Contents lists available at SciVerse ScienceDirect

Electoral Studies

journal homepage: www.elsevier.com/locate/electstud

Performance voting and knowledge of cabinet composition☆

David Fortunato [a],*, Randolph T. Stevenson [b]

[a] University of California, Merced, USA
[b] Rice University, USA

A B S T R A C T

Nearly every empirical study of outcome oriented retrospective voting assumes, though almost always implicitly, that a) every voter knows the composition of the incumbent government, or b) that voters who may not know who is in government do not cast retrospective votes. In this short essay we provide evidence that these assumptions are quite unlikely to hold and discuss how not taking this possibility into account has influenced our understanding of the retrospective voting mechanism and the conditioning effect of political sophistication on the economic vote. In so doing, we advocate for the inclusion of questions regarding cabinet composition in electoral surveys.

© 2013 Elsevier Ltd. All rights reserved.

1. Motivation

While different theoretical explanations for retrospective performance voting can differ dramatically on such issues as the voter's rationality, motivation, emotionality, and sophistication, any coherent explanation of this kind requires that voters know both who is in the incumbent government and how that government performed while in office. If voters do not know who is in government or how they performed, these models suggest either that no performance voting will occur or that performance voting will result in perverse and normatively unappealing outcomes (i.e., rewarding failing incumbents and punishing successful ones). It is no wonder then that political scientists have spent considerable time exploring the question of how voters form their perceptions of incumbent performance

(Wilkin et al., 1997; Palmer and Duch, 2001; Evans and Andersen, 2006), how such perceptions vary within populations (Duch et al., 2000), and whether they are (at least on average) accurate (Lewis-Beck et al., 2004; Duch and Stevenson, 2010). Further, while controversies still exist, all this effort has resulted in a general consensus affirming of the usefulness of models of retrospective performance voting, but also a widespread acknowledgment that mistaken perceptions of performance can have important effects on election outcomes, as well as on how we estimate and interpret empirical models of performance voting (Wlezien et al., 1997; Lewis-Beck, 2006; Kayser and Peress, 2012).

This scholarly effort to understand the prevalence and implications of mistaken perceptions of performance has not, however, been mirrored in the other (arguably prior) requirement that voters know the composition of the incumbent government. Indeed, there is not a single published study (that we could find) that explores either the extent to which voters are systematically mistaken about the composition of the incumbent government or the implications of such mistakes for our understanding of performance voting. Instead, most scholars have been content to assume either that everyone knows who is in government, or that the existence of some voters who do not have such knowledge is inconsequential for estimating the size

* Corresponding author.
E-mail addresses: dfortunato@ucmerced.edu (D. Fortunato), stevenso@rice.edu (R.T. Stevenson).

Proportion of Voters Who
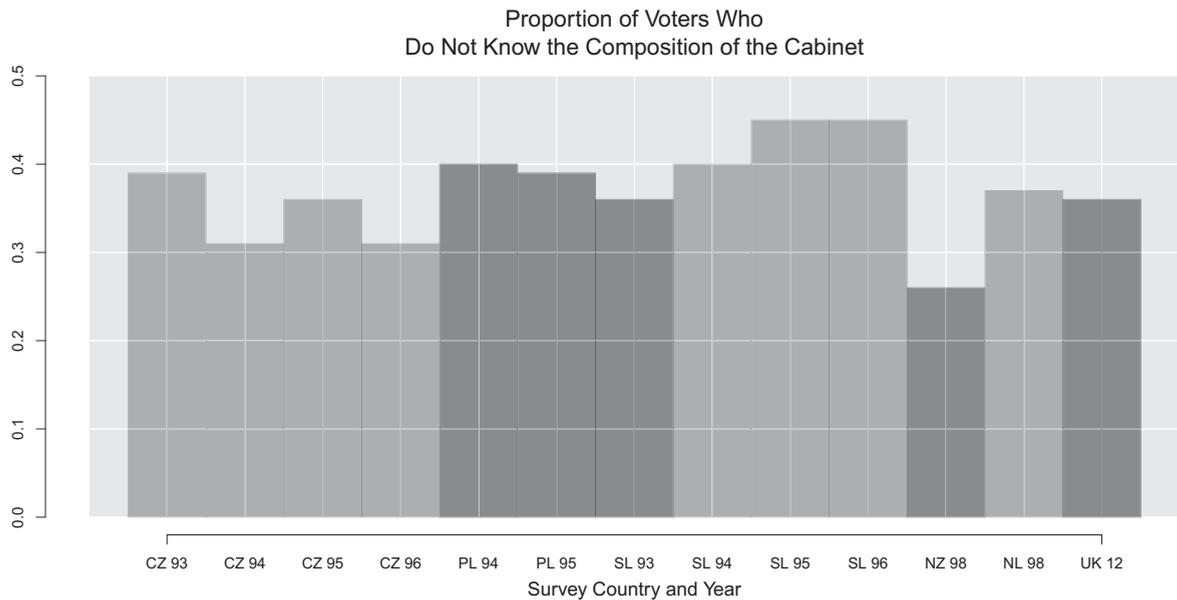Do Not Know the Composition of the Cabinet

