Modelling regimes with Bayesian network mixtures

Marcus Bendtsen and Jose M. Peña
Department of Computer and Information Science
Linköping University

LINKÖPING UNIVERSITY

Hidden Markov Model

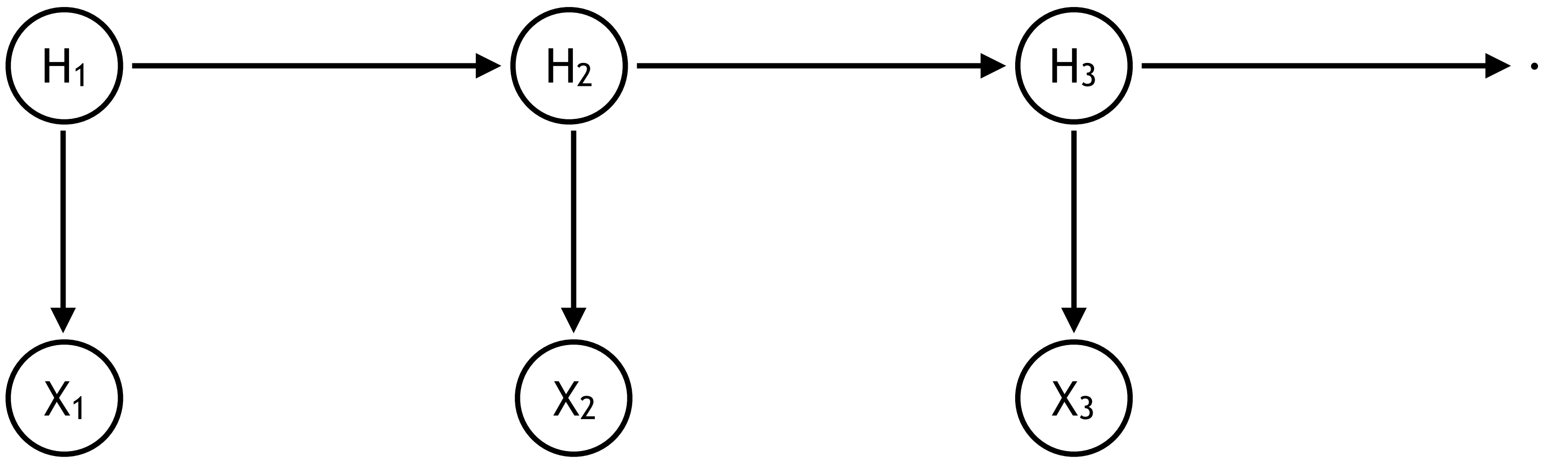

$p(H_1 = fair) = 0.5$

$p(X_1 = heads \mid H_1 = fair) = 0.5$

$p(X_1 = heads \mid H_1 = loaded) = 0.8$

$p(H_2 = fair \mid H_1 = fair) = 0.3$

$p(H_2 = fair \mid H_1 = loaded) = 0.8$

$p(X_3 = heads \mid X_1 = heads, X_2 = heads)$?

$p(H_3 = fair \mid X_1 = heads, X_2 = heads)$?

