii) where it benefits them $\frac{\partial B_j(\cdot)}{\partial e_i} > 0$, as in the case of team incentives; and (iii) where it has no effect $\frac{\partial B_j(\cdot)}{\partial e_i} = 0$, as in the case of a piece rate scheme. We now summarize the main insights from these natural field experiments.

4.2.2. Evidence from the field

Social incentives among bottom tier workers

The firm studied in Bandiera et al. is a leading UK producer of soft fruit. Managerial staff belongs to three classes. The first class consists of a single general manager whom we refer to as the Chief Operating Officer (COO), the second comprises ten field managers, and the bottom-tier of the firm hierarchy consists of workers whose main task is to pick fruit. Field managers are responsible for field logistics, most importantly to assign workers to rows of fruit within the field and to monitor workers. Managerial effort can therefore be targeted to individual workers and is complementary to worker's effort. The main task of the COO is to decide which workers are selected to pick fruit each day, and which are assigned to non-picking tasks. The field experiments described below together provide insights on behavior at each tier of the firm's hierarchy.

In each natural field experiment, the researchers worked closely with the CEO of the firm to engineer exogenously timed changes in monetary incentives to workers or managers. The *same* workers and managers are observed under both incentive schemes and therefore it is possible to control for time invariant sources of heterogeneity across workers, such as their ability, and across managers, such as their management style.⁵⁹ The most important remaining empirical concern is that the estimates of such changes might still reflect naturally occurring time trends in productivity. This is addressed using a battery of tests in each paper. In addition, the time span of study allows the authors to check in each case whether the behavioral response to incentives is long-lasting, or whether they reflect Hawthorne effects, as discussed earlier, whereby individuals respond in the short run to any change in their workplace environment. Being able to use field experiments to estimate short and long run responses to changes in management practice is a theme we will return to below when we present field experimental evidence on gift-exchange in firms, and contrast the evidence from the field and the laboratory.⁶⁰

⁵⁹ Hence this empirical strategy is informed by the evidence that individual "styles" of managers affect firm performance over and above firm level characteristics themselves (Bertrand and Schoar, 2003; Malmendier and Tate, 2005).

⁶⁰ Bandiera et al. study the behavior of nearly all the workers in the firm for each field experiment. However, given the experiment takes place in one firm, to avoid contamination effects across treated and control groups, all workers were simultaneously shifted from one incentive scheme to the other. In contrast, Shearer (2004) exogenously varied the incentive scheme workers were in on each day. In non-experimental studies such as Lazear (2000) on individual pay and Hamilton et al. (2003) on team pay, workers might have had some say on which compensation scheme they would be paid under.

In each natural field experiment, the authors collected primary data on the social networks of each individual worker. With such a precise mapping of the structure of friendship networks in the firm, personnel data providing workers productivity over time, and the field experiment on monetary incentives, the authors are able to shed light on the interplay between monetary and social incentives in this setting.

Finally, they have daily information on the pool of workers available to pick fruit. This allows them to precisely identify the effect of monetary incentives on the selection of workers from this pool. The entire pool of workers is observed in this context because individuals are hired seasonally from Eastern Europe, and they live on the farm for the duration of their stay. This margin of selection—driven by the COO's *demand* for workers—from the firm's internal labor market proves to be an important margin of response to some changes in incentives, particularly in relation to changes in managerial incentives. Still, these field experiments, like Shearer (2004), are silent on the selection effect highlighted by Lazear (2000) in relation to workers choice of which firm to *supply* their labor to.

Another obvious similarity between Shearer (2004) and Bandiera et al. is that they study agricultural environments in which worker productivity is easy to measure, comparable across workers at the same moment in time, and comparable within a worker over time. The fact that worker productivity is measured electronically with little measurement error, also makes analysis of the impact of the field experiment on the distribution of productivity, again as highlighted by Lazear (2000), particularly amenable to quantile regression methods for example. However, it remains true that settings in which worker's output is hard to measure, verify or compare, which might represent the bulk of tasks in the modern service based economy, remain relatively unexplored in field experiments.

In Bandiera et al. (2005) the natural field experiment exogenously changes the monetary incentives to the bottom-tier workers whose primary task is to pick fruit. The study compares the behavior of these workers under a relative incentive scheme to a piece rate scheme. The comparison is revealing because under relative incentives individual effort imposes a negative externality on co-workers' pay whereas under piece rates individual effort has no effect on others' pay. The difference in workers' performance under the two schemes, if any, then provides evidence on whether and to what extent workers internalize the externality they impose on their colleagues. To see this, the framework above is tailored to this specific field experiment as follows.

Consider a group of N workers, each worker i exerts effort $e_i \geq 0$ which determines her productivity. The cost of effort is assumed to be $\frac{\theta_i e_i^2}{2}$. Under relative incentives the benefit from pay depends on the worker's productivity relative to all her co-workers, $B\left(\frac{e_i}{\bar{e}}\right)$, where $\bar{e} = \frac{1}{N} \sum_i e_i$. The relative scheme has the key characteristics that an increase in worker i 's effort—(i) increases her pay; (ii) increases average effort and hence imposes a negative externality by reducing the pay of co-workers. The effort choice under relative incentives then depends on whether workers have social incentives and

therefore internalize this externality. Assuming worker i places the same social weight on all co-workers, so $\pi_{ij} = \pi_i$, the equilibrium effort for worker i solves,

$$\max_{e_i} B\left(\frac{e_i}{\bar{e}}\right) + \pi_i \sum_{j \neq i} \left(B\left(\frac{e_j}{\bar{e}}\right) - \frac{\theta_j e_j^2}{2} \right) - \frac{\theta_i e_i^2}{2}. \quad (31)$$

Assuming worker i chooses her effort taking the effort of others as given, the Nash equilibrium effort for worker i solves,

$$B'\left(\frac{e_i}{\bar{e}}\right) \frac{1}{\bar{e}} \left(\frac{\sum_{j \neq i} e_j}{\left(\sum_i e_i\right)} \right) - \frac{\pi_i}{\bar{e}} \sum_{j \neq i} B\left(\frac{e_j}{\bar{e}}\right) \frac{e_j}{\left(\sum_i e_i\right)} = \theta_i e_i. \quad (32)$$

Under piece rates, individual effort is paid at a fixed rate b per unit and worker i chooses her effort as follows,

$$\max_{e_i} B(b e_i) + \pi_i \sum_{j \neq i} \left(B(b e_j) - \frac{\theta_j e_j^2}{2} \right) - \frac{\theta_i e_i^2}{2}. \quad (33)$$

The equilibrium effort level solves the first order condition,

$$B'(b e_i) b = \theta_i e_i. \quad (34)$$

As worker i 's effort does not affect her co-workers' pay, her optimal choice of effort is independent of π_i . To compare effort choices under the two schemes, evaluate (34) at $b = \frac{1}{\bar{e}}$ so that for a given \bar{e} , the pay per unit of effort is the same under both incentive schemes. The first order condition under piece rates then is,

$$B'\left(\frac{e_i}{\bar{e}}\right) \frac{1}{\bar{e}} = \theta_i e_i, \quad (35)$$

so the difference between the first order conditions (32) and (35) can be ascribed to two sources. The first is the externality worker i imposes on others under relative incentives, the magnitude of which depends on π_i . When $\pi_i > 0$ worker i 's productivity is *lower* under relative incentives compared to piece rates. Second, by exerting more effort, each worker lowers the pay she receives for each unit of effort under relative incentives. This effect, captured by the $\frac{\sum_{j \neq i} e_j}{\left(\sum_i e_i\right)}$ term, also reduces productivity under relative incentives but is negligible in large groups.

The main results from Bandiera et al. (2005) are then as follows. First, the reduced form estimates suggest that the exogenously timed switch from relative incentives to piece rates had a significant and permanent impact on worker productivity. For the average worker, productivity increased by at least 50% moving from relative incentives to piece rates. As in the earlier literature, both the mean and dispersion of productivity significantly increase with the move to piece rates. The productivity gains achieved under piece rates are not found to be at the expense of a lower quality of picking.

The authors then assess whether this productivity change is consistent with the standard assumption that workers ignore the externality they impose on others under the relative scheme ($\pi_i = 0$), or whether they fully internalize it ($\pi_i = 1$). To do this they use the structural model above, imposing a functional form assumption on $B(\cdot)$ and a production function linking effort to observed output, to calibrate the first order conditions of the workers' maximization problem to compute an estimate of each worker's cost parameter, θ_i , under each incentive scheme and behavioral assumption. Since worker's ability is innate, they ought to find the *same* implied distributions of costs across workers under both incentive schemes if the underlying behavioral assumption is correct.

Calibration of the first order conditions for worker's efforts reveals that the observed change in productivity is *too large* to be consistent with the assumption that workers ignore the negative externality they impose on others. At the same time, the observed change in productivity is also *too small* to be consistent with the assumption that workers maximize the welfare of the group and fully internalize the negative externality. The authors then uncover the distribution of social weights π_i across workers that would explain the productivity increases. To do so they assume the true cost of effort θ_i of each worker is that derived under piece rates, and then substitute into the first order condition (32). They find the data is consistent with the average worker placing a weight of $\bar{\pi} = 0.65$ on the benefits accruing to all other co-workers, assuming they place a weight of one on their own benefits.

Further analysis combines the experimental variation induced by the change in incentive scheme, with non-experimental variation of the assignment of workers to work alongside their friends on some days but not on other days. The field experiment method allows the collection of primary data on social networks of each worker on the farm. This reveals that under relative incentives workers internalize the externality more when the share of their personal friends in the group is larger and this effect is stronger in smaller groups. In line with the interpretation that social preferences explain the difference in productivity across the two schemes, the relationship among workers *does not* affect productivity under piece rates. Finally, they find that productivity under relative incentives was significantly lower only when workers were able to monitor each other. Given that monitoring is necessary to enforce collusion while it does not affect altruism, they take this finding to support the hypothesis that workers are able to sustain implicit collusive agreements when relative incentives are in place. Hence, building on

a large body of evidence from laboratory settings, this evidence from the field suggests workers behave as if they have social preferences but do not, in structural form, have social preferences that make them unconditionally altruistic towards others.

The results beg the question of why, given the large gains to productivity and profits, of the move to piece rates, were relative incentives ever employed in the first place. The farm management suggested the relative scheme was mainly adopted to difference out common shocks that are a key determinant of workers productivity in this setting. While this is in line with the predictions of incentive theory, the superiority of relative incentives relies on the assumption that workers ignore the externality their effort imposes on others.⁶¹ This assumption on worker behavior is not supported by this field experiment. Relative incentives led to lower productivity because workers internalized the negative externality to some extent. The results of this natural field experiment then speak directly to Lazear's (1989) observation on how rarely workers are compensated according to rank-order tournaments, and point to new and interesting directions for theory to develop on the optimal provision of incentives under more robust assumptions on worker preferences.

Social incentives among managers

While the evidence from field experiments discussed thus far has focused on the monetary incentives provided to bottom-tier workers, Rosen's (1982) magnification principle implies the incentives provided higher up in the firm hierarchy can have larger effects on firms' performance. Bandiera et al. (2007) present evidence from a field experiment in the same setting as previously described to explore this issue.

They examine the effects of providing bonuses to managers based on the average productivity of their subordinates. They extend the framework above to highlight that, as in most firms, in their context managers can affect worker productivity through two channels—(i) they can take actions that affect the productivity of existing workers, and, (ii) they can affect the identity of the workers selected into employment. A simple theoretical framework indicates that, when workers are of heterogeneous ability and managers' and workers' effort are complements, the introduction of managerial performance pay makes managers target their effort towards the most able workers. This is labeled a “targeting effect” of managerial incentives. In addition, the introduction of managerial performance pay makes managers select the most able workers into employment. This is labeled as a “selection effect” of managerial incentives.

As in Lazear's framework, such targeting and selection effects influence both the mean and the dispersion of workers' productivity. Mean productivity unambiguously rises as managers target the most able workers and fire the least able. The effect on the

⁶¹ See Lazear and Rosen (1981), Green and Stokey (1983) and Nalebuff and Stiglitz (1983). Relative performance evaluation may also be preferred to piece rates as it lowers informational rents to high types (Bhaskar, 2002), and reduces incentives of workers to exert effort in influence activity (Milgrom, 1988).

dispersion is however ambiguous. On the one hand, targeting the most able workers exacerbates the natural differences in ability and leads to an increase in dispersion. On the other hand, if only more able and hence more similar workers are selected into employment in the first place, the dispersion of productivity may fall, depending on the underlying distribution of ability across workers.

They key findings from [Bandiera et al. \(2007\)](#) are as follows. First, the introduction of managerial performance pay increases both the average productivity and the dispersion of productivity among lower-tier workers. The average productivity increases by 21 percent and the coefficient of variation increases by 38 percent.

Second, the increase in the mean and dispersion of productivity is due to both targeting and selection effects. The analysis of individual productivity data reveals that the most able workers experience a significant increase in productivity while the productivity of other workers is not affected or even decreases. This suggests that the targeting effect is at play—after the introduction of performance pay, managers target their effort towards more able workers. The individual data also provides evidence of a selection effect. More able workers, namely those who had the highest productivity when managers were paid fixed wages, are more likely to be selected into the workforce when managers are paid performance bonuses. Least able workers are employed less often and workers at the bottom of the productivity distribution are fired.⁶²

Third, the selection and targeting effect reinforce each other, as workers who experience the highest increase in productivity are also more likely to be selected into employment. The introduction of managerial performance pay thus exacerbates earnings inequality due to underlying differences in ability both because the most able workers experience a larger increase in productivity and because they are selected into employment more often.

Finally, they evaluate the relative importance of the targeting and selection effects through a series of thought experiments. They find that at least half of the 21 percent increase in average productivity is driven by the selection of more productive workers. In contrast, the change in dispersion is nearly entirely due to managers targeting the most able workers after the introduction of performance pay. Namely, the dispersion of productivity would have increased by almost the same amount had the selection of

⁶² The results from this natural field experiment has implications for environments outside the workplace. For example, the provision of teacher incentives based on the average performance of students may have important consequences for the distribution of test scores among students, and the composition of students, and possibly teachers, admitted into schools. For example, [Burgess et al. \(2005\)](#) find that the introduction of school accountability based on test pass rates improved the performance of students in the middle of the ability distribution, at the expense of both high achieving and low achieving students. Similarly, [Hanushek and Raymond \(2004\)](#) and [Reback \(2005\)](#) provide evidence on the distributional consequences on student achievement under the *No Child Left Behind* policy. Finally, [Jacob \(2002\)](#) and [Figlio and Getzler \(2002\)](#) provide evidence on the selection effect. They show that the introduction of accountability schemes lead to an increase in grade retention and special educational placement in Chicago and Florida public schools, respectively.

workers remained unchanged. The reason is that the distribution of ability across workers is such that even when the least able workers are fired, the marginal worker selected to pick is still of relatively low ability. Hence there remains considerable heterogeneity in productivity among selected workers.

