

Against All Odds?

The Red-Green Victory

Helmut Norpoth

Political Science, SUNY Stony Brook

Thomas Gschwend

Mannheimer Zentrum für Europäische Sozialforschung (MZES),
University of Mannheim

Picking winners in electoral contests is a popular sport in Germany, as in many places elsewhere. During the 2002 campaign for the Bundestag, pre-election polls tracked the horse race of party support almost daily. Election junkies were invited to enter online sweepstakes. They could also bet real money, albeit in limited quantity, on the parties' fortunes on WAHL\$TREET, a mock stock market run by *Die Zeit* and other media. As usual, election night witnessed the race of the networks to project the winner the second the polls where voters had cast their ballots closed. But in 2002, there was also one newcomer in the business of electoral prophecy: a statistical forecast based on insights from electoral research.

Three months before the date of the 2002 Bundestag election, we issued a vote forecast based on a model that was published two years earlier.¹ Our prediction was that the coalition parties in office would win the Bundestag election on September 22. SPD and Greens, so the forecast first publicized by dpa on June 23 and confirmed on August 24, were going to get 47.1 percent of the vote.² That would be enough, we said, to defeat the combination of CDU/CSU and FDP, provided all other parties mustered at least 6 percent of the vote. The prediction proved to be a forecaster's dream come true. Not only did the red-green coalition win the election, but the actual vote of those two parties matched our forecast right down to the decimal point. And all that occurred against extremely long odds as set by opinion polls.

To offer a forecast based on a few predictors of the vote was a novel experience in Germany. And while some novelties are swept up eagerly, this one was greeted with a mix of bewilderment, disbelief, and outright disdain. The attempt of forecasting an election outcome struck many as simple-minded, if not altogether wrongheaded; as more an exercise in astrology than astronomy. The only thing certain was the belief that the forecast would badly miss its target. One critic gleefully predicted that our "model will not survive the next [2002] election."³ In June and through much of the summer, of course, polls showed the governing parties trailing far behind and headed for certain defeat. So was it just luck, as detractors claimed afterwards, that our forecast model "survived" the 2002 election? No doubt, getting it dead right was a stroke of luck. That is a once-a-lifetime experience for a forecaster. But having a model that gets the outcome right almost every time and comes close to the margin of victory is not simply a case of luck. What is the *Zauberformel*, or magic formula, as the German media teasingly referred to our forecast model? How did we design it? What are its scientific qualifications? And how has the formula fared in previous Bundestag elections?

No magic trick at all, our forecast formula combines predictors of the vote that are familiar to students of elections in Germany as well as elsewhere. We owe a special debt to forecast models of American elections.⁴ Yet our formula should not be mistaken for a German replica of an American model. The predictors of our forecast model for Bundestag elections, in a nutshell, are long-term partisanship (a normal-vote baseline), short-term chancellor approval, and a medium-term dynamic of declining incumbent support over time, call it the "cost of ruling" or simply government fatigue. Using election returns and measures from opinion surveys, going as far back as 1953, we estimated the statistical influence of those variables on the vote in Bundestag elections. Forecasts after the fact confirmed that the model would have been capable of predicting German election outcomes from 1953 to 1998 with an average error of no more than 1.5 percent and picking the right winner each time. Given that track record, getting the winner right in 2002 should not really be all that surprising. More puzzling perhaps is why the 2002 pre-election polls varied so much when the outcome was so predictable.⁵ We will offer some speculations on this puzzle at the end, but the final resolution

must come from the organizations that poll vote intentions and from the media that report the electoral horse race.

The Forecast Target

Whoever tries to forecast the outcome of an election has to settle first on what exactly it is that is being forecast. The vote shares of all parties? That is what the pre-election polls aim to supply with samples of the total electorate, at least for parties above a minimum threshold of support. Here the quality of the forecast is measured by the average absolute deviation of the actual vote shares from the ones reported in the polls. This is a competition that a forecast model like ours cannot enter. Electoral research has a tough enough time with the analysis of vote choices in a two-party setting. We do not know of any model of voting that can determine with sufficient precision how voters make up their minds in a setting with five relevant political parties. This is not the place to fill this lacuna.