Hidden Markov Model

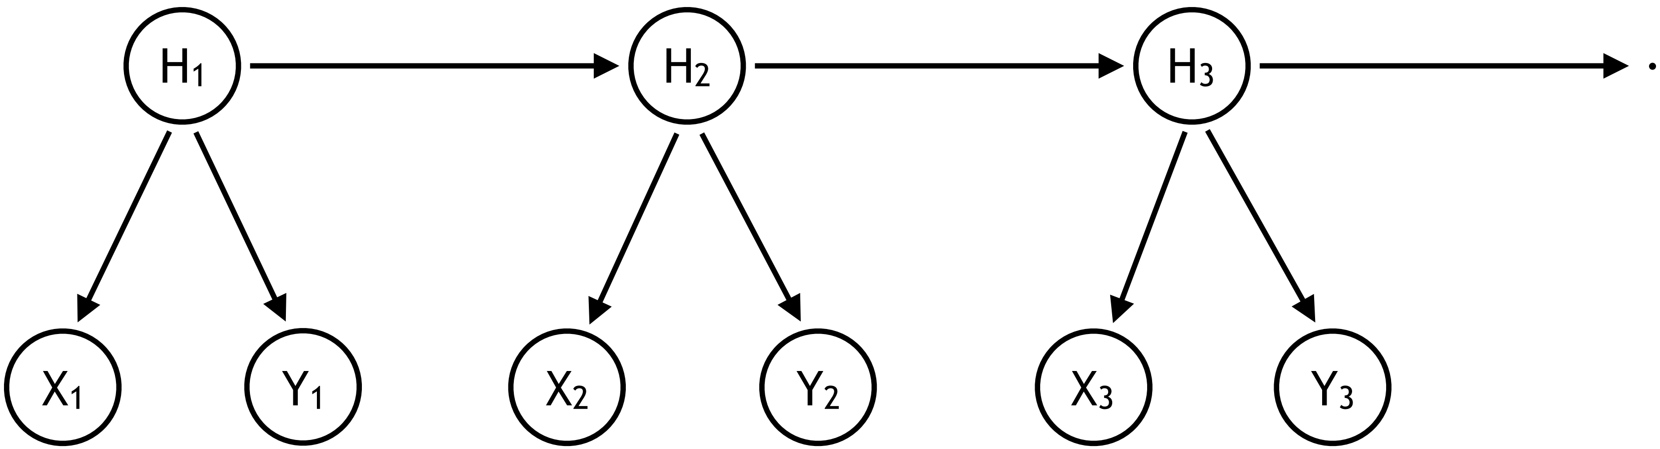

Hidden Markov Model

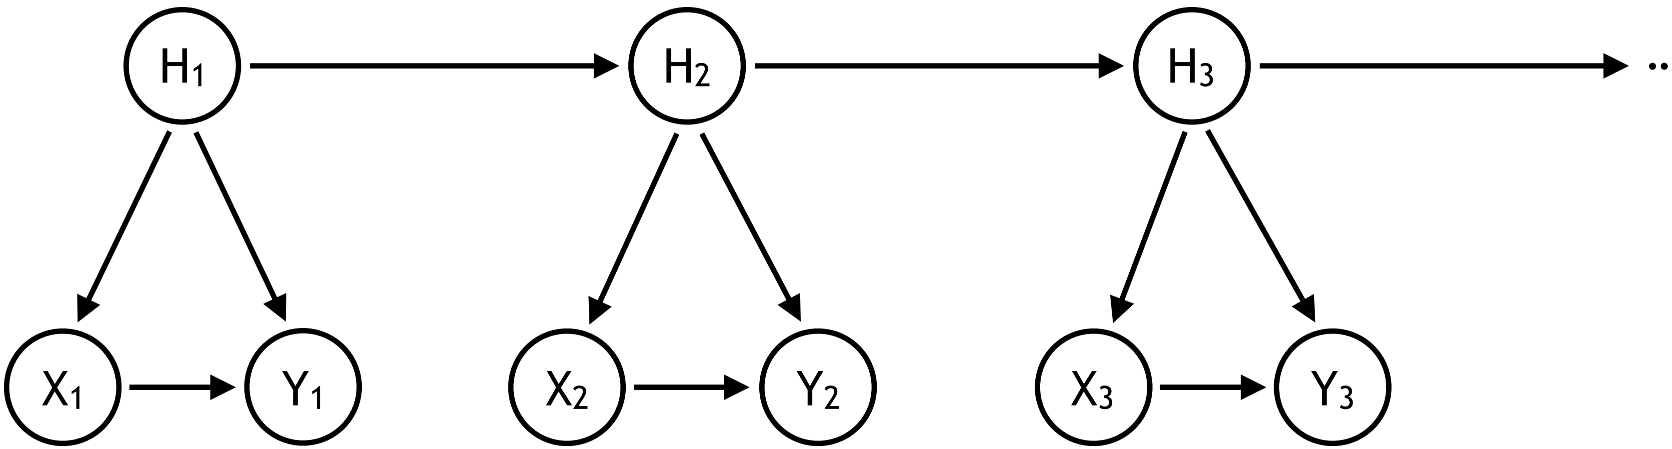

Aim

- The aim of the current project is to extend HMMs with two features:
    - Use different Bayesian networks for each state of the hidden variables.
    - Allow for one of the observable variables to directly influence the value of the next hidden state.

