Recommendations in High-Stake Settings

Nikhil Garg

Cornell Tech

ngarg@cornell.edu

Announcements

- In person lecture on Wednesday! (Weather permitting)

https://networksmarketscornelltech.github.io/Spring2026/1to100/

Matching ←→ Recommenders

THE MATCH
NATIONAL RESIDENT MATCHING PROGRAM®

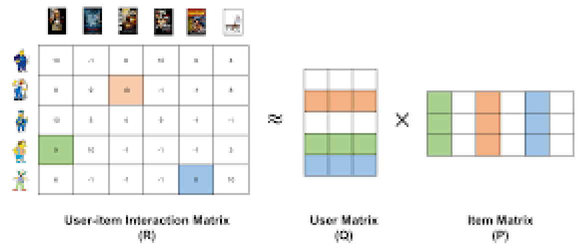

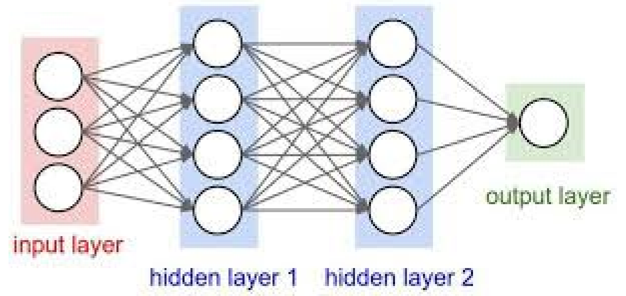

Matching systems from Econ-CS

Residency matching, HS matching

- Offline, ~once a year
- Every participant sees + can rank entire relevant other side
- Central Algorithm
  - Given rankings, decides who goes where
  - Stability (equilibrium) desideratum

Recommendation systems from ML

Netflix, Spotify, news, …

- Real-time
- Millions of participants; can't possibly give full rankings for other side
- Recommender algorithm
  - Estimates user preferences with data
  - ~ naïve suggestions given estimates

Companies have deployed recommenders into platforms for jobs, dating, housing, healthcare
→ How do we incorporate market design ideas into designing recommendation systems?

Market designers have launched stable matching-like systems where humans can't fully rank
← What is the role of noise and ML recommendations in historically "offline" matching systems?

Recommenders: high level agenda

Recommenders are used in online platforms for hiring, dating, healthcare, research dissemination

These settings have desiderata that go beyond preference prediction

- Multi-sided fairness and capacity constraints
- “Mutual” preferences
- Handling strategic behavior for both production and consumption
- Set recommendation and diversity

Opportunity: Tools from market design and economics

Challenge: Take seriously *uncertainty* and *approximation* of machine learning preference prediction

Monoculture in matching markets

Algorithmic monoculture (Kleinberg & Raghavan 2021)

What happens when firms use the same algorithm for decisions?

Applicant has a value v

Firms rank according to v + noise

Monoculture: shared across firms
(e.g., common test scores/algorithm)

Polyculture: Independent across firms
(e.g., independent interviews)

Monoculture in Matching Markets. Kenny Peng and **NG.**
Wisdom and Foolishness of Noisy Matching Markets. Kenny Peng and **NG**

Monoculture in matching markets

The answer from existing literature: monoculture if unequivocally bad

- Firms *can* make worse decisions (compared to independent, “worse” algorithms) [Kleinberg & Raghavan]
- Worse for applicants (increases “systematic exclusion”) [Creel & Hellman; Bommasani et al; Toups et al; Jain et al]

However, this literature ignores two-sided preferences and doesn’t have many participants!

Our work: Many firms and many applicants, incorporates applicant preferences

- Fully strengthen KR result: with many firms, “wisdom of the crowds” (when noise is well behaved)
- Monoculture improves *overall* applicant welfare. Individual applicants’ preferences vary
- Monoculture more robust to disparities in number of applications

Theoretical tool: Azevedo Leshno continuum model of matching markets

Monoculture in Matching Markets. Kenny Peng and **NG.**
Wisdom and Foolishness of Noisy Matching Markets. Kenny Peng and **NG**

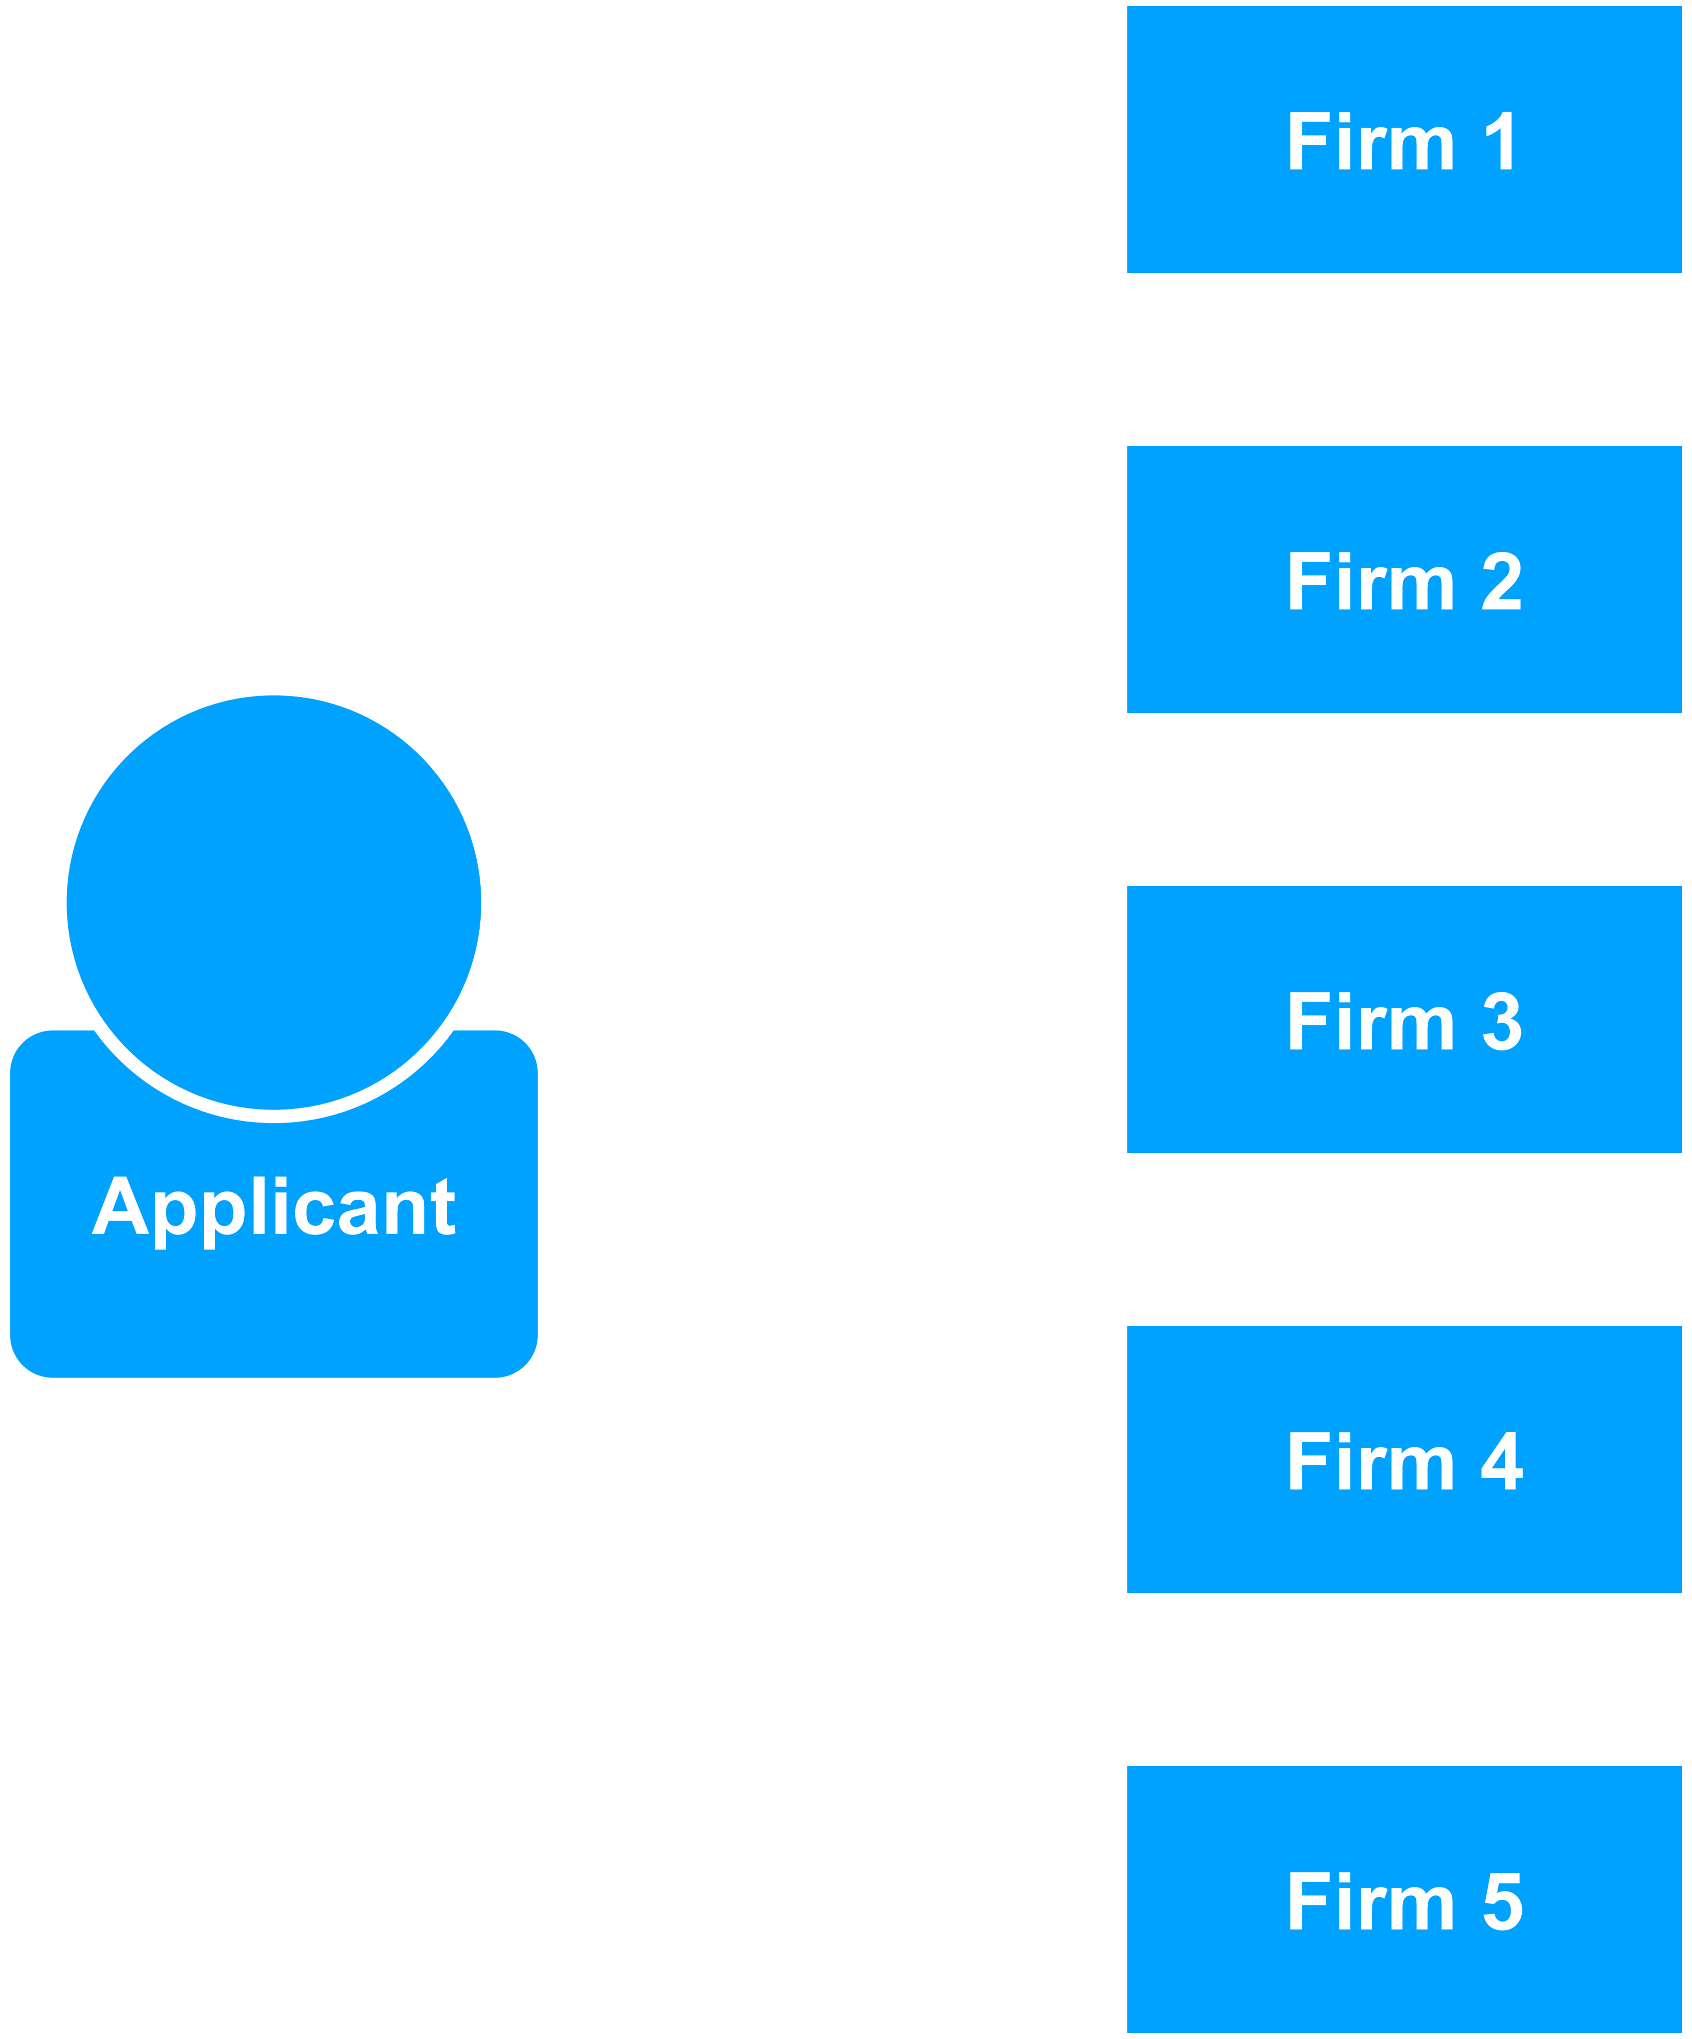
Firm 1
Firm 2
Applicant
Firm 3
Firm 4
Firm 5