Our effort focuses on the politically most telling combination of parties: the parties forming the government before the election. Proportional representation and the resulting multiparty system in Germany have not prevented the creation of stable alignments of parties in and out of government. Politics in Berlin (and previously Bonn) is coalition politics. The electoral process and the policy making process in Germany are joined at the hip by a simple rule: a governing coalition whose parties obtain a majority of seats in an election continues in office, while a coalition that fails in that objective is replaced by another one. Some exceptions notwithstanding, this is a rule that offers valuable guidance for students of German elections. It transforms an unwieldy set of party alternatives into a handy dichotomy of incumbent versus non-incumbent parties. This is not only practical but also theoretically sound, as we argue below.

For each Bundestag election since 1949, we obtained the combined vote of the governing parties. Not included was the vote of any party that belonged to the government at some time during the term preceding the election but that departed before election day. Except for two elections, we were able to ascertain the incumbent vote without any difficulty. The 1969 election is a tough case. CDU/

CSU and SPD entered the election as partners in a grand coalition, but they parted company as soon as the votes were tallied. What is more, the combined share of those two parties (88.8 percent) is an extreme case that requires special consideration or else the exclusion of that case. The most practical solution was to use only the vote share of the CDU/CSU for that election. The CDU/CSU was clearly the dominant party in the coalition, having supplied the chancellor throughout the preceding term (first Erhard, then Kiesinger).

The other tough case is the election of 1983. Technically, the coalition government going into that election consisted of CDU/CSU and FDP. Yet, this was an election that was specially called to let the public decide on the replacement of the previous government of SPD and FDP by means of the constructive vote of no confidence a few months earlier. Therefore, the 1983 election was really about the record of the Schmidt government that had been prematurely terminated in the Bundestag. With the FDP newly aligned with the CDU/CSU, we consider only the vote share of the SPD, the dominant party of the Schmidt government, for that election.

A problem of another sort is posed by the 1990 unification. Should we use the vote in the newly unified Germany for elections since then or stick to the vote in the old Federal Republic? Again, for practical reasons, we decided to keep the continuity of the vote series intact until 1998. So, with an expression of apology to easterners, their votes have been excluded in our analysis of elections from 1990 through 1998. But they won't be excluded any longer. Beginning with the 2002 election, we end this exclusion and start relying on the vote in the unified Germany. That is, of course, the first election for which we offer a true forecast—that is, a forecast before the event actually takes place. In any event, who would have cared about a forecast for just the old Federal republic in 2002?

What then determines voter support for governing parties in any given Bundestag election? What makes the difference between winning and losing? What we need ideally are laws of behavior, but that is something the field of electoral behavior has yet to provide and probably never will. Short of that, we need good theory and reliable data—better yet, theory and data that speak to each other. That certainly sounds more manageable, although not everyone will agree on how to meet those needs. The following is what we

consider to be plausible insights into electoral behavior that are also practical for forecasting.

Theoretical Foundation

Our main point of departure is the premise that the voting decision is determined by the confluence of long-term, short-term, and medium-term components. This is, after all, an exhaustive and mutually exclusive set of categories. The distinction between long-term and short-term factors is well known and widely accepted, although medium-term forces may require some explanation and justification. The influence of long-term forces on electoral decisions is an undisputed fact. Two classics in electoral behavior have documented, each in its own way, the grip of long-term partisan predispositions.⁶ The German electorate is no exception. What determines the vote choices in Bundestag elections to a high extent are long-term stable attachments to political parties, however they may be shaped, whether by psychological or sociological processes.⁷ The aggregate distribution of party identification in the German electorate provides a "normal vote" baseline for the outcome of a given election under normal circumstances.⁸

While the notion of long-term forces, in principle at least, is not hard to grasp, the notion of short-term forces is more elusive. What tips the electoral balance in the short run, as we see it, is a matter of retrospective judgment. The electoral calculus boils down the following: "If the performance of the incumbent party is 'satisfactory' ..., the voter votes to retain the incumbent governing party ...; while if the government's performance is not 'satisfactory', the voter votes against the incumbent ..."⁹ Or, to put it more apocalyptically, the electorate acts as a "rational god of vengeance and reward."¹⁰ Seen from that vantage point, the vote is a retrospective evaluation of incumbent performance.

Our definition of short-term forces in Bundestag elections owes a big debt to that school of thought. The judgment of government performance is the key factor that, in the short-run, decides the outcome of German elections. Since governments in the Federal Republic, of course, are not governed by single parties but by coalitions,

tions of parties, it is imperative to use the plural for incumbent party. Whatever complications that may inject into the calculus of voting cannot be discussed at this point. Suffice it to say that the patterns of coalition politics in Germany have been stable enough to sway most of those concerns.

