

Oxoby, Robert J.

Working Paper

Skill uncertainty and social inference

IZA Discussion Papers, No. 2567

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Oxoby, Robert J. (2007) : Skill uncertainty and social inference, IZA Discussion Papers, No. 2567, Institute for the Study of Labor (IZA), Bonn, <https://nbn-resolving.de/urn:nbn:de:101:1-20080425288>

This Version is available at:

<https://hdl.handle.net/10419/33798>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 2567

Skill Uncertainty and Social Inference

Robert J. Oxoby

January 2007

Skill Uncertainty and Social Inference

Robert J. Oxoby

*University of Calgary
and IZA*

Discussion Paper No. 2567
January 2007

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Skill Uncertainty and Social Inference^{*}

Research in psychology indicates that individuals often make inferences regarding unknown individual qualities based on potentially irrelevant (but socially observable) information. This paper explores occupational choices when individuals receive imprecise signals regarding ability and use the observable characteristics of previously successful individuals to infer own ability. Individuals who fail to observe successful predecessors of their same type may underestimate their potential for success in the occupation. We discuss the role of these biases in light of the literature on affirmative action and firm incentives.

JEL Classification: D63, D83, J64, J70

Keywords: social Inference, occupational choice, discrimination

Corresponding author:

Robert J. Oxoby
Department of Economics
University of Calgary
2500 University Drive NW
Calgary, Alberta T2N 1N4
Canada
E-mail: oxoby@ucalgary.ca

^{*} I would like to thank Ashok Kotwal and Allison Oxoby for helpful comments and suggestions. Financial support was provided by the Institute for Advanced Policy Research and the University of Calgary.

1 Introduction

The economics of asymmetric information typically assumes that (at least) one party has full knowledge with respect to the information in question. Thus, in a principal agent setting, the agent has full knowledge of her ability and the principal must design a compensation mechanism that screens out low ability agents. However, full information regarding one's own ability is not always the case. An agent may have prior beliefs on her abilities as an athlete or academic, but until she undertakes tasks that test these priors, she may be less than certain regarding her skills. As such an individual will not only use her private information to estimate her ability in a task, but may also use public, socially observable information in updating her beliefs regarding own ability.

As an example, consider the study of Steele and Aronson (1995). They find that students from visible minorities under-perform other students when they perceive the task at hand to be one in which their social group is under-represented. However, when the task is not perceived as one reflecting upon the social group, students from visible minorities perform as well as other students. The authors argue that this is due to the students' fear of confirming negative stereotypes and they therefore bias their efforts downward.¹ Similar inferences based on social information are found in the work of Cohen and Zhou (1991), Jemmott and Gonzalez (1989), and Lovaglia et al. (1998).

The key here is that socially observable information (e.g. race or gender) is often inappropriately used to bias individuals' efforts and decisions. These biases (which are not based on individual's ability) may help explain gender differences in higher education in which women account for 60% of college students (Jacob, 2002). Within the economics profession, while a field remains largely male-dominated (Kahn, 1995), Dolado et al. (2005) finds significant gender differences

¹In this environment, confirming a negative stereotype implies working hard yet still failing at the task.

across research fields and, more importantly, that a woman's choice of research field is positively related to the share of women in that field.² This suggests that individuals use information on the observable distribution of "types" (here, gender) in a field or occupation when making inferences about the probability of success in a given area.

These examples fall in line with the extensive literature in psychology documents the ways individuals use information in a manner deviating from the rational choice model (e.g. see Eagly and Chaiken, 1993; Gilovich et al., 2002; Rabin, 1998). On the one hand, experimental evidence indicates that individuals tend to incorporate information in biased ways, attributing positive outcomes to ability and negative outcomes to chance (Eagly and Chaiken, 1993; Langer, 1975). On the other hand, individuals tend to incorporate irrelevant information (via stereotypes or status characteristics) into their subjective beliefs (Berger et al., 1998; Heckhausen, 1991).

