LETHAL INCOMPETENCE: VOTERS, OFFICIALS, AND SYSTEMS

Jonathan Bendor; John G. Bullock

PLEASE SCROLL DOWN FOR ARTICLE

Jonathan Bendor and John G. Bullock

LETHAL INCOMPETENCE:
VOTERS, OFFICIALS, AND SYSTEMS

ABSTRACT: *The study of voter competence has made significant contributions to our understanding of politics, but at this point there are diminishing returns to the endeavor. There is little reason, in theory or in practice, to expect voter competence to improve dramatically enough to make much of a difference, but there is reason to think that officials' competence can vary enough to make large differences. To understand variations in government performance, therefore, we would do better to focus on the abilities and performance of officials, not ordinary citizens.*

Hurricane Katrina was a natural disaster; the incompetent responses of various governments to Hurricane Katrina were man-made. Many observers (e.g., Kweit and Kweit 2006; Lewis 2008; Schneider 2005) have concluded that the conduct of the Federal Emergency Management Agency, the State of Louisiana, and the City of New Orleans was feckless. Their poor performance might have cost hundreds of lives. The point is clear: the competence of officials matters.

What about the capacities of voters? Aren't they fallible, too?

Of course they are. But our thesis is that students of political behavior have spent too much time scrutinizing voter competence. We will argue that in order to understand variations in systemic performance, and to

improve it, we'd do better focusing on the abilities and performance of officials, not ordinary citizens. It is time to shift our gaze.

Our position might be controversial. To avoid misunderstandings that could transform it from controversial to absurd, let us immediately say that, of course, if most voters were Brookings-level policy wonks, the performance of the system would be quite different. But we don't think that average voter competence explains much of a system's temporal variations in performance (or performance differences across countries).[1] Further, we think that trying to improve the competence of officials is generally a more effective way to improve systemic performance than is enhancing voter competence.

Our argument is composed of the following claims. (In this essay we restrict them to the United States, but we suspect that they hold for most democracies.)

- *Claim 1:* (a) The political sophistication of American voters, measured by how much they know about politics and how well they think about what they know, has not changed much in the last fifty years. Neither has their competence, as measured by how often they make mistakes. (b) Their sophistication and competence will probably change little in the foreseeable future.
- *Claim 2:* The system's performance is typically only very modestly responsive to modest across-the-board changes in voter competence.
- *Claim 3:* The competence of elected and appointed officials *can* vary substantially in short time spans, e.g., between administrations.
- *Claim 4:* The system's performance is typically quite responsive to modest changes in the competence of key officials.

Together, Claims 1-4 form the backbone of an argument for why we should pay less attention to the abilities and performance of voters, and more to those of officials.

The rest of the essay will flesh out the argument and then explore several implications. Throughout, our central normative criterion is systemic performance, not individual competence. Of course, the latter affects the former, so we care about individual competence instrumentally. We will not, however, wring our hands about how little the average American voter knows about politics or how often individual voters err. In this respect we adopt the viewpoint of reliability engineers who take component properties more or less as given. If this be minimalism, let us make the most of it!

Sophistication versus Competence

Before we plunge into the heart of the argument, we must distinguish two terms: *sophistication* and *competence*. Sophistication is the extent and organization of political knowledge. A political sophisticate knows a lot about politics and sees many connections among the things that he knows. Competence, by contrast, is the probability that a person chooses, among all the alternatives available to him, the one that best serves his desires. Those desires may be good or bad, selfish or altruistic, typical or eccentric; we pass no normative judgment on them.[2] For brevity, we use p as shorthand for competence; for example, if $p = .6$ for a group of voters, the probability that those voters vote for the candidate who best serves their desires is .6. For two-candidate elections, a voter is *incompetent* if $p < .5$.

Our definition of competence has two immediate implications. First, because competence, as defined here, is not about people's reasons for choosing what they do, we avoid phrases such as "competent analysis" or "competent reasons." Second, competent choices may not be "good" in any absolute sense. They may simply select the best alternative from a set of bad alternatives.

Our focus in this essay is on competence in a specific setting: two-candidate elections for general political offices. Because this choice problem is cognitively easier than multi-candidate primaries or elections (Jackman and Sniderman 2002, 218), it is the context in which voter competence is likely to be greatest. The candidates hail from different parties and stand on different platforms. To be competent in a general election in a two-party system, voters need not determine the best candidate in a crowded field, but only determine which candidate's victory is more likely to serve their desires. Nevertheless, it is important to recognize that Claims 1, 3, and 4 do not presume any particular distribution of competence in the electorate. This issue is relevant for Claim 2, but as we will make clear, even Claim 2 does *not* presume that the average voter has any particular competency level at a given point in time. Voters may be very good at picking the candidate who will better serve their desires or they may be very bad at doing this. Claim 2 will hold either way.

Claim 1: The sophistication and competence of American voters have changed little in the last 50 years, and will probably change little in the foreseeable future.

The modern history of political ignorance is not much of a story: the plot almost never changes. Taking stock of voting behavior in the 1950s, the authors of *The American Voter* found an electorate that was not just ignorant of major policy debates, but unaware of the very existence of those debates (Campbell et al. 1960, 170). Surveying the next thirty-five years' worth of findings, Michael X. Delli Carpini and Scott Keeter (1996, 116-17) found no noteworthy change in aggregate levels of sophistication. Their 1989 Survey of Political Knowledge—still the most extensive survey of what Americans know about politics—included fifteen questions akin to those asked in opinion surveys of the 1940s and 1950s. Nine of the items were answered correctly more often in 1989 than they had been in the immediate postwar era; six were answered correctly less often. The median change in the percentage answering correctly was +4.5 percent, which is more than a trivial increase. But note immediately that it is an increase over half a century, and perhaps because of this, Delli Carpini and Keeter maintain that "the overall similarity [in percentages responding correctly] is more striking than the differences."[3] Martin Gilens, Lynn Vavreck, and Martin Cohen (2007, 1171-72) reach the same conclusion in their analysis of questions about presidential candidates' policy positions that have been asked regularly since 1972 in the American National Election Studies. And both conclusions echo intervening studies by Donald Kinder and David Sears (1985), W. Russell Neuman (1986), Stephen Earl Bennett (1988 and 1989), and Eric R. A. N. Smith (1989).

