Attention and Perception

Lecture 2: A Robust Test of Prejudice

Daniel Martin
danielmartin@ucsb.edu

Recall

Caplin & Martin (2015), "A Testable Theory of Imperfect Perception."

- Question: Can we assume that a decision maker is **updating beliefs correctly** (using Bayes' Rule) and **maximizing utility** for unobservable private information?
  - Key assumptions behind many, if not most, models of perception
  - Includes: drift-diffusion model, race model, signal detection theory

Recall

Caplin & Martin (2015), "A Testable Theory of Imperfect Perception."

- Question: Can we assume that a decision maker is **updating beliefs correctly** (using Bayes' Rule) and **maximizing utility** for unobservable private information?
  - Key assumptions behind many, if not most, models of perception
  - Includes: drift-diffusion model, race model, signal detection theory
- Use choice data (observable) to understand private information (unobservable)
  - Choice data are summarized by $P(a, \omega)$, the joint distribution of actions and states
  - Call this **"state-dependent stochastic choice data" (SDSC)**

Definition

Utility function u satisfies the No Improving Action Switches (**NIAS**) condition with respect to P if,

$$\sum_{\omega \in \Omega} P(a, \omega) u(x(a, \omega)) \geq \sum_{\omega \in \Omega} P(a, \omega) u(x(b, \omega)),$$

for all $a, b \in A$.

Intuition: Cannot improve utility by making a *wholesale* switch from any chosen action

Theorem

P has a Bayesian Expected Utility (BEU) representation if and only if there exists $u : X \to \mathbb{R}$ satisfying the NIAS inequalities.

Moving to Application

- Many possible applications of NIAS
- Key steps:
  1. Determine the data set $P(a, \omega)$ for that application
  2. Then connect back to theory (what does BEU imply in this application?)
- For example:
  - $P(\textit{bundle}, \textit{price})$ NIAS = consumers with price uncertainty
  - $P(\textit{report}, \textit{evidence})$ NIAS = beliefs with bad memory
  - $P(\textit{hire}/\textit{not}, \textit{productivity})$ NIAS across groups = statistical discrimination with subjective perception

Statistical Discrimination with Subjective Perception

- Hiring is rife with subjective perception: hiring managers partially read CVs, partially read cover letters, etc.
- **Intuitive question:** Could it be that differences in choices across groups is due to statistical discrimination with subjective perception?
- Statistical discrimination with subjective perception = information structure $\pi(\gamma|\omega)$ changes with group
- **Technical version:** Can changes in $\pi(\gamma|\omega)$ across groups explain changes in the data $P(a, \omega)$ across groups?

Decision-Problem Dependence

- New idea: The DM might face several different "decision problems" $i \in D$
- Many things can change with the decision problem
  - Actions A_i
  - Prizes $x_i(a, \omega)$
  - Utility $u_i(a, \omega)$
  - Objective prior $P_i(\omega)$
  - Subjective prior μ_i
  - Information structure $\pi_i(\gamma|\omega)$
  - Decision rule σ_i
- And so might get different SDSC $P_i(a, \omega)$ in each decision problem $i \in D$

Decision-Problem Dependence

- New idea: The DM might face several different "decision problems" $i \in D$
- Many things can change with the decision problem
  - Actions A_i
  - Prizes $x_i(a, \omega)$
  - Utility $u_i(a, \omega)$
  - Objective prior $P_i(\omega)$
  - Subjective prior μ_i
  - Information structure $\pi_i(\gamma|\omega)$
  - Decision rule σ_i
- And so might get different SDSC $P_i(a, \omega)$ in each decision problem $i \in D$
- In this application, i indicates the protected category (gender, race, etc.) of the person that the decision is being made about

Martin & Marx (2022), "A Robust Test of Prejudice."

