CS 357: Algorithmic Game Theory

Lecture 9: Markets Without Money

Shikha Singh

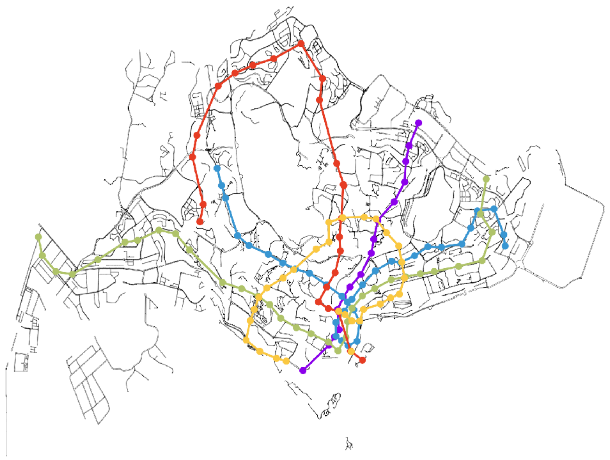

Announcements

- **HW 4** due today
- Lecture notes reviewing topics so far posted
  - Have not been reviewed thoroughly, so there may be typos
  - Double check with lectures and let me know
- **Exam I** on **Friday March 14**
  - Short-"ish" questions, mostly HW style questions with one or two open/ended answers or proofs
  - Cover everything until last week: review HWs, assignments, readings and lectures to prepare
  - Closed book but can bring prepared notes (no more than 5 pages)

Spring Break

Week 6: Centralized Markets w/o Money

Week 5: Centralized Markets w/o Money & Exam I

Week 4: VCG and Sponsored Search Auctions

Week 3: Myerson & Single Parameter Auctions

Week 2: Single-Item Auctions

Week 1: Game Theory

Announcements: Looking Ahead

- Paper Eval #2 next week (March 21)
  - In groups like last time
  - If planning to miss, reach out
- Assignment 3 due week after Spring Break (April 11)
  - No work due over break
  - Expected to finish the week after

Spring Break

Week 6: Centralized Markets w/o Money Pt 2

Week 5: Centralized Markets w/o Money & Exam 1

Week 4: VCG and Sponsored Search Auctions

Week 3: Myerson & Single Parameter Auctions

Week 2: Single-Item Auctions

Week 1: Game Theory

Story So Far: Mechanism Design w Money

Centralized: transactions are decided by a central hub

Centralized Markets

Money Transfer

Goal. Align global objectives (social optimal outcome) with participant objective (maximize utility) using payments.

Next: Mechanism Design without Money

Centralized Markets

~~Money Transfer~~

We will revisit role of money in **decentralized systems**.

Later: Decentralized Systems

Selfish behavior in decentralized markets with or without money

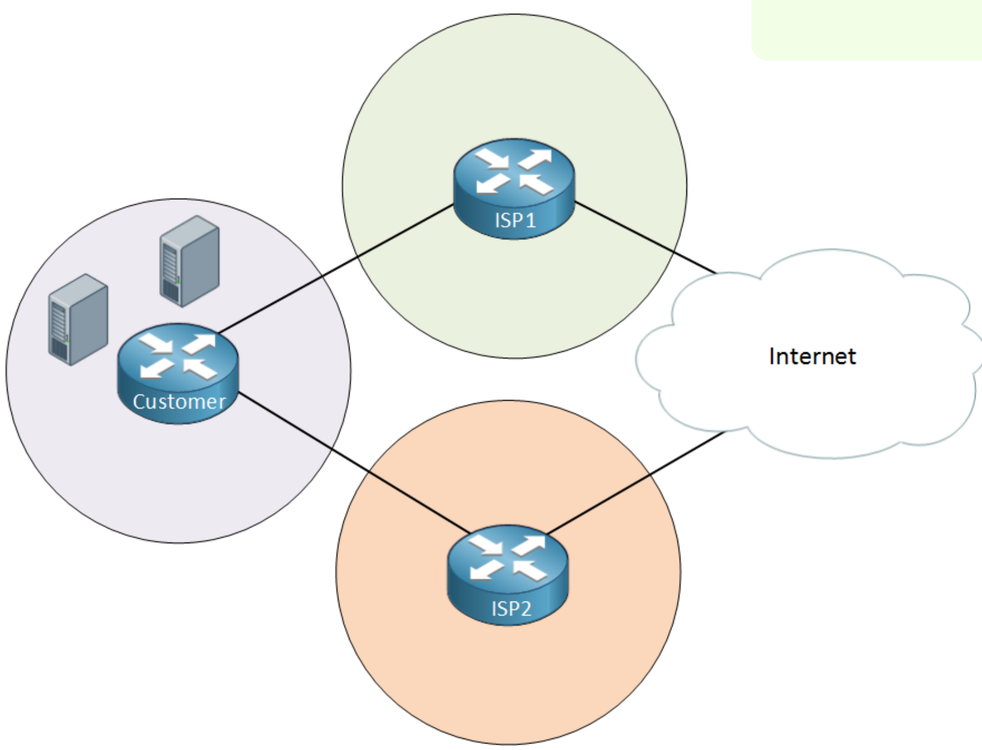

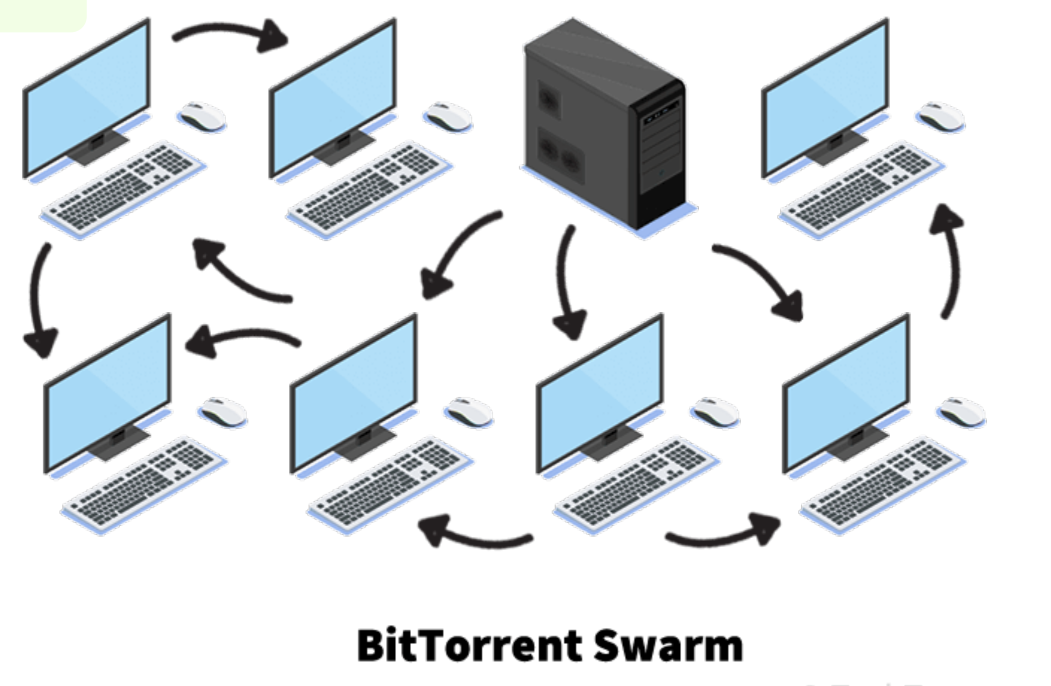

Markets without Money

Designer's Goal: Allocate items to ensure good global guarantees (e.g. welfare)

Agent's Goal: Report **private preferences** so as to maximize their utility.

n agents with private preferences

Goods to be allocated

Markets without Money

Designer's Goal: Allocate items to ensure good global guarantees (e.g. welfare)

Agent's Goal: Report **private preferences** so as to maximize their utility.

Goods to be allocated

n agents with private preferences

Payments so far were a way to enforce strategyproof mechanisms

Markets without Money

Designer's Goal: Allocate items to ensure good **global guarantees**

Agent's Goal: Report **private preferences** that achieve **the best outcome**

Goods to be allocated

n agents with private (**ordinal**) preferences

What are good global guarantees when agents have "ordinal" utilities?

Markets without Money

- Many domains money transfer is either infeasible/inappropriate/illegal
- Problem domains without money?
  - Matching students to courses
  - Matching students to school/ colleges/ dorms
  - Matching doctors to hospitals
- Sharing resources or barter markets:
  - Exchanging goods or services
- Social decision making:
  - Voting to elect a leader, a committee or an outcome

Domain of AGT where TCS truly shines!

Markets without Money

One Sided Markets

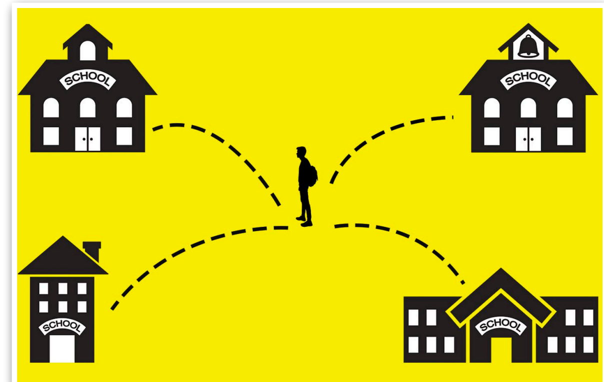

No exchange

Two Sided Markets

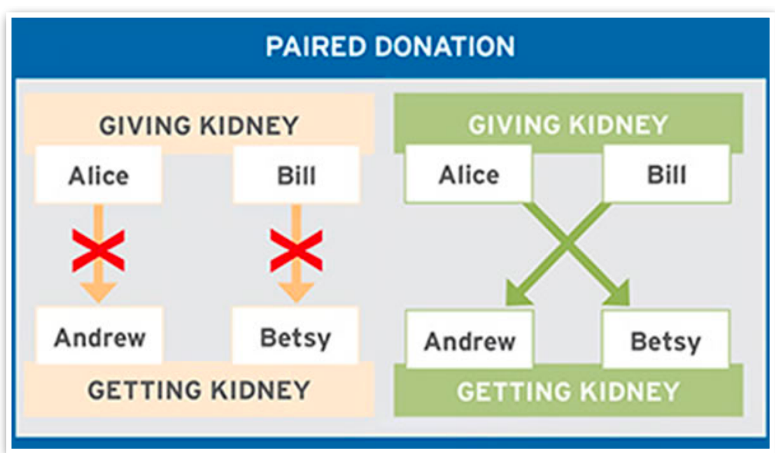

Exchange based

THE MATCH®
NATIONAL RESIDENT MATCHING PROGRAM®

One-Sided Market: Assignment Problems

Designer's Goal: Assignment of items to agents is **Pareto optimal**

Agent's Goal: Report **private preferences** that achieve **the best outcome**

n agents with ordered preferences

College Dorms

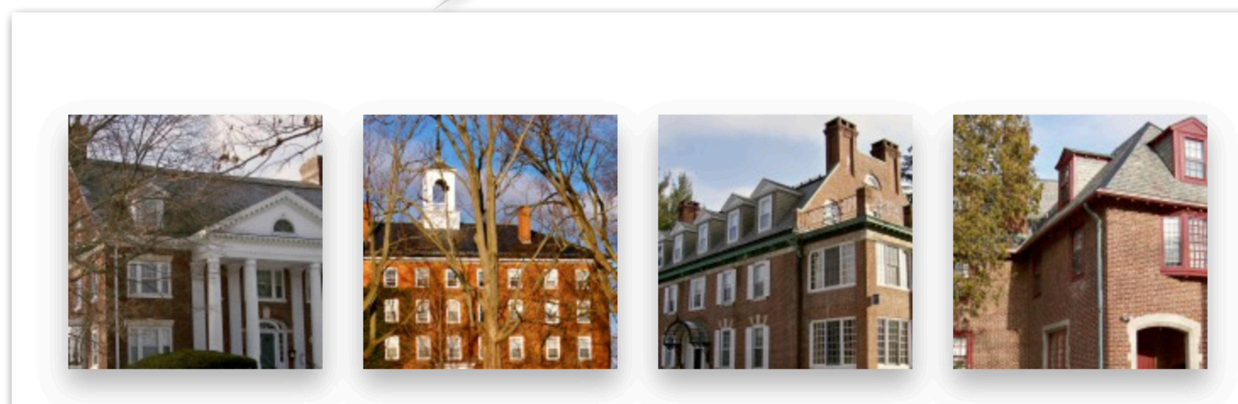

Pareto optimality: An outcome O is Pareto optimal if there is no outcome O' and where every agent does as well as in O and some agent does strictly better.

One-Sided Assignment Problem

- One-sided allocation or assignment problems:
  - Assigning students to dorms
  - Offices to employees
  - Tasks to volunteers
- Model. We have n agents and m items
  - Agents have strict preference ordering over the items
- Feasible assignment: matching between items and agents
- Goal. Find a Pareto optimal assignment (means no other assignment can make an agent better off without making another agent worse off)

One-Sided Market: Assignment Problems

Designer's Goal: Assignment of items to agents is **Pareto optimal**

Agent's Goal: Report **private preferences** that achieve **the best outcome**

n agents with ordered preferences

College Dorms

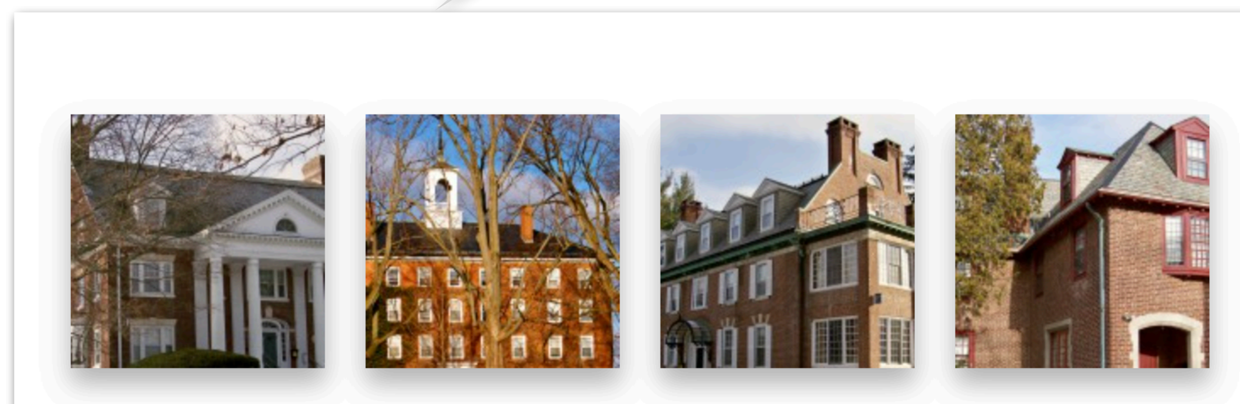

Discussion Question. How is this typically done based on your experience?

