Disentangling components of PTV measures, with an application to the EES 2009

By

Mark Franklin (European University Institute) and

Lorenzo De Sio (LUISS Guido Carli – Rome)

Paper prepared for the ECCER conference, Sophia, December 2011

VERY EARLY DRAFT: Please do not cite.

Why do people vote for a party? Conventional analyses of turnout and party support treat the question of why people vote separately from the question why they support a certain political party. Often those who did not vote are omitted from party support analyses – or votes for them are imputed as though they had simply failed to tell us what party they had voted for. In turnout analyses the question of which party was the party voted for is completely by-passed.

Yet it is most unlikely that voters themselves distinguish the two processes. People vote for a party, they do not vote at random – a fact that most "get out the vote" campaigns (especially prevalent in the US) routinely ignore. Yet, at least one likely reason why a particular individual would fail to vote is because he or she did not know who to vote for. This possibility has, however, been completely ignored in the turnout literature.

In this paper we suggest a framework for addressing the relationship between the two questions. To what extent do features of individual-level party support contribute to our understanding of turnout? Is it meaningful to attempt the development of a unified model that accounts for both turnout and party choice simultaneously? Our general intuition is that having a distinct preference for one party over all others, and pronouncing oneself likely to vote for that party, will not only largely determine which party is voted for but also help to determine whether someone votes at all.

As an analysis strategy, we suggest that a possible useful framework could build on two stepping stones. The first is the use of PTV (propensity-to-vote) item batteries, which report individual orientations towards different parties in a party system. Such batteries have been extensively used in the literature that addresses the determinants of *party preferences* and *choice* (Tillie 1995; Eijk and Franklin 1996; Kroh 2003; Eijk et al 2006; Van der Brug et al 2007), by using analyses that attempted at explaining PTV levels in stacked datasets - based on predictors at the *within-respondent* level (by expressing affinities between individual characteristics and specific parties). However, we observe that such PTV measures also incorporate additional information that is not necessarily pertinent for party choice, but could be relevant for other political orientations, such as a generalized individual attitude towards the political parties in a party system, which could possibly be linked to *turnout*. Such PTV measures indeed incorporate a *systematic individual component*, arising from the fact that different individuals can in general assign higher – or lower – PTVs on average to *all* parties, expressing in this way a general attitude towards the party system – disconnected from a specific party preference. It is on this observation that we base our second stepping-stone: the opportunitity to *disentangle* different *variance components* that together make up the overall variance we observe in PTVs.

We address the analysis of these separate components by employingmulti-level models that permit us to distinguish PTV variance at the intra-individual level of analysis (where preferences are formed) from PTV variance at other levels, presumably connected with the decision to vote at the individual level of analysis, and perhaps with a general perception of legitimacy of the party system at the country level (within which both levels are nested).

The paper is structured as follows. We first present the general approach for effectively disentangling different components of PTV variance. We then proceed with our attempt at developing a unified model for both turnout and party choice. We do so by first demonstrating the relevance of vote propensities to turnout and then move on to the corresponding analysis that splits propensity to vote into within- and between-individual components. However, we should warn the reader that – in terms of this more ambitious goal – our intuition does not lead us to a straightforward understanding of the way in which turnout is integrated with party choice. Rather it leads us to a puzzle that we have not yet been able to resolve.

Different components of PTV scores

As anticipated in the introduction, PTV scores have been frequently used in electoral research for estimating voting models in a *generic* fashion, where the importance of several predictors for preference for a *generic* party is assessed. Such analyses routinely employ *affinity* variables (sometimes party distance measures but mostly Y-hats), which express the specific connection between a particular individual characteristic and a particular party (see De Sio and Franklin 2011). As such, they allow even individual characteristic to have a within-individual variance that can be used for explaining within-individual variance in preference towards different parties.

As a result, it has become customary to (1) only see individual predictors in their Y-hat form in voting models estimated in stacked datasets; (2) only see a focus on party preference, without any attention to the fact that the *overall* level of PTVs could vary among respondents; and that such variance could express something meaningful for political attitudes as already explained.