Fig. 1. Percent of respondents who are Unable to identify the composition of the cabinet.

of (or understanding the nature of) retrospective performance voting.

In this short paper we show that neither of these assumptions is correct. Specifically, we first use survey evidence to show that there are a reasonably large number of voters who do not know who is in the government, or worse, believe parties that are actually in the opposition are in the government. Second, we show that this fact matters for both empirical estimates of the size of economic voting (systemically depressing them) as well as for empirical assessments of the role that political sophistication plays in economic voting. Consequently, we argue that all future election surveys (and studies of performance voting based on them) should incorporate explicit questions asking voters which parties they think are in the government.

2. How often are voters mistaken about which parties are in government?

In this section we present survey evidence that can be used to directly examine the extent to which voters know who is in cabinet. We have collected all available survey data on this topic and have also commissioned a new survey of British voters. The various surveys rely on different questions and ask respondents for different degrees of detail, so we begin with a broad overview and then use the surveys with the most information (the Dutch, 1998; our British, 2012 surveys) to drill down deeper.

Fig. 1 provides an overview of thirteen surveys that asked voters to identify which parties were in the cabinet at various times in six different countries. The specific questions asked in the different surveys are provided in the Supplementary Materials, but the questions are

broadly similar and so support the general conclusion that, across these very different contexts, about a third of respondents cannot correctly identify the parties in the cabinet.[1]

While most of the available surveys only ask respondents whether a party is in the cabinet or not (or to name the cabinet parties), our survey of British voters (conducted in June of 2012) allows us to go beyond these numbers for the cabinet as a whole and probe voter knowledge of the specific roles different parties play in the cabinet. Specifically, we find that 16% percent of our respondents could not identify the Conservatives as the party of the Prime Minister. Likewise, 12% of respondents thought that Labour, the main opposition party, was in the cabinet. While these numbers are relatively modest, we suspect they are close to the lower bound for ignorance of cabinet composition in coalitional systems. After all, this was the first instance of coalition government in the UK's modern era, and, as such, the resulting media coverage was voluminous. Indeed, when we look at another case – the Dutch 1998 pre-election survey – we find a situation much more likely to produce confusion about cabinet composition. Specifically, the Dutch Prime Minister at the time of the survey (who had taken office in 1994) was from the PVDA, but the main opposition party, the CDA, was out of government, and the prime ministry, for the first time since its

---

[1] It is interesting to note that the percentage of voters correctly identifying the parties (as well as the percent who could not identify any) is relatively constant across time in the same country. This may suggest that knowledge of cabinet partnership is not greatly impacted by temporal changes in the composition of cabinets but instead reflects more enduring features of the electorate (like levels of civic education), characteristics of media systems, or perhaps more enduring features of the political system.

formation (from several different Christian parties) in 1977. In this situation we might expect at least some voters, who were not actively monitoring the political situation, to think the CDA was still in cabinet – and, indeed, a quarter of them did.[2]

Overall, the existing survey evidence, though by no means comprehensive – or even plentiful – suggests that about a third of the electorate is unaware of the composition of the cabinet and that a substantial number of those individuals can make fairly grave mistakes – thinking that large and prominent opposition parties are, in fact, in the cabinet when they are not.

Finally, we ask, "who are the voters who are mistaken about the composition of the cabinet?" This is important since it may be that these are the kind of respondents that we can discount as simply being unlikely to vote in general.[3] If this is the case, we might then conclude that the existence of such individuals can have little impact on our estimates or understanding of economic voting. Thankfully, the Dutch study provides the information we need to examine this possibility in more detail.

