Graph Theory

Part Three

Agenda for Today

- ***The Pigeonhole Principle***
  - A simple yet surprisingly effective fact.
- ***Graph Theory Party Tricks***
  - Cool tricks to try at your next group meeting.
- ***A Little Movie Puzzle***
  - Who watched what?

Recap from Last Time

Recap from Last Time

- When there’s an edge between two nodes, we say they are ___________.
- If there’s a path between two nodes, we say they are _____________ from one another.

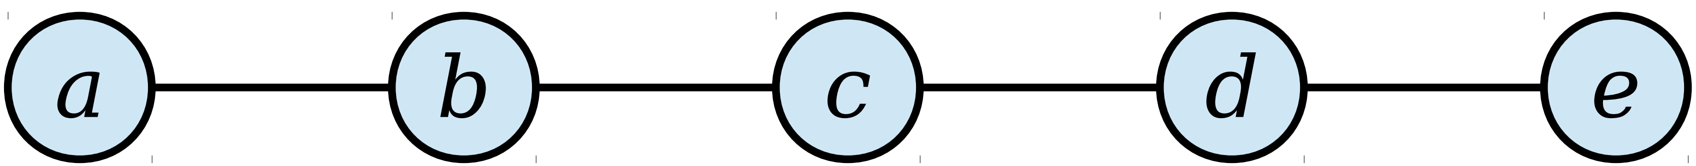

Recap from Last Time

- When there's an edge between two nodes, we say they are __________.
- If there's a path between two nodes, we say they are ____________ from one another.

PollEv.com/cs103spr26

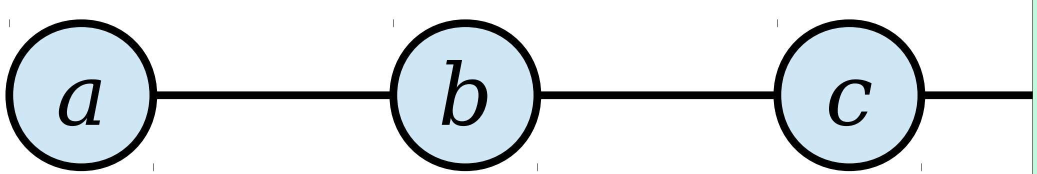

Recap from Last Time

- When there's an edge between two nodes, we say they are ***adjacent***.
- If there's a path between two nodes, we say they are ***reachable*** from one another.

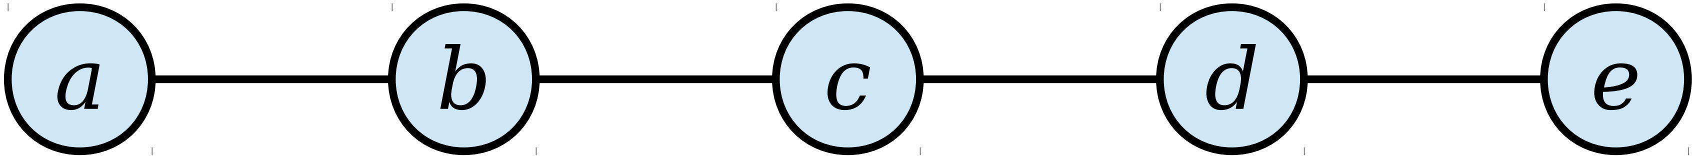

The Pigeonhole Principle

The Pigeonhole Principle

Theorem (The Pigeonhole Principle):
If m objects are distributed into n bins and $m > n$, then at least one bin will contain at least two objects.

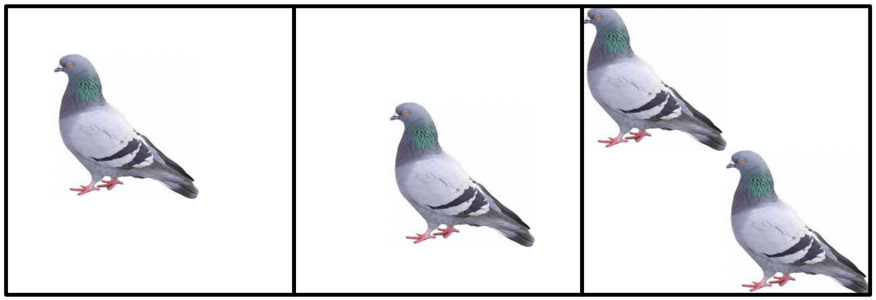

The Pigeonhole Principle

Theorem (The Pigeonhole Principle):

If m objects are distributed into n bins and $m > n$, then at least one bin will contain at least two objects.

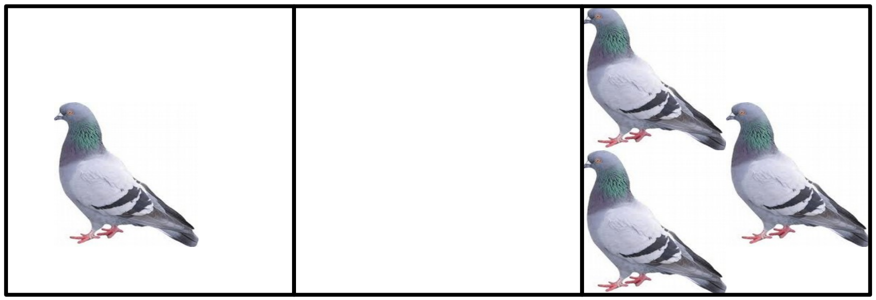

The Generalized Pigeonhole Principle

A More General Version

- The ***generalized pigeonhole principle*** says that if you distribute *m* objects into *n* bins, then
  - some bin will have at least $\lceil m/n \rceil$ objects in it, and
  - some bin will have at most $\lfloor m/n \rfloor$ objects in it.

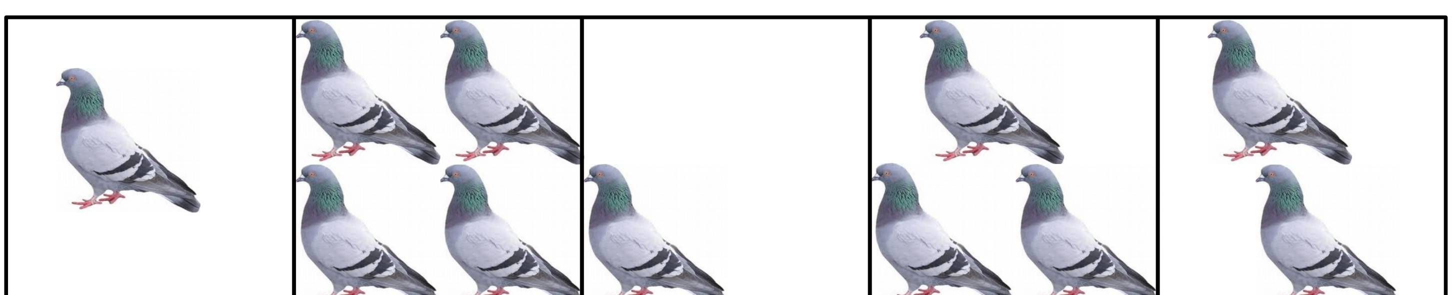

$m = 11$
$n = 5$

$\lceil m / n \rceil = 3$
$\lfloor m / n \rfloor = 2$

New stuff!

An Application: Friends and Strangers

Friends and Strangers

- Suppose you have a party of **six people**. Each pair of people are either friends (they know each other) or strangers (they do not).
- ***Theorem ("Theorem on Friends and Strangers"):*** Any such **6-person** party must have a group of three mutual friends (three people who all know one another) or three mutual strangers (three people, none of whom know any of the others).

