When Samples Are Strategically Selected

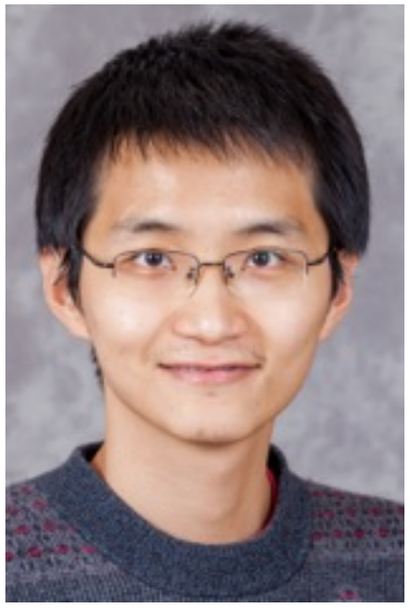

Hanrui Zhang Yu Cheng Vincent Conitzer

Duke University

Academia in 20 years…

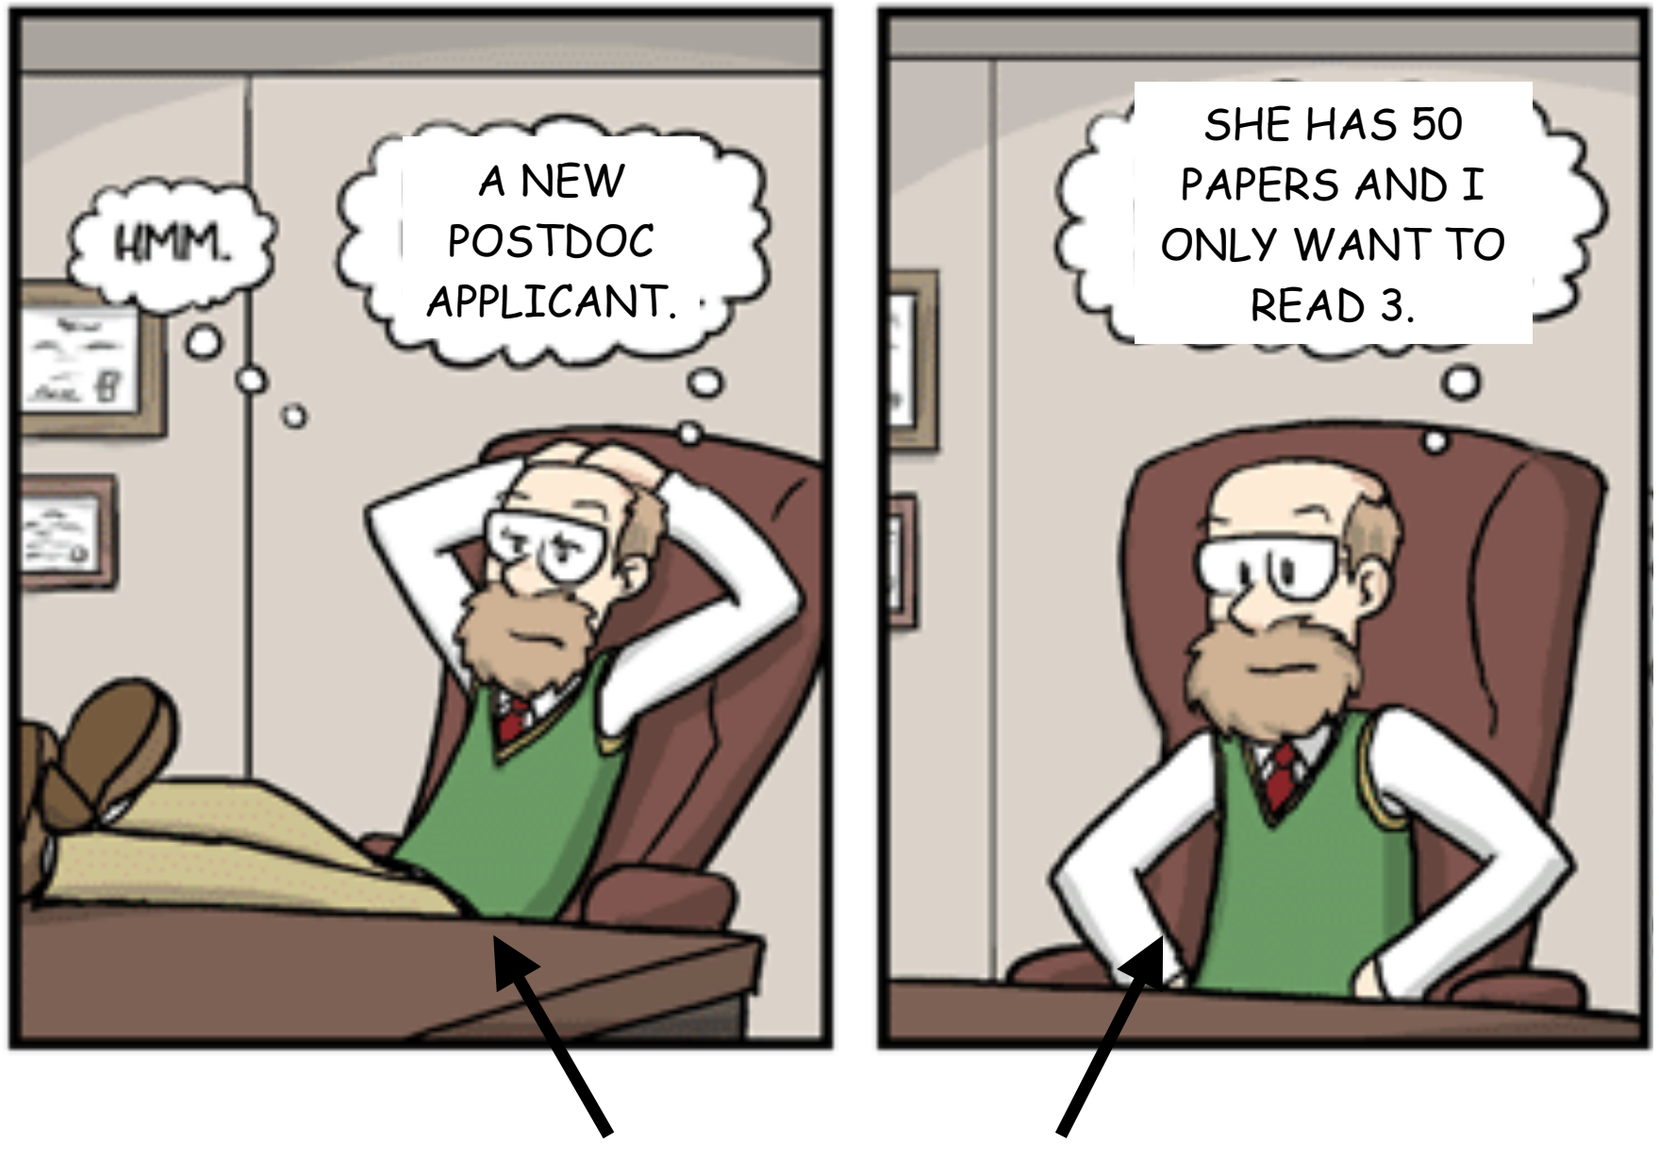

Academia in 20 years…

