Batch Steganography and the Threshold Game

Andrew Ker

adk@comlab.ox.ac.uk

Royal Society University Research Fellow

Oxford University Computing Laboratory

SPIE/IS&T EI 2007

29 January 2007

Batch Steganography and the Threshold Game

Outline

- Batch steganography
- Definition of the "Threshold Game"
- Optimal strategies
- Experimental results
- Conclusions

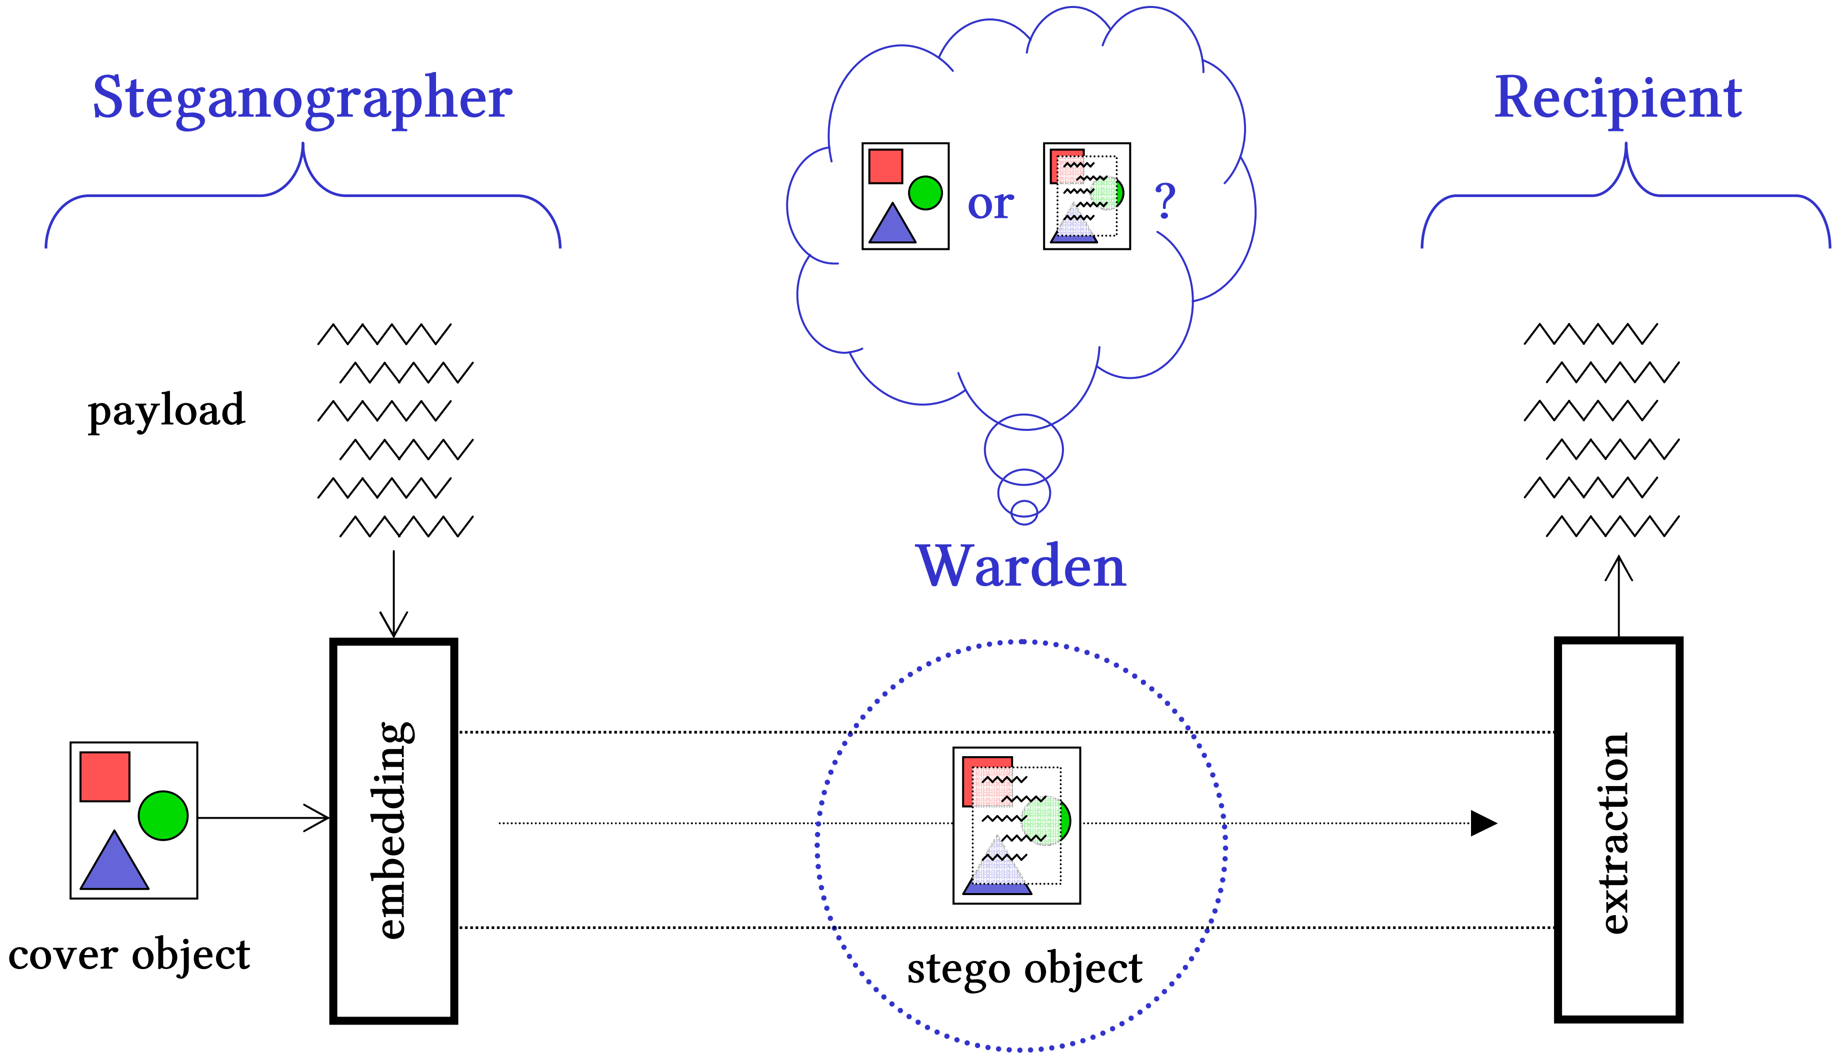
Steganographer
or
?
Recipient
payload
Warden
embedding
cover object
stego object
extraction