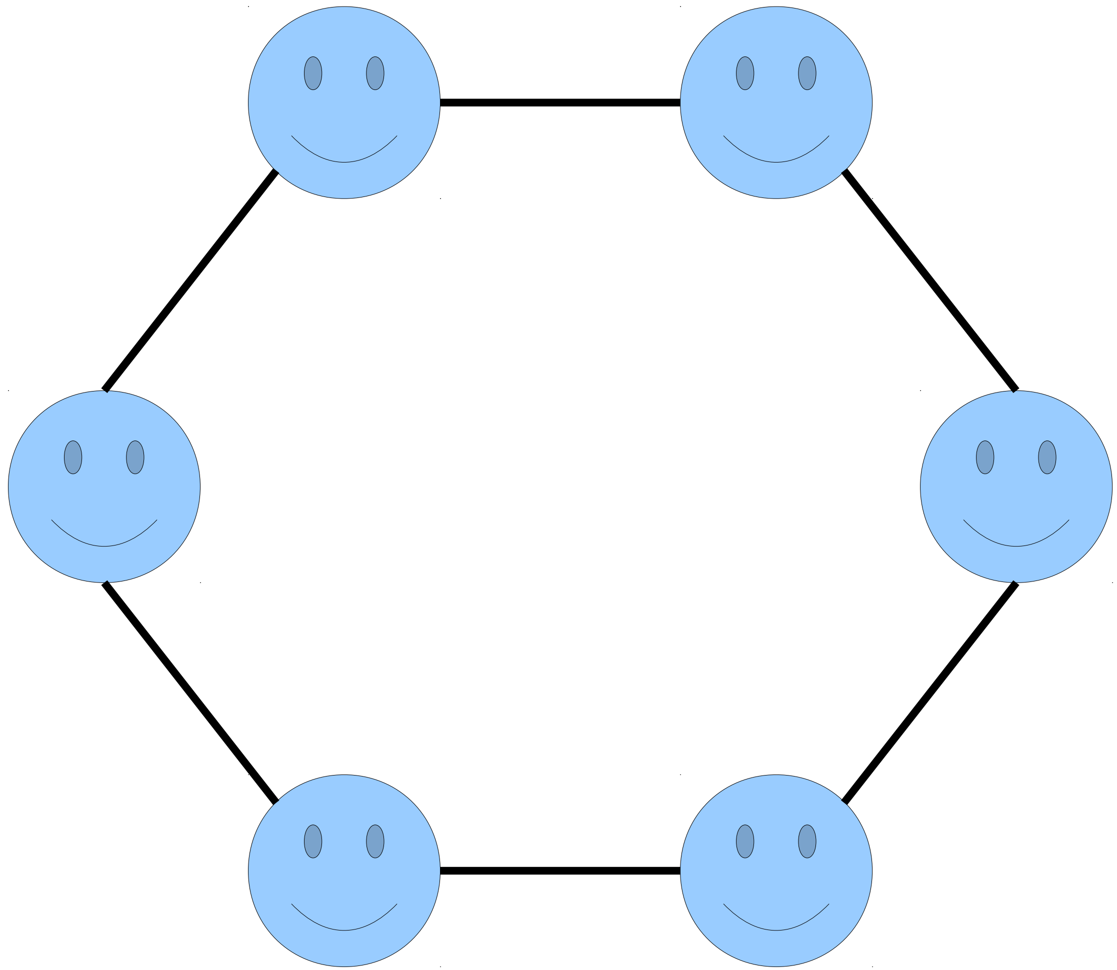

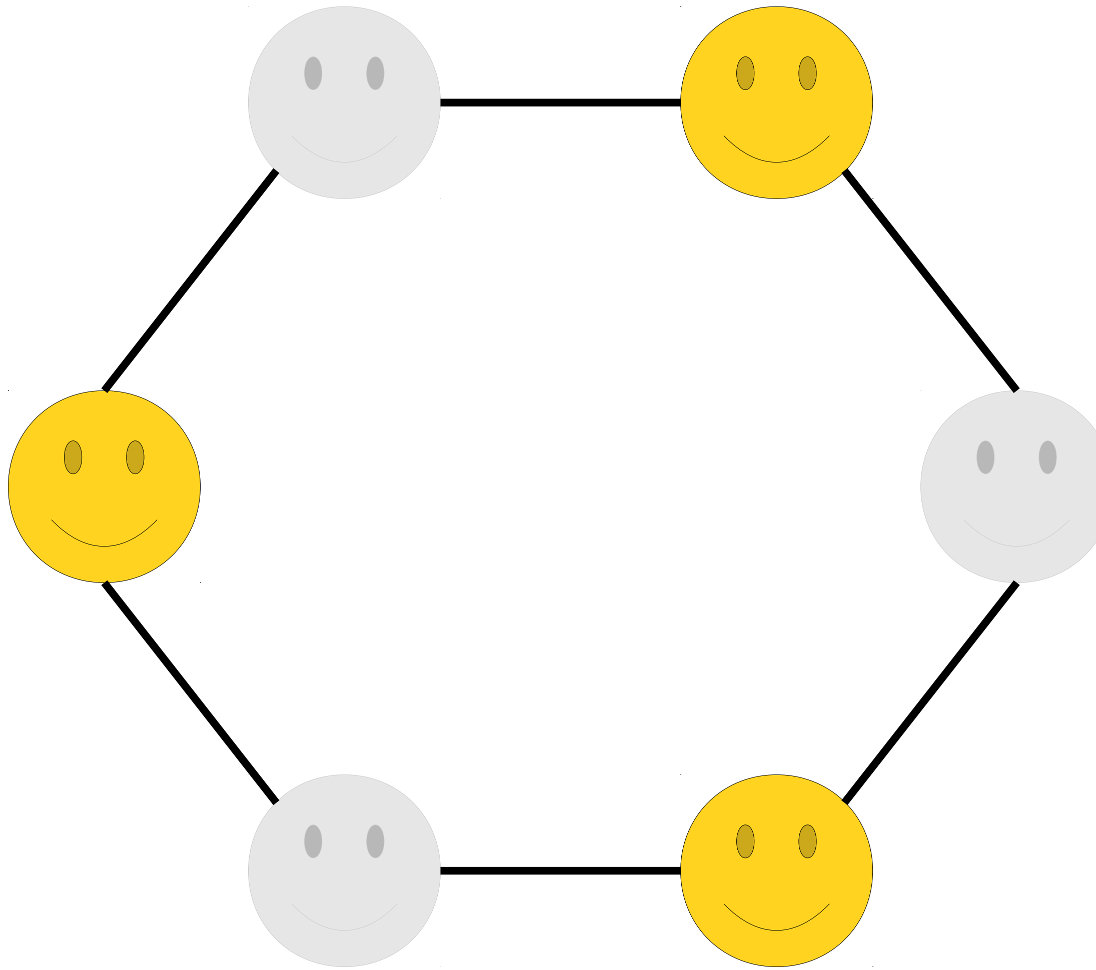

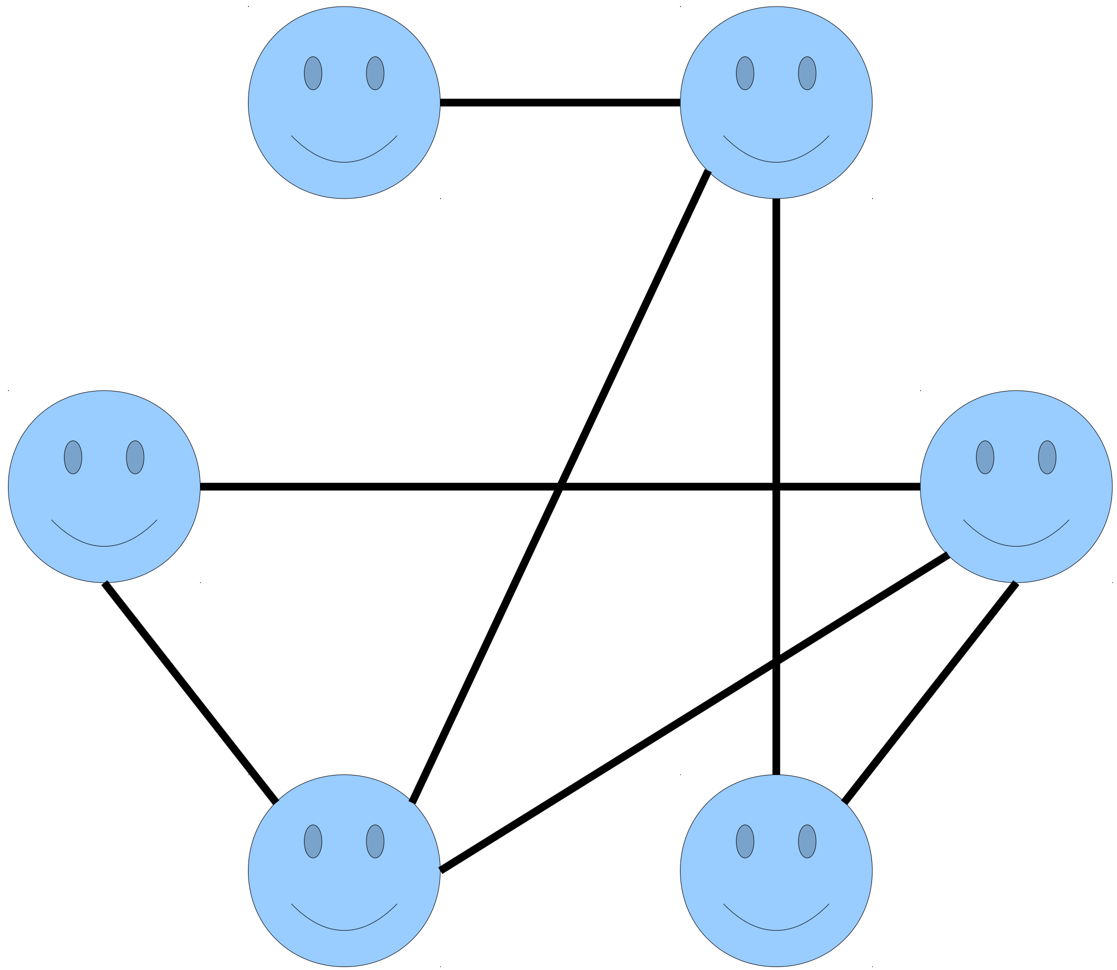

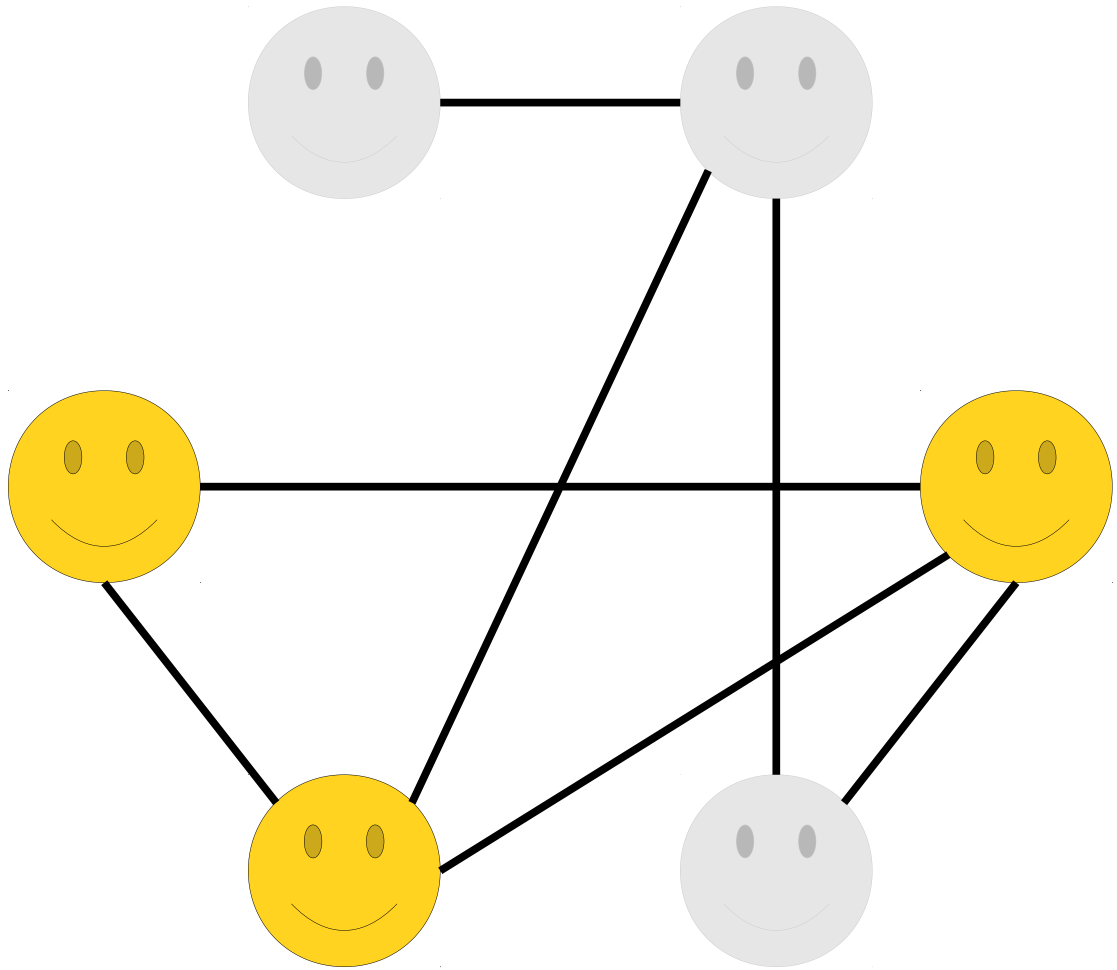

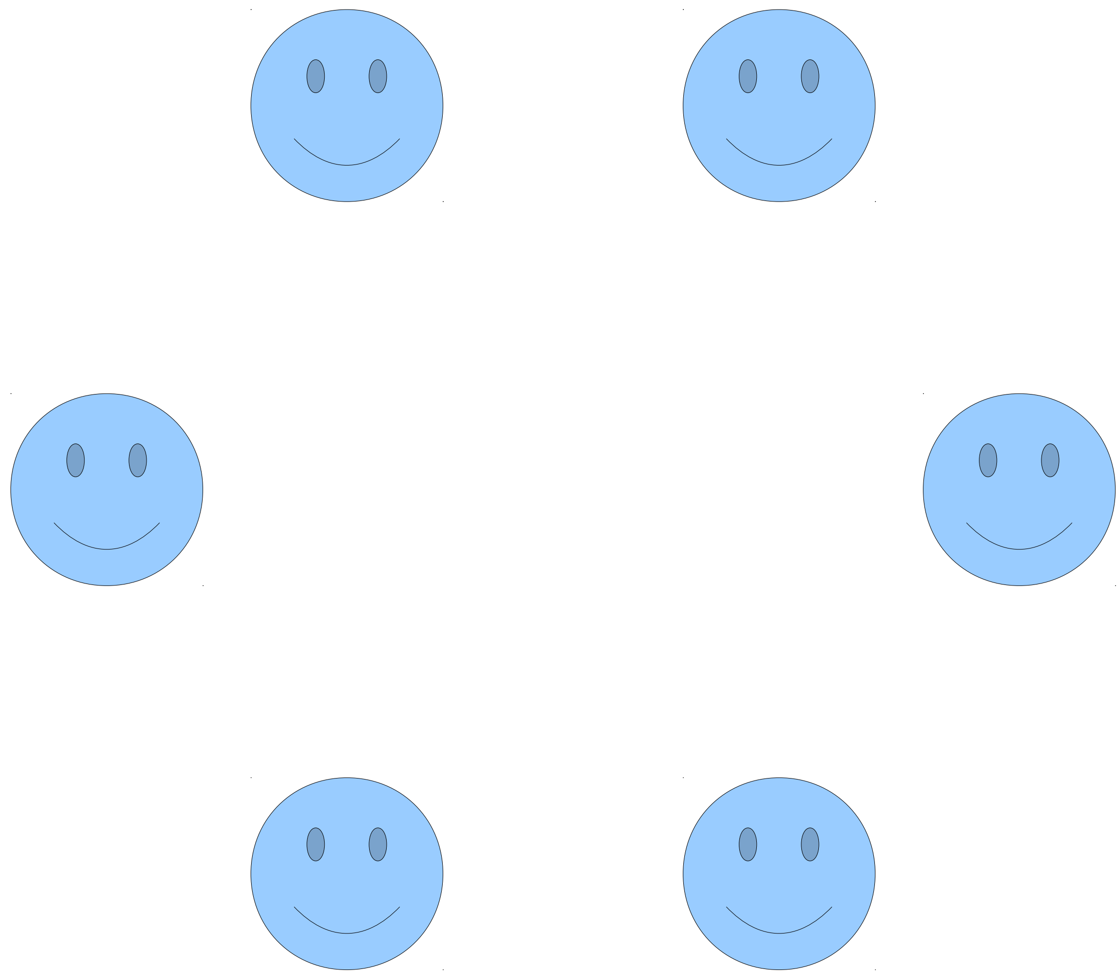

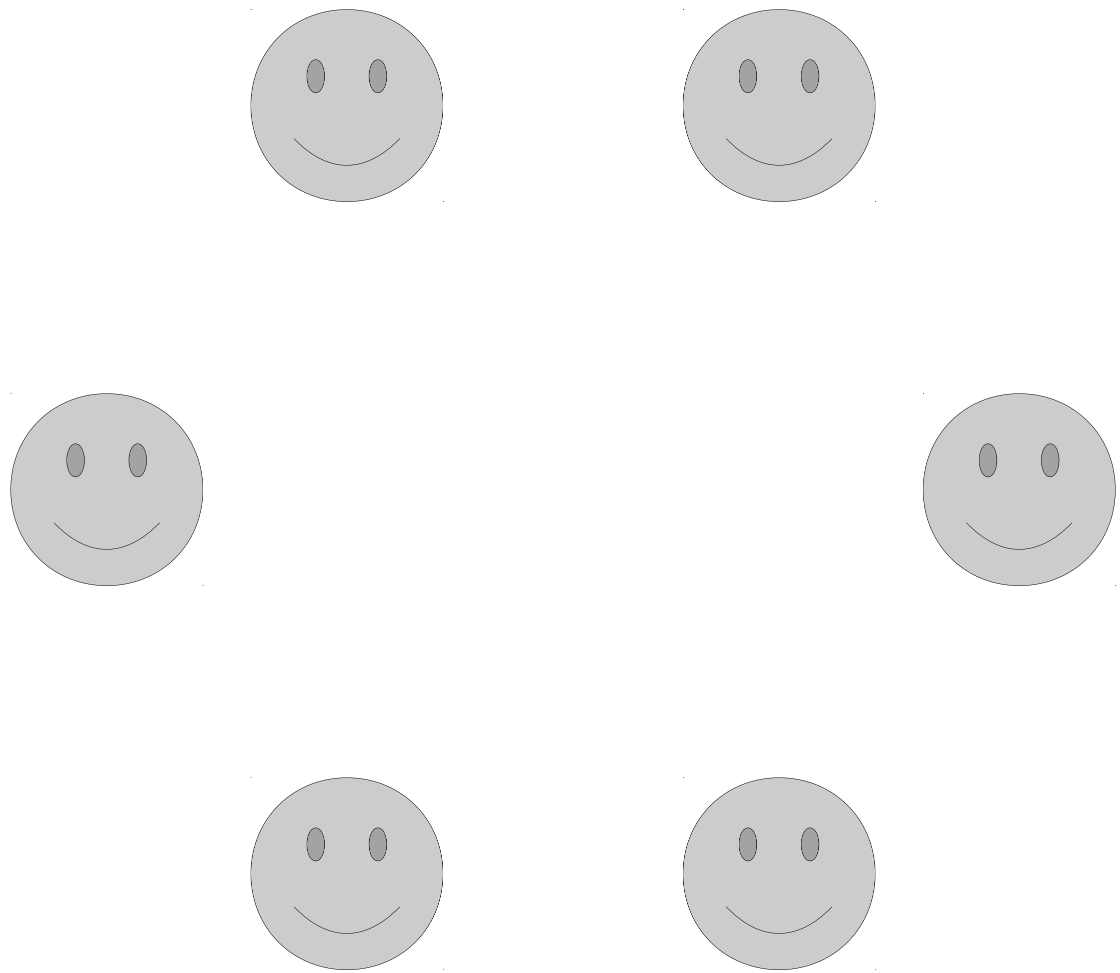

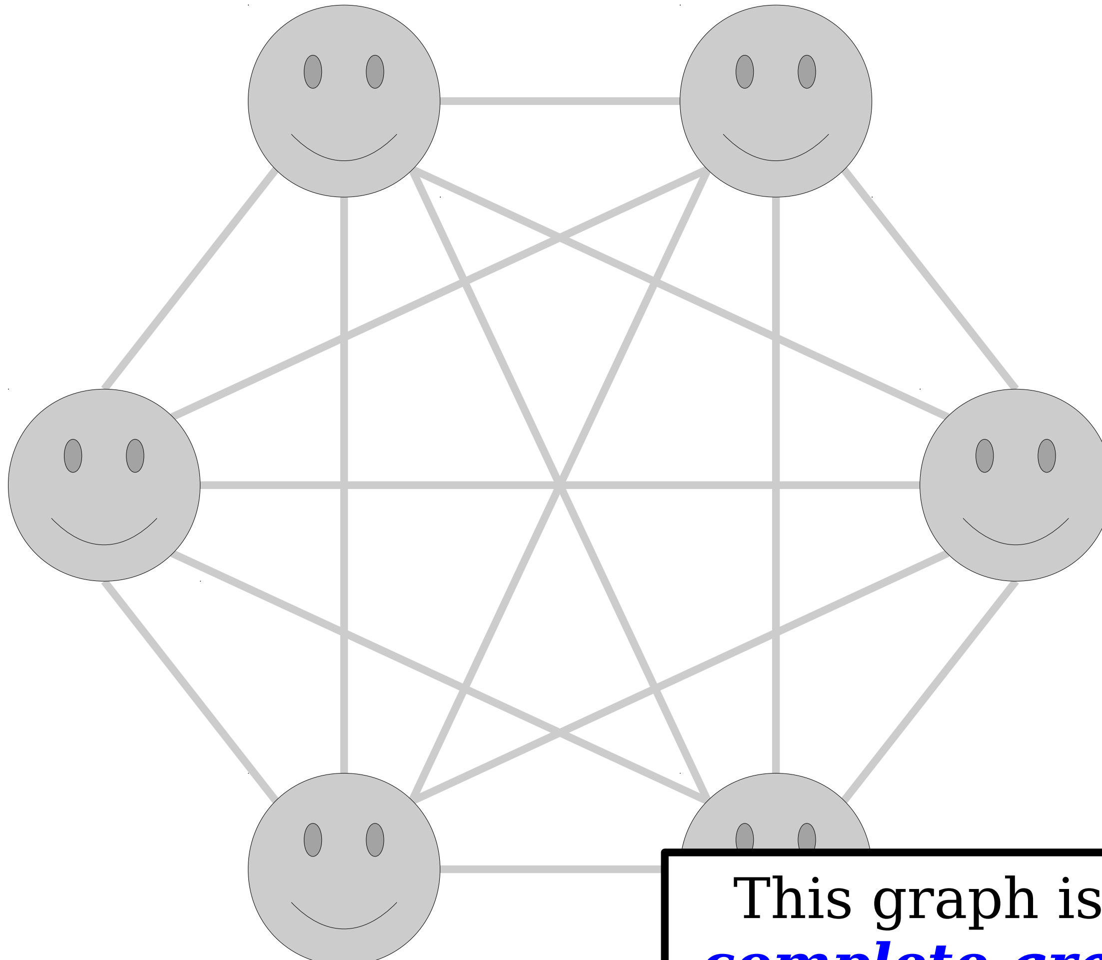

This graph is called K_6, the ***complete graph of order 6***. More generally, the graph K_n consists of n mutually adjacent nodes.