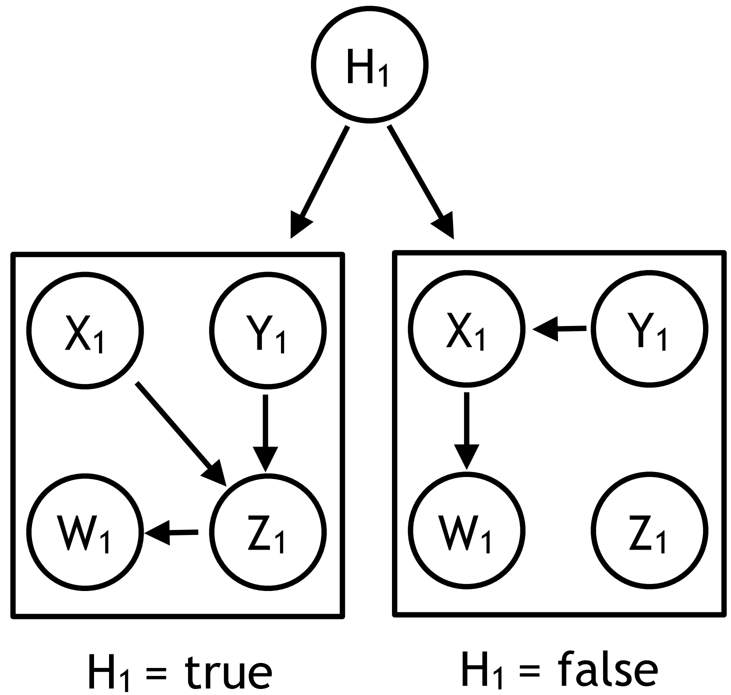
H1
X1
Y1
W1
Z1
X1
Y1
W1
Z1
H1 = true
H1 = false