The interesting point is that – in a multi-level modeling framework – such variance may perhaps, paradoxically, be analyzed by resorting to the original respondent-level predictors that were already employed for the creation of y-hat, within-respondent versions. This is so because – in a stacked data matrix where the lowest level is composed of party support measures *within* respondents – individual respondents become a higher level. Variance that cannot be explained by party-specific predictors (e.g. a higher average PTV for a specific respondents) is then transposed to variance in the values of the respondent-level random intercepts. In most models such variance remains unexplained, as only Yhats are included in the analysis, leaving no predictors at the respondent level. We want to learn whether, in order to model such variance, it is sufficient to include in the model the *plain respondent-level predictors*, which act as higher-level covariates. Table 0 presents an example of such an analysis.

The table presents two models, with the first only including predictors that are routinely employed in PTV analyses (party- and party-respondent-level predictors), and the second including new covariates at higher levels: the individual and the country level. Since this an illustrative analysis, only a few comments are needed to highlight some key features. The first is that R-squared does not increase between the models. This is necessary, as variance between individuals and countries is already accounted for in the first model – in the form of random intercepts. However, in this first model such variance is totally *unexplained*. The fact that the second model adds some explanatory power of these between-individuals (and between-countries) variance is instead shown by the decrease of the BLUPs R-squared. This is a measure of the extent to which *random intercepts alone* have explanatory power for the dependent variable. A decrease of this measure testifies to a *shrinkage* of all random intercepts, showing that between-respondents and between-country variance is increasingly explained by new covariates, and there is less need to rely on the variance of random intercepts to explain average respondent PTVs. This tells us that there is a clear goal for explaining higher-level variance components: bringing the BLUPs R-sq to 0, meaning that all random intercepts are zero, so that all differences between individuals (or countries) are explained by individual or country covariates.

Table 0 Example of a multilevel analysis disentangling different components of PTV variance EES 2009

	(1) Yhats only		(2) Yhats + Plain	
Party-Respondent:				
y_female	0.681***	(0.045)	0.646***	(0.045)
y_minority	0.603***	(0.022)	0.603***	(0.022)
y_married	0.661***	(0.057)	0.657***	(0.057)
y_unionmemb	0.523***	(0.026)	0.520***	(0.026)
y_unemployed	0.636***	(0.058)	0.631***	(0.058)
Left-right distance (0-1)	-4.938***	(0.036)	-4.946***	(0.036)
Party_is_close	2.232***	(0.061)	2.253***	(0.061)
Party_best_for_mip	3.521***	(0.031)	3.485***	(0.032)
Party:*				
Party size_01	3.743***	(0.101)	3.760***	(0.101)
Respondent:				
Female			0.077**	(0.024)
Born			0.008***	(0.001)
Church attendance			-0.025*	(0.010)
Religiosity			0.035***	(0.005)
Family income			0.021*	(0.010)
Political awareness			0.033*	(0.014)
Campaign interest			0.060***	(0.016)
Europe gone too far			0.029***	(0.004)
R feels efficatious			0.124***	(0.021)
R fees close to some party			0.064***	(0.014)
Country:				
Average level of turnout			1.856*	(0.849)
Constant	4.774***	(0.128)	-12.699***	(1.547)
N of responses (level 1)	97,642		97,642	
N of respondents (level 2)	20,787		20,787	
N of countries (level 3)	27		27	
R^2	0.501		0.499	
AIC	480138.424		479777.588	
BIC	480347.350		480109.971	
R-squared (BLUPs)	0.147		0.124	

* Effects of party families not reported. The dependent variable is measured on a 0-10 scale, in contrast to the same variable employed in later analyses (see footnote 1).
Note: Coefficients significant at * 0.05, ** 0.01, *** 0.01.

A second comment is that, clearly, sign and size of the effects are not constrained to be similar for any predictor across both levels. As an example, the Yhats included in the table (the strongest, selected from a list that included many more Yhats) do not correspond to the strongest effects predicting overall individual PTV levels. In general, PTV variance at different levels highlights the manner in which the overall variance in PTVs is made up of different components.