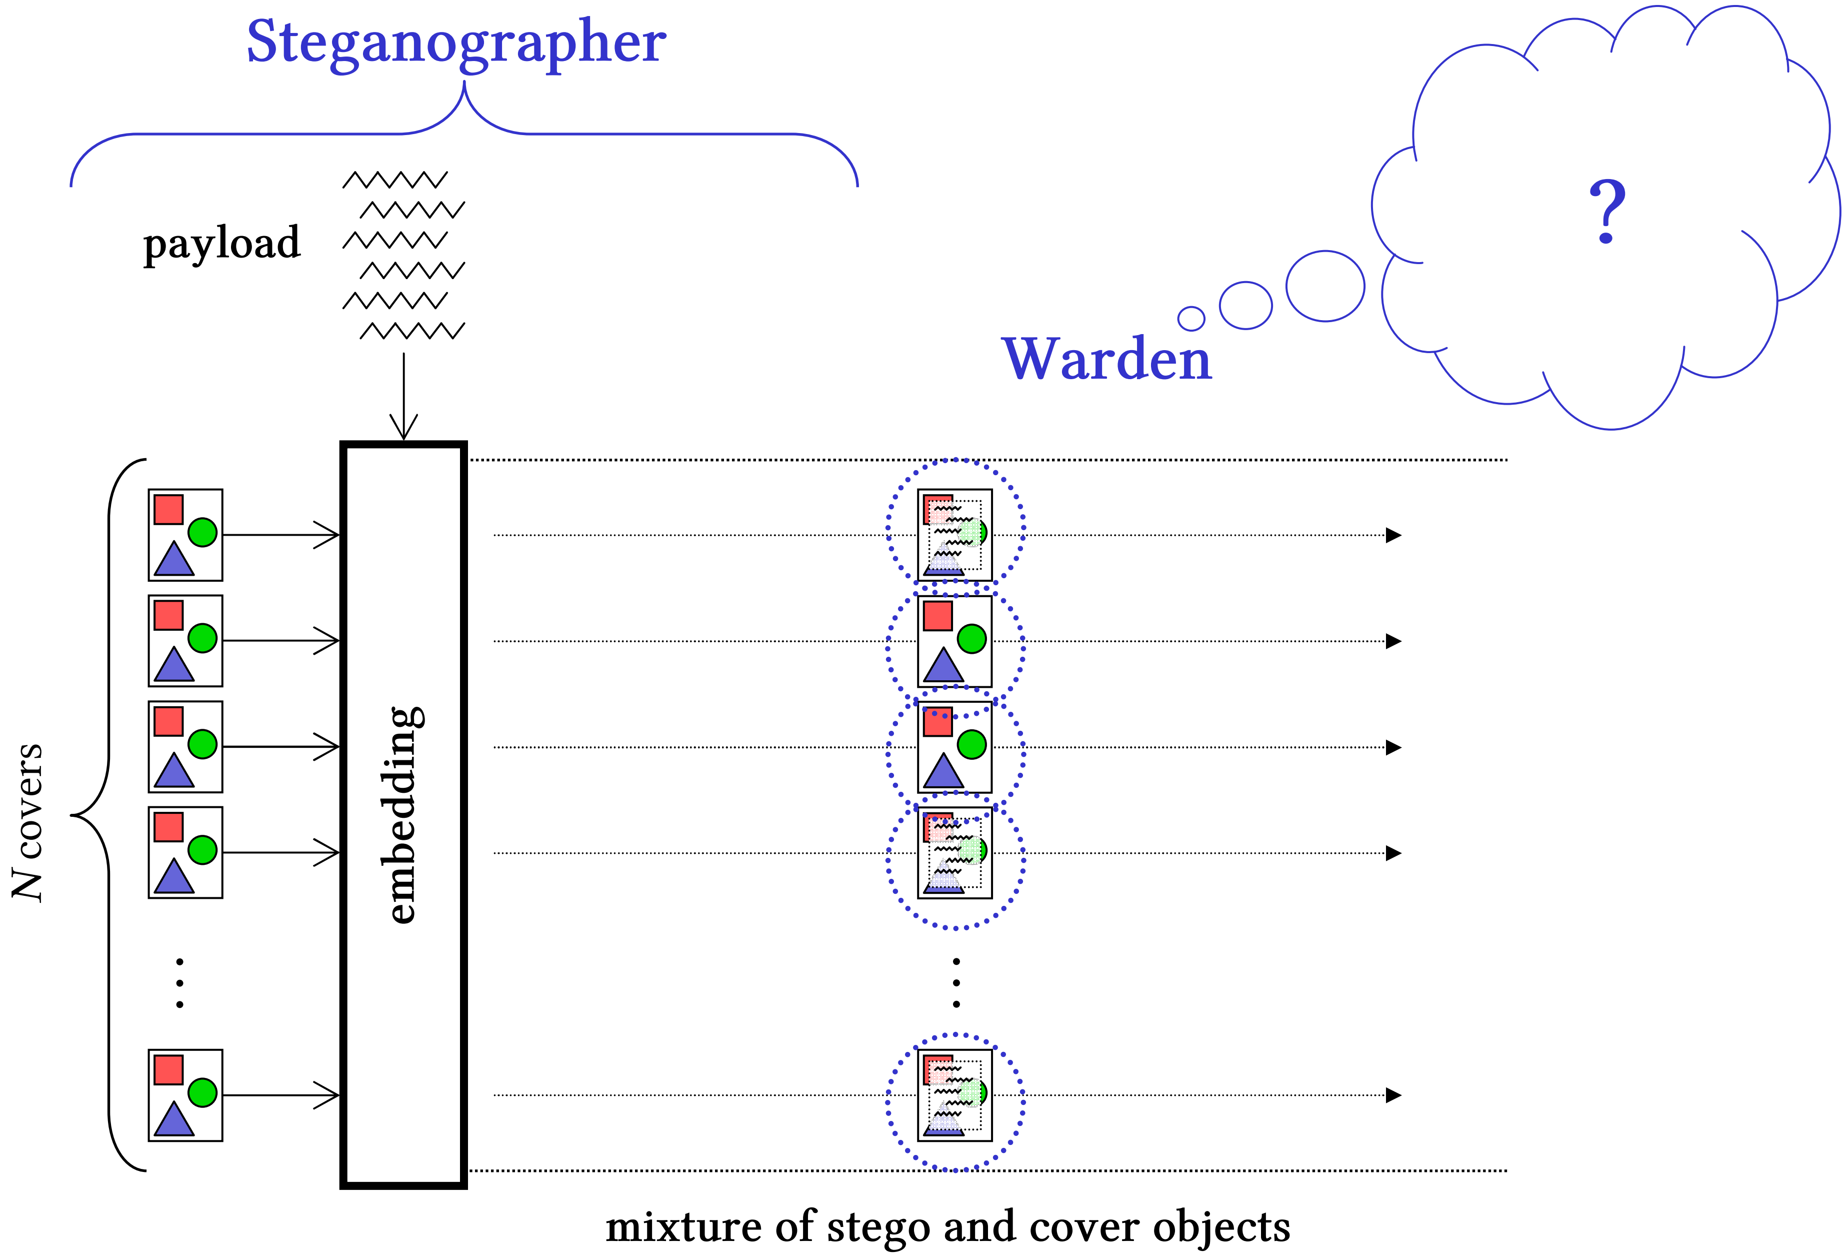

A. Ker, *Batch Steganography & Pooled Steganalysis*, Proc. 8th Information Hiding Workshop, 2006.

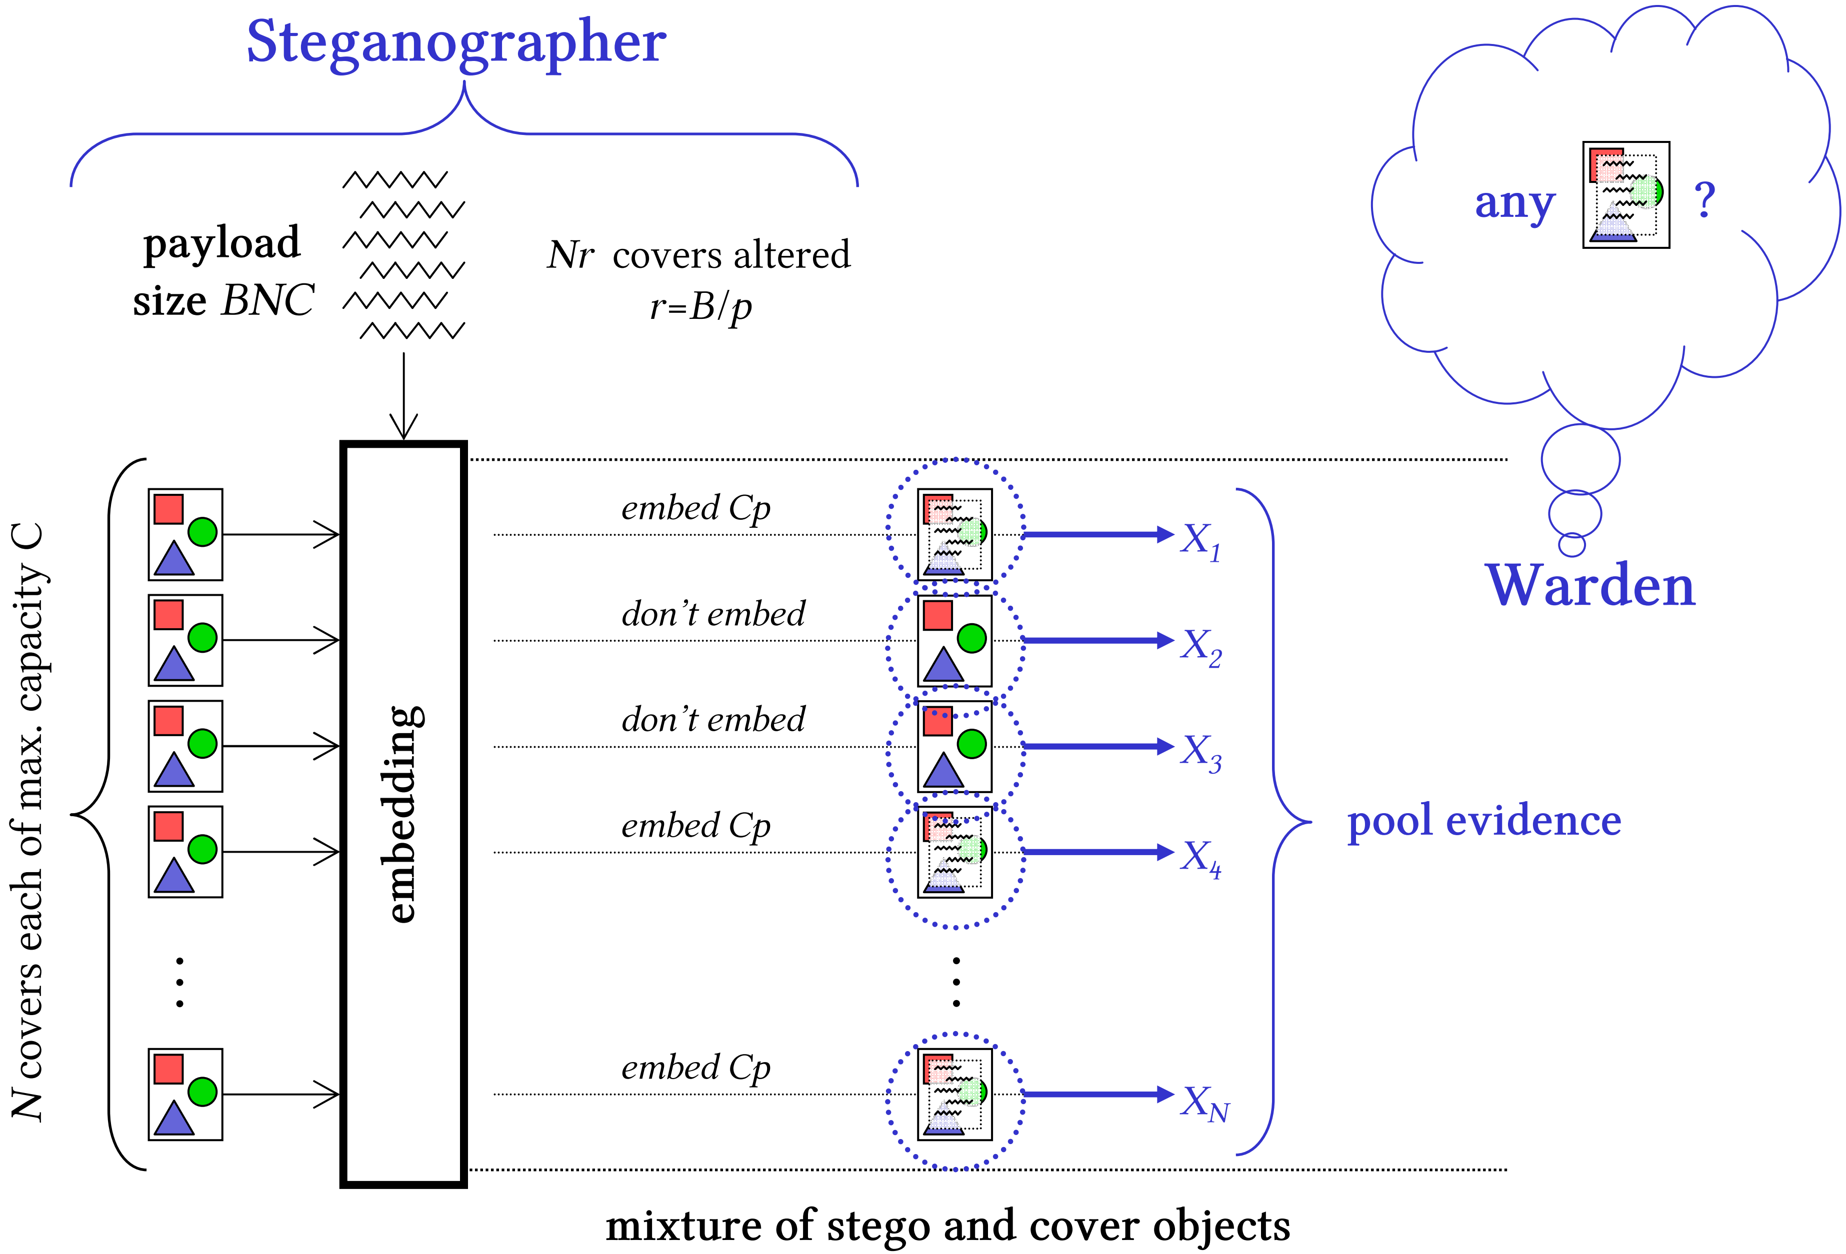

A. Ker, *Batch Steganography & Pooled Steganalysis*, Proc. 8th Information Hiding Workshop, 2006.

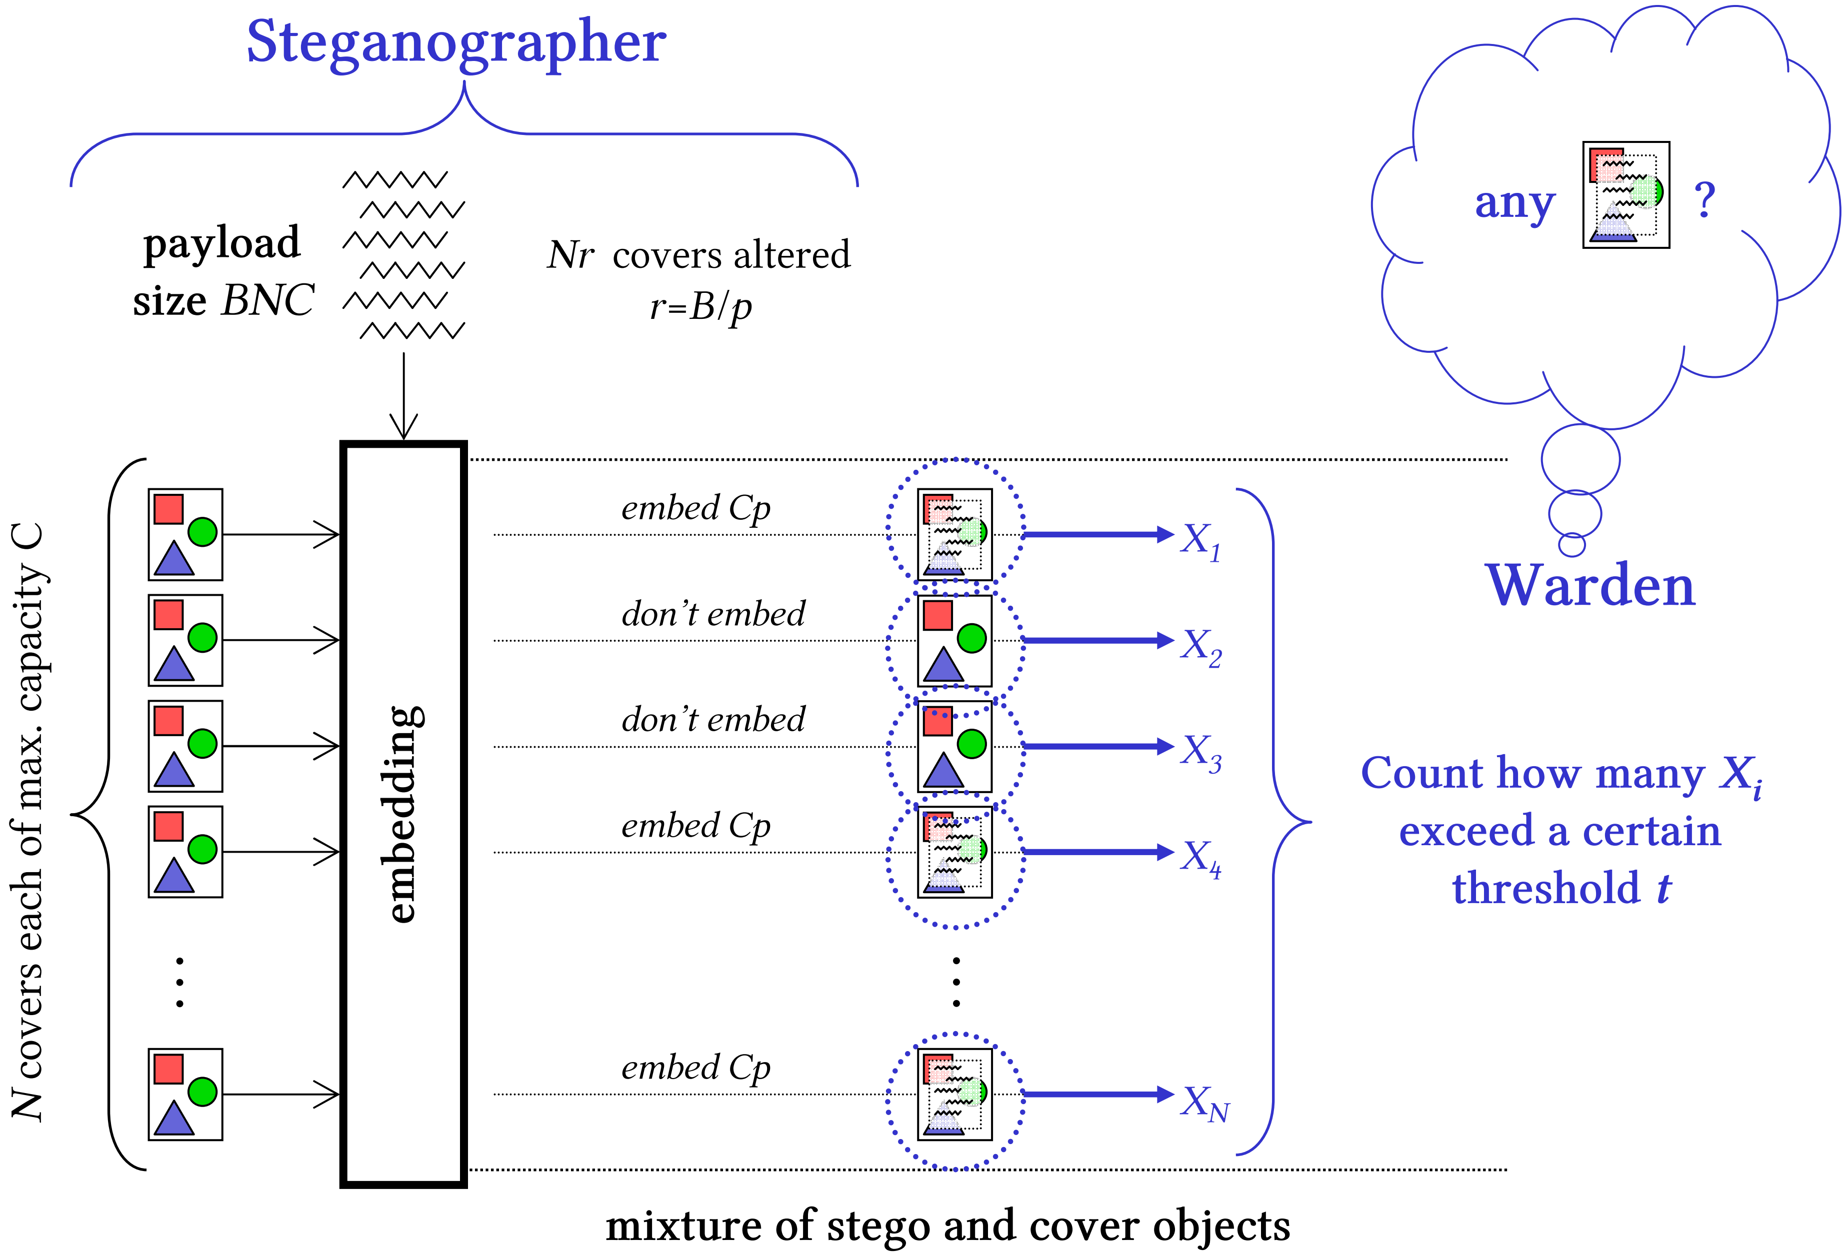

Steganographer
payload
size BNC
Nr covers altered
r=B/p
N covers each of max. capacity C
embedding
embed Cp
don't embed
don't embed
embed Cp
embed Cp
X1
X2
X3
X4
XN
any ?
Warden
Count how many Xi
exceed a certain
threshold t
mixture of stego and cover objects

Assumptions

The results rely on some assumptions:

1. Covers are uniform size;
2. That the Warden applies a **payload estimator** (“quantitative” detector) to individual objects, whose properties are known to both parties;
3. Estimation error is **additive** and **independent** of embedding rate;
4. Estimation error distribution has certain “bell-shaped” properties.