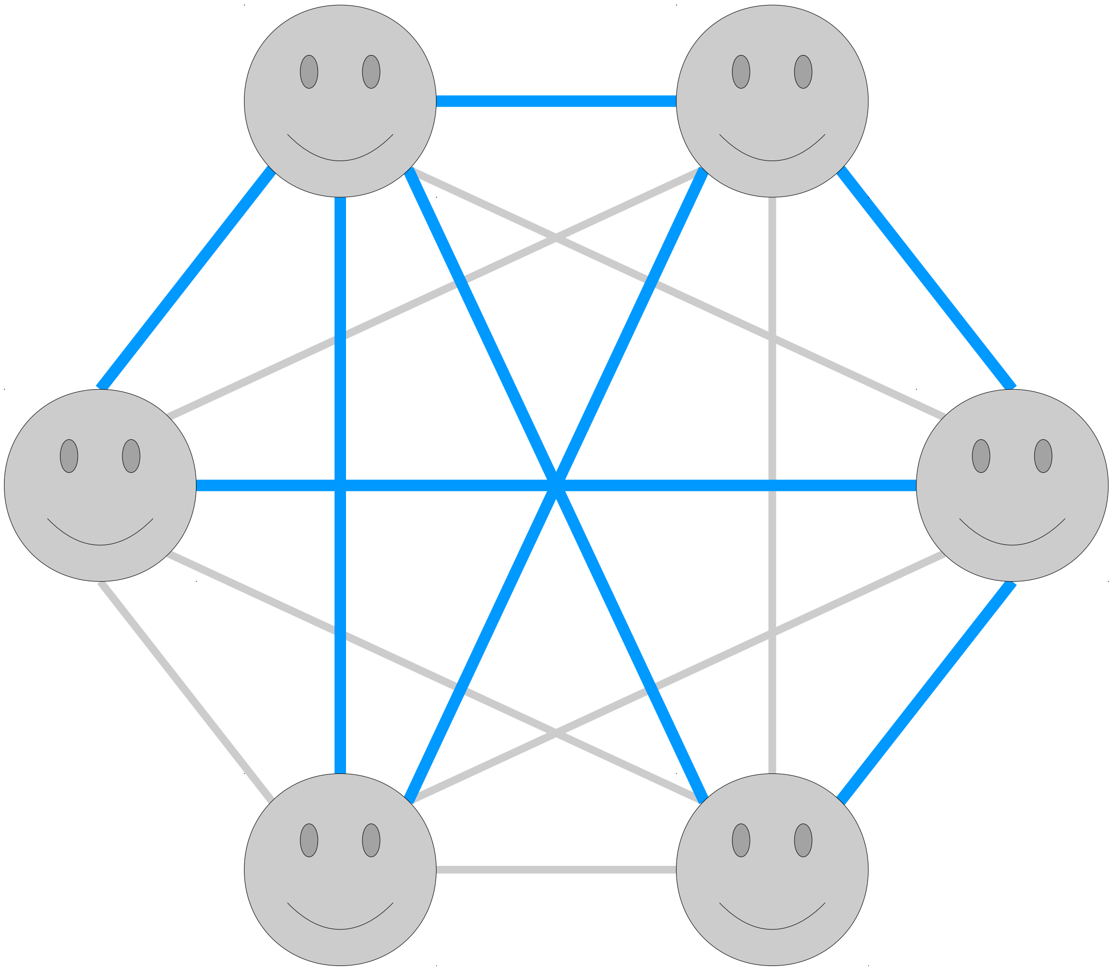

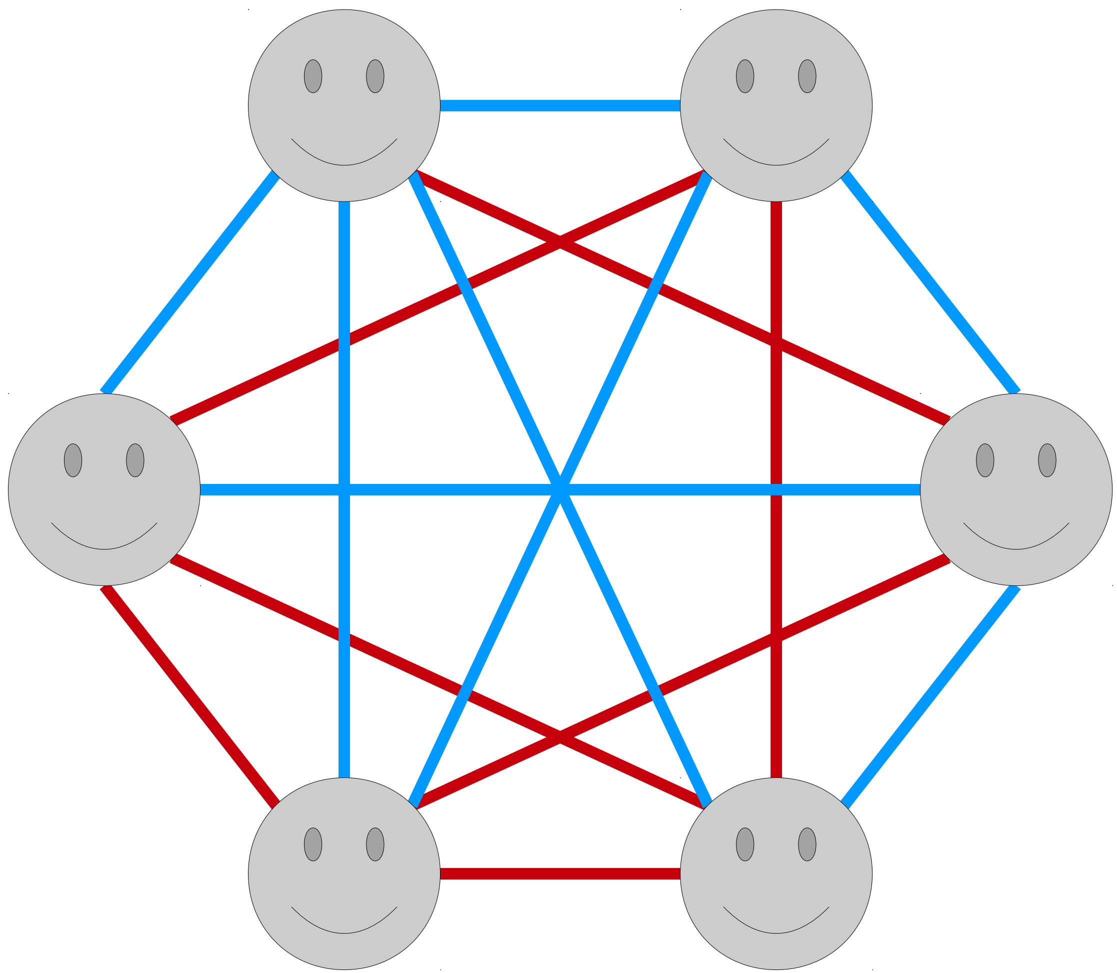

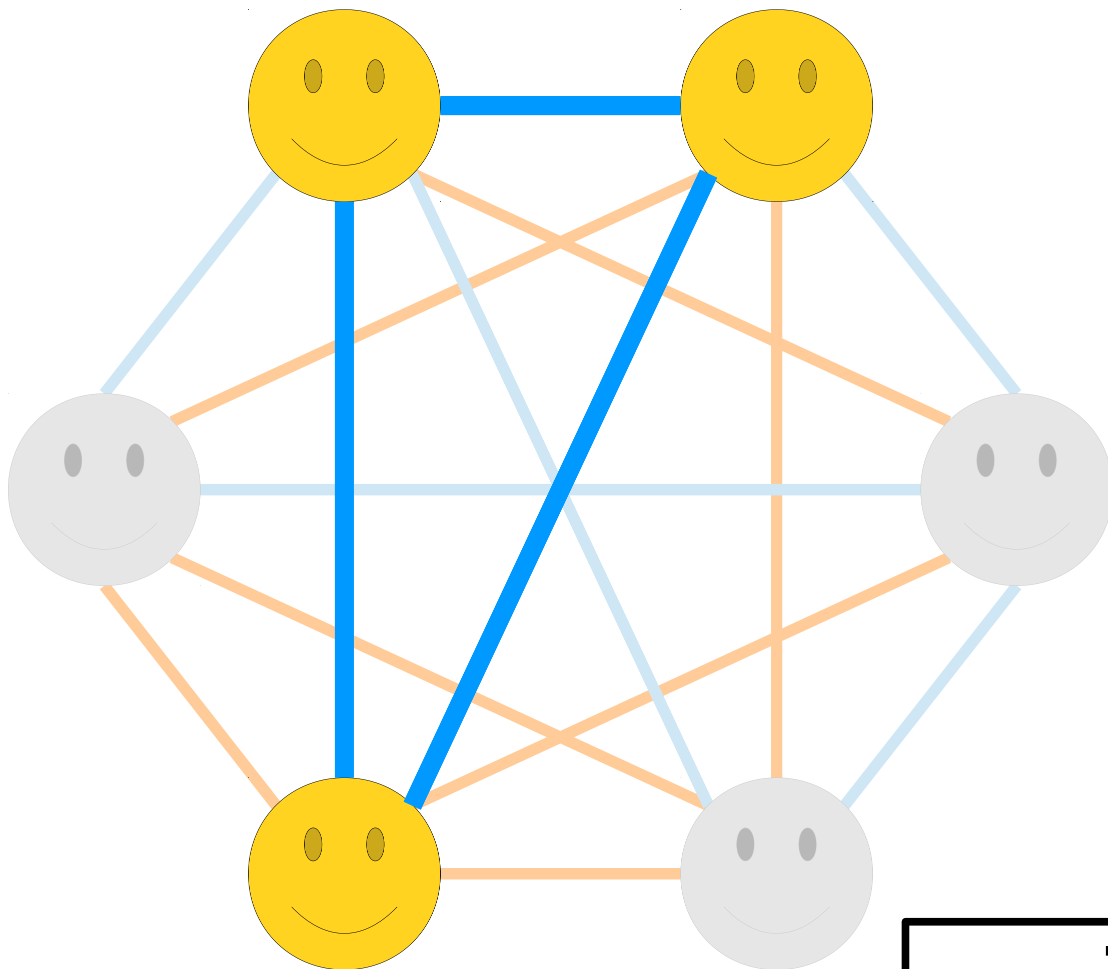

This is a ***monochrome*** (one-color) copy of K_3.

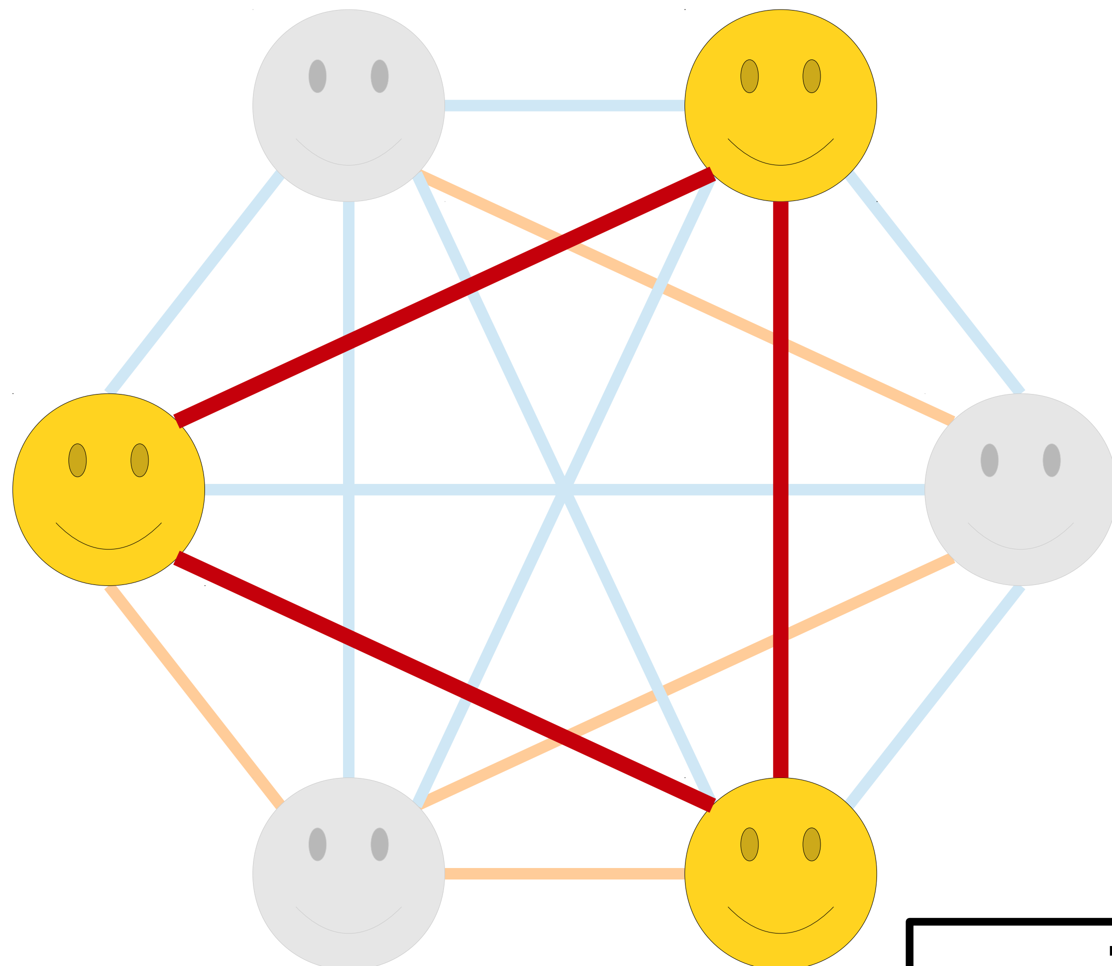

This is a ***monochrome*** (one-color) copy of K_3.

Friends and Strangers Restated

- From a graph-theoretic perspective, the Theorem on Friends and Strangers can be restated as follows:

  Theorem: Color every edge of K_6 either red or blue. The resulting graph always contains a monochrome copy of K_3.

- How can we prove this result?