Propensity to vote and the "standard" turnout model

The link between propensity to vote and turnout is implicit in the very question that gives rise to the vote propensity measures. Respondents are asked, for each party in their party system, "what are the chances that you would ever vote" for that party. Those who report a low likelihood to vote for any party necessarily have a low and relatively undifferentiated set of vote propensities – the chances are low that the respondent in question will vote for any party and so the chances of them turning out to vote are also low. However, as the propensity to vote for a particular party increases so vote propensities become more differentiated – the gap between support for the most preferred party and support for other parties also increases (see Eijk and Franklin 2009). Because of this increase in support for the most favored party, average support for all parties also increases ($r = 0.41$), but parties other than the most favored party hardly participate in this increase ($r = 0.13$). As an inevitable corollary, there is a strong relationship ($r = 0.72$) between the the size of the maximum PTV and the gap between the maximum and average PTV, with the gap increasing closely in step with the maximum PTV given to any party. The strength of this relationship probably varies across party systems, with multi-party systems in which whole party families share specific portions of the left-right spectrum, seeing some tendency for parallel increases in PTV for members of two or more parties in the same party family, but investigating that speculation must wait on future research.

The critical point for our argument is that as the propensity to vote for some party in particular increases, two things happen. First, the vote propensities become more differentiated, providing leverage on our attempt to predict which party will be voted for. Second, the fact that the respondent reports a higher likelihood of voting for some party means *ipso facto* that he or she is more likely to vote. This provides the link between vote propensity and turnout. A respondent who reports a two out of ten chance of voting for any party has a lower chance of turning out than one who reports a 5 out of ten chance of voting for some party. And a respondent who reports a nine out of ten chance of voting for some party is surely extremely likely to vote. Though vote propensities have generally been employed as means of establishing which is the party most likely to be supported, they should implicitly also provide information about the likelihood of voting at all: for any given individual, a higher maximum vote propensity (which also will generally imply higher mean vote propensities) should correspond to more likelihood of actually voting for that person's most preferred party.

In a hierarchical (multilevel) model with vote propensities at the lowest level of analysis, different mean vote propensities provide variance at the individual level (level 2) because the different mean ptv levels for different respondents is accounted for by the random effects at this level, leaving level 1 variance to uniquely concern the relative likelihood of voting for each party. In this regard, a multilevel model with PTVs at the lowest level is analageous to a conditional logit model when the dependent variable is party choice. In such a model the level 1 variance uniquely relates to which party was chosen because inter-individual differences are removed by the fixed effects that conditional logit enforces. Random effects enforce the same focus at the individual level as do fixed effects, but with the critical difference that the variance across individuals is retained as random intercepts to be investigated at the next level up. In this paper we argue that those random intercepts are related to the likelihood that respondents will vote.

In the existing turnout literature three very clear findings are repeated over and over again. Some people vote because this is what they always do – they are *habituated* (Schmitt and Mannheimer 1991, Plutzer 2002, Franklin 2004); others vote because they care which party wins the election and believe that their votes (along with those of like-minded people) can secure the desired outcome – they are *motivated* (Verba Schlozman and Brady 1995; Franklin 1996; Franklin 2004); yet others vote because they are persuaded to do so by parties, candidates, friends, or television pundits and editorial writers – they are *mobilized* (Rosenstone and Hansen 1993; Verba Schlozman and Brady 1995; Franklin 2004). There are also a variety of individual characteristics that are linked to voting, generally seen as relating to the *resources* that individuals bring to the political realm (Verba Schlozman and Brady 1996). A model containing these four elements can be seen as a "standard model" of voter turnout. The variance in mean vote propensities might be related to any or all of these elements and the purpose of the next section of this paper is to determine to what extent this is the case.

