The Detection of Defective Members of Large Populations

November 21, 2019

This is me

PhD Student at Stanford, ex-engineer

NOTES

This section is devoted to brief research and expository articles, notes on methodology and other short items.

THE DETECTION OF DEFECTIVE MEMBERS OF LARGE POPULATIONS

BY ROBERT DORFMAN

Washington, D. C.

The inspection of the individual members of a large population is an expensive and tedious process. Often in testing the results of manufacture the work can be reduced greatly by examining only a sample of the population and rejecting the whole if the proportion of defectives in the sample is unduly large. In many inspections, however, the objective is to eliminate all the defective members of the population. This situation arises in manufacturing processes where the defect being tested for can result in disastrous failures. It also arises in certain inspections of human populations. Where the objective is to weed out individual defective units, a sample inspection will clearly not suffice. It will be shown in this paper that a different statistical approach can, under certain conditions, yield significant savings in effort and expense when a complete elimination of defective units is desired.

It should be noted at the outset that when large populations are being inspected the objective of eliminating all units with a particular defect can never be fully attained. Mechanical and chemical failures and, especially, manfailures make it inevitable that mistakes will occur when many units are being examined. Although the procedure described in this paper does not directly attack the problem of technical and psychological fallibility, it may contribute to its partial solution by reducing the tediousness of the work and by making more elaborate and more sensitive inspections economically feasible. In the following discussion no attention will be paid to the possibility of technical failure or operators' error.

Outline

- The paper
- What makes the paper work?
- How its ideas can be reused

NOTES

This section is devoted to brief research and expository articles, notes on methodology and other short items.

THE DETECTION OF DEFECTIVE MEMBERS OF LARGE POPULATIONS

BY ROBERT DORFMAN

Washington, D. C.

The inspection of the individual members of a large population is an expensive and tedious process. Often in testing the results of manufacture the work can be reduced greatly by examining only a sample of the population and rejecting the whole if the proportion of defectives in the sample is unduly large. In many inspections, however, the objective is to eliminate all the defective members of the population. This situation arises in manufacturing processes where the defect being tested for can result in disastrous failures. It also arises in certain inspections of human populations. Where the objective is to weed out individual defective units, a sample inspection will clearly not suffice. It will be shown in this paper that a different statistical approach can, under certain conditions, yield significant savings in effort and expense when a complete elimination of defective units is desired.

It should be noted at the outset that when large populations are being inspected the objective of eliminating all units with a particular defect can never be fully attained. Mechanical and chemical failures and, especially, manfailures make it inevitable that mistakes will occur when many units are being examined. Although the procedure described in this paper does not directly attack the problem of technical and psychological fallibility, it may contribute to its partial solution by reducing the tediousness of the work and by making more elaborate and more sensitive inspections economically feasible. In the following discussion no attention will be paid to the possibility of technical failure or operators' error.

NOTES

This section is devoted to brief research and expository articles, notes on methodology and other short items.

THE DETECTION OF DEFECTIVE MEMBERS OF LARGE POPULATIONS

By Robert Dorfman

Washington, D. C.

The inspection of the individual members of a large population is an expensive and tedious process. Often in testing the results of manufacture the work can be reduced greatly by examining only a sample of the population and rejecting the whole if the proportion of defectives in the sample is unduly large. In many inspections, however, the objective is to eliminate all the defective members of the population. This situation arises in manufacturing processes where the defect being tested for can result in disastrous failures. It also arises in certain inspections of human populations. Where the objective is to weed out individual defective units, a sample inspection will clearly not suffice. It will be shown in this paper that a different statistical approach can, under certain conditions, yield significant savings in effort and expense when a complete elimination of defective units is desired.

It should be noted at the outset that when large populations are being inspected the objective of eliminating all units with a particular defect can never be fully attained. Mechanical and chemical failures and, especially, manfailures make it inevitable that mistakes will occur when many units are being examined. Although the procedure described in this paper does not directly attack the problem of technical and psychological fallibility, it may contribute to its partial solution by reducing the tediousness of the work and by making more elaborate and more sensitive inspections economically feasible. In the following discussion no attention will be paid to the possibility of technical failure or operators' error.

Group Testing

The setting is World War II…

Group Testing

The setting is World War II…

🤠

Group Testing

The setting is World War II...

🤠 🤠

Group Testing

The setting is World War II...

🤠 🤠 🤠

Group Testing

The setting is World War II...

