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

January 25, 2003

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. Even this rule will not entirely eliminate the incentive to acquire illegal votes.

Indiana University Foundation Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, BU 456, 1309 E. 10th Street, Bloomington, Indiana, 47405- 1701. Office: (812) 855-9219. Fax: 812- 855-3354. Erasmuse@indiana.edu. Php.indiana.edu/∼erasmuse. Keywords: Social Choice, Voting, Elections, Fraud, Illegal Votes, Supermajorities, Bias. JEL Classifications: C11, C44, D70, D72, D81. Copies of this paper can be found at Php.indiana.edu/∼erasmuse/papers/coinflip.pdf.

I thank John Matsusaka and participants in workshops at the business school of Indiana University and the University of Southern California for helpful comments, and Harvard Law School's Olin Center and the University of Tokyo's Center for International Research on the Japanese Economy for their hospitality.

1. Introduction

The 2000 U.S. Presidential Election was the subject of numerous 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. The 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 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 David Rusin [2001] for details.)

Many Democrats were outraged. How could Bush become president when the margin was so close? Surely there should be a revote. A number of journalists, including Stephen Jay Gould (2000), suggested, perhaps humorously, flipping a coin.

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 many of their voters voted by mistake for more than one candidate or for none at all, which cost thousands of votes for Gore. Republicans objected— more quietly, since they were leading— that many of the Democratic votes were cast illegally by felons or people not registered to vote and that Gore was trying to get getting dishonest judges to invalidate overseas absentee ballots.

What about flipping a coin? Would it be good to have a pre-set policy of doing this whenever the margin of victory was small? Whether a policy is good depends on the objective, of course. 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 conventional 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 helps. This will be the case if we can confidently predict that the official count will be subject to fraud of some kind. The model below will show why this is so.

The argument will not be based on the existing literature in economics or political science, because to my knowledge no one has studied optimal voting rules in the presence of fraud. May (1952) long ago showed the optimality of the simple majority victory rule in the basic setting which, except for the presence of illegitimate votes, I will use here. Subsequent articles in the social choice literature, while they may have considered changing the required margin for victory from 0, have not concerned themselves with illegitimate votes as a motivation (e.g., Ferejohn & Grether [1974], Garca-Laprest & Llamazares [2001]). Political science journals do publish articles on vote fraud, but these seem to be more historical, unrelated to the normative literature on voting rules (e.g., Cox & Kousser [1981], Baum & Hailey [1994]). The present article will try to link the two ideas that vote fraud can sway elections and that simple majority voting may not be the best victory rule.

2. The Model

Let us imagine that we are constructing rules for elections between two candidates who we will label as "dishonest" and "honest". 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 often 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 "the dishonest candidate" to mean "the more dishonest," and his illegal votes to be his superiority in number of illegal votes. Also, illegal votes can be interpreted not just as illegitimate additional votes for the dishonest candidate, but as legitimate votes for the honest candidate that have been illegitimately suppressed.

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 his margin of illegal votes by N. Both m and x can be negative, indicating a positive margin for the honest candidate. By definition, $m = x + N$.

Let x be distributed by density $f(x)$ with cumulative density $F(x)$. We will make x a continuous variable for neatness, so the probability of exact ties will be zero and we 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.)