One-Sided Matching Example

1


2
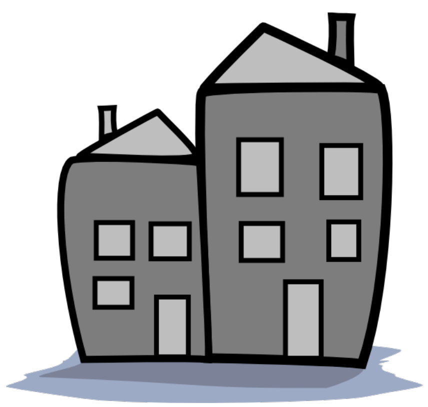

3


How do we matching students to dorms?

Private Preferences

Zoe


1 > 3 > 2

Chris


1 > 2 > 3

1 > 2 > 3

One-Sided Assignment Problem

- Most housing allocation algorithms look something like this:
  - Asks agents to report their preferences over items
  - Choose an ordering of all agents (**lottery order**)
    - Often based on some metrics are considered "fair", e.g., seniority, years of service to college, family size, etc
- Go in order, assign each agent their favorite item that is still remaining
- **Example.** Faculty housing, mini lottery for dorms (at Williams)
- This is a good mechanism?
  - Strategyproof, Pareto optimal?

Serial Dictatorship (SD)

- Each of the n agents submit a ranked ordering over m items
- Each agent is assigned a "lottery index" from $\{1,2,\ldots,n\}$
- For $i = 1,2,\ldots,n$
  - Assign i their favorite choice among options still available
- **Lemma.** SD algorithm is strategyproof & Pareto optimal.

Serial Dictatorship (SD)

- Each of the n agents submit a ranked ordering over m items
- Each agent is assigned a "lottery index" from $\{1,2,\ldots,n\}$
- For $i = 1,2,\ldots,n$
  - Assign i their favorite choice among options still available
- **Lemma.** SD algorithm is dominant strategyproof & Pareto optimal.
- Proof Outline (Truthful reporting is dominant strategy).
  - Cannot control lottery order
  - Given lottery order, truthful reporting obtains the best possible outcome
  - No incentive to deviate (regardless of other's preferences)

Serial Dictatorship (SD)

- Each of the n agents submit a ranked ordering over m items
- Each agent is assigned a "lottery index" from $\{1,2,\ldots,n\}$
- For $i = 1,2,\ldots,n$
  - Assign i their favorite choice among options still available
- **Lemma.** SD algorithm is dominant strategyproof & Pareto optimal.
- Proof Outline (Why Pareto optimal).
  - Idea: show no other assignment can Pareto dominate
  - Any other assignment must make some agent worse off nces)

Serial Dictatorship (SD)

- **Lemma.** SD algorithm produces the **unique** Pareto optimal outcome.
- Proof. Let M' be an assignment where no agent is worse off than in M
  - If any agent is worse off in M' it cannot Pareto dominate M!
  - **Claim**: Any such M' is identical to M
  - Suppose M' is the same as M until step $i = k$
  - Consider agent at step $i = k + 1$, M gives i their favorite among remaining items
  - M' must do the same to make them not worse off
  - Thus M is the unique Pareto optimal outcome

Takeaways

- Serial dictatorship seems great: Pareto optimal and strategyproof
- Any criticism?
  - Can be unfair if priority order between agents is not natural
- **Random-serial-dictatorship** (RSD)
  - Run SD on a ranked ordering that is sampled uniformly at random from all possible orderings

Variants: Shortlists

- What happens if we restrict the # items each agent can rank?
    - Course registration (can only preregister for so many courses)
    - Truthfulness is no longer a dominant strategy
    - Preferences now depending on the order in the lottery
- Strategizing is now all about guessing the lottery order & other's preferences

Variants: House Allocation with Tenants

- What if students have to give up their current dorm to participate
  - (Individually rational.) Participant's utility of outcome is at least as. much as if they did not participate
  - Is this individually rational if participants already have a house?

Variants: School Choice with Different Priorities

School Choice: A Mechanism Design Approach

By ATILA ABDULKADİROĞLU AND TAYFUN SÖNMEZ*

A central issue in school choice is the design of a student assignment mechanism. Education literature provides guidance for the design of such mechanisms but does not offer specific mechanisms. The flaws in the existing school choice plans result in appeals by unsatisfied parents. We formulate the school choice problem as a mechanism design problem and analyze some of the existing school choice plans including those in Boston, Columbus, Minneapolis, and Seattle. We show that these existing plans have serious shortcomings, and offer two alternative mechanisms each of which may provide a practical solution to some critical school choice issues. (*JEL* C78, D61, D78, I20)

School choice is one of the widely discussed topics in education.[1] It means giving parents the opportunity to choose the school their child will attend. Traditionally, children are assigned to public schools according to where they live. Wealthy parents already have school choice, because they can afford to move to an area with good schools, or they can enroll their child in a private school. Parents without such means, until recently, had no choice of school, and had to send their children to schools assigned to them by the district, regardless of the school quality or appropriateness for the children. As a result of these concerns, *intra-district* and *inter-district* choice programs have become increasingly popular in the past ten years.[2] Intra-district choice allows parents to select schools throughout the district where they live, and inter-district choice allows them to send their children to public schools in areas outside their resident districts. In 1987, Minnesota became the first state to oblige all its districts to establish an inter-district choice plan (Allyson M. Tucker and William F. Lauber, 1995). Today,

* Abdulkadiroğlu: Department of Economics, Columbia University, New York, NY 10027 (e-mail: aa2061@columbia.edu); Sönmez: Department of Economics, Koç University, Sarıyer, 80910, İstanbul, Turkey (e-mail: tsonmez@ku.edu.tr). The previous version of this paper was entitled "School Choice: A Solution to the Student Assignment Problem." We are grateful to Michael Johnson for a discussion that motivated this paper. We thank Elizabeth Caucutt, Steve Ching, Julie Cullen, Dennis Epple, Roger Gordon, Matthew Jackson, Tarık Kara, George Mailath, Paul Mil-

Paper Eval #2: Case Study of School Choice

- Will post discussion and analysis questions
- Work in groups and discuss/present next Friday (March 21)
  - Last Friday before Spring Break
  - If planning on missing class, reach out in advance
  - Schedule a time to meet to discuss/present one on one

One-Sided vs Two-Sided

- Schools/colleges may not have "preferences" like individuals
- But they may have "priorities" (milder requirement)
    - Based on ranking on a standardized test
    - Based on institutional priorities
    - Based on distance/socio-economic states, etc
- Still, frequently two-sided algorithms like deferred acceptances is used

School Choice Algorithms

- Most common algorithm: deferred acceptance
- Others: variants of serial dictatorship
- Not used: top-trading cycle
  - As you read the paper, think about why (we'll discuss)
- Where it has proved extremely useful:
  - Exchange markets like kidney exchange

Two-Sided Matching Markets

Two-Sided Markets

- Consider a two-sided market:
  - A set H of n hospitals, a set S of n students
  - Each hospital has a complete and strict preference ranking of students
  - Each student has a complete and strict preference ranking of hospitals
- **Goal.** A perfect matching M that is **stable** (has no blocking pairs)
  - A hospital h and student s form a **blocking pair** (h, s) in a matching M if h prefers s to its current match in M and s prefers h to its current match in M

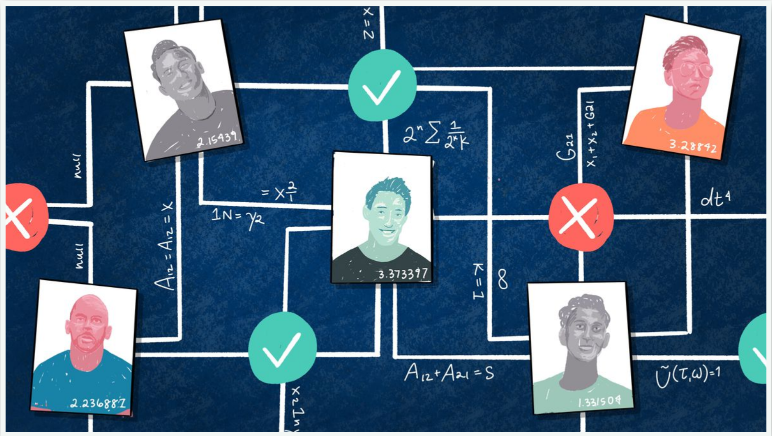

The Tinder algorithm explained: Vox

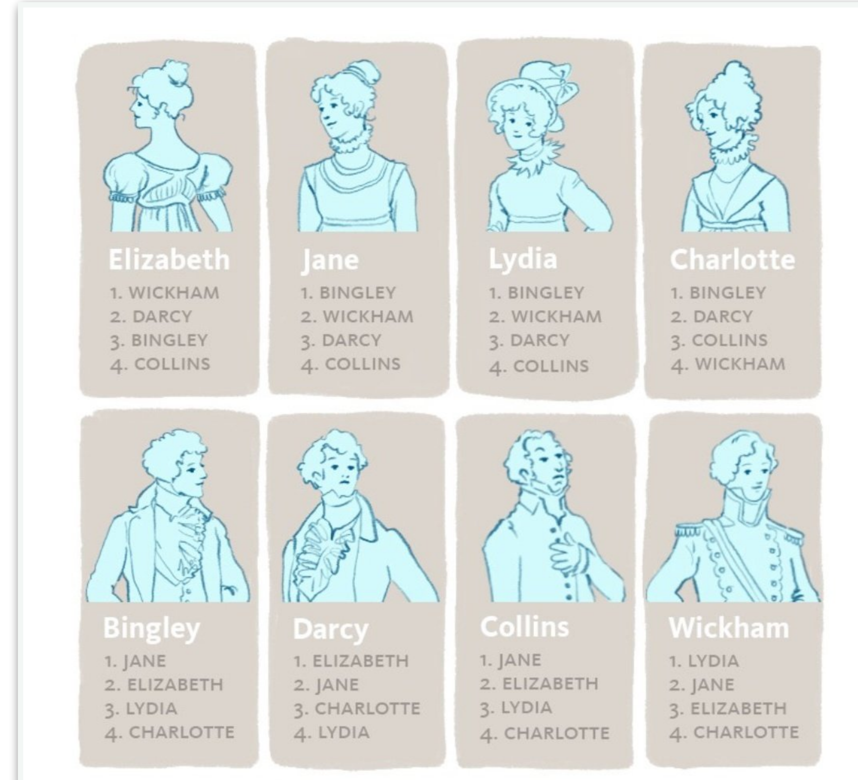

Stylized History:

the **"Stable Marriage"** Problem

The Dating Market: Medium

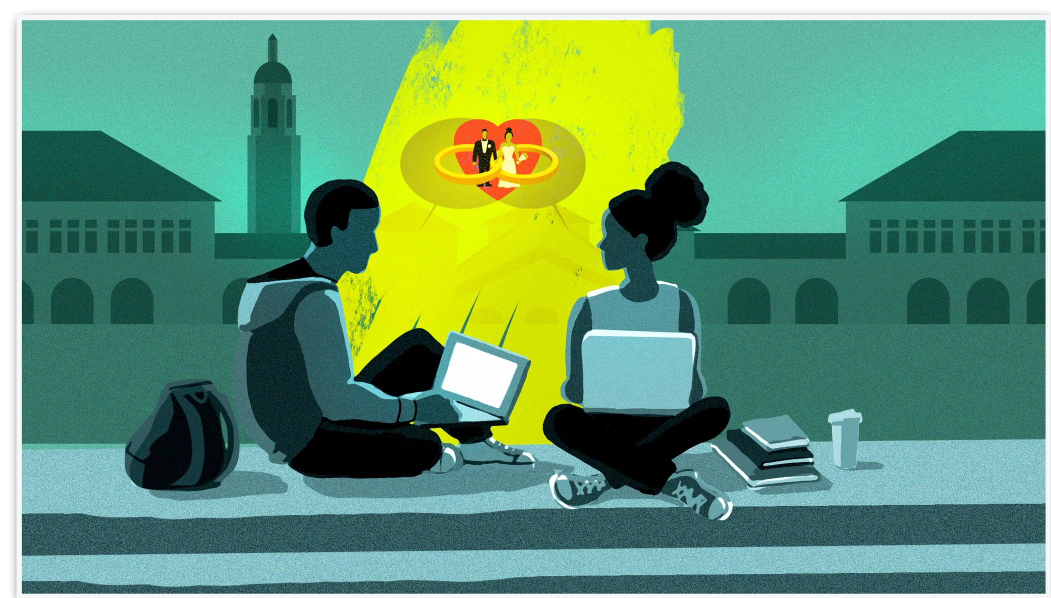

Dating apps are awful. But this algorithm offers just one match: your "backup plan." - Vox

Stylized Model of "Marriage" or "Dating"

1962, The American Mathematical Monthly

COLLEGE ADMISSIONS AND THE STABILITY OF MARRIAGE

D. GALE* AND L. S. SHAPLEY, Brown University and the RAND Corporation

3. Stable assignments and a marriage problem. In trying to settle the question of the existence of stable assignments we were led to look first at a special case, in which there are the same number of applicants as colleges and all quotas are unity. This situation is, of course, highly unnatural in the context of college admissions, but there is another "story" into which it fits quite readily.

1992

Stable Husbands

Donald E. Knuth, Rajeev Motwani, and Boris Pittel
Computer Science Department, Stanford University

2018

A Stable Marriage Requires Communication*

Yannai A. Gonczarowski† Noam Nisan‡ Rafail Ostrovsky§ Will Rosenbaum¶

2008

Sampling Stable Marriages: Why Spouse-Swapping Won't Work*

Nayantara Bhatnagar† Sam Greenberg‡ Dana Randall§

2003

Marriage, Honesty, and Stability

Nicole Immorlica* Mohammad Mahdian*

History of Stable Matching

- In 1900s matching medical residents to hospitals was decentralized
- Increasingly competitive
  - By the 1940s, appointments were often made as early as the beginning of the junior year of medical school

The market for law school graduate is also known for these problems. **Roth** in this article **"Who Gets What And Why"** quotes a law school student who in 2005, on a flight from her 1st interview to 2nd interview, got 3 voicemail messages: the 1st extending an offer from where she just interviewed; the 2nd to urge her to return the call soon; and the 3rd to rescind the offer. Her flight was only 35 mins long!

