Probability Elicitation

Marek J. Drużdżel

Wydział Informatyki

Politechnika Białostocka

m.druzdzel@pb.edu.pl

http://aragorn.wi.pb.bialystok.pl/~druzdzel/

Are probabilities in experts' heads?

William Estes' probability matching experiments

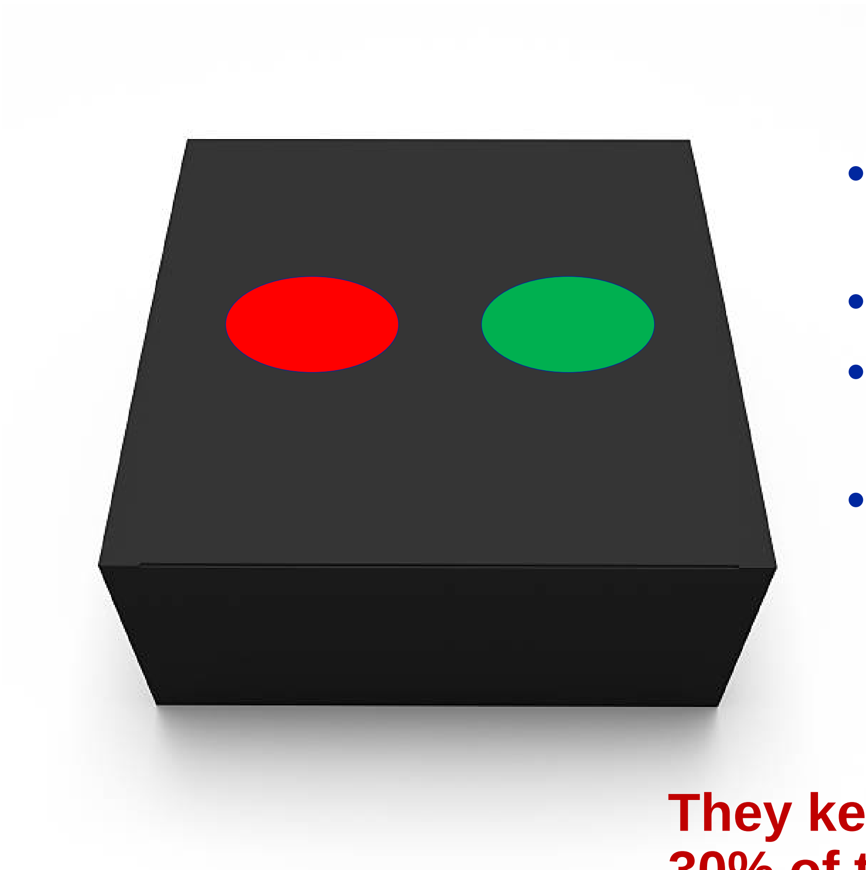

- Correct guessing of the light to be lit carries a small reward.
- Many repetitions.
- What is the optimal strategy if Pr(green)=0.3 and Pr(red)=0.7?
- What do human subjects do?

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

Elicitation of Probabilities

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

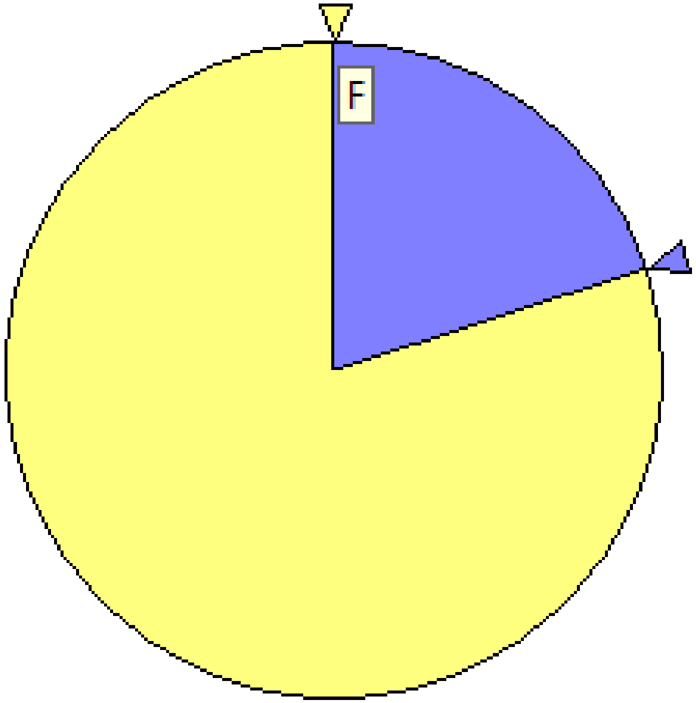

Bar chart

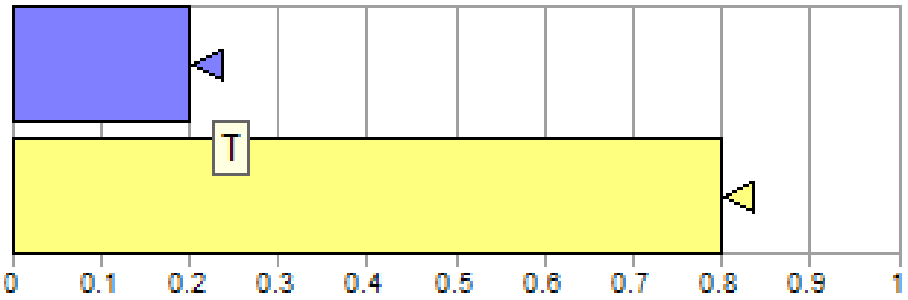

Elicitation of probabilities: Reference lottery

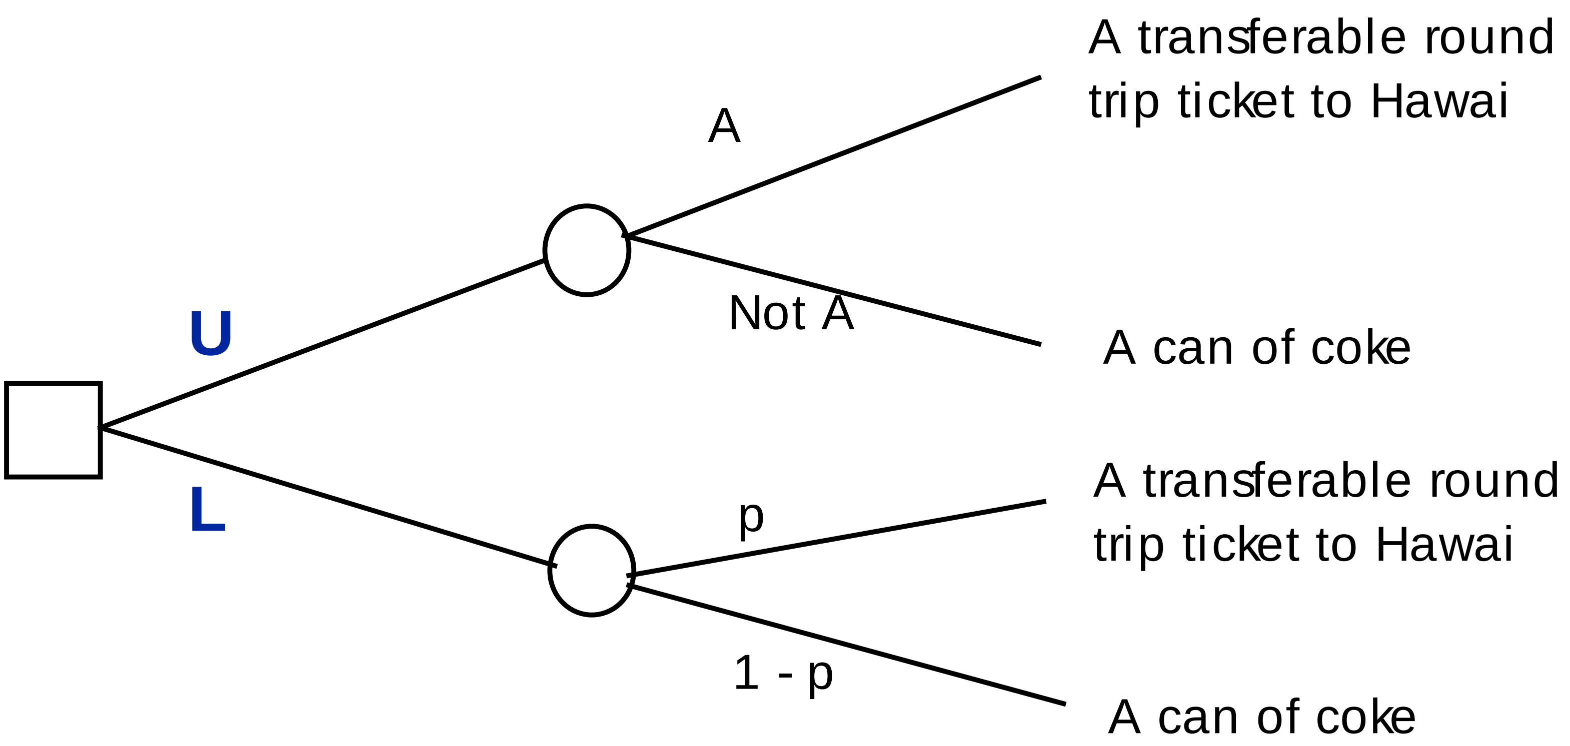

Use a tool like probability wheel (to hide the numbers).

