Information flow in preferential voting

Joint work with Tahiry Rabehaja and Carroll Morgan.

What is a good democratic voting system?

- Everybody must agree that the result is “fair and true”.
- This is usually “achieved” by “scrutineers” watching that the agreed election process is properly carried out:
- Eligible voters get to vote how they want without fear of consequences, and…
- The votes must be counted accurately according to the agreed tallying principles.

Verifiability versus coercion

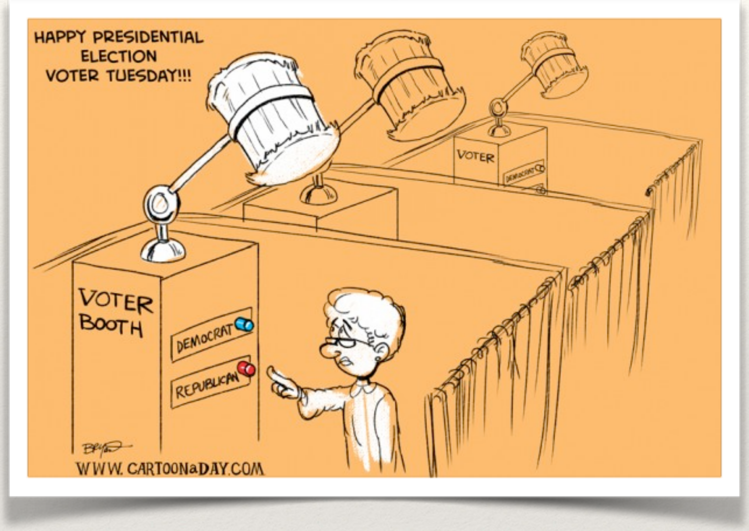

But voting in public is prone to coercion…

Cast votes in public!

- Eligible voters get to vote how they want, and…
- The votes must be counted accurately according to the agreed tallying principles.

Verifiability versus coercion

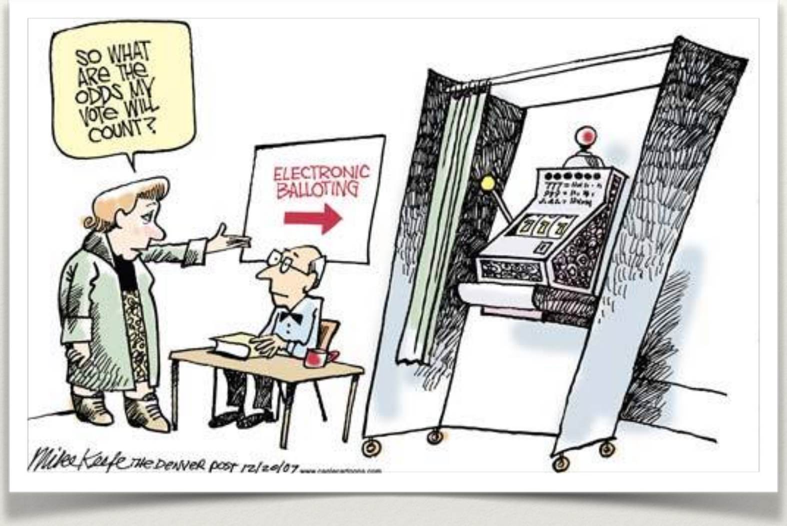
SO WHAT ARE THE ODDS MY VOTE WILL COUNT?
ELECTRONIC BALLOTING
Mike Keefe THE DENVER POST 12/20/07

Votes are cast completely anonymously!

- Eligible voters get to vote how they want, and…
- The votes must be counted accurately according to the agreed tallying principles.

But anonymous votes could be manipulated…

Verifiability versus coercion

We need to have a bit of both. But what's the right balance, and how do we implement it?

Preferential voting

In preferential voting systems:

- Multiple seats
- Multiple candidates
- Verifiability means that (some) information from the votes is (normally) released on a public bulletin board
- In STV systems, there is a function defined by the social context in order to ensure that votes are not "wasted".
- Lots of candidates to choose from
- Lots of room to use "no-hopers" for encoding signatures

Wait - can't that be used to mount a coercion attack?