"Who Gets What and Why" by A Roth

Why have Centralized Markets

- In 1900s matching medical residents to hospitals was decentralized
- Increasingly competitive
    - By the 1940s, appointments were often made as early as the beginning of the junior year of medical school
- In 1945, a variant of deferred acceptance implemented by **AAP (American Associated of Pediatrics)** and NRMP (National Resident Matching program) to match residents to hospitals
- This was the invention of **"the match"**

"The Origins, History, and Design of the Resident Match" by A Roth

Nobel Prize 2012: Shapley & Roth

David Gale
PROFESSOR, UC BERKELEY

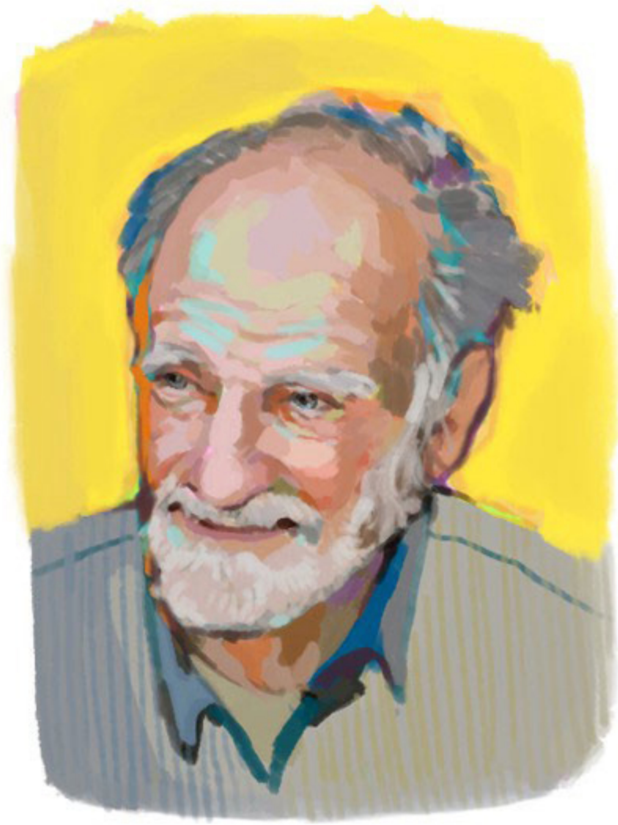

Lloyd Shapley
PROFESSOR EMERITUS, UCLA

Stable matching: Theory, evidence, and practical design

*This year's Prize to **Lloyd Shapley** and **Alvin Roth** extends from abstract theory developed in the 1960s, over empirical work in the 1980s, to ongoing efforts to find practical solutions to real-world problems. Examples include the assignment of new doctors to hospitals, students to schools, and human organs for transplant to recipients. Lloyd Shapley made the early theoretical contributions, which were unexpectedly adopted two decades later when Alvin Roth investigated the market for U.S. doctors. His findings generated further analytical developments, as well as practical design of market institutions.*

Why Stability: The Story of NRMP

- Empirical evidence in support
- In UK in the 60s, residency programs decided to move from a decentralized system to a centralized clearinghouse
- The details of the implementation were left to individual regions
- Roth looked at data from 7 regions
    - Two followed a stable implementation; they remain in use today
    - Five regions implemented unstable variants, 3 of which did not survive long (due to poor participation and negotiations outside the system)

Classic Stable Matching Problem

- **Input**: n applicants and n jobs, complete preference lists

- **Output:** a **perfect matching** M that is **stable** (no applicant and job prefer each other to their match)