Polyculture

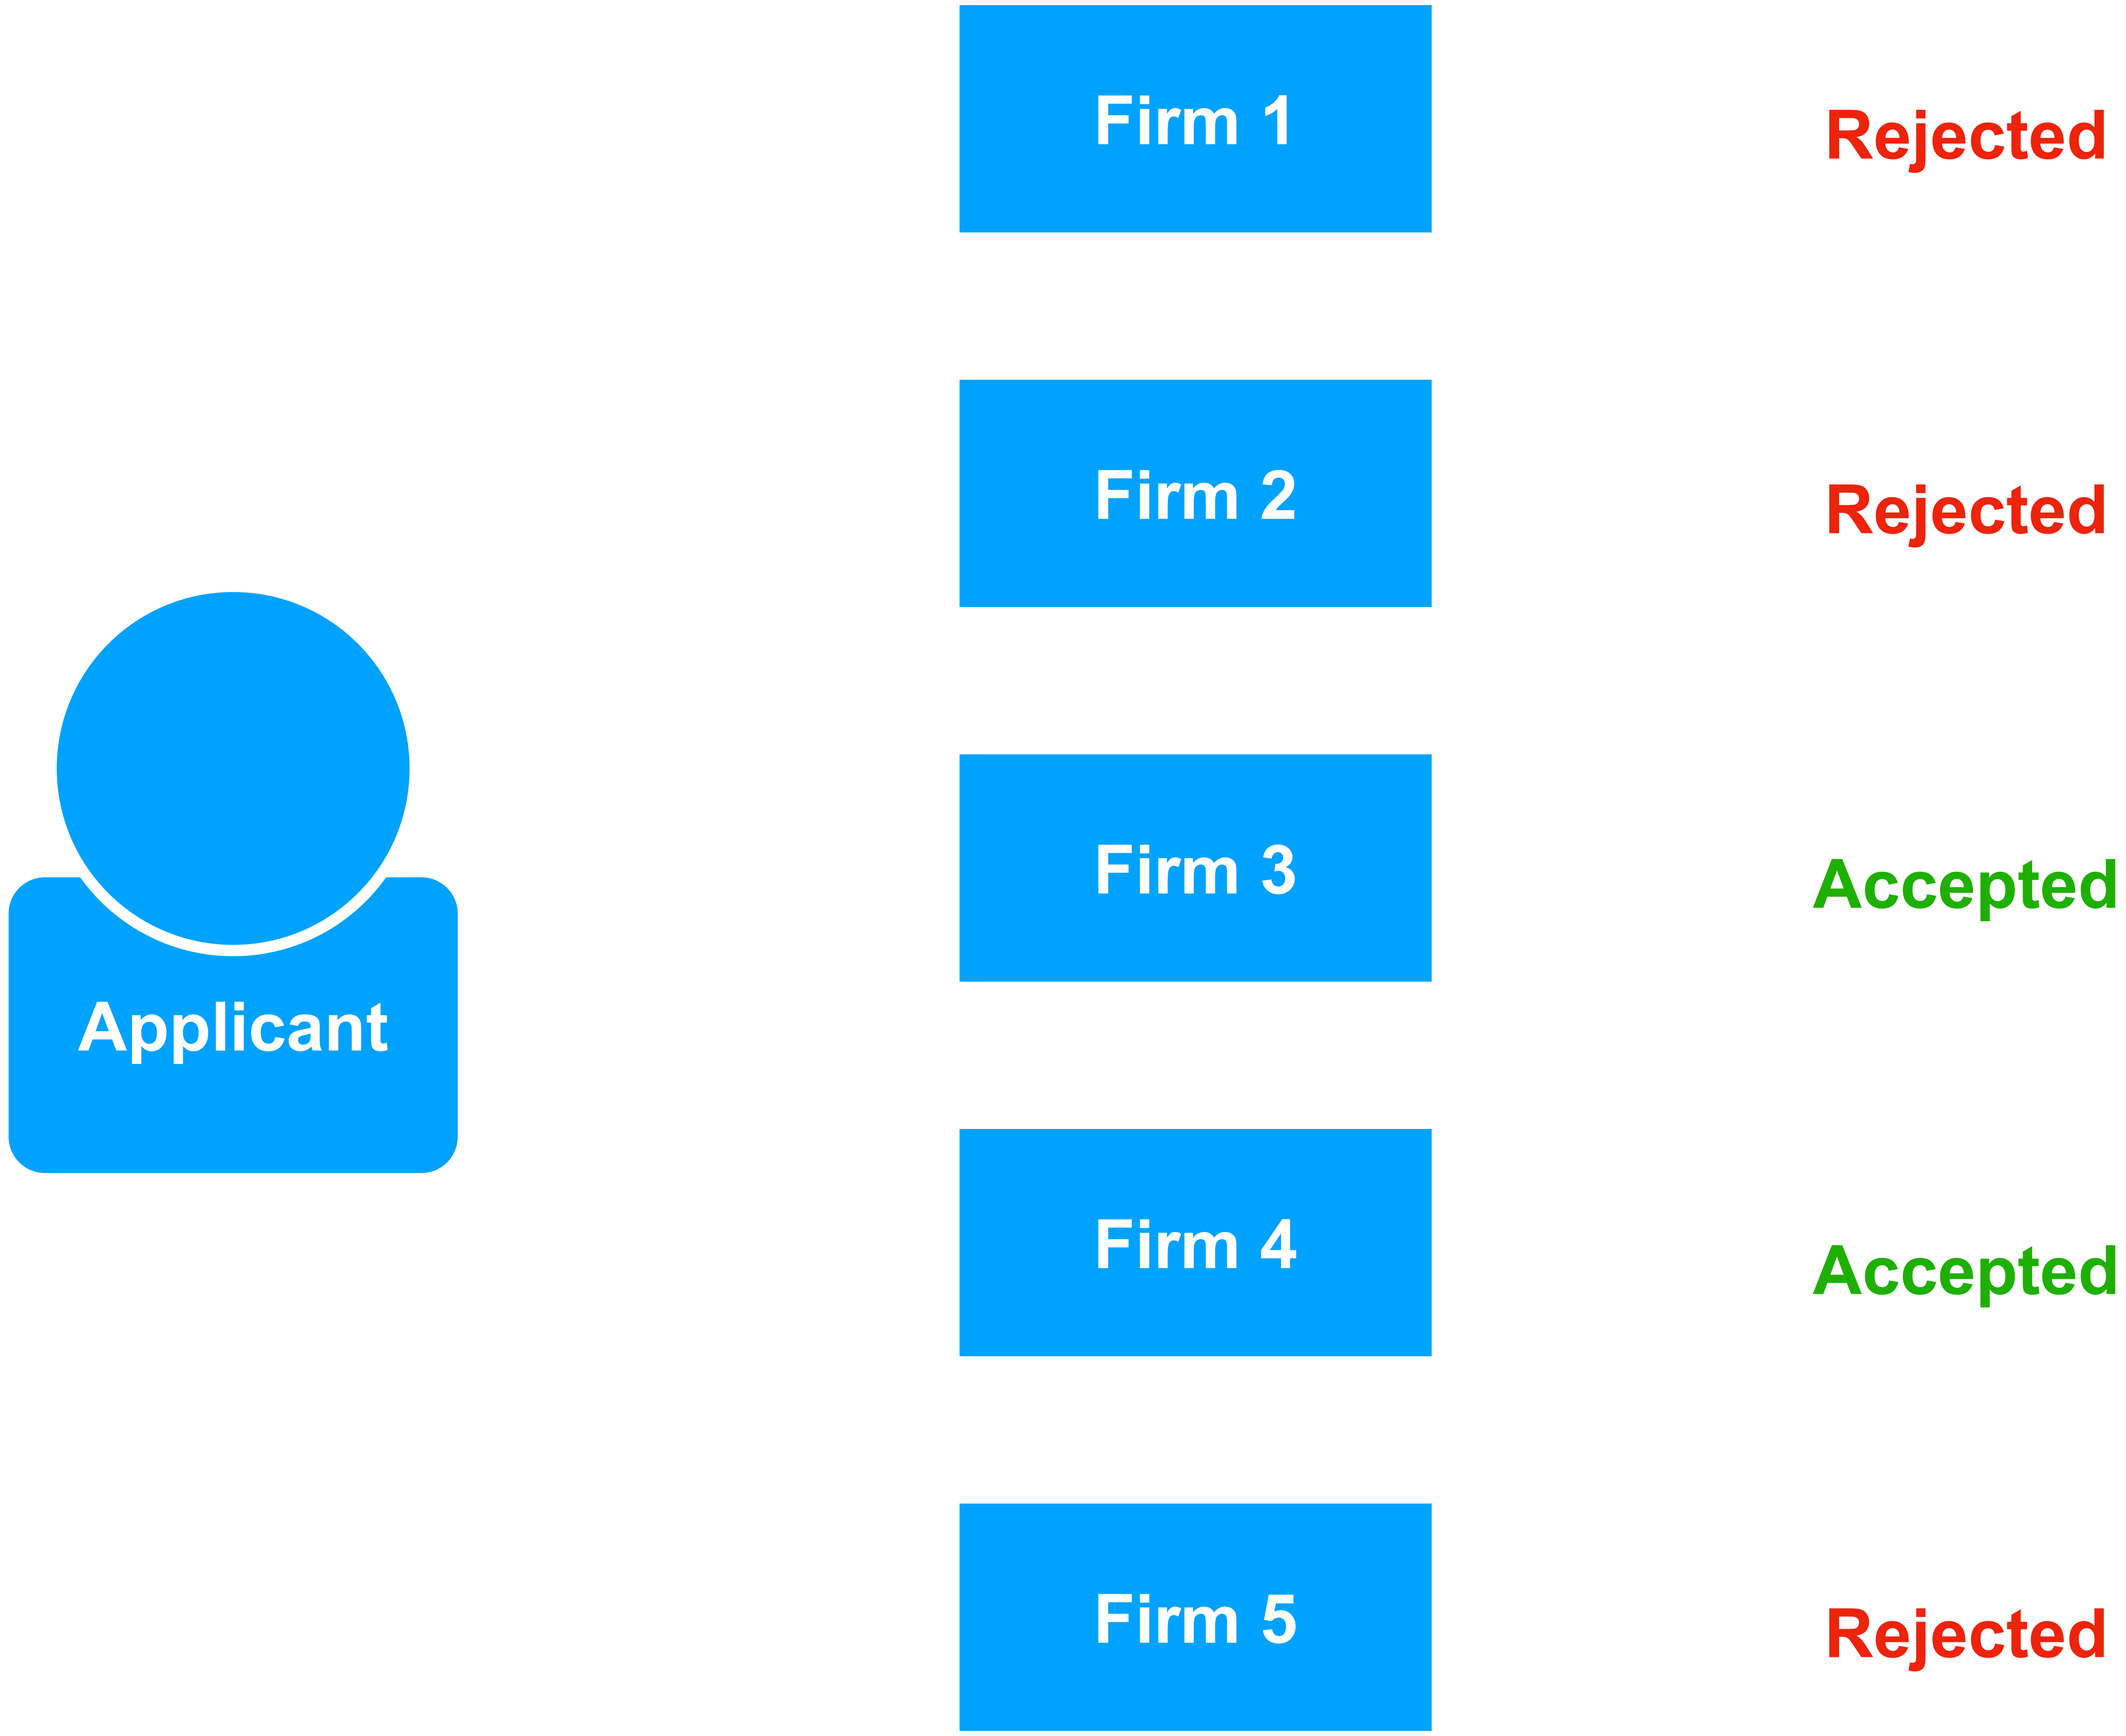

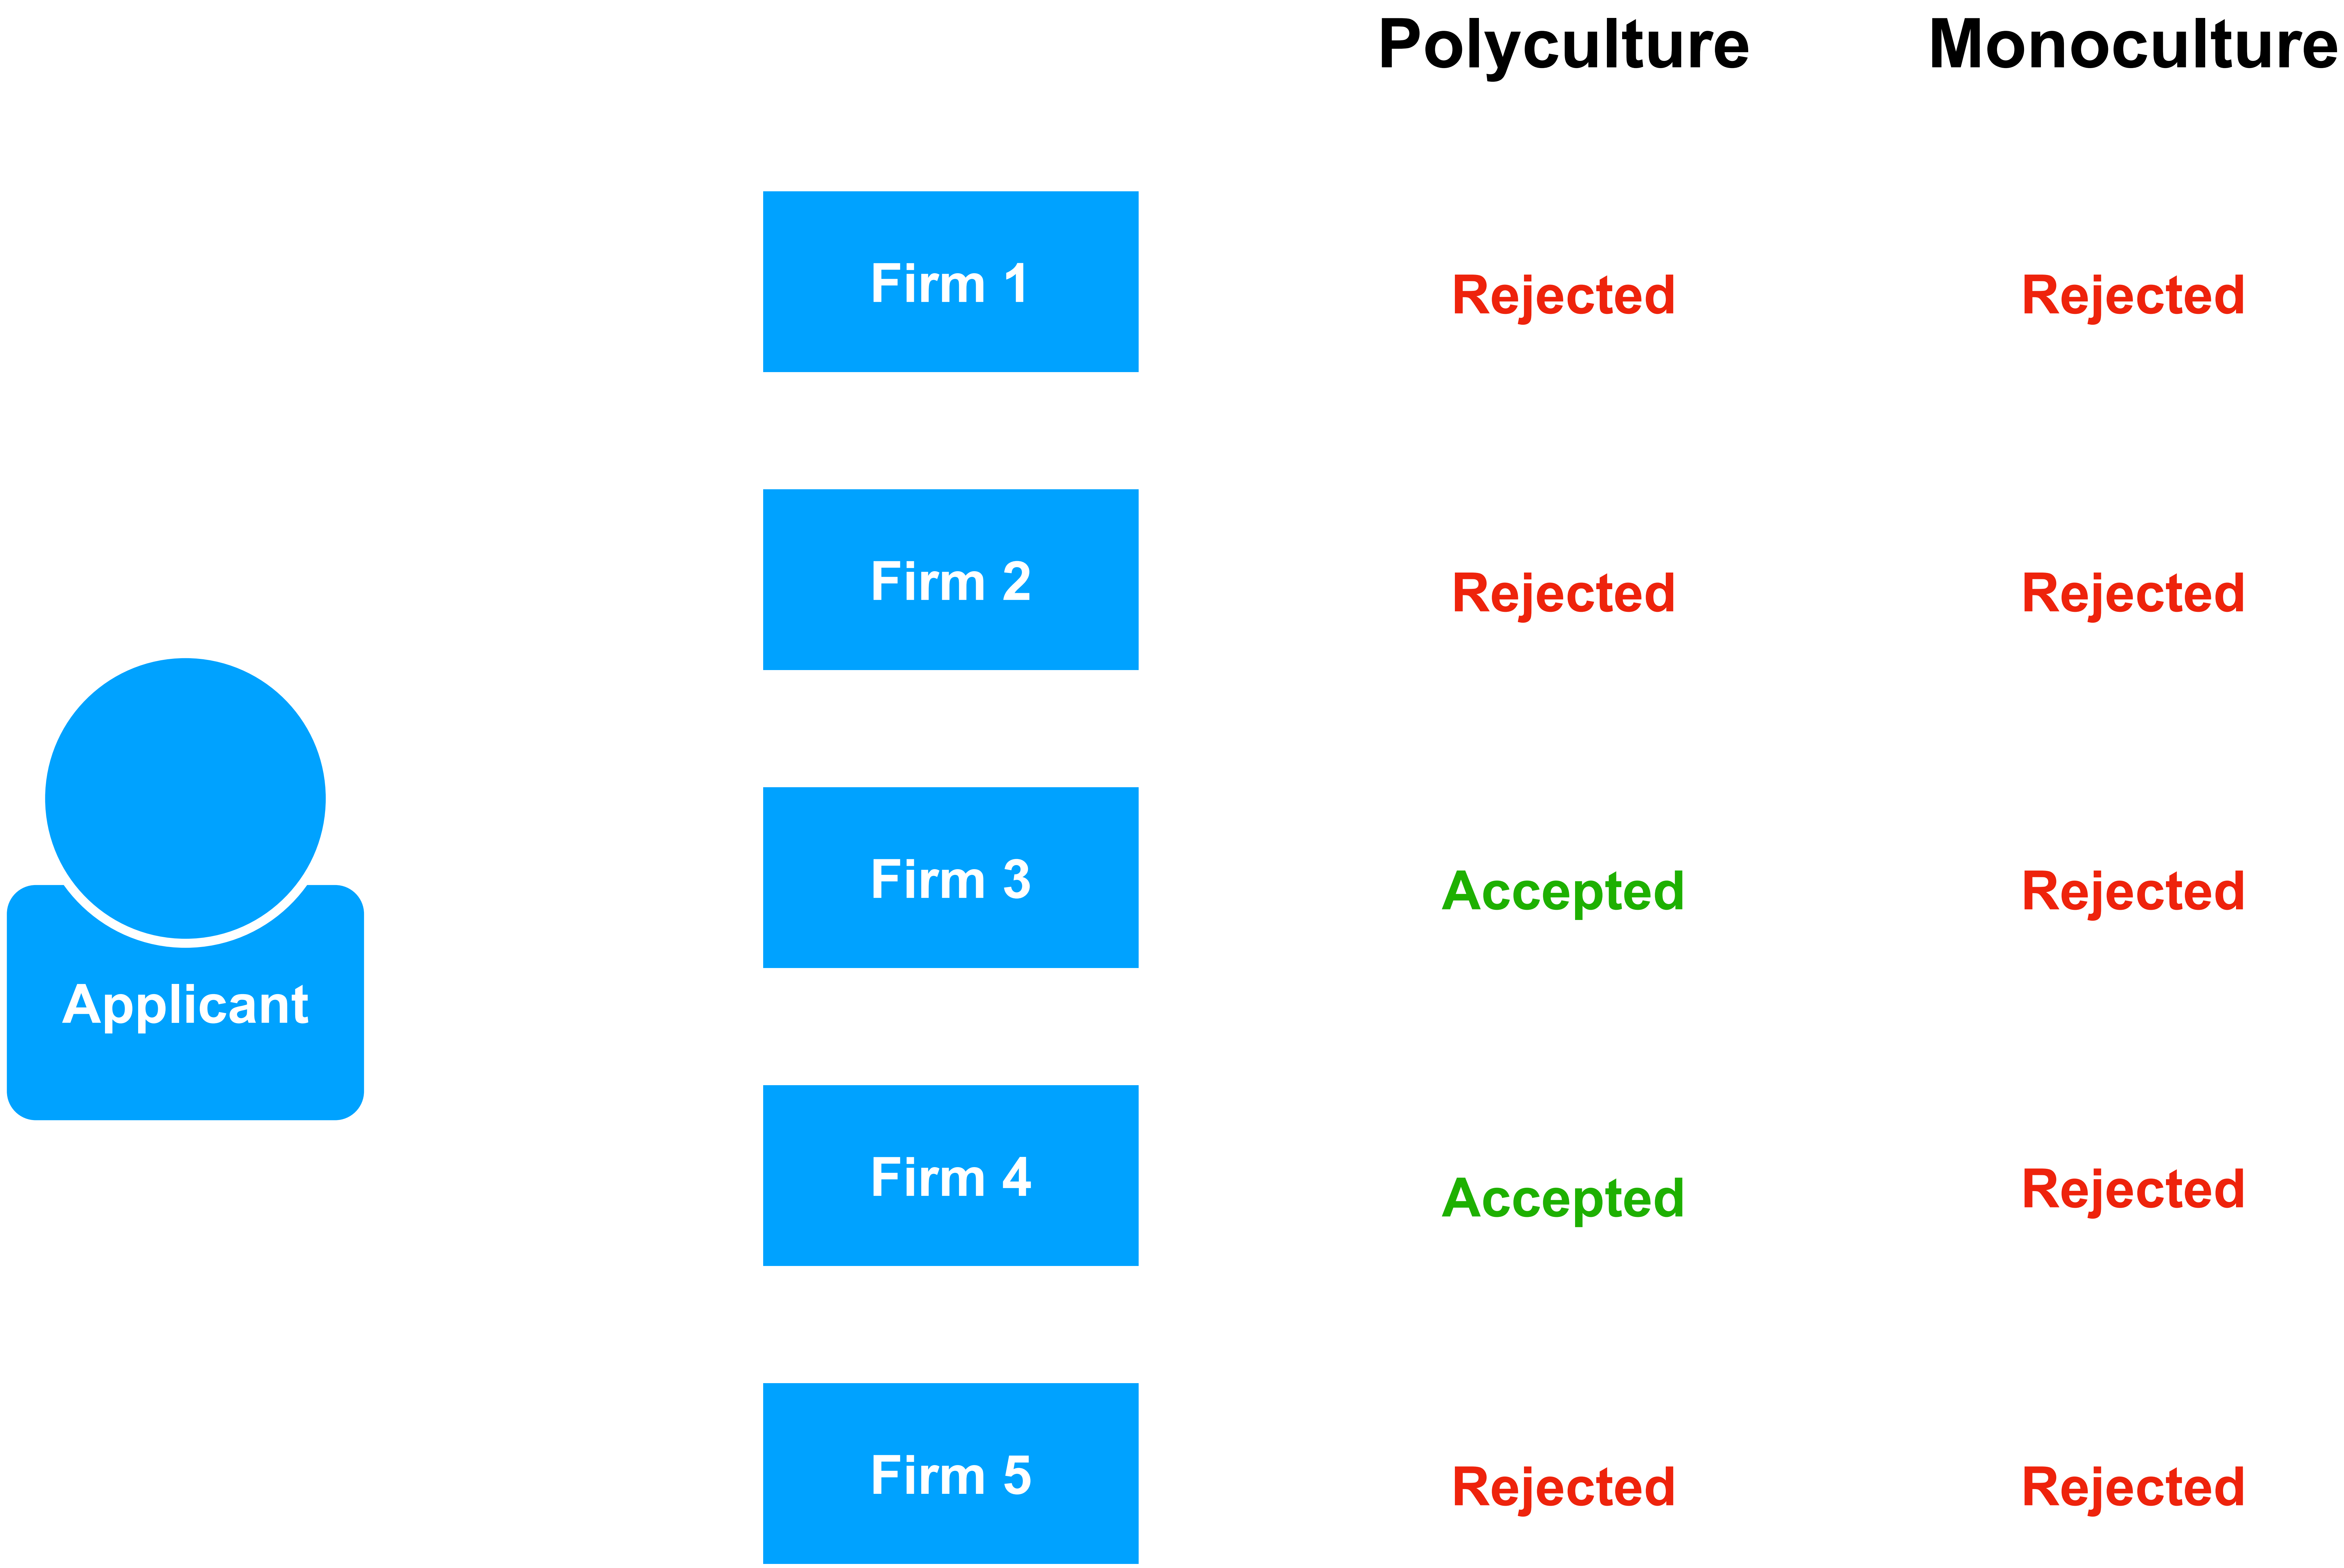
Polyculture
Monoculture
Applicant
Firm 1
Rejected
Rejected
Firm 2
Rejected
Rejected
Firm 3
Accepted
Rejected
Firm 4
Accepted
Rejected
Firm 5
Rejected
Rejected

Monoculture

Rejected

Rejected

More systemic exclusion

Rejected

Rejected

Rejected

[Creel & Hellman; Bommasani et al; Toups et al; Jain et al]

Monoculture

Rejected

Rejected

~~More systemic exclusion~~

Similar number of people should get hired overall in equilibrium!

(Firms do "yield math")

Rejected

Rejected

Rejected

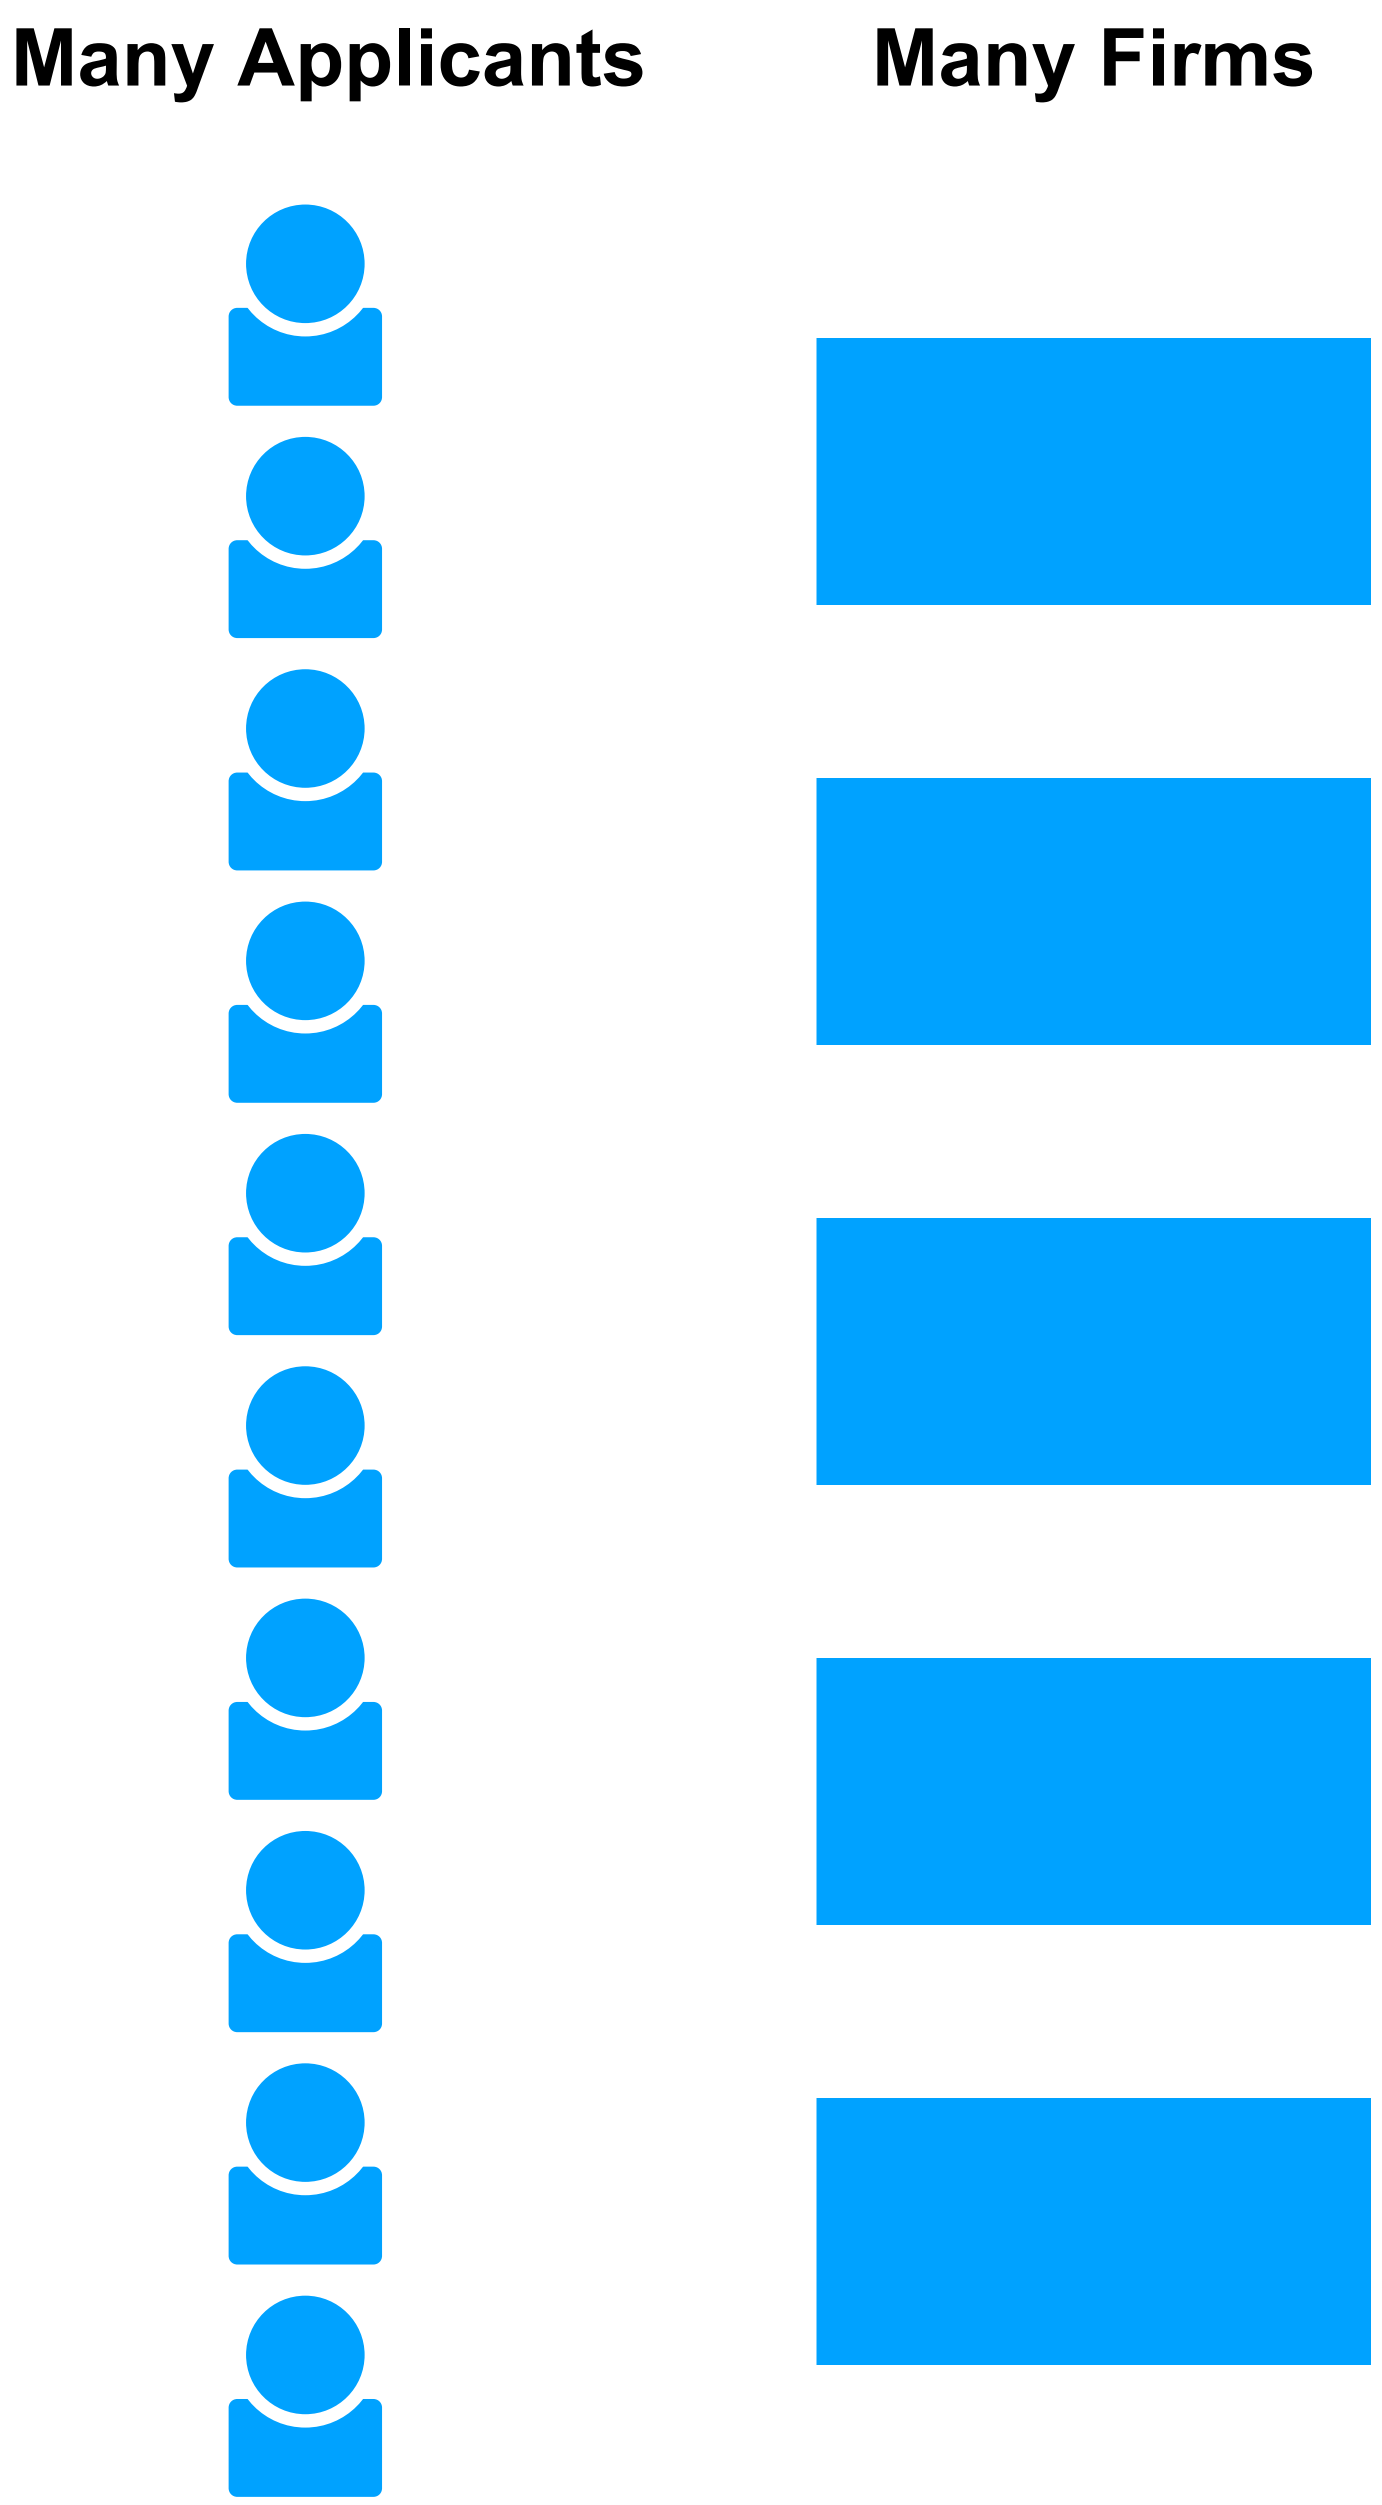

Need for incorporating *market-level* effects (e.g., stable matching as calculated by Gale Shapley algorithm)

Model in one slide

Adapt [Azevedo & Leshno] continuum model for stable matching

- There is a *continuum* of students (uncountably many students)
- Finite number of firms (we will take the limit of number of colleges)
- Students have uniform at random rankings over firms*
- *True preferences* of firms depend on student value v
- Firms *estimated* rankings v + noise, where noise ~ D
- We analyze the stable matching using estimated preferences

Lemma and intuition

In polyculture: whether you get hired ~depends on *maximum* score

In monoculture, only on a single draw

$\Rightarrow$In polyculture, applicants get “more lottery tickets”

$\Rightarrow$Thus, firm cutoffs (admission standards) are higher

Proof strategy: reason about **max order statistic** -- what is the distribution of the *max* score that someone receives?