This paper develops a model to explore the role of the latter type of bias in individual decision-making. Specifically, this paper presents a model of project and effort choice in which individuals are uncertain regarding their ability. In order to estimate own ability, agents glean information regarding the distribution of skills from the observable characteristics of successful predecessors. This information, which may be irrelevant to the task at hand, is used in tandem with an imperfect private signal to form beliefs regarding one's own ability. This biases individuals' beliefs about own ability, leading individuals to make inefficient choices.

Closely related to the approach taken here is the theory of status characteristics and expectation states. According to Berger et al. (1998), individuals form beliefs based on observable status characteristics (race, gender, etc.) where these

²Similar differences are mirrored in graduate program enrollments. See Bartlett (2000) and Booth et al. (2000).

attributes are assumed correlated with power and prestige. That is, individuals use their relative position in social hierarchies as reference points around which they form subjective beliefs, regardless of whether this information is relevant to the task at hand (Webster and Foschi, 1988). The result is a self-fulfilling prophecy in which low status individuals participate less and exert less effort than their higher status counterparts. This confirms individuals' expectations, making future status-based biases more likely in the future. Aside from having been documented in experimental settings (Jemmott and Gonzalez, 1989; Lovaglia et al., 1998), evidence of biases based on observable type have been found in employment settings. In their analysis of research and development teams, Cohen and Zhou (1991) found that external status characteristics (primarily gender but including race, education, institutional affiliations) formed the basis for individuals' status and levels of interaction within the team. Similar results are documented by Gerber (1998) in police teams where status characteristics affected the attitudes of team members. In their study of promotions, Pergamit and Veum (1999) found that men developed greater job attachment and were more likely to be promoted than women and minorities. These differences in promotions in turn produce feedback effects that emphasize the bias.

As a result of biases based on observable type, one may think of social hierarchies as self-reinforcing: individuals with lower relative positions are more likely to bias their beliefs downward and achieve success less often than higher ranked individuals. The presence of these biases may therefore explain "glass ceiling" phenomena in which women and minorities are relatively absent from esteemed occupations or promoted to higher positions less often. If one considers individuals' beliefs about occupational success to be tied to levels of job attachment or human capital investments, these biases may naturally lead to a stratification of individuals by type and occupation that mimics direct discrimination.

Our analysis is closely related to the literature on statistical discrimination (Arrow, 1973; Coate and Loury, 1993; Lundberg and Startz, 1983; Phelps, 1972). However, in our analysis beliefs about different types are held by agents based on the choices and observable characteristics of previous decision makers. Thus, as opposed to the statistical discrimination literature, the discrimination seen here is on the part of employees rather than employers: statistical discrimination arises through agents incorrectly sorting among competing opportunities. In terms of long-run efficiency, our argument aligns with that of Athey et al. (2000). There, affirmative action (characterized as biased promotion practices) may be efficiency enhancing through a “type-based” mentoring technology: individuals gain more firm specific human capital when their mentor is of the same observable type. Our analysis relies not on the presence of a human capital technology, but rather on the simple biases individuals use to infer information about their abilities.

2 The Model

To explore the role of biases based on observable type (henceforth “*type-based biases*”), we consider an ersatz overlapping generation model in which each generation observes the outcomes of the previous generation.³ We assume a population of agents characterized by two parameters: a level of innate ability $s \in \{A, U\}$ and an observable type $t \in \{M, N\}$. The fraction of type t agents in the economy is given by m_t , with $m_M + m_N = 1$.

While each agent knows her own type t , agents are uncertain as to their innate abilities. For simplicity, we assume that agents view $s = A$ and $s = U$ as equally likely.⁴ Each agent receives an imperfect signal θ regarding her ability,

³There is no direct interaction among the generations, only indirect contact via younger generations observing the outcomes and characteristics of older generations.