These findings shed some light on why firms provide performance related pay to managers in the first place. While such incentive schemes are obviously designed to increase unobservable managerial effort, these results suggest another more subtle reason for their use. This stems from the general observation that firms are typically constrained to offer bottom-tier workers the *same* compensation scheme. This may be because of legal, technological or informational constraints (Lazear, 1989; Bewley, 1999; Encinosa et al., 1997; Fehr et al., 2004). To the extent that bottom-tier workers are of heterogeneous ability, however, offering the same compensation scheme to all of them will be sub-optimal. When managers' pay is linked to firm's performance, their interests become more aligned with those of the firm and they have greater incentives to target their effort to specific workers in order to offset the inefficiency that arises because of the common compensation scheme. From the worker's point of view it is then as if they face an individual specific incentive scheme. This opens a broad research agenda to examine whether firms are indeed more likely to offer managers performance pay in settings where lower tier workers are of heterogeneous ability, managers are able to target their effort towards specific workers, and workers are offered the same compensation scheme.

The findings from this field experiment also highlight the interplay between the provision of managerial incentives and the earnings inequality among lower-tier workers. Such a linkage exists whenever managers can target their efforts towards some workers and away from others, and managers choose which individuals are selected into the workforce. Hence that there might be an important interplay between managerial incentives and earnings inequality among workers highlights a possible link between two important trends in labor markets over the past twenty years that have previously been unconnected in the economics literature—the rising use of managerial performance pay, and the rising earnings inequality among observationally similar workers.⁶³

In Bandiera et al. (2009a,b), the authors use the same introduction of managerial bonuses to understand whether managers favor workers they are socially connected to. In general, social connections between managers and workers can help or harm firm performance. On the one hand, social connections may be beneficial to firm performance if they allow managers to provide non-monetary incentives to workers, or help reduce informational asymmetries within the firm. On the other hand, managers may

⁶³ Residual, or within-group wage inequality, is a sizeable contributor of the growth in overall wage inequality in the US. This has been argued to have increased throughout the 1970s and 1980s (Juhn et al., 1993), and into the 1990s (Acemoglu, 2002; Autor et al., 2005).

display favoritism towards workers they are socially connected with, to the detriment of other workers and overall firm performance.⁶⁴

In this experiment, as managerial compensation becomes more closely tied to firm performance, we would expect managers to utilize social connections to a greater extent if indeed, such connections are beneficial for firm performance. On the other hand, if social connections are bad for the firm, we might expect managers to reallocate their effort across workers in response to managerial incentives, towards high ability workers, and away from workers they are socially connected to. To be precise, if the managers' behavior towards connected workers changes once their interests are more closely aligned with the firm's, their previous behavior under fixed wages could have not been maximizing the firm's average productivity.

To measure social connections the authors use a survey they designed to exploit three sources of similarity between managers and workers—whether they are of the same nationality, whether they live in close proximity to each other on the farm, and whether they arrived at a similar time on the farm. The underlying assumption is that individuals are more likely to befriend others if they are of the same nationality, if they are neighbors, or if they share early experiences in a new workplace.⁶⁵

The main findings are as follows. First, when managers are paid fixed wages, the productivity of a given worker is 9% higher when he is socially connected to his manager, relative to when he is not. As workers are paid piece rates, this translates into the same proportionate change in earnings. Second, when managers are paid performance bonuses that tie their pay to the average productivity of workers they manage, being socially connected to the manager has no effect on workers' productivity.

Third, the introduction of managerial performance pay significantly decreases the productivity of low ability workers when they are connected to their manager relative to when they were connected to their manager and she was paid a fixed wage. The introduction of managerial performance pay increases the productivity of high ability workers, especially when they are not connected to their managers. These findings indicate that when managers face low powered incentives, they favor the workers they are socially connected to, regardless of the workers' ability. In contrast, when they face high powered incentives, managers favor high ability workers regardless of the workers' connection status.

⁶⁴ Both the positive and negative effects of social connections have been stressed in the organizational behavior and sociology literatures. Examples of such work includes that on the effect of manager-subordinate similarity on subjective outcomes such as performance evaluations, role ambiguity, and job satisfaction (Tsui and O'Reilly, 1989; Thomas, 1990; Wesolowski and Mossholder, 1997), and on how social networks within the firm influence within firm promotions (Podolny and Baron, 1997).

⁶⁵ Lazear (1989), Kandel and Lazear (1992), and Rotemberg (1994) develop models incorporating social concerns into the analysis of behavior within firms. While they emphasize that individuals have social concerns for others at the same tier of the firm hierarchy, their analysis is equally applicable across tiers of the hierarchy. Bewley (1999) offers extensive evidence from interviews with managers arguing that concerns over fair outcomes for workers and the morale of employees are important determinants of their behavior.

Fourth, an increase in the level of social connections between managers and workers has a detrimental effect on the firms' average productivity when managers are paid fixed wages and has no effect when managers are paid performance bonuses. In this setting, social connections are therefore detrimental for the firm because their existence distorts the allocation of managerial effort in favor of lower ability workers.

This natural field experiment paper contributes to the growing empirical evidence on the interplay between social networks and individual and firm performance. In particular, the design allows the authors to identify not only whether social connections matter within the firm, but also exploit the exogenous variation in incentives to understand whether they are to the benefit or detriment of the firm.

Feedback

In a final natural field experiment from this setting, [Bandiera et al. \(2010, forthcoming\)](#) present evidence to evaluate the effect of performance feedback and monetary prize tournaments, when the workforce is organized in teams. Hence in this set-up workers effort imposes a positive externality on their team members, $\frac{\partial B_j(\cdot)}{\partial e_i} > 0$. They compare the effects of these forms of non-monetary and monetary incentives relative to when teams are paid piece rates, and analyze their effect on two outcomes: how workers sort into teams and team productivity.

This field experiment provides important contributions to the literature along three margins. First, despite the pervasiveness of teams in the workplace, field evidence on team incentives is scarce.⁶⁶ The existing evidence from individual reward schemes provides limited guidance because the margins along which individuals and teams can respond to incentives differ. Specifically, in addition to changes in individual effort, changes in team incentives can lead to changes in team composition. To the extent that workers effort depends on the identity of their team members because of social incentives, changes in team composition can affect the productivity of the individual teams and of the firm as a whole.⁶⁷

Second, tournaments are widely used to provide incentives across diverse organizations such as salespeople competing for bonuses, managers competing for promotions, and politicians competing for vote shares ([Bull et al., 1987](#); [Baker et al., 1988](#)). While several studies have tested whether the response to variation in tournament structure is consistent with theoretical predictions, field evidence on the comparison of

⁶⁶ More than 70% of major US firms use some form of team based rewards ([Ledford et al., 1995](#)). [Lazear and Shaw \(2007\)](#) cite evidence that between 1987 and 1996, the share of large firms that have more than a fifth of their employees in problem solving teams rose from 37 to 66%. The percentage of large firms with workers in self-managed teams rose from 27 to 78% over the same period. In academia, [Wuchty et al. \(2007\)](#) document the increased use of team production in research across disciplines.

⁶⁷ There is only a small literature on selection into teams in laboratory settings ([Weber, 2006](#); [Charness and Yang, 2008](#)), although there is a far more extensive lab-based literature on team production, as reviewed in [Charness and Kuhn \(2011, 2010\)](#).

monetary prize tournaments against alternative monetary and non-monetary incentive mechanisms is scarce.⁶⁸

Third, whenever tournaments are in place, workers inevitably receive some information on their relative performance. This information might have direct effects on productivity if individuals have concerns for their relative position or status (Moldovanu et al., 2007; Besley and Ghatak, 2008), inequality aversion (Fehr and Schmidt, 1999; Charness and Rabin, 2002) or conformity (Bernheim, 1994). The field experiment allows the authors to de-couple the effect of feedback from the effect of monetary prize tournaments. As the provision of feedback is almost costless, measuring its contribution to the overall tournament effect can lead to considerable cost savings if most of the positive effect of tournaments on productivity is actually due to worker responses to feedback.⁶⁹

In the experiment, at the beginning of the season, teams were paid piece rates based on their aggregate productivity. Halfway through the season teams were additionally provided feedback by posting daily histograms of each team's productivity. This feedback makes precise the absolute productivity of each team, and their ranking relative to all other teams. Halfway through the remaining part of the season a monetary prize for the most productive team each week was introduced, in addition to the provision of feedback, and conditional on teams being paid according to piece rates.

When workers first arrive at the farm they are assigned to a team by the general manager for their first week. Thereafter workers are free to choose their own team members at a *team exchange* that takes place every week. A team is formed only if all its members agree. Hence in this setting workers have two choice variables: how much effort to exert into picking, and team composition.

The field experiment is again closely tied to an underlying model. This makes precise two key forces that drive team formation: workers' ability and social connections. As individual earnings are increasing in the ability of team members, workers have incentives to assortatively match by ability. On the other hand, workers might prefer to form teams with friends because this might limit free-riding within teams (Alchian and Demsetz, 1972; Holmstrom, 1982; Kandel and Lazear, 1992), and because they enjoy

⁶⁸ The empirical literature on tournament theory comprises two distinct branches. The first tests whether a particular compensation scheme has a tournament structure. Two specific predictions have been explored—(i) the wage spread should be positively related to the number of workers at the lower job level; (ii) the wage structure should be convex as in Rosen (1986). These tests typically use data from the market for CEOs (Gibbons and Murphy, 1990; Eriksson, 1999; Bognanno, 2001). The second branch of the literature tests whether individual behavior changes with tournament features in a way consistent with theory, using data either from experimental settings (Bull et al., 1987; Nalbantian and Schotter, 1997; Eriksson et al., 2008a,b; Freeman and Gelber, 2008), personnel data (Knoeber and Thurman, 1994; Eriksson, 1999; Bognanno, 2001), or sports (Ehrenberg and Bognanno, 1990). There are few existing field studies—on either individuals or teams—exploring tournament incentives to other incentive schemes such as piece rates or feedback.

⁶⁹ Evidence from the laboratory has tended to focus on feedback to individuals (Freeman and Gelber, 2008). One exception is Sausgruber (2009) who provides experimental evidence on the effects on team performance when told about the performance of one other team, holding team composition constant.

non-pecuniary benefits from interacting with co-workers they are socially connected to Rosen (1986); Hamilton et al. (2003).⁷⁰ To the extent that workers are not socially connected to colleagues of similar ability, a trade-off emerges. The theoretical framework then makes precise how the introduction of feedback and prizes affect this trade-off.

The key empirical results from the field experiment are as follows. First, the introduction of feedback and of monetary prizes leads to significant changes in team composition. Relative to the piece rate regime, the share of team members connected by social ties is lower and team members' ability levels are more similar under the feedback and tournament regimes.

Second, the feedback and tournament schemes have opposite effects on average productivity. Relative to the piece rate regime, the introduction of feedback significantly reduces average team productivity by 14%. The further introduction of a monetary prize tournament, conditional on the provision of feedback, significantly increases productivity by 24%. As made precise in the theoretical framework, the reduction in average productivity when feedback is provided is consistent with workers being better off sorting into teams on the basis of ability rather than friendship as feedback increases the strength of incentives faced, and the firm being worse off because it no longer harnesses the ability of socially connected workers to ameliorate free-riding within the team. Hence the endogenous formation of teams under feedback reduces the firm's productivity overall. In contrast, the tournament incentives are sufficiently high-powered so the increase in worker's effort more than offsets any increase in free-riding within teams. Hence the firm's overall productivity rises.

Third, the dispersion of productivity increases under both regimes because both effects are heterogeneous as indicated by the theoretical framework. Quantile regression results show that the introduction of feedback reduces the productivity of teams at the bottom of the conditional productivity distribution compared to piece rates, while it has no effect on teams above the 40th percentile. In contrast, the introduction of prizes increases the productivity of teams at the top of the conditional productivity distribution compared to piece rates, while it has no effect on teams below the 30th percentile.

Fourth, focusing on the teams that remain intact after each change in incentives, the authors evaluate the effect of feedback and prizes on effort, holding constant team composition. They find that while the effect of feedback on team productivity is positive the magnitude appears small. This emphasizes that the documented negative effect of feedback is primarily due to the endogenous changes in team composition caused by the provision of feedback, rather than changes in behavior of the same team. In contrast

⁷⁰ In line with this, Rotemberg (1994) develops a model showing how altruism between co-workers may endogenously form in the workplace to facilitate cooperation among workers engaged in team production. Empirically, Hamilton et al. (2003) provide non-experimental evidence from the introduction of team production in a garment firm. They find the most able workers sorted first into teams despite a loss in earnings in many cases, suggesting non-pecuniary benefits associated with teamwork.

the additional introduction of monetary prizes increases team productivity by 25% for teams that choose to remain intact. Hence the provision of monetary prizes affects firm performance through both the endogenous changes in team composition and changes in behavior within the same team.

Finally, the authors present qualitative evidence from a worker survey they conducted. As highlighted at the start of this section, this type of primary data collection that is inherent in field experiments, allows the authors to shed light on other margins of behavior between workers that might be affected by the monetary and non-monetary incentives provided, but that the firm does not collect data on *ex ante*. This survey data reveals that relative to the piece rate regime, during the tournament regime significantly fewer workers report pushing their team members to work hard or giving team members instructions. This is consistent with workers being better matched by ability and having fewer social connections with their team members under the tournament regime, so that peer pressure within the team becomes less effective.

By exploring changes in behavior on a range of dimensions, this evidence from the field highlights new directions for research in understanding how agents react to monetary and non-monetary incentives in workplaces characterized by team production where teams form endogenously.