Polyculture

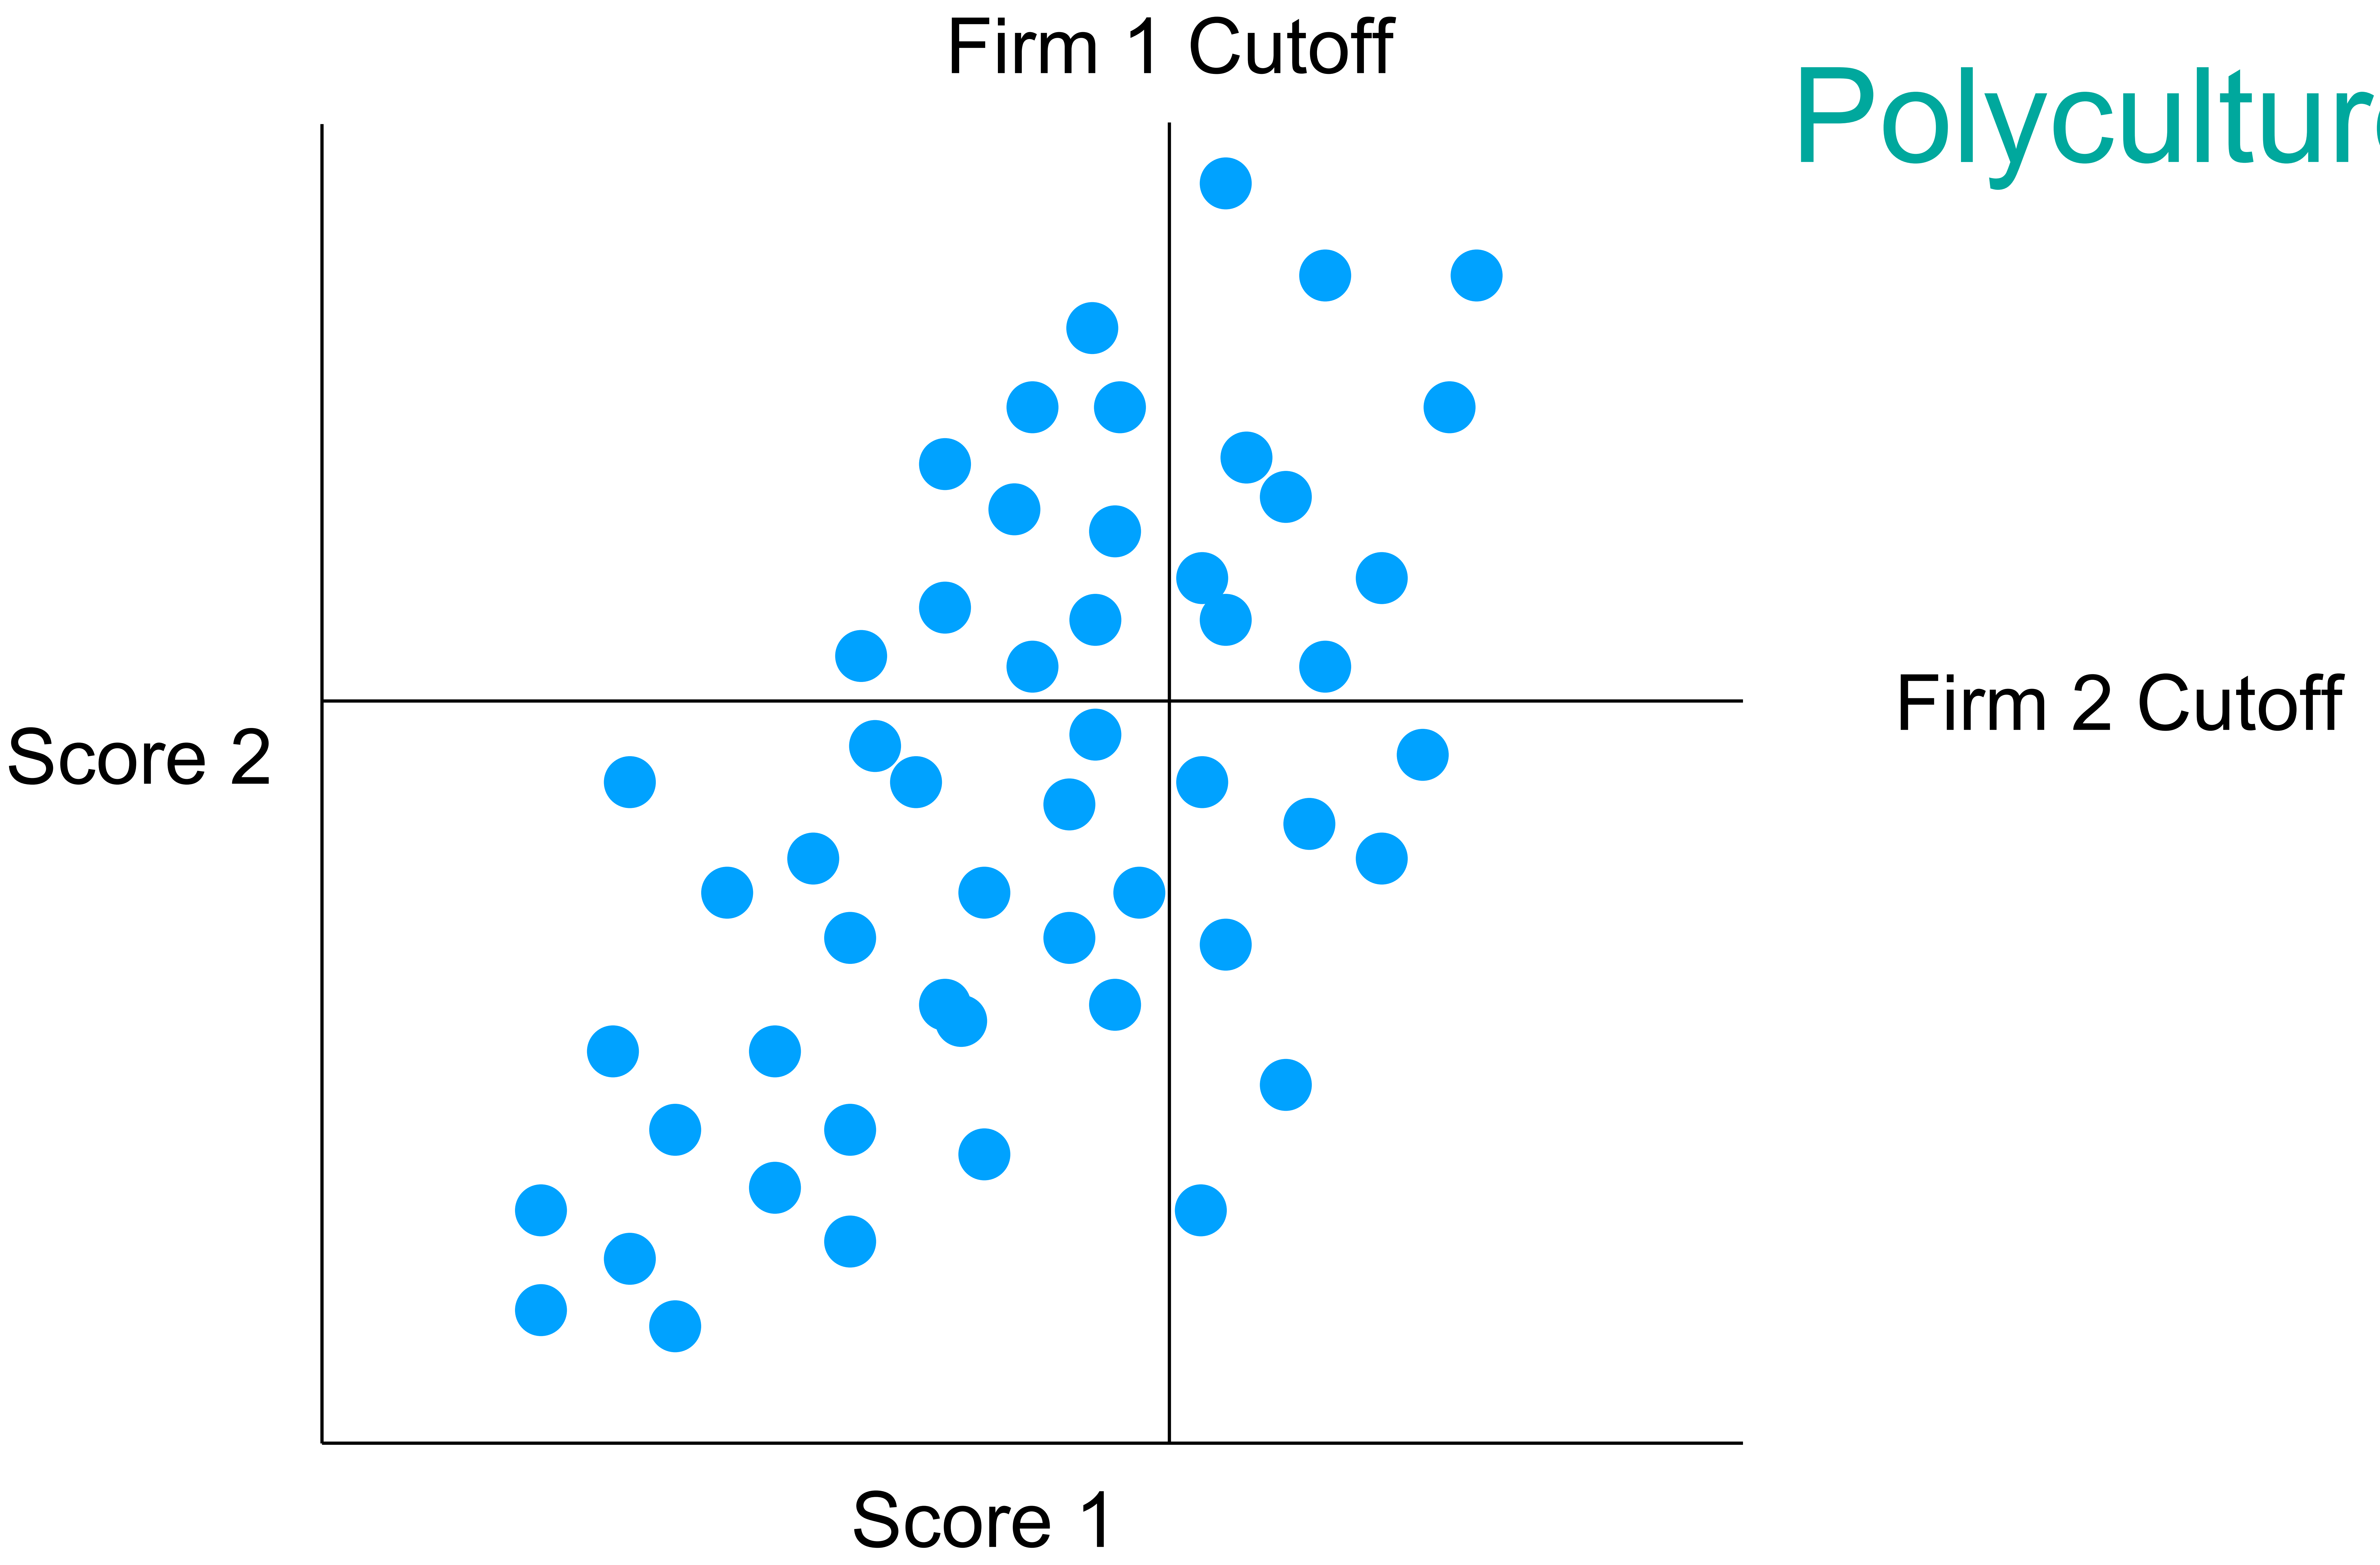

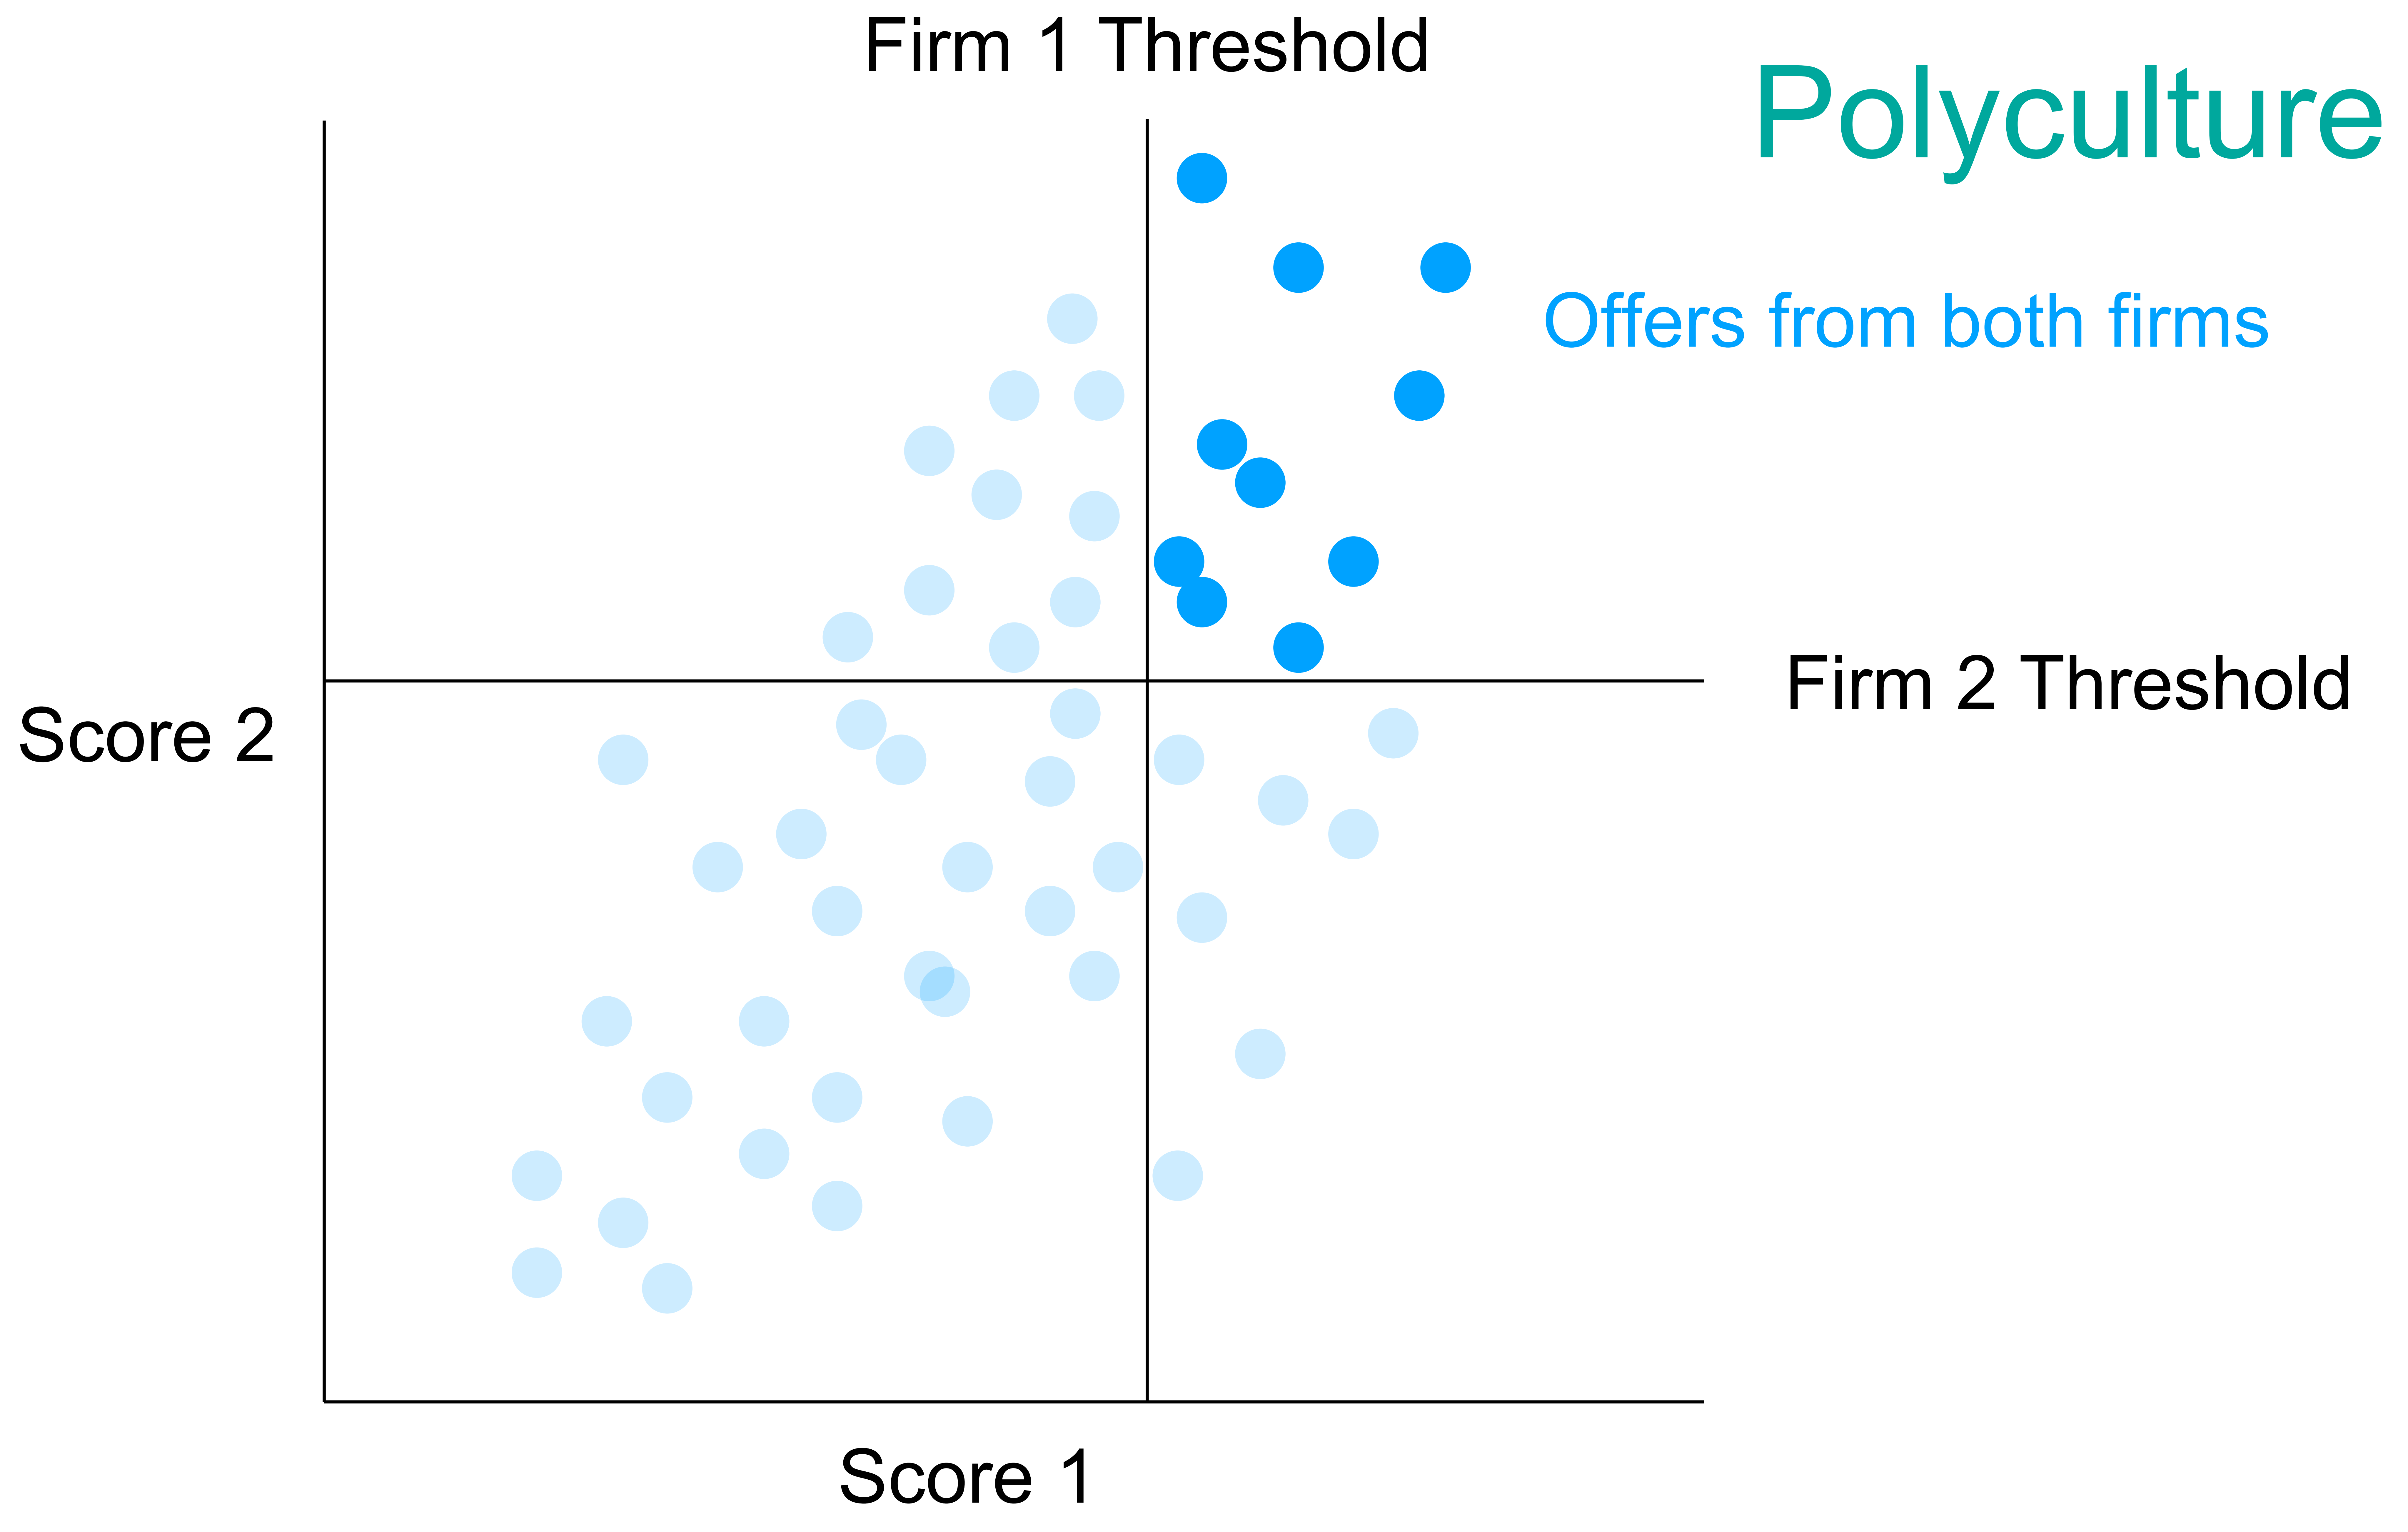
Polyculture
Firm 1 Threshold
Offers from both firms
Firm 2 Threshold
Score 2
Score 1

