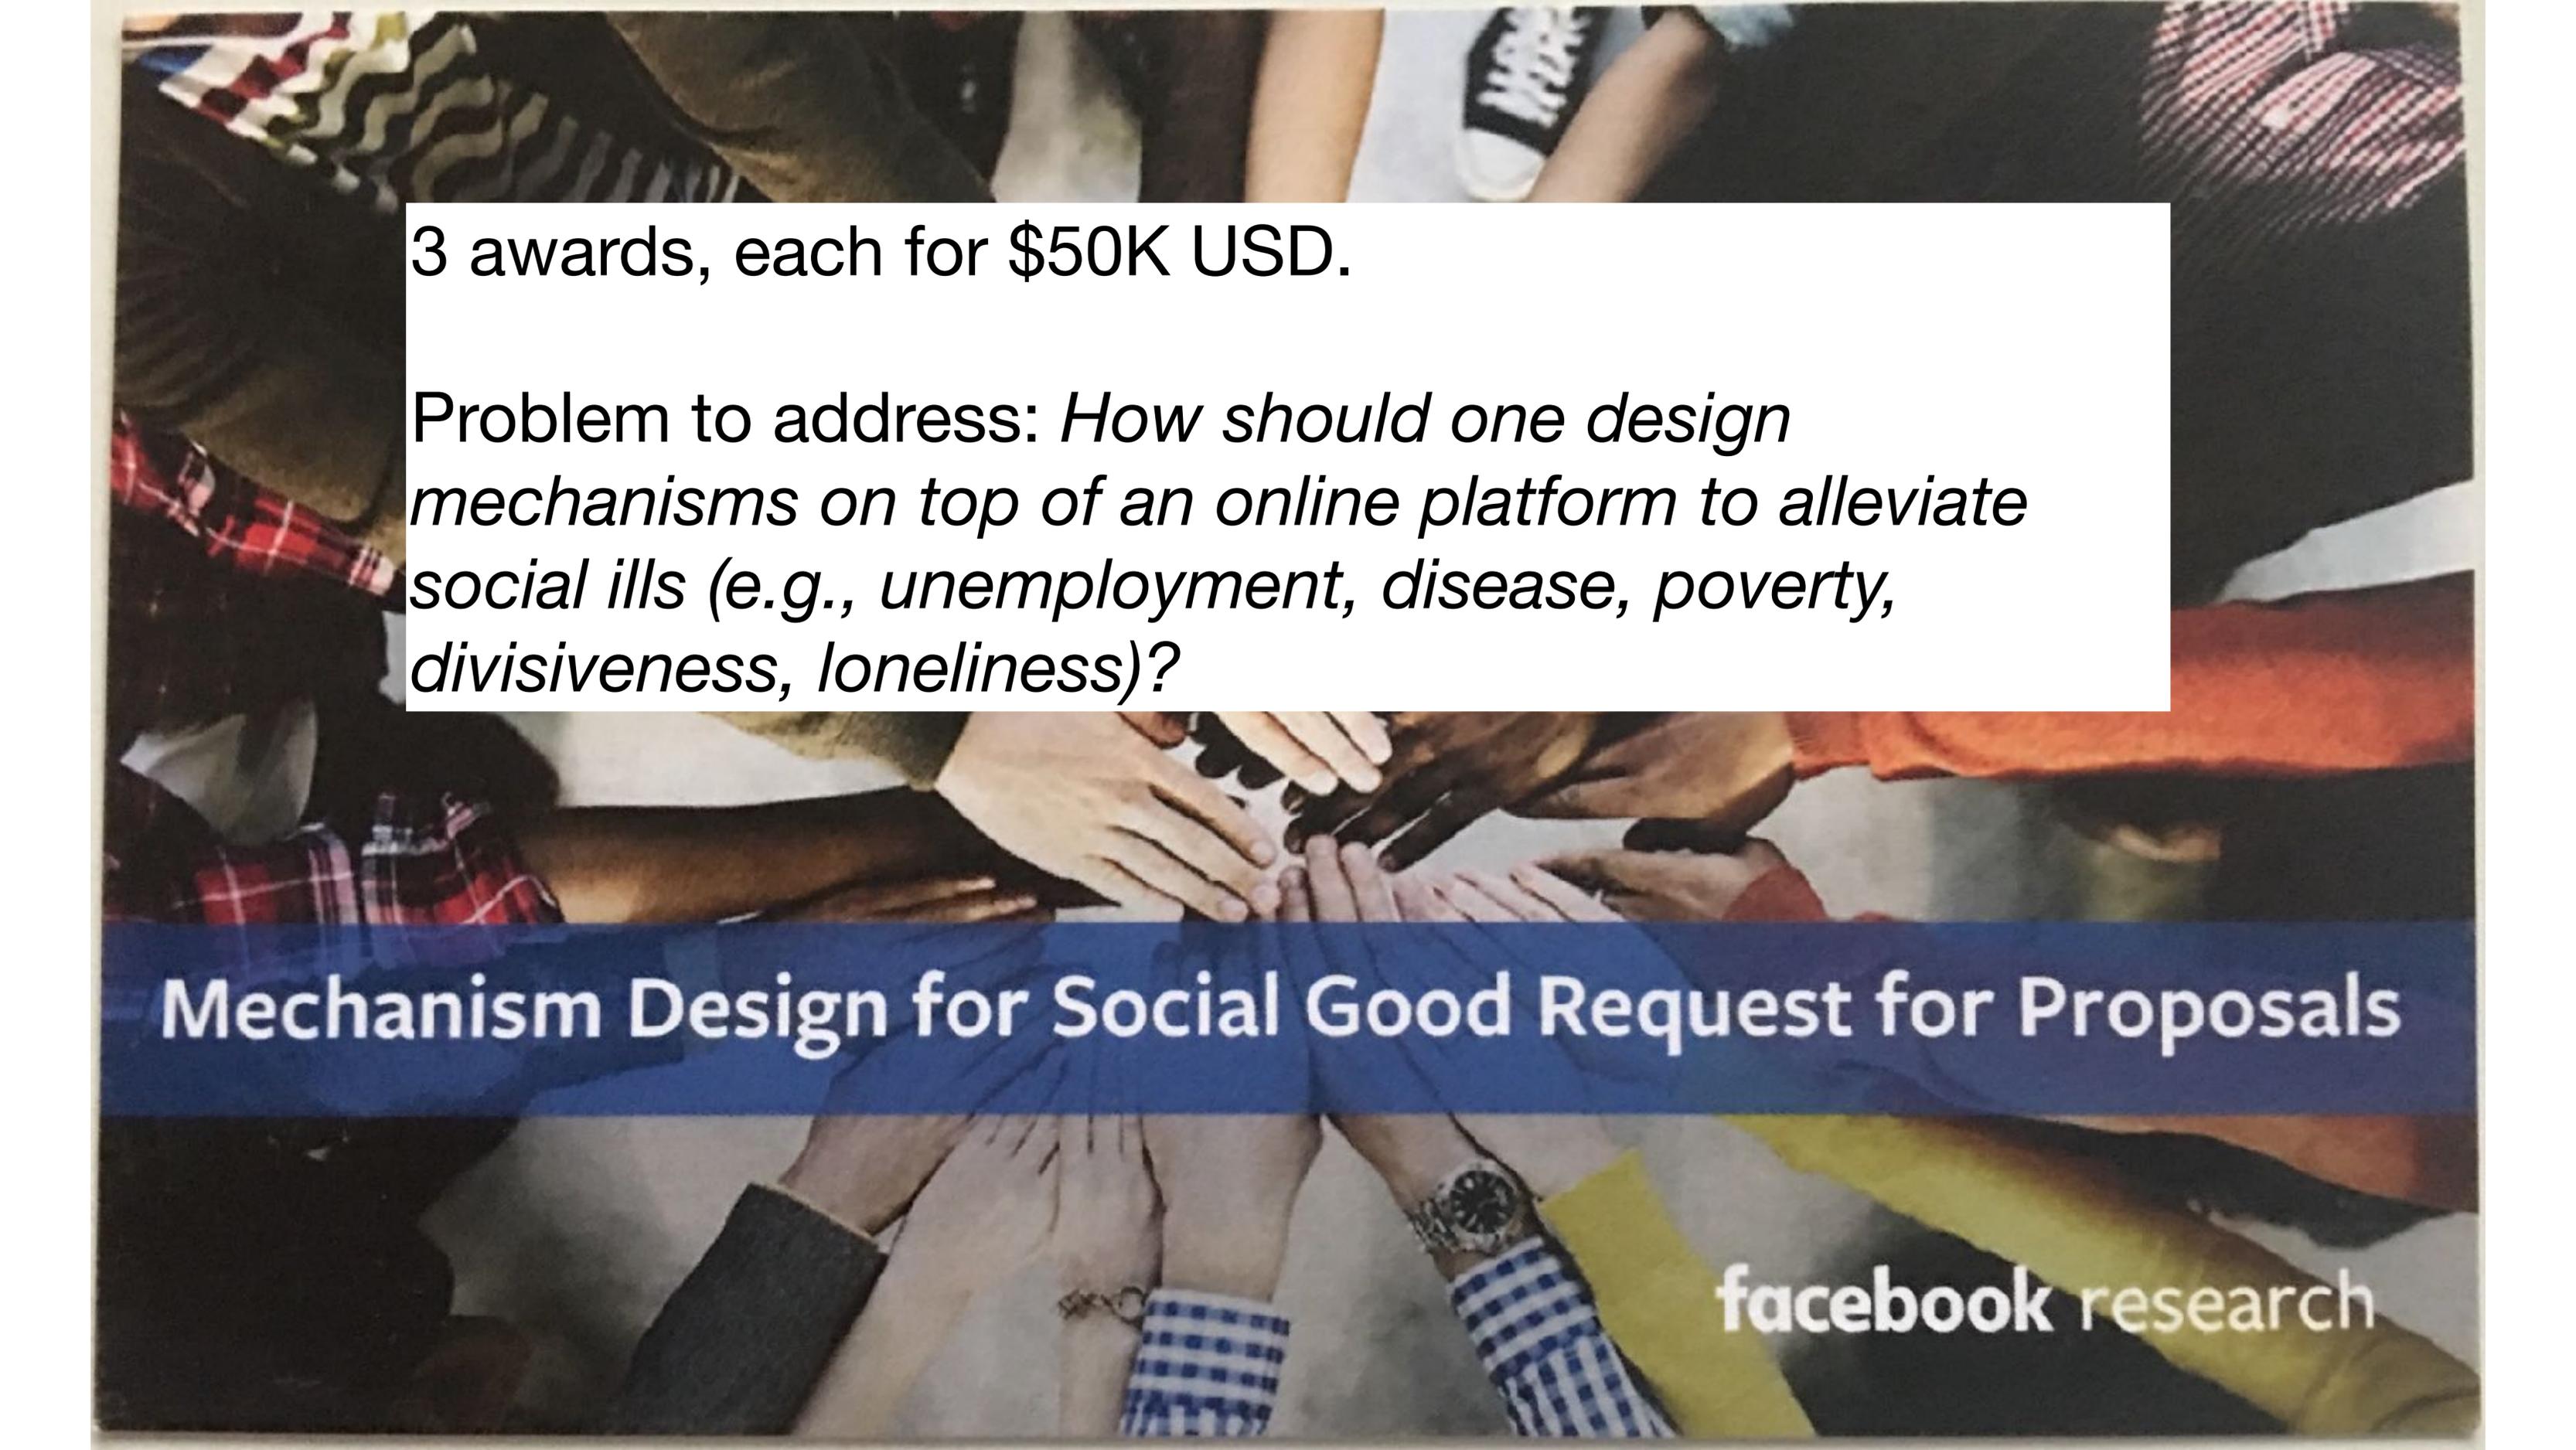


Mechanism Design

Eric Sodomka

Core Data Science @ Facebook

Data Science Africa 2018 - Abuja, Nigeria



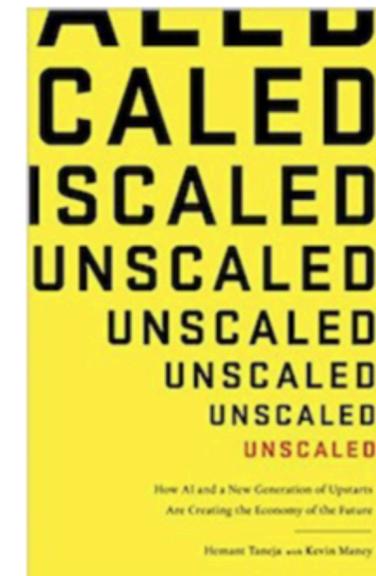
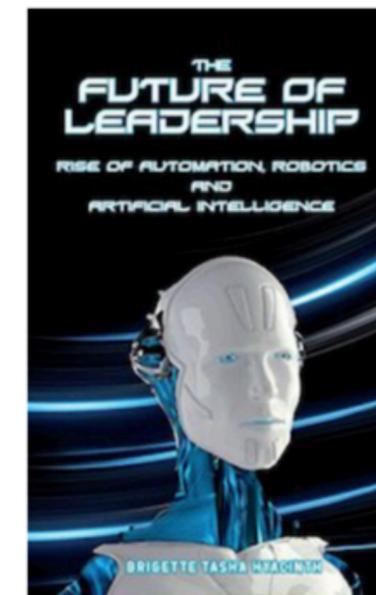
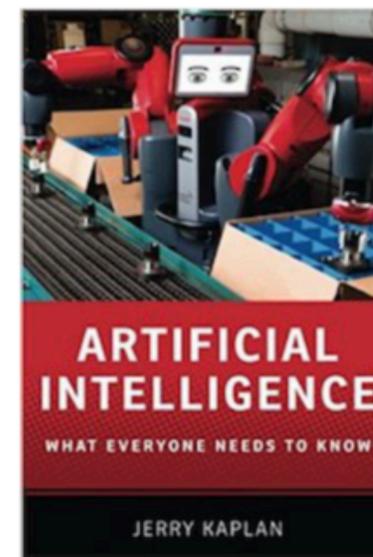
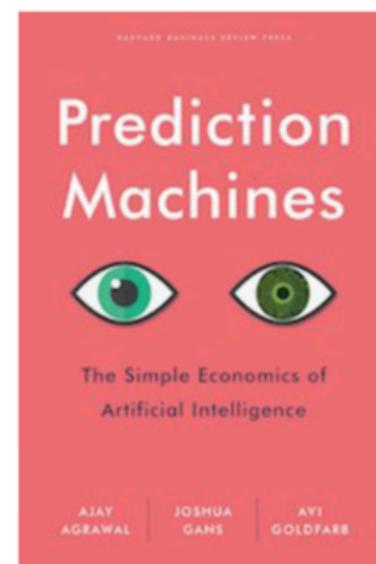
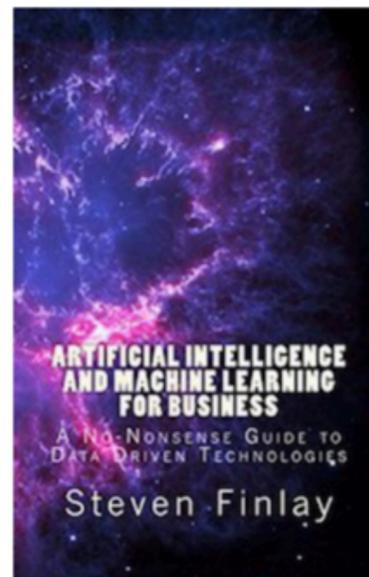
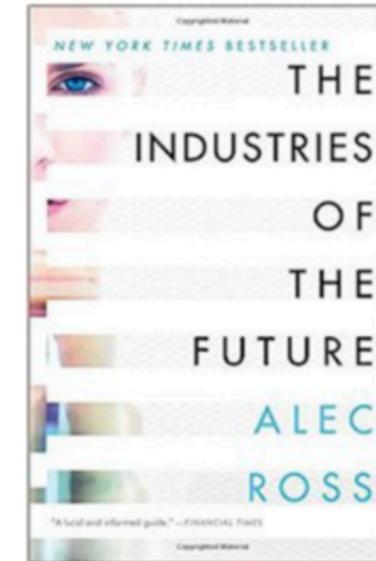
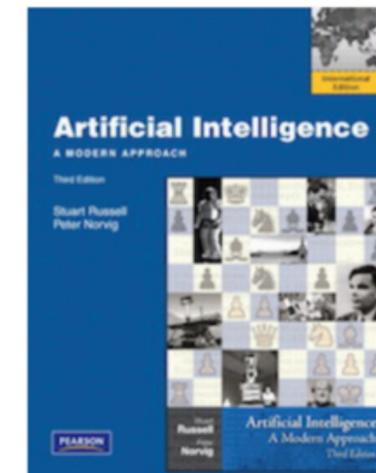
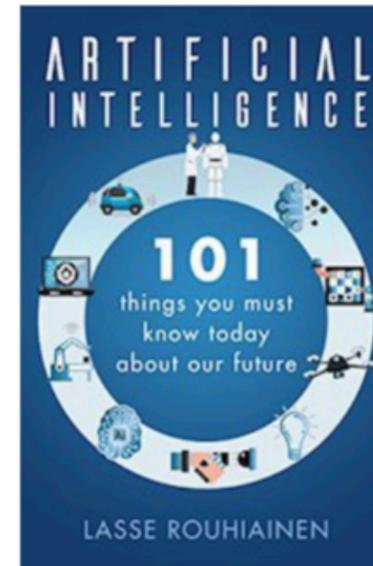
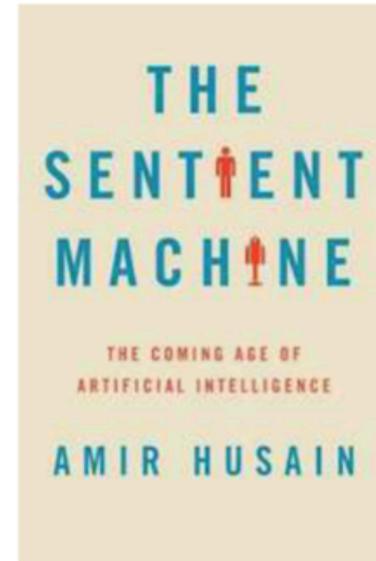
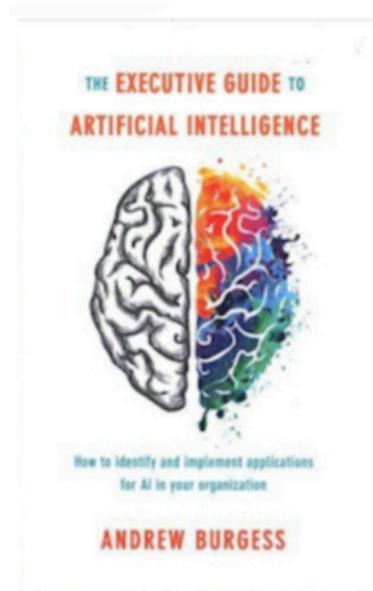
3 awards, each for \$50K USD.

Problem to address: *How should one design mechanisms on top of an online platform to alleviate social ills (e.g., unemployment, disease, poverty, divisiveness, loneliness)?*

Mechanism Design for Social Good Request for Proposals

facebook research

The Problem: Giving out Books



An Initial Proposal

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Discuss with your Neighbor

- Do you have any complaints about Random Assignment?
- Can you come up with anything better?

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Complaints about Random Assignment?

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Complaints about Random Assignment?

- “It doesn’t consider what I want.”
- “I have to trade with people afterwards to get something better.”
- Benefit: “It’s quick and easy.”

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Did you come up with anything better?

- “It doesn’t consider what I want.”
- “I have to trade with people afterwards to get something better.”
- Benefit: “It’s quick and easy.”

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Did you come up with anything better?

Choose Your Favorite

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee **the choice of their favorite book** amongst those remaining.

Random Assignment

1. **Randomly order attendees:** assign each a unique number 1 through 200.
2. In that attendee order, give each attendee a **randomly chosen book** amongst those remaining.

Can Eric allocate books on his own time?

Can Eric allocate books on his own time?

Choose Your Favorite (by proxy)

1. Have each participant submit their **preferences over books**.
2. **Randomly order attendees**: assign each a unique number 1 through 200.
3. In that attendee order, give each attendee **their favorite book** amongst those remaining, **according to their reported preferences**.

Can Eric allocate books on his own time?

Serial Dictatorship

1. Have each participant submit their **preferences over books**.
2. **Randomly order attendees**: assign each a unique number 1 through 200.
3. In that attendee order, give each attendee **their favorite book** amongst those remaining, **according to their reported preferences**.

Evaluating Mechanisms

- How is **Serial Dictatorship** better than **Random Assignment**?
- What does it mean for a mechanism to be “good”?