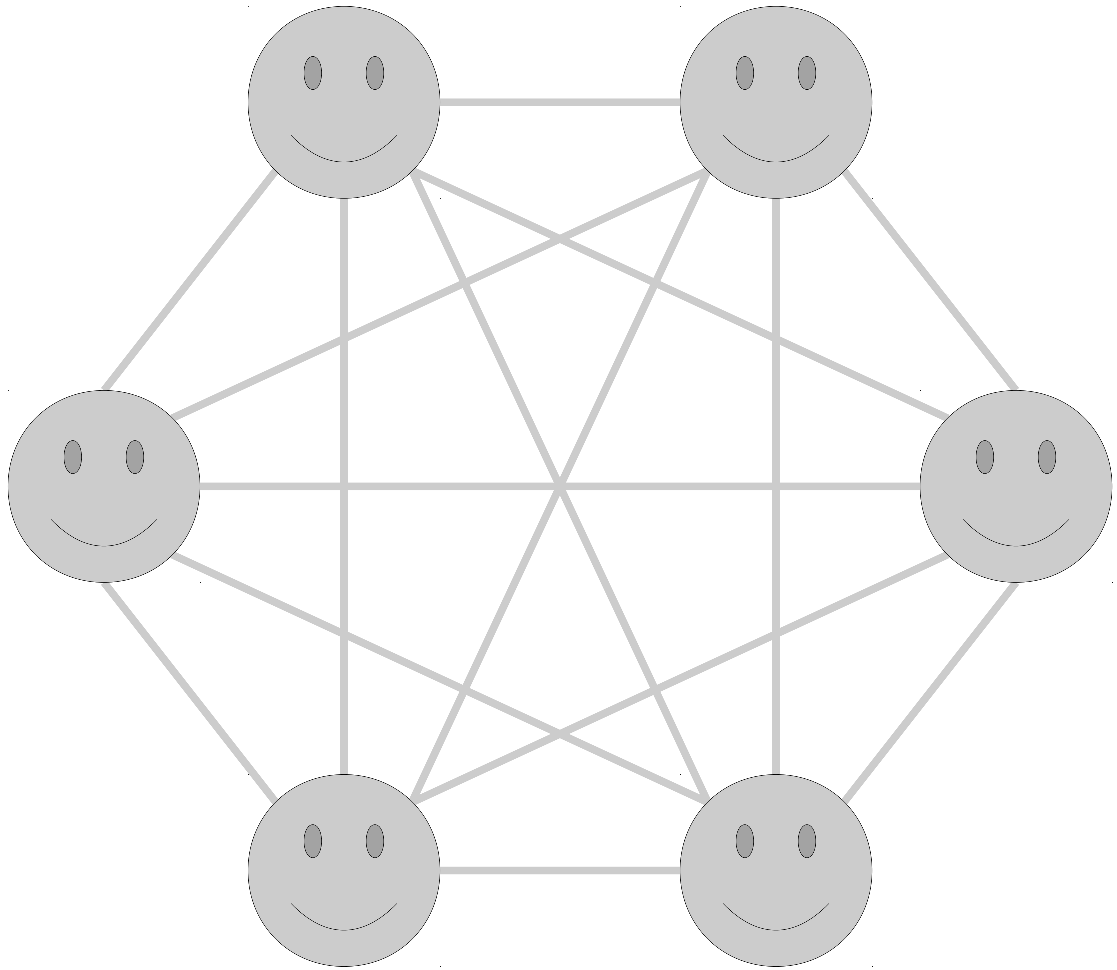

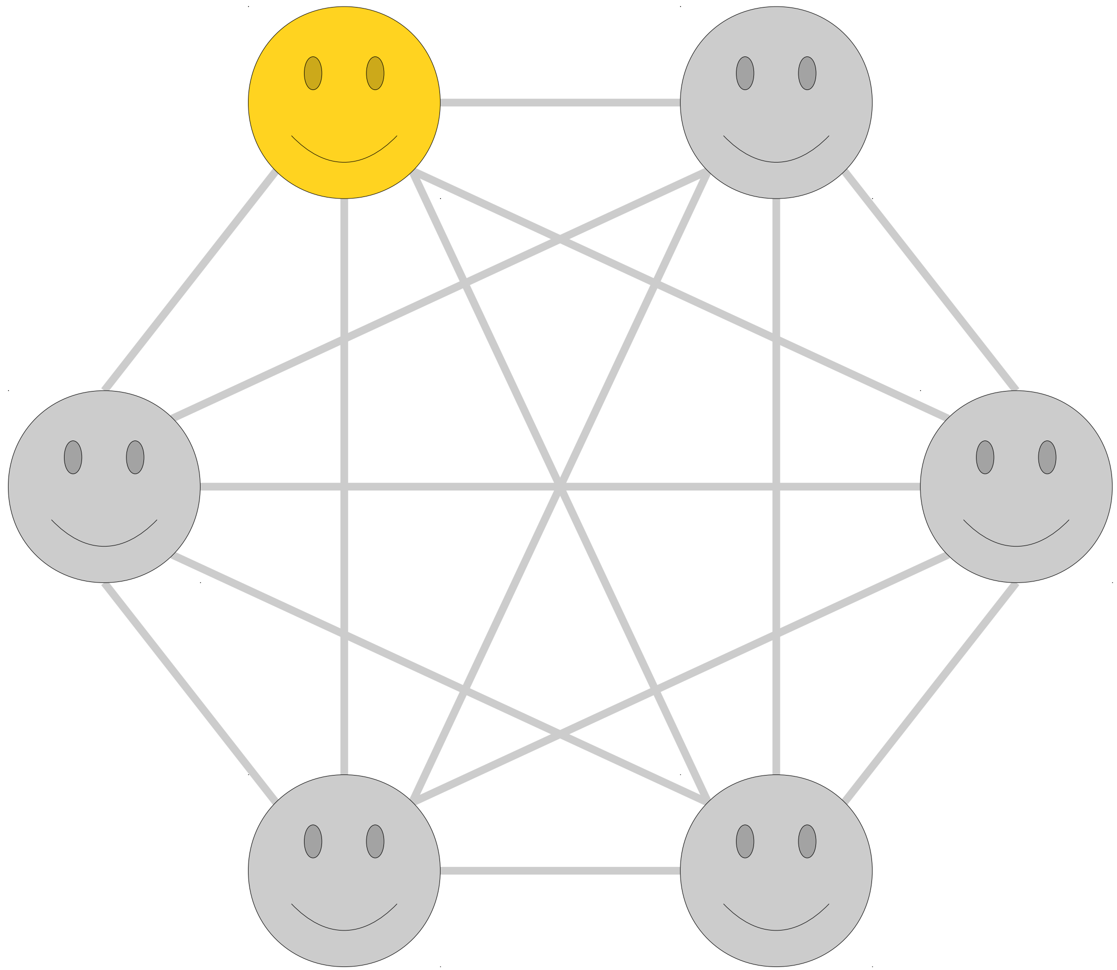

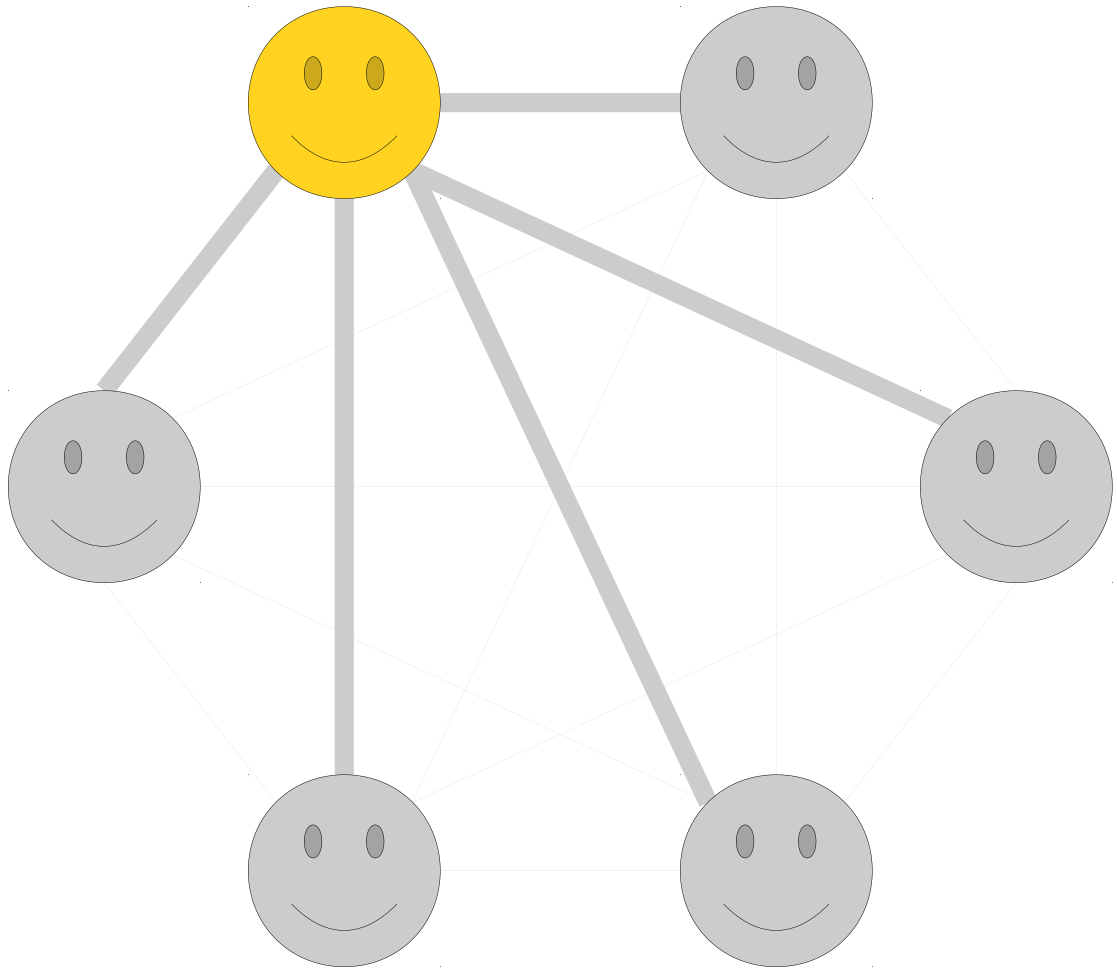

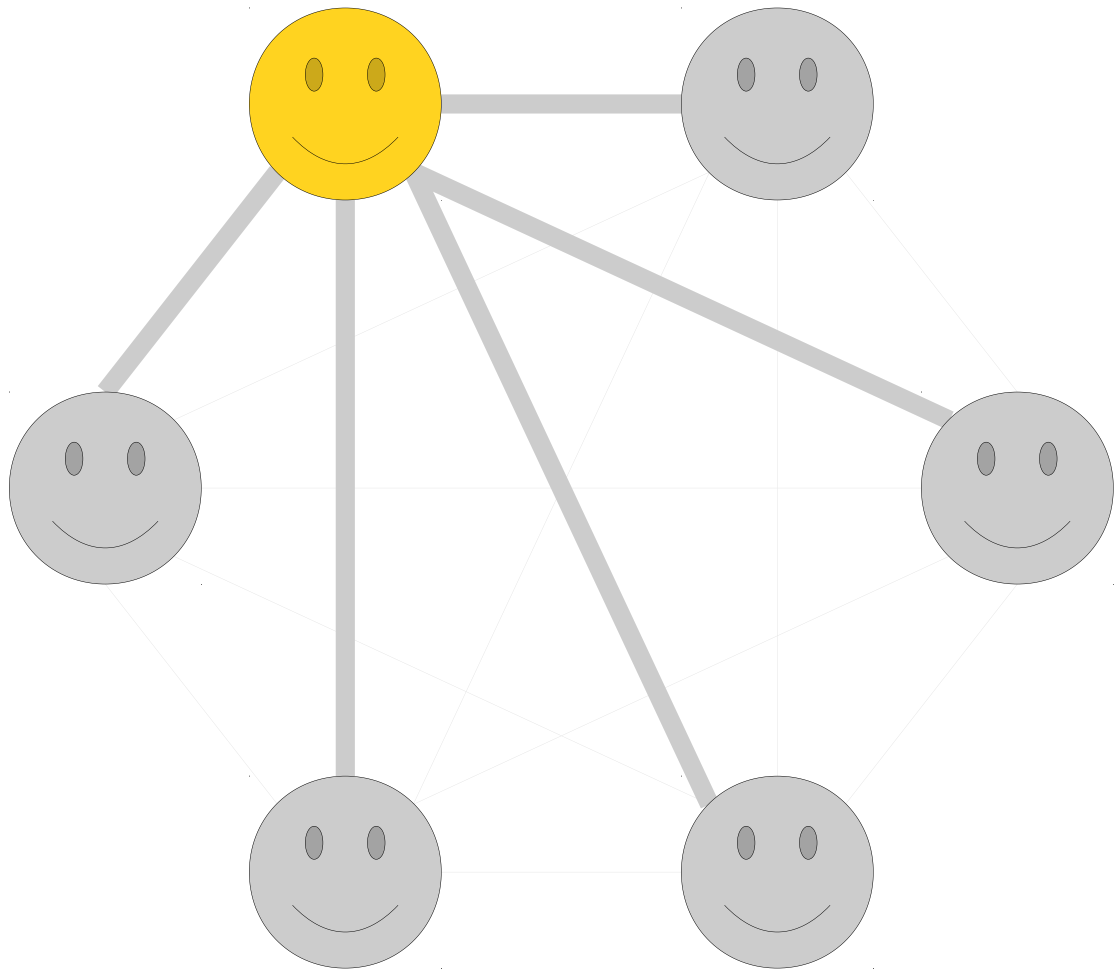

Nodes adjacent to this one: 5
Colors available: 2

Generalized Pigeonhole Principle says that some color must have at least X edges of this node that are that color.

Warmup: what are the bins, what are the birds, in this scenario?

PollEv.com/cs103spr26

- The ***generalized pigeonhole principle*** says that if you distribute *m* "birds" into *n* "bins," then
  - some bin will have at least $\lceil m/n \rceil$ birds in it, and
  - some bin will have at most $\lfloor m/n \rfloor$ birds in it.

