ISSN 1045-6333

SHOULD CANDIDATES FLIP A COIN IF THE DIFFERENCE IN THEIR VOTES IS SMALL?

Eric Rasmusen

Discussion Paper No. 325

05/2001

Harvard Law School
Cambridge, MA 02138

The Center for Law, Economics, and Business is supported by a grant from the John M. Olin Foundation.

JEL Class: C11, C44, D70, D72, D81

Should Candidates Flip a Coin if the Difference in Their Votes is Small?

Eric Rasmusen*

Abstract

A coin flip can be a good way to settle an election if the margin of victory is small and it is known that there is a good chance of fraud by one candidate. In that case, however, an even better rule is to award victory to the apparent loser.

*John M. Olin Senior Research Fellow in Law and Economics, Harvard Law School and Professor of Business Economics and Public Policy, Kelley School of Business, Indiana University

.

Should Candidates Flip a Coin if the Difference in Their Votes is Small?

Eric Rasmusen*

1. Introduction

The 2000 U.S. Presidential Election was the subject of intense cries of unfairness. A week before the election, George Bush seemed likely to beat Al Gore handily, but Gore surprised everyone by catching up in the last few days after intense campaigning and a surprise release of Bush's conviction for drunk driving 20 years before. The entire election turned on who won Florida. Bush was ahead by 1,831 votes, less than one-tenth of one percent of the total Florida vote. After the required recount, his margin had dropped to 784 votes. There then ensued a battle of the lawyers that ended a month later with a state-certified Bush margin of 537 votes and his election as president. (See Rusin [3] for details.)

Many Democrats were outraged. How could Bush become president when the margin was so close? Surely there should be a revote or something. A number of journalists, including Stephen Jay Gould [1], suggested, perhaps humorously, flipping a coin.

*John M. Olin Senior Research Fellow in Law and Economics, Harvard Law School and Professor of Business Economics and Public Policy, Kelley School of Business, Indiana University.

It was just not just the closeness that gave rise to indignation. Both sides claimed that the official count was improper. The Democrats objected that confusion over instructions resulted in many Democratic voters voting for more than one candidate or for none at all, losing tens of thousands of votes for Gore. Republicans objected– more quietly, since they were leading– that many of the Democratic votes cast illegally by felons or people not registered to vote, and that Gore was trying to create votes by inventing new counting rules and getting dishonest judges to invalidate overseas absentee ballots.

Well, what about flipping a coin? Would it be good to have a policy of doing this whenever the margin of victory was small? Whether a policy is good depends on the objective, of course. A mathematics journal is no place to discuss political philosophy, so I will take as given the conventional objective: to maximize the probability that the candidate desired by a majority of those legally voting wins.

The reason usually given for a coin toss is that the voting procedures have random error, so that if the margin were close and the election were repeated, a different candidate might well win. This is a bad reason, as we will see below. If the error is unbiased, then the "official" vote count is an unbiased estimator of the "legitimate" vote, and adding noise to an estimator cannot help, although the higher the variance of the estimator, the less the noise will hurt. And, of course, suggesting a coin toss only *after* the official count is known is hardly playing fair.

If, however, we are setting up a voting rule before we know who will have the winning margin, there are indeed situations where a coin toss could help. This will be the case if we predict that the official count will subject to fraud of some kind.

2. The Model

Let us imagine that we are constructing rules for elections between a dishonest candidate and an honest candidate. In advance, we do not know who will be honest and who will be dishonest, so we cannot use a rule such as "The dishonest candidate wins only if his margin is at least 500 votes. Otherwise the honest candidate wins." We can, however, use a rule such as "A candidate wins if his margin is at least 500 votes. If the margin is less, the election is decided by a coin toss."

It of course would often be realistic that neither or both candidates are dishonest. In the model below, what will matter is the difference between their dishonest vote gains, so the reader should understand "dishonest candidate" to mean "more dishonest," and his illegal votes to be his superiority in number of illegal votes.

