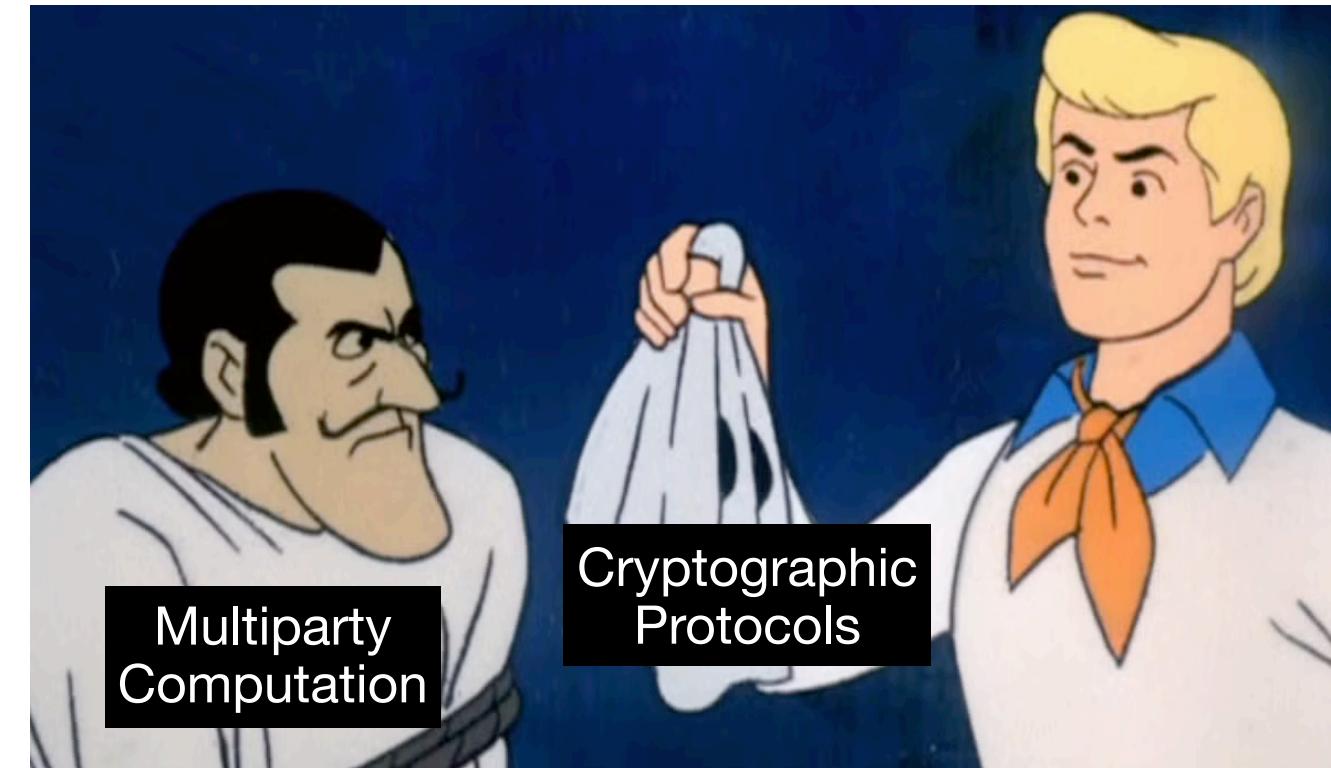# Last time in Cryptographic Protocols

1. Talked about what cryptography is.

2. Unmasked the main character of this course.

3. Listed some real world problems where mutual distrust can be resolved by a trusted third party.

4. Our goal is to emulate the trusted party using a protocol.

CS4501 Cryptographic Protocols
Lecture 2: Adversaries and Simulation

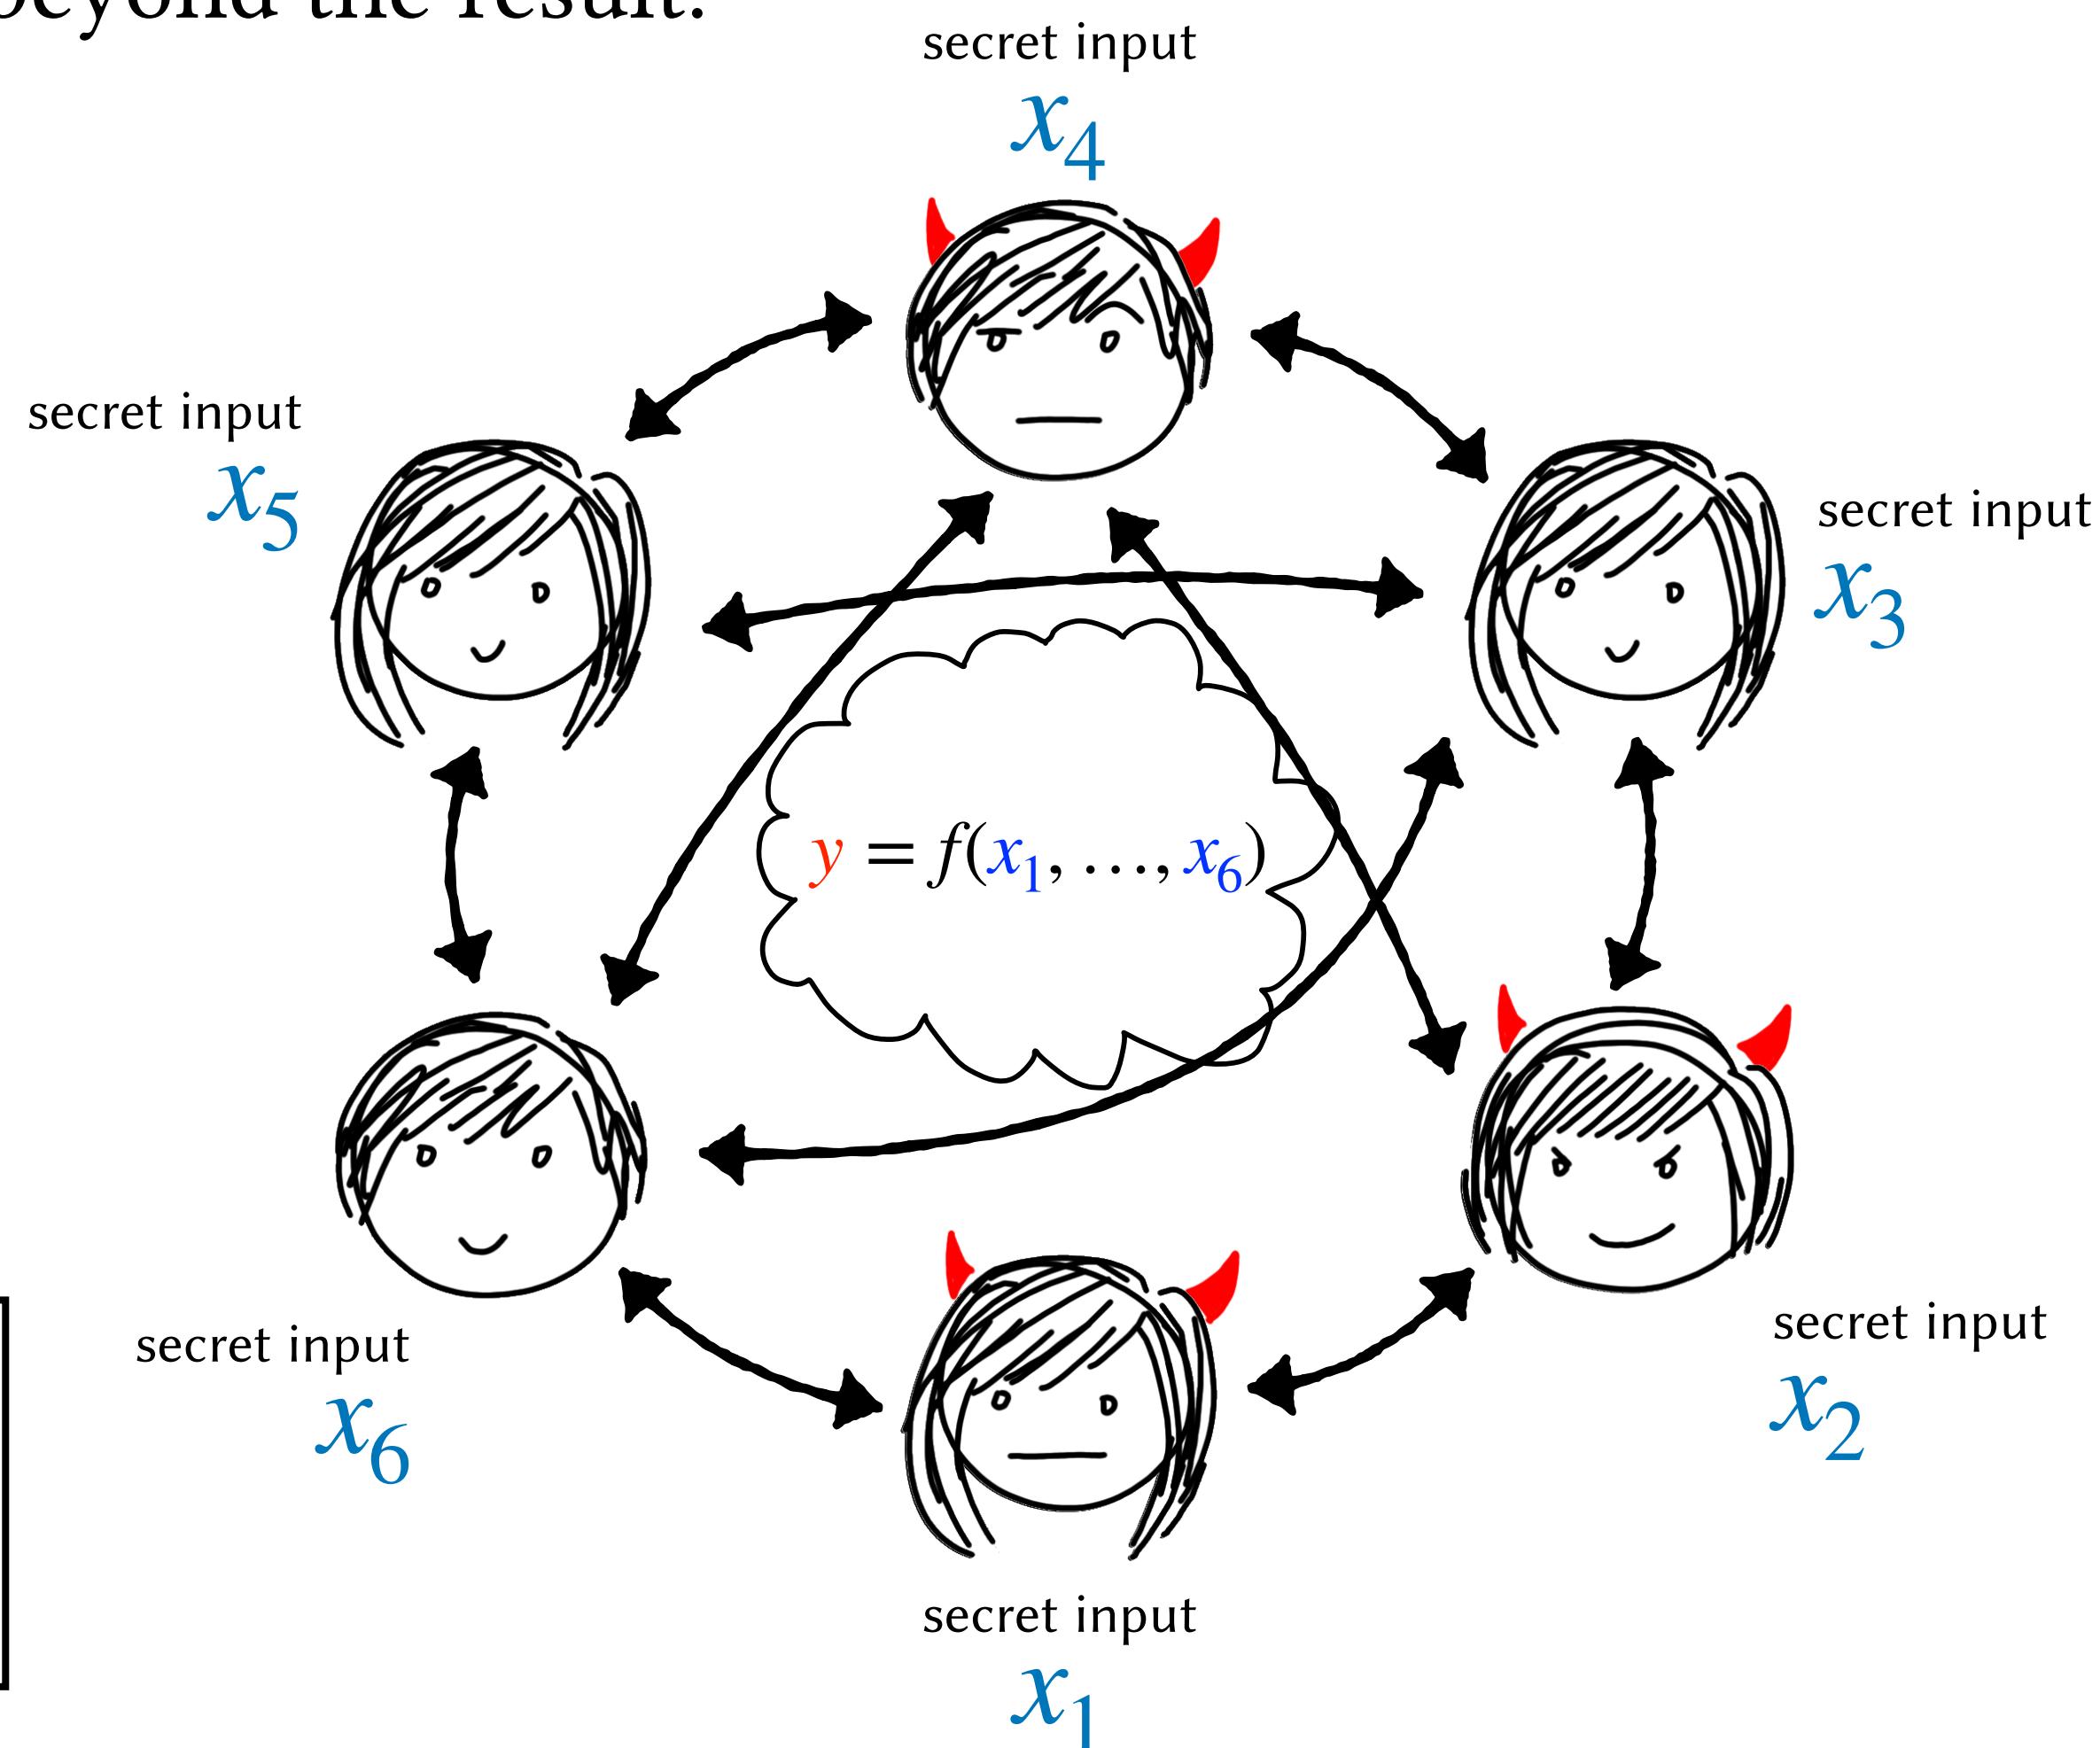https://jackdoerner.net/teaching/#2026/Spring/CS4501

Let's look at MPC another way

Secure Multiparty Computation means **jointly computing** on secret data, while **revealing nothing** about the data beyond the result.

In this example, only y is revealed
The internal workings of f and the
secret inputs $\{x_1, \ldots, x_6\}$ remain hidden

This must be true even if some
parties misbehave in an attempt to
learn more than they should!

Clarification: the code of f is public,
but we don't the internal state it assumes
when it runs on inputs $\{x_1, \ldots, x_6\}$

secret input
x_4

secret input
x_5

secret input
x_3

$y = f(x_1, \ldots, x_6)$

secret input
x_6

secret input
x_2

secret input
x_1

What does it mean to "reveal nothing"

Remember, a protocol includes local instructions and communication. That is, *messages.*

secret input
x_1

output
$y = f(x_1, x_2)$

secret input
x_2
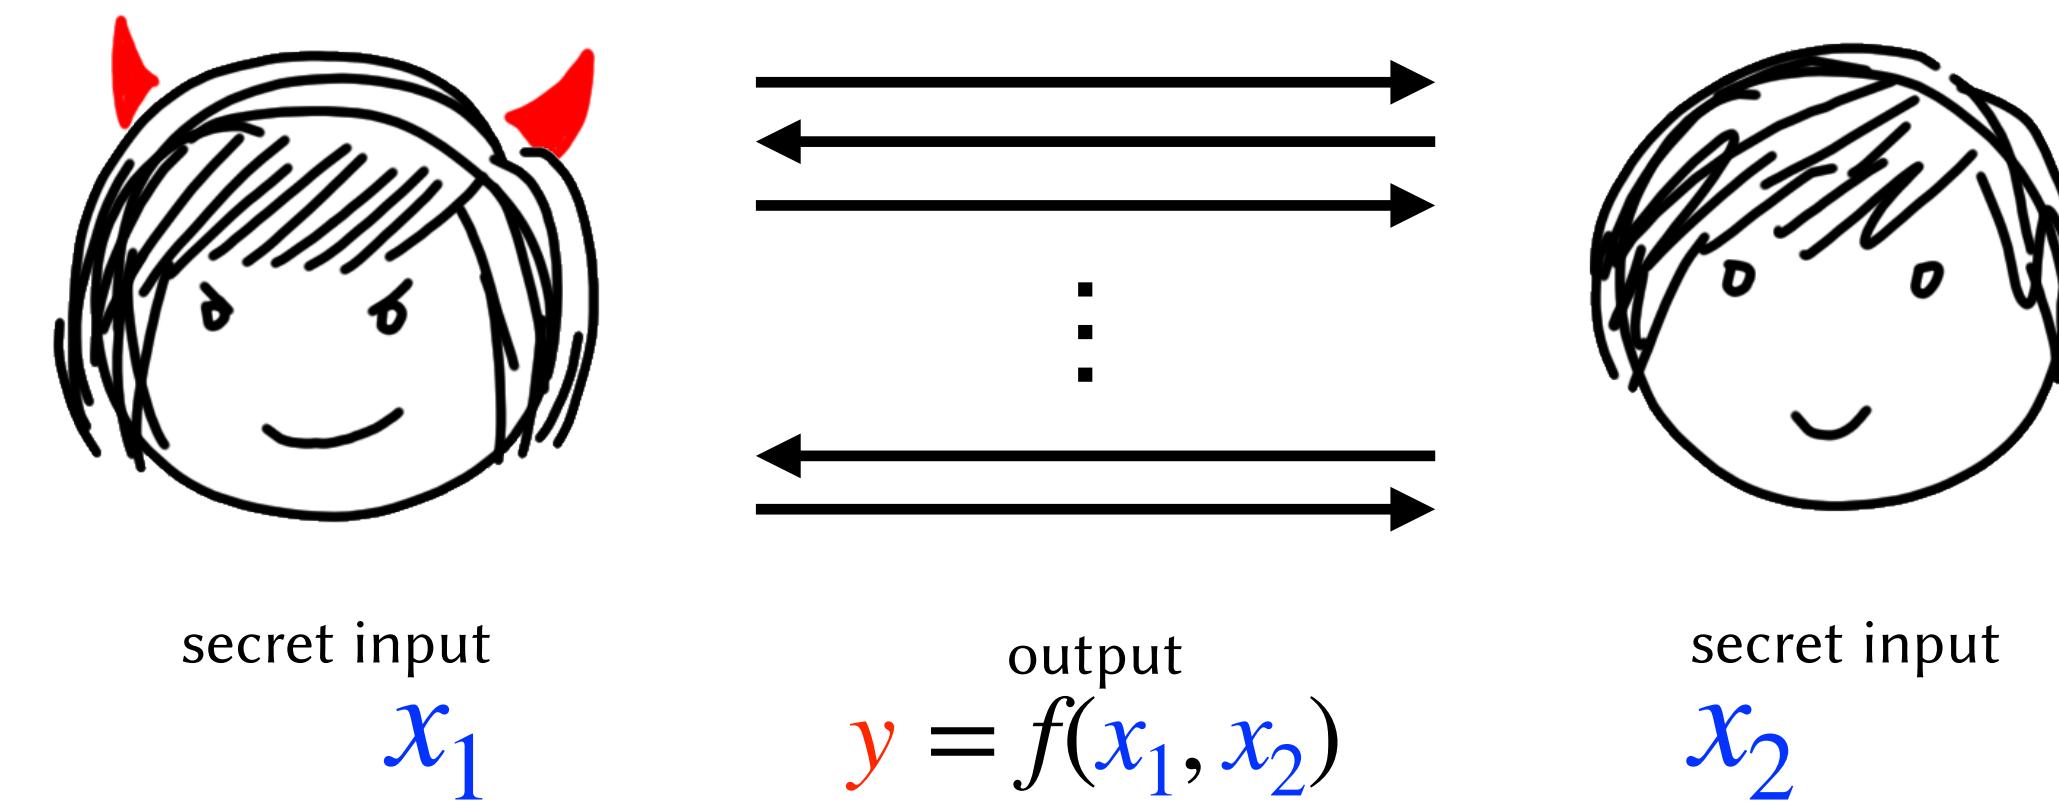
Q: What is revealed to Alice about Bob's input x_2?

A: Whatever information the messages and y convey!

Q: What does it mean that the messages reveal nothing about x_2, beyond y?

A: Using only x_1 and y, Alice could generate messages by herself that are *indistinguishable* from the messages she generated interacting with Bob.

Let's Run a Protocol Together

I need two volunteers.

I am going to embarrass you.

♣♥ Matchmaking ♥♣
(how to go on a date with a cryptographer)

I promise the rest of the class won't be like this.

Does this protocol produce a correct result?

The function we're computing is AND

Yes = ♣♥
No = ♥♣

Yes = ♥♣
No = ♣♥

		Cards (Before Random Rotation)	Result
Yes	Yes	♣♥♥♥♣	
Yes	No	♣♥♥♣♥	
No	Yes	♥♣♥♥♣	
No	No	♥♣♥♣♥	

Does this protocol produce a correct result?

The function we're computing is AND

Yes = ♣♥
No = ♥♣

Yes = ♥♣
No = ♣♥

		Cards (Before Random Rotation)	Result
Yes	Yes	♣♥♥♥♣	Yes
Yes	No	♣♥♥♣♥	
No	Yes	♥♣♥♥♣	
No	No	♥♣♥♣♥	

3 ♥ together

Does this protocol produce a correct result?

The function we're computing is AND

Yes = ♣♥
No = ♥♣

Yes = ♥♣
No = ♣♥

		Cards (Before Random Rotation)	Result
Yes	Yes	♣♥♥♥♣	Yes
Yes	No	♣♥♥♣♥	No
No	Yes	♥♣♥♥♣	No
No	No	♥♣♥♣♥	No

2 ♥ together

What can 🧑 and 🧑 learn from this?

Yes = ♣♥
No = ♥♣

Yes = ♥♣
No = ♣♥

🧑	🧑	Cards (Before Random Rotation)	Result
Yes	Yes	♣♥♥♥♣	Yes
Yes	No	♣♥♥♣♥	No
No	Yes	♥♣♥♥♣	No
No	No	♥♣♥♣♥	No

Rotations of Each Other
Indistinguishable after random rotation!

What can 👩 and 👦 learn from this?

Yes = ♣♥
No = ♥♣

Yes = ♥♣
No = ♣♥

👩	👦	Cards (Before Random Rotation)	Result
Yes	Yes	♣♥♥♥♣	Yes
Yes	No	♣♥♥♣♥	No
No	Yes	♥♣♥♥♣	No
No	No	♥♣♥♣♥	No

As promised, they learn only the result.
We proved that the protocol is *secure.*

What can 🧑 and 🧑 learn from this?

Question 1: what happens if 😈 doesn't rotate the deck randomly?

Question 2: what happens if 😈 reorders the cards instead of rotating.

Question 3: what happens if 😈 inputs YES regardless of her intention.

Why doesn't the last one violate the security of the protocol?

A More Nuanced View of Security

The Setting

- n parties, $P_1, \ldots, P_n$. You can just think of them as computers. Technically modeled as Interactive Turing Machines (ITMs).

- The parties wish to compute a known function $y = f(x_1, \ldots, x_n)$.

- Each P_i has private input x_i. For now they all want the same output y.

- The protocol performing the computation must achieve certain **security properties**, even if some parties collude to attack the protocol.

- This is modeled by introducing an **external adversary** $\mathcal{A}$ that **corrupts** some parties and controls their actions.
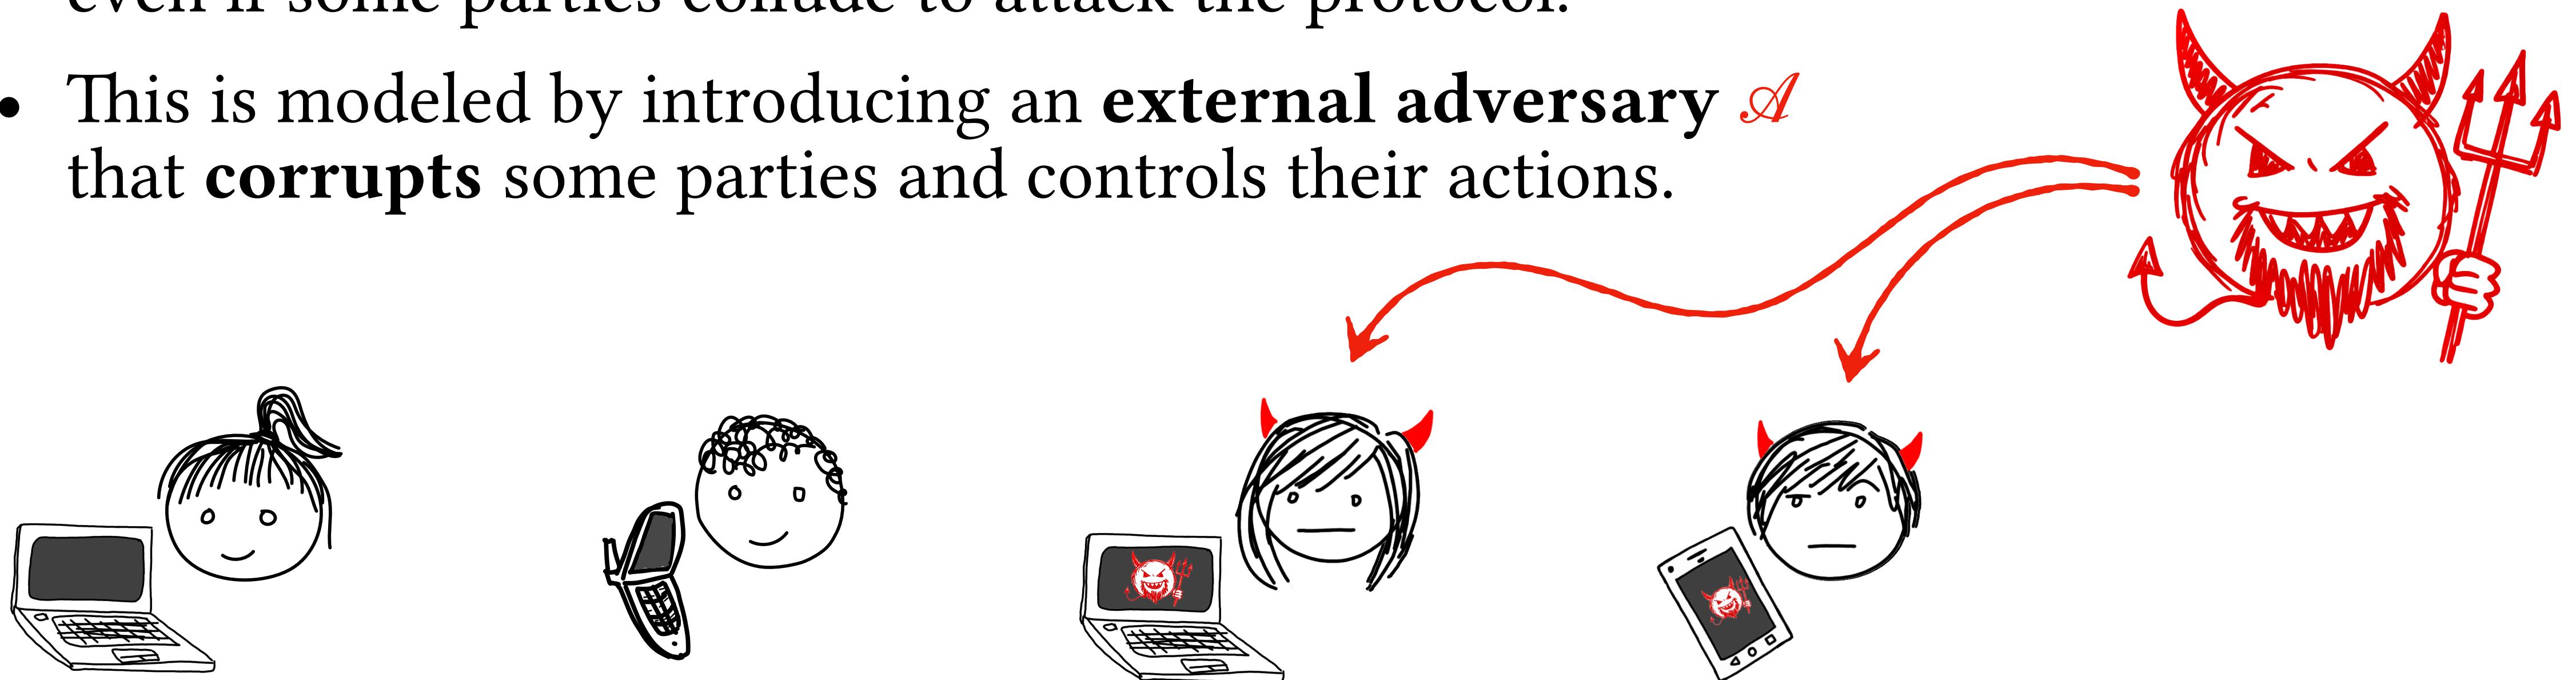
Security Properties (Auction Example)

- **Correctness:** $\mathcal{A}$ can't bid using a bid that isn't the highest bid.

- **Privacy:** $\mathcal{A}$ learns the value of the highest bid, the identity of the highest bidder, and nothing else.

- **Independence of Inputs:** $\mathcal{A}$ cannot bid $1 more than the highest honest bid.

- **Fairness:** $\mathcal{A}$ can't cause the protocol to abort if its bid isn't the highest (i.e. after learning the output).

- **Guaranteed Output Delivery:** $\mathcal{A}$ can't cause the protocol to abort at all (this is stronger than fairness and eliminates DoS attacks).

