Probability Elicitation for Bayesian Networks

Marek J. Drużdżel

Politechnika Białostocka

m.druzdzel@pb.edu.pl
http://aragorn.wi.pb.bialystok.pl/~druzdzel

Session overview

- Elicitation of probabilities
- Learning probabilities from data
- Canonical (a.k.a. ICI, independence of causal influences) models (separate lecture)
- Do parameters matter?
- Other relevant issues
  - Sensitivity analysis
  - Strength of influence
  - Value of information
  - Clarity test

What I want you to know after this session?

- Be familiar with the idea of obtaining subjective probabilities from experts
- Understand how parameters are learned from data
- Do not worry too much about precision of numerical probabilities, use sensitivity analysis
- Know sensitivity analysis, strength of influences, value of information, and clarity test

Elicitation of probabilities

Are probabilities in experts' heads?

William Estes' probability matching experiments

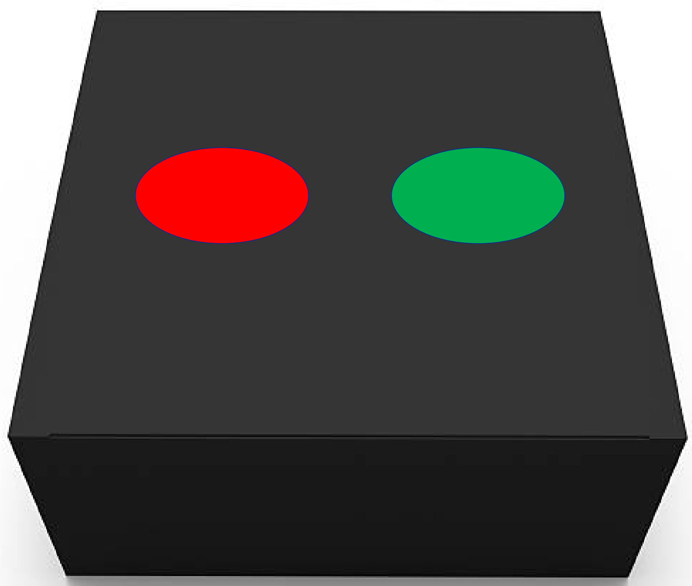

- Correct guessing of the light to be lit carries a small reward.
- Many repetitions.
- What is the optimal strategy if Pr(green)=0.3 and Pr(red)=0.7?
- What do human subjects do?

They keep guessing but guess green 30% of the time and red 70% of the time.

Elicitation of probabilities

Three fundamental methods:

- Ask directly
- Reference lottery
- Symmetric bets

Three additional issues:

- Assessing continuous distributions
- Discretization of continuous distributions
- Decomposition

Elicitation of probabilities: Direct assessment

Example:

"What is your belief regarding the probability that event A will occur?"

Graphical aids that make it indirect: Probability wheel

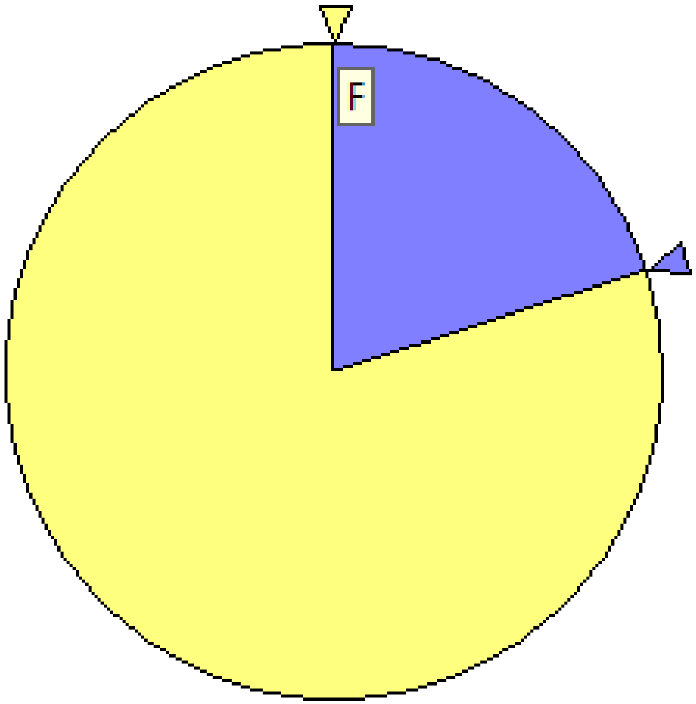

Bar chart

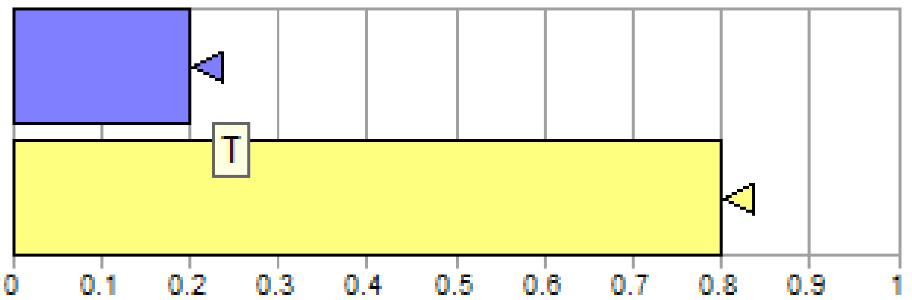

Elicitation of probabilities: Reference lottery

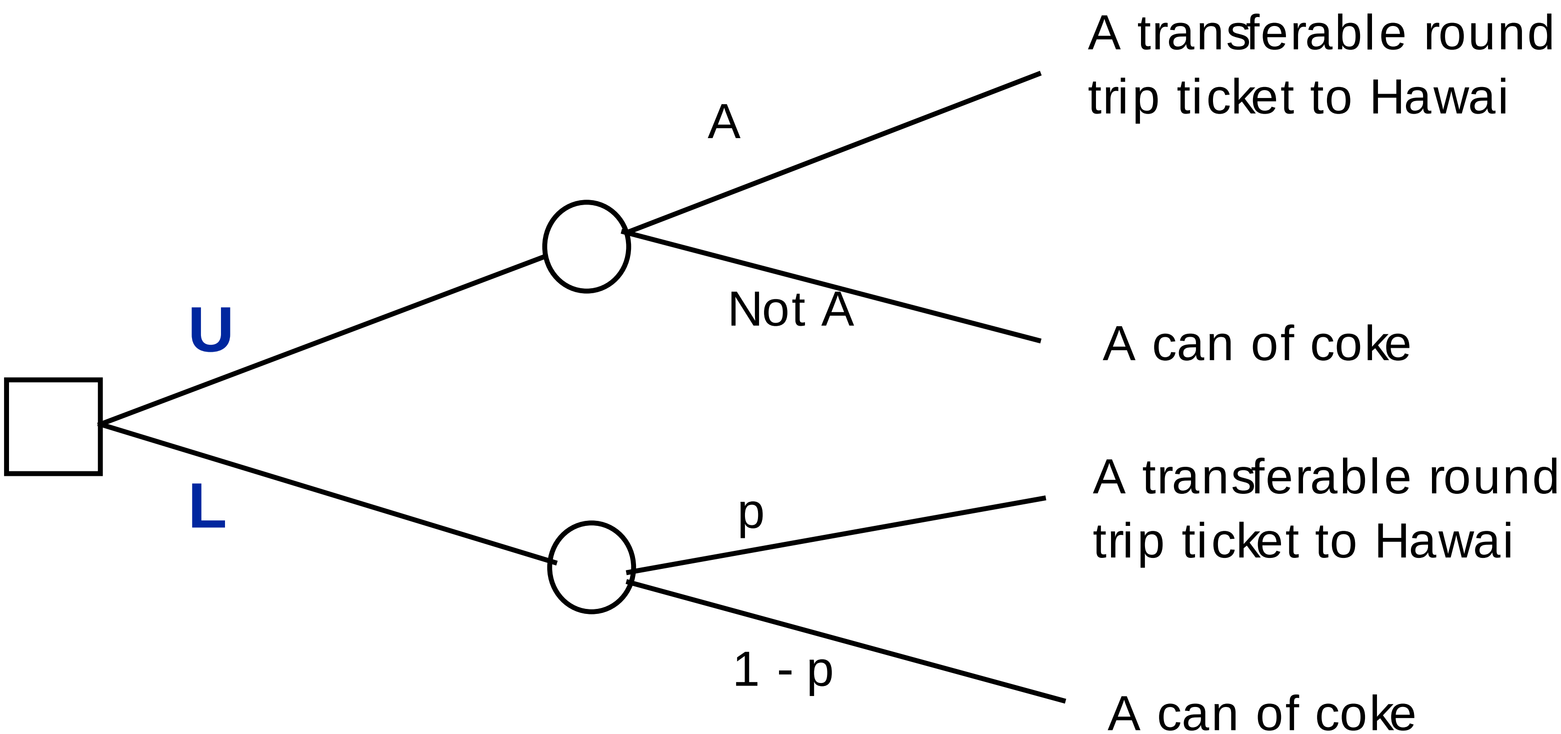

Use a tool like probability wheel (to hide the numbers).

Elicitation of probabilities: Symmetric bets

Offer choice between two lotteries, adjust values until the expert is indifferent between the two lotteries.

Then we have:

$$p X - (1 - p) Y = - p X + (1 - p) Y$$

which yields

$$p = 2 Y / (2 X + 2 Y) = Y / (X + Y)$$

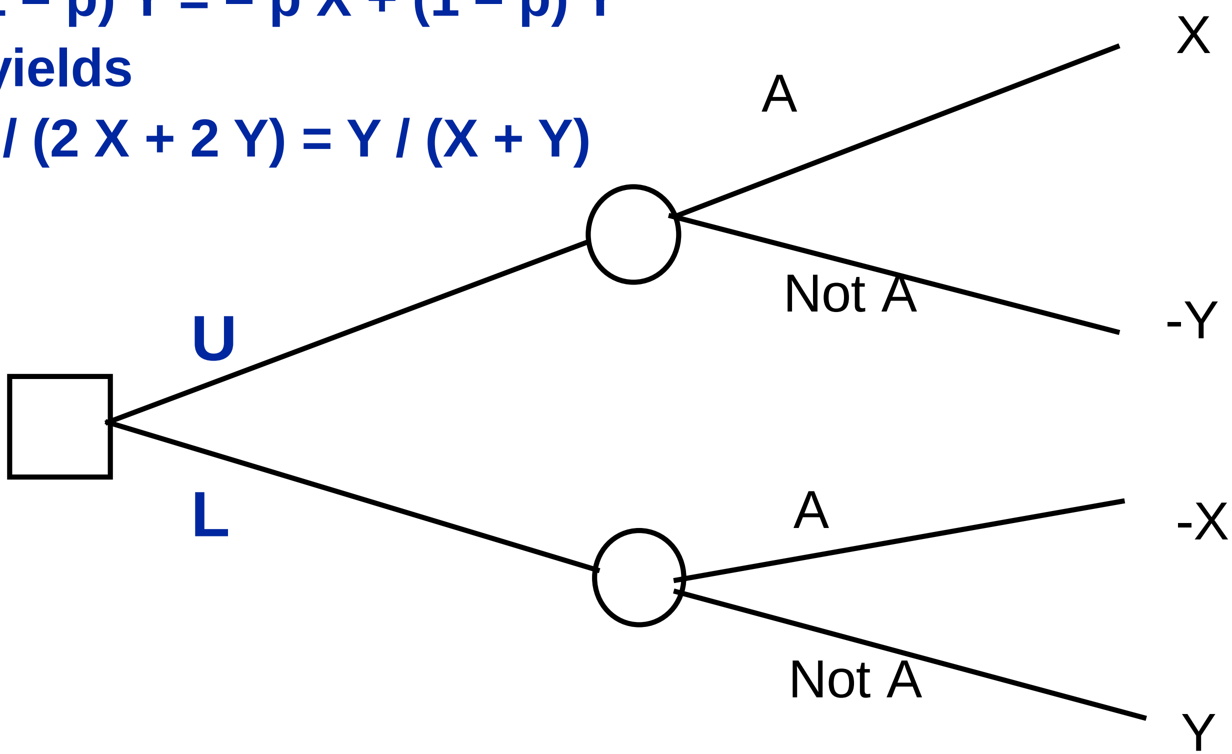

Elicitation of probabilities: Symmetric bets

What is the probability that it will rain tomorrow (in downtown Pittsburgh)?

