Regression Equilibrium

Omer Ben-Porat, Technion

Joint work with Moshe Tennenholtz

Prediction

- Prediction algorithms are extensively studied in the ML literature.
- For commercial companies - another tool that can be exploited to increase revenue.
- Current work assume that prediction is done in isolation!
- Do not address market competition.

Motivation

- Users are interested in the selling values of their apartments.
- Alice offers free prediction services for this purpose on her website.
- Assumption: predictions do not affect the prices.
- After an apartment is sold, its true value is revealed.
- Satisfied users translate to revenue for Alice (user traffic, future services, recommendation etc.).

Illustration (1)

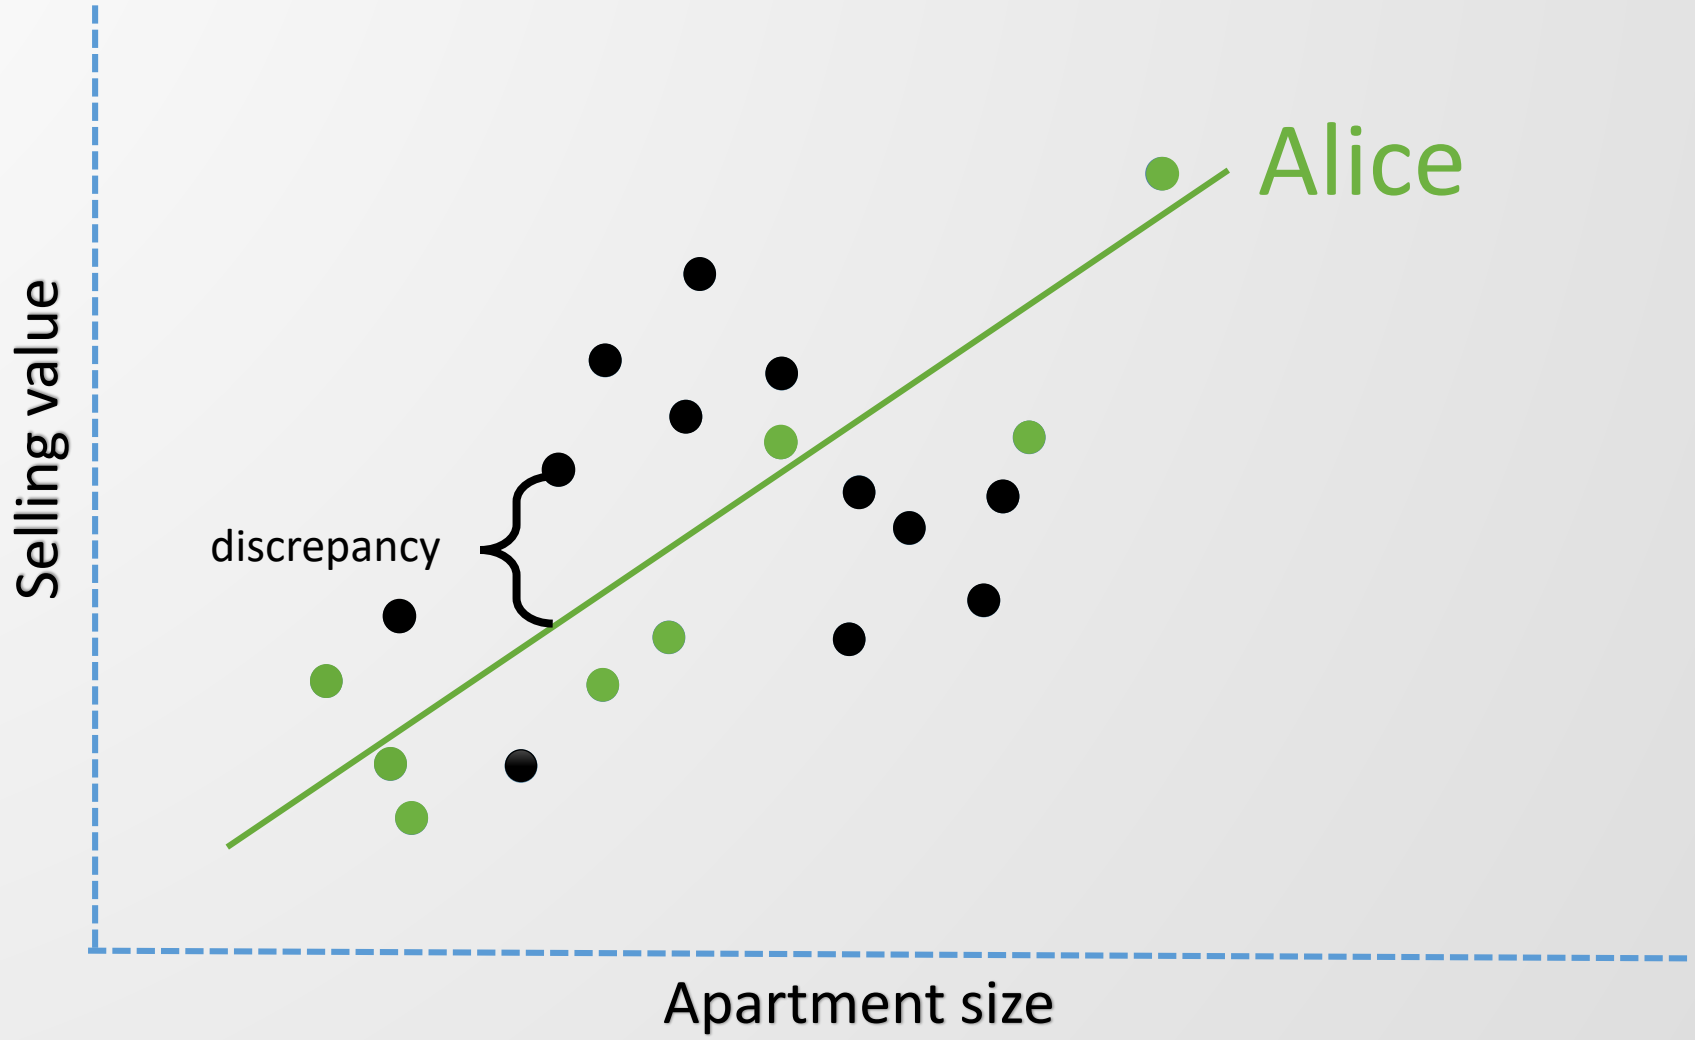

Illustration (2)