The voter turnout model

In this paper we employ a very simple model that incorporates the four distinct mechanisms listed above that are believed to lead people to vote. Because our objective is not so much to improve our explanation of voter turnout as to discover how vote propensities relate to our existing explanation, we chose to proceed by way of a minimal model containing one measure for each of the four concepts encapsulating the elements we have referred to. We derive these measures from an extensive battery of questions asked in 2009 with the purpose of getting a handle on the reasons for low turnout at European Parliament elections. European Parliament elections are particularly suited to an investigation of this kind because turnout at these elections is quite low, providing maximum variance in a dependent variable that measures whether particular individuals vote or not while also (because of the sheer variety of polities included in the study) providing maximum variation in the characteristics that are expected to explain turnout variations.

In order to derive a parsimonious model, the questions in the turnout battery were submitted to a principal components factor analysis that sought to find four factors that could be unambiguously identified with the four elements listed above. Such a factor structure was found to account for 62 percent of the total variance in a group of 12 variables, as shown in Table 1.

Table 1 Factor loadings, Eigenvalues and variance explained for effects on EP election turnout (N = 20,787)

Variable	Motivated	Mobilized[a]	Resourced	Habituated
Trusts EU institutions	0.831			
Effectiveness of EP	0.960			
Cared which candidate won	0.857			
Cared which party won	0.636			
General political interest		0.763		
Campaign awareness		0.786		
Campaign interest		0.625		
Tried to influence others		0.960		
Church attendance			0.556	
Religiosity			0.875	
Years in current neighborhood			0.401	
Voted in previous national election				0.780
Feels close to a political party				0.761
Eigenvalues	3.737	1.700	1.479	1.122
Cumulative variance explained	0.287	0.418	0.531	0.618

a. Coefficients under 0.4 suppressed for clarity.

Used as independent variables in a basic hierarchical model predicting whether members of the sample voted or not, these four factors explained 23.7 percent of the variance, very close to the 25.5 percent explained by the twelve variables individually. So simplifying the data in this way did not cost much in terms of variance explained.

Table 2 shows the basic model referred to above, with the addition of two country-level covariates intended to account for the two most important country-level reasons for differences in turnout: first that some countries employ compulsory voting, ensuring relatively high turnout even at European Parliament elections, and second that some countries – post-communist countries – only recently began their transitions to democratic life and still fail to show all the characteristics of mature democracies.

The coefficients in Table 2, being logit coefficients, are not readily interpretable but, to give some notion of the scale involved, the coefficient for compulsory voting in Model A corresponds to a 17 percent higher turnout in those countries when other variables in the model are held at their mean values – about the same as the effect of mobilization and about twice the effect of habituation.

Table 2 Effects of various stimuli on turnout, hierarchical analysis with logit link (standard errors in parenthases)

Independent variable	Model A		Model B	
	b	s.e.	b	s.e.
Compulsory voting country[a]	0.825	(0.252)**	0.879	(0.258)***
Post-communist country	-0.685	(0.186)***	-0.652	(0.190)***
Motivation factor	0.642	(0.018)***	0.553	(0.020)***
Mobilization factor	0.919	(0.021)***	0.910	(0.022)***
Resources factor	0.232	(0.020)***	0.219	(0.021)***
Habituation factor	0.365	(0.018)***	-0.066	-0.044
Maximum vote propensity			0.099	(0.009)***
Maximum * Habituation			0.051	(0.005)***
Constant	1.46	(0.125)***	0.576	(0.150)***
N of individuals (level 1)	20,787		20,095	
Number of countries (level 2)	27		27	

a. Italy coded 0.5 since compulsory voting was abolished there 20 years ago.
Note: Coefficients significant at * 0.05, ** 0.01, *** 0.01.