Security Properties (More Generally)

- **Correctness**: Corrupt parties can only influence the output by choosing their inputs.

- **Privacy**: Only the output is learned (nothing else).

- **Independence of Inputs**: Parties cannot choose their inputs as functions of other parties' inputs.

- **Fairness**: If one party learns the output, then all parties do.

- **Guaranteed Output Delivery**: All honest parties learn the output.

These aren't *formal* or *rigorous* definitions.
We might not want all of them all of the time.
And we could imagine many *more* security properties!
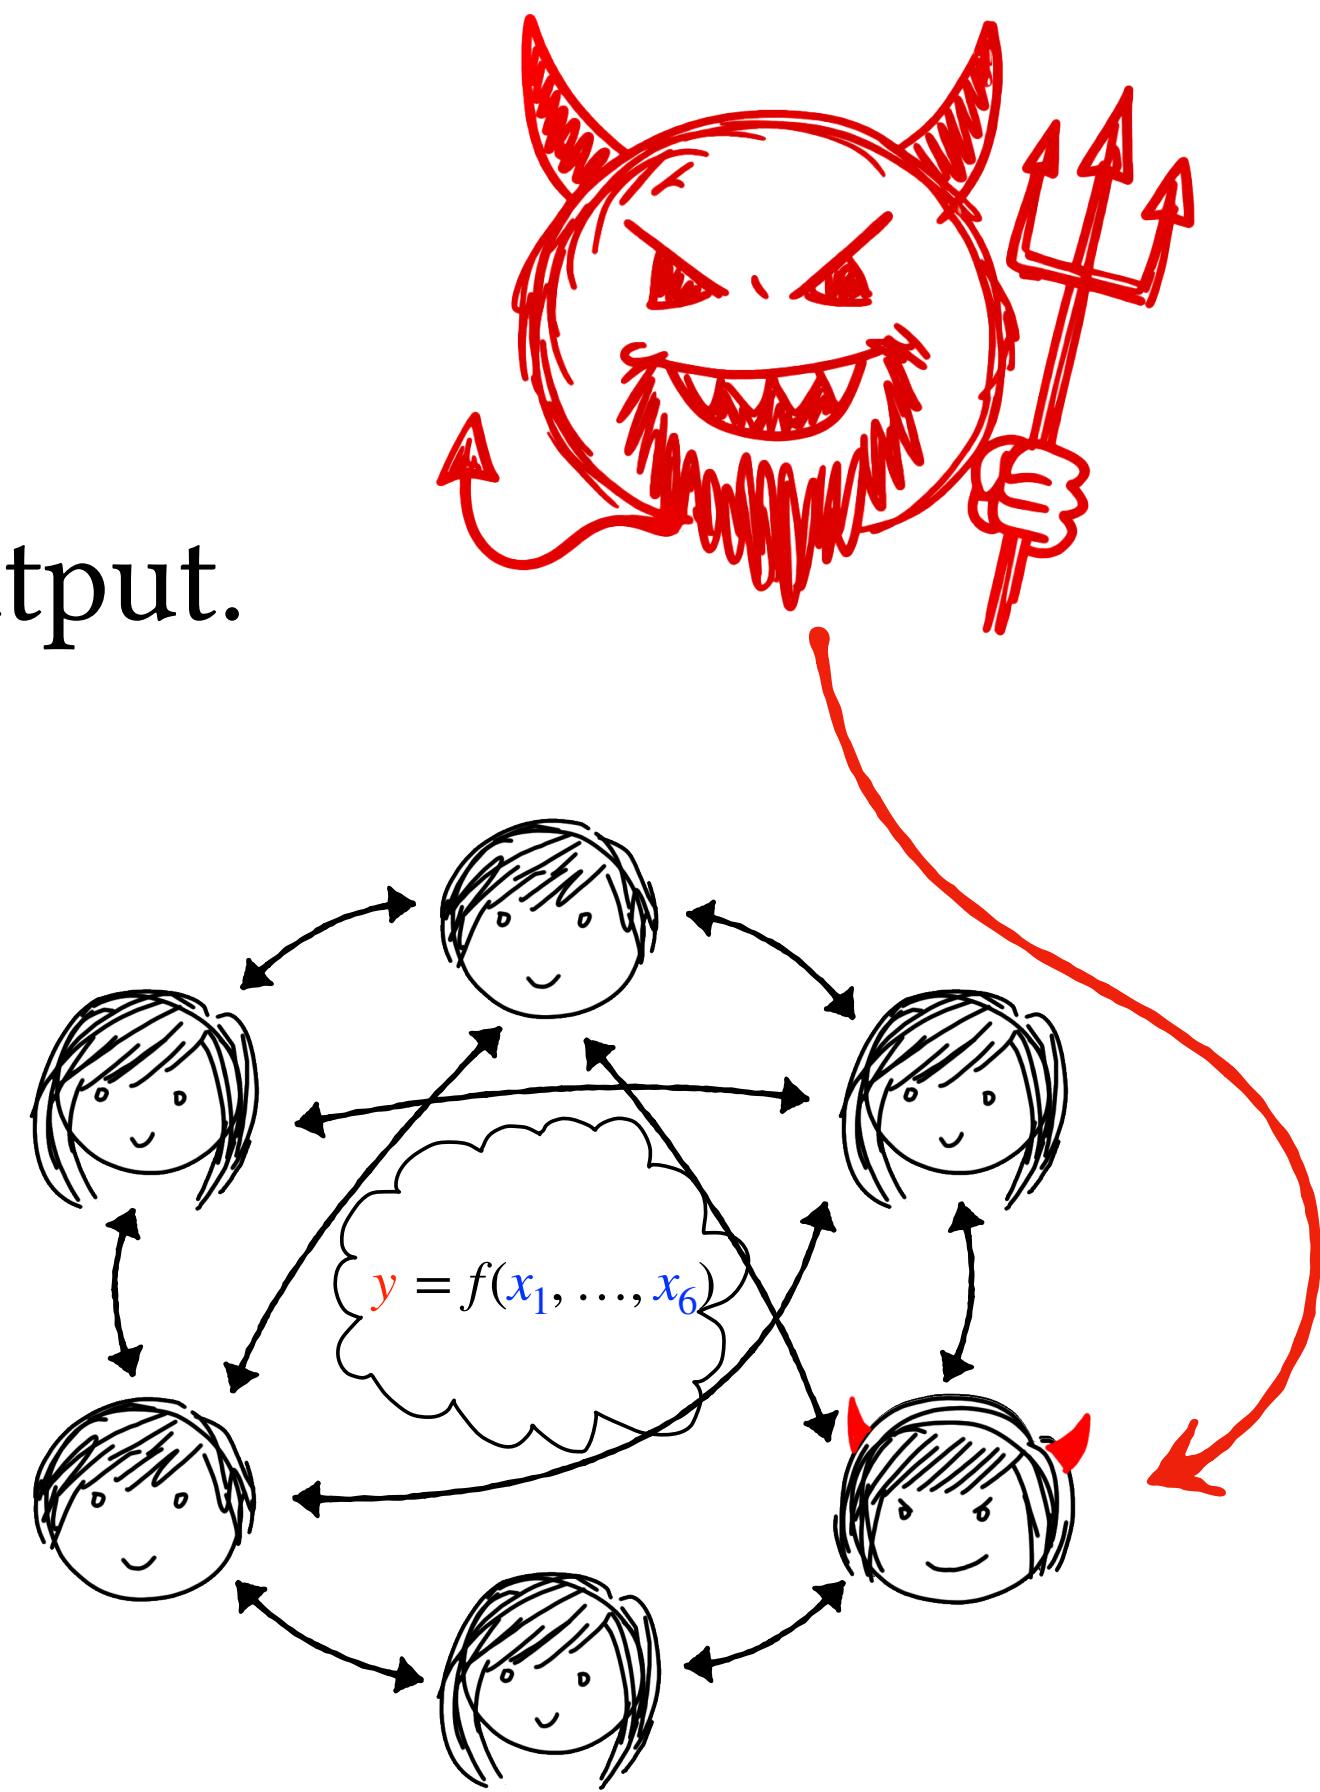
$y = f(x_1, \ldots, x_6)$

Security Properties (More Generally)

- **Correctness**: Corrupt parties can only influence the output by choosing their inputs.

- **Privacy**: Only the output is learned (nothing else).

- **Independence of Inputs**: Parties cannot choose their inputs as functions of other parties' inputs.

- **Fairness**: If one party learns the output, then all parties do.

- **Guaranteed Output Delivery**: All honest parties learn the output.

So far I have told you that *security* is related to $\mathcal{A}$ being able to emulate the protocol messages on it's own. It's intuitive how that's related to privacy, but can it possibly imply *guaranteed output* or even *correctness*?

We need to revisit our definition of security!

There's something else missing from the picture too…

$y = f(x_1, ..., x_6)$
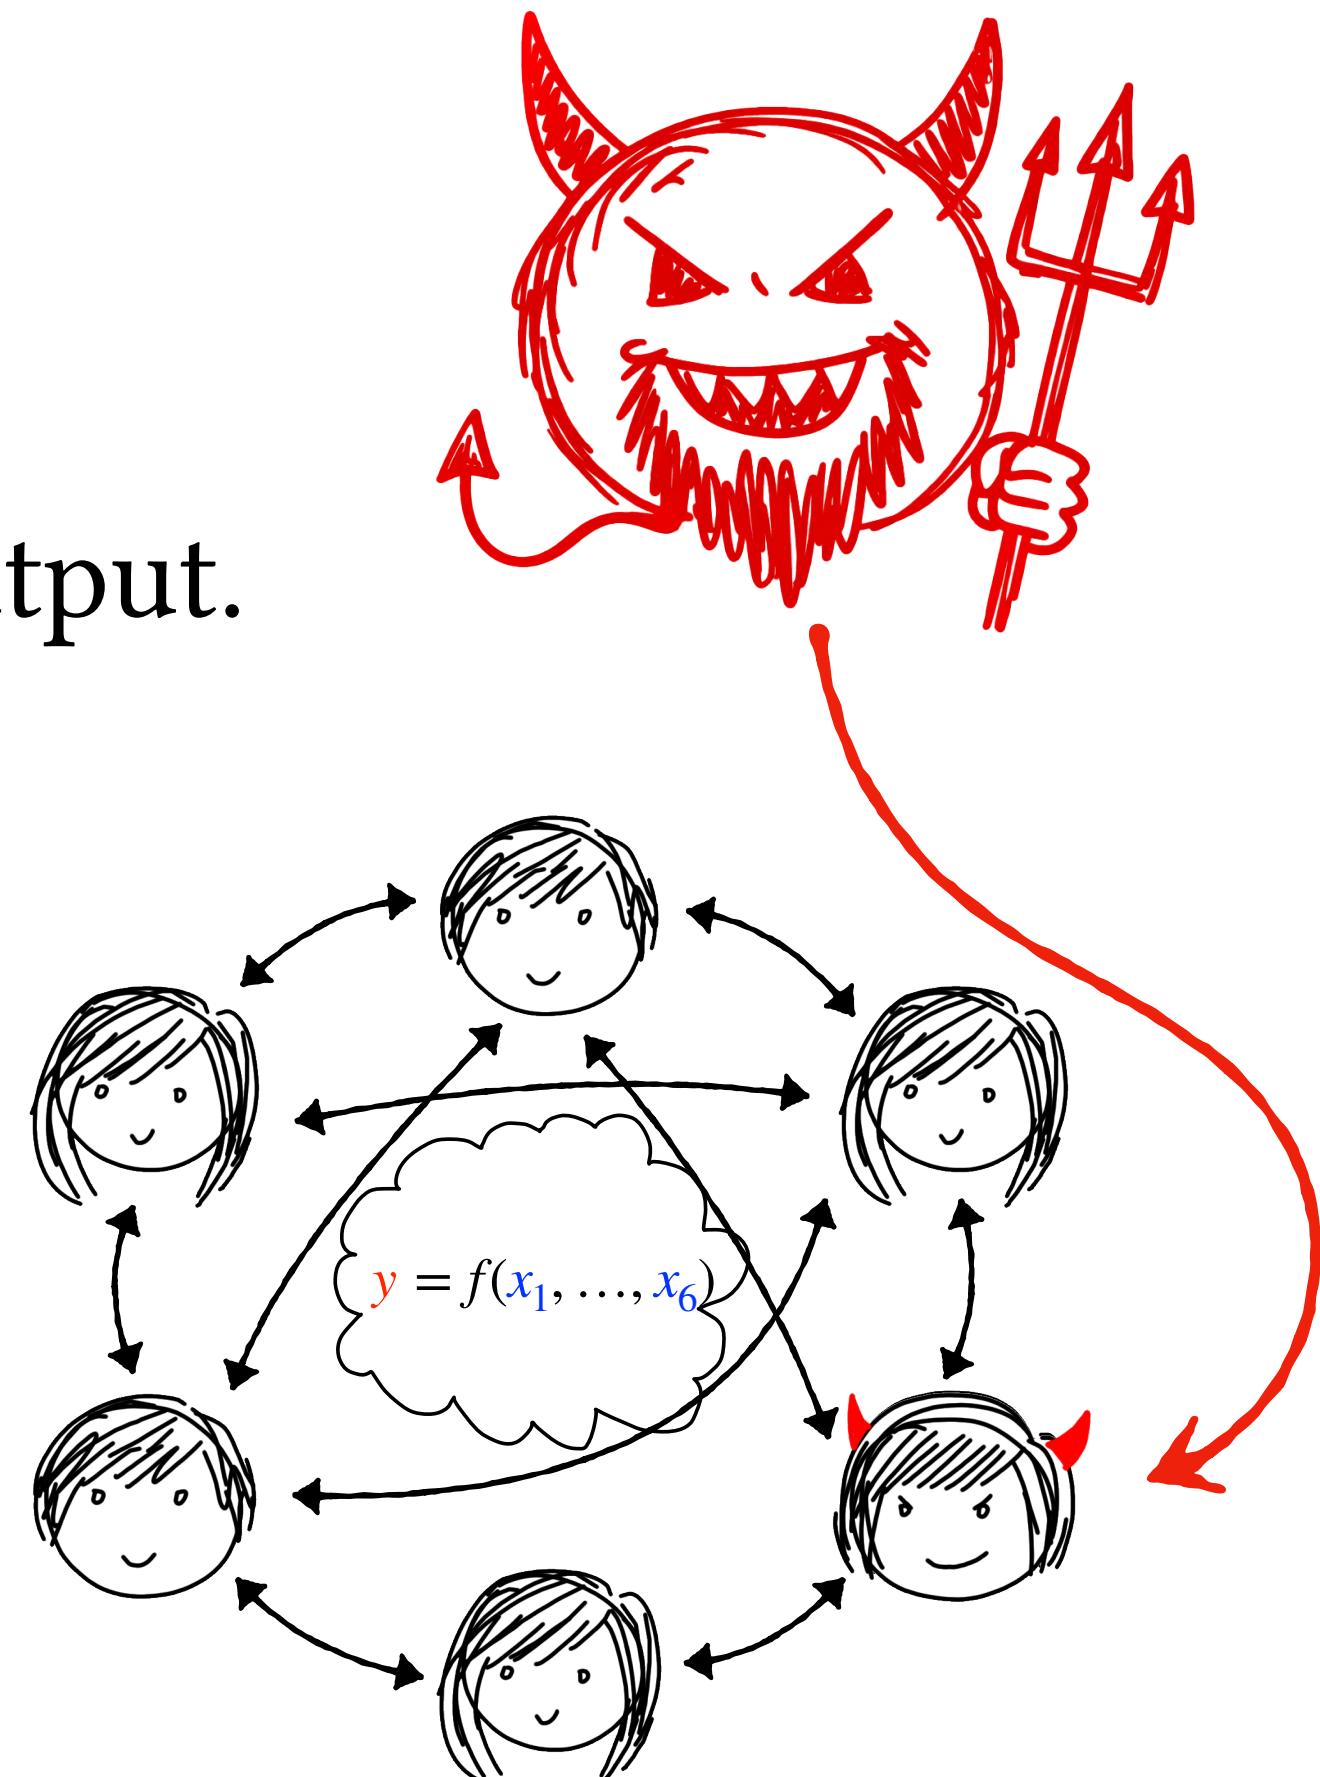
Another Protocol Example

No embarrassment this time.

Example: n-Party Sum

- Let M be some positive integer fixed a priori.

- Each P_i has private input $x_i \leq M/n$.

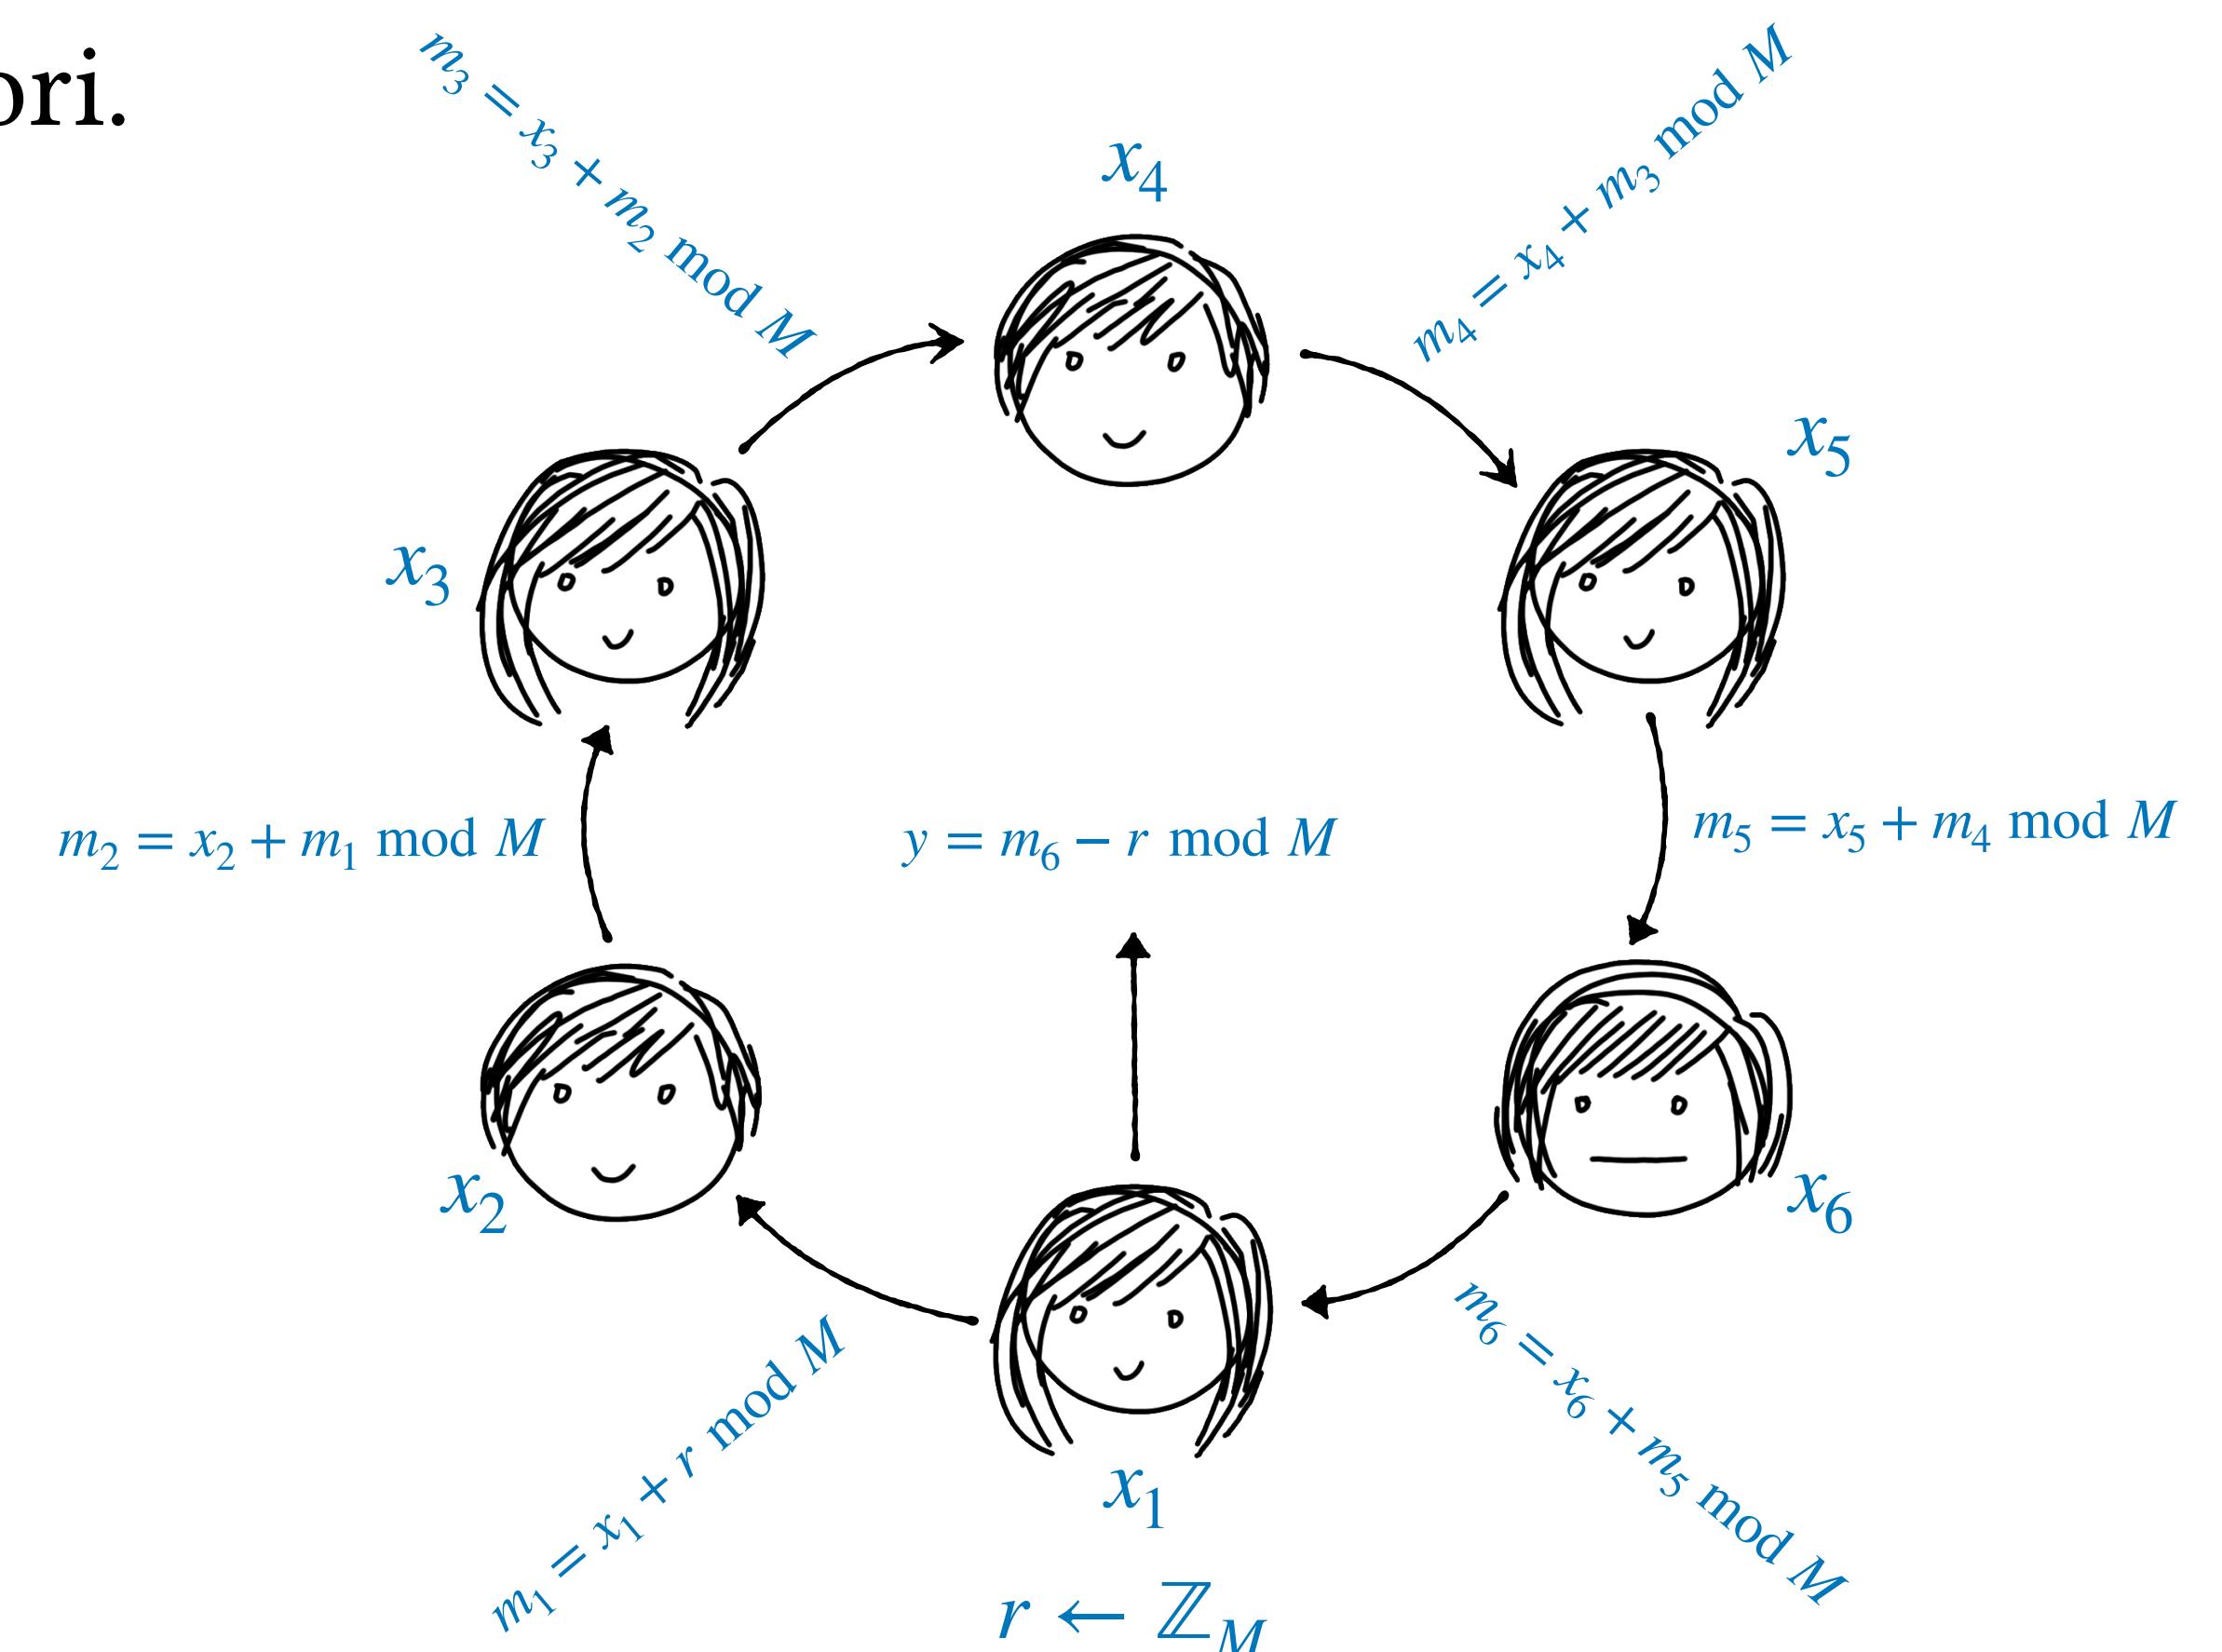- They wish to compute $y = \sum_{i \in [n]} x_i$.

Notation

$[n] = \{1, \ldots, n\}$ $\qquad$ $[n, m] = \{n, \ldots, m\}$

$\mathbb{Z}_M$ is the integers modulo M.
For our purposes, $\mathbb{Z}_M = [0, M-1]$.

$x \leftarrow D$ where D is a distribution means x is a random var distributed according to D

$x \leftarrow S$ where S is a set means x is sampled from the uniform distribution over S

In the diagram:

$m_3 = x_3 + m_2 \bmod M$

$m_4 = x_4 + m_3 \bmod M$

$m_5 = x_5 + m_4 \bmod M$

$m_6 = x_6 + m_5 \bmod M$

$m_2 = x_2 + m_1 \bmod M$

$y = m_6 - r \bmod M$

$m_1 = x_1 + r \bmod M$

$r \leftarrow \mathbb{Z}_M$

In the future we will omit the $\bmod M$ after every operation and just say "they work modulo M"

An Instance with 6 Parties

- Let $M = 100$, so $x_i \leq 16$

- Suppose one party is "honest but curious" (We call this a *semi-honest* adversary)

- $\mathcal{A}$ learns $m_2 = x_2 + x_1 + r$

- $\forall i \in \mathbb{Z}_M \ \Pr[m_2 = i] = 1/100$

- In other words, the distribution of m_2 is independent of x_1 and x_2.

- $\mathcal{A}$ could sample $m_2 \leftarrow \mathbb{Z}_M$ on its own.

- So the protocol is **private**.