The Threshold Game

B is close to zero

Steganographer chooses $\boldsymbol{p} \in [B, 1]$

- Embeds Cp in each of Nr covers ($r = B/p$).

Warden chooses $\boldsymbol{t} \in \mathbb{R}$

- Counts the number of objects whose payload estimate exceeds t.

Payoff

The Threshold Game

Steganographer chooses $p \in [B, 1]$ *B is close to zero*

- Embeds Cp in each of Nr covers ($r = B/p$).

Warden chooses $t \in \mathbb{R}$

- Counts the number of objects whose payload estimate exceeds t.

(Steganographer's) Payoff

- False positive rate at 50% false negatives (i.e. median p-value).

Does depend on B and N, but is always a monotone decreasing function of

$$\chi_p(t) = \frac{\Psi(t) - \Psi(t-p)}{p\sqrt{\Psi(t)(1-\Psi(t))}}$$

Distribution function of payload estimate errors

which we call the "evidence function".

The Threshold Game

The Steganographer (p) wants to MINIMIZE, the Warden (t) to MAXIMIZE,

$$\chi_p(t) = \frac{\Psi(t) - \Psi(t-p)}{p\sqrt{\Psi(t)(1-\Psi(t))}}$$

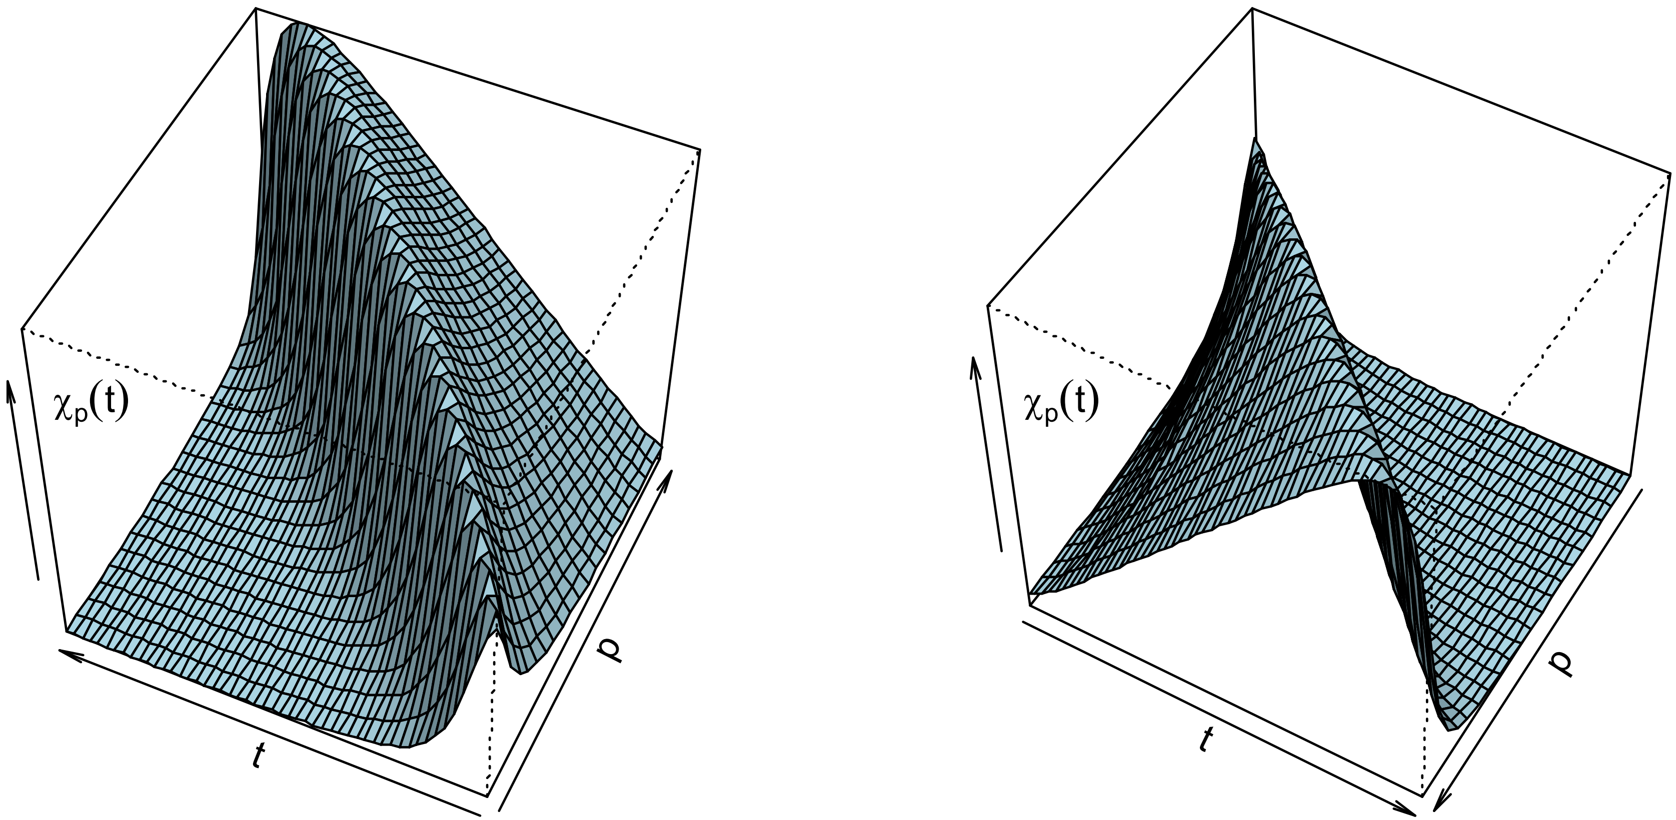

The Threshold Game

The Steganographer (p) wants to MINIMIZE, the Warden (t) to MAXIMIZE,

$$\chi_p(t) = \frac{\Psi(t) - \Psi(t-p)}{p\sqrt{\Psi(t)(1-\Psi(t))}}$$

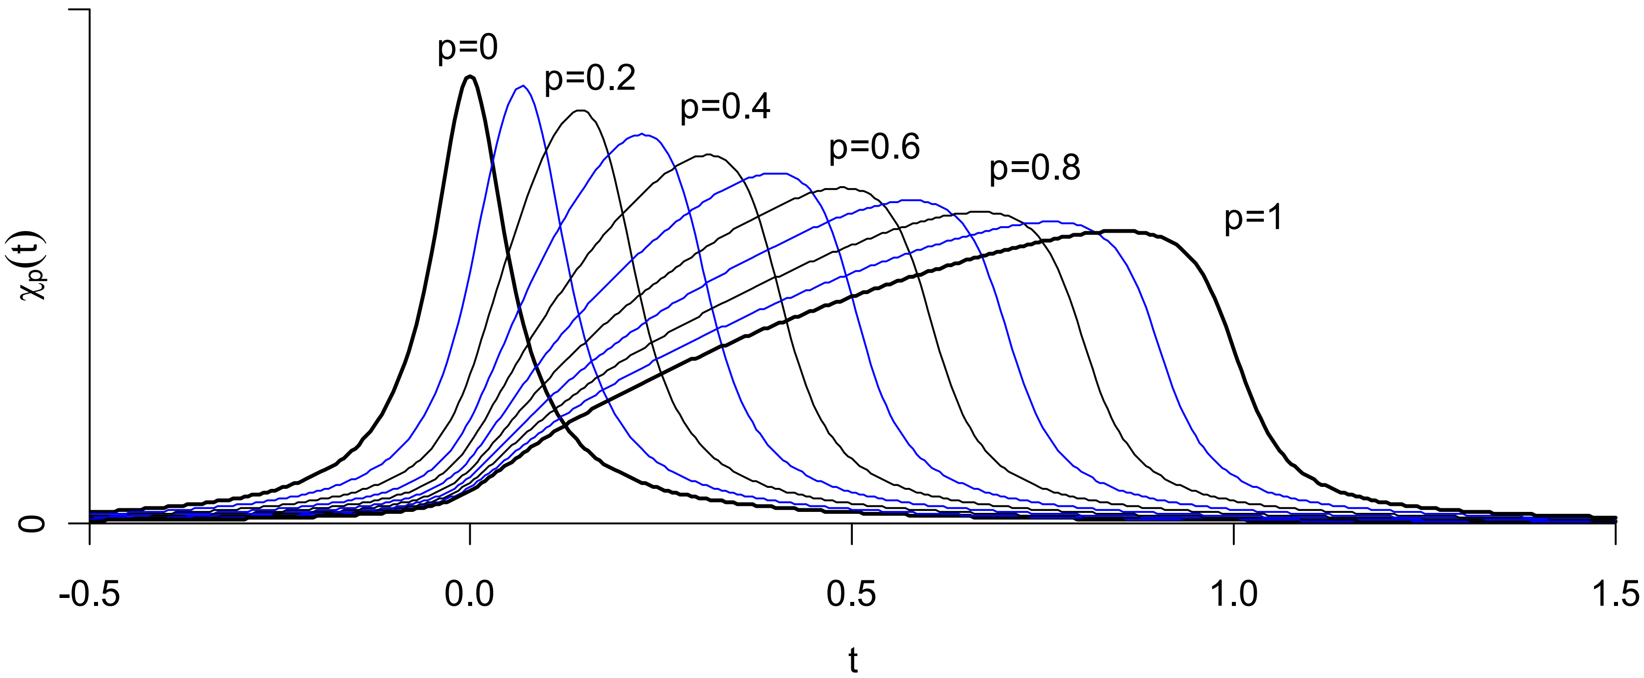

Solving the Threshold Game

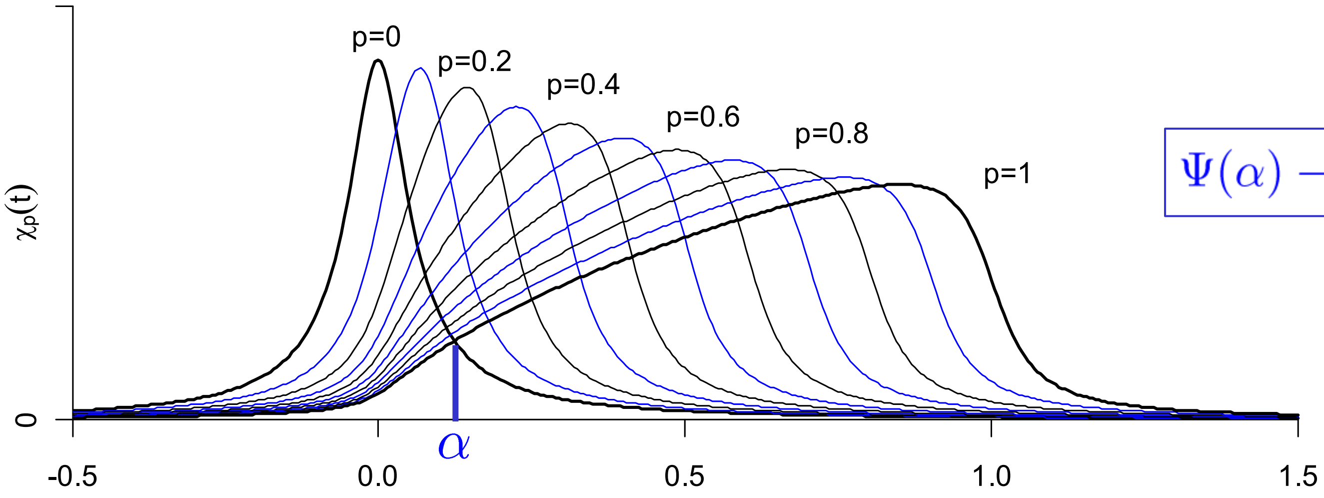

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for	
	Steganographer	Warden
Warden moves first		
Steganographer first		
Simultaneous moves		

Solving the Threshold Game

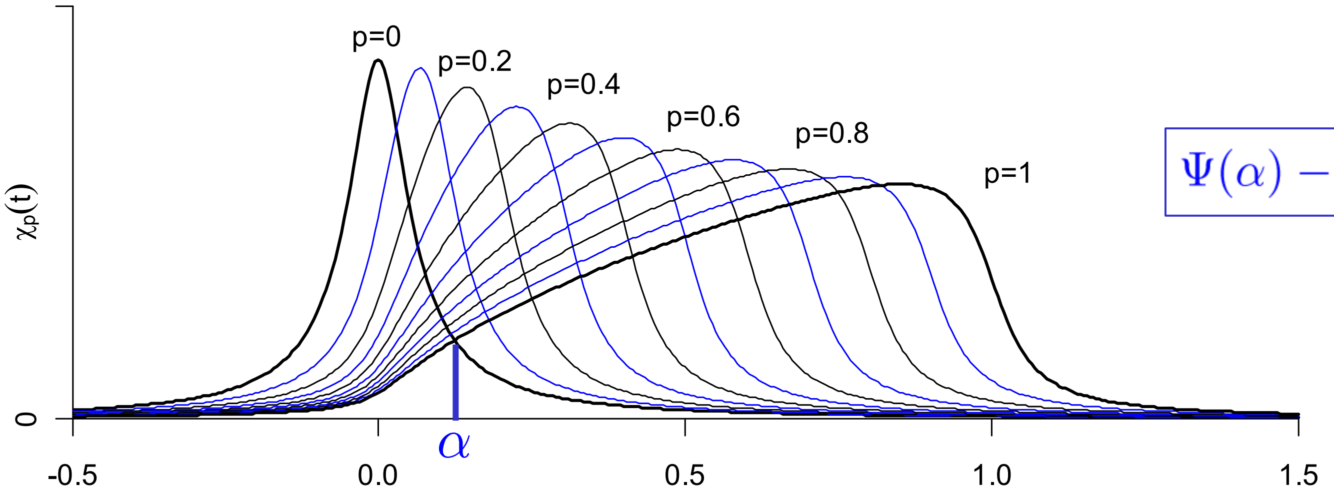

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for	
	Steganographer	Warden
Warden moves first	$p = 0$ or $p = 1$ (doesn't matter which)	$t = \alpha$
Steganographer first		
Simultaneous moves		