Aside from long-term partisanship and short-term performance judgment, our vote model also includes a third factor. To get an inkling for it, consider the fact that voting is not a virgin experience for most citizens. They have done it before, and the act has lost its innocence. The parties involved are familiar, and even the leading personnel has been around for a while. Hence the vote is subject to a fatigue factor that makes not only good sense but is also extremely profitable for forecasting. Some have called it the "cost of ruling": governing is costly for electoral support.¹¹ Parties in office tend to lose rather than gain votes from election to election. Keynes' famous saying ("In the long run, we are all dead") must be adapted to government parties for the medium run. During the Weimar Republic, coalition governments suffered from rapid voter fatigue, with most of them not even surviving a single term.¹² The record of coalition governments is certainly better during the Federal Republic, but it would seem unlikely that they could escape the cost-of-ruling effect.

To sum up thus far, our model for elections in the Federal Republic comprises three plausible predictors: (1) long-term partisan strength; (2) short-term evaluation of incumbent performance; and (3) the medium-term decline of government support over time in office. However elegant or misshapen this model of the German vote may be, it still has to pass statistical muster. And before taking that step, we must design measures for those components.

The Predictors


Measuring electoral variables over a span of fifty years or so, like making sausage, is not something for purists. Opinion surveys track few, if any, questions over such long periods. The business of survey research in Germany is barely that old. For measuring long-term partisanship, we have altogether forsaken survey data and instead derived an estimate from election returns. It is undeniable that the

long-term support of a given party registers in its mean vote percentage over several elections. For the still quite young Federal Republic, it seemed prudent to rely only on a few past elections. Hence the long-term partisan support for the governing parties in the German electorate is measured as follows:

Long-term Partisanship = Average vote in the last three Bundestag elections.

For the early Bundestag elections, it is unavoidable to modify this measure of long-term partisanship. The first three simply have not enough previous ones to fall back on. And there is just no way we can afford dropping them all from the analysis. Those early elections also witnessed a massive transformation of the German party system that realigned the partisan loyalties in the German electorate. That makes it defensible to use only the immediately preceding election for a measure of long-term partisanship for the elections of 1953 and 1957. With nothing to guide us for the first election, we have no choice but to let 1949 go.

Figure 1 Partisanship and the vote



The relationship between long-term partisanship and incumbent vote is quite strong, as can be seen in Figure 1. To be sure, some elections do not fit especially well. One would not expect long-term orientations to offer much guidance for voting in 1953, only the second election of a new political system. This is an incumbent victory that one would predict to derive primarily from short-term forces. Similarly, an incumbent defeat, as in 1998, registers in an electoral showing far below the normal-vote prediction. All in all, long-term partisanship does not correlate with the vote in any given election strongly enough ($r = .50$) to use it as the sole predictor.


What matters in the short-run for the vote, as we have argued above, is public satisfaction with the performance of the incumbent government. It is tempting to equate that with economic performance, as is often done in aggregate vote models. Almost no forecast model of American elections does without the economy as a predictor. And economic conditions have certainly left an imprint on voting in Bundestag elections. Yet the economy is not everything or the only thing that determines government performance. What we need is a more encompassing measure that is not tied to one particular domain. How economics and maybe foreign policy shape that general satisfaction with government performance is a matter explored elsewhere.¹³

For the political system of the Federal Republic, we settled on chancellor approval as the best measure of overall government satisfaction. The federal chancellor is the most visible figure of any federal government and his reputation mirrors that of the government he leads. While purely personal considerations may play a part in evaluations of the chancellor, political criteria would seem to have the upper hand. That partisanship colors chancellor support goes without saying. In the end, what counts is the portion of chancellor approval not due to partisanship.

Lacking a single measure of chancellor approval for the last fifty years, we have fashioned a series of comparable items from several survey sources.¹⁴ For the most part, we used the question in the German Election Studies about chancellor preference (incumbent versus challenger). For the early elections (1953 and 1957) and the special case of 1983, we had to turn to the Allensbach question, "Are you satisfied with the policy of chancellor [name]?" Wherever possible,

our measure of chancellor support for a given election averages the values one month and two months before election day.