However, the most important implication of the table does not require us to know the magnitudes of the effects. We see clearly in Model B that only one coefficient changes its value to any notable extent when compared to the coefficient for the same variable in Model A. The effect of introducing maximum vote propensity and its interaction with habituation is to render the main effect of habituation not significantly different from zero. The comparison of the two models, shown in Figure 1, indicates that, for lower values of maximum vote propensity, the increase in likelihood of turning out with increasing habituation is not significant (the lower bound never exceeds the upper bound). However, with larger values of maximum vote propensity the effect of increases in habituation is very strong. Intermediate values of maximum vote propensity evidently see increasingly strong relationships between habituation and turnout, as the level of vote propensity increases

Figure 1 Effect of habituation on expected turnout with minimum and maximum vote propensity

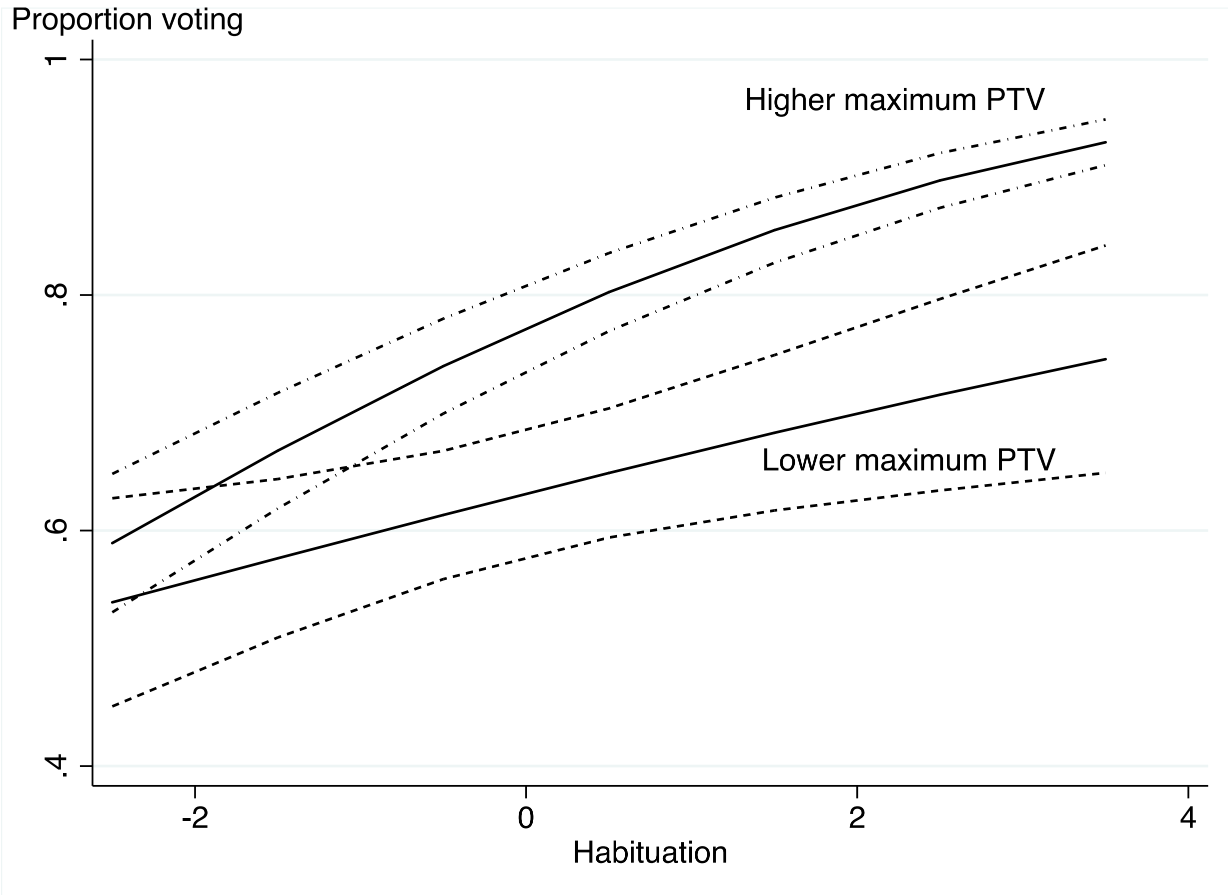

and the slope becomes significant when vote propensity exceeds a value of approximately 4 (not shown). This tells us something relevant about the importance of the party system. Only to the extent that habituality is able to be direct towards some party, it is able to express itself into an actual turnout behavior. The proportions voting shown on the Y axis are, of course exaggerated since what we are dealing with here is reported turnout, which inevitably over-reports actual turnout.