$x_4 = 3$

$m_3 = 3$

$m_4 = 6$

$x_3 = 9$

$x_5 = 2$

$m_2 = 94$

$y = 31$

$m_5 = 8$

$x_2 = 6$

$x_6 = 8$

$m_1 = 88$
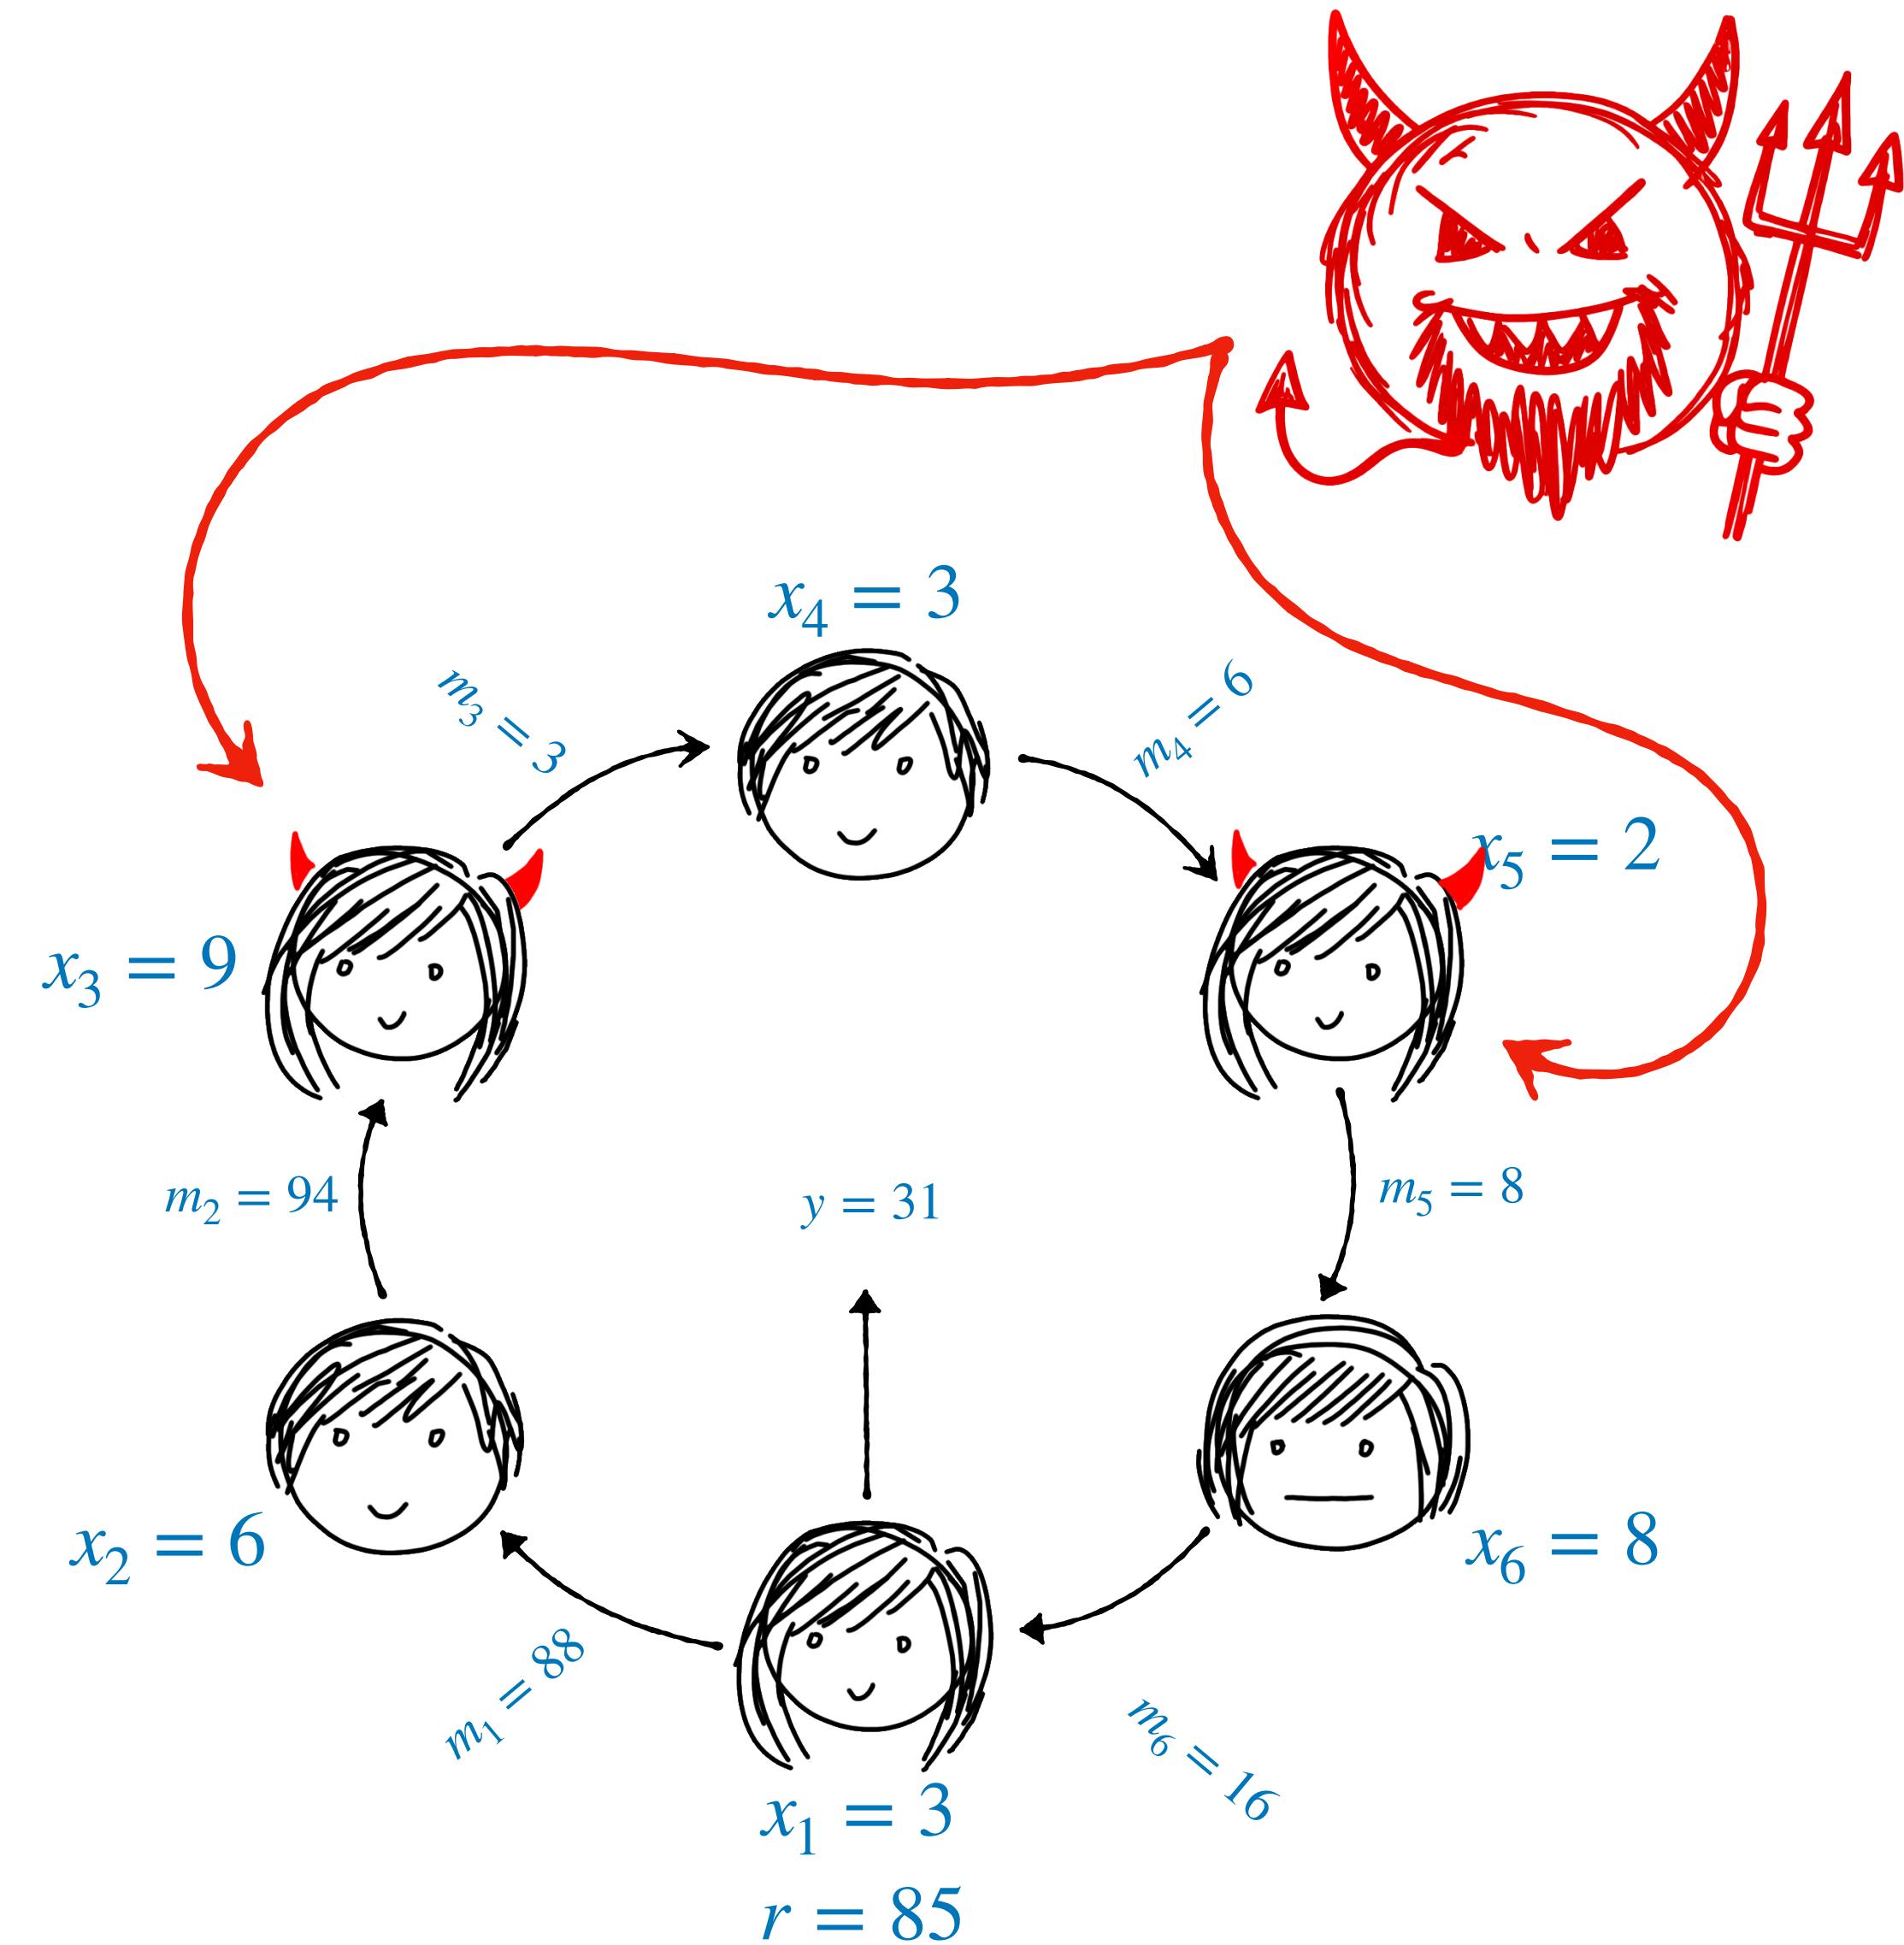
$m_6 = 16$

$x_1 = 3$

$r = 85$
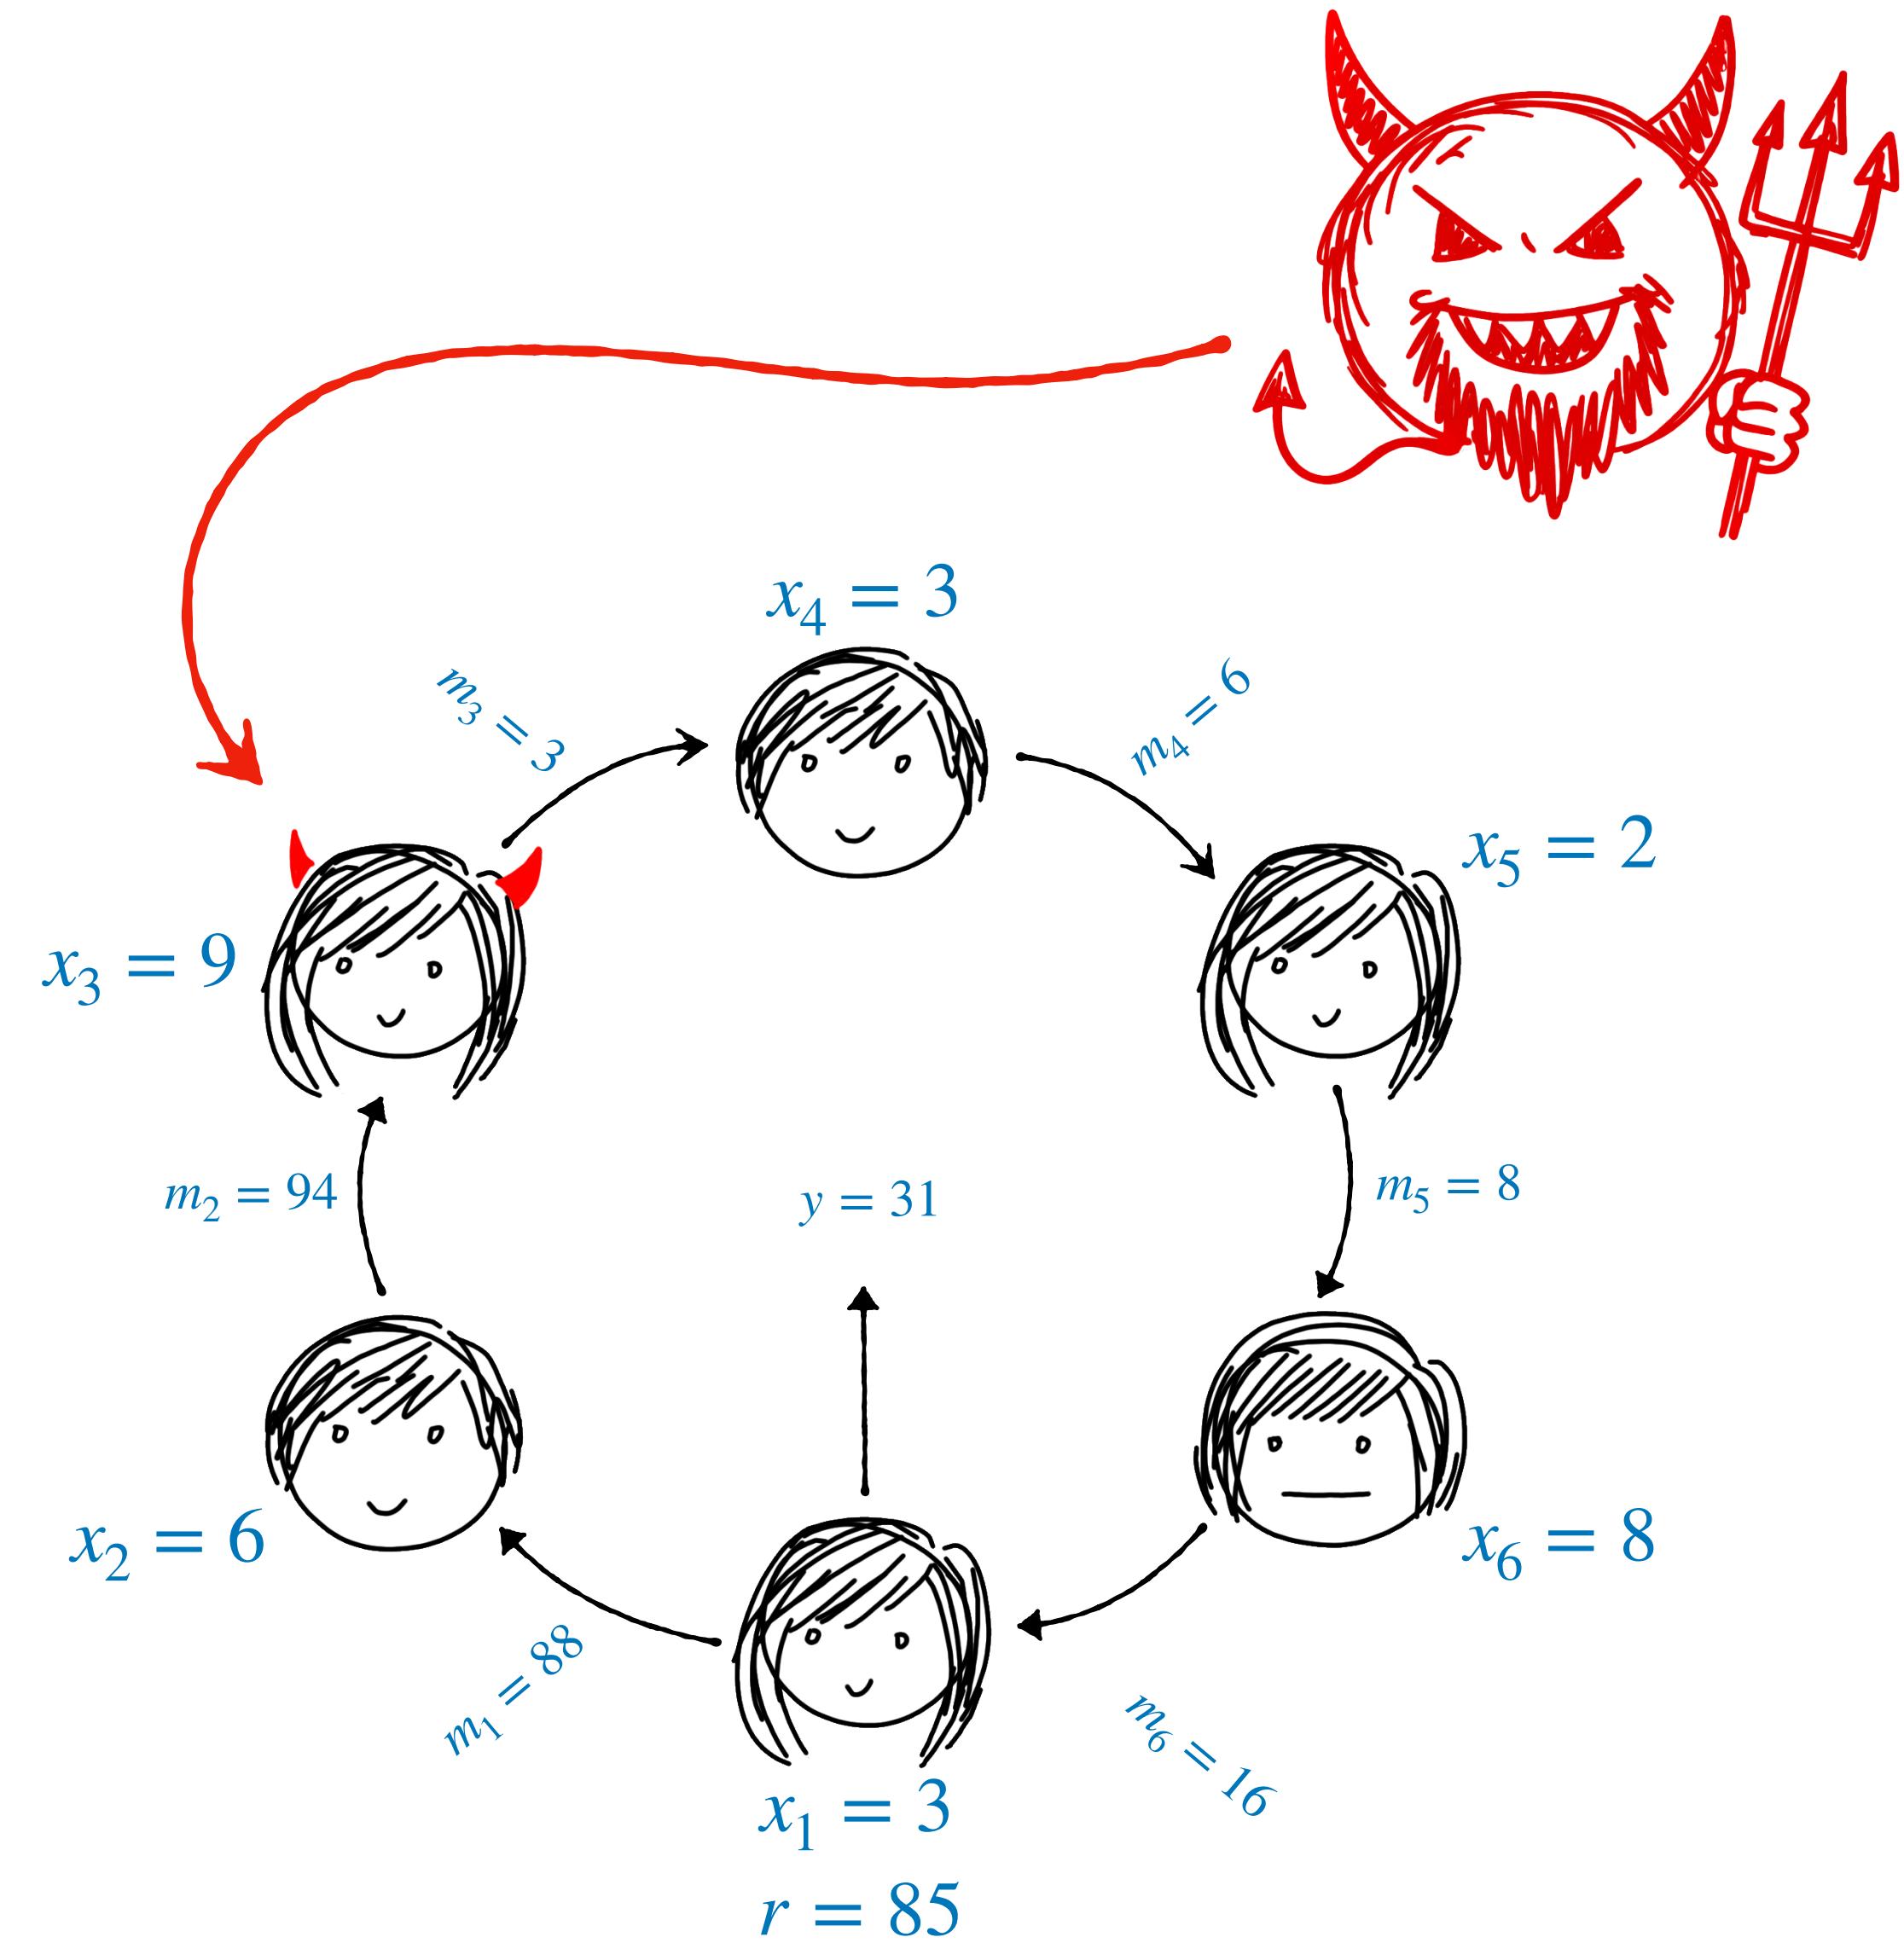
An Instance with 6 Parties

- Let $M = 100$, so $x_i \leq 16$

- Suppose one party is "honest but curious"
  (We call this a *semi-honest* adversary*)*

- $\mathscr{A}$ learns $m_2 = x_2 + x_1 + r$

- $\forall i \in \mathbb{Z}_M \ \Pr[m_2 = i] = 1/100$

- In other words, the distribution of
  m_2 is independent of x_1 and x_2.

- $\mathscr{A}$ could sample $m_2 \leftarrow \mathbb{Z}_M$ on its own.

- So the protocol is **private**.

- But what if *two* parties are corrupted?

$x_4 = 3$

$m_3 = 3$

$m_4 = 6$

$x_5 = 2$

$x_3 = 9$

$m_2 = 94$

$y = 31$

$m_5 = 8$

$x_2 = 6$

$x_6 = 8$

$m_1 = 88$

$m_6 = 16$

$x_1 = 3$

$r = 85$

What can we take from this?

Security is not a property of the protocol *alone.*
It also depends upon the adversarial model.

How to Define Security

Option 1: Property Based

- Define a list of security requirements for the task to be accomplished (e.g. privacy, fairness, guaranteed output)

- Give a unique mathematical formalization for each property. Prove each property individually.

- For simple foundational primitives like *Encryption* or *Digital Signatures*, this is the kind of security definition we usually use. Usually only 1-2 properties we care about. (e.g. Signatures require the property of *Unforgeability*)

- For complex tasks, how do we know if we have covered all of our concerns? Proving many properties is also a huge chore.

- **Deeper Problem**: when we proved the properties of protocol A, we considered it on its own. Will it still have those properties when we use it to build protocol B?

How to Define Security

Option 2: the simulation (a.k.a. real-ideal) paradigm

- Consider the real world in which the protocol is executed.
  Parties have inputs, communicate with each other, produce outputs

- Consider an ideal world in which an additional trusted party exists.
  The trusted party helps the other parties carry out the computation.
  In this world security holds *by definition,* i.e. we define this world to be secure.

- We say that the real protocol is secure if whatever an adversary can achieve by attacking it can also be achieved by attacking the ideal computation involving the trusted party.

- Thus the behavior of the trusted party captures *all* properties of the protocol at once, not just the function it computes. We call this trusted party an ideal functionality. When a protocol is secure with respect to some functionality, we say that it realizes the functionality.
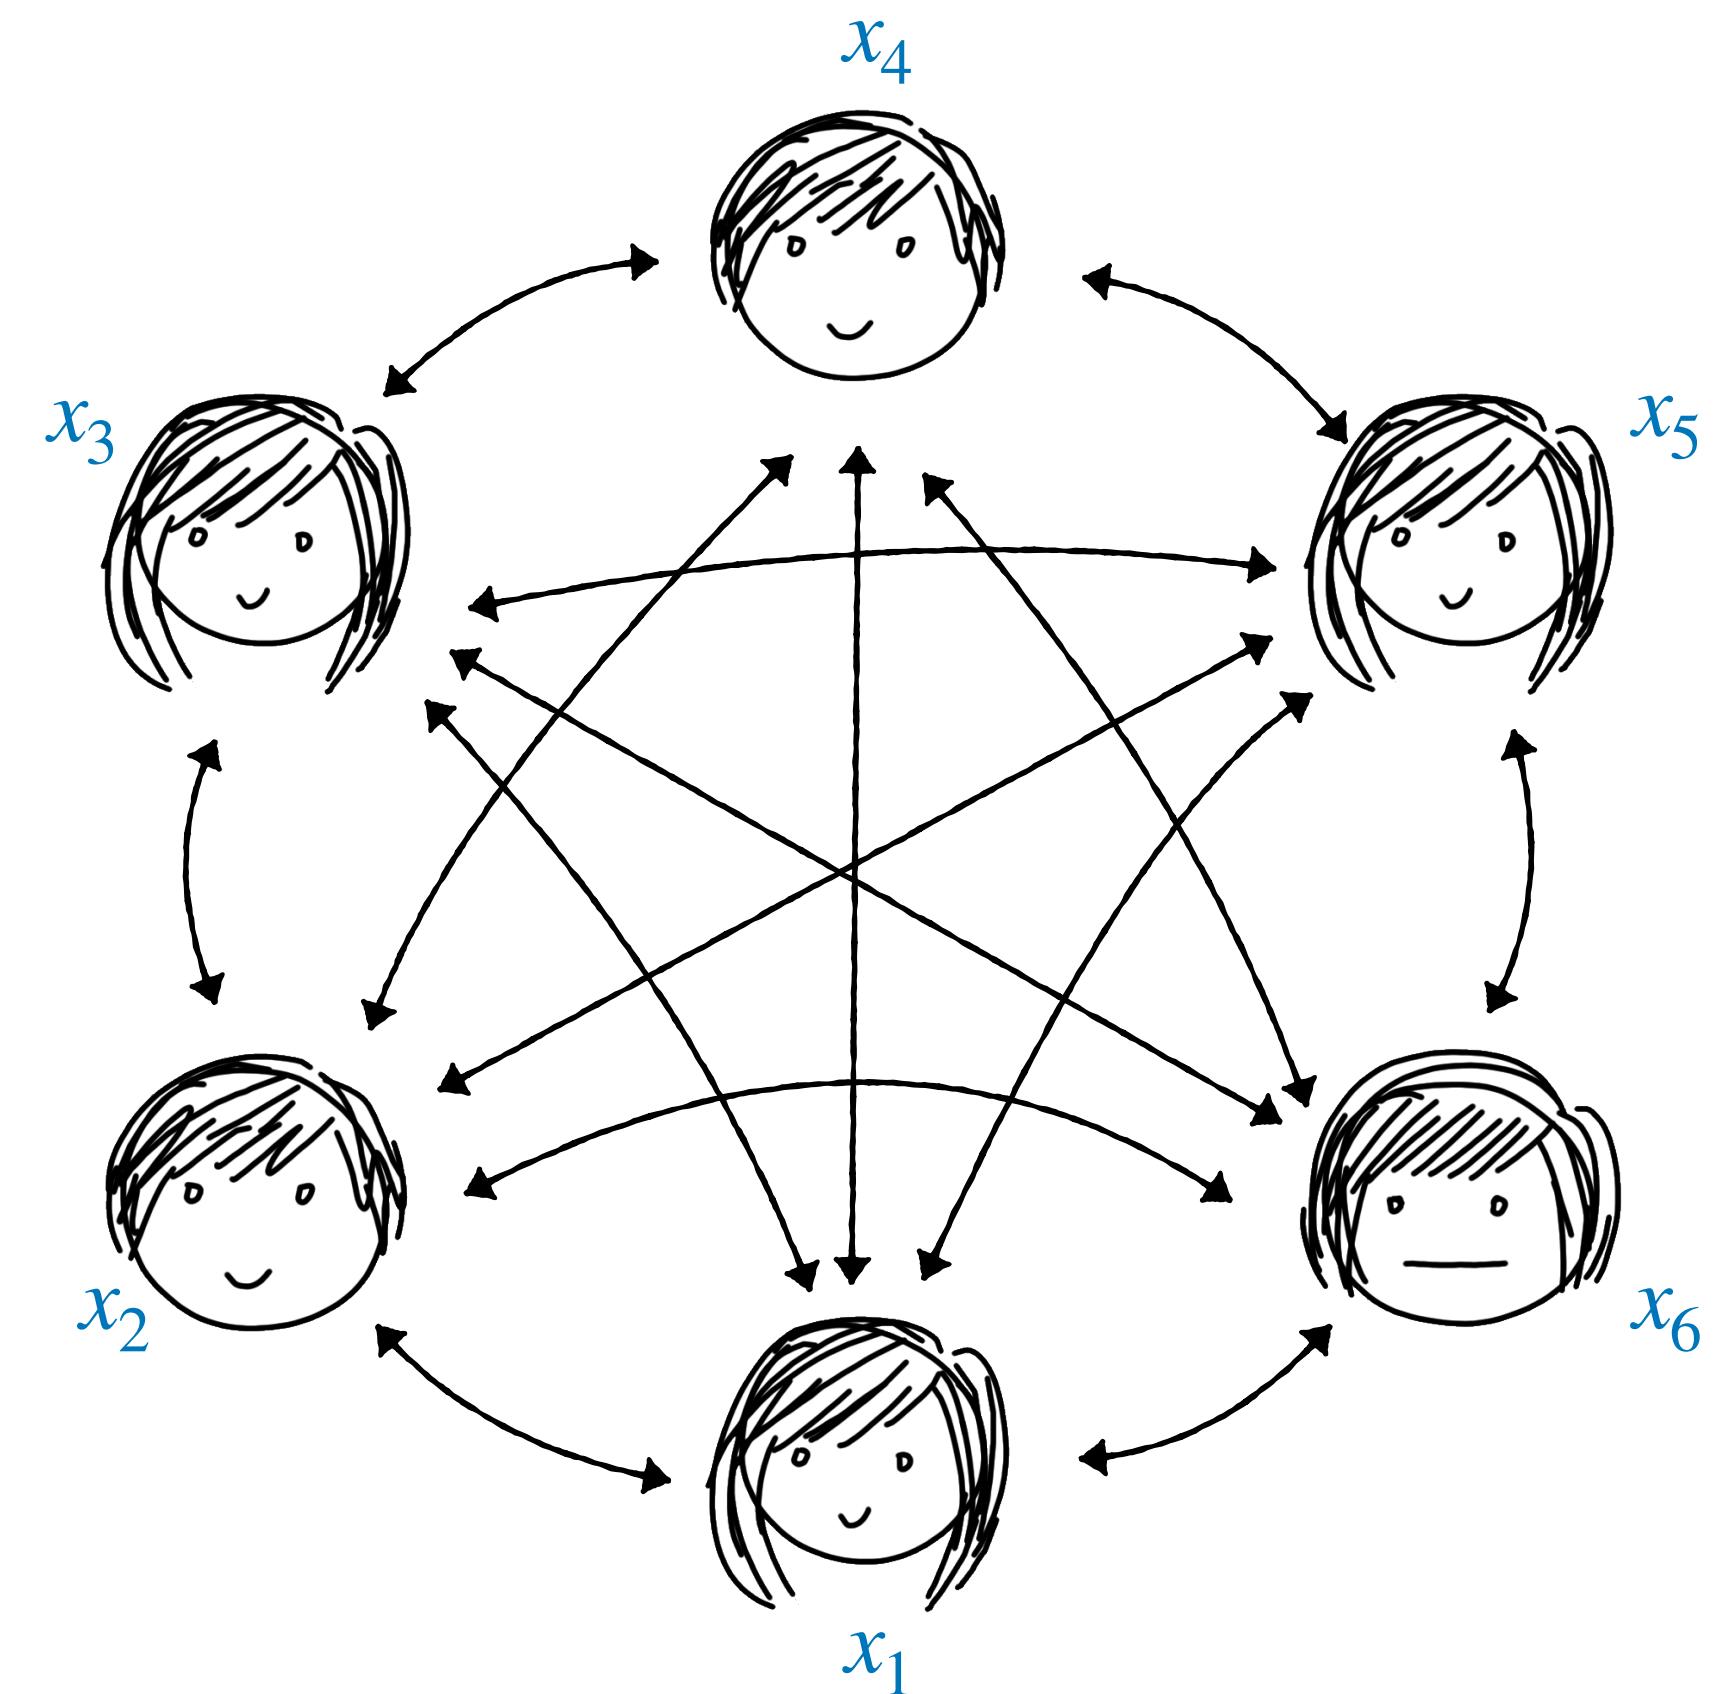
The Ideal World

1. Each P_i sends its input x_i to the ideal functionality ($\mathcal{F}$)

2. $\mathcal{F}$ computes $y = f(x_1, \ldots, x_n)$.