We begin by simply examining whether the respondents who did not know the composition of the cabinet intended to vote (in the pre-election module) and then whether they reported actually having voted (in the post-election module). Of the 606 such respondents who answered the vote intention question, 87% said they intended to vote. While this is lower than the corresponding statistic (95%) for respondents who knew the composition of the cabinet, it is still quite high and certainly does not support the conclusion that voters who are mistaken about the composition of the cabinet are mostly non-voters.[4] A parallel analysis of reported vote shows that 83% of the respondents who did not know the composition of the cabinet before the election reported voting in the election.

Beyond their voting behavior, we also examined a multivariate logistic model in which we tried to predict which respondents (from the whole sample) would not know the cabinet composition based on a set of demographics (gender, education, income, and a non-linear function of age) and measures of the respondent's level of education, political interest, media use, and knowledge of general political facts (other than who is in the cabinet). The detailed results are reported in the Supplementary Materials; but, the main findings are that demographics do not matter at all; except for the usual finding of a

gender gap in knowledge (men are less likely to mistakenly identify the cabinet composition than women).[5] Much more consequential are the measures of education, political interest and media use, and knowledge of general political facts. All of these have large, statistically significant, effects and together (without the demographics) correctly classify about 75% of the cases in the whole sample.

That said, in the sub-sample of 679 respondents who did not know exactly who was in the cabinet, the model only classified 52 percent of these correctly. Clearly, then there are a reasonably large number of respondents who both do not know who is in the cabinet but otherwise score high on our measures of measures of education, political interest and media use, and knowledge of general political facts (and so are predicted, incorrectly, to also know the composition of the cabinet). Indeed, while only 11% (of these 679 respondents) have any university level education, 22% have some formal education beyond secondary education. Likewise 29% have above average interest in politics and 20% are in the top half of the distribution of general political knowledge.[6] Thus, it appears that while most of the respondents who do not know the composition of the cabinet are less educated, less interested in politics, and less politically knowledgeable in general; there are a substantial minority of such respondents that score reasonably high on these values.

3. Consequences of mistaken beliefs about who is in the government

Given the evidence of the extent and nature of mistaken beliefs about who is in cabinet, the question remains whether this matters to either our estimates of the extent of performance voting or our understanding of the nature of performance voting more generally. In the next subsection we take up the first of these questions and show that voters who are mistaken about the composition of the cabinet do appear to cast economic votes, but do so for (and against) the wrong parties.[7] The empirical consequence of this is that such cases can have a disproportionate impact in suppressing estimates of the economic vote, since they seem to be showing exactly the opposite of the expected behavior (i.e., voting for the

---

[2] The Dutch survey does not allow us to know for certain whether these respondents thought the CDA was the PM in the government or simply a partner, though given their long history in that role we are likely safe to assume so.

[3] We might also ask if these voters are generally uniformed and suppose that if they are they would be less likely to cast economic votes. This line of reasoning, however, presupposes that economic voting is more likely among sophisticated, informed individuals. While some studies (e. g., Duch, 2001) make this case, we question this result below. More generally, we do not want to assume that all performance voting is rational voting and so enhanced by sophistication. It is certainly possible, for example, that more emotional or reactive performance voting is more prevalent among less sophisticated voters.

[4] There has not been compulsory voting in Holland since 1970.

[5] Given the coding of cabinet knowledge used here (where to be counted wrong a respondent must actually misidentify the role of party rather than simply say "don't know") this finding holds despite the sometimes reported tendency for males to be less willing to answer "don't know."

[6] We measure interest in politics using an index comprised of the standard political interest question as well as questions about newspaper readership and media use. We measure general political knowledge using questions about party sizes and recognition of photos of political figures.

[7] We assume here that one is interested in accessing the overall impact of economic perceptions on voting behavior and the extent to which this impact is consistent with some underlying theory of individual behavior. Thus, if one's theory implies that support for perceived incumbents should decline when perceived economic performance declines, then the relevant test of this hypothesis (and the relevant level of economic voting) should focus on whether economic perceptions indeed had the proposed effects, conditional on who the voter thought the incumbents were – regardless of whom actual incumbents were.