⁴This assumption is made for notational simplicity. Qualitatively, the results do not change

with θ distributed over the interval $[0, 1]$ according to $F_s(\theta)$ with density $f_s(\theta)$. We assume $F_A(\theta) < F_U(\theta)$ for all θ and $f_U(\theta)/f_A(\theta)$ to be non-increasing. Thus, higher values of θ imply an agent is more likely to be “able” ($s = A$) than “unable” ($s = U$). In the absence of any additional information or biases, the agent’s (objective) probability of being able given θ is

$$p(\theta) = \frac{f_A(\theta)}{f_A(\theta) + f_U(\theta)}. \quad (1)$$

While agents can costlessly observe others’ types t , they have no direct knowledge as to others’ signals θ or abilities s .

It is assumed that there is no productive difference between type M and type N agents. That is, each type has the same distribution $F_s(\theta)$. However, when assessing private ability, individuals incorporate (irrelevant) type-based information into their subjective beliefs about own ability. Let μ_t be the fraction of type t individuals from the previous generation who are viewed as being successful. (In what follows, this will be the share of individuals of type t opting for a risky production decision.) Following theories on role modelling and social learning (Manski, 1993; Chung, 2000), individuals may be viewed as able due to success in a task known to be correlated with ability. Thus, an individual may look to the observable characteristics of successful predecessors in assessing their own ability in the task. Alternately, following Berger et al. (1998), individuals may infer own ability based on diffuse status characteristics (gender, race, etc.) where these characteristics are assumed to be correlated with success. In line with this research, people who have lower socially ranked characteristics (e.g. women, minorities) underestimate their ability by observing few individuals of their type who are deemed status worthy.

Let $\beta_t \equiv \beta(\mu_t, m_t)$ be the bias of a type t agent upon observing the the pre-

if these two events are not equally likely or differ across agent types.

vious generation. We can think of the type-based bias as a function $\beta(\mu_t, m_t) : (0, 1)^2 \rightarrow (0, 1)$ where $\beta(\mu_t, m_t)$ is increasing in μ_t , decreasing in m_t , and satisfies $\beta(\mu, \mu) = 1/2$. We assume that the bias β_t is incorporated into the agents' subjective belief regarding own skill. That is, given her private signal θ and the fractions μ_t and m_t , a type t agent's updated (subjective) probability that she is able is

$$p_t(\theta) = \frac{\beta_t f_A(\theta)}{\beta_t f_A(\theta) + (1 - \beta_t) f_U(\theta)} \quad (2)$$

where (as opposed to $p(\theta)$ in equation 1) the subscript in $p_t(\theta)$ indicates the agent's belief incorporates the bias β_t . Note that if the distribution of population types (m_M, m_N) mirrors the observed distribution from the previous generation (μ_M, μ_N) , then agents make correct inferences about their ability ($p_t(\theta) = p(\theta)$). Thus, for example, if one were to observe 80% of whites in a given profession among a population which is 80% white, a white individual would not develop a biased inference as $\beta(0.80, 0.80) = 1/2$. However, when μ_t and m_t diverge, agents make incorrect inferences about their own ability given their signals ($p_t(\theta) \neq p(\theta)$). Thus if a white individual observed a profession in which 80% of individuals were white among a population which is only 50% white, she would over-estimate her abilities (here, probability of success in the profession) as $\beta(0.80, 0.50) > 1/2$.

To explore the effect of these biases in decision making, we consider a simple scenario in which an agent receives her signal θ and must make a production decision (e.g. an occupational choice). In particular, the agent must choose whether or not to undertake a risky production opportunity or receive a guaranteed level of utility ν . If the agent chooses to undertake the risky opportunity, she must make an effort decision yielding one of two outcomes $x \in \{x_1, x_0\}$ where $x_1 > x_0$. The probability of an agent of type $s \in \{A, U\}$ generating outcome x_1 when exerting effort e is given by $G_s(e) \in (0, 1)$ where $G_A(e) > G_U(e)$ for

all effort levels $e > 0$ and $G'_A(e) \geq G'_U(e)$. This implies that able agents are more likely to realize x_1 for any given level of effort. The timing of the agents' decision are as follows.