Group Testing

The setting is World War II…

Group Testing

The setting is World War II…

Group Testing

The setting is World War II…

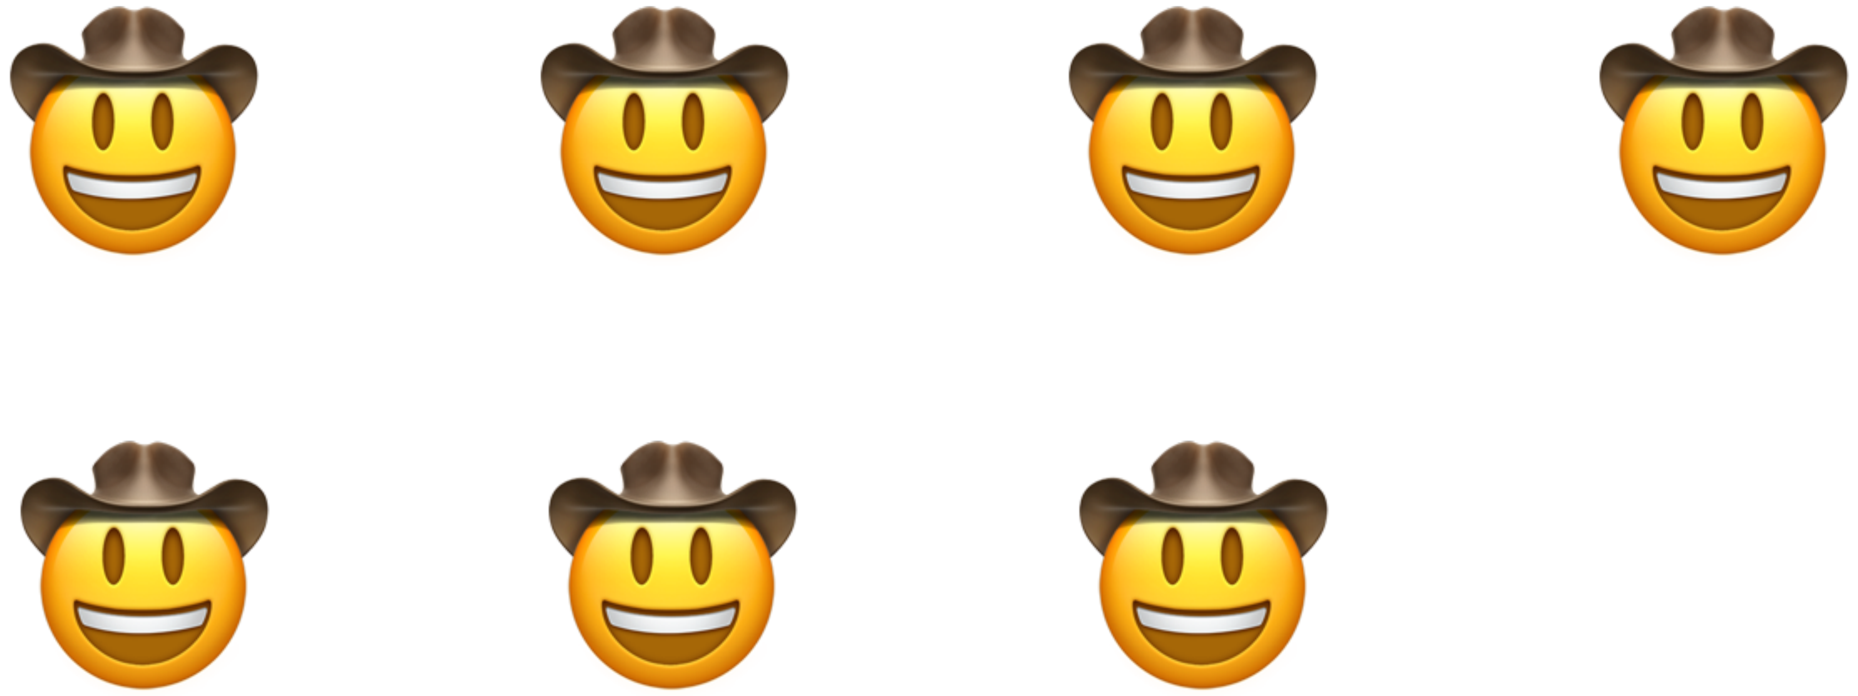

Group Testing

The setting is World War II…

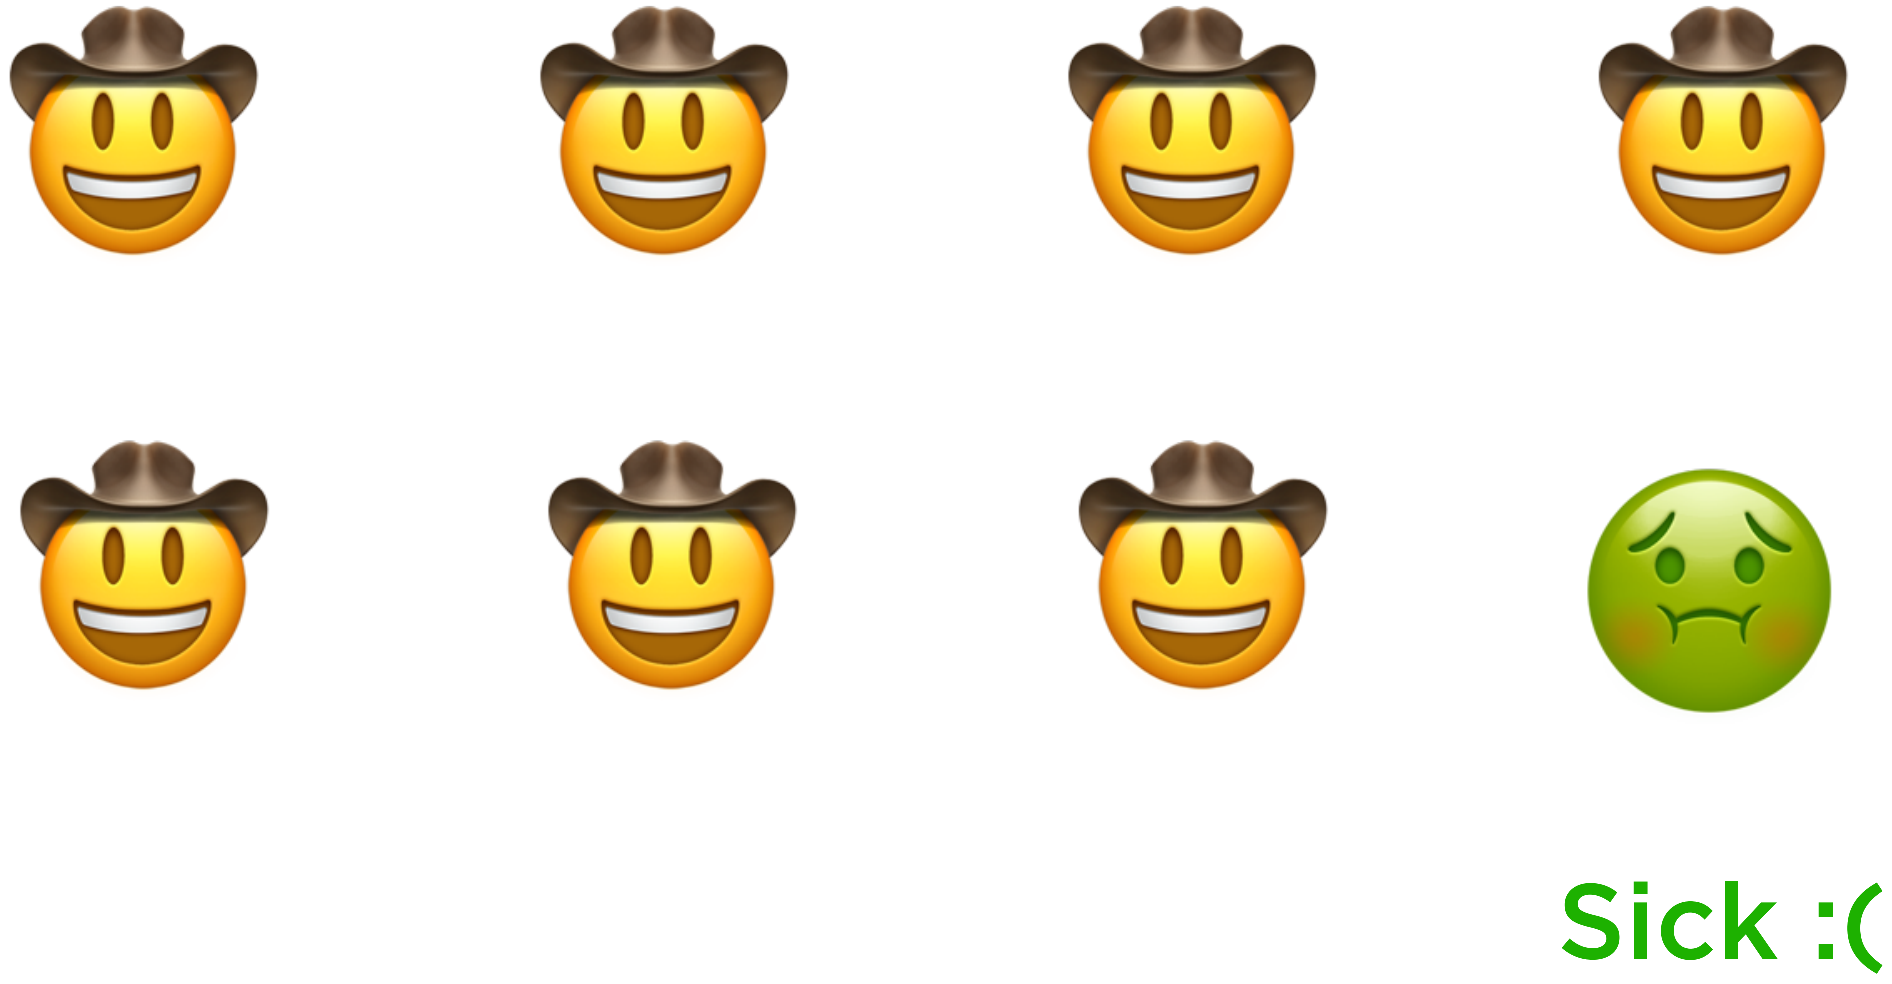

Group Testing

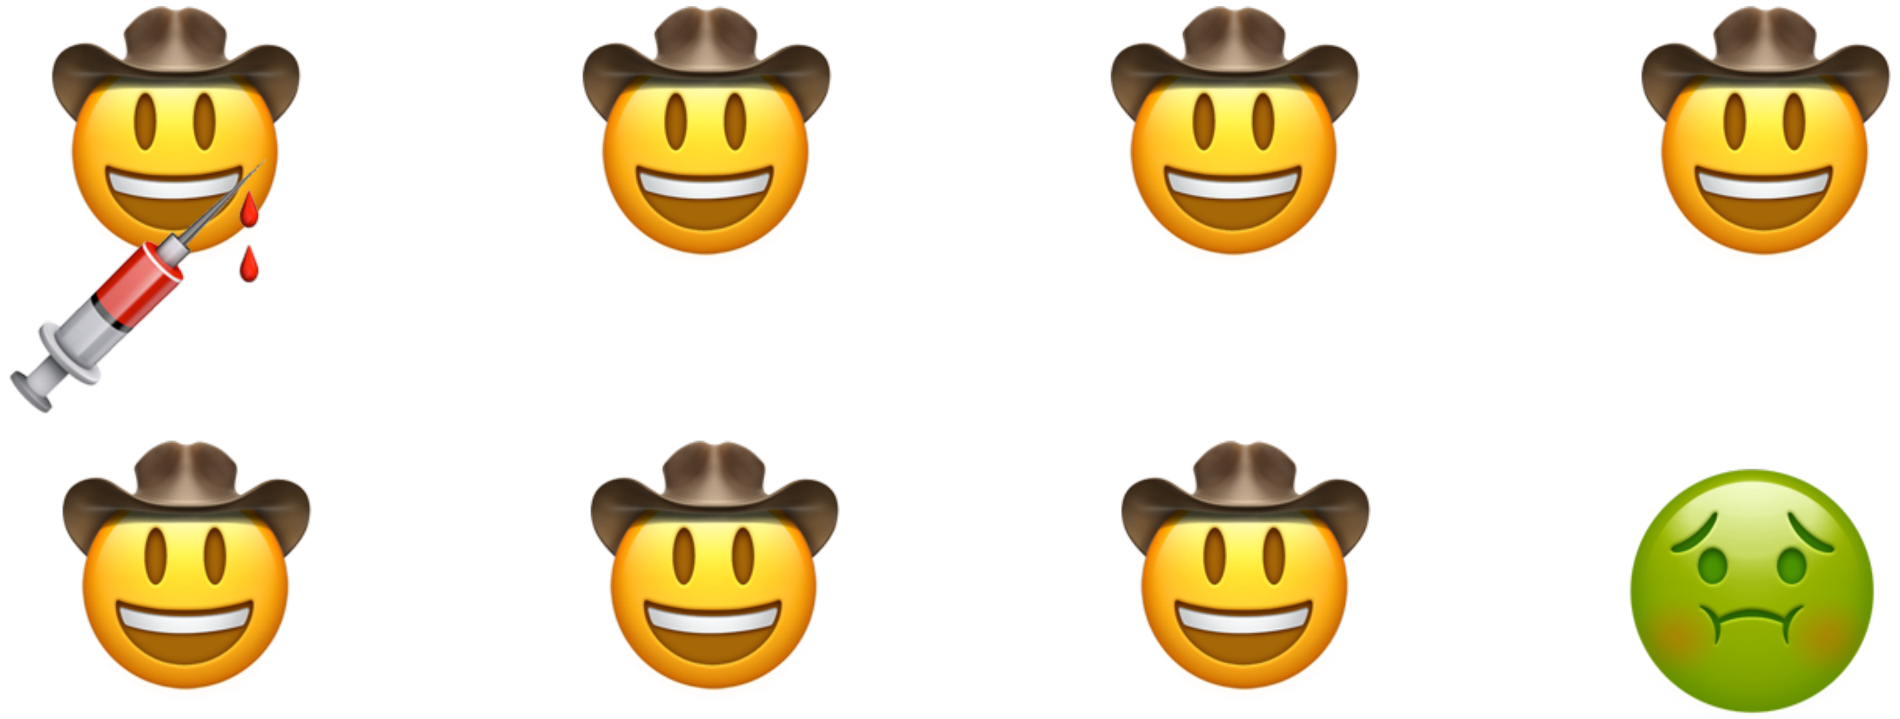

Sick :(

Group Testing

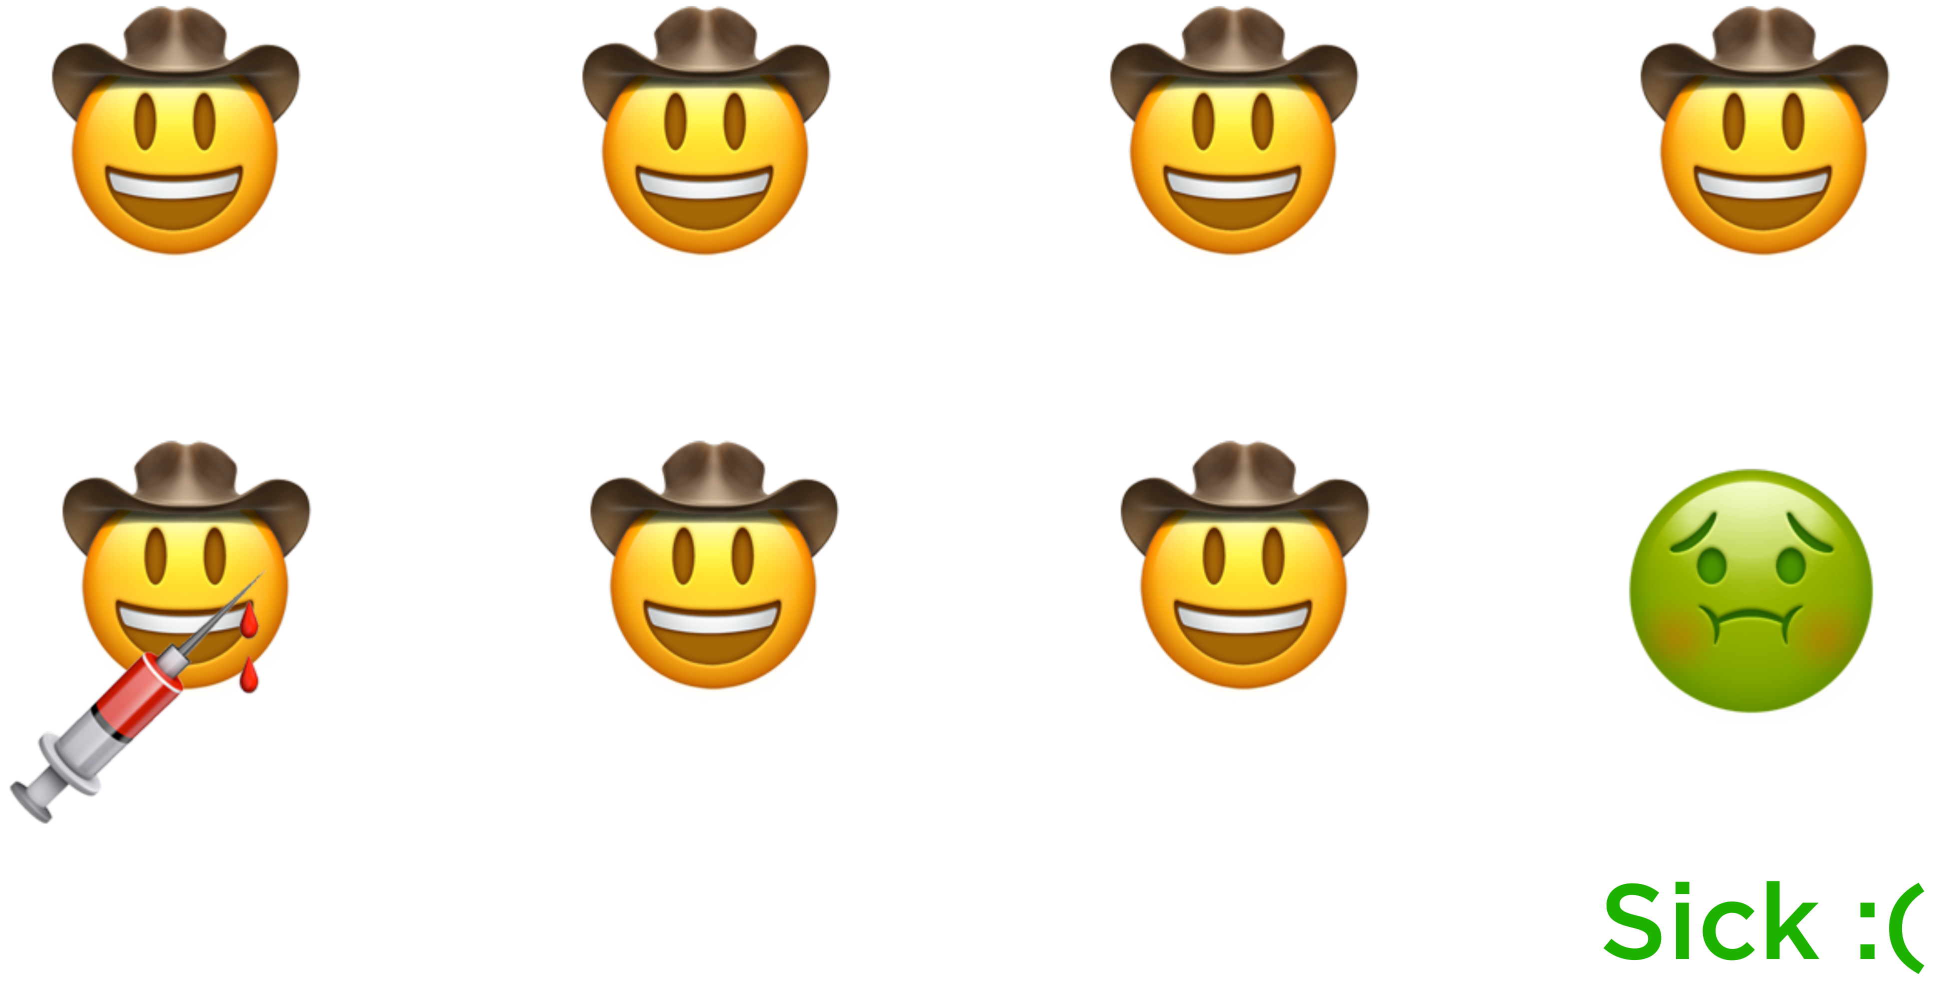

Don’t need individual tests

Don’t need individual tests

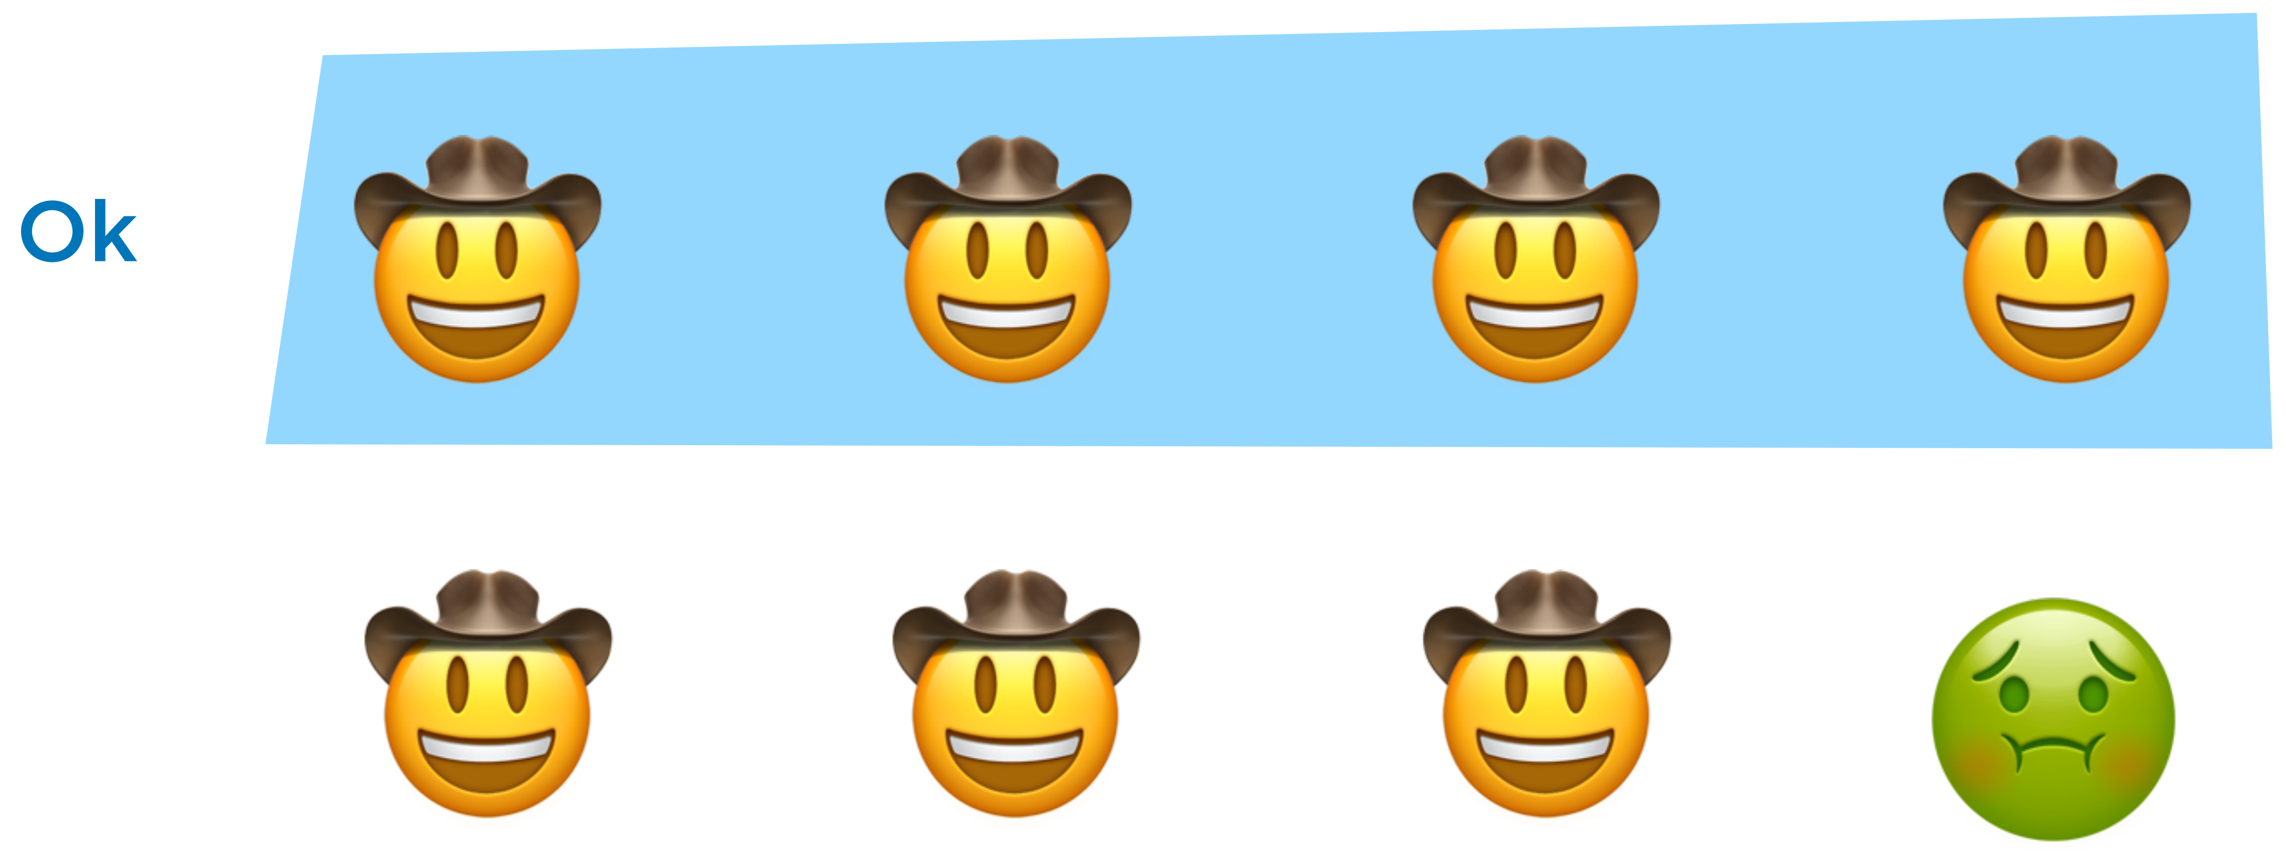

Don’t need individual tests

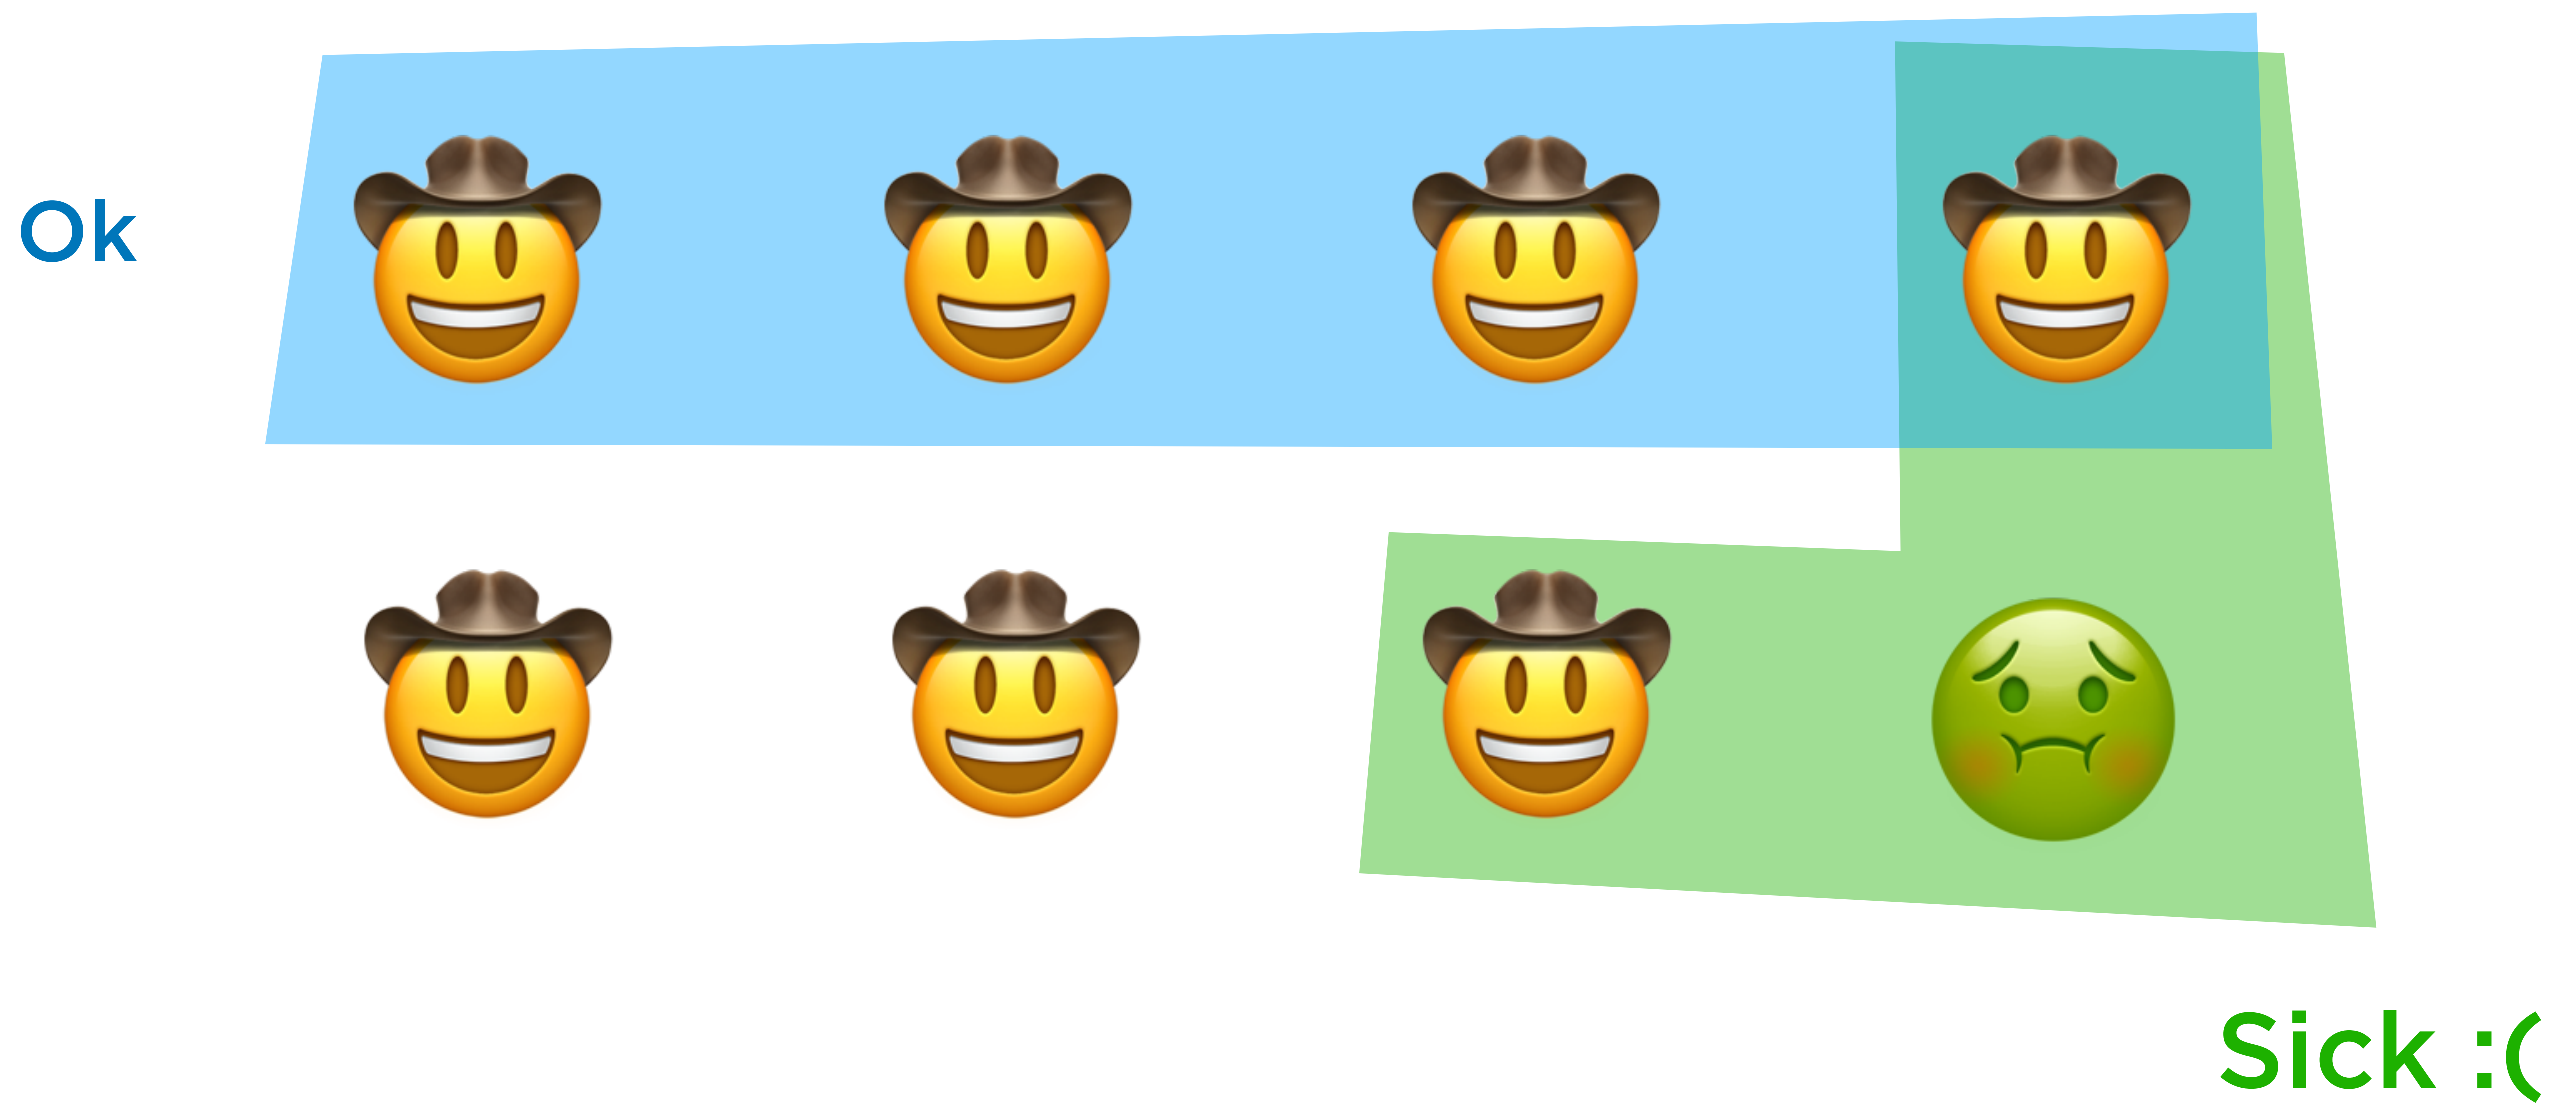

Don’t need individual tests

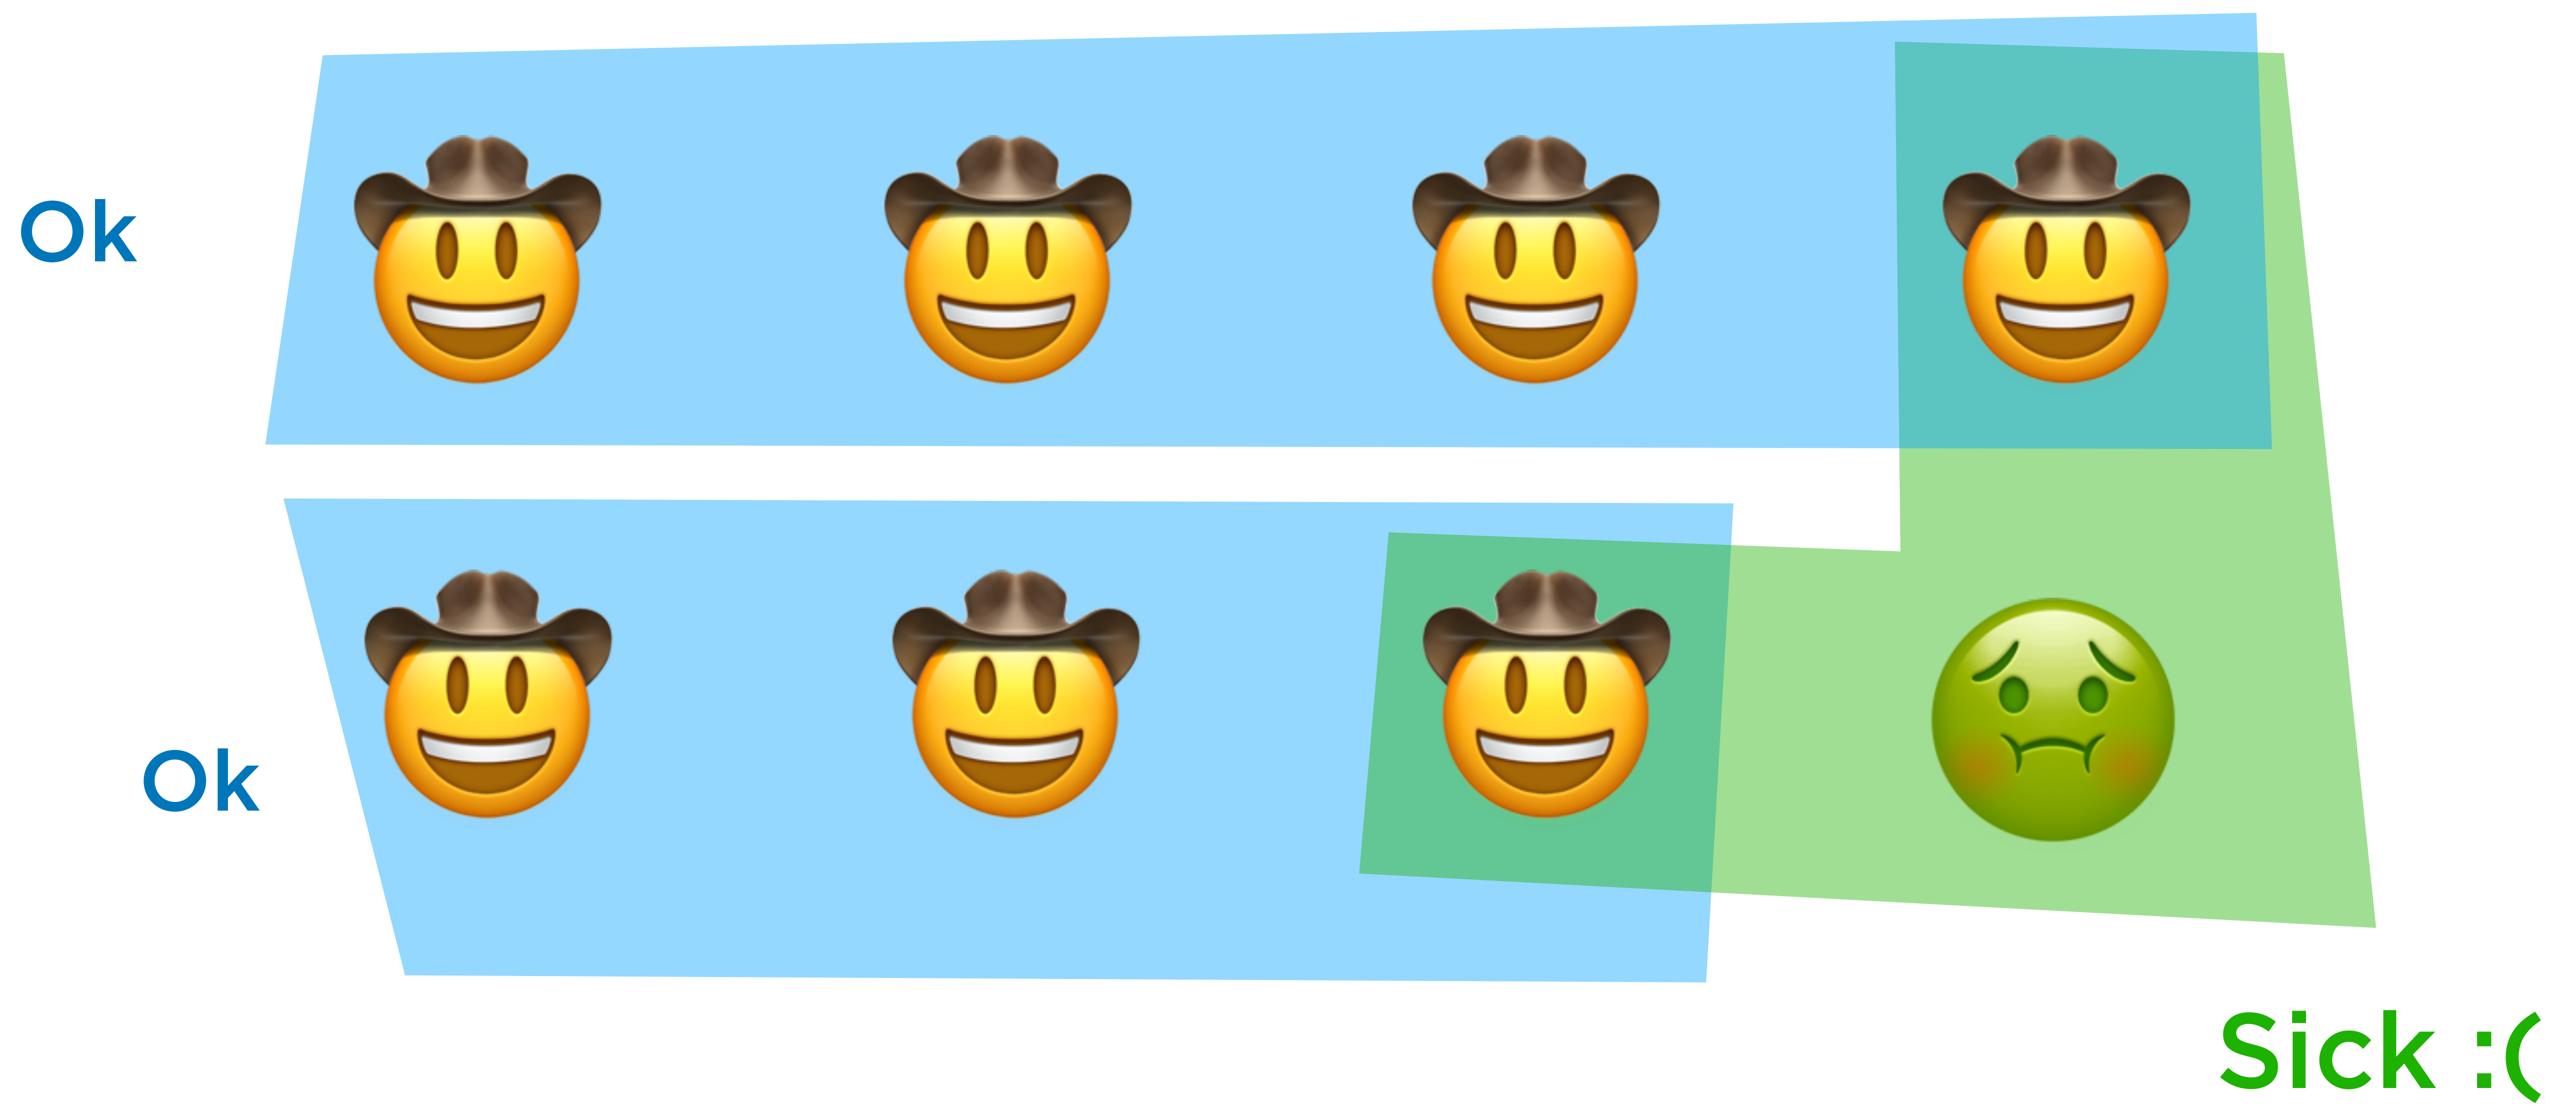

Don’t need individual tests

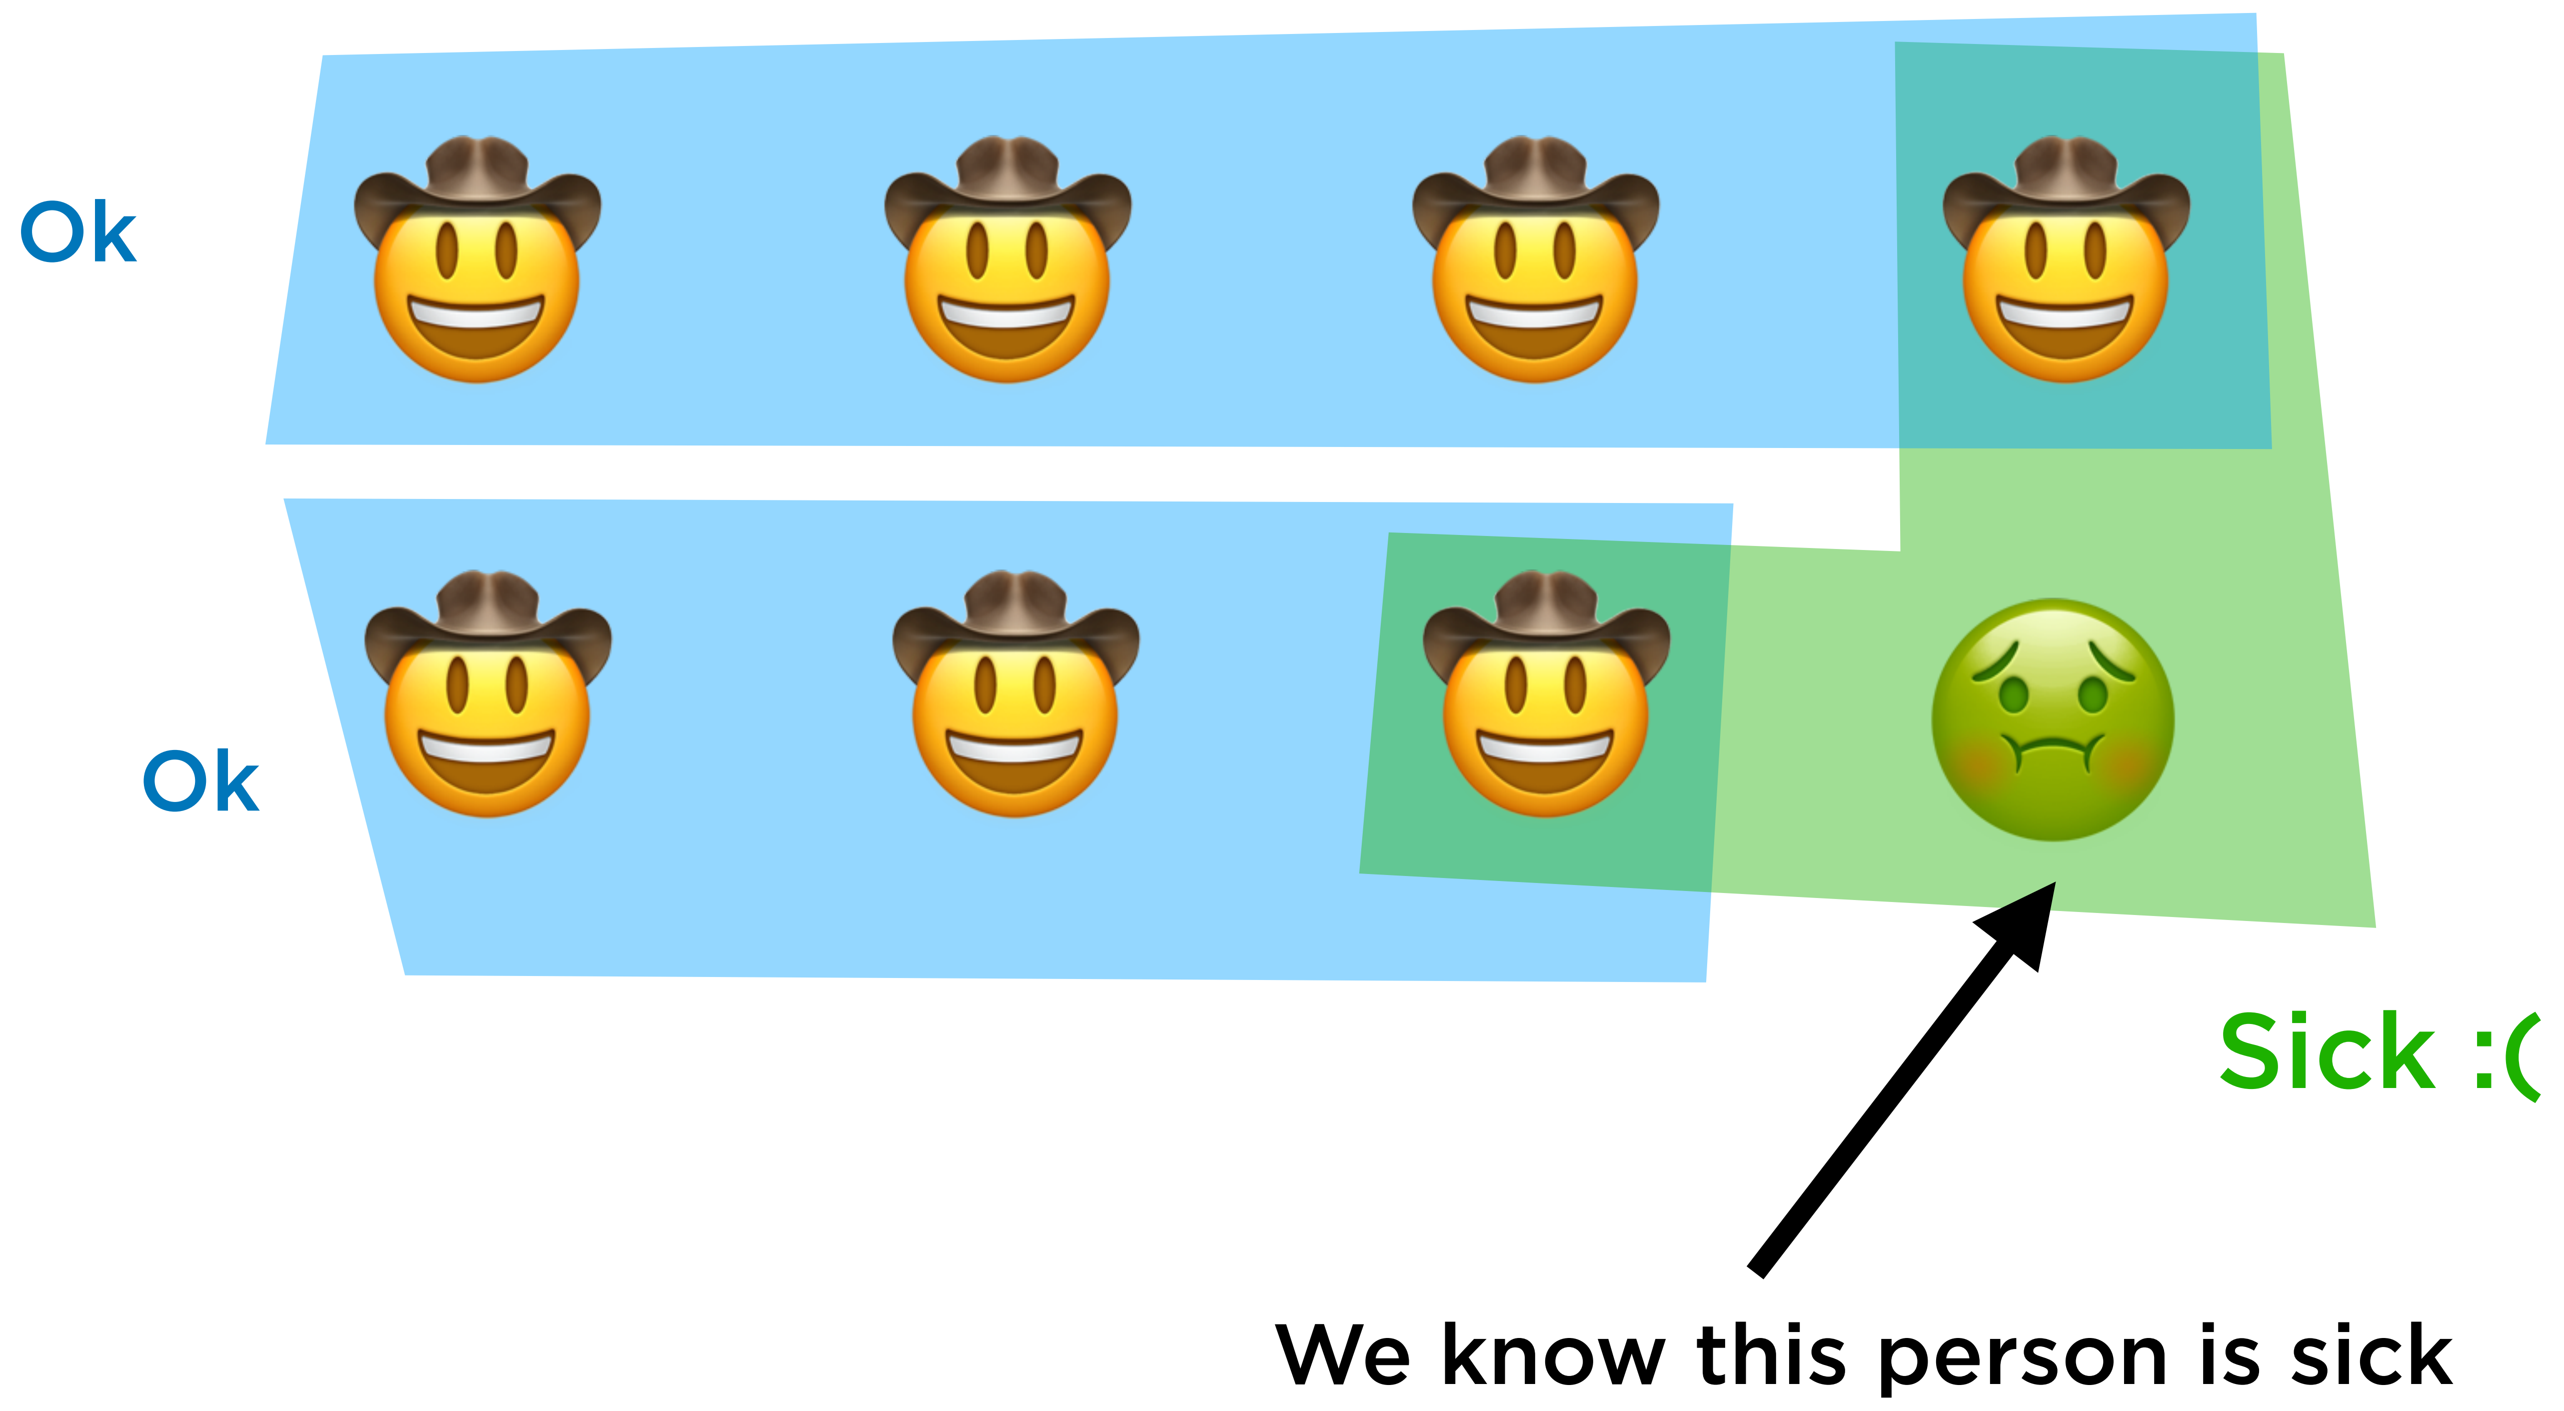

Need to carefully design tests

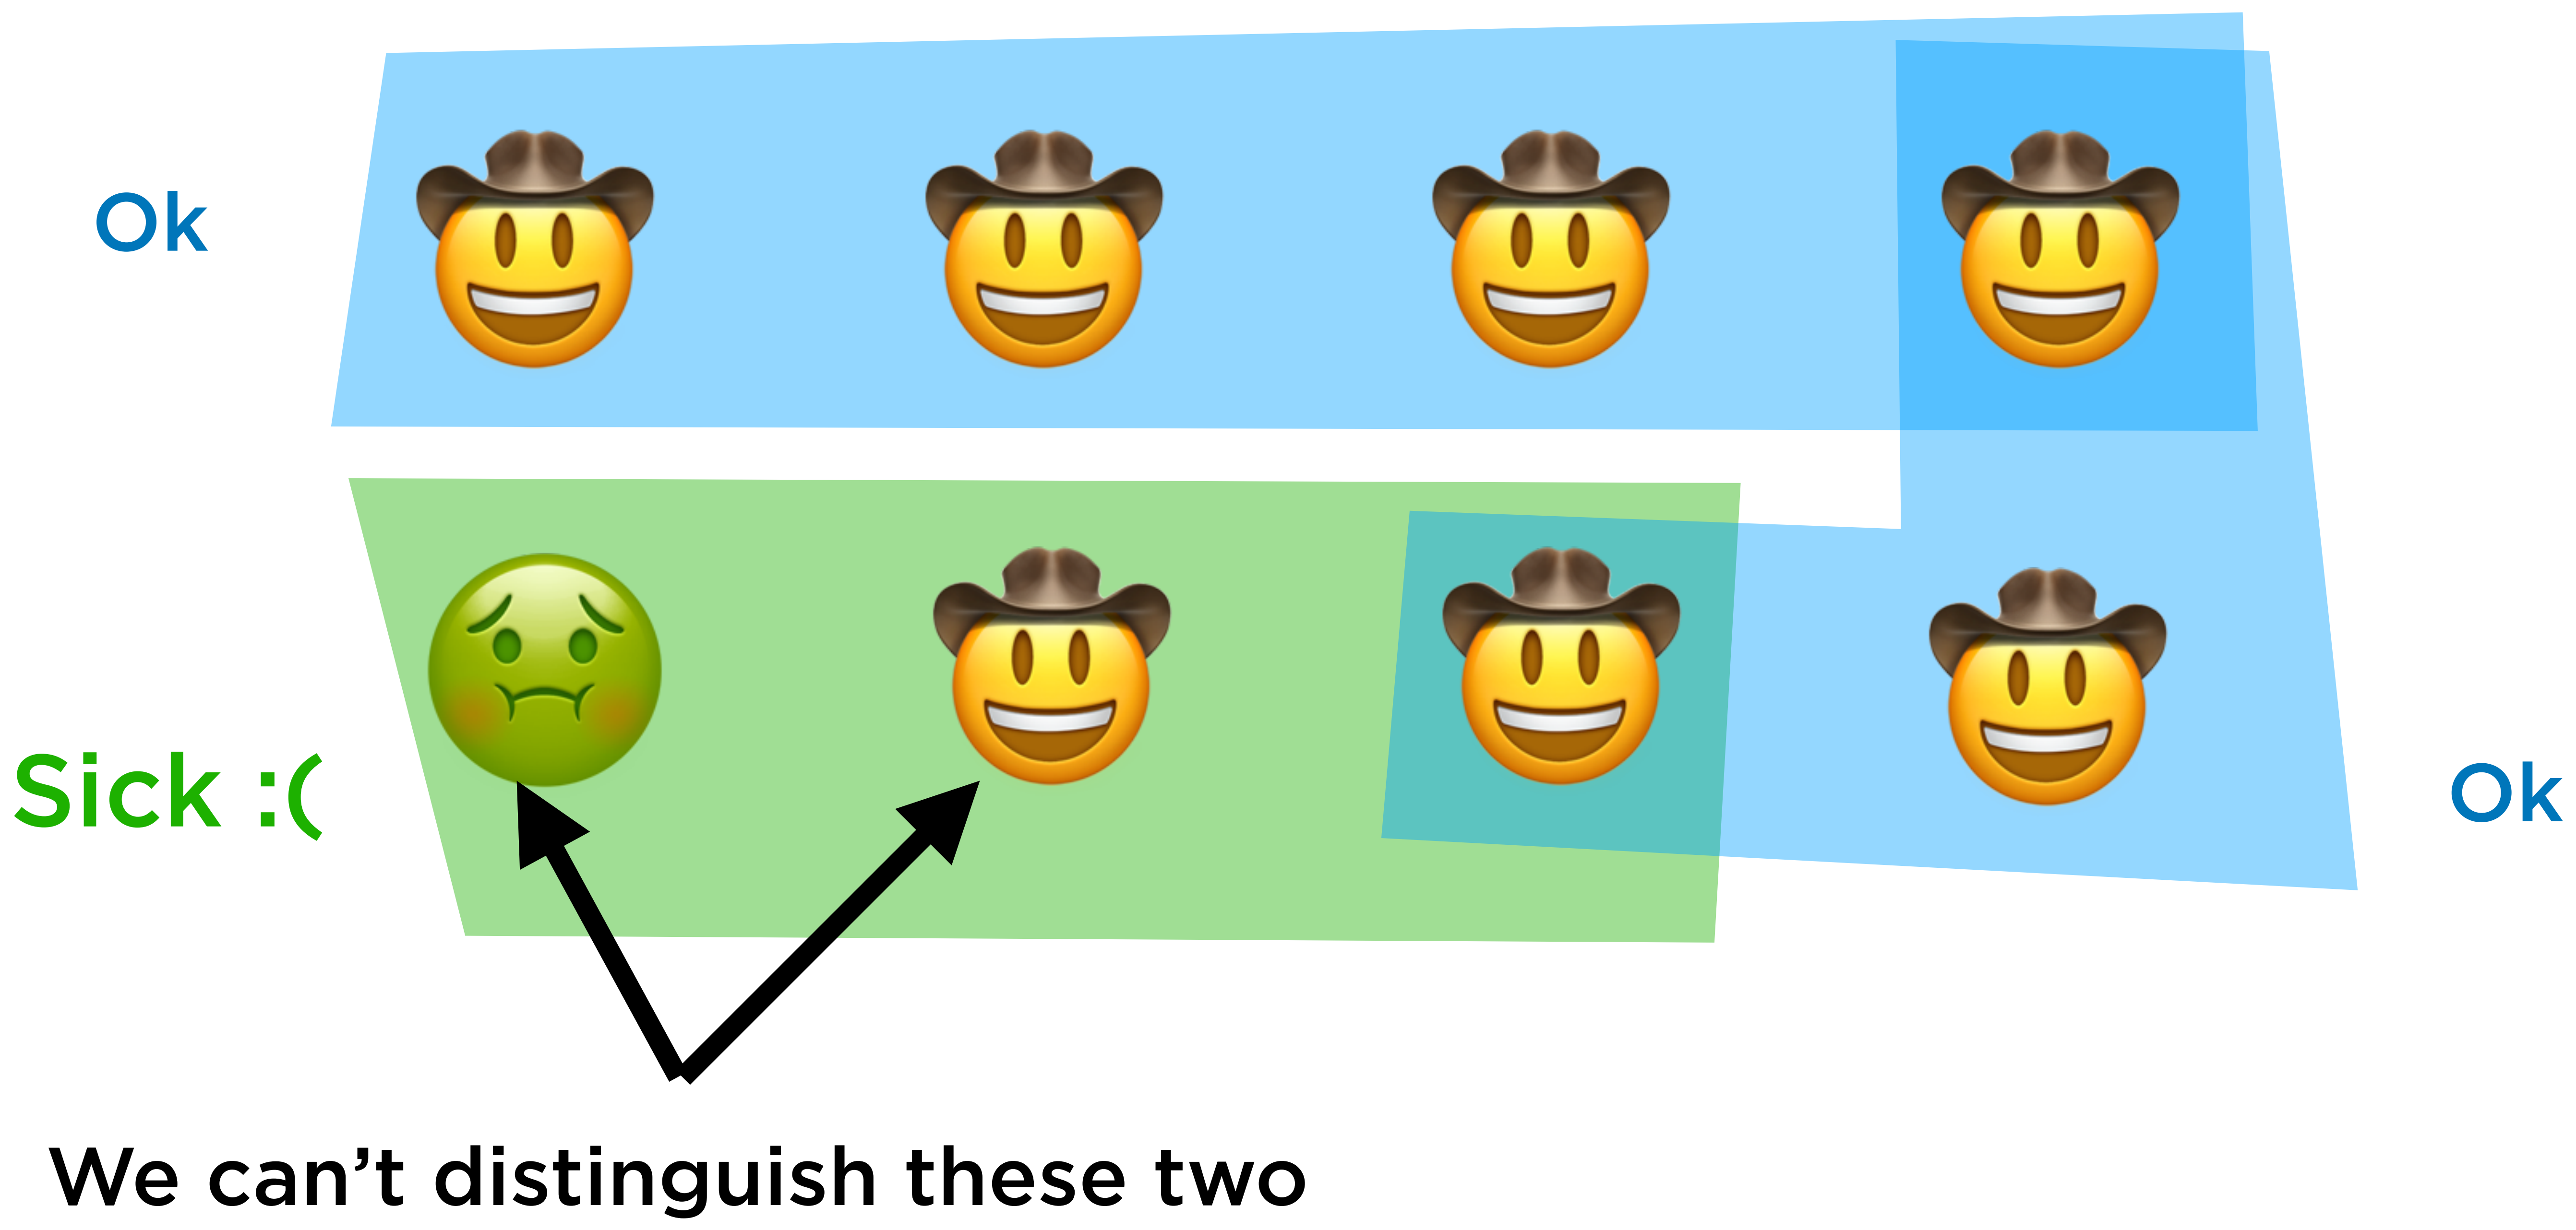

Need to carefully design tests

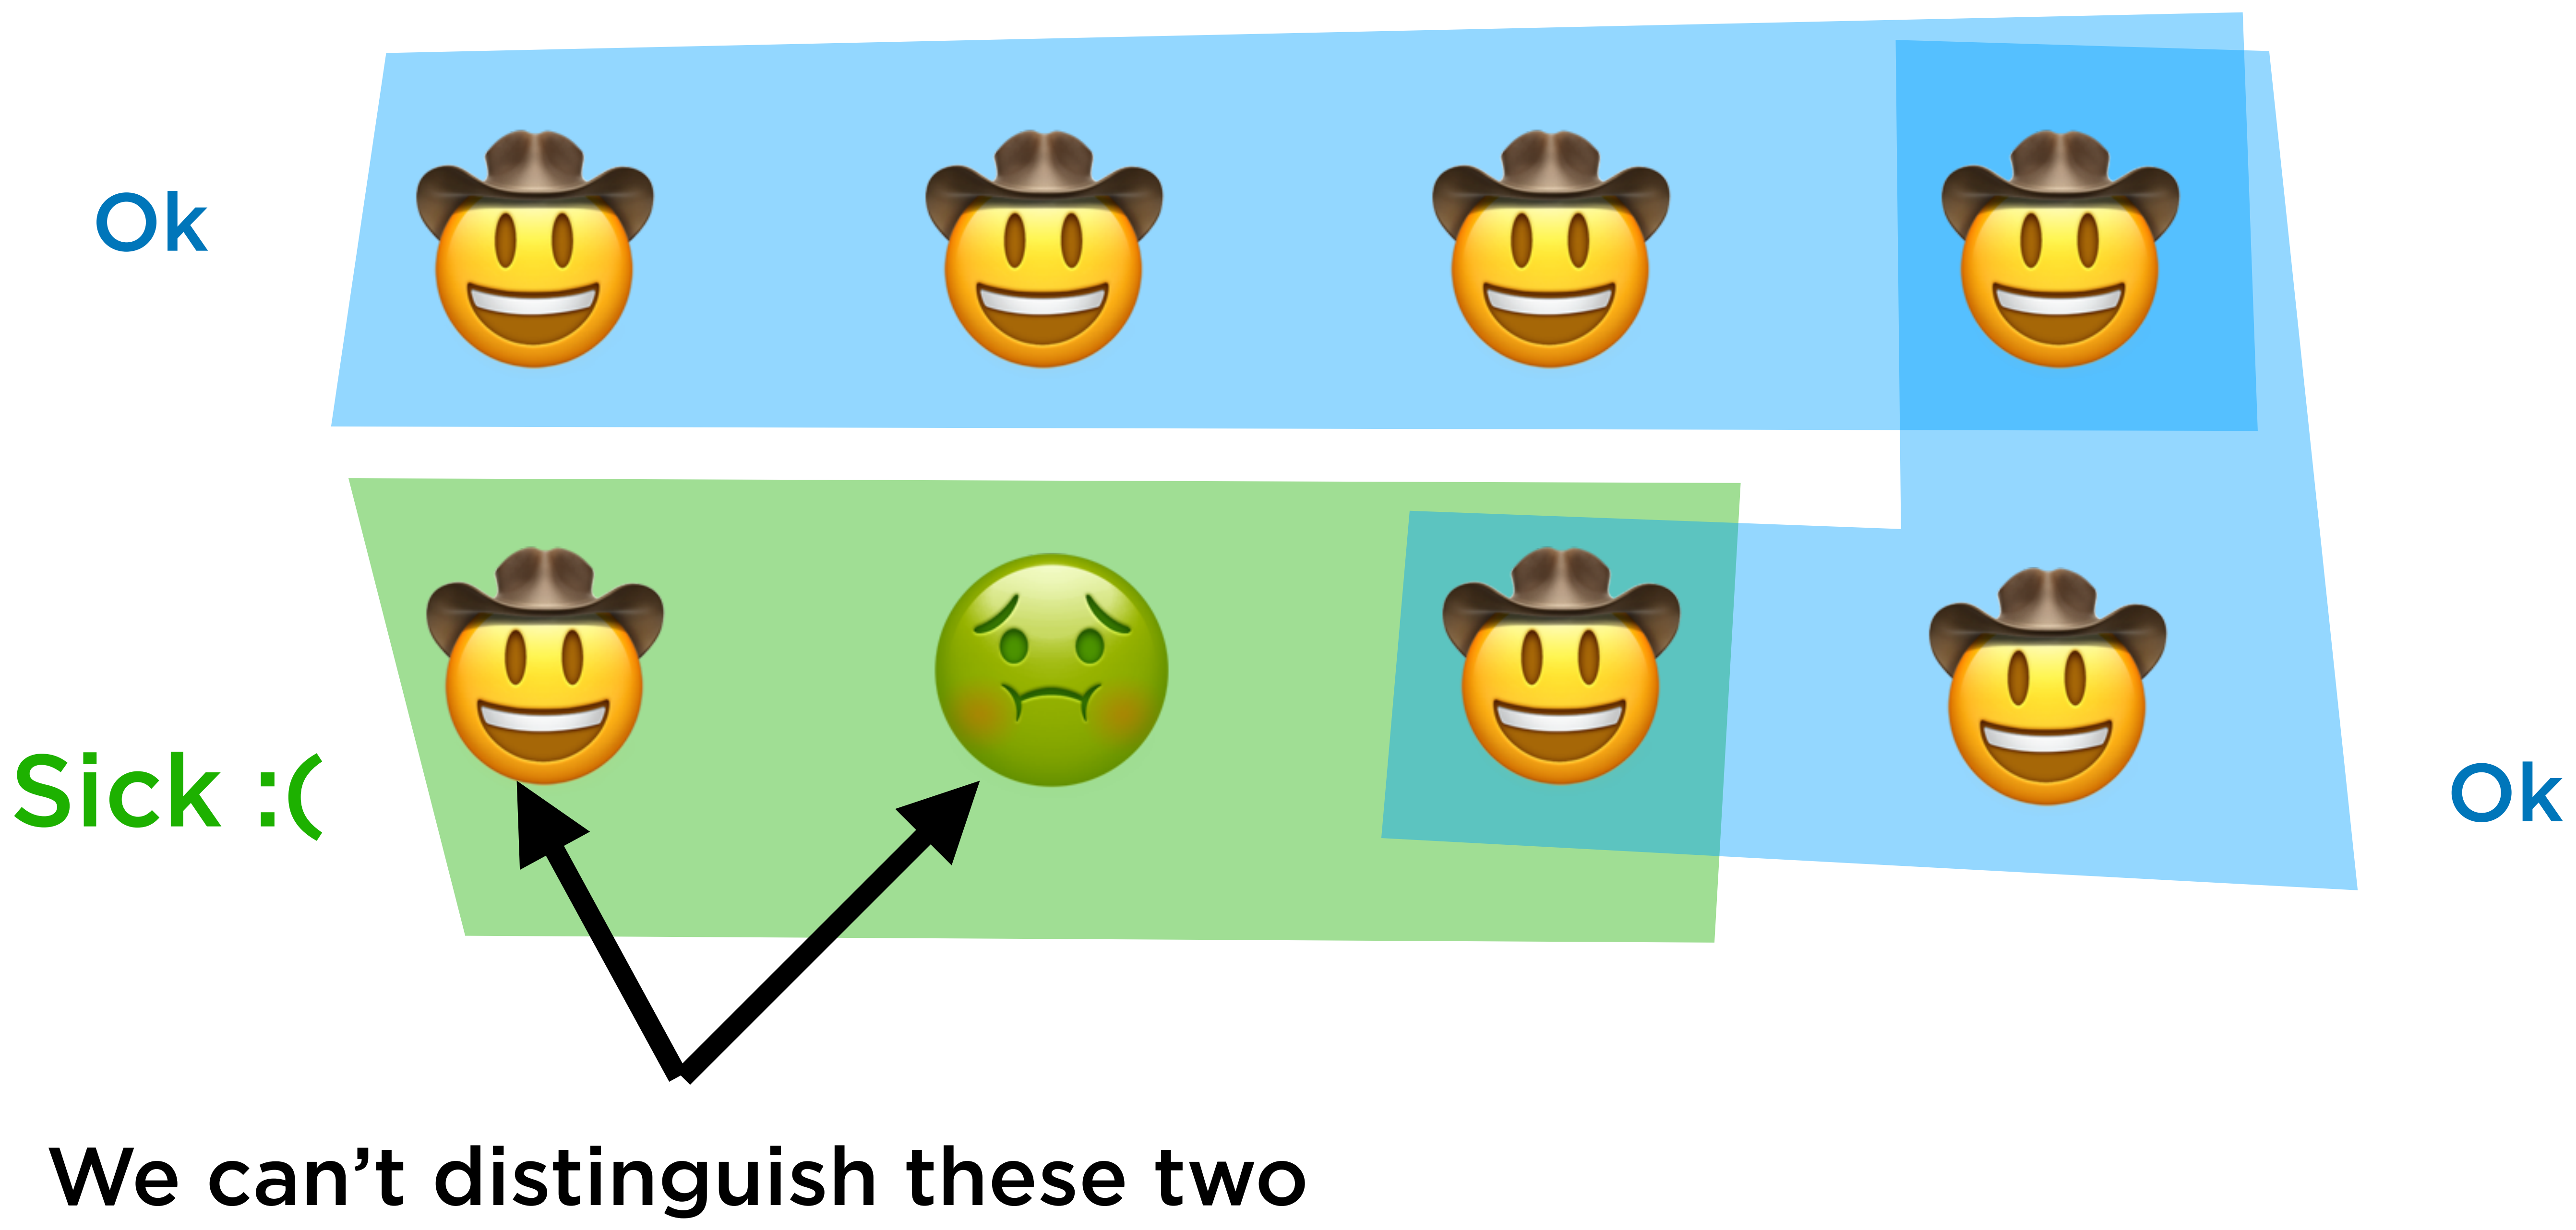

Group Testing Problem

We have n items, at most s of which are "sick."

Definition: A test returns whether a subset of items includes any sick items or not.

Problem: Construct a set of tests which can identify a worst-case set of at most s sick items.

A better design

If every column is unique, we win

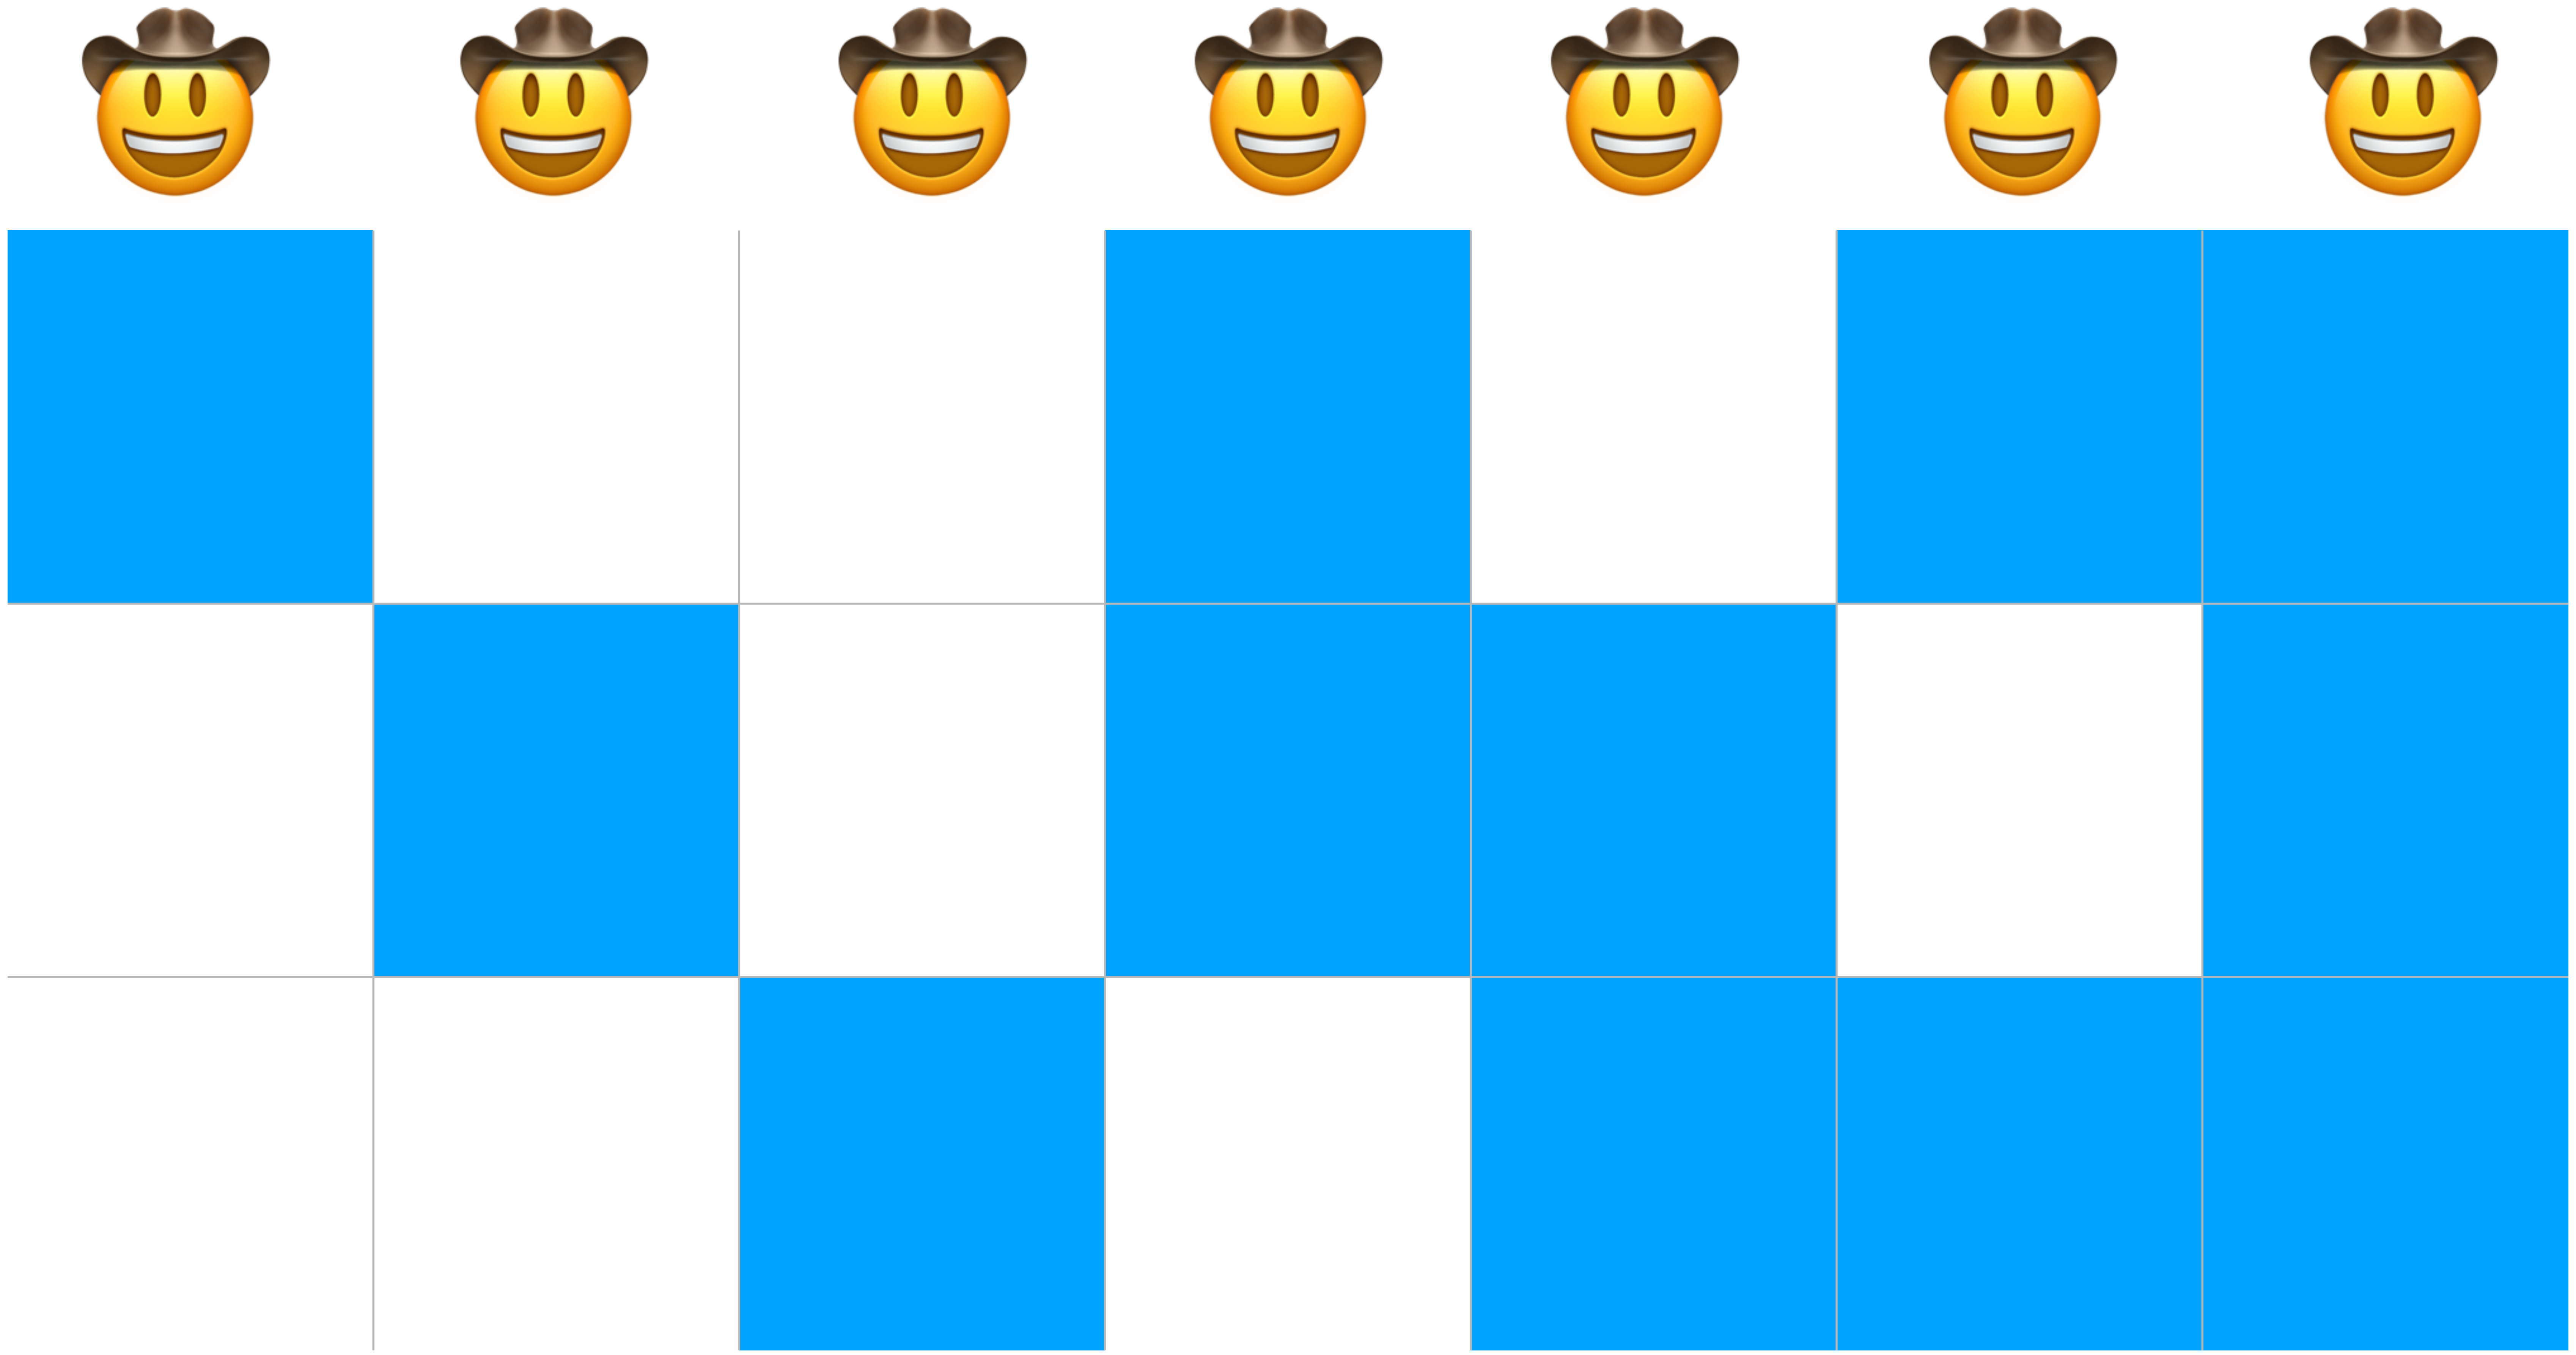

A better design

If every column is unique, we win

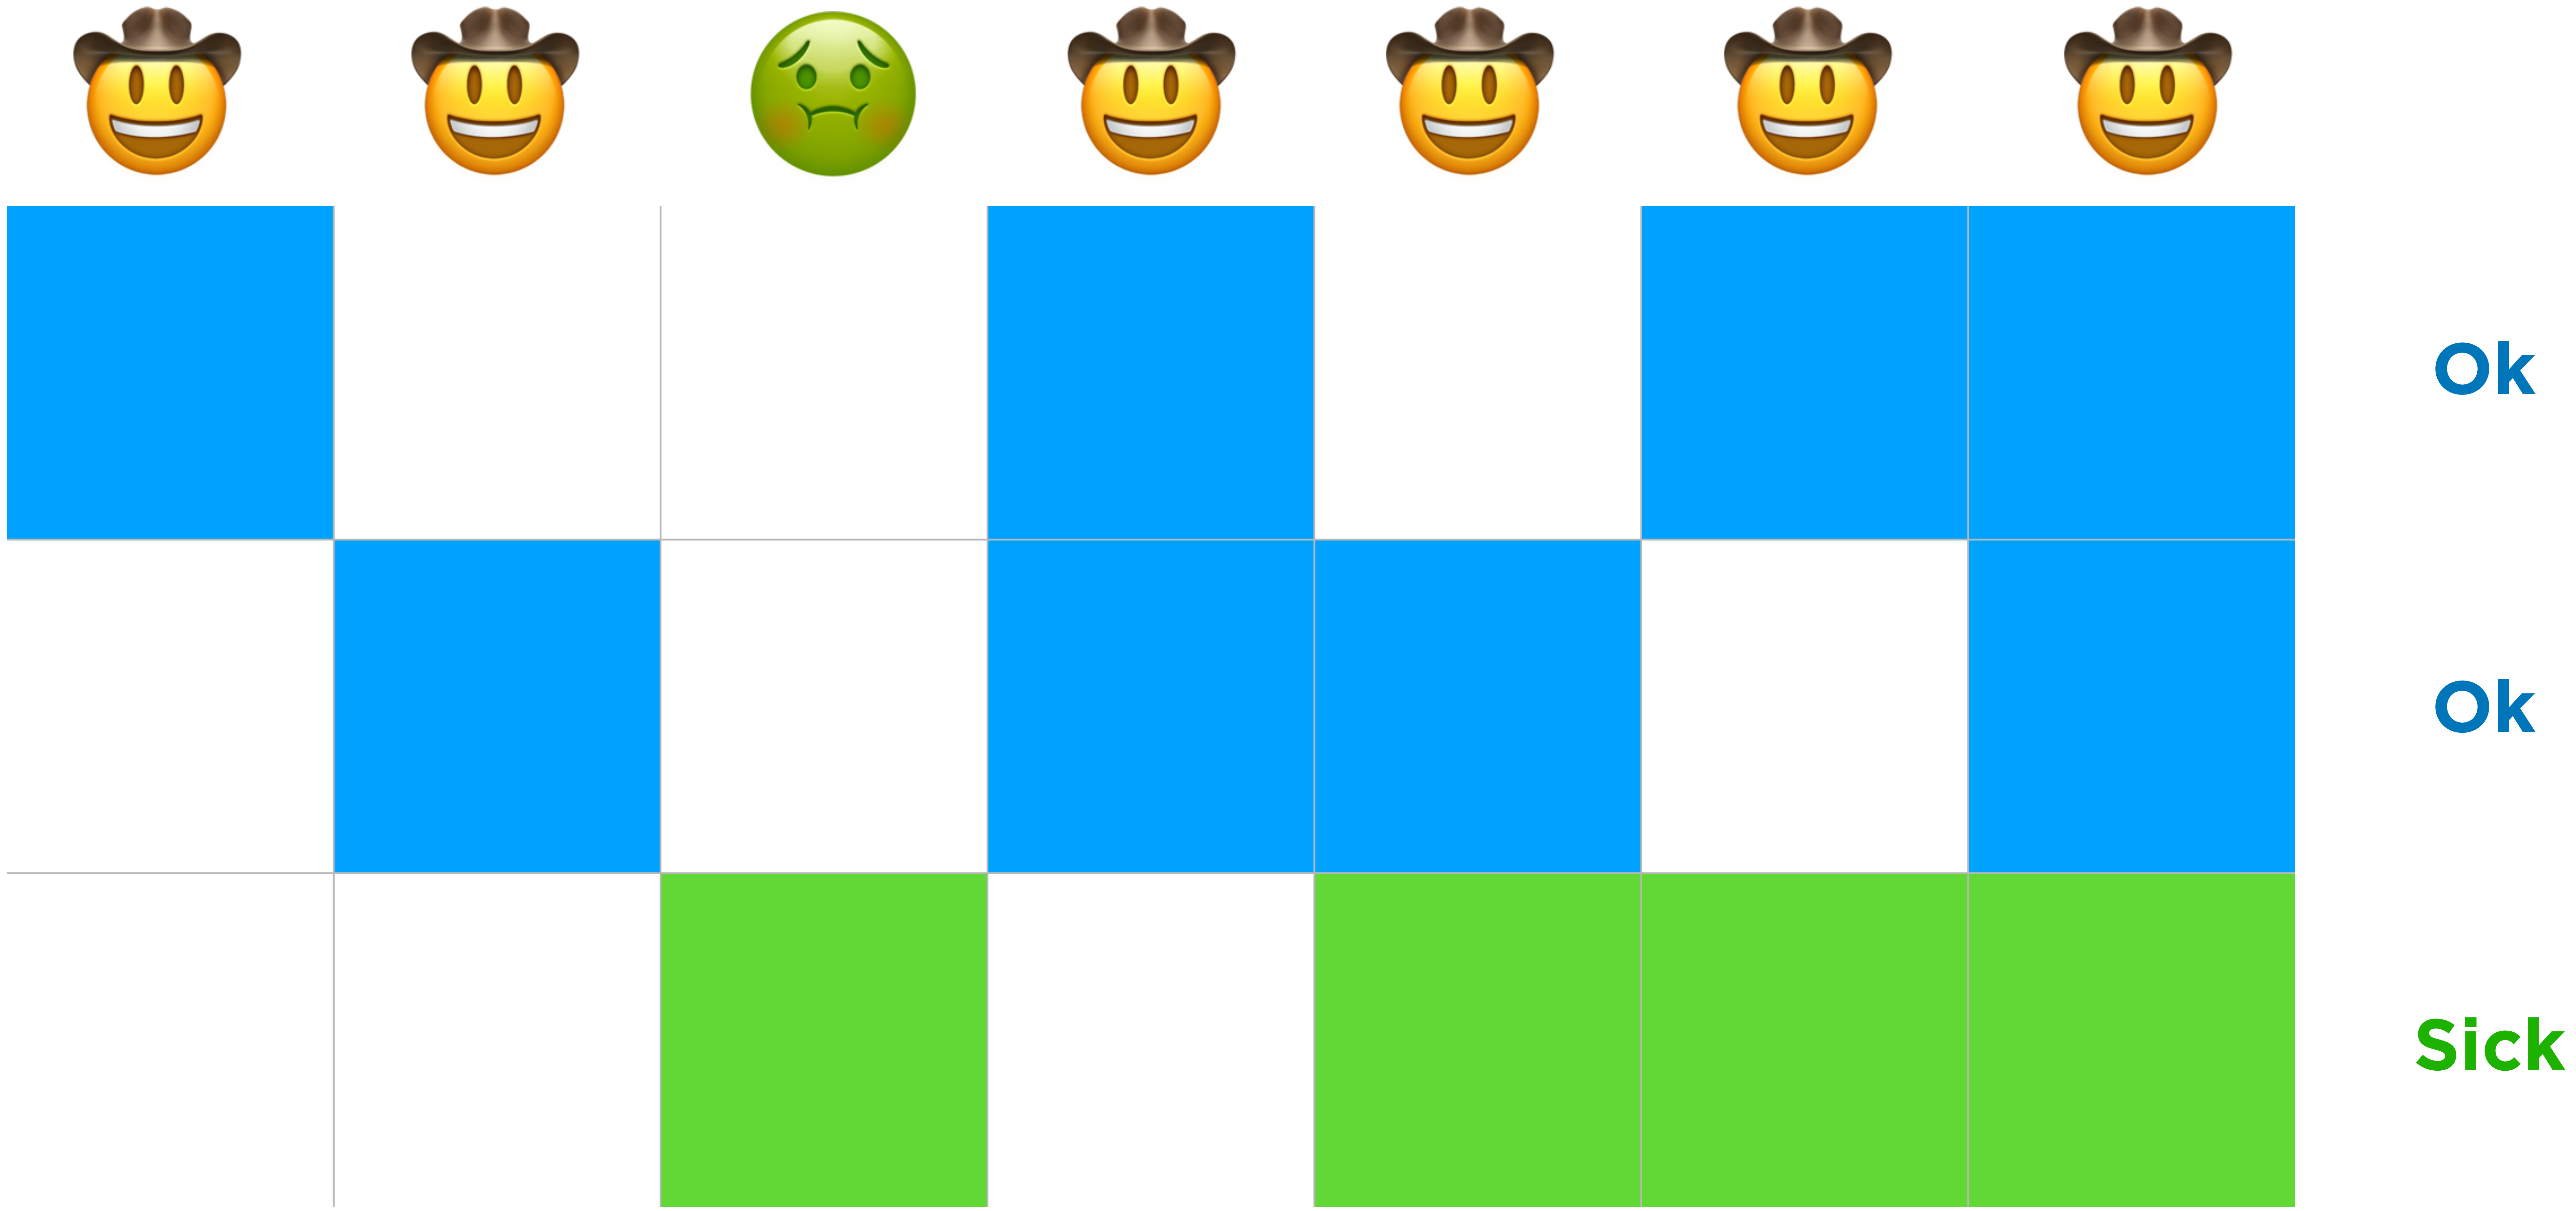

A better design

If every column is unique, we win

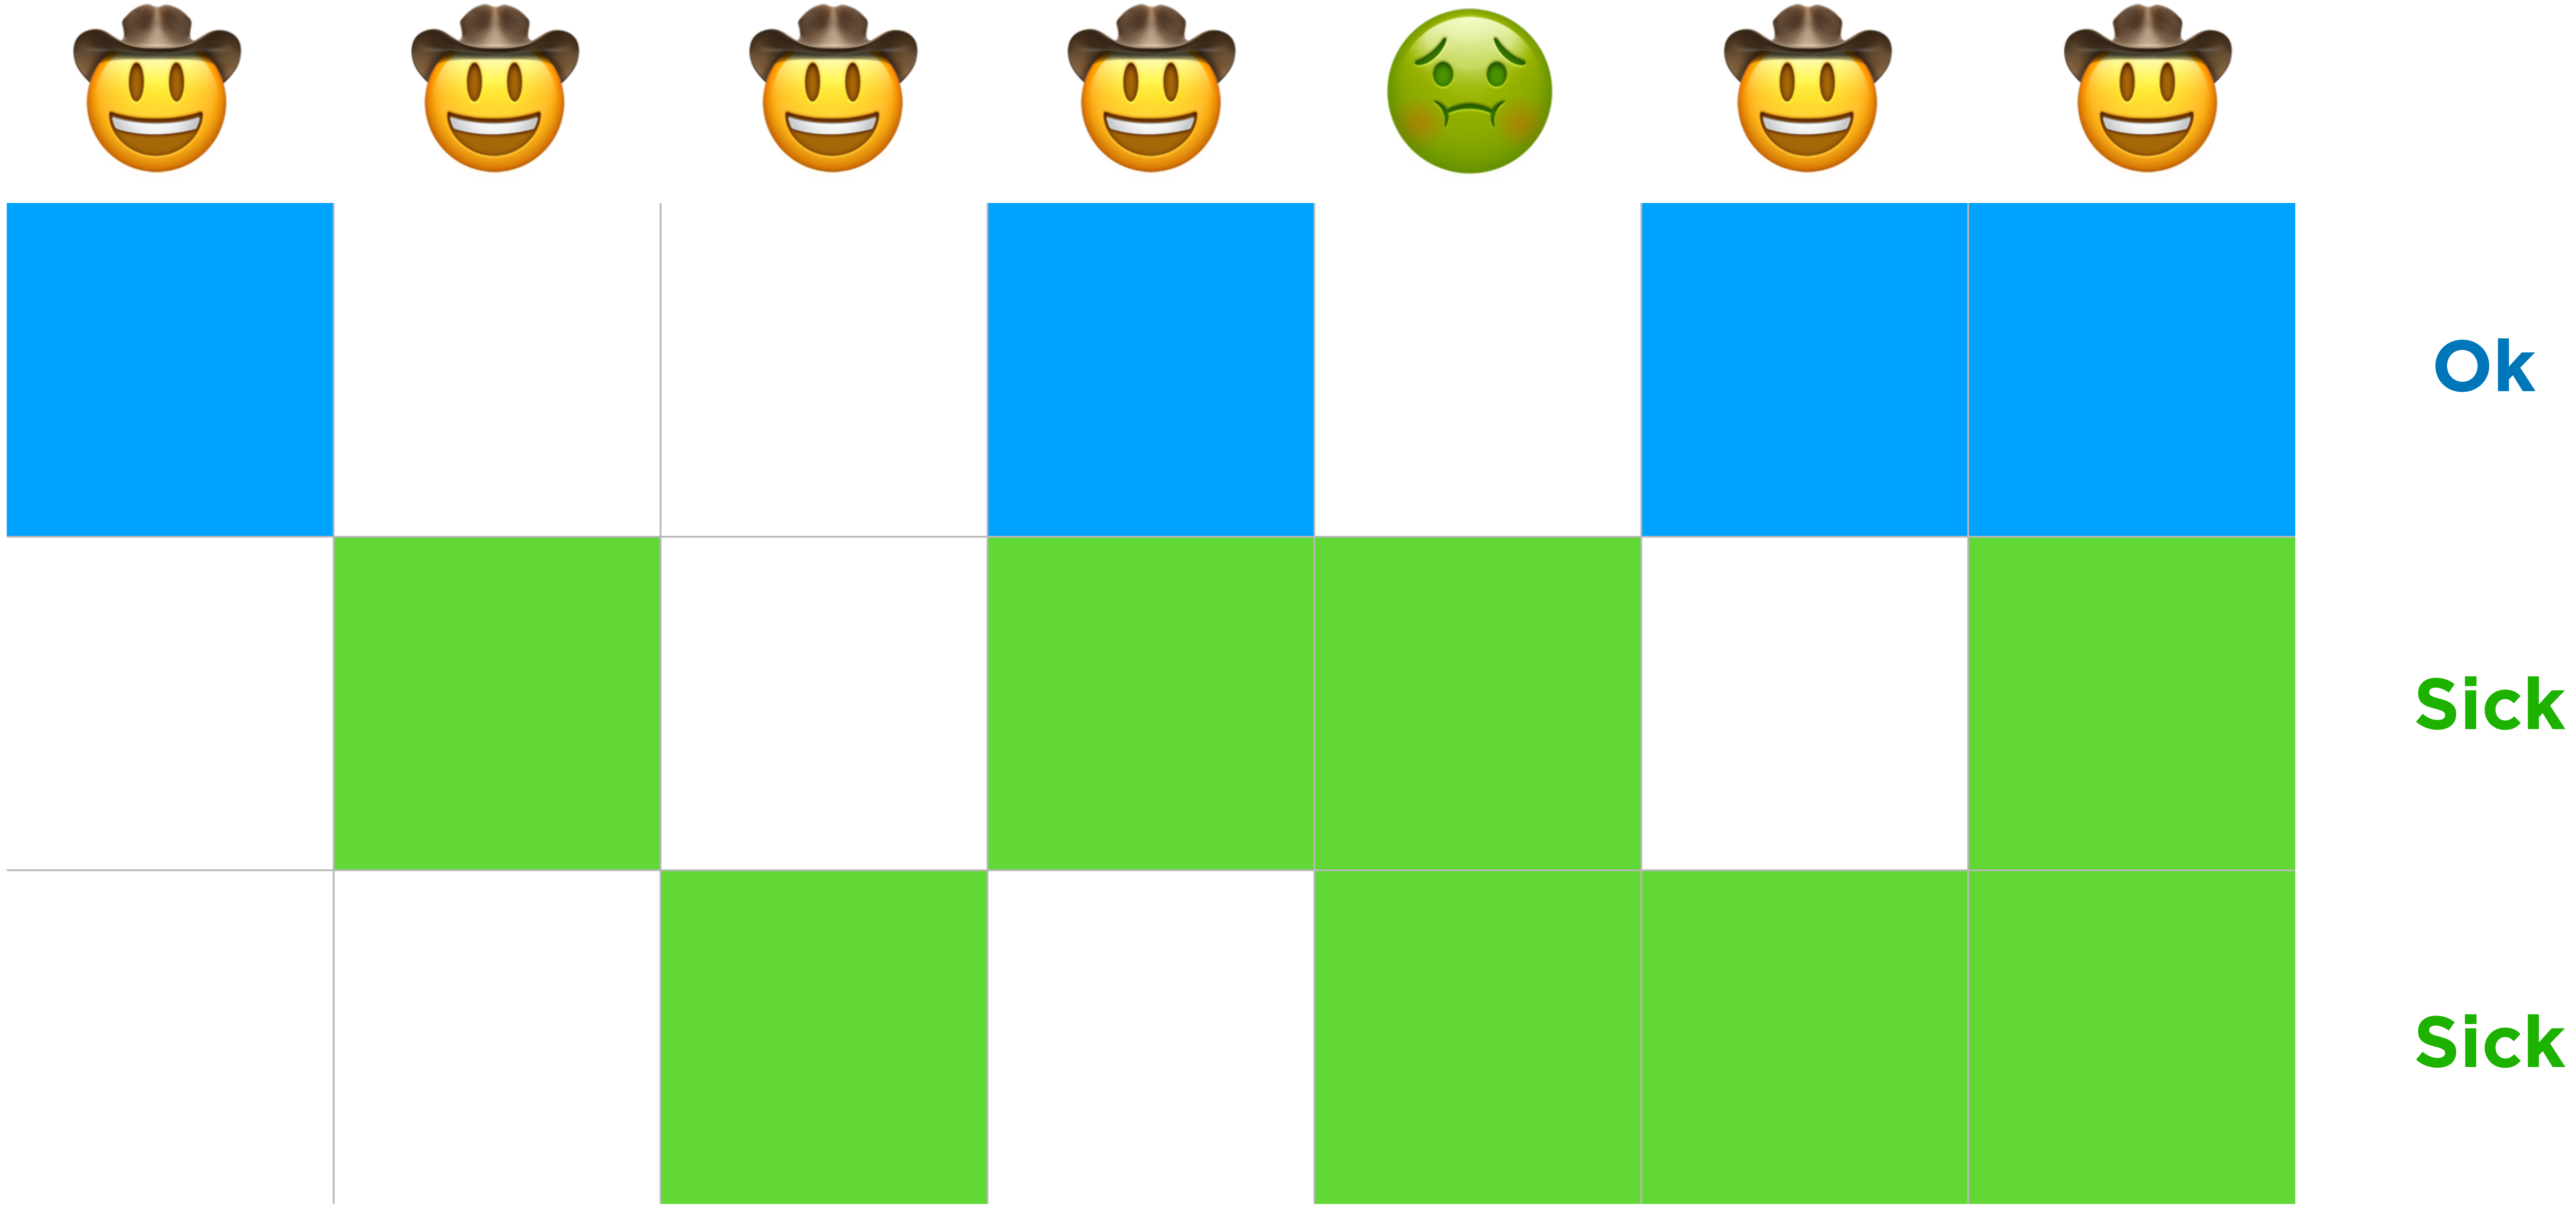

A better design

If every column is unique, we win

🤠	🤠	🤢	🤠	🤠	🤠	🤠	
1	0	0	1	0	1	1	Ok
0	1	0	1	1	0	1	Ok
0	0	1	0	1	1	1	Sick

What we just saw

If there is one sick person, we can find them non-adaptively with **log n** tests!

Dorfman's Construction

This seems hard, so let's just do something totally random

Will show this works with decent probability and **$O(s^2 \log n)$ tests**

Why is $s^2 \log n$ tests cool?

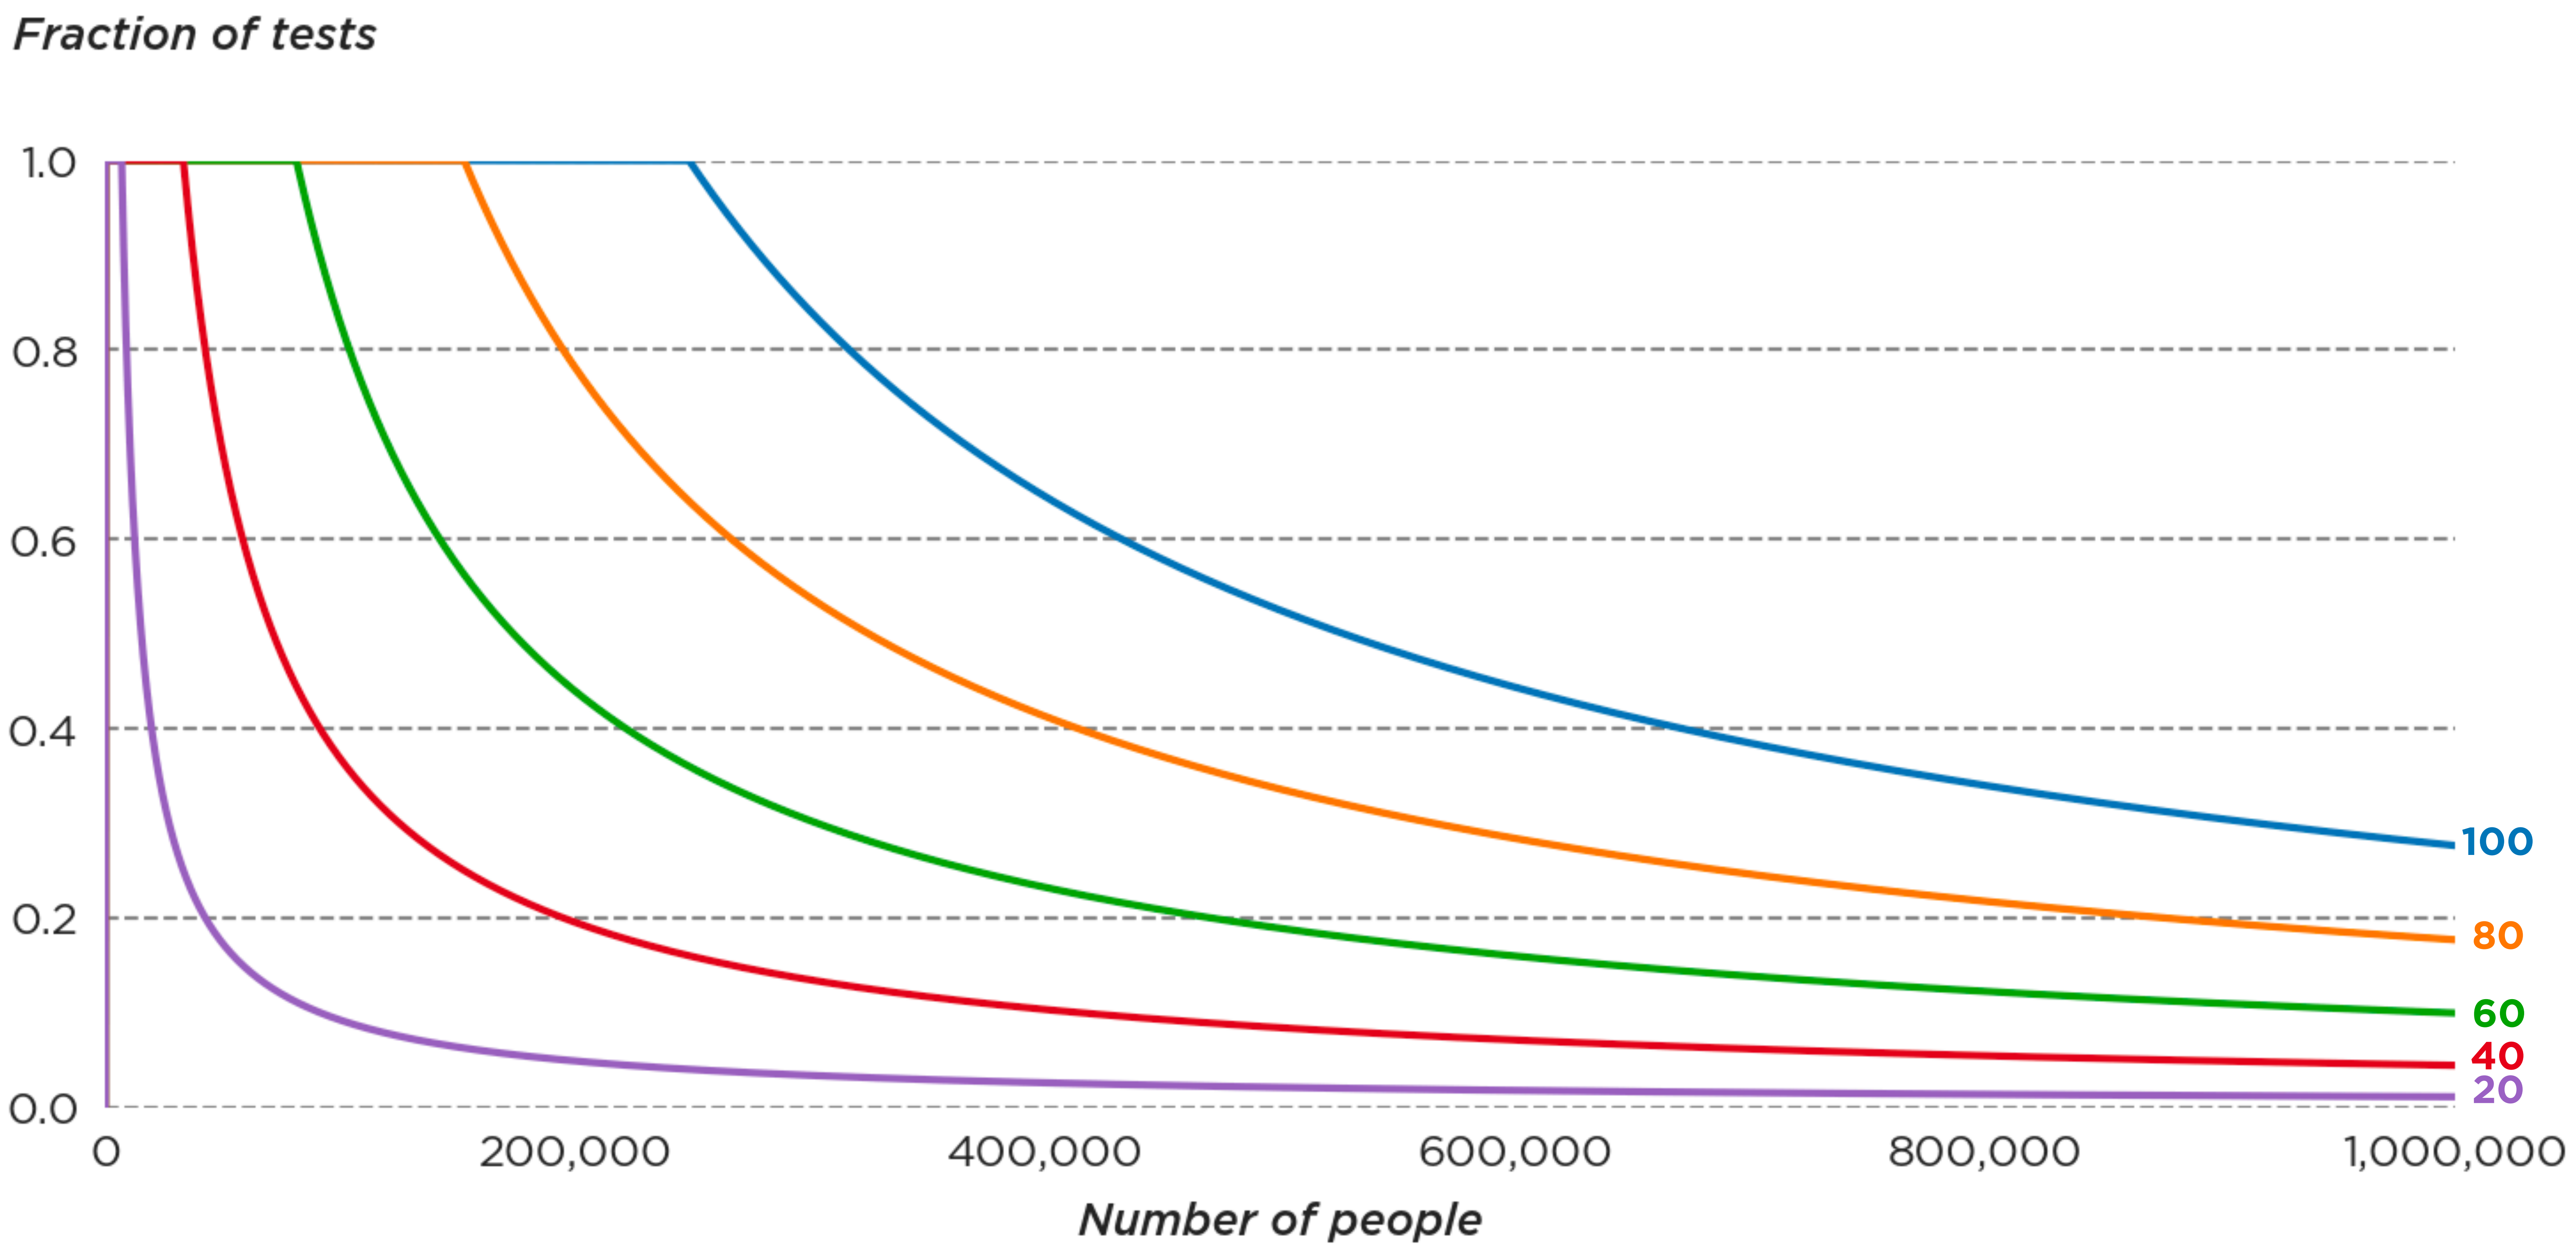

Why is $s^2 \log n$ tests cool?