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

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

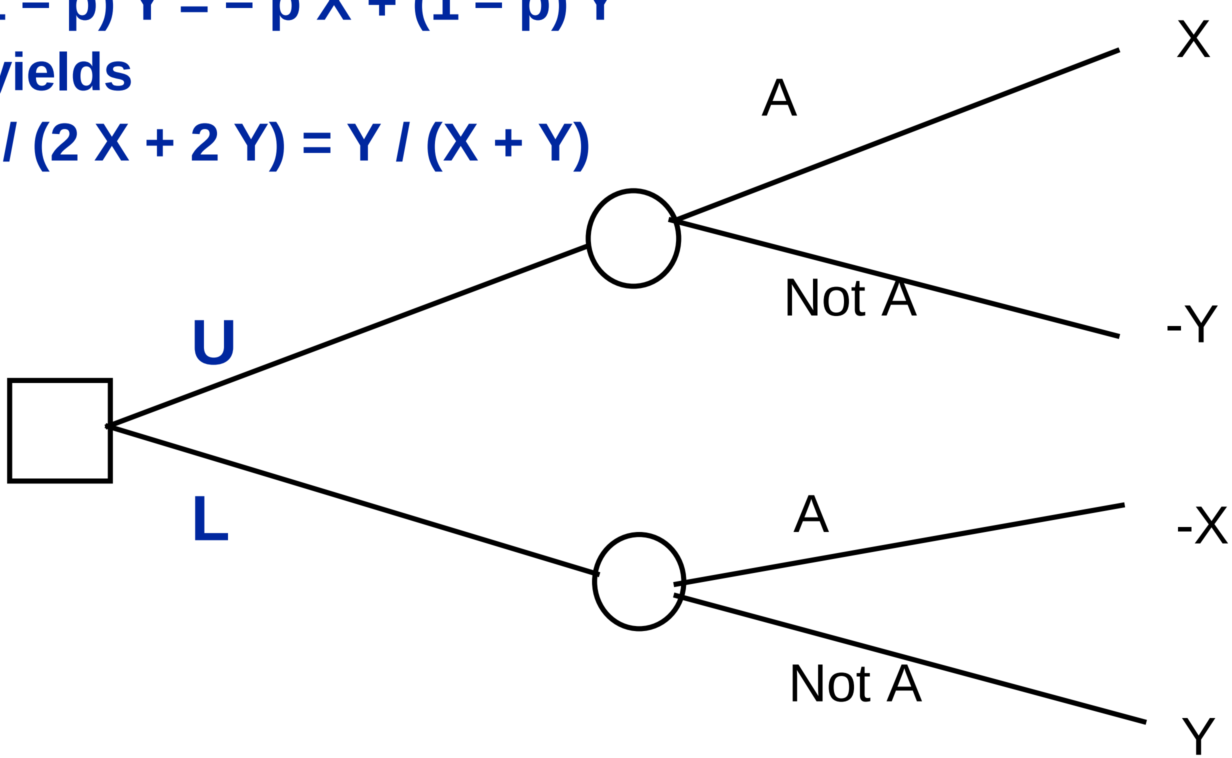

Elicitation of probabilities: Symmetric bets

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

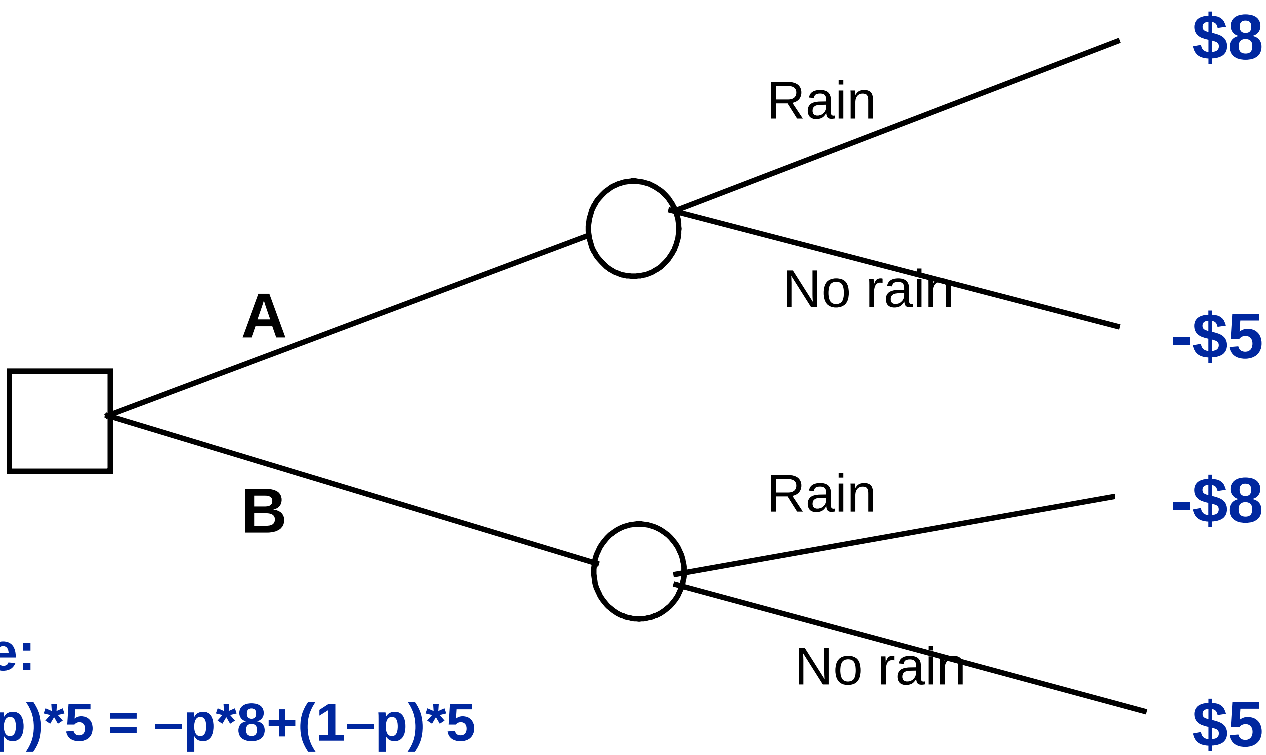

We have:

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.

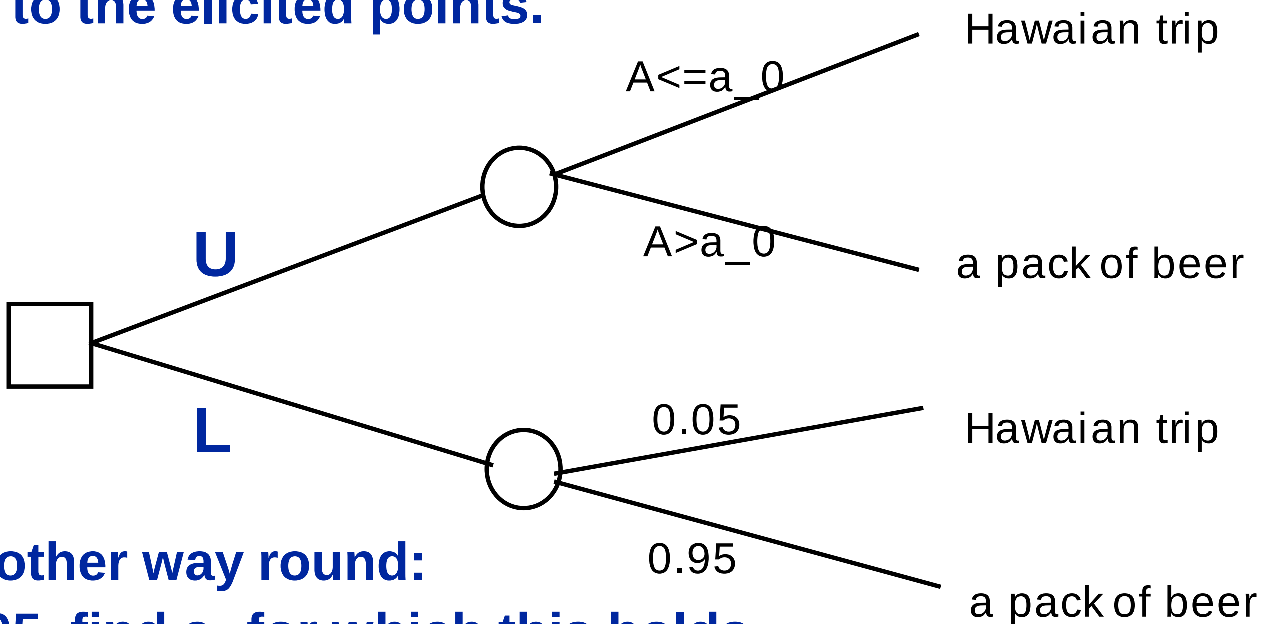

Possible to do the other way round:

