Do Voters Benchmark Economic Performance?

VINCENT AREL-BUNDOCK, ANDRÉ BLAIS AND RUTH DASSONNEVILLE*

Keywords: economic voting, benchmarking, relative evaluations, reference point, performance voting

The conventional theory of economic voting is that voters reward or punish the incumbent government based on how the domestic economy is doing. Recently, scholars have challenged that view, arguing that voters use relative assessments to gauge government performance. From this perspective, what matters is not how well the national economy is doing *per se*, but rather how it performs relative to an international or historical reference point.

This article revisits prominent published works in that emerging tradition, and finds that the available evidence does not support the benchmarking hypothesis. We come to this conclusion after taking a close look at the regression models that are typically used to test benchmarking. We show algebraically that the way in which those models are specified invites a fundamental misreading of the evidence. Finally, we propose an alternative regression equation which can be used to test benchmarking, avoids common misinterpretations, and allows us to assess complex, conditional theories of relative evaluation.

BACKGROUND

Economic voting is one of the most important accountability mechanisms at work in electoral democracies. The fact that voters reward or punish the incumbent government based on how the domestic economy is doing¹ is traditionally viewed as normatively desirable, for it reflects popular control of representatives.²

Recently, a number of scholars have challenged this optimistic view, by pointing out that domestic economic growth is often a weak proxy for government performance. When the local economy moves in sync with secular trends or global shocks, governments may be rewarded or punished for events beyond their control.³ This is especially true in integrated economies, where domestic fortunes are tightly linked to events abroad, and responsibility is blurred.⁴ Democratic accountability thus requires more from voters than a simple response to local economic conditions.⁵

*All three authors are from the Département de Science Politique, Université de Montréal. Contact: vincent.arel-bundock@umontreal.ca, andre.blais@umontreal.ca, ruth.dassonneville@umontreal.ca. We thank Erdem Aytaç, Timothy Hellwig, Michael Lewis-Beck, seminar participants at Texas A&M University, and our regular lunch companions Chez Valère. The data, replication instructions, and the data's codebook can be found at <http://dx.doi.org/10.7910/DVN/OCNIWD>

¹Lewis-Beck 1988.

²Fiorina 1981; Key 1966; Przeworski, Stokes, and Manin 1999.

³Bartels 2012.

⁴Duch and Stevenson 2010; Fernández-Albertos 2006; Hellwig and Samuels 2014.

⁵Achen and Bartels 2016; Anderson 2007.

A new and influential strand of research argues that voters do, in fact, make rational judgements about government performance, because their evaluations are relative.⁶ What matters to rational voters may not be how well the national economy is doing *per se*, but rather how it performs relative to the economies of other countries, or relative to some historical benchmark.

Benchmarking is a powerful idea, which can be traced back to the work of Powell and Whitten. As these authors point out, voters are likely to “evaluate government relative to some expectations about how the economy should have performed”.⁷ But since expectations are difficult to measure, “it seems reasonable to use the international average levels of growth, inflation, and unemployment to estimate a baseline against which each country’s citizens could judge the performance of their own economy.” This approach is intuitive, since “abundant research in other domains of social science supports the proposition that individuals are sensitive to comparative assessments.”⁸

Yet, there are also good reasons to doubt that voters benchmark economic performance. First, the benchmarking hypothesis is at odds with a dominant view on the cognitive limitations of ordinary voters. Indeed, a long tradition of research in political science has depicted the citizenry as poorly informed⁹ and biased.¹⁰ It is difficult to imagine how such an unsophisticated electorate could systematically and accurately compare how well the national economy is performing relative to other countries or past history. Second, even if some authors posit that the media could facilitate benchmarking by making implicit comparisons in their news coverage, the evidence for the underlying mechanism is rather weak. For instance, Kayser and Peress (henceforth, KP) report that high-information voters – those most exposed to media – do not engage in more benchmarking than low-information voters.¹¹ Finally, some empirical studies claim that voters act based on relative economic conditions, but others find that when it comes to evaluating government performance, “the effect of luck is larger than the effect of competence.”¹² In short, the theoretical case for benchmarking is muddled, and the empirical record is mixed.

In this article, we show that the empirical evidence of benchmarking is extremely weak. We argue that the way in which regression models are typically specified to test benchmarking is needlessly complicated, that it invites a fundamental misreading of the evidence, and that it often leads researchers astray. We propose a simpler model specification which can be used to test benchmarking, avoids common misconceptions, and carries powerful intuitions about the theory. We also show how this simple model can be enhanced to test more complicated theories, such as when voters benchmark against multiple reference points, or when the strength of benchmarking depends on the context. We revisit a prominent empirical study of *Benchmarking Across Borders*,¹³ and conduct a faithful replication of all the models reported in that article. When correctly interpreted, the results do *not* support the contention that voters make rational comparative evaluations.

Our findings have important implications for the field of economic voting, and for our understanding of the mechanisms that underpin democratic accountability. More generally, our article makes useful contributions to political science methodology, by highlighting the shortcomings of a widely used empirical strategy, and by proposing a better way to test theories of relative evaluation.

⁶Aytaç 2018; Ebeid and Rodden 2006; Kayser and Peress 2012; Wolfers 2002.

⁷Powell and Whitten 1993, 396.

⁸Kayser and Peress 2012, 664.

⁹Converse 2000; Zaller 1992.

¹⁰Achen and Bartels 2016; Taber and Lodge 2006.

¹¹Kayser and Peress 2012.

¹²Leigh 2009.

¹³Kayser and Peress 2012.

BENCHMARKING VS. CONVENTIONAL ECONOMIC VOTING

The core intuition of benchmarking is illustrated in Figure 1, where we show how domestic growth and a reference point can affect support for the incumbent party. The solid line represents the domestic growth rate during the election year (G_y), and the dashed line represents the growth rate that voters use as a benchmark to evaluate the incumbent government's performance. Depending on the analyst's theory, the reference point could be the international growth rate (G_i) or the historical level of growth in the country under study (G_h).

Figure 1: Marginal effects of domestic economic growth and benchmark growth on votes for the incumbent.

The conventional view of economic voting is that votes for the incumbent (V) are tied to domestic growth. As we move from left to right, G_y increases in Figure 1a but stays constant in 1b. Thus, the conventional prediction is that votes for the incumbent will increase in 1a but stay constant in 1b.¹⁴

In contrast, proponents of benchmarking argue that what matters to voters is not domestic growth *per se*, but rather the difference between domestic growth and the benchmark ($G_y - G_i$).¹⁵ When the solid line is above the dashed line, domestic growth outperforms the benchmark, and voters should reward the incumbent. When the solid line is below the dashed line, domestic growth underperforms relative to the benchmark, and voters should punish the incumbent. As we move from left to right in Figure 1, the “performance gap” or “competence signal” increases in Figure 1a and decreases in 1b. Thus, benchmarking predicts that votes for the incumbent will increase in 1a and decrease in 1b.

¹⁴Many proponents of the conventional view would remain agnostic with respect to Figure 1b.

¹⁵For simplicity, we discuss international benchmarking, but if the analyst is interested in historical benchmarking, the relevant performance gap would be $G_y - G_h$.

These expectations can be restated using the language of multiple regression. When domestic growth increases and the reference point is held constant (Figure 1a), both theories predict that the incumbent's vote share will increase. In other words, both theories predict that the marginal effect of domestic growth will be positive: $\partial V / \partial G_y > 0$. When the reference point increases and domestic growth is held constant (Figure 1b), benchmarking predicts that votes for the incumbent will decrease. In other words, benchmarking predicts that the marginal effect of the reference point will be negative: $\partial V / \partial G_i < 0$.