4.3. The employment relationship

The neoclassical labor market model emphasizes workers behave opportunistically. For example, in the model sketched above from Lazear (2000), when workers compensation is not tied to their performance, as under a fixed hourly wage scheme, all workers exert the minimum effort required to achieve the minimum output requirement, q_0 . There is thus no variation in workers in their output or pay. We now explore the insights field experiments have provided on the existence and nature of such opportunistic behavior in real world settings. We do so through examples related to gift exchange in shirking.

4.3.1. Gift exchange

The standard labor market model assumes in equilibrium firms pay market clearing wages and workers provide minimum effort. This prediction does not receive uniform support empirically. There are numerous cases where employers are observed paying above the market equilibrium wage Akerlof (1982), and where workers exert more than the minimum effort level, as we have already discussed in relation to Lazear (2000) and in many other studies on employee performance under fixed wages. This has led to the development of the gift-exchange model which is based on the assumption of their being a positive association between wages and worker effort (Akerlof, 1982; Akerlof and Yellen, 1988, 1990). In this class of model, employers offer higher than market clearing

wages, and workers are viewed to positively reciprocate by providing higher than the minimum required effort.

Clearly such theories are hard to test using non-experimental data: there might be a host of unobservable factors that create a correlation between wages and worker effort. Hence, there has been a large body of evidence established in laboratory settings on gift-exchange in firm settings, which began with [Fehr et al. \(1993\)](#). In this original study, they constructed a labor market equilibrium with excess labor supply so that the equilibrium wage was low. Employees also had no pecuniary incentive to raise the quality of their work above the minimum required level, so the best response of employers was to pay the low equilibrium wage. Contrary to the prediction, the majority of employers attempted to induce employees to invest greater effort by offering them higher than market-clearing wages. On average, this high wage was reciprocated by greater employee effort. Overall, it was profitable for employers to offer high wage contracts.

[Gneezy and List \(2006\)](#) use a natural field experiment then look for evidence of gift-exchange in similar real world environments in which equilibrium wages are low and workers earnings are not tied to their performance. In moving from the lab to the field, one important comparative static to evaluate is how behavior changes with the duration of the task. In other words, are the types of positive reciprocity observed by workers in the lab, a long run phenomena. The psychology literature provides two reasons why the duration of tasks might matter. First, there is the distinction between hot and cold decision making ([Loewenstein and Schkade, 1999](#); [Loewenstein, 2005](#)). Second, there can be adaptation of behavior over time ([Gilbert et al., 1998](#)).

Two subject pools were utilized for the field experiments. In each a between subject design was used. The first field experiment recruited undergraduate students to participate in an effort to computerize the holdings of a small library at the university. The task was to enter data regarding the books into a computer database. In the no-Gift treatment, individuals were offered a flat wage of \$12 per hour. In the Gift treatment, once the task was explained to participants, they were surprisingly paid \$20 per hour rather than \$12 per hour as advertised. In total 19 workers were hired for six hours each; 10 were randomly assigned to the no-Gift treatment. The second field experiment was part of a door to door fundraising drive to support a university research center. Fundraising solicitors were recruited. All solicitors were told they would be paid \$10 per hour, and those in the Gift treatment were surprisingly told they would actually receive \$20 per hour. In total 23 solicitors were employed over two days, with 10 being randomly assigned to the no-Gift treatment.⁷¹

⁷¹ In all such experiments, it is important to design the set-up to be able to distinguish gift-exchange from the alternative explanation of why there should be a positive relationship between wages and effort—efficiency wages. This hypothesis postulates employers pay above market-clearing wages to motivate workers to increase their effort level so as to avoid being fired, which reduces employer monitoring ([Katz, 1986](#)). Hence in both field experiments subjects were made aware that this was a one-time recruitment opportunity.




Figure 10 Gift exchange and the duration of tasks.

The main results are as follows. First, in line with earlier evidence from laboratory settings, there are signs of significant gift exchange in the first few hours of the task, as measured both by effort in the library task and money raised in the fundraising task. For example, in the library task, effort is around 25% higher for those in the Gift treatment.

Second, there are significant falls in effort over time. After a few hours, there are no longer any significant differences in effort between the no-Gift and Gift treatments in either task. Figure 10 illustrates clearly how any positive reciprocity by workers is a short run phenomena in these two settings. Overall, the results suggest that with the same budget, the employer would have been better off just paying the market clearing wage as in the no-Gift treatments.

While the results go against the standard gift-exchange explanation of the positive association of wages and effort, they are in line with survey evidence on wage rigidity.⁷²

⁷² Clearly this debate on the existence of positive reciprocity in the field remains in need of further study. Charness and Kuhn (2011, 2010) describe more of the related evidence from the laboratory and other field settings.

For example [Bewley \(1999\)](#) considers why wages are downwardly rigid during a recession. He reports that managers are worried that wage cuts might result in decreases in morale that would subsequently result in poor worker performance when the economy recovered, if not immediately. This highlights the importance of fairness considerations in cases of negative reciprocity. With respect to positive reciprocity, as in this field setting, Bewley's evidence is less conclusive. He argues that morale is less important when considering wage increases, but finds that one main consideration when determining raises is the effect on employee turnover once the recession ends. Bewley's work suggests that there appears to be little connection between increasing pay and productivity, except to the extent that higher wages make it possible to attract, and retain, higher quality workers. This ties back to the earlier discussion of [Lazear \(2000\)](#) and the subsequent work on monetary incentives in firms, where it is thought there are qualitatively large selection effects of incentives driven by changes in workforce composition. Again, more evidence on these channels related to employee turnover are required to test more precisely a fuller set of theoretical predictions.

4.3.2. Shirking

Field experiments had also provided insights on research questions related to employee shirking behaviors. The standard economic framework emphasizes employees are rational shirkers: they will slack when the marginal benefits of doing so outweigh the marginal costs. Firms respond to such behaviors by choosing compensation and monitoring policies to reduce shirking. As emphasized above, this view that workers will behave so opportunistically when the marginal returns on their effort are low, is often contradicted by empirical evidence and predictions on behavior in such settings from the psychological and sociological literatures ([Pfeffer, 1996](#); [Kreps, 1997](#); [Baron and Kreps, 1999](#)). While we have earlier studied the role of compensation schemes and wage setting behavior to raise employee effort, we now focus on the effect employer monitoring has on worker behavior.

If employees are rational cheats then, conditional on a given incentive pay arrangement, a reduction in monitoring will lead to an increase in shirking. The most powerful sanction available to employers is typically dismissal. Thus, an increase in shirking resulting from reduced monitoring should be greatest among individuals for whom the ongoing employment relationship is least valuable.

As with most of the research questions posed in this chapter, establishing credible empirical evidence to support the theory is not straightforward. In this case there are two concerns that have plagued the non-experimental literature. First, shirking behavior is by its nature hard to detect. Moreover, the ability of the econometrician to detect shirkers might itself be endogenously related to the employer's monitoring practices. Second, there might be unobserved factors, such as other hiring policies, that cause there to be a correlation between monitoring and shirking.

A carefully designed natural field experiment of Nagin et al. (2002) addresses both challenges. The setting was a telephone solicitation company, with employees dispersed across 16 call centers. At each call center, telephone solicitors were paid according to the same piece rate incentive scheme, one in which salary increased with the number of successful solicitations. This piece rate, together with imperfect information on the outcome of pledges, created incentives for employees to falsely claim that they had solicited a donation.

To curb opportunistic behavior, the employer monitored for false donations by calling back a fraction of those who had responded positively to a solicitation. Employees were informed when hired that their activities would be checked by “callbacks” made by management. The results of each week’s callbacks were communicated to both employees and their immediate supervisors, and the bad calls were deducted from each individual’s weekly incentive pay. Stronger sanctions for bad calls were not generally imposed on employees because the number of bad calls was understood to be a noisy indicator of cheating. For example, donors might sometimes change their mind after agreeing to pledge money.⁷³

To see if the costs of implementing this monitoring system could be reduced, the company conducted a controlled field experiment. This experiment was “double blind” in the sense that neither the employees nor their immediate supervisors were aware of departures from “business as usual.” In the experiment, the employer varied the fraction of bad calls that were reported back to employees and supervisors at each call center. To more precisely estimate the true rate of fictitious pledges, the firm simultaneously increasing the true callback rate from 10% to 25% of pledges. By working closely with the firm, the researchers were able to collect primary survey data on employee attitudes toward the job, their expected job tenure, and the perceived difficulty of finding another, comparable job. These relate closely to the underlying theory of rational shirking behavior. This information is used to test whether those for whom the job was most valuable were also the employees least likely to engage in opportunistic behavior.

The main results are as follows. First, a significant fraction of employees behave according to the predictions of the rational cheater model. In particular, employees respond to a reduction in the perceived cost of opportunistic behavior by increasing the rate at which they shirk. Using the survey data collected, the authors find the employees who responded to reductions in monitoring tended to be those who perceived the employer as being unfair and uncaring. On the other hand, there is no evidence that individuals with good outside options increased shirking by more than other workers when the rate of monitoring declined. Second, a substantial proportion of employees do

⁷³ Olken (2007) presents evidence from a natural field experiment on the effects of top-down monitoring relative to grassroots participation on reducing corruption on road projects organized by village committees in Indonesia. Top-down monitoring via government audits is found to be the far more effective means of reducing corruption.

not appear to respond at all to manipulations in the monitoring rate. As with responses to monetary and non-monetary incentives documented above, there is considerable heterogeneity in how workers respond to employer monitoring. This underlying heterogeneity highlights the need to balance the need to reduce the shirking behavior of some workers inclined to rationally cheat, against those that are unlikely to do so under normal circumstances.

4.4. Moving forward

Economists have only recently begun to exploit field experiments in firms. This nascent literature has already highlighted the strengths of this methodology in being closely linked to testing alternative theories of individual behavior, of utilizing field experiments and structural modelling to make inference on the optimal design of incentive schemes, and to collect primary data to check for non-expected responses on other margins such as the quality of work, or to probe specific tests of the theory. We conclude by highlighting a few key areas for future work to consider.

First, the set of field experiments discussed have focused primarily, although not exclusively, on job tasks in which productivity is easy to observe, measure, compare across workers and time, and the quality of work performed is relatively easily monitored by management and assignable to individuals. Yet many jobs in the economy, or at higher tiers of firms' hierarchies, do not share such characteristics, and more research is required in such settings where performance is evaluated more subjectively, and might therefore be subject to influence activities (Milgrom, 1988), or favoritism (Prendergast and Topel, 1996). As primary data collection is part of the field experimenter's arsenal, this approach might especially help to shed light on these types of evaluations and incentive structures.

Second, most field experiments have been implemented to evaluate the effects of one time changes in management practices. Standard theory suggests history does not matter and that these effects should be equal and opposite to changing incentives the other way. This would be relatively straightforward to test, conditional on being able to control for natural time effects on behavior. A rejection of the standard model might then imply there can be persistent effects of short run changes in management practice. Such effects might operate through habit formation or reference point effects for example, that have been found using non-experimental data from real world settings (Mas, 2006).

Third, given the progression of field experiments exploring the effects of incentives on bottom-tier workers, and then to managers in the middle tier of the firm hierarchy, it is natural to ask whether field experiments might in the future extend to understanding executive pay. The last two decades have seen a surge in the popularity of performance pay for individuals in executive and managerial positions, from CEOs down to middle and lower management (Hall and Liebman, 1998; Hall and Murphy, 2003;

Oyer and Schaefer, 2004). However as yet there remains mostly an unwillingness of organizations to experiment in relation to such high stakes positions.⁷⁴

Broader methodological issues remain to be borne in mind with regards to field experiments in firms. First, there are concerns over whether the set of firms and organizations that allow field experiments to be conducted within them, are selected in some way. For example, those firms that are most likely to gain from changes in management practices might be most amenable to field experiments on these dimensions. Given the potential for such non-random selection, field experiments ought to be designed to precisely measure differential effects, and less weight given to the levels effects.

Second, this body of field experiments offers an intriguing insight into whether firms choose their management practices optimally. Certainly, Shearer's (2004) study highlights why the firm was using piece rates and not fixed wages. For the firm studied in Bandiera et al., in each case the firm followed up on the results of the field experiment by maintaining the incentives that were introduced. However we have to be careful that while field experiments have focused on the effects of carefully engineered interventions on productivity, the firm chooses practices to maximize discounted profits. Productivity increases need to translate in profit increases. An example of this is in the study by Freeman and Kleiner (2005) on a US shoe manufacturer, who find that the move from piece rates to hourly wages reduced productivity, but increased the quality of work to such an extent that profits rose overall. Clearly, there remains scope for experimentation within firms to help them learn the optimal behaviors, and for this to have a large impact economy-wide, and perhaps go some way to explaining large productivity differences across otherwise observationally similar firms.

5. HOUSEHOLDS

Much of an individual's life cycle is spent in some form of partnership or family union. Despite widespread social changes in family structure in Western economies, families and multi-member households remain a key building block of society. Understanding how households make decisions has implications for many of the choices we have already touched upon, such as educational choices for children, labor market participation and labor supply. Shedding light on the household decision making processes also has profound implications for understanding whether, and how, policies such as income transfers and the regulation of marriage and divorce marriage markets, shape these outcomes.

The benchmark model of household behavior has been the unitary model, pioneered by Samuelson (1956) and Becker (1981). While this generates a rich set of predictions

⁷⁴ A similar set of issues arise for field experiments in public economics. In particular, understanding why individuals give to fundraisers or charitable causes. Large scale field experiments have so far focused on how to induce members of the public or those with affinity to the fundraising organization to give. However a disproportionate amount of funds raised come from a few very wealthy donors. No field experiments have been run on them.

for price and income effects on household behaviors, it remains silent on how conflicts between spouses are resolved. Modelling household decision making as the outcome of a bargaining process provides a natural way in which to introduce conflicts (Manser and Brown, 1980; McElroy and Horney, 1981; Chiappori, 1988). Hence, where these approaches differ is in whether households maximize according to a common or dictatorial set of preferences—the unitary approach—or whether they seek to maximize a weighted sum of household member preferences—the basis of the bargaining approach. On the other hand, a key feature of both modelling frameworks is that households are assumed to make efficient decisions.