Table 1

Estimates are from a logistic model (DV = 1 if voted for PM and 0 otherwise). Respondents were asked, "According to you, which party or parties are currently part of the cabinet?" 85 respondents answered "Don't know." these individuals are excluded from all the analysis.

Respondents included in the analysis:	Number of observations	Change in support for PM when economic perceptions worsen
All Respondents other than those answering DK to the party composition question (baseline model)	1387	−.040 (−.077, −.004)
Respondents who correctly identify the cab/opp status of all of the FIVE main parties	1022	−.060 (−.099, −.017)
Respondents who correctly identify the cab/opp status of all of the THREE main parties	1057	−.060 (−.100, −.018)
Respondents who correctly identify the CDA as being in the opposition	1095	−.057 (−.096, −.015)

opposition – who one believes is in government – when the economy is good).[8]

3.1. Underestimating the economic vote

In order to access the impact of voters who do not know the composition of the cabinet on estimates of the size of the economic vote, we rely on data from the 1998 Dutch election study introduced above. In addition to questions about perceptions of cabinet membership, the survey also included questions about the respondent's perceptions of economic performance as well as overall satisfaction with government and the government's economic policy. Following Duch and Stevenson (2008), we focus here on the economic vote for the Prime Ministerial party and so use a binomial logistic model to estimate how economic perceptions, ideology, and various standard demographic variables impact vote choice.[9] We then modify this baseline model to explore the impact on these estimates of excluding voters who are mistaken about the composition of the cabinet (measured in different ways). We expect, of course, that the impact of the economy will become greater when such voters are excluded from the sample; but, it is possible that this will not happen – specifically, it may happen that the kind of voters who do not know the composition of the cabinet are also the kind of voters who do not cast performance based votes.

Table 1 provides the substantive conclusions from the model (detailed model summaries are provided in the Supplementary Materials). Specifically, the table shows the estimated change in the probability of voting for the prime minister's party (i.e., changes in "support") that is due to worsening economic perceptions. Rather than report results for a variety of "typical individuals," we report results that are "averaged over the sample" (see Duch and Stevenson, 2008). These are calculated as follows: The model is estimated and 1000 simulated values of the parameters are drawn from the appropriate multivariate normal distribution (e.g., King et al., 2000). These simulated coefficients are then used to create 1000 simulated changes in probabilities for the case when all the covariates in the model are set at the value of the first observation in the data set – except for economic perceptions, which are, of course, the variable that is changing to create the estimated change in probabilities. Specifically, it is changed from the case in which the individual perceives the economy to have gotten better to one in which she thinks it has stayed the same. This creates 1000 changes in probabilities for the first case in the data. These numbers are stored and the process is repeated using the values of the covariates for the next observation in the data. When this is done for all N cases in the sample we have an N x 1000 matrix of changes in the probability of voting for the PM. These are then averaged over individuals to create 1000 effects that are each "averaged over the sample." Of course, these 1000 estimates still vary because of the variation in the 1000 simulated coefficients vectors used to create them. The values reported are the means of these 1000 simulated "averages over the sample" and the confidence intervals are the 5th and 95th percentiles of the simulated values.

The first row gives the estimates for the full sample, so is the typical sort of estimate reported in the economic voting literature (e.g., Duch and Stevenson, 2008). We find that, averaged over the sample, the support of the PM declined about 4 percentage points when voters' perceptions of the performance of the economy over the last year moved from "got better" to "stayed the same." The magnitude of this estimate is right in line with the range of variation in economic voting that Duch and Stevenson (2008) report for 165 different election studies (average of 5 percent and a standard deviation of 4 percent). Further, it is typical of economic voting in the Netherlands (which is usually smaller in magnitude than in other

---

[8] We have also conducted a Monte Carlo study in which we generate an underlying vector of support for parties that depends on economic perceptions and other variables and then vary the number of voters who are mistaken about the identity of the incumbent cabinet but that cast economic votes for or against their perceived incumbents. We then examine how much the estimates of economic voting from standard empirical models of vote choice change for different numbers of such voters. The results show that for every increase in 10% in the share of mistaken voters, the size of the economic voting estimates (based on an undifferentiated sample of voters) is decreased by about 20%. Thus, if a third of the voters are mistaken about the government composition (as our evidence suggests) and some reasonable percentage of these are economic voters, then we would expect our estimates of economic voting to be cut in half.