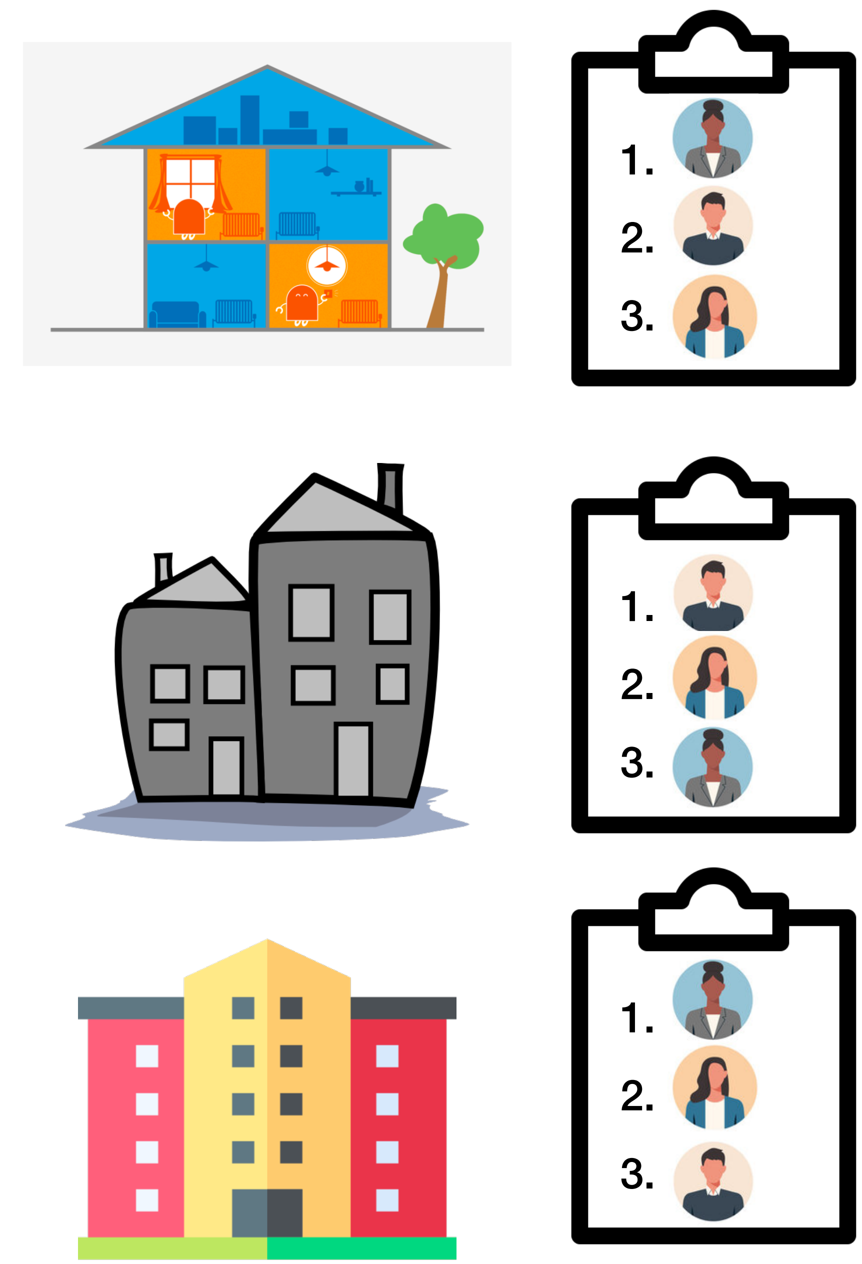

Classic Stable Matching Problem

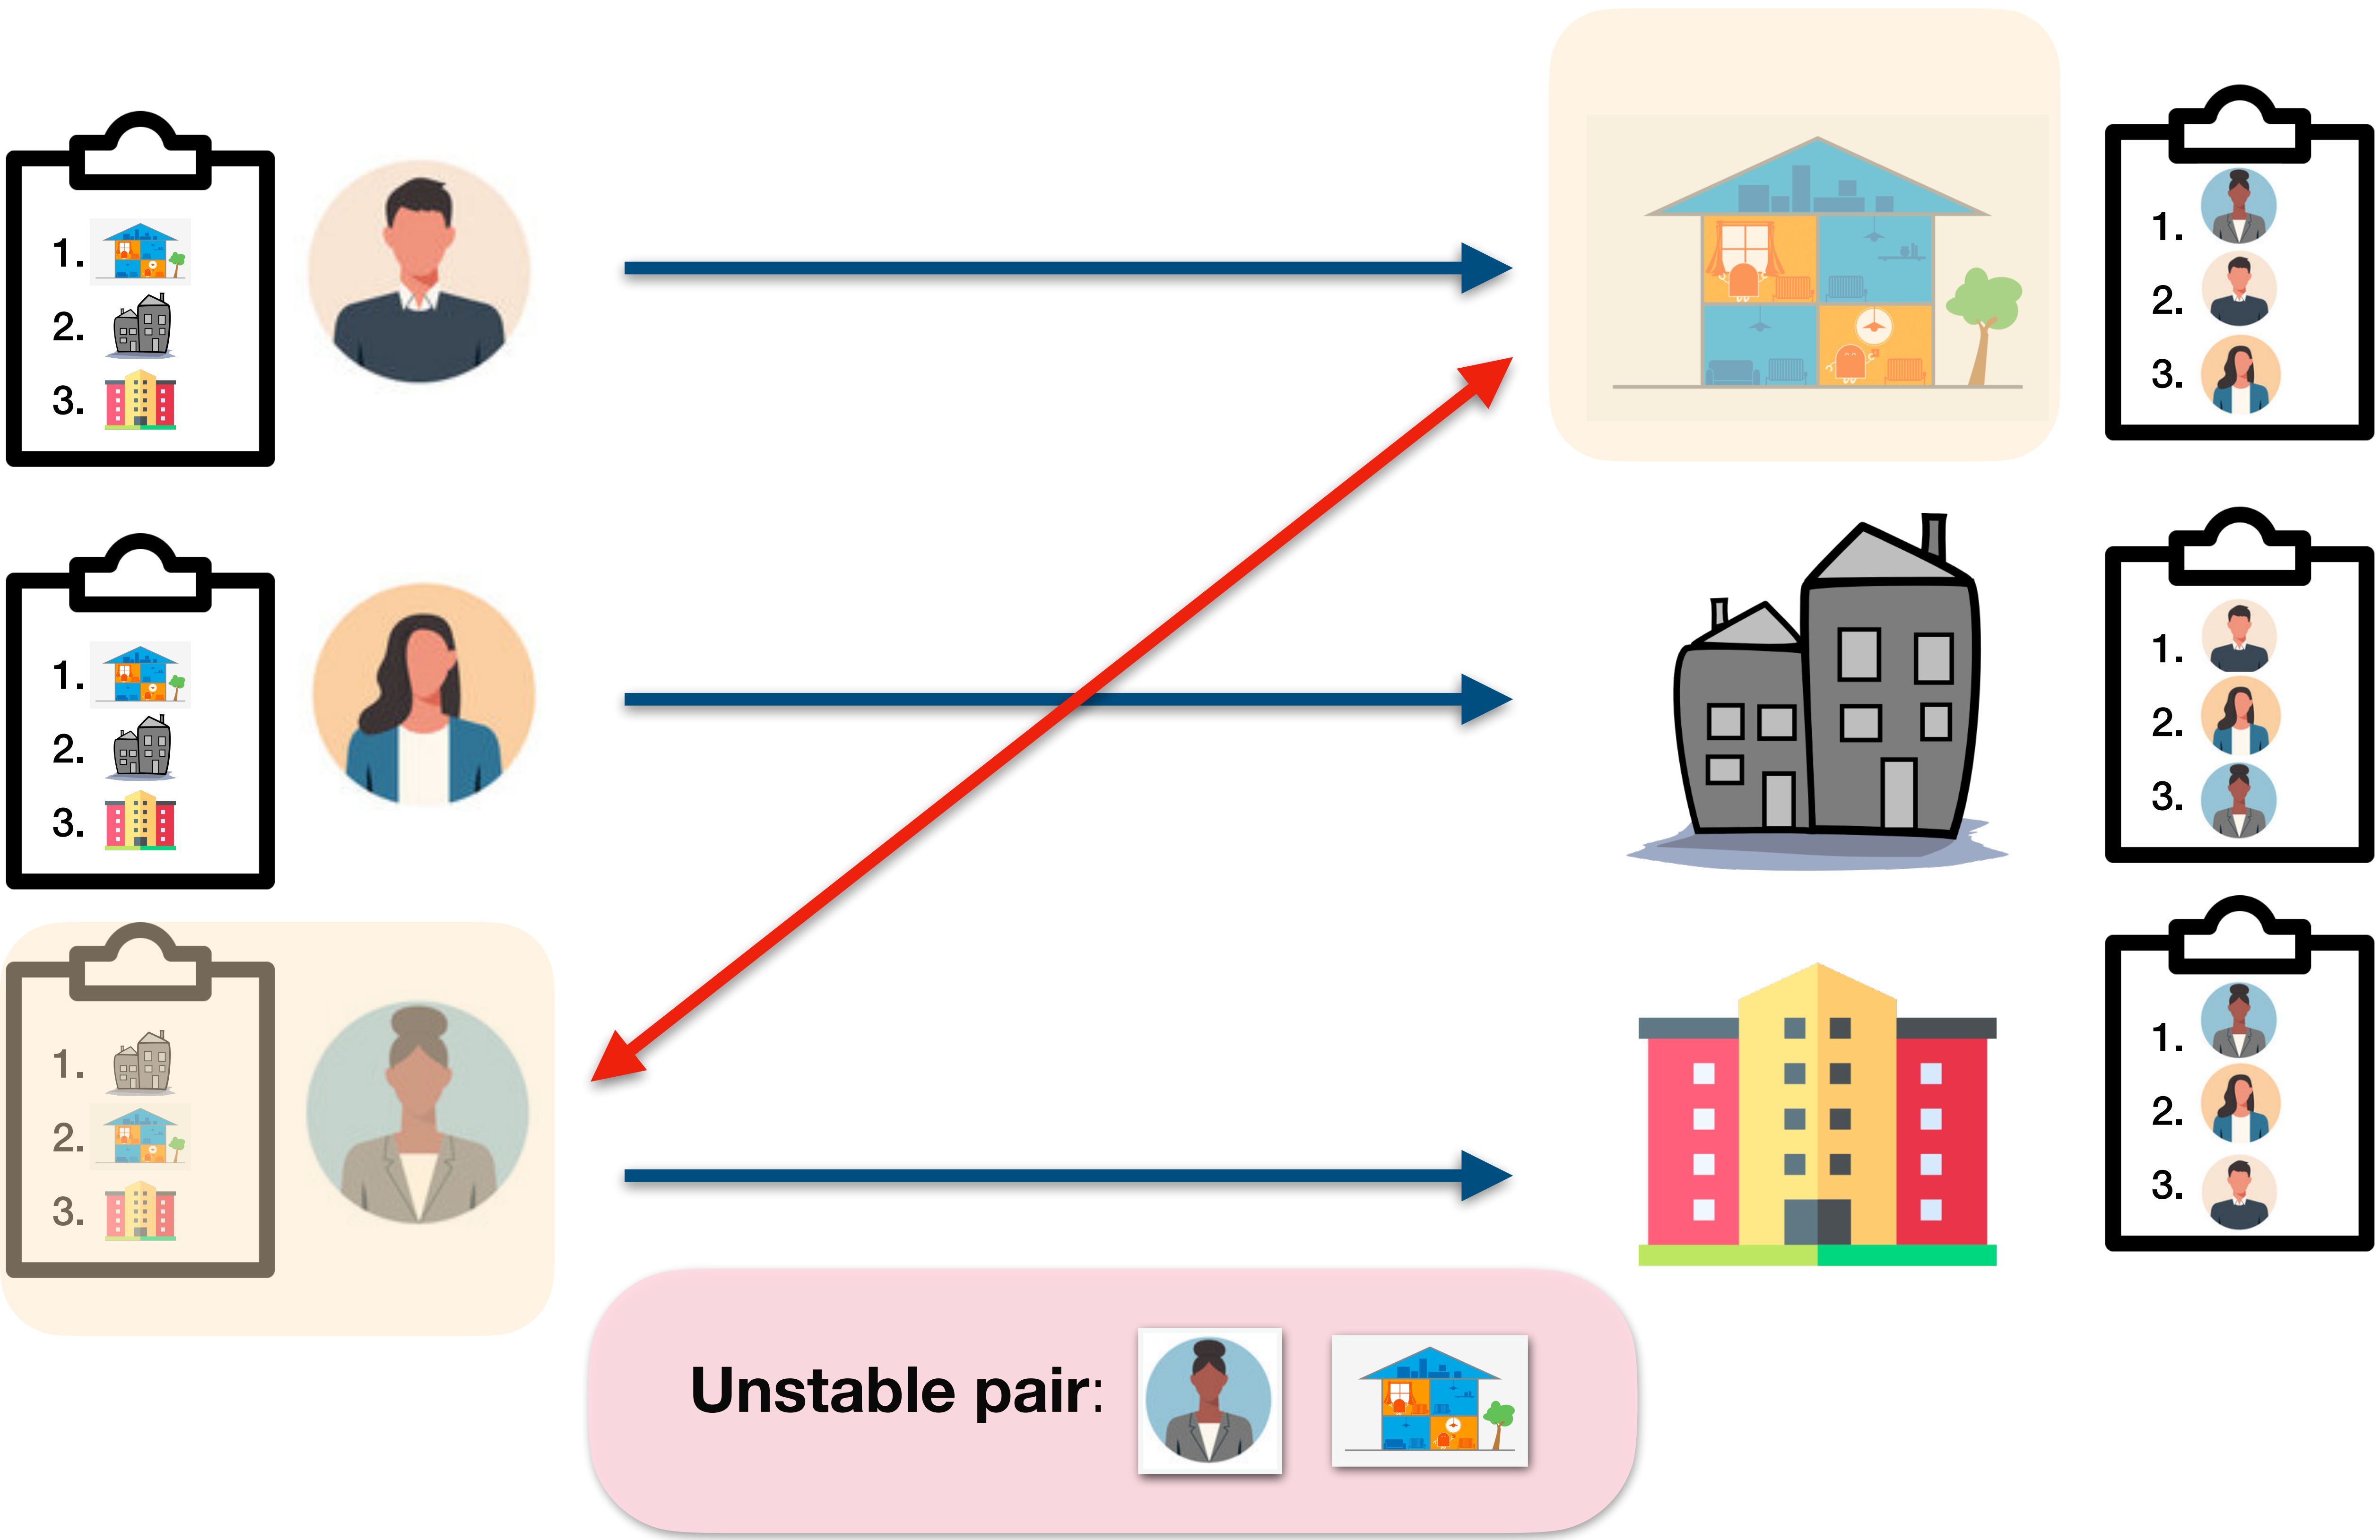

Classic Stable Matching Problem

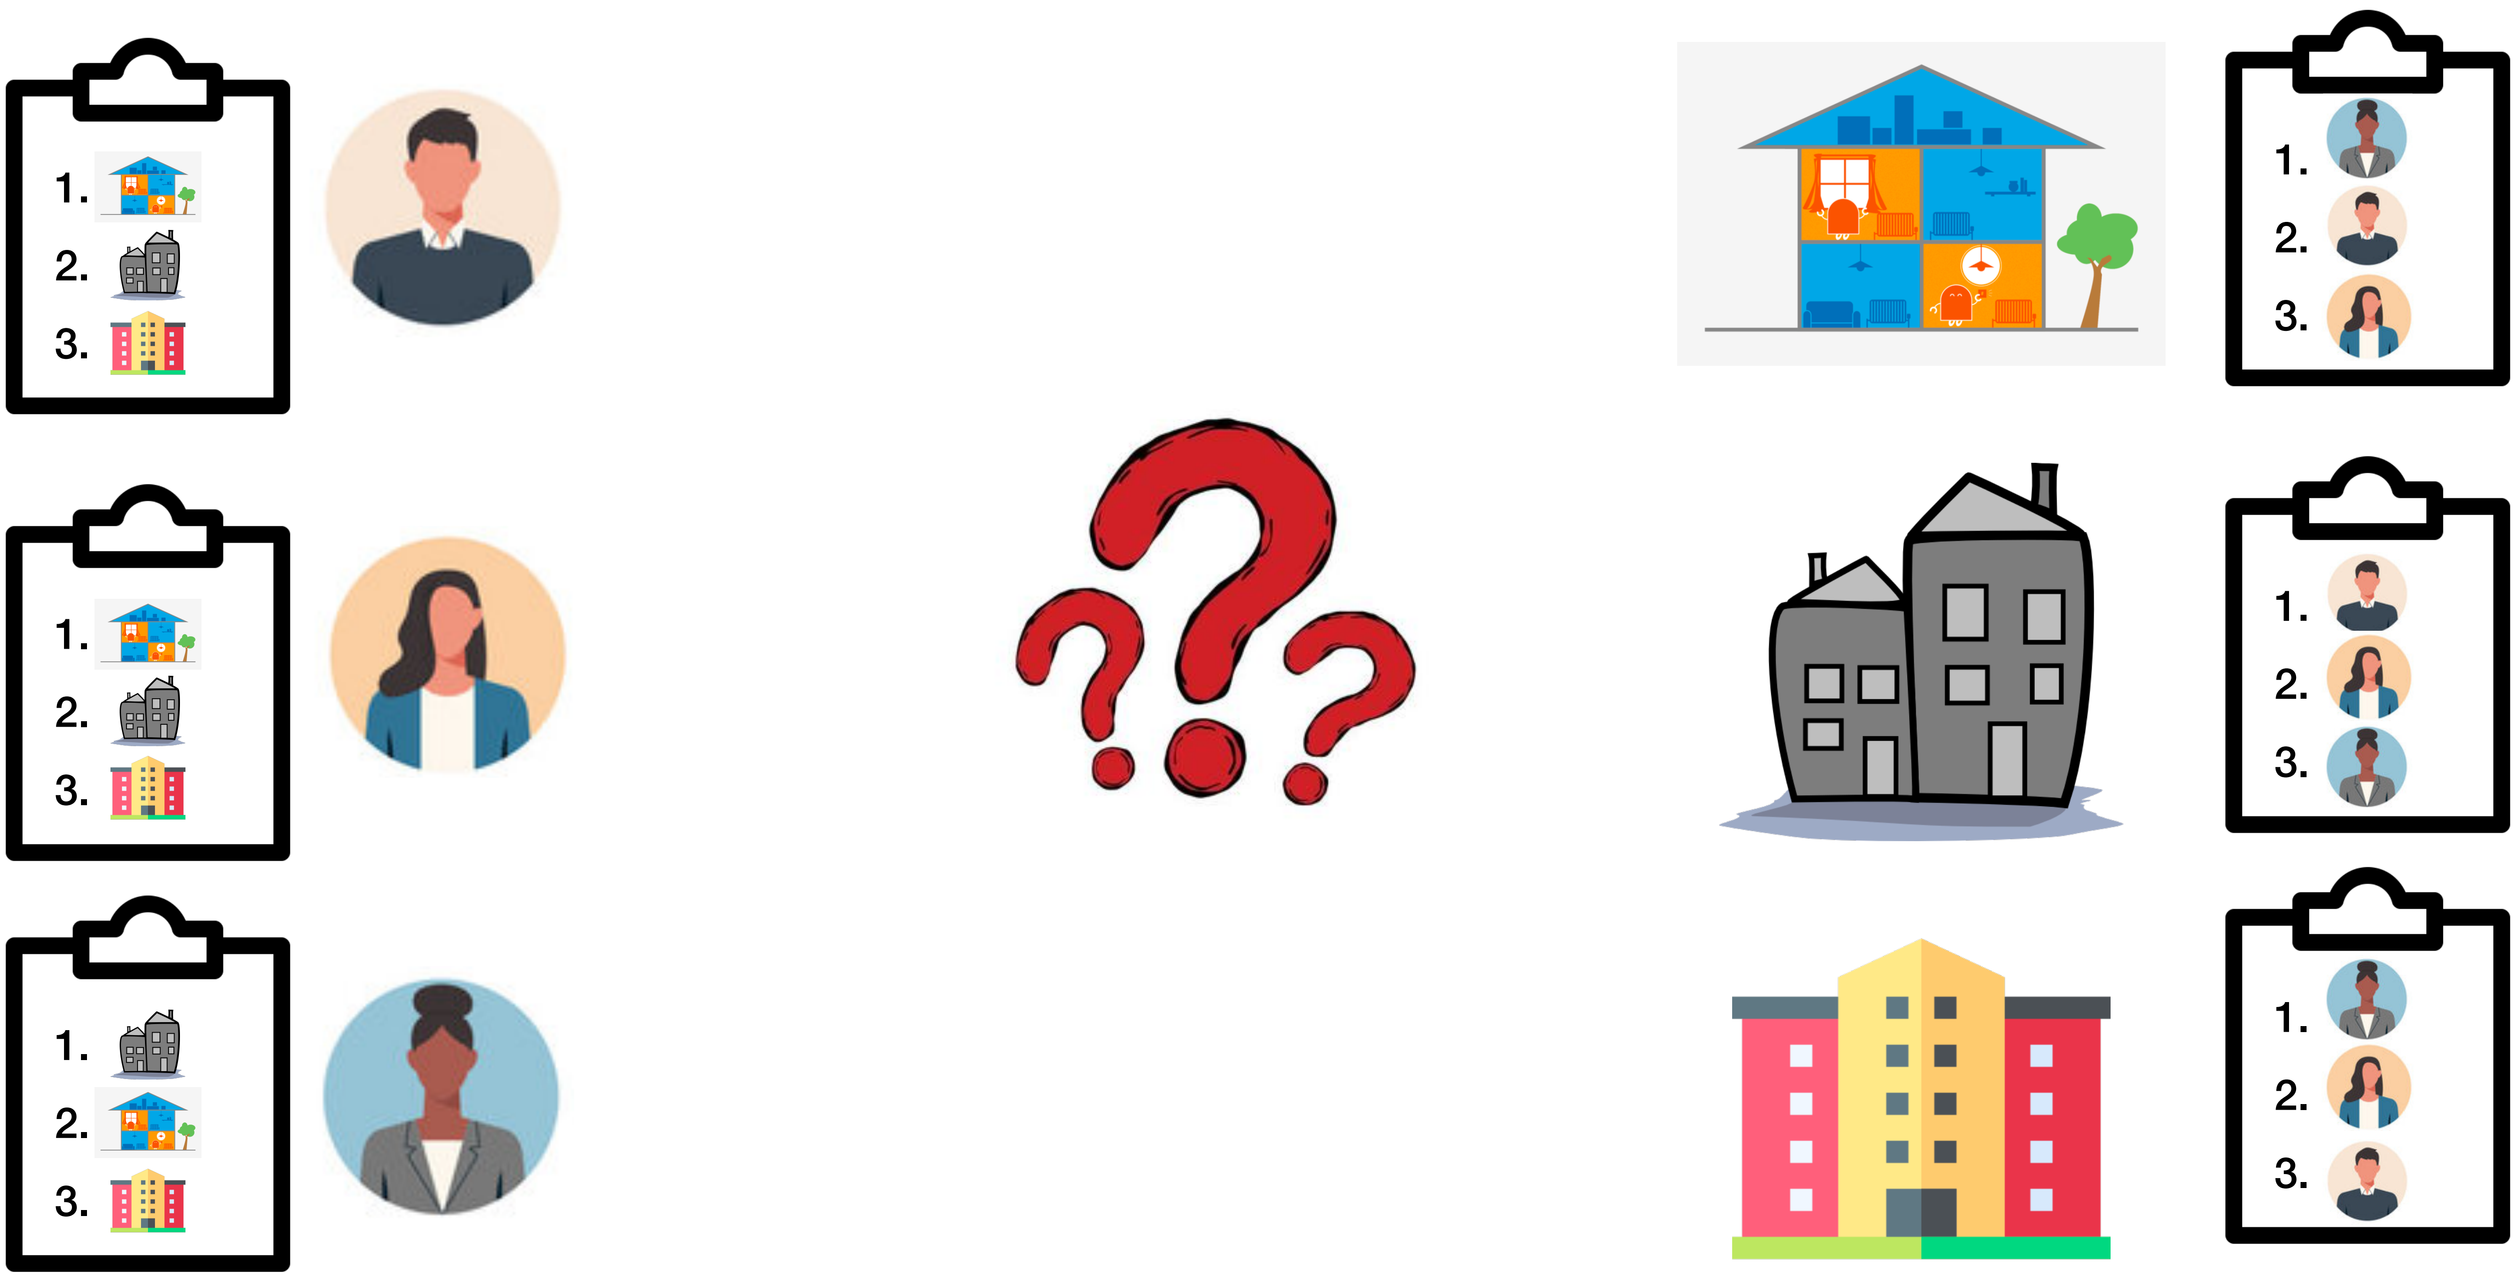

[Gale Shapley 1952] A perfect stable matching **always exists**!