To be sure, specific incidents can produce sharp increases in knowledge of specific facts: witness the massive increase, between 1945 and 1946, in Americans' knowledge of their country's role in the United Nations (Delli Carpini and Keeter 1996, 118, 129). But these are isolated incidents, and the gains are likely to be temporary. Today, just as in the 1950s, most Americans are unaware of the existence of most issues, perforce oblivious of parties' stands on those issues, ignorant of most of the basic rules by which government operates, and unfamiliar with all but the most important handful of events in the nation's history. This is not to say that most Americans are *completely* ignorant of politics; indeed, only a "chemical trace" of the electorate would score absolute zero on a Kelvin scale of political knowledge (e.g., Converse 1990, 372-73). But the overall picture is one of overwhelming ignorance, and it has been this way for half a century.

Now and then, optimistic studies purport to show that the public is becoming more sophisticated (Nie, Verba, and Petrocik 1976) or that

previous research has understated the sophistication of ordinary voters (Mondak 2001), but these arguments have not fared well under scrutiny. (On Nie, Verba, and Petrocik 1976, see Luskin 1987; Smith 1989; and Kinder and Sears 1985, 666-67. On Mondak 2001, see Luskin and Bullock 2006.)

Existing studies also provide no reason to expect change in the near or intermediate future—not in America, nor in any other developed democracy.[4] Optimists may suggest that increasing levels of formal education are likely to produce a more knowledgeable citizenry, but the public grew gradually more educated throughout the twentieth century without any obvious increase in political knowledge (National Center for Education Statistics 1993); years spent in school have only a weak causal link to political sophistication (Luskin 1990; Cassel and Lo 1997; Nie, Junn, and Stehlik-Barry 1996, chs. 7-8). In a similar vein, we might imagine that the proliferation of media outlets carrying political news would increase average political sophistication; but in reality, increasing specialization of media outlets has made it easier for most citizens to ignore politics (Prior 2007), and self-reported media consumption has little connection to political knowledge (Luskin 1990; Delli Carpini and Keeter 1996, 144-45). Writing in these pages, Ilya Somin (1998) suggested that low sophistication is a byproduct of government size: government now does so much that it is difficult for anyone to become sophisticated about more than a small fraction of its activities. This may be right; what is certain is that radically smaller American government is not in the offing.

The studies cited here show that low levels of political sophistication are stable, but Claim 1 is about competence. Might average competence increase substantially in the future, even as average sophistication remains at roughly the same level? We know of few studies that bear directly on the question. Delli Carpini and Keeter (1996, ch. 6, esp. 254-58) argue that the relation between issue positions and vote choice is weakest among those who know least, strongest among those who know most. This suggests, unsurprisingly, that competence and sophistication are related—but how closely? Larry Bartels's (1996) analysis of information effects in the six presidential elections held between 1972 and 1992 speaks to this question. Using national survey data, Bartels estimates the difference between voters' actual probability of voting Republican and the probability that they would have voted Republican if they had been fully informed about which candidate better served their desires.[5] He finds

that some would have been more likely to vote Republican, others less likely; on average, fully informing the voters would have changed their propensity to vote Republican by 8 to 12 percent, depending on the election. These are very substantial increases in competence. But they are the products of colossal increases in the sophistication of the electorate. Much smaller increases in sophistication—the sort of increases that might realistically be expected even from Herculean civic education efforts— would produce far smaller increases in competence. The still smaller changes that may result from the slight over-time variation in sophistication will only produce even smaller changes in competence.

Claim 2: The system's performance is typically only very modestly responsive to modest across-the-board changes in voter competence.

Per Claim 1, we believe that the sophistication and competence of most voters will change little in the foreseeable future. But what will happen if there *are* modest improvements across the board? Both theoretical and empirical considerations are relevant. We start with the former.

Theory.   The theoretical core of Claim 2 is simple: Large numbers of voters amplify the effects of individual levels of competence, for both good and ill. On the benign side, if every voter is better than a fair coin in choosing between two candidates, then odds are good that a large electorate will pick the party that will better serve the desires of a majority of voters. This idea, stated crudely here, is an extension (Miller 1986) of Condorcet's Jury Theorem (CJT). When a group must select one of two alternatives, the CJT and its extensions show how the probability that a majority favors one of two alternatives varies with the size of the group and the competence of its members.[6]

Consider the following example: an electorate of 100,000 voters, each with competence of .6, each trying to decide between two candidates in an election. One candidate will better serve the desires of 51.5 percent of the electorate, while the other will better serve the desires of the other 48.5 percent, but no member of the electorate knows with certainty which candidate is best for him. One might think that this combination of parameter values would produce an electoral system that is quite unpredictable. After all, voters err 40 percent of the time and the minority is nearly as big as the majority, so although the candidate who better serves the majority should win more often than not, this should not occur

with great regularity. Yet it does: given these parameters, one of Nicholas R. Miller's (1986) extensions of the CJT shows that the majority faction's candidate wins 97 percent of the time. The system is very reliable; hence, it is predictable.

And because it is so reliable, Claim 2 will hold: modest improvements in individual competence will not increase systemic reliability much. In this example, boosting the competence of all citizens from .6 to .65—a nontrivial improvement—only raises systemic reliability from .97 to .999. And in an electorate of 150,000, the systemic kick must be still smaller, since with these numbers, the chance that the correct candidate is selected is .99 even with individual competence of only .60. Figure 1 shows how the system reliability in this example rises as the size of the electorate increases.