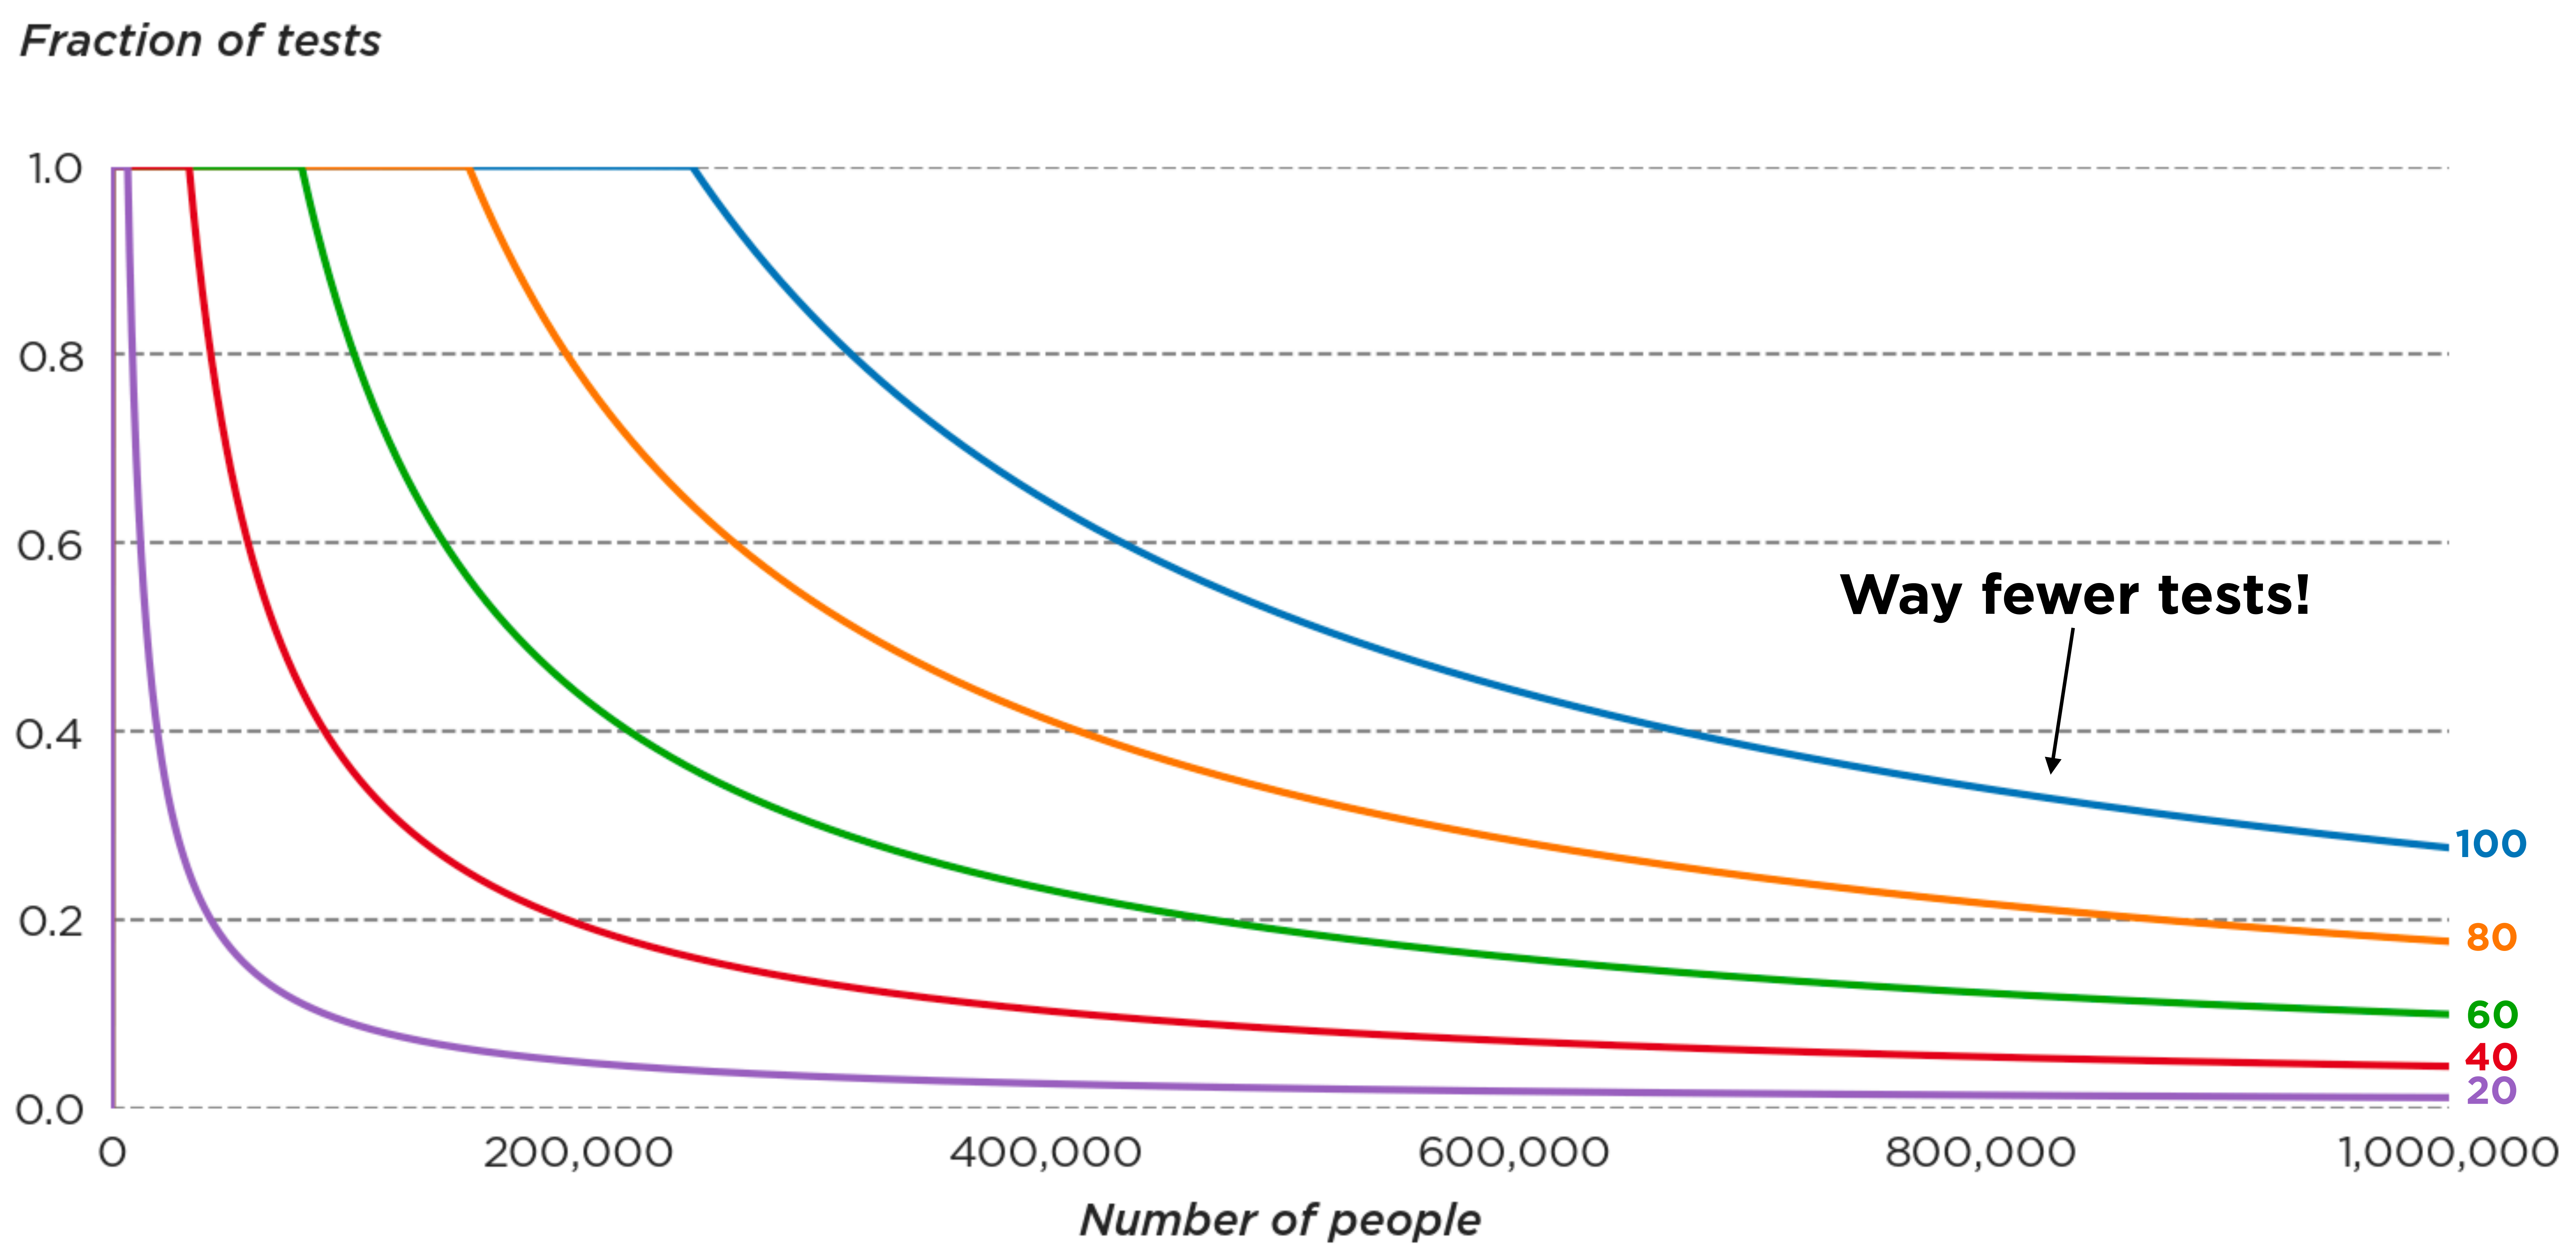

Dorfman's construction

Dorfman's construction

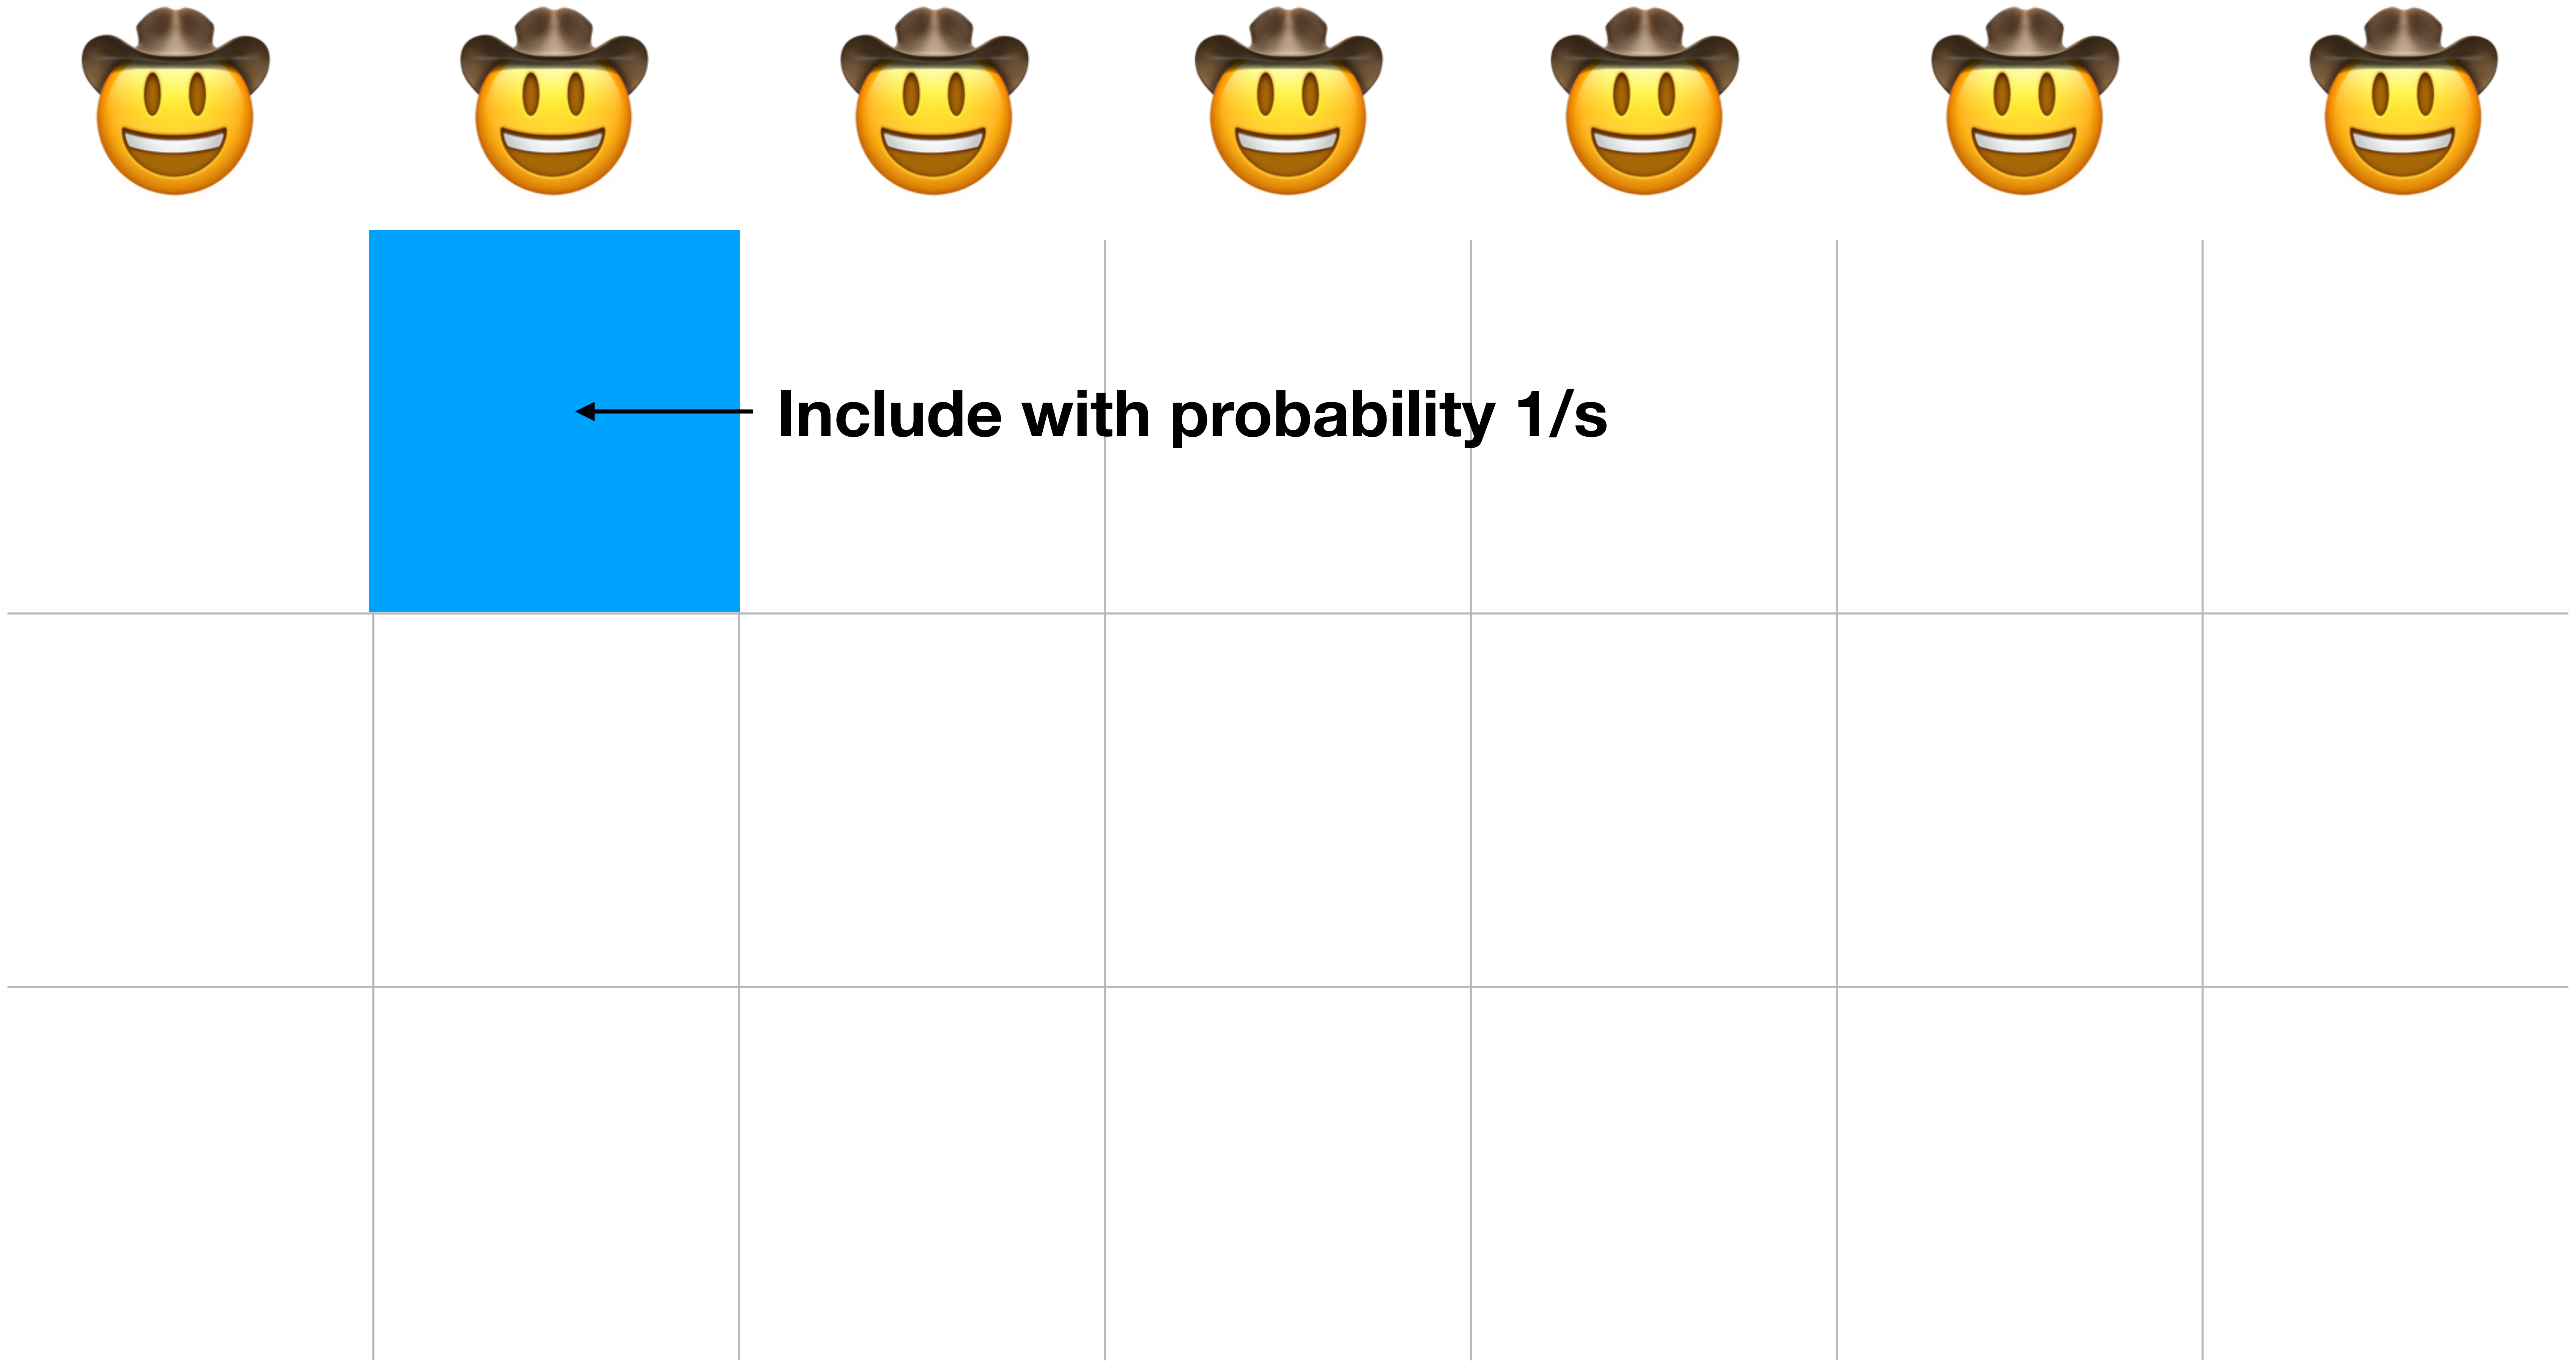

Dorfman’s construction

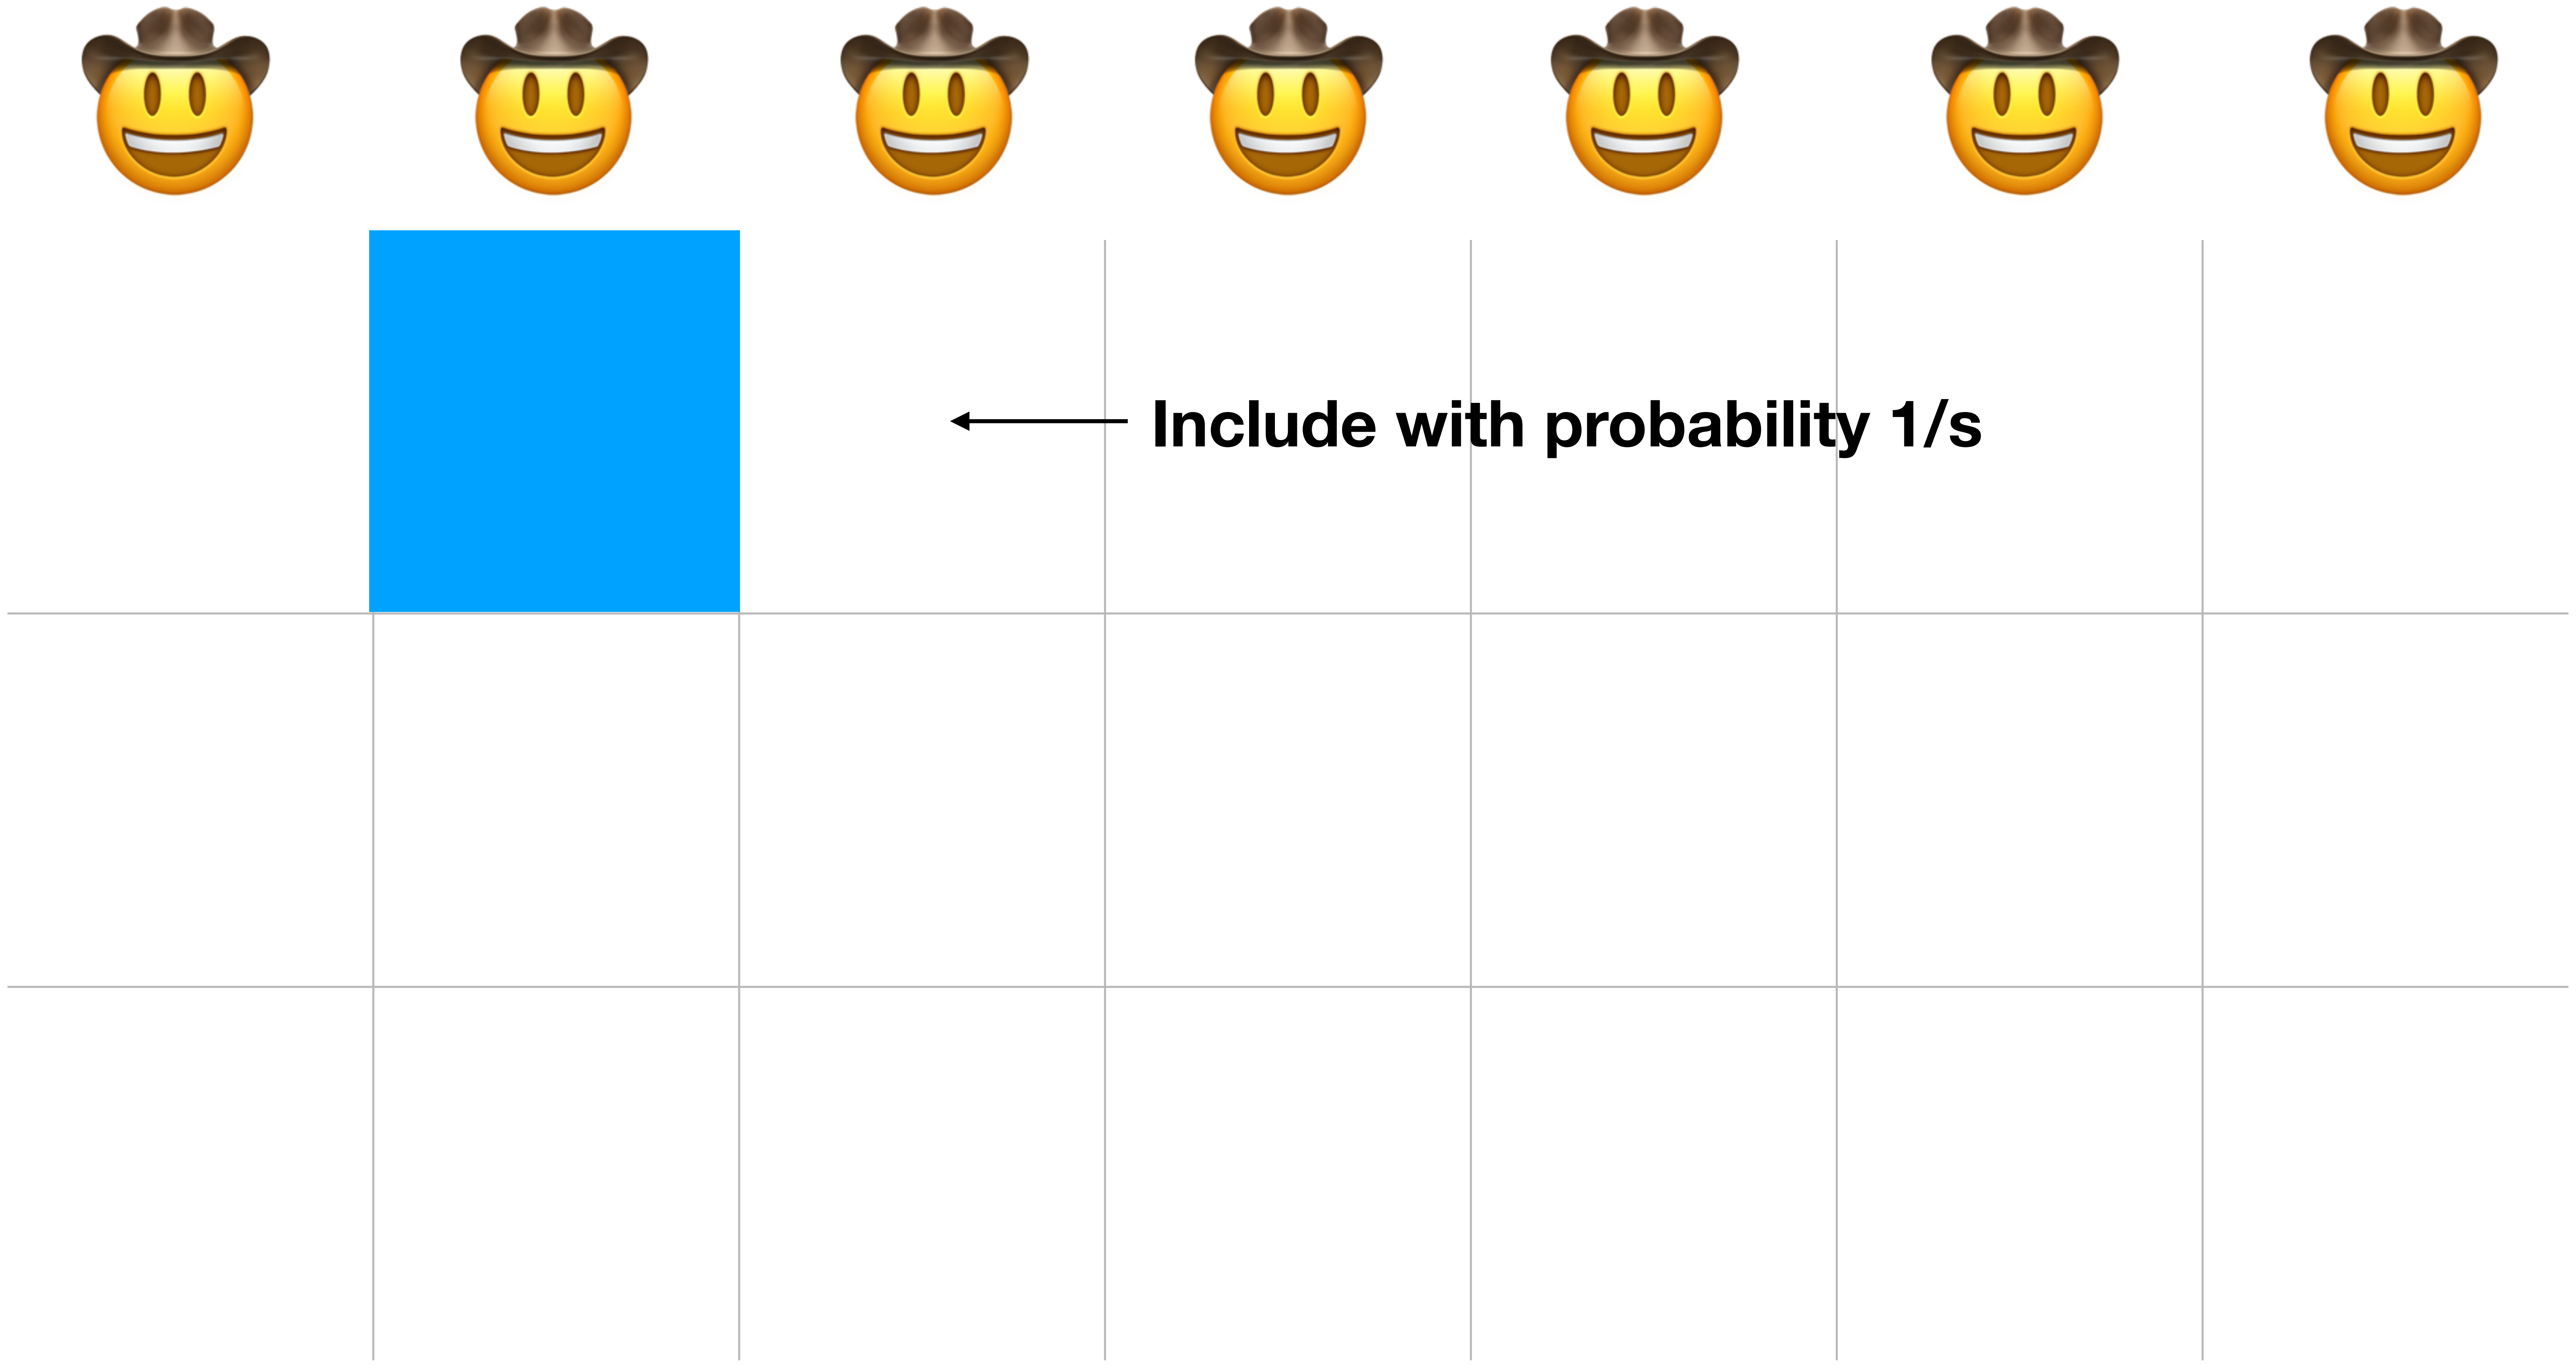

Dorfman's construction

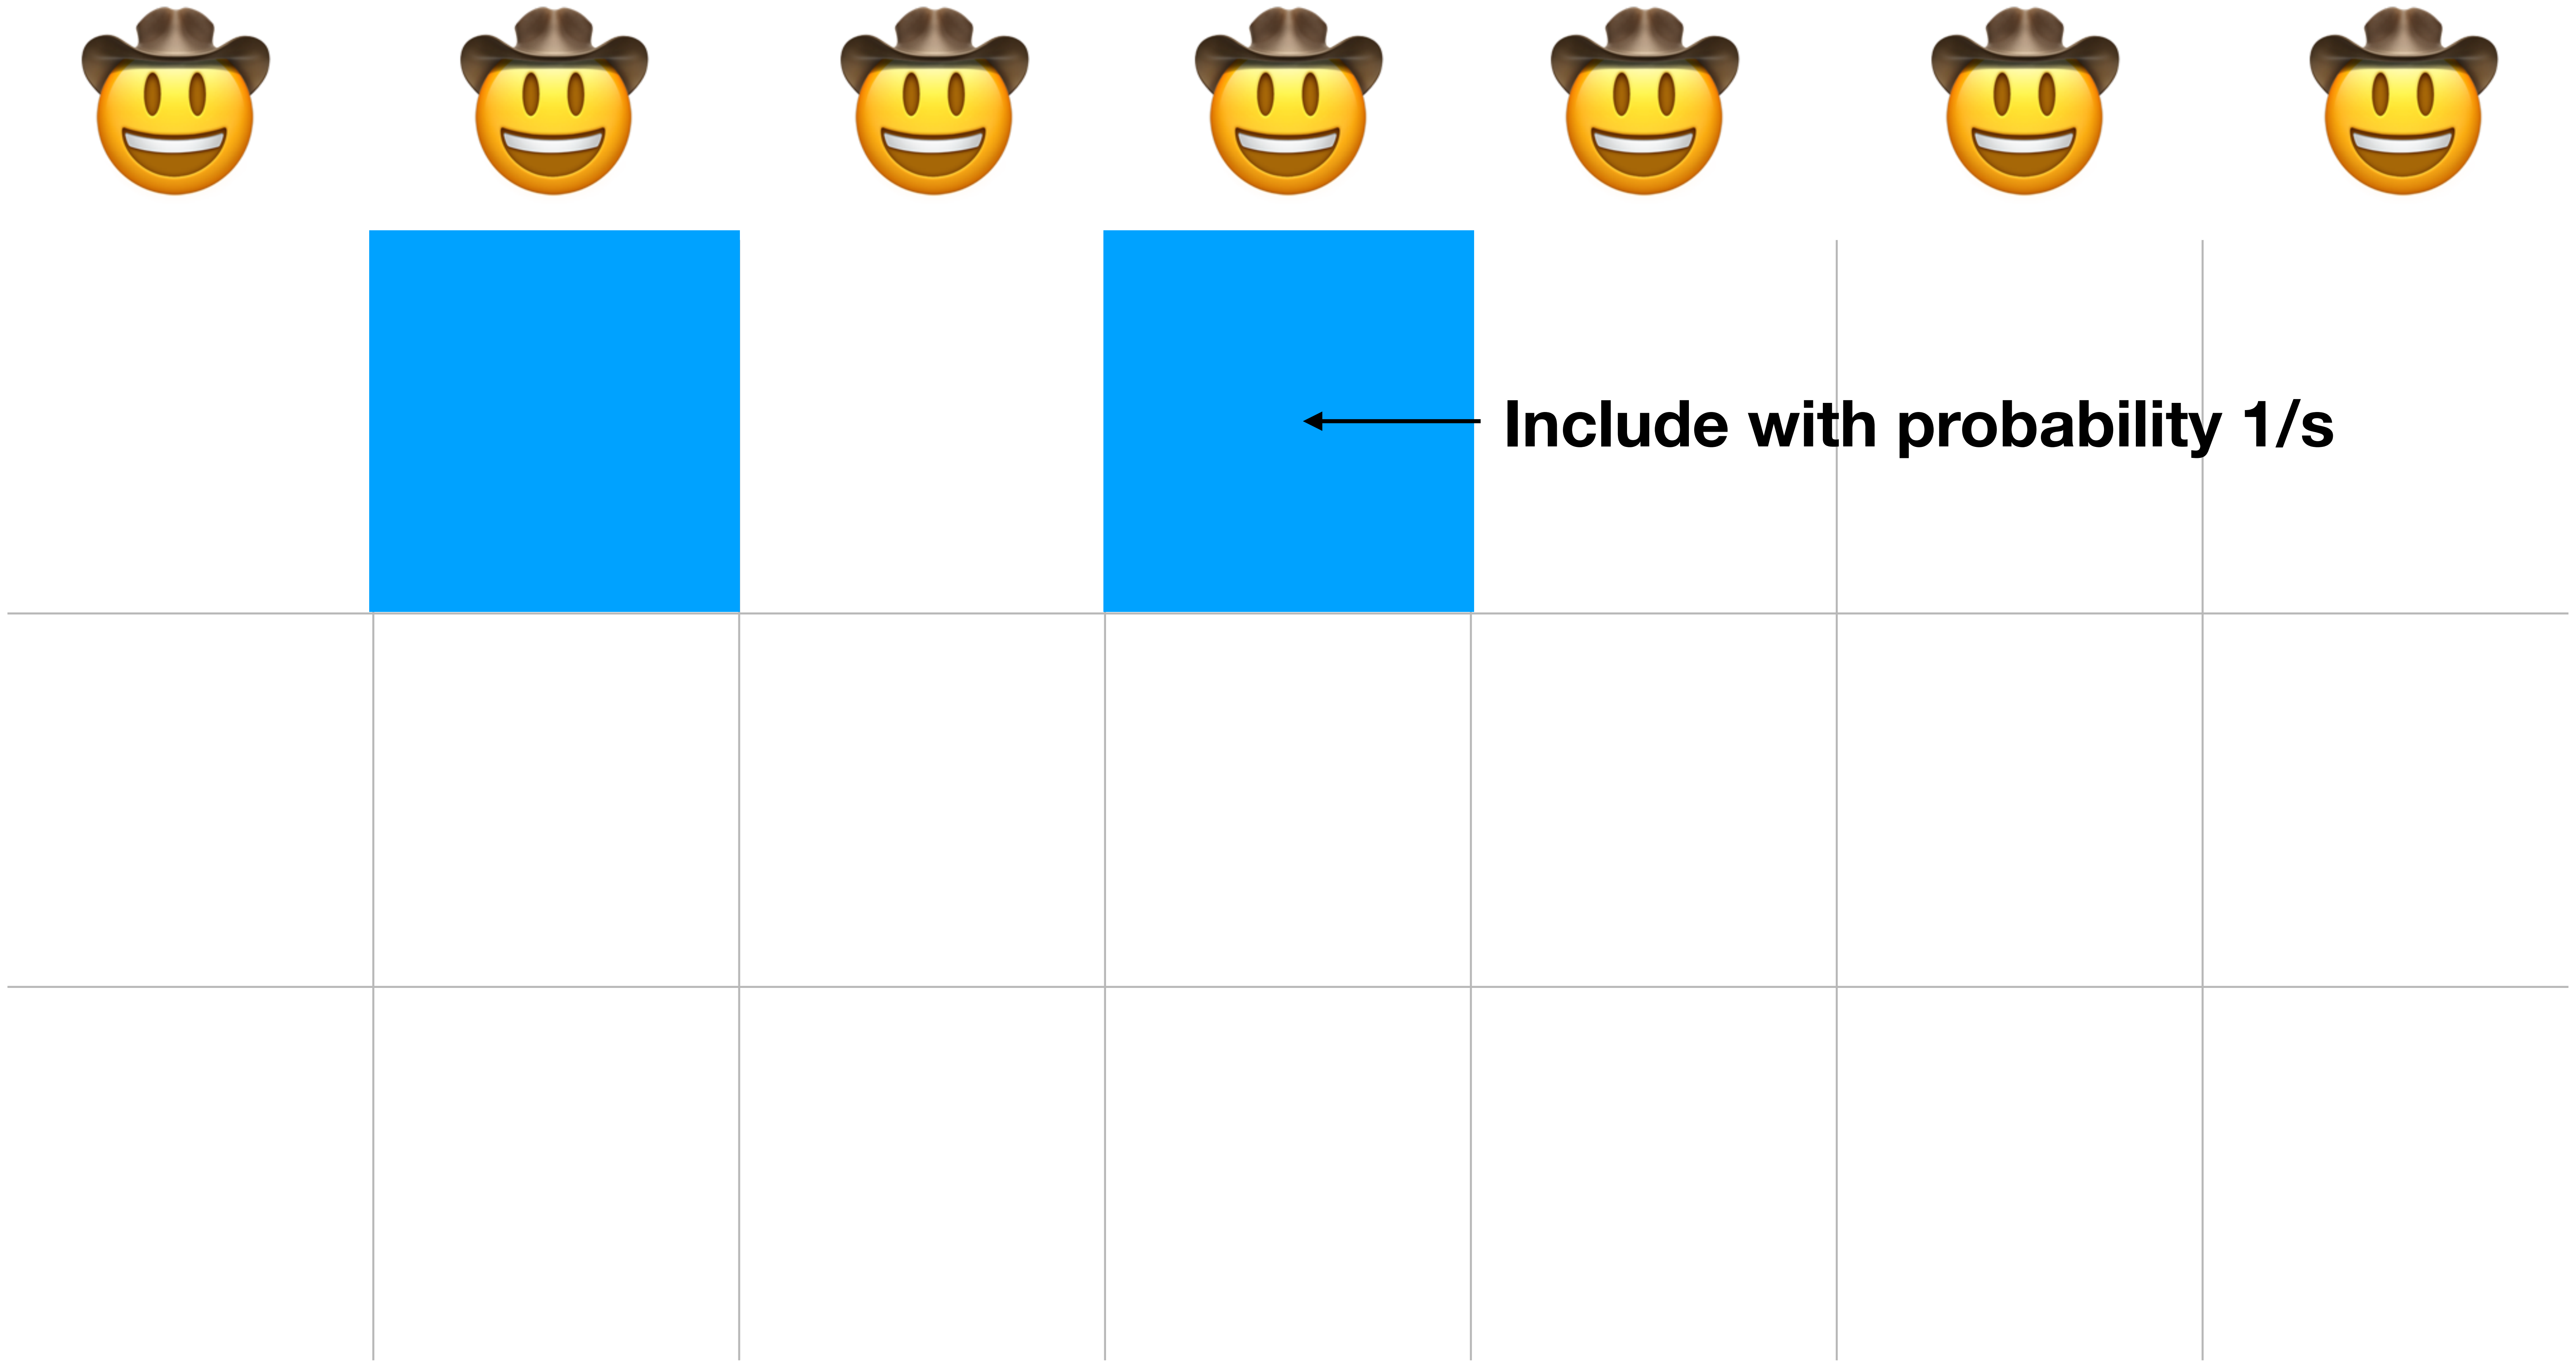

Dorfman's construction

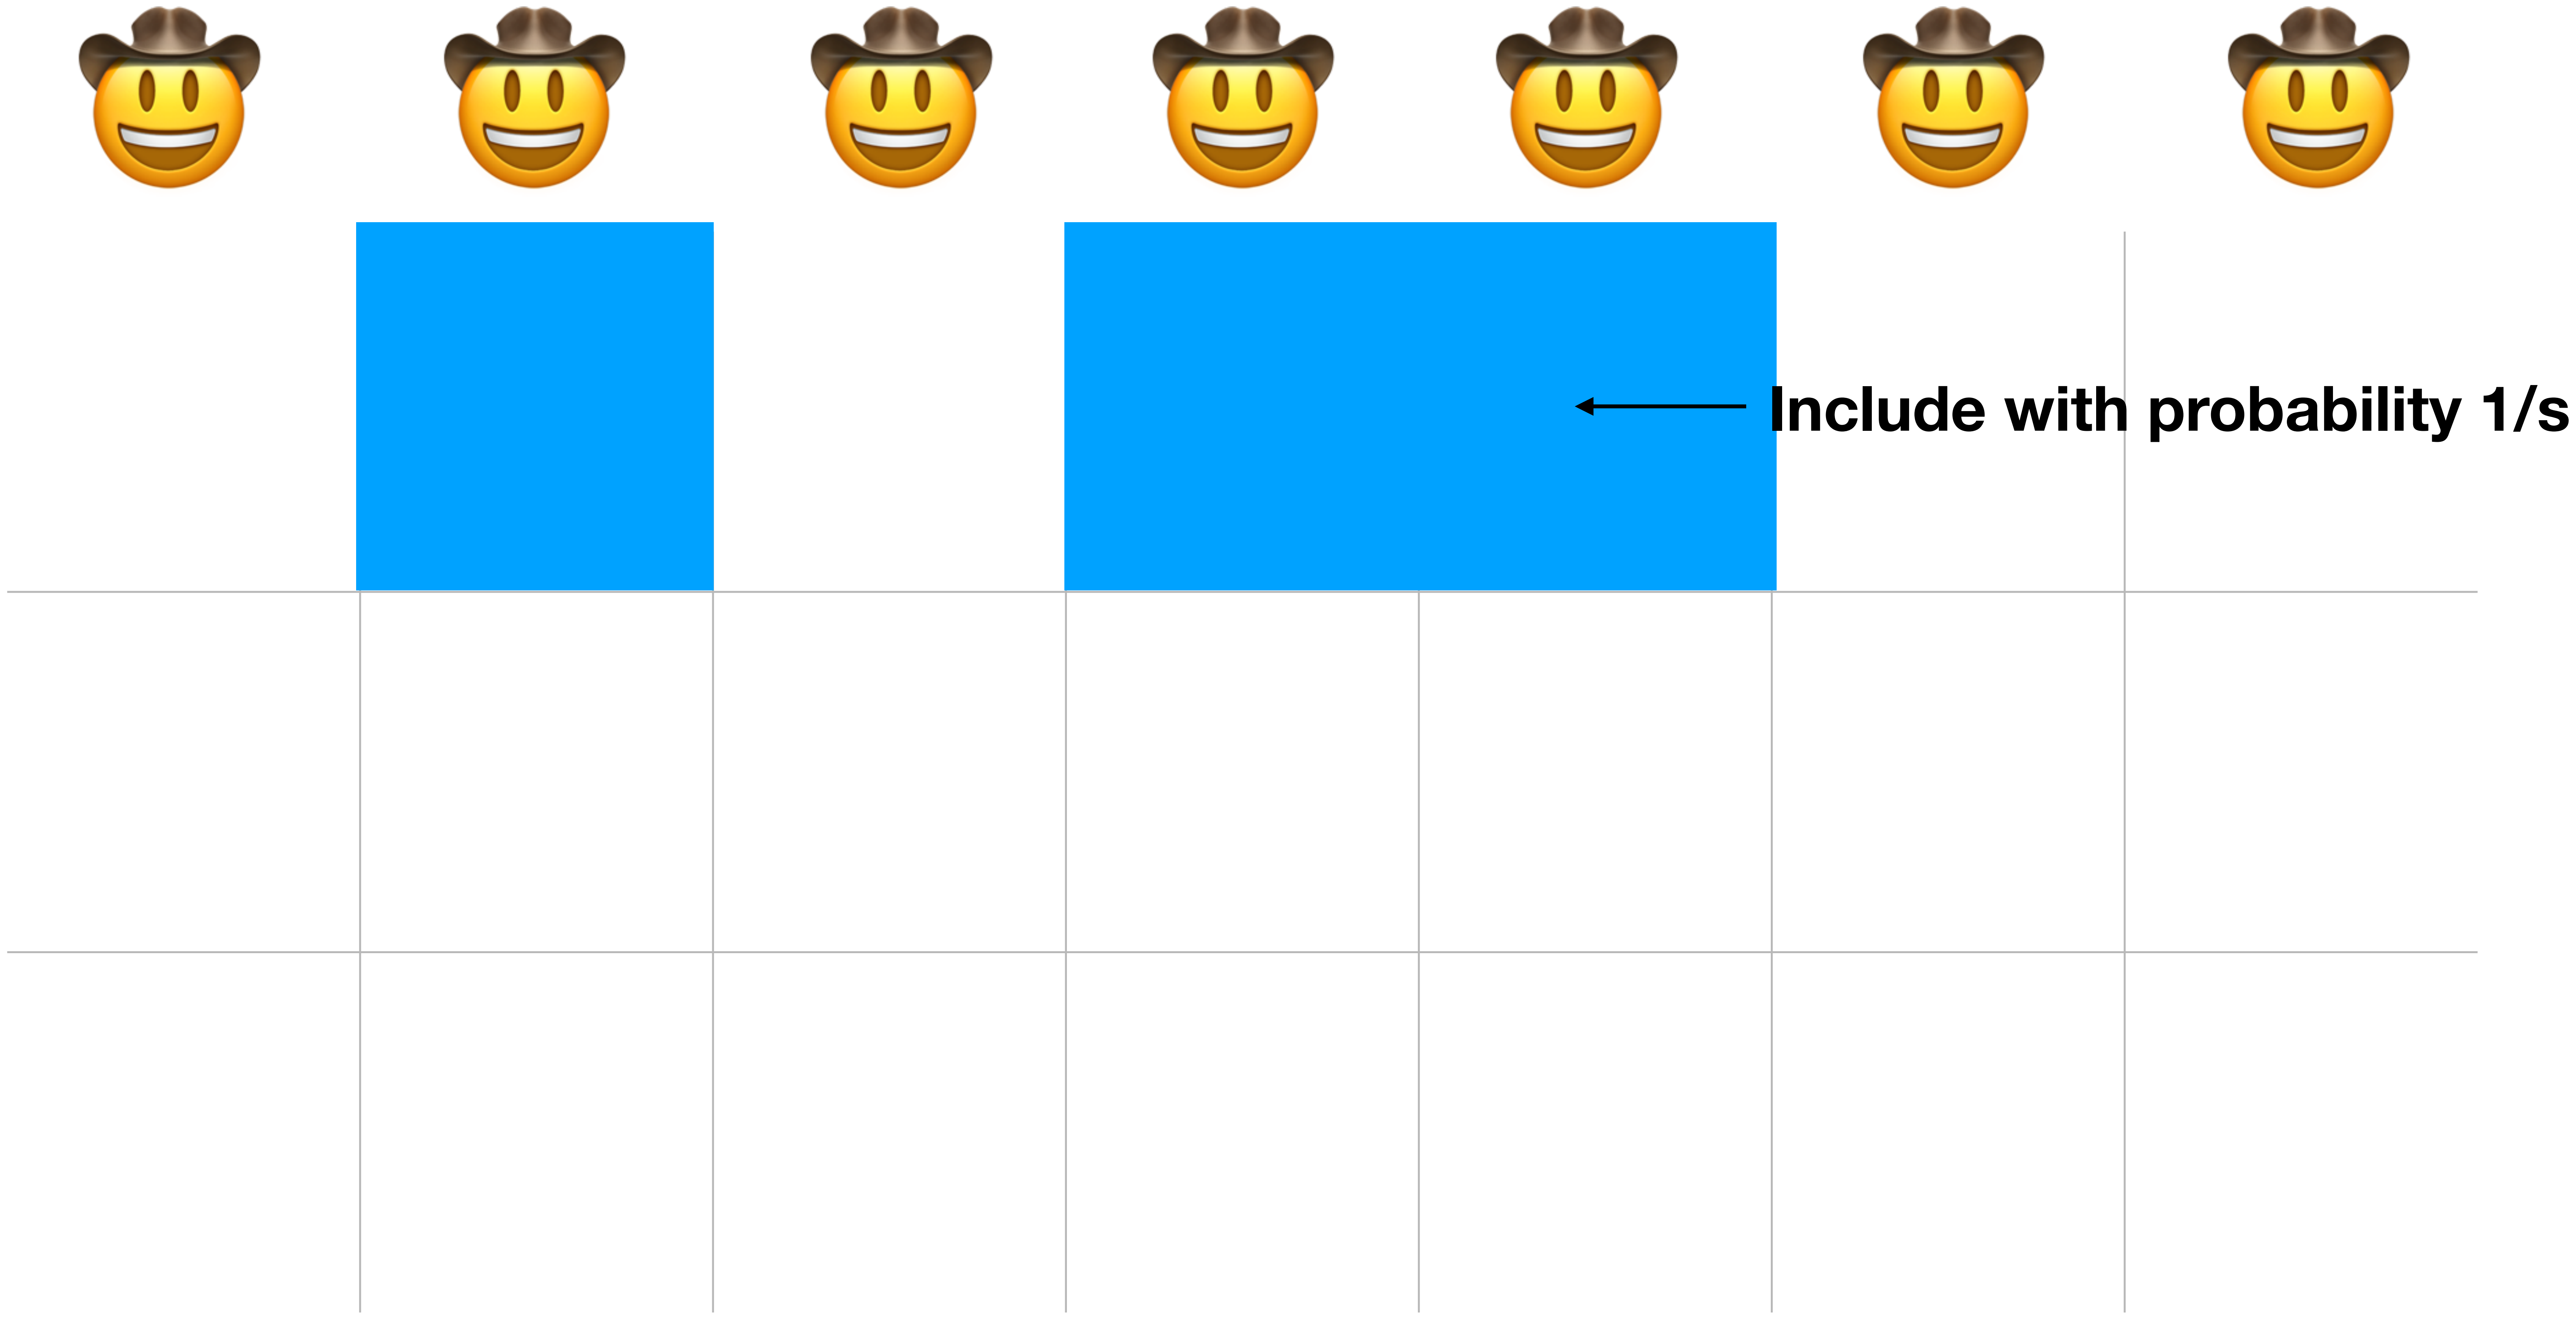

Dorfman’s construction

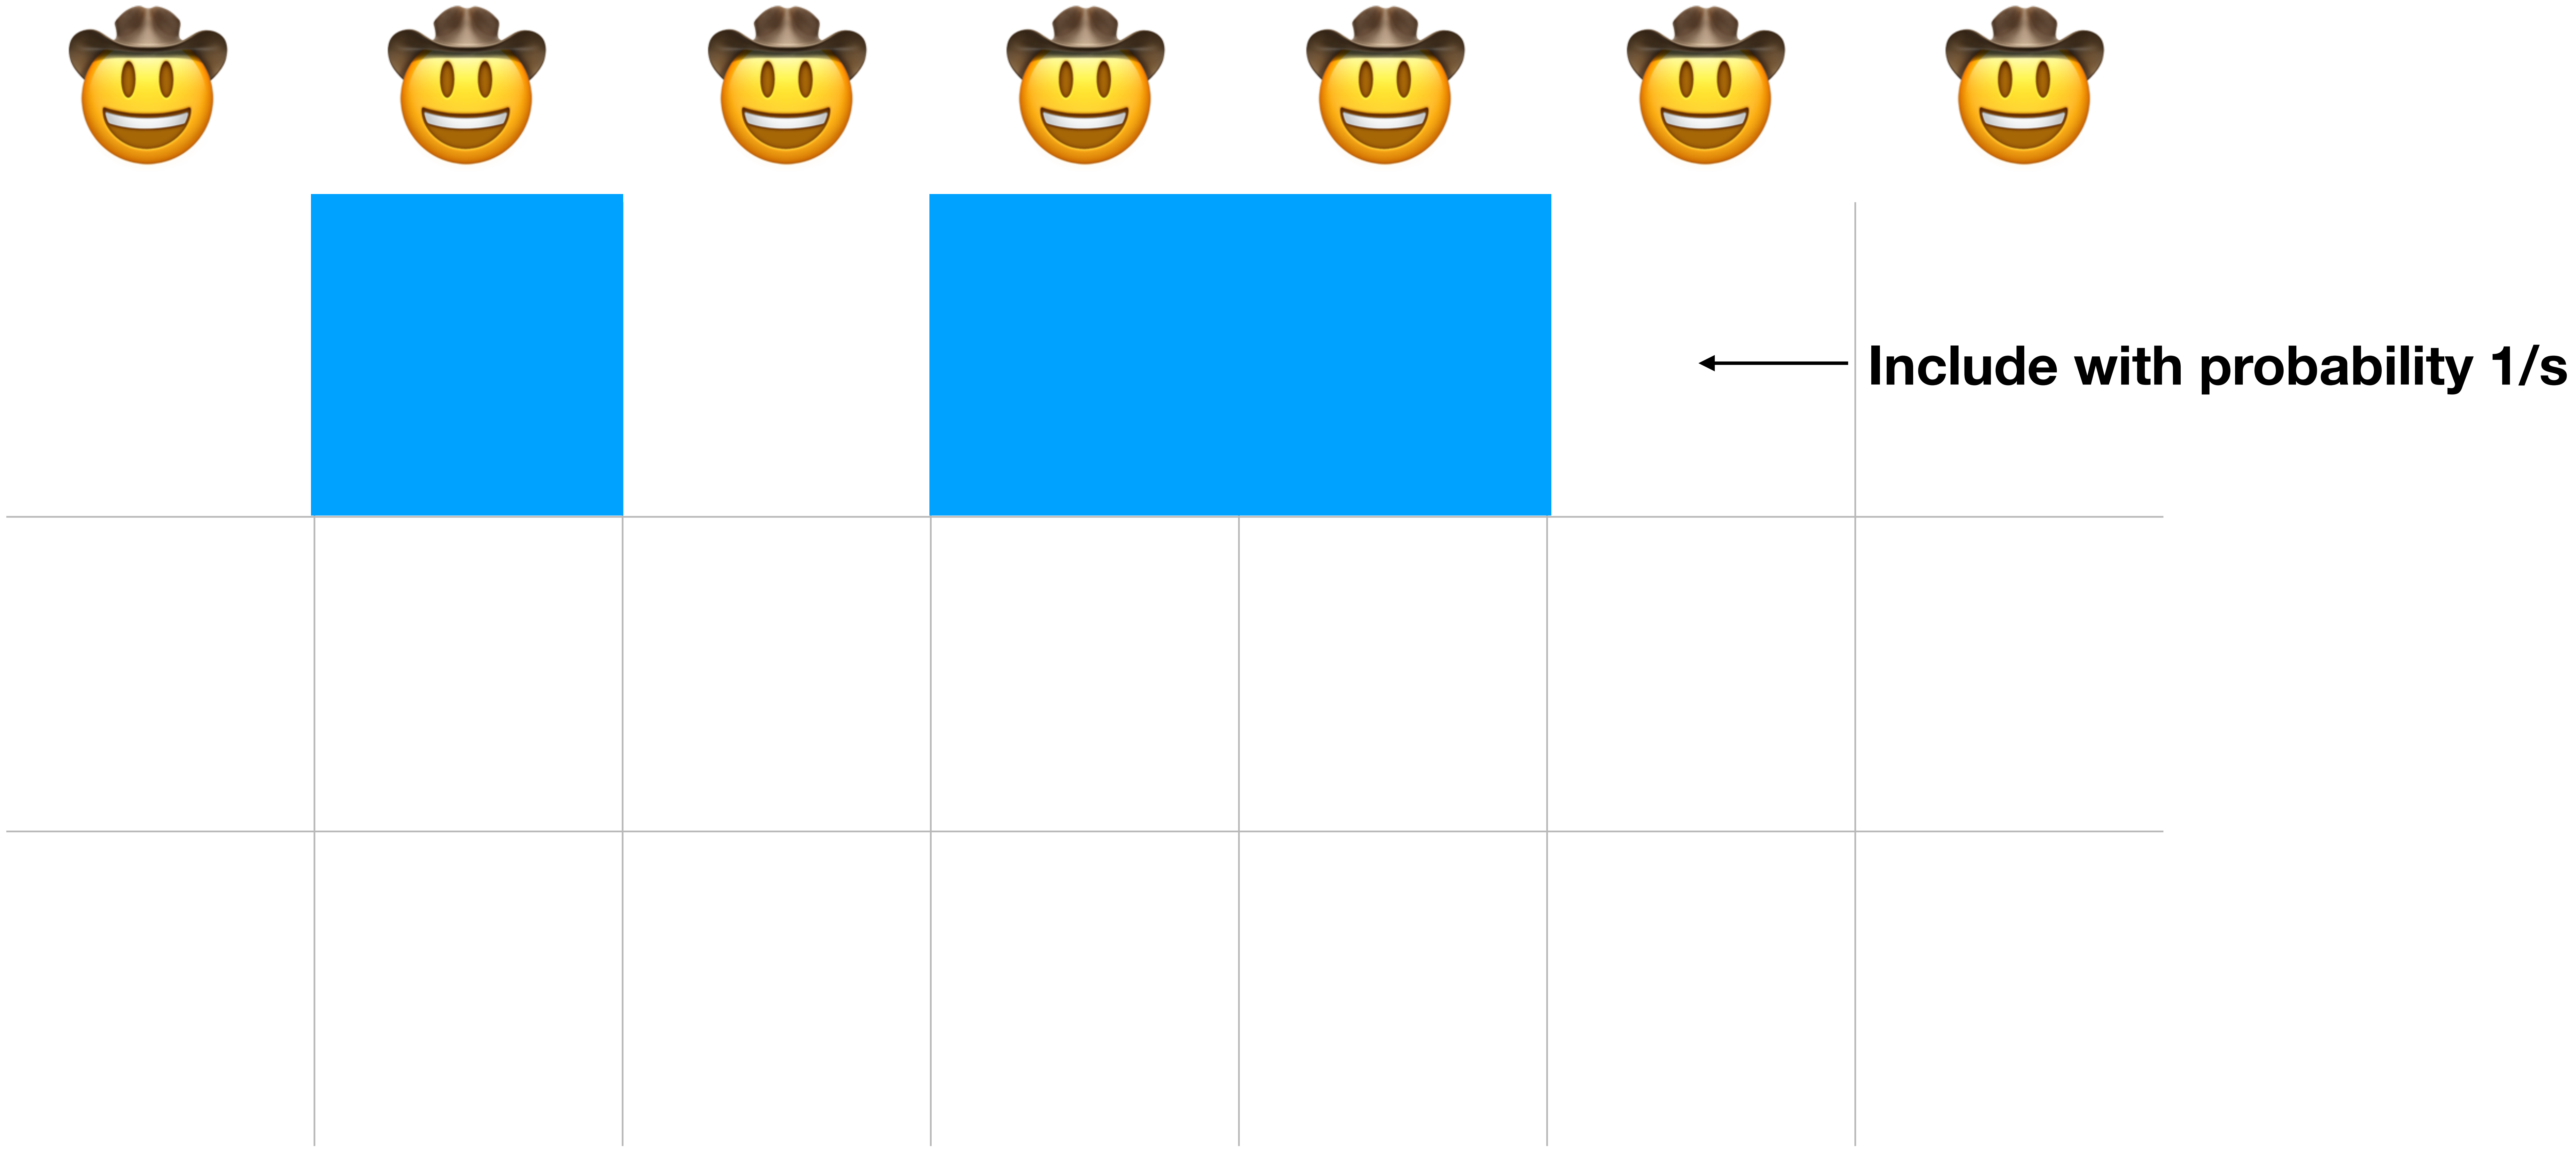

Dorfman’s construction

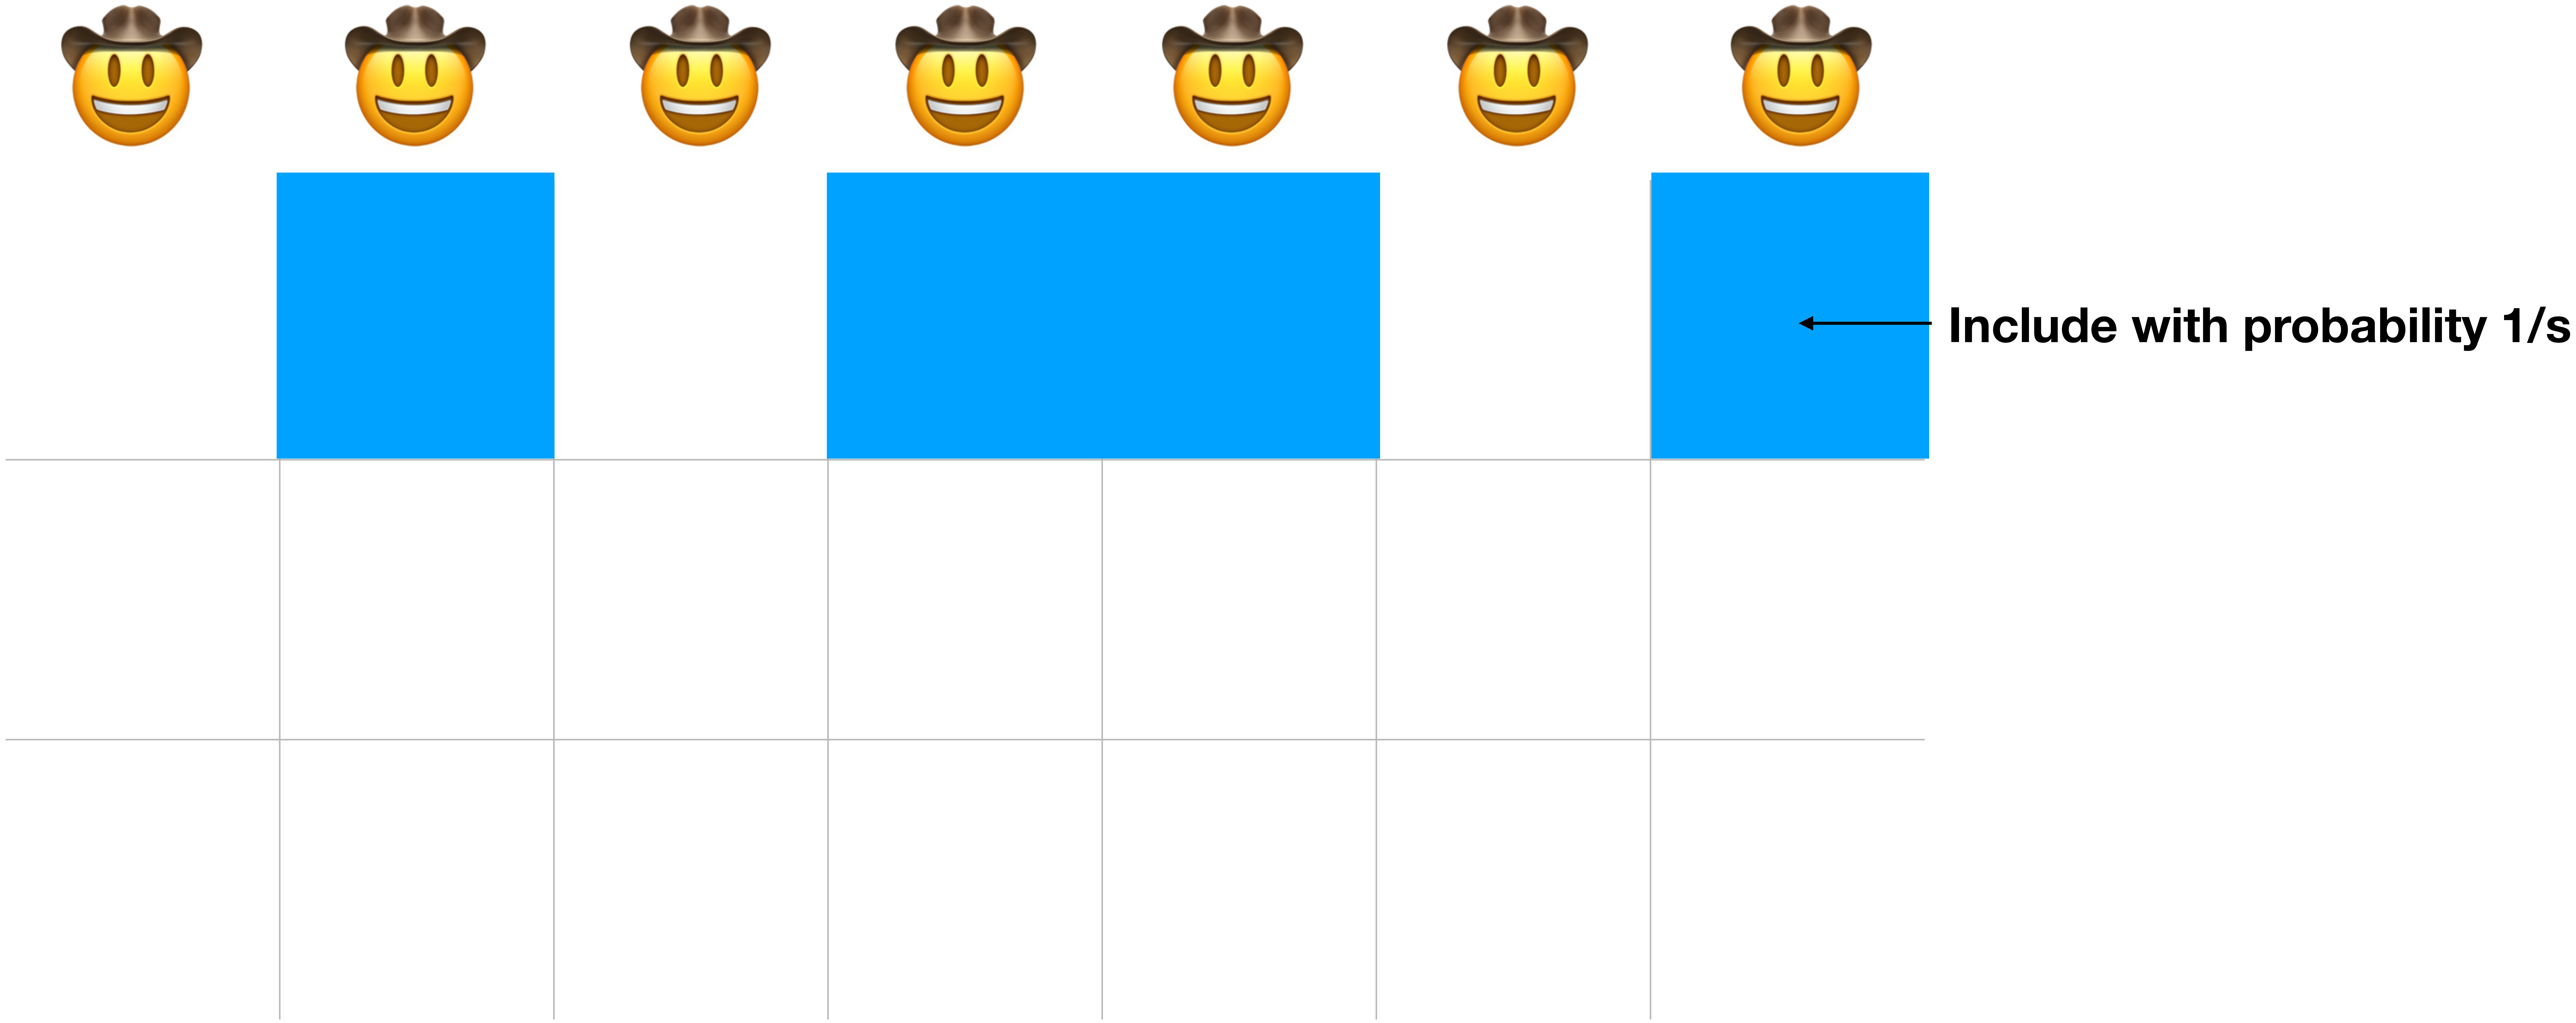

Dorfman's construction

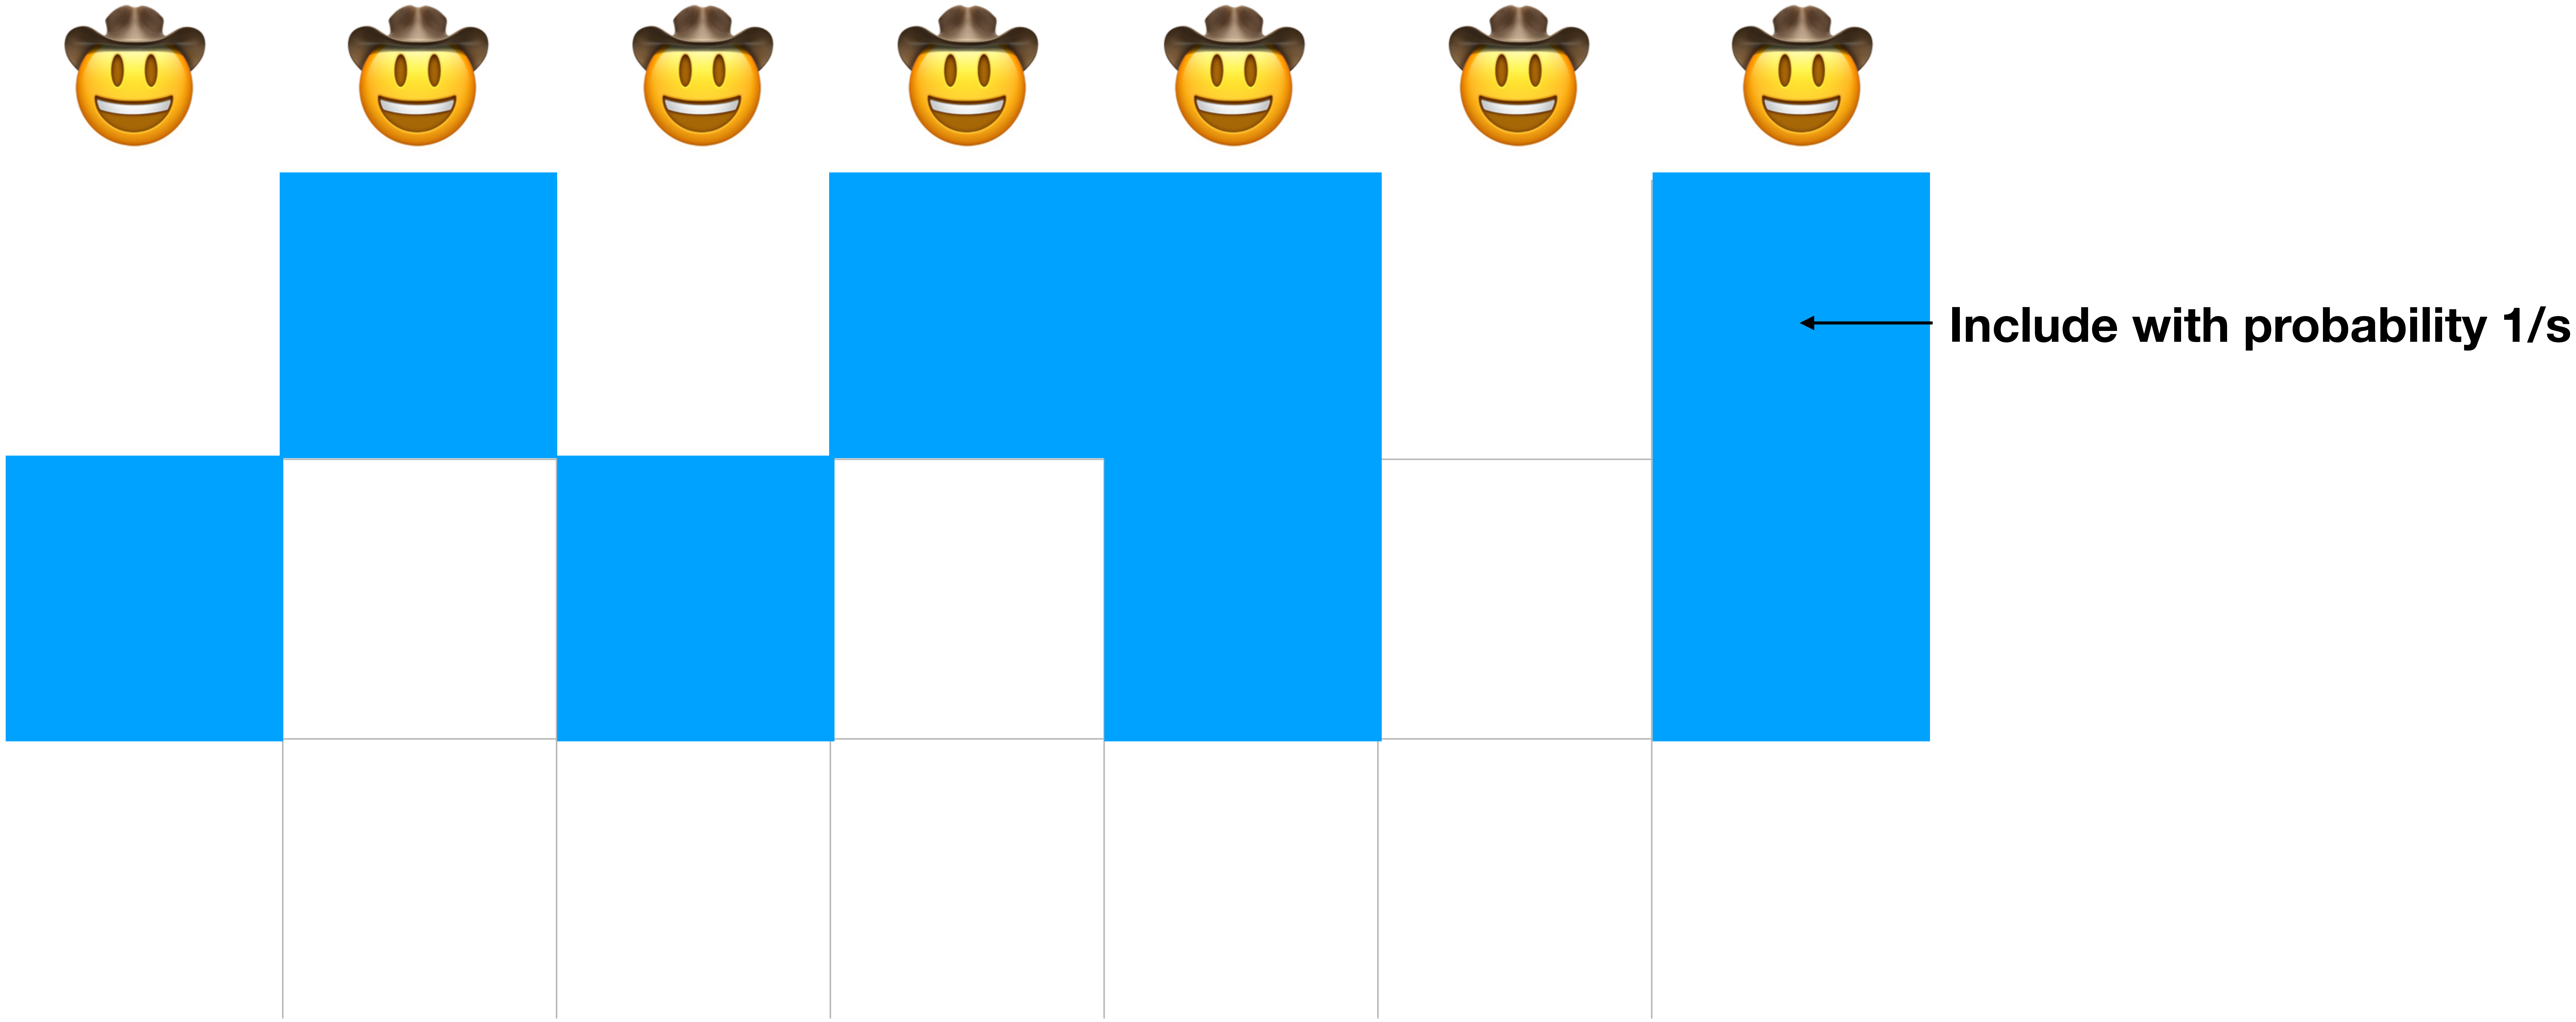