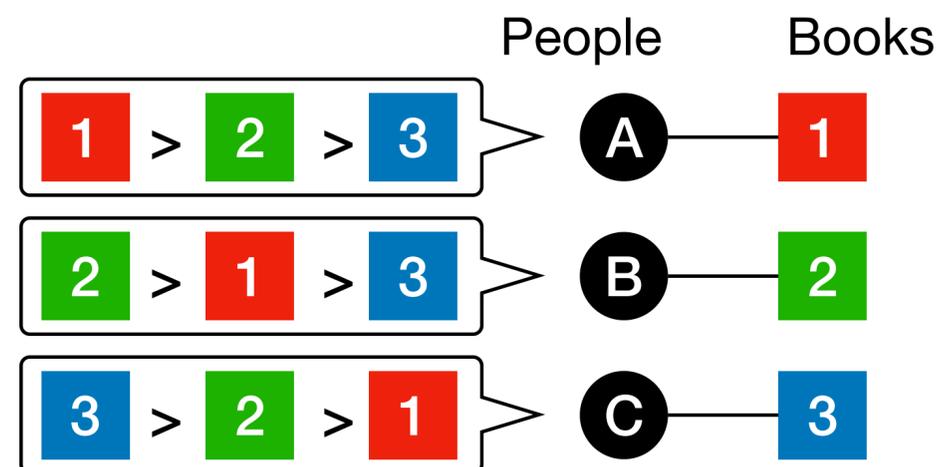
One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

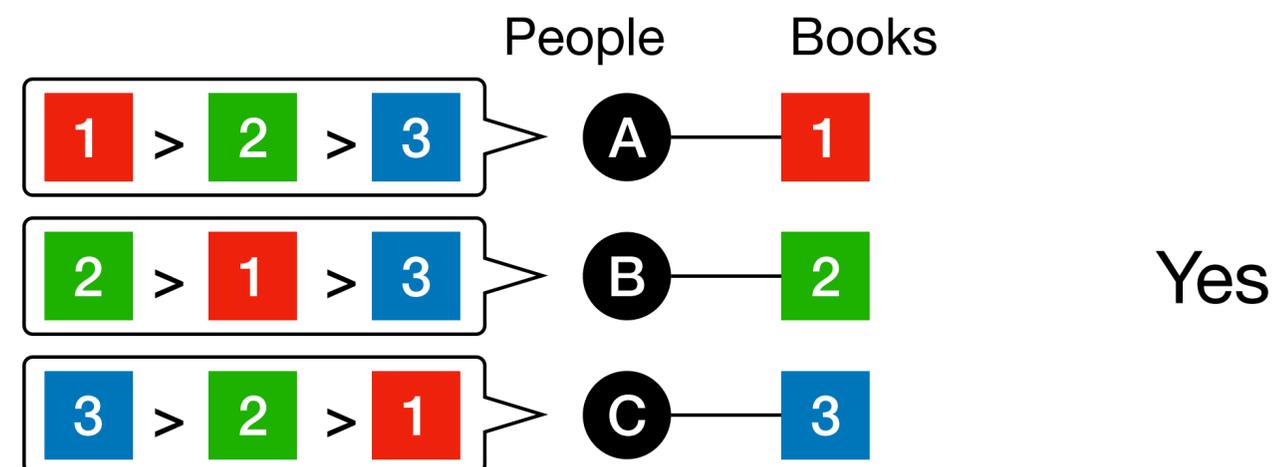
Is this outcome pareto optimal?



One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

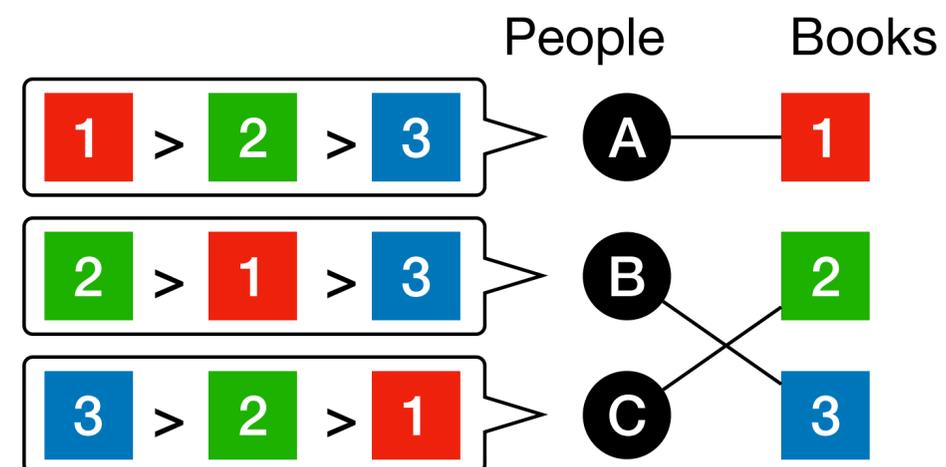
Is this outcome pareto optimal?



One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

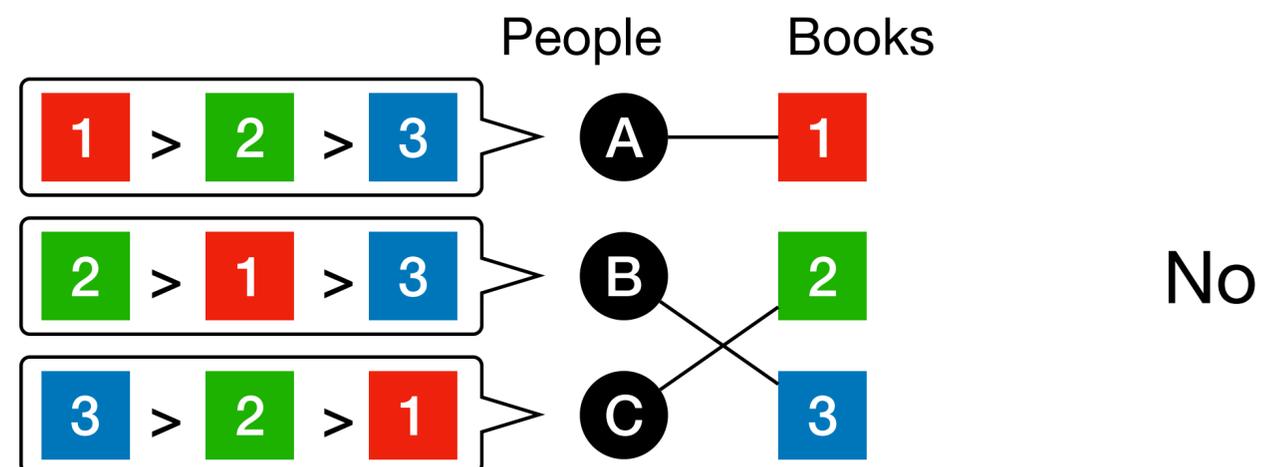
Is this outcome pareto optimal?



One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

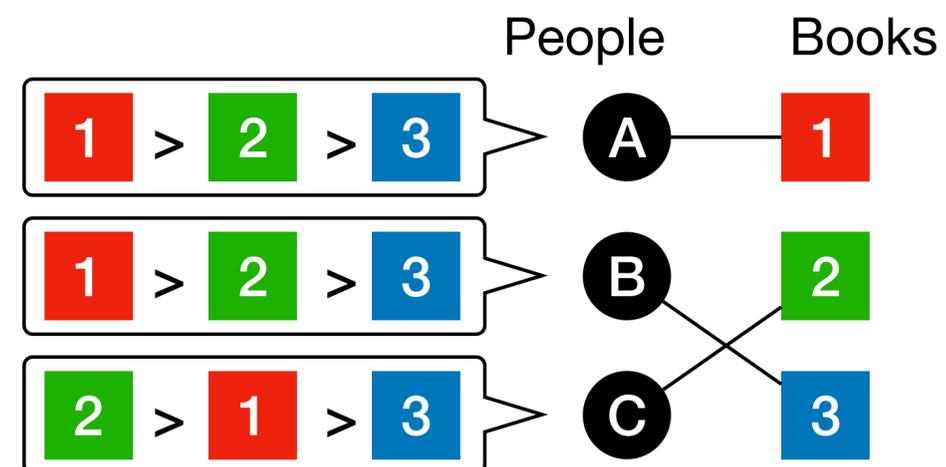
Is this outcome pareto optimal?



One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

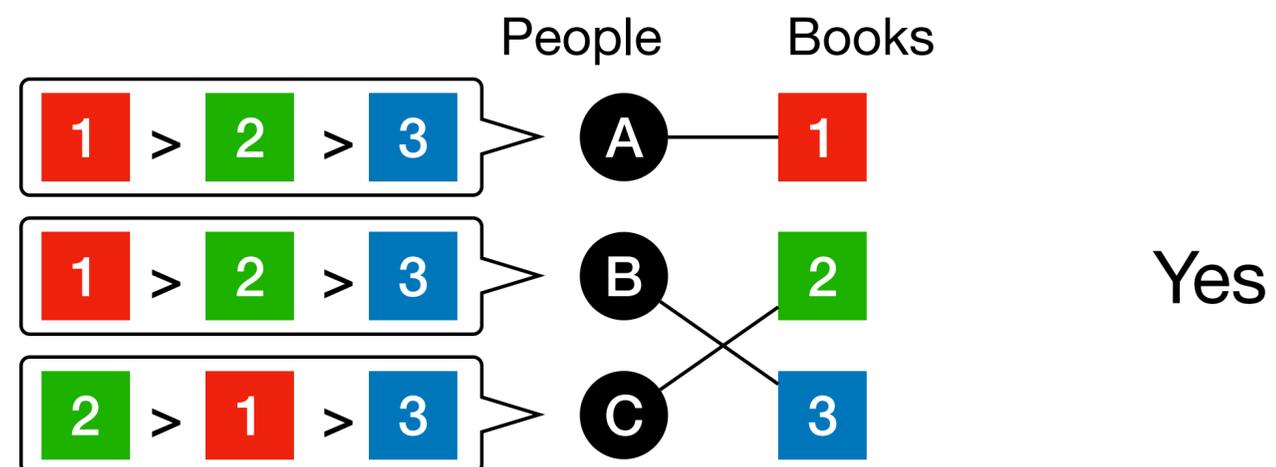
Is this outcome pareto optimal?



One Notion of Goodness

- An outcome is **pareto optimal** if you can't make someone better off without making someone else worse off.
- An outcome is pareto optimal if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

Is this outcome pareto optimal?



Is the Serial Dictatorship outcome always Pareto Optimal?

Is the Serial Dictatorship outcome always Pareto Optimal?

Recall

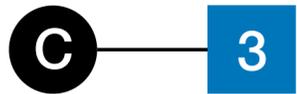
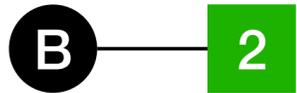
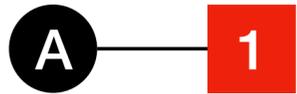
Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



**Outcome from
Serial Dictatorship**

Recall

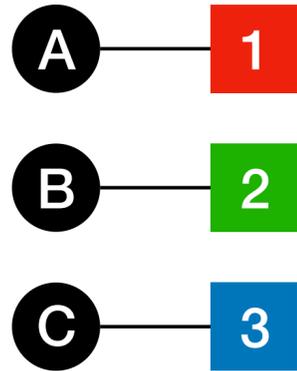
Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



Can we find a “better” outcome where everyone is just as happy, and some are happier?