- Two main explanations for discrimination are prejudice (utility-based) and statistical (belief-based)
- Question: Can we identify when someone is prejudiced?
- Challenge: Don't know the decision maker's utility or their perception!
- Idea: NIAS provides the limits of what can be explained by statistical discrimination, so if violated, this provides strong evidence of prejudice
- Example setting: Reuben et al. (2014)
  - One randomly selected subject is the "employer" and must decide whether to "hire" subjects of different genders
  - Information provided in each treatment: appearance, past performance, and/or candidate predictions for future performance

Statistical Discrimination with Subjective Perception

- Actions: hire or not hire, States: H or L (higher and lower productivity)
- Data: SDSC for both groups (P_w and P_m)
- For members of group w, NIAS requires that for action hire:

$$P_w(\text{hire}, \text{H})u(\text{hire}, \text{H}) + P_w(\text{hire}, \text{L})u(\text{hire}, \text{L}) \geq$$
$$P_w(\text{hire}, \text{H})u(\text{not hire}, \text{H}) + P_w(\text{hire}, \text{L})u(\text{not hire}, \text{L})$$

- Or dividing off the action probability:

$$P_w(\text{H}|\text{hire})u(\text{hire}, \text{H}) + (1 - P_w(\text{H}|\text{hire}))u(\text{hire}, \text{L}) \geq$$
$$P_w(\text{H}|\text{hire})u(\text{not hire}, \text{H}) + (1 - P_w(\text{H}|\text{hire}))u(\text{not hire}, \text{L})$$

Statistical Discrimination with Subjective Perception

- Likewise, for action not hire:

$$P_w(\text{not hire}, \text{H})u(\text{not hire}, \text{H}) + P_w(\text{not hire}, \text{L})u(\text{not hire}, \text{L}) \geq$$
$$P_w(\text{not hire}, \text{H})u(\text{hire}, \text{H}) + P_w(\text{not hire}, \text{L})u(\text{hire}, \text{L})$$

- And dividing off the action probability:

$$(1 - P_w(\text{L}|\text{not hire}))u(\text{not hire}, \text{H}) + P_w(\text{L}|\text{not hire})u(\text{not hire}, \text{L}) \geq$$
$$(1 - P_w(\text{L}|\text{not hire}))u(\text{hire}, \text{H}) + P_w(\text{L}|\text{not hire})u(\text{hire}, \text{L})$$

Utility Bounds

- These NIAS inequalities place bounds on the utility function u
- Normalizing the utility of mistakes in each state to 0 ($u(\text{hire}, \text{L}) = 0$ and $u(\text{not hire}, \text{H}) = 0$), it must be that for w:

$$\frac{P_w(\text{L}|\text{not hire})}{1 - P_w(\text{L}|\text{not hire})} \geq \frac{u(\text{hire}, \text{H})}{u(\text{not hire}, \text{L})} \geq \frac{(1 - P_w(\text{H}|\text{hire}))}{P_w(\text{H}|\text{hire})}$$

- The same bounds hold for m:

$$\frac{P_m(\text{L}|\text{not hire})}{1 - P_m(\text{L}|\text{not hire})} \geq \frac{u(\text{hire}, \text{H})}{u(\text{not hire}, \text{L})} \geq \frac{(1 - P_m(\text{H}|\text{hire}))}{P_m(\text{H}|\text{hire})}$$

A Robust Test of Prejudice

- This holds if and only if

$$P_m(\text{H}|\text{hire}) \geq P_w(\text{H}|\text{not hire})$$

  and the analogous condition with w and m swapped

- Thus, choice data are not consistent with statistical discrimination if:

$$P_w(\text{H}|\text{not hire}) > P_m(\text{H}|\text{hire})$$

  and the analogous condition with w and m swapped

- A simple and robust outcome test of prejudice!