If we hope to discriminate between benchmarking and conventional economic voting, *the main quantity of interest is the marginal effect of the reference point*, since this is where benchmarking theory makes a distinctive prediction.

HOW DO SCHOLARS TEST BENCHMARKING?

Conventional theories of economic voting are typically tested using models of this form:

$$V = \beta_y G_y + \Psi \Omega + \nu, \quad (1)$$

where V is the incumbent's vote share; G_y is the domestic economic growth rate during the election year; Ω is a vector of control variables; and ν is a disturbance term. Clearly, Model 1 cannot be used to test benchmarking, since it ignores relative evaluations altogether.

In their seminal article, Powell and Whitten¹⁶ estimate a regression equation of this form:

$$V = \lambda_{y-i}(G_y - G_i) + \Phi \Omega + \nu, \quad (2)$$

with G_i equal to a "reference point", the international economic growth rate. Model 2 takes us very close to the benchmark story: When the gap between G_y and G_i is positive, the domestic economy outperforms the global economy, and voters should reward the incumbent government for its competence.

As KP note, however, Model 2 cannot be used to distinguish between benchmarking and conventional economic voting, because it suffers from omitted variable bias.¹⁷ Indeed, the composite variable $(G_y - G_i)$ is highly correlated with the level of domestic economic growth (G_y).¹⁸ As a result, we cannot parse out the effect of benchmarking from conventional economic voting, and λ_{y-i} captures both phenomena. Model 2 is thus useful if we want to estimate something akin to the "total effect" of domestic growth and benchmarking on voting behavior, but not if we wish to compare and contrast the two theories.

To solve this problem, KP introduce an additional control for the reference point:

$$V = \theta_{y-i}(G_y - G_i) + \theta_i G_i + \Gamma \Omega + \varepsilon. \quad (3)$$

In this model, $G_y - G_i$ represents a "decomposed" or "local" component of growth, whereas G_i aims to control for changes in the reference point. Model 3 has had tremendous influence in the field. At the time of writing, KP's article has been cited over 150 times, and several other researchers have adopted and adapted their empirical strategy.

¹⁶Powell and Whitten 1993.

¹⁷Kayser and Peress 2012, 663.

¹⁸In KP's dataset the correlation coefficient between G_y and $(G_y - G_i)$ equals .83.

A WIDESPREAD MISCONCEPTION

An intuitive – but ultimately incorrect – way to interpret Model 3 would be to focus on the gap between G_y and G_i , and to treat the θ_{y-i} coefficient as the effect of relative economic performance on votes for the incumbent.

For example, Aytaç argues that a positive estimate of θ_{y-i} provides “evidence for the hypothesis that voters reward (punish) incumbents on whose watch the economy performs relatively better (worse) in domestic and international comparisons.”¹⁹ KP define “local growth” as the gap between G_y and G_i , and interpret θ_{y-i} as measuring the association between “an increase in local growth” and an “increase in the leader party’s vote share”.²⁰ Goplerud and Schleiter follow in those footsteps, and discuss θ_{y-i} as the effect of some “benchmarked” or “local” component of growth on voting behavior.²¹ Using data on the American states, Ebeib and Rodden also interpret θ_{y-i} as the effect of “relative state conditions” on votes.²²

If the domestic economy outperforms a reference point, it may be reasonable for voters to infer that the government is doing good work. In that spirit, Leigh treats θ_{y-i} as the effect of “government competence” on votes for the incumbent.^{23,24} In the American context, Wolfers considers the gap between state and national-level economic growth, and calls θ_{y-i} the “effect of competence.”²⁵

Interpreting θ_{y-i} as the effect of relative economic performance on votes for the incumbent appeals to common sense, but it is a mistake. The root of the problem lies in the fact that G_i appears twice on the right-hand side of Equation 3. This redundancy changes the substantive meaning of our regression coefficients.

To see how, take the partial derivative of Equation 3 with respect to G_y , and find the marginal effect of domestic growth:

$$\frac{\partial V}{\partial G_y} = \theta_{y-i}. \quad (4)$$

This simple exercise demonstrates that the coefficient associated with $G_y - G_i$ is exactly equivalent to the marginal effect of G_y . Against intuitive common sense, θ_{y-i} does *not* measure the effect of relative economic performance on votes for the incumbent. Since θ_{y-i} is the marginal effect of domestic growth, finding a positive coefficient for “benchmarked” or “local” growth is actually supportive of conventional economic voting.

Tests of benchmarking based on Equation 3 have been repeatedly misinterpreted in prestigious scientific journals, by leading scholars of economic voting. The inclusion of duplicate regressors on the right-hand side of Equation 3 has been a source of widespread confusion in the economic voting literature.²⁶

To put this confusion to rest, we need a simpler, more direct test of benchmarking.

¹⁹Aytaç 2018.

²⁰Kayser and Peress 2012, 669.

²¹Goplerud and Schleiter 2016, 444.

²²Ebeid and Rodden 2006, 537-539.

²³Leigh 2009.

²⁴The quantities are interpreted slightly differently in Leigh’s model, since the author uses a conditional logit model, but the methodological problem remains the same.

²⁵Wolfers 2002.

²⁶For a short review of economic benchmarking, see Healy and Malhotra (2013), 296-297.

A SIMPLER TEST OF BENCHMARKING

From Figure 1, we learned that both benchmarking and conventional economic voting predict that the marginal effect of domestic growth should be positive. In contrast, only benchmarking predicts that the marginal effect of international growth should be negative. The simplest and most direct way to test those predictions is to estimate a model of this form:

$$V = \delta_y G_y + \delta_i G_i + \Gamma \Omega + \varepsilon. \quad (5)$$

Since Model 3 includes redundant regressors, it carries no more information than the simpler Model 5. In fact, Models 3 and 5 are perfectly equivalent from a logical standpoint, and they produce identical numerical results: The marginal effect of domestic growth, the marginal effect of international growth,²⁷ the intercept, the control variables' coefficients, the residuals, and all fit statistics are always the same in both models. In the online appendix, we present side-by-side estimates using Models 3 and 5 to illustrate this point numerically.

Yet, even if the two models are formally equivalent, the simpler specification has major advantages in terms of transparency, presentation, and interpretation.

First, the correct interpretation of KP's Model 3 is highly counter-intuitive: The coefficient that they call "Local Component of Growth" in their regression tables (θ_{y-i}) does *not* measure the effect of the local economy's relative performance on votes for the incumbent. As we showed above, this has been a major source of confusion in the field, and benchmarking results have been repeatedly misinterpreted in print. Our simpler specification avoids this problem.²⁸

Second, Equation 5 directly translates the theoretical intuitions conveyed by Figure 1, and it immediately reveals the relevant test statistics. Recall that the discriminating test of benchmarking is that international growth should have a negative marginal effect on votes for the incumbent. In Model 5, the marginal effect of international growth is the δ_i coefficient, and we can simply look at its p-value. In Model 3, the marginal effect of international growth is a linear combination of coefficients ($\partial V / \partial G_i = \theta_i - \theta_{y-i}$), and we must conduct an extra Wald test to know if that combination is negative and statistically significant.

Finally, as we show below, our simpler specification offers a solid foundation on which we can build empirical tests for theories of benchmarking where voters compare multiple reference points, or where the strength of benchmarking is context-dependent.

REPLICATION: BENCHMARKING ACROSS BORDERS

As we explained above, the key quantity of interest for tests of benchmarking is the marginal effect of international growth (holding domestic growth constant). Unfortunately, KP do not consistently report the statistics that are needed to test if that quantity is distinguishable from zero.²⁹ As a result, readers cannot assess the strength of the evidence simply based on the findings printed in *Benchmarking Across Borders*.