**Outcome from
Serial Dictatorship**

Recall

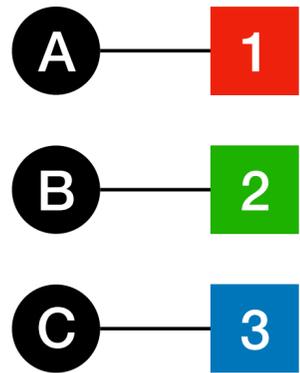
Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



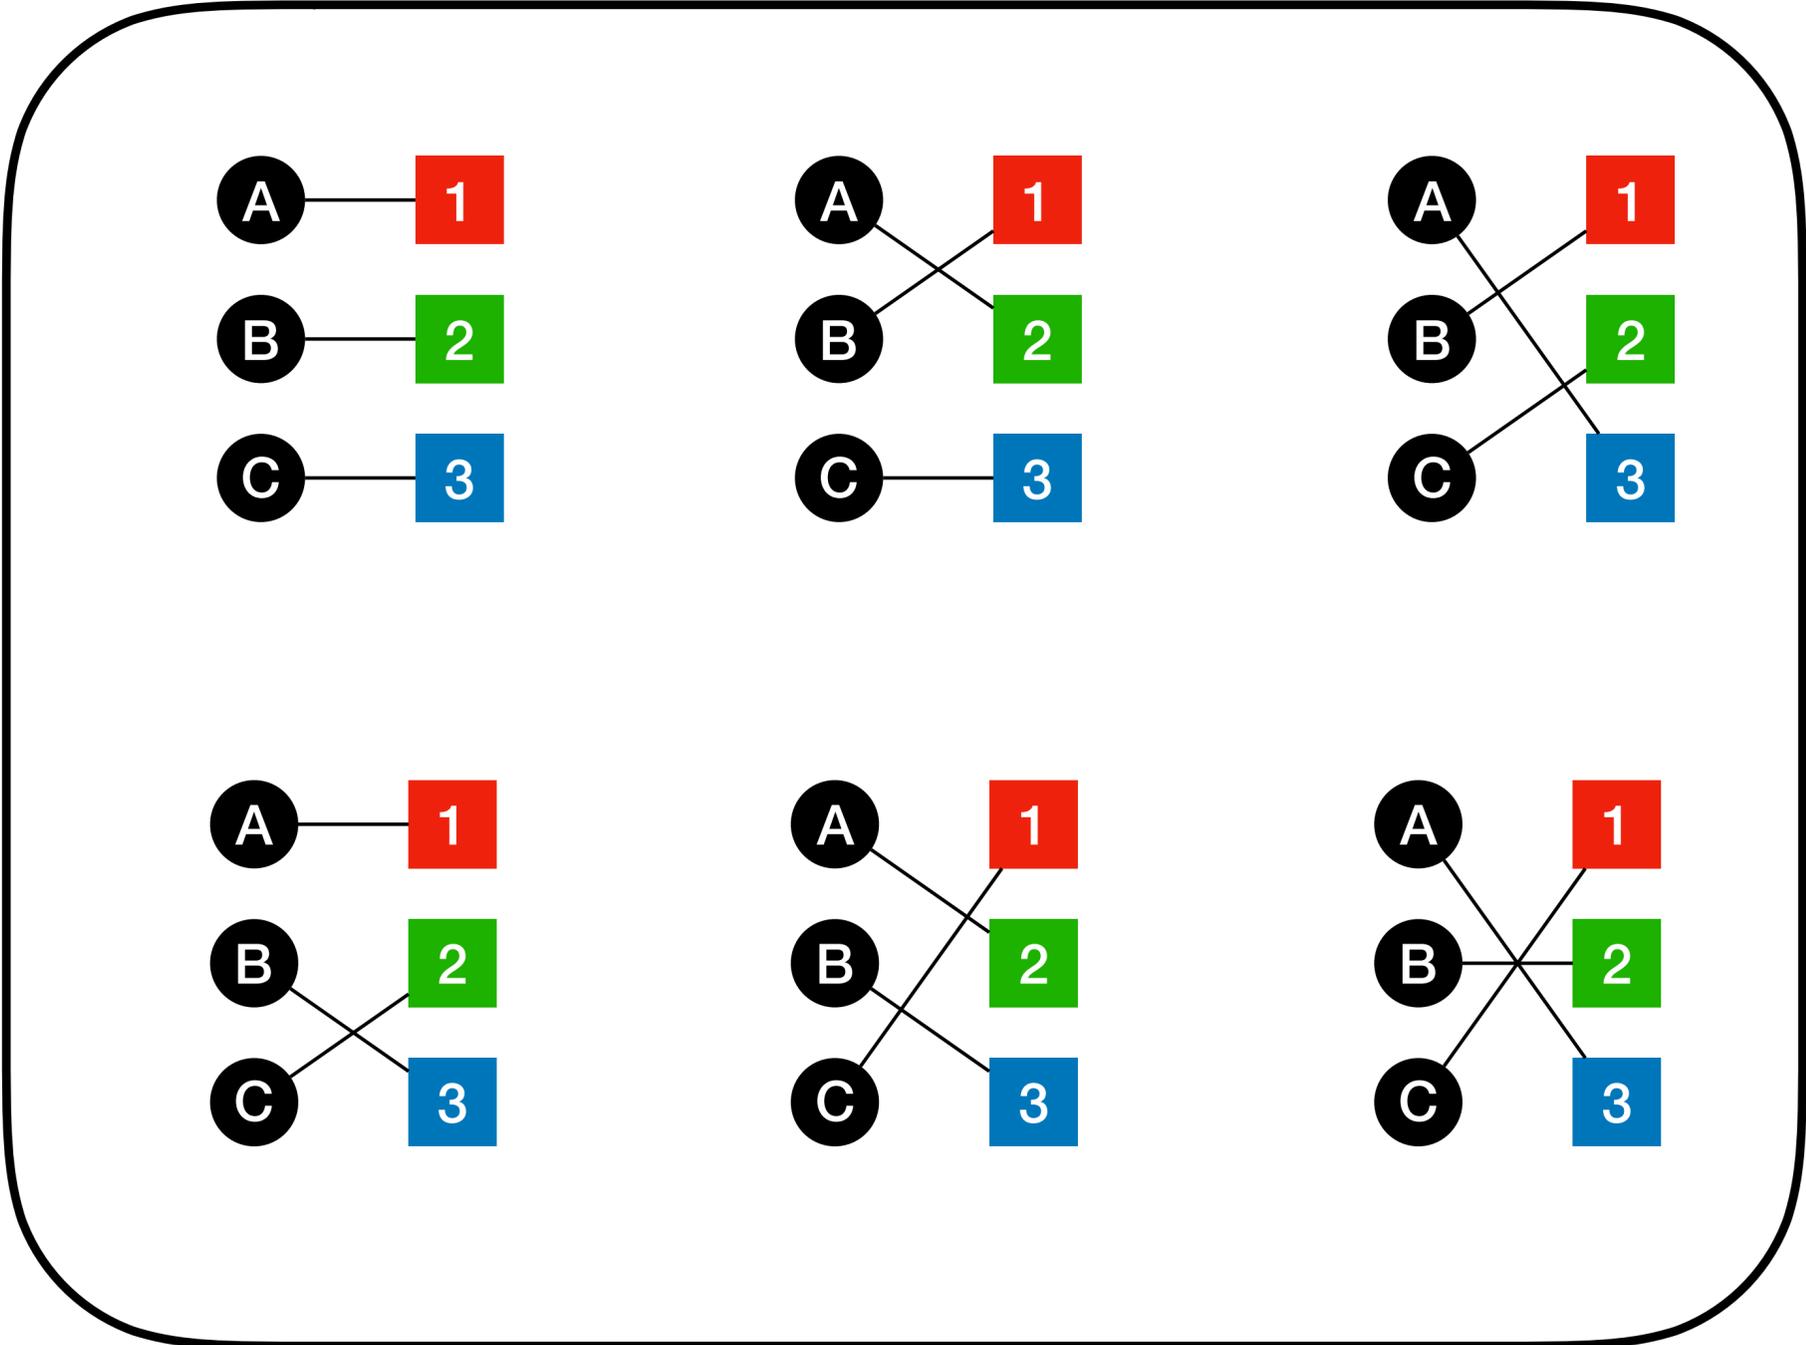
Can we find a “better” outcome where everyone is just as happy, and some are happier?

Outcome from Serial Dictatorship

Recall

Serial Dictatorship mechanism:
(1) Randomly order people. (2) Match each person with their most-preferred available book.

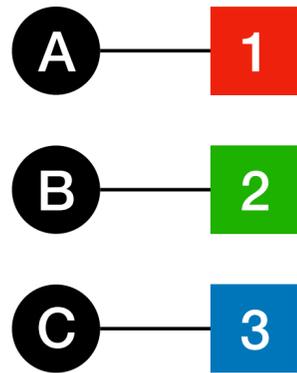
An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.



All possible outcomes
(where everyone is matched)

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



Can we find a “better” outcome where everyone is just as happy, and some are happier?