Dorfman's construction

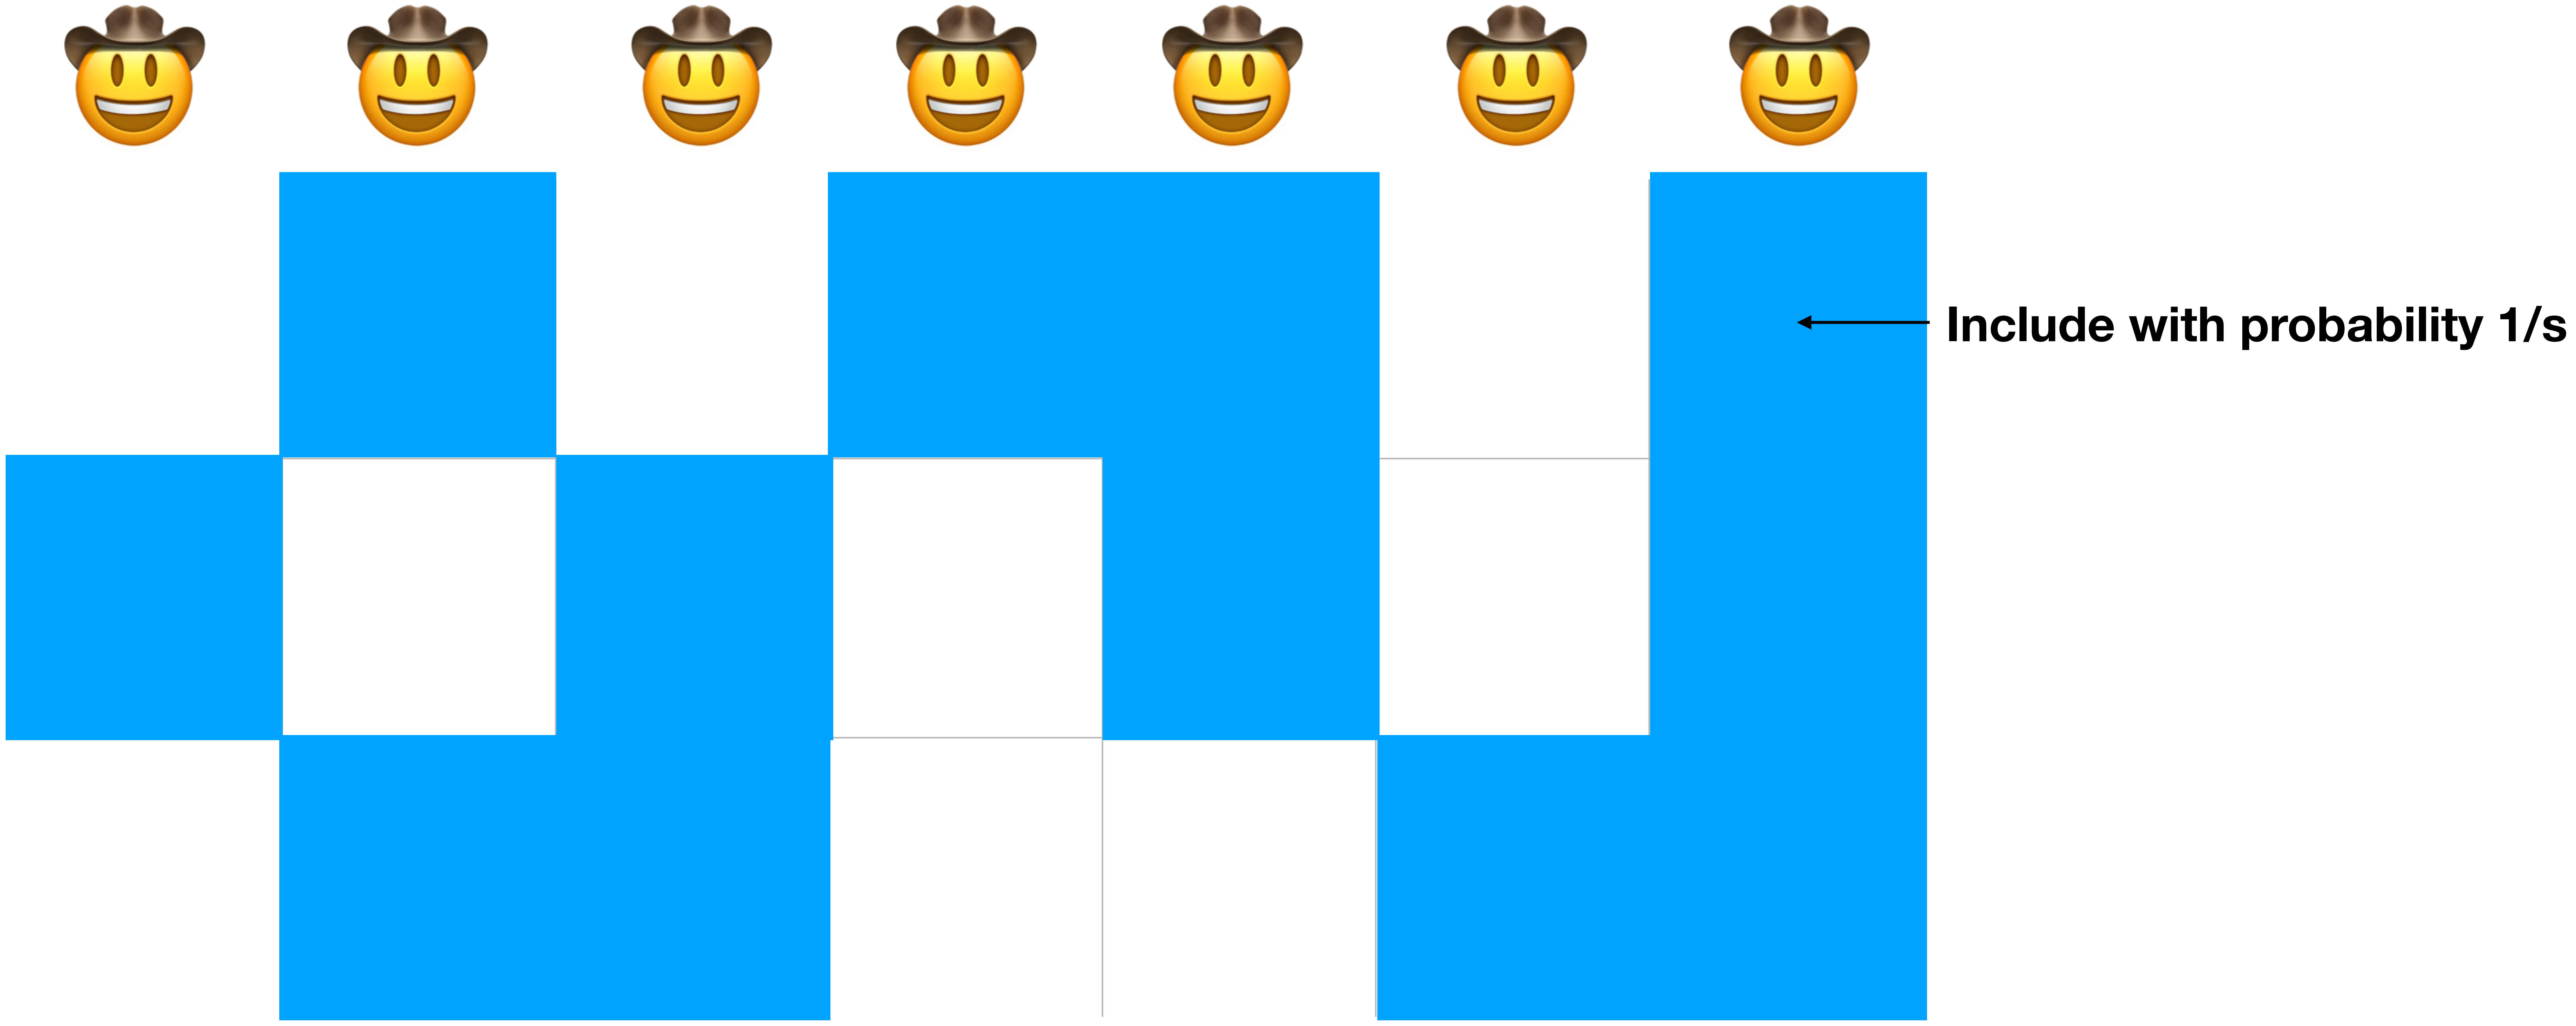

First idea: finding healthy people

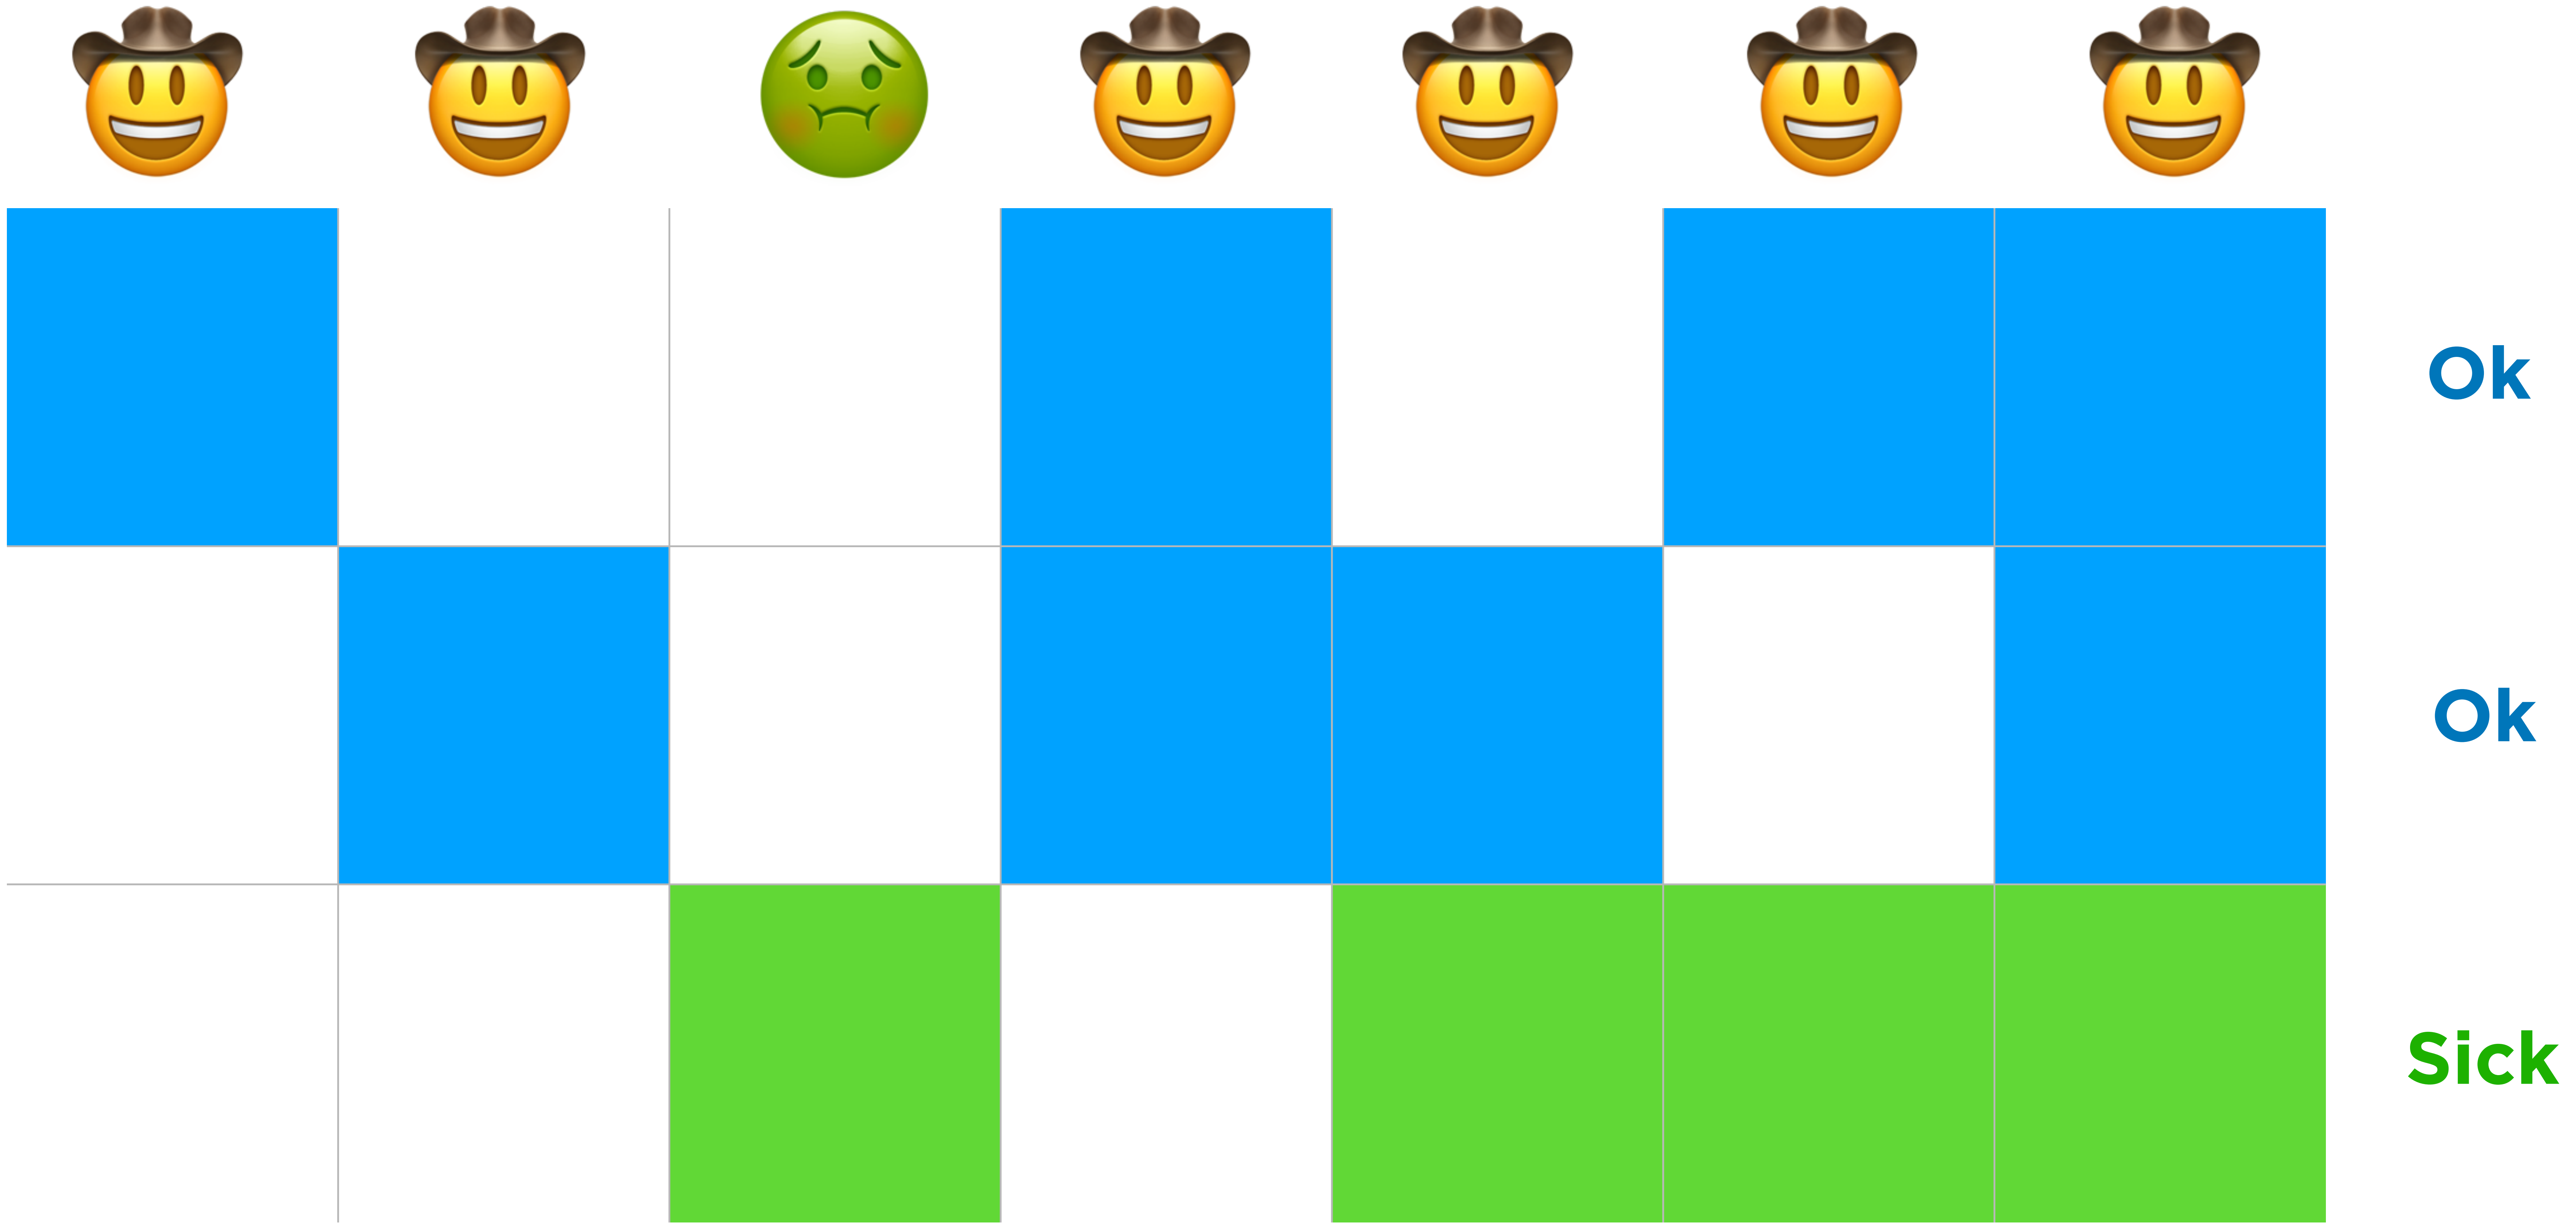

First idea: finding healthy people

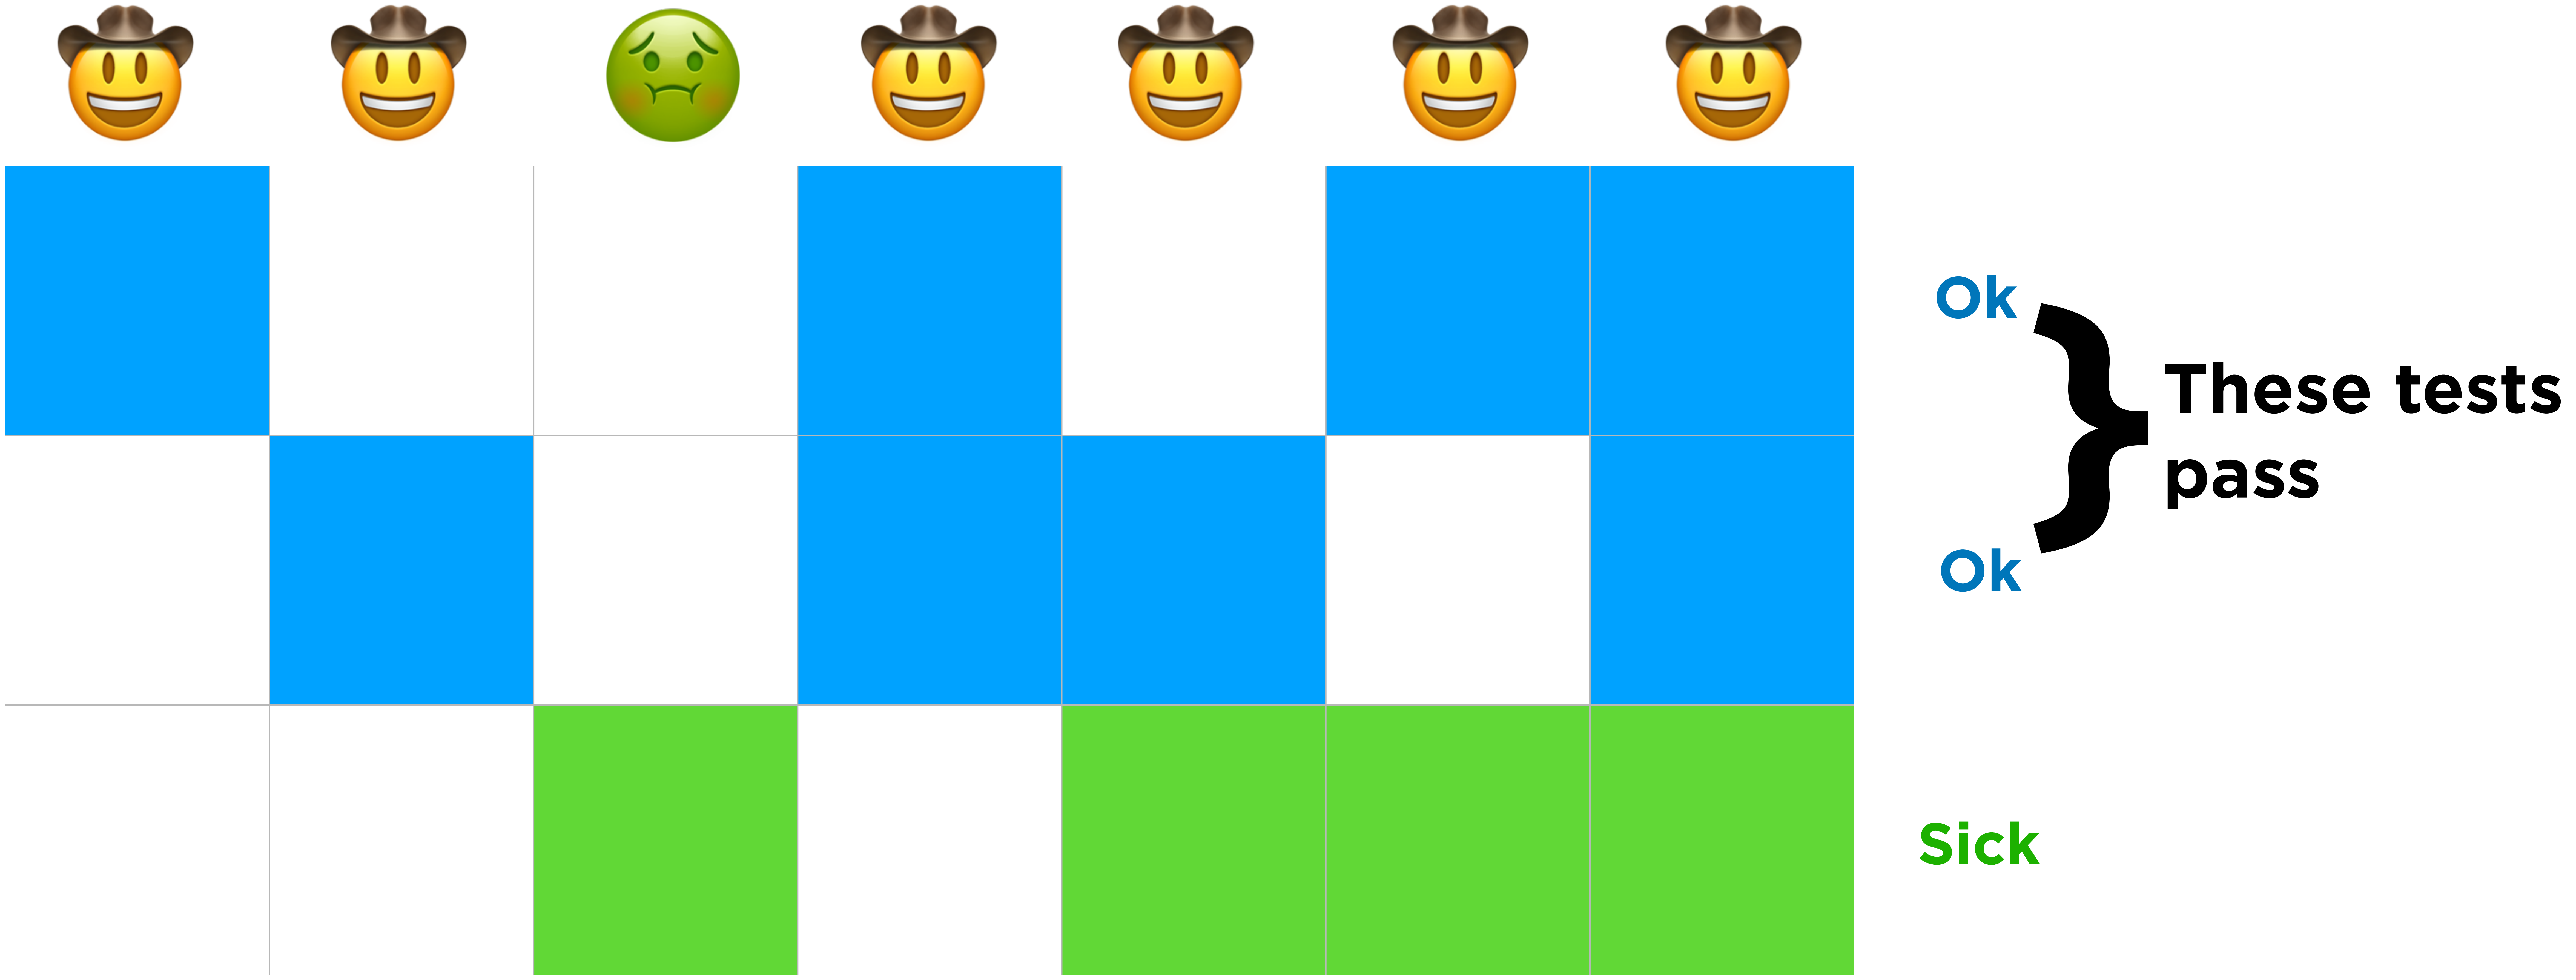

First idea: finding healthy people

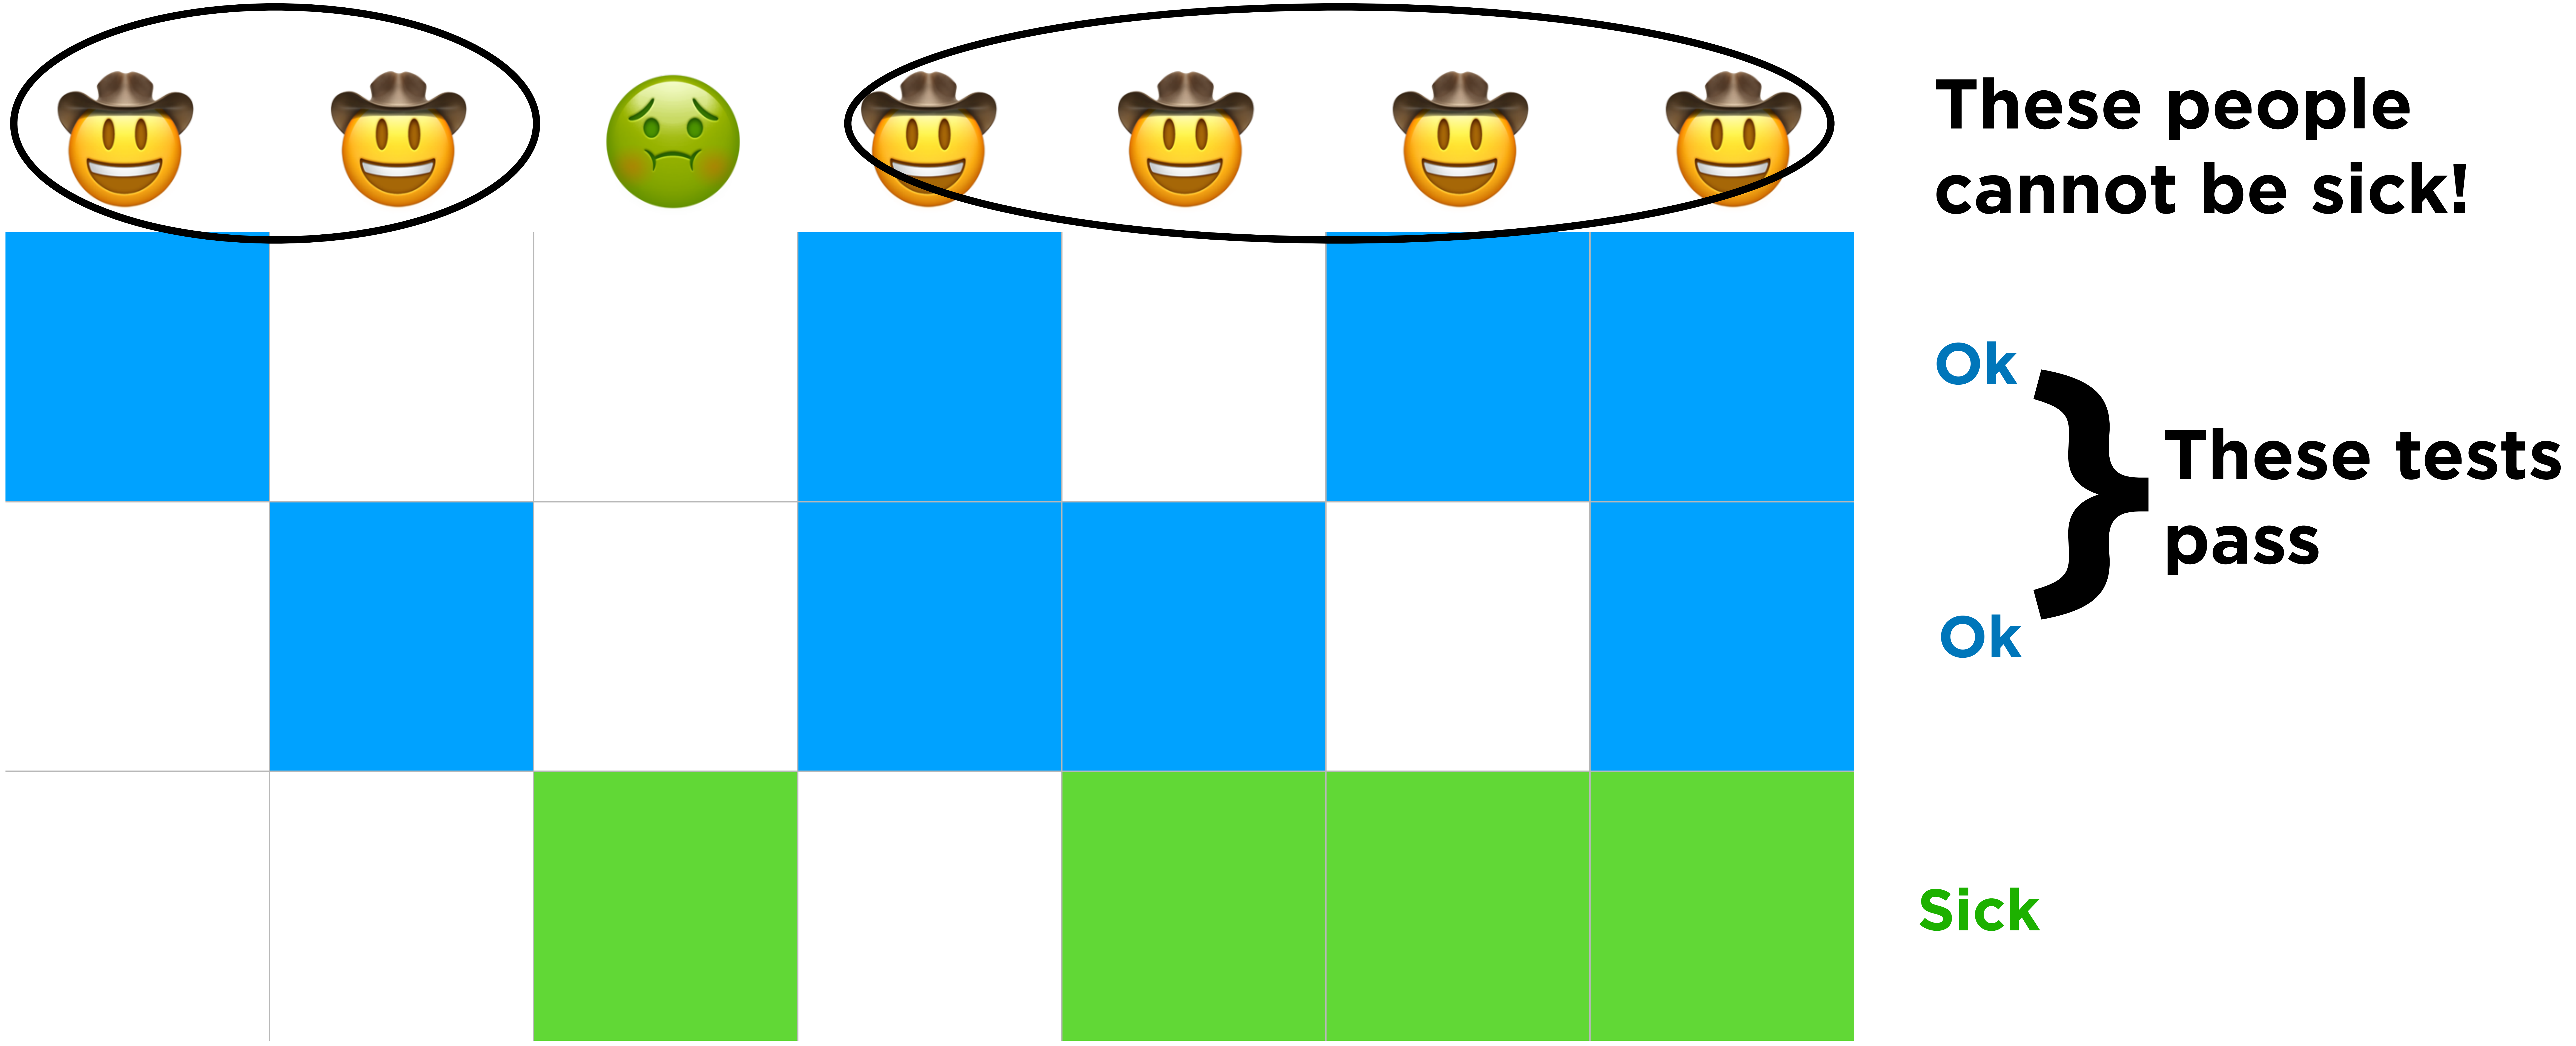

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

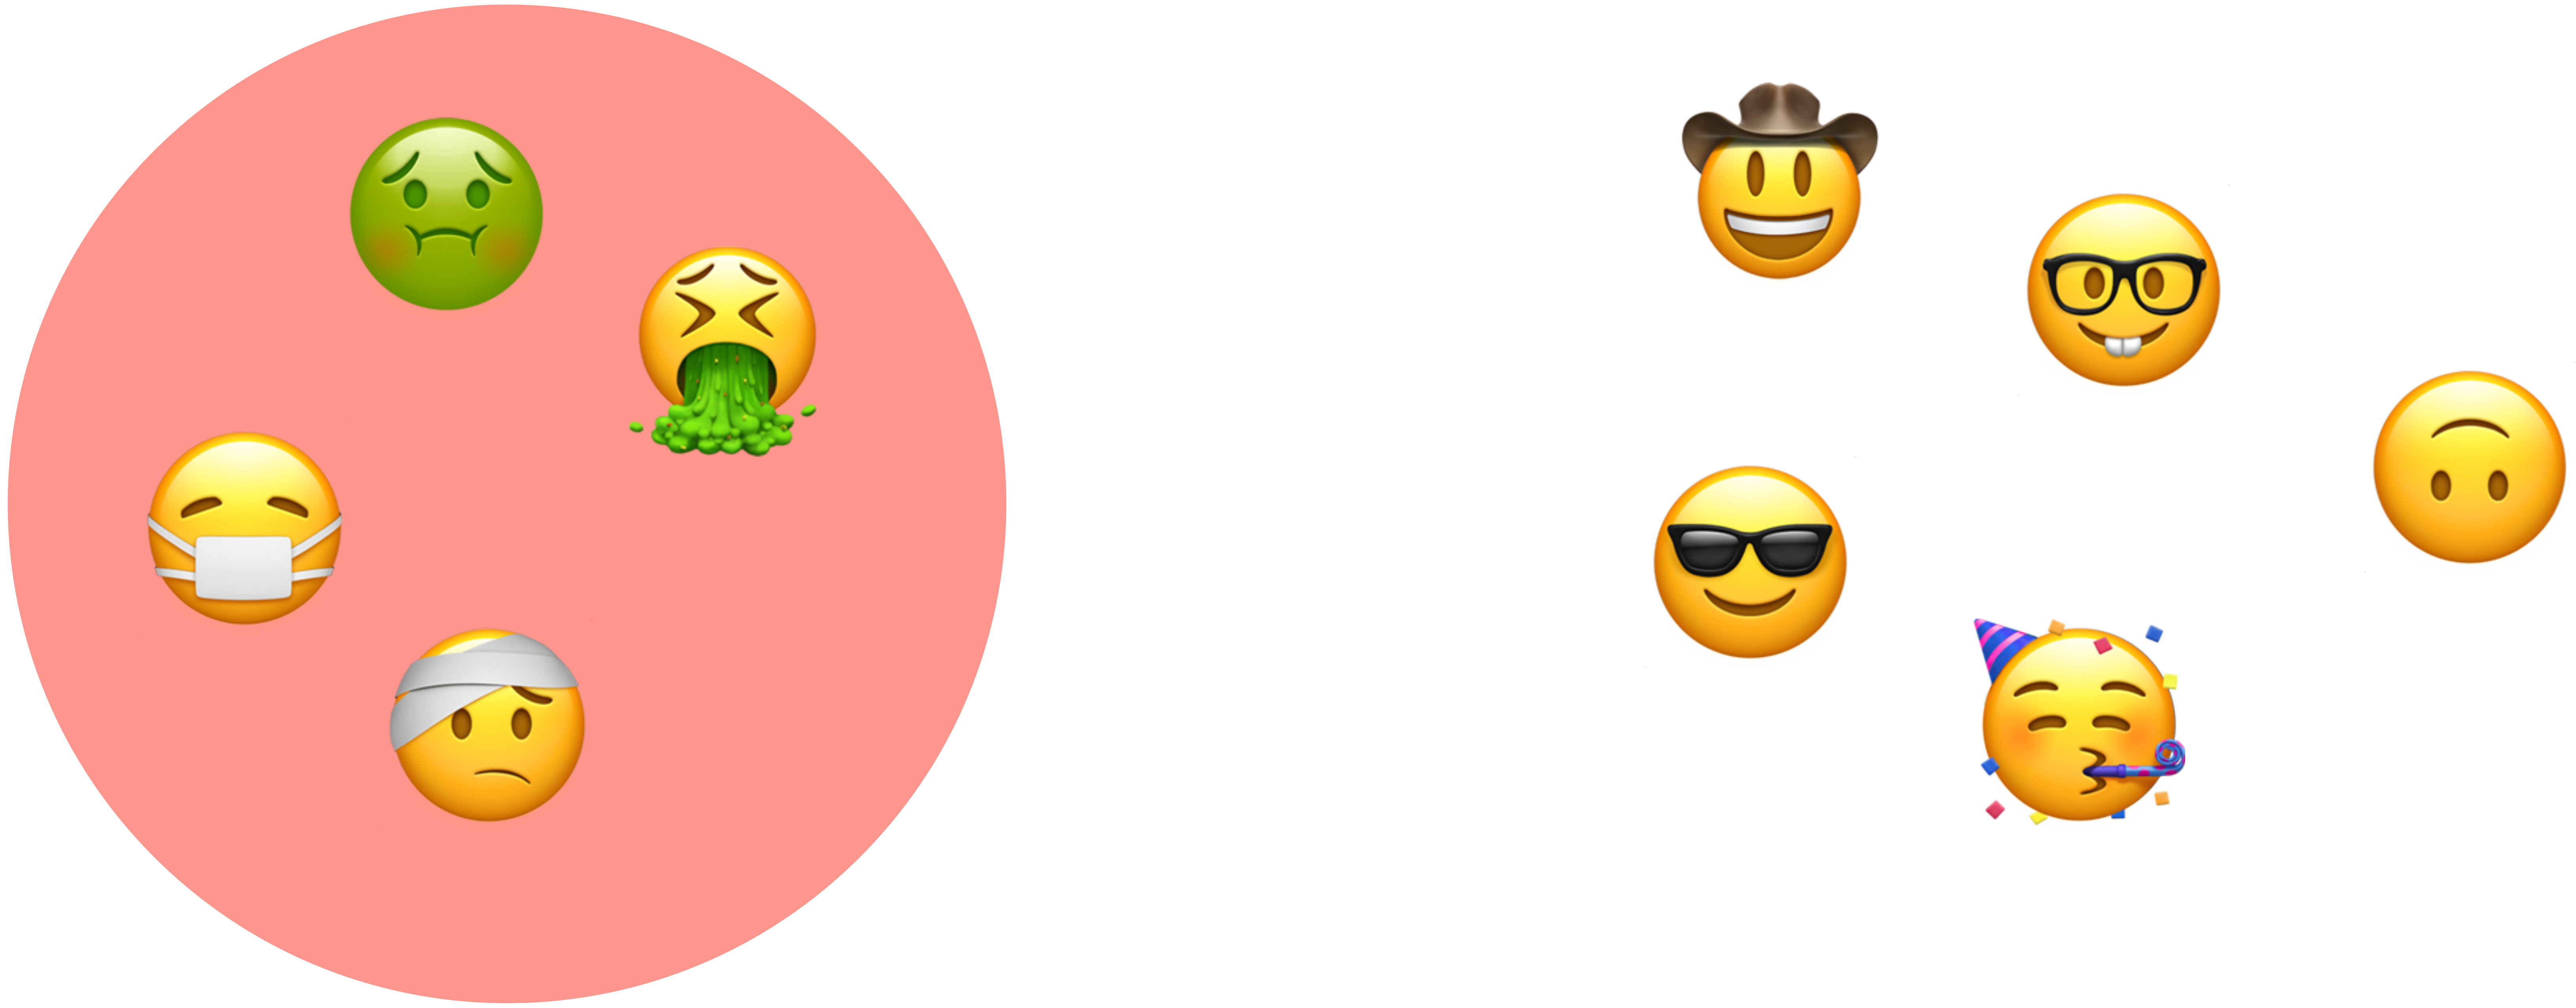

Should not be in the test

Should be in test

First idea: finding healthy people

For each set of sick … to … e
each other person … ea

Math time!