Households might reasonably be expected to reach efficient outcomes because they have repeated and long term interactions, in strategic environments characterized by perfect information, and have the ability to communicate costlessly. Nevertheless, a more recent strand of the literature has developed that takes seriously the idea that either household members behave non-cooperatively within marriage (Ulph, 1988; Chen and Woolley, 2001), have private information or an inability to communicate perfectly (Pahl, 1983; Ligon, 1998; Goldstein and Udry, 1999; Boozer and Goldstein, 2003; Dubois and Ligon, 2004), or cannot make binding agreements (Lundberg and Pollak, 2003; Basu, 2006; Mazzocco, 2004, 2007; Rasul, 2008). In each case, household decisions can then be inefficient.

There are two long-standing strands of the empirical literature on household decision making that stem from these views of the world. First, there have been a number of attempts to uncover whether households bargain efficiently, as is implied by both the unitary and collective choice models. Many of these tests take the form of examining patterns of household demand and consumption (Browning and Chiappori, 1998) or testing for the equality of the marginal product of labor of household members across economic activities (Udry, 1994; Akresh, 2005). A first generation of field experiments on households has begun to shed light on this issue.⁷⁵

Second, there is an older strand of the literature that uses non-experimental approaches to test for the assumption on whether households pool income, consistent with the predictions of the unitary framework, or whether the identity of the income earner matters for outcomes (Thomas, 1990, 1994; Hoddinott and Haddad, 1995; Duflo, 2003; Duflo and Udry, 2004; Rangel, 2006). Rather surprisingly given the roots of field experiments in the social experiments of the 1970s, relatively fewer field

⁷⁵ Tests based on demand patterns exploit the fact that utility maximization by a single consumer subject to a linear budget constraint implies Slutsky symmetry, namely the restriction of symmetry on the matrix of compensated price responses. This prediction is typically rejected in household data (Deaton, 1990; Browning and Meghir, 1991; Banks et al., 1997; Browning and Chiappori, 1998). Browning and Chiappori (1998) derive the counterpart to the Slutsky matrix for multi-member households solely under the assumption of efficient within-household decision making, consistent with Nash bargaining models. They show the assumption of efficiency generates testable restrictions on household demand functions, and distinguish the collective model from both the unitary and the entirely unrestricted case.

experiments have been conducted to help test the specific predictions of either unitary or collective bargaining frameworks.

A parallel stream of literature relates to the use of social experiments to evaluate conditional cash transfer programs, which was touched upon earlier. Two notable studies that have used data from the *PROGRESA* intervention in rural Mexico are [Attanasio et al. \(2006\)](#) and [Todd and Wolpin \(2006\)](#). These both combine the experimental variation in *PROGRESA* transfers across randomly assigned villages with structural estimation of a household's dynamic behavior to shed light on outcomes under alternative policy designs.

5.1. Efficiency

[Ashraf \(2009\)](#) presents evidence from a framed field experiment to understand how information and communication affect household financial decisions. The experiment was conducted with a sample of current or former clients of a rural bank in the Philippines. The main decision each subject had to make was over whether to spend or save income received during the experiment. More precisely, subjects had to choose how to allocate 200 pesos received between: (i) direct deposits into their own or a joint account; (ii) committed consumption using redeemable gift certificates. Each subject was randomly assigned, with his or her spouse, to one of three treatments that varied the privacy of information spouses had, and the ability of spouses to communicate with each other. 149 married couples are involved in the experiment.⁷⁶

In the first treatment, subjects are separated from their spouses at the outset of the experiment. This treatment is referred to as “private information without pre-play communication”. Under this treatment spouses have no information on whether and how much income is received by the spouse, what decisions they have made, or the outcomes obtained. In the second treatment spouses learn each others' payoffs and choice sets. In this treatment, referred to as “public information without pre-play communication”, spouses make simultaneous decisions and so cannot communicate nor observe each others decisions *ex ante*. In the final treatment the procedure is as in the previous treatment except that spouses are able to communicate before making their decisions, and their decisions are observable to each other. This is referred to as the “public treatment”.

Clearly, in the absence of a field experiment, trying to uncover and exploit plausibly exogenous sources of variation in the information available to spouses or their ability to communicate is difficult, and likely to be correlated to factors that affect outcomes

⁷⁶ As in any framed field experiment or laboratory experiment, subjects need to be recruited. Framed field experiments that aim to replicate natural settings—say by working in conjunction with local organizations—might provide data from which to assess whether participants differ from those that choose not to participate. As discussed earlier, the nature of self-selection into experiments is a phenomenon that is only beginning to be understood ([Lazear et al., 2009](#)). Equally important, given the relatively small sample sizes inherent to many framed field experiments, it is crucial to be clear on how large a sample would be required to detect statistically different observable characteristics between participants and non-participants.

directly. Hence this research design allows economists to more carefully scrutinize causal changes in behavior along dimensions that are theoretically important, yet empirically almost impossible to measure in the absence of a field experiment. However, as with the other settings considered throughout, field experiments conducted with households raise important issues that need to be taken into consideration when interpreting results.

First, in common with laboratory experiments, behavior in framed field experiments might not mimic behavior in the real world. Ashraf (2009) addresses this concern by running the experiment in conjunction with a rural bank that all participants were familiar with, and by designing treatments that capture real world differences in communication and information across households in this setting. Second, households are engaged in repeated interactions outside of the context of the field experiment. Hence behavior within an experiment can be undone, or potentially reinforced, by behavior outside of the experiment. To try and address this issue, Ashraf provides payoffs in the form of person-specific gift certificates. Both methodological issues need to be considered in all field experiments with households.

Ashraf's (2009) results shed light on the interplay between information, communication and gender in household decision making. Relative to field experiments in other settings, when the experimenter is engaged in primary data collection with households it is of even greater importance to understand societal norms of behavior within marriage. For example, in the Philippines, women are typically in charge of the financial management of the household, making key decisions on budgeting and allocation. Understanding the context in which the experiment takes place is crucial for designing treatments that reflect real world trade-offs subjects face, and to closely align experimental designs with a theoretical framework. Of course, the cost of this precision in any given context is the limited ability, all else equal, to extrapolate findings to households operating under very different norms.

The three main results are as follows. First, men are found to be more likely to deposit money into their own account under the private treatments, and are more likely to commit it to consumption under the public treatment. Second, the differences in behavior by gender are subtle. A subset of women—those whose husbands normally control household savings decisions—behave in the same way as men whose wives normally control household savings decisions. Third, communication between spouses at the time of decision making induces the majority of men to place the income into their spouses account rather than consume it or put it into their own account. To understand these results, Ashraf discusses a framework of income monitoring within the household where observability of income and communication at the time financial decision making, significantly change the monitor's ability to enforce contracts. The results can then be understood as spouses responding strategically to changes in information and communication and contract enforceability. They suggest a specific channel through

which asymmetric information can create inefficient outcomes in financial decision making, by providing incentives to hide one's additional income from one's spouse.⁷⁷

5.2. Moving forward

Labor economists have sought to explain a far richer set of research questions than just those related to behavior within households. Foremost among these other issues has been the research into the causes and consequences of the formation and dissolution of households. Field experiments have recently begun to help explore issues related to the formation of households or partnerships in the first place. Two examples are [Fisman et al. \(2006\)](#) and [Fisman et al. \(2008\)](#) who conduct a framed field experiment to measure differential preferences in dating across genders and races, respectively. To do so, both studies analyze individual choices of subjects in an experimental speed dating game.⁷⁸

On differential preferences across genders, [Fisman et al. \(2006\)](#) find that women place greater weight on the intelligence and the race of partner, while men respond more to physical attractiveness. Moreover, men do not value women's intelligence or ambition when it exceeds their own.

On racial preferences, [Fisman et al. \(2008\)](#) find that there is a strong asymmetry in racial preferences across genders: women of all races exhibit strong same race preferences, while men—of all races—do not. Second, subjects' background influences their racial preferences: subjects that come from locations that are measured to be more racially intolerant, using data from the General Social Survey and World Values Surveys, reveal stronger preferences for same race preferences. This is despite the subject pool being drawn from individuals that currently reside away from home, and attend a top US university. Third, those exposed to other races in early life—as measured by the fraction of individuals of a given race in the zip code where the subject grew up—are less willing to date someone from this race, suggesting that familiarity might reduce racial tolerance. Finally, physically more attractive individuals are less sensitive to the race of potential partners in the experiment.

This experimental approach provides a nice complement to other non-experimental studies applying structural methods to estimate similar preference parameters in the context of online dating services ([Hitsch et al., 2010](#)). Given the growth in availability of online data in economic research, perhaps in the near future we will witness research methods combining field experiments with interventions akin to audit studies that were previously discussed in relation to the economics of discrimination.

⁷⁷ This strand of field experiments is growing for example, [Robinson \(2008\)](#) presents evidence from a framed field experiment on 142 households in Kenya to test whether intra-household risk sharing arrangements are efficient, and if not, whether limited commitment caused by contractual incompleteness partially explains behavior.

⁷⁸ [Stevenson and Wolfers \(2007\)](#) provide a recent overview of the most pressing issues that are being addressed in research in the economics of the family.

As yet though, on many aspects of the formation and dissolution of families, few research designs have credibly exploited experimental sources of variation from which to identify causal effects. The nature of questions involved might mean these sets of research questions remain outside the domain of field experiments.

6. CONCLUDING REMARKS

Given that complexities of markets severely constrain the ability of traditional economic tools to examine behavioral relationships, it is not surprising that economists have increasingly turned to experimental methods. Within this recent trend is a relatively new approach—field experiments—which have dramatically risen in popularity over the past several years. Since field experiments will likely continue to grow in popularity as scholars continue to take advantage of the settings where economic phenomena present themselves, we view this study as an opportunity to step back and discuss a few of the areas within labor economics wherein field experiments have contributed to our economic understanding. Our central task is to highlight what we view to be the central advantages of the field experimental approach: (i) using economic theory to design the null and alternative hypotheses; (ii) engineering exogenous variation in real world economic environments to establish causal relations and learn the mechanisms behind them; and (iii) engaging in primary data collection and often working closely with practitioners.

A second goal of this study is to draw attention to a methodological contribution of field experiments: complementing other empirical approaches and allowing an exploration of the generalizability of behaviors across settings, such as lab and field behavior. When taking account of the stock of evidence, it becomes clear how field experiments can play an important role in the discovery process by allowing one to make stronger inference than can be achieved from lab or uncontrolled data alone. In this way, the various empirical approaches should be thought of as strong complements—much like theory and empirical modeling—and combining insights from each of the methodologies will permit economists to develop a deeper understanding of our science.

REFERENCES

- Abdulkadiroglu, A., Angrist, J.D., Dynarski, S.M., Kane, T.J., Pathak, P., 2009. Accountability and flexibility in public schools: evidence from Boston's charters and pilots. National Bureau of Economic Research Working Paper No. 15549.
- Acemoglu, D., 2002. Technical change, inequality, and the labor market. *Journal of Economic Literature* 40, 7–72.
- Akerlof, G.A., 1980. The theory of social custom, of which unemployment may be one consequence. *Quarterly Journal of Economics* 94, 749–775.
- Akerlof, G.A., 1982. Labor contracts as a partial gift exchange. *Quarterly Journal of Economics* 97, 543–569.
- Akerlof, G.A., Yellen, J.L., 1988. Fairness and unemployment. *American Economic Review Papers and Proceedings* 78, 44–49.
- Akerlof, G.A., Yellen, J.L., 1990. The fair wage–effort hypothesis and unemployment. *Quarterly Journal of Economics* 105, 255–283.

- Akerlof, G.A., Kranton, R., 2005. Identity and the economics of organizations. *Journal of Economic Perspectives* 19, 9–32.
- Akresh, R., 2005. Understanding Pareto inefficient intrahousehold allocations. Institute for the Study of Labor Discussion Paper 1858.
- Alabarran, P., Attanasio, O., 2003. Limited commitment and crowding out of private transfers: evidence from a randomised experiment. *Economic Journal* 113, C77–C85.
- Alchian, A.A., Demsetz, H., 1972. Production, information costs, and economic organizations. *American Economic Review* 62, 777–795.
- Alevy, J., Haigh, M., List, J.A., 2007. Information cascades: evidence from a field experiment with financial market professionals. *Journal of Finance* 62, 151–180.
- Altonji, J.G., Blank, R., 1999. Race and gender in the labor market. In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 3C. Elsevier Science B.V., pp. 3143–3259.
- Altonji, J.G., Pierret, C., 2001. Employer learning and statistical discrimination. *Quarterly Journal of Economics* CXVI, 293–312.
- Anderson, M.L., 2008. Multiple inference and gender differences in the effects of early intervention: a reevaluation of the Abecedarian, Perry preschool, and Early Training projects. *Journal of the American Statistical Association* 103, 1481–1495.
- Andreoni, J., Petrie, R., 2004. Public goods experiments without confidentiality: a glimpse into fundraising. *Journal of Public Economics* 88, 1605–1623.
- Angrist, J.D., Krueger, A.B., 1999. Empirical Strategies in Labor Economics. In: *Handbook of Labor Economics*, vol. 3, Part A. Elsevier, pp. 1277–1366 (Chapter 23).
- Angrist, J.D., Krueger, A.B., 2001. Instrumental variables and the search for identification: from supply and demand to natural experiments. *Journal of Economic Perspectives* 15, 69–85.
- Angrist, J.D., Lang, K., 2004. Does school integration generate peer effects? Evidence from Boston's Metco program. *American Economic Review* 94 (5), 1613–1634.
- Angrist, J.D., Lavy, V., 1999. Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics* 114, 533–575.
- Angrist, J.D., Lavy, V., 2002. New evidence on classroom computers and pupil learning. *Economic Journal* 112, 735–765.
- Angrist, J.D., Lavy, V., 2009. The effects of high stakes high school achievement awards: evidence from a randomized trial. *American Economic Review* 99 (4), 1384–1414.
- Angrist, J.D., Bettinger, E., Bloom, E., King, E., Kremer, M., 2002. Vouchers for private schooling in Colombia: evidence from a randomized natural experiment. *American Economic Review* 92, 1535–1558.
- Angrist, J.D., Bettinger, E., Kremer, M., 2006. Long-term educational consequences of secondary school vouchers: evidence from administrative records in Colombia. *American Economic Review* 96, 847–862.
- Angrist, J.D., Barski, S.M., Kane, T.J., Pathak, P., Walters, C.R., 2010. Who benefits from KIPP? National Bureau of Economic Research Working Paper No. 15740.
- Aoyagi, M., 2007. Information feedback in a dynamic tournament. mimeo, Osaka University.
- Arabsheibani, G.R., Marin, A., Wadsworth, J., 2005. Gay pay in the UK. *Economica* 72, 333–347.
- Arcidiacono, P., 2005. Affirmative action in higher education: how do admission and financial aid rules affect future earnings? *Econometrica* 73, 1477–1524.
- Arrow, K., 1972. The theory of discrimination. In: Ashenfelter, O., Rees, A. (Eds.), *Discrimination in Labor Markets*. Princeton University Press, Princeton, NJ.
- Ashenfelter, O., 1990. Nonparametric estimates of the labor-supply effects of negative income tax programs. *Journal of Labor Economics* 8, S396–S415.
- Ashraf, N., 2009. Spousal control and intra-household decision making: an experimental study in the Philippines. *American Economic Review* 99, 1245–1277.
- Attanasio, O., Barr, A., Cardenas, J.C., Genicot, G., Meghir, C., 2009. Risk pooling, risk preferences, and social networks. mimeo, Oxford University.
- Attanasio, O., Meghir, C., Santiago, A., 2006. Education choices in Mexico: using a structural model and a randomised experiment to evaluate *progres*a. IFS Working Paper.
- Autor, D., Katz, L., Kearney, M., 2005. Rising wage inequality: the role of composition and prices. mimeo, Harvard University.