- 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

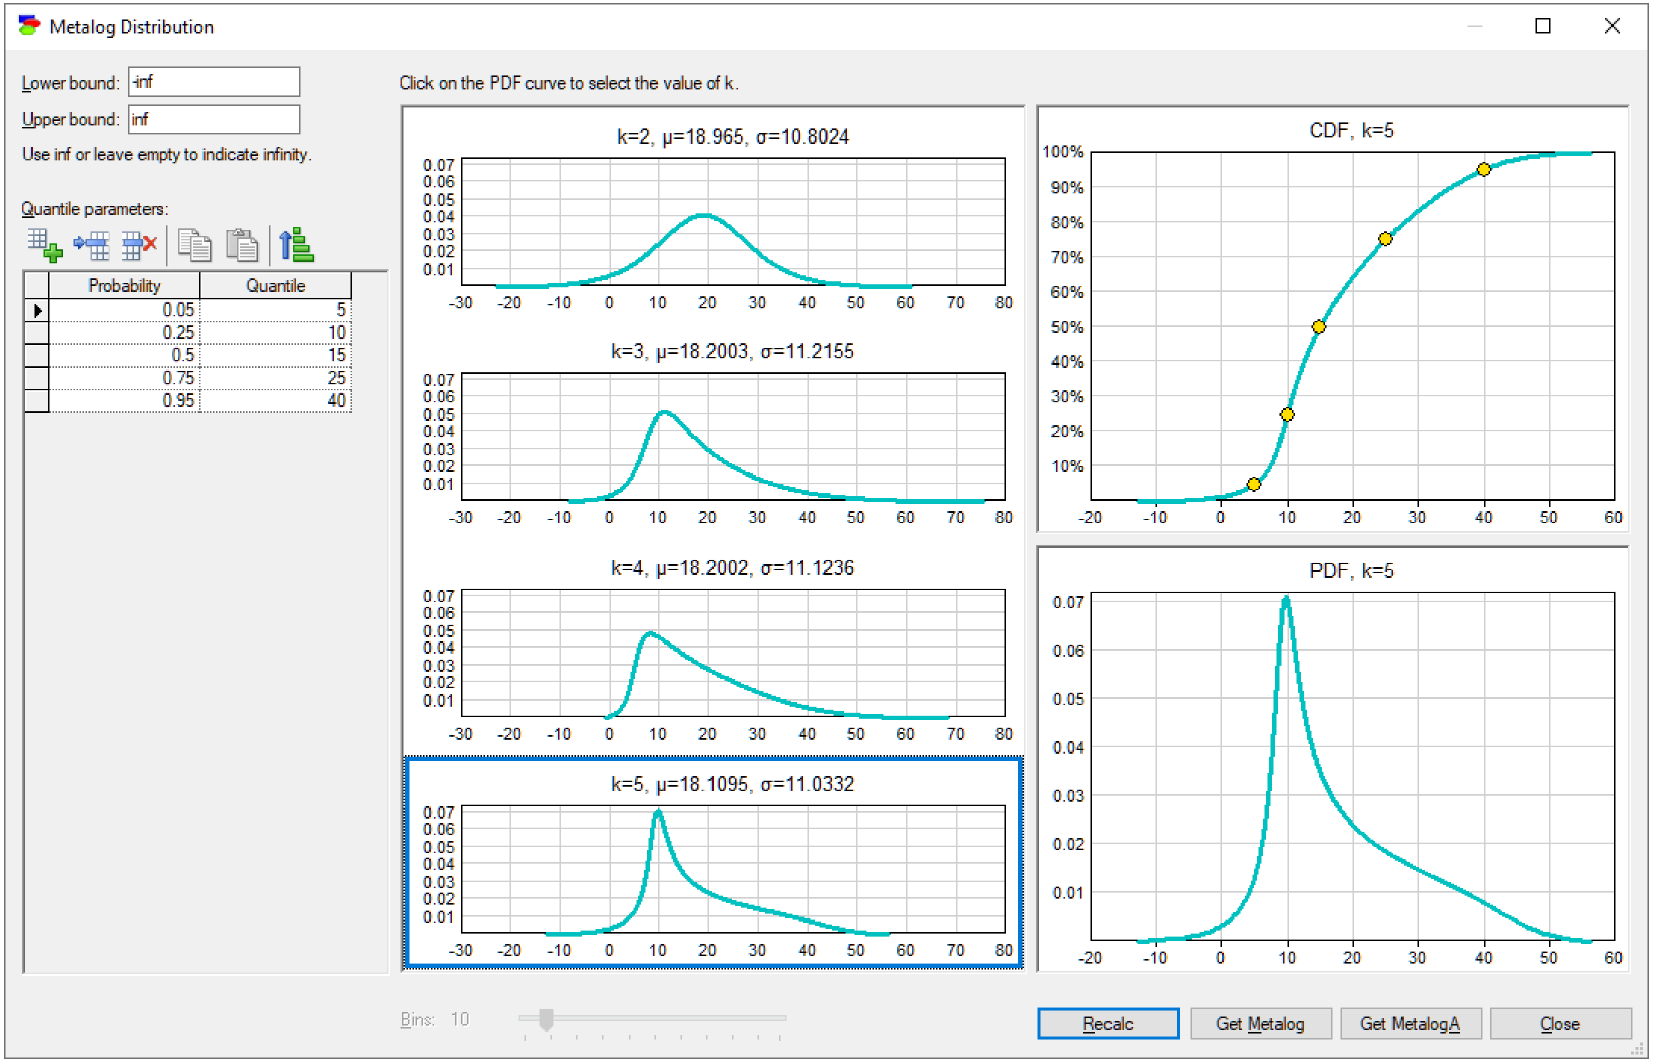

Elicitation of probabilities: QGeNIe

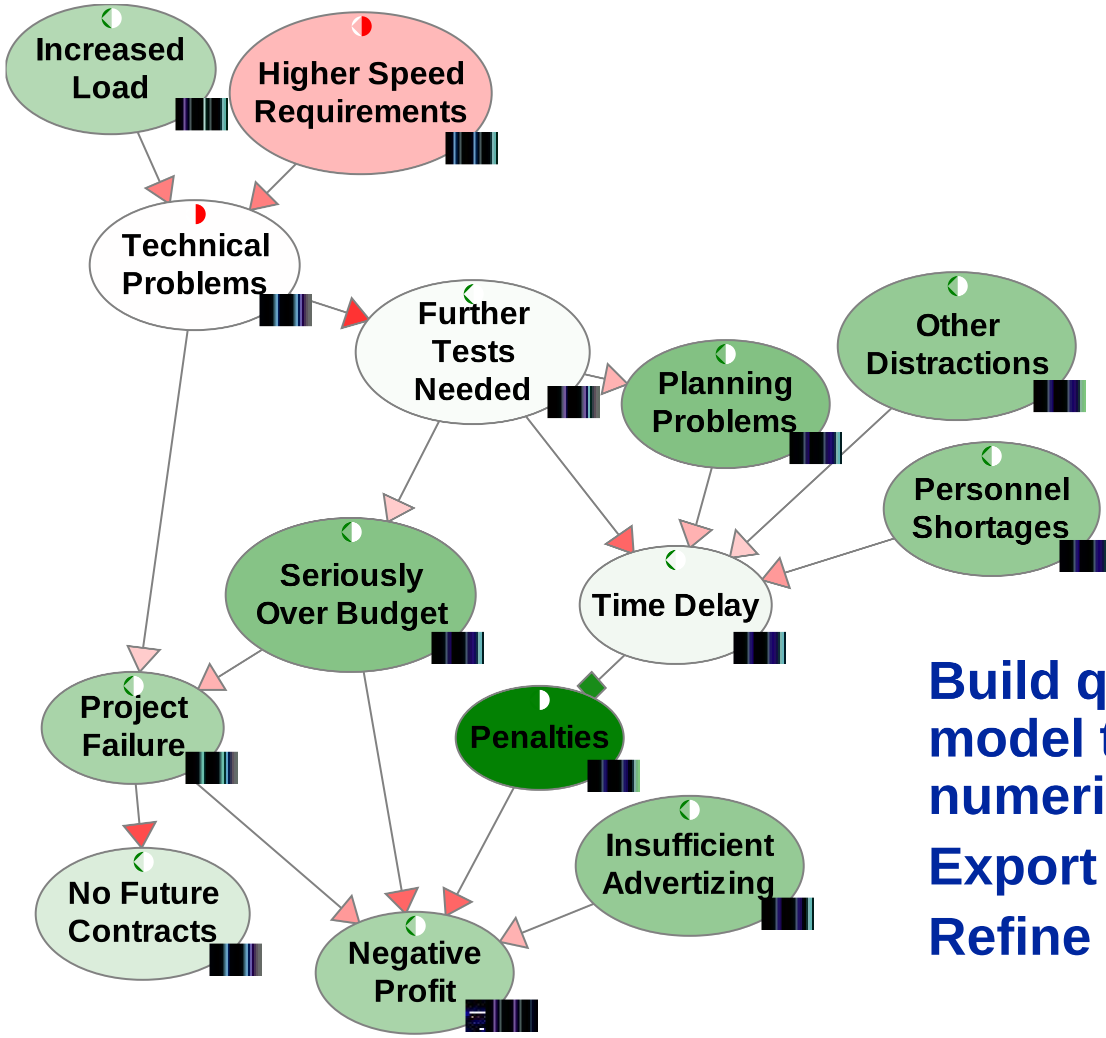

Build quickly a QGeNIe model that requires n+m numerical parameters.

Export it to GeNIe.

Refine the full CPTs.

Elicitation of probabilities: Equations