Should not be in the test

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

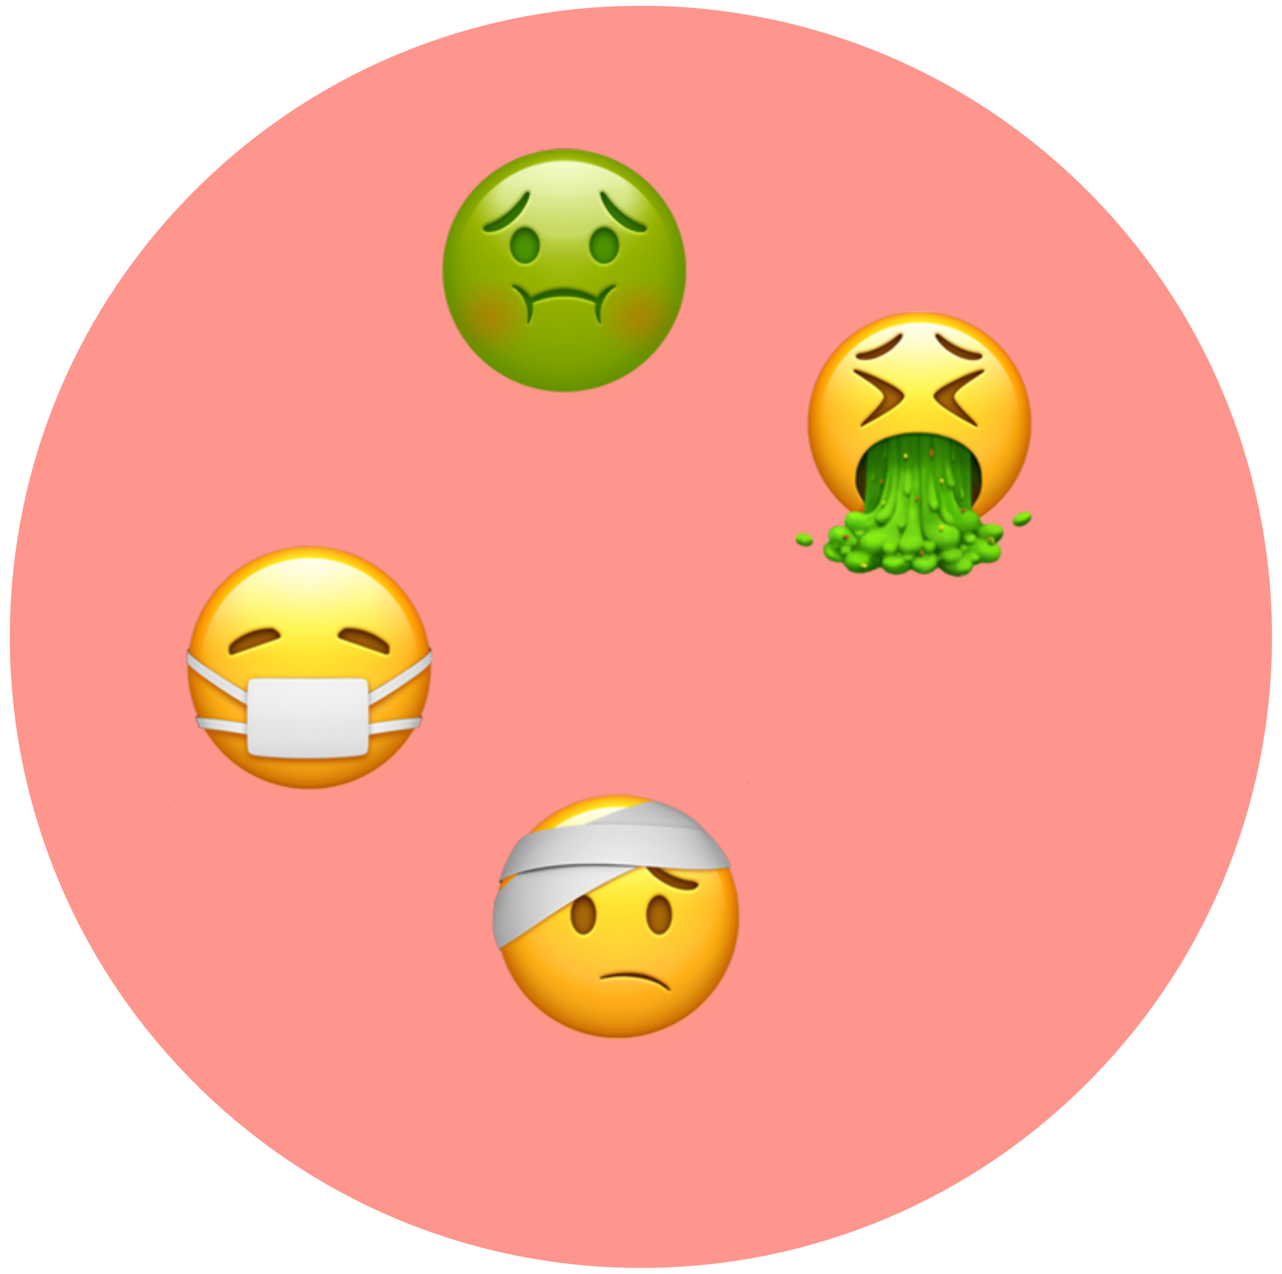

Should not be in the test

What is the probability this happens?

P(none in test) =

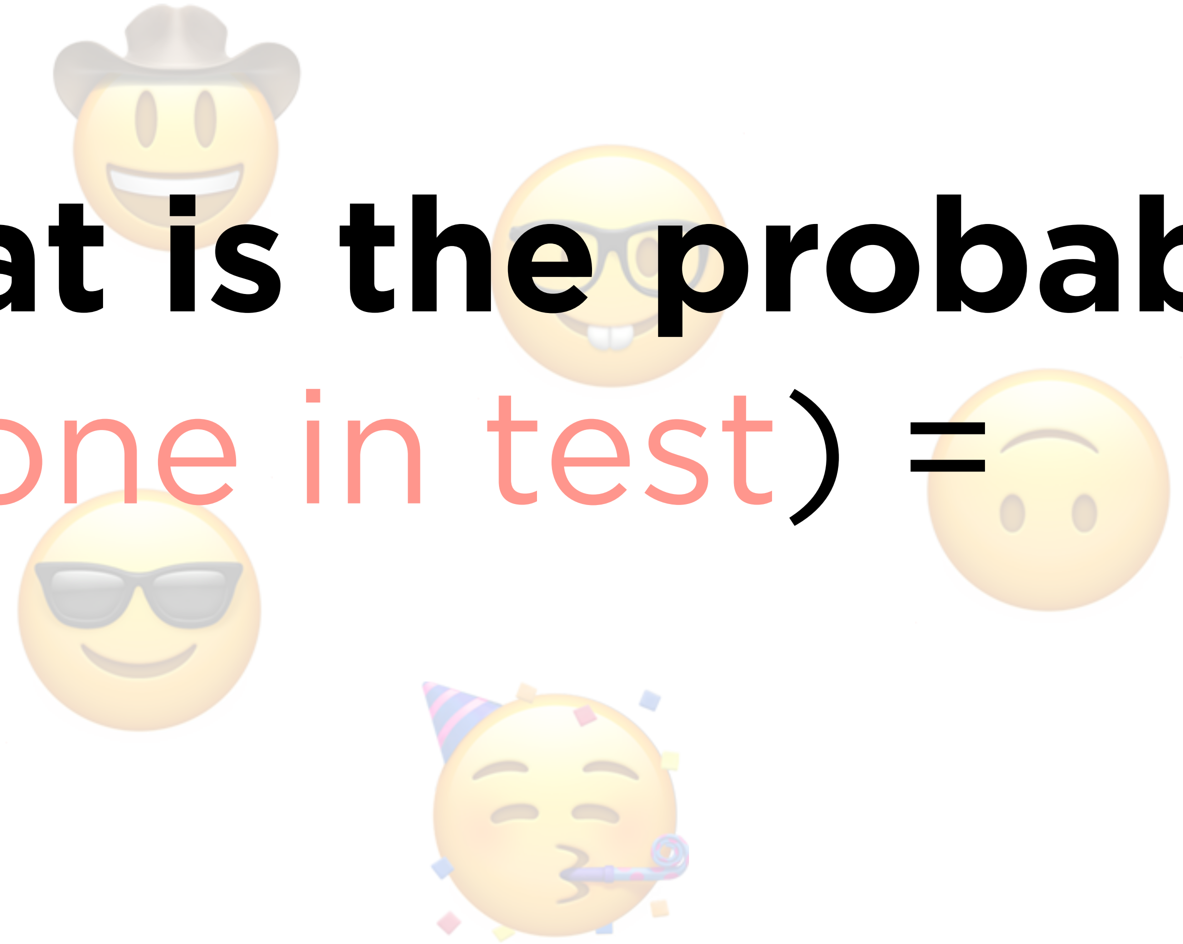

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

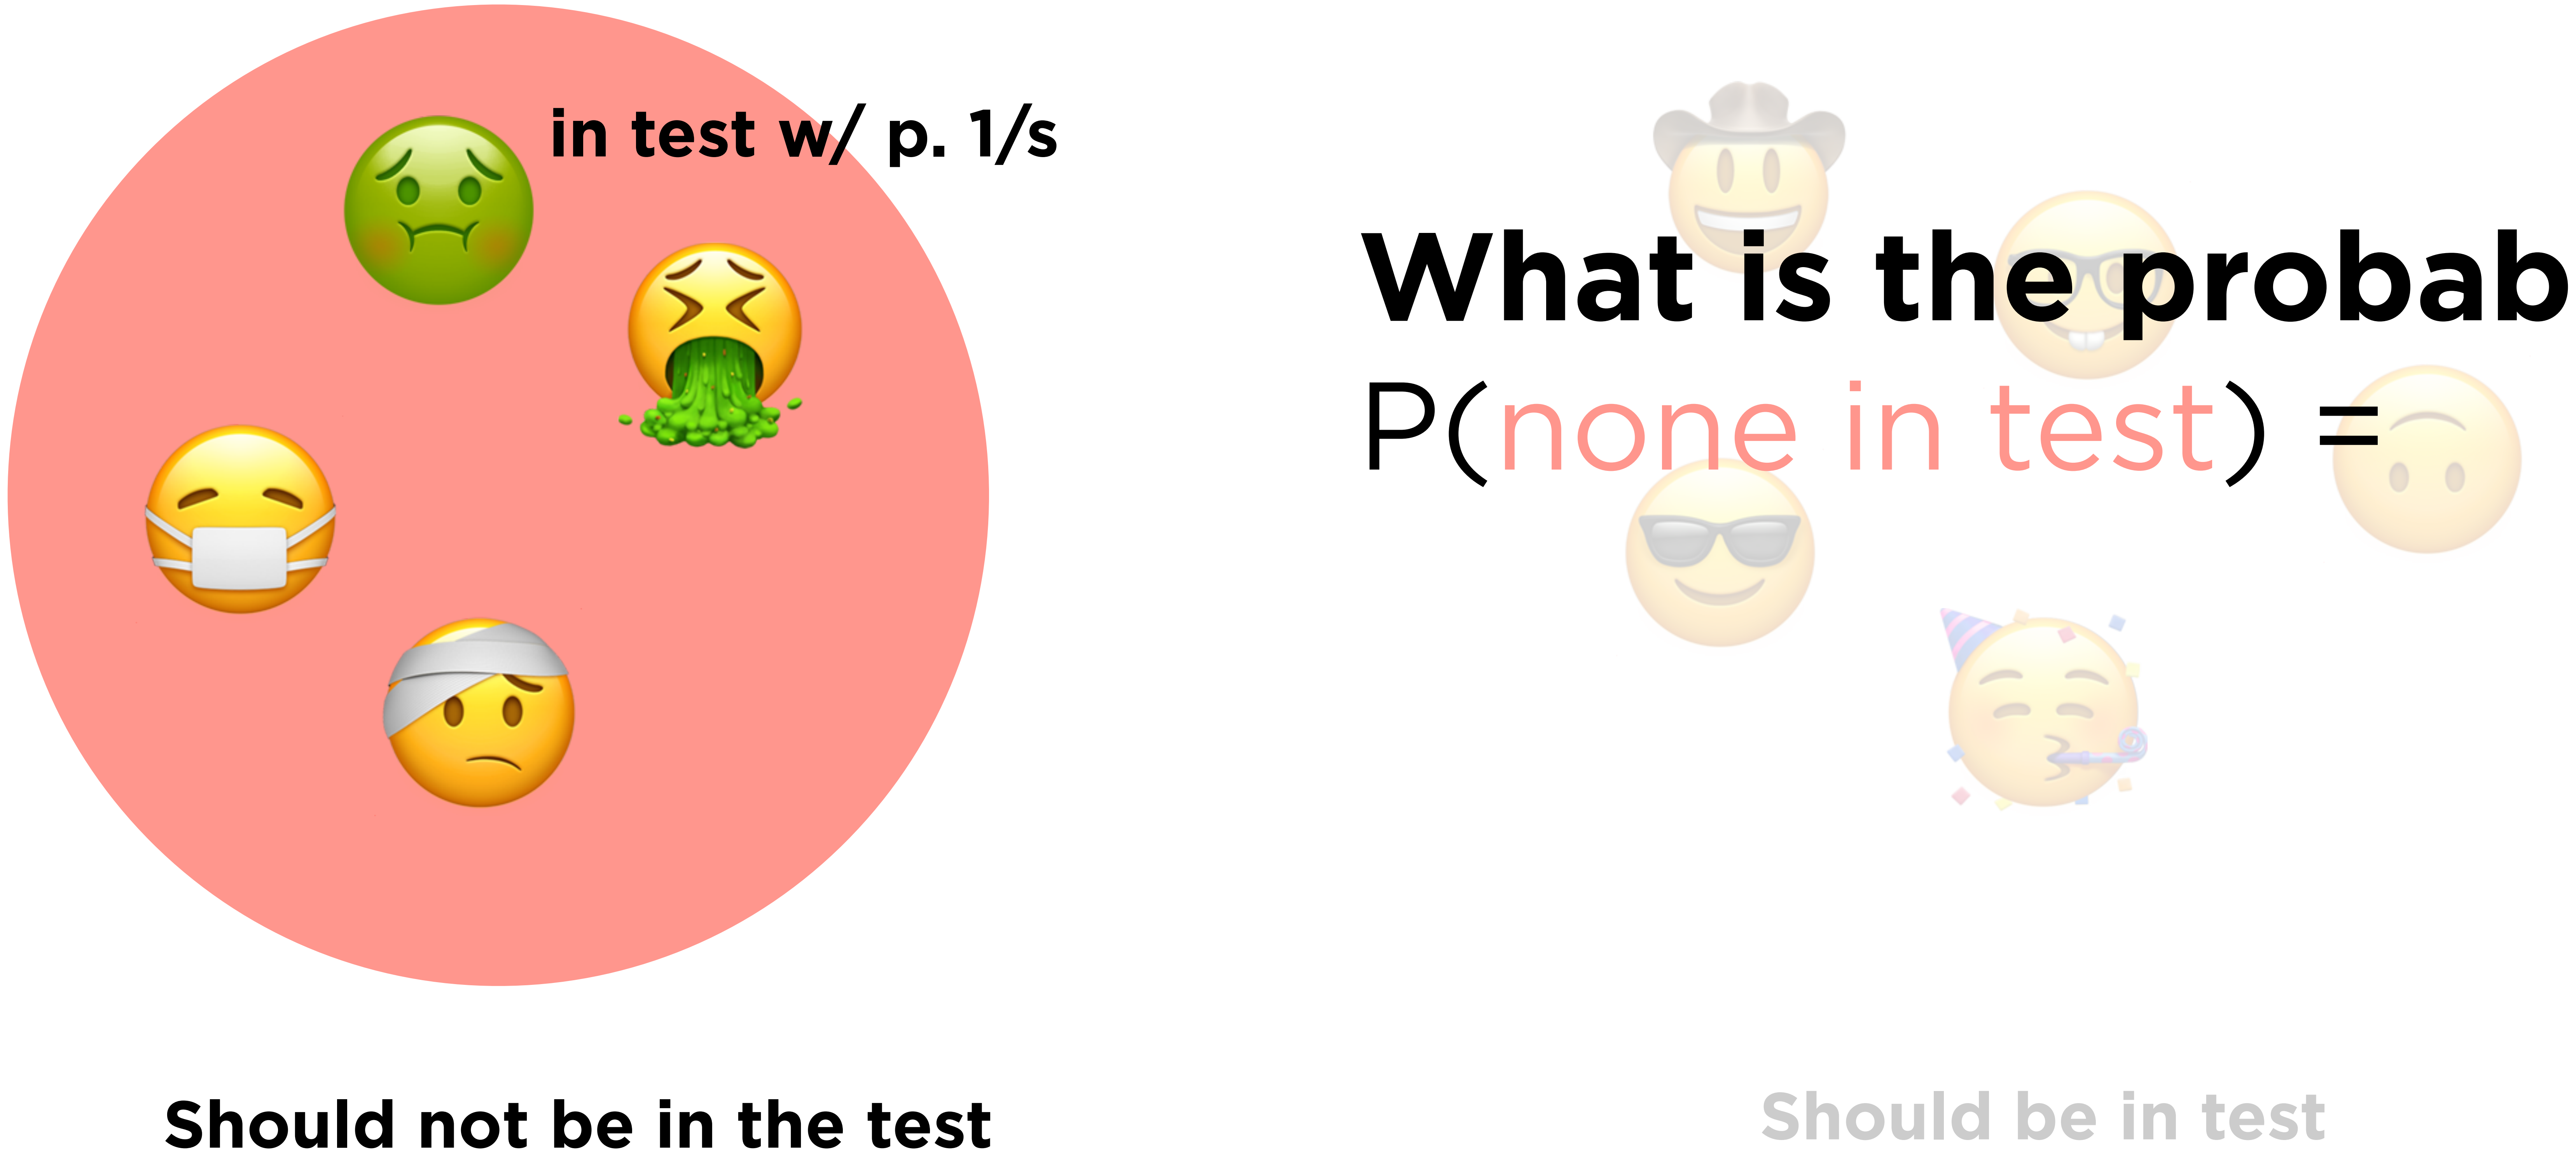

What is the probability this happens?

P(none in test) =

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

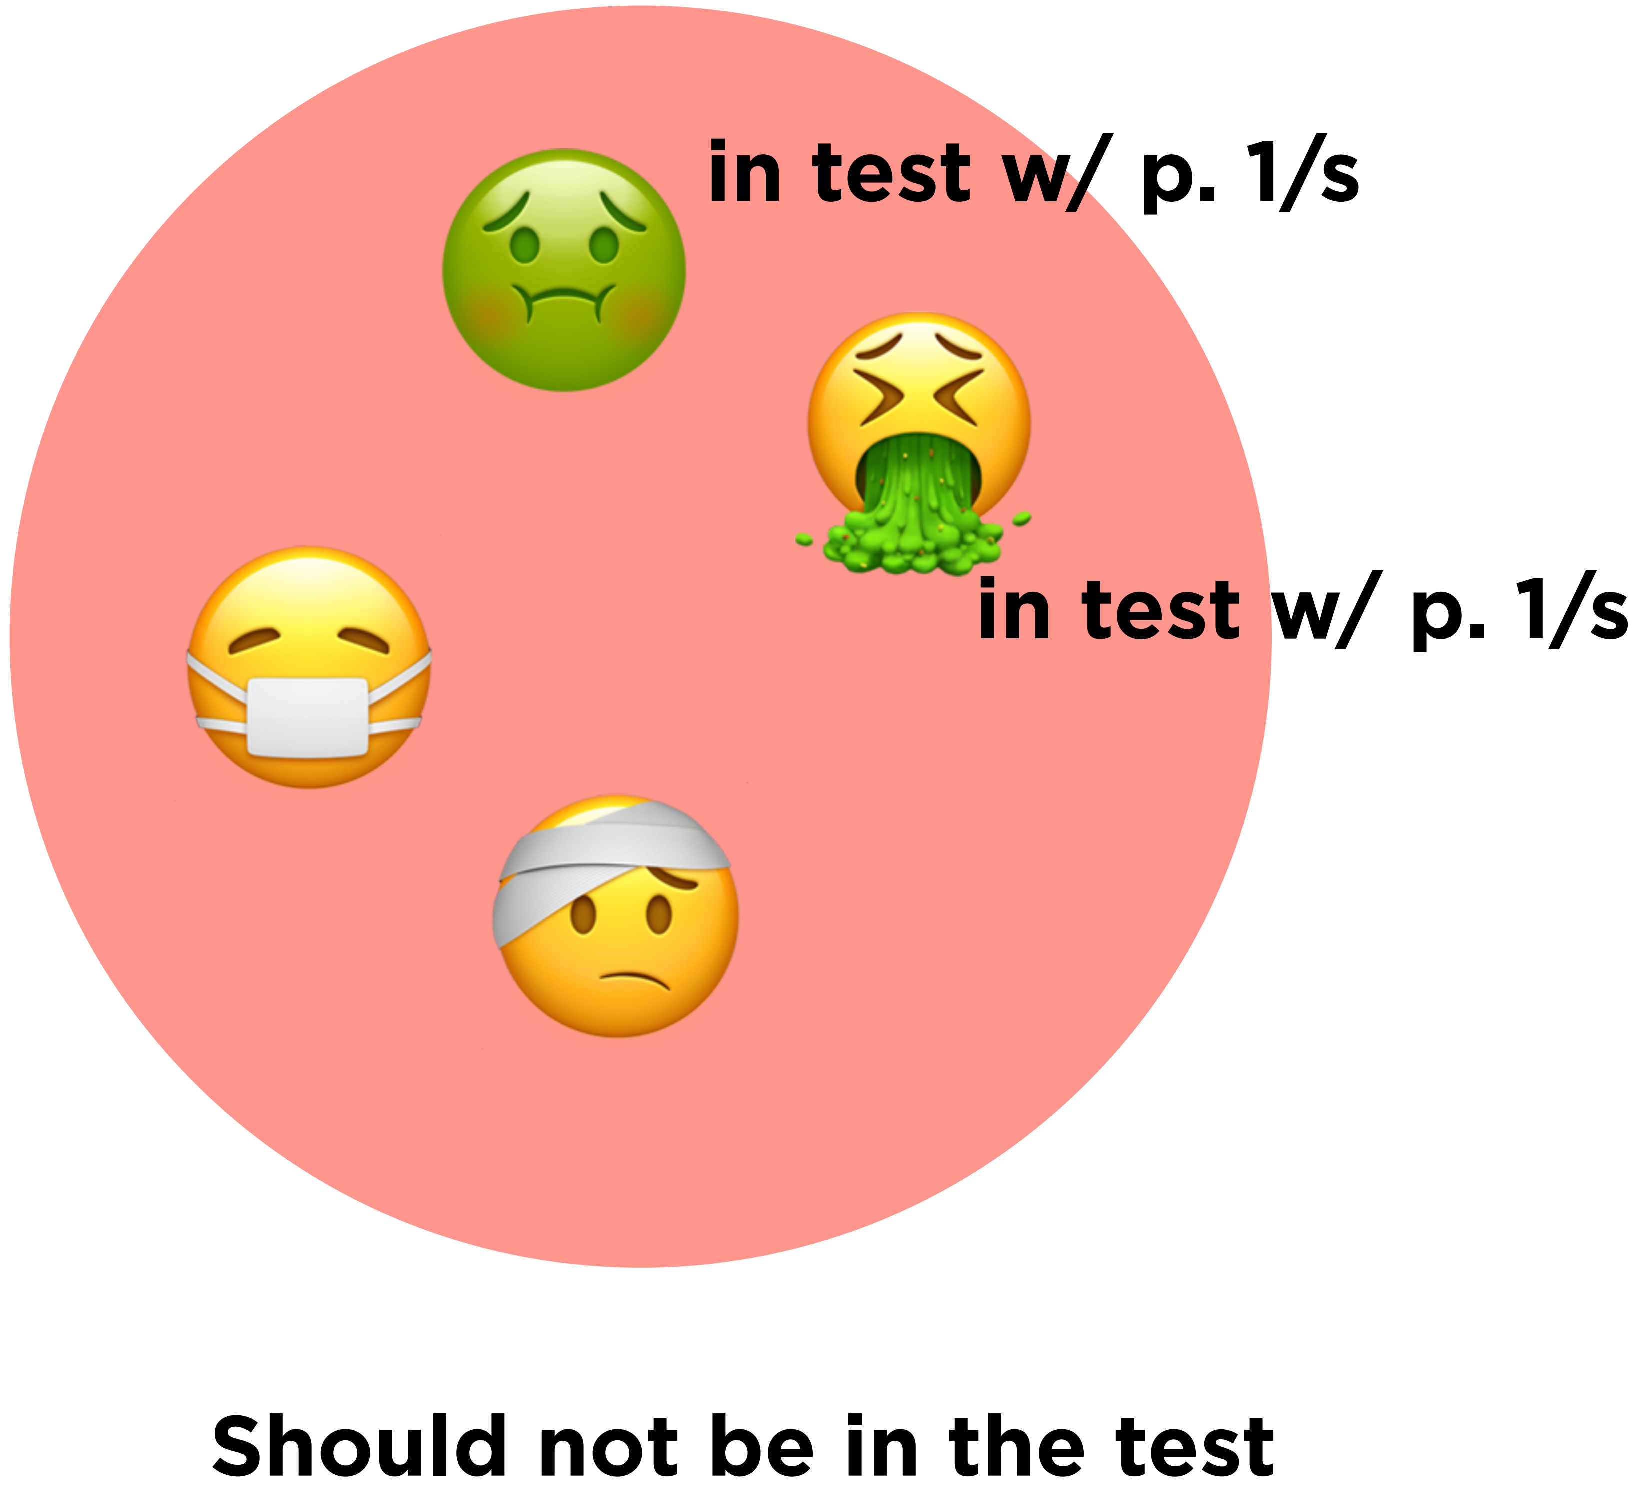

What is the probability this happens?

P(none in test) =

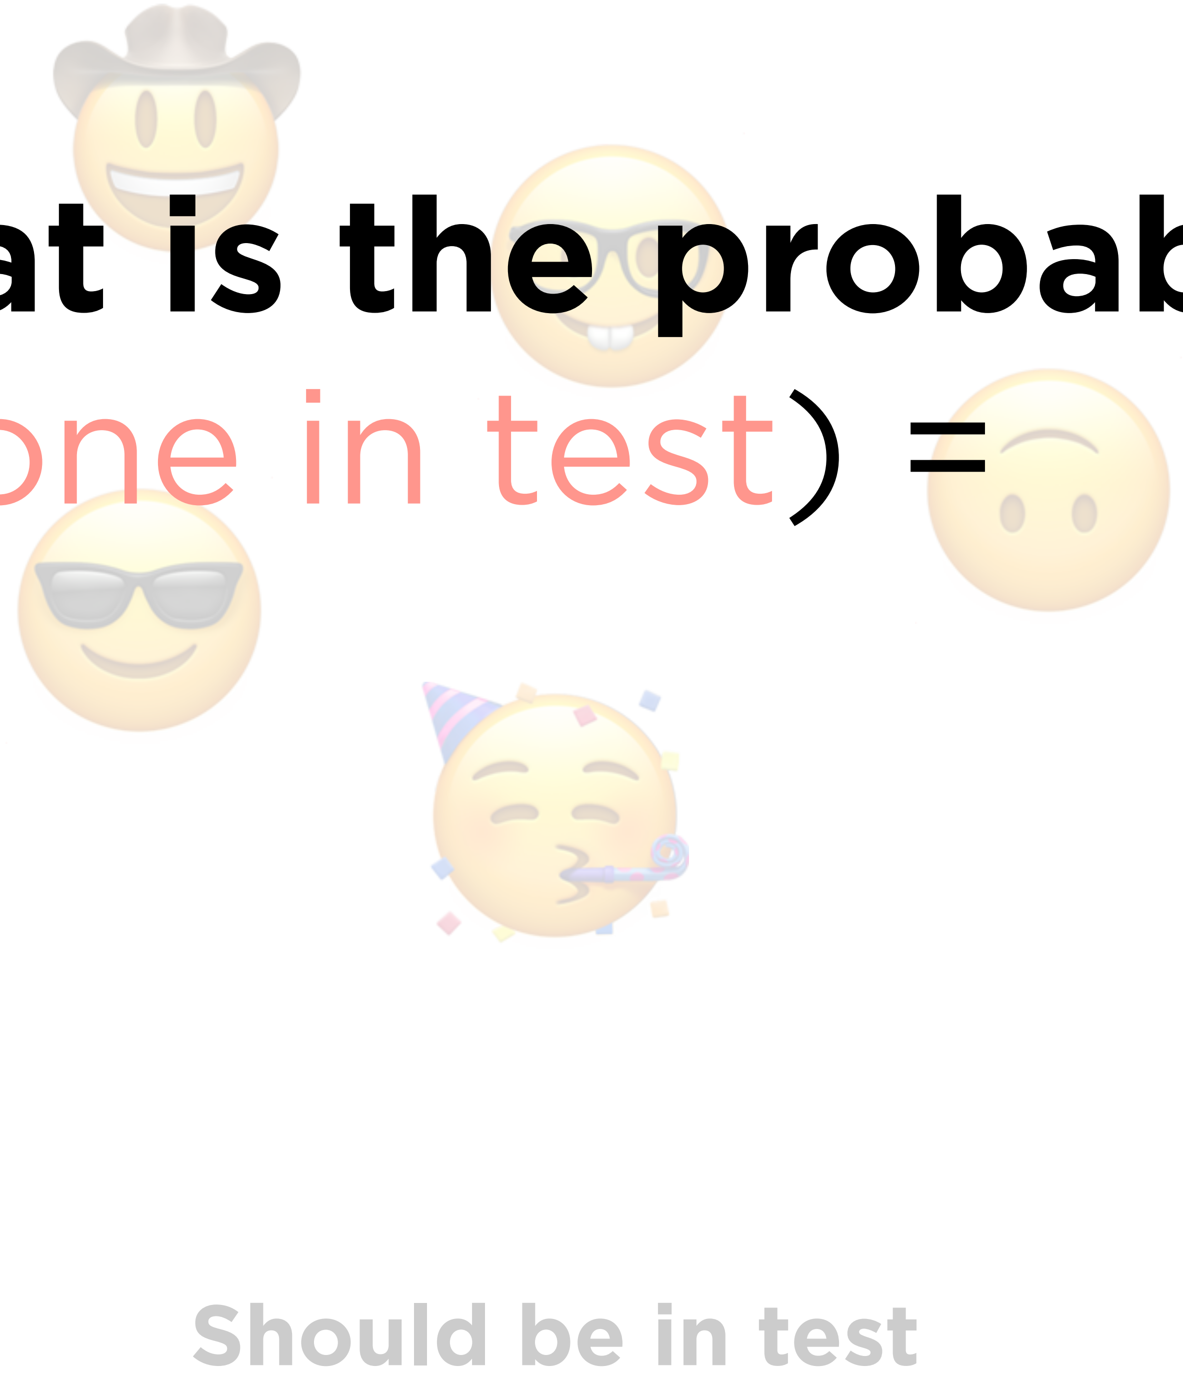

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

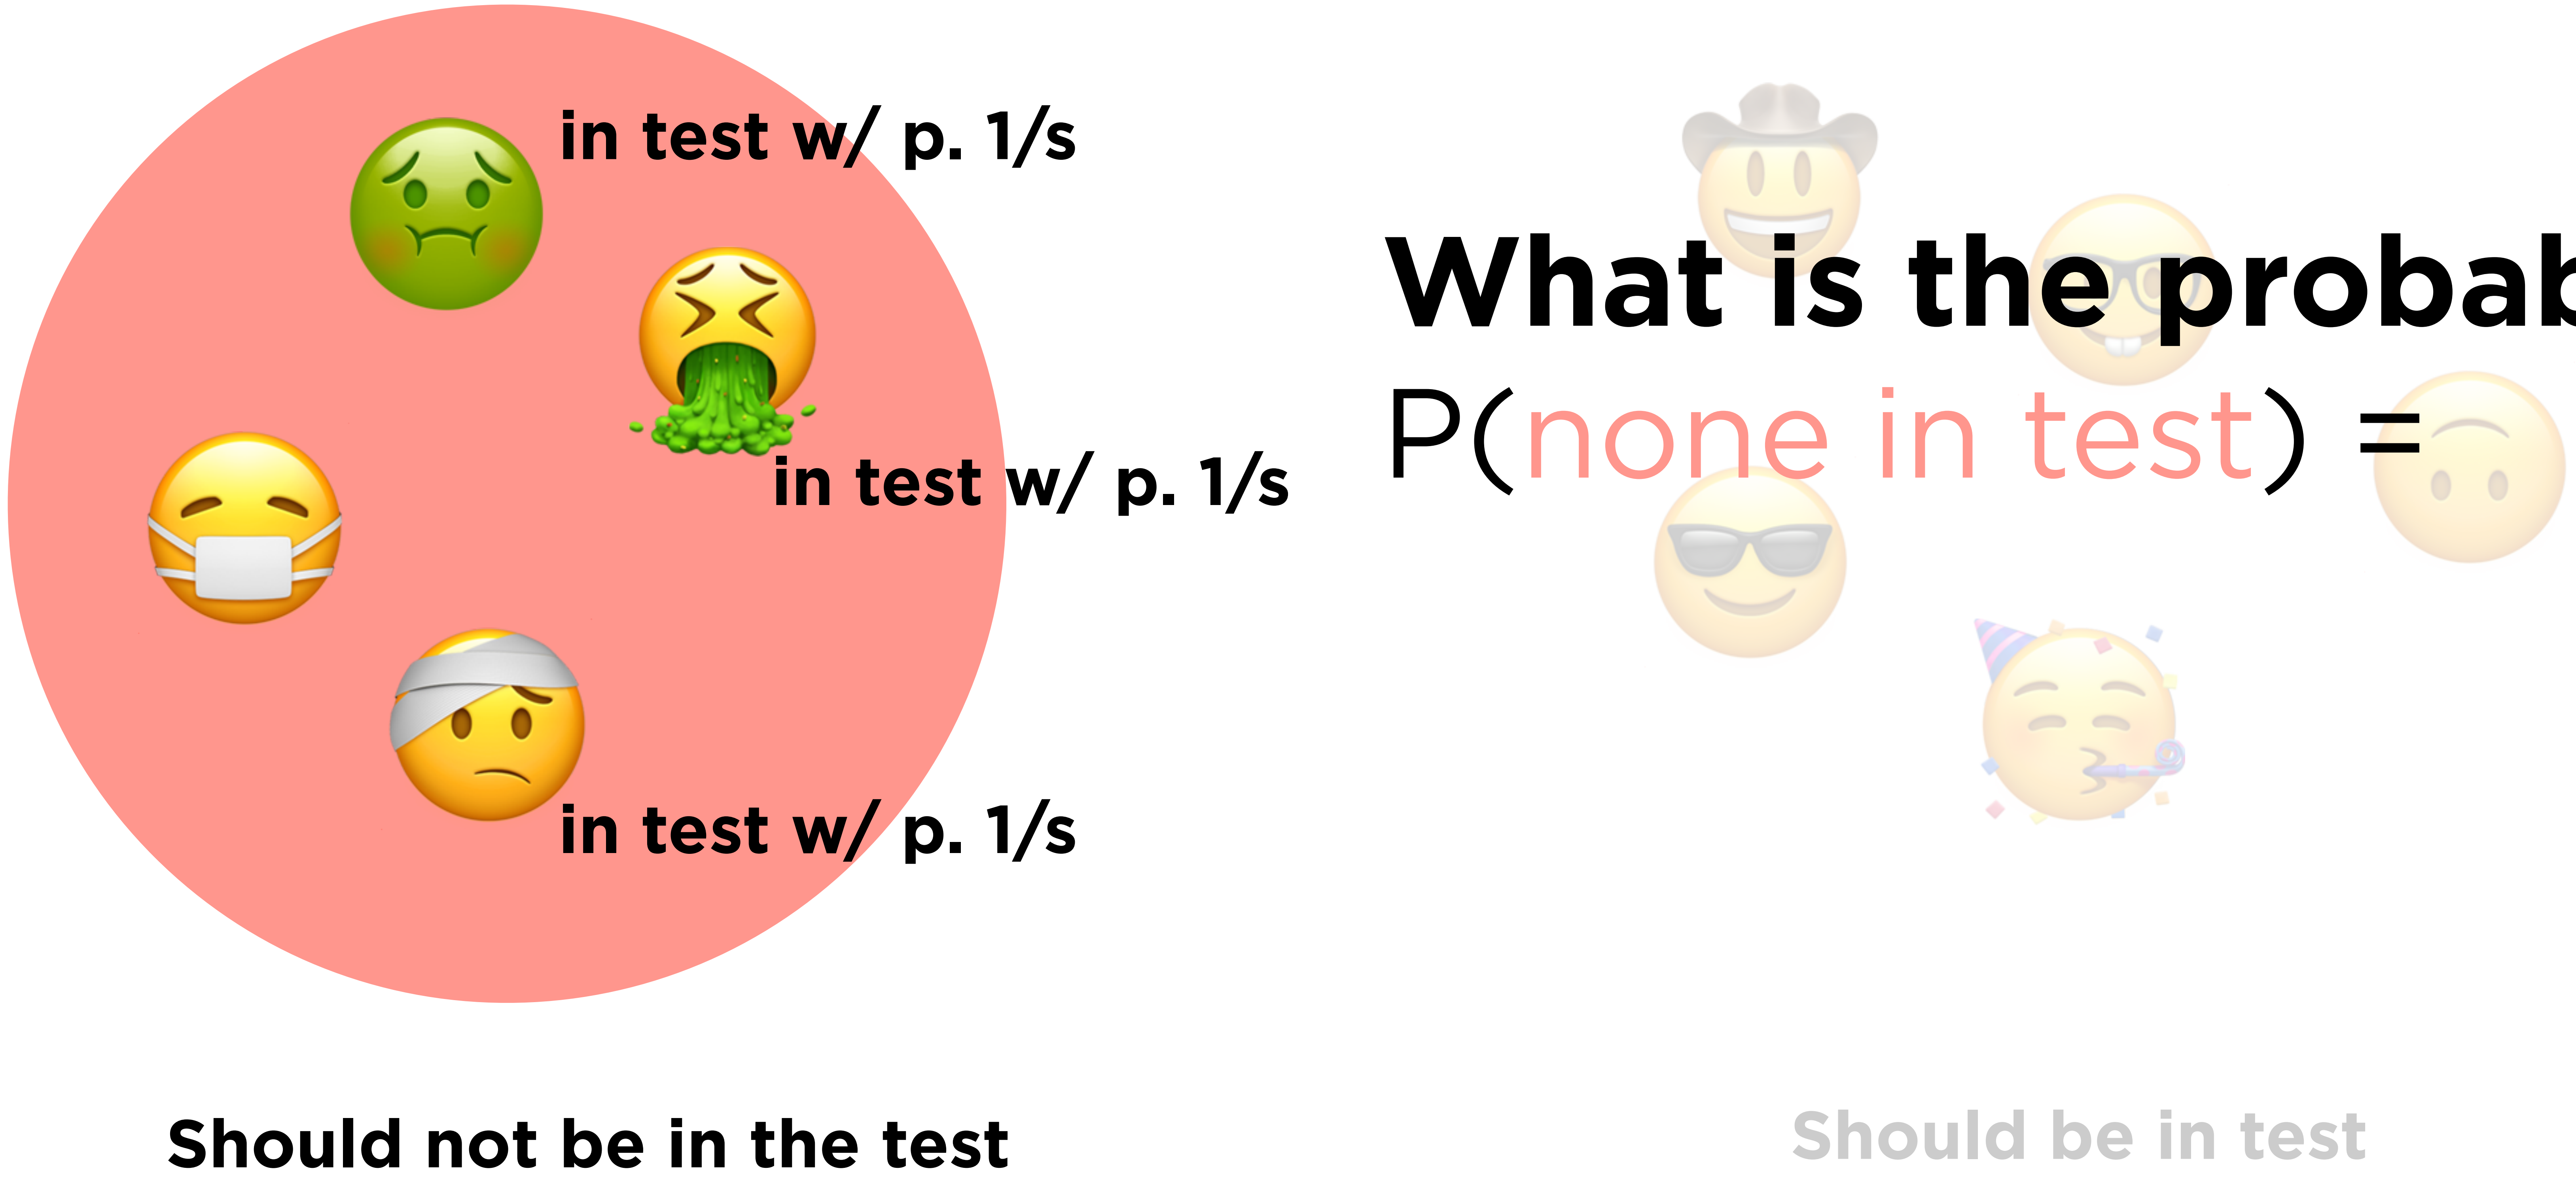

What is the probability this happens?

P(none in test) =

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

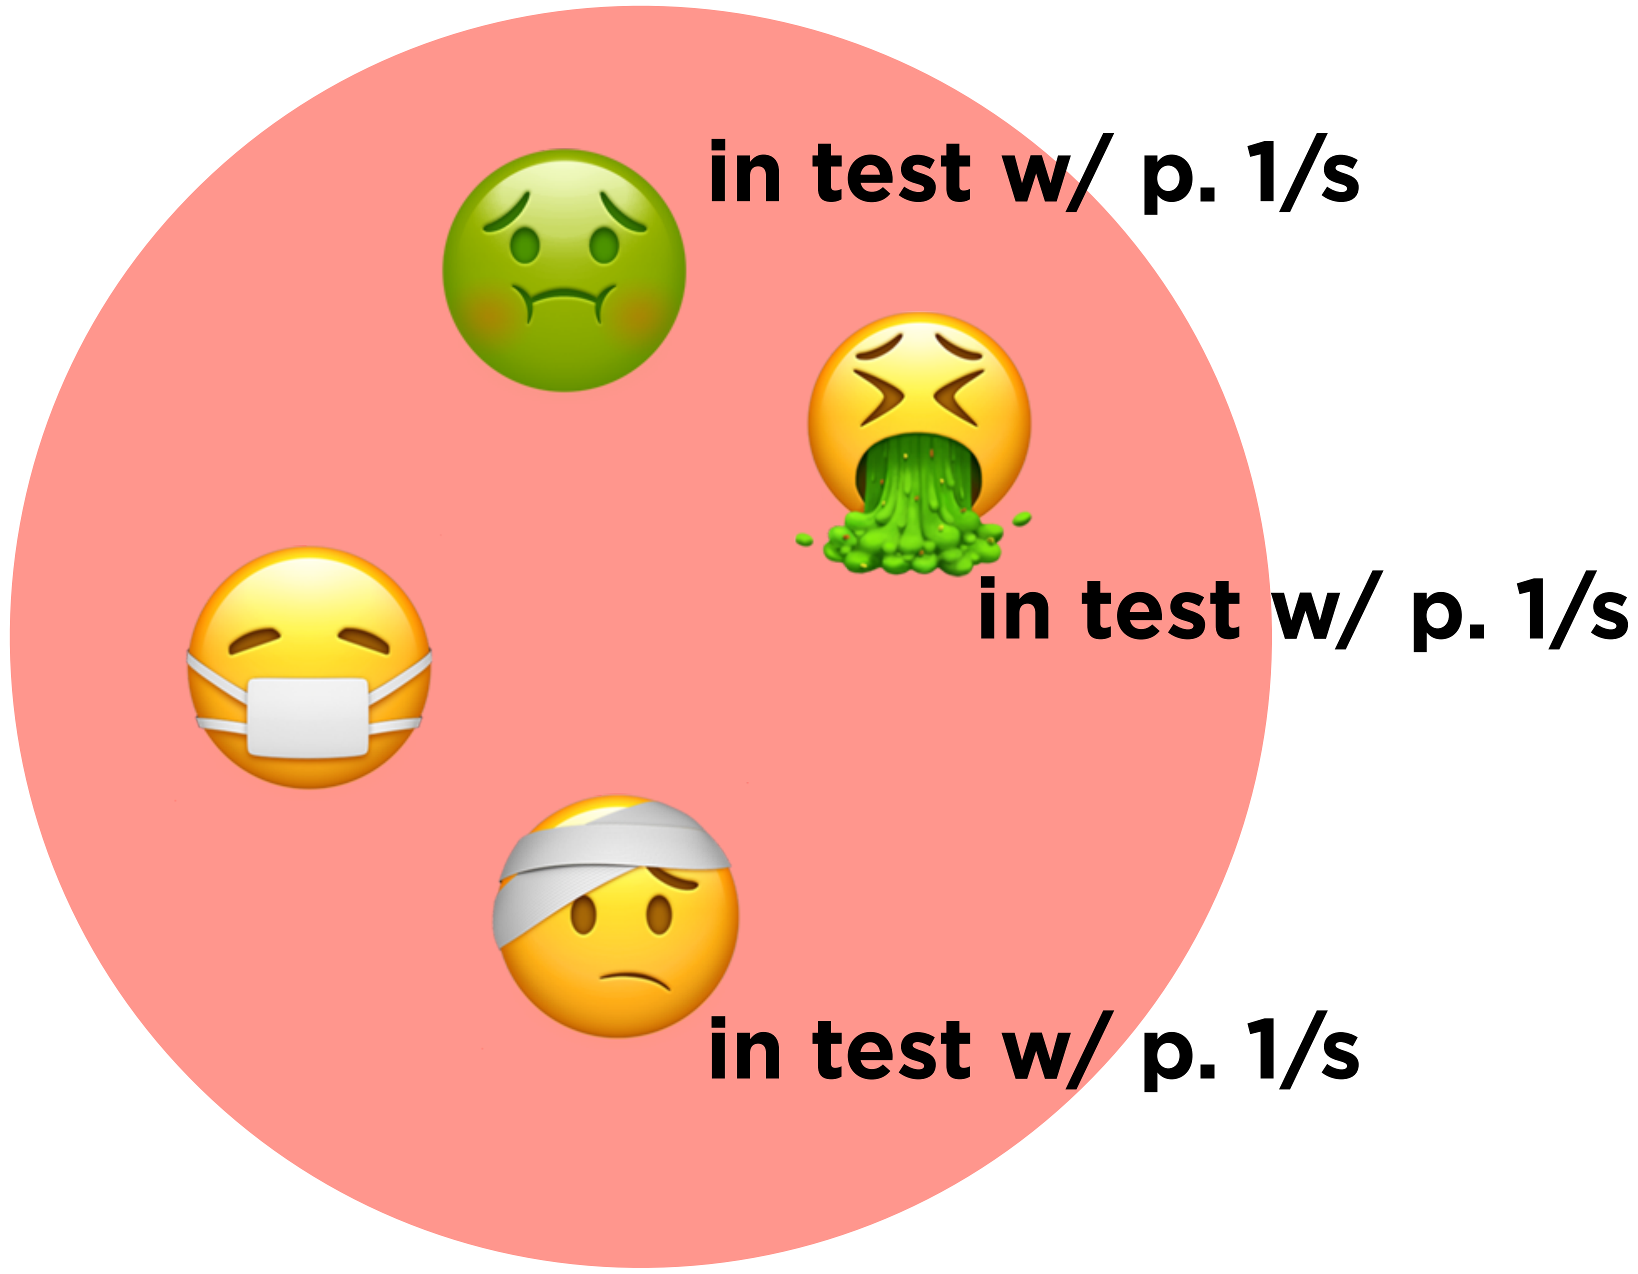

Should not be in the test

What is the probability this happens?

