# Together & Fair:
# On Meritocracy in Optimal Set Selection

Debabrota Basu[a]

Joint work with Christos Dimitrakakis[b],
Thomas Kleine Büning, Meirav Segal and Anne-Marie George[c]

[a] Équipe Scool, Univ. Lille, Inria, UMR 9189-CRIStAL, CNRS, Centrale Lille, France
[b] University of Neuchatel, Switzerland and University of Oslo, Norway
[c] University of Oslo, Norway

# Together & Fair: The Roadmap

# Algorithmic Decision Making

*Selecting Individuals*

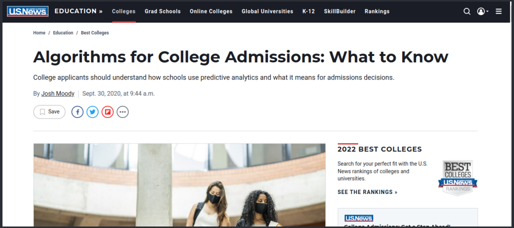


Figure: College Admissions


Figure: Job Recruitment

# A Meritocratic Decision Maker [You94]

*Selecting Individuals*

Alice


Bob


Carla


David


# A Meritocratic Decision Maker [You94]

*Selecting Individuals*

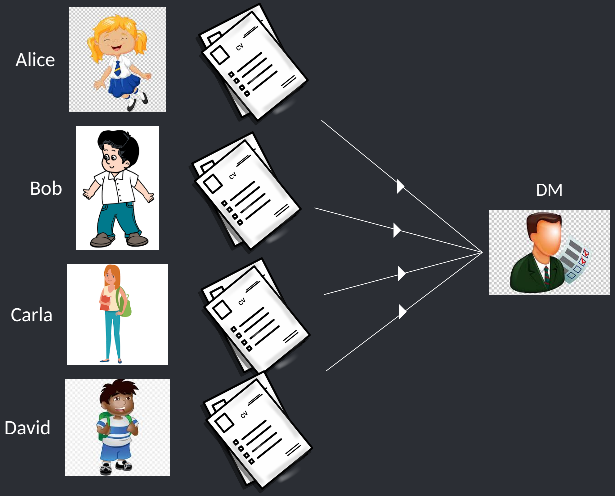

# A Meritocratic Decision Maker [You94]

*Selecting Individuals*

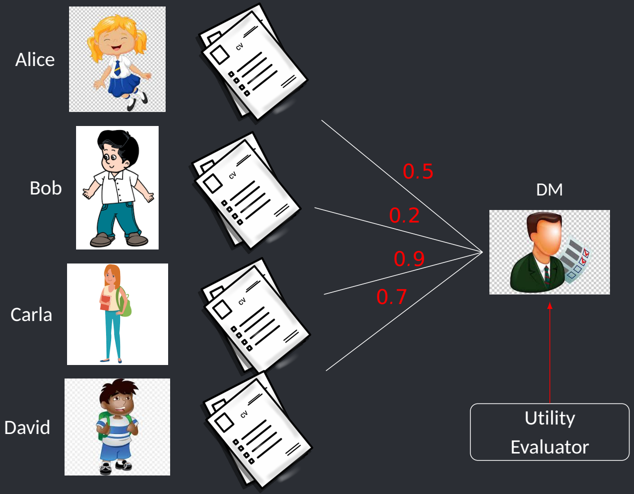

# A Meritocratic Decision Maker [You94]

*Selecting Individuals*

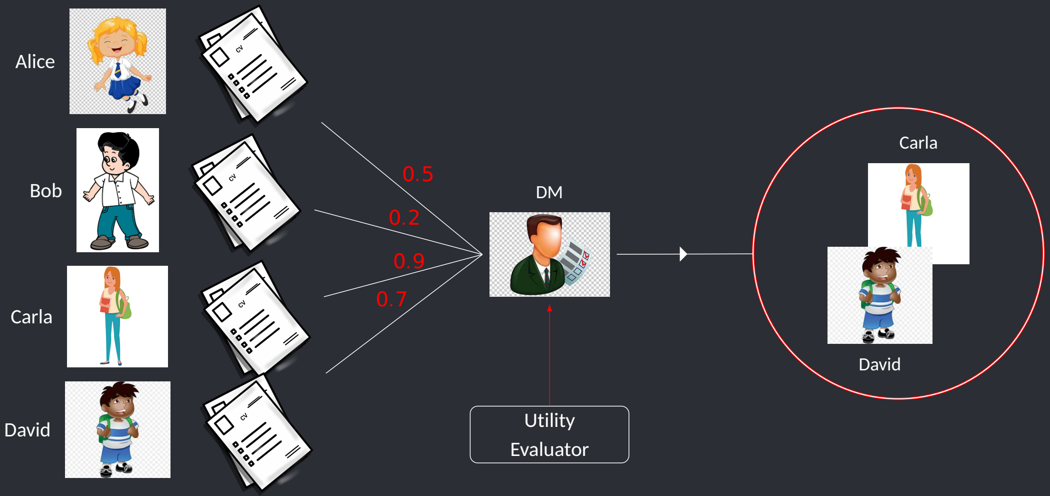

# A Meritocratic Decision Maker [You94]

*Selecting Individuals*

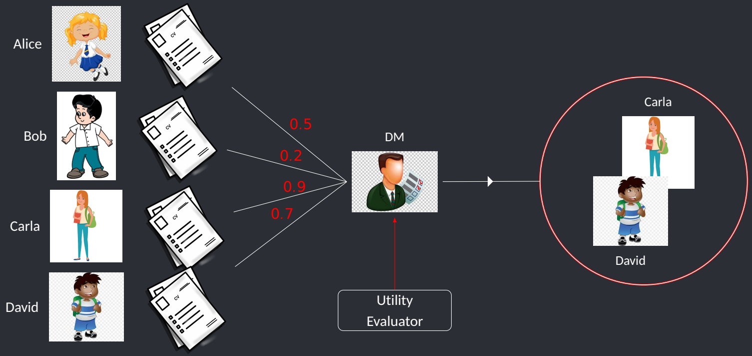


**Assumption:** The success of each candidate is independent from one another.

$\Longrightarrow$ Individual's contribution to DM's utility is the estimated probability of success.

