

# On the trade-off between expressivity & privacy in graph representation learning

Patrick Indri\*, Tamara Drucks\*, Thomas Gärtner



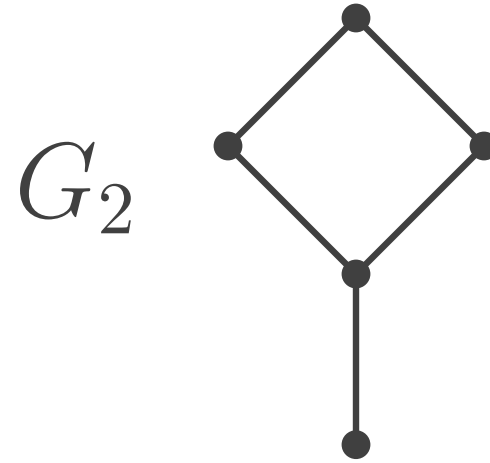
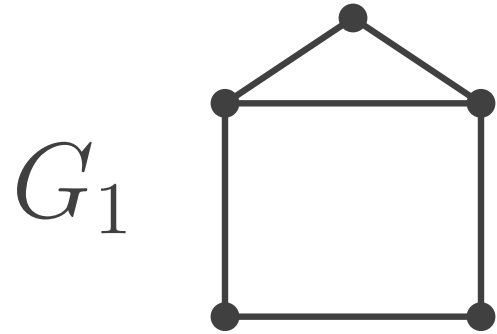
TECHNISCHE  
UNIVERSITÄT  
WIEN



**ICLR**

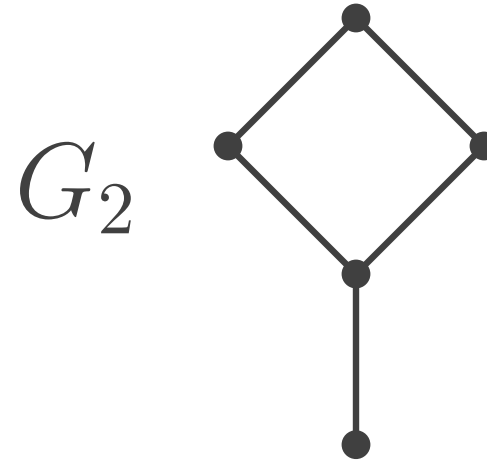
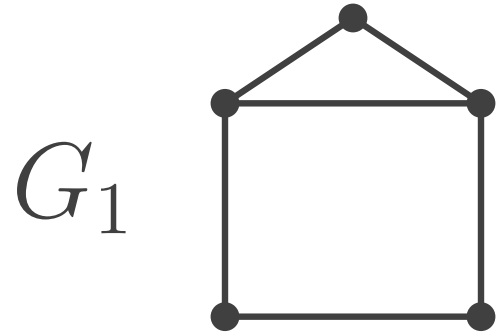
# What's the deal

In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$



# What's the deal

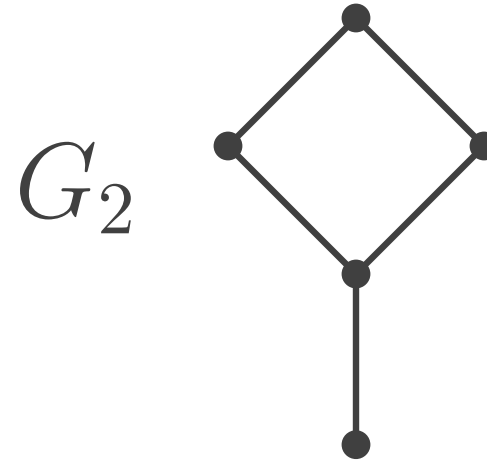
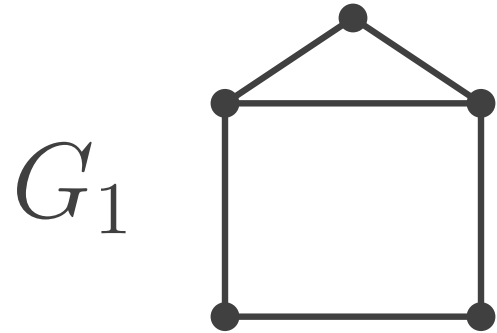
In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$



But what's *good*? Depends on who you ask

# What's the deal

In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$

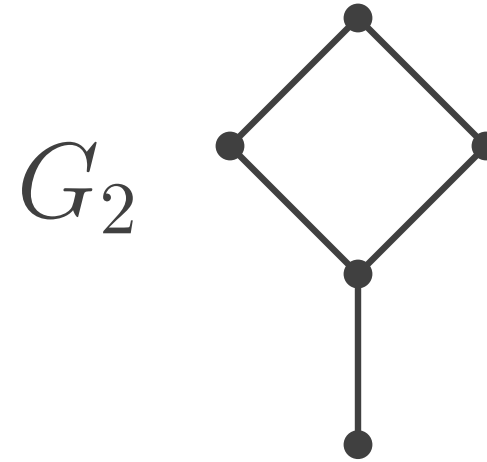
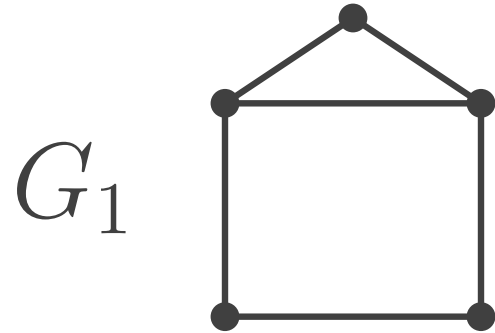


But what's *good*? Depends on who you ask

Let's ask an **expressivity** person!

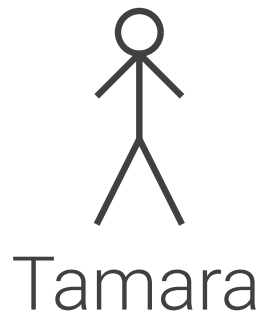
# What's the deal

In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$



But what's *good*? Depends on who you ask

Let's ask an **expressivity** person!

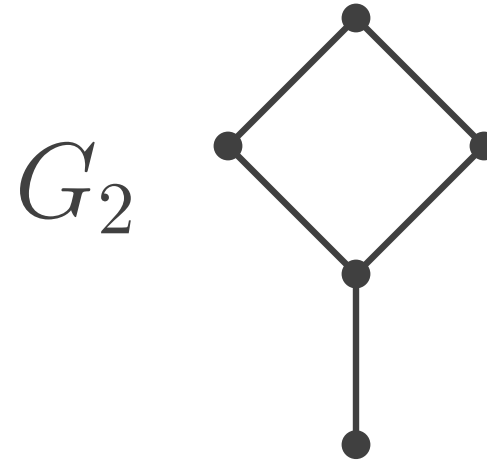
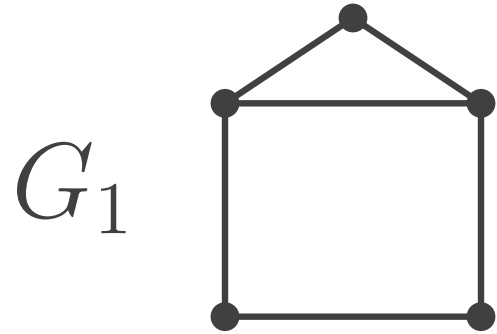


$$\varphi(G_1) = \varphi(G_2) \iff G_1 \simeq G_2$$

Same representations iff the graphs are isomorphic

# What's the deal

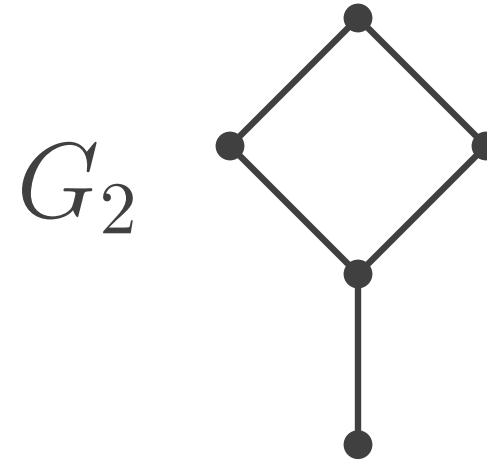
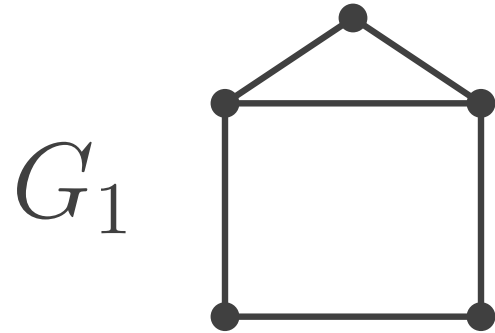
In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$