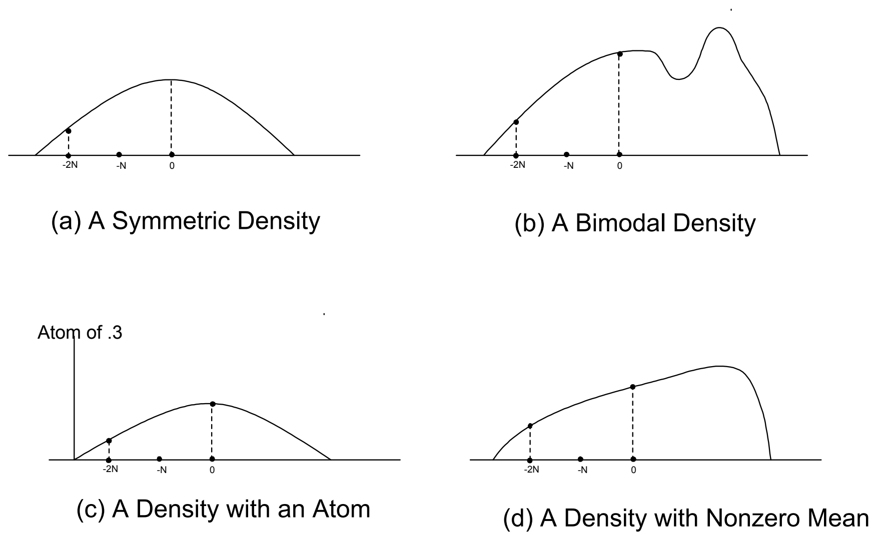

(a) A Symmetric Density (b) A Bimodal Density (c) A Density with an Atom (d) A Density with Nonzero Mean

Figure 1: Densities that Satisfy Assumption A1

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

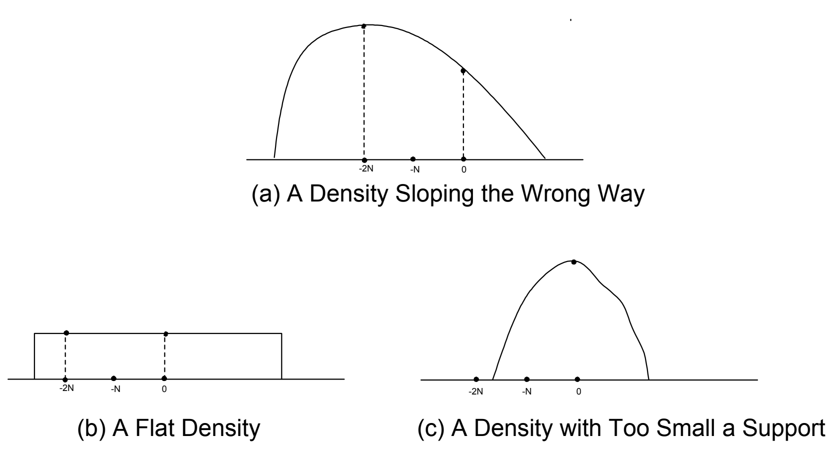

(a) A Density Sloping the Wrong Way (b) A Flat Density (c) A Density with Too Small a Support

Figure 2: Densities that Do Not Satisfy Assumption A1

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. This rule depends on the observed margin of votes, and takes the form $V(m) = p$, where m is the dishonest candidate's margin and p is his probability of victory given that margin.

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 $m \geq 0$; and V= 0 otherwise.

Under the conventional victory rule, the probability that the dishonest candidate wins is

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

The probability that the dishonest candidate is the legitimate winner is

$$Prob(x > 0) = 1 - F(0). \tag{3}$$

The probability that 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 $F(-N)$.

The conventional victory rule is a special case of the following "Coin Flip Rule", with $T = 0$. Figure 3 is a graphic illustration of how the rule works.

The Coin Flip Rule. $V = 1$ if $m \geq T$; $V = 0$ if $m \leq -T$; and $V = .5$ otherwise.