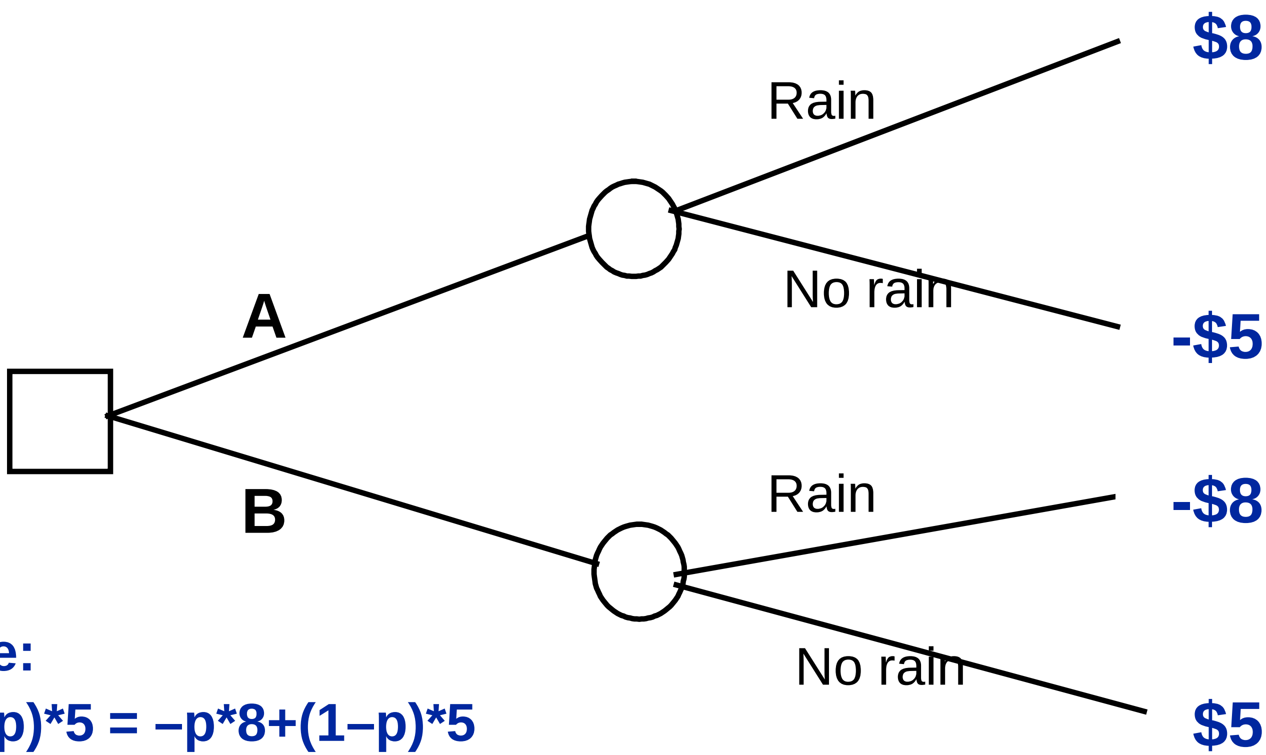

We have:

p*8–(1–p)*5 = –p*8+(1–p)*5

which yields

p=2*5/(2*8+2*5)=5/(8+5)=5/13=0.38

Expert choice: indifferent

Other helpful tools

Elicitation of probabilities: Continuous distributions

- Use methods for elicitation of discrete probabilities but conduct a series of elicitations.
- Reduces each step of elicitation to $P(A \leq a_0)$, where a_0 varies.
- Fit the CDF curve to the elicited points.

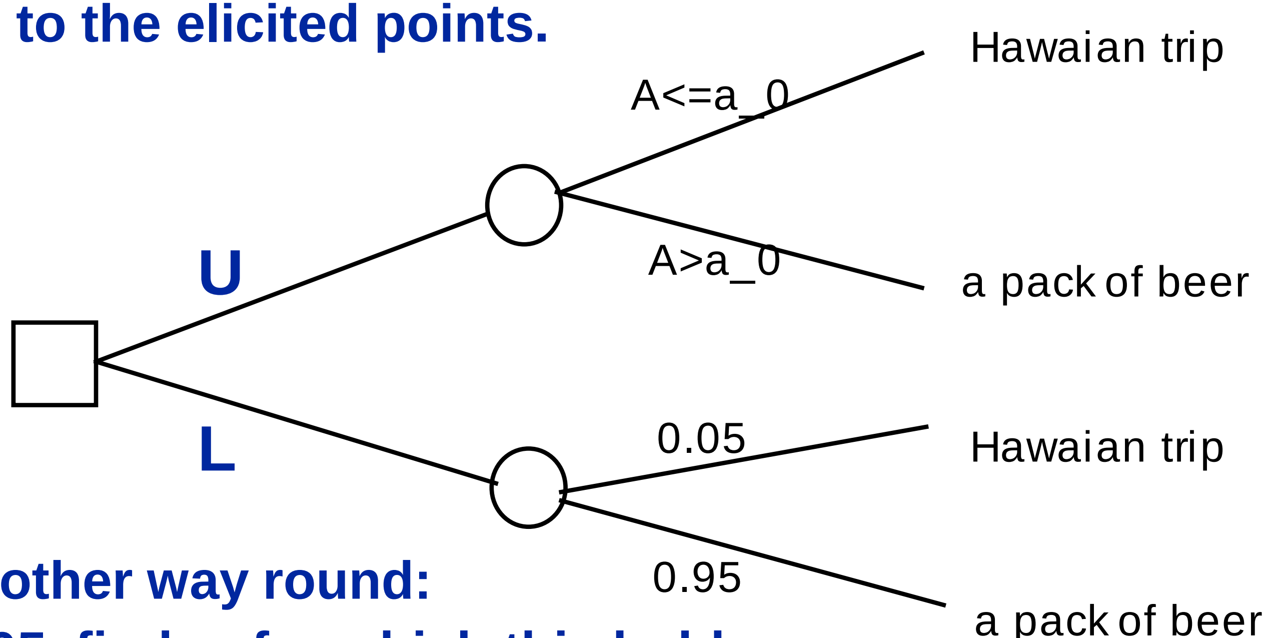

Possible to do the other way round:

- Given $P(A \leq a_0)=0.05$, find a_0 for which this holds.
- Manipulate a_0 until the expert is indifferent between the two options.
- Use the following fractiles: 0.05, 0.95, 0.25, 0.75. 0.5.

Elicitation of probabilities: Metalog Distribution

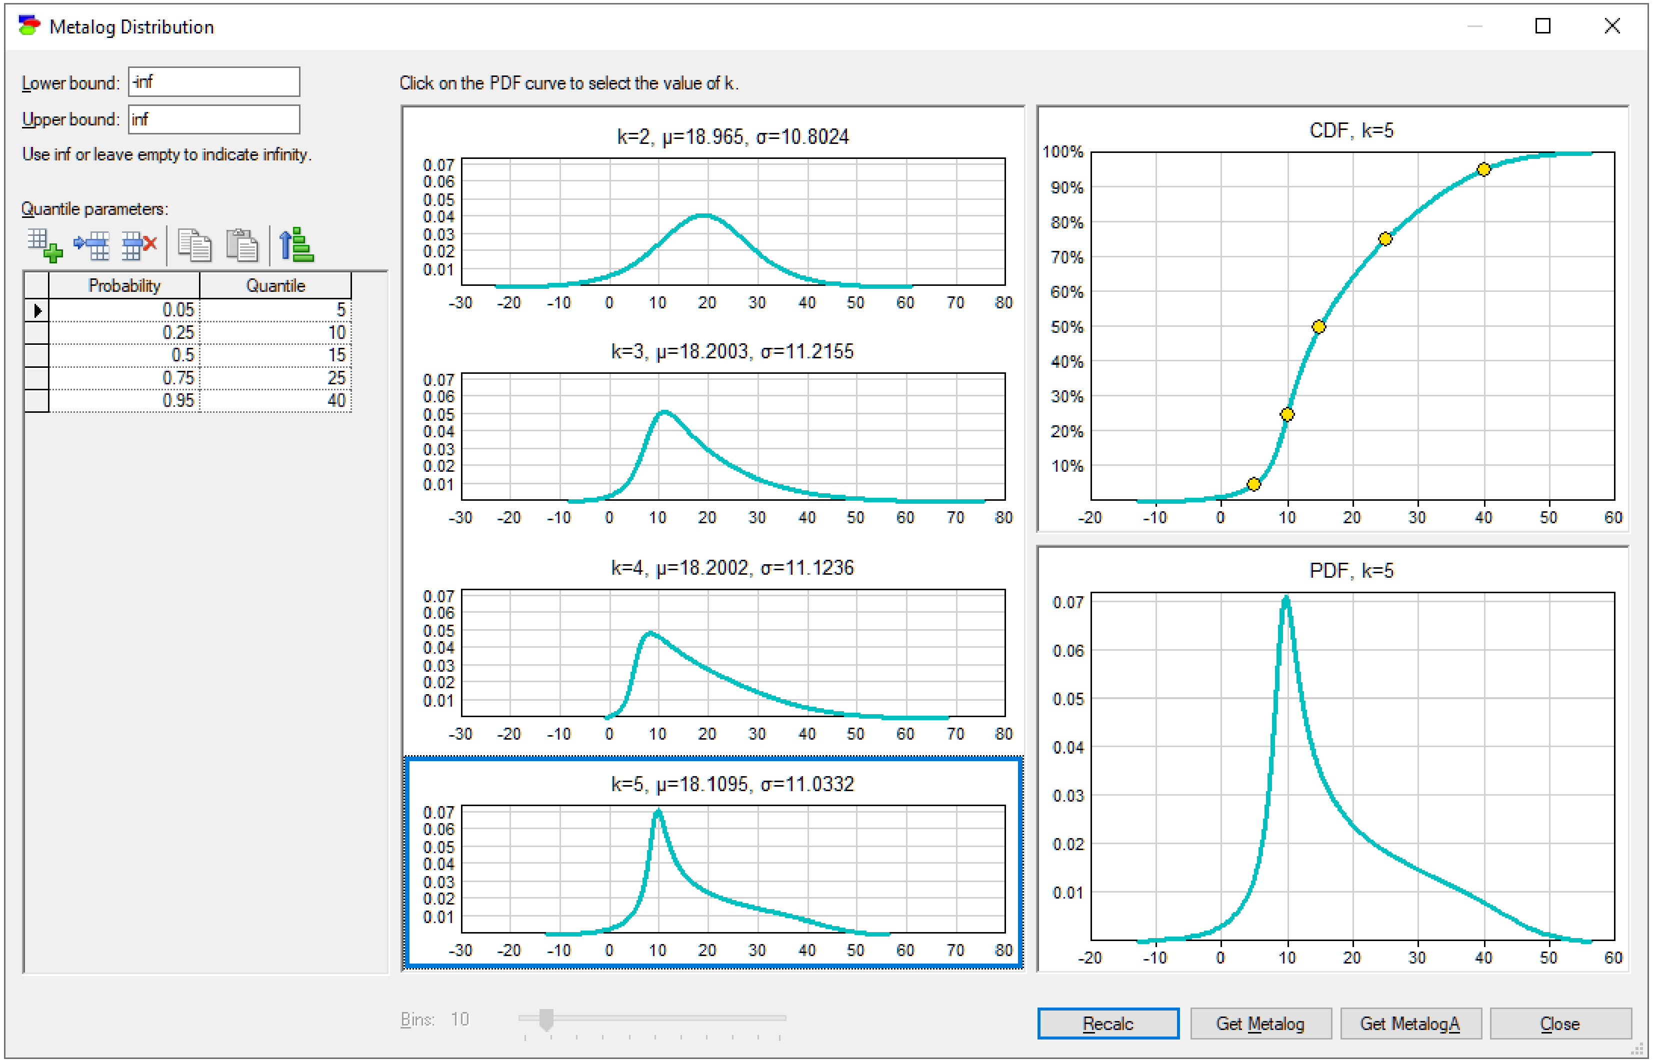

Elicitation of probabilities: Equations

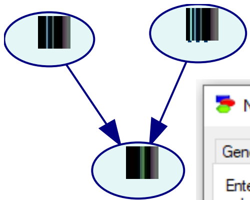

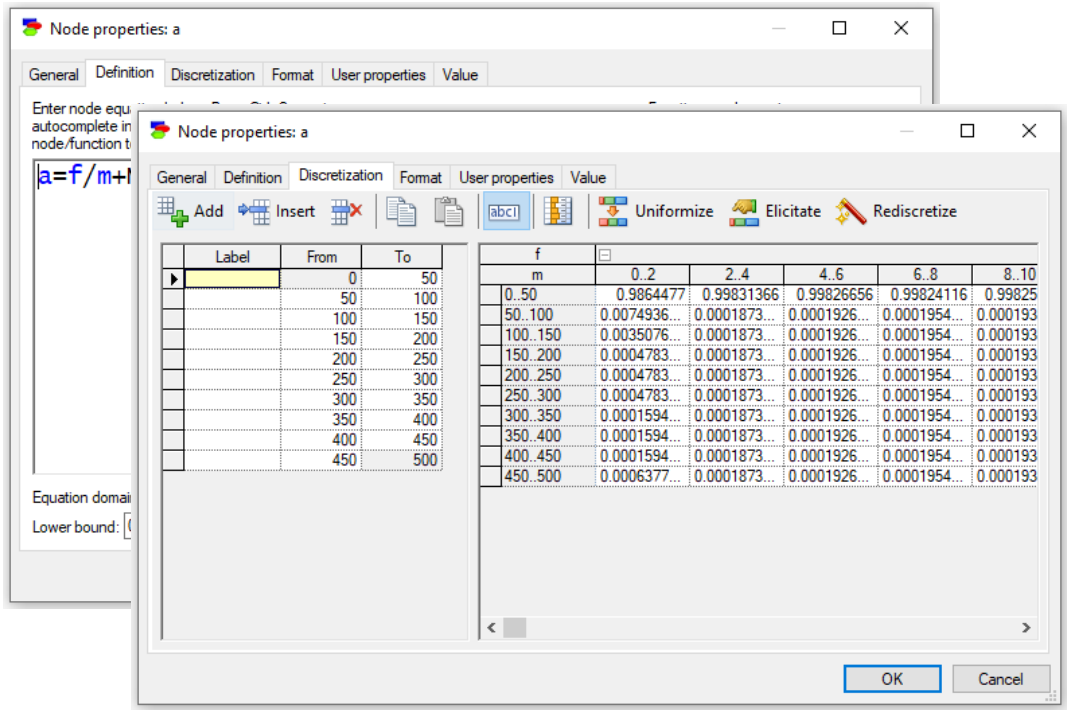

Label	From	To
	0	50
	50	100
	100	150
	150	200
	200	250
	250	300
	300	350
	350	400
	400	450
	450	500