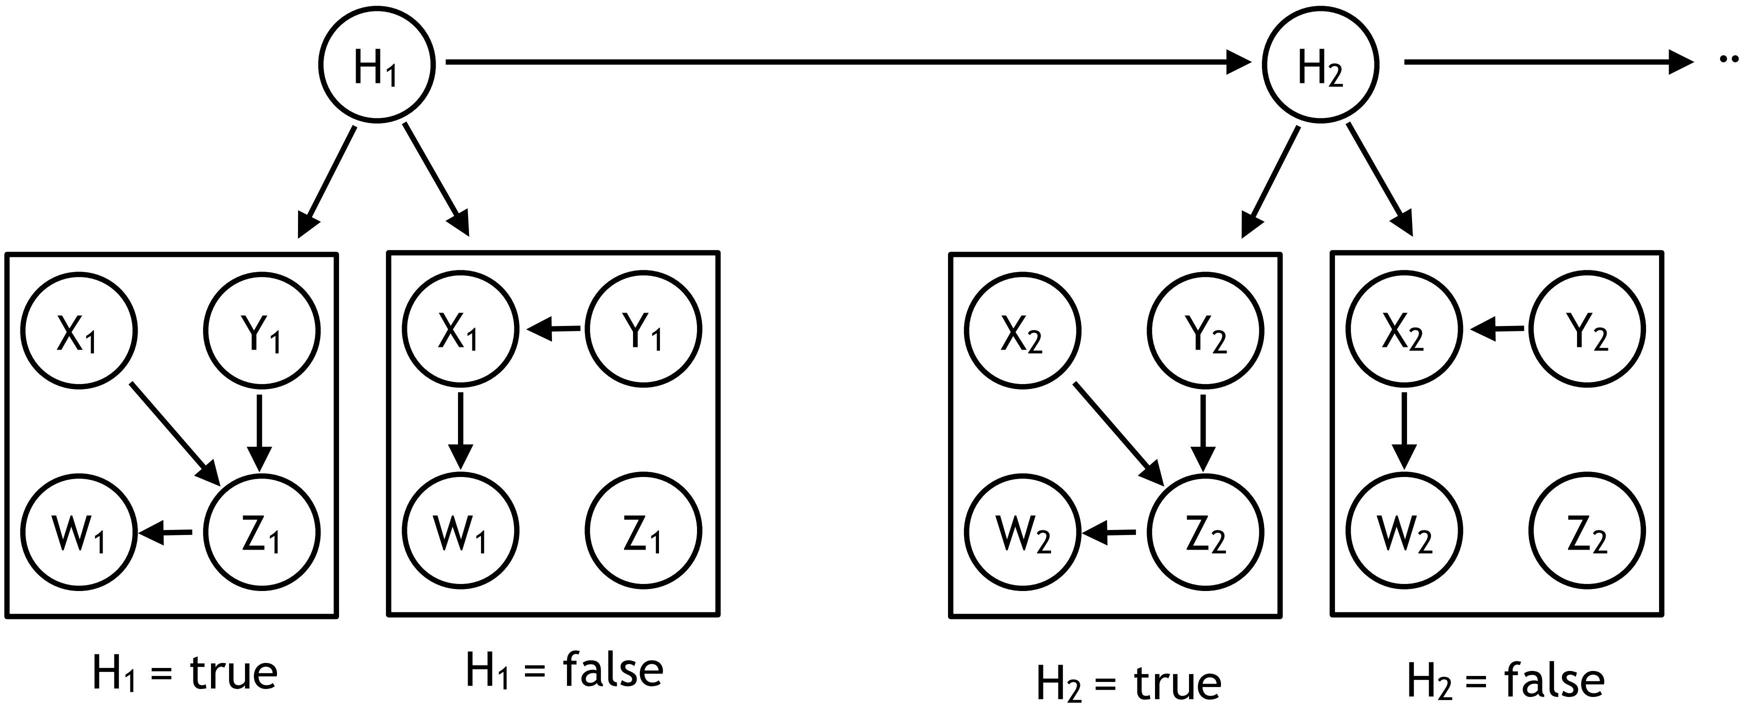

H1
H2
...
X1
Y1
W1
Z1
X1
Y1
W1
Z1
X2
Y2
W2
Z2
X2
Y2
W2
Z2
H1 = true
H1 = false
H2 = true
H2 = false