Shuffle Sum

(Benaloh, Moran, Naishe, Ramchen, Teague, IEEE 2008)

1. Voters cast their votes by listing candidates in order of preference. (I won't be discussing the precise format for this.)
2. First preferences are tallied…
3. Candidates who get over the threshold are elected and therefore eliminated…
4. Candidates who fall below some minimum threshold are deemed "not electable ever" and are eliminated.
5. The voting slips are then "marked" as being partially used by being assigned a weight: a weight of 1 means the slip has not been used to elect anyone and a weight below 1 means the slip has been used to get someone elected and in subsequent counts will be used only in proportion to the remaining weight.
6. Redo steps from 2 for 2nd, 3rd etc. preferences, but with the re-shuffled and re-weighted slips.

Results of steps 3, 4 are made public.

New South Wales state election

(New South Wales Electoral Law, 1901??.)

1. Voters cast their votes by listing candidates in order of preference. (I won't be discussing the precise format for this.)
2. First preferences are tallied…
3. Candidates who get over the threshold are elected and therefore eliminated…
4. Candidates who fall below some minimum threshold are deemed "not electable ever" and are eliminated.
5. The voting slips are then "marked" as being partially used by being assigned a weight: a weight of 1 means the slip has not been used to elect anyone and a weight below 1 means the slip has been used to get someone elected and in subsequent counts will be used ONLY with probability proportional to the weight, and otherwise discarded for evermore.
6. Redo steps from 2 for 2nd, 3rd etc. preferences, but with the re-shuffled and re-weighted slips.

Results of steps 3, 4 are made public.

Information flow

Qualitative definitions of coercion resistance are concerned with individual voters and whether or not they can lie plausibly. We seek quantitative definitions which are relevant even when qualitative definitions hold. This can happen when information leaked during tallying provides additional information for the coercer to use.

Assuming the tallying is done correctly (!) this talk will focus on the published information and how to use some new techniques based on Quantitative Information Flow to measure the trade-off between "verifiability" and "coercion resistance".

Verifiability means that the voters have confidence that their votes have been correctly tallied. In electronic systems this is implemented using cryptography and provides a high guarantee that the cast votes have been counted correctly.

Probabilistic analysis

Our mathematical analysis of information flow shows how to:

- Define a “statistical” measure of coercion resistance in scenarios where the coercer tells victims to use the preference ordering as a signature. The measure can be used to provide strong evidence that the information released during tallying is not sufficient for the coercer to be able to tell whether or not the coercion has occurred or not. The technical details of this include using metrics on probability distributions.
- Even in scenarios where the coercer is potentially able to identify the coercion, we are able to measure the cost to the coercer of mounting a successful attack. The technical details of this include describing the cost function as a non-standard entropy measurement.

Channels for information flow

A channel is a concrete model for reasoning about information flow.

We imagine that a process/program/system which performs some computation on some data which is supposed to be "kept secret" can leak some information about that data during the computation.

An observer of the system can use what he already knows about the likelihood of the value of the secret, together with what he knows about the correlation of secrets and observations, together with his actual observations to make accurate observations about the actual value of the system.

Change in entropy

The extent to which a channel releases information can be measured by looking at the “change in entropy” of a secret modelled as a probability distribution. The probability distribution captures an observer’s state of knowledge before making further observations based on the channel. By making correlations between the observations and the observer’s state of knowledge, the resulting “conditional entropy” can be used to understand how much information has leaked due to the channel.

Channels for information flow

We can view a tallying protocol for elections as a channel as follows.

The secrets in this case are the voting slips after the voters have submitted them to the "election system".

Since the tallying process is public knowledge, and opinion poles taken close to the election give a good indication of how voter preferences sit, the information emitted during the tallying process can in principle be used to form a very accurate idea of how various "voter coteries" actually voted.

Definition and examples

A channel is a mapping from probability distributions over secrets to (joint) probability distributions over secrets and observables.

$$\mathbb{D}\mathcal{X} \rightarrow \mathbb{D}(\mathcal{X}\times\mathcal{Y})$$

This models the observer's initial "knowledge" of the secret

This models the correlation between the initial "knowledge" and the observations.

How does this work in voting?

- What is the secret? It is the preferences on the set of ballots as a whole.
- What is the prior? It is based on the likelihood of the various possible vote-casting patterns informed by e.g. exit polling.
- What are the observations? They are the results of the eliminations during the tallying process (for example, but see later).
- What is the channel? It is the relation between the set of ballots and the observations. We call this a tallying channel.
- The degree to which a tallying channel is able to withstand coercion depends on the information released during tallying, the coercion strategy and the prior (e.g. the exit poll as above).