Deferred Acceptance (DA) Algorithm

- Proceeds in rounds
  - Each **unmatched** applicant **"proposes"** to their most preferred job
  - jobs retain the best proposal they have received & reject others
- Matching is finalized when each applicant is matched

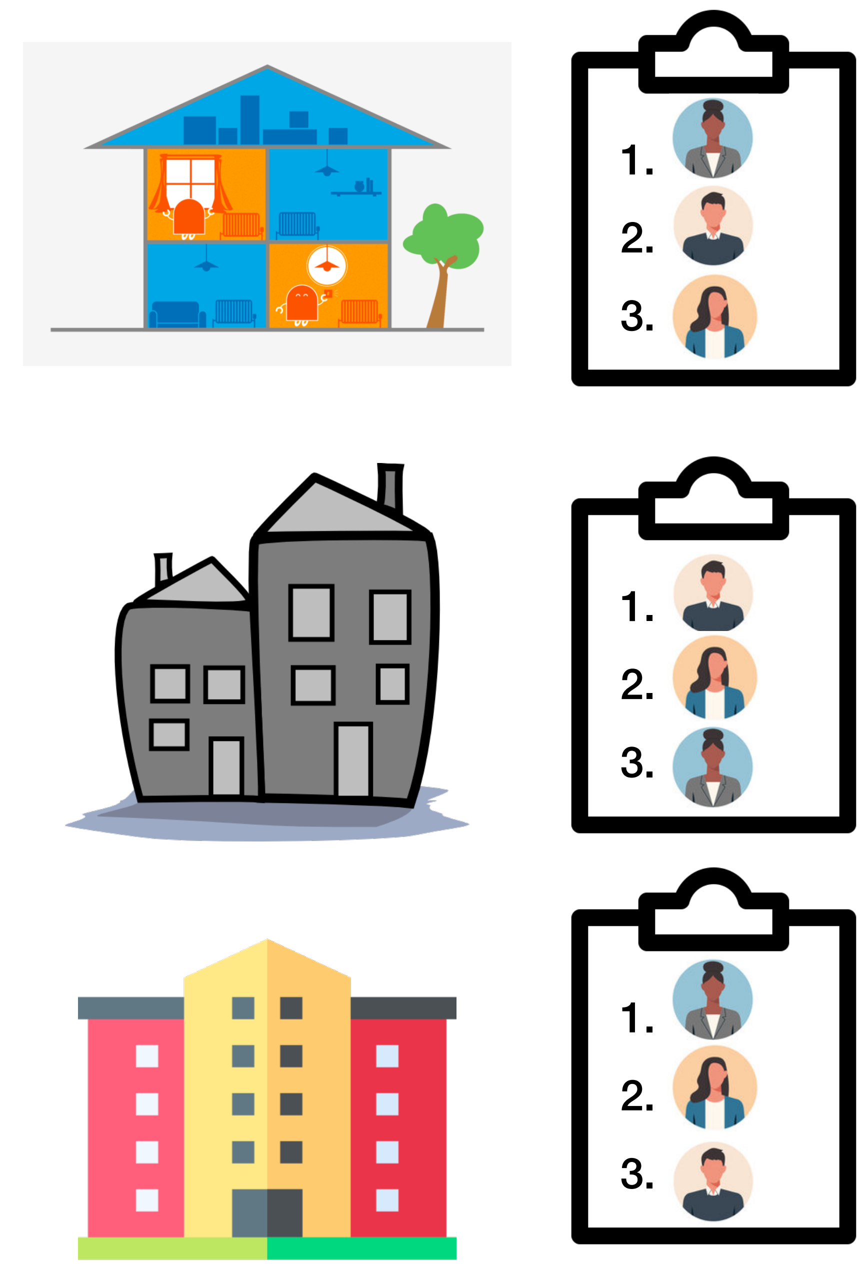

[Gale Shapley 1952] A perfect stable matching always exists!