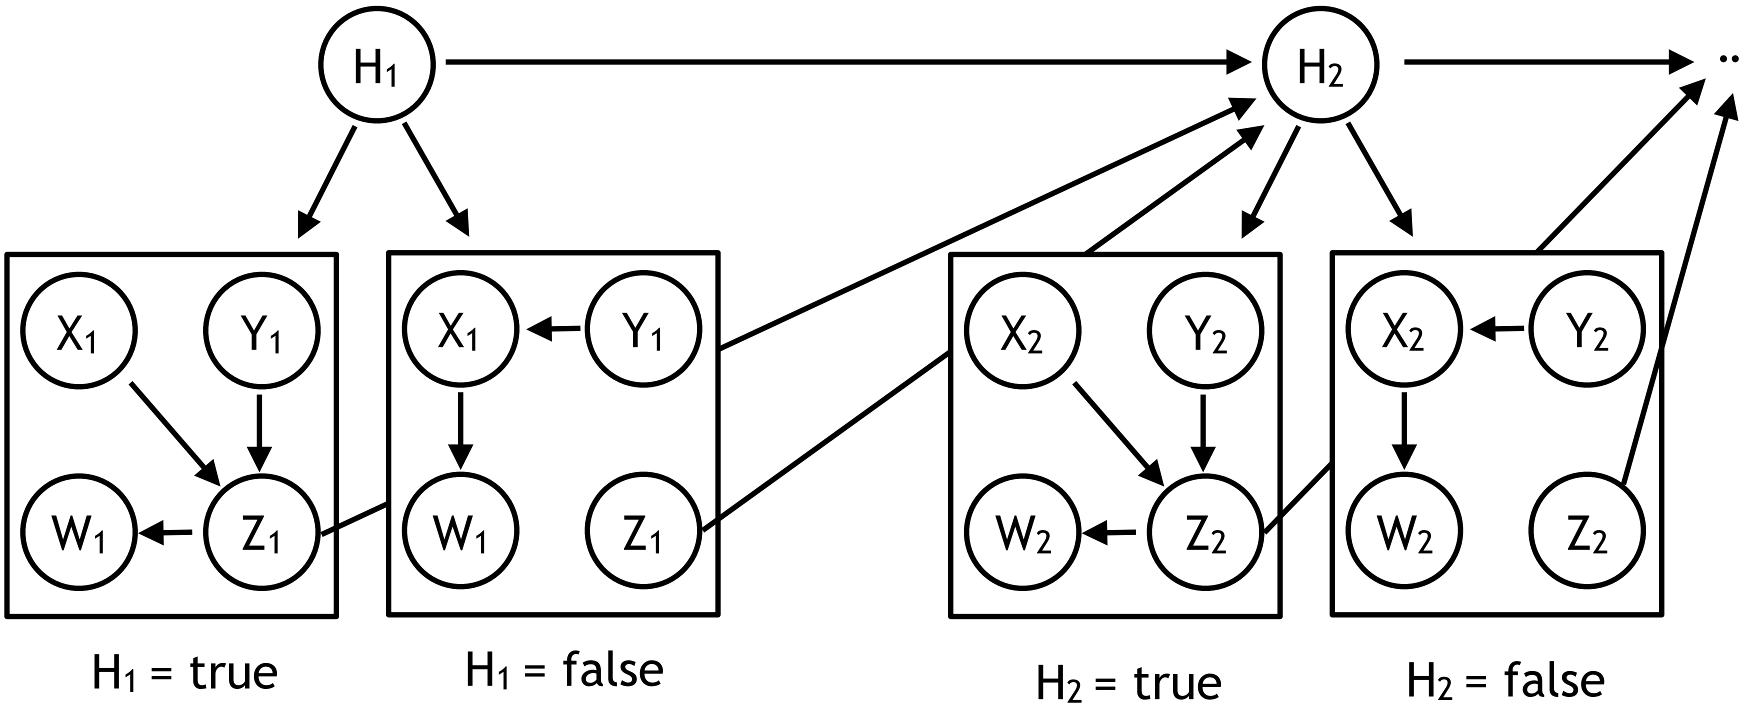

H1
H2
...
X1
Y1
W1
Z1
X1
Y1
W1
Z1
X2
Y2
W2
Z2
X2
Y2
W2
Z2
H1 = true
H1 = false
H2 = true
H2 = false