Lab experiment of Reuben et al. (2014)

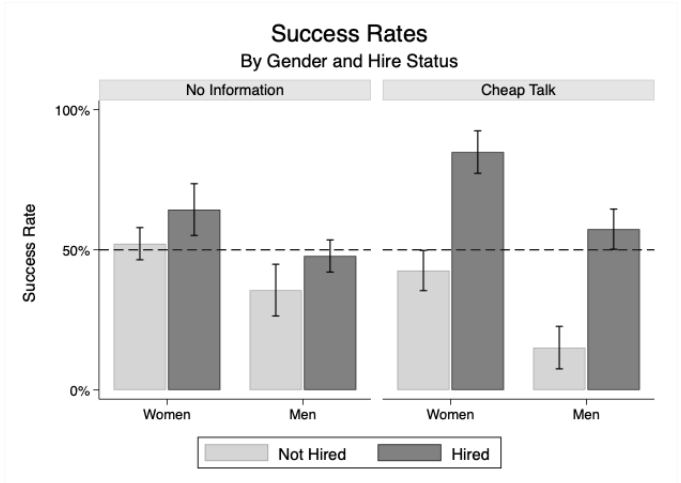

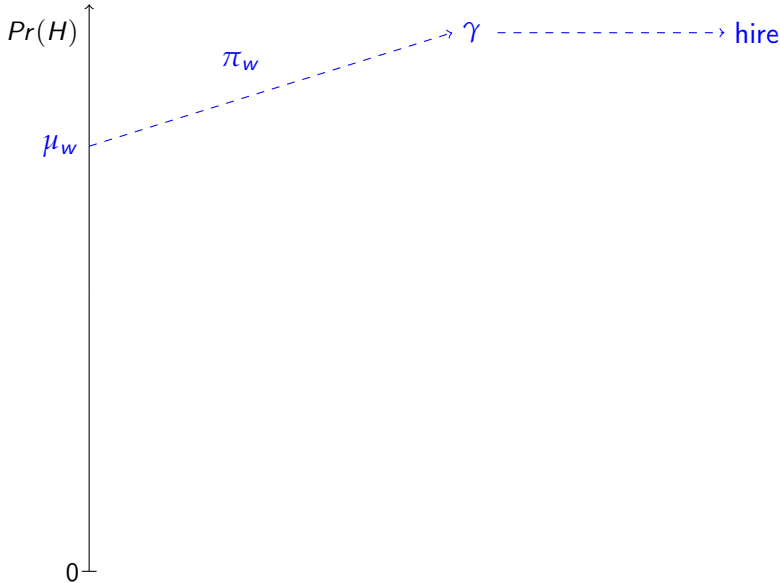
$Pr(H)$
π_w
γ
hire
μ_w
0

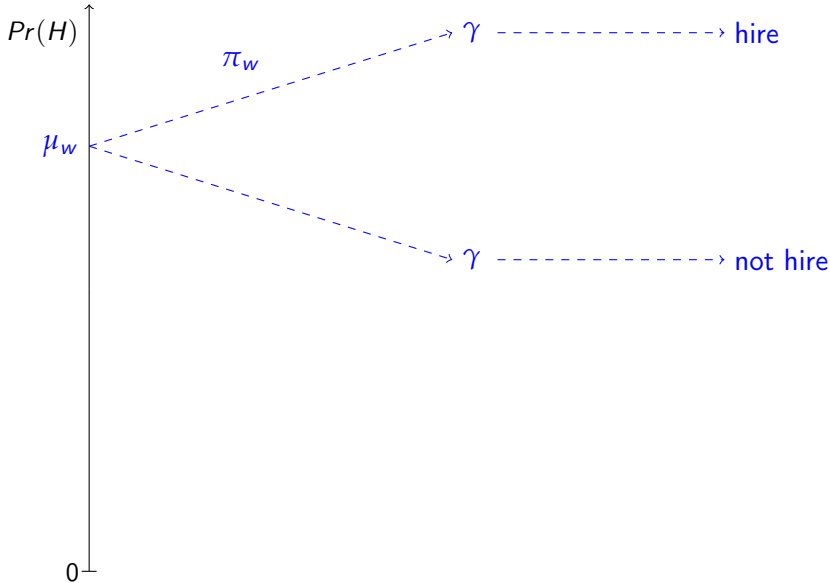
$Pr(H)$
π_w
γ
hire
μ_w
γ
not hire
0