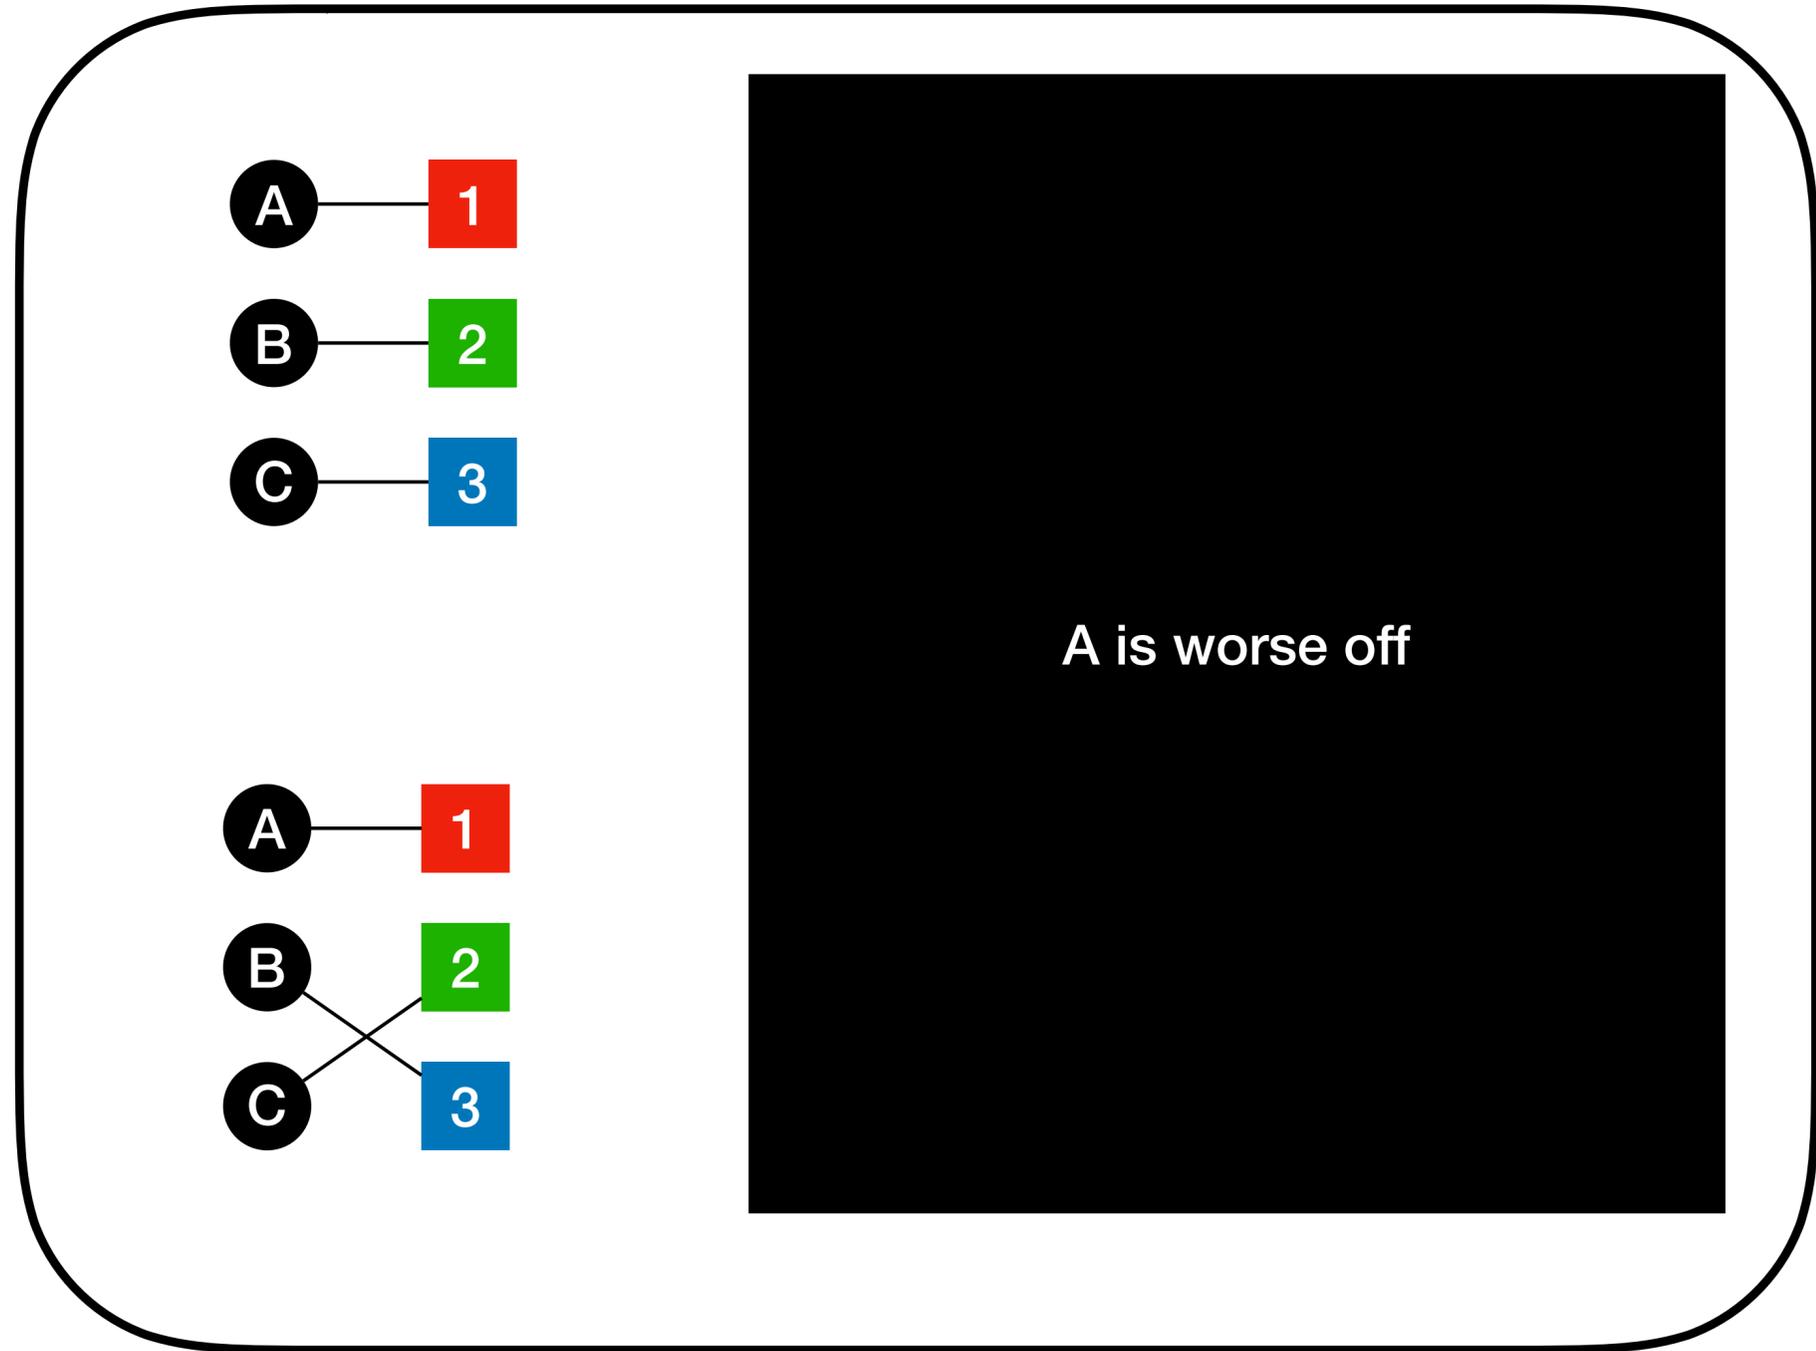
Outcome from Serial Dictatorship

Recall

Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

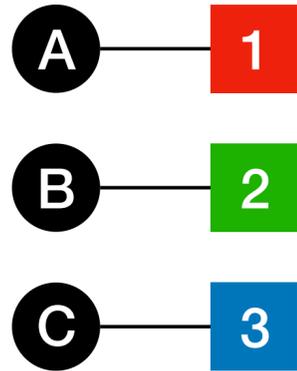
An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.



All possible outcomes (where everyone is matched)

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



Can we find a “better” outcome where everyone is just as happy, and some are happier?