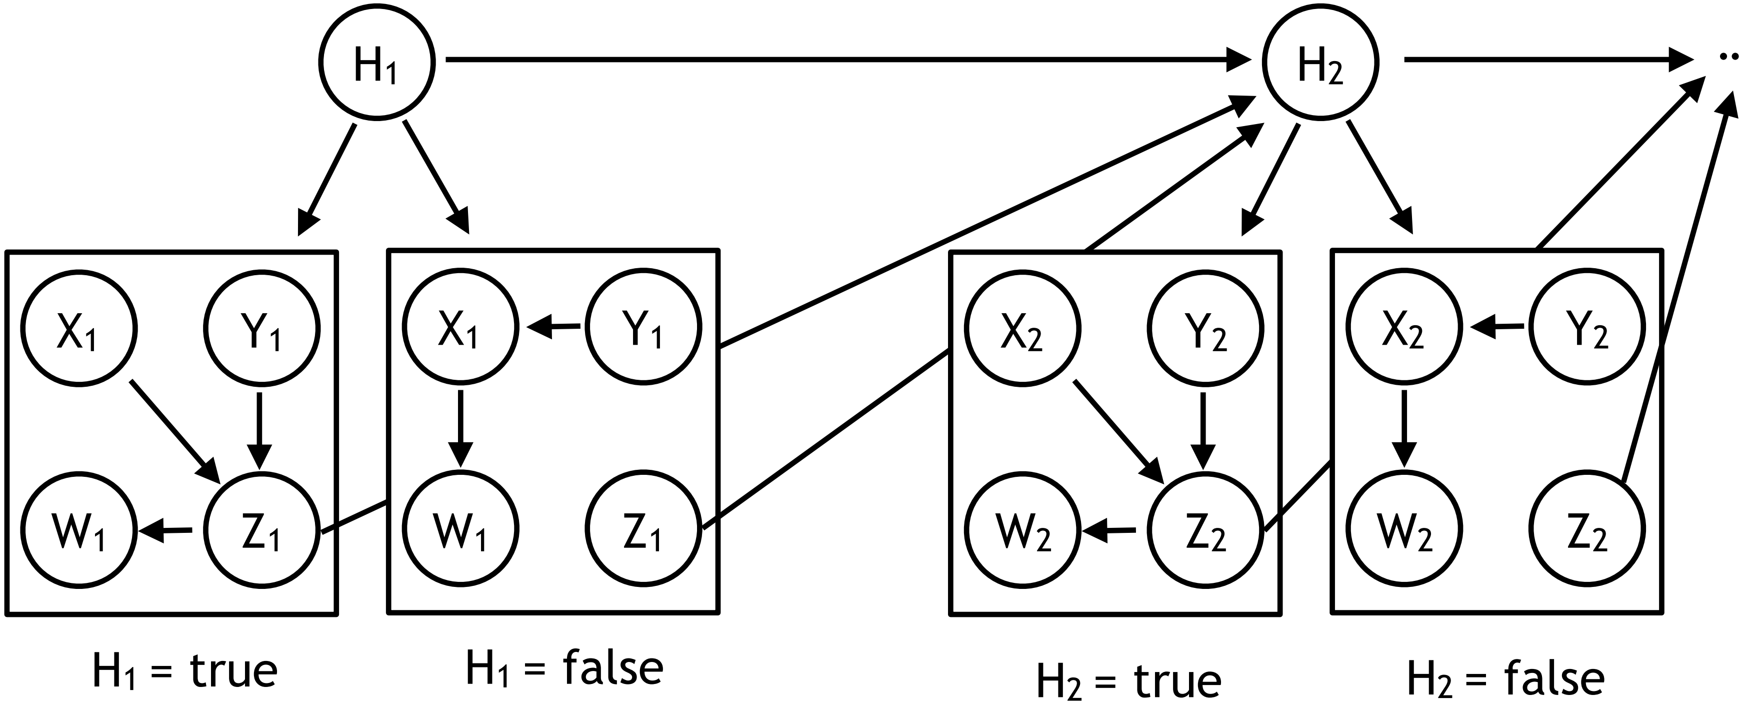

$p(H_1)$

$p(X_1,Y_1,W_1,Z_1|H_1=true)$

$p(X_1,Y_1,W_1,Z_1|H_1=false)$

$p(H_2|H_1,Z_1)$

$p(Z_2 = true \; |X_1,Y_1,W_1,Z_1)$?

$p(H_2 = true \; |X_1,Y_1,W_1,Z_1)$?

Regime mixture model

Learning

- We use **expectation maximisation** to estimate the parameters of the conditional and marginal distributions.
- We include a **BN structure learning algorithm** within each iteration to identify the different BN structures.
- All details are in the paper.