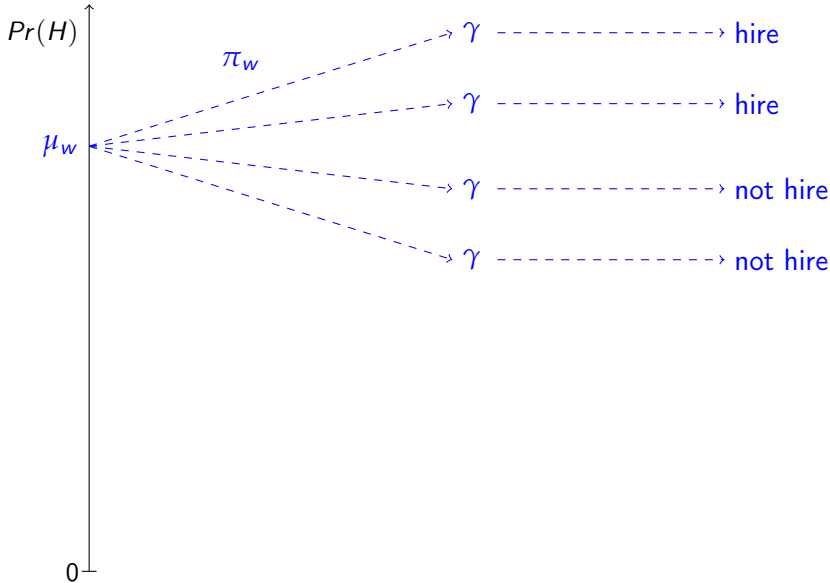
$Pr(H)$
π_w
μ_w
γ
hire
γ
hire
γ
not hire
γ
not hire
0

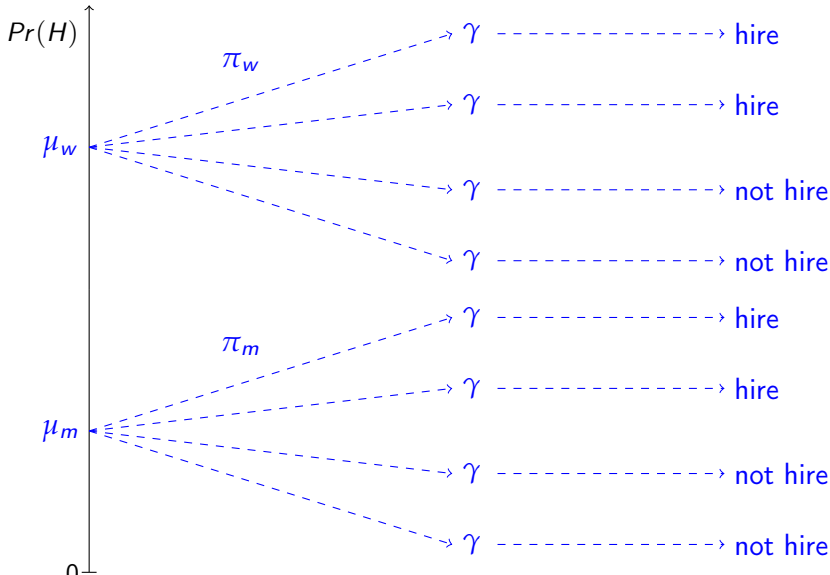
$Pr(H)$
π_w
μ_w
γ
hire
γ
hire
γ
not hire
γ
not hire
π_m
μ_m
γ
hire
γ
hire
γ
not hire
γ
not hire
0