# Algorithmic Decision Making

*Selecting Set of Individuals*

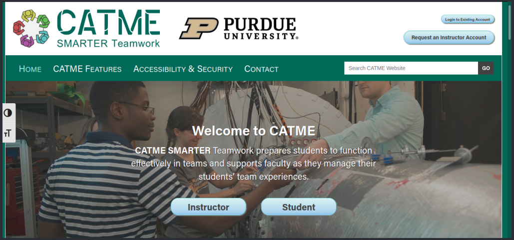


Figure: Student Groups for Projects


Figure: Fantasy Premier League

# Algorithmic Decision Making

*Selecting Set of Individuals*


Figure: Participatory Budget

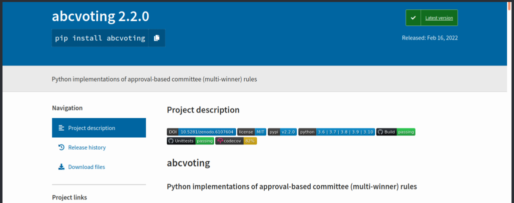


Figure: Committee Selection

# A Meritocratic Decision Maker

*Selecting Sets: Team Building*

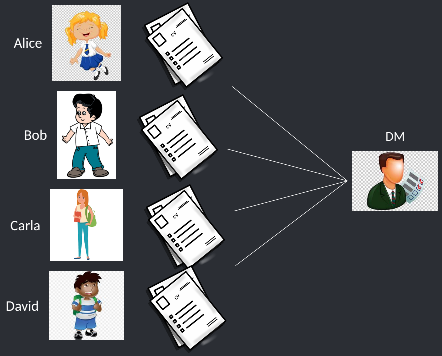


Goal: The DM wants to build a team of two that can do football analytics.

# A Meritocratic Decision Maker

*Selecting Sets: Team Building*

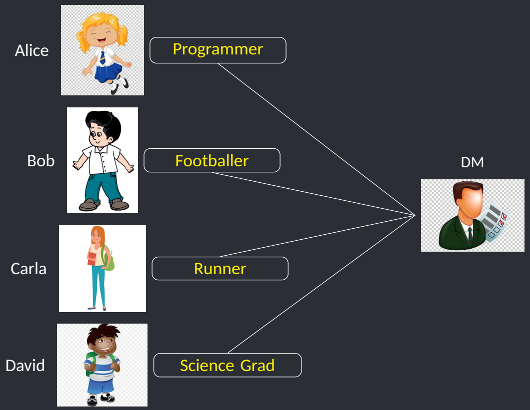


**Observation:** Individuals can have complementary skills or features, and a team is more than an individual.

# A Meritocratic Decision Maker

*Selecting Sets: Team Building*

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

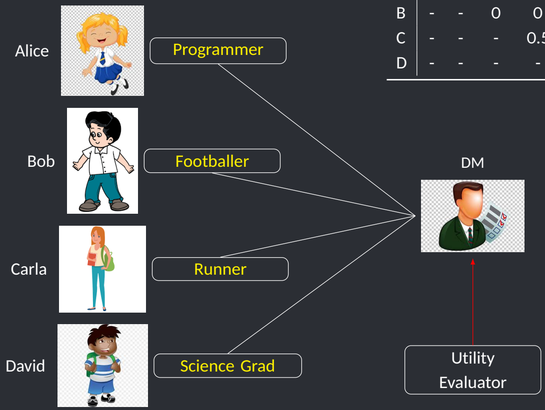

# A Meritocratic Decision Maker

*Selecting Sets: Team Building*

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

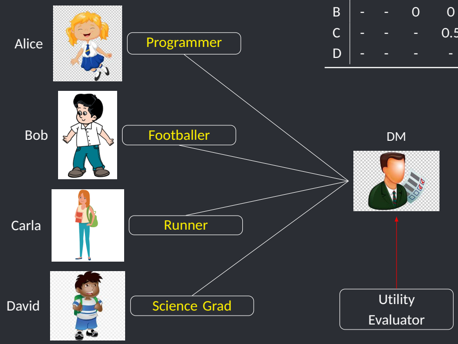


**Observation:** Complementary features lead to higher probability of success for the team.

$\Longrightarrow$ Individual’s contribution to DM’s utility is dependent on the possible teams.