But what's *good*? Depends on who you ask

# What's the deal

In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$

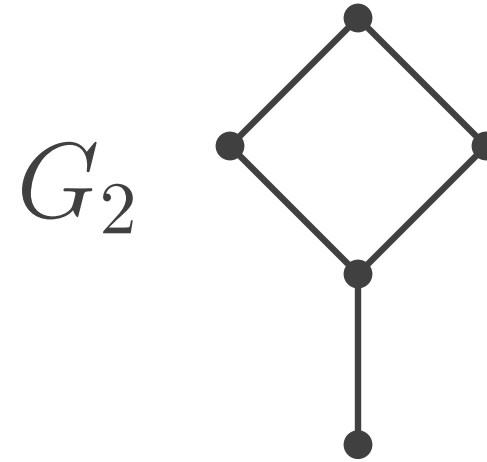
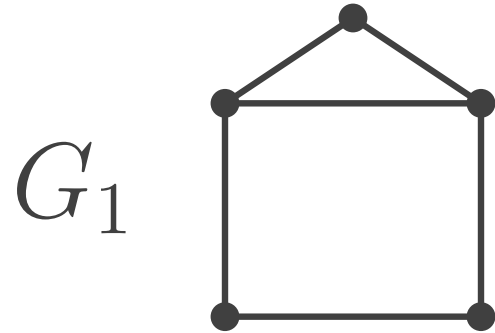


But what's *good*? Depends on who you ask

Let's ask a **privacy** person!

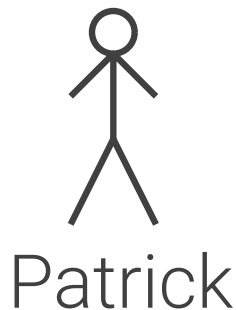
# What's the deal

In graph ML we want *good* representations  $\varphi(G)$  of a graph  $G$



But what's *good*? Depends on who you ask

Let's ask a **privacy** person!

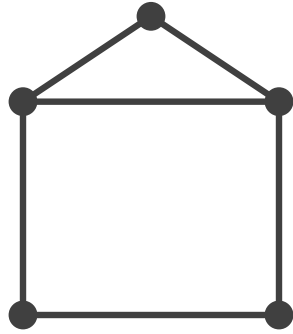


$$d_{\text{edge}}(G_1, G_2) = 1 \implies \varphi(G_1) \approx \varphi(G_2)$$

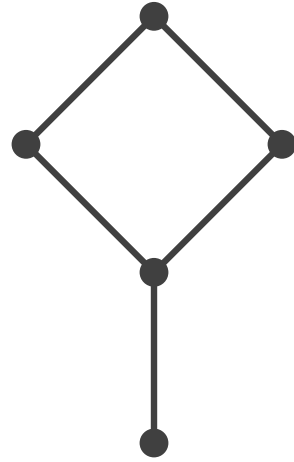
Similar representations for edge-adjacent graphs

# Why not both?

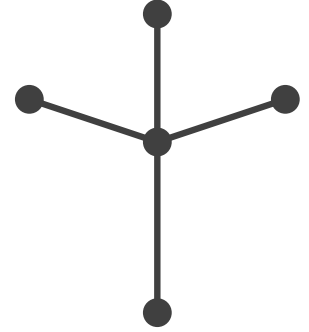
$G_1$



$G_2$

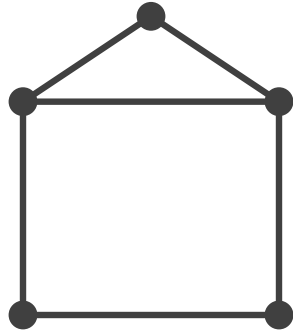


$G_3$

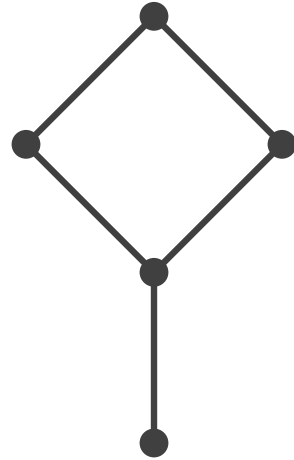


# Why not both?

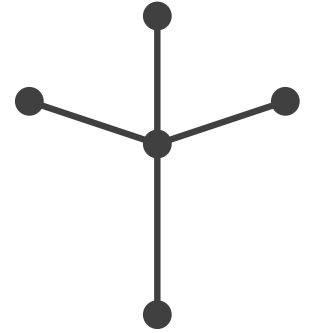
$G_1$



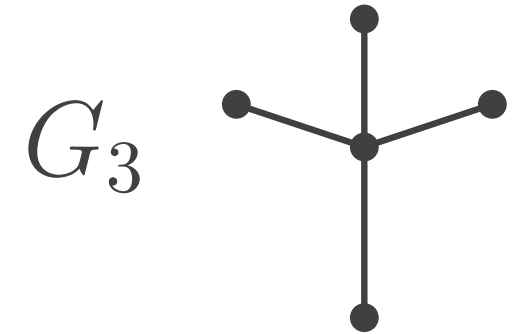
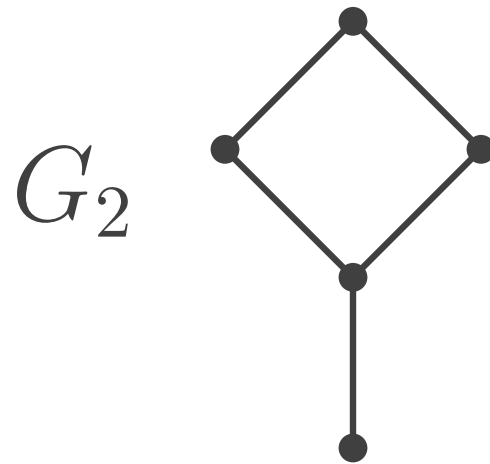
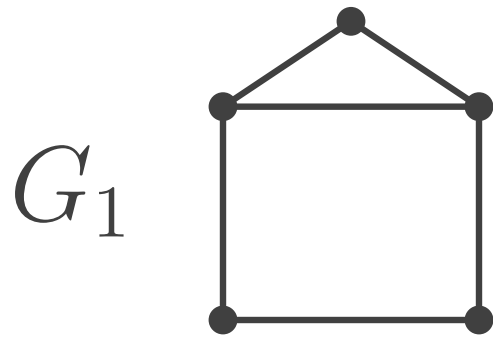
$G_2$