Solving the Threshold Game

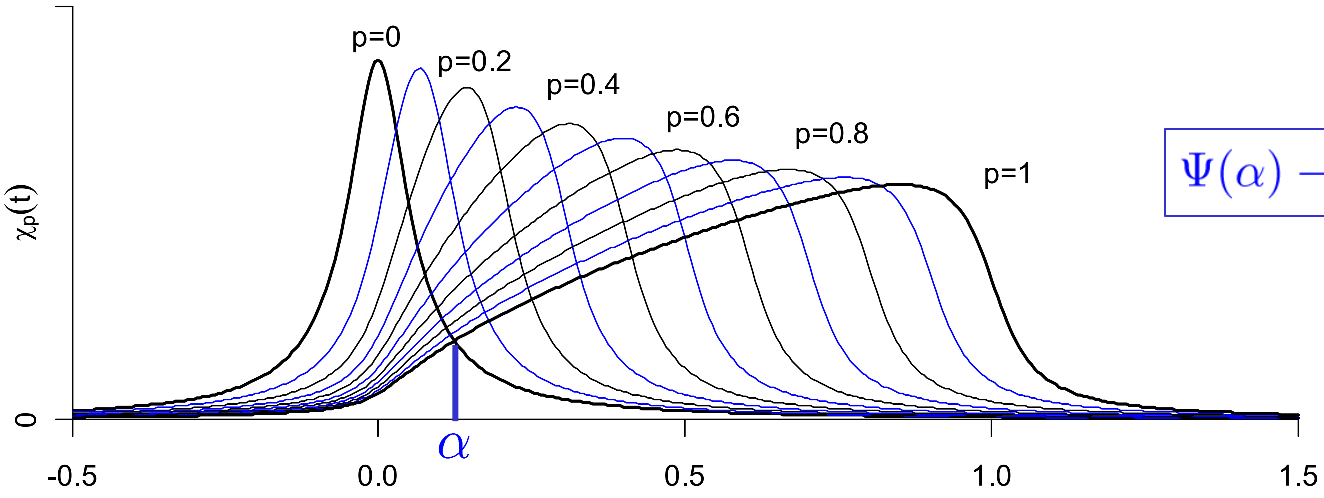

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for	
	Steganographer	Warden
Warden moves first	$p = 0$ or $p = 1$ (doesn't matter which)	$t = \alpha$
Steganographer first	$p = 0$ or $p = 1$ (depending on Ψ)	find corresponding peak
Simultaneous moves		

Solving the Threshold Game

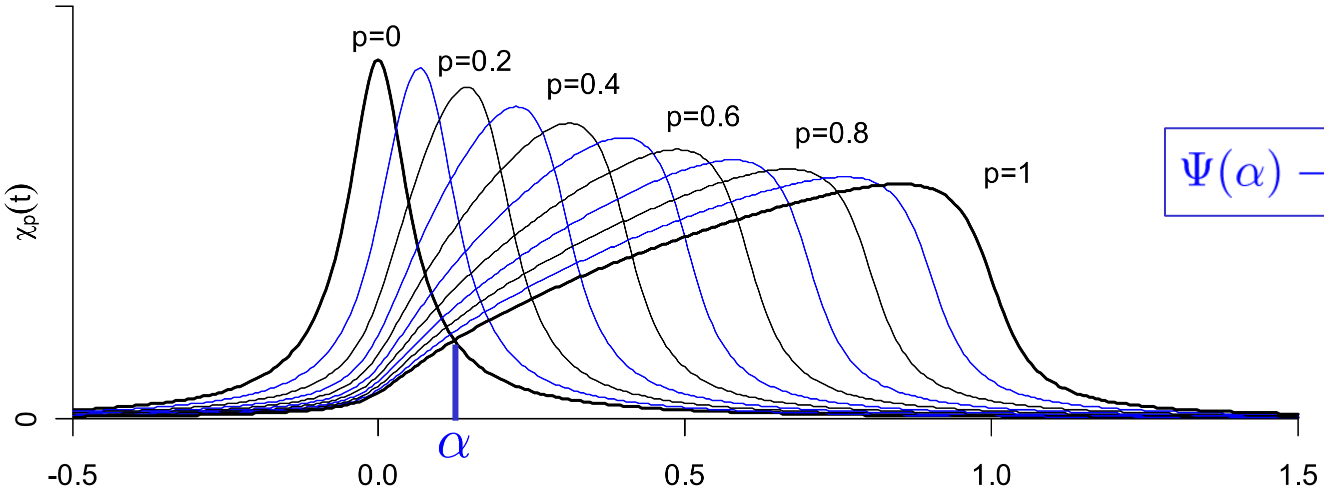

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for: Steganographer	optimal strategy for: Warden
Warden moves first	$p = 0$ or $p = 1$ (doesn't matter which)	$t = \alpha$
Steganographer first	$p = 0$ or $p = 1$ (depending on Ψ)	find corresponding peak
Simultaneous moves "equilibrium"	***No equilibrium in "pure" strategies.***	

Solving the Threshold Game

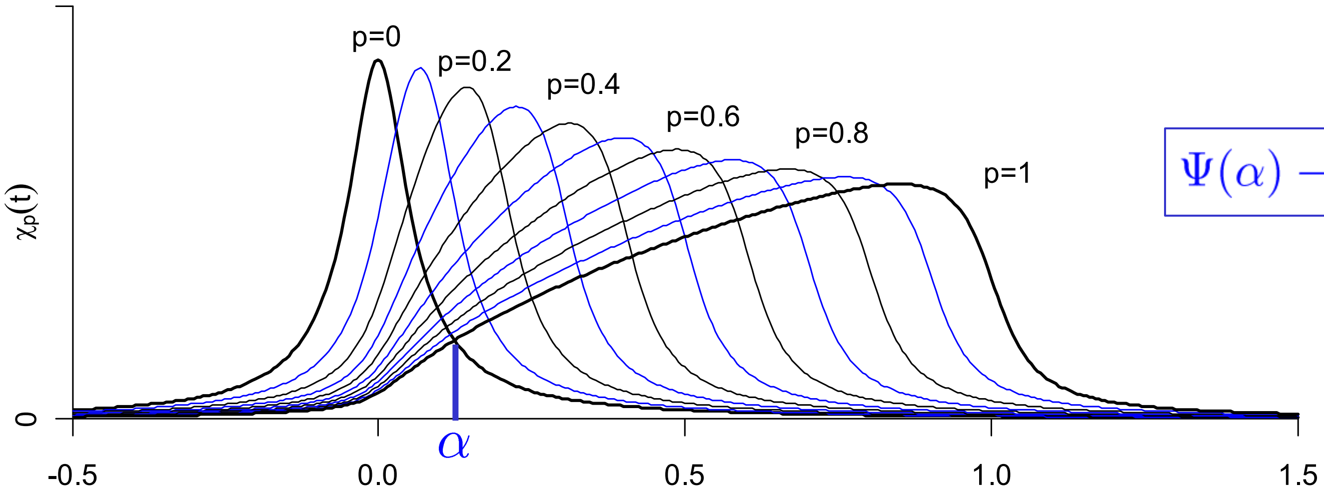

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for	
	Steganographer	Warden
Warden moves first	$p = 0$ or $p = 1$ (doesn't matter which)	$t = \alpha$
Steganographer first	$p = 0$ or $p = 1$ (depending on Ψ)	find corresponding peak
Simultaneous moves “equilibrium” mixed strategies	***Highly problematic: payoff is nonlinear function of mixture.***	