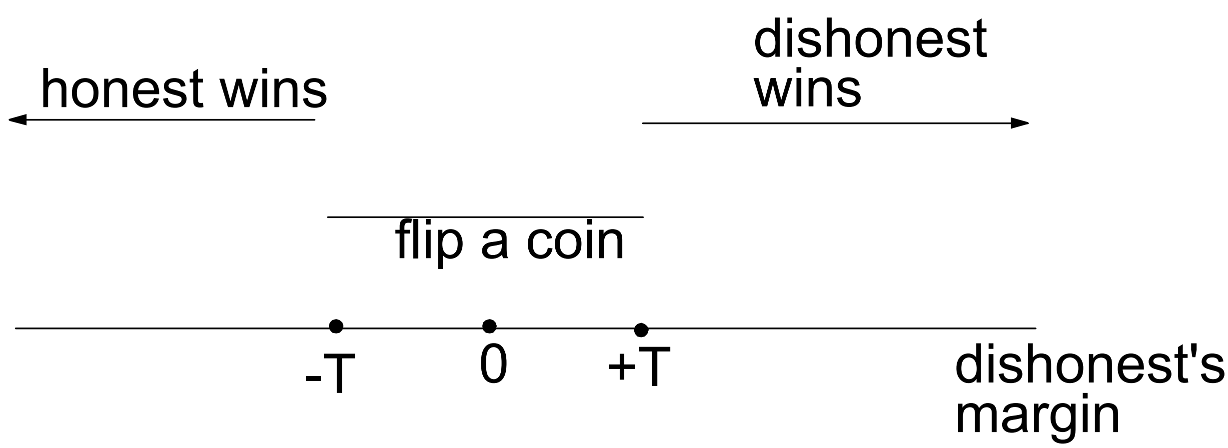

Figure 3: The Coin Flip Rule < f(T-N)

Under 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 that 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 that the honest candidate is the legitimate winner and also wins under the coin flip 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: $T \geq N$, and $T < N$.

(1) $T \geq N$ (threshold greater than the number of dishonest votes). The probability that the dishonest candidate is the legitimate winner and also wins under the victory rule is

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

The probability that 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) $T < N$ (threshold less than the number of dishonest votes). The probability that 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 that the honest candidate is the legitimate winner and also wins under the victory rule is

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

The probability of a legitimate victory is thus

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

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

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

The derivative in (15) is always positive, because $f(-T-N)$ is always less than $f(T-N)$, as shown in Figure 4. 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.

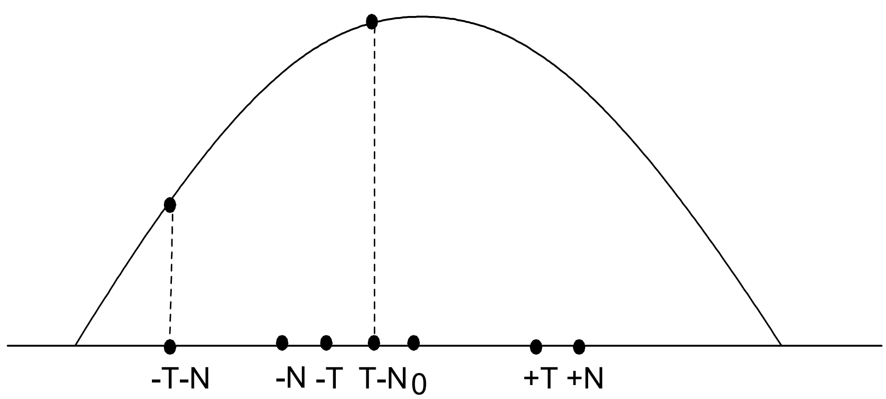

Figure 4: Why f(-T-N) < f(T-N)

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.

One implication of this is that if neither candidate has a margin of illegal votes, so $N = 0$, then a coin should not be flipped— the optimal coinflip rule never flips the coin. This reflects the idea— perhaps trivial here, but still worth pointing out—that adding noise to an unbiased estimator cannot improve it, and, indeed, results in worse decisions. Notice too that this conclusion in no way flows from risk aversion of the decisionmaker, something we have not assumed here. Rather, it flows from the increase in the probability of a wrong decision.

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? 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. Assumption A1 tells us that that is false, however, because both $-m' - N$ and $m' - N$ 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 "reversal rule," illustrated in Figure 5, is superior to the coin flip rule in maximizing the objective function.

