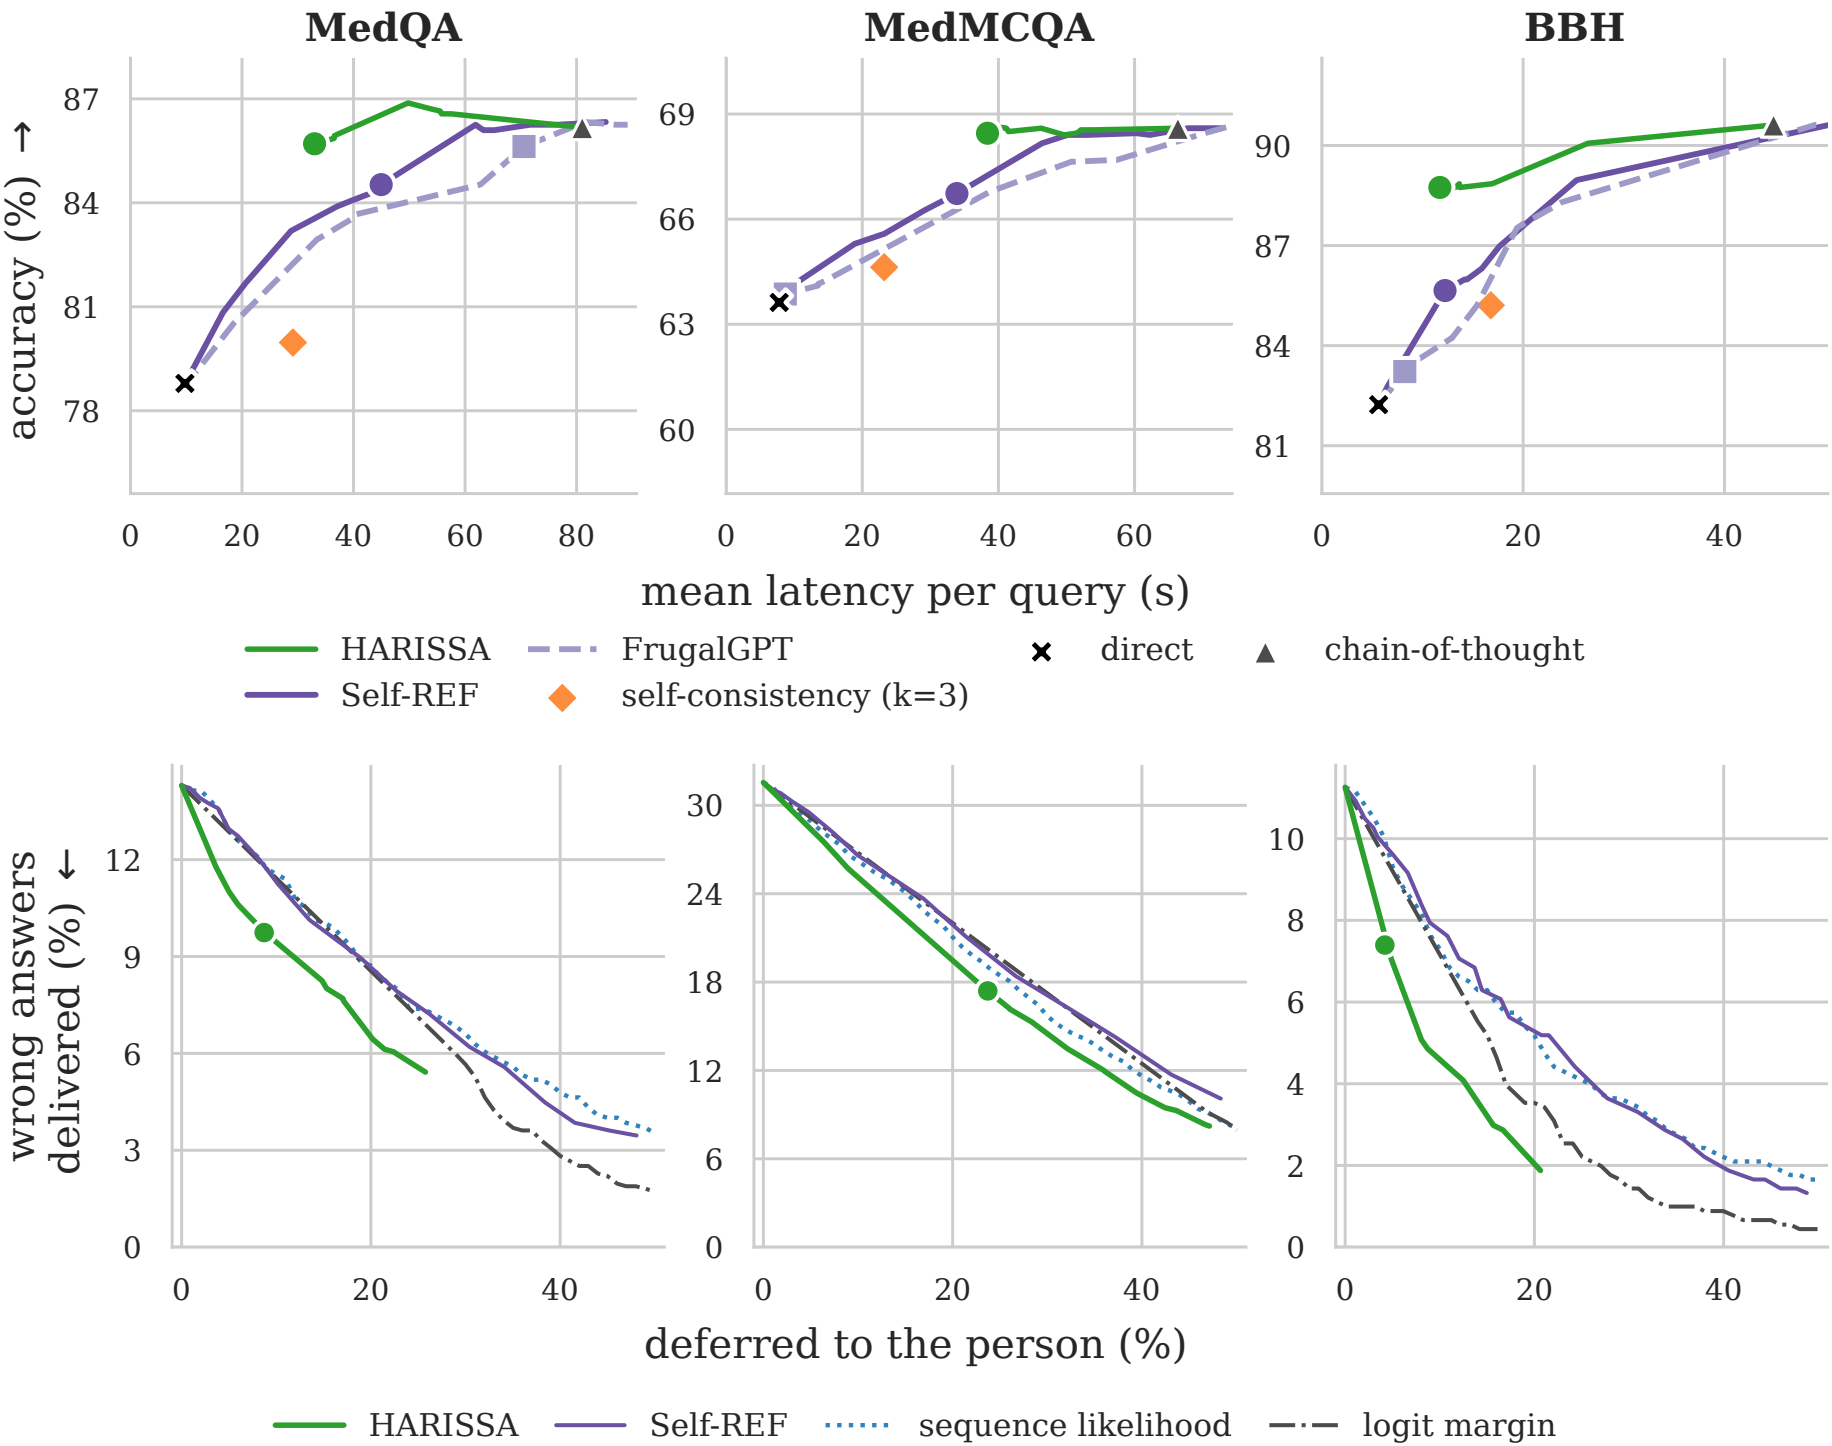

MedQA
MedMCQA
BBH
accuracy (%) ↑
mean latency per query (s)
HARISSA
Self-REF
FrugalGPT
self-consistency (k=3)
direct
chain-of-thought
wrong answers delivered (%) ↓
deferred to the person (%)
HARISSA
Self-REF
sequence likelihood
logit margin