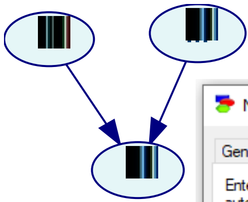

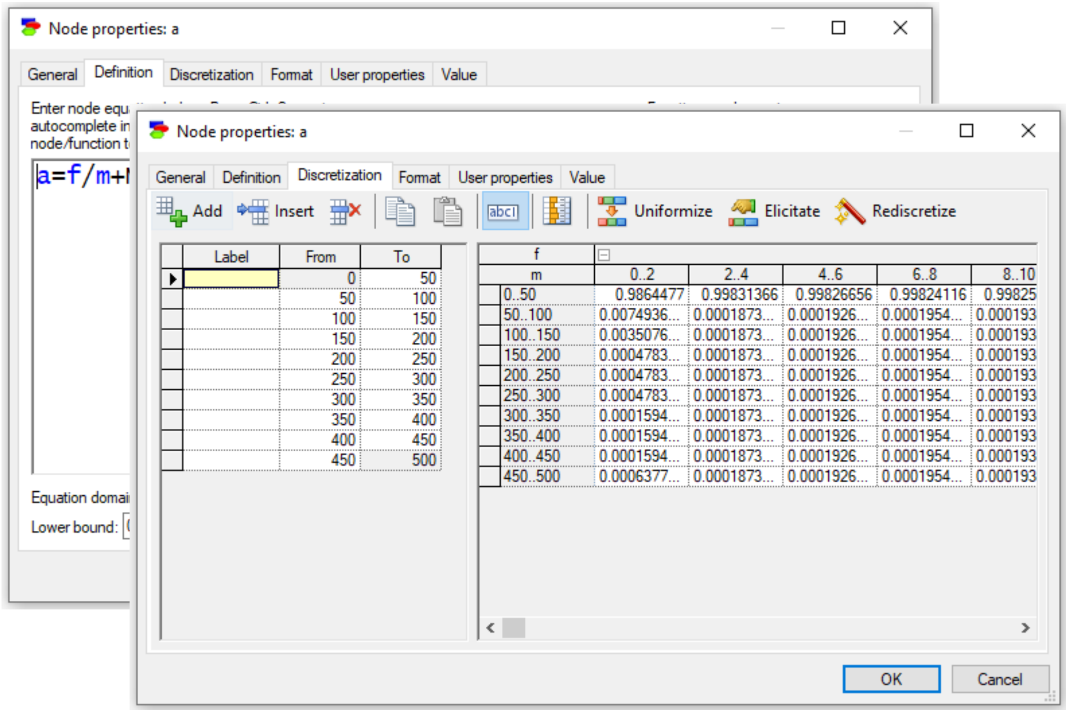

Whenever you know enough about a domain to write the interactions in form of an equation, derive the CPTs by discretizing variables

Prameter reduction: Parent divorcing

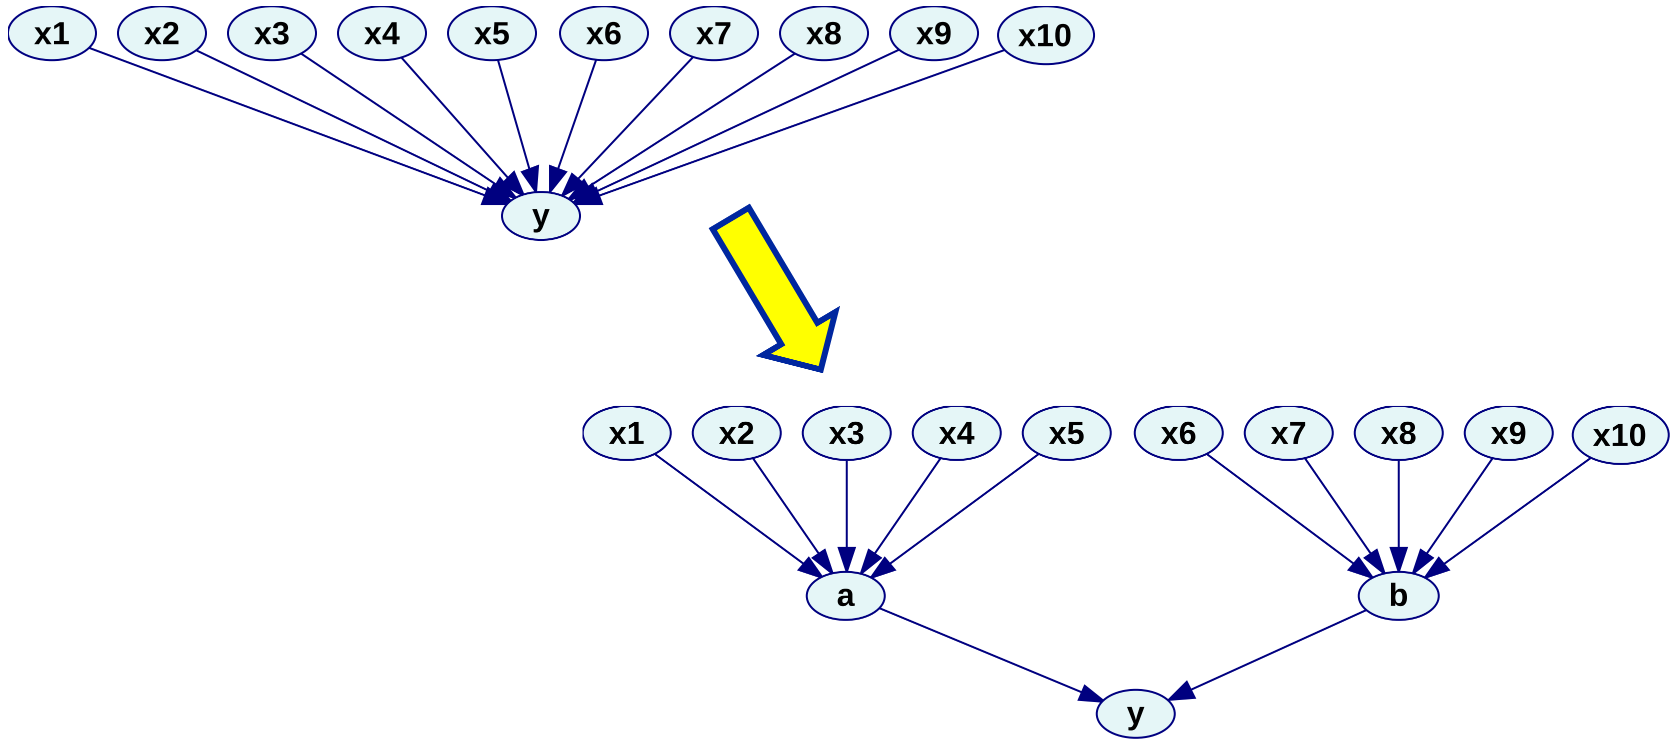

Introduce an intermediate layer of nodes for nodes with many parents.