Classic Stable Matching Problem

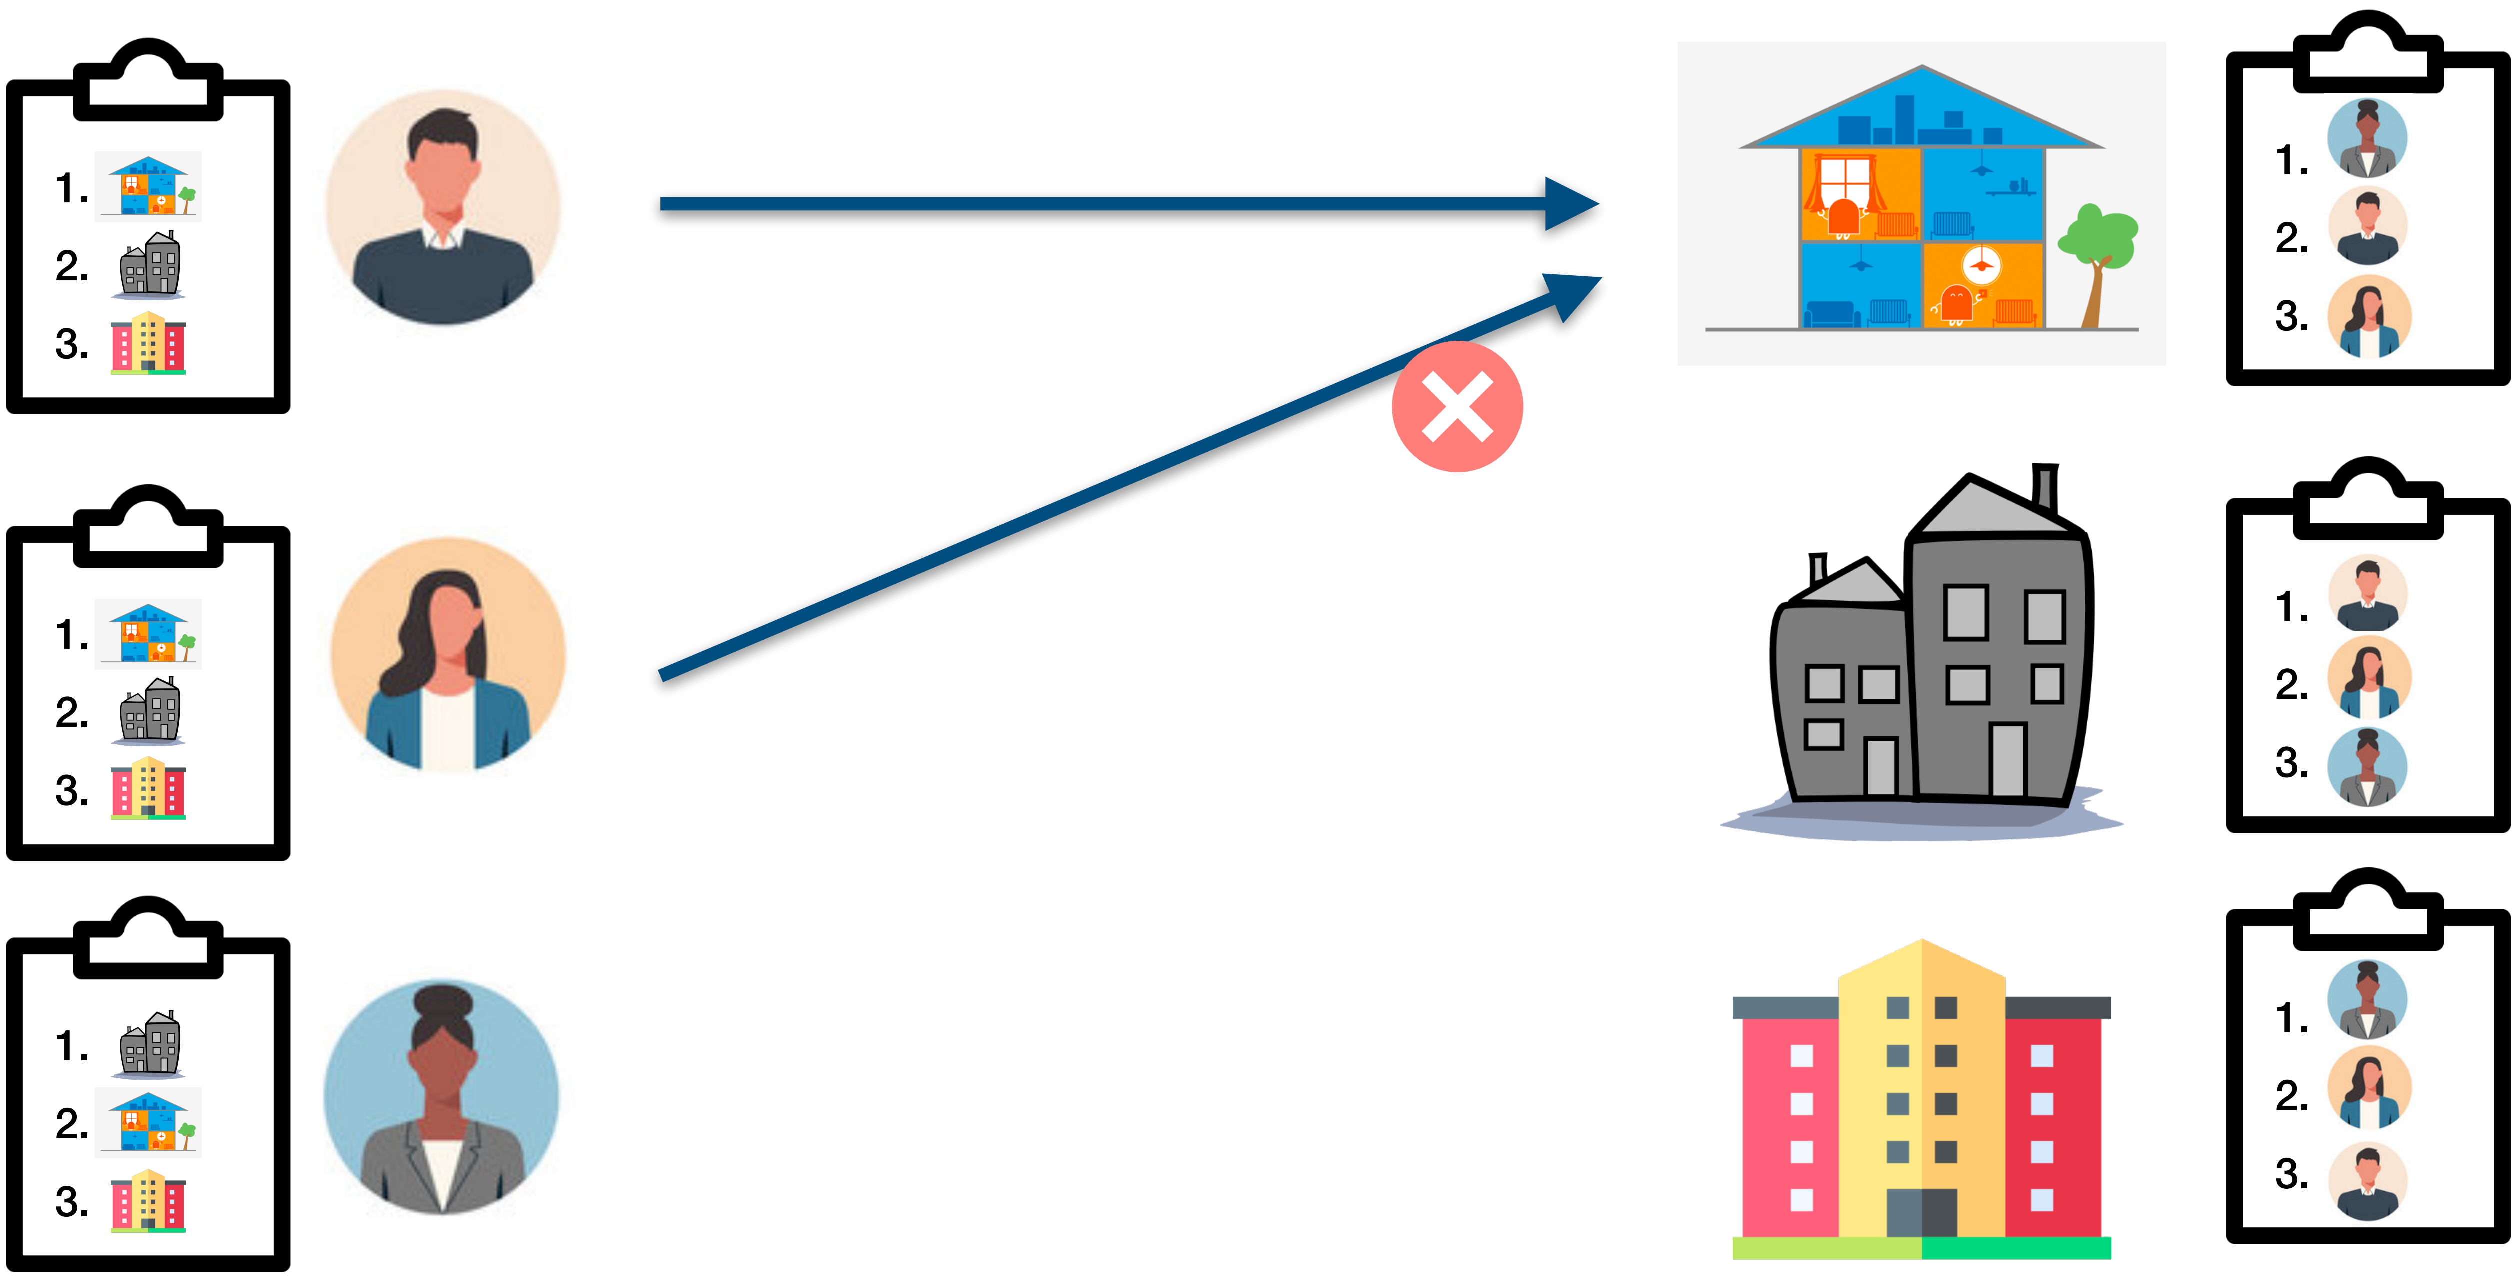

Classic Stable Matching Problem

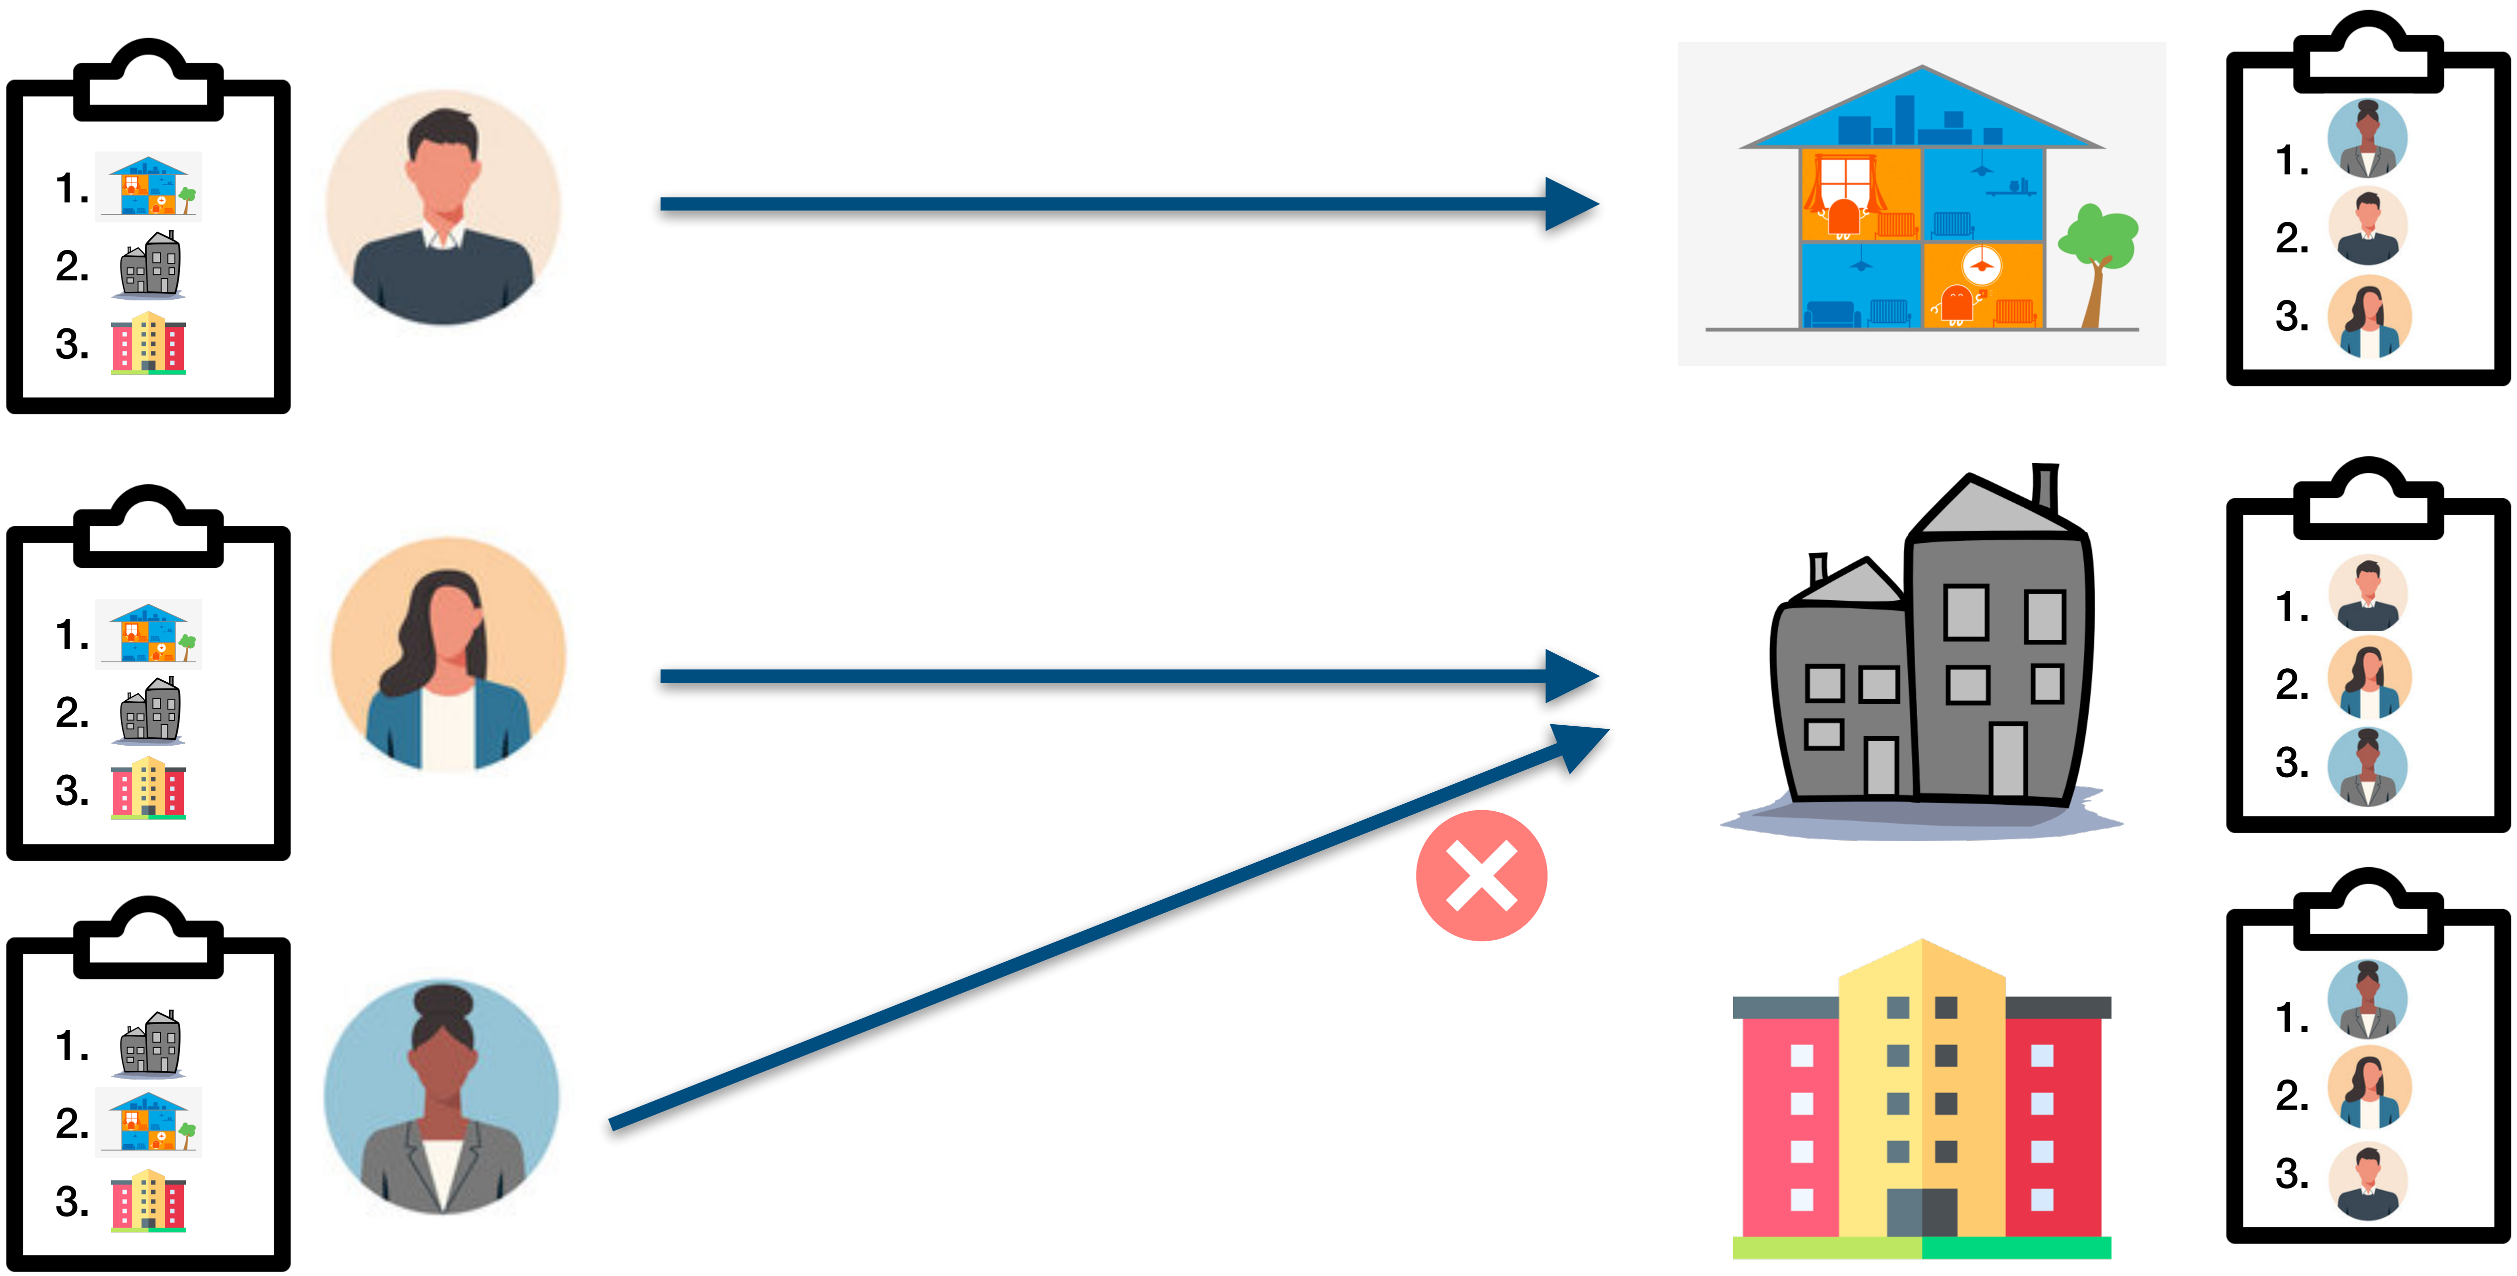

Classic Stable Matching Problem

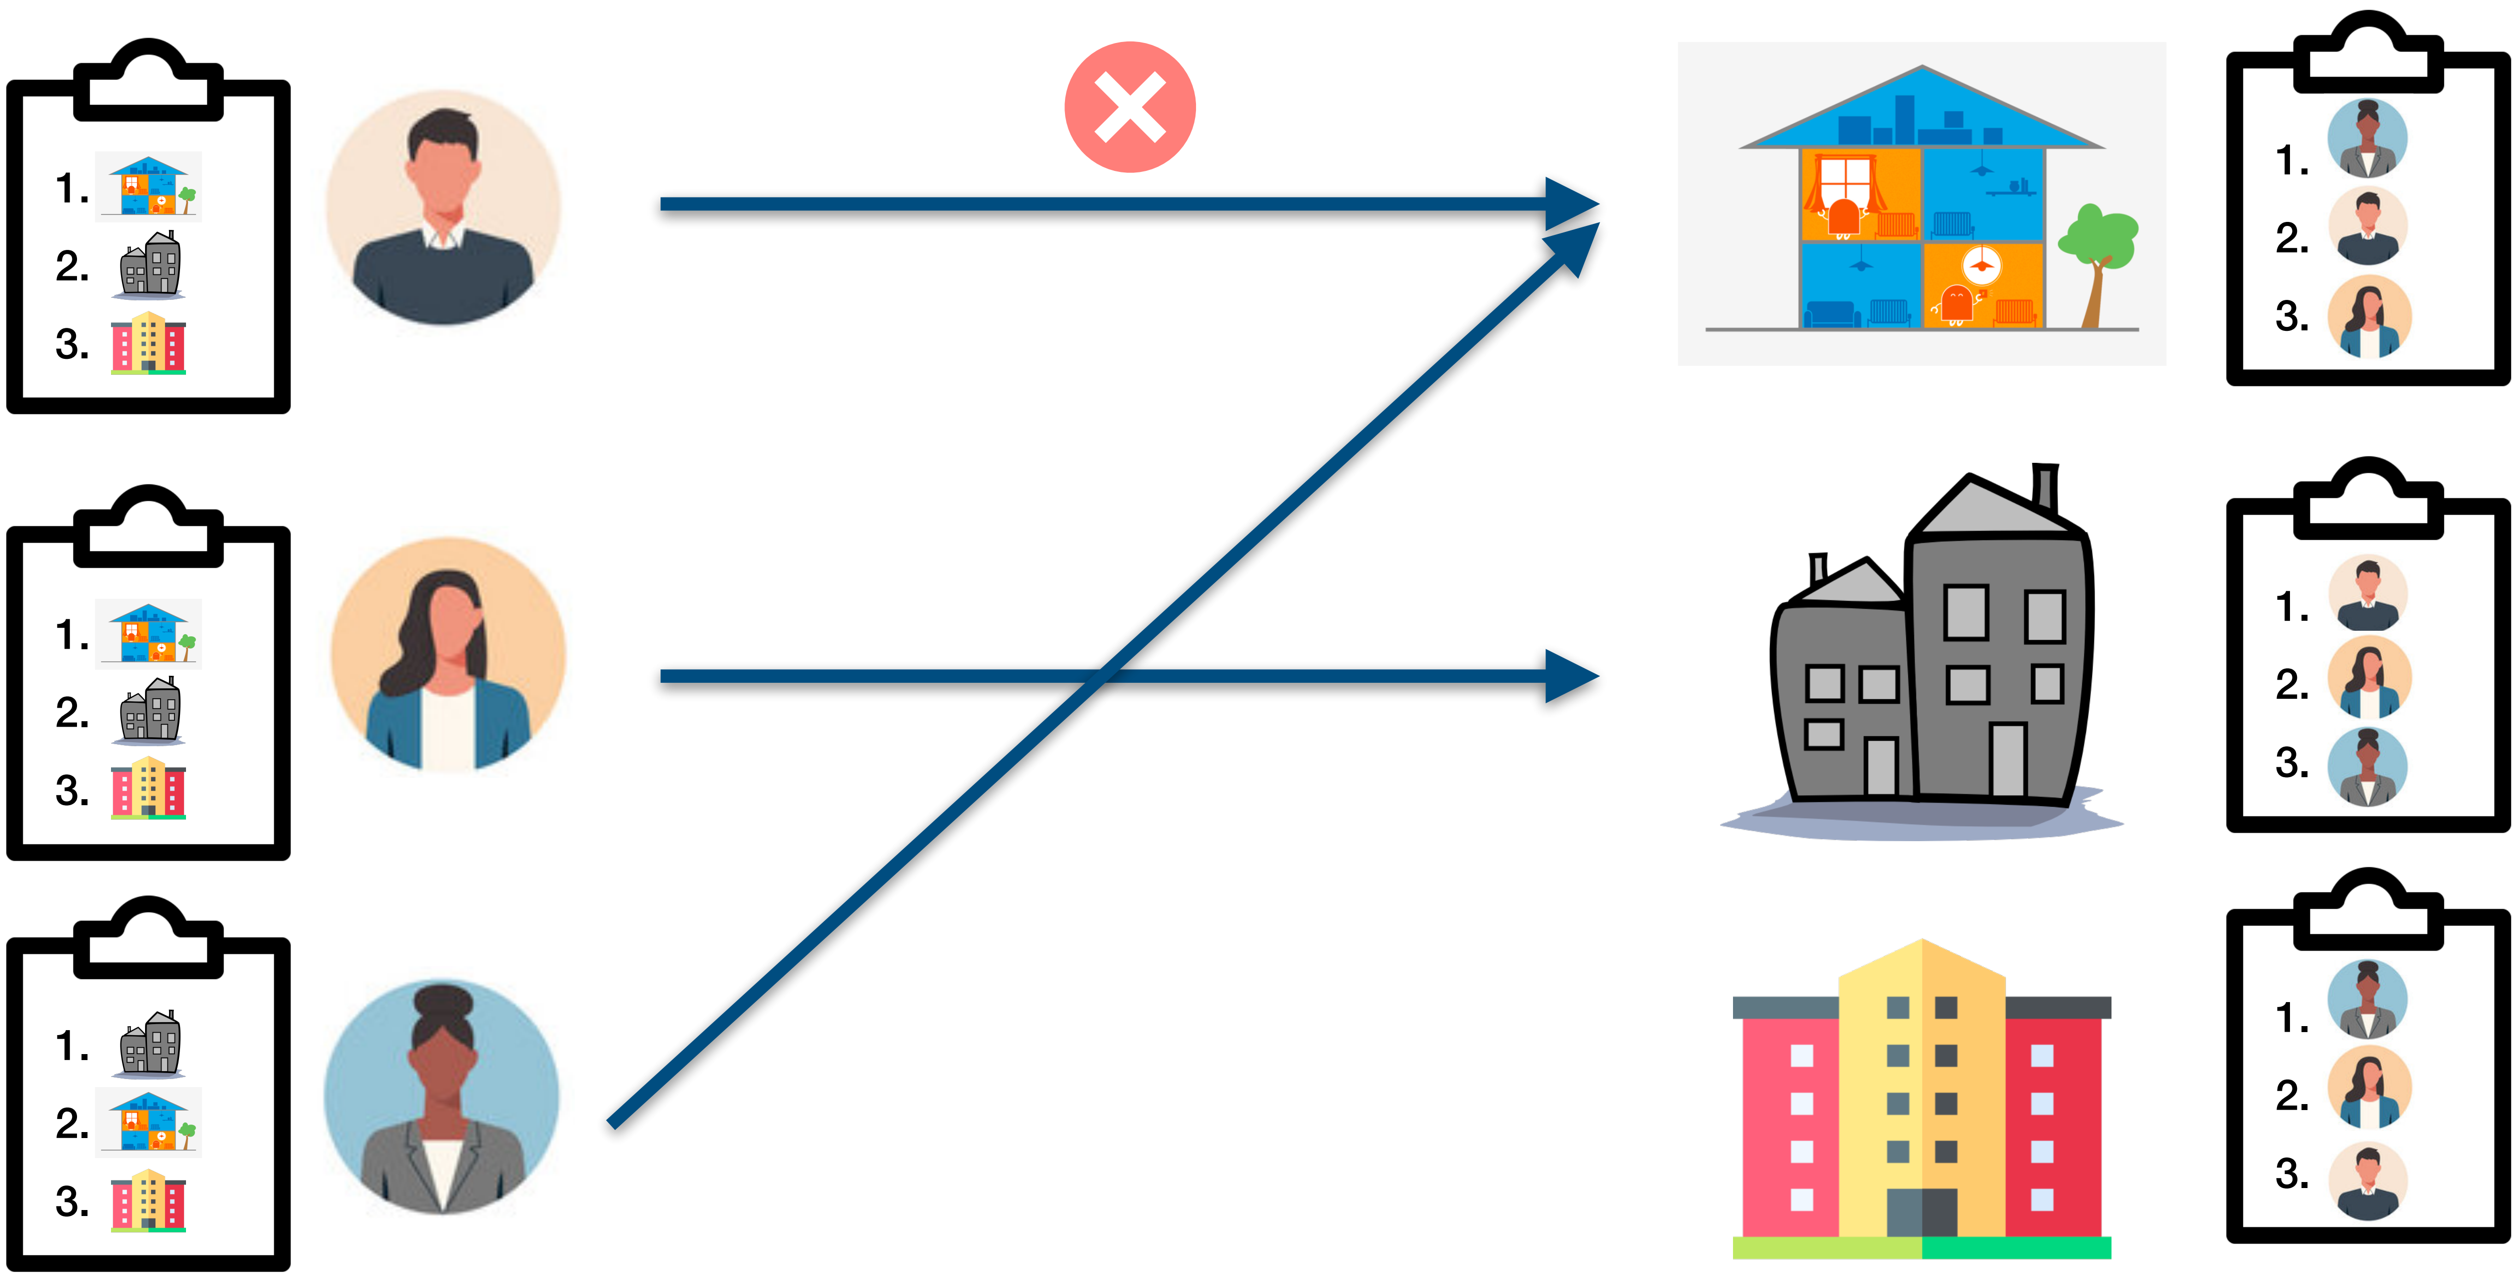

Classic Stable Matching Problem

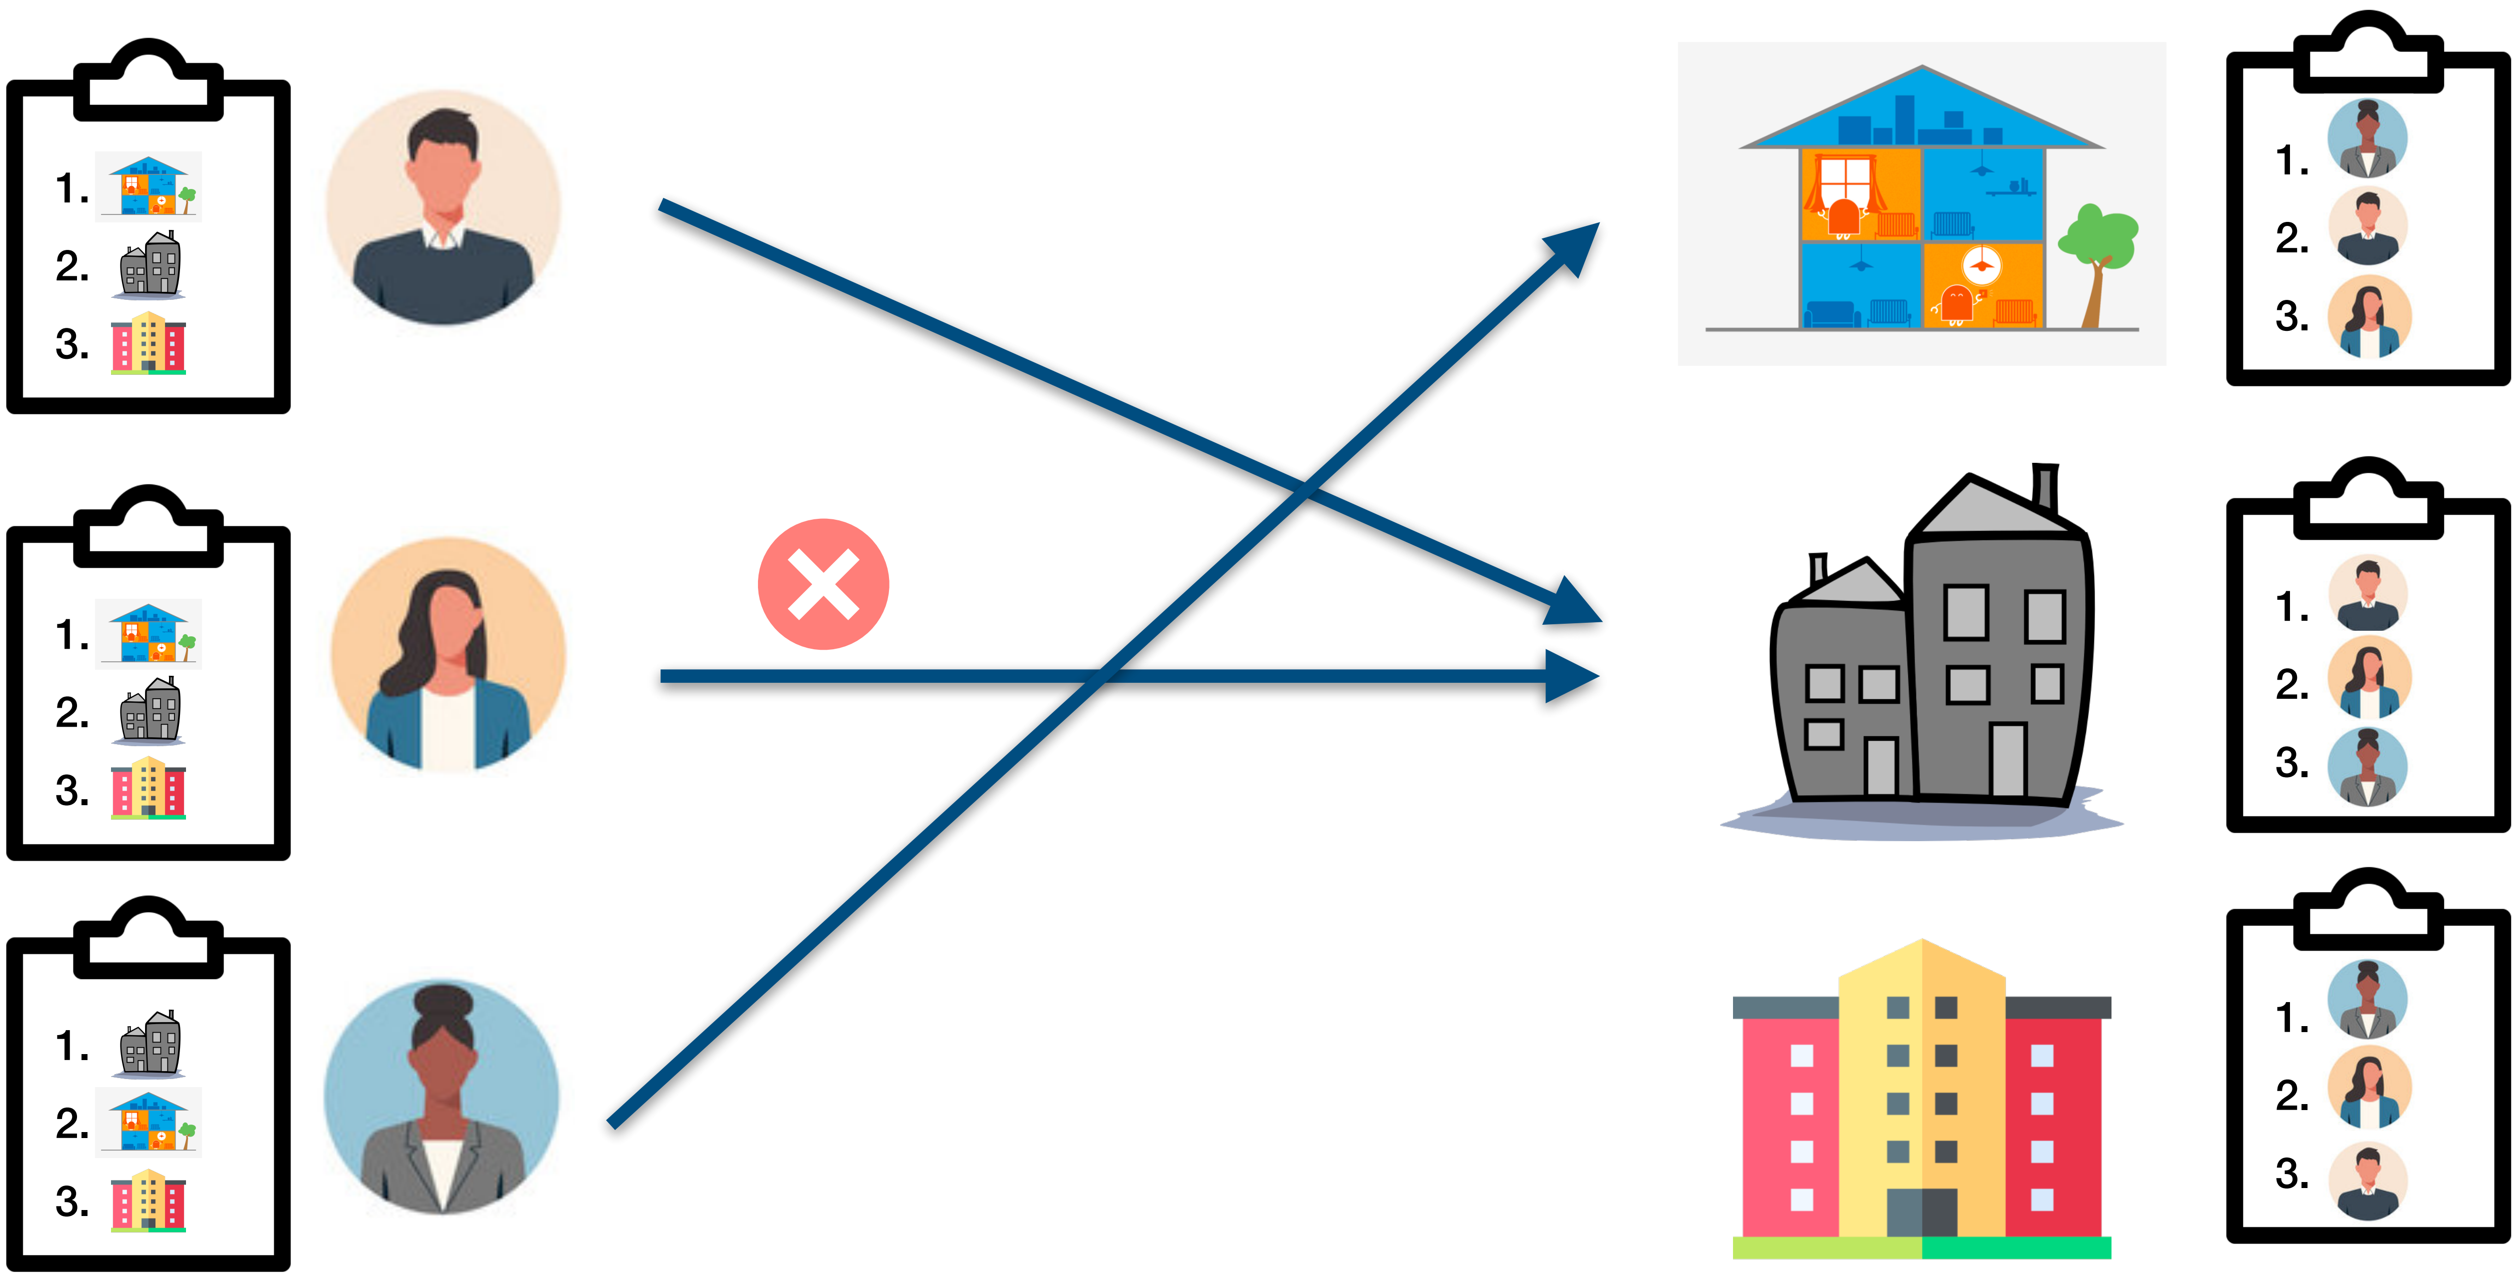

Classic Stable Matching Problem

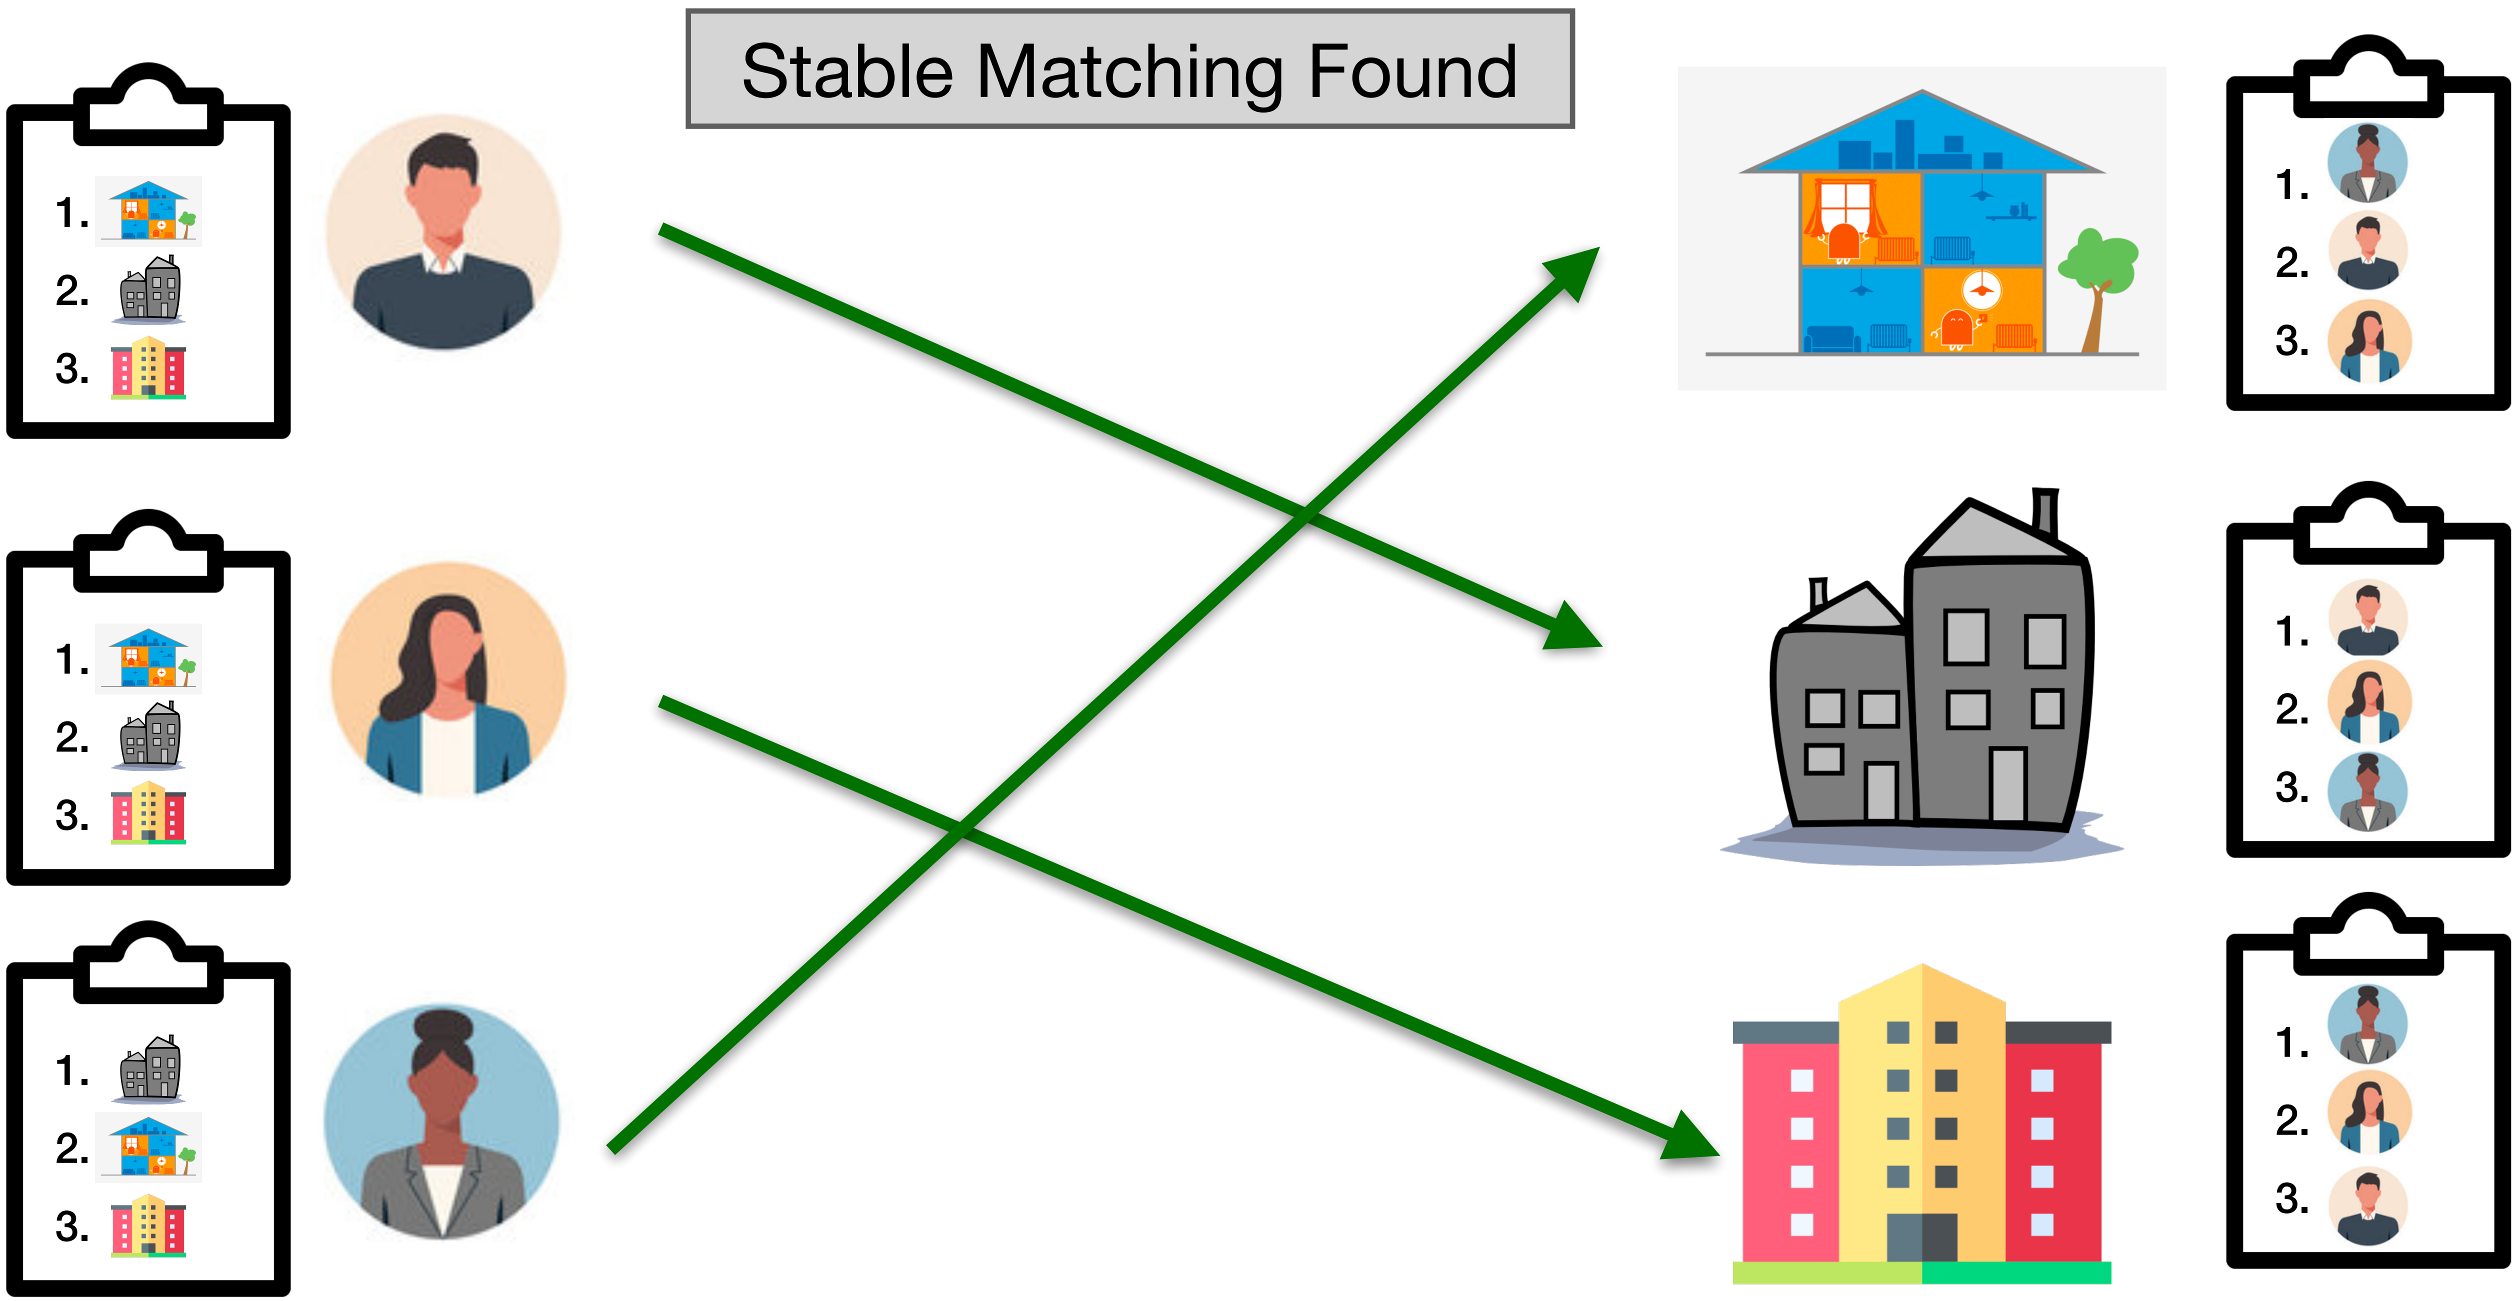

- Output matching is **applicant optimal** and **job pessimal**

Deferred Acceptance Algorithm

GALE–SHAPLEY (*preference lists for hospitals and students*)

INITIALIZE M to empty matching.

WHILE (some hospital h is unmatched and hasn't proposed to every student)

  $s \leftarrow$ first student on h's list to whom h has not yet proposed.

  IF (s is unmatched)

    Add h–s to matching M.

  ELSE IF (s prefers h to current partner h')

    Replace h'–s with h–s in matching M.

  ELSE

    s rejects h.

RETURN stable matching M.

Deferred Acceptance Properties

Lemma 1. DA algorithm always produces a stable matching.

Proof. (By contradiction) Let M be the resulting matching. Suppose $\exists (h, s)$ such that $(h, s'), (h', s) \in M$ and

- h prefers s over s' and s prefers h over h'

Thus h must have offered to s before s'

- Either s broke the match to h at some point for some h'', or s already had a match h'' that s preferred over h