1. Nature chooses each individual's observable type $t \in \{M, N\}$ and ability $s \in \{A, U\}$.
2. Agents receive their private signals of ability θ .
3. Agents observe the previous generation's decisions, subsequently forming their biases β_t and subjective beliefs $p_t(\theta)$.
4. Agents choose between the risky opportunity and ν . If choosing the risky project, agents make effort choices to maximize utility given their subjective beliefs $p_t(\theta)$.
5. The uncertainty is resolved and payoffs are received.

Benchmark case

We begin with the benchmark case in which individuals do not display type-based biases.⁵ To begin, suppose the agent opts for the risky opportunity. Thus, she will choose a level of effort maximizing

$$E[x] - c(e) = H(\theta, e)x_1 + (1 - H(\theta, e))x_0 - c(e) \quad (3)$$

where the cost of effort $c(e)$ is an increasing convex function, and

$$H(\theta, e) \equiv p(\theta)G_A(e) + (1 - p(\theta))G_U(e). \quad (4)$$

⁵For notational simplicity, we drop the subscript t in this section as beliefs are given by equation (1).

The agent's optimal choice of effort, $e^* \equiv e^*(\theta)$ solves

$$[x_1 - x_0] \left(p(\theta) G'_A(e^*) + (1 - p(\theta)) G'_U(e^*) \right) - c'(e^*) = 0. \quad (5)$$

As expected, effort is increasing in the agent's belief about own ability (and hence her private signal).

Given her optimal effort choice based on the values of x_0 , x_1 , and θ , the agent must choose whether or not to undertake the risky opportunity. She will choose the risky opportunity if

$$H(\theta, e^*(\theta))x_1 + (1 - H(\theta, e^*(\theta)))x_0 - c(e^*(\theta)) \geq \nu \quad (6)$$

Let θ^* be the value of the private signal such that equation (6) is satisfied with equality. We can interpret this signal as a *standard* representing the minimum signal of ability necessary to pursue the risky production opportunity. Thus, agents with private signals greater than θ^* should pursue the risky opportunity while those with signals less than θ^* should choose ν . If agents do not display type-based biases, θ^* will serve as a demarcation signal such that agents with signals above (below) θ^* will choose the risky production opportunity (ν).

Type-based biases

We now consider the role of type based-biases on individuals' choices regarding effort and production. Let μ_t be the fraction of individuals from the previous generation who are observed choosing the risky opportunity over that yielding ν .⁶ Upon receiving her signal θ and observing μ_t , each agent chooses a level of

⁶We assume that only choices are observable rather than outcomes. As discussed by Hirszowicz (1982), such an assumption is appropriate in professions where wages are not public information.

effort to maximize

$$H_t(\theta, e)x_1 + (1 - H_t(\theta, e))x_0 - c(e) \quad (7)$$

where, as opposed to the case in the expression (3), agents beliefs about own ability are biased:

$$H_t(\theta, e) \equiv p_t(\theta)G_A(e) + (1 - p_t(\theta))G_U(e). \quad (8)$$

In this scenario, agents' optimal choices of effort $e_t^* \equiv e_t^*(\theta)$ solve

$$[x_1 - x_0] \left(p_t(\theta)G'_A(e_t^*) + (1 - p_t(\theta))G'_U(e_t^*) \right) - c'(e_t^*) = 0. \quad (9)$$

Note that only in the knife-edged case in which $m_t = \mu_t$ will the agents' choices of effort be efficient in the sense of the solutions to equation (5).

We can consider inferences based on observable type as biasing the way individuals interpret their signals θ , resulting in effort choices that fail to maximize expected utility. As effort is increasing in the agent's belief of own ability, it is also increasing in the value of the bias. In particular, assuming $\mu_M > m_M$ (and hence $\beta_M > \frac{1}{2} > \beta_N$), type M agents choose the risky project too often while type N agents choose the risky project too infrequently. As a result, type-based biases are self-fulfilling: over confident type M agents choose the risky project too often, further inflating the prior of type M agents who observe their choices. On the other hand, under confident type N agents choose the risky project too infrequently, biasing downward the beliefs of other type N agents who observe their choices.