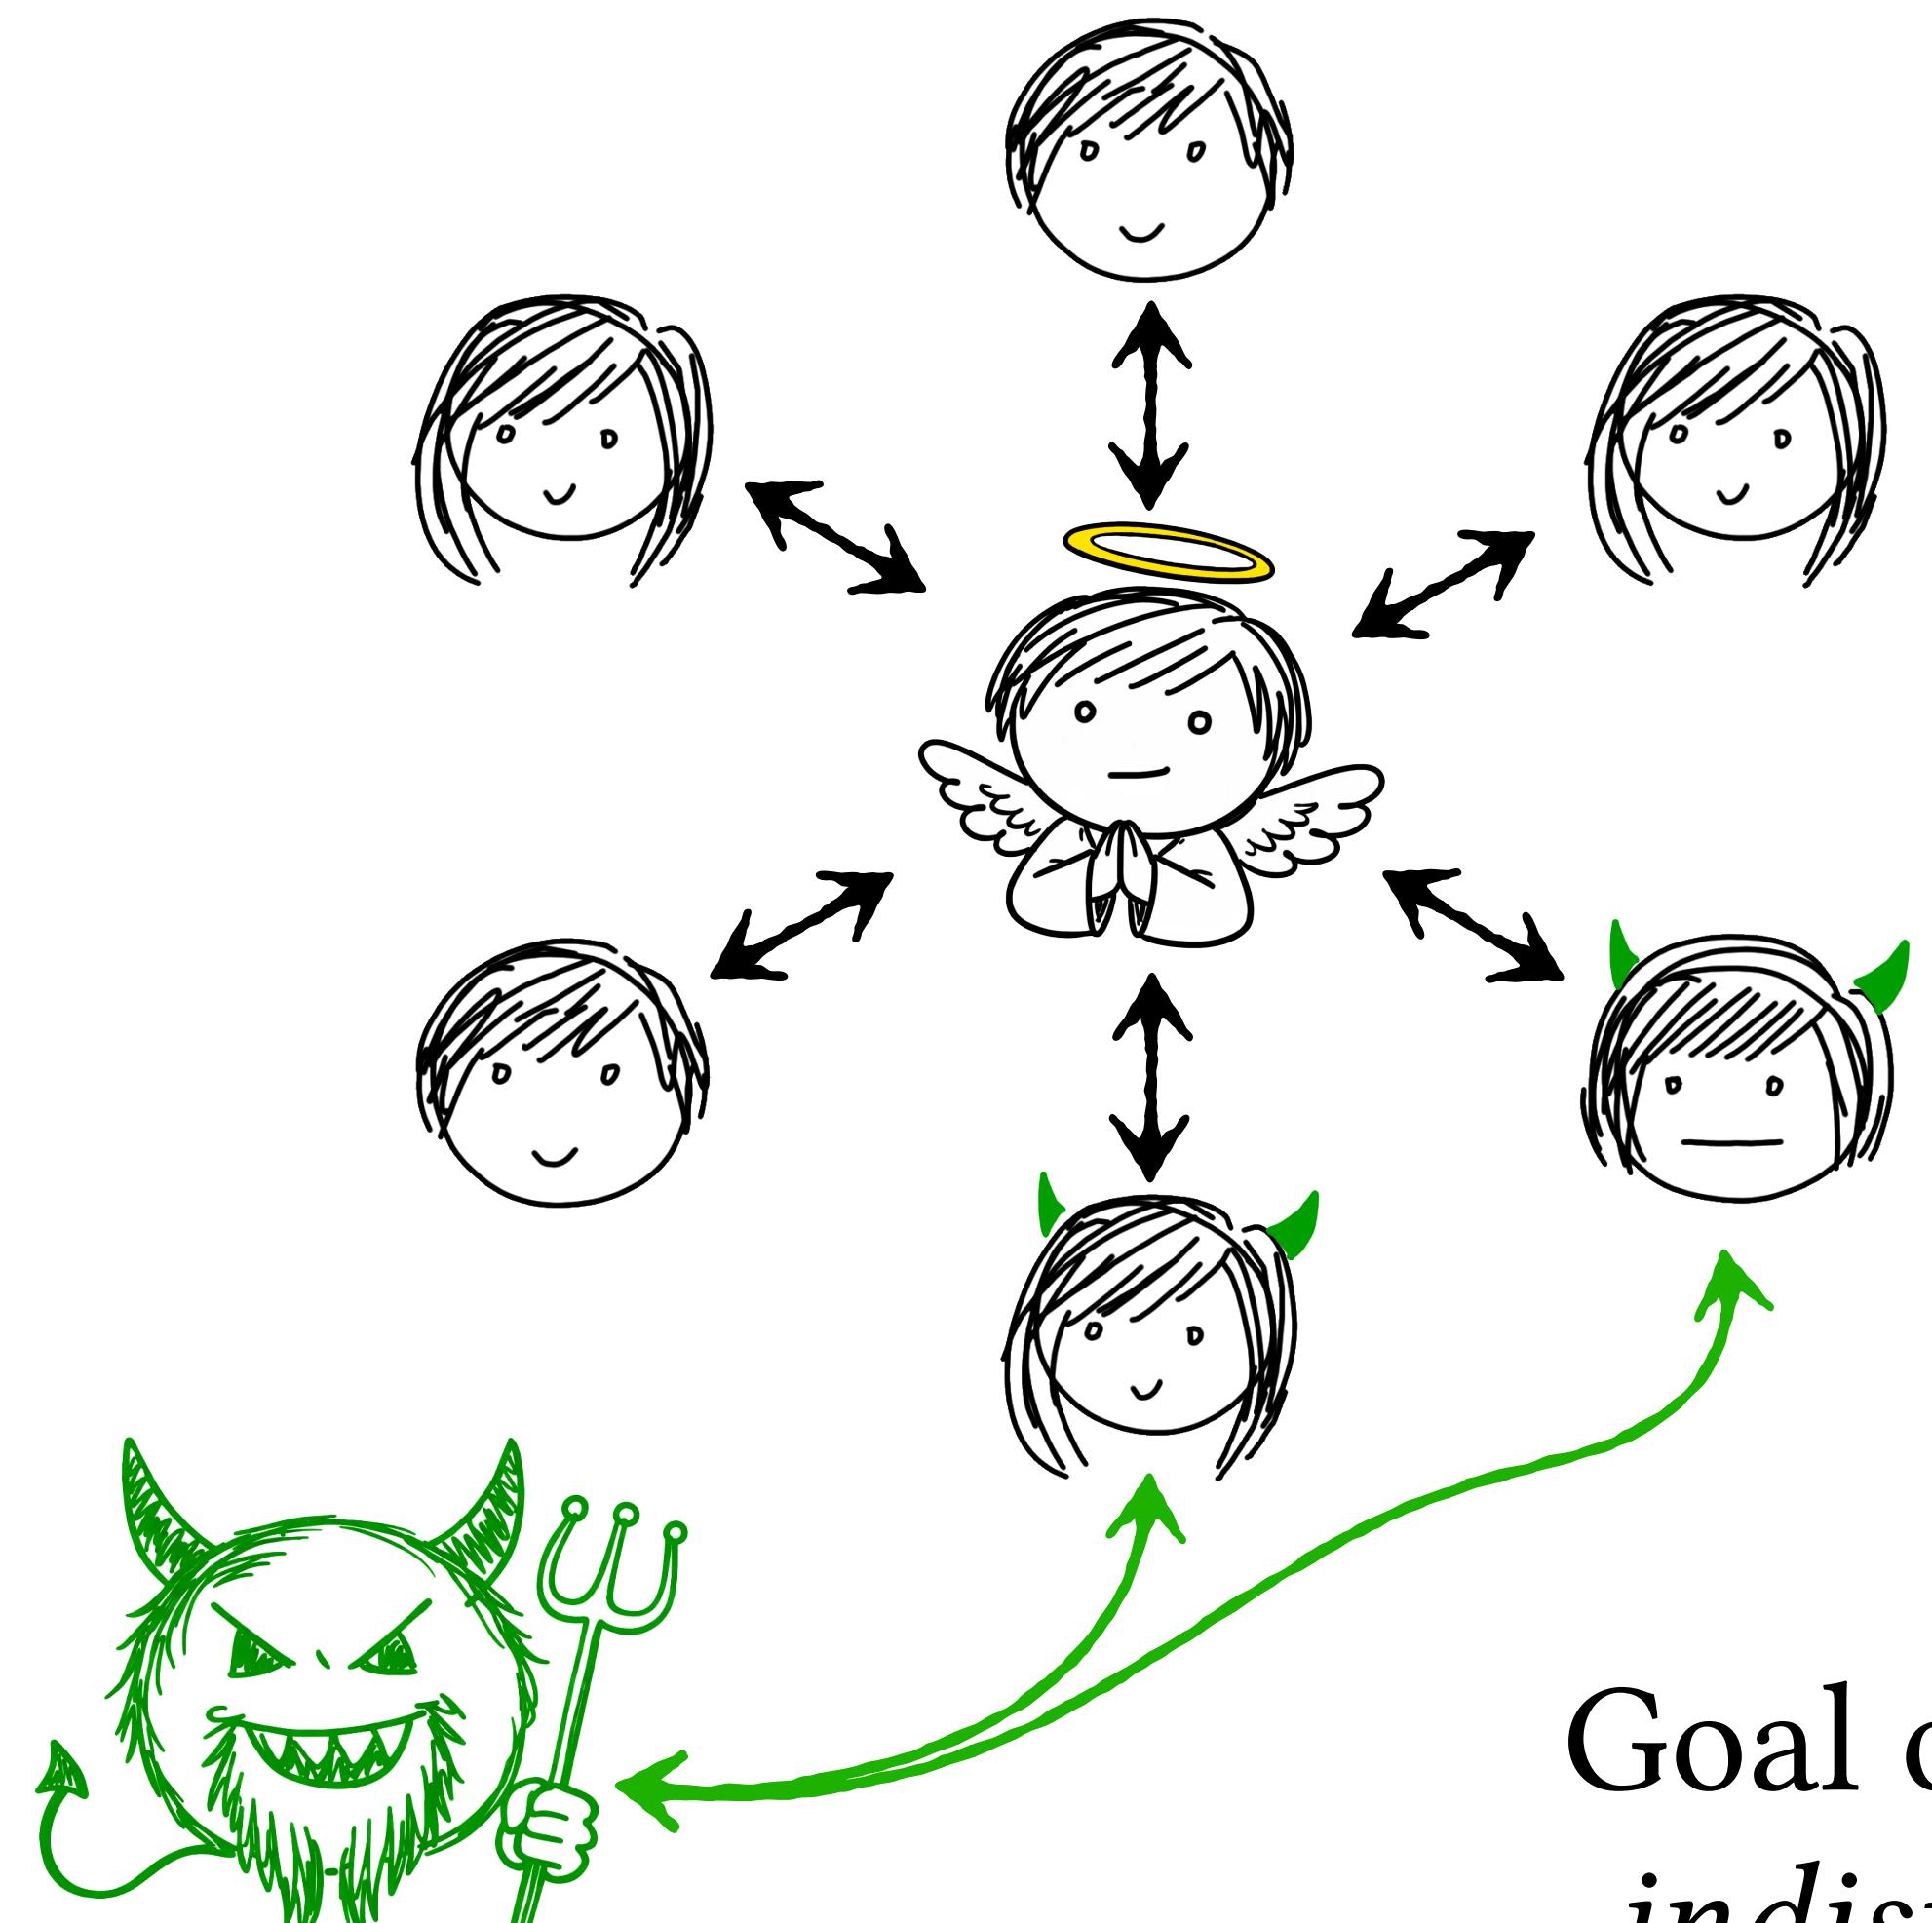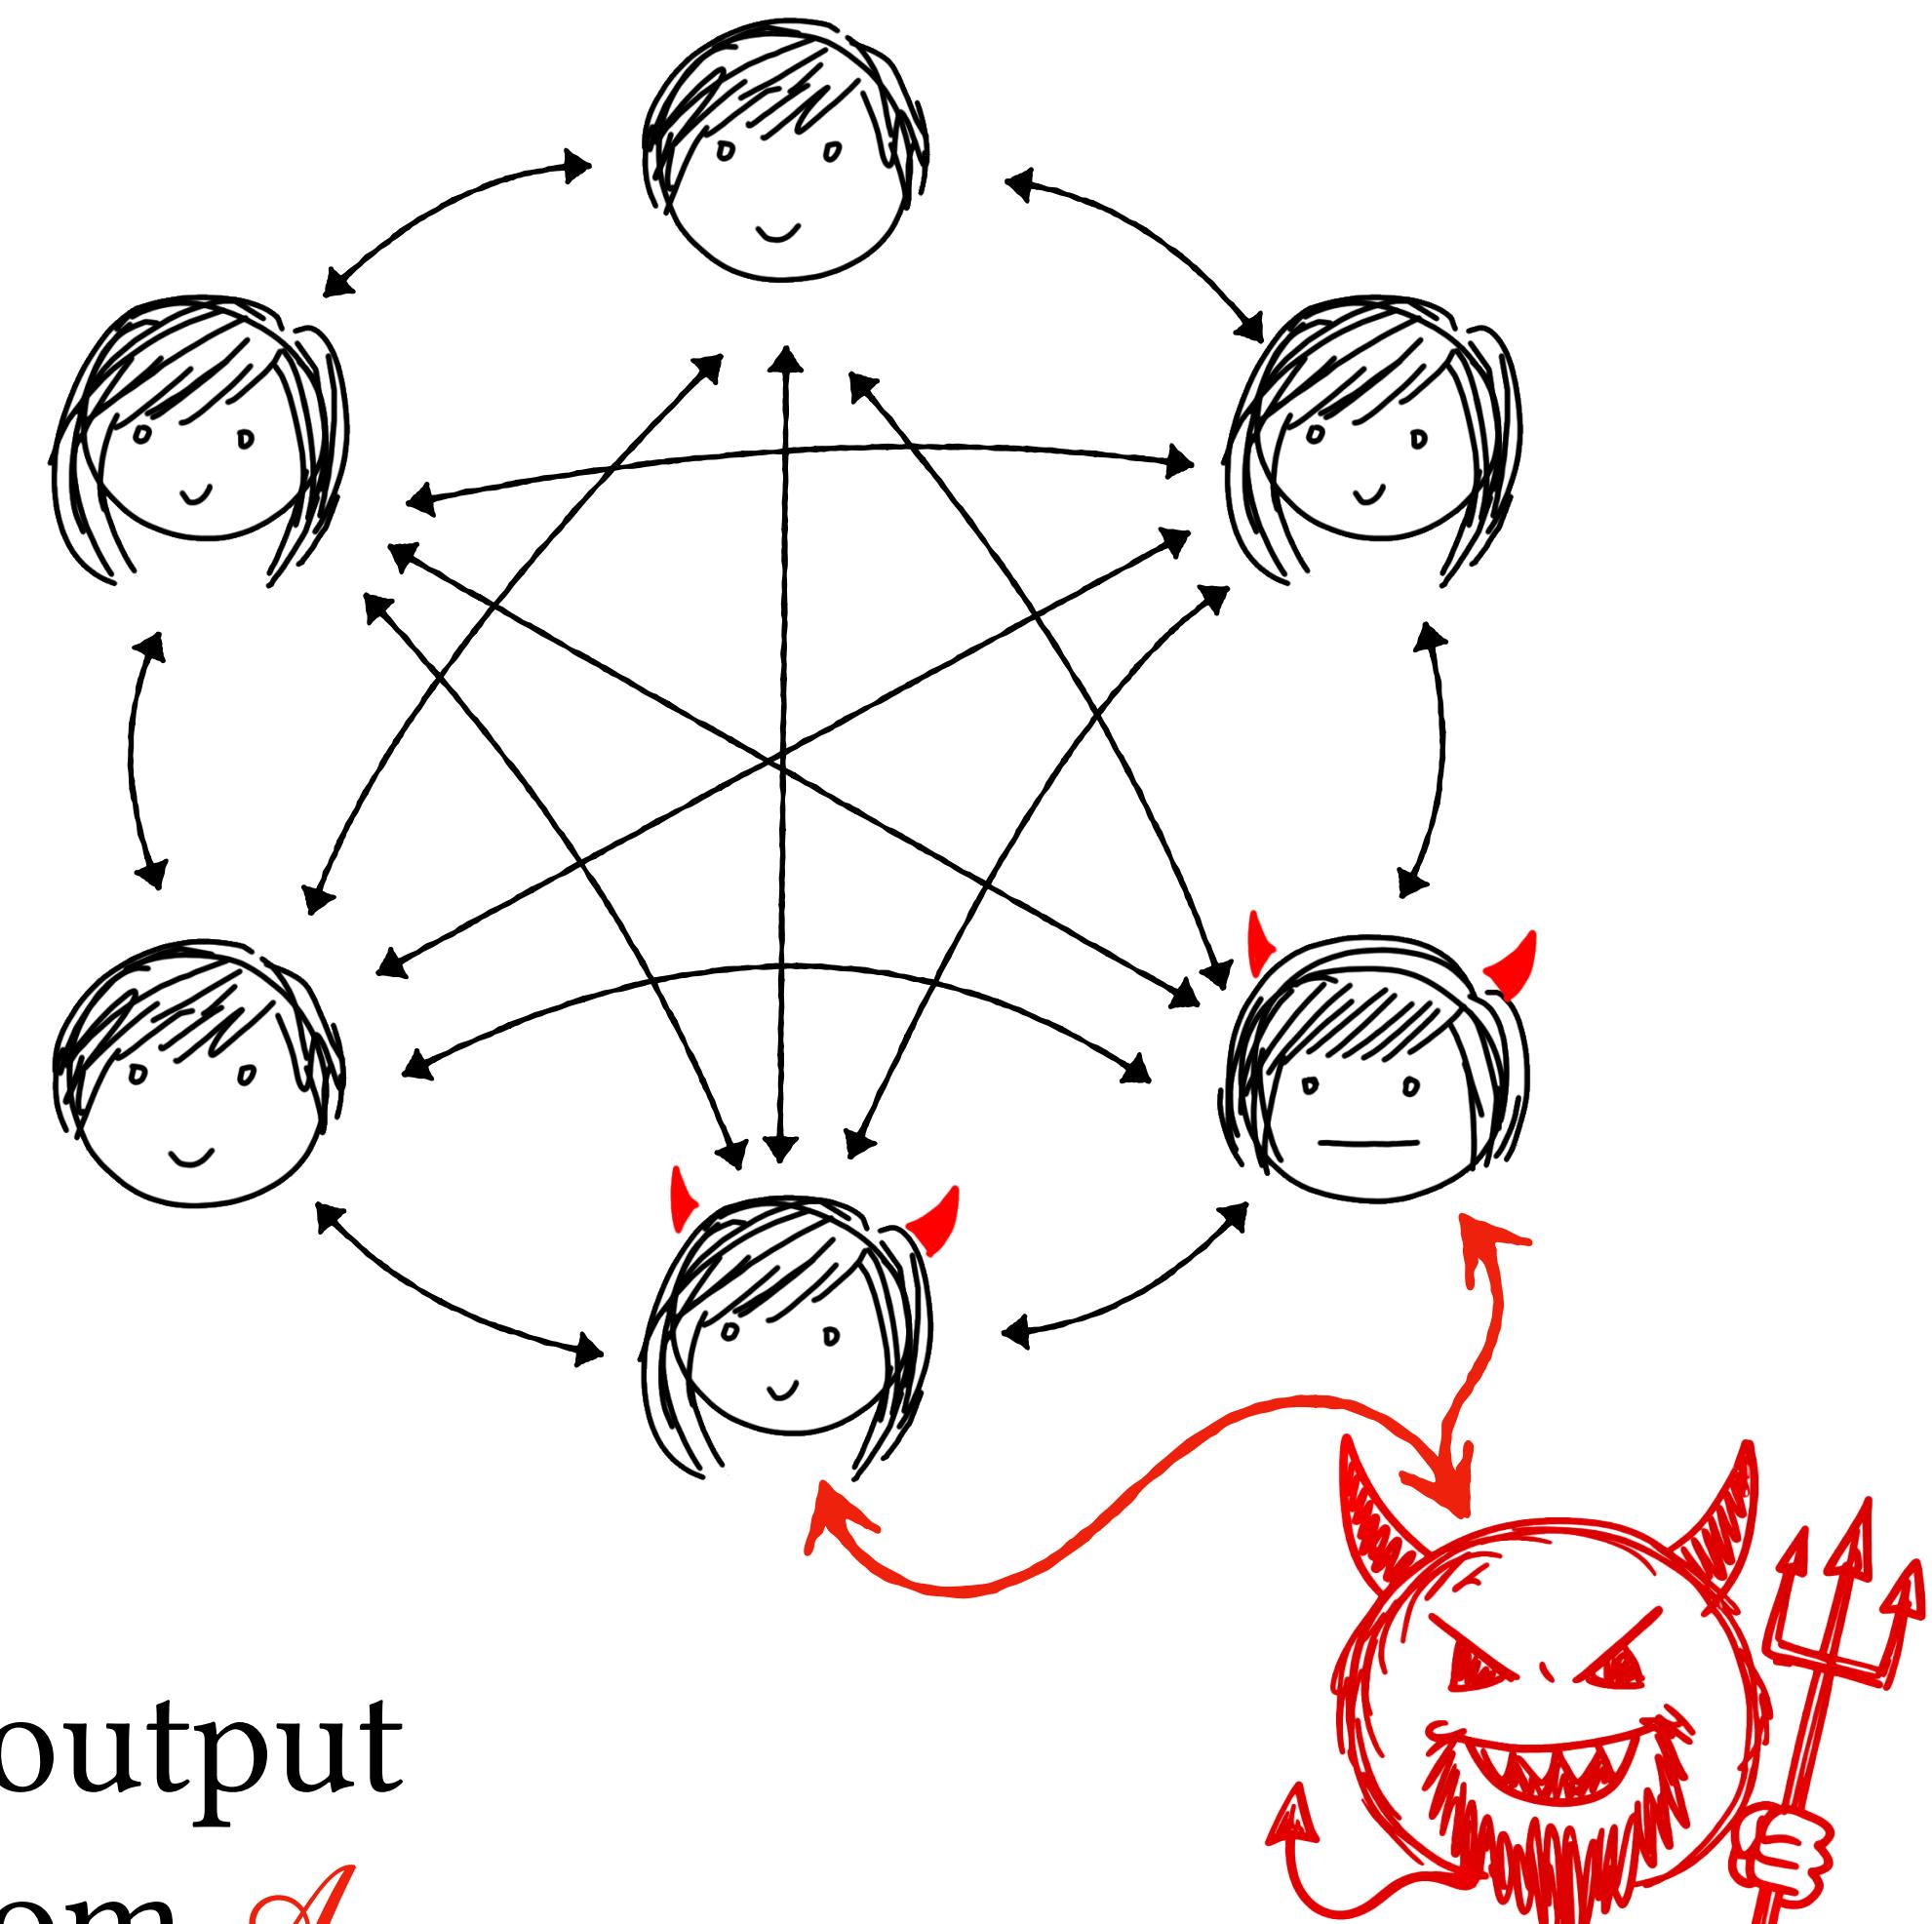3. $\mathcal{F}$ sends y to every party, and they output it.
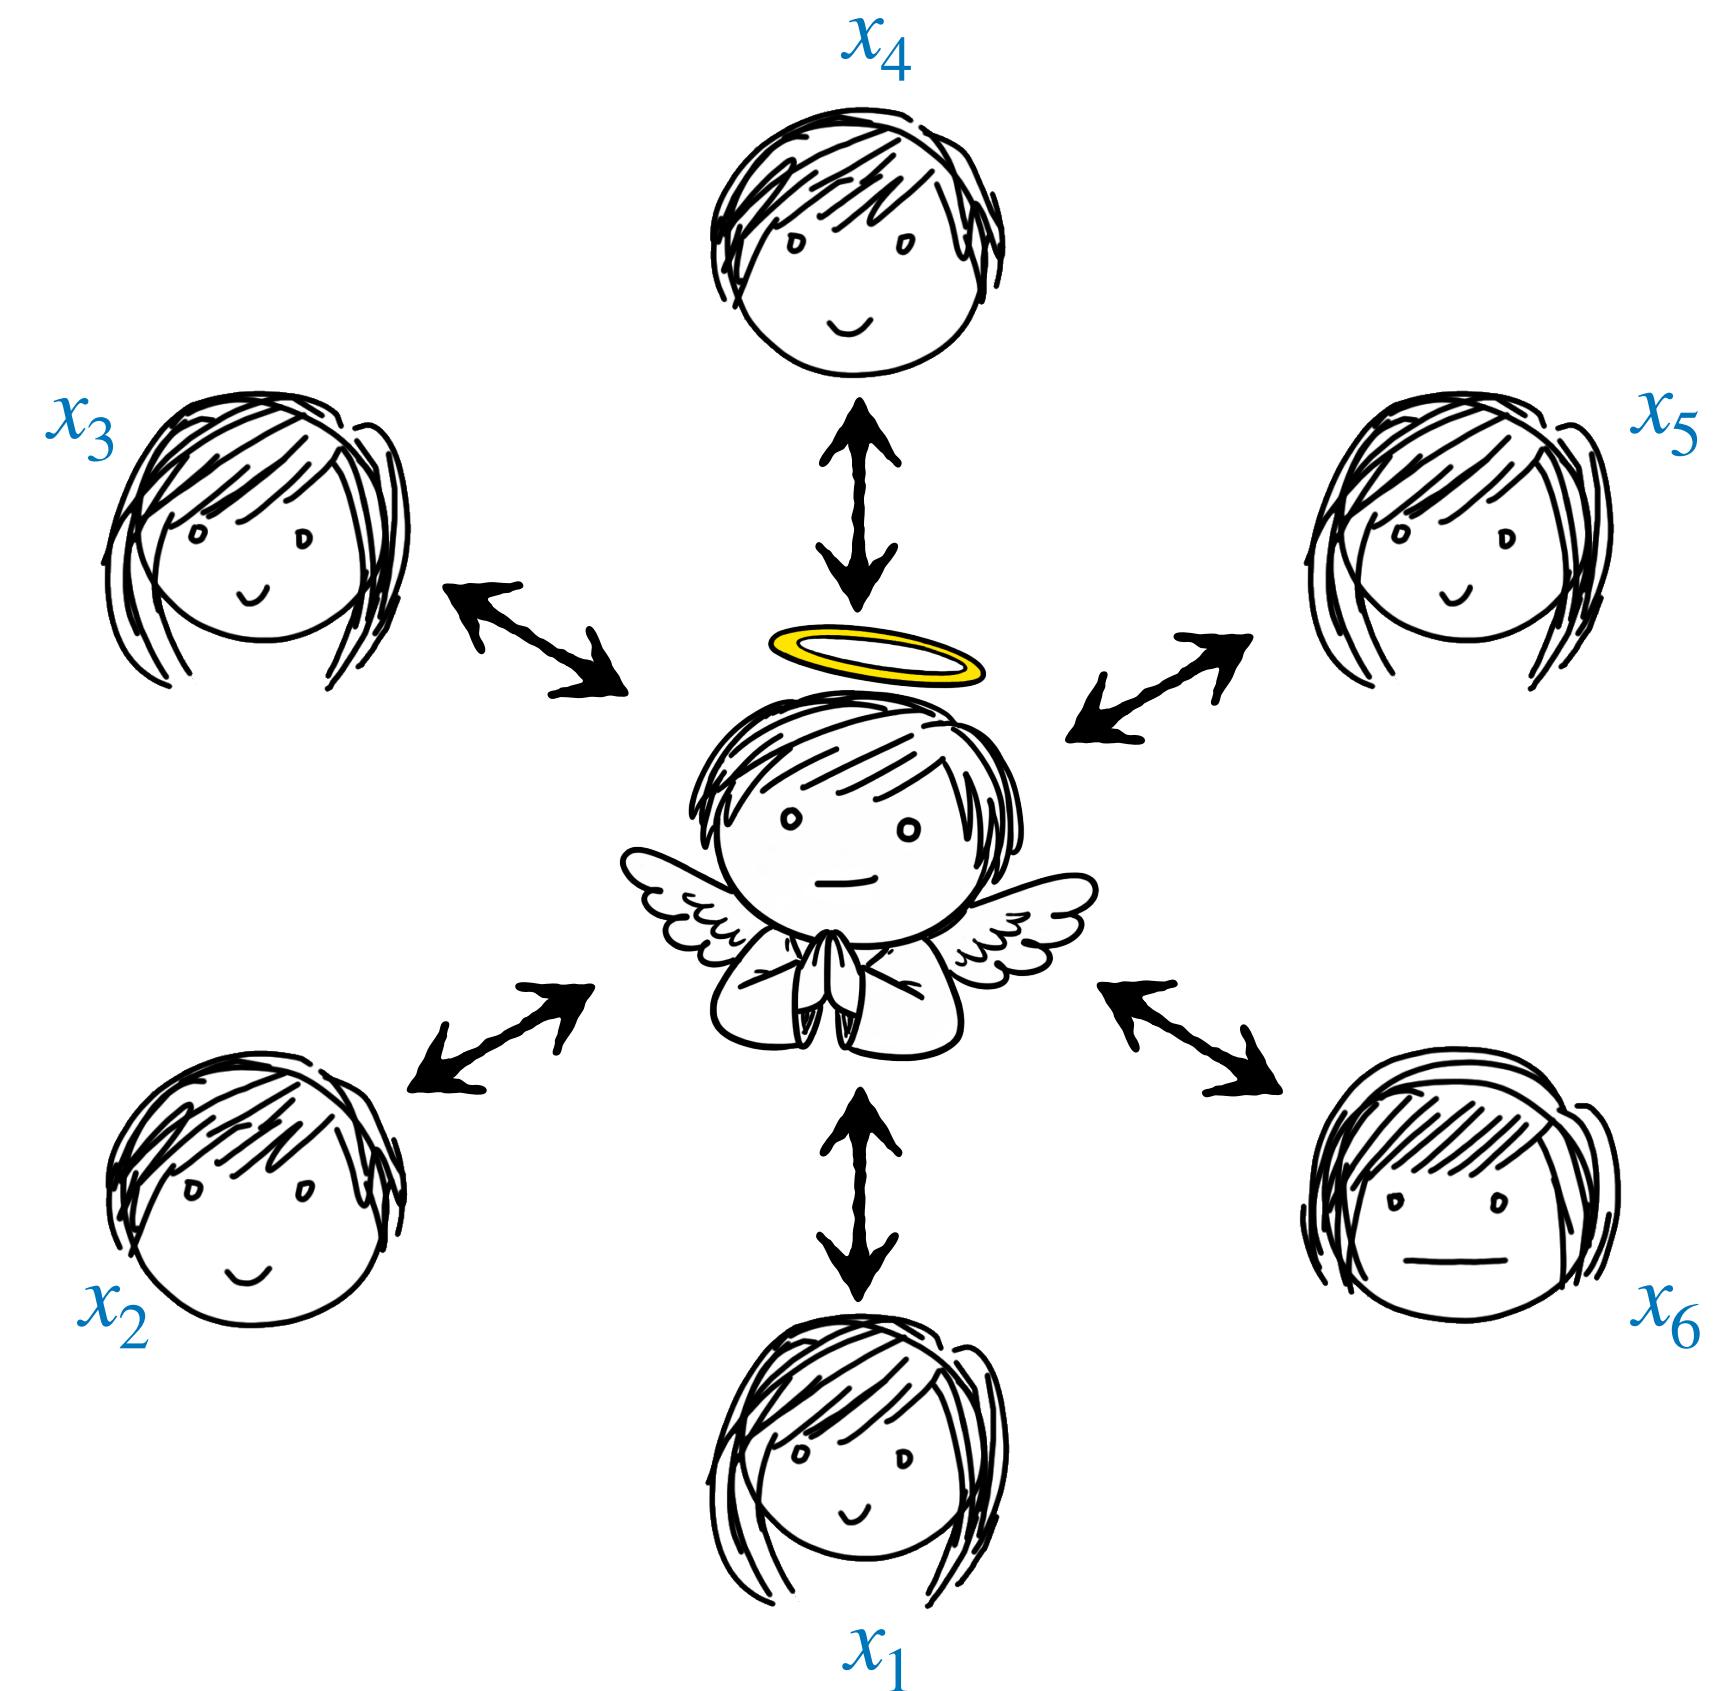
The Real World