f / m	0..2	2..4	4..6	6..8	8..10
0..50	0.9864477	0.99831366	0.99826656	0.99824116	0.99825
50..100	0.0074936...	0.0001873...	0.0001926...	0.0001954...	0.000193
100..150	0.0035076...	0.0001873...	0.0001926...	0.0001954...	0.000193
150..200	0.0004783...	0.0001873...	0.0001926...	0.0001954...	0.000193
200..250	0.0004783...	0.0001873...	0.0001926...	0.0001954...	0.000193
250..300	0.0004783...	0.0001873...	0.0001926...	0.0001954...	0.000193
300..350	0.0001594...	0.0001873...	0.0001926...	0.0001954...	0.000193
350..400	0.0001594...	0.0001873...	0.0001926...	0.0001954...	0.000193
400..450	0.0001594...	0.0001873...	0.0001926...	0.0001954...	0.000193
450..500	0.0006377...	0.0001873...	0.0001926...	0.0001954...	0.000193

Whenever you know enough about a domain to write the interactions in form of an equation, derive the CPTs by discretizing variables

Elicitation of probabilities: Discretization of continuous distributions

Two methods of discretizing continuous distributions:

(1) Extended Pearson and Tuckey:

3 point approximation: 0.05, 0.5, 0.95

Assign them p=0.185, 0.63, 0.185

(2) Bracket medians:

Split the range into intervals, assess the value that corresponds to probability that is median of each interval. Usually borders of intervals are 0.0, 0.2, 0.4, 0.6, 0.8, 1.0.

Elicitation of probabilities: Decomposition

Breaking the assessment into manageable chunks.
The goal is to make the assessment easier (and more reliable!).
Sometimes it is easier to introduce another variable.

For example, instead of assessing P(quadriplegic), i.e., probability that the decision maker becomes quadriplegic, we assess P(quadriplegic|*) P(*), where * are various ways of becoming quadriplegic, e.g., a car accident.

(1) Think how the event in question is related to other events (e.g., P(stock price up | market up)
(2) Think what kinds of alternative uncertain events could eventually lead to the event in question
(3) Think through all of different events that must happen before the event in question occurs.

Learning probabilities from data

Learning probabilities from data: Discrete case

Essentially counting occurrences

Credit20.csv

PaymentHistory	WorkHistory	Reliability	Debit	Income	RatioDebInc	Assets	Worth	Profession	FutureIncome	Age	CreditWorthiness
Without_Reference	Unstable	Unreliable	a0_11100	s70001_more	Favorable	wealthy	High	Medium_income_profession	Promising	a16_21	Negative
Aceptable	Unjustified_no_work	Unreliable	a0_11100	s70001_more	Favorable	average	High	Medium_income_profession	Promising	a66_up	Negative
Aceptable	Unstable	Reliable	a25901_more	s30001_70000	Unfavorable	wealthy	High	Low_income_profession	Not_promising	a16_21	Negative
Excellent	Unstable	Reliable	a25901_more	s30001_70000	Unfavorable	average	Medium	Medium_income_profession	Not_promising	a16_21	Negative
Excellent	Unjustified_no_work	Unreliable	a11101_25900	s0_30000	Unfavorable	average	Low	Medium_income_profession	Not_promising	a66_up	Negative
Without_Reference	Stable	Reliable	a0_11100	s30001_70000	Favorable	average	High	Medium_income_profession	Promising	a16_21	Positive
NoAceptable	Stable	Unreliable	a0_11100	s70001_more	Favorable	wealthy	High	Medium_income_profession	Promising	a66_up	Positive
Excellent	Stable	Reliable	a0_11100	s70001_more	Favorable	wealthy	High	Low_income_profession	Promising	a66_up	Positive
Excellent	Stable	Reliable	a25901_more	s70001_more	Unfavorable	poor	High	Low_income_profession	Not_promising	a16_21	Negative
NoAceptable	Stable	Unreliable	a0_11100	s30001_70000	Favorable	average	Medium	Medium_income_profession	Promising	a22_65	Positive
Without_Reference	Justified_no_work	Reliable	a25901_more	s70001_more	Unfavorable	poor	High	Low_income_profession	Not_promising	a16_21	Negative
NoAceptable	Unstable	Unreliable	a25901_more	s30001_70000	Unfavorable	wealthy	High	Medium_income_profession	Promising	a16_21	Negative
NoAceptable	Justified_no_work	Unreliable	a25901_more	s30001_70000	Unfavorable	wealthy	High	High_income_profession	Promising	a22_65	Negative
Excellent	Stable	Reliable	a11101_25900	s0_30000	Unfavorable	average	Low	Medium_income_profession	Not_promising	a16_21	Negative
Aceptable	Stable	Unreliable	a25901_more	s0_30000	Unfavorable	wealthy	Medium	Medium_income_profession	Not_promising	a66_up	Negative
Without_Reference	Unjustified_no_work	Unreliable	a0_11100	s0_30000	Favorable	poor	Low	Low_income_profession	Not_promising	a66_up	Positive
Aceptable	Unstable	Reliable	a11101_25900	s30001_70000	Unfavorable	average	Medium	Low_income_profession	Not_promising	a66_up	Negative
Without_Reference	Unstable	Unreliable	a0_11100	s30001_70000	Unfavorable	average	High	Medium_income_profession	Promising	a16_21	Negative
Without_Reference	Stable	Reliable	a25901_more	s70001_more	Favorable	poor	High	High_income_profession	Promising	a16_21	Positive
Excellent	Justified_no_work	Reliable	a25901_more	s70001_more	Unfavorable	wealthy	High	High_income_profession	Promising	a66_up	Positive

Row 20 of 20

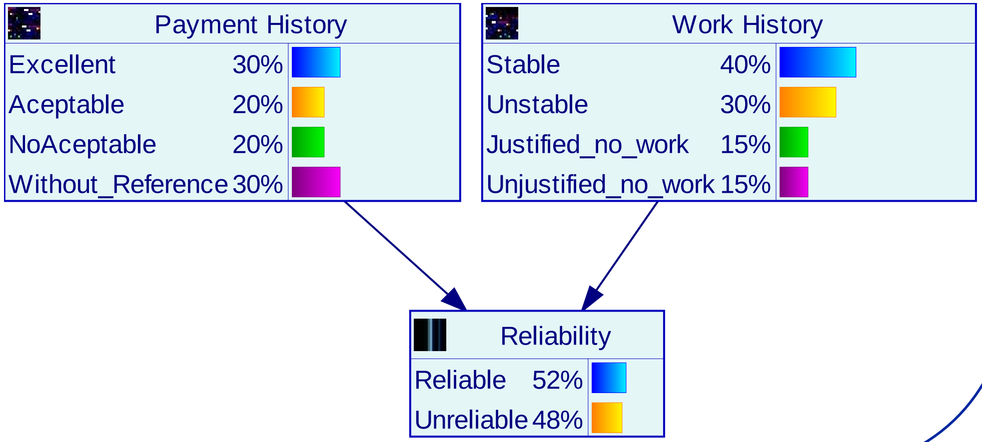

Learning probabilities from data: Discrete case

Prior (unconditional) probabilities

Credit20.csv

Payment History	Work History	Reliability
Without_Reference	Unstable	Unreliable
Aceptable	Unjustified_no_work	Unreliable
Aceptable	Unstable	Reliable
Excellent	Unstable	Reliable
Excellent	Unjustified_no_work	Unreliable
Without_Reference	Stable	Reliable
NoAceptable	Stable	Unreliable
Excellent	Stable	Reliable
Excellent	Stable	Reliable
NoAceptable	Stable	Unreliable
Without_Reference	Justified_no_work	Reliable
NoAceptable	Unstable	Unreliable
NoAceptable	Justified_no_work	Unreliable
Excellent	Stable	Reliable
Aceptable	Stable	Unreliable
Without_Reference	Unjustified_no_work	Unreliable
Aceptable	Unstable	Reliable
Without_Reference	Unstable	Unreliable
Without_Reference	Stable	Reliable
Excellent	Justified_no_work	Reliable

Row 1 of 20

Excellent	0.3
Aceptable	0.2
NoAceptable	0.2
Without_Reference	0.3

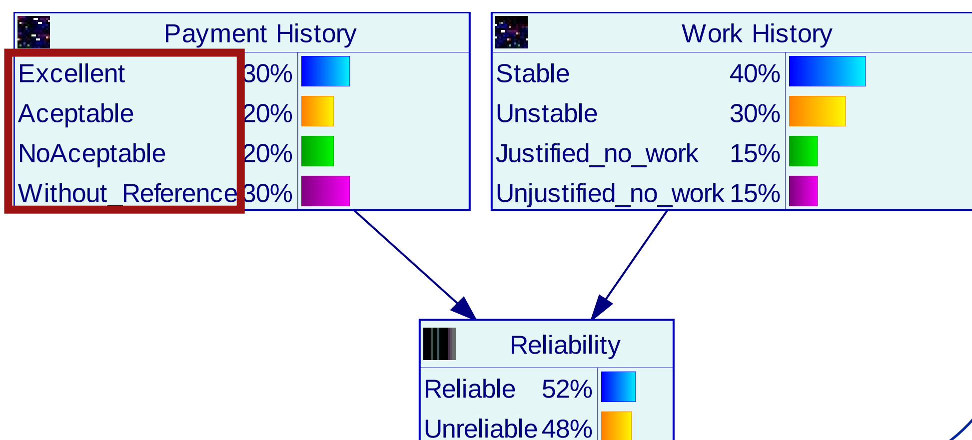

Learning probabilities from data: Discrete case

Prior (unconditional) probabilities

Credit20.csv

	PaymentHistory	WorkHistory	Reliability
▶	Aceptable	Stable	Unreliable
	Aceptable	Unjustified_no_work	Unreliable
	Aceptable	Unstable	Reliable
	Aceptable	Unstable	Reliable
	Excellent	Justified_no_work	Reliable
	Excellent	Stable	Reliable
	Excellent	Stable	Reliable
	Excellent	Stable	Reliable
	Excellent	Unjustified_no_work	Unreliable
	Excellent	Unstable	Reliable
	NoAceptable	Justified_no_work	Unreliable
	NoAceptable	Stable	Unreliable
	NoAceptable	Stable	Unreliable
	NoAceptable	Unstable	Unreliable
	Without_Reference	Justified_no_work	Reliable
	Without_Reference	Stable	Reliable
	Without_Reference	Stable	Reliable
	Without_Reference	Unjustified_no_work	Unreliable
	Without_Reference	Unstable	Unreliable
	Without_Reference	Unstable	Unreliable

Row 1 of 20

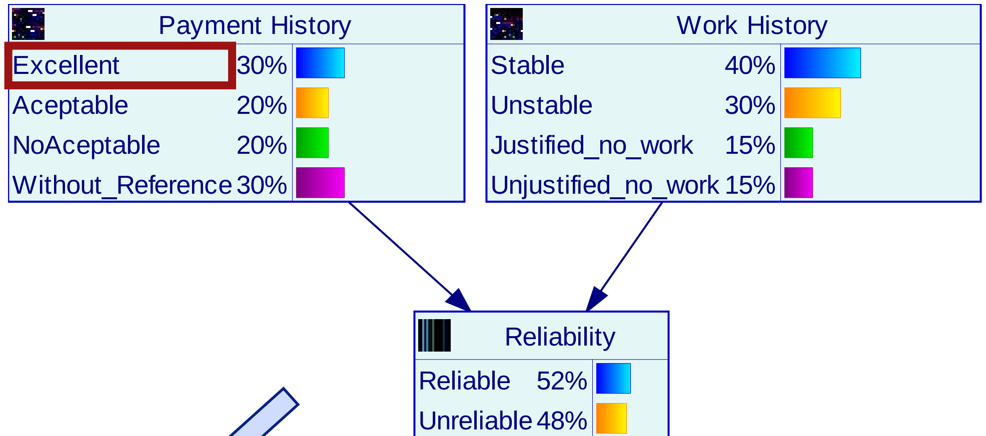

Payment History	Excellent				Aceptable				NoAceptable				Without_Reference			
Work History	Stable	Unstable	Justified...	Unjustified...	Stable	Unstable	Justified...	Unjustified...	Stable	Unstable	Justified...	Unjustified...	Stable	Unstable	Justified...	Unjustified...
Reliable	0.875	0.75	0.75	0.25	0.25	0.83333333	0.5	0.25	0.16666667	0.25	0.25	0.5	0.83333333	0.16666667	0.75	0.25
Unreliable	0.125	0.25	0.25	0.75	0.75	0.16666667	0.5	0.75	0.83333333	0.75	0.75	0.5	0.16666667	0.83333333	0.25	0.75