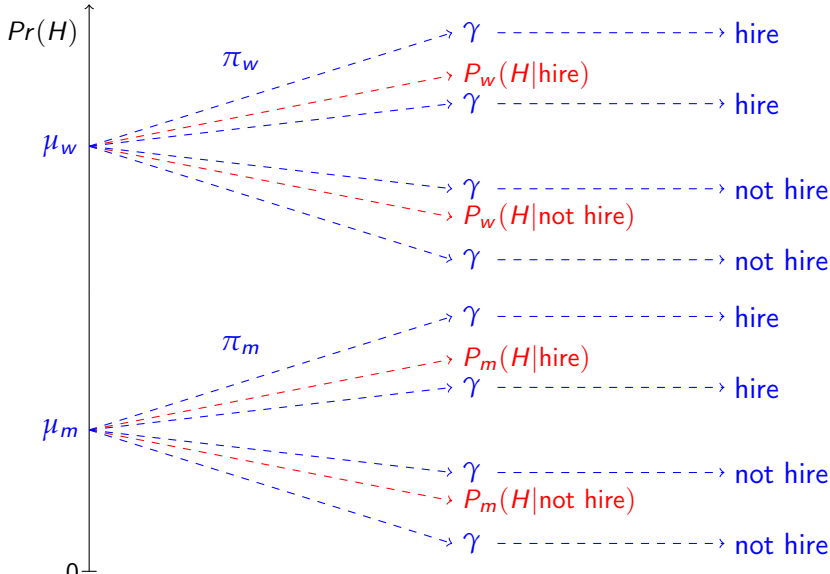
$Pr(H)$
π_w
μ_w
γ
hire
$P_w(H|\text{hire})$
γ
hire
γ
not hire
$P_w(H|\text{not hire})$
γ
not hire
π_m
μ_m
γ
hire
$P_m(H|\text{hire})$
γ
hire
γ
not hire
$P_m(H|\text{not hire})$
γ
not hire
0

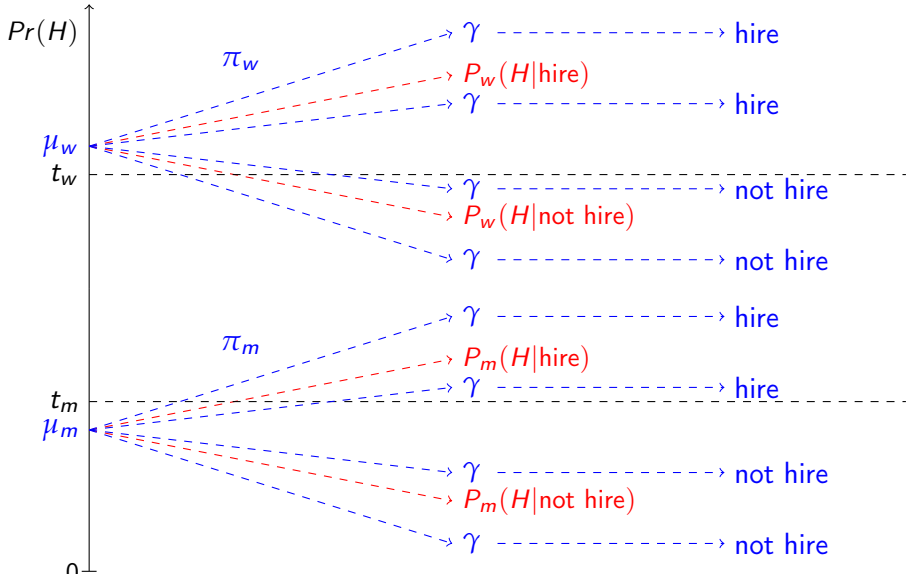
$Pr(H)$
π_w
γ
hire
$P_w(H|\text{hire})$
γ
hire
μ_w
t_w
γ
not hire
$P_w(H|\text{not hire})$
γ
not hire
γ
hire
π_m
$P_m(H|\text{hire})$
γ
hire
t_m
μ_m
γ
not hire
$P_m(H|\text{not hire})$
γ
not hire
0