- Azmat, G., Iriberry, N., 2009. The importance of relative performance feedback information: evidence from a natural experiment using high school students. mimeo, Universitat Pompeu Fabra.
- Baker, G., Jensen, M.C., Murphy, K.J., 1988. Compensation and incentives: practice vs. theory. *Journal of Finance* 43, 593–616.
- Ball, S., Eckel, C., Grossman, P.J., Zame, W., 2001. Status in markets. *Quarterly Journal of Economics* 116, 161–188.
- Bandiera, O., Barankay, I., Rasul, I., 2005. Social preferences and the response to incentives: evidence from personnel data. *Quarterly Journal of Economics* 120, 917–962.
- Bandiera, O., Barankay, I., Rasul, I., 2007. Incentives for managers and inequality among workers: evidence from a firm level experiment. *Quarterly Journal of Economics* 122, 729–774.
- Bandiera, O., Barankay, I., Rasul, I., 2009a. Team incentives: evidence from a field experiment, mimeo, University College London.
- Bandiera, O., Barankay, I., Rasul, I., 2010. Social incentives in the workplace. *Review of Economic Studies* 77 (April), 417–459.
- Bandiera, O., Larcinese, V., Rasul, I., 2009b. Blissful ignorance? Evidence from a natural experiment on the effect of individual feedback on performance. mimeo, LSE.
- Bandiera, O., Larcinese, V., Rasul, I., 2010. Heterogeneous class size effects: new evidence from a panel of university students. *Economic Journal* (forthcoming).
- Banerjee, A.V., Suraj, J., Kremer, M., Lanjouw, J., Lanjouw, P., 2001. Promoting school participation in rural rajasthan: results from some prospective trials. mimeo, Massachusetts Institute of Technology.
- Banerjee, A.V., Cole, S., Duflo, E., Linden, L., 2007. Remedying education: evidence from two randomized experiments in India. *Quarterly Journal of Economics* 122, 1235–1264.
- Banks, J., Blundell, R., Lewbel, A., 1997. Quadratic Engel curves and consumer demand. *Review of Economics and Statistics* 79, 527–539.
- Barnard, C., 1938. *The Functions of the Executive*. Harvard University Press, Cambridge.
- Baron, J.N., Kreps, D.M., 1999. *Strategic Human Resources: Frameworks for General Managers*. John Wiley and Sons, Inc., New York.
- Baron, J.N., Pfeffer, J., 1994. The social psychology of organizations and inequality. *Social Psychology Quarterly* 57, 190–209.
- Barr, A., Serneels, P., 2009. Reciprocity in the workplace. *Experimental Economics* 12, 99–112.
- Barrera-Osorio, F., Bertrand, M., Linden, L., Perez-Calle, F., 2008. Conditional cash transfers in education: design features. *Peer and Sibling Effects: Evidence from a Randomized Experiment in Colombia*, The World Bank: Impact Evaluation Series No. 20.
- Basu, K., 2006. Gender and say: a model of household behavior with endogenously-determined balance of power. *Economic Journal* 116, 558–580.
- Becker, G.S., 1957. *The Economics of Discrimination*, second ed., University of Chicago Press, Chicago.
- Becker, G.S., 1981. *A Treatise on the Family*. Harvard University Press, Cambridge.
- Beer, M., 1990. *Performance Appraisal*. mimeo, Harvard Business School.
- Behrman, J.R., Cheng, Y., Todd, P.E., 2004. Evaluating preschool programs when length of exposure to the program varies: a nonparametric approach. *Review of Economics and Statistics* 86, 108–132.
- Benabou, R., Tirole, J., 2000. Self-confidence and social interactions. NBER Working Paper 7585.
- Benabou, R., Tirole, J., 2003. Intrinsic and extrinsic motivation. *Review of Economic Studies* 70, 489–520.
- Berg, N., Lien, D., 2002. Measuring the effect of sexual orientation on income: evidence of discrimination? *Contemporary Economic Policy* 20, 394–414.
- Bernheim, D., 1994. A theory of conformity. *Journal of Political Economy* 102, 841–877.
- Berry, J., 2009. Child control in education decisions: an evaluation of targeted incentives to learn in India. mimeo, Massachusetts Institute of Technology.
- Bertrand, M., Mullainathan, S., 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* 94, 991–1013.
- Bertrand, M., Schoar, A., 2003. Managing with style: the effect of managers on firm policies. *Quarterly Journal of Economics* 118, 1169–1208.
- Besley, T.J., Ghatak, M., 2005. Competition and incentives with motivated agents. *American Economic Review* 95, 616–636.

- Besley, T.J., Ghatak, M., 2008. Status incentives. *American Economic Review Papers and Proceedings* 98, 206–211.
- Bewley, T.F., 1999. *Why Wages Don't Fall During a Recession*. Harvard University Press, Cambridge.
- Bhaskar, V., 2002. Relative performance evaluation and limited liability. mimeo, University of Essex.
- Bhattacharya, D., Dupas, P., 2010. Inferring welfare maximizing treatment assignment under budget constraints. NBER Working Paper 14447.
- Bjorklund, A., Regner, H., 1996. Experimental evaluation of european labour market policy. In: Schmid, G., O'Reilly, J. (Eds.), *International Handbook of Labour Market Policy and Evaluation*. pp. 89–115 (Chapter 3).
- Blanes-i-Vidal, J., Nossol, M., 2009. Tournaments without prizes: evidence from personnel records. mimeo LSE.
- Bloom, N., Van Reenen, J., 2011. Human resource management and productivity. In: *New Developments and Research on Labor Markets*, first ed., In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 4B. Elsevier, pp. 1697–1767 (Chapter 19).
- Blundell, R., Costa-Dias, M., 2002. Alternative approaches to evaluation in empirical microeconomics. *Portuguese Economic Journal* 1, 91–115.
- Bleakley, H., 2007. Disease and development: evidence from hookworm eradication in the American South. *Quarterly Journal of Economics* 122, 73–112.
- Bobonis, G.J., Miguel, E., Puri-Sharma, C., 2006. Anemia and school participation. *Journal of Human Resources* 41, 692–721.
- Bognanno, M.L., 2001. Corporate tournaments. *Journal of Labor Economics* 19, 290–315.
- Bohm, P., 1984. Revealing demand for an actual public good. *Journal of Public Economics* 24, 135–151.
- Boozer, M., Goldstein, M.P., 2003. Poverty measurement and dynamics. mimeo, Yale University.
- Bramel, D., Friend, R., 1981. Hawthorne, the myth of the docile worker, and class bias in psychology. *American Psychologist* 36, 867–878.
- Brown-Kruse, J., Cronshaw, M.B., Schenk, D.J., 2007. Theory and experiments on spatial competition. *Economic Inquiry* 31, 139–165.
- Browning, M., Chiappori, P.-A., 1998. Efficient intra-household allocations: a general characterisation and empirical tests. *Econometrica* 66, 1241–1278.
- Browning, M., Meghir, C., 1991. The effects of male and female labor supply on commodity demands. *Econometrica* 59, 925–951.
- Bull, C., Schotter, A., Weigelt, K., 1987. Tournaments and piece rates: an experimental study. *Journal of Political Economy* 95, 1–33.
- Burgess, S., Propper, C., Slater, H., Wilson, D., 2005. Who wins and who loses from school accountability? The distribution of educational gain in English secondary schools. CMPO Working Paper Series 05/128.
- Butler, R., 1987. Task-involving and ego-involving properties of evaluation: effects of different evaluation conditions on motivational perceptions, interest, and performance. *Journal of Educational Psychology* 79, 474–482.
- Cadsby, C.B., Song, F., Tapon, F., 2007. Sorting and incentive effects of pay for performance: an experimental investigation. *Academy of Management Journal* 50, 387–405.
- Cahuc, P., Postel-Vinay, F., Robin, J.-M., 2006. Wage bargaining with on-the-job search: a structural econometric model. *Econometrica* 74, 323–364.
- Cameron, J., Pierce, W.D., 1994. Reinforcement, reward, and intrinsic motivation: a meta-analysis. *Review of Educational Research* 64, 363–423.
- Cameron, S.V., Heckman, J.J., 2001. The dynamics of educational attainment for black, hispanic, and white males. *Journal of Political Economy* 109, 455–499.
- Cameron, S.V., Taber, C., 2004. Estimation of educational borrowing constraints using returns to schooling. *Journal of Political Economy* 112, 132–182.
- Card, D., 1999. The causal effect of education on earnings. In: *Handbook of Labor Economics*, first ed., In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 3. Elsevier, pp. 1801–1863 (Chapter 30).
- Card, D., Krueger, A.B., 1992. School quality and black-white relative earnings: a direct assessment. *Quarterly Journal of Economics* 107, 151–200.

- Carpenter, J., Seki, E., 2010. Do social preferences increase productivity? Field experimental evidence from fisherman in Toyama bay. *Economic Inquiry*, doi:10.1111/j.1465-7295.2009.00268.x.
- Case, A., Deaton, A., 1999. School inputs and educational outcome in South Africa. *Quarterly Journal of Economics* 114, 1047–1084.
- Charles, K.K., Guryan, J., 2008. Prejudice and Wages: an empirical assessment of Becker's The Economics of Discrimination. *Journal of Political Economy* 116, 773–809.
- Charness, G., Rabin, M., 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117, 817–869.
- Charness, G., Yang, C.-L., 2008. Endogenous group formation and public goods provision: exclusion, exit, mergers, and redemption. UC Santa Barbara Department of Economics Working Paper 13–08.
- Charness, G., Kuhn, P., 2010. Lab labor: what can labor economists learn from the lab? NBER Working Paper 15913.
- Charness, G., Kuhn, P., 2011. Lab labor: What can labor economists learn from the lab? In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 4A. Elsevier, pp. 229–331.
- Charness, G., Villeval, M.-C., 2009. Cooperation and competition in intergenerational experiments in the field and laboratory. *American Economic Review* 99, 956–978.
- Chen, Z., Woolley, F., 2001. A Cournot-Nash model of family decision making. *Economic Journal* 111, 722–748.
- Chiappori, P.-A., 1988. Rational household labor supply. *Econometrica* 56, 63–90.
- Chiappori, P.-A., Salanie, B., 2003. Testing contract theory: a survey of some recent work. In: Dewatripont, M., Hansen, L.P., Turnovsky, S.J. (Eds.), *Advances in Economics and Econometrics Theory and Applications: Eighth World Congress Volume 1*. Cambridge University Press, Cambridge.
- Cipriani, M., Guarino, A., 2009. Herd behavior in financial markets: an experiment with financial market professionals. *Journal of the European Economic Association* 7, 206–233.
- Clark, D., 2009. The performance and competitive effects of school autonomy. *Journal of Political Economy* 117, 745–783.
- Coate, S., Loury, G.C., 1993. Will affirmative-action policies eliminate negative stereotypes? *American Economic Review* 83, 1220–1240.
- Coleman, J.S., et al., 1966. *Equality of Educational Opportunity*. US Government Printing Office, Washington, DC.
- Cornwell, C., Mustard, D.B., Sridhar, D.J., 2006. The enrollment effects of merit-based financial aid: evidence from Georgia's HOPE program. *Journal of Labor Economics* 24, 761–786.
- Cullen, J.B., Jacob, B.A., Levitt, S.D., 2006. The effect of school choice on participants: evidence from randomized lotteries. *Econometrica* 74, 1191–1230.
- Cunha, F., Heckman, J.J., 2009. The economics of psychology of inequality and human development. *Journal of the European Economic Association* 7, 320–364.
- Cooper, D.J., Kagel, J.H., Lo, W., Gu, Q.L., 1999. Gaming against managers in incentive systems: experimental results with Chinese students and Chinese managers. *American Economic Review* 89, 781–804.
- Currie, J., 2001. Early childhood education programs. *Journal of Economic Perspectives* 15, 213–238.
- Currie, J., Thomas, D., 1995. Does head start make a difference? *American Economic Review* 85 (3), 341–364.
- Currie, J., Thomas, D., 2000. School quality and the longer term effects of head start. *Journal of Human Resources* 35, 755–774.
- Deaton, A.S., 1990. Price elasticities from survey data: extensions and Indonesian results. *Journal of Econometrics* 44, 281–309.
- Dobbie, W., Fryer, R.G., 2009. Are high quality schools enough to close the achievement gap? Evidence from social experiment in Harlem. National Bureau of Economic Research Working Paper No. 15473.
- De Giorgi, G., Pellizzari, M., Woolston, W.G., 2009. Class-size and class heterogeneity. mimeo, Stanford University.
- Deci, E.L., Koestner, R., Ryan, R.M., 2001. Extrinsic rewards and intrinsic motivation in education: reconsidered once again. *Review of Educational Research* 71, 1–27.
- Delfgaauw, J., Dur, R., 2004. Incentives and workers' motivation in the public sector. Tinbergen Institute Discussion Papers 04-060/1.