# A Meritocratic Decision Maker

## *Selecting Sets: Team Building*

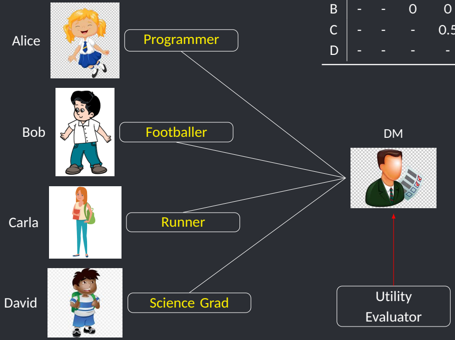


Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

**Result:** Average of marginal contributions of Alice, Bob, Carla and David to all possible teams are: $(1/12, 0, 0, -1/12)$.

# A Meritocratic Decision Maker

*Selecting Sets: Team Building*

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

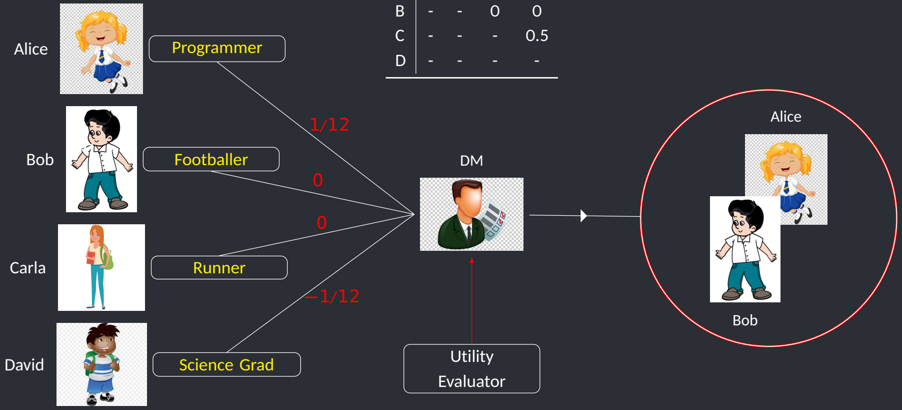

# Questions

How to quantify ‘merit’ of an individual during a set selection?

What are the factors that influence merit?

How to maximise ‘merit’ of a set during selection?

# Our Answers [BSBD21]

### How to quantify 'merit' of an individual during a set selection?

Quantification of 'merit' depends on the expected contribution of the individual given the composition of the teams and their estimated utilities.

### What are the factors that influence merit?

The utility function of the DM, the utility evaluator for the teams, and the probability of selecting a team.

### How to maximise 'merit' of a set during selection?

Computing a policy, i.e. a vector dictating probabilities of selection, which maximises the total expected utility of the selected team.

# Together & Fair: The Roadmap

# Optimal Set Selection

*A Decision Theoretic Formulation*

**Ingredients:**

Population: $\mathcal{N} \triangleq \{1, \ldots, N\}$ denotes the set $n$ individuals

Population's data: $\mathbf{x} \triangleq \{x_1, \ldots, x_N\}$ denotes the set of features of $n$ individuals

Outcomes: $\mathbf{y} \in \mathcal{Y}$ denotes outcomes due to the selected set's performance

Utility: $u(\mathbf{a}, \mathbf{y})$ denotes the utility of the selection $\mathbf{a} \subseteq \mathcal{N}$ w.r.t. the outcomes $\mathbf{y}$

Policy: $\pi(\mathbf{a} \mid \mathbf{x})$, denotes the probability of selecting a subset $\mathbf{a} \subseteq \mathcal{N}$ given the data

# Optimal Set Selection

*A Decision Theoretic Formulation*

## Goodness of a Selection Policy: Expected Utility

$$U(\pi, \mathrm{x}) \triangleq \mathbb{E}_\pi[\,u \mid \mathrm{x}\,] = \mathbb{E}_\pi[\,\mathbb{E}[\,u \mid \mathrm{a}, \mathrm{x}\,]\,]$$
$$= \sum_{\mathrm{a} \subseteq \mathcal{N}} \underbrace{\pi(\mathrm{a} \mid \mathrm{x})}_{\text{Policy}} \sum_{\mathrm{y} \in \mathcal{Y}} \mathbb{P}(\mathrm{y} \mid \mathrm{a}, \mathrm{x})\, \underbrace{u(\mathrm{a}, \mathrm{y})}_{\text{Utility}}.$$

$\mathbb{P}(\mathrm{y} \mid \mathrm{a}, \mathrm{x})$ is a predictive model used by the DM for estimating outcome probabilities.