[9] We include standard demographics, various measures of religious activity and identification, which are strong predictors of support for the Christian parties, and measures of left–right self-placement. We explored various specifications that included partisan attachments, party policy and distance measures and ultimately chose the reported specification because it was robust to all these changes and seems to parsimoniously capture the main drivers of electoral choice.

countries, with statistically insignificant estimates common). For example, of the 12 Dutch elections studied in Duch and Stevenson (2008), none produce confidence intervals for the economic vote of the PM that did not cover zero, though the majority of the point estimates were at least negative. On average, the economic vote of the PM across studies was −.02 percent (though, Duch and Stevenson's estimates for the 1998 case were the largest of all their cases at −.04). This comparatively limited amount of economic voting for the PM in the Dutch context is not a limitation for our analysis, since our concern is not with absolute levels of economic voting, but how this baseline (of little economic voting) changes when we account for perceptions of cabinet composition. Thus, the main result in Table 1 is the comparison of the baseline estimate of economic voting (−.04) with the estimates in the other rows of the table.

The estimates in the last three rows of Table 1 are each based on a sample that excludes a set of voters that lacked some knowledge of the composition of the cabinet. Specifically, we considered three cases. First, we exclude cases in which respondents mistake the cabinet/opposition status of any of the five leading parties.[10] Next, we use the same criteria but only consider mistakes about the three leading parties.[11] Finally, the last row reports estimates in which the only respondents excluded from the sample are those who mistakenly thought the CDA was in the cabinet. About quarter of the respondents did this and it is of interest since this is a particularly relevant sort of error to make if one is an economic voter. After all, Duch and Stevenson (2008) show that the economic vote tends to move between the large, leading parties in the cabinet and opposition, which in this case surely includes the CDA.

The most important comparison is simply between the first row (the traditional estimate) and all the other rows taken together. As we can see, this is the main distinction in the estimates. That is, the subtleties of exactly how mistakes are counted produce only small differences, while the difference between excluding mistaken voters (however measured) or not produces a large difference. The effect of including mistaken voters is to reduce the estimate of economic voting by about 33%. This magnitude is roughly consistent with our Monte Carlo experiment[8] which showed (under ideal conditions) that if 20% of voters were mistaken about the identity of the PM but were nevertheless economic voters, we would see a reduction in the economic voting effect of about 40%. Since we know about 25% of the respondents made the mistake of thinking the CDA was in cabinet (the kind of egregious mistake that is similar to the one used in the simulation) and we would not expect all of these individuals to be economic voters, the finding of a 33% reduction in the estimate makes sense here.

---

[10] "Don't Know" responses to the question of who is in the cabinet are not included in either the baseline model or these models. The reason is clear from our earlier discussion: we are interested in the role of mistaken perceptions, not simply lack of knowledge.
[11] The three leading parties are the PVDA, the CDA, and the VVD. The set of five adds Groen Links and D66.

3.2. Exploring the impact of political sophistication in economic voting

Besides its impact on the empirical assessment of the magnitude of performance voting, accounting for knowledge of the cabinet composition can also be important in understanding how performance voting actually works. As an example, in this section we revisit the question of how political sophistication impacts the ability or willingness of voters to cast performance votes – and show that accounting for knowledge of cabinet composition is absolutely critical for understanding the relationship.

In an influential article, Raymond Duch argued that "[a]s voters become more knowledgeable about the political process, ambiguities regarding the link between government policy and economic outcomes decline, and their level of economic voting rises" (2001: 897). The thrust of his argument, like that of other scholars (e.g., Gomez and Wilson, 2006), is that economic voting is reasonably rational in that it is intended to produce a change in economic policy or outcomes. Thus, it requires a reasonable degree of sophistication about politics and economics.

But, of course, this is not the only possible view. Performance voting in general, and economic voting in particular, may also be expressive – that is, while performance voting is clearly intended to get rid of incumbents, this is not because of a sophisticated understanding about how partisan change leads to policy change (and how this might impact outcomes). Instead, it is an expression of anger, disappointment, or frustration with incumbents who have not performed well. As such, we might even expect that expressive economic voting would be more prevalent among less sophisticated individuals, who, lacking the knowledge needed to cast either prospective or retrospective policy-oriented votes, instead rely on simple rules of thumb: "If I am unhappy then I vote against the incumbent."

