On the Convergence of No-Regret Dynamics in Information Retrieval Games with Proportional Ranking Functions

Omer Madmon, Idan Pipano, Itamar Reinman, Moshe Tennenholtz

Technion - Israel Institute of Technology

Background

- Publishers who publish their content on the web act *strategically* to maximize their exposure.
- Content modification dynamics often leads to *fluctuations*, that harm both the platform and the end users.

Research Question: How can the platform guarantee that publishers will always converge to a stable state?

The Publishers' Game

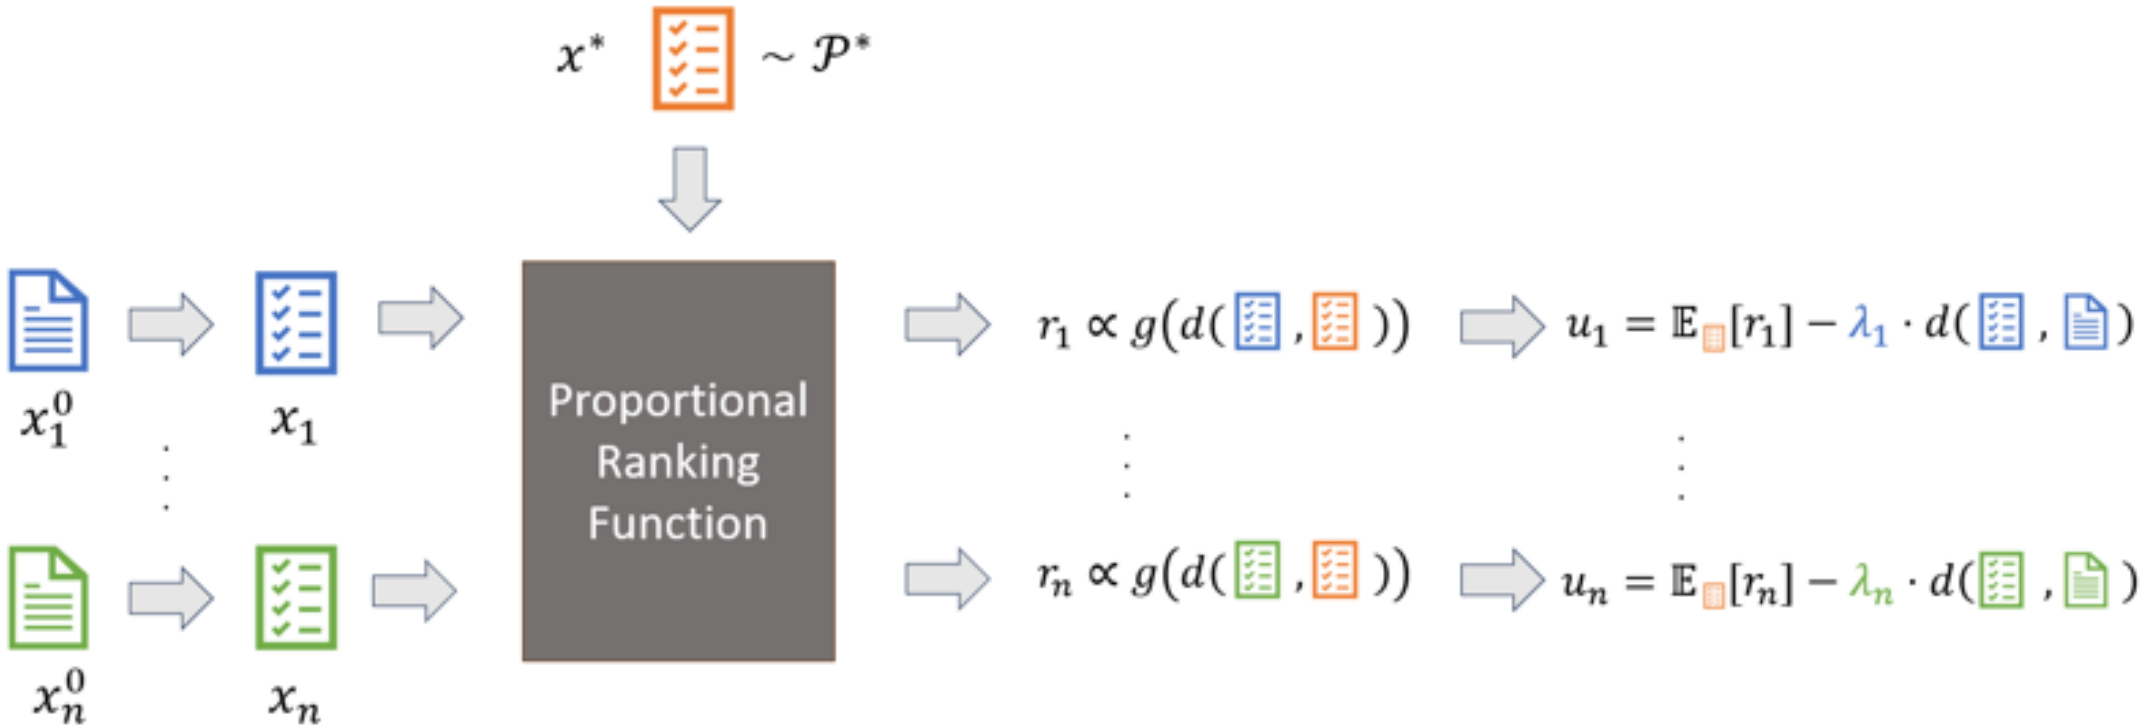

The Publishers' Game

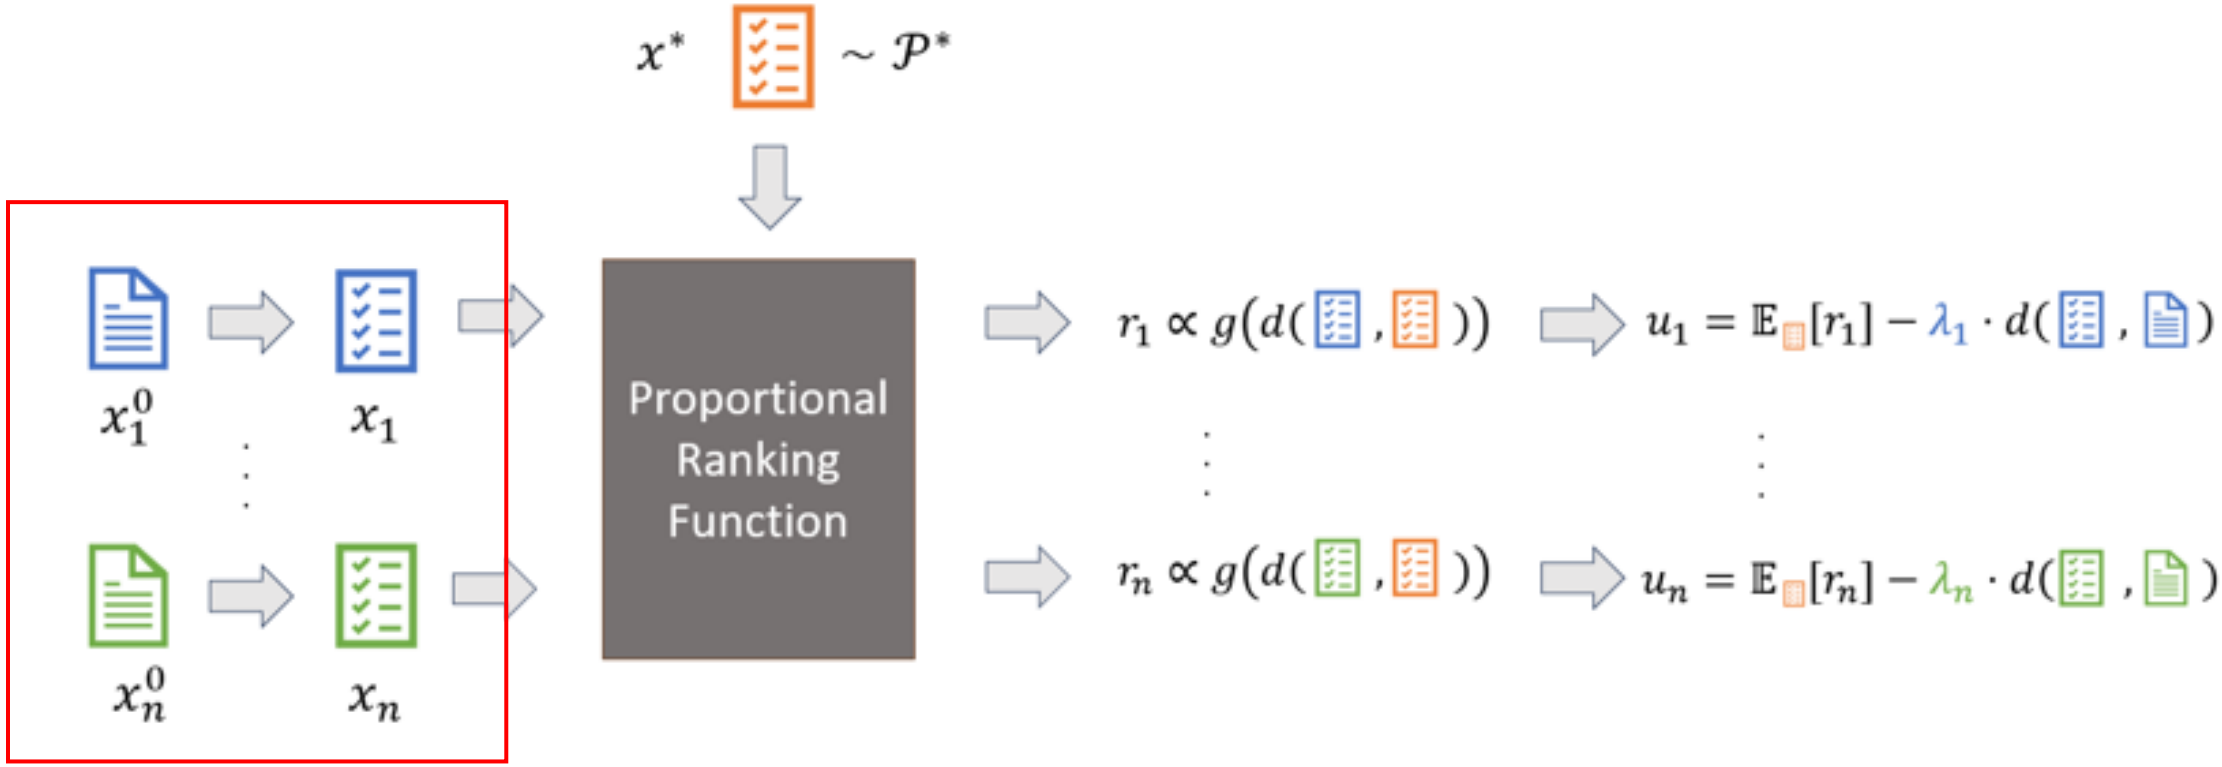

The Publishers' Game

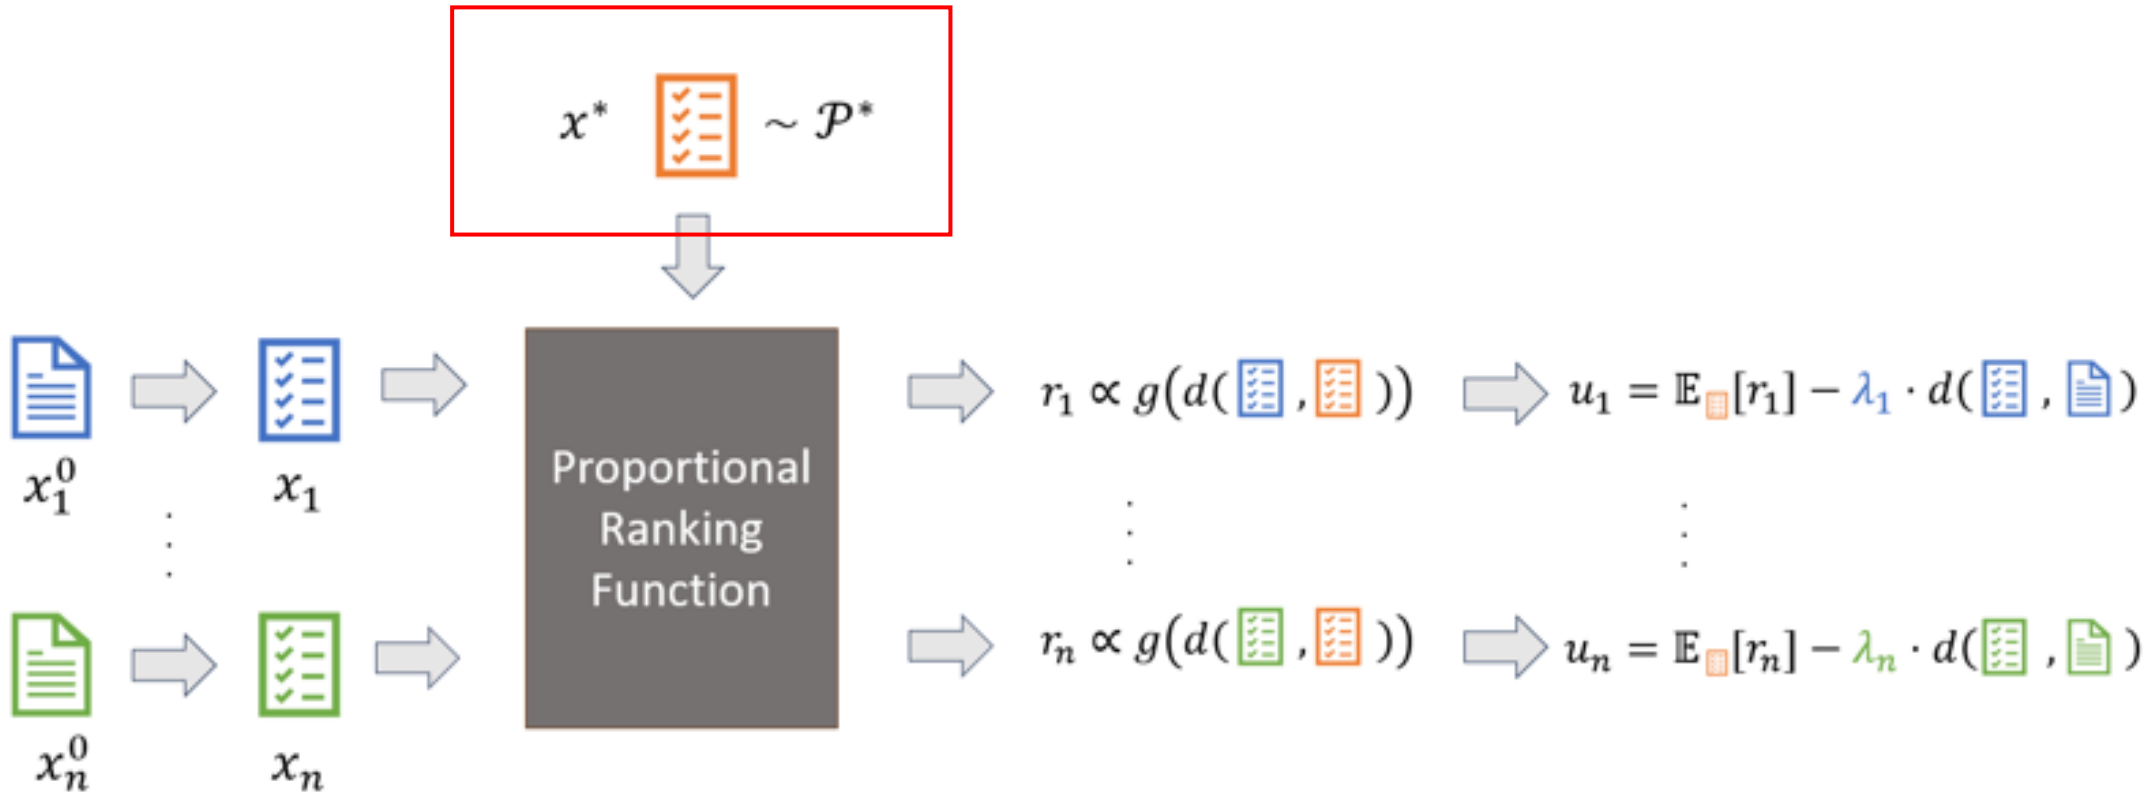