Figure 2 Chancellor support and the vote



The correlation between chancellor approval and the vote for the governing parties is quite strong, as can be gleaned in Figure 2. Many elections fall close to the linear fit, including the 1953 case this time. The fit nonetheless is by no means perfect ($r = .71$). It would be foolish to make a forecast of German elections solely from chancellor support. What is more, the less-than-perfect fit shows that chancellor support and the vote are not two sides of the same coin. It is not a tautology to say that someone votes for one of the governing parties because she approves of the incumbent chancellor. There is enough room for additional factors, quite aside from long-term partisanship.

In the medium run, our vote model specifies a fatigue effect that prophesies electoral decline to incumbent parties. The longer a government is in office, the more its electoral support diminishes, to a point where defeat in an election brings on a new government. Sixteen years is the maximum thus far that any government has held office in the Federal Republic. We measure voter fatigue simply by

the number of terms a government, or the leading governing party if the coalition composition changed, has been in office.

Figure 3 Term of office and the vote

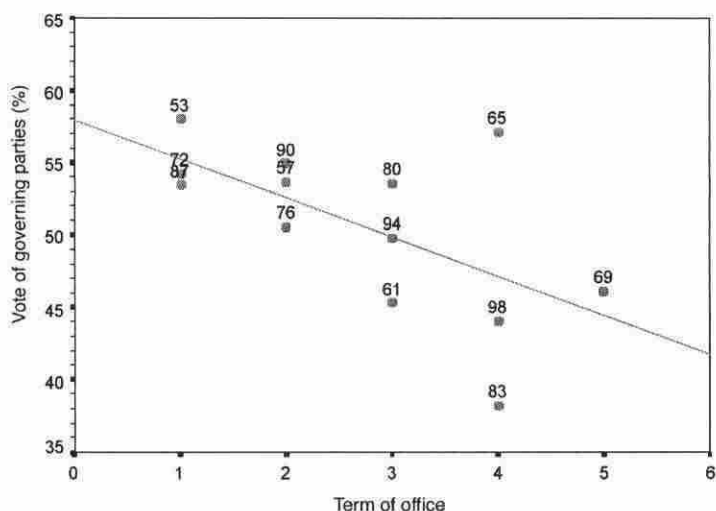


Figure 3 confirms the hypothesis that the vote of governing parties declines as the number of terms of that government increases. The correlation is fairly strong and has the proper negative sign ($r = -.62$). To be sure, some of that electoral decline is prompted by the withdrawal of parties from the government. The departure of several parties from government in the early years (1953-1957) had to diminish the vote of the remaining governing parties at the next election. But since then, voter attrition has been the primary source of the electoral decline of governing parties. That is true for the social-liberal coalition (1969-1982) under Chancellors Brandt and Schmidt as well as the Christian-liberal one under Kohl (1982-1998).

The reader may wonder, however, whether this medium-term fatigue is really something independent from short-term evaluations of incumbent performance or just an artifact of the latter. If such fatigue exists, should it not register in declining satisfaction with the incumbent chancellor, which in turn drives down the vote of govern-

ing parties? Indeed, chancellor support and term of office correlate with each other ($r = -.34$), but not to such a degree that the relationship between them would seem in imminent danger of proving spurious. It remains to be seen how much of the fatigue factor remains with chancellor support held constant, and vice versa. That takes us to the next step where he put all three predictors in one electoral basket.

The Forecast Model

Using all three predictors in a single vote equation lets us determine whether long-term partisan strength, short-term evaluation of incumbent performance, and the medium-term decline of government support each exert an independent effect on the incumbent vote in Bundestag elections. Moreover, it shows how well the three predictors together account for the vote. Is that good enough for forecasting future elections? The combination of the three predictors that provides the best possible fit of the actual vote in past elections (1953 to 1998) is as follows:

	$\text{VOTE} = -6.55 + .76(\text{PAR}) + .39(\text{CHA}) - 1.50(\text{TERM})$
Where	$\text{VOTE} =$ The actual vote of the governing parties (percent)
	$\text{PAR} =$ Average vote of governing parties in last three elections (percent)
	$\text{CHA} =$ Chancellor approval (percent)
	$\text{TERM} =$ Terms in office of governing parties (number)

The coefficients of all three vote predictors are statistically significant beyond any doubt and their signs are all in the expected directions. Thus, each of them brings something distinct to the table. We can rule out the possibility that partisanship has such a tight grip on chancellor approval that the latter has no vote leverage left. It is also not true that the decline of support with increasing number of terms in office manifests itself in chancellor support. However unwelcome to the coalition governments in the Federal Republic, the cost-of-ruling effect proves extremely helpful to us in coming to grips with election outcomes, and hence for forecasting.