Trading the stock market

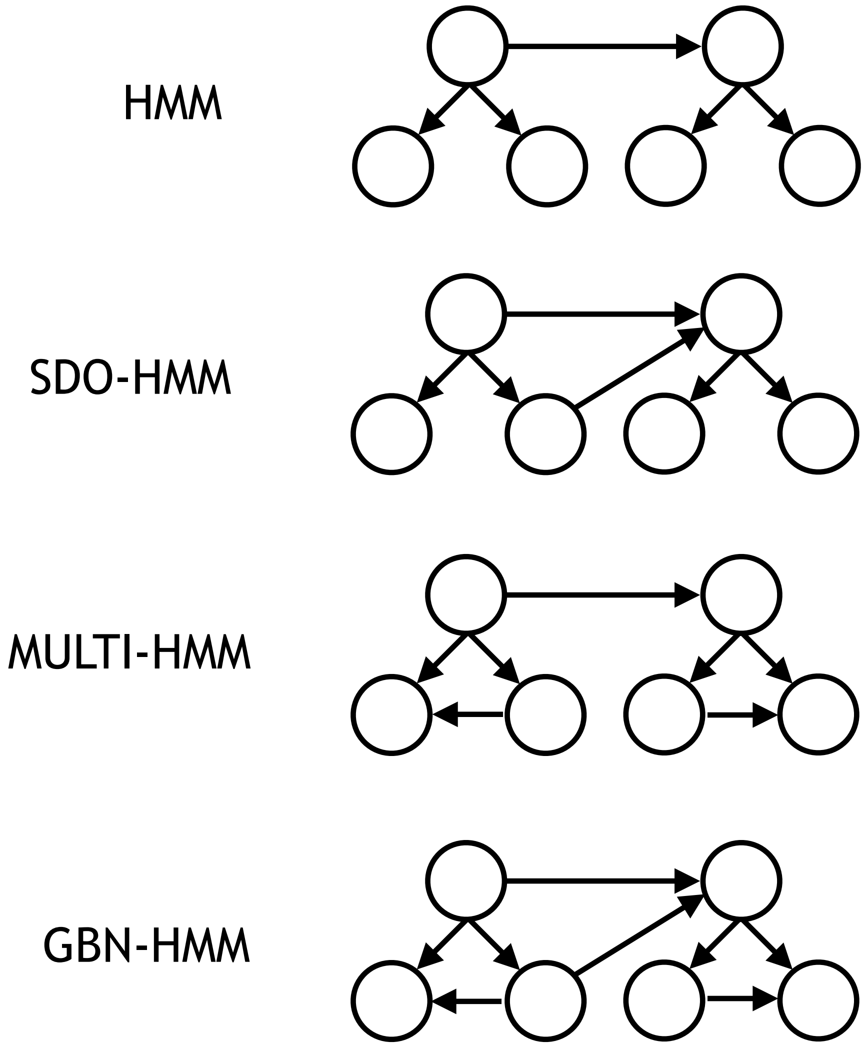

We used four observed variables which all were based on **technical analysis indicators.**

These indicators are computations of price, e.g. the relative strength index which measures which side of the trade (buyers or sellers) have been dominant in the recent past.

The “Z” variable is a comparison between two moving averages with different window size, creating an **indicator of trend**. This variable was discretised into two states (positive and negative).

Data

$t = 0$ $t = T$

Given a fully specified model, we want to generate buy and sell signals for each time point t.

Therefore, at time t we calculate the probability of the trend variable being positive or negative in the next time step.

For each data point at time t:

- Calculate $\alpha = p(Z_{t+1} = \text{positive} \mid D_{1:t})$.
- Calculate $\beta = p(Z_{t+1} = \text{negative} \mid D_{1:t})$.
- Generate a buy signal if $\alpha > \theta$ and a sell signal if $\beta > \theta$.

Over the training data, use cross-validation to decide upon:

- Number of states of the hidden variables.
- Estimate parameters of the conditional and marginal distributions.
- Decide upon the structure of the different BNs (when applicable).
- Decide theta (probability required to generate a signal).

Calculate a set of returns (one return for each held out test block).
Calculate mean(returns) / sd(returns) = Sharpe ratio.

	HMM	SDO-HMM	MULTI-HMM	GBN-HMM
Apple	0.844	0.708	**0.849**	0.718
Amazon	0.466	0.580	0.449	**0.592**
IBM	**0.713**	0.521	0.699	0.616
Microsoft	0.091	-0.189	-0.307	**0.219**
Red Hat	-0.198	**0.111**	-0.780	-0.085
Nvidia	0.113	0.211	0.262	**0.308**
General electric	0.0621	0.362	-0.378	**0.419**

What is in the paper?

- Full description of all the computations necessary for the expectation maximisation iterations.
- Complementary material available for the derivation of these computations.
- Synthetic experiments that show how the different models compare with respect to likelihood of data.
- More details about the stock trading experiments.

www.liu.se

LINKÖPING UNIVERSITY