²⁷It is easy to show algebraically that $\delta_y \equiv \theta_{y-i}$, and that $\delta_i \equiv \theta_i - \theta_{y-i}$.

²⁸Although KP's model includes a single variable to represent the difference between G_y and G_i , it does not offer a single parameter to measure the association between that gap and votes for the incumbent.

²⁹KP only report the results of the relevant Wald test of coefficient equality for 3 of their 24 empirical models.

To see if KP's data support their theory, we re-estimated all of their models using the authors' replication files, and we computed all the quantities of interest.³⁰ Table 1 shows the results for four models,³¹ estimated using KP's preferred measure of international growth (an index constructed via principal components analysis).

TABLE 1: OLS regression models with incumbent vote share as dependent variable.

	Baseline	Controls	Lag	Lag + FE
Domestic Growth (G_y)	1.261*** (0.352)	0.612*** (0.234)	0.529** (0.216)	0.636* (0.323)
Global Growth (G_i)	-1.306*** (0.466)	-0.561 (0.424)	-0.384 (0.382)	-0.274 (0.481)
Local Unemployment	-0.335 (0.229)	-0.041 (0.187)	-0.252 (0.169)	0.186 (0.278)
Global Unemployment	-0.130 (0.264)	-0.328 (0.216)	-0.327* (0.178)	-0.480 (0.320)
Coalition size		-3.333*** (0.714)	-1.506** (0.594)	-1.398** (0.611)
Eff.Num.Parties		-2.774*** (0.599)	0.453 (0.548)	2.417** (0.875)
Population		0.000** (0.000)	0.000 (0.000)	0.000 (0.000)
Year		0.035 (0.056)	0.006 (0.043)	-0.045 (0.062)
Leader vote lag			0.765*** (0.077)	0.756*** (0.084)
Constant	35.069*** (1.997)	-16.506 (111.444)	-2.785 (84.454)	85.467 (119.667)
Observations	213	189	189	189

Robust standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Baseline specification

In column 1 of Table 1, we see that the G_y coefficient is positive. This is consistent with both benchmarking and conventional economic voting. The G_i coefficient is negative and statistically significant. This is consistent with benchmarking.³² However, those results are not credible, because the model in column 1 is fatally under-specified.

³⁰All models were estimated using our simpler Equation 5, but since Models 3 and 5 are strictly equivalent, our conclusions are unaffected when we estimate the original models.

³¹Column 1 corresponds to Table 1 Model 3 in *Benchmarking Across Borders*. Column 2 corresponds to Table 3 Model 2. Column 3 corresponds to Table 3 Model 5. Column 4 corresponds to Table 3 Model 8.

³²Note that there is no evidence of benchmarking with respect to unemployment, even in this model.

Controls, lags, and fixed effects

Ensuring that results are robust to the inclusion of controls and a lagged dependent variable is a minimum standard for most modern research on economic voting. In Column 2 of Table 1, we follow KP and add the same control variables as in their article; column 3 includes the incumbent's vote share in the previous election; and column 4 includes both a lagged dependent variable and country fixed effects.

The three new models are consistent with conventional economic voting: The marginal effects of domestic growth in columns 2 to 4 are all positive and statistically significant. However, none of the three models supports benchmarking: The marginal effects of international growth in columns 2 to 4 are all indistinguishable from zero. As soon as we introduce control variables, a lagged dependent variable, or country fixed effects – widely recognized best practices in the field – the evidence of benchmarking evaporates.

Alternative measures of international growth

The models in Table 1 were all estimated using an index of international growth constructed by principal components analysis. This is KP's preferred measure of G_i , but the authors also consider two alternatives: A trade-weighted average of growth rates around the world, and the international median.

In *Benchmarking Across Borders*, the choice between those three measures is rather inconsequential, because KP conclude that the evidence supports benchmarking, regardless of how they measure G_i . Substantively, the authors take this to mean that “voters respond to their country's deviation from various measures of average international performance”.³³ Moreover, KP do not offer a real theoretical defense of their preferred measure, and fit statistics do not give us strong reasons to favor one measure of G_i over another.³⁴

Nevertheless, access to these two alternative measures of international growth is useful, because it allows us to probe the sensitivity of benchmarking tests to how we measure the reference point. In the online appendix, we replicate the eight regression models that KP estimated using aggregate-level data and their two alternative measures of international growth. None of those eight models shows evidence of benchmarking: The marginal effect of international growth is never distinguishable from zero.

Individual-level survey data

Moving beyond aggregate-level data, KP also study benchmarking using individual-level surveys. Once again, their empirical specification resembles Model 3, and the quantity of interest is the marginal effect of international growth. In the online appendix, we replicate all 12 of KP's individual-level models. None of those models allows us to reject the null of “no benchmarking.”

³³Kayser and Peress 2012, 669.

³⁴Raftery (1995) considers that there is “strong” evidence against a model when its BIC is 6 to 10 points larger than another model. In Table 2 of their article, KP report that the gap in BIC between the principal components and the median growth models is between 3.2 and 4.4.

Three more empirical claims

In the online appendix, we consider three more empirical claims from the original article: (1) A statistically insignificant estimate of θ_i constitutes evidence of “full benchmarking”; (2) the substantive effect of decomposed growth is more important than the substantive effect of domestic growth; (3) at several points in time, the magnitude of the benchmarked economic vote is greater than the magnitude of the non-benchmarked economic vote. Our assessment is that these claims do not add credence to the theory.

Do voters benchmark economic performance?

In their article, KP “argue that previous research has fundamentally misunderstood and hence incorrectly estimated how economic assessments are made.”³⁵ They contend that “voters respond more to national deviations from an international average rate of growth than to the growth rate itself.”³⁶ They claim that their empirical analysis reveals “strong evidence of cross-national benchmarking on economic growth both at the aggregate and at the individual level, across time periods, and across subsamples.”³⁷ Finally, after conducting extensive robustness checks, they conclude that their “main results are not altered”³⁸, and that the evidence is “clearly inconsistent with no benchmarking.”³⁹

We re-evaluated benchmarking on KP’s own terms: Using their original data, logically equivalent statistical models, the same null hypothesis testing framework, and an evaluation criterion that they explicitly endorsed.⁴⁰ Yet, our substantive conclusions are strikingly different.

When models include control variables, a lagged dependent variable, or country fixed effects, we cannot reject the null of “no benchmarking.” When we use alternative measures of international growth, we cannot reject the null of “no benchmarking.” When we test the theory using individual-level survey data, we cannot reject the null of “no benchmarking.” In fact, out of the 24 regression models that we replicated, only one model – without controls or lagged dependent variable – supports the theory. In the 23 other tests, the critical quantity of interest does not cross (or even approach) conventional thresholds of statistical significance.⁴¹ Put simply, the evidence in *Benchmarking Across Borders* amounts to little more than a null result.

HOW TO TEST BENCHMARKING WITH MULTIPLE REFERENCE POINTS

The models considered above show little evidence of benchmarking. This surprising result could be an artefact of several factors. For instance, our models may be too simple to capture the complex processes at work, or KP’s dataset may be too small to conduct well-powered tests. In this section, we consider how to adapt our barebones empirical framework to the more complex

³⁵Kayser and Peress 2012, 680.

³⁶680.

³⁷662.

³⁸670.

³⁹669.

⁴⁰Our hypothesis that the marginal effect of international growth should be negative is formally equivalent to the “partial benchmarking” hypothesis discussed in KP (2012, 668). According to KP, the appropriate way to test partial benchmarking is to check if a Wald test allows us to reject the null that θ_{y-i} and θ_i are equal. The p-value that this Wald test produces is exactly identical to the p-value of the δ_i coefficient in our model.