Now consider the opposite, gloomier possibility: $p < .5$ for every voter. (We will examine systematic differences in voters' competence levels—those correlated with voters' desires—shortly). Then we get a "bad" CJT: the chances are good that a majority of voters will pick the

Figure 1:   Even Under Inauspicious Conditions, the Candidate Who Best Serves Majority Desires Is Likely to Win if p > .5
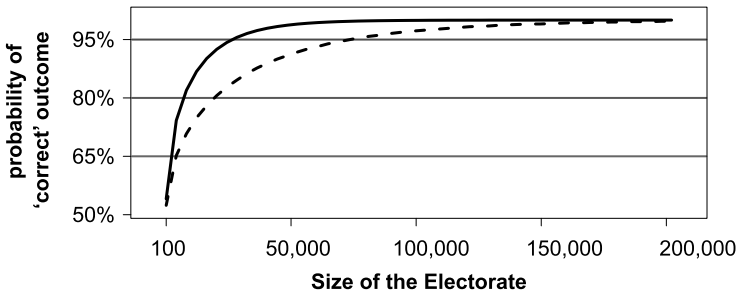
The figure shows how the probability of choosing the candidate who best serves majority desires varies with the size of the electorate. The solid line traces the majoritarian candidate's probability of victory when $p = .6$ for each voter and only 51.5 percent of voters belong to the majority bloc (against 48.5 percent who would be better off with the other candidate). The dashed line traces the majoritarian candidate's chances of victory when p remains at .6 but the majority increases to 52.5 percent of the electorate. Both lines suggest that the candidate who best serves majority desires is likely to prevail even when the electorate is rather small and individual voters are only slightly more likely than not to choose the candidate who better serves their interests.

wrong candidate (Miller 1986). Further, under conditions similar to the assumptions of the original CJT (all voters are equally competent and each makes up his mind independently), the probability of this mistake goes to one as the size of the electorate increases (Miller 1986).

   This is indeed a dismal outcome. But Claim 2 is an empirical assertion, not a normative one. Hence, it is fully consistent with the bad (Miller-extended) version of the CJT. Suppose, for example, that $p = .35$ for all citizens. As before, $N = 100,000$, and the majority faction is 51.5 percent of the total. Then the probability that the electoral mechanism selects the candidate who better serves majority desires is only .001. Now consider the effects of a widespread and rather effective effort in civic education, which boosts p to .4. This does, of course, enhance systemic reliability, but only to .03—not much of an improvement.[7]

   There is, however, a narrow parametric window of opportunity, which occurs when individual competence is close to .5. Suppose, for example, that $p = .49$, with the size of the electorate and its factional split as before. Then civic education boosts p over the critical threshold, to .51. Most people would agree, we believe, that this is indeed a modest change in competence. Yet this *does* produce a significant improvement in systemic reliability: it rises from .42 to .57.

   As noted, however, this is a rather special circumstance: voters must be about as good as fair coins. If they are considerably better or considerably worse, Claim 2 will hold. Moreover, as Figure 2 shows, the "window of opportunity" for small competence improvements to produce big changes in system reliability closes as the size of the electorate increases. The figure shows how the probability that voters choose the "right" candidate varies as a function of voter competence and the size of the electorate, under the assumption that one candidate is best for all voters. When there are only five voters, modest increases in voter competence may produce substantial increases in the probability that a majority will vote for the right candidate, regardless of the level of voter competence. But when there are 51 or 501 voters, modest increases in competence will have almost no effect on this probability unless prior voter competence is close to .5. And for larger electorates, prior voter competence must be almost exactly .5 if small increases in competence are to have any effect. When no one candidate is best for all voters, this window of opportunity closes more slowly, but for city- and state-sized electorates, it remains true that modest increases in competence will have almost no effect on this probability unless prior voter competence is close to .5.[8]

Figure 2:   The Parametric Window of Opportunity
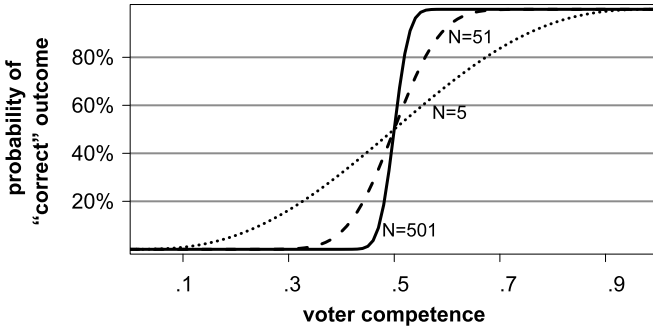
When voters choose between two candidates, and one candidate better serves their desires, the probability that a majority of voters prefer that candidate increases with the competence of those voters. Each line in the figure shows how competence affects the majoritarian candidate's probability of victory according to the Condorcet Jury Theorem. The dotted line shows that when there are only five voters, this probability increases gradually for almost all possible values of voter competence. When there are 51 voters, the dashed line shows that changes in competence within the range of [.4, .6] have large effects on this probability; outside of this range, the effects are small. And when there are 501 voters, the solid line shows that the "window of opportunity" is even narrower: modest changes in voter competence will have substantial effects on the majoritarian candidate's probability of victory only when voter competence is within the competence within the range [.45, .55]. For larger electorates, the window of opportunity will be even smaller.

Thus, there need not be any "miracle of aggregation" (Converse 1990) in order for Claim 2 to hold. Apart from a narrow window—voter competence in the neighborhood of .5—large numbers of voters do the job.[9]

However, because so much attention has given to the benign versions of CJT, it is worthwhile to make a few comments about what these versions say about elections. First, they do *not* imply that large electorates inevitably or even usually select politicians who are good in some absolute sense. Voters can only work with what's offered, and if all the candidates are mediocre then the winner will be mediocre. In this context, statistical aggregation is not magical: it cannot produce high-quality outputs out of poor inputs. (And, of course, even if the best candidate is stellar, the gloomy CJT theorems generally imply that she will lose if the competence of most voters is below .5.)

A second objection to applying standard (benign) CJTs to elections is that, contrary to the assumptions of most of those CJTs, many voters do not make up their minds independently of each other. This is certainly true, but the implications of this fact for American presidential races (for example) are less than one might think. The reason, once again, is the massive size of the electorate. Interdependence of voters' judgments reduces the effective size, but much remains. To see this, consider Figure 1 again. Individual competence is .6 throughout the figure, and the electorates that it depicts are sharply divided, with only 51.5 percent of voters in the majority faction. But the figure shows that in an electorate of 100,000 people, system reliability is .97. This electorate is small compared to that of a presidential election in the U.S.: over 100,000,000 people voted in 2004. And although interdependent judgments (i.e., positively correlated votes) reduce the effective number of independent minds working on a problem (e.g., Ladha 1992, 1995), it would take a large amount of interdependence to drive one hundred million voters down to an effective N of 100,000.