$G_3$

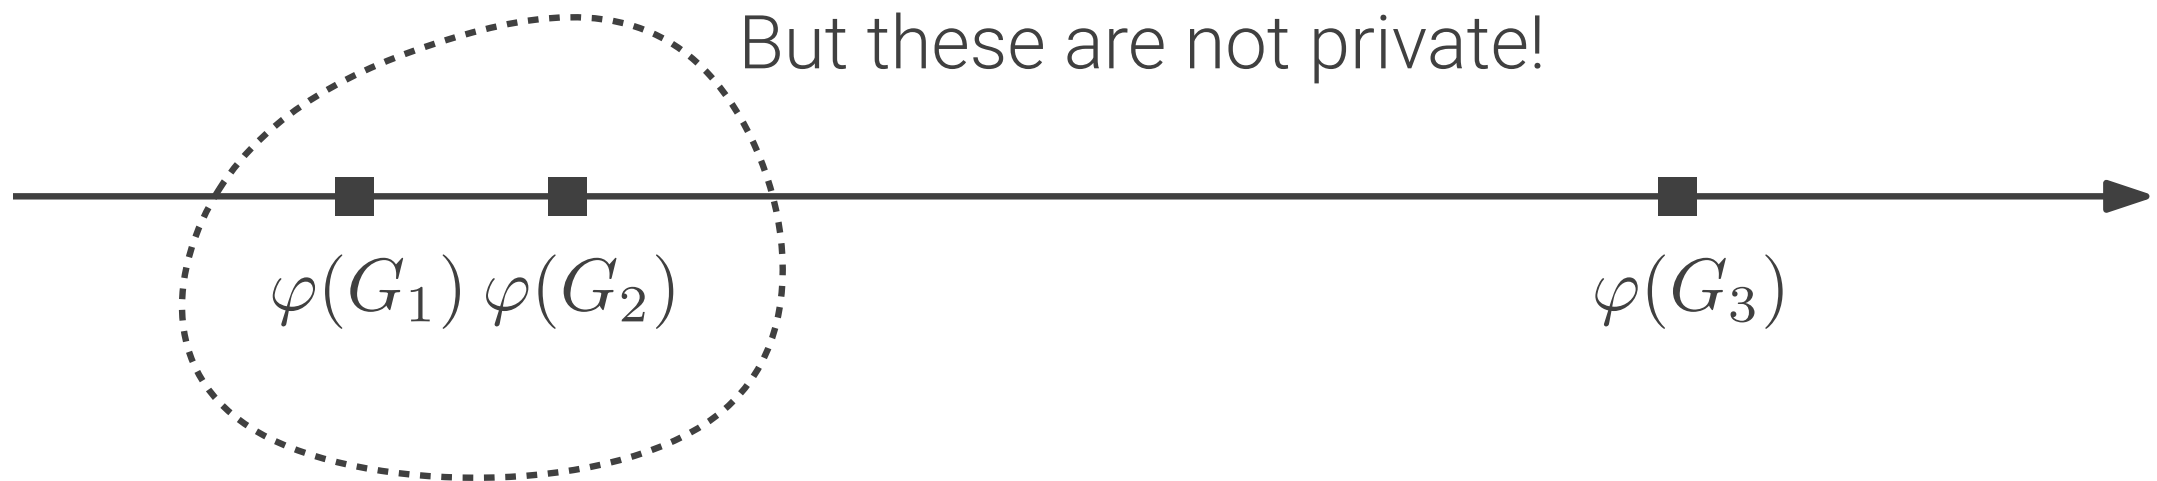
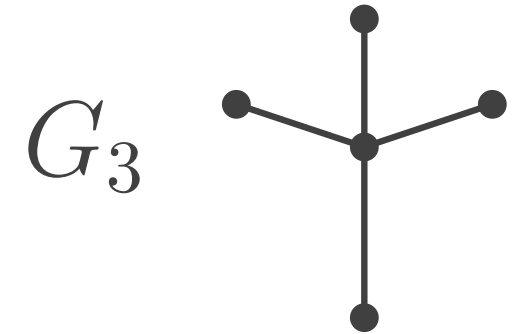
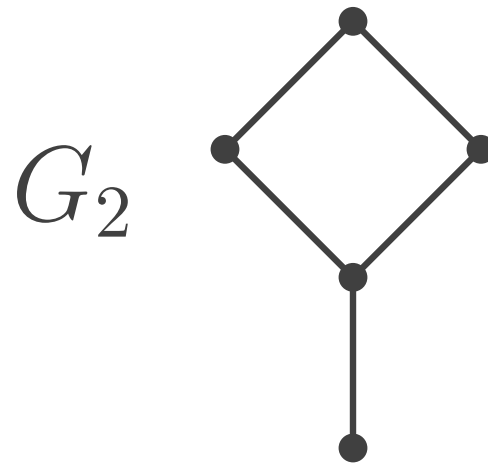
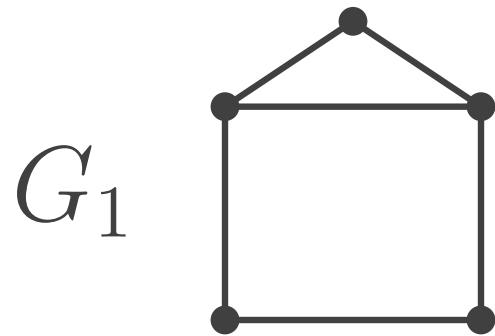


# Why not both?



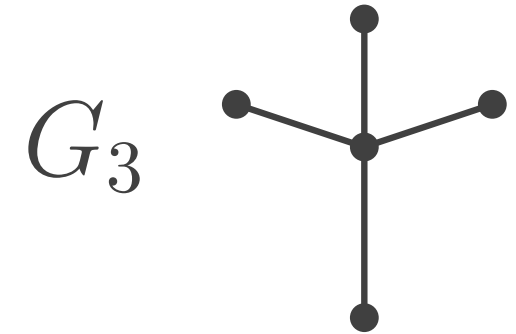
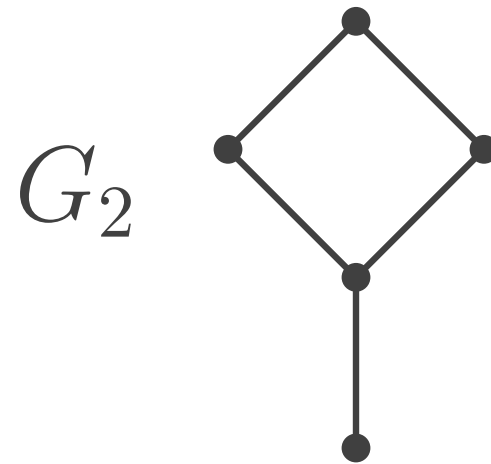
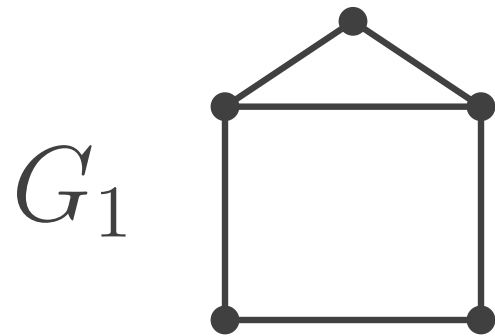
The **expressivity** person wants:  $\varphi(G_1) = \varphi(G_2) \iff G_1 \simeq G_2$

# Why not both?



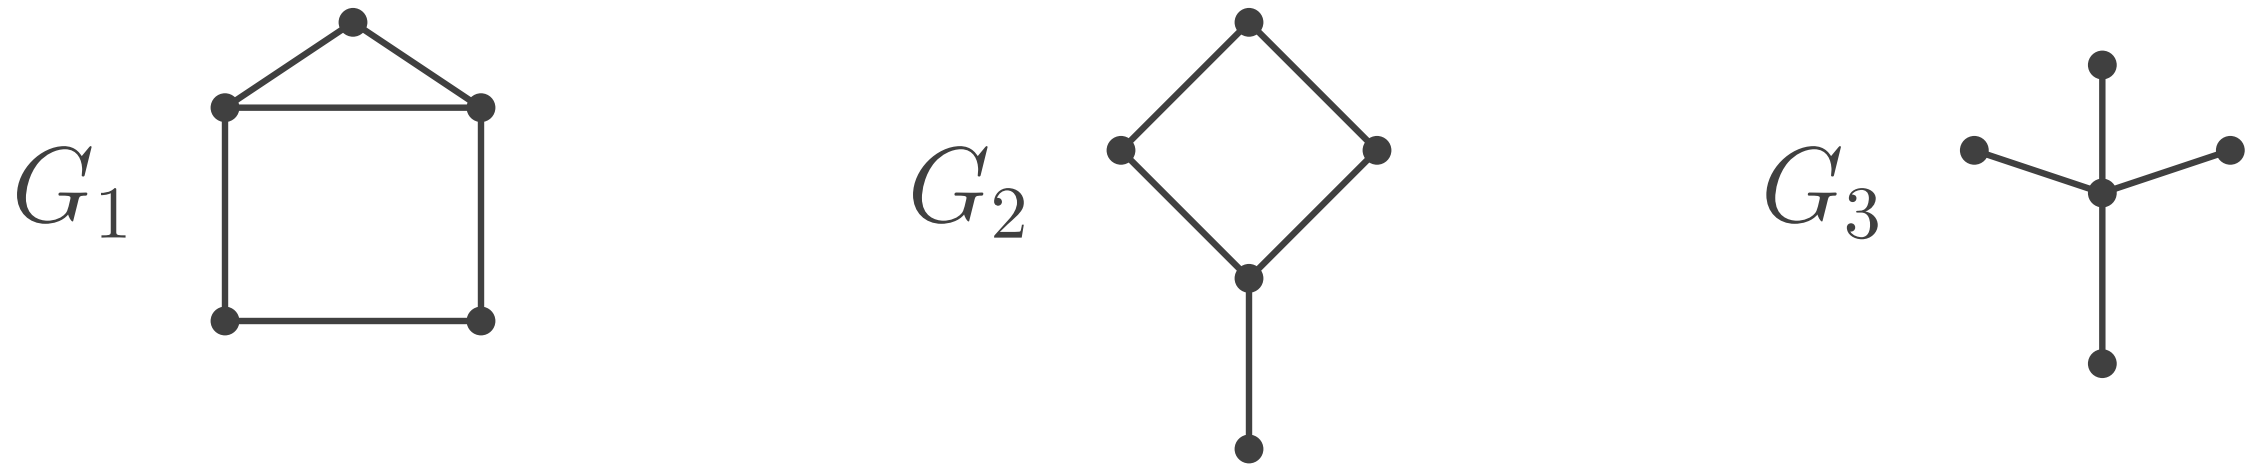
The **expressivity** person wants:  $\varphi(G_1) = \varphi(G_2) \iff G_1 \simeq G_2$

# Why not both?



The **privacy** person wants:  $d_{\text{edge}}(G_1, G_2) = 1 \implies \varphi(G_1) \approx \varphi(G_2)$