# Optimal Set Selection

*A Decision Theoretic Formulation*

## Goodness of a Selection Policy: Expected Utility

$$
\begin{aligned}
U(\pi, \mathrm{x}) &\triangleq \mathbb{E}_\pi[\, u \mid \mathrm{x} \,] = \mathbb{E}_\pi[\, \mathbb{E}[\, u \mid \mathrm{a}, \mathrm{x} \,]\,] \\
&= \sum_{\mathrm{a} \subseteq \mathcal{N}} \underbrace{\pi(\mathrm{a} \mid \mathrm{x})}_{Policy} \sum_{\mathrm{y} \in \mathcal{Y}} \underbrace{\mathbb{P}(\mathrm{y} \mid \mathrm{a}, \mathrm{x})}_{Predictive model} \underbrace{u(\mathrm{a}, \mathrm{y})}_{Utility}.
\end{aligned}
$$

## Optimal Set Selection as Policy Optimisation

Given a family of parameterised policies $\Pi \triangleq \{\, \pi_\theta \mid \theta \in \Theta \,\}$, compute the policy maximising the expected utility

$$
\theta^*(\mathrm{x}) = \operatorname*{argmax}_{\theta \in \Theta} U(\pi_\theta, \mathrm{x}).
$$

# Optimal Set Selection

*A Policy Gradient Algorithm*

**Algorithm** A Policy gradient algorithm

**Input:** a model $\mathbb{P}(\mathbf{y}|\mathbf{a}, \mathbf{x})$, a population $\mathscr{N}$ with features $\mathbf{x}$ and a utility function $u$.
**Initialise:** $\theta^0$, $\delta > 0$, learning rate $\eta > 0$
**while** $||\,\theta^{i+1} - \theta^i\,|| > \delta$ **do**
    Evaluate $\nabla_\theta U(\pi_\theta, \mathbf{x})$ from $u$, $\mathbf{x}$ and $\mathbb{P}$.
    $\theta^{i+1} \leftarrow \theta^i + \eta[\nabla_\theta U(\pi_\theta, \mathbf{x})]_{\theta=\theta^i}$
    $i{+}{+}$
**end while**
**return** $\pi_{\theta^{i+1}}$

# Together & Fair: The Roadmap

1. Algorithmic Decision Making
    1.1 Selecting Individuals
    1.2 Selecting Set of Individuals
2. Optimal Set Selection: A Decision Theoretic Formulation
3. Facets of Meritocracy in Set Selection
    3.1 A Quantifier of Merit: Expected Marginal Contribution (EMC)
    3.2 Two Stability Criteria for Meritocracy: Swap and Local
    3.3 A Connection between Policy Gradients and EMC: Separable Policies
4. A Case Study: Norwegian College Admissions Data

# Expected Marginal Contribution

*A Communitarian Quantifier of Merit*

## Expected Marginal Contribution (EMC)

*Expected Marginal Contribution* (EMC) quantifies the gain (or loss) in expected utility when a policy is constrained to always pick an individual $i$.

$$\mathrm{EMC}_i(U, \pi) \triangleq \sum_{\mathbf{a} \in \mathscr{A}} \pi(\mathbf{a})\big[U(\mathbf{a} + i) - U(\mathbf{a})\big].$$

We quantify the 'merit' of an individual $i$ using **EMC** for a given policy $\pi$, utility $U$, data $\mathbf{x}$.