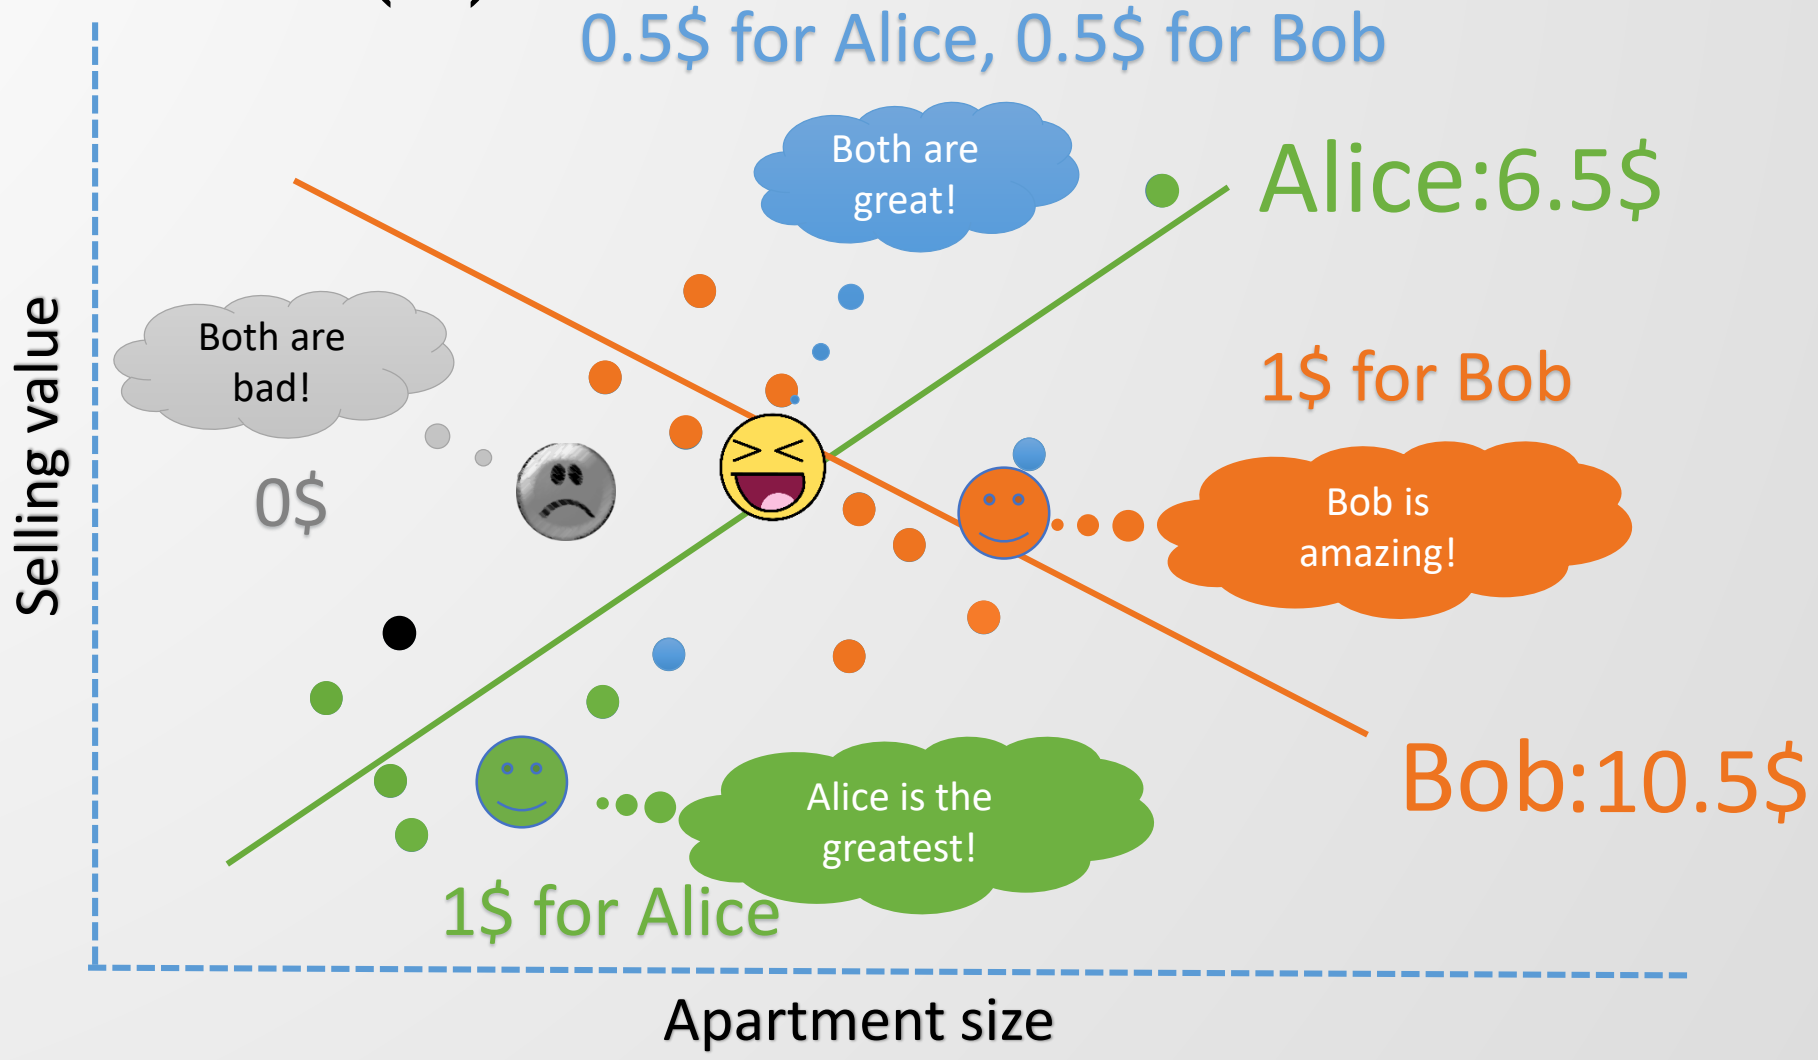

Illustration (2)