Denote the dishonest candidate's margin of votes (votes for him minus votes for the honest candidate) by m, his margin of legal votes by x, and the number of illegal votes by N. Both m and x can be negative, indicating a positive margin for the honest candidate, and $m = x + N$.

Let x be distributed by density $f(x)$ with cumulative density $F(x)$. We will make x is a continuous variable for neatness, so the probability of exact ties will be zero and will not need to clutter the analysis with special rules for tie-breaking.

Assume:

(A1) The true winning margin density $f(x)$ is strictly increasing in the range $[-2N, 0]$.

Figure 1 shows a number of densities which satisfy assumption (A1). Figure 1a is a well-behaved density of the kind I think most applicable. The density is greatest at $x = 0$, meaning a tie is the mode, and declines symmetrically on each side, but not to infinity, since there are only a finite number of voters. Figure 1b shows a bimodal asymmetric density where the mode has the dishonest candidate winning by large margin. Figure 1c shows a density in which the honest candidate has a solid base that enables it to win by a particular large margin 30 percent of the time, a probability atom, but otherwise the candidates are symmetric. Figure 1d shows a density which is unimodal, but with the mode at a win for the dishonest candidate. (All four examples have bounded supports because winning margins cannot exceed the size of the voting population. but bounded support will not be necessary for the conclusions below.)

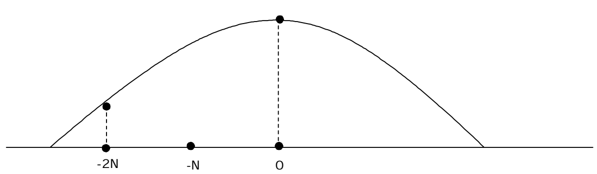

Figure 1a

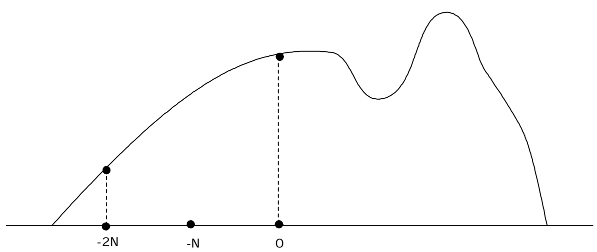

Figure 1b

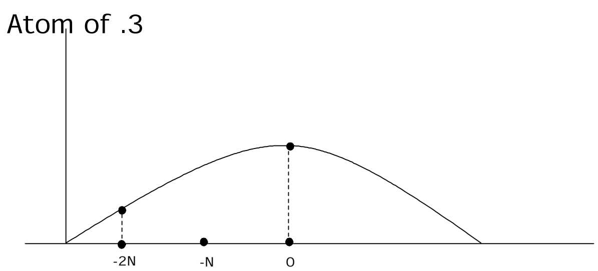

Figure 1c

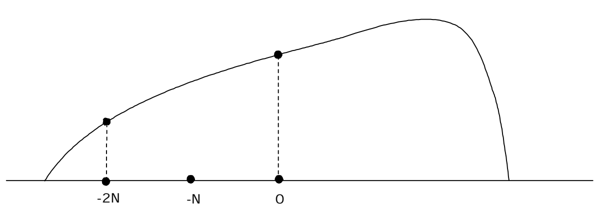

Figure 1d

Figure 2 shows two distributions that do not satisfy assumption A1. In Figure 2a, the distribution is uniform, so $f(x)$ is constant rather than decreasing. In Figure 2b, the distribution's support is less than $2N$, so the density is constant at 0 for part of the interval $[-2N, 0]$.