1. The parties $(P_1, \ldots, P_n)$ run a protocol π on inputs $(x_1, \ldots, x_n)$

2. When π terminates, the parties output y.
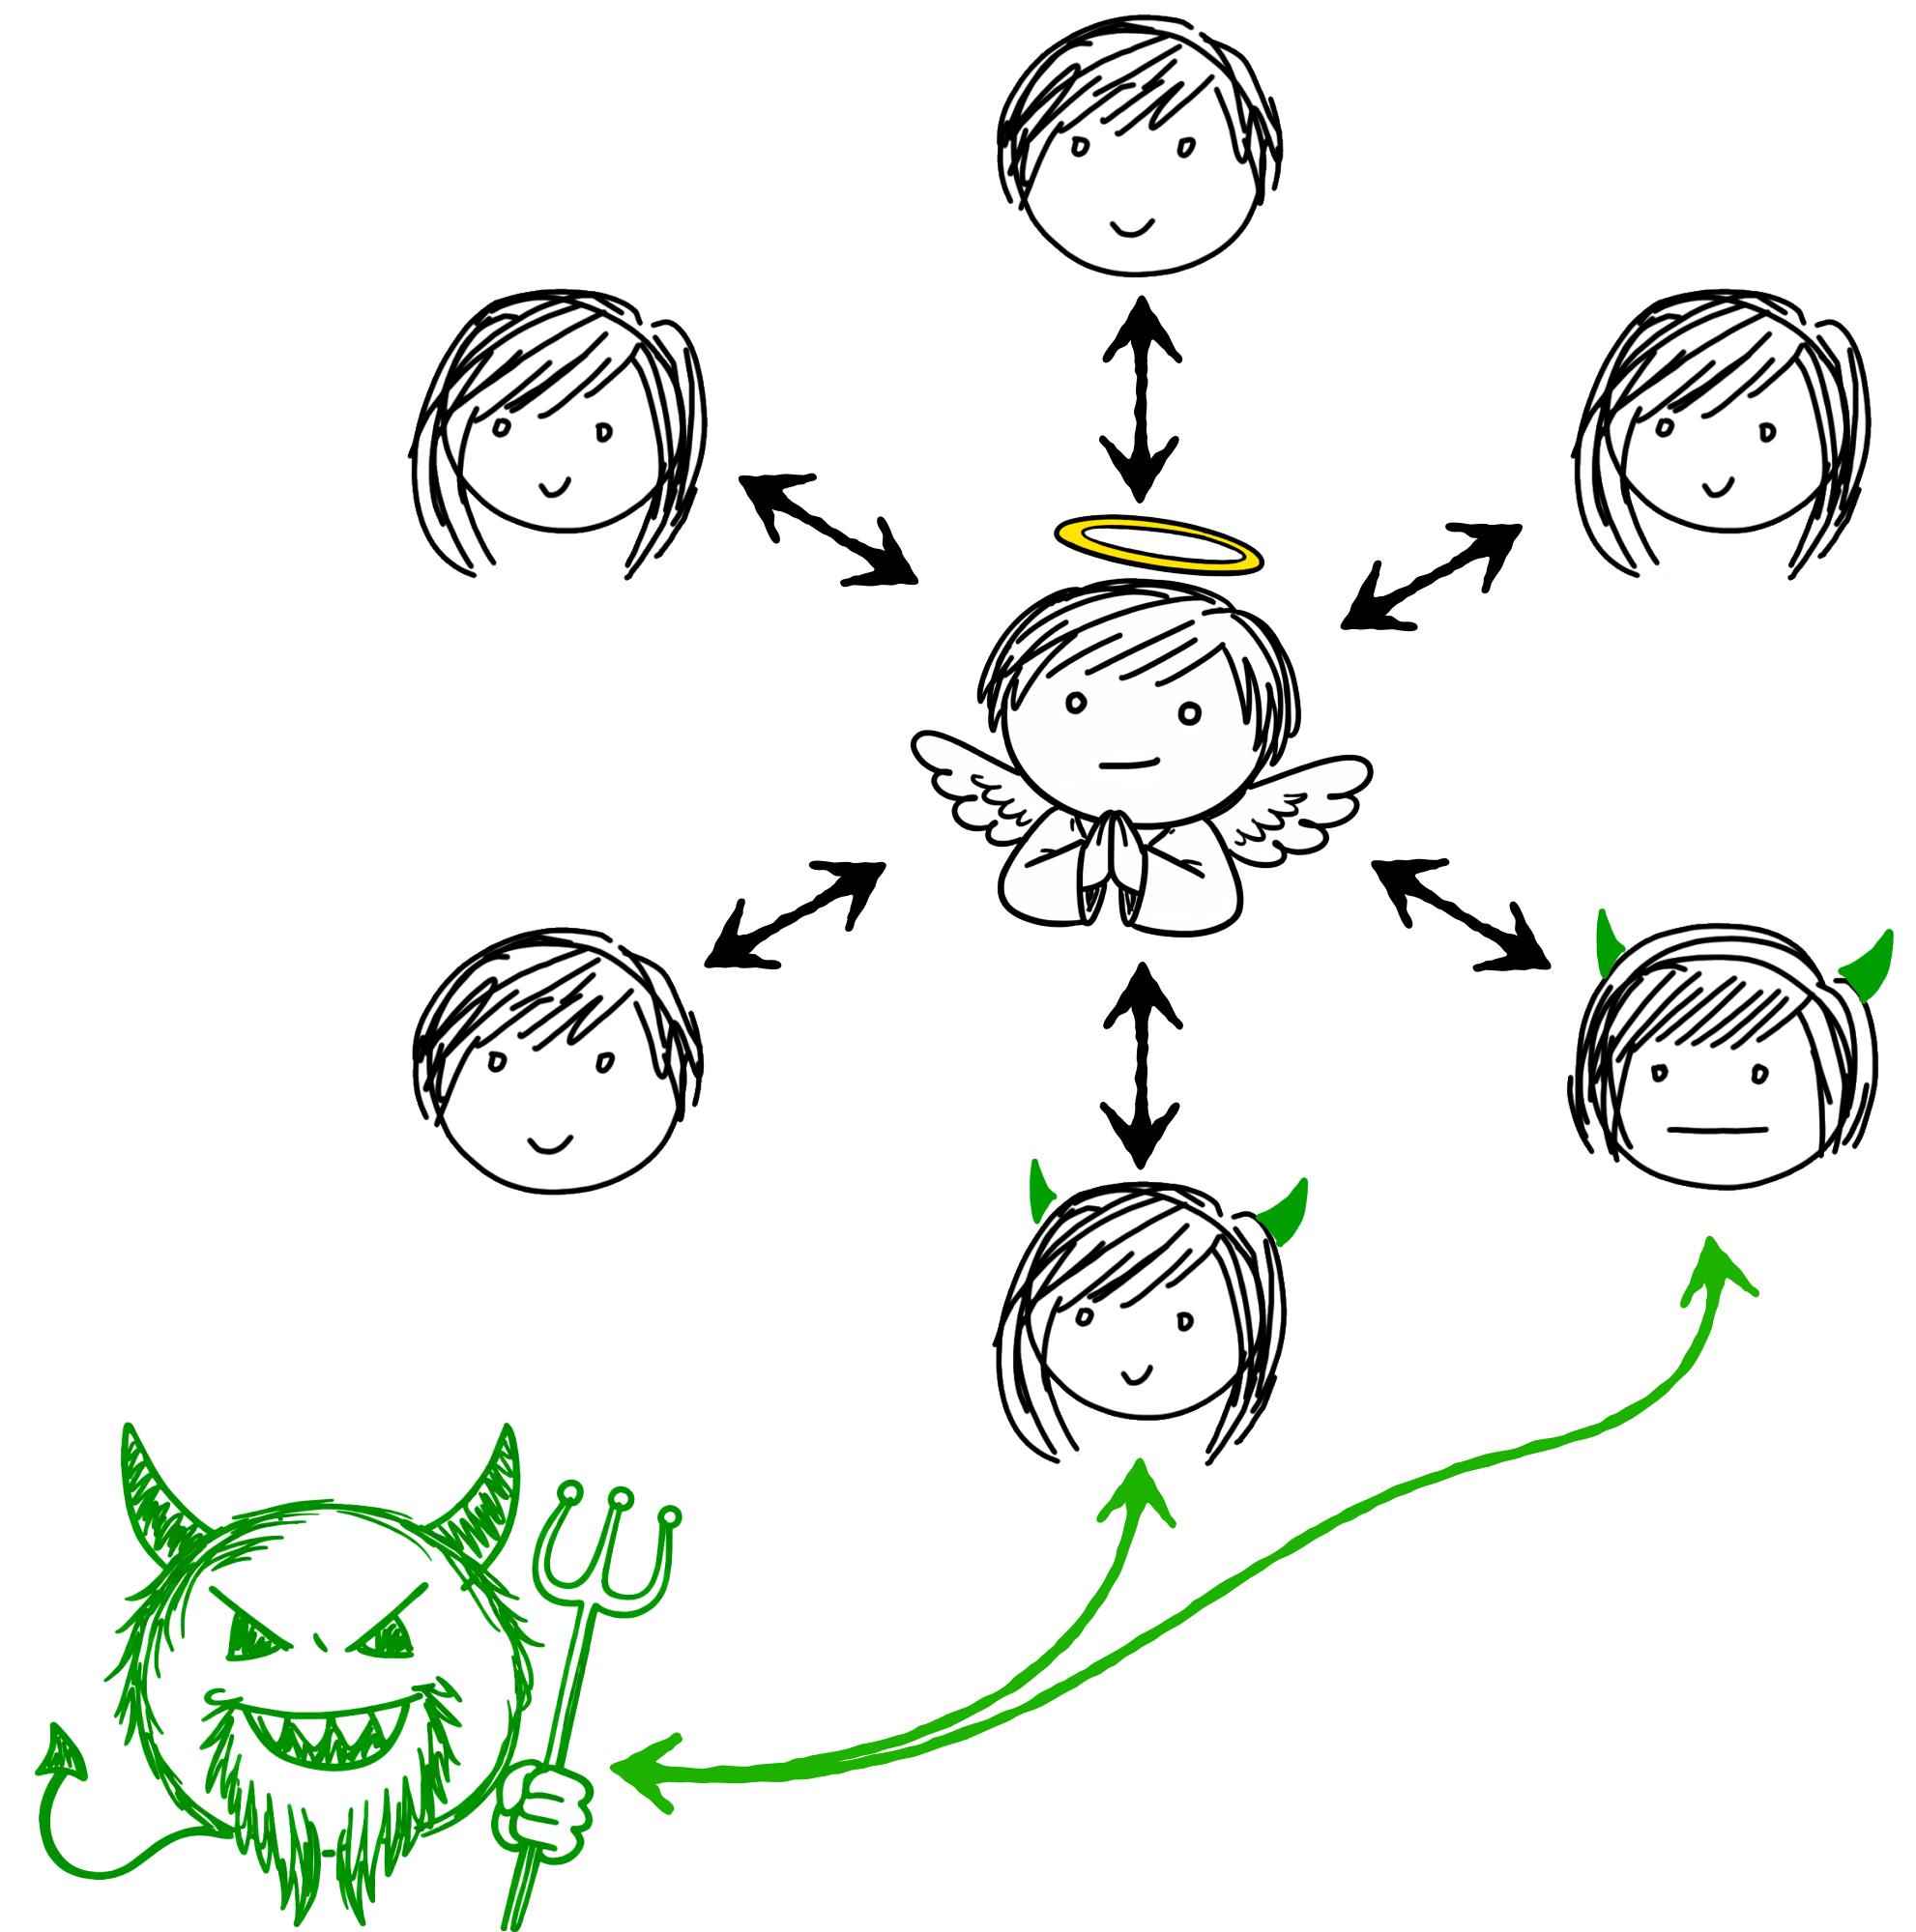
For every real adversary, and ideal adversary

Goal of $\mathcal{S}$: produce output *indistinguishable* from $\mathcal{A}$

the Simulator $\mathcal{S}$

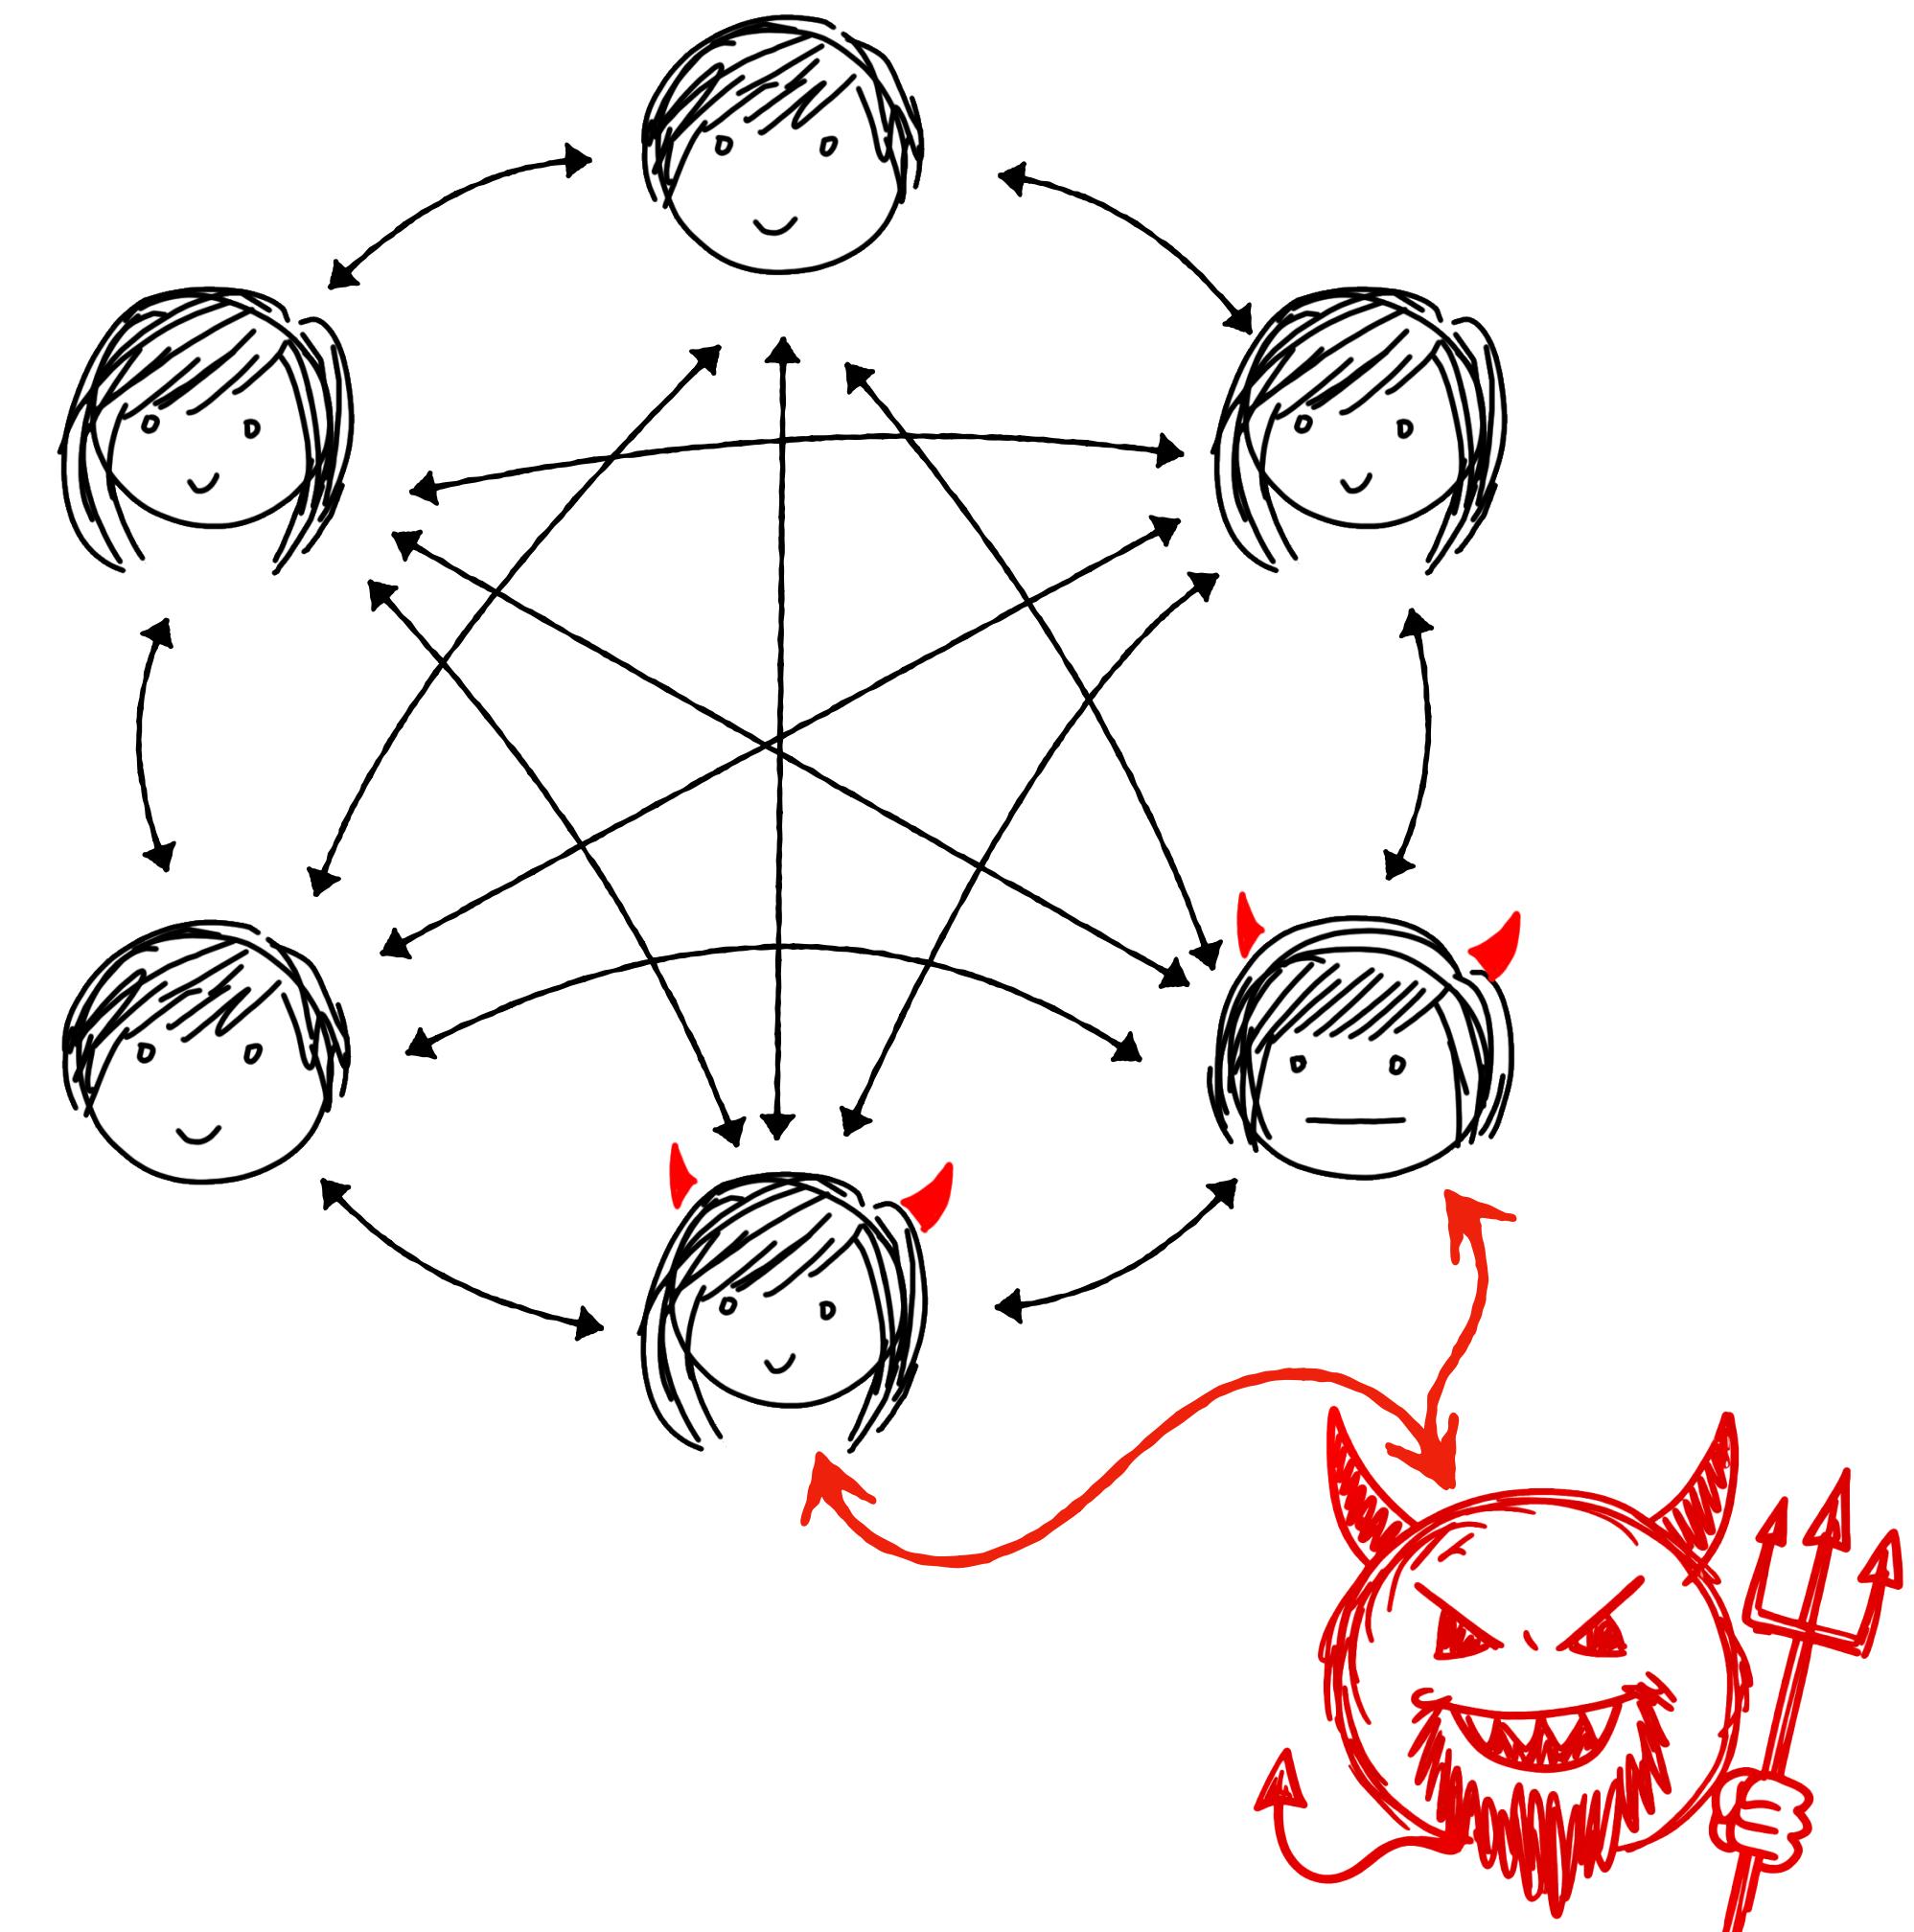the Adversary $\mathcal{A}$

Is $\mathcal{S}$ is Indistinguishable from $\mathcal{A}$? Who will Judge?