In the above example, if all nodes are binary, we reduce the number of parameters from 2^10=1,024 to 2*2^5+2*2=68.

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.

Do parameters matter?

Random noise, Normal(0,σ)

σ = 0.3

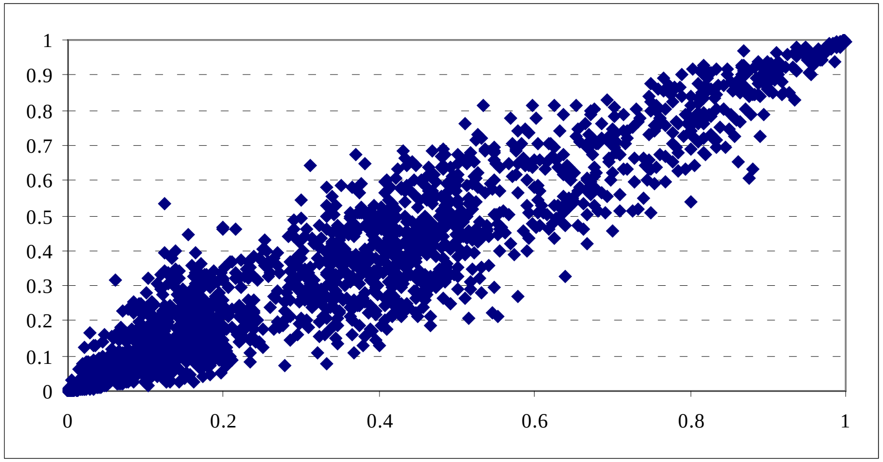

σ = 1.0

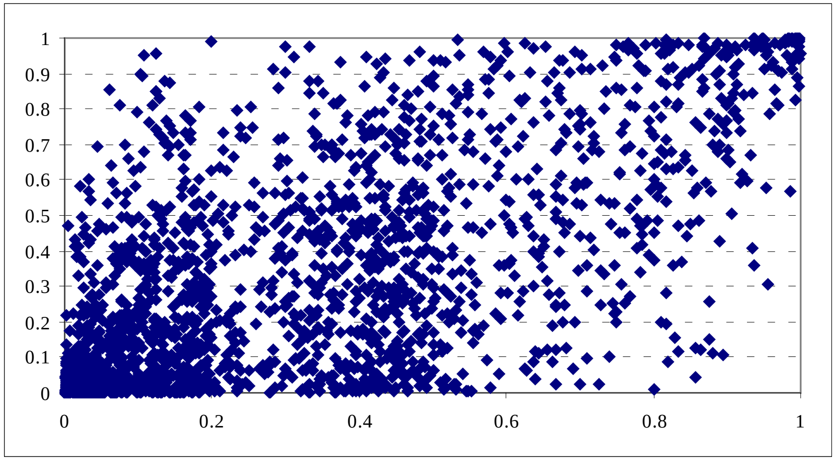

σ = 2.0

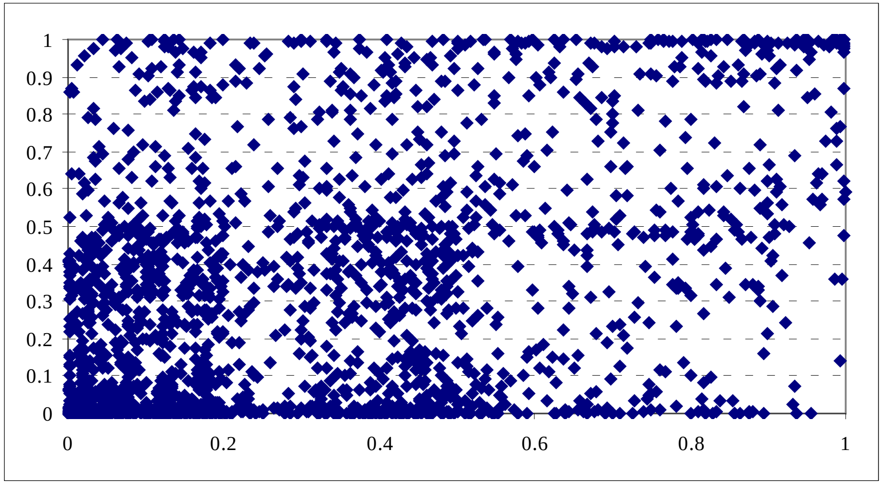

σ = 3.0

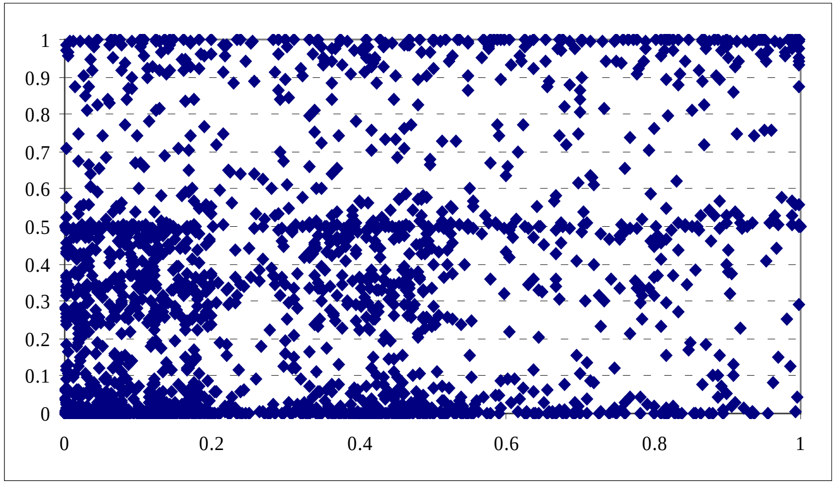

transformed parameters

original parameters

[Onisko & Druzdzel 2011]

