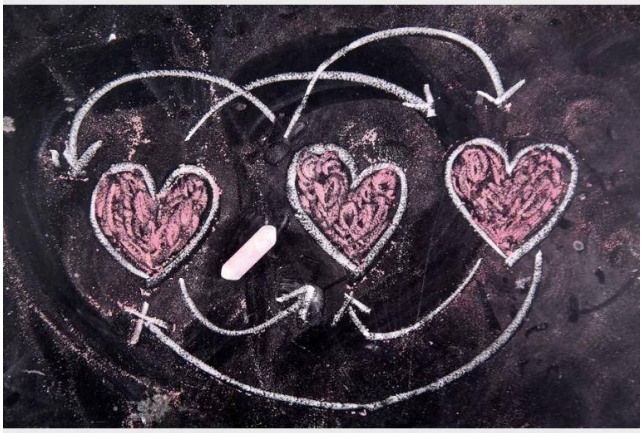
October 10, 2024 | 6 min read Add SciAm

The ‘Stable Marriage Problem’ Solution Underpins Dating Apps and School Admissions

An elegant matchmaking algorithm called Gale-Shapley can find the best possible pairings for everybody

BY MAX SPRINGER EDITED BY ALLISON PARSHALL



Fotografiche/Getty Images

Mathematics

Let’s create a reality dating show unlike any other in one key aspect. First, we’ll rent a villa on a tropical island. Then we’ll fly in five men and five women, each with their own (heterosexual) dating preferences. Our goal, though, is the exact opposite of the *Love Island* franchise: we want absolutely zero drama. Can we ensure that everyone pairs off with a partner and sticks with them, without jealousy rearing its ugly head?

Mathematicians call this dilemma the “stable matching problem” or “stable marriage problem.” And though matters of the heart may be fickle, researchers have proved that by using a simple algorithm, they can *always* find a stable set of matches between all members of two equally sized groups. The late mathematician Lloyd Shapley shared the 2012 Nobel memorial prize in economic sciences for the discovery of this algorithm—and for good reason: it is still used today to pair medical residents with hospitals and children with schools, and it has even inspired dating-app algorithms.

According to mathematicians, a relationship is stable when neither person has a better option—at least, not one that is also interested in them. Instability, then, could look something like this: Imagine Alice is currently paired off with Bob, while Charlie is currently with Darlene. Bob is secretly in love with Darlene, however, and Darlene can’t bear the thought of another day with Charlie. Because Bob and Darlene seem primed to run off

together and leave their partners behind, mathematicians call this situation unstable.

The same dilemma pops up outside of romantic life, too. Shapley and mathematician David Gale first described it as a problem of college admissions: What kind of application process would ensure that colleges and applicants, each with their own sets of preferences, were satisfied with their picks? In 1962 Gale and Shapley showed that for any set of students and colleges (or men and women, in the dating show example), there always exists a set of pairings where every match is stable. What's more, they provided a simple process, or algorithm, that takes everyone's rankings and builds relatively happy pairs.