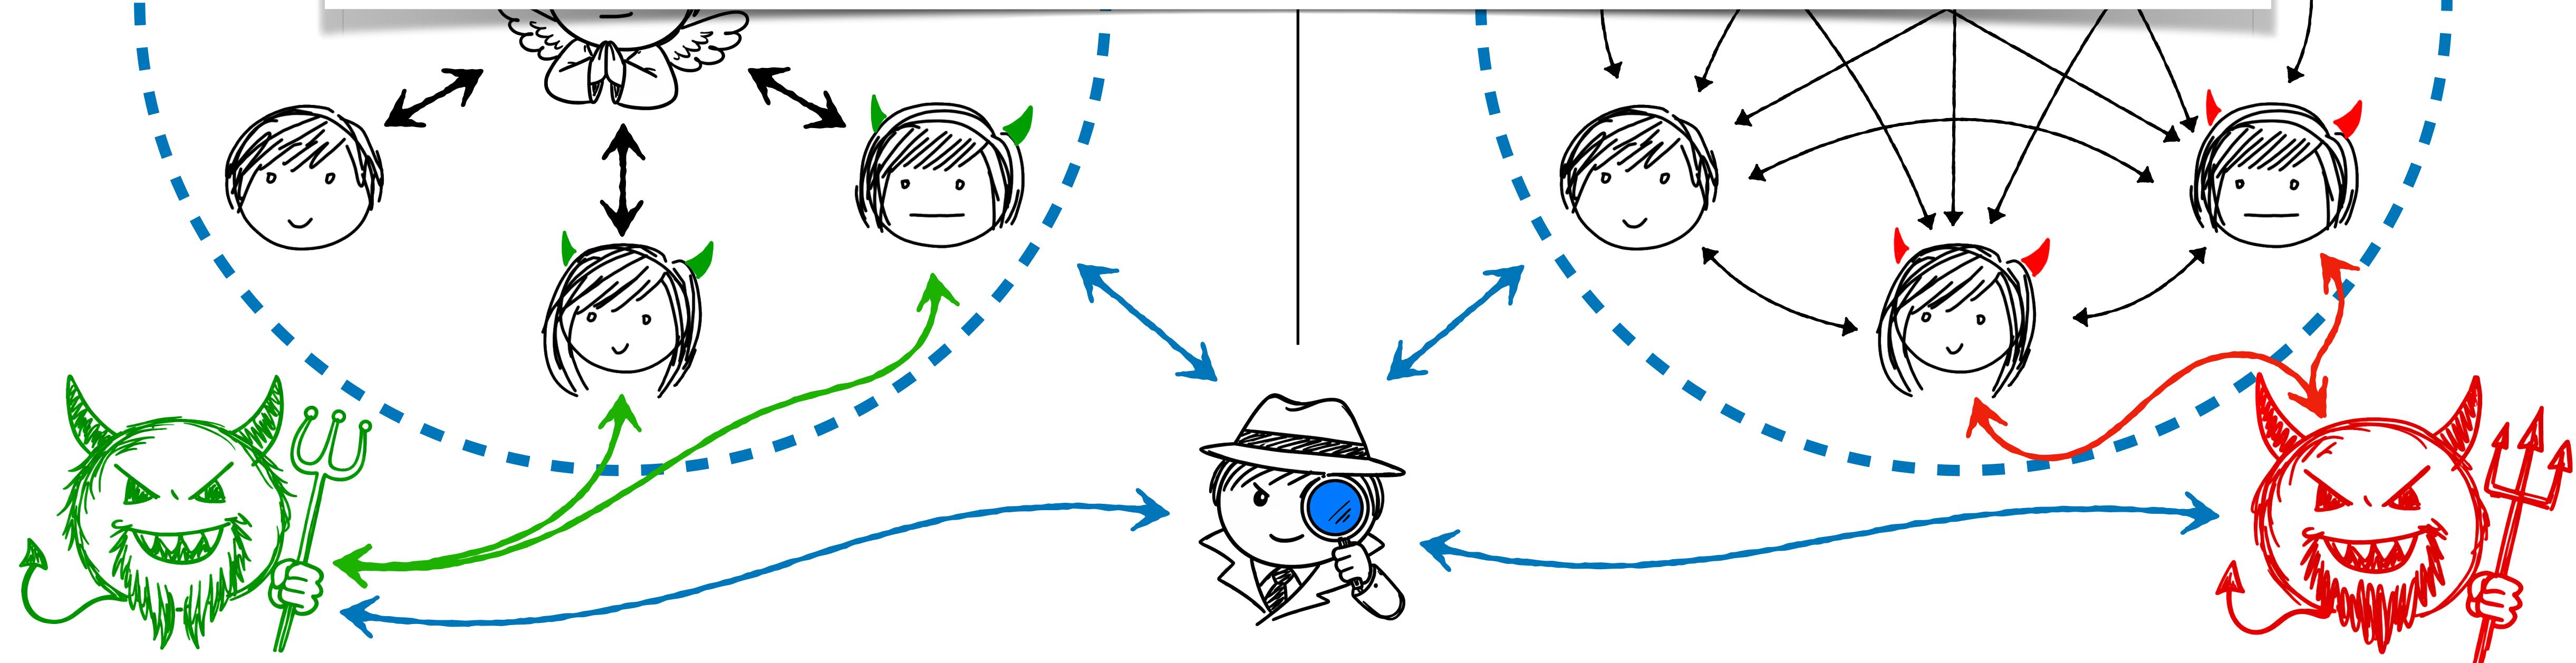the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$
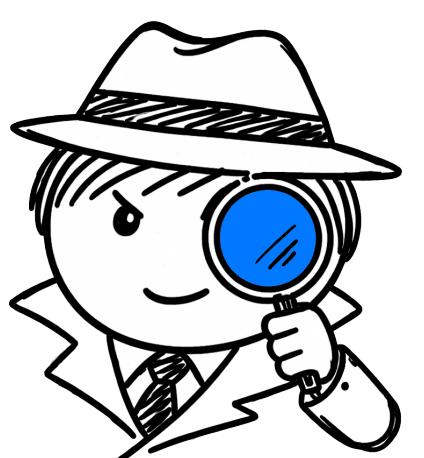
Is $\mathcal{S}$ is dge?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$

Is $\mathcal{S}$ is ...dge?

Finally, How to Define Simulation-Based Security!

The protocol π *realizes* the ideal functionality $\mathcal{F}$ in the presence of some class of adversaries if:

$\forall \; \mathcal{A}$ in the class $\; \exists \; \mathcal{S} \;$ such that $\forall \; \mathcal{D}$
$\Pr[\mathcal{D}$ guesses correctly] is *very small*

(we will define what *very small* means later)

$\approx$
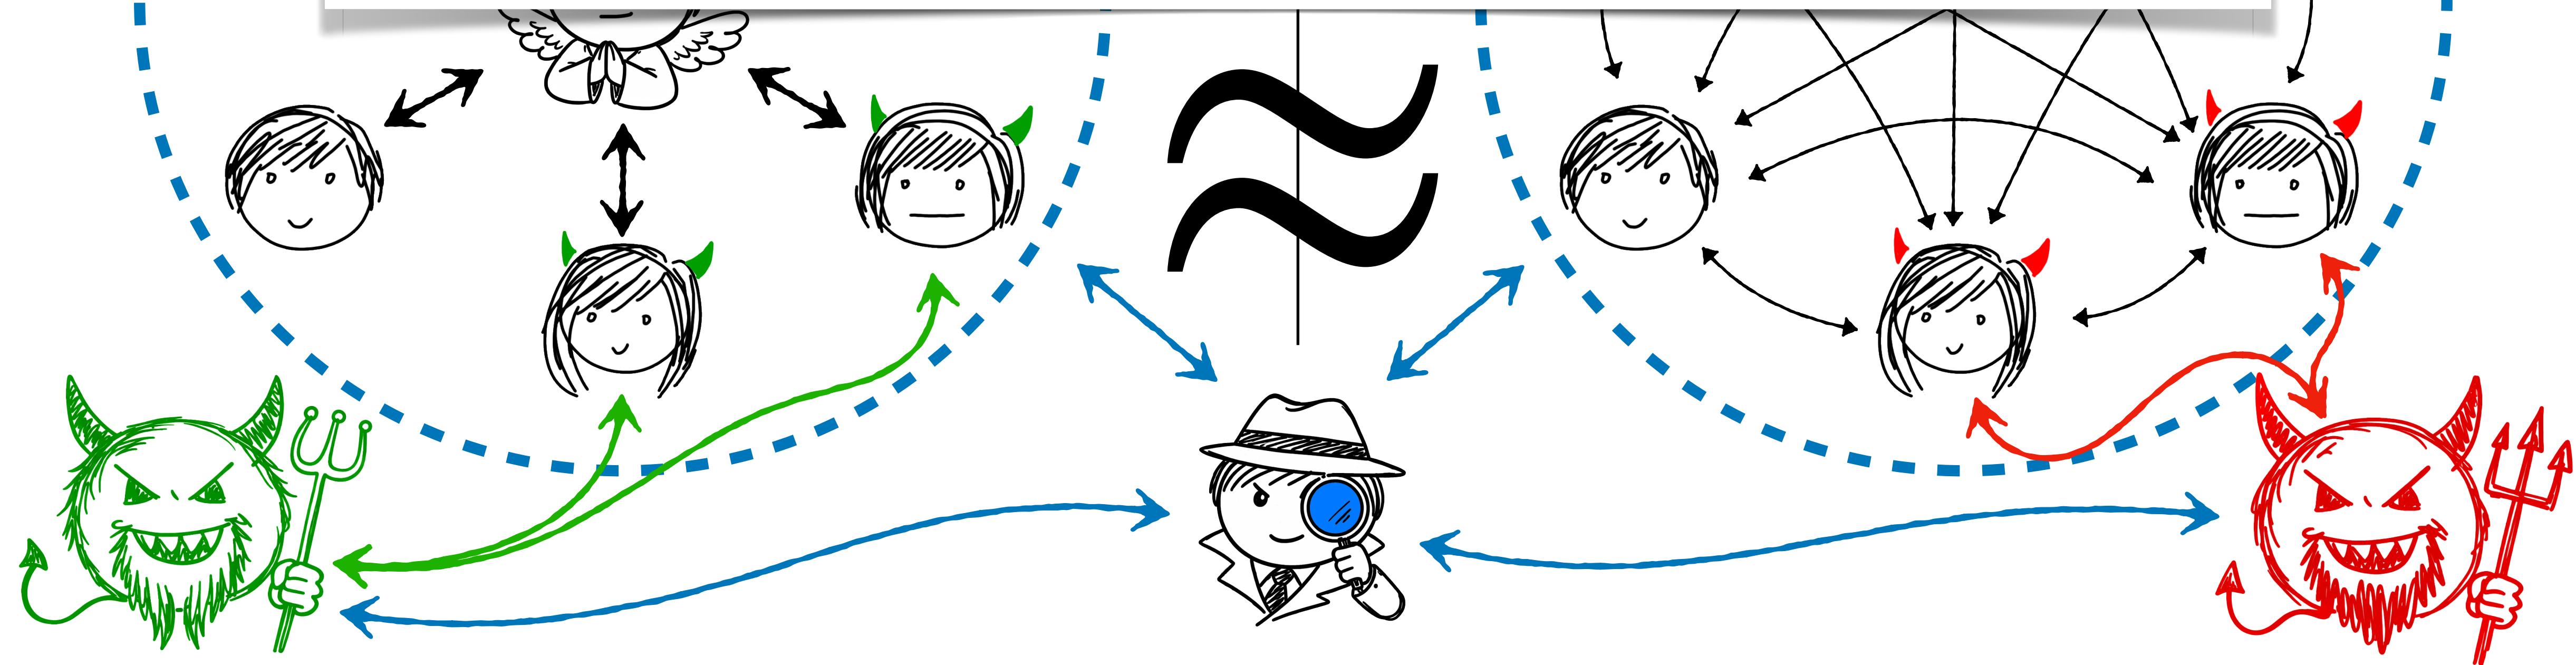
the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$

The Simulation-Based Security Paradigm
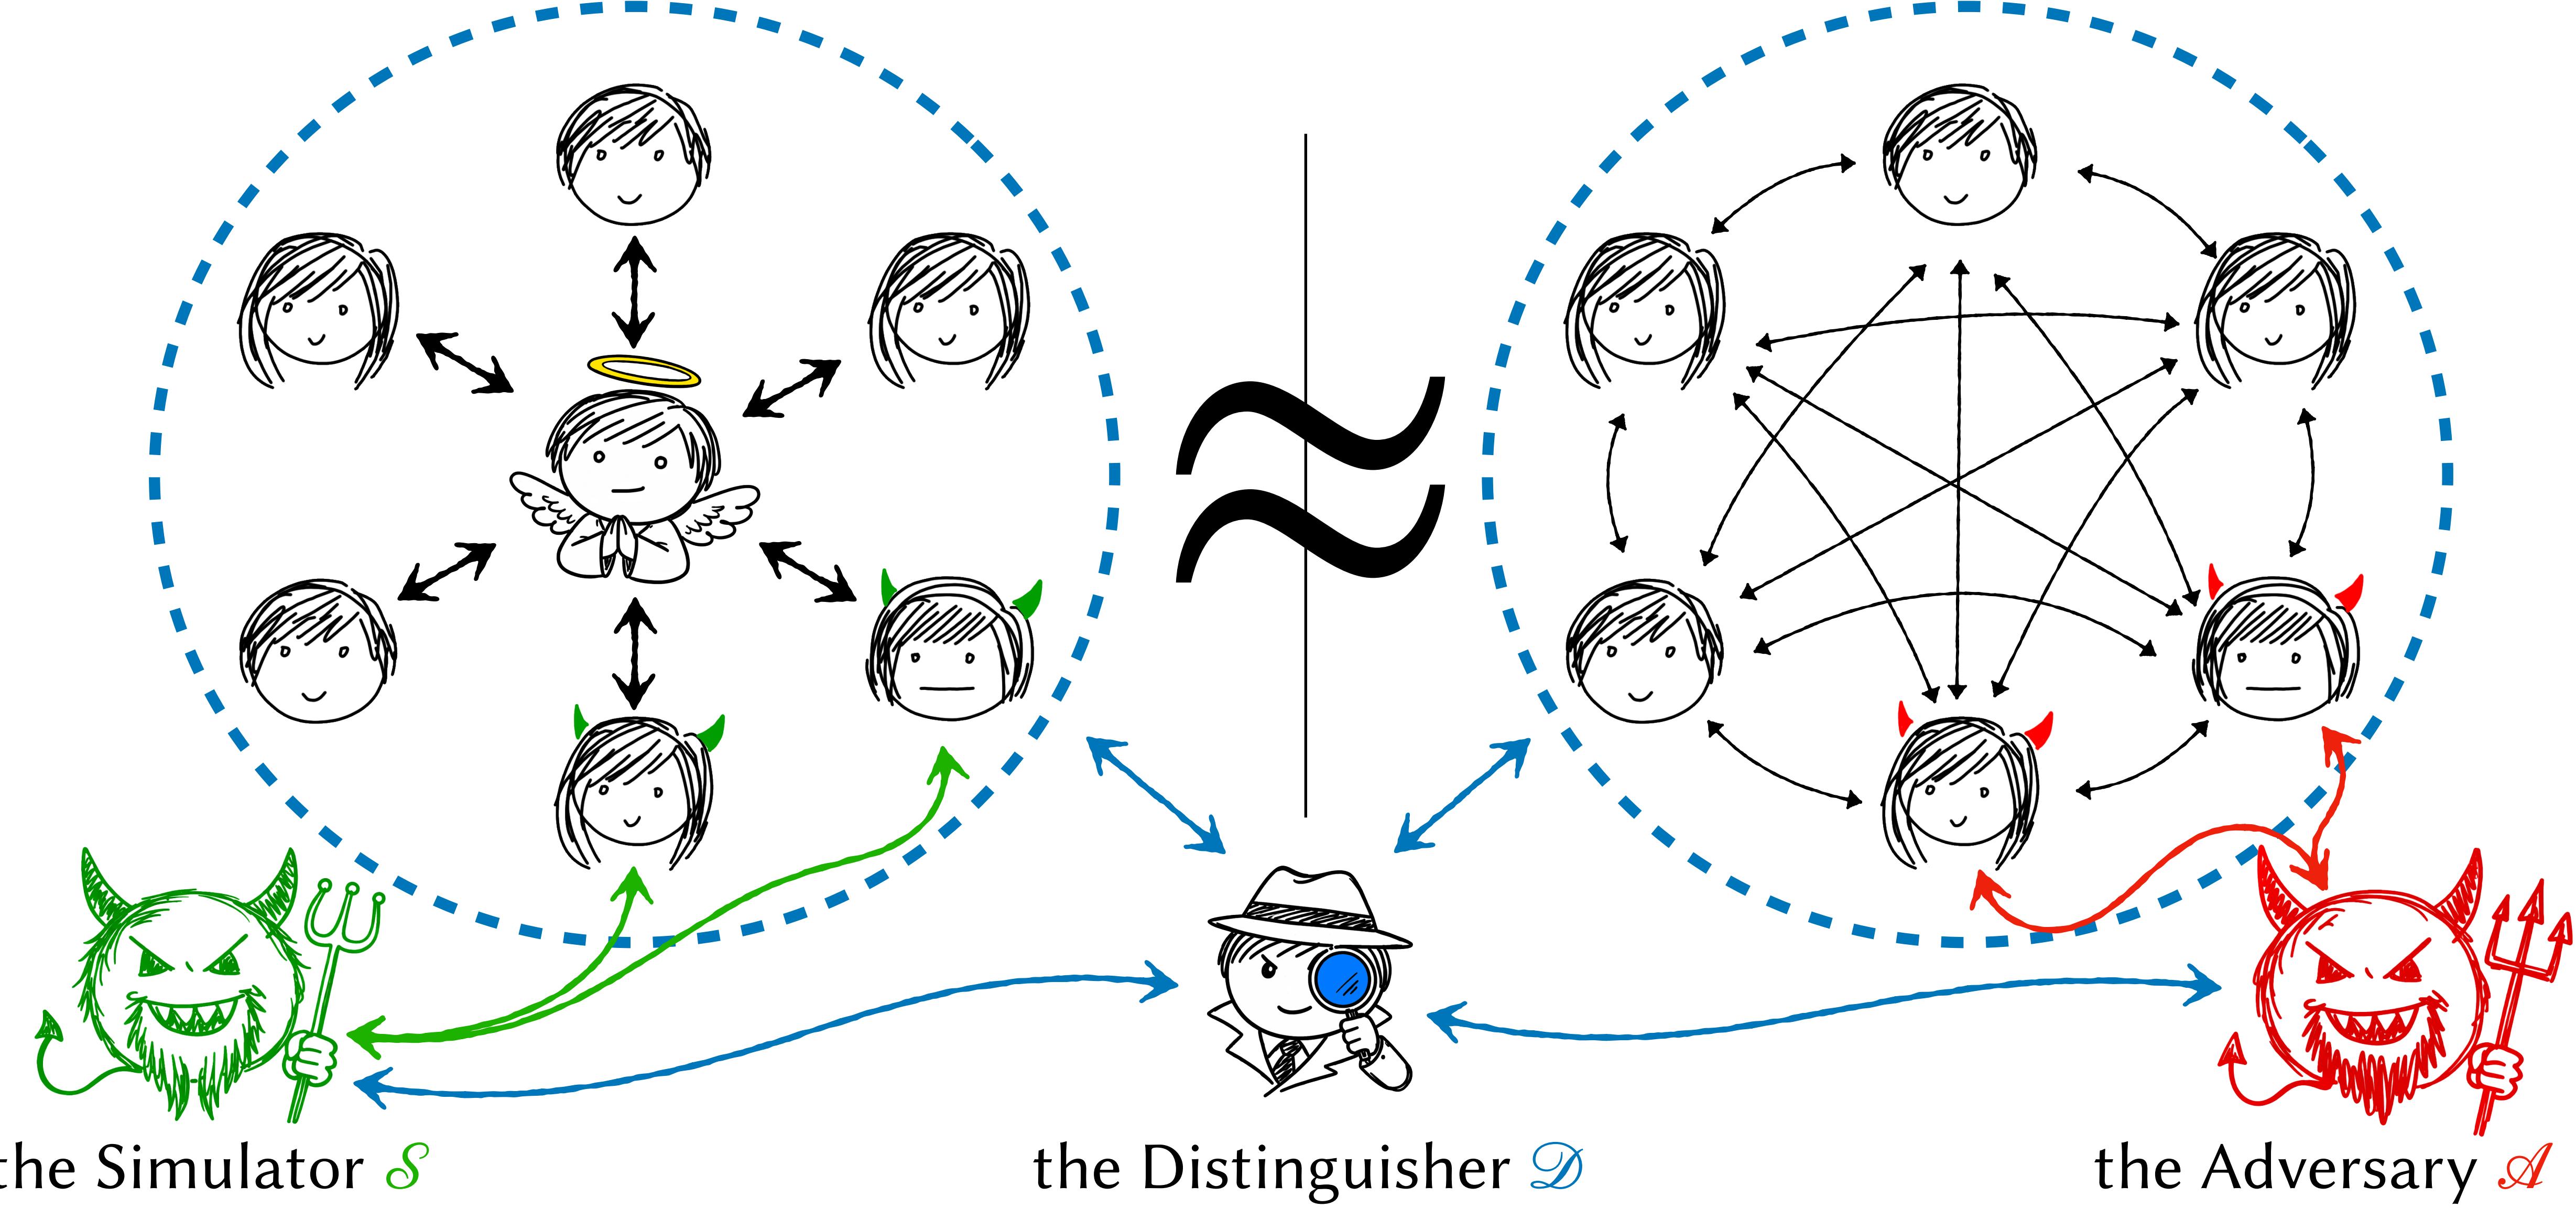
the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$

Sanity Check: Correctness?