An extreme case

Suppose there are 3 candidates A, B, C and 100 voters and two available seats

Suppose that there are three likely orderings for preferences: ABC, ACB and BCA with a (probabilistic) variation as follows:

	$1/2$	$1/2$
ABC	30	32
ACB	60	58
BCA	10	10

There are two possible outcomes: AC are elected or AB are elected. It's very close with the AC option slightly favoured.

What does a coerced profile look like?

If some of the ACB voters can be persuaded to vote ABC instead then only a small shift in preference will see the likelihood shift towards an AB result:

	$1/2$	$1/2$
ABC	35	37
ACB	55	53
BCA	10	10

In this election, first A is elected, then C is eliminated leaving B to be elected. But can the coerced voters plausibly lie? Moreover how much would it cost the coercer in order to be successful?

Observing the signature

What is the chance that the coercer is convinced that the voters have voted in the way so directed? Remember that the coercer does not get to see the voting slips, but only the observations.

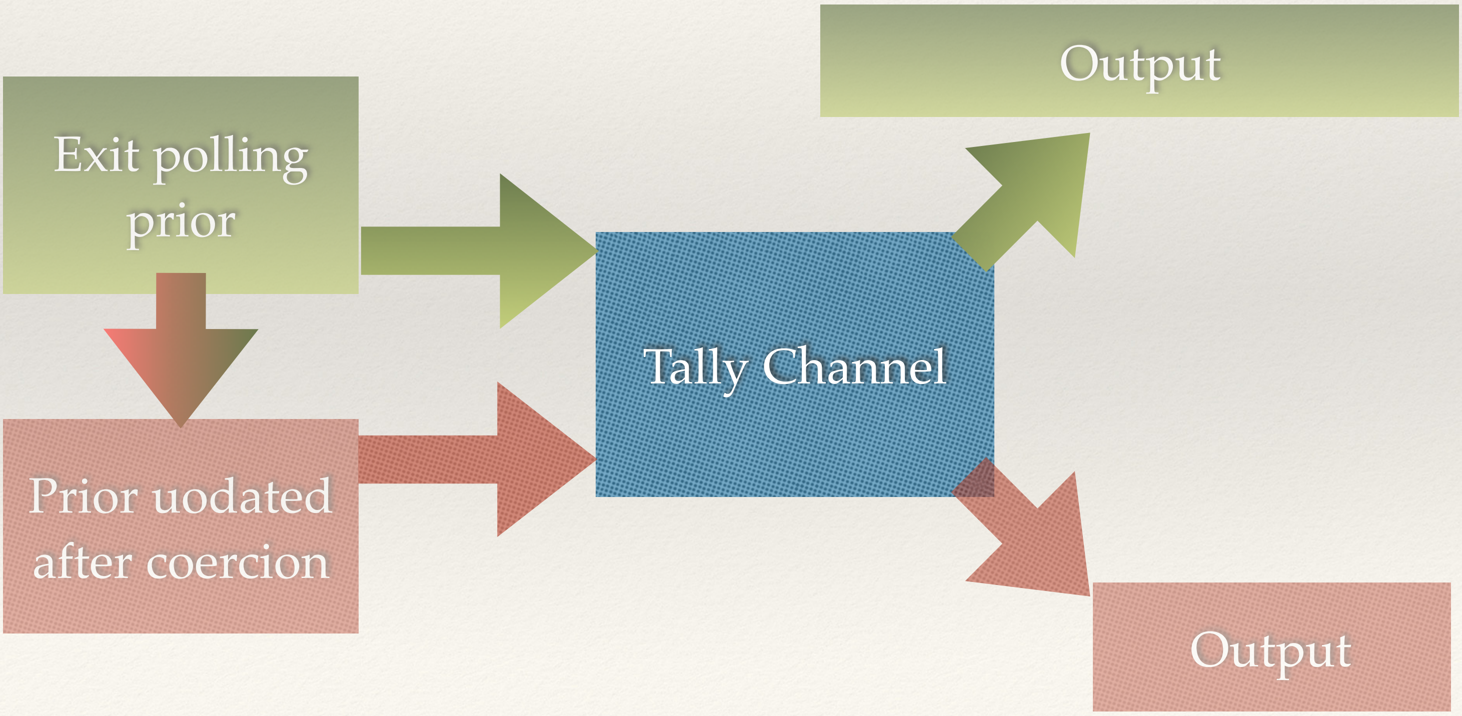

Observing the signature

What is the chance that the coercer is convinced that the voters have voted in the way so directed? Remember that the coercer does not get to see the voting slips, but only the observations.