- Dewatripont, M., Jewitt, I., Tirole, J., 1999. The economics of career concerns, Part II: application to missions and accountability of government agencies. *Review of Economic Studies* 66, 199–217.
- Ding, M., Grewal, R., Liechty, J., 2005. Incentive-aligned conjoint analysis. *Journal of Marketing Research* 42, 67–83.
- Dingwall, R., 1980. Ethics and Ethnography. *Sociological Review* 28, 871–891.
- Dixit, A., 2002. Incentives and organizations in the public sector: an interpretative review. *Journal of Human Resources* 37, 696–727.
- Dohmen, T.J., Falk, A., 2006. Performance pay and multi-dimensional sorting: productivity, preferences and gender. Institute for the Study of Labor Discussion Paper No. 2001.
- Doolittle, F., Traeger, L., 1990. Implementing the national JPTA study. Department of Labor, Washington DC.
- Dubois, P., Ligon, E., 2004. Incentives and nutrition for rotten kids: intrahousehold food allocation in the Philippines. mimeo, UC Berkeley.
- Duflo, E., 2003. Grandmothers and granddaughters: old age pension and intra-household allocation in South Africa. *World Bank Economic Review* 17, 1–25.
- Duflo, E., Udry, C., 2004. Intrahousehold resource allocation in Cote d'Ivoire: social norms. Separate Accounts, and Consumption Choices, NBER Working Paper No. 10498.
- Duflo, E., Gale, W., Iebman, J., Orszag, P., Saez, E., 2006. Saving incentives for low- and middle-income families: evidence from a field experiment with H&R block. *Quarterly Journal of Economics* 121, 1311–1346.
- Duflo, E., Dupas, P., Kremer, M., 2009. Peer effects, teacher incentives and the impact of tracking: evidence from a randomized evaluation in Kenya. National Bureau of Economic Research Working Paper No. 14475.
- Dunford, F.W., 1990. Random assignment: practical considerations from field experiments. *Evaluation and Program Planning* 13, 125–132.
- Duesenberry, J.S., 1949. *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.
- Ederer, F., 2008. Feedback and Motivation in Dynamic Tournaments, mimeo, MIT.
- Ederer, F., Fehr, E., 2007. Deception and incentives: how dishonesty undermines effort provision. IZA Discussion Paper 3200.
- Ellingsen, T., Johannesson, M., 2008. Pride and prejudice: the human side of incentive theory. *American Economic Review* 98, 990–1008.
- Encinosa, W.E., Gaynor, M.S., Rebitzer, J.B., 1997. The sociology of groups and the economics of incentives: theory and evidence on compensation systems. National Bureau of Economic Research Working Paper 5953.
- Engle, P., Black, M., Behrman, J., Cabral De Mello, M., Gertler, P., Kapiriri, L., Martorell, R., Young, M., 2007. Strategies to avoid the loss of developmental potential in more than 200 million children in the developing world. *The Lancet* 369, 229–242.
- Epple, D., Romano, R.E., 1998. Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review* 88, 33–62.
- Epple, D., Romano, R.E., Sieg, H., 2006. Admission, tuition, and financial aid policies in the market for higher education. *Econometrica* 74, 885–928.
- Ehrenberg, R., Bognanno, M., 1990. The incentive effects of tournaments revisited: evidence from the European PGA tour. *Industrial Labor Relations Review* 43, 74–89.
- Eriksson, T., 1999. Executive compensation and tournament theory: empirical tests on Danish data. *Journal of Labor Economics* 17, 262–280.
- Eriksson, T., Poulsen, A., Villeval, M.-C., 2008a. Feedback and incentives: experimental evidence. IZA Discussion Paper 3440.
- Eriksson, T., Teyssier, S., Villeval, M.-C., 2008b. Self-selection and the efficiency of tournaments. *Economic Inquiry* 47, 530–548.
- Ferber, R., Hirsch, W.Z., 1982. *Social Experimentation and Economic Policy*. Cambridge University Press, London.
- Fehr, E., Falk, A., 1999. Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy* 107, 106–134.

- Fehr, E., Fischbacher, U., 2002. Why social preferences matter – the impact of non–selfish motives on competition, cooperation and incentives. *Economic Journal* 112, C1–C33.
- Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* 114, 817–868.
- Fehr, E., Schmidt, K.M., 2000. Fairness, incentives, and contractual choices. *European Economic Review* 44 (4–6), 1057–1068.
- Fehr, E., Kirchsteiger, G., Riedl, A., 1993. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108, 437–459.
- Fehr, E., Goette, L., Huffman, D., 2004. Loss aversion and labor supply. *Journal of the European Economic Association* 2, 216–228.
- Ferreira, M.M., 2007. Estimating the effects of private school vouchers in multidistrict economies. *American Economic Review* 97, 789–817.
- Fershtman, C., Hvide, H.K., Weiss, Y., 2003. Cultural Diversity, Status Concerns and the Organization of Work. CEPR Discussion Paper No. 3982.
- Figlio, D., Getzler, L., 2002. Accountability, Ability and Disability: Gaming the System. National Bureau of Economic Research Working Paper 9307.
- Figlio, D.N., Kenny, L.W., 2007. Individual teacher incentives and student performance. *Journal of Public Economics* 91, 901–914.
- Figlio, D.N., Rouse, C.E., 2006. Do accountability and voucher threats improve low-performing schools? *Journal of Public Economics* 90 (1–2), 239–255.
- Fisher, R.A., 1926. The arrangement of field trials. *Journal of the Ministry of Agriculture of Great Britain* 33, 503–513.
- Fisman, R., Iyengar, S., Kamenica, E., Simonson, I., 2006. Gender differences in mate selection: evidence from a speed dating experiment. *Quarterly Journal of Economics* 121, 673–697.
- Fisman, R., Iyengar, S., Kamenica, E., 2008. Racial preferences in dating: evidence from a speed dating experiment. *Review of Economic Studies* 75, 117–132.
- Francois, P., 2000. Public service motivation as an argument for government provision. *Journal of Public Economics* 78, 275–299.
- Frank, R.H., 1985. *Choosing the right pond: human behavior and the quest for status*. Oxford University Press, New York.
- Franke, R.H., Kaul, J.D., 1978. The Hawthorne experiments: first statistical interpretation. *American Sociological Review* 43, 623–643.
- Freeman, R.B., 1987. *Labour Economics*. The New Palgrave Dictionary of Economics, first edition.
- Freeman, R.B., Kleiner, M.M., 2005. The last American shoe manufacturers: decreasing productivity and increasing profits in the shift from piece rates to continuous flow production. *Industrial Relations* 44, 307–330.
- Freeman, R.B., Gelber, A.M., 2008. Prize Structure and Information in Tournaments: Experimental Evidence. mimeo, Harvard University.
- Friedman, M., Savage, L.J., 1948. The utility analysis of choice involving risk. *Journal of Political Economy* 56, 279–304.
- Fryer, R.G., 2011. Racial inequality in the 21st century: the declining significance of discrimination. In: *New Developments and Research on Labor Markets*, first ed., In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 4B. Elsevier, pp. 855–971 (Chapter 10).
- Gale, E.A.M., 2004. The Hawthorne studies—a fable for our times? *Quarterly Journal of Medicine* 97, 439–449.
- Garces, E., Thomas, D., Currie, J., 2002. Longer-term effects of Head Start. *American Economic Review* 92, 999–1012.
- Gibbons, R.S., 1987. Piece-rate incentive schemes. *Journal of Labor Economics* 5, 413–429.
- Gibbons, R.S., Murphy, K.J., 1990. Relative performance evaluation for chief executive officers. *Industrial Labor Relations Review* 43, 30–52.
- Gibbs, M., 1991. *An Economic Approach to Process in Pay and Performance Appraisals*. mimeo, University of Chicago GSB.
- Gilbert, D.T., Pines, E.C., Wilson, T.D., Blumberg, S.J., Wheatley, T.P., 1998. Immune neglect: a source of durability bias in affective forecasting. *Journal of Personality and Social Psychology* 75, 617–638.

- Gillespie, R., 1991. *Manufacturing Knowledge: A History of the Hawthorne Experiments*. Cambridge University Press, New York.
- Gine, X., Karlan, D., Zinman, J., 2007. *The Risk of Asking: Measurement Effects from a Baseline Survey in an Insurance Takeup Experiment*. Working paper, World Bank.
- Glewwe, P., Kremer, M., 2006. Schools, teachers, and education outcomes in developing countries. In: *Handbook of the Economics of Education*. Elsevier, pp. 945–1017 (Chapter 16).
- Glewwe, P., Kremer, M., Ilias, N., 2003. *Teacher Incentives*. National Bureau of Economic Research Working Paper No. 9671.
- Glewwe, P., Kremer, M., Moulin, S., Zitzewitz, E., 2004. Retrospective vs. prospective analyses of school inputs: the case of flip charts in Kenya. *Journal of Development Economics* 74, 251–268.
- Gneezy, U., List, J.A., 2006. Putting behavioral economics to work: testing for gift exchange in labor markets using field experiments. *Econometrica* 74, 1365–1384.
- Gneezy, U., List, J.A., Price, M.K., 2010. *Clean Evidence of Statistical Discrimination*. Working Paper, University of Chicago.
- Gneezy, U., Rustichini, A., 2000. Pay enough or don't pay at all. *Quarterly Journal of Economics* 115, 791–810.
- Goldin, C., Rouse, C., 2000. Orchestrating impartiality: the impact of “blind” auditions on female musicians. *American Economic Review* 99, 715–741.
- Goldstein, M.P., Udry, C.R., 1999. *Gender and Land Resource Management in Southern Ghana*. mimeo, Yale University.
- Grantham-McGregor, S., Powell, C., Walker, S., Himes, J., 1991. Nutritional supplementation, psychosocial stimulation, and mental development of stunted children: the Jamaican study. *The Lancet* 338, 1–5.
- Grantham-McGregor, S., Cheung, Y., Cueto, S., Glewwe, P., Richter, L., Strupp, B., 2007. Developmental potential in the first 5 years for children in developing countries. *The Lancet* 369, 60–70.
- Green, J.R., Stokey, N.L., 1983. A comparison of tournaments and contracts. *Journal of Political Economy* 91, 349–364.
- Greenberg, J., 1988. Equity and workplace status: a field experiment. *Journal of Applied Psychology* 73, 606–613.
- Greenberg, D., Shroder, M., Onstott, M., 1999. The social experiment market. *Journal of Economic Perspectives* 13, 157–172.
- Greenberg, D., Shroder, M., 2004. *The Digest of Social Experiments*. The Urban Institute Press, Washington.
- Groves, T., Hong, Y., Mcmillan, J., Naughton, B., 1994. Autonomy and incentives in Chinese state enterprises. *Quarterly Journal of Economics* 109, 183–209.
- Hall, B., Liebman, J., 1998. Are CEOs really paid like bureaucrats? *Quarterly Journal of Economics* 113, 653–691.
- Hall, B., Murphy, K.J., 2003. The trouble with stock options. *Journal of Economic Perspectives* 17, 49–70.
- Hamilton, B.H., Nickerson, J.A., Owan, H., 2003. Team incentives and worker heterogeneity: an empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy* 111, 465–497.
- Hanushek, E.A., 1986. The economics of schooling: production and efficiency in public schools. *Journal of Economic Literature* 24, 1141–1177.
- Hanushek, E.A., 2006. School resources. In: *Handbook of the Economics of Education*. Elsevier, pp. 865–908 (Chapter 14).
- Hanushek, E.A., 2007. Some US evidence on how the distribution of educational outcomes can be changed. In: *Schools and the Equal Opportunity Problem*. MIT Press, pp. 159–190 (Chapter 7).
- Hanushek, E.A., Raymond, M.E., 2004. The effect of school accountability systems on the level and distribution of student achievement. *Journal of the European Economic Association* 2, 406–445.
- Harrison, G., List, J., 2004. Field experiments. *Journal of Economic Literature* 152, 1009–1055.
- Hart, O., Holmstrom, B., 1987. The theory of contracts. In: Bewley, T. (Ed.), *Advances in Economic Theory*. Cambridge University Press, Cambridge.
- Hastings, J.S., Weinstein, J.M., 2008. Information, school choice, and academic achievement: evidence from two experiments. *Quarterly Journal of Economics* 123, 1373–1414.