With respect to the sorting of individuals between the risky opportunity and that yielding ν with certainty, given her optimal choice of effort $e_t^*(\theta)$ a type t

agent will choose the risky opportunity if

$$H_t(\theta, e_t^*(\theta))x_1 + (1 - H_t(\theta, e_t^*(\theta)))x_0 - c(e_t^*(\theta)) \geq \nu. \quad (10)$$

Let $\tilde{\theta}_t$ be the value of the private signal of a type t agent such that equation (10) holds with equality. This is the perceived standard (as opposed to the objective standard θ^*) for an agent of type t . Notice that given individuals' biases on own ability and the assumption that $\mu_M > m_M$,

$$\tilde{\theta}_M < \theta^* < \tilde{\theta}_N. \quad (11)$$

Thus, agents are inefficiently sorted between the two production options: type M agents with signals in $[\tilde{\theta}_M, \theta^*)$ will inefficiently choose the risky project over ν . Analogously, type N agents underestimate their ability due to the bias and agents with signals in $[\theta^*, \tilde{\theta}_N)$ will opt for ν rather than undertaking the risky project. Indeed, the presence of type-based biases leads agents to behave in a manner that is doubly inefficient: agents make inefficient effort choices and the distribution of agents between the production opportunities is inefficient.

In a long-term context, the presence of type-based biases not only leads to inefficiency in each period, but the inefficiency is magnified over time. By virtue of the fact that more type M agents undertake the risky project and exert greater effort, the differences between μ_t and μ_{-t} grow with each successive generation. As a result, β_M becomes increases and β_N decreases over time. Thus, subsequent generations' effort choices become more inefficient as does the distribution of agents between the two projects.

Restoring efficiency

In order to restore efficiency, consider the actions of a social planner who is aware of the biases β_t and their effects. The social planner could choose a set of type-based policies $(y_t, z_t) \in \mathbb{R}^2$ with policy parameter y_t such that

$$[x_1 + y_t - x_0]H'_t(\theta, e) = [x_1 - x_0]H'(\theta, e), \quad (12)$$

where

$$H'_t(\theta, e) \equiv p_t(\theta)G'_A(e) + (1 - p_t(\theta))G'_U(e) \quad (13)$$

and

$$H'(\theta, e) \equiv p(\theta)G'_A(e) + (1 - p(\theta))G'_U(e). \quad (14)$$

Thus, the policy parameter y_t restores efficiency in the effort choices of the agents. Further, the policy parameter z_t satisfies

$$\begin{aligned} H_t(\theta, e_t^*(\theta))(x_1 + y_t) + (1 - H_t(\theta, e_t^*(\theta)))(x_0 + z_t) = \\ H(\theta^*, e_t^*(\theta^*))x_1 + (1 - H(\theta^*, e_t^*(\theta^*)))x_0. \end{aligned} \quad (15)$$

Thus, z_t restores efficiency in the distribution of agents between the two production options.

There are several ways to interpret the policies (y_t, z_t) . If we consider the production options facing individuals to be labor market choices (say between the skilled and unskilled sectors of the economy), one interpretation of (y_t, z_t) is as affirmative action, where $y_N > y_M$ and $z_N > z_M$.⁷ In essence, efficiency requires the social planner to adjust incentives as a countervailing force against

⁷Holzer and Neumark (2000, 1999) review the literature on affirmative action, focusing on the efficiency and performance aspects of affirmative action. According to their review, affirmative action programs can be designed with an eye towards efficiency and there is little evidence supporting affirmative action as inefficient.

the influence of the biases. Since agents who have few successful predecessors of their same type will be less confident in their abilities, the policies will provide additional incentives to motivate behaviour.

Alternately, note that there exists N agents with signals $\theta \in (\theta^*, \tilde{\theta}_N)$ who opted for ν . From the standpoint of the firm, these agents can be hired for ν although their expected return is greater than ν . As such, a competitive firm can obtain higher profits by hiring type N agents with signals in this interval. Thus competitive pressures in the labor market can serve as a counterbalancing force against the biases. Further, as firms hire type N agents, the mass of these agents obtaining x_1 increases, thereby increasing μ_N relative to μ_M and reducing the bias (i.e. moving β_t towards $\frac{1}{2}$).