# Revisiting the Meritocratic Decision Maker

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

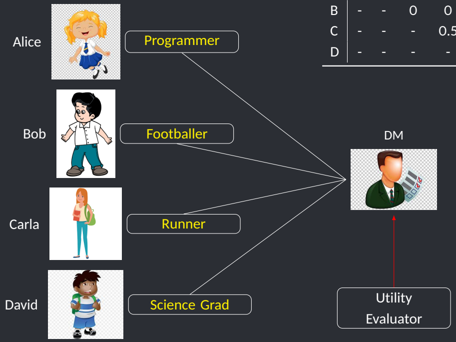

# Revisiting the Meritocratic Decision Maker

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

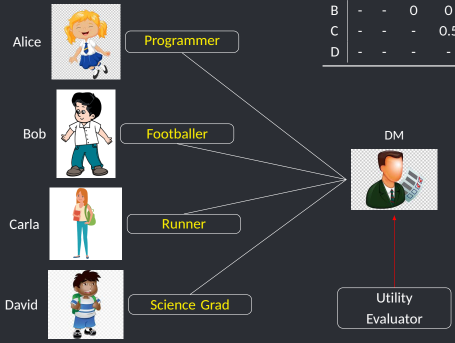


DM's Policy: $\pi = \pi_{\text{egal}}$

$\text{EMC}(U, \pi_{\text{egal}}) = (1/6, 0, 0, -1/6)$

# Revisiting the Meritocratic Decision Maker

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

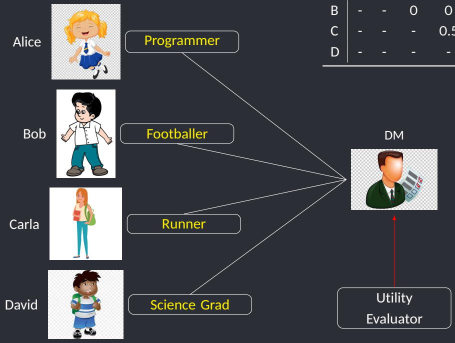


DM's Policy: $\pi = \pi_{\text{egal}}$

$\text{EMC}(U, \pi_{\text{egal}}) = (1/6, 0, 0, -1/6)$

DM's Policy: $\pi =$ never select $A$ and $D$ together, select other sets with equal probability

$\text{EMC}(U, \pi) = (3/24, 1/24, -1/24, -2/24)$

# Properties of EMC

*Generalising Shapley Value*

**Lemma (Axioms of Fair Coalition/Division)**

1) ***Symmetry:*** *If* $U(\mathrm{a} + i) = U(\mathrm{a} + j)$ *for all* $\mathrm{a} \subseteq \mathcal{N}$,

$$\mathrm{EMC}_i(U, \pi) = \mathrm{EMC}_j(U, \pi) \qquad \forall \pi \in \Pi.$$

2) ***Linearity:*** *For all* $\alpha, \beta \in \mathbb{R}$,

$$\mathrm{EMC}(\alpha U_1 + \beta U_2, \pi) = \alpha\, \mathrm{EMC}(U_1, \pi) + \beta\, \mathrm{EMC}(U_2, \pi) \quad \forall \pi \in \Pi.$$

3) ***Null Players:*** *If* $i \in \mathcal{N}$ *has zero contribution to every set,*

$$\mathrm{EMC}_i(U, \pi) = 0 \qquad \forall \pi \in \Pi.$$

# Properties of EMC

*Generalising Shapley Value*

## Lemma (Axioms of Fair Coalition/Division)

1) ***Symmetry:*** *If* $U(\mathrm{a}+i) = U(\mathrm{a}+j)$ *for all* $\mathrm{a} \subseteq \mathcal{N}$,

$$\mathrm{EMC}_i(U,\pi) = \mathrm{EMC}_j(U,\pi) \qquad \forall \pi \in \Pi.$$

2) ***Linearity:*** *For all* $\alpha, \beta \in \mathbb{R}$,

$$\mathrm{EMC}(\alpha U_1 + \beta U_2, \pi) = \alpha\,\mathrm{EMC}(U_1,\pi) + \beta\,\mathrm{EMC}(U_2,\pi) \quad \forall \pi \in \Pi.$$

3) ***Null Players:*** *If* $i \in \mathcal{N}$ *has zero contribution to every set,*

$$\mathrm{EMC}_i(U,\pi) = 0 \qquad \forall \pi \in \Pi.$$

For the egalitarian selection policy $\pi_{\mathrm{egal}}(\mathrm{a}) = \frac{1}{N\binom{N-1}{\|\mathrm{a}\|_1}}$, $\mathrm{EMC}(U, \pi_{\mathrm{egal}}) = \mathrm{Shapley}(U)$ [Sha51].

# Meritocracy and Stability of Selected Set

*Swap Stability*

**Definition (Swap Stability)**

If for any two individuals $i, j \in \mathcal{N}$ with $\pi(\mathrm{a}_i = 1) > \pi(\mathrm{a}_j = 1)$, a *swap stable* policy $\pi$ satisfies $U(\pi + i - j) \geq U(\pi - i + j)$[1]

A policy $\pi$ is swap stable if an individual $i$ is more likely to be selected than $j$, *the expected utility of selecting $i$ but not $j$ is higher than the expected utility of selecting $j$ but not $i$.*

**Lemma (EMC Induces a Swap Stable Ordering)**

$$\mathrm{EMC}_i(U, \pi - i - j) \geq \mathrm{EMC}_j(U, \pi - i - j) \iff U(\pi + i - j) \geq U(\pi - i + j).$$

[1] $U(\pi + i - j) \triangleq \sum_{\mathrm{a} \in \mathscr{A}} \pi(\mathrm{a}) U(\mathrm{a} + i - j)$.

# Meritocracy and Stability of Selected Set

*Local Stability*

**Definition (Local Stability)**

A policy $\pi$ is *locally stable* if for any $i \in \mathcal{N}$, $U(\pi) \geq U(\pi + i)$[2].

Local stability guarantees that the expected utility of selecting $i$ is lower than the expected utility of $\pi$.

**Lemma (Local Stability is Equivalent to Negative EMC)**

*For* $i \in \mathcal{N}$, $U(\pi) \geq U(\pi + i) \iff \mathrm{EMC}_i(U, \pi) \leq 0$.

[2] $U(\pi + i) \triangleq \sum_{\mathbf{a} \in \mathcal{A}} \pi(\mathbf{a}) U(\mathbf{a} + i)$.