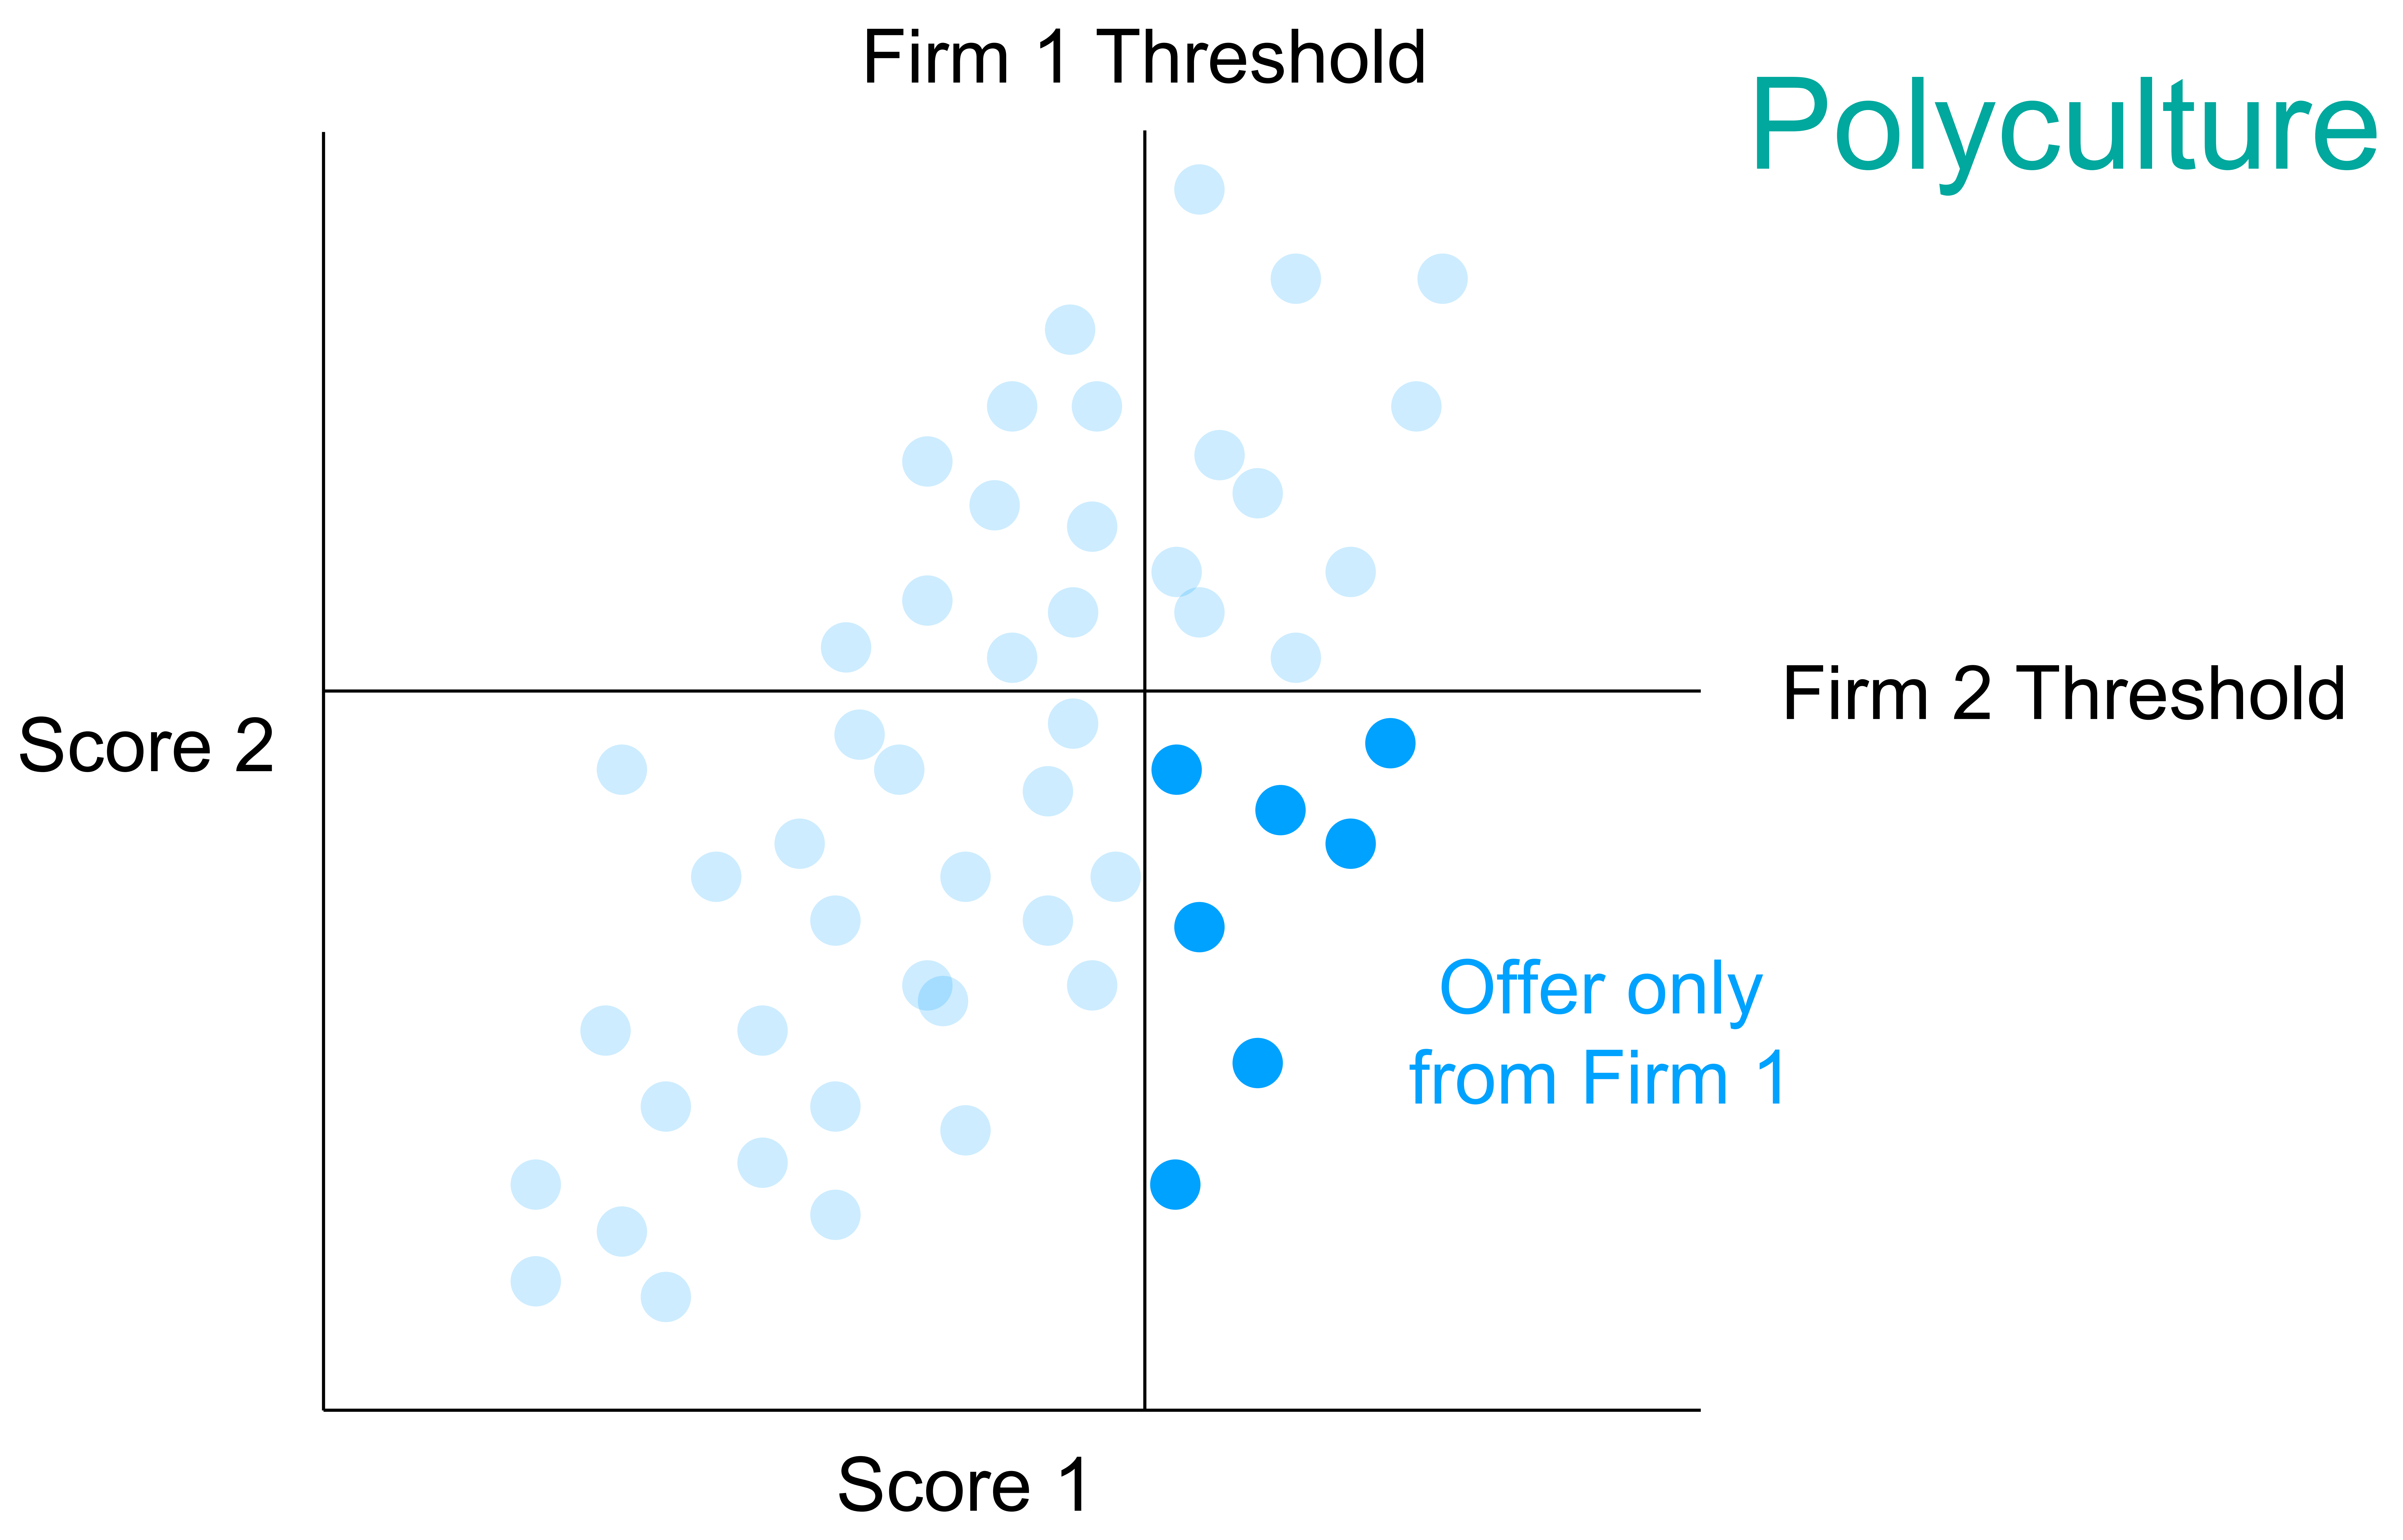
Firm 1 Threshold
Polyculture
Firm 2 Threshold
Score 2
Offer only
from Firm 1
Score 1