3 Conclusion

The economic view of beliefs is that of stark probabilities and the rational use of information. However, much of the research psychology is not in accord with this view. When individuals incorporate irrelevant information into their beliefs, there may be reasons for policies which alter the incentives individuals face. In this way, policies which discriminate based on observable types, even when these distinctions are irrelevant to the task at hand, may be efficiency enhancing.

References

- Arrow, K. J. (1973). Discrimination in labor markets. In O. Ashenfelter and A. Rees (Eds.), *The Theory of Discrimination*, pp. 3–33. Princeton, NJ: Princeton University Press.
- Athey, S., C. Avery, and P. Zemsky (2000). Mentoring and diversity. *American Economic Review* 90(4), 765–786.

- Bartlett, R. L. (2000). Report on the status of women in economics. *American Economic Review, Papers and Proceedings* 90(2), 512–517.
- Berger, J., M. H. Fiesk, R. Z. Norman, and D. G. Wagner (1998). Formation of reward expectations in status situations. In J. Berger and M. Z. Jr. (Eds.), *Status, Power, and Legitimacy*. New Brunswick: Transaction Publishers.
- Booth, A., J. Burton, and K. Mumford (2000). The position of women in u.k. academic economics. *Economic Journal* 110(4), F312–F333.
- Chung, K.-S. (2000). Role models and arguments for affirmative action. *American Economic Review* 90(3), 640–648.
- Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review* 83(5), 1220–1240.
- Cohen, B. and X. Zhou (1991). Status processes in enduring work groups. *American Sociological Review* 56, 179–188.
- Dolado, J. J., F. Felgueroso, and M. Almunia (2005). Do men and women economists choose the saying research fields? Evidence from top-50 departments. IZA working paper number 1859.
- Eagly, A. H. and S. Chaiken (1993). *The Psychology of Attitudes*. Orlando, Florida: Harcourt Brace Jovanovich.
- Gerber, G. (1998). Status in same-gender and mixed-gender police dyads: Effects on personality attributions. *Social Psychology Quarterly* 59, 350–363.
- Gilovich, T., D. Griffin, and D. Kahneman (Eds.) (2002). *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press.
- Heckhausen, H. (1991). *Motivation and Action*. Berlin: Springer-Verlag.

- Hirszowicz, M. (1982). *Industrial Sociology*. New York: St. Martin's Press.
- Holzer, H. and D. Neumark (1999). Assessing affirmative action. NBER Working Papers number 7323.
- Holzer, H. and D. Neumark (2000). Assessing affirmative action. *Journal of Economic Literature* 38(3), 483–568.
- Jacob, B. A. (2002). Where the boys aren't: non-cognitive skills, returns to school and the gender gap in higher education. NBER working paper number 8946.
- Jemmott, J. and E. Gonzalez (1989). Social status, the status distribution, and performance in small groups. *Journal of Applied Social Psychology* 19, 584–598.
- Kahn, S. (1995). Gender differences in academic career paths of economics. *American Economic Review* 83, 52–56.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology* 32(2), 311–328.
- Lovaglia, M. J., J. W. Lucas, J. A. Houser, S. R. Thye, and B. Markovsky (1998). Status processes and mental ability test scores. *American Journal of Sociology* 104(1), 195–228.
- Lundberg, S. J. and R. Startz (1983). Private discrimination and social intervention in competitive labor markets. *American Economic Review* 73(3), 340–347.
- Manski, C. F. (1993). Dynamic choice in social settings. *Journal of Econometrics* 58, 121–136.

- Pergamit, M. R. and J. R. Veum (1999). What is a promotion? *Industrial and Labor Relations Review* 52(4), 581–601.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review* 62, 659–661.
- Rabin, M. (1998). Psychology and economics. *Journal of Economic Literature* 36(1), 11–46.
- Steele, C. M. and J. Aronson (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology* 69, 797–811.
- Webster, M. and M. Foschi (1988). *Status Generalizations*. Stanford, CA: Stanford University Press.