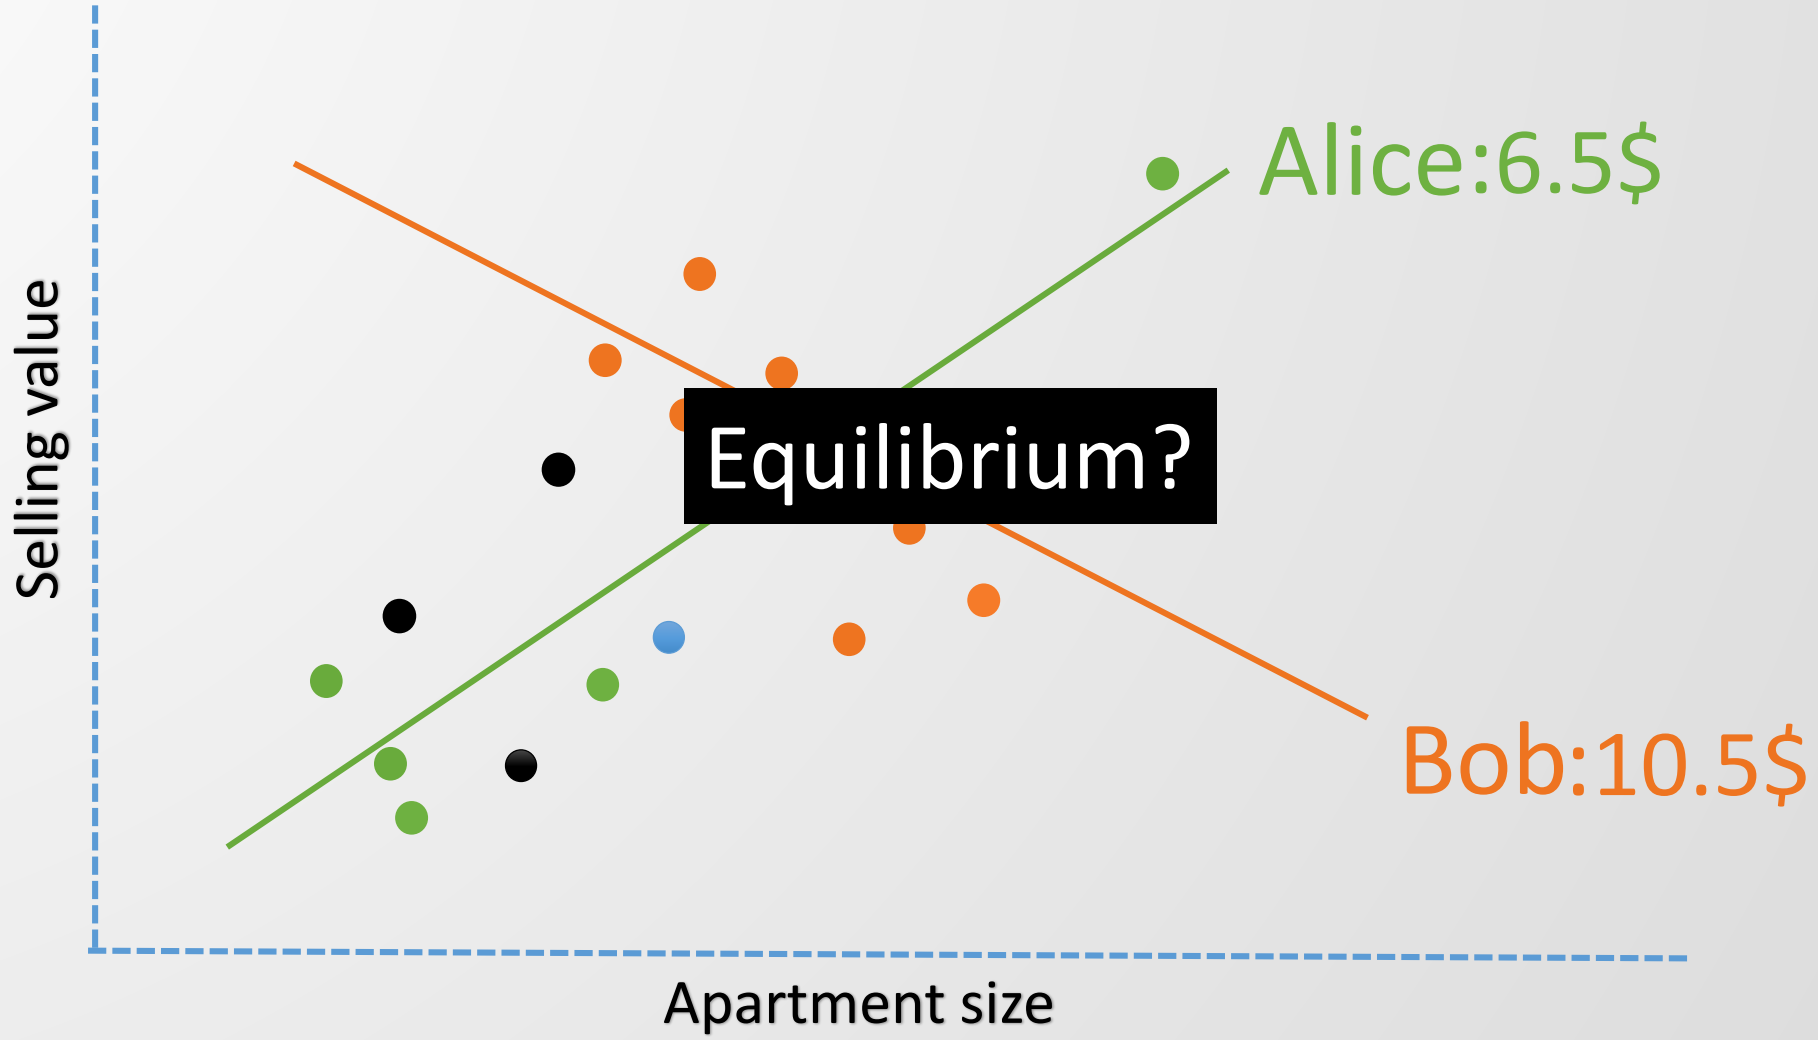

Informal Model (PAC-like)

- A distribution over instances, labels and thresholds.
- N players, play predictive function $x \mapsto y$ (full generality)
- A point (x, y, t) is satisfied with a prediction $\hat{y} = \hat{y}(x)$ if
$$|y - \hat{y}| \leq t.$$
- Several satisfying predictions → u.a.r.
- Player payoff: the expected number of points she satisfies.
- The underlying distribution is unknown.