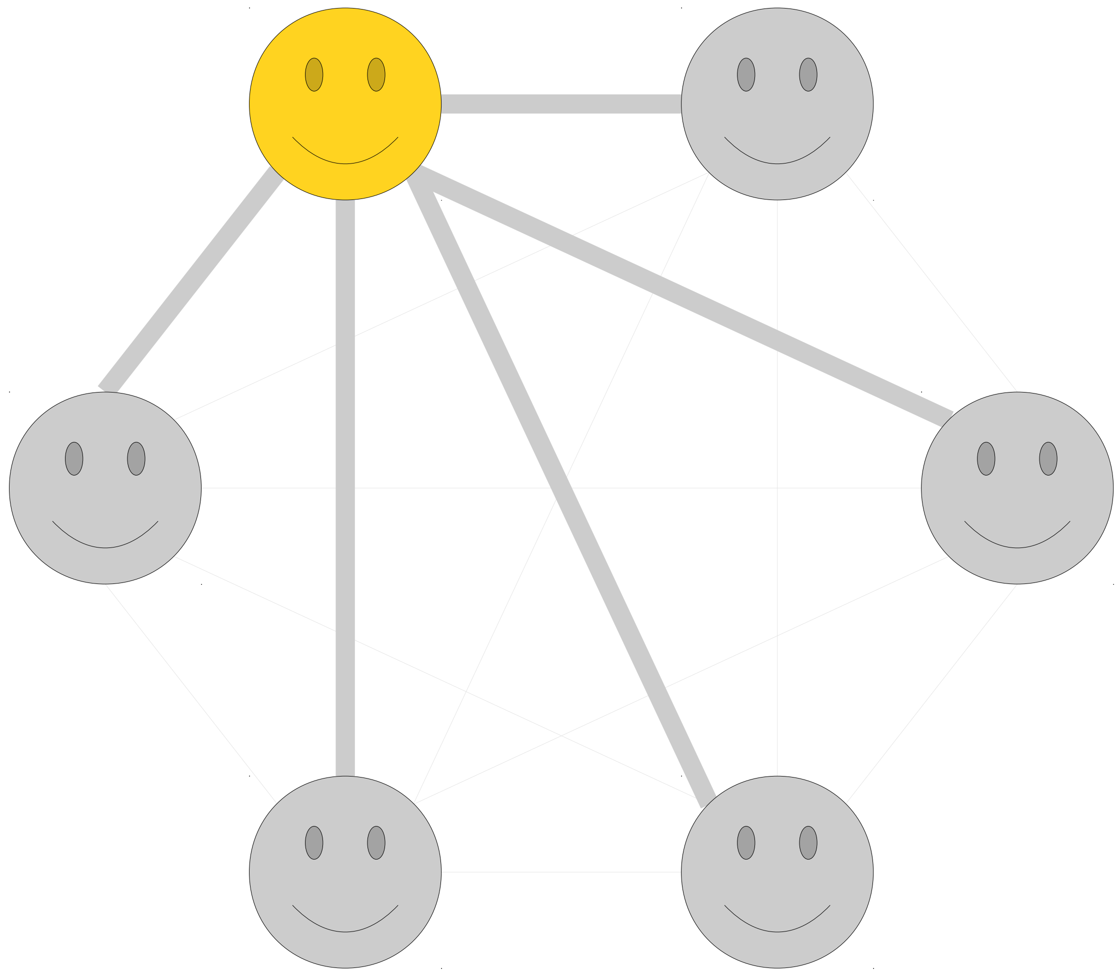

Nodes adjacent to this one: 5
Colors available: 2

Generalized Pigeonhole Principle says that some color must have at least X edges of this node that are that color.

Question: What is X?

PollEv.com/cs103spr26

- The ***generalized pigeonhole principle*** says that if you distribute *m* "birds" into *n* "bins," then
  - some bin will have at least $\lceil m/n \rceil$ birds in it, and
  - some bin will have at most $\lfloor m/n \rfloor$ birds in it.

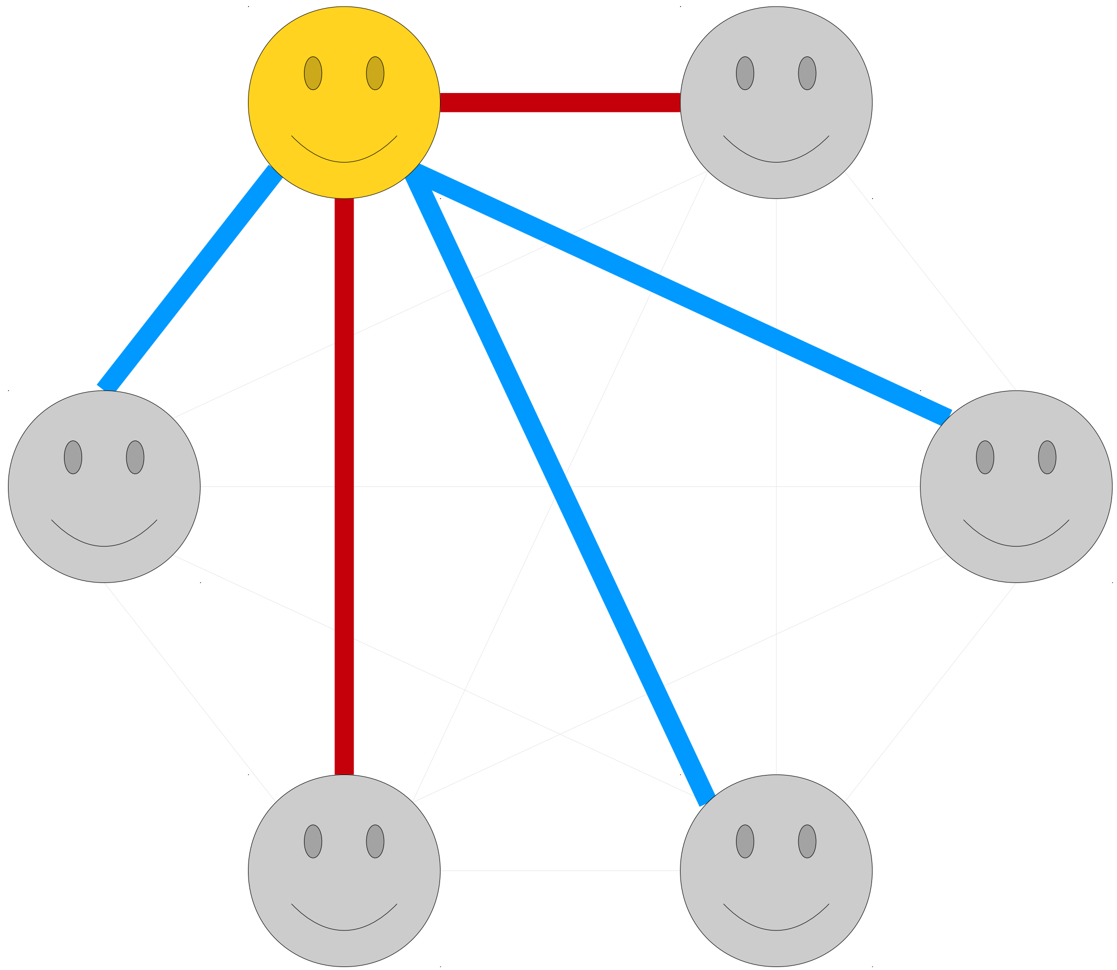

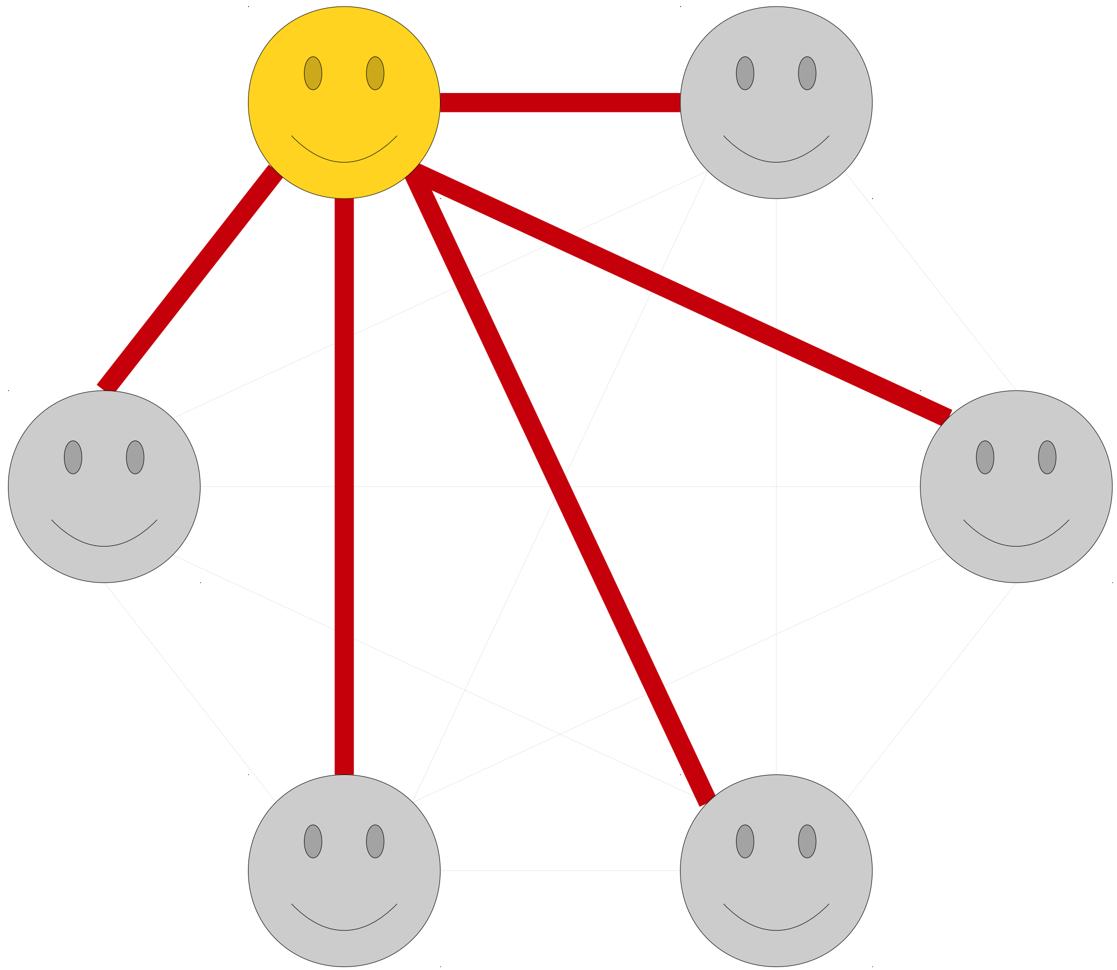

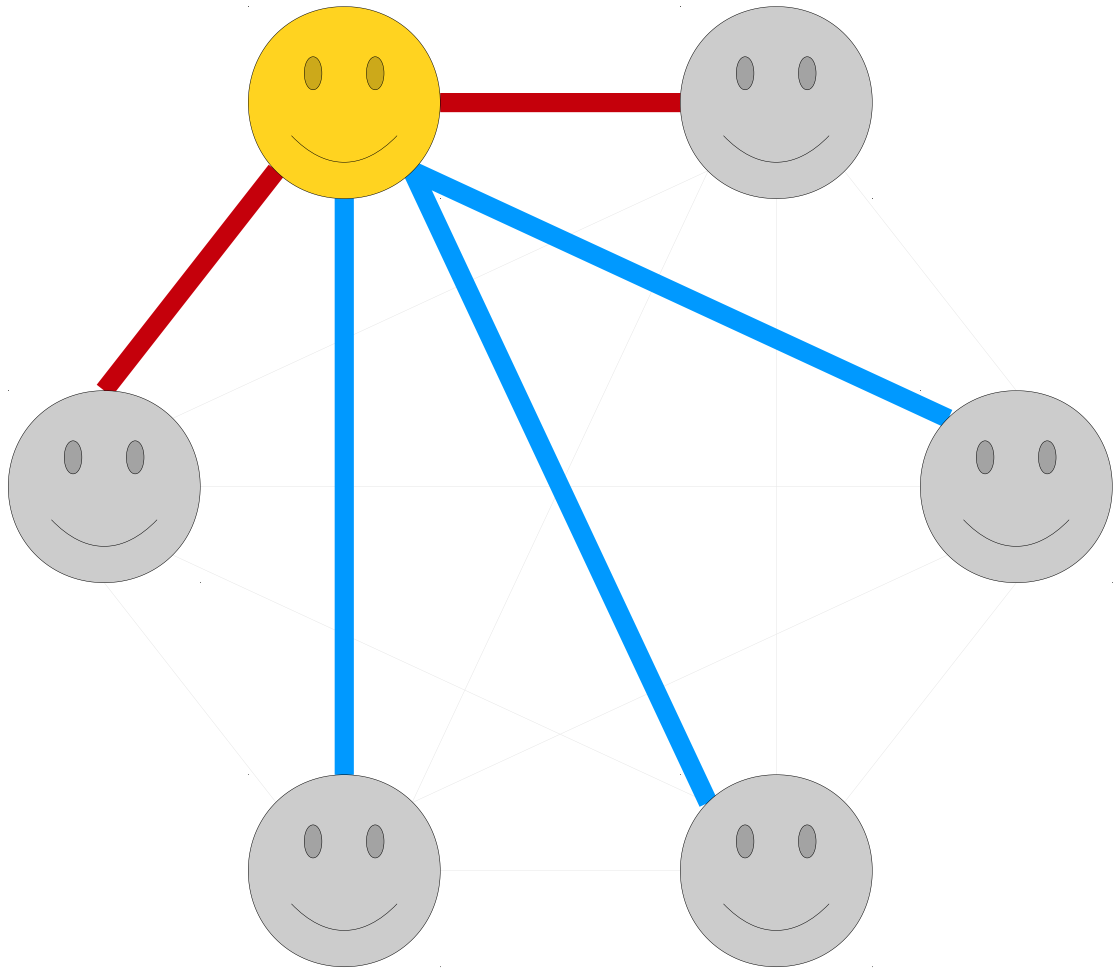

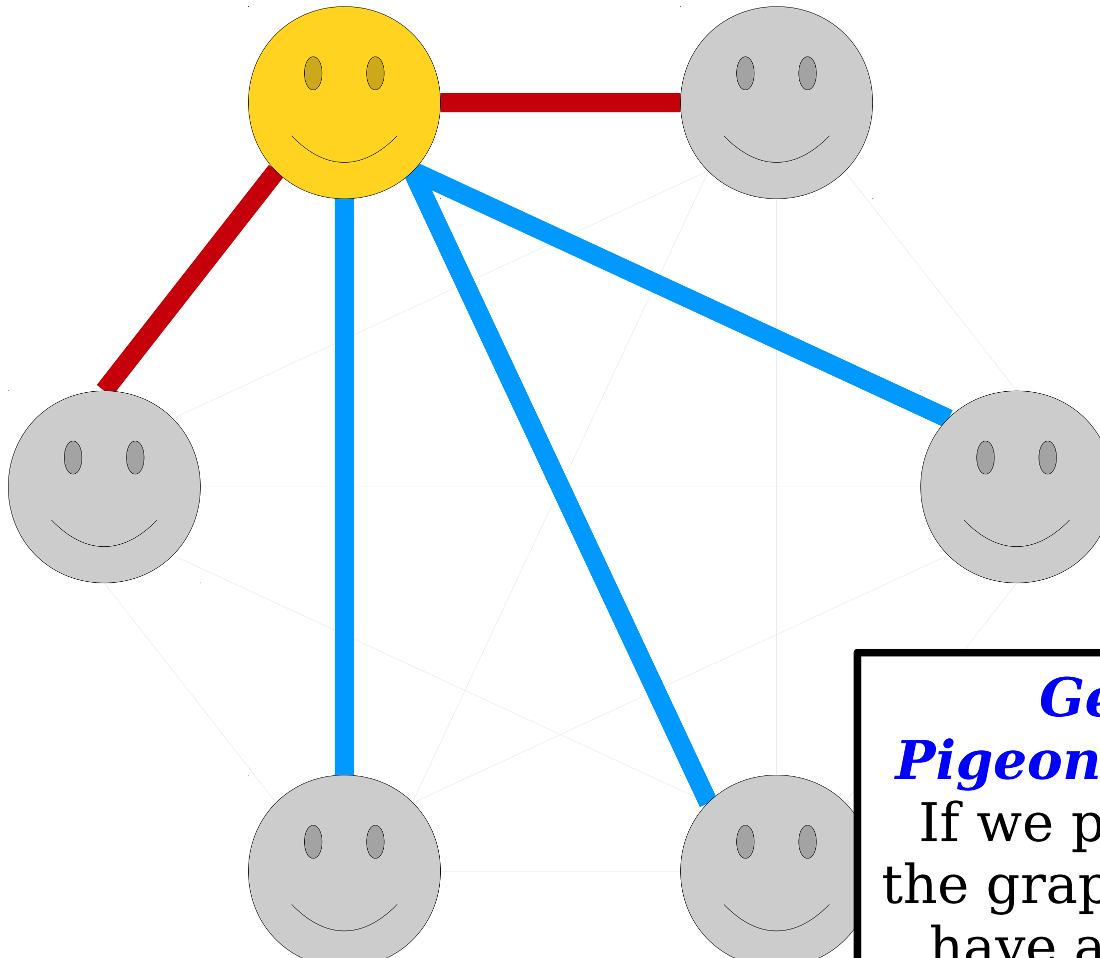

Generalized Pigeonhole Principle:

If we pick any node in the graph, that node will have at least $\lceil 5/2 \rceil = 3$ edges of the same color incident to it.