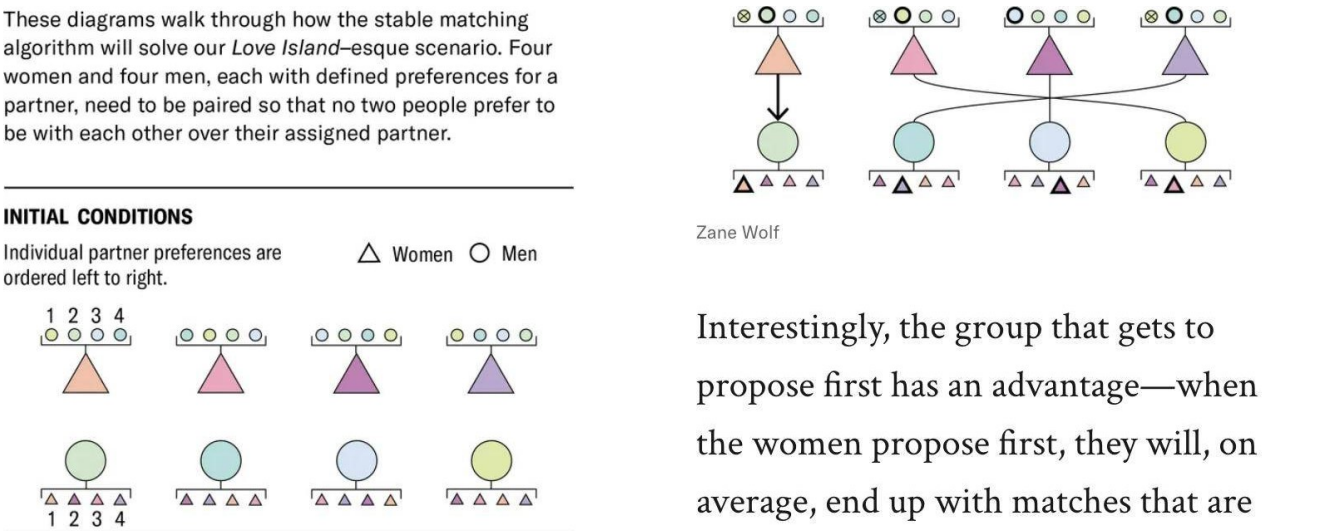
Here's how it works. To find a set of stable, drama-free partnerships for our 10 dating show contestants, we need to first have each contestant rank all members of the opposite gender in order of their preference.

Then, on the first day in the villa, each woman makes a relationship proposal to her top-choice man. Some men receive many proposals, while others might receive none. Each man then rejects all but his more preferred suitor, resulting in a tentative match for some contestants, while others remain unpaired.

On the second day, each rejected woman proposes to her second choice (even if he is already paired up). The men consider the new proposals, and some may abandon their current match if they prefer the new suitor. Then some of those newly heartbroken women would propose to their next possible partner.

This process repeats on the third, fourth and subsequent days, for as many times as is necessary, until everyone has settled on a match. While not everyone will be happy with their pairing using this process, it is mathematically guaranteed that no two people would both prefer to be with each other than who they are currently with (assuming their preferences haven't shifted upon getting to know each other, that is). While this won't guarantee that our *Love Island* spin-off remains a relaxing, drama-free watch, it is probably as good as we will get.

Stable Matching Algorithm



Zane Wolf

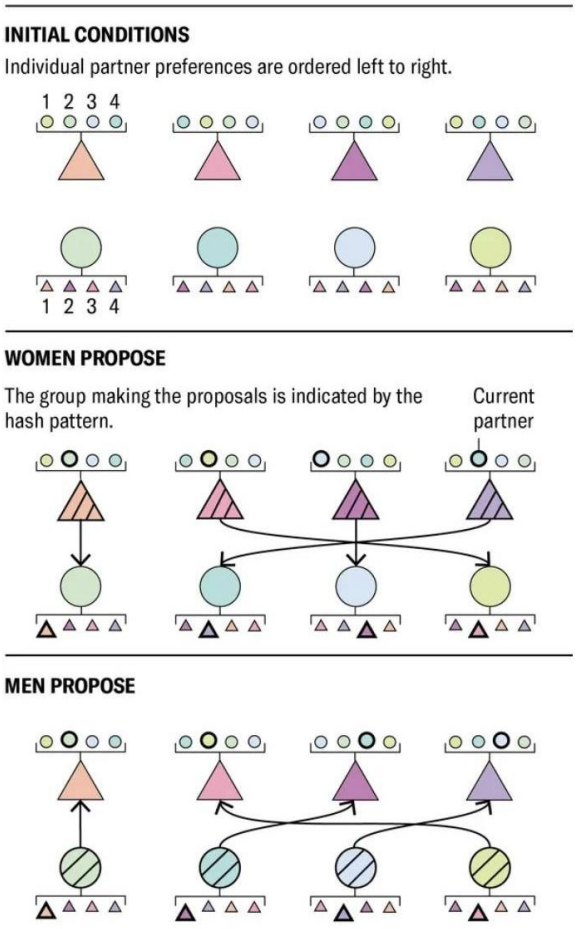
Interestingly, the group that gets to propose first has an advantage—when the women propose first, they will, on average, end up with matches that are more desirable to them than the men will. “The one issue with Gale-Shapley is that it gives you these extreme matches on either side,” explains Vijay Vazirani, a computer scientist at the University of California, Irvine.