Question: How to find an approximate PNE on the distribution?

PAC Learning

Results (1)

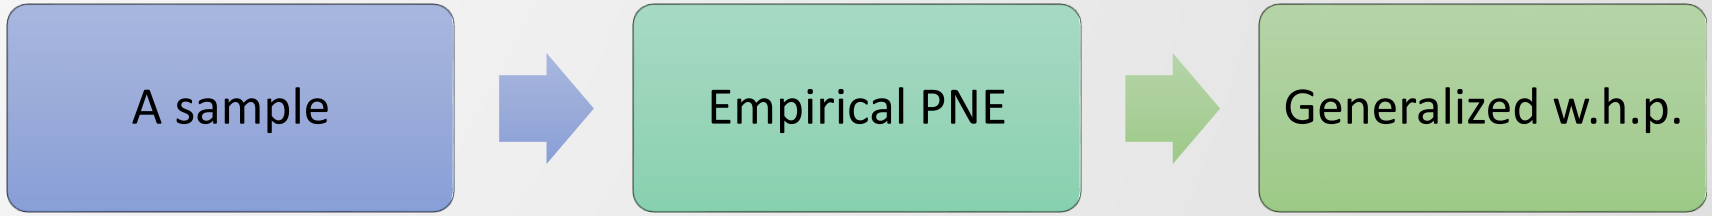

- Proposition: Every empirical game possesses at least one PNE.

Results (2)

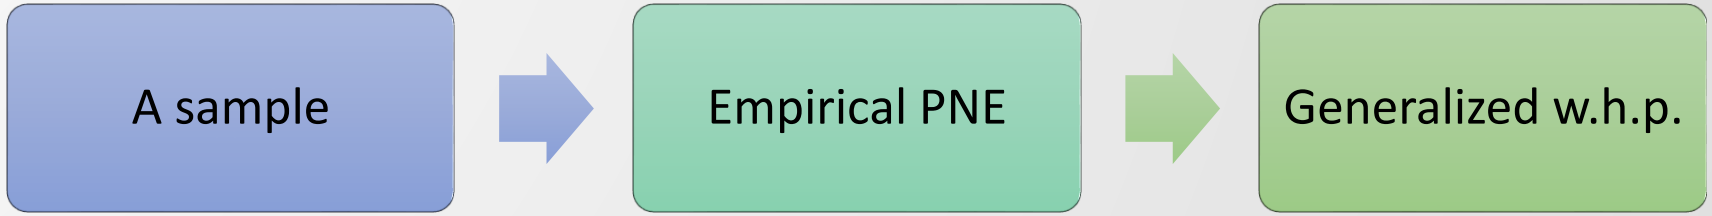

- Proposition: After $O(mN \log N)$ iterations of any better response dynamics, an empirical PNE is obtained.

Results (3)

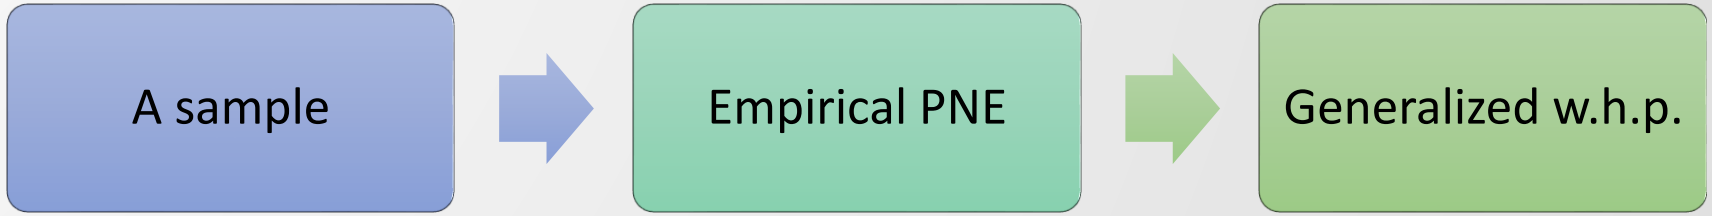

- Lemma: Given a sample of size $poly\left({}^{1}/_{\epsilon}, N, \log {}^{1}/_{\delta}, \sum_{i=1}^{N} d_i\right)$, any player's payoff under any profile is not too distant from its empirical counterpart, w.h.p.