For example, consider the following stylized bloc-model of interdependence. People within each faction fall into blocs of 1,000; the voting of everyone in a given bloc is perfectly correlated. Across blocs, voting is completely independent, and each bloc has a competence level of .6. Then in an electorate of 100,000,000 people there would be 100,000 blocs, and the system's reliability is .97. Thus interdependence could cause a *thousand-fold* reduction in the number of effective minds, yet Condorcet's mechanism would still work surprisingly well.

A similar point holds for the gloomy CJTs. Plausible amounts of positive correlation among incompetent voters would still leave so many independent "minds"—albeit feckless ones—that large numbers would still have their way: the electoral system would be so unreliable that modest changes in the competence of its components would not have big effects. (This assumes, as before, that isn't extremely close to .5.)

Many scholars have focused on problems caused by voters being generally unsophisticated. Others, however, have argued that the more serious issue for democracies is asymmetries in sophistication and competence. Indeed, Miller (1986, 181) showed that if the minority is more competent than the majority and the latter's competency exceeds .5 by a sufficiently small amount, then the chance that the minority will triumph approaches 100 percent as N goes up. This is another bad-news Condorcet Jury Theorem: in this context, increasing N has the wrong effect.[10] There is no miracle of aggregation here; on the contrary.

This bad-news effect relates directly to a point underscored by Converse: the variance in voter sophistication is enormous; indeed, it is "orders of magnitudes of orders of magnitudes" (Converse 1990, 373). For Converse (and e.g., Luskin 2002, 298), this empirical pattern may point to a normative problem.[11] In a two-candidate election, when one candidate better serves the desires of a sophisticated minority and another candidate better serves the desires of a relatively unsophisticated majority, the greater sophistication of the minority can outweigh its demographic disadvantage. Thus, in races in very large electorates, such as U.S. presidential elections, the minority side may triumph. Hence, this bad-news CJT speaks directly to one of Converse's concerns. It also speaks to the unjustly forgotten summary of his 1964 article, in which he speculates on the consequences of electoral competition between sophisticated "elite parties" and unsophisticated "mass parties." Mass parties "enjoy a 'natural' numerical superiority, yet they are cursed with a clientele that is less dependable or solidary in its support." Elite parties have "a natural clientele that is more limited but more dependable" (Converse [1964] 2006b, 56). This is just the setting in which a bad-news CJT may take effect, helping the elite parties' candidates to victory even though a majority's desires would be better served by a candidate from the mass party.

Thus far, however, this point pertains to a system's absolute level of performance—in large electorates a sophisticated minority may beat an unsophisticated majority with high probability—and so it does not bear directly on Claim 2. As we have seen, bad news about absolute electoral performance—the mechanism of elections frequently failing to serve the desires of the majority—may be consistent with Claim 2, i.e., that the system's performance is typically only very modestly responsive to modest across-the-board changes in voter competence. The following numerical example shows why this combination might obtain.

To simplify the calculations, consider an electorate with infinitely many voters. Two candidates are running for office: L and C. L best serves the desires of "liberals," who make up 40 percent of the electorate. C best serves the desires of "conservatives," who are 60 percent of the electorate. If all voters were perfectly competent, C would win the election with 60 percent of the vote. But assume that all voters are imperfectly competent. In this context, a Law of Large Numbers tells us that the fraction of liberal voters who actually vote for L is exactly p_L. Similarly, the fraction of conservative voters who actually vote for C is

exactly p_C. Suppose that, in the status quo ante, $p_C = .6$ and $p_L = .7$. Then 52 percent of the electorate $((.4 \times .6) + (.7 \times .4))$ vote for L, who triumphs with probability 1.0. Then a program of civic education is launched. It is successful: p_C rises to .65 and p_L rises to .75—a substantial improvement in competence.[12] Nevertheless, C's vote share rises by only one percent, from .48 to .49. This in turn implies that in large electorates, C's chances of winning won't rise by much. Indeed, in the population limit, they don't rise at all: L still wins, though more narrowly. (In this example, C wins by means of an equal, across-the-board improvement in voter competence if and only if p_C reaches at least .7. Given an initial competency of .6, this change is arguably more than modest.) Thus, the impact on Claim 2 of the bad news regarding competence asymmetries is less than one might initially think: even when a minority faction is more sophisticated than the majority, modest across-the-board increases in competence are unlikely to change system reliability much.

Moreover, reforms producing roughly equal across-the-board improvements that don't significantly reduce competency asymmetries won't help much. What is needed to reduce competency asymmetries is a specific, targeted effort aimed at improving the competency of the majority. Claim 2 doesn't speak to such possibilities. However, we suspect that this kind of change is quite unlikely. It's not likely to arise as the result of undesigned, decentralized socioeconomic changes (e.g., technological improvements). For example, despite the hopes of techno-logical utopians, the Internet has not created a highly egalitarian virtual democracy. Even if inequality-exacerbating Matthew effects—to him who hath shall be given—don't prevail, such technological changes are available to all voters in real democracies. Hence, if such changes arise they will be across-the-board improvements rather than targeted improvements in the capabilities of the unsophisticated.

Evidence. It would be wonderful if we had a lot of data on how systemic performance responds to changes in voter competence, but unfortunately we don't.

There is some, however. Recall Bartels's (1996) empirical analysis of how voter sophistication (and hence presumably competence) affects electoral outcomes. He finds nontrivial effects on aggregate outcomes: Democratic presidential candidates do almost 2 percentage points better than they would if voters were "fully informed"; incumbents do 5 points better (Bartels 1996, 220). However, it is important to note that Bartels's

standard is a strong one: what would happen if *all* voters were *fully* informed. This is not a modest change from the status quo; it's a big one. And, strikingly, Bartels's results indicate that this awesome increase in voter sophistication would not have changed the outcome of even one of the six elections that he studied (ibid., 216, Table 3). Presumably, the modest increases in sophistication considered in Claim 2 would be even less likely to produce changes in election outcomes. Hence, Bartels's study offers indirect support for Claim 2: if titanic increases in sophistication do not change election outcomes, modest increases aren't likely to do so.