Regardless of whether performance voting is rational and expressive, however, all such voting requires that one identify the incumbents; and, as we have seen, some voters cast performance votes based on mistaken perceptions of who is in the cabinet. As we show below, this fact can undermine empirical attempts to adjudicate between a view of economic voting that sees political sophistication as necessary for economic voting versus one in which political sophistication is unnecessary (and perhaps detrimental).

Duch (2001) tested his argument (that political sophistication is necessary for economic voting) in two developing democracies (Hungary and Poland) and found that in Hungary there was a strong positive relationship between his measure of political sophistication and the importance of economic perceptions in an individual's vote choice. In Poland, using a different measure of political knowledge, he did not find a significant relationship. However, looking closely at the measure of political sophistication he used in Hungary, it turns out that variation in the measure (which is a composite of three kinds of questions) depends greatly on a question asking respondents whether they know who is in the cabinet. Thus, it may simply be that voters who know which parties are in the cabinet cast more economic votes than those who do
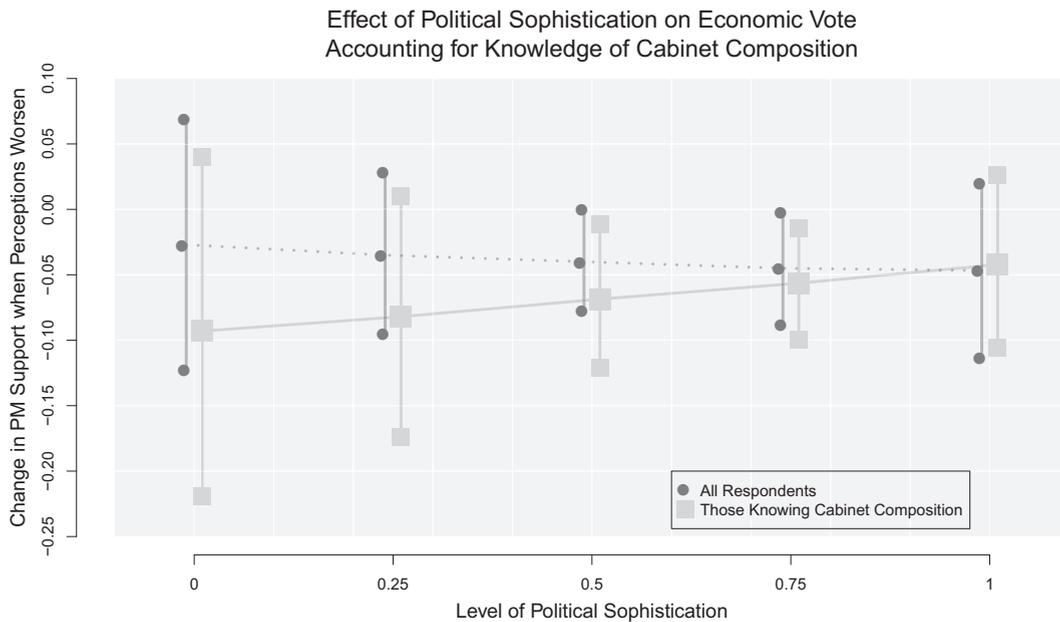
Fig. 2. Effect of general political knowledge on the economic vote Accounting for knowledge of cabinet composition, The Netherlands, 1998. The change in probability comes changing perceived economic performance from "got better" to "stayed the same." The mean estimates are shown with 90% confidence intervals. Circles represent the full sample, squares represent only those respondents that correctly identified the opposition or cabinet status of the five major parties. The points are off-set to ease interpretation.

not know this information – a result that is consistent with our last section, but that cannot distinguish between the rationalist view of economic voting underlying Duch's argument and the alternative expressive view.[12]

Using the Dutch data discussed above, which includes information on respondents' knowledge of cabinet composition as well as more general levels of political sophistication, we can untangle the relationship between these different concepts and economic voting. To do so, we simply add an interactive term to the model that we used to produce our earlier estimates of economic voting. This interaction term conditions the size of economic voting on the respondent's general political knowledge (our measure of general political knowledge, of course, excludes anything about the respondents knowledge of cabinet composition). We estimate two models: one on the full sample and one on the sample of respondents who correctly identified the cabinet/opposition status of the five major parties. If Duch's argument is correct and a general sort of political sophistication is required for economic voting (rather than simply knowing who is in cabinet) then we should see a positive relationship between political sophistication and economic voting in both samples (given we code more economic voting as a larger negative change in vote, the "positive" relationship implied here, will actually be negative in Fig. 2).