The outcomes might be slightly lopsided, but Gale and Shapley’s strategy can’t be beat. And as it turned out, a version of it had already been in use since the 1950s by an organization that matches medical students to residency programs across the country. In 1984 Stanford University economist Alvin Roth used Gale and Shapley’s mathematical language to show that the process used by this organization not only guaranteed stable matches but was also “strategy-proof”—meaning that there’s no way to game the system. This feature, more generally called incentive compatibility, is prized because it

means that everyone will end up with their best option if they report their preferences truthfully.

Results May Vary Based on the Proposing Group

On average, the group making the proposals will finish with more desirable matches than the group receiving the proposals.



Zane Wolf

Roth and economist Elliott Peranson also made a few tweaks to the algorithm. They adapted it to work for medical students who were married to each other and looking to complete their residencies in the same location. They also noted that residents were getting the shorter end of the stick because hospitals proposed first. Roth advocated for residents to propose first to ensure they would get their best outcome. To this day, hospitals and incoming residents provide a ranking of each other, and the mathematics works out to ensure a stable state is achieved. Roth and Shapley won the 2012 economics Nobel for their work.

Roth and his colleagues also used this mathematical language to tackle another gnarly matching problem: assigning kids to public schools in the U.S.'s biggest cities. In 2003 they adapted Gale-Shapley to assign students to New York City's notoriously competitive public high schools. In the first year of operation, the number of students matched with one of their top choices increased from about 50,000 to more than 70,000. Another version of the algorithm is also used to assign students to public schools in Boston.

And back in the romantic sphere, the Gale-Shapley algorithm has even inspired the inner workings of dating apps such as Hinge. While users do not explicitly rank their potential

matches, these apps observe users' history of likes and dislikes, along with their stated dating preferences, to curate a handful of "top matches" that it shows to users first. A like or message sent to a potential match is analogous to a "proposal" in the original algorithm.

"The power of models like [Gale-Shapley] is to abstract an idea across many different settings," emphasizes Jon Kleinberg, a professor of computer science at Cornell University. Problems across different domains "can all have something in common conceptually," he says, and the Gale-Shapley algorithm "gave us a mathematical language to talk about [them]."

Though the algorithm is simple and reliable, it can also amplify existing disparities if there is bias in the rankings. Admissions data acquired by The City, a New York City–based nonprofit newsroom, showed how Black and Latino students are regularly selected for admission at a lower rate by the metropolis's high schools than white and Asian students, especially at many top-performing schools. The residency match program has been shown to have similar shortcomings in racial and gender equity.

"The real issue isn't the algorithms themselves" but the ranking data they use and the environment in which they're deployed, explains Utku Ünver, a professor of economics at Boston College. This dynamic has been visible throughout the ongoing artificial intelligence boom: as complex algorithms learn to reproduce patterns in our data, they often replicate our prejudices, too.

"If humans are biased, then our bias is in the data," says Éva Tardos, a professor of computer science at Cornell University. Researchers have suggested a few ways to counteract the bias in the data. For example, institutions such as hospitals or schools could use larger and more diverse panels of judges to rank the candidates. Additional algorithms could also be used to account for known biases by reweighting the rankings before they are used for matching.

Of course, no algorithm can guarantee a happy match in marriage or in schooling. But the best option

continues to be one that's simple, transparent and incentivizes honesty —so even 60 years later, it seems there's still no beating the Gale-Shapley algorithm.

MAX SPRINGER is a Ph.D. candidate in applied mathematics at the University of Maryland and was a 2024 AAAS Mass Media Fellow at *Scientific American*.