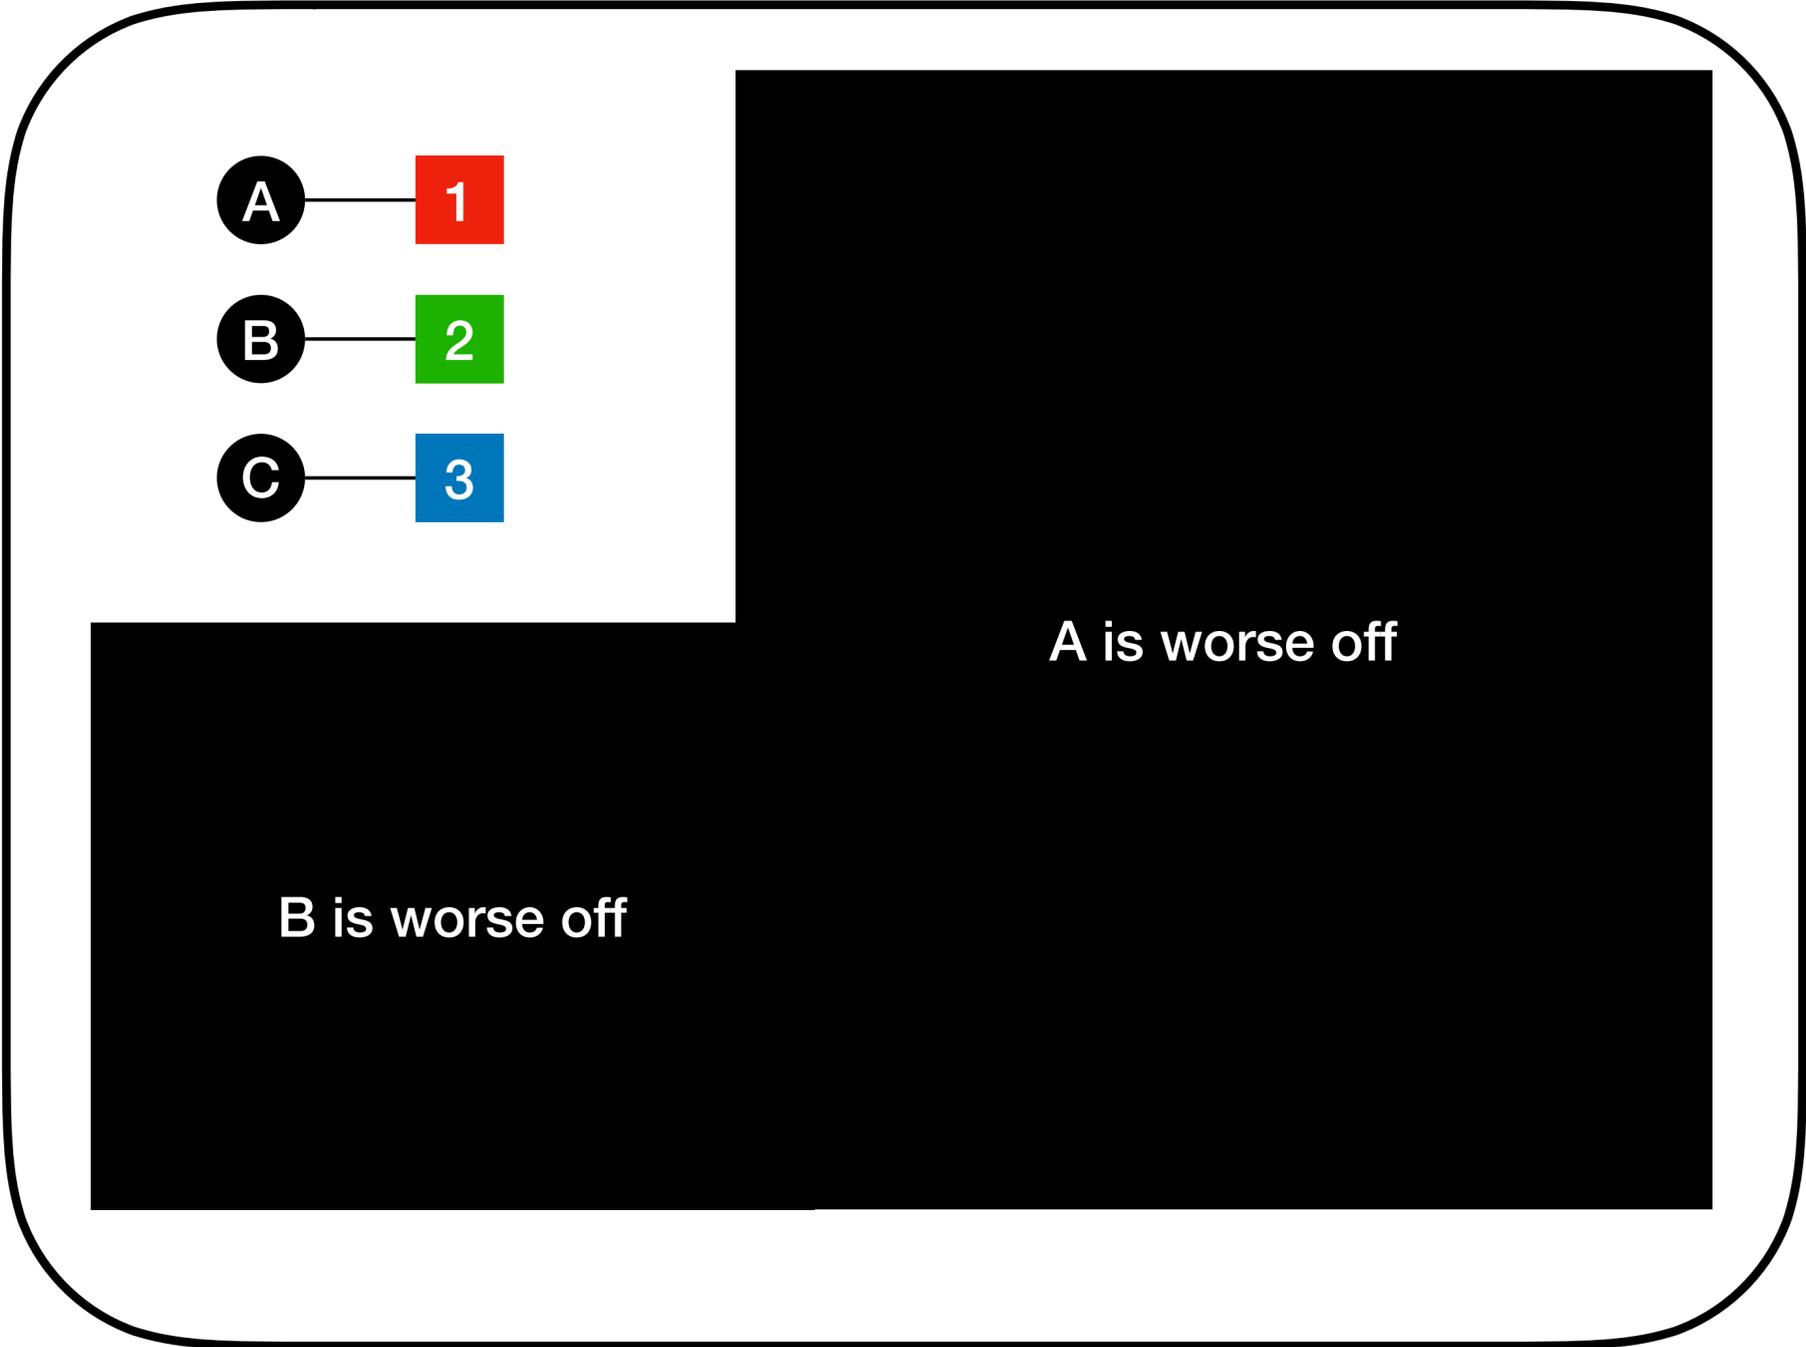
Outcome from Serial Dictatorship

Recall

Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

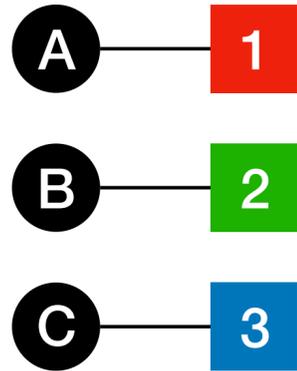
An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.



All possible outcomes (where everyone is matched)

Is the Serial Dictatorship outcome always Pareto Optimal?

People Books



Can we find a “better” outcome where everyone is just as happy, and some are happier?

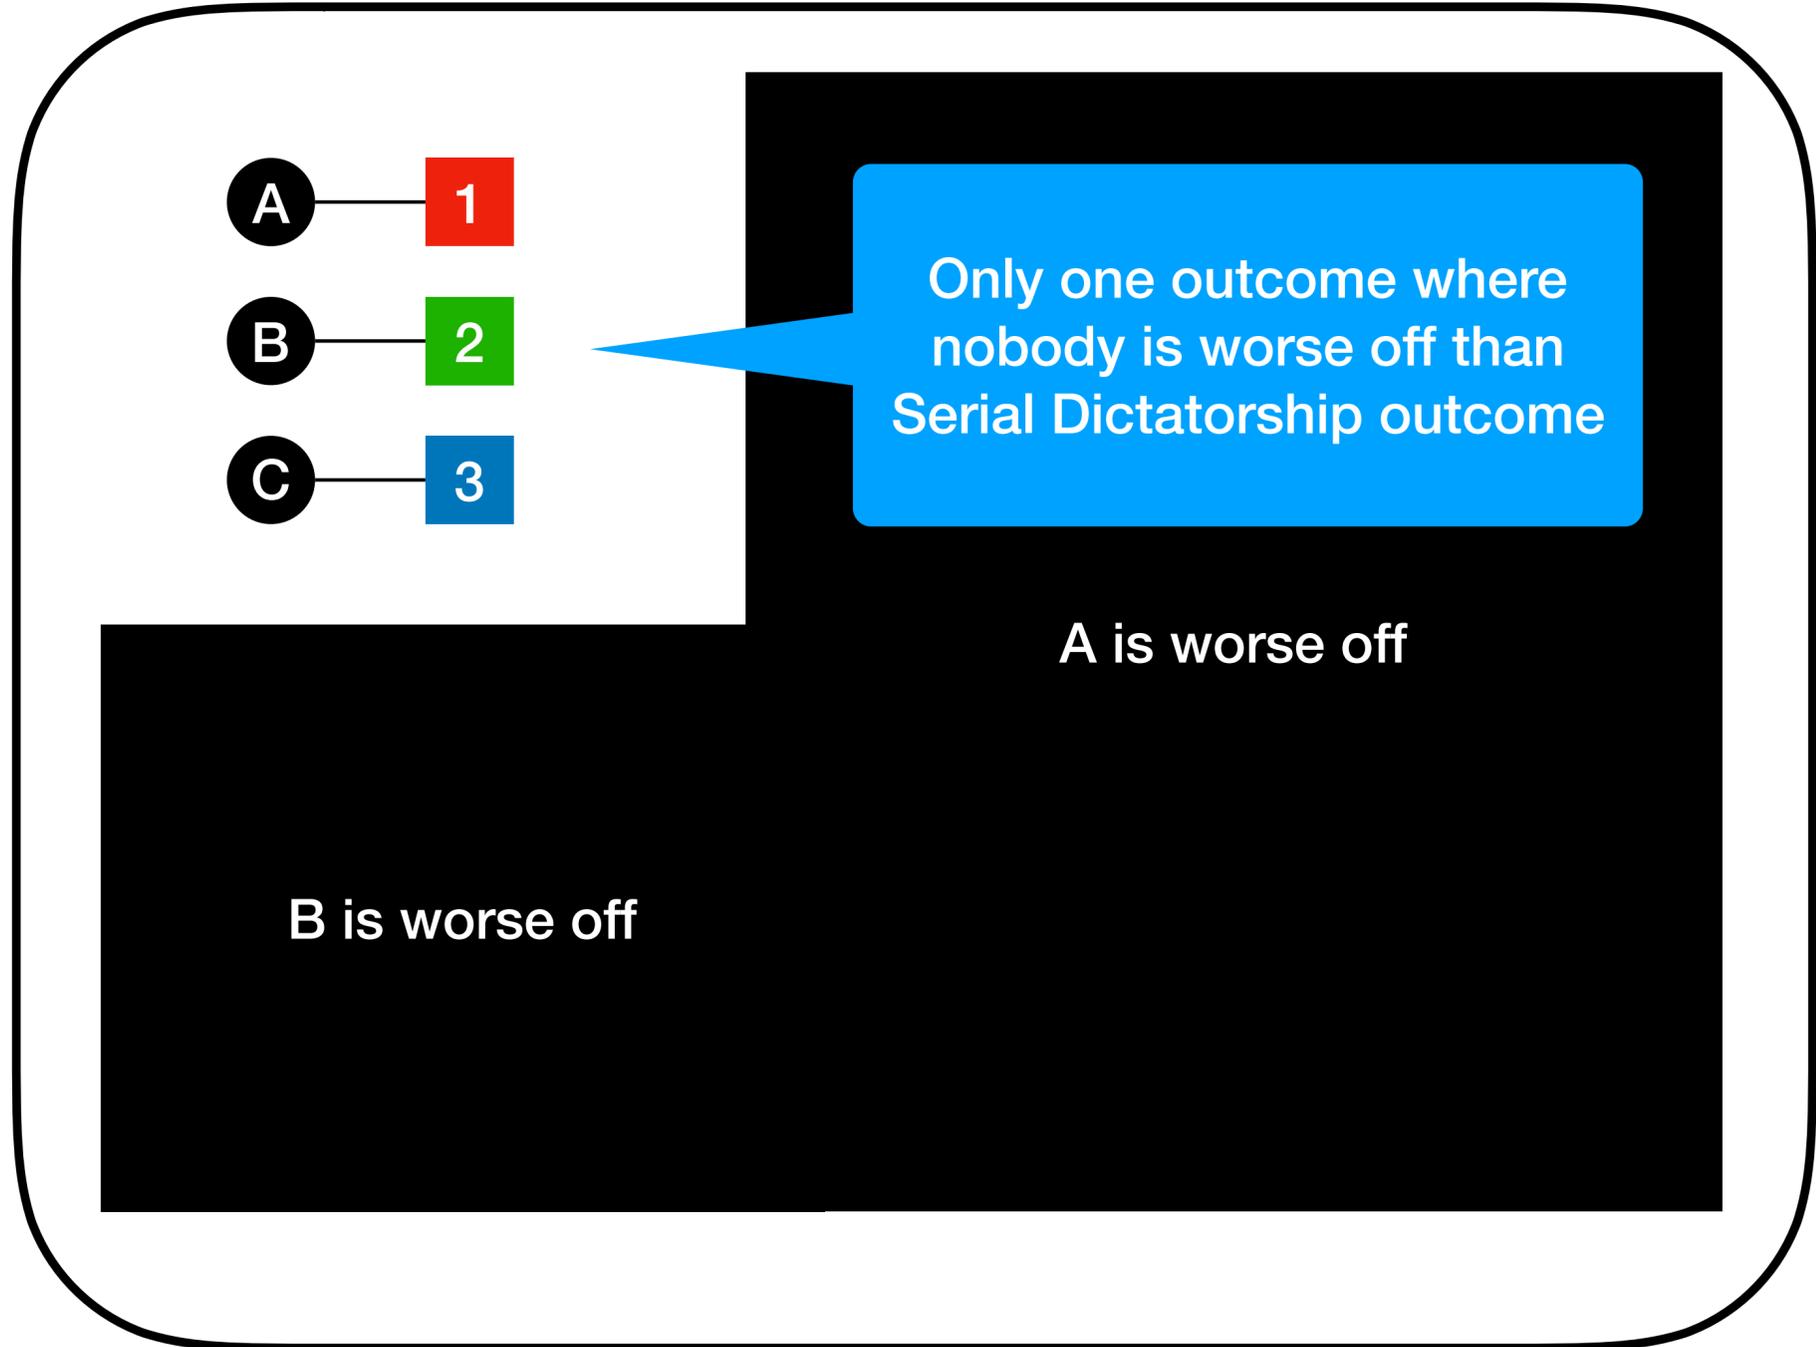
Outcome from Serial Dictatorship

Recall

Serial Dictatorship mechanism:

(1) Randomly order people. (2) Match each person with their most-preferred available book.

An outcome is **pareto optimal** if there is no alternative outcome for which (1) everyone is at least as happy, and (2) someone is happier.



All possible outcomes (where everyone is matched)

Another Notion of Goodness

- A mechanism is **strategyproof** if honesty is the best policy.
- A mechanism is strategyproof if lying about your preferences can't make you better off.

Is the Serial Dictatorship strategy proof?

Is the Serial Dictatorship strategy proof?

- Your reported book preferences don't affect your turn order.

Is the Serial Dictatorship strategy proof?

- Your reported book preferences don't affect your turn order.
- Your reported book preferences don't affect what anyone before you gets.

Is the Serial Dictatorship strategyproof?

- Your reported book preferences don't affect your turn order.
- Your reported book preferences don't affect what anyone before you gets.
- Thus, your reported book preferences don't affect what books are available on your turn.

Is the Serial Dictatorship strategy proof?