⁴¹The p-values of the δ_i coefficient for all 24 models are: 0.233, 0.006, 0.904, 0.472, 0.187, 0.713, 0.329, 0.316, 0.479, 0.454, 0.575, 0.957, 0.330, 0.209, 0.233, 0.690, 0.395, 0.223, 0.702, 0.389, 0.246, 0.798, 0.443, 0.321.

case where voters compare domestic economic performance to multiple reference points. Then, we illustrate by studying a larger dataset drawn from a more recent study of benchmarking.

Aytaç⁴² develops a reference point theory that is highly reminiscent of KP's, but which makes two important substantive changes. First, the author argues that voters use two reference points to assess their government's performance: The level of international growth (G_i) and their own country's historical level of growth (G_h).

Second, Aytaç points out that these reference points could be compared to two alternative measures of the incumbent's performance: Domestic growth during the election year (G_y), or domestic growth during the incumbent's full term in office (G_t). The term-based measure is preferable if we adopt a rational voter model, since such voters can extract more information about the quality of government by observing performance over a longer period. The election year measure is preferable if we take the view – dominant in political psychology – that voters are cognitively limited, myopic, and that they use end-heuristics when engaging in retrospective evaluations.⁴³ Here, we remain agnostic and estimate models using both measures.

In our framework, testing theories of benchmarking with multiple reference points is straightforward: We simply introduce the new reference point variable additively in Model 5. Again, there is evidence of benchmarking if the marginal effect of domestic growth is positive, and if the marginal effects of the benchmarks are negative.

In Table 2, we illustrate this by estimating six models using Aytaç's replication data.⁴⁴ In a first set of three models, we compare international and historical growth to domestic growth in the election year. In a second set of three models, we compare international and historical growth to the average domestic growth rate during the incumbent's full term in office. We include the same control variables as Aytaç.

All six of the models in Table 2 show evidence of conventional economic voting: The coefficients for domestic economic growth (G_y or G_t) are all positive and statistically significant. In contrast, none of the models allows us to reject the null hypothesis of “no international benchmarking”: The G_i coefficient is never statistically significant at the $\alpha = 0.1$ level.

The two right-most models in Table 2 show evidence of historical benchmarking: The G_h is negative and statistically significant. However, it is important to point out that those models rely on a highly unconventional assumption. Indeed, they assume that voters have long enough memories to accurately compare the average level of growth during the incumbent's full term in office, to the average level of growth during the previous government's term. This assumption clashes with common wisdom in the field of economic voting, where “virtually all macro-studies assume a short lag, generally of one year”.⁴⁵ Most studies of benchmarking also use short-term measures of domestic growth.⁴⁶

In sum, the results in Table 2 offer strong support for the conventional theory of economic voting, but evidence of benchmarking is mixed. None of the models allows us to confidently reject the absence of international benchmarking, and the only models that support historical benchmarking require us to jettison the widespread assumption that voters are myopic.

⁴²Aytaç 2018.

⁴³Healy and Lenz 2014.

⁴⁴In the online appendix, we explain why these models do not replicate Aytaç's faithfully.

⁴⁵Lewis-Beck and Stegmaier 2013, 378.

⁴⁶Ebeid and Rodden 2006; Kayser and Peress 2012; Powell and Whitten 1993.

TABLE 2: Benchmarking with multiple reference points and alternative measures of domestic growth. OLS regressions with country-clustered standard errors.

Domestic growth: Reference point(s):	Election year			Term in office		
	International	Historical	Both	International	Historical	Both
G_y	0.676*** (0.162)	0.732*** (0.171)	0.728*** (0.171)			
G_t				1.273*** (0.254)	1.387*** (0.257)	1.492*** (0.273)
G_i	-0.002 (0.273)		0.032 (0.268)	-0.392 (0.291)		-0.410 (0.288)
G_h		-0.308 (0.213)	-0.309 (0.213)		-0.567*** (0.202)	-0.572*** (0.202)
Vote share lag	0.689*** (0.064)	0.691*** (0.064)	0.691*** (0.064)	0.673*** (0.065)	0.669*** (0.063)	0.673*** (0.063)
Coalition	0.216 (1.037)	0.215 (1.029)	0.217 (1.028)	0.161 (1.071)	0.183 (1.025)	0.148 (1.023)
ENP	-1.435*** (0.366)	-1.453*** (0.359)	-1.453*** (0.359)	-1.523*** (0.393)	-1.564*** (0.377)	-1.569*** (0.385)
Presidential	-4.558*** (1.108)	-4.554*** (1.113)	-4.559*** (1.113)	-4.744*** (1.120)	-4.836*** (1.093)	-4.790*** (1.100)
Re-run	12.284*** (2.485)	12.163*** (2.467)	12.165*** (2.468)	12.986*** (2.426)	12.852*** (2.360)	12.869*** (2.387)
Constant	11.697*** (3.354)	12.587*** (3.230)	12.527*** (3.329)	11.853*** (3.475)	12.599*** (3.280)	13.362*** (3.435)
Observations	460	460	460	460	460	460

Robust standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

HOW TO TEST CONDITIONAL THEORIES OF BENCHMARKING

The regression models that we studied so far were relatively under-specified. Indeed, one of the major contributions of Powell & Whitten⁴⁷ was to point out that the level of economic voting depends on the institutional context (e.g., clarity of responsibility). Similarly, there are good reasons to think that benchmarking will vary across populations: some voters – such as those with high information – might engage in more relative economic evaluations than others.

If benchmarking is truly conditional, then the “pooled” models that we estimated above would be inappropriate, and our null results would not be surprising. For this reason, it is extremely important to develop regression models capable of testing conditional benchmarking hypotheses. Again, this is very easy to do in our simple empirical framework.

We use the same starting point as before (Figure 1): Benchmarking predicts that the marginal effect of domestic growth should be positive, and that the marginal effect of the reference point should be negative. If a moderating variable M increases (decreases) the salience of the reference point, then the marginal effect of domestic growth should be more (less) positive, and the marginal effect of the reference point should be more (less) negative where M is high.

This idea can be captured by a simple extension of Model 5:

$$V = \delta_y G_y + \delta_i G_i + \delta_{ym} G_y \times M + \delta_{im} G_i \times M + \delta_m M + \Gamma \Omega + \varepsilon, \quad (6)$$

where M stands for a variable that moderates comparative economic assessments.

As usual, a positive marginal effect of domestic growth ($\delta_y + \delta_{ym} M > 0$) would be consistent with both conventional economic voting and benchmarking. A negative marginal effect of international growth ($\delta_i + \delta_{im} M < 0$) would be consistent with benchmarking. The slopes of those marginal effects (δ_{ym} and δ_{im}) measure the extent to which M moderates relative economic assessments.

To illustrate how one can apply Model 6, we revisit a secondary set of tests from Aytaç (2018), where the author studies if benchmarking is more prevalent in countries with high trade intensity, GDP per capita, or average level of schooling.⁴⁸ We assess the moderating effect of all three variables,⁴⁹ include the same control variables as in Table 2, and use Aytaç’s two alternative measures of domestic growth (G_y and G_t). The full regression results are reported in the online appendix.

Figure 2 shows the estimated marginal effect of international growth in six models. None of the marginal effects is clearly negative, and most lines are nearly flat. These results, estimated using a dataset that is over twice the size of KP’s, offer no evidence of international benchmarking, and no evidence that trade, income, or education increase the salience of comparative economic assessments.

⁴⁷Powell and Whitten 1993.