Polyculture

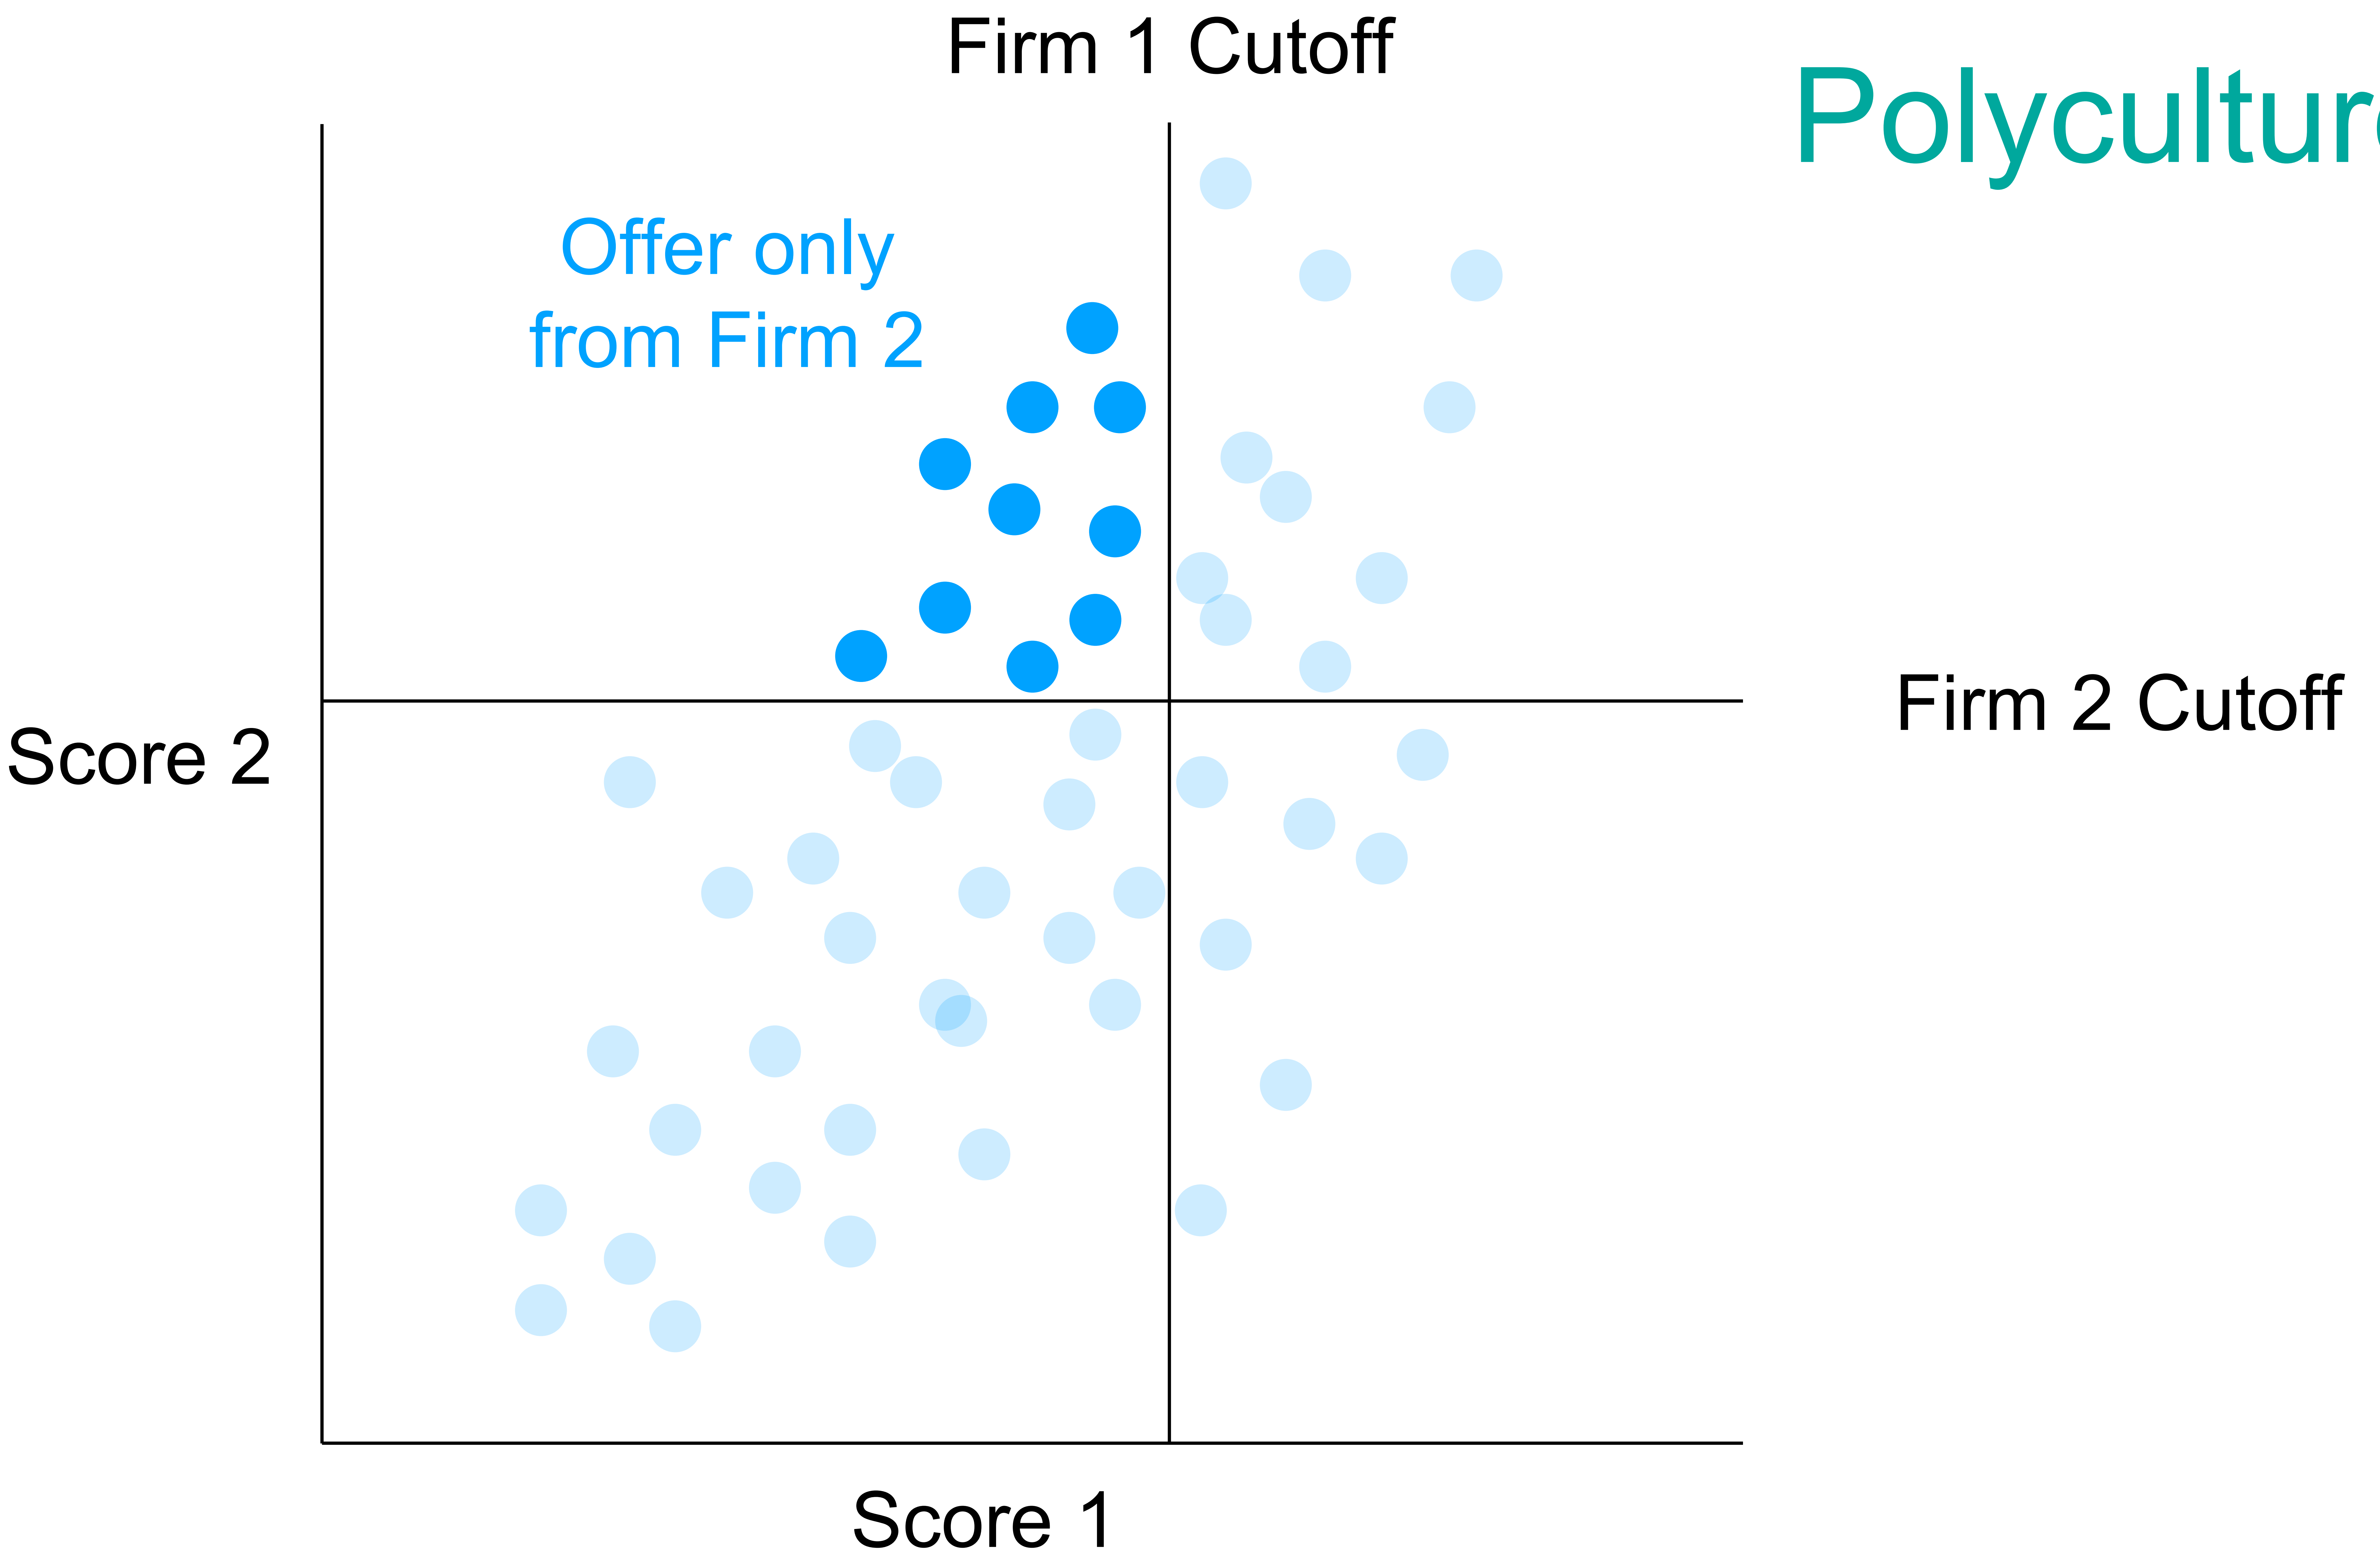

Polyculture

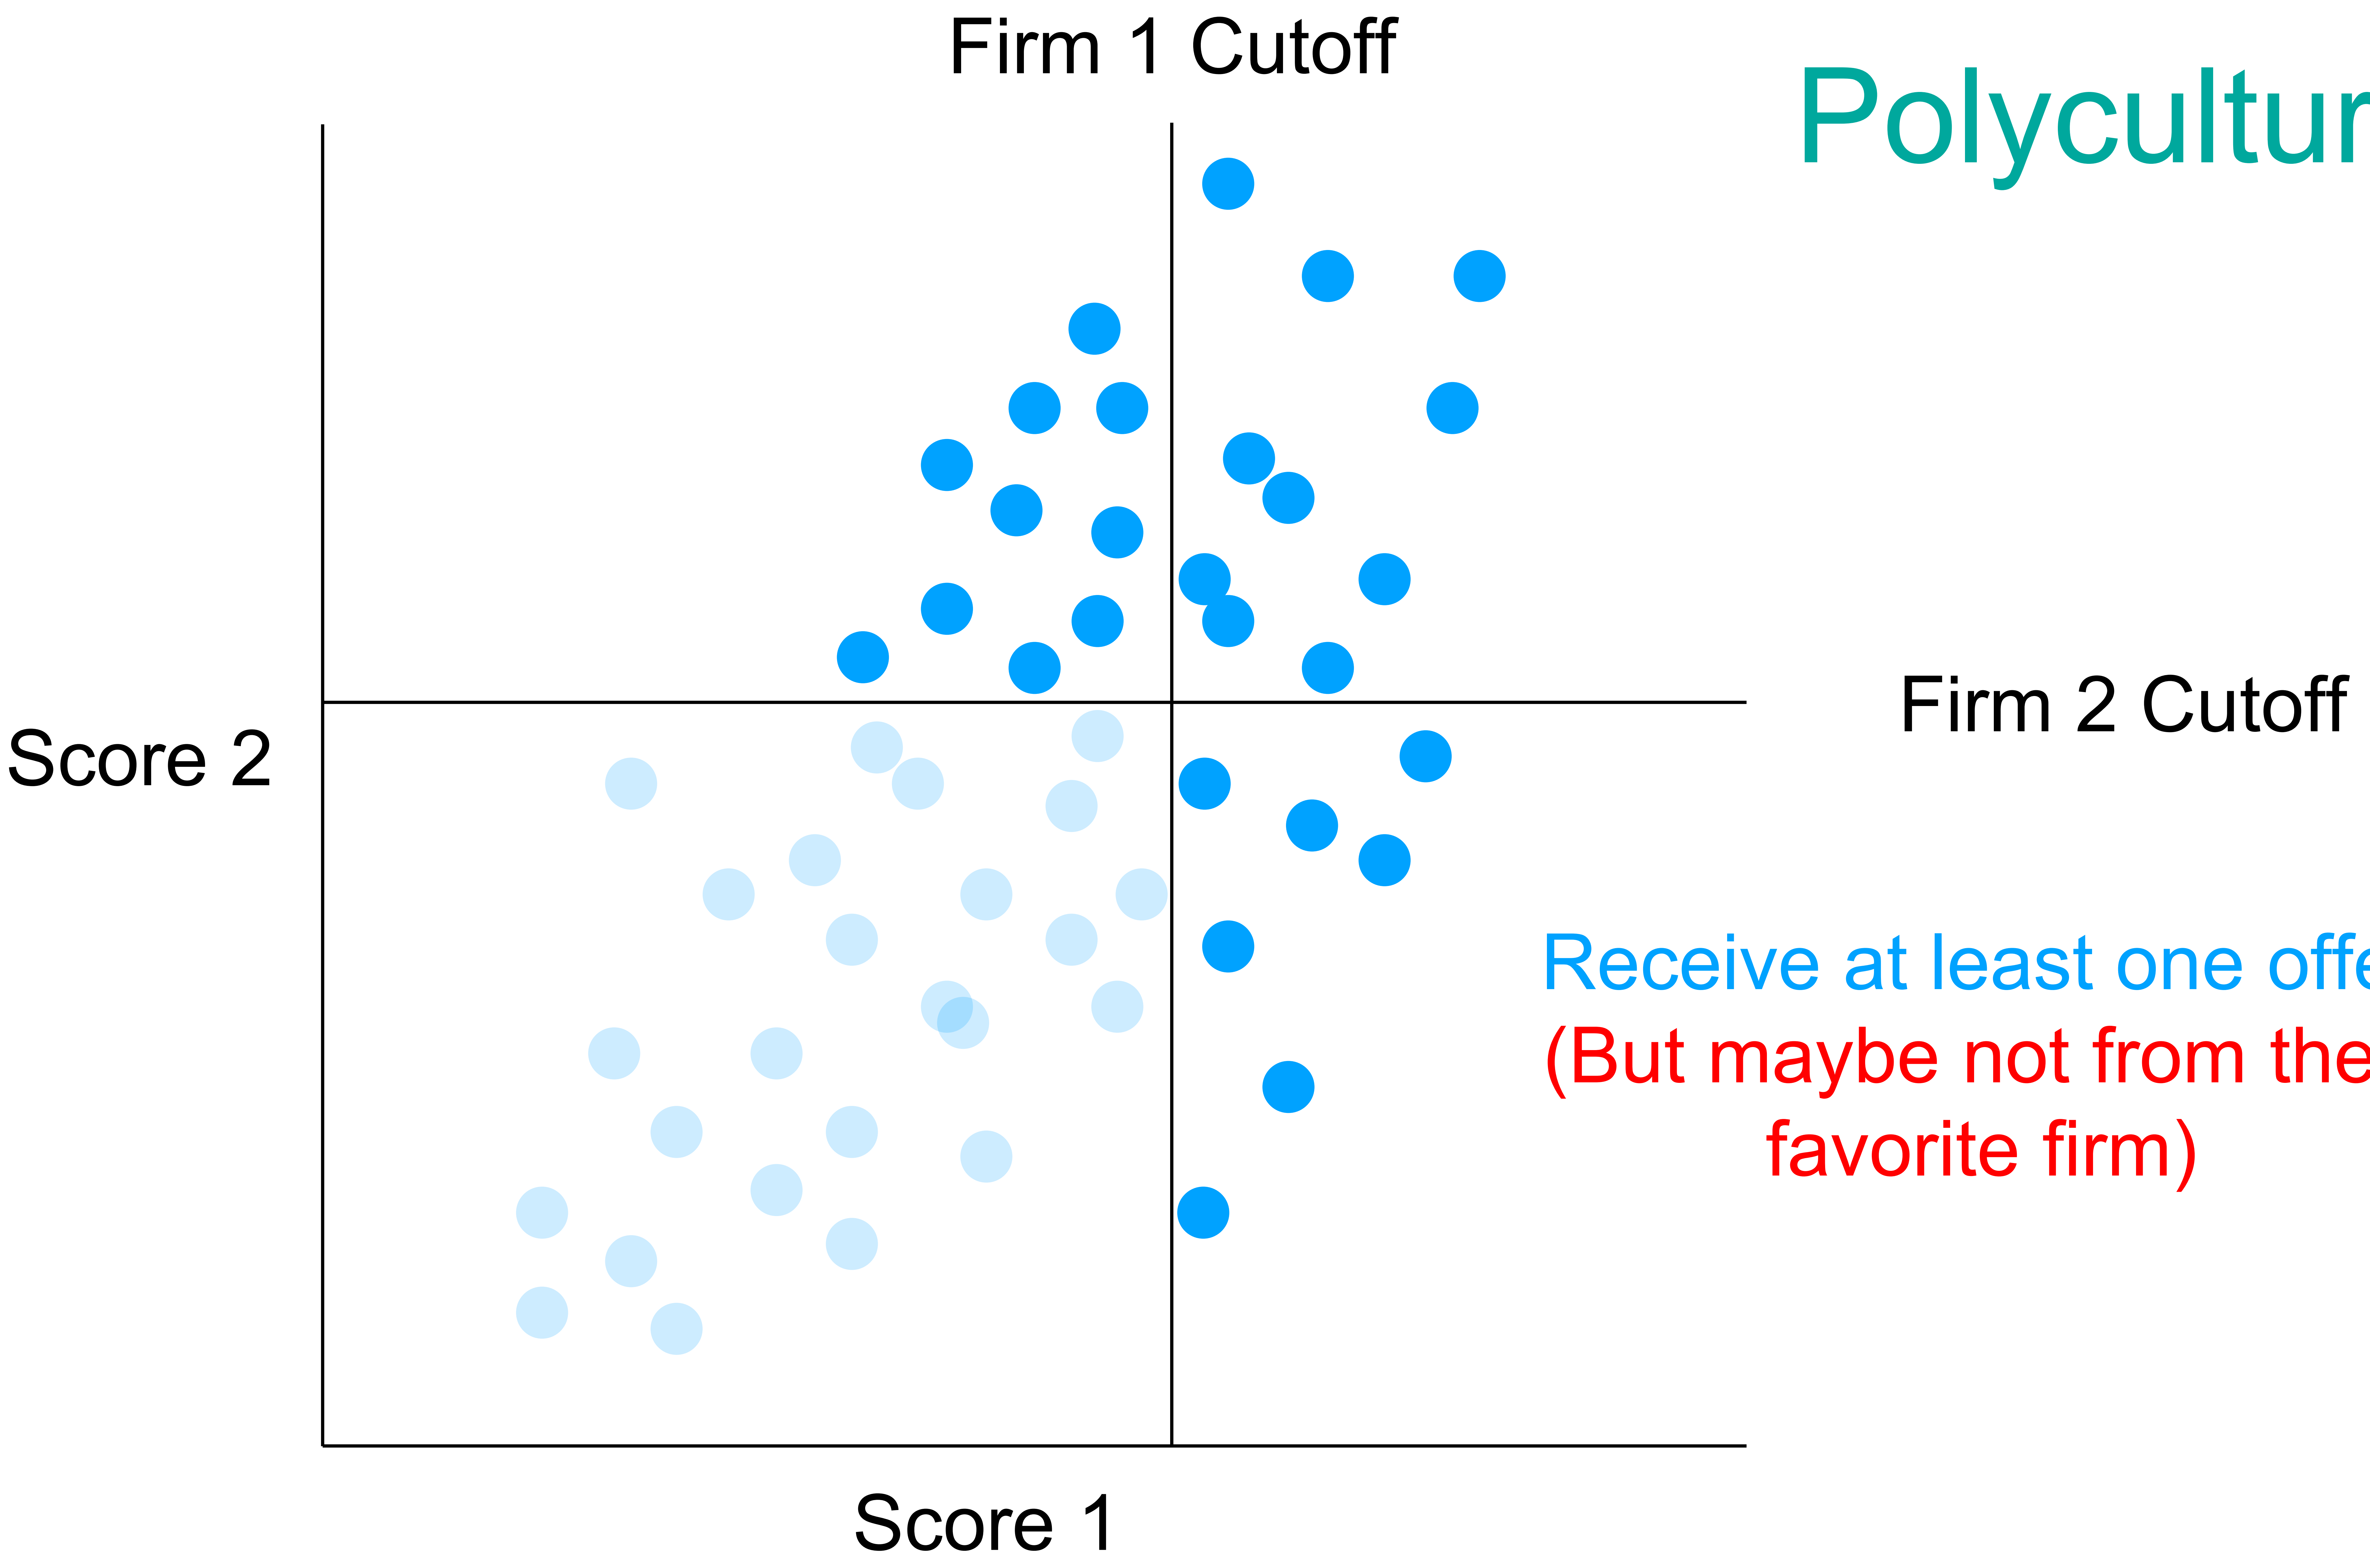

Receive at least one offer

(But maybe not from their favorite firm)

Monoculture

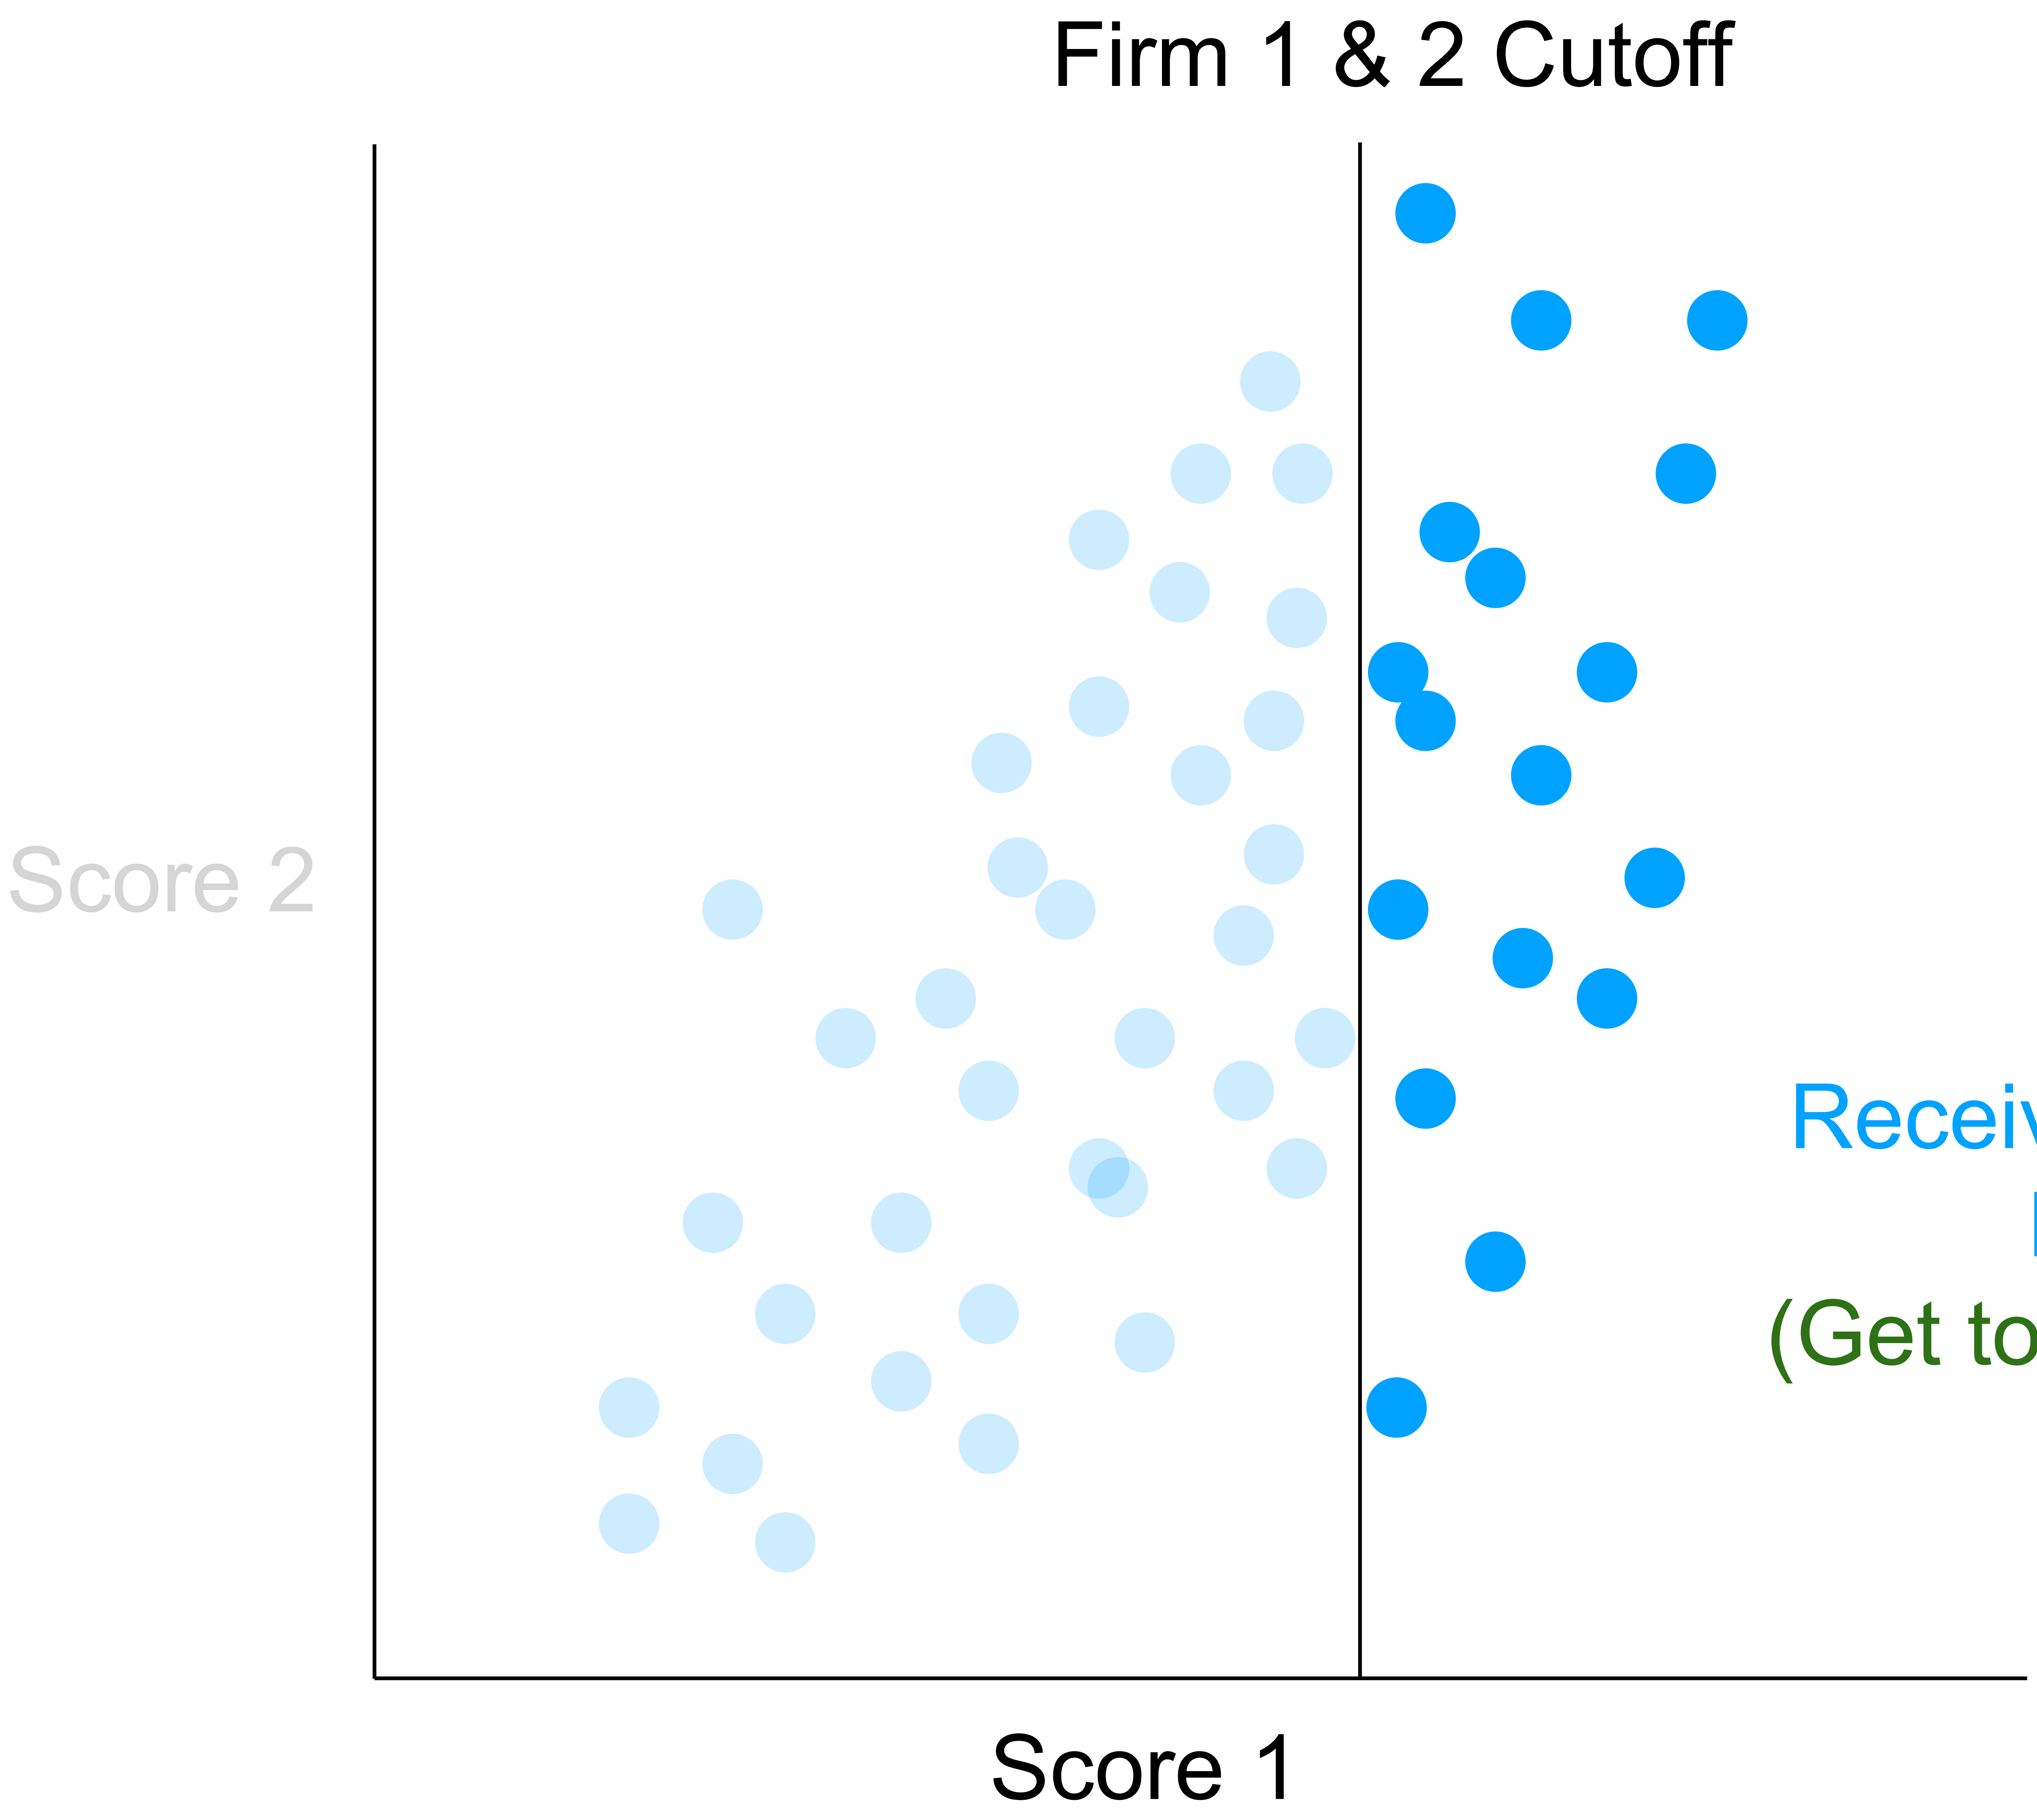

Receive offers from both Firm 1 and 2

(Get to go to favorite firm)

Monoculture

Receive offers from both firm 1 and 2

Cutoff goes down to preserve # accepted

Theorem 1: Wisdom of crowds

Under polyculture in large markets,*
firms "hire the right applicants," but not under monoculture