The implication of this finding is that the effect of habituation is dependent on vote propensity. It is not that habituated people are those with high maximum PTV values (the correlation between maximum PTV and habituation is only 0.4). But habituation appears to play a role in determining voter turnout only for people with relatively high maximum vote propensities.

What about the other way around? Do the factors that determine turnout levels play any role in determining the maximum level of party support? What we see in Table 3 is that very

Table 3 Effects of various factors on maximum propensity to vote, hierarchical analysis (standard errors in parentheses)

Independent variables	Model C b s.e.	Model D b s.e.
Motivation factor	-0.054 (0.001)***	-0.035 (0.002)***
Mobilization factor	0.033 (0.001)***	0.007 (0.003)*
Resources factor	0.006 (0.002)***	-0.215 (0.025)***
Habituation factor	0.095 (0.002)***	0.072 (0.003)***
Campaign interest		0.025 (0.003)***
Party close to respondent		0.028 (0.003)***
Female		0.020 (0.003)***
Single		0.017 (0.004)***
Church attendance		0.074 (0.009)***
Religiosity		0.041 (0.005)***
Constant	0.8293 (0.048)***	0.318 (0.052)***
Number of respondents	20,340	20,340
Number of groups	27	27
Level 1 variance explained	0.255	0.264

a. Italy coded 0.5 since compulsory voting was abolished there 20 years ago.
Note: Coefficients significant at * 0.05, ** 0.01, *** 0.01.

much the same factors affect propensity to vote, explaining very much the same amount of variance, as the factors explaining turnout.[1] When we add individual-level variables customarily associated with turnout but already largely subsumed in these factors they add very little

[1] The dependent variable has been normalized to range from 0 to 1, so as to ensure comparability with effects on turnout.

additional variance.[2]

These findings give us some confidence in the expectation of beinga ble to build a unified model that will provide a role for factors relevant to both turnout and party choice.

Overall versus specific party support

When people are asked “What are the chances that you would ever vote for” each of the parties in their party system, and provided with an 0-10 scale on which to register their answers, those answers will to some extent be ideocyncratic. Some people may conceive of a score of 8 or 9 as indicating virtual certainty, whereas others may feel the need to give a score of 10 to exactly the same sense of certainty. Similarly with low scores. Some may give a score of 0 to parties they think they will never vote for. Others may more cautiously ascribe a score of 1 or 2 when they have precisely the same feeling about the party concerned. More importantly, people may give a generally higher or lower score to all parties, depending on their general orientation towards their party systems. Much of this variation may be random and unaccountable. However, some of it may be systematic and may be used to increase our understanding of political processes. People who live in countries where voting is compulsory would be expected to pronounce themselves more likely to vote for some party even if they do not have a very high liking for any party. The same may well be true of other country-level variables.

More importantly, variables that help to explain differences in party support may also contribute to an explanation of turnout variations. Certain variables may make people feel closer to or further from all parties – they may contribute to our understanding of overall party support as well as to specific party support. In the next section of this paper we will elaborate the turnout model introduced above by placing it at level 2 of a hierarchical model in which level 1 focuses on the within-individual determinants of party support.

The party support model

Why people support a party is a question that can only be asked in the context of a stacked dataset where the parties between which respondents are dividing their support are operationalized at a

[2] But note that one of the factors changes its sign, suggesting spurious effects due to including in Model D variables already included Model C – a consideration that will become relevant for our concluding discussion.

lower level of analysis than the respondents themselves. At that level, the unit of analysis is the party, and the object of investigation is the process by which party support is generated. Independent variables have to be reformulated as affinities, most usually as distances of each respondent from the party concerned, in terms of some metric, or as y-hat predictors of the propensity to vote for each party on the basis of independent variables that do not lend themselves to reformulation as distances (de Sio and Franklin 2011).