Learning probabilities from data: Continuous case

Essentially fitting a distribution to a histogram

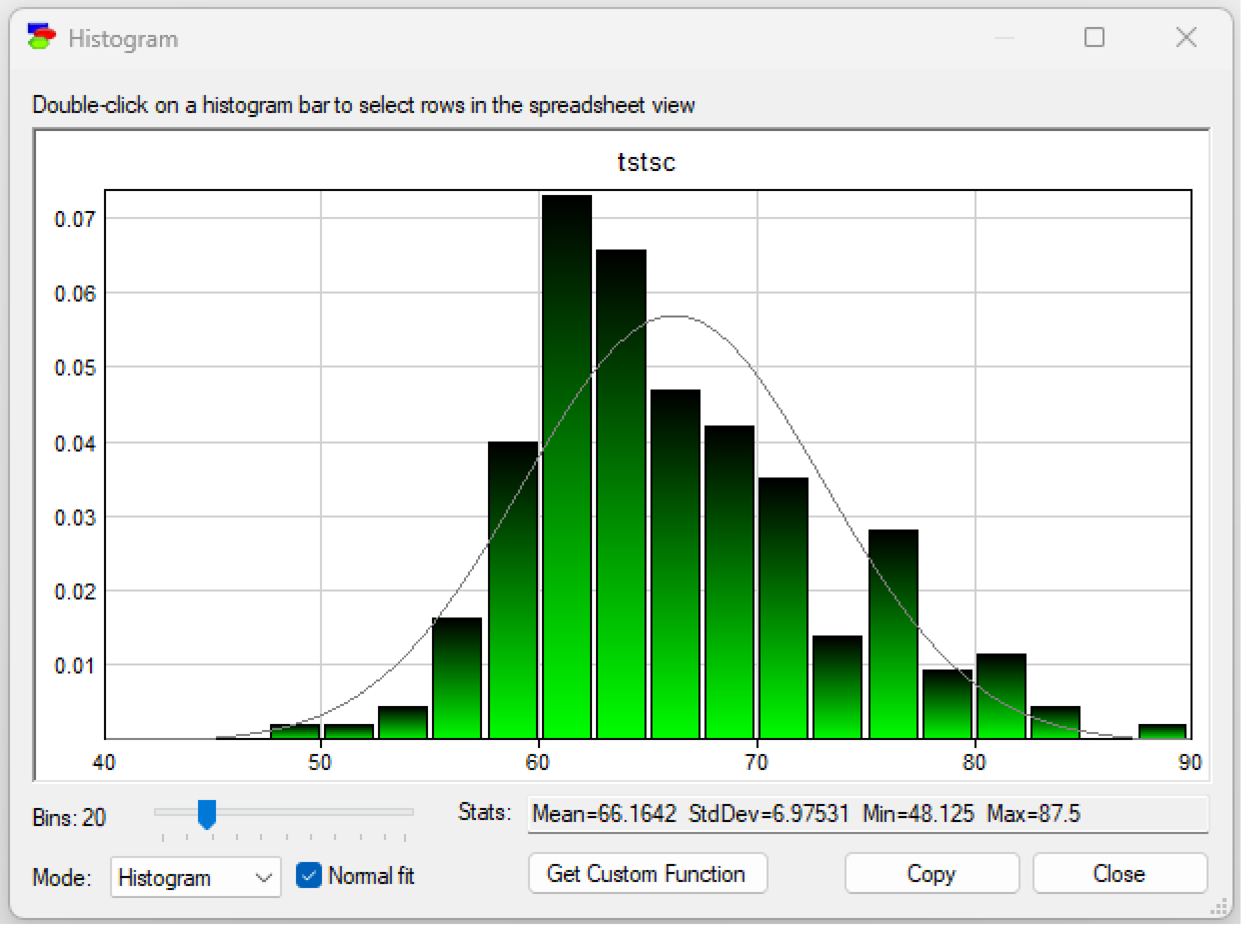

Learning probabilities from data: Continuous case

Essentially fitting a distribution to a histogram

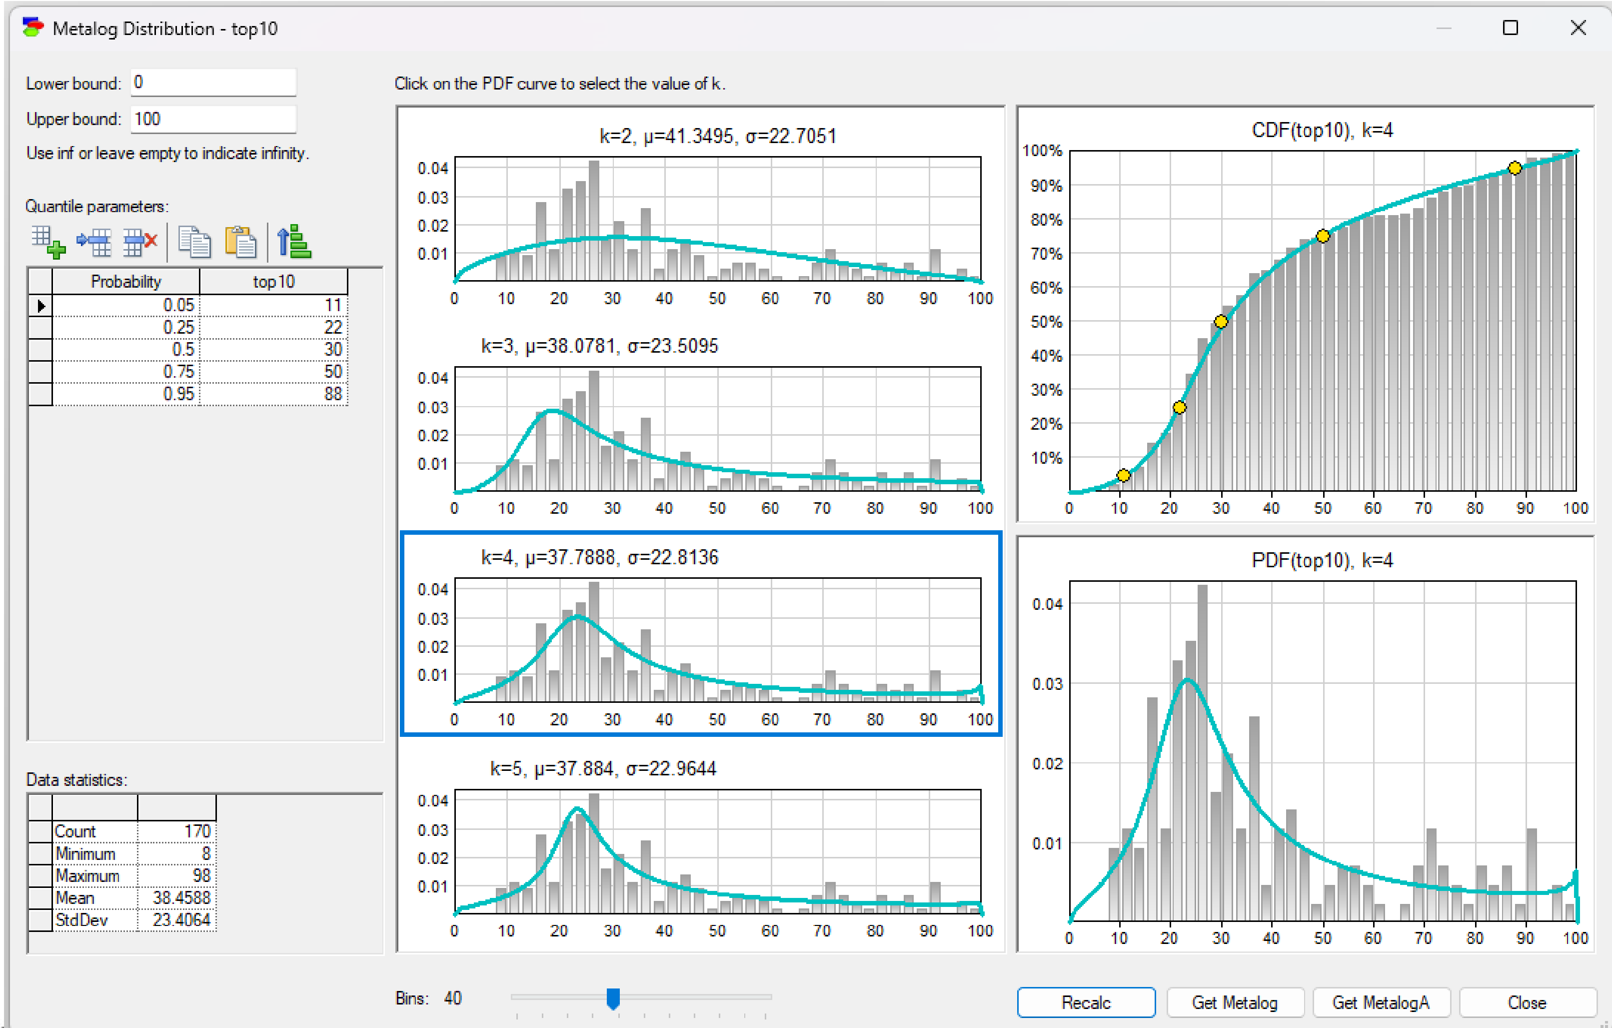

Do parameters matter?

Random noise, Normal(0,σ)