But students always trade up, so s must prefer final match h' over h'', which they prefer over h. $(\Rightarrow\Leftarrow)$ ∎

Deferred Acceptance Properties

- The deferred-acceptance algorithm does not specify the order in which the hospitals should make offers
- Do all orders produce the same unique matching?
- Given an input instance, there may be several stable matchings.
- **Question.** Does Gale-Shapely produce the "best matching" for hospitals or students?
- Turns out hospital-proposing algorithm produces a unique matching that is hospital optimal and student pessimal
    - Matches hospital to "best achievable" student and student to "worst-achievable" hospital among all stable matchings

Best Achievable Partner

Let I be an instance of the stable marriage problem

- A student $s \in S$ is an **achievable partner** for hospital $h \in H$, if (h, s) is part of some stable matching of I.
  - We call the pair (h, s) an achievable pair
- For hospital $h \in H$, let **best(h)** denote the most preferred achievable partner of h (among all stable matchings)
- **Lemma.** $M^* = \{(h, \text{best}(h)) \mid h \in H\}$ is the unique output of the hospital-proposing deferred-acceptance algorithm.
  - True regardless of the order in which different hospitals make offers

Best Achievable Partner

- **Lemma.** $M^* = \{(h, \text{best}(h)) \mid h \in H\}$ is the unique output of the hospital-proposing deferred-acceptance algorithm.
- **Proof (By Contradiction).** Let h be the first hospital rejected by $s^* = \text{best}(h)$
  - s^* instead holds on to offer from some h'
- s^* must be the best achievable partner for h', why?
  - If not h' has already been rejected by $\text{best}(h')$, violates condition that h is the first such hospital
- Let M be a stable matching s.t. $(h, s^*) \in M$
- **Claim.** (h', s^*) is a blocking pair for matching M, why?
  - s prefers h' to h, and h' prefers s^* to whoever they are matched to in M $(\Rightarrow\!\Leftarrow)$ ∎

Takeaways

- The outcome of hospital-offering deferred acceptance is hospital-optimal, among all stable matching
  - There is no tradeoff to make in terms of who offers first!
- What about the accepting side?
  - The outcome of the hospital-offering deferred acceptance is **students-pessimal**, among all stable matchings
  - In particular, students get matched to their worst-achievable partner among all stable matchings
- **Incentive considerations.** Which side of the market has an incentive to misreport their preferences?
- Can misreports be beneficial? Is the mechanism strategyproof?