MPhil Thesis Writing Workshop

Carlos Gonzalez

University of Oxford - Department of Economics

May 7th, 2025

Problem Description

- Consider a set of N **workers** indexed by i characterized by a **reservation wage** r_i and a **productivity** p_i coming from $F_{R,P}$
- **Firm** makes a wage offer w_i to worker i without knowing r_i, p_i, nor $F_{R,P}$ Worker i accepts the offer iff $w_i \geq r_i$
- Observability
  - Always observe $\mathbb{1}(w_i \geq r_i)$, but never r_i fully
  - Conditional on acceptance, firm observes p_i, otherwise it observes nothing
- Firm's goal is to maximize a mix of profits and workers' welfare

$$\sum_{i}^{N} \mathbb{1}(w_i \geq r_i)(p_i - w_i + \lambda(w_i - r_i))$$

Why This Problem? Theory

- When $F_{R,P}$ is known, Akerlof characterized the best reply wage $w^*(F_{R,P})$. This equilibrium wage was shocking at the time as it could feature **market unraveling** dynamics

- This equilibrium showed the effects of **adverse selection** in markets with **quality differentiation** and **imperfect information**

- But what if $F_{R,P}$ is unknown? Is **repeated interaction** enough to converge to Akerlof's equilibrium results? Is there market unraveling?
  - Do we need to recover the entire $F_{R,P}$ to learn the best wage $w^*(F_{R,P})$?
  - Can we learn w^* under limited feedback? $\mathbb{1}(r_i \leq x_i)$ and $p_i \cdot \mathbb{1}(r_i \leq x_i)$
  - Can we learn $w^*(F_{R,P})$ for any $F_{R,P}$ or do we need any restrictions?
  - Do we need any priors on $F_{R,P}$?

Why This Problem? Empirics

- Adverse selection is a powerful mechanism behind the missing markets hypothesis. But does it hold in dynamic settings?
- This (stylized) setting describes the problem of a **wage-setting monopsony** firm, for which there is an increasing interest in Labor Economics
- Counterfactual **policy evaluation**
  - Minimum wage policies
  - Limited information processing capacity
  - Productivity shocks

Results

- I derive a new algorithm which converges to $w^*(F_{R,P})$ without necessarily characterizing $F_{R,P}$
  - Arbitrary (adversarial) distributions
  - Limited feedback
  - No prior requirements
- Algorithm converges at the optimal rate $\mathcal{O}(N^{2/3})$
- (Bandit) learning theory results
  - Extend Cesa-Bianchi, Colomboni and Kasy (ECMA, 2025)
  - Introduction of asymmetric feedback structures
  - Implications of greedy parameter selection
- New insights in policy design

Tips I

- No matter how good you are (you are all VERY good btw!!), your first idea will be VERY bad
- Start early so you can get those bad ideas out of the way and you can start working on less bad ideas
- Fabulous learning opportunity
- Take risks and be creative, but start from a well-established benchmark
- Get out of the “I’m too dumb for my supervisor” loop

Tips II

- Write, write, write (and share, discuss, present!). Slides with a narrative can be a good alternative
- Get things done. The perfect is enemy of the good
- Writing and exposition is key. You should be spending a big share of your time actually writing and polishing your thesis
- Use the summer to read. Create a reading list of papers you are passionate about. It's fine to only understand 20-25% of the paper

Resources

- My thesis here **presidente-carlos.github.io**
- Contact me (Carlos Gonzalez Perez) via email. Happy to talk about anything but I have a better taste for Micro Theory, Machine Learning (and some Econometrics)
- You'll be fine. We were all fine. We were no smarter than you are. It's hard, it's tough, you are harder, you are tougher