transformed parameters

σ = 0.3

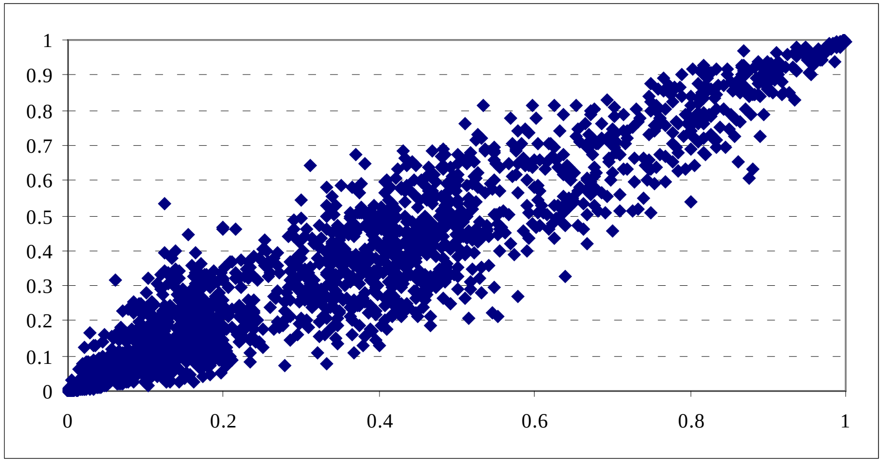

σ = 1.0

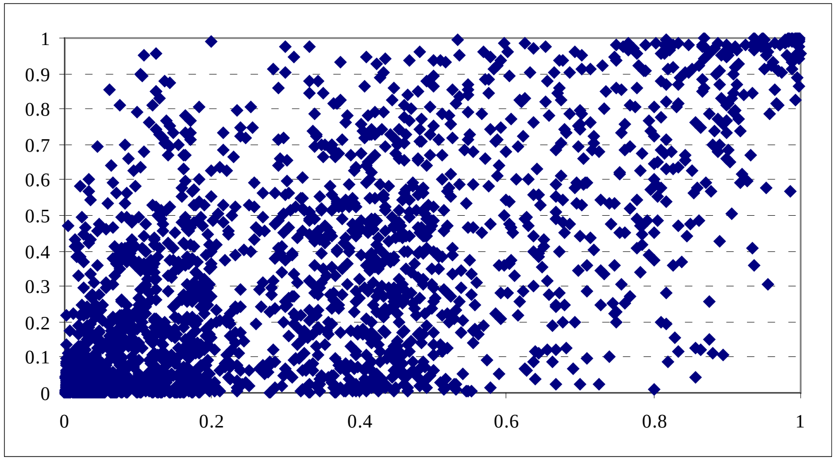

σ = 2.0

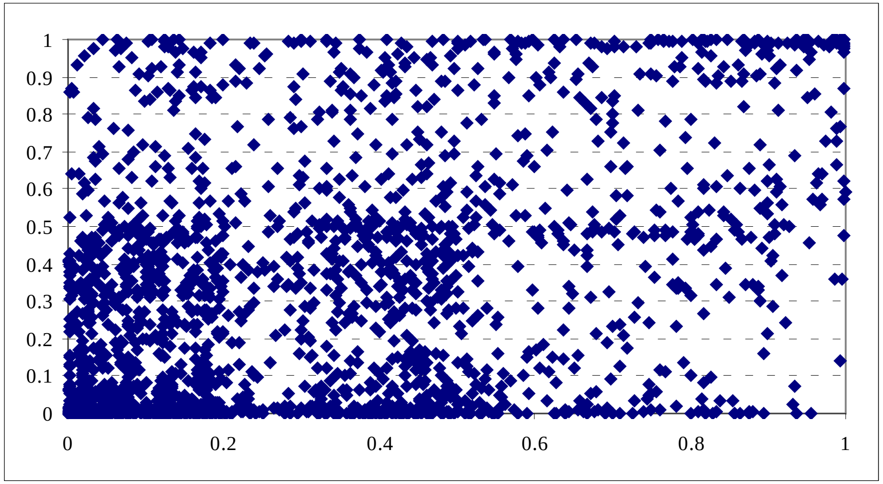

σ = 3.0

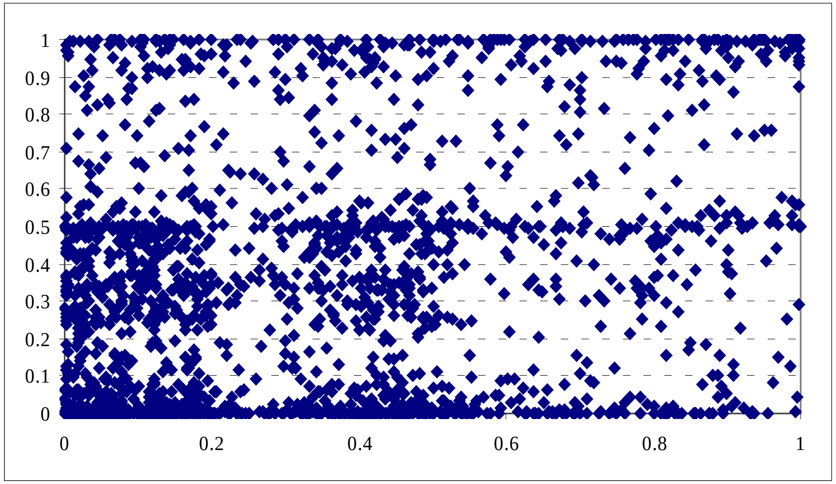

original parameters

[Onisko & Druzdzel 2011]

Biased noise (overconfidence), Normal(0,σ) added to the largest probability in a distribution

transformed parameters

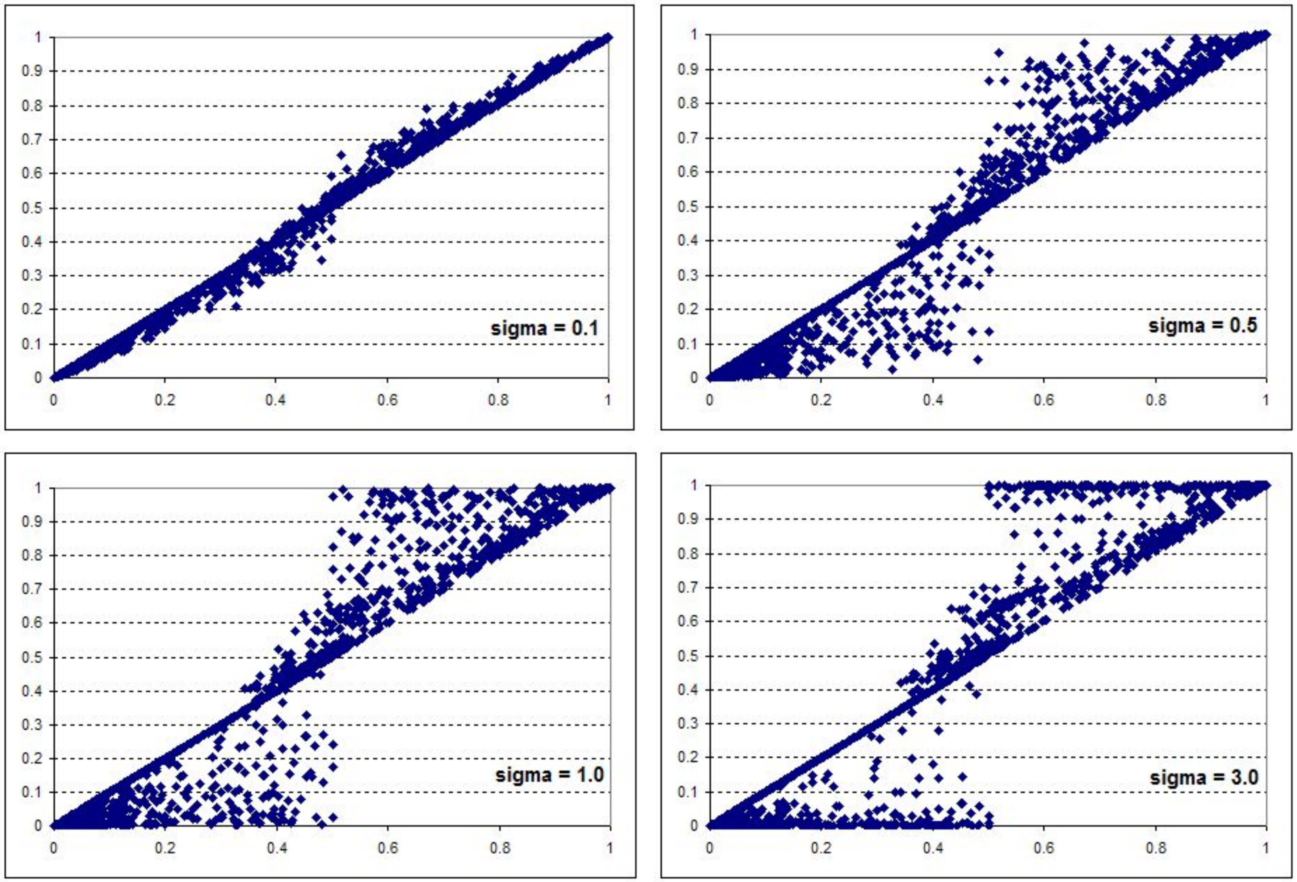

original parameters

Biased noise (underconfidence), Normal(0,σ) subtracted from the largest probability in a distribution