- Your reported book preferences don't affect your turn order.
- Your reported book preferences don't affect what anyone before you gets.
- Thus, your reported book preferences don't affect what books are available on your turn.
- Serial Dictatorship gives you the best available book on your turn (according to reported preferences).

Is the Serial Dictatorship strategy proof?

- Your reported book preferences don't affect your turn order.
- Your reported book preferences don't affect what anyone before you gets.
- Thus, your reported book preferences don't affect what books are available on your turn.
- Serial Dictatorship gives you the best available book on your turn (according to reported preferences).
- Thus, any misreport of preferences could only result in you getting a book you like less.

What if we swapped steps 1 and 2? Is it strategyproof?

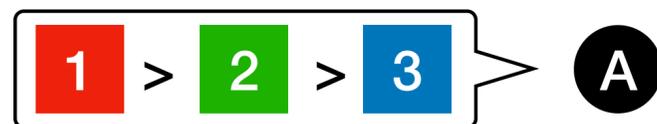
Serial Dictatorship

1. Have each participant submit their **preferences over books**.
2. **Randomly order attendees**: assign each a unique number 1 through 200.
3. In that attendee order, give each attendee **their favorite book** amongst those remaining, **according to their reported preferences**.

The Abridged Serial Dictatorship

Abridged Serial Dictatorship

1. Have each participant submit their preferences over their **top K** books.
2. Randomly order attendees.
3. In that order, give each attendee their favorite book amongst those remaining, according to their reported preferences. **If all their favorite books are gone, give them nothing.**



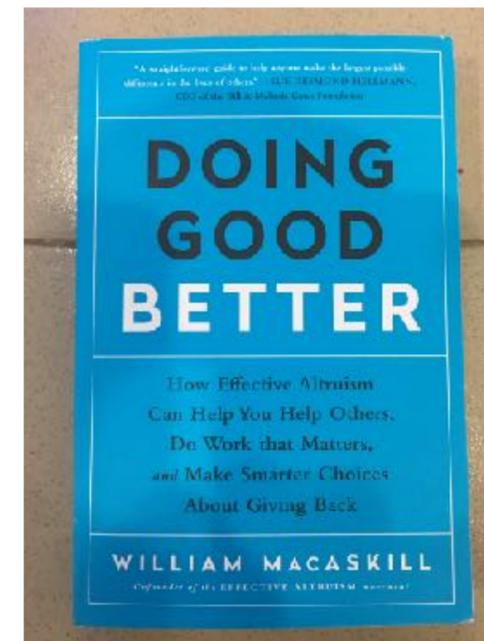
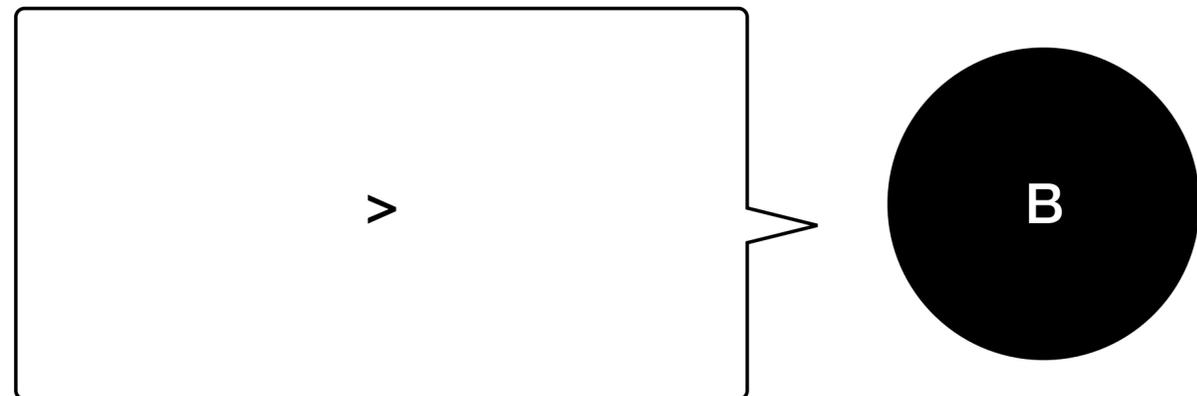
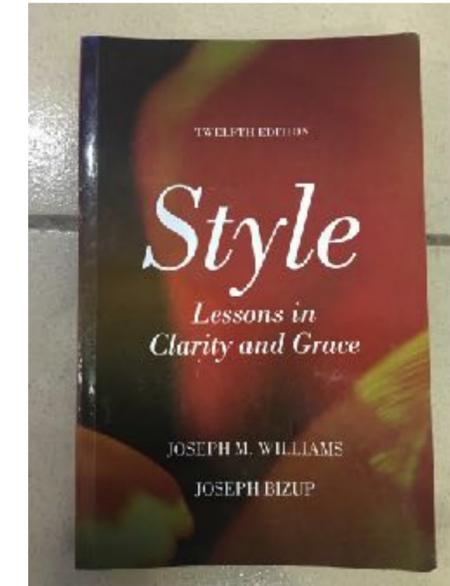
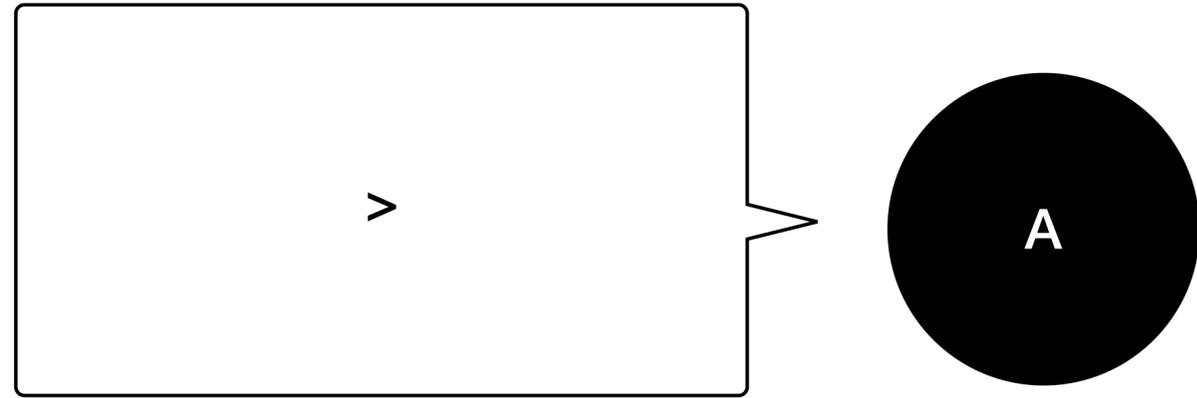
Serial dictatorship: full list



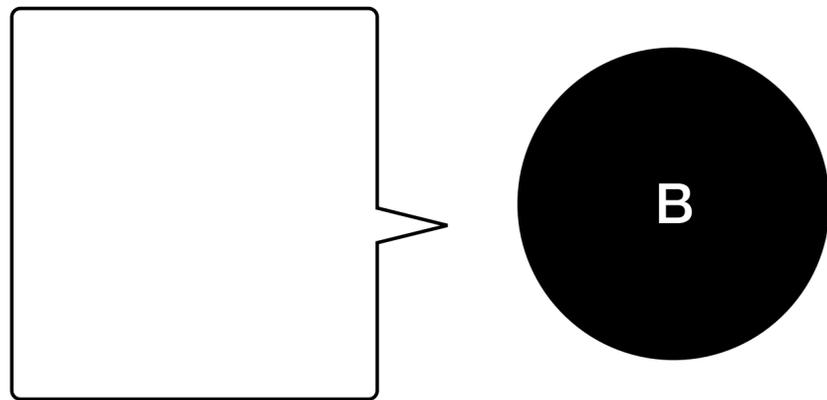
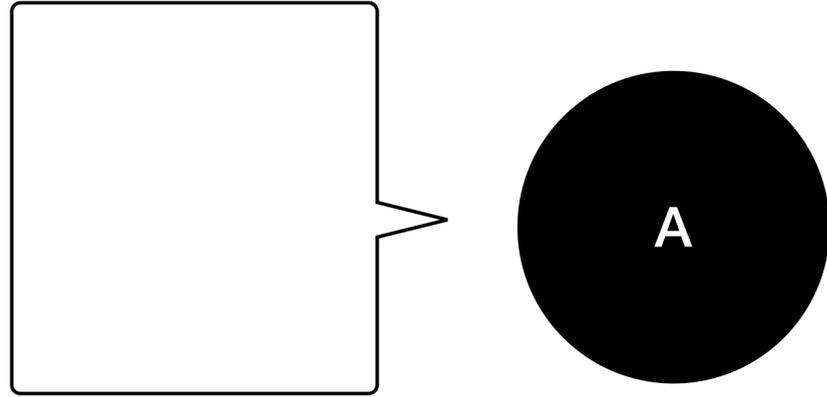
Abridged Serial dictatorship: truncated list

**Is the Abridged Serial Dictatorship strategyproof?
(example)**

Example 1: Serial Dictatorship



Example 2: Abridged Serial Dictatorship



Things to remember

- The rules of the game matter
- Small changes to the rules can make a big difference

Discussion

- Can you think of other systems that you regularly participate in that are clearly not strategyproof, or that do not produce Pareto optimal outcomes?
- What are some problems that feel similar to “book giving” problems? What’s the general version of the problem?