Meta Algorithm

1. Set $m = poly\left(\frac{1}{\epsilon}, N, \log\left(\frac{1}{\delta}\right), \sum_{i=1}^{N} d_i\right)$.
2. Sample $\mathcal{S}$ from $\mathcal{D}^m$.
3. Execute any better-response dynamics on the empirical game until convergence, and obtain a strategy profile $\boldsymbol{h}$.
4. Return $\boldsymbol{h}$.

➢ Theorem: The algorithm returns an ϵ −PNE w.p. of at least $1 - \delta$.

➢ "Meta"?

- Linear best-response oracle, based on BP and Tennenholtz, NIPS 2017.

Simulations: Two-player games

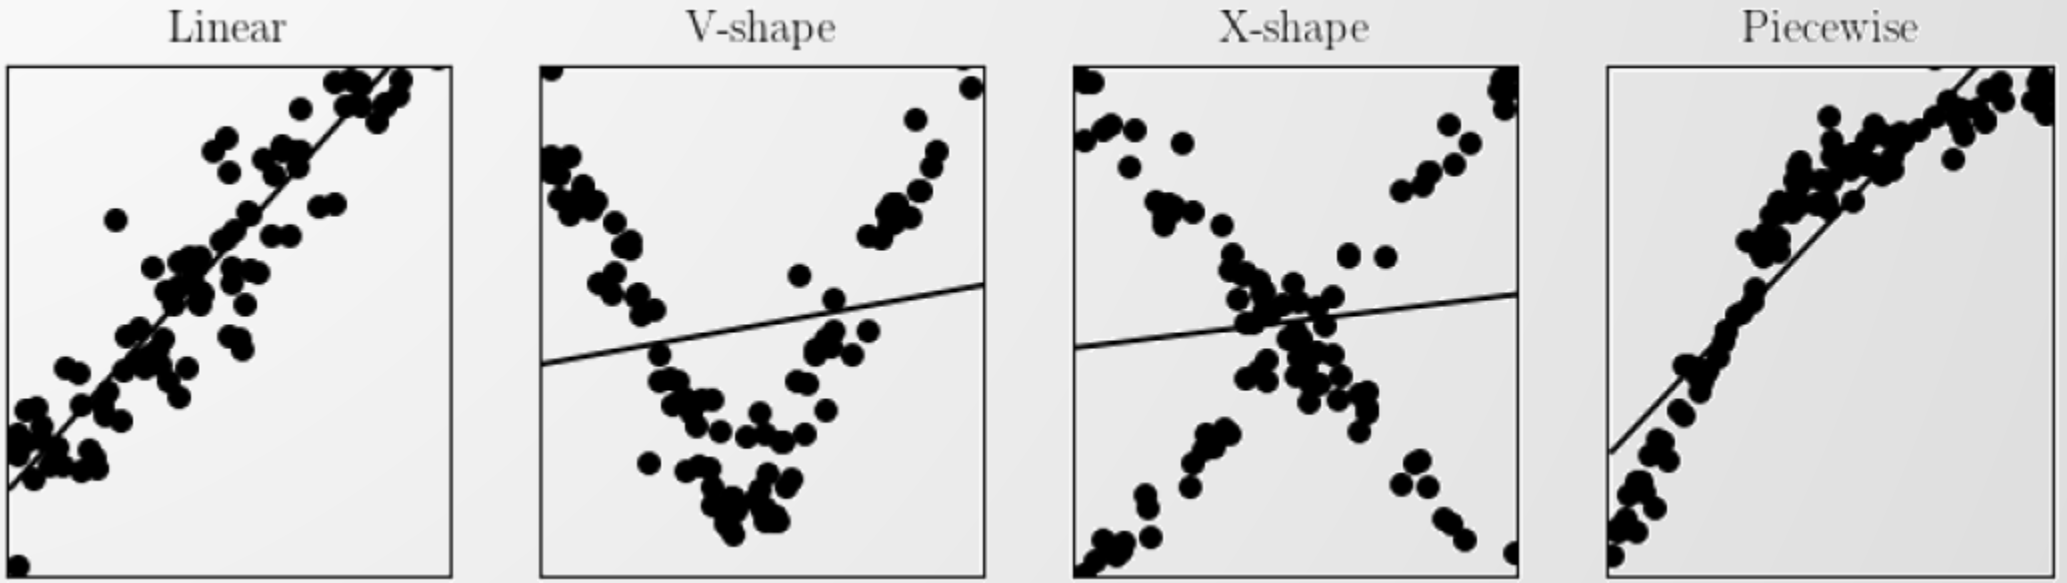

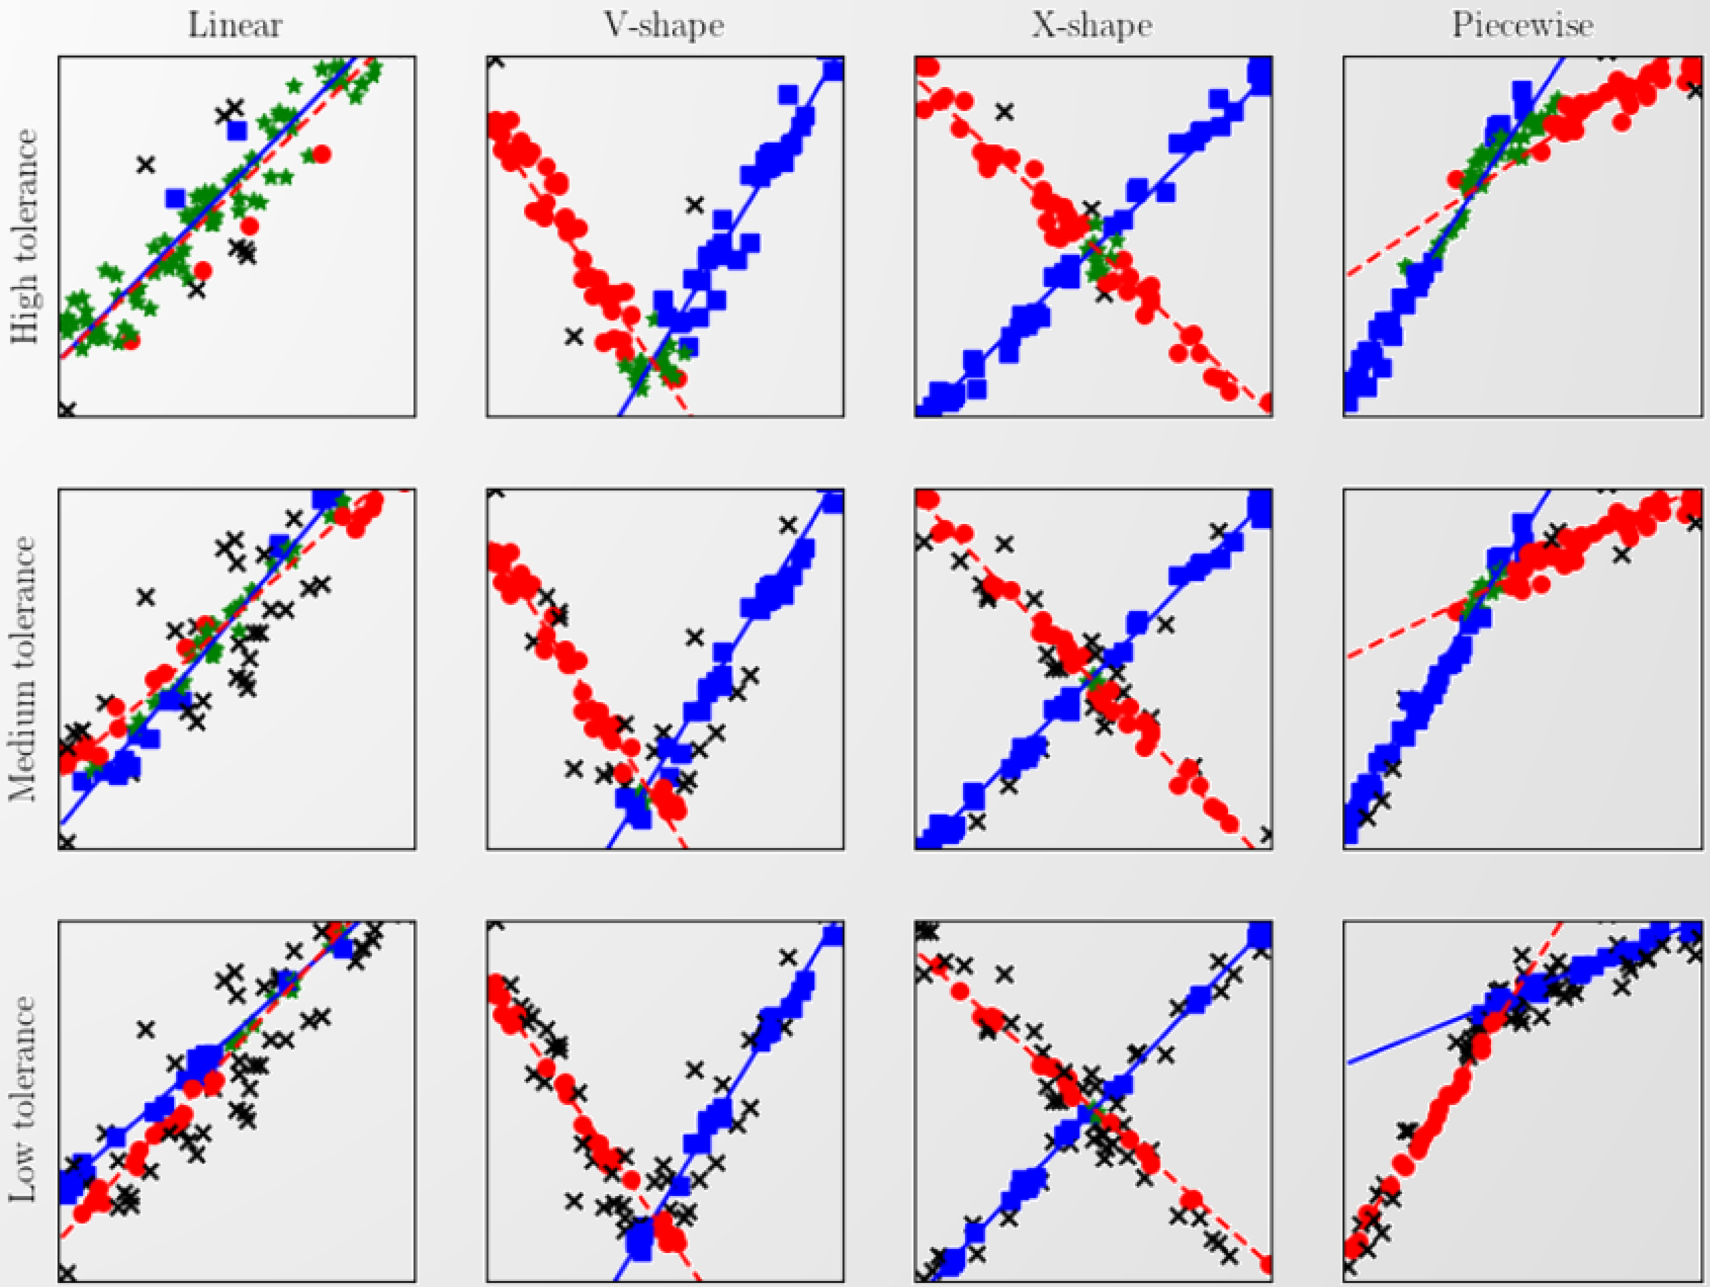
Linear
V-shape
X-shape
Piecewise
High tolerance
Medium tolerance
Low tolerance