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

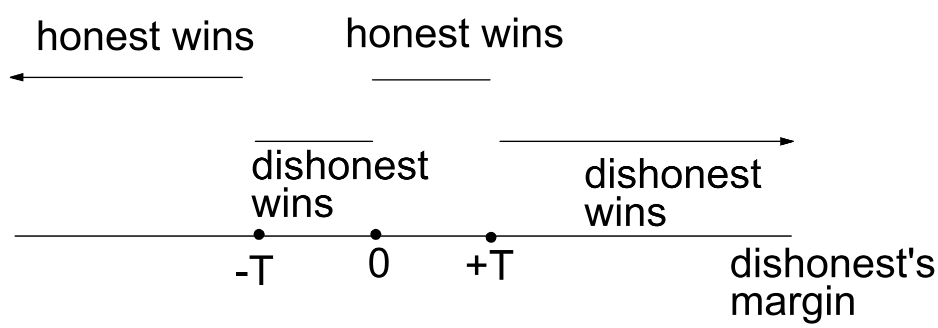

Figure 5: The Reversal Rule

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 that 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 that 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 has the peculiar implication that 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. Endogenous Vote Stealing

Suppose the dishonest candidate decides in light of the victory rule whether to incur some cost necessary to acquire the N illegal votes? What will be the effect of the various rules on his incentives to steal?

Figure 6 shows how the dishonest candidate's margin of observed votes changes as a result of adding the N illegal votes. Adding N votes to his margin shifts the support of the margin distribution to the right. Under the conventional victory rule, this increases his probability of victory, the area under the density to the right of 0, by area B.

Under the coin flip rule, the extra N votes give the dishonest candidate a gain of .5 (A+B) in the probability of winning, because if the margin falls in the interval $[-N, N]$, which has probability A+B, a coin will be flipped, but without vote stealing all of that probability would have gone to the honest candidate. The gain of .5(A+B), however, is less than B because assumption A1 implies that A is less than B. Thus, the coin flip rule has reduced the dishonest candidate's incentive to steal votes.

Under the reversal rule, the extra N votes give the dishonest candidate a gain of A in the probability of winning, because if the margin falls in the interval $[-N, 0]$, which has probability A, he will win, but without vote stealing all of that probability would have gone to the honest candidate. The area B was part of the honest candidate's probability of winning when there was no stealing, and remains part of his probability of winning when there is stealing but the reversal rule is in effect. The gain of A is less than .5(A+B), given assumption A1, so the reversal rule has reduced the dishonest candidate's incentive to steal votes compared to the coinflip rule or the conventional rule. Even the reversal rule, however, still leaves the dishonest candidate better off stealing votes than not stealing them, unless the cost of stealing votes is too high.

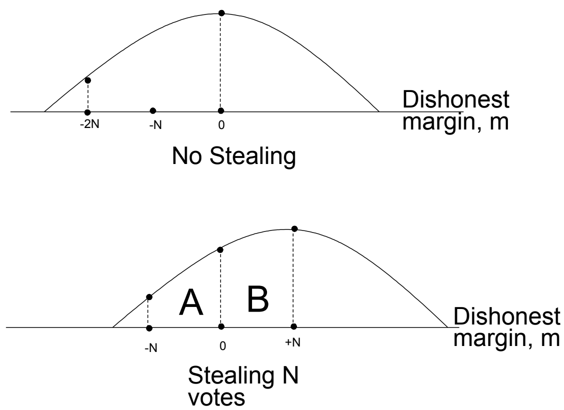

Figure 6: The Benefit of Stealing N Votes

Another possibility would be to steal (or suppress) additional votes after the margin arising from the distribution $f(\cdot)$ is known— as with the much-storied late returns from safe wards in Chicago. This adds to the attractiveness of the coin flip rule compared to either the convention or reversal rules. The coin flip rule has the advantage of weakening incentives to acquire more votes, whether by fair means or foul. The expected payoff still changes discontinuously, at $-N$ and N, but the changes are not as large as with the other rules. Consider the incentive for the dishonest candidate to falsify enough votes to possibly change the outcome under the reversal rule. He can move from zero probability of victory to certainty of victory by adding to his vote total to move from slightly below $-N$ to slightly above or from slightly below N to above; or by reducing his vote total to move from slightly above 0 to slightly below. Under the coin flip rule, on the other hand, there are only two situations in which he can change his probability of victory, and both are smaller changes. He can move from zero probability of victory to a fifty percent probability by adding to his vote total to move from slightly below $-N$ to slightly above, and he can move from fifty percent probability of victory to certainty by adding to his vote total to move from slightly below N to slightly above. Thus, there are vote-stealing cost levels such that he would steal additional votes after the election under the reversal rule but not under the coin flip rule.[1]