# Revisiting the Meritocratic Decision Maker

DM's Policy: $\pi = \pi_{\text{egal}}$

$\pi_{\text{egal}}$ is swap stable but not local stable

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

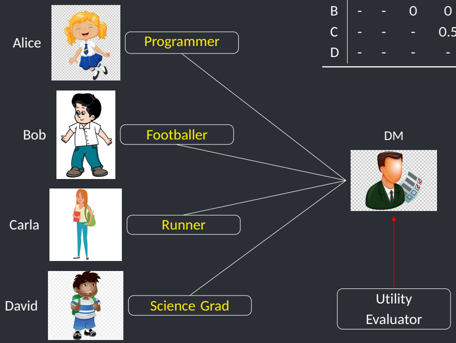

# Revisiting the Meritocratic Decision Maker

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

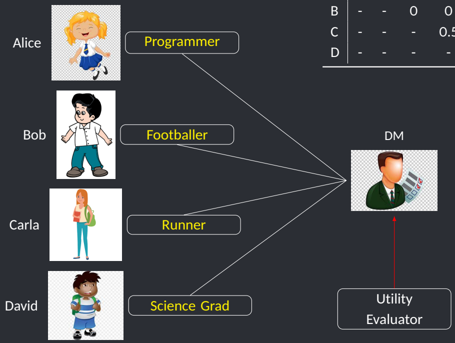


DM's Policy: $\pi = \pi_{\text{egal}}$

$\pi_{\text{egal}}$ is swap stable but not local stable

DM's Policy: $\pi_{C,D}$ selects Carla and David

$\pi_{C,D}$ is locally stable and swap stable but not utility maximising

# Revisiting the Meritocratic Decision Maker

Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

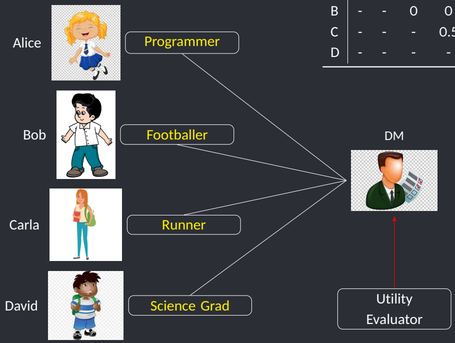


DM's Policy: $\pi = \pi_{\text{egal}}$

$\pi_{\text{egal}}$ is swap stable but not local stable

DM's Policy: $\pi_{C,D}$ selects Carla and David

$\pi_{C,D}$ is locally stable and swap stable but not utility maximising

DM's Policy: $\pi_{A,B}$ selects Alice and Bob

$\pi_{C,D}$ is locally stable, swap stable and maximises utility

# Revisiting the Meritocratic Decision Maker

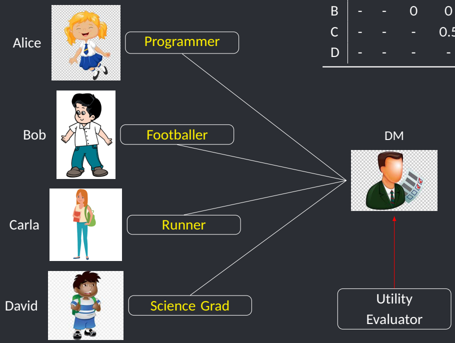


Utilities Together

| | A | B | C | D |
|---|---|---|---|---|
| A | - | 1 | 0.5 | 0 |
| B | - | - | 0 | 0 |
| C | - | - | - | 0.5 |
| D | - | - | - | - |

DM's Policy: $\pi = \pi_{\text{egal}}$

$\pi_{\text{egal}}$ is swap stable but not local stable

DM's Policy: $\pi_{C,D}$ selects Carla and David

$\pi_{C,D}$ is locally stable and swap stable but not utility maximising

DM's Policy: $\pi_{A,B}$ selects Alice and Bob

$\pi_{C,D}$ is locally stable, swap stable and maximises utility

A deterministic policy maximising utility is both swap- and local-stable, i.e. 'meritocratic'.

# Separable Policies

*A Study Relating Policy Gradient and EMC*

A parameterised policy $\pi_\theta$ is *separable over the population* $\mathcal{N}$ if

$$\pi_\theta(\mathbf{a}) = \prod_{i=1}^{N} \underbrace{\pi_\theta(\mathbf{a}_i)}_{\text{Probability of selecting } i} \propto \prod_{i=1}^{N} g(\mathbf{a}_i, \theta_i)$$

for some function $g$ and $\theta = (\theta_1, \ldots, \theta_N)$ .

# Separable Policies

*Separable Softmax Policies*

For separable softmax policies,

$$\pi_\theta(\mathbf{a}) = \frac{e^{\beta\theta^\top \mathbf{a}}}{\sum_{\mathbf{a}' \in \mathscr{A}} e^{\beta\theta^\top \mathbf{a}'}}.$$

$\beta \geq 0$ is the inverse temperature of the distribution.