Fig. 2 gives the substantive results of the model (the detailed estimates are in the Supplementary Materials).

Looking first at the estimates for the full sample (circles), we see that Duch's result replicates in the Dutch Sample. We get the expected relationship between general political knowledge and economic voting (i.e., more knowledge leads to a more negative change in support for the PM when perceptions worsened). However, once we exclude respondents who do not know who is in cabinet, we get a very different picture: indeed the relationship is (weakly) positive – opposite Duch's prediction.

Indeed, if we look more closely at the results, we see that the real differences in the estimates for the different samples are for less sophisticated voters. What is the difference in the samples for these voters? It is simply that there are a large number of these voters in the full sample who do not know who is in cabinet, but none of them are in the restricted sample. Thus, the result suggests that otherwise unsophisticated voters, who nevertheless know who is in cabinet (and so possess the minimum knowledge necessary to cast economic votes) actually do cast economic votes at high rates – indeed, higher than more sophisticated voters. This is consistent with the idea that these voters are acting expressively and with very little information other than what is absolutely required to performance vote. And, clearly supports the message of this paper – that accounting for knowledge of cabinet composition is critical to understanding economic voting.

Finally, the analyses reported above argue strongly for the inclusion of cabinet knowledge questions on electoral surveys. Specifically, we have demonstrated just a few of the benefits of these questions. We were able to get a more direct view of the effects of political sophistication on performance voting and, in so doing, probe the frequency of expressive economic voting. More importantly, we were

---

[12] In support of this conclusion is the fact that his positive result is not replicated in Poland where the measure of political sophistication did not include cabinet knowledge.

able to provide strong evidence that previous studies have systematically underestimated the strength of the retrospective voting mechanism. The evidence suggests that economic voting, expressive or otherwise, is more widespread than we previously believed. We hope that this article encourages our colleagues to consider cabinet knowledge in their future economic voting research and that survey administrators will facilitate this step forward.

Appendix A. Supplementary data

Supplementary data related to this article can be found at http://dx.doi.org/10.1016/j.electstud.2013.05.001.

References

Duch, R.M., 2001. A developmental model of heterogeneous economic voting in new democracies. American Political Science Review 95 (4), 895–910.

Duch, R.M., Stevenson, R.T., 2008. The Economic Vote: How Political and Economic Institutions Condition Election Results. Cambridge University Press, Cambridge.

Duch, R.M., Stevenson, R., 2010. The global economy, competency and the economic vote. The Journal of Politics 72 (1), 105–123.

Duch, R.M., Palmer, H.D., Anderson, C.J., 2000. Heterogeneity in perceptions of national economic conditions. The American Journal of Political Science 44 (4), 635–652.

Evans, G., Andersen, R., 2006. The political conditioning of economic perceptions. The Journal of Politics 68 (1), 194–207.

Gomez, B.T., Wilson, J.M., 2006. Cognitive heterogeneity and economic voting: a comparative analysis of four democratic electorates. American Journal of Political Science 50, 127–145.

Kayser, M.A., Peress, M., 2012. Benchmarking across borders: electoral accountability and the necessity of comparison. American Political Science Review 106, 661–684.

King, G., Tomz, M., Wittenberg, J., 2000. Making the most of statistical analysis: improving interpretation and presentation. American Journal of Political Science 44 (2), 347–361.

Lewis-Beck, M.S., 2006. Does economics still matter? Econometrics and the vote. The Journal of Politics 68 (1), 208–212.

Lewis-Beck, M.S., Nadeau, R., Bélanger, É., 2004. General election forecasts in the United Kingdom: a political economy model. Electoral Studies 23, 279–290.

Palmer, H., Duch, R.M., 2001. Do surveys provide representative or whimsical assessments of the economy. Political Analysis 9, 58–77.

Wilkin, S., Haller, B., Norpoth, H., 1997. From Argentina to Zambia: a world-wide test of economic voting. Electoral Studies 16 (3), 301–316.

Wlezien, C., Franklin, M.N., Twiggs, D., 1997. Economic perceptions and vote choice: disentangling the endogeneity. Political Behavior 19, 7–17.