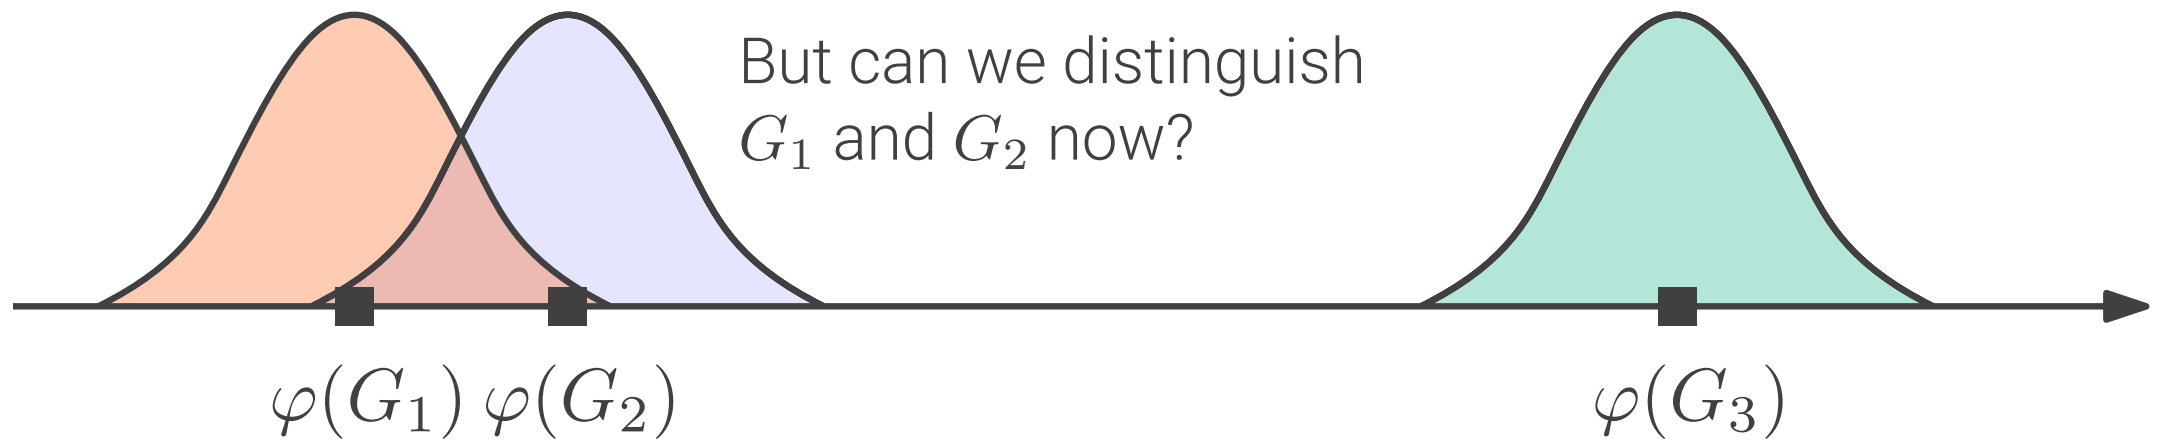
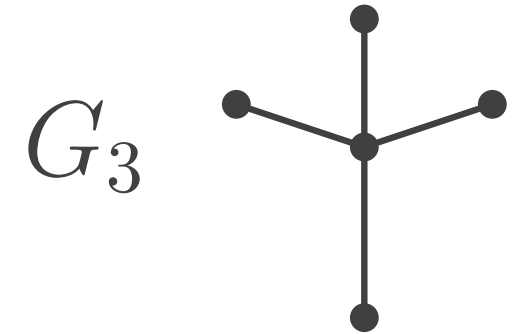
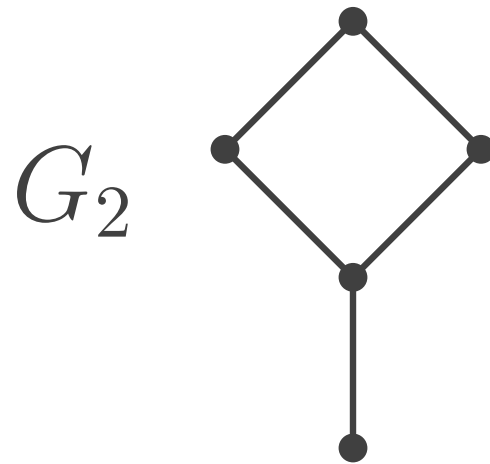
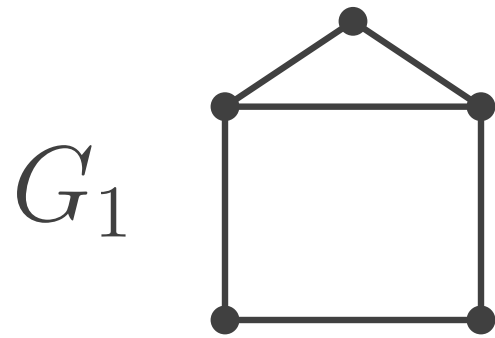
# Why not both?



let's add noise to the embeddings!

The **privacy** person wants:  $d_{\text{edge}}(G_1, G_2) = 1 \implies \varphi(G_1) \approx \varphi(G_2)$

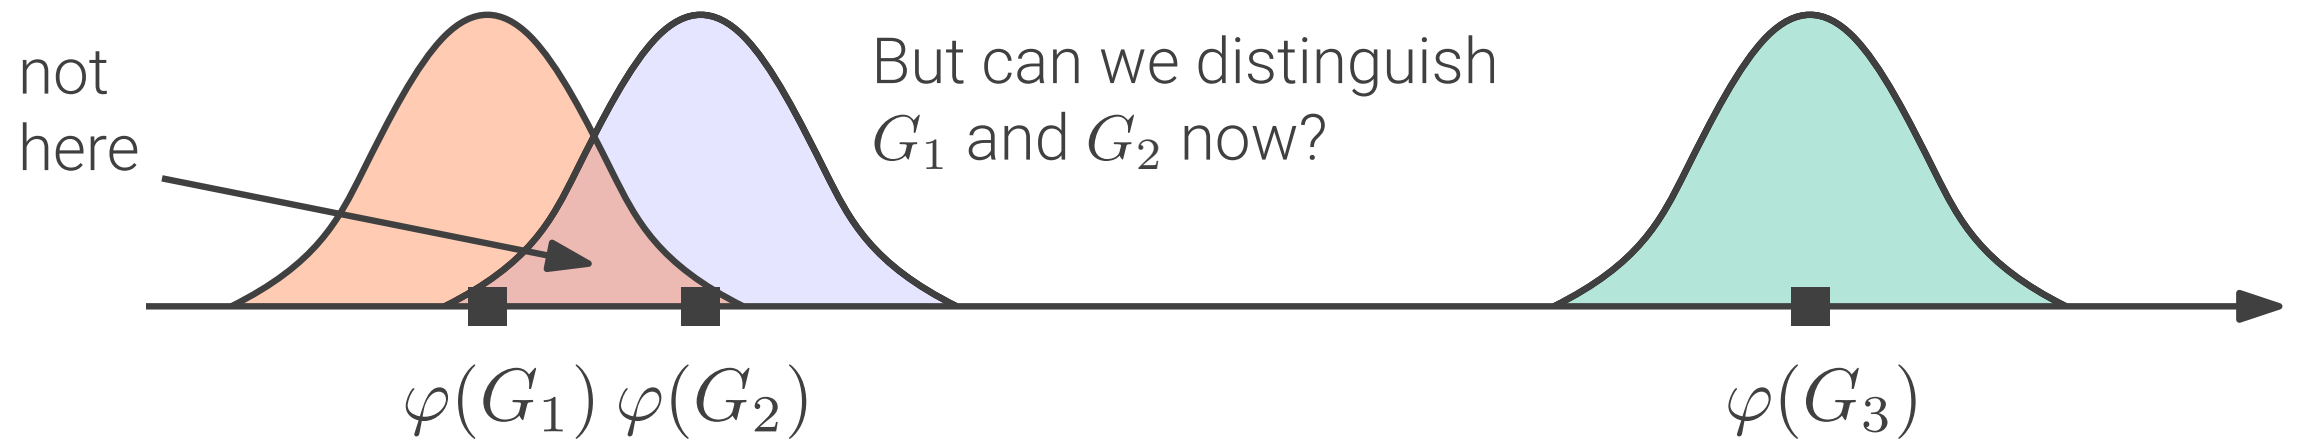
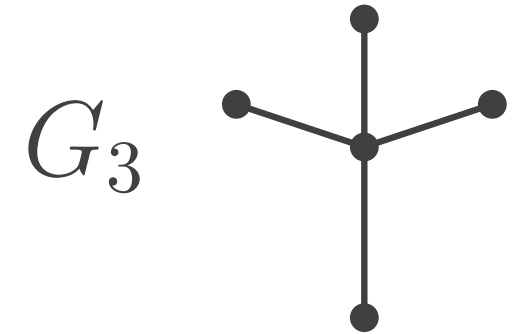
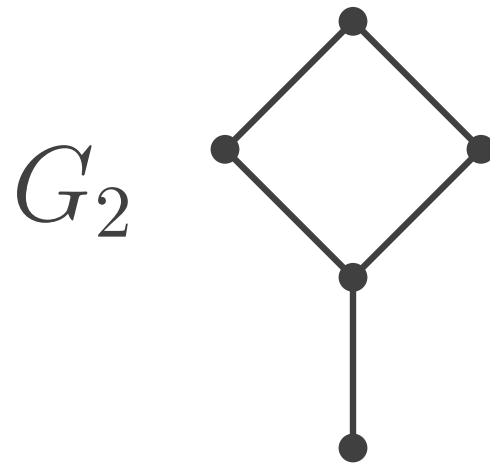
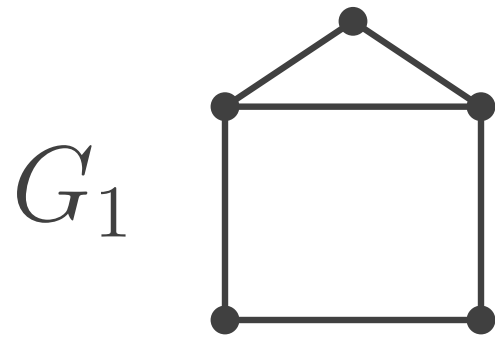
# Why not both?



let's add noise to the embeddings!