Because elections necessarily involve thresholds, it is unlikely that Claim 2 always holds. Counterexamples must occur, at least occasionally. Consider, for example, the 2000 presidential election. Careful empirical research has shown that the butterfly ballot "caused more than 2,000 Democratic voters to vote by mistake for Reform candidate Pat Buchanan" (Wand et al. 2001, 793). Since this exceeded George W. Bush's official margin of victory in Florida, the voters' confusion swung the state, and hence the nation, to Bush. If these voters' desires would have been best served by voting for Gore, as they intended, an improvement in the competence of a very small number of voters (hence a modest change indeed in average competence) would have changed the electoral outcome.

Again, however, this counterfactual involves a *targeted* improvement in a handful of voters. An equal improvement spread uniformly over the entire electorate might not have made any difference at all. In this scenario everyone would be just a tad more competent, which may not have been enough to help the confused voters in Florida. So it is unclear whether the 2000 election constitutes a genuine counterexample to Claim 2.

Claim 3: The competence of elected and appointed officials can vary substantially in short time spans, e.g., between administrations.

The general argument for Claim 3 is simple. Agencies are hierarchies, as is the executive branch in general. Hierarchies are designed to magnify the influence of a few key decision-makers: those at the top of the organization. Because they aren't numerous, there is relatively little redundancy—certainly far less than the redundancy of voters. Hence, the errors of key officials matter more than do the errors of voters. A single inept

decision-maker in a key position—e.g., Michael Brown of FEMA—can really foul things up.[13]

Further, it is vital to remember that redundancy requires independence. Kinder (1998) and others have pointed out that voters make correlated errors. Thus, there aren't 120,000,000 distinct minds trying to select the next president; in effect there are far fewer. But at least in that context we are starting from a very big *N*. Regarding presidential decision-making, the number of physically distinct heads is very small: a dozen or so. If they are like-minded their errors will be correlated, reducing the number of effective minds still further.[14]

The preceding examples were of short-term fluctuations in agencies. One can easily find examples on a different time scale. Perhaps the most significant set of cases involves the gradual improvement of governmental bureaucracies via the installation of merit-based selection and advancement.[15] Although improvements like these usually take decades, it is important to note that *one cannot find any comparable improvements in voter competency over the same time scale.*

Claim 4: The system's performance is typically quite responsive to modest changes in the competence of key officials.

Again, consider the Katrina disaster, which killed more than 1,800 people. Many of these people could have survived had the City of New Orleans evacuated citizens more rapidly. In the opinion of at least some observers (Scanlon 2006), this was feasible. The buses were available; an evacuation plan was ready. But important details weren't done right: the City didn't send enough police door-to-door in relevant neighborhoods, and it didn't give people clear instructions about what to do and where to go (Scanlon 2006, 1, 4-5). Modestly more competent leadership of key agencies might have saved hundreds of lives.[16]

The evidence here is much weaker than is the evidence for Claim 1; we know of no robust empirical regularities uncovered by multiple studies. Given this state of the art, most of our justification for Claim 4 must be theoretical. The logic, which is very simple, is based on the same empirically plausible premises that support Claim 3: governmental organizations are hierarchies, and hierarchies are designed to magnify the influence of a few key decision-makers. True, some governments—e.g., stable democracies—have built-in checks against catastrophic incompetence,[17] and parliamentary systems have procedures (no-confidence

votes) that make it possible to terminate a merely incompetent adminis-
tration that falls short of catastrophic. But no governmental system that
we know of is insulated from modest changes in executive competence.
There simply isn't enough Condorcet-like redundancy at the top of
hierarchies.[18]

A somewhat more conjectural hypothesis is that incompetence at the
top breeds incompetence lower down, via unwise personnel decisions—
an incompetence multiplier effect.[19]

Finally, small performance differences can cumulate over time.
Indeed, as growth economists point out, the cumulative effects of small
annual differences are startling. Normalize the performance of agencies A
and B to 100 at Year 0. Suppose that A picks somewhat more competent
executives than does B. Due to this difference, A's performance index
improves at 3 percent annually; B's, at 2 percent. Initially the organiza-
tions remain close: in Year 10, agency A is at 134 and B, 122. But by Year
20, A has improved to 181 while B is only at 149, and in about seven
decades A's performance index will *double* B's. A small but persistent
competency edge cumulates to impressive performance differences.

Criticisms of Our Position

A natural criticism of our position is that because incompetent leaders are
picked by voters, the latter are responsible for the former.

This statement is true, but it misses its target. We have already
acknowledged that the competence of average voters is important in
explaining a system's *absolute* level of performance. But Claims 1 and 2
are about *changes* in a democracy's performance, and the related possibil-
ity of improving that performance. The desirability of a hypothetical
transformation of the country into a nation of policy wonks doesn't
constitute a critique of either Claim 1 or Claim 2. To be sure, "if men
were angels" (Madison), cognitively as well as motivationally, things
would be different (Simon 1985). A realistic, Madisonian design should
take us largely as we are: busy, distracted, and ill-informed. Most of us are
amateurs in politics, and so we will remain. Designs based on significantly
more ambitious premises are, we believe, doomed to fail.

It is also important to realize that competent voting (as defined
here) does *not* imply that the more competent candidate will always be
selected. A standard result in spatial models of delegation (e.g., Bendor

and Meirowitz 2004) is that principals (here, voters) face a tradeoff between the competence and the ideological proximity of agents (here, politicians). Rational and fully informed principals will, in most contexts,[20] choose from a set of undominated agents: those who are maximally able, for a given degree of difference from the principals' desires;[21] and those who most closely share the principals' desires, for a given degree of competence. (More cannot be said without imposing more structure on the principals' choice problem.) In some circumstances it is rational, when faced with a choice between agent A, who is competent but has very different desires, and agent B, who is a bit less capable but much more in tune with the principal's desires, to select the less competent agent and to do so knowing that one is picking the less able person. Of course mistakes might be made—it might turn out that agent B is much more feckless than the principal had forecasted—but this is just an unlucky *ex post* error, not an *ex ante* mistake, and even fully rational decision makers cannot avoid all *ex post* errors when key parameters are uncertain *ex ante*.[22]

*   *   *

When Converse wrote "The Nature of Belief Systems in Mass Publics" in 1964, it was possible to make a career out of studying political behavior without broaching the subject of citizen competence or being influenced by others who had. No longer. Partly because of Converse, the work of everyone who studies political behavior has been shaped by what we now know about how well citizens think about politics. But precisely because the topic has been heavily studied, it surprises us to see that the field of political behavior has paid so little attention to the flip side of the coin: the competence of key officials. After all, there is no reason to believe that levels of voter competence will change much in the foreseeable future, or that any of the political outcomes that we care about will be affected by the slight variations in voter competence that we have reason to expect. But there is reason to believe that the competence of officials in key government positions *will* vary—substantially—in the foreseeable future, and that government performance will be affected by these variations.