4. Concluding Remarks

This analysis has arrived at a paradox: when it is known in advance that one candidate is stealing votes, but not which one, it is optimal under simple conditions to reverse the election and award victory to the apparent loser. There is a simple intuition that helps one to understand the paradox. If we knew that the dishonest candidate would win with certainty under the conventional victory rule of having a margin of at least one counted vote, then we would do better by awarding victory to the apparent loser. More generally, if the counted margin is less than the number of votes we think have been stolen, then that margin was more likely than not acquired by fraud, and we should also award victory to the apparent loser.

This intuition survives relaxing the model to allow for N, the number of illegal votes, to be stochastic or not known to the social planner with certainty. A more interesting extension of the model would be to allow the number of stolen votes to vary along the continuum rather than just equalling 0 or N. One might then investigate what number of votes would be stolen if they had to be stolen (a) before the victory rule was chosen, or (b) after the victory rule was chosen. I have chosen to keep the present model simple, however, since its setting is plausible, though special; a U.S. state might, for example, have one large city which if controlled by one candidate will at low cost generate a certain number of illegal votes. A state in which the number of votes that might be stolen in equilibrium varies more smoothly is an equally plausible case. I have chosen the simpler case because the reversal

[1]Note that under the conventional rule, the dishonest candidate has incentive to steal additional votes in one situation: if the margin is slightly below 0. This is a move from zero to one hundred percent probability, and the probability of that close a margin may be larger than the sum of the probabilities of margins near $-N$ and $+N$, however, so it is hard to say much about the temptation to steal votes then compared to under the coin flip rule.

rule's optimality is a striking enough result that it ought not to be obscured by other complications.

Is there any chance of the reversal rule being adopted in real elections? One's immediate response is "Of course not!", but it interesting to ask why not. I will not give a satisfactory answer, but I will speculate in the hope that others may give a better answer. A first objection to the reversal rule is that N differs across elections and we do not know at what level to set it in advance. It is not hard, however, to see that the result in this article could survive uncertainty over the size of N. The threshold could be set at our best guess for the typical election. Even if we knew N, however, we might still find the reversal rule objectionable— even if by "we" is meant people who understand the present argument. The problem may be that we do not simply want to maximize the probability of a legitimate victory. Instead, we also wish to avoid sometimes accidentally rewarding a candidate who has stolen votes, even if our victory rule ends up hurting him on average. But why this is so, and whether it is for good reason, I cannot say.

REFERENCES

Baum, Dale & James L. Hailey (1994) "Lyndon Johnson's Victory in the 1948 Texas Senate Race: A Reappraisal," *Political Science Quarterly*, 109:595-613 (Autumn 1994).

Cox, Gary W. & J. Morgan Kousser (1981) "Turnout and Rural Corruption: New York as a Test Case," *American Journal of Political Science*, 25:646-663 (November 1981).

Ferejohn, J.A. & D.M. Grether (1974) "On a Class of Rational Social Decisions Procedures," *Journal of Economic Theory*, 8: 471-482.

Garca-Lapresta, Jos Luis & Bonifacio Llamazares (2001) "Majority Decisions Based on Difference of Votes," *Journal of Mathematical Economics*, 35: 463-481 (June 2001).

Gould, Stephen J. (2000) "Heads or Tails?" *Boston Globe*, November 30, 2000.

May, K.O. (1952) "A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision," *Econometrica*, 20: 680-684.

Rusin, David J. (2001) "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 (February 23, 2001).