the Simulator $\mathcal{S}$

the Distinguisher $\mathcal{D}$

the Adversary $\mathcal{A}$

Sanity Check: Correctness?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$

Sanity Check: Privacy?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$
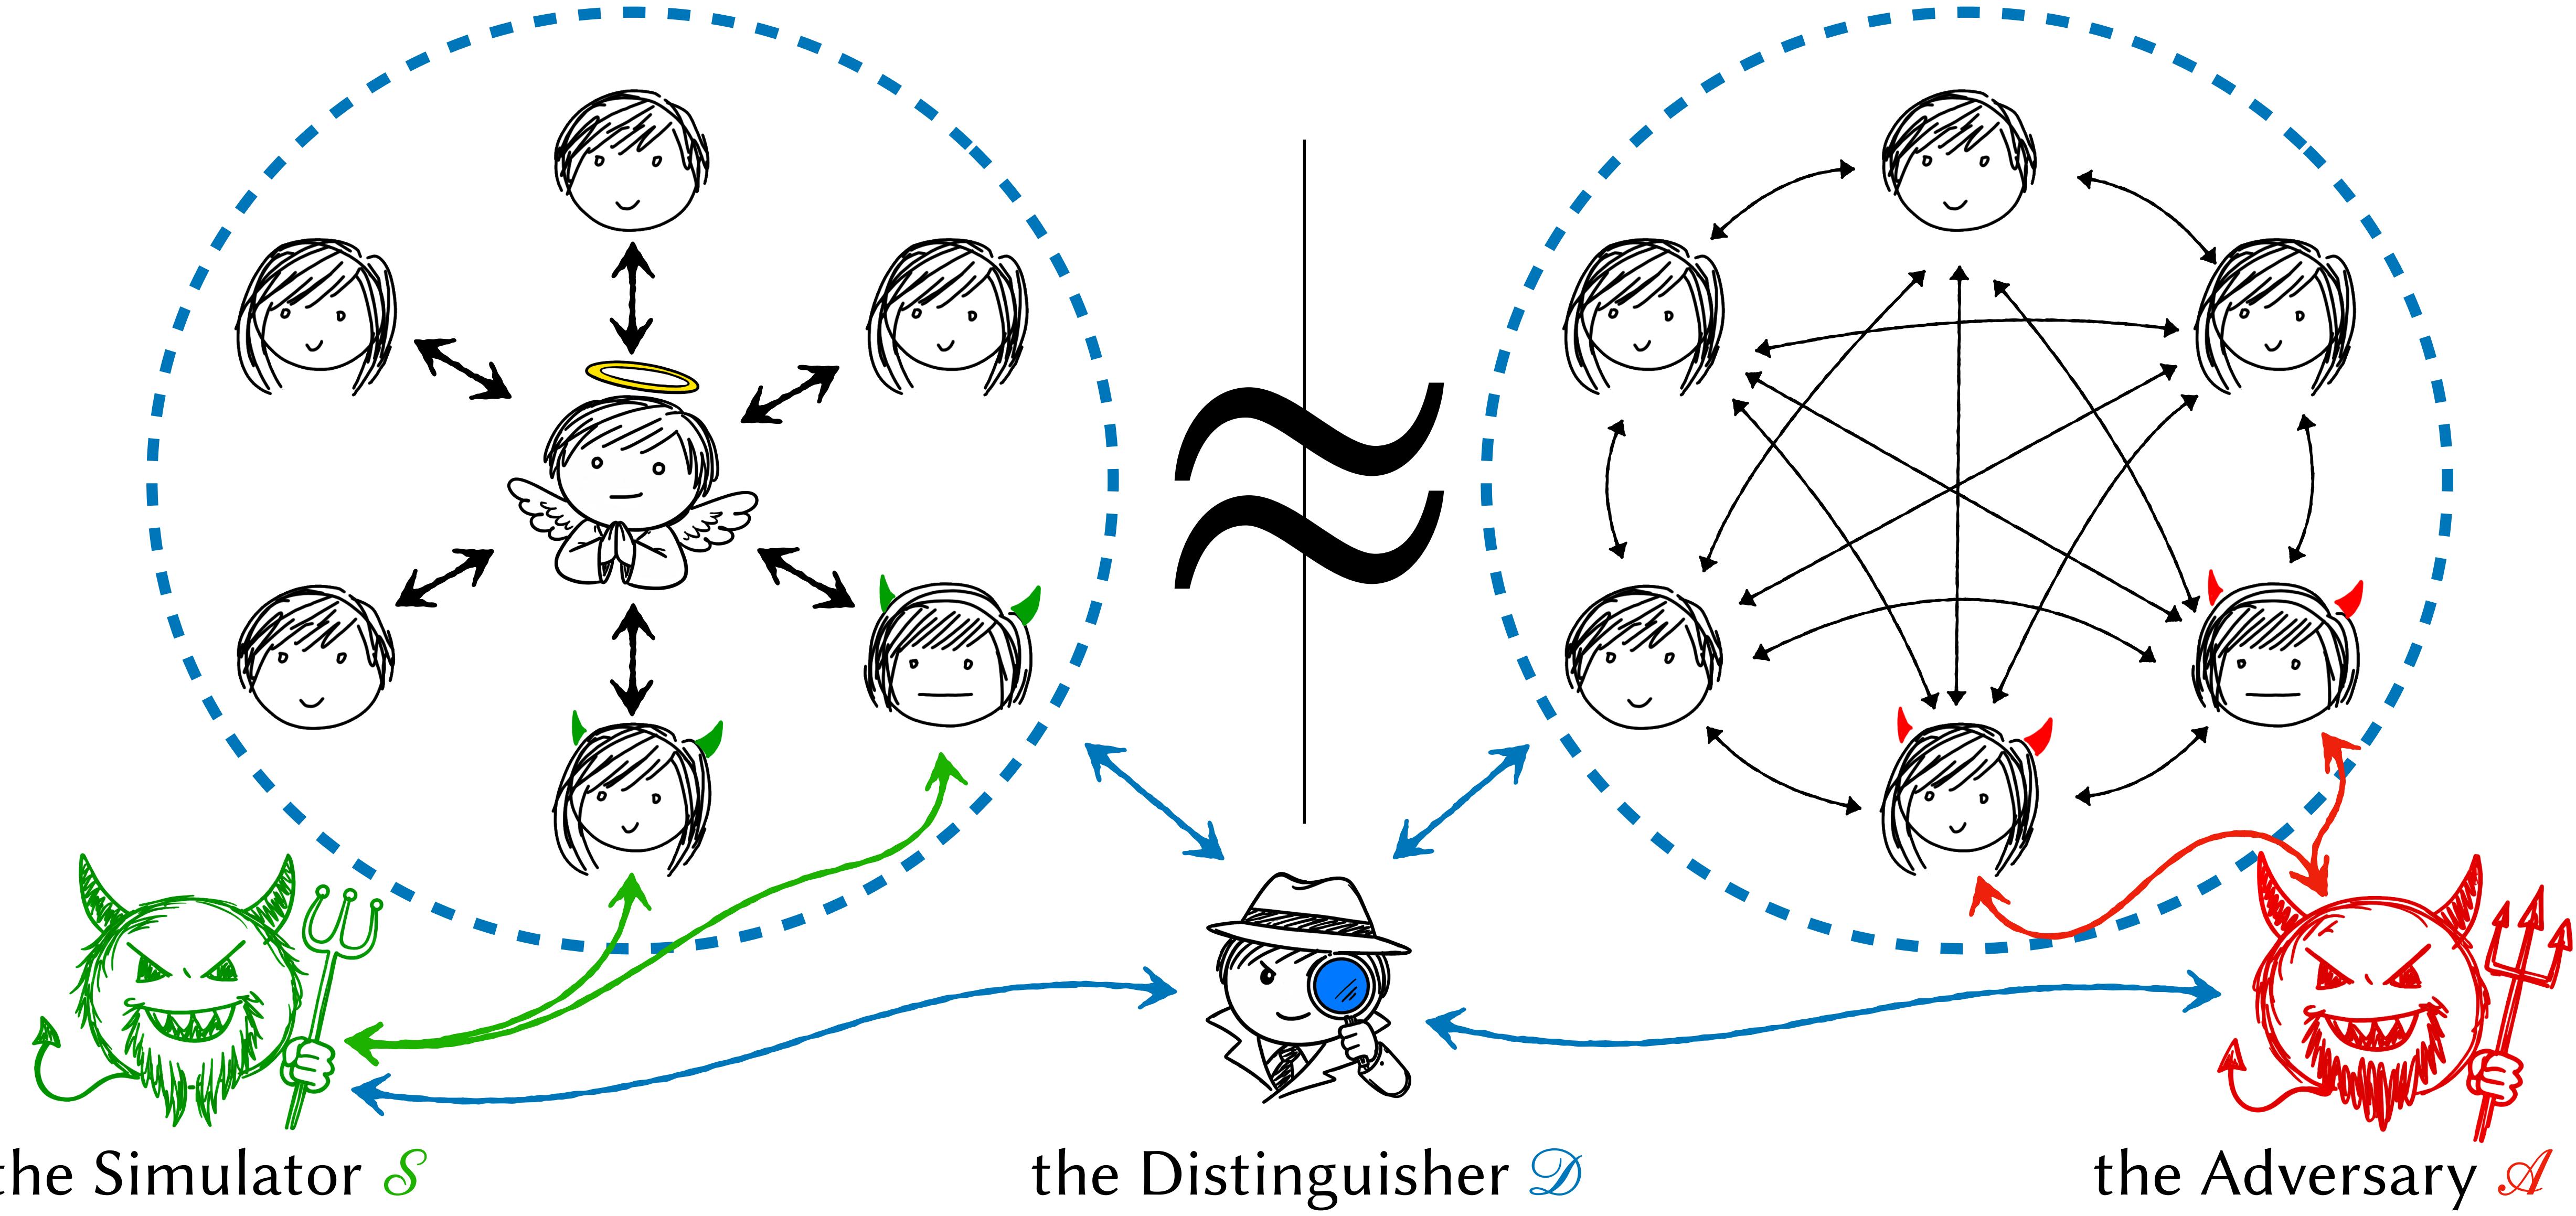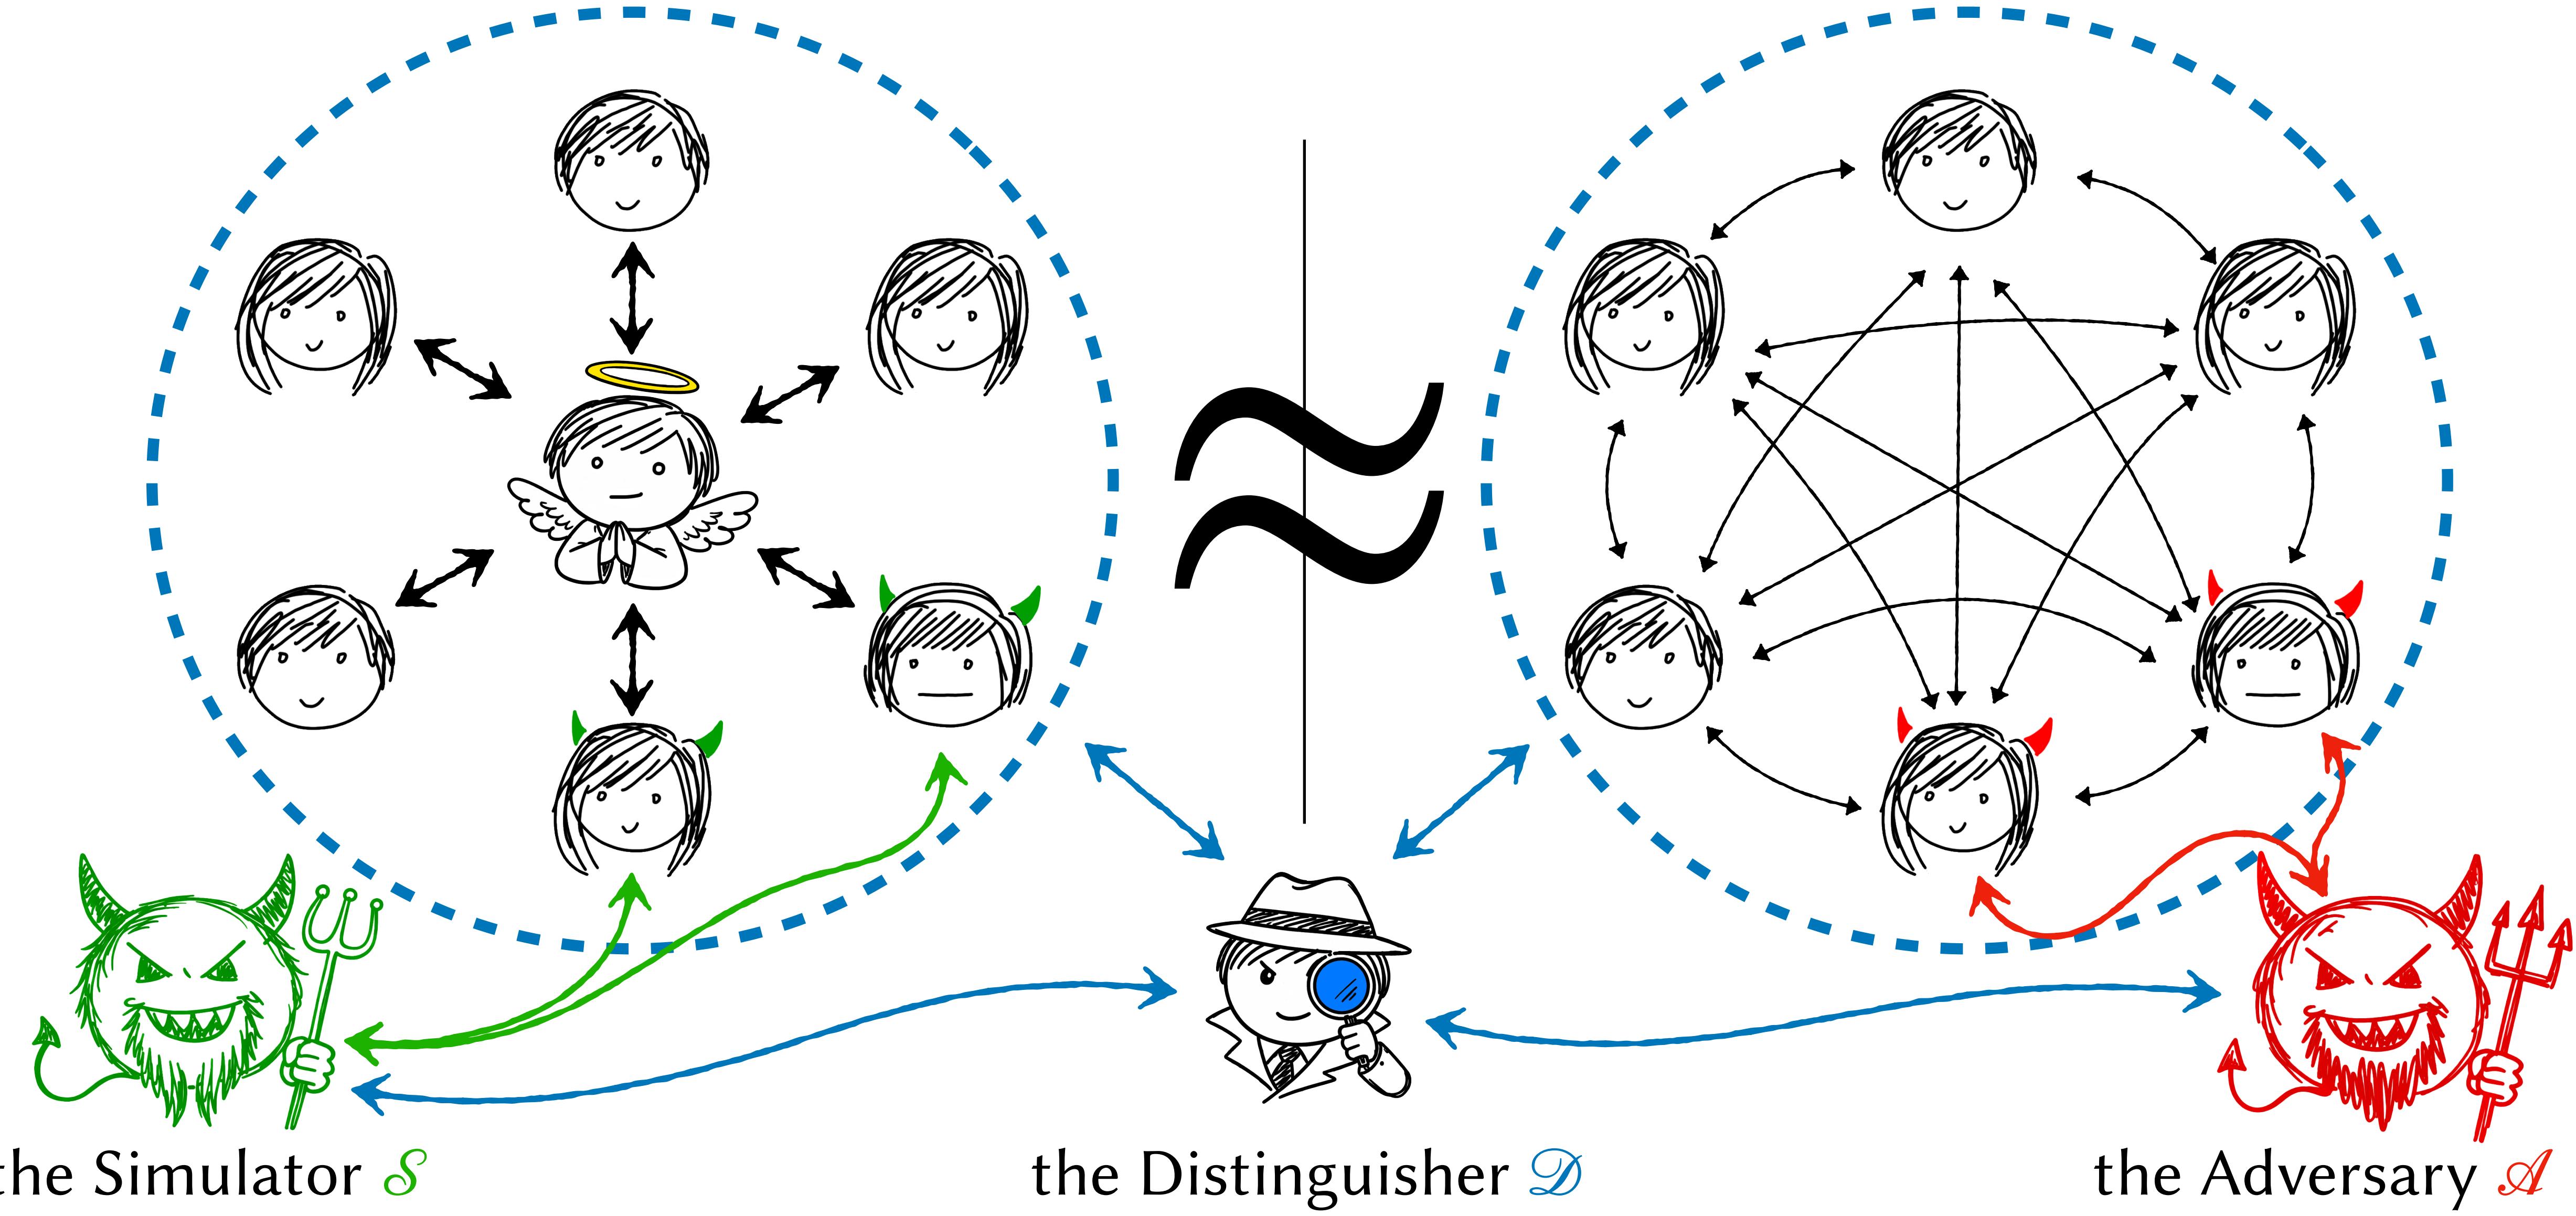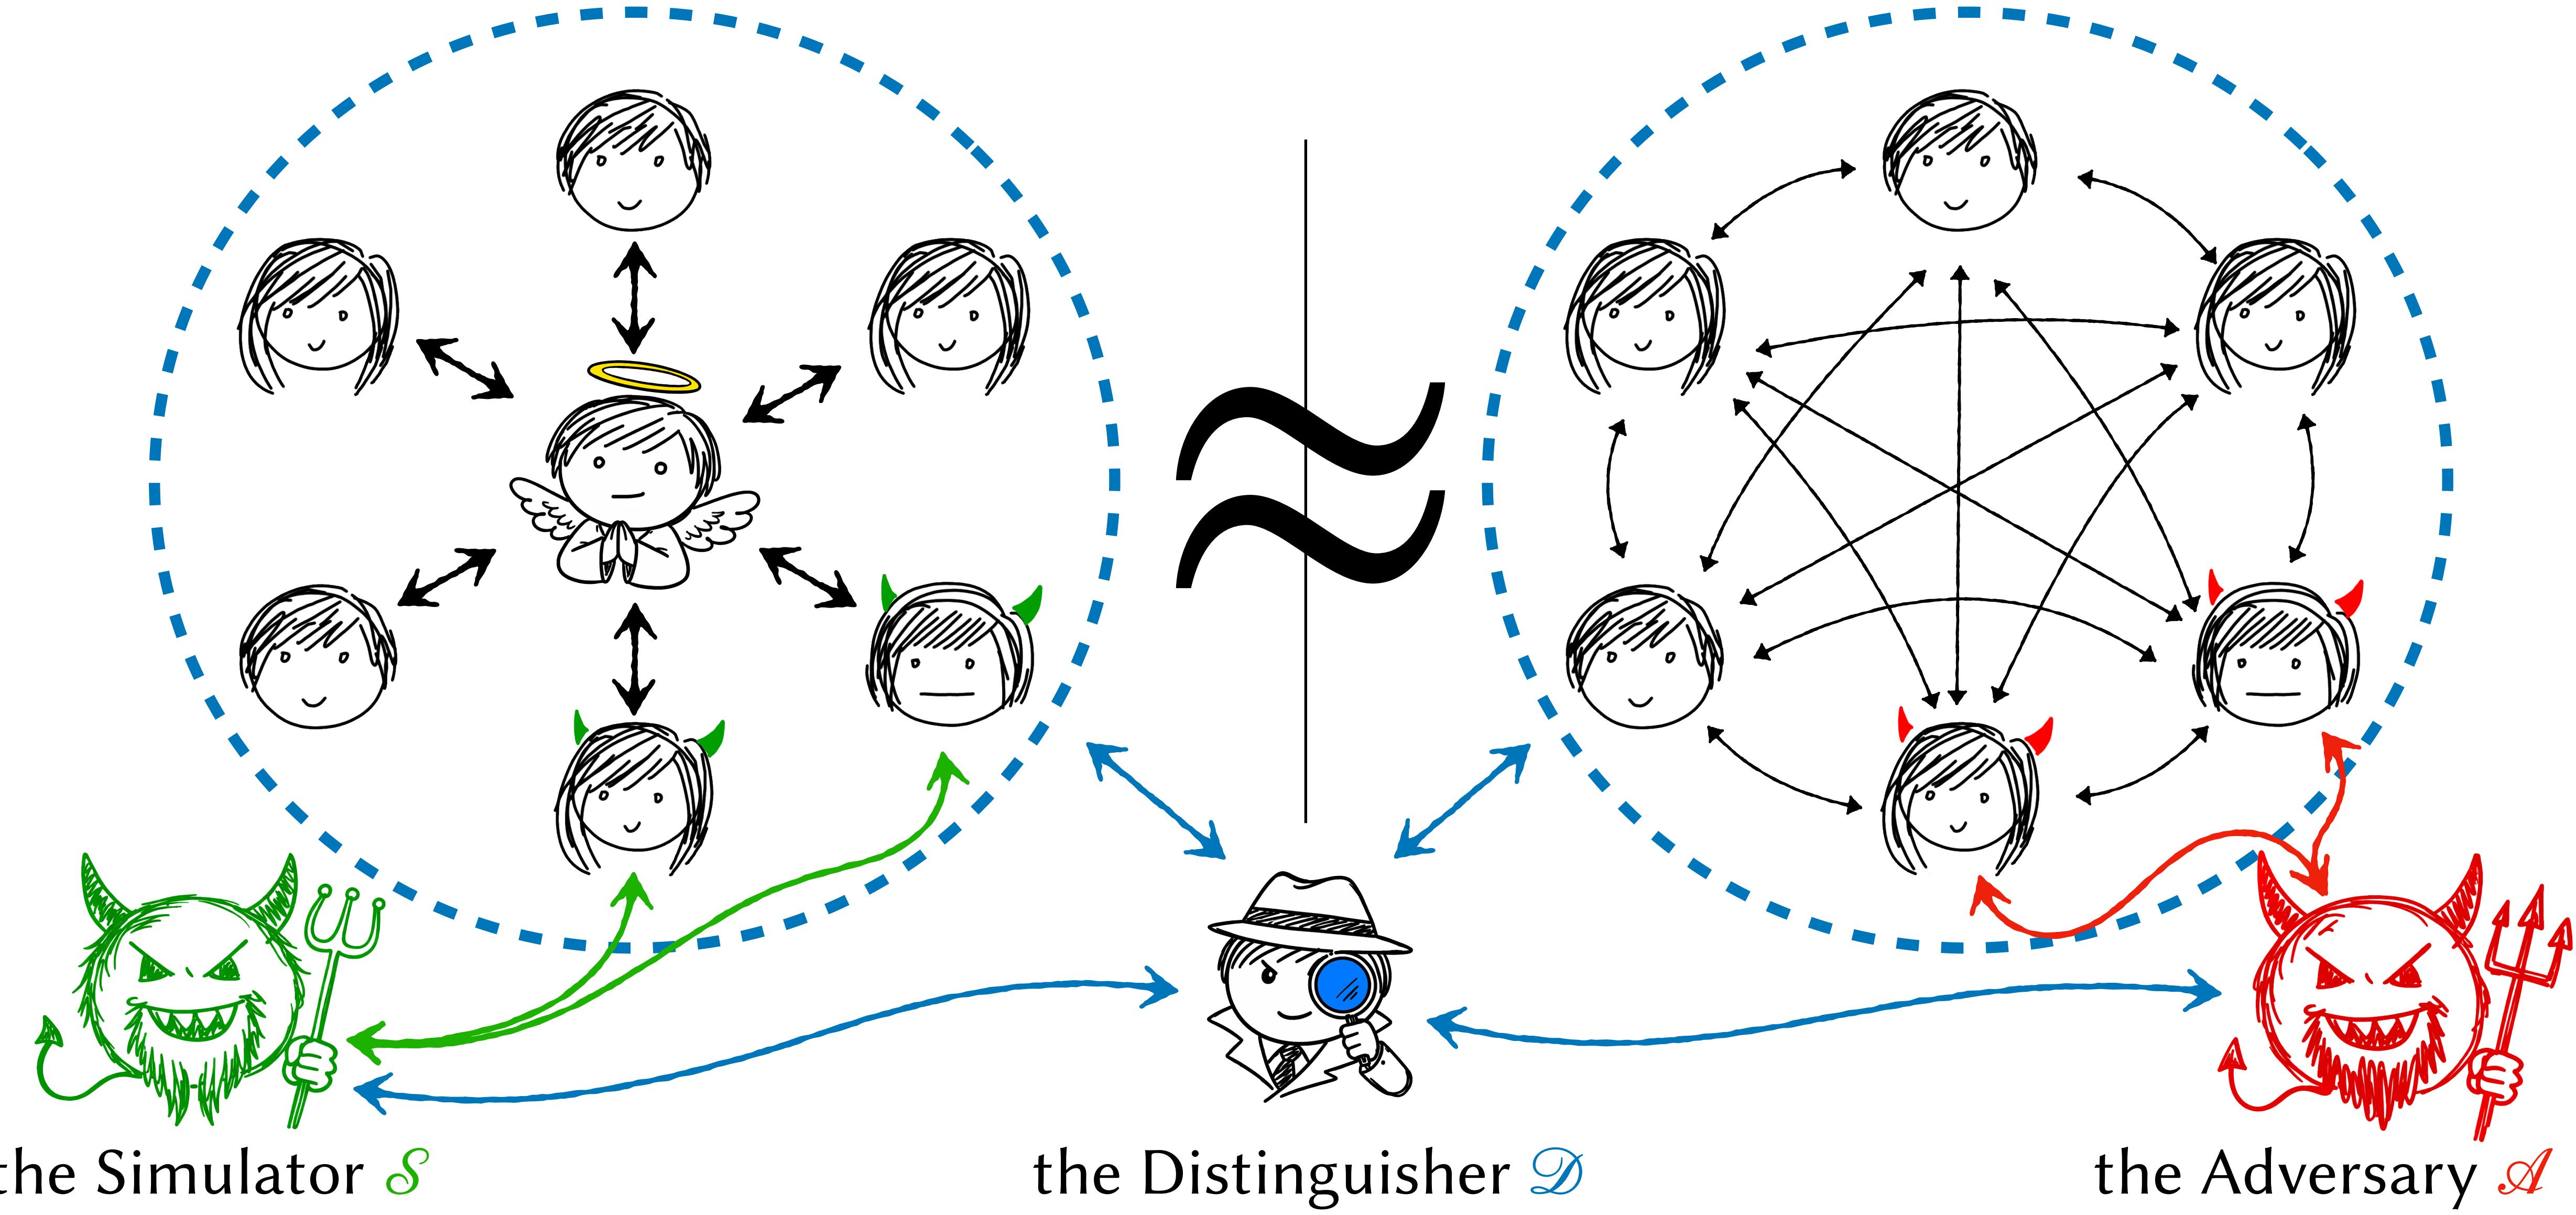
Sanity Check: Input Independence?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$

Sanity Check: Fairness?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$
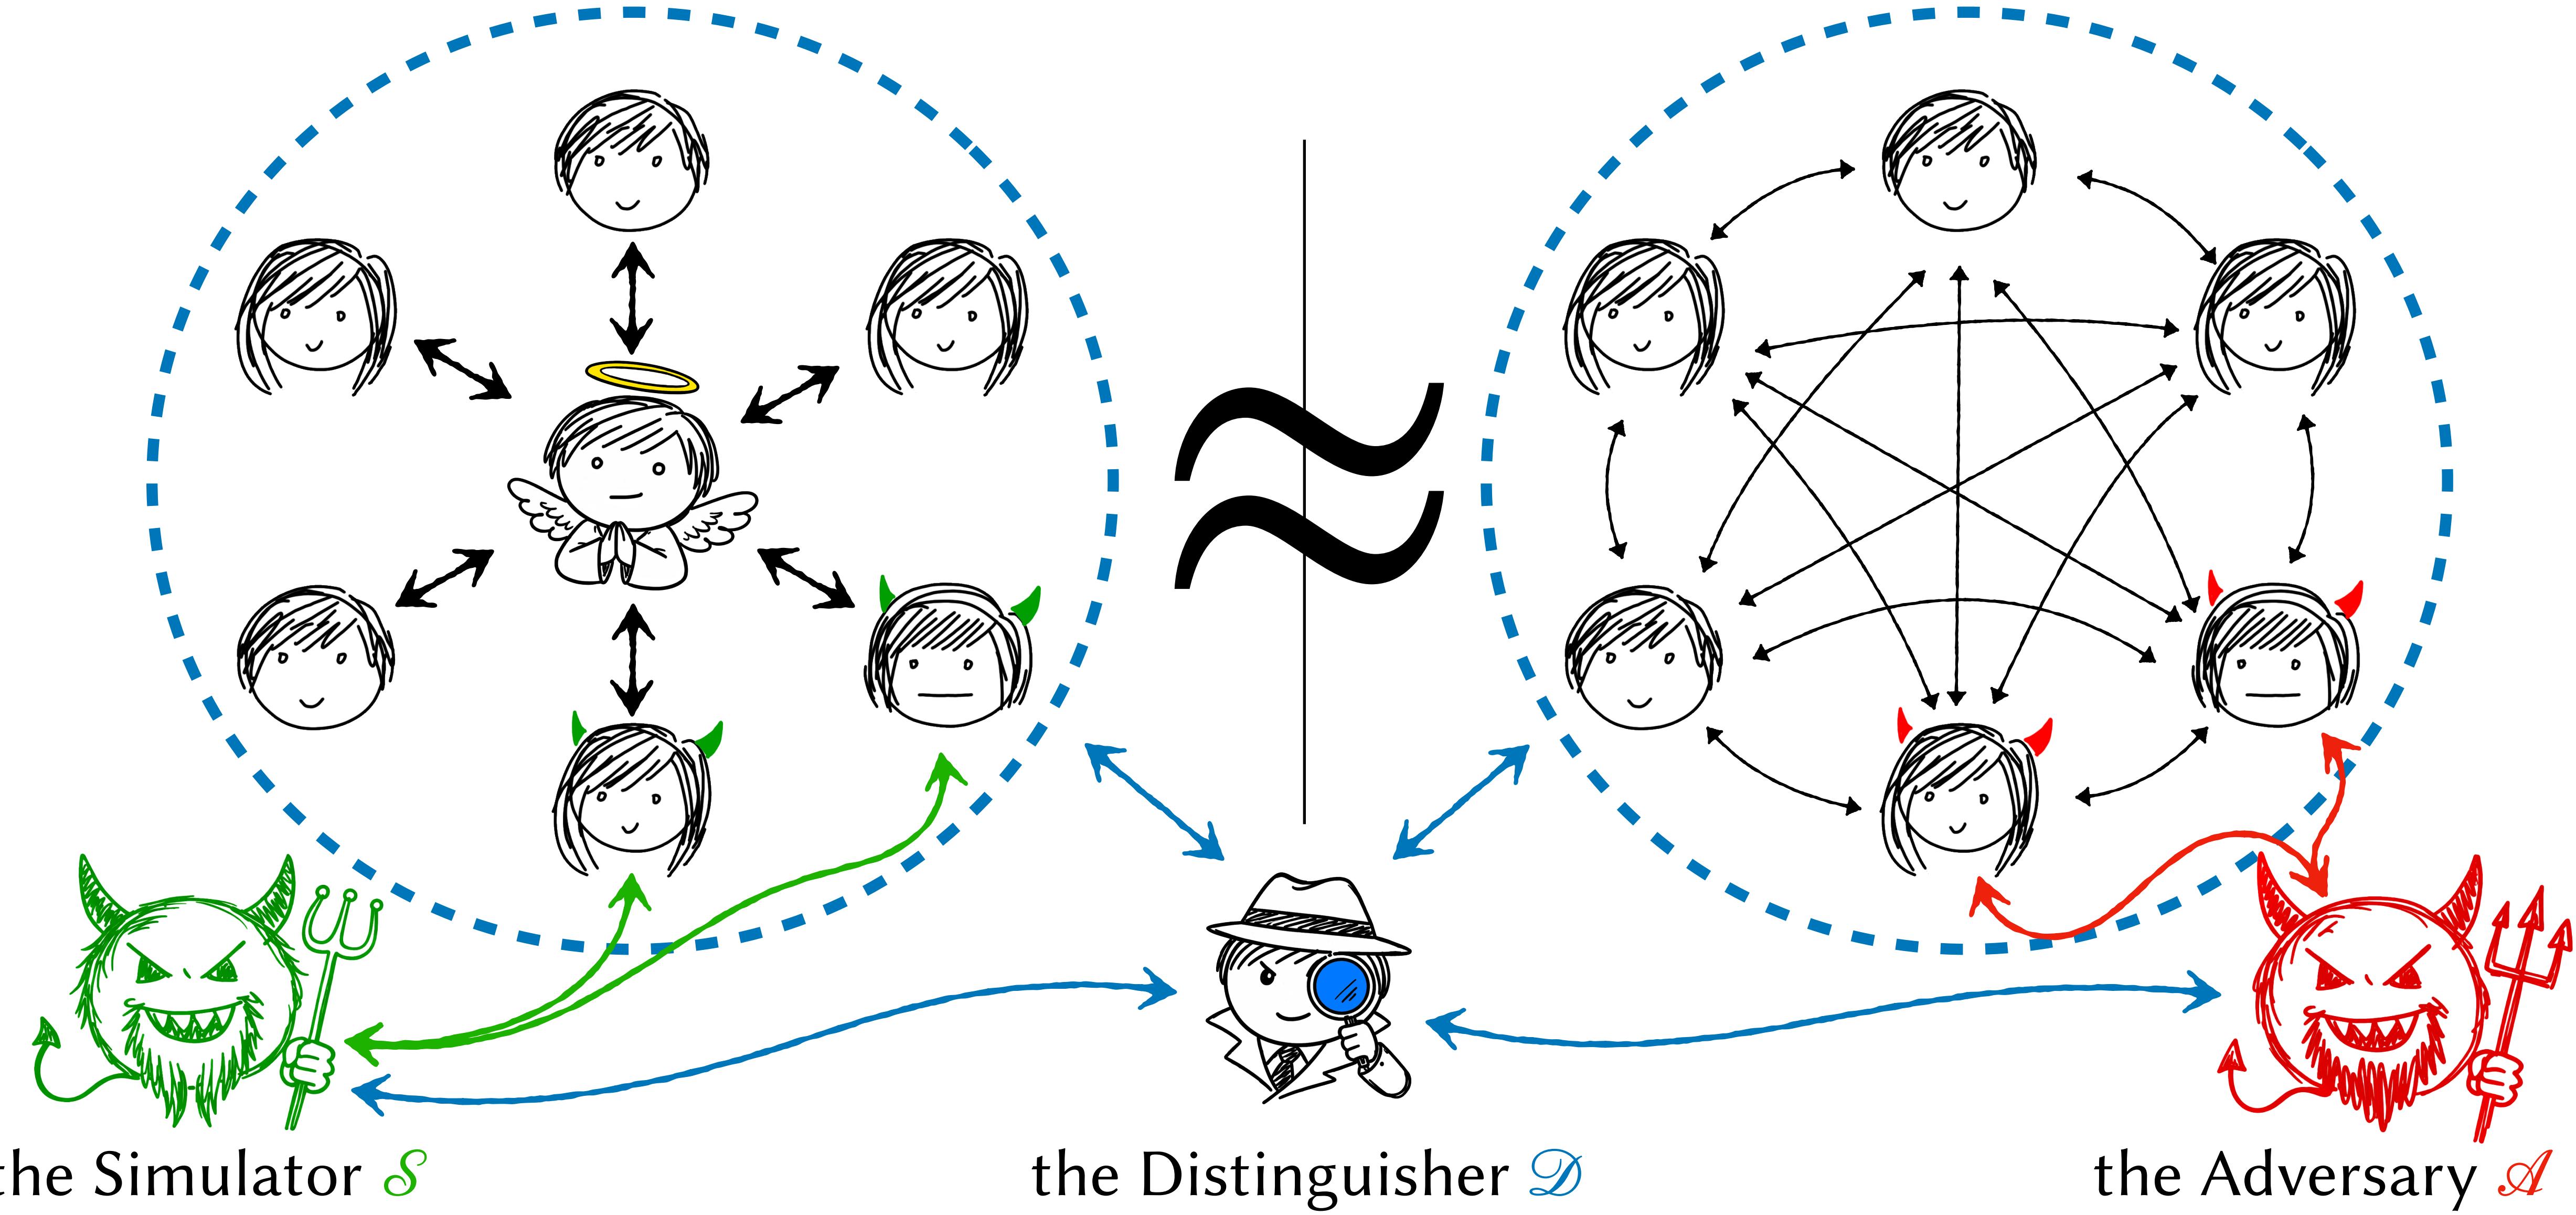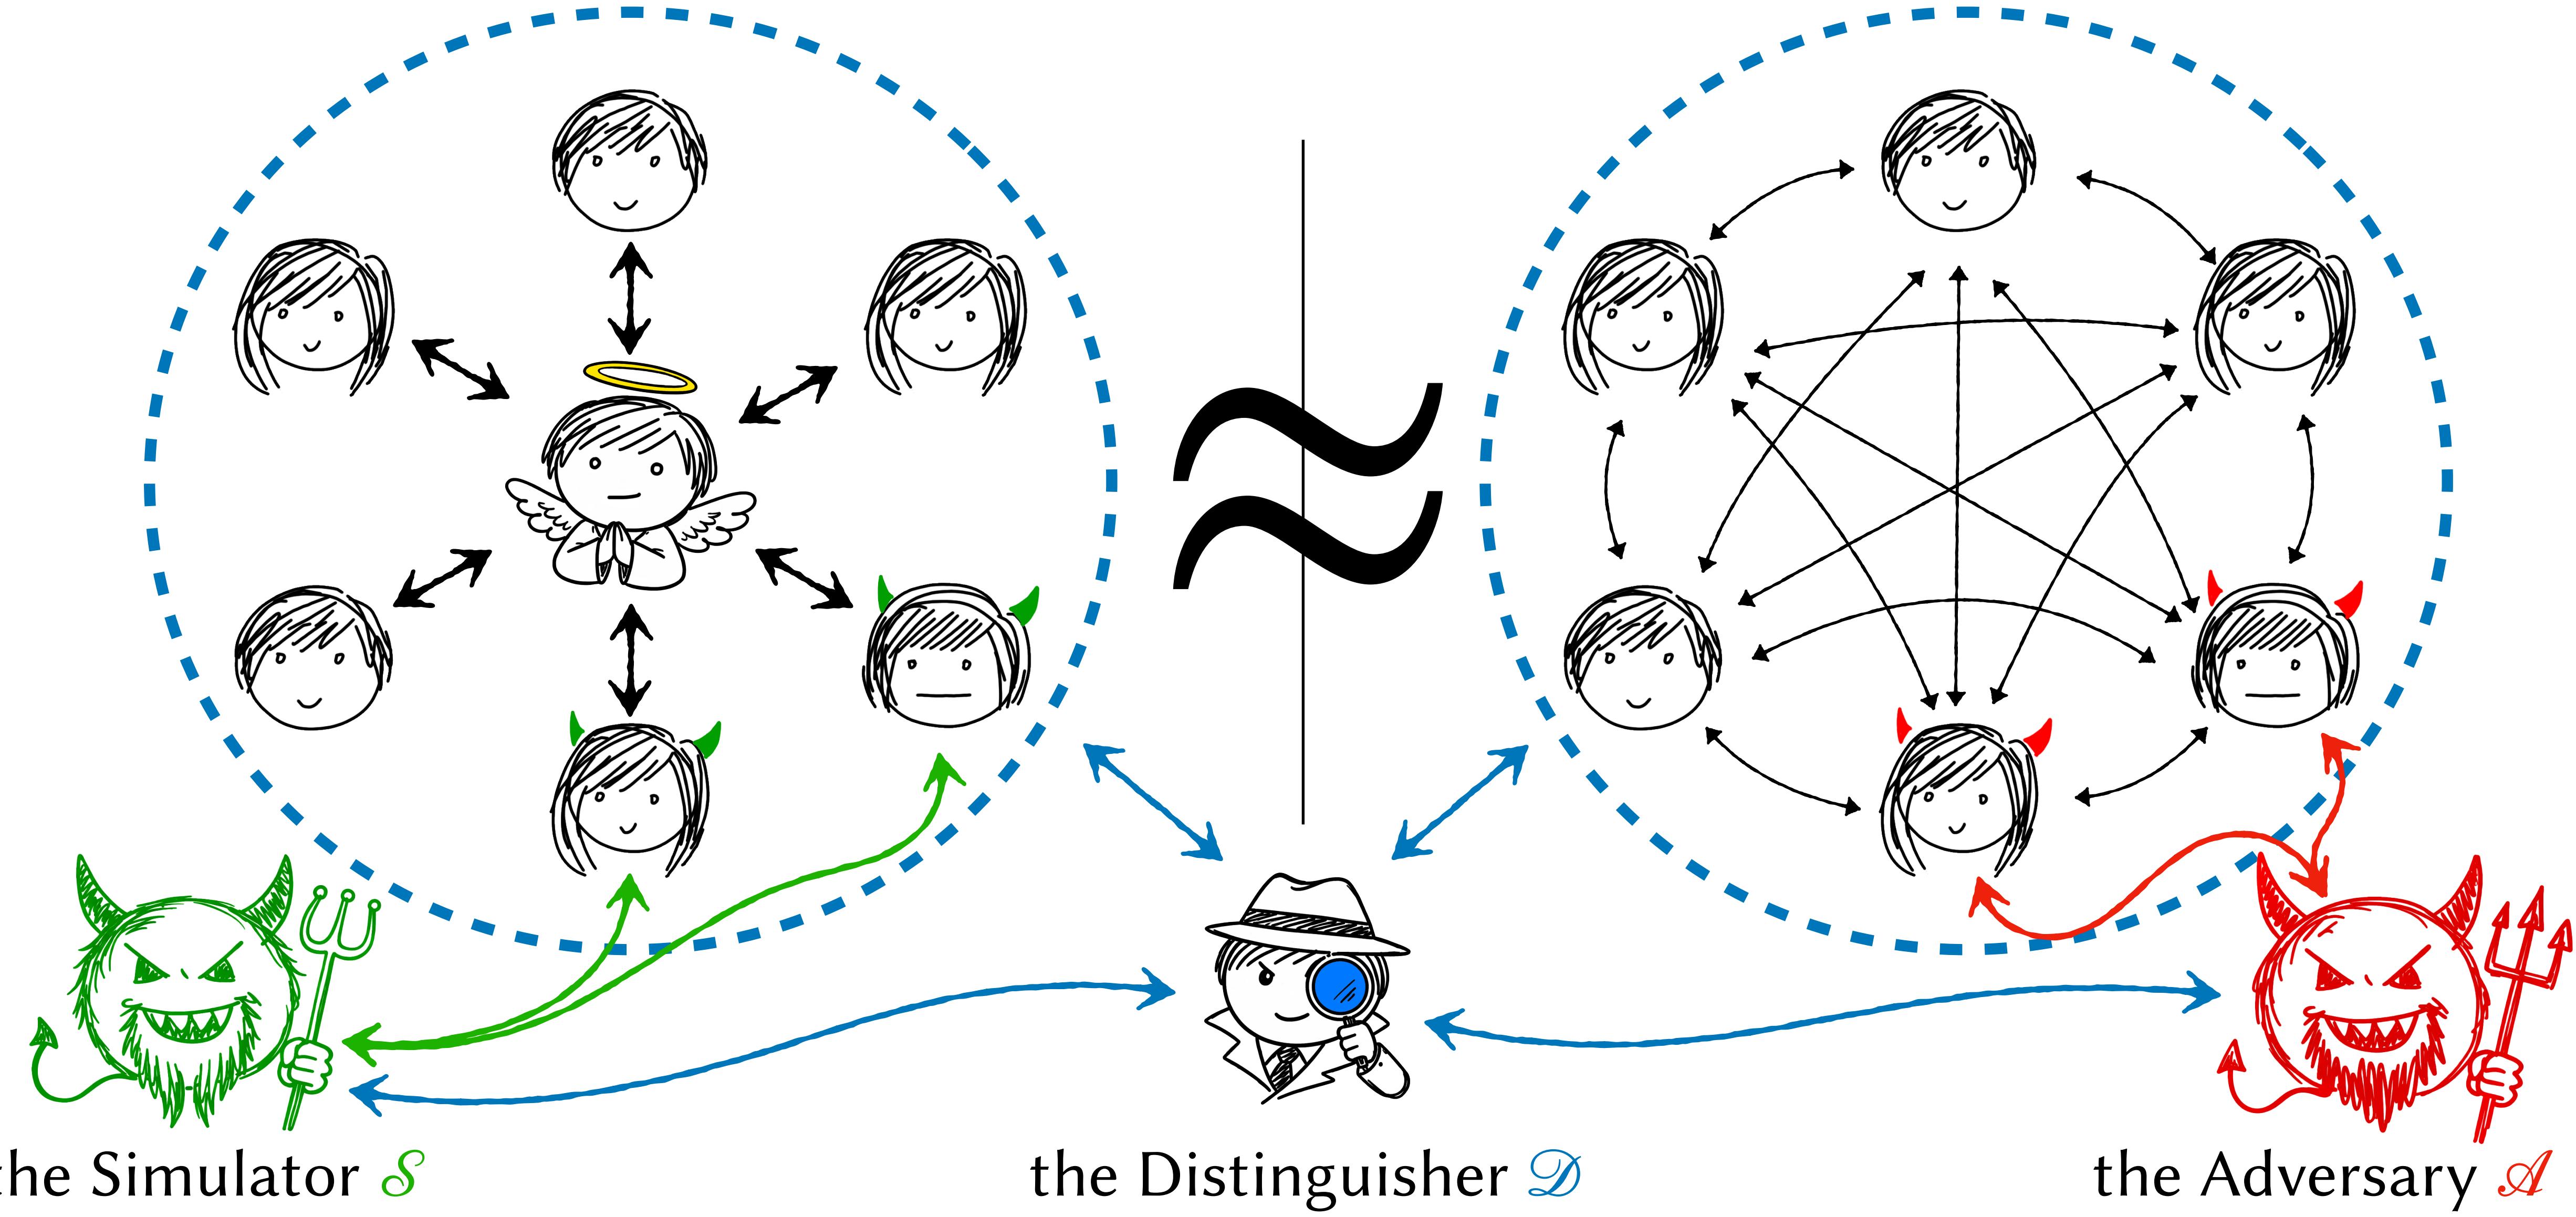
Sanity Check: Guaranteed Output Delivery?