No one should seriously argue—certainly we don't—that students of political behavior should abandon the study of ordinary voters' capacities. Still, to the extent that we study citizen competence because we

care about government performance, we will do well to devote more attention to the competence of those who hold office in government. For voters aren't the risky links in large democracies. Officials are.

NOTES

1. Regarding the latter, we believe that cross-national differences in the competence of particular leaders or in institutional characteristics (democracy versus dictatorship) have more causal clout.

2. We use the term *desires* broadly throughout the article: it denotes values, personal goals, economic self-interest, and preferences (e.g., over policy) that may be unrelated to one's personal fortunes.

3. Note too that the original surveys drew on national samples at a time when response rates regularly topped 70 percent and were sometimes much higher. The rate of response to Delli Carpini and Keeter's survey was 38 percent. On average, survey-takers are unusually knowledgeable about politics (Brehm 1993), and the Survey of Political Knowledge is therefore likely to overstate knowledge levels, both in the absolute sense and relative to the surveys of the 1940s and 1950s (but see Delli Carpini and Keeter 1991 and Keeter et al. 2000 for contrary accounts). This is further reason to believe that changes in political knowledge levels over the last half century are essentially nil.

4. Surprisingly, Converse (2006a, 311-14) is somewhat optimistic about the prospects for long-run improvement in electorates. But like us, he identifies no reasons for optimism and several strong reasons for pessimism.

5. Bartels uses survey-interviewer ratings of interviewees' informedness as a measure of their informedness; in the context of his article, the "fully informed" are those who receive the highest rating, a 5 on a 5-point scale. The implicit assumption is that such people know about the rules of government, the policies that candidates are likely to advance if they win office, and other matters relevant to determining which candidate will best serve their desires. See Bartels 1996, 203-4, for further discussion.

6. There are now many versions of this theorem. Unfortunately, many of these extensions have been ignored by students of voting behavior. As a result, many scholars in this area, including some very eminent ones such as Converse and Larry Bartels (1996), have been rather dismissive of CJTs. For example, in his *Annual Review of Political Science* essay, this is all that Converse had to say about the application of CJTs to elections:

> The Condorcet model may well reflect one force behind gains in apparent competence through aggregation. But it surely is not the most telling model. It assumes, in Bartels's words (1996), that individuals contributing to the group judgment are "modestly and equally well-informed." This does not seem a promising gambit for diagnosing the electorate, given the staggering heterogeneity of informedness across it. (Converse 2000, 349)

Another criticism often used to dismiss Condorcet's Theorem is the following: (a) CJTs presume that decisions are independent, but (b) voters' decisions are correlated, hence (c) CJTs are irrelevant to elections (e.g., Rawls 1999, 314-15). A third criticism is that CJTs are simply irrelevant to cases in which the "correct outcome" is a matter of dispute (e.g., Black [1958] 1987, 163). All of these criticisms reflect views of the CJT that are now obsolete. Most scholars working in this area have understood that Condorcet's assumptions are idealizations: they help to make the analysis tractable, but they are clearly unrealistic. So effort has been made to relax them. These efforts have been productive: it turns out that Condorcet-like conclusions continue to hold under weaker (more realistic) assumptions. Competence need not be homogeneous (Grofman, Owen, and Feld 1983) nor judgment independent (Ladha 1992 and 1995); and Miller (1986) shows elegantly that the CJT can be extended to apply to cases in which one outcome is not best for all voters.

There is a family of Condorcet models. It would be a shame for students of voting to turn their backs on the entire set based on the incorrect belief that every model in the family yields Panglossian conclusions.

7. As with any evaluation, there are different ways to frame the effect. Looking at the ratio of the success probabilities makes things look better; but as an absolute chance of success, .03 is still very poor.

8. Observe that Claim 2 says that the system's performance is *typically* only very modestly responsive to modest across-the-board changes in voter competence. "Typically" allows for the unusual effects produced by unsystematic heterogeneities in voter competence. Because such variation—i.e., uncorrelated with voters' desires—obscures the central systemic tendencies produced by large N, we do not study such differences here. It is worth noting, however, that in Condorcet's original context of common-interest problems, certain kinds of heterogeneity are relatively innocuous. In particular, decision-makers can vary significantly in competence without affecting the thrust of the benign theorem's conclusion: it is easy to show that so long as all agents are more competent than random chance, the probability that the group-majority decides correctly gets arbitrarily close to 1 as N increases without limit (e.g., Grofman et al. 1983, 268). It is not heterogeneity *per se* that perturbs the conclusion: it is variation big enough to include incompetence.

9. As noted, this part of the argument holds most crisply when competence-heterogeneity is restricted to one or the other side of the .5 threshold. This condition is probably violated in most electorates. In, e.g., American presidential races, probably *some* of the voters are incompetent, while others are undoubtedly competent. We suspect that the thrust of the argument in the text will still hold—systemic reliability will improve significantly in response to modest across-the-board improvements in voters' competence only if, *ex ante*, there are many voters who are in the competency-neighborhood of .5—but this awaits demonstration.

10. Note that this is true even though the benign competency condition—decision-makers are better than fair coins—holds for all voters.

11. However, Luskin's summary of what's known indicates that the partisan or ideological effects of errors isn't clear.

12. As Miller's work indicates, it is the *ratio* of the two competency probabilities that is decisive, and an equal across-the-board improvement must improve the ratio of the less sophisticateds' competence to that of the more sophisticated. (Of course, this exercise in comparative statics must respect the constraint that probabilities cannot be less than 0 or more than 1.)