⁴⁸The interaction models that we report here are slightly different from those estimated by Aytaç (2018). In the online appendix, we take a close look at Aytaç’s interaction specifications. Our discussion highlights some of the pitfalls of testing conditional benchmarking using composite variables and redundant regressors.

⁴⁹For comparison, the moderators are all rescaled to the [0, 1] interval.

Figure 2: Marginal effect of international growth on votes for the incumbent in six regression models with three different moderators and two alternative measures of domestic growth. 95% confidence intervals in grey.

CONCLUSION

In this paper, we re-interpreted the theory of benchmarking and explained that, all else equal, it predicts that votes for the incumbent should be *positively* related to domestic growth, but *negatively* related to reference points. By re-casting the theory's predictions in terms of the marginal effects of domestic growth and the reference points, we showed that benchmarking could be tested using a simpler linear model which excludes duplicate regressors, immediately produces the relevant discriminating statistics, and greatly facilitates interpretation.

We re-analyzed data from prominent studies which have claimed to present evidence clearly supportive of benchmark. Across a range of models, we found robust evidence that domestic growth affects voting behavior, but very little sign of benchmarking. We therefore conclude that benchmarking is an interesting hypothesis, but that it is *not* supported by the available evidence.

These results should not be interpreted as a wholesale rejection of the benchmarking hypothesis. Indeed, it seems reasonable to expect that some populations may be more responsive to international or historical comparisons than others. For example, voters may be particularly attuned to the economic performance of neighboring countries or rivals.⁵⁰ At the individual level, some types of voters (e.g., politically sophisticated) may also be more prone to compare domestic with global economic performance or present economic conditions with previous ones. The idea that some voters evaluate economic performance in relative terms has some intuitive appeal, but the idea that most voters systematically and accurately compare with an “objective” benchmark seems rather implausible given citizens' cognitive limitations.

⁵⁰Jérôme, Jérôme-Speziari, and Lewis-Beck 2001; Hansen, Olsen, and Bech 2015.

Perhaps most importantly, we have shown that there are great risks in using composite measures when it comes to testing theories of relative evaluation. We have demonstrated that there is a straightforward way to test the benchmarking hypothesis, which is to avoid composite variables, and to simply enter each term additively in the regression equation. With such an approach, we can formulate clear tests of the benchmarking hypotheses: all else equal, the benchmark should have a negative marginal effect on support for the incumbent party (or other relevant dependent variables). This simple empirical framework can also be extended in straightforward fashion to test theories of benchmarking with multiple reference points or context-conditionality. We hope to have provided clear guidelines for further research on this complex and important question.

FUNDING

We acknowledge the financial support of the Social Sciences and Humanities Research Council of Canada, and of the Fonds de Recherche Société et Culture du Québec.

REFERENCES

- Achen, Chris H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections do Not Produce Responsive Government*. Princeton: Princeton University Press.
- Anderson, Christopher J. 2007. "The End of Economic Voting? Contingency Dilemmas and the Limits of Democratic Accountability." *Annual Review of Political Science* 10:271–296.
- Aytaç, Selim E. 2018. "Relative Economic Performance and the Incumbent Vote: A Reference Point Theory." *Journal of Politics* 80 (1): 16–29.
- Bartels, Larry. 2012. "Elections in Hard Times." *Public Policy Research* 19 (1): 44–50.
- Converse, Philip E. 2000. "Assessing the Capacity of Mass Electorates." *Annual Review of Political Science* 3:331–353.
- Duch, Raymond M., and Randolph Stevenson. 2010. "The Global Economy, Competence, and the Economic Vote." *Journal of Politics* 72 (1): 105–123.
- Ebeid, Michael, and Jonathan Rodden. 2006. "Economic Geography and Economic Voting: Evidence from the US States." *British Journal of Political Science* 36 (3): 527–547.
- Fernández-Albertos, José. 2006. "Does Internationalisation Blur Responsibility? Economic Voting and Economic Openness in 15 European Countries." *West European Politics* 29 (3): 28–46.
- Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven: Yale University Press.
- Goplerud, Max, and Petra Schleiter. 2016. "An Index of Assembly Dissolution Powers." *Comparative Political Studies* 49 (4): 427–456.
- Hansen, Kasper M., Asmus L. Olsen, and Michael Bech. 2015. "Cross-National Yardstick Comparisons: A Choice Experiment on a Forgotten Voter Heuristic." *Political Behavior* 37 (4): 767–789.

- Healy, Andrew, and Gabriel L. Lenz. 2014. "Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy." *American Journal of Political Science* 58 (1): 31–47.
- Healy, Andrew, and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16:285–306.
- Hellwig, Timothy, and David Samuels. 2014. "Voting in Open Economies: The Electoral Consequences of Globalization." *Comparative Political Studies* 40 (3): 283–306.
- Jérôme, Bruno, Véronique Jérôme-Speziari, and Michael S. Lewis-Beck. 2001. "Évaluation économique et vote en France et en Allemagne." In *L'opinion européenne*, edited by Bruno Reynié D. et Cautrès. Paris: Presses de Sciences Po.
- Kayser, Mark Andreas, and Michael Peress. 2012. "Benchmarking across Borders: Electoral Accountability and the Necessity of Comparison." *American Political Science Review* 106 (3): 661–684.
- Key, V.O. 1966. *The Responsible Electorate*. Cambridge (MA): Harvard University Press.
- Leigh, Andrew. 2009. "Does the World Economy Swing National Elections?" *Oxford Bulletin of Economics and Statistics* 71 (2): 163–181.
- Lewis-Beck, Michael S. 1988. *Economics and Elections: The Major Western Democracies*. Ann Arbor: University of Michigan Press.
- Lewis-Beck, Michael S., and Mary Stegmaier. 2013. "The VP-function Revisited: A Survey of the Literature on Vote and Popularity Functions after over 40 Years." *Public Choice* 157 (3/4): 367–385.
- Powell, Bingham G., and Guy D. Whitten. 1993. "A Cross-National Analysis of Economic Voting: Taking Account of the Political Context." *American Journal of Political Science* 37 (2): 391–414.
- Przeworski, Adam, Susan C. Stokes, and Bernard Manin. 1999. *Democracy, Accountability, and Representation*. 2nd ed. Cambridge: Cambridge University Press.
- Raftery, Adrian E. 1995. "Bayesian Model Selection in Social Research." *Sociological Methodology* 25:111–163.
- Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50 (3): 755–769.
- Wolfers, Justin. 2002. *Are Voters Rational?: Evidence from Gubernatorial Elections*. Technical report 1730. Stanford: Graduate School of Business, Stanford University, March.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.

Supplementary Materials:
Do Voters Benchmark Economic Performance?

VINCENT AREL-BUNDOCK, ANDRÉ BLAIS AND RUTH DASSONNEVILLE*

*vincent.arel-bundock@umontreal.ca, andre.blais@umontreal.ca, ruth.dassonneville@umontreal.ca

TWO EQUIVALENT MODELS OF BENCHMARKING

Kayser and Peress (2012) estimate this model:

$$V = \theta_{y-i}(G_y - G_i) + \theta_i G_i + \Gamma\Omega + \varepsilon.$$

We recommend this simpler model:

$$V = \delta_y G_y + \delta_i G_i + \Gamma\Omega + \varepsilon.$$

In the main text, we explain why the second model, which excludes duplicate regressors, is preferable. In Table A1, we show numerically that the two models are logically equivalent. In particular, we see that $\theta_{y-i} \equiv \delta_y$, that $\theta_i - \theta_{y-i} \equiv \delta_i$, and that all the other coefficients are exactly identical in the two models.

TABLE A1: The simple and complicated models of benchmarking are equivalent.