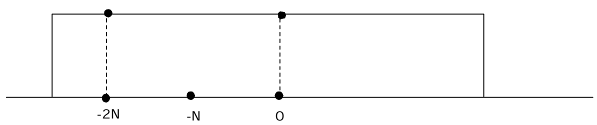

Figure 2a

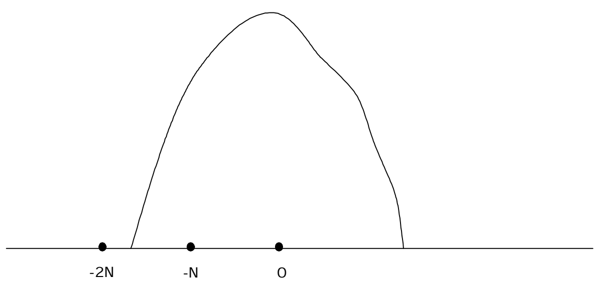

Figure 2b

Let us denote a victory for the dishonest candidate by $V = 1$ and a victory for the honest candidate by $V = 0$. Our problem is to choose a "victory rule": a rule which awards victory to one candidate or the other.

Assume society's objective is to maximize the probability of a legitimate victory, defined as the candidate with the most legal votes being declared the victor. We will denote a legitimate victory by L, where

$$\begin{aligned} L &= 1 \; if \; x \geq 0 \; and \; V = 1 \\ &= 1 \; if \; x < 0 \; and \; V = 0 \\ &= 0 \; otherwise. \end{aligned} \tag{1}$$

If society knew which candidate was dishonest, which we have ruled out, the optimal victory rule would simply replicate the objective by subtracting N votes from the dishonest candidate's margin and declaring as winner whoever had the most legal votes, i.e.,

The Full-Information Rule. $V = 1$ if $m - N \geq 0$; and V= 0 otherwise.

We will require, however, that any victory rule be symmetric, since we do not know the identity of the dishonest candidate in advance:

Symmetry Requirement. If $V(m) = p$, then $V(-m) = 1 - p$.

The conventional victory rule is:

The Conventional Rule. $V = 1$ if $x + N \geq 0$; and V=0 otherwise.

This is a special case, with $T = 0$, of the following:

The Coin Flip Rule. $V = 1$ if $T \leq x + N$; $V = 0$ if $T \leq -T$; and $V = .5$ otherwise.

The probability the dishonest candidate wins under the conventional victory rule is

$$Prob(m > 0) = Prob(x + N > 0) = Prob(x > -N) = 1 - F(-N). \qquad (2)$$

The probability the dishonest candidate is the legitimate winner is

$$Prob(x > 0) = 1 - F(0). \qquad (3)$$

The probability the honest candidate wins is

$$Prob(m < 0) = Prob(x + N < 0) = Prob(x < -N) = F(-N). \qquad (4)$$

Expression (4) is also the probability that the honest candidate wins legitimately, since he never wins except by having a majority. The probability of a legitimate victory is thus

$$1 - F(0) + F(-N). \qquad (5)$$

The probability of a legitimate victory decreases in N, since bigger N means smaller $F(-N)$.

Now let's look at the coin flip rule. The probability the dishonest candidate is the legitimate winner is not just the probability he is legitimate, because sometimes, due to the coin toss, he fails to win even if he is legitimate. The probability he is the legitimate winner and also wins under the victory rule is

$$\begin{aligned}&Prob(x>0, m>T) + .5prob(x>0, -T<m<T)\\&= Prob(x>0, x+N>T) + .5prob(x>0, -T<x+N<T)\\&= Prob(x>0, x>T-N) + .5prob(x>0, -T-N<x<T-N)\end{aligned} \tag{6}$$

The probability the honest candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned}&Prob(x<0, m<-T) + .5prob(x<0, -T<m<T)\\&Prob(x<0, x+N<-T) + .5prob(x<0, -T<x+N<T)\\&= Prob(x<0, x<-T-N) + .5prob(x<0, -T-N<x<T-N)\end{aligned} \tag{7}$$

We need to consider two cases: $T>N$, and $T<N$.