13. The 2006 report of the Senate Committee on Homeland Security and Governmental Affairs says both that many parties bear blame for the disaster that followed Hurricane Katrina and that Michael Brown's exceptionally poor performance made matters worse (U.S. Senate 2006, esp. chs. 1 and 14). Brown, the director of FEMA during Katrina, refused to coordinate efforts with his immediate superior and the head of the Department of Homeland Security, Michael Chertoff, instead insisting on speaking directly to White House staff (U.S. Senate 2006, 311; see also ibid., 8 and 214). He was unaware that even a Category 3 hurricane would threaten thousands of lives in New Orleans, even though this was well known to FEMA professionals and had been suggested by FEMA exercises held shortly before Katrina (U.S. Senate 2006, 4 and ch. 8). In contrast to previous practice, he did not position FEMA staff in New Orleans before Katrina hit (Adams 2005; Hsu 2005). On the day that the hurricane made landfall, he promised the governor of Louisiana that a convoy of evacuation buses would arrive within hours, but neither he nor anyone else at FEMA ordered buses to be sent for two more days (U.S. Senate 2006, 9). Two days after the hurricane hit, he was dismissive of dire warnings from a FEMA staff member in New Orleans (Brown 2005). One of Brown's predecessors, James Lee Witt, was evidently much more competent; the abrupt change—an extreme example of Claim 3—was striking. See Lewis (2008, 157-68) for an analysis of how this competency-difference at the top affected FEMA's performance.

14. For an excellent exposition of the value of cognitive diversity in group decision-making, see Page 2007.

15. We are not claiming that merit-based personnel systems work flawlessly. Our argument doesn't require such an absurdly strong assertion. It suffices to note that there have been big changes in many systems. Consider, for example, the forty-year transformation of the Japanese navy, which used a host of merit-based reforms—hiring expert British advisers, adopting competitive exams to select officers, using battle exercises rather than personal hunches to determine how to enter battles—to transform itself from a virtual nonentity in the middle of the nineteenth century to a force capable of defeating the powerful Russian navy in 1905 (Bendor n.d.).

16. We were tempted to cite the example of the drop in FEMA's performance that occurred partly as a result of the change at the top, from Witt to Joe Allbaugh to Brown, but this is not really a case in point. As we noted, the consensus opinion of most FEMA observers is that this transition in directors entailed a big change in competence, not a modest one. (One might argue that the system's performance was degraded so badly by the loss of competency at the top that the episode confirms a more extreme version of Claim 4—big changes in the competence of key officials produce big shifts in the system's performance—but we won't undertake a defense of that variant of Claim 4 here.)

17. The 25th Amendment to the U.S. Constitution offers at least partial protection against a clearly senile president from having his finger on the nuclear button.

18. There are other types of redundancies, such as the multiple veto points of a separation-of-powers system (Landau 1973), but these can produce tradeoffs between type-1 errors (adopting bad policies) and type-2 errors (rejecting good ones). A central feature of Condorcetian redundancy is that it can suppress both types of errors simultaneously.

19. More than a few people believe this accurately describes the Bush administration's approach to civilian administration in Iraq (see especially Chandrasekaran 2006). We thank Josh Bendor for suggesting the idea of the incompetence multiplier effect.

20. Not always: in some contexts a principal will rationally violate the ally principle by selecting an agent who is ideologically more different than one who is more similar, ceteris paribus (Bendor, Glazer, and Hammond 2001). But these are rather special circumstances; the claim in the text holds most of the time.

21. As most of these models are spatial, in the Downsian sense, ideological distance is defined as the distance between the ideal point of the principal and that of the agent. (Complexities arise if there are multiple principals; these are sometimes resolved by referring to the ideal point of the median voter.)

22. Another potential criticism is that many voters have mistaken factual beliefs, and that these voters may therefore be unlikely to vote for the candidate who would better serve their desires (i.e., $p < .5$ for these voters). The premise and conclusion of this criticism may be correct, but they do not affect our argument. As we have noted, our argument does not require a minimum level of competence in the electorate; in particular, Claim 2 is likely to hold even if many or all voters fail to choose the candidate who would better serve their desires.

REFERENCES

Adams, Rebecca. 2005. "FEMA Failure a Perfect Storm of Bureaucracy." *CQ Weekly,* 12 Sept. 2378.

Bartels, Larry M. 1996. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40 (February): 194-230.

Bendor, Jonathan. N.d. "Institutions and Individuals." In idem., *Bounded Rationality and Politics.* Berkeley: University of California Press. Forthcoming.

Bendor, Jonathan, and Adam Meirowitz. 2004. "Spatial Models of Delegation." *American Political Science Review* 98 (May): 293-310.

Bendor, Jonathan, Amihai Glazer, and Thomas H. Hammond. 2001. "Theories of Delegation." *Annual Review of Political Science* 4: 235-69.

Bennett, Stephen Earl. 1988. "'Know-Nothings' Revisited: The Meaning of Political Ignorance Today." *Social Science Quarterly* 69 (June): 476-90.

Bennett, Stephen Earl. 1989. "Trends in Americans' Political Information, 1967-1987." *American Politics Quarterly* 17: 422-35.

Black, Duncan. [1958] 1987. *The Theory of Committees and Elections*. Norwell, Mass.: Kluwer.

Brehm, John O. 1993. *The Phantom Respondents: Opinion Surveys and Political Representation*. Ann Arbor: University of Michigan Press.

Brown, Michael D. 2005 August 31. E-mail to Marty Bahamonde. http://i.a.cnn.net/cnn/2005/images/11/03/brown.emails.pdf (accessed June 8, 2008).

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. Chicago: University of Chicago Press.

Cassel, Carol A., and Celia C. Lo. 1997. "Theories of Political Literacy." *Political Behavior* 19 (December): 317-35.

Chandrasekaran, Rajiv. 2006. *Imperial Life in the Emerald City*. New York: Knopf.

Converse, Philip E. 1990. "Popular Representation and the Distribution of Information." In *Information and Democratic Processes,* ed. John A. Ferejohn and James H. Kuklinski. Urbana: University of Illinois Press.