	KP	ABD
Local Growth ($G_y - G_i$)	0.611774*** (0.234206)	
Domestic Growth (G_y)		0.611774*** (0.234206)
Global Growth (G_i)	0.050281 (0.374609)	-0.561493 (0.423542)
Local Unemployment	-0.040779 (0.187222)	-0.040779 (0.187222)
Global Unemployment	-0.327835 (0.216280)	-0.327835 (0.216280)
Coalition size	-3.332975*** (0.714253)	-3.332975*** (0.714253)
Eff.Num.Parties	-2.774487*** (0.599097)	-2.774487*** (0.599097)
Population	0.000039** (0.000015)	0.000039** (0.000015)
Year	0.034884 (0.056109)	0.034884 (0.056109)
Constant	-16.505970 (111.444462)	-16.505970 (111.444462)
Observations	189	189
R^2	0.516	0.516

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

BENCHMARKING ACROSS BORDERS

In the text, we evaluated KP's evidence according to one criterion: There is evidence of benchmarking if the marginal effect of international growth, controlling for domestic growth, is negative. This condition is exactly identical to the hypothesis that KP call "partial benchmarking" in their article. In Tables A2 to A5, we present full replication results for all of KP's models. Those tables show that there is essentially no evidence of partial benchmarking in KP's own models and data.

Here, we consider three additional empirical claims from *Benchmarking Across Borders*.

Claim #1: A statistically insignificant estimate of θ_i constitutes evidence of "full benchmarking."

KP write:¹

If "voters fully benchmark, we would expect the coefficient on local growth to be positive and the coefficient on global growth to be zero."

In terms of Model 3, this hypothesis can be translated as $\theta_{y-i} > 0$ and $\theta_i = 0$. The $\theta_{y-i} > 0$ condition is trivial because θ_{y-i} is the marginal effect of domestic growth; both benchmarking and conventional theories make the same prediction with respect to that parameter. The full benchmarking criterion thus boils down to $\theta_i = 0$.

Unfortunately, tests of full benchmarking cannot add much credence to the theory, because they suffer from three fundamental flaws.

1. Full benchmarking is a more restrictive special case of partial benchmarking. If full benchmarking ($\theta_{y-i} > 0$ and $\theta_i = 0$), then partial benchmarking ($\theta_{y-i} > \theta_i$). The contrapositive: If no partial benchmarking, then no full benchmarking. We have shown that KP's data offer little evidence in support of the *weak* partial benchmarking condition. Thus, readers should be extremely wary of claims that the *stronger and more restrictive* full benchmarking condition finds empirical support.
2. As KP concede (footnote 11), their tests of full benchmarking are under-powered. This is because the benchmark measure of global growth was constructed by combining growth rates from many countries. A natural side-effect of this strategy is that KP's measure of global growth has low variance, and that confidence intervals around the θ_i parameter are very large. Low statistical power and high uncertainty are not problems *per se*, but they become vexing when combined with the third flaw of the full benchmarking hypothesis.
3. The $\theta_i = 0$ proposition does not offer a well-founded statistical test of benchmarking, because it would require that the analyst *accept a null hypothesis*. Accepting the null when our model fails to reject it is a well-known statistical fallacy. That trap is especially dangerous when our tests are under-powered and confidence intervals are wide.² If the proposed test was valid, scientists could conclude that any null hypothesis is true, provided that their sample was small enough, or that their independent variable showed sufficiently low variation. At the very least, proponents of benchmarking who hope to convince that

¹Kayser and Peress 2012, 668.

²Most introductory econometrics texts discuss this issue. See for example Wooldridge (n.d.) or the amusing discussion of the consequences of "micronumerosity" (i.e., small samples) for null hypothesis tests in Goldberger (n.d.).

$\theta_i = 0$ should demonstrate that the confidence intervals for that parameter are narrow and centered around zero. In fact, they are not.

In short, finding a statistically insignificant estimate of θ_i should not be treated as evidence of benchmarking, because that result is overdetermined by the under-powered nature of KP's test. Furthermore, since "full benchmarking" is a more restrictive special case of "partial benchmarking", any claim that there is evidence of the former but not the latter would be dubious. Thus, we must reject the "full benchmarking" hypothesis, and focus on the criterion that we used in our paper.

Claim #2: The substantive effect of "decomposed" growth is more important than the effect of domestic growth.

The foregoing discussion was largely focused on statistical significance, ignoring the information about substantive importance which can be gleaned from coefficient sizes. In their article, KP estimate a version of Equation 1, and compare the G_y coefficient in that model to the $(G_y - G_i)$ coefficient in Equation 3:

"the difference in the effect sizes between column (1) and column (3) (our preferred benchmarking model) is worth noting. According to model (1), a 1% increase in growth is associated with a 0.604% increase in the leader party's vote share. According to model (3), a 1% increase in [the local component of] growth is associated with a 1.261% increase in the leader party's vote share. The estimated effect size therefore more than doubles when moving from the conventional model to the benchmarked model."

As we explained in the text, the coefficient associated with the $G_y - G_i$ composite variable is the marginal effect of domestic growth on votes for the incumbent; it has no meaningful interpretation in terms of the effect of relative economic performance on votes. As such, a comparison between β_y and θ_{y-i} seems to be of little relevance.

Moreover, we note that the coefficient estimate described in the quote above is extreme. In the 12 aggregate-data models that KP consider, the coefficient estimate associated with $G_y - G_i$ takes on values of: 1.261, 0.612, 0.529, 0.646, 0.818, 0.820, 0.538, 0.683, 0.474, 0.687, 0.582, 0.756. The first value is clearly an outlier which, incidentally, was produced by the only model (out of 24) that allowed KP to reject the null hypothesis of "no benchmarking."

Claim #3: At several points in time, the magnitude of the benchmarked economic vote is greater than the magnitude of the non-benchmarked economic vote.

For Figure 2 of their article, KP estimate Models 1 and 2 over different time windows, and plot the evolution of the β_g and λ_{g-r} coefficients over the 1990-2010 period. Considering that KP explicitly reject Model 2 as an inappropriate test of benchmarking,³ this comparison seems irrelevant. Moreover, the models used in Figure 2 lack the control variables that make KP's core model specification more credible.

³Kayser and Peress 2012, 663.

RELATIVE ECONOMIC PERFORMANCE AND THE INCUMBENT VOTE

Aytaç's additive specification

In our article, we estimated a series of models closely related, but not exactly identical, to those used by Aytaç.⁴ In his article, Aytaç estimates this model (Table 1, Model 2):

$$V = \kappa_{th}(G_t - G_h) + \kappa_{ti}(G_t - G_i) + \kappa_i G_i + \kappa_y G_y + \Lambda\Omega + \nu \quad (7)$$

The author defines “Relative Domestic Growth” as $G_t - G_h$ and “Relative International Growth” as $G_t - G_i$. Then, he interprets the results as follows:

Model 7 “introduces my measures of relative performance, Relative Domestic Growth and Relative International Growth. Both variables have positive and statistically significant effects on the incumbent’s vote, providing evidence for the hypothesis that voters reward (punish) incumbents on whose watch the economy performs relatively better (worse) in domestic and international comparisons.”⁵

This is precisely the mistake that we warned against in our article. When we control for “Relative International Growth” ($G_t - G_i$), the coefficient associated with “Relative Domestic Growth” ($G_t - G_h$) loses its obvious substantive interpretation. κ_{th} does not measure benchmarking. Likewise, by “partialling-out” one of the bounds of the $G_t - G_i$ interval, we change the substantive meaning of its associated coefficient. As a result, κ_{ti} no longer has a straightforward interpretation in terms of the gap between domestic and international growth.