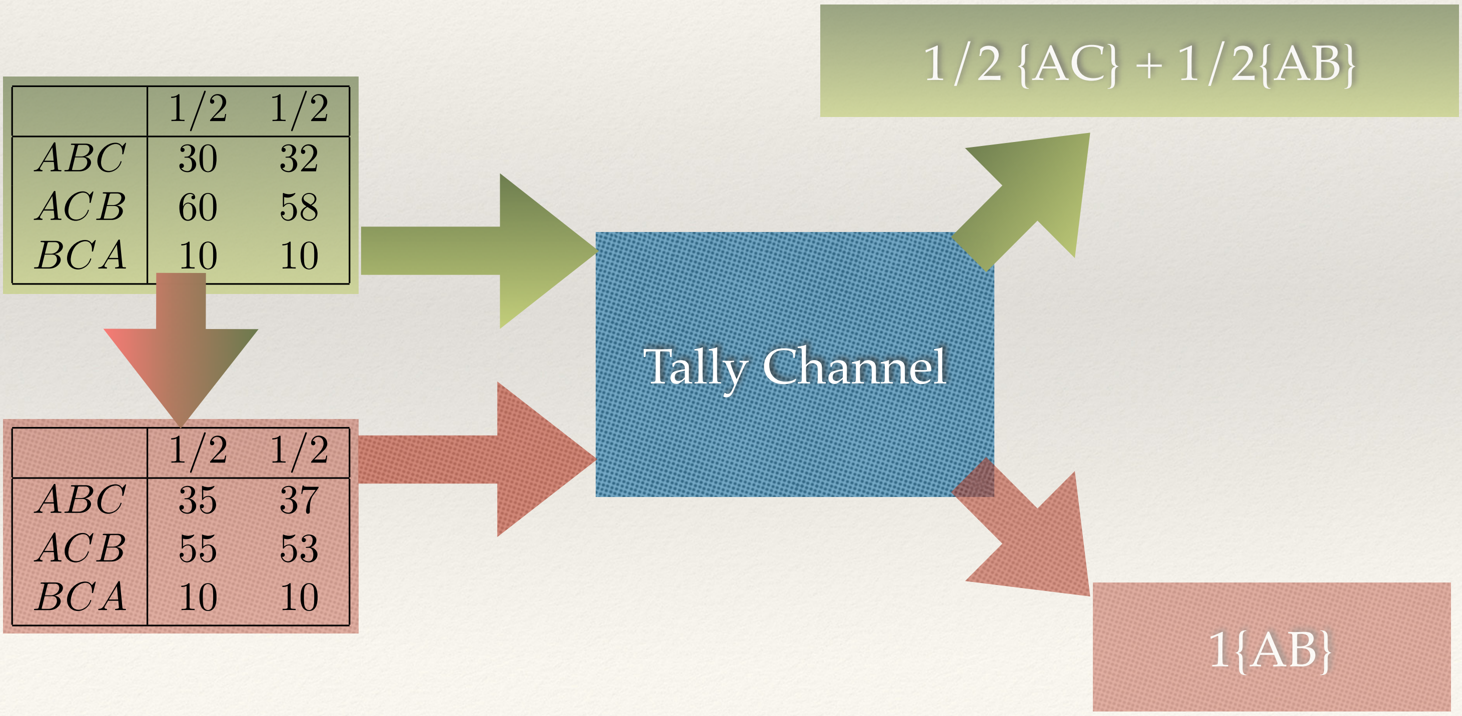

	1/2	1/2
ABC	30	32
ACB	60	58
BCA	10	10

	1/2	1/2
ABC	35	37
ACB	55	53
BCA	10	10

Tally Channel

1/2 {AC} + 1/2{AB}

1{AB}

Observing the signature

The observer can "tell the difference" between these marginals by "testing" them on different "random variables" from Observed state to reals…

1/2 {AC} + 1/2{AB}

Theorem: Given two (hyper)distributions Δ , $\Delta'$and any gain function V_g we have that:

$$|\mathcal{E}_{\Delta}(V_g) - \mathcal{E}_{\Delta'}(V_g)| < K(\Delta, \Delta')$$

1{AB}

Observing the signature

The observer can “tell the difference” between these marginals by “testing” them on different “random variables” from Observed state to reals…

1/2 {AC} + 1/2{AB}

In this case the relevant Kantorovich distance is on the marginals: is 1/2.