Converse, Philip E. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3: 331-53.

Converse, Philip E. 2006a. "Democratic Theory and Electoral Reality." *Critical Review* 18 (1-3): 297-329.

Converse, Philip E. [1964] 2006b. "The Nature of Belief Systems in Mass Publics." *Critical Review* 18 (1-3): 1-74.

Delli Carpini, Michael X., and Scott Keeter. 1991. "Stability and Change in the U.S. Public's Knowledge of Politics." *Public Opinion Quarterly* 55 (Winter): 583-612.

Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.

Gilens, Martin, Lynn Vavreck, and Martin Cohen. 2007. "The Mass Media and the Public's Assessments of Presidential Candidates, 1952-2000." *Journal of Politics* 69 (November): 1160-75.

Grofman, Bernard, Guillermo Owen, and Scott L. Feld. 1983. "Thirteen Theorems in Search of the Truth." *Theory and Decision* 15 (September): 261-78.

Hsu, Spencer S. 2005. "Chertoff Vows to 'Re-engineer' Preparedness." *Washington Post,* 20 Oct.: A2.

Jackman, Simon, and Paul M. Sniderman. 2002. "Institutional Organization of Choice Spaces: A Political Conception of Political Psychology." In *Political Psychology,* ed. Kristen Renwick Monroe. Mahwah, N.J.: Lawrence Erlbaum Associates.

Keeter, Scott, Carolyn Miller, Andrew Kohut, Robert M. Groves, and Stanley Presser. 2000. "Consequences of Reducing Nonresponse in a National Telephone Survey." *Public Opinion Quarterly* 64 (Summer): 125-48.

Kinder, Donald R. 1998. "Opinion and Action in the Realm of Politics." In *The Handbook of Social Psychology,* 4th ed., ed. Daniel T. Gilbert, Susan T. Fiske, and Gardner Lindzey. New York: McGraw-Hill.

Kinder, Donald R., and David O. Sears. 1985. "Public Opinion and Political Action." In *The Handbook of Social Psychology,* 3rd ed., ed. Gardner Lindzey and Elliot Anderson. New York: Random House.

Kweit, Mary, and Robert Kweit. 2006. "A Tale of Two Disasters." *Publius: The Journal of Federalism* 36 (3): 375-92.

Ladha, Krishna K. 1992. "The Condorcet Jury Theorem, Free Speech, and Correlated Votes." *American Journal of Political Science* 36 (August): 617-34.

Ladha, Krishna K. 1995. "Information Pooling through Majority-Rule Voting: Condorcet's Jury Theorem with Correlated Votes." *Journal of Economic Behavior and Organization* 36 (May): 353-72.

Landau, Martin. 1973. "Federalism, Redundancy and System Reliability." *Publius* 3 (Autumn): 173-96.

Lewis, David E. 2008. *The Politics of Presidential Appointments: Political Control and Bureaucratic Performance*. Princeton: Princeton University Press.

Luskin, Robert C. 1987. "Measuring Political Sophistication." *American Journal of Political Science* 31 (November): 856-99.

Luskin, Robert C. 1990. "Explaining Political Sophistication." *Political Behavior* 12 (December): 331-61.

Luskin, Robert C. 2002. "From Denial to Extenuation (and Finally Beyond): Political Sophistication and Citizen Performance." In *Thinking about Political Psychology,* ed. James H. Kuklinski. New York: Cambridge University Press.

Luskin, Robert C., and John G. Bullock. 2006. "Taking 'Don't Know' for an Answer: DK Responses and Political Knowledge." University of Texas at Austin. Typescript.

Miller, Nicholas R. 1986. "Information, Electorates, and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem." In *Information Pooling and Group Decision Making,* ed. Bernard Grofman and Guillermo Owen. Greenwich, Conn.: JAI Press.

Mondak, Jeffery J. 2001. "Developing Valid Knowledge Scales." *American Journal of Political Science* 45 (January): 224-38.

National Center for Education Statistics. 1993. "120 Years of American Education: A Statistical Portrait." National Center for Education Statistics. http://nces. ed.gov/pubsearch/pubsinfo.asp?pubid=93442 (accessed June 7, 2008).

Neuman, W. Russell. 1986. *The Paradox of Mass Politics: Knowledge and Opinion in the American Electorate*. Cambridge, Mass.: Harvard University Press.

Nie, Norman H., Jane Junn, and Kenneth Stehlik-Barry. 1996. *Education and Democratic Citizenship in America*. Chicago: University of Chicago Press.

Nie, Norman H., Sidney Verba, and John R. Petrocik. 1976. *The Changing American Voter*. Cambridge, Mass.: Harvard University Press.

Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.

Prior, Markus. 2007. *Post-Broadcast Democracy*. New York: Cambridge University Press.

Rawls, John. 1999. *A Theory of Justice*. Revised ed. Cambridge, Mass.: Harvard University Press.

Scanlon, Joseph. 2006. "Two Cities, Two Evacuations: Some Thoughts on Moving People Out." http://understandingkatrina.ssrc.org/Scanlon/ (accessed June 7, 2008).

Schneider, Saundra K. 2005. "Administrative Breakdowns in the Governmental Response to Hurricane Katrina." *Public Administration Review* 65 (September/October): 515-16.

Simon, Herbert A. 1985. "Human Nature in Politics: The Dialogue of Psychology with Political Science." *American Political Science Review* 79 (June): 293-304.

Smith, Eric R.A.N. 1989. *The Unchanging American Voter*. Berkeley: University of California Press.

Somin, Ilya. 1998. "Voter Ignorance and the Democratic Ideal." *Critical Review* 12 (4): 413-58.

U.S. Congress. Senate. Committee on Homeland Security and Government Affairs. 2006. 109th Cong., 2nd sess.

Wand, Jonathan N., Kenneth W. Shotts, Jasjeet S. Sekhon, Walter R. Mebane, Jr., Michael C. Herron, and Henry E. Brady. 2001. "The Butterfly Did It: The Aberrant Vote for Buchanan in Palm Beach County, Florida." *American Political Science Review* 95 (December): 793-810.