As we explained in our article, the correct way to test benchmarking is to look at the marginal effects of our variables of interest. If voters engage in benchmarking, we should find:

$$\frac{\partial V}{\partial G_t} = \kappa_{th} + \kappa_{ti} > 0 \quad (\text{Marginal effect of full term growth}) \quad (8)$$

$$\frac{\partial V}{\partial G_y} = \kappa_y > 0 \quad (\text{Marginal effect of election year growth}) \quad (9)$$

$$\frac{\partial V}{\partial G_i} = \kappa_i - \kappa_{ti} < 0 \quad (\text{Marginal effect of international growth}) \quad (10)$$

$$\frac{\partial V}{\partial G_h} = -\kappa_{th} < 0 \quad (\text{Marginal effect of historical growth}) \quad (11)$$

Based on these criteria, the results in the original article no longer reject the null hypothesis of “no international benchmarking”, but there is still support historical benchmarking. This is consistent with the results we reported in our article, where we concluded that none of the models showed evidence of international benchmarking, but that there seemed to be evidence of historical benchmarking in the models where voters are assumed to rationally consider the average level of growth over the incumbent’s full term in office.

The only minor difference between the results that Aytaç reports⁶ and ours, is due to the fact that he includes both G_y and G_t in the same model, whereas we treat those two measures of

⁴Aytaç 2018.

⁵24.

⁶Aytaç 2018.

domestic growth as competing alternatives, in separate models. We believe that our approach is more principled: Given that the election year is part of the incumbent's term in office, it makes little sense to talk about the effect of changes in G_y , holding G_t constant. Still, we concede that reasonable people could disagree on this point.

Aytaç's multiplicative specification

To test conditional benchmarking, Aytaç estimates a model of this form:

$$\begin{aligned} V = & \kappa_{ti}(G_t - G_i) + \kappa_{mti}M(G_t - G_i) + \\ & \kappa_{th}(G_t - G_h) + \kappa_{mth}M(G_t - G_h) + \\ & \kappa_y G_y + \kappa_{my}M G_y + \\ & \kappa_i G_i + \kappa_m M + \Lambda\Omega + \nu, \end{aligned} \quad (12)$$

where M stands for a variable that moderates comparative economic assessments.

By now, the next step should be familiar. If voters benchmark, these conditions should hold:

$$\frac{\partial V}{\partial G_t} = \kappa_{ti} + \kappa_{th} + (\kappa_{mti} + \kappa_{mth})M > 0 \quad (\text{Marginal effect of full term growth}) \quad (13)$$

$$\frac{\partial V}{\partial G_y} = \kappa_y + \kappa_{my}M > 0 \quad (\text{Marginal effect of election year growth}) \quad (14)$$

$$\frac{\partial V}{\partial G_i} = \kappa_i - \kappa_{ti} - \kappa_{mti}M < 0 \quad (\text{Marginal effect of international growth}) \quad (15)$$

$$\frac{\partial V}{\partial G_h} = -\kappa_{th} < 0 \quad (\text{Marginal effect of historical growth}) \quad (16)$$

The equations above highlight some of the pitfalls of using composite variables to test conditional theories of benchmarking. Indeed, the marginal effects in Equation 12 are non-trivial expressions. In the original article, Aytaç (2018) fails to consider this complexity: He plots $\kappa_{ti} + \kappa_{mti}M$ and assesses the evidence of benchmarking on the basis of that graph. However, Equation 15 makes clear that this is not actually the quantity of interest. Specifically, the expression that Aytaç considers is missing the κ_i component. And since all the estimates of κ_i in his model are positive, Aytaç's plot systematically overstates the evidence of international benchmarking.

TABLE A2: Replication of KP Table 1 – Aggregate-level results for benchmarking in the economic vote

	Model 1	Model 2 Median	Model 3 PC	Model 4 Trade
G_{dt-1}	0.6044 (0.2673)	0.8185 (0.3315)	1.2612 (0.3525)	0.8196 (0.4037)
G_{gt-1}		-0.7207 (0.6019)	-1.3056 (0.4661)	0.0846 (0.7039)
$Unem_{dt-1}$	-0.2485 (0.2055)	-0.2095 (0.2095)	-0.3359 (0.2287)	0.0274 (0.2113)
$Unem_{gt-1}$		-0.3806 (0.5268)	0.2058 (0.2554)	-0.3046 (0.5433)
(Intercept)	34.0951 (1.6989)	37.4457 (3.9438)	35.0705 (1.9972)	32.7637 (4.3722)
N elections	213	213	213	146
N countries	22	22	22	22
R^2	0.029	0.035	0.061	0.034
Adj. R^2	0.020	0.017	0.043	0.007
RMSE	10.4858	10.5025	10.3616	10.4145

OLS estimates with heteroskedasticity robust standard errors in parentheses.

TABLE A3 : Replication of KP Table 3 – Robustness checks for aggregate-level models

	Model 1 Median	Model 2 PC	Model 3 Trade	Model 4 Median	Model 5 PC	Model 6 Trade	Model 7 Median	Model 8 PC	Model 9 Trade
G_{dt-1}	0.5375 (0.2184)	0.6118 (0.2342)	0.6826 (0.2844)	0.4744 (0.1902)	0.5288 (0.2163)	0.6869 (0.2304)	0.5824 (0.2264)	0.6360 (0.3235)	0.7561 (0.2970)
G_{gt-1}	-0.4402 (0.6113)	-0.5615 (0.4235)	-0.2326 (0.6307)	-0.4690 (0.4794)	-0.3837 (0.3819)	-0.3442 (0.4850)	-0.3578 (0.4682)	-0.2741 (0.4815)	-0.0250 (0.4581)
$Unem_{dt-1}$	-0.1480 (0.1610)	-0.0411 (0.1872)	-0.0653 (0.1584)	-0.2413 (0.1333)	-0.2524 (0.1686)	-0.1812 (0.1388)	0.0025 (0.2030)	0.1848 (0.2782)	-0.1581 (0.2734)
$Unem_{gt-1}$	-0.3196 (0.4228)	-0.2868 (0.2265)	-0.1552 (0.4309)	-0.4388 (0.3323)	-0.0748 (0.2104)	-0.1757 (0.4065)	-0.6272 (0.4139)	-0.6649 (0.4539)	-0.3462 (0.6718)
$coalsize$	-3.2305 (0.7206)	-3.3331 (0.7143)	-3.6446 (0.9056)	-1.3574 (0.6003)	-1.5058 (0.5941)	-1.7489 (0.7101)	-1.4885 (0.5393)	-1.3986 (0.6114)	-1.5110 (0.6041)
$enep$	-2.7303 (0.6086)	-2.7742 (0.5991)	-2.0543 (0.6761)	0.4177 (0.5620)	0.4537 (0.5481)	0.5802 (0.6896)	2.3533 (0.8824)	2.4167 (0.8758)	3.4912 (1.0286)
pop	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)	0.0000 (0.0000)	0.0002 (0.0001)	0.0002 (0.0001)	0.0006 (0.0006)
$elecyr$	0.0559 (0.0591)	0.0349 (0.0561)	0.0761 (0.1033)	0.0258 (0.0467)	0.0056 (0.0428)	-0.0293 (0.0664)	-0.0526 (0.0577)	-0.0455 (0.0625)	-0.1427 (0.1108)
Vote (previous election)				0.7753 (0.0770)	0.7650 (0.0765)	0.7284 (0.0959)	0.7697 (0.0739)	0.7562 (0.0842)	0.8054 (0.1163)
(Intercept)	-58.1209 (116.7971)	-16.5549 (111.4470)	-103.3335 (206.5107)	-41.2479 (91.8039)	-2.8184 (84.4547)	67.4097 (131.1914)	99.8943 (110.7509)	85.4243 (119.6557)	263.3620 (211.0814)
N elections	189	189	131	189	189	131	189	189	131
N countries	21	21	21	21	21	21	21	21	21
R^2	0.508	0.516	0.481	0.710	0.709	0.671	0.529	0.532	0.543
Adj. R^2	0.487	0.494	0.447	0.696	0.695	0.646	0.505	0.509	0.509
RMSE	7.8379	7.7784	7.8773	6.0339	6.0431	6.2998	5.5732	5.5552	5.5645

OLS estimates with heteroskedasticity robust standard errors in parentheses.