(1) $T>N$. The probability the dishonest candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned}Prob(Dis.\ leg.\ win) &= Prob(x>0, x+N>T) + .5prob(x>0, -T<x+N<T)\\&= Prob(x>0, x>T-N) + .5prob(x>0, -T-N<x<T-N)\\&= Prob(x>T-N) + .5prob(0<x<T-N)\\&= [1-Prob(x<T-N)] + .5[Prob(x<T-N) - Prob(x<0)]\\&= 1 - F(T-N) + .5[F(T-N) - F(0)]\\&= 1 - .5F(T-N) - .5F(0)\end{aligned} \tag{8}$$

The probability the honest candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned}Prob(hon\ leg.\ win) = \; & Prob(x < 0, x + N < -T) + .5prob(x < 0, -T < x + N < T) \\ & = Prob(x < 0, x < -T - N) + .5prob(x < 0, -T - N < x < T - N) \\ & = Prob(x < -T - N) + .5prob(-T - N < x < 0) \\ & = F(-T - N) + .5[F(0) - F(-T - N)] \\ & = .5F(-T - N) + .5F(0).\end{aligned} \tag{9}$$

The probability of a legitimate victory is thus

$$\pi = [1-.5F(T-N)-.5F(0)]+[.5F(-T-N)+.5F(0)] = 1-.5F(T-N)+.5F(-T-N) \tag{10}$$

The optimal T maximizes this. The first order condition is

$$d\pi/dT = -.5f(T - N) - .5f(-T - N) = 0. \tag{11}$$

Expression (11) cannot be solved. The derivative is negative for all T in the interval $[N, \infty]$ that we are considering so the smaller T is, the better. Thus, the optimum is $T^* = N$ if it is in this interval.

(2) $T < N$. The probability the criminal candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned}Prob(dis\ legit\ winner) \; & = Prob(x > 0, x + N > T) + .5Prob(x > 0, -T < x + N < T) \\ & = Prob(x > 0, x > T - N) + .5prob(x > 0, -T - N < x < T - N) \\ & = Prob(x > 0) + 0 \\ & = [1 - Prob(x < 0)] \\ & = 1 - F(0)\end{aligned} \tag{12}$$

The probability the honest candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned} Prob(Honest\ legit\ winner) &= Prob(x < 0, x + N < -T) + .5prob(x < 0, -T < x + N < T) \\ &= Prob(x < 0, x < -T - N) + .5prob(x < 0, -T - N < x < T - N) \\ &= Prob(x < -T - N) + .5prob(-T - N < x < T - N) \\ &= F(-T - N) + .5[F(T - N) - F(-T - N)] \\ &= .5F(-T - N) + .5F(T - N). \end{aligned} \quad (13)$$

The probability of a legitimate victory is thus

$$\pi = 1 - F(0) + .5F(-T - N) + .5F(T - N). \quad (14)$$

The optimal T maximizes this. The first order condition with respect to N is

$$d\pi/dT = -.5f(-T - N) + .5f(T - N) = 0. \quad (15)$$

The derivative in (15) is always positive, because $f(-T - N)$ is always less than $f(T - N)$, as shown in Figure 3. Both winning margins x are negative numbers in this case, but $T - N$ is closer to 0, where the density is greater under our assumptions.

Thus, $T^* = N$ is the optimum.