Solving the Threshold Game

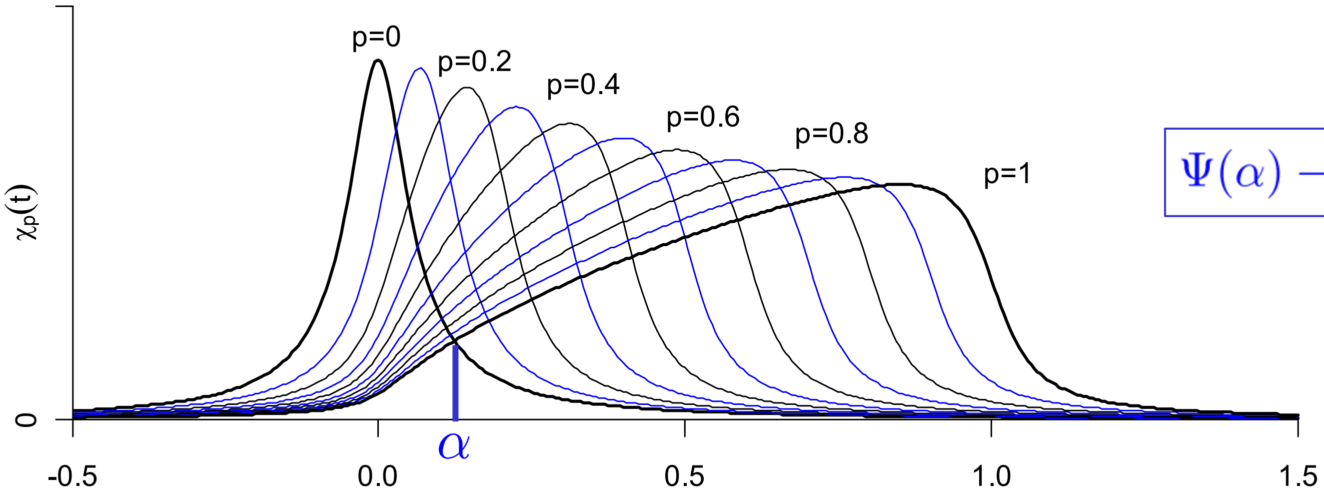

$$\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$$

	optimal strategy for	
	Steganographer	Warden
Warden moves first	$p = 0$ or $p = 1$ (doesn't matter which)	$t = \alpha$
Steganographer first	$p = 0$ or $p = 1$ (depending on Ψ)	find corresponding peak
Simultaneous moves "equilibrium" mixed strategies	$p = 0$ and $p = 1$ (pick at random with probabilities specified by Ψ)	$t = \alpha$

Experimental Results

Performed experiments using:

- Simple LSB replacement embedding;
- "Sample Pairs" payload estimation [Dumitrescu et al, IHW 2002];
- Large library of cover images.

Estimate Ψ from observations, then find the approximate solution for $\Psi(\alpha) - \Psi(\alpha - 1) = \Psi'(\alpha)$ numerically.

For this type/size of cover & detector, α=0.0585.

Experimental Results

For N=10000, B=0.005, all $t \in \{0,0.001,0.002,\ldots,1\}$ and all $p \in \{0,0.001,0.002,\ldots,1\}$

Theory predicts optimal t at α=0.0585; observed optimal t=0.057

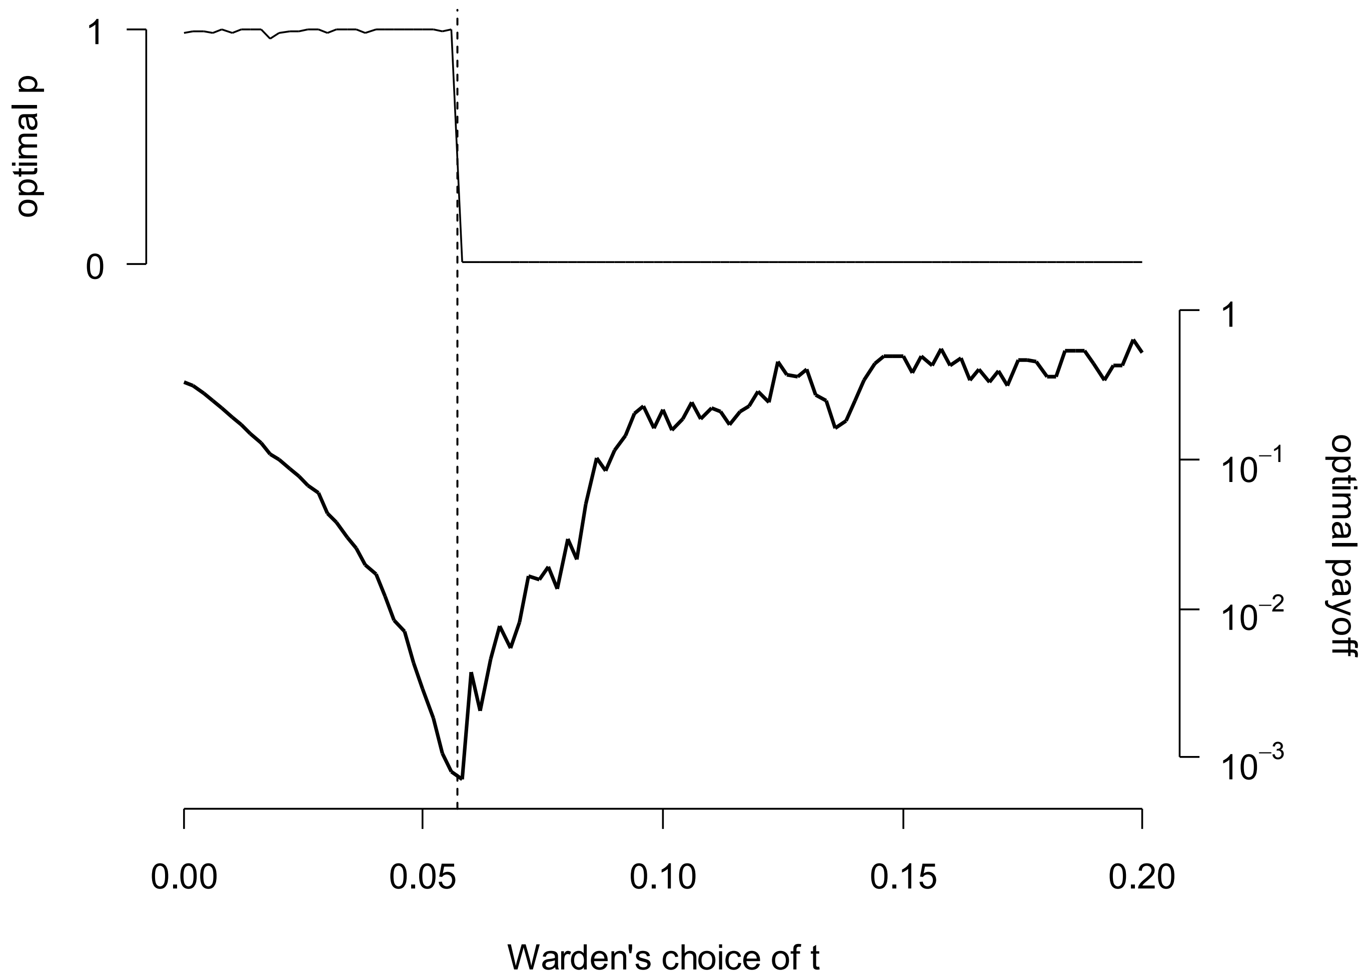

Conclusions

- The Threshold Game is not a completely implausible scenario.

The Game has solution/equilibrium which is surprisingly simple.

- The Warden should choose $t = \alpha$, the root of an equation which depends (only) on the shape of the error distribution.

When comparing quantitative estimators, this is the "optimum" point on the ROC curve: useful for benchmarking.

- The Steganographer should **mix** between total concentration of steganography and total spreading, but never adopt an intermediate option.

It is not clear whether this is informative for practical steganography.

End

`adk@comlab.ox.ac.uk`