Selection Motives

- Another possibility is prejudice in the selection motive
  - When the decision-maker positively selects for a trait in one group and negatively selects for the same trait in another group
- Demonstrate using the correspondence field experiment of Bertrand & Mullainathan (2004)
  - Names that strongly signal gender and race were randomly added to fictitious resumes of subjectively high and low quality
  - They observed a strong disparity in job callbacks depending on the race of the name that was added to a resume

Our Test

- Proposition 2 shows there is prejudice in the selection motive against group m if

$$P_m(H|\text{call back}) < P_m(H|\text{not call back})$$

- Equivalently, if high quality candidates are called back less often than low quality candidates in that group:

$$P_m(\text{call back}|H) < P_m(\text{call back}|L)$$

- Why might this happen? Many possible reasons!
- To the best of our knowledge, this form of prejudice is new to the discrimination literature

Bertrand and Mullainathan (2004)

- Proposition 3 shows our test remains true as long as observed outcomes correlate sufficiently with true outcomes
  - Allows us to leverage exogenous variation in resume quality
- Based on variation in resume quality, Martin & Marx (2022) find robust evidence of prejudice in selection motive at the intersection of gender and race
- Show the probability of a callback for an African-American male *decreases* with resume quality: from 7.4% for low-quality resumes to 4.3% for high-quality resumes

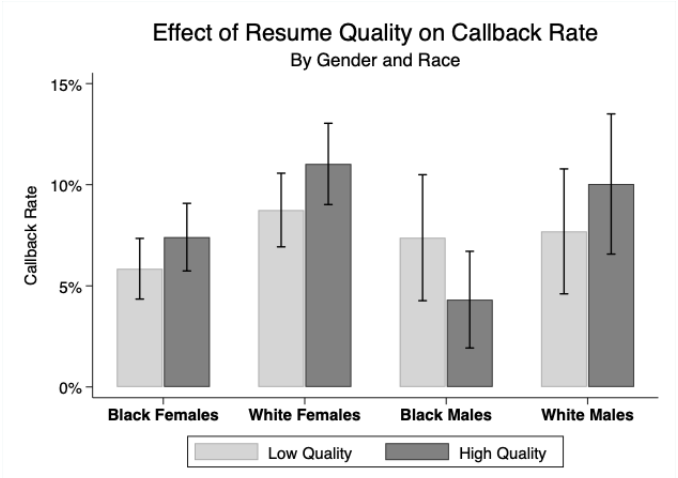

Figure 1: Similarly to their Table 5, standard errors are corrected for clustering at the employment-ad level in a probit regression.

Bertrand and Mullainathan (2004)

- Same pattern for all job types (manager, sales representative, retail)
- Martin & Marx (2022) appear to be the first to document this pattern
- This pattern implies that:
  1. Employers are paying attention to resume details
  2. Employers are prejudiced in their selection motive regardless of what information they gleaned from the resumes they received

Contributions

1. Martin & Marx (2022) provide a test of prejudice that is general and simple
   - Demonstrate with two well-known experiments: Reuben et al. (2014) and Bertrand & Mullainathan (2004)
2. It is very robust to selective attention and unobservable learning
   - Makes no assumptions about how informativeness varies across unobservables or how informativeness varies across decision-makers
   - We relax assumptions of correct subjective prior in the paper (important for recent discrimination literature)
3. Can be applied to average outcomes
   - Becker (1957) showed that a decision-maker is prejudiced if there are differences in outcomes across groups *for marginal decisions*
   - Limitation: Requires knowing which decisions are marginal

References I

Becker, G. S. (1957), *The Economics of Discrimination*, University of Chicago Press, Chicago.

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**(4), 991–1013.

Caplin, A. & Martin, D. (2015), 'A testable theory of imperfect perception', *The Economic Journal* **125**(582), 184–202.

Martin, D. & Marx, P. (2022), 'A robust test of prejudice for discrimination experiments', *Management Science* **68**(6), 4527–4536.

Reuben, E., Sapienza, P. & Zingales, L. (2014), 'How stereotypes impair women's careers in science', *Proceedings of the National Academy of Sciences of the United States of America* **111**(12), 4403–4408.