the Simulator $\mathcal{S}$ the Distinguisher $\mathcal{D}$ the Adversary $\mathcal{A}$
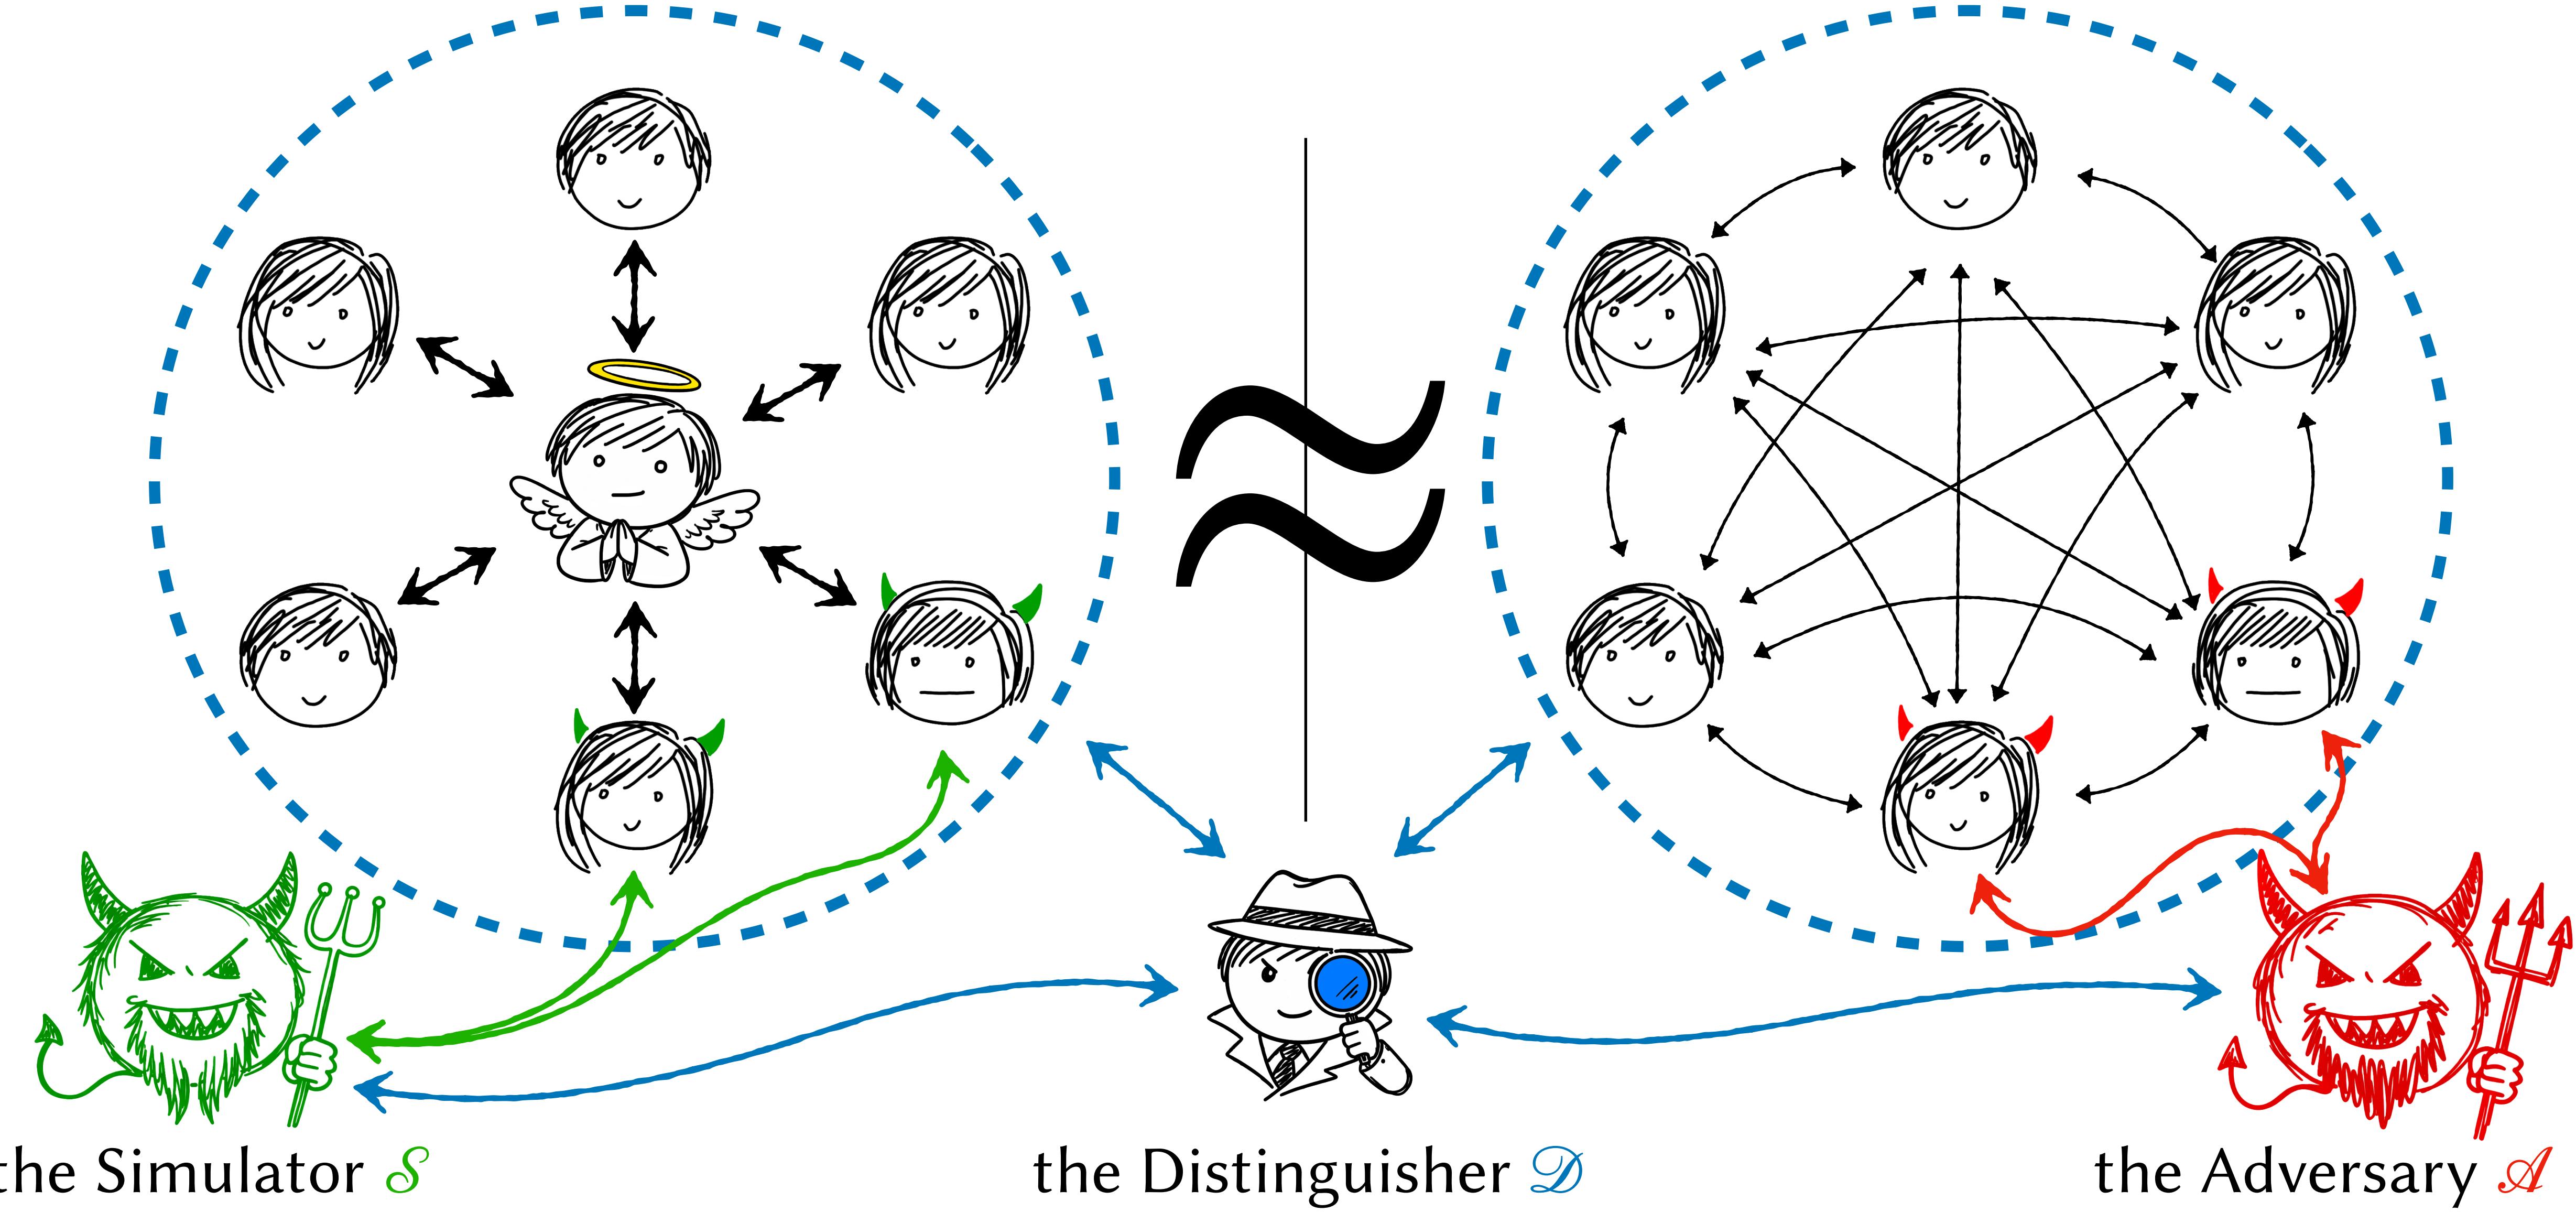
This is complicated. Why do we do it?

- This way of defining security is *extremely* general. We can capture any computational task and security behavior using the code of the functionality.

- While the *model* is complicated, the *guarantee* is simple to understand: Just imagine a trusted party performing the task that you desire.

- The functionality gives a complete picture of the security that the protocol achieves. Nothing is accidentally missed.

- Supports *composition*: use functionalities as parties in larger protocols. Reason about security in a modular and reusable fashion!

Specifying the Details

We've met the players, but in order to understand what we're achieving, we must know more about them.

- **Functionality:** what do we want to compute, and with what IO behavior? A functionality can also capture *vulnerabilities* by directly taking inputs from or leaking information to the simulator $\mathcal{S}$.

- **Adversarial Model:** what kinds of behavior and interaction with the system do we want to protect against? E.g. who can be corrupted?

- **Security Type:** how strong should our protection be? How much computing power does $\mathcal{D}$ have? How much better than random is the guesswork of $\mathcal{D}$ allowed to be?

- **Network Model:** Who is connected to who in the real world? Are those connections private? Do the parties have a shared clock? Is there a way to *broadcast* reliably?

Facets of Adversarial Models

Behavior:

- *Semi-honest* (a.k.a. *Passive*): corrupted parties follow the protocol honestly, but share their internal state with $\mathscr{A}$, who tries to learn more than is allowed.

- *Malicious* (a.k.a. *Active*): corrupted parties can deviate from the protocol instructions in arbitrary ways. $\mathscr{A}$ coordinates their actions.

Adversarial Power:

- *Unbounded* (i.e. "all-powerful"): $\mathscr{A}$ has unlimited computing power. Can break any cryptographic assumption. We can still achieve security using information theory!

- *Computationally Bounded*: $\mathscr{A}$ runs in Probabilistic Polynomial Time (PPT).

Behavior:

- *Semi-honest* (a.k.a. *Passive*): corrupted parties follow the protocol honestly, but share their internal state with $\mathcal{A}$, who tries to learn more than is allowed.

- *Malicious* (a.k.a. *Active*): corrupted parties can deviate from the protocol instructions in arbitrary ways. $\mathcal{A}$ coordinates their actions.

Adversarial Power:

- *Unbounded* (i.e. "all-powerful"): $\mathcal{A}$ has unlimited computing power. Can break any cryptographic assumption. We can still achieve security using information theory!

- *Computationally Bounded*: $\mathcal{A}$ runs in Probabilistic Polynomial Time (PPT).

Corruption Strategy:

- *Static*: corruptions are determined at the beginning of the experiment. Honest parties always stay honest.

- *Adaptive*: $\mathcal{A}$ can dynamically corrupt parties during the protocol (security is very hard to achieve in this setting).

Security Types

Perfect:

- $\mathcal{D}$ and $\mathcal{A}$ have unbounded computational power.

- The real and ideal experiments must be identically distributed from the perspective of $\mathcal{D}$.

- $\mathcal{D}$ must be able to do no better than a random guess.

Statistical:

- $\mathcal{D}$ and $\mathcal{A}$ have unbounded computational power.

- The real and ideal experiments must be *statistically indistinguishable*. (Their *statistical distance* must be *negligible* relative to the *security parameter*)

Perfect:

- $\mathcal{D}$ and $\mathcal{A}$ have unbounded computational power.

- The real and ideal experiments must be identically distributed from the perspective of $\mathcal{D}$.

- $\mathcal{D}$ must be able to do no better than a random guess.

Statistical:

- $\mathcal{D}$ and $\mathcal{A}$ have unbounded computational power.

- The real and ideal experiments must be *statistically indistinguishable.* (Their *statistical distance* must be *negligible* relative to the *security parameter*)

Computational:

- $\mathcal{D}$ and $\mathcal{A}$ are *efficient.* They run in PPT.

- We can make *cryptographic* assumptions. That is, we can assume certain computational problems can't be solved by $\mathcal{D}$ or $\mathcal{A}$.

- The real and ideal experiments must be *computationally indistinguishable.* This means that the outputs of $\mathcal{D}$ are *statistically close* when it interacts with the real and ideal experiments, even though the experiments themselves might have *statistically far* distributions.

We need to formalize all of this mathematically!

But first, let's look at what we will achieve in this class.

The "Fundamental Theorem" of Multiparty Computation

Every computable function f can be securely computed
*(even if $n - 1$ parties are corrupted)*

The "Fundamental Bound" of Multiparty Computation

If a majority of parties are maliciously corrupted, there exist
functionalities $\mathcal{F}$ that cannot be realized regardless
of what assumptions are made about the adversary's power.

Assuming an Honest Majority

Let n be the number of parties and let t be the number of corruptions

Behavior	Corruption Bound	Network Assumptions	Crypto Assumptions	Security
Semi-Honest	$t < n/2$	private, authenticated point-to-point channels	—	Perfect
Malicious	$t < n/3$	private, authenticated point-to-point channels	—	Perfect
Malicious	$t < n/2$	private, authenticated p2p channels + broadcast	—	Statistical
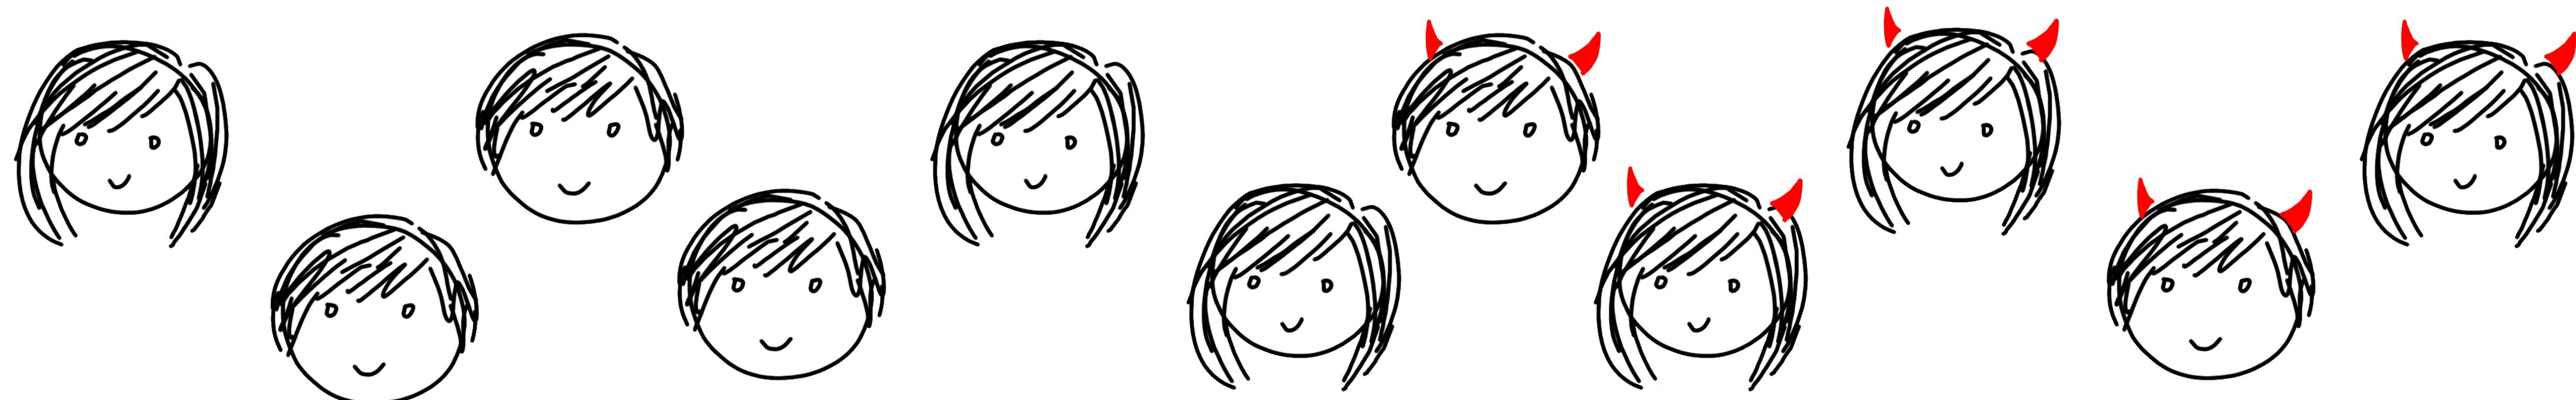
Assuming a *Dis*honest Majority

Let n be the number of parties and let t be the number of corruptions

Behavior	Corruption Bound	Network Assumptions	Crypto Assumptions	Security
Semi-Honest	$t < n$	authenticated point-to-point channels	Oblivious Transfer	Computational
Malicious	$t < n$	authenticated p2p channels + broadcast	Oblivious Transfer	Computational (*no fairness or GOD*)
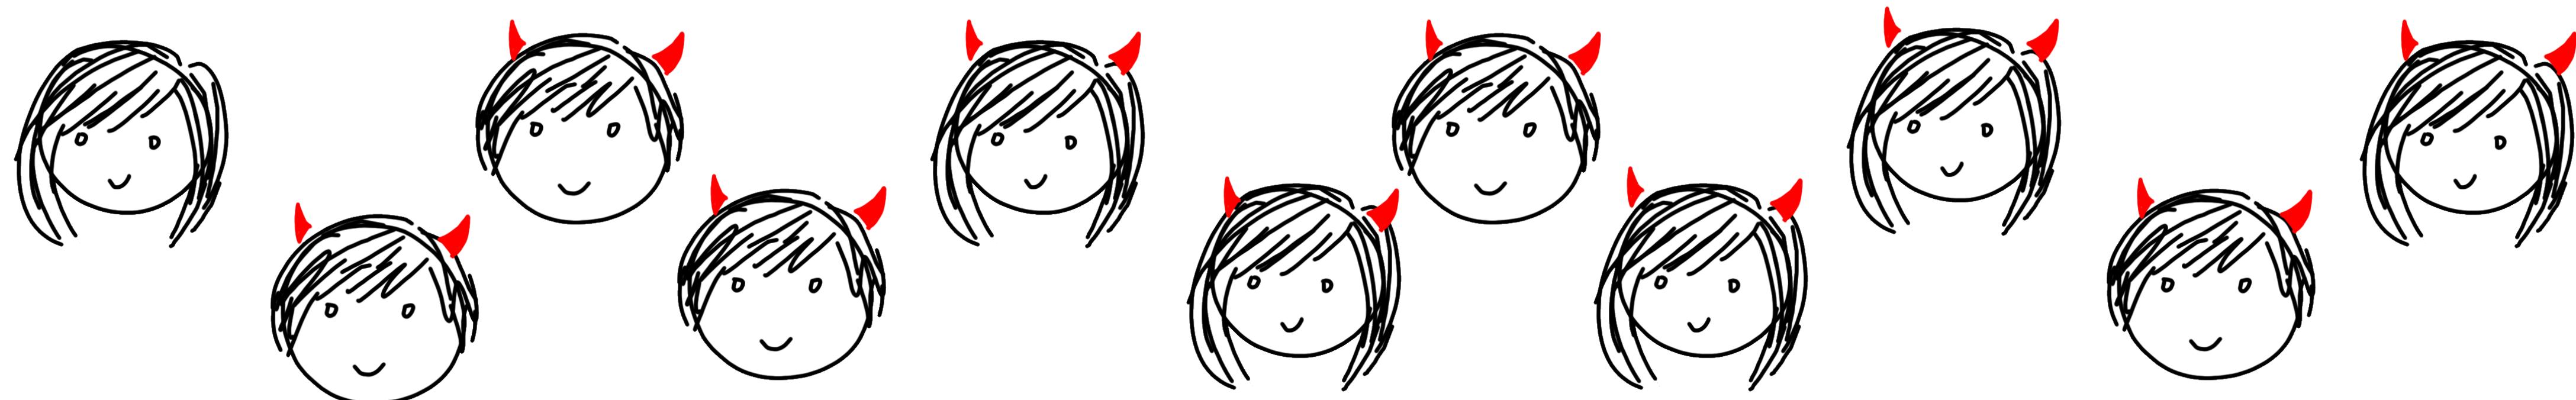
Syllabus (tentative):

A taxonomy of adversaries; a variety of techniques
(now the taxonomy should be clearer than it was before)

	Part 1: *Information-theoretic* techniques. Adversaries with unbounded power	**Part 2:** *Cryptographic* techniques. Adversaries with bounded power
Semi-honest Adversaries: *follow* the rules of the protocol	Secret Sharing BGW protocol for an honest majority	Oblivious Transfer GMW protocol for a dishonest majority Yao's protocol for two parties Fully Homomorphic Encryption
Malicious Adversaries: *break* the rules of the protocol	Verifiable Secret Sharing BGW protocol for honest supermajority	Coin Tossing Zero-Knowledge Proofs GMW Compiler Byzantine Agreement + Broadcast

Overarching Questions:
How do we characterize unknown adversaries? How do we formalize intuitive security notions?
What kinds computation can we perform securely in each setting?

CS4501 Cryptographic Protocols
Lecture 2: Adversaries and Simulation

https://jackdoerner.net/teaching/#2026/Spring/CS4501