

# **Rational Self-Doubt: The Re-calibrating Bayesian**

Sherri Roush

If one is highly confident that #3 in the line-up is the murderer from having seen the crime, and then learns of the substantial experimental psychology evidence that human beings are unreliable at eyewitness testimony, is one thereby obligated to reduce one's confidence? How far, and why? I generalize 1st-order Bayesian rationality constraints away from idealizations in a principled way, to give a rule for revising 1st order beliefs on the basis of 2nd-order evidence about one's reliability. It is a conditionalization rule that re-calibrates the subject and sidesteps standard objections to calibration. It shows why taking doubt about one's own judgment seriously does not end up in incoherence or runaway skepticism, and what the added value of this kind of evidence is. I'll discuss some applications to disagreement and testimony, and some preliminary results from its implementation in an AI program.