1{AB}

Observing the signature

The observer can "tell the difference" between these marginals by "testing" them on different "random variables" from Observed states to reals. This suggests a novel definition for quantitative coercion resistance.

1/2 {AC} + 1/2{AB}

Definition: A tallying channel $C : \mathcal{X} \to \mathcal{Y}$ is private at $\epsilon > 0$ and prior π and coercion strategy σ if:

$$|C[\pi]\!\downarrow - C[\pi']\!\downarrow| < \epsilon \qquad\qquad \pi' = \sigma(\pi)$$

1{AB}

The coercion strategy

The previous example assumed a coercion strategy where the coercer could locate the voters that need to have their minds changed. In a more realistic coercion attack the coercer would not know whether the voter needs to have his/her mind changed. This attack would select some k voters to coerce with uniform probability — in this scenario only some proportion of the coerced voters will change their minds.

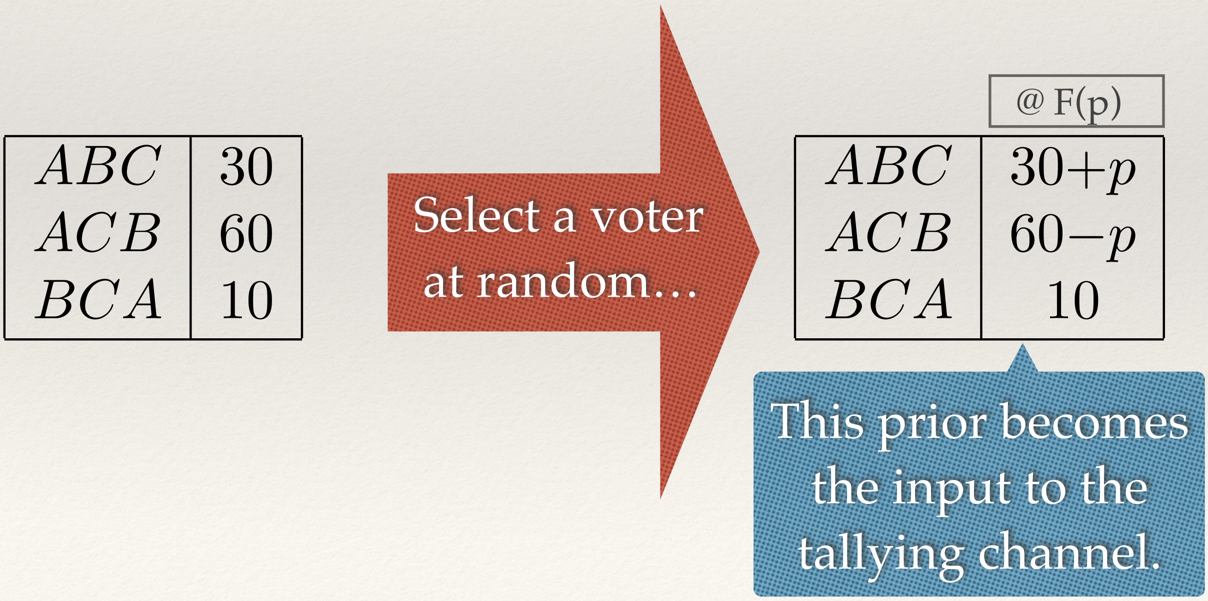

ABC	30
ACB	60
BCA	10

Select a voter at random…

@ F(p)

ABC	$30+p$
ACB	$60-p$
BCA	10

This prior becomes the input to the tallying channel.

Computing the cost

Finally, even if we find that a tallying channel is not coercion resistant, we can measure the cost of a successful coercion in a similar manner and possibly argue that the coercion is too expensive to be worthwhile!

Future work

- Scale up the analysis for real data!
- Complete the modelling of Shuffle Sum and the NSW election
- Use available election data to measure the information released
- Use Carroll's NSW election prototype, and test it on scenarios where we inject a coercion attack.