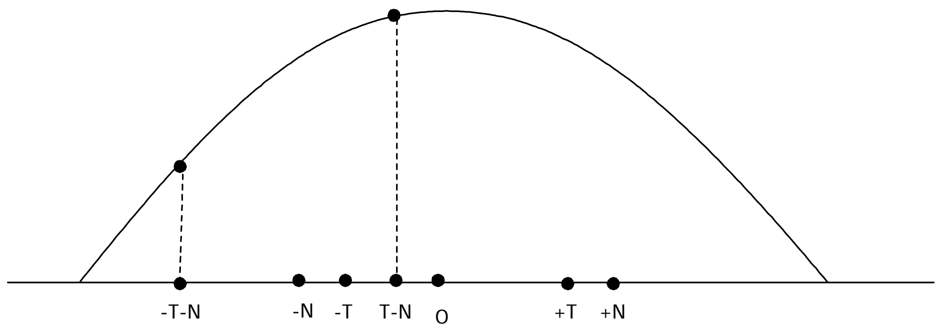

Figure 3

We can conclude that when we think that one candidate will have N illegal votes, the optimal coin flip rule flips a coin if the margin of victory is less than N.

*A Bayesian Approach.*This result can also be interpreted in Bayesian terms. If society observes margin m, what should its posterior belief be of the probability that the legal margin x is also positive? In this model, on observing $m = m'$, society knows that either (a) $x = m' - N$ or (b) $x = -m' - N$, depending on which candidate is the dishonest one. If $m' > N$, then in case (a), $x > 0$, and in case (b), $x' < 0$, so the posterior should be that with probability 1 the apparent winner is the legitimate winner. This is why T^* should not exceed N.

If $m' \in [-N, N]$, then society cannot deduce with certainty who was the legitimate winner. In that case, if the apparent winner is the dishonest candidate, the apparent winner is not legitimate, but if it is the honest winner, the apparent winner is indeed legitimate. The posterior probability that the apparent winner is the legitimate winner is, by Bayes's Rule,

$$P(m') = \frac{f(-m' - N)}{f(m' - N) + f(-m' - N)}. \tag{16}$$

If $P(m)$ is greater than .5– i.e., if $f(-m' - N) > f(m' - N)$– then victory ought to be awarded to the apparent winner, and otherwise to the other candidate. Assumption A1 tell us that is false, however, because both $-m' - N$ and $f(m' - N)$ are in the interval $[-2N, 0]$ over which the density is increasing. Thus, for margins between 0 and N, our posterior is that the apparent winner is probably *not* the legitimate winner!

Specific numbers may make this clearer. Suppose $N = 500$, and the winning margin is 100. If the dishonest candidate is the apparent winner, with $m = 100$, then $x = -400$, and we would like a rule that reverses his victory. If the honest candidate is the apparent winner, with $m = -100$, then $x = -600$, and we want a rule that confirms the apparent winner. Which is more probable, $m = 100$ or $m = -100$? It is $m = 100$ that is more probable, because it arises when $x = -400$, which is more probably than $x = -600$ given assumption (A1). In short: if a candidate wins by too few votes, the most likely explanation is that he actually lost the legal vote and only flipped the result by virtue of illegal votes.

This suggests that the following victory rule is superior to the coin flip rule both in maximizing the objective function and in perversity.

The Reversal Rule. $V = 1$ if $m \in [-N, 0]$ or $m > N$; $V = 0$ otherwise.

We have seen that the optimal rule has $T^* = N$. Let us compare the optimal Coin Flip Rule with the Reversal Rule using the following general rule (called "general" only for convenience; note that it takes the threshold

$T = N$ as given).

The General Rule. $V = z$ if $m \in [-N, 0]$; $V = 1$ if $m > N$;$V = 1 - z$ if $m \in [0, N]$; $V = 0$ if $m < -N$.

If $z = .5$, the General Rule is identical to the optimal Coinflip Rule; if $z = 0$, it is identical to the Reversal Rule. Let us determine the optimal level of z.

The probability the dishonest candidate is the legitimate winner and wins under this victory rule is

$$\begin{aligned} & zProb(x > 0, -N < m < 0) + (1-z)Prob(x > 0, 0 < m < N) + Prob(x > 0, m > N) \\ & = Prob(x + N > N) \\ & = Prob(x > 0) \end{aligned} \tag{17}$$

Equation (17) is telling us that if the dishonest candidate wins legitimately, the General Rule always awards him victory, so z is irrelevant to his probability of being the legitimate winner and also winning under this victory rule.