The **privacy** person wants:  $d_{\text{edge}}(G_1, G_2) = 1 \implies \varphi(G_1) \approx \varphi(G_2)$

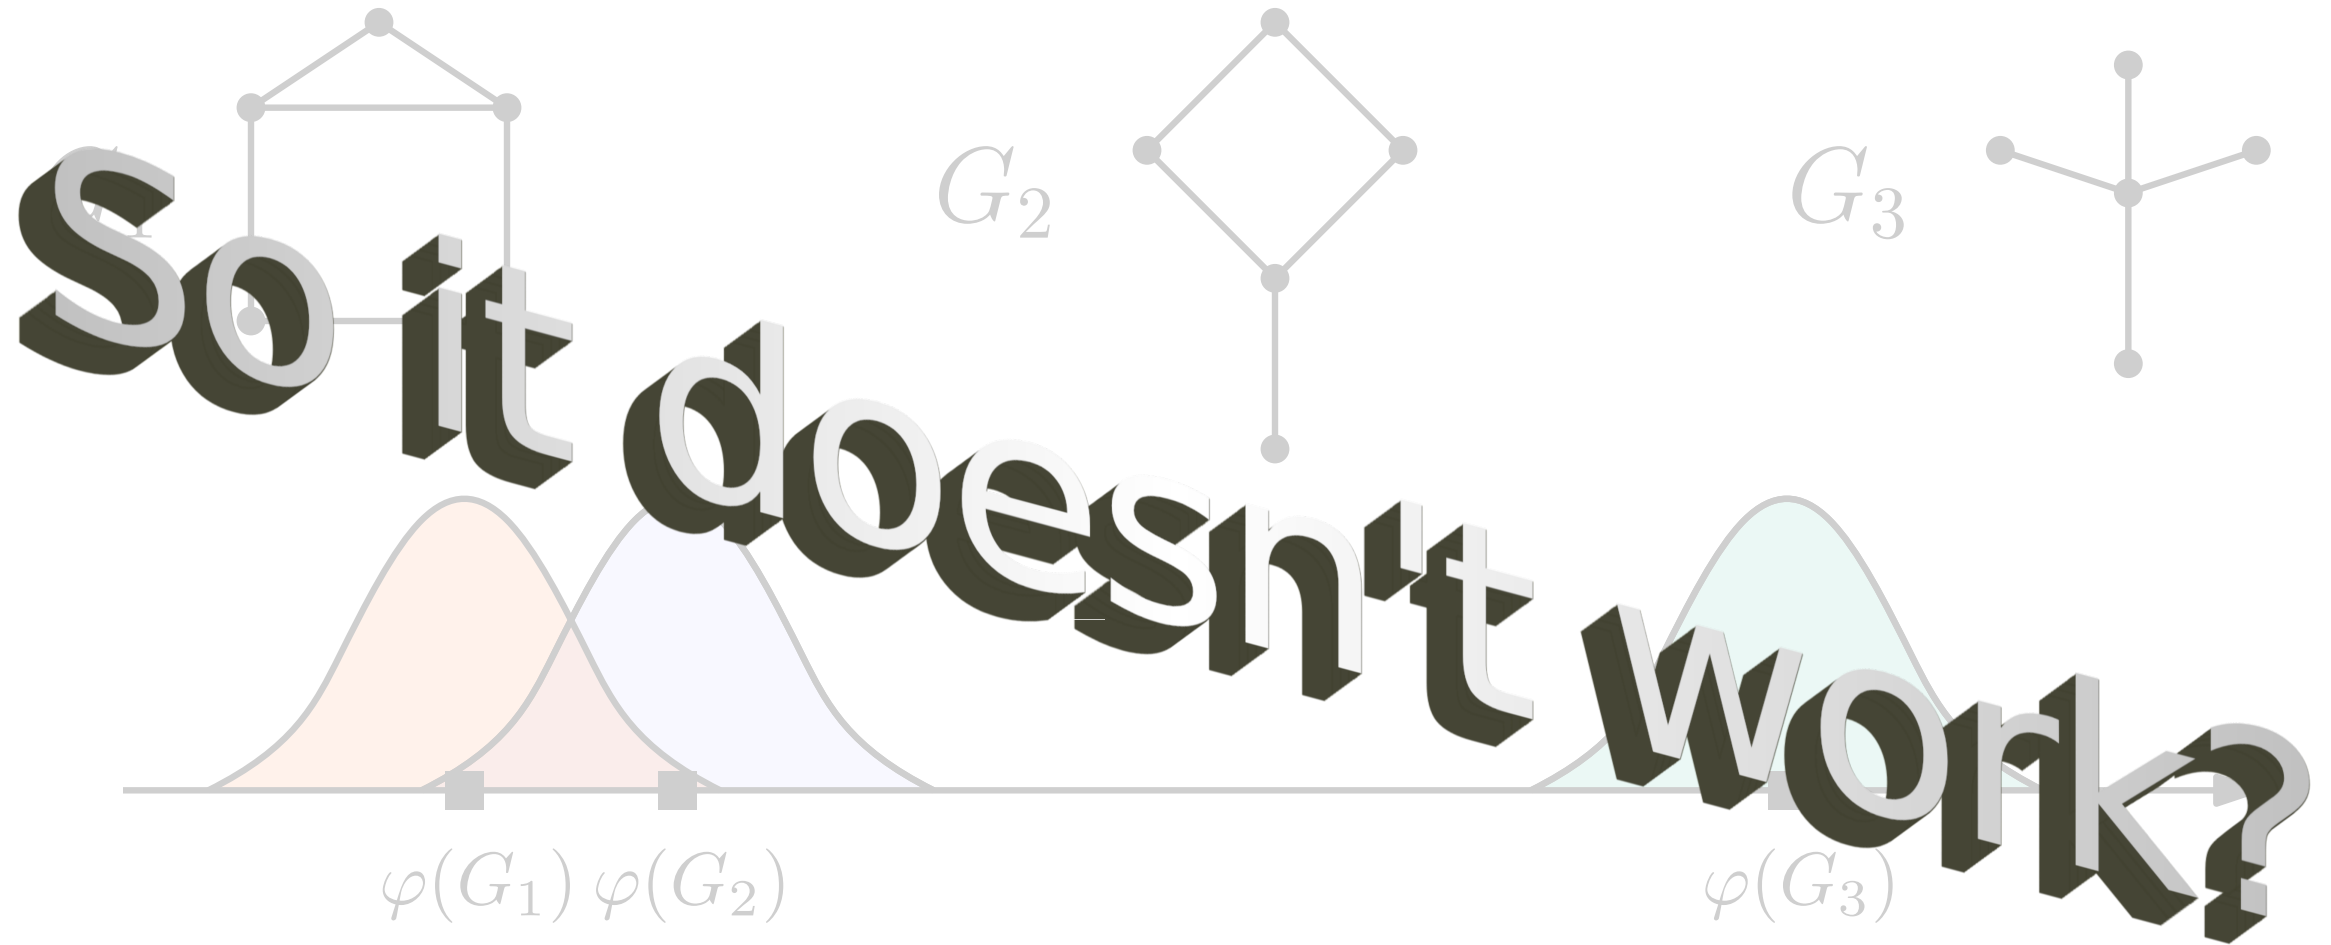
# Why not both?