**Lemma**

*The gradient of the softmax policy $\pi_\theta$ is a linear transformation of the EMC. Specifically,*

$$\nabla_{\theta_i} U(\pi_\theta) = \beta \pi_\theta(\mathbf{a}_i = 1)\ \mathrm{EMC}_i(U, \pi_\theta) \qquad \forall i \in \mathscr{N}.$$

# Separable Policies

*Separable Linear Policies*

Separable linear policies select individual $i$ with probability $\theta_i$, i.e.

$$\pi_{\theta_i}(\mathrm{a}_i) = \theta_i \mathbb{I}\{\mathrm{a}_i = 1\} + (1 - \theta_i)\mathbb{I}\{\mathrm{a}_i = 0\}.$$

**Lemma**

*For separable linear policies, if* $\pi_{\theta_i}(\mathrm{a}_i = 0) > 0$,

$$\nabla_{\theta_i} U(\pi_\theta) = \frac{\mathrm{EMC}_i(U, \pi_\theta)}{\pi_{\theta_i}(\mathrm{a}_i = 0)} \qquad \forall i \in \mathcal{N}.$$

# EMC, Shapley Value, and Policy Gradient

*A Visual Comparison*

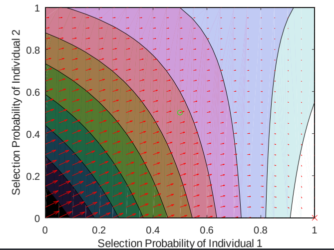


Figure: EMC

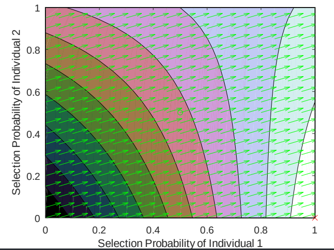


Figure: Shapley value

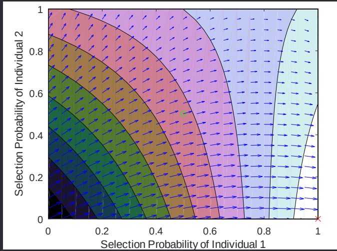


Figure: Policy gradients (Linear)

# Together & Fair: The Roadmap

# Experimental Analysis

*Dataset*

- **Application data**[3]: Applicants to all Norwegian university programs with features: birth date, semester of application, gender, citizenship, country of educational background, high school grades in form of GPA and summarised language/science points, other points, admission decision, each applicant's preference for a program.

- **Exam data**[4]: All students at Norwegian universities for all their taken exams including: courses, study program, and achieved grades.

*Goal:* DM is interested in maximising good course results of admitted students across all three considered disciplines/courses and demographic parity among selected students.

---

[3] `https://dbh.nsd.uib.no/dokumentasjon/tabell.action?tabellId=379`
[4] `https://dbh.nsd.uib.no/dokumentasjon/tabell.action?tabellId=472`

# Experimental Analysis

*Utility Function of the DM*

- **Utility (Log-linear):** $u(\mathbf{a}, \mathbf{y}) = \sum_{j=1}^{3} \log\left(\sum_{i \in \mathcal{N}} \mathbf{a}_i \cdot \mathbf{y}_{i,j}\right) - c \cdot \|\mathbf{a}\|_1$.[3]
  $c$ is the cost associated with admitting a student.

- **Demographic fairness (Statistical Parity):**

$$|\pi(\mathbf{a}_i = 1 \mid i \text{ is male}) - \pi(\mathbf{a}_j = 1 \mid j \text{ is female})| \leq \varepsilon.$$

- **Predictive model:** DM uses a regression model for estimating $\mathbb{P}(\mathbf{y} \mid \mathbf{x}, \mathbf{a})$, i.e. the course results $\mathbf{y}$ of this year's applicants $\mathbf{x}$.

[3] The potential outcomes of non-admitted applicants do not contribute to the utility