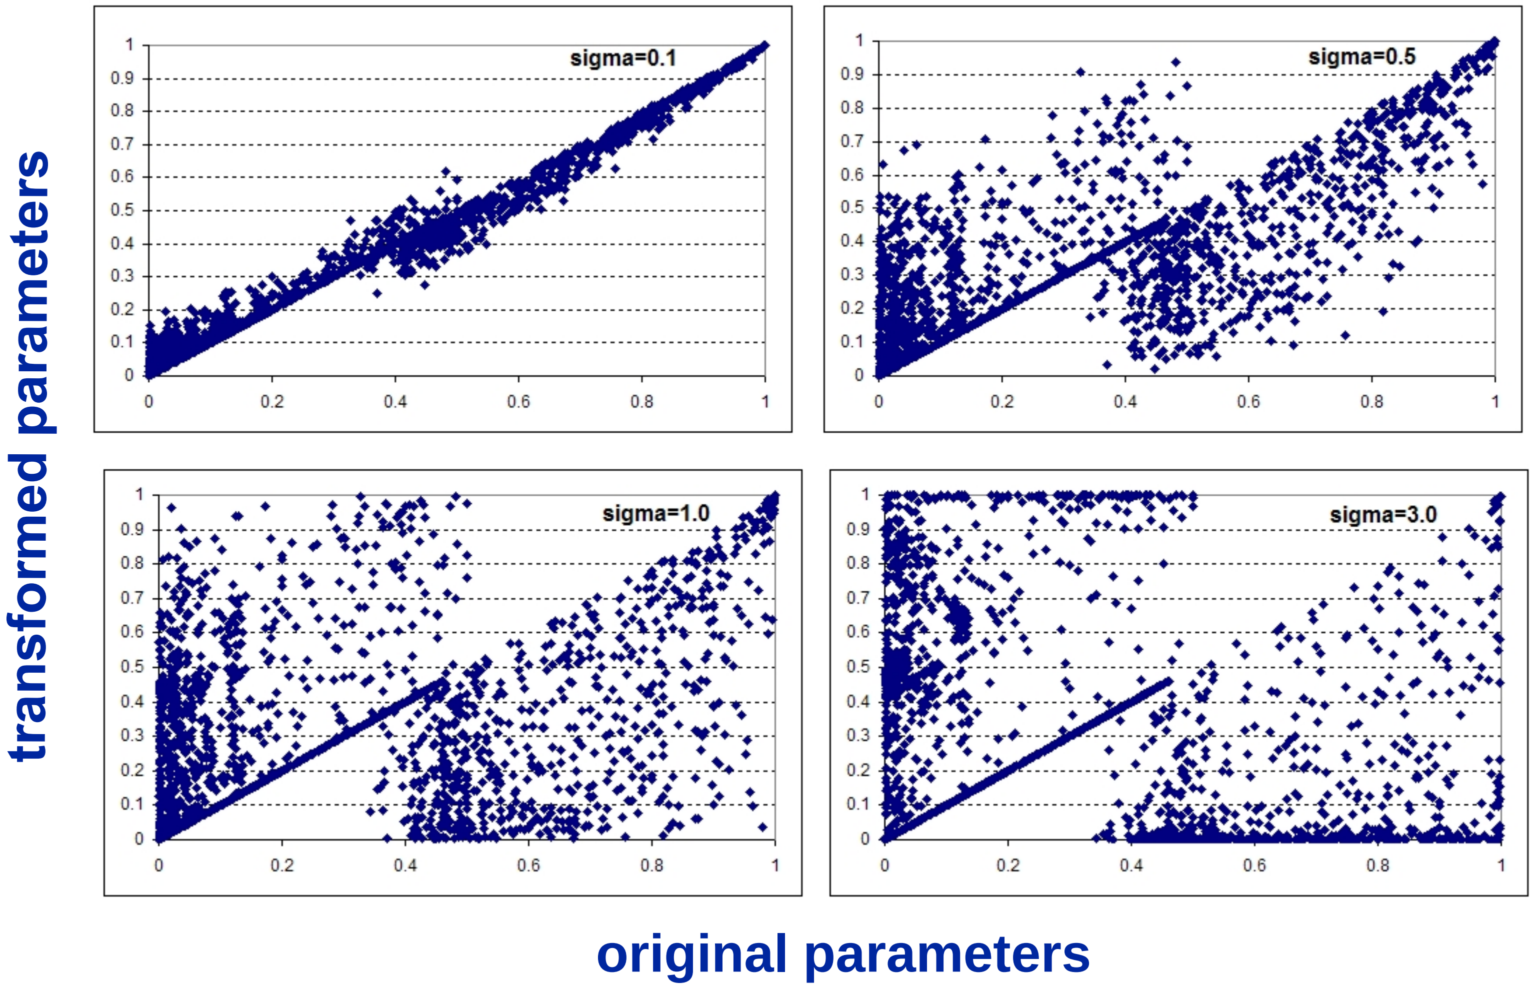

Diagnostic performance as a function of parameter accuracy

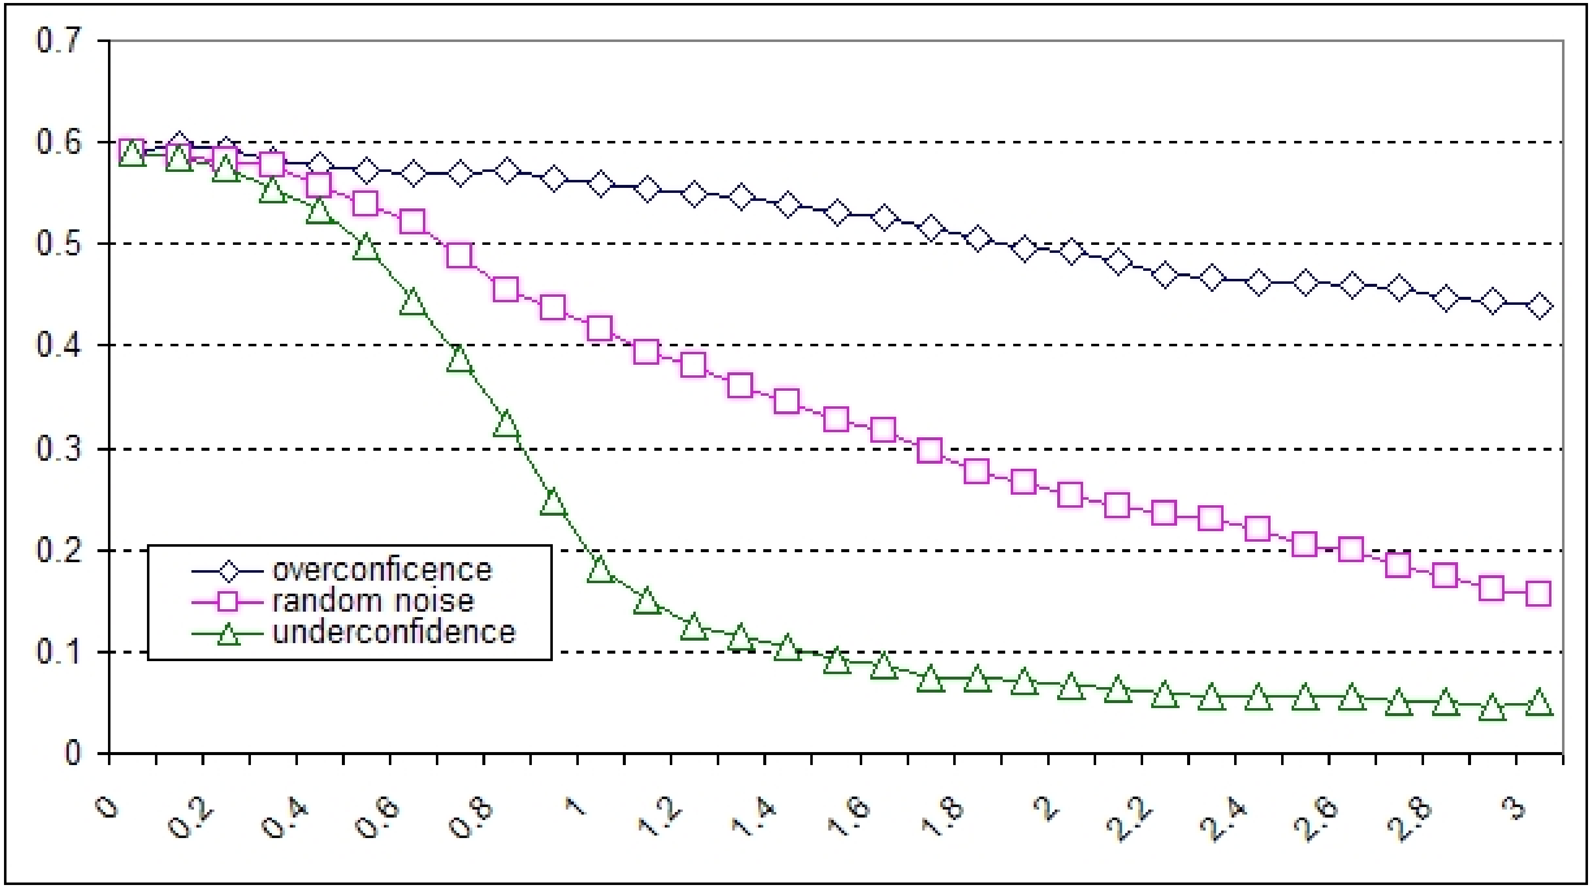

Rounded vs. original probabilities for various levels of rounding accuracy

transformed parameters

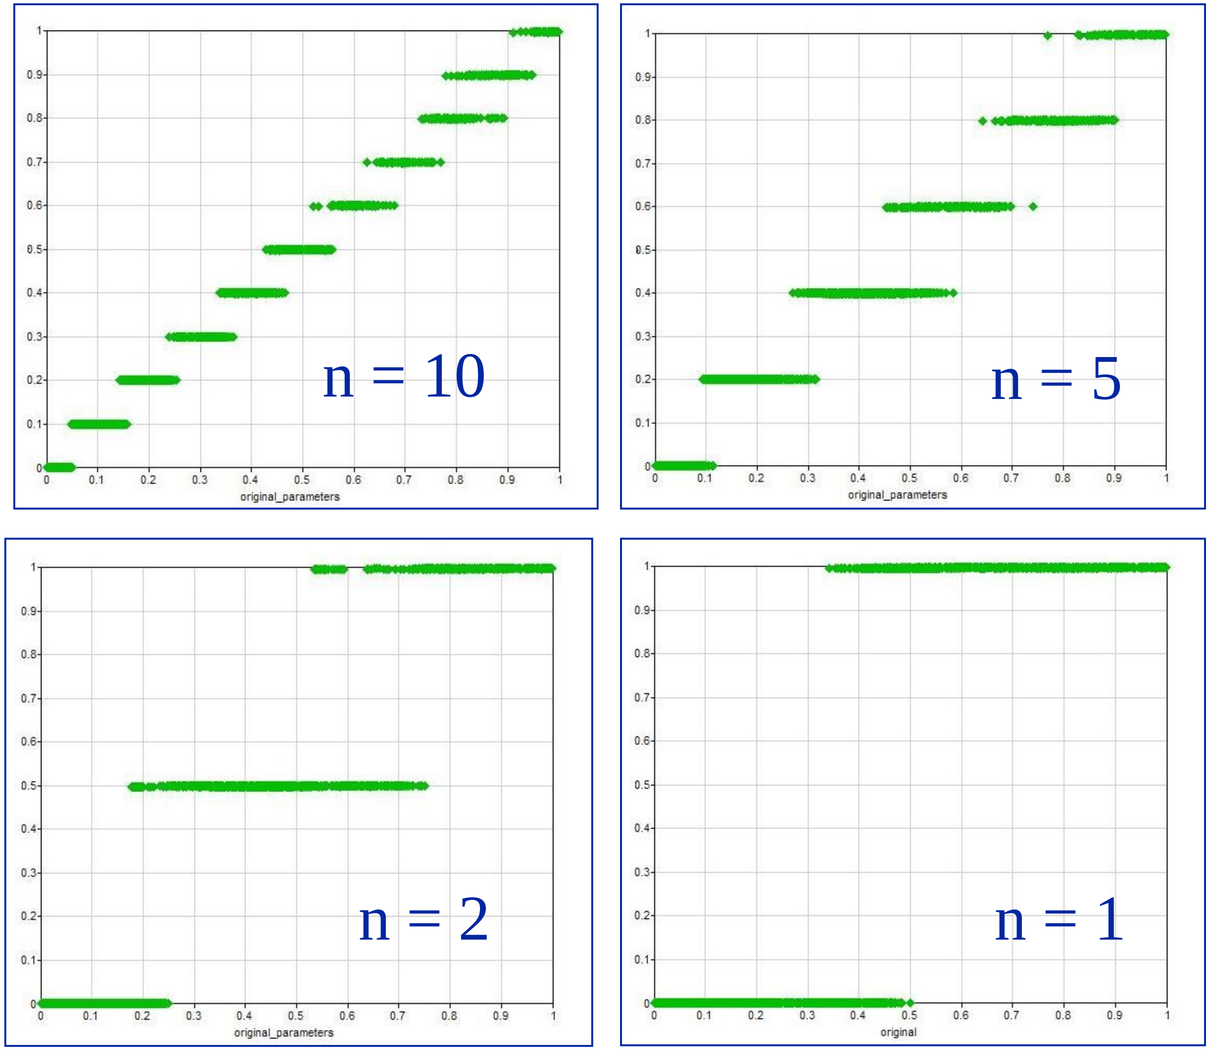

original parameters

Histograms of original and rounded probabilities for various levels of rounding accuracy

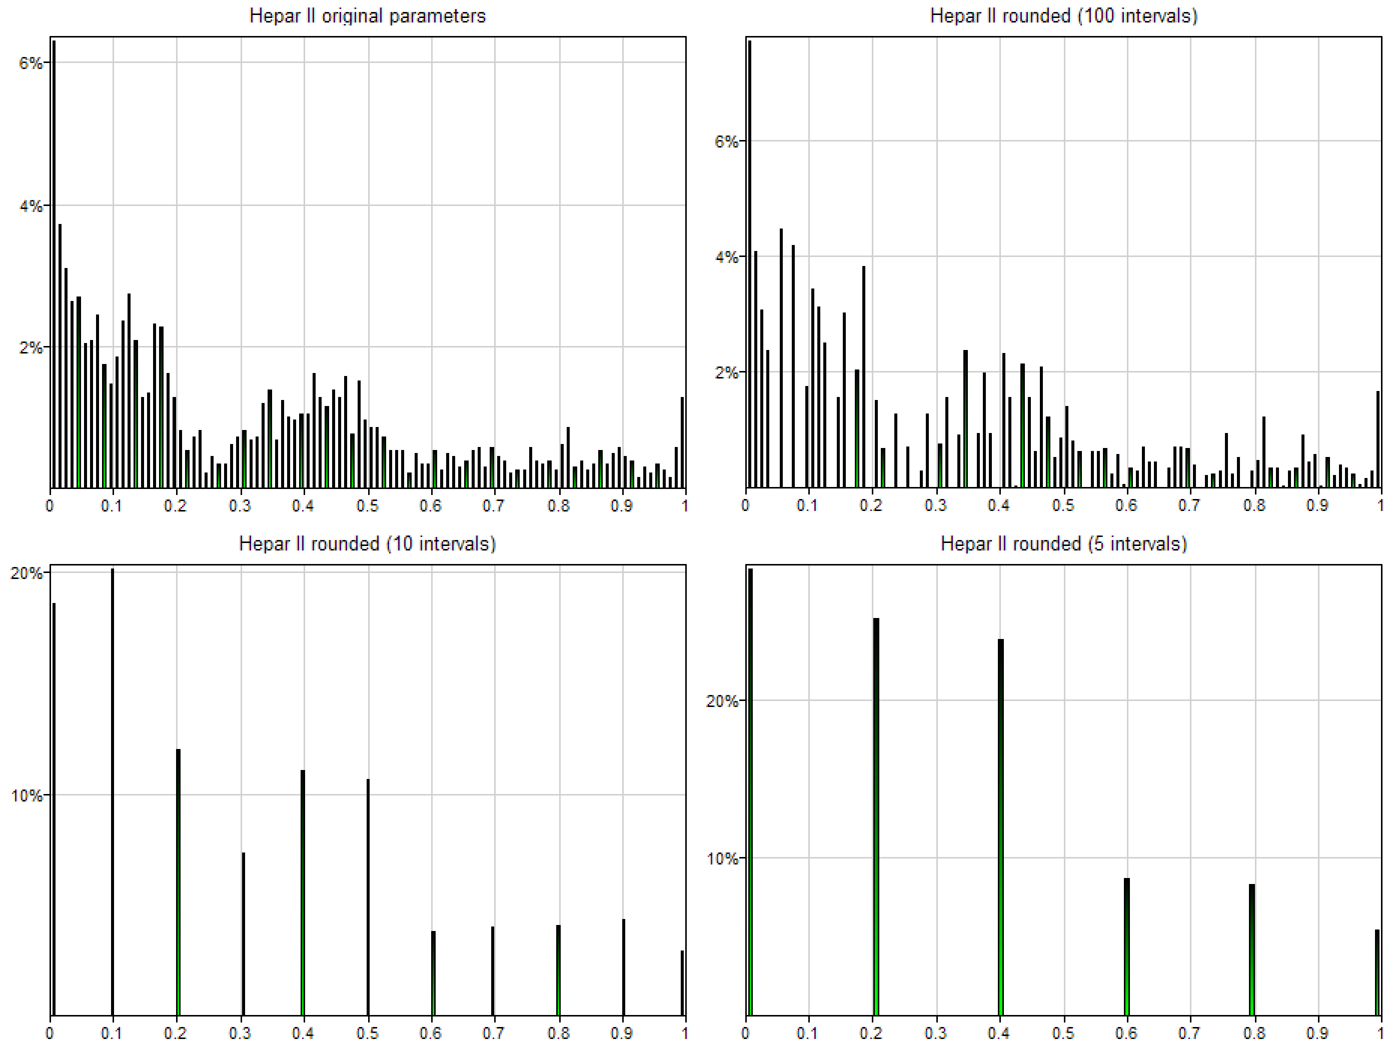

Diagnostic performance as a function of parameter accuracy (w=1)

P(A|B) = P(B|A) / P(B) P(A)

	acc	#zeros	%zeros
n=100	0.595	116	5%
n=10	0.295	400	19%
n=5	0.219	605	28%
n=4	0.230	754	35%
n=3	0.136	869	41%
n=2	0.100	1056	49%
n=1	0.161	1453	68%