What do these results tell us about election outcomes in the Federal Republic? For one, long-term partisan strength does help a governing coalition get re-elected, but it does not guarantee it. The governing parties can expect to retain about three fourth of their "normal" vote in a given election. Contrary to claims about dealignment, long-term partisanship is a powerful electoral factor in Germany. Still, it is not strong enough to secure re-election. Second, chancellor approval has a strong effect, above and beyond long-term partisanship. For every point in chancellor approval, the incumbent parties can expect to gain close to one-half of a percent in votes. Granted, the incumbent parties cannot count on the vote of everyone who favors the chancellor. Let us not forget that partisanship counts for a lot, too, in German elections. But chancellor approval adds a critical margin to the base support of the governing parties. A chancellor whose approval rises from, say, 40 to 60 percent adds eight percentage points to the vote total of the parties in his government. That spells the difference between winning and losing. And finally, government fatigue associated with increasing terms of office proves costly for the governing parties at election time: they have to expect to drop about 1.5 percentage points for every term in office, independent of the incumbent chancellor's approval and partisanship.

Taken together, the three predictors capture the actual vote of incumbent parties in Bundestag elections from 1953 to 1998 with an average error no larger than 1.46 percent. That is close enough in most elections to pick the right winner. Luckily, in close elections (such as 1976), the margin of error is far smaller than the average to allow us to make the right call. By another measure of fit, the explanatory power of the vote equation with the three predictors reaches 94 percent. Very little of the variance of the actual vote in elections between 1953 and 1998 is left unexplained. Such a performance inspires strong confidence that the model can make reasonably accurate forecasts about future elections.

As often is the case, however, betting on the future leaves not only chronically doubting souls uneasy. Isn't all this of the "too good to be true" variety? We ourselves probably would be suspicious if others came up with bold claims about being able to make vote predictions. Only three explanatory variables! And only thirteen cases! Buyer beware! Are there not any number of factors that would cor-

relate as highly with the vote in Bundestag elections, and many without any rhyme or reason, such as the outcome of the baseball world series. Many Germans have heard of the notorious “Mierscheid law.” It predicts the SPD vote based on the size of crude steel production in Germany.¹⁵ Or take more titillating predictors such as Andreas Heij’s “breast barometer.”¹⁶ This Munich psychology professor claims that the bigger the breasts of Playboy models, the better the chance of a CDU/CSU victory. Based on that “law,” he predicted (as a joke, one assumes) that the Schröder government was doomed. The longer the series of Bundestag elections, the harder it will be to come up with such “laws.” If there is no good causal story between alleged predictors and the vote, their luck will run out sooner or later. None of those alternatives correlated as accurately over so many trials with the Bundestag vote as did our model. Our predictors certainly have plausibility on their side.

Figure 4 Forecasts and results



What is more, tests have shown that our model is able to “predict” with great accuracy each of the elections from 1953 to 1998 from an analysis of all the others.¹⁷ Of course, all those elections were in the book before our model arrived. So what we did was to exclude each of those elections, one at a time, find the vote equation, and use it to “forecast” the excluded election. So for the election of 1998, we would have based our analysis on the elections from 1953 to 1994, pretended that we went back in a time machine to the summer of 1998, and made our forecast for that election before the event. And ditto for all the other previous Bundestag elections. Figure 4 displays the forecasts based on that strategy, along with the actual vote for all those elections.

Apart from the notoriously difficult case of 1953—the first instance of a federal government seeking re-election—there is only one other election (1965) where our model prediction is off by more than two percentage points. Overall, these simulated forecasts of the combined vote of governing parties in past Bundestag elections (1953–1998) miss the actual outcome on average by only 1.5 percentage points. And we would have correctly picked the winner for every single Bundestag election since 1953. Our best result is the 1976 election: the prediction for SPD/FDP of 50.6 percent is within one-tenth of a point of the outcome (50.5 percent). To come that close requires a bit of luck, but to be close consistently is not just luck. That is a sign of a good model and also reassurance for forecasting.