$P(\text{none in test}) = (1-1/s)^s$

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

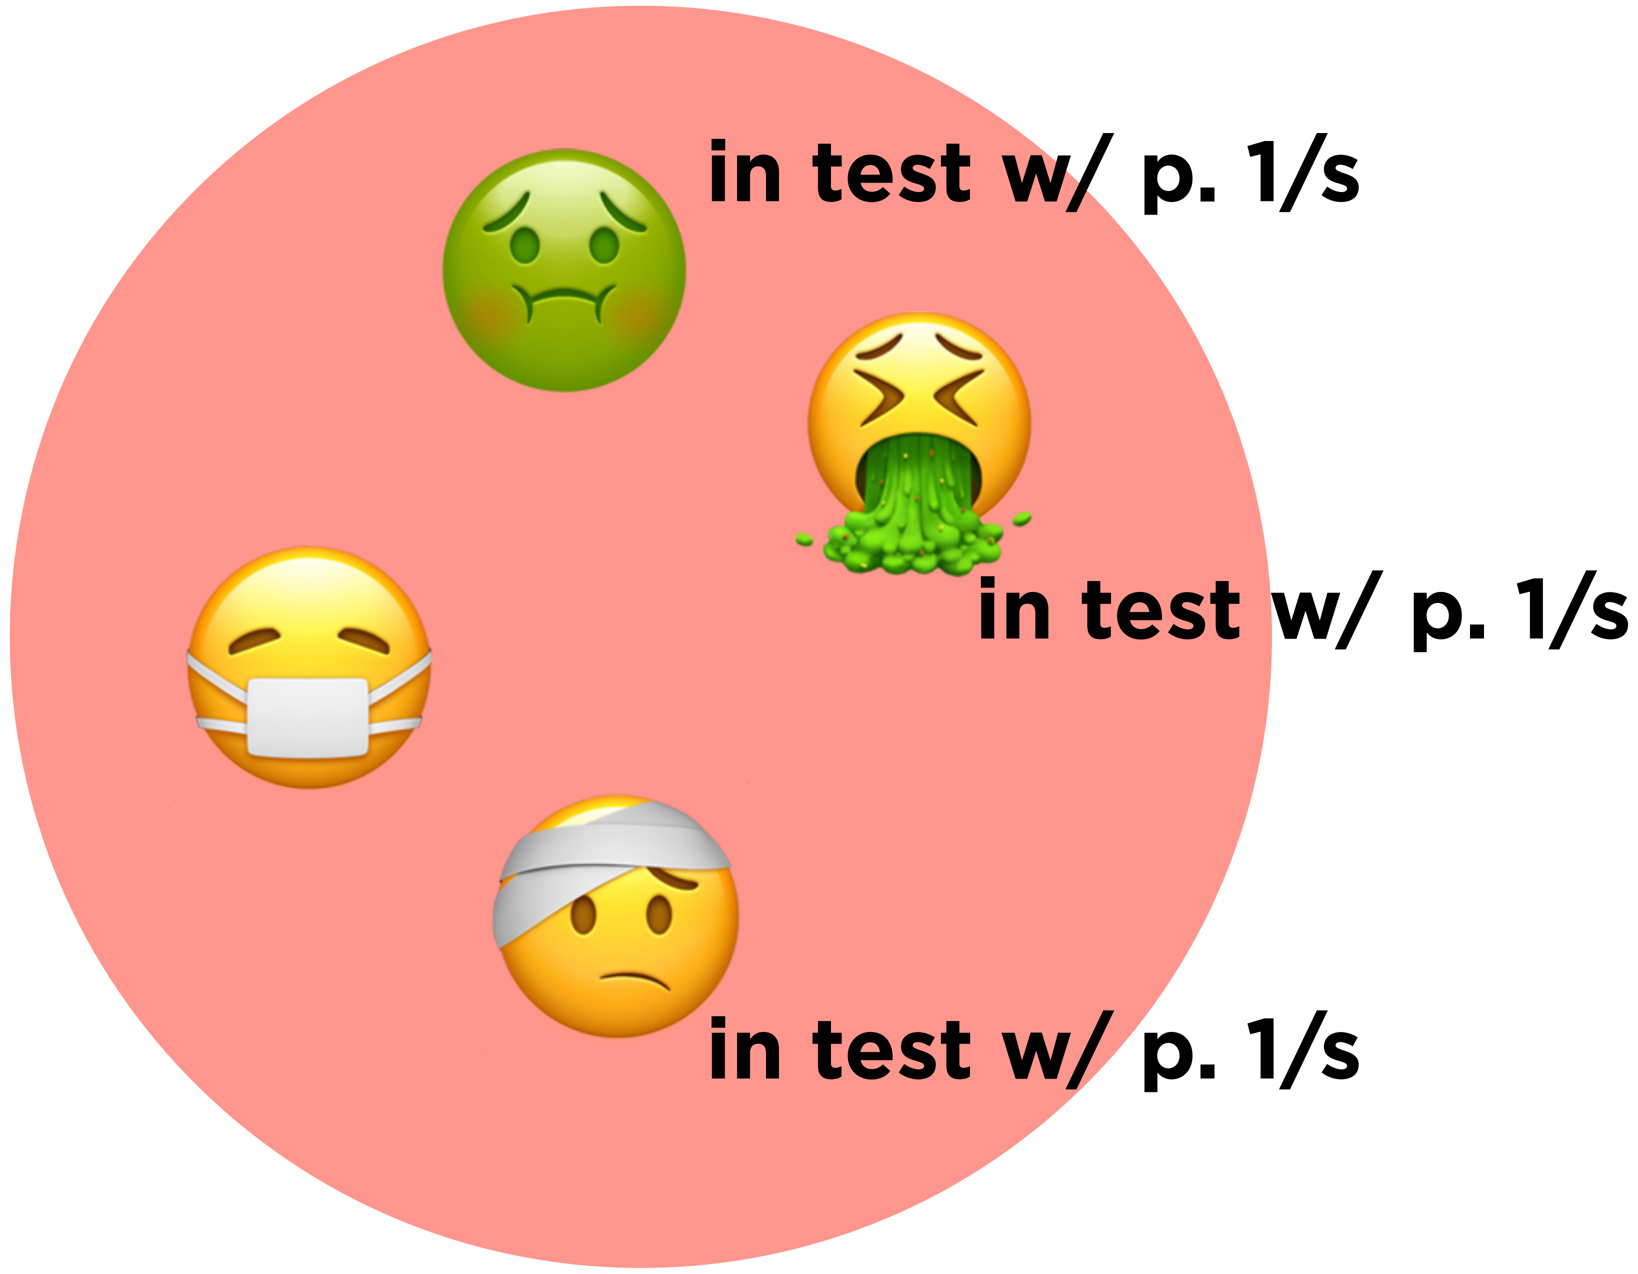

Should not be in the test

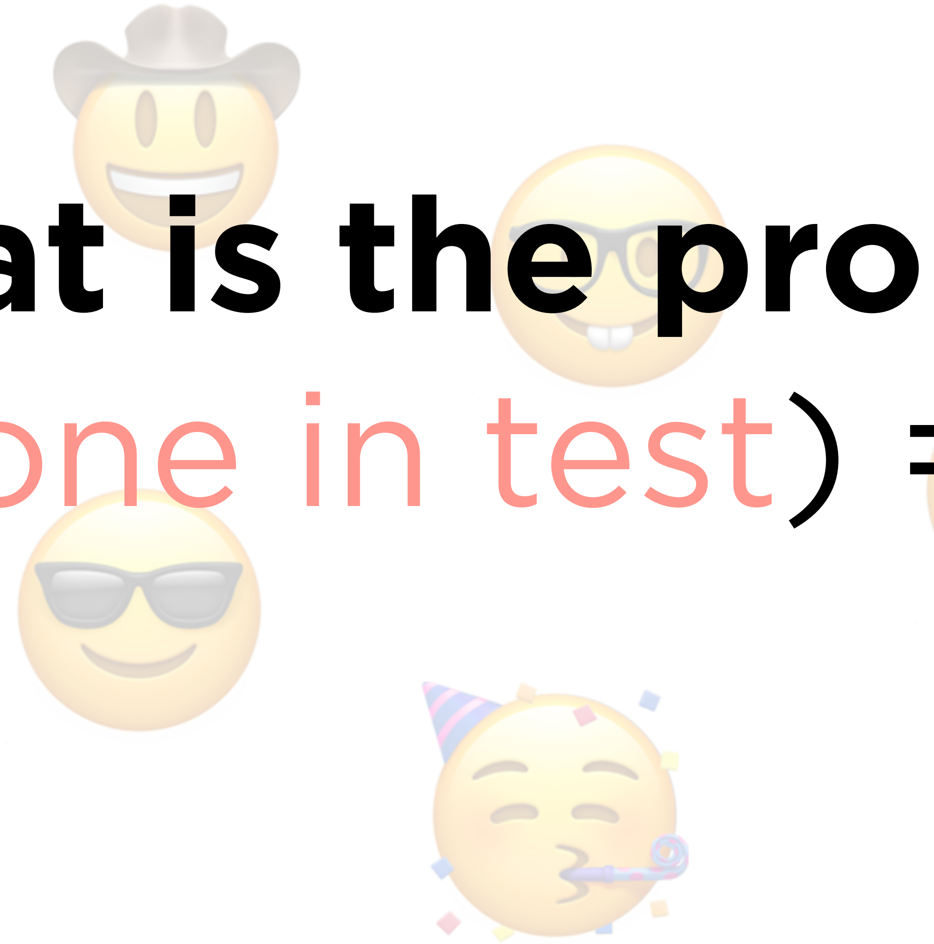

What is the probability this happens?

$P(\text{none in test}) = (1-1/s)^s \approx e^{-s/s}$

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

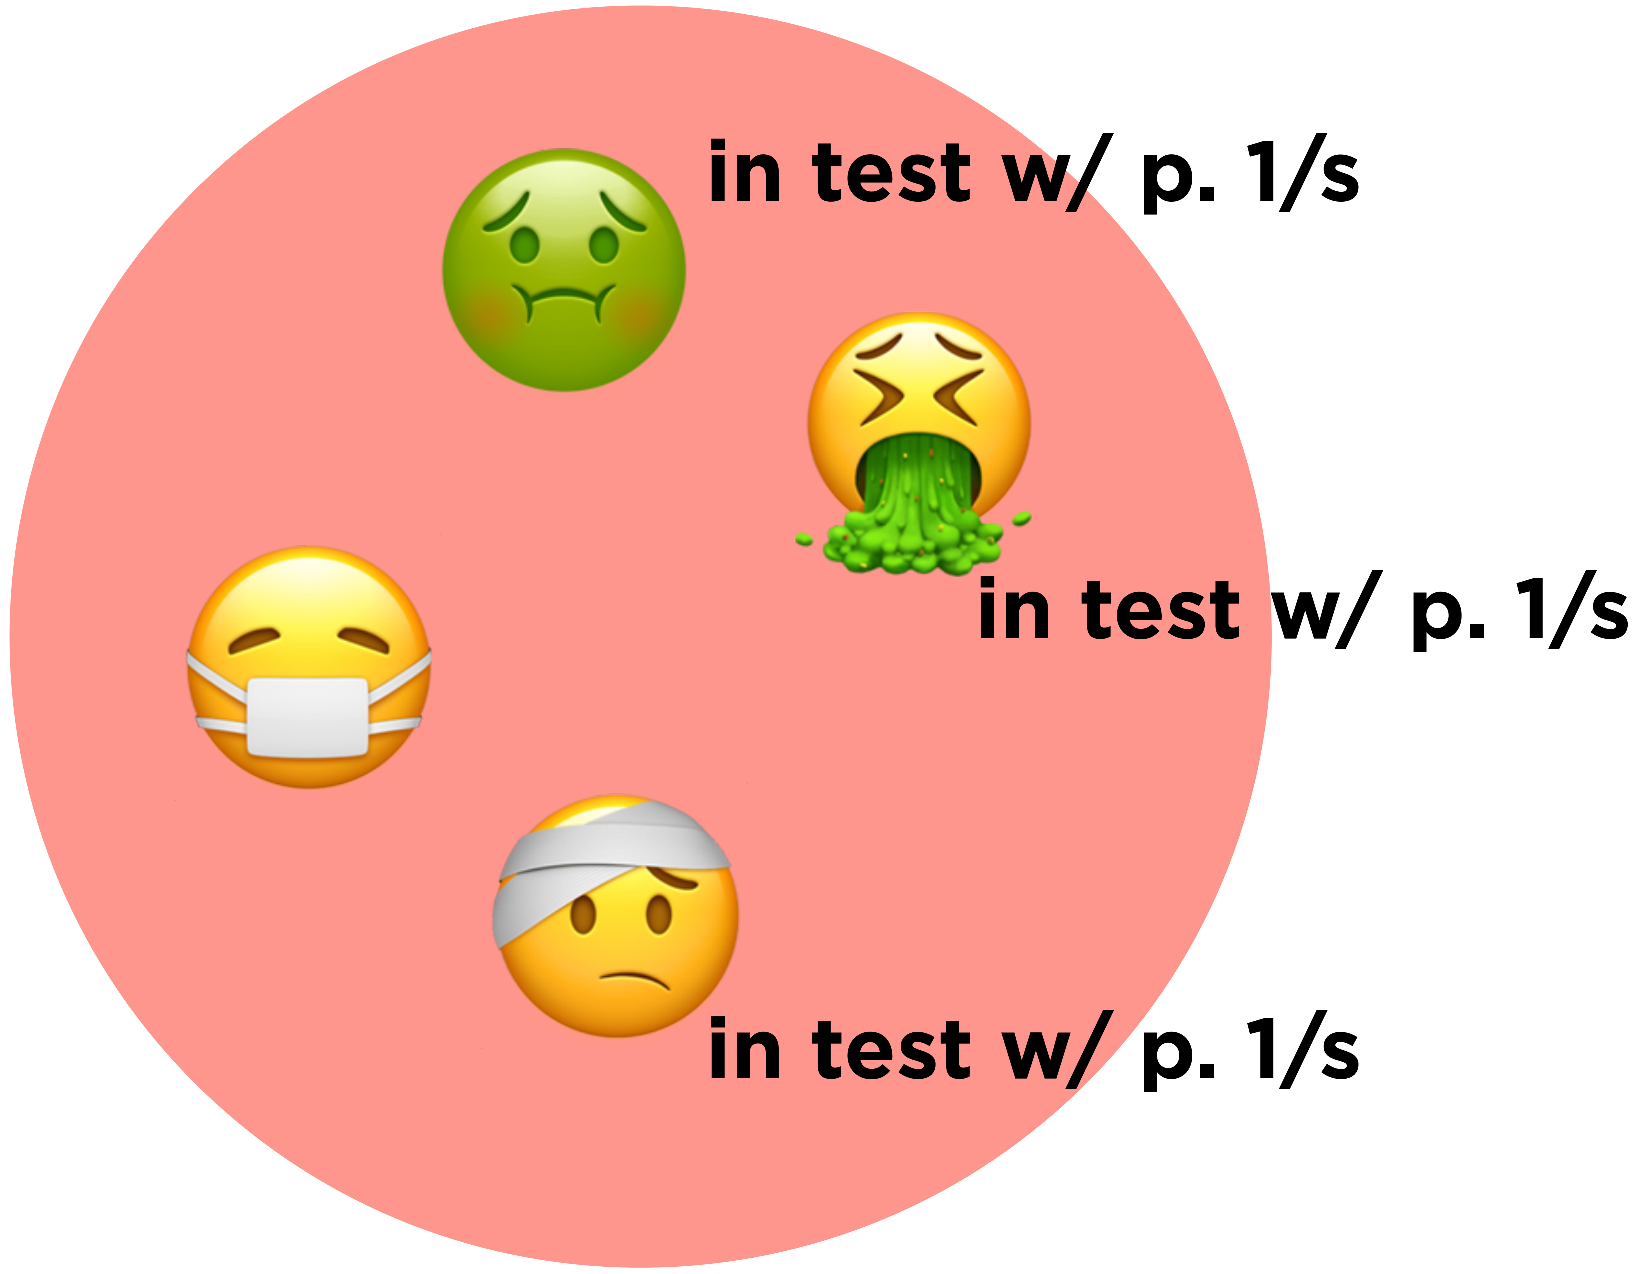

Should not be in the test

What is the probability this happens?

$P(\text{none in test}) = (1-1/s)^s \approx e^{-s/s} \approx 1/3$

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

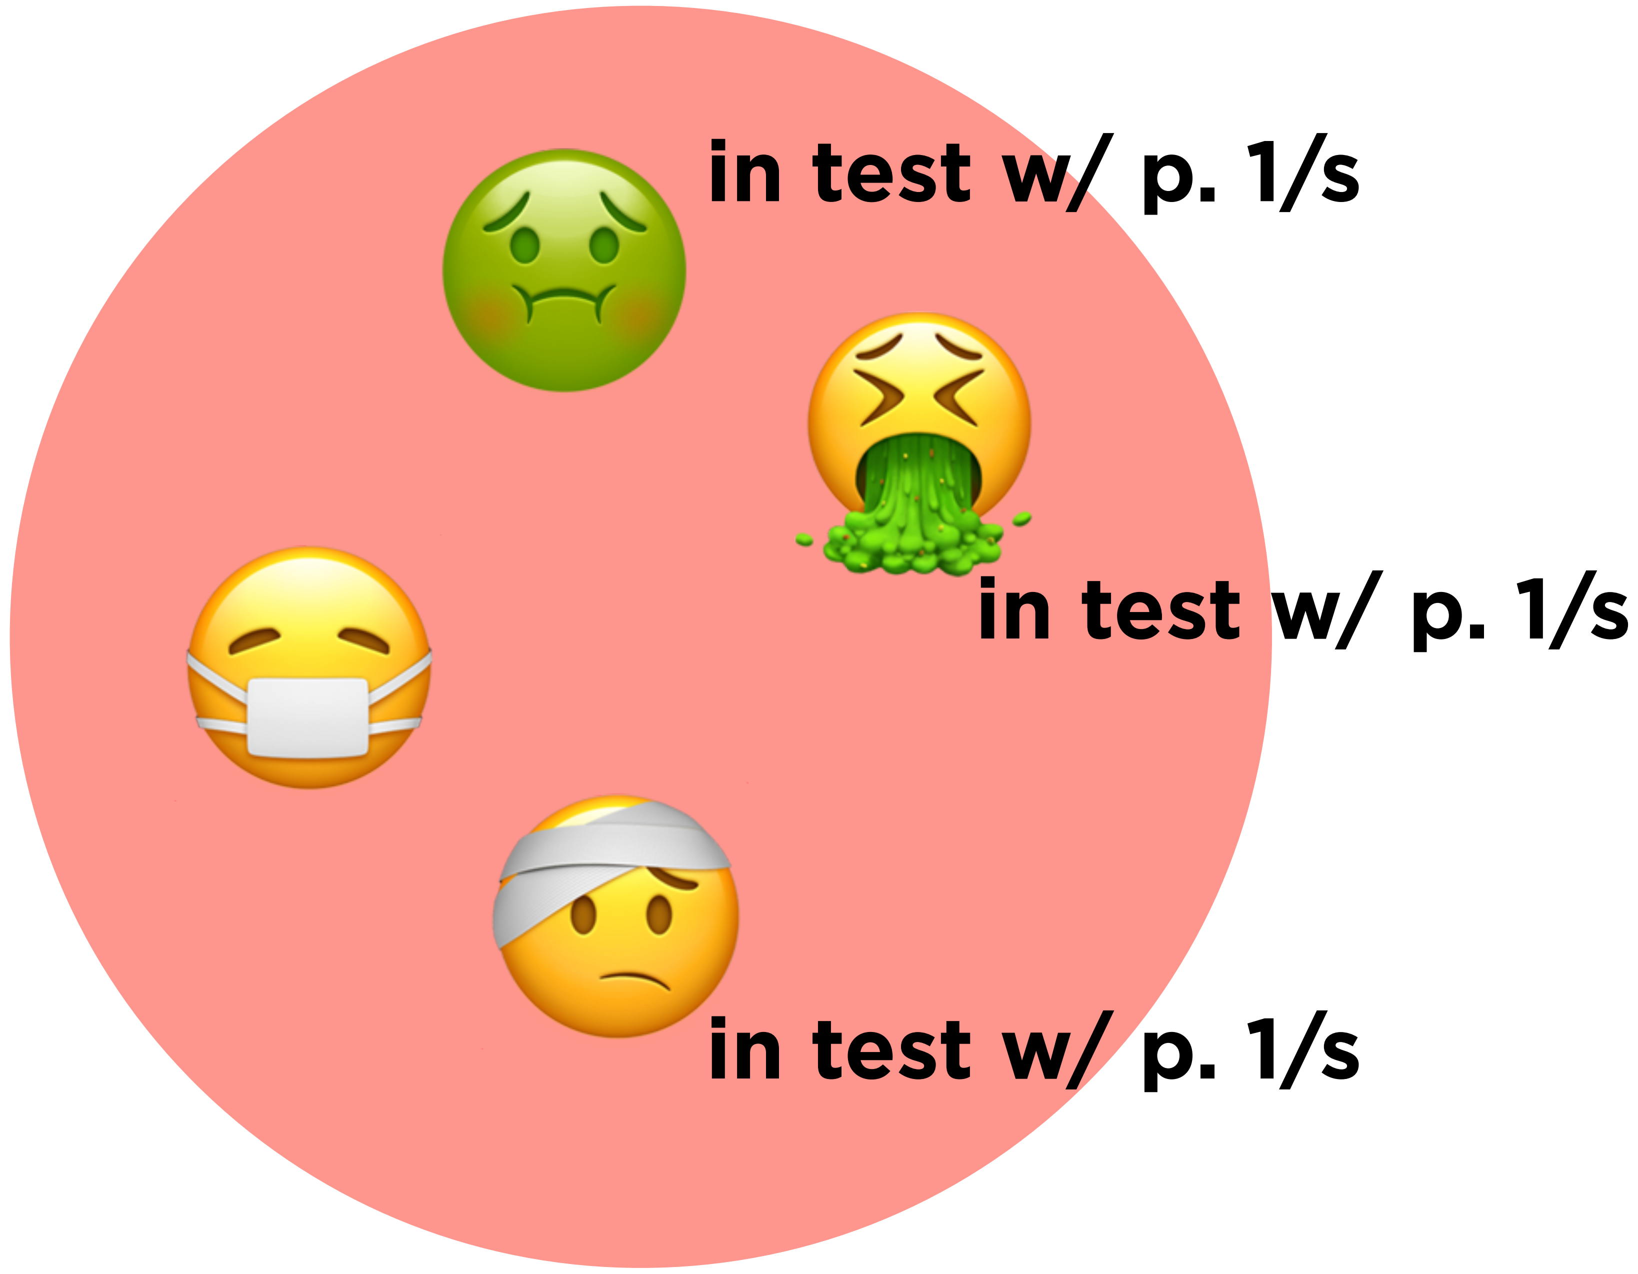

Should not be in the test

What is the probability this happens?

$P(\text{none in test}) = (1-1/s)^s \approx e^{-s/s} \approx 1/3$

Idea: not too many sick people, so pretty good probability of missing 'em all

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

What is the probability this happens?

P(🙃 in test) = 1/s

Not in test w/ probability 1/3

Need this person in test

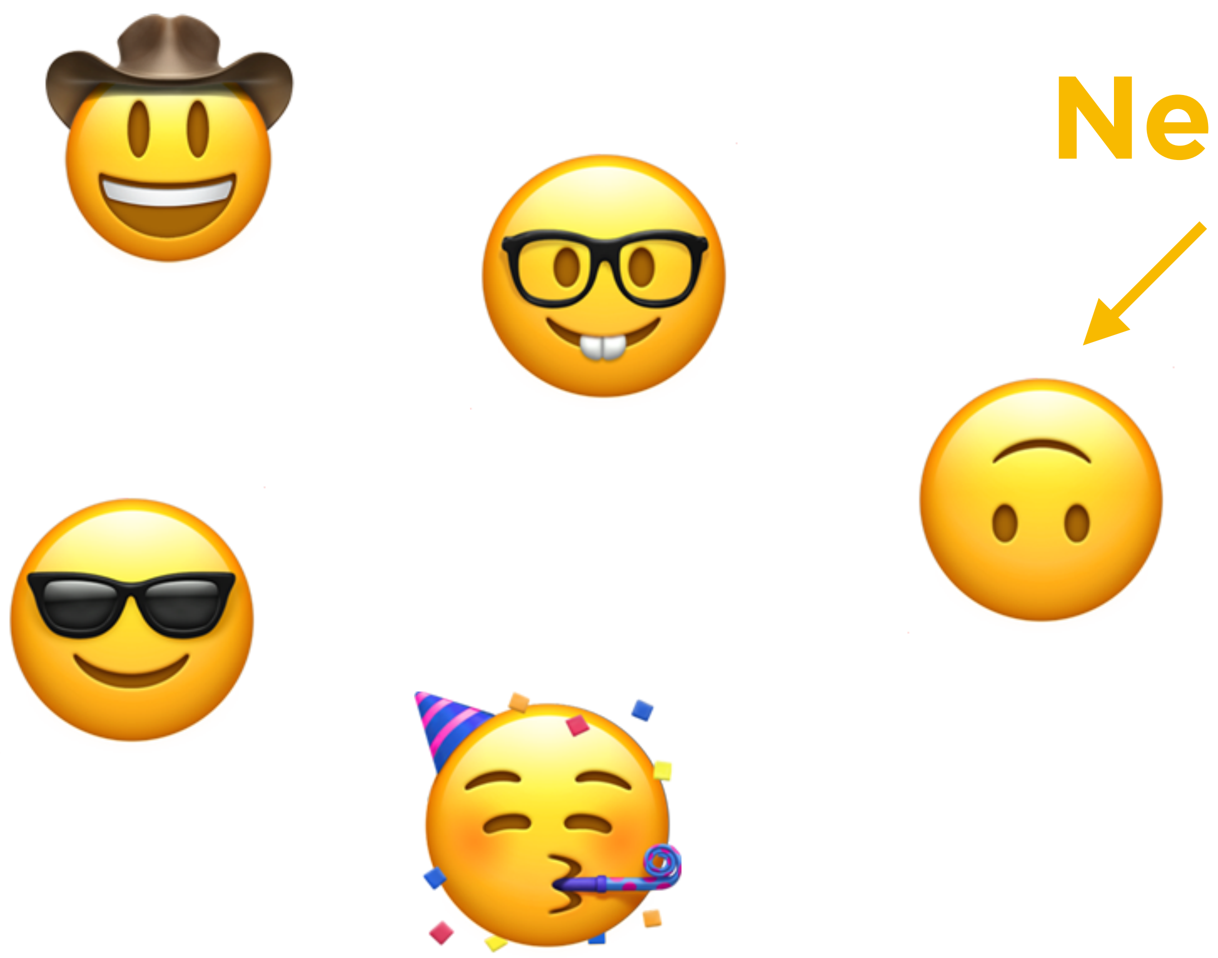

Should be in test

First idea: finding healthy people

For each set of sick people, need to be able to prove each other person is healthy

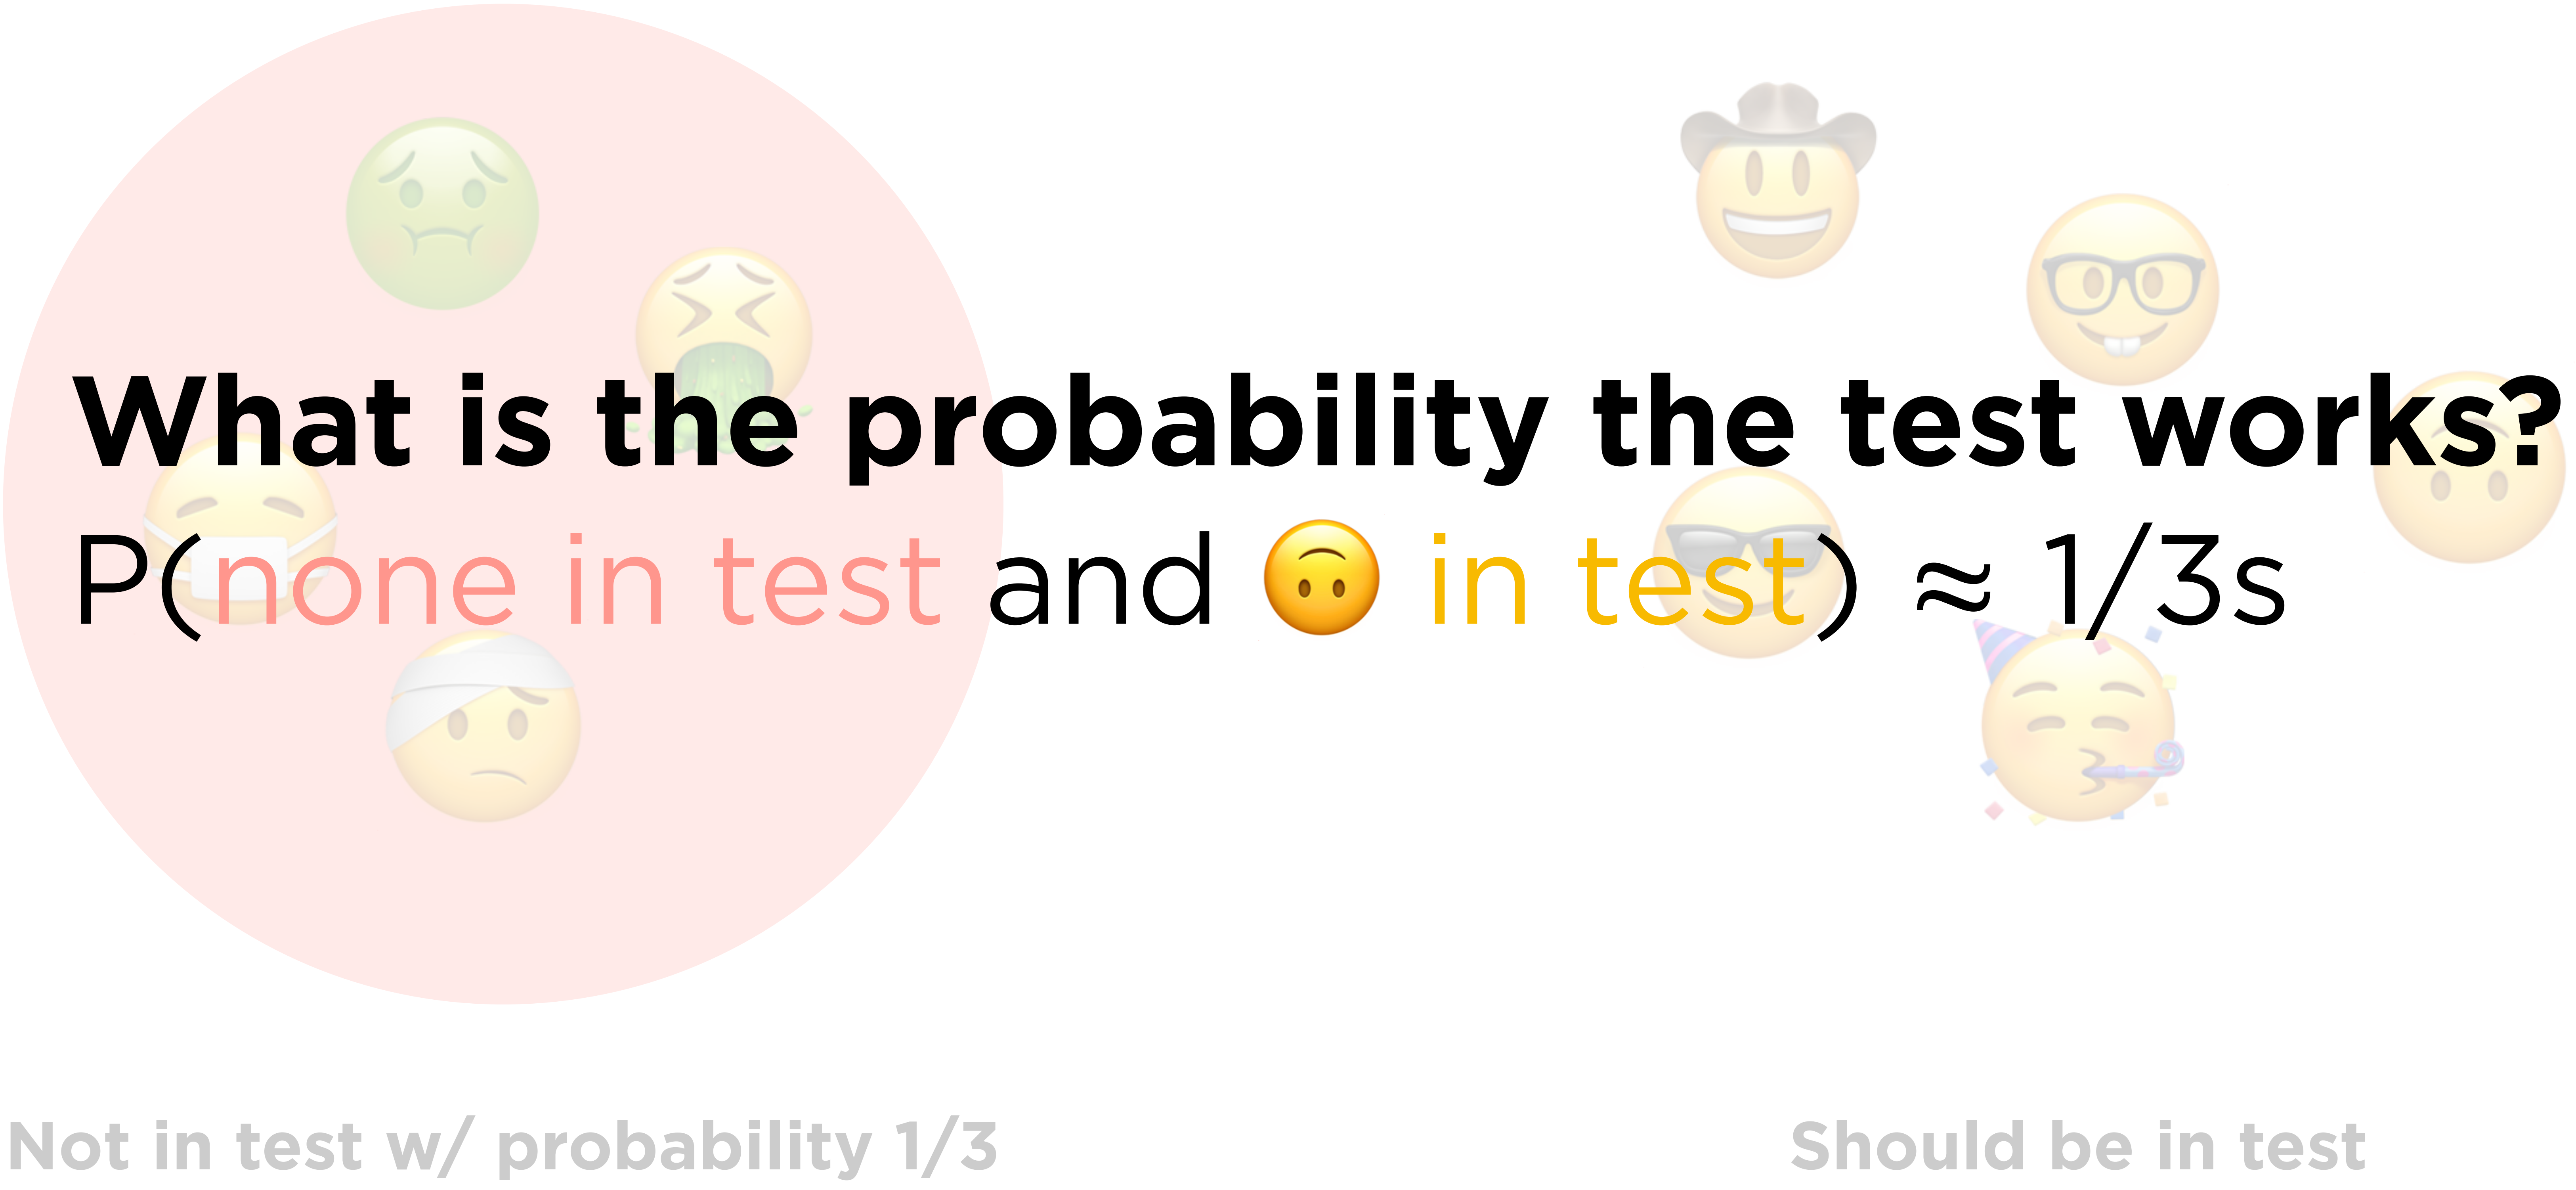

What is the probability the test works?

P(none in test and 🙃 in test) ≈ 1/3s

Repeating tests

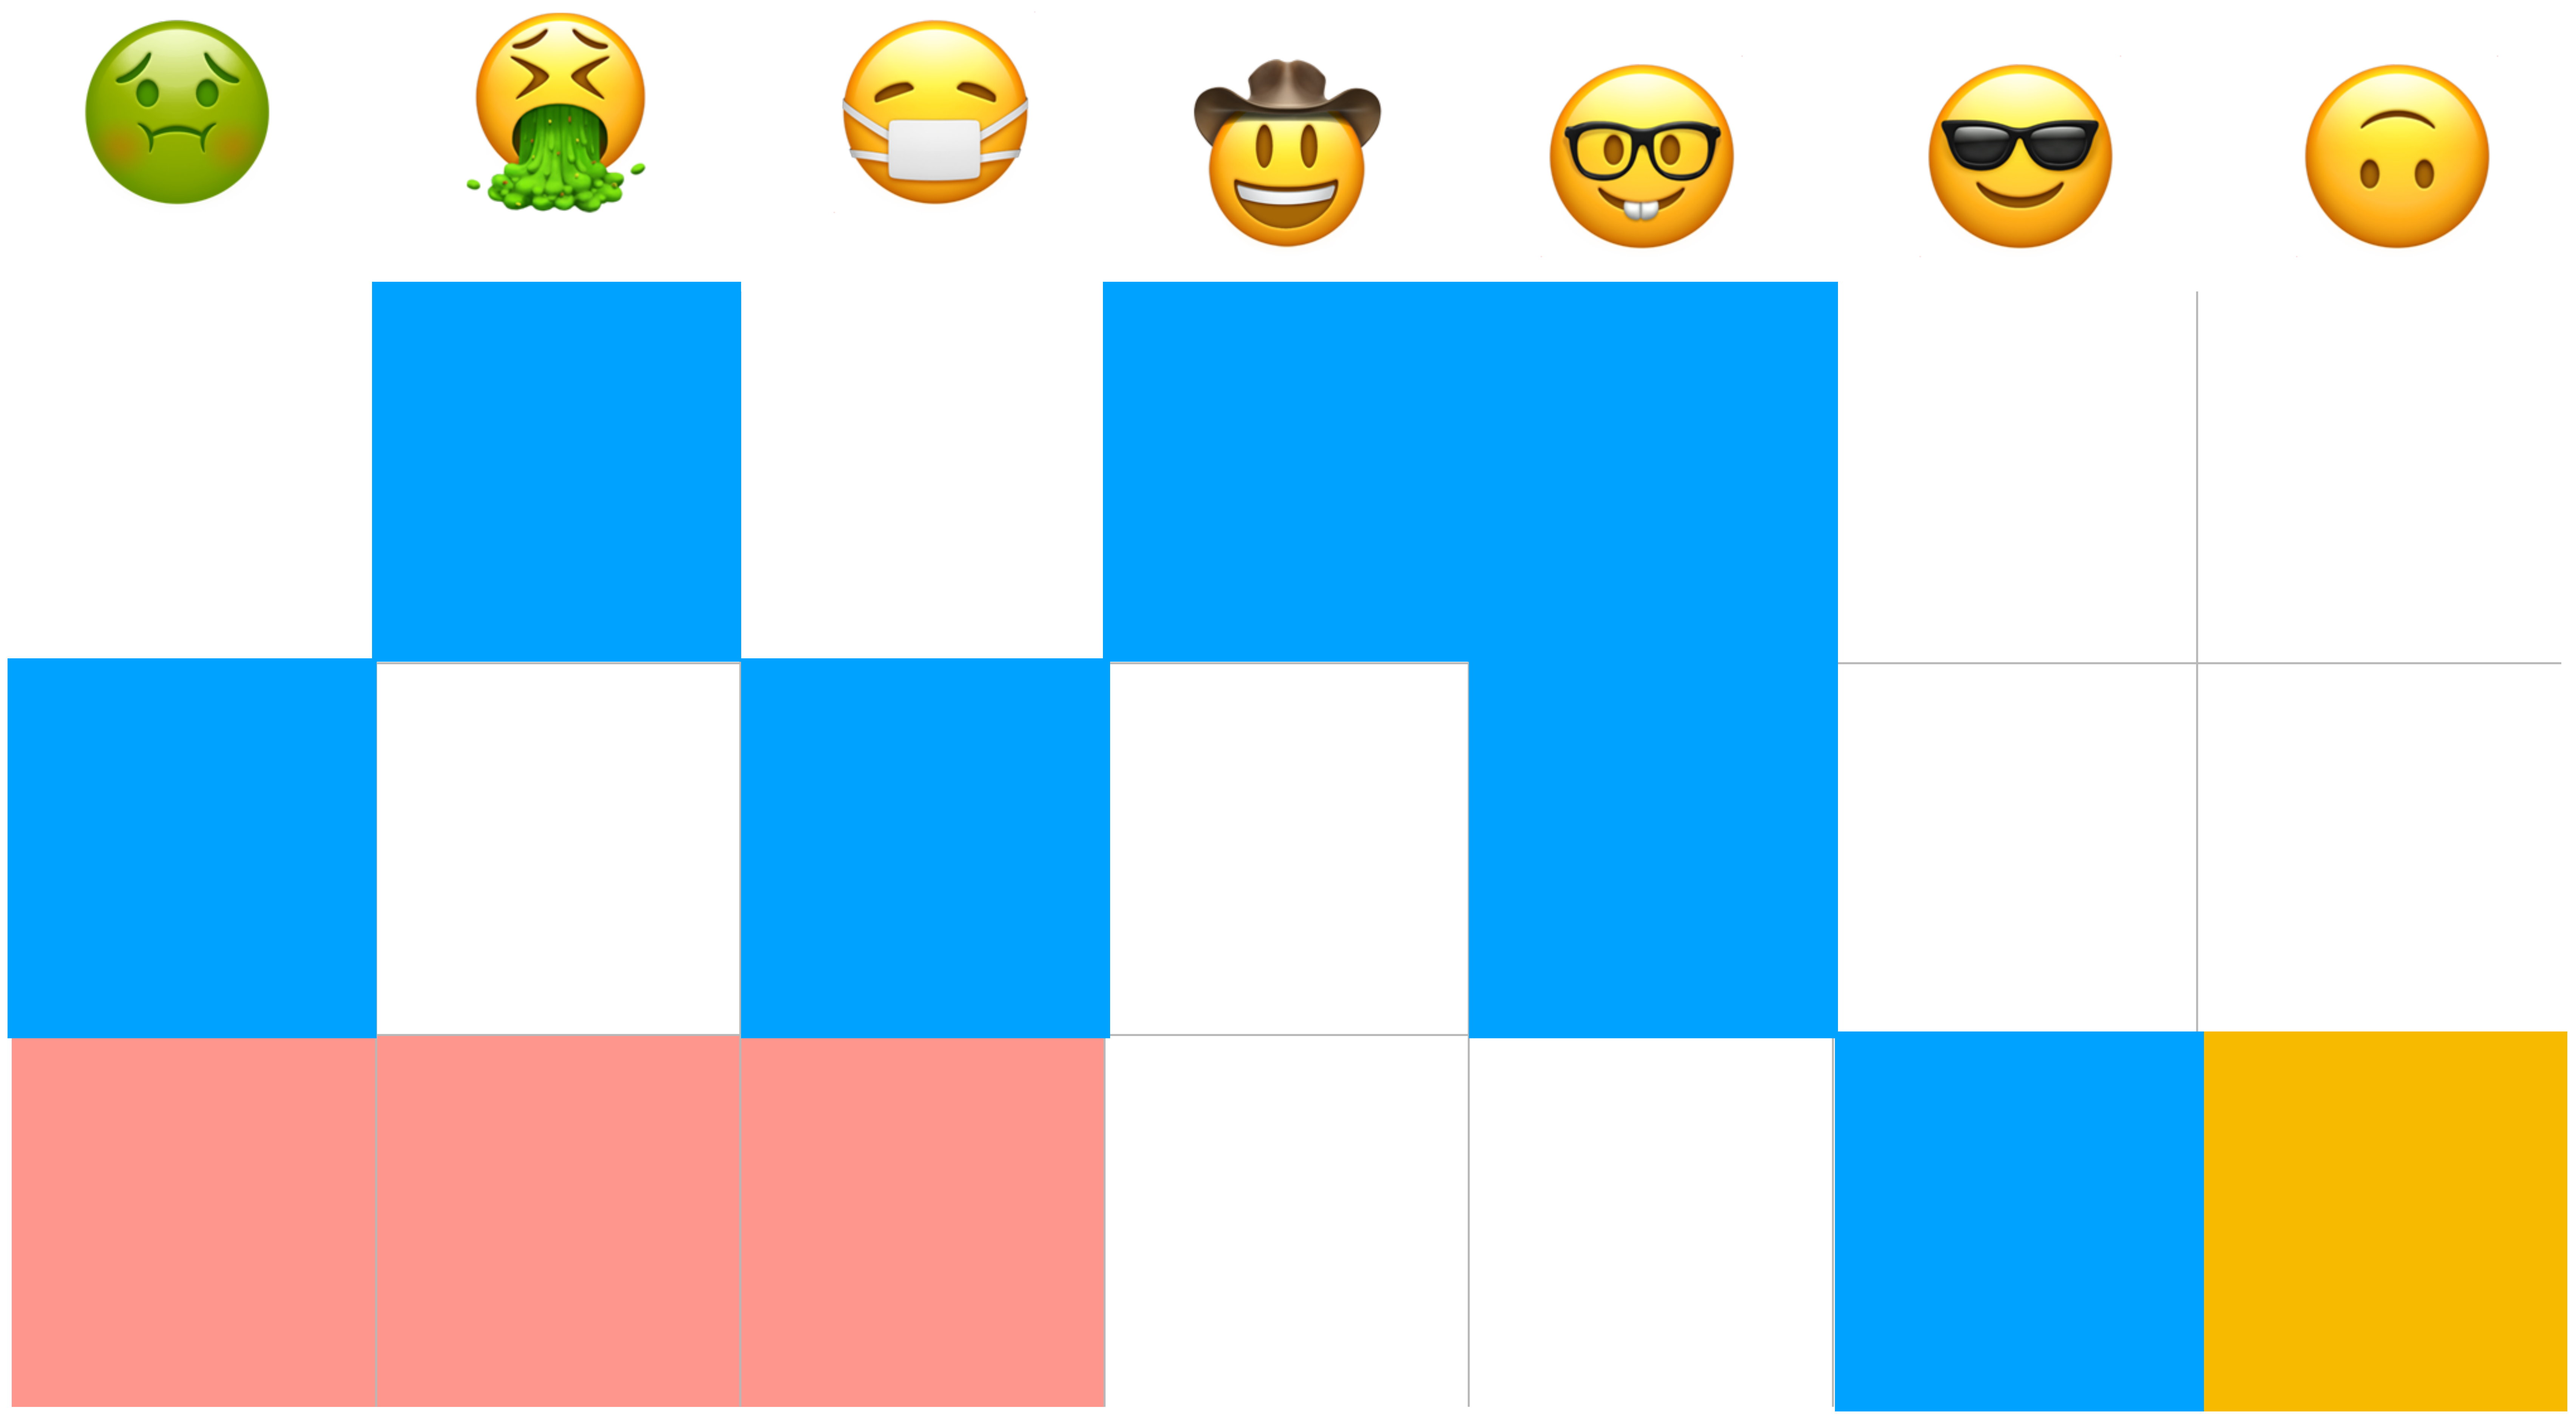

Works with probability 1/3s

...

Works with probability 1/3s

Repeating tests

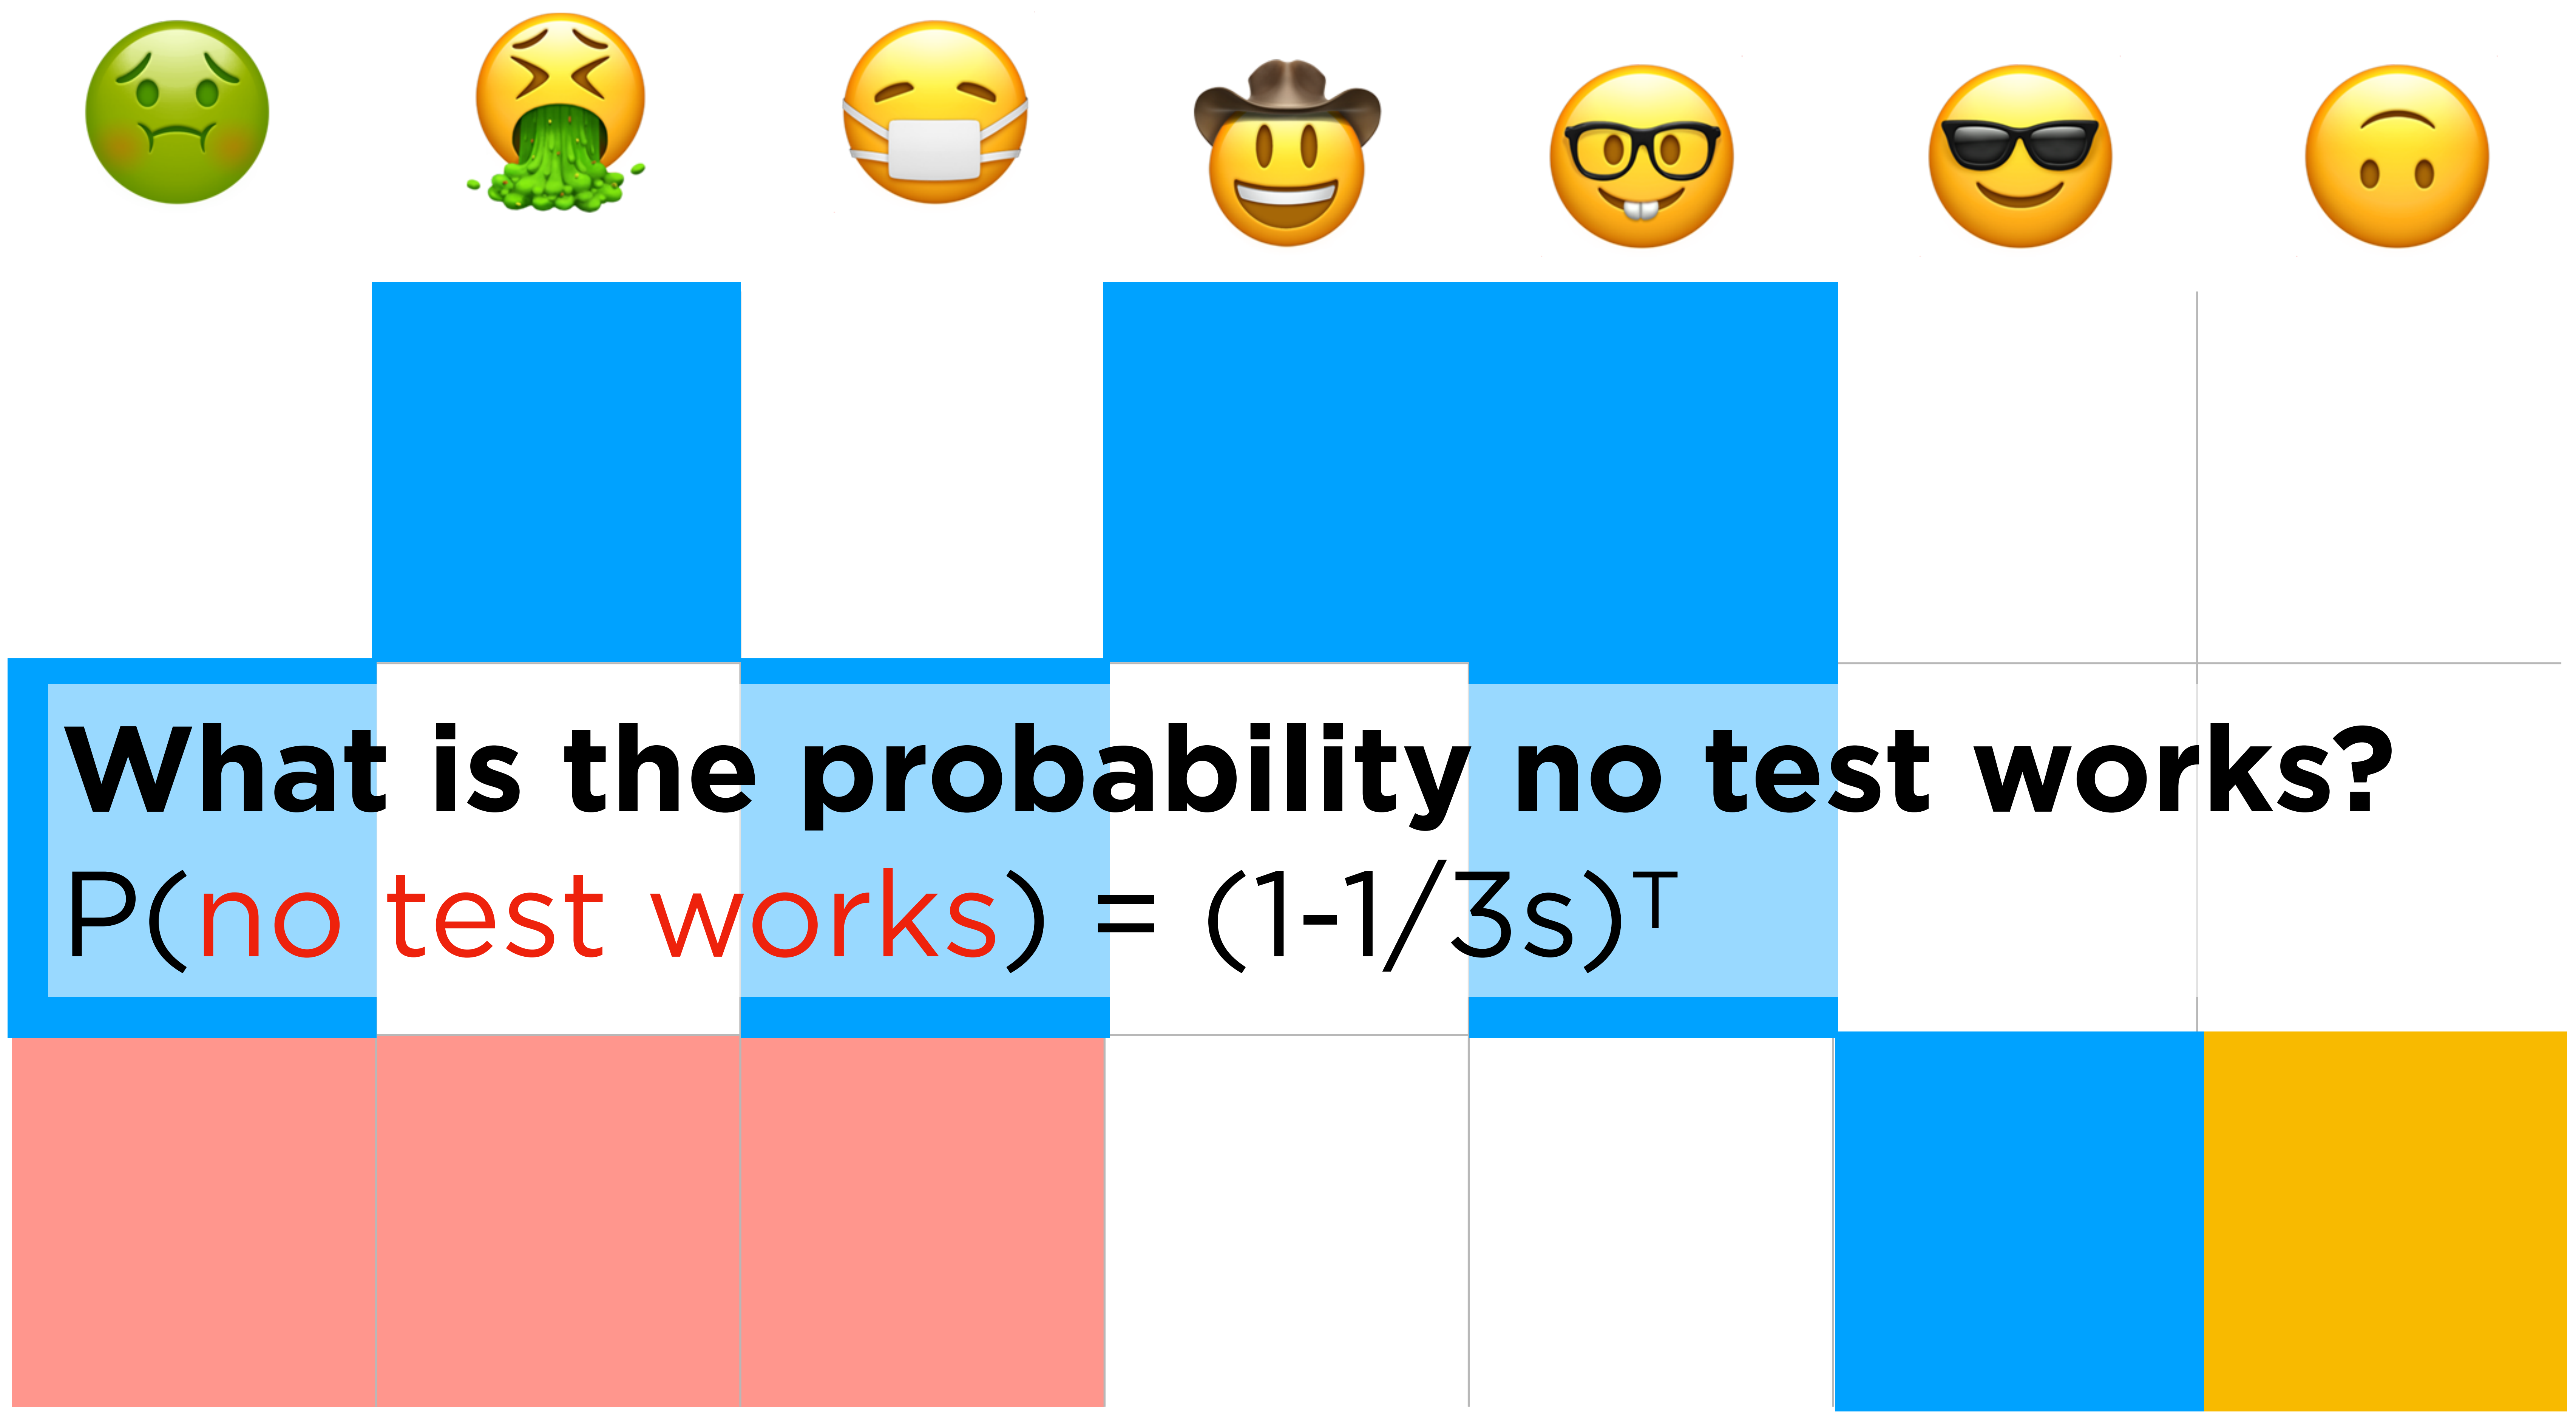

Works with probability 1/3s

...

Works with probability 1/3s

What is the probability no test works?