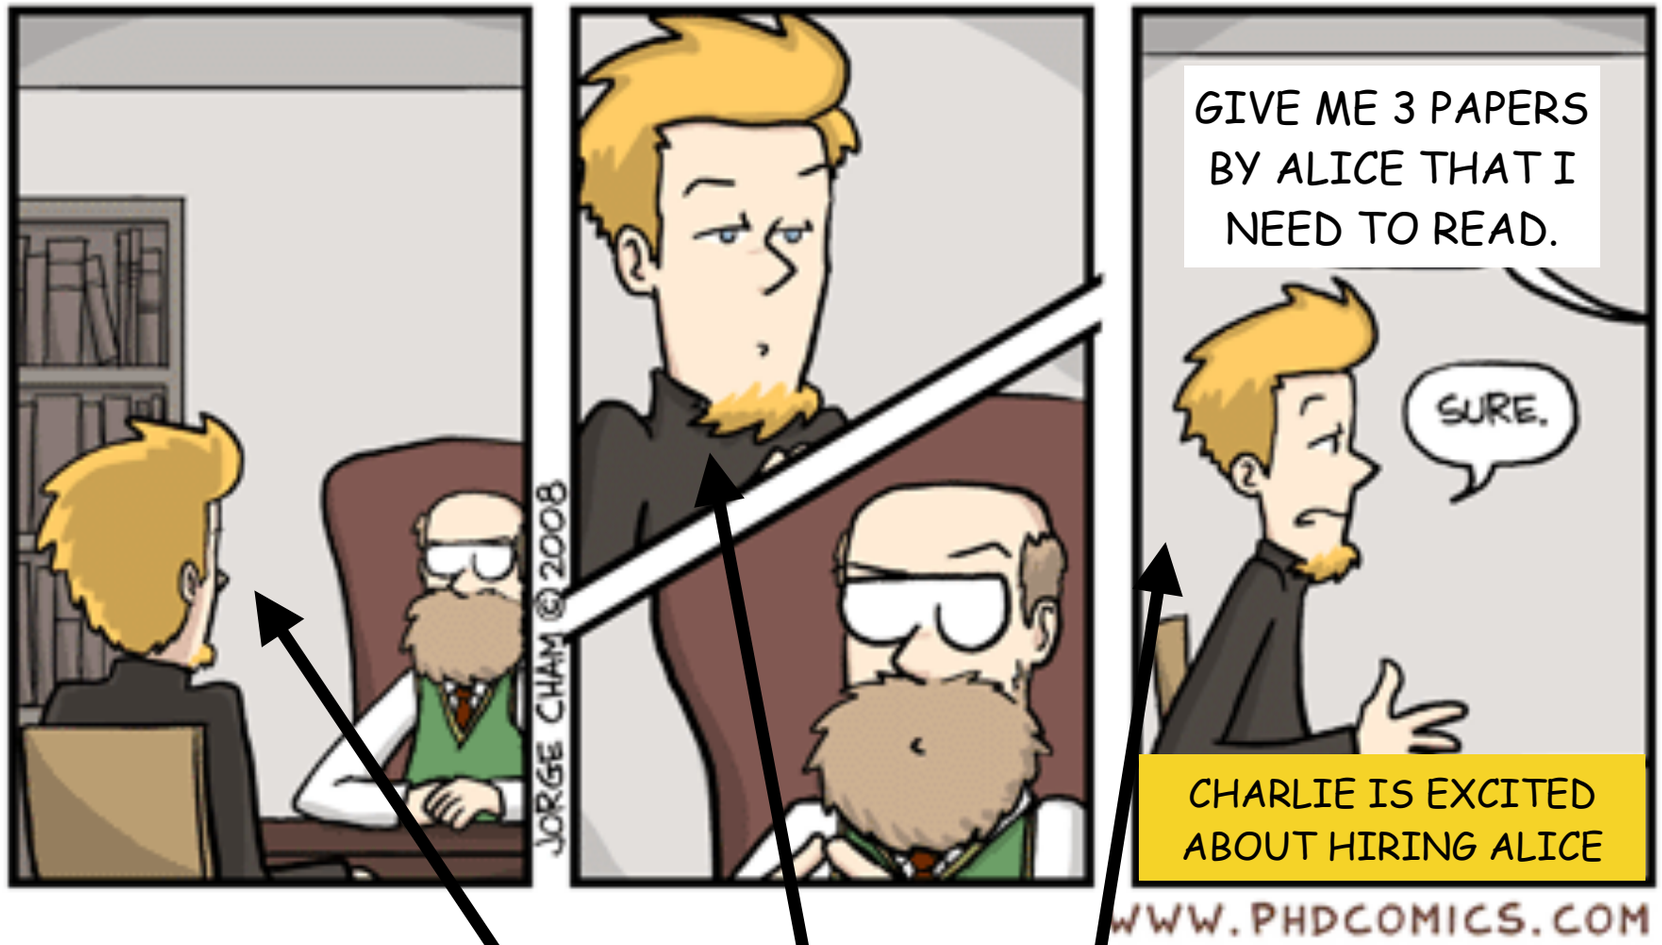

Charlie, Bob's student

Academia in 20 years…

The general problem

A **distribution (Alice)** over paper qualities $\theta \in \{g, b\}$ arrives, which can be either a good one ($\theta = g$) or a bad one ($\theta = b$)

Alice, the postdoc applicant

The general problem

The **principal (Bob)** announces a **policy**, according to which he decides, based on the **report** of the **agent (Charlie)**, whether to **accept θ (hire Alice)**