The 2002 Vote Forecast

Besides a “magic formula,” one needs information about the predictors included in the formula to forecast a future event like the vote in an upcoming election. It does us no good if that information is not available until after the event. We were fortunate that information on the predictors of our vote model were readily available before election day. In the case of German elections, however, there is a catch. While our forecast target was the 2002 vote for the whole Federal Republic, old and new *Länder*, the forecast formula only applied to the old Federal Republic. We could, of course, have made a forecast just for the old Federal Republic. But who would have cared for that in 2002? By the same token, a separate model for the five new

Länder is out of reach with only three elections. So taking the formula derived from voting in the old Federal Republic until 1998 and feeding it with data from the unified Germany in 2002 to get a forecast for the whole electorate is a risky step. How risky? Will that substantially throw off the forecast? We think not. Granted, voters in east and west do not make up their minds in the same way, but these differences are not likely to prove too serious in the aggregate, particularly since only one out of five voters is from the new *Länder*. It is a price we decided to pay for being able to make an electoral forecast for the entire Federal Republic in 2002. So what we need is information from the whole German electorate about each of three predictors in 2002, stick that into the formula, and turn the crank. That is all the magic there is in this *Zauberformel*.

As for long-term partisanship, recall that it is measured by the average vote in the last three Bundestag elections. That comes to a combined 43.3 percent for SPD and Greens in 2002 (the average of 38.5 in 1990, 43.7 in 1994, and 47.6 in 1998). Such a partisan base did not augur well for the re-election prospects of the red-green coalition in 2002. Turn back to Figure 1 and note that 43.3 percent is the lowest of any incumbent government in the Federal Republic. If long-term partisanship was all that mattered for the vote, the red-green coalition was doomed in 2002. This disadvantage perhaps makes clear why those parties fared so poorly in the horse-race polls throughout the election year until the campaign finally heated up.

Our second predictor, chancellor approval, required information about the standing of Chancellor Gerhard Schröder vis-à-vis his challenger Edmund Stoiber two months and one month before the election. As in other recent elections, we relied on the surveys of the Forschungsgruppe Wahlen. The average of their polling in July and August, while excluding respondents without a chancellor preference, yielded a measure of chancellor approval for Chancellor Schröder of 57 percent. It is reassuring that this is also what the other major polling organizations in Germany (Infratest-Dimap, EMNID, Forsa, and Allensbach) showed, on average, in July and August. Such a high rating for the incumbent chancellor augured well, as can be seen by turning back to Figure 2, for the re-election prospect of the governing parties. A mark of almost 50 percent of the vote is in sight if chancellor approval were all that mattered for the vote.

The final predictor is the simplest to ascertain. The red-green coalition was completing its first term. At that stage, as Figure 3 makes clear, incumbent governments are hitting electoral home runs. Of course, given the dismal partisan base of the red-green government in 2002, no landslide victory should have been expected. Still, the prospects were far rosier than they would have been after two or three terms in office.

Feeding our forecasting formula with these data for the three predictors generated the following forecast, which was posted on August 26, nearly one month before the day of the 2002 election:

$$\begin{aligned}\text{Red-Green Forecast 2002} &= -6.55 + .76(43.3) + .39(57) - 1.50(1) \\ &= 47.1\%\end{aligned}$$

This forecast confirmed the preliminary forecast that we issued on June 23, three months before election day. It also predicted 47.1 percent for the red-green coalition (because Schröder enjoyed the same approval then, too). Hence the 2002 election was as predictable one month before the election as it was three months before. A vote of 47.1 percent for the governing parties would be enough to ensure victory over the opposition of CDU/CSU and FDP, so long as all other parties (including the PDS) mustered at least 6 percent. That appeared an easy target to meet, given that over the last three elections the average vote of all the other parties exceeded 8 percent. But would 47.1 percent of the vote be enough for the red-green coalition to capture a majority in the Bundestag? To be sure about that would have required a forecast of the PDS vote, actually two forecasts. One whether the PDS would surpass the 5 percent hurdle and the other whether it would capture at least three districts. No magic formula here, and no guesses either. What we did forecast was that the red-green coalition would defeat the black-yellow opposition. That is what happened on election day 2002, and with exactly the same vote as we predicted.

Predictable Election, Variable Polls

At the time when we went public for the first time with a forecast of a red-green victory, that prospect looked utterly farfetched to just about

all observers. The German media were painting exactly the opposite picture. It was one of doom and gloom for the Schröder government. Quoting opinion polls, the headlines practically called the election for the opposition CDU/CSU and FDP. The governing parties were trailing so far behind in the horse race that defeat was inevitable. And that fate was prophesied for our forecast model as well.