Related Work

- Competing optimization algorithms:
  - Dueling Algorithms (Immorlica et al., 2011).
- Competing solution concepts for machine learning tasks:
  - Competing schedulers (Ashlagi et al., 2010), (Ashlagi et al., 2013).
  - Competing bandits (Mansour at el., 2018).
  - Online prediction (Schrijvers and Roughgarden, 2017).
  - Clustering/segmentation – (Hotelling 1929).
- Strategic input:
  - Strategyproof classification/regression. (Dekel et al. 2008 , Meir et al. 2012, Chen et al. 2018).
  - Segmentation (Nissim et al. 2018).
- Prediction with several entities:
  - Collaborative (Blum et al. 2017), competitive (BP and Tennenholtz 2017).

Extension: Direct Attraction

- Variant: each user grants 1$ to the player with the most accurate prediction (breaking ties uniformly).
  - In the spirit of Dueling Algorithms (Immorlica et al., 2011).
- Empirical PNE may not exist!

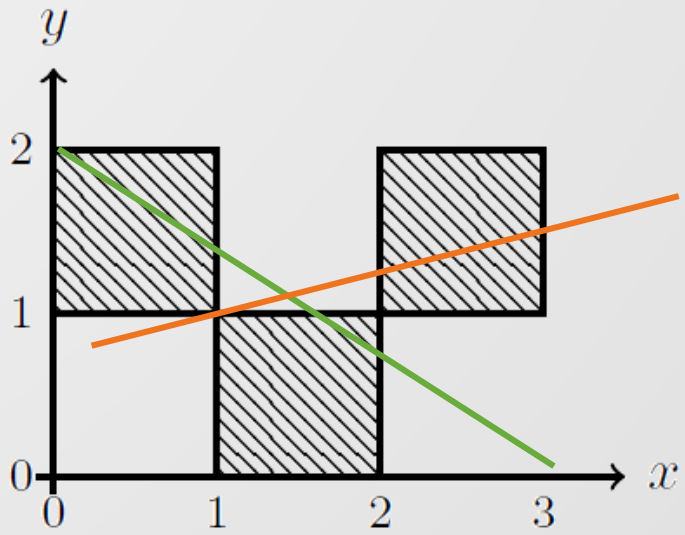

Credit: Yakov Babichenko

Future work

- Best response oracles.
- Variant: direct attraction
  - In the spirit of Dueling Algorithms (Immorlica et al., 2011).
- Different monitoring: players are not aware/partially aware of the strategies of the other players.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement n° 740435).