# Expected Utility and Deviation from Meritocratic Stability

*Demographically Oblivious Policy*

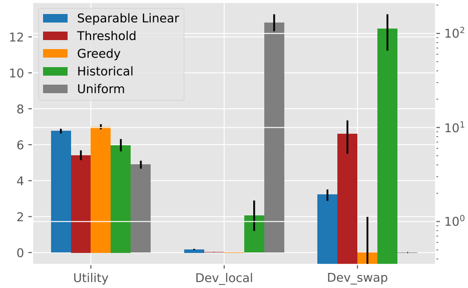


Figure: True Outcome

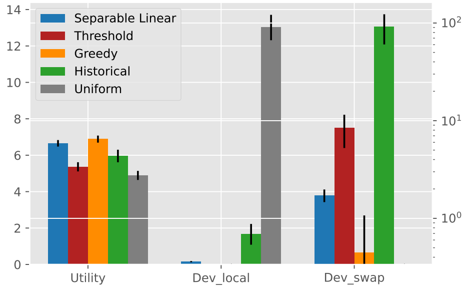


Figure: Predicted Outcomes

# Expected Utility and Deviation from Meritocratic Stability

*Demographic Fairness Constrained Policy*

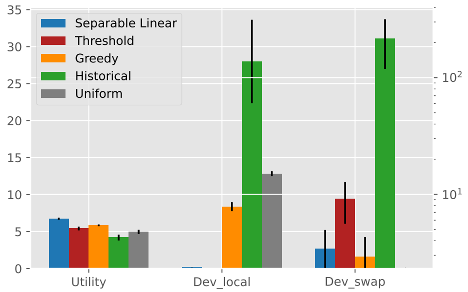


Figure: True Outcome

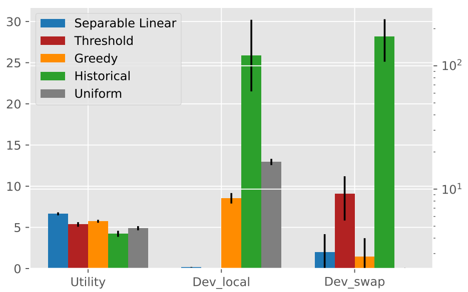


Figure: Predicted Outcomes

## Take-away: Summary of Contributions

- Problem Formulation: Optimal set selection can be reduced to a policy optimisation problem given the utility of DM and a predictive evaluation model.
- Meritocracy in Set Selection:
  -> Expected marginal contribution (EMC) is the quantifier of an individual's 'merit' or contribution to DM's expected utility.
  -> EMC generalises fairness axioms of team building obtained for Shapley values.
  -> A meritocratic policy should satisfy swap and local stabilities.
- Policy Optimisation and EMC:
  -> Policy gradient and EMC for separable policies are proportional.
  -> A deterministic utility maximising policy satisfies meritocratic stabilities but a stochastic utility maximising policy might not.
- Case of College Admissions: Historical and egalitarian policies deviate quite a bit from meritocracy, while linear separable poicies reach closest to meritocracy.

# Take-away: What's Next?

- Stochastic utility maximising policies might not satisfy swap stability.

  -> The proof shows that we might require *individually smooth* policies in the sense that similar individuals (in terms of compatibility across optimal sets) are being selected similarly often.

  Studying the family of individually smooth policies!

- Meritocratic fairness conflicts with demographic fairness constraints [Bin20].

  -> Our experiments show that constrained policy optimisation can solve this setting and EMC for constrained policies can quantify deviation from meritocracy.

  Studying the EMC for constrained policies, and the corresponding trade-off between meritocracy and demographic fairness.

# Bibliography I


[Bin20] Reuben Binns.
On the apparent conflict between individual and group fairness.
In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, page 514–524, New York, NY, USA, 2020. Association for Computing Machinery.

[BSBD21] Thomas Kleine Buening, Meirav Segal, Debabrota Basu, and Christos Dimitrakakis.
Fair set selection: Meritocracy and social welfare.
*CoRR*, abs/2102.11932, 2021.

[Sha51] Lloyd S. Shapley.
Notes on the n-person game II: The value of an n-person game.
Technical report, Rand Corporation, 1951.

[You94] Michael Dunlop Young.
*The rise of the meritocracy*.
Transaction Publishers, 1994.

**How to define and ensure meritocracy in optimal set selection?**

Use expected marginal contribution given DM's policy and utility function.

**Our Paper:** `https://arxiv.org/pdf/2102.11932.pdf`

## Performance Metrics

To measure violations of swap stability, we suggest to use

$$\mathrm{Dev}_{\mathrm{swap}}(\pi, \mathrm{x}) = \sum_{i,j \in \mathcal{N}} (\pi(\mathrm{a}_i = 1 \mid \mathrm{x}) - \pi(\mathrm{a}_j = 1 \mid \mathrm{x}))^+ \left(U(\pi - i + j, \mathrm{x}) - U(\pi + i - j, \mathrm{x})\right)^+,$$

where $(x)^+ \triangleq \max\{0, x\}$. Here, large values of $\mathrm{Dev}_{\mathrm{swap}}(\pi, \mathrm{x})$ indicate large deviations from swap stable decisions. Note that our choice of $\mathrm{Dev}_{\mathrm{swap}}$ not only accounts for the number of infringements, but also the magnitude of the deviation from swap stability. For instance, if $U(\pi + i - j) \ll U(\pi - i + j)$ while $\pi(\mathrm{a}_i = 1) \gg \pi(\mathrm{a}_j = 1)$, the measured deviation from swap stability is accordingly large. In particular, if $\mathrm{Dev}_{\mathrm{swap}}(\pi, \mathrm{x}) = 0$, the policy $\pi$ is swap stable. To measure the deviation from local stability, we use the cumulative positive EMCs under policy $\pi$:

$$\mathrm{Dev}_{\mathrm{local}}(\pi, \mathrm{x}) = \sum_{i \in \mathcal{N}} \left(\mathrm{EMC}_i(U, \pi, \mathrm{x})\right)^+.$$