prediction markets

voting

online labor markets

two-sided matching

cryptocurrency

online learning

crowdsourcing

auctions

reputation systems

team formation

money-burning mechanisms

Next Step: Join Facebook Group

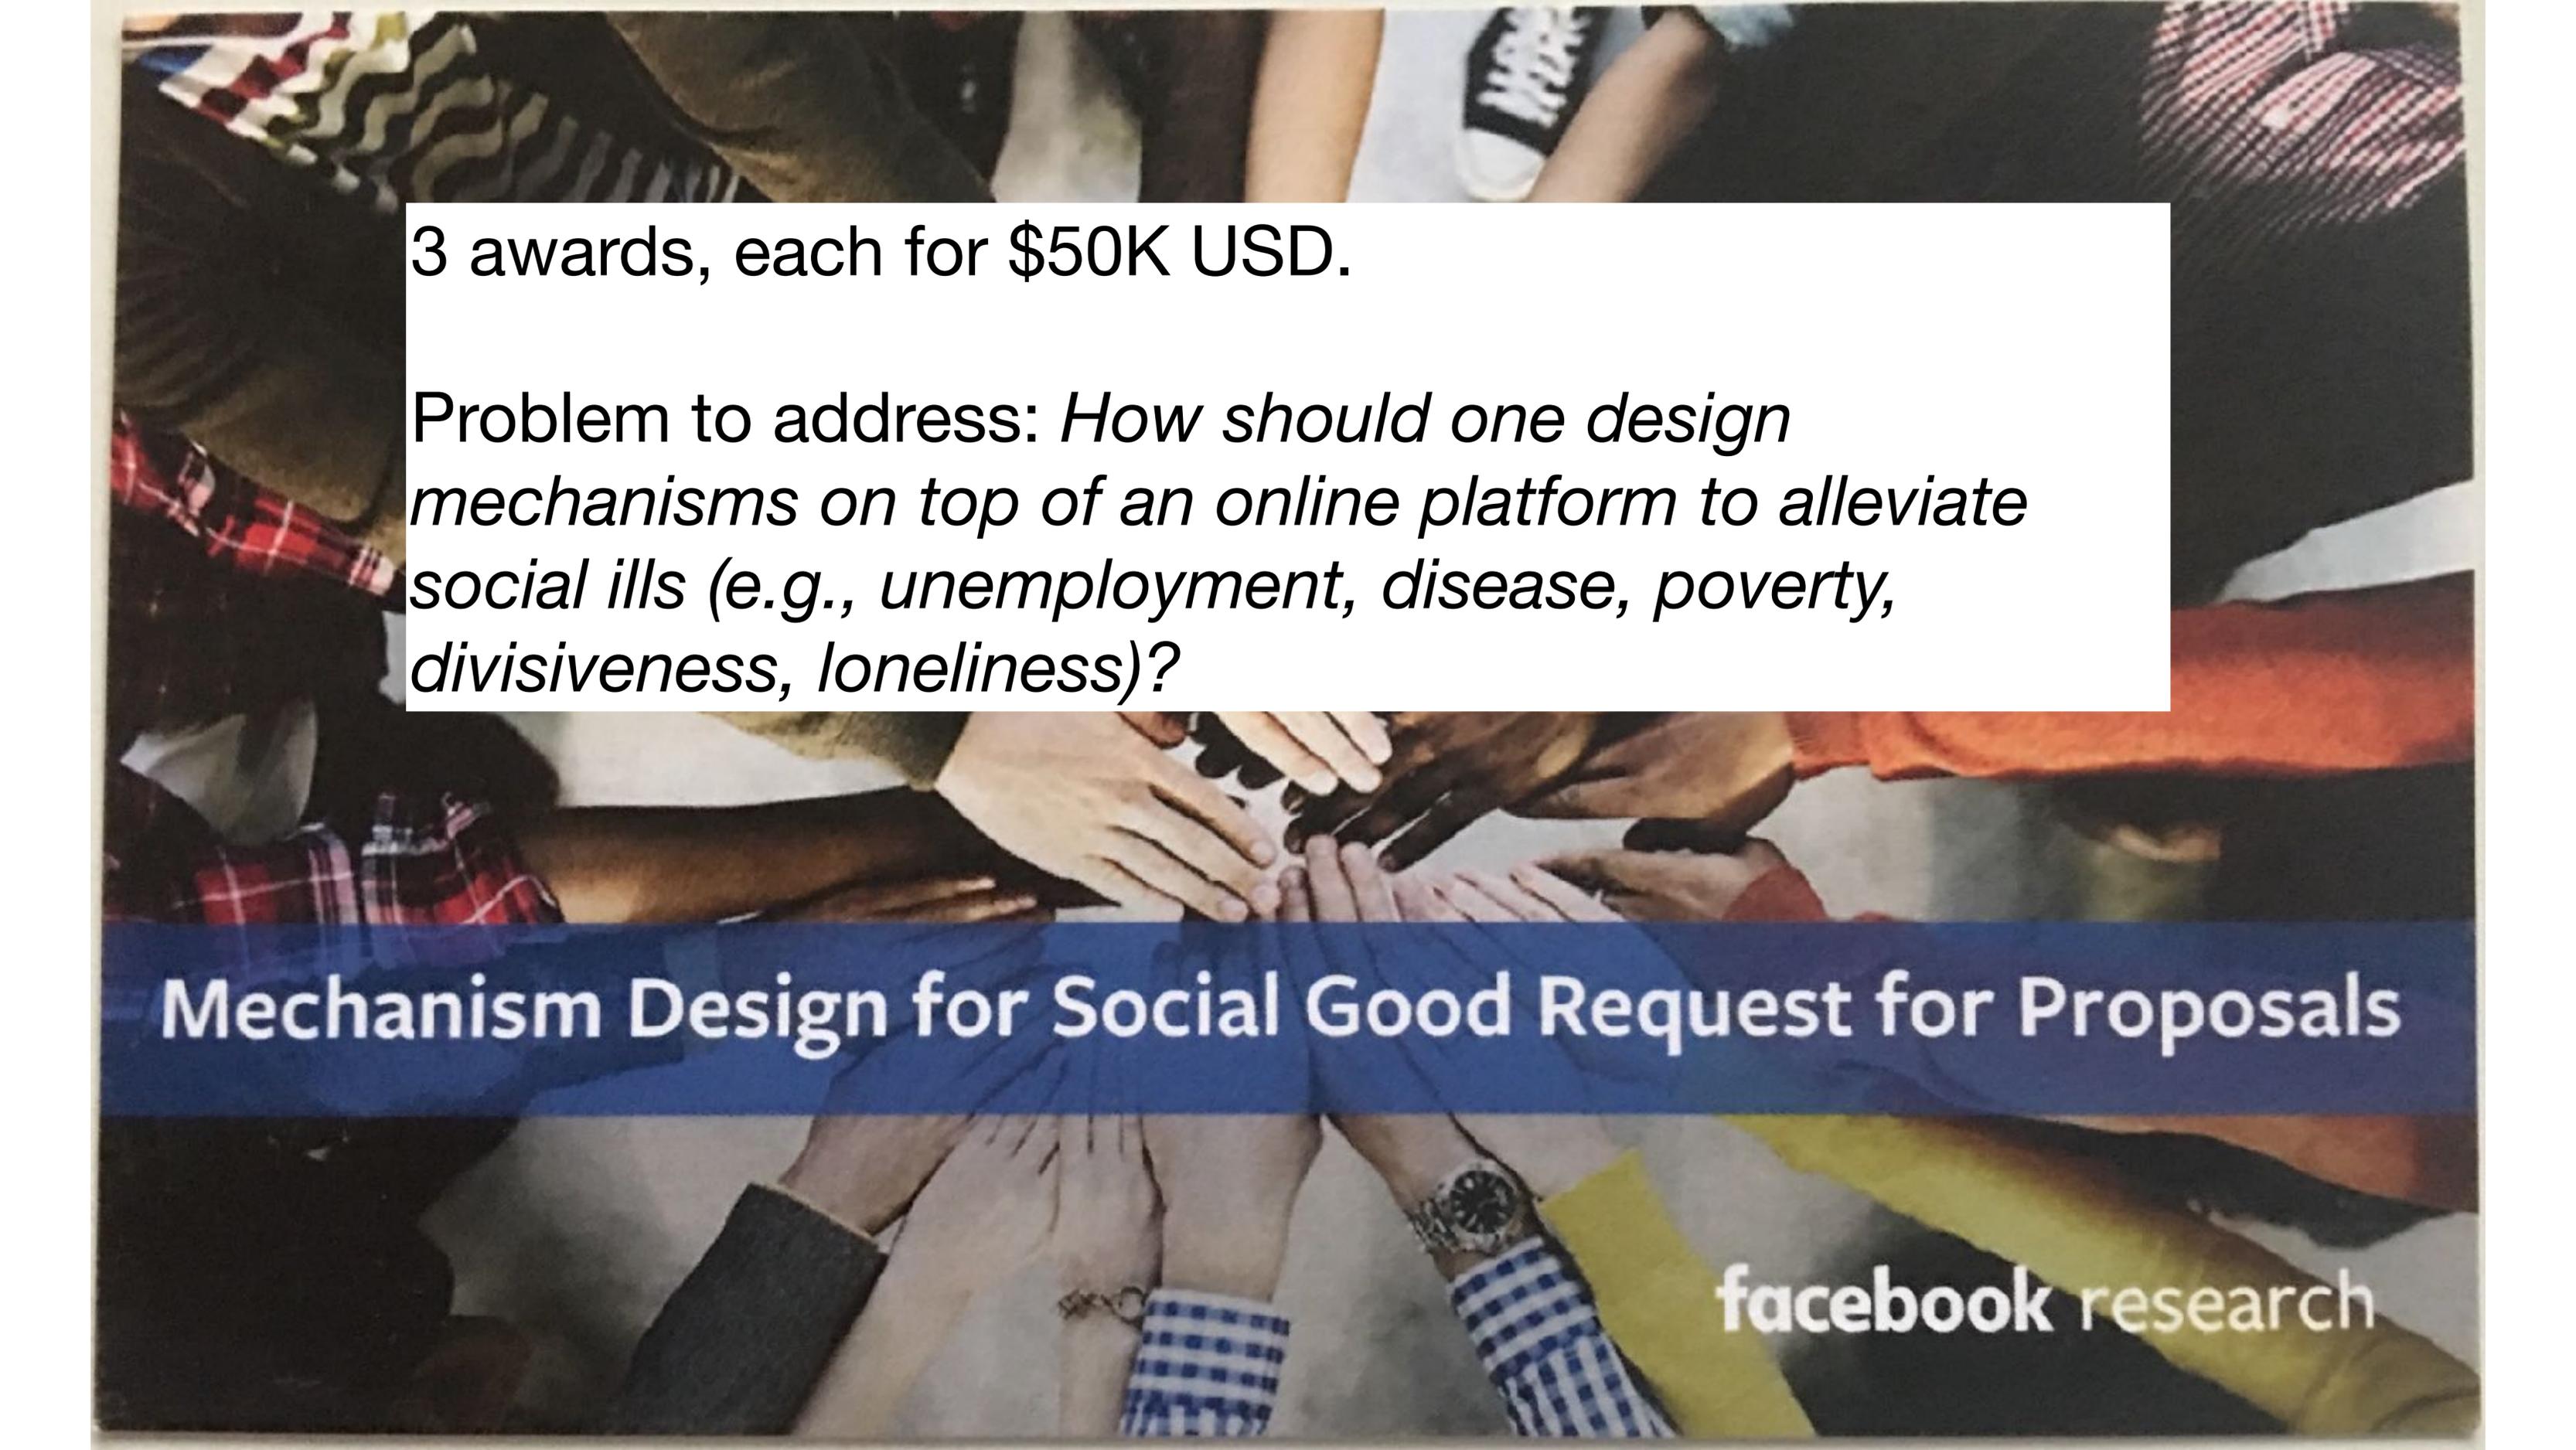
Mechanism Design @ Data Science Africa [goo.gl/bGFDhU](https://www.facebook.com/groups/bGFDhU)

- This talk based on first lecture of course by Tim Roughgarden: *Incentives in Computer Science*. Available online!
- Resources for learning more about mechanism design
- Resources for getting involved in the mechanism design community
- Announcements for Mechanism Design for Social Good proposals (\$150K in grants in 2018)

WINE 2018: The 14th Conference on Web and Internet Economics

December 15-17, 2018, Oxford, United Kingdom





3 awards, each for \$50K USD.

Problem to address: *How should one design mechanisms on top of an online platform to alleviate social ills (e.g., unemployment, disease, poverty, divisiveness, loneliness)?*

Mechanism Design for Social Good Request for Proposals

facebook research

Next Step: Join Facebook Group

Mechanism Design @ Data Science Africa [goo.gl/bGFDhU](https://www.facebook.com/groups/bGFDhU)

- This talk based on first lecture of course by Tim Roughgarden: *Incentives in Computer Science*. Available online!
- Resources for learning more about mechanism design
- Resources for getting involved in the mechanism design community
- Announcements for Mechanism Design for Social Good proposals (\$150K in grants in 2018)