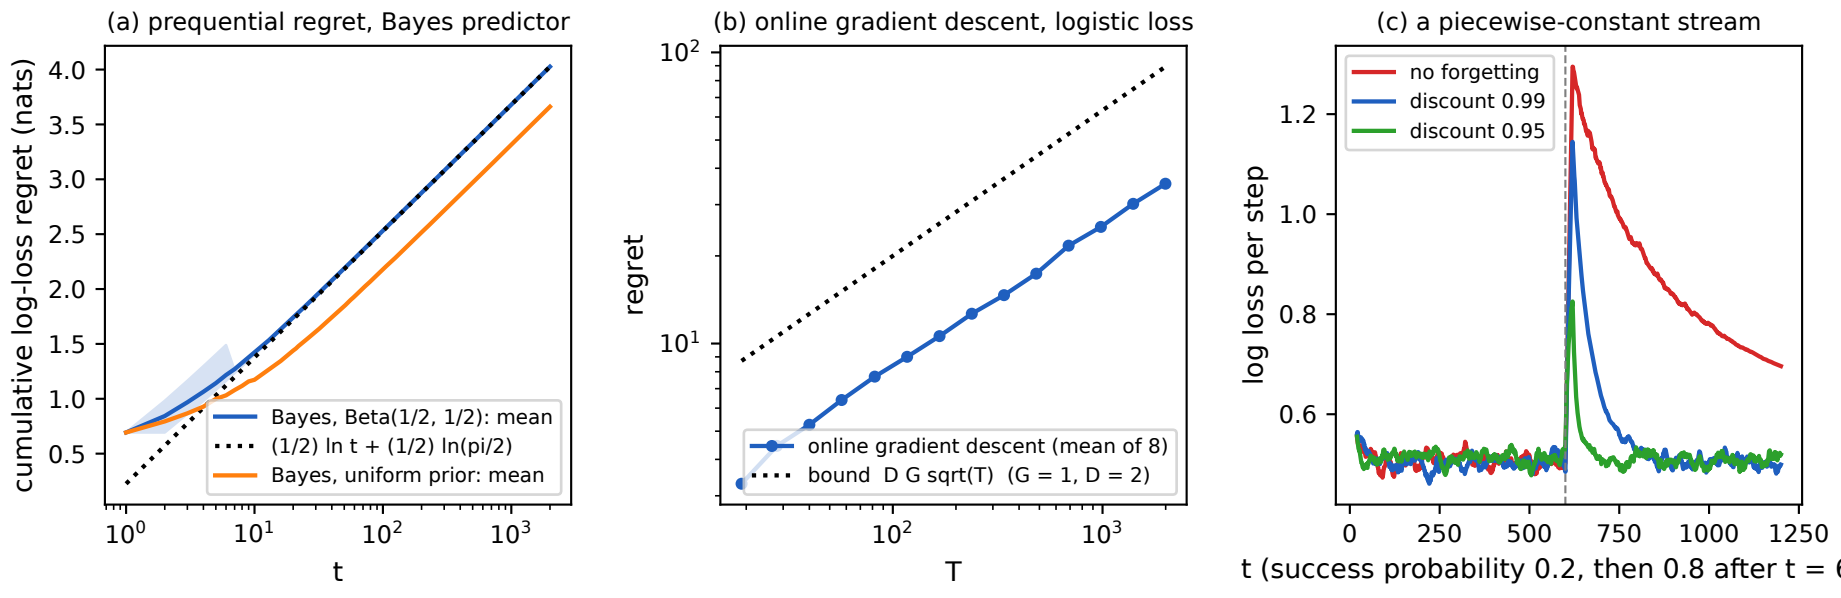

(a) prequential regret, Bayes predictor
cumulative log-loss regret (nats)
t
Bayes, Beta(1/2, 1/2): mean
(1/2) ln t + (1/2) ln(pi/2)
Bayes, uniform prior: mean
(b) online gradient descent, logistic loss
regret
T
online gradient descent (mean of 8)
bound D G sqrt(T) (G = 1, D = 2)
(c) a piecewise-constant stream
log loss per step
no forgetting
discount 0.99
discount 0.95
t (success probability 0.2, then 0.8 after t =