Question: if we reason through this with blue, will our analysis be any different than if we reasoned through with red?

PollEv.com/cs103spr26

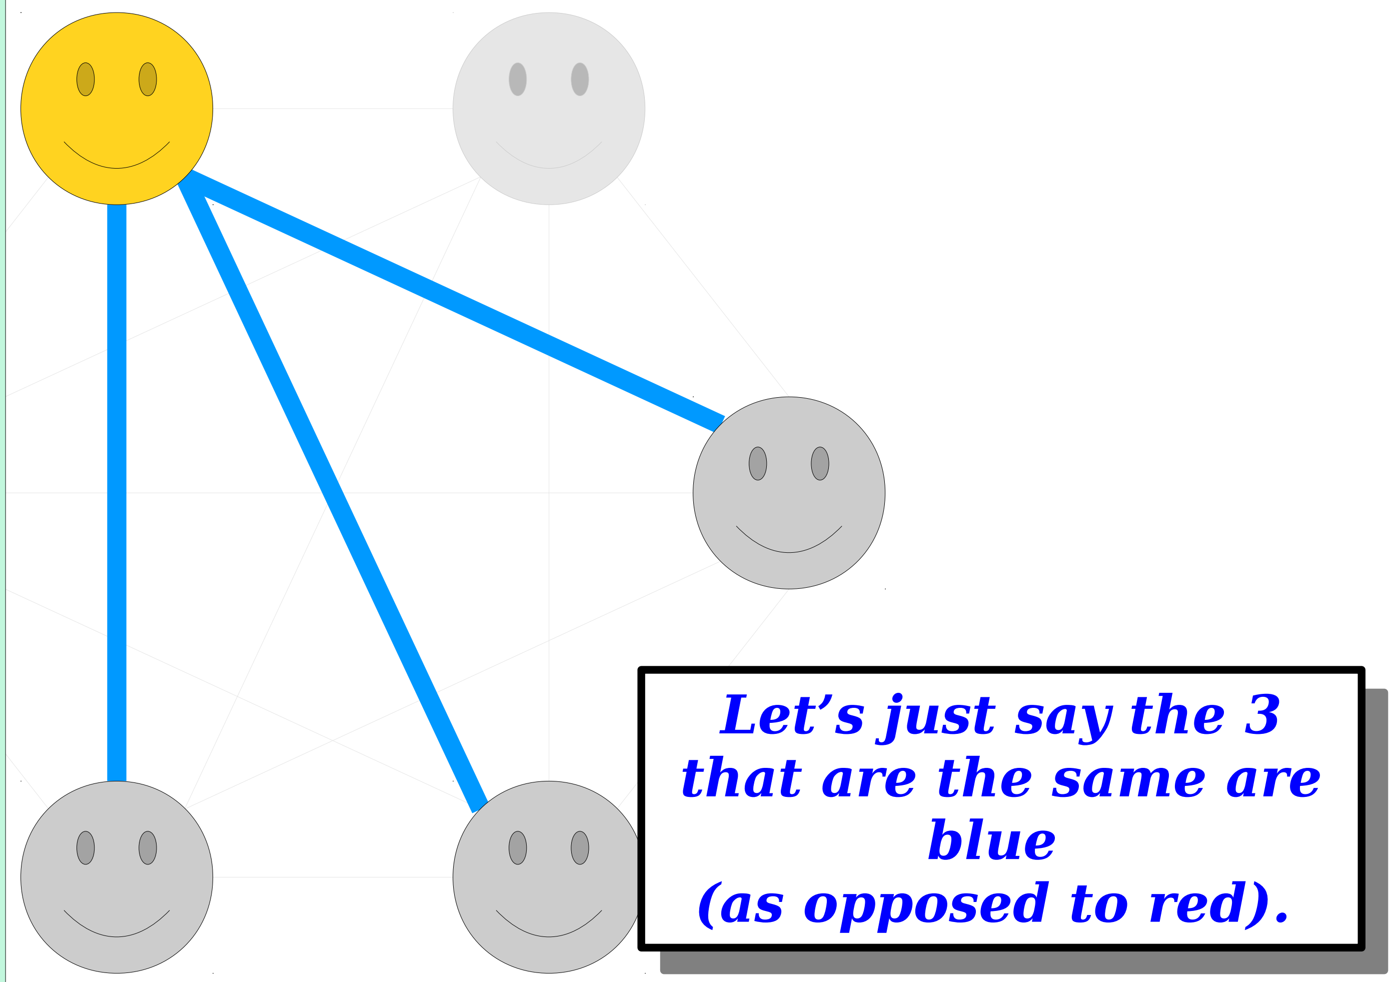

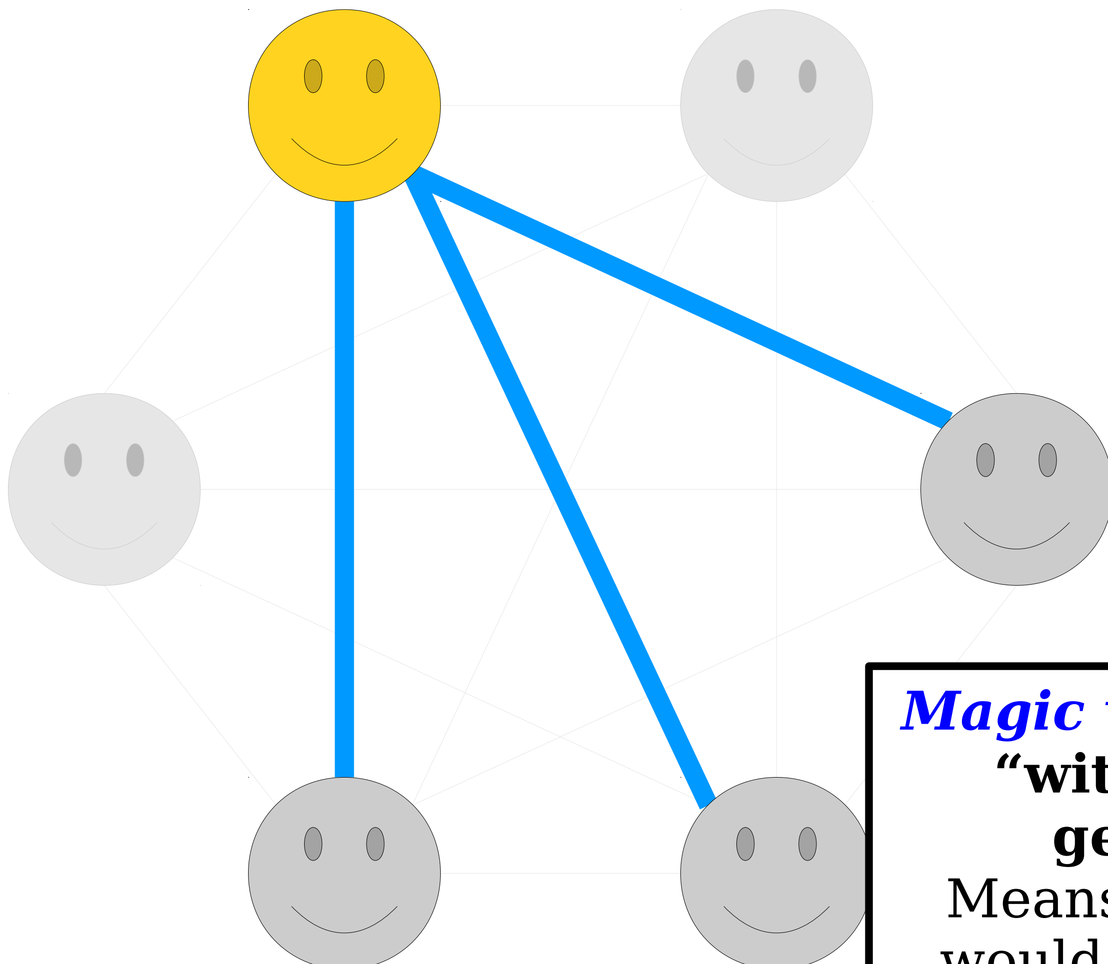

Magic words in proof:
“without loss of generality”
Means the argument would be the same in two different cases, so you just pick one and focus on that.