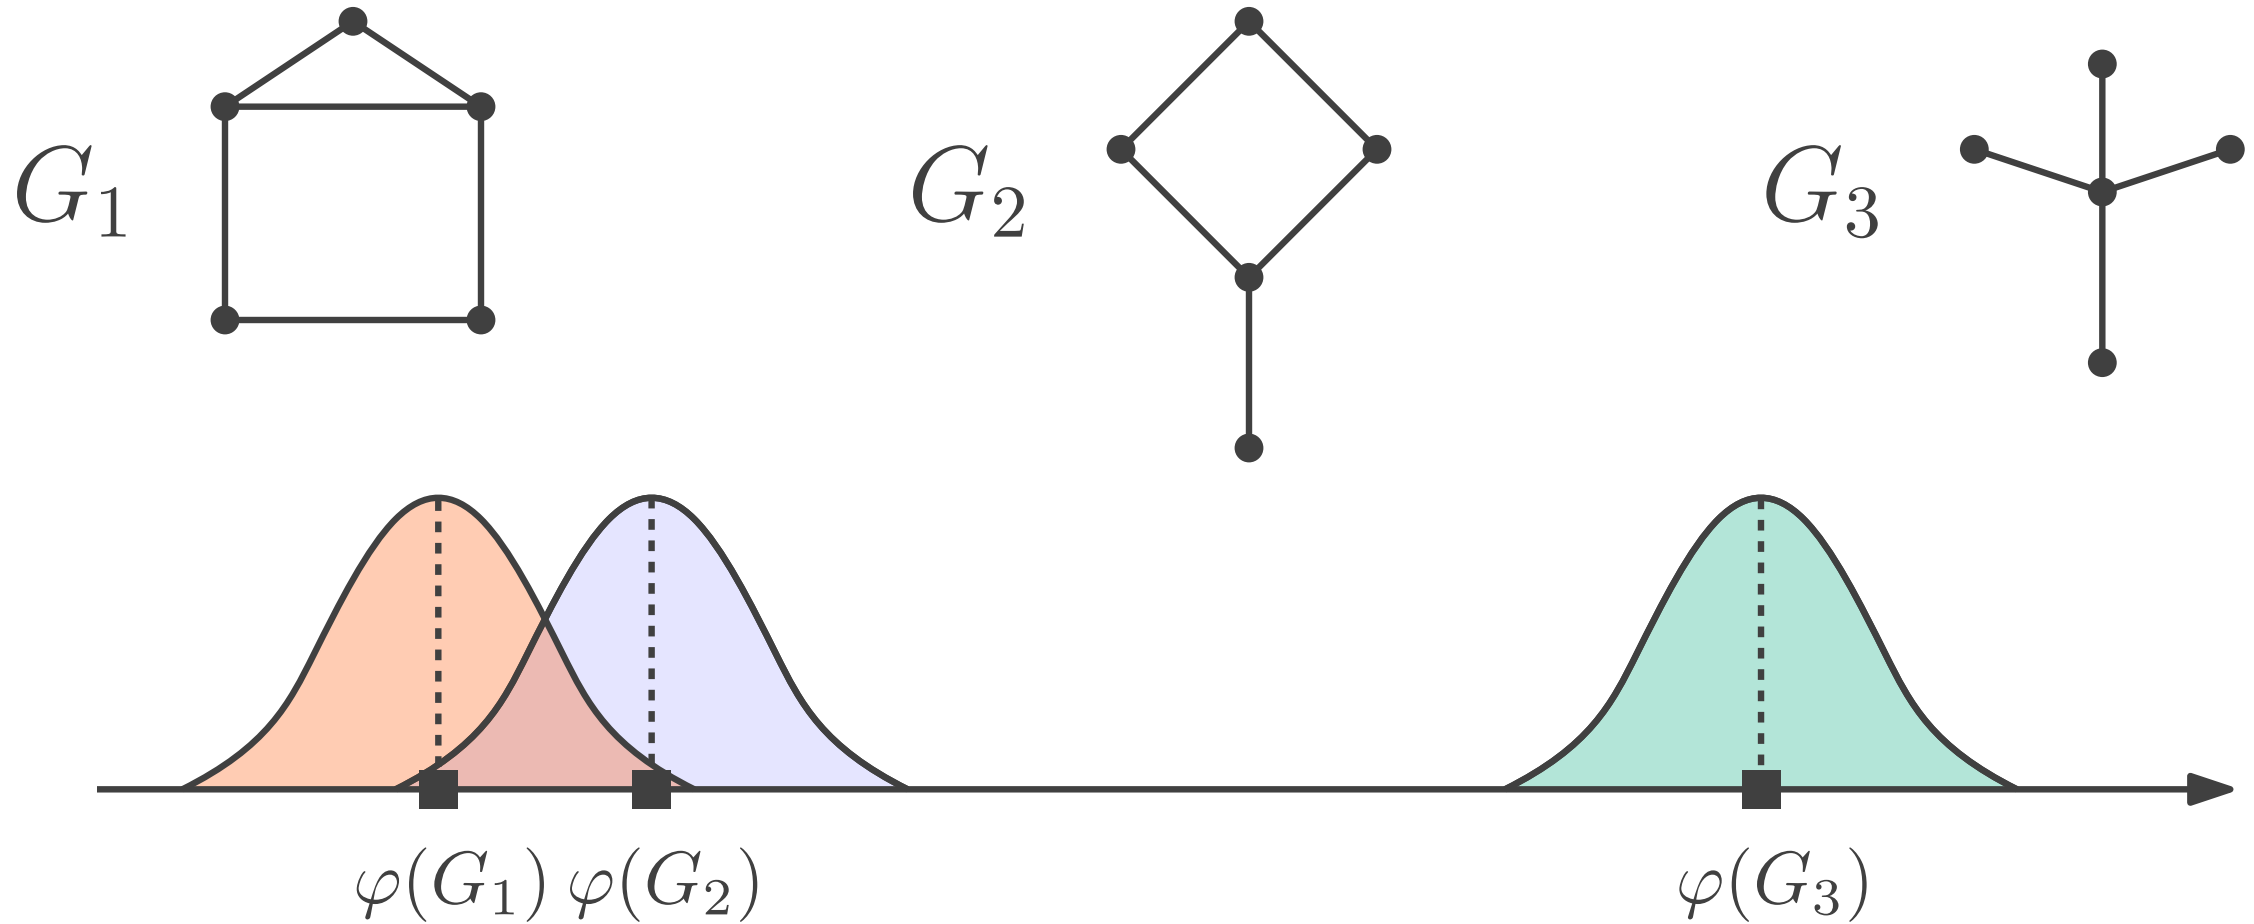
let's add noise to the embeddings!

The **privacy** person wants:  $d_{\text{edge}}(G_1, G_2) = 1 \implies \varphi(G_1) \approx \varphi(G_2)$

Why not both?