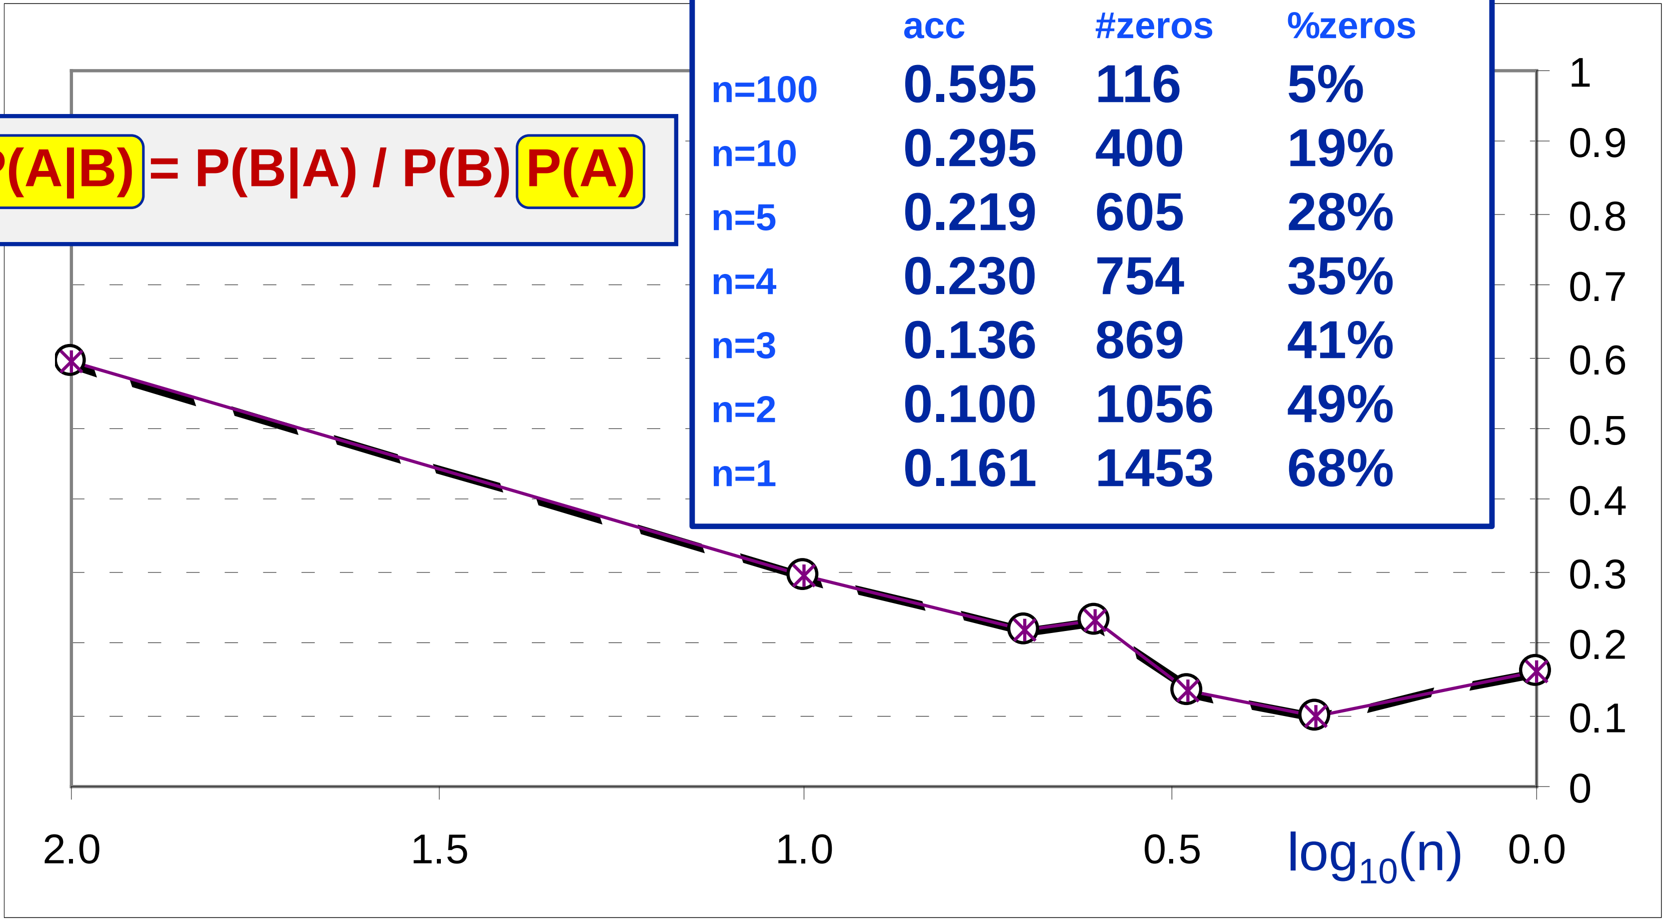

Diagnostic performance as a function of parameter accuracy and ε (w=1)

What if we replace all zeros by some small number ε ?

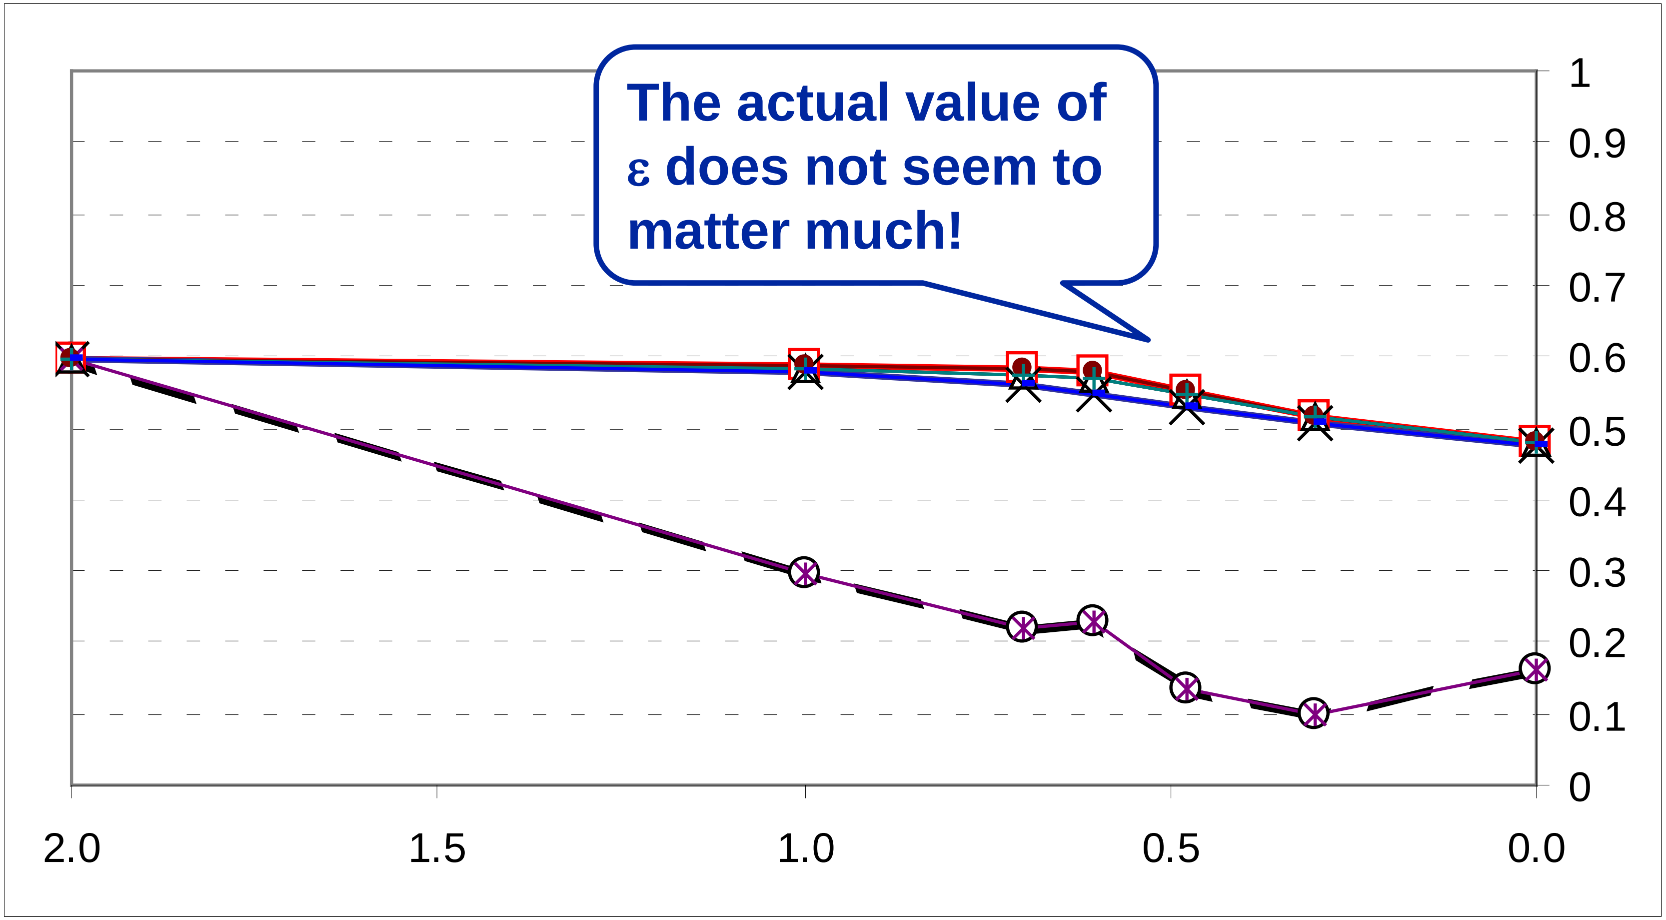

Other Relevant Issues

Is precision real or illusory?

- When getting the parameters from experts, we may well get better models when eliciting fewer parameters.
- When learning, the same may happen!

Sensitivity analysis in Bayesian networks

- **Given a target node (or a set of target nodes) and a possible set of evidence nodes, we can identify the parameters that matter most for those target(s)'s posteriors.**
- **We compute essentially the derivative of the posterior over each of the parameters.**