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

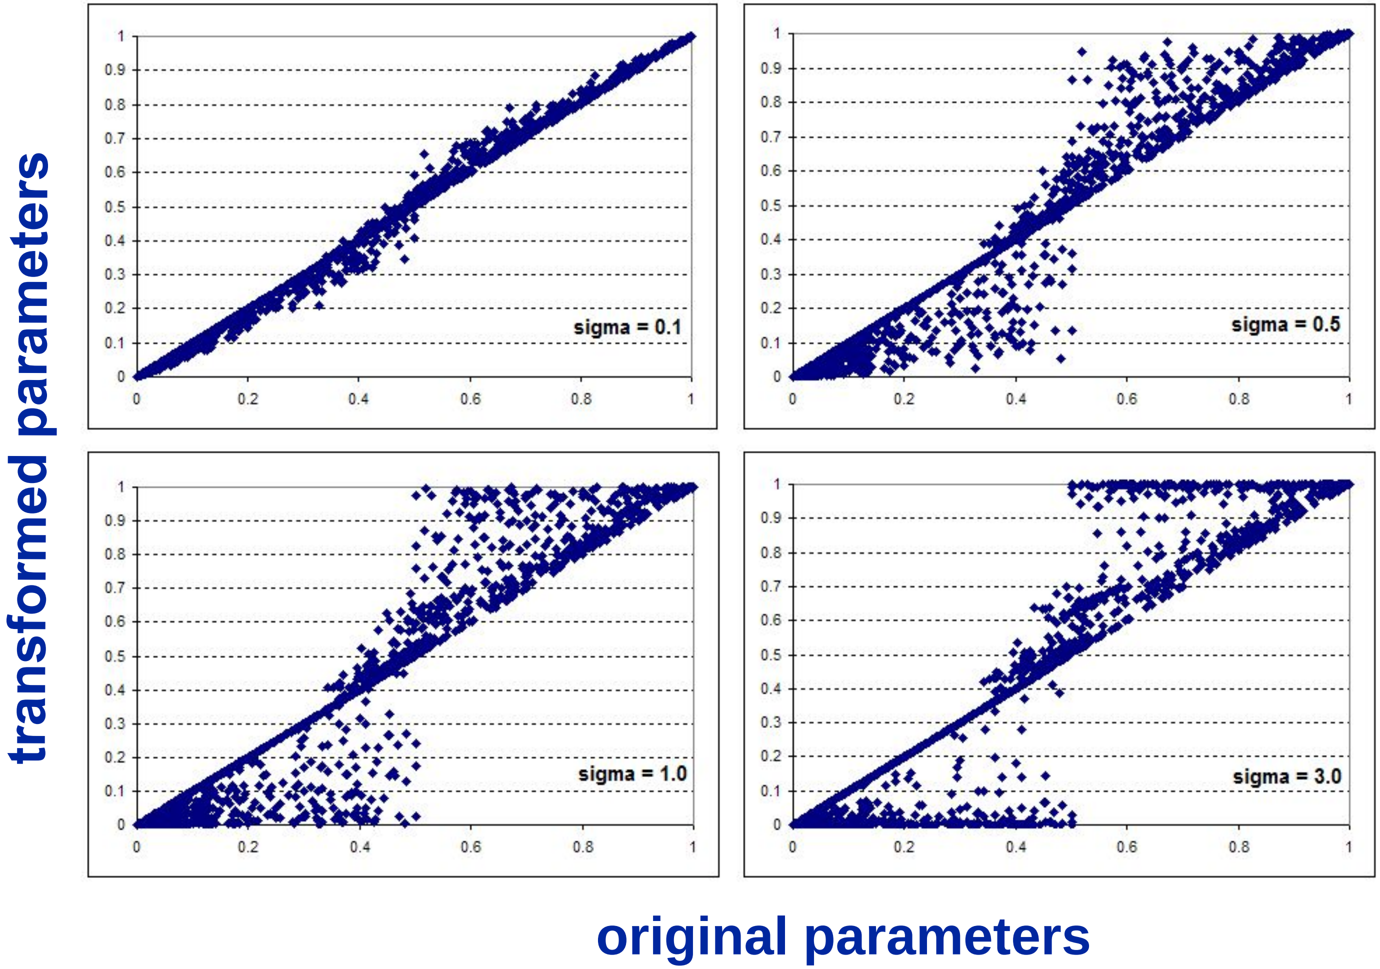

Biased noise (underconfidence), Normal(0,σ) subtracted from the largest probability in a distribution