$P(\text{no test works}) = (1-1/3s)^T$

Repeating tests

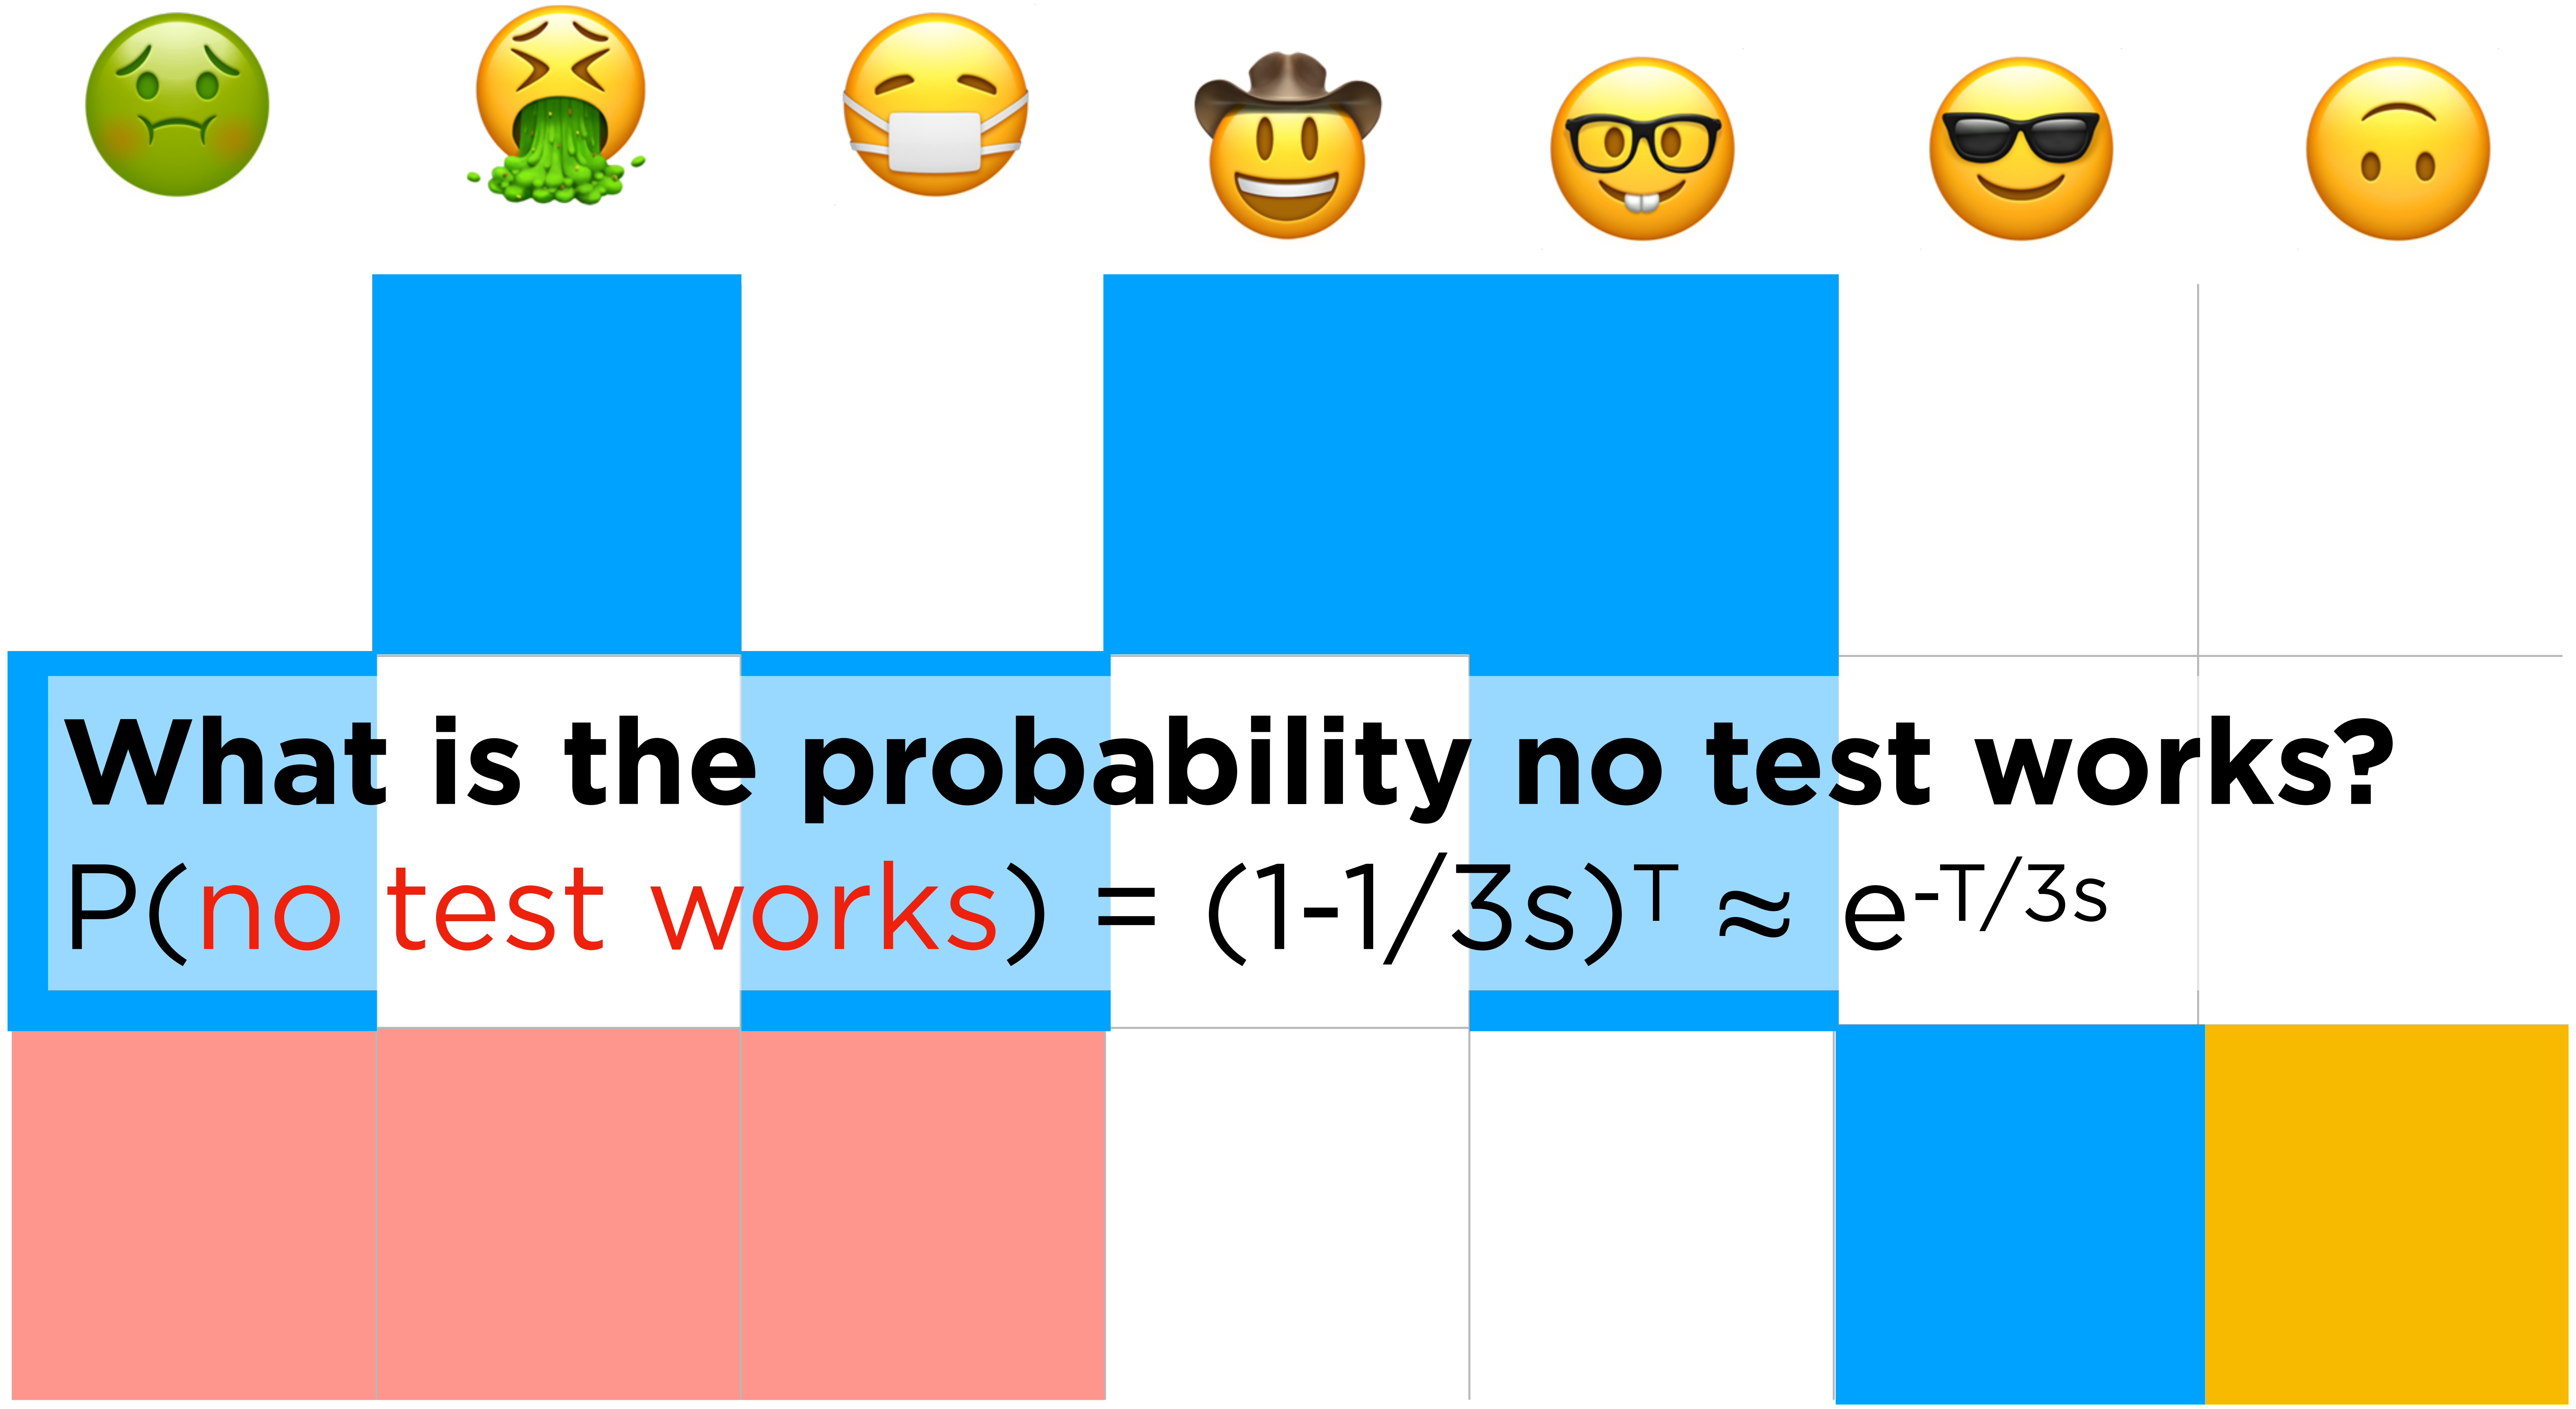

Works with probability 1/3s

...

Works with probability 1/3s

What is the probability no test works?

$P(\text{no test works}) = (1-1/3s)^T \approx e^{-T/3s}$

Repeating tests

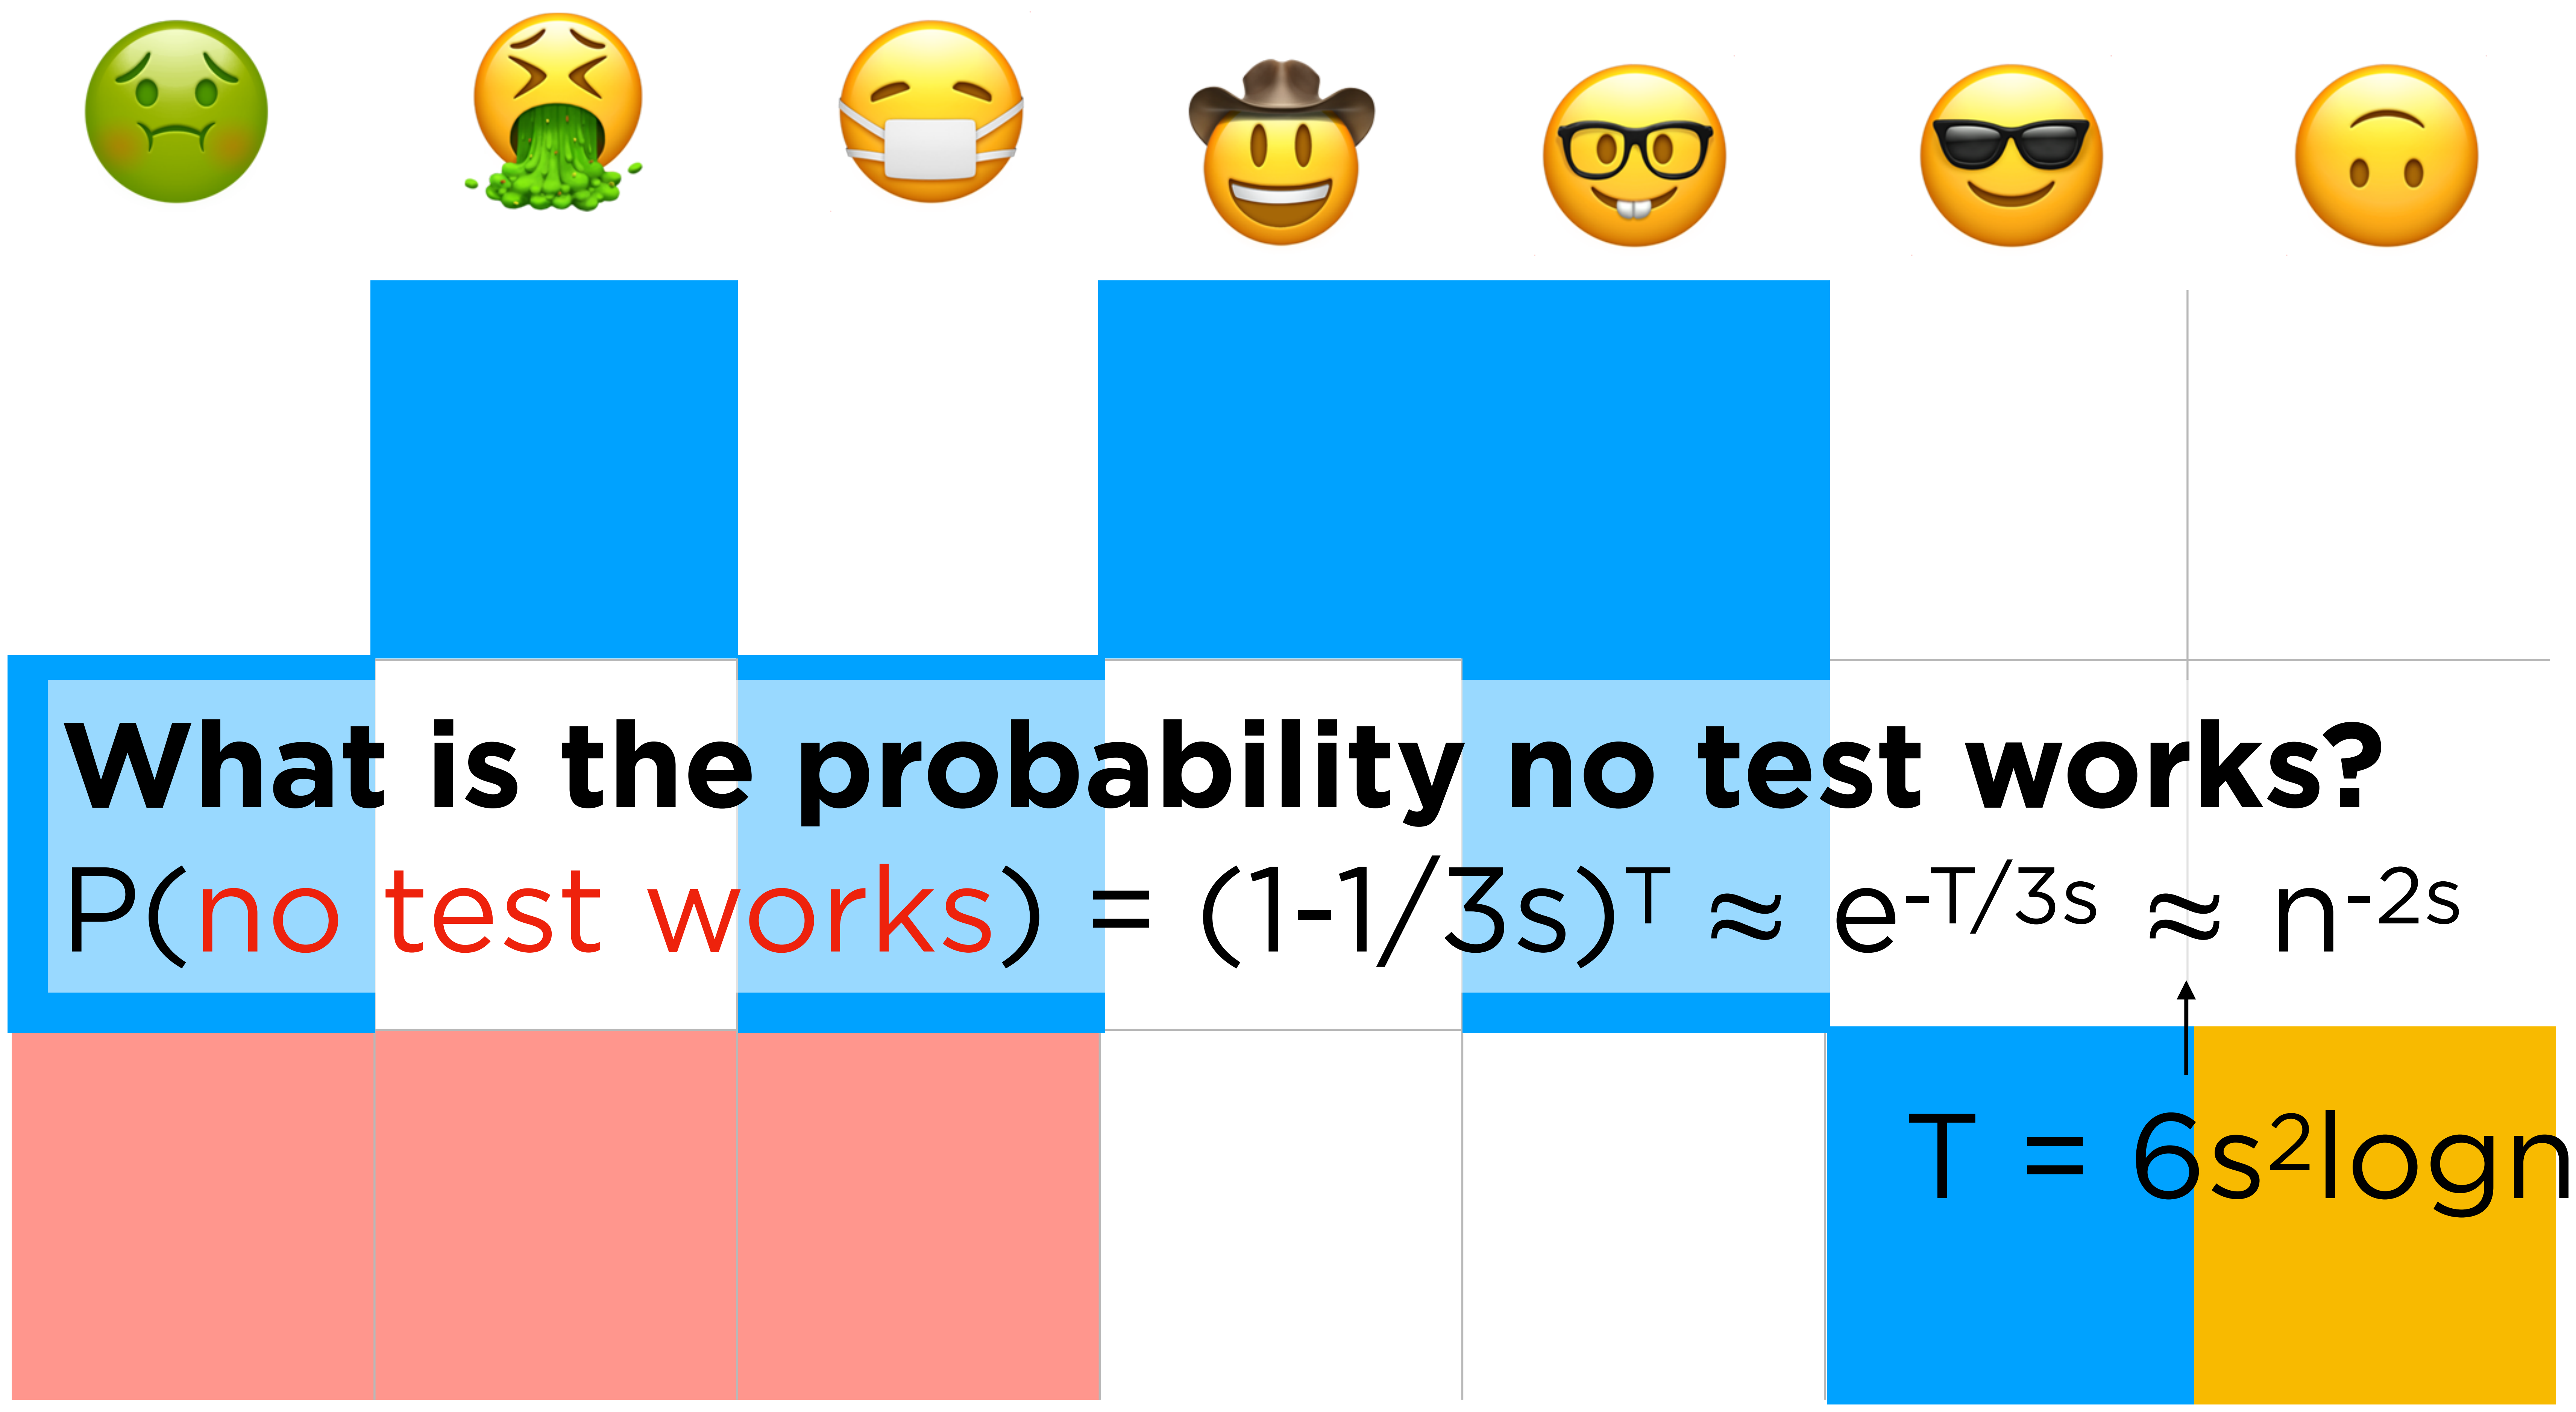

What is the probability no test works?

$P(\text{no test works}) = (1-1/3s)^T \approx e^{-T/3s} \approx n^{-2s}$

$T = 6s^2 \log n$

Works with probability 1/3s

...

Works with probability 1/3s

Union bound

🤢 🤮 😷 🤠 🤓 😎 🙃

We just saw

P(no test works **for** 🙃 **and** 🤢,🤮,😷) $\approx n^{-2s}$

Union bound

🤢 🤮 😷 🤠 🤓 😎 🙃

We just saw

P(no test works **for** 🙃 **and** 🤢,🤮,😷) ≈ n^{-2s}

But we haven’t dealt with 🤠,🤓,😎

Dorfman’s Construction

It works!

With good probability!

And very few tests!!

Why did this work?

Key components

1. Not too many things to find
2. Get information about many things in each test
3. Can do something random

Coin weighting Compressed Sensing Traitor Tracing Streaming Algorithms Johnson-Lindenstrauss Mastermind Network Tomography Finding wifi users IP Traceback Error correcting codes Multicast Message Authentication

Coin weighting Compressed Sensing Traitor Tracing Streaming Algorithms Johnson-Lindenstrauss Mastermind

Network Tomography

Finding wifi users IP Traceback Error correcting codes Multicast Message Authentication

Joint work with Mary Wootters

Statistical Science
2004, Vol. 19, No. 3, 499–517
DOI 10.1214/088342304000000422

Network Tomography: Recent Developments

Rui Castro, Mark Coates, Gang Liang, Robert Nowak and Bin Yu

Abstract. Today's Internet is a massive, distributed network which continues to explode in size as e-commerce and related activities grow. The heterogeneous and largely unregulated structure of the Internet renders tasks such as dynamic routing, optimized service provision, service level verification and detection of anomalous/malicious behavior extremely challenging. The problem is compounded by the fact that one cannot rely on the cooperation of individual servers and routers to aid in the collection of network traffic measurements vital for these tasks. In many ways, network monitoring and inference problems bear a strong resemblance to other "inverse problems" in which key aspects of a system are not directly observable. Familiar signal processing or statistical problems such as tomographic image reconstruction and phylogenetic tree identification have interesting connections to those

Tomography from Ro
Measurements

g, Mehmet Yildiz, Paul Barford, Ro

ristics is important
networks. Network
easuring character-
-internal links us-
rk has established
ographi
However
twork t
ding an
paths
paper, w
ased on
nates th
rs. Our
and SYN
he share

monitoring; C.4[Perfo
niques
General Terms: M
tation
Keywords: Network
measurement and mon

Network Loss Inference Using Unicast End-to-End Measurement

Mark Coates and Robert Nowak *
Department of Electrical and Computer Engineering, Rice University
6100 South Main Street, Houston, TX 77005–1892
Email: {mcoates, nowak}@ece.rice.edu, Web: www.dsp.rice.edu

Abstract

The fundamental objective of this work is to determine the extent to which unicast, end-to-end network measurement is capable of determining internal network losses. The major contributions of this paper are two-fold: we formulate a measurement procedure for network loss inference based on end-to-end packet pair measurements, and we develop a statistical modeling and computation framework for inferring internal network loss characteristics. Simulation experiments demonstrate the potential of our new frame-

is easily carried out on most networks and is scalable. Our approach employs unicast, end-to-end measurement of single packet and back-to-back packet pair losses, which can be performed actively or passively. By back-to-back packet pairs we mean two packets that are sent one after the other by the source, possibly destined for different receivers, but sharing a common set of links in their paths. Throughout the remainder of the paper we work with "success" probabilities (probability of non-loss) instead of loss probabilities. This provides a more convenient mathematical parameterization

Network Tomography on General Topologies *

Tian Bu
Department of Computer Science
University of Massachusetts
Amherst, MA 01003, USA
tbu@cs.umass.edu

Nick Duffield
AT&T Labs-Research
180 Park Avenue
Florham Park, NJ 07932, USA
duffield@research.att.com

Francesco Lo Presti
Dipartimento di Informatica
Università dell'Aquila
Via Vetoio, Coppito (AQ), Italy
lopresti@univaq.it

ABSTRACT

In this paper we consider the problem of inferring link-level loss rates from end-to-end multicast measurements taken from a collection of trees. We give conditions under which loss rates are identifiable on a specified set of links. Two algorithms are presented to perform the link-level inferences for those links on which losses can be identified. One, the *minimum variance weighted average (MVWA) algorithm* treats the trees separately and then averages the results. The second, based on *expectation-maximization (EM)* merges all of the measurements into one computation. Simulations show that EM is slightly more accurate than MVWA, most likely due to its more efficient use of the measurements. We also describe extensions to the inference of link-level delay, inference from end-to-end unicast measurements, and inference when some measurements are missing.

1. INTRODUCTION

As the Internet grows in size and diversity, its internal behavior becomes ever more difficult to characterize. Any one organization has administrative access to only a small fraction of the network's internal nodes, whereas commercial factors often prevent organizations

all links
ure 1(a)
sources
terest is
are need
ure 1(b)
of a ne
measur
from m
problem

(a) Ne

Network Delay Tomography

Yolanda Tsang, *Member, IEEE*, Mark Coates, *Member, IEEE*, and Robert D. Nowak, *Member, IEEE*

paper, we introduce a new methodology for network
y, specifically, estimating the probability distribu-
e queuing delay on each link based on end-to-end
acket pair measurements. Our approach employs
nd-to-end measurement of back-to-back packets.
o-back packets, we mean two packets that are sent
usly by the source, possibly destined for different
but sharing a common set of links in their paths. The
s should experience approximately the same on each
k in their path.

bution

nference methodologies focused on multicast routing.
st routing, packets are delivered from sender to the re-
one send operation. Along the path, probe packets are
as needed as the paths diverge [2], [6]. Although mul-
ods show promise for network performance inference,
niques are often impractical in real networks. Many

Network Tomography: Estimating Source-Destination Traffic Intensities From Link Data

Y. VARDI

The problem of estimating the node-to-node traffic intensity from repeated measurements of traffic on the links of a network is formulated and discussed under Poisson assumptions and two types of traffic-routing regimens: deterministic (a fixed known path between each directed pair of nodes) and Markovian (a random path between each directed pair of nodes, determined according to a known Markov chain fixed for that pair). Maximum likelihood estimation and related approximations are discussed, and computational difficulties are pointed out. A detailed methodology is presented for estimates based on the method of moments. The estimates are derived algorithmically, taking advantage of the fact that the first and second moment equations give rise to a linear inverse problem with positivity restrictions that can be approached by an EM algorithm, resulting in a particularly simple solution to a hard problem. A small simulation study is carried out.

KEY WORDS: Communication network; Computer network; EM algorithm; Linear inverse problems; Maximum-likelihood; Moment method; Origin-destination tables; Poisson traffic; Positivity constraints; Trip matrix.