- Hausman, J.A., Wise, D.A., 1979. Attrition bias in experimental and panel data: the Gary income maintenance experiment. *Econometrica* 47, 455–473.
- Hausman, J., Wise, D. (Eds.), 1985. *Social Experimentation*. University of Chicago Press for National Bureau of Economic Research, Chicago, pp. 1–55.
- Heckman, J.J., 1992. Randomization and social policy evaluation. In: Manski, C.F., Garfinkel, I. (Eds.), *Evaluating Welfare and Training Programs*. Harvard University Press, Cambridge.
- Heckman, J.J., Masterov, D., 2005. The Productivity Argument for Investing in Young Children. National Bureau of Economic Research Working Paper Number 13016.
- Heckman, J.J., Siegelman, P., 1993. The urban institute audit studies: their methods and findings. In: Fix, M., Struyk, R. (Eds.), *Clear and Convincing Evidence: Measurement of Discrimination in America*. The Urban Institute Press, Washington, DC.
- Heckman, J.J., Smith, J.A., 1995. Assessing the case for social experiments. *Journal of Economic Perspectives* 9, 85–110.
- Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P.A., Yavitz, A., 2010. The rate of return to the high/scope Perry preschool program. *Journal of Public Economics* 94, 114–128.
- Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P., Yavitz, A., 2010. Analyzing social experiments as implemented: a reexamination of the evidence from the high scope Perry preschool program. *Quantitative Economics* 1 (1) (forthcoming).
- Hitsch, G., Hortacsu, A., Ariely, D., 2010. Matching and sorting in online dating. *American Economic Review* 100, 130–163.
- Hoddinott, J., Haddad, L., 1995. Does female income share influence household expenditures? Evidence from Côte d'Ivoire. *Oxford Bulletin of Economics and Statistics* 57, 77–96.
- Holmstrom, B., 1982. Moral hazard in teams. *Bell Journal of Economics* 13, 324–340.
- Holmstrom, B., Milgrom, P., 1991. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization* 7, 24–52.
- Homan, R., 1991. *The Ethics of Social Research*. Longman, London.
- Hossain, T., List, J.A., 2009. The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations. NBER Working Paper 15623.
- Hotz, V.J., 1992. Designing an evaluation of JTPA. In: Manski, C.F., Garfinkel, I. (Eds.), *Evaluating Welfare and Training Programs*. Harvard University Press, Cambridge.
- Hoxby, C.M., 2000a. The effects of class size on student achievement: new evidence from natural population variation. *Quarterly Journal of Economics* 116, 1239–1286.
- Hoxby, C.M., 2000b. Does competition among public schools benefit students and taxpayers? *American Economic Review* 90, 1209–1238.
- Hoxby, C.M., 2000c. Peer Effects in the Classroom: Learning from Gender and Race Variation. National Bureau of Economic Research Working Paper No. 7867.
- Hoxby, C.M., Muraka, S., 2009. Charter Schools in New York City: Who Enrolls and How They Affect Their Students' Achievement. National Bureau of Economic Research Working Paper No. w14852.
- Hoxby, C.M., Weingarh, G., 2006. Taking Race Out of the Equation: School Reassignment and the Structure of Peer Effects. mimeo, Harvard University.
- Hsieh, C.T., Urquiola, M., 2006. The effects of generalized school choice on achievement and stratification: evidence from Chile's voucher program. *Journal of Public Economics* 90 (8–9), 1477–1503.
- Ichniowski, C., Prennushi, G., Shaw, K., 1997. The effects of human resource management practices on productivity. *American Economic Review* 86, 291–313.
- Ichniowski, C., Shaw, K., 2008. Insider econometrics: a roadmap to estimating models of organizational performance. In: Gibbons, R., Roberts, J. (Eds.) *Handbook of Organizational Economics* (forthcoming).
- Jacob, B.A., 2002. Accountability, Incentives and Behavior: The Impact of High-Stakes Testing in the Chicago Public Schools. National Bureau of Economic Research Working Paper 8968.
- Jacob, B.A., 2004. Public housing, housing vouchers, and student achievement: evidence from public housing demolitions in Chicago. *American Economic Review* 94, 233–258.
- Jacob, B.A., 2005. Accountability, incentives and behavior: the impact of high-stakes testing in the Chicago public schools. *Journal of Public Economics* 89, 761–796.
- Jones, D., Kato, T., 1995. The productivity effects of employee stock-ownership plans and bonuses: evidence from Japanese panel data. *American Economic Review* 85, 391–414.

- Jowell, R., Prescott-Clarke, P., 1970. Racial discrimination and white-collar workers in Britain. *Race* 11, 397–417.
- Juhn, C., Murphy, K.M., Pierce, B., 1993. Wage inequality and the rise in returns to skill. *Journal of Political Economy* 101, 410–442.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decisions under risk. *Econometrica* 47, 313–327.
- Kandel, E., Lazear, E., 1992. Peer pressure and partnerships. *Journal of Political Economy* 100, 801–813.
- Katkar, R., Reiley, D.H., 2006. Public versus secret reserve prices in eBay auctions: results from a Pokémon field experiment. *Advances in Economic Analysis and Policy* 6, Article 7.
- Katz, L.F., 1986. Efficiency wage theories: a partial evaluation. *NBER Macroeconomics Annual* 1, 235–276.
- Katz, L.F., Kling, J.R., Liebman, J.B., 2001. Moving to opportunity in Boston: early results of a randomized mobility experiment. *Quarterly Journal of Economics* 116, 607–654.
- Keane, M.P., 2010. A structural perspective on the experimentalist school. *Journal of Economic Perspectives* 24 (2), 47–58 (forthcoming).
- Keane, M., Wolpin, K.I., 1997. The career decisions of young men. *Journal of Political Economy* 105, 473–522.
- Kluger, A.N., Denisi, A.S., 1996. Effects of feedback intervention on performance: a historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin* 119, 254–284.
- Knoeber, C., Thurman, W., 1994. Testing the theory of tournaments: an empirical analysis of broiler production. *Journal of Labor Economics* 12, 155–179.
- Kramer, M.S., Shapiro, S.H., 1984. Scientific challenges in the application of randomised trials. *Journal of the American Medical Association* 252, 2739–2745.
- Kremer, M., 2003. Randomized evaluations of educational programs in developing countries: some lessons. *American Economic Review* 90, 102–106.
- Kremer, M., 2006. Expanding educational opportunity on a budget: lessons from randomized evaluations. In: *Improving Education Through Assessment, Innovation, and Evaluation*, American Academy of Arts and Sciences: Project on Universal Basic and Secondary Education. MIT Press, Cambridge.
- Kremer, M., Moulin, S., Namunyu, R., 2002. Unbalanced Decentralization. mimeo, Harvard University.
- Kremer, M., Miguel, E., Thornton, R., 2009. Incentives to learn. *Review of Economics and Statistics* 91 (3), 437–456.
- Kreps, D.M., 1997. The interaction between norms and economic incentives intrinsic motivation and extrinsic incentives. *American Economic Review Papers and Proceedings* 87, 359–364.
- Krueger, A.B., 1999. Experimental estimates of education production functions. *Quarterly Journal of Economics* 114, 497–532.
- Krueger, A.B., 2003. Economic considerations and class size. *Economic Journal* 113, F34–F63.
- Krueger, A.B., Zhu, P., 2004. Another look at the New York City school voucher experiment. *American Behavioral Scientist* 47 (5), 658–698.
- Kruglanski, A., 1978. Endogenous attribution and intrinsic motivation. In: Greene, D., Lepper, M.R. (Eds.), *The Hidden Costs of Reward*. Erlbaum Pub., Hillsdale, NJ.
- Ladd, H.F., 1998. Evidence on discrimination in mortgage lending. *Journal of Economic Perspectives* 12, 41–62.
- Lavy, V., 2002. Evaluating the effect of teachers' group performance incentives on pupil achievement. *Journal of Political Economy* 110, 1286–1317.
- Lavy, V., 2009. Performance pay and teacher's effort, productivity, and grading ethics. *American Economic Review* 99, 1979–2011.
- Lavy, V., Paserman, D., Schlosser, A., 2008. Inside the Black Box of Ability Peer Effect: Evidence from Variation of Low Achievers in the Classroom. National Bureau of Economic Research Working Paper No. 14415.
- Lazear, E.P., 1995. *Personnel Economics*. MIT Press, Cambridge, Mass.
- Lazear, E.P., 1989. Pay equality and industrial politics. *Journal of Political Economy* 87, 1261–1284.
- Lazear, E.P., 2000. Performance pay and productivity. *American Economic Review* 90, 1346–1361.
- Lazear, E.P., 2001. Educational production. *Quarterly Journal of Economics* 116, 777–803.
- Lazear, E.P., 2005. Output-based pay: incentives or sorting? In: Polachek, S.W. (Ed.), *Research in Labor Economics – Accounting for Worker Well-Being*, vol. 23. pp. 1–25.

- Lazear, E.P., Rosen, S., 1981. Rank order tournaments as optimum labor contracts. *Journal of Political Economy* 89, 841–864.
- Lazear, E.P., Shaw, K.L., 2007. Personnel economics: the economist's view of human resources. *Journal of Economic Perspectives* 21, 91–114.
- Lazear, E.P., Malmendier, U., Weberz, R.A., 2009. Sorting and Social Preferences. mimeo, Stanford University.
- Ledford, G.E., Lawler, E.E., Mohrman, S.A., 1995. Reward innovations in Fortune 1000 companies. *Compensation and Benefits Review* 27, 76–80.
- Levine, D.K., Pesendorfer, W., 2002. The Evolution of Cooperation Through Imitation. mimeo, UCLA.
- Levitt, S.D., List, J.A., 2007a. What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives* 21, 153–174.
- Levitt, S.D., List, J.A., 2007b. Viewpoint: on the generalizability of lab behavior to the field. *Canadian Journal of Economics* 40, 347–370.
- Levitt, S.D., List, J.A., 2009. Field experiments in economics: the past, the present, and the future. *European Economic Review* 53, 1–18.
- Levitt, S.D., List, J.A., 2010. Was there really a hawthorne effect at the hawthorne plant? an analysis of the original illumination experiments. *American Economic Journal: Applied Economics* (forthcoming).
- Levitt, S.D., List, J.A., Sadoff, S., 2009. Checkmate: exploring backward induction among chess players. *American Economic Review* (forthcoming).
- Levitt, S.D., List, J.A., Sadoff, S., 2010. The Effect of Financial Incentives on High School Achievement: Evidence from Randomized Experiments. Working Paper.
- Ligon, E., 1998. Risk-sharing and information in village economies. *Review of Economic Studies* 65, 847–864.
- List, J.A., 2001. Do explicit warnings eliminate the hypothetical bias in elicitation procedures? evidence from field auctions for sports cards. *American Economic Review* 91 (5), 1498–1507.
- List, J.A., 2002a. Preference reversals of a different kind: the 'more is less' phenomenon. *American Economic Review* 92 (5), 1636–1643.
- List, J.A., 2002b. Testing neoclassical competitive market theory in the field. *Proceedings of the National Academy of Science* 99 (24), 15827–15830.
- List, J.A., 2003a. Using random n th price auctions to value non-market goods and services. *Journal of Regulatory Economics* 23, 193–205.
- List, J.A., 2003b. Does market experience eliminate market anomalies? *Quarterly Journal of Economics* 118, 41–71.
- List, J.A., 2004a. Young, selfish and male: field evidence of social preferences. *Economic Journal* 114, 121–149.
- List, J.A., 2004b. The nature and extent of discrimination in the marketplace: evidence from the field. *Quarterly Journal of Economics* 119, 49–89.
- List, J.A., 2004c. Neoclassical theory versus prospect theory: evidence from the marketplace. *Econometrica* 72, 615–625.
- List, J.A., 2006. Field experiments: a bridge between lab and naturally occurring data. *Advances in Economic Analysis and Policy* 6, Article 8.
- List, J.A., Livingstone, J., 2010. Exploring the Link between Market Power and the Nature and Magnitude of Discrimination. Working Paper, University of Chicago.
- List, J.A., Price, M.K., 2005. Conspiracies and secret price discounts in the marketplace: evidence from field experiments. *Rand Journal of Economics* 36, 700–717.
- List, J.A., Reiley, D.R., 2008. Field experiments. In: Durlauf, Steven N., Blume, Lawrence E. (Eds.), *The New Palgrave Dictionary of Economics*, second ed., Palgrave Macmillan.
- List, J.A., Sadoff, S., Wagner, M., 2010. So You Want to Run an Experiment, Now What? Some Simple Rules of Thumb for Optimal Experimental Design. NBER Working Paper No. 15701.
- List, J.A., Millimet, D.L., Fredriksson, P.G., Warren McHone, W., 2003. Effects of environmental regulations on manufacturing plant births: evidence from a propensity score matching estimator. *Review of Economics and Statistics* 85 (4), 944–952.
- Lizzeri, A., Meyer, M., Persico, N., 2002. The Incentive Effects of Interim Performance Evaluations CARESS. Working Paper 02-09.

- Loewenstein, G., 2005. Hot-cold empathy gaps and medical decision making. *Health Psychology* 24, 49–56.
- Loewenstein, G., Schkade, D., 1999. Wouldn't it be Nice? Predicting Future Feelings. In: Kahneman, D., Diener, E., Schwartz, N. (Eds.), *Well Being: The Foundations of Hedonic Psychology*. Russell Sage Foundation.
- Longnecker, C.O., Sims, H.P., Gioia, D.A., 1987. Behind the mask: the politics of performance appraisal. *The Academy of Management Executive* 1, 183–193.
- Lucking-Reiley, D., 1999. Using field experiments to test equivalence between auction formats: magic on the internet. *American Economic Review* 89, 1063–1080.
- Ludwig, J., Miller, D., 2007. Does head start improve children's life chances? Evidence from a regression discontinuity design. *Quarterly Journal of Economics* 122, 159–208.
- Lundberg, S., Pollak, R., 2003. Efficiency in marriage. *Review of Economics of the Household* 1, 153–167.
- Lusk, J.L., Fox, J.A., 2003. Value elicitation in laboratory and retail environments. *Economics Letters* 79, 27–34.
- Macleod, W.B., Malcomson, J.M., 1989. Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica* 57, 447–480.
- Malcomson, J., 1984. Work incentives, hierarchy, and internal labor markets. *Journal of Political Economy* 92, 486–507.
- Malmendier, U., Tate, G., 2005. CEO overconfidence and corporate investment. *Journal of Finance* 60, 2661–2700.
- Manser, M., 1999. Existing labor market data: current and potential research uses. In: Haltiwanger, J., Manser, M., Topel, R. (Eds.), *Labor Statistics Measurement Issues*. The University of Chicago Press, Chicago.
- Manser, M., Brown, M., 1980. Marriage and household decision-making: a bargaining analysis. *International Economic Review* 21, 31–44.
- Manski, C.F., 1995. *Learning About Social Programs from Experiments with Random Assignment of Treatments*. University of Wisconsin-Madison: Institute for Research on Poverty, Discussion Paper 1061-95.
- Manski, C.F., Garfinkel, I., 1992. Introduction. In: Manski, C.F., Garfinkel, I. (Eds.), *Evaluating Welfare and Training Programs*. Harvard University Press, Cambridge.
- Mas, A., 2006. Pay, reference points, and police performance. *Quarterly Journal of Economics* 121, 783–821.
- Mas, A., Moretti, E., 2009. Peers at work. *American Economic Review* 99, 112–145.
- Maxfield, M., Schirm, S., Rodriguez-Planas, N., 2003. *The Quantum Opportunities Program Demonstration: Implementation and Short-Term Impacts*. Mathematica Policy Research Report 8279–093.
- Mayo, E., 1933. *The Human Problems of an Industrial Civilization*. Macmillan, New York.
- Mazzocco, M., 2004. Saving, risk sharing, and preferences for risk. *American Economic Review* 94, 1169–1182.
- Mazzocco, M., 2007. Household intertemporal behaviour: a collective characterization and a test of commitment. *Review of Economic Studies* 74, 857–895.
- McElroy, M., Horney, M., 1981. Nash-bargained decisions: towards a generalization of the theory of demand. *International Economic Review* 22, 333–349.
- Meyer, H.H., Kay, E., French, J.R., 1965. Split roles in performance appraisal. *Harvard Business Review* 21–29.
- Meyer, B.D., 1995. Natural and quasi-natural experiments in economics. *Journal of Business and Economic Statistics* 13, 151–161.
- Miguel, E., Kremer, M., 2004. Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72, 159–217.
- Milgrom, P.R., 1988. Employment contracts, influence activities, and efficient organization design. *Journal of Political Economy* 96, 42–60.
- Milgrom, P.R., Roberts, J., 1990. The efficiency of equity in organizational decision processes. *American Economic Review Papers and Proceedings* 80, 154–159.
- Moffitt, R.A., 1981. The Negative Income Tax: Would it Discourage Work? *Monthly Labor Review*.
- Moffitt, R.A., 1999. *Econometric Methods for Labor Market Analysis*. In: *Handbook of Labor Economics*, Elsevier (Chapter 24).
- Moldovanu, B., Sela, A., Shi, X., 2007. Contests for status. *Journal of Political Economy* 115, 338–363.