Now imagine a German Rip van Winkle going to sleep in June 2002, just after watching the *heute journal* with the latest horse race numbers, and waking up no sooner than the day after the election. He would have been utterly dumbfounded. In just three months the political world in Germany had turned upside down. The victory for CDU/CSU and FDP that was prophesied by opinion polls had slipped away, while the red-green coalition, given up for dead, had triumphed against all odds. But then imagine Rip's cousin Dot going to sleep with our forecast in mind but oblivious to the polls, and waking up the day after the election, too. She would have found the political world in perfect order. What could she have said to her cousin Rip to clear up his bewilderment?

In other words, the question is, How come the polls swung so wildly from one side to the other when the election was so predictable? In the end, only a thorough scrutiny of the polling procedures can resolve that puzzle. We are not privy to those operations, which are guarded like trade secrets. All we can offer is some reasonable speculation based on a general understanding of opinion polling and some clues provided by published reports. To begin with, German pollsters will admit that a vote intention in an opinion poll is not the same as a vote cast at the polls. The problem is that they do not act accordingly when publicizing their horse race numbers in the media. Those numbers pretend that everyone has reached a voting decision when a large proportion of voters remains undecided. At best, the pre-election polls may indicate how voters with a firm decision plan to vote. But that is of little value so long as close to one-third of the electorate professes to be undecided or unsure about its vote. With so many voters undecided until very late, as was hinted in the coverage of polls in 2002, polling projections of victory for one side or defeat for the other were misleading, if not irresponsible. In the total electorate the combination of CDU/CSU and FDP, we suspect, never enjoyed anything close to a true major-

ity of support months before the election. So these parties did not lose something in the end they really commanded before. They only believed they did, thanks to polling mechanics, especially the treatment of undecided, and the poll coverage in the media.

If there is a pressing story here, it is how the undecided in 2002 came to rally in disproportionate numbers to the red-green coalition rather than the black-yellow opposition. That takes us to the effect of our key short-term predictor: chancellor approval. Consider a voter who is undecided months before the election but who wants Schröder as chancellor, not Stoiber. Most of the undecided must have fit that profile. As the date of the election approaches, such a voter must resolve the indecision: either vote for one of the parties that delivers Schröder as chancellor or stay home. In the end, while some of those voters may have stayed away from the polls, the rest mostly voted SPD or Green, we suspect. To facilitate and maybe force this type of decision is what electoral campaigns are all about. They activate and reinforce latent predispositions to bring about consistency, as the classic campaign study, *The People's Choice*, told us long ago.¹⁸

But what about last-minute events like the flood and the Iraq issue? Didn't they sway voters toward the red-green side? No doubt those German equivalents to October surprises proved helpful to the incumbent parties. Were they a necessary condition without which that swing would not have happened? We do not believe so. Chancellor Schröder enjoyed a huge advantage in the chancellor's race long before the flood and Germany's role in a possible war with Iraq made headlines. One way or another, the hot phase of the election campaign would have pressured undecided voters to bend their voting decision to their chancellor preference. The only way for the red-green coalition to have lost the 2002 election was for the chancellor's approval to collapse. That did not happen before the election. Instead Schröder enjoyed a steady lead over Stoiber throughout 2002, typically in double digits. Had the polls and media concentrated on the chancellor race, there would have been little movement or few surprises to report. The headline would have read ad nauseam: Schröder ahead, eyeing re-election. That is how it turned out on election day. Not against the odds at all.

Appendix

Estimates of the Forecast Model of Bundestag Elections

Independent Variables	Dependent Variable: Vote of Governing Parties (%)
	Parameter (SE)
Long-term partisanship	.76*** (.10)
Chancellor support	.39*** (.05)
Term	-1.50*** (.35)
Constant	-6.55 (6.61)
\bar{R}^2	.936
Standard error of residuals	1.46
(N)	(13)
Durbin-Watson d	1.76

Note: Model estimation based on elections 1953-1998.