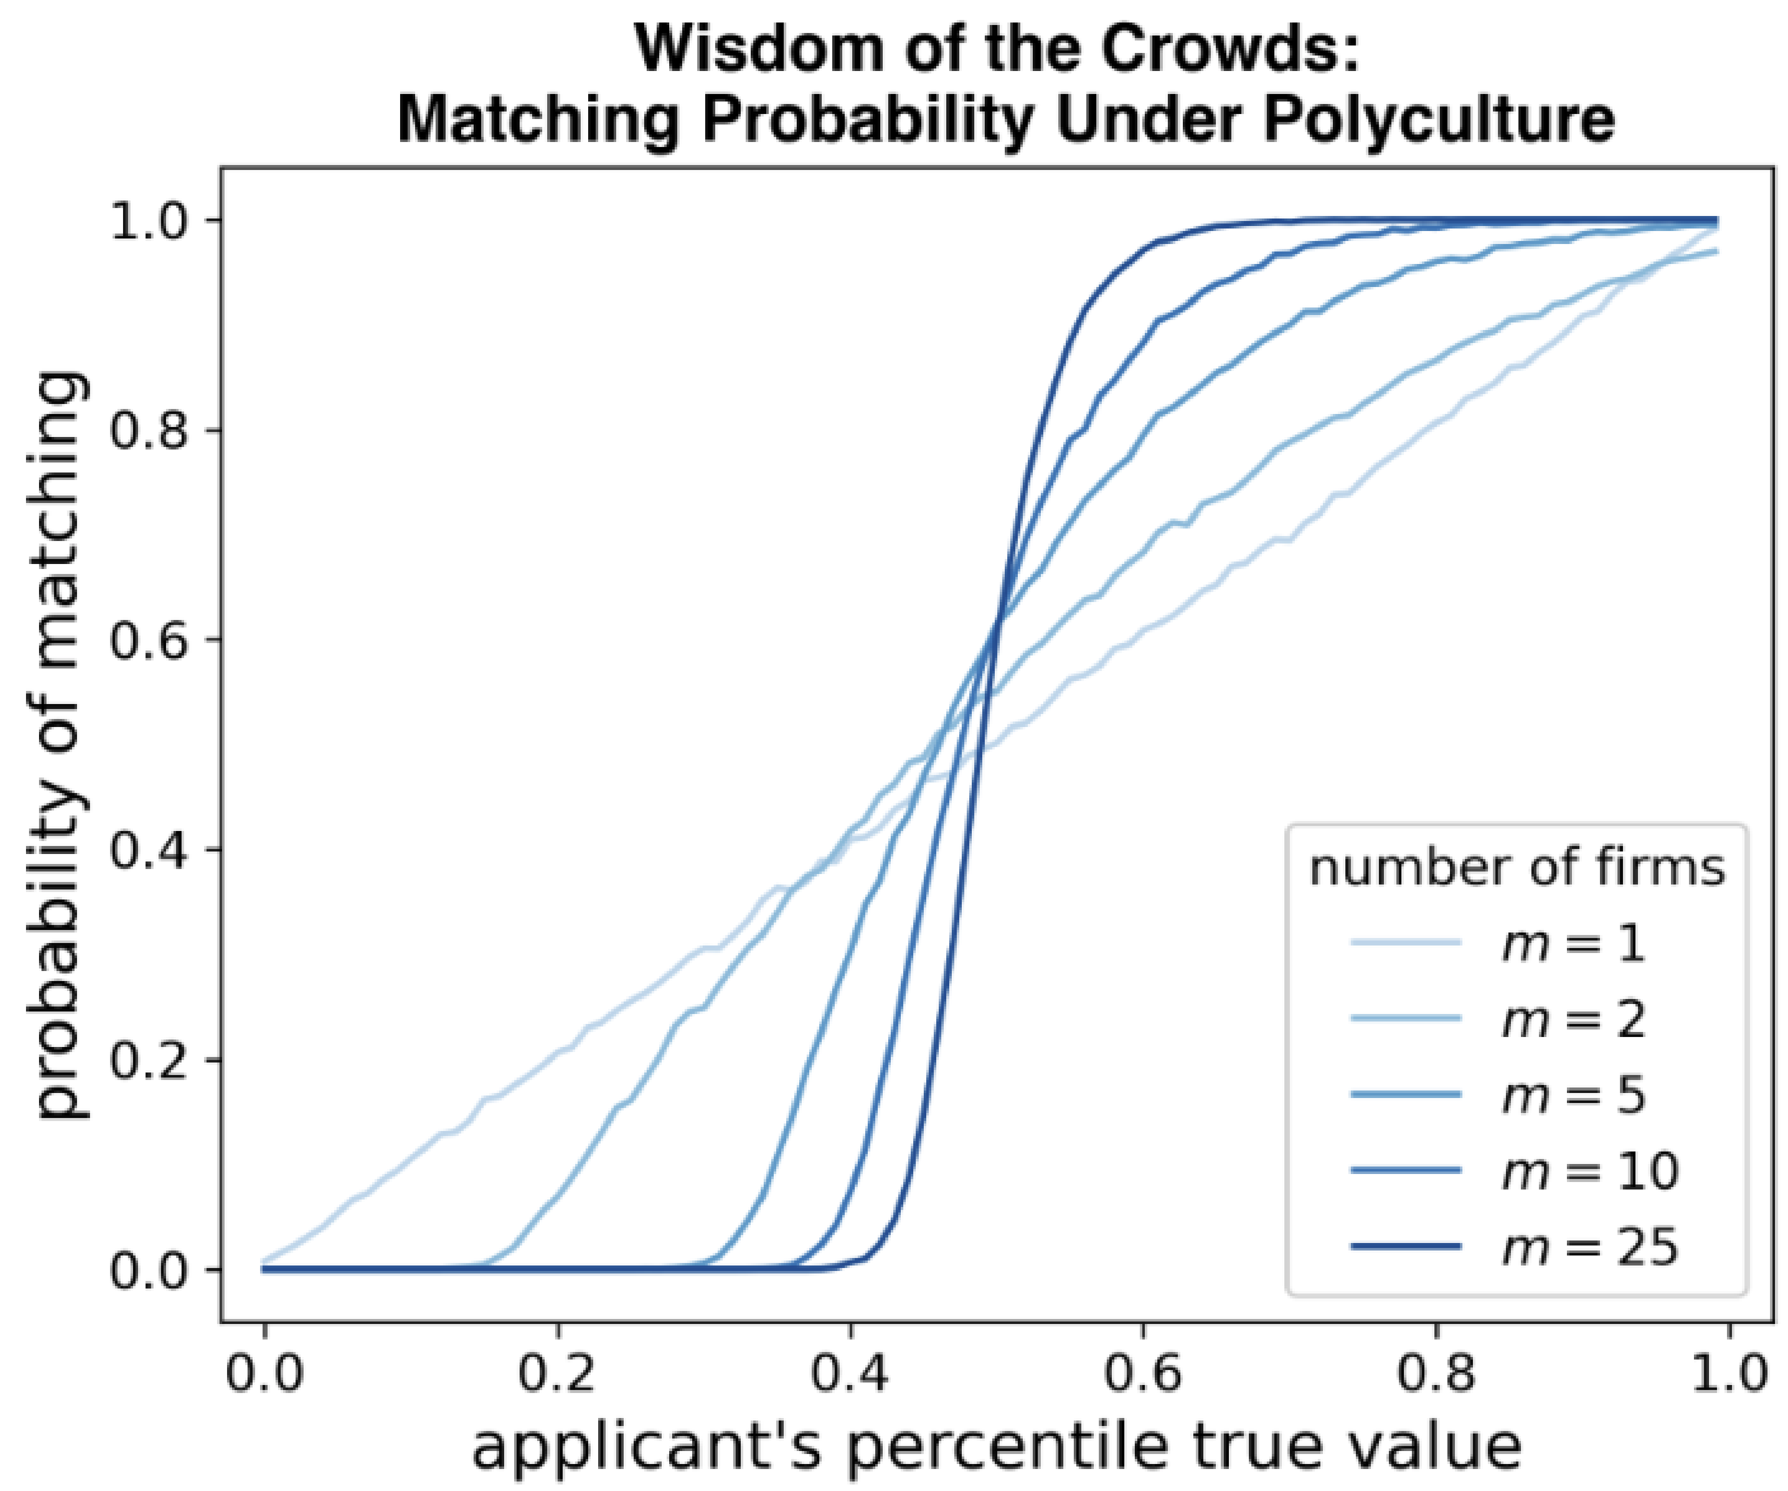

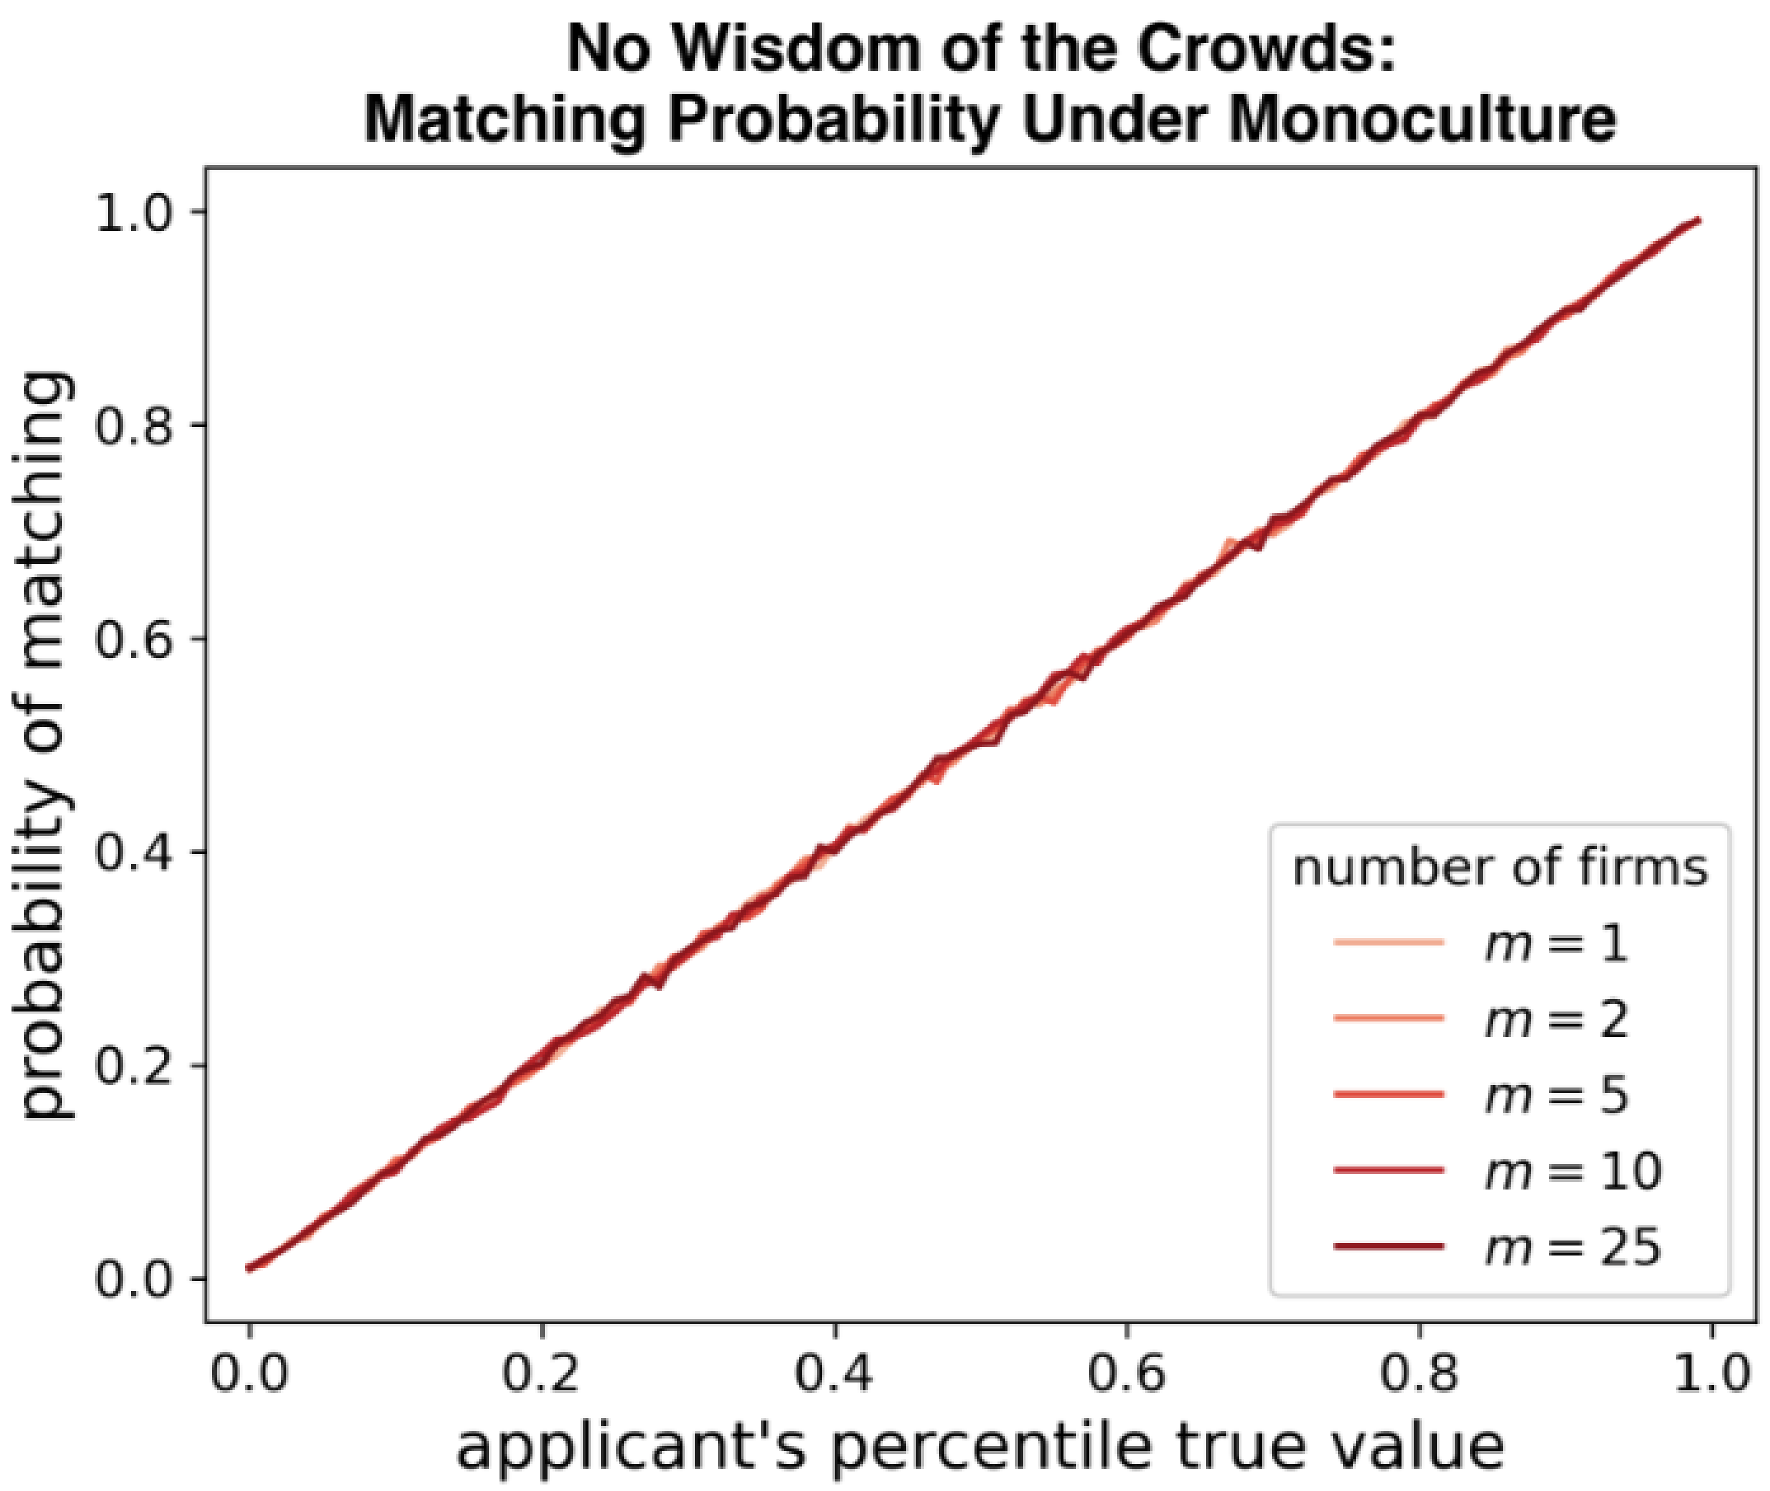

But is only true with “max concentrating” noise

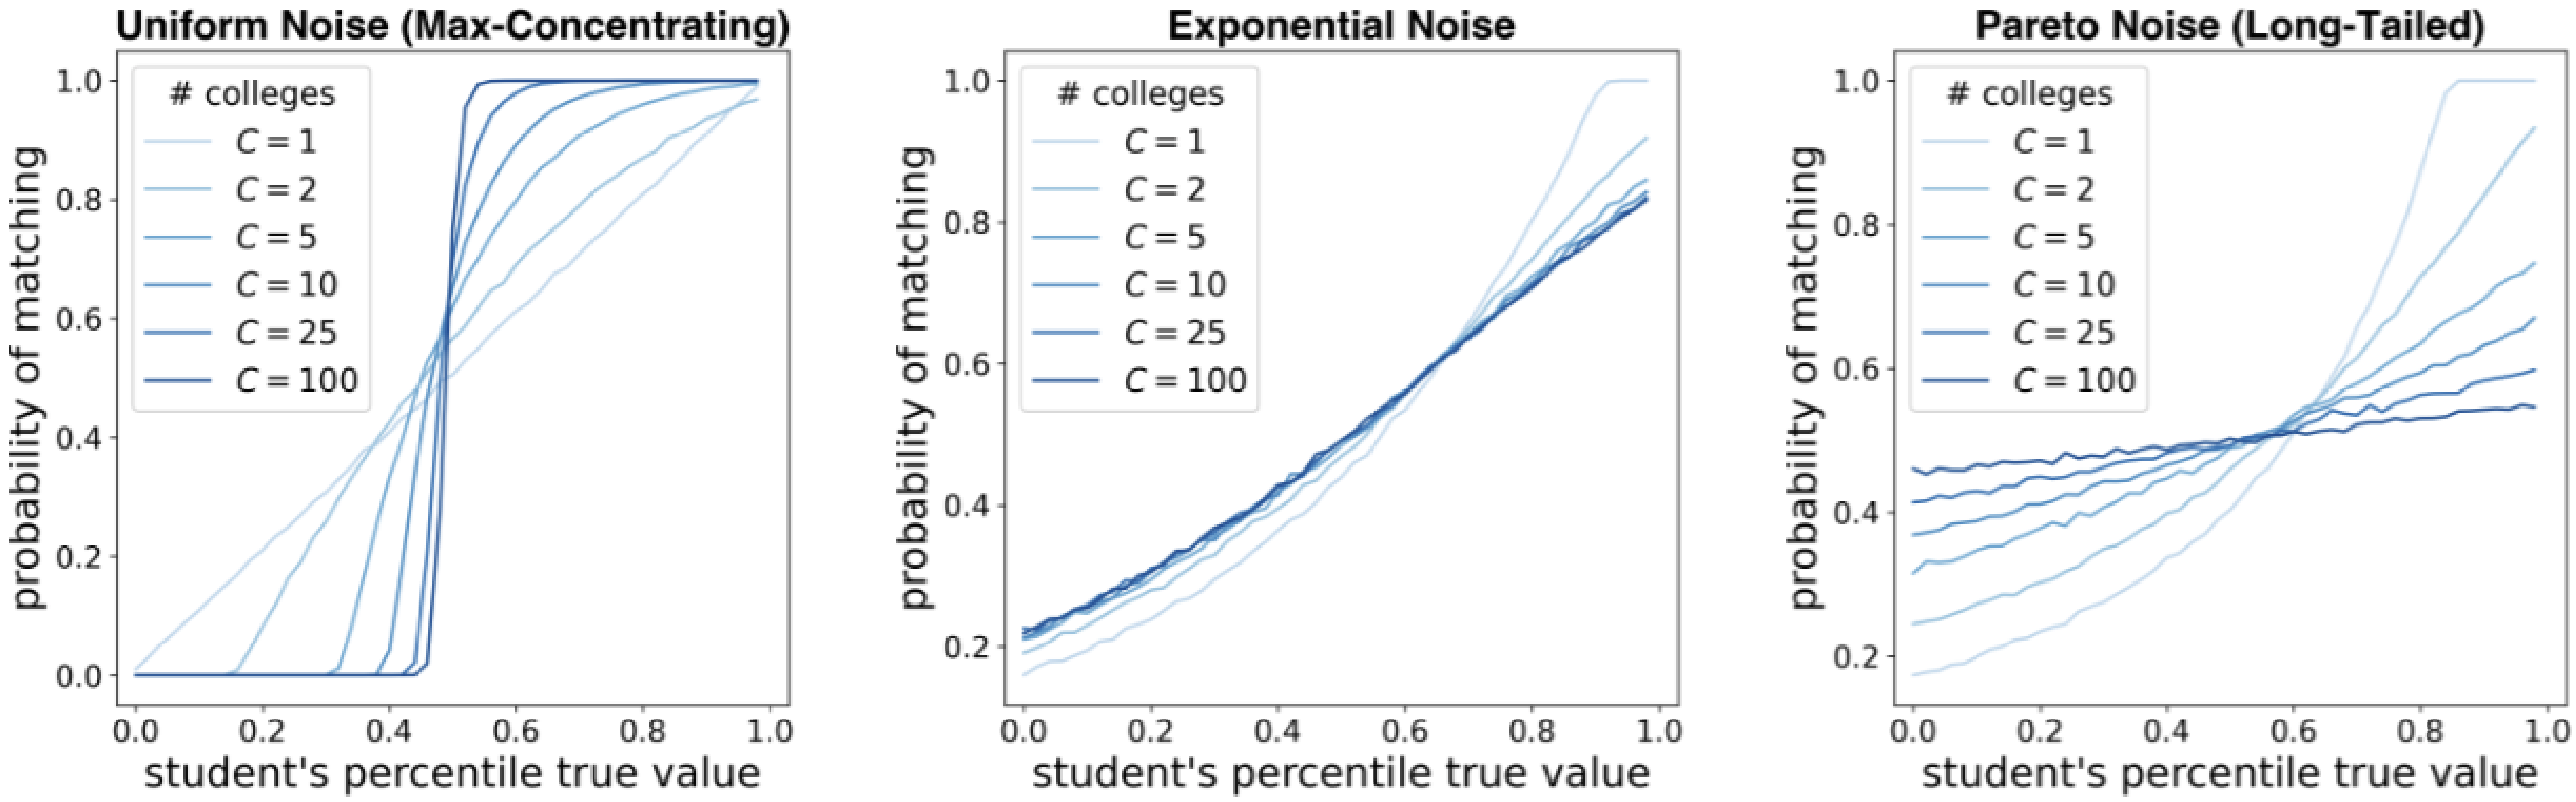

(a) Noise Attenuation as $C \to \infty$ (Wisdom of Crowds)

(b) Noise Stays $\approx$ Same

(c) Noise Amplification as $C \to \infty$ (Foolishness of Crowds)