A network

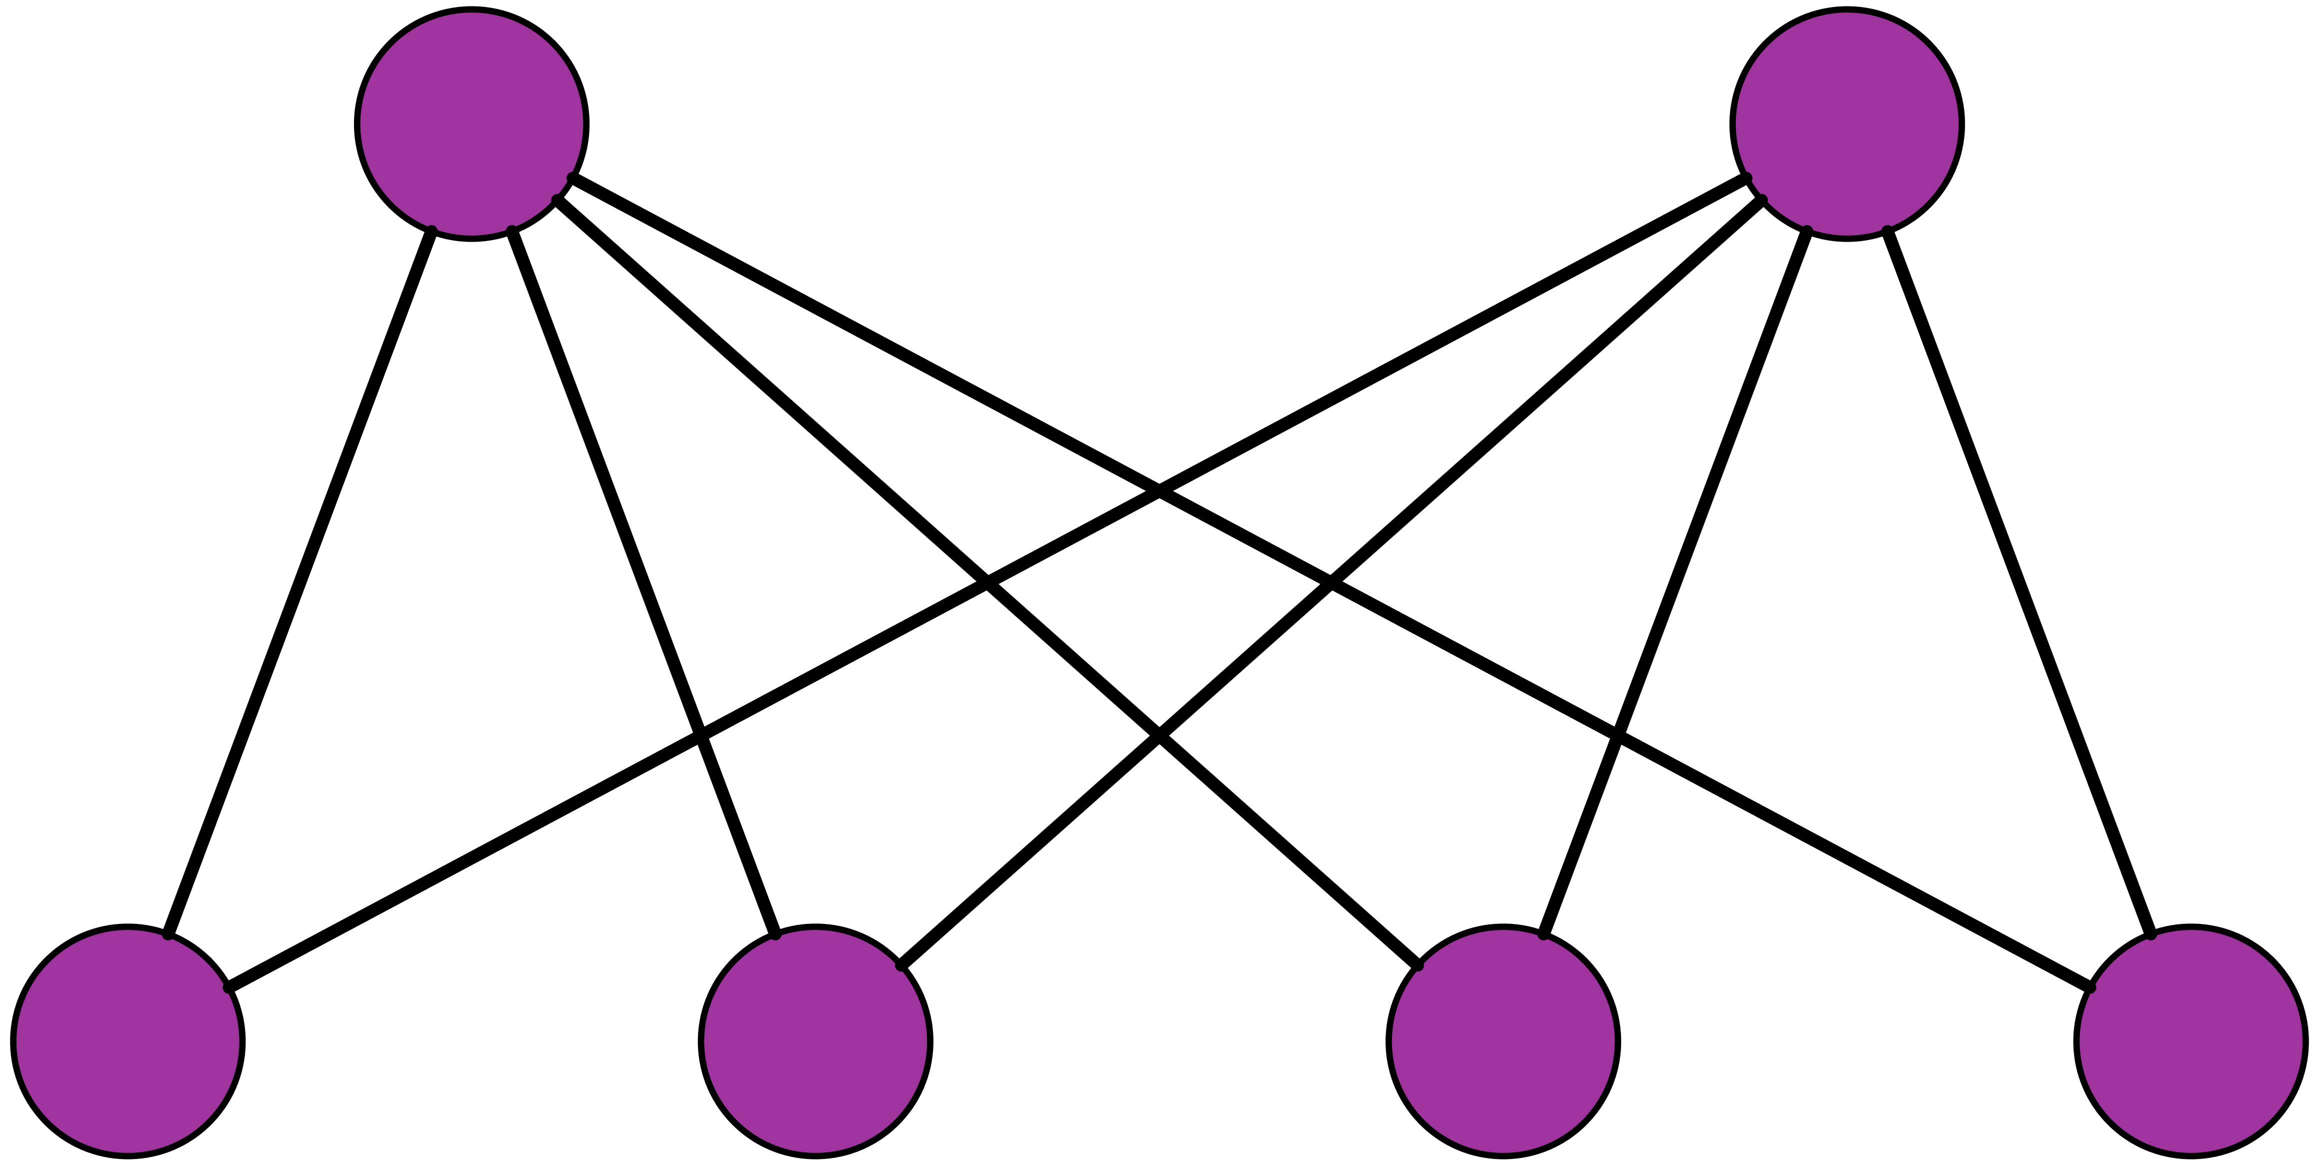

A network, failing

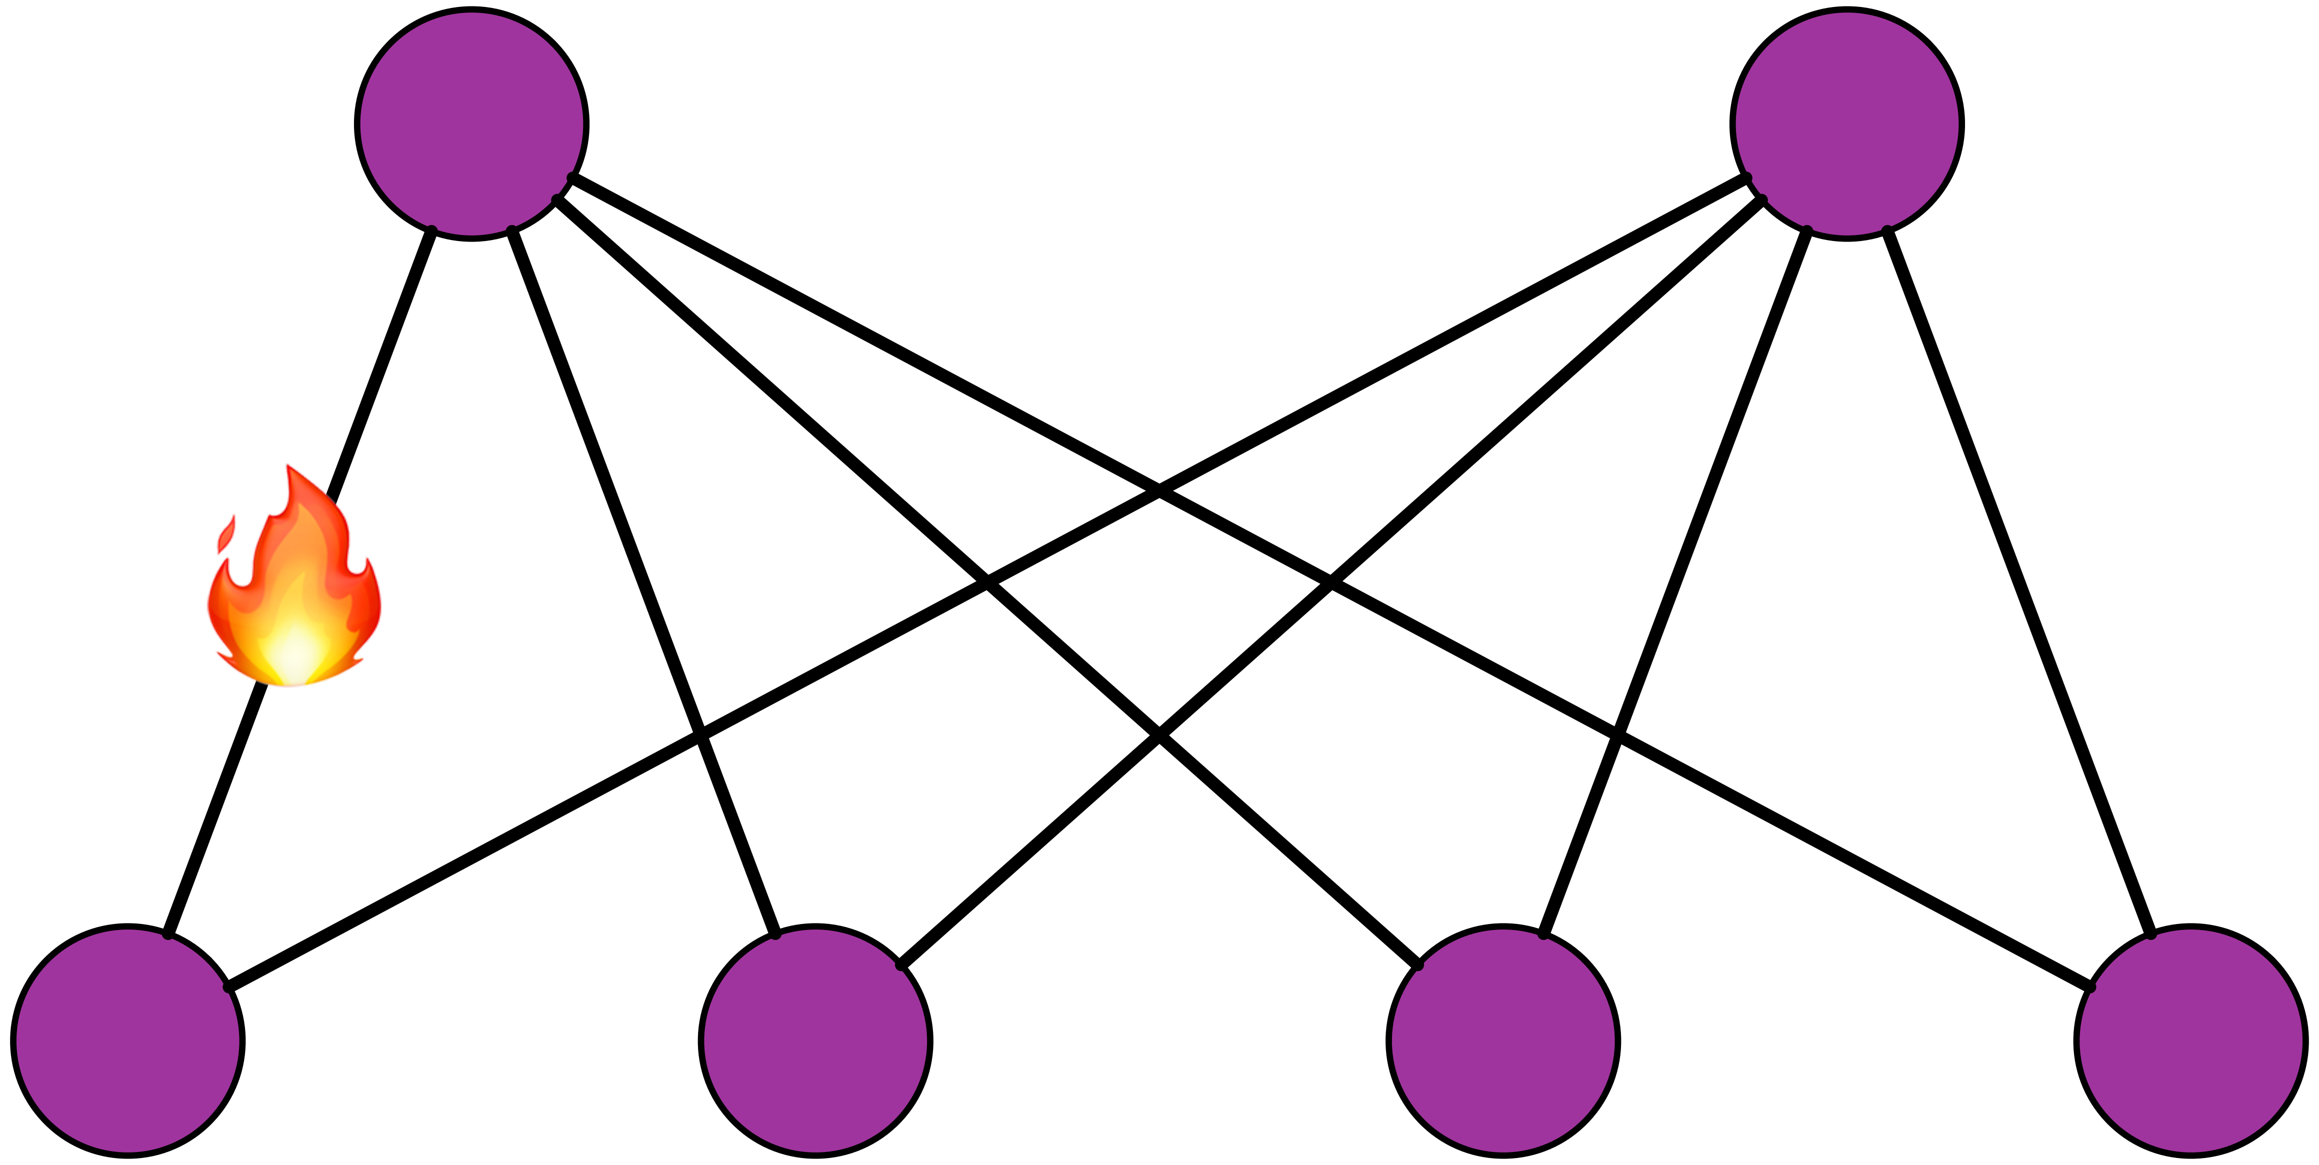

Finding failures

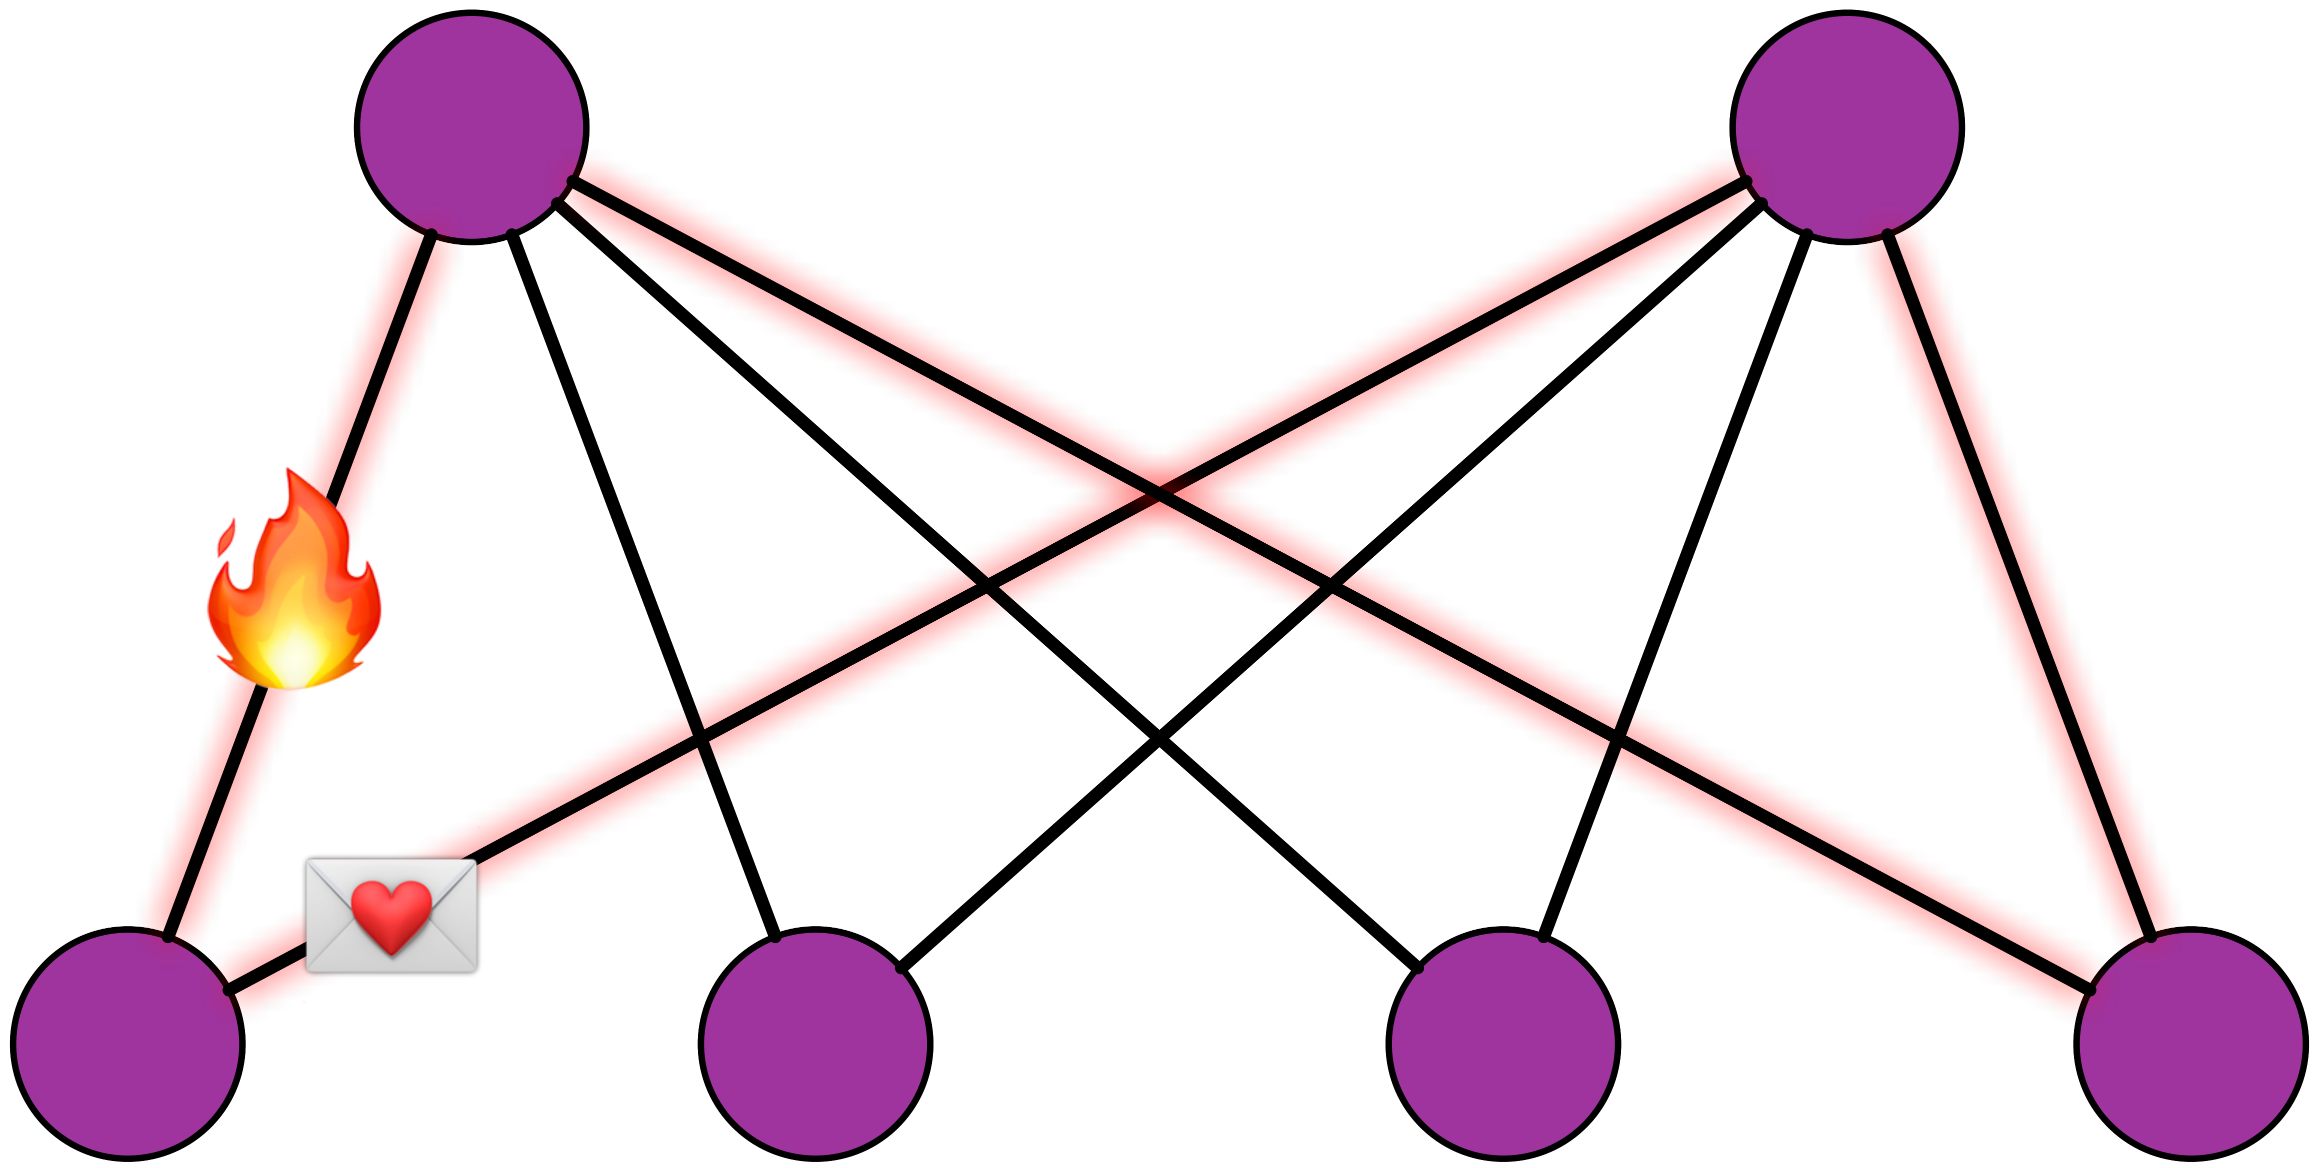

Finding failures

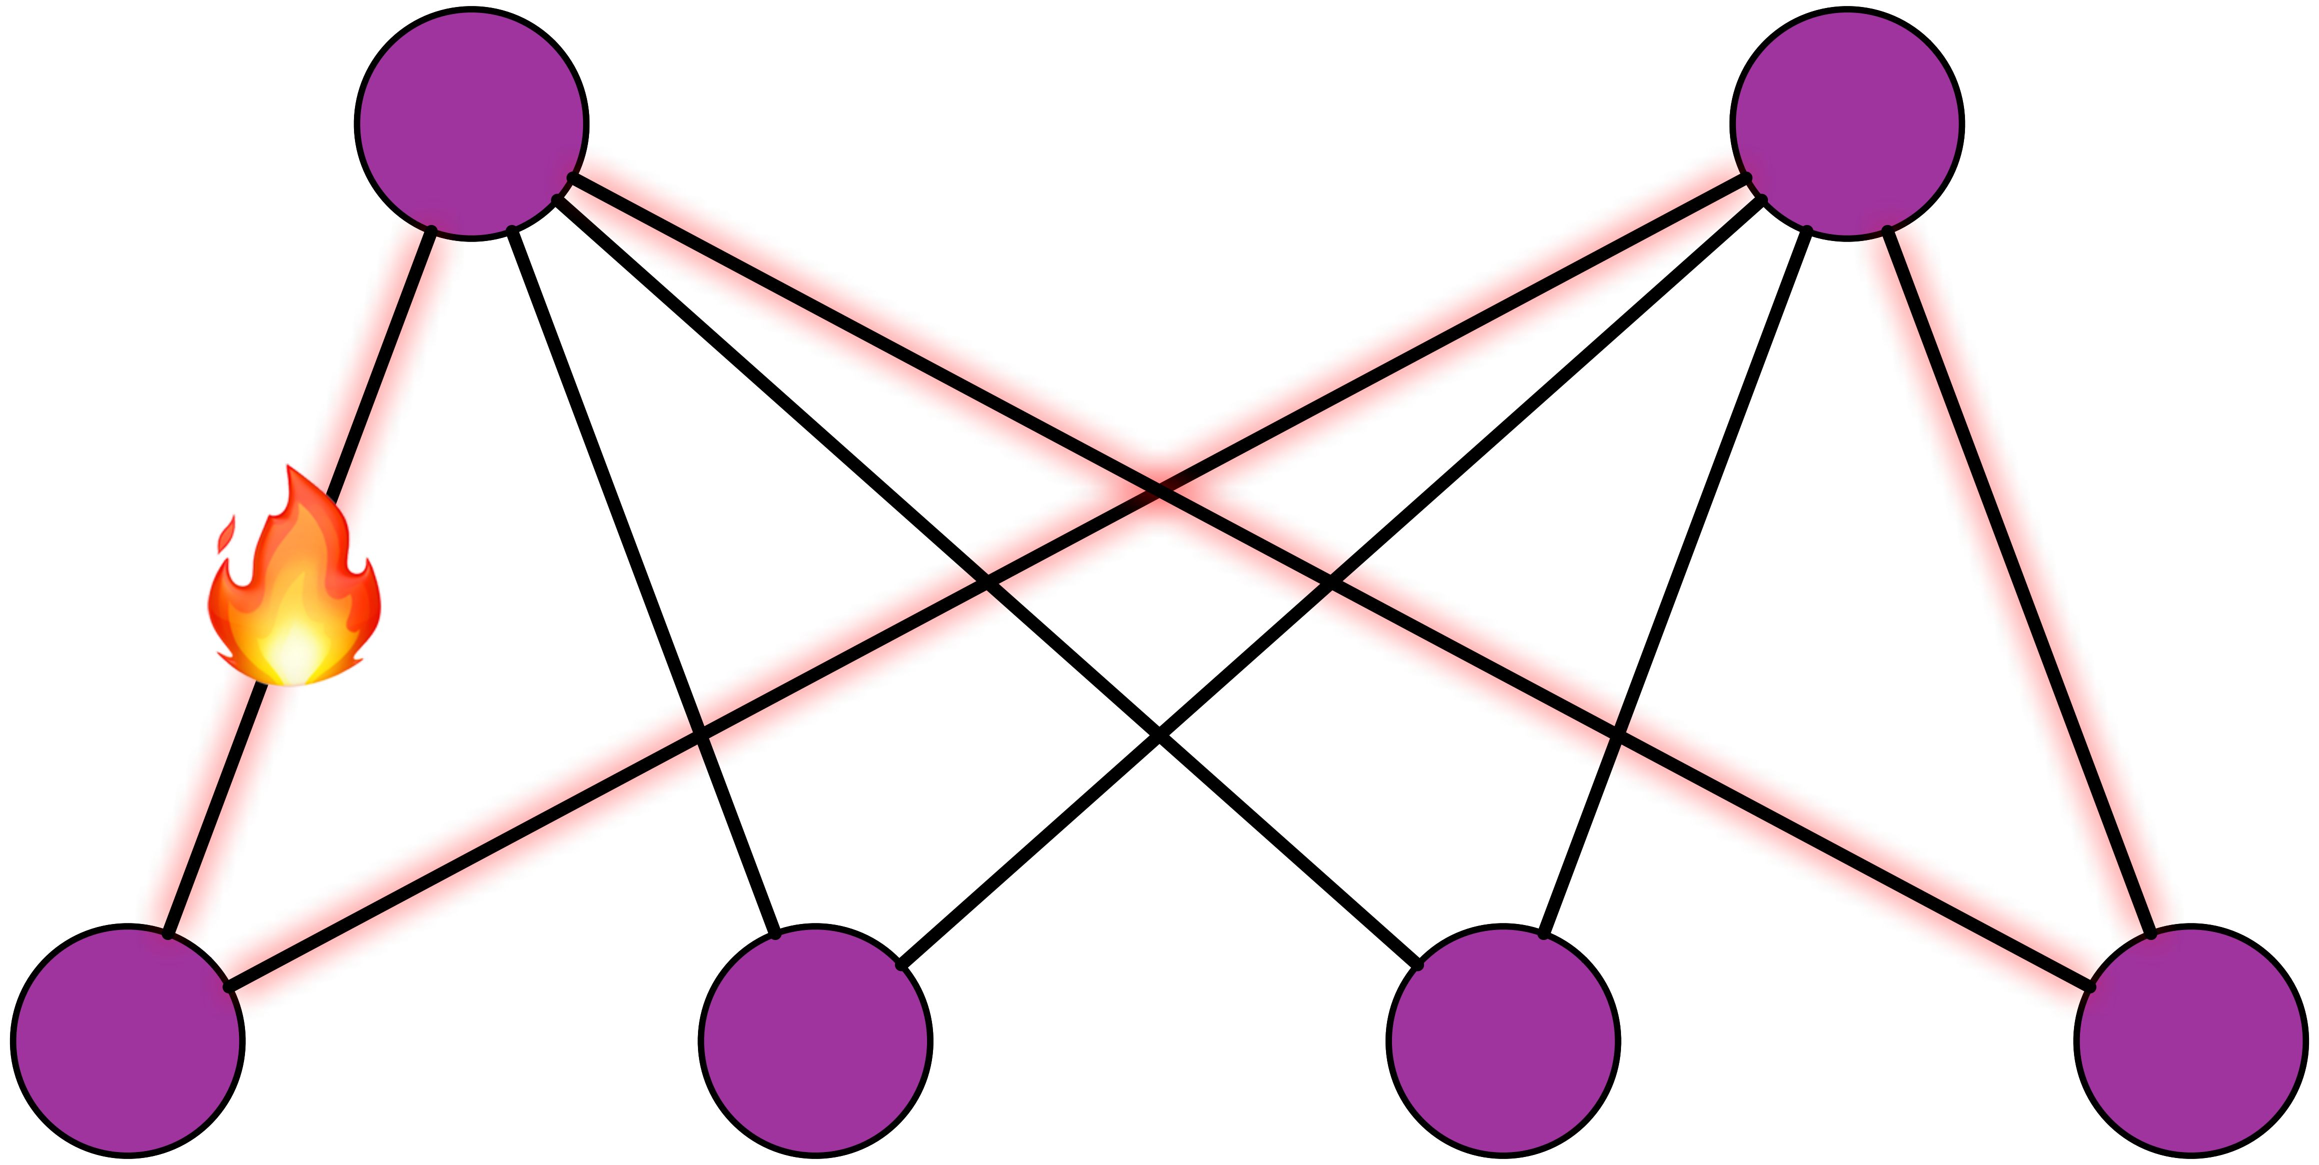

Finding failures

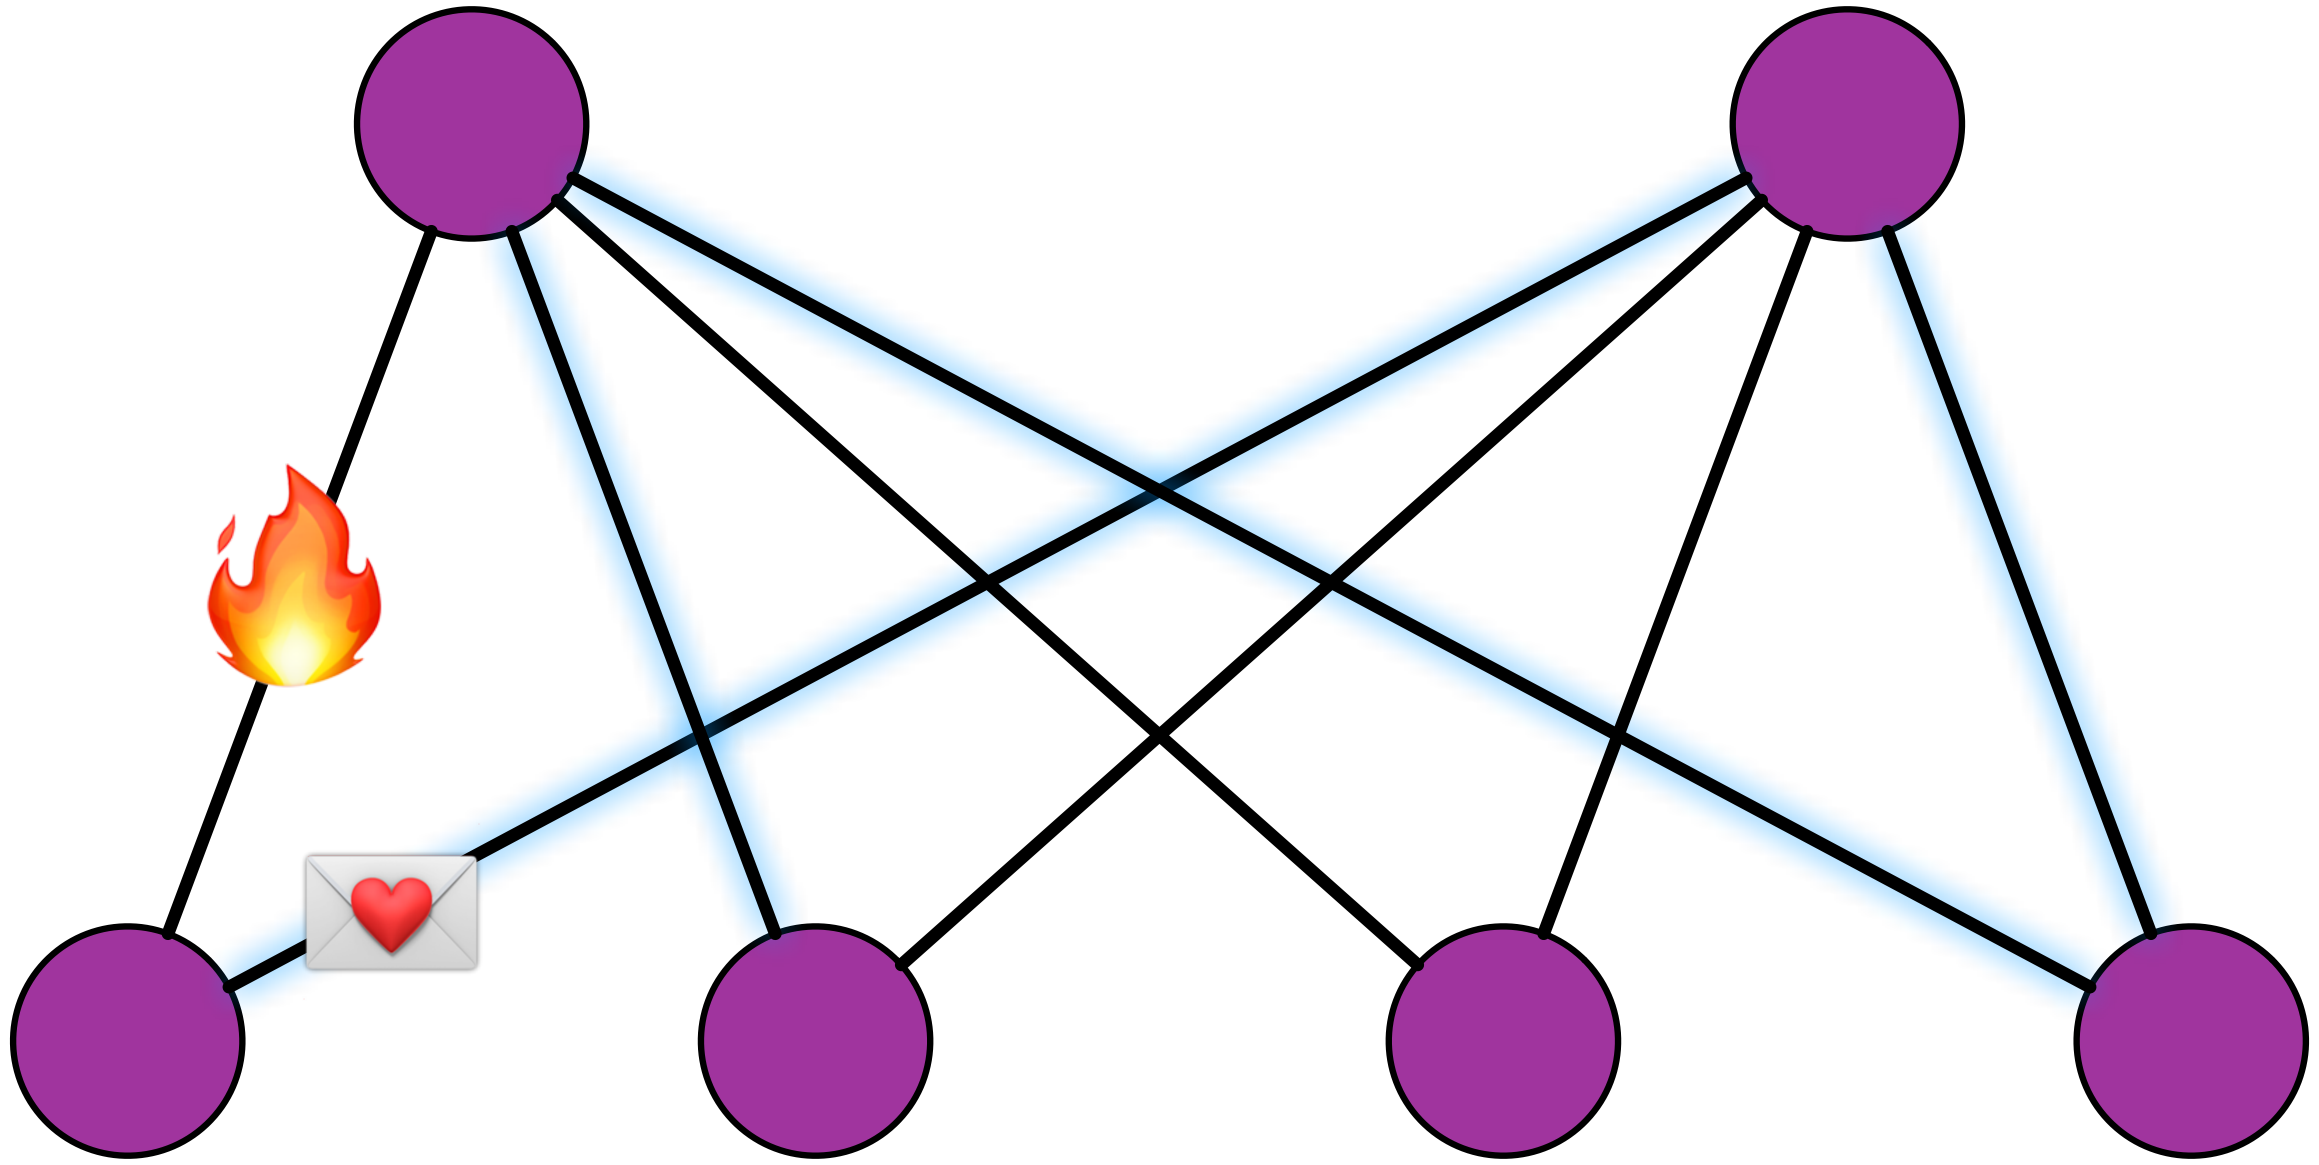

Finding failures

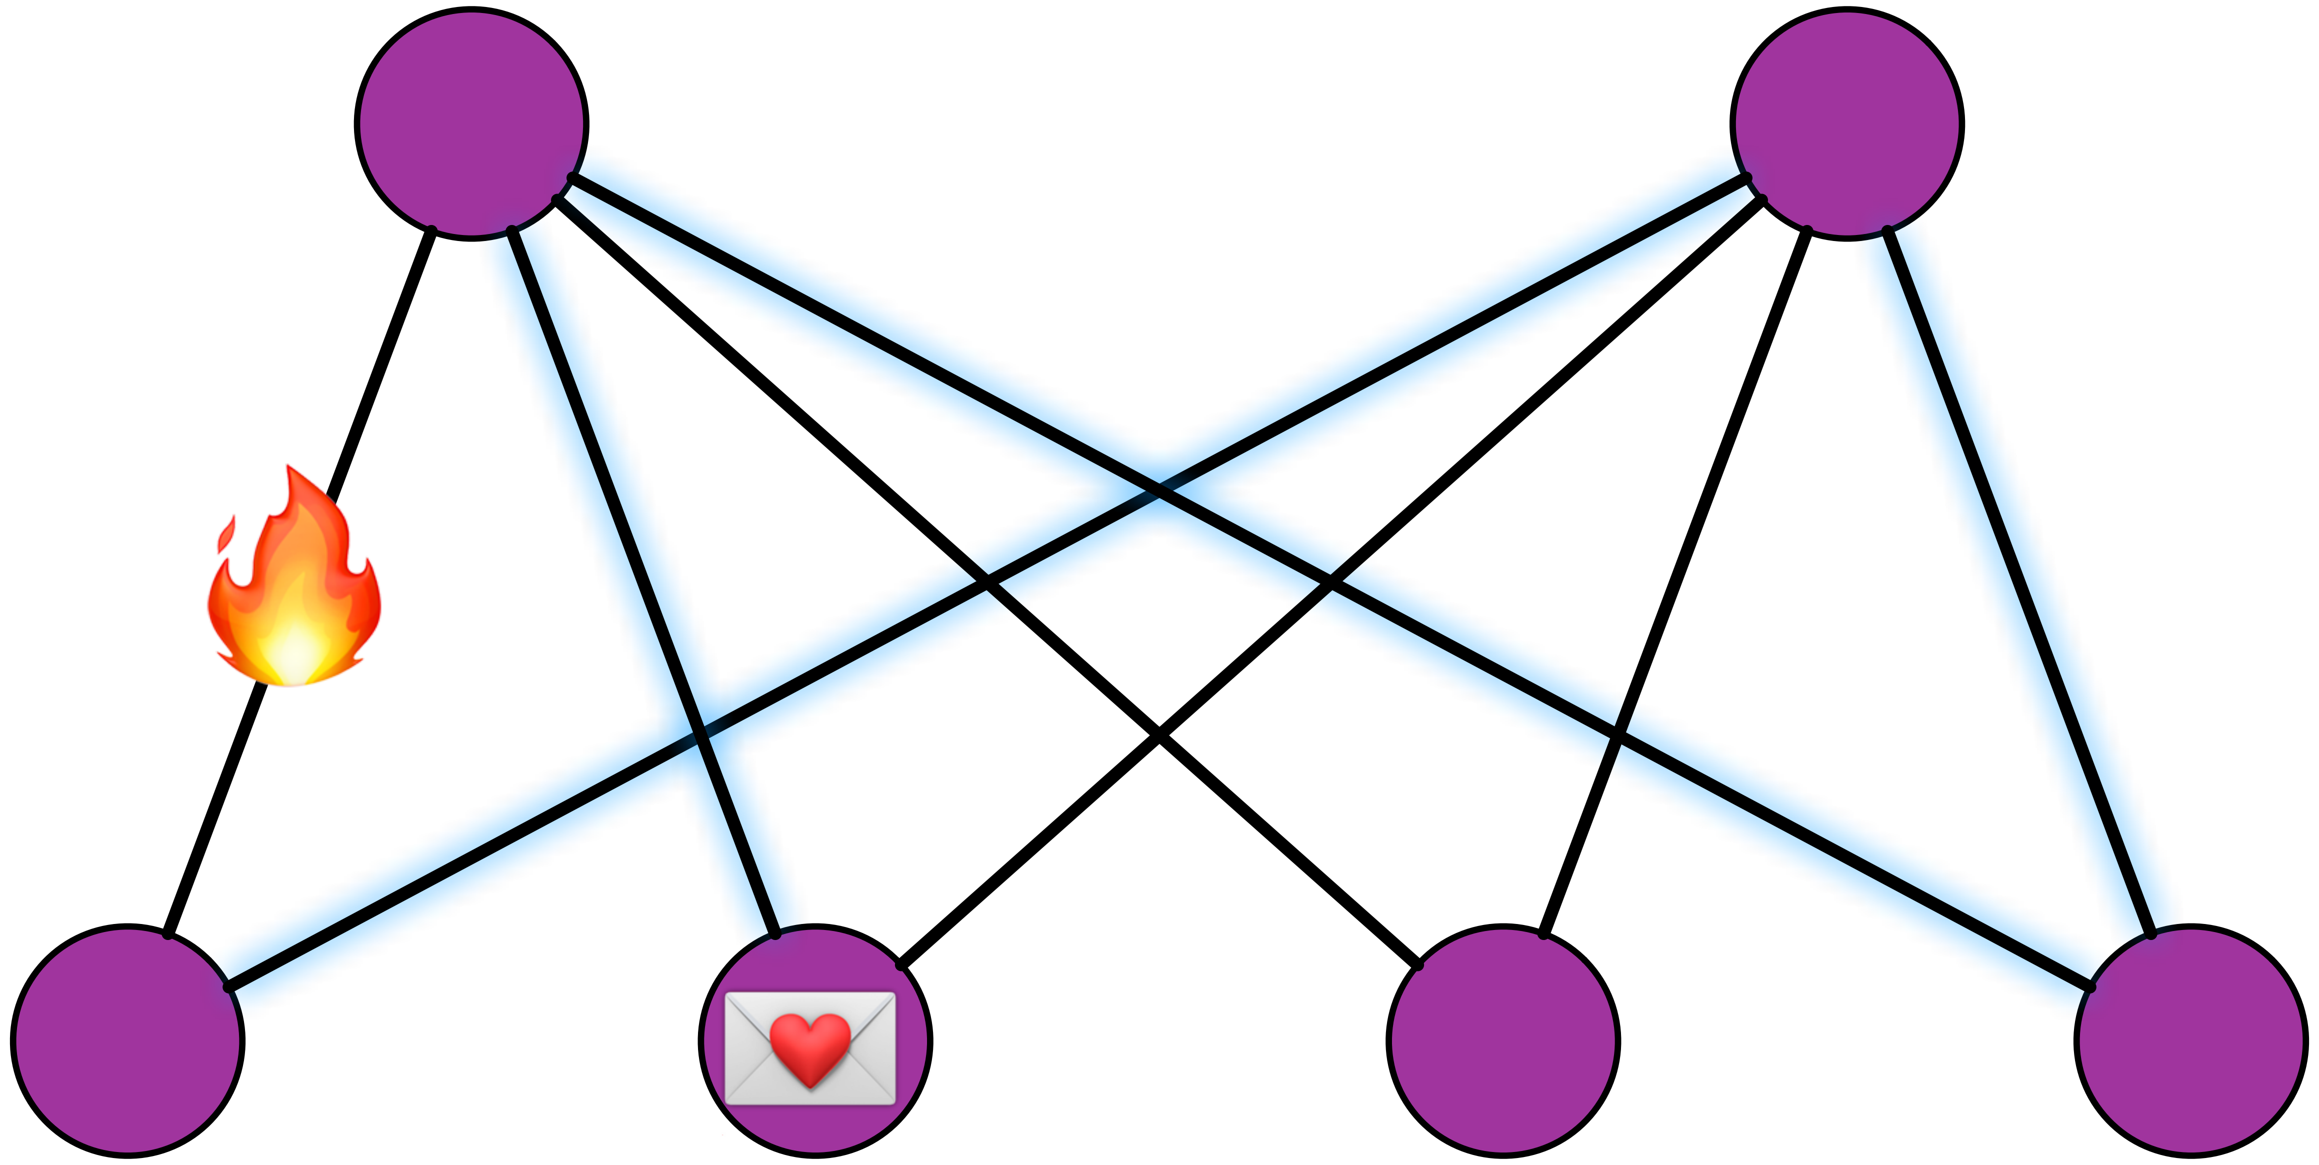

Tomography problem

We have a graph $G=(V,E)$ with n edges, at most s edges are sick.

Definition: A *graph-constrained test* returns whether any edges in a *connected subset* of edges are sick or not.

Problem: Construct a set of graph-constrained tests which can identify any set of at most s sick edges.

This seems tricky

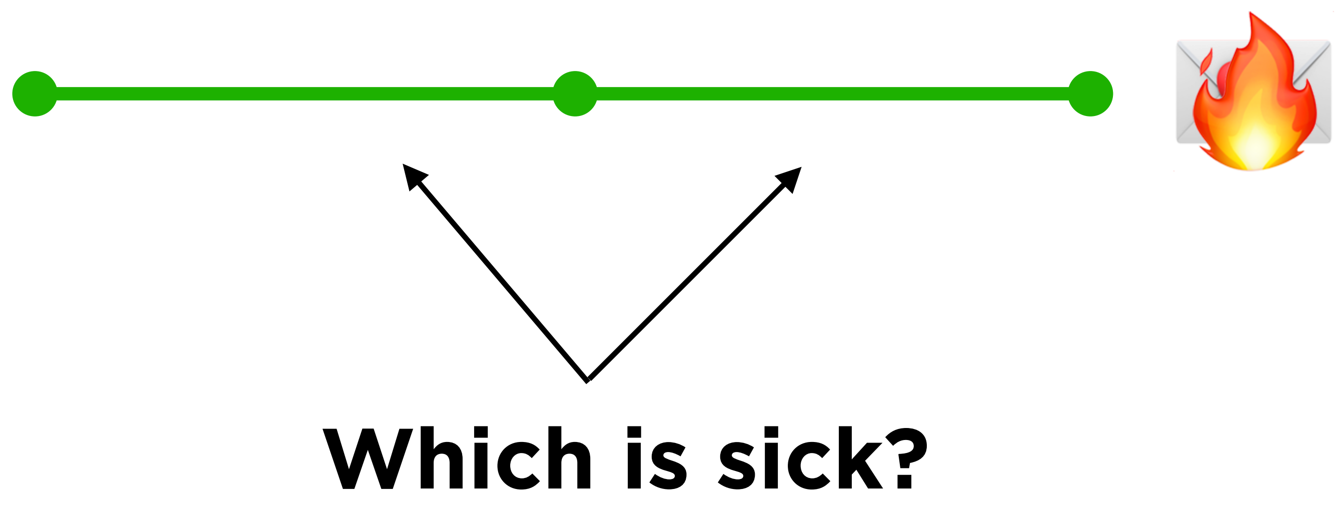

This seems tricky

Theorem [Harvey et al 2007]: For the line graph on *n* nodes, about *n*/2 tests required

Proof: Each neighboring pair of edges must be separated by some test. Each test is a path and can only separate two pairs. There are about n pairs.

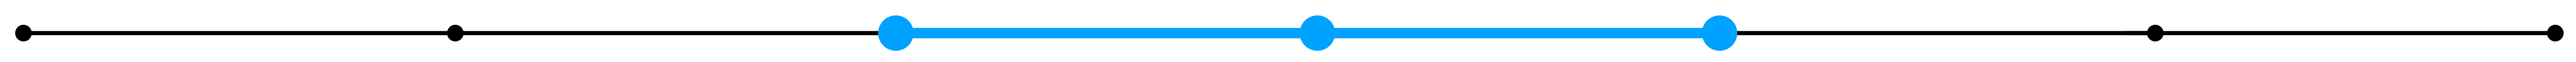

Key components?

1. Not too many things to find

2. Get information about many things in each test

3. Can do something random

Key components?

👍 Not too many things to find

2. Get information about many things in each test

3. Can do something random

Key components?

👍🏼 Not too many things to find

👍🏼 Get information about many things in each test

3. Can do something random

Key components?

- 👍 Not too many things to find
- 👍 Get information about many things in each test
- 👍 Can do something random

Our informal result

If a graph is *sufficiently well-enough connected*, we can find any set of *s* sick edges using **$O(s^2 \log n)$** tests

Our informal result

If a graph is *sufficiently well-enough connected*, we can find any set of *s* sick edges using
$O(s^2 \log n)$ tests

Same as group testing

Algorithm

For 1…$2s^2 \log n$:

- Include each edge with probability $p \sim 1/s$
- Use large connected components as tests

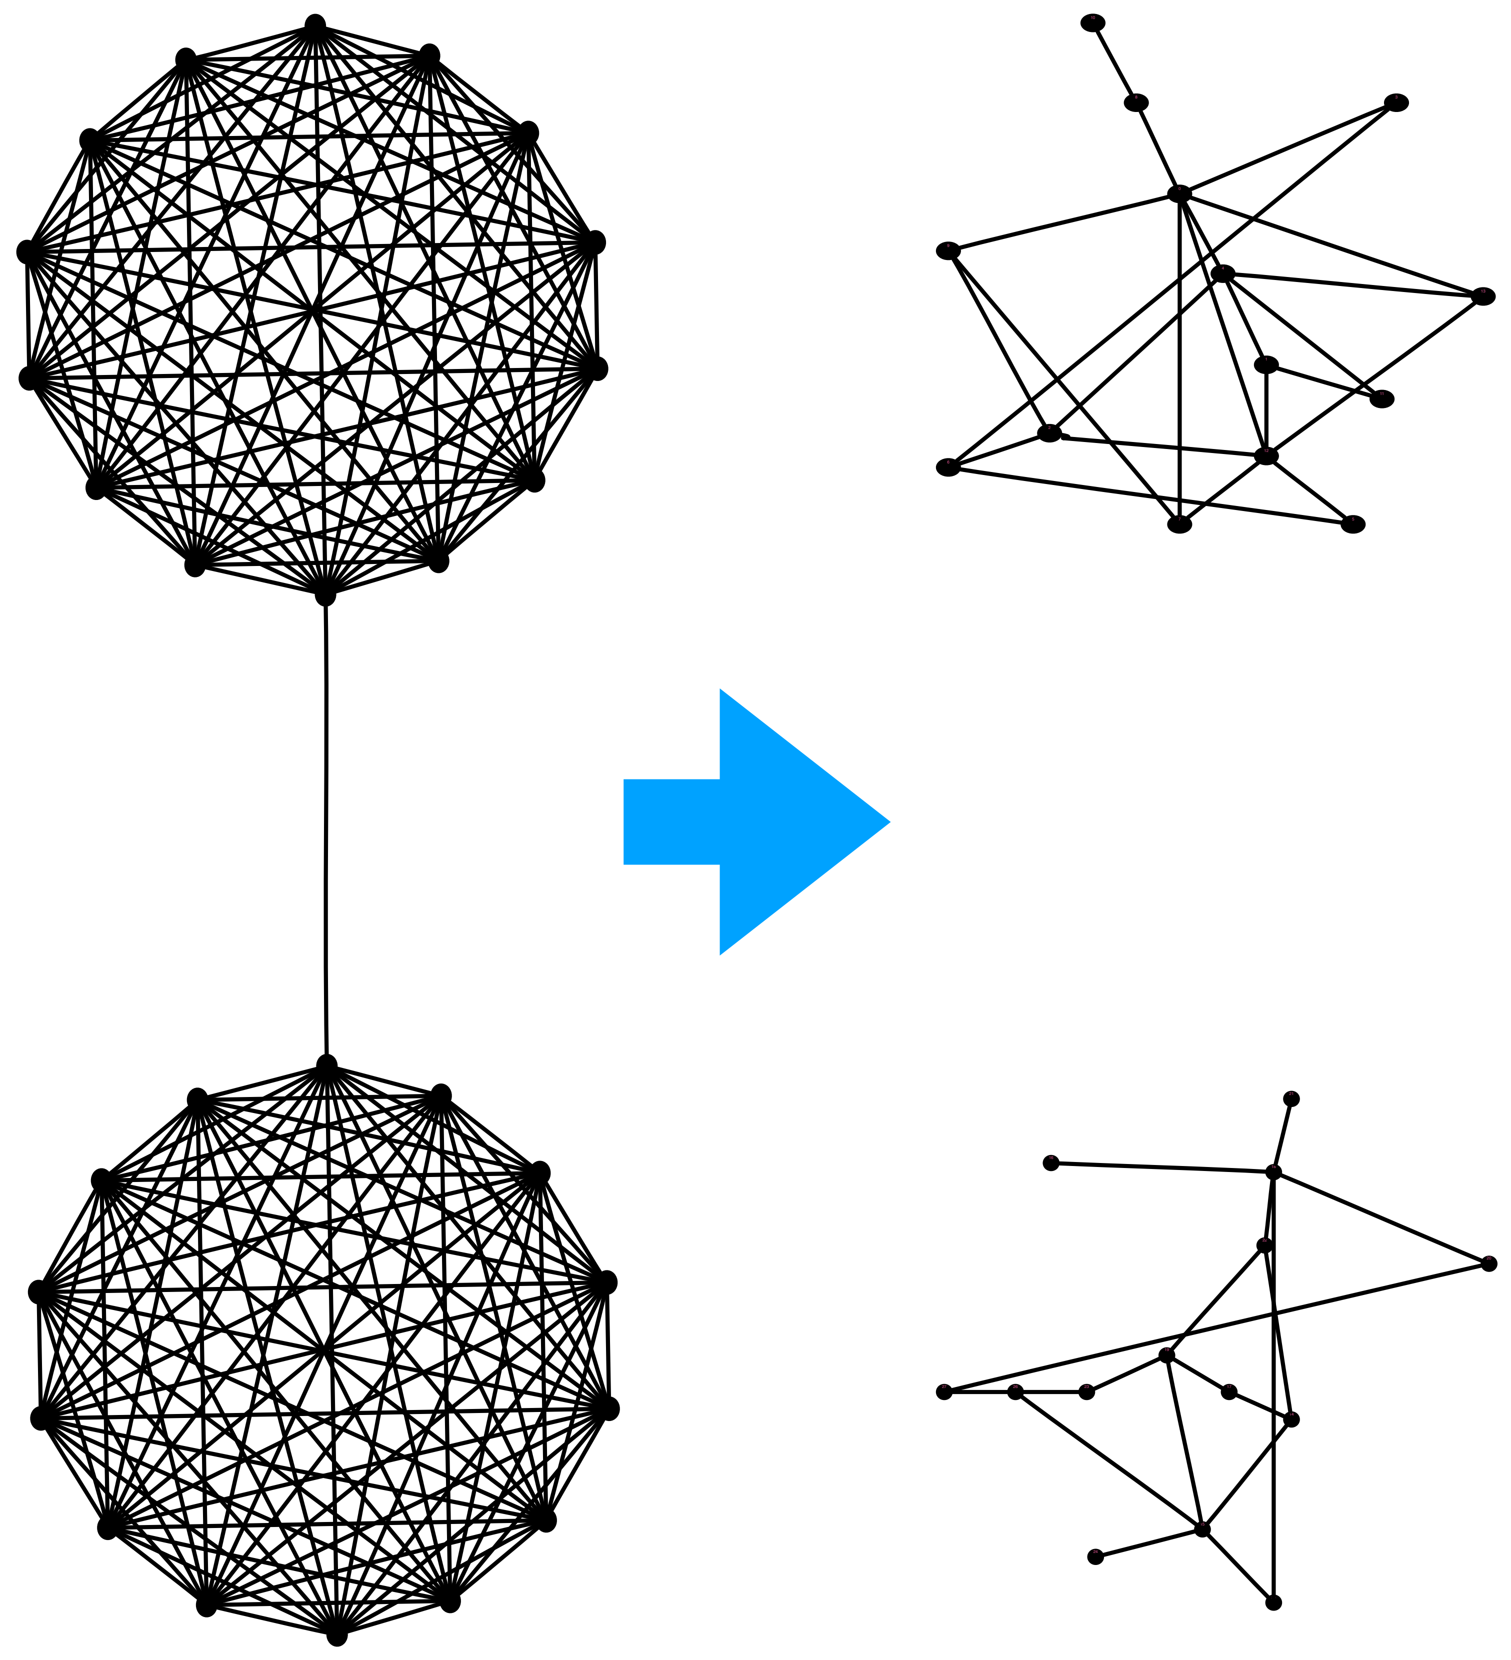

Example: K_4 (6 edges)

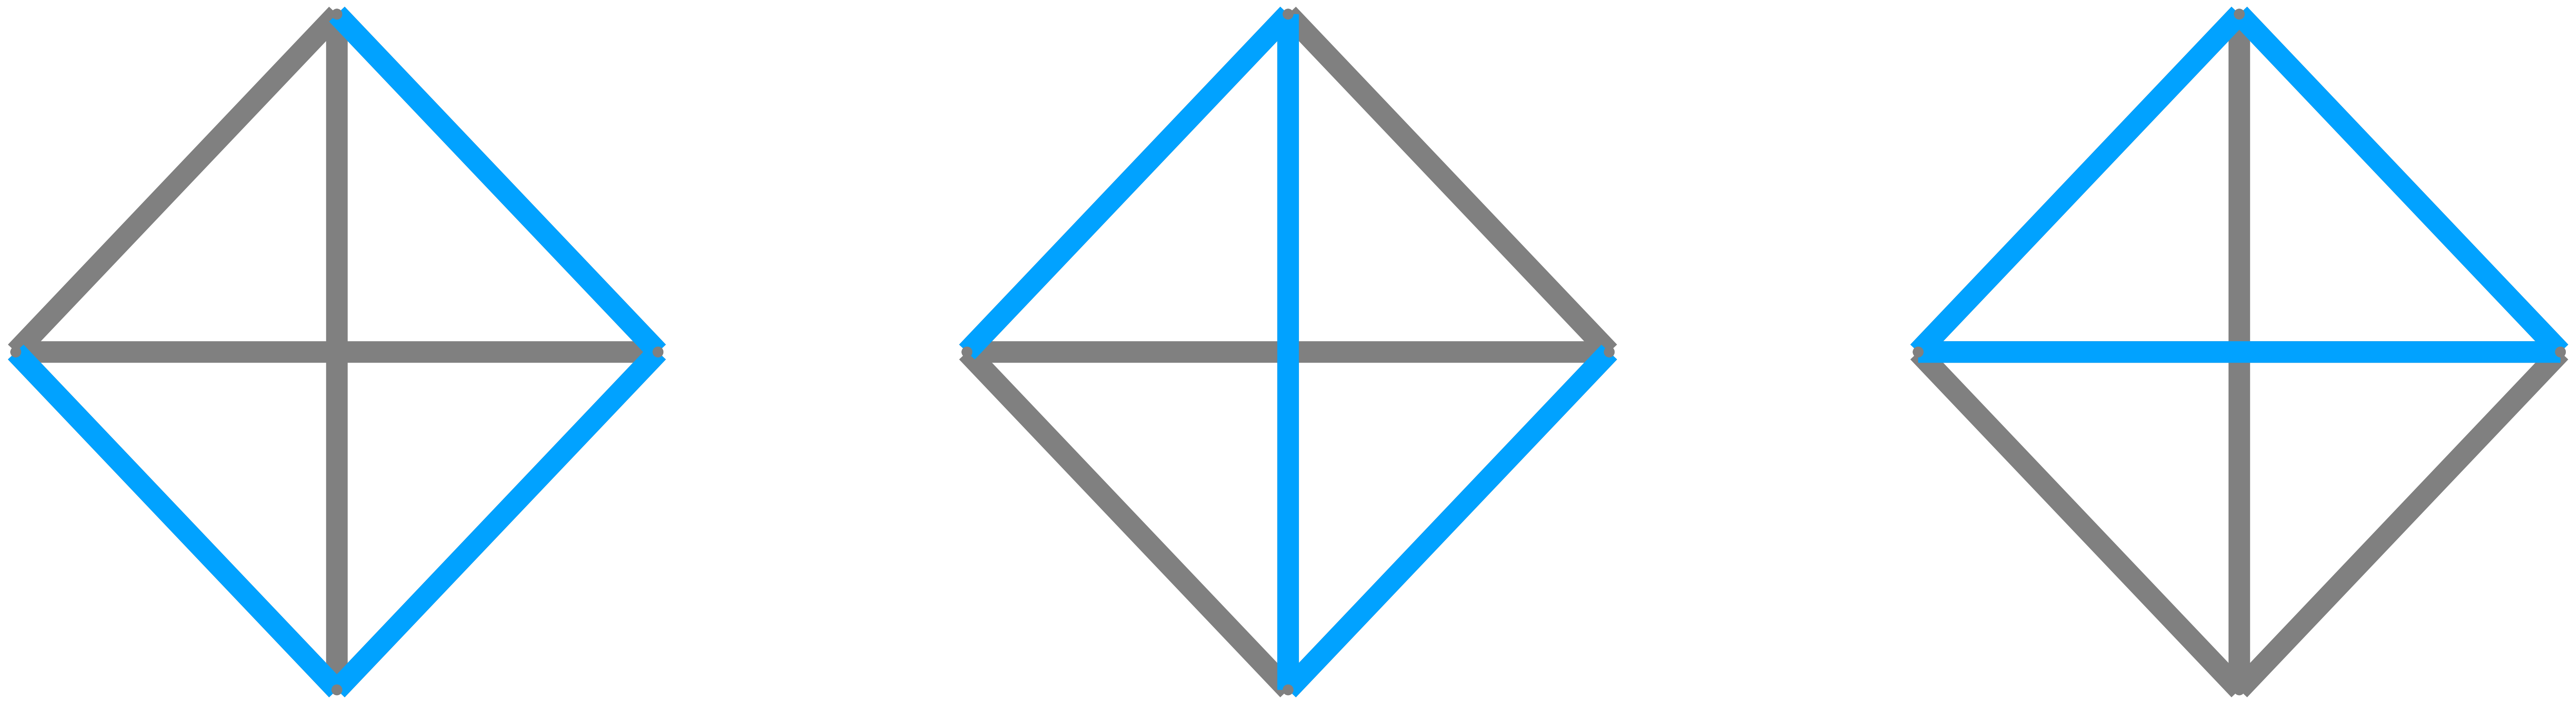

Example: K_4 (6 edges)

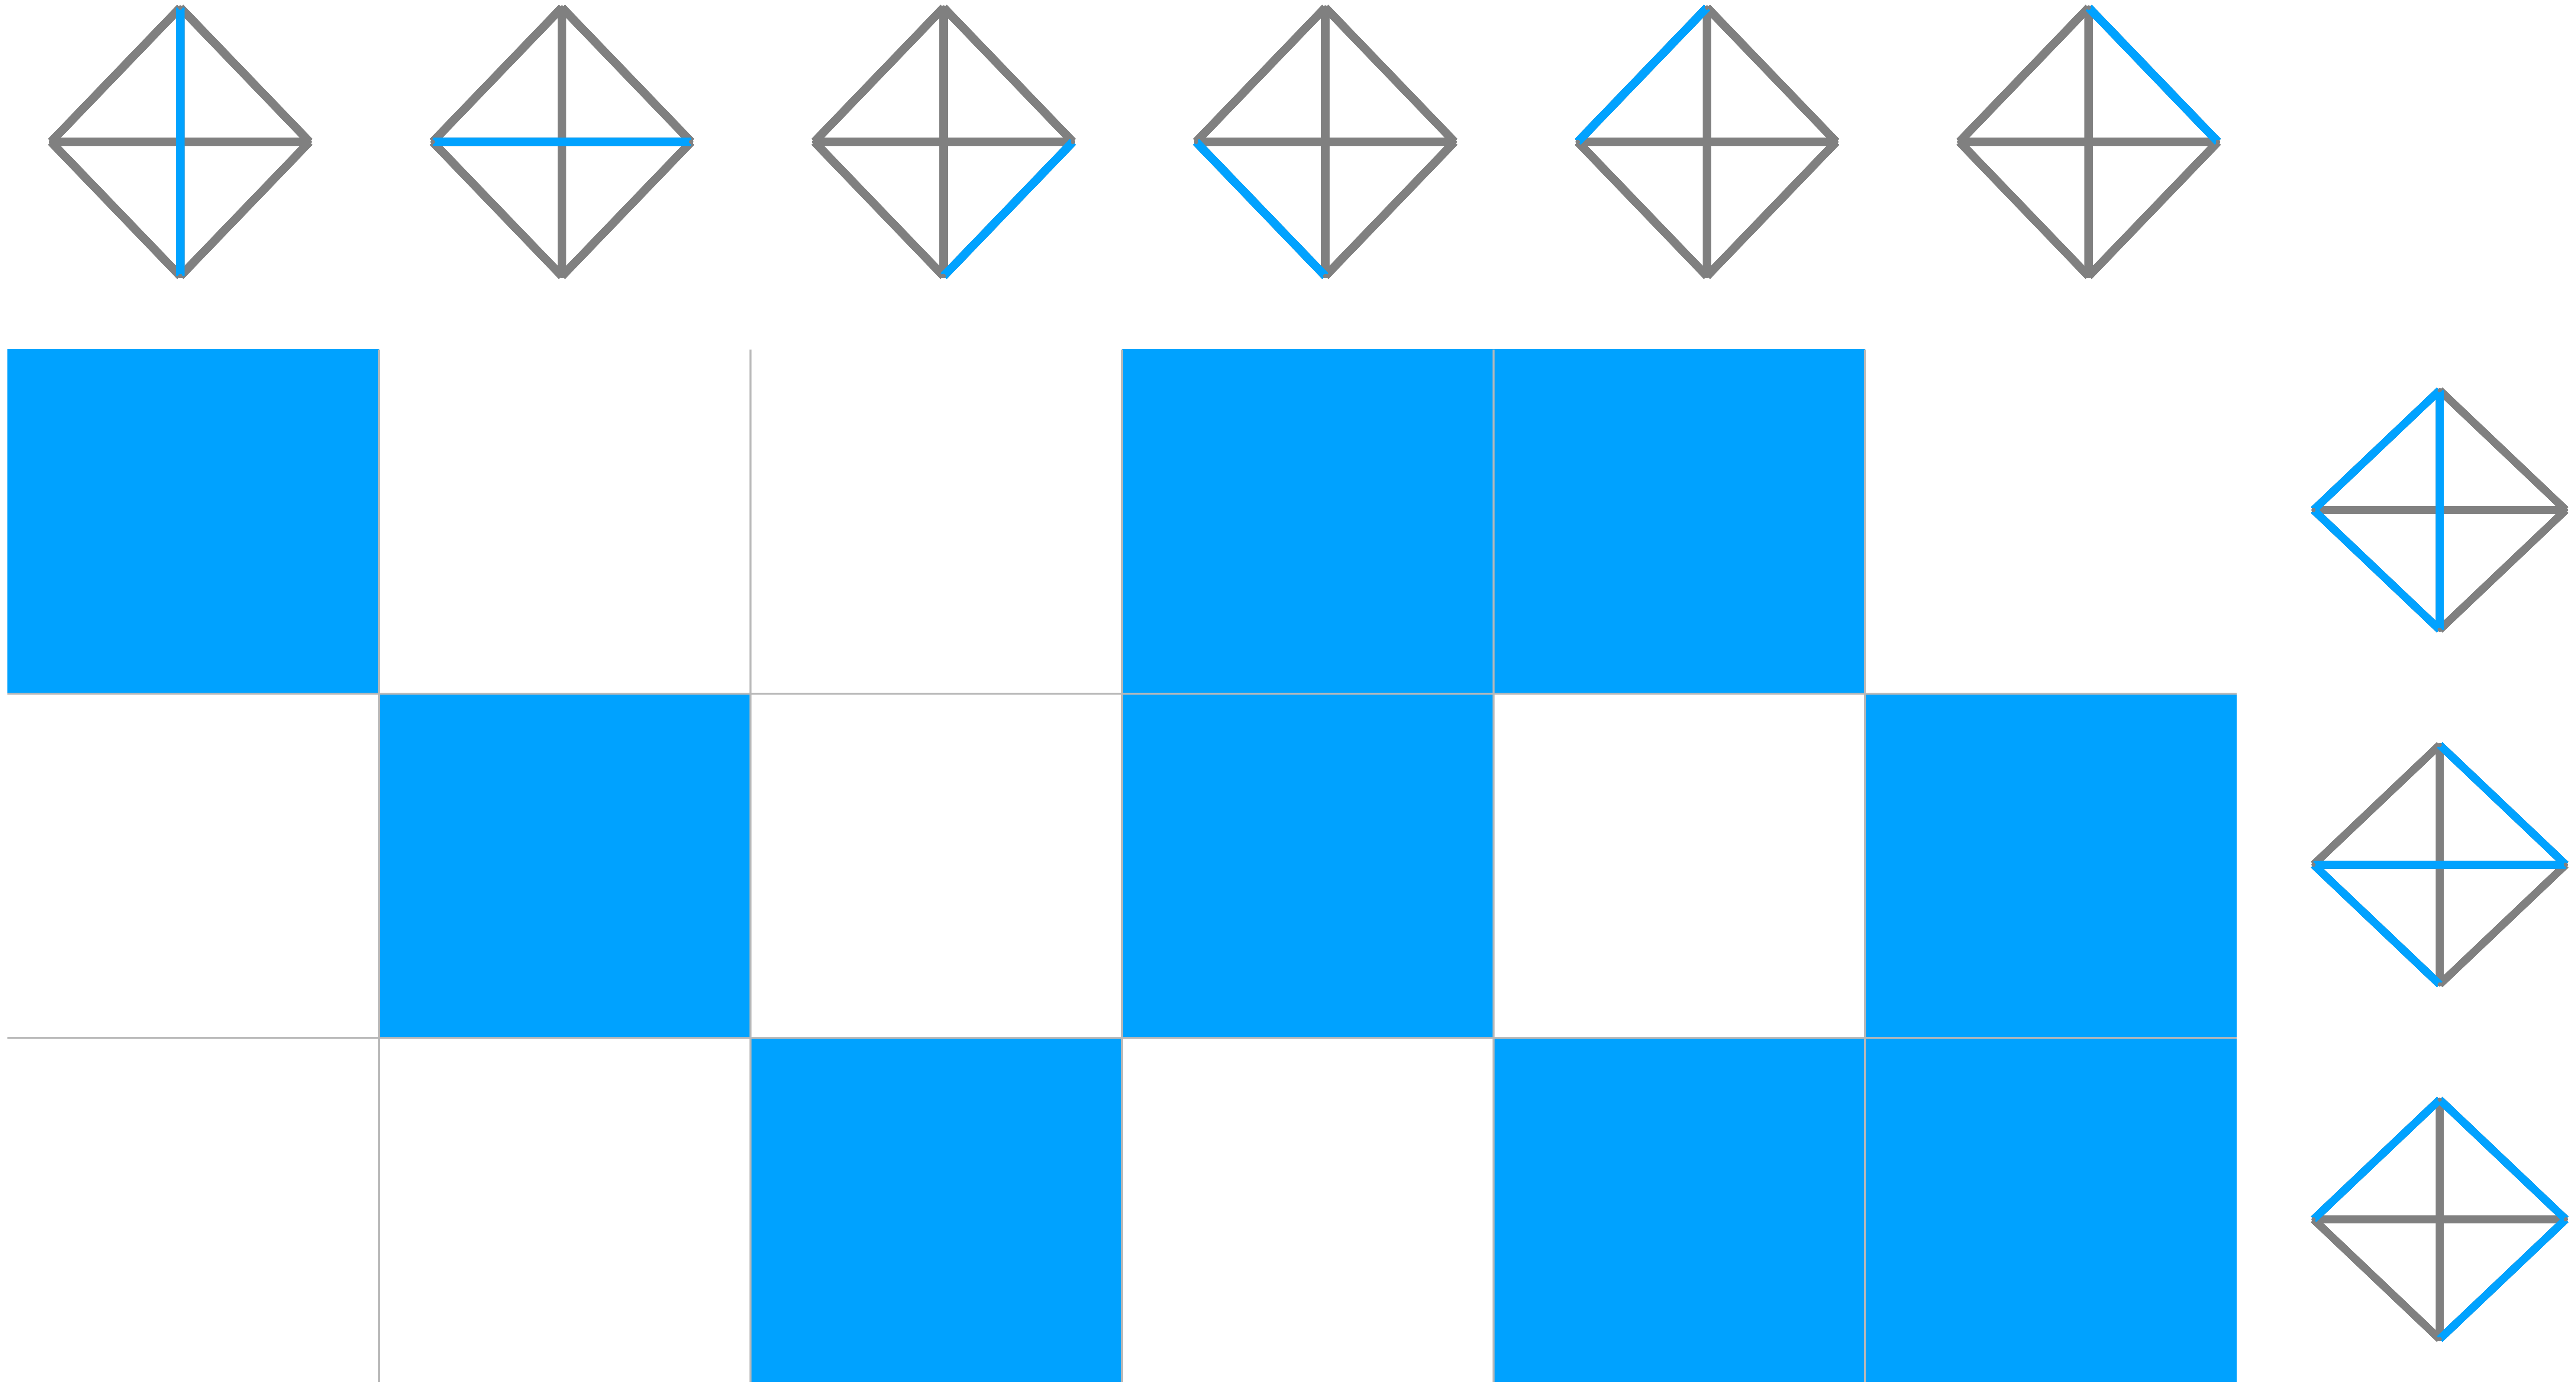

Example: K_{10} (45 edges)

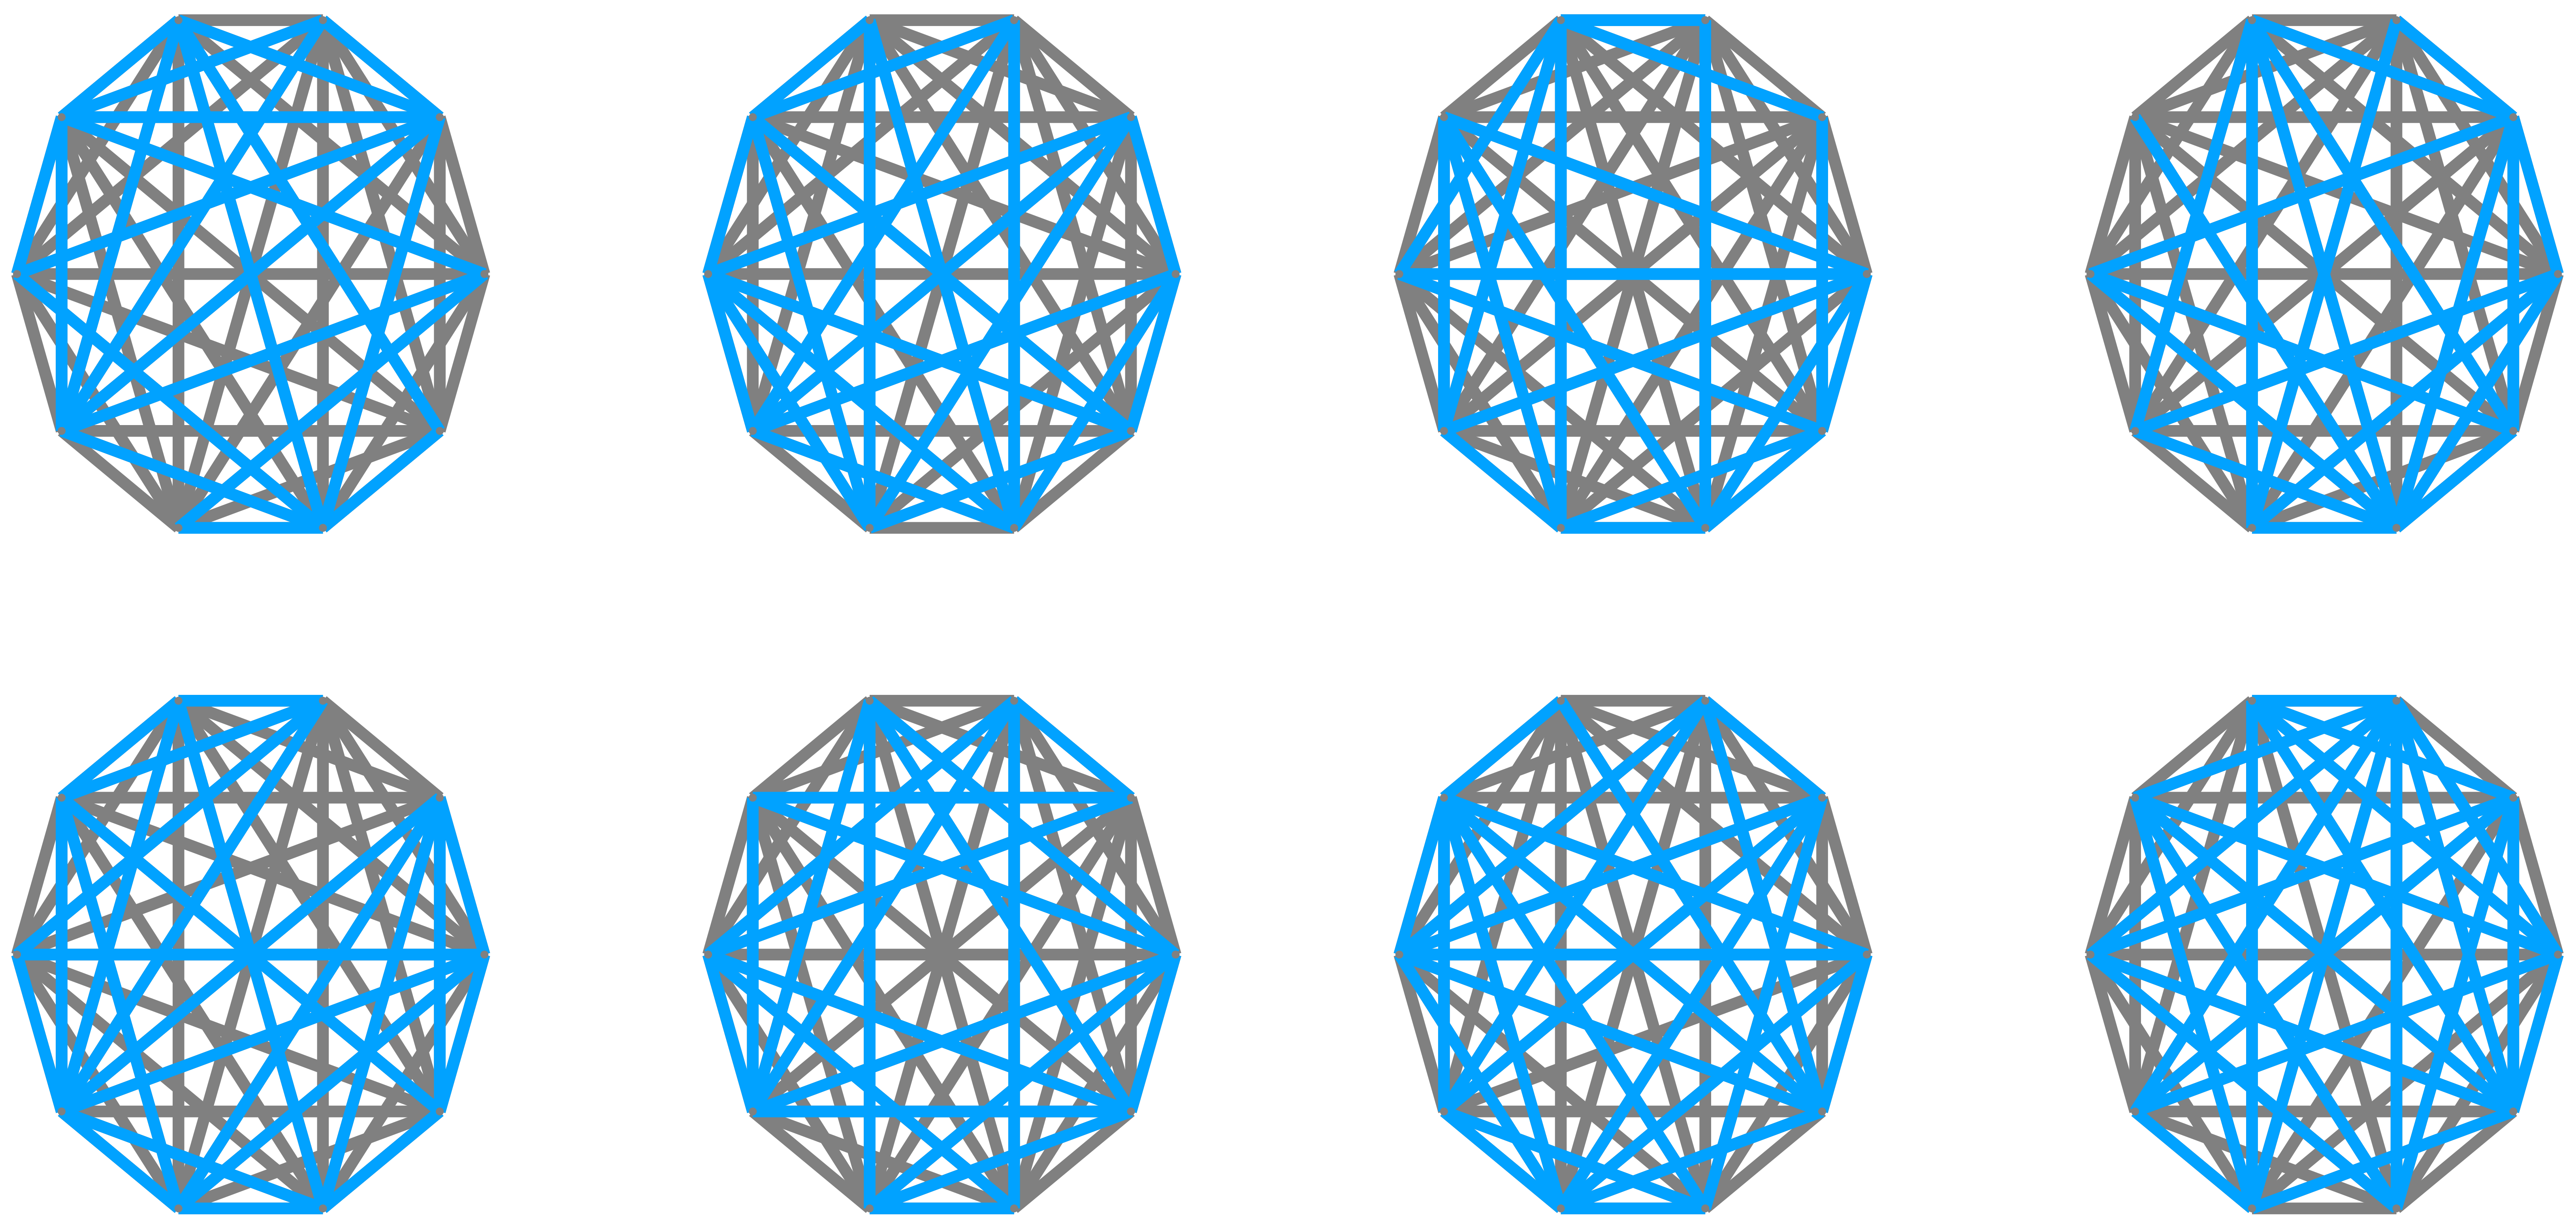

Coin weighting Compressed Sensing Traitor Tracing Streaming Algorithms Johnson-Lindenstrauss Mastermind

Network Tomography

Finding wifi users IP Traceback Error correcting codes Multicast Message Authentication

Coin weighting Compressed Sensing Traitor Tracing Streaming Algorithms Johnson-Lindenstrauss Mastermind Network Tomography Finding wifi users IP Traceback Error correcting codes Multicast Message Authentication

Key components

1. Not too many things to find
2. Get information about many things in each test
3. Can do something random

NOTES

This section is devoted to brief research and expository articles, notes on methodology and other short items.

THE DETECTION OF DEFECTIVE MEMBERS OF LARGE POPULATIONS

BY ROBERT DORFMAN

Washington, D. C.

The inspection of the individual members of a large population is an expensive and tedious process. Often in testing the results of manufacture the work can be reduced greatly by examining only a sample of the population and rejecting the whole if the proportion of defectives in the sample is unduly large. In many inspections, however, the objective is to eliminate all the defective members of the population. This situation arises in manufacturing processes where the defect being tested for can result in disastrous failures. It also arises in certain inspections of human populations. Where the objective is to weed out individual defective units, a sample inspection will clearly not suffice. It will be shown in this paper that a different statistical approach can, under certain conditions, yield significant savings in effort and expense when a complete elimination of defective units is desired.

It should be noted at the outset that when large populations are being inspected the objective of eliminating all units with a particular defect can never be fully attained. Mechanical and chemical failures and, especially, manfailures make it inevitable that mistakes will occur when many units are being examined. Although the procedure described in this paper does not directly attack the problem of technical and psychological fallibility, it may contribute to its partial solution by reducing the tediousness of the work and by making more elaborate and more sensitive inspections economically feasible. In the following discussion no attention will be paid to the possibility of technical failure or operators' error.

Thanks!