TABLE A4: Replication of KP Table 5 – Individual-level results for benchmarking in the economic vote

	Model 1	Model 2 Median	Model 3 PC	Model 4 Trade
Outside option	-1.9415 (0.1939)	-1.9435 (0.1932)	-1.9411 (0.1943)	-1.9171 (0.1988)
Leader party	0.6706 (0.2271)	0.9432 (1.6838)	0.7031 (0.2454)	1.2971 (0.4029)
Policy distance	-0.4984 (0.0243)	-0.4986 (0.0236)	-0.4984 (0.0242)	-0.5022 (0.0250)
G_{dt-1} x Leader party	0.1009 (0.0406)	0.1139 (0.0396)	0.1206 (0.0485)	0.1110 (0.0452)
$Unem_{dt-1}$ x Leader party	-0.0083 (0.0289)	-0.0036 (0.0275)	-0.0030 (0.0296)	0.0058 (0.0293)
G_{gt-1} x Leader party		-0.1058 (0.1087)	-0.0564 (0.0449)	-0.0732 (0.0614)
$Unem_{gt-1}$ x Leader party		0.0030 (0.2640)	0.0043 (0.0153)	-0.0794 (0.0432)
N voters	40260	40260	40260	40260
N elections	31	31	31	31
N countries	17	17	17	17
pseudo- R^2	0.166	0.166	0.166	0.168

Coefficients of conditional logit models explaining vote for the lead party.

Standard errors in parentheses are clustered by election.

TABLE A5: Replication of KP Table 6 – Robustness checks for individual-level models

	Model 1 Median	Model 2 PC	Model 3 Trade	Model 4 Median	Model 5 PC	Model 6 Trade	Model 7 Median	Model 8 PC	Model 9 Trade
Leader party	0.4349 (1.4502)	0.1991 (0.1980)	0.5450 (0.3953)	0.7500 (1.4234)	0.4101 (0.2905)	0.8585 (0.4196)	0.6655 (1.4788)	0.5339 (0.3176)	0.9419 (0.4546)
Policy distance	-0.5035 (0.0244)	-0.5036 (0.0245)	-0.5036 (0.0261)	-0.5042 (0.0247)	-0.5041 (0.0248)	-0.5045 (0.0265)	-0.5026 (0.0245)	-0.5028 (0.0247)	-0.5028 (0.0264)
G_{dt-1} x Leader party	0.0835 (0.0592)	0.1305 (0.0588)	0.1184 (0.0530)	0.0870 (0.0567)	0.1268 (0.0589)	0.1172 (0.0501)	0.0898 (0.0586)	0.1273 (0.0626)	0.1190 (0.0527)
G_{gt-1} x Leader party	-0.0425 (0.1066)	-0.0501 (0.0588)	-0.1006 (0.0826)	-0.0394 (0.1028)	-0.0510 (0.0592)	-0.0914 (0.0787)	-0.0268 (0.1045)	-0.0466 (0.0607)	-0.0803 (0.0809)
$Unem_{dt-1}$ x Leader party	-0.0163 (0.0242)	-0.0002 (0.0241)	-0.0133 (0.0201)	-0.0212 (0.0233)	-0.0068 (0.0242)	-0.0184 (0.0193)	-0.0259 (0.0231)	-0.0121 (0.0235)	-0.0226 (0.0188)
$Unem_{gt-1}$ x Leader party	-0.0035 (0.2089)	-0.0228 (0.0210)	-0.0140 (0.0490)	-0.0166 (0.2047)	-0.0172 (0.0217)	-0.0215 (0.0482)	0.0067 (0.2147)	-0.0159 (0.0231)	-0.0194 (0.0492)
Additional controls									
Party characteristics	X	X	X	X	X	X	X	X	X
Demographics x Leader party				X	X	X	X	X	X
Demographics x Party characteristics							X	X	X
N voters	38177	38177	36229	37602	37602	35658	37602	37602	35658
N elections	31	31	30	31	31	30	31	31	30
N countries	17	17	17	17	17	17	17	17	17
pseudo-R ²	0.214	0.215	0.210	0.215	0.215	0.211	0.222	0.223	0.219

Coefficients of conditional logit models explaining vote for the lead party. Standard errors in parentheses are clustered by election.

TABLE A6: Conditional benchmarking. OLS regressions with country-clustered standard errors.

	Education	Education	Trade	Trade	Income	Income
G_i	0.376 (1.158)	0.515 (1.071)	0.621 (1.402)	0.806 (1.290)	-0.219 (1.043)	0.094 (1.026)
G_i X Moderator	-1.356 (1.714)	-0.913 (1.573)	-1.578 (1.966)	-1.211 (1.769)	-0.210 (1.366)	-0.064 (1.361)
G_t	0.864 (0.575)		1.350* (0.795)		2.076*** (0.722)	
G_t X Moderator	0.811 (1.013)		-0.016 (1.331)		-1.099 (0.912)	
G_y		0.737* (0.373)		0.762** (0.351)		1.339*** (0.417)
G_y X Moderator		-0.090 (0.719)		-0.107 (0.629)		-0.978 (0.594)
Moderator	3.576 (5.379)	4.117 (5.409)	8.652 (6.465)	6.691 (6.067)	7.519 (5.395)	5.474 (4.830)
Vote share lag	0.672*** (0.065)	0.691*** (0.065)	0.683*** (0.064)	0.698*** (0.064)	0.678*** (0.065)	0.695*** (0.064)
Coalition	0.022 (1.112)	0.153 (1.056)	-0.069 (1.058)	0.058 (1.031)	0.021 (1.105)	0.118 (1.072)
ENP	-1.557*** (0.396)	-1.464*** (0.367)	-1.600*** (0.374)	-1.489*** (0.351)	-1.546*** (0.379)	-1.434*** (0.354)
Presidential	-4.417*** (1.150)	-4.272*** (1.146)	-4.278*** (1.220)	-4.237*** (1.210)	-3.931*** (1.184)	-3.987*** (1.140)
Re-run	12.571*** (2.586)	11.841*** (2.616)	12.592*** (2.477)	11.942*** (2.519)	12.644*** (2.479)	11.773*** (2.479)
Constant	9.838* (5.585)	9.293* (5.301)	5.794 (5.994)	6.988 (5.698)	5.974 (5.606)	7.308 (5.084)
Observations	457	457	458	458	457	457
R^2	0.602	0.585	0.603	0.586	0.600	0.586

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

REFERENCES

- Aytaç, Selim E. 2018. "Relative Economic Performance and the Incumbent Vote: A Reference Point Theory." *Journal of Politics* 80 (1): 16–29.
- Goldberger, Arthur Stanley. n.d. *A Course in Econometrics*. Cambridge (MA): Harvard University Press.
- Kayser, Mark Andreas, and Michael Peress. 2012. "Benchmarking across Borders: Electoral Accountability and the Necessity of Comparison." *American Political Science Review* 106 (3): 661–684.
- Wooldridge, Jeffrey M. n.d. *Introductory Econometrics: A Modern Approach*. Scarborough (ON): Nelson Education.