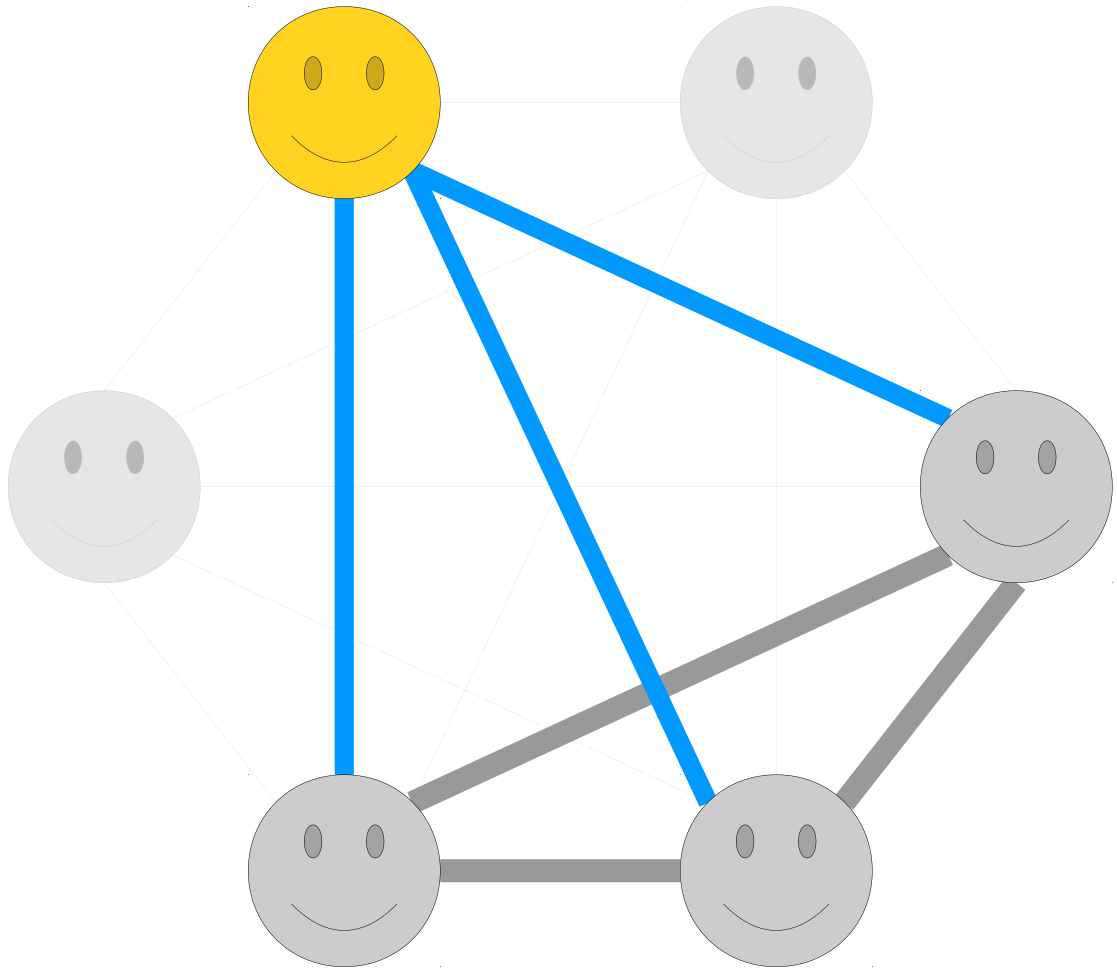

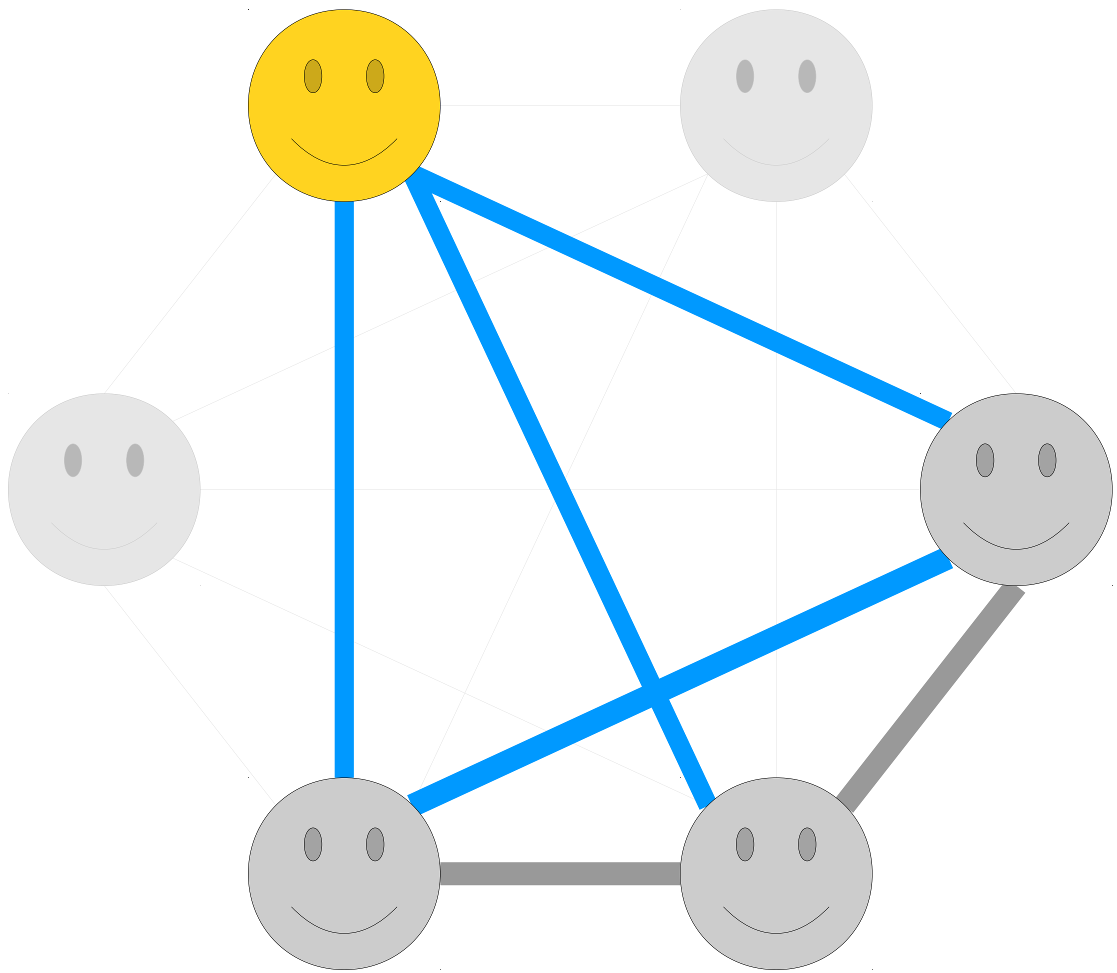

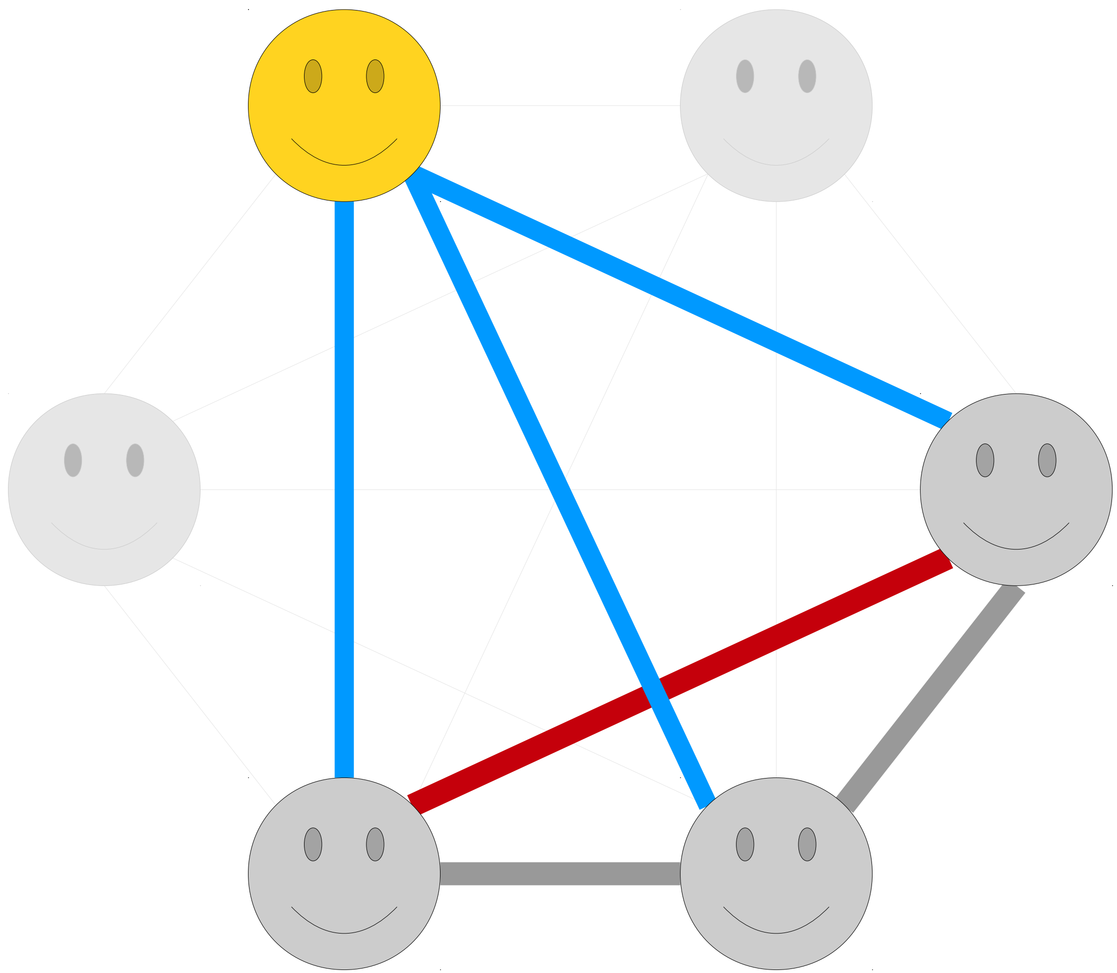

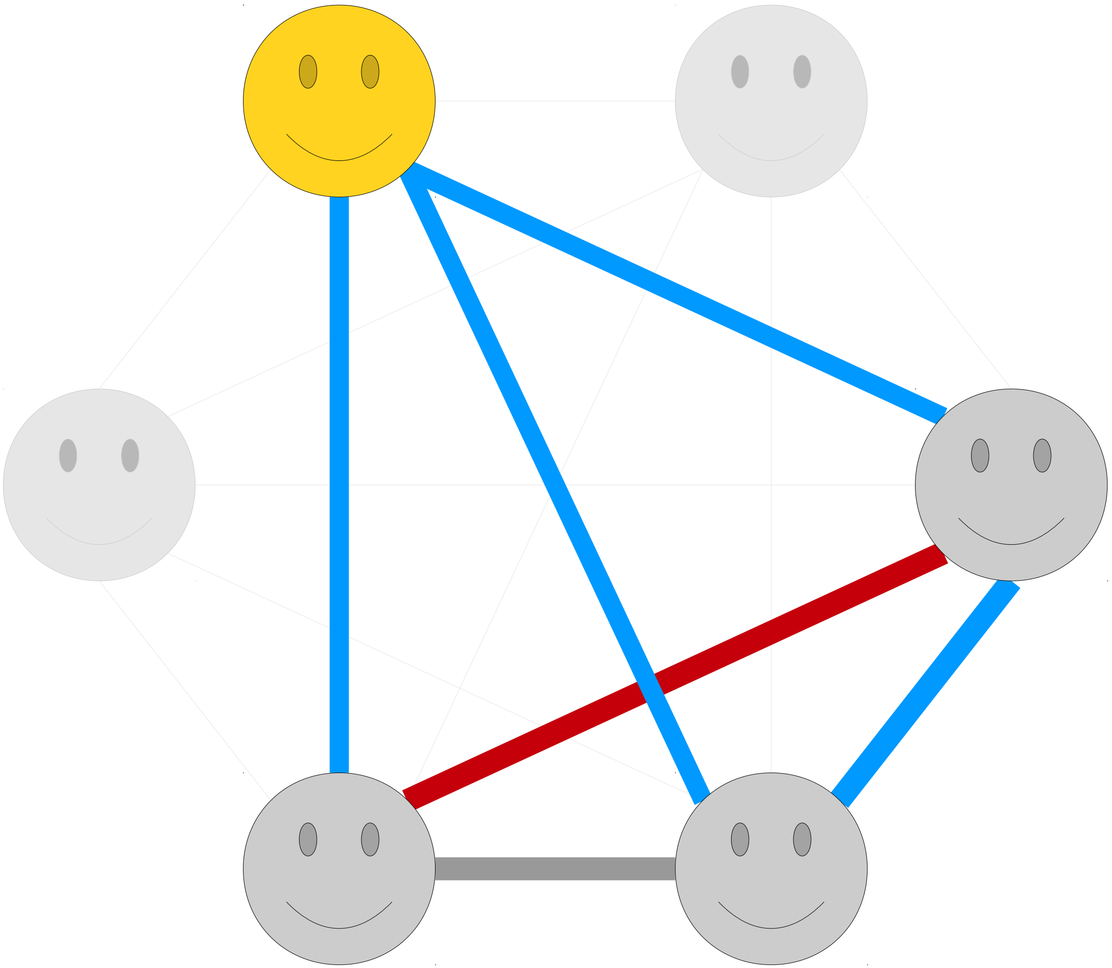

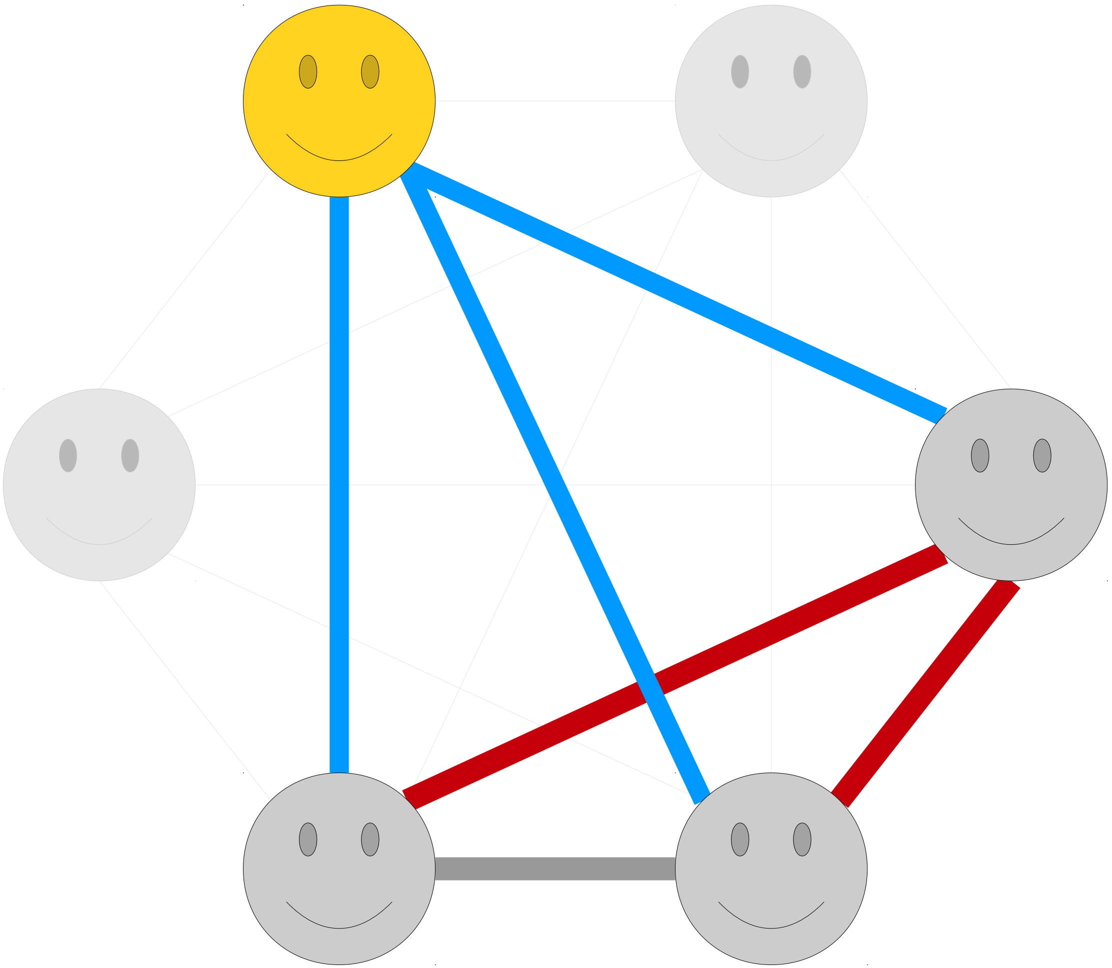

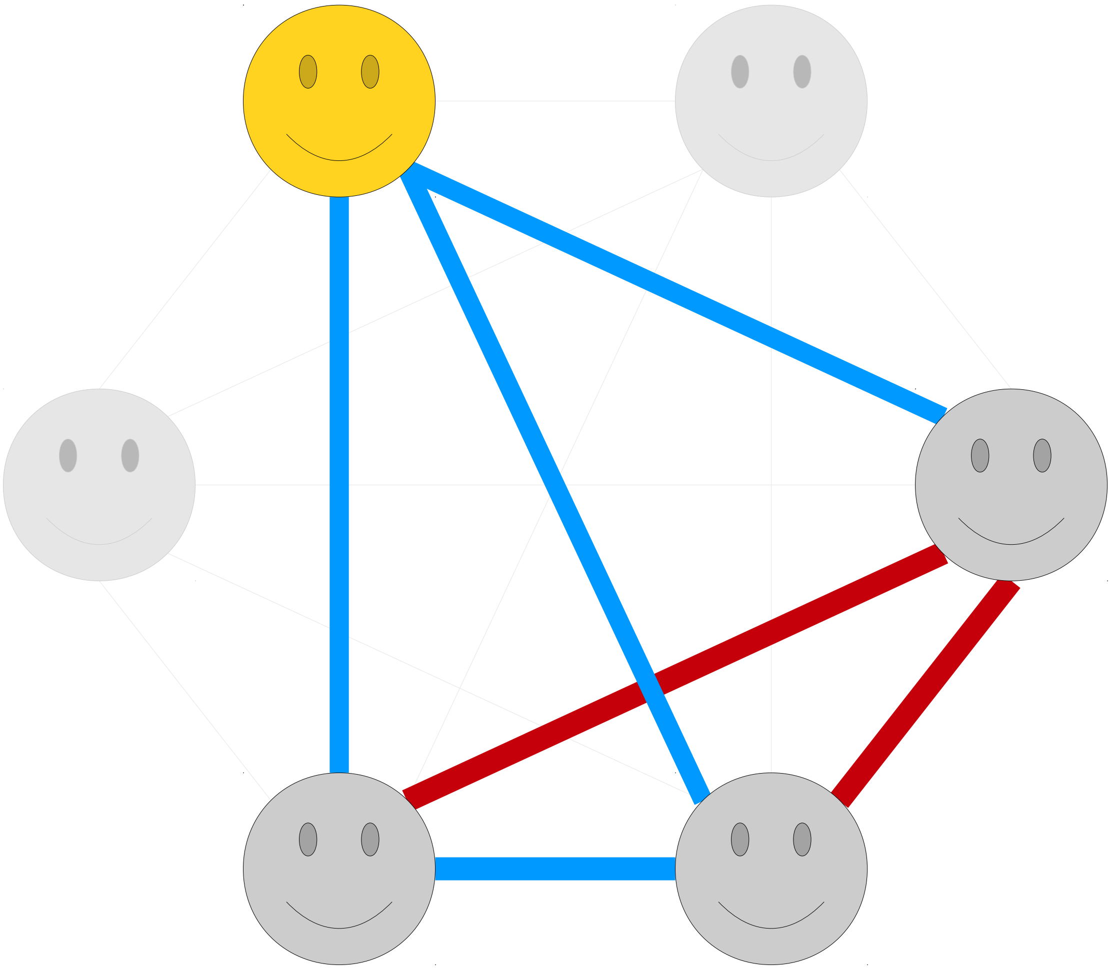

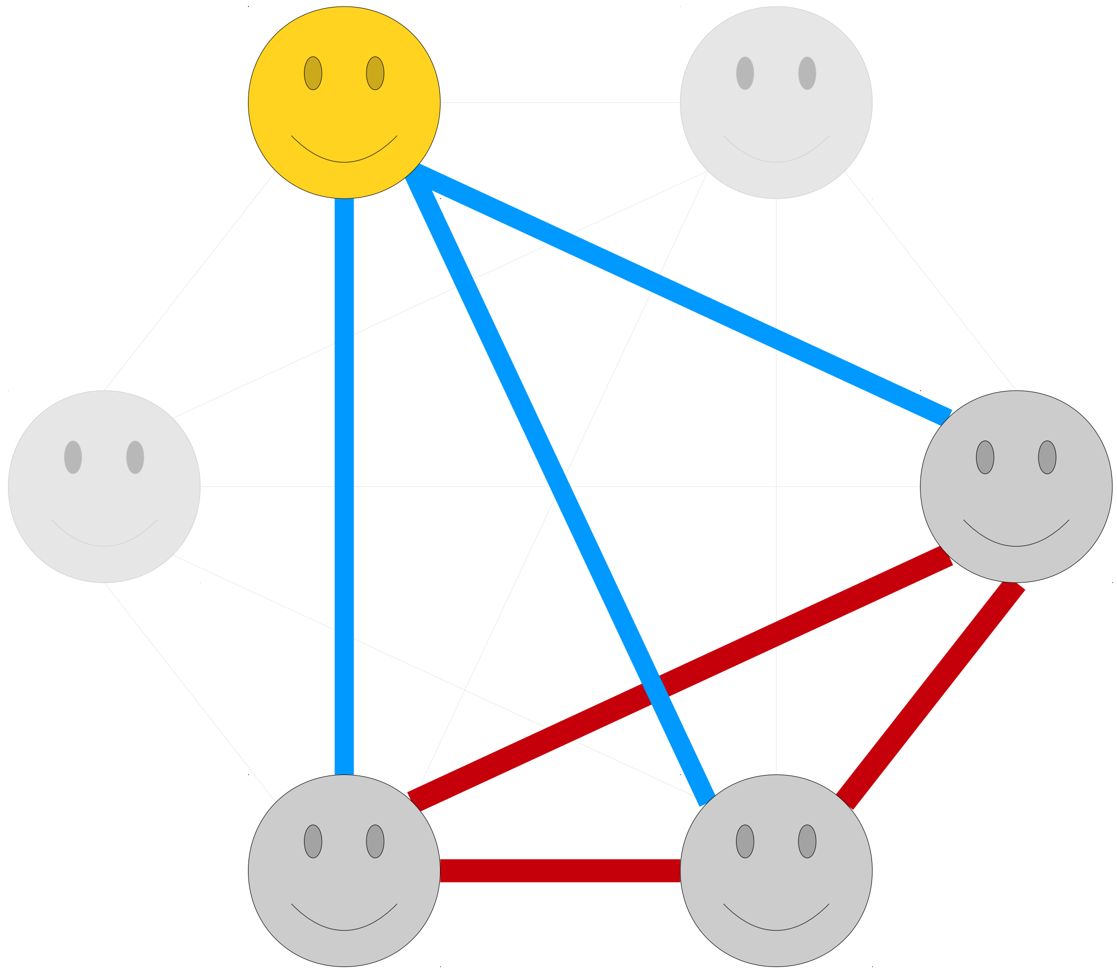

Theorem: Color each edge of K_6 red or blue. The resulting graph contains a monochrome copy of K_3.

Proof: We need to show that the colored K_6 contains a red copy of K_3 or a blue copy of K_3.

Pick some node x from K_6. It is incident to five edges and there are two possible colors for those edges. Therefore, by the generalized pigeonhole principle, at least $\lceil 5/2 \rceil = 3$ of those edges must be the same color. **Without loss of generality**, assume those edges are blue.

Let r, s, and t be three of the nodes adjacent to node x along a blue edge. If any of the edges $\{r, s\}$, $\{r, t\}$, or $\{s, t\}$ are blue, then one of those edges plus the two edges connecting back to node x form a blue K_3. Otherwise, all three of those edges are red, and they form a red K_3. Overall, this gives a red K_3 or a blue K_3, as required. ■

Ramsey Theory

- This proof is a special case of a broader family of results called ***Ramsey theory***.
- ***Theorem (Ramsey):*** For any natural number s, there is a number $R(s)$ such that
  - for all $n < R(s)$, there's a way to color the edges of K_n red and blue so there are no monochrome copies of K_s, and
  - for all $n \geq R(s)$, every way of coloring the edges of K_n red and blue always has a monochrome copy of K_s.
- Take Math 108 (combinatorics) to learn more!
- A more philosophical (and less literal) take on this theorem: true disorder is impossible at a large scale, since no matter how you organize things, you're guaranteed to find some interesting substructure.

The Game of Sim

- Here's a game you can play with two players.
  - One players plays as red, the other as blue.
  - Begin with six disconnected points.
  - Each turn, a player draws a line of their color.
  - The first to make a triangle of their color loses.
- The theorem we just proved means the game can't end in a draw: someone must win and someone must lose.
- The strategy is more subtle than it looks. Try playing this with a friend to see why!

A Little Math Puzzle

“In a group of $n > 0$ people …

- 90% of those people enjoyed *CODA*,
- 80% of those people enjoyed *Nomadland*,
- 70% of those people enjoyed *Parasite*, and
- 60% of those people enjoyed *Knives Out*.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?”

(Adapted from here.)

Other Pigeonhole-Type Results

If m objects are distributed into n boxes, then [condition] holds.

If m objects are distributed into n boxes, then some box is loaded to at least the average $^m/_n$, and some box is loaded to at most the average $^m/_n$.

If m objects are distributed into n boxes, then [condition] holds.

Theorem: If m objects are distributed into n bins, then there is a bin containing more than m/n objects if and only if there is a bin containing fewer than m/n objects.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Lemma: If m objects are distributed into n bins and there are no bins containing more than $^m/_n$ objects, then there are no bins containing fewer than $^m/_n$ objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than $^m/_n$ objects, yet some bin has fewer than $^m/_n$ objects.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$m = x_1 + x_2 + x_3 + \ldots + x_n$$

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< m/n + x_2 + x_3 + \ldots + x_n \end{aligned}$$

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$
\begin{aligned}
m &= x_1 + x_2 + x_3 + \ldots + x_n \\
&< m/n + x_2 + x_3 + \ldots + x_n \\
&\leq m/n + m/n + m/n + \ldots + m/n.
\end{aligned}
$$

Lemma: If m objects are distributed into n bins and there are no bins containing more than $^m/_n$ objects, then there are no bins containing fewer than $^m/_n$ objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than $^m/_n$ objects, yet some bin has fewer than $^m/_n$ objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than $^m/_n$ objects, meaning that $x_1 < {}^m/_n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< {}^m/_n + x_2 + x_3 + \ldots + x_n \\ &\leq {}^m/_n + {}^m/_n + {}^m/_n + \ldots + {}^m/_n. \end{aligned}$$

This third step follows because each remaining bin has at most $^m/_n$ objects.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< m/n + x_2 + x_3 + \ldots + x_n \\ &\leq m/n + m/n + m/n + \ldots + m/n. \end{aligned}$$

This third step follows because each remaining bin has at most m/n objects. Grouping the n copies of the m/n term here tells us that

$$\begin{aligned} m &< m/n + m/n + m/n + \ldots + m/n \\ &= m. \end{aligned}$$

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< m/n + x_2 + x_3 + \ldots + x_n \\ &\leq m/n + m/n + m/n + \ldots + m/n. \end{aligned}$$

This third step follows because each remaining bin has at most m/n objects. Grouping the n copies of the m/n term here tells us that

$$\begin{aligned} m &< m/n + m/n + m/n + \ldots + m/n \\ &= m. \end{aligned}$$

But this means $m < m$, which is impossible.

Lemma: If m objects are distributed into n bins and there are no bins containing more than m/n objects, then there are no bins containing fewer than m/n objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than m/n objects, yet some bin has fewer than m/n objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than m/n objects, meaning that $x_1 < m/n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< m/n + x_2 + x_3 + \ldots + x_n \\ &\leq m/n + m/n + m/n + \ldots + m/n. \end{aligned}$$

This third step follows because each remaining bin has at most m/n objects. Grouping the n copies of the m/n term here tells us that

$$\begin{aligned} m &< m/n + m/n + m/n + \ldots + m/n \\ &= m. \end{aligned}$$

But this means $m < m$, which is impossible. We've reached a contradiction, so our assumption was wrong, so if m objects are distributed into n bins and no bin has more than m/n objects, no bin has fewer than m/n objects either.

Lemma: If m objects are distributed into n bins and there are no bins containing more than $^m/_n$ objects, then there are no bins containing fewer than $^m/_n$ objects.