* $p < .05$

** $p < .01$

*** $p < .001$

Notes

1. See Thomas Gschwend and Helmut Norpoth, "Soll und Haben: Die deutsche Wählerschaft rechnet mit den Parteien ab," in *50 Jahre Empirische Wahlforschung in Deutschland. Entwicklung, Befunde, Perspektiven, Daten*, Markus Klein, Wolfgang Jagodzinski, Ekkehard Mochmann and Dieter Ohr, ed. (Wiesbaden, 2000), 389-409; also Thomas Gschwend and Helmut Norpoth, "Wenn am nächsten Sonntag ...": Ein Prognosemodell für Bundestagswahlen," in *Wahlen und Wähler: Analysen aus Anlass der Bundestagswahl 1998*, Max Kaase und Hans-Dieter Klingemann, ed. (Wiesbaden, 2001), 473-99.
2. See Jochen Neumeyer, "Zauberformel für Wahl: Forscher entwickeln Prognosemodell," dpa, June 23, 2002.
3. As quoted by Neumeyer.
4. See Michael S. Lewis-Beck and Tom W. Rice, *Forecasting Elections* (Washington, DC, 1992); James Campbell and James C. Garand, *Before The Vote. Forecasting American National Elections* (Thousand Oaks, CA, 2000); and Helmut Norpoth, "Is Clinton Doomed? An Early Forecast for 1996," *PS* 28 (1995), 201-207.

5. For an analysis of this paradox in US elections, see Andrew Gelman and Gary King, "Why Are American Presidential Election Campaign Polls So Variable When Votes Are So Predictable?" *British Journal of Political Science* 23 (1993): 409-451.
6. See Angus Campbell, Philip E. Converse, Warren E. Miller and Donald E. Stokes, *The American Voter* (New York, 1960); and Paul Lazarsfeld, Bernhard Berelson and Hazel Gaudet, *The People's Choice* (New York, 1944).
7. See e.g., Franz Urban Pappi, "Klassenstruktur und Wahlverhalten im sozialen Wandel" in Max Kaase and Hans-Dieter Klingemann, ed., *Wahlen und Wähler. Analysen aus Anlass der Bundestagswahl 1987* (Opladen, 1990); and Helmut Norpoth, "Party Identification in West Germany: Tracing an Elusive Concept," *Comparative Political Studies* 11 (1978): 36-61.
8. See Jürgen W. Falter and Hans Rattinger, "Parteien, Kandidaten und politische Streitfragen bei der Bundestagswahl 1980: Möglichkeiten und Grenzen der Normal-Vote-Analyse," in Max Kaase and Hans-Dieter Klingemann, ed., *Wahlen und politisches System. Analysen aus Anlass der Bundestagswahl 1980* (Opladen, 1983), 320-421.
9. See Gerald Kramer, "Short-Term Fluctuations in U.S. Voting Behavior, 1896-1964," *American Political Science Review* 65 (1971): 131-43.
10. See V.O. Key, *Politics, Parties and Pressure Groups*, 5th ed. (New York, 1964), 568; also Morris Fiorina, *Retrospective Voting in American National Elections* (New Haven, 1981).
11. See Martin Paldam, Martin, "How Robust is the Vote Function? A Study of Seventeen Nations over Four Decades," in Helmut Norpoth, Michael S. Lewis-Beck and Jean Dominique Lafay ed., *Economics and Politics. The Calculus of Support* (Ann Arbor, 1991), 9-31. While the Federal Republic came up as an exception to that rule in that analysis, our replication, including elections since 1983, shows the German case to be a perfect fit, with an estimate for the key parameter of -1.7, which is close to the average of -1.6 in Paldam's analysis.
12. See Heino Kaack, *Geschichte und Struktur des deutschen Parteiensystems* (Opladen, 1971), 124-28.
13. See Gschwend and Norpoth 2000 and 2001.
14. For the design of the chancellor-support measure, see Helmut Norpoth, 1977. "Kanzlerkandidaten. Wie sie vom Wähler bewertet werden und seine Wahlentscheidung beeinflussen," *Politische Vierteljahresschrift* 18 (1977): 551-72, especially note 34. Given the wide variation in the percentage of undecided/no opinion from survey to survey, we excluded those categories in calculating chancellor support. That makes all incumbent chancellors appear more popular than the raw percentages indicate. Since this bias benefits all of them it does not distort the estimation of the chancellor effect on the vote.
15. This law is attributed to a Jakob Maria Mierscheid, a phantom SPD Member of the Bundestag of considerable fame. His political life, in all its fictitious glory, was recently documented by Dietrich Sperling and Friedhelm Wollner, *Jakob Mierscheid: Aus dem Leben eines Abgeordneten* (Baden-Baden, 1998). The Mierscheid Law was quite successful in predicting the SPD vote since 1961, with the exception of the 1990 election.
16. See Berliner Kurier, "Männer kreuzeln nach BH-Grösse." July 17, 2002. (www.berlinonline.de/wissen/berliner-kurier/)
17. See Gschwend and Norpoth, 2001, 489-91.
18. See Lazarsfeld et al., 1944, chs. VIII-IX.