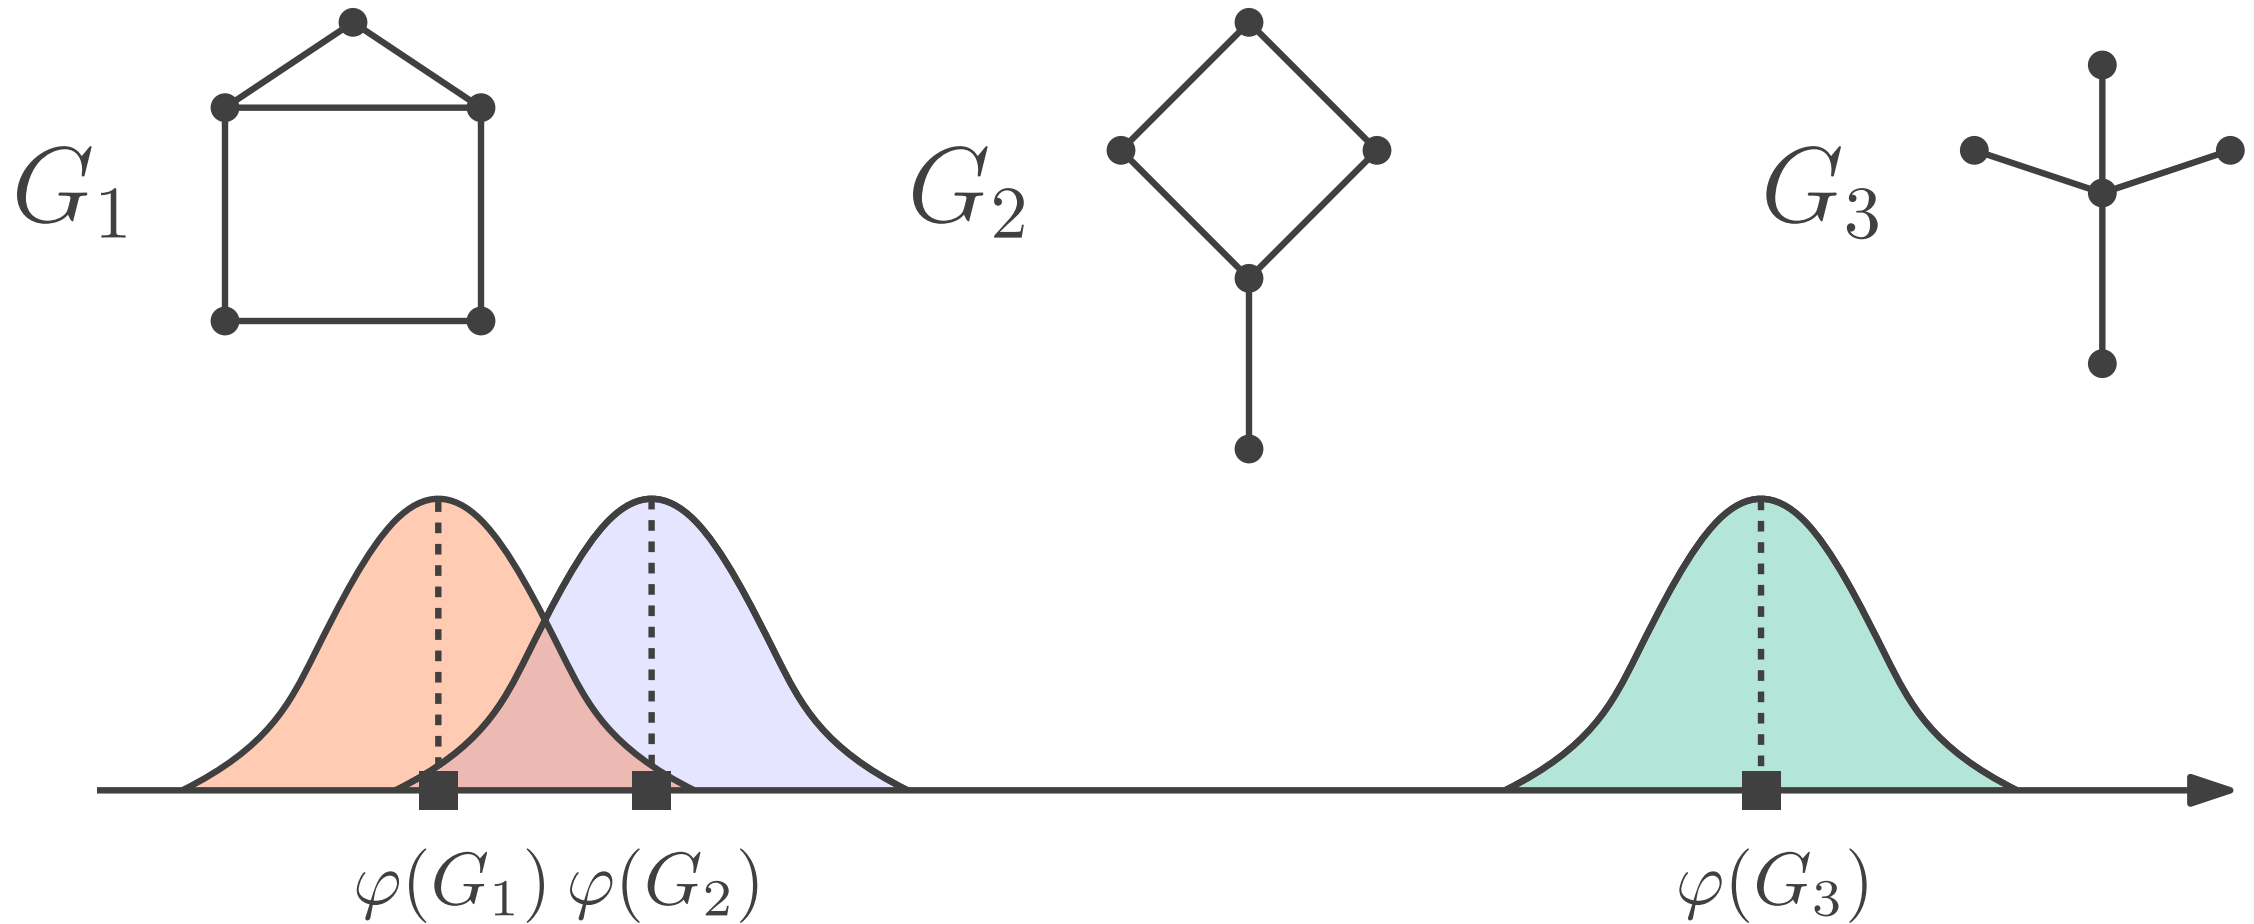


# We can make both work! Sort of



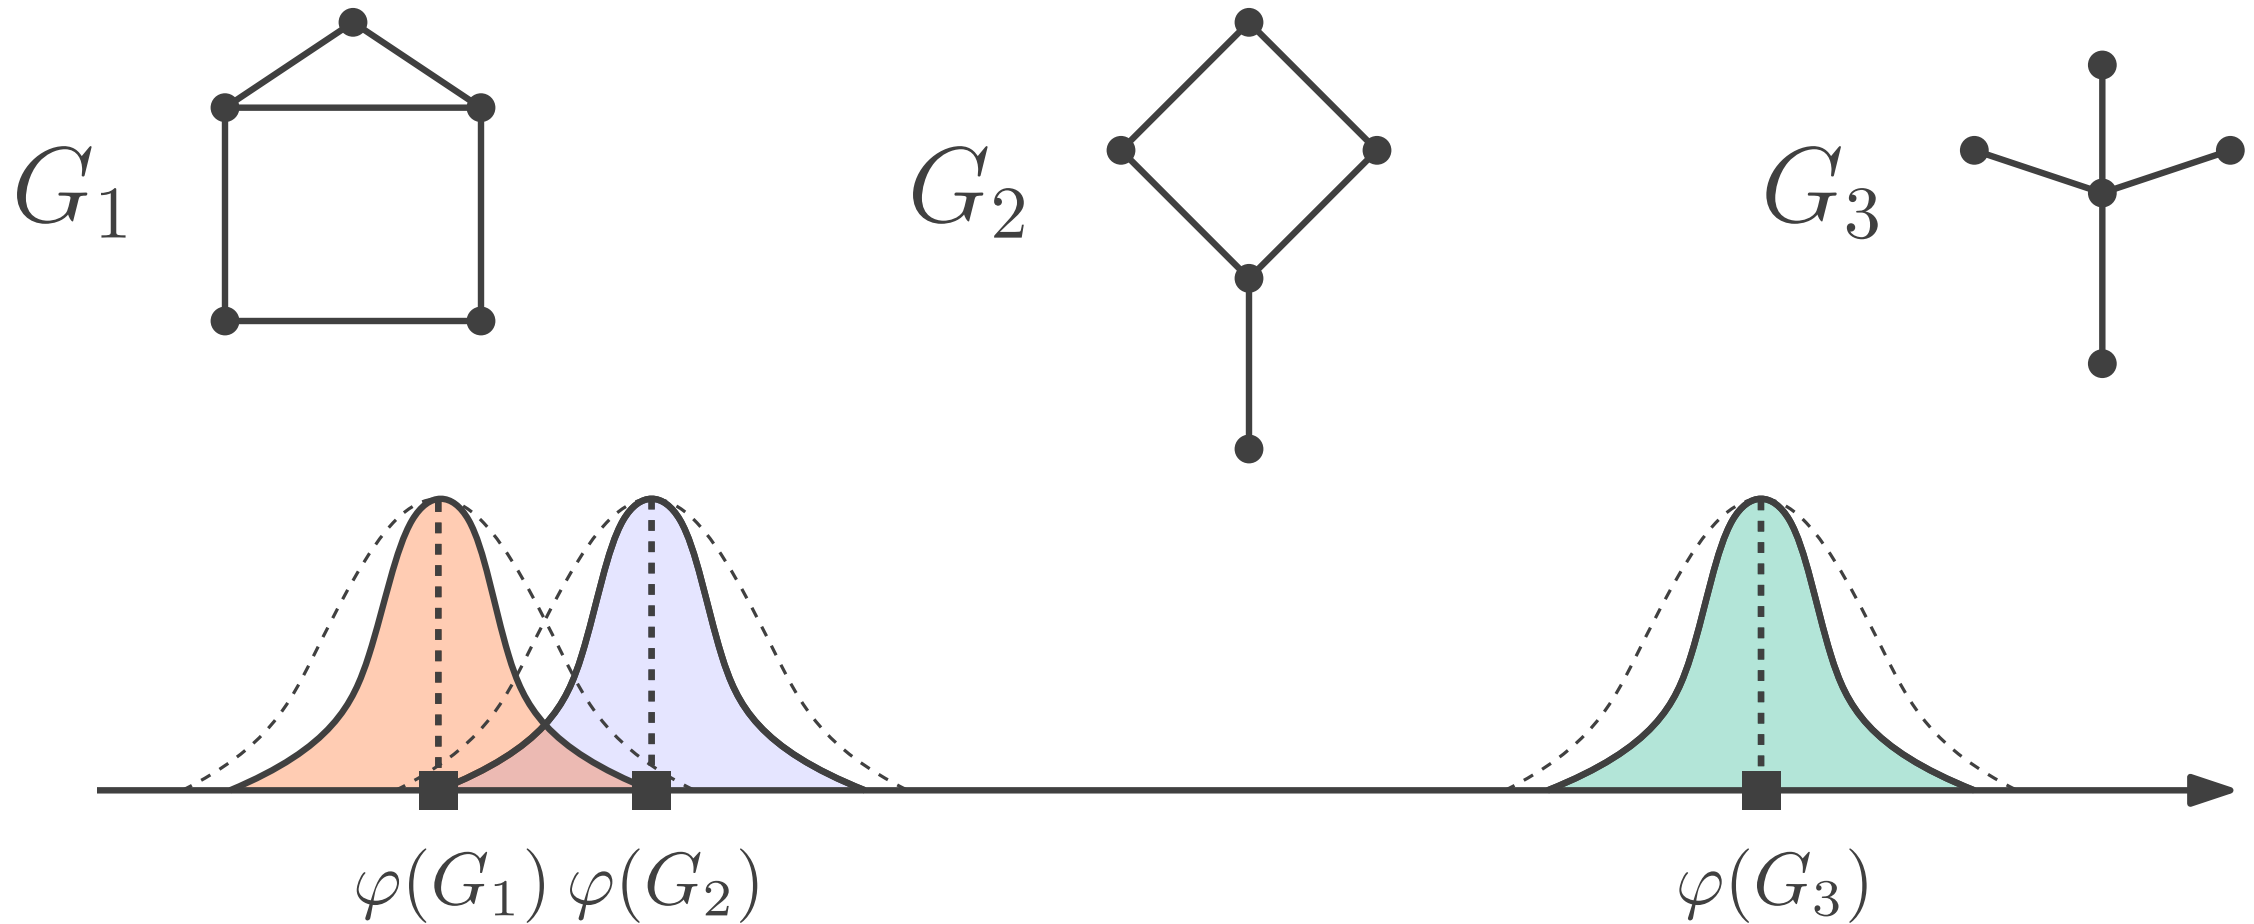
The **expected value** is still what the expressivity person wants