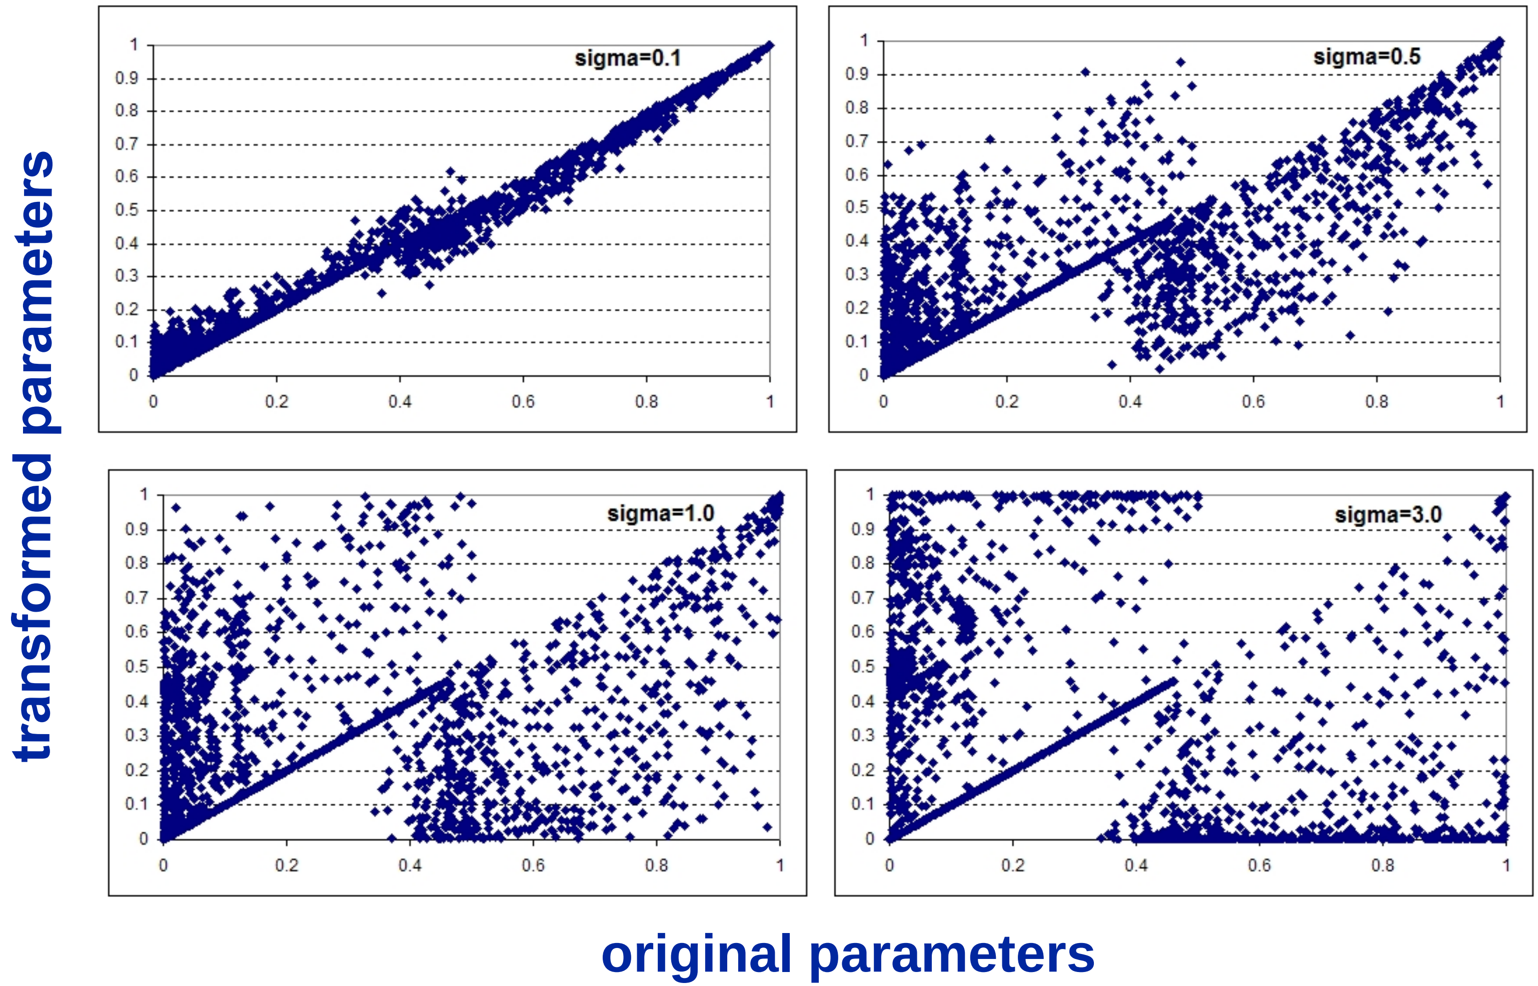

Diagnostic performance as a function of parameter accuracy

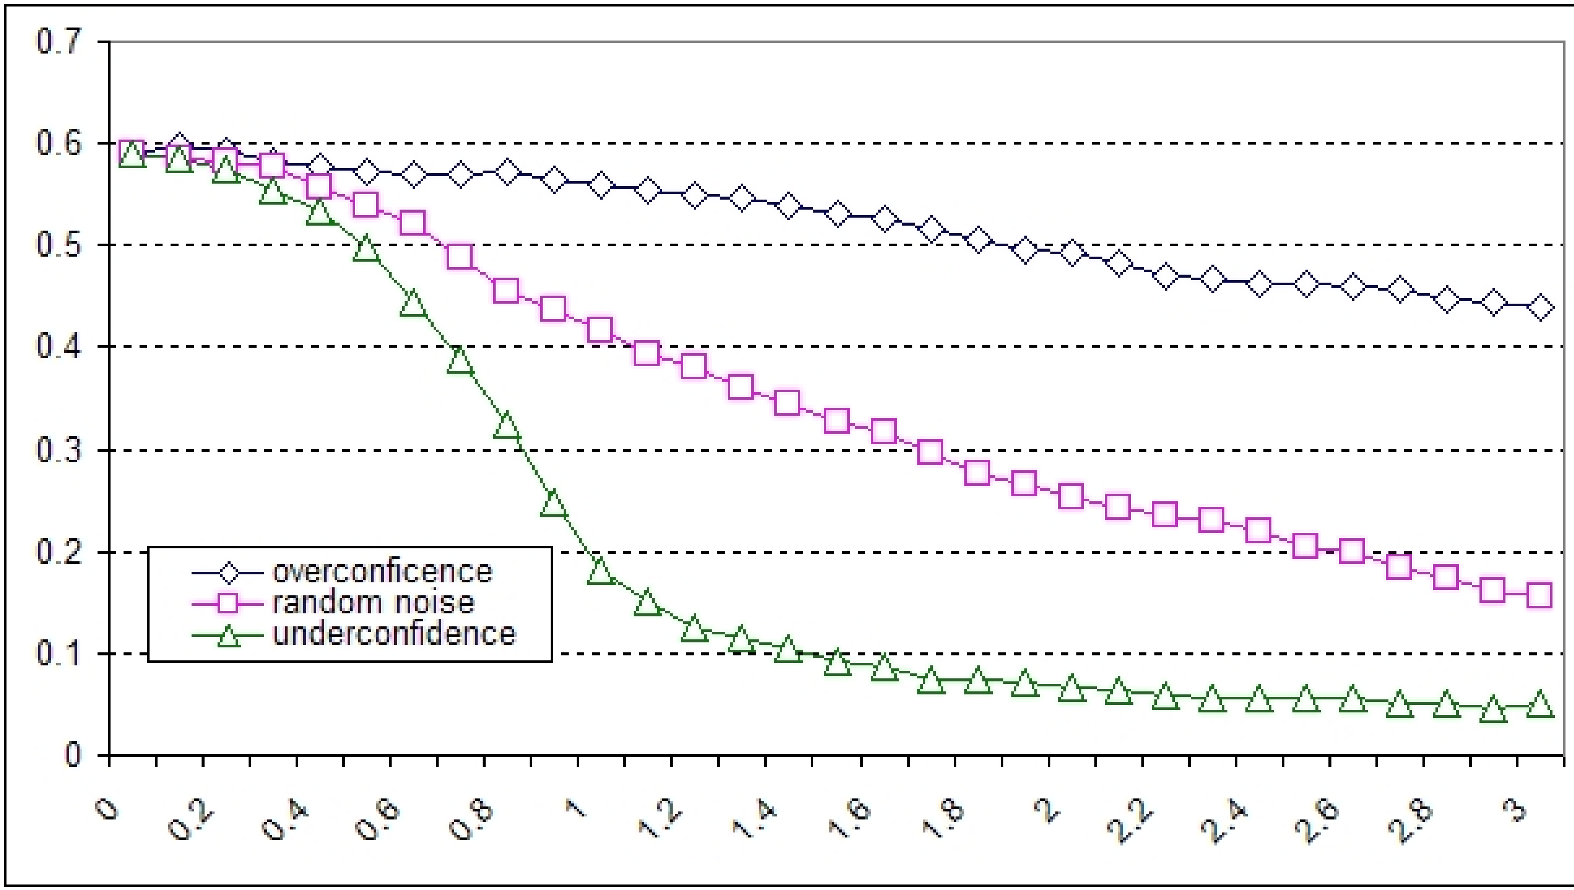

Rounded vs. original probabilities for various levels of rounding accuracy

transformed parameters

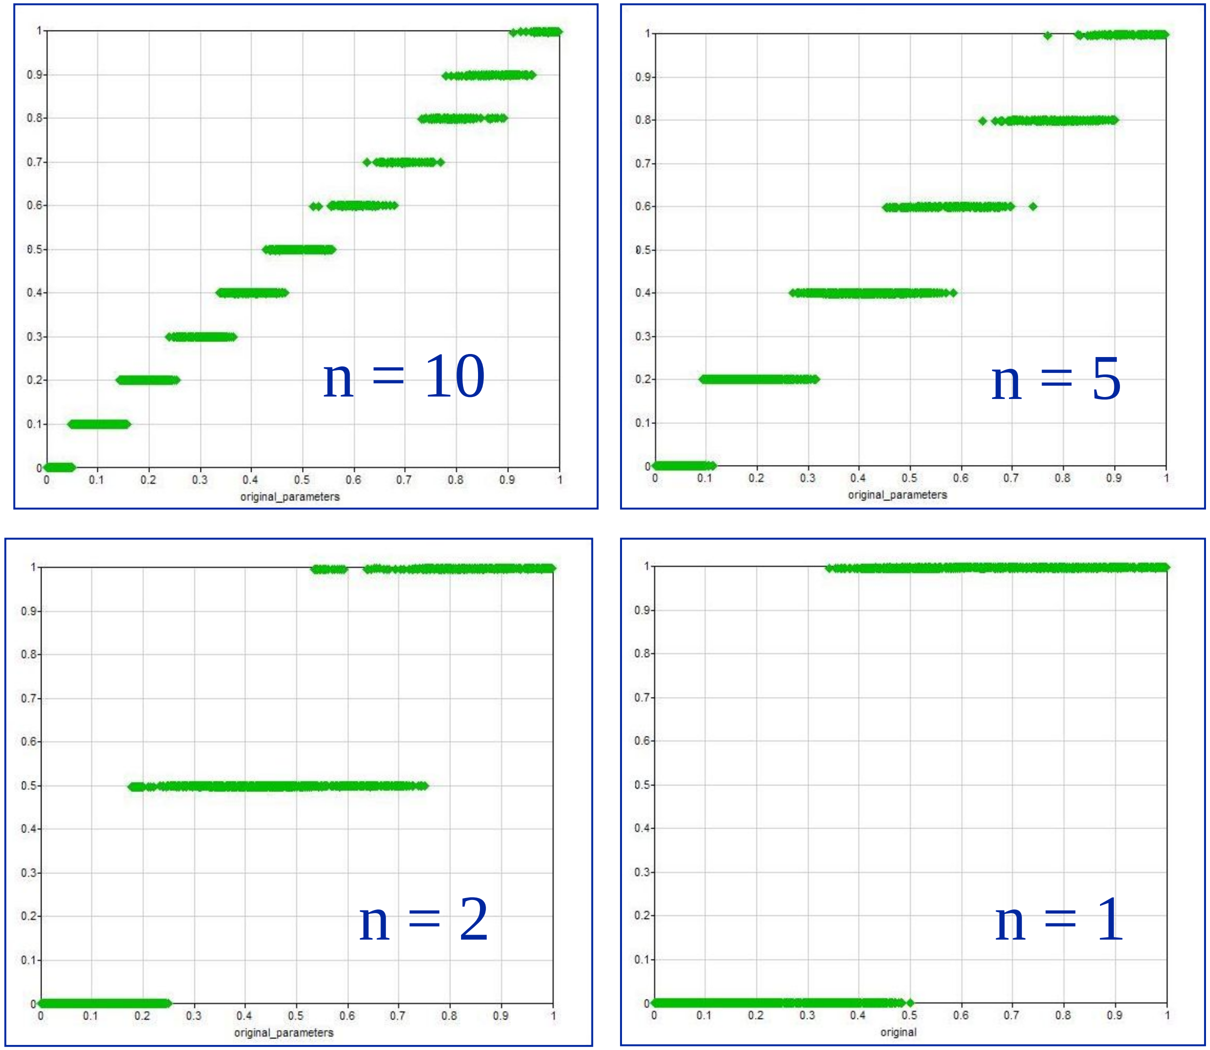

original parameters

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