Proof: Assume for the sake of contradiction that m objects are distributed into n bins such that no bin contains more than $^m/_n$ objects, yet some bin has fewer than $^m/_n$ objects.

For simplicity, denote by x_i the number of objects in bin i. Without loss of generality, assume that bin 1 has fewer than $^m/_n$ objects, meaning that $x_1 < {}^m/_n$. Adding up the number of objects in each bin tells us that

$$\begin{aligned} m &= x_1 + x_2 + x_3 + \ldots + x_n \\ &< {}^m/_n + x_2 + x_3 + \ldots + x_n \\ &\leq {}^m/_n + {}^m/_n + {}^m/_n + \ldots + {}^m/_n. \end{aligned}$$

This third step follows because each remaining bin has at most $^m/_n$ objects. Grouping the n copies of the $^m/_n$ term here tells us that

$$\begin{aligned} m &< {}^m/_n + {}^m/_n + {}^m/_n + \ldots + {}^m/_n \\ &= m. \end{aligned}$$

But this means $m < m$, which is impossible. We've reached a contradiction, so our assumption was wrong, so if m objects are distributed into n bins and no bin has more than $^m/_n$ objects, no bin has fewer than $^m/_n$ objects either. ■

“In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?”

Insight 1: Model movie preferences as balls (movies) in bins (people).

Insight 2: There are n total bins, one for each person.

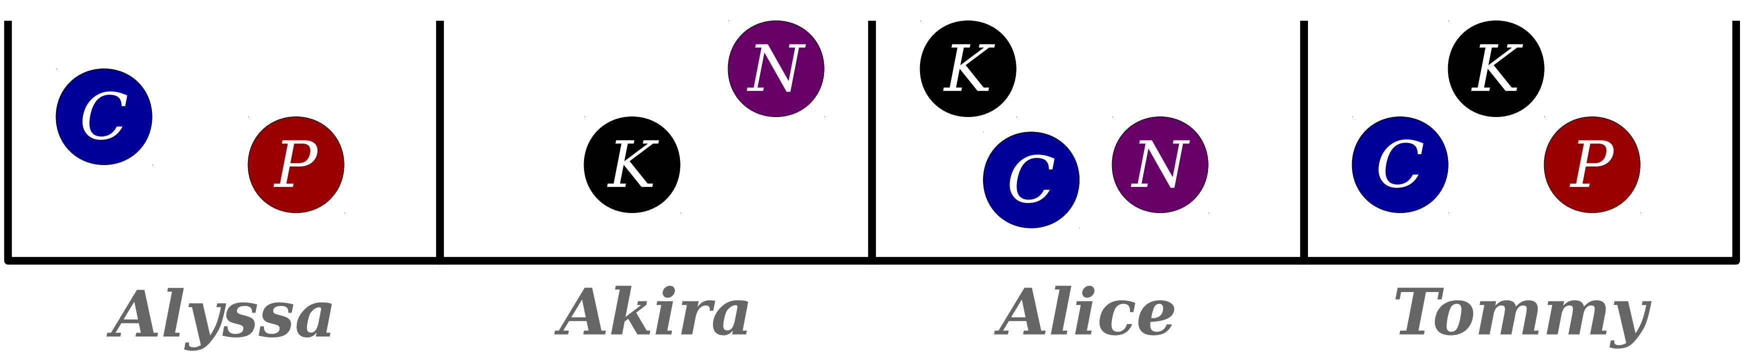

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

$$.9n + .8n + .7n + .6n$$
$$= 3n$$

Insight 3: There are 3*n* balls being distributed into *n* bins.

Insight 4: The average number of balls in each bin is 3.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Insight 5: No one enjoyed more than three movies…

Insight 6: … so no one enjoyed fewer than three movies …

Insight 7: … so everyone enjoyed exactly three movies.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Insight 8: You have to enjoy at least one of these movies to enjoy three of the four movies.

Conclusion: **Everyone** liked at least one of these two movies!

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins. Since no person liked all four movies, no bin contains more than $3 = {}^{3n}/_{n}$ balls, so by our earlier theorem we see that no bin contains fewer than three balls.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins. Since no person liked all four movies, no bin contains more than $3 = {}^{3n}/_{n}$ balls, so by our earlier theorem we see that no bin contains fewer than three balls. Therefore, each bin contains exactly three balls.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins. Since no person liked all four movies, no bin contains more than $3 = {}^{3n}/_{n}$ balls, so by our earlier theorem we see that no bin contains fewer than three balls. Therefore, each bin contains exactly three balls.

Now suppose for the sake of contradiction that someone didn't enjoy *CODA* and didn't enjoy *Parasite*.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins. Since no person liked all four movies, no bin contains more than $3 = {}^{3n}/_{n}$ balls, so by our earlier theorem we see that no bin contains fewer than three balls. Therefore, each bin contains exactly three balls.

Now suppose for the sake of contradiction that someone didn't enjoy *CODA* and didn't enjoy *Parasite*. This means they could enjoy at most two of the four movies, contradicting that each person enjoys exactly three.

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Theorem: In the scenario described here, all n people enjoyed at least one of *CODA* and *Parasite*.

Proof: Suppose there is a group of n people meeting these criteria. We can model this problem by representing each person as a bin and each time a person enjoys a movie as a ball. The number of balls is

$$.9n + .8n + .7n + .6n = 3n,$$

and since there are n people, there are n bins. Since no person liked all four movies, no bin contains more than $3 = {}^{3n}/_{n}$ balls, so by our earlier theorem we see that no bin contains fewer than three balls. Therefore, each bin contains exactly three balls.

Now suppose for the sake of contradiction that someone didn't enjoy *CODA* and didn't enjoy *Parasite*. This means they could enjoy at most two of the four movies, contradicting that each person enjoys exactly three.

We've reached a contradiction, so our assumption was wrong and each person enjoyed at least one of *CODA* and *Parasite*. ■

"In a group of $n > 0$ people …

- 90% of those people enjoyed ***CODA***,
- 80% of those people enjoyed ***Nomadland***,
- 70% of those people enjoyed ***Parasite***, and
- 60% of those people enjoyed ***Knives Out***.

No one enjoyed all four movies. How many people enjoyed at least one of *CODA* and *Parasite*?"

Going Further

- The pigeonhole principle can be used to prove a *ton* of amazing theorems. Here's a sampler:
  - There is always a way to fairly split rent among multiple people, even if different people want different rooms. *(Sperner's lemma)*
  - You and a friend can climb any mountain from two different starting points so that the two of you maintain the same altitude at each point in time. *(Mountain-climbing theorem)*
  - If you model coffee in a cup as a collection of infinitely many points and then stir the coffee, some point is always where it initially started. *(Brower's fixed-point theorem)*
  - A complex process that doesn't parallelize well must contain a large serial subprocess. *(Mirksy's theorem)*
  - Any positive integer *n* has a nonzero multiple that can be written purely using the digits 1 and 0. *(Doesn't have a name, but still cool!)*

More to Explore

- Interested in more about graphs and the pigeonhole principle? Check out…
  - … ***Math 107*** (Graph Theory), a deep dive into graph theory.
  - … ***Math 108*** (Combinatorics), which explores a bunch of results pertaining to graphs and counting things.
  - … ***CS161*** (Algorithms), which explores algorithms for computing important properties of graphs.
  - … ***CS224W*** (Deep Learning on Graphs), which uses a mix of mathematical and statistical techniques to explore graphs.
- Happy to chat about this in person if you'd like.

Next Time

- ***Mathematical Induction***
  - Reasoning about stepwise processes

- ***Applications of Induction***
  - To numbers
  - To anticounterfeiting
  - To modern art