

DOI:10.1145/3416266

Dennis Shasha

Upstart Puzzles

Privacy-Preserving Polling

Considering the probability of a flip of a coin determining the winner of a national election.

WHEN PEOPLE ARE asked whom they will vote for, they might not want to say. After all, people might judge them, ask for contributions, or publish the answer. Suppose there are two candidates, randomly called B and T. Suppose, again for the sake of this hypothetical, that there is a slight stigma against people who support T.

You say to them: “Please flip a coin. If the coin comes up tails, please tell us whom you like best. If it comes up heads, then always say T.” That way, even if a person says they will vote for T, nobody knows for sure.


Warm-Up: Suppose the true probabilities are 60% for B and 40% for T. Suppose 200 people are polled. How many will say T in response to the poll with the coin-flip rule and how many will say B?

Solution to Warm-Up: Approximately half the people—100—will flip heads and will say T, regardless of their preferences. Of the other half—60—will say B and 40 will say T. So 140 will say T and 60 will say B.

Warm-Up 2: Suppose 70% want T and 30% want B. We poll 200 people again. How many will say T in response to the poll with the coin flip rule and how many will say B? Solution: 170 for T and 30 for B.

Solution to Warm-Up 2: So to find the true support for T and B, simply subtract from the T score half of the total number of people polled. Leave the B score alone.

But now suppose a country is so di-



We understand that your choice of candidate may be something you want to keep private. At the end of this process, only you will know for sure whether the choice you mention is your real choice or not.”

vided that, depending on whom you talk to, there might be a stigma to vote for either candidate. Can the pollsters still do their job?

Question: Can you think of a protocol that will protect privacy for supporters both of B and T?

Solution: Here is one possibility. Tell the pollees (the people asked): “Please flip a coin twice. If it comes up heads both times, then please say T. If it comes up tails both times, please say B. With any other combination, please

tell us the truth.” Suppose again for the purposes of example 60% want B and 40% want T. If 200 people are polled, B will get 60 true answers and 50 because of double tails. The remaining 90 will go to T. Thus, we subtract a quarter of the total number of people polled (50 in this example) from B (yielding $110 - 50 = 60$), a quarter from T (yielding $90 - 50 = 40$) and we get the correct answer.

The only trouble with this privacy-preserving approach is that it requires doubling [CONTINUED ON P. 104]