Wisdom and Foolishness of Noisy Matching Markets. Kenny Peng and **NG**

Can extend more generally

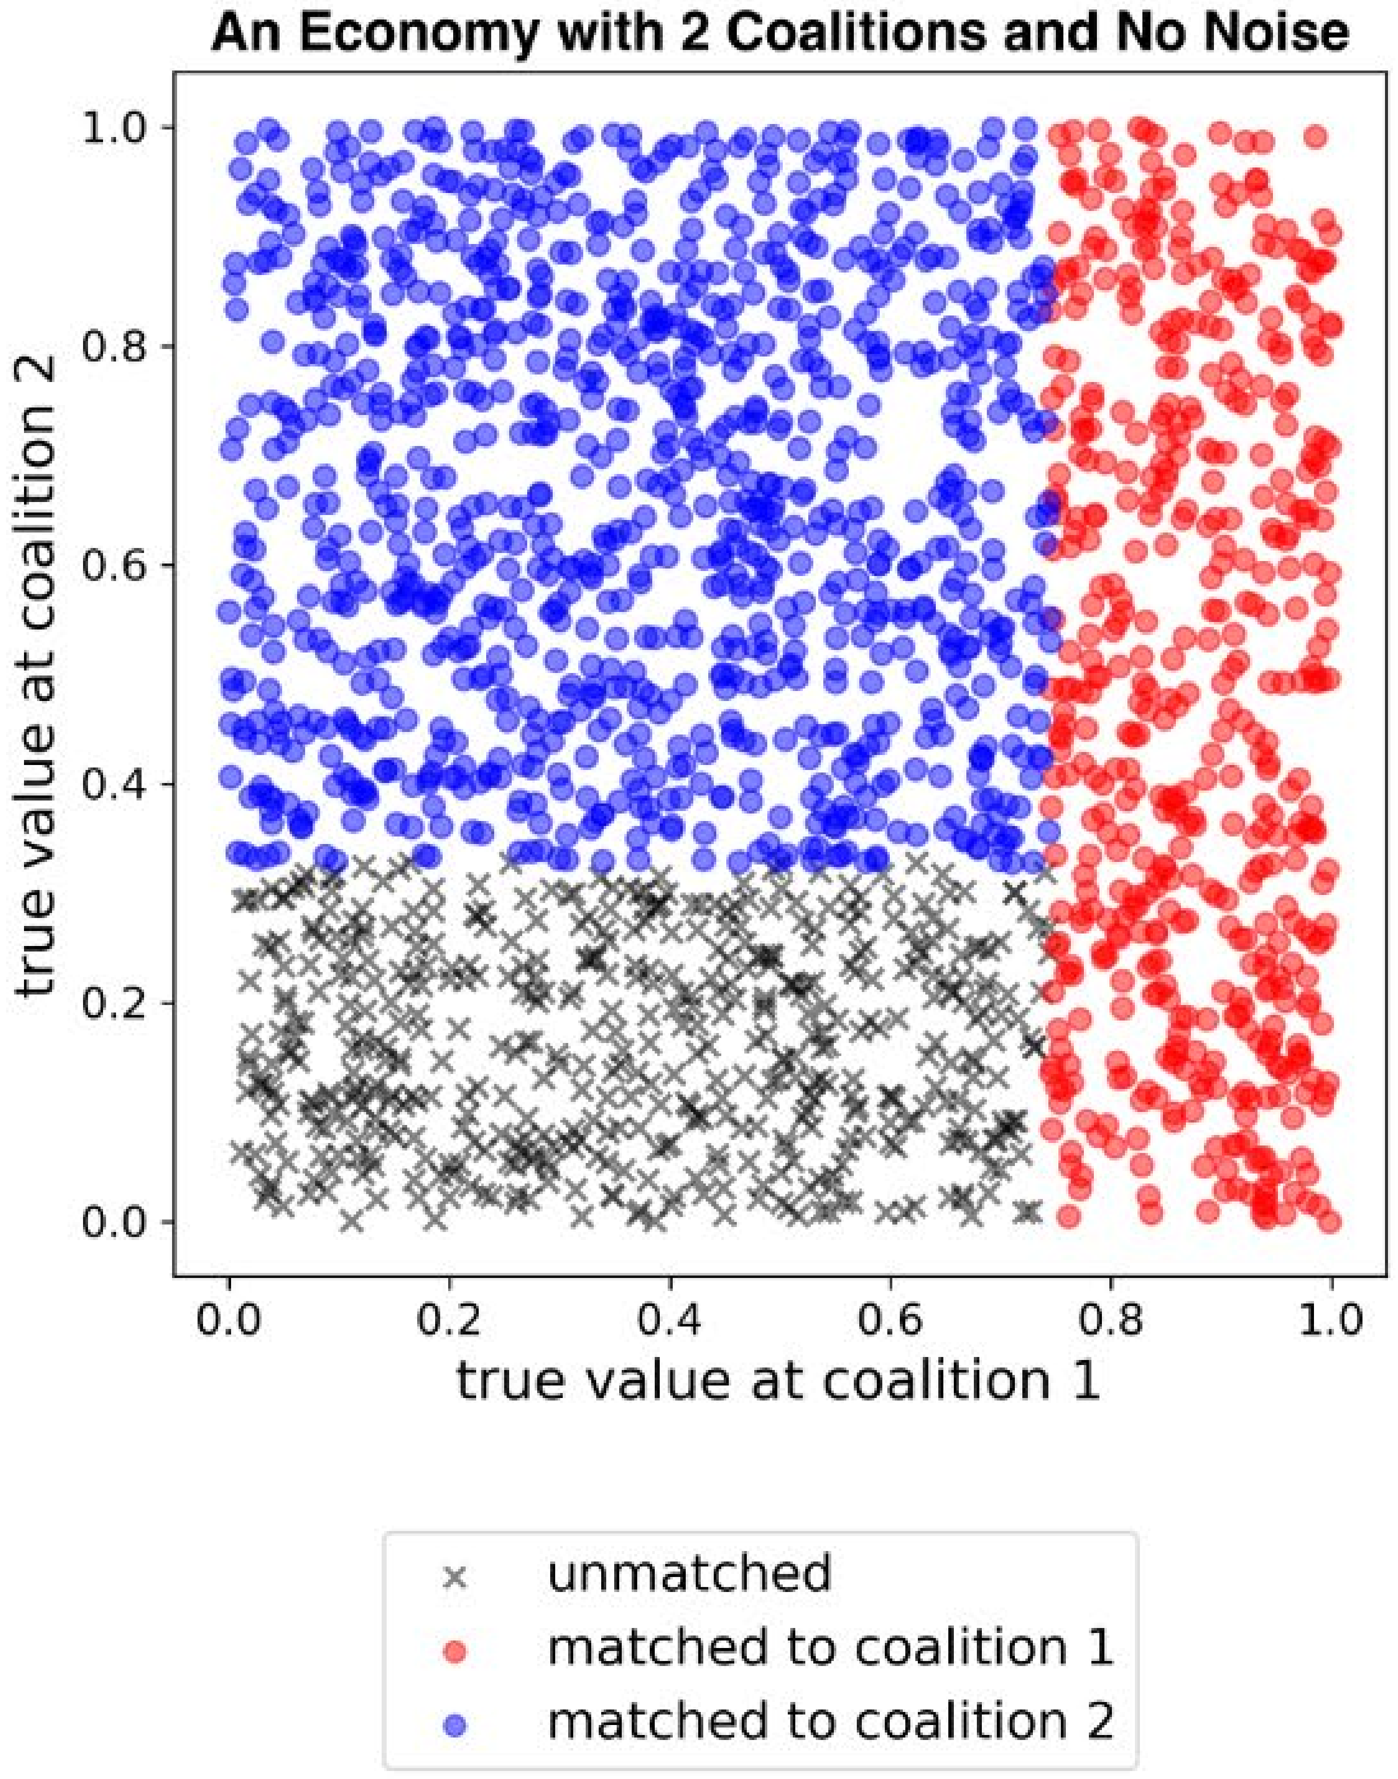

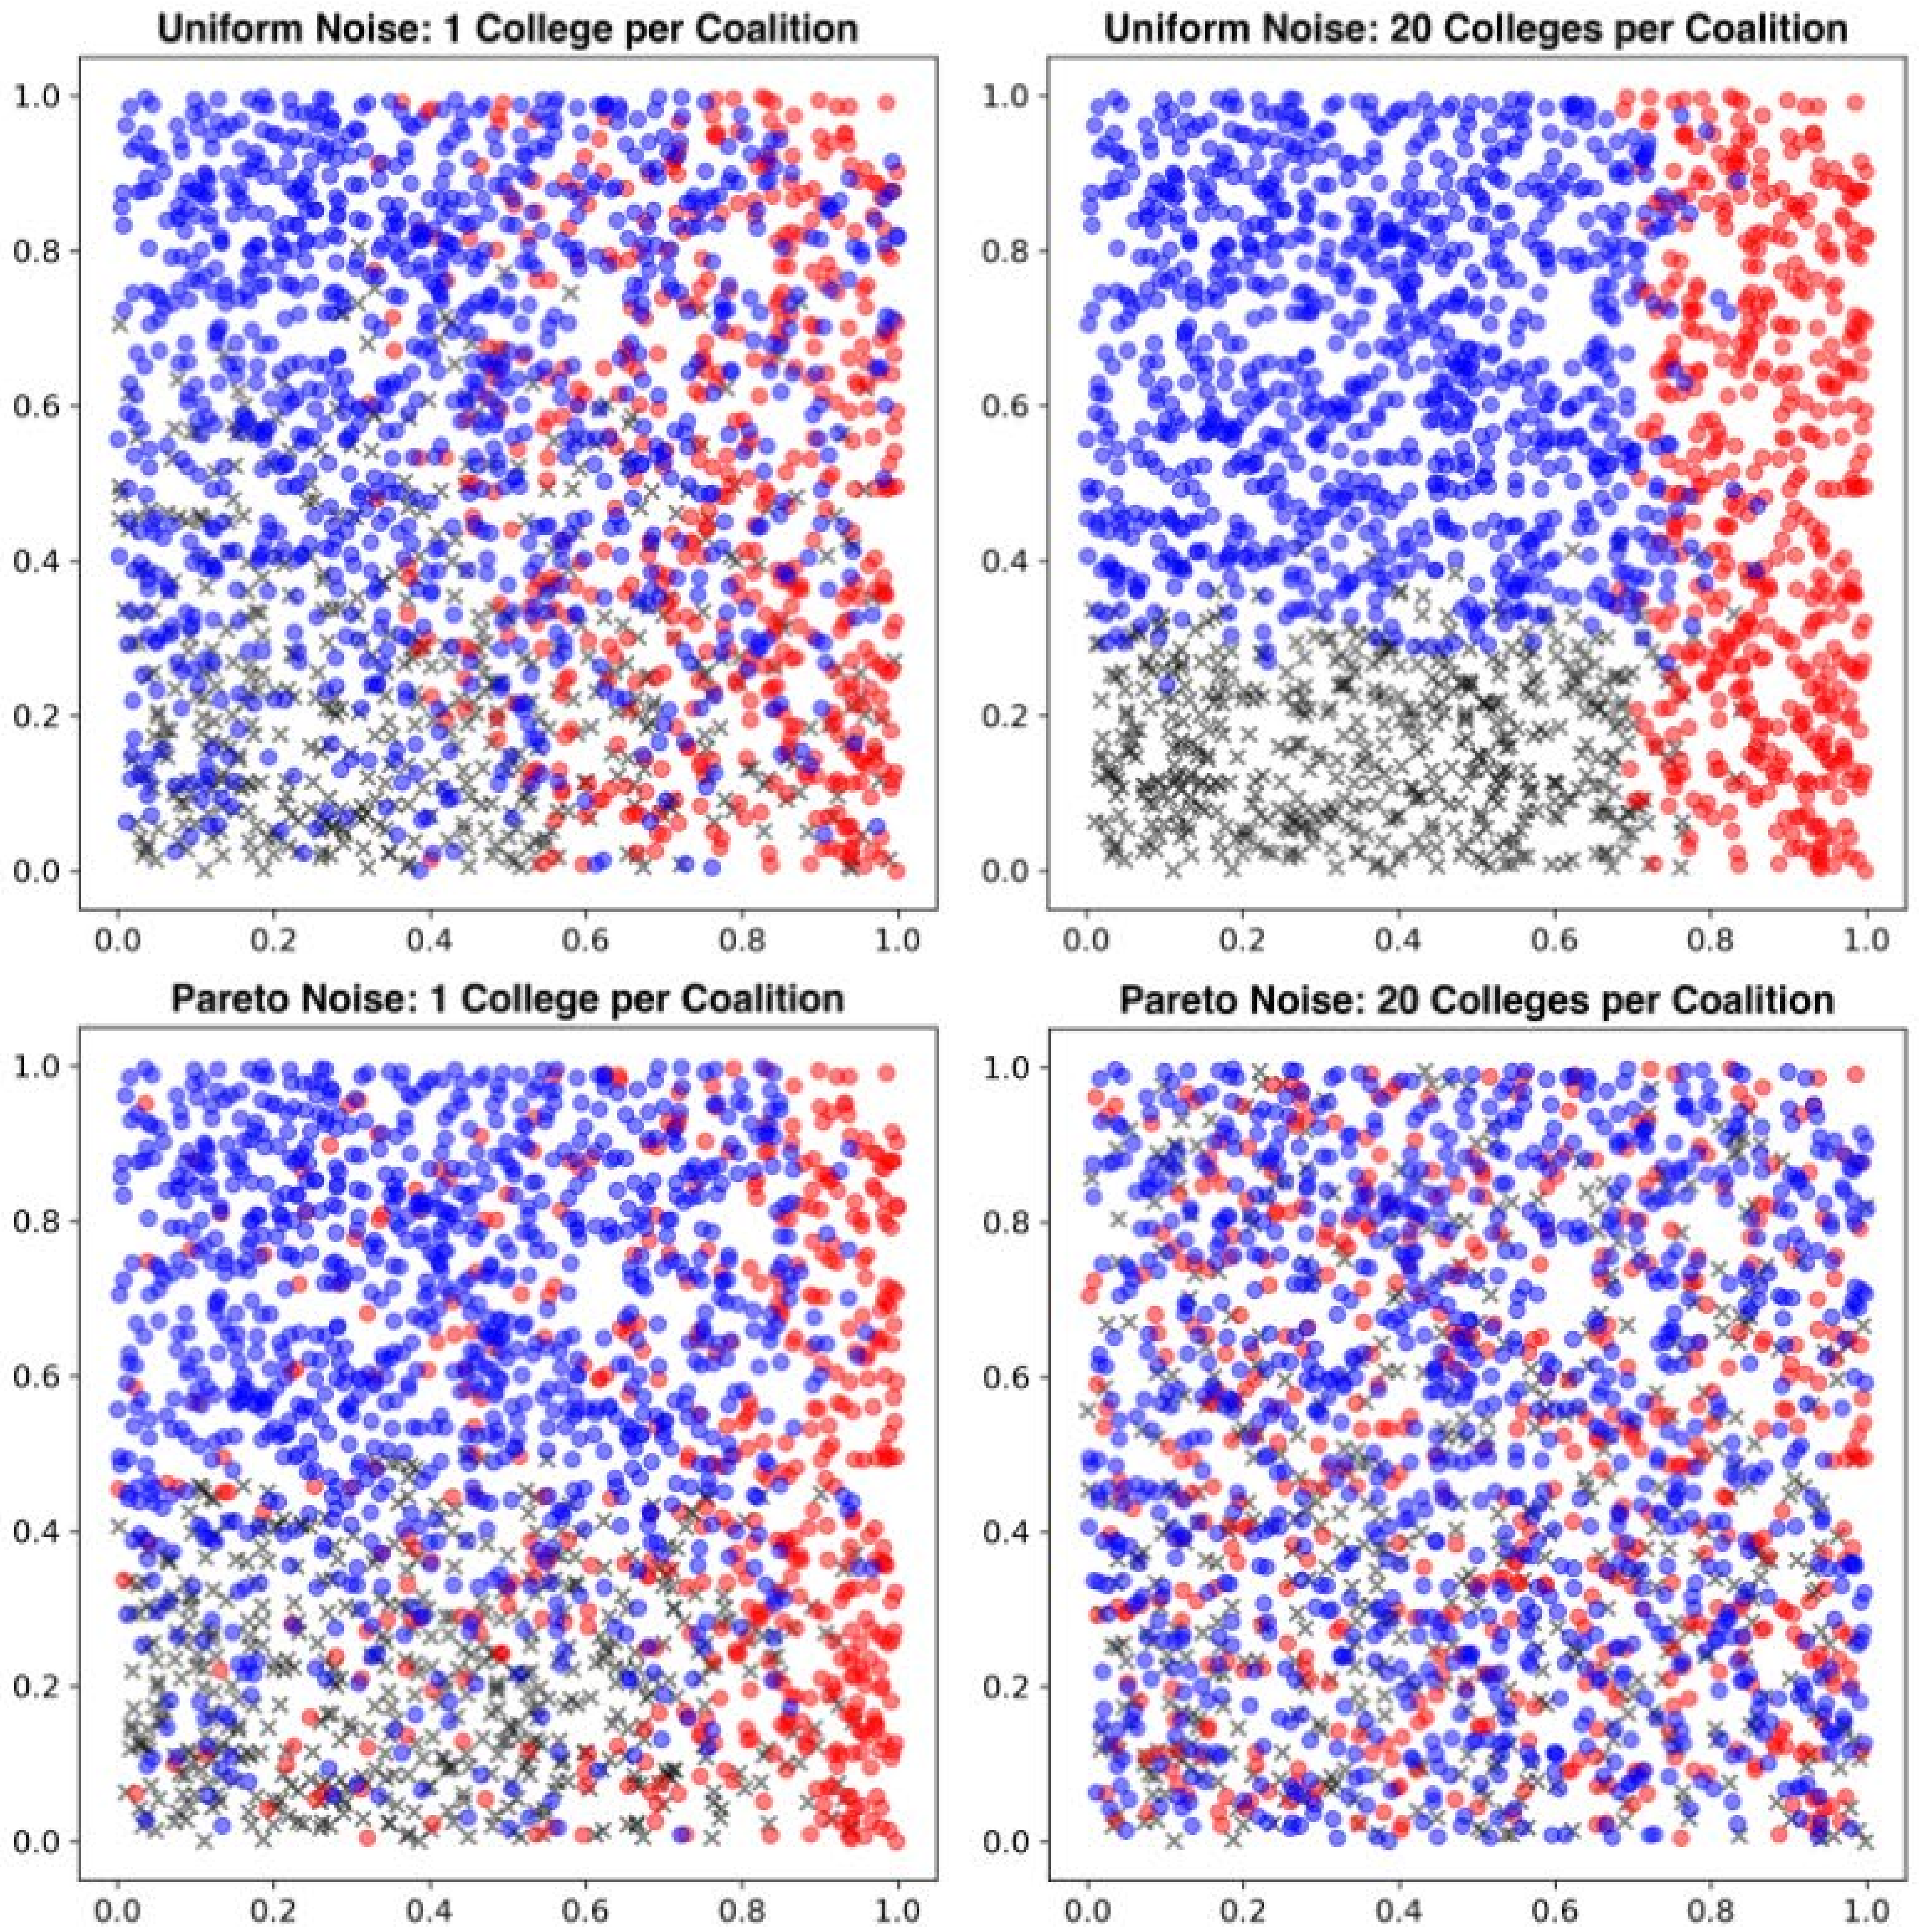

Wisdom and Foolishness of Noisy Matching Markets. Kenny Peng and **NG**

Theorem 2: Applicant welfare

Overall applicant welfare *higher* under monoculture!

By assumption, *same number of applicants receive a match*

But, conditional on matching, more likely to match with favorite firm!

Welfare also depends on true applicant value

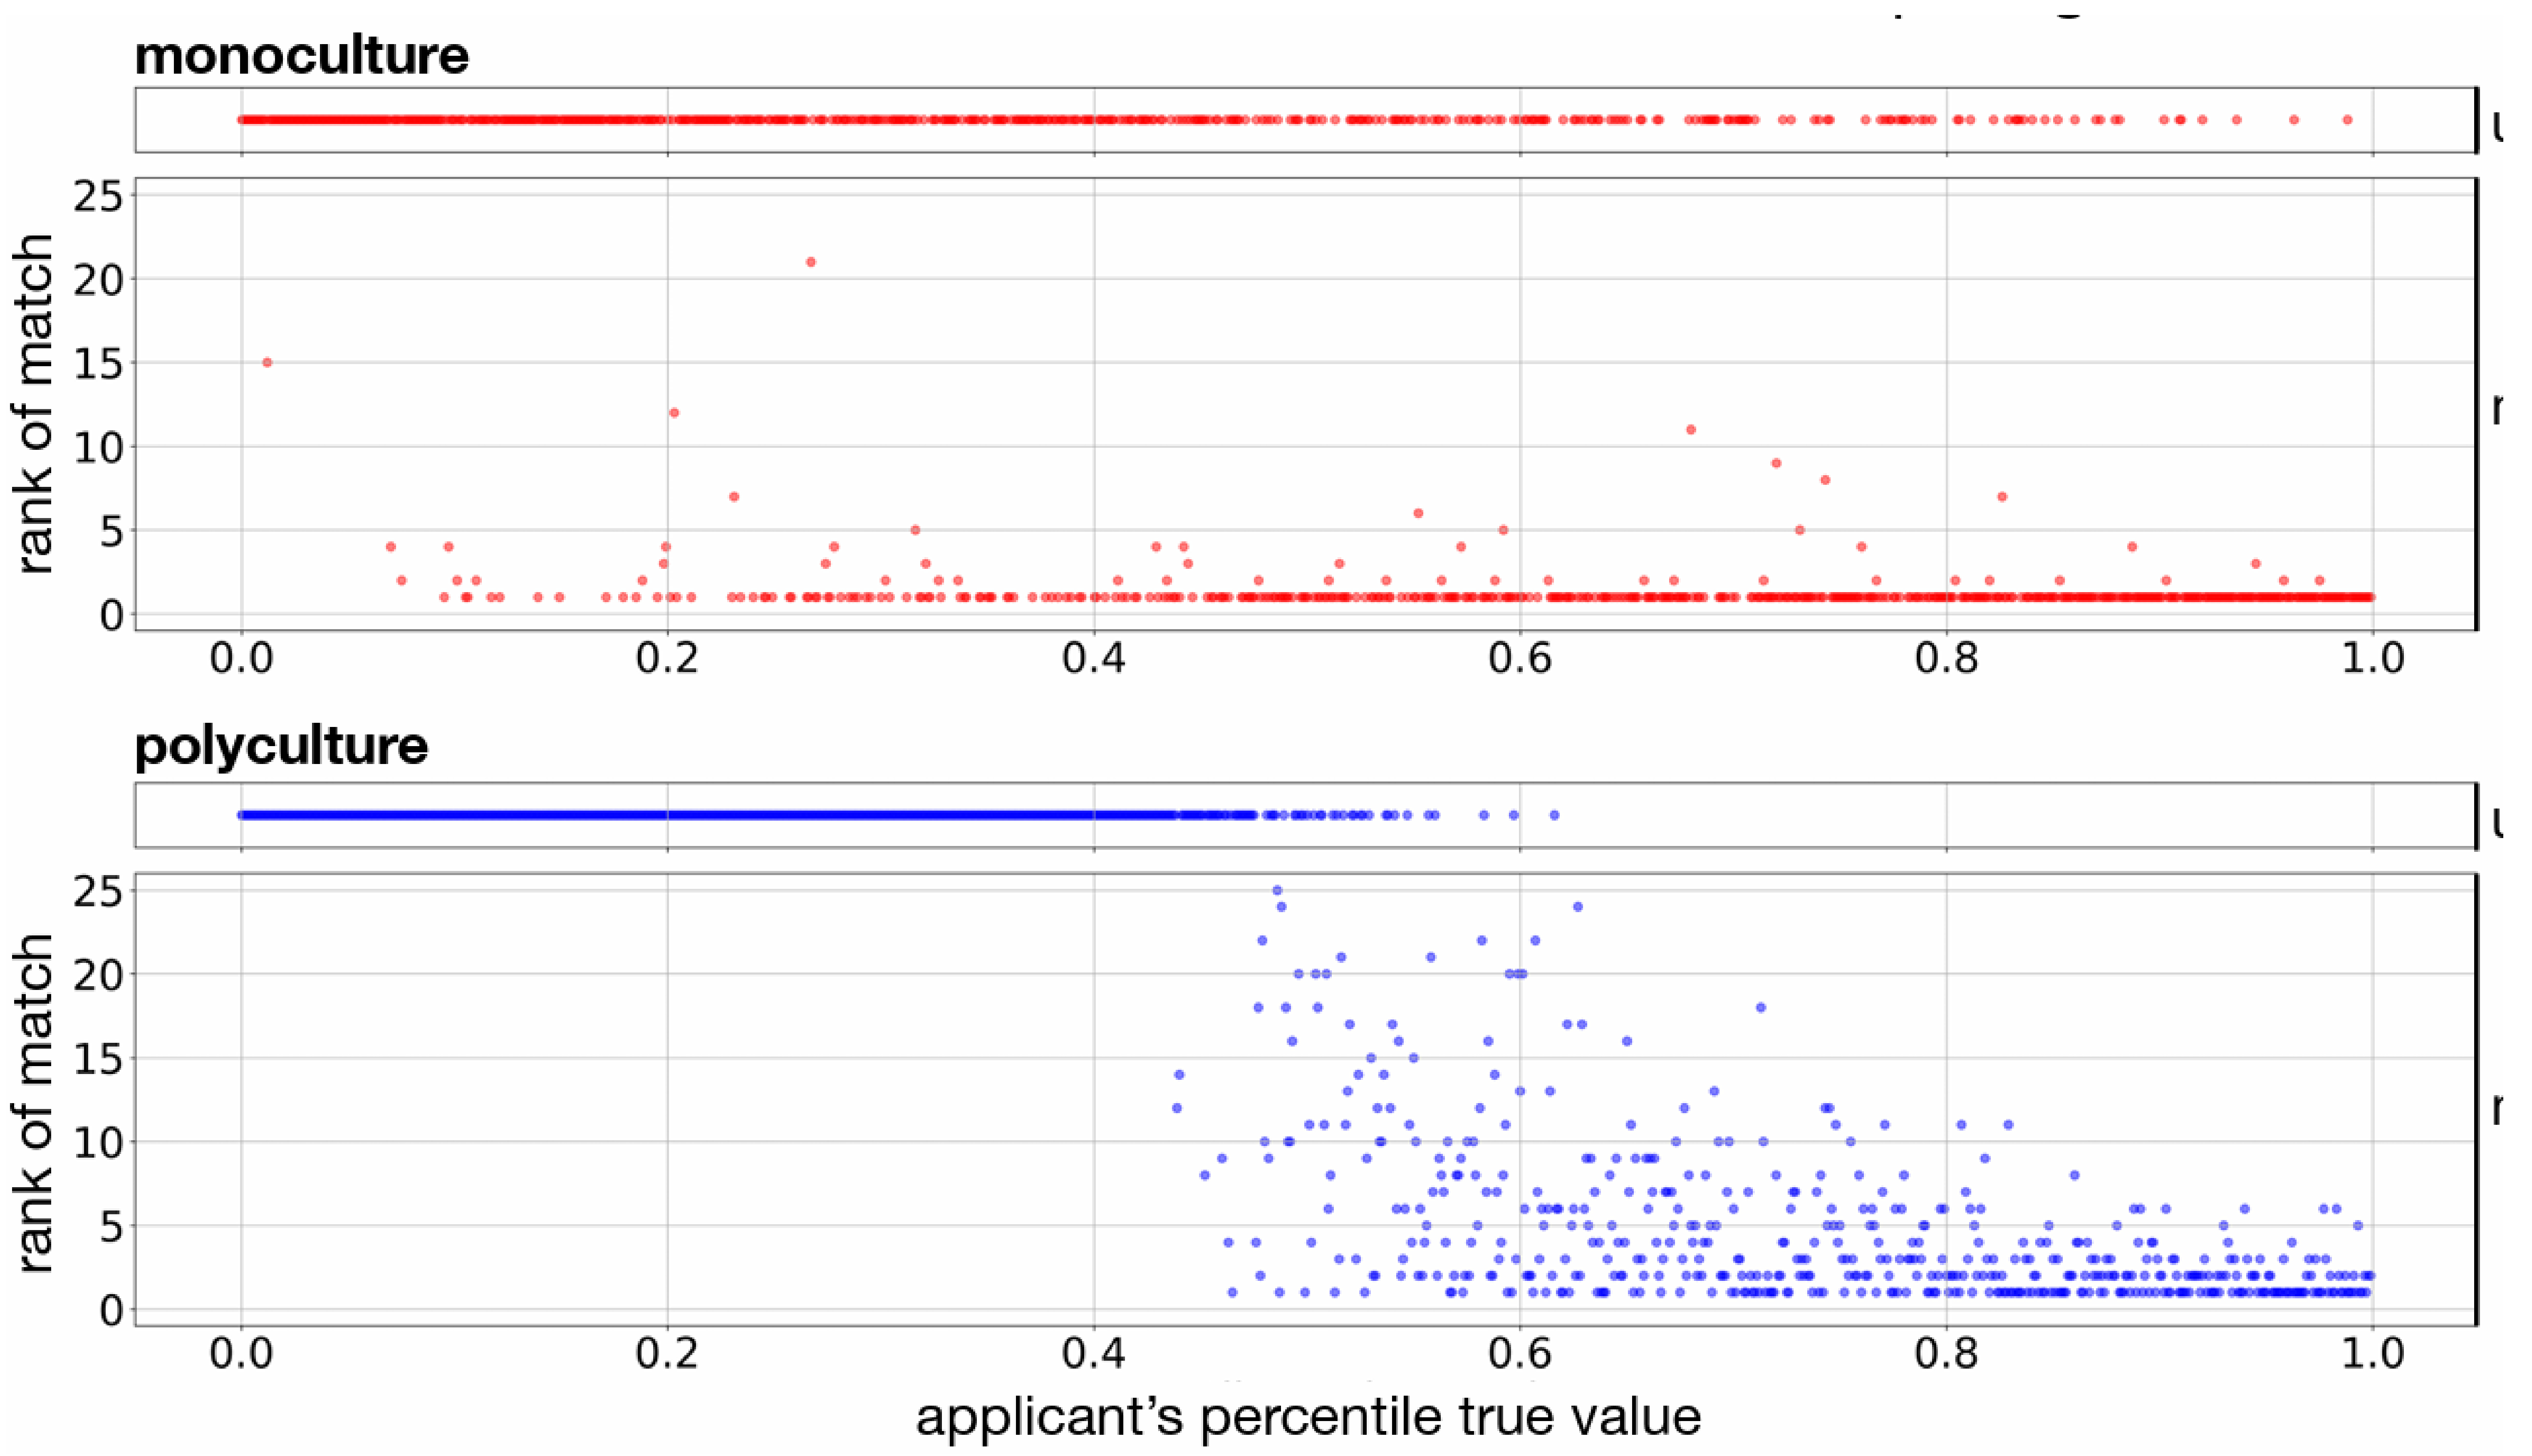

Theorem 3: Disparities

In practice, some people submit many more applications than others

Monoculture is more robust to these disparities!

Why? Polyculture gives *some* people “more lottery tickets”