The general problem

The **agent (Charlie)** has access to **n(=50) iid samples (papers)** from θ (Alice), from which he **chooses m(=3)** as his **report**

The general problem

The **agent (Charlie)** sends his **report** to the principal, **aiming to convince the principal (Bob) to accept** θ **(Alice)**

The general problem

The **principal (Bob)** observes the **report** of the **agent (Charlie)**, and makes the decision according to the policy announced

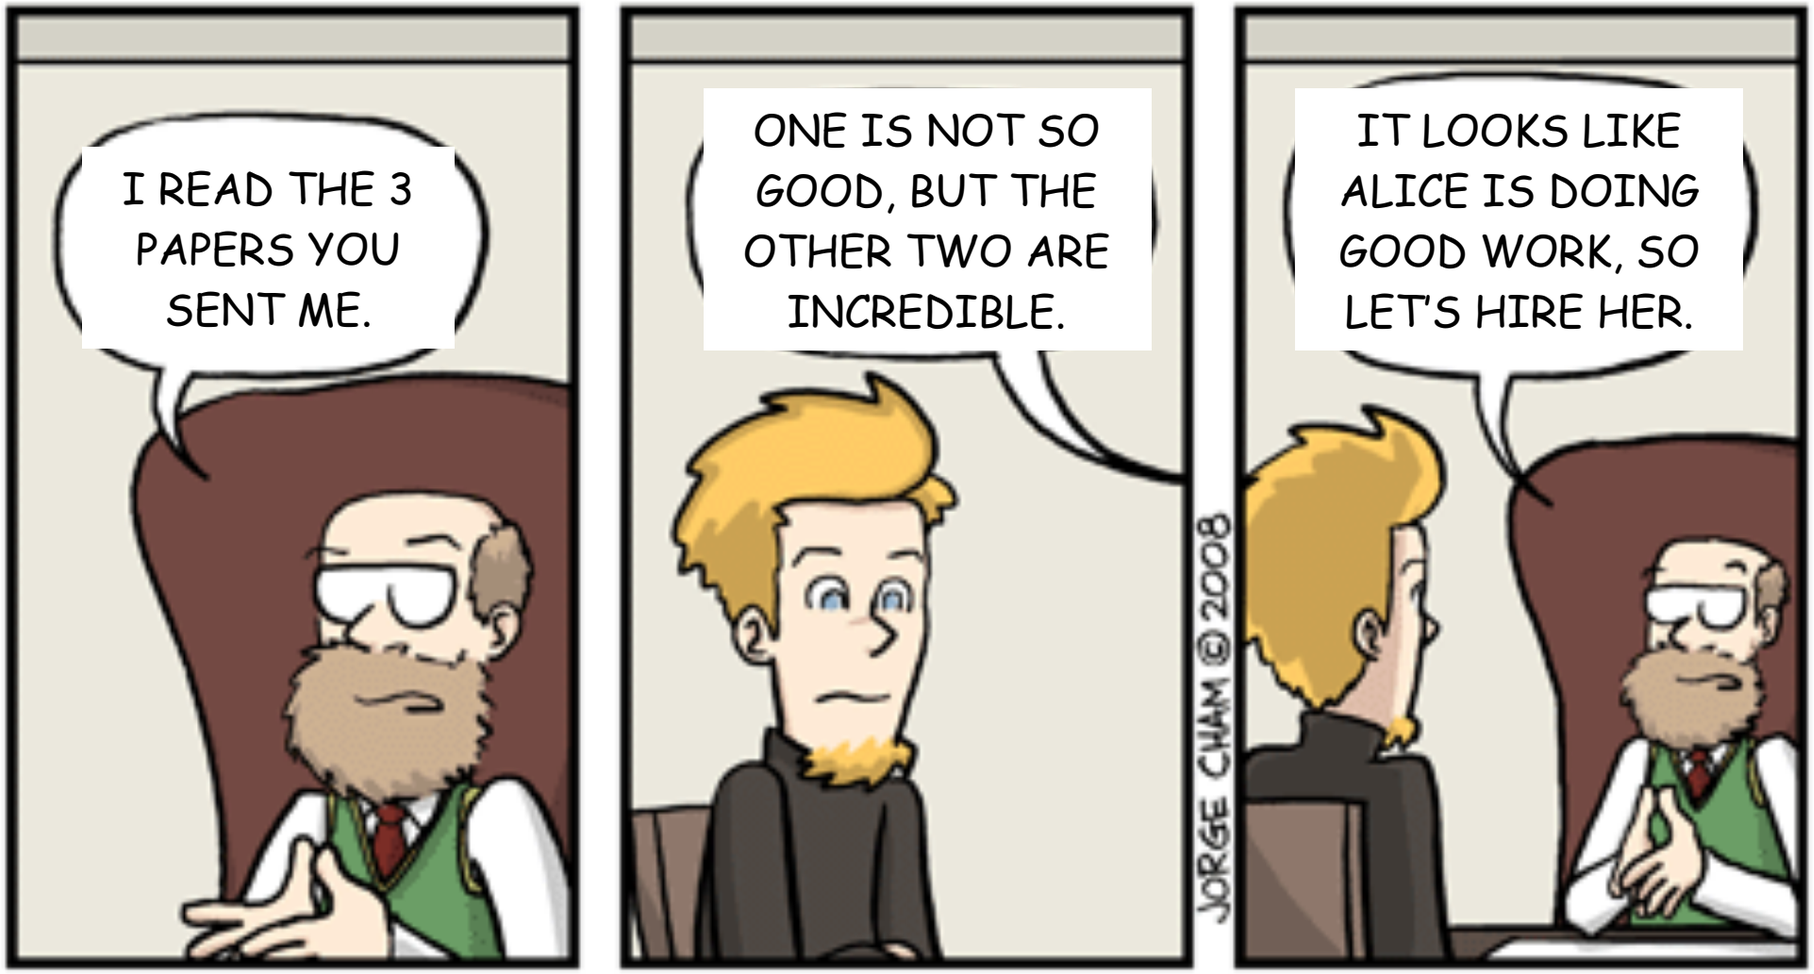

Questions

- How does strategic selection affect the principal's policy?

- Is it easier or harder to classify based on **strategic samples**, compared to when the principal has access to **iid samples**?

- Should the principal ever have a **diversity** requirement (e.g., at least 1 mathematical paper and at least 1 experimental paper), or only go by total quality?

A “hard” world

- A good candidate writes a good paper w.p. 0.05
- A bad candidate writes a good paper w.p. 0.005
- All candidates have n = 50 papers, and the professor wants to read only m = 1

A “hard” world

- A good candidate writes a good paper w.p. 0.05
- A bad candidate writes a good paper w.p. 0.005
- All candidates have n = 50 papers, and the professor wants to read only m = 1
- A reasonable policy: **accept** iff the reported paper is **good**