The Publishers' Game

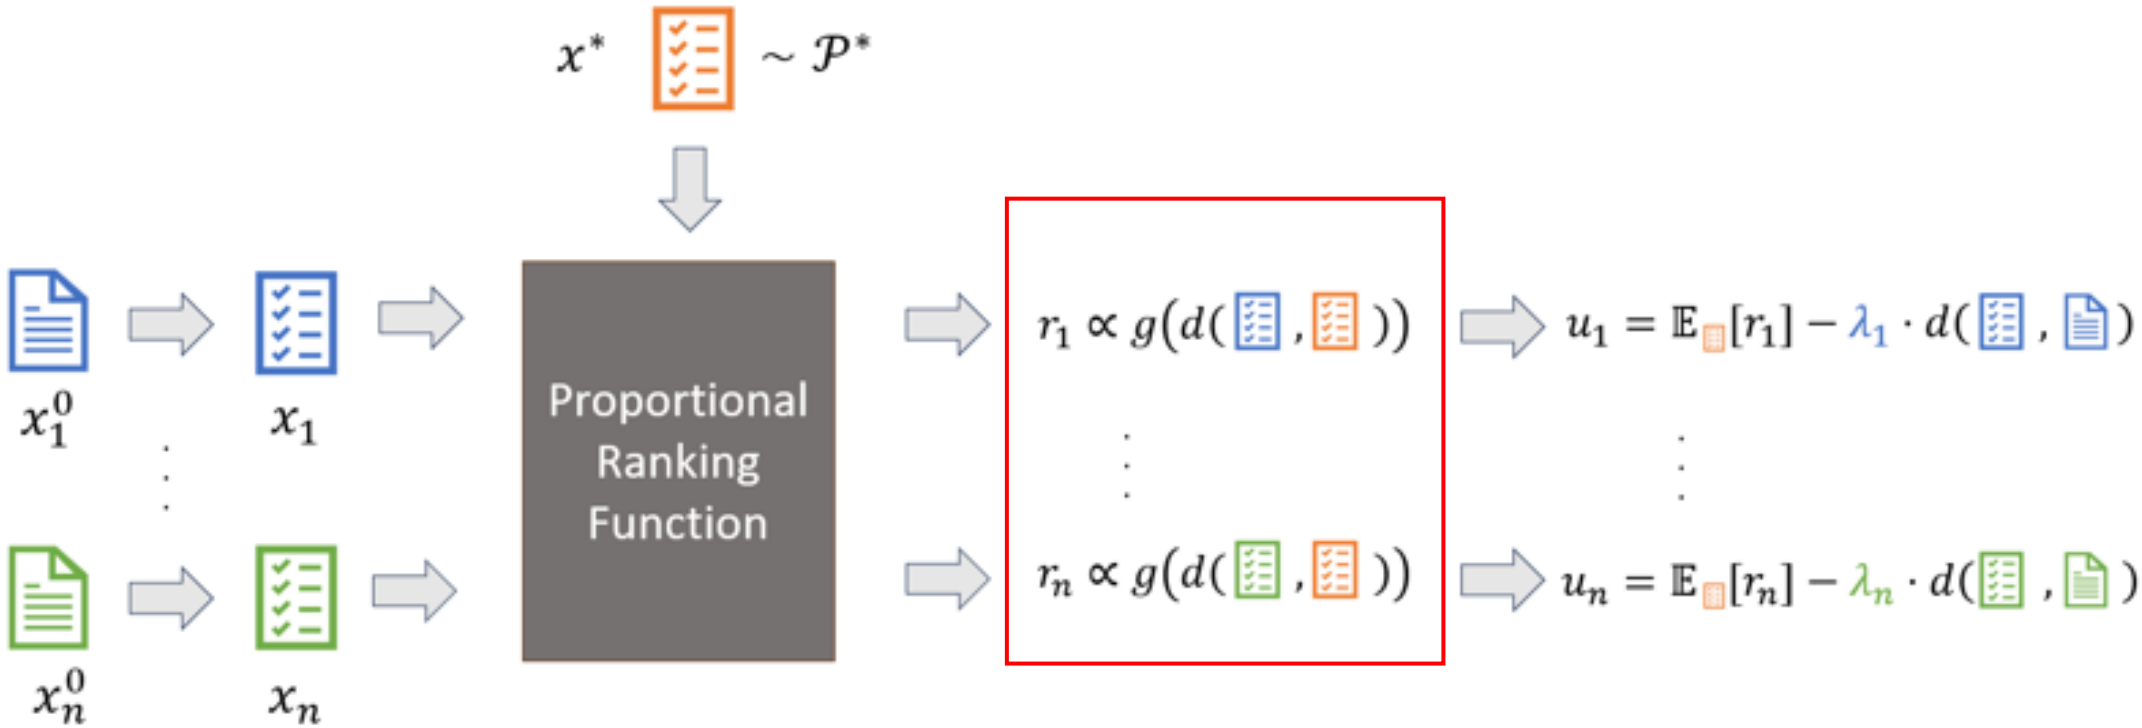

The Publishers' Game

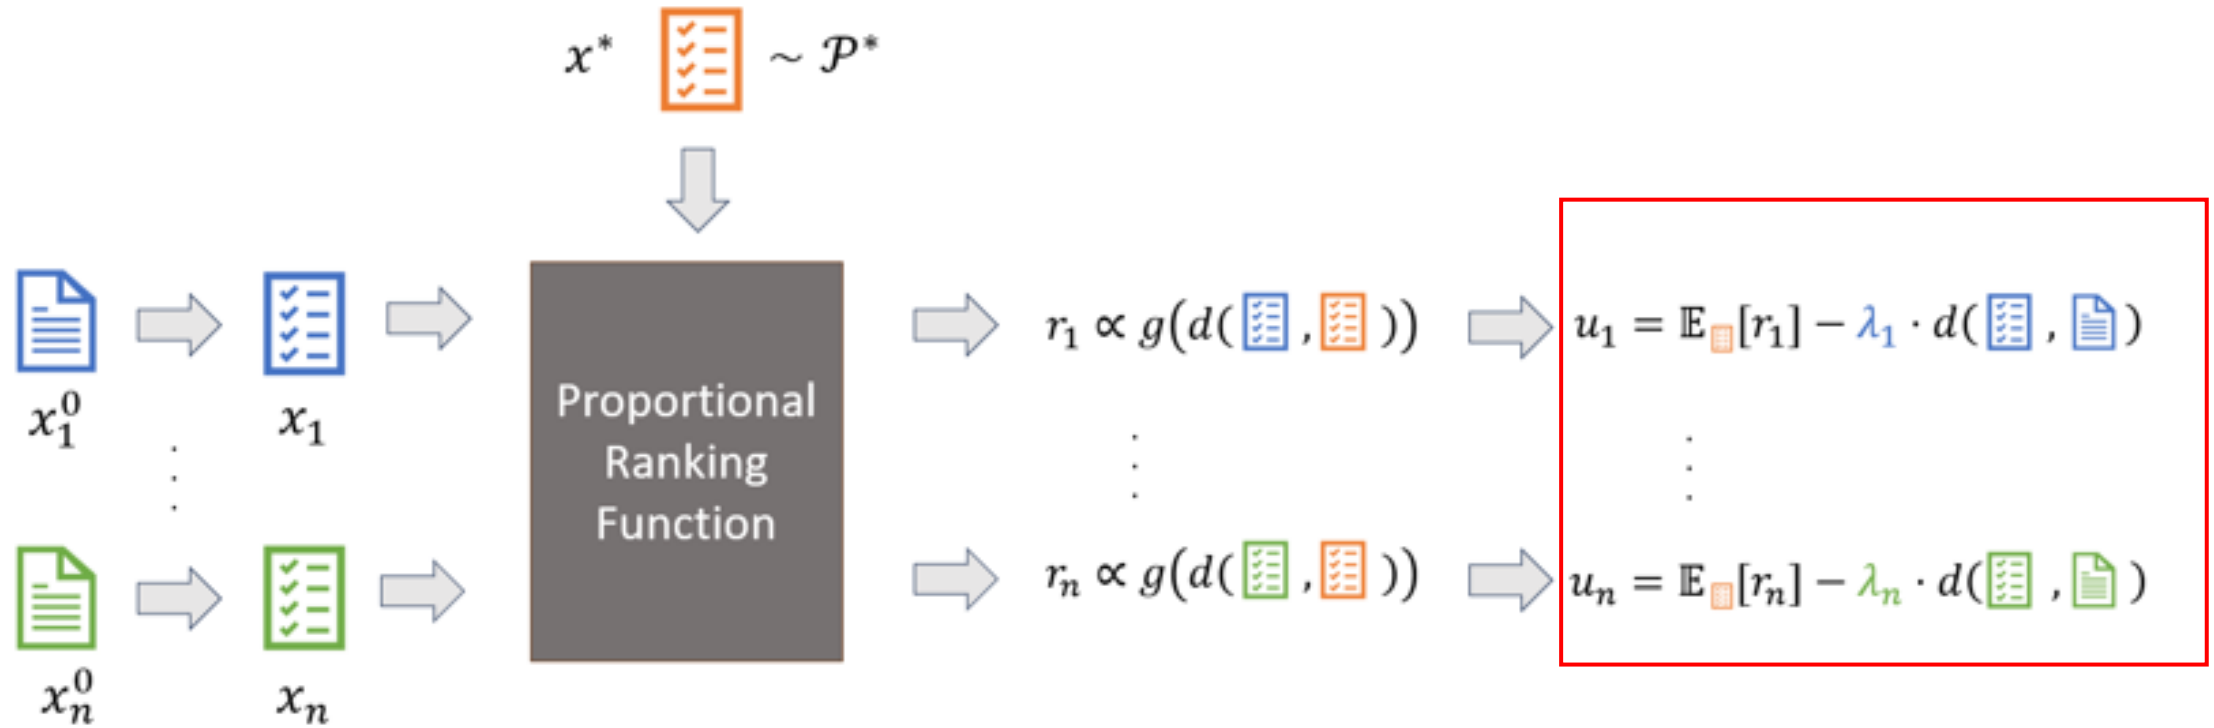

No-Regret Dynamics

Publishers aim to minimize their *external regret*:

$$Reg_i^T := \max_{x_i \in X_i} \left\{ \sum_{t \in [T]} u_i \left(x_i, x_{-i}^{(t)} \right) \right\} - \sum_{t \in [T]} u_i \left(x_i^{(t)}, x_{-i}^{(t)} \right)$$

In no-regret dynamics, we assume that all publishers' regret grows slowly (sub-linearly).

Main Result

- **Theorem (informal).** Let r be a PRF with activation function g. Then, the following claims are equivalent:
  - The activation g is concave
  - All games induced by r are concave (as in Rosen 1965)
  - All games induced by r are socially-concave (as in Even-Dar et al. 2009)
- **Corollary (informal).** Using concave activation guarantees the convergence of no-regret dynamics to a Nash equilibrium.
- In the paper: some more experimental results.

Thank you for listening!