Table 4 Effects of individual-specific variables on level 2 intercepts in a hierarchical model of party support, compared to effects of the same variables on turnout (from Table 2)*

Independent variable	First differences from Model B	Model E		Model F	
		b	s.e.	b	s.e.
Compulsory voting country[a]	0.12	0.024	(0.039)	0.038	(0.034)
Post-communist country	-0.09	-0.079	(0.029) **	-0.041	(0.025)
Motivation factor	-0.10	-0.011	(0.001) ***	-0.014	(0.001) ***
Mobilization factor	0.14	0.004	(0.001) ***	0.008	(0.001) ***
Resources factor	0.02	0.005	(0.001) ***	0.004	(0.001) **
Habituation factor	0.04	0.000	(0.001)	0.009	(0.001) ***
Left-right distance				-0.499	(0.004) ***
Party for most important problem				0.348	(0.003) ***
Party close to respondent				0.226	(0.006) ***
Female y-hat				0.068	(0.005) ***
Minority y-hat				0.060	(0.002) ***
Married y-hat				0.067	(0.006) ***
Union member y-hat				0.052	(0.003) ***
Unemployed y-hat				0.065	(0.006) ***
N of responses (level 1)		97,642		97,642	
N of Individuals (level 2)		20,787		20,787	
N of countries (level 3)		27		27	
Level 1 variance explained		0.500		0.500	

* Effects of party families not reported. For scaling of the dependent variable, see footnote 1.
Note: Coefficients significant at * 0.05, ** 0.01, *** 0.01.

In Table 4, however, we start with a model that contains no within-respondent predictors, and compare that with the first differences that correspond to the coefficients already presented for Model B in Table 1. What this comparison shows is that effects of country and individual-level variables on intercepts are about ten times less than the same effects on turnout from Table 2. Even more unexpected to us is the fact that these effects on the whole become greater when party-level variables are controlled for in Model F (and the habituation factor, which failed to reach significance in Model E, does so in Model F).

Lack of significance for the habituation factor in Model E might somehow have been explained by the fact that Table 3 contains an analysis of PTVs, for which habituation was found, in Table 2, to play no role. But the role of habituation is restored in Model F, for reasons that currently elude us.

Discussion

This paper set out to investigate whether a unified model could be constructed that would unite an analysis of individual party support at the lowest level of a hierarchical model with an analysis of overall party support at the next level up. The expectation of some degree of success in this endeavor was engendered by preliminary analyses of the roles played by vote propensities in the explanation of turnout variations and of the role played by the determinants of turnout variations in determining the maximum level of vote propensities. However, the hierarchical model does not give us findings that we can make sense of in these terms. Variations in individual-level intercepts can be explained to some extent by the same factors that explain turnout variations, and these variables are indeed the same as those that explained variations in maximum vote propensities in Table 3. But the relative importance of variables deriving from factor analysis as compared to simple demographics is reversed in the model that includes affinity and other party-level independent variables.

Clearly, even though we know a lot about the determinants of individual vote propensities (and our findings do not cast doubt on these more traditional uses of PTV analysis), there is much we do not know about how hierarchical models of vote propensity data relate to turnout.

One thing we discovered, which perhaps should have been obvious but escaped us until it hit us between the eyes, is that the introduction of affinity measures into a multi-level model itself causes perturbations in the higher level intercepts which can then give rise to spurious effects of individual-specific variables. Table 5 shows what happens when the mean of the dependent variable is introduced as a predictor (the first independent variable in the table). We see that in the first model (Model G), which contains no affinity or party-level variables, the result of introducing this independent variable is to (not unexpectedly) remove virtually all the variance in level 2 intercepts. Not only do none of the individual-specific variables included in that model prove statistically significant, but the residual unexplained variance (R-squared BLUPs in the bottom row of the table) proves virtually non-existent.