- Moretti, E., 2004. Workers' education, spillovers and productivity: evidence from plant-level production functions. *American Economic Review* 94, 656–690.
- Muralidharan, K., Sundararaman, V., 2007. Teacher Incentives in Developing Countries: Experimental Evidence from India. mimeo, Harvard University.
- Nagin, D., Rebitzer, J.B., Sanders, S., Taylor, L.J., 2002. Monitoring, motivation, and management: the determinants of opportunistic behavior in a field experiment. *American Economic Review* 92, 850–873.
- Nalbantian, H.R., Schotter, A., 1997. Productivity under group incentives: an experimental study. *American Economic Review* 87, 314–341.
- Nalebuff, B.J., Stiglitz, J.E., 1983. Prizes and incentives: toward a general theory of compensation and competition. *Bell Journal of Economics* 14, 21–43.
- Nechyba, T., 2000. Mobility, targeting and private school vouchers. *American Economic Review* 90 (1), 130–146.
- Niederle, M., Vesterlund, L., 2007. Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics* 122, 1067–1101.
- Olken, B.A., 2007. Monitoring corruption: evidence from a field experiment in Indonesia. *Journal of Political Economy* 115, 200–249.
- Orne, M.T., 1962. On the social psychological experiment: with particular reference to demand characteristics and their implications. *American Psychologist* 17, 776–783.
- Orcutt, G.H., Orcutt, A.G., 1968. Incentive and disincentive experimentation for income maintenance policy purposes. *American Economic Review* 58, 754–773.
- Oyer, P., Schaefer, S., 2004. Why do some firms give stock options to all employees? An empirical examination of alternative theories. *Journal of Financial Economics* 76, 99–133.
- Oyer, P., Schaefer, S., 2011. Personnel economics: hiring and incentives. In: *New Developments and Research on Labor Markets*, first ed., In: Ashenfelter, O., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 4B. Elsevier, pp. 1769–1823 (Chapter 20).
- Paarsch, H.J., Shearer, B., 1999. The response of worker effort to piece rates: evidence from the British Columbia tree-planting industry. *Journal of Human Resources* 643–667.
- Paarsch, H.J., Shearer, B., 2000. Piece rates, fixed wages, and incentive effects: statistical evidence from payroll records. *International Economic Review* 41, 59–92.
- Pahl, J.M., 1983. The allocation of money and the structuring of inequality within marriage. *The Sociological Review* 31, 237–262.
- Peterson, P., Howell, W., Wolf, P., Campbell, D., 2003. School vouchers: results from randomized experiments. In: *The Economics of School Choice*. University of Chicago Press, pp. 107–144.
- Pfeffer, J., 1996. *Competitive Advantage Through People: Unleashing The Power of The Work Force*. Harvard Business Press, Cambridge, Mass.
- Phelps, E., 1972. The statistical theory of racism and sexism. *American Economic Review* LXII, 659–661.
- Pigou, A.C., 1920. *The Economics of Welfare*.
- Plug, E., Vijverberg, W., 2003. Schooling, family background, and adoption: is it nature or is it nurture. *Journal of Political Economy* 111, 611–641.
- Podolny, J.M., Baron, J.N., 1997. Resources and relationships: social networks and mobility in the workplace. *American Sociological Review* 62, 673–693.
- Prendergast, C., 1999. The provision of incentives in firms. *Journal of Economic Literature* 37, 7–63.
- Prendergast, C., 2001. Selection and Oversight in the Public Sector with the Los Angeles Police Department as an Example. National Bureau of Economic Research Working Paper No. 8664.
- Prendergast, C., Topel, R.H., 1996. Favoritism in organizations. *Journal of Political Economy* 104, 958–978.
- Puma, M., Burstein, N., Merrell, K., Silverstein, G., 1990. Evaluation of the Food Stamp Employment and Training Program: Final Report, Bethesda, Md. Abt Associates, Bethesda, MD.
- Punch, M., 1985. *The Politics and Ethics of Fieldwork*. Sage, London.
- Raaum, O., Torp, H., 1993. AMO-kurs: Hvem søker, hvem får plass - og hvem får jobbetterpå? Søkelys på arbeidsmarkedet.
- Rabin, M., Schrag, J.L., 1999. First impressions matter: a model of confirmatory bias. *Quarterly Journal of Economics* 114, 37–82.
- Rangel, M.A., 2006. Alimony rights and intrahousehold allocation of resources: evidence from Brazil. *Economic Journal* 116, 627–658.

- Rasul, I., 2008. Household bargaining over fertility: theory and evidence from Malaysia. *Journal of Development Economics* 86, 215–241.
- Reback, R., 2005. School Accountability and the Distribution of Student Achievement. mimeo, Columbia University.
- Rege, M., Telle, K., 2004. The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics* 88, 1625–1644.
- Riach, P.A., Rich, J., 2002. Field experiments of discrimination in the market place. *Economic Journal* 112, F480–F518.
- Riach, P.A., Rich, J., 2006. An experimental investigation of sexual discrimination in hiring in the English labor market. B. E. *Journal of Economic Analysis & Policy* 6, Advances Article 1.
- Rice, B., 1982. The Hawthorne effect: persistence of a flawed theory. *Psychology Today* 16, 71–74.
- Rivkin, S.G., Hanushek, E.A., Kain, J.F., 2005. Teachers, schools, and academic achievement. *Econometrica* 73, 417–459.
- Robinson, J., 2008. Limited Insurance Within the Household: Evidence from a Field Experiment in Western Kenya. mimeom, UC Santa Cruz.
- Rockoff, J., 2009. Field experiments in class size from the early twentieth century. *Journal of Economic Perspectives* 23, 211–230.
- Roethlisberger, F.J., Dickson, W., 1939. *Management and the Worker*. Harvard University Press, Cambridge.
- Rosen, S., 1982. Authority, control, and the distribution of earnings. *Bell Journal of Economics* 13, 311–323.
- Rosen, S., 1986. The theory of equalizing differences. In: Ashenfelter, Orley, Layard, Richard (Eds.), *Handbook of Labor Economics*, vol. 1. North-Holland, Amsterdam.
- Rosenbaum, P., Rubin, D., 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 41–55.
- Rosenzweig, M.R., Wolpin, K.I., 2000. Natural ‘natural experiments’ in economics. *Journal of Economic Literature* 38, 827–874.
- Ross, H.L., 1970. An Experimental Study of the Negative Income Tax, *Child Welfare*, December.
- Rotemberg, J.J., 1994. Human relations in the workplace. *Journal of Political Economy* 102, 684–717.
- Rouse, C., 1998. Private school vouchers and student achievement: an evaluation of the Milwaukee parental choice program. *Quarterly Journal of Economics* 133 (2), 553–602.
- Roy, D., 1952. Quota Restriction and Goldbricking in a Machine Shop. *American Journal of Sociology* 57, 427–442.
- Rozaan, A., Strenger, A., Willinger, M., 2004. Willingness-to-pay for food safety: an experimental investigation of quality certification on bidding behavior. *European Review of Agricultural Economics* 31, 409–425.
- Rubin, D.B., 1990. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Statistical Science* 5, 472–480.
- Samuelson, P.A., 1956. Social indifference curves. *Quarterly Journal of Economics* 70, 1–22.
- Samuelson, W., Zeckhauser, R., 1988. Status quo bias in decision making. *Journal of Risk and Uncertainty* 1, 7–59.
- Sausgruber, R., 2009. A note on peer effects between teams. *Experimental Economics* 12, 193–201.
- Schultz, T.P., 2004. School subsidies for the poor: evaluating the Mexican *Progres*a poverty program. *Journal of Development Economics* 74, 199–250.
- Seabright, P., 2002. Blood, Bribes, and the Crowding-out of Altruism by Financial Incentives. mimeo, Toulouse University.
- Sen, A., 1999. *Development as Freedom*. Knopf, New York.
- Sethi, R., Somanathan, E., 1999. Preference Evolution and Reciprocity. mimeo, University of Michigan.
- Shapiro, C., Stiglitz, J.E., 1984. Equilibrium unemployment as a worker discipline device. *American Economic Review* 74, 433–444.
- Shearer, B.S., 2004. Piece rates, fixed wages and incentives: evidence from a field experiment. *Review of Economic Studies* 71, 513–534.
- Slovic, P., Lichtenstein, S., 1971. Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance* 6, 649–744.
- Snow, C.E., 1927. Research on Industrial Illumination. *The Tech Engineering News* November, 257–282.

- Sobel, J., Takahashi, I., 1983. A multi-stage model of bargaining. *Review of Economic Studies* 50, 411–426.
- Soetevent, A.R., 2005. Anonymity in giving in a natural context – a field experiment in 30 churches. *Journal of Public Economics* 89, 2301–2323.
- Splawa-Neyman, J., 1923a. On the application of probability theory to agricultural experiments. Essay on principles, Section 9. *Statistical Science* 5, 465–472; Translated and edited by D. M. Dabrowska and T. P. Speed from the Polish original, which appeared in *Roczniki Nauk Rolniczych, Tom X (1923): 1-51 (Annals of Agricultural Sciences)*.
- Splawa-Neyman, J., 1923b. Contributions of the theory of small samples drawn from a finite population. *Biometrika* 17, 472–479; The note on this republication reads, These results with others were originally published in *La Revue Mensuelle de Statistique, Publ. Parl'office Central de Statistique de la Republique Polonaise, Tom. vi. pp. 1–29, 1923*.
- Stanley, M., 2003. College education and the mid-century GI bills. *Quarterly Journal of Economics* 118, 671–708.
- Stafford, P., 1986. Forestalling the demise of empirical economics: the role of microdata in labor economics research. In: Ashenfelter, O., Layard, R. (Eds.), *Handbook of Labor Economics*, Elsevier (Chapter 7).
- Stevenson, B., Wolfers, J., 2007. Marriage and divorce: changes and their driving forces. *Journal of Economic Perspectives* 21, 27–52.
- Stiglitz, J.E., 1975. Incentives, risk, and information: notes towards a theory of hierarchy. *Bell Journal of Economics* 6, 552–579.
- Street, D., 1990. Fisher's contributions to agricultural statistics. *Biometrics* 46, 937–945.
- Taber, C., Weinberg, B.A., 2008. Labour economics (new perspectives). In: Durlauf, Steven N., Blume, Lawrence E. (Eds.), *The New Palgrave Dictionary of Economics*, second ed., Palgrave Macmillan.
- Thaler, R., 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1, 39–60.
- Thomas, D.A., 1990. The impact of race on managers' experiences of developmental relationships. *Journal of Organizational Behavior* 11, 479–492.
- Thomas, D.A., 1994. Like father, like son or like mother, like daughter: parental education and child health. *Journal of Human Resources* 29, 950–989.
- Thorndike, E.L., 1913. *Educational Psychology*. Oxford: Columbia University.
- Todd, P., Wolpin, K.I., 2006. Assessing the impact of a school subsidy program in Mexico: using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American Economic Review* 96, 1384–1417.
- Tsui, A.S., O'Reilly, C.A., 1989. Beyond simple demographic effects: the importance of relational demography in superior-subordinate dyads. *Academy of Management Journal* 32, 402–423.
- Udry, C., 1994. Risk and insurance in a rural credit market: an empirical investigation in northern Nigeria. *Review of Economic Studies* 61, 495–526.
- Ulph, D.T., 1988. A General Non-cooperative Nash Model of Household Consumption Behaviour. University of Bristol Working Paper 88/205.
- Vandegrift, D., Yavas, A., Brown, P.M., 2007. Incentive effects and overcrowding in tournaments: an experimental analysis. *Experimental Economics* 10, 345–368.
- Van Den Steen, E., 2004. Rational overoptimism (and other biases). *American Economic Review* 94, 1141–1151.
- Veblen, T., 1934. *The Theory of the Leisure Class: An Economic Study of Institutions*. Modern Library, New York.
- Vollmann, J., Winau, R., 1996. Informed consent in human experimentation before the Nuremberg code. *British Medical Journal* 313, 1445–1449.
- Vollmeyer, R., Rheinberg, F., 2005. A surprising effect of feedback on learning. *Learning and Instruction* 15, 589–602.
- Weber, R.A., 2006. Managing growth to achieve efficient coordination in large groups. *American Economic Review* 96, 114–126.
- Weichselbaumer, D., 2003. Sexual orientation discrimination in hiring. *Labour Economics* 10, 629–642.
- Wesolowski, M.A., Mossholder, K.W., 1997. Relational demography in supervisor-subordinate dyads: impact on subordinate job satisfaction, burnout, and perceived procedural justice. *Journal of Organizational Behavior* 18, 351–362.

- Williams, K.T., O'Reilly, C.A., 1998. Demography and diversity in organizations: a review of 40 years of research. *Research in Organizational Behavior* 20, 77–140.
- White, M., Lahey, J., 1992. Restart effect: does active labour market policy reduce unemployment? Policy Studies Institute.
- Wuchty, S., Jones, B.E., Uzzi, B., 2007. The increasing dominance of teams in production of knowledge. *Science* 316, 1036–1039.
- Yinger, J., 1998. Evidence on discrimination in consumer markets. *Journal of Economic Perspectives* XII, 23–40.
- Zimmerman, D.J., 2003. Peer effects in academic outcomes: evidence from a natural experiment. *Review of Economics and Statistics* 85, 9–23.