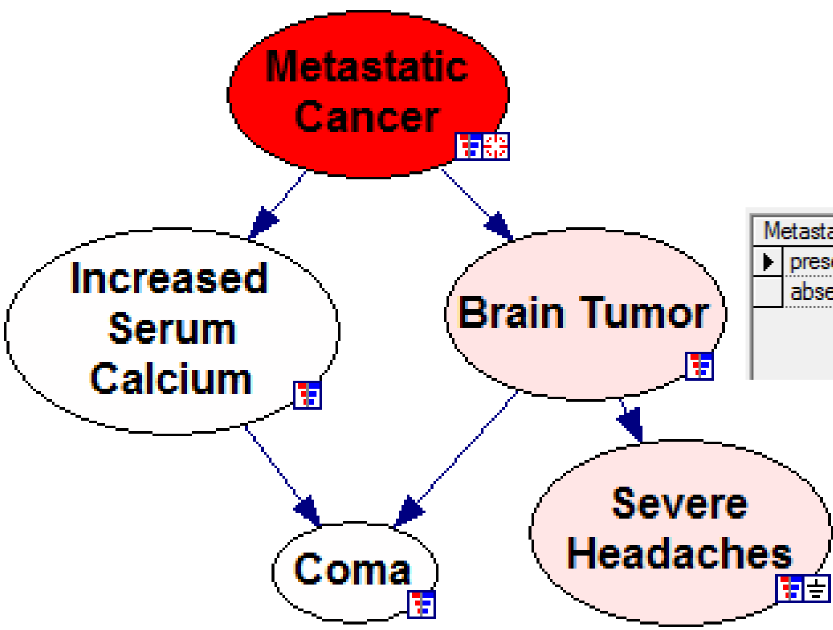

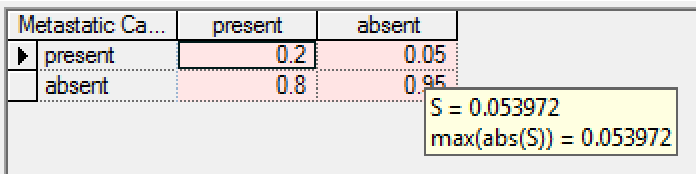

Metastatic Ca...	present	absent
present	0.2	0.05
absent	0.8	0.95

S = 0.053972
max(abs(S)) = 0.053972

Tornado diagrams: Sensitivity analysis in BNs

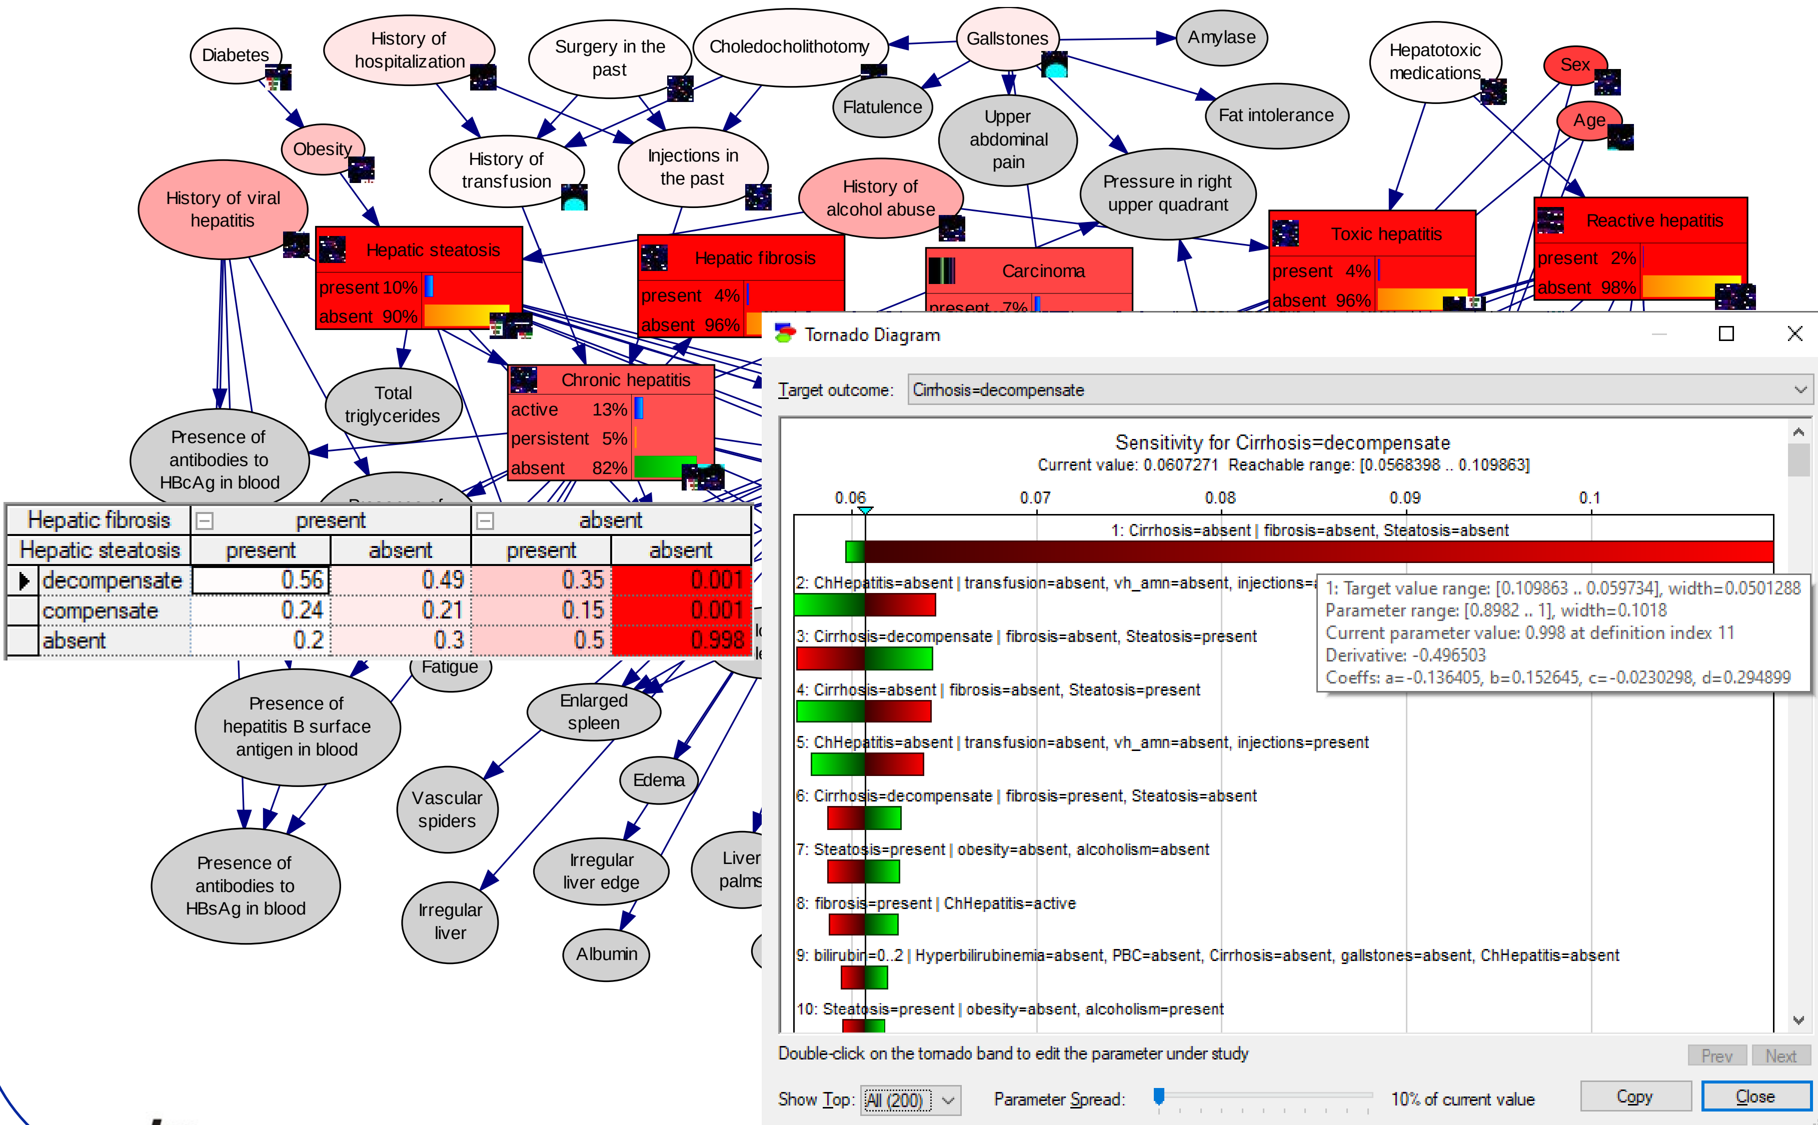

Sensitivity analysis in Ids (a preview ☺)

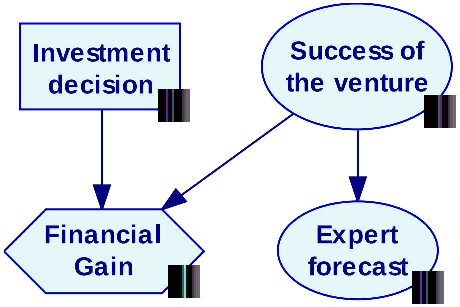

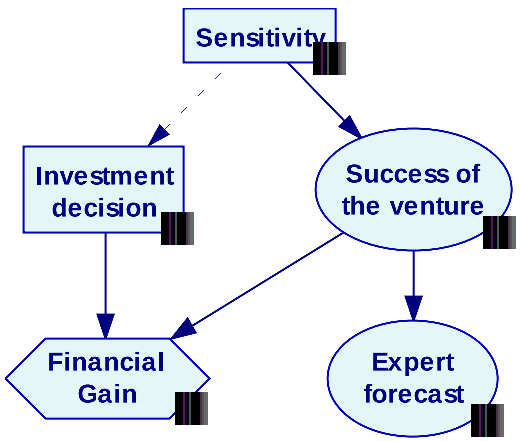

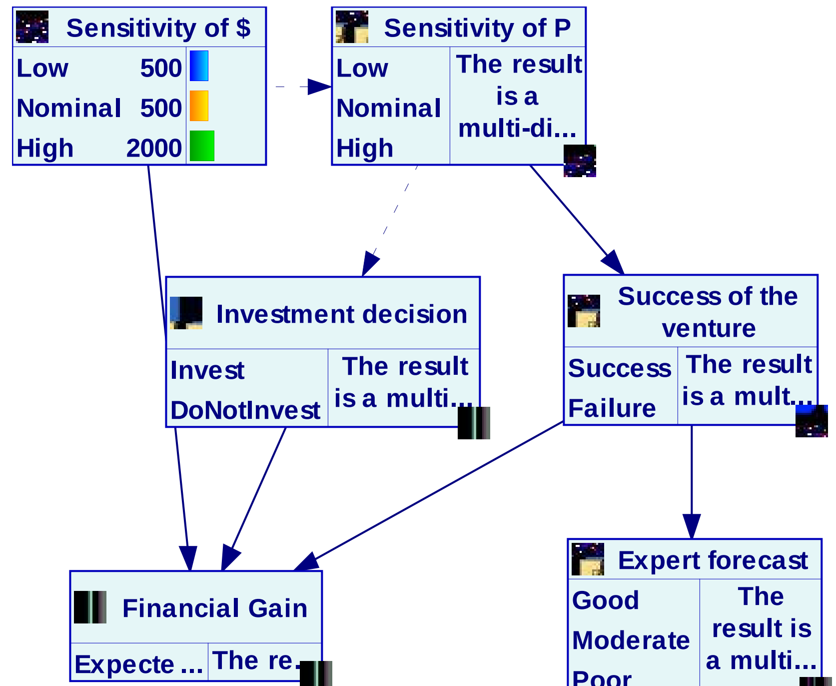

Strength of influence in Bayesian networks

- Because we have the entire joint probability distribution, we can compute individual strengths of influences between nodes
- It is essentially a measure of difference between conditional probability distributions in a CPT.

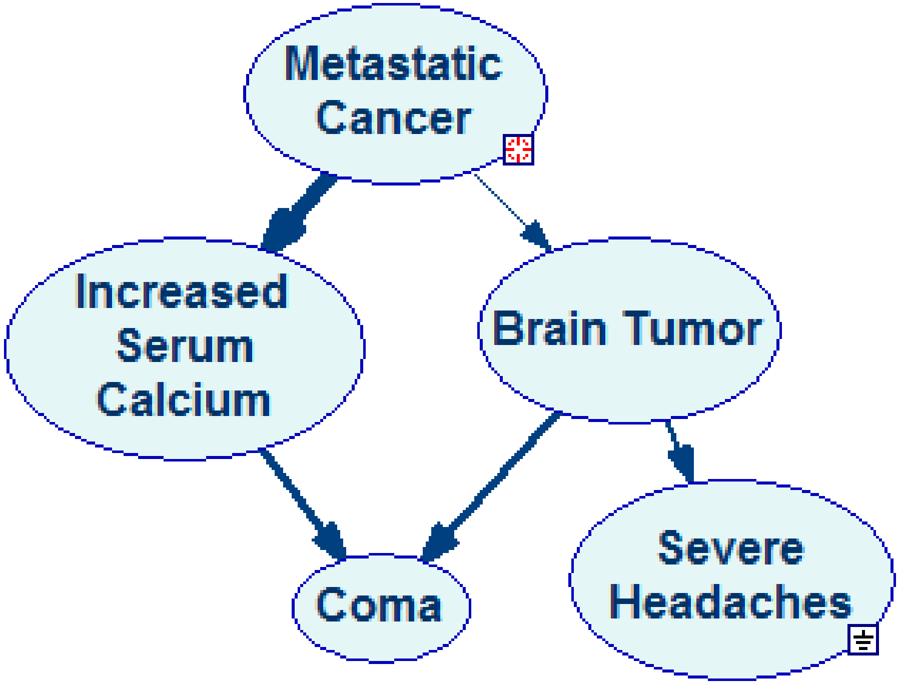

Value of information in Bayesian networks

- Because we have the entire joint probability distribution, we can compute the expected value of observations and rank-order them from the most to least informative.
- It is essentially expected cross-entropy between the targets and the individual observations.

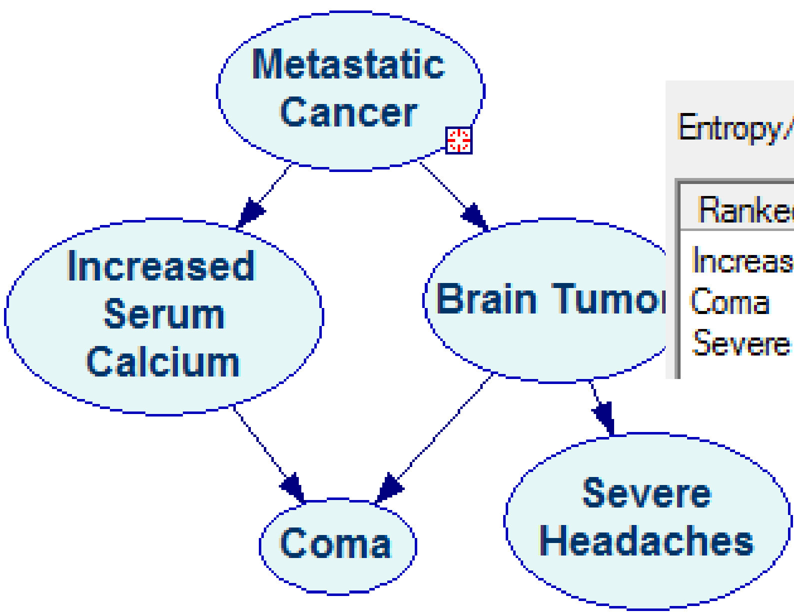

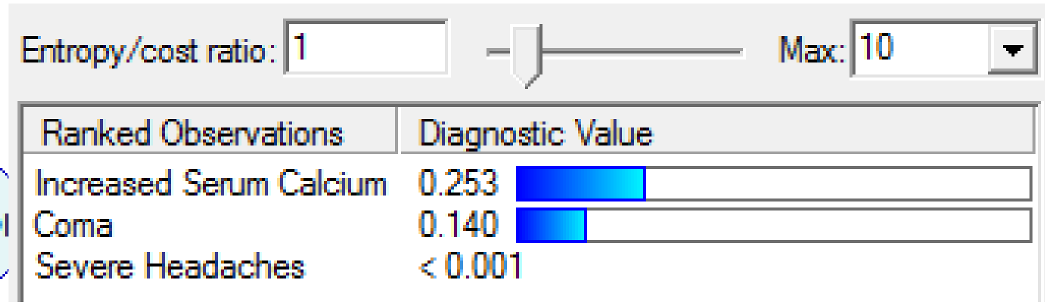

Clarity test

- "Gas price in 2026" vs. "average regular unleaded gas price taken over all gas stations within the city of Pittsburgh on January 1 2026".
- "Market up or down" vs. "the market goes up means that the Standard & Poor's 500 Index rises in 2026".
- The matter of clarifying definitions of alternatives, outcomes, and consequences is absolutely crucial in real-world decision problems. The clarity test forces us to define all aspects of a problem with great care.