However, when we introduce party-level and affinity measures in Model H, all but one of the individual-specific variables become significant, suggesting that variance in the level 2

Table 5 Effects of individual-specific variables on level 2 intercepts in a hierarchical model of party support that contains the mean level of vote propensity as an independent variable.

	Model G		Model H	
	b	s.e.	b	s.e.
Mean propensity to vote	1.000	(0.001) ***	0.890	(0.007) ***
Female y-hat			0.056	(0.004) ***
Minority y-hat			0.061	(0.002) ***
Married y-hat			0.058	(0.005) ***
Union member y-hat			0.053	(0.002) ***
Unemployed y-hat			0.056	(0.005) ***
Partysize (scaled 0-1)			0.374	(0.010) ***
Let-right distance (scaled 0-1)			-0.402	(0.003) ***
Party_close to respondent			0.213	(0.006) ***
Party for most important problem			0.346	(0.003) ***
Motivation factor	0.001	(0.001)	0.004	(0.001) ***
Mobilization factor	0.000	(0.001)	0.003	(0.001) **
Resources factor	-0.001	(0.001)	-0.001	(0.001)
Habituation factor	0.000	(0.001)	0.006	(0.001) ***
Constant	-0.003	(0.045)	-0.174	(0.102)
N of responses (level 1)	97,642		97,642	
N of respondents (level 2)	20,787		20,787	
N of countries (level 3)	27		27	
R^2	0.218		0.507	
R-squared (BLUPs)	0.001		0.011	

* Effects of party families not reported. For scaling of the dependent variable, see footnote 1.
Note: Coefficients significant at * 0.05, ** 0.01, *** 0.01.

intercepts has been introduced by the inclusion of party-level covariates. That such variance should be introduced makes sense, as these covariates will be more successful at accounting for the behavior of some individuals than of others, but the important point is that the variance introduced in Model H is not random in relation to the individual-specific covariates. This might well be because those covariates (or at least some of them) were represented among the variables that generated the individual-specific variables in the factor analysis reported in Table 1. At all events, the variance introduced in this way appears not to be random in relation to those particular covariates since they do explain some of it at a relatively high level of significance. This implies that the coefficients we see for the individual-specific factor variables in Model H will be spurious, much like the change in sign that is evident in Table 3 when covariates are inserted into a model that already contained variables built on those covariates (see Footnote 2).

Coming back to the analyses reported in Table 5, a next step in our investigations must evidently be to see what happens if we rigorously segregate the individual-specific variables from the affinity variables in any analysis of this kind. If that does not move us forward we might need to consider the possibility that a multi-level model does not provide an appropriate vehicle for disentangling the different sources of variance in a complex causal framework.

References

Brug, Wouter van der, Cees van der Eijk and Mark N. Franklin (2007). *The Economy and the Vote: Electoral Responses to Economic Conditions in 15 Countries,* (New York: Cambridge University Press).

De Sio, Lorenzo, and Mark Franklin (2011) "**Generic variable analysis:** climbing the ladder of generality with social science data". Paper prepared for the ECCER Conference, Sofia, December.

Eijk, Cees van der, and Mark N. Franklin. 2009. *Elections and Voters*. Houndmills: Palgrave.

Eijk, Cees van der, Wouter van der Brug, Martin Kroh, and Mark N. Franklin. 2006. "Rethinking the dependent variable in electoral behavior - on the measurement and analysis of utilities." *Electoral Studies*.

Eijk, Cees van der, and Mark N. Franklin. 1996. *Choosing Europe? The European electorate and national politics in the face of union*. Ann Arbor: University of Michigan Press.

———. 2009. *Elections and Voters*. Houndmills: Palgrave.

Kroh, M. 2003. Parties, Politicians, and Policies: Orientations of Vote Choice Across Voters and Contexts, University of Amsterdam, Amsterdam.

Oppenhuis, Eric. 1995. *Voting Behavior in Europe: A Comparative Analysis of Electoral Participation and Party Choice*. Amsterdam: Het Spinhuis.

Tillie, Jean. 1995. *Party Utility and Voting Behaviour*. Amsterdam: Het Spinhuis.