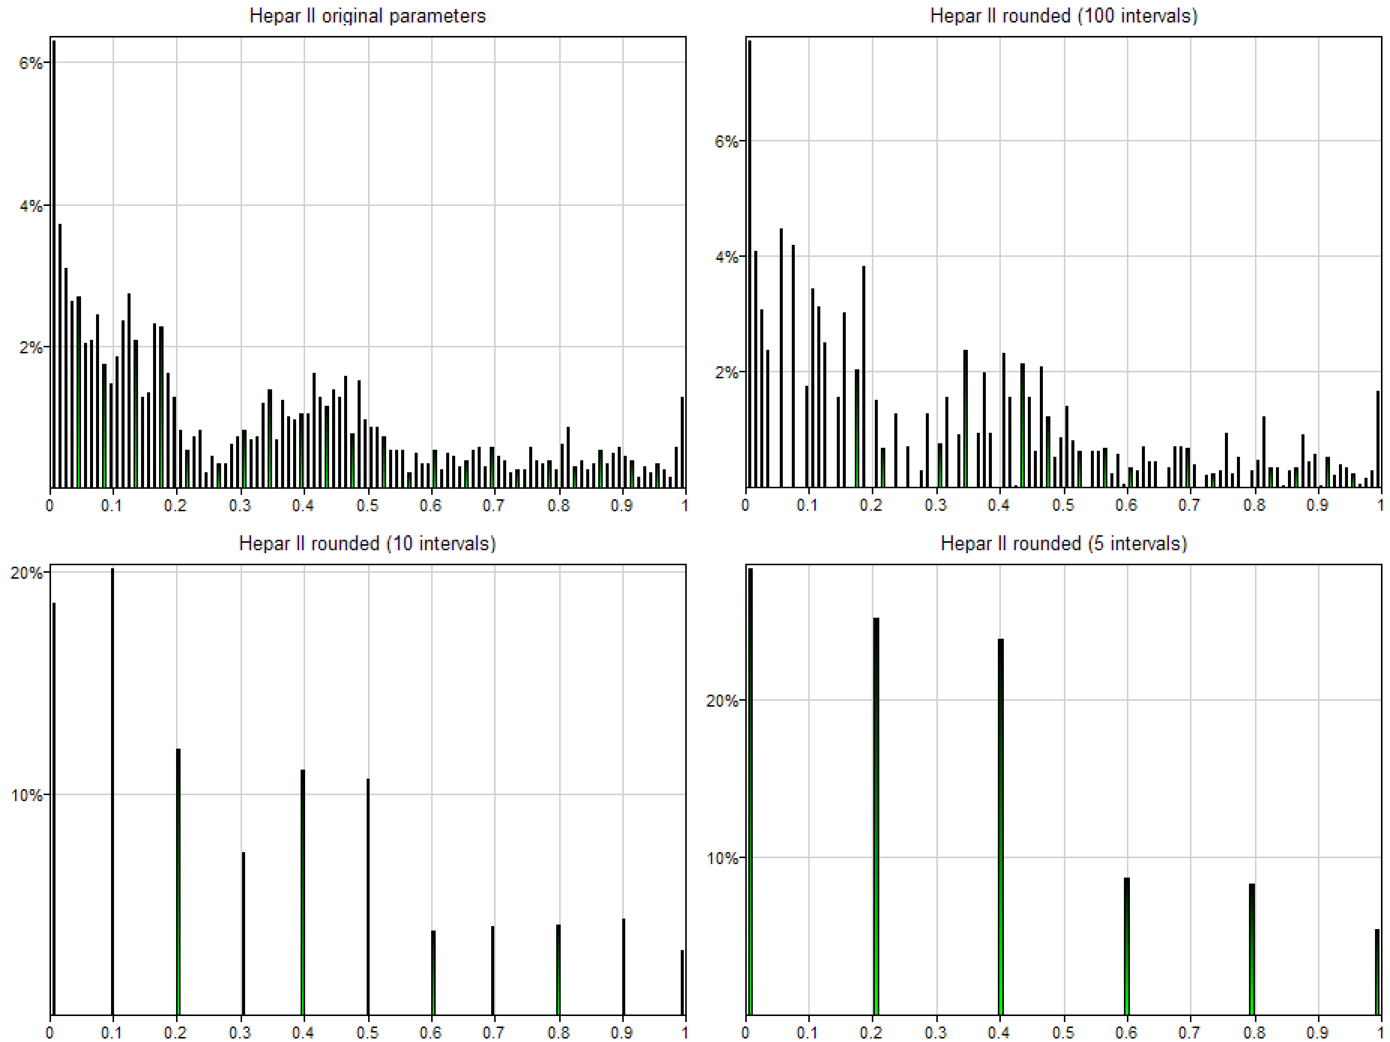

Diagnostic performance as a function of parameter accuracy (w=1)

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

	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%

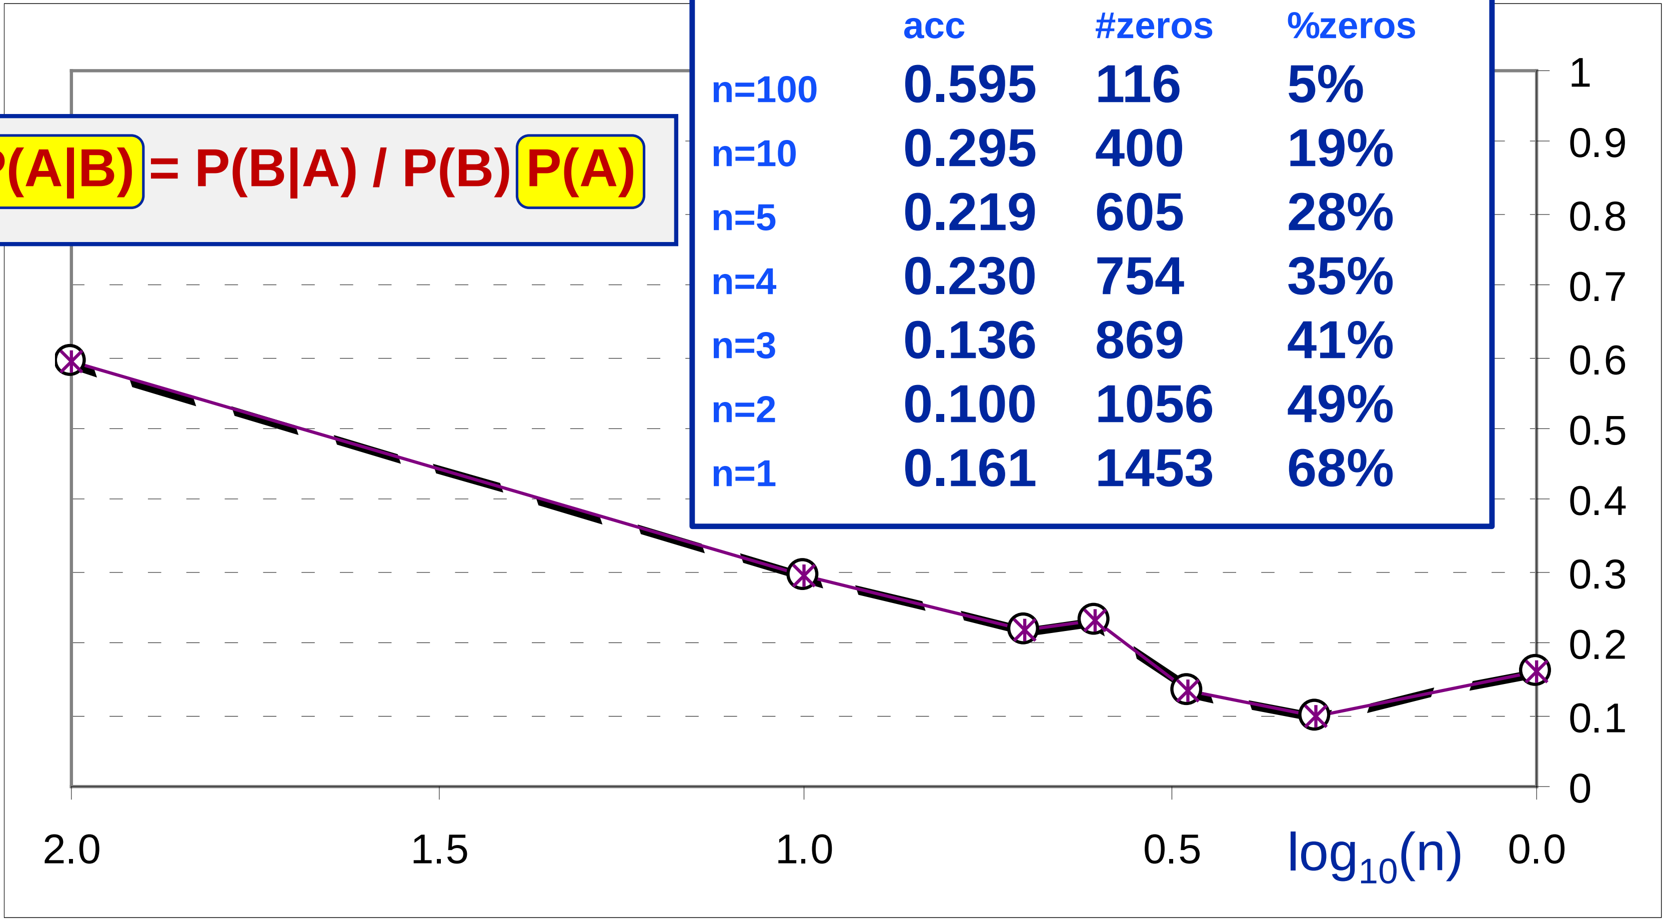

Diagnostic performance as a function of parameter accuracy and ε (w=1)

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