# We can make both work! Sort of



The **expected value** is still what the expressivity person wants  
And with a good **noise calibration** we can still be private

# We can make both work! Sort of



The **expected value** is still what the expressivity person wants  
And with a good **noise calibration** we can still be private

# What we do

We consider expressivity **in expectation** and a **local** notion of differential privacy.

# What we do

We consider expressivity **in expectation** and a **local** notion of differential privacy.

**Result 1.** We obtain graph embeddings that are **expressive** in expectation and differentially **private**.

# What we do

We consider expressivity **in expectation** and a **local** notion of differential privacy.

**Result 1.** We obtain graph embeddings that are **expressive** in expectation and differentially **private**.

determined by  
*pattern class*



# What we do

We consider expressivity **in expectation** and a **local** notion of differential privacy.

**Result 1.** We obtain graph embeddings that are **expressive** in expectation and differentially **private**.

determined by  
*pattern class*



**Result 2.** **More expressive** pattern classes often require **more noise** to be added to guarantee privacy.

# How?

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

Pattern graphs  $(F_1, \dots, F_n)$

$$\varphi(G) = \begin{array}{c} F_1 \quad F_2 \quad F_3 \quad \dots \quad F_n \\ \hline \begin{array}{|c|c|c|c|c|} \hline 3 & 27 & 12 & & 9 \\ \hline \end{array} \end{array}$$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

Pattern graphs  $(F_1, \dots, F_n)$

$$\varphi(G) = \begin{array}{c|c|c|c|c} F_1 & F_2 & F_3 & \dots & F_n \\ \hline 3 & 27 & 12 & & 9 \end{array}$$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

$G \simeq G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all** simple graphs  $F$

number of homomorphisms from  $F$  to  $G'$

Pattern graphs  $(F_1, \dots, F_n)$

$F_1$	$F_2$	$F_3$	...	$F_n$
3	27	12		9

$\varphi(G) =$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

$G \simeq G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all** simple graphs  $F$

number of homomorphisms from  $F$  to  $G'$

Homomorphism counts give us a **complete** embedding

Pattern graphs  $(F_1, \dots, F_n)$

$F_1$	$F_2$	$F_3$	...	$F_n$
3	27	12		9

$\varphi(G) =$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

number of homomorphisms  
from  $F$  to  $G'$

$G \simeq G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all** simple graphs  $F$

Homomorphism counts give us a **complete** embedding

Pattern graphs  $(F_1, \dots, F_n)$

$$\varphi(G) = \begin{array}{|c|c|c|c|c|} \hline & F_1 & F_2 & F_3 & \dots & F_n \\ \hline & 3 & 27 & 12 & & 9 \\ \hline \end{array}$$

**complete in expectation**

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

$G \simeq_{\mathcal{F}} G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all**  $F \in \mathcal{F}$  ← restrict pattern class

Pattern graphs  $(F_1, \dots, F_n)$

	$F_1$	$F_2$	$F_3$	...	$F_n$
$\varphi(G) =$	3	27	12		9

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

$G \simeq_{\mathcal{F}} G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all**  $F \in \mathcal{F}$  ← restrict pattern class

Homomorphism counts give us an **expressive** embedding

Pattern graphs  $(F_1, \dots, F_n)$        $\varphi(G) =$ 

$F_1$	$F_2$	$F_3$	...	$F_n$
3	27	12		9

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

## Why homomorphisms?

$G \simeq_{\mathcal{F}} G' \iff \text{hom}(F, G) = \text{hom}(F, G')$  **for all**  $F \in \mathcal{F}$  ← restrict pattern class

Homomorphism counts give us an **expressive** embedding

Pattern graphs  $(F_1, \dots, F_n)$

$F_1$	$F_2$	$F_3$	...	$F_n$
3	27	12		9

**expressive in expectation**

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

We add noise to the normalized counts to be **differentially private**

Pattern graphs  $(F_1, \dots, F_n)$

$$\varphi_{\text{DP}}(G) = \begin{array}{c} F_1 \quad F_2 \quad F_3 \quad \dots \quad F_n \\ \boxed{0.2 \quad 0.8 \quad 0.3 \quad \dots \quad 0.1} \end{array}$$

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

We add noise to the normalized counts to be **differentially private**

## Why differential privacy?

$\varphi_{\text{DP}}(G) \approx \varphi_{\text{DP}}(G')$  **for all**  $G, G'$  with  $d_{\text{edge}}(G, G') = 1$

Pattern graphs  $(F_1, \dots, F_n)$        $\varphi_{\text{DP}}(G) =$

$F_1$	$F_2$	$F_3$	...	$F_n$
0.2	0.8	0.3		0.1

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

We add noise to the normalized counts to be **differentially private**

## Why differential privacy?

$\varphi_{\text{DP}}(G) \approx \varphi_{\text{DP}}(G')$  **for all**  $G, G'$  with  $d_{\text{edge}}(G, G') = 1$

## How much noise?

We add **small-ish** noise that depends on a **local** notion of privacy

Pattern graphs  $(F_1, \dots, F_n)$        $\varphi_{\text{DP}}(G) =$ 

$F_1$	$F_2$	$F_3$	...	$F_n$
0.2	0.8	0.3		0.1

# How?

We count **homomorphisms** from **sampled** pattern graphs  $F_i$  to  $G$

We add noise to the normalized counts to be **differentially private**

Pattern graphs  $(F_1, \dots, F_n)$

$$\varphi_{\text{DP}}(G) = \underbrace{\begin{array}{|c|c|c|c|c|} \hline & F_1 & F_2 & F_3 & \dots & F_n \\ \hline 0.2 & 0.8 & 0.3 & & & 0.1 \\ \hline \end{array}}_{\text{expressive in expectation}}$$

# To wrap up

**TLDR:** We sample **noisy homomorphism densities** to obtain **provably** expressive and private graph representations

By kinda relaxing privacy a bit and looking at a *local* notion of it

By kinda relaxing expressivity and ensuring it in *expectation*

# To wrap up

**TLDR:** We sample **noisy homomorphism densities** to obtain **provably** expressive and private graph representations

By kinda relaxing privacy a bit and looking at a *local* notion of it

By kinda relaxing expressivity and ensuring it in *expectation*

We also run **experiments** and they work!

# To wrap up

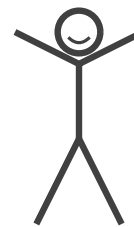
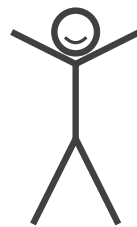
**TLDR:** We sample **noisy homomorphism densities** to obtain **provably** expressive and private graph representations

By kinda relaxing privacy a bit and looking at a *local* notion of it

By kinda relaxing expressivity and ensuring it in *expectation*

We also run **experiments** and they work!

**THANKS:)**



Tamara Patrick

Both moderately happy