- A good candidate is accepted w.p. $1 - (1 - 0.05)^{50} \approx 0.92$
- A bad candidate is accepted w.p. $1 - (1 - 0.005)^{50} \approx 0.22$

A “hard” world

- A good candidate is accepted w.p. $1 - (1 - 0.05)^{50} \approx 0.92$
- A bad candidate is accepted w.p. $1 - (1 - 0.005)^{50} \approx 0.22$

Strategic selection helps the principal!

An “easy” world

- A good candidate writes a good paper w.p. ~~0.05~~ 0.95
- A bad candidate writes a good paper w.p. ~~0.005~~ 0.05
- All candidates have n = 50 papers, and the professor wants to read only m = 1

An “easy” world

- A good candidate writes a good paper w.p. ~~0.05~~ 0.95
- A bad candidate writes a good paper w.p. ~~0.005~~ 0.05
- All candidates have n = 50 papers, and the professor wants to read only m = 1
- A reasonable policy: **accept** iff the reported paper is **good**
- A good candidate is accepted w.p. $1 - (1 - 0.95)^{50} \approx 1$
- A bad candidate is accepted w.p. $1 - (1 - 0.05)^{50} \approx 0.92$

An “easy” world

- A good candidate is accepted w.p. $1 - (1 - 0.95)^{50} \approx 1$
- A bad candidate is accepted w.p. $1 - (1 - 0.05)^{50} \approx 0.92$

Now strategic selection hurts the principal!

More questions

- What does the optimal policy look like?

- What parameter(s) determine its performance?

And answers…

Come to our poster!