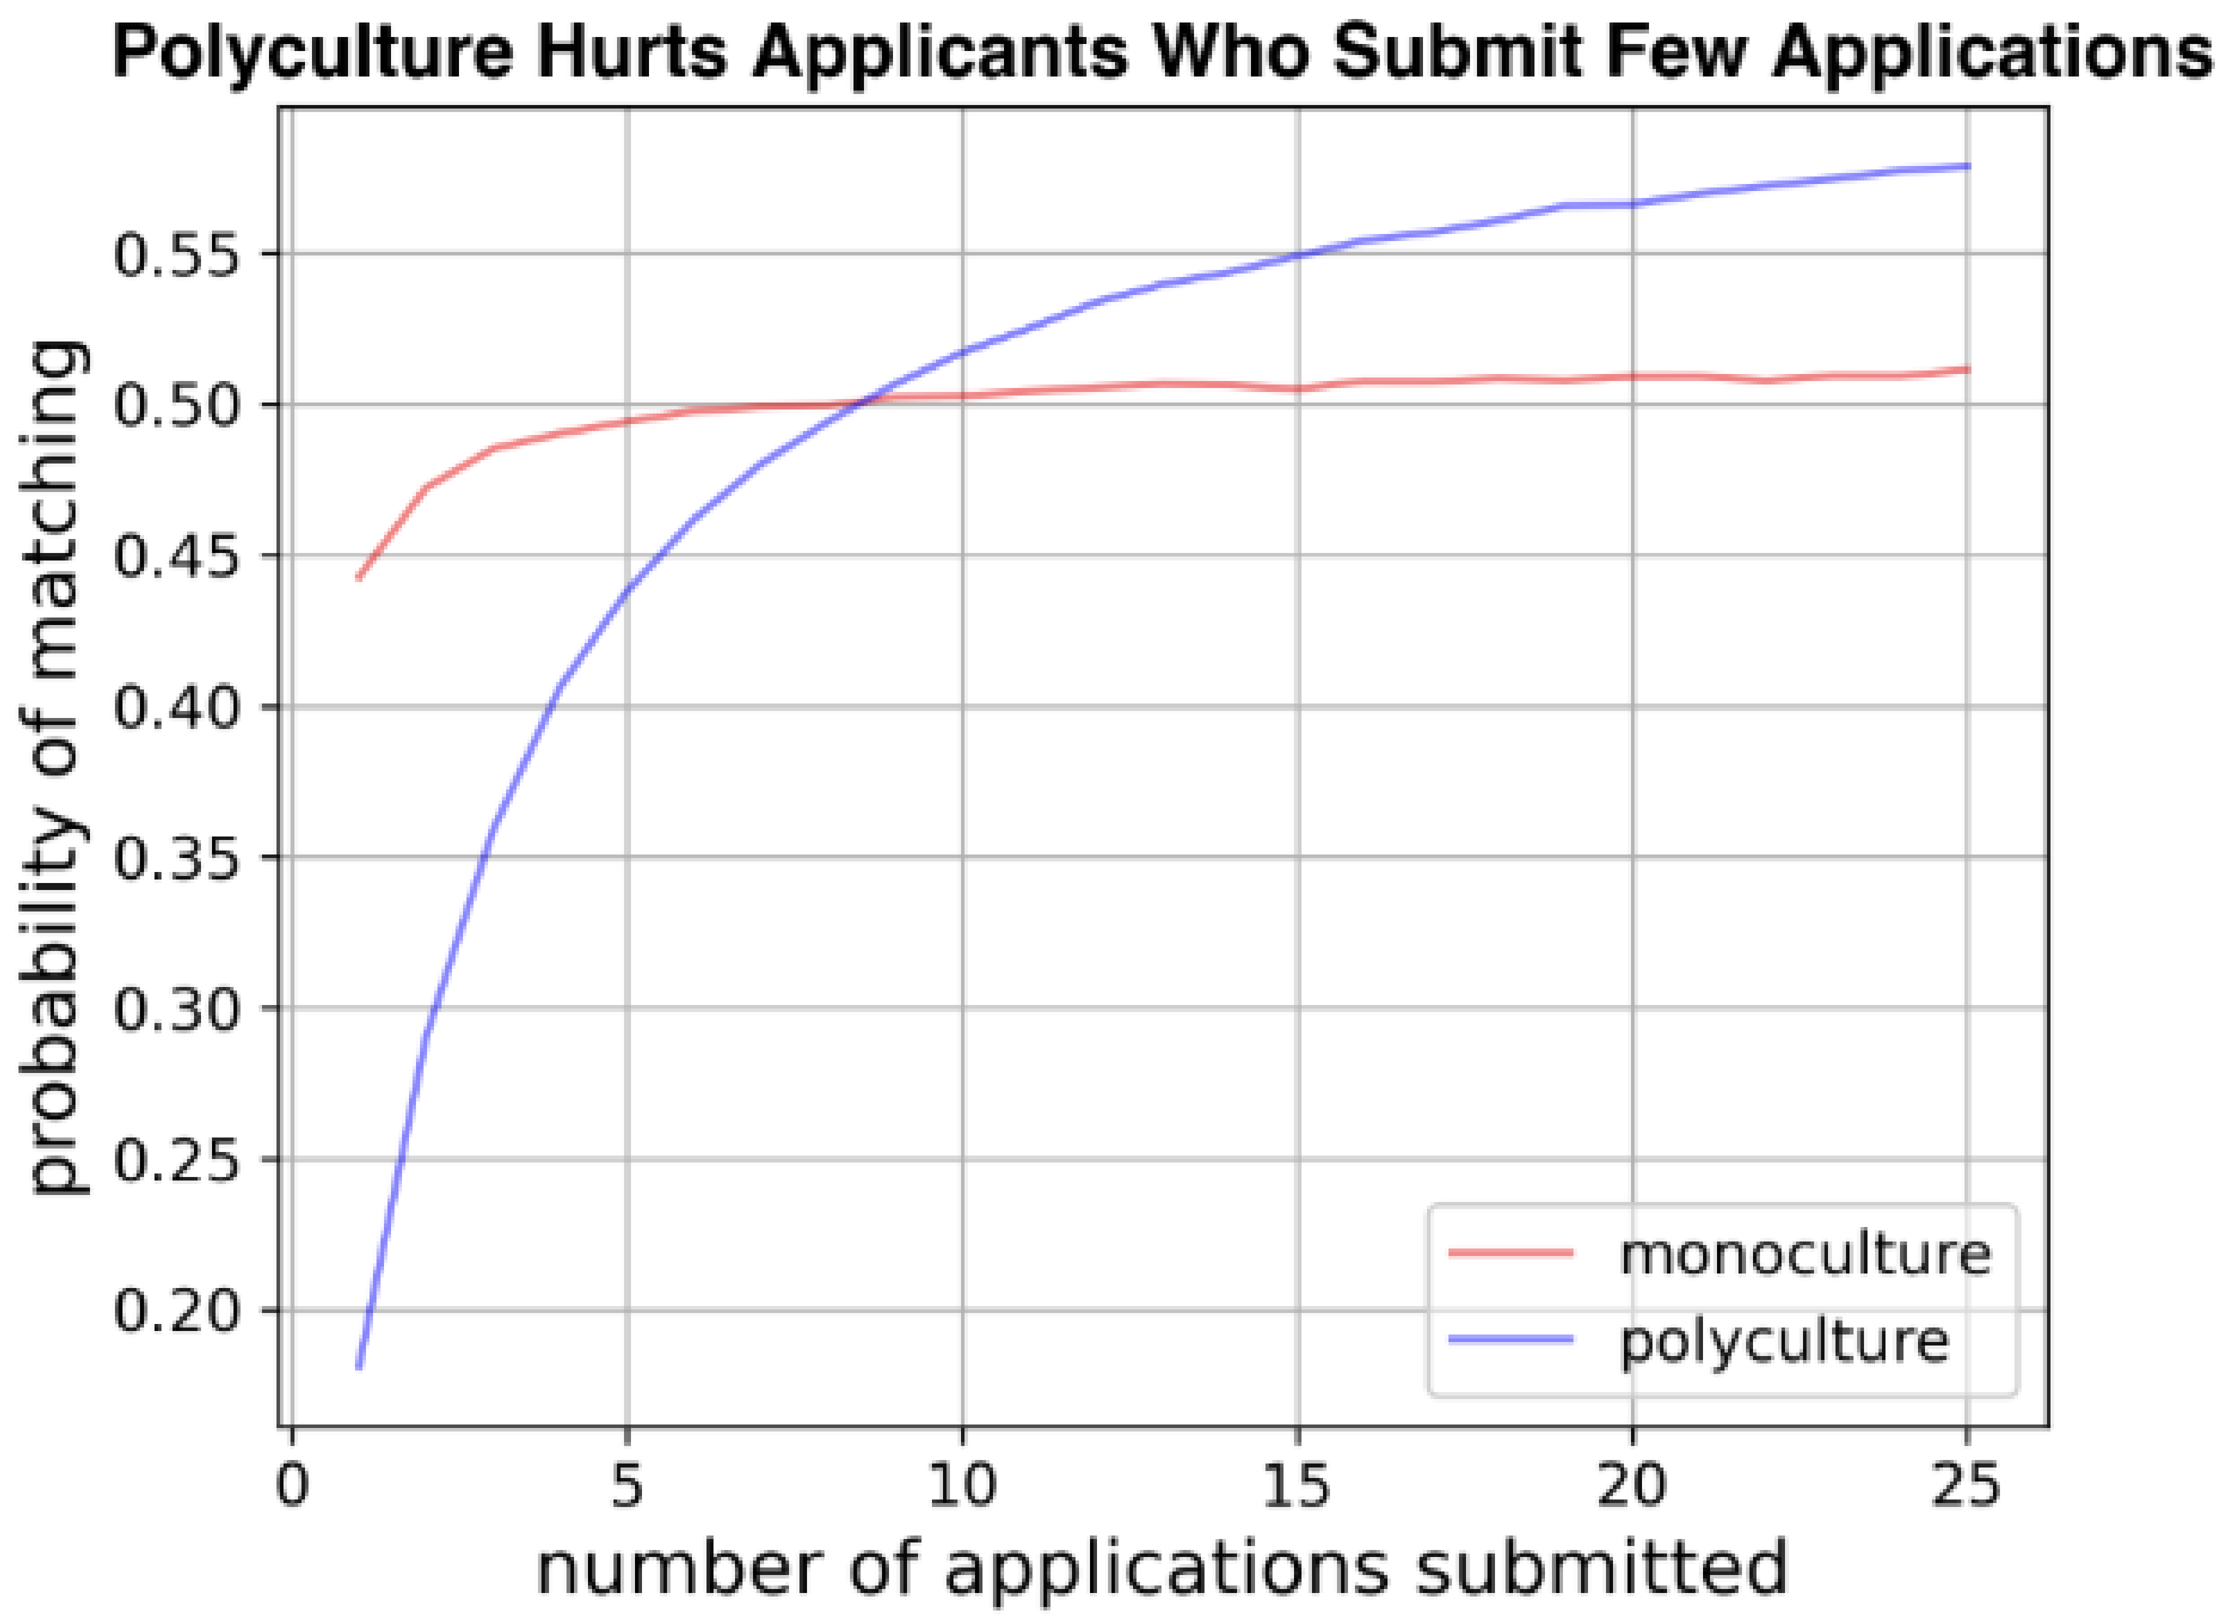

Polyculture

- High firm efficiency (welfare)
- Benefits highest-quality applicants
- Lower "variance" outcomes

Monoculture

- Higher overall applicant welfare
- Benefits lower-quality applicants
- "Fairer" under disparities
- Systematize bias*

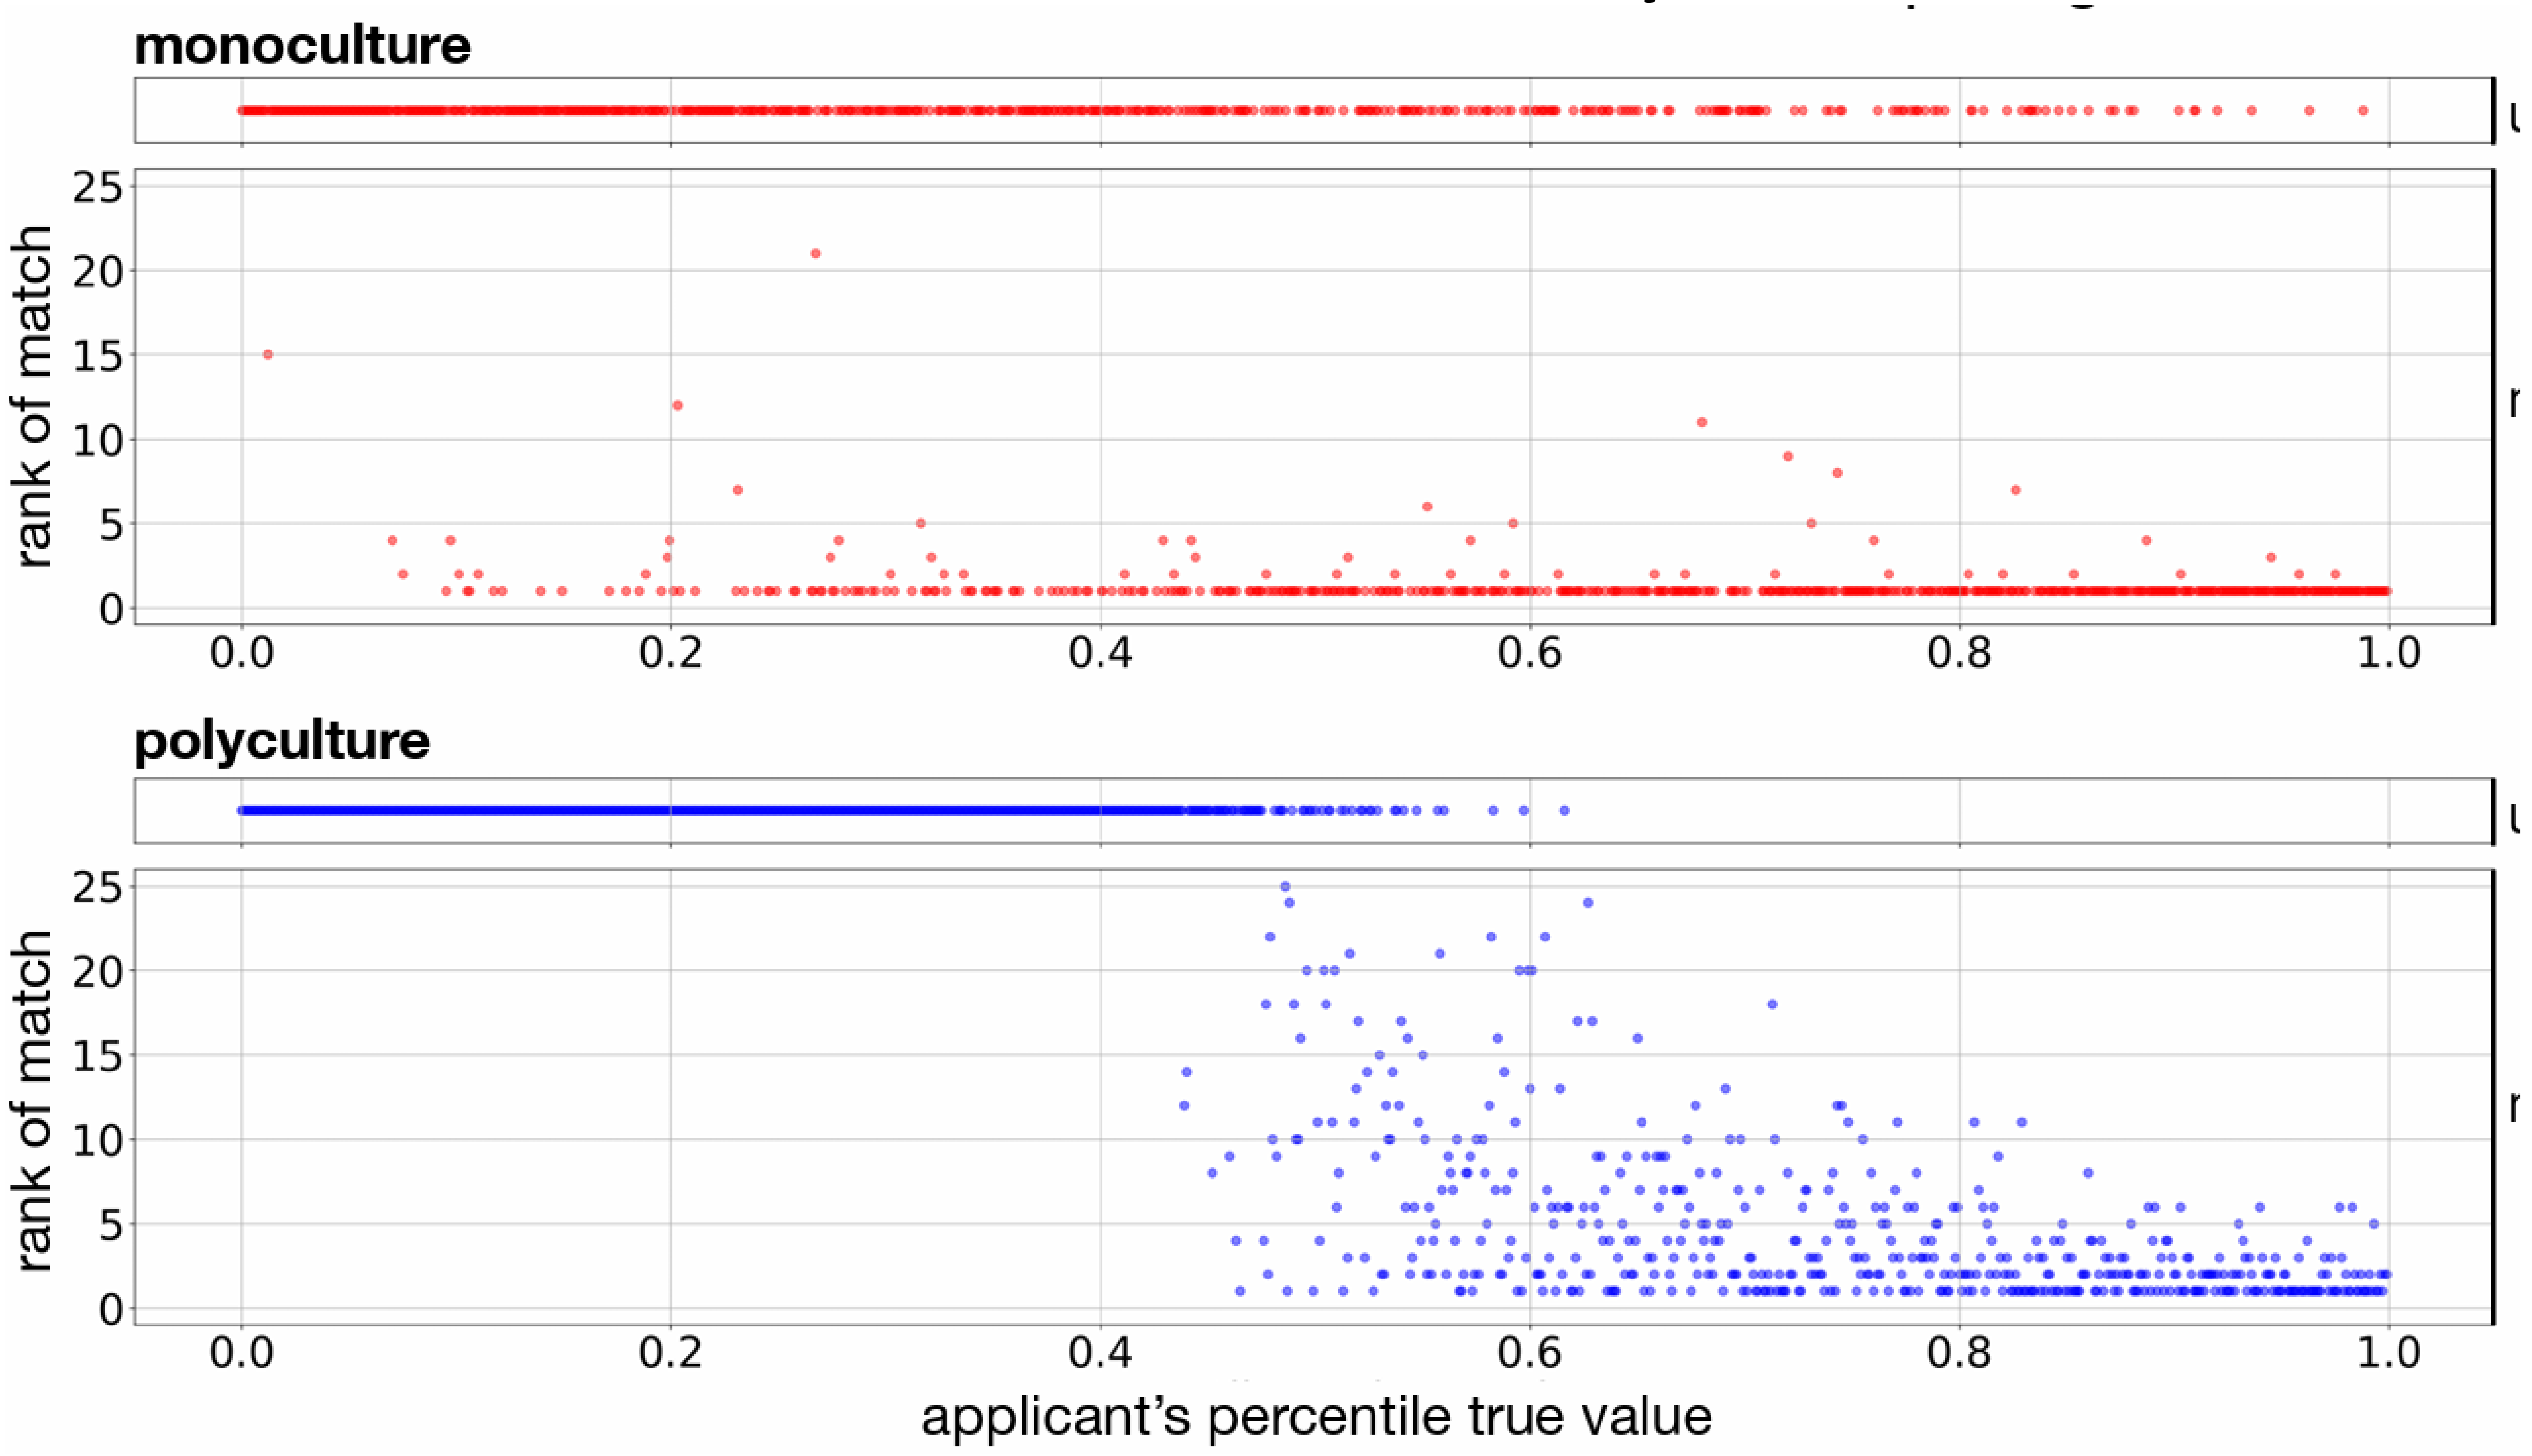

Correlated Errors in LLMs

Are LLMs correlated (even conditional given ground truth)?

- Models by same companies
- Larger models

Yes!

Elliot Kim, Avi Garg, Kenny Peng and **NG**, ICML 2025

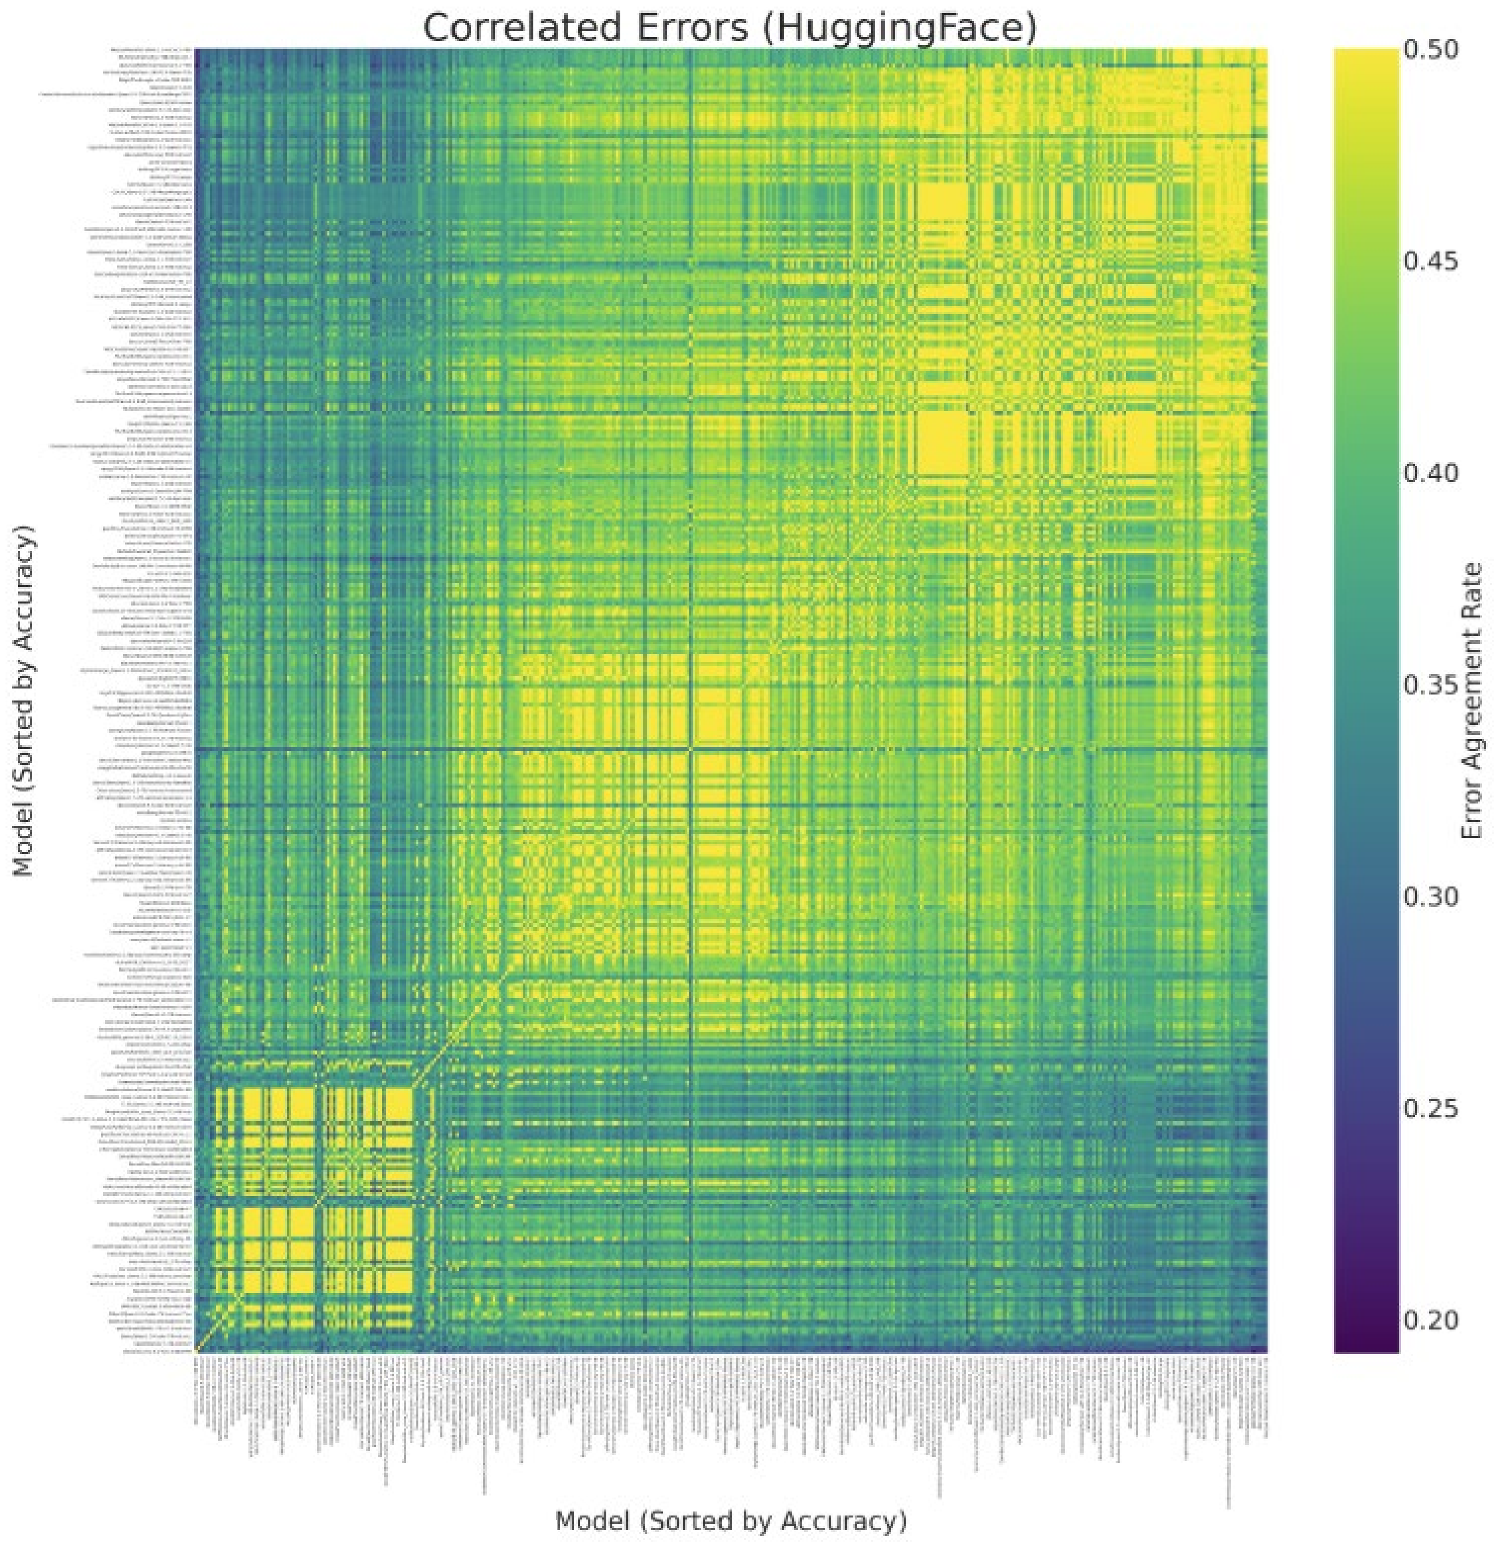

(a) HuggingFace

Table 1. Agreement on Errors

	HUGGINGFACE	HELM	RESUMES
Intercept	0.398**	0.602**	0.653**
	(0.001)	(0.001)	(0.007)
Same Company	0.066**	0.022**	0.021
	(0.003)	(0.005)	(0.012)
Same Architecture	0.076**		
	(0.001)		
Acc. 1	0.014**	0.055**	0.015**
	(0.000)	(0.001)	(0.006)
Acc. 2	0.013**	0.054**	0.028**
	(0.000)	(0.001)	(0.006)
Acc. 1: Acc. 2	0.023**	0.026**	0.043**
	(0.000)	(0.001)	(0.005)
# models	349	71	20
# responses/model	12,032	14,042	1,800
R^2	0.340	0.613	0.415

https://networksmarketscornelltech.github.io/Spring2026/1to100/