The probability the honest candidate is the legitimate winner and also wins under the victory rule is

$$\begin{aligned} & Prob(x < 0, m < -N) + (1-z)Prob(x < 0, -N < m < 0) + zProb(x > 0, 0 < m < N) \\ & = Prob(x + N < -N) + (1-z)Prob(-N < x + N < 0) \\ & = Prob(x < -2N) + (1-z)Prob(-2N < x < 0) \end{aligned} \tag{18}$$

Thus, the probability of the legitimate winner winning under the General Rule is

$$Prob(x > 0) + Prob(x < -2N) + (1-z)Prob(-2N < x < 0), \tag{19}$$

which is clearly maximized by setting $z = 0$ and using the Reversal Rule.

The Reversal Rule is odd because if $N = 500$ and the dishonest candidate knew he was going to have a "winning" margin of 100 votes, he would do well to throw away 150 votes. But in our model, the dishonest candidate cannot do that. He obtains the N illegal votes before he discovers the winning margin on election day, and he cannot give them back.

The optimality of the Reversal Rule is also counterintuitive because the objective function in this problem is out of the ordinary. Voting is a winner-take-all tournament, not an attempt to measure the winning legal margin with minimal mean squared error. This is best seen by comparison with a similar problem. Suppose we have a scale that we know is either 40 or -40 milligrams off, with equal probability, and we are measuring an object from a population whose weights are unimodally and symmetrically distributed with mean 5000 milligrams. Our measurement is 5010 milligrams. We deduce that the true weight is therefore either 5050 or 4070 milligrams. Typically, our objective is to come up with an estimate for the weight which is unbiased with minimum variance, or perhaps which might be biased but has minimum mean squared error. In both cases, the estimate would be somewhere between 4070 and 5000 milligrams, since 4070 is more probable than 5050 as the true weight, but 5050 also has positive probability. If, however, our objective was to maximize the probability of estimating the weight absolutely correctly, or to maximize the probability of choosing an estimate in the correct interval $[0, 5000]$ or $[5000, \infty]$ our best estimate would be 4070. It is this second kind of objective that was assumed for the election problem.

3. Conjectures

I leave the following conjectures to readers who find this model interesting.

Conjecture 1. If the dishonest candidate chooses how many illegal votes N to buy after the victory rule is chosen, that could result in either higher or lower T^* and N, depending on functional form and parameter values.

In the model of this note, the illegal vote N was fixed. Conjecture 2 makes the situation game-theoretic. This requires specifying a payoff function $U(V, N)$ for the dishonest candidate that he can use in choosing N, with U increasing in victory V and decreasing in the number of votes stolen N. The most reasonable order of play is for society to set the rule first– knowing that the dishonest candidate will react to it in choosing N– and then for the candidate to choose N for the particular election.

Conjecture 2. If the dishonest candidate chooses how many illegal votes N to buy before the victory rule is chosen, that could result in either higher or lower T^* and N, depending on functional form and parameter values.

Similarly, but less realistically, it could be that the meta-rule is that the dishonest candidate sets up his campaign to fix N first, and only the day before the election does society choose the victory rule. The order of moves in game matters (see, e.g., Rasmusen, 2, pp. 90-108), so the victory rule chosen will generally be quite different from in Conjecture 1.

REFERENCES

1. Gould, Stephen "Heads or Tails?" *Boston Globe*, November 30, 2000.

2. Rasmusen, Eric *Games and Information: An Introduction to Game Theory*, Third Edition, Oxford: Blackwell Publishers, 2001.

3. Rusin, David J. “Likelihood of Altering the Outcome of the Florida 2000 Presidential Election by Recounting,” Northern Illinois University Dept. of Mathematics, January 5, 2001, http: //www.math.niu.edu/ rusin/uses-math/recount/index.html (viewed February 23, 2001).